

Part

**1**

# Introduction to Network Technologies and Performance



# Open Systems Interconnection (OSI) Model

**Justin S. Morrill, Jr.**

*Hewlett-Packard Co., Colorado Springs, Colorado*

A *protocol* is an agreed-upon set of rules and procedures that describe how multiple entities interact. A simple example of a protocol in everyday life is the motoring rule specifying that the vehicle to the right at an intersection has the right-of-way, other things being equal. If this traffic protocol is violated, the result might be a serious problem.

When the entities are network devices, protocols are necessary for interaction to happen at all. If two devices follow different protocols, their communication will be no more successful than a conversation between a person speaking French and a person speaking Chinese. As there is more and more essential data traffic over a wide variety of networks, the ability to guarantee protocol interoperability has become increasingly vital. A number of standards have been developed to make that possible.

Among these standards, one has been designed to facilitate complete interoperability across the entire range of network functions: the *Open Systems Interconnection* (OSI) Reference Model, published by the International Standards Organization (ISO).

In computing and communications, *open* refers to a nonproprietary standard. An *open system* is one in which systems from different manufacturers can interact without changing their underlying hardware or software. The OSI model is such a standard and is a useful framework for describing protocols. It is not a protocol itself, but a model for understanding and defining the essential processes of a data communications architecture.

Since its conception, the OSI model has become a vital tool in two ways:

1. As a point of reference for comparing different systems or understanding where and how a protocol fits into a network.
2. As a model for developing network architectures that are maximally functional and interoperable.

## 4 Introduction to Network Technologies and Performance

### 1.1 Data Communications Protocols

In data communications, all interaction between devices is specified by protocols. These protocols are an agreement between sender and receiver defining conventions such as:

- When a device may transmit.
- The order of an exchange.
- What kind of information must be included at any given point in the transmission (such as which sections of a data package contain addressing, error control, message data, etc.) or which wire is reserved for which type of information, as in the interface described below.
- The expected format of the data (such as what is meant by a given sequence of bits).
- The structure of the signal (such as what pattern of voltages represents a bit).
- The timing of the transmission (for example, the receiving device must know at which points to sample the signal in order to correctly separate the bits).

The EIA 232 (also known as RS-232) physical connection, commonly found on the back of data terminals and personal computers, is specified by a protocol. This protocol is defined by the Electrical Industries Association (EIA), a standards-setting organization that assigns, numbers, and publishes the standards for manufacturers. The protocol includes the pin assignments for each signal and the loading and voltage levels that are acceptable. When a data communications connection fails, this protocol is usually the first to be analyzed for violations or problems that may impair the link operation.

As data communications have evolved, many manufacturers have decided to comply with standard protocols in order to ensure that their equipment will interoperate with that of other vendors. On the other hand, there are still proprietary protocols used that limit interoperability to devices from the same vendor. In either case, protocols provide the descriptions, specifications, and often the state tables that define the procedural interactions that allow devices to communicate properly.

#### 1.1.1 Layered protocols

Because of the complexity of the systems that they define, data communications protocols are often broken down into *layers*, also called *levels* (so called because they are schematically stacked on top of one another in order of use). The functions at each layer are autonomous and encapsulated so that other layers do not have to deal with extraneous details, but can concentrate on their own tasks. Encapsulation also provides a degree of modularity so that protocols at the same layer can be interchanged with minimum impact on the surrounding layers.

### 1.2 The OSI Reference Model

The OSI model, shown in Figure 1.1, consists of seven layers: Physical, Data Link, Network, Transport, Session, Presentation, and Application. The upper layers are

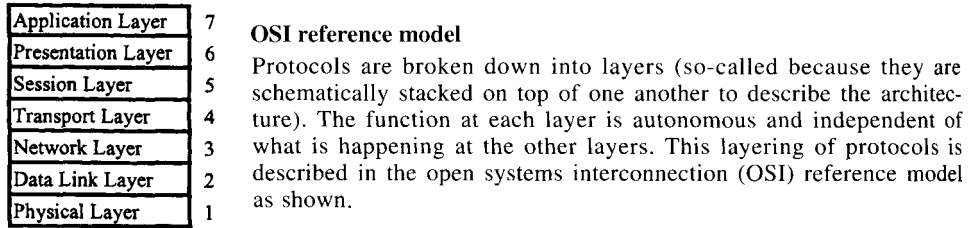


Figure 1.1

implemented in software, whereas the lower layers are implemented in a combination of software and hardware. Network test and measurement is concerned primarily with the functions of the lower layers and not with the content of the message, but with how well it is delivered.

**Note:** The layers of the OSI model may not be distinct in a specific protocol; in the TCP/IP protocol suite, for example, the popular *File Transfer Protocol* (FTP) includes functions at the Session, Presentation, and Application layers of the OSI model. Rather, the OSI model represents a theoretical superset of what is generally found in practice.

### 1.2.1 The Physical layer (layer 1)

The Physical layer in a data communication protocol (also known as layer one or level one) deals with the actual transmission of bits over a communication link. A loose analogy for the physical layer is the function of the internal combustion engine and the resulting source of mechanical motion in an automobile. The engine system performs on its own as long as its lubrication, ignition, cooling, fuel, and oxygen supply elements are functioning properly, and as long as the operator avoids actions that would damage the engine.

Protocols at layer one define the type of cable used to connect devices, the voltage levels used to represent the bits, the timing of the bits, the specific pin assignments for the connection, how the connection is established, whether the signal is electrically balanced or is single-ended, and so on. The specifications of EIA 232 in North America, or its V.24 European equivalent, are examples of Physical layer protocols.

**Note:** Numbering of protocols is done by the various standards bodies. The *X* and *V* series are defined by the International Telecommunications Union (ITU) in Europe; the EIA standards are published by the Electrical Industry Association in the United States. Other examples of Physical layer standards are the X.21 interface, EIA 449 interface, V.35 modem, 10Base-T Ethernet LAN, and Fiber Distributed Data Interface (FDDI) LAN.

The Physical layer elements interoperate with the media of connection and with the next layer of abstraction in the protocol (layer 2, the Data Link layer). Its specifications are electrical and mechanical in nature.

### 1.2.2 The Data Link layer (layer 2)

The Data Link layer provides error handling (usually in the form of error detection and retransmission) and flow control from one network node to the next. It provides

## 6 Introduction to Network Technologies and Performance

error-free transmission of a data parcel from one network link to the next. Using the automobile analogy, the Data Link layer might be compared to sensing changing conditions and modifying the inputs to the engine system to control it (for example, slowing the engine by limiting fuel and ignition).

In most protocols, the Data Link layer (layer 2) is responsible for providing an error-free connection between network elements. This layer formats the data stream into groups of bytes called *frames* of data for transmission and adds framing information to be interpreted by the remote device to which the frames are sent. Data Link layer functions generally exchange acknowledgment frames with the peer processes (Data Link layer functions) of the device to which it is directly connected. This interaction confirms the receipt of data frames and requests retransmission if an error is detected. Another major function of this layer is flow control, a provision for pacing the rate of data transfer to prevent a fast sender from overrunning a slow receiver.

### 1.2.3 The Network layer (layer 3)

The Network layer provides error-free transmission of a single data parcel end-to-end across multiple network links. Again with the automobile analogy, the Network layer might be compared to the operator's subliminal steering, which keeps the car on the road, and negotiating turns at appropriate corners. Additionally, decisions to change speed and make detours to avoid traffic congestion and even emergency avoidance of accidents also equate to layer 3 functions. The driver controls these functions, but does so automatically without thinking consciously about them, and can deal simultaneously with many other details that can be associated with higher-layer functions.

In data communication, the Network layer, layer 3, is responsible for the switching and routing of information and for the establishment of logical associations between local and remote devices, the aggregate of which is referred to as the *subnet*. In some cases, this layer deals with communication over multiple paths to a specific destination. The Network layer also can deal with congestion through flow control and rerouting information around bottlenecked devices or links. Information pertinent to layer 3 is appended to the frame from the Data Link layer. Once this addition is made, the result is a *packet* (named after a packet of mail that might be sent through a postal service).

### 1.2.4 The Transport layer (layer 4)

The Transport layer is responsible for the end-to-end delivery of the entire message. With the automobile analogy, this layer might be compared to the plan that the driver executes in getting from the origin to the destination of the trip. Often this plan requires using a map and choosing the most appropriate path based on the time of day, the urgency of the arrival, and so forth.

Transport layer (layer 4) responsibilities include the integrity of the data, the sequencing of multiple packets, and the delivery of the entire message—not just to the appropriate machine but to the specific application on that machine for which the data is intended (i.e., *port-to-port* delivery). While the lower three layers tend to be technology-dependent, the Transport layer tends to be independent of the end users' communications device technologies. This independence allows it to mediate

between the upper and lower layers, and to shield the upper layer functions from any involvement with the nuts and bolts of data transport.

### 1.2.5 The Session layer (layer 5)

The Session layer is responsible for establishing, maintaining, and terminating sessions between users or applications (if they are peer-to-peer). This layer might be very loosely compared to traffic laws that establish right-of-way.

The Session layer (layer 5) protocols establish conversations between different machines and manage applications on them with services of synchronization and mutual exclusion for processes that must run to completion without interruption. Protocols at this layer are responsible for establishing the credentials of users (checking passwords, for example), and for ensuring a graceful close at the termination of the session. An example of a graceful close mechanism is one that guarantees that the user of an automatic teller machine actually receives the money withdrawn from his or her account before the session terminates. Another example is the behavior of a printer with a paper jam. The function that causes the printer to reprint the damaged page, rather than going on from the jam point, is a Session layer protocol.

### 1.2.6 The Presentation layer (layer 6)

The Presentation layer ensures that the data is in a format acceptable to both communicating parties. It creates host-neutral data representations and manages encryption and decryption processes. In the automobile analogy, functions at this layer can be compared to a system that mediates geographically localized differences between automobiles, such as speedometer calibration in miles per hour or kilometers per hour, or steering wheel placement on the right or left side.

The Presentation layer (layer 6) is concerned with the syntax and semantics of the information that passes through it. At this layer, any changes in coding, formatting, or data structures are accomplished. Layer 6 is typically the layer used to accomplish encryption, if any, to prevent unauthorized access to the data being transmitted.

### 1.2.7 The Application layer (layer 7)

The Application layer provides the user or using process with access to the network. In the automobile analogy, it is roughly comparable to the mission of the trip and to the interface between car and driver (speedometer, odometer, gearshift, etc.). The mission sets the context of operation, including the urgency and the conservativeness or aggressiveness of the trip.

This layer is concerned with network services for a specific application, such as file transfer between different systems, electronic mail, and network printing.

### 1.2.8 User data encapsulation by layer

User data is formed and presented to the Application layer. From there it is passed down through the successively lower layers of the model to the Physical layer, which sends it across a link. At layers 7 through 2, information used by processes at each

8 Introduction to Network Technologies and Performance

layer is appended to the original message in a process called *encapsulation*. This information is added as headers at layers 7 through 2, and as a trailer at layer 2 (see Figure 1.2).

When the encapsulated transmission reaches its destination, it is passed up through the layers in a reverse of the sending process. Each layer removes and processes the overhead bits (header and/or trailer) intended for it before passing the data parcel up to the next layer. This activity requires the precise exercise of a number of parameters and procedures, providing multiple opportunities for processing error.

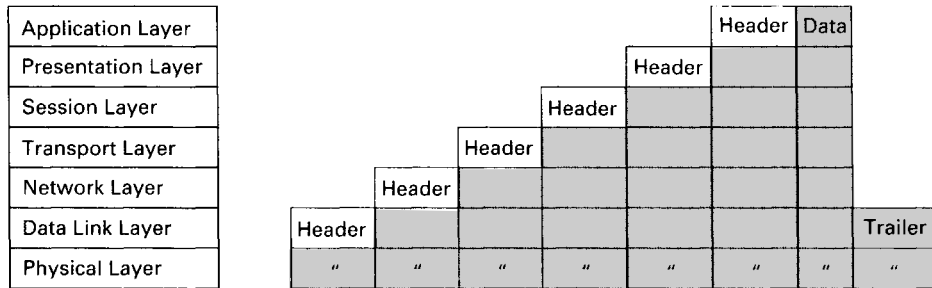


Figure 1.2 Encapsulation of data.



# Data Communications Basics

**Marc Schwager**

*Hewlett-Packard Australia Ltd., Victoria, Australia*

## 2.1 Introduction

The purpose of this chapter is to provide a basic understanding of the major components of a data communications network. This chapter focuses on the most common elements likely to be encountered in a data communications network. Voice networks, wireless networks, and proprietary networks such as those used in process control applications are not discussed. The treatment is necessarily brief; references listed at the end of the chapter for further information.

### 2.1.1 The network fabric

The network fabric is the combination of devices, wires, computers, and software that interact to form a data communications network. There are many of these that are brought together to create the *local area network* (LAN) and *wide area network* (WAN) environments that are in common use. There are three interlinked concepts that this chapter addresses: the protocol stack (TCP/IP, SNA, etc.), network topologies (ring, star, etc.), and the interconnects. The latter are the devices that do most of the work in the network, such as routers, hubs, and switches. These three aspects of networking will determine a large part of how network testing is approached.

### 2.1.2 A brief history of data networks

Data networks evolved from three areas: mainframe communications, personal computer (PC) networks that share peripherals, and workstation networks that share data.

The early data networks were built around point-to-point networks, that is, one mainframe was connected directly to another. IBM created protocols such as Remote Job Entry (RJE) to facilitate load sharing and job sharing between computers. The

## 10 Introduction to Network Technologies and Performance

minicomputer companies in the late 1970s and early 1980s expanded these capabilities considerably. With the widespread adoption of Ethernet and the proliferation of PCs, small networks emerged that enabled a workgroup to share expensive peripherals like laser printers. Engineering workstations were being developed that had integral networking capabilities, which were used for data and task sharing. The end of the 1980s saw the widespread adoption of networking and the creation of internetworks. These large corporate, university, and government networks were essentially a consolidation and interconnection of the “islands” of networking that had evolved.

These networks still carry many different protocols, and they connect many types of computer equipment. The network fabric must be extremely flexible and adaptable to handle the task. This is one reason that there are so many different interconnects. It makes the job of managing today’s networks challenging, and to make things worse, traffic in a typical corporate network grew at around 40 percent per year in the 1990s. The great intermeshing of networks will continue through the foreseeable future, with the major focus on the consolidation of voice, data, and video over a worldwide, high-speed fiber infrastructure.

## 2.2 Protocols

### 2.2.1 Common protocol stacks

Protocols are the language by which computers and other devices communicate on the network. A standard model, which takes a layered approach, has evolved to describe these protocols. Defined by the International Standards Organization, (ISO) it is called the Open Systems Interconnect (OSI) Reference Model. It has seven layers, each of which has a function to perform. A collection of these layers is called a *protocol stack*. Interconnects will base routing decisions on the lower layers. Some common protocol stacks are profiled here, with comments on their use.

**The OSI model.** Table 2.1 shows the Open Systems Interconnect model. Note that functions such as error detection can occur in more than one layer of the protocol stack. While the OSI model covers seven layers in a complete implementation, there are many protocol stacks that are focused at the Network layer and below. This is the case in most of the following examples.

**X.25.** Table 2.2 shows X.25, which is common in wide area networks. X.25 is a transport protocol stack, being defined only up through the Network layer. The use of hop-to-hop error recovery at both the Data Link layer and the Network layer makes X.25 a very robust protocol stack, and therefore a good choice when line quality is poor. Unfortunately this also makes it slow: X.25 can add 40 to 60 ms in traffic delay per network hop. Frame relay is preferable for connecting LANs over a wide area network.

**Frame relay.** Like X.25, frame relay (described in Table 2.3) is a WAN transport protocol stack, being defined only up through the Network layer. The absence of hop-to-hop error recovery makes frame relay much faster than X.25. Error recovery is handled by the upper-layer protocols such as TCP/IP in a typical LAN environment. Due to its low latency, frame relay is often used for connecting LANs over a wide area network. Frame relay can deal gracefully with traffic bursts, and can specify quality

**TABLE 2.1 The Open Systems Interconnect (OSI) Model.**

OSI layer	Function
Application	Provides common application service elements (CASEs) such as file transfer, virtual terminals, message handling, job transfer, directory services.
Presentation	Creates host neutral data representations, manages encryption and compression.
Session	Manages setup and orderly teardown of conversations, synchronization to coordinate data transfers.
Transport	Connection management, fragmentation management, flow control, priority control, error detection and correction, multiplexing data flows over one physical segment.
Network	Controls the topology and access to the network. This layer links logical (or network) addresses to physical addresses.
Data Link	Detects and corrects errors in the received bit stream. Physical addresses are in this domain.
Physical	Transmits and receives the data. Specifications deal with the wire or fiber (known as the <i>media</i> ), connectors, as well as the optical or electrical signals that are carried on the medium, including signal quality.

**TABLE 2.2 The X.25 Protocol Stack.**

Layer	Service	Notes
Network	X.25PLP	X.25 Packet Layer Protocol—Includes error recovery mechanisms
Data Link	LAPB	Link Access Procedure—Includes error recovery mechanisms
Physical	X.21	X.21bis is the spec for V-series interfaces (typically RS232). X21 has it's own physical interface as well.

**TABLE 2.3 The Frame Relay Protocol Stack.**

Layer	Service	Notes
Network	T1.606	This is the ANSI std, the CCITT equivalent is I.622
Data Link	T1.618	Link Access Procedure—No error recovery mechanisms (LAPF)
Physical	I.430/431	CCITT

of service (QoS). This is accomplished by having the user specify a committed information rate (CIR), which the network agrees to deliver, and some burst parameters that allow excess traffic in small amounts to pass through the network.

**ISDN.** Integrated Services Digital Network (ISDN), described in Table 2.4 has been around for years. In the 1980s it was something of a holy grail in wide area networking. It only broadly maps to the OSI model, so Table 1.4 should be treated as an approximation. It is designed to integrate voice and data traffic. Primary Rate ISDN (PRI) has been well accepted as a WAN service in Europe. In the United States, Basic Rate

## 12 Introduction to Network Technologies and Performance

**TABLE 2.4 The ISDN Protocol Stack.**

Layer	Service	Notes
Network	Q.931	Network Termination 2 (NT2), Error correction, segmentation.
Data Link	LAPD Q.921	Network Termination 2 (NT2) switching, layer 2 & 3 multiplexing, switching, concentration.
Physical	BRI, I.4xx PRI, G.703	Network termination 1 (NT1). Line maintenance, timing, layer 1 multiplexing, physical, electrical termination.

**TABLE 2.5 Transmission Control Protocol/Internet Protocol (TCP/IP).**

Layer	Service	Notes
Transport	TCP/UDP	Transmission Control Protocol: connection-oriented, used by services such as X Window, electronic mail, file transfer protocol (FTP), and Telnet. User Datagram Protocol: connectionless, used by services such as simple network management protocol (SNMP).
Network	IP, ARP	Internet protocol used for routing and addressing. Address Resolution Protocol (ARP) maps physical addresses to IP addresses.
	ICMP	Internet Control Message Protocol (ICMP) supplies control and error-handling functions.
Data Link	LLC/MAC	Link-Level Control/Media Access Control: This is typical for LANs.
	802.3	Each LAN device has its own unique address known as the MAC address. Other Data Link layer services such as Serial Line Internet Protocol (SLIP), and Point to Point Protocol (PPP) are common.
Physical	Various	802.3 is for Ethernet, Token-Ring is 802.5, others possible.

**TABLE 2.6 The Novell Netware Protocol Stack.**

Layer	Service	Notes
Transport	NCP/SPX	NetWare Core Protocol uses Service Advertisement Protocol to link clients and servers. Sequenced Packet Exchange (SPX) used for peer-to-peer networking.
Network	IPX	Internetwork Packet Exchange
Data Link	LLC/MAC 802.2/3	Link Level Control/Media Access Control; this is typical for LANs. Each LAN device has its own unique address, known as the MAC address. Other Data Link layer services such as Serial Line Internet Protocol (SLIP) are common.
Physical	LAN	802.3 is for Ethernet, Token-Ring is 802.5, others possible.

**TABLE 2.7 The SNA Protocol Stack.**

Layer	Service	Notes
Application	Function Mgt Data Services (FMDS)	Provides application mapping such as application files. Access to appropriate Network Addressable Units.
Presentation	NAU Service Manager	Network Addressable Unit (NAU) services manager. Manager Supports data compression and session services.
Session	Data Flow Control	Manages connection flow (full, or half duplex, etc.)
Transport	Transmission Control	Manages end-to-end transmission for sessions.
Network	Path Control	Manages logical channel links, virtual route control.
Data Link	SDLC	Synchronous Data Link Control.
Physical	Physical	Physical connections.

ISDN (BRI) is finding broad acceptance for home office and Internet access applications. The next generation of ISDN, called Broadband-ISDN or B-ISDN, generally refers to the Asynchronous Transfer Mode (ATM) protocol stack.

**TCP/IP.** TCP/IP (Table 2.5) is the protocol of the Internet. Above the transport, many common services such as FTP, e-mail, Telnet, SMTP, and SNMP exist. TCP/IP was developed by DARPA to be an extremely reliable transport (i.e., survive a nuclear war). It accomplishes this by allowing many different routes to a given endpoint, and by allowing for retransmissions if a packet fails to reach an endpoint.

**Novell NetWare.** NetWare is built around IPX, a Network layer protocol roughly analogous to IP (Table 2.6). Novell also supplies some higher-layer services (not shown) relating to server-based file sharing and other workgroup functions. NetWare is one of the most widely used LAN protocol stacks. The challenge with Novell has always been how to scale it up across a WAN. This has to do with the way NetWare advertises its services (frequently, and to almost everyone)—making for lots of WAN traffic. Novell has added burst mode to improve performance, and also the option of replacing IPX with IP in the stack to improve routing scalability.

**The SNA model.** IBM's Systems Network Architecture (SNA), shown in Table 2.7, is a hierarchical architecture. It is broken into domains, each controlled by a System Services Control Point (SSCP), most likely a mainframe. The SSCP deals with Physical Units (PUs) and Logical Units (LUs), which are defined based on capability. Different LUs have different upper-layer network services available to them; for example, LU1 is for application-to-terminal communications, while LU6 is for program-to-program communications. PUs come in different types, including terminals (PU1), hosts (PU5), and a variety of others.

### 2.2.2 Framing

Data generally moves in *frames*, *packets*, or *cells*. These packets are assigned address fields, which are used by various devices on the network for routing, bridging, and so on. Let's examine how the packets are formed and addressed. As a piece of

## 14 Introduction to Network Technologies and Performance

data moves from a computer into the top of the protocol stack, it gets wrapped in a series of headers and trailers that allow each layer of the stack to do its job. A simplified conceptual example of data moving from a computer through an IP stack onto an Ethernet LAN is shown in Figure 2.1. This describes the basic elements, with many detailed fields left out in order to reduce confusion.

Data starts on the local computer. As it is passed along, moving from the top of the protocol stack down to the network interface card, it is broken into the correct size for the protocol by the network driver. The *network driver* is a small piece of software that communicates between the computer system and its network card. As the data progresses down the TCP/IP stack from the top, service information is added at the TCP level. In the case of TCP, services are mapped to a logical entity called a *port number*. Following this, the IP layer adds the Network layer addressing information (in this case the IP address). The IP layer then hands the packet down to the Data Link layer, where the media access control (MAC) address or physical address is appended. A *cyclical redundancy check* (CRC) is added to the end of the packet to ensure packet integrity.

The packet is now fully assembled and ready to be passed to the Physical layer, where it is turned into electrical or optical signals on the physical media. In some cases the packet may be further processed by an interconnect. In the example, for instance, the completed packet might move to a router to be transported across a wide area network using the frame relay protocol. In this case, a frame relay header and trailer would be appended by the sending router, and then stripped off at the receiving end by the receiving router. The process that happens at each layer of the protocol stack, which treats anything passed down from above as data and appends appropriate headers and/or trailers to it, is known as *encapsulation*.

## 2.2.3 Data forwarding functions

This section describes five key packet forwarding functions and their relationship to the network stack. The network equipment that makes use of each function will be discussed later.

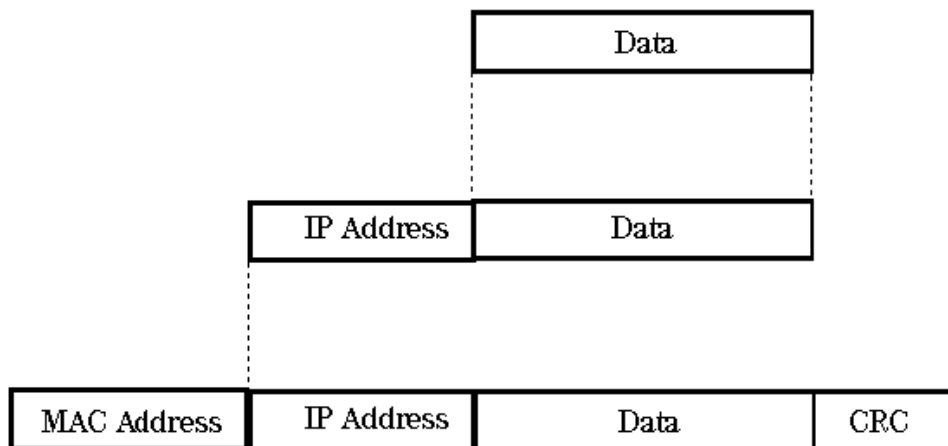


Figure 2.1 Data framing.

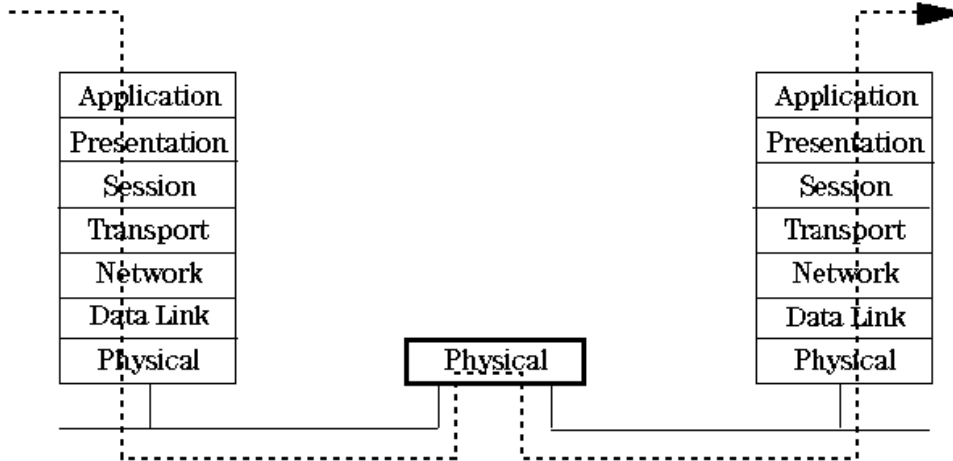


Figure 2.2 The function of a repeater.

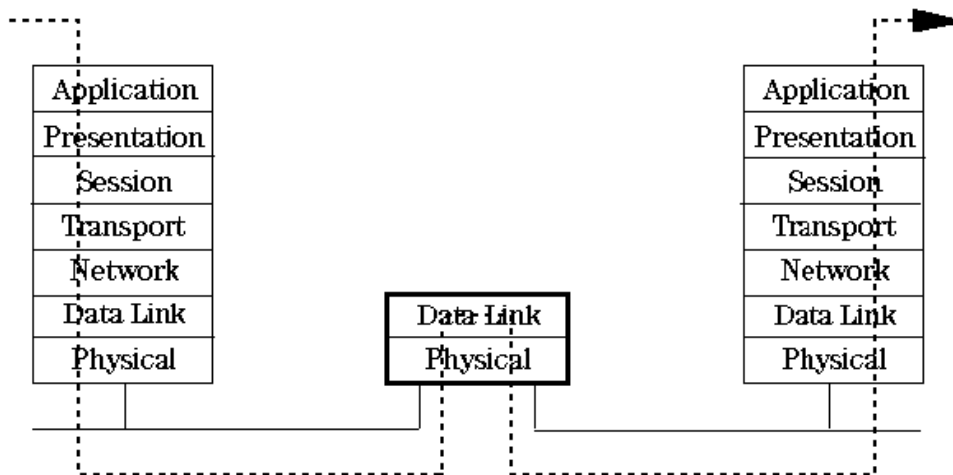


Figure 2.3 The function of a bridge.

**Repeating.** Repeating occurs at the physical layer. Repeating is used to extend cable distances and to isolate noise. As shown in Figure 2.2, only the Physical layer of the protocol stack is involved in repeating. A repeater simply looks at the electrical (or optical) signals on the media, and recreates those signals on a second piece of media. The new signals are regenerated and cleaned up to meet the physical specification of the Physical layer protocol. All traffic is repeated to all connections. No destination decisions are made.

**Bridging.** Bridging is accomplished at the Data Link layer (Figure 2.3). It can be used to connect two different physical media, such as the commonly used Ethernet

## 16 Introduction to Network Technologies and Performance

LAN cabling Thinnet (10Base2) and twisted-pair (10Base-T). Packets are forwarded from one link to another as needed, based on the Data Link layer address. LAN switching also works in this fashion, but at much higher speed. Network layer addressing is irrelevant for bridging.

**Routing.** Routing (Figure 2.4) operates at the Network layer; one use of routing is to connect networks that have different Data Link layers. Common examples would include connecting a LAN using Ethernet to a FDDI backbone, or connecting a LAN to a WAN. Routing can be very complex, but with the complexity comes flexibility and power. The most common Network layer protocol used for routing is IP, but Novell's IPX and other protocols also are routed. Routing relies on careful configuration in order to operate correctly. When configured correctly it provides secure, efficient communications that can scale up to very large networks. For example, Hewlett-Packard maintains a routed network with over 110,000 hosts worldwide.

**Gateways.** Gateways (Figure 2.5) are used when two entirely different network stacks need to exchange data. Computers can be configured to act as gateways by installing a card for each type of network, along with some appropriate software. To connect a TCP/IP Ethernet network to an SNA network would require a gateway due to differences at all levels in the protocol stack. Connecting an Ethernet network to a Token-Ring LAN would require only a bridge, provided the upper layers of the protocol stack are the same.

**ATM switching.** Asynchronous Transfer Mode (ATM), shown in Figure 2.6, is a Data Link protocol. It deserves special mention, however, both for its notoriety and for the way it operates. Data is transmitted in small, fixed-size packets (53 bytes long) called *cells*. The small cell size gives ATM the ability to interleave voice, data, and video traffic and deliver deterministic performance. End stations have ATM ad-




Figure 2.4 The function of a router.



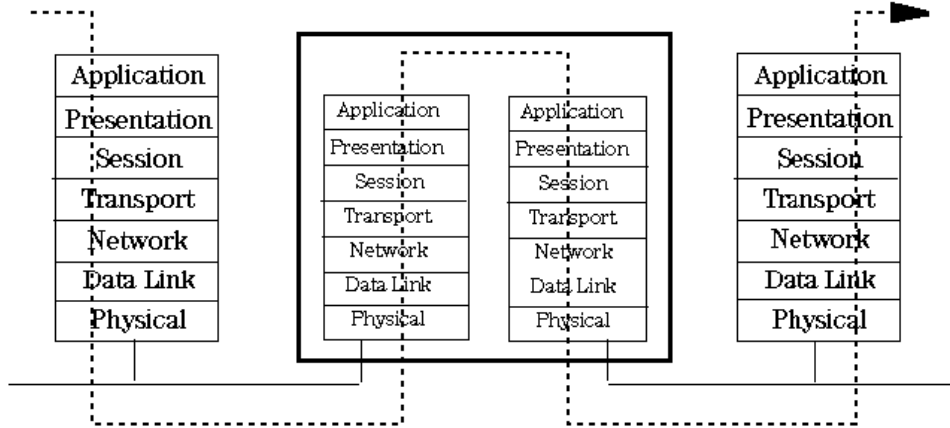


Figure 2.5 The function of a gateway.

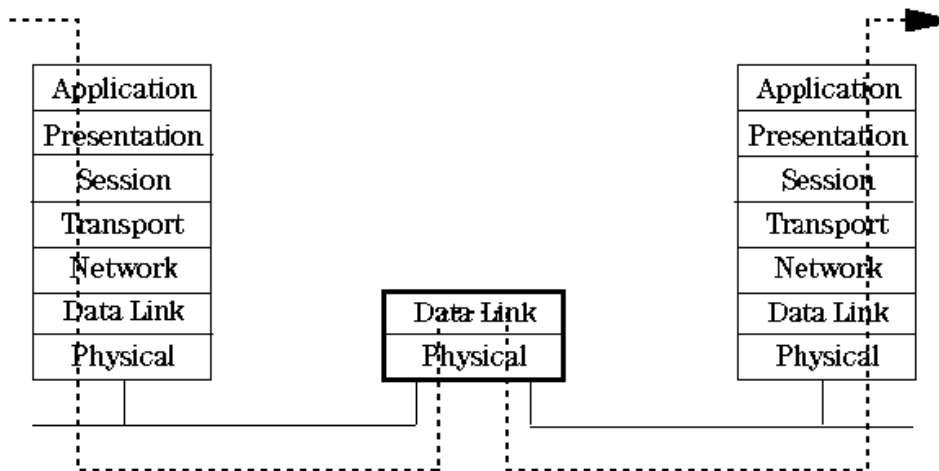


Figure 2.6 The function of an ATM switch.

dresses. ATM is connection-oriented, and a connection must be set up between the stations prior to beginning communications. Connections are set up either manually for permanent connections, or automatically for temporary connections.

ATM cells are forwarded by devices called *ATM switches*. To set up the connection, each switch in the path maps the input data stream to a specific output stream. These are designated as unique virtual path identifier/virtual channel identifier (VPI/VCI) pairs. Note that these change as they pass through each switch (Figure 2.7). When data is sent, the only address information in the cell is the VPI/VCI, which may be different depending on where the cell is examined. While ATM can be used directly by computers in an end-to-end fashion, it is more commonly used as a way to carry IP or frame relay traffic in a transparent fashion.

## 18 Introduction to Network Technologies and Performance

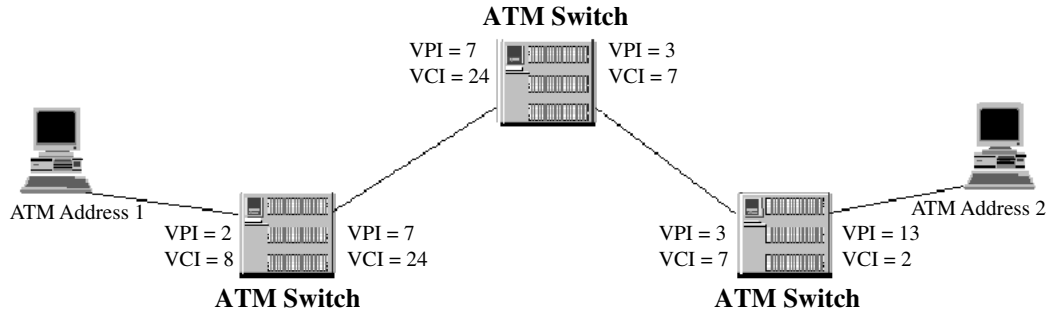


Figure 2.7 ATM VPI/VCI pairs.

## 2.3 Topologies

Networks are organized in different physical ways. These are called *topologies*. Table 2.8 gives an overview of topologies. Included in the table are:

- A diagram of the topology
- Devices commonly found on this type of network
- Protocols commonly used on the topology
- General attributes of the topology
- Notes on troubleshooting
- General comments

### 2.3.1 Point-to-point

These were historically the first networks. Point-to-point networks are used for a wide variety of situations, from connecting a PC to a server via a modem, to very high-speed links connecting supercomputers. Failures are easily isolated to a single link. Point-to-point networks do not scale gracefully. The number of links to connect a given number of nodes is given by the equation

$$L = \frac{N \times (N - 1)}{2} \quad (2.1)$$

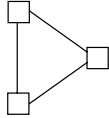
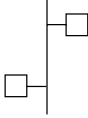
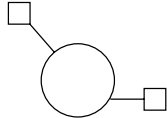
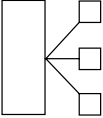
where

L = number of links

N = number of nodes

As N gets large, link creation and maintenance becomes difficult. For example, a 5-node network requires 10 links, while a 100-node network would require 4950 links!

**TABLE 2.8 Network Topologies.**

Topology	Devices	Protocols	Attributes	Troubleshooting	Comments
 <p>Point-to-Point</p>	<p>Mainframes Minicomputers Modems Interface cards SNA hardware SLIP Dial-up connections PADs PCs, terminals, Workstations</p>	<p>X.25 Frame relay ISDN SNA SLIP PPP Analog modem (many speeds and styles)</p>	<p>Static addressing Fixed routes WAN links owned and maintained by public carriers</p>	<p>Failures easily isolated to a link</p>	<p>First historic networks Many links required to connect all nodes. The formula is <math>L = N*(N - 1)/2</math> for complete coverage, where L is the number of links required and N is the number of nodes to be connected.</p>
 <p>Bus</p>	<p>Mainframes (recently) Minicomputers Print, file servers PCs, workstations Transceivers NICs Repeaters Bridges Routers</p>	<p>Mostly Ethernet: 802.2, 802.3, LocalTalk Rare but still existent -Arnet, 802.4 Typically IPX, IP, AppleTalk, Banyan VINES</p>	<p>Thin or thick LAN coax for Ethernet: 802.3, 10Mbps UTP daisy-chained for Apple LocalTalk</p>	<p>Physical fault domain spans entire cable Physical faults are a major failure mode</p>	<p>These were the first LAN networks Distance, number of host limitations spawned interconnect market. Poor physical security Bus topologies are being rapidly replaced by star topologies in private networks.</p>
 <p>Ring</p>	<p>Mainframes Minicomputers Print, file servers PCs, workstations NICs Token-Ring only: -MAU/MSAUs -CAUs -Source route bridges Bridges Concentrators (FDDI) Routers Multiplexers</p>	<p>802.5, FDDI Token-Ring: Typically SNA, 3270, IPX FDDI: IP, DECnet, IPX Encapsulated TR on FDDI not uncommon SONET/SDH in the WAN/MAN</p>	<p>Type 1, Type 3 Token-Ring connections CDDI is Cat 5, Multimode fiber for FDDI 4 or 16 Mbps for TR, 100 Mbps for FDDI 155 Mbps - 2.4 Gbps for SONET/SDH</p>	<p>Physical fault domain limited by protocol in TR Physical faults a major failure mode FDDI dual attach mitigates failures; look for ring wraps Mixing TR and other protocols can be a problem source</p>	<p>Driven by IBM, Token-Ring was one of the first; FDDI followed. Token-Ring, CDDI look like star topologies -physically; FDDI on fiber looks more like a ring TR : Source routing allows growth without routers, up to a point. Distance, number of hosts, source hop limitations drive topology limits MANs use SONET/SDH rings This is the most widely used LAN technology today by a factor of 2. It is quite inexpensive and typically deployed in a hierarchical fashion</p>
 <p>Star</p>	<p>Mainframes Minicomputers Print, file servers PCs, workstations Transceivers NICs Hubs, stackable and modular (concentrators) Bridges, routers Switches</p>	<p>10Base-T, 100Base-X 802.3 Ethernet ATM Typically IP, IPX</p>	<p>Typically Cat 3 or Cat 5 wiring 10Base-T for 10 Mbps 100 Mbps LANs include 100Base-T, 100VG-AnyLAN</p>	<p>Component and wiring failures easily isolated to a single link. Violating distance or configuration specs can cause problems</p>	<p>MANs use SONET/SDH rings This is the most widely used LAN technology today by a factor of 2. It is quite inexpensive and typically deployed in a hierarchical fashion</p>

## 20 Introduction to Network Technologies and Performance

### 2.3.2 Bus

The use of a “bus” created the first LAN networks. Because any device on the network can talk, a method was developed to minimize collisions on the network. The scheme employed on Ethernet networks is Carrier Sense Multiple Access with Collision Detection (CSMA/CD). A station will listen to the network to see if any other station is transmitting; if not, it will try to send its message. If by some chance two stations do this simultaneously, a collision occurs. When one is detected, each station waits a random interval and tries again. Collisions are a normal part of the Ethernet world, tending to limit performance to around 60 percent of the theoretical bandwidth, with throughput degrading under rising load.

Bus networks were easy to install in a small work area, and in small-scale usage provided an easy way to add users. They were developed for office as well as industrial use. Their use has been waning for a number of important reasons. One is component cost. Bus networks tend to be based on coaxial cable, which is more expensive than the twisted-pair wiring used in newer, hub-based networks such as 10Base-T Ethernet. A second reason is that the newer structured wiring designs (star topologies) have isolated fault domains. When a bus network fails, it takes down the entire segment, affecting all other users connected to the same physical cable. Cable faults are a common failure with this style of network.

### 2.3.3 Ring

A ring network can appear physically like a star network. The ring configuration often only manifests itself in the path that data follows in the network. (See Token-Ring MAUs below, for an example of this.) Ring LANs like Token-Ring and FDDI are generally based on token passing, where each station can take its turn on the network only when it has possession of a special packet called a *token*.

The advantage of this method is seen as the network utilization increases. Unlike the CSMA/CD-based Ethernet networks, there are no collisions in a token scheme. Token-passing networks therefore can maintain very high utilizations with little performance degradation. The tradeoff is that the ring protocols have a higher overhead, which cuts down the available bandwidth. Ring topologies such as Token-Ring, FDDI, and SONET (used in the wide area) have built-in fault resiliency. FDDI networks have found wide application in campus backbones. The downside of ring networks has been the higher historic costs associated with them due to the extra hardware required to implement the token protocols.

### 2.3.4 Star

While star networks have been used in the wide area for some time, it wasn't until the invention of the 10Base-T Ethernet hub that they became widespread in the local area. The combination of low cost and structured wiring have made this topology the most widely installed in LANs today. As in point-to-point networks, physical failures are easily isolated. These networks can be deployed hierarchically, avoiding the scaling issues associated with point-to-point. Star networks can be interconnected by a routing mesh, which looks similar to a point-to-point network. In a meshed net-

work, each router is connected to at least two other points. This gives a measure of fault tolerance in case one path fails, as well as the opportunity to balance the network load.

### 2.3.5 Virtual networks

Virtual networks (Figure 2.8) have appeared relatively recently. The physical topology of these networks is usually a hierarchical star or a routed mesh. Virtual networking allows you to gather arbitrary collections of nodes into a group for administrative purposes even if they are on different physical subnetworks. For example, you might put the members of an engineering team together in a group. The advantage of this approach is administrative, and requires that the network interconnects have enough bandwidth to make any rerouting transparent.

## 2.4 Interconnects

Interconnects are the devices that comprise the network. There are many categories, and the distinction between them becomes blurred as networking companies become more clever in their engineering and marketing. Some of the major interconnects are profiled in this section. The first section covers LAN devices and the second section covers WAN devices.

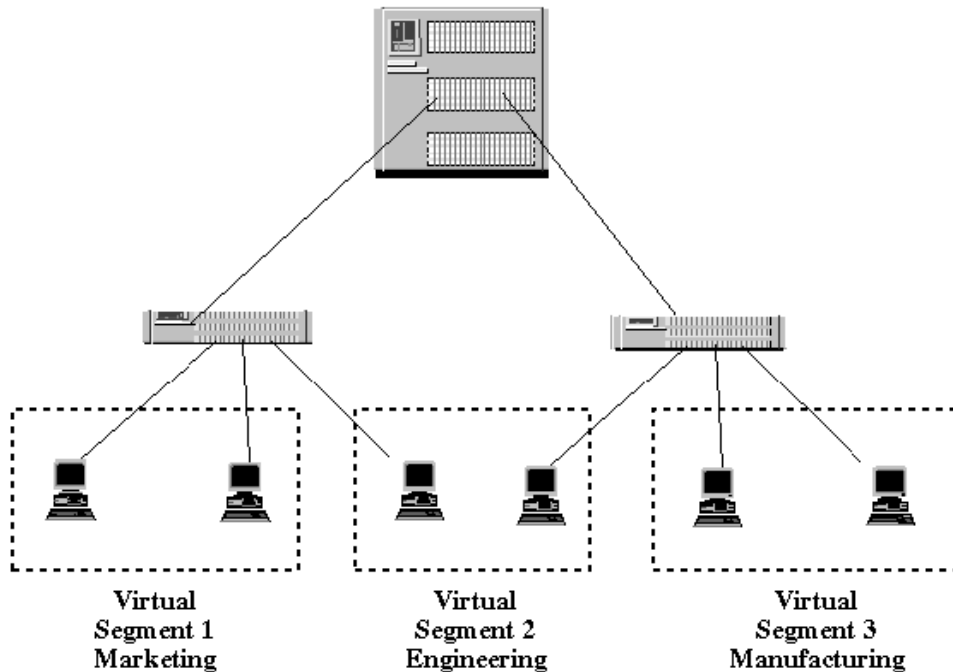


Figure 2.8 Virtual networks.

## 22 Introduction to Network Technologies and Performance

### 2.4.1 LAN interconnects

This section contains descriptions of and comments about devices commonly found on local area networks. Tables 2.9 and 2.10 contain the following information on LAN interconnects:

- Common name
- Device function
- Device limitations
- Designing for reliability
- Deployment hints
- Troubleshooting issues
- General comments

**Transceivers.** Transceivers (Figure 2.9) are used to connect the Attachment Unit Interface (AUI) port of a computer or peripheral to the physical medium. Most of today's computers come with a 10Base-T port (RJ-45 connector) built in. A transceiver might be used if you wanted to use a different medium, such as fiber. Transceivers are inexpensive, making it worthwhile to keep spares on hand, as they occasionally fail dramatically.

**Repeaters.** Repeaters (Figure 2.10) are used to extend cable length. They work by replicating the signals at the physical level. A repeater can be used to switch media types, in similar fashion as a bridge. Unlike a bridge, however, a repeater will not limit Ethernet collision domains, that is, two workstations on different cables connected by a repeater will still produce a collision if they transmit simultaneously. Repeater use is limited both by performance considerations (i.e., how many stations are to be squeezed into a segment), as well as protocol dependencies such as inter-frame gap preservation. A repeater will partition the network into two physical fault domains, so cable tests must be done on each side if a physical fault is suspected. For protocol problems, an analyzer can be hooked up anywhere. Repeaters generally will not filter out protocol errors.

**Hubs.** Hubs (Figure 2.11) are the most widely used interconnect today. They are used to connect end stations to a network. They may be connected in a hierarchical fashion, up to a limit of three for Ethernet. Note that a different cable (or a switch on the hub) is needed to connect two hubs together. If you need to configure the network so that traffic passes through more than three hubs, a bridge, router, or a LAN switch (discussed later) will be needed. The hub's structured wiring approach limits physical fault domains to a single wire.

There are two common hub packages: stackable hubs, and modular hubs or *concentrators*. The least expensive are stackables, which can be purchased by mail for less than \$100. The more expensive hubs come with built-in management capabilities. Ethernet hubs act as multiport repeaters, so any traffic sent to one port is repeated to

TABLE 2.9 LAN Interconnects (Part 1).

Name	Function	Limitations	Design for reliability	Troubleshooting	Comments
Transceivers Also called Media Access Units of MAUs	Used to connect computers and peripherals to a local area network.	Can be bulky, and can be knocked accidentally when handling off the back of a computer. Occasionally fail catastrophically. Keep a few extra transceivers on hand.	Transceivers can cause network problems. Look for runts and jabbers (short/long packets with bad CRCs) failures to an address.	Failures easily isolated to a link in 10Base-T. With bus topologies, use an analyzer to localize.	There are switch settings on transceivers that can increase reliability or hinder performance. Make sure these are set right for your network configuration. The "sqe" switch can clobber repeaters.
Repeaters	Used to extend cable length or adapt different cable types. Copies all traffic from one link to another. Multiport repeaters can connect a number of coax links.	Typically operates at the physical layer. If it sees a signal, it will copy it.	There is a limit in Ethernet of 3 repeaters for a segment. Exceeding this will cause problems with the interframe gap.	Will propagate errors. If Ethernet specs are violated, will cause errors. If cables are suspect, cable tests should be run on each individual wire.	These were some of the first LAN devices. They generally have little or no SNMP management capabilities. Hubs are also referred to sometimes as repeaters.
Hubs	Hubs are used to connect end stations to the network. They may also connect other hubs in a hierarchical configuration. These are repeaters that operate in a star topology. They form the basis of the majority of Ethernet networks today.	Stackable hubs are generally limited in their flexibility. A series of hubs connected together will create one large segment (i.e., collision domain), which can become congested.	Very reliable in general. Failures are easy to isolate in the star topology. Design the network to conform to specification. Watch cable lengths. No point-to-point signal should pass through more than 3 hubs before encountering a router or a bridge.	Hubs have varying degrees of SNMP management capabilities. This tends to vary with price. The most expensive will have embedded RMON agents. You can see all the traffic in a segment by hooking an analyzer up to a port in the hub. Collisions will vary by port. Do not be concerned about this.	These are plug-and-play devices for the most part. When connecting two hubs together in a 10Base-T environment, you must have a twisted cable. Most of these hubs are built around a single chip! The repeater MIB will map MAC address to port number. The more expensive units will do some internal bridging to create virtual LAN segments.
Token Ring MAUs	The Media Access Unit looks like a hub, or star, but operates in a Token Ring environment like a ring. It provides a way for a computer to hook into a Token-Ring with	Used in Token-Ring environments. The older connections are bulky and unwieldy.	Token-Ring has a robust protocol which isolates fault domains quickly. To take advantage of it, you should have a copy of IBM LAN Manager or equivalent, and take	You can see all the traffic in a segment by hooking an analyzer up to a port in the MAU. Depending on what you are looking for, you may	In a controlled environment, Token-Ring is a stable protocol. When connected to Ethernet via a router and routing common LAN protocols such as Novell or AppleTalk, problems are not uncommon.

TABLE 2.9 LAN Interconnects (Part 1) (Continued).

Name	Function	Limitations	Design for reliability	Troubleshooting	Comments
Bridges	<p>what looks like a single wire; it is actually two wires that form a piece of the ring.</p> <p>Bridging operates at layer 2 of the network. It connects one or more segments and passes traffic between them based on the destination MAC address. These were invented to overcome distance limitations and traffic congestion introduced by repeaters. Also used to extend LANs over the Wide Area</p>	<p>Passes all broadcast traffic. Can be limited in forwarding and filtering configurations. Some "bridges" are actually providing frame translation (such as from Token-Ring to Ethernet). Others can do protocol level (such as IPX) filtering. Bridge forwarding rates can limit LAN performance it can vary by packet size protocol mix, number of hosts, and protocol type.</p>	<p>the time to understand the protocol mechanism and errors. There is a mechanism for removing offensive nodes from the ring.</p> <p>Check your segment traffic against the bridge forwarding rate. Spanning tree capability resolves potential looping. Bridged networks are very susceptible to broadcast storms. These are hard to pin down, and can drastically reduce network performance. If your network is growing, consider moving to routers which, while harder to configure provide more flexibility, security, and manageability.</p>	<p>want to monitor without inserting (in a protocol sense) into the ring. Make sure that your analyzer can accomplish this.</p> <p>A bridged network can appear to an analyzer as a single segment unless filtering is going on. This can make troubleshooting problematic when the analyzer is trying to track 2500 hosts! To troubleshoot broadcast storms, you will need the ability to capture packets in a protocol analyzer. The general technique is to capture continuously, and set a trigger to freeze the buffer when a certain level of broadcast traffic occurs.</p>	<p>Bridges come with many different forwarding and filtering capabilities. For local area uses, LAN switches often provide a higher-performance solution, although they provide no filtering. It is hard to generalize, however, if you have a multiprotocol environment, consider taking the step to routing.</p>



TABLE 2.10 LAN Interconnects (Part 2).

Name	Function	Limitations	Design for reliability	Troubleshooting	Comments
Source route bridges	Used to link Token-Ring networks together.	Can handle a maximum of seven hops. For larger networks, routers are preferable.	Beware of overloading intermediate rings carrying transit traffic.	Watch the hop count limitation.	This is a simple way to extend Token-Ring without resorting to routers.
Routers	Link groups of computers and other network devices together using network-level addressing. Includes security and firewall features.	Expertise needed for correct configuration. Proprietary routing protocols can hinder interoperability.	Make sure they are configured properly. Stay current with firmware upgrades.	Monitor ICMP traffic for IP routing information. Compression must be turned off to use an analyzer.	Routers allow network segmentation to reduce congestion.
Servers	Servers are computers that may be acting as gateways, proxy servers for security, or routers.	Packet forwarding speed, and limited feature sets.	Configuration, especially of security firewalls.	If the server is also used for data storage, monitor performance. Networking tasks can consume large amounts of server resources.	These are general-purpose machines. They are not designed to be high-performance interconnects, and are thus suitable for smaller networks. The exception to this is when a server is configured to be a proxy server for security reasons.
Lan switches	These are fast replacements for hubs. They forward traffic like bridges, working at the physical address level.	Same as bridges. Subject to broadcast storms. Limited to a maximum of 3 layers hierarchically, due to address buffer requirements.	Store-and-forward switches will check CRC and not forward bad packets. Cut-through switches do not do this, and are faster.	Unlike hubs, traffic is not repeated on all ports. Visibility is limited to one link at a time. Some vendors allow port mirroring to a test port.	These are plug-and-play devices for the most part.
ATM switches	Very fast interconnects that are used anywhere from the workgroup to the backbone, based on size and features.	ATM standards are still evolving. Single-vendor solutions are more practical.	Stay with a single vendor until standards mature.	Interoperability can be suspect across ATM devices. ATM is often used as a transport for frame-based protocols. These systems can be complex. Testing requires sophisticated gear. Chances are good	Connection-oriented networks are fundamentally different from LANs. ATM goes a step farther and uses small cells to transfer data.

TABLE 2.10 LAN Interconnects (Part 2) (Continued).

Name	Function	Limitations	Design for reliability	Troubleshooting	Comments
DTCs	Used to connect data terminals to a LAN.	Inflexible, designed for one purpose.		that your stock protocol analyzer will not handle ATM well in detail, or at speed.	These are being phased out by the widespread use of PCs.
Concentrators (also called modular hubs)	Large, consolidated, industrial-strength, general-purpose interconnects. A backplane with plug-in cards. Cards typically include Token-Ring, Ethernet, FDDI, and others.	The backplane locks you into one vendor for additional capabilities.	Excellent SNMP-based management capabilities for configuration and troubleshooting.	RMON agents largely available and worth the investment.	Very flexible, expandable, and reliable units.




Figure 2.9 A transceiver.




Figure 2.10 A repeater.




Figure 2.11 Hubs in the network.

## 28 Introduction to Network Technologies and Performance

all the ports on the hub. This allows you to hook up an analyzer to any port on the hub to monitor all hub traffic. Note that collisions occur on a wire-by-wire basis, so each different port will show different numbers for collisions. Most hubs will have an indicator LED on each port to indicate the port status.

**Media Access Units (MAUs).** A MAU (rhymes with “cow”) is basically a hub for Token-Ring networks (Figure 2.12). Note in the diagram that while a MAU looks like the nexus of a star topology, the data actually travels on a ring. Each port on the MAU typically has an insertion LED that lets you know whether a station is inserted into the ring. Token-Ring will automatically remove stations from a ring, and heal the ring if a physical fault is observed. MAUs also have Ring In (RI) and Ring Out (RO) ports that allow them to be connected together to form larger rings (up to the limit of the Token-Ring specifications for number of stations per ring and total ring distance). An analyzer may be connected anywhere in the ring to observe the network.

**Bridges.** Bridging (Figure 2.13) allows you to scale up a network. Bridges can be used to solve a number of problems. The most common reasons to use a bridge are to connect different media types, reduce congestion on a segment, and to extend a LAN over longer distances. A bridge works by creating a table of MAC addresses for each of its links. It creates the table by listening to the network for packets and keeping track of which source addresses are on which link. When a packet reaches a bridge, it is compared to the table. If the MAC address of the destination is not on that segment, then the packet is forwarded. If the MAC address of the destination is on that segment, then the packet is not forwarded. This keeps local traffic in one collision domain from congesting other portions of the network. The exception to this is broadcasts. These are generally passed through the bridge. Large bridged networks are notorious for excessive broadcast traffic and broadcast storms.

Some bridges can filter by protocol (e.g., AppleTalk, DECnet, IPX), which is handy for keeping traffic separate and reducing global congestion. Bridges can be linked together in such a way as to inadvertently cause loops, where packets could travel

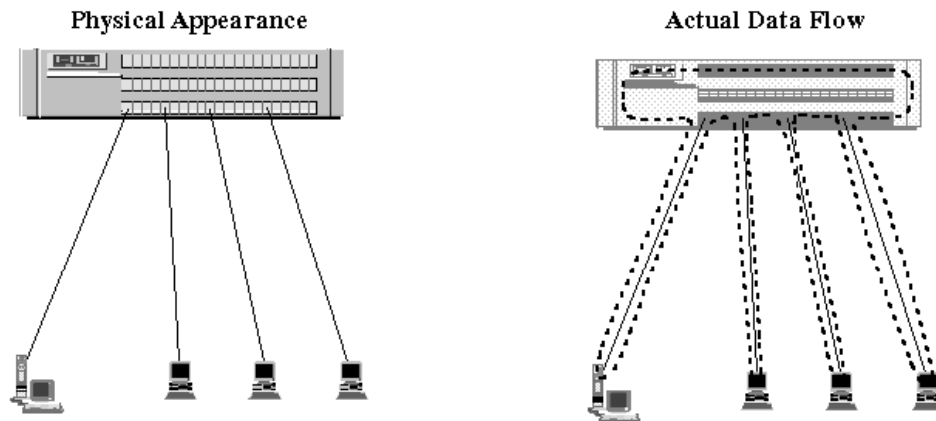
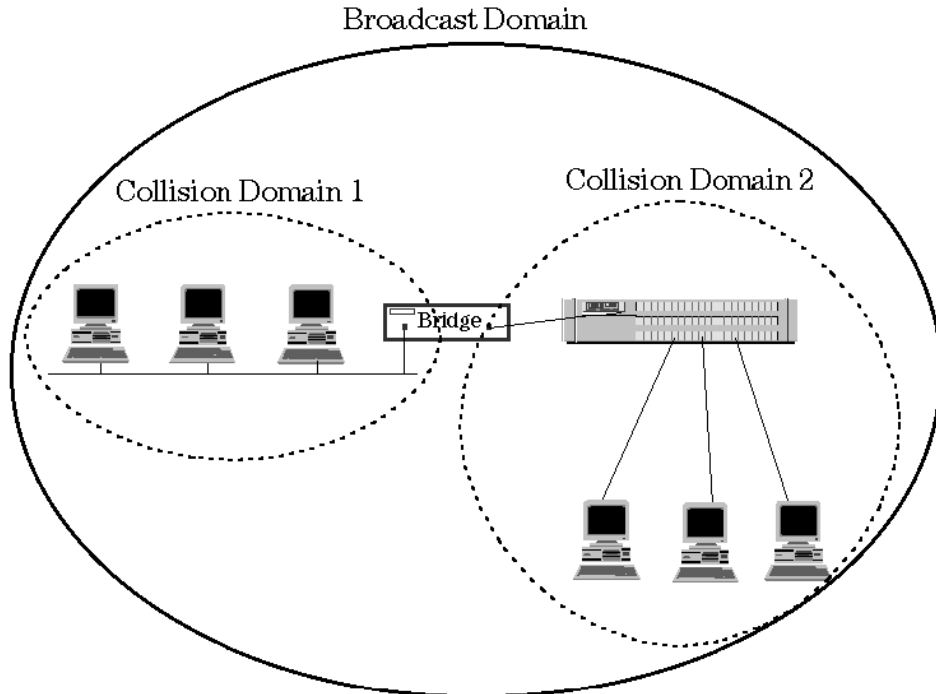


Figure 2.12 Token Ring Media Access Units (MAUs).



**Figure 2.13** A bridge in the network.

around in circles endlessly. This is avoided by the use of the *spanning tree algorithm*, common on most bridges today. Bridge forwarding rates can limit LAN performance. These rates vary by packet size, number of nodes, and protocol mix, so beware of best-case test data from vendors.

Bridges limit the physical fault domains and the protocol fault domains. If you are having a problem within a segment of your network, you must hook the analyzer up to the same segment, or you will not find the problem.

**Source route bridges.** Source route bridges (Table 2.10 and Figure 2.14) are Token-Ring devices that use a feature of the Token-Ring protocol to route traffic between rings. In Figure 2.14, source routing would be used to communicate between a station on ring 1 and a station on ring 3. Note that if there is a lot of traffic like this, ring 2 is going to get fairly busy just passing traffic between rings 1 and 3. Source route bridges are fairly easy to install and configure compared to a router. Source routing is limited to 7 hops (ring transits). If your network is that large, you should consider buying a router or a Token-Ring LAN switch.

**Routers.** Routers (Figure 2.15) are the workhorses of the public and private inter-networks today. They link different subnetworks using the Network layer address, typically IP but sometimes IPX. They use routing protocols (OSPF, RIP, IGRP) to communicate with one another, to keep routing tables up to date, and make decisions

## 30 Introduction to Network Technologies and Performance

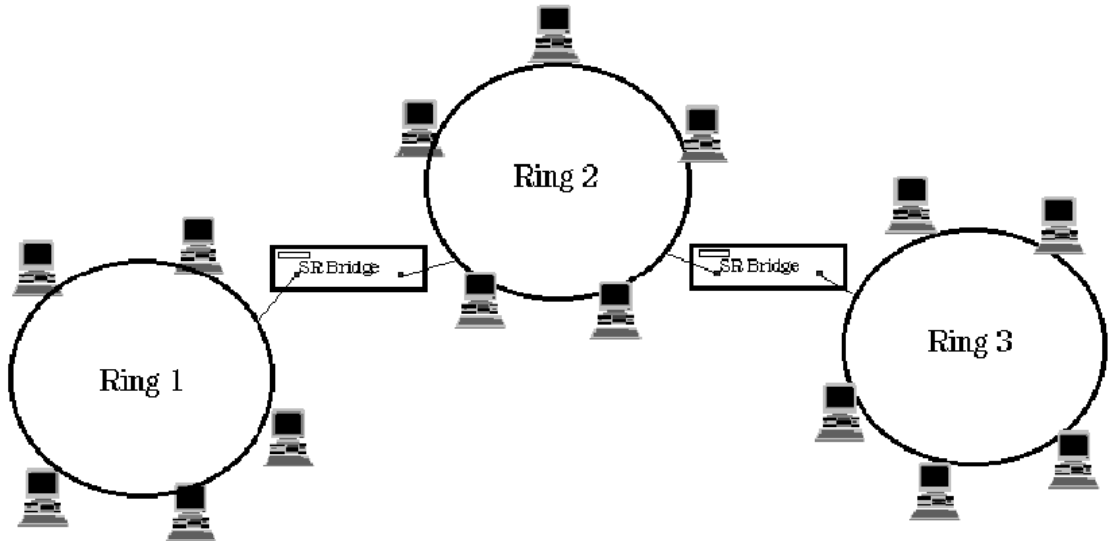


Figure 2.14 Source route bridges.

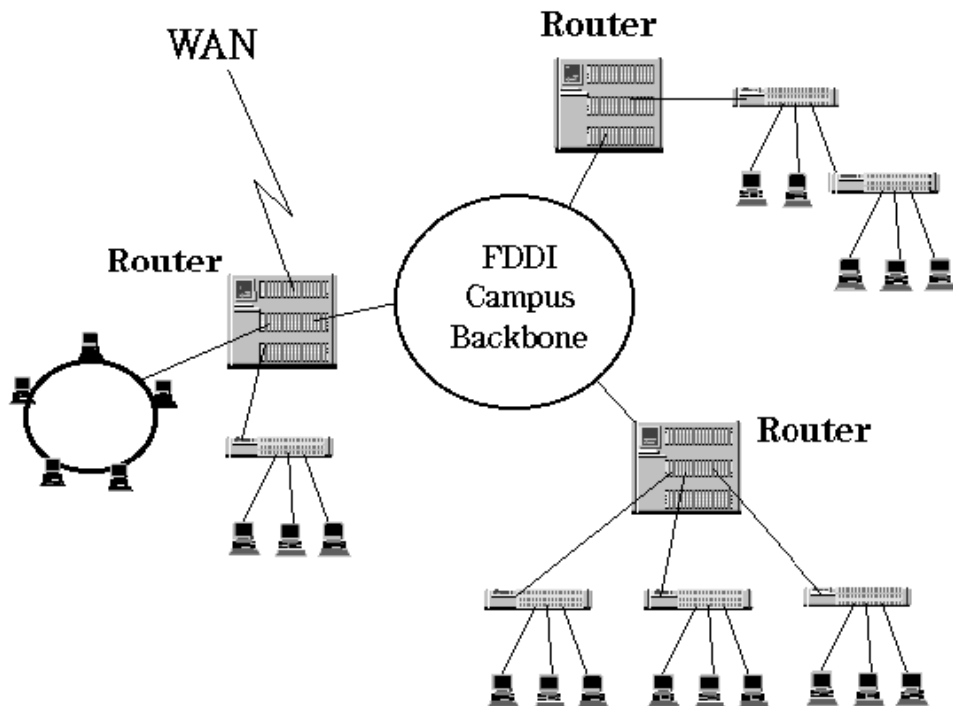


Figure 2.15 A routed network.

on how to route packets. Routers generally have high capacities for traffic forwarding; like bridges, however, their forwarding rates will vary by packet size and by protocol. They are used extensively to link to WANs, and also as *collapsed backbones*. As Network layer devices, they can route between a wide variety of protocols: Token-Ring, Ethernet, FDDI, etc. Routers often have plug-in interface for many of these media.

When testing routed networks, you must be on the segment of interest or you will not see the traffic you are looking for. ICMP messages give a lot of information about what is going on in an IP-routed network. If you have compression turned on for your WAN links, you will not be able to view the traffic with a protocol analyzer (which can be considered a feature if you like secure networks). If you want to do accounting with the router, check that performance is still adequate when accounting is turned on. Routers can be complex to configure properly. They typically offer sophisticated SNMP-based management tools for configuration and monitoring.

**Servers.** Servers (Figure 2.16) can be used as gateways, security filters, proxy servers, and routers. Routing is a common function provided by IPX servers, but it can impact the server's performance. Server-based routing performs poorly compared to dedicated routers, but is fine for small networks. (This is not universally true, however; a portion of the Internet backbone runs on IBM RS6000 servers.) As firewalls, unless they are configured properly, servers can compromise network security; they must be treated with caution when used as an interconnect device. As routers, they can have unpredictable results in larger networks; an AppleTalk server, if started with routing turned on, will inform all other routers in the vicinity that all traffic should be forwarded to it.

**LAN switches.** LAN switches (Figure 2.17) are basically very fast multiport bridges. Full media bandwidth is supplied to each port on the device, and a very fast back-




Figure 2.16 A server.

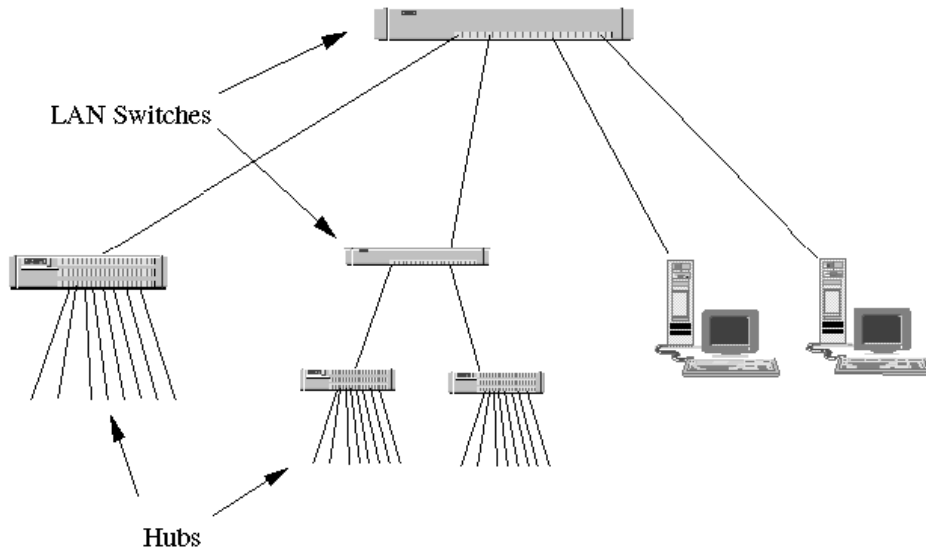


Figure 2.17 LAN switches in the network.

plane reduces or eliminates congestion. Unlike a hub, where bandwidth is shared, each connection to the switch has dedicated bandwidth at the speed of the media. LAN switches are used to increase performance. A typical configuration is shown with a switch connecting directly to two servers and aggregating traffic from a number of hubs. Mixed-media LAN switches are common, a typical device having a few high-speed ports (such as 100Base-X) for connecting to routers and servers, and many normal-speed ports (such as 10base-T) for other connections.

Like a learning bridge, a LAN switch develops a table of which addresses are associated with which ports. When a packet arrives, the switch examines the destination MAC address and forwards the packet only to the correct port. There are two methods of switching, *cut-through* and *store-and-forward*. Cut-through switching makes the routing decision as soon as the MAC address has been decoded. Store-and-forward switches read in the entire packet and check for CRC alignment errors before forwarding the packet. Cut-through advocates claim their method is faster, while store-and-forward advocates claim that they are more reliable.

From a testing point of view, since packets are only routed from source to destination, promiscuous monitoring must be done along the data path. Unlike a hub, where any port may be monitored to see traffic, you must connect in line between the stations being monitored. Some LAN switches have a special port to aid in network monitoring. LAN switches have the same spanning tree features and broadcast issues discussed for bridges, but not necessarily the filtering capabilities.

**ATM switches.** ATM switches generally fall into four categories: workgroup, enterprise, edge, and central office. ATM is aimed at being an end-to-end unifying technology, which is one of the reasons there is such a broad range of them. Workgroup switches are used to bring high-speed (greater than 155 Mbps) information, often



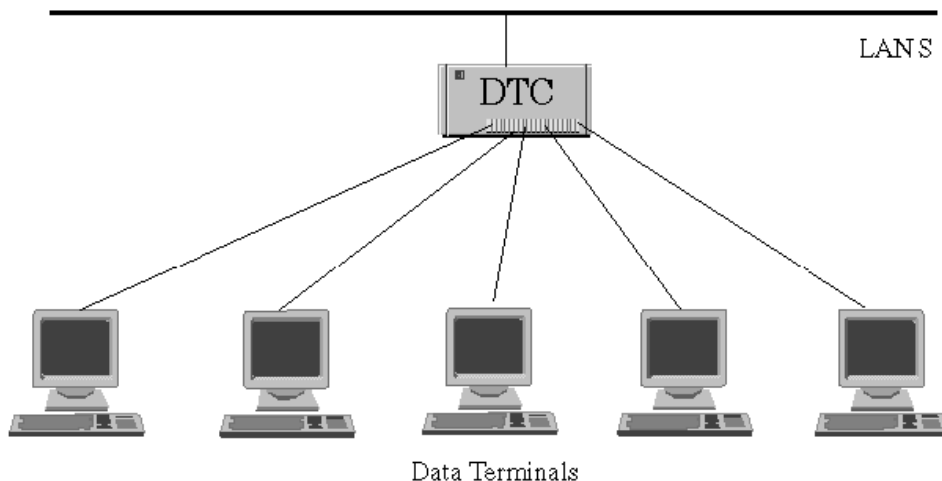
multimedia, to a desktop system. Enterprise switches are generally focused on creating faster backbones. Edge switches are used in the WAN, at the edge of the carrier networks, and central office or core switches are very large switches used to consolidate and transport traffic in the carrier networks.

ATM brings four main features to networking:

- A scalable network built around fast, hardware-based switching
- Bandwidth on demand
- Quality of service (QoS) guarantees
- Consolidation of voice, video, and data

As with LAN switching, in ATM there are no simple monitoring points to hook on an analyzer. Since ATM can transport many different protocols, test gear must be able to characterize the ATM at the Data Link layer (at very high speed), and also any traffic being carried above it. This could include encapsulated frame relay, TCP/IP, and others. ATM also has protocols for configuration and management that may require monitoring and analysis. For the foreseeable future, interoperability will continue to be a challenge for ATM switches.

**Data terminal concentrators.** Data terminal concentrators (DTC), shown in Figure 2.18, are designed to connect serial terminals to a LAN. These can be placed out in the workplace, simplifying the installation and wiring for the terminals. The terminal believes it is talking to a computer via a serial link, but the computer (usually a mini-computer) is receiving the data through its LAN port. As data terminals are replaced by personal computers and workstations, DTC use should decrease. Testing one of these links will require both a LAN analyzer and a serial line analyzer.



**Figure 2.18** Data terminal concentrators (DTCs).

**34 Introduction to Network Technologies and Performance**

**Concentrators or modular hubs.** Modular hubs are large, rack-mounted devices with fast backplanes that hold a series of different plug-in cards. They are the Swiss Army knife of private networks. There typically is a wide selection of cards that include most LAN and WAN functions. A typical configuration would include a redundant power supply, a number of 10Base-T hub cards (with many ports each), a LAN switch card, and a routing card. Other common interfaces include Token-Ring, FDDI, ATM, and others. Each of these cards will perform its associated functions of bridging, routing, switching, or repeating. Modular hubs are industrial-grade units, and thus more expensive than rack-and-stack solutions for small networks. Modular hubs normally have extensive SNMP-based network management systems for configuration and monitoring. The distinction between a large router that holds hub cards, and a large hub that holds router cards, can be somewhat blurry.

**2.4.2 WAN interconnects**

This section contains descriptions of and comments about devices commonly found on wide area networks. Table 2.11 includes the following information on WAN interconnects:

- Common name
- Device function
- Device limitations
- Designing for reliability
- Deployment hints
- Troubleshooting issues
- General comments

**DSU/CSU.** These devices provide the demarcation point between the public and private networks. Generally when a WAN service is purchased from a carrier, the provider is responsible up to the point where this device is located. Testing to this point provides a reasonable way to determine where a fault is occurring in the network. There is a range of products available in this category, from cheap basic units to more expensive units that are SNMP-manageable.

**Multiplexers.** Multiplexers allow you to aggregate different traffic streams up to a higher-speed link before you place it on the WAN. They accomplish this predominantly through time division multiplexing (TDM). A private operator may combine some voice traffic, some SNA traffic, and some other protocols before shipping them off to a remote site. Demultiplexing must be done at the receiving location. The advantage of multiplexing is that it allows the purchase of WAN bandwidth in bulk, and therefore lowers the network costs. Carriers routinely use multiplexers to combine and extract different traffic in the network infrastructure. Telcos are making large investments today in high-speed SONET/SDH multiplexers that provide this capability. It is not unusual for these to operate at speeds of 622 Mbps and beyond.

TABLE 2.11 WAN Interconnects.

Name	Function	Limitations	Design for reliability	Troubleshooting	Comments
DSU/CSU	Provides the connection between the local networks and the WAN provided by the carrier. Provides electrical isolation and loopback capabilities.	Very minimal management capabilities. Some will have SNMP MIBs with limited functions.	Static Addressing, fixed routes, WAN links owned and maintained by public carriers.	A good access point for testing.  Can be configured to loopback to check network continuity of the WAN or private side connection.	Typically delimits where the carrier network begins and the private network ends.  Check bit error rate of your carriers service from here.
Multiplexer	Provides for mixing different traffic streams into one network link. Can also add or selectively demultiplex one or more traffic streams.	Provides only capability to combine or uncombine traffic. No routing or protocol encapsulation capabilities.		Testing protocols typically requires demuxing first.  A typical test is to generate some traffic and verify muxing/demuxing properly by comparing input and output.	Typical usage will combine voice traffic from a PBX along with data traffic from a router and multiplex them into a T1/E1 line.
Modem	Takes a digital signal (typically RS-232) and allows it to be carried across an analog voice line.	Generally low-speed devices ranging from 9600 to 56K.	Buy from reputable vendors that implement standard command stack. Make sure modem has good diagnostics and data displays.	A serial data analyzer is needed for testing. Hook between the computer and the modem. Noisy lines can cause problems.	Now a home appliance, more robust modems (called "long haul" modems) were used for point-to-point wide area connections. Modems represent a security threat to the network.
Compression	Will take a digital traffic stream and compress the data, thereby lowering the cost of transmitting the data. Data is uncompressed at the destination. 50% compression typical.	Often proprietary compression schemes are used. Mixing and matching vendors can be an issue. Standards do exist.	Because compressed streams cannot be analyzed, it is often turned off so that faults can be quickly analyzed.	Compressed traffic is usually unintelligible to a protocol analyzer.	Can be found as a feature in a number of different wide area interconnects. Compression provides a low-cost means of encryption, though should not be used in situations requiring true encryption.
ISDN access devices	Will provide connections to ISDN lines from computers, ISDN	ISDN service is limited in certain areas. Basic rate limited to 128 Kbps.	Matching end-user home equipment to carrier equipment is vital.	ISDN-capable analyzer is required for serious troubleshooting.	ISDN used widely in Europe. After being pronounced "dead" by the press, BRI has now found great

TABLE 2.11 WAN LAN Interconnects. (Continued)

Name	Function	Limitations	Design for reliability	Troubleshooting	Comments
	phones, and other ISDN devices.	Primary rate available to DS1/E1 on an nx64 basis.	Check with your carrier before purchasing gear. Configuration can be difficult.		utility for telecommuting. The typical device has an Ethernet local connection and an ISDN WAN connection with built-in NT1.
Packet switch: X.25	Beaks the data stream into packets, which are sent to a X.25 network address via a public switched network.	Does not transport bursty, traffic well. Limited to DS1/E1 and slower.	X.25 is a very reliable protocol which does error checking at each packet switch in the link. Very good for noisy lines. This is a well established, stable protocol.	X.25-capable analyzer is required for serious troubleshooting.	X.25 is being rapidly displaced by frame relay. Midstream error correction makes X.25 very robust, though somewhat slow.
Packet Switch: Frame relay access device (FRADs)	Carries data in frames that are sent through frame relay switches to a destination address.	Currently limited to DS1/E1 speeds. DS3 emerging.	FR handles bursty traffic well, and is gaining broad acceptance carrying LAN and SNA traffic in U.S.	A frame relay-capable analyzer is required for serious troubleshooting. Frame relay performs no mid-stream error correction like X.25.	Frame relay has low overhead and can handle traffic bursts. It is one of the fastest growing WAN technologies in the 1990s.

There is also a category of multiplexer called an *inverse multiplexer*. These take a high-speed network and split it into a number of lower stream speeds for transport across the WAN. These streams are then reassembled at the destination location. This technique may be used if high-speed WANs are not available in a given location.

**Modems.** Modems today have become familiar devices due to the explosion of Internet access by PCs. There are many speeds and standards from which to choose. There is a de facto standard command set to control them. They can be purchased with many features, such as dial-back security features and data compression.

Modems represent a serious security threat to a network and should be managed carefully. For management, front-panel LEDs can give reasonable status indications. Most problems are caused by noisy lines or configuration. A serial line analyzer can be used between the computer and the modem to solve difficult problems.

**Compression units.** These add-on units can be used to save money on wide area traffic charges. They will typically lower the traffic on a link by up to 50 percent using a variety of standard and proprietary schemes. They have the side benefit of scrambling the data en route, which enhances data security but is no substitute for real encryption if it is needed. Compression capability is often built into routers. Once compressed, data must be decompressed before it can be interpreted by a WAN analyzer.

**ISDN devices.** Primary Rate ISDN devices at E1 data rates have been widely used in Europe and Japan for some time. ISDN allows the integration of voice and data. Basic Rate (BRI) usage has exploded in the US of late, typically using two 64 kbps data channels to facilitate home office connections to a corporation, and to gain faster Internet access. A typical home office device will accept the ISDN signal from the carrier (the NT1 is built in) and provide a 10Base-T Ethernet port to the user.

**WAN packet switches.** These devices come in two major flavors, X.25 and frame relay. X.25 packet switches have been in widespread use for over a decade. They grew out of the need to connect computers across the wide area. The packet switching function gave a good degree of flexibility to the network that previously required point-to-point connections. The X.25 protocol performs error detection and correction at each switch in the network path, which makes it very useful for areas with noisy lines (i.e., developing countries).

Frame relay is similar to X.25, but it does away with the per-hop error management, and thus is quite a bit faster. Frame relay handles bursts well, and is rapidly gaining wide acceptance as a means to transport LAN traffic across a WAN. Frame relay access devices (FRADs) are widely available, and sport features such as voice transport. Many frame relay switches today have ATM capabilities, and a number of carriers offer Frame relay services that are transported by ATM. It will not be unusual in the future to see LAN traffic being carried over frame relay that is in turn being transported by ATM. Protocol analyzers for this application will need to decode these encapsulations cleanly.

**38 Introduction to Network Technologies and Performance**

**2.5 Further Reading**

- Comer, Douglas A.* Internetworking with TCP/IP, Vol. 1: Principles, Protocols, and Architecture. (Englewood Cliffs, N.J.: Prentice Hall, 1991.)
- McDyson, David E., and Spohn, Darren L.* ATM Theory and Application. (New York: McGraw-Hill, 1994.)
- Miller, Mark A.* Troubleshooting TCP/IP Networks. (San Mateo, Calif.: M&T Books, 1992.)
- . Troubleshooting LANs. (San Mateo, Calif.: M&T Books, 1991.)
- Minoli, Daniel.* Enterprise Networking: Fractional T1 to SONET, Frame Relay to BISDN. (Norwood, Mass.: Artech House, 1993.)
- Naugle, Matthew G.* Network Protocol Handbook. (New York: McGraw-Hill, 1994.)
- Pearlman, Radia.* Interconnections: Bridges and Routers. (Reading, Mass.: Addison-Wesley, 1992.)
- Smythe, Colin.* Internetworking: Designing the Right Architectures. (Wokingham, England: Addison-Wesley, 1995.)

## Digital Telecommunications Basics

**Hugh Walker**

*Hewlett-Packard Ltd., South Queensferry, Scotland*

### 3.1 The Existing Telecommunications Network

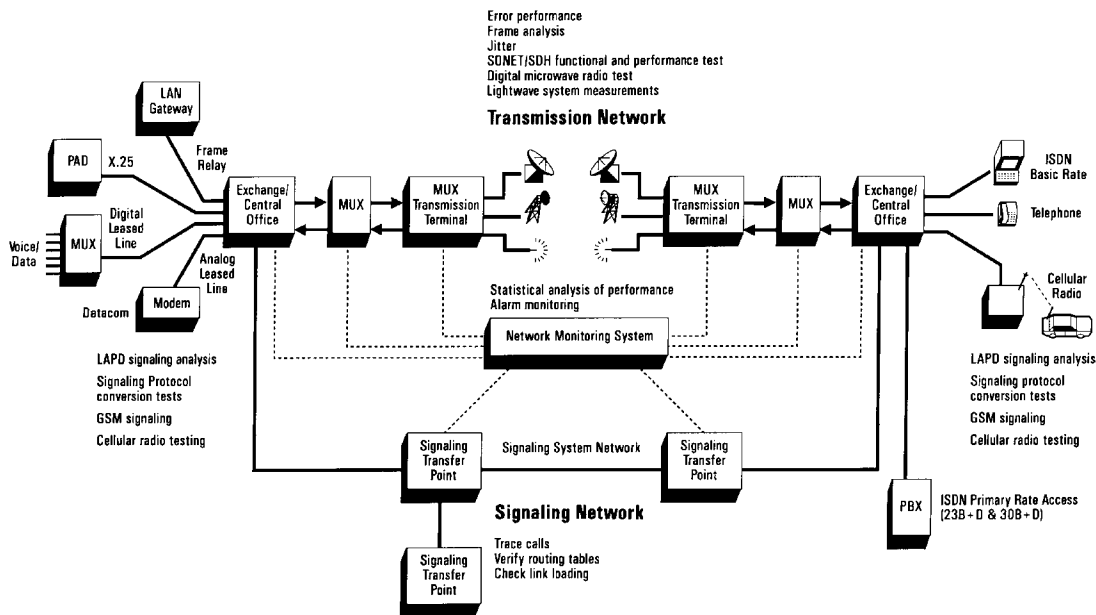
Telecommunications networks have existed for more than 100 years, but the rate of change has accelerated since the 1970s with the introduction of semiconductor technology and computers. With the rapid growth of services such as mobile telephone, cable television, and Internet and World Wide Web communication, it is easy to forget that we still rely on a great deal of equipment and fixed plant that was installed many years ago—and in the case of copper-wire local loop, perhaps decades ago.

In reviewing the elements of a digital communications network, it therefore is significant that many of today's developments are in fact an evolution of past network technology. A good example is the 3.1 kHz bandwidth voice channel, which in digitized form is the 64 kbps Pulse Code Modulation (PCM) signal, that is, 8-bit bytes at an 8 kHz sampling rate. PCM, invented by Reeves in 1937, was first used in the public network in 1962, but even the latest broadband communications equipment uses 8-bit bytes and a basic frame repetition rate of 125  $\mu$ s (8 kHz). In other words, the operating parameters of the network were defined for voice traffic, yet increasingly the network is being used for data communications. The circuit-switched telephone network is not optimized for bursty data traffic, but because of the very large investment in plant and equipment, the telecommunications industry has to find a way of adapting it to new uses.

Figure 3.1 shows a model of the existing telecommunications network. The three main areas are:

1. Customer premises or end-user environment, and the local loop access.
2. Switching and signaling (Central Office or exchange).
3. Multiplexing and transmission.

40 Introduction to Network Technologies and Performance



**Figure 3.1** A simplified block diagram of the telecom network showing the end-to-end connection between a variety of customer premises and the related measurements. The access network connects subscribers to the exchange switch, and the core network of transmission and signaling carries telecommunications traffic from source to destination.

**3.2 Customer Premises and Local Access**

Equipment at the customer premises ranges from a telephone handset in a house, to complex onsite systems such as the private branch exchange (PBX), LAN, X.25 network, and digital multiplexers for private network operations that might be found in a factory or business. Each of these end users connects to the switching center or exchange via one of the standard analog or digital interfaces called the *User Network Interface* (UNI). These interfaces include the traditional 2/4 wire analog telephone channel (the local loop described below), or a primary rate digital multichannel signal at 1.544 Mbps (a T1 line with 24 channels in North America) or 2.048 Mbps (E1 with 30 channels elsewhere).

**3.2.1 Local loop**

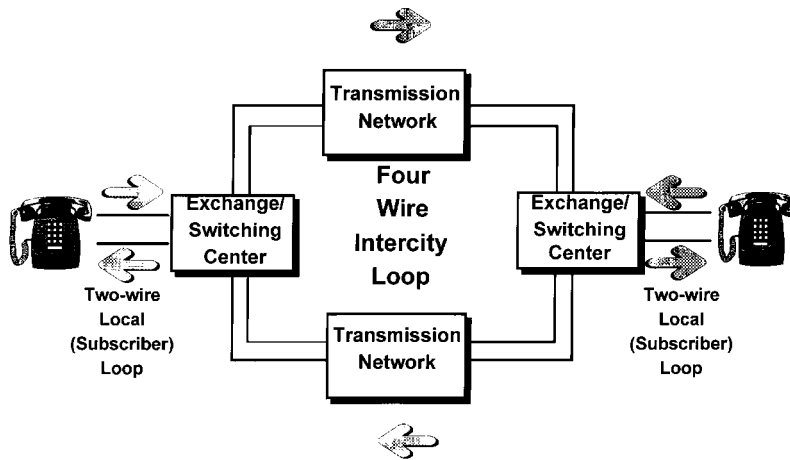
Figure 3.2 shows the simplest form of customer premises interface. The handset is connected to the Central Office (CO) or exchange by a pair of copper wires called a *2-wire circuit*. This circuit may be several miles long and is referred to as the *local loop* or *outside plant* (OSP).

Both directions of transmission are carried simultaneously using a hybrid transformer that acts as a directional coupler to separate go and return signals (see Figure 3.3). Isolation of 10 – 20 dB is obtained, depending on impedance matching. The 2-wire circuit also carries dc power (–48 V) to the handset for signaling. In the industry these are referred to as *wet lines*. Switching and inter-exchange transmission are

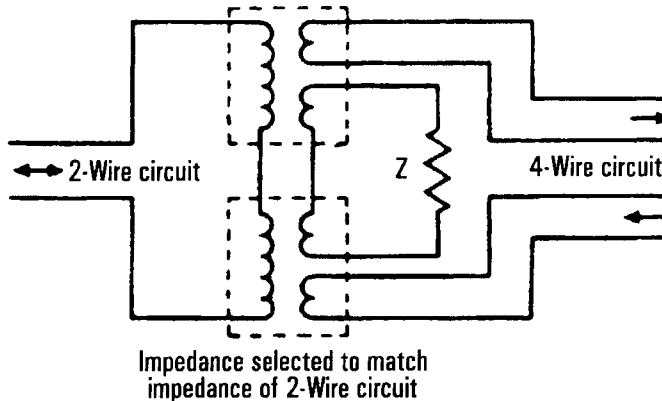


accomplished at the CO. At this point the two directions are handled separately using a 4-wire circuit.

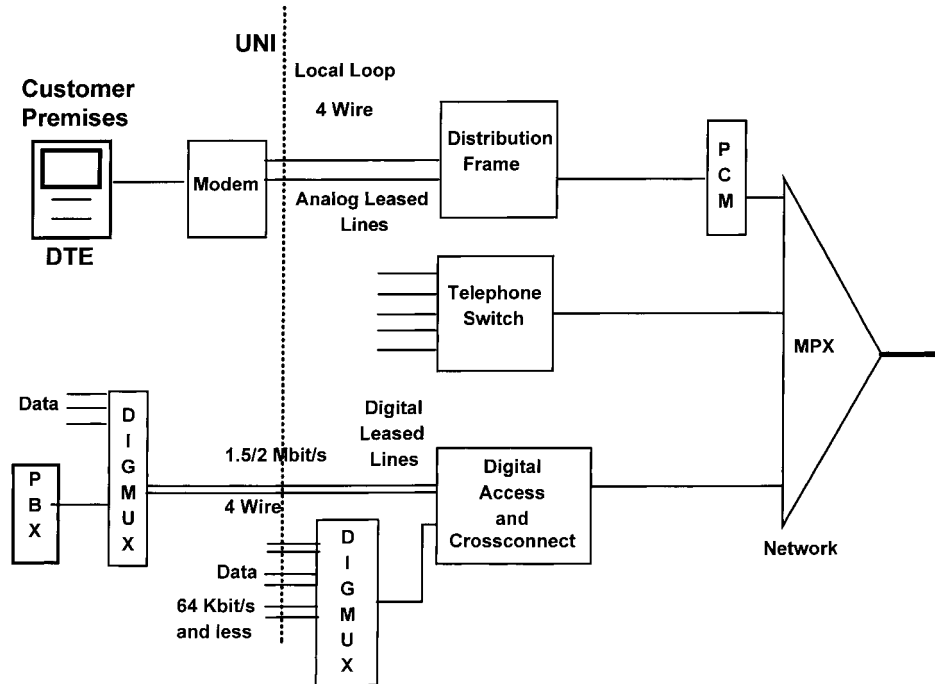
Because by 1996 there were approximately 800 million 2-wire local loops worldwide, growing at 5 percent annually, there is great interest in finding ways to use this embedded infrastructure for a range of new services such as data and video transmission. At the same time, the Plain Old Telephone Service (POTS), or Public Switched Telephone Network (PSTN), in which a simple handset is powered from the exchange, has the great advantage of robustness and reliability. Some developments, such as Asymmetrical Digital Subscriber Line (ADSL), superimpose the new digital access on the existing analog line. The latest technology can achieve bandwidths of several megabits per second over a 2-wire circuit that was originally installed to handle 3 kHz



**Figure 3.2** The local loop access is usually carried on a twin copper wire called a *2-wire circuit*, which handles both directions of transmission. In the core network, the two directions are carried by separate physical paths, termed a *4-wire circuit*.



**Figure 3.3** The hybrid transformer that separates the go and return signals on a 2-wire circuit at the telephone and exchange.



**Figure 3.4** Apart from the main telephone switched network, operators also provide analog and digital leased lines that are permanent connections between the customer's premises and the core network, bypassing the telephone switch. Leased-line access is economical for business customers with large volumes of telecommunications traffic.

telephony, by taking advantage of the higher bandwidth requirement from network to user (downstream) than in the upstream direction.

### 3.2.2 Leased lines

When an end user, such as a business, has a lot of data traffic, it is often more economical to lease a private point-to-point circuit from the network operator on a permanent basis. Leased lines use the common multiplex and transmission system but bypass the telephone switch. Both analog and digital leased lines are available (analog or digital in the local loop), as shown in Figure 3.4.

A leased line offers several advantages to the end user. It guarantees circuit quality, thereby allowing higher-speed data modems to be used on analog lines. Furthermore, because it is a permanent connection that bypasses the telephone switch, it eliminates the problem of gaining access during busy periods. With further advances in modem technology providing dial-up speeds of 28.8 kbps and above, and the improved quality in the circuit-switched network, the speed advantages of analog leased lines have largely been overtaken, particularly with the advent of Integrated Services Digital Network (ISDN).

Digital leased lines provide direct access to the digital network. The most popular are at the primary rate of 2.048 Mbps (E1) or 1.544 Mbps (T1 in North America).

Both use a 4-wire circuit with regenerators. The end user connects the line to a digital multiplexer, which also can act as a concentrator for telephone and data traffic. For lower-speed lines (64 kbps or less), such as Digital Data Service, the digital multiplexer may be sited at the exchange.

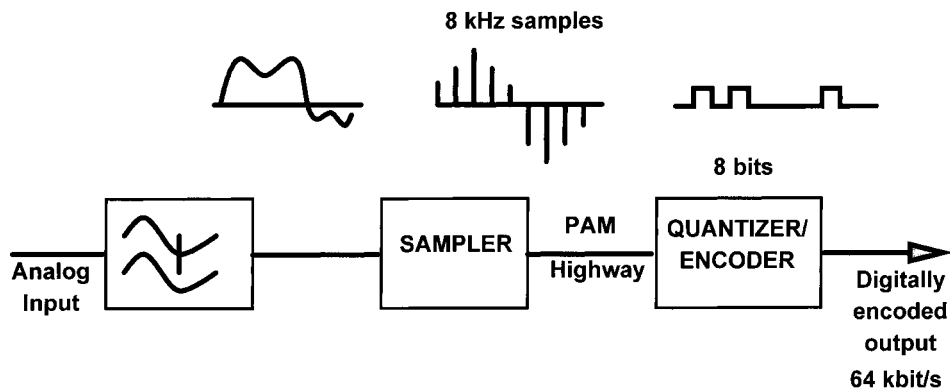
To reconfigure leased lines, analog lines are normally terminated on a wiring or distribution frame. Digital lines are connected to a Time Division Multiplex switch called a *digital crossconnect*.

### 3.2.3 Pulse Code Modulation (PCM)


In the digital communication network, the analog voice signal is converted to digital at the earliest point possible. In the existing network, most local loops are analog, but when the access line is terminated at the exchange, the analog voice signal is immediately converted to digital PCM. In private or ISDN networks, the voice signal will be digitized either at the telephone handset or in the customer's premises. Thus PCM in the modern network is really part of the access technology and over time will move further towards the end-user equipment in a fully digital access network.

Pulse Code Modulation (PCM) is the mechanism used for converting the analog voice channel to a digital signal. Figure 3.5 shows the three stages of the process. With a maximum input frequency of 3.4 kHz, Nyquist's sampling theorem indicates a minimum sampling rate of 6.8 kHz. In fact, practical systems use a stable ( $\pm 50$  ppm) 8 kHz sampling frequency. The filter before the sampler removes signals above 4 kHz, which could cause aliasing.

**Quantizing error.** The final stage of the PCM process is the quantization of the sampled values into one of several discrete levels. This process results in a quantizing error (see Figure 3.6). In this example, the analog signal samples are encoded as one of eight possible levels midway between the decision thresholds. (The analog signal could lie anywhere between three thresholds and be encoded as the same digital value.) Therefore there is a maximum disparity equal to  $\pm 0.5$  times the quantizing interval between



**Figure 3.5** The PCM process digitizes the analog telephone signal. To prevent aliasing, the analog bandwidth is restricted to a maximum frequency of 3.4 kHz before being sampled with a very stable 8 kHz clock. These samples are then digitized into an 8-bit coded word, resulting in a 64 kbps signal.



**Figure 3.6** Digitizing a continuously variable analog signal always results in quantization error or noise, because the discrete digital values never exactly match the analog signal.

the true value of the signal and its quantized level. This random difference is called *quantization noise* or *quantization distortion*. The smaller the signal, the more severe the problem.

Quantization error can be reduced by using smaller intervals and encoding each sample with a longer digital word. (A 12-bit ADC would be required for good quality.) For a given sampling rate, however, this technique increases the bit rate of the system. A better solution is to *compand* (compress and expand) the signal to improve the signal-to-noise ratio at low levels.

**Compression and expansion.** Figure 3.7 shows a typical companding curve that encodes the sample value (horizontal axis) into  $\pm 128$  levels using the logarithmic curve. Notice how many of the levels are used for small signals to maintain an acceptable signal-to-noise ratio. The  $\pm 128$  levels are represented by an 8-bit word (a *byte* or *octet*). As bytes are produced at the sampling rate of 8 kHz, the result is a bit stream at the familiar 64 kbps. Slightly different companding equations are used in the U.S. and Europe. However, all aspects of the PCM process are fully standardized by the ITU-T (Recommendation G.711, first issued in 1972).

Figure 3.8 shows the variation of the signal-to-noise ratio as a function of signal level. The linear portion at low signal levels is equivalent to the Least Significant Bits (LSBs) of a 12-bit analog-to-digital (A/D) converter. The flat portion (constant signal-to-noise) is the companded 8-bit conversion.

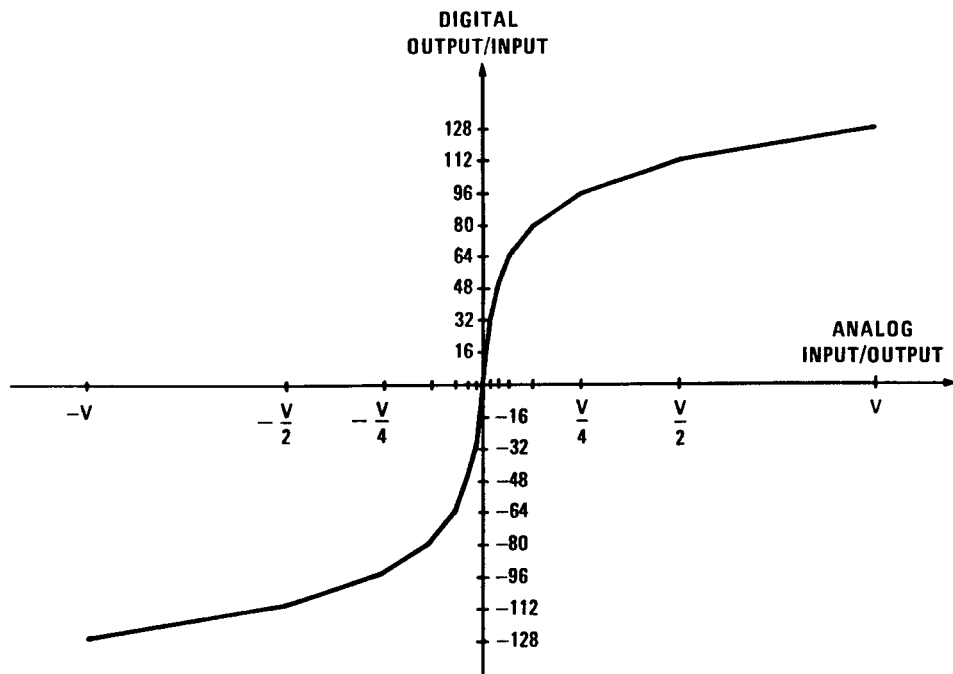
Companding is the simplest technique to reduce the data rate of digitally encoded analog signals such as telephony, video signals, and audio. More sophisticated tech-

niques take advantage of the dynamics of the signal by noting that successive samples are closely correlated, and that, for this reason, only differences need to be transmitted. Examples of these techniques are *adaptive differential PCM (ADPCM)* and *continuously variable-slope Delta Modulation (DM)*.

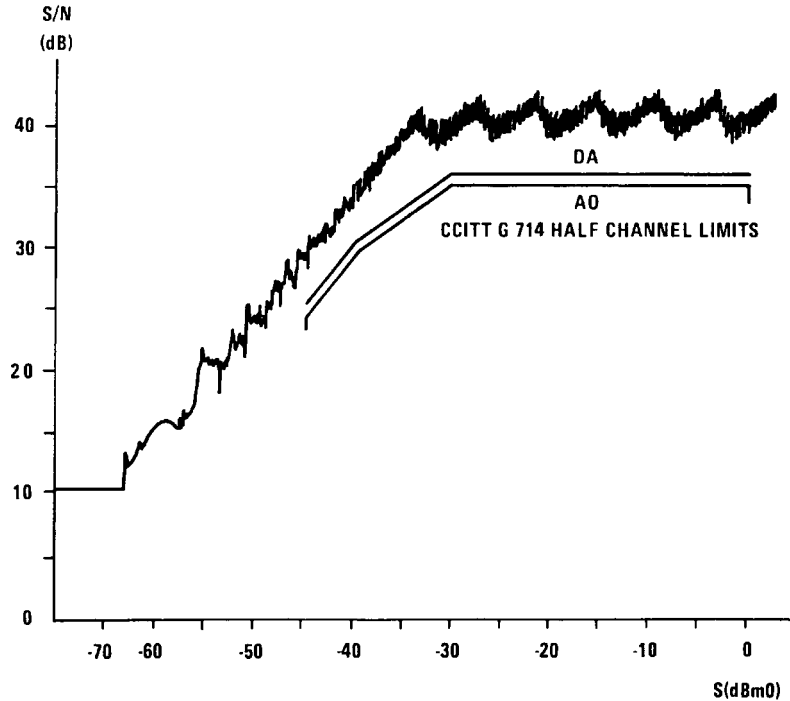
PCM at 64 kbps is now a world standard. Unfortunately, however, there is a great deal of redundancy inherent in the coding. Good quality can be obtained with 32 kbps, and digital cellular radio uses 16 kbps or less for speech transmission. Voice messaging typically uses 16 kbps as well, requiring much more complex processing of the signal. ADPCM at 32 kbps is standardized in ITU-T Recommendation G.721, and 64 kbps ADPCM (which gives higher-quality speech with 7 kHz bandwidth) is standardized in ITU-T Recommendation G.722.

### 3.2.4 Analog data modems


Digital data represents an increasingly large percentage of telephone network traffic. Although much of the transition to ISDN is already underway, most local loops and switching services remain analog. Modem traffic has increased considerably with the popularity of Internet access from home personal computers. The difficulties in transmitting digital data over analog lines are shown in Figure 3.9. The standard telephone channel has a bandwidth of 300 to 3400 Hz. Binary data, on the other hand, has a  $(\sin x)/x$  spectrum extending up from dc. To transmit data through a



**Figure 3.7** To minimize the effects of quantization error, nonlinear coding is used so that more levels are used for small signals (where quantization effects are most pronounced) than large signals. This is called *companding* or *compression /expansion*.



**Figure 3.8** For 8-bit encoding, A-law PCM encoding produces an S/N ratio vs. signal level graph like this. For low-level signals the coder operates as if it were a 12-bit rather than 8-bit encoder; for higher levels, the compander maintains a constant signal-to-noise ratio by progressively compressing the signal. A-law is used in the international market; North America uses the mu-law coding, which is slightly different.



**Figure 3.9** Digital data is unsuitable for direct baseband transmission through the telephone network because of the limited bandwidth assigned to telephony. Modems modulate the data onto a carrier signal with a spectrum that matches the voice-channel bandwidth.

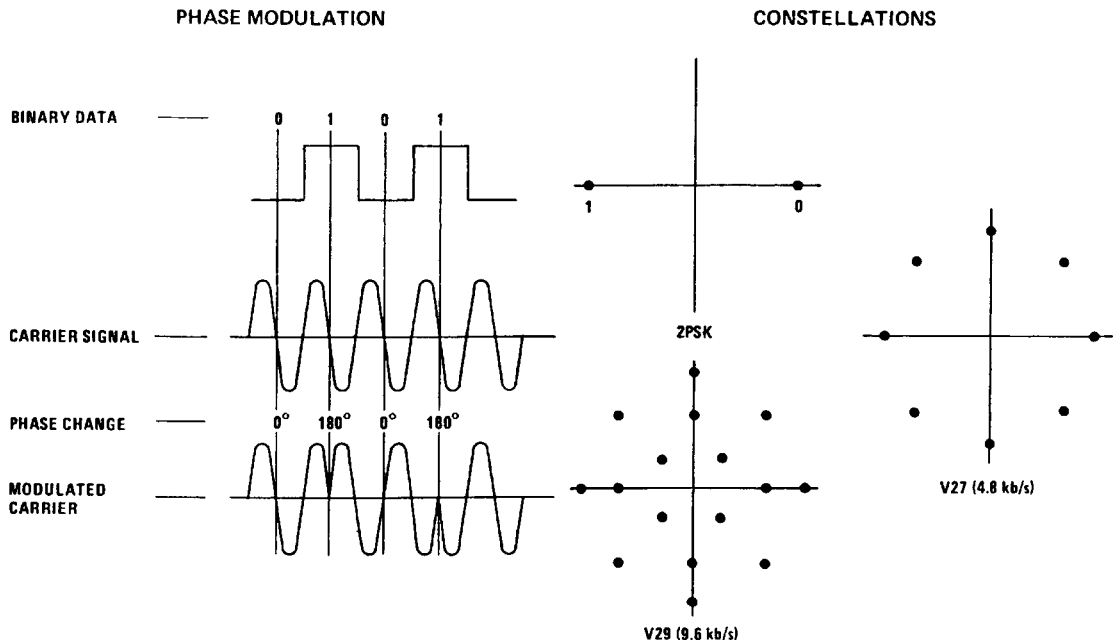
telephone channel, therefore, it must be modulated onto a carrier frequency (e.g., 1700 or 1800 Hz) so that its spectrum matches the channel bandwidth. This is the function of a modem. Paradoxically, as soon as this analog signal reaches the exchange central office, it is converted to a PCM digital signal for transport through the network.

**PSK Modulation.** Most higher-speed modems use *phase shift keying* (PSK) to modulate the carrier. In PSK, different phases of a waveform represent different bits (Figure 3.10). In its simplest form, 2-PSK, two phases are used (usually 0 and 180 degrees) to represent the bits 0 and 1. By increasing the number of shifts used, more information can be represented by each shift. For example, 8-PSK uses eight phases. Because  $8 = 2^3$ , each shift therefore can represent not just one but three bits at a time (000, 001, 010, 011, 100, 101, 110, and 111). An 8-PSK modem can provide a data rate three times that of a 2-PSK modem. Over standard phone lines, this rate works out to 4.8 kbps.

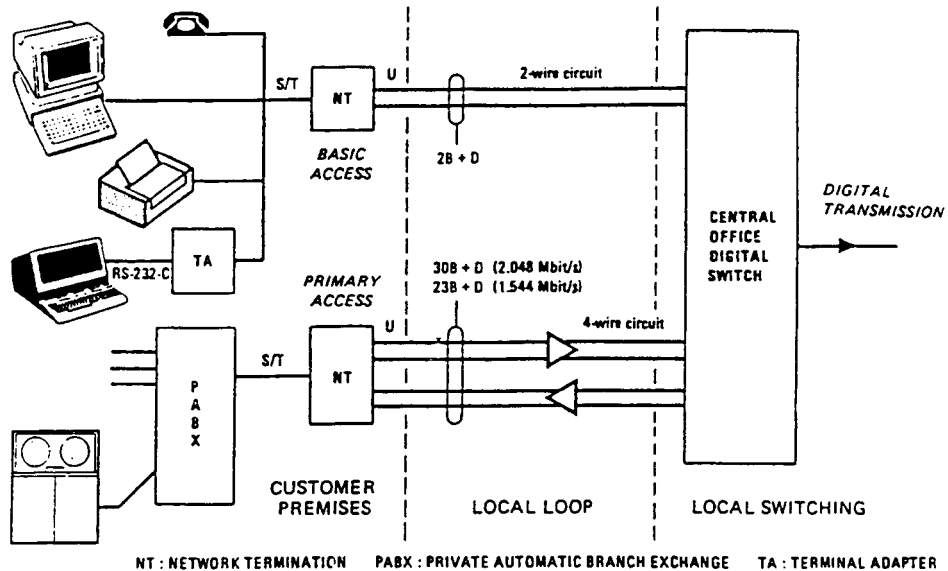
**QAM Modulation.** By shifting both phase and amplitude, using a technique called *Quadrature Amplitude Modulation* (QAM), even more shift options can be created and further increase the number of bits that can be represented by each shift. The 16-QAM technique allows four bits to be represented by each shift and provides data rates of 9.6 kbps over standard phone lines. Other QAM levels are used to achieve the data rates used by digital radio transmission systems (discussed below). The V.29 constellation in Figure 3.10 is an example of QAM modulation.

### 3.2.5 ISDN

The Integrated Services Digital Network (ISDN) is an evolving telecommunications network standard that pulls a wide variety of consumer services into a single, high-speed access package. An ISDN supports data, voice, image, facsimile, video, and telemetry over the same set of wires (see Figure 3.11). Its main purpose, however, is to give users ready access to digital services.



**Figure 3.10** A variety of modulation schemes have been developed for modems; the more complex the scheme, the higher the bit-rate that can be handled within the restricted telephone channel.



**Figure 3.11** ISDN provides direct digital subscriber access to the network over existing 2-wire and 4-wire circuits. Basic rate access effectively replaces the familiar analog telephone connection, while primary rate access over a 4-wire loop is used by business customers. ISDN combines both voice and data services in a common access, with sophisticated user-network signaling.

The ISDN standard divides digital services into three categories: *bearer services*, *teleservices*, and *supplementary services*.

**Bearer services.** Bearer services support the transfer of information (data, voice, and video) without the network knowing or manipulating the content of that information. Bearer services act at the first three layers of the OSI model. They can be provided using circuit-switched, packet-switched, frame relay, or cell relay facilities.

**Teleservices.** Teleservices add a level of complexity to the bearer services. They act at layers 4–7 of the OSI model and accommodate more complex data transfer needs than those that use the bearer services alone. Teleservices include telephony, Telex, teleconferencing, Telefax, etc.

**Supplementary services.** Supplementary services add an additional layer of functionality to the bearer and teleservices. These services include reverse charge, call waiting, message handling, etc.

Each service type has its own transmission requirements. Telephone services use connection-oriented systems that transfer a constant-bit-rate stream of data at low bandwidth (64 kbps) with a controlled lag and delay variance. Computer data networks vary in characteristics. Some are connectionless and some are connection-oriented. The bandwidth required varies greatly. Typically data services are much more tolerant of delay variation than any other communications system, but more delay-insensitive. Cable TV (CATV) is connectionless and delay-intolerant.



**Access to ISDN.** Much of the ISDN standard deals with interfaces between the subscriber and the network. To allow users the most possible flexibility, three types of channels and two basic interfaces are defined. The channels are labeled B, D, and H. The interfaces that use them are called BRI and PRI.

**B, D, and H channels.** A bearer channel (B channel) is the primary user channel. It is designed to carry data in full duplex mode end-to-end and has a data rate of 64 kbps.

A data channel (D channel) carries control signaling, not user data (although it can be used for low-rate data transfer, telemetry, and alarm transmission). It is responsible for call setup, maintenance, and termination between the user and the network at either end of a connection. A D channel can be either 16 or 64 kbps, depending on the interface.

Three hybrid channels (H channels) are defined to support high data-rate applications such as video and teleconferencing. The first, H0, has a data rate of 384 kbps; the second, H11, a data rate of 1536 kbps; and the third, H12, a data rate of 1920 kbps.

**BRI and PRI interfaces.** There are two kinds of ISDN local loop access, Basic Rate Interface (BRI) and Primary Rate Interface (PRI).

The BRI specification calls for a 192 kbps digital pipe consisting of two B channels (64 kbps each), one 16 kbps D channel, and 48 kbps of operating overhead. This interface is geared to the needs of residential subscribers and small businesses. In most cases, basic rate access can be implemented on existing 2-wire circuits (local loop).

The PRI specification requires a digital pipe of 1.544 Mbps divided among 23 B channels, one 64 kbps D channel, and 8 kbps of operating overhead. Outside North America, the PRI standard is 2.048 Mbps, which provides 30 B channels and one 64 kbps D channel and 64 kbps operating overhead. PRI is intended for larger offices and can support LAN and PBX traffic. Primary rate access is carried by metallic local loop using conventional 4-wire PCM technology and framing (discussed in section 34.3).

### 3.3 Central Office or Exchange

The switching center is usually called an *exchange* or *Central Office (CO)*. It contains a variety of equipment, the most important elements of which are the *circuit switch* for interconnecting telephone subscribers, and a *packet data switch* for supporting packet switching technologies such as X.25. Other equipment is designed to handle the so-called *private circuits*, which are semipermanent connections leased by business users to bypass the normal circuit switching and provide higher-quality communications, described earlier.

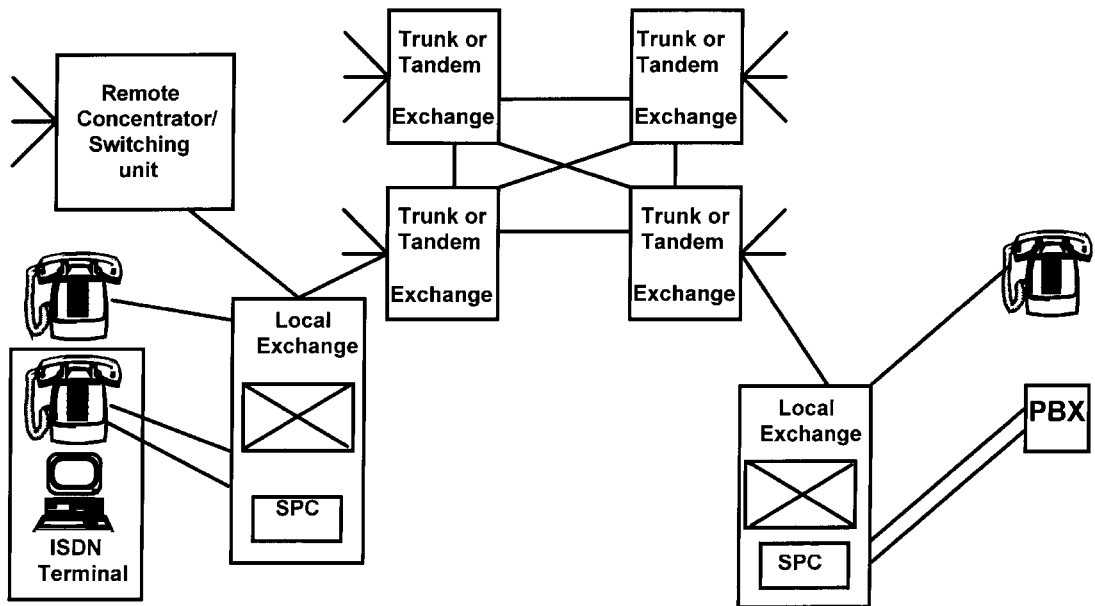
Whenever there are multiple devices with a need to communicate one-on-one, there is also the problem of how to connect them to make that communication possible. Point-to-point links are impractical and wasteful when applied to larger networks. The number and length of the links requires too much infrastructure to be cost-efficient, and the majority of those links would be idle for most of the time.

A better solution is *switching*. Switches are hardware and/or software devices capable of creating temporary connections between two or more devices linked to the switch but not to each other. In a switched network, some of these nodes are connected to the communicating devices. Others are used only for routing.

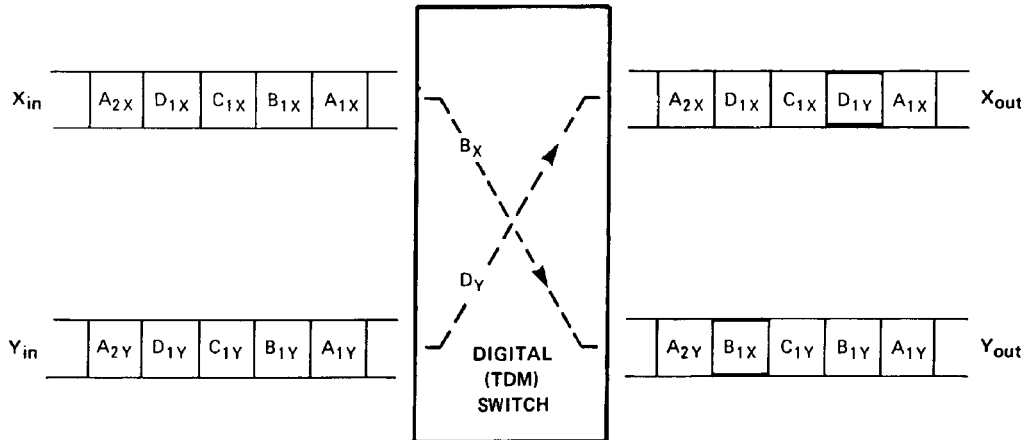
Most networks use a two-level approach to switching, *local* and *trunk* (Figure 3.12). The local loops, which can be analog or digital, connect subscribers to the Central Office, where local switching occurs. Higher levels of switching occur in trunk, or long-haul, networks. Trunk switches are fully interconnected by light-wave or microwave transmission systems. Traffic in the trunk networks is more concentrated than that in the local exchanges. Typically there are more than ten times as many local switches as trunk switches (a figure that also indicates the amount of telecommunications traffic that is completed within the local exchange area).

In general, telephone switching is the process of automatically connecting an input to an output in a flexible manner to enable the routing of telephone calls from an originating location to any one of myriad destinations. Digital switching is the switching of telephone channels contained within a PCM/TDM multiplexed stream without the need to demultiplex down to the original 3.1 kHz analog voice format.

The *primary multiplex* is closely associated with switching. It uses TDM to multiplex the individual telephone channels to the primary rate of 1.544 Mbps or 2.048 Mbps. A modern digital switch operates directly on these multiplexed signals, making it the standard interface in an exchange. For example, a cellular radio base station typically connects to the exchange at this primary rate.



**Figure 3.12** Typically telecommunication networks can be divided into two sections: the local access and switching associated with a town or city, and the main long-distance trunk transmission and switching network that is accessed via the local exchanges.



**Figure 3.13** Digital switching involves the reordering of timeslots in the digital bit stream. This requires incoming data to be stored for a short time in a buffer store, ready to be inserted in the right position in the outgoing stream. This is called *time-switching*.

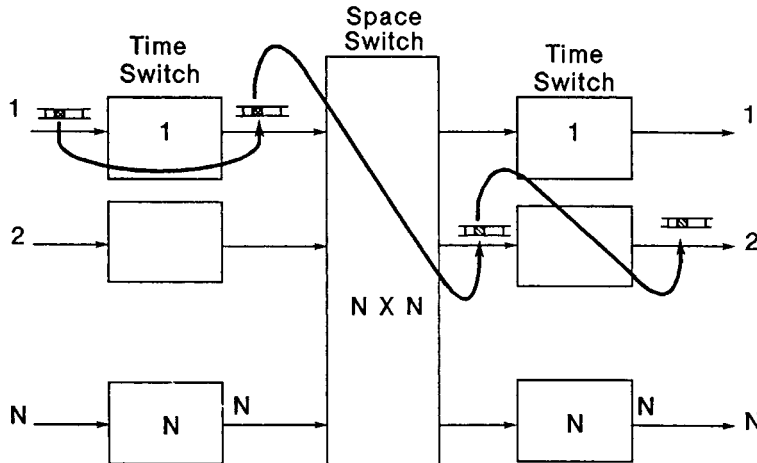
With a digital switch, as many as 100,000 to 200,000 subscriber lines may be connected via subscriber line cards. The switch is controlled by a computer called *stored program control* (SPC). The switching hardware controlled by the computer may not be physically all in the same building; modern systems allow for remote concentrators or switches under the control of the exchange. In addition, with digital systems it is sometimes economical to route even local calls back to a central exchange. This technique is referred to as *back hauling*.

Figure 3.13 shows an example of digital (TDM) switching. Within the switch, the data bits contained within the B timeslots of the incoming X stream are transferred to the D timeslots of the outgoing Y stream. In the same manner, the data bits contained within the D timeslots of the incoming Y stream are transferred to the B timeslots of the outgoing X stream. If the TDM signals are carrying PCM-encoded voice channels, then the above process is the equivalent of switching call  $Bx$  from incoming route X to outgoing route Y (and vice versa), without first reverting back to the original analog format. Each timeslot represents a separate physical telephone channel.

**Note:** Switching data between timeslots usually involves storing the data bytes in a memory for a short time, ready to be inserted in the outgoing data stream at the appropriate timeslot. TDM switching therefore has an inherent delay. ITU-T Recommendations Q.41 and G.114 give details of acceptable delays through network elements. For a digital switch, an average delay not exceeding 450  $\mu$ s is recommended.

### 3.3.1 Practical digital switch

In theory, all switching could be achieved by timeslot interchange. In reality, however, systems are limited by the speed of their logic circuits. A commercial digital switch employs a combination of *time-switching* and *space-switching*, or *matrix-switching*. Figure 3.14 shows the basic structure and operation of a Time-Space-Time switch commonly used in a large digital exchange.



**Figure 3.14** A practical telephone switch usually uses a combination of time and physical space switching to accomplish the connection of up to 100,000 customers.

A large Central Office (CO) or exchange might handle 100,000 subscriber lines with perhaps 60,000 outgoing lines or trunks to other local switches or to the long-distance network. The traffic capacity of a switch is measured in *Erlangs*.

If you use your telephone 10 percent of the time, you generate a traffic level of 0.1 Erlangs. A typical domestic subscriber might generate 0.02 to 0.05 Erlangs, while an office telephone might generate 0.1 to 0.25 Erlangs. With the considerable amount of data traffic expected from future ISDN applications, a figure of 0.25 Erlangs per line is often given as the norm. Assuming uniform, random traffic density on all lines, that gives a total capacity of 25,000 Erlangs in a typical 100,000-line switch. The rapid increase in Internet access in the 1990s has put a severe load on telephone switches that were dimensioned primarily for voice traffic. A typical voice call lasts about 3 minutes, whereas Internet accesses average 20 minutes and can last several hours.

A major measure of performance is the ability of the stored program control (SPC) computer to accept and set up a large number of telephone calls. This figure is measured in *Busy Hour Call Attempts* (BHCA). A fully equipped modem switch may handle one million BHCA, meaning that each subscriber line makes an average of 10 calls per hour.

### 3.3.2 Packet switching

In a packet-switched network, the user's data stream is segmented into *packets*, or *frames*, at the Data Link and Network layers. Each packet contains the destination address, making it possible to send packets separately instead of in a single stream. Figure 3.15 shows the frame and packet structure of an X.25 packet and how each part relates to the OSI model (See Chapter 5, Section 5.6). At level 2, the frame is delimited by start and stop flags. The Frame Check Sequence is a Cyclic Redundancy Check (CRC) calculated on the bits in the frame. If an error is detected, the frame is retransmitted to provide error correction. The Information Field, at level 3, contains

a Header and the User Data Field. The Header contains the Logical Channel Identifier (LCI), which defines the virtual circuit for the duration of the call and routes the packet to the correct destination. The LCI is defined during the call setup procedure. (Virtual circuits are discussed in a subsequent section.)

Packets are passed from source to destination via a number of intermediate nodes. At each node, the packet is stored briefly (store-and-forward transmission) and any necessary overhead processing (such as error detection) is completed before it is sent on.

Figure 3.16 shows a public packet-switched network (PPSN) that consists of interconnected packet-switched exchanges (PSEs) providing multiple routes between any two points. Multiple routes allow failed or congested sections to be bypassed and make packet-switched networks extremely reliable.

**Datagram switching.** Packet-switched services can take either of two forms: *datagram* or *virtual circuit*.

In datagram switching, each packet is treated as an independent unit with no connection to other packets in the transmission. Packets with the same destination therefore do not need to follow the same route. Instead, each can be sent by whichever path is the most efficient at that instant. This flexibility is the major advantage of the datagram approach. Disadvantages are the amount of overhead required by each packet (each packet must carry complete addressing and sequencing information), and the amount of processing required at each node (new routing decisions must be made for each packet at every node).

**Virtual circuit switching.** In virtual circuit switching, a route (called a *virtual circuit*) is established before any data packets are sent. All of the packets of a transmission are

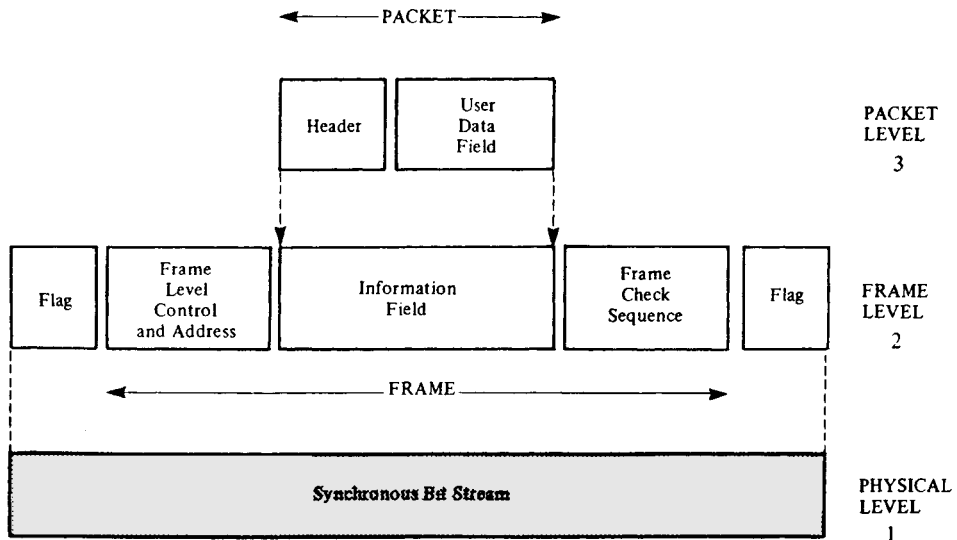



Figure 3.15 The structure of the X.25 packet frame.

## 54 Introduction to Network Technologies and Performance



**Figure 3.16** An X.25 packet network incorporating packet switch exchanges (PSE) and packet assemblers and disassemblers (PAD).

then sent, in order, along that same route. The virtual circuit is kept in place until terminated by either the sender or receiver. This route changes from transmission to transmission—unlike a dedicated circuit, which is the same for every exchange between two given nodes—but does not change during a transmission. Advantages of this method include lower overhead, because only the first packet needs to carry complete addressing information; subsequent packets can be identified by an abbreviated address or short *virtual channel identifier* number (VCI). Virtual circuit services also tend to be faster because routing decisions are made once, at setup, instead of being required for each packet at each node. In addition, packets arrive in sequence. Not only do they not need to be reordered by the receiver, but the receiver can detect lost packets and request retransmission over the same route.

The highly successful X.25 packet-switched network is an example of virtual circuit packet switching. Key features of this service are: variable packet length; in-band signaling (call setup, control, and release packets are sent in the same channel and virtual circuit as the data packets); and inclusion of flow- and error-control mechanisms at both layer 2 and layer 3. Each node checks, acknowledges, and, if necessary, requests retransmission of errored packets.

**Connection-oriented and connectionless services.** X.25 is defined as a *connection-oriented* service. All packets in a call are transmitted in sequence through the defined virtual channel, allowing error correction by retransmission. In a connectionless service, each packet is sent independently with a full address label. In this case, packet sequence is not guaranteed and level 2 (node-to-node) error correction is not provided. Most LANs provide connectionless service.

The high level of reliability offered by X.25 is important when there is a high probability of errors being introduced during transmission (such as the existence of links that are highly susceptible to noise). With advances in media quality, in particular optical fiber, however, this much overhead becomes redundant. Two outgrowths of X.25 are providing major improvements in efficiency: frame relay and cell relay, also known as *Asynchronous Transfer Mode (ATM)*.

Both are connection-oriented services using virtual circuits. While X.25 is designed to work at about 64 kbps, however, frame relay strips out most of the error-control overhead and achieves data rates of up to 2 Mbps. ATM also provides only minimum error control but makes an additional change. Whereas frame relay *frames*, like X.25 packets, can be of variable length, ATM uses fixed-length packets called *cells*. This change further reduces the necessary overhead and not only improves routing speed, but simplifies multiplexing and demultiplexing. The result is data rates in the hundreds to thousands of Mbps.

### 3.3.3 Signaling

In order to route traffic through a network, it is necessary to send messages to the switches so that the right connections are made (such as for establishing a virtual path). These messages are called *signaling*. In an ordinary telephone call, signaling originates from the handset. Signaling examples include connection establishment, ringing, dial tones, busy signals, and connection termination.

Early technology uses dial pulses created by interrupting the dc supply for the exchange. This process has its origins in the actuation of step-by-step uniselector switches in the electromechanical exchanges. While a modern digital exchange can still accept this type of signaling, the preferred method is now *dual-tone multifrequency signaling (DTMF)*. This is faster and can be used while the call is in progress, for example to control a voice-messaging system.

Generally, DTMF dialling would also be used by customer premises equipment such as computers, data terminals, and PBXs. High-capacity customer premises signaling may use the Common Channel Signaling capability of the ISDN D channel at basic or primary rate (see section 3.2.5). Once the call has been received and processed by the local exchange, the ongoing call setup and control, including billing, is handled by the network's own internal signaling system. This can be either *Channel Associated Signaling* or *Common Channel Signaling*.

Existing systems may use Channel Associated Signaling (CAS), whereby the signaling instructions are sent through the same circuit as the voice and data traffic. This mechanism is also called *in-band signaling*. CAS systems are limited by available rates and the fact that no signaling is possible while the call is in progress. To satisfy the requirements of ISDN for comprehensive and interactive signaling, a second system is

preferred: Common Channel Signaling (CCS). The high speed of CCS greatly reduces call setup time and improves network efficiency, particularly for calls of short duration.

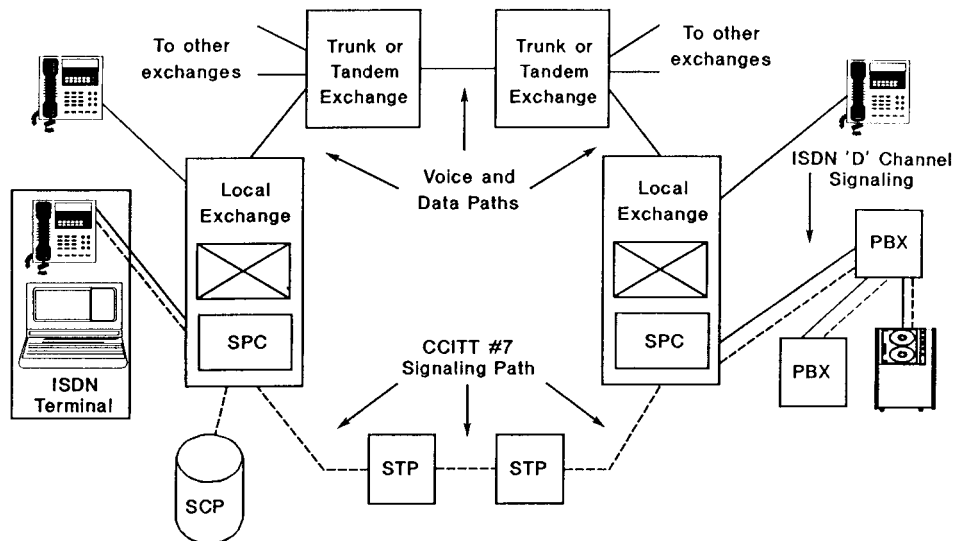
Figure 3.17 shows the structure of a CCS system. In the figure, the solid lines represent subscriber use and the dotted lines represent links reserved for signaling use.

In CCS, a separate signaling path is used to interconnect the computer systems that control the digital switches. The call therefore goes by one path, while the setup and monitoring information goes by another. This special-purpose data network is designed according to ITU-T standards for signaling system No. 7 (SS7). It uses robust error correcting data communications protocols similar to X.25 packet switching, with SS7 messages called *Signaling Transfer Points* (STP). By separating the traffic and signaling systems, signaling can operate completely independently of the traffic being carried in the normal telecommunication system and, therefore, be more efficient.

### 3.4 Multiplexing and Transmission

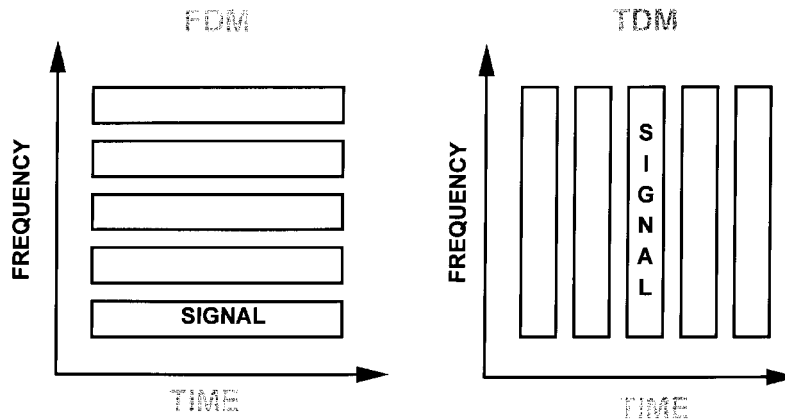
In the early days of telecommunication, every connection was carried by a separate physical circuit. With current technology, however, many thousands of channels are multiplexed together on a common optical fiber or radio carrier. Without multiplexing, long-distance and international telecommunications would be impossibly expensive. In the early days, a call from London to Edinburgh (400 miles) took place over wires that weighed 800 lbs. per mile. While you spoke, you had 250 tons (250,000 kgs) of copper all to yourself.

Multiplexing enables multiple signals to share a path simultaneously. This sharing is accomplished by dividing up the capacity of the path and allocating each portion



**Figure 3.17** The common channel signaling network interconnects the computers and databases that control the switching of customer traffic in the network. This data communications network is quite separate from the traffic-carrying voice and data paths.





**Figure 3.18** Multiplexing allows the transmission of many independent telecommunications channels through a common path such as a fiber optic cable or radio channel. The time-frequency domain can be divided either along the frequency axis, for frequency division multiplexing (FDM), or along the time axis, for time division multiplexing (TDM).

to a different signal. The path portions allocated to a single signal constitute that signal's channel. Every transmission path is considered to have two dimensions: *bandwidth* (frequency) and *time*. Multiplexing methods effectively slice up the path in one or the other of these dimensions (Figure 3.18).

Analog transmission uses *Frequency Division Multiplexing* (FDM). In FDM, the individual signals are allocated a portion of the frequency spectrum, i.e., a frequency *band*. Each signal has unlimited use of its band in terms of time, but the transmitted signal spectral components must never lie outside the allocated frequency band.

Digital transmission uses *Time Division Multiplexing* (TDM). In TDM, the individual signals are allocated the entire frequency bandwidth, but only for a limited portion of time, called a *timeslot*. There are two types of TDM, *synchronous* and *asynchronous* (statistical).

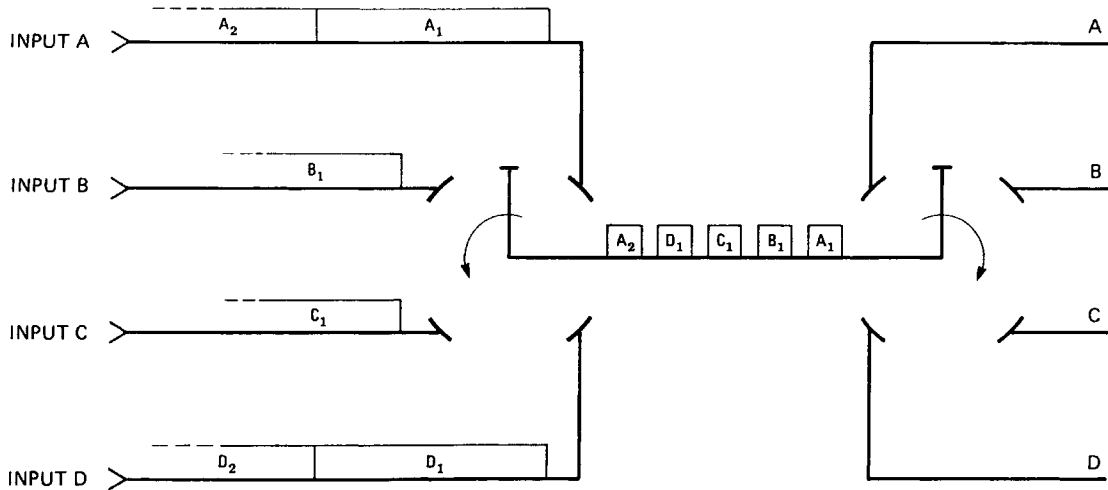
### 3.4.1 Synchronous TDM

Figure 3.19 shows a simple, commutator-based synchronous time division multiplex system. In this example, a commutator allows a fixed number of bits from each input stream onto the path in turn. The timeslots belonging to one input stream constitute a channel. Timeslots are of equal length and are preassigned. If an input does not have data to send, its slots go empty.

In this type of TDM, four rules apply:

1. The bit rates of the input streams must be equal.
2. The multiplexed data bit rate equals the number of tributary inputs times the input bit rate. This relationship implies the existence of a multiplex clock at the transmitter.

## 58 Introduction to Network Technologies and Performance



**Figure 3.19** In this simple commutator analogy of TDM, sequential samples of each of the incoming streams is inserted in the high-speed multiplexed stream.

3. Commutator rotation rate must be exactly the same at the transmitter and receiver. This requirement implies a clock recovery mechanism at the receiver.
4. Commutator phasing must be the same at transmitter and receiver and requires some form of tributary or channel identification within the multiplexed signal.

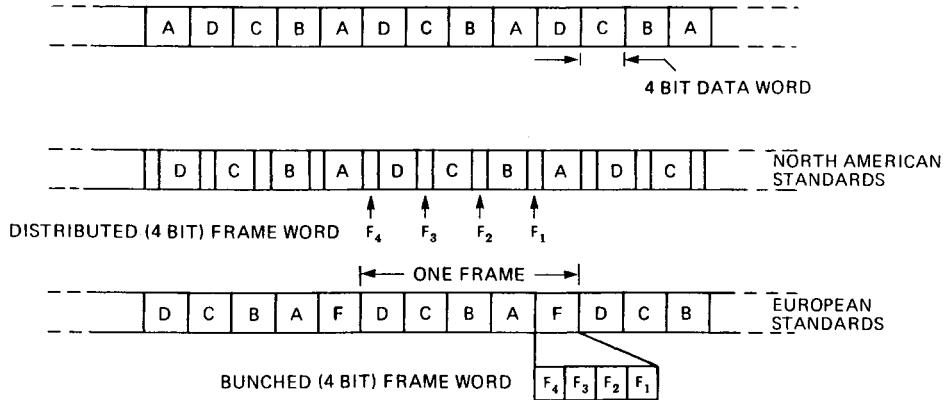
The interleaving of data streams can be done on a bit-by-bit or byte-by-byte basis. (*Byte interleaving* allows parallel processing using lower-speed logic and is finding favor in broadband synchronous systems.)

One rotation of the “commutator” constitutes a frame (Figure 3.20). To help the receiving equipment demultiplex the interleaved bit stream accurately, *framing bits* are added to each frame. These bits identify the starting and ending points of each cycle (i.e., which timeslot belongs to channel 1, which to channel 2, etc.). They are distributed in either of two ways. In one method, the framing bits are inserted between frames, and a complete frame word is built up over a multiframe. This method is the standard for North American systems. In the other method, the framing bits are grouped at the beginning of the frame. This method is called *bunched frame word* and is the standard for European systems.

As mentioned previously, synchronous TDM assumes a constant bit rate for all input streams. But what if our input streams have different rates? In that case, we bring all tributary inputs to an equal, higher, bit rate before multiplexing by a technique called *positive justification* (Figure 3.21). Positive justification means adding redundant (non-data-carrying) bits to the tributary input stream. The rate at which the redundant bits are added depends on the chosen higher bit rate (justified bit rate), and the actual input (tributary) bit rate.


Imagine that the justified bit rate is 6 bps where the input tributary rate is 3 bps (data stream B). This difference means that the nominal justification ratio is 1:2, and that every second bit in the justified output will be a redundant bit.

### CHANNEL IDENTIFICATION METHODS (FRAME ALIGNMENT TECHNIQUES)



**Figure 3.20** Because the individual channels are interleaved in a TDM signal, the beginning of the sequence needs to be marked with a frame word so that the individual channels can be identified and extracted in a particular timeslot or distributed over several frames.

### ASYNCHRONOUS TO SYNCHRONOUS INPUTS (CONVERSION VIA POSITIVE JUSTIFICATION)



**Figure 3.21** In order to multiplex separate digital channels together, they need to be synchronized to the same bit rate. When the incoming streams are asynchronous, they are synchronized to a common higher bit rate by adding additional dummy bits called "stuffing bits." This process is called *positive justification*.

In a real system, the tolerance on tributary clock rate is tight, and far less justification is required than in this simple example. When the bit rates are close but not exactly the same, they are said to be *plesiochronous*. Thus the digital hierarchy based on the “bit stuffing” or justification process is called the *Plesiochronous Digital Hierarchy* (PDH), in contrast to the fully synchronous, byte-interleaved *Synchronous Digital Hierarchy* (SDH). The big drawback with the plesiochronous multiplexing is that to grab a particular tributary or channel in a high-level stream requires the systematic demultiplexing process of dejustification, in which the redundant stuffed bits are identified and removed. Clearly a fully synchronous system avoids this problem.

### 3.4.2 Asynchronous (statistical) TDM


Asynchronous TDM is similar to synchronous TDM except that timeslots, rather than being assigned, are allocated dynamically on a first-come first-served basis. The multiplexer scans the input lines in order, filling timeslots until a frame is filled. If a line does not have data, it is skipped on that round and the slot is filled by bits from the next line that has data to send. If only one input line has data, an entire frame can be filled by its slots.

The advantages of this mechanism are that fewer frames are transmitted with empty slots, so a greater proportion of the link’s capacity is used. In addition, dynamic slot allocation means that asynchronous TDM does not require a 1:1 ratio of timeslots per frame to input lines. In fact, asynchronous TDM assumes that not all of the input lines will have data 100 percent of the time. For this reason, the number of timeslots in each frame is always less than the number of input lines. This difference allows more of the capacity of the channel to be used at any given time and allows a lower-capacity line to support a greater number of inputs. The actual number of timeslots in a frame is based on a statistical analysis of the number of input lines that are likely to be transmitting at a given time, or the aggregate bandwidth required.


The disadvantage of asynchronous TDM is the amount of overhead it requires. In synchronous TDM, channels are indicated by the order of the timeslots in the frame. For example, the first slot of each frame belongs to channel A, the second slot to channel B, etc. Asynchronous TDM frames, however, have no set slot assignments. For the demultiplexer to know which slot belongs to which output line, therefore, the multiplexer must add identifying information to each timeslot. To minimize the impact of this overhead on transmission rate and processing, identifiers are kept small, usually only a small number of bits. In addition, it is possible to append the full address only to the first segment of a transmission, with abbreviated versions to identify subsequent segments.

### 3.4.3 Standard multiplex hierarchies

Figure 3.22 shows the two PDH multiplexing hierarchies (and their associated bit rates and interface codes) specified by the ITU-T (Recommendation G.703). The Bell hierarchy (T lines) is based on the 1.544 Mbps primary rate and is used in North America. The Committee of European PT&Ts (CEPT) hierarchy (E lines) is based on the 2.048 Mbps primary rate, and is used in Europe and most other parts of the world. (In Japan, a third set of standards is employed above the primary level.)



### European Digital Multiplex Hierarchy



**Figure 3.22** The North American or Bell hierarchy and the European or International hierarchy are the two most commonly used systems in the telecommunications network. The North American system is based on the 24-channel primary rate of 1.544 Mbps, while the European system is based on the 2.048 Mbps primary rate. These are the plesiochronous digital hierarchies, which are based on the concept of positive justification for all levels above the primary rate.

Because of the limitations of PDH multiplexing (outlined previously), a new Synchronous Digital Multiplexing (SDH) hierarchy was developed in the late 1980s, initially in the USA as the standard *Synchronous Optical Network*, or *SONET*. The SONET (North American) and SDH (international) standards are very similar, which have simplified the design and manufacture of equipment for the worldwide market.

In SDH/SONET, the hierarchical bit rates are intended to be fully synchronous with byte-interleaved multiplexing, so the bit rates are exact multiples of each other (Figure 3.23). Thus much simpler add/drop multiplexing is possible because there is no bit stuffing. The SDH/SONET standards also provide much more powerful network management capability built into the frame structure than was available with PDH, and the physical parameters of the optical line interface also are specified to allow interconnection of multivendor networks at the lightwave interface. The hierarchical interfaces are specified as Optical Carrier (OC- $n$ ) or Synchronous Transport Signal (STS- $n$ ) in SONET, and as Synchronous Transport Module (STM- $n$ ) in SDH.

**Primary rate switches.** In the Bell system, 24 voice channels are encoded and time-division multiplexed to form a 1.544 Mbps digital signal. This signal is the digital equivalent of two 12-channel groups used in an analog system and is sometimes referred to as a *digroup*. In the CEPT system, 30 voice channels are encoded and time-division multiplexed to form a 2.048 Mbps digital signal. These are the two primary rate multiplex standards.

62 Introduction to Network Technologies and Performance




Figure 3.23 The SDH/SONET multiplex hierarchy is based on synchronous byte interleaving without positive justification. Each hierarchy rate is exactly four times the rate below.

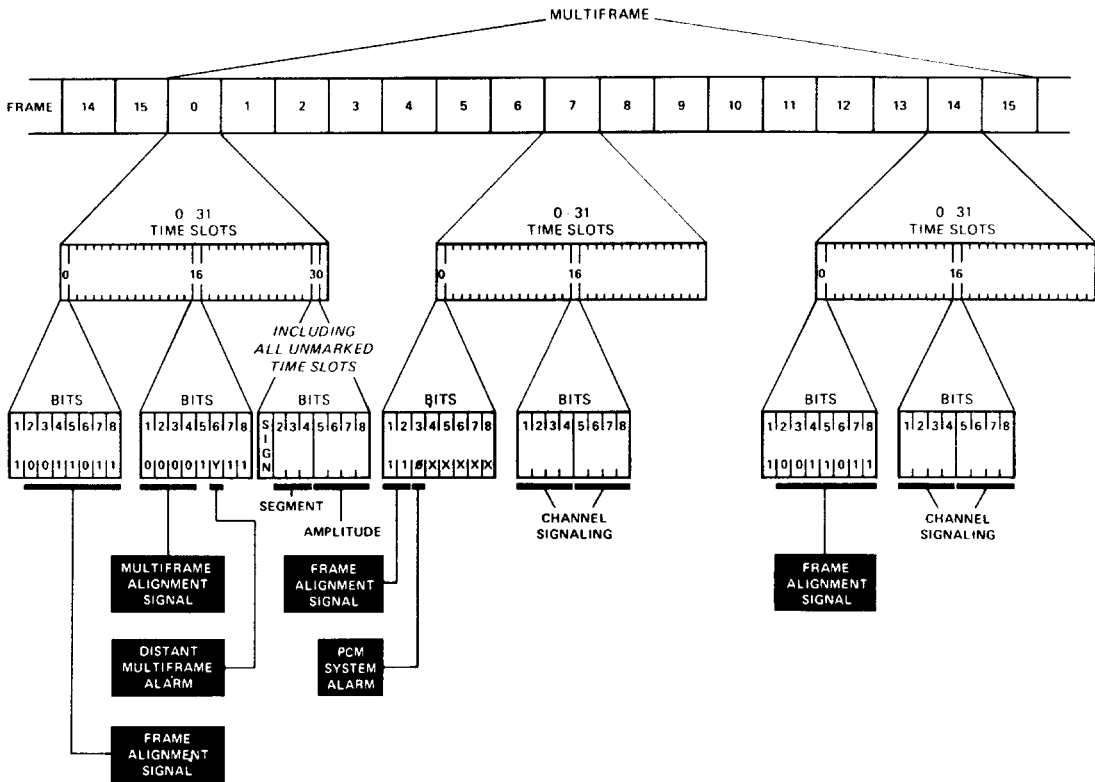


Figure 3.24 The structure of the E1 2.048 Mbps PCM frame and multiframe.

Figure 3.24 shows the structure of an ITU-T (CEPT) standardized frame. As we have seen, each channel is encoded into an 8-bit word at a sampling rate of 8 kHz. Put another way, every 125  $\mu$ s there is an 8-bit word representing a sample of the analog signal in a single channel.

If we combine 30 such channels and allocate each 8-bit word a timeslot within the 125  $\mu$ s frame, we have a time-division multiplexed, word-interleaved signal. The 8-bit words for each of the 30 channels are carried in timeslots 1 through 15 and 17 through 31.

To enable the receiving equipment to identify the start position of each timeslot, an additional 8-bit timeslot (TS0) is allocated as a frame alignment signal. Another 8-bit timeslot (TS16) is added too, to carry signaling information associated with these 30 channels.

Over the years, the earlier 2 Mbps frame format has been enhanced to provide more powerful in-service performance monitoring, and also to report remote-end alarms and system performance. This enhancement is called the *Cyclic Redundancy Checksum* (CRC-4) frame, which is defined in ITU-T G.704. The original 16-frame multiframe is divided into two sub-multiframes. A CRC-4 remainder is calculated on the bits in the sub-multiframe; this is sent in the next sub-multiframe for comparison at the receiving end. Any discrepancy would indicate an error in the preceding block of data (sub-multiframe). This information, along with alarm status, is sent back to the transmitting end via CRC-4 frames in the opposite direction.

The Bell standard, used in North America, specifies a similar process of interleaving 8-bit words (Figure 3.25). In this case, however, 24 channels are multiplexed rather than 30. This mechanism results in a frame of 192 bits, to which is added one framing bit, for a total of 193 bits per frame. Thus several frames together are needed to form a full frame alignment word. This is called a *multiframe*.

In practice, only five out of every six frames use the full 8-bit encoding. In one out of every six frames, only 7-bit coding is used for the speech sample. The remaining (least significant) digit of the channel carries signaling information (routing, on-hook/off-hook, etc.). This technique is known as *bit stealing* and results in a signaling capacity of 1.3 kbps.

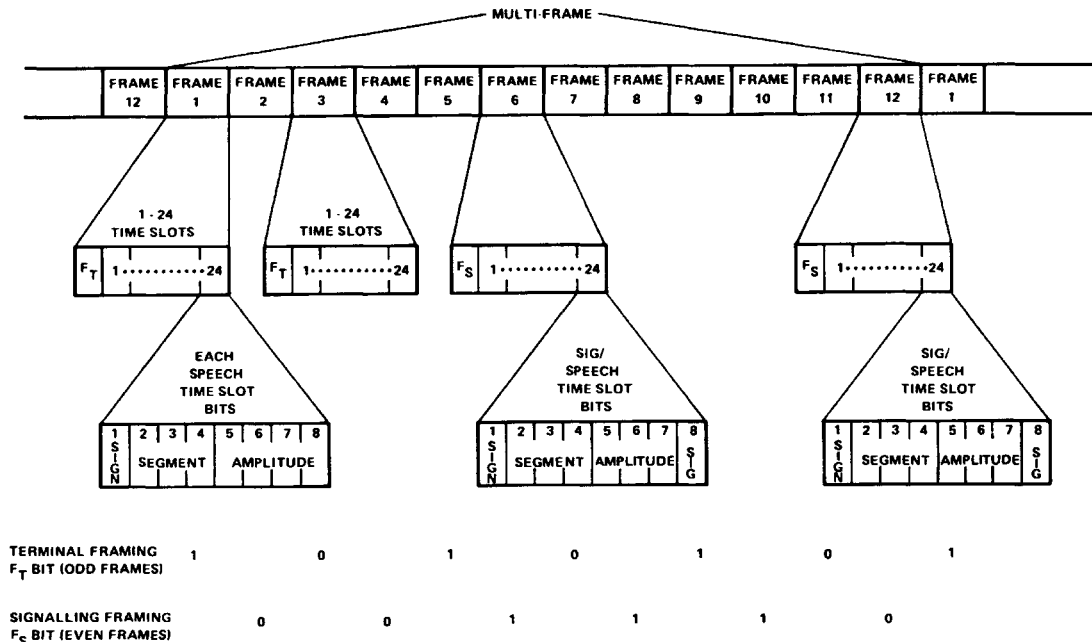


Figure 3.25 The structure of the North American D4 frame and multiframe at 1.544 Mbps.

Bit stealing, coupled with the requirements of AMI encoding (see section 3.5), means that digital data can use only seven bits of each byte of a Bell frame. This restriction results in a data rate of 56 kbps, compared to the 64 kbps available in CEPT systems. (The full 64 kbps service is sometimes called a *clear channel* and is the basis of ISDN, the Integrated Services Digital Network.)

**Note:** The 12-frame arrangement shown here is sometimes referred to as *D-4 framing*. A new standard, called *Extended Superframe Format (ESF)* uses 24 frames in a multiframe. **Bit 1** of the frame is used for the frame alignment word, as well as to carry a 6-bit word for CRC-6 error detection (ITU-T G.704) and as a 4 kbps data channel for maintenance. (For more details see Chapter 7, Section 7.2.2)

### 3.4.4 Digital transmission systems

The four commonly used digital transmission mechanisms are:

- Optical fiber
- Satellite
- Microwave radio
- Coaxial cable


**Optical fiber.** Optical fiber is the most popular high-capacity medium for network operators (PTTs, telcos, and common carriers) where existing routes (or way leaves) exist. The enormous potential bandwidth of optical fiber is gradually being exploited and is responsible for the much lower cost of long-distance and international telecommunications. In the mid 1990s, high-capacity optical fiber systems operate at 2.5 Gbps using the synchronous hierarchy, equivalent to over 30,000 telephone channels per fiber. Some international undersea systems have twice this capacity, operating at 5 Gbps.

The next stage of capacity expansion will exploit *Wavelength Division Multiplexing (WDM)*, wherein several gigabit lightwave carriers are transmitted down a single fiber at slightly different wavelengths, or optical frequencies. The high bandwidth of lightwave systems and the very low attenuation per kilometer of optical fiber (requiring a regenerator or optical amplifier only every 50–100 km), have completely obsoleted the earlier coaxial cable transmission systems that were installed in the late 1970s.

**Satellite.** Satellite systems fall into two broad categories. The first includes large international systems that use Time Division Multiple Access (TDMA), which is digital, and Frequency Division Multiple Access (FDMA), which is analog. The second includes the smaller multichannel systems found in private telecom networks that use either TDMA or Single Channel per Carrier (SCPC), which can be analog or digital.

**Microwave radio.** Microwave radio and satellite systems often are preferred for lower-capacity routes, difficult terrain, and for private and military communications networks where radio's advantages of flexibility, security, and speed of installation are particularly valuable. In the increasingly deregulated telecommunications market, short-range microwave radio provides a convenient way of giving access to customers and bypassing the hard-wired local loop.





**Figure 3.26** Components of digital transmission terminal equipment.

**Media advantages and disadvantages.** In public networks, lightwave transmission accounts for 70–80 percent of circuit capacity, and microwave radio for 20–30 percent, although the ratio depends very much on traffic density and terrain. Some transmission networks are exclusively fiber optic. Satellites are unsurpassed for providing connections to remote or sparsely populated countries.

**Digital transmission process.** Figure 3.26 shows the components that make up digital transmission terminal equipment. The digital traffic comes from the multiplexer and is converted from the standardized interface code (see section 3.4.3) to a media code determined by the system designer for optimum transmission. The suitably encoded data is then fed to the media transmitter for output through the medium.

At the receiver's terminal equipment the process is reversed, but not until the received information has been regenerated in its original, distortion-free format. Regeneration removes the noise and pulse distortion by sampling the incoming signal in a decision circuit and then reconstructing a new digital output signal. To sample the received signal, a stable synchronized clock must be recovered from the incoming signal using a narrow-band clock recovery circuit. An important attribute of the media coding is to ensure reliable clock recovery, independent of the transmitted bit stream sequence.

### 3.4.5 Digital transmission encoding

Transmission of digital information in a digital format requires the translation of the 1s and 0s into a sequence of voltage pulses that can be propagated over a physical medium. This translation is known as *digital encoding*.

The encoding mechanisms in current use were developed to satisfy two basic requirements:

1. Adequate timing content for regenerator operation.
2. Suitable spectrum shaping for media transmission.

The signal must contain enough changes to allow the receiver to differentiate accurately between bits. Binary data can include long strings of 0s or 1s. In an encoding system that uses one voltage level to mean 1 and another to mean 0 (the latter often a zero voltage), a long string of similar bits contains no changes and therefore no timing information that the receiver clock recovery can use for synchronization. The objective is to create a transmission system that is *bit-sequence-independent*.


The frequency spectrum of random binary data (extending to dc) is not ideally suited for simple transmission through any of the common media employed today. For example, with metallic or optical cables, a signal that does not invert equally above and below the 0-V line results in a dc component that degrades the performance of the link. Requirements vary by media. Metallic or optical encoding must result in no dc or low-frequency components. Radio encoding must stay within tight bandwidth restrictions.

Many media codes have been designed to address these requirements. The choice of solution is left to the system designer.


**Metallic media encoding.** The requirement for no dc or low frequencies is satisfied by employing a pseudo-ternary code containing an equal number of positive voltage and negative voltage pulses. These codes are called *pseudo-ternary* because, although they employ three voltage levels (positive, zero, and negative), two of those values are given the same meaning. True ternary codes, such as 4B3T described below, assign a distinct meaning to each voltage level.

The simplest pseudo-ternary codes are the Alternate Mark Inversion (AMI) codes (Figure 3.27). In Return-to-Zero AMI (AMI-RZ), the signal inverts at the beginning of each 1 and returns to zero at the middle of the bit. A 0 is indicated by an unchanging zero voltage. Inverting on each 1 avoids the buildup of a dc factor. Returning to zero at the middle of each bit allows the receiver to synchronize with the bit stream—at least every time the bit is a 1. To avoid losing synchronization during long strings of 0s, AMI-RZ generally calls for scrambling.

An alternative, used in earlier U.S. systems, places an arbitrary limit on the run of 0 bits (e.g., no more than 15); while acceptable for PCM voice traffic, this process affects the clear channel capability required for digital data transmission and ISDN.



**Figure 3.27** Coding for cable transmission and electrical interfaces. The coding rules ensure a minimum number of transitions for all data sequences so that reliable clock recovery is guaranteed.



**Figure 3.28** Examples of media coding for optical line systems.

More convenient are the zero-substitution pseudo-ternary codes, such as *High Density Bipolar 3* (HDB3). In this code, any string of four consecutive 0s is replaced by a predictable pattern of violations containing positive or negative voltage pulses (in other words, timing content). Other examples of zero-substitution codes include the Bipolar N-Zero Substitution series, standardized in the United States: B3ZS, B6ZS, and B8ZS (replacing the earlier AMI implementation.)

In cases where bandwidth restriction is required, a true ternary code such as 4 Binary 3 Ternary (4B3T) may be employed. In this code, 4-bit binary segments are substituted with 3-symbol ternary words. The 4B3T code uses three meaningful voltage levels: zero, positive, and negative. Each level represents a pattern of bits instead of an individual bit. These ternary words (patterns) are chosen to guarantee that the spectral shaping and timing content requirements are met within a limited bandwidth.

**Optical media encoding.** Optical fiber transmission systems have media encoding requirements similar to those of metallic cable systems. At present, however, optical systems are restricted to transmitting 2-level binary symbols (+ and -), as shown in Figure 3.28. For this reason, a zero-disparity binary coding scheme such as Coded Mark Inversion (CMI) or 5 Binary 6 Binary (5B6B) is normally employed. *Zero-disparity* means that an equal number of 1s and 0s appear in the encoded signal.

CMI encoding effectively doubles the bit rate of the signal by replacing each uncoded binary digit with a pair of pulses. Binary 1s are replaced alternately by 11 and 00. Binary 0s are replaced by 01. CMI also is used in Europe as the interequipment interface code for 139 Mbps equipment.

In 5B6B encoding, 5-bit binary segments are replaced by 6-bit binary words. In other words, a distinct 6-bit sequence is used to represent a distinct 5-bit sequence.


The alphabet of 6-bit words is chosen to guarantee that adequate timing content and zero-disparity are achieved within a signal.

Binary coding schemes such as 5B6B and 7B8B increase the bit rate required by the transmission system. The latest designs for very high-speed lightwave systems, based on the SONET/SDH standard, use a scrambler (rather than line-encoding schemes) with a higher-Q clock recovery circuit. This reduces the amount of extra bandwidth required for reliable transmission.

**Radio media encoding.** With digital microwave radio systems, the overriding requirement is to restrict the transmitted signal to within an allocated frequency bandwidth (Figure 3.29). To fit within this constraint and still meet the demand for increased data rate throughput, digital microwave radio uses *multilevel symbol encoding*. In this type of media code,  $n$  input data bits are converted to one symbol (which may take  $2^n$  levels), thereby giving a direct reduction in the bandwidth required for transmission. Each signal change represents multiple bits. Fewer changes therefore are required to represent the same number of bits. For example, three (2-level binary) bits can be encoded as one (8-level) symbol, requiring only one-third of the bandwidth for transmission. Adequate timing content is achieved by scrambling the uncoded binary data.

Note also that in Quadrature Amplitude Modulation (QAM) systems, two multilevel signals are transmitted simultaneously by modulating the multilevel signals on carriers 90 degrees out of phase with one another.

### MEDIA CODING EXAMPLES



**Figure 3.29** Multilevel coding for digital radio systems is designed to minimize the occupied bandwidth for a given data rate.

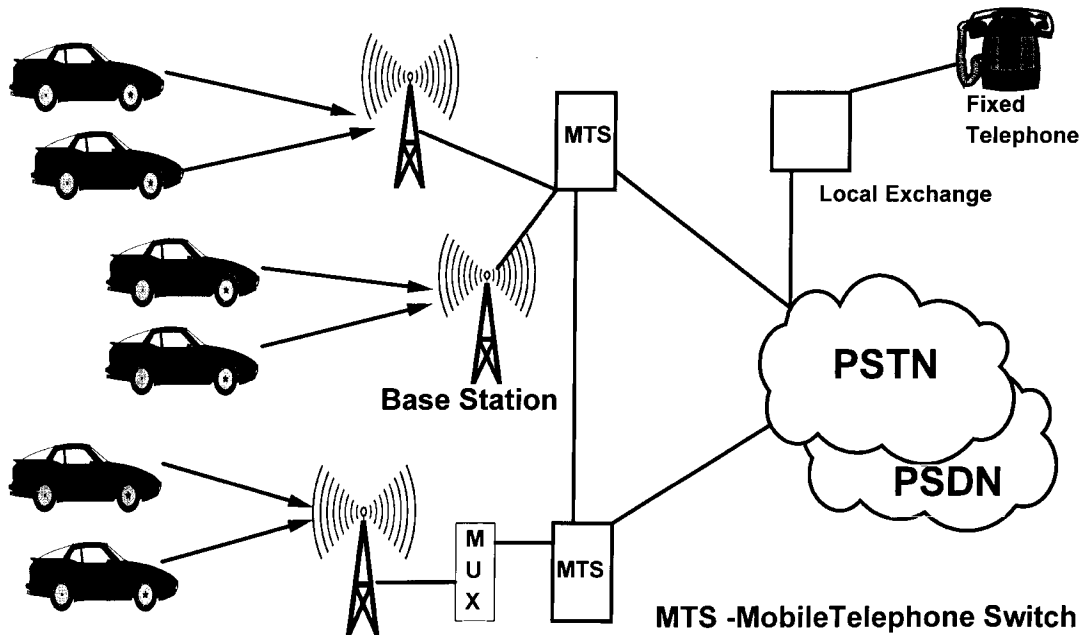


Figure 3.30 The cellular radio network.

High-capacity line-of-sight microwave radio systems generally use high-density modulation schemes such as 64-QAM and 256-QAM; short-range, low-capacity systems operating at high microwave frequencies typically use 4-PSK (4-level phase shift keying, also referred to as QPSK).

### 3.5 Cellular Radio

Cellular radio can be thought of as a complex local loop access mechanism that replaces the usual 2-wire telephone line with the equivalent of a 4-wire connection through UHF (ultrahigh-frequency) radio. The complexity arises from the RF communications hardware required (called the “air interface”), and from the need to keep track of the mobile telephones as they move around a geographical area. Figure 3.30 shows the basic elements of a cellular radio system.

Cellular radio differs from ordinary mobile radio in that it is a true extension of the public telephone network and provides full-duplex communication. (Mobile radio is half-duplex, based on press-to-talk.) Cellular radio systems operate in frequency bands around 450 and 900 MHz. For example, two 25-MHz bands have been allocated at 900 MHz (890–915 MHz and 935–960 MHz) to provide two directions of transmission. The new generation of PDS and DCS systems use 1800- and 1900-MHz bands. Current analog cellular systems typically use a channel bandwidth of 20–30 kHz with frequency modulation (FM), though new systems reduce this to 12.5 kHz.

**Cellular radio standards.** A number of different cellular telephone standards are in use. In the analog era, the Nordic Mobile Telephone (NMT) standard was the first to be introduced, in 1979. This was followed by AMPS in North America and TACS in the United Kingdom. These are all FM FDMA (Frequency Division Multiple Access) systems.

New digital mobile systems entered the market in the 1990s, notably the GSM *Time Division Multiple Access* (TDMA) system, which has become the standard in Europe and the Asia-Pacific region. In North America, the digital standards use TDMA and *Code Division Multiple Access* (CDMA), which allows multiple users to occupy the same spread-spectrum bandwidth, but identifies each with a specific coding incorporated into the modulated RF signal. Digital standards achieve a higher number of subscribers per megahertz of allocated bandwidth, provide better service quality, and offer improved security compared to earlier analog systems because of the coding and frequency hopping.

**Cells.** The term *cellular radio* derives from the system's frequency plan, which divides a geographical area into a honeycomb of cells. The size of a cell depends on traffic density. Each cell in a city might be 3.3 km in diameter; in a downtown business area, *microcells* as small as 500 m<sup>2</sup> might be used. Out in the sparsely populated countryside, cells could be 10–30 km in diameter. Each cell is served by a *base station* that operates with a small portion of the allocated frequency band, and transmitter power is set according to the size of the cell to be covered. Each cell's portion of the frequency band differs from those of its neighbors to avoid interference. By setting the transmitter power at only the level required to cover a particular cell, the operating frequency can be reused again for another cell some distance away. This is called the *frequency plan*.

As a mobile telephone roams from one cell to another, it retunes to new frequencies. At the same time, the cellular radio network updates its database (the Vehicle Location Register) to show the changing location of the particular mobile unit. The database enables a subscriber in the public switched telephone network (PSTN) to call into a mobile unit. These connections are the function of the mobile telephone switch (MTS) or mobile telephone exchange (MTX), which uses *Signaling System 7* (SS7) for database messages and controlling the call. The database transactions not only keep track of the mobile unit's movements, they also check that legitimate calls are being made into the network by comparing details of the mobile's registration with information in the database called the *Home Location Register* (HLR).

The mobile telephone segment has been one of the fastest growing parts of the telecommunications network. Since the late 1980s, the average annual growth in the number of mobile subscribers has been around 50 percent per year, reaching a total of over 120 million in 1996, of which around 30 percent were digital. This represents a significant fraction of the total network traffic.

### 3.6 Bibliography

- Mazda, Fraidoon (Ed.). *Telecommunications Engineers Reference Book*. (Newton, Mass.: Butterworth-Heinemann, 1993.)
- Siemens A.G. *International Telecom Statistics*. (Munich: Siemens A.G., 1996.)

---

Part

**2**

# Network Test and Measurement





# Testing in the Life Cycle of a Network

Robert L. Allen

Michael J. Cunningham

*Hewlett-Packard Communications Test Solutions Group*

## 4.1 Introduction

Communications network test and measurement are activities undertaken to characterize the operation of networks and network elements. *Measurement* determines performance parameters, either on an as-needed basis or continuously via dedicated monitoring equipment. *Test* adds the comparison of measured parameters to accept/reject thresholds, or the application of controlled stimuli. Testing and measurement contribute to advancements in communications networks by providing the quantitative indications of whether network elements and services are performing as expected during the various phases in a network's life.

### 4.1.1 Network life cycle

The life of the network consists of four major phases:

1. A *development phase*, wherein new network equipment and services are designed and debugged.
2. A *production phase*, wherein the design is replicated in substantial volume.
3. An *installation and commissioning activity*, wherein new equipment is put into operation, usually expanding an existing network.
4. An *operational phase*, wherein the network equipment or service is in the field and needs to be kept in good working order. This phase has two test activities, one to monitor network health and one to repair problems.

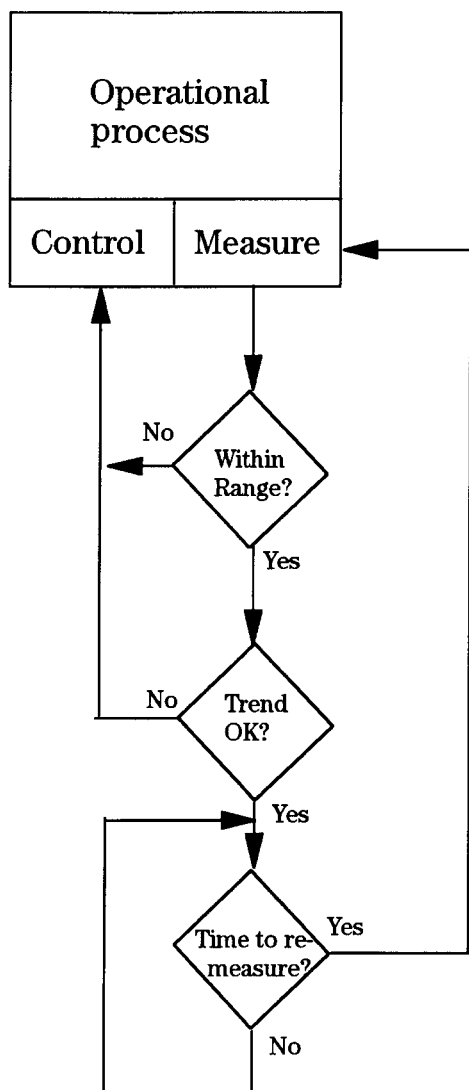
Test and measurement instruments and systems provide important information needed to manage processes in all of these phases.

## 74 Network Test and Measurement

## 4.1.2 Process control

Generally speaking, each phase consists of complex processes that need to be controlled so their outputs meet expectations. Test and measurement determine the parameters that indicate the health of these processes.

Sometimes the parameters will indicate that the process is operating as expected and no action is needed. At other times, it will not be operating suitably and intervention is needed to set it right. Or perhaps the process is still within tolerance, but the trend on a parameter might indicate that it eventually will cause trouble and some preemptive action is wise. Finally, process measures can detect when basic capacity may be less than what is needed to accommodate the predicted demand for the process. Test and measurement generally are the means to get precise data about the process to trigger action and verify that the action has been effective (Figure 4.1).



**Figure 4.1** This flowchart depicts what is done with test and measurement information to decide whether a process requires intervention.

**TABLE 4.1 Phases/Activities of Communications Network Test and Measurement.**

Phase	Development	Production	Installation	Operational	
Activities				Monitor	Repair
	—Simulation	—Board Fabrication	—Performance	—Data	
	—Prototyping*	—Component Fab	Verification*	collection*	—Fault Location*
	—Performance	—Assembly	—Installation	—Threshold	—Module
	Verification*	—Adjustment*	—System	Comparison	Replacement
	—Environmental	—Performance	Verification*	—Root Cause	—Performance
	Characterization*	Verification*		Determination	Verification*
	—Reliability				—Module Repair*
	Demonstration*				

\*Denotes activity with significant test and measurement content

The overall objective of any process is to satisfy the needs and expectations of the customers who purchase the output of the process. This generally requires that the quality and cost of each process output are as good or better than alternative sources, making quality and cost the goal of process control—and hence the reason for the underlying metrics provided by test and measurement.

#### 4.1.3 Test objectives at life cycle phases

Test and measurement provide the information needed to control the processes during the network life cycle. During the development phase, it assures that the design of the network element or service has been accomplished correctly. During the production phase, it verifies that the product was assembled correctly using good parts. During the installation and commissioning phase, it assures that nothing was damaged in transit and that it has been connected to the network correctly. During the operational phase, it verifies that the network is providing adequate performance, isolating what elements are responsible for any problems, and verifying that any repairs are completed correctly. See Table 4.1 for a summary of the activities during each phase and the ones that depend heavily on test and measurement.

These activities answer several important questions at each life cycle phase:

- Development Phase
  - Does the design meet performance goals?
  - Does the design conform to industry standards?
  - Has reliability testing demonstrated an acceptable mean-time-to-failure?
  - Has overall performance been verified over the intended environmental range?
  - Does the product meet applicable requirements for electromagnetic compatibility?
  - Does the product conform to required product safety standards?
- Production Phase
  - Have internal adjustments been set to design values?
  - Have sufficient functions been verified to assure complete operation?
  - Has performance been verified at demanding environmental extremes?
  - Are all guaranteed performance parameters within test limits?
  - Is performance stable?

**76 Network Test and Measurement**

- Installation and Commissioning Phase
  - Do network elements function completely?
  - Have all network connections been made properly?
  - Is network performance per applicable specifications?
  - Does performance under load meet requirements?
  - Have sufficient benchmarks been taken for future reference?
- Operational Phase
  - Is the network performing at level that meets customer expectations?
  - What is the root cause of situations that are below acceptable performance?
  - Where is the problem located?
  - Have exchange modules repaired the problem?
  - Is network performance now per applicable specifications?
  - Has the faulty module been repaired properly?

**4.1.4 Special-purpose communications test equipment**

Special-purpose test equipment for communications networks combines specific functions and data reduction /display capability tailored to the items under test. This lets the measurements be made more quickly and gets the results in a more usable form. Several factors relating to the network also contribute to the need for special-purpose test equipment.

1. The need for interoperability between network elements, making conformance to industry standards very important and giving rise to test equipment that is able to verify this conformance.
2. The network's standardized way of combining the traffic of several users on a single path, requiring that test equipment interface at the various levels of the multiplex hierarchy, stimulating and measuring performance at all levels.
3. The network's geographic dispersion, requiring most test equipment to work with information that can be determined at a single location.
4. The network's use to carry computer traffic, creating a class of instrument to analyze data communications packets traveling on a shared medium.
5. The need to determine root causes of operational problems by comparing the information gained at several points in the network.

These requirements result in test equipment functions that are specialized to the communications network. One is generating, modulating, and demodulating carrier signals that match industry-standard interfaces, which functions are provided by integrated test sets at UHF, microwave, or lightwave frequencies. Another is generating formatted serial bit streams at the clock rates corresponding to the standardized digital hierarchy, complemented by error detection based on recovery of the same bit stream after it has been looped back at the appropriate point in the network, which are provided by the *bit error ratio test set* (or BERTS). A third is generating, capturing, and decoding data communications packets that conform to industry-standard interfaces and message formats, which are as provided by the protocol analyzer.

All the life cycle phases also require general-purpose test equipment such as oscilloscopes, voltmeters, spectrum analyzers, logic analyzers, etc. General-purpose test equipment applies to several types of electronic equipment, including communications network equipment. This handbook covers special-purpose test equipment only; for similar treatment on general-purpose test equipment, see *Electronic Instrument Handbook, Second Edition*, by Clyde F. Coombs, Jr. (McGraw-Hill, 1995).

## 4.2 Development Phase Test and Measurement

Development of communications products is complicated by three factors:

1. The fast pace of technology: *Design to perform*.
2. The demands of national and international standards: *Design to conform*.
3. The organic nature of communications networks: *Design to interoperate*.

The fast pace of technology is a double-edged sword. While new and better technologies enable higher performance in new products, they also render last year's products obsolete. The communications product designer must stay abreast of new technologies and understand how these can be used to enhance product performance for competitive advantage. The first imperative for a communications equipment designer is to design to perform—perform better than the competitive product, at a lower price.

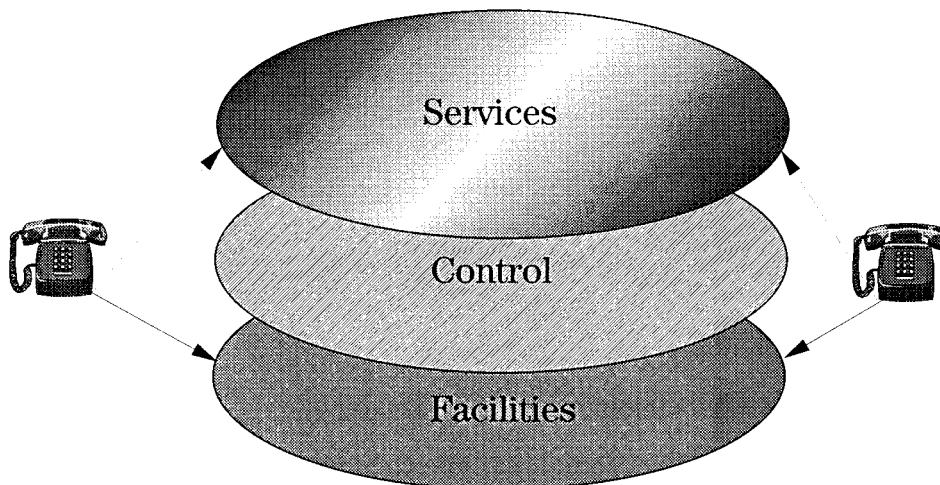
While standards continue to be a long-term stabilizing factor in communications, the short-term situation for new technologies can be rather competitive as different factions promote their ideas in standards bodies, or others attempt an end run, hoping to establish a de facto standard. Often the designer is working ahead of the final version of the standard and must decide which of perhaps several competing approaches will win the day. In the end, the design must conform to whatever standard wins in the competitive marketplace. Design to conform, yes, but conform to which standard?

“The organic nature of communications networks” means that they are constantly evolving. The new product must work in the old network while driving its evolution to become the more capable network of tomorrow. Designing to interoperate demands interoperation with all the old network equipment, as well as the new generation from other developers designing to the same standard.

This section reviews how test and measurement equipment enables the three imperatives: design to perform, design to conform, and design to interoperate.

### 4.2.1 The organic nature of communications networks

Clearly there was a time when communications networks did not exist. Now that we have them all around us, however, it seems they have always been here. Furthermore, communications networks seem never to die, nor for that matter are we particularly aware of their birth. New services offered via ever-growing, ever-changing communications networks are the most tangible evidence most users have of the installation of some new network capability, but the physical network is so ubiquitous that we seldom notice the physical changes.



**Figure 4.2** The physical network that can be seen and touched is the facilities network. This is controlled by an overlay packet-switching network. On this combination of facilities and control, various services are implemented. In voice telephone service, the telephone is physically connected to the facilities network, but logically to the service. The control network sets up and tears down the call and keeps track of usage for billing.

A good example of this is the public telecommunications network, which is actually a network of networks on which a variety of services are deployed (Figure 4.2). Services are implemented on the physical network, and it is with these services that we interact. The most common service is POTS, or Plain Old Telephone Service—people talking to people, or people accessing the Internet via a modem and a dial-up line. Data services are implemented on the same network facilities, as well as on private-line services or private switched virtual network services.

The underlying networks are called *facilities networks*. These consist of several types of network elements: access, switching, and transport facilities comprising new and old network technologies, some over half a century old. As various services are implemented on the facilities networks, the combination of facilities and services is controlled by yet another network, the *signaling network*, overlaid on (and sometimes partly utilizing) the facilities network (see Figure 4.3).

This composite of networks, control, and services is probably the most immense and complex technological development of humankind. In some respects it is a worldwide, distributed, real-time computing system. In other respects it is almost organic. It is like a rain forest with an underlying structure of trees and plants (facilities networks) supporting a variety of populations of living creatures (services) in its canopy. Life goes on in the canopy as trees and plants sprout, grow up, and die below. If all of a certain species of plant life dies out, one or more populations might die as well. Introduction of new types of plants may support entirely new populations.

So it is with the public telecommunications network. As old facilities are taken out of service, and this may take decades to accomplish, some services are terminated,

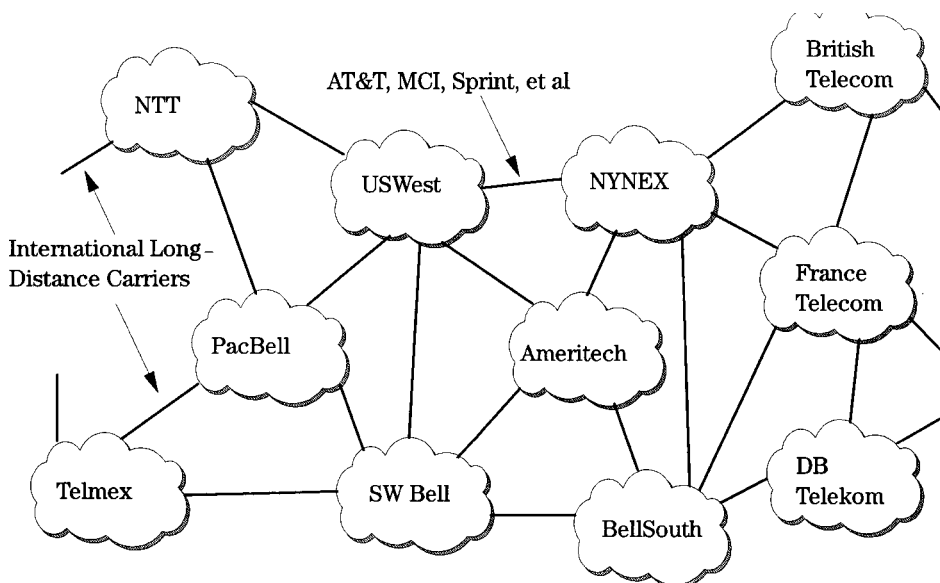
and some live on, surviving on new facilities. As new facilities are added, new services spring up that can only exist on the new facility.

It is into this ever-changing milieu that a network equipment manufacturer must introduce new network elements. The new product must conform to the standards set for the particular type of facility it is part of, and must support the new services envisioned by the network service provider. At the same time, it might have to support some variety of existing services, and it must interoperate with the older network equipment making up the facilities networks at the moment. Thus it is a game of fitting in—and, at the same time, standing out sufficiently to merit installation in support of new services and additional revenue for the service provider.

#### 4.2.2 The nature of communications network new product development

A generic new product development cycle might consist of the following phases:

- *Defining*: What to build? What functions? What performance level?
- *Designing and prototyping*: Prove feasibility. Does the design concept work?
- *Performance testing*: Prove required performance for each function.
- *Reworking as necessary*: Correct discrepancies between design and practice.
- *Transferring to manufacture*: Quantity production.



**Figure 4.3** The public switched telecommunications network is a worldwide network of networks. A variety of services may be implemented on these networks, some local to a particular network; others, like telephone service, are ubiquitous wherever the network reaches.

**80 Network Test and Measurement**

The new product development cycle for communications equipment must include a phase for trying the product in the actual field environment, as well as more extensive testing of conformance to standards and interoperability with existing equipment and systems:

- *Defining*: What to build? What functions? What standard? What performance level?
- *Designing and prototyping*: Prove feasibility. Does it work?
- *Performance testing*: Prove required performance for each function.
- *Conformance testing*: Test against functional constraints of the applicable standards.
- *Interoperability testing*: Combine performance and conformance testing to assure interoperability in real networks.
- *Field testing*: Prove performance, conformance, and interoperability in a real network setting.
- *Reworking as necessary*: Correct discrepancies between design and practice.
- *More performance, conformance, and interoperability testing*.
- *Transferring to manufacture*: Quantity production.

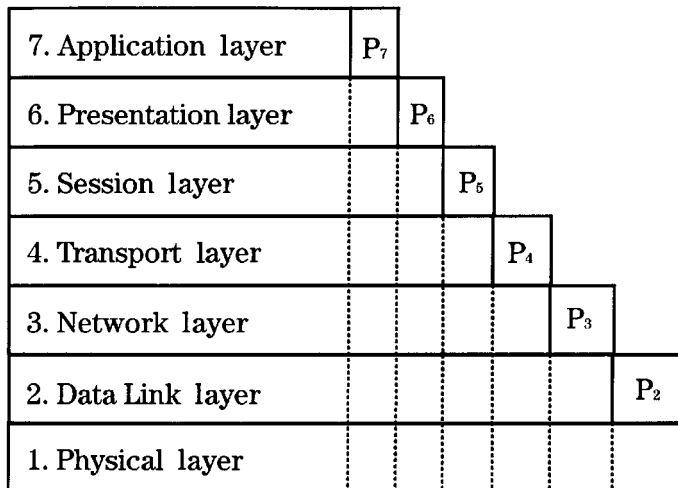
**4.2.3 Design to perform**

*Design to perform* means first of all to perform the function as called for in the definition. Second, and most important for the success of the product, it must perform better at a lower cost than competing products. As the prototype is built, the functionality and performance are continually tested using test and measurement instruments specifically designed for communications equipment development. The prototype is first built with sufficient functionality to support the Physical layer (layer 1) of the protocol stack (see Figure 4.4). A protocol tester is used to generate test patterns to send to the prototype equipment and to analyze the resulting output.

As layer 1 is completed, attention turns to the Data Link layer (layer 2). At this time, the same protocol tester should be capable of setting up layer 1 easily, while allowing comprehensive testing of layer 2. This pattern repeats as the design moves up the protocol stack. With minimal effort from the operator, the protocol tester sets up the lower layers where testing is complete, yet allows comprehensive testing of the layer of interest. At any time, of course, the operator might want to retest a lower layer, and that layer must be available quickly for full testing, even though it had been part of a shortcut set up initially.

Once basic functionality is proven, *stress testing* begins. This fully loads the network element with traffic generated by the protocol tester while watching for traffic dropouts. Errors are inserted to test error recovery functions. Timing jitter is added to input signals and the output is monitored for errors. In developing a new router, the key performance parameter might be the number of packets routed per second. In this case, the test set must be able to send packets at a rate high enough to stress the





**Figure 4.4** A network element operating in an ISO seven-layer protocol stack is designed and tested layer by layer to ensure the signal is properly processed at each layer, in accordance with the specifications for the protocol standard. Notice that each layer adds protocol bits to control the protocol at that layer.

router. A switch designer, on the other hand, might be concerned with loading every port with a constant bit rate, so the test set needs to have many parallel outputs.

#### 4.2.4 Design to conform

As functionality and performance are proven, test scripts are run on the protocol tester to try all the details of the relevant standard for the product. These scripts are procured from protocol test centers, which specialize in providing comprehensive test suites for the various telecommunications standards. Conformance testing methodology itself is defined in an international standard, ISO 9646. There are four basic steps:

1. From the specification for the standard, an Abstract Test Suite (ATS) is defined.
2. The protocol test center writes an Executable Test Suite (ETS) from the ATS.
3. The developer, or protocol test center in the case of certification, runs the ETS on a protocol tester connected to the new product or its prototype.
4. A test report is generated to guide further development or justify the certificate of conformance.

Relevant portions of the test suites can be run at each stage of development to demonstrate the conformance to the standard as the design progresses. This discipline will assure that anomalies observed at one protocol layer are not caused by failure to conform at a lower layer.

After the new network element has passed all the conformance tests in the developer's lab, the product is submitted to the test center for official certification.

## 82 Network Test and Measurement

By designing to conform and running conformance tests in the developer's lab, the certificate usually can be obtained with a single visit to the protocol test center.

### 4.2.5 Design to interoperate

Interoperability testing answers the question: Will this new product or network element work with another particular product, or with a class of products built to the same standard? Interoperability testing aims to give confidence that the new product will work in a real network.

It is a matter of strategically combining conformance testing with stress testing to test the full range of parameter variations against the limits of the standard under realistic traffic loading. The conformance testing assures the product meets all the timing and signal interactions demanded of the standard, and theoretically met by equipment already operating in the network. Stress testing assures it can do this when fully loaded with network traffic. As in designing to conform, designing to interoperate is most effective when testing is done as the design progresses.

## 4.3 Production Phase Test and Measurement

Most communications equipment is a complex assembly of sophisticated electronic components. A PCM multiplexer, for example, contains many printed circuit boards, each holding a large number of integrated circuits and other components. Manufacturing such a piece of equipment is not easy, but it is made somewhat easier if problems are discovered early in the build cycle.

This helps pinpoint the process that needs optimization, and it minimizes the cost of recovering from the problem (i.e., the cost of rework). Hence ICs are tested, unloaded boards are tested, loaded boards are tested, subassemblies are tested, and the final item is tested. Early tests are usually basic and are done to verify that components are free of manufacturing flaws. Later ones are more advanced and attempt to determine functional integrity and overall product quality. Tests on the final item demonstrate compliance with published specifications.

### 4.3.1 Production test

Testing during the production phase can be categorized by the nature of the item under test.

- Component test
  - PC boards (unloaded), to determine that the board is free of undesired connections or opens between traces or layers.
  - ICs, to determine that all functions perform adequately.
  - Others (such as expensive or key components that will mount on boards), to test general functionality.
- Subassembly test
  - PC boards (loaded), statically tested to verify proper loading, and dynamically tested for functional aspects.
  - Modules (such as UHF filter, fiber optic transceiver, SHF mixer), functionally tested to verify critical performance specifications.

- End product test
  - Parametric test, to set adjustable internal parameters at design levels.
  - Functional test, to determine that guaranteed performance is provided.
- Module troubleshooting and repair
  - Offline tests, to repair or scrap problem assemblies.

At the component and subassembly level, most of the test equipment needed is not specially designed for communications requirements. End product testing is where the special-purpose test equipment is required. It provides stimuli as expected in the network and analyzes the unit's response for its compliance with specifications. Troubleshooting and repair also involve some special-purpose models, particularly when the troubleshooting begins with a substantial portion of the end product.

#### 4.3.2 Specification budget

End product accept/reject criteria have to allow for effects of environment (especially ambient temperature) on the performance of the item, as well as aging and possible test error. For example, consider verifying that a cellular radio base station transmitter produces sufficient UHF output power. If the guaranteed output is 1 W or greater, the production accept/reject criterion may be 1.3 W. This difference is determined from output power variations with operating environment (characterized during development phase), expected aging effects (characterized during development phase), and the uncertainty of the production test (calculated from instrument specifications, connection effects, and calibration method). These variables are generally combined linearly in a specification budget that establishes a test limit that is consistent with the guaranteed performance. Figure 4.5 shows this, including a probable distribution of test results for a large number of units.

#### 4.3.5 Automatic test

Test automation is often used to reduce overall test cost and to get comprehensive data to uncover subtle process problems. Although initially more costly because of the more sophisticated equipment involved (computers and programmable instruments), and the development work to create the application software that runs the tests, the reduced test time per item often will recover this investment over the expected volume of items to be produced. The more complete capture of data allows correlation of poor yields with underlying process problems, often in processes that feed this one. In the most automated situations, information is sent from a variety of test systems over a data communications network to a single computer system, where correlation between earlier test results and subsequent process problems can be sought.

#### 4.3.6 Design for test

A significant consideration in the design of sophisticated electronic products is how they will be tested both in production and during field repairs. A part of the answer is access to test points. A loaded printed circuit assembly, for example, might need important nodes brought to edge connectors so a test system can access them without




Figure 4.5 Specification budgets establish test limits that are consistent with guaranteed product performance.

probing or special fixtures. Often a test interface connector is included on the printed circuit assembly to provide complete access. In general, design goals should include some consideration of the production and repair processes and their need for efficient access to important nodes.

#### 4.4 Installation/Commissioning Test and Measurement

Installation includes everything from pulling cable to setting up equipment as small as a modem or as large as a Central Office switch. Installation may be confined to a single building or manhole, or it may stretch across the nation or halfway around the world. It is almost always part of network growth, expansion, or upgrade. Brand-new standalone network installations are in the minority.

In every case, the new equipment must be tested to be certain it meets specifications and will not bring the rest of the network down when it's turned on for service. Tests guide the installer in the step-by-step process of installation, and satisfy the supplier that the product has been delivered and connected according to the customer's requirements. Commissioning tests are more formal, and typically apply to major installations. These tests, usually called for in the contract, assure the buyer that the network or network element performs as specified.

#### 4.4.1 Installation testing

Installing new network elements, whether a multimillion-dollar switch or several miles of fiber optic cable, involves a process, and that process adds value to the network element. Just shipping the equipment to the installation site adds value (cost), value that is lost if the equipment must be shipped back to the factory for repair, or if a specialist must be sent to the site. Just burying a cable adds value, value that is lost if it must be dug up for repair or replacement. Said another way, each step that adds value also increases the cost of finding and repairing a fault. This cost can increase dramatically as the installation progresses.

How much testing ought to accompany installation? Consider two examples on either side of the range of possibilities. Twisted-pair copper cables have been manufactured and installed for many decades. The industry has millions of miles of experience with it. It is relatively inexpensive, so that spare capacity is always installed to allow for growth and any irrecoverable faults in individual pairs. Pretesting copper cables in the field probably won't pay for itself. On the other hand, a communications satellite is very expensive and, once in orbit, repairs are nearly impossible. Prelaunch testing of satellites is probably the most comprehensive of any communications test activity.

Thus the type and degree of installation testing depends on:

- The complexity of the technology being installed.
- The maturity of the technology.
- The cost of post-installation fault location and repair.

#### 4.4.2 Installer's test and measurement tools

Installation test sets differ according to the complexity and maturity of the network technology being installed. For more mature technologies, test sets are more likely to be multifunction, covering more than one network technology. For newer or more complex technologies, the test sets are more likely to be close derivatives of those used in the development of the technology itself. In either case, the test capability must be enough to give confidence that the installation has been done correctly with properly functioning components, and that the installed facilities will support the service intended. The following features will speed the process and reduce training required for operators:

- Go/no-go indications of parameters tested
- Preprogrammed or "canned" test suites
- Multifunction testing in one box, with a user interface common to all functions

#### 4.4.3 Commissioning test and measurement sets

Because commissioning tests are more formal and often are specified in the installation contract, capability to provide a record of the test results usually is required. This might be in the form of a hardcopy record or, more effectively, stored in memory

from which a formal report can be derived. These results, if made available to the network operator, can provide important benchmarks for future performance monitoring and troubleshooting. Because it is sometimes not economically practical to use separate test sets for commissioning tests, installation test sets for more mature network technologies are often provided with commissioning test features as well. As with installation test tools for newer network technologies, commissioning test sets for these may be derivatives of test sets used in the development phase.

#### 4.5 Operational Phase Test and Measurement

A network manager is constantly balancing operational cost versus network performance, all the while dealing with the inevitable problems that bring parts of the network down from time to time. Benchmarking normal performance parameters and measuring ongoing network performance and traffic load are key to maintaining the balance, and to successful troubleshooting. A strategy for effective network management should include the following:

- A carefully planned and regularly performed benchmarking program
- Proactive efforts to discover negative performance trends before service is affected
- Means to restore service to mission-critical applications immediately.
- Rapid fault isolation and repair tools and procedures

##### 4.5.1 Benchmarking

A good benchmarking program consists of measuring top-level performance parameters, and the parameters that drive the top-level parameters, at key points throughout the network on a regular basis. Such benchmarks, when stored and compared with current data, reveal performance and traffic load trends, as well as provide important clues for troubleshooting. Benchmarking can be done manually or semiautomatically for a small network, but is far more effective if automated. On a large network, automation is essential.

##### 4.5.2 Proactive performance management

A well-managed program of benchmarking the network, then monitoring performance vis-a-vis the benchmarks, allows the network manager to track performance trends and take action before performance degrades enough to seriously affect users. Such a program is called *proactive performance management*, and depends on nonintrusive testing and automatic monitoring of performance trends.

**Nonintrusive testing.** After network elements are put into an operating network, test and measurement activities are of necessity nonintrusive. The objective is to monitor key signal characteristics or message contents without interfering with the normal operation of the network. In other words, the network must continue to carry live traffic while the monitoring process provides information on how well the traffic is being handled. This generally is done by looking at signal elements that are independent of the information that is being carried.

In the case of a digital multiplex hierarchy, for example, the framing bits of the data stream provide a pattern that is supposed to conform to certain rules. If these bits are noncompliant with those rules, it indicates performance problems. Generally it is possible to watch the bit error ratio of these framing bits and to base service decisions on thresholds that are more sensitive than those that would cause a drop in perceived quality by a user.

**Automatic monitoring of performance trends.** Online monitoring instruments can be connected to small computers (at the same site or at a remote site via a data communications link) to facilitate the analysis of the gathered data. This provides for logging, comparison to benchmark data, and statistical summaries; it also automates triggering alarms when a monitored parameter exceeds its threshold.

In most cases, the instrument and computer can be viewed as a monitoring system, especially when the computer is controlling and gathering data from several instruments that may be deployed at different spots in the network. To reduce the cost associated with such systems, the measurement equipment generally will not have the full control and display functions necessary to operate as a standalone. These faceless instruments can be thought of as probes that act as the data acquisition front end of the monitoring system.

A key contribution of the computer in an automatic monitoring system is transforming the measured data into a more useful form. Without this transformation, the network operator is overwhelmed with data, making it difficult to see where the most important problems are and what actions to take. In highly developed systems, for example, data is taken from hundreds of probes, and graphical summaries in bar charts or radar diagrams indicate whether any threshold has been exceeded. Clicking on the bar or axis explodes the diagram into the underlying data (Figure 4.6).

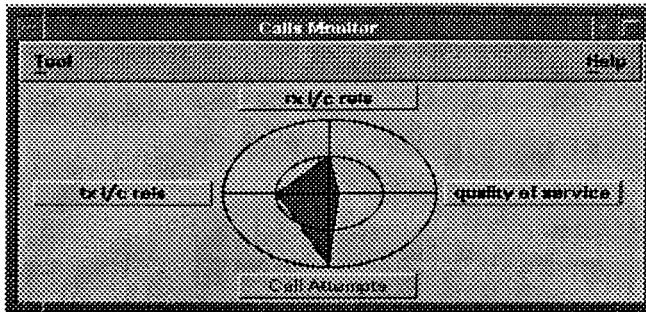
Two other major benefits accrue from automating network monitoring. One is the broad view that results from bringing together the measurement results from several spots in the network. This network-wide view makes it much easier to determine the true root cause of problems, so repair personnel are dispatched to the correct site with the correct resources to effect the repair. Another benefit is that only a few people's skills must be at the high level necessary to make troubleshooting decisions. These are the people at the central site; the field crews only need know how to make the repairs, not how to decide what to repair.

#### 4.5.3 Restoring service to mission-critical applications


Immediate restoration of service requires either spare facilities that can be patched in, or the ability to reroute the traffic around the failed node or segment. Of course, the value of rapid service restoration is highest when combined with a monitoring system that allows early detection of performance degradation, so that restoration procedures can be accomplished before the customer complains.

#### 4.5.5 Rapid fault location and repair tools


Detecting a fault and locating it are two different things. Automatic monitoring is excellent for detecting faults and usually can provide sufficient localization to initiate



Screen A



Screen B



Screen C

**Figure 4.6** Screen A shows four parameters (one per radial axis on the polar charts) that a monitoring system tracks from many points on the network. The circles represent thresholds (two per axis) as defined by the user. The color of the display changes from green to yellow if any parameter crosses the first threshold, and to red if the second. Clicking on any radial axis produces Screen B, showing the data from the monitoring points for the selected axis. Clicking on the reporting location that exceeds the threshold (top bar) produces Screen C, a time history of that location.

rapid restoration via spare facilities or rerouting. Often, however, someone with portable test tools actually pinpoints the problem and carries out the repair.

Dispatching a trained person to a local site is time-consuming and can be expensive. It is therefore important that he or she has a complete set of tools to complete the job with one visit. In many cases, the trouble site is inside a building, perhaps through security checkpoints and an elevator ride away from the parked repair van. This calls for a multifunction tester that can facilitate troubleshooting and confirm



satisfactory repair/restoration of service—and all this under a single handle in a package about the size of a briefcase.

The toolbox for the troubleshooter or repair person should include the following:

- The equivalent of a continuity tester for the service involved.
- An interface breakout box for the service involved.
- A full-service tester that facilitates troubleshooting and allows the operator to confirm that service has been restored.

## 4.6 Communications Standards

The design and production testing of communications network elements, as well as installation and ongoing maintenance and management, depend heavily on national, regional, and international standards. The following information provides a brief introduction to how these standards are set and how to get more information.

Communications standards come from three sources:

- A dominant provider or user of services and equipment.
- A consortium or forum of providers.
- Accredited standards bodies (usually government-sponsored).

Standards set by a single dominant commercial organization are called *de facto*. Others are often called *de jure*, although few carry sufficient weight of law to bring compliance. In fact, most truly de jure standards relate to safety, environmental effects, etc. These are usually referred to as *mandatory standards*. Most standards emanating from industrial consortia or from accredited standards bodies are voluntary consensus standards. A vendor is free to offer products outside the voluntary consensus standard, but may find few takers. Likewise, a service provider may choose to build a network of nonstandard elements, but in so doing also chooses to remain an island with no interconnection to standard networks.

### 4.6.1 Some important communications standards bodies

Almost every nation in the world has at least a national standards organization; developed countries may have more than one, plus industry associations. These national standards bodies may be grouped into regional organizations before connecting to the world communications standards of the ITU, the ITU-T (telecommunications), and the ITU-R (radio).

This is by no means an exhaustive list. The World Wide Web is such a rich source of links to other relevant sources, however, that from one of these starting points, any standards organization ought to be locatable.

**ANSI** American National Standards Institute (New York)  
<http://www.ansi.org> (links to T1 and X3 committees)

**ATM Forum** Asynchronous Transfer Mode Forum (Mountain View, Calif.)  
<http://www.atmforum.com>

## 90 Network Test and Measurement

**ATSC** Australian Telecommunications Standardization Committee  
*http://www.standards.com.au/~sicsaa*

**CTIA** Cellular Telecommunications Industry Association (USA)  
*http://www.wow-com.com*

**DAVIC** Digital Audio Video Council (Geneva)  
*http://www.davic.org*

**DVB** Digital Video Broadcasting  
*http://www.dvb.org*

**ETSI** European Telecommunications Standards Institute (France)  
*http://www.etsi.fr*

**IEEE** Institute of Electrical and Electronic Engineers (New York)  
*http://www.stdsbbs.ieee.org*

**IETF** Internet Engineering Task Force (Reston, Va.)  
*http://www.ietf.org*

**ITU** International Telecommunications Union (Geneva)  
*http://www.itu.ch* (links to ITU-T and ITU-R)

**RCR** R&D Center for Radio Systems (Japan)  
*http://www.phsmou.or.jp*

**TIA** Telecommunications Industry Association (USA)  
*http://www.industry.net/tia*

**TSACC** Telecommunications Standards Advisory Council of Canada  
*http://www.tsacc.ic.gc.ca*

# Introduction to Telecommunications Network Measurements

**Hugh Walker**

*Hewlett-Packard Ltd., South Queensferry, Scotland*

## 5.1 Introduction

Network technology is changing at an increasing rate. In the past, major investments in transmission and switching technology took many years to depreciate. Today, however, the pressures of the market and the advances in technology demand more rapid turnover. The unrelenting rollout of new technology creates challenges for new test equipment and maintenance strategies.

The business climate in telecommunications is changing, too. Because of competition and deregulation, combined with the increasing importance of telecommunications for business activities, network operators are becoming more service- and customer-focused. Network performance is measured in terms of quality of service (QoS). Successful delivery is measured by the highest quality at the lowest prices. Network operators also need to bring new technology into service very quickly to create competitive advantage.

## 5.2 Quality of Service

Network quality of service (QoS) can be characterized by five basic performance measures:

1. Network availability (low downtime).
2. Error performance.
3. Lost calls or transmissions due to network congestion.
4. Connection setup time.
5. Speed of fault detection and correction.

**92 Network Test and Measurement**

Network availability and error performance are usually the parameters that service providers guarantee in terms of QoS. Generally these parameters need to be checked while the network is in service (i.e., carrying traffic), using a network management system.

Lost calls and call setup time are the main criteria for measuring performance in switched networks, often indicating whether network planning is keeping up with traffic growth and changing traffic types. The move to common-channel signaling has greatly reduced call setup time, while also increasing system flexibility for offering new services. The growth of Internet traffic, however, with long holding times on the circuit-switched network, has again called into question network performance.

Some network operators now guarantee that they will fix faults within a specified time or pay compensation to the customer. This requires good processes for troubleshooting, well-trained technicians with access to powerful test equipment, and probably the use of centralized automatic test systems for rapid fault finding.

**5.3 Testing Objectives**

An initial reaction to network testing might be that it is something to be avoided if possible because it costs time and money. On reflection, however, effective testing can add value rather than being an expense, and can enhance the network operator's business.

There are three major business problems that are driving operators today:

1. Time-to-market of new products and services.
2. Reducing the cost of delivering a service.
3. Improving and guaranteeing service quality.

One could add a fourth consideration: the need for reassurance about the security of the network and the ability to detect problems before customers find them. No unpleasant and embarrassing surprises. In a volatile and rapidly changing market, it can be just as challenging to retain existing customers as to find new ones.

Thus testing is at the heart of this new business environment. Unless one has good measurements of QoS, one cannot assess competitive strength or make significant improvements.

As discussed later in this chapter, monitoring the traffic on the network yields not only QoS measures but also can provide a useful source of additional information on the business, such as market information and costs.

Network testing can be divided into three application areas:

1. Bringing new equipment and systems into service.
2. Troubleshooting and detecting network degradation.
3. Monitoring and ensuring quality of service.

**5.3.1 Bringing new equipment and systems into service**

When new equipment is installed and brought into service, the installer (who may be the equipment manufacturer) makes a comprehensive series of tests. These tests

usually are made to more restrictive limits than normal performance expectations (expressed as a fraction of the reference performance objective). These limits are specified in ITU-T Recommendation M.2100 (formerly M.550), "Performance Limits for Bringing Into-Service and Maintenance of International PDH Paths, Sections and Transmission Systems." (See Table 5.1.) Because the measurements are made out-of-service, a more extended acceptance test of certain factors, such as error performance, can be carried out over a period perhaps as long as one month.

Once a system is in use, performance standards must be maintained. When a service degradation occurs, it must be determined whether the fault exists within a particular vendor's network or elsewhere. This information is determined most effectively by in-service testing or performance monitoring. Many test instruments also provide some degree of nonintrusive testing.

### 5.3.2 In-service maintenance

Once equipment is in service, long periods of downtime are unacceptable, so maintenance strategy has to be carefully thought out. ITU-T Recommendation M.20, "Maintenance Philosophy for Telecommunications Networks," defines three types of maintenance strategy:

1. Preventive maintenance
2. Corrective maintenance
3. Controlled maintenance

*Preventive maintenance* is carried out at predetermined intervals to reduce the probability of failure or degradation of performance. This method was commonly

**TABLE 5.1 Performance Limits Specified in ITU-T M.2100.**

DIGITAL LINE SECTION		DIGITAL PATH SECTION	
Limit (Relative number impairments)	Performance for staff	Limit (Relative number of impairments)	Performance for staff
Bringing into service	0.1	ACCEPTABLE	ACCEPTABLE
Performance after repair	0.125		
Degraded	0.5	Bringing into service	0.5
Reference performance objective		Performance after repair	
		Degraded	0.75
		Reference performance objective	1
Unacceptable	>10	UNACCEPTABLE	UNACCEPTABLE

## 94 Network Test and Measurement

applied to older analog systems that needed periodic adjustment to compensate for drift.

*Corrective maintenance* is carried out after a failure or degradation is reported by a monitoring system or user.

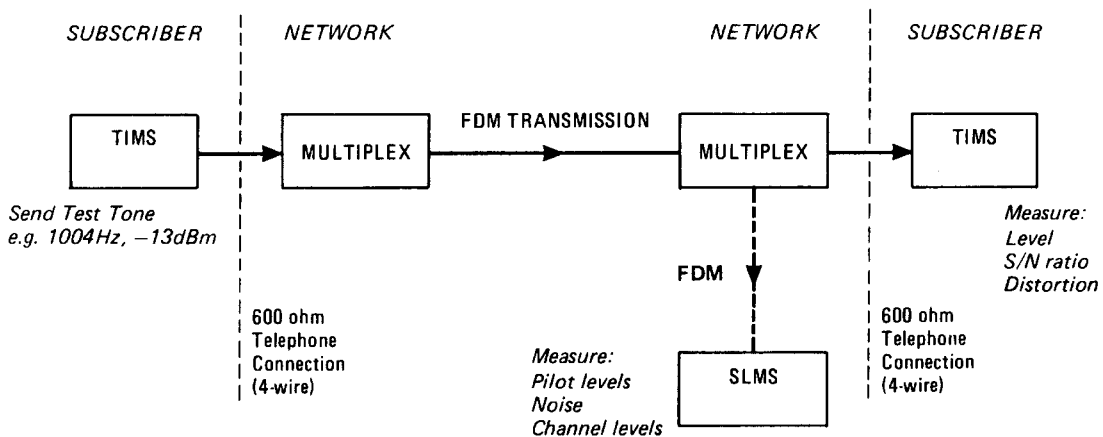
*Controlled maintenance* involves centralized network monitoring and identification of degraded network performance. Centralized monitoring can be supplemented by field maintenance teams using portable test equipment.

Of these methods, controlled maintenance is preferred for maintaining high levels of QoS. It provides early warning of degradations and potential failures, thereby reducing system downtime. Repair work and adjustments can be anticipated and scheduled for quiet periods. In this way, disruption also is minimized.

#### 5.4 Analog Performance Testing

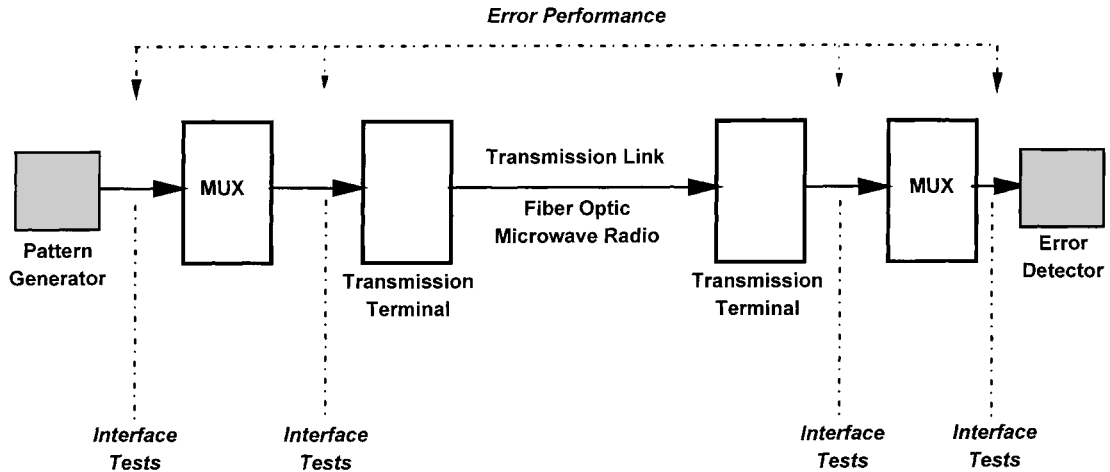
Figure 5.1 shows the major measurable elements of an analog transmission system. The simplest analog test is to measure the system gain and signal-to-noise (S/N) ratio between the end-to-end telephone connections. The test is usually made with a portable *Transmission Impairment Measuring Set* (TIMS). The test operator sends a fixed tone into the system and makes measurements at the opposite end to check for signal level and signal-to-noise ratio (noise with tone). When an analog data modem is to be used on the path, various data-impairment measurements may be specified, such as impulse noise, phase jitter and gain/phase hits (Walker 1989).

Although telecommunications networks are now largely digitized, the connection between a telephone and the exchange continues in most cases to be analog. TIMS measurements are therefore still important. In the 1990s, wideband TIMS measure-



**FDM = FREQUENCY DIVISION MULTIPLEX**

**Figure 5.1** Analog system performance measurements can be made either at the local loop access voice band frequencies using a Transmission Impairment Measuring Set (TIMS), usually at the 4-wire 600-ohm line, or at the Frequency Division Multiplex (FDM) line level using a Selective Level Measuring Set (SLMS). At the multiplex level the frequencies are much higher, typically starting at 60 kHz and going up to 65 MHz. A range of analog transmission measurements can be made with both test sets, as shown in this diagram.



**Figure 5.2** Digital transmission measurements fall into two main categories. *Interface tests* check the compatibility of the electrical or optical interfaces of equipment to ensure error-free interconnection. *Error performance measurements* are usually made with a digital transmission analyzer or BER tester, with the objective of detecting any bit errors in the transmitted data stream. Error performance tests can be made either in-service or out-of-service.

ments up to 200 kHz have been used to evaluate local loops for ISDN and digital data transmission.

Similar kinds of measurement can be made in the analog multiplex system using a *Selective Level Measuring Set (SLMS)*. Because this is a Frequency Division Multiplex (FDM) system, possibly carrying several thousand separate telephone channels in a bandwidth of up to 65 MHz, the SLMS has to be able to select and measure individual channels as well as the pilot tones inserted to control system levels. FDM systems operate either over coaxial cable or microwave radio.

An FDM multichannel traffic signal resembles white noise; system impairments create degraded signal-to-noise ratio in individual telephone channels due to intermodulation distortion, particularly under heavy traffic loading. To evaluate this, the out-of-service noise-loading test is made using a *notched white noise test stimulus*. By measuring the noise level in the notch at the receiving end, the equivalent signal-to-noise degradation can be estimated as a function of traffic load. In the analog era this was a very important test, particularly for microwave radio, because impairments are additive in analog systems.

## 5.5 Digital Performance Testing


The tests made on digital transmission systems can be divided into several categories, as shown in Figure 5.2 and Table 5.2.

### 5.5.1 Interface specifications and tests

Anyone familiar with RF and microwave knows the importance of *matching* at interfaces so that the performance of cascaded networks equals the sum of the parts. The same is true in digital communications. If instrument parameters do not match

**TABLE 5.2 Categories of Digital Transmission Tests and Appropriate ITU-T Recommendations.**

Type of test	Typical tests	Relevant ITU-T standards
Interface tests	PCM Codec	G.712/713/714 (O.131-133 measurement)
	Pulse shape	G.703
	Clock frequency	
	Voltage/impedance	
	Coding	
	Framing	G.704/706/708
Out-of-service error performance tests (Installations and commissioning)	Jitter wander	G.823/824/825 (O.171 measurement)
	BER using PRBS patterns	G.821/826 (O.151 measurement)
In-service error performance tests (maintenance, fault finding, quality of service)	Code errors	G.821/826
	Frame errors	M.2100/2110
	Parity errors	

**Figure 5.3** Interface specifications, with a sample of a standard pulse mask for checking the transmit pulse shape to the ITU-T recommendation G.703.

equipment parameters, bit errors appear when they are connected. This matching is defined in a series of interface specifications contained in ITU-T Recommendation G.703; Recommendations G.823/824 and G.825 address timing jitter.

Electrical interface specifications (Figure 5.3) are usually measured during equipment design and manufacture to ensure compatible interconnection between network elements at a Network Node Interface (NNI) and User Network Interface (UNI).




The ITU-T specifications include pulse height (voltage level); pulse shape (rise time, fall time, overshoot, and duty cycle); and equality of positive and negative pulses in a ternary signal. These measurements usually are made with an oscilloscope to check that the pulse shape falls within a prescribed mask.

The physical interface is usually 75-ohm coaxial cable with a return loss of 15–20 dB, although with higher-speed SONET/SDH equipment the physical interface may be fiber optic. In addition, bit rates must be maintained within strict limits (see Table 26.9), and the tester must check that receiving equipment can operate correctly within this tolerance. Interface coding specifications include algorithms for AMI, HDB3, CMI, B3ZS, etc. (For a discussion of encoding mechanisms, see Chapter 3, “Telecommunications Basics.”)

*Timing jitter* is defined by ITU-T as short-term variations of a digital signal’s significant instants from their ideal positions in time. The *significant instant* might be the rising or falling edge of a pulse. Figure 5.4b shows the occurrence and impact of jitter; at certain points in time, the pulse is significantly offset from its correct position. If this offset becomes large, then there will be an error when the receiver attempts to sample and decode the digital signal.

The simplest way to measure jitter is with an oscilloscope and *eye diagram*; the “eye” area is shown in Figure 5.4a. Jitter appears as a spread or “muzziness” in the vertical transitions. Most telecommunications systems, however, require more precise



**Figure 5.4** Timing jitter disturbs the pulse from its ideal position in time, and the perturbations cause a narrowing of the eye area as shown in (a). Examined in real time (b) at instants  $T_1, T_2, T_3$ , and so on, one can see that the bit pattern is displaced from the ideal positions in time. The instantaneous offsets  $t_1, t_2, t_3$  form the Jitter Function  $J(t)$ . If jitter becomes excessive, the eye opening will be closed sufficiently to cause errors when sampling the data. Sampling is usually timed to occur at the center of the eye, at the point of greatest eye height.

## 98 Network Test and Measurement

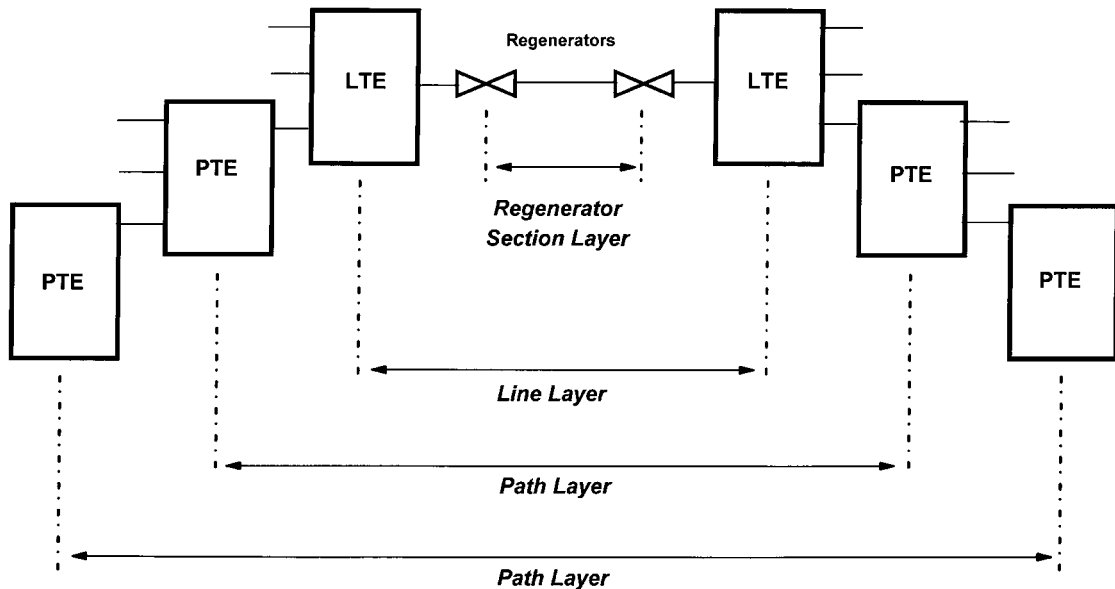
measurements. In these cases it is essential to know how the level of jitter varies with jitter frequency. This relationship is measured with a jitter test set that demodulates the jitter signal.

Jitter itself must be checked at the output of the equipment under test. The system's tolerance to jitter must be checked at the input by gradually increasing the level of jitter on a test signal until bit errors occur. In addition, the tester must check that jitter present at the input is not magnified by the equipment; otherwise, problems can arise when several pieces of equipment are cascaded in a network. This measurement is called *jitter transfer*.

If equipment conforms fully to all the interface specifications, in principle it should be possible to construct any arbitrary network without generating bit errors. Problems still can arise, however, if the live traffic signals have very extreme pattern densities that are not fully simulated by the out-of-service PRBS test.

## 5.5.2 Error performance tests

Digital network error performance can be measured over a complete end-to-end connection called a *path*, or over parts of the network called *lines* and *sections*. These network segments are illustrated in Figure 5.5. Path measurements indicate the overall quality of service to the customer. Line and section measurements are used for troubleshooting, installation and maintenance, and for assuring transmission objectives are met.



**Figure 5.5** A digital transmission system can be viewed as an overall end-to-end path terminated by Path Terminating Equipment (PTE). The system is made up of lines and sections terminated by Line Terminating Equipment (LTE). Sometimes paths are described as *low-order paths*, implying they are the end-to-end service provided to the customer. *High-order paths* exist within the network at a multiplexing level and indicate the extent of a path for error performance monitoring, such as a virtual container in an SDH system.

The fundamental measure of performance or quality in digital systems is the probability of a transmitted bit being received in error. With the latest equipment, the probabilities of this occurrence are very low, on the order of  $10^{-12}$  or less. It is still necessary to measure the performance of these systems, however, and in particular to analyze the available margins of safety, and to explore potential weaknesses that later could lead to degraded performance.

### 5.5.3 In-service and out-of-service measurements

In-service error performance measurements rely on checking known bit patterns in an otherwise random data stream of live traffic. As discussed in Chapter 27, some in-service measurements are more representative than others of the actual error performance of the traffic signal. Furthermore, some are applicable to the path measurement, provided the parameters are not reset at an intermediate network node. Others are only useful at the line or section level. The most commonly used error detection codes (EDCs) are *frame word errors*, *parity errors*, or *cyclic redundancy checksum errors*.

Out-of-service measurements involve removing live traffic from the link and replacing it with a known test signal, usually a *pseudorandom binary sequence* (PRBS). These tests are disruptive if applied to working networks, but are ideal for installation and commissioning tests because they give precise performance measurement. Every bit is checked for error. Although the PRBS appears random to the digital system, the error detector (Figure 5.2) knows exactly what it should receive and so detects every error. The error detector calculates the probability of error as the *bit error ratio* (BER). BER is defined as the number of errors counted in the measurement period, divided by the total number of bits received in the measurement period.

Thus the bit errors or error events can be detected by out-of-service or in-service techniques. These are sometimes referred to as the *performance primitives*. To be useful for assessing quality of service (QoS), however, they must be analyzed statistically as a function of time according to the various error performance standards specified in Table 5.2. This analysis yields percentages for the availability of a digital communication link, and the portion of time that it exceeds certain performance criteria that are acceptable to the customer. One of the most important standards is the ITU-T Recommendation M.2100/2110.

## 5.6 Protocol Analysis in the Telecommunications Network

Up to this point we have discussed the capability of the telecom network to transmit digital bits or analog signals over a path without errors or quality degradation. Testing BER, for example, assumes that the traffic carried by the network is completely random data, or at least that the payload within a frame structure is random. This apparently random traffic signal will, in fact, always have a structure. It might be a PCM voice signal, a data signal, a signaling message for controlling network switching, or possibly an ISDN signal or an ATM cell data stream for broadband services.

When telecom networks were predominantly carrying voice traffic using in-band signaling, there was little interest in checking the information content of the traffic

or knowing if it conformed to the rules or protocols of data communications, except for the X.25 packet-switched network.

Networks today are much more sophisticated and carry a wide range of different services among many vendors. Rather than being just the transporter of telecommunications traffic, increasingly the network is an integral part of the information structure created by the convergence of computers and communications. The most significant example of this is the *common-channel signaling system (SS7)*, which interconnects the network switches and databases and controls all aspects of service delivery and billing. A large amount of analysis and monitoring is required, not so much of the data transmission itself, but of the messages and transactions taking place. An important example of signaling transactions occurs in a cellular telephone network, when constant reference to databases is necessary for tracking the location of mobile phones during handover from one cell to the next, and for billing and verifying legitimate users.

Protocols are based on the seven-layer Open System Interconnection Reference Model, shown in Figure 5.6 (ITU-T Recommendations X.200 and X.700). Each layer in the model has a specific, independent function that provides a service to the layer above and is supported by the layer below. Protocols are the rules that govern transactions within and communication between layers. In theory, observing these rules should allow the free interconnection of equipment from different vendors and between different networks, because each will interpret the messages in the same way.

## THE OPEN SYSTEM INTEGRATION (OSI) MODEL

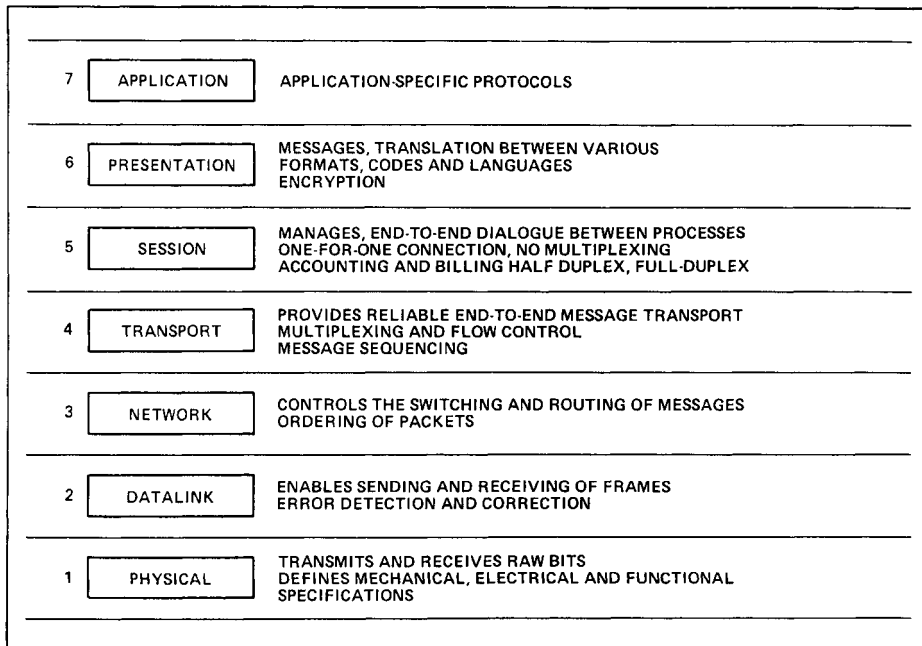
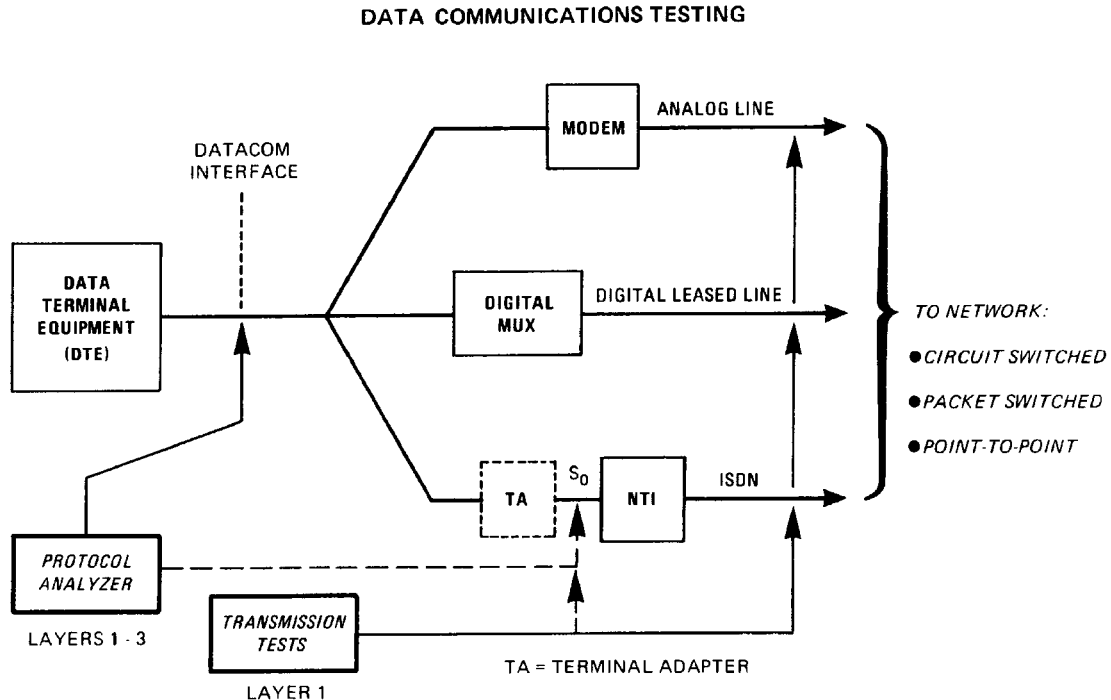


Figure 5.6 The seven-layer OSI protocol stack.



**Figure 5.7** Protocol analyzers traditionally have been used for testing data communications networks at datacom interfaces, usually at the customer's premises. As protocol testing has moved into the telecommunication network, protocol analyzers now often have standard telecom interfaces for ISDN and ITU-T G.703. Protocol analyzers provide checks of the transactions taking place within the OSI layers, as well as communication between layers.

Digital performance measurements discussed in section 5.5 relate to the Physical layer (layer 1), which is concerned with the error-free movement of data through the network. The standard frame structures and error-detection codes discussed in section 5.5.3 perform some of the Data Link layer (layer 2) functions, but not the confirmation of receipt and retransmission of errored frames normal in a data communications protocol. The basic telecom network is thus not truly protocol-oriented—but increasingly is carrying services that are.

The SS7 signaling system, for example, uses this infrastructure but builds a complete protocol stack of signaling transactions on top of it. ISDN and ATM switching are other examples of protocol-heavy services now carried by the telecom network. With such services, network behavior and performance will be influenced by the message content of the traffic, which will include routing information, priority, and so on.

In order to analyze and troubleshoot these systems, a protocol analyzer is required, in some cases dedicated to the particular application such as SS7 or ATM. Protocol analyzers have been in use for many years in local and wide area networks, predominantly in enterprise networks (see Figure 5.7). This traditional data communications test tool is now finding its way into the telecom network as traffic becomes more data-intensive.

## 102 Network Test and Measurement

Often in data communications there is the need to observe and analyze, or even to simulate, the interactions between network devices interconnected by WANs or LANs. The need may be in the context of one or more of the following scenarios:

- The developer of network equipment or of the network itself needs to analyze and simulate operation under a number of circumstances.
- The network planner needs to measure current levels of network use and then anticipate future needs as new services and capabilities are added.
- The installer (of computers, communications equipment, and/or networks) needs to commission and test a network system's devices and their interactions.
- Field service personnel for a computer and communications equipment vendor, or for a service provider, are faced with troubleshooting an intermittent problem.
- The network manager of a private network operates with system elements from several vendors and uses multiple service providers in order to get the best performance, reliability, and price. When a problem arises, a tool is needed to determine its source so as to avoid finger-pointing among the vendors.

In each of these scenarios, there is need for an instrument that can observe non-intrusively and help the user interpret the complex interactions within the data communications protocols that control the behaviors of the devices. In some cases there is need to simulate network elements to test for problems. In other situations, there is an application to measure the performance and utilization of the network and of the devices within it.

These tests and measurements may be made reactively when a problem occurs, or may be made proactively when looking for trends that indicate developing problems. When new services are being introduced, or new equipment is being installed or system software upgraded, it is necessary to emulate specific messages or protocols to confirm correct operation of the network. In all these cases, an appropriate tool for solving network problems is the protocol analyzer, described in more detail in Chapter 27.

In the case of SS7 signaling networks, a number of recommendations have been issued by ITU-T for maintenance and protocol testing. Some of the more commonly used recommendations are shown in Table 5.3. Monitoring and analysis is usually done at the Signaling Transfer Points (STPs), which are the packet-switching hubs

**TABLE 5.3 ITU-T Recommendations for Testing and Maintenance of SS7 Networks.**

Maintenance strategy	Protocol testing	
	Protocol layer	Recommendation
M.4100	Level 1/2 Message Transfer Part (MTP 2)	Q.781
Q.750	Level 3 Message Transfer Part (MTP 3)	Q.782
Q.752	Level 4 Telephone User Part (TUP ISDN User Part (ISUP))	Q.783 Q.767

in the SS7 network. These hubs contain the routing tables for calls and also are the most likely sites of congestion under heavy traffic (McDowall 1994).

## 5.7 Centralized Control

As discussed earlier in this chapter, the need to reduce operating costs and improve quality of service is driving network operators toward centralized maintenance and network management. The benefits are twofold. First, it allows skilled technical staff to be more productive as they can control maintenance and troubleshooting over a wider area. Second, network monitoring can provide a real-time view of the state of the network, allowing the operator to assess quality of service and detect hot spots and system degradation, ideally before customers see a reduction in service quality.

At the first level it is possible to control test instrumentation remotely, so that measurements can be made by a skilled technician from a central site. This is sometimes referred to as *virtual remote operation* using modern PCs or workstations; the front-panel operation of the remote instrument can be replicated and controlled from the display screen at the central site. Increasingly the demand is for full network monitoring or *operational support systems* (OSSs), which provide real-time monitoring at hundreds or even thousands of test points across the whole network. These computer-controlled systems can acquire measurement data either from system monitors built into the operational equipment, or by add-on measurement probes sited adjacent to network nodes.


### 5.7.1 Virtual remote capability

The scope and power of all the transmission and protocol measurements discussed previously can be expanded by linking the test equipment into a control center (Figure 5.8). Individual instruments can be left to monitor network nodes. When these instruments are connected sequentially to each transmission path, accumulated data can be transferred back to a central PC under command of the network operator. Any time problems are reported, the instruments can be instructed to monitor particular channels continuously.

A remote testing development, one that takes advantage of PC/workstation hardware and software, is the so-called virtual remote capability. Using the Microsoft Windows® environment, lifelike representations of remote instrument front panels can be displayed on a centralized desktop computer, as shown in Figure 5.9.

The instruments themselves communicate with the workstations via a simple data link. By using the mouse, instrument function keys (“soft keys”) can be pressed on the workstation and the effect replicated at the remote instrument. Measurement results displayed at the instrument are likewise relayed back to the workstation screen.

In current systems, communication and control is two-way. If the engineer at the remote site operates the instrument and makes measurements, these actions are relayed back to the centralized control point. This capability is different from conventional remote-control systems, which usually lock out local control.



**Figure 5.8** A simple remote test system in which a PDH frame analyzer is connected to many different test monitor points in a 140 Mbps transmission network via an access switch matrix. The PC can scan all the test points sequentially looking for trouble, or can connect to a specific point for more detailed analysis and troubleshooting. Via dial-up modem links, the PC can control geographically dispersed test sets, considerably increasing productivity of technical staff.

The power of the new generation of computers can be harnessed to make possible centralized measurements while requiring very little extra training; the operator is working with a user interface that looks and feels like the real instrument. For network operators with scarce resources of expert, highly trained engineers, this option reduces time wasted on travel and spreads expertise over a larger number of sites. Furthermore, the remote instruments can be used at any time as normal, portable field test sets just by disconnecting the data link.

### 5.7.2 Network monitoring

The solutions described here represent methods of improving efficiency in stand-alone instrument applications, rather than being true network monitoring systems. Concepts such as virtual remote are a way of improving productivity in maintenance strategies that are based on portable instruments. Network monitoring systems require more investment and planning, but provide greater potential for quality improvement. Advantages of network monitoring include:



- More complete and continuous monitoring at multiple nodes.
- Concentration, reduction, and graphical display of measurement data.
- Use of databases for network configuration and traffic statistics (such as network maps, historical statistics, etc.).

Network monitoring systems for many years have been of proprietary or nonstandard design. Some of these systems were supplied by the original network equipment (NE) manufacturers and might operate only with a specific class or model of equipment. Similarly, the management reporting and user interfaces often differ among proprietary systems. As a result, current network monitoring might rely on nonstandard physical or measurement interfaces, communication protocols, and operating software. Furthermore, the monitoring system might use some level of built-in measurement, or might rely on an overlaid measurement system.

### 5.7.3 Operational support systems

The move to integrated digital networks means that network monitoring and testing can be standardized and centralized as part of a network management system, called



**Figure 5.9** A typical display on the PC screen of a virtual remote test system. The front panel of the remote test set is replicated at the controlling PC. The front panel can be operated by pointing and clicking the mouse, while the displayed results from the remote tester are relayed back to the PC and appear as they would at the remote site. Virtual remote systems allow both ends to control the test set for maximum flexibility when troubleshooting.

an *operational support system* (OSS), also referred to as an *operations system* (OS). An OSS uses the built-in testing capabilities of the transmission and switching equipment, which ITU-T refers to as the *Telecommunications Management Network* (TMN), described below. An OSS/TMN reduces the amount of manual testing. It does not, however, eliminate altogether the need for standalone equipment. Portable test tools remain essential for corrective maintenance, network engineering, and stress testing.

#### 5.7.4 The Telecommunications Management Network (TMN)

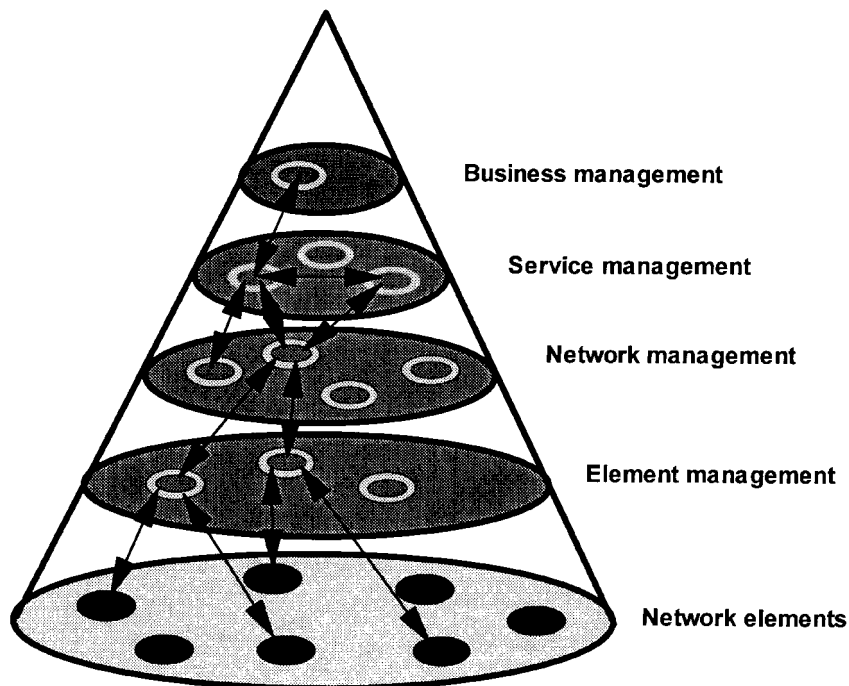
The ITU-T recommendation for developing a managed telecommunications network for the controlled maintenance strategy is described in Recommendation M.3010. The Telecommunications Management Network (TMN) provides a framework of standards for network structure, interconnection protocols and interfaces, and management applications such as performance, administration, provisioning, and maintenance. The objective is to establish an open system architecture for network management, similar to the OSI model for data communications protocols, to facilitate the integration of multivendor systems.

TMN uses a layered model similar to the OSI Reference Model. In it, each layer supports the one above in ascending application order. The five layers are: the Network Elements layer, the Element Management layer, the Network Management layer, the Service Management layer, and the Business Management layer (Figure 5.10). Data is collected from many different network elements (probably manufactured by different suppliers) and processed to provide uniform management information on network, service, and business applications.

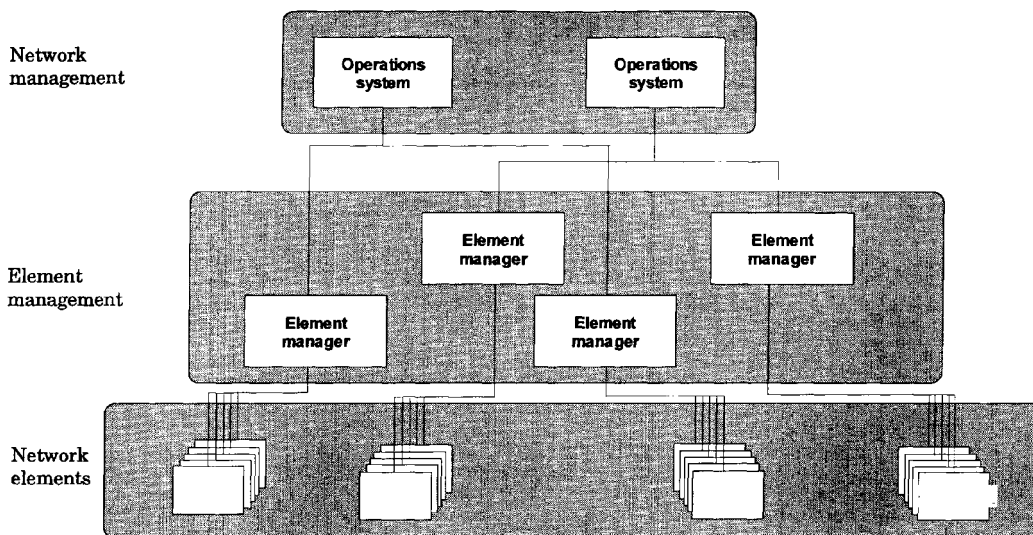
The TMN model includes specified  $q$  interfaces and reference points that provide isolation of the Network Management layer (the management applications) from the network elements. The Element Management layer adapts the built-in measurements of multivendor network equipment and isolates network element technology development from the operations support systems (Figure 5.11).

The diagrammatic representation of the TMN layered model as a triangle or cone (Figure 5.10) implies that the raw data from many network elements converges on one or two service support or business support systems. From the network operator's perspective, the triangle might be inverted; the importance of the business support and service support is far greater than the individual network elements of the telecommunications infrastructure.


Network management systems or OSSs were installed initially to enhance fault and performance analysis—in other words, to help guarantee Quality of Service. Network operators now are also taking advantage of powerful, real-time data collection engines to provide information for business-related applications such as marketing, customer service, billing, and fraud detection, all of which become increasingly important in a competitive and deregulated market (Figure 5.12). In parallel with the move to protocol-oriented telecommunications networks, there is the realization that a gold mine of information is waiting to be tapped from what previously was a test and measurement QoS performance monitoring system (Urquhart 1996).



**Figure 5.10** The five layers of the TMN management hierarchy. As with the OSI stack, each layer provides a service to the layer above, but is functionally separated through standard interfaces and protocols. This should allow the development of multivendor network management systems that can provide unified network management applications with a variety of network elements from different manufacturers.



**Figure 5.11** The network elements that could be supplied by many different manufacturers interface to the network management layer via the element management layer. The network management software is thus quite separate from the individual network elements and overcomes the problem of different proprietary network management systems supplied with each different type or make of network equipment.

**Management Levels****business management****service management****network management****element management****network elements**

- ▶ Interoperator Billing
- ▶ Transit Billing
- ▶ Marketing Statistics
- ▶ Billing Verification
- ▶ Fraud Management
- ▶ Quality of Service
- ▶ Service Assurance
- ▶ Fault Management
- ▶ Performance Management

INCREASING VALUE

**Figure 5.12** The data gold mine. Collecting the raw data on network performance can yield a whole range of useful information for revenue generation and protection. This diagram shows some of the higher-level capabilities of the HP *acceSS7* monitoring system which extracts data from the SS7 network. This can be used for conventional applications such as fault/performance management and QoS, but also can be used to drive billing systems and fraud detection at the service and business management levels.

**5.8 References**

- McDowall, Ron. "When SS7 Is Put to the Test." *Global Telephony* (April 1994).  
 Urquhart, Reid. "Mining for Gold." *Telephony* (June 24, 1996).  
 Walker, Hugh. "Testing and Troubleshooting Datacom Circuits." *Evaluation Engineering* (May 1989).

# Conformance and Interoperability Testing

**Jean Boucher**

*Hewlett-Packard (Canada) Ltd., Montreal, Canada*

In this era of network communications, the quality and marketability of new products depends on their ability to interoperate. To meet this need, international organizations define standards that are adopted by all member countries. Because they must be ratified by a consensus of countries with different needs, standards may end up as vague and incomplete frameworks. This looseness leaves room for countries to adapt the standard to their national needs.

Standards consist of specifications such as electrical, mechanical, and functional. Companies then use these specifications as guidelines for product development. If a specification can be understood precisely, the products developed from it will be compatible. If specifications are ambiguous, inconsistent, or incomplete, however, two similar products developed by two different companies from the same specification might not work together at all. The primary role of testing is to uncover these problems.

## 6.1 Testing Methodologies

Several methodologies exist for verifying compliance to specifications. Among them are:

- Conformance testing
- Interoperability testing
- Regression testing
- Acceptance testing

Each of these mechanisms assures a product's compliance to a standard or contract at a different stage in its development.

### 6.1.1 Conformance testing

The first step in verifying the correctness of a product implementation is to ensure that it performs in accordance with the specification upon which it was built. This process is called *conformance testing*. In practice, the role of conformance testing is to increase confidence that a product conforms to its specification, and to reduce risk of malfunctioning when the product is put into place (for example, into an ATM network).

Conformance testing is a well-established testing methodology based on the multipart ISO 9646 international standard. Figure 6.1 depicts the major steps involved.

The process starts with a specification. For each characteristic in the specification, a *Test Purpose* is written. An example of a Test Purpose from the ATM specification is as follows:

Verify that the IUT supports point-to-point VC [Virtual Channel] connectivity, where the IUT is the Implementation Under Test.

Next, each Test Purpose is turned into an *Abstract Test Case* (ATC). An ATC explains (with all necessary details) what is sent to the IUT, what is expected from the

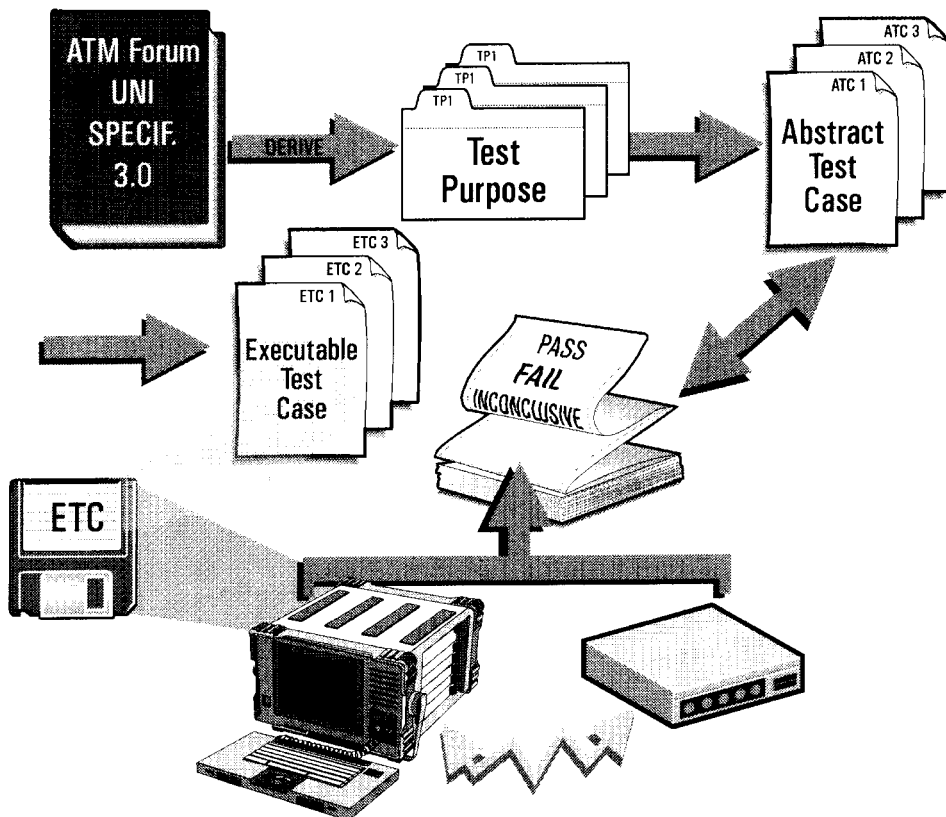


Figure 6.1 The conformance testing process.

IUT, what the IUT must do in order to pass the test case, how the IUT can fail the test case, etc.

Multiple test cases are required to cover all aspects of a given protocol specification. When all aspects have been covered, the resulting collection of ATCs form an *Abstract Test Suite* (ATS). For example, the ATM Forum Cell Layer Conformance Test Suite contains 46 test cases. ATSs for other well-known protocols are even larger. The ATM Forum Unit 3.1 Signaling Conformance Test Suite for the Network Side contains 66 test cases.

Once it is implemented on a particular piece of hardware (a protocol analyzer, for example), the ATS becomes an *Executable Test Suite* (ETS). An ETS offers the user the ability to select one or more test cases from the whole test suite, and to run those test cases against an IUT and generate a Test Report. If the Test Report uncovers any IUT error, the product designer can fix the problems found and rerun the ETS.

An important feature of conformance testing is that each test case, when run against an IUT, gives a clear and single verdict: either *Pass*, *Fail*, or *Inconclusive* (which means that the same test should be rerun to get a Pass or Fail verdict). These verdicts appear in the Test Report, along with a detailed trace of each test case run. Aspects of test suite execution and the contents of a Test Report are described in section 6.6.3.

### 6.1.2 Interoperability testing

*Interoperability testing* is the next logical step after conformance testing. While conformance testing can increase confidence that system A conforms to specification X and that system B also conforms to specification X, interoperability testing evaluates the extent to which systems A and B actually can work with one another. Figure 6.2 illustrates the major steps of the interoperability testing process:

1. By extrapolating the specification under real-life situations, Test Purposes are defined and an ATS is written.
2. From this ATS, an Executable Test Suite (ETS) is implemented.
3. The ETS is executed on a protocol analyzer against two or more *Systems Under Test* (SUTs).
4. A Test Report is generated from this test campaign (process) which again might uncover errors in an SUT.

Interoperability testing has several points in common with conformance testing. At first glance the ATCs and ATSs of both test types appear similar. The main difference between them is that interoperability testing verifies several SUTs at the same time (Figure 6.3). The processes of implementing the ATS and running the ETS are the same. An interoperability test results in one of the same three verdicts: Pass, Fail, or Inconclusive. Finally, Test Reports present results in a similar fashion.

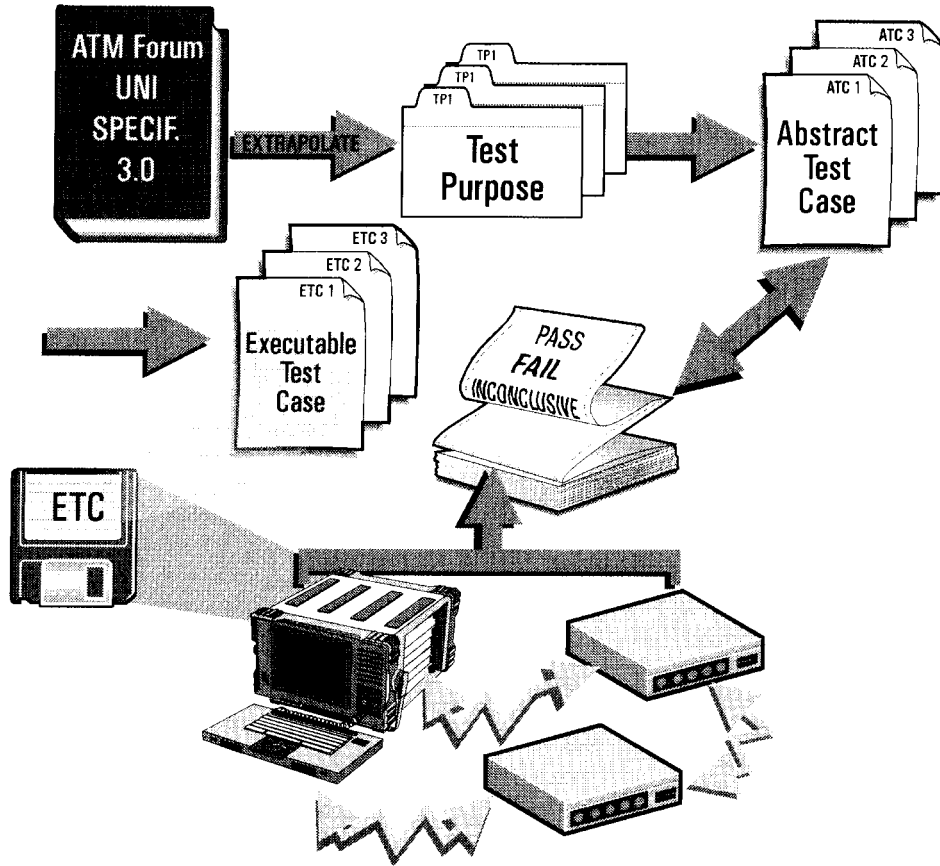


Figure 6.2 Interoperability testing.

- Conformance:  
Tester against 1 IUT.

- Interoperability:  
Tester against  
2 or more SUTs.




Figure 6.3 Conformance vs. interoperability differences in practice.



### 6.1.3 Conformance versus interoperability testing

Despite the similarities just noted, there are fundamental differences between the two testing techniques.

Conformance testing is comparable to an absolute measurement. It looks at one aspect of the specification at a time and provides one test case (or a very few) to verify the behavior of an IUT in regard to this aspect (Figure 6.4a). Interoperability testing, on the other hand, is comparable to a relative measurement. It takes one aspect (or a few related aspects) of the specification and aims to verify how these aspects are handled by two communicating network components (Figure 6.4b). This difference often leads to more comprehensive test cases than does conformance testing.

For example, using the ATM specification sample shown above, we see that the test purpose is to verify that the IUT supports point-to-point VC connectivity. In order to satisfy this purpose, the Conformance ATC sends only one cell to a VC, while




Figure 6.4a Conformance will test VPI or VCI, minimum or maximum, one at a time.




Figure 6.4b Interoperability will test VPIs in a given range.

## 114 Network Test and Measurement

the corresponding Interoperability ATC sends multiple cells into the originating endpoints of a Virtual Channel Connection (VCC).

In another example, the Conformance Test Suite verifies that the IUT relays cells on a VC connection for the minimum (or maximum) nonreserved Virtual Channel Identifier (VCI) value supported by the IUT. Interoperability testing, on the other hand, verifies that two SUTs can communicate over the overlapping ranges of VCI values common to both SUTs.

Table 6.1 compares the scope of conformance and interoperability testing for several aspects of an ATM system. As can be seen, conformance testing tends to verify basic protocol features, while interoperability testing is designed to replicate real-life scenarios. Although interoperability looks more appealing than conformance testing, the latter is a first and mandatory step. If this were not true, it would not be possible to determine whether an interoperability test case failure could be attributed to SUT A or SUT B.

### 6.1.4 Regression testing

It is rare that only one version of a product is put on the market during its life cycle. Subsequent versions of a product can include bug fixes, enhancements, additional features, etc. *Regression testing* is a technique to ensure that the existing features of an IUT migrate properly as the products evolve (Figure 6.5).

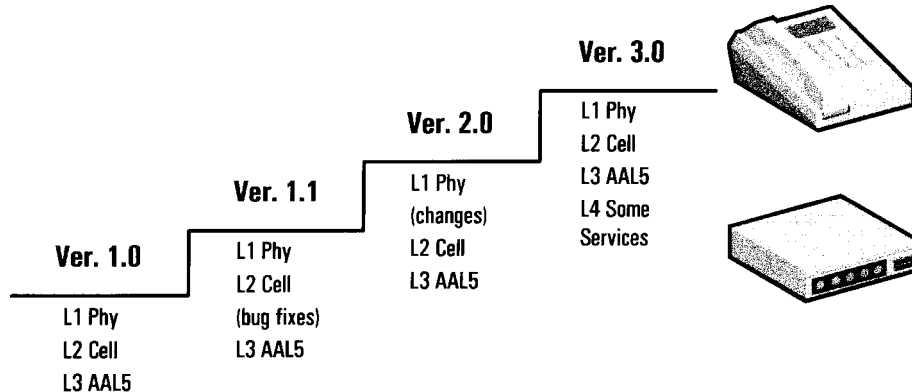
For example, if version 2.0 of an ATM product fixes some known errors at the ATM Cell Layer, regression testing ensures that the existing (and unchanged) Adaptation Layer still performs as well as it had with the previous Cell Layer.

### 6.1.5 Acceptance testing

*Acceptance testing* groups a series of predefined requirements, often culled from a procurement document. A series of tests is performed (such as conformance to a protocol specification, interoperability between vendors, performance, and tests of specific in-house features required by the manufacturer). Usually these tests are performed at the customer premises after the product has been ac-

**TABLE 6.1 ATM Conformance vs. Interoperability.**

• What can Conformance testing do?	• What can Interoperability testing do?
—Verify that the IUT supports point-to-point VP connectivity.	—Verify that two SUTs can communicate over the overlapping range of VCI values common to both SUTs.
—Verify that the IUT relays cells for a given VC while preserving cell sequence integrity.	—Verify the transparency of F5 OAM end-to-end cells when both SUTs support VCC service.
—Verify that the IUT supports VP management, it has the capability to identify and decode OAM F4 flow cells.	—Verify the ability to pass traffic in both directions simultaneously at different cell rates.



**Figure 6.5** Regression testing, a series of existing tests that make sure the existing features of an IUT are migrating with the product's evolution.

cepted. In normal practice, acceptance testing is performed at the buyer's premises with the understanding that the product will not be accepted until all conditions are met.

Acceptance testing can include regression testing whenever a customer receives modifications to previously purchased products.

## 6.2 When to Test

One testing strategy used before introducing a new product into the market is to perform minimal in-house testing after product development is complete, then move to a controlled test laboratory for full-scale testing. In a field such as networking, where time-to-market is critical for product introduction, it can be tempting to save time by skipping full-scale testing prior to manufacture, and to ensure product quality by completing the testing at external sites. The major drawback of this strategy is that it often uncovers problems only during the last phases of a development cycle, resulting in costly and time-consuming redesign and re-implementation.

A better solution is to perform conformance and interoperability testing during product development. This scenario allows testing to be performed:

- *By application engineers*, to test each feature of the IUT as it is implemented. Some providers' test suites allow the user to select only a portion of the series of test cases in order to facilitate this type of testing. This strategy uncovers problems early in the design cycle, where they can be fixed at the least cost, and allows engineers to add features and build atop a solid product.
- *By test and quality assurance engineers*, who can run all applicable test cases against an IUT, either before it goes into production or before it is shipped to customers. Saving test setup and results can help systematize testing from one IUT to another, and facilitate regression testing.

## 6.3 Test Suite Development, Execution, and Results

So how do we design and run an ETS against a particular IUT? The steps fall into three categories: test suite development, execution, and results.

### 6.3.1 Test suite development

The basic steps in the development of a test suite are to read the specification, derive the Test Purposes, write Abstract Test Cases, complete the Abstract Test Suite, and implement an Executable Test Suite.

**Read the specification.** This step is not always as easy as it sounds. In many cases specifications are long, written in an unfamiliar language, and include complex formulations. Specifications also might be ambiguous, incomplete, and inconsistent. Any of these factors can result in two similar products, developed from the same specification by two different companies, being unable to interoperate.

**Derive test purposes.** The objective of the Test Purpose is to verify that a feature or characteristic described by the specification has been implemented correctly in the product. A series of Test Purposes is derived from a Protocol Specification to cover all features of potential IUTs. For example, an ATM Conformance Test Purpose might be to verify that the IUT supports point-to-point virtual channel (VC) connectivity.

**Write Abstract Test Case.** Each Test Purpose is turned into an Abstract Test Case (ATC) that clearly states the expected behavior of the IUT in order to meet that Test Purpose. An ATC gives details about:

- Protocol Data Unit (PDU) sent to the IUT.
- PDU expected from the IUT.
- Time frame in which the IUT must respond.
- What the IUT must do to get a Pass verdict.
- How the IUT can get a verdict of Fail or Inconclusive.

A *Protocol Data Unit* (PDU) is a complete unit of information sent or received by the IUT. For example, a PDU can be a “cell” at the ATM Cell Layer, a “frame” in frame relay, or a “message” in narrow- or broadband ISDN signaling protocol.

Figure 6.6 shows an example of a Conformance ATC scenario designed to test an ATM Cell Layer function. Conformance testing is a stimulus-and-response methodology. Cells in this example are sent to the IUT; in reply, the IUT sends them back. Upon receipt, these cells are compared to the expected cells. These events are driven by a protocol tester. The scenario is called a Test Case. Figure 6.7 shows an ATM Cell Layer interoperability test scenario; the concepts are the same as those of conformance testing.

**Complete the ATS.** Once they have been written, the Abstract Test Cases are grouped together and definitions are added. Examples of such definitions include

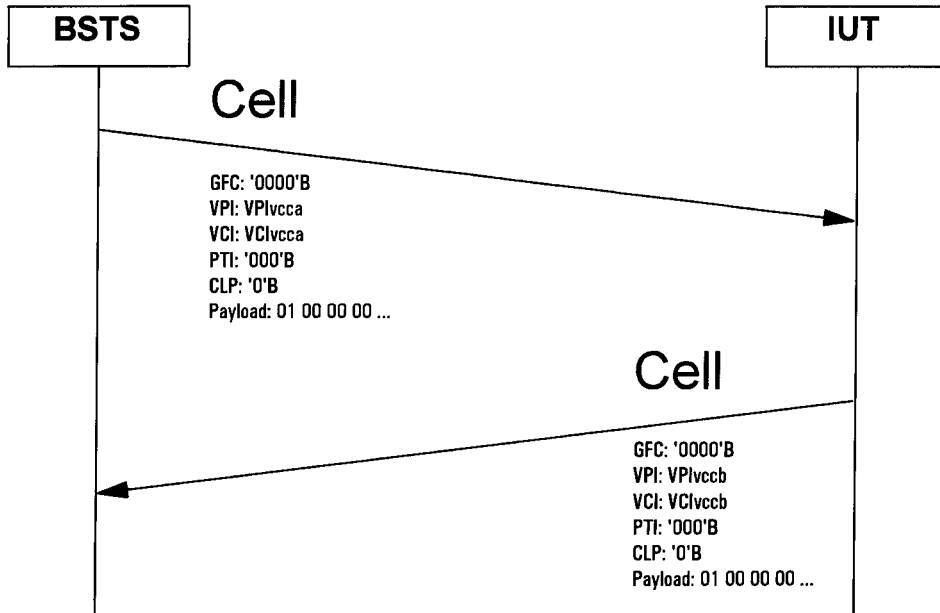


Figure 6.6 Conformance Abstract Test Case (ATC) and an example of an ATM Cell layer test scenario.




Figure 6.7 Interoperability Abstract Test Case (ATC) and an example of an ATM Cell layer test scenario.

PDU structure, ports, timer duration, etc. The result is an Abstract Test Suite (ATS). An ATS can be compared to sheet music. The notes have been determined and written down. A further step is needed, however, before they can be played.

**Implement an ETS.** Once implemented on a particular piece of hardware, each Abstract Test Case becomes an Executable Test Case (ETC). A series of Executable Test Cases is called an Executable Test Suite (ETS). An Abstract Test Suite is a design. An Executable Test Suite is an instance of the design that can run inside a tester.

Once an ETS is loaded in a tester, the IUT can be connected directly and some ETCs run. As soon as the results of these test cases have been received, the IUT implementation can be modified and the same test cases (or a subset of them) run again for comparison.

### 6.3.2 Test execution

Test execution includes three basic steps: parameterization, selection, and running the test.

**Test suite parameterization.** An ATS is designed to adapt its behavior to all sorts of IUTs. If the specification says that a given feature can be implemented optionally, for example, and a number of test cases have been included in the ATS to test such a feature, those test cases should not be run against an IUT that does not have this optional feature. For another example, if the IUT can react in two different valid ways in a given situation, the test suite must know before the execution begins which way has been chosen for this particular IUT so that it can adapt its sequence of events appropriately.

This customization is accomplished through a *Protocol Implementation Conformance Statement* (PICS) and a *Protocol Implementation Extra Information for Testing* (PIXIT) document. Each document contains a series of questions for parameterizing the ATS to a particular IUT. A PICS document is oriented to a protocol specification. It indicates which capabilities and options have been implemented in the IUT. A PIXIT document, however, is oriented to IUT dynamic behavior. It contains additional parameters used to tailor the execution of certain test cases (such as timer duration, the value of specific octets in some cells, etc.).

**Note:** In the case of ATM, both the ATM Cell Layer Conformance PICS and PIXIT proformas (prescriptions and descriptions) are created by ATM Forum Testing SWG (Sub-Working Group), along with (and sometimes before) the related ATS.

In preparation for a test campaign, a test operator normally will go through the following steps:

1. Get the PICS and PIXIT proformas and fill in the questions on paper. This first step is required by testing laboratories before conformance testing can commence. If the ETS is to be run privately in an R&D lab, however, the test operator can go directly to step 2.
2. Load the ETS on the protocol analyzer and fill in online the same PICS and PIXIT questions. The answers to these questions become the parameters (called Test Suite Parameters) of the ETS.
3. Customize the test campaign by answering additional questions as required (such as the number of times each test case should be run, delay between test cases, how detailed the test case traces should be, etc.).

Typically, an ETS package allows the operator to save all PICS, PIXIT, and additional settings. This feature saves time and adds consistency to further test campaigns. It also facilitates regression testing at a later stage.

**Test case selection.** The last step before testing can begin is test case selection. The operator must tell the ETS which cases should be run during the current test campaign. Some test cases might not be selectable, such as those where the IUT has not implemented an optional feature. Among those test cases that are selectable, the operator might opt to run only those that concentrate on particular aspects of the IUT. Figure 6.8 shows a Test Suite Browser screen for the selection of test cases.

This notion of selectability is clearly defined in conformance testing. For each test case, there is a boolean expression (called a *test case selection expression*) that must produce the value True before the test case can be selected. Each selection expression is based on one or more test suite parameters.

Consider the test case selection example shown in Figure 6.9. As you can see, test case *Ver\_VC\_Only\_F4* can be selected only if the IUT supports VC service but not VP service. In addition, most ETS packages allow the operator to disregard the test case selection expressions and to select any test case desired (such as for testing abnormal situations).

**Running the test.** As soon as one or more test cases have been selected, the test suite can be run. During execution, the tester sends PDUs to the IUT, analyzes its reaction (the contents of the cells sent by the IUT, time when these PDUs are sent, etc.), compares the expected and the observed behavior of the IUT, assigns a verdict to the test case, displays and records on disk all protocol data exchanged, and proceeds to the next test case. This type of testing is called *stimulus-and-response testing* because the tester expects a response from the IUT each time a stimulus is sent.

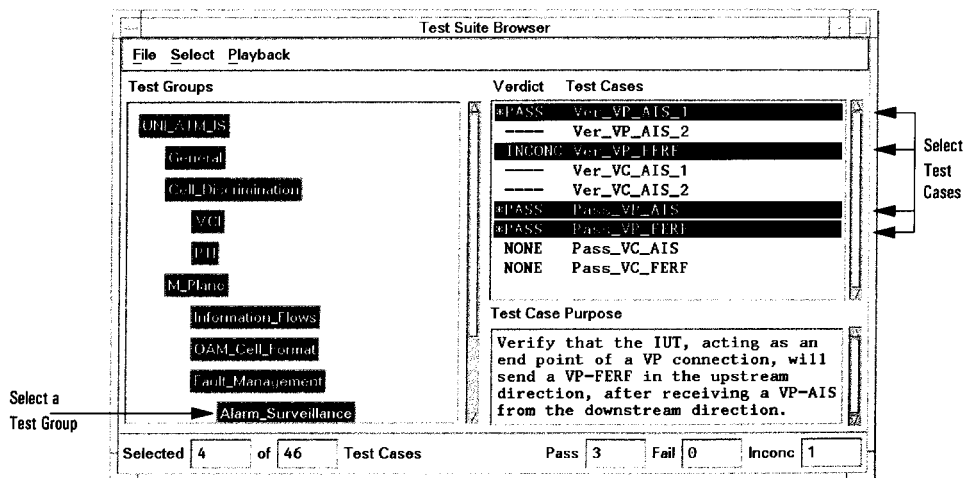


Figure 6.8 Test Case selection.

## 120 Network Test and Measurement

Test Case: Ver\_VC\_Only\_F4  
 Selection Expression: VC\_SERV\_ONLY  
 Def. of VC\_SERV\_ONLY: VC\_SERVpar AND NOT VP\_SERVpar  
 Parameter VC\_SERVpar: Answer to the PICS question:  
 Does the IUT supports VC service?  
 Parameter VP\_SERVpar: Answer to the PICS question:  
 Does the IUT supports VP service?

Figure 6.9 Test Case example.

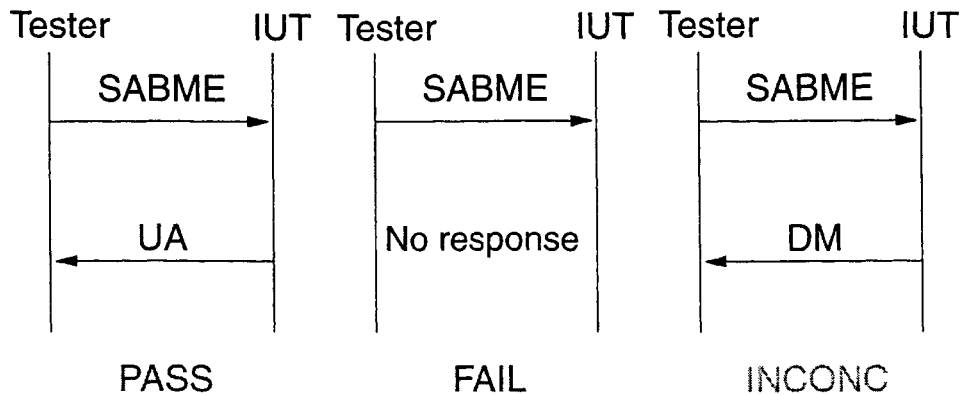


Figure 6.10 Test Case verdicts: some examples from ISDN LAPD.

In most cases the test suite runs without operator intervention. The only exception is when the test suite requires the operator to trigger a particular IUT event manually.

### 6.3.3 Verdicts and results

A conformance test campaign produces a verdict for each test case run. Figure 6.10 shows some sample verdicts from an ISDN LAP-D test. The three possible verdicts are Pass, Fail, and Inconclusive:

- *Pass* is given when the IUT has met the Test Purpose. This verdict indicates that the IUT has behaved exactly as specified in the Abstract Test Case (that is, the right PDUs have been sent with the right contents at the right moment).
- *Fail* is given when the IUT has not met the Test Purpose. This verdict means that an event other than that stated in the Abstract Test Case has occurred at least



once (for example, the wrong PDU was sent, incorrect contents were found, or the right PDU was sent too late).

- *Inconclusive* is given when something has gone wrong but the tester is unable to verify whether the IUT met the Test Purpose or not. For example, operator intervention was required to trigger an IUT event but the operator failed to proceed, or incorrect behavior was encountered in the early stages of the test case, called the *preamble*, before the Test Purpose could be verified.

**Test Case Trace.** Results of a test campaign also include a trace of each test case run. A Test Case Trace contains:

- The test case identifier, plus the date and time when execution began.
- All PDUs sent to and from the IUT, in their original transmission order, time-stamped, and with all fields decoded.
- Optionally, statements (as coded in the Abstract Test Case) to help the operator follow the course of events and detect where a problem has occurred and what the expected result should have been (according to the ATC).
- The verdict assigned to the test case.
- The date and time when the test case execution ended.

Figures 6.11a and 6.11b give an example of a Test Case Trace.

**Test Report.** At this stage a detailed Test Report is generated, which includes all results produced by the tester during a test campaign. A detailed Test Report provides:

- A summary table that indicates:
  - The date and time when a test report was produced.
  - The test suite name and version number.

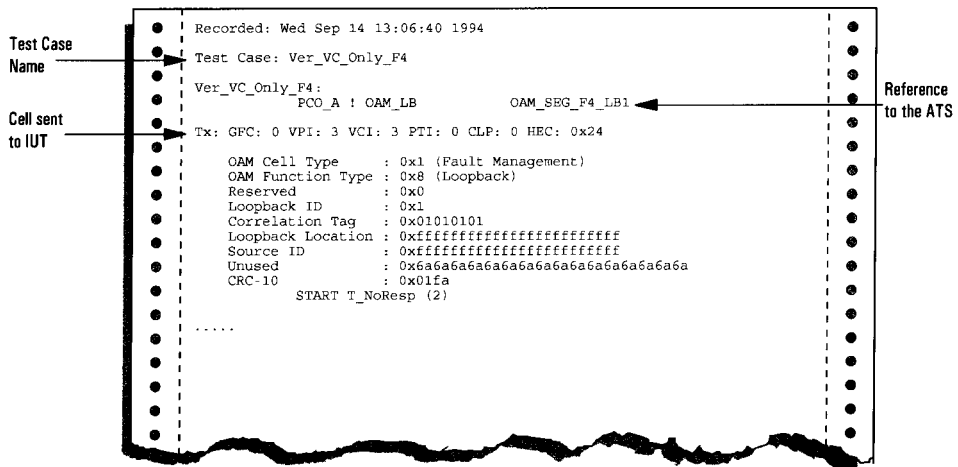


Figure 6.11a Beginning section of a Test Case Trace.

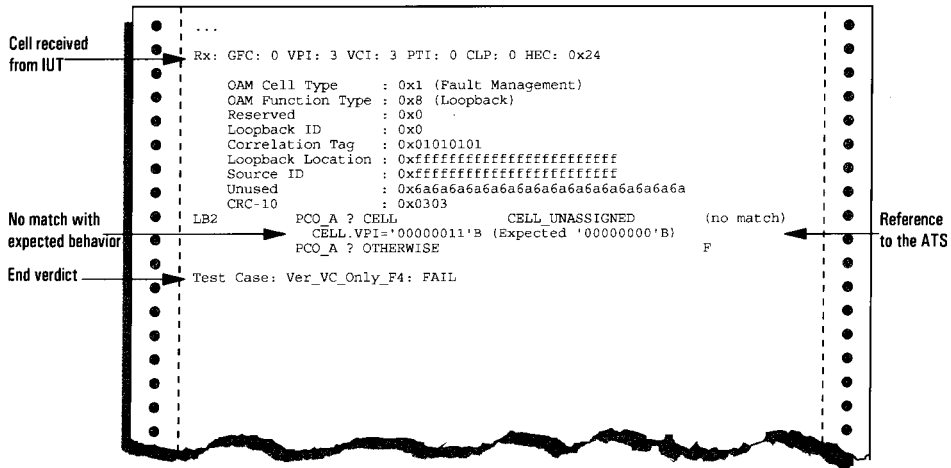


Figure 6.11b Ending section of a Test Case Trace.

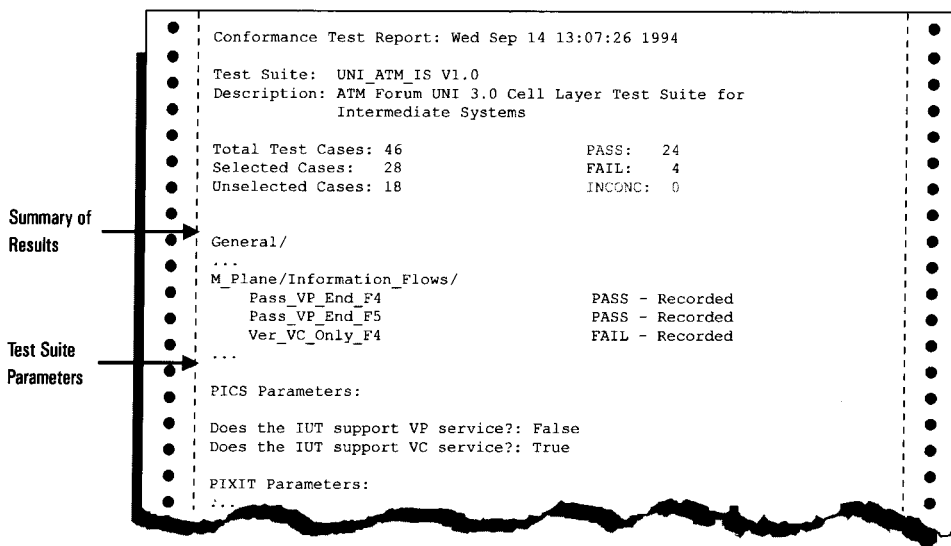


Figure 6.12 An excerpt from a conformance Test Report.

- The environment, including type of IUT, company name, and location (optional).
- The number of test cases selected.
- The total number of Pass, Fail, and Inconclusive verdicts.
- The name and its verdict for each test case.
- A list of all PICS and PIXIT questions, with the answers given by the test operator.
- A complete trace for each test case run.

Figure 6.12 shows an excerpt from a detailed test report.

## 6.4 Tree and Tabular Combined Notation

*Tree and Tabular Combined Notation* (TTCN) is a precisely defined notation for specifying test scenarios (ATSS). Defined by the ISO as Part 3 of ISO 9646 (and accepted by the ITU-T), it is independent of test methods, protocols, layers, and test tools. Different versions of the notation include TTCN DIS (Draft International Standard), TTCN IS (International Standard), and Concurrent TTCN (Amendment 1). Other new documents were under development at the time of this writing.

As an international standard, TTCN is recognized throughout the world by standards bodies and testing committees. It is used widely in lower-layer protocols such as the ATM Cell Layer and signaling (ATM Forum Conformance ATS), the frame relay UNI and NNI (ACT-FR), the ISDN LAP-D, the ISDN layer 3 (NI-1 BCC, NI-1 SS, TBR3, TBR4), X.25 DTE layers 2 and 3, and MTP (SS7). In addition, it is used in upper-layer protocols such as FTAM, MHS, SCCP (SS7), TCAP (SS7), and ISO Session.

What does TTCN do? TTCN provides a formal notation that describes test scenarios in a complete and unambiguous fashion. It allows comprehensive coverage of:

- The sequence of all possible events during a test case.
- The contents of PDUs sent to the IUT.
- The contents of PDUs expected from the IUT.
- The time frame allowed for the IUT to respond.
- The events when verdicts (Pass, Fail, Inconclusive) are assigned.

Other notations are informal and incomplete by comparison. An ATC given in another notation might state informally what the IUT must do in order to get a Pass verdict, but might not indicate how it can Fail.

### 6.4.1 TTCN features

TTCN offers two forms, *Graphical Representation* (TTCN-GR) and *Machine Processable* (TTCN-MP). TTCN-GR is designed for editing, printing, and publishing ATSS. Figure 6.13a shows a sample TTCN-GR screen.

TTCN-MP provides testing functionality equivalent to that of TTCN-GR but is used to exchange ATSS between developers. It uses ASCII representation and follows a strict syntax. For this reason, it can be used as input to software tools such as a TTCN Translator. Figure 6.13b shows an example of TTCN-MP.

A TTCN ATS is divided into four major sections:

- *Test Suite Overview* provides a table of contents and index to all test cases and test steps.
- *Declarations* provides all definitions, PDU structure, timer duration, etc.
- *Constraints* provides the exact contents of PDUs sent and expected.
- *Dynamic Behavior* provides actual test scenarios.

- 1st form: Graphical Representation (TTCN-GR).
  - Used for editing and printing/publishing ATSS.

Test Case Dynamic Behaviour				
<b>Test Case Name</b> : Ver_VC				
<b>Group</b> : UNI_ATM_IS/General/				
<b>Purpose</b> : Verify that the IUT supports point-to-point VC connectivity.				
<b>Default</b> :				
<b>Comments</b> : Requires a VC connection Ref. 3.1				
			<b>t</b>	<b>s</b>
LB1	PCO_A!CELL_NR	CELL_SQ(VPIvcca, VCIvcca, '01'0)		
	START T_Test			
	PCO_B?CELL_NR	CELL_SQ(VPIvccb, VCIvccb, '01'0)	P	
	PCO_B?CELL	CELL_UNASSIGNED		
	GOTO LB1			
	?TIMEOUT T_Test		F	
	PCO_B?OTHERWISE		F	

Figure 6.13a Tree and Tabular Combined Notation (TTCN) has two possible forms. Shown here is the Graphical Representation (TTCN-GR), used for editing and printing/publishing Abstract Test Suites.

- Example of TTCN-MP:

```

$Begin_TestCase
$TestCaseId Ver_VC
$TestGroupRef UNI_ATM_IS/General/
$TestPurpose
/* Verify that the IUT supports point-to-point VC
connectivity. */
$DefaultsRef
$Comment
/* Requires a VC connection Ref. 3.1 */
$BehaviourDescription
$BehaviourLine
$LabelId
$Line [0]PCO_A!CELL_NR
$CRef CELL_SQ(VPIvcca,VCIvcca,'01'0)
$VerdictId
$Comment /* */
$End_BehaviourLine
...

```

Figure 6.13b The second form of TTCN is Machine Processable (TTCN-MP), used to exchange Abstract Test Suites among developers.

#### 6.4.2 How to read an ATS in TTCN

Figure 6.14a shows an example Test Case Dynamic Behavior screen using TTCN. Most test cases (as in this example) start with the tester sending a PDU to the IUT. When the test scenario requires that the IUT send a cell first, however, the syntax `<IUT!cell_type>` is used. This statement asks the test operator to trigger this event from the IUT. The term otherwise stands for any type of PDU with any contents.

Figure 6.14b shows an example of a PDU Type Definition screen using TTCN. The notation `CELL_NR` refers to cell type. Figure 6.14c shows an example of a PDU Constraint Declaration using TTCN. The notation `CELL_SQ` refers to cell contents. The `CE; ; _NR` and `CELL_SQ` are referenced in the first line of Figure 6.14a.

**Basic semantics.** The basic semantics of an ATS expressed in TTCN follow these general rules:

- **TTCN Statements in Sequence:** TTCN statements in sequence are indented once from each other. When a statement is successful, control passes to the next statement in sequence (Figure 6.14d).
- **Alternative TTCN Statements:** Statements at the same indentation level are possible alternative events. Control loops from one alternative to the other until a

• Sending/receiving cells to/from the IUT.

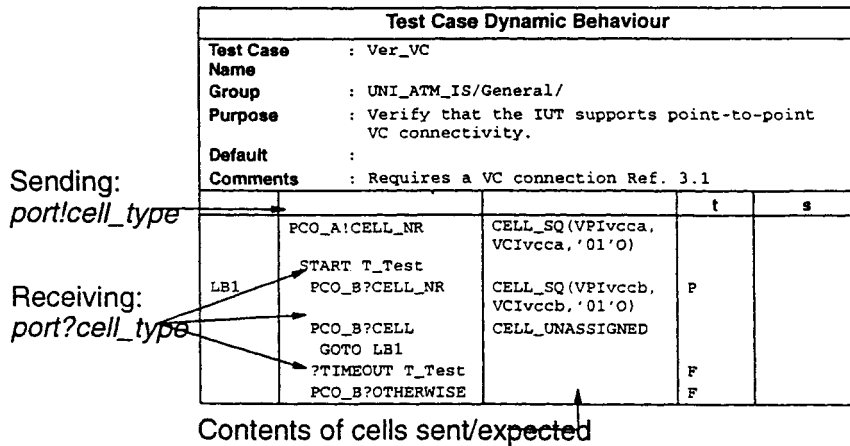


Figure 6.14a Dynamic behavior example, using TTCN.

• Cell Type: **CELL\_NR**

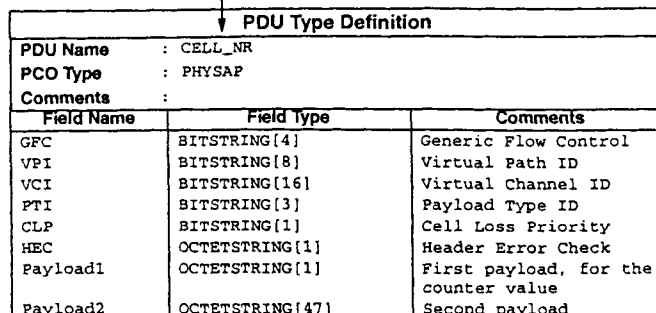


Figure 6.14b PDU type definition screen using TTCN.

- Cell Contents: CELL\_SQ

PDU Constraint Declaration		
<b>Constraint Name</b>	: Cell_SQ(VPI_val,VCI_val:BITSTRING;N:OCTETSTRING)	
<b>PDU Type</b>	: CELL_NR	
<b>Derivation Path</b>	:	
<b>Comments</b>	:	
VPI	VPI_val	
VCI	VCI_val	
PTI	'000'B	
CLP	'0'B	
HEC	Valid_HEC	
Payload1	N	
Payload2	PadOctet('00'0,47)	

Figure 6.14c PDU constraint declaration using TTCN.

- Basic semantics: TTCN Statements in Sequence.
  - TTCN statements in Sequence are indented once from each other.
  - When a statement is successful, control goes to the next statement in sequence.




Figure 6.14d TTCN statements in sequence.

statement is successful. At that point, control moves to the next statement in sequence following the successful event (Figure 6.14e).

- **Verdicts:** From the example in Figure 6.14b, if a proper cell with proper contents is received, the IUT gets a Pass verdict. If the IUT does not respond within T\_Test seconds or sends an incorrect cell, it gets a Fail (Figure 6.14f). Verdict types include:
  - Preliminary Verdicts:* (P), (F), and (I)
  - Final Verdicts:* P, F, and I, plus R (for “keep highest preliminary verdict assigned”), where verdicts in their order of precedence (from lowest to highest) are Pass (P), Inconclusive (I), and Fail (F).
- **End of Execution:** Execution stops when there are no more statements in sequence following a successful event, or when a final verdict is met. A verdict must be assigned before execution stops.

**Timers.** Figure 6.14g shows timers set in an ATS. The `START T_Test` is the `T_Test` start-timer, and `?TIMEOUT T_Test` is an “if” timer showing where the `T_Test` expires. Timers also can be canceled, and read using the `READTIMER` function.

**Labels and GOTO statements.** Figure 6.14h shows a label and `GOTO` statement in an ATC. They provide a simple mechanism to force a loop.

**Assignments and boolean expressions.** Figure 6.14i shows examples of TTCN assignments and boolean expressions. Assignments are enclosed within parentheses and include the symbol `:=` used in languages such as Pascal and Ada. Boolean expressions are enclosed in brackets. Two alternative expressions provide a simple mechanism to implement an “if-then-else” statement, as in this example:

```
[VCI_NR>=Max_VCI]
```

- **Basic semantics: Alternative TTCN Statements.**
  - Statements at the same indentation level are possible alternative events.
  - Control loops from one alternative to the other until one of them is successful.
  - Then control goes to next statement in sequence following the successful event.




Figure 6.14e Alternative TTCN statements.

• **Verdicts**

Test Case Dynamic Behaviour				
Test Case Name : Ver_VC				
Group : UNI_ATM_IS/General/				
Purpose : Verify that the IUT supports point-to-point VC connectivity.				
Default :				
Comments : Requires a VC connection Ref. 3.1				
			t	s
LB1	PCO_A!CELL_NR	CELL_SQ(VPIvcca, VCIvcca, '01'0)		
	START T_Test			
	PCO_B?CELL_NR	CELL_SQ(VPIvccb, VCIvccb, '01'0)	P	
	PCO_B?CELL	CELL_UNASSIGNED		
	GOTO LB1			
	?TIMEOUT T_Test		F	
	PCO_B?OTHERWISE		F	

If proper cell with proper con is received, IUT get a PASS

If IUT does not respond within T\_Test seconds or sends an incorrect cell, IUT get a FAIL

Figure 6.14f Verdicts.

• Timers

Test Case Dynamic Behaviour				
<b>Test Case Name</b> : Ver_VC				
<b>Group</b> : UNI_ATM_IS/General/				
<b>Purpose</b> : Verify that the IUT supports point-to-point VC connectivity.				
<b>Default</b> :				
<b>Comments</b> : Requires a VC connection Ref. 3.1				
			<b>t</b>	<b>s</b>
	PCO_A!CELL_NR	CELL_SQ(VPIvcca, VCIvcca, '01'0)		
Start timer T_Test	→			
	LB1	START T_Test PCO_B?CELL_NR	CELL_SQ(VPIvccb, VCIvccb, '01'0)	P
		PCO_B?CELL	CELL_UNASSIGNED	
If timer T_Test expires	→	GOTO LB1		
		?TIMEOUT T_Test		F
		PCO_B?OTHERWISE		F

Timers can also be CANCELED

Figure 6.14g End of execution.

• Labels and GOTO statements.

Test Case Dynamic Behaviour				
<b>Test Case Name</b> : Ver_VC				
<b>Group</b> : UNI_ATM_IS/General/				
<b>Purpose</b> : Verify that the IUT supports point-to-point VC connectivity.				
<b>Default</b> :				
<b>Comments</b> : Requires a VC connection Ref. 3.1				
			<b>t</b>	<b>s</b>
		PCO_A!CELL_NR	CELL_SQ(VPIvcca, VCIvcca, '01'0)	
Label	→	LB1	START T_Test PCO_B?CELL_NR	CELL_SQ(VPIvccb, VCIvccb, '01'0)
			CELL_UNASSIGNED	P
Used in GOTO	→	PCO_B?CELL		
		GOTO LB1		
		?TIMEOUT T_Test		F
		PCO_B?OTHERWISE		F

Figure 6.14h Timers.

**Tree attachments (function calls) to test steps.** Figure 6.14j shows examples of tree attachments, which start with a + symbol. The test steps (whose names follow the + sign) are defined in separate tables, using the same Dynamic Behavior syntax and semantics.

**Link between test case traces and ATS in TTCN.** Figure 6.15a shows the beginning of a Test Case Traces printout. This section identifies the test case using the Test Case Name and indicates which cells were sent to the IUT.

Figure 6.15b shows the end of the same Test Case. It identifies the cell received from the IUT, gives the complete TTCN statement (including label, event, and constraint), and explains why the event was unsuccessful.



### 6.4.3 Benefits of TTCN

TTCN allows an operator or test engineer to know exactly what to expect from test scenarios. There are no surprises or ambiguities. In certain cases it can help clarify otherwise unclear sections of the specification. Learning a single notation allows test operators and customers to read ATSs for all major protocols around the world.

Thanks to TTCN-MP's strict syntax and semantics, it greatly automates the translation of ATSs in TTCN into executable code (using C source code, for example.) This facility allows the production of ETSs that reflect precisely the desired ATS (resulting in a better product sooner), and the embedding of diagnostic trace statements to help pinpoint problems in an IUT quickly and accurately.

- Assignments and Boolean Expressions.

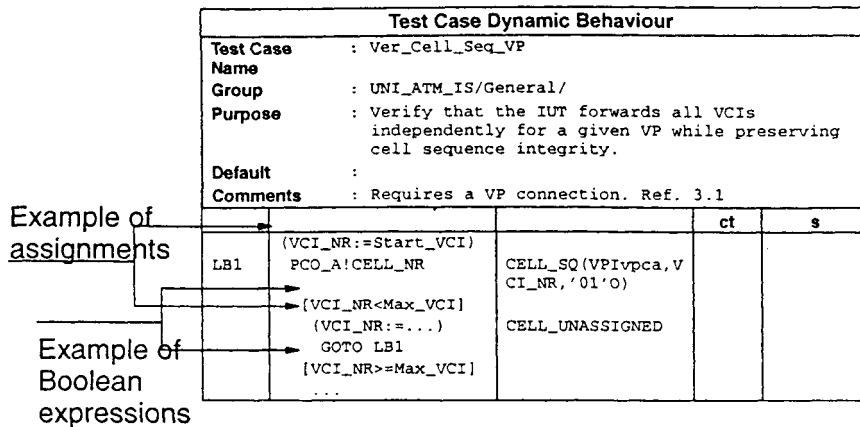


Figure 6.14i Labels and GOTO statements.

- Tree attachments (i.e. function calls) to Test Steps.

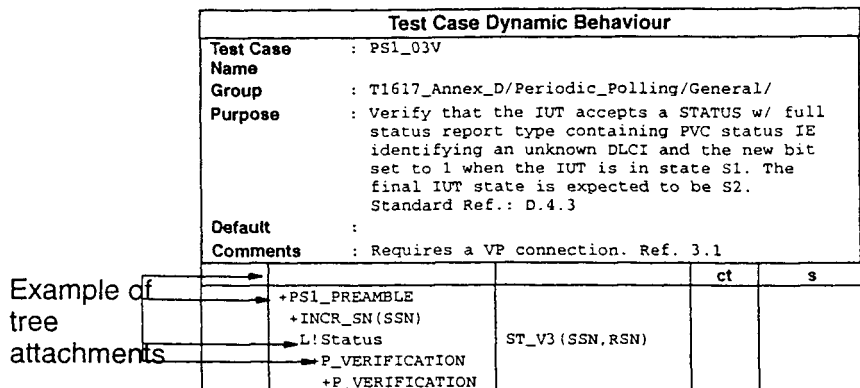


Figure 6.14j Assignments and boolean expressions.

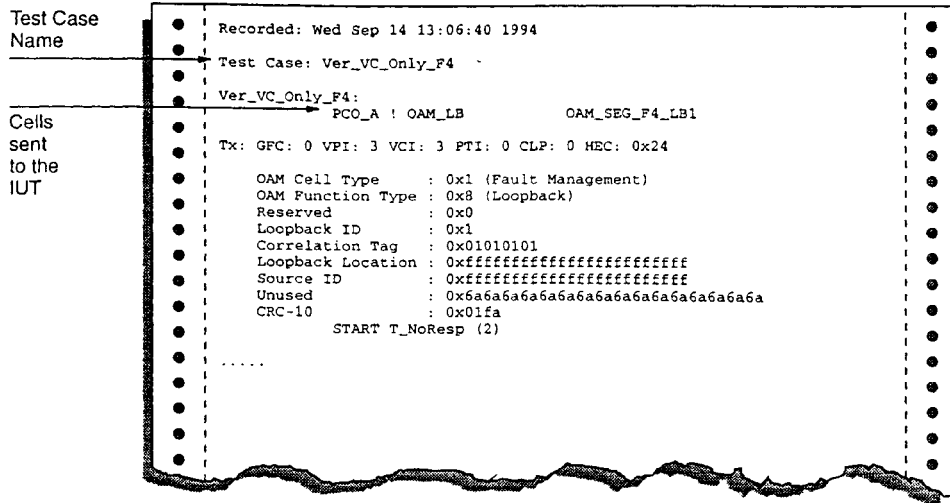


Figure 6.15a Section identifying the test case.




Figure 6.15b Test case traces printout.

---

Part

**3**

# Wide Area Networks



# PDH Networks

## Principles of Digital Transmission

**Doug Conner**

*Hewlett-Packard Ltd., Westfield, Massachusetts*

**Hugh Walker**

*Hewlett-Packard Ltd., South Queensferry, Scotland*

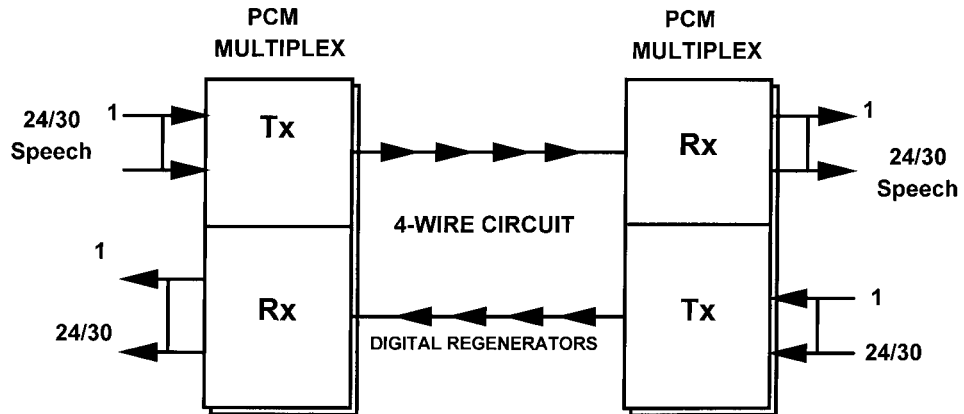
### 7.1 Introduction to Plesiochronous Digital Networks

The term *Plesiochronous Digital Hierarchy* or *PDH* refers to a multiplexing system that is not fully synchronous. Plesiochronous, according to the ITU-T recommendations, means nominally at the same bit rate but not synchronized to a common master clock. The variation from “nominal bit rate” allowed in a plesiochronous telecom system is typically between 15 and 50 parts per million (ppm) offset from the specified clock frequency.

PDH multiplexing and transmission systems comprised the first generation of digital telecommunications network technology, developed in the 1960s and 1970s. PDH has now been superseded by the synchronous SDH and SONET hierarchy developed in the late 1980s. A great deal of PDH equipment exists in the world’s telecommunications networks, however, and the new synchronous system is also designed to interwork with it. Testing PDH networks thus will continue to be an important issue for many years to come.

The digital telecommunications network had its origins with the development of pulse code modulation (PCM), invented by Reeves in 1937 and patented in 1939. As described in Chapter 3, PCM involves sampling, quantizing, and coding the analog telephone voice signal to produce a compressed binary digital signal. When Reeves invented PCM, the traffic on the telecommunications network was almost entirely voice telephony, apart from a very small amount of Telex and telegraph. The practical application of PCM had to wait, however, until the development of solid-state technology in the 1950s and 1960s.

In 1962, the Bell System in the U.S. installed the first point-to-point multiplexed digital transmission system, shown schematically in Figure 7.1. The main purpose of this



**Figure 7.1** An early point-to-point PCM system operating at primary rate over twisted pairs previously used for voice band telephony. The digital regenerators are required every 2 km (1.25 mi). These early systems were deployed mainly to increase the traffic capacity of existing trunks between exchanges by taking advantage of the higher noise immunity of digital systems.

early PCM system was to increase the capacity of trunks between main exchanges or Central Offices. It operated at the T1 rate of 1.544 Mbps, carrying 24 telephone channels over a 4-wire circuit that previously handled just one analog voice channel. Digital regenerators were necessary every 2 km (approximately 1.25 mi) to overcome the losses in the twisted-pair cable. Conveniently enough, this was the approximate spacing of the loading coils previously used to condition the lines for voice-frequency traffic.

Early PCM systems in Europe also operated with 24 multiplexed channels, but the standard soon became the 30-channel system at the 2.048 Mbps E1 primary rate. The next step was to take several of these primary-rate T1 or E1 multiplexed signals and combine them into a single, high-capacity transmission path, which in the 1970s would be either microwave radio or (more likely) coaxial cable. In 1972, the ITU's International Telegraph and Telephone Consultative Committee (CCITT, now the ITU-T) issued the first version of the Recommendation G.703, "Physical/Electrical Characteristics of Hierarchical Digital Interfaces." This document defines the interconnection requirements for PDH equipment. The equivalent North American standard is ANSI T1.102, "Digital Hierarchy—Electrical Interfaces."

Two main PDH standards are in use; one is based on the 1.544 Mbps primary rate for North America, and the other based on the 2.048 Mbps primary rate found in most other countries of the world. In addition, Japan has a different hierarchy for the higher levels, also based on 1.544 Mbps, as shown in Table 7.1 (taken from ITU-T Recommendation G.702 and G.703 and ANSI T1.102).

The fundamental idea about plesiochronous multiplexing is that each multiplexer and demultiplexer is a standalone island within the network. It has its own internal clock source, which needs to have moderate stability to meet the limits specified in Table 7.1, but there is no need to synchronize these internal clock sources to a master clock. Most networks are synchronous at the primary T1/E1 rate, however, because the 24/30-channel assembly is a fully synchronous structure to allow the digital circuit switches to operate directly on the timeslots, as described in Chapter 3. The

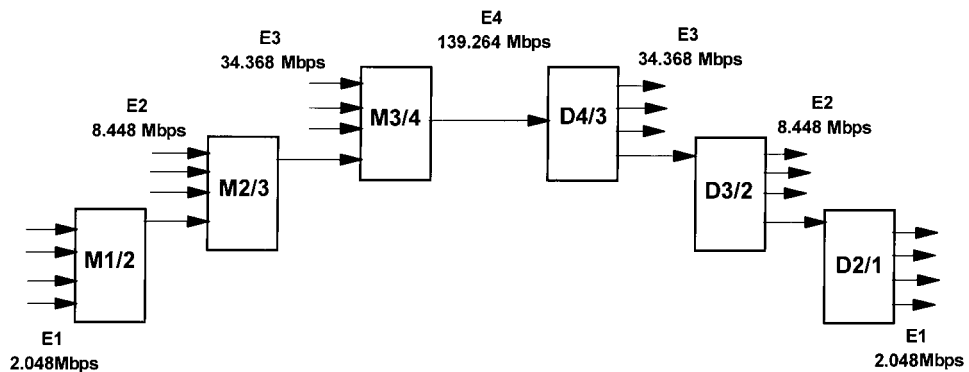
digital switches in the network therefore are synchronized to an atomic standard reference clock, so that traffic can be exchanged across the network and between different networks with the minimum number of slips in the data stream, as specified in ITU-T recommendation G.822. Above primary rate, the PDH multiplex and transmission hierarchy is asynchronous.

Consider the PDH multiplexing and transmission hierarchy shown in Figure 7.2. Each multiplexer and demultiplexer block is autonomous, with its own internal ref-

**TABLE 7.1 PDH Hierarchy Rates.**

Digital hierarchy level	Hierarchical bit rates (Mbps) for networks with the digital hierarchy based on the primary rate of:		
	1.544 Mbps North America	1.544 Mbps Japan	2.048 Mbps International
1	1.544 Mbps (T1/DS1) +/- 50 ppm (G.703) +/- 32 ppm (T1.102) 3.152 Mbps (DS1C) +/- 30 ppm	1.544 Mbps (J1) +/- 50 ppm	2.048 Mbps (E1) +/- 50 ppm
2	6.312 Mbps (DS2) +/- 30 ppm (G.703) +/- 33 ppm (T1.102)	6.312 Mbps (J2) +/- 30 ppm	8.448 Mbps (E2) +/- 30 ppm
3	44.736 Mbps (T3/DS3) +/- 20 ppm	32.064 Mbps (J3) +/- 10 ppm	34.368 Mbps (E3) +/- 20 ppm
4	139.264 Mbps (DS4NA) +/- 15 ppm 274.176 Mbps (DS4) +/- 10 ppm (Note 1)	97.728 Mbps (J4) +/- 10 ppm	139.264 Mbps (E4) +/- 15 ppm

Note 1. The fourth level of North American hierarchy at 274 Mbps is rarely used and is not included in the standards.



**Figure 7.2** The PDH multiplexer/demultiplexer pyramid for the international 2 Mbps standard, showing the nomenclature used to describe the stages of multiplexing and demultiplexing. Each multiplexer in the PDH system is a standalone element with its own internal clock, and there is no requirement for synchronization. Extracting a 2 Mbps (E1) tributary from a high-capacity 140 Mbps (E4) path, however, requires three stages of demultiplexing as shown here.

erence clock. The multiplexer receives input tributary streams (which in turn might have come from different sources with slightly different clock rates) and recovers a separate clock for each incoming stream. It synchronizes all of these to its own internal clock using *positive justification* or *bit stuffing* (see Chapter 3, section 3.4.1). Then the synchronized streams can be interleaved and the higher-order stream transmitted at the rate of the multiplexer's internal clock.

The PDH demultiplexer reverses this process by de-interleaving the stream, having locked on to the frame alignment signal. It recovers a clock from the incoming data signal, effectively synchronizing itself to the internal clock of the transmitting multiplexer. It then examines the bit stream and extracts all the redundant stuffed bits used for justification at the transmit end and generates a *gapped data signal*. By means of phase-locked loops, these recovered tributaries are reclocked at a steady rate, so that if the destuffing has been done correctly, the regenerated tributaries will reappear at exactly the same rate as they entered the PDH transmission system. In other words, the system is transparent to clock frequency variations of individual multiplexers through which the streams pass.

The PDH system is thus very robust, and has served the world's telecom network well. Because of the need to operate asynchronously, however, additional hardware is required for shift registers, buffer, stores and phase-locked loops at each mux/demux device. Recovering a low-order tributary from a high-order stream is only possible by executing stage-by-stage demultiplexing to eliminate the arbitrary stuffed bits at each level. As networks have needed more flexibility to meet commercial pressures, PDH gradually has been superseded by synchronous multiplexing (SDH/SONET), which allows easy drop and insert of traffic streams and bandwidth grooming. It also provides better in-service monitoring.

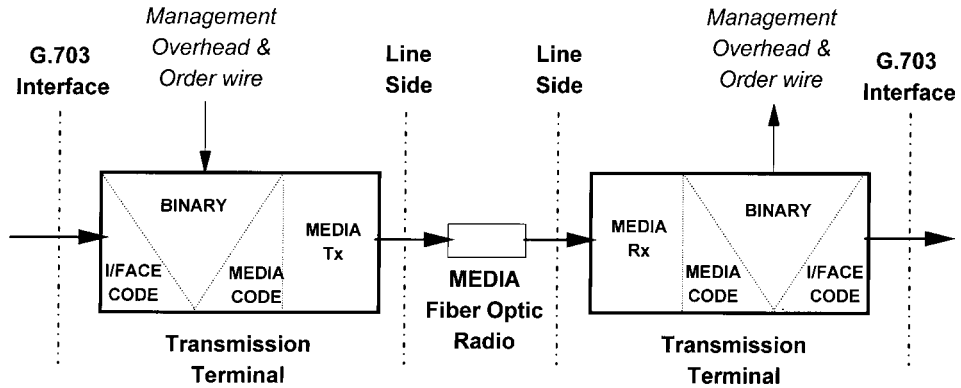
As mentioned earlier, PDH interface characteristics are defined by ITU-T Recommendation G.703. These interfaces will be found on multiplexers and demultiplexers and also on transmission terminals; the "line side" of PDH transmission systems (Figure 7.3) is proprietary, however, and will vary from one manufacturer to another. Signal levels, line coding, frequencies, wavelengths, and management overheads are not specified in the standards, in contrast to the SONET/SDH. Additional bits may be added in the transmitting terminal for error detection and forward error correction, network management, engineer's order wire, scrambling, framing, and line coding; thus the gross line bit rate will vary from one design to the next. For this reason it is usual to test network performance only at the PDH interfaces.

In the next section, the PDH frame structures used in the international and North American systems are described in more detail. The PDH frame structure is important for in-service performance monitoring.

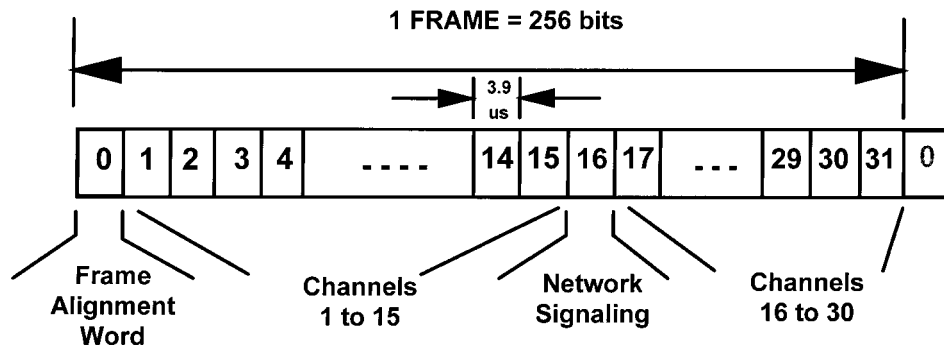
## 7.2 PDH Multiplexing Hierarchies, Frame Structures, and Standards

Because the international or European PDH hierarchy and the North American hierarchy evolved separately with completely different frame structures and bit rates, in this section they will be considered separately. This also means that the test requirements of the two systems are different and usually require different test equipment, as discussed later.





**Figure 7.3** PDH transmission systems are standardized only at the ITU-T G.703 hierarchical interfaces at the input and output of the transmission terminal. On the line side, every manufacturer may use a different media code and management overhead structure, so standardized testing can be done only at the G.703 interface. In the media transmitter, a clock is recovered from the incoming coded PDH hierarchical signal, and the G.703 interface code (e.g., HDB3 or CMI) is removed and binary data and clock streams derived. The terminal then adds its own overhead for network management, error control, framing, and order wire, before the composite signal is encoded for media transmission over microwave radio or fiber. The receiver reverses this process to regenerate the PDH signal.



**All Timeslots contain 8 bit words  
(PCM, Data, Subrate, n x 64 kb/s)**

**Network Signaling - CAS or CCS**

**Figure 7.4** The 32-timeslot frame at the 2.048 Mbps (E1) primary rate, which begins with the frame alignment word. The frame, which repeats every 125  $\mu$ s, is a fully synchronous frame because there is no provision for additional justification bits for synchronization.

**7.2.1 European 2.048 Mbps primary rate PDH hierarchy**

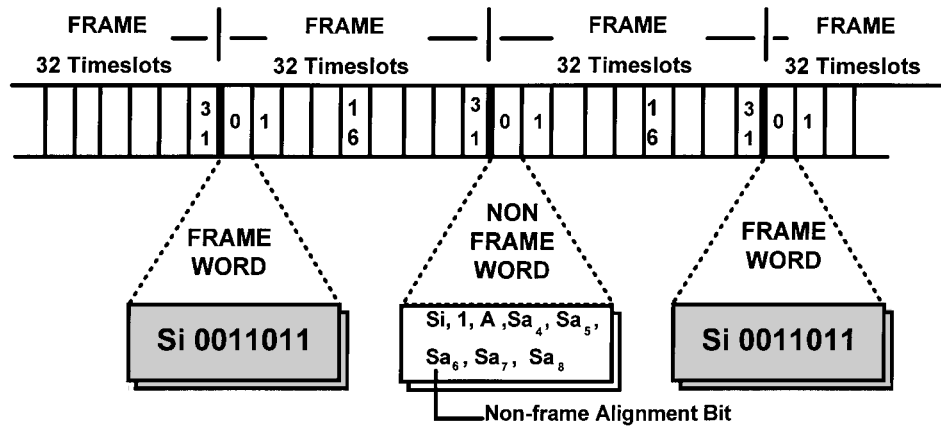
**Primary rate frame structure.** The fundamental building block for the European or international PDH hierarchy is the 2.048 Mbps E1 frame structure, comprising the 32-timeslot synchronous frame defined by ITU-T Recommendations G.704 and G.732 (shown in Figure 7.4).

Strictly speaking, no frame structure is required; a feature of the 2 Mbps hierarchy, defined in ITU-T Recommendation G.703, is that transmission can be *bit sequence independent*. In other words, 2 Mbps and 64 kbps facilities are “clear channel” and do not require any particular signal structure to pass through the network. Although this transparency can be useful for transmission of wideband signals, sending an unstructured signal into the network can have drawbacks.

An apparently random signal cannot be monitored in-service by the service provider for transmission errors, and it is impossible to provide bandwidth grooming or switching of channels. It is likely that the network operator will not be able to guarantee network performance with unstructured 2 Mbps traffic. In view of this, most private and public networks operate with the standard 2 Mbps frame structure defined in ITU-T Recommendation G.704.

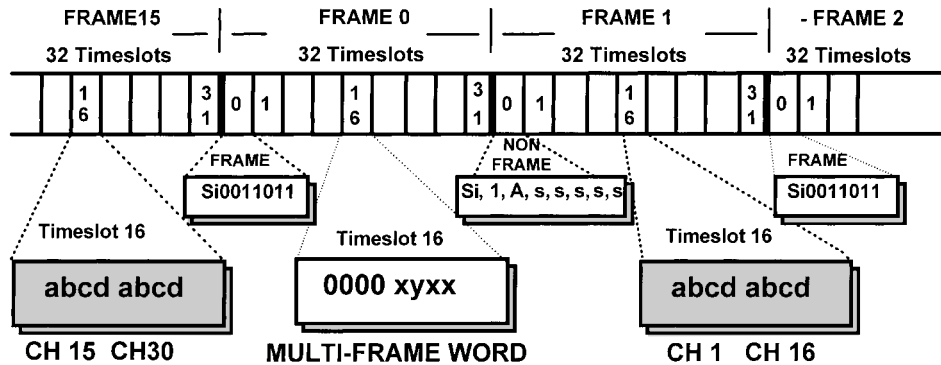
Each 2 Mbps frame contains 256 bits (32 timeslots, each of 8 bits) at a repetition rate of exactly 8 kbps. The first timeslot (timeslot zero, TS0) is reserved for framing, error-checking, and alarm signals; the remaining 31 can be used for traffic. The individual channels can be used for 64 kbps PCM, subdivided further for low-rate data or voice compression such as ADPCM (Adaptive Differential PCM), or aggregated for wideband signals such as videoconferencing or LAN interconnection. Sometimes a timeslot (such as TS16) is reserved for signaling (ISDN primary rate D channel signaling such as Q.931, for example, or channel-associated ABCD signaling).

The start of the 32-timeslot frame is signified by the frame alignment word (0011011) in the TS0 of alternate frames, as shown in Figure 7.5. In the other frame,



- Si** : Reserved for INTERNATIONAL    **A** : Remote Alarm
- Sa<sub>N</sub>** : Reserved for NATIONAL routes ("1" for INTERNATIONAL routes)
- LOSS of frame** : 3 bad frame words in 4
- Frame RECOVERY** : 2 good frame words + 1 Non-frame word

**Figure 7.5** The 7-bit frame word is transmitted in alternate frames as shown. The frame alignment criteria are defined in ITU-T G.704.



**x** : Reserved for NATIONAL routes ("1" for International)  
**y** : MULTI-FRAME Alarm  
**abcd** : CHANNEL ASSOCIATED SIGNALING digits

**Figure 7.6** Timeslot 16 (TS16) contains alternately the multiframe alignment signal or pairs of signaling bits for channel-associated signaling (CAS). The multiframe provides a reference so that the receiving equipment can decode the 4-bit signaling word for each of the 30 PCM channels. When common-channel signaling (CCS) is used, TS16 can be used for traffic or to carry the CCS signal.

bit 2 is set to 1 and bit 3 contains the A-bit for sending an alarm to the far end. The S-bits are all intended for international and national use and, when unused, are set to logical 1.

Once the demultiplexer has achieved frame alignment, it can separate the individual 64 kbps channels in the frame. If three out of four frame alignment words are received in error, the terminal declares loss of frame alignment and initiates a resynchronization process. The recovery criterion is one correct frame alignment word, one nonframe word bit 2 (logical 1), followed by one correct frame alignment word.

When the 2 Mbps frame was used exclusively for PCM voice transmission, the frame alignment criterion was very reliable. With data transmission, however, the traffic can simulate the frame alignment and nonframe alignment words, meaning false framing is possible. A new, more robust standard has been developed, building on the earlier framing standard; this will be discussed shortly.

Once the multiplexer has gained frame alignment, it searches in TS16 for the multiframe alignment signal (0000) in bits 1–4. This marks frame 0 of the group of 16 frames, called the multiframe (shown in Figure 7.6).

The multiframe is necessary only when channel-associated signaling (CAS) is used. Timeslot 16 then contains pairs of 4-bit ABCD signaling words. Over a complete multiframe, all 30 channels are serviced. If common-channel signaling (CCS) is used, then multiframe alignment is unnecessary; TS16 is used simply as a 64 kbps data channel for CCS messages, or it can be turned over to revenue-earning traffic (giving a total of 31 channels for the payload).

The 2 Mbps frame structure just described is in widespread use. It has some limitations, however, particularly with increased data transmission and demand for on-line performance monitoring.

As already mentioned, there is a risk of false framing, which would have serious effects on data. Performance monitoring of the received signal is limited to checking for errors in the frame alignment signals. With only a total of 7 bits out of 512, it gives a poor indication of errors in the payload. There is no way for the remote end to send back this rudimentary error-performance data, so only one direction of transmission can be monitored at each location.

The new CRC-4 (Cyclic Redundancy Checksum 4) frame structure is defined in ITU-T Recommendation G.704. The CRC-4 remainder is calculated on complete blocks of data, including all payload bits; the 4-bit remainder is transmitted to the far end for comparison with the recalculated CRC-4 remainder. If the two 4-bit words differ, then the receiving equipment knows that one or more errors are present in the payload block. Every bit is checked, so an accurate estimate of block error rate (or errored seconds) is made while the link is in service. The CRC-4 framing algorithm is more complex and is extremely unlikely to be fooled by payload data patterns.

Figure 7.7, taken from ITU-T G.704, shows the sequence of bits in the frame-alignment (TS0) position of successive frames. In frames not containing the frame alignment signal, the first bit is used to transmit the CRC multiframe signal (001011), which defines the start of the SMF. Alternate frames contain the frame alignment word (0011011) preceded by one of the CRC-4 bits. The CRC-4 remainder is calculated on all 2048 bits of the previous sub-multiframe (SMF), and the 4-bit word sent as C1, C2, C3, C4 of the current SMF. (Note that the CRC-4 bits of the previous SMF are set to zero before the calculation is made.)

At the receiving end, the CRC remainder is recalculated for each SMF and the result is compared to the CRC-4 bits received in the next SMF. If they differ, then it is

	Sub-multiframe (SMF)	Frame number	Bits 1 to 8 of the frame (TS0)								TS 1-31
			1	2	3	4	5	6	7	8	9-256
Multi-frame	I	0	C1	0	0	1	1	0	1	1	P A Y L O A D
		1	0	1	A	S	S	S	S	S	
		2	C2	0	0	1	1	0	1	1	
		3	0	1	A	S	S	S	S	S	
		4	C3	0	0	1	1	0	1	1	
		5	1	1	A	S	S	S	S	S	
		6	C4	0	0	1	1	0	1	1	
	7	0	1	A	S	S	S	S	S		
	II	8	C1	0	0	1	1	0	1	1	
		9	1	1	A	S	S	S	S	S	
		10	C2	0	0	1	1	0	1	1	
		11	1	1	A	S	S	S	S	S	
		12	C3	0	0	1	1	0	1	1	
		13	E	1	A	S	S	S	S	S	
		14	C4	0	0	1	1	0	1	1	
15		E	1	A	S	S	S	S	S		

Note 1 E = CRC-4 remote error indication bits  
 Note 2 S = Spare bits

Note 3 C1 to C4 = Cyclic Redundancy Check-4 (CRC-4) bits  
 Note 4 A = Remote alarm indication

**Figure 7.7** The CRC-4 multiframe structure, where bit 1 of the frame alignment word carries the four CRC remainder bits for sub-multiframes 1 and 2. The E-bits provide a far-end block error (FEBE) indication, flagging that the remote terminal has detected a CRC block error.

assumed that the checked SMF is in error. What this tells us is that a block of 2048 bits had one or more errors. Each second, 1000 CRC-4 block error checks are made. Note that this in-service error detection process does not indicate bit error ratio (BER) unless one assumes a certain error distribution (random or burst) to predict the average errors per block. Rather, it provides a block error measurement.

This is very useful for estimating *percentage errored seconds* (%ES), which is usually considered the best indication of quality for data transmission—itself a block transmission process. CRC-4 error checking is very reliable; at least 94 percent of errored blocks are detected even under high BER conditions, according to ITU-T Recommendation G.706.

Another facet of the CRC-4 frame structure is the ability to transmit remote error detection back to the sending end. When an errored block is detected at the receiving end, the E-bit is set in the return path. This is termed a *Far End Block Error* (FEBE) or *Remote End Block Error* (REBE). By checking CRC-4, E-bits (FEBE), and A-bits (Alarms), an indication of both go and return paths is possible.

**Handling non-64 kbps traffic.** When the primary rate frame structures were conceived, it was assumed that nearly all the traffic would be standard 64 kbps PCM; in the 1990s, however, an increasing proportion has become data traffic. It is likely that these services will require more or less bandwidth than the 64 kbps channels available in the 2 Mbps frame.


Wideband services (such as videoconferencing, LAN interconnection, and high-speed computer links) usually require a bandwidth greater than 64 kbps but perhaps less than the full 2 Mbps (384 kbps, for example). These wideband signals can be sent in 30-channel, 2 Mbps frame by “sharing” the signal among several “aggregated” 64 kbps channels or  $N \times 64$  kbps bearer services (128–1920 kbps if  $N$  ranges from 2 to 30).

When aggregating 64 kbps channels, it is essential to guarantee bit sequence integrity, especially if the circuit passes through a switch. In other words, all  $N$  channels must undergo the same time delay. According to ITU-T Recommendation G.704, the  $N \times 64$  kbps signal is accommodated in  $N$  contiguous timeslots (omitting TS16), each timeslot taking consecutive octets of the traffic signal (Figure 7.8). If the remaining timeslots are unused for traffic, they should be filled with 1s. Of course, more than one  $N \times 64$  kbps signal may be carried in the 2 Mbps frame, depending on the bandwidth.

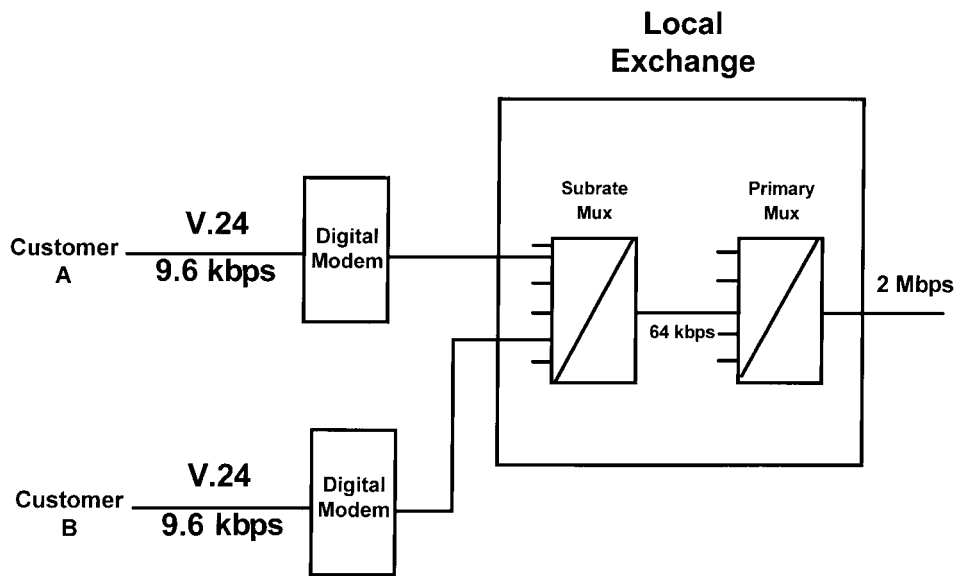
In practice it is not necessary to use contiguous timeslots, provided they are filled in an agreed-upon sequence and demultiplexed sequentially at the far end. An example of a noncontiguous plan is the recommendation for five 384 kbps channels (six timeslots each) given in ITU-T G.735. Sequences could be (1-2-3) + (17-18-19), (4-5-6) + (20-21-22), and so on.

Sometimes the full 64 kbps bandwidth is unnecessary, for example in applications that previously used analog data modems at 2.4 kbps or 9.6 kbps. *Subrate framing* allows a service provider to split up 64 kbps bandwidths into still lower-rate sections.

Before subrate, the choice was either low-rate over analog lines via a modem, or 64 kbps digital. Analog lines are expensive to maintain, however, and are incompatible with the modern integrated digital network. Nor do they offer the quality of ser-



**Figure 7.8** For transmission of wideband signals within the standard G.704 frame, individual 64 kbps channels are tied together to provide aggregate bandwidth of  $N \times 64$  kbps. The wideband signal is spread across several timeslots, so it is important that these are received in the right sequence at the far end; otherwise the payload will not be reconstructed properly.



**Figure 7.9** Subrate data multiplexing allows several low-rate channels at, say, 9.6 kbps to be packed into a single 64 kbps timeslot. This requires a digital modem at the customer's premises and a subrate mux at the local exchange. In North America this is referred to as the Digital Data Service (DDS).

vice offered by digital lines. Now a service provider can offer any of the following rates as well as 64 kbps: 0.6, 2.4, 4.8, 7.2, 9.6, 14.4, 19.2, 24, 32, 48 and 56 kbps, depending on the subrate standard.

Low-speed, widely distributed data networks, such as ATMs (automatic teller machine bank terminals) and EFTPOS (electronic fund transfer at point of sale), need little bandwidth to service each customer. For many such applications, 64 kbps bandwidth is too large.

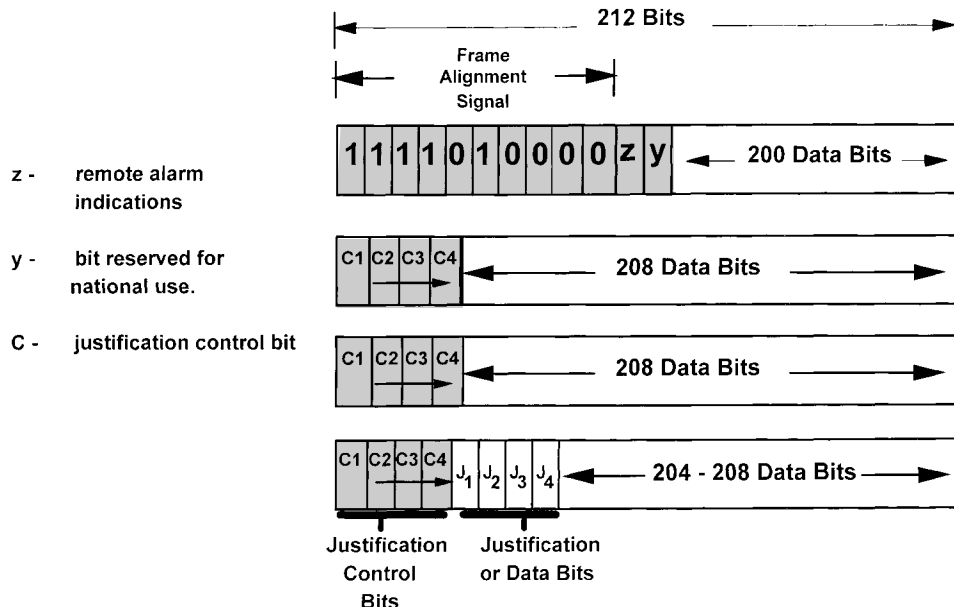
In the example shown in Figure 7.9, customers A and B “share” 64 kbps bandwidth with 9.6 kbps of data each. In turn, their 64 kbps signal becomes one timeslot of a 2 Mbps signal. In practice, up to five customers with data at 9.6 kbps would share one

64 kbps signal. Subrate data multiplexing is defined in ITU-T Recommendations X.50 and X.58. Subrate data services have been overtaken to some extent by the demand for wideband local loop access, with dial-up modem speeds of 28.8 kbps and above.

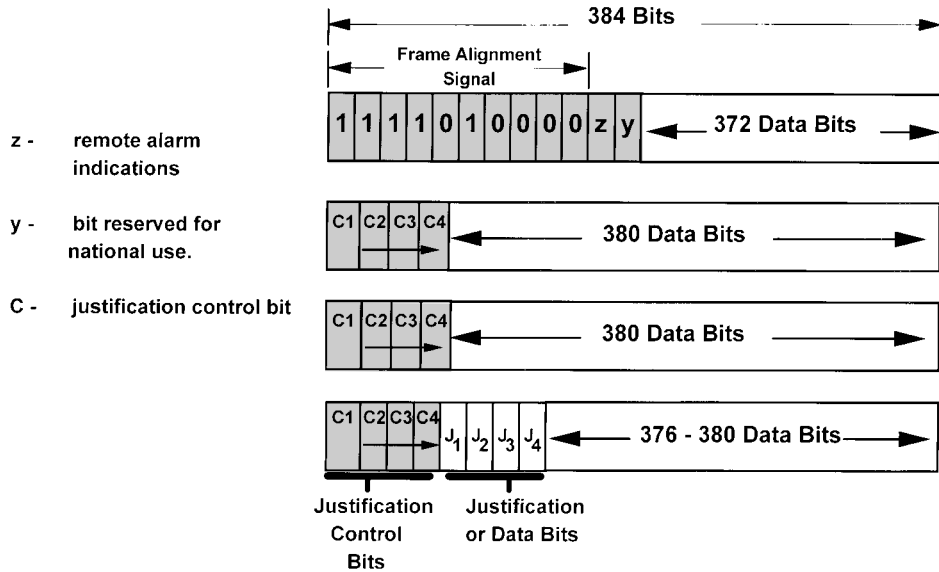
**Higher-order PDH multiplexing.** Above the primary rate, the PDH multiplex structure is complicated by the need to include positive justification for synchronizing the tributary streams. A full description of the multiplexing process can be found in **Reference 1** and also in the appropriate ITU-T recommendations. The description here will be limited to that necessary for understanding the measurement of higher-order multiplex signals.

The second-order multiplex process takes four tributaries at nominally 2.048 Mbps and multiplexes them together to form a transmit stream at 8.448 Mbps ( $\pm 30$  ppm). The frame structure of the 8.448 Mbps (E2) signal is shown in Figure 7.10, following ITU-T Recommendation G.742. The start of the frame is signified by the 10-bit frame alignment signal 1111010000. The 848-bit frame also contains 12 justification control bits and 4 optional positive justification data bits (one per tributary) that may be added, depending on the relative bit rates of the data streams and the need for bit stuffing.

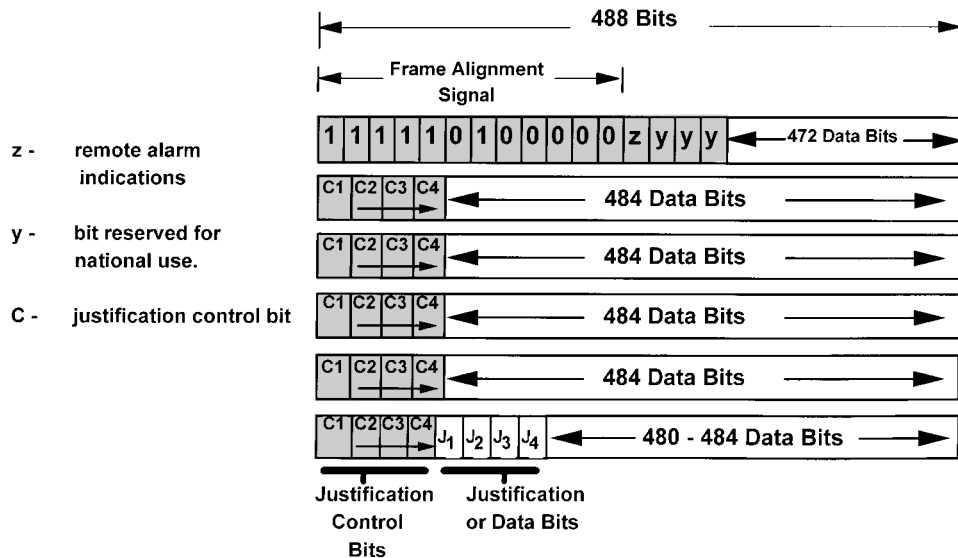
If all the control bits C1 are set to 1, then justification bit J1 has been added; if set to zero, then no justification bit is present. The same applies to the other justification bits. Note that the control bits are repeated three times throughout the frame (five times in the E4 frame). The demultiplexer uses a majority vote on the settings of the control bits (2 out of 3, or 3 out of 5) so that the justification process is very robust to bit er-



**Figure 7.10** The frame structure for the 8.448 Mbps (E2) rate according to ITU-T G.742, showing the position of the frame alignment word and the triplicated justification control bits C1–C4 (one per tributary). The frame length is 848 bits; the final section may contain 204 to 208 data bits, however, depending on the state of justification control bits. In the demultiplexer, the justification control bits indicate whether the justification bit positions should be ignored or read as valid data bits in the tributary system.



**Figure 7.11** The frame structure for the 34.368 Mbps (E4) rate according to the ITU-T G.751. The frame length is 1536 bits.



**Figure 7.12** The frame structure for the 139.264 Mbps (E4) rate according to ITU-T G.751. Note that in this case the justification control bits are replicated five times. The frame length is 2928 bits.

rors. If there were a mistake in the justification process, then the subsequent demultiplexer would lose frame alignment and a considerable amount of data would be lost.

Similar frame structures for E3 at 34.368 Mbps and E4 at 139.264 Mbps are specified in ITU-T G.751 (Fig 7.11 and 7.12, respectively). The E3 frame has a 10-bit



Frame Alignment Signal (FAS) and a 1536-bit frame; the E4 frame has a 12-bit FAS and a 2928-bit frame length.

The FAS bits are the only fixed information in these higher-level PDH frame structures, so in-service testing is limited to checking for errors in these bits.

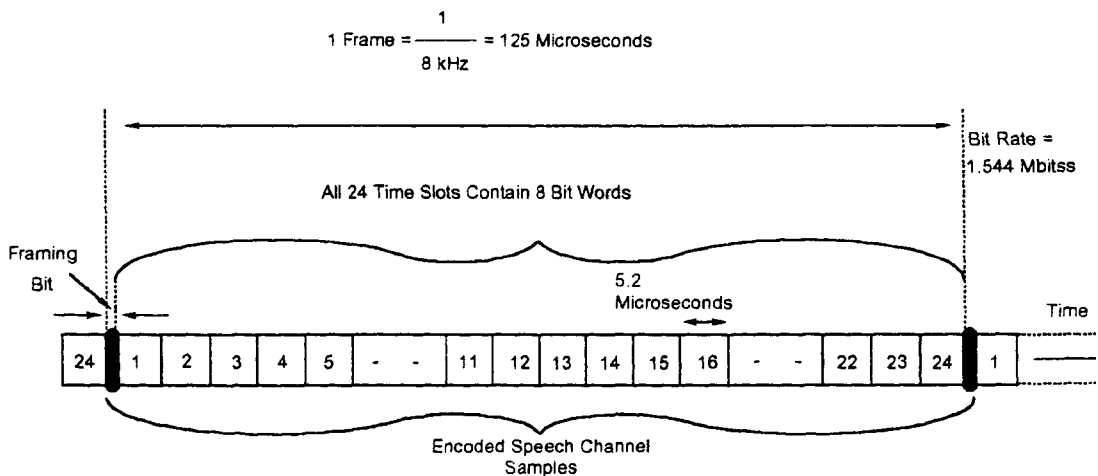
If a PDH multiplexer or demultiplexer loses the input signal or frame alignment, it sends out an “all-ones” Alarm Indication Signal (AIS), which is picked up downstream to set network alarms.

**7.2.2 North American 1.544 Mbps primary rate PDH frame structures**


The following section covers the basic elements of North American frame formats D4 (SF) and D5 (ESF), as well as fractional T1.

**Basic elements of the North American 24-channel frame.** As discussed earlier, a T1 frame is composed of 24 multiplexed timeslots (see Figure 7.13) with a framing bit to signify the beginning of the frame. Each timeslot contains an 8-bit word. Each bit occupies 648 ns, meaning each timeslot is 5.2 μs in duration. By adding 24 timeslots together, the total time of one T1 frame becomes 125 μs. Each timeslot can contain either sampled voice or digital data.

**SF/D4 framing.** Each T1 frame begins with a *frame bit* (see Figure 7.14), which enables the network to maintain synchronization and determines where the first timeslot begins. The D4 format uses every frame bit to verify frame synchronization; if two consecutive frame bits out of five are in error, the network equipment declares *Frame Synchronization Loss*. Subtracting the frame bits (8000 per second) from the T1 rate of 1.544 Mbps, the maximum payload rate is 1.536 Mbps for a full T1. Many users require only a fraction of that, which allows service providers to multiplex many users onto a single T1.



**Figure 7.13** A DS1 frame consists of 24 DS0 timeslots multiplexed together with a frame bit at the beginning. Each frame is 125 μs long and may carry voice or data.



- Every f-bit used for framing
- In-service testing looks for errors in pre-determined f-bit sequence 100011011100
- Payload data rate is 1.536 Mbit/s

**Figure 7.14** A D4 Super Frame (SF) consists of 12 DS1 frames multiplexed together. The 12 framing bits, one from each frame, occur every 192 bits.

**North American D4 Superframe (SF) framing.** By grouping 12 T1 frames together, a D4 *Superframe* (SF) is created (see Figure 7.15). Using the framing word 100011011100, the network is able to separate each frame for demultiplexing. The framing bits are broken into two types, *Framing Terminal* (Ft) and *Framing Signaling* (Fs). Although all bits are used for frame synchronization, the network uses the Fs bits to indicate where the voice channel signaling bits (AB) are. The AB bits indicate the status of a voice call; for example, On/Off Hook, Ring, Busy, and Wink overwrite the least significant bit in each timeslot in the 6th and 12th frame only.

**SLC-96® frame format.** The original T-carrier systems served the purpose of trunk communications (Central Office to Central Office). The SLC-96® system was introduced in 1979 to capitalize on the advantages of T-carrier within the local loop. The system provides for 96 subscriber channels to be transmitted over multiple T1 lines between a remote terminal and the Central Office. Thus was the acronym *SLC-96®* derived, meaning “Subscriber Loop Carrier for 96 channels.”

The SLC-96® system was an extension of the D4 channel bank family, and initially used many of the same plug-in circuit packs and mechanics. The SLC-96® system was implemented to integrate the T1 line interface and span powering into the same shelf, saving physical space and providing single-ended system maintenance (because craftspeople would not be available at both ends of the link).

There are three operational modes for SLC-96®:


- *Mode 1* Four T1 lines each carry signals from 24 subscribers without concentration. A fifth T1 line is used for protection switching (redundancy).

- *Mode 2* Two T1 lines are used for main communications, with a third being used for protection switching. Forty-eight subscribers share each main T1 line on a first-come, first-served basis.
- *Mode 3* There are two main T1 lines, plus a third for redundancy. The system is used for special services (nailed-up connections) or pay phone use, and is limited to a maximum of 48 channels (24 on each T1 line).

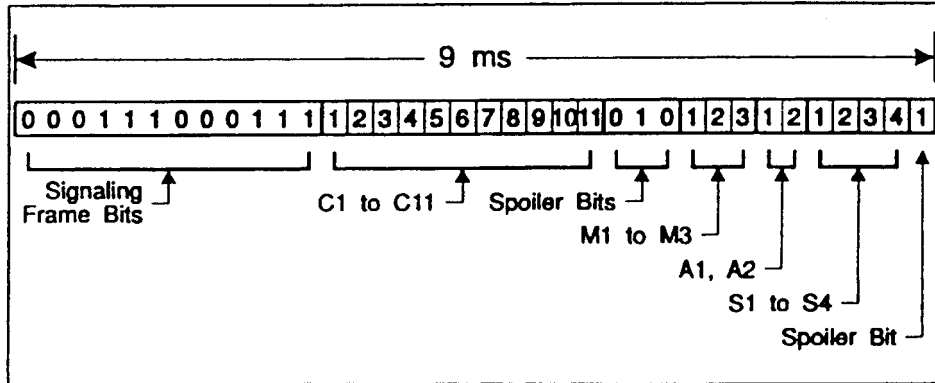
Since 1986, the Series 5 Digital Loop Carrier system has gained popularity as an alternative for new installations instead of SLC-96®.

**Frame structure of SLC-96.** The Ft bits are retained as in D4 format, and alternate 1, 0. The Fs bits form a data link. Information between the remote terminal and the central office terminal regarding the status of the system is carried in this data link. Figure 7.16 shows the data link frame structure. The first 12 bits are used for synchronization, and the remaining 24 are grouped in order of their transmission into six fields:

1. *Concentrator field (bits 1–11)* This is used to control channel assignment/de-assignment.
2. *First spoiler field (bits 12–14)* This field contains the fixed pattern 010 and is used to prevent the receiver from misframing.
3. *Maintenance field (bits 15–17)* This field is used to control channel and drop testing.



**Figure 7.15** D4 Super Frame utilizes Framing Terminal (Ft) bits for network synchronization and Framing Signaling (Fs) bits to identify the timeslots that carry the AB signaling bit indicators, which are used for On/Off Hook indications.



**Figure 7.16** The SLC-96 Data Link frame structure, which carries information between the remote terminal and the central office. The first 12 bits are used for synchronization; the remaining 24 bits are grouped into six information fields.

4. *Alarm field (bits 18–19)* This field is used to carry alarm information and control commands.
5. *Protection line switch field (bits 20–23)* This field is used to control switching of the protection DS1 line.
6. *Second spoiler field (bit 24)* The final field of every data link frame consists solely of a single bit set to 1 and again is used to prevent the receiver from misframing.

**Extended Super Framing (ESF).** In the SF format, the network uses all 12 frame bits for frame synchronization. With the advent of more reliable T1 spans, the *Extended Super Framing* (ESF) format evolved, in which not all the overhead is needed for synchronization. The D5 (ESF) frame format consists of 24 frames combined to make one Extended Superframe (ESF). The line rate remains the same (1.544 Mbps), but the frame synchronization word changes from 12 bits to 24 bits, and not all 24 bits are used for frame synchronization.

Just like the SF format, ESF moves 1.536 Mbps of customer data, but the 8000 bits of framing overhead is divided up. The overhead is divided into three parts: 4000 bits of user data link, 2000 bits for CRC, and 2000 bits for framing. The user data link can be used by the customer or by the network for Performance Report Messages (PRM), which are used for far-end reporting of the T1. The CRC is a mathematical calculation done on the previous T1 frame and passed to the far end, where it is compared to the calculation done there. This process is approximately 99 percent accurate. The frame bits are used the same way as in SF.

**North American D5 (ESF) framing.** Figure 7.17 illustrates how an ESF T1 is defined. Also similar to SF are the signaling bits for voice services: the ABCD bits in the 6th,


12th, 18th, and 24th frames. Note that these formats do not determine the line rate, but are used by the network for synchronization, in-service error checking, and identifying where customer data resides.

**ZBTSI frame format.** The *ZBTSI frame format* was introduced in 1987 as an extension of the ESF format. It was developed to allow transmitting of strings of 0s in the payload without the need for BZS line coding. This bit-sequence independence permits network operators to obtain clear-channel capability at 64 kbps without having to replace line equipment.

The ZBTSI frame format identifies strings of 0s that would cause a signal to violate the pulse density requirement. Those octets so identified are replaced with nonzero octets. Half of the Frame Data Link (FDL) bits contain flags, referred to as *Z-bits*, used to indicate whether a substitution has taken place.

The ZBTSI frame format still provides the advantages of ESF because the 6-bit CRC error detection remains. A disadvantage of the ZBTSI frame format is that a signal incurs a delay of four frames when encoded and then decoded. Speech traffic may be noticeably impaired if the signal passes through many digital-access and crossconnect switches (DACS).

**Fractional T1 traffic.** Fractional T1 service allows customers to purchase only the amount of bandwidth they require, rather than a whole T1. Fractional T1 services are



**Figure 7.17** The North American D5 Extended Super Frame (ESF) format combines 24 DS1 frames and uses the frame bits for a CRC check, Frame Alignment, and a 4 kbps data link. It also has two additional signaling bits C and D, for a total of 4 ABCD bits.


sold in contiguous or noncontiguous sections of bandwidth. In the contiguous mode, the customer purchases DS0s in adjacent timeslots. Applications that require uninterrupted bandwidth (such as video) use this type of service. The noncontiguous mode allows the service provider to multiplex other customers into the unused bandwidth. Applications that are not time-sensitive (reassembled with a small amount of delay after being divided up across the T1) can be transported over this type of service.

**Multiplex hierarchy in North America.** In the DS1 Network there are 24 DS0 time-slots that can carry voice or data. By combining multiple DS1 signals together, we form a DS3. In order to accomplish this, the network samples 28 DS1 signals and first combines them into a DS2 running at 6.312 Mbps. As shown in Figure 7.18, the DS1 signals are combined in groups of 4 DS1 signals to make a DS2. These DS2 frames are then combined to form the DS3 signal running at 44.736 Mbps.

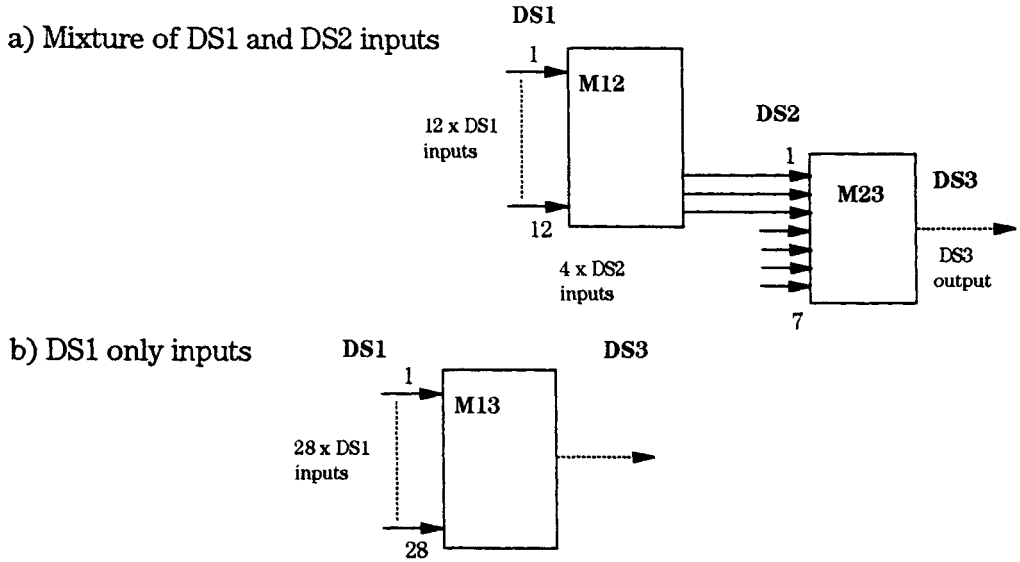
**Multiplexing from DS1 to DS3.** Depending on the multiplexer, 28 DS1s may be combined to form the DS3 or, in some cases, a DS2 may be input directly into an M23 multiplexer to be combined to make the DS3 (see Figure 7.19). Video applications may use this because video signals run at 6 Mbps. When combining DS1 signals, the multiplexer will bit-interleave the 28 DS1 signals to form the DS2 frame, which then will be interleaved with the other DS2 subframes to form the DS3.

**M13 Justification.** During the M13 multiplexing scheme, it might not be possible for all the DS2 signals to run at the same rate. In this case the multiplexer must “stuff” extra bits to ensure that all the DS2 streams run at the same rate. So that the bits can be inserted and removed, the framing overhead will indicate whether or not stuffing has occurred. The C-bits in the DS3 overhead indicate the presence or absence of stuffing bits. If all three are 1, stuffing has occurred; if all three are 0, no stuffing has occurred.


**DS2 signal-framing structure.** Figure 7.20 shows an example of the stuffed bit positions, where the last data field contains the stuffed bit. As mentioned earlier, the value of all three C-bits will indicate whether the stuffed bits are being used for the network or customer data. This diagram also shows the M-bits, which are used for



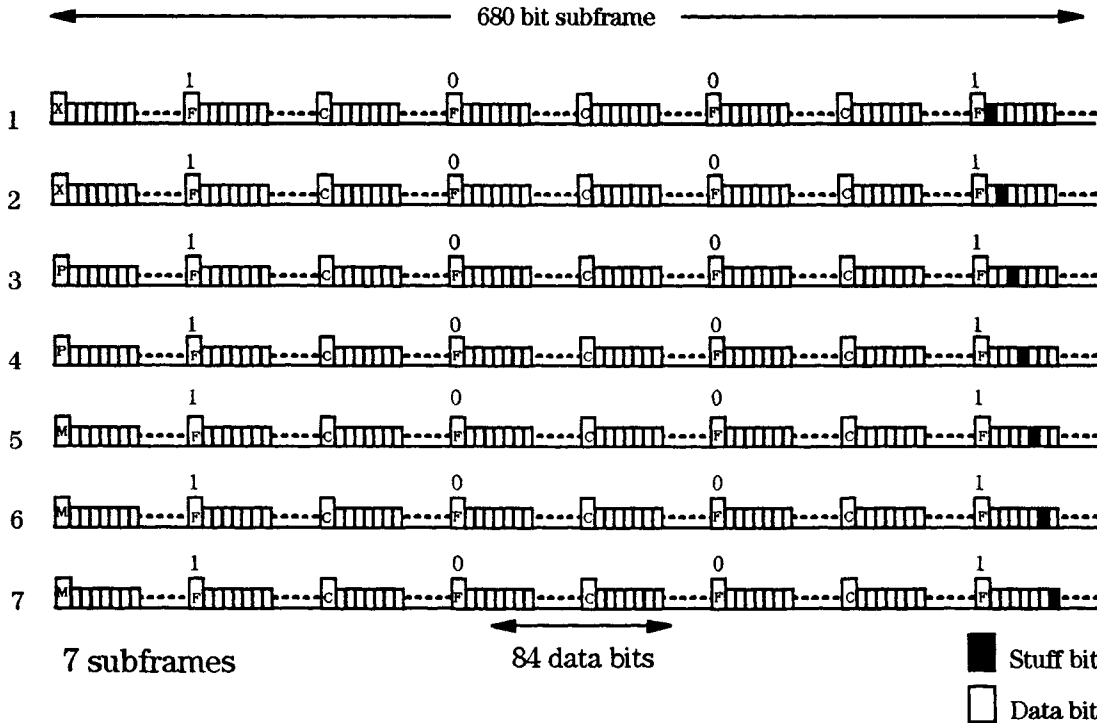
**Figure 7.18** The typical North American multiplexing hierarchy is to multiplex twenty-four 64 kbps DS0s into a DS1, which is transmitted at 1.544 Mbps. Then a two-stage multiplexing scheme is used to multiplex four DS1s together to form a DS2 at 6.312 Mbps. Then seven DS2s are multiplexed together to form the DS3, which is transmitted at 44.736 Mbps.



**Figure 7.19** In some cases (such as video services) there might be a requirement to multiplex directly at the DS2 rate. In these cases an M23 multiplexer can be used. The more common practice is to multiplex DS1s to DS2s, then to DS3 as shown in Fig. 7.18.



**Figure 7.20** A DS2 frame is constructed of customer data interleaved with network overhead information. These bits maintain the synchronization as well as indicate the DS2 subframe position.



**Figure 7.21** A DS3 frame consists of the DS2 subframes interleaved together with additional network overhead. In much the same way as with DS2, this network overhead is used to maintain frame synchronization but also is used for stuffing indicators, parity checks, and subframe indicators.

multiframe indicators, and the F-bits, which are used for DS2 frame synchronization. Between each of the overhead bits are 48 bits of data. These 48 bits are pieces of the 4 DS1s, which are bit-interleaved together to form the data portion of the DS2.

**DS3 signal-framing structure.** Figure 7.21 is an example of the DS3 M13 frame structure. The X-bits are used for user-defined message sets. Moving down the first column, the P-bits are parity bits, which are set after the mux looks at all the information bits in the DS3 frame. There are two P-bits per DS3 frame, and their state will always be identical. If there is an odd number of 1s in the DS3 frame, then the bits will be set to 1; if the number is even, then the bits will be 0. The M-bits (multiframe bits) are used to maintain multiframe alignment in the DS3 frame, the F-bits are used for frame alignment, and the C-bits will indicate stuffing. The data portion of the DS3 frame is made up of bit-interleaved samples of the DS2's frame and data fields.

### 7.3 Measurement Issues

#### 7.3.1 T-carrier services: What can go wrong

A number of common problems can crop up in T1 services installation. When conditioning analog copper to T1 copper, all bridge taps and load coils must be removed from the pair; if they are not, the T pulses may be misshapen and unrecognizable to



other network devices. If network equipment is not set to drive a strong enough signal out, or is expecting a low-level signal in and is overdriven, problems also will arise.

When regenerating T1 signals, network equipment must be set to the correct timing source, i.e., loop, internal, or external. If it is not, timing slips may occur.

All network equipment must be set for the correct frame type and line coding. If an element is set to B8ZS in an AMI system, this causes data errors and therefore low throughput. Faulty network equipment (poor grounding, bad copper, or drifting oscillators) will cause errors and down time.


### 7.3.2 Loopback types

When testing a PDH service, a common practice is to have a *loopback* at the far end. This can be accomplished by a hard loop connecting transmitting ports to receiving ports, or by sending a loop code in the data stream (see Figure 7.22). This will bring the transmitted pattern back to the tester, where a BERT (Bit Error Rate Test) can be performed.

There are two ways to send loop codes in the T1 network, in-band or out-of-band. The in-band loop codes are transmitted over the 1.536 Mbps of customer data and cause the whole pipe to be looped back for BERT testing. SF circuits can use only in-band loop codes; ESF circuits can use out-of-band loop codes as well. With out-of-band loop codes it is possible to loop either the data and framing, or loop the data only and recalculate the framing and CRC. This technique can be used to determine if the CSU is causing a framing problem.


Although out-of-band loop codes will initiate a loop up/down more quickly, the main advantage is shown in Figure 7.23. By using the out-of-band line loopback, which loops data and overhead, CRC errors on the transmit side will be looped back to the tester. Switching to the payload loopback will cause the CRC to be recalculated,

- Controlled by in-band codes (over all timeslots)
- Controlled by out-of-band (ESF) codes - two types:



**Figure 7.22** In ESF circuits two types of CSU loopback codes can be sent. *Line* loops back all information, while *Payload* loops back the data only and allows the network overhead to be recalculated.

- Can be set up/torn down quicker
- CRC errors from line loopback but not from payload isolate problem to "go" part of the circuit



**Figure 7.23** One advantage of ESF loopbacks is that, with a Line loopback it is not possible if an error occurs to determine which direction the error is in. By using Payload loopback, the network overhead will reframe and recalculate the CRC, which might remove the error from the return path.

so CRC errors would not be present at the tester. The line and payload loopbacks can be used in this function to determine which side has the problem. Another method of determining which leg (Tx or Rx) has the problem is to run a full-duplex or back-to-back test. In that case, a test set would be used at each end of the circuit; seeing which tester receives the errors will indicate which leg has the problem.

### 7.3.3 Code errors (bipolar violations)

The following section will review some common measurements of PDH services. We will cover T1 first and finish with T3. With an AMI line-coded signal, every other 1 must be of the opposite polarity. Two consecutive pulses with the same polarity (Figure 7.24) are known as a *bipolar violation*, or BPV. In a test using a known pattern, a 1 that does not occur when it is expected is considered a bit error. If a 1 occurs in the same polarity as the previous one, it will be considered a BPV because this violates the AMI rule. The network will try to correct the event, which might cause a bit error. By making an in-service measurement of BPVs, we can test the quality of the service. This is one of the advantages of the T1 network.

### 7.3.4 Error measurement in digital transmission systems

Listed below are some typical errors that can be measured in a T1 system:

- Bit errors (logic errors)
- CRC errors
- Frame errors
- BPV

Some error types are in-service, while others are out-of-service. The best example of out-of-service testing is a bit error test (BERT). This is simply transmitting a known set of 1s and 0s, receiving them, and verifying that they match the transmit-

ted set. Out-of-service testing is the most accurate, but is not always available. Sometimes in-service testing is the most desirable, however. If used correctly, in-service error verification can be a good indicator of the performance of the T1 circuit.

### 7.3.5 Jitter

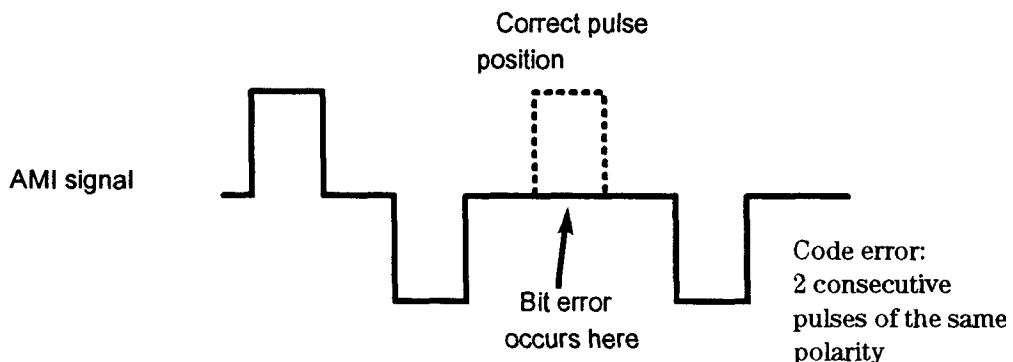
The following section will cover T1 jitter. A small amount of jitter is present in all T1 networks because every regenerating device adds some degree of jitter to the network. Jitter measurements have two parameters, frequency and amplitude. Amplitude is expressed in U.I. (Unit Intervals) and defines how far the pulse has moved from where it was expected to be in time. Frequency defines how quickly the pulses are moving.

**Jitter sources.** Jitter is caused by network equipment such as repeaters, regenerators, and higher-order multiplexers. T1 office and line repeaters will cause high-frequency jitter because the T1 pulses are regenerated at 1.544 Mbps, while higher order multiplexers such as M13 or SONET multiplexers will stuff bits at a low speed and cause low-frequency jitter. Random jitter (discussed in the next section) occurs sporadically and can be caused by various events.


**Random and systematic jitter.** Random jitter is sporadic and may have both positive and negative components. It can be caused by network equipment failures and noise, as well as atmospheric conditions. Systematic jitter may be of high amplitude (because it accumulates throughout the network resulting from pattern-dependent effects), and presents a more serious problem. If jitter values become too high, the bit, BPV, or frame errors can occur, causing low throughput.

**672411 jitter mask.** Figure 7.25 shows a Bellcore 672411 jitter mask. If a network device is not able to operate without errors and has jitter input at values inside the mask, the device is said to “fail the mask” and might require service. If the device operates error-free with jitter values that exceed the mask, the device is considered robust.

**Eye diagram and jitter.** Figure 7.26 is an example of what jitter looks like on an oscilloscope. It is evident that the pulses are moving back and forth in time, which makes



**Figure 7.24** Bipolar Violations (BPV) occur when consecutive pulses of the same polarity appear on the line.



**Figure 7.25** Here is an example of the Bellcore 62411 Jitter Mask. Note that wideband jitter is from 10 hz to 40 KHz, while highband jitter is from 8 KHz to 40 KHz.

it difficult for network equipment to recover and regenerate the T1 pulses at the correct point. Also shown is an example of frequency and amplitude and how it affects the pulse placement, which in turn affects the next regenerating device. If an excessive amount of jitter exists in the service, the end user may experience intermittent errors because the network cannot reproduce the pulses at the correct point in time.

**Jitter summary.** Jitter is the relative phase difference of the received pulse to a reference pulse, and can accumulate throughout the network. This can cause bit, BPV, or frame errors. Properly designed, quality network equipment can help to minimize network jitter.

## 7.4 Out-of-Service Testing

Out-of-service testing requires revenue-earning traffic to be removed from the transmission system so that a wideband test signal can be applied. This is disruptive, so this type of testing is usually applied in production test and when installing new equipment, though it may be used briefly when checking a system after repair in the field. The advantage of out-of-service testing is that it tests fully the performance of the transmission, since every transmitted bit is checked for errors. Furthermore, a variety of test patterns can be used to explore the limits of performance, and the equipment's in-service performance monitors and alarms can be checked by applying known faults and degradations to the test signal.


Because the North American and international PDH standards are different, they have different test requirements, which are described separately below.

### 7.4.1 Testing international 2.048 Mbps primary rate PDH systems

Out-of-service test equipment is defined in ITU-T Recommendations O.150, O.151, O.152, and O.153. These are concerned primarily with defining the pseudorandom binary sequence (PRBS) test patterns to be used at different bit rates. These are summarized in Chapter 27, Table 27.1. Most out-of-service tests use PRBS test patterns with or without framing. When framing is used, the PRBS fills the payload area, and is stopped momentarily while the framing bits are inserted.

Because European standard PDH systems are bit-sequence-independent, no frame structure is required to check transmission quality or error performance. For these tests it is necessary only to have a pattern generator (conforming to the ITU-T standards) at one end of the link, and a matching error detector at the other end. The design of these test sets and the application of error performance standards are described in detail in Chapter 27. As discussed earlier, the standard interface in PDH transmission networks is defined in ITU-T G.703, so the pattern generator transmitting output and error detector receiving input also should conform to these specifications for the bit rate under test.

On the receiving side, G.703 specifies that equipment should operate error-free with a maximum amount of cable attenuation following a  $\sqrt{f}$  characteristic, as shown



**Figure 7.26** The usefulness of eye diagrams can be seen by examining a DS1 pulse with jitter on an oscilloscope. The pulse may appear to have a ghosting effect. The DS1 pulse stream may be moving in time and have frequency and amplitude values that indicate how fast and how much the DS1 is shifting.

**TABLE 7.2 Allowable Cable Losses at Receiver Input.**

Bit Rate	Maximum Cable Loss at Half Bit Rate ( $\sqrt{f}$ law characteristic)
2.048 Mbps (E1)	6 dB @ 1.024 MHz
8.448 Mbps (E2)	6 dB @ 4.224 MHz
34.368 Mbps (E3)	12 dB @ 17.184 MHz
139.264 Mbps (E4)	12 dB @ 70 MHz

in Table 7.2. The test set receiver therefore should incorporate a fixed or variable cable equalizer to compensate for this roll-off. A variable equalizer is preferable because it optimizes the signal-to-noise ratio for different cable lengths and avoids overcompensation on short cables, which would cause overshoots, potentially creating errors on pseudoternary codes like HDB3.


If the test signal is framed, and the error detector is capable of decoding a framed signal, then a number of additional out-of-service tests for checking the alarm operation and error performance monitoring are possible (such as CRC block error checking in the G.704 2 Mbps frame). If tests are required across a multiplexer or demultiplexer, transmitting at one bit rate and receiving at another, then a framed tester (Figure 7.27a) is needed to inject or extract a tributary signal. If only unframed testers are available, however, then the test can be made by “looping around,” as shown in Figure 7.27b.

Two levels of framing capability may be provided for the error detector receiver. At the simpler level, it may be capable of detecting the frame word at a particular hierarchy level and checking for errors. This is a useful in-service test capability. Alternatively, the receiver may have full demultiplexing capability, in which case it will be able to extract a complete tributary stream from a higher-level signal (for example, a 2 Mbps signal from a 140 Mbps carrier), and do a full analysis on the tributary. This is very useful for checking an end-to-end, low-order digital path as it passes through high-capacity network nodes and crossconnect switches. In effect, the test set becomes like a programmable demultiplexer; usually it can only deal with one tributary at a time, however, in contrast to the complete function of the operational equipment.

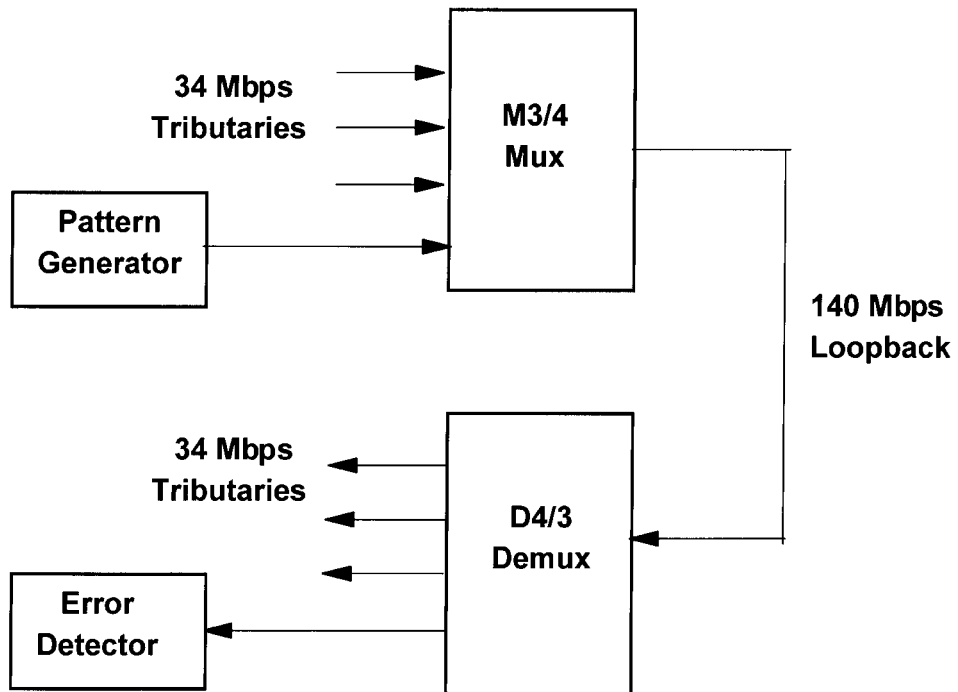
A fully framed test signal at 2 Mbps is particularly useful for analyzing primary rate leased lines. For example, one can send a PRBS test signal in one or more 64 kbps timeslots and check for timeslot integrity at the output of a crossconnect switch. Similarly, a multiplex can be stimulated with a 64 kbps signal and a check made in the appropriate timeslot of the outgoing 2 Mbps stream.

A fully framed 2 Mbps test signal also allows one to check the proper operation of alarms and performance monitoring within the multiplexer. Send errors in the frame alignment signal either continuously or in burst mode; check the loss of frame criteria and resynchronization in the multiplexer. Simulate CRC block errors and alarms and check how the multiplexer responds locally and through E- and A-bits on the outgoing stream.

An interesting application of PRBS test signals is the measurement of round-trip delay. The circuit is “looped back” at the far end, and a long-sequence PRBS such as  $2^{23}-1$  is sent from the test set. The time delay is computed by correlating the received signal with the transmitted pattern, and, for example, delays up to 1 second can be measured to a resolution of 1 ms. Round-trip delay is becoming more important with



**Figure 7.27a** Testing across a multiplexer requires a framed or structured receiver so that the test tributary can be demultiplexed for error checking. The PRBS transmitted from the pattern generator then can be recovered and checked for errors in the receiver.



**Figure 7.27b** If a structured test set is not available, a loopback test is possible by recovering the test tributary through a demultiplexer. Because mux and demux usually are co-located, this is a frequently used procedure.

increased use of real-time data services, particularly frame relay applications for interconnecting LANs.

**Tests on  $N \times 64$  kbps circuits and subrate data multiplexing.** The foremost requirement here is to check the integrity of a wideband signal spread across several timeslots, providing an aggregate channel. Some network equipment may treat the individual 64 kbps channels as independent entities, switching them to different timeslot positions or even rerouting them through different paths. If this happens, the multiple timeslots of the  $N \times 64$  kbps signal will not arrive at the destination in the right sequence, and some could be missing altogether. It will be difficult or impossible to reconstruct the wideband signal.

To test an  $N \times 64$  kbps channel, one can inject a separate PRBS pattern in each allocated timeslot, and check for continuity and error-free reception at the far end on a channel-by-channel basis. It is more realistic to “spread” a single PRBS across the sequence of timeslots allocated to the wideband channel, however, just as the live signal would. In this approach, the first octet of the PRBS would go to the first timeslot, the second octet to the second timeslot in the plan, and so on. The unused timeslots usually are filled with all 1s, though a PRBS such as  $2^6-1$  also can be used.

For the wideband signal to be received without error at the far end, not only would the timeslot allocations have to be maintained, they would also need to arrive in the right sequence. The integrity of the  $N \times 64$  kbps circuit thus would be proved. If a timeslot had been misplaced, then one would need to send an identifiable pattern (an 8-bit word, for example) in each timeslot and search for it at the receiving end. When doing a loopback test at  $N \times 64$  kbps, it is possible that the return path might use a different timeslot allocation. In such a case, the tester would need to have independent settings of transmitter and receiver timeslots. An example of a 384 kbps channel based on the G.735 recommendation mentioned earlier is shown in Figure 7.28.

For subrate data multiplexing, the test set needs to implement the ITU-T X.50 and X.58 frame structures within the 64 kbps channel. This capability is sometimes included in testers designed for installing and maintaining 2 Mbps digital leased lines. For more information, consult **Reference 2**.

#### 7.4.2 Testing North American PDH systems

When access to the T1 is available, out-of-service testing can be performed by running a BER test. Faults can be isolated as either a transmitting or receiving problem by dividing the circuit in half and working back to the testing location. Another advantage to out-of-service testing is to establish operating parameters such as signal levels, pulse shapes, and power levels. By running specific data patterns, a BERT baseline may be established for future out-of-service testing.

The same impairments occur on DS3 transmission lines as occur on DS1 lines and, just like DS1 out-of-service testing, may be used to isolate problems. By analyzing the receiving-end results (parity errors, bit errors, FEAC codes) it is possible to determine whether the problem is in transmitting or receiving. The ability to baseline the service upon circuit turn-up, prior to any problems, may save time on finding and solving future problems. Finally, qualifying the service before turning it over is a must, and a BER test will prove that the circuit is acceptable.



```

APPLICATION [ 2Mb/s ]
INTERFACE [ TERNARY ] LINECODE [ HDB3 ]
FRAME (G.704) [ CAS MFM ]
THROUGH MODE [ OFF ]
TX CLOCK SOURCE [ INTERNAL ]

PATTERN [ 16 BIT WORD ]
[1100101010010101]
RECEIVE TIMESLOT [ SELECT ]
Tx [ F***.....S***..... ]
Rx [ F...***.....S...***..... ]
    0 |-----|-----|-----|-----|-----|-----|-----|-----|-----| 31
      Bandwidth Tx 384kb/s Rx 384kb/s

ALARM GENERATION [ OFF ]
ERROR ADD [ BIT ] [ SINGLE]
STATUS:
2Mb/s 704kb/s 8Mb/s 64kb/s MORE

```

**Figure 7.28** Testing a 384 kbps bandwidth facility ( $6 \times 64$  kbps) showing different timeslot allocation in the transmit and receive direction according to ITU-T Recommendation G.735 (1, 2, 3 & 17, 18, 19 in the transmit direction, 4, 5, 6 & 20, 21, 22 in the receive direction). In this case the active timeslots are identified by asterisks, and the test pattern is a user-defined 16-bit word.

Most DS3 BER tests are performed using one of three types of patterns:

- PRBS
- Fixed
- User-defined (see below)

Pseudorandom Bit Sequences (PRBS) such as  $2^{23}-1$ ,  $2^{15}-1$ , and  $2^{20}-1$ , are used most commonly. Fixed patterns such as all 1s or AIS also can be used to verify alarms. In certain instances, a specific sequence of 1 and 0 may cause errors to occur, in which case users might want to build their own patterns. Although different patterns can cause faults on the service, it should be remembered that DS3 is always B3ZS-encoded so patterns with long strings of 0s may not stress the circuit more than others.

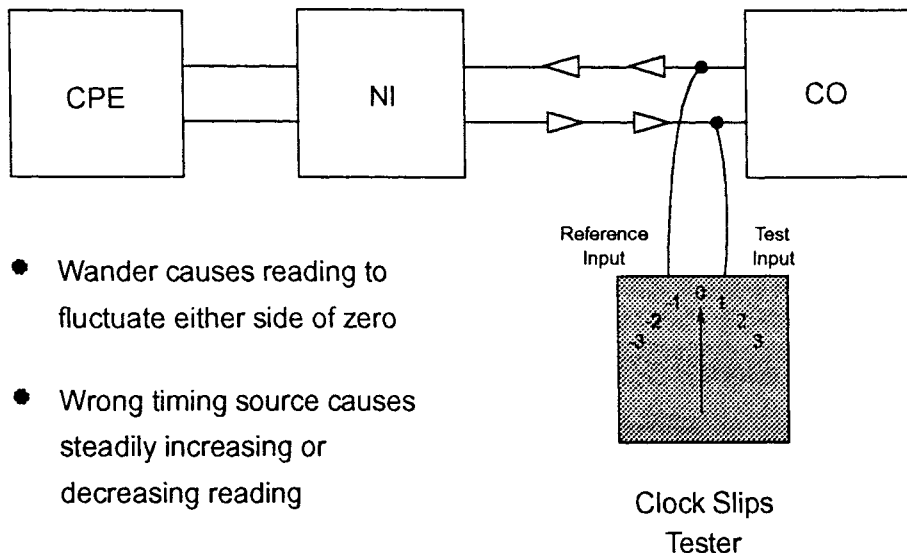
**Test patterns.** The patterns below create the multipattern and commonly are used to qualify T1 lines. A *Quasi-Random Signal Source* (QRSS) is a good digital simulation of voice traffic and has a string of 14 consecutive 0s, which stresses repeating net-

work equipment. An all-ones pattern will force office and line repeaters to run at full power and can help find marginal equipment. A 3-in-24 pattern can force repeaters to run at the minimum power level and stress the clock-recovery circuits. By running 2-in-8 and 1-in-8 patterns, equipment that is misconfigured for B8ZS/AMI can be identified. The 2-in-8 will not enable B8ZS line coding, but the 1-in-8 will. This causes errors to be counted on the 1-in-8.

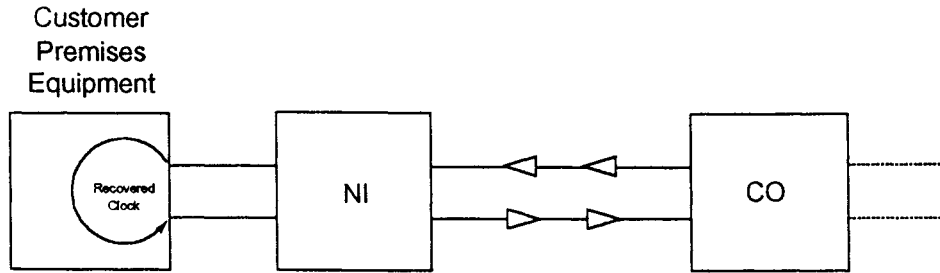
**Comparing two clocks.** When a CSU is suspected of being misconfigured for timing, a *timing slip analysis* should be performed. In order to accomplish this, an operable T1, such as the inbound T1 from the service providers, is used as a reference (see Figure 7.29). This is an in-service test and will not affect customer data; if slips occur, however, the CSU might have to be replaced or reconfigured, which will affect service. If the slips count in both directions, this can indicate T1 wander or jitter.

**Timing problems.** Timing in the T1 network is critical. If a network device is not configured correctly, timing slips and errors can occur. The CSU example in Figure 7.30 should be configured for loop timing. By forcing a loop at the NI (Network Interface) or “Smart Jack,” the service can run error-free. When the CSU is looped back and the circuit counts errors, the CSU could be bad or misconfigured. Typically in T1 services, the timing is supplied by the service provider, which may be a Long Haul Carrier, Regional Bell Operating Company (RBOC), or Competitive Access Provider.

**Applications: Testing T1 services.** If an out-of-service test needs to be performed on the entire T1 from mid-span, a round-robin test will check both directions (Figure 7.31). In this instance, a loop must be present at both ends to receive the signal back




**Figure 7.29** Many times during an installation the timing source on a CSU is incorrectly set. By comparing the transmit and receive lines, a clock slip test can be performed. Any slip measurement indicates a problem on the line and is undesirable.



- CPE should transmit using received clock (loop timing)
- Testing to NI loopback doesn't show up the problem

**Figure 7.30** When testing a T1, looping the far end CSU may show errors. If looping the far end NIU makes the circuit run error-free, this could indicate a timing or even a cabling problem between the devices.



**Figure 7.31** When testing T1 services from mid-span, a round-robin test can be performed. This will verify the service from end to end. If any errors occur, a full duplex test may be required to isolate the problem.

at the test set. Although a BER test is run and lines can be verified, there is no way to identify from which direction possible errors are originating. If errors do occur, a full-duplex end-to-end test can determine if a transmitting or receiving problem ex-

ists. The problem can be isolated by installing loops and working backward to the test site.

## 7.5 In-Service Testing

### 7.5.1 Monitoring in-service QoS

Monitoring the Quality of Service (QoS) while a system is in operation and has traffic has become one of the most important issues in an increasingly competitive telecommunications market. Being able to guarantee performance levels and detect degradations before customers notice them are differentiators for the operator, and are an important ingredient of the Controlled Maintenance strategy (ITU-T Recommendation M.20) described in Chapter 5.

In the deregulated competitive environment, in-service monitoring is very desirable because it allows the operator to monitor quality of service continuously and determine whether degraded performance is being caused within the operator's network or by another vendor. Often the customer's traffic will traverse several different networks between source and destination. The different parts of the route are referred to as network *sections*, while the end-to-end connection is referred to as the *path*. Monitoring overall path performance provides an indication of the service the customer receives, while section measurements are important for troubleshooting and countering "finger-pointing."

In-service measurements cannot rely on a bit-by-bit error check, as is possible with an out-of-service PRBS test; the monitoring equipment has no way of knowing the data content of random customer traffic. Instead, in-service testing must check for errors in any known fixed patterns (such as frame words in the random data stream), or must apply error detection codes (EDCs) to blocks of data. The most powerful detection processes are based on computing parity or checksums on blocks of data, including the payload bits, and transmitting the result to the far end for re-computation and comparison.


A number of technical developments on the transmission side have specifically aimed to improve in-service monitoring, notably the revised primary rate frame formats at T1 (ESF) and E1 (CRC-4). One of the driving forces for developing the new SDH and SONET standards was to improve the built-in network monitoring and control of high-capacity transmission links. As mentioned earlier, the transmission of wideband unformatted data, while possible on clear-channel E1–E4 paths, creates problems in performance monitoring because there is no recognizable data pattern to check for errors. A contribution of SDH systems is that these unformatted E1–E4 data streams now can be packaged in a virtual container, with a path overhead that does a parity check on the tributary payload.

Measurements can be made with a portable test set equipped with framing and demultiplex capability in its receiver. Alternatively, the measurements can be made by a network monitoring system; in some cases, the detectors can be built into the operational equipment itself, in which case the network monitoring system needs to check status indicators and alarms. Because these are nonintrusive measurements, they should be made at a protected monitor point or by using a resistive bridging

probe at an unprotected  $75\Omega$  T-junction, as shown in Figure 7.32. The resistive probe results in a loss of 20 to 30 dB, so additional receiver gain is necessary in the test set. In North America, nonintrusive tests usually are made at a protected cross-connect point, which additionally requires equalization of the  $\sqrt{f}$ -law cable losses.

This section will consider what in-service measurements are possible on traditional PDH systems designed according to the European standards and North American standards. Some basic in-service errors, such as line code errors, are useful for checking the performance of a particular transmission link, while others may provide a quality measure over a complete end-to-end transmission path. The major benefit of in-service tests is that they allow the user's traffic to flow normally without interruption.

This means that error performance statistics can be collected over a longer period, and with the storage available on modern test sets, weeks of data can be stored and timestamped for multiple in-service parameters. These might include CRC-4 block errors, FAS errors, HDB3 code errors, and alarm history. Long-term monitoring is useful for catching that elusive burst of errors that only seems to occur at the busiest time of the day! It also helps to confirm that the overall quality of the circuit meets



**Figure 7.32** In-service tests require nonintrusive bridging of the active transmission path. Usually this is available at a “protected point.” Alternatively, a high-impedance bridging probe can be used, resulting in a signal loss of 20 to 30 dB made up for by amplification in the test set. Useful in-service tests can be made only by a test receiver capable of recognizing the hierarchical frame signals and checking for errored bits in the frame word. More sophisticated test sets can demultiplex low-level tributaries, or even 64 kbps channels, from a 140 Mbps high-capacity link.

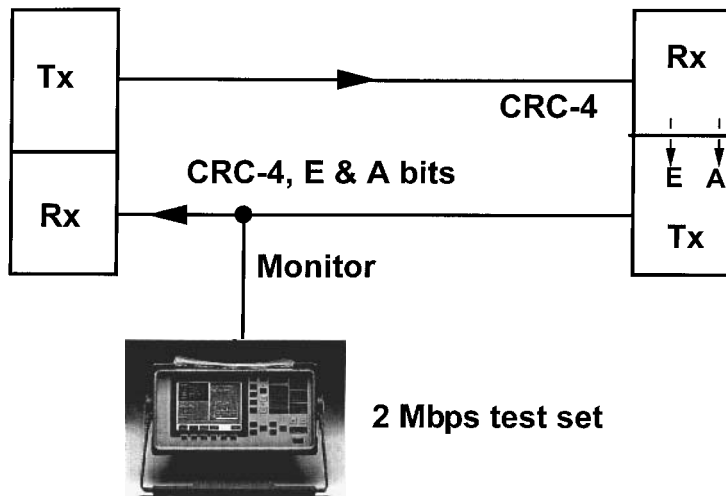
specification. For more information on technical aspects of in-service measurement parameters and performance analysis, please refer to Chapter 26, Section 26.3.3.

### 7.5.2 European in-service testing

**In-service testing at the 2 Mbps (E1) rate.** The importance of in-service tests at the 2 Mbps (E1) level was discussed earlier, and especially the value of CRC block error detection (and E-bits) for estimating errored seconds. The E1 level is the basic building block of the switched telecom network, and also is the most commonly used rate for digital leased lines in private enterprise networks. The CRC-4 error-detection process checks all the payload bits, whether they are PCM voice channels, compressed voice encoding, video, or data. The Far End Block Error (E-bit) allows complete analysis of both transmission directions from a single, nonintrusive monitoring point (Figure 7.33).

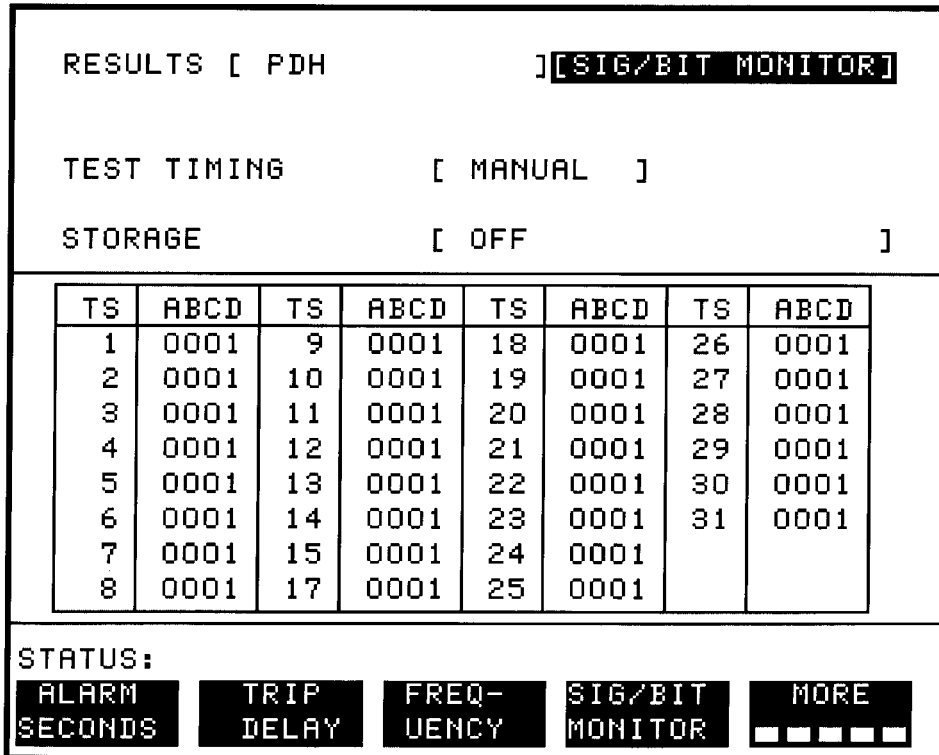
As discussed earlier, this in-service error detection process does not indicate bit error ratio (BER) unless one assumes a certain error distribution (random or burst) to predict the average errors per block. Rather, it provides a block error measurement. This is very useful for estimating *percentage errored seconds* (%ES), which usually is considered the best indication of quality for data transmission—itself a block transmission process. CRC-4 error-checking is very reliable; at least 94 percent of errored blocks are detected even under high BER conditions, according to ITU-T Recommendation G.706.

The E1 test set must be able to decode a CRC-4 frame and analyze and store the measurement results. These are divided into Anomaly Events (AE) such as frame or



**Indicates performance of both go and return paths**

**Figure 7.33** The CRC-4 frame structure at 2 Mbps (E1) provides complete in-service error checking of the traffic payload. Errors and alarms detected at the far end receiving terminal are relayed back to the transmitting end using the E and A bits. A test set monitoring in-service on either the transmit or receive paths thus will have the complete picture of both directions of transmission.



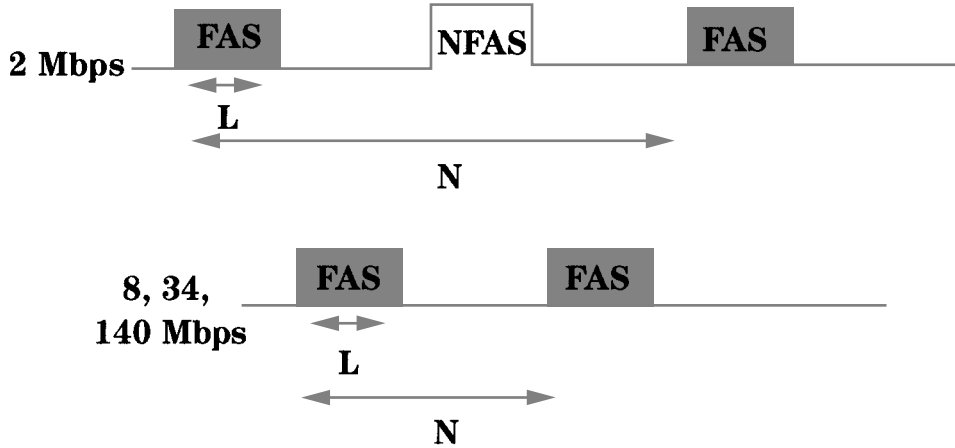
**Figure 7.34** A display of the 4-bit ABCD signaling words carried in timeslot 16 for channel-associated signaling.

CRC errors, and Defect Events (DE) such as loss of signal, loss of frame synchronization, etc., which will set the alarm A-bit in the backward direction (Figure 7.33). Normally these error events, called *performance primitives*, would be accumulated and analyzed statistically according to the error performance standards ITU-T M.2100 and G.826. (See Chapter 27, section 27.3.4.)

With a framed 2 Mbps set, other useful in-service checks can be made. One can monitor the individual 64 kbps channels to check if they are carrying voice or data. A demultiplexed channel at 64 kbps can be fed to a protocol analyzer or decoded to provide a voice-frequency output.

TS16 usually is assigned to signaling. It is possible to demultiplex TS16 and display the 30 ABCD words for channel-associated signaling (Figure 7.34) to investigate permanently idle channels or “stuck bits.” Before taking a channel out of service, one can check that it is idle and so avoid unnecessary outage.

One further application of the framed 2 Mbps test set capable of “through-data” mode is to drop and insert single or multiple 64 kbps test channels while the remaining channels carry revenue-earning traffic. This requires the tester to be placed in circuit with the 2 Mbps line; having done that, a much more detailed analysis is possible on one or more 64 kbps channels while still providing a partial service.



Bit Rate (Mbps)	L (bits)	N (bits)	FAS Bit Rate (kbps)
2.048 (E1)	7	512	28
8.448 (E2)	10	848	99.6
34.368 (E3)	10	1536	223.7
139.264 (E4)	12	2928	570.7

**FAS = Frame alignment signal**  
**NFAS = Non Frame alignment signal (2 Mbps only)**

**Figure 7.35** A table showing the relationship between the frame word length and the length of the PDH frame at different hierarchical rates, enabling the equivalent FAS bit rate to be calculated. This is the notional rate of the in-service channel available for error rate checking; it represents only a tiny fraction of the overall bit stream, however, so the important payload area remains unchecked.

**In-service testing at higher PDH rates (E2 to E4).** While the CRC-4 frame provides very good in-service performance at 2 Mbps, the capability at the higher rates of 8, 34, and 140 Mbps is much less satisfactory. The only guaranteed fixed pattern is the frame alignment signal (FAS) at the beginning of each frame. This provides a tiny snapshot of the overall performance, and makes a big assumption that errors in the 10- or 12-bit frame alignment word are representative of the remaining payload bits, which could number up to 2900 in a 140 Mbps frame!

Over a long measurement period, the errors in the frame bits probably give a reasonable approximation to the average bit error ratio (BER) when the errors are evenly distributed according to a Poisson distribution; the prediction becomes very unreliable in the presence of burst errors, however, and is a poor indication of block errors such as %ES. Nevertheless, it remains the best way of in-service testing PDH transmission systems unless these signals are carried in the virtual container of SDH transmission.

As already described, at each rate the FAS is a fixed sequence of L-bits (see Figure 7.35), which repeats every N transmitted bits. By taking the ratio of L/N relative to



the bit rate, it is possible to calculate the FAS bit rate at each level. This is the rate at which this fixed, known sequence is transmitted within the overall bit stream and is the bandwidth available for in-service error testing.

An error detector capable of recognizing these frame words in the random traffic signal can detect Anomaly Events as frame bit errors. It can also detect alarms and *remote alarm indication* (RAI) bits as Defect Events. A test set with full demultiplexing capability can extract a 2 Mbps tributary from a 140 Mbps stream and make a CRC-4 error analysis as described above. It also can display the full alarm status for the 140 Mbps composite signal, as shown in the example in Figure 7.36.

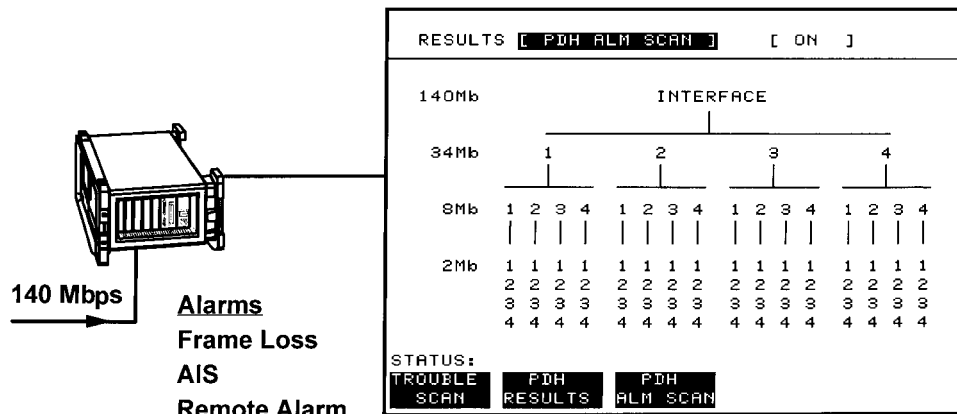
The various frame and alarm bits available for checking in-service in the PDH E1–E4 hierarchical rates are summarized in Figure 7.37.

### 7.5.3 North American hierarchy

**Maintenance layers and performance primitives.** To review the various in-service measurements available, it is first necessary to consider the layered maintenance model shown in Figure 7.38. In this model (fully described in Bellcore FR-NWT-000475, “Network Maintenance: Transport Surveillance”), two types of maintenance are defined:

- Path layer
- Line or section layer

**Path layer.** An end-to-end digital service such as DS1 or DS3 is defined as a path and may traverse many different digital sections at different multiplex bit rates and using different technologies (such as lightwave or microwave). Path-layer monitoring therefore gives an indication of the overall DS1 or DS3 service performance being provided to customers.



**Figure 7.36** With a demultiplexing test set, the complete alarm picture can be scanned and displayed for all hierarchy levels by monitoring the 140 Mbps signal. On this test set, an alarm would show as reverse video on a particular number.

170 Wide Area Networks

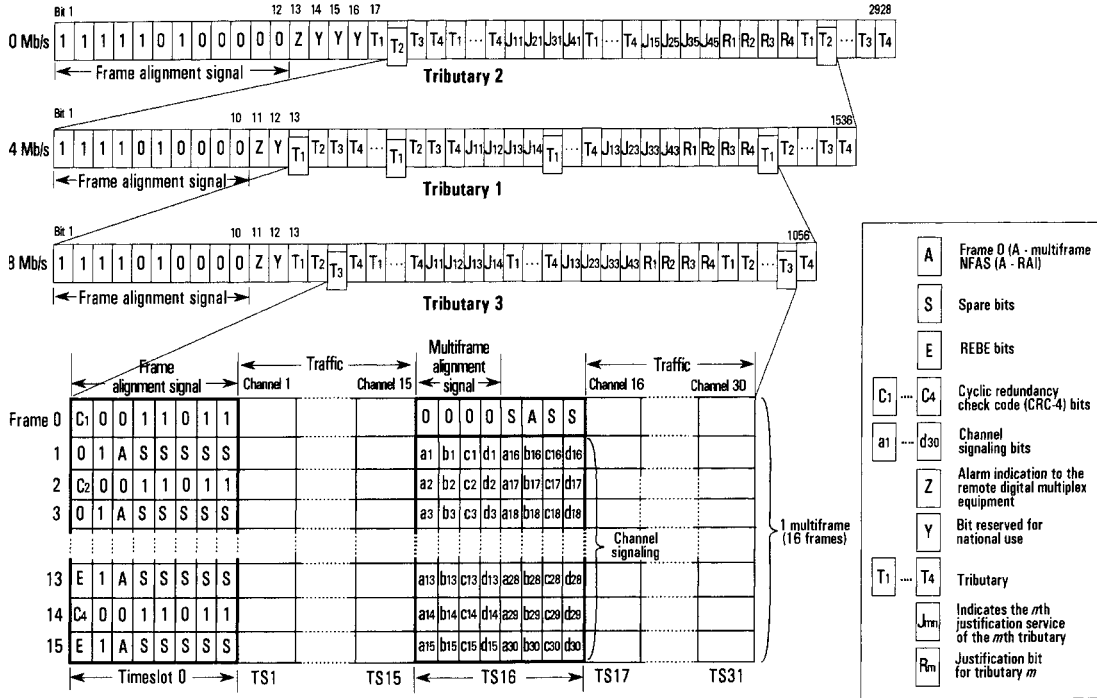
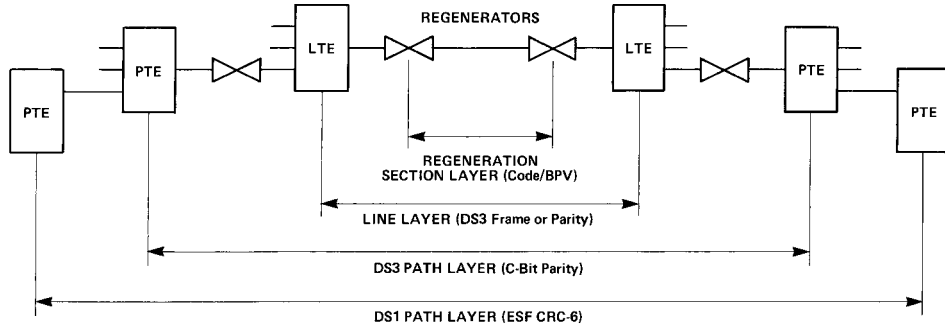


Figure 7.37 A summary of the signal structures for the four hierarchy levels of the international PDH standard.



LTE: LINE TERMINATING EQUIPMENT  
 PTE: PATH TERMINATING EQUIPMENT

Figure 7.38 Layered maintenance model used in North America, showing the Path layer and Line or Section layer. The Path layer provides an indication of error performance for the end-to-end service provided to the customer at either DS1 or DS3, while the Section layer indicates the performance of the previous maintenance section only, and thus is useful for troubleshooting and fault location.

**Line or section layer.** Monitoring at the line or section layer provides maintenance information on a facility in the network and is helpful in sectionalizing problems. Degradation detected at this level might contribute only part of the result for overall

path layer performance, however; within the transmission network there are typically many lines and regenerator sections between line terminating equipment (LTE). At the line layer, for example, there might be a 135 Mbps microwave radio link carrying three DS3 streams with several hops (or regenerative repeater sections) en route.

Performance primitives are basic error events or other performance-related occurrences detected by monitoring the frame format code or the line code of a digital signal. These performance primitives are grouped into categories of anomalies and defects. *Anomalies* generally are degradations in performance, whereas a *defect* is a limited interruption in the ability of a system to perform a required function. Examples of defect primitives are loss of signal (LOS), out-of-frame (OOF), severely errored frame (SEF), and alarm indication signal (AIS).

Examples of performance primitives are:

- Line code violations
- Bipolar violations (BPV) in B3ZS, B8ZS, or AMI
- Frame errors
- Parity errors
- C-bit parity (DS3)
- Extended Superframe (ESF) Cyclic Redundancy Checksum (CRC-6) (DS1)

Some performance primitives, such as BPV and parity errors, are corrected by regenerators or line terminal equipment before the signal is passed on to the next section. Errors detected by these primitives can have occurred only in the previous line section and not elsewhere in the network. They are useful for sectionalizing problems but cannot be used for assessing overall service performance.

To measure overall path layer performance, it is necessary to use an error-detection process that will pass through the various network hierarchy and technology sections transparently. Two such primitives have been devised: ESF CRC-6 and C-bit parity.

**CRC-6.** For DS1 services, the preferred frame format is Extended Superframe (ESF), consisting of 24 standard DS1 frames (a total of  $24 \times 193 = 4632$  bits). A Cyclic Redundancy Checksum (CRC) is computed over each superframe and the 6-bit CRC word inserted in the next superframe. The CRC-6 calculation is made only in the path terminating equipment (PTE) and is not recalculated at the line or section level; errors will accumulate along the path in exactly the same way as bit errors in the payload or customer data. Recalculating CRC-6 at the receiving PTE and comparing it with the value sent from the transmitter will detect an error occurrence in the path.

**C-Bit parity.** For end-to-end DS3 performance, a new path layer measurement called *C-bit parity* has been introduced.

Traditionally, DS3 in-service error performance has relied on conventional DS3 parity bits (P-bits) and bipolar code violations. Typically these are recalculated or corrected at each item of terminal equipment and do not allow a cumulative system measurement to be made over the whole path. On the other hand, C-bit parity provides such a measurement and, as described later, has a number of other very useful features for network monitoring.

**DS1 ESF CRC-6 and performance criteria.** The DS1 frame consists of one framing bit (F-bit) at the beginning, followed by 192 bits ( $24 \times 8$ ) of traffic or payload. Frames are assembled into a superframe and, in the earlier D4 or SF framing standard, the superframe consists of 12 frames. In Extended Superframe Format (ESF), 24 frames are used to form a superframe containing a total of 4632 bits. (One second contains about 333 superframes.) Of the 24 F-bits, 6 are used for the framing pattern sequence, 6 for the CRC word, and 12 for a 4 kbps data link. This is shown in Table 7.3.

**TABLE 7.3 Extended Superframe Format.**

FRAME NO	F BITS			
	BIT NO	FPS	CRC	DL
1	0	-	-	X
2	193	-	C1	-
3	386	-	-	X
4	579	0	-	-
5	772	-	-	X
6	965	-	C2	-
7	1158	-	-	X
8	1351	0	-	-
9	1544	-	-	X
10	1737	-	C3	-
11	1930	-	-	X
12	2123	1	-	-
13	2316	-	-	X
14	2509	-	C4	-
15	2702	-	-	X
16	2895	0	-	-
17	3088	-	-	X
18	3281	-	C5	-
19	3474	-	-	X
20	3667	1	-	-
21	3860	-	-	X
22	4053	-	C6	-
23	4246	-	-	X
24	4439	1	-	-

Notes:

Frame 1 is transmitted first

FPS – Framing Pattern Sequence  
(...001011...)

CRC – Cyclic Redundancy Check channel  
(bits C1–C6)

DL – 4-kbps Data Link; X indicates bit  
assigned to DL

The CRC word is calculated on all the payload bits of the preceding superframe. This is important, since an error in any of these bits will cause a change in the CRC calculated at the receiving end and the error will be detected upon comparison with the transmitted CRC word. The detection probability is very high: 99 percent for error rates less than  $10^{-3}$ . A detailed description of ESF can be found in ANSI T1.107.

One important factor to remember about CRC-6 is that it does not indicate the number of errors in a superframe; only that one or more has occurred. Therefore it truly is a block error measurement; it is not possible to estimate bit error ratio (BER) unless one makes assumptions about the distribution of the errors. CRC-6 is, however, an excellent way of estimating error seconds and the corresponding percentage of error-free seconds and availability—the basis of “tariffed” performance.

**Using error-detection results.** Having detected CRC error events or severely errored framing events (SEF) such as Out of Frame (OOF) or Change of Frame Alignment (COFA), how are the results used to set performance criteria?

The requirements for DS1 performance monitoring are given in an ANSI standard for In-Service Digital Transmission Performance Monitoring (ANSI T1.231 – 1993). This document defines some path performance parameters based on CRC and SEF:

- *Errored Seconds* (ES) containing one or more CRC events or one or more SEF events.
- *Errored Seconds, Type A* (ESA) containing only one CRC event.
- *Errored Seconds, Type B* (ESB) containing 2–319 CRC events and no SEF events.
- *Severely Errored Seconds* (SES) containing 320 or more CRC events or an SEF event. This is intended to approximate the SES definition in ITU-T G.821 for a BER of  $10^{-3}$ .

*Consecutive SES* (CSES) is a period of 3–9 consecutive SES. *Unavailable Seconds* (UAS) occur after 10 or more consecutive SES. The definition is similar to the ITU-T G.821 availability criteria.

*Degraded Minutes* (DM) are defined in G.821 as 1-min periods exceeding BER of  $10^{-6}$  after subtracting SES. ANSI T1.231 indicates that this parameter is for further study. It is not clear how the  $10^{-6}$  BER would be interpreted from the block error rate provided by CRC-6.

All of these measurements can readily be made on the incoming data stream, either by the terminal equipment itself or by a test instrument. In order to monitor the performance of the transmit path from the near end, measurement data needs to be sent back from the far end. This is the purpose of the 4 kbps ESF data link mentioned earlier. The data link bits can be used in a bit-oriented mode for higher priority alarm and control information, or in a message-oriented mode for sending the longer-term path performance parameters listed above. Details of these messages can be found in ANSI T1.403, “Network-to-Customer Installation—DSI Metallic Interface.”

Performance criteria normally are presented as a percentage of time, so the monitoring unit or external test equipment would accumulate the performance parameters listed earlier over a period of time (e.g., an hour, a day, a week, etc.) and express the result as a percentage. ANSI T1-231 recommends 15-min and 1-day accumulation

periods. The acceptable performance threshold would depend on the grade of service and the length of the route, as described in Chapter 27, Section 27.3.3.

**DS3 C-Bit parity and performance parameters.** In-service error detection and performance monitoring of DS3 traditionally uses either Bipolar Violations in the B3ZS interface code or the conventional parity check on the DS3 frame. The drawback with these two measurements is that errors in parity and interface code normally are corrected at intermediate points in the network, so they are not readily usable for end-to-end path monitoring, only for the previous span.

Because there are very few frame bits in the DS3 frame, however, frame errors give an unreliable indication of overall payload performance. Conventional DS3 framing is designated *M23* in the standards.

A new framing application has been defined, called *C-bit parity*. Initially proposed by AT&T, the new DS3 frame structure is defined in ANSI T1.107.

The idea behind C-bit parity is to provide a parity measurement that will pass through existing DS3 transmission equipment transparently, so that error performance data will accumulate along the DS3 path. The C-bits are not reset at intermediate points.

In the C-bit parity application, the C-bits no longer are used for stuffing control; asynchronous DS2 tributary operation is not allowed. When carrying direct DS3 services or 28 DS1s, the C-bits are free and can be used for other purposes.

Table 7.4 shows the disposition of frame overhead bits in the DS3 M-frame, which comprises 7 subframes, each of 680 bits. Each square in the matrix of Table 7.4 represents the designated frame overhead bit followed by 84 payload bits. The C-bits are carried in columns 3, 5, and 7. The remaining columns are the same as for M23 applications. The two P-bits in subframes 3 and 4 provide the conventional parity check on the payload bits (set to 11 for even and 00 for odd parity).

**TABLE 7.4 C-bit Parity Framing at DS3.**

SUBFRAME	BLOCK#							
	1	2	3	4	5	6	7	8
1	X	F1	AIC	F0	N <sub>r</sub>	F0	FEAC	F1
2	X	F1	DL <sub>a</sub>	F0	DL <sub>a</sub>	F0	DL <sub>a</sub>	F1
3	P	F1	CP	F0	CP	F0	CP	F1
4	P	F1	FEBE	F0	FEBE	F0	FEBE	F1
5	M0	F1	DL <sub>t</sub>	F0	DL <sub>t</sub>	F0	DL <sub>t</sub>	F1
6	M1	F1	DL <sub>1</sub>	F0	DL <sub>1</sub>	F0	DL <sub>1</sub>	F1
7	M0	F1	DL <sub>a</sub>	F0	DL <sub>a</sub>	F0	DL <sub>a</sub>	F1

Notes:

Total M-Frame = 4760 bits

Each Subframe = 680 bits

Each of the 8 blocks in each subframe = 85 bits

(1 bit frame overhead + 84 bits payload)

The two X-bits in subframes 1 and 2 normally are set to 11, but change to 00 for indicating a real-time alarm to the far end in case of AIS (Alarm Indication Signal) or loss of frame.

The first C-bit in subframe 1 is the Application Indication Channel (AIC). This bit is set permanently to 1 for the C-bit parity application. (It is random for M23 application due to stuffing.) The second C-bit in subframe 1 is reserved for network use, and the third is the Far End Alarm and Control channel (FEAC). The sequence of FEAC bits can be used for loopback control or for sending back to the near end a menu of alarm information.

The CP-bits carried in subframe 3 indicate odd or even parity in the same way as the P-bits. The FEBE bits in subframe 4 indicate a C-bit parity or loss-of-frame event from the far-end PTE. In this way, both directions of transmission can be monitored from the near end.

The remaining 12 C-bits are used for data links and, if unused, are set to 1. Of course, normal P-bit parity and frame-error detection are still available for monitoring the span.

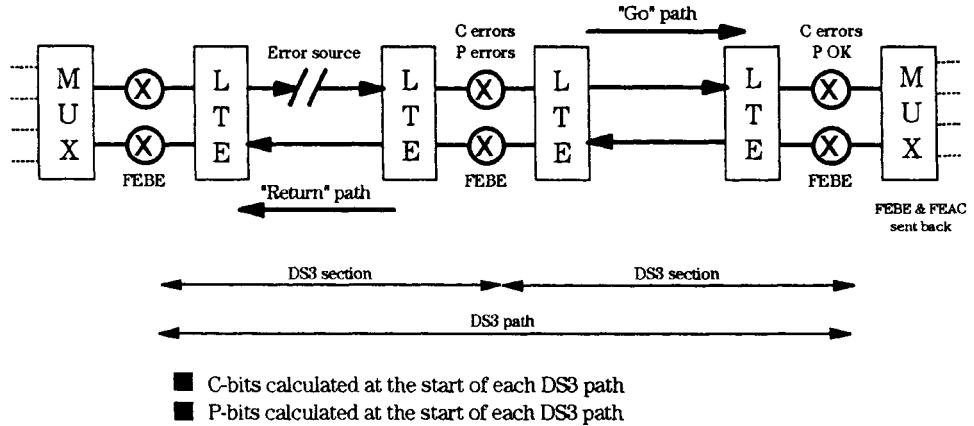
As with the CRC-6 at DS1, the parity error detection process is a block-error measurement made on the payload bits in the DS3 frame. Unless the error distribution is known, it is difficult to calculate the equivalent BER. One drawback with parity is that multiple errors in the frame could cancel out the parity calculation and indicate error-free performance. This underestimate is most noticeable at high-error rates.

The performance parameters for DS3 services are defined in ANSI T1.231. C-bit parity and frame errors are accumulated in the following categories:

- *OOF Seconds*, containing an out-of-frame or AIS event.
- *Errored Seconds Type A (ESA)*, containing one parity error but no OOF, SEF, or AIS event.
- *Errored Seconds Type B (ESB)*, containing 2–44 parity errors but no OOF, SEF, or AIS events.
- *Severely Errored Seconds (SES)*, containing more than 44 parity errors or an OOF, SEF, or AIS event.

Exactly the same parameters can be derived from the FEBE bits for the near- to far-end path performance. As with DS1 performance objectives, DS3 error events are accumulated over a period of time and expressed as a percentage.

**C-bit Parity: end-to-end bidirectional monitoring.** Figure 7.39 shows a typical testing scenario of a DS3 service. After the signal has had the parity check done at LTE 1, an error occurs. LTE 3 counts errors and sends an FEBE upstream. Since this is C-bit framing, the CP errors are passed downstream, but the P-bit parity errors are removed. LTE-C recalculates the P-bit parity and sends it downstream with the CP parity incorrect. The far end mux then sends an FEBE upstream, indicating a CP parity error.



LTE: Line Terminating eqpt.      MUX: DS1/DS3 multiplex  
 FEBE: Far End Block Error      FEAC: Far End Alarm and Control  
 Channel

**Figure 7.39** By utilizing the network overhead bits in a DS3, C-Bit frame error isolation can be performed in-service. The point where both the C-bit and P-bit errors occur indicates the faulty leg.

## 7.6 References

- "PCM and Digital Transmission Systems" by Frank F.E. Owen, McGraw-Hill 1988, ISBN 0-07-047954-2.
- "Testing Sub-rate Data Services," Hewlett-Packard Application Note 1211-3 (Publication number 5091-2072E).

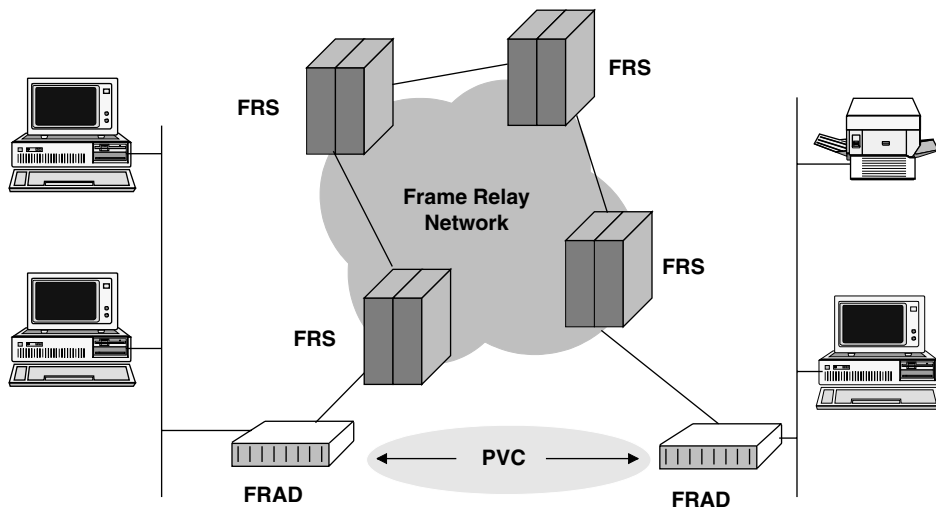


# Frame Relay

**Bill Risley***Hewlett-Packard Co., Colorado Springs, Colorado*

## 8.1 Introduction

Frame relay is a highly efficient technology for routing data over wide area networks (WANs). Users from all segments of industry implement frame relay to improve performance, cut cost, simplify their networks, and improve reliability. Typical frame relay applications link remote local area networks (LANs) to central databases and other resources (Figure 8.1). Both ANSI and the ITU-T have established recommendations



**Figure 8.1** The typical application for frame relay involves internetworking local area networks. To accomplish this, frame relay makes use of frame relay access devices (FRADs) and a set of dedicated frame relay switches (FRS). In most applications, frame relay uses permanent virtual circuits (PVCs).

for frame relay, and the Frame Relay Forum has ratified a number of Implementation Agreements that have been widely accepted.

### 8.1.1 Frames

Frame relay is one of several Data Link layer (OSI layer 2) protocols, which are based on the High-Level Data Link Control (HDLC) bit-oriented protocol. These protocols are employed to organize user data in such a way that information of any type can be transferred error-free through the network.

To do this, the user data are encapsulated between header octets (which provide addressing and control information) and error-checking octets. This procession of octets—the address, control field, user data, and error-check sequence—is called a *frame*. Each frame is separated from other frames by the use of a special *idle* character. When no frames are being transmitted on the link, this idle character is transmitted continuously to indicate the idle condition. The appearance of a character other than an idle character normally signals the start of a frame. When idle characters again appear, it indicates the end of a frame.

### 8.1.2 Frames relayed

*Frame relay* is a term that arises out of the *Fast Packet Concept*. This concept embraces the integration of voice, video, and data, and is the conceptual framework for the high-speed digital communications that emerged in the 1990s. In this concept there are two chief ways of transferring, or relaying, information in networks:

1. *Frames*, as defined above. These may be of sufficient length to contain the entire user message.
2. *Cells*, which are of fixed length. These often require the user information to be fragmented for transfer and reassembled at the destination.

The term *frame relay* thus refers to the passing of link-layer frames from one point to another in an information network; it is in contrast to *cell relay*, which refers to the transfer of fixed-length messages and is the basis for communications methods such as Asynchronous Transfer Mode (ATM), discussed elsewhere in this book.

Since frame relay usually transfers user information intact, the process of fragmentation and reassembly can be avoided. This reduces processing requirements and makes it easier for frame relay to be adapted to existing network architectures and infrastructures. The simple format of the frame relay frames also means that the overhead can be very low, resulting in a very efficient use of the available bandwidth. The simplicity and efficiency of frame relay leads to easy, inexpensive implementation, and is the key to frame relay's rapid deployment and widespread use.

### 8.1.3 Overview of frame relay

Frame relay offers a high-performance alternative to leased-line bridging. Its implementation differs from that of leased lines in significant ways. Rather than relying on dedicated facilities to connect users, frame relay calls for:

- A switching backbone to connect sites.
- A unique link-layer protocol to transfer and route frames.
- Permanent virtual circuits (PVCs), sometimes called *permanent logical links* (PLLs).

Note that the term *virtual circuit* is used to indicate that a method other than a physical electrical path is used to connect two points.

The frame relay standards define three layers of functionality (although, unlike X.25, the third layer is used for signaling only). These layers correspond to the Physical and part of the Data Link layers of the OSI model. At the Physical layer, frame relay supports any interface recognized by ANSI or the ITU-T. (Specific interfaces used are discussed later.) The Data Link layer defines relay, switching, and congestion notification services. In addition, frame relay defines virtual circuit routing at the Data Link layer (rather than the Network layer, where such services usually are implemented).

Frame relay does not provide the error and flow control usually associated with link-layer protocols. It assumes high-quality, low-noise links and leaves error control to the users' Transport layer functions. The frame relay frame does include a frame check sequence field for error detection. When errors are detected, however, the frame is merely discarded. It is left to other protocols at the would-be receiver to discover the loss and notify the sender for retransmission.

### 8.1.4 Physical layer interfaces

The interface provided to the user is referred to as the *User-Network Interface* (UNI). Interfaces between service providers (internal to the network, as far as the user is concerned) are referred to as *Network-Network Interfaces* (NNIs).

As noted previously, frame relay supports the various physical interfaces recognized by ANSI and the ITU-T:

- *CEPT-E1* refers to circuits that conform to the ITU-T Recommendations G.703 and G.704 for a primary rate operating at 2.048 Mbps.
- *T1* refers to circuits that conform to the same ITU-T recommendations and to ANSI T1.403, operating at 1.544 Mbps.

In Europe and other countries that conform to the CEPT-E1 primary rate recommendations, frame relay commonly is delivered to the subscriber through a *Channel Service Unit* (CSU), which terminates the circuit from the *Central Office* (CO) and conditions the signals. Access by the user can be through CEPT-E1 directly, or through

V-series connections (such as V.35 or V.11) if an intermediary *Data Service Unit* is used. (The DSU converts the primary rate signals.) In North America and Japan, frame relay often is delivered over T1 to a combination DSU/CSU; access by the user is through V.35. Two other connection types used in North America are DDS and ISDN.

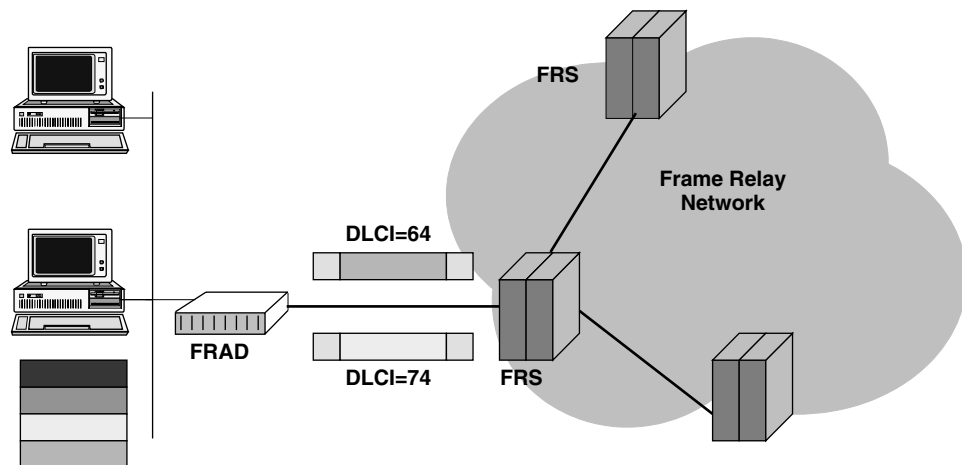
### 8.1.5 Data Link Layer Interface

The *Data Link Layer Interface* (DLLI) is the frame relay link layer. Figure 8.2 shows the basic structure of a frame relay frame. Each data segment is encapsulated by a header and a frame check sequence.

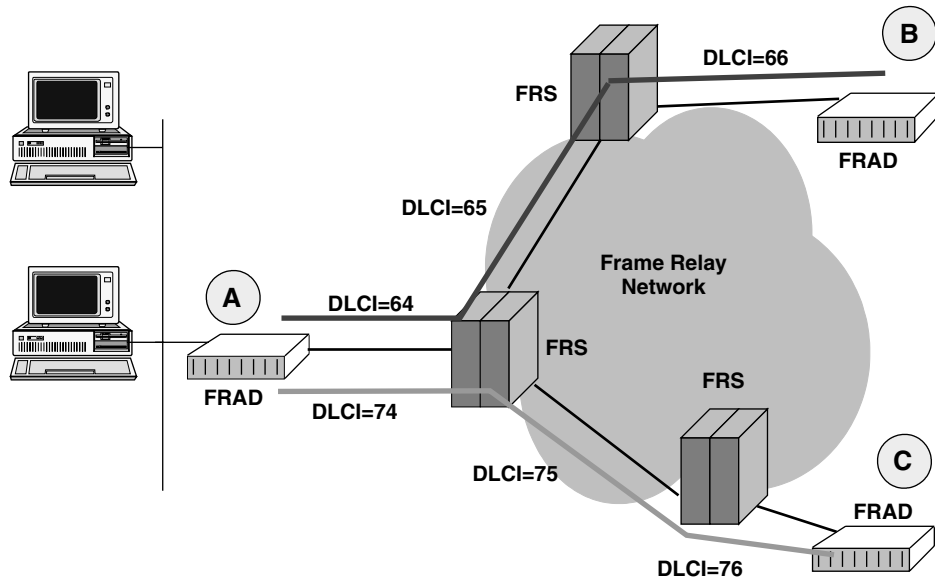
The header contains a *Data Link Connection Identifier* (DLCI) to identify the *permanent virtual circuit* (PVC), which is the logical connection to a distant site. This means that frames routed to one location contain a DLCI associated with that destination. It is the function of the frame relay switches in the network to use the DLCI to route the traffic to the configured destination.

Since frame relay permits the use of multiple PVCs on a single physical connection, multiple PVCs may be configured. When the network must send a frame to one location, that frame is launched into the network with the appropriate DLCI. To reach a different destination, a different DLCI is used. In this way frames going to different sites can be multiplexed onto the same physical link.

In fact, statistical multiplexing is an inherent feature in frame relay because of this characteristic of launching frames to different destinations on an as-needed basis. This also gives frame relay the characteristic of providing *bandwidth on demand*. That is, when a given application must send many frames to a destination, more of the available bandwidth on the link will be consumed by the application at that time.



**Figure 8.2** As each frame is launched into the frame relay network, it is assigned a Data Link Connection Identifier (DLCI) that associates it with each PVC. Frames sent on one PVC have a different DLCI than frames sent on another PVC. By using different DLCIs, many different PVCs can be accommodated on a single physical connection. Because the different circuits can be used as needed, frame relay can allocate the access-channel bandwidth dynamically and on demand.



**Figure 8.3** To construct a PVC, each switching node must map information arriving on one port to a different port. Each of these ports will use a different DLCI to designate the frame's destination. This means that many different DLCIs are required to identify a PVC.

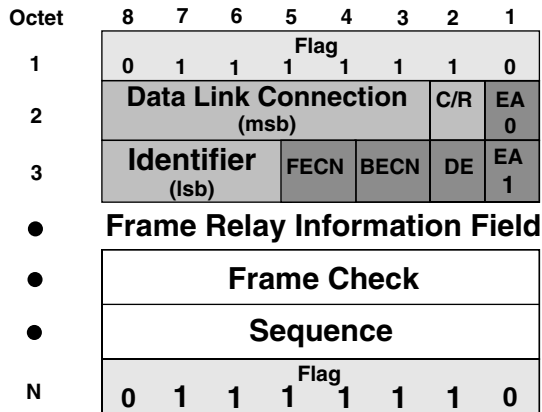
Each PVC usually is identified by a different DLCI at each point in the network (Figure 8.3). At one end, DLCI  $x$  is used to identify the PVC. Between switching points, DLCI  $y$  might be used, and at the final segment connecting to the destination router, DLCI  $z$  might be used. This means that the DLCI for a given PVC has only local significance; a different DLCI is assigned to a transmission by each router or switching point. As data segments are routed through the network, they are assigned one DLCI after another. At the assigning router, the DLCI indicates through which exit port a frame should be sent. At the receiving router, the DLCI is used to find the new, corresponding DLCI in a switching table.

(Note that the DLCI should not be confused with network addressing, such as a call reference or logical channel number.)

### 8.1.6 The frame relay frame

Figure 8.4 shows the frame relay frame structure and header.

- **Flag** As in HDLC, frames start and end with at least one flag.
- **DLCI** The DLCI is a number between 0 and 1023 that is used to identify the PVC being used.
- **C/R** The Command/Response field is not used by the network, but is passed transparently to the user.
- **EA** The header contains two Extended Address bits. The first is 0 and the second is 1. Frame relay allows the header to be extended so that the range of DLCI



**Figure 8.4** The frame consists of a frame relay header, the user data, and the frame check sequence (FRS). Each frame is separated from its predecessor and successor by at least one idle character, or flag.

values can be increased. As long as the EA bit is 0, it indicates that there is another header byte to follow.

- **FECN and BECN** The Forward and Backward Explicit Congestion Notification bit fields are explained in section 8.1.7.
- **DE** The Discard Eligibility bit may be set by the user to indicate that a frame has lower priority than others. If network resources become overloaded, these frames are the first to be dropped by the network.

**8.1.7 Network congestion**

Congestion occurs when the network runs out of capacity to pass all data submitted. Figure 8.5 shows the uses of forward and backward congestion notification.

When congestion occurs, the FECN bit is set to 1 by a routing node to warn downstream devices that the path is congested and that ensuing frames might be delayed or discarded. This bit warns the receiver of the congestion, not the sender.

The receiver then can set the BECN bit on frames that it is returning to the sender to alert the sender to the congestion. The sender then can choose to send frames more slowly or to halt transmission and resume at another time. This option works only if the exchange is full- or half-duplex. If data are being sent in one direction only (simplex), the sender will remain unnoticed of the congestion and will be unable to counteract it.

**8.1.8 In-channel signaling procedures**

In frame relay, local in-channel signaling messages provide information about:

- PVC availability and status
- Link integrity
- Error occurrence in received data

These signaling messages are layer 3 messages that are structured like Q.931 (ISDN) messages and are transferred using Q.922 (layer 2 frame relay) frames that have a DLCI reserved for signaling. The signaling procedures also:

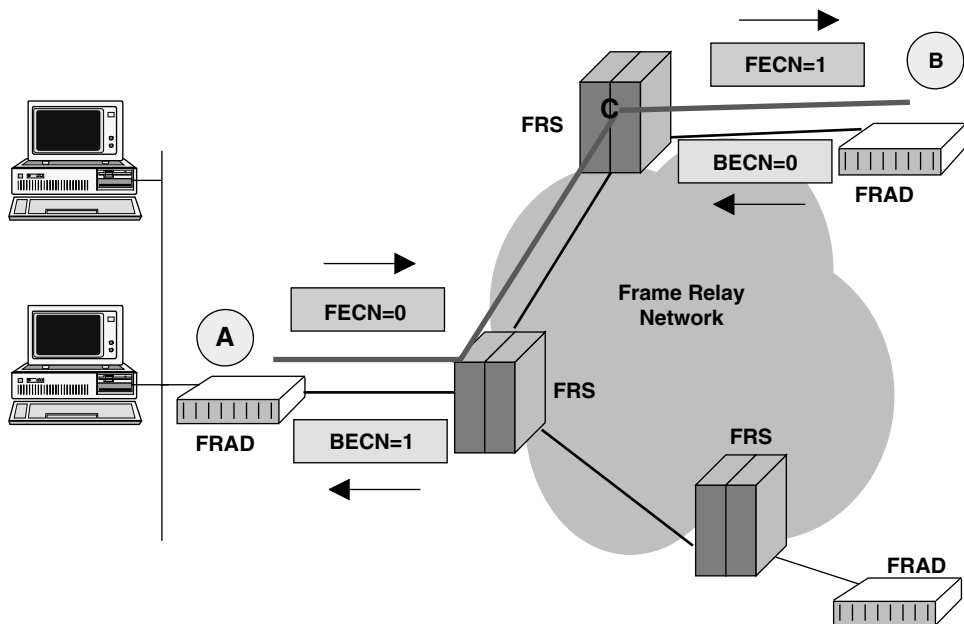
- Support the addition and deletion of Permanent Virtual Circuits (PVCs).
- Report on the active or inactive state of DLCIs.
- Indicate link reliability and protocol errors.

The principal feature of these procedures is the periodic issuing of Status Enquiry messages from the user and the required Status Message Response from the network (Figure 8.6).

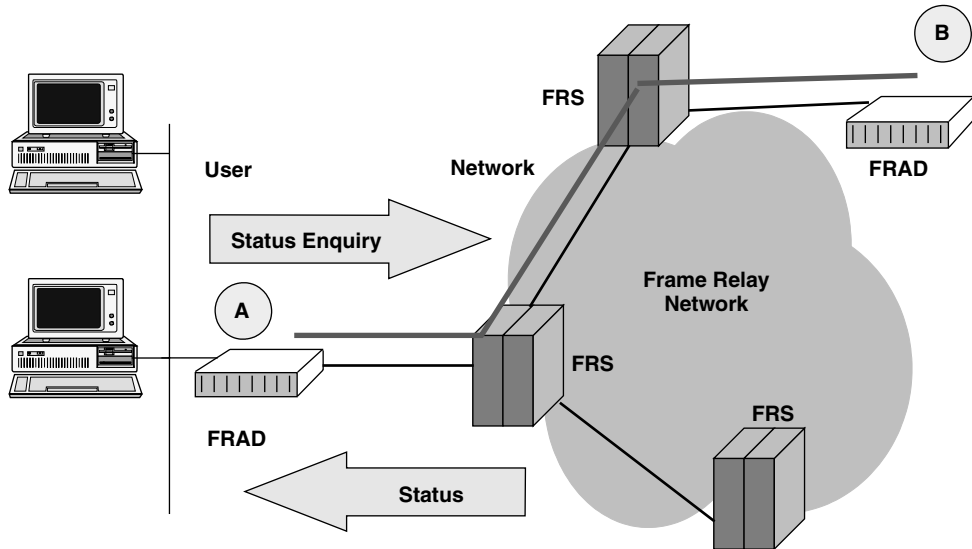
(Note that since frame relay began, these messages have been referred to as *Local Management Interface* (LMI). This term is still used, although *Local In-Channel Signaling* is the term now used in specifications.)

### 8.1.9 Link Integrity Verification (LIV)

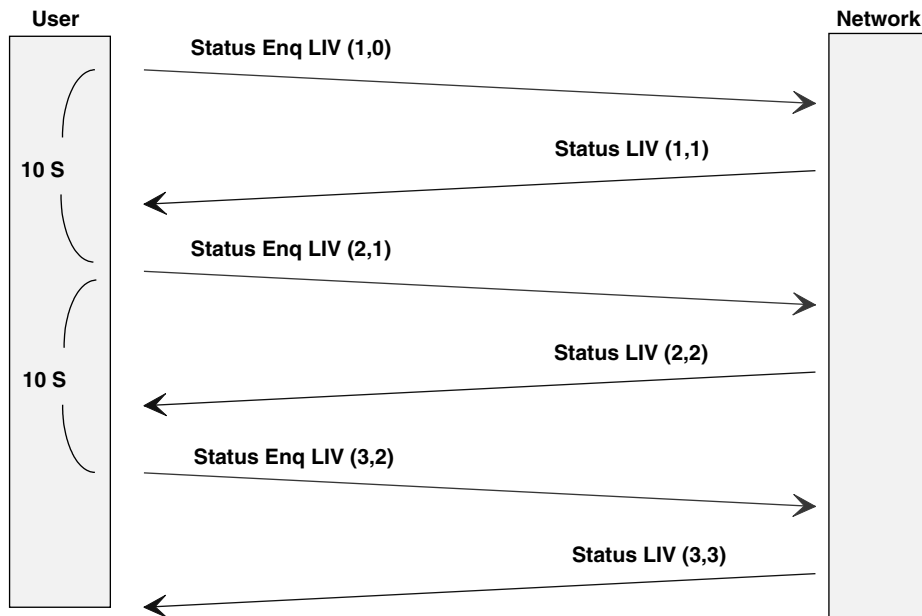
LIV is a procedure by which the user and the network periodically exchange polling sequence numbers to determine the integrity of the link. Examples of polling sequence numbers are indicated in Figure 8.7 and Figure 8.8 in the parentheses of the LIV and Full Status Report (FSR) messages. The user sends out a Status Enquiry message containing a send polling sequence number and the value of the last received polling sequence number.



**Figure 8.5** Because frame relay allows the access bandwidth to be used on demand, it is possible for demand to exceed the available bandwidth. When this occurs, frame relay switches are able to notify the user that congestion has occurred and that data might have been lost.

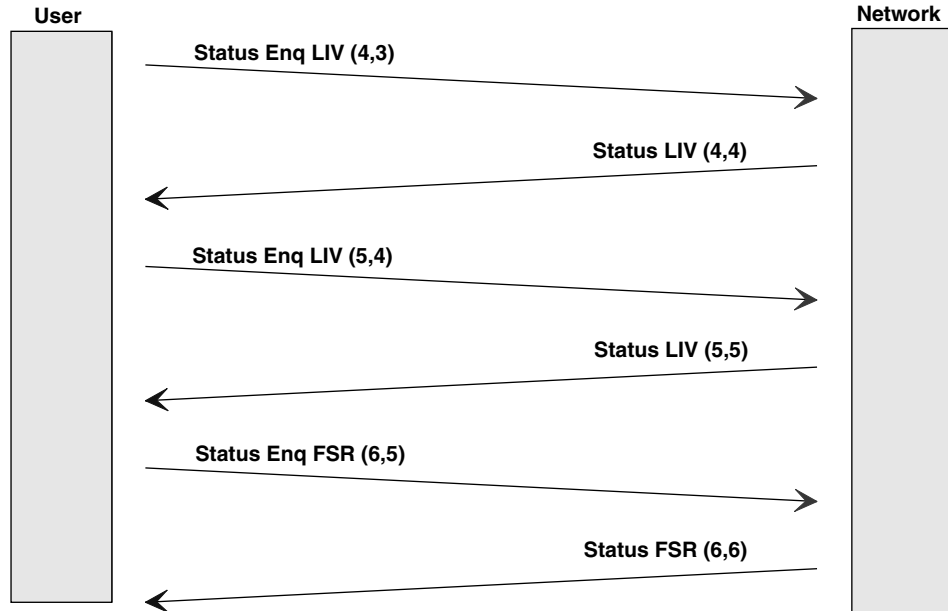


**Figure 8.6** Since the network establishes virtual circuit connections, it also needs to communicate to users the status of connections. To do this, OSI Network layer (layer 3) messages are used. In the interchange shown here, the user requests status from the network using in-channel signaling procedures.



**Figure 8.7** The first level of status has to do with the integrity of the link. If the user still wants to use the link, he or she must send Status Enquiry messages to the network. The network responds with Link Integrity Verification (LIV) messages, which allow for exchanging polling sequence numbers to indicate that the link is active.





**Figure 8.8** To verify the assignment and active state of the PVCs, the user must request a complete status report from the network. This OSI layer 3 message results in a Full Status Report (FSR) message being sent by the network to the user equipment

To indicate that the link has integrity, the network must return a Status Message Response within a few seconds. This response contains the sequence number received from the user and a new send polling sequence number.

For the link to maintain its integrity, the messages must be exchanged according to predetermined timers. The default value for the sending Status Enquiry messages is once every 10 seconds.

### 8.1.10 Full Status Report (FSR)

As a default setting, one out of every six message exchanges consists of a Full Status Report (FSR). In this instance, the user requests a Full Status Report and the network must respond with a message that provides information about each available PVC. This report may indicate whether a PVC is new, active, or inactive. The user must detect that a PVC is deleted when it no longer appears in the Full Status Report. Status Enquiry and Status messages that include a Full Status Report also contain the Link Integrity Verification polling sequence number exchange.

## 8.2 Overview of Test Strategies

The test strategy followed in this chapter is straightforward. By breaking a problem or task down into smaller problems and tasks, it is possible to isolate any element or aspect of the service or equipment. Using the same logical approach, larger problems

are broken down into smaller, more manageable tasks across all areas, including installation, maintenance, management, and troubleshooting. While other approaches could be used, such as inductive problem solving or input/output analysis, procedural analysis proves the most effective.

### 8.2.1 Installation

Installation procedures exist to demonstrate service access and service provisioning according to administratively determined reliability, performance, and management expectations. The first step is to break an installation procedure down into smaller, practical steps:

1. *Verify the physical interface*, showing that it is possible to connect to the service.
2. *Verify the physical layer*, showing that it is possible to send and receive data at the expected line rate.
3. *Verify the link layer*, showing that the established link is reliable and that no unacceptable error indications exist.
4. *Verify in-channel signaling*, ensuring that in-channel signaling procedures are being followed correctly.
5. *Verify PVC assignments*, ensuring that PVCs are assigned as agreed to.
6. *Verify data transfer*, showing that the service will perform data transfers as expected.

Along the way, any indications of reliability problems, protocol problems, and other errors or instances of congestion should be collected.

### 8.2.2 Service characterization

Service characterization serves to assess and measure service performance according to administratively determined factors such as access rate, PVC configuration, and Committed Information Rate (CIR).

Service performance characterization depends on making measurements over a time period to establish baselines for comparing trends and patterns. This process also can be broken down into a sequence of steps:

1. *Link usage and reliability*: Understanding the link usage, reliability, and error indication characteristics of the system under test.
2. *PVC usage and activity*: Assessing PVC usage, activity, and performance relative to expectations.
3. *LAN stack usage and activity*: Characterizing the internetwork by LAN type in order to tune the service. (This step is not always necessary, but may be desirable.)

### 8.2.3 Troubleshooting

Troubleshooting allows isolating the cause of a fault and makes it possible to resolve network trouble as quickly as possible. Some faults are catastrophic, while others are

intermittent. Troubleshooting catastrophic failures parallels the procedure used for installation. Troubleshooting intermittent problems more often follows the procedure for service characterization.

Frame relay is a link-layer protocol service that depends on a physical connection, a Physical layer access port, the Data Link layer, in-channel signaling procedures, and proper PVC configuration. Given this, the process for troubleshooting frame relay can be summarized:

1. Check the connections and cabling.
2. Check the clocking and line rate.
3. Check the integrity of the link itself, making sure that there are no unacceptable error indications.
4. Check that in-channel signaling procedures are being followed correctly, that PVCs are assigned as agreed to, and that equipment is configured correctly.
5. Make note of reliability problems and other error indications or congestion indications.

### 8.3 Testing to Verify T1 or CEPT-E1 Circuits

One of the first and most important tasks in installing or maintaining a frame relay network is to verify correct operation in the Physical layer of the T1 or CEPT-E1 access port. An inactive access port in the equipment, CSU, or network may indicate a bad cable, a bad connector, or a bad port.

#### 8.3.1 Signal presence and frame synchronization

The first step is to verify that the T1 or CEPT-E1 circuits are correctly terminated, either by user equipment, a CSU, network repeaters, or line termination units.

**Signal loss.** If signal loss is indicated, the connections to the circuit may be wrong, the circuit itself may be disconnected or broken, the access ports may be down or bad, or the signals transmitted may be too small to be received. Once the connections to the circuit are eliminated as the cause of the problem, the circuit cable must be checked for continuity. If it checks out, then testing can be done at the various access ports to isolate the problem to one of them.

**Frame loss with signal present.** If frame loss or frame synchronization error is indicated, signal levels may be marginal, or the framing type may be other than that expected for frame relay. After checking the framing (ESF for T1 and CRC-4 for E1), if framing errors still occur it may be necessary to check the signal conditioning and the basic reliability of the line.

Signal levels in T1 DSX-1 are controlled according to a parameter called *Line Build Out*. Both the CSU and DSU may have adjustments for this.

**Checking reliability using BERT.** Reliability generally is expressed in terms of the bit error rate and is determined by running a bit error rate test (BERT). BERT is

usually run when the line (the physical wire) is pulled, in order to qualify it for service provisioning.

### 8.3.2 Verifying test equipment channelization and port access rate

Once good frame sync is established, the data link itself can be verified. Frame relay supports fractional channel services based on using the T1 or CEPT-E1 timeslots, so it may be necessary to check the channelization of the data. The number of timeslots used to form the data channel also determines the port access rate. For example, if six timeslots are used, the port access rate is 384 kbps ( $6 \times 64$  kbps) for most frame relay services.

**Stats and counters.** Determine the number of frames and the number of bad, aborted, and short frames present. If almost all frames are aborted, the channelization of the test equipment must be checked. If all or the vast majority of frames counted are good frames, then the equipment channelization and line channelization are set correctly.

An additional test must be made to determine whether the test equipment shows different equipment and line rates. If so, the fractional channels selected in setup must be changed. If the equipment is set properly, then an administrative error or an error in the network and/or DSU configuration is likely.

### 8.3.3 Verifying V-series circuits

The user-side access to frame relay service is a V-series interface. As with T1 and CEPT-E1 interfaces, the V-series access port must operate properly for the frame relay service to meet expectations. If the access port of the equipment or DSU is not operating, it may indicate a bad cable, a bad connector, or a bad port.

**Monitor or simulate.** The testing procedure varies based on the type of information desired. Test equipment is connected differently depending on whether the service is being monitored or simulated. When installing or troubleshooting the frame relay network interface, the equipment side is simulated. When testing the user equipment, the network is simulated.

**Check results.** If a problem is indicated, it means that the signal is stuck in one state or the other, or is not connected. The connections to the circuit may be wrong, the circuit itself may be disconnected or broken, the access ports may be down or bad, or the signals transmitted may be too small to be received. Once the connections to the circuit have been eliminated as the cause of the problem, the circuit cable must be checked for continuity. If it checks out, tests must be run at the access ports to isolate the problem to its port.

## 8.4 Characterizing Service Performance

A service characterization procedure serves to assess and measure service performance according to such administratively determined factors as access rate, PVC configuration, and Committed Information Rate (CIR).

Testing that characterizes network performance yields important information regarding the value received by the end user, the distribution of usage across PVCs, and the demand for bandwidth by service type and network application.

Service performance characterization depends on making measurements over a period of time in order to establish baselines for comparison, trends, and patterns. Two steps lay the foundation for the process:

1. Identify link usage and reliability levels, and error indication characteristics.
2. Assess PVC activity, usage, and performance relative to expectations.

The primary application of frame relay is to internetwork LANs. For this reason, it is also common to characterize internetwork usage by LAN type to tune the service.

Because frame relay multiplexes many PVCs using different DLCIs, it provides a convenient point to get an overview of what is happening in the network. With LAN-over-WAN monitoring capability, you can check to see if the LAN traffic is passing across the WAN as expected. For service providers, this may be the only point of access to the UNI available for resolving internetwork trouble.

#### 8.4.1 Assessing access port usage and reliability

Performance can be assessed from a number of points, among them the DLCI and LAN stack. Once collected, statistical information can be logged and accumulated to allow for long-term characterization.

**Usage by DLCI.** A good frame relay analyzer can detect the number of DLCIs active in a system and identify the level of usage for each one. Usage includes the information rate, sometimes called *throughput*, and utilization. These values can be compared directly to the previously negotiated CIR.

**Usage by LAN stack.** A frame relay analyzer also can identify which LAN stacks are active and the utilization rates for each one. Some LAN packets are encapsulated according to RFC 1490 and use Q.922 Unnumbered Information (UI) frames of type 03. The analyzer must be set specifically to search for these packets.

In addition to these two primary methods of encapsulating data on frame relay, many others also are in use. A given router or bridge may have its own unique (proprietary) interlayer or sublayer between the frame relay header and the LAN protocols.

## 8.5 Internetwork Troubleshooting

Frame relay is a relatively simple protocol. As outlined above, it is a link-layer protocol service. It depends on a physical connection, a physical layer access port, Data Link layer conformance, in-channel signaling procedures, and proper PVC configuration.

Internetwork troubleshooting commences by checking the connections and cabling, the clocking, the line rate, and any channelization. It is also useful to check the integrity of the link itself to verify that no unacceptable error indications exist. Conveniently, frame relay has link integrity verification built in, in the form of the in-channel signaling procedures. If these procedures are followed correctly, it is possible to see if the PVCs that are assigned are set up as agreed and that equipment is con-

figured correctly to make use of them. Reliability problems, protocol problems, and other error indications or congestion indication should be noted.

Every installation comes down to troubleshooting the connections and the line, the clocking and channelization, the data link, the in-channel signaling, and the router and switch configuration. Depending on the information gathered, troubleshooting proceeds in one of four ways:

1. If there are no clocks or the data is not being framed correctly, the most likely cause is the connections or the cabling itself. A BERT can pinpoint the cause of line problems.
2. If there are clocks and the data is being framed correctly, but there are no in-channel signaling polling sequences, the most likely cause is that the equipment signaling type is not configured correctly or that the switch itself is not responding.
3. If the polling sequence is correct, but there are no PVCs designated, then the switch and the network are suspect.
4. If the expected PVCs are available but data cannot be interchanged, then the switch, PVC implementation, or the far-end LAN may be down.

For all intents and purpose, frame relay is about interconnecting LANs with logical private lines. Unlike leased lines, where each link is a physical point-to-point connection, frame relay multiplexes several circuits onto one physical connection, and this has the potential to make troubleshooting more complex as network complexity increases.

# Integrated Services Digital Network

**Mark Powell**

*Hewlett-Packard Co., Colorado Springs, Colorado*

## 9.1 Introduction

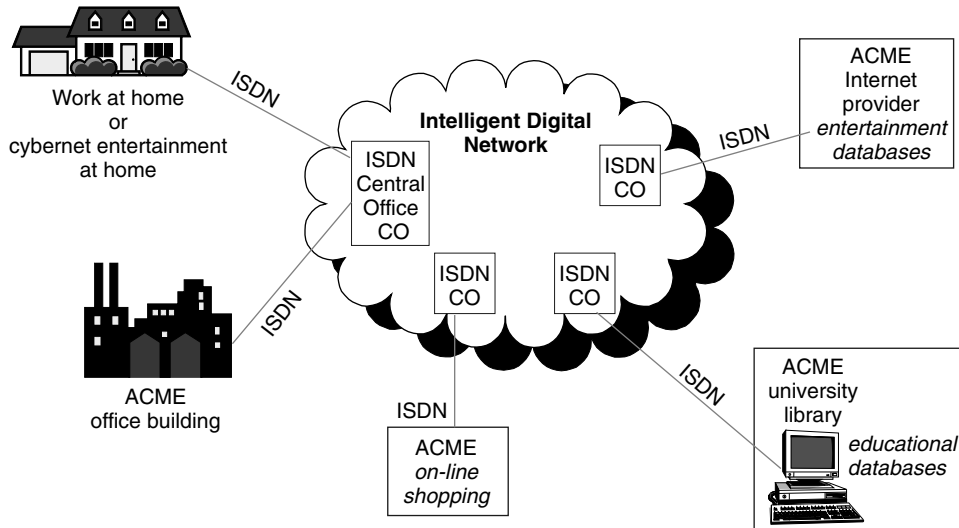
Integrated Services Digital Network (ISDN) provides a higher-bandwidth access for network subscribers and users of the Intelligent (or Integrated) Digital Network (IDN). Modern applications and information sources have pushed users of networking services to demand higher performance and higher access speeds. This includes people working from home, accessing cyberspace entertainment facilities from home, and working from offices and commercial locations. Older access technologies (analog modems, low-speed data services, packet-switched services) simply do not meet modern demands.

### 9.1.1 Introduction to Integrated Services Digital Network

For a user working from home or at an office, ISDN provides the means of access to the IDN. Figure 9.1 shows the variety of locations from which users might need to access remote databases, connect to the local area network (LAN) at the office, or use any of the services available to users over the IDN. ISDN satisfies some important user demands placed upon these services:

- Greater bandwidth for higher-speed applications
- Greater reliability, better performance, and greater security
- Simple connection and interaction with the network
- Cost savings

**Greater bandwidth.** The connection from the user to the network often is referred to as the *local subscriber loop*. By far the most of the installed local loop access is via common twisted-wire pairs, designed for standard telephone circuits, sometimes referred to as “Plain Old Telephone Service” (POTS). This is the traditional, voice-



**Figure 9.1** ISDN provides the local access into the IDN, making available a variety of communication and computing services to different user profiles.

grade analog telephony service supporting one voice conversation over one pair of twisted wires extending from the user location to the telephone company or service provider Central Office (CO). For a private residence, there might be from one to three pairs of wires installed. Commercial premises that have a Private Branch Exchange (PBX) or LAN could have many pairs of twisted wires.

The move from older analog POTS local access to ISDN allows subscribers to achieve much greater rates of data transfer (greater bandwidth) using newer digital techniques, rather than older analog modem techniques. ISDN is a digital line. There is no dial tone and no ringing voltage. Each existing pair of wires is capable of being converted to support a Basic Rate ISDN (BRI) or a Primary Rate ISDN (PRI) access. A BRI can provide two separate 64 kbps data streams that can be used separately, such as with simultaneous digital voice and digital data, or combined together to form a 128 kbps data stream. Commercial users often can justify the costs of stepping up to a PRI access that can provide 23 separate 64 kbps data streams, or a single aggregate 1.536 Mbps data stream.

**Greater reliability.** Modern networks provided by the telephone companies and other service providers have evolved at tremendous rates. There has been a remarkably rapid shift from older, analog-based networks to entirely digital networks, providing both digital local access and digital transport. Greater reliability is possible as a result of networks becoming all-digital. Digital systems are also much less susceptible to noise and other imperfections that introduce errors, thus providing better performance. Modern digital signal processing techniques also allow for greater levels of security.

**Simple connection.** When a user has a telephone, fax machine, and computer to connect to a network, there could be three separate interfaces involved. If interna-



tional travel is required, the user might find that each country has a different requirement. Through international standards, ISDN has an overall objective to provide a small, standard set of user and network interfaces that allow users standard access to a variety of network services. ISDN can integrate virtually all forms of communication, including voice, data, and video.

**Cost savings.** Any change to a new technology such as ISDN, or adoption of a new technique requires significant cost justification. By converting and using the very large existing installed plant of twisted-pair local loops, the service provider—and ultimately the end user—can achieve significant levels of cost savings. A single pair of wires capable of a single voice or data conversation can be converted to a BRI or PRI ISDN line, capable of carrying more than one conversation. The cost-of-service for a BRI line capable of two voice or data conversations typically is tarified comparably to that of two separate analog voice-grade lines.

Although ISDN technology has been available since the mid-1980s, only in the last couple of years has there been market demand and notable growth in ISDN. A number of factors have fueled this growth. End users continue to ask for and expect to receive ever-growing performance and features from their network and communication systems. With ISDN, end users now have the ability to set up applications that combine both voice and data over a single network. ISDN can be used to tie remote terminals or personal computers into other computers, LANs, or private wide area networks (WANs). Important services and applications include:

- **Voice telephony** Voice communication is supported simultaneously with data services, allowing a user to talk and access data at the same time.
- **Facsimile** Older data transmission methods are not very fast. ISDN supports services for the transmission and reception of graphics, images, etc., at data rates as high as 64 kbps.
- **Images** Medical images, interactive video, and in-home movies are examples of images that can be digitally transferred over the ISDN by using 64 kbps data streams or aggregate multiples of 64 kbps data streams to achieve higher-bandwidth performance.
- **LAN connectivity** A PC using ISDN can log in remotely to a local area network. In this way, the user has all the same features as if connected directly to the LAN, such as file transfer, database access, and electronic mail. This application is a great benefit to telecommuters, a growing segment of the working population.
- **Internet access** The Internet is the prime example of an interlinked web of networks, penetrating over 100 countries with 11,000 separate networks feeding into it, containing up to 1.7 million host computers. Home-based access and rising Internet use are creating demand for faster access to the multimedia presentation of the World Wide Web (WWW). “Surfing the Web” becomes agonizingly tedious if not accomplished at the higher data rates made possible by ISDN.
- **Videoconferencing** Combining simultaneous voice, data, and video capabilities enables practical videoconferencing to take place. Videoconferencing or desktop

conferencing can preclude the need for travel, and allows immediate attention to be given to problems that arise on a day-to-day basis.

- **Automatic Number Identification** ANI allows the calling party's phone number to be passed to the party called. Companies that adopt this approach can use this information to search a database and retrieve caller profile data. The caller's information is displayed on a screen as an operator answers the phone. This reduces the time spent on each call, and allows each caller to be treated with more personalized service.

**Narrowband versus broadband ISDN.** Emerging telecommunication networks can provide digital bandwidth up to 2.4 Gbps. Frame relay, Asynchronous Transfer Mode (ATM), and other broadband technologies allow users to access the broadband services available on the Integrated Services Digital Network. B-ISDN is required to deliver video telephony and high-speed data services. For the purposes of this discussion, narrowband ISDN can be differentiated loosely from B-ISDN on the basis of data rates; narrowband ISDN applies to data rates of 2.048 Mbps and lower. The discussion in this chapter focuses on traditional or narrowband ISDN.

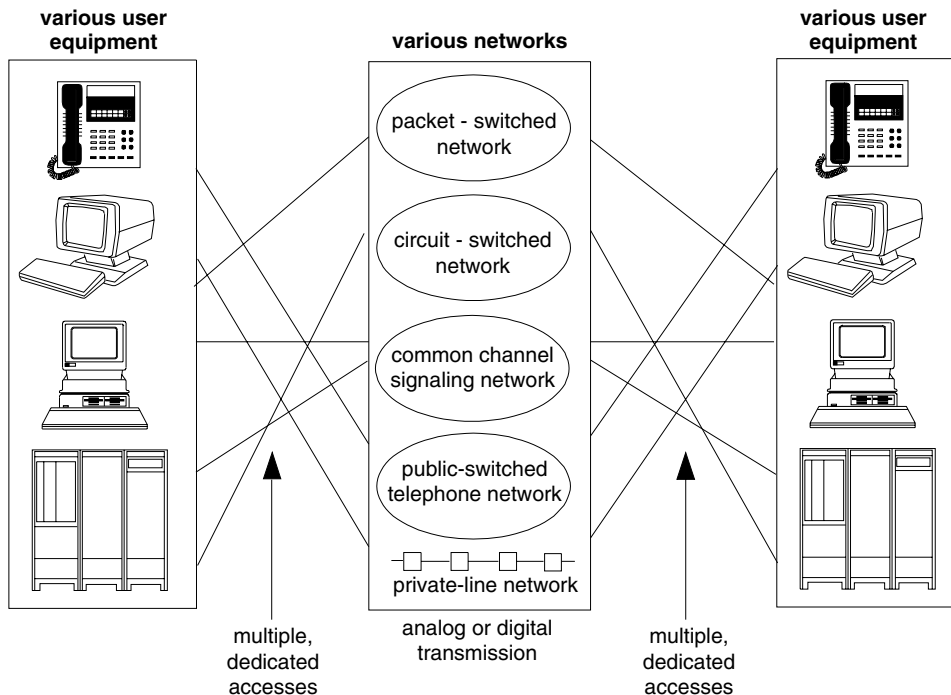
### 9.1.2 ISDN and the intelligent network

In older, analog networks, access to the transport networks was accomplished using a variety of methods that depended on the equipment users wished to connect. Users typically needed an assortment of physical interface and access procedures to connect their devices (Figure 9.2). With the push by service providers to offer better service and more features, a rapid migration from older, analog-based networks to digital networks has been occurring. In the analog systems, signaling information (dialing a number) travels over the same channel as voice; this is known as *in-band signaling*. The signaling information consists of either electrical current pulses or tones.

Figure 9.3 illustrates modern digital networks, where the local subscriber access is provided by ISDN and the service provider transport network has migrated to an all-digital system. With a digital network, the signaling information is sent in a separate channel from the voice/data information; this is known as *out-of-band signaling*. The digital signaling information consists of protocol-based messages that provide signaling or connection control and management. This message-oriented signaling method has the advantage of not consuming valuable information channel bandwidth, leaving a clear channel for voice and data traffic. One signaling channel can control one or many traffic channels, thus increasing the efficiency of the networks.

In addition to serving as a transport medium for digitized voice and data, these new digital networks are being designed to process information within the network, thus becoming "intelligent." The evolving Intelligent (or Integrated) Digital Network has moved rapidly toward centralized, high-speed databases that control network call routing. The IDN accomplishes this routing and database access by utilizing high-speed signaling links between CO switches and the various regional and national switching centers.

The IDN can be divided into two distinct portions. The first consists of ISDN, which provides a standard user-to-network interface (UNI) point. The second part of the



**Figure 9.2** Before ISDN was available, network users had to have a separate connection for each service required. In particular, access to a telephone system for voice was wired separately from data connections. Connection into a wide area network was separate from that required for access into an organization's private intranetworking. This resulted in the need for different connectors, service providers, and user interfaces. Maintaining all these links was time-consuming and costly.

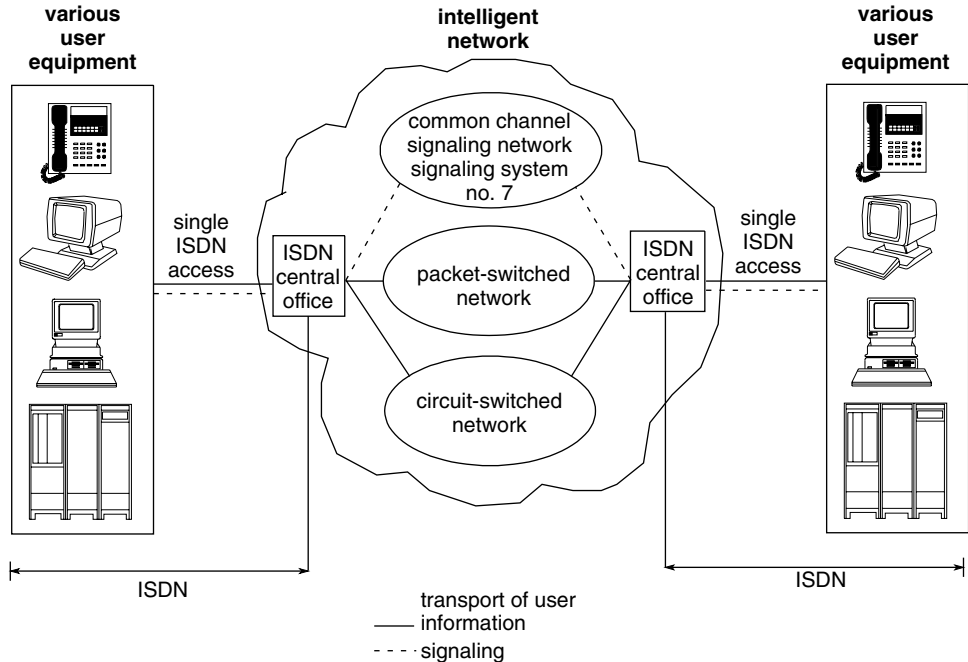
network consists of the transport networks that are controlled by Signaling System 7 (SS7). SS7 is a message-based protocol signaling method. SS7 is the backbone of the IDN. Without it, the advanced and rich features of the IDN would not be possible. ISDN uses a different and separate protocol from SS7. Translation between the two control protocols occurs at the ISDN CO or switch. Strict adherence to standards, particularly when dealing with international networks, obviously is very important.

### 9.1.3 ISDN functions

ISDN can be divided into two distinct functional areas:

1. Signaling
2. Transport of user data

ISDN uses out-of-band, message-oriented signaling to control everything regarding a call. In message-oriented signaling, every operation, from telling the system the phone is off-hook to disconnecting the call, is performed by sending messages back and forth. The signaling channel then assigns separate channels to transport the



**Figure 9.3** The control system for the Intelligent Network is common channel signaling. At the end-user side, ISDN provides the link between the customer premise and the Central Office (or ISDN exchange). Within the Intelligent Network there are high-speed databases controlling network call routing and other functions. There also is high-speed Data Link signaling between the Central Office switches and between the switches and the databases. This inter-exchange common channel signaling method is Signaling System 7 (SS7).

users' data. Signaling is a key element of both ISDN access and the Intelligent Network. If the signaling system goes down, the network cannot function. Signaling actions are classified in these four categories:

- Supervision information
- Control information
- Addressing (dialing)
- Alerting (ringing)

Supervision information, such as on-hook or off-hook, refers to the state of the interface. Control information, such as hold or forward, can refer to a variety of value-added features, such as 800 or free phone services. The major contribution of the Intelligent Network and ISDN beyond older networks is primarily in the area of control. Addressing and alerting are familiar and therefore somewhat self-explanatory concepts.

When a call is placed, the signaling system is responsible for:

- Routing the call through the network,
- Overseeing the connection,

- Handling the billing and financial administration, and
- Eventually disconnecting the call.

#### 9.1.4 ISDN channel types

Information is transferred between a user and the Central Office (or ISDN station) via channels. A *channel* is defined as a specific portion of the total digital bandwidth of the transmission line. ISDN standards define B, D, and H channels (Table 9.1), but the most widely used are the B and D channels.

The B, or *bearer channel*, is a 64 kbps digital channel. It does not carry signaling (control) information. Digitized voice or data transmissions (including video) in either circuit-switched or packet-switched formats can be transported, however. Older, standard data terminals may be adapted to the B channel through well-defined rate adaption algorithms (like V.110 and V.120). B channels also may be combined to achieve greater aggregate speeds. *Multilink Point-to-Point Protocol* (MLPPP) or *Bandwidth on Demand* (BONDing) are two major methods for achieving higher aggregate speeds. For example, the two 64 kbps B channels of a BRI may be combined to achieve 128 kbps aggregate data speed.

The D, *demand or data channel*, is a separate 16 or 64 kbps channel used primarily for signaling information. Signaling information establishes, maintains, and clears ISDN network connections. The nature of the signaling functions cause signaling to occur in bursts. When the D channel is not carrying signaling information, provisions have been made to allow packet-switched (X.25) data to be transmitted. Signaling information, however, has priority on the D channel at all times.

The H channel has been designed for high-bandwidth applications and bonds multiple B channels. H channels provide greater aggregate bandwidth in PRI applications. This capability of channel aggregation allows multi-rate communications on a dynamic basis through inverse multiplexing over multiple B channels.

Table 9.1 summarizes the functions of the B, D, and H channels.

#### 9.1.5 Basic Rate versus Primary Rate ISDN

The ISDN standards define user access to ISDN using B and D channels to create different channel configurations. These channels are then Time Division Multiplexed to create an aggregate signal on the transmission line. The implementation

**TABLE 9.1 ISDN Channel Types. The B and D channels of ISDN are specific portions of the total bandwidth of the transmission line. The B channel carries the voice/data information and the D channel the signaling information. The H channels are aggregated B channel, providing higher bandwidth to the users.**

Channel Type	Data Rate	Function
B	64 kbps	user information-data, voice, video
D	16 kbps or 64 kbps	signaling/control information (for the B channels) and user data (packet-switched)
H	H <sub>0</sub> —384 kbps H <sub>1</sub> —1.536 Mbps	equivalent to B channels

of ISDN has been approached in two different ways, Basic Rate Interface and Primary Rate Interface.

**Basic Rate Interface (BRI).** The Basic Rate Interface (BRI) consists of two B channels and one D channel. This configuration is often called  $2B + D$ . The two B channels may be independently accessed. For example, one B channel can carry voice information while the other B channel is carrying data. In this manner, voice and data can be integrated over the same transmission facilities. The D channel carries the signaling information controlling the two B channels, as well as being used to transfer packet-switched data, like X.25, in the extra bandwidth.

A single BRI can support up to eight devices (telephones, fax machines, PCs, modems, etc.). While BRI supports as many as three simultaneous calls, only one can be a voice conversation. BRI typically is implemented using an 8-pin RJ-45 connector. Full-duplex connectivity is accomplished over a twisted-pair local loop through the application of special carrier electronics.

**Primary Rate Interface (PRI).** There are two versions of Primary Rate Interface (PRI). In North America and several other locations in the world, the primary rate interface consists of 23 B channels, a D channel, and overhead. The second version, used in Europe and throughout the rest of the world, consists of 30 B channels, a D channel, and overhead. The standards specify that a D channel can support up to five PRI connections. PRI provides a full-duplex point-to-point connection.

Table 9.2 summarize the offerings of both BRI and PRI.

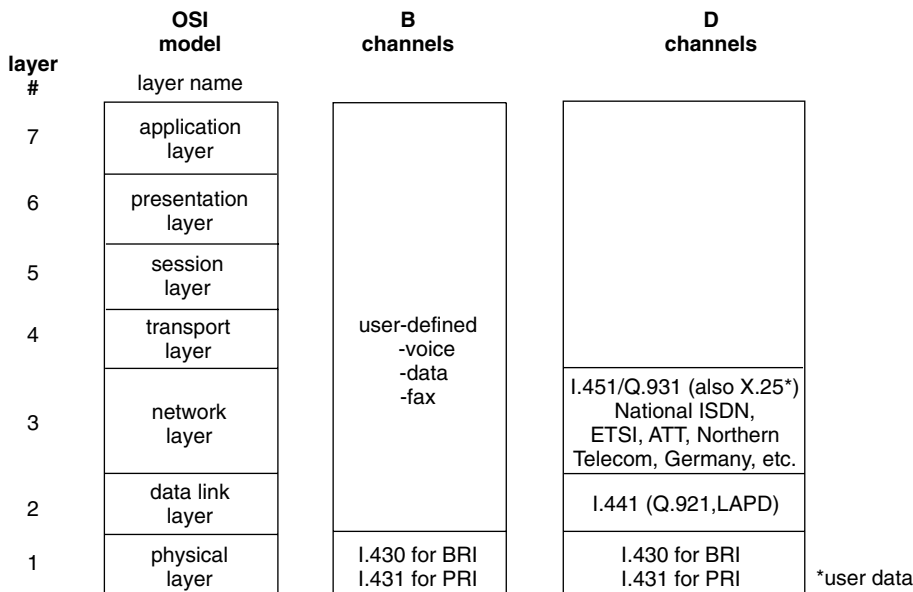
### 9.1.6 ISDN standards and the OSI Model

ISDN is defined for the first three layers of the Open Systems Interconnect (OSI) seven-layer Reference Model. Worldwide definition and approval of the ISDN standards, referenced as *I-series* and *Q-series* sections of the CCITT standards, is carried out by the International Consultative Committee for Telephone and Telegraph (CCITT) standardization body. A summary of the ISDN protocols and their relationship to the OSI model is shown in Figure 9.4.

National variations and extensions, as well as ISDN switch-specific variations and extensions, have occurred. These variations are based primarily upon the CCITT set of standards. These variations and extensions have led to many of the problems and

**TABLE 9.2 Summary of BRI and PRI Channel Capacity. The two basic types of ISDN, Basic Rate Interface and Primary Rate Interface, provide aggregate data rates that range from 144 kbps to 2.048 Mbps. S1 is equivalent to North American digital transmission speeds of T1 (1.544 Mbps), and S2 is equivalent to CCITT digital transmission speeds of E1 (2.048 Mbps).**

Type of interface	Referred to as	#B channels	Capacity of each	+ #D channels	Capacity of each	+ # misc channels	Capacity of each	= Total capacity
BRI	S0	2	64 kbps	1	16 kbps	none		144 kbps
PRI v1	S1	23	64 kbps	1	64 kbps	1	8 kbps	1.544 Mbps
PRI v2	S2	30	64 kbps	1	64 kbps	1	64 kbps	2.048 Mbps



**Figure 9.4** Side-by-side correlation of the seven-layer OSI Reference Model and the appropriate CCITT specifications as they are defined for both the B and the D channels are easily identified in this illustration. Specifications for the B channel address layer 1. Specifications for the D channel address layers 1, 2, and 3.

issues involving interoperability and incompatibility between different vendors' ISDN equipment, ISDN connections between different countries, or between different Regional Bell Operating Companies (RBOCs) within the U.S.

**B channel.** Recall that the B channels carry only user information: voice, data, facsimile, or video. Because of this, the only ISDN protocol specified for the B channel is at the OSI Physical layer (layer 1). B channels carry voice as Pulse Code Modulation (PCM) or other digitizing schemes, and can carry data.

Rate adaption for data is defined for the B channel by the International Telecommunication Union Telecommunication Standardization Sector (ITU-T). B channel procedures are defined for Terminal Adapters (TA) in Europe under V.110, and under V.120 for TAs in North America.

Video is defined by the ITU-T as H.320, the umbrella standard for videoconferencing that addresses narrowband visual telecommunications systems and terminal equipment.

If the configuration is a BRI, the protocol is specified by section I.430 of the CCITT and is similar to the U.S. ISDN specification. The remaining layers of the B channel (layers 2 through 7) are user-defined.

If the configuration is a PRI, the Physical layer is specified by CCITT-I.431 and also is the same as in the U.S. National ISDN specification. The remaining layers of the B channel (layers 2 through 7) are user-defined, depending on the type of traffic.

**D channel.** Since the D channel is time-multiplexed onto the same transmission media as the B channels, the specifications for layer 1 are similar to those for the B

channels. On the D channel, however, layers 2 and 3 are also specified to provide signaling. The Data Link layer (layer 2) protocol is defined by CCITT-I.441 or Q.921 and is commonly referred to as the Link Access Protocol for the D channel (LAPD). This also is similar to the U.S. ISDN specifications.

The Network layer (layer 3) for the D channel is specified by CCITT-I.451 or Q.931. It is at the Network layer for the D channel where proprietary variations among switch vendors and countries may occur. These implementations use Q.931 as the base, but add enhancements and new features that were not defined in the CCITT specifications. This is where problems can arise. In the U.S., the national ISDN specifications define the variations and additions to Q.931. ATT, Northern Telecom, and others have their own custom versions that are slightly different, however, potentially introducing interoperability and compatibility problems.

## 9.2 ISDN Architecture and Operation

This section will describe the types of ISDN equipment and how the equipment is interconnected to create ISDN networks. On the user's premise there are two types of functional blocks:

- Network Termination Equipment (NT)
- Terminal Equipment (TE)

*Functional blocks* are logical representations that perform specific functions. Functional blocks may be combined when designing real equipment. Depending on the user's needs and network configuration, some functional blocks might not be necessary.

The interfaces between functional blocks are called *reference points*. Reference points also are logical rather than physical; there might not be a physical interface at a given reference point. This is the case when the functions of one piece of equipment are provided in another piece of equipment. By interconnecting functional blocks and reference points, ISDN networks can be constructed.

**Network Termination (NT) Equipment.** Network Termination (NT) equipment handles communication between the ISDN exchange and the customer premises. NT equipment typically is the demarcation point ("demarc") between the customer premises and the network administration. There are two types of NT equipment, NT1 and NT2.

NT1 devices provide functions equivalent to the Physical layer (layer 1) of the OSI model. These functions include signal conversion, timing, maintenance of the physical transmission line, and the physical and electrical termination of the network at the user end. Sometimes the NT1 is built into another piece of equipment and therefore might not exist physically as a separate device. The functionality of the NT1 must be present in an ISDN network, however.

NT2 devices are more intelligent than NT1 devices. NT2 devices perform Data Link layer (layer 2) as well as Network layer (layer 3) functions. Whenever the NT2 does not provide layer 3 capability, then the NT2 will pass the original layer 2 and layer 3 data received from NT1 to the Terminal Equipment. NT2 equipment provides local premises distribution functions, like controlling multiple BRIs feeding into a



single PRI. NT2 examples include PBXs, concentrators, terminal controllers, front-end processors, and T1 multiplexers.

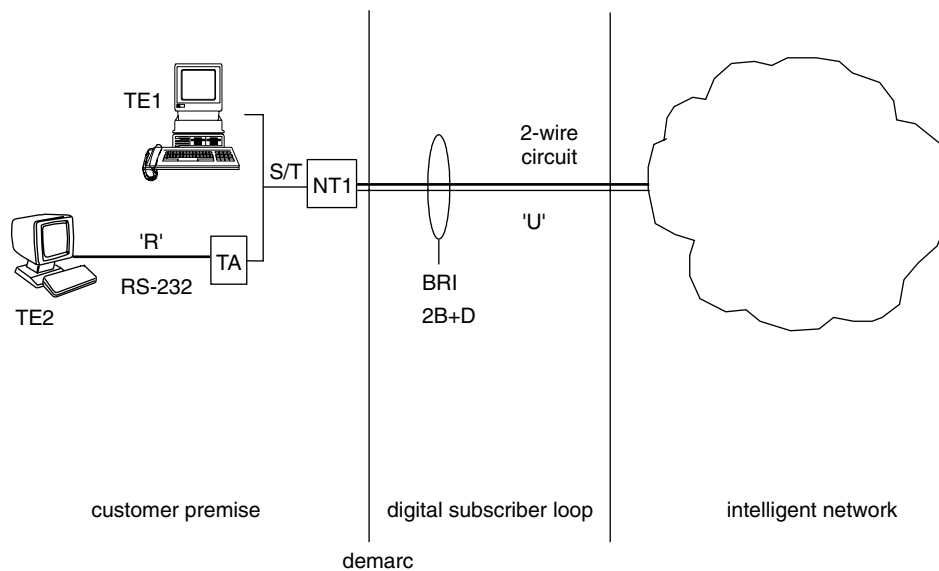
**Terminal Equipment (TE).** Terminal equipment handles communication on the customer premises. Examples of terminal equipment include data terminals, telephones, personal computers, and digital telephones. TE devices provide protocol handling, maintenance functions, interface functions, and connection functions to other equipment.

Terminal Equipment type 1 (TE1) devices perform the functions listed above, as well as containing an interface that is compatible with the ISDN network interface recommendations. Examples of TE1s include voice/data terminals, digital telephones, and computers with ISDN cards and software.

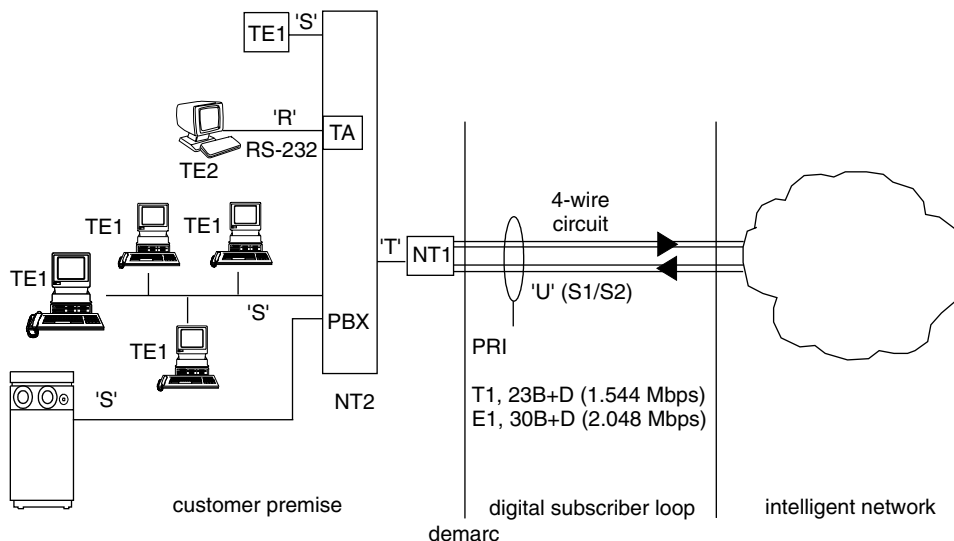
Terminal Equipment type 2 (TE2) devices also perform the TE function as listed above, except for the signaling protocol. TE2s do not contain an ISDN-compatible interface. Instead, they have a non-ISDN-compatible interface, such as RS-232, V.35, or X.21. TE2s must be connected to ISDN through a Terminal Adapter (TA). Today's standard personal computers and telephones are examples of TE2s.

Terminal Adapters (TA) allow TE2 devices to interface to an ISDN network. TAs perform such functions as converting non-ISDN transmission rates and protocols to ISDN standards. TAs also provide the D channel signaling. TAs may be separate devices, or they may be integrated into an NT2 or a TE2.

**Reference points.** Figure 9.5 shows a typical Basic Rate network; Figure 9.6 shows an ISDN Primary Rate user-network interface. The various functional blocks previ-



**Figure 9.5** This Basic Rate Interface (BRI) user-network interface has an information carrying capability of 144 kbps. At the S/T BRI interface, there are additional overhead bits (control, framing, etc.), and the total transfer rate for this interface is 192 kbps. At the U BRI interface there is a different configuration of overhead bits and the total transfer rate for the U interface is 160 kbps.



**Figure 9.6** This Primary Rate Interface (PRI) user-network interface shows that the primary rate line is used between the Central Office and the customer premise. An NT2 uses the primary rate line as a trunk to service the many basic rate lines feeding into it. The NT2 takes care of all the tasks associated with maintaining the basic rate lines, as well as setting up calls to the Central Office via the primary rate line.

ously defined (NT, TE, TA) are interconnected by reference points. Reference points are conceptual and do not always have a physical interface. They are the connection points between functional blocks. ISDN reference points are referred to simply as R, S, T, U, V.

*R reference point* is a non-ISDN interface (such as RS-232, V.35, or X.21) between a non-ISDN terminal (TE2) and a TA.

*S reference point* is a four-wire interface (one pair to send, one pair to receive) between a TE1 and an NT, or between a TA and an NT. Up to eight TE1s or TAs may be connected by an S reference point to an NT. An NT2 effectively splits the T reference point into several S reference points. The S reference point is described in CCITT section I.440 (basic rate) or CCITT section I.441 (primary rate), as well as other national standards.

*T reference point* is a four-wire interface between a TA and a TE1, or between an NT2 and an NT1. Physically this interface is identical to the S reference point. In some cases, such as a PBX (NT2), the NT1 is built into the NT2 and there is no physical T reference point.

*U reference point* is the transmission line between the Customer Premise Equipment (CPE) and the ISDN exchange. Specifically it is between the NT1 and the exchange's line-termination equipment (LT). For a BRI, the U reference point is a full-duplex interface over a single pair of twisted wires. (The same wires are used to send and receive.) The PRI utilizes a four-wire interface.

*V reference point* divides the LT equipment from the Exchange Termination (ET) equipment. In actual practice, the LT and ET may be the same equipment, and the V reference point would not exist.

### 9.2.1 ISDN Physical layer

The purpose of the Physical layer of the OSI stack is to provide the electrical and functional procedures needed to transmit the data onto the physical media. Two CCITT specifications exist for layer 1 of ISDN:

- I.430 for the BRI
- I.431 for PRI

U.S. ISDN specifications are similar to these. Since the B and D channels are time-multiplexed onto the same transmission line, the specifications for the B and D channel are the same. These standards define and provide for the following capabilities:


- Transmission capability, timing, and synchronization functions.
- The necessary procedures to allow Network Service Provider or Customer Premise Equipment to be activated or deactivated.
- Signaling capability and the necessary procedures to allow terminals to gain access to the common D channel in an orderly way.
- The necessary procedures to perform maintenance functions.
- An indication of the layer 1 status to higher layers.
- Point-to-point capability, as well as point-to-multipoint arbitration capability.
- Determination of the bit formats in layer 1.
- Voltage levels on the physical media.

**BRI INFO signal/states.** In the process of establishing, maintaining, and disconnecting a BRI at the S or T reference points, there is a handshaking or communication between the Terminal Equipment (TE) and the Network Termination (NT). There are four specific signals that occur on the S and T reference points during interactions between the TE and NT. These signals are called INFO (information) signals and are defined in the specifications.

The INFO signals depend on the state of the link and may occur in any order. Either the TE or NT may initiate a connection. The TE and NT move progressively from INFO 0 to INFO 3, and from INFO 0 to INFO 4, respectively. INFO 3 and INFO 4 are the states that signify that the Physical layer link is established and synchronized with the flow of proper frames. Figure 9.7 illustrates this interaction. Monitoring and observation of the INFO states provides important and useful diagnostic information when troubleshooting a BRI connection. The BRI U interface also has similar handshaking processes and signals that provide information regarding the status of the layer 1 connection.

The full-duplex BRI data stream between a TE and the NT (S reference point) is 192 kbps and consists of two B and one D channel. It also has additional overhead (control) bits that allow the BRI to support both point-to-point (single endpoints) or point-to-multipoint. Point-to-multipoint (or passive bus) allows for up to eight independent ISDN stations, each capable of two B channels.

Purpose: to establish, maintain or terminate an ISDN layer 1 link



**Figure 9.7** Establishment of the physical link for BRI involves handshaking between the Terminal Equipment (TE) and the Network Termination (NT). Either side may initiate the connection; then communication progresses through the indicated INFO states until the layer 1 link is established, signaling that data transfer can begin.

### 9.2.2 ISDN Data Link layer

Layer 2, the Data Link or Frame layer interface, is responsible for the reliable transfer of information across the physical links. The Data Link layer:

- Ensures error-free data transmission between layer 3 entities across the user-to-network interface by providing error detection and correction.
- Receives services from layer 1 and provides services to layer 3.
- Provides the form of the bit stream (frame format) and provides flow control.

The protocol running over the D channel at the Data Link layer is defined as CCITT-I.441 (Q.921) and is commonly known as Link Access Procedure for the D channel (LAPD). The United States ISDN specifications are similar. The B channel protocols for the Data Link layer can vary from High-level Data Link Control (HDLC) to voice. The B channel protocols are not defined by the ISDN standards and can consist of whatever the user wants to transmit, as long as the protocols conform to the layer 1 standards.

LAPD provides layer 2 addressing, flow control, and error detection for the D channel. The error detection of layer 2 is responsible for finding transmission errors that might have occurred. In the areas of flow control and error detection, LAPD is very similar to Link Access Procedure-Balanced (LAPB), which is layer 2 for X.25. LAPD differs, however, in the addressing capability that it provides. LAPD allows for multiple logical connections at the Data Link layer. This is needed because the D channel controls all of the B channels that can operate independently and requires different logical connection on the interface.

The LAPD layer 2 uses a frame structure with fields that include:

- **Flags** These are used for frame synchronization; the pattern equals 01111110 (7E hexadecimal).

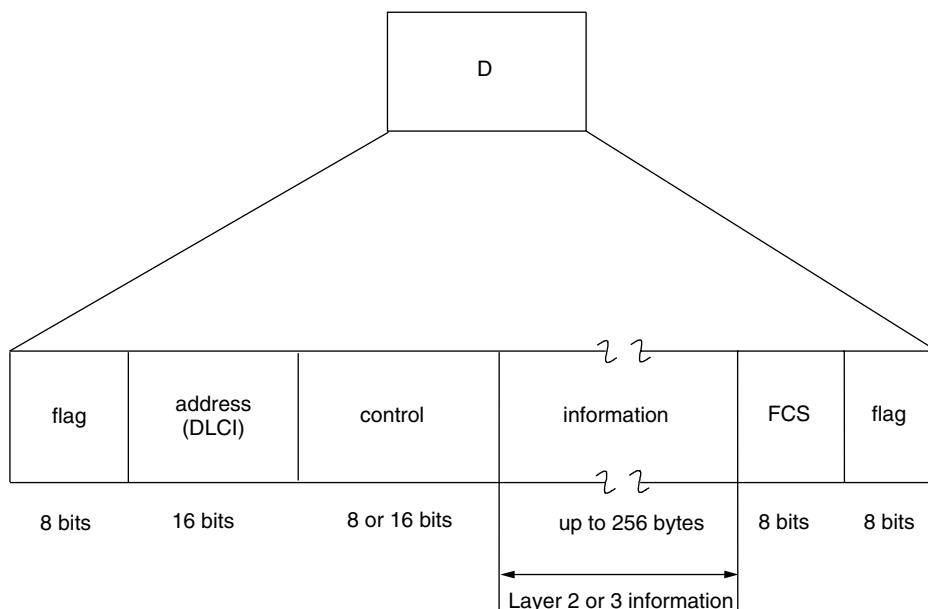
- **Address** The address field is the Data Link Control Identifier (DLCI) that provides the multiplexing required to support multiple Data Link connections.
- **Control** The control field is for controlling information transfers and for supervisory functions.
- **Information** If present, this is a variable-length field containing the actual information (message packet) for layer 2 or layer 3.
- **FCS** This is a Frame Check Sequence for error checking.

Figure 9.8 illustrates the LAPD frame structure.

LAPD frames are defined by the control field formats. These include numbered information frames (I frames), supervisory frames (S frames), and unnumbered information transfers and control frames (U frames).

- *I frames* control the transfer of layer 3 information to layer 3 entities.
- *S frames* handle layer 2 flow control management, such as acknowledging I frames, etc.
- *U frames* provide additional transfer capabilities and Data Link control functions.

One recent significant development that allows users to dynamically change bandwidth as the need changes is Multilink Point-to-Point Protocol (MLPPP). The IETF RFC 1990, "Multilink Point-to-Point Protocol (MLPPP)," will extend the use of ISDN. MLPPP takes advantage of the ability of switched WAN services



**Figure 9.8** The LAPD frame structure contains five fields. The flag, address, control fields and FCS are of fixed length; the information field can vary in length up to 256 bytes to accommodate varying message sizes.

to open multiple virtual connections between devices to give users extra bandwidth as it is needed.

With MLPPP, routers and other access devices can combine multiple PPP links connected to various WAN services into one logical data pipe. MLPPP is independent of the actual physical links and the WAN services that run over them. It functions as a logical link layer, dynamically adding or removing links between two communicating devices as bandwidth needs change. It allows the additional bandwidth to be added without disrupting the existing WAN infrastructure. With MLPPP, different WAN services (such as ISDN, frame relay, and ATM) can be used together.

### 9.2.3 ISDN Network layer

As with the layer 2 discussion, the protocols involved at layer 3 are split between B channel protocols and D channel protocols. On the B channel, ISDN standards do not define a protocol. The D channel has two protocols currently defined: CCITT's X.25 and I.451 (more commonly referred to as Q.931).

**X.25 functions.** The X.25 protocol is used to transport user data over the D channel when the channel is not being used for signaling.

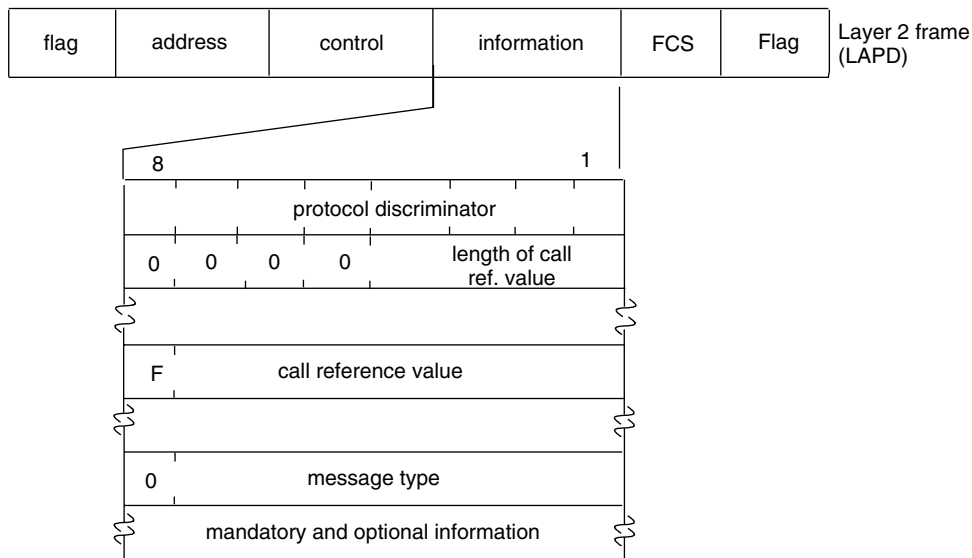
**Q.931 functions.** The Q.931 protocol performs signaling in the ISDN environment that is used to establish, maintain, and terminate network connections. The U.S. ISDN specifications vary from Q.931 and other implementations in the addition of Information Elements beyond the Q.931 specification.

The main purpose of layer 3 and Q.931 is to establish, maintain, and terminate connections across the ISDN and Intelligent Network (via the SS7 network). In addition, Q.931 also is in charge of allocating resources, such as B channels and X.25 connections on the D channel. Q.931 also has numerous timers and counters used to ensure that the signaling information is transmitted correctly and arrives error-free. The Q.931 error recovery ensures that:

- Packets of information arrive in the proper order.
- Information packets are appropriate for the state of the connection.
- Messages are properly acknowledged.

**Q.931 message structure.** Q.931 uses messages to convey information between two layer 3 entities. The key elements of the frame, illustrated in Figure 9.9, are:

- *Protocol discriminator* distinguishes between messages for signaling/Q.931 and other protocols, such as X.25.
- *Call reference value length* is the number of octets (length) of the actual call reference value.
- *Call reference values* (CRVs) are assigned for a call by the side originating the call.
- *Call reference flag* identifies which end of the data link originated the call.



**Figure 9.9** The components of the Q.931 message are called *information elements*. Depending upon the type of message, some information elements are mandatory, while others are optional.

Q.931 groups messages into four categories:

- *Call Establishment* (examples Connect, Setup, Alerting).
- *Call Information* (examples Resume, Suspend).
- *Call Disconnection* (examples Disconnect, Release, Restart).
- *Miscellaneous* are used to maintain and control the network connection; examples are Facility, Notify, Status.

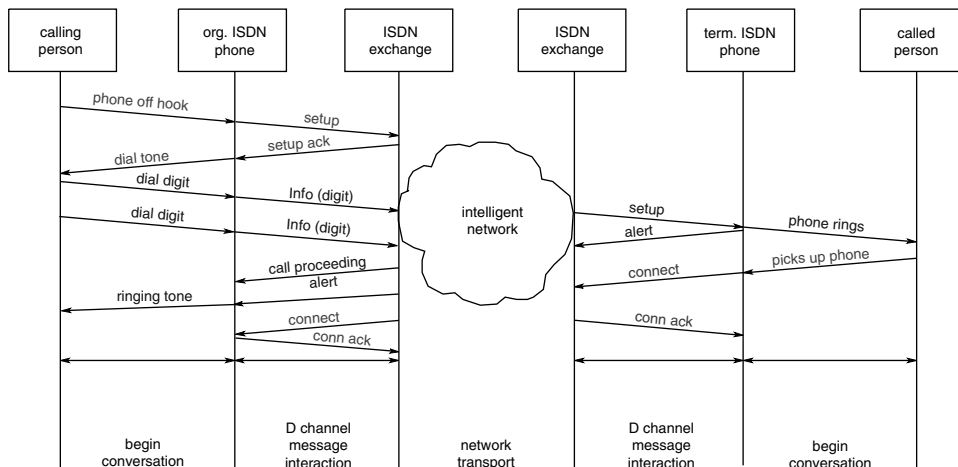
Figure 9.10 illustrates a possible ISDN call setup procedure with Q.931.

### 9.3 Tools and Measurements for ISDN Networks

Testing of ISDN devices and networks occurs in three main areas relating to the design and deployment of these devices: research and development, installation and commissioning, and maintenance.

#### 9.3.1 Users of ISDN test tools

The first category of ISDN test-tool users includes ISDN equipment and network developers, who are involved in the design of Customer Premise Equipment (CPE) such as terminals, TAs, NT devices, and CO switches. Usually the design process is divided into stages or pieces; development testing occurs at many stages in this. Each stage or piece must be tested for functionality and to verify that the design is operating properly.



**Figure 9.10** This diagram describes a typical sequence of Q.931 messages transmitted on the D channel to establish a voice call on a B channel. Depending on the implementation of Q.931, additional or fewer messages may be sent.

Toward the end of the development cycle, implementation testing is performed on the entire design. Implementation testing consists of verifying that the design meets the ISDN standards it incorporates: performance, conformance, and certification tests. This ensures that the design will interoperate with other equipment when it is installed. Protocol testing is a key part of this, using conformance test suites, etc. As a design moves to manufacturing, other tests may be performed to verify functionality and ensure that quality units are shipped to customers.

The second and third categories of people in ISDN testing includes network and equipment installers and maintenance groups. These people are responsible for installing the ISDN networks and equipment, debugging the networks, and keeping them up and running. Most organizations have an escalated or tiered support philosophy. Relatively simple tasks of installing equipment are accomplished by less-skilled personnel who use simple, low-cost tools; this often is termed *Tier 1 support*. If problems occur that are beyond the scope and skill set of the technician, then the technician would call the next level of support, often called *Tier 2*. These people will have more comprehensive training and experience and more sophisticated test tools, such as protocol analyzers. There might even exist a Tier 3 support team, who are considered experts in their disciplines.

During equipment installation, tests can be performed to verify the functionality of the equipment being installed. These tests can consist of matching the equipment configurations to the network parameters. Tests also may be performed to resolve interoperability and compatibility issues with new equipment. After the equipment has been installed and is running, the network maintenance group will troubleshoot problems by isolating the cause of the problem and checking out the faulty equipment. People involved with installation and maintenance testing are usually the network or data technicians.



### 9.3.2 ISDN problem solving

Equipment and network installation and maintenance for ISDN involves similar concepts and problems as research and development, particularly when problems are protocol-related. For this discussion, the ISDN problem solving will be mainly focused on installation, commissioning, and maintenance of ISDN equipment and networks.

Problems with ISDN equipment and networks can occur at layers 1, 2, or 3. It is useful to separate these problems into *connection problems* and *configuration problems*. As with non-ISDN circuits, there are three basic areas of testing for ISDN circuits.

- Line testing
- Transmission testing
- Protocol testing

*Line testing* normally means testing the metallic line itself, if required or desired. It is used to determine loading, proper impedance and continuity of signal path. *Transmission testing*, as the name implies, checks the quality of the transmission facilities, examining types and numbers of errors that occur within specified time periods. *Protocol testing* checks and verifies the proper logical flow of information according to the rules of the specified protocol.

These testing areas for ISDN lines are approached somewhat differently compared to traditional lines. Protocol testing becomes a much more important part of testing ISDN circuits than with non-ISDN lines. Protocols are used for every function, including voice communications. Experience has shown that a significant portion of ISDN problems and trouble are involved with layer 2 and layer 3 protocols.

**Physical layer problems.** Table 9.3 summarizes some of the potential causes and diagnostic tests that can be performed to isolate and identify the problems. Bit Error Rate Testing (BERT) and INFO state analysis are two key tests for determining

**TABLE 9.3 ISDN Physical Layer Testing. Typical connection and configuration problems for Physical layer testing can be addressed with an analyzer that has both BERT and INFO State analysis capabilities.**

Connection Problems	
Potential Cause:	A break in the physical connection A downed subscriber line Improper or incorrect cabling Failure to plug into correct jack
Diagnostic Tests:	BERT Attenuation Continuity
Configuration Problems	
Potential Cause:	The interface of CPE or switch may not be operating correctly
Diagnostic Tests:	Info State Status Power Supply sources

Physical layer problems. An analyzer that can monitor the status of the INFO states for BRI ISDN lines, activity on the B and D channels, and status of the power states will help to isolate these problems. Using the BERT capabilities of the analyzer, line quality and conditioning also can be determined.

**Data Link Layer problems.** The most basic layer 2 tests look for Physical layer problems that did not show up in layer 1 testing. Layer 2 information such as bad Frame Check Sequences (FCS) indicates bit errors during transmission. Frame reject reports of an error condition indicate poor digital subscriber line quality.

The next aspect of layer 2 testing looks at configuration issues and errors. These include:

- Proper and consistent assignments of Terminal Endpoint Identifiers (TEIs) and Service Access Point Identifiers (SAPIs). The LAPD protocol utilizes a Data Link Control Identifier (DLCI) that contains a SAPI and a TEI.
- Proper configuration of the Subscriber Profile Identifier (SPID) in the ISDN equipment and in the ISDN switch. (SPIDs are required only in the United States). The SPID interaction occurs after the TEI and SAPI negotiations.
- Timing measurements. When a prompt is sent out, is the appropriate response returned within the correct amount of time? Does this vendor implement the layer 2 handshaking in the same as another?
- Protocol conformance testing that verifies that the proper Q.921 (LAPD) procedures are followed, such as link setup, information-frame transfer, and link disconnection.

Table 9.4 summarizes these points.

**Network Layer problems.** Although the core of layer 3 is defined by the CCITT Q.931 standard, different switch manufacturers have gone beyond the basic definitions and implemented different extensions. For Network layer testing, consider the connection and configuration aspects of layer 3.

*Connection testing* should verify that the proper procedures or protocol for Q.931 are occurring. This testing will uncover incompatible implementations of layer 3 message interactions including call establishment, message transfer, and call disconnect.

**TABLE 9.4 Data Link Layer Testing. Typical connection and configuration problems for Data Link layer testing could be handled by an analyzer that can detect frame reject errors and allows for protocol testing to verify proper adherence to Q.921.**

- 
- |                  |  |
|------------------|--|
| 1. Connection    | –Data Link indicators of physical problems; Bad FCS indicate bit errors during transmission, Frame rejects report potential poor line quality.   |
| 2. Configuration | –Proper and consistent assignments of TEIs and SAPIs (DLCI) and SPID at both the CPE equipment and the ISDN switch side<br>–Timing measurements; time-outs and message interchanges occurring properly?<br>–Protocol conformance; all of the multi-vendor ISDN equipment following the LAP-D procedures correctly. |
-

**TABLE 9.5 Network Layer Testing. Typical connection and configuration problems for Network layer testing require equipment that can verify conformance to Q.931. It should also provide timestamping so that disconnects can be investigated.**

---

1. Connection	<ul style="list-style-type: none"> <li>-Verify the proper call procedures or protocol for Q.931 are occurring</li> <li>-Verify that the B channel data or voice is working correctly</li> </ul>
2. Configuration	<ul style="list-style-type: none"> <li>-Check that the SPID (Subscriber Profile Identifier) is configured correctly and accepted.</li> <li>-Timing; interaction of prompts (Alerting) and responses (Setup)</li> <li>-Protocol conformance and interoperability variance among ISDN equipment manufacturers and international implementation</li> </ul>

---

Once the D channel signaling has established a connection over a B channel or over X.25 on the D channel, an analyzer may be used to diagnose problems with the protocol being used on the B channel, or with the D channel X.25 link. A protocol analyzer can:

- Verify that voice connections on the B channel are working in both directions.
- Make sure B channel circuit-switched data transfer is functioning properly.
- Verify the operation of the protocol on the B channel (e.g., X.25 packet data).

*Configuration testing* should verify that the SPID is configured correctly in the ISDN CPE and that the ISDN switch accepts it. It should also verify timing; if a particular response (such as alerting) is not received within the required amount of time from the prompt (setup), the call may be disconnected. A protocol analyzer will display timestamps along with the decoded messages. Finally, configuration testing should verify that the protocol implementation is set up properly and operating correctly in both the ISDN CPE and the ISDN switch. Protocol analysis will reveal any interoperability and configuration incompatibilities.

Table 9.5 summarizes these points.

### 9.3.3 Categories of ISDN testing

The preceding discussion of the various layers now can be put together to formulate a higher-level view of ISDN testing. Consider three categories of problems encountered by ISDN users and support personnel: installation and commissioning problems, maintenance problems, and user data problems.

**Installation and commissioning problems.** When ISDN service and equipment is installed and commissioned (turned up) at a user's location, there usually are first-time problems encountered. Table 9.6 is a summary table for a few typical ISDN installation and commissioning problems and tests to identify and isolate these problems.

**Maintenance problems.** An installed and operational ISDN equipment and system might simply stop working altogether. Table 9.7 is a summary table for some typical maintenance problems and tests to identify and isolate these problems.

**TABLE 9.6 ISDN Testing: Installation and First-Time Problems. This table summarizes typical problems, potential causes, and tools required to address problems of ISDN installation and first-time problems experienced as the network is turned up.**

Problem	Potential Causes	Tools Required
1. Call rejected by local carrier	Misconfiguration at CPE and/or switch, such as Subscriber Profile Identifier	a. Monitor with full decode b. Place call (substitute)
2. Call rejected by long distance carrier	Various	a. Full decodes including Cause Codes
3. Call rejected by called country	a. Addressing differences b. Other incompatibilities across countries	a. Monitor with full decode b. Place call (substitute)
4. Able to place basic calls, but some features don't work	Misconfiguration at CPE and/or switch	a. Flexible simulation with user-modifiable scripts
5. User can't get CPE to place a call	a. Misconfiguration b. New user/user error	a. Monitor with full decode b. Place call (substitute)
6. Verify installation, including line quality	Normal operating procedure	a. Pre-defined (canned tests) b. Simulate (place calls) with BERT c. Customize canned tests

**TABLE 9.7 ISDN Testing: Maintenance. This table summarizes typical problems, potential causes, and tools required to address problems of ISDN maintenance, or "what to do if the network stops running."**

Problem	Potential Causes	Tools Required
1. No longer able to place calls for no apparent reason	Software upgrades at telco switch causes interoperability problems	a. Monitor will full decodes b. Analysis of different vendor and country specific ISDN standards
2. No longer able to place calls after service upgrade	a. Incompatible services indicators b. Other configuration changes	a. Monitor with full decodes b. Place calls (substitute)
3. No longer able to place calls after change to "standard"	Different implementation of ISDN with different manufacturers or different countries	Monitor with full decodes
4. First generation ISDN equipment is "buggy"	Explosive growth of ISDN, many new vendors of CPE equipment	Monitor with full decodes
5. Look for intermittent problems	Various	a. Flexible triggering b. Statistical analysis c. Remote control/testing
6. Resolve different manufacturer implementation	Interoperability/implementation incompatibilities	Reputable, unbiased test vendor

**User data problems.** Problems can arise when user data is not making it through to the endpoint or is somehow being adversely affected. Table 9.8 is a summary table for some typical user data problems and tests to identify and isolate these problems.

Of all the types of problems, the most commonly encountered will usually be either Physical layer wiring issues or configuration issues, involving ISDN equipment parameters or subscriber profiles (like SPID). Once the ISDN service is operational, then most problems will occur at the Network level (layer 3) or involve in-depth analysis of the user's data.

### 9.3.4 Tools and measurements

Depending on the size of the organization and the expertise of either the users or the support personnel, different test tools and measurements are available for troubleshooting ISDN equipment and services.

- Equipment swapping
- Embedded diagnostics
- Handheld testers
- Protocol analyzers
- Personal computers/laptops

The easiest and quickest method of troubleshooting is simply to swap out the suspected ISDN equipment and see if the new equipment will work. Through the process of trial and elimination, the faulty equipment can be isolated and identified.

The ISDN switch in the Central Office is likely to have some level of diagnostics built in to the switch. An example of these diagnostics could be continuous monitoring of individual BRI lines for statistical performance parameters. This information could be accessed by support personnel from a computer terminal for diagnosis. The CPE might have built-in diagnostics that communicate status via LEDs or LCD displays. A quick look at the LEDs could reveal the source of the problem.

**TABLE 9.8 ISDN Testing: User Data Problems. This table summarizes typical problems, potential causes, and tools required to address problems of ISDN user data problems experienced when the end user attempts to access the network or experiences performance problems once connected.**

Problem	Potential Causes	Tools Required
1. Cannot establish IP connection over ISDN	a. Wrong encapsulation b. Protocol mismatch	LAN decodes over the B channel
2. Unable to establish X.25 connection to far end	X.25 Parameters not configured properly	Monitor with X.25 decodes over the D channel or over the B channel.
3. Lower than expected data throughput on an ISDN router	The PPP or MLPPP connections are not operating correctly	Monitor with PPP or MLPPP decodes over the B channel(s)

Low-cost, simple-to-use, rugged handheld testers, sometimes called *butt-in sets*, are typically provided to the Tier 1 installation and commissioning team. These devices are analogous to the analog butt-in set used by telephony technicians to listen for dial tone. With ISDN handhelds, there is no analog dial tone, but the sets are capable of placing simple voice or data calls, performing BERT tests, and determining simple configuration issues.

If the problem cannot be identified by the handheld testers, then the problem is passed up to the next level of support personnel; generally a protocol analyzer is required at that point. A protocol analyzer is capable of connecting to the BRI or PRI lines and performing monitoring functions of the B and D channels using comprehensive decodes, filter, triggers, and searching capabilities. If necessary, it is also capable of simulating or placing ISDN voice or data calls. This will allow isolating the ISDN device and determining if it is implementing the ISDN protocol correctly. Bit Error Rate Testing (BERT) can be used to determine if the ISDN link is meeting performance specifications. The protocol analyzer also can access the user data on the B channel and determine if and where in the user data the problem may be originating.

Testing personnel also might carry a PC for logging into databases to obtain and report on trouble tickets on which they are working. This PC might have ISDN capabilities that would allow it to be substituted for suspect ISDN equipment.

The primary requirement for ISDN test tools is that they be reliable, rugged, portable, and offer a comprehensive range of capabilities that will address layer 1, 2, and 3 problems. Tools should be well-suited to flexibly accommodate problems of installation, maintenance, and user data. A comprehensive list of ISDN test tool requirements should include:

- Basic and Primary Rate access
- Simultaneous monitoring and simulation
- Monitoring with full decodes
- INFO state (layer analysis)
- Full-featured Bit Error Rate Tests (BERT)
- Statistical analysis
- Prewritten tests (for ease of use)
- Simulation to provide ISDN call placement on D channel
- Comprehensive protocol analysis on the B channel
- Voice access to handset on B channels for voice quality monitoring
- Wide range of protocol testing support for other datacomm needs, such as LAN, WAN, and ATM

## 9.4 Summary

ISDN provides a higher bandwidth access to the Intelligent Digital Network. ISDN satisfies some important user demands:

- Greater bandwidth for higher-speed applications
- Greater reliability, better performance, and security
- Simple connection and interaction with the network
- Cost savings

There are many competing services available that can provide similar functionality as ISDN. This is the motivating force for providers of ISDN to price and provide ISDN services competitively.

ISDN is widely available now from service providers. Problems persist, however, not only with getting service installed, but also ensuring interoperability and compatibility among diverse ISDN components from diverse vendors. ISDN is not yet plug- and-play. Service providers and carriers must meet performance criteria and ensure that interoperability and compatibility. End users want to make sure they get the kind of service they pay for, in turn providing the level of performance their internal corporate customers expect.

Test tools for ISDN networks can range from hand-held Bit Error Rate Testers to high-end conformance protocol testers. Having the proper test tools will help to ensure optimum operation of ISDN networks, as well as user satisfaction.





# Broadband Communications and Asynchronous Transfer Mode

**Stewart Day**

*Hewlett-Packard Australia Ltd., Victoria, Australia*

## 10.1 Introduction to Asynchronous Transfer Mode (ATM)

Asynchronous transfer mode (ATM) is the base switching technology that the ITU-T (formerly CCITT) chose for the Broadband Integrated Services Digital Network (B-ISDN). Though it was originally envisioned that ATM be deployed once the new broadband services appeared, industry (particularly through the ATM Forum) has accelerated the practical application of the technology to the point where ATM itself is becoming the enabler for new services. In addition, ATM is being used to solve existing problems such as LAN interconnection and the continued evolution of MANs and WANs (metropolitan area and wide area networks). ATM can be deployed in all types of local and wide area networks.

### 10.1.1 ATM basics

ATM is a cell-based architecture. Cells are short, fixed length-packets. ATM cells, each with a fixed length of 53 bytes, can be processed more efficiently than bit streams or variable-length packets. This efficiency allows high-bandwidth switching and multiplexing in both local and wide area environments.

### 10.1.2 Multiplexing incompatible traffic types

Cell technology provides a powerful mechanism for transporting different traffic types on a common communications infrastructure. By keeping cells small and regular, real-time traffic (such as voice and video) can be segmented and packetized, then multiplexed onto pathways carrying other traffic types (such as data) without loss of quality.

### 10.1.3 Switching, multiplexing, and bandwidth-on-demand

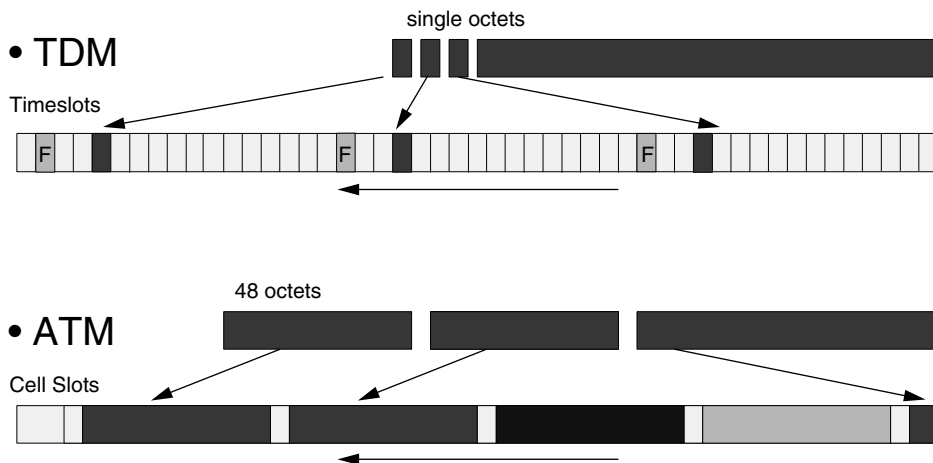
A second advantage of ATM is the ability to maximize switching efficiency and bandwidth usage. In existing Time Division Multiplexed (TDM) networks, data is divided into octets and placed in fixed timeslots in the transmission stream. Such a framing structure provides only fixed-bandwidth services and is inefficient to switch.

In an ATM network, the ATM layer uses a contiguous stream of fixed-length cells (Figure 10.1). Fixed-length cells enable faster switching, multiplexing, and bandwidth-on-demand. Cell switching and multiplexing systems are implemented directly into hardware instead of having to be manipulated in software. Each cell also contains addressing and control information, so that switching functions can also be implemented in hardware. From this architecture are coming increases in performance by an order of magnitude over older systems that use software to figure out where to route data.

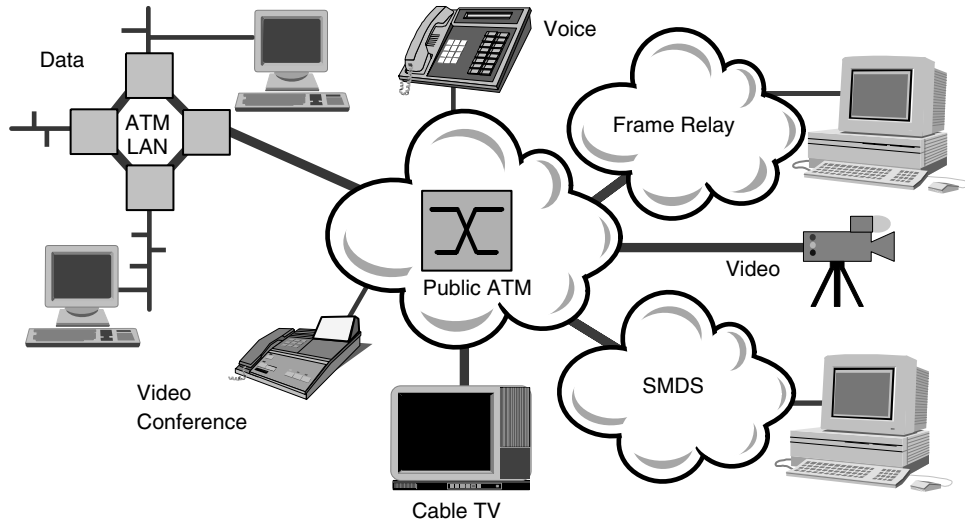
### 10.1.4 Broadband services

*Broadband services* are those requiring bandwidths above 2 Mbps. These services include many existing and future communications areas:

- Interconnection of LANs, MANs, and WANs
- High-speed data transfer
- Video phoning and videoconferencing
- Broadcast of standard and high-definition television
- Broadcast of digital audio



**Figure 10.1 ATM compared to TDM.** In traditional Time Division Multiplexed (TDM) networks, each service is segmented into single octets, which are allocated to a timeslot within every frame of the digital stream. This is ideal only for low-bandwidth, constant bit rate services. In an ATM network, each service is segmented into 48 octet packet payloads, which are multiplexed into the digital stream only as required. This is much more efficient for services with variable bit rates and is scalable from very low to very high bandwidths.



**Figure 10.2 The ATM network.** ATM can be deployed effectively in the core of public carrier networks, carrier access networks, residential access networks, enterprise backbones, and all the way to the desktop. ATM also is capable of transporting and interworking with existing LAN and WAN technologies such as Ethernet, frame relay, and SMDS. ATM has been designed to be an integrated network technology that can meet the diverse quality needs of data, voice, video, and multimedia services.

- Library access to video and audio material
- Database access
- Interactive multimedia

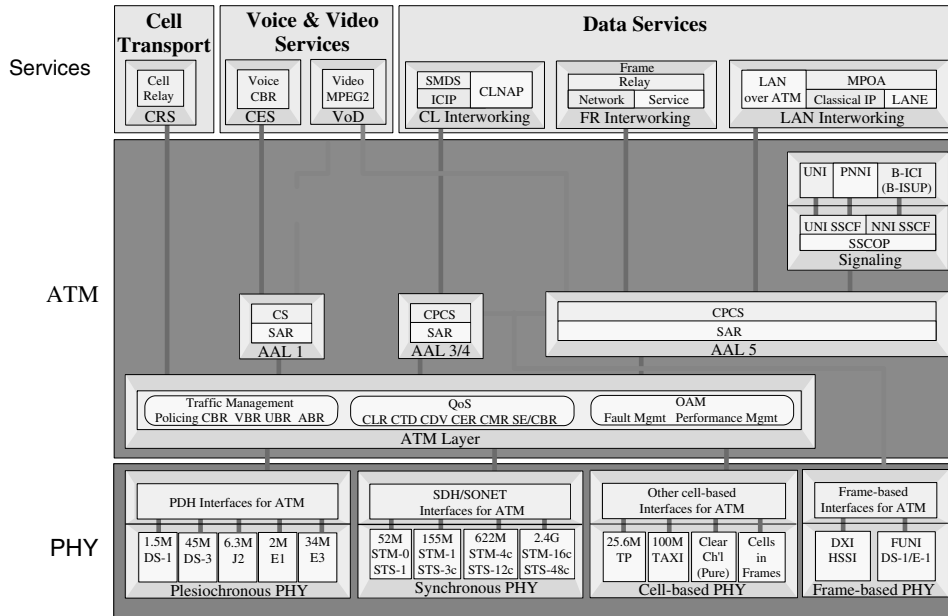
In addition to supporting all of these and other broadband services, ATM also can also handle efficiently services with bandwidths below 2 Mbps.

## 10.2 The ATM Communications Network

The broadband environment requires interconnecting existing private and public networks with the new networks that will provide the services of the future (Figure 10.2). As a communications infrastructure, ATM is ideally suited to be used both as a backbone network (interconnecting LANs, MANs, and WANs), as a means for high-bandwidth user connections such as multimedia applications—or even ATM LANs themselves.

### 10.2.1 ATM protocol map

Figure 10.3, the ATM communications map, shows the interrelationship of ATM with the various network services. The map is split into three major hierarchical levels: Physical, ATM, and Services. The Services level includes the key technologies planned or being carried over ATM networks, primarily data, voice, and video services. In the center of the map are the key technologies of ATM: the ATM Cell layer, which provides transport of service data; the AALs (ATM Adaptation layers), which



**Figure 10.3 The ATM protocol map.** ATM is a flexible technology that can carry a wide range of both new and existing service types over a wide range of new and existing physical interface rates and types. The protocol map shows the ATM layer to be the common denominator in this diverse network environment.

adapt service data into and out of ATM cells; and the network control procedures, which allow operation and management of the network. At the bottom of the map are the key physical network technologies specified for carrying ATM cells, split into public telecom technologies and private enterprise technologies. Also included are frame-based interfaces that allow the transport of ATM data without segmenting it into cells.

**Physical level.** The Physical level of the map shows the most common standard physical interfaces on which ATM can be used. These include not only cell-based specifications to carry standard 53-byte ATM cells, but also frame-based interfaces that allow the variable-length AAL frame structures to be transmitted directly without segmentation into cells.

ATM standards are defined for the most popular public network interfaces throughout the world, including both the latest optical SONET/SDH interfaces and the older electrical PDH interfaces used in each world region. Currently these interface specifications range from DS1 at 1.5 Mbps to OC-48c/STM-16c at 2.488 Gbps, but it is feasible for ATM to be carried at both lower and higher rates as well.

In addition to public network technologies, a series of lower-cost interfaces have been specified for the private enterprise. With the flexible nature of ATM, almost any interface rate and media can be used. Standardized interfaces include lower-cost variants of SONET at 51 and 155 Mbps over multimode fiber and UTP (unshielded twisted-pair) copper. Also standardized are interfaces based on LAN technologies, such as a 100 Mbps (TAXI) based on FDDI, and a 25.6 Mbps UTP based on Token-

Ring. Many other interface types, such as Cells in Frames and Clear Channel, are also being proposed in an effort to find the most appropriate method of access for particular services. Other initiatives, such as wireless ATM, will result in the appearance of further ATM interfaces.

Finally, there also are frame-based interfaces specified to transport ATM structures, such as HSSI (High-Speed Serial Interface) and FUNI (Frame User Network Interface). Instead of transporting cells, these interfaces transport the variable-length AAL Common Part Convergence Sublayer (CPCS) structures used with data services, which avoids the processing overhead of segmenting and reassembling the frames into ATM cells.

**ATM level.** The ATM level of the map contains the core ATM protocols: the ATM layer, ATM Adaptation layer (AAL), and connection control through signaling.

*The ATM layer* is responsible for managing the transport of ATM cell streams through the network of ATM switches. The 53-byte cells contain a 5-byte header for identification, routing, and control information, and a 48-byte payload to carry the service data. Services are allocated bandwidth on demand by using only the number of cells they require, as opposed to reserving a fixed bandwidth as with a TDM network. Specific protocols at this layer are discussed in more detail in subsequent sections.

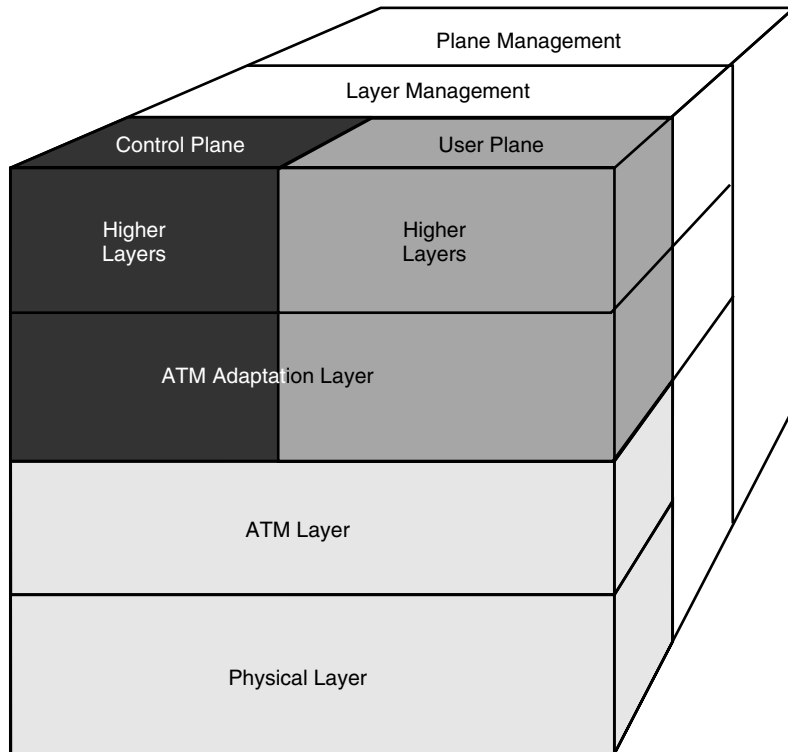
*The ATM Adaptation layer (AAL)* is responsible for the segmentation and reassembly (SAR) of service data to fit the fixed-length cell payloads. Different AALs are defined to meet the quality of service (QoS) requirements of the different service types that can be carried on ATM. Each AAL is optimized for a particular type of service. As new services are developed, a need might develop for new AALs. The specific AAL protocols are described later.

Before any ATM cells can be sent through the network, an *ATM connection* must be set up to allocate a virtual channel (VC) and virtual path (VP) for the cell stream. The VP and VC are hierarchical, with each VP containing a large number of VCs. This gives the network flexibility to group channels together in the same path and allows switching to be done at either level. ATM signaling and connections are described in subsequent sections.

**Services level.** The services level of the map shows the main services being specified to use an ATM network. These range from the basic cell relay services to methods for carrying voice and video over ATM and methods of using ATM in data networks. With data services, ATM can be used for LAN connections to the desktop, as a LAN backbone and in the WAN. The protocols also allow ATM to interwork with other LAN and WAN protocols in a mixed network environment.

## 10.2.2 The B-ISDN protocol architecture

The B-ISDN protocol reference model is defined in ITU-T Recommendation I.121. It calls for a three-dimensional, layered architecture of both planes and layers (Figure 10.4). The higher, or service, layers are applications such as frame relay, SMDS, LAN, TCP/IP, SNA, video, or voice.



**Figure 10.4 The ATM protocol architecture.** The protocol architecture of ATM was developed by the ITU-T (formerly CCITT). It shows that while network control and user data have individual higher-layer and Adaptation layer protocols, they are both integrated at the ATM layer and transported over a single physical network.

The planes of the model consist of:

- User Plane
- Control Plane
- Management Plane

The *User Plane* provides for the transfer of user application information. It contains the Physical layer, ATM layer, and ATM Adaptation layers that enable ATM to support different services and applications. Protocols in this plane provide control information, including flow control and error control and correction.

The *Control Plane* includes the functions for providing switched services. It performs the signaling necessary to setup, supervise, and release calls and connections.

The two-part *Management Plane* provides Layer and Plane management functions to support the User and Control planes. Plane Management coordinates functions between all of the planes. Layer Management handles the flow of operation and maintenance information between the layers.

The User Plane is the transport mechanism for higher-layer (user) information. It is made up of the Physical, ATM, and ATM Adaptation Layers, which correspond to the map in Figure 10.4.

### 10.3 The Physical Layer

#### 10.3.1 Cell based physical layer

Cell-based framing is the purest form of ATM. Unframed ATM is completely scalable, meaning that it can be carried over any transmission medium at any suitable rate the user desires. The ITU has specified transmission at the same rates and interfaces as SDH/SONET at 155 Mbps and 622 Mbps.

#### 10.3.2 SDH/SONET interfaces


Along with cell-based ATM, SDH/SONET (Synchronous Digital Hierarchy and Synchronous Optical Network) was specified originally as the standard carrier of ATM cells. The standards specify concatenated SDH/SONET payloads (STS-3c, STM-4c, etc.) to allow the full bandwidth capacity to be available for a single service, allowing 600 Mbps on a single channel over STM-4c/STS-12c.

Figure 10.5 shows how ATM cells can be carried by an SDH/SONET frame. ATM cells are carried within 8 kHz frame structures at the standard SDH/SONET bit rates. Overhead octets provide framing and OAM, and the HEC (checksum) is used to identify cell boundaries.

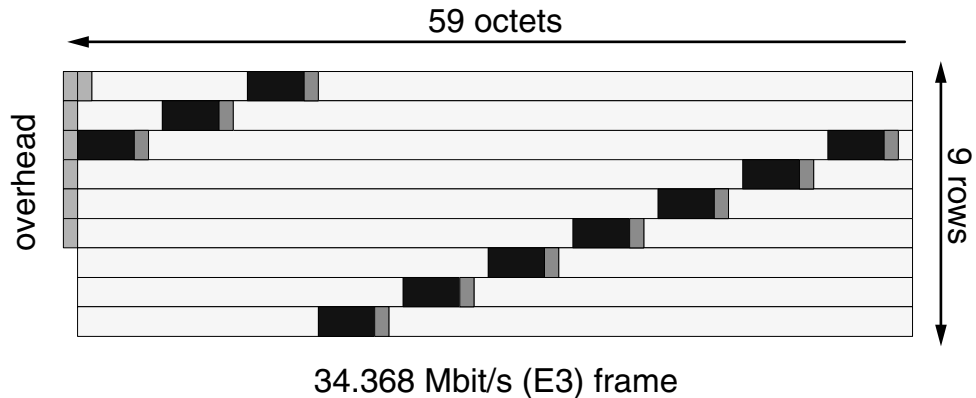
#### 10.3.3 PDH interfaces

The transfer of cell streams over PDH networks was developed by Bellcore for SMDS/DQDB over DS1 and DS3 systems. Because ATM cells are the same length as SIP L2 PDUs, an identical frame structure has been used for ATM: the Physical Layer Convergence Protocol (PLCP). PLCP, however, has a large framing overhead, which reduces considerably the available cell bandwidth.

When it came to specifying a framing method for the European PDH rates, for that reason it was decided to use a more efficient synchronous frame, similar to that used in SDH. In addition, the new frame structures include frame octets, so the usual PDH



**Figure 10.5 SDH/SONET 155 Mbps framing.** ATM cells can be transported in standard concatenated SDH/SONET frames at rates from STS-1/STM-0 (51 Mbps) to STS-48c/STM-16c (2.5 Gbps) and beyond. The contiguous cell stream is transported in the SPE/VC payload. The cell stream is octet-aligned with the SONET/SDH frame, but the cells are not in a fixed position within the frame.



**Figure 10.6 PDH 34 Mbps framing.** ATM cells can be transported over the most common PDH interface rates in North America, Europe, Japan, and the rest of the world from DS1 (1.544 Mbps) to E4 (140 Mbps). The E3 (34 Mbps) frame has a payload plus overhead that contains framing and management data, whereas with SONET/SDH the cell stream is octet-aligned but the cells are not in a fixed position in the frame.

framing bits are not needed (unlike PLCP, where they are still used). All of these framing structures, including direct mapping for DS-1 and DS-3, are specified in ITU-T recommendations G.8041/G.832. Figure 10.6 shows ATM cells being transmitted within a 34.368 Mbps (E3) frame.

### 10.3.4 Private network interfaces

The ATM Forum has defined interfaces for private ATM networks. Except for the IBM UTP interface, the following are defined in the Forum's UNI version 3.1/4.0:

- ATM cells at private UNI
- 25.6 Mbps UTP, asynchronous
- 51.8 Mbps UTP (STS-1), synchronous
- 100 Mbps multimode fiber (TAXI), asynchronous
- 155 Mbps UTP (STS-3c), synchronous
- 155 Mbps multimode fiber (Fiber Channel), synchronous

Originally, block-encoded multimode fiber interfaces were specified with a 100 Mbps interface using the FDDI (TAXI) PMD and a 155 Mbps interface using the Fiber Channel PMD (Figure 10.7). In both systems the bit stream is block-encoded, using line codes, to a higher line rate. To reduce deployment costs, UTP (unshielded twisted-pair) interfaces recently have been specified for 52 Mbps and 155 Mbps using SONET framing, and 25.6 Mbps using block encoding.

### 10.3.5 Frame based interfaces

The ATM forum has also defined interfaces for ATM at the frame level. These transport data at the AAL CPCS level of AAL 3/4 or AAL 5 as variable length frames.




High Speed Serial Interface (HSSI) uses DXI (Data Exchange Interface) to transfer frames at rates up to 50 Mbps. HSSI was developed for short connections in SMDS networks and also can support ATM.


Frame User Network Interface (FUNI) transports service data across low speed data links (DS-1, E-1 and below, including  $N \times 64$ ) without the overhead of the ATM cell header or interworking with Frame Relay protocols. FUNI (see Figure 10.8) is specified for AAL 3/4 and AAL 5 and supports VBR and UBR traffic types. The FUNI header contains a restricted version of the cell level VPI/VCI and supports congestion and cell loss priority indicators. Because the AAL structure is used, ATM network management and signaling procedures can be supported directly.

### 10.4 The ATM Layer

The ATM layer provides cell-level management, routing, traffic control, and multiplexing.




**Figure 10.7 Block Encoded framing at 100 Mbps.** ATM cells can be transported over any digital network technology. To reduce cost for enterprise networks over PDH and SONET/SDH, standards have been developed to carry ATM cells over multimode fiber and twisted-pair cabling. These are closely based on existing LAN interfaces such as Token-Ring and FDDI. The 100 Mbps Block Encoded interface (also known as TAXI) was one of the first of these private interfaces. It uses the multimode fiber and physical line coding of FDDI. Instead of being transported contiguously, idle and unassigned fill cells are removed and the remaining cells are transported asynchronously with an additional beginning of cell line code. When cells are not being transported, idle line codes are transmitted through the network. The 25.6 Mbps Block Encoded interface uses a similar technique, but over unshielded twisted-pair (UTP) cabling.



**FUNI** Frame User Network Interface  
**CPCS** Common Part Convergence Sublayer  
**SSCS** Service Specific Convergence Sublayer  
**FH, FT** Frame Header, Trailer

**Figure 10.8 Frame UNI for ATM.**

## 226 Wide Area Networks



**Figure 10.9 The ATM cell.** The ATM cell is a fixed-length packet with a 5-octet header and 48-octet payload to carry service data.

#### 10.4.1 The ATM cell

The ATM layer is a constant stream of cells. Each cell consists of a 48-byte payload and a 5-byte header (Figure 10.9). The header contains the switching and routing information required by the ATM layer. The payload carries the higher-layer service information, which previously has been segmented into fixed-length blocks using the ATM Adaptation layer protocols.

ATM layer protocols are concerned solely with the information contained in the header, with the exception of OAM cells used to send management information.

#### 10.4.2 ATM layer interfaces

The ATM Forum specifies multiple ATM layer network interfaces:

- *Public User Network Interfaces (UNI)*, which connect a user or service to a public carrier's ATM network.
- *Private User Network Interfaces (UNI)*, which connect a user or service to ATM LANs and private ATM networks.
- *Network-to-Network Interfaces (NNI)*, which enable two B-ISDN network nodes to be interconnected.
- *Private-Network Node Interfaces (PNNI)*, which enable two B-ISDN nodes to be interconnected in a private network.
- *Broadband Inter-Carrier Interfaces (BICI)*, which enable interconnection between multicarrier ATM networks or operators.

#### 10.4.3 The cell header

The cell header format used at the UNI and NNI differs. The UNI format, shown in Figure 10.10, has a Generic Flow Control field that can be used at the UNI for local functions. The NNI uses these four bits for extra Path addressing (added to the VPI), because an NNI is expected to support many more Virtual Paths than a UNI. Other control fields used in the header are:

- *VPI/VCI* (Virtual Path Identifier/Virtual Channel Identifier), used to identify cells for routing.
- *PT* (Payload Type), used to indicate whether the cell contains user or management information.

- *CLP* (Cell Loss Priority), which identifies which cells the user would prefer to lose if network congestion or policing occurs.
- *HEC* (Header Error Control) is a checksum used for cell delineation.

#### 10.4.4 Cell delineation


*Cell delineation* is the process of aligning cell boundaries within the transmission system (Figure 10.11). There are three states to the process:

- In HUNT, a check is performed every 40 bits. If a valid HEC (checksum) pattern is found, the process moves to PRESYNC. If no valid HEC patterns are detected, the process remains in HUNT.
- In PRESYNC, the 40 bits at each subsequent cell header position are checked for valid HEC. After DELTA valid HECs, the process moves to SYNC. If an invalid HEC is detected before DELTA is reached, the process returns to HUNT.
- In SYNC, every cell is checked. If ALPHA consecutive invalid HECs are detected, the process returns to HUNT. Otherwise, synchronization is maintained. *Single Error Correction Double Error Detection* (SECCDED) also can be used in SYNC.

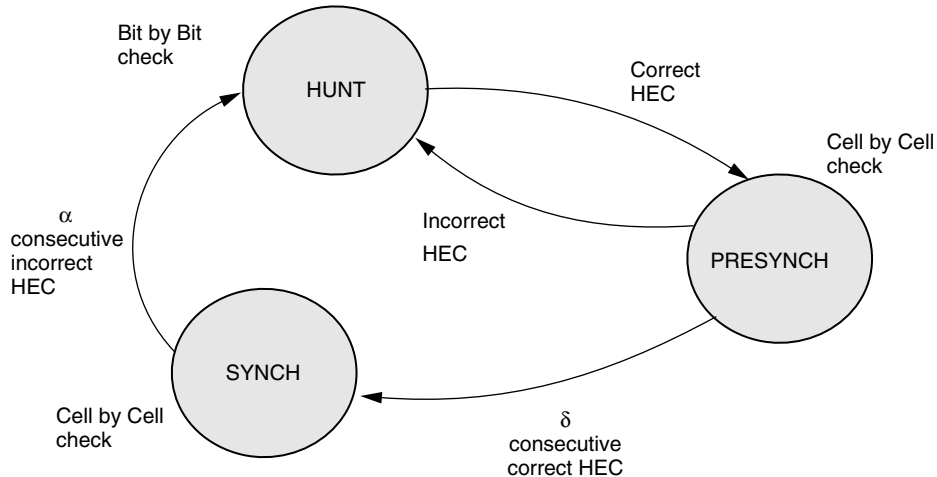
#### 10.4.5 Special-function cell types

Along with those used for service data transfer, there are several cell header patterns reserved for specific functions:

- Assigned
  - VPI/VCIs with user data
  - VPI/VCI OAM cells
  - Signaling cells
  - Management cells
  - Other reserved VPI/VCIs
- Unassigned



**Figure 10.10 The cell header.** The 5-octet cell header contains connection identifiers (VPI, VCI), control information (PT, CLP, GFC), and a header checksum (HEC) that can correct single-bit errors anywhere in the header.



**Figure 10.11 The cell delineation process.** Cell delineation is used to find the start of each cell in the contiguous bit stream. The receiver starts in state HUNT, performing a bit-by-bit check until it finds 40 bits with a valid HEC. The receiver then moves to state PRESYNCH, where it looks 53 bytes along for another valid HEC. If  $\delta$  (usually 6) contiguous cells with good HECs are found, the receiver assumes it has synchronized with the cell stream and moves to state SYNCH. If at any time during PRESYNCH a bad HEC is found, the receiver moves back to state HUNT and restarts the bit-by-bit check.

Once in state SYNCH, the receiver is able to process the cell data, extracting the higher-level AAL and service data for processing. The receiver continues to check the HEC of each cell, with the ability to correct cells with a single bit error in the header. If  $\alpha$  (usually 7) contiguous cells with HEC errors are found, the receiver assumes that synchronization has been lost and moves back to state HUNT.

- Physical layer cells
  - Idle
  - Cell-based physical OAM

OAM flows are specified for VP (F4) and VC (F5) connections. Signaling uses reserved VPI/VCI channels. Management functions also use reserved channels. Some other channels are reserved for future use.

Unassigned cells have a fixed header value. The network also reserves some header values for Physical layer functions such as rate adaptation (Idle) and OAM for cell-based transmission networks.

## 10.5 ATM Traffic Management


### 10.5.1 Cell multiplexing

The fundamental cell stream consists of unassigned cells. Each source can replace unassigned cells with assigned data or control cell channels as required.


When assigned cells are multiplexed from several sources, they must wait in queues for unassigned cell slots not being used by higher-priority cell streams. This results in variable delay on a channel, caused by the traffic distributions of other channels (Figure 10.12). In extreme cases, it can also lead to cell loss, where no unassigned cell slots are available before queue overflow occurs.

Each channel can be identified by its VPI/VCI routing identifier. The VPI/VCI allow two hierarchical levels of switching (Figure 10.13).

At the VP crossconnect level, only VPI values are considered when routing cells and all virtual channels within each virtual path are maintained across the switch. Note that the output VPI value will not necessarily be the same as the input VPI.



**Figure 10.12 Multiplexing cell streams.** Cell streams from multiple sources are multiplexed together at switches and access devices to form a mixed stream of cells. Cells must wait in queues at these devices for an unassigned cell slot before they can proceed; this leads to variable delay. As the number of user cells increases, the number of unassigned cell slots will decrease, causing cells to wait longer in queues. This leads to congestion, increasing cell delays. Under extreme conditions this can cause cell loss when queues overflow.



**Figure 10.13 Virtual paths and virtual channels.** Each service is allocated a virtual connection through the ATM network. Connections can be provided at two levels of hierarchy, virtual paths (VP) and virtual channels (VC). The ATM cell header allocates 8 bits for the VPI (at the UNI), allowing 256 paths (4096 at the NNI). Sixteen bits are allocated for the VCI, allowing 65,536 VCs on each VP. The hierarchy allows ATM switching to be done at both the VP and VC levels, allowing efficient connection allocation and helping network management by enabling, for example, the grouping of VCs with similar quality requirements into the same VP for optimal routing through the network.

**230 Wide Area Networks****10.5.2 Virtual paths and virtual channels**

At the VC switch level, both VPI and VCI are used to determine to which output ports cells will be routed. Both VPI and VCI values may change, but the cell payloads will not be altered.

**10.5.3 Connections and signaling**

Connections can be created either statically as permanent virtual connections (PVCs) or dynamically as switched virtual connections (SVCs) using signaling protocols. ATM connections can be either point-to-point, point-to-multipoint (as with a broadcast service), or multipoint-to-multipoint (as in a videoconference). Routing decisions must find not only the shortest route, but more important, a route through the network that can guarantee all the QoS needs of the service.

Connection setup results in a series of VPI/VCI values allocated to transport the cell stream through the network. The VPI/VCI values are part of the 5-byte cell header and are uniquely allocated to particular services between each switch in the network. A particular service therefore may use different VPI/VCI values between different switches on its route through the network. Each switch maintains routing tables that allow cells to be switched in hardware at high speed, rather than having to be reassembled to the higher-layer service to decide where each should be routed next. Inside an ATM switch, no processing is done on the cell payload, with all management and routing decisions based purely on the cell header.

With ATM's potential to be implemented in both public and private networks, a series of different signaling protocols is being developed to meet the different needs. The UNI protocols support user access to the network, both public and private. The PNNI protocols support signaling between switches in a private network. The Broadband ISDN User Part (B-ISUP) protocols, part of Signaling System 7, support signaling inside public networks and between different public carriers via B-ICI.

**10.5.4 Traffic types**

Once the connection is set up, the service's cells are multiplexed with cells from other services and with unassigned and idle cells into a contiguous stream on each physical link. ATM supports a mixture of constant-bit-rate (CBR) and variable-bit-rate (VBR) services, with the goal of allowing the variable-rate cell streams to multiplex together in such a way that the bandwidth required to transport them is less than the sum of their maximum bandwidth requirements (See Figure 10.14).

This process is known as *statistical multiplexing* and should allow much more efficient utilization than is common in networks that reserve a constant bandwidth even for services that need it only for a portion of the time. In an effort to make negotiating bandwidth and quality guarantees easier for services in which it is difficult to predict future bandwidth requirements (LAN services, for example), other traffic classes, such as unassigned bit rate (UBR) and available bit rate (ABR), have been defined. In the case of ABR, a flow-control mechanism is implemented which tells sources they should send data at a higher or lower rate depending on network conditions.

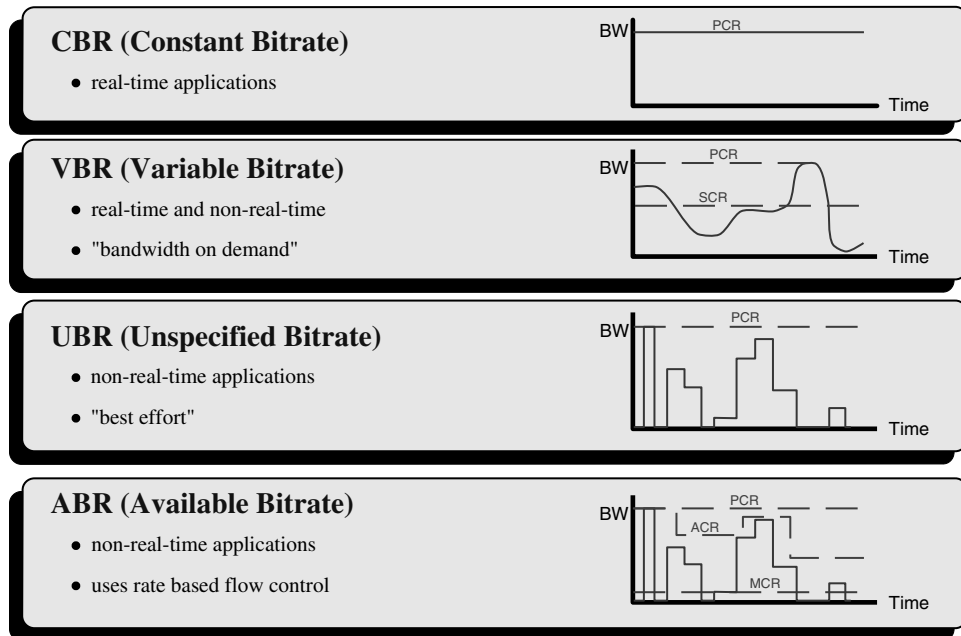


Figure 10.14 ATM traffic classifications.

### 10.5.5 Quality of service

*Traffic classes* define the profile of the traffic that the user's service generates. The users' service also requires guarantees of the quality of service provided by the network. The key parameters defined in ATM are:

- Cell Loss Ratio (CLR)
- Cell Transfer Delay (CTD)
- Cell Delay Variation (CDV)
- Cell Error Ratio (CER)
- Cell Mis-insertion Rate (CMR)
- Severely Errored Cell Block Ratio (SECBR)

As in any other network technology, ATM cells are affected by transmission and switching delays and transmission bit errors. In addition, however, ATM has some unique impairments as a direct consequence of congestion occurring from the statistical multiplexing process. CDV is caused by cells being buffered for varying lengths of time in each switch as they wait for access to the switch fabric or output port. As the load on network links increases, the available capacity decreases; cells spend longer times in the switch buffers. In extreme cases, buffers will fill up and overflow, leading to cell loss. Effects of these impairments will differ on each different type of traffic; delay will be significant to real-time services, for example,

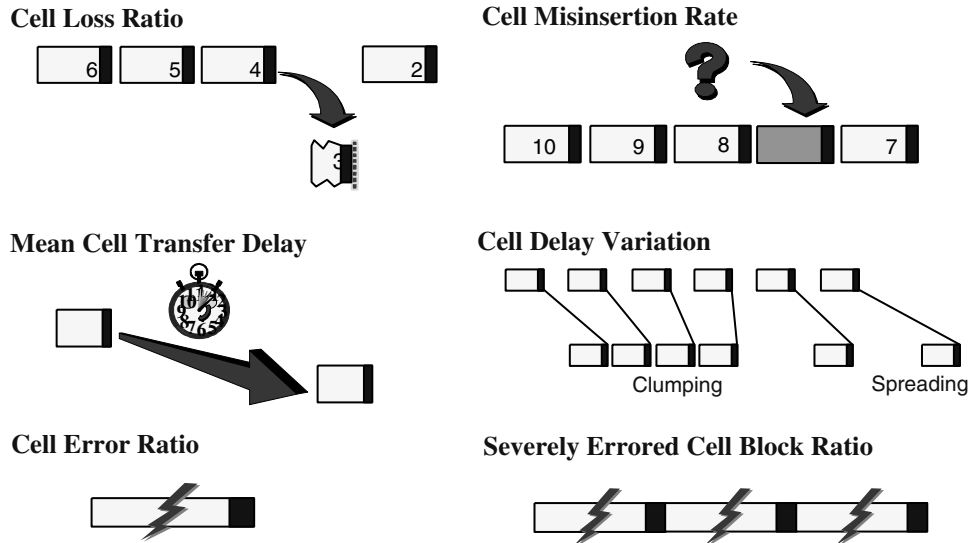


Figure 10.15 ATM quality of service.

whereas cell loss will be much more significant to data services (see Figure 10.15).

The network must guarantee certain QoS parameters for each service, but the user's service also must adhere to the traffic profile that was contracted for. In order to help prevent congestion, the network implements policing checks on the conformance of user traffic to the parameters requested. Cells that do not conform are either tagged or discarded. Tagging occurs if the nonconforming cell is of high priority, in which case it is changed to low priority. Cells are discarded from the network if they already are of low priority prior to tagging, or if tagging is not being used. Cell priority is indicated using the Cell Loss Priority (CLP) bit in the cell header.

### 10.5.6 Operation and maintenance at the ATM layer

At the ATM Layer, Operation and Maintenance (OAM) flows are defined to allow fault and performance management at both the virtual path and virtual channel levels. There are OAM functions at each level of the protocol stack (Physical layer, Transport layer, Services layer). Predefined ATM cells are used for OAM functions at the Virtual Path and Virtual Channel levels. An OAM cell has recognizable header information alerting the network manager to what type of OAM cell it is, and what information will be contained in the payload.

OAM functions include:

- Network congestion point detection
- Defect and failure detection
- Performance measurements
- System protection
- Fault location detection

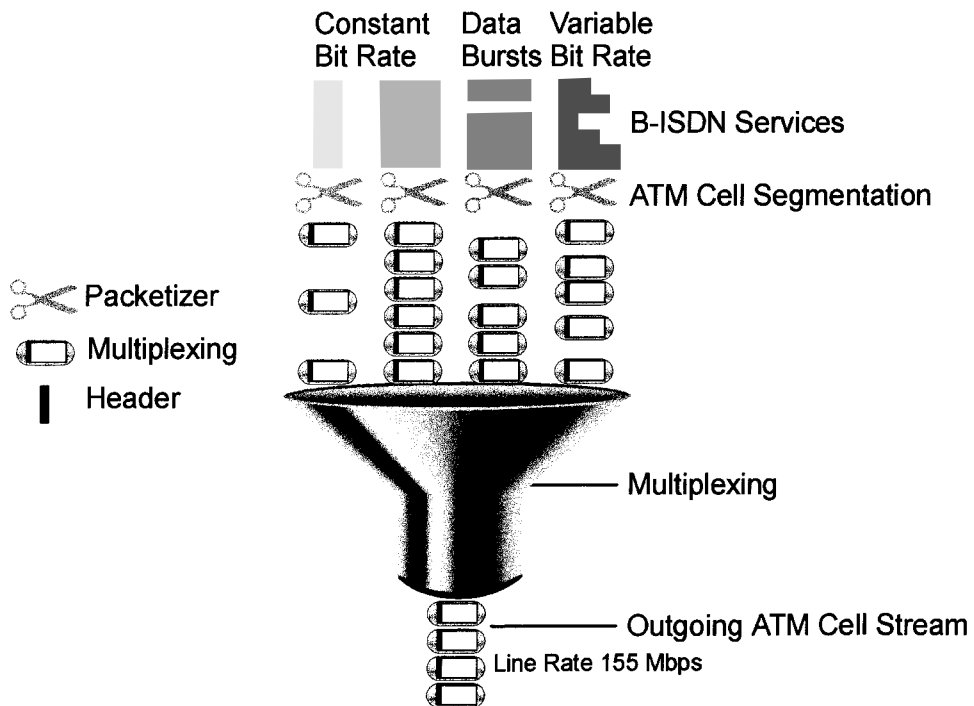


Fault management techniques are based on the OAM functions of SONET/SDH networks that provide alarm notification of faults through *Alarm In Service* (AIS) and *Remote Defect Indication* (RDI) signals. In addition, performance management techniques are defined to allow measurement of bit errors, cell loss, and cell delay for particular paths and channels in the network.

This fault and performance information is gathered by switches at path and channel termination points. It is made available for use by network management systems, together with traffic statistics and information from other measurement devices. As with signaling, the management needs are very different in the private and public parts of ATM networks. Consequently, a whole series of protocols, management information bases (MIBs), and applications are evolving to measure, collect, process, and use the information required to manage an ATM network.

## 10.6 ATM Adaptation Layer (AAL)

The ATM Adaptation layer is responsible for adapting the different types of service data into the fixed-length payload of ATM cells (Figure 10.16). The AAL is the protocol layer between the ATM and services layer. The primary function of this layer is segmentation and reassembly: taking information from the service (such as frame relay) being carried by the network, which can be either variable-length or fixed-rate bit



**Figure 10.16 The ATM Adaptation layer.** The ATM Adaptation layer is responsible for adapting the different types of service data into the fixed-length payload of ATM cells, then checking and reassembling the original data at the far side of the network. Different AAL types are defined to handle different service types.

**TABLE 10.1 ATM Adaptation Layer**

Service class	Class A	Class B	Class C	Class D
Timing compensation	Required		Not required	
Bit rate	Constant		Variable	
Connection mode		Connection-oriented	Connectionless	
AAL type	Type 1	Type 2	Type 3	Type 4
Service example	Circuit emulation	Compressed video	CO data transfer	CL data transfer

- Type 2 defined in I.363.2 (Feb. 1997)
- Type 3 and 4 merged to Type 3/4
- Type 5 also covers Classes C and D

streams, and splitting the information into cells without loss of integrity to the original message. At the receiver, the AAL extracts the information from the cells and turns it back into its original form. In addition, the Adaptation layers take care of adapting rates and cell jitter, as well as error-checking and removing corrupted information.

The ITU has defined four main classes of service (classes A through D) to support the different requirements of the different services; see Table 10.1. From these classes, four AAL types were defined initially. In implementing AAL 3 and AAL 4, however, it was found that the same functions were required for both types and the two were merged into AAL 3/4. Subsequently the ATM Forum decided that AAL 3/4 was unnecessarily complex for many applications (such as simple point-to-point ATM links), so the much simpler AAL 5 was developed.

AAL 2 was initially reserved for variable bit-rate real-time services, such as coded video. These services are now specified to use MPEG 2 over AAL-5 (see Chapter 11). AAL-2 has now been defined for carrying low-bandwidth delay-sensitive applications such as compressed voice over ATM.

### 10.6.1 AAL 1

AAL 1 is defined to handle constant-bit-rate (CBR) data, such as from a PDH or leased line service. Figure 10.17 shows the SAR (segmentation and reassembly) treatment of the received data stream.

The original frequency of the CBR service can be recovered after the cells have been transported through the network using a technique such as *Synchronous Residual Time Stamp* (SRTS). SRTS uses the CSI field of cells with odd sequence numbers to transfer an offset to a master network clock. The sequence number and timestamp are protected from errors during transmission with an error-checking and correcting code.

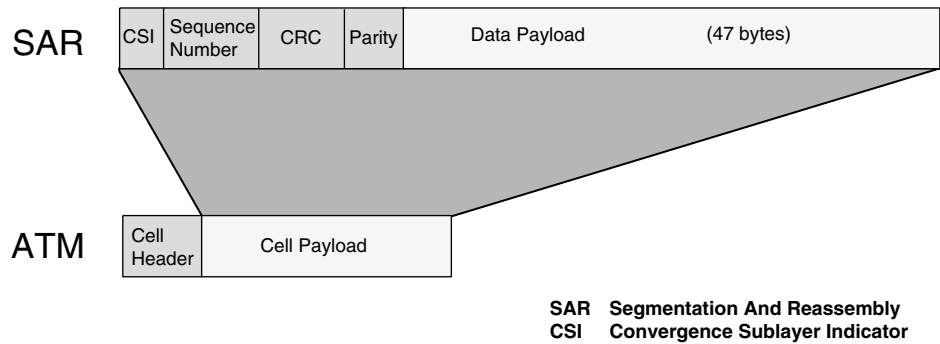
### 10.6.2 AAL 3/4

AAL 3/4 is defined to handle variable-bit-rate data transfer, either connection-oriented or connectionless. Variable-length packets (up to 64K) from services such as SMDS are encapsulated with a header and a trailer to form the *Convergence Sub-layer PDU* (CS-PDU), and then segmented into cells. The CS-PDU is 44 bytes long

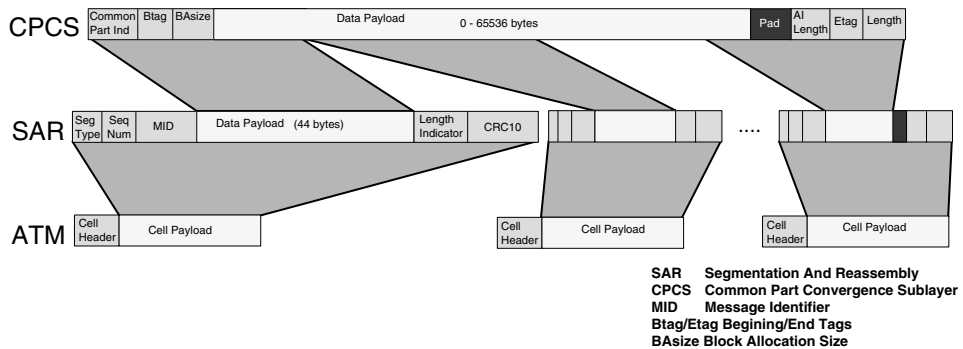
and is further encapsulated with another header (2 bytes) and trailer (2 bytes) to become a *Segmentation and Reassembly PDU* (SAR-PDU), which is inserted into cell payloads (Figure 10.18). Building sequence numbers and other protection into the SAR-PDU makes it possible to handle error conditions during the transfer of the cells.

The protocol incorporates support for shared media by allowing multiple messages to be multiplexed onto a single virtual channel using the MID (Message Identifier)


- Constant Bit Rate - Connection Oriented
- PDH & Leased Line Emulation
- Clock recovery via Synchronous Residual Time Stamp (SRTS) and common network clock



**Figure 10.17 AAL 1.** This is used to transport a continuous constant bit rate stream of bytes. AAL 1 service data is segmented into 47 bytes, with an additional overhead byte to detect cell loss/misinsertion and transfer timing information.



**Figure 10.18 AAL 3/4.** This is used to transport variable-length frames. AAL 3/4 service data is encapsulated in a CPCS frame with header and trailer bytes for protection; it then is segmented into 44 bytes for each cell. Four additional overhead bytes are added at the SAR level to detect the beginning and end of the frame, cell loss, and data corruption. The SAR level also allows multiplexing data from multiple services onto a single ATM connection using the MID.



**Figure 10.19 AAL 5.** This is used to transport variable-length frames. AAL 5 service data is encapsulated in a CPCS frame with trailer bytes for protection. The CPCS frame is padded to a multiple of 48 bytes and then is segmented directly into 48-byte cell payloads with no additional overhead. The last cell of an AAL 5 PDU is identified by a bit of the PT field in the cell header.

field. In addition, the CLNAP, AAL 3/4, and ATM Cell protocol is virtually identical to the SMDS SIP-3 and SIP-2 protocol, allowing straightforward interworking.

### 10.6.3 AAL 5

The ATM Forum has defined the *Simple Efficient Adaptation Layer* (SEAL) as a simpler alternative to AAL 3/4 where shared media support and the high level of protection (against mis-sequencing, for example) are not required, such as over simple point-to-point ATM links.

AAL 5 procedures are straightforward. There is no AAL-level multiplexing and all cells belonging to an AAL-5 CPCS-PDU are sent sequentially so no mis-sequencing protection is needed (Figure 10.19). The CPCS-PDU has only a payload and a trailer; the trailer contains padding, a length field, and a CRC-32 field for error detection. A PTI bit in the ATM header is used to indicate when the last cell of a PDU is transmitted, so that a PDU can be distinguished from the one that follows.

Due to the simplicity of AAL 5, it has been chosen as the AAL type for frame relay and LAN Emulation interworking. AAL 5 also has been chosen as the AAL for the lower part of the signaling protocol stack.

## 10.7 ATM Services Layer

ATM Services are those protocols that can be carried over ATM. These include not only end-subscriber services such as video, audio, and data transfer, but also other network technologies, such as PDH circuit transfer, frame relay, SMDS, and LANs.

The challenge for the ATM network is to provide each of these different services with their individual quality of service requirements, while managing the operation of the network in the most economical manner. The ATM network control protocols, such as those for management and signaling are also transported over the ATM network and can be considered at the Services layer.

# ATM Testing

## Deployment of ATM-Based Services

Stewart Day

*Hewlett-Packard Australia Ltd., Victoria, Australia*

### 11.1 Introduction

Asynchronous Transfer Mode (ATM) is a network technology designed to meet the needs of the world's communications services, both today and into the future. It has been designed to allow integrating a wide range of services having diverse traffic characteristics and service quality requirements into a single network infrastructure. It is a highly scalable technology and can be used over physical interfaces with rates from as low as 1.5 Mbps to 2.4 Gbps and beyond.

ATM can be deployed in both local area network (LAN) and wide area network (WAN) environments, and therefore is suitable for both private enterprise and public carrier networks. With standards defining the internetworking of ATM with other technologies, such as Ethernet and frame relay, ATM can be deployed as part of a gradual migration to high-bandwidth integrated networks.

#### 11.1.1 ATM technology

ATM is a connection-oriented technology using fixed-length packets. It effectively sits at layer 2 of the OSI protocol stack as a Data Link technology, although it also can be used to carry other layer 2 technologies (such as LANs), and additionally could be used as a layer 3 network technology (Figure 11.1). In its most popular current use, ATM is playing a significant part in the infrastructure growth of the Internet, where it is firmly placed as a layer 2 technology for IP.

Each ATM packet, or *cell*, is 53 bytes long, with a 5-byte header and 48-byte trailer. In an ATM network, cells are used to carry service data on a *virtual connection* (VC) over a predefined route through a network of switches. Cells from multiple services are multiplexed together with management and empty cells into a *contiguous cell stream*, i.e., the first bit of one cell follows immediately after the last bit of the

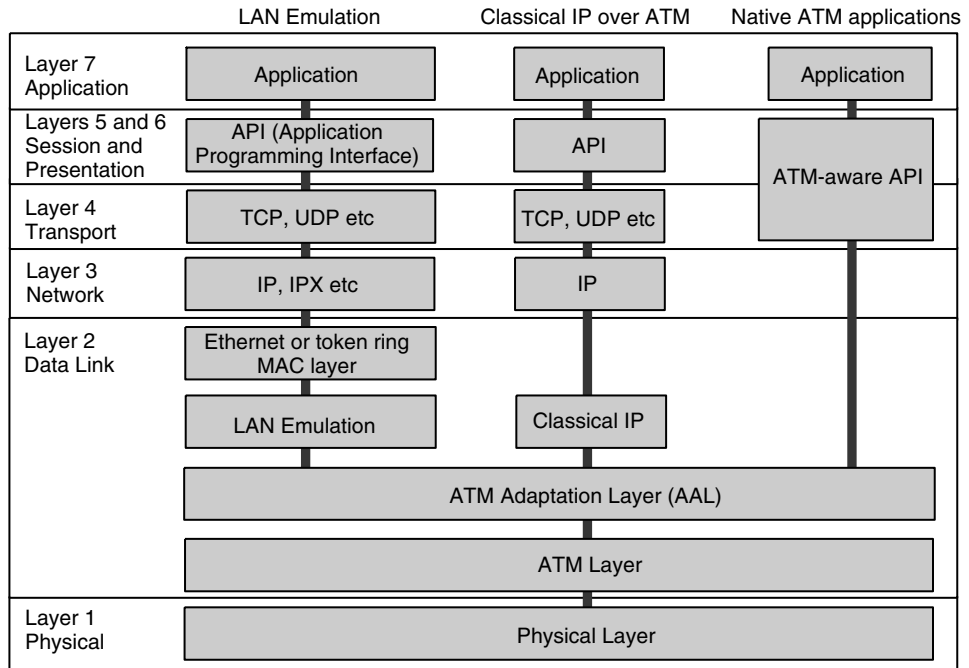


Figure 11.1 ATM in the OSI model.

preceding cell. Service data is adapted, using the ATM Adaptation Layer (AAL) to fit into as many cells as required.

The short, fixed-length cell structure allows fast switching in hardware, a key performance benefit over the software processing of router-based networks. The cell also is short enough to allow delay control on real-time services (such as voice), while at the same time long enough to transfer data services efficiently (such as long, variable-length packets from a LAN).

An ATM service uses cells only when it has data to send, a process referred to as *bandwidth on demand*. This is the asynchronous part of ATM and contrasts with synchronous Time Division Multiplexing (TDM) networks, where a fixed bandwidth is reserved for the duration of a connection.

Bandwidth on demand results in *variable bit rate* (VBR) traffic profiles in addition to services (particularly many real-time services) that remain *constant bit rate* (CBR). VBR profiles allow a resource allocation effect known as *statistical multiplexing*. Statistical multiplexing works by assuming that, over time, the high-bandwidth periods on some services will correspond with the low-bandwidth periods of other services. Rather than reserving bandwidth at the maximum rate for each service, the network operator can reserve a lower average rate, allowing more services to be accommodated. This results in the sum of the maximum bandwidths of all services on a link actually being greater than the bandwidth of the link itself. This process has to be managed carefully if congestion is to be avoided.

### 11.1.2 ATM service deployment

While simple in concept, ATM has evolved into a complex series of protocols to accommodate the needs of the many services it is designed to carry (Figure 11.2). Deployment, in both public and private network environments, therefore presents some key challenges.

As demand for ATM accelerates, it is vitally important that new services be installed and commissioned as quickly as possible. It is equally important that both providers and users be highly confident that the new ATM service will operate reliably, particularly because users will be unable to tolerate disruption of mission-critical services. Transition to ATM must be as painless as possible.

With the complexity and interconnection of different technologies, problems will occur. It is vital to diagnose and solve these problems quickly. This need places great emphasis on the quality of the diagnostic tools used and their ability to guide maintenance staff to the source of the fault.

In addition to diagnosing faults, tools must allow network operators to optimize network performance. Competitive pressures mean that each network operator must gain optimum performance from each network to provide attractively priced service packages to customers. The complexity of ATM, along with its statistical nature, makes this task increasingly difficult as these networks grow. Again, the diagnostic tools must be sophisticated in their characterization of the network and its services.

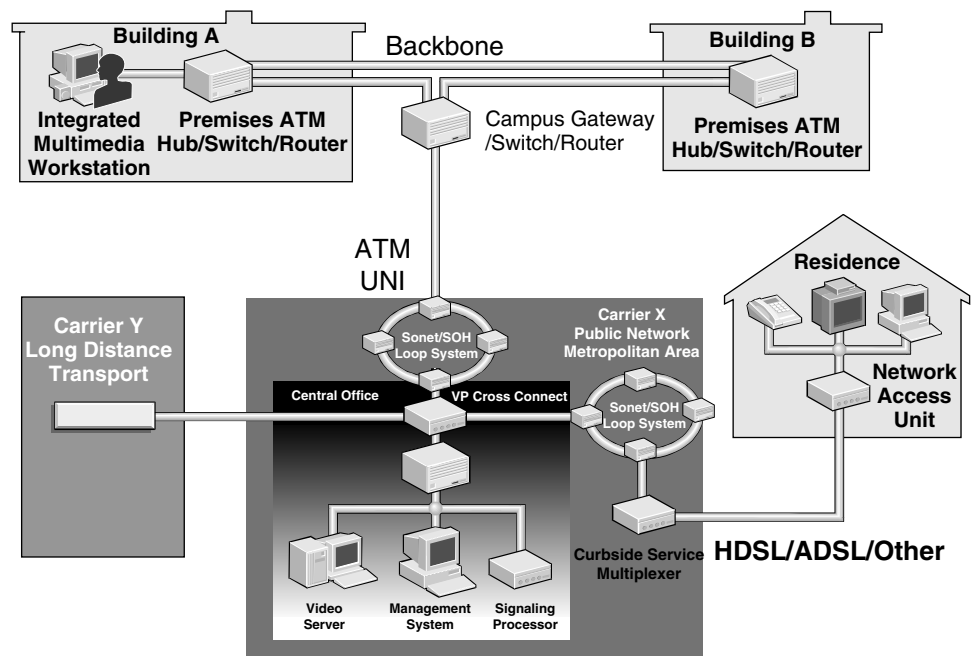


Figure 11.2 ATM internetworking.

## 11.2 ATM Testing

Because the intent of those who developed ATM was integrating existing and future services on a single network technology, the subject of ATM testing is naturally broad and diverse. Segmentation can be done based on the levels of the ATM protocol model, i.e., from physical transport testing to service protocol testing, and equally based on the stages of technology deployment, from R&D through to operational monitoring and troubleshooting.

Additionally, ATM is (and is likely to remain) only one of many network technologies in almost all real networks. ATM testing therefore must also be integrated with testing of these other technologies (such as LAN, WAN, and TDM) if end-to-end network performance is to be understood successfully and managed in order to guarantee the quality of service (QoS) that the entire network can provide to the users' applications.

### 11.2.1 ATM protocol segmentation

ATM testing essentially requires the combination of three key test areas (Figure 11.3):

- ATM Physical layer testing
- ATM protocol testing
- ATM service protocol testing

**ATM physical layer testing.** ATM can use virtually any digital transmission method available today or under development (ranging from twisted-pair through coax and

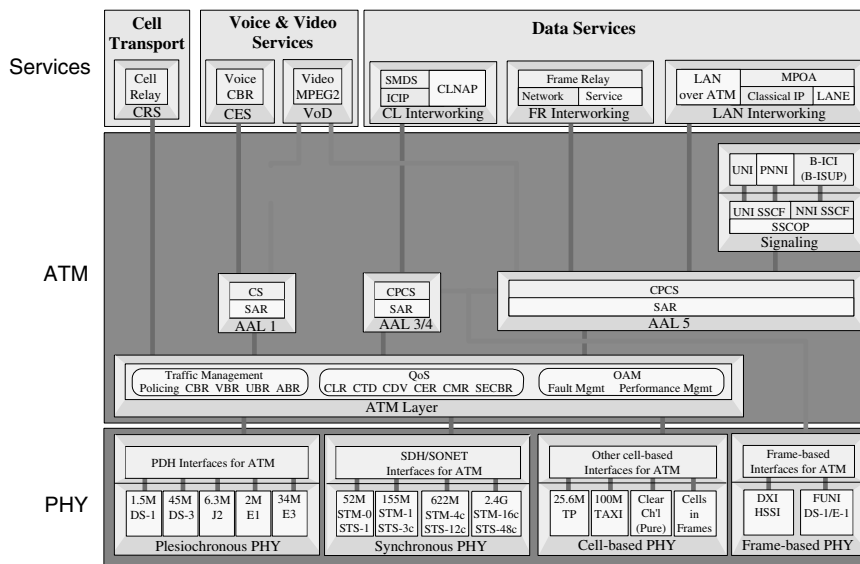


Figure 11.3 ATM protocol segmentation.



fiber optic to wireless), at speeds ranging from kbps to Gbps. Physical transmission test requirements are therefore equally diverse. Regardless of the media, the needs of ATM-based services are no different than those of other services, i.e., transporting digital signals between network equipment with as few errors as possible. This test requirement essentially boils down to ensuring that the transmission system meets its specified bit error rates and keeps clock jitter within specified limits.

Additionally, testing must ensure that physical framing procedures function correctly and that any other physical overhead functions, such as error and alarm detection and notification, comply with the appropriate standard. Testers designed for verifying physical transmission in ATM networks therefore are usually based on existing equipment developed to verify digital physical transport for other services, such as PDH, SDH, or SONET TDM networks, or LAN/WAN packet networks.

In addition to testing the physical transport system, the other key Physical layer function is *transmission convergence* of the fixed-length ATM cells into the framing structure of the physical network. Different convergence methods are specified for each defined transport system, but each essentially involves mapping the cell stream (octet-aligned) into the transport frame payload. In most cases the ATM cell stream uses its own framing mechanism, not fixed to the transport framing; this is known as *cell delineation*.

Cell delineation uses an algorithm to find valid cell headers, consisting of 5 bytes with a good Header Error Control (HEC) checksum in the 5th byte, starting from a serial stream of unframed bits. The algorithm relies on cells being contiguous; once it has found a valid 5-byte pattern, it checks for another valid pattern 53 bytes later. After a number of consecutive valid cell headers have been detected, the cell stream is synchronized, and processing of cell headers and payload data can proceed. While synchronized, HEC checking continues and has the ability to correct a single-bit error anywhere in the header. If two or more bit errors are detected in a cell header, the cell is discarded by the detecting switch. If a number of consecutive cells occur with cell header errors, synchronization is lost and the cell delineation algorithm starts again.

Cell delineation testing and HEC analysis are the primary functions of transmission convergence for most transport systems; they are a basic function of all ATM testers.

**ATM protocol testing.** The ATM protocols define a series of procedures to transport service data across a network of switches. The protocols can be split into these categories:

- Connection management
- Service adaptation
- Cell transport

**Connection management.** *Connection management* covers the protocols used to set up and manage virtual connections through the network of ATM switches. Connections can be set up permanently (Permanent Virtual Circuit, or PVC), or dynamically set up on demand using signaling protocols (Switched Virtual Circuit,

or SVC). In both cases, decisions must be made on which route to use through the network to guarantee the quality of service requested by the service user, while maintaining the service quality of the other user connections through the network. For this to take place, the users must specify certain details about their service application, choosing from a range of standard traffic types and classes. If the requested traffic parameters and service quality cannot be provided, the connection request will be rejected.

The key ATM technology that must be tested here is the set of ATM signaling protocols. Different protocol variants have been developed to meet the needs of:

- UNI: *User Network Interface*, the interface between the user's terminal equipment and a network switch (ITU-T Q.2931/Q.2971, ATM Forum UNI 3.1/4.0).
- PNNI: *Private Network Network Interface*, the interface between switches inside a private enterprise network (ATM Forum PNNI 1.0).
- NNI: *Network Network Interface*, the interface between switches inside a public carrier network (ITU-T Q.2761-Q.2764).
- B-ICI: *Broadband Inter Carrier Interface*, the interface between different public carrier networks. (ATM Forum B-ICI 2.0).

Note that NNI and B-ICI have been specified by ITU and ATM Forum to use B-ISUP (Broadband Integrated Services User Part ITU-T Q.2761-Q.2764), a part of the SS7 (Signaling System 7) system used in the world's public telephone networks.

ATM signaling is carried in-band on predefined ATM virtual connections through the network. It uses its own adaptation layer, *Signaling ATM Adaptation Layer SAAL*), to provide guaranteed delivery of the protocol messages. Signaling testing requirements range from verifying protocol conformance, a vital task in ensuring multivendor interoperability, through to performance stress testing and operational monitoring—where, for example, network operators could gain an understanding of network utilization patterns at different times of the day.

**Service adaptation.** Because ATM cells are short, fixed-length packets, adaptation procedures are required to allow the wide variety of service data structures to be carried across the network. The ATM Adaptation Layer (AAL) defines a series of adaptation techniques to segment service data into cell payloads at the entry point to the ATM network, and to reassemble the received data back to its original format at the exit point from the ATM network. This process is known as *Segmentation and Reassembly* (SAR). Three different AAL types are currently specified (ITU-T I.363), with differing levels of functionality to accommodate the needs of different types of service most effectively.

AAL 1 is designed for constant-bit-rate (CBR) services with real-time delay requirements, such as a 64 kbps telephone signal or uncompressed digital television signal. The AAL 1 protocol incorporates a sequence number, to allow the ATM end point to detect cell loss; CRC and parity bits, to guard against sequence number bit errors; and a clock synchronization mechanism to allow the frequency of the original CBR signal to be recovered across the network. The AAL 1 functions require a single byte of overhead from every cell payload, leaving 47 bytes for service data.

AAL 3/4 is designed for variable-bit-rate (VBR) services without real-time requirements, such as LAN or WAN data services. It incorporates sequence numbers, CRCs across the cell payload, and length and multiplexing fields; it is able to detect and recover from the error conditions introduced by the ATM network (primarily bit errors and cell loss). The AAL 3/4 functions require 4 bytes of overhead from every cell payload, leaving 44 bytes for service data.

AAL 5 was a late addition to the standards. It was designed for the same services as AAL 3/4, but with greatly reduced error-handling capabilities. This reduction in capability allows a gain in efficiency over AAL 3/4 because no overhead is required in the payload of all but the last cell of the segmented service data packet. This gain in efficiency over AAL 3/4, and the expectation (hope!) that error events should be rare, is the primary reason why AAL 5 is defined in the standards for encapsulation of nearly all LAN and WAN services over ATM.

AAL testing requirements center on ensuring that SAR functions correctly and that errors are correctly detected and handled. Additionally, AAL analysis can be used in the measurement of service “goodput,” the proportion of packets successfully reassembled versus those which cannot due to bit errors or cell loss.

**Cell transport.** Finally, once the connection has been set up and the service data has been segmented into the cell payloads using an AAL, a cell header is added and the services cells are multiplexed into a cell stream (with cells from other services), and sent across the network. On each physical link between switches, the services cell header carries connection identifiers: Virtual Path Identifier (VPI) and Virtual Channel Identifier (VCI). These are used in the switch to reference lookup tables (set up during connection establishment) that define the next physical link (and VPI/VCI for this connection over that link). The cell stream is therefore switched through the network over the predefined route.

In order to allow successful management of ATM cell transport, traffic management procedures are defined in the standards (ITU-T I.371 and ATM Forum UNI 3.1/4.0). These define procedures for traffic and congestion control, including traffic parameters and descriptors, quality of service, and network performance. Traffic can be classified into four major types, each of which could be routed and managed differently in the network. These are:

- CBR (Constant Bit Rate)
- VBR (Variable Bit Rate)
- UBR (Unspecified Bit Rate)
- ABR (Available Bit Rate)

*CBR* traffic is real-time, constant-bandwidth service defined by a *peak cell rate* (PCR) value. The tolerance to cell jitter around PCR is defined by the *cell delay variation tolerance* parameter (CDVT). The network guarantees transport quality as long as the transmitted cell stream complies with these parameters. CBR will have the highest priority in a switch.

*VBR* traffic is real-time and non-real-time, variable-bandwidth service defined by PCR (with tolerance CDVT) and *sustained cell rate* (SCR). SCR is the average

bandwidth over time and has a tolerance parameter of *maximum burst size* (MBS). The network guarantees transport quality as long as the transmitted cell stream complies with these parameters. VBR will have next highest priority after CBR. *UBR* traffic is non-real-time, variable bandwidth service with no guarantees.

UBR is a “best-effort” service designed for cases where the bandwidth is bursty and unpredictable, such as from a LAN or IP network. UBR has the lowest priority in a switch; UBR services therefore are the most likely to suffer cell loss if congestion occurs. UBR service can be improved by adding a technique known as *early packet discard* (EPD). EPD blocks transmission of the remaining cells of a segmented packet if an error or congestion is detected, thus avoiding sending further cells that will not be successfully reassembled at the far side of the network.

*ABR* traffic is non-real-time, variable-bandwidth service designed for the same bursty traffic as UBR, but adding a flow control feedback mechanism to help prevent congestion and cell loss while using the network bandwidth as effectively as possible. Cell sources that adjust their transmission rate to comply with the feedback from the network should experience minimum cell loss. With ABR, a *guaranteed minimum cell rate* (MCR) parameter is defined, which ensures that even in a congested network, a service will get this level of traffic through.

To ensure that service traffic complies with the selected traffic type and parameters, the *usage parameter control* (UPC) procedure, also known as *policing*, is used. Policing uses a *generic cell rate algorithm* (GCRA), also referred to as the “leaky bucket algorithm,” which tests each cell as it passes a point in the network (usually the UNI) against the specified parameters for the traffic type. Cells that comply are passed on unchanged to the network. Cells that do not comply are either tagged or discarded.

Cell tagging involves changing the *cell loss priority* (CLP) bit in the cell header from high priority to low priority. Any time congestion is experienced in a switch, low-priority cells will be discarded first. Cells already set with low priority will be discarded if policing decides they do not comply. Each network operator will have his or her own policy over whether cell tagging is supported in the network.

The service agreement between the user and network operator will guarantee quality of service (QoS) performance for cells that comply with the traffic contract. The primary QoS parameters (defined in ITU-T I.356) are:

- CLR (cell loss ratio)
- CTD (cell transfer delay)
- CDV (cell delay variation)

*Cell loss ratio* (CLR) is the ratio of lost cells to transmitted cells. Cell loss is caused primarily by severe congestion, causing buffer overflow. It also can be caused by bit errors in the cell header.

*Cell transfer delay* (CTD) is the mean transfer delay through the network. Cell delay is caused by signal propagation through physical media, switch fabric, and buffers.

*Cell delay variation* (CDV) is the variation in cell delay through the network. CDV is specified as both a 1-point measurement (for CBR services) and a 2-point

measurement. CDV also is referred to as *cell jitter*. CDV is caused by cells being queued in buffers for varying lengths of time as they wait to be transferred through the switch fabric and multiplexed onto the output link. As loading increases in the network, CDV is likely to increase, and in extreme cases of congestion buffer overflow might occur, causing cell loss.

Other specified QoS parameters include *cell error ratio* (CER), *cell misinsertion rate* (CMR), and *severely errored cell block ratio* (SECBR). These parameters are not affected by congestion in the network.

Cell transport testing centers on verifying that the correct VPI/VCIs are being used, ensuring the correct operation of policing, and measuring the QoS parameters for cell streams through the network. Additionally, the ability to characterize network performance and identify points of congestion is particularly useful when determining if the most appropriate routes are being used for each service, and whether the network links have been dimensioned correctly. Finally, traffic characterization can be used to verify that appropriate bandwidth parameters have been selected and that resources are not being reserved needlessly.

**ATM service protocol testing.** ATM service protocol testing covers the different specialized requirements of each service using the ATM network, whether directly or through interworking of ATM with other network technologies such as LAN or WAN. The particular testing needs of the key ATM services being deployed are discussed in detail later in this chapter, but essentially they can be split into:

- Verification of the protocol encapsulation and function mapping of the service data to the ATM network.
- Measurement of the performance of the service across the ATM network.

Additionally, service protocol testing might also require some means of measuring the effect of ATM impairments (such as cell loss or cell delay) on the end-user service; one example might be monitoring the effect of cell loss on a compressed digital video signal transported over ATM. Guaranteeing service quality will be assisted greatly by understanding the interaction of protocols through the protocol stack and correlating that with end-to-end measurements across the ATM connection and service endpoints. This capability is called *Service Analysis*.

### 11.2.2 ATM test market segments

The ATM market can be segmented based on the stages of the technology life cycle. At each stage there are differing requirements for test functionality, equipment flexibility, cost, and portability. Additionally, a range of differing test methods will be appropriate at each stage.

**Research and development.** ATM network equipment development requires test equipment that covers all functionality being designed into the equipment; it is flexible enough to allow R&D engineers access to the bits and bytes of data structures at all levels of the protocol stack. Key test functions during development include

verification of protocol conformance to standards, verification of the performance of data transfer through the equipment under normal and extreme load conditions, and verification that the equipment can be configured correctly and responds appropriately to alarm and error conditions.

Due to the degree of flexibility and detail required in this environment, ATM R&D test equipment tends to be large and modular, with the ability to be configured to suit the particular functions under development. These testers are based largely on programmable logic platforms that allow new capabilities and new modules to be added over time, as new standards evolve. They also include comprehensive programming capabilities, allowing designers to create custom scripts to test or simulate particular features where implementation might not be complete. Typically these testers are not portable and would be used only by technology gurus with a great deal of ATM knowledge.

**Manufacturing test.** In manufacturing, ATM test capability is required as part of the test systems used to perform functional tests (such as error and alarm tests) for boards, subsystems, and systems. Applications include board functional test, burn-in, system test, and integration test. For these applications, flexibility is vital to allow the test system to access the necessary test points, particularly those that are internal to the network equipment and might not be in a standard physical format. Equally important is the ability to integrate with the equipment used to verify other aspects of the system under test, such as specialized Physical layer verification; this can be aided if the various parts of the test system can be controlled from a common programmable scripting environment.

**Field trial.** The field trial stage of ATM deployment usually consists of building a small network of new equipment, with detailed testing under controlled conditions to verify that the real service can be operated and managed successfully. In general, field trials still require the detailed level of test capability found in R&D, only in a more portable platform. Continuity with the tests and procedures performed in R&D is useful because it allows the expertise of those who designed or evaluated the equipment to be transferred to the field trial team.

**Installation and commissioning.** During installation, the primary need is for test equipment that is quick and easy to operate and can verify that the installation procedure has been correctly followed, and that the basic functionality of the equipment operates correctly. Ideally, equipment for installation testing should be rugged and portable, and it must be easy to understand for the technicians performing this task, no matter how much or little they might actually know about ATM.

Network and service commissioning is the final step before the network becomes operational. Effective testing at this stage is vital to ensure that the services can be provided as expected. Tests must check that each element in the network operates correctly with every other, and that no protocol conflicts occur. These are particularly likely in multivendor networks, where each vendor might have interpreted a newly developed standard slightly differently. Essentially this stage of deployment requires a focused subset of the capability required during the field trial stage; con-

tinuity of test procedures can help by ensuring that the knowledge and experience of the development and field trial teams can be used effectively during commissioning.

**Operations, maintenance, and troubleshooting.** Operational ATM networks will rely on a combination of network management capability, for normal operation, and additional test equipment to perform maintenance and troubleshooting tasks. These are examined in a following section, titled “Testing in Operational ATM Networks.”

### 11.2.3 Out-of-service vs. in-service testing

*Out-of-service testing* is used during equipment development, and installation and commissioning of equipment and services. It also is used in operational networks where in-service techniques have not successfully found the cause of a problem (Figure 11.4).

Test equipment is connected to a network link and used to generate test traffic to the *system under test* (SUT). The SUT can be a piece of equipment or section of the network. The responses of the SUT are then measured and analyzed for protocol correctness and acceptable performance.

With individual equipment or protocols, this approach can be formalized as conformance testing, to verify precisely through a series of standard test cases whether a standard is being conformed to or not. The standard cases are called *automated test suites* (ATS) and are defined by ATM Forum. The test equipment must emulate remaining elements of the network in order to test how the SUT reacts to both normal and errored protocol behavior. This technique is used particularly with protocols such as ATM signaling and is vital, particularly where multivendor interoperability is required in the network.

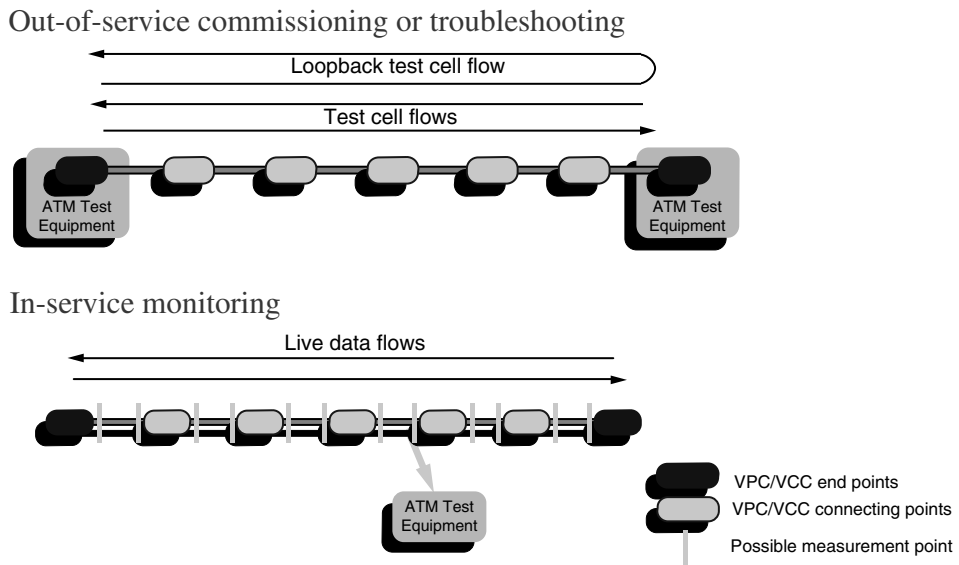


Figure 11.4 Out-of-service and in-service testing.

Network performance can be measured by simulating traffic with ATM cell streams containing test data such as sequence numbers and timestamps. These can help create a detailed characterization of the SUT's behavior; the tests can be enhanced by adding background traffic or loading other network ports. The ITU-T currently is specifying a test cell (ITU-T O.191) for precisely this purpose; key issues remain to be resolved, however, including how to realistically model the traffic distributions for the test traffic.

*In-service testing* is used in operational ATM networks to verify operation, gather statistics, and troubleshoot problems. In-service test methods center on monitoring the signals at one or more points around the network, and measuring performance and faults through detecting alarms and errors and decoding traffic. Ideally, in-service testing should be nonintrusive, i.e., test access should be passive, causing no change to the monitored signal. As explained in a later subsection, "Test Access," this is not always possible.

Intrusive test methods, such as injection of additional test traffic onto a connection, will modify the real traffic profile and could cause resource problems that invalidate the measurements. Other, less intrusive techniques could include measuring the performance of a dedicated test connection that uses the same route and assuming that the performance will be similar; these also might not be accurate, depending on how switches buffer and prioritize different traffic streams. Ideally nonintrusive methods can be used for in-service performance measurement. A key issue here is how to measure delay between two remote points accurately without passing timestamped cells between them.

#### 11.2.4 Testing in operational ATM networks

Discussion of testing actual, operating ATM networks will address four major topics: built-in switch statistics and MIBs, operation and maintenance (OAM), test equipment for operational networks, and test access.

**Built-in switch statistics and MIBs.** Ideally, the network management system will be able to gather sufficient traffic, error, and alarm indications from the switches to successfully operate and troubleshoot the network. To ensure that each switch gathers compatible information that can be accessed by network management systems, the ATM Forum and IETF are defining MIBs (Management Information Base), which codify the information required to manage the network. For private networks, MIBs can be accessed using SNMP (Simple Network Management Protocol) from the IETF. In public carrier networks, *Common Management Information Protocol* (CMIP) from ITU-T is mostly used. The major MIBs for ATM networks are:

- ATM MIB for managing ATM devices, networks, and services. (*IETF RFC 1695*)
- ATM-RMON MIB for ATM remote monitoring extensions to the RMON MIB for LANs. (*IETF and ATM Forum draft*)
- ILMI MIB for Physical and ATM layer configuration and statistics in each ATM switch, including ATM address registration. (*ATM Forum*)



- M1 to M5 MIBs defining the interfaces between network management systems and network elements in both public and private networks, and between them. (*ATM Forum drafts*)

At this stage in ATM development, the functionality available from these MIBs remains limited. Much more work remains to be done before they are finalized and can be used effectively.

**OAM.** ATM, and many of the physical transport technologies it can be carried over, have built-in fault and performance management functionality called *OAM* (operations and maintenance, ITU-T I.610). OAM functions can be used to detect specific conditions automatically, such as links that are physically broken or have a high level of bit errors, and report them to the network operator. With some newer technologies, such as SONET/SDH, it also is possible for the switches to rectify the fault automatically by switching to a backup link or a new physical route through the network.

OAM techniques were developed primarily for networks, such as TDM and SONET/SDH, where signals have constant reserved bandwidth (e.g., 64 kbps time-slots), and the only performance measures relate to bit errors and clock jitter. In ATM networks, however, where traffic can be variable and suffer unique impairments (such as cell loss and cell delay variation), performance management requires many more statistics to be gathered.

The standard for ATM OAM defines performance management cell flows at both the virtual path (F4) and virtual channel (F5) levels. These cell flows are added to the VP (using the same VPI) or VC (using the same VPI/VCI) at a predetermined rate, such as one cell for every 128, 256, 512, or 1024 user cells. Each OAM cell contains fields such as sequence number, user cell count since last OAM cell, parity check on the user cells, and optionally a timestamp. OAM cells are generated and analyzed by the ATM switches either end-to-end or across specific segments of the route through the ATM network. Analyzing them can give measures of bit errors, cell loss, and delay.

Implementation of performance OAM cells, however, requires additional hardware—and therefore expense—on each port of an ATM switch. Consequently, very few switch vendors have thus far implemented these features. Even once they become more widely available, the information they provide will be only part of what is required to successfully understand and troubleshoot the operation of an ATM network.

**ATM test equipment in operational networks.** There is, and therefore will continue to be, a need for ATM test capabilities in operating networks. In general, this equipment is likely to be either:

- Dispatched portable field service test boxes, for maintenance and troubleshooting, or
- Measurement probes that can be distributed around the network and remote-controlled from a central site to augment the switches' own statistics and assist the network management system.

In the early stages of ATM network deployment, it is possible for test equipment to incorporate elements of both, such as dispatched (standalone) portable testers that can be distributed around the network at several points and remote-controlled from a central site. This gives the advantage of extending the same test techniques from dispatched through to distributed testing. This can add significant value where remote synchronization capability is added to allow integrated tests to occur across multiple points in the network, with centralized correlation of measurements at these multiple points, and through the entire protocol stack of the service.

**Test access.** A significant issue for any testing in operational ATM networks, particularly those using fiber optic transport systems such as SONET/SDH, is how to gain in-service test access to the network. The fundamental problem is that switches do not, in general, provide test access ports—and even if they did, the monitor port might not provide an accurate enough copy of the actual traffic to allow the desired measurements to be done. In-line test access is possible in the case of fiber, but it requires adding optical splitter devices which, by tapping off a portion of the optical signal power, attenuate the remaining signal to the next network element. This signal degradation could be enough to cause transmission signal problems. Additionally, splitters add cost to each link on which they are placed.

Despite the problems of test access, there is no alternative unless the switches implement all the required test functionality themselves. The ATM Forum currently is working on standardizing test access techniques, including ATM circuit steering (ACS), where a copy of the signal under test is steered over a second route from the switch back to the test equipment for analysis.

### 11.2.5 Core ATM tester features for service deployment and operation

Having looked at both the protocol and lifecycle segmentation of the ATM market, it is evident that there is a wide range of features requiring testing, as well as differing levels of flexibility and depth for each application area. For service deployment and operations, the core features are:

- Protocol testing through the layers of the protocol stack for each service in the network.
- ATM layer test capability, including traffic analysis, QoS measurement, and signaling test (if used by any of the deployed services).
- Physical layer support for each interface used in the network, including multiple ports to allow simultaneous measurement across both directions of a link, or between an input and output of a network element.
- Synchronization of tests and correlation of measurements through all layers of the protocol stack, and between multiple ports.
- Remote-control operation, allowing the tester to be left in the network, log measurements over time, and be accessed from a central location such as a network operations center.
- Ease of use, particularly an intuitive graphical user interface (GUI), comprehensive help system, and canned tests (both predefined and created through test

scripts), allowing test procedures to be automated and understood by all levels of user, no matter how much or how little knowledge of ATM they actually might have.

- Portable and rugged equipment, allowing easy transportation to where it is needed in the knowledge that it will operate reliably once it gets there.

Additional features that also may be useful for deploying and operating specific services are identified in the following section.

### 11.3 ATM Service Testing

This section will now look at the most common ATM-based services and their testing needs.

#### 11.3.1 Cell relay service

The basic ATM service is called *cell relay*. Characteristics of a cell relay service include:

- Cell-based information transfer
- Connection-oriented transmission
- Support of both CBR and VBR applications
- AAL usage as appropriate

This service essentially involves the transport of data, in 48-byte groups, over a connection across an ATM network (see Figure 11.5). Initially, permanent connections (PVCs) are specified, but as signaling standards develop, there should be no barrier to use of switched connections (SVCs). Definition of the cell relay service allows for both constant and variable bit-rate applications, and the use of the Adaptation layer as appropriate in the terminal equipment.

Responsibility for the actual user service and applications will be with the user. For the network operator of a cell relay service, testing will focus on fault handling

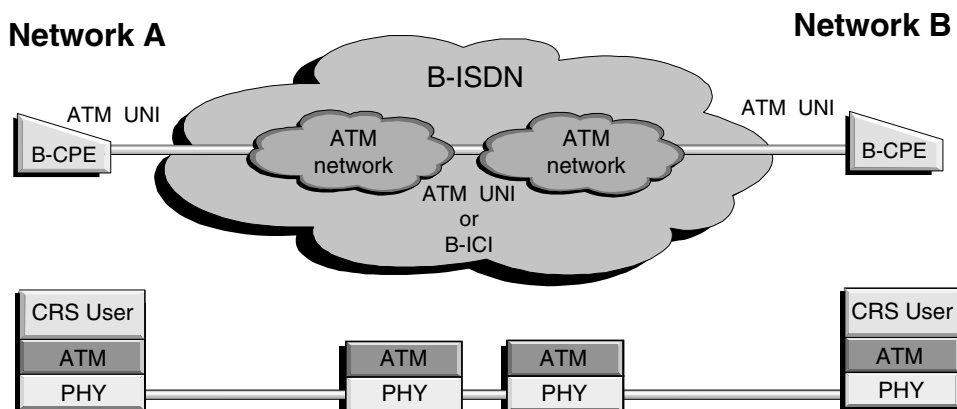


Figure 11.5 Cell relay service.

verification and performance testing of the allocated ATM connections. In particular, this will include verifying the use of the correct VPI/VCI, characterizing the bandwidth and traffic distribution of the service, and ensuring that it meets its QoS guarantees.

A tester suitable for cell relay, therefore, must be concerned primarily with fault and performance management at the ATM and Physical layers. To make these tests possible, the tester must include:

- One or two full-duplex test ports at the interface rate or rates in the network.
- ATM and Physical layer test capability that includes alarm and error generation and measurement; QoS measurement, both in-service and out-of-service; and ATM layer traffic characterization.
- AAL test capability as appropriate.

Many tests, particularly of a general transmission nature, can be done by accessing a single monitor point for in-service analysis, or performing out-of-service stimulus/response tests. These tests then can be enhanced by adding a second port, perhaps where several different ATM interface types are used, to allow correlation of measurements at multiple points in the network and testing of the switching functions between the two interfaces. If appropriate, AAL support also should be included.

### 11.3.2 ATM WAN backbone services

**Frame relay interworking.** Frame relay is a public WAN service used predominantly for LAN interconnection. It is currently the networking technology experiencing the most growth; services are available from most network operators throughout the world. It also is playing a key role in the growth of the Internet, being a significant portion of the network infrastructure of many Internet service providers (ISPs).

Frame relay is a connection-oriented technology based on multiplexing and switching variable-length frames of data and control information. In general, connections are permanently set up (PVC), but switched (SVC) connections are now being offered. Frame relay was developed as a simplified form of the X.25 packet network, taking advantage of the higher quality of today's transmission lines, which allows removing most of the error checking and retransmission techniques that had been necessary to give reliable connections over poor-quality analog circuits. Frame relay therefore is designed to give major performance improvements over X.25 and use most of the infrastructure already in place, making it relatively inexpensive to install.

Frame relay also can be used to provide more effective network usage than that provided by dedicated leased-line services. By providing frame multiplexing and switching, with the concept of bandwidth on demand (similar to ATM), it is possible to offer a comparable service to many users at lower cost than their existing connections.

Frame relay's X.25 and leased-line origins have allowed network operators to offer connections at rates between 56 kbps to 2 Mbps. The latest network equipment now supports Frame Relay at rates up to 45 Mbps; several network providers, particularly some ISPs, have taken advantage of this in preference to using ATM at these rates.

Frame relay standards are available from ITU-T, ANSI, TTC, and the Frame Relay Forum industry interest group.




Figure 11.6 Frame relay Network Interworking.

**ATM in frame relay networks.** The key role of ATM in frame relay networks is to provide a high-bandwidth backbone in the core of a carrier's network, while leaving frame relay for customer access. There also are requirements for interconnecting users between ATM and frame relay network segments. The ATM Forum, in liaison with the Frame Relay Forum, have agreed upon a set of interworking standards for each scenario; these are referred to as *Network Interworking* and *Service Interworking*.

**Network Interworking.** Network Interworking sets a standard for using an ATM backbone to connect two frame relay endpoints. (See Figure 11.6. For further information also see FRF.5, available from the Frame Relay Forum.) These endpoints may be frame relay networks, terminal equipment, or ATM terminals supporting the frame relay protocol stack. The specified technique encapsulates the frame relay frames over AAL 5 and carries them transparently across the ATM network. This standard is of most use to network operators who wish to upgrade their frame relay internally to ATM, but do not wish to change the services offered to their customers. The relatively high cost of supporting frame relay protocols in ATM terminals makes it unlikely that Network Interworking will be used outside this application.

Interworking ATM and frame relay involves mapping the variable-length frame relay frames into the fixed-length ATM cell payloads. This process is specified to use the AAL 5. Translation of frame relay DLCI (Data Link Connection Indicators) into ATM VPI/VCI values can be done either by mapping one-to-one, or multiplexing several DLCI values onto a single ATM channel. Frame relay congestion and discard priority fields also have specified mappings to equivalent features in the ATM network. Finally, traffic management must also be coordinated across the interworking unit.

**Service Interworking.** Service Interworking (Figure 11.7) is aimed at those wishing to use ATM to upgrade high-bandwidth segments of their networks to ATM while keeping frame relay for other segments whose networking needs can be met without ATM. Service Interworking (FRF.8) specifies translation mechanisms to transfer service data between frames and cells, and vice versa. This standard removes the need for end stations to know on what type of network the destination station resides, therefore avoiding the need for expensive frame relay protocol support in the ATM network. Service Interworking is seen to offer a gradual ATM migration path for large corporate frame relay networks, and is expected to grow in popularity as key user bandwidth needs increase beyond frame relay's standard rates.

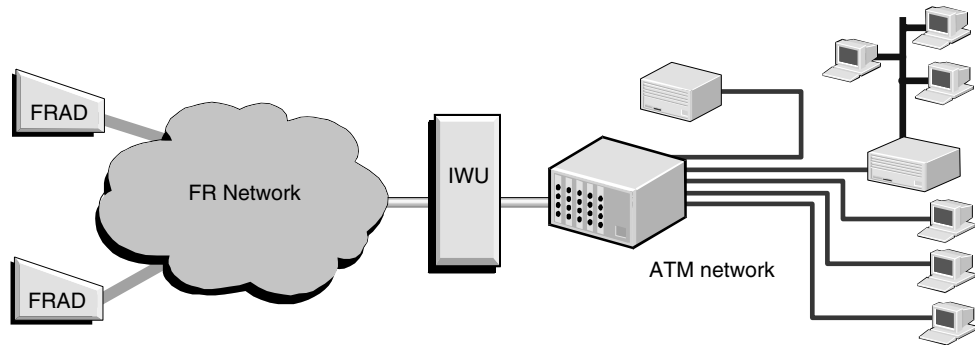


Figure 11.7 Frame relay Service Interworking.

Service Interworking does not require mapping of frame relay into ATM. Instead, the service data being carried over the frame relay link will be extracted and encapsulated directly into ATM; for example, IP encapsulated over frame relay would be converted into IP encapsulated over ATM.

**What and where to test?** During installation and commissioning, key testing focuses on protocol conformance, whether for frame encapsulation or frame-to-cell translation. Appropriate mapping of frame relay congestion and discard priority indicators must be checked to ensure transparency in the case of Network Interworking, and provide suitable translation to ATM indicators for Service Interworking. Equally, functions relating to mapping fault handling, dealing with traffic, and managing connections through the Interworking Units must be viewed along with performance measurements across Interworking Units and between network endpoints.

Once interworking is operational, the testing focus moves to troubleshooting and performance optimization. Troubleshooting must be quick and easy, to allow pinpointing of problems and fast recovery of user services. This includes monitoring end-to-end service performance.

The two main test locations in an FR/ATM internetwork are over the end-to-end frame relay connection, and locally across the FR-to-ATM interworking units. Because the frame relay portions of the network are likely to be more established, suitable frame relay test equipment might already be available for testing the end-to-end link. A dedicated tester supporting interworking, however, is needed for the links to the ATM network.

**Out-of-service protocol test.** As a first step in the test process, it is important to check the protocol mappings as frames pass from the frame relay network to the ATM network and vice versa (Figure 11.8). This measurement checks that the encapsulation, segmentation, and reassembly processes are implemented correctly and that the appropriate mappings of DE, BECN, FECN, and OAM functions occur.

The best way to test these mechanisms is to perform an out-of-service test using the tester to generate and analyze traffic on both sides of the IWU. This allows controlled traffic to be generated with either good or bad data in order to discover whether or not the IWU performs each step of the process correctly. Equally, alarm

and error conditions can be forced to ensure that appropriate action occurs in both networks. Multichannel generation and analysis allow connection multiplexing to be tested to ensure that appropriate mapping of channels occurs.

**In-service performance monitoring.** Following out-of-service protocol testing, the IWU can be inserted in a live network and the tester can move on to monitor the passage of real traffic. Monitoring on both sides of the IWU makes it possible to correlate events as they cross between networks and see how successful this is. It should also be possible to quantify the effect of the ATM link on the frame relay service, and to characterize the FR-sourced ATM traffic itself. Such a test could include correlation of Quality of Service parameters at the ATM layer to the performance of the frame relay service. The resulting data should help greatly to optimize the ATM

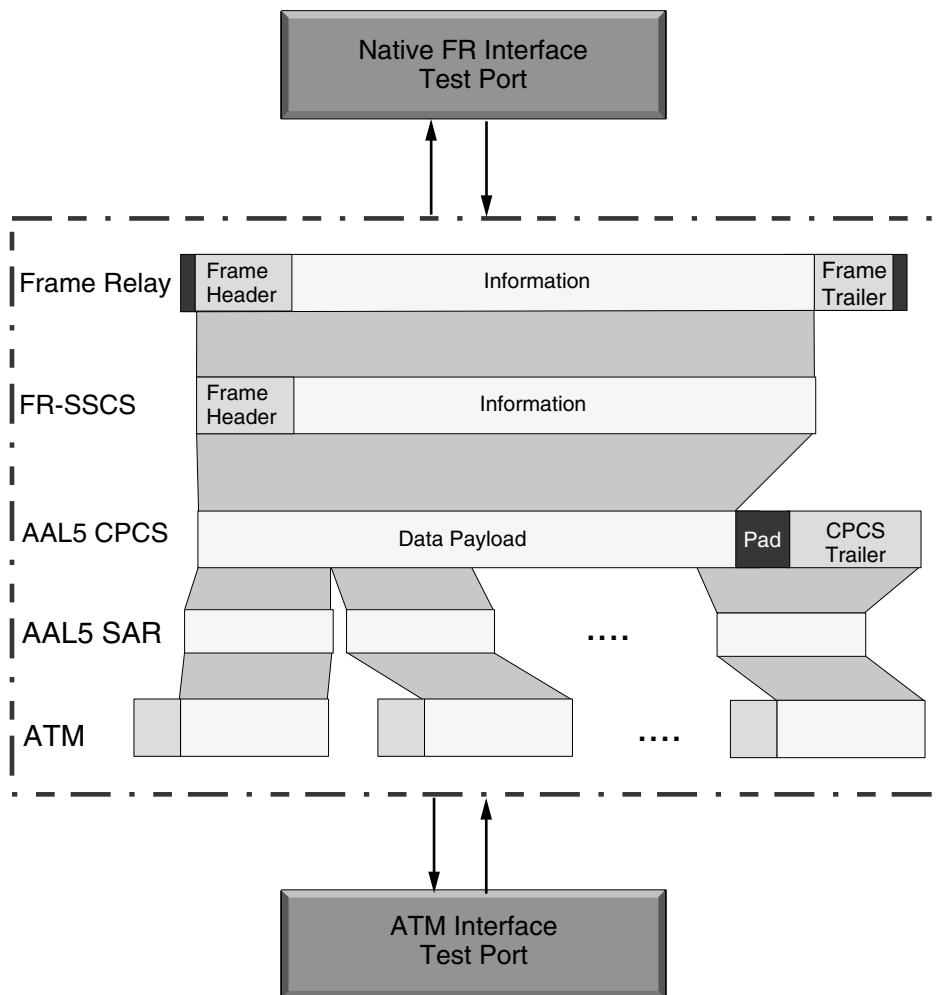



Figure 11.8 Frame relay interwork protocol testing.



**Figure 11.9** Frame relay interwork performance monitoring.

network to cope with multiple frame relay connections, and with the many other types of traffic that might be using it simultaneously. Figure 11.9 shows the relationship of the performance monitor to the FR/ATM internetwork.

**Test case.** As an example of a typical problem situation encountered by the installers of an ATM link into an existing frame relay network, imagine the following situation:

An ATM link has been installed to connect two existing frame relay network segments. Communication is fine within each segment but not possible between them.

Possible causes include:

- Interworking equipment not functioning properly.
- ATM network not functioning properly.

Because the IWU appears to be switched on and working, test equipment is needed to diagnose the fault.

*In-service problem diagnosis:* The first step is to monitor the network at various points to determine where traffic is present and where it is not. Next, an attempt can be made to confirm that traffic is crossing the IWU and which channels it is using, and whether the protocol mapping is correct. If no ATM traffic is being sourced by the IWU, the next step is to check that appropriately routed frames are reaching it. Checks can be done on each part of the link to see where errors and alarms that might indicate network malfunctions are present. During this test process, it is possible to correlate what is being seen between different segments of the network.

This procedure would require test equipment with both FR-over-ATM and native frame relay monitoring and decoding capabilities. The ability to monitor simultaneously on both the frame relay and ATM links would be useful, although this testing still could be done by monitoring each network separately.

*Out-of-service problem diagnosis:* In this example, the in-service diagnosis indicates that both frame relay and ATM segments appear to be working but that the IWU seems to be transmitting incomplete protocol data units (PDUs) into the ATM network. This traffic is rejected by the IWU at the far side of the ATM network. The



next step, therefore, is to take the IWU out of service and run detailed tests with data patterns generated by the tester.

What if this step shows data being lost when frames are longer than a certain size? In this example it turns out that most of the traffic the network is generating exceeds this size, which explains the incomplete PDUs. Re-examining the installation procedures for the IWU reveals the source of the incorrect setting. Modifying the configuration and replacing the IWU in the network solves the problem.

This procedure would require test equipment with not only FR-over-ATM and native frame relay monitoring and decoding capabilities, but also the ability to simulate each protocol through a series of test cases in order to diagnose the fault. Ideally, automated test scripts could be used to create repeating stimulus/response tests with PDUs of different lengths to detect this particular fault.

**Tester requirements.** A tester suitable for frame relay interworking tests must be concerned primarily with verifying protocol conformance and network performance at the service level. It also should be capable of ATM and Physical layer testing (as for cell relay service), to allow correlating cell-level Physical and ATM layer behavior to frame-level behavior of the frame relay interworking service. To make these tests possible, the tester must include:

- Two full-duplex test ports of the interface rate(s) in the network, providing ATM port protocol support for frame relay over ATM, and native frame relay port test capability.
- Data verification for FR-over-ATM to native frame relay.
- Simulation of frame relay and FR-over-ATM traffic, alarms, and errors.
- ATM and Physical layer test capability, including alarm and error generation and measurement, QoS measurement (both in-service and out-of-service), and ATM layer traffic characterization.
- Correlated analysis of data, both on frame relay and ATM ports, and through all levels of each protocol stack to the user service being carried.

**SMDS interworking.** Switched Multimegabit Data Service (SMDS) is a public WAN service developed by Bellcore, which (like frame relay) is used primarily for LAN interconnection. SMDS was developed as a high-speed alternative to frame relay, particularly for metropolitan area networks (MANs). It has been widely deployed, particularly in several European countries and with some U.S. carriers, but now is losing favor to ATM and particularly frame relay.

SMDS is a connectionless technology based on the *Distributed Queue Dual Bus* (DQDB) transport and multiplexing of IEEE 802.6. In Europe, the European Telecom Standards Institute (ETSI) modified the Bellcore standards to support European interfaces and called it *Connectionless Broadband Data Service* (CBDS). SMDS standards are available from Bellcore, ETSI (for CBDS), the SMDS Interest Group (SIG), and its European equivalent, ESIG.

Using a DQDB format gives a 53-byte cell structure that is nearly identical to that of the ATM cell. In fact, the PDU structures for connectionless services directly over

ATM (using AAL 3/4) are taken largely from those developed for SMDS, with the intention that ultimately SMDS networks will be migrated easily to ATM.

SMDS is specified to be carried over PDH lines at DS1, DS3, E1, and E3. This specification allows data rates of between 1.5 Mbps and 45 Mbps and places SMDS in direct competition with certain ATM applications.

**ATM in SMDS networks.** The key role for ATM-to-SMDS interworking is the interconnection of users between the networks. While ultimately SMDS networks may be upgraded to ATM, the process will be gradual and, in the interim, the two systems must be able to coexist in different parts of the network.

Initial ATM Forum interworking standards allow for ATM to transfer data transparently between two SMDS network users or CPE (see Figure 11.10). Direct connection between SMDS and ATM network users ultimately will be possible, but the method of doing so is still under investigation.

Interworking SMDS and ATM involves encapsulating variable-length *SMDS Interface Protocol (SIP)* level-3 PDUs into *Inter-Carrier Interface Protocol Connectionless Service (ICIP\_CLS)* PDUs, which are then carried on an ATM channel using AAL 3/4. Encapsulation first discards the SIP L3 trailer; the information it contains can be regenerated easily across the ATM network, and its functions duplicate those already being done (such as CRC checks). Despite the similarity of their structures, it is not possible to map SIP level-2 PDUs directly to ATM cells because the SMDS routing information is carried in the level-3 PDUs instead of in the cell headers.

The mapping function includes transfer of routing information, carrier identification, group address resolution, and mapping of SMDS-related QoS information into equivalent ATM functionality.

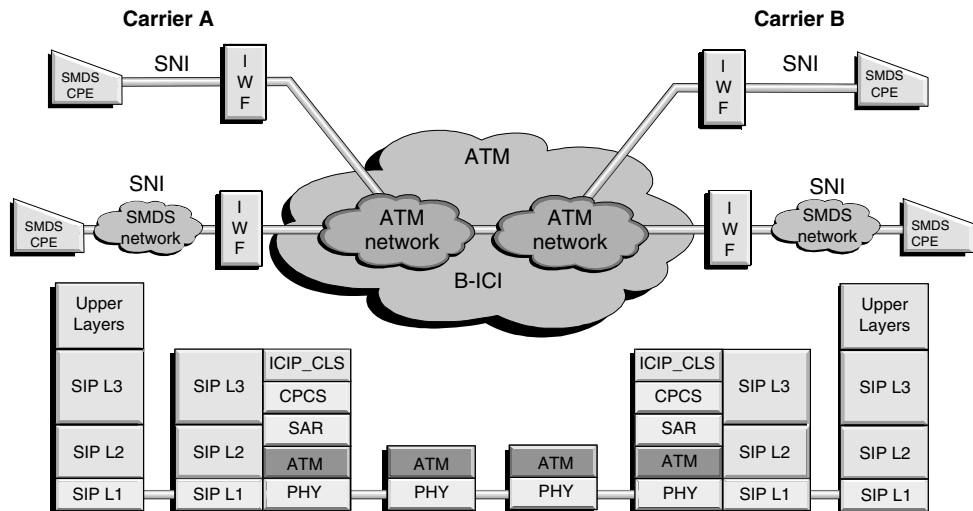


Figure 11.10 SMDS interworking.

**Out-of-service protocol test.** As a first stage in commissioning equipment, it is important to ensure that the protocol encapsulation and mapping occur correctly across the IWU at full data rates. The IWU also includes many protocol layers, with error-handling capabilities in its encapsulation of SMDS in ATM. It is important that a tester be capable of verifying as many of these capabilities as possible to ensure that appropriate action is taken.

The AAL 3/4 SAR multiplexing ID (MID) field allows multiple encapsulated SMDS PDUs to be multiplexed onto a single ATM channel. Alternatively, individual SMDS PDUs can use individual ATM channels. These mappings must be tested as well.

**In-service performance monitoring.** Once the operation of the IWU equipment is verified, the network itself can be tested with real traffic (Figure 11.11). It is important here that the tester be able to monitor points in both the SMDS and ATM networks so that correlation can occur, allowing analysis of traffic as it crosses the IWU. As with other service types, a key monitoring task is to characterize the attributes of the SMDS traffic in the ATM network so that a clear picture of how other ATM traffic might affect the SMDS transfer can be drawn, and optimization can occur.

**Tester requirements.** A tester suitable for SMDS-to-ATM interworking, as with frame relay, must be focused primarily on verifying protocol conformance and network performance at the service level. Similarly, it also should be capable of ATM and Physical layer testing in order to allow correlation through the protocol stack from the Physical and ATM Cell layers, through to the SMDS Interworking Service layer. To make these tests possible, the tester must include:

- Two full-duplex test ports of the interface rate(s) in the network, including ATM port protocol support for SMDS over ATM, and native SMDS port test capability.
- SMDS-over-ATM to native SMDS data verification.
- Simulation of SMDS and SMDS-over-ATM traffic, alarms, and errors.
- ATM and Physical layer test capability, including alarm and error generation and measurement, QoS measurement (both in-service and out-of-service), and ATM layer traffic characterization.
- Correlated analysis of data on both SMDS and ATM ports, and through all levels of each protocol stack to the user service being carried.

### 11.3.3 ATM LAN interworking

There are several initiatives for connecting LANs and the higher-layer networking protocols (such as IP) over ATM. The ATM Forum has done extensive work on LAN emulation to allow ATM to be used in existing LAN environments; the Forum is now also expanding this work to cover the multitude of higher-layer networking protocols in the multiprotocol-over-ATM (MPOA) area.

At the same time, the Internet Engineering Task Force (IETF) has developed its own methods of encapsulating LAN and higher protocols over ATM. In particular,

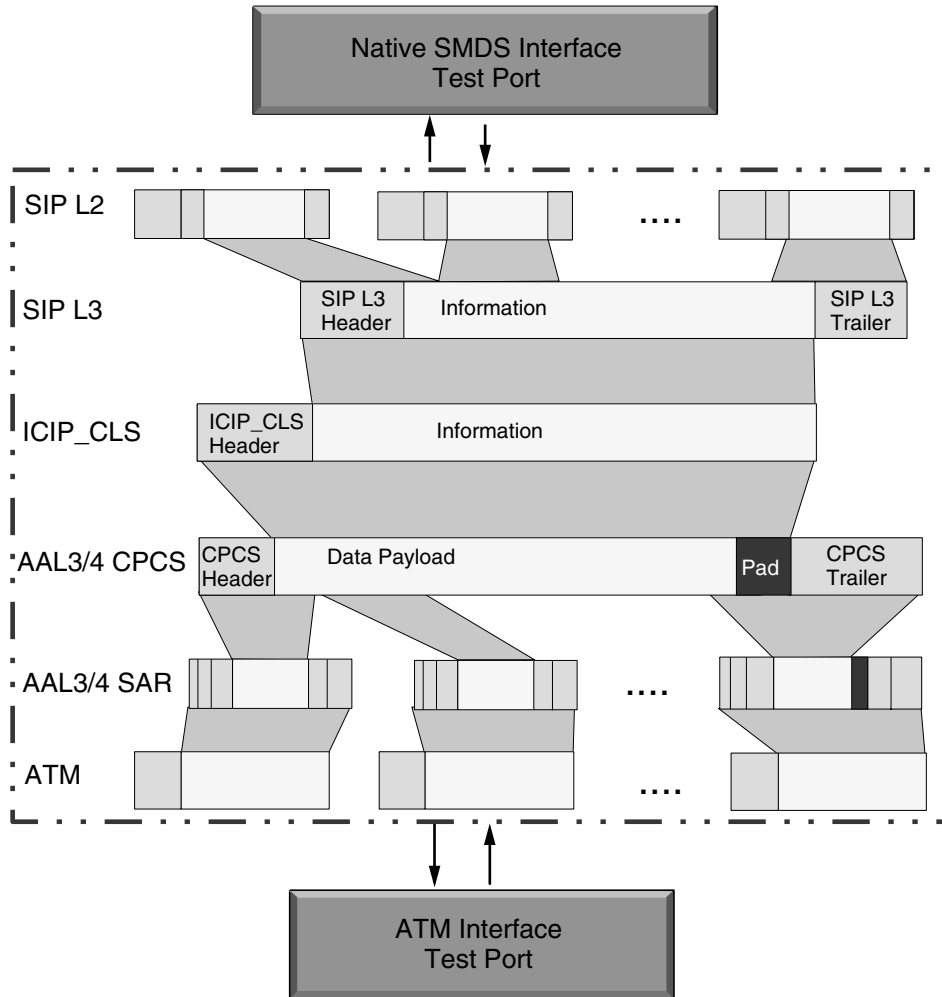


Figure 11.11 SMDS interwork testing.

IETF has developed techniques for using IP over ATM, enabling IP subnets of ATM stations (known as *Classical IP*) over ATM.

**Layer 2: LAN Emulation.** LAN Emulation (LANE) over ATM networks is designed to allow ATM to interwork with the legacy LAN technologies of Ethernet, Token-Ring, and FDDI. ATM, being a connection-oriented technology, is quite different in structure to the connectionless shared media upon which these legacy LANs are built. To allow ATM to become a compatible LAN technology that can be connected via bridges and routers to these other LANs, the ATM Forum has been developing the LANE specification. LANE allows not only interworking LANs over ATM, but also running existing LAN applications and protocols directly on ATM workstations.

The LAN Emulation (LE) architecture revolves around LE Clients (LECs) supporting the LE Service on a virtual LAN (Figure 11.12). The procedure starts with LEC initialization to determine which services are available and whether or not to join the emulated LAN. Registration then follows, letting the LE service know which MAC addresses and routing descriptors the LEC represents and how to handle unregistered addresses.

When data transmission is requested at the MAC layer, an address resolution request is made to translate the LAN address to the appropriate ATM address. A signaling procedure is then used to set up an ATM virtual channel to the destination LEC. Once established, the VC is used to transfer the MAC packets. The connection will be kept open for a certain period of time so that subsequent packet transfers can avoid further address-resolution and connection-setup stages, thus improving performance. If no further data transfers occur before the period runs out, the connection is released and any future data transfer has to use a new connection.

If a broadcast or multicast address is specified, or addresses are unknown, the packets are broadcast to all LECs using the Broadcast and Unknown Server (BUS). Note that the BUS can be used to provide connections for packet transfer while the address resolution process takes place, resulting in a change of connection once resolution is complete.

The data transfer itself involves encapsulating the MAC packet with an LE header and checksum in a PDU, which is then segmented using AAL 5 into ATM cells. Different LE data frames are specified for Ethernet-style LANs and Token-Ring LANs, to account for differences such as the reversed transmission order of bits in each octet.

LAN Emulation places heavy demands on the ATM network. In addition to the actual transfer of the data across the network, transfer of a LAN packet to an ATM station

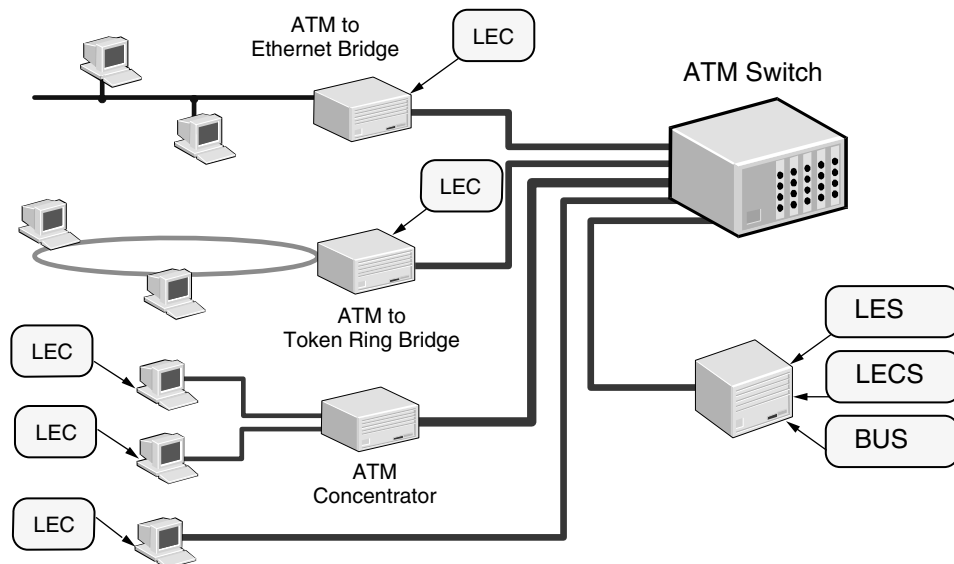


Figure 11.12 LAN Emulation (LANE).

involves a protocol exchange between the LEC and various servers (LECS, LES, and BUS) to register, resolve addresses, and set up the connection. Key implementation issues include accommodating faults in any of these server functions, and developing strategies for distributing them around the network (as opposed to locating them on the one machine).

Performance issues include the efficiency of the BUS and other server processes, along with encapsulation and data transfer performance. With many LAN protocols being “chatty” and sending a regular variety of broadcast traffic, a major implementation issue is how well the network handles the broadcast load—a factor that could be significant as more network segments are included). Routers might well be required to help this problem.

**Layer 3: IP over ATM.** With the widespread use and availability of TCP/IP protocols for networking workstations and PCs using existing network technologies, the IETF had a strong incentive for devising ways of using IP over ATM networks. This initiative (RFC 1577) allows existing applications designed to be used over IP to run directly over ATM networks without modification. The ATM network becomes an IP subnet and can be connected to other IP subnets using conventional router devices (Figure 11.13).

With Classical IP over ATM, the ATM end station maps its ATM address to an IP address and then can communicate with other IP stations on the network using ATM connections. These connections can be set up either permanently (PVC) or dynamically (SVC), using ATM address resolution protocol (ATMARP) and signaling. IP and ATMARP packet encapsulation uses AAL 5, as specified in RFC 1483. Note that RFC 1483 also specifies LAN encapsulation over AAL 5, but this is not the same encapsulation used for LAN Emulation.

**Layer 3: Multiprotocol over ATM.** MPOA is the ATM Forum’s initiative to provide a unified approach for the use of layer 3 protocols, such as IP and IPX, over ATM. It is an evolution of the LANE work and specifies LANE to be used when layer 2 bridging is required. MPOA is being designed to support fully routed environments, and in-




Figure 11.13 Classical IP over ATM.

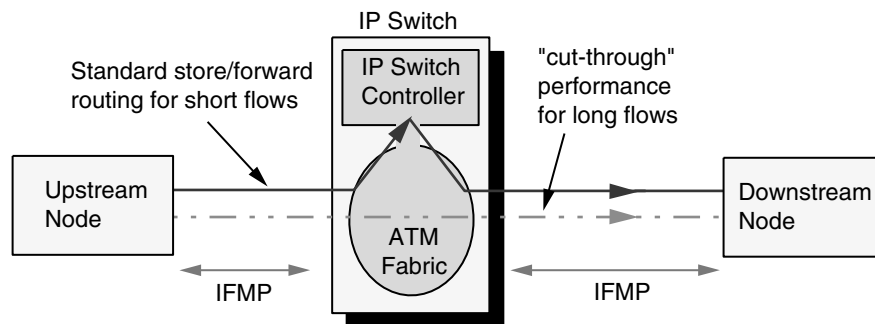



Figure 11.14 IP switching.

incorporates concepts for separating switching and routing functions (unlike traditional routers) by using separate dedicated connections to gain speed and efficiency. Unlike LANE or Classical IP over ATM, MPOA aims to make use of core ATM features such as ATM layer QoS capabilities. This work is still in progress.

**Layer 3: IP switching.** With the most intense battle for ATM acceptance being in the LAN and corporate backbone, some new technological developments have started to appear; they aim to make use of the speed of ATM, without the complexity of ATM protocols (such as signaling, LANE and MPOA), and therefore gain acceptance as a viable alternative to ATM.

Foremost among these developments is *IP switching* (from Ipsilon), where ATM switches classify flows of IP data transfer as either long or short. Short flows, such as those from SNMP and other IP protocols, are routed as normal by each IP switch. Long flows, such as those from FTP, are allocated direct ATM connections and are switched at ATM speeds through the network of IP switches (Figure 11.14).

Other rival techniques also have been announced, such as Tag switching (from Cisco). While these new technologies do appear to offer advantages over ATM technology, particularly in today's IP networks, they currently do not address the need for QoS capabilities. As users start to demand real-time service guarantees over the

Internet from a technology designed for “best-effort” nonprioritized transfer, ATM, having been developed specifically to handle this, may well return to favor.

**LAN-over-ATM tester requirements.** The deployment and operation of LAN over ATM services will require a variety of test techniques encompassing protocol verification and performance monitoring. It is important to be able to monitor the entire packet transfer process through the various stages of the communication. This requires monitor ports not only on the LAN itself, but also in both directions of ATM cell flow between the various end station clients and network servers.

Correlation will be required across the various monitor ports to decode and analyze the LAN data, the control protocols (including address resolution), and the signaling and data encapsulation for conformance to the appropriate standards. Emulation testing techniques are useful for verifying these functions during commissioning.

Performance testing also is required as part of this process to analyze the behavior of servers, bridges, and routers under stress conditions, and to ensure that fault-handling procedures behave as required. Additionally, the impact of ATM on the end station applications must be determined; for example, it is necessary to understand the effect of cell loss or CDV on transfer performance and network degradation with retransmissions. Continuity checks, (using “pings”) are needed; it is vital that data analysis and presentation be effective, so that users and support personnel can understand the real problems through the complexity of the technologies.

#### 11.3.4 ATM voice and video services

Video and audio services will appear both in corporate desktop and residential/consumer environments. Applications will range from high-bandwidth use (such as medical imaging), to broadcast video and desktop video conferencing, where compression will allow lower bandwidths to be used. The quality demands will vary widely, with low cost being a vital element of any mass-market application.

Despite the development of Circuit Emulation standards, existing voice services are not widely seen as a prime candidate for ATM initially, although audio services will appear as part of such applications as video conferencing and networked multimedia.

**Voice over ATM.** *Circuit emulation* is the name given to the transport of traditional TDM circuits across an ATM network in such a way that the ATM network appears to the TDM circuits as just another CBR link. This service is the simplest way to transfer voice circuits across an ATM network; because the data is real-time, however, delays must be kept to a minimum.

Of particular importance is controlling cell delay variation (CDV) or cell jitter. By the nature of ATM's multiplexing and switching, ATM networks introduce variable delay into a cell stream. Because circuit emulation traffic is CBR, it will accumulate CDV as it transfers across the network; this must be removed before it continues back onto the CBR TDM link. Buffering at the output of the ATM network could be used, but this adds to the absolute delays already experienced by the cells. It is important to make sure these parameters are controlled appropriately in the network, especially as the network grows and its utilization increases (Figure 11.15).



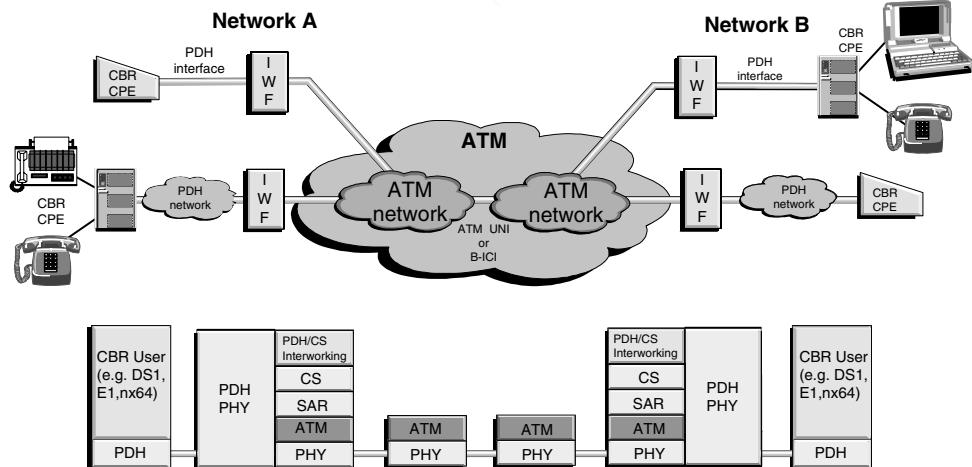


Figure 11.15 Circuit emulation service.

Circuit emulation transfers voice services over ATM by reserving the full TDM link bandwidth as a high-priority CBR connection. Other variable techniques have been developed that compress the voice signals into a VBR connection and allow better utilization of the network. As this technology matures, and ATM bandwidth becomes more fully utilized, circuit emulation might turn out to be no longer required.

In both cases, testing procedures must focus on ensuring that the delay and jitter requirements of the TDM signals are within required tolerances following transfer over the ATM network. In particular, it will be important to determine the effect of congestion in the ATM network on these values and the effectiveness of the traffic management functions designed to handle them.

**MPEG2 video over ATM.** The Motion Picture Experts Group (MPEG) is responsible for a set of specifications for the transfer of audio and video over digital transport systems. MPEG1 was developed for the transfer of VCR-quality signals; MPEG2 addresses broadcast-quality signals. The MPEG specifications include compression schemes, with the coded signal bandwidths giving the required quality for MPEG1 at about 1.5 Mbps, and for MPEG2 at about 6 Mbps. MPEG2 is seen as the technology suitable for video-on-demand (VoD) applications. With ATM being an obvious transport candidate, the ATM Forum (among others) has been working on the transfer of MPEG2 over ATM (Figure 11.16).

In essence, MPEG2 compresses and then packetizes the encoded video and audio signals for transport over a network. At the decoder, it synchronizes the transport stream and decodes the signals. Transport-stream packets are a fixed 188 bytes long and can include timing information for use in the synchronization process. A key element of the MPEG design is that a constant delay is expected between the encoder and decoder. ATM is highly prone to CDV (Cell Delay Variation), which means that successful implementation of MPEG2 over ATM must address this issue. Control of

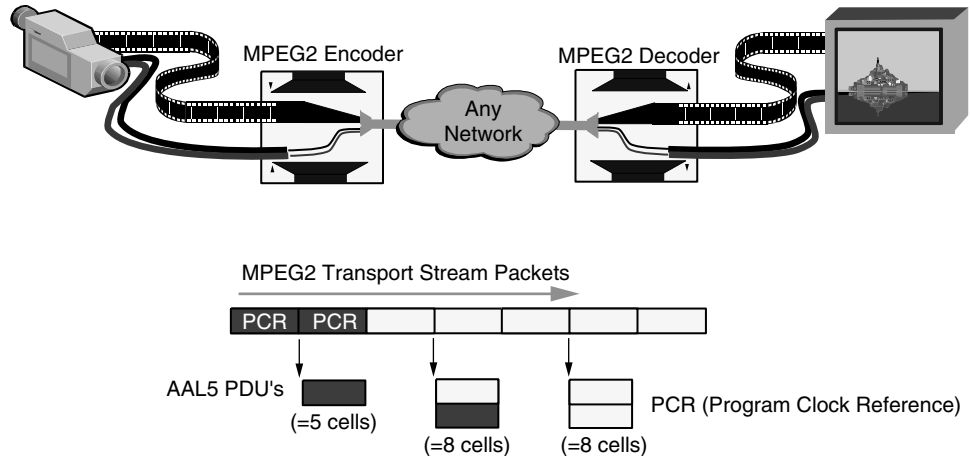


Figure 11.16 MPEG2 over ATM.

CDV requires buffering in the MPEG2 decoder, which in the case of VoD services would be the set-top box.

Other key challenges for deployment include detecting and handling errors from ATM impairments such as cell loss. The fixed-length transport packet most efficiently fits into AAL 1 payloads or AAL 5 payloads (if used in multiples of 2 packets per AAL 5 PDU). A major consideration in choosing the AAL type has been to decide which protocol layer should handle which function.

MPEG2 includes timing synchronization, so the equivalent function in AAL 1 is not required. In addition, AAL 5 is the more widely used, and is therefore cheaper to implement. It has been decided to use AAL 5 initially.

Note that although the default is to put two MPEG2 packets in each AAL 5 PDU, the timing requirements force packets containing the program clock reference (PCR) to be transmitted immediately if they are the first packet lined up for an AAL 5 PDU. Inefficient 5-cell PDUs therefore will be transmitted in these cases.

**Internet voice and video over ATM.** With the growth in popularity of the Internet and World Wide Web (WWW), the Internet is now being viewed as the preferred transport for new interactive voice and video services. New internetwork protocols, such as IPv6, *Real-Time Protocol* (RTP), and *Resource Reservation Protocol* (RSVP), have been developed with a view to supporting the real-time prioritized data flows that interactive voice and video services will require.

With ATM deployment increasing in the core of the Internet, deployment of these services will introduce a whole new set of interworking issues that test equipment must help solve. Service quality will depend not only on the performance of the ATM network, but also on the other network technologies used to carry these services to the end users. Effective testing techniques will have to be developed for these new services, but are likely to be heavily based on the techniques described through this chapter.

## 11.4 Summary

Deployment of ATM-based services is a complex task. In order to meet the unique challenges that ATM network operators will face, ATM test equipment dedicated to installation, commissioning, and troubleshooting tasks will be required. The following are important characteristics of such equipment:

- It must support multiprotocol testing through the layers of the protocol stack for the wide range of services in the network.
- It must have ATM layer test capability to enable the user to see the effect of service traffic on the ATM network and the effect of ATM impairments on service performance.
- It must provide Physical layer support for each interface used in the network.
- Synchronization of tests and correlation of measurements must be possible through all layers of the protocol stack and between multiple ports.
- The tester must be easy to use, allowing complex test procedures to be carried out by all levels of user skill, no matter their level of understanding of ATM.
- Test equipment must be portable and rugged, allowing easy transportation to the source of a problem. Equally, it must be capable of being left at the remote site and controlled back at base.
- Finally, it is important that test equipment be able to track the fast pace of technical standards development by being upgradable as new features are required.



---

Chapter  
**12**

## ATM Layer Testing

**David J. Terzian, P.E.**

*Hewlett-Packard Co., Westford, Mass.*

### 12.1 Introduction

Installing and maintaining Asynchronous Transfer Mode (ATM) networks presents new challenges for network operators as they strive to provide the highest-quality network for data, video, and voice applications. These challenges include provisioning ATM switches to provide the correct traffic parameters, quickly troubleshooting the network to isolate problems that might arise, and maintaining service quality as demand for network resources increases with the addition of new users.

This chapter provides an overview of installing and maintaining ATM networks. It begins by providing testing objectives, followed by the protocol stack model, and finally a description of the types of test equipment and practices used to install and troubleshoot networks.


The section on traffic management provides information relating to the types of services offered to customers and introduces the traffic parameters related to those classes of service.

There is a section examining the Quality of Service (QoS) measurements that can be made at the ATM layer to track network performance and help the service provider maintain the most reliable network. Switched Virtual Circuit (SVC) testing describes the types of tests that can be performed when signaling protocols are used to negotiate traffic and QoS parameters prior to establishing a connection between two end stations.

Rounding out the chapter is a section on Operations, Administration, and Maintenance (OAM) cells and their function in an ATM network, and a troubleshooting summary to help technicians determine the source of network problems.

#### 12.1.1 Testing objectives

At the ATM layer, 53-byte cells are transmitted through the network (Figure 12.1). ATM network testing will help ensure that the ATM layer is functioning properly. A



**Figure 12.1** The 53-byte ATM cell consists of a 5-byte header and a 48-byte payload. The header contains fields for the Generic Flow Control (GFC) used at the User-to-Network Interface (UNI) and the Virtual Path/Channel Identifiers (VPI/VCI), which are used to route cells through the network. In addition, the Payload Type Identifier (PTI) is used to denote what is contained in the payload, the Cell Loss Priority (CLP) is used to indicate whether the cell can be dropped during network congestion, and the Header Error Control (HEC) is used to maintain the integrity of the preceding information via a cyclic redundancy check (CRC). The payload is used to transfer user data or network status information.

well-functioning network provides end users with a trouble-free system, which in turn means that customers can focus on their main businesses rather than concern themselves with why applications might not be functioning properly.

A network that isn't working properly might cause symptoms such as choppy-appearing video conferences, repeated file retransmissions, and remote application timeouts caused by long delays between end-to-end communications. These problems might be caused by transmission delay, congestion, or an error-prone circuit.

Measurements at the ATM layer help identify network problems such as those described, and are useful for isolating their sources. Examples of parameters measured end-to-end include:

- Cell Transfer Delay
- Cell Delay Variation
- Bit Error Ratio
- Cell Loss Ratio

Once these parameters are measured, the network provider can compare the results to determine whether they exceed the requirements of particular applications. In addition, the network operator can track these parameters over time to ensure that network quality doesn't degrade as additional user traffic finds its way to the network.

## 12.2 The B-ISDN/ATM Protocol Stack Model

The ATM layer lies in the middle of the B-ISDN/ATM protocol stack model (Figure 12.2). This model is useful to differentiate the various functions that must be provided in an ATM network, and to help understand the functions a particular type of network equipment or workstation ATM adapter card performs.

At the top of the stack lie Services such as frame relay, which is used for TCP/IP transport. In addition, signaling services for call negotiation are done at this layer.

Just below the top of the stack is the ATM Adaptation layer (AAL). At this layer, frames representing Services are segmented into 53-byte ATM cells prior to transmission to the ATM network; at the destination they are reassembled into the original format. This layer enables information from variable-length frames to be neatly packaged in fixed-length cells for efficient routing through the network.

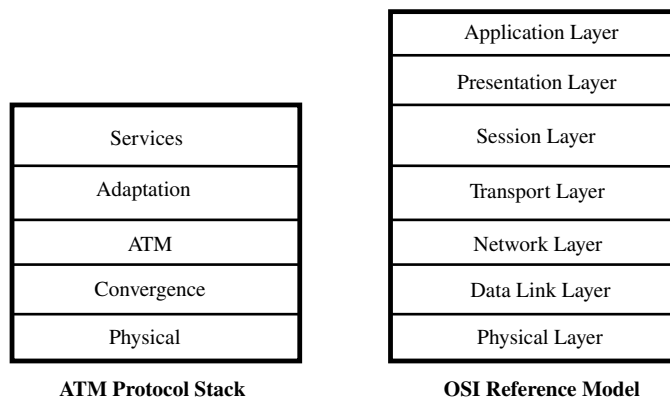
The adaptation function can be done at workstation ATM adapter cards, at ATM switches performing LAN emulation, or at video coders/decoders (codecs). Codecs segment Motion Pictures Experts Group (MPEG) streams into ATM cells and reassemble them at the destination so they are in the proper format for video applications.

The most common AAL types in use (one layer above the ATM layer) are AAL 1, AAL 3/4, and AAL 5. AAL 1 is used for services such as video also used as AAL 5 and circuit emulation, while the other types are used primarily for data transport. Each AAL type has a different coding scheme, which can include error-checking, cell sequencing, and clock recovery.

The ATM layer lies in the middle of the model. At this layer, 53-byte cells are sent through the network. ATM switches multiplex cell streams, route them to their correct destinations, and transmit them over the physical network.

The Convergence layer takes the fixed-length ATM cells and maps them onto the physical medium, a network that might be a Plesiochronous Digital Hierarchy (PDH), a Synchronous Digital Hierarchy (SDH), or a Synchronous Optical Network (SONET). At this layer, idle cells are inserted to compensate in cases where the full rate of the physical media is not being used by live traffic.

At the Physical layer, bits are transported from one point in the network to another. Due to the variety of LAN and WAN transmission equipment, in many cases these bits will travel over a variety of copper and fiber-based media.



**Figure 12.2** The B-ISDN/ATM Protocol Stack Model provides a reference to understand how the ATM layer is related to other layers in the B-ISDN/ATM protocol stack. At the top of the stack is the Services layer, which, for example, can represent frame relay traffic transmitted over an ATM network. This traffic is segmented into 53-byte cells at the Adaptation layer. At the ATM layer, the fixed-length cells are multiplexed and routed through the network. At the Convergence layer, cells are mapped into SONET/SDH or PDH frames for transmission at the Physical layer over fiber- or copper-based media. The lower three layers of the OSI model have similar functions to those of the ATM protocol stack.

There is not a direct correlation between the ATM protocol stack and the OSI Reference Model, but the lower three layers of the OSI model have functions similar to the ATM protocol model. The Physical layer in each model performs the same function of transmitting bits over a physical link and includes functions such as timing, signal levels, etc.

The Data Link layer in the OSI model describes the exchange of protocol data units (PDUs) and performs an error-detection function. The Network layer describes the delivery of reliable, in-sequence PDUs, or *packets*, and might perform data segmentation and reassembly. In the ATM protocol model, the segmentation and reassembly of data and error detection are performed at the AAL layer.

### 12.3 Functions of an Analyzer

ATM analyzers are used to verify the transport integrity of a network and to troubleshoot problems once they arise. In similar fashion to traditional BERT testers, which are used to check bit error ratios at the Physical layer (SONET/SDH or PDH), ATM analyzers can run a BERT at the ATM layer and monitor bit errors in the cell payload.

In addition, ATM analyzers are used to inject cell traffic into the network to make sure that services are provisioned correctly for customers prior to turn-up. ATM testers have the ability to transmit profiles of cell streams in order to emulate customer traffic.

ATM analyzers are used to characterize an ATM network for parameters such as cell transfer delay and cell loss (see section 12.8). These QoS tests are used to monitor ATM network performance to determine whether they exceed thresholds established by carriers for internal operations or in contracts with customers.

### 12.4 Test Equipment Overview

During the various phases of a technology lifecycle, the type of test equipment purchased changes. When a new technology such as ATM is introduced, test equipment for R&D typically is used widely, because standard equipment might not be available yet, and because people unfamiliar with ATM want to have the most features to help ensure that a robust ATM product is designed. During the growth phase, more intermediate test equipment is purchased at moderate prices. During the maturity phase, simple testers for large-scale field deployment are purchased.

A handheld for field use usually will contain the following features:

- Battery power
- Single physical interface
- Bit Error Ratio Test at the Physical and ATM layers
- Alarm and error generation and detection at the Physical layer
- Cell stream transmission to emulate various classes of service
- ATM traffic monitoring



- OAM capabilities
- QoS measurements
- Auto discovery of ATM virtual circuits, including bandwidth utilization
- Cell capture buffer

The PC-based tester for use in the CO or at the customer premises usually will contain these features, plus:

- Protocol analysis to determine traffic types and utilization
- AAL-type monitoring
- Graphical traffic display of particular virtual circuits
- Multiple physical interfaces for use in LANs or WANs

R&D analyzers are used by equipment manufacturers to design and test ATM switches and transmission equipment. These products are rich in functionality and provide comprehensive testing from the physical through the protocol layers. This equipment includes features from the PC-based tester, and in addition usually will have:

- Many physical interfaces for connecting to a wide variety of network types.
- Extensive physical layer tests, including the ability to control and monitor all overhead bytes for SONET/SDH and PDH frames.
- Full protocol decodes and the ability to transmit specific Protocol Data Units (PDUs).
- The ability to impair ATM transmission to simulate cell delay, cell delay variation, etc.
- AAL-type decoding and error monitoring.

Table 12.1 summarizes the advantages of each type of tester. Usually a manufacturer or service provider would employ a range of test equipment to fit its particular

**TABLE 12.1 Capabilities by Tester Type.**

Feature	R&D Analyzer	Laptop or Central Office Test Set	Field Handheld
Protocol Layer Testing	Extensive	Moderate to Extensive	None
AAL Layer Testing	Extensive	Moderate	Low
ATM Layer Testing	Extensive	Extensive	Moderate
Physical Layer Testing Features	Extensive	Extensive	Moderate
Physical Interfaces	Many	Several	Usually single
Price	High	Moderate	Low
Portability	Low	High	High

needs. For example, a manufacturer would require an R&D system to design and troubleshoot its ATM switches, and might also require some ATM handheld testers for its technicians to use during switch installation in the field.

On the other hand, a service provider might purchase an R&D system for its network operations center to troubleshoot difficult problems. In addition, it might purchase a handheld tester or a PC-based protocol tester for each Central Office (CO) containing an ATM switch.

## 12.5 Background for Testing ATM Networks

In order to install and maintain an ATM network, network managers and technicians must understand when in-service and out-of-service testing is appropriate, how to gain test access, and how Permanent Virtual Circuit and Switched Virtual Circuit testing differ.


### 12.5.1 In-service and out-of-service testing

In the traditional networks of the past, when a problem that could not easily be identified arose on a circuit, the traffic was rerouted through another path of the network and out-of-service testing was conducted to determine the cause of the problem. In contrast, ATM is unique in that virtual circuits can have varying bandwidths and do not necessarily have to occupy the full bandwidth of a particular physical link (such as fiber). That means testing can be conducted (albeit carefully) while live traffic is running over another virtual circuit on the same physical link (Figure 12.3). On the other hand, if a problem affects the whole physical link and cannot be easily resolved—such as with a noisy line—then traffic might be rerouted and out-of-service tests conducted in the traditional manner.

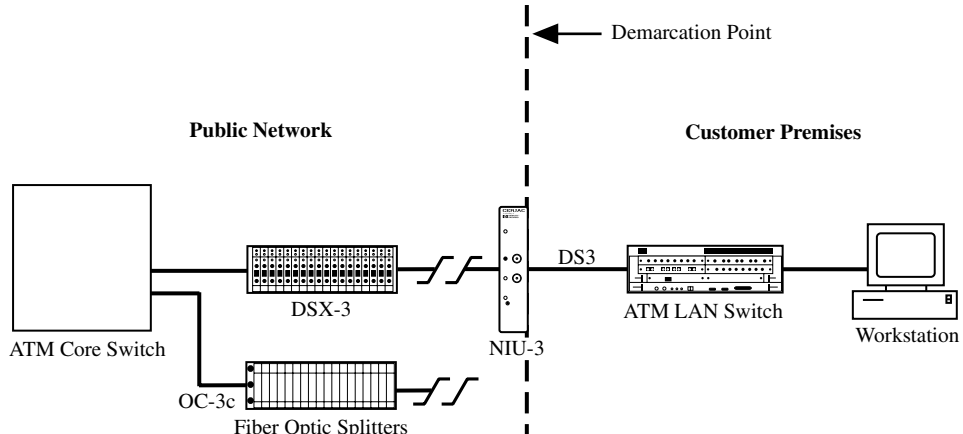
### 12.5.2 Test access points

There are several types of equipment in the network where technicians can gain access to ATM traffic. These include spare ports in an ATM switch, splitters in optical-carrier networks, cross connects (DSXs) for PDH networks, and an ATM DSU/CSU or Network Interface Unit (NIU) at the customer premises (Figure 12.4).

Nonintrusive access for monitoring network activity can be achieved at splitters and monitor ports in network cross connects and customer premises equipment. During installation or after network traffic has been rerouted, intrusive access can be obtained at an ATM switch port or at the customer premises. Particular care must be



**Figure 12.3** Testing can be conducted on a particular virtual circuit while traffic is running on another virtual circuit if customer traffic is not utilizing the full available bandwidth of the physical link. In this figure, there are three virtual channels (VCs) contained within the virtual path (VP). Note that bandwidth is reserved for management traffic.



**Figure 12.4** In this figure, depicting network access points for testing, ATM traffic is generated at the workstation and is transmitted over the LAN at 25 Mbps. Once traffic arrives at the ATM LAN switch, it is transmitted over the WAN using a DS3 interface in the switch. The traffic then travels through a DS3 cross connect and arrives at an incoming DS3 port in the ATM core switch. The traffic then is switched to an outgoing OC-3c port and is transmitted through the SONET network.

exercised, however, to ensure that test traffic injected into the network does not adversely affect other customer traffic running on the network.

### 12.5.3 PVC vs. SVC tests


The initial deployment of ATM networks consisted primarily of Permanent Virtual Circuits (PVCs). Similar to a leased line, PVCs are provisioned in advance and provide a fixed path for traffic to travel through the wide area network. Traffic and QoS parameters are provisioned in the network according to the contract the carrier has with each customer. Although these parameters might change somewhat with customer needs, they typically remain constant for a long period of time. The goal of the service provider is to ensure that the QoS is maintained, even as new customers are added to the network.

Testing PVCs is straightforward in that the circuit route and customer traffic and QoS parameters are predefined. Testing can be conducted easily over a particular path with more predictable results.

With the adoption of the ATM Forum's ATM User-Network Interface (UNI) Signaling Specification, Version 4.0, Switched Virtual Circuits (SVCs) over the WAN will become more prevalent. SVCs use signaling protocols to determine whether the network and end user can allocate appropriate bandwidth and QoS parameters in order for traffic to be transported from one point in the network to another.

SVC connections are established based on the most convenient path currently available in the network, similar to the way long-distance telephone calls are made. The advantage of SVCs is that they make more efficient use of the ATM network; unused capacity can be made available to other users.

SVCs add a layer of complexity to testing, since traffic and QoS parameters must be negotiated in advance to determine whether the network and the called user have



**Figure 12.5** With Switched Virtual Circuits (SVCs), cell traffic may travel from Boston to Chicago via switch B for one transmission and then may travel via switch C, located in a different city, for the next call transmission.

adequate resources to accept the transmission. With SVCs, the route traffic takes through the network might be changing constantly; as a result, troubleshooting becomes more difficult (Figure 12.5).

## 12.6 Installation Testing

The primary challenges that occur during network installation involve making sure that traffic can travel from one end of the network to the other. Therefore, prior to customer service turn-up, the following tests should be conducted to ensure that the network can handle the expected traffic conditions:

1. Physical layer BERT to ensure end-to-end transport integrity.
2. End-to-end ATM continuity test to ensure correct VPI/VCI ATM switch mappings for Permanent Virtual Circuits (PVCs).
3. Cell Header testing to protect against misdelivery of cells.
4. ATM BERT to ensure ATM layer transport integrity.
5. Transmission test to emulate customer traffic conditions.
6. QoS tests to determine if ATM switches have been correctly provisioned for the customer application.
7. OAM cell testing to determine if the ATM network responds appropriately.

The first four tests are addressed in this section; the last three are addressed in subsequent sections because they require more detailed explanations and also are appropriate for provisioning additional circuits and maintaining a high-quality network.

### 12.6.1 Physical layer Bit Error Ratio Test (BERT)

The Physical layer BERT should be conducted end-to-end prior to running ATM traffic over the network. This BERT will help determine if parameters are set properly throughout the network (for example, C-bit parity for DS3 networks), and will help isolate transmission problems if they exist on a particular segment of the network.

To run a physical layer BERT, the far end of the segment under test must be looped back to the tester. The tester then compares the outgoing BERT pattern with the pattern received to compute the Bit Error Ratio. Particular segments can be isolated until the source of the problem is determined (Figure 12.6).

Physical layer problems can be the source of problems at the ATM layer. To minimize problems at the Application layer, Bit Error Ratios at the Physical layer should not exceed a rate typically on the order of  $10^{-9}$  for PDH networks and  $10^{-12}$  for SONET/SDH networks. These tend to be rules of thumb; critical applications might require lower error ratios.

### 12.6.2 End-to-end ATM connectivity testing


Each ATM cell header has a *Virtual Path Identifier* and *Virtual Channel Identifier* (VPI/VCI), which directs cells through the network. These VPIs/VCIs have local significance only and can change from switch to switch. In addition, virtual connections are valid for one direction only.

ATM switches use lookup tables to determine how cells should be routed through the network. Cells arrive on an incoming port of an ATM switch and are delivered to the appropriate outgoing port in accordance with the switch's lookup table. The cells might have their VPI/VCIs changed when they are switched to an outgoing port.

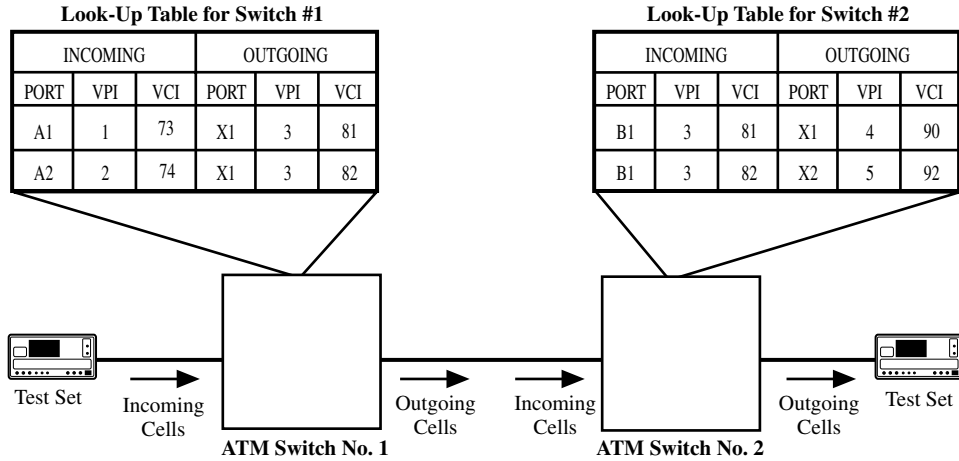
End-to-end ATM continuity tests help ensure that ATM switches have been provisioned to route cells correctly to their destinations (Figure 12.7). If cells do not arrive at their appropriate destination, then network segments can be tested separately to determine which switch might have incorrect VPI/VCI port mappings. Transmitting a stream of ATM cells from one point of the network to a destination will verify that end-to-end connectivity has been achieved.

### 12.6.3 Cell header testing

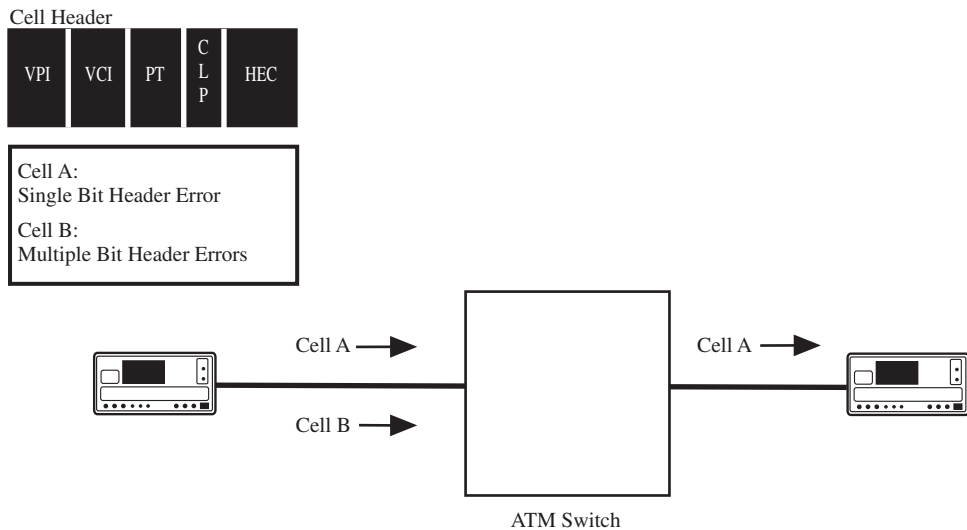
Cell header testing will verify that, where applicable, switches respond appropriately to errored cell headers. An error correction field (CRC-8) exists in the fifth byte of the header, which is called the *Header Error Control* (HEC). The purpose of this header error checksum is to provide protection against cells being delivered to incorrect locations. If the information in the preceding four bytes does not match the value computed in the fifth byte, then the switch should discard these cells when there is more than one bit in error and correct the header when there is a single-bit error (Figure 12.8).



**Figure 12.6** A Physical layer BERT run during installation to determine circuit quality between CO and customer premises.



**Figure 12.7** In this figure, illustrating cell VPI/VCI verification for PVCs, a test set transmits ATM cells with VPI=1 and VCI=73 into Switch No. 1, port A1. Switch No. 1 checks its lookup table, which indicates that the cell headers should be changed to VPI=3, VCI=81 and sent out port X1. The test set at the far end confirms that the outgoing cells of Switch No. 2 have VPI=4 and VCI=90.



**Figure 12.8** The Header Error Control (HEC) field is used to maintain the integrity of the first four bytes of the header by calculating a CRC-8 for these bytes. Cell A has a single-bit error, which can be corrected at the ATM switch, and is transmitted to the next switch in the network. Cell B has multiple errored bits, which has corrupted the VPI/VCI. This cell therefore is dropped by the ATM switch and does not get delivered to an incorrect location.

### 12.6.4 ATM Bit Error Ratio Test (BERT)

The purpose of the ATM BERT is to verify the ATM transmission integrity between two points of the network. An ATM BERT is run from a tester whereby cells are looped back at the far-end ATM switch. These cells—which contain the BERT pattern in the cell payload—might be returned on a different VPI/VCI.

During installation, a full-bandwidth (100 percent) ATM BERT can be run using the total capacity of the physical link. Once live traffic is running on the network, the bandwidth of the ATM BERT must be reduced to less than the available capacity of the link and is dependent on how much bandwidth is being used by live customer traffic and management traffic (See Figure 12.9).

## 12.7 Traffic Management

The ATM Forum Traffic Management Version 4.0 has expanded the classes of service to include both real-time and non-real-time Variable Bit Rate (VBR), and has added Available Bit Rate (ABR) service. ABR service allows service providers to make more efficient use of bandwidth in their networks by adjusting traffic levels to maximize the use of the network. In addition, QoS parameters have been expanded in the above specification; these will be discussed in section 12.8.


### 12.7.1 Classes of service

The ATM Forum has defined five classes of service that carriers can offer to customers. These services are tailored for particular customer applications, and each is specified by particular parameters. The five classes are:

1. Constant Bit Rate (CBR)
2. Variable Bit Rate, real-time (rt-VBR)
3. Variable Bit Rate, non-real-time (nrt-VBR)
4. Unspecified Bit Rate (UBR)
5. Available Bit Rate (ABR)

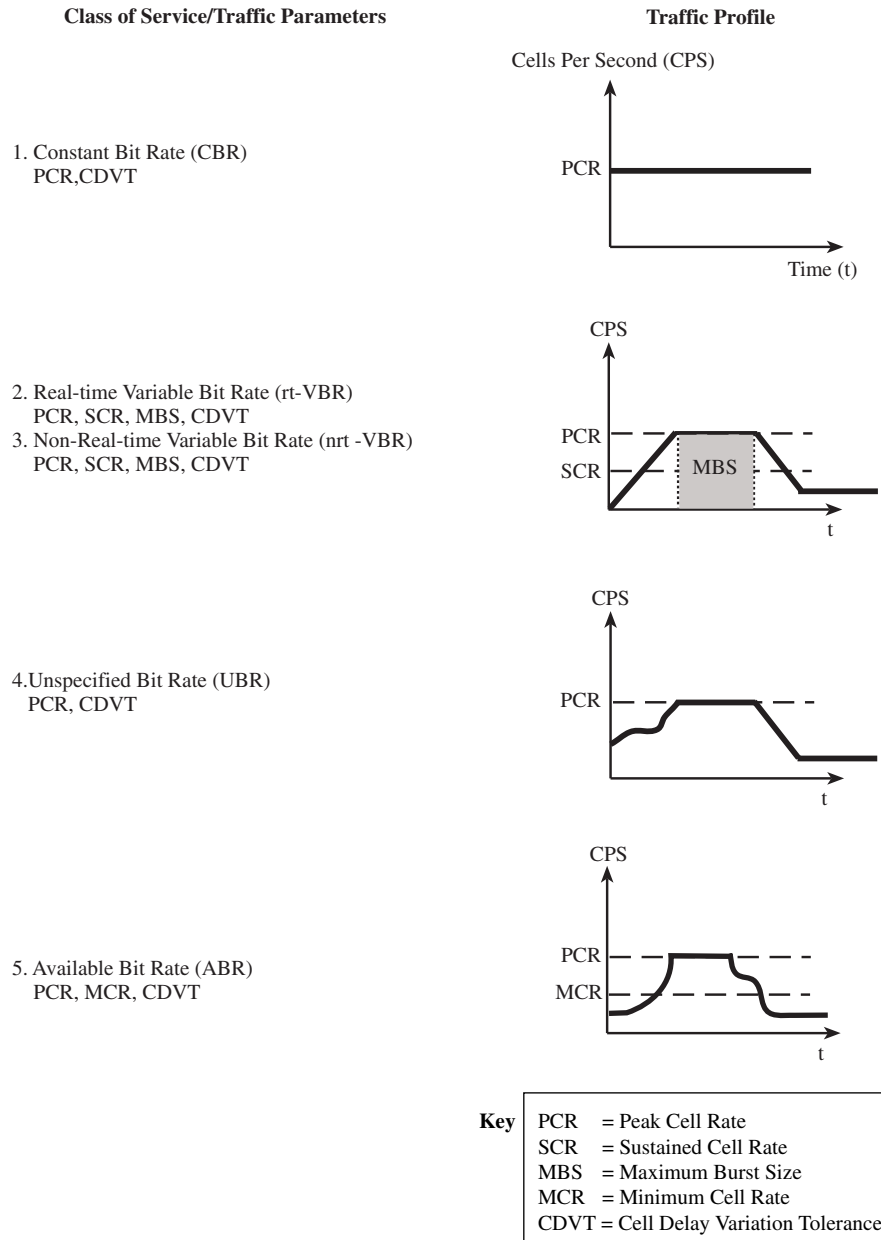
These classes of service are described by two or more of the following traffic parameters:

- Peak Cell Rate (PCR)
- Sustained Cell Rate (SCR)
- Maximum Burst Size (MBS)
- Minimum Cell Rate (MCR)
- Cell Delay Variation Tolerance (CDVT)



**Figure 12.9** A virtual circuit ATM BERT can be run at various percentages of bandwidths, ranging from very small up to 100 percent, which would consume the entire capacity of the physical link. This figure shows an ATM BERT being run at 20 percent bandwidth (every fifth cell). The other cells on the link might consist of switch management traffic, customer traffic, or idle cells.

Figure 12.10 illustrates which traffic parameters correspond to each class of service and provides traffic profiles for each. Subsequent sections describe each class of service in more detail.



**Figure 12.10** Each class of service has a set of traffic parameters, which characterizes its profile (solid line). Constant Bit Rate (CBR) is the simplest to characterize because the traffic rate remains constant. On the other hand, Available Bit Rate (ABR) uses signaling protocols to regulate traffic and maximize the efficiency of the network.



**Constant Bit Rate (CBR).** CBR services are used for video, voice, and distance learning applications where timing is important. This type of traffic has the highest priority in the network and is defined by the PCR, which remains constant, and by the CDVT.

The CDVT, usually specified in milliseconds, is the maximum cell jitter applications can tolerate in the network end-to-end. This might be caused, for example, by some cells in a stream getting delayed by multiplexing cell streams at a switch during congested traffic conditions.

The ATM network must be provisioned so the maximum cell delay variation during live traffic conditions will be less than the CDVT. Conservatively, this means that the sum of the CDVT parameters set at each switch in the transmission path must be less than the acceptable end-to-end cell jitter. In practice, however, the measured end-to-end cell jitter is likely to be less, since traffic traveling through the network is not likely to reach the maximum CDVT for each switch along the way.

**Real-time Variable Bit Rate (rt-VBR).** Real-time VBR can be substituted for some CBR applications having end systems that can recover from variable traffic rates and small cell losses. This service, in addition to being defined by the PCR and CDVT, uses MBS and SCR. The MBS is the maximum number of cells that can be transmitted at the PCR; the SCR represents the sustained traffic rate the customer is likely to use over time.

**Non-real-time Variable Bit Rate (nrt-VBR).** Non-real-time VBR can be used for applications such as transaction processing which do not have strict timing requirements. It is characterized by the same traffic parameters as those for real-time VBR. This service does have less stringent QoS requirements, however.

**Unspecified Bit Rate (UBR).** UBR is akin to flying standby; if a seat in coach becomes available, then the destination will be reached. This type of service can be used where no service guarantees are required, such as for e-mail and file transfer. It is defined by the PCR and the CDVT. In practice this service is not very common, since there are no guarantees for the subscriber.

**Available Bit Rate (ABR).** ABR is the newest class of service and is akin to having a seat in coach and waiting for an upgrade to first class. It was designed to make more efficient use of available capacity of the network, while at the same time providing a minimum guaranteed bandwidth, specified by the MCR. If the end-user application requires additional bandwidth—and such bandwidth is available in the network—then the PCR may be realized. ABR can be used for applications that require data transfer and distributed file service.

ABR implementations vary in their degree of sophistication. In its most basic form, the Explicit Forward Congestion Indication (EFCI) bit in the Payload Type (PT) within the cell header gets set to 1 to indicate congestion at a particular switch. The end station is notified that congestion exists in the network and, if this feature is implemented, may decrease the cell rate of the connection.

In an advanced implementation, ABR makes use of signaling to negotiate traffic and QoS parameters prior to establishing a connection (see section 12.9). Parame-

ters negotiated when establishing connections include the PCR, MCR, the Initial Cell Rate (ICR), Cells in Flight (CIF), and the Round Trip Time (RTT).

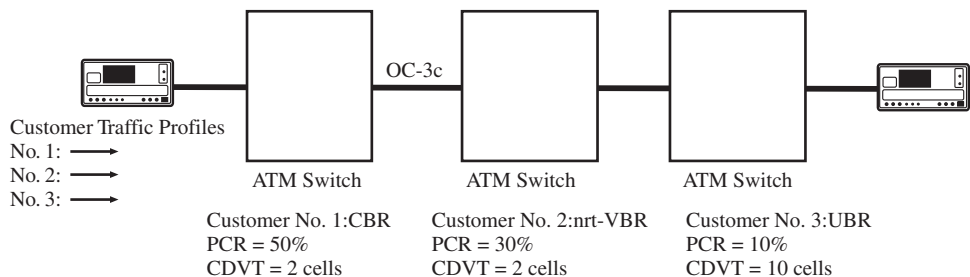
The CIF refers to the maximum number of initial cells that can be sent by the source prior to the return of the first Resource Management (RM) cell. RM cells, generated at switches and end stations, contain ABR parameters. These cells traverse the network and provide feedback on network conditions to the traffic sources so they can adjust transmission parameters accordingly. The maximum time required for an RM cell to go from the source to the farthest destination and back is the RTT. In addition to the MCR, some other traffic information contained in the RM cells include the Explicit Cell Rate (ER) and the Current Cell Rate (CCR).

### 12.7.2 Verifying service provisioning

Once a service provider meets with an end user and decides which type of service will best meet the customer's needs, an internal order will be placed to establish that service for the customer. Sometimes the specific service will be straightforward, while at other times it will be described in detail in a contract between the customer and the carrier. In either case, the service must be provisioned so that it will meet the specific needs of the customer. Provisioning might include providing new lines to the customer premises and customizing ATM switch settings.

Figure 12.11 illustrates three customers that have subscribed to different classes of service from the same carrier. In this example, all three customers use PVCs that have been provisioned in advance to meet their needs. Customer 1 is a medical center running a video telemedicine application between a medical school and a teaching hospital. Customer 2 is a large travel agency that uses the network to access reservation systems from airlines, hotels, and car rental companies. Customer 3 is a small engineering firm that uses the network for file transfer and e-mail.

Each of these customers requires specific traffic and QoS parameters, which must be met by the network. In order to verify that each PVC has been provisioned correctly, traffic profiles with the subscribed parameters can be transmitted through



**Figure 12.11** Each customer in this example has different applications with unique requirements from which traffic profiles can be derived. Customer 1 is a medical center running a telemedicine video application which requires low CDVT end-to-end. Customer 2 is a large travel agency that accesses reservation systems via an ATM network. Customer 3 uses its ATM network for non-time-critical applications such as e-mail and file transfers. With an ATM tester, the technician can preprogram three different transmit streams with profiles that match each class of service. QoS measurements are made to ensure that each customer will have acceptable traffic performance once it is brought online.

the network and corresponding measurements made to ensure the network will meet the needs of each customer.

The verification of service provisioning prior to customer turn-up serves two purposes. First, it provides the network operator with the confidence that ATM switch parameters have been set correctly for a particular customer according to its contract. Second, if customer application problems arise immediately after a customer is brought up, then troubleshooting can focus on whether the customer has specified appropriate traffic and QoS parameters for its particular application.

If customer problems persist even when the network has been provisioned correctly and the customer has specified the right parameters, then troubleshooting can focus on other areas. For example, the cause of the problem could be the result of network degradation from heavy traffic conditions.

### 12.7.3 Congestion testing

The performance of an ATM network depends on many factors, including the amount of traffic, the routes used to transmit traffic, the capabilities and performance of the switches, and the type of transmission equipment used in the network. ATM switches may use any of the following means to deal with traffic congestion:

- Prioritizing traffic according class of service
- Policing traffic for customers who violate their traffic parameters
- Dropping cells with Cell Loss Priority (CLP) of 1
- Discarding frames
- Shaping traffic

Service providers can test their networks in advance to determine how the network will respond to traffic congestion and whether this response was expected. Traffic priority schemes can be verified by transmitting two streams of traffic with different priorities and confirming that when congestion is reached, the high-priority cell stream remains unaffected.

Traffic policing for a virtual circuit can be checked easily by transmitting traffic at a rate exceeding the customer contract rate and ensuring that the excess traffic is discarded at the switch. In a similar fashion, cell discard can be checked by transmitting a traffic stream with CLP=1 while simulating congested conditions and verifying that cells from the stream with CLP=1 are discarded.

Frame discard raises the efficiency of ATM networks. Rather than discarding random cells during congested conditions, frame discard will drop AAL 5 frames at the switch. Then, at the Application layer, these frames can be retransmitted. In contrast, if an equivalent number of cells were dropped randomly, a greater number of frames would require retransmission.

Traffic shaping modifies the characteristics of a cell stream in order to enhance network efficiency; it can be provided at end stations or within ATM network elements. If a particular connection has a high cell delay variation on the incoming port of a switch, for example, the outgoing stream might be transmitted with lower cell delay variation to improve the traffic profile.

## 12.8 Quality of Service

Quality of Service (QoS) parameters are used with traffic parameters to establish a class of service to meet a particular customer's needs. Table 12.2 summarizes the traffic and QoS parameters applicable for each class of service. QoS parameters represent objectives for the network end-to-end.

QoS measurements are important because they allow service providers to determine whether they are meeting their contracts with customers and whether they are able to maintain or improve network performance as additional customers are added. QoS requirements obviously would be more stringent for telemedicine applications (CBR) than they would be for e-mail (nrt-VBR).

Four of the QoS measurements ATM analyzers make that can be used to track network performance over time and to troubleshoot network problems once performance begins to degrade are:

1. Cell Delay Variation
2. Cell Transfer Delay

**TABLE 12.2 ATM Service Categories.**

Traffic Parameters	ATM Service Layer Categories				
	CBR	rt-VBR	nrt-VBR	UBR	ABR
PCR and CDVT	■	■	■	■	■
SCR, MBS, CDVT		■	■		
MCR					■
QoS Parameters					
Peak-to-peak CDV	■	■	☐	☐	☐
Mean CTD	☐	☐	■	☐	☐
Maximum CTD	■	■	☐	☐	☐
CLR	■	■	■	☐	■
Feedback	☐	☐	☐	☐	■
Key					
CBR = Constant Bit Rate					
rt-VBR = Real-time Variable Bit Rate					
nrt-VBR = Non-real-time Variable Bit Rate					
UBR = Unspecified Bit Rate					
ABR = Available Bit Rate					
PCR = Peak Cell Rate					
SCR = Sustained Cell Rate					
MBS = Maximum Burst Size					
MCR = Minimum Cell Rate					
CDV = Cell Delay Variation					
CTD = Cell Transfer Delay					
CLR = Cell Loss Ratio					
CDVT = Cell Delay Variation Tolerance					
■ Specified					
☐ Unspecified					

Source: ATM Forum

**TABLE 12.3 QoS Objectives for MPEG-2 Applications Running over ATM.**

QoS Parameter	End-to-End Maximums
2-point CDV	1 ms
CTD	1 second for noninteractive video services
CLR	1 cell loss every 30 minutes
BER	$1 \times 10^{-10}$

Source: Bellcore GR - 2901 - CORE

### 3. Cell Loss Ratio

### 4. Cell Error Ratio

The first three QoS parameters are provisioned in advance for PVCs or are negotiated during call setup for SVCs. The fourth parameter is not negotiated during SVC setup.

The following sections describe these QoS parameters and how QoS measurements can be made in ATM networks. Table 12.3 summarizes QoS requirements for MPEG-2 applications running over ATM networks.

#### 12.8.1 Cell Transfer Delay


Cell Transfer Delay (CTD) measures the peak and mean delay cells experience while traveling from one point in the network to another. Transfer delay might be caused by transmission delay and switch processing delay. To run a CTD test, a stream of traffic (with cell timestamps) is transmitted from one point of the network to another. In practice—due to clock synchronization issues—the traffic is looped back to its source and the total time is divided by two to provide the CTD between two points in the network (Figure 12.12).

#### 12.8.2 Cell Delay Variation

Peak-to-peak Cell Delay Variation (CDV), or *cell jitter*, is the variation in delay from one to another that cells experience while traveling through an ATM network. Cell delay variation might result from cell queuing or multiplexing at switches, transmission through multiplexers, or video encoding. CDV is particularly critical for CBR applications where timing is important.

There are two measurements for CDV described by the ATM Forum and the American National Standard Institute (ANSI) for Telecommunications: *one-point CDV* and *two-point CDV*. The one-point CDV is measured in reference to the PCR at a particular point in the network and is the difference between the reference arrival time ( $c_k$ ) and actual arrival time ( $a_k$ ). Here is the equation for one-point CDV when  $c_0 = a_0 = 0$ :

$$c_{k+1} = \begin{cases} c_k + T & \text{if } c_k \geq a_k \\ a_k + T & \text{otherwise} \end{cases} \quad (12.1)$$



**Figure 12.12** Cell Transfer Delay (CTD) measures the peak and mean delay cells experience as they travel from one point of the network to another. In practice, cells usually are looped back to their source so that, for a given cell, the outgoing timestamp can be compared with the incoming timestamp.

where

$c_k$  = cell reference arrival time

$a_k$  = actual arrival time

T = Inter-arrival time between cells at the PCR (inverse of PCR)

Cells that arrive earlier than their expected time cause cell clumping (positive one-point CDV values), while cells that arrive later than their expected time cause gaps to occur in the cell stream (negative one-point CDV values).

The two-point CDV is the difference between the absolute transfer delay (CTD) of a cell between two measurement points in the network and a reference cell transfer delay between these same two measurement points.

### 12.8.3 Cell Loss Ratio

The Cell Loss Ratio (CLR) measures the percentage of cells lost between two points in the network.

$$\text{CLR} = \frac{\text{Number of Cells Lost}}{\text{Number of Cells Transmitted}} \quad (12.2)$$

Cell loss might result from traffic congestion at switches, traffic policing, protection switching, header errors, or physical media problems. In order to measure cell loss, a stream of traffic with cell sequence numbers is transmitted through the network. At the receiving end, the number of cells lost is measured to arrive at the CLR (Figure 12.13).

### 12.8.4 Cell Error Ratio

The Cell Error Ratio (CER) measures the accuracy of cell transmission. Errored cells are caused by cell payload errors in one or more bits.

$$\text{CER} = \frac{\text{Errored Cells}}{\text{Total Number of Cells Transmitted}} \quad (12.3)$$

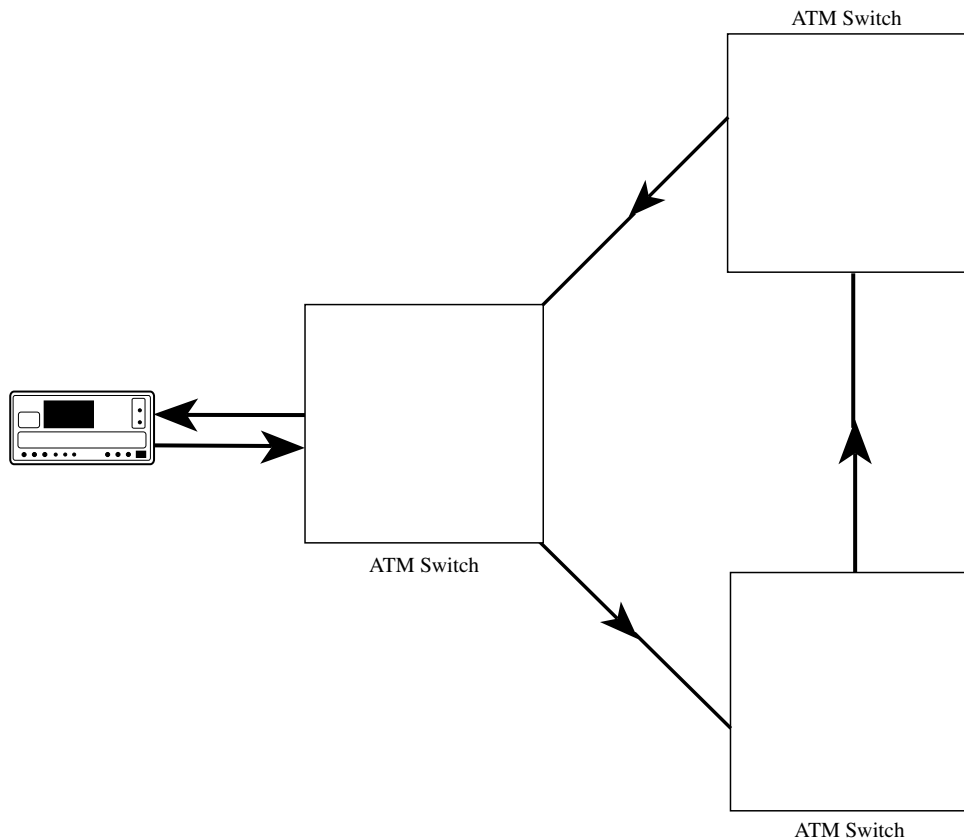
Errored cells usually result from problems at the Physical layer. It should be noted that the CER refers to the percentage of cells with one or more bit errors, while an ATM BERT will indicate the bit error ratio for all bits transported in the cell payload.

## 12.9 Switched Virtual Circuit Testing

As mentioned previously, SVCs use signaling protocols to negotiate traffic and QoS parameters in advance, to determine whether adequate resources exist to establish a connection between two end stations. Due to the dynamic nature of traffic patterns through the network as calls are established and taken down, troubleshooting becomes more difficult.

The following steps are required at the UNI to establish a connection between two endpoints:

1. The calling user transmits a *setup* message to the network.
2. The network must have the capacity to accept at least the minimum or alternative ATM traffic descriptors for the process to continue.
3. The network transmits a *setup* message to the called user.
4. The called user must have the capacity to provide at least the minimum or alternative ATM traffic descriptors for the process to continue.



**Figure 12.13** Cell Loss Ratio measures the percentage of cells lost within a path in the network. Cell loss can result from traffic congestion, traffic policing, or physical media problems.

5. The called user responds back to the network with a *connect* message with the traffic characteristics that have been accepted by the network and the called party.
6. The network transmits a *connect* message with the traffic parameters that have been assigned to the connection to the calling user.

There might be cases where network congestion precludes additional traffic from being transmitted onto the network. In this case, the message “user cell rate unavailable” is likely to be transmitted from the network to the calling user. In another case, the called user might be busy accessing a distance learning application while another user wants to set up a video conferencing connection; the message “resources not available, unspecified” might be returned to this calling user.

The following types of tests can be conducted for SVCs:

1. Transmit call *setup* messages to determine if the network and called user respond appropriately, either by proceeding with call establishment if resources are available, or returning an appropriate message describing the “otherwise” condition. The *setup* messages include the requested ATM traffic descriptors and either the Minimum Acceptable ATM Traffic Descriptor or the Alternative ATM Traffic Descriptor.
2. Transmit call *connect* messages to determine if network connections are established appropriately between two endpoints. *Connect* messages can be checked to determine whether they include the appropriate traffic parameters assigned to the connection.
3. Confirm connection by transmitting traffic profile with traffic and QoS parameters permitted by resources available in the network and at the called user station.

## 12.10 Operations, Administration, and Maintenance (OAM) Cell Testing

OAM cells are used to support fault management and performance monitoring at the ATM layer. This enables the exchange of information between different nodes in the network and alerts network operators of problems. Physical layer operations for SONET/SDH are referred to as F1–F3 Flows, while at the ATM layer F4 Flows are used for Virtual Path Connection operations and F5 Flows are used for Virtual Channel Connection operations.

At the Physical layer, this exchange of information is accomplished through the use of *overhead fields* associated with *signal frames*. At the ATM layer, network information exchange is achieved through the use of special cell formats, the most common of which is shown in Figure 12.14.

The OAM Cell Type (4 bits) distinguishes whether the function of the cell is activation/deactivation, fault, performance, or system management. The OAM Function Type (4 bits) distinguishes whether the cell is for notifying an alarm, performing a continuity check or loopback, or for reporting network performance information.

The Functions-specific field is 45 bytes and is unique for a particular cell type. The Reserved field (6 bits) is reserved for future use and the Error Detection Code (10 bits) is used for a CRC-10 error detection code to detect errored OAM cell payloads, thereby preventing a switch from processing corrupted information.






Figure 12.14 OAM cell formats.

TABLE 12.4 OAM Cell Format Summary for ATM Alarms.

Alarm	Flow Type	VCI	PTI	OAM Cell Type	Function Type
VP AIS	F4	4		0 0 0 1	0 0 0 0
VP RDI	F4	4		0 0 0 1	0 0 0 1
VC AIS	F5		1 0 1	0 0 0 1	0 0 0 0
VC RDI	F5		1 0 1	0 0 0 1	0 0 0 1

Table 12.4 shows the cell format summary for the alarms Alarm Indication Signal (AIS) and Remote Defect Indication (RDI). The AIS alerts downstream nodes of an ATM or Physical layer failure at the upstream node, while the RDI is generated at the termination node of the failed connection and alerts upstream nodes of a failure downstream. These alarms are generated at the rate of one cell per second.

During installation, ATM networks can be tested to determine whether switches respond appropriately to alarms generated by test equipment. The AIS and RDI alarms are used on PVCs but not SVCs.

It is expected that over time additional OAM cell functions will be defined and implemented for ATM networks.

## 12.11 ATM Troubleshooting Summary

Although ATM networks tend to be reliable, sometimes they can be affected by different types of problems that originate from various sources. In fact, sometimes a symptom can be caused by more than one source. Establishing hard-and-fast, reliable rules for troubleshooting network problems therefore becomes difficult.

There are basic tests that can be conducted to help isolate the source of problems, however. Some of these tests are described in Table 12.5 and provide the network operator with possible sources of problems, given particular symptoms experienced by customer applications.

In order to minimize potential problems once live customer traffic is running over the network, the best policy is for carriers to emulate customer traffic prior to service turn-up. This will help carriers anticipate possible problems and make necessary adjustments to improve service.

**TABLE 12.5 Troubleshooting Common Customer Service Problems.**

Customer Symptoms	Suggested Tests	Possible Causes of Symptoms
Slow response time of applications	Bandwidth and congestion	Traffic exceeds allocated bandwidth on a regular basis or “bursty” overloads. IP Packet loss causing retransmission of data.
Constant retransmission of data necessary	Bandwidth and congestion, BERT	Excess traffic, ATM switch buffer overflow, noisy circuit (ATM switch or transmission equipment).
Loss of service	Loopback tests to customer premises, Physical Layer Testing	CPE equipment, ATM switch, or transmission equipment failure.
Unidentified traffic arriving at customer site	ATM cell traffic scan, VPI/VCI verification, and Cell Misinsertion Rate.	ATM switch routing tables not configured correctly for customer site, cell misinsertion at ATM switch.
CBR or circuit emulation applications don't work properly (AAL 1).	End-to-end CDV, CLR	End-to-end CDVT not acceptable. CDVT set too high in switches, provisioned bandwidth inadequate.

# An Introduction to Synchronous Signals and Networks

**Doug Moore**

*Hewlett-Packard Ltd., South Queensferry, Scotland*

## 13.1 General

This chapter provides basic information on the synchronous signal structure, and to help the reader become familiar with the new telecommunications terminology that has emerged with arrival of synchronous systems. First, however, it is necessary to give some details about the older plesiochronous networks, and to describe the evolution of the new synchronous networks that replaced them.

More detailed material on synchronous network standards can be obtained from the documents listed in section 13.13.1 near the end of the chapter.

## 13.2 The Plesiochronous Network

Before the late 1980s, most high-capacity transmission networks were based on a hierarchy of digital multiplexed signals. Lower-rate tributary signals, such as the ITU-T 2.048 Mbps (E1) or North American 1.544 Mbps (DS1) were multiplexed in fixed asynchronous steps into a higher-rate signal for transmission.

Access to the individual tributary signals at each level in the hierarchy, for signal routing and test purposes, was provided by signal crossconnect points at the appropriate level in the multiplexing structure. Notice that because of the asynchronous nature of the multiplexing, gaining access to a lower-level tributary signal for rerouting or test purposes meant demultiplexing the whole line signal structure step-by-step down to the lower-level tributary data rate.

At each multiplexing step, the bit rate of the individual tributary signals was controlled within specified limits, but was not synchronized with the multiplex equipment. Because the tributary bit rates were controlled, this type of multiplexing is often referred to as being *plesiochronous*, that is to say, “nearly synchronous.” This

type of system is often referred to as a *Plesiochronous Digital Hierarchy* (PDH). Figure 13.1 depicts a network built upon PDH, with multiplexers at each node.

PDH networks were developed at a time when point-to-point transmission was the predominant network requirement. To support this requirement, the standard approach to network management and maintenance was to use manual distribution frames for access to individual signals. By the late 1980s this scenario was out-of-date. In addition, the PDH networks then in place had been found to severely limit the ability of the network operators to respond to the demands of an evolving telecommunications market. PDH networks are limited because:

- They are inflexible and expensive for telecommunications networking,
- They offer extremely limited network management and maintenance support capabilities, and
- Higher-rate line systems were proprietary.

For telecommunications networking purposes, flexibility is assessed in terms of how accessible an individual tributary signal on a particular line system is, so that it may be rerouted. PDH high-capacity line systems are not viewed favorably in this respect because access to any tributary signal cannot be obtained without demultiplexing the whole line signal, step-by-step, down to the appropriate level. From a cost perspective, gaining access to and rerouting a tributary signal covered only half of the equipment bill; the other half was incurred after rerouting, in remultiplexing

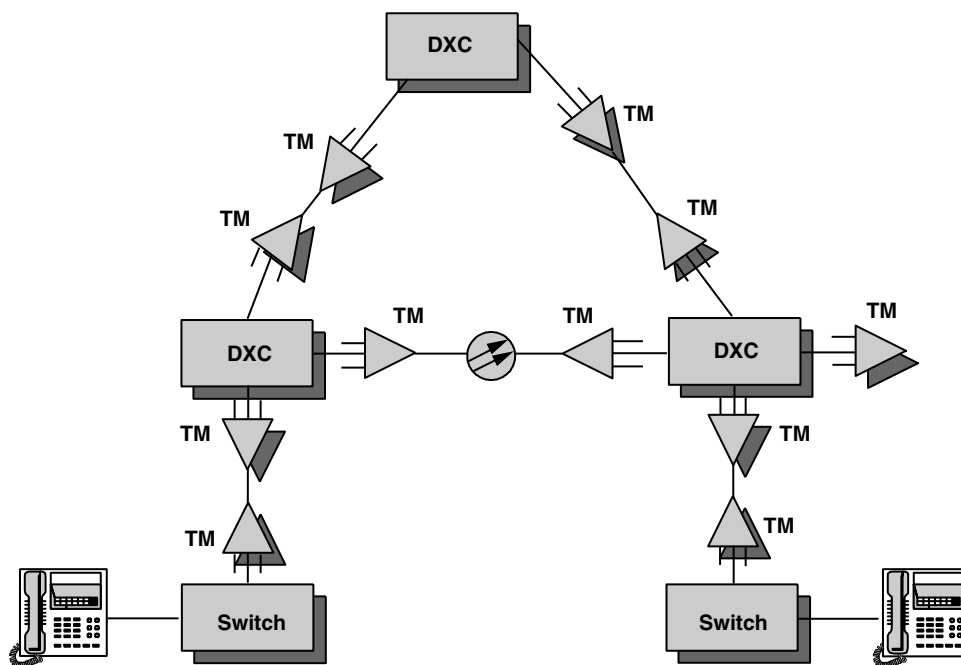


Figure 13.1 Example PDH network.

step-by-step back into the line signal for transmission. This makes plesiochronous multiplexing technology an expensive solution for telecommunications networking.

When originally conceived, network management and maintenance practices in PDH high-capacity networks were based on manual signal crossconnection and out-of-service testing techniques. There was no need to add extra capacity to the frame structures of the multiplexed signals for management and maintenance. As the complexity of the networks increased, however, and automatic computer management techniques became available, the lack of spare signal capacity in these signal frame structures severely limited the improvements that could be made.

A further limitation of PDH high-capacity line systems was that there was no common standard. Individual manufacturers of network equipment had their own proprietary designs. Both ends of the line system therefore had to be purchased from the same manufacturer. There was no possibility of interworking among components supplied by different manufacturers.

### 13.3 The Synchronous Network

The arrival of optical fiber as a transmission medium led to a rapid rise in achievable data rates and consequent available bandwidth within the telecommunications network. Coupled with the proliferation of automatic (microprocessor) control within the network, these developments opened the prospect of building extremely flexible and complex networks operating at high data rates. The limitations of PDH systems meant that it would have been very expensive, if not impossible, to take full advantage of these changes using the existing techniques.

To answer this need, committees within ANSI and ITU-T in the mid-1980s started to define a new network standard. The objective was to produce a worldwide standard for synchronous transmission systems that would provide network operators with a flexible and economical network. In 1985 the ANSI-accredited T1X1 committee started work on the *Synchronous Optical Network* (SONET) standard for North America; in June of 1986, ITU-T's Study Group XVIII started work on the Synchronous Digital Hierarchy (SDH) standards. SDH is the recognized international standard. The two standards were first published in 1988 and are broadly similar, the major difference being the base data rate used for multiplexing (see section 13.6).

Unless otherwise stated, the information in this chapter applies to both standards. A list of documents relevant to each standard can be found in section 13.13.1 near the end of the chapter.

The synchronous standards that were defined have the following advantages over the previous PDH standards:

- Flexible tributary extraction
- Built-in signaling
- Future-proofing
- Multivendor network capabilities

Designed for cost-effective, flexible telecommunications networking, the synchronous standards are based on the principles of direct synchronous multiplexing.

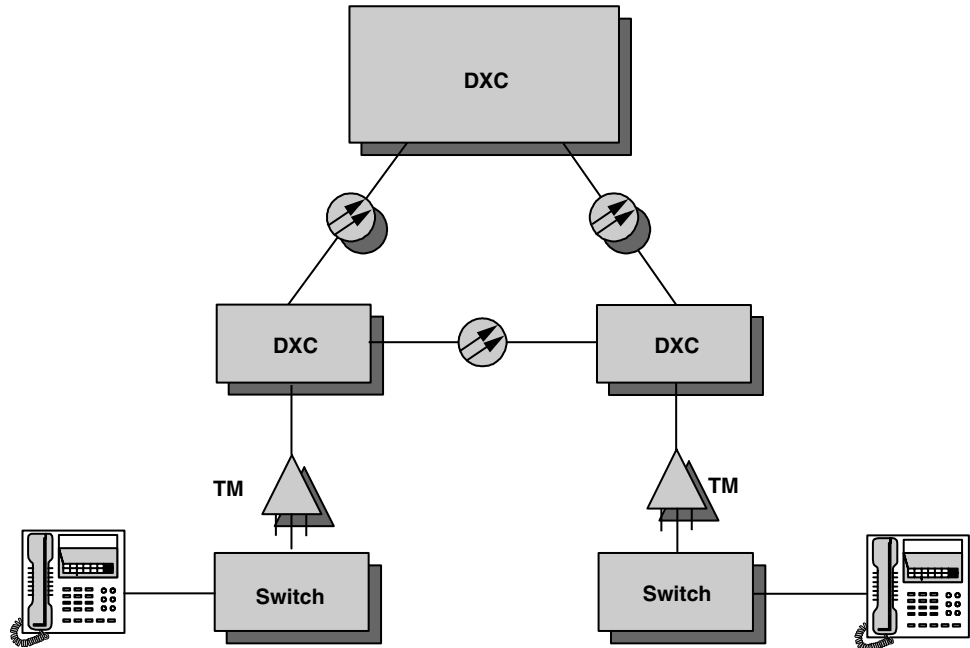


Figure 13.2 An example SDH network.

In essence, this means that individual tributary signals can be multiplexed directly into a higher-rate synchronous signal without intermediate stages. Synchronous *Network Elements* (NEs) then can be interconnected directly, with obvious cost and equipment savings compared to the existing network. Figure 13.2 shows an example SDH network; contrast its simplicity with the PDH network in Figure 13.1.

The signal structure provides built-in signal capacity for advanced network management and maintenance. Such capabilities are required in a flexible network in order to manage and maintain that flexibility effectively. Approximately 5 percent of the SDH signal structure is allocated to supporting network management and maintenance procedures and practices.

The synchronous structure provides a flexible signal transportation capability. The signal is capable of transporting all the common tributary signals found in the plesiochronous telecommunication networks. This means that the synchronous network can be deployed as an overlay to the plesiochronous network, and, where appropriate, provide enhanced network flexibility by transporting existing signal types. In addition, the standards have the flexibility to accommodate new types of customer service signals that network operators will wish to support in the future. Indeed, since the first standards documents were published in 1988, many new data formats have been “mapped” into the synchronous payloads. Probably the most important of these is Asynchronous Transfer Mode (ATM).

Synchronous structures can be used in all three traditional telecommunications application areas, namely long-haul, local network, and loop plant network. This

therefore makes it possible for a unified telecommunication network infrastructure to evolve. The fact that the synchronous standards provide a single common standard for this telecommunications network means that equipment supplied by different manufacturers may be interconnected directly.

### 13.4 Synchronous Signal Structure

A synchronous signal comprises a serial data stream of octets (bytes) that are organized into a frame structure. Within this frame structure, the identity of each byte is known and preserved with respect to a framing or marker byte. These frames are transmitted sequentially at a defined number per second, which is the frame rate.

For clarity, a single frame in the serial signal stream is represented by a two-dimensional map (Figure 13.3). The map comprises  $N$  rows and  $M$  columns of boxes that represent individual bytes of the synchronous signal; a  $B$  represents an information byte, and an  $F$  represents a framing byte (seen in the upper left corner of the  $N \times M$  matrix).

The signal bits are transmitted in a sequence starting with those in the top left corner byte (the  $F$  byte), followed by those in the 2nd byte in row 1, and so on, until the bits in the  $M$ th (last) byte in row 1 are transmitted. Then the bits in the 1st byte of row 2 are transmitted, followed by the bits in the 2nd byte of row 2, and so on, until the bits in the  $M$ th byte of the 2nd row are transmitted. The sequence continues through the remaining rows until the bits in the  $M$ th byte of the  $N$ th row are transmitted. Then the whole sequence repeats.

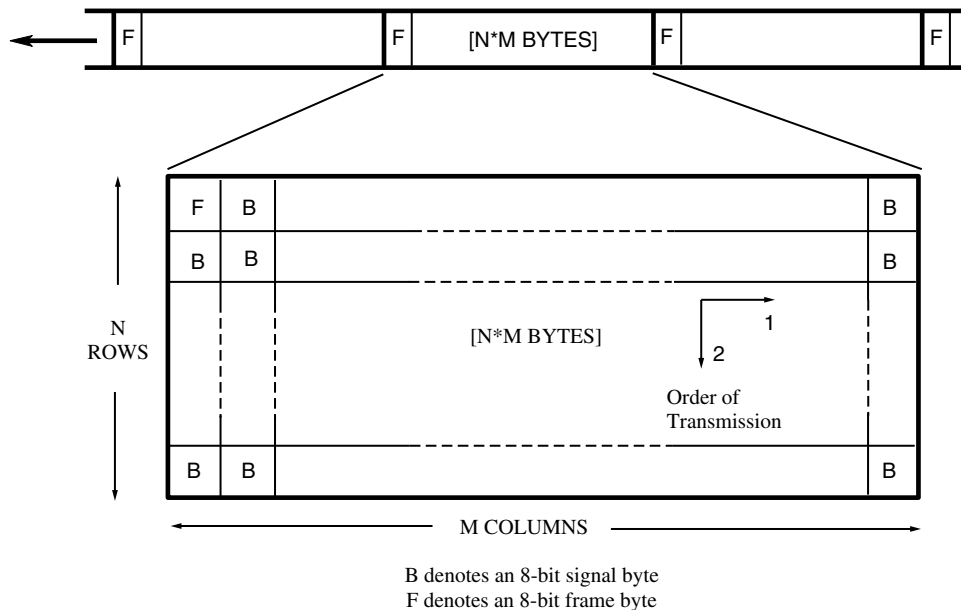


Figure 13.3 Basic synchronous 2-dimensional map.

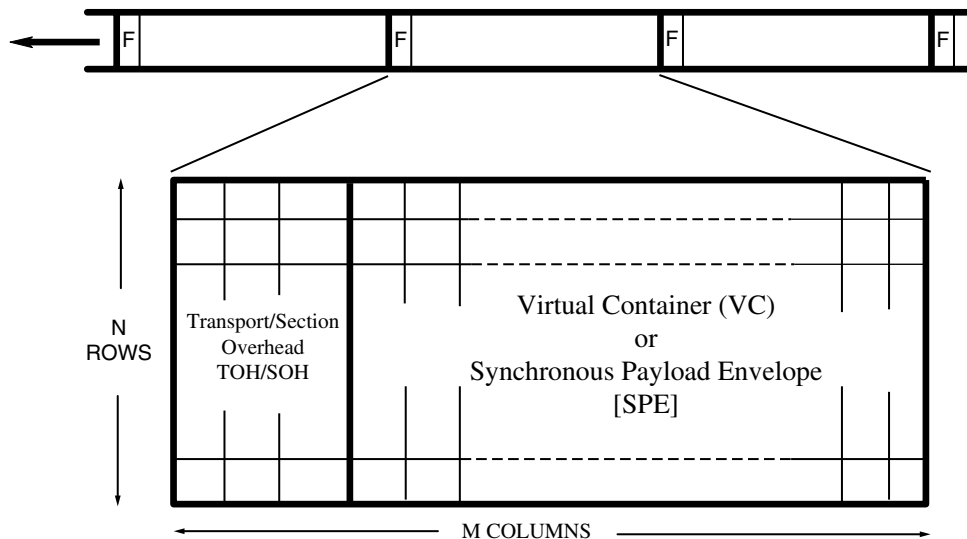


Figure 13.4 Synchronous transport frame.

### 13.5 Synchronous Transport Frame

The concept of transporting tributary signals intact across a synchronous network has resulted in the term *synchronous transport frame* being applied to such signal structures. More important, however, is that signal capacity is set aside within a synchronous transport frame to support network transportation capabilities. A synchronous transport frame therefore comprises two distinct and readily accessible parts within the frame structure, a *payload envelope* part and an *embedded overhead* part (Figure 13.4).

#### 13.5.1 Payload envelope

Individual tributary signals ( $DS_n$  in SONET, for example, or *Exx* signals in SDH) are arranged within a payload envelope, which is designed to traverse the network from end to end. Although it may be transferred from one transport system to another many times on its route through the synchronous network, this signal is assembled and disassembled only once. In SONET it is called the *Synchronous Payload Envelope* (SPE) and in SDH notation the *Virtual Container* (VC).

#### 13.5.2 Embedded overhead

Some signal capacity is allocated within each transport frame to provide the facilities (such as alarm monitoring, bit-error monitoring, and data communications channels) required to support and maintain the transportation of an SPE/VC between nodes in a synchronous network. The information contained in this overhead pertains only to an individual transport system and is not transferred with the SPE/VC between transport systems. This is called the *Transport Overhead* (TOH) in SONET and the *Section Overhead* (SOH) in SDH.



### 13.6 Base-Level Frame Structure

The major difference between the SDH and SONET standards is the base-level signal, from which all other signals are byte-multiplexed. The next two sections deal with these two frame types.

#### 13.6.1 SONET STS-1 frame structure

The base level SONET signal is called the *Synchronous Transport Signal level 1* (STS-1). The two-dimensional map for the STS-1 signal frame (Figure 13.5) comprises 9 rows by 90 columns, giving a total signal capacity of 810 octets, or 6480 bits, per frame. The frame repetition rate, or *frame rate*, is 8000 frames per second,<sup>1</sup> making the duration of each frame 125  $\mu$ s. With these frame dimensions and repetitions, the basic SONET signal structure bit rate works out to 51.84 Mbps:

$$810 \text{ bytes/frame} \times 8 \text{ bits/byte} \times 8000 \text{ frames/sec} = 51.84 \text{ Mbps}$$

The Transport Overhead occupies the first three columns of the STS-1 frame, a total of 27 bytes. The remaining 87 columns of the STS-1 frame, a total of 783 bytes, are allocated to the Synchronous Payload Envelope signal.<sup>2</sup> This provides a channel




Figure 13.5 SONET STS-1 frame structure.

1. At 8000 frames/second, each byte within the SONET signal structure represents a channel bandwidth of 64 kbps (i.e., 8 bits/byte  $\times$  8000 bytes/second = 64 kbps). This is the same bit rate as a PCM voice channel or a DS0 timeslot.


2. The SPE capacity of 50.11 Mbps ensures that the basic SONET signal frame may be used to transport the DS3-level tributary signal (at 44 Mbps) of the existing PDH networks.

capacity of 50.11 Mbps in the STS-1 signal structure for carrying tributary payloads intact across the synchronous network.

**13.6.2 SDH STM-1 frame structure**

The base-level SDH signal is called the *Synchronous Transport Module level 1* (STM-1). The two-dimensional map for the STM-1 signal frame (Figure 13.6) comprises 9 rows by 270 columns, giving a total signal capacity of 2430 bytes (19,440 bits) per frame. The frame rate is 8000 frames per second,<sup>3</sup> making the duration of each frame 125 μs. With these frame dimensions and repetitions, the bit rate of the basic SDH signal structure is 155.52 Mbps.

Transport Overhead occupies the first nine columns of the STM-1 frame, a total of 216 bytes. The remaining 263 columns of the STM-1 frame, a total of 2403 bytes, are allocated to the Virtual Container signal.<sup>4</sup> This provides a channel capacity of 150.34 Mbps in the STM-1 signal structure for carrying tributary payloads intact across the synchronous network.



**Figure 13.6** STM-1 Virtual Container (VC-4) frame structure.

3. At 8000 frames/second, each byte within the SDH signal structure represents a channel bandwidth of 64 kbps (i.e., 8 bits/byte × 8000 bytes/second = 64 kbit/s). This is the same bit rate as a PCM voice channel.

4. The VC capacity of 150.34 Mbps ensures that the basic SDH signal frame may be used to transport the E4-level tributary signal (at 139.264 Mbps) of the existing PDH networks. The virtual container associated with an STM-1 frame is referred to as a *Virtual Container level 4*, or VC-4. Virtual container levels 1, 2, and 3 are obtained by subdividing the VC-4. More details are provided in the relevant standards documents.

### 13.7 Synchronous Byte-Interleaved Multiplexing

To achieve data rates higher than the basic rates, groups of synchronous transport frames may be packaged for transportation as a higher-order synchronous transport signal. Higher-order grouping is achieved by the process of *byte-interleaved multiplexing*, whereby input transport signals are mixed together on a fixed byte-by-byte basis. The input signals are required to have the same frame structure and bit rate; they also must be frame-synchronized with one another.

For example, four parallel and frame-synchronized STM-1 signals may be byte-interleaved to form an STM-4 signal at 622.08 Mbps, four times the STM-1 bit rate. (This process is illustrated in Figure 13.7.) Similarly, three parallel and frame-synchronized STS-1 SONET signals may be byte-interleaved to form an STS-3 SONET signal at 155.52 Mbps (three times the STS-1 bit rate). Not all possible STS/STM-*n* signals are used, however; the most commonly accepted line rates are shown in Table 13.1.

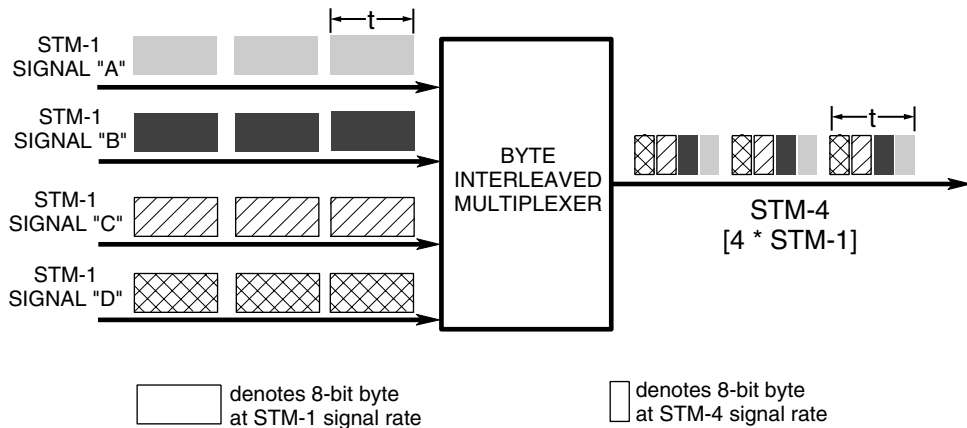


Figure 13.7 Synchronous byte-interleave multiplexing (STM-1 to STM-4).

TABLE 13.1 Byte Interleaving  
STM-1 to STM-4

SONET	Line Rate Mbps	SDH
STS-1	51.84	STM-0
STS-3	155.52	STM-1
STS-12	622.08	STM-4
STS-48	2,488.32	STM-16
STS-192	9,953.28	STM-64

Commonly used synchronous line rates


### 13.8 A Useful Analogy

The overhead-with-payload concept of the synchronous network can be thought of in terms of a more familiar road/rail containerized transport system (Figure 13.8). The container is analogous to the synchronous payload (SPE or VC), which stays intact along with its shipping information (the Path Overhead) from source to destination. The additional overhead added to the payload can be considered to be the “truck,” which ferries the container from node to node within the synchronous network. The truck might change at each node, but the container will stay intact.

Figure 13.9 shows an analogous network built up using these trucks and containers. The signal arrives at the network boundary at node 1, and is assembled into a container with the shipping information attached (in this case, destination node A). The container progresses through the network, changing its transportation at each node. It moves from truck to truck, and even is assembled (byte-interleaved) with other containers, each with its own destination information, into the multiple-container cargo of a train. The container is disassembled only when it reaches its destination at node A and leaves the transport network. In the same way, the SPE can be moved throughout the synchronous network, either at the same data rate (a truck) or at a higher data rate (a train) incorporated with other SPEs.

The Section/RS and Line/MS overheads (explained in the next section) then can be viewed as the actual mechanisms of the transport system—trucks, trains, etc.—ferrying the container from node to node on its journey through the network. The information required for carrying the container through each point is contained in the overhead, which can be changed at each node of the route through the network.

Though this is a somewhat crude analogy, it helps demonstrate the nature of the synchronous transport mechanism. As with all analogies, it has its weaknesses: In this case, it represents the network as unidirectional. The synchronous network, however, is intrinsically bidirectional. It uses this property to transmit information



**Figure 13.8** The synchronous “truck.”

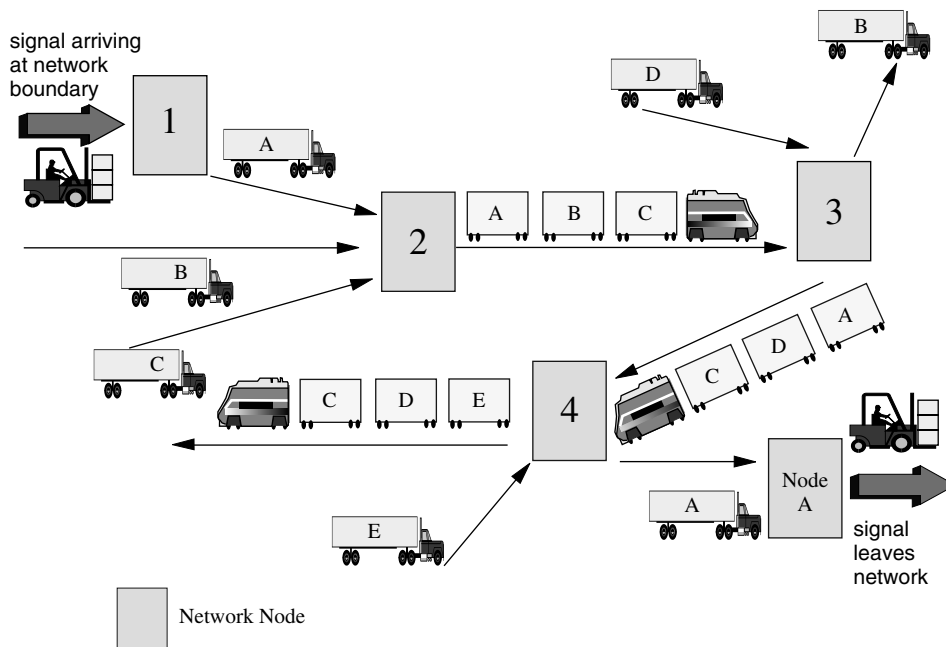


Figure 13.9 Synchronous network analogy.

about the incoming data stream back to the source NE, by altering the bytes in the outgoing signal's overhead. The next section deals with the embedded overhead that makes this control possible.

### 13.9 Embedded Overhead Capabilities

The synchronous transport frame carries two classes of data, namely the revenue-generating tributary signals plus the supporting network signals, the latter referred to as *embedded overhead*. Embedded overhead signals provide the functions needed by the network to transport the tributary signals efficiently across the synchronous network.

#### 13.9.1 Network spans

For network management and maintenance purposes, a synchronous network is similarly subdivided into three *spans* (the nomenclature for which differs slightly from SONET to SDH):

- The *Path span*, which allows network performance to be maintained from a service end-to-end perspective, i.e., from the point at which a tributary signal is assembled into its SPE/VC, to the point at which it is disassembled.
- The *Line or Multiplex Section (MS) spans* (in SONET and SDH, respectively), which allow network performance to be maintained between transport nodes.

302 Wide Area Networks

- The *Section or Regenerator Section (RS)* spans (in SONET and SDH, respectively), which allow network performance to be maintained between line regenerators, or between a line regenerator and a SONET Network Element.

An example of a simple, point-to-point network is shown in Figure 13.10 and Figure 13.11, with the three section types highlighted. Each span is provided with its own overhead, hence there are three categories of overhead. Each overhead provides the support and maintenance signals associated with transmission across that segment.

Embedded Overhead is split into three categories:

- SONET and SDH *Path Overhead*
- SDH *Multiplexer Section (MS) Overhead* or SONET *Line Overhead*
- SDH *Regenerator Section (RS) Overhead* or SONET *Section Overhead*

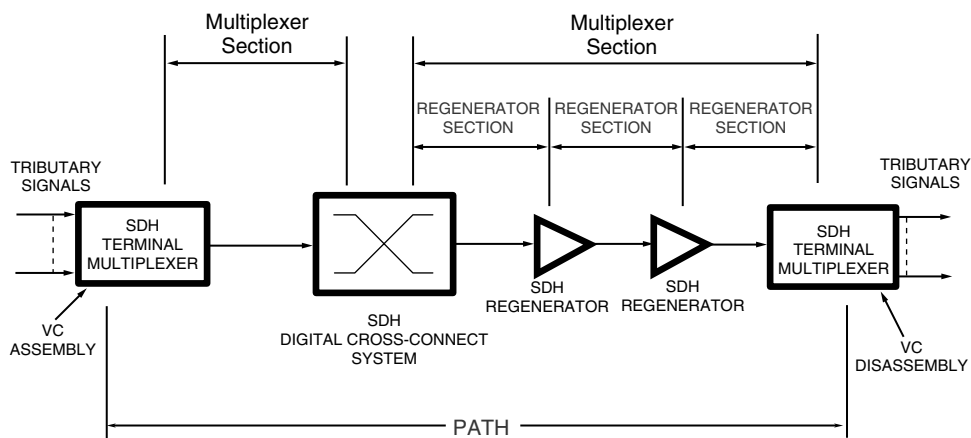


Figure 13.10 SDH network segments.

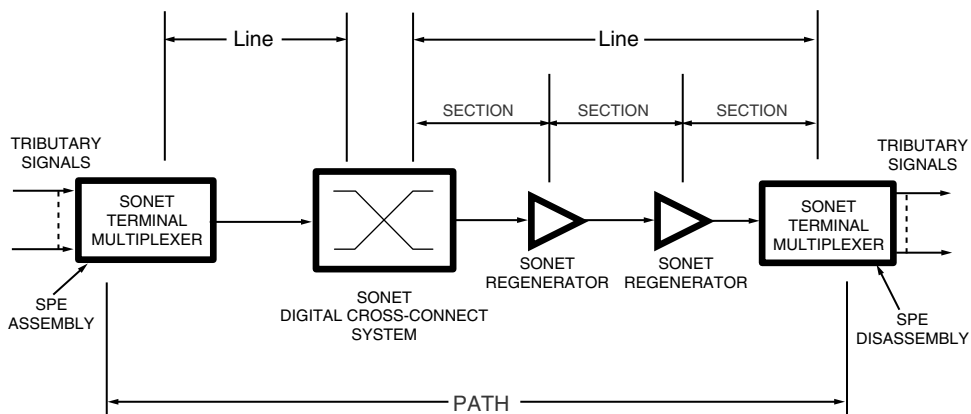


Figure 13.11 SONET network spans.

### 13.9.2 The overhead areas

The Path Overhead comprises 9 bytes and occupies the first column of the payload envelope. Path Overhead is created and included in the SPE as part of the SPE assembly process, and it remains as part of the SPE for as long as the SPE stays assembled. The Path Overhead provides the facilities required to support and maintain the transportation of the SPE between path-terminating locations, where the SPE is assembled and disassembled.

The Line/MS and Section/RS Overheads provide facilities to support and maintain the transportation of the SPE between adjacent nodes in the synchronous network. These facilities are included within the Transport Overhead part of the transport frame, the exact size being determined by the format.

Transport Overhead in the SONET STS-1 frame, comprising the first three columns of the frame, is split between Section Overhead and Line Overhead (Figure 13.12). The Section Overhead occupies the top three rows of the Transport Overhead, for a total of 9 bytes in each STS-1 frame. The Line Overhead occupies the bottom six rows of the Transport Overhead for a total of 18 bytes in each STS-1 frame.


Transport Overhead in the SDH STM-1 frame, comprising the first nine columns of the frame, is split between Regenerator Section Overhead and Multiplex Section Overhead (Figure 13.13). The RS Overhead occupies the top three rows of the Transport Overhead, for a total of 27 bytes in each STM-1 frame. The MS Overhead occupies the bottom six rows of the Transport Overhead, for a total of 54 bytes in each STM-1 frame.

### 13.9.3 The overhead bytes

The bytes within each overhead area perform specific functions inside the synchronous network. This section will deal with each overhead area in turn, giving a brief description of the active bytes. The general behavior is common to SDH and SONET; exceptions and divergences will be pointed out.

**The Path overhead bytes.** There are nine bytes of Path overhead carried in the VC-4/SPE, which are shown in Figure 13.14. Their functions are as follows:

- J1** The J1 byte supports continuity testing between any receiving terminal along the path and the path source. It is used to repetitively transmit either:
  - Mode 1: A 64-byte, fixed-length string, or
  - Mode 2: A 16-byte message consisting of a 15-byte string and 1-byte header containing a CRC-7 checksum.
- B3** The B3 byte provides a Bit Interleaved Parity (BIP-8) “path” error monitoring function. The path BIP-8 is calculated over all bits of the previous VC-4/SPE, the computed value being placed in the B3 byte before scrambling.
- C2** The C2 byte indicates the construction of the associated container by means of a label value assigned from a list of 256 possible values.



**Figure 13.12** Path overhead bytes.

- G1** The G1 byte is used to send status and performance monitoring information from receiving path terminating equipment to the originating equipment. This allows status and performance of a two-way path to be monitored at either end, or at any point along the path.
- F2** Allocated for network operator communications between path terminations.
- H4** Multiframe phase indication for VT/TU structured payloads.
- F3** *SDH*: Allocated for network operator communications between path terminations.
- K3** *SDH*: Provides the protocol sequences that control Automatic Protection Switching (APS) of the path. This functionality provides further support for SDH networking capabilities.



**N1** SDH: Tandem Connection and Path data byte.

**Z3, Z4** SONET: Reserved for growth.

**Z5** SONET: Same as N1 in SDH.

**Line and Multiplexer Section overhead bytes.** The defined bytes of the MS/Line section overhead are made up as follows:

**B2** SDH: The three B2 bytes provide a BIP-24 “multiplexer section” error monitoring function. The MS BIP-24 is calculated over all bits of the previous STM-1 frame except those located in the Regenerator Section overhead. B2 bytes are provided for all STM-1s in an STM-*n* frame structure.

**B2** SONET: The B2 byte provides a BIP-8 “line” error monitoring function. The line BIP-8 is calculated over all bits of the line overhead and payload envelope capacity of the previous STS-1 frame before scrambling, and the computed value is placed in the B2 byte before scrambling. This byte is provided for all STS-1s in an STS-*n* frame structure.

Framing A1	Framing A1	Framing A1	Framing A2	Framing A2	Framing A2	Ident C1/J0		
BIP-8 B1			Orderwire E1			User F1		
Datacom D1			Datacom D2			Datacom D3		
Pointer H1	Pointer H1	Pointer H1	Pointer H2	Pointer H2	Pointer H2	Pointer H3	Pointer H3	Pointer H3
← B2	BIP 24 B2	→ B2	APS K1			APS K2		
Datacom D4			Datacom D5			Datacom D6		
Datacom D7			Datacom D8			Datacom D9		
Datacom D10			Datacom D11			Datacom D12		
Synch S1	Growth Z1	Growth Z1	Growth Z2	Growth Z2	MS FEBE M1	Orderwire E2		

Figure 13.13 STM-1 SOH bytes.

Framing A1	Framing A2	Ident C1/J0
BIP-8 B1	Orderwire E1	User F1
Datacom D1	Datacom D2	Datacom D3
Pointer H1	Pointer H2	Pointer H3
BIP-8 B2	APS K1	APS K2
Datacom D4	Datacom D5	Datacom D6
Datacom D7	Datacom D8	Datacom D9
Datacom D10	Datacom D11	Datacom D12
Synch S1	Line FEBE M0	Orderwire E2

Figure 13.14 STS-1 SOH bytes.

- K1, K2** K1 and K2 control Multiplexer Section or Line Protection switching. They are defined for the first STM-1 in an STM- $n$  frame and for STS-1 number 1 in an STS- $n$  structure.
- D4–D12** Bytes D4 to D12 provide a 576 kbps data communication channel between Multiplexer Section termination equipment. This message-based protocol channel is used to carry network administration and maintenance information. These bytes are defined for STM-1 number 1 of an STM- $n$ , and for STS-1 number 1 of an STS- $n$ .
- S1** Bits 5–8 of the S1 byte are the Synchronization Status Message (SSM), a 4-bit code indicating the quality level of the synchronization clock used to generate the signal.

- M0** *STS-1 only:* Line FEBE byte.
- M1** Line FEBE byte
- Z1, Z2** The Z1 and Z2 bytes are reserved for functions not yet defined.
- E2** The E2 byte provides an express order wire channel for voice communications between Multiplexer Section terminating equipment and is only defined for STM-1 number 1 of an STM-*n* signal.
- H1–H3** *SDH:* The AU pointer bytes are associated with, but not actually part of, the MS overhead. H1 and H2 contain the pointer information. The three H3 bytes are the “pointer action” bytes. H3 bytes are used to carry “live” information from a VC during the STM frame in which a negative pointer adjustment occurs. AU pointers are provided for all VC-3/4s in an STM-*n*.
- H1–H3** *SONET:* The three bytes H1, H2, and H3 facilitate the operation of the STS-1 payload pointer and are provided for all STS-1s in an STS-*n*.

**Section and Regenerator Section overhead bytes.** The bytes of the RS/Section overhead are made up as follows:

- A1, A2** A1 and A2 provide a frame alignment pattern (11110110 00101000). These bytes are provided in all STM-1s within an STM-*n*, and all STS1s in a STS-*n*.
- C1** In older network equipment the C1 byte is set to a binary number corresponding to its order of appearance in the byte-interleaved STM-*n* frame. It can be used in the framing and de-interleaving process to determine the position of other signals. This byte is provided in all STM-1s within an STM-*n*, and all STS-1s within an STS-*n*, with the first STM/STS-1 being given the number 1 (00000001).
- J0** In more modern equipment the J0 byte transmits repetitively a 16-byte message consisting of a 15-byte string and 1-byte header containing a CRC-7 checksum. This byte supports continuity test between sections.
- B1** An 8-bit-wide, bit-interleaved parity (BIP-8) providing error performance monitoring at the RS/Section level. This even parity check is computed over all bytes of the previous STM/STS-*n* frame (after scrambling). The computed value is placed in the B1 byte before scrambling. These bytes are defined for the first STM-1 in an STM-*n* frame, and the first STS-1 in an STS-*n* frame.
- E1** The E1 byte provides a local order wire channel for voice communications between regenerators, hubs, and remote terminal locations. These bytes are defined for the first STM-1 in an STM-*n* frame, and the first STS-1 in an STS-*n* frame.

- F1** The F1 byte is allocated for user's purposes and is terminated at all re-generator section-level equipment. These bytes are defined for the first STM-1 in an STM- $n$  frame, and the first STS-1 in an STS- $n$  frame.
- D1–D3** A 192 kbps message-based data communications channel providing administration, monitor, alarm, and maintenance functions between RS/Section termination equipment. These bytes are defined for the first STM-1 in an STM- $n$  frame, and the first STS-1 in an STS- $n$  frame.

### 13.10 In-Service Maintenance Signals

The ability of the synchronous network to generate alarm and performance monitoring data, and to propagate this information throughout the network, is one of the keys to the efficiency and flexibility of this system.

The wide range of alarm signals and parity checks built into the synchronous signal structure support effective in-service testing. Major alarm conditions such as *Loss of Signal* (LOS), *Loss of Frame* (LOF), and *Loss of Pointer* (LOP) cause an *Alarm Indication Signal* (AIS) to be transmitted downstream. Different AIS signals are generated depending upon which level of the maintenance hierarchy is affected.

In response to the different AIS signals, and detection of major receiver alarm conditions, other alarm signals are sent upstream to warn of trouble downstream. *Far End Receive Failure* (FERF) is sent upstream in the MS/Line overhead after MS/Line AIS, or LOS, or LOF has been detected by equipment terminating in a Multiplexer Section span. A *Remote Alarm Indication* (RAI) for a high-order path is sent upstream after Path AIS or LOP has been detected by equipment terminating a path. Similarly, a Remote Alarm Indication (RAI) for a low-order path is sent upstream after low-order Path AIS or LOP has been detected by terminating equipment. Figures 13.15 and 13.16 depict the alarm flow in SONET and SDH networks, respectively.

Performance monitoring at each level in the maintenance hierarchy is based on *Bit-Interleaved Parity* (BIP) checks calculated on a frame-by-frame basis. These BIP checks are inserted in the overhead associated with each of the three network maintenance spans. The FEBE signals are sent upstream to the originating end of a path.

Section 13.13.2 gives brief descriptions of alarms generated by the synchronous system at RS/Section, MS/Line, AU/STS Path levels.

### 13.11 Subdivision and Concatenation

The frame structures described above are tailored to carry a specific PDH data signal, namely DS3 for the SONET SPE and E4 for SDH VC4. The obvious question is, "How does the synchronous network carry payloads that differ from the rates that fill the SPE/VC?" The answer for lower-rate signals is by use of Virtual Tributaries (VTs) or Tributary Units (TUs).

#### 13.11.1 SONET Virtual Tributary structure

The SONET STS-1 SPE, with a channel capacity of 50.11 Mbps, has been designed specifically to provide transport for a DS3 tributary signal. Transport for a tributary

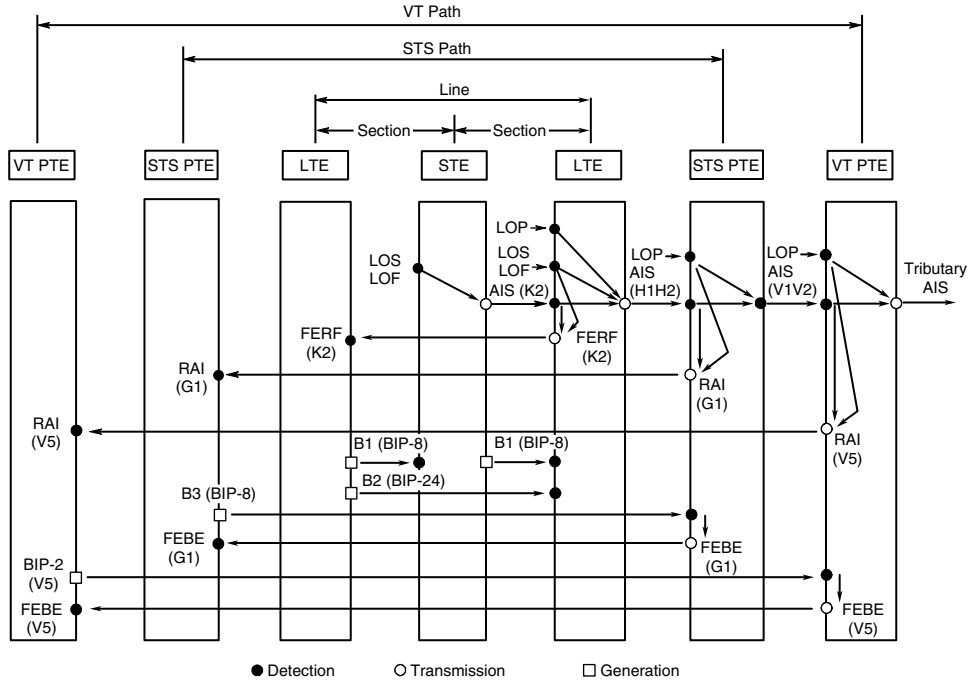


Figure 13.15 SDH alarm flow.

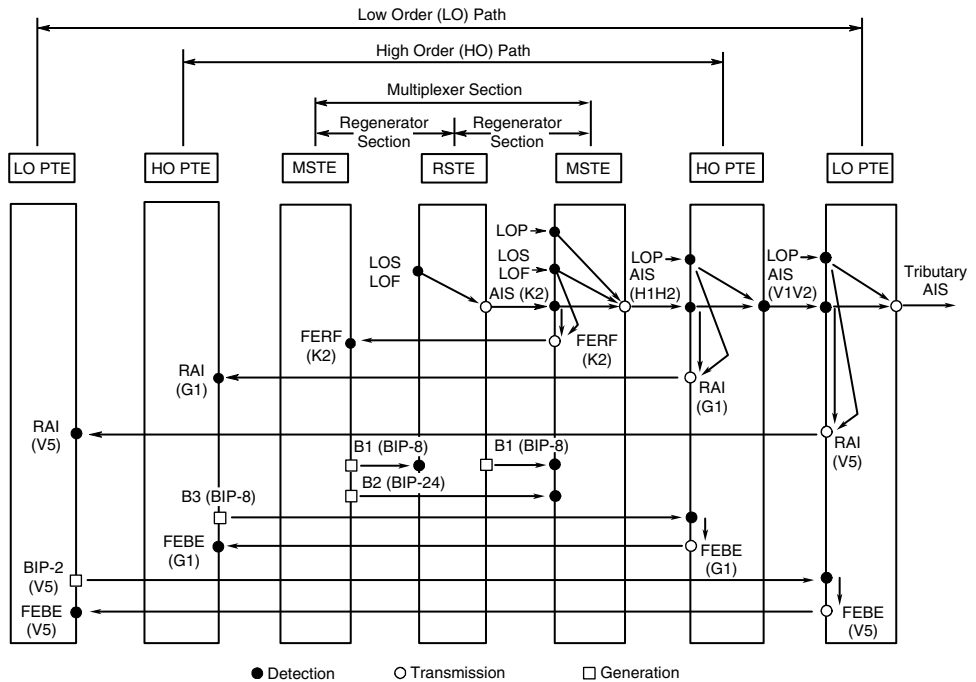


Figure 13.16 SONET alarm flows.

signal with a rate lower than that of a DS3 (such as a DS1, for example) is provided by a *Virtual Tributary* (VT) frame structure. VTs are specifically intended to support the transport and switching of payload capacity that is less than that provided by the STS-1 SPE. By design, the VT frame structure fits neatly into the STS-1 SPE in order to simplify VT multiplexing capabilities. A fixed number of whole VTs may be assembled within the STS-1 SPE.

**Virtual Tributary frame sizes.** A range of different VT sizes is provided by SONET.

- **VT1.5** Each VT1.5 frame consists of 27 bytes, structured as 3 columns of 9 bytes each. At a frame rate of 8000 Hz, these bytes provide a transport capacity of 1.728 Mbps and will accommodate the mapping of a 1.544 Mbps DS1 signal. Twenty-eight VT1.5s can be multiplexed into the STS-1 SPE.
- **VT2** Each VT2 frame consists of 36 bytes, structured as 4 columns of 9 bytes. At a frame rate of 8000 Hz, these bytes provide a transport capacity of 2.304 Mbps and will accommodate the mapping of a CEPT 2.048 Mbps signal. Twenty-one VT2s can be multiplexed into the STS-1 SPE.
- **VT3** Each VT3 frame consists of 54 bytes, structured as 6 columns of 9 bytes. At a frame rate of 8000 Hz, these bytes provide a transport capacity of 3.456 Mbps and will accommodate the mapping of a DS1C signal. Fourteen VT3s can be multiplexed into the STS-1 SPE.
- **VT6** Each VT6 frame consists of 108 bytes, structured as 12 columns of 9 bytes. At a frame rate of 8000 Hz, these bytes provide a transport capacity of 6.912 Mbps and will accommodate the mapping of a DS2 signal. Seven VT6s can be multiplexed into the STS-1 SPE.

**VT1.5 packaged STS-1 SPE.** The VT1.5 (Figure 13.17) is a particularly important virtual tributary size because it is designed to accommodate a DS1 tributary signal, which has the highest density of all the tributary signals that appear in the existing PDH networks. The VT1.5's  $3 \times 9$  column/row structure fits neatly into the same nine-row structure of the STS-1 SPE. Thus, as noted previously, 28 VT1.5s can be packed into the 86 columns of the STS-1 SPE payload capacity. This leaves two columns in the STS-1 SPE payload capacity as spares. These spare columns are filled with fixed stuffing bytes, which allow the STS-1 SPE signal structure to be maintained.

**Virtual Tributary structure.** The Virtual Tributary frame represents, in essence, a miniature transport frame structure. It has the attributes of a SONET transport frame, but it is carried within the standard SONET STS-1 frame. Thus a low-rate tributary signal can be mapped into the VT payload capacity. VT Path overhead is added to this payload capacity to complete the VT Synchronous Payload Envelope (VT SPE). The VT SPE is linked to the VT frame by a *VT Payload Pointer*, which is the only component of VT transport overhead. The VT frame then is multiplexed into a fixed location within the STS-1 SPE.

Although the VT frame structure is illustrated here as residing in one STS-1 SPE, it actually is distributed over four consecutive STS-1 SPE frames. It is,




Figure 13.17 VT1.5 packaged STS-1 SPE.

therefore, more accurate to refer to the structure of the VT as a *VT multiframe structure*. The phase of the multiframe is indicated by the functionality provided in the Path overhead.

### 13.11.2 SDH Tributary Units (TUs)

The channel capacity provided by an STM-1 VC-4 is 149.76 Mbps. This has been designed specifically to provide transport for a 140 Mbps E4 tributary signal. Transport for lower-rate tributary signals, such as 2 Mbps, is provided by a Tributary Unit (TU) frame structure. TUs are specifically intended to support transporting and switching payload capacity of less than that provided by the VC-4. By design, the TU frame structure fits neatly into the VC-4, thereby simplifying TU multiplexing. A fixed number of whole TUs may be assembled within the C-4 container area of a VC-4.

**Tributary Unit frame sizes.** A range of different TU sizes is provided by SDH.

- **TU11** Each TU11 frame consists of 27 bytes, structured as 3 columns of 9 bytes. At a frame rate of 8000 Hz, these bytes provide a transport capacity of 1.728 Mbps and will accommodate the mapping of a North American DS1 signal (1.544 Mbps). Eighty-four TU11s can be multiplexed into the STM-1 VC-4.
- **TU12** Each TU12 frame consists of 36 bytes, structured as 4 columns of 9 bytes. At a frame rate of 8000 Hz, these bytes provide a transport capacity of 2.304 Mbps and will accommodate the mapping of a CEPT 2.048 Mbps signal. Sixty-three TU12s can be multiplexed into the STM-1 VC-4.




Figure 13.18 TU12 packaged VC-4.

- **TU2** Each TU2 frame consists of 108 bytes, structured as 12 columns of 9 bytes. At a frame rate of 8000 Hz, these bytes provide a transport capacity of 6.912 Mbps and will accommodate the mapping of a North American DS2 signal. Twenty-one TU2s can be multiplexed into the STM-1 VC-4.
- **TU3** Each TU3 frame consists of 774 bytes, structured as 86 columns of 9 bytes. At a frame rate of 8000 Hz, these bytes provide a transport capacity of 49.54 Mbps and will accommodate the mapping of a CEPT 34 Mbps signal or a North American DS3 signal. Three TU3s can be multiplexed into the STM-1 VC-4.

**TU12 packaged VC-4.** The TU12 (Figure 13.18) is a particularly important size of tributary unit. This is because it is designed to accommodate a 2 Mbps tributary signal, the most common tributary signal in existing CEPT networks. The TU12's  $4 \times 9$  column/row structure fits neatly into the same nine-row structure of the STM-1 VC-4. Sixty-three TU12s can be packed into the 260 columns of payload capacity (i.e., the C-4 container) provided by a VC-4. This leaves eight columns in the C-4 container as spares. These spare columns result from intermediate stages in the "TU12 to VC-4" multiplexing process, and are filled by fixed stuffing bytes.

**Tributary Unit frame structure.** The Tributary Unit Frame essentially represents a miniature transport frame structure. It has the attributes of an SDH transport frame but is carried within the standard SDH STM-1 frame structure.

A TU frame is created by mapping a low-rate tributary signal in to the TU's container, adding "low order path overhead" to create the TU's virtual container (VC-11, VC-12, VC-2, or VC-3, depending on TU type), linking this VC to the TU frame by



means of a TU pointer, which is the only element of TU Section overhead. The TU frame then is multiplexed into a fixed location within the VC-4.

Although the TU frame structure is illustrated here residing in one VC-4, it actually is distributed over four consecutive VC-4 frames. It is therefore more accurate to refer to the structure as a *TU multiframe*. The phase of the multiframe is indicated by one of the nine VC-4 path overhead (H4) bytes.

**13.11.3 Concatenation**

Payloads with data rates greater than that provided by the SPE/VC are dealt with by a technique called *concatenation*. These signals are denoted by the suffix “c” after the usual STS/STM-*n* signal notation. The data rates that are concatenated are the rates above the base rate. In SONET the rates are STS-3c, STS-12c, STS-48c, etc; in SDH they are STM-4c, STM-16c, etc. Examples of the lowest level of concatenated signals for SDH and SONET follow.

**SONET STS-3c signals.** A higher-rate STS-3 transport signal is normally assembled by byte-interleave multiplexing three STS-1 transport signals that contain tributary signals at the DS3 signal rate (44.74 Mbps) or less. In the SONET context, concatenation means that the higher rate STS-3 transport signal in effect provides one single SPE with a larger payload capacity (Figure 13.19). A higher-rate (greater than 50 Mbps) tributary signal is mapped directly into the larger payload capacity of the STS-3c transport signal (where “c” denotes the concatenation). The STS-3c SPE is assembled without ever going through the STS-1 signal level.




Figure 13.19 SONET STS-3c concatenated SPE.




Figure 13.20 STM-4c concatenated VC.

Once assembled, a concatenated SPE is multiplexed, switched, and transported through the network as a single entity.

**SONET STS-3 concatenated SPE.** The STS-3c signal frame has the same overall dimensions, 9 rows by 270 columns, the same frame repetition rate, 8000 frames per second, and therefore the same signal rate, 155.52 Mbps, as the standard STS-3 signal. Also in common with the standard STS-3 signal, the first 9 columns of the STS-3c frame, a total of 81 bytes, are allocated to Transport overhead.

The STS-3c payload capacity comprises 260 columns of 9 bytes, for a total of 2340 bytes. These bytes provide a transport capacity of 149.76 Mbps at a frame repetition rate of 8000 Hz. Signal capacity for Path overhead is allocated in the first column of the STS-3c SPE, a total of 9 bytes per frame.

**SDH STM-4c signal.** An STM-4 transport signal (Figure 13.20) is normally assembled by byte-interleave multiplexing four STM-1 transport signals. This multiplexing process results in the VC area being occupied by four individual VC-4s. Each VC-4 consisting of Path overhead and a container capable of carrying mapped tributary signals at rates up to 149.76 Mbps.

In the case of a concatenated STM-4 (denoted STM-4c), the virtual container area is entirely filled by a single VC-4-4c. This VC-4-4c consists one Path overhead and a single container capable of carrying a tributary signal operating at rates up to approximately 600 Mbps. Once assembled, a VC-4-4c (or any other concatenated VC structure) is multiplexed, switched, and transported through the network as a single entity.

**STM-4c frame structure.** The STM-4c signal frame has the same overall dimensions as an STM-4 (9 rows by 1080 columns), the same frame repetition rate (8000 frames per second), and therefore the same signal rate (622.08 Mbps). The SOH area of an STM-4c is identical in structure as that of the STM-4 frame; the first 36 columns are allocated to Section overhead.

The STM-4c's container comprises 1043 columns of 9 bytes each, for a total of 9387 bytes. These bytes provide a transport capacity of 600.77 Mbps at a frame repetition rate of 8000 Hz. Signal capacity for Path overhead is allocated in the first column of the VC-4-4c (i.e., a total of 9 bytes per frame).

### 13.12 Payload Pointers

As explained previously, the embedded overhead of a synchronous signal contains a number of bytes designated as *payload pointers*. These pointers (Figure 31.21) are crucial to the synchronous system's efficient operation, and perform the following functions:

- Facilitate asynchronous operation
- Aid efficient mapping
- Minimize network delay in the network

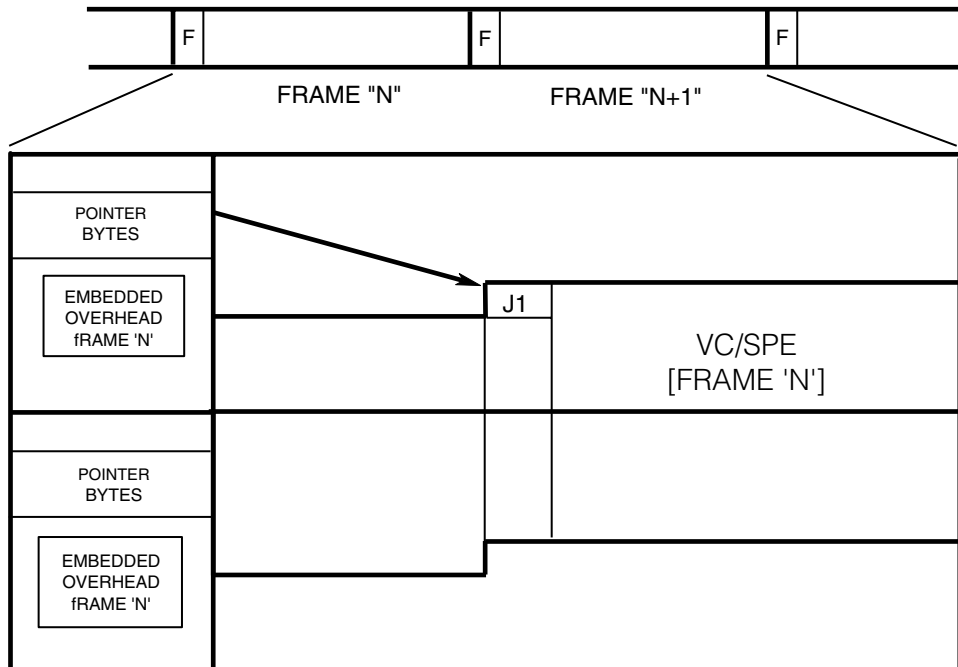


Figure 13.21 Pointers in action.

In a synchronous network ideally all network nodes should derive their timing signals from a single master network clock. In practice this will not always be the case. Timing differences can be caused by a node losing the network timing reference and operating on its standby clock, or by differences at the boundary between two separate synchronous networks. With this in mind the synchronous standards are designed to handle such asynchronous operation within the network.

To accommodate timing differences (clock offsets), the VC-4/SPE can be moved (justified) positively or negatively,  $n$  bytes at time, with respect to the transport frame. (The value of  $n$  is 1 in SONET and 3 in SDH.) This is achieved simply by recalculating or updating the pointer at each synchronous network node. In addition to clock offsets, updating the pointer also will accommodate any other timing phase adjustments required between the input signals and the timing reference of the node.

To facilitate efficient multiplexing and crossconnection of signals in the synchronous network, the VC-4/SPE is allowed to float within the payload capacity provided by the STM-1/STS-1 frames. This means that the payload envelope can begin anywhere in the synchronous payload capacity and is unlikely to be wholly contained in one frame. More likely than not, the VC-4/SPE will begin in one frame and end in the next.

When a VC-4/SPE is assembled into the transport frame, additional bytes are made available in the embedded overhead. These bytes, referred to as the *pointer*, contain a value that indicates the location of the first byte (J1) of the VC-4/SPE. The payload is allowed to float freely within the space made available for it in the transport frame, so that timing phase adjustments can be made as required between the payload and the transport frame.

Another approach to overcoming network timing issues is to use 125  $\mu$ s slip buffers at the inputs to synchronous multiplexing equipment. This type of buffer corrects frequency differences by deleting or repeating a payload frame of information as required. These slip buffers are undesirable because of the signal delay they impose and the signal impairment that slipping causes. Using pointers avoids these unwanted network characteristics.

Pointer processing does, however, introduce a new signal impairment known as *pointer adjustment jitter*. This jitter impairment appears on a received tributary signal after recovery from a payload envelope that has been subjected to pointer changes. Excessive jitter on a tributary signal will influence the operation of the network equipment processing the tributary signal immediately downstream. Great care therefore is required in the design of timing distribution for the synchronous network. This is done to minimize the number of pointer adjustments and, therefore, the level of tributary jitter that results from synchronous transport. Consequences and effect of jitter on the network is dealt with in Chapter 24.

### 13.13 Additional Information

This chapter represents a very brief overview of the synchronous telecommunication standards. The next two subsections contain a list of some of the more important standards documents for both SDH and SONET, as well as a more detailed examination of the alarms mentioned in section 13.10.

### 13.13.1 Synchronous standards documents

Many standards documents exist for both SONET and SDH. These are the two most useful documents for starting to understand the data formats. Each of the documents cross-references to other standards for further study.

- Synchronous Optical Network (SONET) Transport Systems Common Generic Requirements GR-253-CORE
- ITU-T Recommendation G.707 Network node interface for synchronous digital hierarchy (SDH)
- European Telecommunications Standards Institute (ETSI) requirement ETS 300-417-1-1. Generic functional requirements for Synchronous Digital Hierarchy (SDH) equipment.

### 13.13.2 Alarm definitions

***SDH RS and MS Alarms.*** The Regenerator Section and Multiplexer Section alarms in SDH are as follows:

- Loss of Signal (LOS)
  - LOS state entered when received signal level drops below the value at which an error ratio of 1 in  $10^3$  is predicted.
  - LOS state exited when two consecutive valid framing patterns are received; during this time no new LOS condition is detected.
- Out of Frame (OOF)
  - OOF state entered when four (or five in some implementations) consecutive SDH frames are received with invalid (errored) framing patterns. Maximum OOF detection time is therefore 625 ms.
  - OOF state exited when 2 consecutive SDH frames are received with valid framing patterns.
- Loss of Frame (LOF)
  - LOF state entered when OOF state exists for 3 ms. If OOFs are intermittent, the timer is not reset to zero until an in-frame state persists continuously for 3 ms.
  - LOF state exited when an in-frame state exists continuously for 3 ms.
- Loss of Pointer (LOP)
  - LOP state entered when  $n$  consecutive invalid pointers are received or  $n$  consecutive NDFs are received (other than in a concatenation indicator), where  $n$  is 8, 9, or 10.
  - LOP state exited when three equal valid pointers or three consecutive AIS indications are received. (AIS indication is an all-ones pattern in pointer bytes. Concatenation indicator is pointer bytes set to

1001 xx 1111111111

that is, NDF enabled (H1H2 bytes for AU LOP).

- Multiplexer Section AIS (MS-AIS)
  - Sent by Regenerator Section Terminating Equipment (RSTE) to alert downstream MSTE of detected LOS or LOF state. Indicated by STM-N signal containing valid RSOH and a scrambled all-ones pattern in the rest of frame.
  - Detected by MSTE when bits 6 to 8 of received K2 byte are set to 111 for three consecutive frames. Removal is detected by MSTE when three consecutive frames are received with a pattern other than 111 in bits 6 to 8 of K2.
- Far End Receive Failure (FERF or MS-FERF)
  - Sent upstream by Multiplexer Section Terminating Equipment (MSTE) within 250 ms of detecting LOS, LOF, or MS-AIS on incoming signal. Optionally transmitted on detection of excessive BER defect (equivalent BER, based on B2 BIPs, exceeds threshold of  $10^{-3}$ ).
  - Indicated by setting bits 6 to 8 of transmitted K2 byte to 110.
  - Detected by MSTE when bits 6 to 8 of received K2 byte are set to 110 for three consecutive frames. Removal is detected by MSTE when three consecutive frames are received with a pattern other than 110 in bits 6 to 8 of K2.
  - Transmission of MS-AIS overrides MS-FERF.
- High-Order Path AIS
  - Sent by MSTE to alert downstream High Order Path Terminating Equipment (HO PTE) of detected LOP state or received AU Path AIS. Indicated by transmitting all-ones pattern in entire AU-3/4 (i.e., all-ones pattern in H1, H2, H3 pointer bytes, plus all bytes of associated VC-3/4).
  - Detected by HO PTE when all-ones pattern received in bytes H1 and H2 for three consecutive frames. Removal is detected when three consecutive valid AU pointers are received with normal NDFs (0110), or a single valid AU pointer is received with the NDF enabled (1001).
- High-Order Path Remote Alarm Indication (HO Path RAI, also known as HO Path FERF)
  - Generated by High-Order Path Terminating Equipment (HO PTE) in response to received AU Path AIS. Sent upstream to peer HO PTE.
  - Indicated by setting bit 5 of POH G1 byte to 1.
  - Detected by peer HO PTE when bit 5 of received G1 byte is set to 1 for 10 consecutive frames. Removal detected when peer HO PTE receives 10 consecutive frames with bit 5 of G1 byte set to 0.

***SONET Section and Line Span alarms.*** The Section and Line Span alarms in SONET, analogous to the RS and MS alarms in SDH, are as follows:

- Loss of Signal (LOS)
  - STS- $n$  with all-zeros pattern lasting 10–100 ms or longer. NE must enter LOS state within 100 ms of onset of all-zeros pattern.
  - LOS exited when two consecutive valid framing patterns received and during this time no new LOS condition is detected.

- Loss of Frame (LOF)
  - STS- $n$  which is Out of Frame (OOF) for 3 ms or longer. (**Note:** OOF entered when four consecutive frames are received with invalid/errored framing patterns. OOF exited when two consecutive frames are received with valid framing patterns.
  - LOF exited when STS- $n$  remains continuously in-frame for 3 ms or longer (objective 1 ms).
- Loss of Pointer (LOP)
  - STS LOP state is entered when no valid pointer is received in eight consecutive frames, or when eight consecutive NDFs are received (other than in a concatenation indicator).
  - Incoming STS path AIS shall not cause LOP.
  - LOP exited when a valid pointer with normal NDF, or a concatenation indicator, is received in three consecutive frames (STS) or three consecutive superframes (VT).
- Line AIS
  - Sent by STE to alert downstream LTE of received LOS or LOF state. Generation of Line AIS must be done within 125 ms of trigger event. Line AIS is indicated by a STS- $n$  signal consisting of valid section OH and a scrambled all-ones pattern in rest of frame.
  - LTE detects Line AIS when bits 6, 7, and 8 of K2 byte are set to 111 for five consecutive frames.
  - STE shall deactivate Line AIS within 125 ms of exiting failure state.
  - Removal of Line AIS is detected by downstream LTE when five consecutive frames are received with a pattern other than 111 in bits 6 through 8 of K2 byte.
- STS Path AIS
  - Sent by LTE to alert downstream STS PTE that a failure has been detected upstream.
  - STS Path AIS indicated by all-ones pattern being sent in H1, H2, and H3 bytes plus entire STS SPE within 125 ms of failure being detected.
  - STS PTE detects STS Path AIS when all-ones pattern is received in bytes H1 and H2 for three consecutive frames.
  - STS Path AIS is deactivated within 125 ms of LTE exiting failure or Line AIS state. On removal, a valid pointer is created with NDF set, followed by normal pointer operations.
  - STS PTE detects removal of Path AIS when a valid pointer is received with NDF set or when a valid pointer is received for three consecutive frames.





---

Part

**4**

# Local Area Networks



# Private Network Technologies

**Michael A. Pozzi**

*Hewlett-Packard Co., Colorado Springs, Colorado*

## 14.1 Introduction

This chapter provides a brief description of each of several common LAN (local area network) and WAN (wide area network) technologies, comparing and contrasting inherent performance characteristics, and presenting the advantages and disadvantages of each from a performance point of view. The following network technologies are covered in this section:

- LAN technologies, including
  - Ethernet/IEEE 802.3
  - Fast Ethernet/100Base-T
  - Token-Ring/IEEE 802.5
  - FDDI
  - Switched networks
- WAN internetworking technologies, including
  - Private leased lines
  - Packet-switched public networks (X.25 and frame relay)

## 14.2 Local Area Network Technologies

Each of the three most common LAN technologies has its own advantages, and all are widely deployed on a worldwide basis. Table 14.1 presents a brief summary of Ethernet, Token-Ring, and FDDI.

**TABLE 14.1 Each of the three most common LAN topologies has its own advantages, and all are widely deployed on a worldwide basis.**

	Speed	Cost	Advantage
Ethernet	10Mbps	low	inexpensive reliable ubiquitous
Token Ring	4 or 16Mbps	medium	deterministic access good throughput
FDDI	100Mbps	high	high speed fault tolerant

### 5.2.1 Ethernet and IEEE 802.3

Ethernet is the most widely used local area networking topology in the world. Its popularity is due to the fact that Ethernet delivers fast, reliable connectivity that is inexpensive and easy to install and maintain. Ethernet commonly is used to connect individual desktops to the site LAN.

The original Ethernet specification was written in the early 1980s by a consortium composed of Digital Equipment Corporation, Intel Corporation, and Xerox. It specified a 10 Mbps data rate on a coaxial cable bus topology, and a contention resolution process called *Carrier-Sense Multiple Access with Collision Detection*, which is often abbreviated CSMA/CD. The CSMA/CD process allows any station to transmit on the network if no other station already is transmitting. In the event that two or more stations begin transmitting simultaneously (the likelihood of which depends on network load), there will be a *collision*. Both transmitting stations will detect that a collision has occurred, cease their transmission, and wait some specified amount of time before attempting to transmit again.

A few years later the Ethernet specification was adopted by the IEEE (Institute of Electrical and Electronic Engineers) and rewritten as the IEEE 802.3 standard. The original Ethernet frame format was modified by the IEEE in the 802.3 specification, with the result that 802.3 is very similar but not the same as Ethernet.

Ethernet and IEEE 802.3 are very similar local area network protocols providing connectivity at 10 Mbps media speed. Initially designed for use on coaxial cable, Ethernet/802.3 also is used on UTP (unshielded twisted-pair), STP (shielded twisted-pair), and optical fiber. The frame formats for Ethernet and IEEE 802.3 are virtually identical (Figure 14.1), which allows both to coexist on the same network. Subtle differences prohibit interoperability, however.

Though often wired in a physical star configuration, Ethernet/IEEE 802.3 is a logical bus, and all devices share the same transmission media (Figure 14.2). Only one device can transmit at a time. The media access method, as mentioned previously, is CSMA/CD: Each device must wait for the media to become quiet before transmitting, and must listen for other devices that might transmit at the same time. In the event of a collision, both devices will back off a certain period of time and try again.


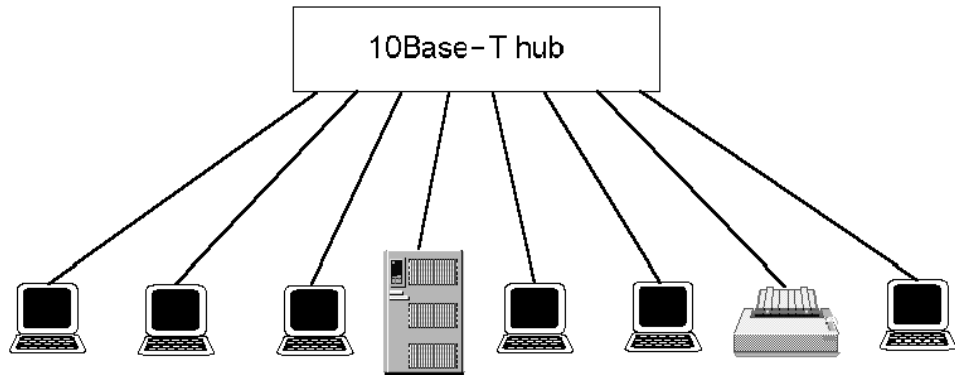
Modern Ethernet networks are wired in a physical star topology using the 10Base-T standard, as shown in Figure 14.3. The term *10Base-T* refers to 10

**Ethernet Frame Format**

Preamble	Destination MAC Address	Source MAC Address	Ethertype	Data	Frame Check Sequence
8 bytes (including Start of Frame Delimiter)	6 bytes	6 bytes	2 bytes	minimum 46 bytes maximum 1500 bytes	4 bytes

**IEEE 802.3 Frame Format**

Preamble	Destination MAC Address	Source MAC Address	Length	Data	Frame Check Sequence
8 bytes (including Start of Frame Delimiter)	6 bytes	6 bytes	2 bytes	minimum 46 bytes maximum 1500 bytes	4 bytes

**Figure 14.1** Comparison of Ethernet and IEEE 802.3 frame structures.**Figure 14.2** Ethernet and IEEE 802.3 networks connect devices in a logical bus topology, where each has access to the same transmission media.**Figure 14.3** The 10Base-T star-wired topology uses twisted-pair cabling and a multiple-port repeater called a *hub* to implement Ethernet and IEEE 802.3 networks.


Mbps using baseband signaling over twisted-pair cabling. A 10Base-T *hub* is a repeater that connects all networked devices together using twisted-pair cabling. The star-wired topology of 10Base-T improves the network's tolerance to physical faults, and is easier and less expensive to install and maintain. A 10Base-T network is still a logical bus, however, in that all devices must share the same 10 Mbps of transmission bandwidth.

Ethernet networks perform best under light to moderate traffic loads, generally under 30 to 40 percent of the available media bandwidth. The CSMA/CD media access method is inherently random and non-deterministic, meaning that the media access time increases exponentially under heavy traffic loads (Figure 14.4).

The 1518-byte maximum data payload per packet specification also limits the maximum data throughput of Ethernet networks, particularly when Ethernet is used to interconnect Token-Ring or FDDI ring segments, which have larger allowable data packet sizes. In these cases, the larger packets must be fragmented or broken up into smaller packets so they can be transmitted over the Ethernet/802.3 segment. The packet fragmentation process usually is implemented in the routers that interconnect the two different network types. Reassembling the fragmented packets is the responsibility of the destination network node. Packet fragmentation and reassembly consumes computing resources in the routers and in the end nodes, with a negative impact on network performance.

Key performance parameters for Ethernet networks include utilization percentage, frame rate (frames per second), collision rate, packet deferral rate, error rate, and average frame size. These are explained briefly in subsequent paragraphs, and enumerated in more detail in Table 14.2.

**Collision rate.** Collisions are a regular Ethernet occurrence and happen when two or more nodes try to send on the media at the same time. When a collision occurs,



**Figure 14.4** The media access time for Ethernet is random and depends on the traffic volume. At high traffic levels, the amount of time required to access the media increases exponentially with traffic volume.

**TABLE 14.2** There are no standards for acceptable levels of the performance parameters listed below. What is acceptable for one network may not be on another because of the number of attached stations, the amount and type of data traffic, and many other factors. However, these guidelines can be applied for many typical Ethernet/802.3 network implications.

Parameter	Used to Indicate	Guidelines
Utilization %	Network (transmission media) congestion	<40% sustained utilization <70% peak (1 second)
Frame Rate	Device congestion	Device dependent (typically <5,000 frames/second)
Collision Rate	Network (transmission media) congestion	<10% of Frame Rate
Packet Deferral Rate	Network (transmission media) congestion	<10% of Frame Rate
Error Rates		
Runts	Collision fragments; faulty NIC	None, except collision-related
Jabbers (Giants)	Faulty NIC; misconfigured router	None
Bad FCS Frames	Electrical noise; Collision fragments	None, except collision-related
Misaligned Frames	Collision fragments; faulty NIC	None, except collision-related
Broadcast, Multicast Frame Rate	Misconfigured routers, nodes or applications	Network-dependent (generally <20 to 30 per second)
Protocol Distribution by frames	Consumption of interconnect device bandwidth by application	Network and application dependent
by kbytes	Consumption of network (transmission media) bandwidth by application	Network and application dependent
Frame Size Distribution	Efficiency of networked applications	Application dependent (generally, larger frames are more network-efficient than smaller ones)
Top Talkers by frames	Consumption of interconnect device bandwidth by node	Network dependent
by kbytes	Consumption of network (transmission media) bandwidth by node	Network dependent

the sending nodes sense the collision, perform a random timeout count, and then retry. As the network becomes more heavily loaded with frames, more collisions will occur. High collision rates also can be caused by faulty adapters or out-of-control nodes generating frames that saturate the network.

**Packet deferral rate.** A *packet deferral* occurs when any node attempts to transmit a frame and senses that another node is already transmitting (i.e., carrier is sensed on the media). The node must *defer*; or wait until the network is idle before it can proceed with the transmission. The packet deferral rate is a statistic often

kept by nodes for each of their Ethernet ports. It is not possible to measure the packet deferral rate with a protocol analyzer or other test tool.

**Runts.** *Runts* are frames that are shorter than 64 bytes, and therefore are invalid on an Ethernet network. They can be the result of collisions on the network, or can be a sign that a node is generating short frames without padding them up to 64 bytes. Because runts often are too short to include a source address, they are very difficult to associate with a particular node.

**Jabbers.** *Jabbers* are frames that are longer than 1518 bytes, and therefore are invalid on an Ethernet network. Jabber frames are also often called “Giants”. Jabbers usually are the result of a node generating frames outside Ethernet specs, or a faulty transceiver on the network.

**Bad FCS frames.** These are frames containing a *Frame Check Sequence* that does not match the checksum calculated when the frame is received. Frames with a bad FCS contain one or more bit errors and are discarded by the receiving node. Bit errors can be caused by electrical noise or faulty terminations, or by a faulty transceiver or cable system component. Collisions also can cause frames to have a bad FCS.

**Misaligned frames.** *Misaligns* are frames whose length in bits is not divisible by eight (in other words, a noninteger number of bytes). These frames usually also have bad FCSs and are usually the result of electrical problems on the network cabling, a faulty workstation, or a collision.

### 14.2.2 Fast Ethernet

Fast Ethernet (100Base-T) is fundamentally the same technology as Ethernet, using CSMA/CD and operating at 100 Mbps over fiber or high-grade unshielded twisted-pair. Fast Ethernet connections often are used to connect high-usage nodes, such as servers or routers, to an Ethernet switch in order to relieve traffic congestion that otherwise might tend to occur there.

The data throughput of 100 Mbps Fast Ethernet is ten times as great as 10 Mbps Ethernet. The performance parameters for Fast Ethernet are the same as those of Ethernet with one exception: When either Ethernet or Fast Ethernet is used to connect only two nodes in a switched environment, either type can be run in a full-duplex mode with no device contention or collisions.

### 14.2.3 Switched networks

*Ethernet switches* are devices used to divide an Ethernet network into smaller *collision domains*. The switch device has multiple ports that can be used to connect multiple end nodes or multiple Ethernet segments to form a switched Ethernet network. The Ethernet switch is a selective repeater; in effect, it acts like a very fast multiple-port bridge, switching packets between ports based on destination Ethernet/802.3 address.

While a 10Base-T hub repeats each received packet to all of the attached ports, an Ethernet switch will forward each received packet only to the intended destination port. The result is that the Ethernet switch provides dedicated 10 Mbps bandwidth



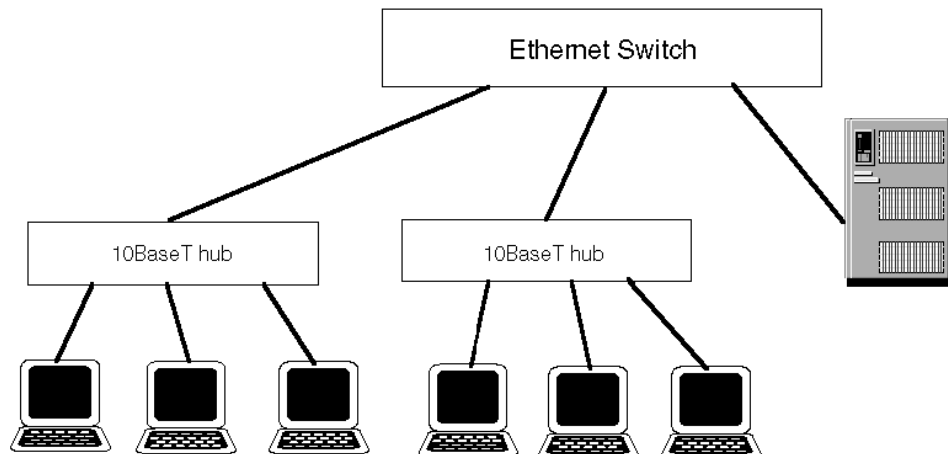
on each of its ports. At each port there exist only two stations contending for transmission bandwidth, the switch and the end node. Only in cases where multiple end nodes share the same switch port (through the use of a standard 10Base-T hub) will there be contention for the same transmission bandwidth. Transactions can occur independently on all the other switch ports.

Ethernet switches are most often configured with one or two Fast Ethernet ports for connections to servers, routers, or other heavily used devices. Without these 100 Mbps ports, such devices would experience frequent congestion on their connection ports to the switch. In fact, without these high-speed ports there often is little benefit derived from the switch alone because of that congestion. Ethernet switches configured with Fast Ethernet ports to servers and routers can greatly improve the performance of a congested 10Base-T network (Figure 14.5).

Single stations with dedicated switch ports can be configured to operate in a full-duplex mode, effectively doubling the bandwidth available at that connection. Full-duplex operation is available at both 10 Mbps and 100 Mbps data rates, and often requires an upgrade to the station's network interface card.

The *switching latency* of the Ethernet switch can become a limiting factor for network performance in some cases. Ethernet switches are of two basic varieties: *Cut-through switches* minimize switching latency by passing frames through before they are completely received. *Store-and-forward switches* do not forward frames until they have been completely received and the frame check sequence has been verified. While the switching latency is greater, the store-and-forward technology better isolates errored frames and collision fragments to individual collision domains.

Monitoring data traffic with a protocol analyzer on a switched Ethernet network can be a challenge because (unlike standard Ethernet or 10Base-T) there is no single physical location where all traffic flows. Some switches can be configured to route selected traffic to a monitor port for analysis, but this requires dedicating a



**Figure 14.5** An Ethernet switch separates the Ethernet/802.3 network into multiple collision domains by selectively repeating received packets only to the intended destination switch port, based on the destination Ethernet/802.3 address.

spare switch port. The additional frame forwarding process also can have a negative impact on the performance (latency) of the switch. Another solution is to connect an analyzer in series between the switch and one of the file servers, where most traffic flows.

#### 14.2.4 Token-Ring (IEEE 802.5)

Token-Ring is another very popular local area networking technology for desktop connectivity. Its main advantages include fault tolerance and a deterministic access method. Larger maximum frame size (relative to Ethernet) also allows for greater data throughput rates. Its main disadvantage relative to Ethernet is that it typically is more expensive to install. As 802.3 differs slightly from original Ethernet, so too does the 802.5 specification differ slightly from its most prevalent implementation in IBM's Token-Ring product.

A *token ring* consists of a number of devices serially connected in a ring topology (Figure 14.6). Token-Ring networks are designed to operate at either 4 or 16 Mbps, and typically use unshielded or shielded twisted-pair copper cables (although the IEEE 802.5 specification differs in not mandating a particular physical medium or topology). The frame format for IEEE 802.5 token ring frames is shown in Figure 14.7.

The IEEE 802.5 token ring specification defines no maximum length for the data field. However, the time required to transmit a frame may not be greater than the token holding period defined for the transmitting device. In practice, the maximum frame size on a 4 Mbps ring generally is 4096 bytes; on a 16 Mbps ring it generally is 18 Kbytes.


The media access method used in IEEE 802.5 is *token passing*. A *token* is the “symbol of authority,” which is passed from station to station to indicate which device currently is allowed to transmit. After a device completes its transmission, it passes the token to the next device on the ring so that each station takes its turn.

One station is designated the Active Monitor, and must constantly monitor the ring to insure that the token continues to circulate on a regular basis. The Active Monitor can be any Ring Station that wins the *Monitor Contention* process. It is primarily responsible for keeping the token alive on the network. In networks with managed hubs or routers, the Active Monitor usually is one of those devices; they often are the first stations on the ring to come up, or have been on the longest.

Most Token-Ring networks are physically wired in a star topology using a media access unit (MAU) to construct the logical ring internally (Figure 14.8). This increases the tolerance of the network to faults in the physical wiring and in the individual station transceivers.

The token-passing scheme of Token-Ring is a very ordered, deterministic process—most of the time. This process is disrupted, however, each time a device enters or exits the ring (including station power on/off). These disruptions, and the ensuing ring restoration process, can limit network performance.

Token-Ring networks achieve fault tolerance by defining *fault domains* on the ring. If a station fails to receive a signal from its *Nearest Active Upstream Neighbor* (NAUN), that station will begin transmitting a distress signal, or *beacon*. The fault




**Figure 14.6** Token-Ring networks are based on a ring topology. A special frame called a *token* is passed from station to station around the ring, and each station can transmit data only when it possesses the token.

**Token Ring Frame Format**

Starting Delimiter	Access Control	Frame Control	Destination MAC Address	Source MAC Address	Data	Frame Check Sequence	Ending Delimiter	Frame Status
1 byte	1 byte	1 byte	6 bytes	6 bytes	0+ bytes	4 bytes	1 byte	1 byte

**Figure 14.7** The IEEE 802.5 Token-Ring frame structure.

domain in this case will include everything between the transmitter of the NAUN and the receiver of the beaconing station. In most cases, either the MAU or the individual nodes involved are able to restore ring integrity by isolating and removing the faulty hardware from the ring.



**Figure 14.8** Token-Ring networks typically are wired in a physical star topology, while remaining a logical ring. The star-wired physical topology increases the network's tolerance to physical faults.

The NAUN is the closest station on the ring to the reporting station. The NAUN will change as stations insert and remove themselves from the ring. Most Token-Ring MAU ports are numbered in sequence. If there are stations active on port 6 and port 8 only, then the station in port 8 will have its NAUN as the station in port 6. Should a station in port 7 become active, then it would become the new NAUN for the station in port 8. NAUNs are discovered by each Ring Station during the *Neighbor Notification* process.

Each Token-Ring network elects one of its attached stations to serve as the Active Monitor through a process called *claiming*. The station with the highest numerical Token-Ring address present on the ring at the time of claiming is chosen to be the Active Monitor. The purpose of the Active Monitor is to ensure orderly and efficient data interchange on the ring by ensuring regular circulation of the token, controlling transmission priority, maintaining appropriate ring delay, and several other functions.

Token-Ring networks often make use of a bridging method called *source routing*. Source routing, more correctly termed *source route bridging*, makes extensive use of broadcast messages in order to establish routes between nodes across multiple-ring networks. This broadcast traffic can become excessive in large networks, and can have negative impact on network performance.

Token-Ring networks carrying SNA data traffic also can be very susceptible to time delays. Sessions often can be dropped if the frame transmission latency exceeds the timeout value, which can occur regularly on congested networks or on networks that are interconnected by WAN links over long distances. *Data link switching* (DLSw) and other methods are often used to counter this potential performance problem.

Key performance parameters for Token-Ring networks include utilization percentage, frame rate, average frame size, and hard and soft error rates. These performance parameters are summarized in Table 14.3; the major types of hard and soft errors are described in the following subsections.

**TABLE 14.3** There are no standards for acceptable levels of the performance parameters listed below. What is acceptable for one network may not be on another because of the number of attached stations, the amount and type of data traffic, and many other factors. However, these guidelines can be applied for many typical token ring network implementations.

Parameter	Used to Indicate	Guidelines
Utilization %	Network (transmission media) congestion	<70% sustained utilization <90% peak (1 second)
Frame Rate	Device congestion	Device dependent (typically <5,000 frames/second)
Hard Error Rates		
Ring Purge Frames	Ring resets; station insertion or removal	Minimize
Ring Beacon Frames	Hard failure of NIC, MAU or cabling	Minimize transient beacons No streaming beacons
Claim Token Frames	Ring resets; station insertion or removal	Minimize
Soft Error Rates		
Isolating Soft Errors internal errors, burst errors, line errors, abort errors, address recognized/copied errors	Marginal timing, electrical noise Station addresses identify a fault domain	Minimize
Non-Isolating Soft Errors frequency errors, frame copy errors, token errors, receiver congestion errors	Marginal timing, electrical noise Cannot be isolated to a fault domain	Minimize
Broadcast, Multicast Frame Rate	Misconfigured routers, nodes or applications	Network-dependent (generally <20 to 30 per second)
Protocol Distribution by frames	Consumption of interconnect device bandwidth by application	Network and application dependent
by kbytes	Consumption of network (transmission media) bandwidth by application	Network and application dependent
Frame Size Distribution	Efficiency of networked applications	Application dependent (generally, larger frames are more network-efficient than smaller ones)
Top Talkers by frames	Consumption of interconnect device bandwidth by node	Network dependent
by kbytes	Consumption of network (transmission media) bandwidth by node	Network dependent
Source Routing Distribution	Distribution of traffic by source and destination network (local, remote)	Network dependent

**Token-Ring hard errors.** Token-Ring networks are subject to the following types of errors, which are considered “hard” errors:

- Ring purge frames
- Ring beacon frames
- Claim token frames

**Ring purge frames.** *Ring purge frames* are frames that are generated by the Active Monitor when a claim token process is completed, or if a token error occurs (often a lost frame). The purpose of a ring purge frame is to reinitialize the ring by removing any data frames or tokens that might be circulating. Ring purges are normal when stations insert into the ring, but should not be common otherwise. Counts higher than just a few usually indicate cabling problems.

**Ring beacon frames.** *Ring beacon frames* are issued when a serious fault occurs, such as a break in the physical cable. Ring beacon frames are sent by a Ring Station (present on all Token-Ring adapters) reporting the address of its nearest active upstream neighbor and indicating that it is no longer receiving the token that the NAUN should send. Ring beacon frames usually indicate cabling or adapter faults between the station generating the ring beacon frame and its NAUN.

**Claim token frames.** *Claim token frames* are frames that are sent by any station on the ring that detects the absence of the Active Monitor, or that is trying to contend for the position of a new Active Monitor. The purpose of the claim token frame is to initiate the claiming process.

**Token-Ring soft errors.** Token-Ring networks also are subject to errors that are considered “soft” errors. This term refers to a class of abnormal events that affect only one station and generally do not impact the operation of the entire ring. Soft errors include the following:

- Isolating soft errors
- Non-isolating soft errors
- Internal errors
- Burst errors
- Line errors
- Abort errors
- A/C errors
- Frequency errors
- Frame copy errors
- Token errors
- Receiver congestion errors
- Lost frame errors

**Isolating and non-isolating soft errors.** *Isolating soft errors* are those that can be traced to specific neighbor stations. *Non-isolating soft errors* are those that cannot be traced to specific stations.

**Internal errors.** *Internal errors* are abnormal events detected by a station whenever it recognizes a recoverable internal error.

**Burst errors.** *Burst errors* are reported by a station whenever it detects the absence of transitions in the received signal (loss of signal) for more than two and one-half bit times between starting and ending delimiters (i.e., within the received frame).

**Line errors.** *Line errors* are frames generated by a station to report a code violation either in a token, a frame check error sequence, or between starting and ending frame delimiters.

**Abort errors.** A frame in which the ending delimiter immediately follows the starting delimiter, with none of the required fields between, is called an *aborted frame*. An *abort error* is registered each time an aborted frame is observed on the ring.

**A/C errors.** An *Address Copied error* (A/C error) is an abnormal event detected in the Address Recognized and Frame Copied procedure on the ring. An A/C error may indicate the presence of more than one Active Monitor on the ring.

**Frequency errors.** These errors occur when the ring clock and a physical ring station's crystal clock differ significantly in frequency.

**Frame copy errors.** These error frames indicate that a station has recognized a frame addressed to it, but the frame does not have the Address Recognized Indicator bit set to 00 as it should. This condition can indicate a transmission problem from the sending station, or can be the result of stations with duplicate addresses.

**Token errors.** *Token error* frames are generated by the Active Monitor to indicate that one of several token protocol errors occurred. If there are more than just a few such frames, it often indicates that a token or frame was not received within the 10 ms limit. If the token had a nonzero priority and a Monitor Count of one, it means the frame passed by the Active Monitor twice. In either case, a token error is often an indication of cabling or NIC problems.


**Receiver congestion errors.** *Receiver congestion errors* indicate that a station is not capable of copying a frame directed to it out of its buffer and into memory. This can be a result of the station partially crashing, leaving the Token-Ring interface operational but main memory corrupted, in which case the device should be cold-booted. In other cases, congestion errors can indicate that the receiving station could be too busy to accept frames from its buffer and has become a bottleneck, as can happen with low-performance Token-Ring bridges and routers.

**Lost frame errors.** These are frames indicating that a transmitting station did not receive the end of its last frame.

### 5.2.5 FDDI

FDDI (Fiber Distributed Data Interface, sometimes pronounced “fiddy”) is a high-bandwidth, general-purpose LAN technology commonly used for backbones. A *backbone* is the central network segment that is used to interconnect all parts of the network. The backbone typically must carry high volumes of data traffic, and must be very reliable. FDDI often is chosen as a backbone technology because of its ability to carry high traffic volume very efficiently and with a high degree of fault tolerance.

Like Token-Ring, an FDDI ring consists of a number of devices serially connected in a ring topology (Figure 14.9). Unlike Token-Ring, FDDI specifies a dual-ring topology: a primary ring for data transmission, and a secondary ring for fault tolerance (redundancy). Designed to operate at 100 Mbps, FDDI networks typically operate



**Figure 14.9** FDDI (Fiber Distributed Data Interface) uses a dual-ring topology for fault tolerance. The primary ring carries data traffic, while the secondary ring is an active standby in case of a fault in the primary ring.



FDDI Frame Format

Preamble	Starting Delimiter	Frame Control	Destination MAC Address	Source MAC Address	Data	Frame Check Sequence	Ending Delimiter	Frame Status
8 bytes	1 byte	1 byte	6 bytes	6 bytes	0+ bytes*	4 bytes	1 byte	1 byte

Figure 14.10 The FDDI frame structure.

on multimode or single-mode fiber, or on UTP copper cables using the *Twisted-Pair Physical layer Media-Dependent* (TP/PMD) specification, sometimes called “FDDI over copper.”

FDDI networks generally provide excellent performance, high data throughput, and high reliability (including fault tolerance). The biggest limitations of the FDDI technology are its complexity and expense. For these reasons, FDDI rings typically are used only as high-speed backbone connections and not normally for desktop connectivity.

The maximum size of a FDDI frame is 4500 bytes. The frame format for FDDI frames is shown in Figure 14.10.

As in Token-Ring, the media access method used in FDDI is *token passing*. A token is a special packet that grants the station possessing it the permission to transmit on the network. After a device completes its transmission, it passes the token to the next device on the ring, so that each station takes its turn.

FDDI networks are often designed using a combination of a main ring and a number of “trees” branching from *concentrators* connected to the main ring (Figure 14.11). This further improves the fault tolerance of the network to potential problems occurring with stations connected via the concentrators. Stations that insert or remove themselves from the ring frequently (like end-user workstations) often are connected using concentrators in order to minimize the disruptions on the main ring.


Critical nodes (like file servers) connected through concentrators often are provided with two separate paths to the main ring. A secondary path is normally inactive, and is used as a backup in case of a fault on the primary connection. In this way, fault tolerance is achievable on the tree structure. This type of connection is called *dual homing*.

Key performance parameters for FDDI include utilization percentage, frame rate, average frame size, and hard and soft error rates. These parameters are summarized in Table 14.4.

**FDDI hard errors.** Like Token-Ring, FDDI is subject to conditions that are considered “hard” errors. The hard errors include:

- Beacon frames
- Claim frames

**Beacon frames.** *Beacon frames* are sent by FDDI stations when they stop receiving frames or tokens from their upstream neighbors. When a station receives a beacon frame from its upstream neighbor, it stops sending its own beacon, until the only station left is the one reporting its own address and that of its upstream neighbor, between which the problem lies. Counts in this field usually represent cabling or



**Figure 14.11** Large FDDI networks often use tree structures to attach multiple end stations through devices called *concentrators*. Minimizing the number of stations directly attached to the main ring further increases the reliability and fault tolerance of the FDDI network.

adapter faults between the station generating the beacon frame and its upstream neighbor.

**Claim frames.** *Claim frames* are sent by any station that is inserting into the ring, is removing from the ring, has not received a frame or token within a reasonable time, or has detected a high bit-error rate. Claim frames also are sent when a station wishes to send frames more frequently than the already negotiated *Token Rotation Time (TRT)*.

**FDDI soft errors.** The errors considered “soft” errors in FDDI include the following:

- Bad FCS frames
- Violations

- E-flag set
- Short preambles
- Long frames

**TABLE 14.4 There are no standards for acceptable levels of the performance parameters listed below. What is acceptable for one network may not be on another because of the number of attached stations, the amount and type of data traffic, and many other factors. However, these guidelines can be applied for many typical FDDI network implementations.**

Parameter	Used to Indicate:	Guidelines
Utilization %	Network (transmission media) congestion	<80% sustained utilization <90% peak (1 second)
Frame Rate	Device congestion	Device dependent
Hard Error Rates Claim Frames	Ring resets; station insertion or removal	Minimize
Beacon Frames	Ring resets; station insertion or removal	Minimize
Soft Error Rates Long frames, short preambles, E-flag set, violations, bad FCS frames	Marginal timing, electrical noise (on TP/PMD copper rings)	Minimize
Broadcast, Multicast Frame Rate	Misconfigured routers, nodes or applications	Network-dependent (generally <20 to 30 per second)
Protocol Distribution by frames	Consumption of interconnect device bandwidth by application	Network and application dependent
by kbytes	Consumption of network (transmission media) bandwidth by application	Network and application dependent
Frame Size Distribution	Efficiency of networked applications	Application dependent (generally, larger frames are more network-efficient than smaller ones)
Top Talkers by frames	Consumption of interconnect device bandwidth by node	Network dependent
by kbytes	Consumption of network (transmission media) bandwidth by node	Network dependent
Ring State LEM reject count, LEM count, SMT transmit frames, ring op count	Ring resets; transmission errors	Minimize
Token Rotation Time	Media access time	Network dependent

**Bad FCS frames.** These are frames with a Frame Check Sequence that does not match the checksum calculated by the receiver. This is usually an indication of a faulty transceiver or cable system component.

**Violations.** This is an invalid symbol in the MAC sublayer of the Data Link layer. There can be multiple causes, typically noise on the line, a faulty transceiver, or an errored MAC implementation.

**E-Flag set.** The *E-flag* is a single bit at the end of the frame that is set whenever an error is detected in the frame received from the upstream neighbor, meaning that the frame check sequence is incorrect.

**Short preambles and long frames.** *Short preambles* are those that are less than the specified 14 symbols in length. *Long frames* are those that exceed the 4500-byte maximum specification.

**FDDI ring state.** A third class of parameters shown in Table 14.4 have to do with the status of the FDDI ring as a whole. These include:

- Ring op count
- SMT transmit frames
- LEM count
- LEM reject count
- Token Rotation Time (TRT)

**Ring op count.** This is the ring operation counter, which registers the number of times the ring has been initialized per second.

**SMT transmit frames.** This is the number of SMT (station management) frames sent on the network. SMT frames are used to implement the station management process within FDDI stations, which is a process of monitoring and reporting the operation of the FDDI protocols.

**LEM count.** The aggregate *Link Error Monitor* (LEM) count is a tally of abnormal events observed on the FDDI link.

**LEM reject count.** The *LEM reject count* is a tally of the number of times the FDDI link is reset, usually because of a high number of link errors.


**Token Rotation Time (TRT).** The *Token Rotation Time* is the actual time taken for the FDDI token to circulate around the ring. It normally is expressed in nanoseconds.

### 14.3 Wide Area Network Technologies

The three most common WAN technologies are private leased lines, and the X.25 and frame relay public packet-switched networks. Each has its own advantages and disadvantages, summarized in Table 14.5, and all are widely deployed.

**TABLE 14.5 Each of the three most common WAN topologies has its own advantages, and all are widely deployed on a worldwide basis.**

	Speed	Cost	Advantage
Private Leased Lines	56kbps – 45Mbps	high	end user control customized
X.25	up to 64kbps	medium	error-free switched virtual circuits
Frame Relay	up to 2Mbps	low	inexpensive high speed

**Figure 14.12** Private leased lines provide constant bandwidth over time, while the typical requirements of most data traffic include great variations between peak and average bandwidth.

### 14.3.1 Private leased lines

Private leased lines are dedicated transmission facilities that are leased from a provider on a monthly basis. The amount of bandwidth available is fixed, generally in increments of 56 or 64 kbps. The physical locations of the transmission line endpoints are fixed. The framing (ESF, D4, etc.) and line code (AMI, B8ZS, etc.) also are fixed for each leased line.

Private leased lines can carry voice, data, video, or any combination thereof. There are no restrictions on which protocols can be used for data transmission over the leased line, except that they must be compatible with the chosen framing and line speed.

Link-layer protocols are used on wide area networks to frame data for transmission, for addressing, frame sequencing, error detection and recovery, and several other purposes. Private-line wide area transmission links often use *High-level Data Link Control* (HDLC), *Synchronous Data Link Control* (SDLC), *Point-to-Point Protocol* (PPP), or a proprietary, vendor-specific link-level protocol.

Private leased lines deliver a constant amount of bandwidth over time. The bandwidth requirements of most data connections are for a large peak-to-average data throughput ratio, as shown in Figure 14.12. The challenge when using a private leased line is to choose a bandwidth that is high enough to accommodate the

peak data throughput while still making economic sense given the longer-term average data throughput.

Many different factors affect the performance of a private leased network connection. The most obvious is the restricted amount of available bandwidth, relative to the amount of data traffic present on the attached local area networks. Beyond that, leased lines can suffer from bit errors (or transmission errors), which will cause re-transmissions of user data and a resulting delay in network response time.

**Clock speed latency.** Even error-free transmission facilities with plenty of available bandwidth can cause two kinds of delays that impact network performance. The first is *clock speed latency*, or the delay due to the slow (relative to the attached LANs) clock rate used to place a data frame on the line. For example, consider the delay in transmitting a 1000-byte frame across a leased 64 kbps line. The clock speed latency can be calculated as:

$$\frac{1,000 \text{ bytes}}{64 \text{ kbps}} \times \frac{8 \text{ bits}}{1 \text{ byte}} = 0.125 \text{ seconds} \quad (14.1)$$

**Transmission latency.** The second kind of delay is *transmission latency*, which is the time delay for the frame to propagate from the source end of the private leased line to the destination. The transmission latency also can be quite significant, particularly for long connections or for connections using satellite links, and must be added to the clock speed latency in order to calculate the total latency for each frame traversing the private leased line.

It is important that the network interconnect devices (remote bridges and routers) and networked applications using private leased lines be configured to provide optimum performance over these time-delayed and limited-bandwidth connections. End-node receiver buffer size and transmission window size should be set large enough to avoid delays while waiting for acknowledgments.

If they are set at too large a value, the maximum frame size, router buffer size, and buffer flushing timeouts can result in dropped connections for timeout-sensitive protocols like DEC's LAT or IBM's SNA (*Local Area Transport* protocol and *Systems Network Architecture*, respectively). If set too small, the maximum frame size will limit the efficiency of the transmission facility and could result in packet fragmentation and reassembly overhead, which in turn has negative impact on overall network performance.

**Leased-line performance parameters.** Performance parameters for private leased-line WAN interconnections are summarized in Table 14.6. They include:

- Error rates and line status parameters
  - BPVs* or *bipolar violations* occur when two consecutive pulses have the same polarity.
  - A frame slip* indicates a temporary loss of synchronization on a T1 link.
  - A code violation* indicates an error in the transmission of the E1 line code.

**TABLE 14.6** There are no standards for acceptable levels of the performance parameters listed below. What is acceptable for one network may not be on another because of the data rate of the line, the amount and type of data traffic, and many other factors. However, these guidelines can be applied for many typical leased line private WAN implementations.

Parameter	Used to Indicate	Guidelines
Utilization %	Network (transmission media) congestion	<50% average utilization <100% peak utilization
Frame Rate	Device congestion	Device dependent
Error Rates/Line Status signal loss, frame sync loss, yellow alarm, bipolar violation, frame slips, code violations	Transmission errors, clock synchronization problems, hardware faults	Minimize
Quality of Service Good frames, bad frames, abort frames, short frames, % good frames, % errored frames, % information frames, % information bytes	Impact of transmission errors on the user data frames; efficiency of the data link (non-information bearing protocol overhead)	Minimize % errored frames; % information frames and % information bytes is protocol and network dependent
Broadcast, Multicast Frame Rate	Misconfigured routers, nodes or applications	Network-dependent (generally <20 to 30 per second)
Protocol Distribution by frames	Consumption of interconnect device bandwidth by application	Network and application dependent
by kbytes	Consumption of network (transmission media) bandwidth by application	Network and application dependent
Frame Size Distribution	Efficiency of networked applications	Application dependent (generally, larger frames are more network- efficient than smaller ones)
Top Talkers by frames	Consumption of interconnect device bandwidth by node	Network dependent
by kbytes	Consumption of network (transmission media) bandwidth by node	Network dependent
Bit error rate (BERT) bit error rate, block error rate, errored seconds, error-free seconds, severely errored seconds, degraded minutes, available and unavailable time	Transmission errors	Minimize error rates; Maximize available time

## 344 Local Area Networks

- Quality of Service (QoS) parameters
  - The *Info frames* and *non-info frames* statistics represent the ratio of Data Link layer information frames to total frames.
  - The *Info bytes* and *non-info bytes* statistics represent the ratio of Data Link layer information bytes to total bytes.
- Bit Error Rate (BERT) parameters
  - Bit error rate* is the ratio of bit errors to the total bit count within the measurement sample interval.
  - Block error rate* is the ratio of block errors to the total block count within the measurement sample interval.
  - Errored seconds* is the total number of seconds that contained at least one bit error within the measurement sample interval.
  - Error-free seconds* is the total number of seconds within the measurement sample interval that did not contain any bit errors.
  - Severely errored seconds* is the number of one-second intervals where the bit error rate is greater than  $1 \times 10^{-3}$  within the measurement sample interval.
  - Degraded minutes* is the number of one-minute intervals where the bit error rate is greater than  $1 \times 10^{-6}$  within the measurement sample interval.
  - Available time* is the amount of time the circuit was able to transmit data reliably within the measurement sample interval.
  - Unavailable time* is the amount of time the circuit was not able to transmit data reliably within the measurement sample interval.

### 14.3.2 About packet-switched networks

Packet switched networks are public wide area data transmission facilities and services that offer an alternative to private leased lines for establishing wide area networks. Public packet-switched networks can provide data connectivity between multiple locations, handling all of the required addressing, data switching, and error recovery services internally and in a manner that is transparent to the end users.

Packet-switching networks are designed to better accommodate the high peak-to-average throughput requirements of typical data traffic by sharing wide area network resources among many users, and charging users on a per-packet basis. Users of public packet-switching networks in effect pay only for the data they transmit, rather than pay for a fixed amount of bandwidth that is needed only for times of peak transmission, but which most of the time is underused or not used at all.


The two most popular packet switched networks in worldwide use today are X.25 and frame relay.

### 14.3.3 X.25 networks


X.25 is a protocol specification for the user-to-network interface for a packet-switched network. X.25 was developed in the early 1980s, and is widely deployed. Figure 14.13 is a schematic of how users are connected transparently to the X.25 “cloud.”

The X.25 network is connection-oriented, providing both switched and permanent virtual circuits. A *permanent virtual circuit* (PVC) operates much like leased lines





**Figure 14.13** X.25 specifies the protocols used to connect each user to a public X.25 packet-switching network. The internal operation of the packet-switching network (including switching, routing, and error detection and correction) is completely transparent to the user.



### X.25 physical connection

**Figure 14.14** Multiple X.25 logical connections, called *virtual circuits*, are carried by a single physical link. Each virtual circuit is identified with a unique logical channel number (LCN).

in that the transmission endpoints are fixed. On the other hand, a *switched virtual circuit* (SVC) is a temporary connection, and must be re-established using a call setup protocol for each connection desired. Packets belonging to a particular virtual circuit are identified using a *logical channel number* (LCN), and multiple logical connections typically share a single physical connection to the network (Figure 14.14).

The X.25 protocol guarantees error-free delivery of packets using a store-and-forward process between X.25 switches within the network fabric. This store-and-forward technology limits data rates to 64 kbps in most implementations. Other factors impacting network performance include the bandwidth of the access line connecting the customer premise with the X.25 network, the transmission latency, the Data Link-level window size, and the maximum frame size. Table 14.7 summarizes an X.25

**TABLE 14.7 There are no standards for acceptable levels of the X.25 performance parameters listed below. What is acceptable for one network may not be on another because of the data rate of the line, the amount and type of data traffic, and many other factors. However, these guidelines can be applied for many typical X.25 implementations.**

Parameter	Used to Indicate	Guidelines
Utilization %	Network (transmission media) congestion	<50% average utilization <100% peak utilization
Frame Rate	Device congestion	Device dependent
Quality of Service Unsuccessful calls, reset requests, restart requests	Configuration errors; network congestion	Minimize
Efficiency % information frames, % information bytes, % data packets, % non-data packets	Efficiency of the data link (non-information bearing protocol overhead); efficiency of the network layer	Minimize % non-information frames and % non-data packets
Broadcast, Multicast Frame Rate	Misconfigured routers, nodes or applications	Network-dependent (generally <20 to 30 per second)
Protocol Distribution by frames	Consumption of interconnect device bandwidth by application	Network and application dependent
by kbytes	Consumption of network (transmission media) bandwidth by application	Network and application dependent
Frame Size Distribution	Efficiency of networked applications	Application dependent (generally, larger frames are more network-efficient than smaller ones)
Top Talkers by frames	Consumption of interconnect device bandwidth by node	Network dependent
by kbytes	Consumption of network (transmission media) bandwidth by node	Network dependent

network's performance parameters, which fall into two general categories, quality of service (QoS) and network efficiency.

**Unsuccessful calls.** This is a QoS parameter that expresses the number of Call Request packets that cannot be matched to a Call Accept packet.

**Reset requests.** Another QoS parameter, this one tracks the number of Reset Request packets sent. A PVC is always in the data transfer state, so an end station using these circuits has no need to send call setup packets. There still is a set procedure for starting a data transfer on a PVC, however, which might be necessary if a PVC has been down but is now available again. The procedure is called the *reset procedure* and it is started by the station sending a *Reset Request* packet and specifying the

channel that needs resetting with the appropriate logical channel number. The interface (for the specified channel) must be in the data transfer state to be able to accept the Reset Request.

**Restart requests.** This QoS parameter tracks the number of Restart Request packets issued. *Restart packets* are used for clearing all the SVCs and resetting all the PVCs currently held by the end station that issues the Restart Request. The logical channel identifier of a Restart Request packet is always set to 0 due to this packet's indiscriminate action on all the subscribed virtual circuits. The essential point concerning the restart procedure is that the station can at any time issue a Restart Request to initiate the restart procedure on all the currently active logical channels. Thus the restart procedure provides the only means of placing all the virtual circuits of an interface into a known state.

**Data and nondata packets.** These efficiency parameters express the ratio of Network-layer data packets to total packets, providing an indication of how many packets are actually carrying user data, as opposed to supervising link operation, establishing or tearing down virtual circuits, and other overhead functions. Measuring efficiency by counting packets is meaningful for interconnect devices (bridges and routers), which must make packet forwarding decisions on each data packet received, and whose performance often is specified in terms of the number of packets that can be handled per second.


Since these ratios can vary significantly between different network implementations, their meaning is most significant in the context of a network baseline, where the values are observed at regular intervals over time.

**Data and nondata bytes.** The ratio of Network-layer data bytes to total bytes provides an indication of how many of the network transmitted bytes are actually user data bytes, as opposed to ones supervising link operation, establishing or tearing down virtual circuits, and performing other overhead functions. Measuring efficiency by bytes is the best indication of bandwidth consumption of the physical transmission media, which has a capacity specified in terms of bytes per second.

As with the packet parameters, these ratios can vary significantly between different network implementations. Their meaning therefore is most significant in the context of a network baseline, where the values are observed at regular intervals over time.

#### 14.3.4. Frame relay networks

Like X.25 networks, frame relay networks also use connection-oriented, packet-switching technology that is designed to maximize efficiency by sharing wide area transmission facilities. Unlike X.25, frame relay provides no error checking or end-to-end guarantee of error-free transmission. Eliminating the store-and-forward switching technology necessary to provide error-free service allows frame relay networks to operate at higher data rates, up to T1 (1.544 Mbps) or E1 (2.048 Mbps). Figure 14.15 depicts transparent user access to the frame relay "cloud" using *frame relay access devices* (FRADs).



**Figure 14.15** A frame relay fast packet-switching network can be used for efficient high-speed interconnection of multiple locations. In this illustration, “FRS” means *Frame Relay Switch* and “FRAD” means *Frame Relay Access Device*.

Frame relay also provides for both switched and permanent virtual circuits, although PVCs are by far the more common. Each logical channel in a frame relay network is identified by a number called the *Data Link Connection Identifier* (DLCI). As in X.25, multiple logical channels (multiple DLCIs) typically are present on each physical connection.

Each logical channel in a frame relay network is assigned a *Committed Information Rate* (CIR), which specifies the data rate guaranteed by the network to the user for that virtual circuit. Frame relay networks will deliver short bursts of user data beyond the CIR on a “best-effort” basis. In this way, frame relay networks can accommodate the high peak-to-average data throughput requirements of most data traffic.

User data in excess of the CIR that is sent to the frame relay network can be marked *discard eligible*, or DE. DE-marked data frames may be discarded if the network experiences congestion at any point along the path. The network will indicate congestion conditions in either the forward or backward directions using the *Forward Explicit Congestion Notification* (FECN) and *Backward Explicit Congestion Notification* (BECN) bits. Other factors impacting network performance include the bandwidth of the access line connecting the customer premise with the frame relay network, the transmission latency, the Data Link-level window size, and the maximum frame size.

**Traffic and congestion parameters.** Table 14.8 summarizes the performance parameters of frame relay networks, the most significant of which deal with traffic and congestion. These include:

- *CIR Utilization percentage* represents the data throughput of a particular DLCI expressed as a percentage of the CIR for that channel. This statistic can be of great assistance when optimizing or reconfiguring a frame relay network.
- *DE* (Discard Eligibility) is a frame relay mechanism that allows the source of a data stream to prioritize frames, indicating those preferred to be discarded in the event of network congestion. If the DE bit of a frame is set to 1, the frame is a preferred candidate for discard.

**TABLE 14.8** There are no standards for acceptable levels of the Frame Relay performance parameters listed below. What is acceptable for one network may not be on another because of the data rate of the line, the amount and type of data traffic, and many other factors. However, these guidelines can be applied for many typical Frame Relay implementations.

Parameter	Used to Indicate	Guidelines
CIR Utilization %	User data rates relative to the committed information rate (CIR)	<100% CIR utilization
DE (discard eligible)	Data marked "Discard Eligible"	Bursts above the CIR are eligible for discard by the network
FECN, BECN	Congestion in the Frame Relay network	Bursts above CIR with FECN or BECN congestion may indicate packet loss
Frame Rate	Device congestion	Device dependent
Broadcast, Multicast Frame Rate	Misconfigured routers, nodes or applications	Network-dependent (generally <20 to 30 per second)
Protocol Distribution by frames	Consumption of interconnect device bandwidth by applications	Network and application dependent
by kbytes	Consumption of network (transmission media) bandwidth by application	Network and application dependent
Frame Size Distribution	Efficiency of networked applications	Application dependent (generally, larger frames are more network-efficient than smaller ones)
Top Talkers by frames	Consumption of interconnect device bandwidth by node	Network dependent
by kbytes	Consumption of network (transmission media) bandwidth by node	Network dependent

- *FECN* (Forward Explicit Congestion Notification) is a frame relay flow control flag bit used to notify the receiving node that there is incoming network congestion.
- *BECN* (Backward Explicit Congestion Notification) is a frame relay flow control flag bit used to notify the sending node (source node) that there is network congestion on the outbound path. The suggested response is to reduce the frame rate into the network.



# Private Networks Performance Testing

**Michael A. Pozzi**

*Hewlett-Packard Co., Colorado Springs, Colorado*

## 15.1 Introduction

This chapter examines performance testing in private networks, including customer premise local area networks (LANs) and wide area internetwork (WAN) connections. It includes criteria for evaluating network performance, factors that affect network performance, network performance metrics, tools for monitoring network performance, and baselining and benchmarking techniques for characterizing network performance.

## 15.2 Network Performance Criteria

“Network performance” can mean different things to different people. A user on a networked computing system ideally should be completely unaware that the network exists. Network performance from a user’s perspective can be thought of as the degree to which the network is completely invisible, as if the user were directly connected to any resource he or she chooses to access.

Network performance can be defined as reliability, availability, data throughput, error rate, response time, application performance, or in many other different ways. As each network is unique, so too are the criteria that define performance for each individual network. It is important for each network manager to understand what constitutes good performance for the network being managed, so that results can be measured and compared against a goal.

### 15.2.1 Reliability

One very basic measure of network performance is reliability: Can the network be depended upon? Reliability is a perception held by the network’s users. It is based

## 352 Local Area Networks

upon history: past downtime, application performance, and response time. While each of these factors is quantifiable, the interpretation of what is or is not a reliable network depends on the situation. Every network is different. The real value of downtime, application performance, and response time measurements comes with regular network *baselining*, where a history of these measurements can be examined for changes and for trend analysis. The baselining process is examined in more detail later in this chapter.

Networks that are perceived by users as unreliable will not be used, at least not to the extent that they otherwise would be. User productivity will suffer as a result.

### 15.2.2 Availability

Network *availability* is the percentage of time that the network is available for use measured over some fixed time period. Network availability is defined as:

$$\frac{\text{total elapsed time} - \text{total downtime}}{\text{total elapsed time}} \times 100\% \quad (15.1)$$

where *downtime* is defined as the amount of time during which the network was not available for users.

A low availability percentage certainly indicates a network performance problem, in that the network is unavailable to users for some significant percentage of time. Even though a high availability does mean that the network does not go down, it does not necessarily equate to a well-performing network.

### 15.2.3 Data throughput

*Data throughput* is a measure of traffic volume actually being carried by the network, typically expressed in kilobytes per second. Networks that are capable of carrying a higher data throughput, i.e., higher-speed networks, are sometimes thought of as higher-performance networks.

Data throughput also can be defined on a *per transaction* basis. If it takes 1 second to transfer a 500,000-byte file, then the data throughput for that transaction is 500 kbytes/sec. Data throughput per transaction is one measure of network performance that is very representative of actual end-user experience.

### 15.2.4 Error rate

The *error rate*, expressed as an average number of error events per second measured over some time interval, indicates the degree to which errors are impacting network performance. Errors can include data transmission errors (bit errors), protocol or syntax errors, timing errors, or errors resulting from faults in the physical transmission medium.

Error types are different for different network technologies. (Refer to Chapter 14 for definitions of common errors for Ethernet, Token-Ring, FDDI, private leased WAN lines, X.25, and frame relay networks.)



### 15.2.5 Network response time

Network *response time* is the amount of time that passes between when a request is issued and when a response to that request is received. Response time often is used to characterize the performance of a network.

$$\text{Response Time} = T_{\text{response}} - T_{\text{request}} \quad (15.2)$$

In order to measure response time on a TCP/IP-based network, a *ping* request (ICMP echo request packet) is issued to a target network (IP) address and a response time is measured. Similar request/response message pairs can be used in networks running other protocols.

Network response time is a direct measure of the performance of the network itself, including the time it takes for the network to deliver a single packet to the target station and receive an acknowledgment. Response time alone cannot be used to completely characterize network performance, however, as it does not include the effects of interactions between the network and networked applications. For example, the priority given to ICMP echo request and response packets by network forwarding devices (bridges, routers, or switches) might not be the same as that given to actual user traffic generated by the applications.

### 15.2.6 Application performance

The most meaningful measure of network performance from the network user's point of view is the degree to which networked applications operate effectively over the network, compared to expectations. It typically is measured in terms of the response time experienced by a user when executing a task requiring interaction with another machine also connected to the network.

Application performance is impacted by the response time of the network itself, error rates, data throughput, and network availability. In the best case, the performance of a networked application can approach that of the same application running completely locally (not requiring interaction with other networked machines). The goal of the network manager should be to provide network services such that users running networked applications do not perceive the impact of the network on application performance. This is done by maintaining network performance at a level high enough that users do not perceive the added delay caused by transactions across the network. The primary metric for application performance is response time as experienced by the user/application for each transaction.

## 15.3 Factors Affecting Network Performance

Although there are many factors that have impact on network performance, the most common ones are traffic congestion, device overload, network topology, broadcast traffic, bandwidth constrictions, device configuration, and application software design.

### 15.3.1 Traffic congestion

The performance of any network at some point will be limited by the amount of traffic it can carry. The raw bandwidth of the transmission medium has a maximum capacity that is defined by the bit rate or clock rate. The usable bandwidth is further restricted by overhead functions such as data frame formatting, addressing, routing, error checking, receiver synchronization, and media access schemes, to name a few.

A useful rule-of-thumb for estimating the traffic-handling capacity of a network segment is the total media transmission capacity divided by the number of devices sharing that bandwidth. For example, if 10 users are connected to a 10 Mbps Ethernet segment, each user is provided an average of

$$\frac{10\text{Mbps total network bandwidth}}{10 \text{ users}} = 1\text{Mbps available bandwidth per user} \quad (15.3)$$

if the traffic were split equally. Such a network would provide an average throughput capacity of 1 Mbps (125,000 bytes/sec) for each user. This capacity then must be evaluated in light of the anticipated requirements, which will depend on the application. For example, 125 kbytes/sec is more than adequate for word processing applications, but probably falls far short for graphics-intensive computer-aided design work.


### 15.3.2 Device overload

Nodes attached to the network also have a finite traffic-handling capacity. End-user devices are limited by the rate at which they can transmit and receive data. Interconnect devices (such as bridges, routers, switches, and servers) must carry traffic for many users. Since these devices must make packet forwarding decisions independently for each frame they receive, their traffic capacity is defined in terms of *maximum frame rate*. When interconnect devices reach their maximum traffic capacity, they could begin to discard any additional data traffic. Discarded data packets ultimately must be detected and retransmitted by the end stations, degrading application performance.


### 15.3.3 Network topology

The choices of network technology and access method can limit the performance of the network, as described in Chapter 14. Of equal or even greater impact, however, is the segmentation of a large network into individual rings or collision domains. These segments are connected to one another using interconnect devices such as bridges, routers, or switches. Transactions within each domain are most efficient when the size of the domain is kept to a minimum (i.e., the smallest number of devices). Interactions between domains, however, are slowed by the necessary interconnect devices. Interdomain interactions are most efficient when the number of domains is minimized. The key to optimizing network performance is to design the topology to fit the required traffic patterns.

Consider the 10Base-T network shown in Figure 15.1. The “A” clients communicate mostly with server A, and “B” clients communicate mostly with server B. The



**Figure 15.1** In this example, all clients and all servers are attached to the same Ethernet network segment. Sometimes called a “flat” network, this configuration often results in undesirably high levels of traffic volume, broadcasts, and collisions, which have a negative impact on network performance.



**Figure 15.2** By splitting the network in half with a router, and placing clients in the same segment as their most often used server, the levels of traffic volume, broadcasts, and collisions on each segment are greatly reduced, resulting in improved network performance.

single-segment topology shown is not optimal, however, because all traffic must traverse the entire network.

By segmenting the network into two separate collision domains and isolating the A client-server traffic from the B client server traffic, the contention for the 10 Mbps available Ethernet bandwidth is reduced by half (Figure 15.2). Only a minimal amount of traffic must pass between segments.

The device used to segment the network in Figure 15.2 is a *router*, which is a device used to forward data packets from one destination to another based on the addressing information contained in the Network layer protocol, layer 3 of the OSI model.


### 15.3.4 Broadcast traffic

Several types of operations generate broadcast traffic, which consists of data packets sent to multiple devices simultaneously (Figure 15.3). These operations include packet routing, address resolution, service advertisements, and booting of diskless workstations across the network, to name a few. Broadcast traffic can be a significant load on network performance because it consumes not only transmission bandwidth, but also processing bandwidth within the receiving devices.

The effect of broadcast traffic can be minimized by network segmentation, without compromising those operations that do require the use of broadcast traffic. MAC-level broadcasts are confined to individual segments and do not pass through routers. Referring again to the example of Figure 15.1 and Figure 15.2, the broadcast traffic in the first case goes to each station on the network. In the second (segmented) case, broadcast traffic is restricted to each segment separately, effectively reducing the amount of broadcasts at any point in the network by half.

### 15.3.5 Bandwidth restrictions in WAN links

Just as fluid flow is limited by the diameter of a pipe, the flow of data packets is limited by the available transmission bandwidth. In a large private network, the most se-



**Figure 15.3** Broadcast traffic, as opposed to unicast, is sent to all destinations on the network, which consumes computing bandwidth in each device because each must examine and potentially respond to the broadcast message.

vere bandwidth restrictions normally are encountered on wide area internetwork links. Geographically dispersed LANs running at 10 Mbps or more often are interconnected via wide area links at 64 kbps, 1.544 Mbps, or 2.048 Mbps. Obviously, only a small portion of the total LAN traffic can be carried over the WAN link.

Wide area network bandwidth very often is an expensive and therefore precious commodity. The amount of WAN bandwidth you allocate for each wide area connection should be high enough so that it doesn't significantly restrict data throughput during peak loading, yet no so high that too much of the traffic handling capacity goes unused most of the (non-peak) time. Optimally sizing WAN connections in terms of their bandwidth, and managing the use of that bandwidth over time, are two processes that can have a very great impact on the overall performance of a geographically dispersed network.

### 15.3.6 Device configuration

The configurations of end-user nodes and segment interconnect devices also can affect network performance. Common examples of configuration parameters that have performance impact include:

- Maximum packet size
- Sliding window size
- Subnet mask
- Default gateway

**Maximum packet size.** Configuring the maximum allowable data packet size for each network segment can affect data throughput and application response time. Choosing a large maximum data packet size minimizes the amount of protocol overhead (routing and error checking) and minimizes delays associated with packet fragmentation and reassembly. But allowing large packets over relatively slow transmission media (such as wide area serial links) increases the probability that transmission errors will result in packet retransmissions, resulting in session timeouts for delay-sensitive data traffic such as LAT or SNA. Choosing the optimum maximum data packet size for each wide area network link means balancing between these opposing effects.

**Sliding window size.** A *sliding window* protocol allows for transmitting several data packets before an acknowledgment from the receiver is required, rather than waiting for individual acknowledgment of each packet. The size of the sliding window is the number of outstanding packets allowed before the transmitting station must wait for a receipt acknowledgment. Setting the sliding window size high can improve the performance of transmission links with long delays, such as satellite WAN links, but a large sliding window can have negative performance effects in those cases when high bit error rates necessitate frequent retransmissions.

**Subnet mask.** An *IP subnet mask* is used by each station on an Internet Protocol network to separate the network portion of a destination IP address from the node portion of the IP address. Misconfigured subnet masks can result in local segment

traffic being sent through routers unnecessarily, further loading these devices and potentially affecting network performance.

**Default gateway.** The *default gateway* is the router used by an end node to deliver IP traffic destined for a remote network segment. Designation of a particular IP router as the default router will increase the IP traffic through it and potentially affect the performance of that device if it becomes overloaded.

### 15.3.7 Application software design and configuration

User applications that operate across the network should be designed and configured to do so efficiently. A poorly designed application can make even the most efficient network appear to the user to perform poorly. Unfortunately, it is not easy for a network administrator to evaluate the network efficiency of an application without actually trying it out on the network.

For example, file accesses done over the network should maximize the data block size of each read request to minimize packet routing and error detection overhead associated with multiple smaller packets. Devices transmitting data should wait for minimum available receive buffer space before proceeding to transmit many small and inefficient data packets.

Such design limitations of networked applications are difficult for the network manager to evaluate and often do not become apparent until after they are installed. When selecting a new major application for use on the network, it is good practice to first evaluate reference sites where the application has been installed successfully on a similar network. Be very cautious if no such references are offered.

## 15.4 Network Performance Metrics

Given the many different criteria for evaluating network performance, it is no surprise that there are many different ways to measure how well a network is operating. Network engineers need to evaluate network performance in concrete, measurable terms. The overall performance of a network can be characterized by frame error rates, frequency of collisions, node-to-node response time, average frame length distribution, or data throughput. Each of these is a very measurable quantity, and each reveals something about the overall performance of the network as perceived by the users. Several different statistical measurements can be made on data traffic in order to evaluate network performance.

### 15.4.1 Traffic rate

Traffic rate measurements can uncover congestion-related problems, one of the most common limitations to network performance. Traffic rate is most often measured in one of three ways:

- Data throughput
- Percentage utilization
- Frame rate

*Data throughput* is the measure of traffic volume actually being carried by the network, typically expressed in kilobytes per second (kbytes/sec). The device measuring data throughput simply counts the number of data bytes transmitted on the network over some measurement time interval, then calculates and reports the average data rate for that interval (total number of bytes divided by the length of the measurement time interval).

The maximum data throughput of a network is limited by the clock rate used. For example, a 10 Mbps Ethernet LAN segment has a maximum data throughput of

$$\frac{10 \text{ Mbps}}{8 \text{ bits per byte}} = 1250 \text{ kbytes/sec} \quad (15.4)$$

Other factors, including minimum interframe spacing specification, the required frame preamble, and the collision detection and retransmission process, will reduce the practical maximum data throughput on an Ethernet network still further.

The usable data throughput from the user's perspective is further restricted by the overhead added to each data packet by the lower-layer network protocols. Data throughput can be defined at each of the various layers in the protocol stack. The data throughput as defined in this section refers to the ISO Data Link layer, or the MAC (Media Access Control) sublayer thereof.

Data throughput also can be measured on each individual connection at the Network layer (such as IP), Transport layer (such as TCP), or Application layer (such as FTP) by counting only user data bytes for that protocol layer over the measurement time interval. While this requires a more sophisticated measurement tool, it more closely matches the actual data throughput experienced by a network user.

*Percentage utilization* measurements indicate how much of the available transmission bandwidth is being consumed. *Transmission bandwidth* in this case refers to the bit rate or clock rate of the network. Percentage utilization is defined as the data throughput expressed as a percentage of the maximum traffic handling capacity of the transmission medium:

$$\text{utilization \%} = \frac{\text{measured data throughput}}{\text{raw bandwidth of the transmission medium}} \times 100\% \quad (15.5)$$

A 10 Mbps Ethernet segment with 625 kbytes/sec (5 Mbps) of measured data throughput would have a utilization percentage of

$$\frac{5 \text{ Mbps}}{10 \text{ Mbps}} \times 100\% = 50\% \quad (15.6)$$

In another example, a leased 64 kbps transmission line has a maximum traffic handling capacity of 8 kbytes/sec (in each direction). If the actual data throughput on that line measured over some time period is 4 kbytes/sec, then the utilization percentage is 50 percent for that period.

Percentage utilization is the best indication of traffic congestion in the network transmission media. It is the percentage of available bit times or timeslots that are actually being used.

*Frame rate* is a more useful indication of congestion in interconnect devices, which must examine each frame for packet-forwarding decisions. Frame rate is simply the

number of frames being transmitted across the network in a given time interval. In this case the measurement device counts frames transmitted on the network over some interval, and then calculates and reports the average frame rate for that interval (total number of frames divided by the length of the measurement time interval).

Frame rate is expressed in terms of frames per second. It is the best indication of device overload due to traffic congestion.

All three measures of traffic rate (data throughput, percentage utilization, and frame rate) are best evaluated in the context of regular network baselining, rather than in an absolute sense. Changes in traffic volume can be observed over time, and trends can be used to predict congestion before it occurs. Methods and tools for measuring data throughput, percentage utilization, and frame rate are covered in section 15.5; a complete description of the baselining process also appears in section 15.6.

#### 15.4.2 Errors

Error rate measurements reveal the overall health and integrity of the physical transmission media and the attached devices. Error rate is calculated by counting each error type over some measurement time interval and reporting the average number of such events over that interval.

The format of error types varies for each network technology. In Ethernet networks, transmission-related problems typically exhibit relatively high numbers of misformed frames, including *runts* (frames that are too small), *jabbers* (frames that are too large), misaligned frames (frames that do not end on an 8-bit character boundary), and frames with bad frame check sequences (indicating a bit error in the transmission). For a more detailed description of error types for Ethernet and other common LAN and WAN networks, refer to Chapter 14.

Reporting of error rates sorted by MAC (Media Access Control) source address (top error sources) can quickly isolate the offending node or nodes in cases of faulty network interface cards.

#### 15.4.3 Collisions

*Collision rate* counts the total number of collision events over some measurement time interval and reports an average collision rate in terms of events per second for that interval. Collision rate is specific to Ethernet technologies, including 10Base-T and Fast Ethernet/100Base-T.

The collision rate on a network is a useful indication of the degree to which the transmission medium is saturated with traffic. As the collision rate increases, so does the probability that a node will experience a delay due to traffic congestion when transmitting a frame. The collision rate must always be evaluated with respect to the packet rate for the network over the same time interval. A useful (though debatable) rule of thumb is that the collision rate should not exceed 10 percent of the packet rate.

#### 15.4.4 Broadcast, multicast, and unicast frame rates

*Broadcast frames* are sent by a node to every other node on the network, as defined by the destination MAC address FF-FF-FF-FF-FF-FF. Broadcast frames are most of



ten used by nodes advertising their existence, by nodes looking for a service, by nodes doing source route bridging, or by nodes or routers looking for the physical address of a given destination network address. *Multicast frames* are sent to groups of addresses as indicated by the least significant bit of the first byte in the MAC address. *Unicast frames* are those sent to a single node identified by the destination MAC address.

While some amount of broadcast and multicast traffic is normal on most networks, excessive numbers of such frames can degrade network performance significantly. In addition to consuming transmission bandwidth, broadcast and multicast frames will consume precious CPU cycles in many (all) network-attached devices, because each must evaluate the contents of these frames. Proper use of multiprotocol routers can reduce greatly the reach and impact of broadcast/multicast traffic by filtering these frames to and from adjacent segments.

Broadcast, multicast, and unicast traffic levels are measured by counting each of these frame types, identified by destination MAC address, over some measurement time interval. The results are reported as average frames rates for each measurement interval over time.

#### 15.4.5 Traffic distribution

Measurements of traffic distribution by node, by connection, or by protocol can be used to determine which users and which applications are consuming network bandwidth. Understanding how network bandwidth is being consumed allows for intelligent allocation of precious bandwidth in situations where traffic congestion might otherwise limit network performance.

Traffic distribution measurements are made both by frame counts and by byte counts, expressed either as an average rate per second over each measurement time interval, or as a cumulative total since the beginning of the measurement period. Frame rate totals are valuable for evaluating forwarding device overload (in bridges, routers, etc.), while byte rate is more useful when dealing with congestion of the transmission media.

*Node statistics* are counts of frames and bytes transmitted and received by each node or station active on the network over the measurement time interval. Monitoring traffic by node is useful in determining which nodes are responsible for generating or receiving the most data. When node statistics are reported in order of frames or bytes transmitted, it is referred to as a *top talkers* measurement.

*Connection statistics* keep track of the number of frames and bytes sent and received between each pair of stations communicating over the network over the measurement time interval. Connection statistics can be tracked at different protocol levels, including MAC-level connections (by Data Link layer or MAC address) and Network-level connections (by network address, such as IP or IPX).

*Protocol statistics* measure data traffic by protocol type. Each frame is counted according to the protocol type of the information it is carrying, as determined by the protocol type field that is present in the frame. These statistics track data traffic both by frames and by bytes.

Protocol statistics can be measured and reported at various levels within the different protocol layers. At the MAC level, i.e., a network protocol being carried in the

MAC frame, examples include IPX (Novell NetWare's Internetwork Packet Exchange protocol), as well as IP. At the Network layer, i.e., a transport or application protocol being carried within the packet, examples from the IP stack include FTP (File Transfer Protocol) and Telnet.

#### 15.4.6 Frame length distribution

As a general rule, the largest average frame size will result in maximum network efficiency. This is because each frame transmitted over the network must carry with it a certain amount of overhead for addressing, error checking, and other necessary functions. Maximizing the amount of data carried in each frame minimizes the number of frames needed and hence the amount of overhead used to communicate a given amount of information. Minimizing the frame rate also will reduce the burden on interconnect devices, which must make a forwarding decision on each frame received.

There are some situations where using large frame sizes can actually degrade network performance. On a transmission facility with significant bit error rates, for example, the probability of retransmission increases with increased frame length.

Using the maximum frame size also might result in significant delay for time-sensitive applications using the same transmission facility, particularly if that transmission facility is a natural bandwidth bottleneck (as most wide area links typically are). Large frame buffers at either end of such WAN transmission facilities can exacerbate this problem further and eventually result in disconnected conversations due to expired timeouts. Where bulk data transfers must coexist with time-sensitive conversations on the same network, some compromise in setting the maximum allowable frame size will be required.

Frame sizes used by stations communicating over a network can be adjusted by configuring network interface cards, applications, and interconnect devices such as routers. By observing the average frame length distribution for each of the protocol stacks in use on a network, a network engineer can observe the efficiency of each protocol.

Frame length distribution is measured by counting the number of frames observed on the network that fall into various length ranges (0–63 bytes, 64–127 bytes, 128–511 bytes, etc.) over some measurement time interval. The results can be reported as a frame rate for each range over time, as an average frame size for each protocol over time, or as a cumulative total number of frames in each length range for each protocol over the entire measurement period.

#### 15.4.7 Response time

*Response time* measures the round-trip time delay experienced by a transaction across the network. Response time most often is measured by the *packet internet groper* or *ping* utility on TCP/IP networks, and other similar utilities on other network types. A *ping transaction* consists of an Internet Communications Message Protocol (ICMP) echo request from one IP network node to another and the echo reply sequence sent back. As such, the ping measures the response time of the target

**TABLE 15.1 Network Performance Metrics. This is a summary of the network performance metrics discussed in this chapter sorted by protocol layer in the OSI model, including a listing of how each is used to evaluate network performance.**

OSI Protocol Layer	Function	Network Performance Measurement	Use
Application Transport	Session Management Data Sequencing	Data throughput	Efficiency of Application
		Frame rate	Load on Interconnect Devices
		Connection Statistics	Monitor Application Connections
		Protocol Statistics	Bandwidth Consumption
		Frame Length Distribution	Application Efficiency
		Application Response Time	User Response Time
Network	Addressing (Routing) Packet Fragmentation	Data Throughput	Efficiency of Network Layer
		Frame Rate	Load on Interconnect Devices
		Node Statistics	Bandwidth Consumption
		Connection Statistics	Bandwidth Consumption
		Protocol Statistics	Bandwidth Consumption
		Frame Length Distribution	Network Protocol Efficiency
		Node-to-Node Response Time	Network Response Time
Data Link (MAC)	Media Access Control Addressing (Physical)	Data Throughput	Congestion of Network Media
		Percentage Utilization	Congestion of Network Media
		Frame Rate	Load on Interconnect Devices
		Errors	Health of Transmission Media
		Broadcast and Multicast Frame Rates	Network/CPU Resource Drain
		Note Statistics (Top Talkers)	Bandwidth Consumption
		Connection Statistics	Bandwidth Consumption
		Frame Length Distribution	Network Efficiency


IP nodes plus the network itself (meaning the Physical, Data Link, and Network protocol layers) between the two nodes.

The *actual delay* experienced by a network user or application is the network response time plus the delay through the target device, plus the delay induced by higher-layer protocols and their interactions with the Network layer. *Application response time* can be calculated by measuring the time delay from a transaction request to the corresponding reply, at the highest protocol layer in use. Application response time can be very difficult to measure and may vary considerably from application to application.

Table 15.1 presents a summary of the network performance metrics discussed in this chapter.

## 15.5 Methods and Tools

The two principal means of assembling the network performance metrics are *protocol analyzers* and *distributed monitoring systems*; some examples of the latter type effectively can be made part of the overall network structure.



**Figure 15.4** This simplified block diagram of a protocol analyzer includes a capture filter for selective capture of data traffic, a circular data buffer for continuous storage of the most recent data traffic, and a display filter for selective decoding and display of captured data.

### 15.5.1 Protocol analyzers

A *protocol analyzer* is a standalone unit that can be moved from one network segment to another relatively easily. It simply attaches to the network, captures data, and analyzes information contained in the frames it captures.

A protocol analyzer is used as an in-depth troubleshooting tool. Its primary function is to capture, decode, and display data frames and all of the information they contain at each of the various protocol layers (Figure 15.4). Basic functionality includes *capture filtering* (selectively capturing frames based on address, protocol type, pattern match, and other criteria); *display filtering* (selectively displaying captured frames based on address, protocol type, pattern match, and other criteria); *triggering* (taking a specified action based on the occurrence of some specified event); and *post-capture searching and analysis* functions.

Most protocol analyzers also analyze data traffic and report various statistics about that traffic. Common statistical measurements include percentage utilization, data throughput, packet rate, error rate for a number of different error types, collision rate, top talkers, and protocol distribution. These statistical measurements are particularly valuable for characterizing network performance. Besides statistical analysis and data capture, some protocol analyzers provide expert analysis and other applications that are designed to help troubleshoot network problems quickly.

Protocol analyzers should be capable of connecting to many different network interfaces, such as Ethernet, FDDI, T1, DS3, and so on. The analyzer should be able to examine all the traffic seen on the network under heavy traffic load. The most stressful condition for an analyzer is determined by the frame rate, not percentage utilization, because each frame must be captured and analyzed individually. An analyzer should be capable of capturing (or selectively capturing) and analyzing all the data present on the network and saving that data to a trace file.

### 15.5.2 Distributed monitoring systems

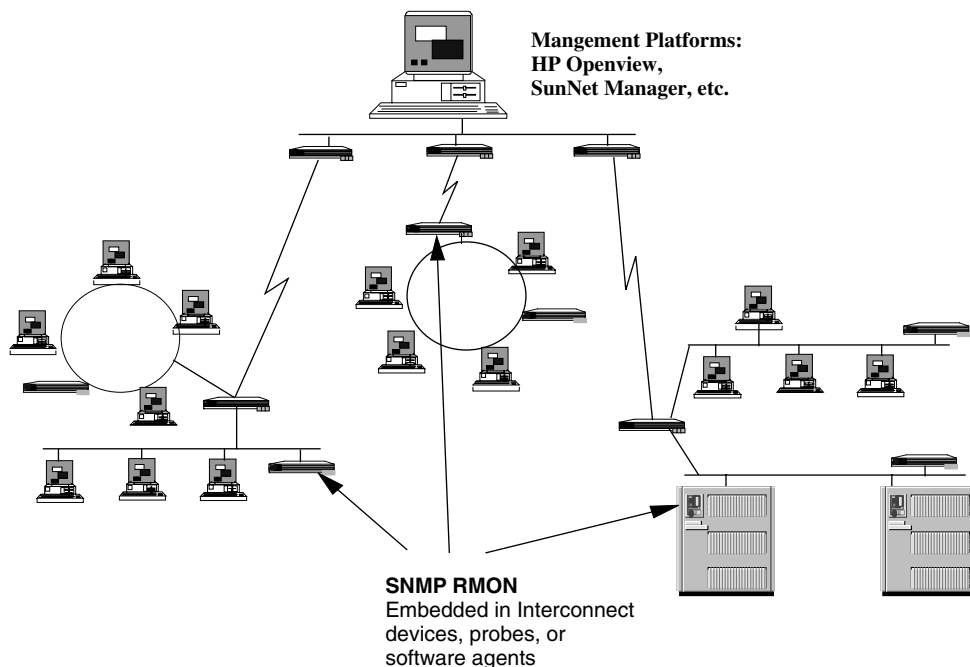
Distributed monitoring systems are available in different varieties, but usually have one important similarity: They are based on the remote monitoring (RMON) standard, which defines a set of statistics to collect.

Some distributed monitoring agents are standalone hardware probes designed to connect to a network, passively monitor it, and send analysis data back to the management console. Other distributed monitoring agents are built into network interconnect devices such as routers and hubs. The data they collect also must be transmitted back to a management console. A distributed monitoring agent also could be software that resides on network nodes, using the node's network interface card (NIC) to monitor the network passively and send data back to the management console (Figure 15.5).

The manager application can communicate with hundreds of managed devices. This communication is accomplished using a protocol known as SNMP (Simple Network Management Protocol). For ease of use, the manager application generally provides a graphical user interface. A *management information base* (MIB) resides on the managed device and stores information, such as number of connections or speed of transmission. The agent resides on the managed device and communicates with both the manager and MIB.

A distributed monitoring system should provide access to enough statistical information to perform baselines and benchmarks. At a minimum, this should include all network analysis information provided in the SNMP RMON and RMON 2 standards. Additional capabilities might be provided by the manufacturer using proprietary extensions to the MIB.

Many network managers use distributed monitoring systems to baseline multiple network segments simultaneously for long sample periods with large sample intervals.



**Figure 15.5** Distributed monitoring systems extract network performance data collected by remote monitoring (RMON) agents placed throughout the network and present them on a central console for interpretation by the network manager.

This yields valuable intersegment traffic analysis and long-term trends information. Protocol analyzers often are used to perform complementary baselines and component benchmarks on a single network segment or component using smaller sample periods and small (1-second) sample intervals. The resulting analysis is more narrowly focused and detailed, providing the information necessary for fault isolation and performance tuning of components and applications.

## 15.6 Network Baselineing

*Baselineing* is a process for network performance characterization, from which grows the process of network optimization. A *baseline* is a set of statistical measurements made over a period of time that characterizes network performance. A complete baseline will include all of the network performance metrics listed previously, and perhaps more. A baseline is a comprehensive snapshot of a network's overall health.

### 15.6.1 Benefits of network baselining

Doing a network baseline often exposes inefficiencies in network operation, providing immediate opportunities for improving network performance. In most cases the baseline uncovers inefficiencies in the network that are not serious enough to prevent communication, but degrade overall network performance. For example, low average packet sizes might be caused by insufficient buffer memory, or routing errors caused by misconfigured workstations.

Routine network baselining provides many other benefits in addition to uncovering inefficiencies. By characterizing network operation on a regular basis, an operator will gain a much deeper understanding of exactly how the network functions. A baseline provides the information needed to understand and manage network operation. Timely alerts help track down device congestion, transmission media capacity limits, and other traffic-related problems.

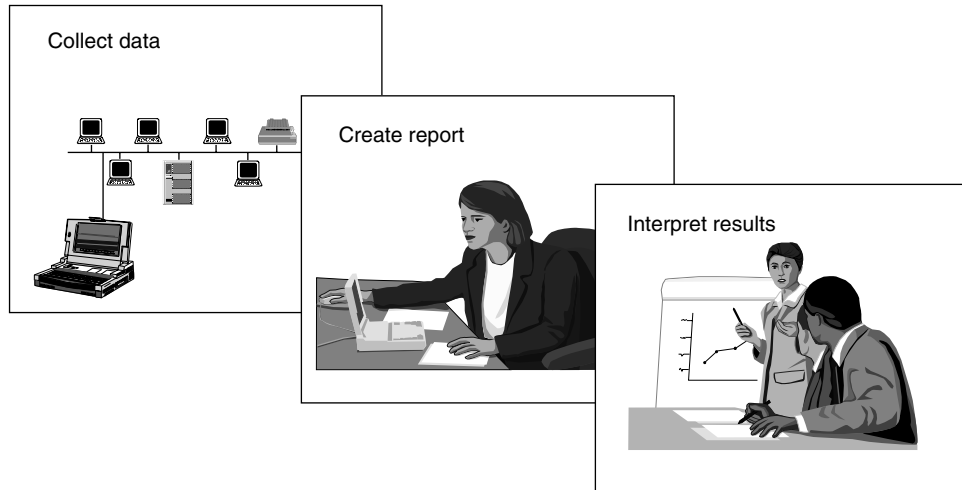
Reports generated by the baselining process can be used to justify hardware or software upgrades. Baselining both before and after upgrades will allow their impact on overall network performance to be evaluated.

Each network baseline is a complete characterization of network operation, and is useful for uncovering inefficiencies and isolating faults. But the real value in regular network baselining comes with the analysis of several baseline reports done at regular intervals on the same network segment.

Comparing baseline measurements done at regular intervals allows the operator to recognize changes and observe trends. Some problems can be anticipated and resolved before they become apparent to users. The information provided in baseline reports also is essential in order to plan for future growth.

### 15.6.2 The baselining process

There are three basic steps in the baselining process: collecting the data, creating the report, and interpreting the results (Figure 15.6).



**Figure 15.6** There are three fundamental steps to the baselining process. First the data is collected, and then a report is generated. The third and most important step is results interpretation, which becomes easier over time as each successive baseline report teaches more about the network.

**Collecting the data.** The first step in the baselining process is data collection. Data for the baseline typically is collected using either a distributed monitoring system or a high-performance protocol analyzer, attached directly to the LAN segment or WAN link to be characterized. Although any network segment can be characterized, network baselines are most often run on critical backbone LAN segments and high-speed WAN interconnections.

Baseline data should be collected for a fixed period of time at regular intervals; a typical baseline consists of data collected for a 24-hour period once per month. The time period for data collection should be chosen to represent typical network traffic, perhaps at a time of moderate to heavy utilization. In order for comparisons of different baselines to be meaningful, the data should be collected over similar time periods.

A sampling interval for the data collection process also must be chosen. The sample interval is the minimum resolution at which data samples are stored by the analyzer. While the minimum sample interval yields the best resolution, it also produces the maximum-size data file, which can become difficult and time-consuming to process.

**Creating the report.** Data collected by the analyzer is imported into an industry-standard spreadsheet program, such as Microsoft Excel, to create a baseline report. Professional-quality tables and charts can be generated with most full-featured spreadsheets; this can be done manually, or done automatically with a reporting tool designed specifically for consolidating protocol analyzer or monitoring system data into network baseline reports. In the case of Excel, these tables and charts are automatically processed into a report in Microsoft Word for Windows format, including a glossary of terms and an explanation of the measurements made.

A complete baseline report should include data on:

- Traffic volume
- Errors
- Broadcast traffic
- Top talkers
- Protocol distribution

**Interpreting the results.** The third and final step in the process is the most important: interpreting the results. Although the graphical presentation of network performance statistics greatly facilitates interpretation, it still requires a certain degree of networking expertise to draw appropriate conclusions from the data.

Every network is different and network performance data must be analyzed on a case-by-case basis. Nonetheless, there are three general rules that apply to most situations:

1. *Look for abnormalities.* To begin with, look for high levels of network utilization, low average data packet size, or a high level of errored frames as indications of poor network health.
2. *Look for changes.* Compare successive baselines and question any significant changes in traffic patterns or error levels. Be sure that you understand and are comfortable with these changes. Notice long-term trends consisting of a series of gradual changes, and try to anticipate the significance of these trends over time.
3. *Learn what is normal for your network over time.* Use baseline reports to set measurement thresholds on the analyzer, so that a future change in network behavior will trigger the analyzer to alert you to the change.

## 15.7 Network Benchmarking

Network benchmarking is a process for evaluating application or network component performance. A *network benchmark* is defined as a set of statistical measurements that characterize the performance of a networked application or network component. A benchmark is a detailed analysis of the performance of a specific application or a specific network component, and is much more narrowly focused than a baseline. A complete benchmark will include all of the network performance metrics of a baseline, and perhaps more. Data for a benchmark study is often filtered, based on the application or on the address of the component.

The objective of a benchmark is to understand and/or predict how an application or a device performs, or how it may impact the performance of the entire network. Benchmarks can be used to evaluate various alternatives for network upgrade, or to anticipate user service levels before a major hardware or software change. Data collected during a benchmark study often can be used for troubleshooting later if a problem arises with the targeted application or component.



The benchmarking process consists of the same three steps as baselining: collecting the data, creating the report, and interpreting the results. The fundamental differences when benchmarking are:

- Apply a filter when collecting data to capture only frames to or from the targeted application or component.
- Use the minimum available sample interval for maximum resolution over a limited measurement period.
- Use a test network when possible in order to isolate the device under test from other fluctuations in unrelated network traffic.
- Perform each benchmark several times and average the results, especially when testing on an operational network.



# Network Interconnection Technologies and Testing

**Marc Schwager**

*Hewlett-Packard Australia Ltd.*

## 16.1 Why Test Networks?

Networks are tested to make sure that they perform properly. Most sophisticated networks require testing in some form, a condition that's not likely to change. There are four different types of tests that are performed, depending on the stage in the cycle of network deployment. The four types are *conformance testing*, *installation testing*, *troubleshooting*, and *network monitoring*. (The subject of testing during research and development is not discussed here.)

### 16.1.1 Types of testing

Here is an overview of the four types of network testing. Each of these is explored in more detail later in the chapter.

**Conformance testing.** When a network device is being designed, or when it is purchased against a specification, conformance testing commonly is performed. Two different aspects of conformance testing can be considered. Does the device implement protocols correctly? Does it meet the performance requirements for the intended application? Generally a high-end protocol analyzer with sophisticated simulation capabilities is used for conformance testing.

**Installation.** When a device is first installed in a network, a series of tests may be run to verify that it is functioning normally. In some cases a standard turn-up process has been established, which can check out different aspects of the device and complete any initialization and configuration needed for operation. It is a good idea to perform some kind of installation tests, even if they are merely staged on a small pilot network. If a device has any defects, it is better to find out before hooking it on a live network.

**Troubleshooting.** Troubleshooting is, unfortunately, the most widely practiced type of network testing that occurs today. Routine application of the other three types of testing can keep fire-fighting to a minimum. Generally you will want the most powerful tools you can obtain for this task. You also will be ahead of the game if you have information from remote monitors to help pinpoint the problem. The expert systems that ship with today's advanced protocol analyzers can be a help, but they are no substitute for an intimate knowledge of the network, and the protocols being carried.

**Monitoring.** Monitoring your network is equivalent to having dashboard gauges in your car. It can alert you to trouble, and help you regulate performance and plan for upcoming changes. You wouldn't drive without some form of instrumentation, nor should you try and operate a large network without it. The basis of network measurement and testing is a capability called *promiscuous monitoring*. This is a standard configuration capability of many networking chips. When in promiscuous mode, a network chip will accept any and all traffic that it sees on the network and make it available to the measuring software. This is the foundation for today's protocol analyzers and network monitors.

### 16.1.2 Categories of testing

Since today's network interconnects provide more management capability than ever, do you still need test equipment today? If so, for what would this test equipment be used? There are basically three classes of testing that you may need to do at some point, regardless of the capability of your network devices.

**Physical testing.** In general, an interconnect will give a pass/fail indication of a cable problem. Problems with cable breaks are still an issue with networks today. Of course there are more subtle problems as well, such as distance and crosstalk limitations. The cable tester will tell how far away the cable break is, allowing the repair spot to be pinpointed. In the case of distance and crosstalk problems, cable testing is indispensable. It is a good practice to measure LAN cables to ensure conformance to specification.

**Transmission.** Transmission problems, such as noisy lines, are a fact of life. The equipment used for this class of measurement includes capabilities such as bit error rate testing (BERT), and transmission impairment measurements (TIMS). Although these capabilities do not exist in interconnects, the interconnect device might give an indication of a problem based on performance statistics.

**Protocol testing.** The protocol layers above the Physical layer are the most complex to troubleshoot and require the ability to capture frames and decode them, often tracing conversations across the network to look for configuration problems. The most sophisticated tool for this is the protocol analyzer. Often these tools have expert systems that can help classify problems.

### 16.1.3 Analyzers vs. built-in monitors

From a network monitoring perspective, interconnects are gaining more monitoring power than ever. Many now have sophisticated statistical monitoring as well as

packet capture and alarm generation capabilities, thanks to the *Simple Network Management Protocol Remote Monitoring Management Information Base* (SNMP RMON MIB, an 11-letter acronym!).

Does this mean that it is time to put the protocol analyzer away? In some cases yes, but the analyzers generally have much more comprehensive monitoring and analysis capabilities than even the most robust RMON solution. In addition, analyzers are portable and can be moved to remote segments where RMON agents might not exist. For most network problems you can start with RMON, but if you don't have an RMON MIB on the segment in question, or if the going gets tough, a protocol analyzer is indispensable. One final note on monitoring is that if the network fails, you may not be able to reach the monitor unless a special, out-of-band line has been installed.

## 16.2 Conformance Testing

When testing for conformance, you will want to verify that the device has implemented the protocol stack accurately and completely, and that it will meet your performance needs.

### 16.2.1 Protocol conformance

Even with the most well-used protocols, there are occasions when a device will fail to interoperate. This could be due to a bug in the device, or a difference in interpretation of a specification—especially in a new specification. Full-blown conformance testing is usually left to the vendor. It requires a sophisticated set of tools with full simulation capabilities.

If you are interested in conformance testing, there are conformance test suites available from a variety of sources. You usually can find these through the Internet without too much trouble. The general technique for the test involves generating a well-known set of protocol messages, monitoring the response from the device, and comparing the output to a known-good reference. It is tedious. The conformance test for ATM signaling contains well over 300 tests!

### 16.2.2 Performance

Performance is a more practical area for a user to test. It is a good idea to check certain aspects of device performance before installation if you are going to be stressing the device in an unusual way. Vendors often publish performance tests that may be used as a guideline, but they have to be interpreted in terms of your own network traffic.

For instance, a router vendor could publish a performance specification stating the number of packets per second that may be routed. This specification might not take into account variable packet sizes, the number of different protocols being routed, or the number of different source and destination addresses being routed. Each of these can dramatically effect device performance when together they reach a certain critical size or mix. This generally varies by device and is dependent on how the device was designed. A common technique is to capture some representative traffic from the network and use this in conjunction with the traffic generation

capabilities of the protocol analyzer. To perform the test, replay the traffic on a test network to which the device is connected and observe the results.

### 16.3 Installation Testing Procedures

Installation testing is usually done to verify that a device about to be installed works well. If it is the first time the device has been installed, conformance testing could be in order. In many cases, similar devices from the same vendor have been installed many times before, so this is a routine check out before turning the device over to the live network. Most vendors will have in their manuals a procedure and a checklist for installation. Many devices have sophisticated diagnostics that will verify correct operation.

In cases where a new type of device is being installed, or one from a new vendor, *staging* might be useful. This involves setting up an isolated pilot network and observing device performance. It is also a useful way to become familiar with the operation of the device in a nonthreatening environment.

### 16.4 Troubleshooting a Network

So you have an analyzer. How do you hook it up and then what do you do with it? Before disconnecting potentially important cables, read these sections.

#### 16.4.1 LAN connections

**10Base-T.** Testing 10Base-T is easy. Connect an analyzer to any point on the hub. You will immediately be monitoring all the traffic on the segment. Connection will be via RJ45 connector, or a transceiver if you are plugging into an AUI on the hub.

**10Base2 (Thin Ethernet).** With this type of cabling you will need to find an exposed *T* connector to use as a connection point. If there isn't one, look for a device like a noncritical PC that can be unplugged for a while and use that tap. Failing that, get an extra *T* connector and some cable and connect at the end of the network. When you break the cable to insert the tester, you will cause a massive blast of collisions that will abate once you have hooked on your wire and replaced the end-cap terminator. If you insert in the middle of the cable, keep in mind that there must be a few meters of cable minimum between taps.

**Token-Ring.** Hook the analyzer into an open port on the MAU. You will immediately be monitoring all the traffic on the segment. Depending on the analyzer, you may choose to monitor the network without becoming an active part of the protocol on the net. This is important if you are looking at the behavior of the transfer of Active Monitor responsibility and you don't want the analyzer to become the Active Monitor.

**LAN switches.** A switch might have a monitoring port that you can use. If this is not the case, you need to connect between the switch and the other node of interest. This requires the analyzer to act as a physical repeater, and it must have both an input and an output port in order to hook in correctly. If this is not available, an easy

trick is to obtain an inexpensive hub and connect the switch, the node, and the analyzer into it. Because of the nature of a switch, you will see only the traffic to and from the device you are monitoring.

#### 16.4.2 LAN troubleshooting hints

Here are a few basic hints on troubleshooting a LAN. There are many problems that can occur. If the network worked once and suddenly stops, however, this indicates either a change (new gear, configuration) or a piece of equipment has failed. Methodical isolation in combination with instrumentation normally can isolate a problem fairly quickly.

When setting up the network keep, a list of MAC addresses, upper-layer (IP, IPX) addresses, locations, and owners. The result of troubleshooting is often the MAC address of the offending node. Without the list, fixing things is difficult. Remember that the first portion of the MAC address is the vendor ID, which can help somewhat if you have lost the list, but it is often obscure. (The “vendor address” of a PC is the vendor that made the network interface card inside it.)

Recommended tools include a cable tester and a continuity checker. With a little engineering a quick go/no-go cable checker for 10Base-T can be cobbled up out of a transceiver and a 9-volt battery: Connect the battery to the power pins on the AUI side of the transceiver. Plug the media access side of the transceiver into the cable and check if the traffic light becomes active.

In bus-topology, coax-based LANs, physical problems are a large portion of problems found. This is why most people move to structured wiring (star topologies) as soon as practical. With a bus, the fault domain spans the entire cable. A typical failure in an office environment is caused by someone moving furniture around and disturbing the cable. Here are some things to look for:

- Improper cable length
- Cable crimps
- Damaged or missing terminators
- Card or transceiver failure
- Malfunctioning printers, servers, or workstations

It is easy to add cable to expand the LAN, but if the length specifications in the connection rules are exceeded, problems can occur. The cable length specifications are derived from the electrical properties of the medium (such as impedance) as they affect signal levels and timing, the goal being that the interframe gap will separate packets into discrete events for a given run of cable.

Crimps in the cable can cause reflections and decrease the signal quality, and in extreme cases cause physical breaks. Also cast a suspicious eye on coax running near heat sources such as baseboard heaters. The inner insulation can soften enough to allow the center conductor to contact the braid—without any visible evidence on the outer insulation.

End caps (terminators) should be firmly in place, and along with the T connectors should be physically isolated. They also must be of the proper impedance for the

type of coax. Removing an end cap will generate continuous collisions because of signal reflections back toward the origin.

A good cable checker will show length, as well as reflections caused by crimps. Remember also that adding devices on coax requires a certain minimum spacing between devices.

Other problems can occur when a card or transceiver fails. If a problem goes away when a particular device is shut down, the problem source is well on its way to identification. Likewise, reboot printers and server before bringing out the analyzer, to see if the problem is a transient software anomaly and clears up.

Protocol-level problems generally require a protocol analyzer.

For LANs with repeating devices (e.g., standard 10Base-T), each cable is an isolated physical fault domain. The protocol layers will affect the entire network. If a problem is isolated to a single cable/user, look for physical problems. For a twisted-pair hub, or a Token-Ring MAU, the light on the device next to the cable will give an immediate indication of cable failure. Token-Ring will isolate physical problems and drop the offending device out of the ring. Hubs generally recover very quickly after resetting, so if you suspect that the hub has locked up, a quick reset might bring it back.

If the trouble spans all cables in a coax environment, however, or involves multiple hubs, then it probably is a protocol problem. One exception to this occurs if the Ethernet specs are exceeded with regard to the number of repeaters in a connection: A maximum of three is allowed before you need a bridge or a router.

In bridged LANs, broadcast problems escalate, passing through the bridge transparently and affecting the entire network. If you have performance problems with your LAN, connect an analyzer and look at the broadcast levels. A rule of thumb for an Ethernet is fewer than 100 broadcasts per second. If you see broadcasts spiking every 30 seconds or so, you probably have *broadcast storms*. Use the protocol analyzer to evaluate the problem. One cause of broadcast storms is a failed network interface card that has a source address of all Fs. It might be possible to isolate this if the IP address is still intact.

Routed LANs almost certainly require an analyzer to find problems. A common cause of problems is duplicate IP addresses from improperly configured PCs. This is avoidable through good configuration control. An analyzer will only see traffic that is routed to a specific subnet; unlike a bridged or repeated network, you are only viewing a portion of the network.

IP routing messages are transmitted via the ICMP protocol. This can be a rich source of information if you are encountering problems reaching a node. Install a capture filter in your analyzer to capture all ICMP traffic. Of course, the most widely used tool to determine whether an end node is reached and alive is the ICMP echo message, also known in IP as a *ping*. (There are equivalents in DECnet, AppleTalk, and others.) Ping will verify that the network routing is intact, and that the destination card is alive and responding. It will not give you any information about what is happening above the card; the server might be locked up completely, but the network card is alive.

Managing across the entire routed network requires the use of distributed monitors, such as RMON probes.



### 16.4.3 WAN connections

Because of the point-to-point nature of WANs, connecting a monitoring device usually requires a Y cable and breaking the link. Hook up at the CSU/DSU or the router. Normally you will need to break the link in order to insert the cable. If the link is already down due to a problem, then this is not a problem.

### 16.4.4 WAN troubleshooting hints

WAN troubleshooting can become complicated for the simple reason that you don't own your WAN—the carrier does. Finger-pointing often can ensue. There are two ways to avoid this. First, find an exceptionally good WAN supplier, and second, purchase a WAN analyzer so you can verify suspected problems with the WAN before calling your provider.

WAN problems fall into two areas. The first is the physical transmission. If the line quality deteriorates and the bit error rate goes up, severe performance problems can occur. In order to check the physical line quality, you will need an analyzer capable of TIMS and BERT testing. You also will need to take the line out of service in order to check this! Both of these tests require traffic be generated on the line. Your service provider usually has extensive test capability and normally can do this for you. These problems can be intermittent (when it rains, for example, and the lines get noisy); you might want the capability to test on short notice. These analyzers also can provide a rich set of statistics concerning your WAN performance.

Other sources of WAN problems include clocking-related issues and jitter. These can cause intermittent signal loss, as well as total link failures. Timing on DS1 circuits typically is supplied by the carrier. Isolating timing problems requires specialized test equipment. In-service testing such as timing slip analysis can be performed to verify correct operation of a device. These tests often will trace problems back to misconfigured equipment, or device faults. Jitter is caused mainly by network equipment such as repeaters and multiplexers. If jitter becomes too high, bit errors or frame errors can occur, leading to lowered throughput.

At the higher layers, it is useful to examine the type of traffic that is going across a WAN. If you are using NetWare protocols, for example, you may be sending broadcast traffic (SAPs) advertising services such as file and print servers across the WAN. Proper router configuration generally can filter out unwanted traffic.

A complex source of problems comes from the encapsulation of protocols as a packet moves from the LAN through the WAN. For instance, you may be encapsulating AppleTalk in IP in order to route it, and sending that over frame relay. This is known as *tunneling*. In order to troubleshoot this, you need a high-quality WAN analyzer capable of decoding LAN-over-WAN. You also will need some detailed documentation on your protocols to understand what they should look like when encapsulated. The normal procedure when a problem is found here is to consult your router vendor.

If you have baselines for your network, these will guide you in deciding whether you are experiencing problems with the network. There is no substitute for information on the normal statistical operating envelope of your network when faced with an apparent abnormality.

### 16.4.5 Intermittent problems

Intermittent problems are the hardest ones to find. A common technique to isolate problems like this requires an analyzer that can trigger and perform an action based on a network event. The event usually is a network error condition, occurrence of a certain type of frame, or a network statistic (such as broadcasts) reaching a certain level.

To perform the test, set up a reasonably large circular buffer in the analyzer that continually captures all packets, and wraps around when it becomes full. (This is a common capability in analyzers.) Next, set up a trigger that is based on an event that defines the problem (such as a high broadcast level). Configure the trigger so that when the event occurs the trigger will stop the packet capture. At this point the problem event and all the traffic leading up to it will be in the capture buffer. This then can be analyzed using protocol analysis to determine what is causing the problem. It is best if the analyzer can capture at full bandwidth, because some intermittent problems (such as broadcast storms) can be caused by only one bad packet!

## 16.5 Network Monitoring

Sources of network trouble have changed over time. In the beginning physical problems were dominant. The cabling was unreliable. The fault domains were limited only by the repeaters. If a cable problem arose from someone kicking a wire under a desk, the entire LAN was affected. The first troubleshooting tools were wire testers that checked only for continuity. Today's structured wiring environment has limited the cable fault domain to a single node. Cable testers still play an important role today, checking critical parameters such as distance and crosstalk, which can adversely effect today's high-speed networks.

As networks grew more complicated from a protocol perspective, the source of problems migrated from the cables up the stack. Physically the networks are composed of cables and interconnects, and many of today's problems originate at the interconnect devices. Encapsulation and routing of diverse protocols has created a complex environment that requires a protocol analyzer for serious troubleshooting. Specific problems can occur between "compatible" interconnects. Standards implementations can have varying interpretations, and it is not unusual to need a router software patch to fix a problem.

Network problems can go undetected for some time if not checked for. This happens because the networks, at least in the local area, historically have had excess capacity. This is changing and adverse performance from a suboptimum network is a problem today. The only way to understand the health of the network is to measure it. Generally this means collecting more than just throughput statistics from an interconnect.

### 16.5.1 Preventing problems

Network monitoring on LANs and WANs has become standard procedure for large networks. From an economic point of view, the question becomes one of who is using the network bandwidth and for what. From a more practical point of view, net-

work monitoring can help manage the performance and the health of the network. A simple example of this is tracking traffic levels over time to determine long-term trends. Using this information, you can predict when your network traffic will surpass the packet forwarding rate of your bridge or router and plan appropriately. Network monitoring systems based on standards like RMON can be configured to provide alarms based on error conditions, providing the capability for management by exception.

### 16.5.2 Planning for growth

One of the first questions to occur when setting up a monitoring system is, “How do I know what is normal for my network?” There are some rules of thumb, but the most effective method is to utilize a long-term monitoring program and create a baseline from which to operate. Typical items to baseline include error rates, traffic by protocol type, packet size distribution and overall traffic levels. Other items, such as the ratio of collisions or broadcasts to overall traffic also can be useful, as can the WAN bit error rate as a function of traffic levels.

In order to create a baseline, a month’s worth of data should be collected. Keep in mind that network traffic has some specific drivers; in the case of a LAN it is the working hours of the employees, or the time when backups occur. A baseline should take note of those items, to prevent comparing the traffic at 4:00 AM, when nobody is using the network, to the peak hours of the week. Having baselines is an effective way to plan for the growth of the network.



---

Part

**5**

# Cellular Networks



# Introduction to Cellular Radio Networks

**Tom Walls**

*Hewlett-Packard Ltd., South Queensferry, Scotland*

## 17.1 Introduction to Cellular Systems

Although cellular radio is a relative newcomer to mobile communications, the concept of a cellular system dates back to the early 1940s, and the principle of a mobile phone system goes back still earlier. The first mobile phone system was an experimental system installed in 1921 by the Detroit police department. It used transmission frequencies at 2 MHz band and was capable only of one-way transmission. The called police officer had to use the public telephone network to answer the call. But it was the predecessor of today's mobile phone networks.

The cellular concept was conceived in the 1940s. Planners imagined a system of transmitters that could cover certain geographic areas and provide very efficient communications for a large number of users. Unfortunately, however, the technology required to realize such a system wouldn't exist for almost another 40 years. The objective of providing two-way communication became more realizable as the basic concepts were moved upwards in frequency, and the development of VLSI made it possible to reduce the physical size and power consumption of the devices, thus opening up new opportunities in the deployment of cellular radio systems.

Today's mobile telephone systems began in earnest with the Mobile Telephone System (MTS) in 1964. It operated at 150 MHz band and its major advancement was automatic channel selection. It was followed five years later by the Improved Mobile Telephone System (IMTS) at 450 MHz. This system set standards for current mobile phone systems.

The first truly cellular radio system was introduced in the Scandinavian countries (Norway, Denmark, Sweden, and Finland) in 1979. This was the Nordic Mobile Telephone (NMT) system. NMT was followed three years later by the Advanced Mobile Phone System (AMPS) in the US and Canada, followed a year later by the Total Access

Communications (TACS) system in the United Kingdom. The current systems being installed are based on digital technology such as GSM, NADC and CDMA, as described in the sections following.

## 17.2 Cellular Radio Concepts

Cellular radio is based on the concept of *frequency reuse*, in which available channels are assigned in groups to each of many different locations. This frequency reuse allows a much higher subscriber density per megahertz of spectrum than previous systems. These locations, actually geographic regions, are known as *cells*.

The need for frequency reuse stems from the nature of the early transmitter systems, which were very powerful, meaning that the chosen frequencies could not be reused for a radius of several miles. This led to major limitations on the capacity of the systems; once a particular frequency was in use, the channel was busy for the entire coverage area of the cell, even when the requirement for a usable channel was confined to a small portion of the total coverage area.

The reduction in transmitter power, and hence the reduction in cell size, created the environment for low-power transmitters, called *base stations*, specifically designed to cover only that area. The base station transmitter is connected to the mobile network's telephone exchange (MTX). The MTX then is connected to the local telephone exchange to gain access to phones worldwide, as shown in Figure 17.1.

A mobile radio gains access to the cellular system through the base stations. That call is then routed by the MTX to standard telephone lines or to another mobile. The link from the base station to the MTX can be either land lines or a microwave link.

Depending on the terrain, the antennas may be omnidirectional, bidirectional, or focused beams. In many parts of the United States, for example, omnidirectional antennas are used because the terrain is flat. In Hong Kong, however, where there are skyscrapers on almost every street, focused beam antennas are used to gain maximum coverage. Some cells are as small as 500 meters square.

When a user reaches the fringe of a cell and the base station starts to lose the signal, the handset begins to look for an adjacent cell to which to transfer the call. The

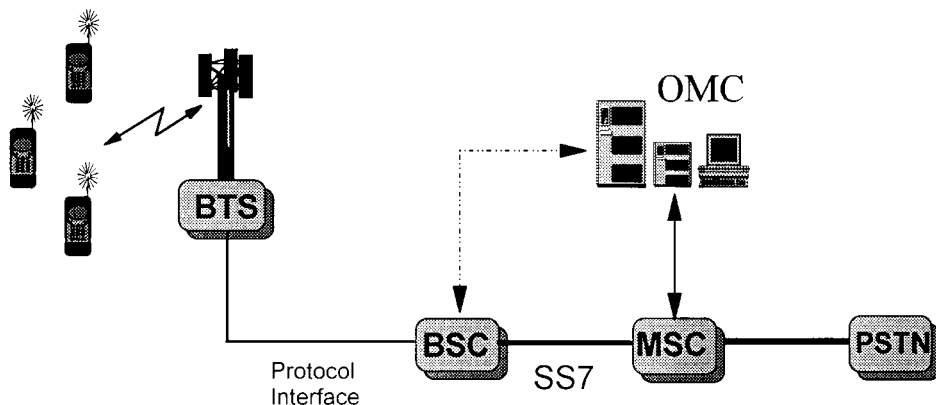


Figure 17.1 Typical digital cellular network topology.



signaling information carried via the protocol layers assists in this decision-making process by providing a conduit for the handset measurement results. Meanwhile, the traffic or voice channel is maintained by using a different section of the protocol channel.

This process is called a *hand-off*, and decisions are made in real time on measurements performed by the mobiles and network elements. If an adjacent cell is found, the call is transferred without losing the conversation. In this way, a car and driver can travel anywhere in the coverage zone and continue a conversation. This is why cellular systems have become popular. Hand-offs have been made possible by the increase in network intelligence. Mobile stations are now able to upload and download information to or from the network to enable mobile tracking.

Today, customer expectations of these networks include the ability to make and maintain a call anywhere within the subscribed operator's coverage area. The network operators market their services in this emerging market as extensions to the fixed line system. In addition, they also offer mobile data services, primarily for the business user. These service expectations and provisions force the operators to provide inherent network abilities such as call continuity, coverage, and speech quality at an acceptable cost.

Cellular systems have come a long way since 1921. How they are realized and what challenges face the operators and equipment providers, from a test and measurement perspective, will be discussed in the remainder of this chapter.

### 17.3 Cellular Network Technology

There are two major technology issues that present a major challenge for cellular systems, *air interface* (radio transmission) and *mobility management*. The world cellular markets are served by the following air interface technologies:

- **CDMA** Code Division Multiple Access
- **GSM** Global System for Mobile communications (formerly Groupe Speciale Mobile, now SMG)
- **PDC** Personal Digital Cellular (JDC, Japanese personal communication system, 800/1500 MHz)
- **AMPS** Advanced Mobile Phone System, U.S. cellular standard
- **TACS** Total Access Cellular System, U.K. analog cellular standard
- **NMT** Nordic Mobile Telephone system, Scandinavian analog cellular standard for 450 and 900 MHz
- **NADC** North American Digital Cellular

One might assume from the preceding list that cellular technologies are found in both analog and digital form. Table 17.1 summarizes the analog cellular standards, and Table 17.2 the digital.

*Time Division Multiple Access* (TDMA) and *Code Division Multiple Access* (CDMA) technologies have developed as alternatives to *Frequency Division Multiple*

TABLE 17.1 Analog Cellular Systems.

	AMPS	TACS (NTACS/ETACS)	NMT-450	NMT-900
Principal Geography	North America	Europe	Europe	Europe
Introduction	1982	1983	1979	1985
Frequency Range	896–894 MHz down 824–849 MHz up	860–870/916–949 MHz down 915–925/871–904 MHz up	463–468 MHz down 453–458 MHz up	935–960 MHz down 890–915 MHz up
Data Structure	FDMA	FDMA	FDMA	FDMA
Channel Spacing	30 kHz	25 kHz	25 kHz	12.5 kHz
Number of Channels	832	400/1240	200	1999




Figure 17.2 Cellular access methods.

*Access* (FDMA), as shown in Figure 17.2. With TDMA, the usage of each radio channel is partitioned into multiple timeslots, and each user is assigned a specific frequency/timeslot combination. Thus, only a single mobile in a given cell is using a given frequency at a particular time. With CDMA (direct sequence spreading), a frequency channel is used simultaneously by multiple mobiles in a given cell, and the signals are distinguished by spreading them with different codes. One obvious advantage of TDMA and CDMA is the sharing of the radio hardware in the base stations among multiple users. Figure 17.3 illustrates the more efficient frequency reuse of CDMA.

Each of these networks has been built, deployed and tested in line with a specific standard that is essential in guaranteeing interoperability between different vendors' equipment. In addition to these standards, the network operator and equipment

TABLE 17.2 Digital Cellular Systems.

	GSM900	DCS1800	NADC	PDC	CDMA	PCS
Principal Geography	Europe	Europe	North America	Japan	North America, Japan	North America
Introduction	1992	1993	1992	1993–1994	1995–1996	1996–1997
Frequency Range	925–960 MHz down 880–915 MHz up	1710–1785 MHz down 1805–1880 MHz up	869–894 MHz down 824–849 MHz up	810–826 MHz down 940–956 MHz up 1777–1801 MHz down 1429–1453 MHz up	824–849 MHz (US) 869–894 MHz (US) 832–834, 843–846, 860–870 MHz (Japan) 887–889, 898–901, 915–925 MHz (Japan)	1930–1990 MHz down 1850–1910 MHz up
Data Structure	TDMA	TDMA	TDMA	TDMA	CDMA	Multiple technologies, including
Channels per Frequency	8	8	3–6	3–6	32–64 (dynamic adapt)	<ul style="list-style-type: none"> <li>• PCS TDMA</li> <li>• PCS CDMA</li> <li>• PCS 1900</li> <li>• Wideband CDMA</li> </ul>
Modulation	0.3 GMSK (1 bit/symbol)	0.3 GMSK (1 bit/symbol)	$\pi/4$ DQPSK (2 bits/symbol) $\alpha = 0.35$	$\pi/4$ DQPSK (2 bits/symbol) $\alpha = 0.5$	Mobile: QPSK Base: OQPSK (1 bit/symbol)	
Speech CODEC	REL-P-LTP 13 Kbits/s	REL-P-LTP 13 Kbits/s	VSELP 8 Kbits/s EFR	VSELP 8 Kbits/s	8 Kbits/s var rate CELP 13 kbit/s var rate CELP	
Mobile Output Power	3.7 mW to 8W	250 mW to 2W	2.2 mW to 6W	0.3 W to 3 W	10 nW to 1 W	
Modulation Data Rate	270.833 kbps (1 bit/symbol)	270.833 kbps	48.6 kbps (2 bits/symbol)	42 kbps (2 bits/symbol)	9.6/14.4 kbps data, 1.2288 Mbps spreading	
Filter	0.3 Gaussian	0.3 Gaussian	SQRT raised cosine $\alpha = .35$	SQRT raised cosine $\alpha = .50$	615 kHz Chebychev low pass (FIR)	
Channel Spacing	200 kHz	200 kHz	30 kHz	50 kHz 25 kHz interleave	1.23 MHz	
Number of Channels	124 frequency channels 8 timeslots per channel (1000)	124 frequency channels 9 timeslots per channel (3000)	932 frequency channels w/3 users per channel (2496)	1600 frequency channels w/3 users per channel (4800)	19–20 frequencies	
Source	ETSI GSM Standard	ETSI/GSM Standard	IS-54/135	RCR Spec Std 27B	IS-95	




Figure 17.3 Cellular frequency reuse patterns.

providers must design and deliver equipment that will adhere to regulatory standards on spurious frequency emissions, safety, and telecommunication network compliance. This ensures successful interworking with other systems on the assigned frequency bands, and with other network systems.

Cellular networks differ from fixed-line networks in several areas. The air interface is the most obvious area of difference, followed closely by the mobility management that enables the network to know where a mobile station is currently located, and to track it irrespective of whether the mobile is making a call or in idle mode. The air interface design or Physical layer requires careful choices of analog or digital, modulation format, frequency channel selection, and type of code.

### 17.3.1 Air interface

The RF environment provides many challenges for the system designer and planners. *Multipath propagation effects* add a dimension of system design complexity not encountered by the fixed-line operator. The fixed-line environment is much easier to control since the signals are carried on point-to-point cabling. RF signals, however, are reflected from buildings, particularly in the urban environment. This creates multiple, uncontrolled paths from transmitter to receiver. (Before the advent of CATV, “ghosting” on TV sets in urban areas was an extremely common manifestation of multipath.)

Multipath fading and other factors make modeling the transmission medium extremely difficult and introduces inaccuracies in the RF planning phase. Before a cellular network is deployed, a considerable level of planning is required to ensure that adequate coverage is provided. The choices of air interface technology are many; these system design tradeoffs (such as speech and data service requirements, data bandwidths, operating environment, fading performance) are beyond the scope of this book.

The choice of analog or digital system is a basic tradeoff in system capacity and efficiency. The analog systems use techniques such as FDMA (FSK) for signaling and FM for speech (Table 17.1). In an analog system, hand-over decisions usually are

based on received signal strength at the base stations surrounding the mobile. The development of low-rate digital speech coding techniques and the continuing penetration of VLSI have made completely digital systems viable.

Digital systems can support more users per base station per megahertz of spectrum, allowing wireless system operators to provide service in high-density areas more economically. Digital architectures such as TDMA and CDMA provide a more natural integration with the current digital wireline network, enabling mixed voice and data applications. The digital architecture also provides potential for further capacity as reduced-rate speech codecs and encryption for communication privacy are introduced (Table 17.2).

### 17.3.2 Mobility management

Mobility management requires an intelligent network, complete with a set of protocols that can pass information about the mobile station around the network to ensure traceability as it moves from cell to cell. This requires a hierarchy of network protocols to ensure that messages are passed efficiently and effectively around the network.

Mobility management is a sublayer in the protocol and handles the tasks that are specific to the mobile network, such as:

- Verifying the user and equipment identity
- User security
- User confidentiality
- Proper service provision

Procedures are defined to ensure that these tasks can be performed in the network; e.g., location updating and authentication must occur periodically.

## 17.4 Summary

It is generally accepted that the air interface in the cellular system is more unpredictable than any other aspect of the systems design. In this chaotic environment, the radio engineer must design and test the system to ensure the air-interface standards compliance. The differing modulation, coding, and wide bandwidth signals make the new systems difficult to test and verify. Test equipment makes a crucial contribution to the overall verification of the system performance, and this will be the focus of the remainder of this part of the book.



---

Chapter  
**18**

# Cellular Measurement Strategies

**Bob Irvine**  
**Gordon Innes**

*Hewlett-Packard, Ltd., South Queensferry, Scotland*

## 18.1 Cellular Network Life Cycle

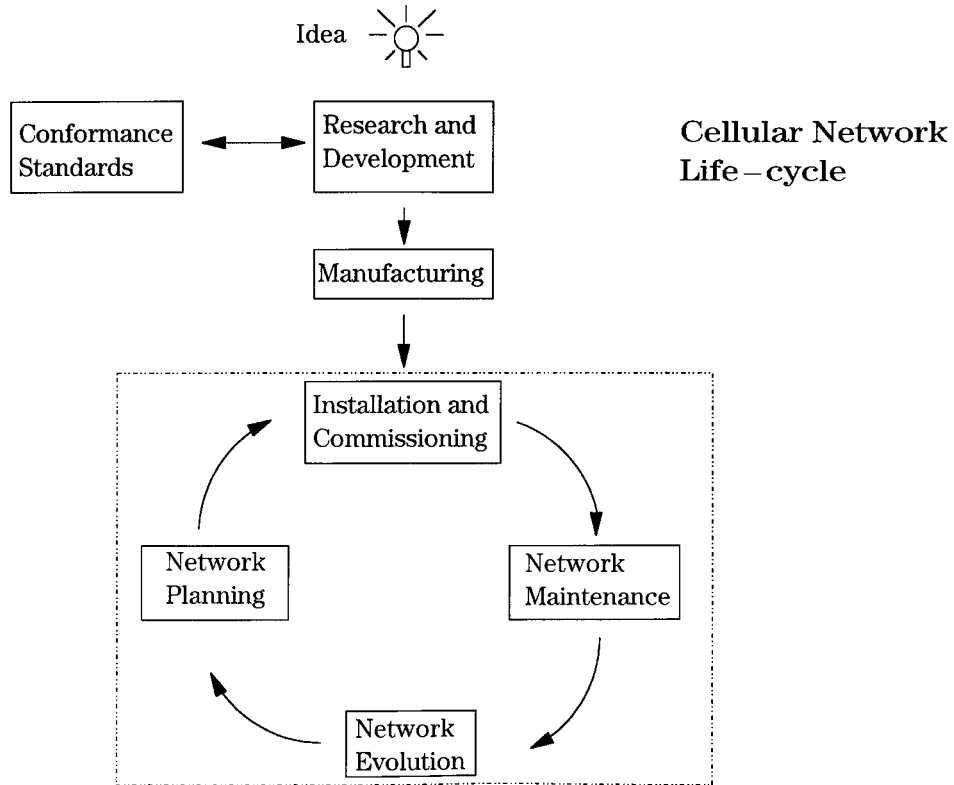
Like any product, a cellular network has a life cycle that includes:

- Research and development
- Manufacturing
- Installation and commissioning
- Maintenance
- Evolution and planning for improvement

In addition to this traditional life cycle, however, networks are based on international, government, and industry standards that dictate performance. These form the foundation for any network specification. As a result, defining and understanding the applicable standards before establishing the network performance expectations, and therefore the network specification, is also a step in the overall development and deployment cycle. These standards will establish the tests and measurements required to ensure that a new network or network element will function as part of a greater network operation. Figure 18.1 shows this total life cycle for a cellular network.

## 18.2 Research and Development, and Type Approval, and Confidence Testing

Test equipment is required initially in the research and design laboratories. The equipment usually is very generic and flexible. Test accuracy is crucial to ensure design margins in the system. Very often the R&D engineer will verify performance on several pieces of equipment to provide confidence that the results are repeatable and consistent. Speed of testing is not an issue because the early tests will almost always be car-



**Figure 18.1** Cellular network life cycle.

ried out manually. The user interface need not be highly polished; it is more important to have appropriate functionality available than a simple and intuitive interface.

The focus during the system's design phase is on verifying component performance and ensuring that the chosen parts are meeting or exceeding the system specifications across the extremes in environmental conditions. These specifications are updated continually during the design phase, and test equipment performance must provide an order-of-magnitude accuracy margin in the hardware domain.

Once these specifications have been verified, any anomalies or unachievable specifications must be reported to the system designers to calculate the effects on the overall system. In Europe many of these standards are discussed at ETSI meetings; in the USA they are discussed in TIA meetings. Any changes will be reflected in the base standards, with test standards following. Test methods also are verified to ensure accuracy and repeatability across acceptance test labs.

Software designers also need test equipment to guarantee that their measurement algorithms and protocols are operating as predicted during the initial system modeling. Protocol testing requires powerful tracing tools that respond in real time to hardware triggers. The various protocol layers must be tested to ensure that a high-integrity communication channel is provided end-to-end.



Hardware designers also are faced with many challenges in handset design for the consumer market. The size, cost, and battery lifetime are a few of the additional constraints that make it very difficult to achieve the system performance requirements. Handsets must be designed to withstand extremes of temperature yet provide consistent performance.

Most cellular network systems have been driven forward by the combination of two forces: the development of *formal, standardized Type Approval* processes to regulate the equipment that will be deployed in a network, and the *research and product development* activities that result in actual implementations.

Commercial pressure to develop a system rapidly requires overlap in these activities.

### 18.2.1 Making the standards

The goal of a Type Approval process is to ensure the successful interoperability of equipment within a network. Although this is often associated with the support of multivendor networks, the requirement also holds for the multiple elements of a single-vendor solution.

Type Approval development for an open system is by its very nature a cooperative affair. To be successful it requires a balance of inputs from a variety of sources, and a strong controlling body to provide the framework. An academic, or *core research description* of the basic properties of the system is required, along with a *solid architectural design* that will maintain the consistency and cohesion of the complete network. Once this is in place, much of the detailed specification is pushed forward by representatives of four groups: network operators, network equipment manufacturers, approvals agencies, and test equipment manufacturers.

The dynamics of such a diverse group working together can at times produce some unusual alliances between competitors. It also occasionally results in apparently bizarre, contradictory, or overlapping requirements that will satisfy the needs, and gain approval, of all interested parties. The initial research may require intensive, high-specification test equipment and suitably flexible test beds with which ideas can be subjected to rigorous test.

In general, however, the most significant requirement at this stage is for powerful computer simulation and modeling facilities. These can be used to explore the operation of the complete network at various abstraction levels, or to investigate specific subcomponents. The most common requirement is for effectively modeling the air interface signal propagation behavior in complex, real-world environments. It is also important to simulate the operation of the network protocols that will provide support for the system's basic functional capabilities. This includes the ability to track and route calls and other services, to and from the end users. It is at the point when the specific functional requirements of each system element have been defined that the Type Approval test requirements can be generated.

**Type Approval Test.** These normally can be classified in two groups:

- *Protocol tests* verify the correct operation of the software components.
- *Physical tests* ensure compliance to a minimum acceptable set of physical and environmental characteristics.

Depending on the policy of the authority promoting the system, the defined tests will provide input into one of the following scenarios. It could be decided to commission one (or more) groups to develop a Type Approval Test System. This will be used to provide definitive arbitration of whether a particular piece of equipment can be certified as compliant to the required standards. Alternatively, it could be left to individual organizations (governmental, system test houses, network operators, or even equipment manufacturers) to decide how the Type Approval requirements should be verified. In practice, elements of each usually can be found in any particular approvals process.

### 18.2.2 Testing to the standards

The test equipment required to verify protocol operation has similar needs to that required during development of the protocol software modules. It should provide the capability to generate, monitor, and respond appropriately to various protocol sequences. This gives a flexible test bed environment that facilitates expression of the expected, and allowable, protocol exchanges. This should be in a form similar to that used within the specifications.

It also should be amenable to change, since the requirements of this part of the system often are the least stable. Because it covers the software functionality, it also is the portion most subject to change as further development of a system takes place, and as new capabilities are introduced. It often is useful for the test equipment to be programmable using a number of different notations. This allows for differences in the way various parts of a standard might be specified, and permits the test functions themselves to be verified using alternative implementations.

Physical testing can be split into two parts, *functional verification* and *environmental stability*. Similarly, the test equipment requirements can be split into two parts. Equipment is required that will observe the performance of the unit to be tested. Usually the system designers will have identified which particular characteristics of an element are most critical for any system. Most tests will be focused on ensuring that operation is happening correctly. In a GSM network, for example, the RF power burst (which carries the digital information across the air interface) is subject to many test variations for both base station and mobile phone elements.

The second requirement of the test equipment is the ability to generate the specified test environments. This varies from the relatively simple temperature and vibration/shock variations, through the introduction of interfering signals, and in some cases simulation of the complex interactions and perturbations of the air interface signals due to fading and Doppler effects.

When selecting test equipment for these purposes, it is important to pay special attention to its specified capabilities. The performance of the unit under test is what must be measured, not the characteristics of the test equipment. Accurate error estimates at this point are particularly important when testing to specific absolute performance criteria. To reduce the range of error bounds, it is vital to use equipment that has tightly controlled and well-defined capabilities.

In general, the requirement for a Type Approval system for a particular element is to provide a functional capability that mimics the requirements defined in the stan-

standard for the other elements with which the tested element communicates. For some functions this can require simulating the operation of large portions of the network, not just the limited functionality of the immediate neighbor elements in the structure.

The result of the development and Type Approval activities will be a set of designs. When fitted together correctly, these designs will provide network elements capable of communicating successfully with other approved elements.

### 18.3 Manufacturing

Once the initial system verification and handset and base station testing is complete, the next hurdle is high-volume production. Is it possible to mass-produce the design in an economic and repeatable manner? Testing mobile radios on the production line is in many ways similar to any other high-volume production process. The process must ensure that the tests are focused on the critical few parameters that guarantee the integrity of the product.

At the outset it is very difficult to know these parameters until a significant quantity has been manufactured. The manufacturing engineers must rely on the R&D team to advise on critical parameters; usually the manufacturing test engineer will be an integral part of the design team, making knowledge transfer efficient. Repeatability and accuracy are important characteristics of the test equipment, and measurement reporting usually is performed over a high-speed computer interface. This allows the manufacturing test supervisor to collate large amounts of data and watch for component variations that could increase test time.

The test process will consist of component, subassembly, and final assembly test bays; at each stage the test requirements will be different. Typically the manufacturer will request a generic set of functionality to ensure commonality across the process, keeping test technician training to a minimum. Final test should always consist of making a phone call either in a special test mode or in loopback mode. Since the end user will be concerned about audio quality, this requires a special test station to verify the lack of resonance and audio leakage. Special test jigs also are required when the tests are exercised over the air interface and coupling of the signal to the antenna is the only means of communicating to the device under test (DUT).

Faulty components must be detected early in the manufacturing process to avoid wasting valuable time at a later stage. Often the DUT requires on-board adjustments to meet specification, usually power or frequency. With modern designs many of these parameters can be adjusted using an on-board D/A converter and computer control. The adjustment value is often calculated in the test system controller and then downloaded over a test protocol link to the DUT. Speed of test is crucial if the manufacturer is to maximize the throughput of the test line, so design of the test line is important for performance optimization. Often manufacturers will start in a new technology with a nonoptimized line and then redesign when significant experience has been gained.

#### 18.3.1 Objectives of testing in manufacturing

Testing during the manufacture of cellular handsets or mobiles and base stations (BS) is very similar for both products and can be split broadly into two categories,

*calibration* and *manufacturing process control*, which are linked throughout the test process. Calibration is needed to align the complicated radio frequency (RF) circuitry to take account of component tolerance variances. Manufacturing process control is needed to ensure that consistent quality of product is being produced, quality that will meet the design specifications.

In a competitive marketplace, reducing the cost of test is a real motivator for manufacturers. The cost of testing is usually dependent on the time it takes, and many production lines have a high degree of automation. Test instrumentation is generally controlled by a networked PC so that control programs and measurement data can be distributed and collected easily. The key to an efficient production strategy is to perform just enough testing for calibration and to test only parameters that vary. The goal of the testing process is to give a high degree of confidence that any mobile or BS selected at random from the production line would pass a full Type Approval test.

### 18.3.2 Manufacturing test elements

In cellular mobile and BS, manufacturing testing usually focuses on the RF parameters because the digital portions of circuits tend to work either completely or not at all. It is usual to store measurement samples from some (if not all) products to build up a picture of how the manufacturing process varies over time. This allows production engineers to take corrective action long before problems have started to affect the final product quality. In general, the earlier in the test process that a problem is found, the less cost is incurred to correct it. Finding a problem at the raw circuit board stage of an assembly will cost much less to correct than finding that the fully assembled phone has a faulty component.

The manufacturing process for most modern mobile and BS equipment tends not to be fixed at the start of production, but evolves continuously throughout the product's life, becoming more efficient. It begins as a lengthy procedure with perhaps an excessive amount of testing. As experience and confidence in the product build up, test points or complete tests may be removed (or inserted) to refine and mold the process, testing only the critical areas where a high degree of variability occurs or extensive calibration is required. Usually a sample of products is subjected to a fuller test on a regular basis to ensure that the overall testing strategy is not flawed.

The initial design of a mobile or base station plays an important part in its "testability," and there is increasing pressure to design products with testing and calibration in mind. This will include adding special test modes in which the products can be controlled by an external computer that commands them into known conditions for measurement, and electrically erasable/programmable read-only memories (EEPROM) to hold on-board calibration data. These data will be used as a lookup table when the final product is in use.

Calibration takes up a significant part of the testing performed in manufacturing. Even with careful design, there are tradeoffs to be made. For mobiles in particular, there is a desire to keep the manufacturing cost low; this can mean that lower-cost components with wider tolerances force a certain degree of calibration. Testing cost has to be weighed against using more expensive components with tighter tolerances. As a general rule of thumb, the more expensive the final product, the more testing

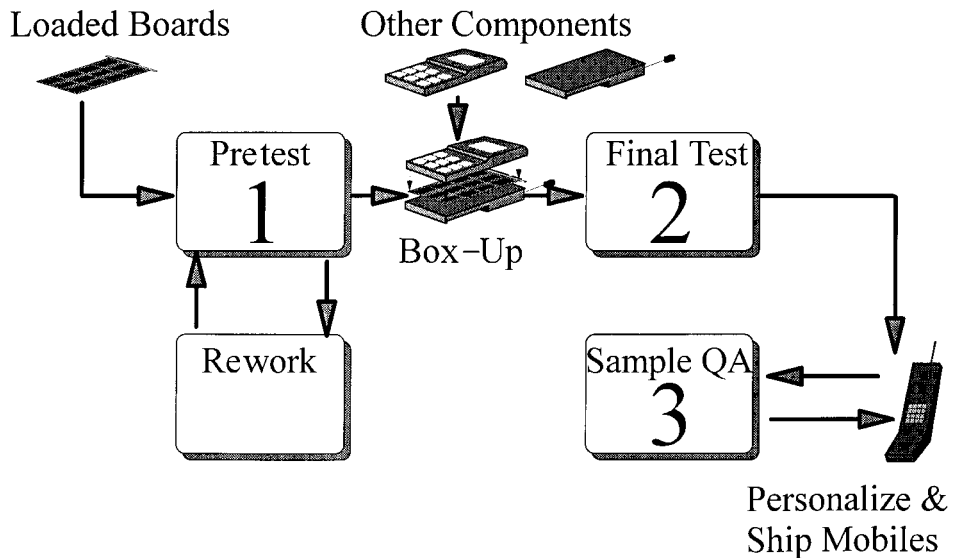
can be carried out in the production process without adversely affecting the end-user price.

In addition to calibration, product performance optimization must be considered. It might be desirable to set some performance criteria at the edge of their tolerance bands to allow gains in other areas. The decisions are generally made in the R&D labs, but are implemented in the manufacturing process. For example, one tradeoff would be the output power accuracy of a mobile versus its battery life. Setting the power output close to the lower tolerance limit would allow the mobile to operate longer with the same battery than if the power were set to the middle of the spec. The more accurate the test equipment, the closer the limits can be approached with confidence that the mobile will be within specification.

### 18.3.3 Manufacturing process flow

The manufacturing process flow can be split broadly into three sections, *pretest*, *final test*, and *sample quality assurance (QA)*, which are shown in Figure 18.2. The type of testing performed and the equipment used differs at each stage of the process. Because both mobile and BS are essentially different-sized packages containing similar circuitry for transmission and reception, many of the tests and processes are similar. We will now examine each stage of the manufacturing process in turn.

**Pretest.** The first stage in the process, after the circuit boards are loaded with components, is the pretest. This is performed to check basic circuitry operation and try



**Figure 18.2 Manufacturing Flow.** Raw circuit boards enter at Pretest and faults are repaired at Rework. Then the circuit boards are assembled into the case and move to Final Test, where they are calibrated and performance is verified. Finally, some of them may be retested in Sample Quality Assurance before being packed and shipped to customers.

to find severe defects as early as possible in the test process. Faults like completely nonfunctional integrated circuits and missing components ideally should be identified at the pretest stage, where they can be replaced most easily.

A certain amount of calibration may take place at pretest, though usually most of this is done at final test due to the influence of external coverings on sensitive RF circuitry. Pretest usually is fully automated, with a manual rework loop to correct any identified faults. There are two basic strategies for pretest, illustrated in Figure 18.3. Mobiles generally are manufactured using either strategy, but base stations tend to be manufactured using strategy 2.

**Pretest Strategy 1.** This is a top-down approach, which starts from the assumption that the circuit board is likely to be working. An integrated cellular test set is the ideal choice of test equipment for this strategy. An attempt is made to establish a call with the basic mobile circuit board and measurements are made with wide limits. The test development time with this approach is quick because much circuit board functionality is implied by the fact that it can operate well enough to establish a call.

The type of measurements carried out on the transmitter module are carrier power, modulation quality, and (in digital TDMA systems) power versus time. The receiver sensitivity or bit error ratio in digital mobiles also may be checked. During testing the mobile is controlled using over-the-air signaling from the test set. Measurements are made either at the mobile antenna connector or accessory connector. What is looked for in pretest is functionality and adjustability. This technique has the advantage that it is reasonably easy to implement, and a large portion of the circuitry is exercised and tested in the process of establishing a call.


**Pretest Strategy 2.** The second method of implementing pretest uses discrete test instruments such as spectrum analyzers, voltmeters, and signal generators.

### Cellular Test Set

- Faster Test Development
- Common Pre / Final Test Bays
- Consistent Test Methods

### Discrete Instruments

- Multi-format Pre-test for many different circuit board types.
- Longer test development time



**Figure 18.3 Pretest options.** A cellular test set offers a quick, easy way to check if a circuit is working. The fact that it operates gives a high degree of confidence that it will be able to be calibrated into a final-quality unit. Discrete instruments offer superior flexibility because there is no dedicated protocol built in, as there is in the cellular test set. The flexibility is offset by the fact that each section of the circuit board has to be tested individually, which slows the test development process.

These typically are set up in a 19-inch rack that also houses a power supply and a computer to automate the measurements. This is a bottom-up approach that assumes no circuit components to be working and checks each part of the circuit separately. Test development is usually slower, as there is a need to test each section individually.

RF signals and voltages are injected at various circuit test points and the responses measured using a spectrum analyzer and voltmeter. This process closely mirrors the way the circuits were designed on the R&D bench, only now the process is automated. This technique has the advantage that it allows manufacturing flexibility; many different boards can be tested with a basic set of equipment, simply by changing the fixturing and control software. This has to be traded off with the complexity of testing separately each part of the circuit board and not seeing how it all performs as a whole.

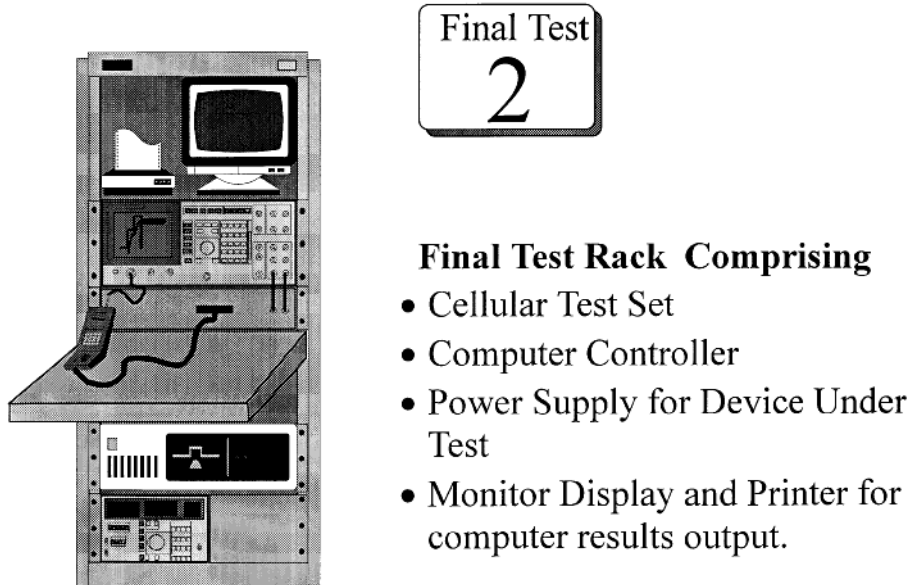
Because simulating a network of mobiles is difficult, this second method is more commonly used in base station pretest. The mobile or BS under test is controlled by an external computer connected to a special test bus. In the case of a BS, the test bus may be a restricted functional simulation of the real-life base station controller.

Once the product has passed the pretest stage, it moves on to final test. At final test, full calibration is performed and functionality is checked. Because RF components are prone to being affected by their immediate surroundings, the raw circuit boards are usually assembled into their final cases before final test commences.

**Final test.** The first part of final testing usually involves calibrating the mobile or BS transmitter. In general, the transmitter power output levels must be calibrated and the modulation quality must be checked. The method of measuring transmitter carrier power varies depending on the type of mobile or base station; several examples of this type of measurement are given in section 18.3.2. In addition to transmit power, the power-versus-time mask will be checked in TDMA systems. Modulation quality is affected by several factors, again depending on system type. Examples of this type of measurement include phase and frequency error for GSM, and SINAD for the analog systems. The receiver also must be aligned and checked for sensitivity. Modern digital receivers usually are tested with a bit error ratio test (BERT). A typical test station for cellular mobile final test is shown in Figure 18.4.

During the calibration stage of final test, the mobile or BS is controlled by a computer using a special test bus. Usually the only other connection to the test equipment is the antenna port. Once calibration is complete, a parametric functional test is used to verify correct operation. For a mobile, this will involve establishing a call with a cellular test set; for a base station, a special test mobile is used that is able to report information about the RF signaling and link parameters.

The functional test will check the transmitter's ability to change RF channel, vary output power, and (in TDMA systems) change time slot. The receiver operation will be checked at a low signaling level to verify adequate sensitivity. There is likely to be some degree of audio testing because that is a key function of the final product. It is interesting to note that there has been no mention of any protocol testing so far. This is because protocol problems are not amenable to correction during manufacturing; all protocol testing, whether for mobile or base, should have been carried out during the R&D and Type Approval design stages.



**Figure 18.4 Final test station.** The final test station is generally an equipment rack and, in the case of mobile test, will include: a cellular test set to establish a call with the mobile and make measurements, a computer to control the test station, a power supply to power the mobile, and display and printing devices for measurement results.

**Quality assurance.** When the final test is complete, theoretically the mobile or BS is ready for packing and shipping to the customer. Most manufacturers do perform a certain amount of QA testing on a portion of their products, however. The QA test process will normally be a more extensive version of the final functional test, perhaps taking more measurement points or testing on more channels. Because mobile network service providers often inspect newly manufactured mobiles or base stations, the manufacturer attempts to simulate this incoming inspection test to give a high degree of confidence that products will not be rejected by the customer. Faults found at the QA stage indicate serious defects in the manufacturing process and need immediate investigation to prevent the creation of further faulty units.

In many cases, failures at final test and sample QA arise from measurement problems rather than any real defect in the mobile or BS. In some instances, problems are caused by poor accuracy and repeatability in the test equipment being used. Even if manufacturing test software has been carefully designed to catch all product defects at pretest, it still is possible for poorly specified test equipment to fail good units later in the process.

### 18.3.4 Specification budgets

A *specification budget* is a listing of the uncertainties in each measurement a piece of test equipment can make, such as how accurately it can measure power or frequency. This usually is obtained from the instrument data sheet. If properly specified test equipment is used at each stage of the production process, the budget can be used to calculate testing pass/fail limits for each stage of the manufacturing process. These




limits directly affect the number of mobiles that pass the final test with a high degree of confidence that they will meet their design specifications. This is known as the process *yield*. Using a specification budget ensures that failures occurring at final test and sample QA do not result from measurement problems. The specification budget can be used during the design of a manufacturing line as a tool to help select test equipment.

Figure 18.5 illustrates the principles of a specification budget. Budgets sometimes are single-ended (as shown), often symmetrical (simply add a reflection of the diagram to the left of the page), and occasionally asymmetric.

The room temperature specification usually is the value determined by the regulatory authorities controlling the standard. In the case of GSM900 mobiles, one example would be transmitter output power, which for most levels has to be within  $\pm 3$  dB of the nominal value. This specification is driven by the ETSI GSM 05.05 (ETS 300 577) specification for the RF interface. The full environmental specification is generally looser than that for room temperature ( $\pm 4$  dB in our example) and is also found in the GSM 05.05 specification. The exact environmental conditions of temperature, humidity, and vibration are also specified in ETSI documents.

To pass Type Approval, the mobile must be capable of meeting both the room temperature and the full environmental specifications. If a particular manufacturer's environmental drift is greater than that anticipated by GSM 05.05, it will be necessary for the manufacturer to tighten the manufacturing room temperature specifications to ensure all mobiles, if tested, could meet the full environmental specification. The value for environmental drift will have been obtained from a sample of mobiles tested during the development phase.

The manufacturing test pass/fail limit is used to reject unacceptable mobiles during the manufacturing process. It is calculated from the room temperature specification by subtracting the measurement uncertainty. Measurement uncertainty is determined



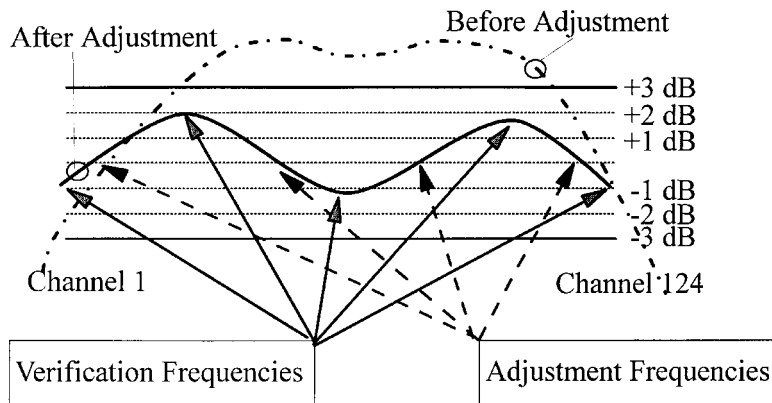
**Figure 18.5 Specification budget to set manufacturing test limit.** There will be a natural spread in the performance of manufactured products, caused by component tolerance variations. The measurement uncertainty affects where the manufacturing test line limit is set. It is important to guarantee passing the *full environmental* and *room temperature* specifications. A bigger measurement uncertainty means that for a given spread of product performance, fewer are guaranteed to pass.

from the test equipment specifications; it often depends on the specifications of test equipment being used at more than one stage in the manufacturing process.

The curve in the diagram indicates the proportion of mobiles tested with measured performance falling at any given value. The curve forms a probability distribution. It's often possible to approximate this to a normal distribution and predict yield from a relatively small sample of units. The yield is the proportion of units with performance within the manufacturing test limit. If the yield is unacceptably low, either the manufacturing limit has to be moved by using better test equipment with lower measurement uncertainties, or the distribution of the product's performance has to be improved. This can be achieved by improving the product's design, or by increasing the amount of adjustment and calibration to improve performance.

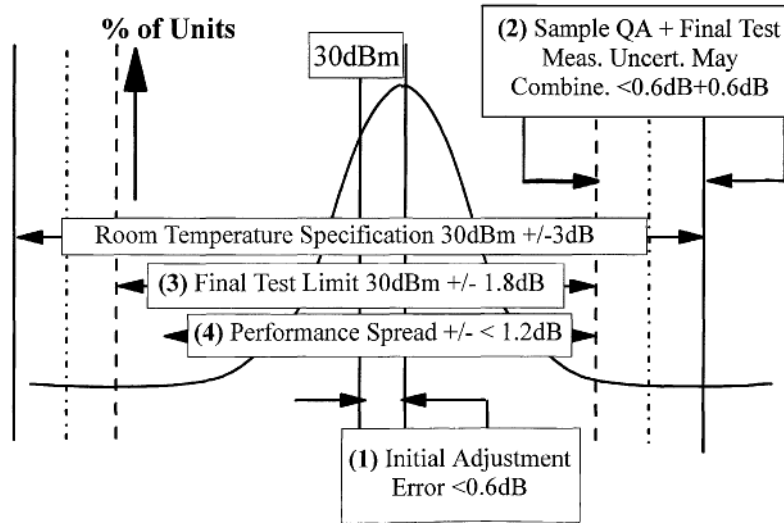
Most mobile phones have their output power adjusted to meet specification. (Refer to Figure 18.6, an example from GSM.) The mobile output power before adjustment might look like the upper rounded trace. The unadjusted mobile falls well outside the acceptable  $\pm 3$  dB limits. At each adjustment frequency, the output power is measured during final test, and correction coefficients are stored in EEPROM. The mobile firmware then uses the stored coefficients to interpolate between adjustment points.

At the adjustment points, the output power clearly will be very close to the correct nominal value, the only error coming from the repeatability of the mobile and the accuracy of the test equipment being used to make the adjustment. The performance will be closest to breaking outside the acceptable limits *between* adjustment points. When the performance of the mobile is being verified (during the second phase of the final test process), measurement points between adjustment points are chosen, making it more likely that the worst parts of the mobile's power performance will be exposed.



Adjustment values stored in mobile's EEPROM

**Figure 18.6 Mobile phone power calibration.** The upper dotted-line curve shows the mobile's output power performance on the Y axis, versus frequency channel along the X axis. It is clearly outside the  $\pm 3$  dB specification permitted in GSM. The solid curve trace shows the calibrated power performance curve. Notice that the verification frequencies (shown by solid arrows) are placed in between the adjustment frequencies (shown by dotted arrows). This is because the frequencies most likely to fail the specification are those furthest from the adjustment frequencies. The adjustment frequency measurements will be stored in the mobile's EEPROM to be used as a lookup table for output power.



**Figure 18.7 Specification budget applied to carrier power measurement.** This example illustrates how measurement uncertainties from several instruments have to be taken into account when setting final test pass/fail limits. Item (1) is the potential error introduced by the initial adjustment of the power level. Item (2) shows how errors from sequential measurement instruments may be combined. Item (3) is the final test pass/fail limit: the room temperature spec minus the sum of sample QA error and final test error. Finally, item (4) shows permitted mobile performance spread: pass/fail limit less initial adjustment error. Refer to the text for a full explanation.

**Combining specification budgets and measurements.** Figure 18.7 shows how a specification budget can be used for the mobile transmitter output power. The diagram is similar to the earlier specification budget shown in Figure 18.5, but shows the symmetrical budget appropriate for output power, which can be greater or less than the required specification.

Assume that the test equipment used has a power measurement accuracy specification of  $\pm 0.6$  dB. The normal distribution curve is for measured values of output power at the verification frequencies. The distribution is centered on the adjustment value, which could be in error by as much as 0.6 dB, which in this example is the accuracy of the test equipment being used for adjustment. The diagram also illustrates how measurement uncertainty from more than one test station has to be considered when setting the final test limits.

**Effects of measurement uncertainties from multiple test stations.** For the sample QA station to emulate a customer's incoming inspection test, it should use the  $\pm 3$  dB room temperature pass/fail limit. The manufacturing process must be designed so that 100 percent of units passing final test also pass sample QA. If the sample QA test set has a measurement accuracy of  $\pm 0.6$  dB, mobiles must be better than  $\pm 2.4$  dB to guarantee passing the  $\pm 3$  dB limit at this station. If the final test station also has a measurement accuracy of  $\pm 0.6$  dB, the final test limit there should be set at  $\pm 1.8$  dB due to the cumulative effect of the uncertainties at the two stations. Assuming an adjustment error, also of 0.6 dB, the allowed performance spread of mobiles is less than  $\pm 1.2$  dB.

Combining the adjustment and verification stages of final test into one station, and therefore one test set, can be used to gain measurement advantages. The absolute-level accuracy of test equipment nearly always will be poorer than the relative-level accuracy. If the same test set is used to make adjustments as to make verification measurements, the 0.6 dB adjustment error will not adversely affect the acceptable performance spread of units. The test line limit, and so the product quality, is not affected by this assumption. For a  $\pm 3$  dB room temperature specification, the acceptable performance spread would become  $\pm 1.8$  dB. This emphasizes the value of combining adjustment and verification at final test. If adjustments were performed at the pretest station, this assumption would not have been valid.

An important observation is the relationship between measurement uncertainty and throughput. If measurement uncertainty can be reduced, the acceptable performance spread of units for a given yield can be allowed to increase. This in turn will lead to a reduction in the number of adjustment and verification points, saving test time and therefore cost.

### 18.3.5 Manufacturing test summary

Manufacturing test is about process control and calibration. Being selective in choosing what to test and what equipment to test it with is the key to a streamlined process. As technologies have advanced, the techniques used by mobiles and base stations to transfer information across the RF air interface have changed, as have the measurements.

We have moved from analog FM-based systems through to pulsed digital encoded systems, and multiformat capabilities. Manufacturers are looking to the future as they invest in test equipment. Flexibility is becoming key, particularly in cellular test sets that so far have tended to be dedicated to one technology. Cellular test sets support multiple technology measurements in one box, with ever-increasing speed and accuracy.

## 18.4 Base Station Installation and Commissioning

Installation and commissioning of the system begins when reasonable numbers of the system components are available. Out-of-service testing is often possible as the infrastructure is installed since there are no customers expecting service.

The process of installing and making operational a new base station site usually is a complex and long-winded affair. It can be traced back to computer simulations of the RF coverage, and call density predictions for a particular part of the network. An ideal site description, in terms of antenna height and position, will be generated from the model to give optimal coverage. The search then will begin for a suitable site; this involves identifying land and building owners prepared to accept a site, or in some cases the applicability of existing sites. After further work to establish antenna height and placement acceptable to the local planning authority, the possible sites will be put back into the computer model, sometimes along with actual propagation measurements made at the sites.

#### 18.4.1 Radio frequency (RF) site survey

The equipment for such measurements will normally consist of a temporary transmitter placed at the site, and a mobile receiver used to take field strength measurements of the test transmissions and other potential interfering signals in the area. The transmissions can be either fixed-power, unmodulated signals, or sometimes a fully modulated simulation of a normal base station capability for the network type proposed.

The receiver will normally consist of a sensitive and accurate spectrum analysis system, along with a location tracking capability, that can be driven or carried around the projected coverage area. As well as measuring the intended transmission, it is important to quantify the level of any interfering signals, either from other network base stations or from outside sources. Where a site will be shared, or where there is other existing transmission equipment nearby, it also is necessary to measure the effect that the proposed new base station will have on the other installations.

#### 18.4.2 Site preparation

After a site has been selected and approved, all the equipment required will be installed: antenna mounted, power supplies connected, transceivers installed, and fixed or microwave network connections established. Each of these activities will require some (usually relatively simple) power-up and test. Often this will take the form of running self-test procedures on the equipment, or using simple standard tools such as oscilloscopes, power meters, and multimeters.

#### 18.4.3 Commissioning

The major testing will take place when the site is commissioned. At this point it is not unusual for acceptance tests to involve performing, at least partially, some of the critical Type Approval tests. This often will include providing a simulation of the controlling portion of the network. Since it is undesirable to have untested equipment put on trial using a live system, this simulator can be used to provide a well-controlled and repeatable test environment. Since the installation sites are often remote and unstaffed, the test equipment requirements differ significantly from the requirements of previous phases.

Installation is usually carried out by the equipment manufacturer; occasionally network operators choose to install base stations. Installation testing differs from the many other times a base station is tested in that it is performed only once. This often means that the BS has to be tested in isolation, before it is connected to the network.

When the test is being run by the equipment manufacturer, who already has reasonable confidence in the equipment's performance, the testing has two main objectives. First, a function check is necessary to make sure that the base station has been correctly installed; this generally requires little or no parametric testing. Second, sufficient parametric testing is necessary to ensure that the BS is capable of passing the network operator's acceptance test. In many cases installers carry out their own version of the customer's acceptance test, using identical equipment and procedures.

This gives the opportunity to rectify any problems or make adjustments before handing over the base station to the customer.

Because the base station is not yet in service, there is no need to make nonintrusive measurements. Because the BS usually is not connected to the network, it is necessary to find some way to duplicate any functions needed for the test process. There are two frequently used techniques for accomplishing these objectives.

**Functional Test.** The first technique uses a specialized tester. The functional tester is connected to both the BS protocol and RF interfaces. By taking control of the BS over the protocol interface, it emulates some of the BSC control functionality. By connecting to the RF interface the functional tester mimics some of the functions of a mobile handset. Very often these protocol interfaces contain proprietary messages for controlling the BS. The use of proprietary messages forces the functional tester to be extremely specialized.

Functional testers typically perform little or no parametric testing. What they do is provide a solution to the base station manufacturer's first test objective: verifying that the installation was performed properly.

**Parametric Test.** These testers typically do not address the second objective of verifying that the base station will pass the network operator's acceptance test. To meet this objective it usually is necessary to perform a variety of in-channel and out-of-channel transceiver tests. It often is necessary to tune combiners and balance the power level from several transmitters.

The suite of tests will not be as extensive as the R&D or manufacturing tests, but will focus on the critical few that verify correct operation. Accuracy is important because system margins are small and errors can lead to substandard operation.

The testing often will take two forms, simple performance checks and network operation checks. The first part involves measuring key parameters such as RF power level, current drain, broadcast frequency, and distortion or phase noise. The exact measurements made depend on the cellular system being deployed, and usually will be used to ensure that basic equipment features are performing adequately after transportation and installation. These tests can be performed with a dedicated system tester, which will provide the stimulus and control required for the base station to operate, and will measure the responses.

There is a trend towards less acceptance testing. As the reliability of the deployed equipment improves, and the on-site customization and tuning requirements decrease, there is less need to test. In addition, the need to rapidly roll out many micro-BSs increases the demand to simplify the test requirements. As an alternative, the base station may be activated, possibly in a restricted way, on the network. By using tap-ins to the generated signals, less sophisticated test equipment can be used. In particular, a spectrum analyzer or other receiver can be used without requiring any protocol capabilities.

**Coverage and Network Test.** Network operation tests do require the base station to be activated. At this point it is possible to perform "drive tests" of routes in the coverage area using a specially enhanced mobile. This mobile, sometimes in conjunction with in-network base station monitoring, logs data about the receiving performance and signal quality as it would be seen by a network user. Specifically, the mobile can be used to test actual signal strength against predictions, as well as the expected operation during hand-off from the new BS to others and back again.

Somewhat coincidentally, but nonetheless crucially, the operation of the whole system is tested in the new coverage area to ensure that calls can be established and maintained between the mobile and the network.

Microwave links are often used in the network as the backbone with which to link the remote BS sites to the central BSC. The amount of testing required on these links at installation varies a great deal. Low-capacity links often require little or no testing. High-capacity links, found deeper in the network, often require complex alignment and optimization of group delay and amplitude flatness. Low-capacity links often have built-in BERTs that provide a basic check on system operation.


## 18.5 Service, Incoming Inspection, and Repair

Once a cellular network has been installed and commissioned, there is the lifelong maintenance to consider. Both the handsets and base stations are subject to failures through general wear and tear; mobile handsets in particular are subject to user misuse such as being dropped or thrown around. For base stations, some parametric RF testing is carried out proactively as part of general maintenance to try to track down potential problems before they cause network outages.

A good service and repair strategy is a network operator's key to keeping the installed customer base loyal. In service and repair for both mobiles and base stations, we generally are fault-finding for a known set of faults. Faults tend to be repeated over time due to gradual wear-out of components or latent manufacturing defects. This means the repair process should become easier as time goes on; thus there is good justification for keeping accurate records of previously detected faults and solutions.

**Mobile incoming inspection.** Most network operators subject mobile handsets new from the manufacturers to an incoming inspection test before releasing them for use on the live network (Figure 18.8). This testing makes use of a cellular test

- ⊗ Usually carried out by network operator
- ⊗ Use cellular test set
  - Automated test as high number of mobiles to check
  - Consistency with manufacturing final test
- ⊗ Guarantee customer satisfaction
  - Ensures mobile will operate
  - Cuts the risk of network pollution



**Figure 18.8 Mobile incoming inspection summary.** The network operator checks mobiles for functional and RF performance before allowing them to be used on the network.

**TABLE 18.1 Mobile Incoming Inspection Testing: What and Why.**

Test	Reason
Ability to make and receive a call	This is the fundamental test; if it cannot make or receive calls with the test network, it is unlikely to operate with the real network.
Ability to perform a channel hand-off	A fundamental operation in a cellular network is hand-off. If hand-overs do not work properly, calls will be dropped.
Transmitter output power and power control	This affects the operating range and battery life. If power is too low, it might not cover larger cells. If it is too high, battery life will be shortened. If adaptive power control is used in a network, it is vital that the mobile outputs the correct power levels. Failure to do so will result in dropped calls or unnecessary channel hand-offs.
Modulation quality	Poor modulation quality will result in reduced operating range and may cause interference for other network users.
Receiver sensitivity	This affects the operating range of the mobile and the quality of the speech received.

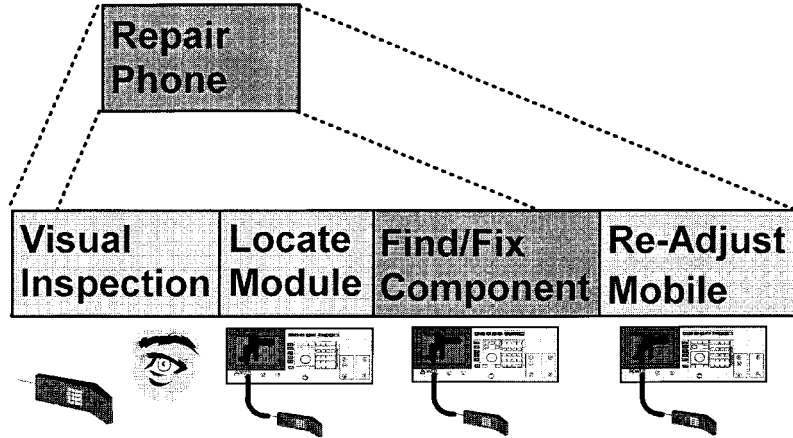
set to check the ability of the mobile to establish a call with the test set's simulated network, and also performs some parametric testing. Typically carrier power, power control, modulation quality, receiver sensitivity, and some form of spurious emissions testing is carried out. For a summary of the tests performed, and why, refer to Table 18.1.

The aim of incoming inspection is to make sure customers are not given a mobile that is obviously not working, and also to minimize the risk of polluting the network with badly radiating mobiles that may cause interference to others while seeming to operate normally. Usually the test sequence is automated using an external computer, or it could be part of the cellular test set's built-in firmware. Automation means all the mobiles are tested to the same pass/fail limits in a similar way. This type of incoming inspection test usually is repeated early in the service and repair process if a mobile is suspected of being faulty.

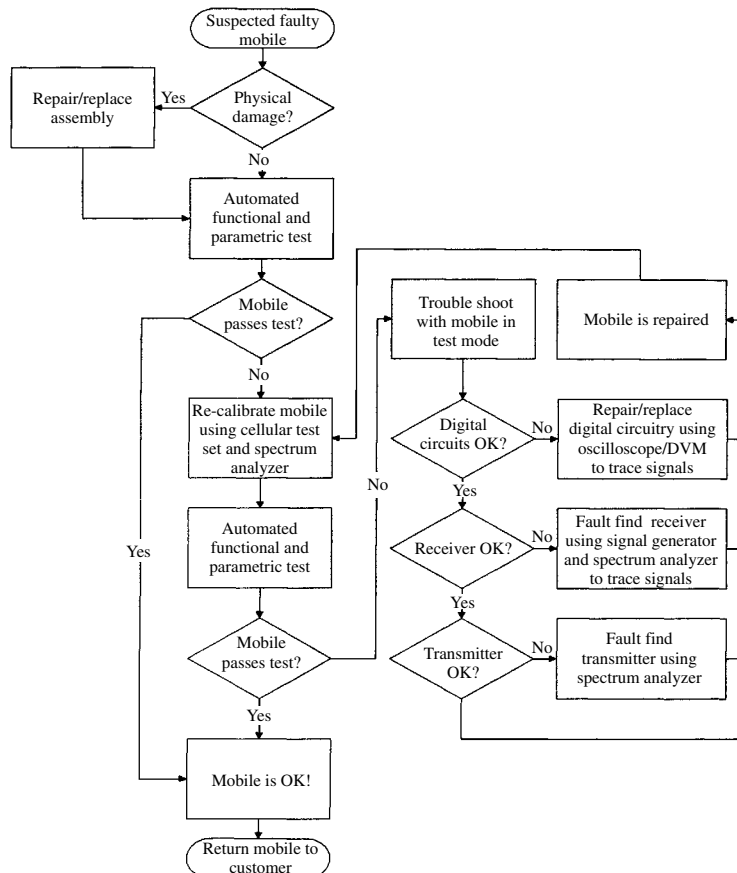
### 18.5.1 Mobile service organization

Mobile service and repair organizations can be divided into three categories ranging from simple test and exchange, through board-level repair, to component-level repair. Analog cellular phones traditionally have been repaired down to component level in field service shops. With the move from analog to digital technology, however, there has been a change in the skill sets needed to repair mobile phones. Much of the knowledge existed only in the manufacturers' factories. As the digital mobile market grows, driven mostly by the enormous success of GSM and IS95, there is a drive to move the repair operations for these phones from the factories to local repair shops. Figures 18.9a and 18.9b show a general flow for cellular mobile repair.





**Figure 18.9a Finding problems.** First identify obvious damage with a physical inspection. Then, if necessary, trace the fault to a particular module. Swap the module or repair the faulty component before recalibration and final functional test.



**Figure 18.9b Mobile repair flow diagram.** This shows a typical flow of a faulty mobile through troubleshooting to repair.

**Level 1 repair.** As shown in Table 18.2, the simplest repair strategy is referred to as *Level 1 repair*. In this type of operation, the suspected faulty mobile is tested using a process similar to the network providers' incoming inspection processes. A cellular test set is used to run an automated test sequence that checks basic call processing operations and measures key RF parameters. The mobile either passes or fails the test. This type of operation is sometimes known as Go/No-Go, referring to the pass/fail results, and automation means no special skills are needed to carry out the testing. A hard copy of the measurement results is usually obtained on a printer connected to the cellular test set. If the mobile fails this test, then the customer will be given an exchange mobile and the faulty one is passed to either a Level 2 or Level 3 repair shop.

**Level 2 repair.** A Level 2 repair shop will carry out repairs and calibration to circuit assembly level and also replace any of the modules used in the phone's assembly (Figure 18.10). First of all, cosmetic repairs are carried out to items like the antenna, battery, keypad, and case. Most mobiles contain only two circuit boards, one for the RF and one for the digital circuitry, although with miniaturization these sometimes are combined into a single PCB.

A cellular test set is used to check functionality and measure the mobile's basic RF parameters, including transmit carrier power, modulation quality, and receiver sensitivity. If the unit fails these tests and cannot be brought back into alignment, then circuit boards are swapped until the faulty one is identified. The faulty circuit board could be exchanged in a parts pipeline process with the mobile manufacturer or a larger service shop that is able to carry out a full component-level repair.


The high cost of maintaining a pipeline filled with high-value RF and digital circuit assemblies, coupled with import and export problems, is driving most manufacturers to move component-level repairs out into field service shops. It is important that mea-

**TABLE 18.2 Mobile Service Level Classifications.**

Class	Type of Repair	Function	Test Equipment
Level 3	Component-level repair	Performs full troubleshooting down to faulty components on circuit boards. Can carry out a full calibration and test to the same specifications as the original manufacturer.	Cellular test set (high functionality), spectrum analyzer, <sup>1</sup> oscilloscope, <sup>1</sup> digital voltmeter <sup>1</sup>
Level 2	Module-level repair	Performs testing (manual and automated) to trace fault to a replaceable module or circuit board. Failures passed to Level 1 for repair. Can perform some recalibration.	Cellular test set (high functionality)
Level 1	Go/No-Go	Automated testing that checks basic mobile performance. Failures are passed to Level 2 or Level 3 for repairs.	Cellular test set (limited functionality)

<sup>1</sup>May be included within the functionality of the cellular test set.

## Typically, Replace Whole Module



## Typically, Repair to Component Level



**Figure 18.10 Cellular phone modules.** Keyboard, battery, and antenna would be completely replaced if faulty. The circuit boards usually are repaired down to the component level. This might happen at the service shop if it is suitably equipped, or the circuit boards might be exchanged in a parts pipeline process with the manufacturer.

measurements performed on mobile RF modules are traceable in accuracy to the manufacturing process. Manufacturers will use the full allowable spread of the specification when producing the mobiles. This can lead to no-fault-found loops, where Level 3 repair shops and manufacturers waste time looking for nonexistent faults due to the fact that the Level 2 shop might have been testing with less accurate test equipment.

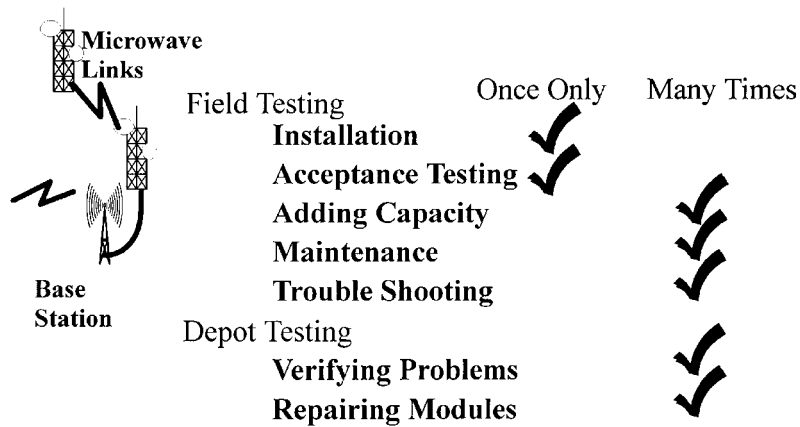
**Level 3 repair.** Level 3 service shops troubleshoot and repair mobiles down to the component level. They use well-equipped cellular test sets with additional toolkit functions, such as spectrum analysis and RF signal generator capabilities, to diagnose almost any fault that could occur. Traceability to the manufacturing process is often an issue in this type of repair; very often a reduced portion of the manufacturer's production test software is used to perform complex calibrations and checks on repaired mobiles when key components have been changed.

In Level 3 repair operations, after an initial visual inspection to identify obvious defects, the strategy is first to isolate the fault to a specific assembly. Multichip boards such as digital control or RF assemblies will be repaired down to component level. First the fault must be isolated to the transmitter or receiver section. A cellular test set equipped with a spectrum analyzer and signal generator may be used to inject and trace signals along the RF paths. An oscilloscope may be of use in checking out the digital portions of the circuitry. In general, specific test points are identified and documented by the mobile manufacturer to aid and guide the repair technician. Certain measurement results also are a good guide. Low output power or poor modulation quality points to a transmitter problem. Poor sensitivity points to a receiver problem. After repair, an automated test of the key RF parameters and general functionality is performed to ensure that the mobile meets its specifications.

18.5.2 Base station service testing

While a faulty mobile only affects a single user, a faulty base station can affect many users; network operators therefore are keen to avoid BS failures. Cellular networks perform much continuous monitoring, with mobiles reporting received power level and modulation quality data back to the network control center, which gives a good indication of general network well-being. Most operators perform some degree of routine maintenance testing to try to find faults before they occur. Figure 18.11 shows when RF testing is performed.

**In-service and out-of-service testing.** Table 18.3 shows the main characteristics of base station testing, which can be split into two types, in-service and out-of-service. The main difference between the two is whether or not the link is made between the base station and its controller attached to the network. Once a base station is connected to the network and commissioned, it is a problem to bring it out of service again because it means reducing the cellular coverage and potentially losing any established calls. Troubleshooting and maintenance tend to use



**Figure 18.11** When base station RF testing is needed. Some testing is carried out in the field on-site. Other testing is performed at the network operator’s service depot.

**TABLE 18.3** Base Station “In-Service” and “Out-of-Service” Test Characteristics.

In-Service	Out-of-Service
Little or no disruption to normal service.	Service has not yet started.
BS remains connected to its controller/network.	BS is disconnected from its controller/network.
Nonintrusive test methods can be used.	Nonintrusive or intrusive test methods can be used.
Typically used during: <ul style="list-style-type: none"> <li>▪ Acceptance Testing</li> <li>▪ Adding capacity</li> <li>▪ Maintenance</li> <li>▪ Troubleshooting</li> </ul>	Typically used during: <ul style="list-style-type: none"> <li>▪ Installation</li> <li>▪ Soak testing to find intermittent faults at network operator’s offices.</li> </ul>

nonintrusive test methods on base stations already in service. Nonintrusive methods rely on making transmitter measurements from test ports on the antenna feed and directly off the air.

**Maintenance testing.** Once the base station has been installed and handed over to the network operator, the lifelong challenge of maintaining a high quality of service begins. One of the operator's first jobs is to ensure that the base station meets requirements. In the months and years that follow, there are many situations requiring RF measurements at the base station site. As the operator's subscriber base increases, it will be necessary to add network capacity. If problems occur with the base station, they need to be tracked down and fixed. Sometimes, even when the base station is performing perfectly, service quality can suffer due to interference from other users of the radio spectrum. Network operators often face the challenge of tracking down sources of interference. In order to maintain the highest quality of service, many operators carry out periodic maintenance programs to anticipate and find problems before they affect service quality.

For acceptance testing (and, in fact, for practically all other testing during the life of the network), the BS is connected to the BSC via the protocol interface. In most cases when testing is required, it will not be possible to remove the base station from network service for any length of time. The resulting gap in coverage would adversely affect the quality of service for the end users. For these reasons, most performance measurements must be made nonintrusively on network equipment while it is in service.

Most digital systems gather a great deal of performance data during everyday operation. Mobiles and base stations report the power level and quality of the signal they are receiving. Many base stations have extensive built-in self-monitoring capabilities. Logging this information at the network operations and maintenance center (OMC) provides a convenient way to monitor alarms on the base station. Once major performance changes are detected in the base station transmitter or receiver, a technician can move on-site to investigate and repair the problem.

This OMC data provides limited information about in-channel operation. It generally will not indicate the transmitter's out-of-channel performance. It is possible for the base station to be completely functional while still generating spurious signals, or adjacent channel leakage that interferes with other cells or networks. For this reason it is important to supplement the OMC data logging with periodic maintenance programs designed to find transmitter problems that can affect the performance of other cells or networks.

Adverse conditions in the field mean that it might be necessary to exchange several modules in a base station to fix a problem, when only one is faulty. It also is generally desirable to exchange several modules when an intermittent problem is suspected. Finding the source of the intermittent problem and tracing the fault to a particular area is more conveniently carried out at a central repair depot. New, repaired, or problem modules can be run for extended periods at the repair depot to uncover any problems before being installed in an operational BS in the field.

For regular maintenance it is desirable to have nonintrusive maintenance techniques that do not disrupt regular service. Test equipment such as a spectrum analyzer

may be used to make power, timing, and modulation quality measurements on the BS transmitter without any interruption to its normal operation. There also are dedicated functional test sets specifically designed for different types of base stations. These test sets are equipped with measurement hardware appropriate for the type of base station under test and will be capable of demodulating the specific carrier signals for which they are designed.

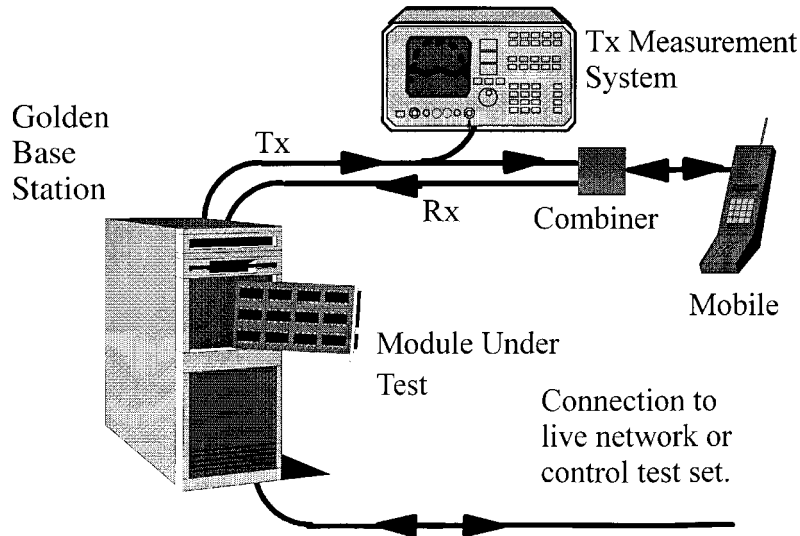
BS receiver testing generally is not possible while it is in service; this normally would involve breaking the connection between the BS and its controller to access the data the base station has received, thus rendering it out of service. In some cellular systems it is possible to make receiver sensitivity measurements using a *test mobile*. This is a specially made mobile that can set up a loopback call, over the air, with the base station. Data is sent from the mobile to the base station and then looped back to the mobile. An indication of the receiver sensitivity is gauged by how accurate the looped-back data compares with what was sent.

Table 18.4 lists some of the key types of measurements typically performed on a cellular base station. The exact nature of the measurement will depend on the type of cellular system being implemented. The priority column is a general guide only and will vary with network operator. The fact that a call can be made gives a good indication that most of the components in the base station must be working to a degree. Once this is established, it is possible to concentrate on ensuring that no interference is generated in the serving or adjacent cells and finally tune the base station to give optimum performance in its coverage area.

Most cellular base stations are located in inaccessible places, often making it difficult to perform detailed testing in the field. If a transmitter (TX) or receiver (RX)

**TABLE 18.4 Base Station Parametric Testing: What and Why.**

Priority	Test Parameter	Why
1	Call setup functionality	Basic requirement for operation
2	Spurious emissions	Any spurious signals could cause interference for the cell and other adjacent cells.
3	Intermodulation attenuation	A cell typically will have multiple users (equating to multiple TX/RX pairs active) at one time. They must not create interference for one another.
4	Transmitter carrier power	Verifies that the planned coverage is achieved.
5	Power versus time (TDMA systems)	Ensures that pulsed RF transmitters operate within the correct power time template. Noncompliance can lead to interference between calls on the same channel.
6	Modulation quality	Poor modulation quality will affect the coverage area of a cell. It may also generate interference for cell users.
7	Receiver sensitivity	Distant mobiles will not be received if sensitivity is poor. This affects coverage at the edges of cells.



**Figure 18.12 Depot testing of base station modules.** This shows the “Golden Base Station” that may be used to soak test suspected faulty TX/RX modules, as well as transmitter measurement equipment and a test mobile to check receiver sensitivity.

module is suspected of being faulty, it will be exchanged in the field for a known-good one. The suspected faulty module is taken back to the network operator’s service department for further investigation and repair.

**Troubleshooting faulty TX/RX modules.** At the network operator’s offices there usually is set up a “Golden Base Station” (Figure 18.12), connected to the real network but used as a test bed for suspected faulty modules. Comparative parametric measurements can be made between known working modules and the suspected faulty ones. The “Golden Base Station” also may be controlled by a special test set that can simulate the real network and extract received data from the RX modules to enable receiver sensitivity measurements.

### 18.5.3 Service and repair summary

Both mobiles and base stations contain transmitter and receiver components and are subjected to a similar set of tests. Cellular network operators strive to detect faults as early as possible, hence mobiles are subject to incoming inspection and base stations are monitored to verify continued correct operation. If faults occur, both base stations and mobiles are repaired down to component level.





# Cellular Measurement Descriptions

**Bob Irvine**  
**Tom Walls**

*Hewlett-Packard Ltd., South Queensferry, Scotland*

Cellular measurement activity can be divided into three major test categories: *transmitter*, *receiver*, and *functional*. The details of the test methods can vary with technologies or the device under test (DUT). Base stations and mobile stations share most measurement methods, but some are unique.

## 19.1 Test Equipment

Measuring so diverse a set of parameters as those outlined in the preceding chapter requires a range of test equipment. Standard measurement tools, such as spectrum analyzers, are customized using downloadable software to perform a specific suite of measurements. Other equipment, such as oscilloscopes and power meters, are used in their standard forms.

Growth in network subscribership has created a need for focused test solutions aimed at high-volume manufacturing, in turn creating the need for faster and more efficient test sets dedicated to checking the critical parameters. The way to achieve this has been a one-box tester in which internal interfaces are optimized for fast data transfer and real-time measurements. This poses a challenge to the test equipment manufacturer to ensure that the correct suite of tests is made available and measurement accuracy isn't sacrificed for speed.

## 19.2 Transmitter Tests

### 19.2.1 Carrier power

The method of measuring carrier power will vary with the type of system being tested. In a TDMA system like GSM, the power in one burst must be measured and averaged. This is quite different from measuring the carrier power in a traditional analog system


like AMPS or TACS, where the carrier is a continuous-wave signal that can be measured simply with a broadband power meter. We will examine two power measurement techniques, one for the GSM system and one for the IS95 CDMA system.

**Burst carrier power.** Measuring the output power from a GSM transmitter is complicated by the TDMA multiplexing scheme. The mobile transmitter only turns on during its active timeslot (Figure 19.1). The absolute output power is defined as the average measured during the middle or “useful” part of the burst when data is transmitted.

The GSM power measurement can be made conveniently with tuned or wideband power meters, provided they are capable of averaging only during the useful portion of the burst. Most cellular test sets have this type of power measurement built in, with the power meter synchronized to the base station simulator. Once a call is established with the mobile, the test set has a reference with which to make carrier power measurements.

Readings on conventional peak power meters, not specifically designed for GSM signals, will be affected by the overshoot or undershoot of the burst. Thermal power meters, or other devices with long-term averaging properties, sometimes can be used with TDMA systems by taking into account the 1:8 duty cycle of the signal being measured. This technique is generally not recommended for GSM signals. The relatively slow rising and falling edges of the burst, and the variation from phone to phone in pulse rise shape and burst length, can cause large changes in actual duty cycle. This, combined with the effects of overshoot and undershoot, can lead to poor measurement results.

- ▷ **Average Power During Useful Part =  $X$  dBm  $\pm$ 3 dB**
- ▷ **Maximum Output for Power Class =  $X$  dBm  $\pm$ 2 dB**
- ▷ **Instantaneous Power  $\pm$ 1 dB Relative to Average**



**Figure 19.1 Measuring the carrier power in a GSM burst.** The power is measured during the useful part of the burst when the 147 data bits are transmitted. It must remain flat to within  $\pm 1.0$  dB from the average level during this time. The output power levels associated with a range of GSM TX Levels are shown in the right-hand column. Level 15 is the lowest power level for a Phase 1 mobile and Level 5 the highest for a handheld transceiver.




Figure 19.2 Power versus time measurement. This shows the GSM PvsT mask.

### 19.2.2 Typical power control

Like many modern digital systems, GSM uses *dynamic power control*, which means that the mobile (and sometimes base station) is capable of varying the output power depending on the path loss that has to be overcome.

To test the power control capability of a mobile, it is necessary to send a signal that tells the mobile the power level on which to transmit, then make a carrier power measurement as previously described. The mobile may be conveniently signaled to change output power levels while on a simulated call with a cellular test set. Alternatively, most cellular phones have test mode commands that allow an external device to be used to command it to a particular transmit power level.

### 19.2.3 Power versus time

In TDMA systems it is necessary to verify that mobiles and base stations only transmit in their allocated timeslots. This is achieved by comparing the output carrier power burst against a power template mask; to pass, the burst must lie completely within the mask. The mask shown in Figure 19.2 is for the GSM system. Similar masks exist for other TDMA systems.

Power versus time can be measured conveniently with a time-gated spectrum analyzer, set to zero frequency span and tuned to the channel center frequency. The spectrum analyzer can be triggered to take a measurement either with an RF envelope detector that gives a digital output pulse on detection of the rising edge of the

RF pulse, or with a cellular test set that will supply the trigger signal based on its control of the call.

The spectrum analyzer settings must be chosen carefully. The resolution bandwidth is chosen to be narrow enough to give a signal-to-noise performance necessary to display the burst's full dynamic range. The resolution bandwidth also needs to be wide enough not to distort the profile by slowing down transitions or displaying ripple induced by modulation during the useful part of the burst. Once the burst is captured, it can be compared to the mask profile defined by the particular cellular standards. The burst usually is divided into three segments: rising edge, falling edge, and middle or useful part where the modulation takes place.

Some cellular test sets also can make this measurement without the use of an external spectrum analyzer. They use digital sampling and signal processing to make high-quality measurements of the burst profile.

#### 19.2.4 Burst timing accuracy


When a cellular test set is used to measure power versus time, there's usually a bonus: burst timing accuracy information. This is a measure of how accurately the mobile has timed the transmission of the burst. Since the cellular test set is a simulated base station, it knows exactly when to expect transmissions from the mobile. If they occur early or late, the test set can detect it and report it as a burst timing error. The burst timing error is related to specific bits in the modulated or "useful" part of the burst.

### 19.3 Modulation Quality, Phase, and Frequency Error


The modulation quality of an RF carrier will directly affect the ability of a receiver to decode the transmitted information correctly. Many of the digitally encoded cellular systems, including the Global System for Mobile Communications (GSM) and North American Digital Cellular (NADC), use modulation schemes that rely on accurately controlling the phase of the carrier to encode the binary sequence being transmitted. Near-perfect modulation would be ideal but requires complex and expensive transmitter design. A balance must be struck between cost-effective design and the desire for high-quality modulation. In the GSM system, for example, the peak phase error must be less than  $20^\circ$ , the RMS phase error must be less than  $5^\circ$ , and the frequency error must be less than 90 Hz for a mobile.

Before the process of calculating phase and frequency error can begin, a sampled record of the transmitter's phase trajectory during one TDMA burst is captured (Figure 19.3a). A number of techniques are available for obtaining this phase trajectory. One method uses high-speed sampling and digital signal processing to ensure high accuracy and repeatability. The incoming RF burst is down-converted and digitized directly; the sampled data is processed to extract the phase trajectory. Obtaining the phase trajectory using digital processing avoids accuracy and repeatability problems often associated with techniques using analog I/Q demodulators prior to digitizing.

Understanding these concepts requires thinking of the phase trajectory as being relative to the phase of the carrier center frequency. Streams of 1 bits will cause a phase decrease of  $90^\circ$  each, while 0 bits cause  $90^\circ$  phase increases. In the GSM system, where a Gaussian premodulation filter is used, the filtering stops the phase trajectory from meeting its  $90^\circ$  target points.




**Figure 19.3a Measuring phase and frequency error.** One TDMA burst is captured with high-speed digital sampling. It is demodulated and the ideal phase trajectory computed.



**Figure 19.3b Overlaying actual phase trace with computed phase trace.** The actual demodulated phase trajectory is compared with the computed ideal phase trajectory for the decoded bit stream. The difference between the actual and the computed trajectory represents the phase error across the bit stream.

The sampled phase trajectory is processed to produce a demodulated data pattern. The data pattern is used by the digital signal processor to synthesize a perfect phase trajectory.

Overlaying the sampled trajectory with the perfect trajectory highlights the imperfections in the measured modulation (Figure 19.3b). Subtracting the two waveforms produces a plot of phase error at each point across the TDMA burst.



**Figure 19.3c** Calculating the peak and RMS phase error and the frequency error. The phase error trace has two ingredients: slope and roughness. A best-fit straight line is used to calculate the slope. The slow change of phase across the burst, shown by the dotted line, is removed from the phase error calculation and expressed separately as frequency error. The remaining phase error trace, shown by the jagged line, is summarized by calculating its peak error and RMS error.

While these examples have concentrated on the GSM system, this technique of calculating the phase and frequency error is applicable to most of the digitally encoded modulation schemes that rely on a relative change in phase to convey a bit pattern.

The entire process of sampling a burst, calculating its phase trajectory, demodulating, producing a perfect trajectory, and calculating frequency error, peak, and RMS phase error, can be carried out using high-speed digital signal processors in a second or less (Figure 19.3c).

## 19.4 Interference Generation Tests

### 19.4.1 Spurious emissions

Spurious emissions tests are designed to protect other radio spectrum users from unwanted emissions from transmitter or receiver circuitry in mobiles or base stations. Specifications vary for mobiles on a call or in idle mode, and for different cellular systems. Depending on the design of the mobile, conducted and radiated spurious emissions must be checked. Testing spurious output over a variety of extreme power supply and temperature conditions sometimes can be revealing. When a mobile's battery voltage droops, the circuitry should switch off cleanly, rather than getting stuck in an unpredictable mode with unwanted RF outputs.

Spurious emissions tests can be made conveniently using a spectrum analyzer (Figure 19.4). The resolution bandwidth and sweep frequency range are chosen for the particular cellular system and type of spurious signal that are being checked. In

some cases, screened RF measurement rooms are needed to keep out other radio energy that could interfere with the measurement.

### 19.4.2 Output Radio Frequency Spectrum (ORFS)

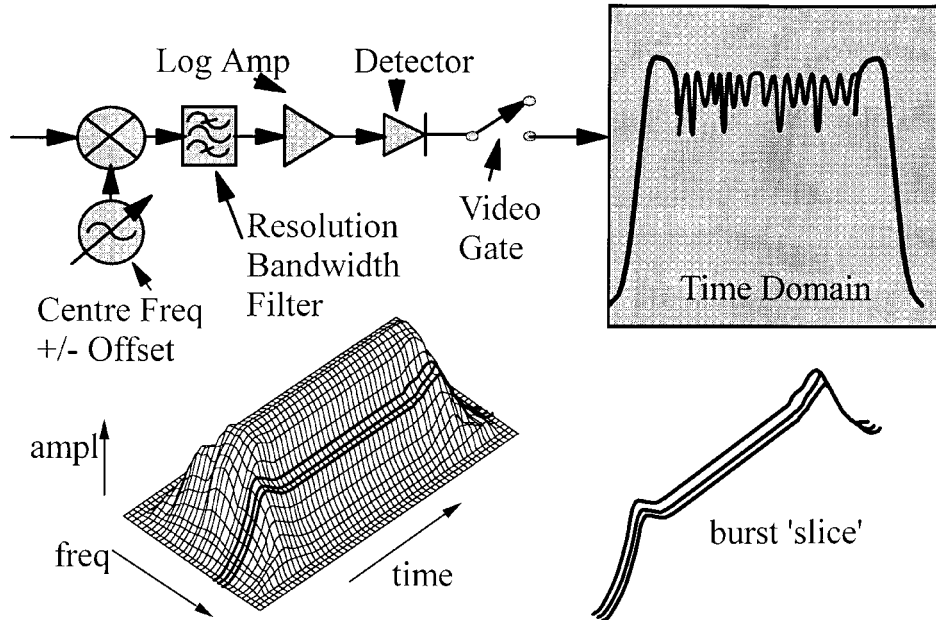
Output RF Spectrum (ORFS) is a test for spurious RF signals generated in channels adjacent to the active transmitting channel. This test is appropriate for both GSM mobiles and base stations; variants exist for other TDMA bursted cellular systems. There are two distinct measurements, *output RF spectrum due to modulation* and *output RF spectrum due to ramping* (or switching). The modulation ORFS test is designed to detect spurious energy generated by the carrier being MSK modulated. The ramping/switching ORFS is designed to detect spurious energy generated by the pulsed nature of the RF bursts.

Output RF spectrum can be one of the most difficult GSM measurements to visualize or understand. Matters are further confused by the fact that most pieces of measurement equipment display output RF spectrum traces as amplitude versus time at a particular frequency offset, not (as most would expect) amplitude versus frequency. See Figure 19.5.

The measurement is made using a time-gated spectrum analyzer, set to zero frequency span, and tuned to the channel center frequency plus or minus an offset. The offset frequencies allow the analyzer to take amplitude-vs-time slices from the measured bursts at the GSM specified frequency offsets. A reference measurement begins the sequence by establishing the amplitude at the center frequency (zero offset). The reference measurement is used to convert the results at each offset to relative or dBc values.



**Figure 19.4 Spurious emissions.** These are measured with a spectrum analyzer. The test frequency bands will vary depending on the radio system. Any unplanned or unwanted RF signals generated by the transceiver device are designated as spurious.



**Figure 19.5 Measuring output RF spectrum.** This shows the block diagram components of a spectrum analyzer, tuned to take a “slice” of a burst in the frequency domain at one carrier frequency offset. The trace in the box shows the time domain display.


Within the spectrum analyzer, the resolution bandwidth filter defines the width of each time domain slice taken from the burst. The log amp improves display dynamic range, the detector converts the down-converted and filtered input signal to a video waveform suitable for display. The video gate provides synchronization by selecting the correct portion of each TDMA frame for display and postprocessing.

The ripple displayed during the center of the burst is an expected by-product of the measurement technique. The instantaneous input frequency will be varying approximately  $\pm 67$  kHz due to the 0.3GMSK modulation and data pattern. Since the resolution bandwidth filter is tuned to select a narrow slice of the burst, the instantaneous input signal will move backwards and forwards across the selected frequency many times during the burst. The energy gathered by the filter during each crossing depends on the exact data pattern being transmitted. This produces the random ripple pattern shown.


The modulation ORFS is measured by averaging the trace at a particular offset frequency across the useful part of the burst (Figure 19.6). The average is compared to the value obtained at the channel center frequency to provide a dBc result. Results at individual offsets are compared with the GSM specification limits. Stepping away from the channel center in very small offsets would reveal the smooth, arch-shaped modulation spectrum shown.

The ramping ORFS is measured using a similar process (Figure 19.7). The resolution bandwidth and video postprocessing are modified to reveal the humped characteristics at the ends of the burst. Instead of averaging across the trace, as in the modulation





**Figure 19.6 Output RF spectrum due to modulation.** This is measured during the “useful part” of the burst. Time-gated spectrum analysis is used to eliminate the effects of the rising and falling edges of the RF burst.



**Figure 19.7 Output RF Spectrum due to ramping.** In the ORFS due to ramping, we are interested in the maximum peaks of RF energy generated by the rising and falling edges of the RF burst. The final measurement dBc value is converted to an absolute value in dBm for comparison with the GSM specifications.

case, the highest peak value is used as the result. The results at each frequency offset are converted from a dBc value to an absolute dBm value for comparison with the GSM specifications. Typically, phones with faster amplitude ramps produce poorer spectrum due to ramping performance. In some base stations the rate the power amplifier turns on can be adjusted to minimize the effects of ORFS due to ramping.

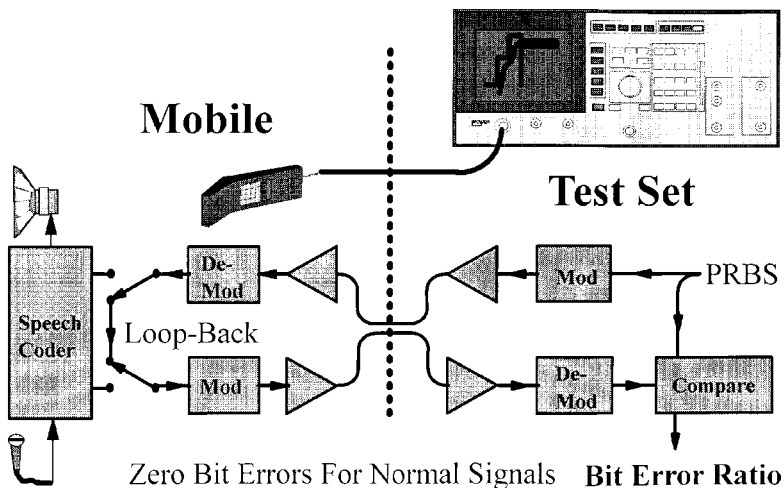
## 19.5 Receiver Tests

### 19.5.1 Bit error ratio test (BERT)

A bit error ratio test typically is used to assess the receiver sensitivity of a mobile or base station operating one of the digital cellular systems. The measurement goal is to determine how well the receiver can demodulate the digitally encoded data. The received data bits usually are not available directly from the receiver's integrated circuits for testing, so a special *loopback* mode usually is implemented: The device under test decodes and then retransmits the decoded bit stream back to the measurement test set. One cannot simply measure the audio output of the speech coder; the use of error correction schemes in the coded data would mask the true receiver sensitivity.

A typical test setup for mobile receiver sensitivity is shown in Figure 19.8. A cellular test set with a high-quality signal source is used to establish a call with the mobile, which is then signaled to go into loopback mode. Parameters that must be set are the amplitude of the stimulus signal from the test set and the number of bits over which the test will run.

A pseudorandom bit sequence (PRBS) normally is used as the stimulus data stream. The stimulus signal amplitude usually is set to a low level ( $-102$  dBm for GSM, for example) because the intent is to stress the mobile's receiver. The mobile will decode the bit stream and retransmit it back to the test set at a much higher level; a comparison is made between what was sent and what comes back to deter-



**Figure 19.8 Mobile receiver sensitivity test using a cellular test set.** The mobile loops the received data stream back to the test set for comparison with what was sent.

mine the bit error ratio. The signal is transmitted back to the test set at a high level in order to stress only the receiver of the device under test, not that of the test set.

The output from a bit error ratio test usually is a percentage bit error rate or a count of the number of bit errors.

### 19.5.2 Frame erasure

Some systems (GSM, for example) incorporate data coding schemes that cause data bursts to be erased if they contain a certain number of bit errors. This keeps spurious noise due to corrupt data entering the speech decoder from being generated in the earpiece. If an entire burst is erased, this is known as *frame erasure*; it is useful to measure this because it gives an indication of clumped bit errors.

## 19.6 General Tests

In addition to the RF-specific tests, there are more general tests that are performed on mobiles and base stations.

### 19.6.1 Current drain and battery life

The current drain of a cellular mobile will directly affect its battery life. Mobiles use a lot more power when transmitting than when receiving. In the traditional nonpulsed analog systems, measuring battery drain is a matter of using a simple ammeter that monitors the current drain depending on the output power of the transmitter.

In the modern bursted digital systems, the current drain now varies dynamically as the transmitter pulses on and off. Many different schemes have been designed to extend the battery life, such as *discontinuous transmission* (DTX) and *discontinuous reception* (DRX). DTX causes the mobile to stop transmitting when there is no voice detected in the mouthpiece. DRX means the mobile goes into a low-power standby mode while it is switched on but not actually on a call. It “wakes up” periodically to check if it has been paged.

To measure the current drain in these pulsed systems accurately, it is necessary to link the operation of the current meter to the cellular test set’s protocol controller. The link allows measurement of peak currents during transmitter turn-on and average current when the receiver is in idle mode.

### 19.6.2 Protocol testing

Protocol testing is used to verify that mobiles and base stations of a similar system type can interact properly. It generally requires specialized test equipment that is dedicated to a particular cellular system. Protocol testing emphasizes testing the mobile or base station in as many real-life scenarios as possible. Depending on whether a mobile or a base station is being tested, some form of simulator for the other part of the system will be required.

Testing protocol is usually a very time-consuming process because it requires setting up *all* possible combinations of signaling parameters that are likely to occur in the everyday use of the device under test.

## 19.7 Code Division Multiple Access (CDMA)

While many facets of CDMA are quite different from analog or even TDMA digital cellular systems, a few key aspects drive many of CDMA's unique testing requirements. Unlike most TDMA digital cellular systems, the CDMA wideband transmission system allows powerful error correction codes to be applied to all of the encoded voice data bits. CDMA therefore does not need tests designed to examine transmission quality for different types of data bits (classes).

In addition, the processing gain of CDMA (error correction codes plus spreading codes) makes the CDMA system very tolerant of transmission errors. What would appear to be gross errors in any other cellular system's transmitted signal are normal for a properly operating CDMA mobile. Traditional tests that examine modulation quality (error vector magnitude) and receiver performance (bit error rate) do not provide meaningful insight into a CDMA mobile's performance. For transmitter measurements, this leads to a new measurement specific to the CDMA modulation format.

Finally, the fact that CDMA is designed to operate with high levels of interference also drives new measurements that must duplicate the normal interference levels experienced by a CDMA mobile.

Developed by industry members of the TIA, IS-98 is designed to be an open industry standard to promote equipment interoperability. This document specifies minimum standards of performance for environmental, protocol, transmitter, and receiver characteristics for both the AMPS analog and CDMA digital modes. The number and complexity of these tests prevent a detailed discussion of all of them here. Many, while important for initial Type Acceptance, will not be performed regularly in typical manufacturing, incoming inspection, or service applications due to their costly nature. Accordingly, this section will concentrate on the key tests that will be used most often.

While test modes can be used to facilitate testing, industry members rejected them and opted for a more general approach to testing via a simulated over-the-air link. This will allow any mobile to be tested with any device that follows the IS-98 standard.

To simplify testing, IS-98 specifies that all CDMA mobiles must support a special service option. The CDMA standard allows for multiple service options to handle future requirements such as data services. Service option 001 is the normal speech transmission mode for CDMA. Service option 002 is the data loopback mode called out in the IS-98 standard. Service option 002 provides a convenient method to test a CDMA mobile under a simulated over-the-air link.

Service option 009 also is a data loopback mode, but is for testing the new 14.4 kbps traffic channel used with the improved vocoder developed by the CDMA Development Group (CDG). In both data loopback modes, the CDMA mobile demodulates the signal it receives from the base station simulator and then retransmits the same data back to the simulator. This allows accurate characterization of the CDMA mobile receiver performance.

### 19.7.1 Getting a CDMA mobile on a simulated link

To establish a simulated link, a base station simulator must:

- Provide a pilot channel for short code timing and frequency reference.
- Transmit a sync channel to provide system time (fine synchronization).

- Call the mobile via a paging channel requesting service option 002 (the mobile autoanswers).
- Direct the mobile to a traffic channel.
- Pass protocol messages to the mobile on the traffic channel.
- Maintain the link during required measurements.

In order to test a CDMA mobile on a simulated link, the test equipment functioning as a base station simulator must provide specific signals and protocol messages to establish and maintain a CDMA link. The simulator must provide a *pilot channel* to allow the mobile to get short code timing alignment and frequency alignment, and a *sync channel* that broadcasts the state of the long code and system time to establish proper time alignment. To create a link, the simulator must call the mobile via a paging channel and direct the mobile to activate service option 002. Once on a simulated traffic channel, the base station simulator must maintain the link by passing any required protocol message to the mobile during testing.

In addition to supporting pilot, sync, paging, and traffic channels, the base station simulator must provide other channels to simulate the nominal interference presented to a CDMA mobile. Two noise sources are required: an OCNS source to simulate the noise from other users in the same cell, and an AWGN source to simulate the noise from users in adjacent cells.

OCNS stands for *Orthogonal Channel Noise Source*. Since other users in the same cell are encoded with orthogonal Walsh codes, OCNS noise must use a different Walsh code than the one used for the simulated traffic channel link. AWGN stands for *Additive Gaussian Noise*. The interference from users in adjacent cells is not orthogonal, but is uncorrelated since they are encoded with the short sequence ( $2^{15}$ PRBS) that is offset in time. The AWGN source provides uncorrelated noise that accurately simulates the interference from users in adjacent cells.

All of these sources must be accurately calibrated and support relative amplitude resolution and accuracy of  $\pm 0.2$  to  $\pm 0.1$  dB. This performance is necessary to accurately set the desired signal-to-noise ratios required for tests called out in IS-98. The sensitivity of CDMA phones at their performance limit translates into the fact that a 0.8 dB change in  $E_b/N_t$  (signal to noise ratio) can alter the FER performance from 0.5 percent to 5 percent! This is why the relative accuracy of the test equipment is vital in order to get good measurement results.

### 19.7.2 CDMA transmitter tests

This section examines some of the transmitter tests suggested in the TIA IS-98 document. The IS-98 tests concentrate on transmitted waveform quality, power control performance, absolute power characteristics, and spurious emissions. CDMA transmitter tests include:

- Frequency Accuracy
- CDMA Hard Hand-off
- Time Reference Accuracy
- Waveform Quality ( $\rho$ )

- Range of Open-Loop Power Control
- Time Response of Open-Loop Power Control
- Access Probe Output Power
- Range of Closed-Loop Power Control
- Maximum RF Output Power
- Minimum Controlled Power
- Standby and Gated Output Power
- Conducted TX Spurious Emissions
- Radiated TX Spurious Emissions

This section concentrates on waveform quality, open- and closed-loop power performance, maximum RF output power, and gated power.

**Waveform quality.** The figure of merit specified in IS-98 for the quality of a OQPSK-modulated transmission from a CDMA mobile is called  $\rho$  (Greek letter rho). The  $\rho$  measurement is also referred to as the *power correlation coefficient*. The concept of the  $\rho$  measurement is fairly simple. Although the CDMA system is designed to operate with high levels of interference, the ultimate capacity of any given cell is limited by the total interference (number of active users). For adjacent cells that are equally loaded, this limit is about 32 callers per cell or sector.

If any mobile station's transmitter is not properly encoding each user's data into the required code, some of the transmitted power will appear as increased noise to other users. The  $\rho$  measurement computes the power of a CDMA transmitted signal that correlates to the desired code. Thus  $\rho$  gives an indication of the increased interference that will be caused by modulation errors in a CDMA transmitter.

A  $\rho$  value of 1.00 indicates that all of the transmitted power correlates with the ideal transmission code. The specified performance level that a CDMA mobile must meet is 0.944, indicating that 94.4 percent of the transmitted energy correctly correlates into the ideal code. At this level of  $\rho$  performance, the increased noise to other users caused by a CDMA transmitter will be an additional 0.25 dB.

In the test equipment the  $\rho$  measurement is performed by downconverting a CDMA modulated signal to an IF low enough to allow the waveform to be digitized, the signal is analyzed by a DSP processor. By processing the captured waveform data, the test equipment accurately computes the power correlation coefficient  $\rho$ . In addition to performing the  $\rho$  measurement on an active traffic channel, the test equipment can also perform the test mode  $\rho$  measurement. To use this mode, the CDMA phone under test must support a special firmware test mode.

**Frequency accuracy and static time offset.** Two other important mobile station parameters derived from the  $\rho$  measurement are *transmitted frequency error* and *static time alignment*. Since the transmitted CDMA waveform is spread using pseudorandom codes, the resulting RF waveform appears as a block of random noise. A conventional frequency counter cannot accurately measure the center fre-

quency of an OQPSK modulated signal. The value of the frequency error used to maximize the measured value of  $\rho$  provides the estimate of the carrier frequency error. In a similar manner, during the calculation of  $\rho$  the DSP must derive an estimate for the static time offset. This is a measure of how accurately the CDMA mobile has aligned its timing to the reference signal broadcast by a CDMA base station.

As a part of the  $\rho$  measurement, the test equipment uses its DSP to calculate and report both frequency accuracy and static time alignment. In addition, the test equipment reports the parameters of carrier feedthrough, amplitude error, and phase error. The carrier feedthrough parameter is a measure of I/Q modulator DC offsets that result in degraded  $\rho$  performance. If the carrier feedthrough is higher than  $-25$  dBc, this could be a major source of  $\rho$  degradation. The I/Q modulation parameter's magnitude error and phase error help pinpoint possible sources of poor  $\rho$  performance.

**Open-loop power tests.** Open-loop power control causes a CDMA mobile to monitor the received power from the base station and continuously adjust its output power accordingly. The mobile ideally must raise or lower its output power linearly for every change in the received power from the base station. Open-loop power control follows the following equation, with the powers expressed in dBm:

$$\text{Mobile TX Power} = -73 - \text{Received Base Power} \quad (19.1)$$

Measuring the accuracy with which a CDMA mobile performs open-loop power control requires that the mobile be actively transmitting and monitoring the signal level from a CDMA base station simulator. Then, by changing the output level of the pilot channel and measuring the response of the CDMA mobile, one can verify the mobile's open loop power control performance. This test is performed at three different power levels of the base station simulator's pilot channel. Note that the power measuring instrument must be able to measure power accurately over an 80 dB dynamic range. By specification, the CDMA mobile must follow the open-loop power control equation with  $\pm 9.5$  dB accuracy.

Measuring open-loop accuracy is relatively easy. Here is an at-a-glance summary of the test parameters:

- Verifies open-loop power control estimate accuracy
- Measure over an 80 dB dynamic range
- Measured at:
  - Base  $-105$  dBm – Mobile  $+32$  dBm
  - Base  $-65$  dBm – Mobile  $-8$  dBm
  - Base  $-25$  dBm – Mobile  $-48$  dBm
- Mobile should be accurate within  $\pm 6$  dB, and must be within  $\pm 9.5$  dB.

First establish a service option 002 call; then set the internal CDMA source to the specified level and measure the mobile's power. The average power detector should be used for the  $-105$  and  $-65$  dBm/1.23 MHz test points, and the channel power detector is required to measure the  $-25$  dBm/1.23 MHz test point because of the very low signal level returned by the mobile under test for this condition.

**Closed-loop power tests.** For closed-loop power control, the base station directs the mobile to fine-tune its output level. Based on the received level, the base station commands the mobile to increase or decrease its output power by 1 dB every 1.25 ms (800 times per second). The standard method of testing closed-loop power performance involves verifying the overall range and linearity of the mobile's closed-loop power control range.

A CDMA mobile station must demonstrate a  $\pm 24$  dB closed-loop power dynamic range, as well as have a well-defined slew rate as it changes power. To verify performance, the test equipment first must establish a call with the CDMA mobile, then command the mobile to increase its power by over 24 dB and measure that the mobile has, in fact, increased power at least 24 dB. The mobile also must be commanded to lower its power by at least 24 dB to verify that the mobile can decrease its power by at least that amount. Here is an at-a-glance summary of closed-loop power tests:


- Verifies closed-loop power control range and linearity.
- Measured over a  $\pm 24$  dB dynamic range.
- The mobile must offer at least  $\pm 24$  dB of closed loop power control around the open-loop power control estimate.
- Measured by sending 100 up and then 100 down power control bits.

**CDMA power measurements.** One of the TIA power tests involves measuring the maximum output power of a CDMA mobile. Based on the class of the CDMA phone, the maximum output must be at least 200 mW for class 3, 500 mW for class 2, or 1.25 W for class 1. (Most will be class 3, i.e., small handheld units.) To make this measurement, a service option 002 call must be established with the access channel parameters set to drive the phone to the highest possible output power.

Once the call is established, the power of the test equipment's source is lowered to  $-105$  dBm/1.23 MHz, and the power is measured using the Average Power measurement. Traditional power detectors found in analog one-box test sets will not provide an accurate power measurement on a CDMA signal due to the wide, fast amplitude variation caused by the CDMA system's modulation format. The test equipment's average power detector is designed to capture these fast modulation fluctuations. Once the CDMA signal is detected by the power detector, the test equipment sends the captured data to a DSP processor that computes the actual power contained in the CDMA signal. To summarize CDMA power measurements:

- Maximum output power test:
  - Set CDMA source to  $-105$  dBm/1.23 MHz.
  - Set access channel parameters to produce full power.
  - Make a service option 002, full-rate call.
  - Measure power.
- Maximum power specifications:
  - Class 1 mobiles: 1.25 to 6.3 W
  - Class 2 mobiles: 0.5 to 2.5 W
  - Class 3 mobiles: 0.2 to 1.0 W
- Test requires the accurate measurement of a wideband signal with high crest factor.





**Figure 19.9** Gated output power.

**Gated output power.** In order to provide maximum system capacity, the CDMA cellular system uses a variable-rate voice coder. The coder varies the data rate according to the activity in the voice channel. When the voice coder drops below full rate (9600 bps), a CDMA mobile pulses its output on and off proportionally with the data rate reduction. Thus, at half rate a CDMA mobile transmits 50 percent of the time, and at one-eighth rate (1200 bps), it transmits 12.5 percent of the time.

To minimize interference caused by pulsing the RF carrier, IS-98 specifies a time-versus-amplitude template to which a CDMA mobile must conform. Figure 19.9 shows the required rise and fall times with which a CDMA mobile must comply when it pulses its output. Unlike many TDMA systems, the CDMA time-versus-amplitude specification only specifies a 20 dB dynamic range. This is possible since all CDMA mobiles use the same frequency anyway, and they are designed to operate at that level of interference.

### 19.7.3 Receiver tests

IS-98 describes all of the CDMA-specific receiver tests and the minimum acceptable performance for each test. These tests concentrate on demodulation performance under various transmission conditions:

- Demodulation of Paging Channel in AWGN
- Demodulation of Forward Traffic Channel in AWGN
- Demodulation of Forward Traffic Channel in Multipath Fading Channel

- Soft Handoff Power Control Bit Tests
- Receiver Sensitivity and Dynamic Range
- Single-Tone Desensitization
- Intermodulation Spurious Response Attenuation
- Receiver Spurious Emissions

This section examines in detail receiver sensitivity and dynamic range, demodulation of the forward traffic channel with AWGN and fading, and intermodulation spurious response attenuation.

**Frame error rate.** Each CDMA frame contains the digitized voice bits for 20 ms of speech. When a frame has been so corrupted that error correction cannot fix all the errors, a frame error has occurred. Because of the processing gain of the CDMA system (redundancy in the transmitted waveform), individual bit errors in the received waveform are of little consequence.

Bit errors on the received signal usually are repaired by the error-correcting action of the CDMA codes. Because of this, the traditional test of digital receiver performance, bit-error-rate, has no usefulness in CDMA applications. A more meaningful test for CDMA is *frame error rate*. FER is the true measure of CDMA receiver performance. Corrupted frames in CDMA are not retransmitted; the voice decoder must either interpolate the missing data or mute the audio output. The acceptable level of frame error rate for acceptable speech quality is about 3 percent. Frame error rate at a glance:

- Every 20 ms of digitized speech (9600 bps or less) constitutes a CDMA frame.
- When a frame cannot be correctly decoded, a frame error has occurred.
- Individual chip errors (over-the-air) do not significantly degrade CDMA performance.
- CDMA voice quality is acceptable with frame error rates up to 3 percent.

**Receiver sensitivity tests.** The CDMA receiver sensitivity test is performed without any AWGN interference. Only OCNS noise is required to simulate other users in the same cell. The next two paragraphs describe a typical setup for a sensitivity test for all of the required channels.

**Base station simulator configuration for sensitivity tests.** The total cell power  $I_{or}$  is specified at  $-105$  dBm in a 1.23 MHz bandwidth. The power in each channel is specified in terms of dB below the total cell power. For example, the pilot channel always has the most power and is specified to be  $-7$  dB below the total cell power. Since the total power must add  $\mu$  to the total cell power, the OCNS source is set to produce the remaining power. Although this places a large amount of power in OCNS, it accurately simulates real conditions where up to 30 other users may be active in a cell. The test is repeated with a total cell power of  $-25$  dBm per 1.23 MHz bandwidth to ensure that the receiver does not overload with strong signals.

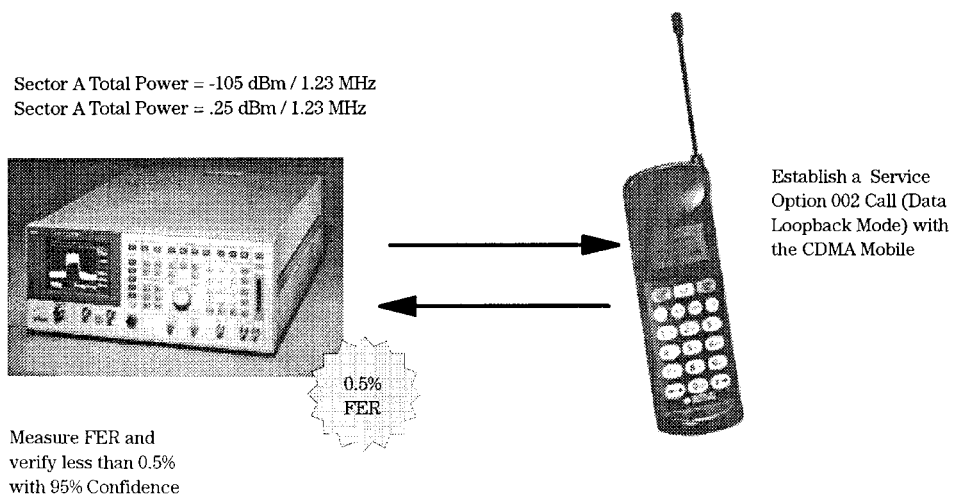
The CDMA source in the test equipment is fully user-configurable and can be set up to conform to the IS-98 recommended settings. The test equipment allows entry

of individual CDMA source settings in the same manner as IS-98. In the test equipment, the sector power ( $I_{or}$ ) for a sensitivity test thus is set to  $-105$  dBm/1.23 MHz, the pilot channel is set to  $-7$  dB, the paging channel is set to  $-12$  dB, the sync channel is set to  $-16$  dB, and the traffic channel is set up  $-15.6$  dB. Once these channel levels are specified, the test equipment automatically calculates the level of OCNS to produce the set sector total power.

A typical measurement setup for performing the CDMA receiver sensitivity and dynamic range test is shown in Figure 19.10. The test equipment is set up at the specified total cell power, with each individual channel at its proper relative power. A call is placed to the mobile with the test equipment configured for a service option 002 connection. The test equipment transmits a PRBS data pattern to the CDMA mobile; the mobile demodulates the signal and then transmits the data back to the test equipment. The test equipment compares the returned data to the transmitted data to calculate the frame error rate.

Since the mobile retransmits the data to the test equipment at a high level due to open-loop power control, any returned frames that have errors are considered bad frames. For the receiver sensitivity test, the minimum performance level 0.5 percent FER. The sensitivity test is performed only on full-rate traffic channels. The dynamic range test is similar to the sensitivity test except that the test equipment source is set to produce a high-level signal instead of a low-level signal. This test verifies the ability of the mobile station's receiver to handle a strong signal without problems.

**Receiver selectivity tests.** The selectivity test for CDMA, called *demodulation of the forward traffic channel in AWGN*, is similar to the sensitivity test, but adds the AWGN noise source. The AWGN test is the CDMA equivalent of an analog phone selectivity test because it measures the ability of the CDMA phone to extract the desired signal in the presence of other users (simulated by the AWGN source). The



**Figure 19.10** Receiver sensitivity and dynamic range.

AWGN noise source is set to  $-74$  dBm in a 1.23 MHz bandwidth, while the total cell power is set to  $-75$  dBm in a 1.23 MHz bandwidth. The intermodulation spurious response attenuation test and the single-tone desensitization test measure FER performance in the presence of CW tone interference sources rather than using an AWGN noise source.

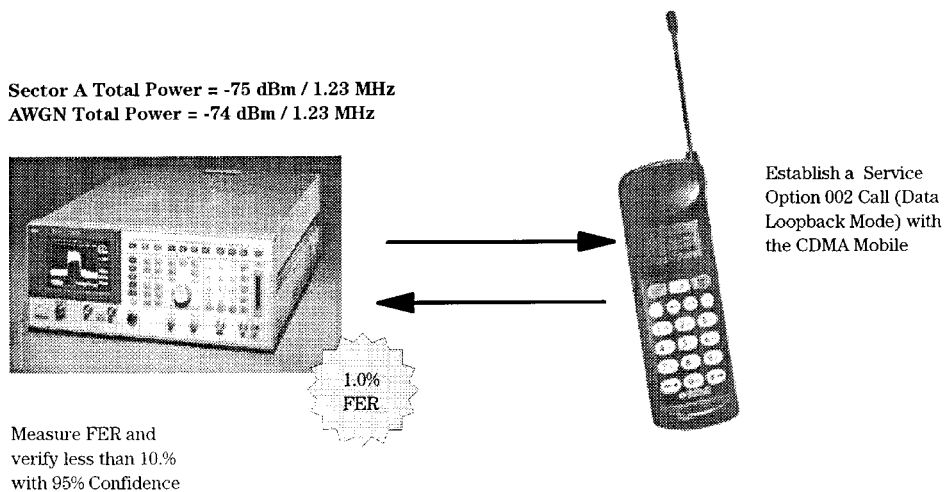
Like the sensitivity test, the demodulation of the forward traffic channel in AWGN test cannot be performed without a simulated CDMA link with the mobile under test. Service option 002 is required to measure the FER performance accurately. As shown in Figure 19.11, this test is performed at all four possible traffic channel data rates: 9600 bps, 4800 bps, 2400 bps, and 1200 bps. For the 9600 bps case, three traffic channel relative levels are called out, with the minimum FER performance ranging from 0.5 to 3 percent. These three test setups correspond to three different signal-to-noise ratio conditions presented to the mobile's receiver. Similar test conditions are specified for the other three traffic channel data rates.

The fading tests for CDMA are similar to the FER tests in AWGN, but use a channel simulator to distort the cell simulator's signals in a manner similar to real-world conditions. This test is repeated under a variety of traffic channel signal-to-noise ratio conditions at all of the traffic channel data rates. Three main multipath fading profiles are used:

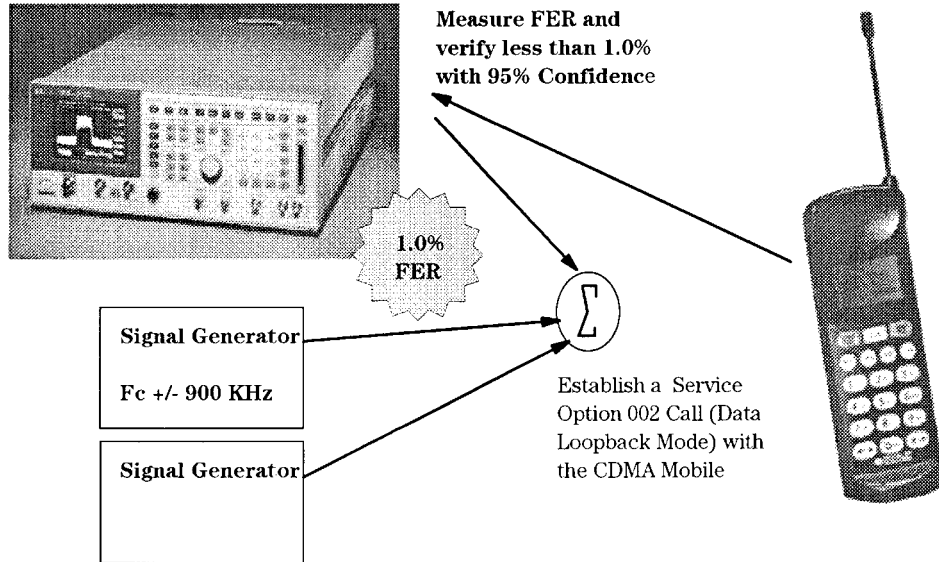
- *Case 1:* 8 km/h velocity with two-ray fading at 9600 bps only
- *Case 2:* 30 km/h velocity with a single ray at all four data rates
- *Case 3:* 100 km/h velocity with three-ray fading.

The FER minimum standards for these tests range from 3.0 down to 0.349 percent.

The test equipment can be used for fading tests, but requires a number of external supporting products to complete the measurement. First, the IS-98 standard re-



**Figure 19.11** Demodulation of the forward traffic channel with AWGN.



**Figure 19.12** Intermodulation spurious response attenuation.

quires that the traffic channel be the only channel to undergo the fading process. This means that the test equipment's AWGN noise source cannot be used because there is no method available to separate the test equipment's traffic channel and AWGN source. A fading simulator is used to apply the proper fading profile to the test equipment's CDMA source. Once the fading profile is added, the AWGN source can be summed to provide the proper amount of noise. This composite signal is routed to the unit under test. The mobile's reverse link is then sent to the test equipment for FER calculation.

The intermodulation spurious response attenuation test is similar to the FER with AWGN test, except that CW tones are used as interference sources instead of AWGN (Figure 19.12). The tones are placed at +900 kHz, +1700 kHz, and -1700 kHz relative to the carrier frequency. The placement of these tones is designed to test the ability of the CDMA receiver to reject potential intermodulation products. These signals are set to a power level of -40 dBm, and the total cell simulator power is set to -102 dBm in a 1.23 MHz bandwidth. Under these conditions, the minimum performance level must be 1.0 percent FER. A similar test that uses only a single CW interference source is called the *single-tone desensitization test*.

#### 19.7.4 CDMA functional tests

Some possible application areas for functional testing include manufacturing final checkout, incoming inspection of mobiles for service providers, or repair verification in service shops. Examples of these functional tests include placing a phone call to

**438 Cellular Networks**

verify overall functionality, a soft handoff check to verify rake receiver performance, and voice quality evaluation by a human operator. Here is a summary:

- Not specifically called out in IS-98
- Designed to quickly verify CDMA mobile operation for:
  - Manufacturing final checkout
  - Incoming inspection
  - Verification of repair
- Examples:
  - Call processing check
  - Soft/softer handoff check
  - Voice quality evaluation

**19.7.5 Softer handoff check**

If the base station simulator can support two CDMA cells or sectors, both the phone call functional test and the soft handoff functional check can be combined into a single test. For this test, a phone call is originated either by the mobile or by the base station simulator. Once the link is established, the base station simulator is set so that its second cell or sector presents an attractive level to the mobile for soft or softer handoff. After the mobile detects the second cell and requests soft handoff, the base simulator directs the second cell to transmit to the mobile and directs the CDMA mobile to listen to the second cell. If these actions are completed successfully, the handoff capabilities of the mobile have been verified.

In summary, the CDMA softer handoff check verifies softer handoff capability, and includes these steps:

- Bring up a service option 001 call in echo mode.
- Activate sector B, then raise its power to above  $T_{Add}$ .
- Check to see if mobile sends a pilot strength message.
- Activate softer handoff.
- Speak into phone.
- Listen to delayed voice in phone.
- End soft handoff.

**19.7.6 Voice quality check**

Another excellent functional test is possible with the test equipment because of its voice echo mode. In voice echo mode, the test equipment echoes the voice data it receives from a CDMA mobile back to itself. In essence, the voice echo mode is the opposite function of service option 002. After a link is established, an operator can speak into the CDMA mobile and then hear his or her voice return through the mobile after a slight delay. The echo mode test has a cost advantage

in that a voice coder is not needed in the test equipment to perform this test. The steps are:

- Bring up a service option 001 call.
- Select echo mode.
- Speak into phone.
- Listen to delayed voice in phone.
- Terminate the call.

This quick functional test is accomplished by placing a service option 001 call from the mobile under test. This action verifies that the DUT can properly lock to a pilot, decode the sync and paging channels, and perform call processing while also verifying keypad operation. Once a link is established, the operator speaks into the phone and waits to hear his voice echoed back in the phone. This step verifies overall voice quality and the ability of the microphone and speaker to operate properly. The call is then terminated from the test equipment. This final step checks for proper call release.





# Cellular Network Life Cycle Testing

**David Bonner**

*Hewlett-Packard Ltd., South Queensferry, Scotland*

Cellular networks are evolving continually as new cells are added and new services are introduced. Managing an ever-changing network is difficult enough, but the difficulty is compounded by the pressures of having to improve the overall quality of the service provided to the customer. Due to this dynamic nature, a cellular network can be seen as having its own life cycle. This chapter describes the use of test equipment during all phases of that life cycle.

## 20.1 The Drivers for Continual Improvement

A network's success depends on its customers' ability to make their calls successfully. There are three drivers that determine whether customers can make calls and that drive every aspect of engineering work on the cellular network. These three drivers are:

- Coverage
- Capacity
- Call quality

Customers increasingly are expecting the service on a cellular network to be comparable to fixed-line telephony. For the customers this means the ability to make a call whenever they want, from wherever they are within the coverage area, with high quality throughout the duration of the call. The implications for a cellular network operator therefore are:

- Coverage where subscribers expect it
- Sufficient capacity to satisfy customer needs
- Call quality meeting customer expectations

## 442 Cellular Networks

### 20.1.1 Coverage

In the initial buildout of cellular networks, coverage is the primary engineering concern for the network operators. At the initial build-out of a cellular network, cells are introduced that cover a large geographic area. This coverage typically is for built-up areas and major road routes. These cells will be used to provide the initial cellular service.

The next stage in increasing coverage is what is termed *infill*. This is the stage in which the smaller towns are covered and extra cells are placed in previously covered areas to fill in the “holes” in the coverage. This can include in-building coverage.

### 20.1.2 Capacity

As well as being a method of increasing coverage, infill is typically the first activity aimed at increasing the capacity of the network. Other popular ways of increasing capacity are:


- Use of additional frequencies
- Dual-mode systems
- Dual-band working
- Cell sectorization

At initial build-out, cells normally are introduced into the network in an omnidirectional fashion. This means that the cells transmit in all 360 degrees of the compass. *Cell sectorization* is a technique whereby the cell is split into different sectors, with the sectors transmitting in different directions. Each sector has the same capacity as an entire cell previously had. A normal split is to have three sectors each transmitting in separate 120° directions. This has the effect of increasing the capacity to three times its previous value (Figure 20.1).

The most popular method to increase capacity would be to use more frequency channels for carrying the calls. A practical example of this is the dual-band systems that operate at 800 MHz and 1900 MHz in the United States. The use of extra frequencies is strictly controlled by government agencies, however, and therefore extra frequencies might not be available. If this is the case, other methods of increasing capacity will have to be used.

*Dual-mode systems* are systems that operate using different air interface standards for different functions within the network. A dual-mode terminal is required to make use of these systems. One example of these types of systems would be a dual-mode GSM/DECT system, where the users have access to the GSM portion of the network when on the move and access to the DECT portion of the network when in the office.

*Dual-band systems* are systems that operate using the same air interface but in two separate frequency bands. The GSM and DCS-1800 systems of Europe, which operate at 900 MHz and 1800 MHz respectively, are two systems that operate together in dual-band mode.



**Figure 20.1** The effect of sectorization on increasing wireless network capacity.

### 20.1.3 Call quality

Once coverage and capacity issues have been addressed, the focus of the cellular network operator turns to quality of service (QoS) optimization. Quality of service from the customer's point of view is:

- The ability to set up a call
- Adequate speech quality during the call
- Closing down the call cleanly when desired

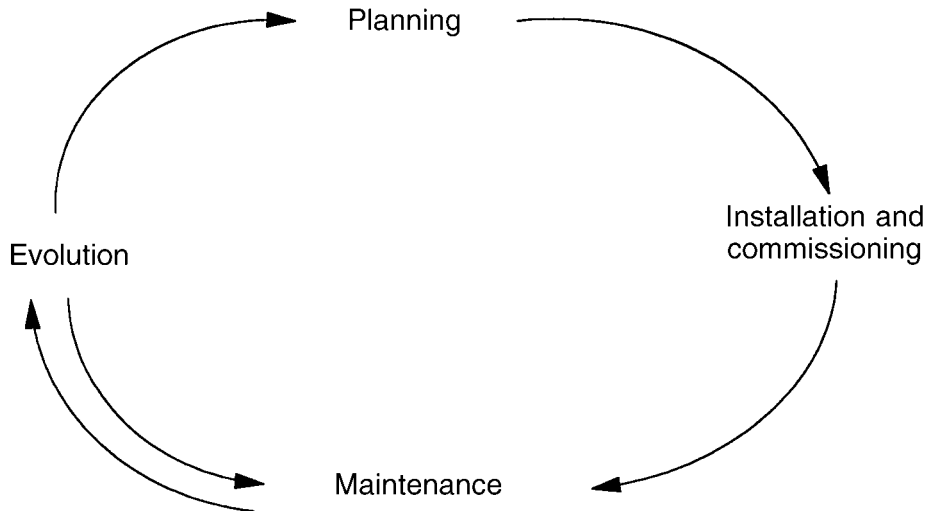
Quality of service is monitored on the network using various pieces of equipment:

- Operations and maintenance centers from the equipment vendor
- Drive systems
- Third-party management systems

These systems are described later in the chapter.

## 20.2 The Life Cycle Phases

Due to the customer expectations of coverage, capacity and quality, wireless networks are ever-changing. These networks require ongoing management. To understand the problems associated with testing and managing a wireless network, the network can be thought of as having a life cycle with various phases; with wireless networks, essentially there are four such phases: planning, installation and commissioning, maintenance, and evolution (Figure 20.2).



**Figure 20.2** The life cycle of a wireless network.

**TABLE 20.1** Wireless System Life Cycle: Tasks in Each Phase.

<ul style="list-style-type: none"> <li>• Planning               <ul style="list-style-type: none"> <li>–Survey</li> <li>–Frequency planning</li> <li>–Site planning (tower heights etc.)</li> <li>–Planning for new frequencies</li> <li>–Planning for expansion (more switching centers, extra coverage, etc.)</li> <li>–Spectrum clearing</li> </ul> </li> <li>• Installation and Commissioning               <ul style="list-style-type: none"> <li>–Installing base stations and switching centers</li> <li>–Commissioning base stations and switching centers</li> </ul> </li> </ul>	<ul style="list-style-type: none"> <li>• Maintenance               <ul style="list-style-type: none"> <li>–Fixing faults (including getting fixes from vendors)</li> <li>–Dealing with customer complaints</li> <li>–Testing new products</li> <li>–Testing new software releases</li> <li>–Data transcript into switching centers</li> </ul> </li> <li>• Evolution               <ul style="list-style-type: none"> <li>–Definition of required future features</li> <li>–Introduce new infrastructure</li> <li>–Additions of extra capacity</li> </ul> </li> </ul>
---	--

The wireless network is in all phases of the life cycle at any given time; it is the relative emphasis the network operator places on each phase that is important. Early on in the development of a network, most of the time is spent in the planning phase; as the initial plans are completed, the emphasis changes to the installation and commissioning phase. The transition between the phases shows the shifting emphasis by the operator and the movement of resources to concentrate on the next task.

Particular tasks are carried out in each phase of the life cycle. These are summarized in Table 20.1.

### 20.2.1 Planning

The planning phase starts with the award of a cellular network operator's license. The network operator will choose a network planning tool that will be used to deter-

mine where the base stations should be sited. The major testing effort at this phase involves drive testing of the proposed base station sites to make certain that the actual coverage provided by the cell is equivalent to that predicted by the network planning tool.

A proposed site is drive-tested by erecting at the site a mast that radiates an unmodulated carrier. The surrounding area is driven and monitored to determine the extent to which coverage will be possible from the proposed site. Due to radio interference from many sources, further tests are required to determine which frequencies are best for use by each base station. The use of frequencies with low interference is crucial in providing high speech quality to the end user.

The HP E4900 spectrum monitoring system is an example of equipment that can be used for this purpose. This system, built around a spectrum analyzer, provides a variety of measurements to characterize the RF spectrum including spectrum carrier measurement, occupancy statistics, and demodulation and recording. The system also provides a number of preprogrammed reports that allow the user to make frequency allocation decisions quickly. The E4900 system is ideally suited to repeated monitoring of an area, as when refinements are made to the coverage plan. Previous results can be compared to ensure that base station siting and frequency choice are as close to the ideal as possible.

Further tasks that take place in the planning phase have no direct test needs. Most of the testing comes in the later phases of the life cycle.

## 20.2.2 Installation and commissioning

This phase of the life cycle is when the network is actually built. The time to install and commission each piece of equipment is always short, so testing for correct equipment implementation must be carried out efficiently. Automated testing is essential to reducing the time taken to install and commission a base station. Because implementations vary from vendor to vendor, the tests required also will vary. To ensure efficient installation and commissioning, the automated testing must be tailored to each vendor's implementation.

Specialized test equipment is available to meet these needs. The HP 8921 series of test sets provide base station site test solutions for a number of wireless systems, including AMPS, CDMA, and TDMA. The base station testing carried out in this phase typically is a subset of the testing carried out in the manufacturing test (see Chapter 19) of the equipment. The tests carried out at installation and commissioning are specific to the type of wireless system being installed. Typical tests for a CDMA base station include waveform quality, frequency error, time offset, and code domain power. (The code domain power measurement calculates the power in each Walsh-coded signal to ensure that all of the codes can be used to carry calls.)

Base stations typically are connected to the test set for testing; if this goes well, some tests are run with the BS connected to the live network. These tests typically are various call scenarios, such as mobile-originated calls, mobile-terminated calls, and handovers. At this point some network- and vendor-specific features come in to play. Some base station implementations allow for the BS to be connected to the network, but with only specialized mobiles allowed access. This is particularly useful for

installation and commissioning because subscribers are protected from the BSs that have not been fully tested. If the test calls prove successful, the network partitioning features are turned off and the BS is then live in the network.

There is a multitude of other equipment in addition to base stations that must be installed and commissioned: transmission equipment, switching centers, and IN (Intelligent Network) platforms. Each piece of equipment is tested before connecting it to the network. The testing typically takes place in three stages:

- The system under test (SUT) is connected to a test harness (TH).
- The live network is connected to the TH, which is simulating the SUT.
- The SUT is connected to the live network for in-service tests.

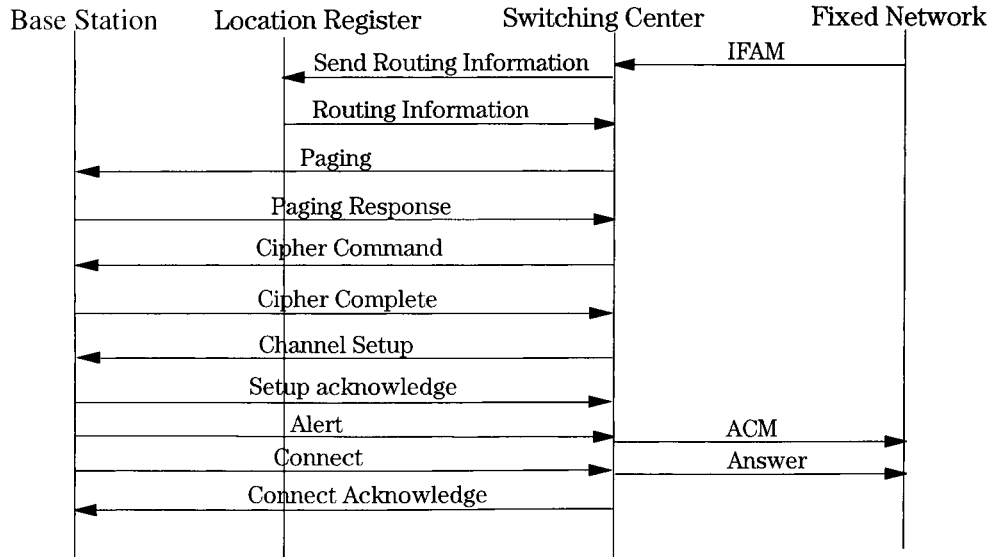
Testing switching centers and IN platforms in a wireless network is similar to testing of these pieces of equipment in fixed networks. Additional testing is required for the functionality specific to wireless systems. Transmission equipment testing, normally carried out by the third-party provider, is covered elsewhere in this book.

Switching centers and IN platforms are tested using the three-stage process previously described. All layers of the communications protocols are tested. The lower levels of the protocol typically are tested to specifications written by the operators of the fixed-line network with which the cellular network interconnects. In turn, these test specifications are based on international standards, such as the Signaling System 7 CCITT test specifications Q.781 and Q.782.

The upper layers of the protocol are tested using specific protocol decodes implemented on the Test Harness. These tests are based upon the real-world situations where the network equipment will be used. Most tests are based on various call scenarios, such as incoming calls, outgoing calls, emergency calls, and diverted calls. In addition, overload situations are tested to determine the actual throughput of the equipment. The test harness is programmed to respond to signaling messages initiated by the network equipment under test. The responses from the TH can be altered to determine how the system under test reacts under certain failure scenarios. Figure 20.3 is an example showing the testing of a switching center for a successful incoming call setup.

The test harness is capable of simulating all the other equipment with which the system under test communicates. As such, the TH must be capable of generating signaling messages for different protocols within a single test sequence. If a test harness with these facilities is not available, multiple test instruments can be used to simulate the various pieces of equipment. The test harness must be programmable, to allow the response of the system under test to be checked at all stages in the test sequence. For example, in Figure 20.3 the switching center must be capable of releasing all the allocated resources if, at any time, either the originating or terminating party closes down the call.

An independent protocol analyzer is placed on the link between the two pieces of equipment being tested. This instrument acts as the verifier of correct operation and as the arbiter if an error situation occurs. The HP 37900 can perform this protocol analysis function (in addition to having the ability to perform as a test harness).



**Figure 20.3** A protocol diagram for an incoming call setup.

Testing all protocol layers is of prime importance when the wireless network is to be built using equipment from multiple vendors. Due to different implementations of certain features, the interconnection of equipment from multiple vendors has to be tested more rigorously than if two pieces of equipment from a single vendor were to be tested. In these cases the test harness must be flexible enough to emulate equipment from various manufacturers. The HP 37900 is a test instrument capable of this function.

### 20.2.3 Maintenance

Maintenance is an ongoing task for the network operator. There are essentially three ways in which a network operator identifies problems in the network: customer complaints, periodic maintenance, and alarms generated by the network management system.

The customer care department within a wireless operation seeks to gather the customer complaints and feed them to the engineering departments, in order to identify any faults in the network. Regular preventive maintenance is carried out on the network to catch any faults that the customer and the network monitoring system have not seen. Due to the increasing reliance on network monitoring systems, however, regular proactive maintenance is on the decline.

A major effort has been put forth to increase the usefulness of network management systems because they are seen as a way of noticing network faults before the customer does, and as a means of reducing the costly regular maintenance programs.

One type of network management system, which is provided by equipment vendors, is the Operations and Maintenance Center (OMC). OMCs are provided to monitor the

performance of that specific equipment provider's gear. OMCs gather statistics generated from the network equipment distributed throughout the network. The OMC then consolidates the information to provide a network-wide view of how the system is performing, while flagging areas of the network that have specific problems.

The OMC is in an ideal position to provide specific information on current quality of service. Certain OMCs also provide failure counters that provide information to help identify specific problems. The major downfall of OMCs comes if a network operator decides to implement a multivendor network. The network equipment provider will then have to purchase an OMC from each vendor so that the whole network is managed in as coherent a manner as possible.

OMCs are not the only type of network management system available. Third-party companies have started to produce network monitoring systems that concentrate on certain areas of the network for which the OMCs traditionally have not provided effective solutions.

The HP AcceSS7 signaling monitoring system is an example of such a system. AcceSS7 monitors the SS7 signaling links in a wireless network and provides early warning of problems that will affect the customer. QoS measures in the AcceSS7 system include:

- Link status
- Channel assignment success rate
- Disconnect failure rate

The link status measurement shows the current status of each link in SS7 parlance: out of service, in service, processor outage, etc. The channel assignment success rate shows the percentage of channel assignments that have been unsuccessful for each monitored linkset. The disconnect failure rate identifies the linksets that have the highest ratio of disconnects indicating a failure has occurred. These key measurements, and others like them, are used to provide information on various types of failures as quickly as possible; they are updated typically every 15 seconds.

Once a problem has been identified, tools are required to determine the cause. There are two main methods of carrying out the problem diagnosis: *protocol analysis* and *drive testing*.

Protocol analysis is described elsewhere in this book. The application of protocol analysis in wireless networks is the same as in fixed networks. The major difference is that the RF interface in wireless networks can cause effects not seen in fixed networks, such as the sudden loss of a user because he or she has moved out of coverage. The nature of the RF interface has to be borne in mind when applying protocol analysis methods on wireless systems.

Drive-around systems typically consist of a mobile phone, a positioning mechanism (GPS, for example) and a PC. The phone is connected directly to the PC so that the computer can see all the activity taking place on the air interface when the mobile is in a call. The PC typically monitors all air interface activity including the received signal strengths received at the mobile station, as well as the call success rates and the reasons for any call failures. The positioning equipment is used to pinpoint the exact locations from which calls are made. In this manner the network can



be checked to ensure that QoS is maintained and to determine whether changes made to the network have resulted in QoS improvements.

The system provided by Necsy is a good example of this type of tester, though it differs slightly from the “generic” system just described. With the Necsy system there is an additional fixed unit as well as the mobile. Test calls are initiated between the fixed and mobile units and several parameters are collected, including the signal level and signal quality of the radio links used in the calls.

Drive-testing systems are particularly well suited to solving problems in a localized area.

### 20.3 Evolution

The evolution of the network is largely dependent on customer requirements. Evolving the wireless system to increase the coverage area or capacity is driven by the customer needs in those areas. Additionally, measures to improve quality of service are driven by the customers’ interpretation of current service quality.

The range of tasks that have to be managed during the evolution phase is huge. From the addition of a single new base station to the migration to a dual-band system, all tasks have to be managed to ensure that they have a net benefit to the customer.

The work that takes place in this phase is generally definition and specification of the wireless system’s future needs. As such, the testing requirements are minimal.

Defining the future needs of the system typically is carried out in two ways:

- Specifying system enhancements through the standards body
- Communicating with vendors to ensure implementation of the enhancements

If a wireless operator has specific future needs from the system, the process would involve an attempt to standardize the feature through the appropriate standards body. Standardization eases the path towards implementation, because vendors usually are reluctant to implement proprietary features. Once the feature is standardized, the wireless operator is in a strong position to request the feature from any vendor.



---

Part

**6**

# **Basic Telecommunications Technologies**



# Transmission Media Characteristics and Measurements

**Ronald D. Lowe**

*Hewlett-Packard Company, Colorado Springs, Colorado*

## 21.1 Introduction

Transmission media, the interconnect pathway that ties together all associated communications hardware, must be considered an integral part of the communications network. The media, commonly referred to as *cabling*, is that part of the communications network that is buried in the ground, hidden behind walls, or run between floors at a relatively high installation cost. Swapping out a computer or workstation is fairly easy, but changing the wiring in a building is difficult and expensive. State-of-the-art building designs provide for communications network requirements and adhere to building wiring standards. [1–4]

This chapter reviews the two dominant transmission types, copper and fiber optic cabling. The copper cable types considered are coaxial shielded and unshielded twisted-pair. Fiber optic cable, a flexible strand of glass, is either of the single-mode or multimode type. These cable-specific types, along with their physical characteristics, are defined subsequently in this chapter. The transmission line theory for each cable type is discussed and referenced.

Wireless also is a means of communication. The phrase *wireless media*, or communication via radio waves through air or a vacuum, is confusing; “wireless cabling” is a contradiction in terms. The reader should refer to the specific wireless technology of interest to gain an understanding of its inherent transmission characteristics.

## 21.2 Copper Media

Copper media will be subdivided into three categories. First is twisted-pair copper wire. Electrically balanced, meaning the conductors are equally isolated from ground potential, it is available both in shielded and unshielded forms, referred to as STP

and UTP (shielded and unshielded twisted-pair, respectively). Second is coaxial cable, an electrically unbalanced media. Third is the general “other” category, where variations of the first two are discussed. Examples of “other” are twinaxial cable (“twinax”) and the single twisted-pair in full duplex mode.

### 21.2.1 Network copper cable

How well a cable functions in a particular network depends on the data transmission rate imposed upon it. Consider this example as a means of clarification. Fast Ethernet and 10Base-T both are CSMA/CD (Carrier Sense Multiple Access with Collision Detection) networks. Fast Ethernet, with its 100 Mbps (megabits per second) data rate is 10 times faster than 10Base-T. Both networks use a UTP cable that is balanced with 100 $\Omega$  characteristic impedance. Fast Ethernet cabling must be of a higher quality with respect to certain other transmission parameters than the cable for 10Base-T. Because the cabling infrastructure is not easily changed, the network designer should develop a cabling plan that includes future upgrades to faster data rates since, as noted, the cabling infrastructure is not easily changed (Table 21.1).

### 21.2.2 Transmission line theory for copper cable

The transmission characteristics of copper cable are best understood from the underlying transmission line theory. Any copper cable has four basic parameters. All copper wire has resistance  $R$  measured in ohms ( $\Omega$ ), and inductance  $L$  measured in henrys. Between any two conductors there is capacitance  $C$  measured in farads. Also between two conductors (since no insulator is perfect) is a parameter called conductance  $G$  measured in mhos. All other parameters can be derived from the four parameters  $R$ ,  $L$ ,  $C$ , and  $G$ . Cable manufacturers use the design criteria for these four parameters to develop cable suitable for communications networks.

**TABLE 21.1 Typical Networks and Associated Cables.** These cryptic names have been adopted by the standards organizations; some are acronyms, while others have evolved. The 10Base $x$  and 10Base-T networks are commonly referred to as Ethernet; the 100Base-TX has acquired the moniker “Fast Ethernet.” WAN and LAN refer to wide area and local area networks, respectively. As the cost of fiber optic cable comes down, many of the network types have fiber (as an alternative medium) being drafted into the standards—although the fiber-based FDDI has released a copper twisted-pair standard listed in the table as TP-PMD (twisted-pair, physical medium-dependent).

Network Name	Std. Ref.	Copper Media Type	Transmission Mode	Cable Characteristic Impedance
10Base2	IEEE 802.3	Coax	Unbalanced	50 $\Omega$
10Base5	IEEE 802.3	Coax	Unbalanced	50 $\Omega$
10Base-T	IEEE 802.3	UTP	Balanced	100 $\Omega$
Token-Ring	IEEE 802.5	UTP/STP	Balanced	100/150 $\Omega$
100Base-AnyLAN	IEEE 802.12	UTP	Balanced	100 $\Omega$
100Base-TX	IEEE 802.3u	UTP	Balanced	100 $\Omega$
WAN (voice)	IEEE 743	UTP	Balanced	600, 900 $\Omega$
WAN (T1 data)	ANSI T1	UTP	Balanced	135, 150 $\Omega$
FDDI TP-PMD	ANSI X3T9.5	UTP	Balanced	100 $\Omega$

Before embarking on a study of cable characteristics, a few statements about transmission lines are in order. The reader is advised to review a reference work on microwave, antenna, or transmission-line theory and associated measurements if these concepts are new. [5, 7, 10] The following three statements are general truths about transmission lines:

1. Electrical energy (a network data packet, for instance) travels at the speed of light in a vacuum and at less than the speed of light in other media, such as copper wire. This speed can be computed and is referred to as a *velocity of propagation*,  $V_p$ . As an analogy, observe the effects of a pebble dropped into a quiet pool of water. The ripples produced propagate away from the pebble at some measurable velocity.
2. A cable that is not properly terminated (proper termination will be defined) is subject to returning reflections of the incident electrical energy. Again using the pebble and pool analogy, if you observe long enough you will see reflections returning from the edge of the pool. The amplitude and timing of these reflections are mathematically predictable.
3. Under certain conditions it is possible that the returning electrical reflections will be of the proper phase and amplitude to the incident energy to make the ripples appear stationary; this is called a *standing wave* condition. The audible tone produced by striking a tuning fork is a suitable analogy for a standing wave or resonance condition.

We will now define the computed parameters, which will be followed by their respective equations. The quantity  $Z_0$  is known as *characteristic impedance*. Consider a long cable that is terminated at the far end in some impedance  $Z$ . If a signal (electrical energy) is applied at the near end, propagates through the cable to the far end, and is completely absorbed, then that cable is said to be terminated in its characteristic impedance  $Z = Z_0$ . Also known is that the cable has the same value of characteristic impedance  $Z_0$ . No energy will be reflected and there will be no standing waves. A cable terminated with its  $Z_0$  is said to be properly “matched” or “loaded.”


A shorthand way of referring to the velocity of propagation of electrical energy as a factor or percentage of the speed of light is the *propagation constant*, or  $\gamma$  (gamma). A signal that propagates on a cable at 70 percent of the speed of light,  $c$ , is said to have a  $\gamma$  of 0.7. With the above as a premise, the fundamental equations for cable are as follows:

$$Z_0 = \sqrt{\frac{R + j\omega L}{G + j\omega C}}, \quad \omega = 2\pi f, \quad j = \sqrt{-1} \quad (21.1)$$

$$\gamma = \sqrt{(R + j\omega L)(G + j\omega C)} \quad (21.2)$$

where:

- $R$  = resistance
- $L$  = inductance
- $G$  = conductance
- $C$  = capacitance.



**Figure 21.1 Cable-equivalent circuit.** This circuit is shown in an unbalanced configuration for the sake of simplicity. When the number of circuit elements approaches infinity in the limit, equation 21.1 is the result.

Note that the equations are dependent upon the frequency,  $f$ , of the signal. Any cable transmission line can be represented by an equivalent circuit made up of  $R$ 's,  $G$ 's,  $L$ 's, and  $C$ 's, as shown in Figure 21.1

The quality of the cable is directly related to how well the manufacturer can control the tolerances on the  $R$  and  $L$  for the copper wire and the  $G$  and  $C$  for the insulating dielectric. High-precision cable, i.e., that with a narrow tolerance range of  $Z_0$  and  $\gamma$  throughout its length, will cost more. As requirements for more bandwidth (higher data rates) and longer distances increase, it becomes more economical to change media, from copper wire to fiber optics. Fiber optics will be addressed subsequently in this chapter.

Another phenomenon of interest is the electrical wavelength of the signal and how it compares to our cable length. Equation 21.3 represents the relationship of velocity of propagation to frequency and wavelengths. Cables with lengths that are appreciably shorter than a quarter wavelength of the signals carried can be analyzed with conventional circuit analysis techniques. Cables that are approximately a quarter wavelength long, depending upon the termination value, can act more like antennas than transmission lines.

$$\lambda = \frac{V_p}{f} \quad (21.3)$$

where:

- $V_p$  = velocity of propagation of signal in meters/second
- $f$  = frequency of source
- $\lambda$  = wavelength in meters.

Transmission line theory is used to analyze cable characteristics when the cable lengths of interest are appreciably longer than a quarter wavelength of the frequency of interest. Cable installations for data communication networks should be considered as multiple wavelengths for analysis purposes. This claim is substantiated with the following calculation. Most cable has a  $V_p$  of 0.7 to 0.8 of the speed of light. This is determined primarily by the insulating dielectric used in the cable. For 10Base-T



with a 10 Mbps data rate (10 Mbps has a fundamental frequency of 10 MHz), the following applies:

$$\begin{aligned}
 \lambda &= \frac{V_p}{f} & (21.4) \\
 &= \gamma \times \left( \frac{c}{f} \right) \\
 &= \frac{(0.7 \times 3 \times 10^8)}{10 \times 10^6} \\
 &= 21 \text{ m (68 ft)}
 \end{aligned}$$

where  $c = 3 \times 10^8$  mps, the speed of light.

A quarter wavelength in 10Base-T cable application is approximately 5.2 m (17 ft). For T1 networks at 1.544 Mbps,  $\lambda = 136$  m (446 ft) and  $\lambda/4 = 34$  m (111 ft). These quarter-wavelength values of 17 ft for 10Base-T and 111 ft for T1 are significantly shorter than typical cable lengths employed by the respective technologies.


The foregoing is an oversimplification since a data signal is a stream of electrical pulses made up of many harmonics of a fundamental frequency. The mathematics requires Fourier analysis for an accurate assessment of the effective wavelength of the signal in a cable. Such an exercise will not negate the simplifying assumption that network cable plants should be treated as multiwavelength systems. The primary purpose of the preceding discussion is to emphasize the importance of having proper and accurate terminations for cable in network applications. Indeed, as higher-speed networks are being deployed, greater than 100 Mbps, the problem with the continued use of copper is that the cable tends to be a better antenna than a transmission line.

Figure 21.2 shows the cross-section of two types of copper cable, along with their associated electric and magnetic fields. It can be observed readily that the electric and magnetic fields of UTP cable can interfere with adjacent cable pairs. This susceptibility to interference must be weighed against the lower cost of UTP compared to STP or coaxial cable. The twist in UTP helps minimize the susceptibility to noise interference. Coaxial cable has the additional problem of being unbalanced and susceptible to externally imposed ac and dc ground currents.

### 21.2.3 Twisted-pair cable

Consider a two-pair cable of UTP, such as depicted in Figure 21.3, in a network in which one of the pairs is designated as transmit and the other pair as receive. Coupling between pairs, called *crosstalk*, limits both the data rate and length for acceptable error rates in network use.


Cable manufacturers have made significant reductions in crosstalk by bundling wire pairs with different twist rates. Each of the four pairs in a Category 5 LAN UTP cable, for example, will have a unique number of twists per foot. (The “5” in “Cat 5” is a quality grading established by the 568A Standards body, listed in the chapter



**Figure 21.2 Media fields.** The fields surrounding the UTP cable extend out past the physical boundaries of the wire. These fields will impinge upon adjacent cable pairs, causing interference and noise; likewise, fields in adjacent pairs will cause interference and noise in this pair. Coaxial cable is self-shielding; the fields are fully contained within the outer shield. UTP is used where economic factors and/or the need for a balanced mode outweigh the advantages of coax.



**Unshielded Twisted-Pair Cable (UTP)**



**Shielded Twisted-Pair Cable (STP)**

**Figure 21.3 UTP and STP cables.** The UTP cable depicted is the most common 4-pair variety, though 2- to 24-pair bundles also are quite common. STP cable exists most often in 2-pair configuration with shield, allowing a transmit and receive pair to be shielded as a unit. (Courtesy Mohawk/CDT)

references.) The wire standards specify the acceptable crosstalk and the cable manufacturers will vary the twist rates and use other proprietary schemes to meet the standards. Shielding the individual pairs as depicted in Figure 21.3 is another means of improving the quality of the cable.

The parameters of interest for twisted-pair cable are six in number; they can be determined from three measurements. These measurements are described briefly in a more thorough discussion in Chapter 25 of this book.

- $Z_0$ , the characteristic impedance
- *Length*, the physical length of the cable
- *Crosstalk*, the coupling between adjacent pairs in the cable
- *Resistance*, the dc resistance of the copper conductors
- *Attenuation*, the signal loss at a specified frequency
- *Wire map*, the connector-to-connector wiring by pin numbers

Length and  $Z_0$  are measured most easily with a TDR (time domain reflectometer), a special-purpose instrument that transmits a calibrated pulse of electrical energy into a terminated cable and then detects and measures the time of any returned or reflected energy. As mentioned previously, if no energy is reflected, then the cable's characteristic impedance matches the value of the termination.

With a calibrated termination one knows, by substitution, the  $Z_0$  of the cable. If the cable is left unterminated, i.e., an open circuit, the energy is reflected back to the TDR. The round-trip time divided by two (since the pulse traverses the cable twice) is used by the TDR to compute the cable's length:

$$\text{Length} = V_p \times \left( \frac{\text{time}}{2} \right) \quad (21.5)$$

where  $V_p$  is specified by the cable manufacturer.

Attenuation and crosstalk are determined by using a frequency signal generator and an ac voltmeter. All measurements must be made on properly terminated cable to avoid the standing wave problem. Attenuation is the amount of signal lost or dissipated by the cable. The measurement is frequency-dependent, so the measurements are made repeatedly over a wide frequency range. The standards for each specific network will specify the required frequency range. Table 21.2 is a sample of the specifications.

**TABLE 21.2 Sample of Media Specifications.**

Network	Maximum Length	Maximum Attenuation	$Z_0$	DC Resistance
10Base-T	100 m	11.5 dB @ 10 MHz	100 $\Omega$	9.2 $\Omega$
Token-Ring	100 m	10 dB @ 16 MHz	100 $\Omega$	28.6 $\Omega$
T1	3200 m	7.8 dB @ 1.0 MHz	135/150 $\Omega$	25.3 $\Omega$

The attenuation is the level of signal measured at the termination divided by the signal measured at the source, is represented as the logarithm of this ratio (see footnote below), and is expressed in the unit decibels, dB. For instance, if half the signal is dissipated in the cable, then

$$\begin{aligned}\text{loss} &= 20 \log_{10} \left( \frac{1}{2} \right) \\ &= -6.02 \text{ dB}\end{aligned}\tag{21.6}$$

Crosstalk is the amount of signal from the transmit pair that “leaks” into the receive pair due to imperfect cable. In this case, the signal generator is on one end of a terminated cable on the transmit pair and the ac voltmeter is on the same end but attached to the receive pair. For a perfect cable with no coupling between the pairs, the voltmeter would read 0 V. A 1 percent level of signal leakage represents  $20 \log_{10}(1/100)$ , or  $-40$  dB of isolation or crosstalk. This measurement is called NEXT or *near-end crosstalk*, since both measuring devices are at the same (near) end of the cable.

Resistance and wire-mapping measurements are performed with an ohmmeter. The resistance measurement is straightforward if the cable is still on the manufacturer’s reel and both ends are available. An installed cable’s ends usually are separated by great distances, however, a loopback measurement can be performed. One simply ties together temporarily the far-end pairs of wires. At the near end, the measurement then can be made on the pair of wires and the reading divided by two to obtain the one-way resistance. Wire mapping is the exercise of identifying which wires in the cable are connected to which pins on the standard 8-pin modular jack (RJ-45) or LAN connector used in UTP networks. The 8-pin jack is referred to as having four pairs of pins, as depicted in Table 21.3

Note that the 10Base-T cable is referred to as a *straight-through* cable in that pair 2 at one end is connected to pair 2 at the other end. The Token-Ring and T1 cables are crossover cables, where the Tx (transmitter) connects to Rx (receiver) and vice versa. There are variations of these customs, making it prudent to document and verify the connectivity of cabling to network components. The wire-mapping measurement uses an ohmmeter to verify that the cable wiring requirements have been met for the specified network (Table 21.4).

---

#### Attenuation Level

In this usage, signal level is assumed to be a voltage measurement. The equation comes from simplifying the power equation:

$$\begin{aligned}\text{Because } P &= v^2/Z, \text{ therefore} \\ 10 \log_{10} \frac{P_2}{P_1} \\ 10 \log_{10} \frac{\left( \frac{V_2^2}{Z} \right)}{\left( \frac{V_1^2}{Z} \right)} &= 10 \log_{10} \left( \frac{V_2}{V_1} \right)^2 \\ &= 20 \log_{10} \left( \frac{V_2}{V_1} \right)\end{aligned}$$

Stated in words, the 3 dB half-power point is a 6 dB half-voltage point in constant impedance.

**TABLE 21.3 Pair-Naming Standards for 8-Pin Jack.** The 8-pin plug and jack, referred to by its designation RJ-45, has become the connector of choice for UTP wiring systems. This also makes a visual differentiation from the standard modular telephone jack, which usually is a 4-pin or 2-pair wiring system.

Pair No.	Pins
1	4 & 5
2	1 & 2
3	3 & 6
4	7 & 8

**TABLE 21.4 UTP Pairings for Cable Media.** This table shows the cable by pair numbers; cabling documents will translate the pair numbers to a color code, the colors depending on the specific vendors' color designations. Most common is a solid color paired with the same color plus a contrasting stripe, such as green and green/white. This technique allows the installer to maintain the pair integrity of the cable.

	One End		Other End	
	Tx	Rx	Tx	Rx
10Base-T	2	3	2	3
Token-Ring	3	1	1	3
T1	2	1	1	2
FDDI TP-PMD	2	4	4	2


### 21.2.4 Coaxial cable

Coaxial cable is a well-behaved medium for networking (Figure 21.4). The cable's impedance can accurately be controlled, usually to  $50\Omega$ , and it will carry very high frequencies. Its primary disadvantages are that it is physically bulky (a single-conductor strand is about the same diameter as a four-pair UTP bundle), is more fragile, and is more expensive. The current frequency limit for UTP is about 150 MHz, whereas coaxial cable will work well to several hundred megahertz over respectable distances. Close attention must be paid to grounding the shield to satisfy a requirement by building codes.[3]


The measurements for coaxial cable are the same for twisted-pair cable, but crosstalk and wire mapping have no meaning in a coax environment.

- $Z_0$
- Length

## 462 Basic Telecommunications Technologies



**Figure 21.4 Coaxial cable.** In its simplest form coaxial cable will have a single shield around a center conductor. Multiple outer shields can be used for further improvement in field isolation and high-frequency performance. (Courtesy Mohawk/CDT)



**Figure 21.5 Twinaxial cable.** Some grades of twinaxial cable might provide a twist in the two center conductors. Twinax provides a high-quality transmission path, but at a relatively high cost.

- Resistance
- Attenuation

### 21.2.5 Other copper cable

As noted at the outset of this chapter, the two main types of “other” copper cable are twinaxial (“twinax”) and single-pair full-duplex.

**Twinaxial copper cable.** Twinaxial cable is a legacy cable used primarily in mainframe computer rooms to connect the various peripherals (usually terminals) to the computer, front-end processor, or cluster controller (Figure 21.5). Twinax looks like coax except that it has two parallel conductors instead of one, somewhat like STP but with a single pair of wires embedded in the dielectric and supported by an outer shield. It is relatively expensive and has not been specified for deployment by the LAN or WAN standards bodies. It does an excellent job of providing for high-speed data connections and is found in proprietary installations.

**Single-pair full-duplex (POTS cable).** This type of cable is UTP but with full-duplex (simultaneous transmit and receive) capability over a single pair of wires. Referred to as *POTS* (“plain old telephone service”) cable, its use has been limited primarily to analog voice-band service applications. Given that average consumers and small businesses have buildings wired with this cable for telephone service, the carrier companies have developed schemes to use this same installed cable for digital subscriber loop service. This provides increased bandwidth without requiring the installation of new cable.

## 21.3 Optical Fiber Media

Fiber optic cable media are widely deployed in higher-rate (>100 Mbps) data communication channels, and also where long distances between repeaters are the norm. The cable fibers, in addition to providing very wide bandwidth, are free from electromagnetic interference and offer high security against intrusion. In its simplest form an optical fiber is a flexible glass filament that carries a modulated light signal.

(Note that the term *light* no longer means that the signal consists of visible light; it is a legacy term carried over from the early days of fiber optic transmission. Today the term *light* means that a signal is present or the fiber is in service, while *dark* refers to a spare fiber that is not in service.)

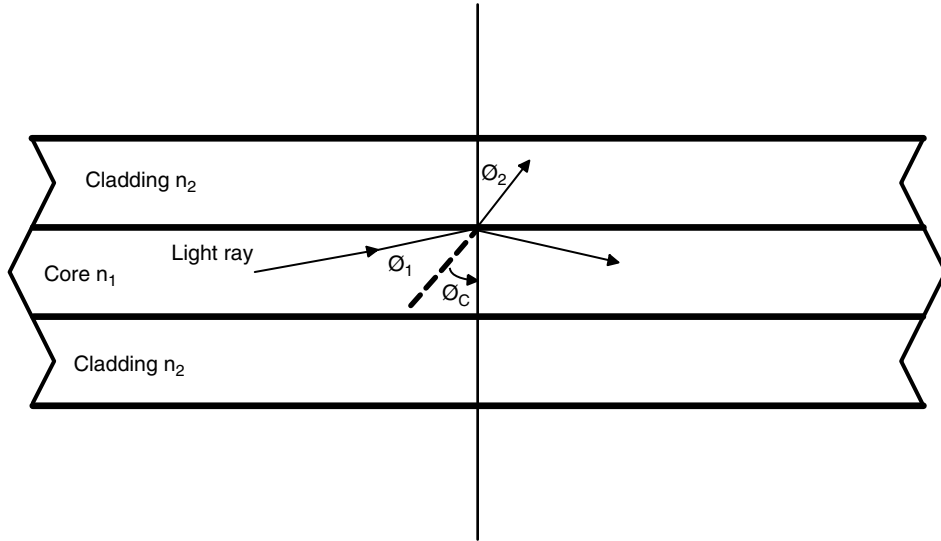
Fiber optic signals are more commonly referred to by their wavelength, expressed in nanometers (nm). The longest visible wavelengths, at the bottom of the red portion of the spectrum, are in the 600–700 nm range. Fiber optic signals normally are infrared; 900 and 1300 nm are the most common wavelengths. This section provides a description of fiber optic media and will review the two most common types deployed in data communications, multimode and single-mode fiber. For purposes of this discussion, the terms *light* and *signal* will be used interchangeably.

### 21.3.1 Fiber optic media technology

Propagation of light in fiber optic media is best understood from waveguide theory. Energy is being transmitted as an optical wave in a cylindrical waveguide made of a glass dielectric. How the signal propagates is a complex issue that is dependent upon the wavelength, dimensions of the fiber, and method of signal injection. [8,9] The references listed at the end of the chapter can be consulted for detailed discussions of how different wave equations can be used to describe the various modes of light propagation.

More important to this chapter is some discussion about the different *cladding* methods used in the cable and some of the loss and distortion mechanisms in this medium. The key to keeping the transmitted light inside the fiber, i.e., preventing it from leaking out the side where the cable bends, is to control the index of refraction of the media core and the boundary cladding.

The *index of refraction* of a material is defined as the ratio of the speed of light in a vacuum to the speed of light in the material, or, in our case, glass fiber. Figure 21.6 depicts a ray of light that is incident upon a boundary between two materials



**Figure 21.6 Light ray incident upon cable cladding.** The rays incident upon the cladding boundary are shown most easily in a planar diagram (though the fibers are, of course, cylindrical). It is correct to assume that if the cable is subjected to a sharp bend, light will escape when the critical angle criterion is violated. This is not a common occurrence in practice, however, because such a sharp bend is likely to fracture the glass and render the cable useless.

with different indices of refraction. How the incident ray is transmitted through or reflected from the wall of the fiber core is determined by Snell's Law:

$$n_1 \sin \theta_1 = n_2 \sin \theta_2 \quad (21.7)$$

where  $n_1$  and  $n_2$  are indices of refraction and  $\theta_1$  and  $\theta_2$  are the respective angles of ray incidence.[8]

By rewriting equation 21.7 into the following form, we can find a special condition called the *critical angle*.

$$\sin \theta_1 = \left( \frac{n_2}{n_1} \right) \sin \theta_2 \quad (21.8)$$

This is where  $\theta_2 = 90^\circ$  and equation 21.7 becomes  $\sin \theta_c = n_2/n_1$ , where  $\theta_c$  is the critical angle for total reflection. This means that a ray of light that is incident upon a boundary layer in the glass fiber will be totally reflected back into the fiber if the angle of incidence  $\theta_1$  is greater than  $\theta_c$ . The cladding or outer coating in a glass fiber is designed specifically with an appropriate  $n_2$  different from  $n_1$  in the core so that the light does not leak out the sides of the fiber.

There are two popular types of outer cladding used in fiber optic cabling, the *step-index* and the *graded-index*. Figure 21.7 shows a representation of fiber media and the refractive index vs. radius of fiber for the two types of cladding.

Referring to Figure 21.7, one must realize that the core, cladding, and primary coating all are of the same flexible glass material. During manufacture, an extruding

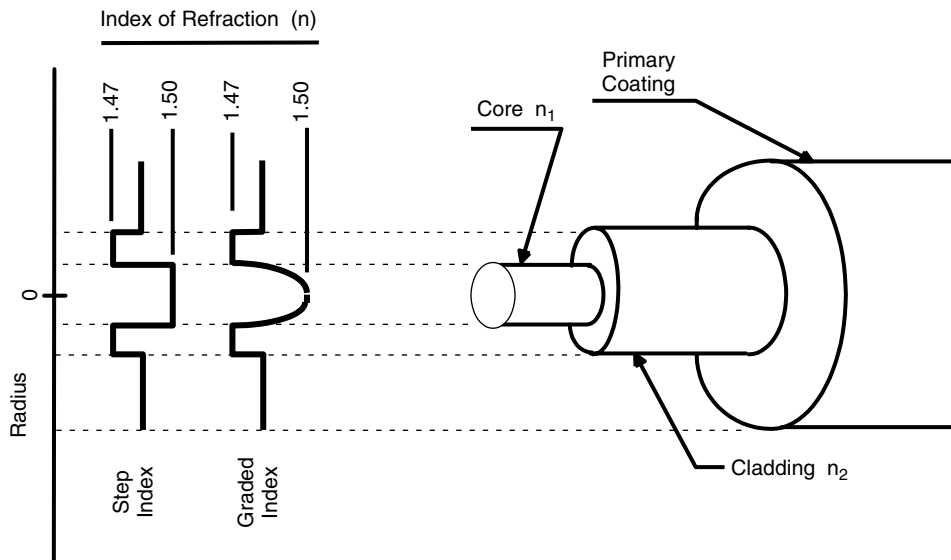


process, the glass is doped with materials that cause the change in index of refraction, thus creating the boundaries. Since glass is a silicon oxide, the doping materials also are oxides, but of other metals such as germanium, boron, and phosphorous. The difference between the step- and the graded-index is related to how the doping material is introduced. Early fibers were of the step-index type, but as manufacturing processes improved, graded-index fiber became available. Graded-index fiber has less light leakage at the boundary than does step-index, but is more expensive.


### 21.3.2 Multimode fiber

*Multimode fiber* takes its name from the fact that the transmitted light can transverse the media in many modes. Consider for a moment a microwave signal that is made up of electric and magnetic fields. [7] Each possible orientation of these fields with respect to one another and to the propagation of the combination constitutes a *mode*. Light transmission in fiber can be thought of in analogous terms. [8] Modes can be controlled by both the dimensional characteristics of the fiber and the quality of the energy source. A *noncoherent* (lacking spectral purity) light source transmits a multiple-mode signal (Figure 21.8).

Multimode transmission is a compromise between distortion and signal loss versus cost effectiveness in a wide-bandwidth media system. Such systems have been designed and currently are the dominant fiber optic media deployed for local area networks. The modulation technique for fiber is simply light on and light off. The multimode condition gives rise to a distortion called *dispersion* in this type of fiber. Since each transmission mode propagates along a different path, at the end of the



**Figure 21.7 Fiber media cladding types.** The graded index cladding causes the critical angle to change as a function of radius. The nature of this change improves the reflection ability of the fiber at the cladding boundary.



**Figure 21.8 Light ray in multimode fiber.** For the sake of simplicity the light ray is shown entering multimode fiber from a point source, but in reality it originates from an infinite number of points and could follow an infinite number of paths. There tend to be dominant modes, however, which makes this type of transmission medium technically and economically viable. (Check the end-of-chapter references for further information on modes of propagation.)

cable the light pulse will suffer broadening or smearing. This limits the fiber's usable distance and bandwidth. Manufacturers use a graded cladding technique to optimize the fiber performance for specific application bandwidths per customer needs.

Multimode fiber is specified by the diameters in microns of the fiber core and cladding. For example, "62.5/125" fiber cable has a core diameter of 62.5 microns and a cladding diameter of 125 microns. The light sources for multimode fiber are infrared LEDs (light-emitting diodes).


Other losses in this type of fiber come from connecting, splicing, and reflections. Imperfect connecting and splicing are obvious losses. Reflections can be caused by impurities in the fiber and by damage from excessive bending or flexing of the cable. Losses can be detected by using a fiber optic power meter. Every fiber installation should be accompanied by a documented end-to-end loss measurement per fiber. This becomes a benchmark for that fiber and is used by the network designer to determine the acceptability of the fiber for a specific segment. Reflections, cable damage, and length are best measured with an OTDR (optical time domain reflectometer).

### 21.3.3 Single-mode fiber

Single-mode fiber takes its name from its dimensional design and light source, which allows only one mode of light transmission (Figure 21.9).

Single-mode design is most nearly like the traditional single-mode waveguide transmission systems for microwave radio. With a core diameter on the order of 8 microns, this fiber and its associated hardware tend to have a raw material cost an order of magnitude higher than does multimode cable. Sources of coherent light are required, which today are laser diodes.

Single-mode is used where long-haul requirements are the dominant design factor. Current applications cover the long distances in WANs, MANs, and all undersea communication cable installations. Lengths of up to 20 km can be achieved between repeaters. The cost penalty is offset by the much wider bandwidths available and by the reality that the only alternatives for the required distances are microwave or satellite links. Measurements are the same as for multimode, namely power, loss, and reflections. Instruments are power meters and OTDRs with single-mode interfaces.



**Figure 21.9 Light ray in single-mode fiber.** This simplistic view of single-mode propagation has the ray of light confined to a centered, noninterfering flight through the cable. Single-mode cable is not as widely used as multimode because it is difficult and expensive to manufacture the cores so small, and because laser technology is required to generate a coherent light ray to propagate through the core.

## 21.4 References

- EIA/TIA 568A, *Commercial Building Telecommunications Wiring Standard*. (Washington, D.C.: Electronic Industries Association.) [1]
- EIA/TIA 569, *Commercial Building Pathway and Space Standard*. (Washington, D.C.: Electronic Industries Association.) [2]
- EIA/TIA 607, *Commercial Building Grounding and Bonding Standard*. (Washington, D.C.: Electronic Industries Association.) [3]
- EIA/TIA TSB-67, *Services Bulletin on Testing of Category 5 Cable*. (Washington, D.C.: Electronic Industries Association.) [4]
- Johnson, Walter C. *Transmission Lines and Networks*. (New York: McGraw-Hill, 1950.) [5]
- \_\_\_\_\_. *Transmission Systems for Communication, 4th Ed.* (Murray Hill, N.J.: Bell Telephone Laboratories, 1971.) [6]
- Jordan, E.C., and Balmain, K.G. *Electromagnetic Waves and Radiating Systems, 2nd Ed.* (Englewood Cliffs, N.J.: Prentice-Hall, 1968.) [7]
- Yeh, Chai. *Handbook of Fiber Optics: Theory and Applications*. (San Diego: Academic Press, 1990.) [8]
- Tosco, Federico. *Fiber Optic Communications Handbook, 2nd Ed.* (Blue Ridge Summit, Pa.: TAB Books, 1990.) [9]
- Coombs, Clyde F. Jr. *Electronic Instrument Handbook, 2nd Ed.* (New York: McGraw-Hill, 1995.) [10]



## Fiber Optic Network Elements

**Waguih Ishak**

*Hewlett-Packard Laboratories, Palo Alto, California*

### 22.1 Introduction

Fiber optic communication links are considered the backbone of the global information network, with more than 15,000,000 km of fiber installed worldwide. These fiber optic links traditionally have been used for connecting major metropolitan areas, such as Los Angeles to San Francisco and New York to Chicago. They also have been used for undersea connections such as the transatlantic (6000 km) and transpacific (9000 km) links. More recently, fiber optic links have found wide use in the *access network* between a major city and its suburban towns, and also within residential areas (fiber to the curb and fiber to the home). One of the fastest-growing areas for fiber optic are LANs within an enterprise or a building, where multimode fiber can be used for distances below about 2 km.

#### 22.1.1 Capacity

State-of-the-art fiber optic systems operate at 2.4 Gbps and 10 Gbps. For comparison, note that one 10 Gbps link can carry 160,000 telephone calls (each requiring 64 kilobits bandwidth)! At the same time, research laboratories around the world are developing multi-Gbps components and systems that have potential for terabit per second communications links, expected by the late 1990s.

With more users and computers connecting via the Internet, there will be increasing demands on the bandwidth and data rates of global and metropolitan communications networks. Higher-capacity long-haul networks (>10 Gbps) are needed. This can be accomplished by increasing the modulation rates of the optical sources and using time division multiplexing (TDM) techniques, or by developing wavelength division multiplexing (WDM) networks capable of carrying multiple wavelengths on the same fiber, each operating at multi-Gbps data rates.

Beyond the increase in fiber optic network data rates, the performance of the devices, components, and subsystems used in such networks is improving at a rapid rate. For example, a semiconductor laser used in an optical amplifier system must have a mean time between failure (MTBF) of more than 100,000 hours (Bellcore specifications). If the amplifier is used in submarine cable (transatlantic or transpacific), the laser must be sufficiently reliable to withstand extreme operating conditions such as temperature, humidity, and pressure.

In addition to developing high-performance components, the trend continues toward lowering the effective cost of information per bit and per mile. As a result, designers of lightwave devices, components, and subsystems are faced with a challenge. They need to maximize the performance of each system building block, minimize the adverse interactions between these blocks—and at the same time design for manufacturability and cost-effectiveness.

Lightwave instrumentation is important in helping the designers of optical sources (lightwave signal analyzers), optical components (lightwave component analyzers), and systems and links (optical time domain reflectometers). Further, other instruments such as tunable laser sources, power meters, and communication analyzers are useful tools to test and characterize lasers, photodetectors, amplifiers, and optical signals.

The small size, large bandwidth, and very low attenuation of optical fibers make them attractive as alternatives to conventional copper cables in telecommunications applications such as telephone and CATV systems, and in data communications applications such as computer interconnection.

### 22.1.2 Basic fiber optic system

A basic communications system consists of a transmitter, a receiver, and an information medium (Figure 22.1). The transmitter is a system capable of generating the information to be sent to the receiver. The information can be in analog or digital format. The transmission medium carries the information over some distance and delivers it to the receiver. The receiver interprets the information and transforms it into an accessible form.

In the case of the fiber optic communications system, the transmitter is an optical source, either a laser or light-emitting diode (LED), which is modulated with the information to be transmitted. Optical sources can be modulated internally by varying their operating currents, or externally by using an external optical modulator fol-


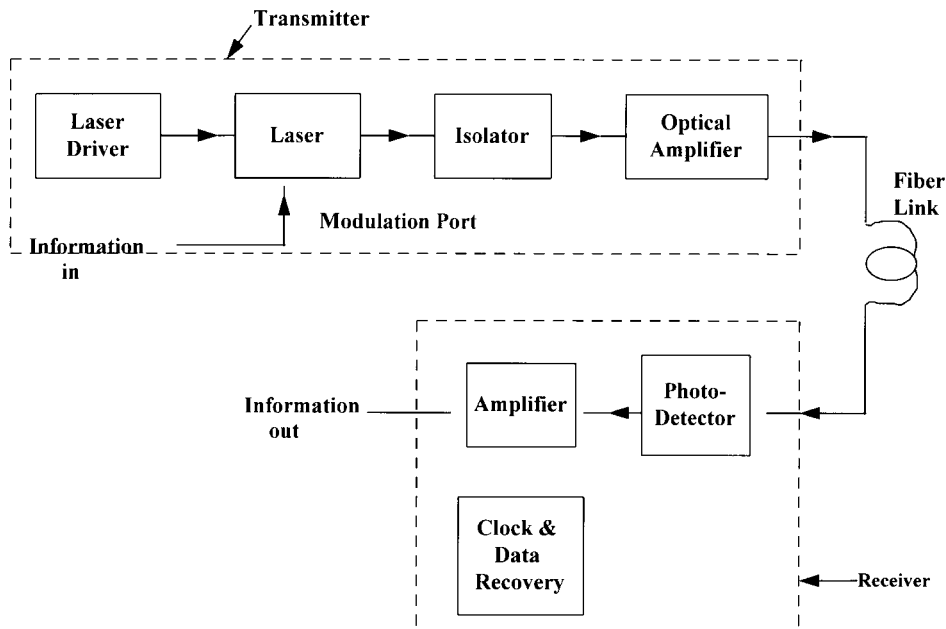


Figure 22.1 Simplified block diagram of a fiber optic communications system.



**Figure 22.2** Expanded block diagram of a basic fiber optic communications system.

lowing the optical source. The transmission medium is an optical fiber and the receiver is a photodetector.

### 22.1.3 Fiber optic system components

An actual fiber optic communications network uses many more electrical and optical components than mentioned in the preceding section. Examples include optical isolators, optical polarization controllers, fiber couplers, optical filters, optical connectors, electronic demodulators and amplifiers, and electrical signal generators (Figure 22.2). Additionally, a fiber optic network will contain many computers and software programs to control, manage, and maintain the network.

State-of-the-art fiber optic communications systems use semiconductor lasers as the optical source. If the data rate is below a few Gbps, the lasers can be modulated by varying their bias current. If the data rate exceeds 10 Gbps, however, an external optical modulator such as a lithium niobate Mach-Zehnder modulator (see section 22.1.6) is coupled to the laser source. The fiber can span long distances: The transatlantic fiber link, for example, is more than 6000 km.

Repeaters will be needed (every 50 km of span, for example) to boost the signal level above the noise. In most of the installed fiber systems, these repeaters are electronic devices that transform the optical signal to an electrical signal using a photodetector (O/E converter), amplify the electrical signal using electronic amplifiers, and then convert the electrical signal back to an optical signal using a laser (E/O converter). These repeaters limit the data rate on the fiber because of limited bandwidth on the electronic amplifiers.

At the end of the communications channel, the receiver consists of a photodetector, which converts the optical signal to electrical signal, followed by an amplifier to boost the signal level, and a demodulator to decode the information. In such a system, the information determines the data rate and the laser wavelength determines the carrier frequency. For example, a typical system will use a laser operating at 1300 nm wavelength, corresponding to a carrier frequency of about 176 THz (176,000 GHz), and the information will modulate the laser at a rate of 560 MHz.

The invention of optical fiber amplifiers has changed the way fiber optic communications systems are configured. Commercially available erbium-doped fiber amplifiers (EDFAs) can boost an optical signal level by 30 dB or more over a significantly wide band (1530–1570 nm). The use of EDFAs, as signal boosters at the transmitter and as preamplifiers across the fiber and at the receiver, should eliminate many of the repeaters needed in the conventional systems described previously.

This fact is significant because the existing repeaters do not allow the option of varying the data rate over the communications system. EDFAs, with their extremely wide band (about 30 nm), will make the fiber optic communications system upgradable without major changes in the configuration. In addition, since the EDFA can be coupled easily to the transmission fiber, it will be possible to use a large number of EDFAs in the system, allowing longer transmission distance. AT&T has demonstrated a 9000 km undersea system using 300 EDFAs uniformly spaced along the fiber length.

## 22.2 Time Division Multiplexed Networks

The most commonly used architecture for fiber optic networks is called *time division multiplexing* (TDM), in which several channels carrying different information are multiplexed in time and then propagated on the fiber (Figure 22.3). At the receiver end, the channels are demultiplexed and processed to retrieve the information for each channel.

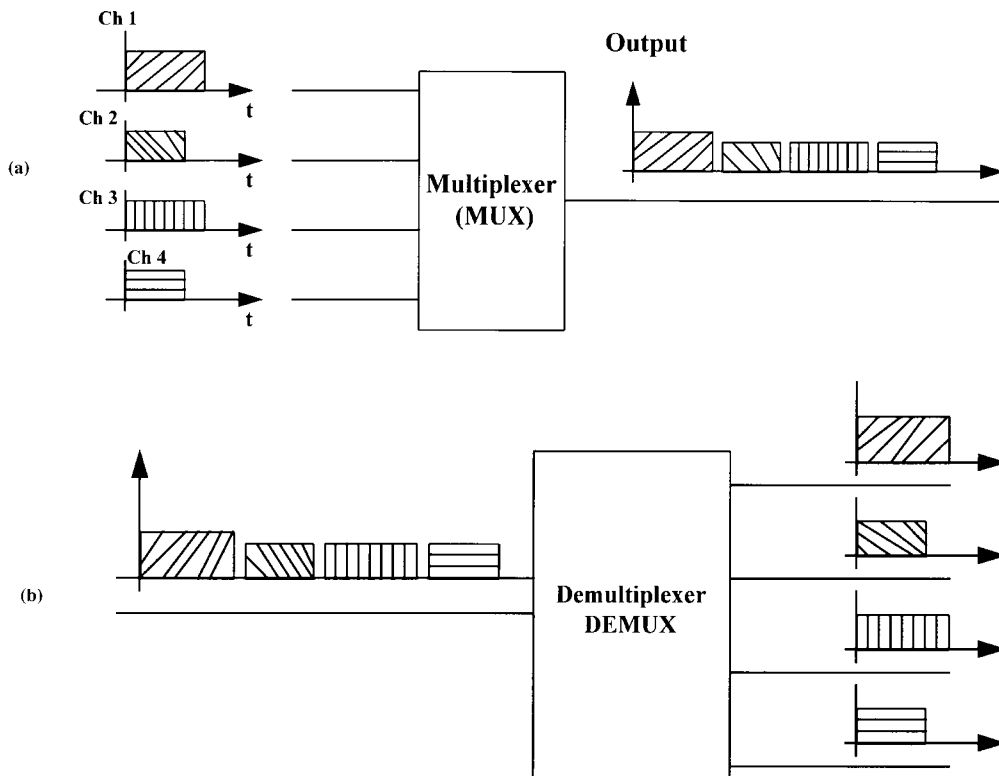
In the case of voice information, for example, a telephone call requires 64 kilobits of bandwidth. Thousands of telephone calls are multiplexed at the local and regional switching offices and combined as a one bit stream of data at a very high data rate, equal to the product of the number of calls times 64 kbps. (For example, 40,000 telephone calls will result in about 2.5 Gbps aggregate data rate).

When this bit stream reaches the fiber optic transmitter, it is used to drive the modulator port of the transmitter. The laser source therefore will be modulated at the same rate, and the optical bit stream is sent on the fiber toward the receiver. As the optical signal arrives at the receiver end, a photodetector converts it into an electrical signal with the same data rate. A demultiplexer retrieves and separates the original channels, redirecting them to their appropriate destinations.

Although this example used voice data to illustrate how TDM networks work, the same principle is used to transmit other signals, such as video and computer data. The important aspect of TDM is the fact that each channel is allowed a certain time-slot on the fiber.

Almost all of the installed base of fiber optic networks uses TDM with data rates of 45, 155, 622, and 2400 Mbps. Starting in 1995, new TDM networks were installed at 10 Gbps, allowing more than 160,000 telephone calls (or 1000 textbooks of 500





**Figure 22.3** Time division multiplexing (TDM) transmitter (a) and receiver (b).

pages each) to be transmitted on the same fiber. As the demand on bandwidth increases, higher-rate TDM networks will be needed. Although lightwave devices have been demonstrated to operate at extremely high data rates (well in excess of 50 Gbps), it is extremely difficult to design and build electronic circuits (such as laser drivers and receiver amplifiers) at speeds above 10 Gbps.

### 22.3 Wavelength Division Multiplexed Networks

Another important development has been the emergence of *wavelength division multiplexing* (WDM) as a means for increasing the capacity of fiber optic communications systems. In WDM systems, several TDM channels operating at different optical wavelengths are optically multiplexed and sent on the same fiber (Figure 22.4). Four 2.5 Gbps TDM channels, operating at 1535, 1538, 1541, and 1544 nanometers, could be combined on a single fiber—with even more channels operating at other wavelengths. (The fiber actually can carry thousands of these 2.5 Gbps channels before it reaches its bandwidth limit.)

The main advantage of WDM networks is that they have a high overall data rate (10 Gbps in the preceding example) while using lower-frequency electronic circuitry for each channel. In addition, the wavelength of each channel can be used as a pa-

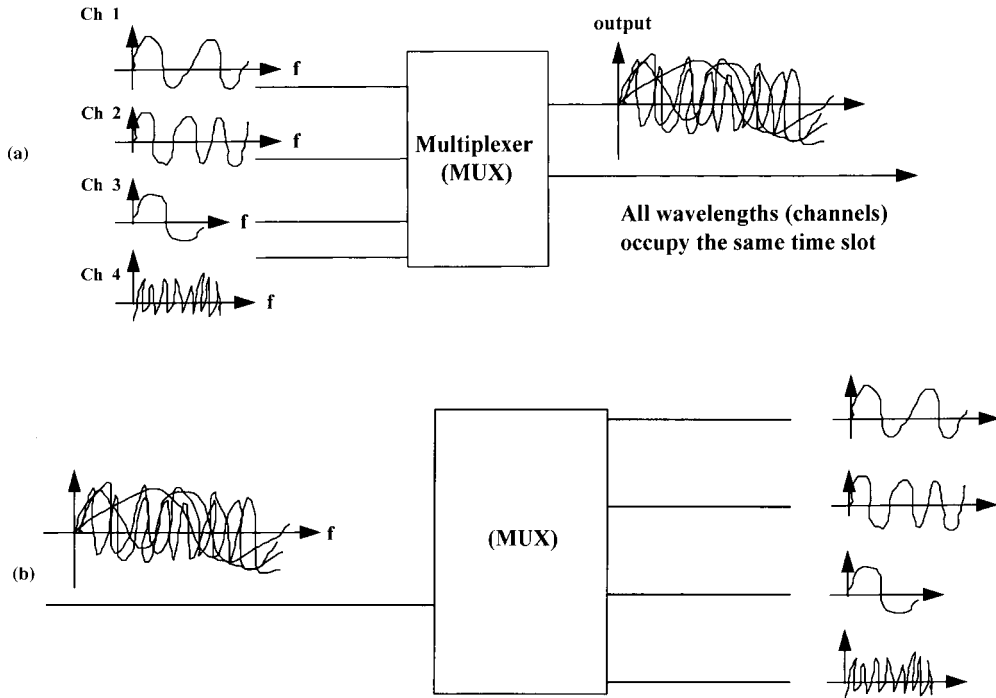


Figure 22.4 Wave division multiplexing transmitter (a) and receiver (b).

parameter to route the channel to a specific destination, providing for an important degree of freedom in future telecommunications and data communication networks. In 1996, several WDM fiber optic networks were in operation in the United States and Europe, and many more were planned.

In a WDM system, the transmitter consists of several laser sources operating at different wavelengths. The laser outputs are modulated and multiplexed for transmission over the same optic fiber. At the receiver, a demultiplexer separates the channels for the appropriate photodetector circuit, as shown in the lower part of Figure 22.4. In a typical system, single-mode semiconductor lasers are used with wavelength separation of a few nanometers (corresponding to hundreds of GHz) and EDFAs are used as signal boosters along the fiber. It is expected that future systems will use channel separation of a few GHz.

The design of the communications systems relies heavily on the characteristics of the lightwave signal propagating through the various components of the system. A lightwave signal is an electromagnetic wave in the wavelength range between 200 and 2000 nm, i.e., a frequency range between 150 and 1500 THz. The transmission section of the system generates this signal, which passes through many optical components for modulation, amplification, and coupling before it reaches the optical fiber. When a WDM system is used, the transmitter generates multiple signals, each

at a different wavelength, modulates each signal independently, amplifies the signals, and finally multiplexes the signals and couples them to the same fiber.

Once inside the fiber, the signal (single wavelength or multiple wavelengths) propagates over many kilometers before its amplitude is reduced to the point at which amplification is needed, typically after about 50 km. State-of-the-art repeaters are fiber-based erbium-doped amplifiers with more than 30 dB gain. In a point-to-point link, the signal moves from the transmitter to the receiver without rerouting or switching, as shown in Figure 22.5 part (a). In switched systems, optical or electronic switches are used to route the signal(s) to various branches within the fiber network and eventually to the receiving end, as shown in Figure 22.5 part (b).

After propagation through the fiber, the signal reaches the receiving end, where the necessary information is extracted. In a TDM system, a photodetector converts the optical signal into an electrical signal, which is processed further by an electronic amplifier and other electronic circuitry such as filters and discriminators. If a WDM system is used, the multiple-wavelength signal is first demultiplexed into its components, and each component passes through a photodetector, an amplifier, and other electronic circuits.

The sections that follow describe the basic elements of a fiber optic network: transmitters, amplifiers, distributors (multiplexers and demultiplexers), and receivers. The focus will be on devices and components used in the high-speed networks (faster than 1 Gbps) that are considered to be the backbone of the global information network infrastructure.

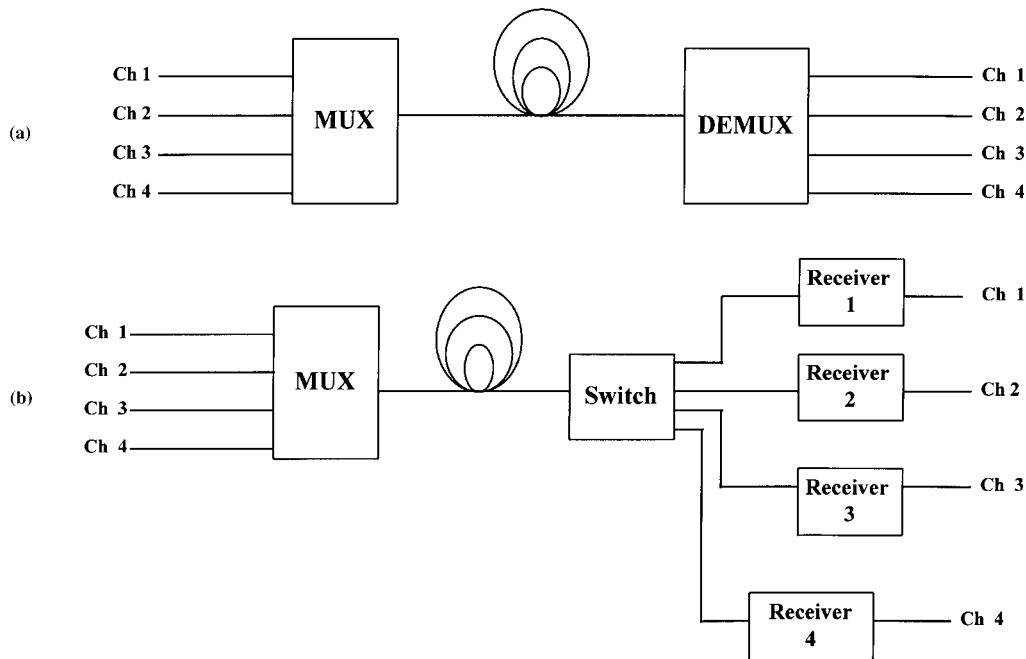


Figure 22.5 Point-to-point (a) and switched fiber (b) systems.

WDM networks have been growing rapidly with many installations in the United States, Europe, Japan, and undersea. There have been demonstrations of state-of-the-art R&D test beds with 16 channels at 20 Gbps per channel. At that rate, the overall data rate of the WDM system will be 320 Gbps, allowing the transmission of more than 5 million telephone calls (or 30,000 textbooks) on the same network. (In contrast, a 20 Gbps TDM network will have one-sixteenth that capacity.)

## 22.4 Transmitters

A transmitter is a complex system that contains a large number of electronic and optical components and circuits, as was shown in Figure 22.2. It is placed at the originating end of a fiber optic link and usually is kept under controlled temperature and other environmental parameters inside a building. A simple block diagram for a transmitter used in a TDM network was shown in Figure 22.3 part (a). The input to the transmitter is an electrical signal that contains the information to be carried through the network. This signal is used to modulate the optical source inside the transmitter. The modulation can take one of three forms:

1. For relatively low data rates (less than 622 Mbps), the laser source can be directly modulated by applying the input signal to an electrode on the laser that controls the light output.
2. Higher data rates (above 622 Mbps), require an external modulator, which can be a separate optical component following the laser.
3. The modulator can be integrated with the laser chip itself.

The modulated light signal is, therefore, an electromagnetic signal that has a very high center frequency (approximately 200 THz) with a modulation (amplitude, frequency, or phase) rate equal to the input (information) electrical signal.

This modulated signal then is passed through an optical isolator, a unidirectional device that allows light to go from the laser toward the fiber but rejects any signal coming from the opposite direction. This isolator is important because it protects the laser from being damaged by undesired reflections. The modulated optical signal then is passed through an optical amplifier (semiconductor or fiber amplifier) to compensate for the losses from the modulator and the isolator. The amplified signal is coupled to the optical fiber using a standard optical connector or a pigtail connection.

The term *laser* stands for *light amplification by stimulated emission of radiation*. This section's focus is on semiconductor laser diodes (LDs). Because of their small size, high efficiency and reliability, and excellent control of wavelength, power, and spectrum characteristics, they are being used routinely in FO systems. Because laser light output is coherent, LDs are used for applications requiring high data rates (in excess of 100 Mbps). On the other hand, the incoherent nature of the output signal from light-emitting diodes (LEDs) limits their applications to data rates of less than 100 Mbps. While the structure and the fabrication process of LEDs and LDs are similar, their light output characteristics are quite different.

### 22.4.1 Light-emitting diodes (LEDs)

LEDs produce light with a wide spectral width; when used in FO communications systems, they can be modulated at frequencies up to 100 MHz. LEDs have the advantages of low temperature sensitivity and no sensitivity to back reflections. Furthermore, LEDs produce incoherent light output that is not sensitive to optical interference from reflections.

LEDs generate light by spontaneous emission. This occurs when an electron in a high-energy conduction band changes to a low-energy valence band, as shown in Figure 22.6. The energy lost by the electron, the bandgap energy,  $E_g$ , is released as a photon, the entity of light. The released photon's energy is equal to the energy lost by the electron, and the wavelength of the emitted photon is a function of its energy.

Because different materials have different orbital states that determine the energy levels of the various electrons, the wavelength of the emitted photon is determined by the material used to make the LED. The wavelength of the emitted photon is given by:

$$\lambda = \frac{hc}{E_g} = \frac{1.24 \mu\text{m}}{E_g(\text{eV})} \quad (22.1)$$

where:

$h$  is Planck's constant,  $6.62 \times 10^{-34}$  Ws<sup>2</sup>

$c$  is the speed of light,  $2.998 \times 10^8$  mps

$E_g$  is the material bandgap in joules

For a semiconductor material with  $E_g = 0.9$  eV, the wavelength of the emitted photons will be about  $1.38 \mu\text{m}$ . The most commonly used materials for LEDs are gallium arsenide (GaAs) with  $E_g = 1.42$  eV and gallium phosphide (GaP) with  $E_g = 2.24$  eV. By adding other materials to the GaAs or the GaP, such as aluminum or indium, it is possible to tailor the bandgap energy to achieve any wavelength in the  $0.5$  to  $2.0 \mu\text{m}$  range.

With appropriate n and p doping, these materials can be used to form a simple pn diode that can function as an LED. Conduction-band electrons are generated by forward-biasing the pn junction of the diode. For a better confinement of the output optical power, a *double heterostructure* (DH) LED is used (Figure 22.6). In a DH-LED, the junction is formed by dissimilar semiconductor materials with different bandgap energy and refractive index values; the free charges are confined to recombining in a narrow, well-defined semiconductor layer, called the *active layer*.




Figure 22.6 Spontaneous emission from light-emitting diode (LED).




Figure 22.7 LED spectrum

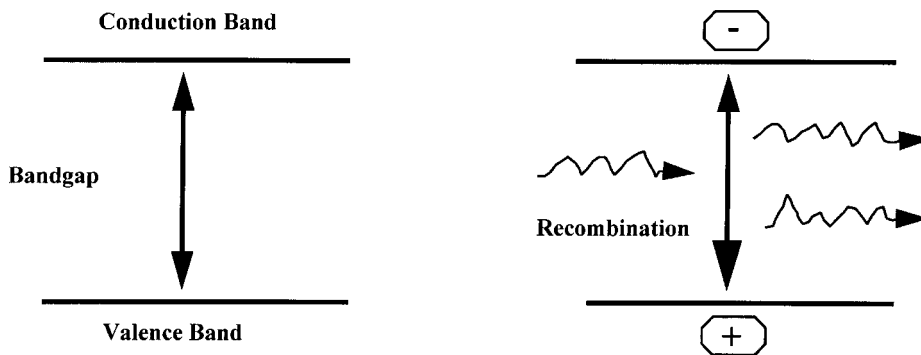


Figure 22.8 Stimulated emission for semiconductor laser.

The spectrum of an LED has a broad distribution of wavelength centered around a wavelength calculated from the foregoing equation. The spectral width often is specified as the full-width at half-maximum (FWHM) of the distribution. Typical values for the FWHM range from 20 to 100 nm (Figure 22.7).

#### 22.4.2 Fabry-Perot (FP) laser diodes

Lasers are capable of producing high output power and directional beams. When used in fiber optic communication systems, semiconductor lasers can be modulated at rates up to about 10 Gbps. Lasers are sensitive to high temperature and back reflections, and the coherent emitted light is sensitive to optical interference from reflections. Two types of semiconductor lasers commonly are used in FO communications systems: *Fabry-Perot* (FP) and *distributed feedback* (DFB) lasers.

The Fabry-Perot laser differs from an LED in that it generates light by *stimulated* (not spontaneous) emission: Photons trigger additional electron-hole recombinations, resulting in additional photons as shown in Figure 22.8. A stimulated photon travels in the same direction and has the same wavelength and phase as the photon that triggered its generation. Stimulated emission can be thought of as amplification of light.

As a photon passes through the region of holes and conduction band electrons, additional photons are generated. If the material is long enough, enough photons can be generated to produce a significant amount of power at a single wavelength.

An easier way to build up power is to place a reflective mirror at each end of the region where the photons multiply. Because the photons can travel back and forth between the two mirrors, their number increases with each trip. The mirrors form the resonator needed for the operation of the laser. For the laser action to occur, a greater number of conduction band electrons than valence band electrons must be present. Called *population inversion*, this state is achieved by forcing a high current density into the active region of the laser diode. The possible wavelengths produced by the resonator are given by:

$$f_{\text{res}} = \frac{mc}{(2ln)} \quad (22.2)$$

where:

$m$  is an integer

$c$  is the speed of light

$l$  is the length of the resonator

$n$  is the refractive index of the laser cavity.

The mode spacing, which is the separation between the different wavelengths, is determined as:

$$\text{mode spacing} = \frac{c}{(2ln)} \quad (22.3)$$

### 22.4.3 Distributed feedback lasers (DFB)

DFB lasers are similar to FP lasers, except that all but one of their spectral components are significantly reduced. Because its spectrum has only one line, the DFB laser's spectral width is much less than that of a FP laser. This greatly reduces the effect of chromatic dispersion in fiber optic systems, allowing for greater transmission bandwidths.

The distributed feedback laser utilizes a grating (a series of corrugations) along the active layer of the semiconductor, as shown in Figure 22.9. Rather than using only the two reflecting surfaces at the ends of the diode, as does a Fabry-Perot laser, the distributed feedback laser uses each ridge of the corrugation as a reflective surface. At the resonant wavelength, all reflections from the different ridges add in phase. Because the DFB laser has much smaller spacings between the resonator elements compared to the Fabry-Perot laser, the possible resonant wavelengths are much farther apart in wavelength, and only one resonant wavelength is in the region of laser gain. This results in the single laser wavelength.

The ends of the diode still act as a resonator, however, and produce the lower-amplitude side modes. Ideally the dimensions are selected so that the end reflections add in phase with the grating reflections. In this case the main mode will occur at a wavelength halfway between the two adjacent side modes; any deviation is called a




Figure 22.9 Distributed feedback (DFB) laser.

*mode offset.* Mode offset is measured as the difference between the main-mode wavelength and the average wavelength of the two adjacent side modes.

#### 22.4.4 Optical modulators

The two most commonly used schemes to modulate the output of a laser diode are direct modulation and external modulation. *Direct modulation* is achieved by applying a high-frequency signal to the laser's drive current electrode (Figure 22.10). This method has been used reliably for modulation rates of up to 1 Gbps in installed fiber links and at much higher data rates (approximately 10 Gbps) in research laboratories. The advantage of direct modulation is the simplicity of the required circuitry, but there are many disadvantages, such as changes in the peak wavelength and spectral bandwidth of the laser output, and limited modulation rates.

*External modulation* of the laser output generally is used for modulation rates above 1 Gbps. It can be achieved by using optical modulator devices at the output of the laser diode, or by fabricating a modulator section that is directly integrated with the laser diode chip, usually called an *integrated laser/modulator*. Both schemes proved to be viable for the 2.4 Gbps and 10 Gbps fiber links that have been installed in the 1990s. This discussion is limited to the two most commonly used modulator techniques: the Mach-Zehnder (MZ) integrated optic lithium-niobate ( $\text{LiNbO}_3$ ) external modulator and the electro-absorption integrated laser/modulator.

**Mach-Zehnder modulators.** The  $\text{LiNbO}_3$  MZ modulator is based on the electro-optic effect in which the influence of an electric field changes the material's refractive index. Figure 22.11 is a schematic diagram of an MZ modulator. An x-cut  $\text{LiNbO}_3$  crystal is used to fabricate an optical interferometer, which consists of an input optical waveguide that splits into two waveguides (at the Y-junction in Figure 22.11). Each arm (waveguide) of the interferometer has an electrode to apply an ac voltage. The two arms combine at the output Y-junction to form the output waveguide. The wave-



guides are formed by diffusing titanium into the  $\text{LiNbO}_3$ , resulting in regions of higher refractive index to guide the optical signal.

The laser output is coupled into the input waveguide of the modulator. The light is split equally into the two arms of the interferometer. By applying a voltage to the two electrodes atop the two arms, the electro-optic effect will cause the optical velocity

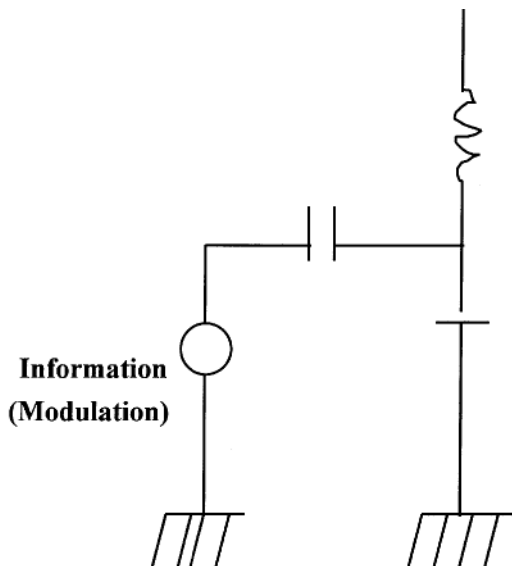


Figure 22.10 Direct laser modulation.

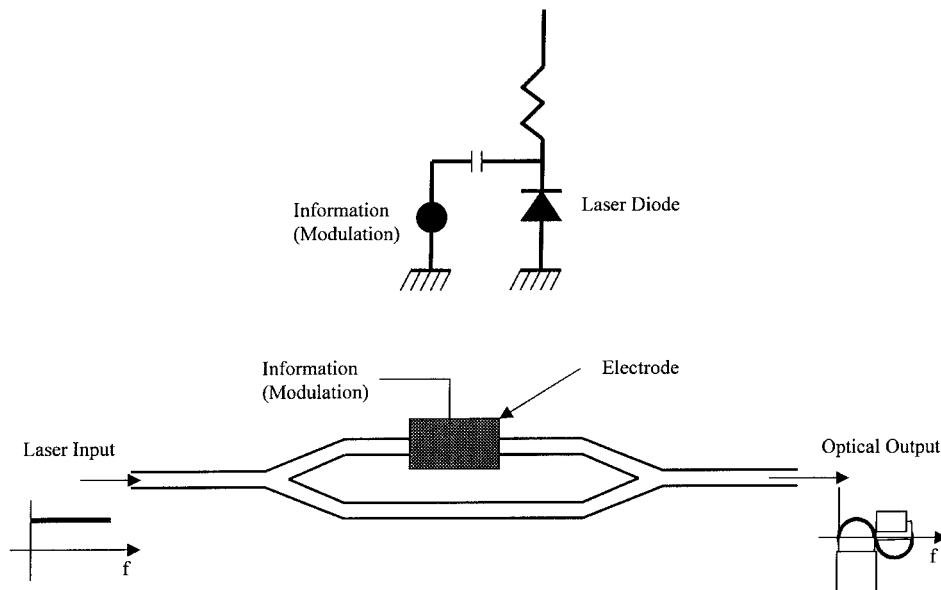


Figure 22.11 Mach-Zehnder optical modulator.

to be higher in one waveguide and lower in the other. After combining at the output Y-junction, total transmission occurs if there is no voltage applied to the electrodes; extinction occurs for  $180^\circ$  optical phase difference.

Mach-Zehnder  $\text{LiNbO}_3$  modulators are being used in many installed fiber optic links at 2.4, 5 and 10 Gbps and can be purchased from several vendors. They typically perform with the following specifications:

- Wavelength 1300 and 1550 nm
- Drive voltage 5–10 V
- Modulation rates 1–20 Gbps
- Insertion loss 5–10 dB

**DFB lasers with integrated electro-absorption modulators.** To reduce the cost of fiber optic transmitters, integrated laser/modulator devices have been developed for multi-Gbps data rates. The most commonly used structure combines a distributed feedback (DFB) laser with an electro-absorptive (EA) modulator section. The device is fabricated using conventional lithography techniques in III-IV epitaxial materials such as indium gallium arsenate phosphide (InGaAsP) epitaxially grown on indium phosphide (InP) substrates. The device consists of a laser section and a modulator section on the same chip (Figure 22.12).

The modulator section is a simple cavity with an electrode deposited on top of the semiconductor material; its operation is based on the Stark effect: An applied electric field causes the semiconductor material to absorb specific wavelengths of an optical signal. The light output from the laser section passes through the modulator section with essentially no losses if a certain voltage (above a certain threshold value) is applied to the modulator's electrode. Without an applied voltage, the light will be absorbed and converted to heat. The combined device there-




Figure 22.12 DFB with integrated electro-absorptive modulator.




Figure 22.13 Optical isolator.

fore acts as a compact digital modulator/switch. Typical performance specifications of commercial DFB/EA laser/modulator devices are:

- Operating wavelength 1550 nm
- Drive voltage 3 V
- Modulation rate 2.4 Gbps

#### 22.4.5 Optical isolators

An isolator is a nonreciprocal device that transports optical power in one direction only. Isolators currently are used to protect laser diodes from undesired reflections, which can cause additional noise. Optical isolators are constructed using a combination of optical components such as polarizers and analyzers, and magnetic components such as a Faraday rotator, to produce the nonreciprocal characteristics. A schematic diagram of a basic optical isolator is shown in Figure 22.13.

The optical signal first passes through a polarizer, oriented in parallel to the incoming state of polarization. The Faraday rotator will rotate the polarization of the optical signal by  $45^\circ$ . The signal then passes through the analyzer, which is oriented at  $45^\circ$  with respect to the input polarizer. The isolator passes an optical signal from left to right and changes its polarization by  $45^\circ$ . In addition, an isolator produces about 2 dB of insertion loss. In the opposite direction (right to left), an optical signal experiences a  $45^\circ$  change in polarization caused by the analyzer and another  $45^\circ$  change caused by the Faraday rotator, resulting in  $90^\circ$  of polarization rotation; it will, therefore, be cross-polarized with (and hence blocked by) the polarizer.

**484 Basic Telecommunications Technologies**

Isolator technology has progressed rapidly in the 1980s, resulting in outstanding performance in terms of low insertion loss, high isolation, and extended wavelength coverage. An example of these isolators has the following specifications:

- Operating wavelength 1530–1580 nm
- Isolation loss < 1.5 dB
- Isolation > 60 dB
- Input reflection < -60 dB

**22.4.6 Optical amplifiers**

Current state-of-the-art fiber optic networks employ optical amplifiers to boost the light signal power without the need for electronic repeaters. These optical amplifiers can provide more than 30 dB of gain with very low noise and extremely wide frequency band.

Optical amplifiers are used in three different applications:

- In-line amplification to boost the signal level
- Front-end preamplification before the photodetector to enhance the s/n ratio
- Booster amplification to compensate for coupler and splitting insertion losses

There are two amplifier types in common use, the *semiconductor laser amplifier* (SLA) and the *erbium-doped fiber amplifier* (EDFA). SLAs are devices fabricated from indium gallium arsenide materials grown epitaxially on indium phosphide substrates using conventional lithography techniques (similar to those used in fabricating integrated circuits). These devices can be of the resonant type or a traveling wave amplifier (TWA) type. A resonant amplifier is similar to the FP laser described previously in this chapter. They are biased below the laser threshold current, provide high optical gain; extreme care must be taken to control their temperature behavior, however. The TWA is a device similar to a semiconductor laser with a special antireflective coating material at both ends of the device. TWAs provide moderate optical gains with extremely low noise.

**Erbium-doped fiber amplifiers.** The invention of the EDFA has revolutionized fiber optic networks because of the simplicity (fiber-in, fiber-out), high optical gain (greater than 30 dB), large bandwidth (several THz), low noise, and low polarization sensitivity. Recently installed fiber optic networks use EDFAs at spans of 40–90 km. For example, the newly installed TAT13 network under the Atlantic Ocean (more than 6000 km), uses EDFAs every 40 km, for a total of 150 amplifiers without any electronic repeaters or regenerators. Similarly, the “Fiber optic Link Around the Globe” (FLAG), which is 27,000 km in length, will use EDFAs every 85 km at a wavelength of 1550 nm.

An EDFA consists of several optical components (Figure 22.14). A short (about 10–30 m) length of erbium-doped silica fiber is pumped using a 1480 nm semiconductor laser at one end of the fiber. The coupling changes the electronic structure of

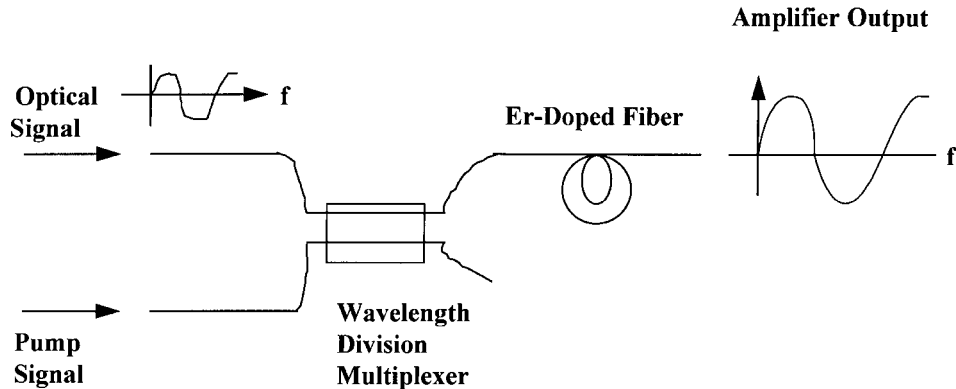


Figure 22.14 Fiber amplifier.

the erbium-doped fiber, allowing more electrons to reach high energy level. As these electrons relax into lower energy levels within the fiber band structure, they emit photons in the 1550 nm range that add to the optical signal at the fiber input, causing the necessary optical gain.

A coupler (usually called a WDM component) is used to combine the optical signal carrying the information with the pump signal. At the other end of the erbium-doped fiber, the optical signal is output with a gain of 30–40 dB. The length of the fiber is chosen to provide the necessary gain and to decouple the pump signal from the optical signal.

Several versions of EDFAs exist. Some are pumped from both ends of the fiber to achieve higher optical gain; some are pumped using 980 nm (instead of 1480 nm) for higher reliability and lower cost. Construction of the amplifier, however, is basically the same, with fiber-in/fiber-out that allows extremely easy coupling to the fiber network. Additionally, erbium-doped fluoride fiber amplifiers are very promising for applications requiring flat gain over a wide wavelength window. These amplifiers are still under investigation, but they promise to provide flat gain (within several tenths of a dB) over a range greater than 15 nm (2 THz).

Other kinds of fiber amplifiers are under investigation. *Praseodymium-doped fiber amplifiers* (PDFAs) provide optical gains in the 1300 nm wavelength domain, but they are not as efficient as EDFAs in the 1550 nm region.

## 22.5 Optical Fibers

An *optical fiber* is a cylindrical dielectric waveguide made of low-tech materials such as silica glass. Its central core, through which the light is guided, is embedded in an outer cladding of slightly lower refractive index. As a result of recent technological advances in fabrication, light can be guided through 1 km of glass fiber with loss as low as about 0.16 dB.

When the core diameter is small (less than 10 microns), only a single mode is permitted and the fiber is said to be single-mode fiber. Fibers with large core diameters (50 and 62 microns are standard diameters) are multimode fibers.

## 22.6 Receivers

After an optical signal has been launched into the fiber, it will become progressively attenuated with increasing distance. This signal will be detected at the receiver end by a photodetector (also called photo diode), which converts the optical signal into an electric current output—the two most commonly used types of photodiodes are p-i-n and avalanche (APD) photodiodes.

After the photodiode, the electric current is amplified and reshaped to bring the signal integrity to a high enough level for further post-processing.

### 22.6.1 P-I-N photodetectors

The p-i-n photodetector is a simple, stable, and wideband device. Figure 22.15 shows the theory of operation for the device. A light signal (photons) impinging on a photodiode is absorbed inside the device. If the energy of the photon is greater than the bandgap energy of the material, an electron-hole pair is generated in the i-region of the detector, and hence an electric current. For wavelengths in the 800–900 nm range, silicon is the ideal material for p-i-n photodiodes. For the 1300–1600 nm range, InGaAs is the material of choice. The p-i-n photodetectors have been designed, fabricated, tested, and used worldwide.




Figure 22.15 P-I-N photodetector.




Figure 22.16 Avalanche photodetector.

### 22.6.2 Avalanche photodetectors (APDs)

The function of an APD (Figure 22.16) is similar to that of the p-i-n photodetector in which incident photons are absorbed and converted to electron-hole pairs. The main difference is its operation near reverse breakdown, which causes carrier multiplication and leaves high gain. The main disadvantage of an APD is its reliability. To obtain a very high-gain bandwidth product (better than 100 GHz), a high reverse voltage (approximately 40 V) usually is applied to the APD, which shortens the life of the device.





## Timing and Delay Jitter

David Robertson

*Hewlett-Packard Ltd., South Queensferry, Scotland*

### 23.1 Introduction

There has been a rapid evolution of integrated digital networks throughout the world. As these networks have developed, there have been several studies on their robustness to digital impairments. This chapter considers just one of these impairments: timing jitter. It shows how jitter is generated, the importance of controlling jitter accumulation, and the effect jitter has on various services carried.

The philosophy behind setting standards for jitter is described, particularly the link between timing jitter and network synchronization control. This chapter then describes some practical jitter measurements to ensure that network performance standards are met.

### 23.2 Jitter Defined

One of the best definitions of timing jitter has been provided by the ITU-T. Timing jitter is “short-term variations of the significant instants of a digital signal from their ideal positions in time.” For the purposes of this definition, a *significant instant* is any convenient, easily identifiable point on the signal, such as the rising or falling edge of a pulse.

We can see what is happening more clearly in Figure 23.1. In this illustration, an ideal timing signal is compared at a succession of instants  $T_n$  with a real timing signal that has some timing jitter. At instant  $T_1$ , the trailing edge of the real signal is displaced by a period  $t_1$  from its ideal position. Similarly, at  $T_2$  to  $T_6$ , there are displacements  $t_2$  to  $t_6$ . If we plot these displacements on a graph, we can see that they form a cyclical pattern that can be called the *jitter function*. Represented in this way, the jitter function can be quantified in terms of frequency and amplitude.

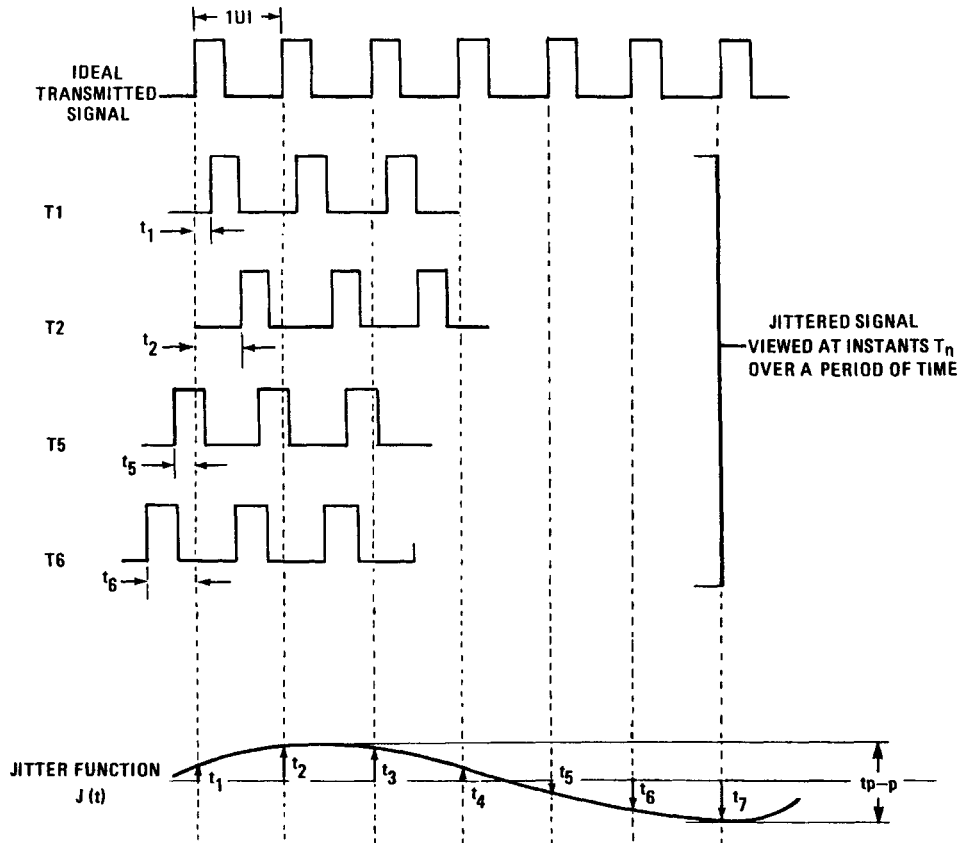



Figure 23.1 Deriving the jitter function.

It is important to be aware that the jitter function is usually complex in practice, made up of a range of components at different frequencies and amplitudes. We will discuss the sources of these components in more detail later.

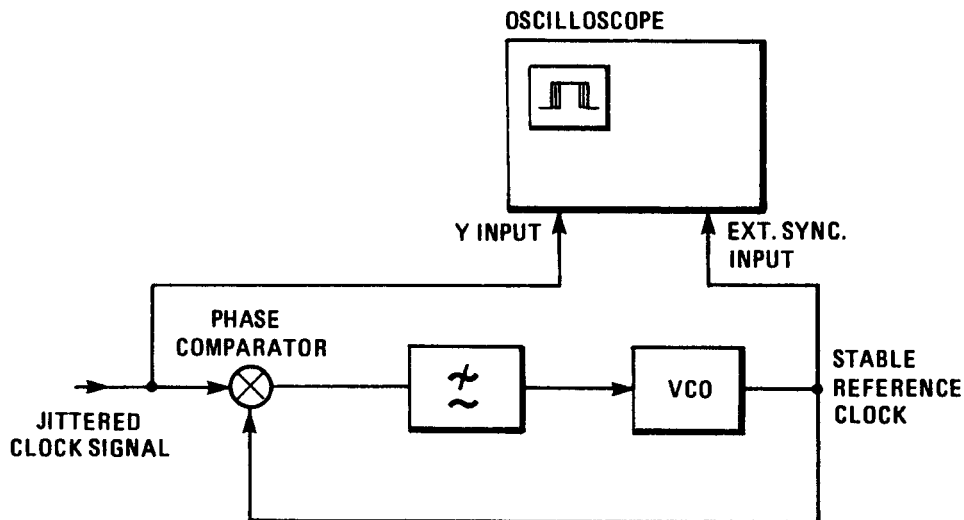
Now that we know what jitter is, how can we observe it and measure it? One common way of evaluating the quality of a digital signal is to set up an *eye diagram* on an oscilloscope (Figure 23.2); a typical arrangement is shown in Figure 23.3. A stable reference clock is derived from the jittered clock signal by means of a phase-locked loop and is used to trigger the oscilloscope externally so that the jittered clock symbols are superimposed on one another. The display can be analyzed as shown in Figure 23.4.

It can be seen from this that jitter has the effect of closing the eye in the horizontal axis. It is worth noting that while an eye diagram can give a qualitative indication of the performance of a digital link, it might not show randomly intermittent eye closures and therefore is not suitable for quantitative analysis.

A method of obtaining a quantitative measurement commonly used in jitter measuring instruments is shown in Figure 23.5.



**Figure 23.2 Eye diagram.** The eye diagram of a random digital signal displayed on a digital sampling oscilloscope. Over successive sweeps, the random bit patterns build up a composite picture of all the possible pattern sequences and transitions. The large open area in the center of the pulse is called the eye opening, with the 1 value at the top and the 0 value at the bottom. The distance between top and bottom at the center of the pulse is the eye height, while the distance between the transitions is called eye width. The eye diagram is a useful qualitative measure of digital system performance. An engineer can spot immediately if the eye is degraded by noise, timing jitter, pulse degradation, or intersymbol interference (ISI). ISI is pattern-dependent, so different test patterns will create different eye diagrams.



**Figure 23.3** Producing an eye diagram.




Figure 23.4 Looking at jitter on an eye diagram.

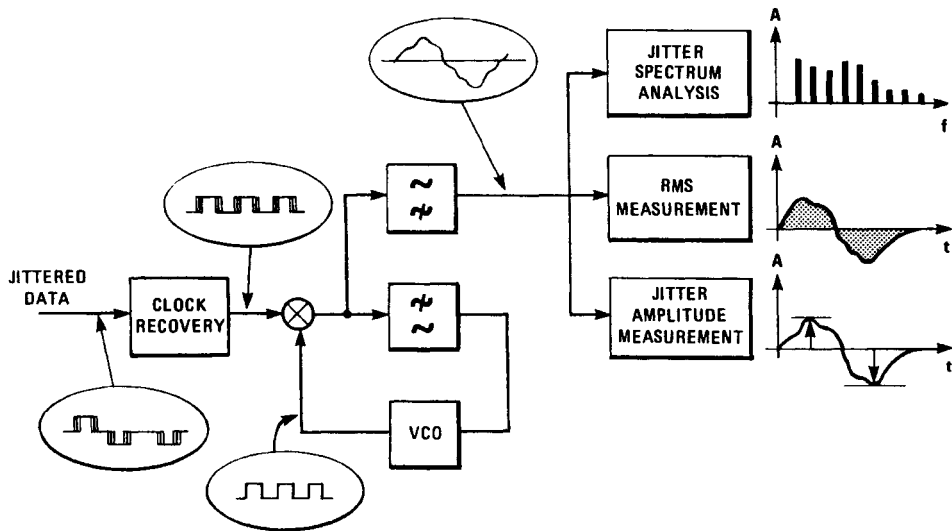


Figure 23.5 Jitter measurements.

Jitter measurements are always made on a clock signal. This may be derived from the data signal in the case of plesiochronous or synchronous systems, or may be taken directly from the operational equipment. In either case, it is essential that the bandwidth of the clock recovery circuit is wide enough not to filter out any of the jitter function from the signal.

The jittered clock signal is fed into a phase comparator, where the phase of the significant instants is compared with a stable reference clock. The phase comparator produces an output voltage that is proportional to the phase difference between the input jittered clock and the reference clock. This output voltage is sometimes called the *demodulated jitter waveform*; the stable reference is derived from it by filtering with a very narrow-band low-pass filter (cutoff frequency typically 5 Hz or lower) to produce a long-term averaged control voltage for a voltage-controlled oscillator. The rest of the demodulated jitter waveform is selected by a high-pass filter and made available for further analysis.

Three main types of measurements can be performed on jitter waveforms. These are *jitter frequency* spectrum analysis, *jitter amplitude* measurements, and *root-mean-square* measurements. Spectrum analysis is simply splitting the waveform up into its frequency components; this may be a way of identifying the major sources of jitter in operational equipment under development.

Jitter amplitude measurements use a peak detector to record the maximum amplitude of jitter that occurs during a specified period. There may be peak-to-peak, positive-peak, and negative-peak values, since jitter waveforms are rarely symmetrical.

Root-mean-square measurements are really a measurement of the power in the demodulated jitter waveform using an RMS voltmeter; power would seem to be a meaningless concept, however, when associated with a phase difference.

What are the units with which jitter is measured? Jitter frequency is measured in the familiar unit hertz. Jitter amplitude is measured in a variety of units, the most common of which is the *Unit Interval* (UI), which is defined as the nominal time for the transmission of one binary digit. Suppose the nominal bit rate of a transmitted signal is 2048 kbps. The nominal time for the transmission of a binary digit is  $1/2048 \times 10^3$  or 0.488 ms. If there is jitter present on this signal with a phase deviation of 1 ms peak-to-peak, the amplitude could be expressed in Unit Intervals as 1/0.488, or 2.05 UI.

Notice that we also expressed the jitter amplitude in units of time in this example. The advantage of using the Unit Interval as the unit of measurement is that it is normalized—in other words, not dependent on the bit rate—and we therefore can compare jitter amplitude at different levels in a digital hierarchy directly when they are expressed in this way.

### 23.3 Why Is Jitter a Problem?

It is clear that jitter degrades the performance of digital networks, but what effect does it have on the services carried? Jitter can degrade communications signals in three distinct ways. It can cause:

- Bit errors in the decision process of network elements such as regenerators, multiplexers, digital crossconnects, etc.
- Uncontrolled slips due to the overflow or underflow of buffer stores.
- Distortion of reconstructed analog signals due to jittered data on the input of digital-to-analog (DA) converters.

## 494 Basic Telecommunications Technologies

The effect of bit errors clearly is more damaging to data traffic than to voice traffic, but the effect of slips is not so obvious. The results of a study on the impact of slips on different services can be summarized thus:

- Voice traffic: 5 percent of slips are audible.
- Voiceband encoded data: error bursts occur in received data.
- Digital data: at least two blocks are lost per slip.
- Facsimile: complete page degradation if error correction techniques are not employed.

Distortion of recovered analog signals is particularly troublesome in the case of video.

### 23.4 Sources of Jitter within the Network

The primary sources of jitter within the telecommunications network are the network elements themselves, such as SDH add-drop multiplexers and SDH regenerative repeaters. The processes which cause jitter and the different types of jitter can be briefly summarized as follows:

- *Mapping/demapping jitter*: This is caused by the phase-smoothing process associated with the read/write clock of elastic stores, and the bit stuffing/destuffing process used to compensate for frequency offsets or variations of the mapped tributary signal.
- *Pointer jitter*: Pointers allow offsets and clock variations between the SDH frame and the virtual container (VC). The stability of reference clocks must be carefully controlled in synchronous networks; otherwise excessive pointer movements will occur (due to clock noise, frequency offsets, and wander), which in turn introduces jitter at the PDH tributary ports.
- *Systematic jitter*: This is caused by misaligned timing recovery circuits, finite pulse width, and clock threshold offsets.
- *Wander*: This is caused by reference clock instability, as well as noise.

### 23.5 Jitter Standards

We have seen that timing jitter must be controlled during the development of integrated digital networks if serious degradation of the quality of service is to be avoided. It is worth considering who needs jitter standards, and why, before the existing standards are described.

It is desirable to have jitter standards for the use of national network operators and manufacturers of telecommunications operational equipment. Network operators have to be able to offer a certain minimum quality of service regardless of the type of traffic carried by the network. They also have to ensure that international communication is possible between customers in different national networks. Operational equipment designers and manufacturers need to be able to produce equip-

ment that will interface directly with other equipment in the network and that will have a specification compatible with the network operators' performance objectives.

Fortunately there is an organization whose role it is to create such standards, particularly to ensure that international communication is possible. The ITU-T has developed consensus standards that establish a minimum acceptable performance. These standards are necessarily a compromise between different network operators' performance objectives and the limitations of equipment already installed or under development.

Recommendation G.823, "The control of jitter and wander in digital networks which are based on the 2048 kbit/s hierarchy," defines a jitter control philosophy for PDH networks based on:

- A maximum network limit at any PDH hierarchical interface that should not be exceeded
- A method of specifying the performance of individual equipments
- Guidelines for the study of jitter accumulation in digital networks

Individual equipment is treated as a "black box" whose jitter performance is defined in terms of three measurements.

- A minimum jitter tolerance on the input port
- A maximum permitted output jitter on the output port in the absence of jitter on the input port
- A jitter transfer characteristic between input and output ports

The recommendation describes the test methods and test signals that should be used to verify equipment performance. The recommended limits are defined in terms of jitter amplitude over a range of jitter frequencies and are presented in the form of graphs or *masks*.

SDH recommendations use a similar philosophy. ITU-T G.825, "The control of jitter and wander within digital networks which are based on the Synchronous Digital Hierarchy (SDH)," covers:

- A maximum network limit at SDH interfaces, which should not be exceeded
- A minimum jitter/wander tolerance on the SDH network element's input port

ITU-T G.958, "Digital line systems based on the Synchronous Digital Hierarchy (SDH) for use on optical fiber cables," covers optical line systems with the following limits:

- A maximum permitted output jitter on the optical line equipment's output interface in the absence of jitter on the input port
- Network equipment limits for the jitter transfer characteristic between optical line equipment's input and output ports
- Jitter tolerance of optical line equipment's input ports

TABLE 23.1 Key ITU-T Recommendations for Jitter.

ITU-T Recommendation	ITU-T Recommendation Description	Key Elements for Jitter
G.823	<b>Digital Networks</b> The Control of Jitter and Wander within Digital Networks which are based on the 2048kbitps hierarchy.	PDH Output Jitter Specification PDH Jitter Tolerance Limits/Mask
G.958	<b>Digital Sections and Digital Line Systems</b> Digital Line Systems based on the Synchronous Digital Hierarchy for use on Optical Fiber Cables	SDH RMS. Output Jitter SDH Jitter Transfer Limits/Mask SDH Jitter Tolerance
G.825	<b>Digital Networks</b> Control of Jitter and Wander within Digital Networks which are based on the Synchronous Digital Hierarchy (SDH)	SDH Network Limits for Jitter SDH Network Equipment Jitter/Wander Tolerance
G.783	<b>General Aspects of Digital Transmission Systems; Terminal Equipments</b> Characteristics of Synchronous Digital Hierarchy (SDH) Equipment Functional Blocks	Tributary Jitter due to combined Mapping Process/Pointer Test Sequences SDH RMS. Output Jitter
O.171	<b>Specifications of Measuring Equipment</b> Timing Jitter Measuring Equipment for Digital Systems	Specification for Jitter/Wander Measuring Equipment

ITU-T G.783, “Characteristics of Synchronous Digital Hierarchy (SDH) equipment functional blocks,” covers many functional aspects of SDH network elements, but its main impact on jitter covers the pointer processes in synchronous network elements, e.g., add-drop multiplexers, etc. In this respect, ITU-T G.783 recommends the following:

- Pointer test sequences that should be applied to the SDH network element, and the accompanying specification for maximum permitted tributary output jitter in the presence of combined demapping jitter and pointer jitter, to ensure error-free interoperability with existing PDH networks.
- A maximum permitted output jitter on the network element’s output interface in the absence of jitter on the input port.

The ITU-T also has created recommendation O.171, “Specification for instrumentation to measure timing jitter on digital equipment,” to ensure that jitter measurements are repeatable and that results obtained using test equipment produced by different manufacturers can be compared directly. Table 23.1 presents key ITU-T recommendations for jitter.

## 23.6 The Integrated Digital Network

PTT experience has shown that it is not economically feasible to provide a separate network for each new service offered. Instead, PTTs have developed integrated digital networks consisting of digital switches interconnected by digital transmission



links and controlled by a common-channel signaling system. These networks are capable of carrying a wide range of services.

The digital network presents some problems for the network operator, not the least of which is how to synchronize the components of a vast, distributed digital machine in a robust and reliable manner.

## 23.7 Synchronizing Integrated Digital Networks

Synchronizing a digital system as large and widely dispersed as an integrated digital network is a daunting task. In practice, the approach taken is to provide a highly accurate and stable National Reference Clock, or use the Global Positioning System (GPS), and distribute timing information to all the switching nodes in the network by means of master-slave clock distribution links. A local clock, synchronized in phase and frequency to the National Reference Clock or GPS, is generated at each switching node. The clock distribution network is quite separate from the transmission links that carry communications traffic between switching nodes.

### 23.7.1 PDH synchronization

PDH transmission equipment (such as multiplexers, line termination units, and regenerators) usually is self-timed. In other words, it extracts timing information from the data stream entering its input port and uses this to synchronize its own clock to the data. Special line coding techniques are adopted to maintain timing information in the data stream during long runs of 0s in the data.

The transmission links connect switching nodes in the network and, as we have seen, derive their timing from the transmitting node. Therefore there is a synchronization problem at the receiving node, which can be split into two components:

- A small, fixed, and stable difference in frequency between the transmitting and receiving node clocks.
- Transient frequency and phase variations in the received data stream due to electrical disturbances and jitter accumulation along the transmission link.

The effect of these timing impairments is minimized by coupling the transmission link to the receiving node through an elastic buffer. The incoming data is read into the buffer by a recovered clock and read out by the local clock. The elastic buffer needs to store enough bits to absorb the largest expected transient variation in data stream timing.

A permanent stable difference between the average value of the recovered clock and the local clock eventually will cause the elastic buffer to overflow or underflow so that some bits are lost or read twice. Any impairment that changes the number of bits in the data stream is called a *slip* and can cause loss of frame synchronization. Slips cannot be eliminated but can be reduced to an acceptable level by suitable design methods.

The impact of slips on the services carried by the network is minimized by a sophisticated elastic buffer, which controls the slip so that an entire frame is dropped or repeated and frame synchronization is not lost.

### 23.7.2 SDH synchronization

SDH/SONET network elements are supplied with a reference clock (unlike PDH network elements, which are self-timed). The reference clock may be 2 Mbps for SDH networks or 1.5 Mbps for SONET networks; more frequently, the favored clock distribution method is to use the synchronous line rate (1.55, 622 Mbps, etc.).

Synchronous networks ideally should derive their timing signals from a single master network clock. SDH/SONET, however, is designed to be able to handle asynchronous operation within the network. This is essential to accommodate timing differences resulting from an SDH/SONET network element losing the network reference clock and switching to its standby clock or to accommodate timing differences at the boundary between two separate networks (i.e., at international boundaries or between different national operators).

In order to accommodate timing differences (clock offsets), the virtual container (VC- $n$ ) can be moved (justified) positively or negatively three bytes at a time (for a VC-4) with respect to the SDH/SONET transport frame. This is accomplished by altering the pointer value at the relevant node. The change to the pointer value, known as *pointer processing*, introduces a new signal impairment, however, known as *pointer adjustment jitter*. This jitter impairment appears on the plesiochronous tributary recovered from the VC- $n$  and can cause a 24 UI jitter transient at the input to the network element's desynchronizers. The function of the desynchronizers is to minimize these jitter transients; this is accomplished by means of low-frequency phase-lock loops and elastic buffers that leak out the jitter transient gradually over time.

## 23.8 Jitter in Transmission Systems

We have identified the self-timed operational equipment in transmission systems as the major source of jitter in an integrated digital network. The types of transmission equipment in the network can be divided into the following main categories:

- Multiplexers
- Coaxial line systems
- Microwave radio systems
- Fiber optic line systems
- Satellite systems

The sources of jitter that we will consider here are, first of all, the electronic components used in equipment design, and second, some particular aspects of regenerator and digital multiplexer design that cause jitter to be generated or to accumulate. It should be noted at this point that sources of jitter in the network tend either to depend on transmitted data pattern content, or be random and pattern-independent. The pattern-dependent sources cause jitter to accumulate on the digital signal as it passes along the transmission system; these therefore are more troublesome.




Figure 23.6 Transition uncertainty.

### 23.9 Jitter from Electronic Components

There are two main mechanisms in electronic components that introduce jitter into timing signals in digital equipment. The first is random electrical noise originating from contact and surface irregularities in components. It is sometimes called *flicker* or *1/f noise* because it tends to increase at very low frequencies. This results in intrinsic phase noise on the signal output from oscillators. The second mechanism is the phase noise in logic circuits resulting from transition uncertainties. This is illustrated in Figure 23.6.

Although these sources are measurable, they are of low amplitude compared to the other sources we will consider. A typical amplitude for this type of jitter is less than 0.01 UI peak.

### 23.10 Jitter from Digital Regenerators

Digital regenerators potentially are a major source of jitter impairments in a digital transmission system. Hundreds of regenerators might be used along the length of a coaxial cable line, so we have to consider not only the sources of jitter within each unit, but also the mechanisms of jitter accumulation. The main components of a digital regenerator are shown in Figure 23.7.

In essence, a fundamental frequency component at the nominal bit rate is extracted from the data signal by means of a nonlinear process such as full-wave rectification, and is used to reproduce the system clock in the clock recovery circuit. The recovered clock is used to sample the incoming data stream at the expected center of each degraded data pulse position and decide whether the pulse is a 1 or a 0, so that the original data stream can be recreated.

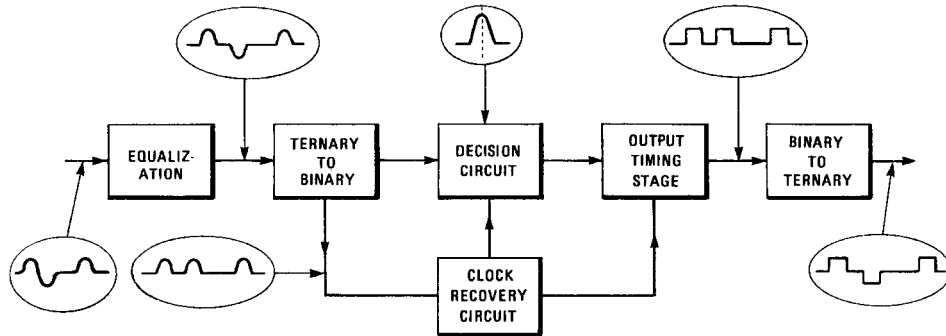


Figure 23.7 Digital regenerator system.

One method of extracting the fundamental frequency is to use an LC tank circuit tuned to resonate at the nominal bit rate of the line system. The tank circuit is an electronic flywheel that will reproduce the system clock, provided the number of 0s in the bit stream is limited by some coding method.

Note from Figure 23.7 that jitter on the recovered clock affects the data stream in two ways. Errors can be made in the decision process due to displacement of the sampling position from the center of the data pulse. Second, since the recovered clock is used to retime the regenerated bit stream, any jitter present on the clock will be transmitted down the line.

### 23.10.1 Pattern-dependent jitter sources

Both pattern-dependent and random sources of jitter can be identified in regenerators. The three most important pattern-dependent sources are:

- Intersymbol interference
- Finite pulse width effects
- Amplitude-to-phase conversion impairments

**Intersymbol interference.** Incorrect equalization of the incoming bit stream could change the shape of the equalized signal in a way that is dependent on the data pattern. When this mismequalized signal is applied to the nonlinear element of the clock recovery circuit, it causes a phase variation in the extracted timing signal.

**Finite pulse width effects.** It can be shown that the signal used to drive the tank circuit should pass through 0 V at the same time the tank circuit output does. If this does not happen, there will be a phase variation on the output signal that is dependent on the pulse shape and pattern content of the input signal.

**Amplitude-to-phase conversion impairments.** A limiting amplifier in the threshold detector following the tuned circuit should produce an output that is completely independent of the input signal amplitude. In practice, due to aging and temperature effects, the amplifier may exhibit a threshold offset that causes jitter to be introduced onto the timing signal (Figure 23.8). The output amplitude of the tuned circuit changes with the pulse density of the incoming digital signal.

### 23.10.2 Random jitter sources

Random jitter sources are those that are not strongly dependent on the pattern content of the incoming bit stream. The most important are:

- Tuned circuit mistuning
- Differential pulse delay
- Crosstalk from adjacent channels in the same line system.

**Tuned circuit mistuning.** Mistuning of the LC circuit causes two effects. Static phase shift causes a change in transmission delay through the regenerator, which is insignificant if the degree of mistuning is retrieved. Dynamic phase shift varies with the pulse density of the bit stream. When the pulse sequence is random, the RMS value of jitter introduced is proportional to the product of the square root of the Q factor and a factor dependent on the degree of mistuning. In a long string of regenerators, some tuned circuits are likely to be mistuned in opposite directions, so that the effects of mistuning tend to cancel and are not highly cumulative.

**Differential pulse delay.** The outputs in some regenerator designs process the positive- and negative-going pulses of the bit stream through separate physical paths. The output circuitry requires transistors to be driven into saturation; any variations in the junction capacitance of these transistors therefore will cause a phase shift in the position of the output pulse. This mechanism generates mainly high-frequency components of jitter, which are removed by the filtering effect of the clock recovery circuit of the next regenerator in the transmission line.

**Crosstalk.** Signal crosstalk from other digital transmission systems operating on the same cable can cause phase shifts in the regenerator timing signal. Crosstalk




Figure 23.8 Amplitude/phase conversion impairment induced jitter.

## 502 Basic Telecommunications Technologies

coupling will be different from one regenerator to another, so it is an uncorrelated and less significant source of jitter.

**23.10.3 Jitter accumulation in a string of regenerators**

Studies have shown that if the content of the bit stream is not completely random, pattern-dependent sources predominate in the jitter accumulation process. For a string of more than 100 regenerators using LC-type clock recovery circuits, the following approximation has been derived:

$$J_n \sim \theta B \frac{n}{2} \quad (23.1)$$

where:

- $J_n$  = jitter output after  $n$  regenerators
- $\theta$  = mean square amplitude from a single regenerator
- $B$  = bandwidth of timing recovery circuit.

If, however, the data content of the bit stream is completely random, it has been shown that jitter accumulates in a way that is proportional to the fourth root of the number of regenerators.

**23.10.4 Jitter reduction**

When the network is carrying a range of services, in practice it is not desirable to place restrictions on the pattern content of the data. Since the jitter accumulation performance of the transmission lines is considerably improved if the bit sequence is random, however, one way of achieving this is to pass the data through a digital scrambler. A simple type of scrambler is shown in Figure 23.9 together with the descrambler, which is required to recover the original data pattern at the other end of the transmission link. One disadvantage of using scramblers is that they can cause error extension effects. In the scrambler shown, if one binary error occurs in the scrambled data, the descrambler will convert it into three errors in the decoded bit stream.

**23.10.5 Jitter tolerance in regenerators**

We have seen how jitter is generated and the way it accumulates in a string of regenerators, but how does jitter create errors in the transmission process? Consider the operation of a simple decision circuit based on a D-type bistable as shown in Figure 23.10.

The D-type makes a decision on whether a 1 or 0 is present on the data input on the rising edge of the recovered clock pulse. The possibility of making an error does not depend on the absolute amplitude of jitter on the incoming bit stream, but on the difference between the amplitude of the jitter on the recovered clock and the data signal. This is sometimes called *misalignment jitter*.

If there were no other impairments affecting the shape of the waveform (such as noise or incorrect equalization), the maximum misalignment jitter that could be tolerated would be 0.5 UI. In practice, other impairments can reduce this to 0.2 UI or less.

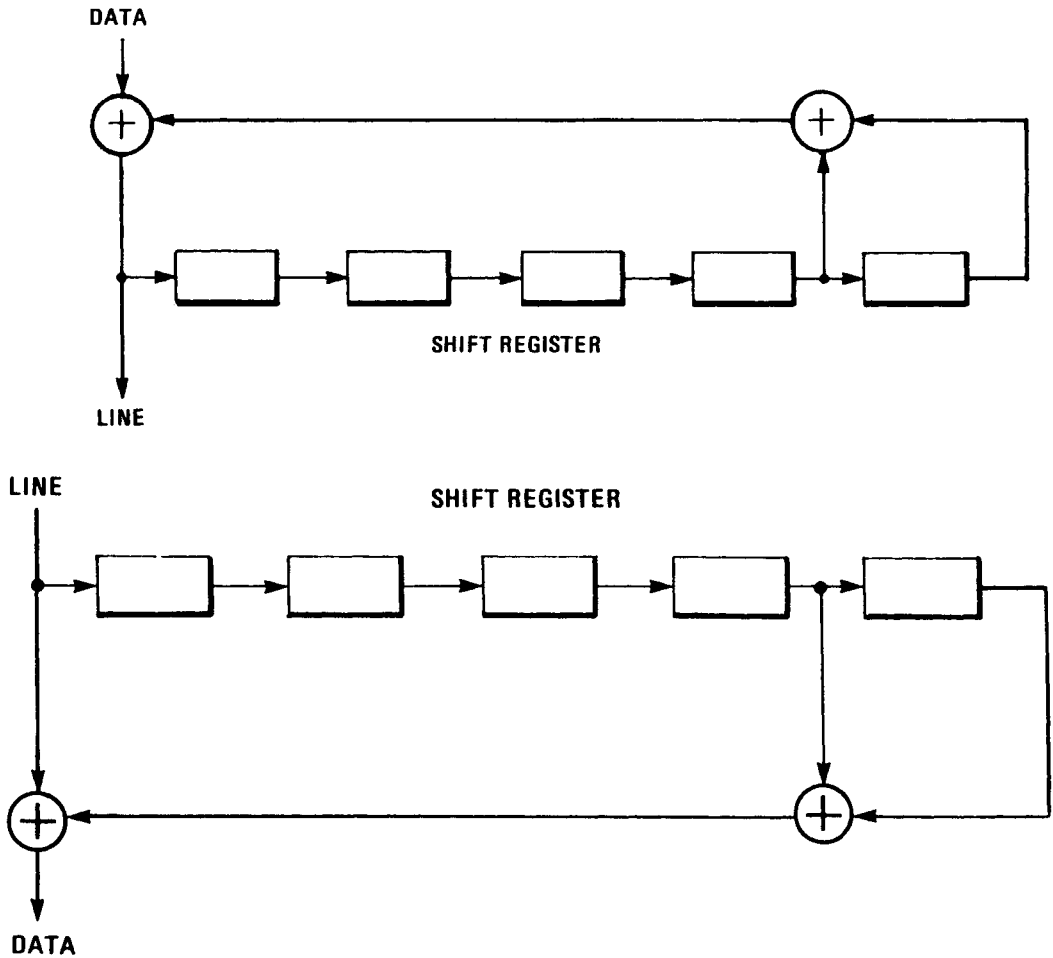


Figure 23.9 Data scrambling.

If we think of jitter as a frequency modulation process, we can see that it will cause a frequency deviation of the nominal bit rate that is proportional to the jitter amplitude. The bandwidth of the tank circuit in the clock recovery section therefore will determine the amplitude of jitter that is passed onto the recovered clock.

It can be shown that the cutoff frequency of the jitter tolerance characteristic corresponds to the half bandwidth between 3 dB points of the tank circuit:

$$F_c = \frac{F_m}{2Q} \quad (23.2)$$

where:

$F_m$  = nominal bit rate and

$Q$  = quality factor of the tank circuit at its resonant frequency.

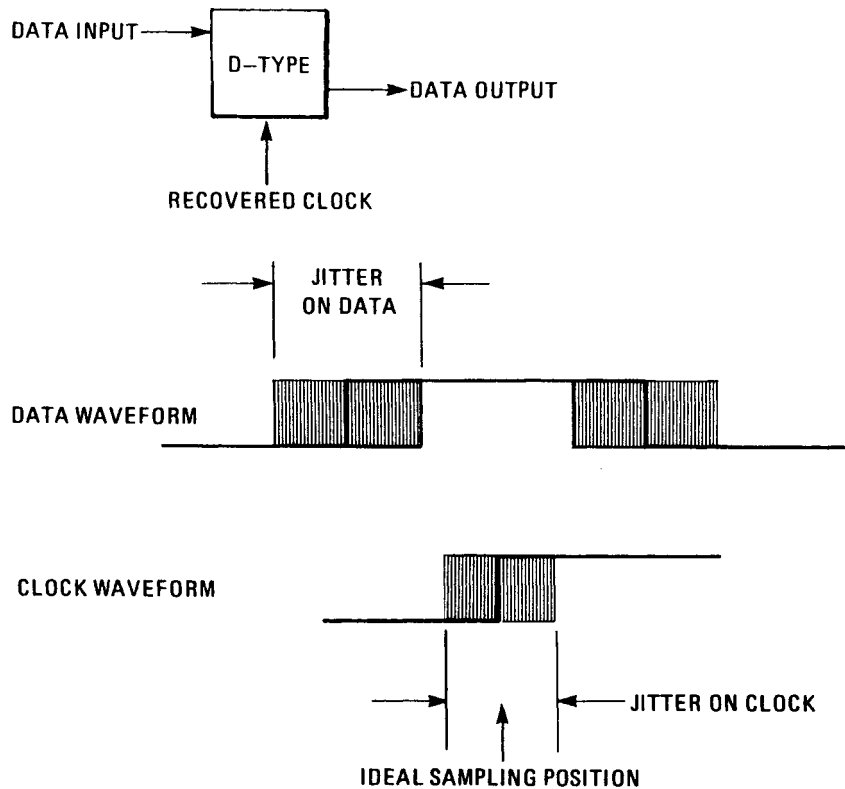


Figure 23.10 Decision circuits.

This indicates one of the compromises in clock recovery circuit design where a tradeoff has to be made between high  $Q$  for good noise immunity and low  $Q$  for wide jitter tolerance.

### 23.11 Jitter Test Categories and Methods

Testing jitter performance of PDH, SDH, and ATM network elements, as well as the telecommunications network itself, is vital in order to ensure error-free performance. The key jitter test categories are:

- *Jitter tolerance testing* of SDH line inputs and PDH tributary inputs
- *Jitter transfer* or *jitter gain* of SDH regenerative repeaters
- *Output jitter* of SDH line interfaces
- *Tributary jitter* performance of SDH network element's tributary outputs

Table 23.2 describes the various jitter test categories and attempts to provide an indication of which type of network equipment should be tested and where each test normally would be performed:



**TABLE 23.2 Jitter Test Categories.**

Test Categories	Network Element under Test	Design Verification & Field Trials	Manufacture	Installation & Commissioning	Operation & Maintenance
Jitter Tolerance or MTIJ	ADM, DXC, Regenerator Input Ports	Essential	Essential		Recommended
Jitter Transfer or Jitter Gain	Regenerators	Essential	Essential	Recommended	
Output Jitter – Line Rate	Network Test			Essential	Recommended
Tributary Jitter	Network Test			Essential	Recommended
Tributary Jitter (due to pointer adjustments)	ADM, DXC Tributary O/P ports	Essential	Essential	Recommended	
Tributary Jitter (due to demapping process)	ADM, DXC Tributary O/P ports	Essential	Essential	Recommended	


- In the laboratory as part of the design or qualification process (Design Verification), or
- As part of a manufacturing test procedure (Production), or
- During the installation and commissioning process by the network equipment manufacturer (NEM) installation teams, or
- By the telecommunications network operators (TNO) themselves.

Finally, “Maintenance” refers to tests performed by the TNO as part of a routine maintenance strategy. This table is not part of any ITU-T recommendation; it is intended for guidance only and may differ slightly from some NEM’s or TNO’s practices.

### 23.11.1 Jitter tolerance or maximum tolerable input jitter (MTIJ)

Jitter is present to a certain extent in all telecommunications networks. Bit errors or data loss will occur in a digital signal if jitter at the input port of a network element exceeds a threshold value. It is therefore vital that the network element be designed to tolerate a sufficient level of jitter, i.e., not lose lock or introduce errors when certain values of jitter are present. The ITU-T specifies the lower limit of maximum tolerable input jitter (MTIJ) in the form of an ITU-T jitter tolerance mask. All network elements should be rigorously tested against this minimum standard. The measured input jitter tolerance of the network element must be greater than the ITU-T standard mask.

The ITU-T defines the lower limit of maximum tolerable input jitter (MTIJ) in terms of the amplitude and frequency of sinusoidal jitter which, when modulating a test pattern, should not cause any significant degradation in network equipment operation. The ITU-T has a number of standards relating to jitter tolerance.



PDH Rate	A2	A3	f1	f2	f3	f4	Pattern
2Mbps	1.5	0.2	20 Hz	2.4 kHz (93 Hz)	18 kHz (700 Hz)	100 kHz	$2^{15}-1$
8Mbps	1.5	0.2	20 Hz	400 Hz (1.7 kHz)	3 kHz (80 kHz)	400 kHz	$2^{15}-1$
34Mbps	1.5	0.15	100 Hz	1 kHz	10 kHz	800 kHz	$2^{23}-1$
140Mbps	1.5	0.075	200 Hz	500 Hz	10 kHz	3500 kHz	$2^{23}-1$


Figure 23.11 ITU-T G.823 PDH jitter tolerance specifications.

ITU-T G.823 (Figure 23.11) defines the MTIJ for PDH interfaces based on 2 Mbps hierarchy, i.e., 2, 8, 34 and 140 Mbps are the key plesiochronous interfaces specified here. ITU-T G.823 has two different masks, one for high-Q and one for low-Q systems. High-Q refers to network equipment optimized for a narrow jitter transfer function, which is a design trade-off versus jitter tolerance. High-Q systems normally are regenerators that must control jitter accumulation along the link. Low-Q refers to network equipment designed for excellent jitter tolerance and a wide jitter transfer function, where jitter accumulation is not the main concern.

ITU-T G.958 (Figure 23.12) defines the MTIJ for SDH line systems, regenerators, etc., at the synchronous rates, i.e., STM-1, STM-4 and STM-16. ITU-T G.958 also has two different masks, one for type A and one for type B systems. Type A/type B is the SDH term analogous to high-Q/low-Q in PDH terminology. Type A refers to wide bandwidth, i.e., low-Q systems optimized for jitter tolerance. Type B refers to narrow bandwidth, i.e., high-Q systems with a reduced jitter tolerance specification.

ITU-T G.825 (Figure 23.13) defines the specifications for the MTIJ for SDH network elements, e.g., ADMs, DXCs, etc., and covers modulating frequencies from very low frequencies of 12  $\mu$ Hz at high jitter amplitudes of thousands of UI up to modulating frequencies of 20 MHz at STM-16. As an installation test, jitter tolerance normally is evaluated against the ITU-T G.958 mask. Testing to the full ITU-T G.825 mask is more appropriate during design verification.

Jitter tolerance tests normally are performed by an SDH jitter analyzer transmitting a PRBS test pattern modulated either internally or externally with sinusoidal jitter. The jitter is increased until errors are detected in the analyzer's receiver. The




**Type A Regenerators**

Type A	A2	A1	f0	f1
STM-1	1.5	0.15	6.5 kHz	65 kHz
STM-4	1.5	0.15	25 kHz	250 kHz
STM-16	1.5	0.15	1000 kHz	100 kHz

**Type B Regenerators**

Type B	A2	A1	f0	f1
STM-1	1.5	0.15	1.2 kHz	12 kHz
STM-4	1.5	0.15	1.2 kHz	12 kHz
STM-16	1.5	0.15	1.2 kHz	12 kHz

Figure 23.12 ITU-T G.958 SDH jitter tolerance specifications.



	A0	A1	A2	A3	A4	f0	f12	f11	f10	f9	f8	f1	f2	f3	f4
STM-1	2,800	311	39	1.5	0.15	12u	178u	1.6m	15.6m	125m	19.3	500	6.5k	65k	1.3M
STM-4	11,200	1,244	156	1.5	0.15	12u	178u	1.6m	15.6m	125m	9.65	1k	25k	25k	5M
STM-16	44,790	4,977	622	1.5	0.15	12u	178u	1.6m	15.6m	125m	12.1	5k	100k	100k	20M

Figure 23.13 ITU-T G.825 SDH jitter tolerance specifications.

applied jitter should be increased beyond the jitter amplitude specified in the ITU-T mask. The term *auto jitter tolerance* refers to the test technique where the analyzer automatically increases the jitter until the point of failure, plots the point on the graph, and then proceeds to the next jitter frequency. The analyzer displays the plotted graph alongside the ITU-T mask. If the plotted graph is above the mask, then the jitter tolerance exceeds the ITU-T standard; any incursion below the mask indicates a failure.

There are two alternative techniques for performing jitter tolerance tests:

- 1 dB Power Penalty (or BER Penalty)
- Onset of Errors

The *1 dB Power Penalty* method (Figure 23.14) is the recommended test method and involves setting up the PRBS BER test under no-jitter conditions and reducing the optical power level at the network element's optical input, using the optical attenuator, until a BER of  $10^{-10}$  is measured by the SDH Jitter analyzer's receiver. (**Note:** The optical receiver's sensitivity is the input power that results in an error rate of  $10^{-10}$ .) The optical power is then increased by 1 dB, again using the optical attenuator.

A BER value is then entered into the SDH jitter analyzer's receiver. The applied jitter is increased until the specified BER ratio is detected in the receiver. A  $10^{-10}$  BER is specified by the ITU, but this results in very long test times; in practice, a BER in the range  $10^{-7}$  to  $10^{-10}$  is used and generally provides valid results. Refer to Supplement No. 3.8 of the ITU-T O-Series Recommendations or ITU-T O.171 Appendix A. The BER Penalty technique is used to evaluate the ability of a network element's clock recovery input circuitry to accurately recover the clock from a jittered data signal.

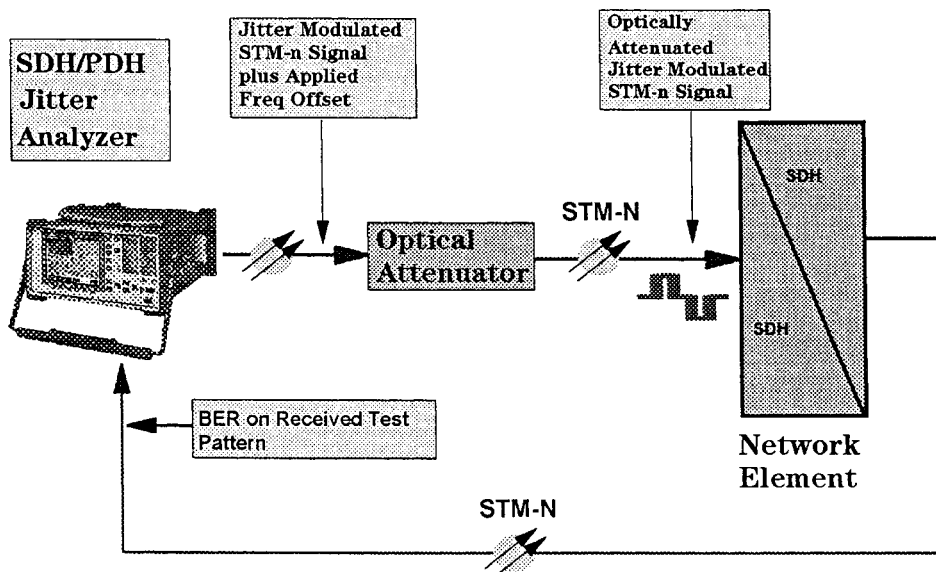


Figure 23.14 1 dB Power Penalty jitter tolerance test setup.

The *Onset of Errors* technique involves increasing the applied jitter until errors or alarms of any kind are detected by the analyzer's receiver. This technique is used to evaluate electrical interfaces and the ability of circuits, such as synchronizers and desynchronizers in SDH network equipment, to accommodate dynamically varying input data rates.

In order to perform jitter tolerance testing accurately and to the required ITU-T standards, the following test equipment requirements must be satisfied. The analyzer should contain fixed jitter tolerance masks covering low- and high-Q systems with peak-to-peak jitter amplitudes and modulating frequencies to ITU-T G.823. For SDH line systems, the fixed jitter tolerance masks are those specified in ITU-T G.958 with a choice of Type A or B masks.

The analyzer also must be able to generate the low jitter modulation frequencies and high amplitudes specified in ITU-T G.825, and indeed meet the more stringent requirements of ITU-T O.171 (the jitter measuring equipment standard). It is important that the analyzer be able to perform automatic jitter tolerance testing at the maximum frequency offsets (specified in ITU-T G.703). Clock recovery circuits normally can tolerate large values of jitter at nominal frequency offsets, but this is not usually the case when the recovery circuit is pulled to the end of its frequency range. This provides a realistic measure of an interface's MTIJ by simulating the expected worst-case conditions normally found in a transmission network.

ATM traffic also is carried on PDH and SDH physical interfaces, so it is a requirement that ATM networks and network equipment meet the same ITU-T standards for jitter as PDH and SDH network equipment. In order to test jitter and wander tolerance of ATM network equipment, it is essential that the jitter analyzer be capable of generating an ATM signal structure, complete with jitter-generation capability. Jitter problems in ATM networks may manifest themselves in different ways than the traditional PDH/SDH networks. Jitter on an ATM cell stream can corrupt the cell header, causing cell loss and cell misinsertion, and also corrupt the payload, leading to errored cells.

### 23.11.2 Jitter transfer

*Jitter transfer* defines the ratio of output jitter amplitude to input jitter amplitude, versus jitter modulation frequency, for a given bit rate. This jitter gain usually is expressed in decibels (dB):

$$\text{Jitter gain: } 20 \log_{10} \left( \frac{J_{\text{out}}}{J_{\text{in}}} \right) \text{ dB} \quad (23.3)$$

where:

$J_{\text{out}}$  = jitter at regenerator output port

$J_{\text{in}}$  = jitter applied at regenerator input port.

In real network equipment, some proportion of the jitter present at the input port of a regenerative repeater will be transmitted to the output port. The function of a regenerator is to receive a low-amplitude signal and to attempt to reconstruct the

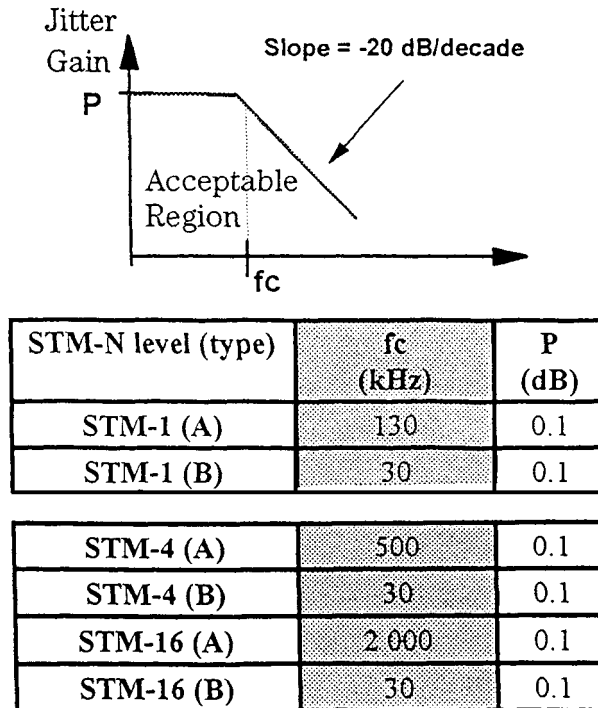


Figure 23.15 ITU-T G.958 SDH jitter transfer specifications.

original signal as closely as possible for retransmission. In order to reconstruct the original signal, the clock must be recovered. Clock recovery is less than perfect, however, and the extracted timing signal will be jittered. Therefore the regenerated output signal also will have jitter added to it.

Regenerators are by definition used in long links and often are cascaded to form a long-haul link. When this happens, an accumulation of jitter occurs. With cascaded digital equipment, it therefore is particularly important to restrict and also measure the value of jitter gain in order to minimize jitter accumulation over such long-haul links.

The ITU-T specifies the jitter transfer, i.e., jitter gain versus jitter modulation frequency, in the form of a transfer characteristic template. Because the design or network goal is to minimize jitter gain, the jitter transfer results plot must be under the transfer mask—unlike the jitter tolerance results, which should be above the tolerance mask. ITU-T G.958 defines the Jitter Transfer Pass masks for SDH systems at STM-1, STM-4, and STM-16 (Figure 23.15) for type A and type B systems. ITU-T G.958 (Figure 23.15) also specifies the input jitter amplitudes in the same recommendation. ITU-T G.823 defines the Jitter Transfer Pass masks for PDH systems. The actual values for individual PDH components (such as PDH MUXs, etc.) are specified in ITU-T G.742 and ITU-T G.751.

Network equipment's jitter transfer can be measured using an SDH jitter analyzer with a digital signal modulated by sinusoidal jitter and a narrow band, that is, a selective receiver. A problem exists with a large number of jitter analyzers, however, in

that their receivers are wideband and are unable to measure within a sufficiently narrow bandwidth. These instruments are designed to measure peak-to-peak jitter in the transmission network for troubleshooting purposes; they are not designed to make selective jitter measurements. Such jitter analyzers just measure the peak-to-peak value of the incoming jitter over a wide frequency range.

Problems occur when testing the jitter transfer of real network equipment such as SDH regenerators. The regenerator produces intrinsic jitter, which disturbs the measurement because the receiver cannot determine whether it is measuring the injected sinusoidal jitter from the SDH analyzer's transmitter or the intrinsic jitter generated at a different frequency by the regenerator. The problem is greatest at higher frequencies, when the amount of jitter injected (according to ITU-T G.958) is much smaller. The measurement is then corrupted by the higher-amplitude intrinsic jitter generated by the regenerator at lower frequencies, leading to incorrect results.

The accurate method for measuring jitter transfer requires a selective measurement. Experience has shown that a filter bandwidth of approximately 10 Hz is required in order to provide sufficiently accurate and repeatable results.

### 23.11.3 Output jitter

*Output jitter* refers to the amount of jitter present at the output ports of the network element. The ITU-T specifies the maximum permissible amounts of output jitter that should be present in the telecommunications network. It is important that the amount of jitter present in the network be measured regularly as part of a preventive maintenance strategy (Figure 23.16).

The overriding reason for this is to ensure that the amount of jitter never exceeds the specified lower limit of maximum tolerable input jitter (MTIJ) specified for the network element's input ports. In other words, if the jitter level is excessive, the network element's input circuitry (clock recovery circuits, etc.) might not have been designed or qualified to work error-free under such jittered conditions. ITU-T G.825 (Table 23.3) specifies the output jitter limits for SDH networks at STM-1, STM-4, and STM-16. ITU-T G.823 (Table 23.4) specifies the output limits for PDH networks based on 2 Mbps, i.e., 2 Mbps, 8 Mbps, 34 Mbps, and 140 Mbps.

The ITU-T also specifies the measurement bandwidth. The PDH jitter measuring equipment standard ITU-T O.171 reinforces the network equipment recommendations by defining the appropriate high-pass (HP) and low-pass (LP) filters (Table 23.5), which jitter analyzers should use to make such output jitter measurements at the PDH rates. The ITU-T O.171 recommendation also specifies the accuracy that




Figure 23.16 ITU-T measurement arrangements for output jitter.

## 512 Basic Telecommunications Technologies

**TABLE 23.3 ITU-T G.825 SDH Maximum Output Jitter Limits and Measurement Bandwidths.**

	Network Limit B1 (UIpp)	Network Limit B2 (UIpp)	HP1 f1	HP2 f3	LP f4
<b>STM-1</b>	1.5	0.15	500 Hz	65 kHz	1.3 MHz
<b>STM-4</b>	1.5	0.15	1000 Hz	250 kHz	5 MHz
<b>STM-16</b>	1.5	0.15	5000 Hz	1 MHz (under study)	20 MHz

**TABLE 23.4 ITU-T G.823 PDH Maximum Output Jitter Limits and Measurement Bandwidths.**

	Network Limit B1 (UIpp)	Network Limit B2 (UIpp)	HP1 f1	HP2 f3	LP f4
<b>2 Mbps</b>	1.5	0.2	20 Hz	18 kHz	100 kHz
<b>8 Mbps</b>	1.5	0.2	20 Hz	80 kHz	400 kHz
<b>34 Mbps</b>	1.5	0.15	100 Hz	10 kHz	800 kHz
<b>140 Mbps</b>	1.5	0.08	200 Hz	10 kHz	3500 MHz

**TABLE 23.5 ITU-T O.171 Measurement Filters.**

	HP1	HP2	LP
<b>2 Mbps</b>	20 Hz	18 kHz	100 kHz
<b>8 Mbps</b>	20 Hz	80 kHz	400 kHz
<b>34 Mbps</b>	100 Hz	10 kHz	800 kHz
<b>140 Mbps</b>	200 Hz	10 kHz	3500 kHz

**TABLE 23.6 ITU-T G.825 Measurement Filter Requirements.**

	HP1	HP2	LP
<b>STM-1</b>	500 Hz	65 kHz	1.3 MHz
<b>STM-4</b>	1000 Hz	250 kHz	5.0 MHz
<b>STM-16</b>	5000 Hz	1 MHz (under study)	20.0 MHz

jitter analyzers must meet to make output jitter measurements. The SDH network equipment standard ITU-T G.825 (Table 23.6) specifies the measurement bandwidth and, therefore, the filters that are required at SDH rates. (SDH jitter measuring equipment will be covered by a new standard, ITU-T O.175.)



In addition, ITU-T G.958, the optical SDH line system standard, specifies the maximum acceptable jitter at a network element's optical output as an RMS value of 0.01 UI RMS (measured with a jitter-free STM- $n$  signal applied at the input). This output jitter measurement is referred to as *Jitter Generation* in ITU-T G.958. For jitter generation, the measurement bandwidth is still under study, although the suggested method is to use a 12 kHz high-pass filter.

#### 23.11.4 Tributary jitter

In order to ensure reliable and error-free interworking between PDH and SDH networks, SDH equipment must control within specified limits the level of jitter present at PDH output tributaries from the synchronous network. Testing a network element's tributary output jitter performance therefore is essential when evaluating new SDH equipment.

The jitter present on a PDH signal output from an SDH network element results from two primary sources, namely:

- *Pointer adjustments*, which compensate for asynchronous operation between different nodes within an SDH network.
- The *bit-stuff justification* process, which is performed when mapping an asynchronous signal into the synchronous transport signal.

By far the most serious (that is, causing the largest amount of tributary jitter) is pointer adjustment jitter. The ITU-T G.783 and ETSI TM-1015 have generated specifications designed to limit the amount of jitter resulting from each of these sources. Verifying an SDH network element's compliance with these standards therefore is essential and requires two different test approaches:

- Pointer Adjustment Jitter test
- Mapping Jitter test

**Tributary jitter due to pointer adjustments.** Pointer adjustments compensate for clock differences between different nodes in an SDH network, e.g., due to the failure of the reference clock, clock noise, frequency offsets and wander, etc. High transient levels of tributary jitter in the PDH output signal are caused by these pointer adjustments. Jitter resulting from pointer adjustments is totally different in character from that previously experienced in PDH networks:

- It is transient in nature.
- It is relatively high in amplitude: 24 UI at the desynchronizer input. (Figure 23.17 shows typical pointer jitter characteristics.)
- Most of the energy is contained in low-frequency components.

It is essential to use a jitter analyzer designed for accurate and reliable results when measuring pointer adjustment jitter, avoiding "traditional" PDH jitter analyzers and some SDH analyzers that measure only repetitive sinusoidal jitter. This is particularly




Figure 23.17 Pointer jitter characteristics.

important when the large phase step associated with an SDH pointer movement can cause up to 24 UI of jitter per AU-4 pointer movement prior to the network element's desynchronizer (Figure 23.17). It also is a requirement in ITU-T O.171 February 1996 (Jitter Measuring Equipment Standard) that the jitter analyzer has low intrinsic jitter performance,  $<0.075$  UIpp for 10–3500 kHz measurement bandwidth.

The ITU-T/ETSI standards define four pointer sequences for use when evaluating a network element's pointer adjustment jitter performance. These sequences are designed to emulate the pointer activity that results from degradation or failure within a network's synchronization. The Pointer Test Sequences shown in Figure 23.18 are used to test 34 Mbps and 140 Mbps tributary outputs on the associated TU-3 and VC-4 pointers respectively. Sequences A and D emulate the network situation where there is no overall frequency offset, but mimic the situation where loss of synchronization or excessive wander occurs. Sequences B and C emulate a network situation with a frequency offset of 4.6 ppm between ends of the path, the worst-case offset of a Stratum 3 clock.

When testing pointer adjustment jitter on a 2 Mbps tributary output, similar pointer sequences are generated on the TU-12 pointer. The key differences are that Sequence D is invalid for 2 Mbps; the time separating regular adjustments (referred to as  $T_2$ ) in Sequences B and C is  $>750$  ms (not 34 ms), and the time separating the double adjustments (referred to as  $T_3$ ) in Sequence B is 2 ms (not 0.5 ms).

As can be seen from Table 23.7, the maximum acceptable pointer adjustment jitter varies depending on:

- Which pointer sequence is used to stress the network element's desynchronizer
- The jitter measurement bandwidth

Connect the SDH jitter analyzer as shown in the typical test setup shown in Figure 23.19. The SDH analyzer generates the mapped payload containing a PRBS test pattern, adds the ITU-T G.783 pointer sequences, and measures the jitter and BER on the demapped PDH signal. It also provides independent generation of a PDH offset as well as Pointer Sequence Generation, an important aspect because ITU-T G.783 allows the presence of any in-range PDH offset during pointer adjustment tests. Note that is vital when making pointer jitter measurements that the measurement period be long enough to ensure that a complete pointer sequence is generated and received, ensuring that the maximum peak-to-peak value of jitter is detected.

**Tributary jitter due to demapping process.** Demapping jitter, as mentioned earlier, is a less serious impairment than pointer adjustment jitter. This jitter impairment is

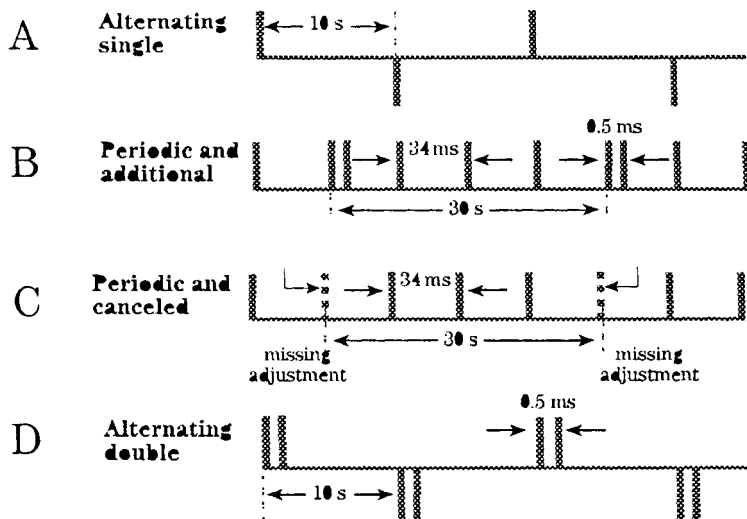


Figure 23.18 ITU-T G.783 pointer adjustment sequences.

TABLE 23.7 ITU-T G.783 Pointer Adjustment Sequences, Jitter Limits, and Measurement Bandwidths.

Tributary Output	Pointer	Pointer Test Sequence	Measurement Bandwidth	Maximum Jitter (UIpp)
2 Mbps	TU-12	A, B, C	0.02–100 kHz <sup>1</sup>	0.4
		A, B, C	18–100 kHz <sup>2</sup>	0.075
34 Mbps	TU-3	A, B, C	0.1–800 kHz <sup>1</sup>	0.4
		D	0.1–800 kHz <sup>1</sup>	0.75
		A, B, C, D	10–800 kHz <sup>2</sup>	0.075
140 Mbps	AU-4	A, B, C	0.02–3500 kHz <sup>1</sup>	0.4
		D	0.02–3500 kHz <sup>1</sup>	0.75
		A, B, C, D	10–3500 kHz <sup>2</sup>	0.075

1. Equivalent to measurement filter “LP + HP1”

2. Equivalent to measurement filter “LP + HP2”

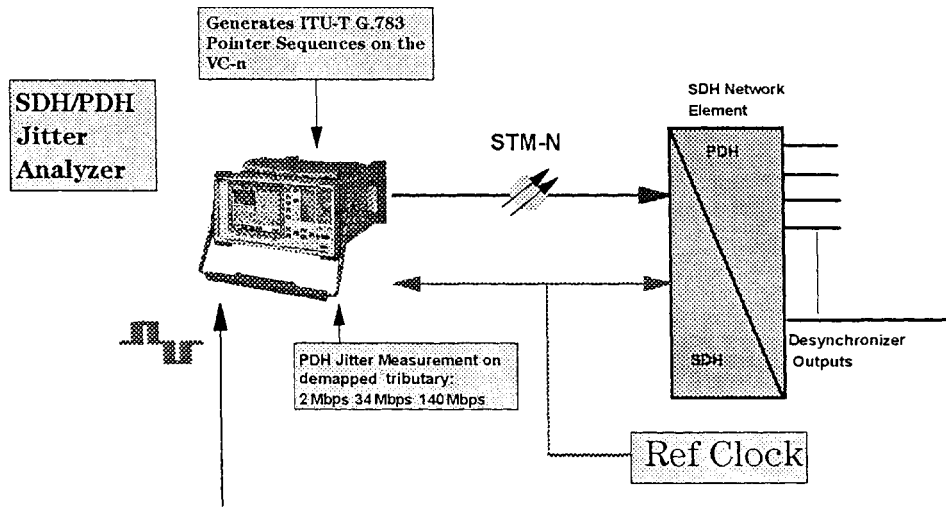


Figure 23.19 Pointer jitter test setup.




Figure 23.20 Demapping jitter characteristics.

caused by the “bit-stuff justification” process used when mapping an asynchronous (2/34/140 Mbps) payload into the SDH transport signal. Demapping jitter is of low amplitude (ITU-T G.783 specifications are shown in Table 23.7) and relatively high frequency, and therefore can be suppressed by the SDH network desynchronizer. Its amplitude varies as the PDH tributary frequency is offset relative to the VC- $n$ . (This is due to changes in the mapping’s bit-stuff justification ratio to compensate for such offsets.) Finally, the peak demapping jitter occurs at a small offset from 0 ppm, the PDH tributary rate from the VC- $n$ . See Figure 23.20, which depicts peak mapping jitter versus PDH payload offset in ppm.

The demapping jitter test finds the maximum peak-to-peak jitter caused by offsetting the frequency of the mapped PDH signal relative to the VC- $n$  (virtual container) used to transport the payload. In order to ensure that the actual maximum jitter value is found, the jitter measurements must be performed for a large number of closely spaced offsets (both positive and negative).

The basic test configuration used to evaluate a network element's demapping jitter performance is the same as the pointer adjustment jitter test setup (see Figure 23.21). The SDH jitter analyzer generates an SDH signal containing a frequency-controlled PDH (2, 34, and 140 Mbps) payload, containing a PRBS test pattern, mapped into the VC-*n* channel associated with the PDH output under test. The SDH jitter analyzer receives the demapped PDH signal and performs a jitter measurement on the PDH output signal.

It is essential that the SDH jitter analyzer and the network element are timed from the same clock. This ensures that there are no pointer adjustments and that it is demapping jitter and not pointer jitter that is being measured. In order to make accurate demapping jitter tests, the SDH jitter analyzer must provide a precise frequency offset control of the mapped PDH test pattern relative to the VC-*n* test channel. The frequency of the mapped payload is adjusted while the frequency of the VC-*n* remains unchanged. At each offset value, the jitter present on the demapped PDH tributary signal is measured on the jitter analyzer.

ITU-T G.783 specifies the maximum acceptable demapping jitter as less than 0.075 UIpp (Table 23.8). In order to obtain accurate and reliable results, the jitter

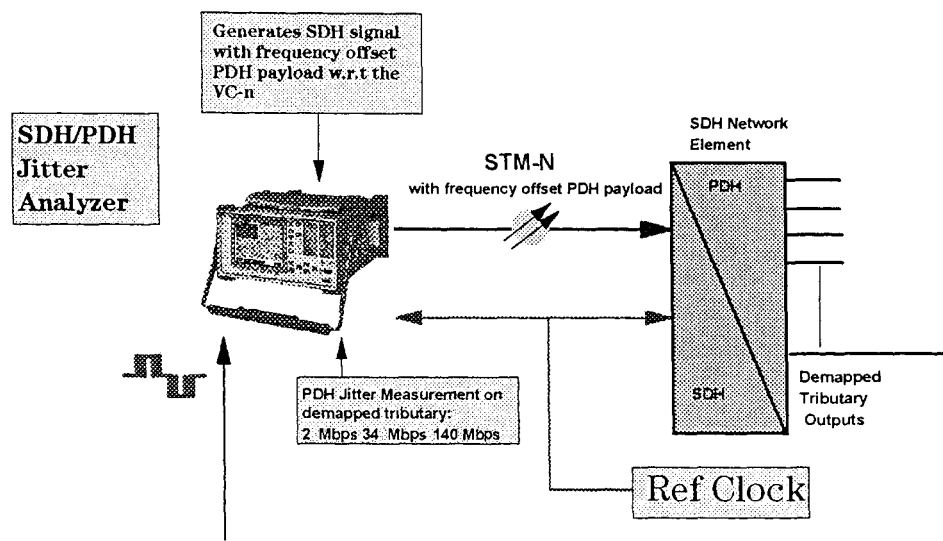


Figure 23.21 Demapping jitter test setup.

TABLE 23.8 ITU-T G.783 Demapping Jitter Specification.

Payload	Offset Range (ppm)	Measurement Bandwidth	Maximum Jitter (UIpp)
2 Mbps	±50	18–100 kHz <sup>1</sup>	0.075
34 Mbps	±20	10–800 kHz <sup>1</sup>	0.075
140 Mbps	±15	10–3500 kHz <sup>1</sup>	0.075

1. Equivalent to measurement filter “LP + HP2”

## 518 Basic Telecommunications Technologies

analyzer's measurement accuracy and intrinsic jitter specifications must be significantly less than this value (preferably a minimum of three times better).

### 23.12 Wander

Wander causes great concern in SDH networks. Wander on the 2 Mbps timing reference can lead to synchronization problems and, therefore, pointer movements. This in turn leads to jitter on the tributary outputs, which can lead to bit errors, frame slips, and loss of data. The trouble with wander is that it is extremely difficult to eliminate. This is because the source of wander is environmental change; as such, wander cannot be eliminated simply by careful network design and equipment selection. Appropriate criteria can minimize its effect, such as by increasing the size of elastic buffers in network equipment, or by employing other strategies.

Wander is defined as phase variations that occur with frequencies less than about 10 Hz. The ITU-T definition of wander is, "The long-term variations of the significant instants of a digital signal from their ideal position in time (where long-term implies that these variations are of frequency less than 10 Hz)." Slowly varying climatic effects, such as temperature, result in a "wandering" of the received pulse stream. Pulse streams varying with a period of one day are referred to as *diurnal wander*, and over one year as *annual wander*. Examples of potential sources of wander are changes in optical fiber temperature and drift in a regenerator's laser wavelength.

The ITU-T draft recommendation G.813, previously referred to as G.81s, specifies the wander performance of SDH equipment slave clocks. The specification limits are for *Maximum Time Interval Error* (MTIE) and *Time Deviation* (TDEV). MTIE is the maximum peak-to-peak delay variation of a given timing signal with respect to an ideal timing signal within an observation time ( $t$ ) for all observation times of that length within the measurement period ( $T$ ). TDEV is a measure of the expected time variation of a signal as a function of integration time.

TDEV also can provide information about the spectral content of the phase noise of a signal. This specification covers wander generation, wander transfer, and wander tolerance. Wander/noise generation, as specified by the ITU-T, represents the amount of noise produced at the output when there is an ideal reference input or the clock is in holdover state. Wander/noise tolerance is the minimum phase noise or wander level at the input of the clock while maintaining the clock within prescribed performance limits.

### 23.13 Additional Reading

Trischitta, Patrick R., and Varma, Eve. *Jitter in Digital Transmission Systems*. (Norwood, Mass.: Artech House, 1989.)

Cook, Tommy. "SDH: Pointer Problems." *Telecommunications Magazine*, August 1994.

Hewlett-Packard. *Tributary Jitter Testing of SDH Network Equipment using ITU-T G.783 Pointer Sequences*. (HP 37717B/C Product Note: 5965-4862E.)

———. *Automatic Verification of Network Equipment to ITU-T Jitter Tolerance Standards*. (HP 37717B/C Product Note: 5965-4863E.)

### 23.14 Standards Documents

ITU-T O.171, "Timing Jitter and Wander Measuring Equipment for Digital Systems."

ITU-T G.783, "Characteristics of SDH Equipment Functional Blocks."

ITU-T G.823, "The Control of Jitter and Wander within Digital Networks based on the 2048 kbps Hierarchy."

ITU-T G.825, "The Control of Jitter and Wander within Digital Networks based on SDH."

ITU-T G.958, "Digital Line Systems based on SDH for use on Optical Fiber Cables."

ETSI TM-1015





## Protocol Analysis

**Stephen Witt**

*Hewlett-Packard Co., Colorado Springs, Colorado*

### 24.1 Introduction

What is a protocol? Why does it need analysis? Don't networks just work? Why are there so many different protocols, standards, and networking technologies? The field of computer networking is complex and becoming increasingly so.

Computer networks are made up of many different computer systems, applications, and network topologies. The capital investment in cabling and transmission infrastructure is massive. The number of users demanding access to computer networks is ever-increasing, and these users are demanding more bandwidth, increased performance, and new applications. There is a constant stream of new equipment and services being introduced in the marketplace. In this complex environment, computer networking is possible only when equipment and services vendors adhere to standards covering protocols, physical connectors, electrical interfaces, topologies, and data formats. Protocol analysis is used to ensure that the products implemented according to these standards behave according to specification.

#### 24.1.1 Protocol definition

In general usage, a protocol is a code or a set of rules specifying the correct procedure for a diplomatic exchange. In computer network terminology, a *protocol* is a specific set of rules, procedures, and conventions defining the format and timing of data transmission between devices connected to a network. Protocols are defined so that devices communicating on a network can exchange information in a useful and efficient manner. Protocols provide a means for exchanging information so that computer systems can provide services and applications to end users. The wide range of protocols stems from the wide range of transmission media in use, the wide range of applications and services available to end users, and the many independent organizations and vendors creating protocol standards.

Protocols handle synchronization, addressing, error correction, header and control information, data transfer, routing, fragmentation and reassembly, encapsulation, and flow control. Protocols can be categorized in many ways, with each specific protocol having several applicable classifications. Some of the more important distinctions in protocols include:

- Message framing
- Delivery mechanism
- Timing
- Control

**Message framing.** A *message* is logical unit of information that is transferred between two computer systems. Message framing can be bit-oriented, byte-oriented, or character-oriented. *Bit-oriented protocols* use specific bit patterns to delimit the start and end of message transfers, with the control information embedded in specific bit patterns. *Byte-oriented protocols* use a protocol header that includes a byte count that is transferred in the header. *Character-oriented protocols* use special control characters to signal the start and end of message transfers. The control characters depend on the data code being used on the network, such as ASCII (American Standard Code for Information Interchange) or EBCDIC (Extended Binary Coded Decimal Interchange Code).

**Delivery mechanism.** The delivery mechanism is either *connectionless* or *connection-oriented*. Connectionless service, also known as *datagram service*, does not use an established connection or circuit to transfer a frame through the network. Frames are sent to a destination computer system through any number of existing network paths, without any acknowledgment of receipt being returned. There is low overhead associated with connectionless services, but reliability can suffer if the underlying transmission network is unreliable. Connection-oriented service is based on a connection being established before information is exchanged; acknowledgment of received frames occurs. Connection-oriented service is inherently more reliable.

**Timing.** Timing in a computer network is either synchronous or asynchronous. *Synchronous networks* are clocked by a master clock and transmit data at repeated intervals. *Asynchronous networks* transfer information a byte at a time, with delays between transmissions. Synchronous networks are more pervasive. They are more expensive to implement but provide more efficient transmission.

**Control.** Control in a computer network is either master/slave or peer-to-peer. A *master/slave* network is controlled by a master node on the network that controls when other nodes on the network can transmit information. In a *peer-to-peer* environment, however, no permission is granted or required and any node can transmit information at any time. *Contention schemes* are implemented in peer-to-peer networks to handle conflicts.

### 24.1.2 The need for protocol analysis

In order for two applications running on two different computer systems to communicate with one another (such as a database application executing on a server and a

client application performing database queries), meaningful information must be continually, efficiently, and correctly exchanged. This requires a physical connection to exist, either twisted-pair copper wire, coaxial cable, optical fiber, or wireless transmission (radio or infrared). The physical characteristics and specifications of the transmission media must be standardized so that different computer systems can be electrically connected to one another.

The bit streams exchanged over the physical medium or media must be encoded such that the analog stream of information can be converted to digital signals. For two people to communicate effectively, they must speak the same language. Similarly, two computer systems must speak the same “language” in order to communicate. The bit stream therefore must conform to a standard that defines the encoding scheme, the bit order (least significant or most significant bit first), and the bit sense (a high value defined either as a 1 or a 0). Errors in the transmission must be detected and recovered; if necessary, the data must be retransmitted. A protocol analyzer is used to examine the bit stream and ensure that it conforms to the protocol standards that define the encoding schemes, bit sequences, and error conditions.

Once a bit stream can be transmitted and received, physical communication is established and information exchange can be accomplished. Information is exchanged in logical units; a protocol frame, packet, message, or cell is the logical unit transmitted on the physical infrastructure of a computer network. (In this chapter, the term *frame* will be used to mean any or all of these protocol-specific units.) Depending on the type of network and protocols, these frames of data are either fixed in size, such as the 53-byte cells used in ATM, or they can be variable in size, such as the 64- to 1518-byte frames used by Ethernet networks. The most fundamental aspect of protocol analysis is the collection and analysis of these frames.

In networks there usually is more than one path between devices. The frames containing the data therefore must be addressed properly so that they can traverse single or multiple routes through the network. Fragmentation and reassembly issues must be handled: frames are often disassembled and reassembled so that they can be of proper size and be encapsulated with proper header information, to ensure that intermediate and end devices in the network can manipulate them properly. The network also must handle error conditions such as nodes that stop responding, transmit errored frames or signals, or use excessive bandwidth. A protocol analyzer is used to examine the addresses of the frames, check fragmentation and reassembly, and investigate errors.

Connections are established so that communication is efficient. This prevents the communication channel from being redundantly set up each time a frame is transmitted. This is similar to keeping a voice line open for an entire telephone conversation between two people, rather than making a phone call for each sentence that is exchanged. To ensure this efficiency, connections or conversations are established by devices on the network so that the formal *handshaking* (i.e., negotiation of mutually acceptable communications parameters) doesn't have to be repeated for each information exchange. Protocol analysis includes scrutinizing protocol conversations for efficiency and errors.

The data that is transferred in the frames to the host system application must conform to an agreed-upon format, and if necessary it must be converted to an architecture-independent format so that both computer systems can read it. Each time a user enters a command, downloads a file, starts an application, or queries a database, the

preceding sequence of processes is repeated. A computer network is continuously performing these operations in order to execute an end user's application. Protocol analysis involves critically examining the formatted data that is exchanged between host system applications.

### 24.1.3 Protocol standards

Communication between devices connected to computer networks is controlled by transmission and protocol standards and recommendations. These standards are necessary so that equipment and services from different vendors can interoperate with one another. While standards can be defined and implemented in the private sector by computer and network component vendors (such as Cisco Systems, Hewlett-Packard, and IBM), most standards and recommendations are created by organizations including (but not limited to) ANSI, CCITT, ETSI, IEEE, ISO, and the ITU. The ATM Forum and the IETF are technical working bodies that develop standards for networking products.

**ANSI.** The American National Standards Institute (ANSI), located in New York, develops and publishes voluntary standards for use in the United States. ANSI is the U.S. representative to the ISO. ANSI publishes the specification for FDDI and Fiber Channel. ANSI works with the IEEE on ISO approval for IEEE LAN specifications.

**ATM Forum.** The ATM Forum is a technical working body located in Foster City, CA. It is an industry group working on ATM specifications.

**CCITT.** The Consultative Committee on International Telegraphy and Telephony (CCITT) is part of the ITU located in Geneva, and now is also known as the ITU-T. The CCITT makes recommendations about telephone and data communication interfaces. Members of the CCITT include private companies, international organizations, and the national telecommunication authorities. In many countries these authorities are the Post, Telegraph, and Telephone (PTT) administration. The CCITT is responsible for X.25, X.75, X.21, and ISDN.

**ETSI.** The European Telecommunications Standards Institute (ETSI) is located in Sophia Antipolis, France. It is the European equivalent of ANSI, defining telecommunications specifications for the common European community.

**IEEE.** The Institute of Electrical and Electronic Engineers (IEEE, pronounced "I-triple-E") is located in New York. It publishes telecommunication and computing standards. The IEEE is responsible for the LAN Physical and Data Link layer specifications such as IEEE 802.3 CSMA/CD and the IEEE 802.5 Token-Ring. The IEEE is a member of ANSI.

**IETF.** The Internet Engineering Task Force is a technical working body made up of industry representatives. They publish the standards for the TCP/IP suite of protocols. The standards are available on the Internet.

**ISO.** The International Standards Organization (ISO), located in Paris, publishes international data communication standards such as the OSI (Open Systems Interconnection). It is the prevailing international standards body.

**ITU.** The International Telecommunications Union (ITU) was established by the United Nations and publishes telecommunications standards. The ITU has replaced the CCITT as the world's leading telecommunications standards organization. The ITU specifies many of the ATM standards.

#### 24.1.4 Semantic and syntactic rules

Communication is possible only when there is an agreed-upon format for the exchange of information and when there is a common understanding of the content of the information being exchanged. Protocols therefore must be defined by both semantic and syntactic rules.

**Semantics.** Semantics refers to the meaning of the information in the data frame, including control information for coordination and error handling. An example of the semantic information in a frame is a request to establish a connection, initiated by one computer and sent to another computer.

**Syntax.** Syntax refers to the structure, arrangement, and order of the protocol, including data format and signal levels. An example of the syntax of a protocol is the relative position in the frame of a protocol field such as the network address.

#### 24.1.5 Protocol analysis

*Protocol analysis* is concerned with both the syntax and the semantics of the protocols. A *protocol analyzer* is a dedicated, special-purpose computer system that acts as a node on the network. Unlike a typical node, however, it monitors and captures all of the network traffic for analysis and testing.

The term “protocol analyzer” was created in the early 1980s to describe a new class of products dedicated to testing serial data communication networks. These early products provided a set of features that focused on examining the communication protocols and ensuring that they conform to the standards.

This class of products has evolved to include support for all types of computer and communications networks, architectures and protocols. They provide capabilities to compare data frames with the protocol standards (protocol decodes), load and stress networks with traffic generation, and monitor network performance with statistical analysis. The measurement capabilities have expanded to focus on a broad set of applications such as network troubleshooting, network performance monitoring, network planning, network security, protocol conformance testing, and network equipment development. These new applications go far beyond simply examining network traffic with protocol decodes.

Protocol analysis consists of specific tasks a user performs in order to troubleshoot network problems or monitor network performance. The user invokes specific measurements such as protocol decodes or protocol statistics in order to perform these tasks. For example, performance monitoring requires gathering statistical measurements such as network utilization, error activity, network layer utilization, and key node activity. To gather statistics on a particular node or set of nodes, it is necessary to set up filters based on specific network addresses. Since the human brain is excellent at multitasking and performing tasks by gathering different pieces of information and assimilating them into a problem solution, protocol analysis is implemented by simultaneously executing various measurements and presenting the results to the user.

The people who design, implement, and maintain large computer networks are constantly faced with network failures, performance bottlenecks, configuration problems, and network downtime. Network managers use a variety of tools and methodologies

to install and maintain networks. Protocol analysis tools give network managers a “window in to the network,” allowing them to view and analyze the network traffic.

As new networking technologies are introduced and higher-performance network components are delivered to the marketplace, there is constant pressure to drive costs down and increase performance. While new network architecture designs are inherently more reliable, the driving factors of cost and performance exclude building extensive test capability into the network itself. Additional testing tools such as network management systems, distributed monitoring systems, and dispatched protocol analyzers therefore are necessary to guarantee that network downtime is kept to a minimum. All of these tools employ protocol analysis to implement the necessary testing capabilities.

## 24.2 Protocols and Network Architectures

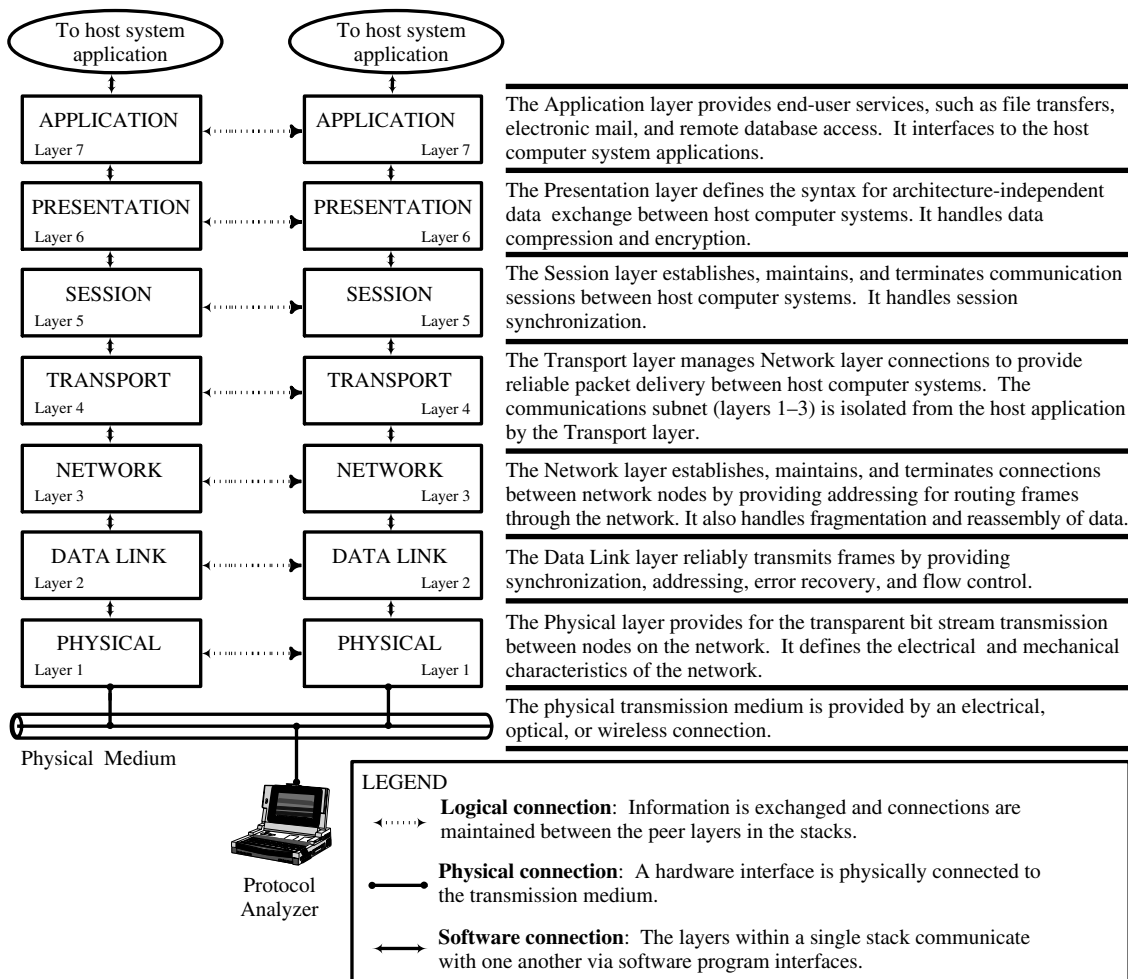
A *network architecture* is the design and structure of a computer network and is defined by its physical implementation, protocols, protocol stacks, and topology. Each of these is defined by specifications that are either proprietary, public, or standard. *Proprietary specifications* are developed by vendors and might or might not be open for use by others. An example is IBM’s Systems Network Architecture (SNA). *Public specifications* are implementations that have been widely deployed and have been absorbed into the public domain. The TCP/IP network architecture is an example of a public specification. National and international standards bodies define *standard specifications*. An example is the specification for the X.25 network, defined by the CCITT.

### 24.2.1 The OSI Reference Model

The International Organization for Standardization (ISO), located in Geneva, is responsible for many of the international computer networking standards. The ISO defined a model for computer communications networking called the Open Systems Interconnection Reference Model. This model, commonly called “the OSI model,” defines an open framework for two computer systems to communicate with one another via a communications network. The OSI model (Figure 24.1) defines a structured, hierarchical network architecture.

The OSI model consists of the *communications subnet* (protocol layers 1–3) and the services that interface to the applications executing in the host computer systems (protocol layers 4–7). The combined set of protocol layers 1–7 is often referred to as a *protocol stack*. Layer 7, the Application layer, is the interface to the user application executing in the host computer system. The protocol stack executes in a host computer system; each layer has a software interface to the layers below and above it. The only actual physical connection between network devices is at the Physical layer, where the interface hardware connects to the physical medium.

There is a logical connection, however, between corresponding layers in the two communicating protocol stacks. For example, the two Network layers (layer 3 in the OSI reference model) in two protocol stacks operate as if they were communicating directly with one another, when in reality they communicate by exchanging information through their respective Data Link layers (OSI layer 2) and Physical layers (OSI layer 1). Current network architectures are hierarchical, structured, and based



**Figure 24.1 The OSI Reference Model.** The OSI Reference Model defines an internationally accepted standard for computer network communications. Host-system applications exchange information in a hierarchical and structured manner.

in some manner on the OSI Reference Model. The functionality described in the OSI model is embodied in current network architectures, albeit at different layers or across multiple layers.

### 24.2.2 Network architectures

Most networks can be broadly categorized as either a local area network (LAN) or a wide area network (WAN). A LAN is limited by distance to a building or campus environment. It typically is used to connect client computers with file and application servers, as well as peripherals such as printers and plotters. A WAN typically covers a large physical distance and requires the communications technologies of telephone

companies to interconnect networks; a WAN is usually used to interconnect LANs. All of the network connections in a LAN environment are owned by the end user, while in a WAN the network connections are owned by the service provider.

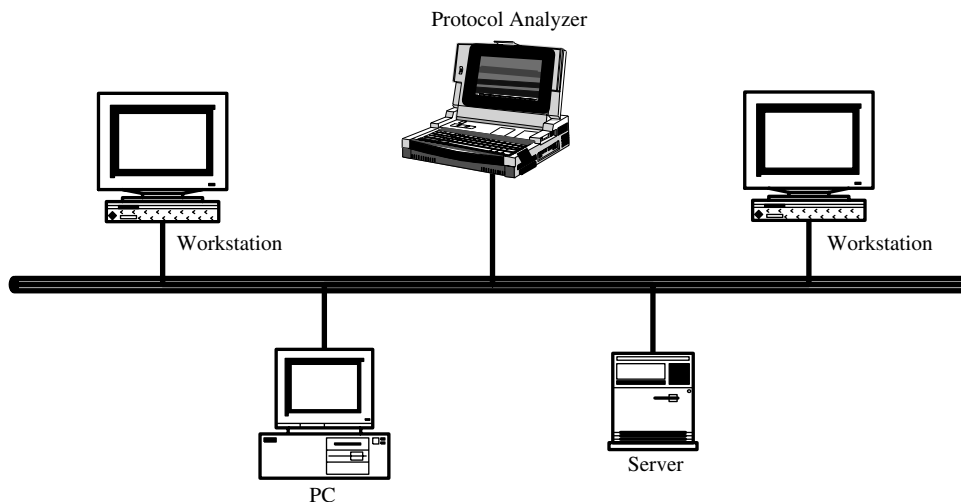
### 24.2.3 Network topologies

The network topology is the actual physical layout of the network. It is defined by the electrical configuration and the logical path of information flow. A protocol analyzer actually makes an electrical connection to the network under test. In certain network topologies, most notably a ring, the protocol analyzer optionally can be an active participant in the normal protocol at the MAC level (the lower part of the Data Link layer), or it can be a passive observer. Network topologies can be categorized as:

- Bus
- Star
- Ring
- Mesh
- Point-to-point

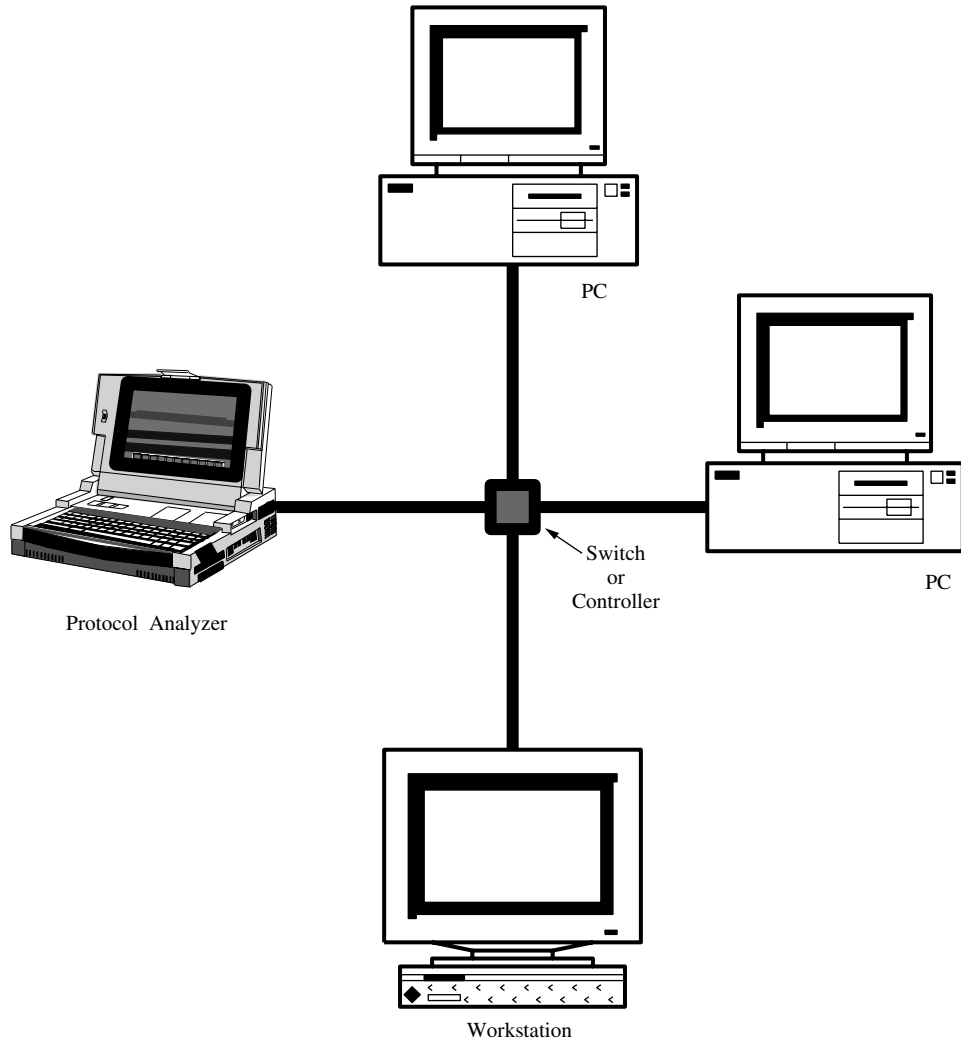
Today's complex internetworks typically are a combination of two or more of these five network topologies. For example, two LANs with a bus topology might be interconnected with a packet-switched WAN implemented with a mesh topology.

**Bus.** In a *bus topology*, all of the devices on the network share a common physical medium (Figure 24.2a). They operate much like a telephone party line, where any device on the network can receive, or listen, to all of the communications. This is



**Figure 24.2a Protocol analysis in a bus topology.** All devices have access to all of the network traffic in a bus topology. The protocol analyzer connects the same as any other node and can capture all traffic on the network.






**Figure 24.2b Protocol analysis in a star topology.** Network traffic in a star topology is switched from one node to another by a central controller or switching device, often called a hub. The protocol analyzer can capture the traffic that is switched to a particular node.

sometimes referred to as a *promiscuous environment*. In a bus environment the logical information path is the same as the electrical path; everything is shared and accessible by all of the nodes on the network. The most prevalent bus topology is 10Base2 (coax) Ethernet.

**Star.** In a *star topology*, all of the devices connected to the network are connected to one central device (Figure 24.2b). The electrical connections are switched so that devices on the network can communicate with one another. Ethernet running on twisted-pair (10Base-T) exemplifies a star topology.



**Figure 24.2c Protocol analysis in a ring topology.** All devices have access to all of the network traffic in a ring topology. The protocol analyzer connects to the ring in the same manner as any other node and can capture all traffic on the network.

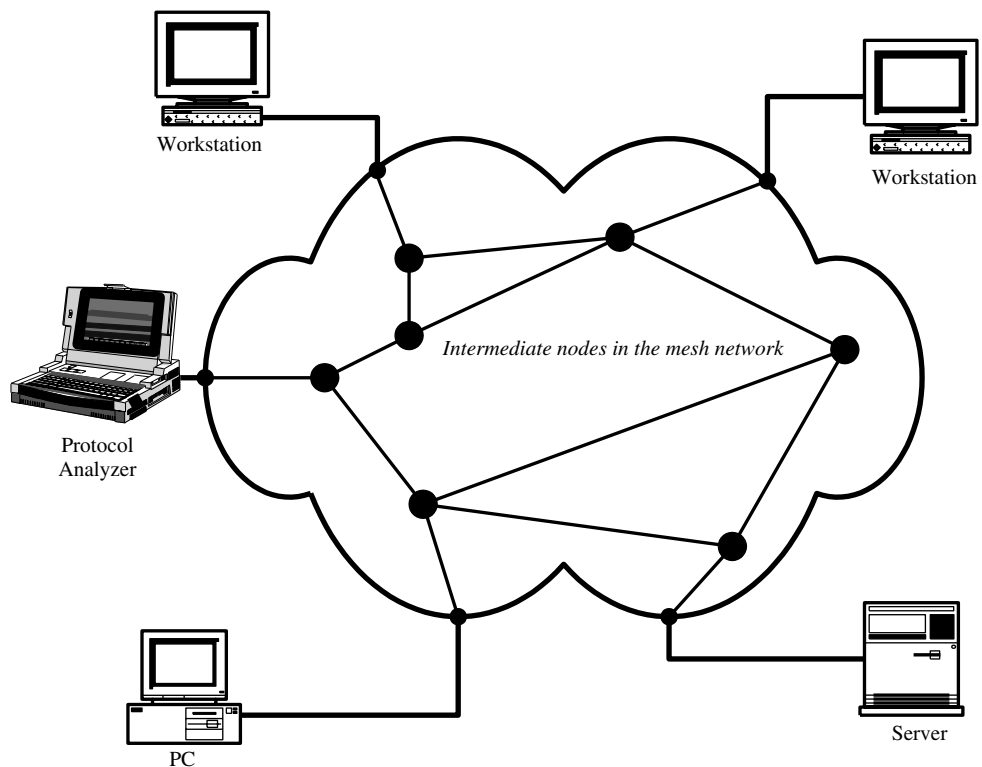
**Ring.** In a *ring topology*, all of the devices are connected to a common medium, and the medium is wired back on itself to form a ring (Figure 24.2c). In a ring environment all the devices have electrical access to one another. Typically the logical exchange of information is coordinated by passing a special frame called a *token*, which grants permission to transmit on the ring. When the token arrives at a device, the device can transmit information and then pass the token on to the next device, or, if it has nothing to send, it simply passes the token. FDDI and IBM's Token-Ring are examples of ring topology (though to outward appearances Token-Ring resembles a star).

**Mesh.** A mesh network is made up of many nodes interconnected so as to form multiple paths through the network (Figure 24.2d). Because multiple paths exist, a mesh network is inherently reliable. This reliability comes at a cost, however; meshed networks are more expensive and more difficult to implement. Mesh networks must establish routes through the network and switch information along these routes. There are three types of switching:

- Circuit switching
- Message switching
- Packet switching

**Circuit switching.** To allow two devices to communicate, a dedicated path or circuit is set up through the network. The two devices maintain the communications path until the transfer is complete. Circuit switching is reliable and provides good bandwidth, but it is more costly than packet switching because it does not make efficient use of the network bandwidth.

**Message switching.** An entire message is sent through the network in the same manner as a packet-switched network, i.e., with no dedicated path being established.




**Figure 24.2d Protocol analysis in a mesh topology.** Many different paths exist in a mesh network. A protocol analyzer can capture all of the traffic on the segment of the mesh to which it is connected.

At each intermediate node, however, the message is received and stored before being sent to the next node. Message switching sometimes is referred to as *store-and-forward* technology.

**Packet switching.** The data to be transferred is broken into a series of frames, or *packets*, that are individually transferred through the network. Packet switching is widely used to implement efficient internetworks, taking advantage of the bursty nature of data traffic and only establishing a network path for the particular packet being sent.

**Point-to-point.** In a *point-to-point topology*, two devices are connected directly with one another (Figure 24.2e). The electrical path and the logical flow of information is one simple and direct path.

The type of network topology not only determines where a protocol analyzer is connected to a network, it also determines what information is actually available for the protocol analyzer to analyze. Table 24.1 summarizes what information can be observed in each topology.



**Figure 24.2e Protocol analysis in a point-to-point network topology.** In a point-to-point topology it is necessary to disconnect the network connection and insert the protocol analyzer in series. The analyzer will capture all traffic sent between the two nodes.

**TABLE 24.1 Network Topology Analysis.**

Network Topology	Examples	What an Analyzer Can Analyze
Bus	10Base2 Ethernet	The analyzer accesses all of traffic on the segment.
Star	10Base-T Ethernet 100Base-T Fast Ethernet	The analyzer accesses only the traffic on the segment of the star to which it is attached.
Ring	Token-Ring FDDI	The analyzer sees all the traffic on the ring. The analyzer can be an active participant in the protocol, or it can electrically connect to the ring but not participate.
Mesh (switched)	X.25 frame relay ISDN ATM	The analyzer sees all traffic on the segment to which it is attached.
Point-to-point	modem 100Base-T Fast Ethernet	The analyzer accesses only the traffic between the two devices. The analyzer is connected in series with the network connection via a Y-cable, patch panel, or monitor point.

Other factors affecting a protocol analyzer's ability to analyze all of the traffic on a network are the interconnect devices embedded in the network. In a bridged LAN, for example, the bridges will pass all Network layer traffic, giving an analyzer access to all of the protocol information on the network. In a routed environment, however, the routers filter the traffic by protocol and by destination network address, so an analyzer will see only traffic that passes through the router onto the segment to which the protocol analyzer is attached. This concept is treated more fully in Chapter 16.

#### 24.2.4 Physical network implementations

The Physical layer in a protocol stack is responsible for the physical transmission of the bit stream. Physical layer specifications include mechanical, electrical, functional, and procedural aspects. The mechanical characteristics of the Physical layer include the construction of the actual network connectors. The electrical characteristics specify the voltage levels and timing parameters. The functional characteristics specify how analog signals are converted to digital signals, including data, control, timing, and grounding. The procedural characteristics specify how the data is actually transmitted and received. The Physical layer determines the network bandwidth, determines the network topology, and greatly influences the types of errors that will be encountered.

The Physical layer determines how a protocol analyzer is physically connected to the network under test. In point-to-point networks, for example there are three ways to connect such a device. The cable can be disconnected and the protocol analyzer can be connected serially between the two network connections. A second method is to use a patch panel to connect the protocol analyzer to the network under test. A patch panel is permanently configured into the network and allows connections to be made between the electrical interfaces of incoming and outgoing lines. Finally, some devices in a point-to-point network provide monitor ports specifically for testing. Table 24.2 describes the specifications of the most common physical network implementations in use today.

#### 24.2.5 Protocol stacks

The OSI model defines how networks and protocols are designed in order to operate in an "open" environment. The ISO suite of protocols is implemented to conform specifically to this model. Today, however, there are many different network architecture implementations from many different vendors, and although most are based on hierarchical models of interacting protocols, most do not conform exactly to the OSI model.

Multiple network architectures can coexist on one physical network, operating as multiple logical networks that typically do not interact with one another. Each different network architecture performs the same functions as described in the ISO model, but the different functions may be done at different layers in the protocol stack or multiple layers may be combined into one layer.

For example, Table 24.3 defines the protocol stacks required to implement a file transfer application in the ISO network architecture and the TCP/IP network

**TABLE 24.2 Common Network Physical Implementations.**

Name	Standard/Specification	Implementation	Data Rate
Carrier Sense Multiple Access with Collision Detect (CSMA/CD)	ANSI/IEEE 802.3, ISO 8802/3	1Base5, unshielded twisted-pair, baseband (StarLAN) 10Base2, 50Ω coaxial cable, baseband (thin Ethernet) 10Base5, 50Ω coaxial cable, baseband (thick LAN) 10Base-T, unshielded twisted-pair, baseband 10Broad36, 75Ω coaxial cable, broadband	1 Mbps 10 Mbps 10 Mbps 10 Mbps 10 Mbps
Token-passing bus	ANSI/IEEE 802.4, ISO 8802/4	Phase continuous carrier band Phase coherent carrier band Broadband	1 Mbps 5, 10 Mbps 1, 5, 10 Mbps
Token-Ring	ANSI/IEEE 802.5, ISO 8802/5	Shielded twisted-pair Category 3 or 5	1, 4, 16 Mbps
VG AnyLAN	IEEE 802.12	100Base-VG AnyLAN	100 Mbps
Fast Ethernet	IEEE 802.14	100Base-T4, unshielded twisted-pair 100Base-TX, shielded twisted-pair or Category 5, unshielded twisted-pair/full-duplex 100Base-FX, dual multimode fiber/full-duplex	100 Mbps / 200 Mbps 100 Mbps / 200 Mbps
FDDI (Fiber Distributed Data Interface)	ANSI ASC X329.5/ISO 9314-1	Single-mode fiber, multimode fiber, electrical/dual ring	100/200 Mbps
MAN (Metropolitan Area Network)	IEEE 802.6, ANSI X3T.9	DS-1 Physical Layer Convergence Procedure (PLCP), DS-1	1.544 Mbps
ITU-T G.804 PDH (Plesiochronous Digital Hierarchy)	ITU-T G.703, ITU-T G.704, ANSI T1.403 ITU-T G.704 ITU-T G.703 ITU-T G.832, ITU-T G.704 ITU-T G.704 ITU-T G.832	T1 (DS-1), electrical CEPT E1, electrical J2, electrical, single-mode fiber CEPT E3, electrical T3 (DS-3), electrical CEPT E4, electrical	1.544 Mbps 2.048 Mbps 6.312 Mbps 34.368 Mbps 44.736 Mbps 139.264 Mbps
T1.105 SONET, Synchronous Optical Network or G.708 SDH, Synchronous Digital Hierarchy	OC-1, STS-1 OC-3, STM-1, STS-3c OC-12, STM-4c, STS-12c OC-48, STM-16c, STS-48c	Single-mode fiber Unshielded twisted-pair, shielded twisted-pair, multimode fiber, single-mode fiber, electrical Single-mode fiber Single-mode fiber	51.840 Mbps 155.52 Mbps 622.08 Mbps 2.5 Gbps
TAXI	FDDI 4b/5b PMD (Physical Medium-Dependent)	Multimode fiber	100 Mbps 140 Mbps

Fiber Channel	Fiber Channel 8b/10b PMD (Physical-Medium-Dependent)	Shielded twisted-pair, multimode fiber	155 Mbps
IBM Block Encoded	IBM 4b/5b PMD (Physical Medium-Dependent)	Unshielded twisted-pair Category 3	25.6 Mbps
HSSI, High-Speed Serial Interface	ANSI X3T.9	Electrical	45 Mbps
V-series	V.35	Electrical	64 Kpbs
	RS-232	Electrical	64 Kpbs
	RS-449	Electrical	64 Kpbs

**TABLE 24.3 Comparison of ISO and TCP/IP Protocol Stacks.**

OSI Reference Model Protocol Layer	ISO Protocol Stack Protocols and Protocol Standards	TCP/IP Protocol Stack Protocols and Protocol Standards
Application	FTAM (File Transfer Access and Management) ISO 8571/8572	FTP (File Transfer Protocol)
Presentation	COPP (Connection-Oriented Presentation Protocol) ISO 8823	
Session	COSP (Connection-Oriented Session Protocol) ISO 8327	
Transport	COTP (Connection-Oriented Transport Protocol) ISO 8073	TCP (Transmission Control Protocol)
Network	CLNP (Connectionless Network Protocol) ISO 8473	IP (Internet Protocol)
Data Link		
Logical Link Control	LLC (Logical Link Control) IEEE 802.2	LLC (Logical Link Control) IEEE 802.2
Media Access Control	CSMA/CD Ethernet IEEE 802.3	CSMA/CD Ethernet IEEE 802.3
Physical	10Base-T	10Base-T

architecture, across a 10Base-T connection. TCP/IP, the protocol widely used on the Internet, originated with the U.S. Defense Advanced Research Projects Agency (DARPA). The TCP/IP protocol stack is a proven and practical way to implement a network architecture. The OSI protocol stack represents a very thorough and pedantic approach to designing a network architecture.

A network architecture is based on a network operating system (NOS), which makes use of a protocol stack. An NOS is a software program, typically used in a LAN environment, that provides computer networking services, interfacing the computer network and the computer system. Table 24.4 describes some of the most common protocol stacks.

### 24.2.6 Protocol frame example

At each layer in a network architecture, additional protocols are encapsulated within the frame. This concept is shown in Figure 24.3, which details one possible top-layer application protocol, the FTP or File Transfer Protocol, that can execute on the TCP/IP protocol stack. The FTP data is the actual file being transferred to a host system application. FTP is the data or *payload* portion of the Transmission Control Protocol, which is the layer 4 transport protocol in this particular protocol stack. TCP is in turn the data or payload for IP, which is the layer 3 network protocol. IP is the data or payload in the IEEE 802.2 Logical Link Control (LLC) protocol, which is the upper portion of the layer 2 Data Link control layer. Finally, the LLC is the data or payload in the IEEE 802.3 Media Access Control (MAC) layer.

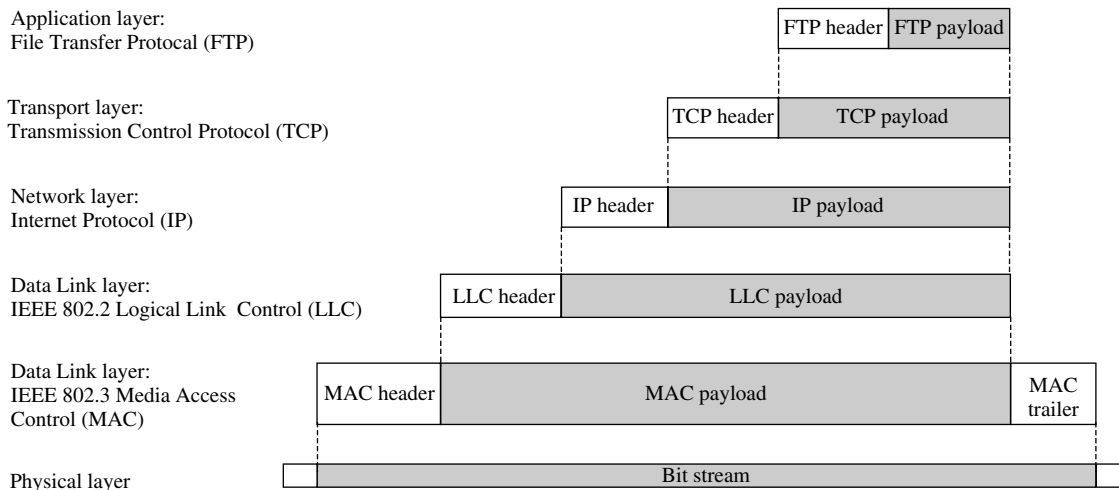
This example is one of literally hundreds of different protocol implementations that are detailed in numerous protocol handbooks and specifications. It is not appropriate in this chapter to fully describe all of the different protocol fields as this



**TABLE 24.4 Common Protocol Stacks.**

Protocol Stack	Description
AppleTalk	Apple Computer's proprietary LAN stack for connecting Macintosh computers and peripherals. It is similar to the OSI model. It was originally offered on a proprietary CSMA/CD interface called LocalTalk, but since has been integrated with mainstream network interfaces.
ATM	ATM, Asynchronous Transfer Mode, provides Transport-level services over a cell-based switching network.
Banyan VINES	VINES is a proprietary network operating system designed for large, heterogeneous, enterprise networks. It was developed by Banyan Systems, Inc., and is based on Unix.
DECnet	Digital Equipment Corporation's open network architecture, based on the ISO protocol stack. DECnet Phase IV is based on DEC protocols and DECnet Phase V is based on ISO protocols.
Frame relay	Frame relay, defined by the Frame Relay Forum and originally based on the ANSI and ITU-T specification, provides Network-layer services for a packet-switched network.
IBM/SNA	IBM's proprietary network architecture, Systems Network Architecture, is widely used in mainframe applications. It was defined before the OSI reference model. While it is a structured and hierarchical architecture, it does not correspond to the OSI model.
ISDN	ISDN, Integrated Services Digital Network provides Network-layer services for voice and data transmission.
ISO	The ISO protocol stack is an open system strictly adhering to the OSI Reference Model. All seven layers are implemented to the ISO specifications.
LAN Manager	A network operating system for LANs, designed by Microsoft and 3Com.
MAP	Manufacturing Automation Protocol (MAP) was designed for use in manufacturing environments. It is based on the ISO protocol stack running on the IEEE 802.4 token-passing bus.
Novell NetWare	NetWare is a network operating system for PC LANs and enterprise networks, produced by Novell. Originally it was derived from XNS.
Sun	Sun Microsystems Network File System (NFS) provides a set of network file system services that execute on the TCP/IP stack.
TCP/IP	TCP/IP was originally developed by the Department of Defense's Advanced Research Projects Agency (DARPA) to create a multivendor network connecting different computers across a variety of networks. It maps into to the OSI model with applications running directly on the Transport layer. It has become the predominant network architecture and forms the basis for the Internet.
X Window	Developed by the MIT X Consortium (a working group of computer companies) to implement a networked graphical windowing system running on Unix workstations. X Window executes on the TCP/IP stack.
X.25	A packet-switched network standard developed by the CCITT to provide Network layer services. It is widely used in public data networks.
XNS	Xerox Network System (XNS) was originally developed by Xerox. It provides a set of transport and network services that have been used by other vendors, such as Novell, as a basis for applications. XNS was used as the basis for the OSI Reference Model.

## 538 Basic Telecommunications Technologies



**Figure 24.3 A typical frame encapsulation.** All protocol information from a particular layer in the protocol stack is encapsulated as data in the next lower layer of the stack. The payload (shown shaded at each layer) is the encapsulated protocol from the next higher layer. This example, though specific to layers 1–4 and 7 of the ISO model, generalizes to other protocol stacks.

particular example is given to show a representation of how frames are constructed and how protocols are encapsulated within one another.

### 24.2.7 Protocol functions

Each network architecture has its own protocol stack, made up of different protocols that in many cases are performing equivalent functions to those in the OSI model. The functionality of most stacks is represented in the OSI Reference Model.

Despite the wide variety of protocols, however, some basic functions performed by protocols are common to all. These perform the “real work” in a network architecture and often are the points where failure occurs. Thus, they are the points of interest for performance monitoring. Protocol analysis is used to examine the protocol information that performs these functions. Functions that are of particular interest in performing protocol analysis include:

- Synchronization
- Addressing
- Error correction
- Header and control
- Payload (or data)
- Routing
- Fragmentation and reassembly
- Encapsulations

- Flow control
- Synchronization

**Synchronization.** Synchronization is used to determine the beginning and end of a frame of information. It is performed with signaling information in the form of bit fields, character sequences, and timeslots. There is a wide variety of schemes in use to synchronize frame transmission on networks.

**Addressing.** Frames contain source and destination addresses. In the Ethernet frame example in the preceding section, this is done at the Data Link layer. Addressing also is performed at the Network layer with logical addresses called *network addresses*, which are the addresses of the source and destination computer systems. At the Data Link layer (layer 2), addressing is used to ensure point-to-point transfer of information from one physical interconnect device on the network to the next; at the Network layer (layer 3), addressing is used to ensure the end-to-end transfer of frames from the source computer system to the destination computer system. Many interconnect devices, such as routers, manipulate the addresses to facilitate the routing of frames through the network. Additional protocols are used to resolve the address mappings.

**Error correction.** Frame transmissions are always susceptible to noise that makes its way into the signal. If this occurs, the protocols must detect the error and then either correct it automatically, request a retransmission, or ignore the frame. The detection often is accomplished by using *Frame Check Sequences* (FCS) that are added to the data in the frame. The FCS is calculated using a *Cyclic Redundancy Check* (CRC), which is a polynomial equation calculated on the incoming data to determine if all the bytes in the transmitted frame match the FCS value. The FCS is calculated by the sender and then, upon receipt of the frame, recalculated by the receiver; the receiver then makes a comparison to determine if errors have occurred. In the Ethernet frame example, the FCS is calculated at the Data Link layer.

**Header and control.** The header and control fields of the protocol frames contain information necessary to manage the protocol operation. These include priority, routing, address, command and response information, and type fields.

**Payload (data).** Most frames are sent so that information can be transmitted between two applications. Therefore most protocol frames on a network contain a payload or data field that ultimately is received by the host system application. Each protocol layer treats the information received from the layer above as data. There also are frames sent to facilitate network operation. These contain network-specific information and do not include a payload or data field.

**Routing.** Routers are used to connect different local and wide area networks together and provide access to a larger internetwork. Internetworks are made up of LANs and WANs that often are implemented with dissimilar protocols and network architectures. Routing is implemented to provide a translation from one protocol to another. Frames transmitted on the network contain routing information to facilitate this transfer.

**Fragmentation and reassembly.** In order to ensure that frames can be efficiently transmitted through an internetwork, it is often necessary to limit their size.

There also are many cases where the information to be transmitted is larger than what a single frame can handle. In either case, it is necessary to break a frame into several fragments, which are then transmitted through the network and reassembled at the destination. The frames contain information that allows the destination device to determine whether all of the fragments have been received and reassemble them in the order they were sent.

**Encapsulation.** Many frames on a network have to traverse several intermediate networks before they arrive at their destination. In order to accomplish this, the frame must be encapsulated within another frame of a different protocol type. For example, TCP/IP frames may be encapsulated in frame relay frames, or ISO transport frames may be encapsulated in X.25 frames.

**Flow control.** Different nodes on a network often will be capable of receiving and transmitting data at different rates; therefore the two devices will negotiate a transfer rate that prevents data loss. It also may be necessary to wait for the destination computer system to become available to receive information.

### 24.3 Protocol Analysis and the Life Cycle of Networks

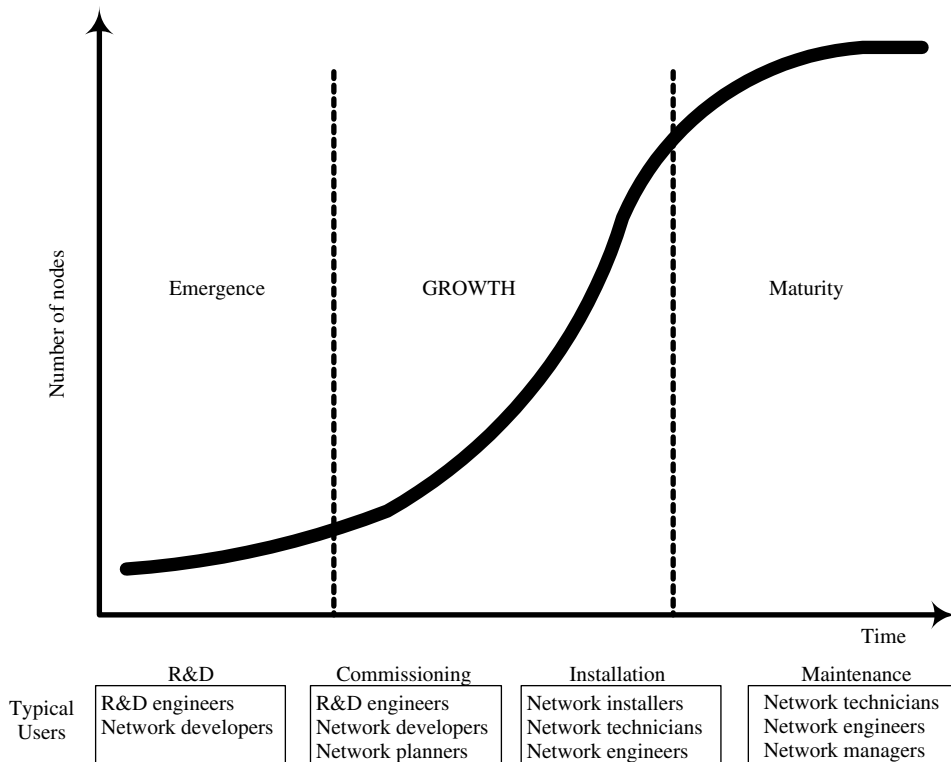
There is a life cycle for the design and deployment of networking technologies (Figure 24.4a). Network technologies begin in the *emerging phase*. At this point there are few network installations and very few users. The networking technology is still under development while network equipment vendors develop products. Once the technology proves to be viable it enters the *growth phase*, where it is rapidly adopted by end users. The number of installed networks and end user nodes grow rapidly during this phase. Once the technology is widely deployed it enters the *mature phase*, where it is widely accepted as a safe, proven technology.

#### 24.3.1 Network life cycles and analysis needs

As far as it concerns protocol analysis, there are four distinct stages, each with different testing needs (Figure 24.4b):

1. Research and development
2. Commissioning
3. Installation
4. Maintenance

**Research and development.** The R&D stage typically occurs during the emerging and growth phases of the network life cycle. During this stage engineers and network developers are implementing new network architectures and network equipment. Protocol analysis tools are used extensively for these design activities; therefore they must excel in many test applications. For example, protocol analyzers must provide conformance testing that verifies that the new designs meet the protocol standards. Protocol analyzers also are used in bench-top design applications for interactive simulation scenarios for testing new designs, which can include creating



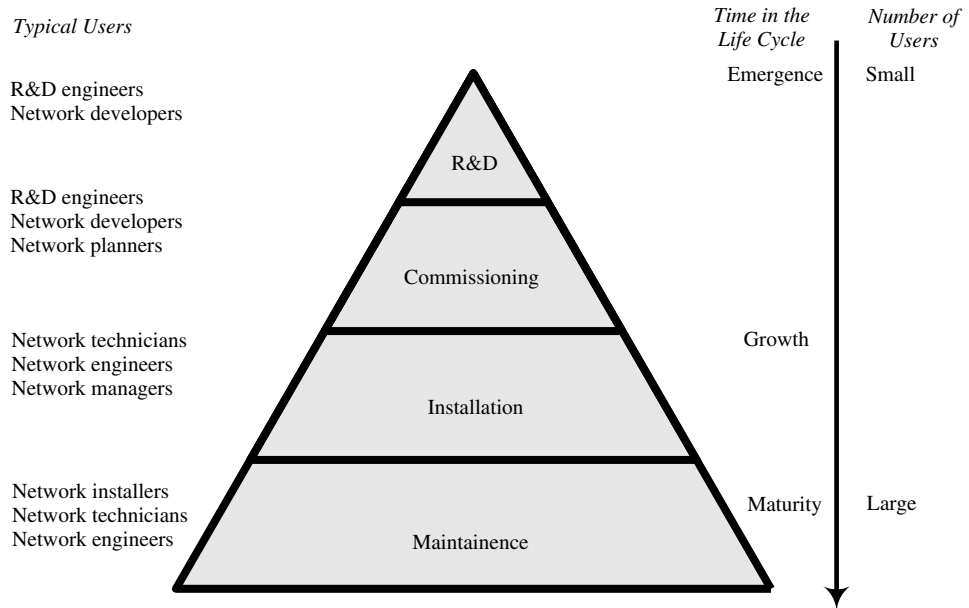
**Figure 24.4a Network technology deployment.** The number of installed nodes increases dramatically as a new network technology transitions from the emergence phase to the growth phase. As the technology matures, the increase in new installations levels off.

different error, traffic, and bandwidth profiles. Finally, the R&D designers use protocol analysis tools to execute extensive test suites of specific error and traffic scenarios to verify the performance of hardware and software implementations.

**Commissioning.** When new network equipment is installed (a large ATM switch, for example), it goes through a commissioning stage. This involves verifying the operation of the switch in the actual network. The protocol analysis tools used for commissioning usually are the same as those used for research and development. The focus is on stress testing the new piece of equipment in a live situation and verifying that it interoperates with the rest of the network. In many network commissioning scenarios there is no existing network to use for testing the new equipment; the network is still in its infancy. Protocol analysis tools therefore are critical for simulating a typical network.

**Installation.** When networks are installed, a new cabling infrastructure is put in place, modified, or extended. In addition, the switches, routers, bridges, repeaters, and other interconnect devices necessary to interconnect the network are installed. These devices are configured, the necessary software is installed, and finally the network is brought online. Protocol analysis tools are used to verify the proper operation of the

## 542 Basic Telecommunications Technologies



**Figure 24.4b The life cycle of networks.** As a new network technology is deployed, the number of users increases at each stage of deployment.

new network. Typically the physical medium or media are tested with transmission test tools; then protocol services are tested with a protocol analyzer. The protocol testing focuses on verifying proper protocol exchanges and simulating network scenarios.

**Maintenance.** Once installed, the network must be maintained, extended, and optimized for performance. Protocol analysis capabilities are included in network management systems, distributed monitoring tools, and dispatched analyzers. Testing in the maintenance phase concentrates on gathering and examining protocol frames to find errors and on gathering performance statistics. Often “finger-pointing” situations develop between different vendors who have implemented different parts of the network. These situations can be resolved by collecting data that isolates the problem with a protocol analyzer.

## 24.4 Tools for Managing Networks

Protocol analysis measurement capability is embedded into many commercially available products, providing users with the ability to monitor network performance and to troubleshoot network problems. Figure 24.5 shows the four classes of tools used to troubleshoot networks:

- Network management systems
- Distributed monitoring systems
- Protocol analyzers
- Handheld test tools

**Network management systems.** Network management systems are the most comprehensive and sophisticated tools for maintaining networks. They provide a network-wide view, showing all the segments and interconnections on the internet-work, and allowing the user to troubleshoot the entire enterprise. Network management tools are the most expensive of the four types of tools. Network management systems integrate many different applications that help manage the network and the systems connected thereto. Examples include device configuration and computer system resource monitoring. Protocol analysis applications are also one of many such applications.

**Distributed monitoring systems.** Distributed monitoring systems are a hybrid of dispatched protocol analyzers and network management systems. They typically are dedicated to performance monitoring and also provide troubleshooting capabilities. They consist of instrumentation that is distributed through the network, monitoring critical network segments to capture traffic statistics. This instrumentation performs some protocol analysis functions and sends the results back to a management application running on a Unix workstation or a PC. The information is sent via the network under test or through a telemetry network (typically a modem or a LAN) that is dedicated to the remote monitoring process.

**Protocol analyzers.** Protocol analyzers are portable tools that are “dispatched” to the troubleshooting site. They are carried to a segment or connection point on the network that is experiencing problems or requires testing. Dispatched protocol analyzers must be lightweight, compact, and easy to configure. They usually are built on a PC platform; they can be software applications running on the PC

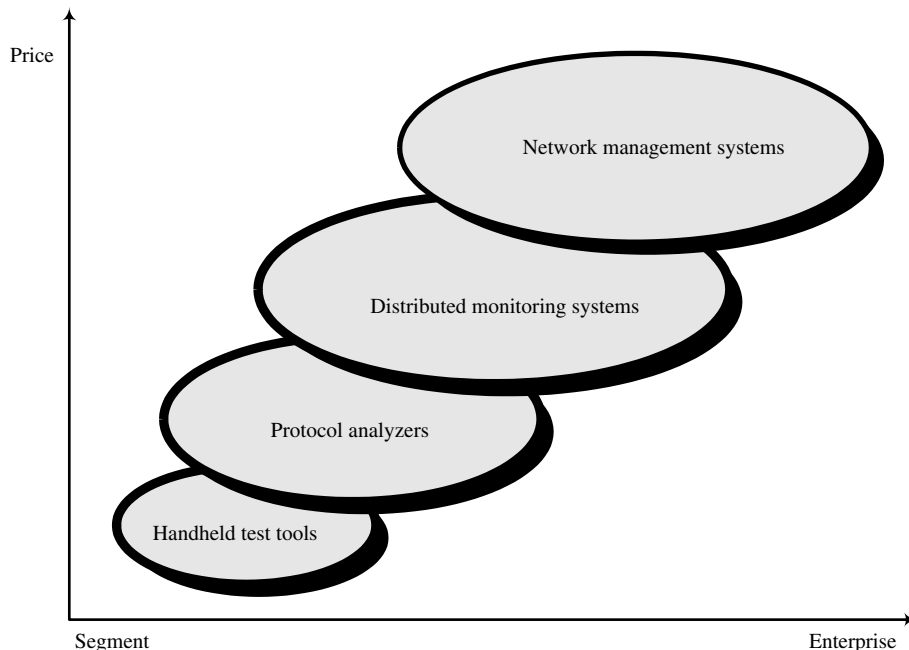


Figure 24.5 Network troubleshooting tools.

**544 Basic Telecommunications Technologies**

with a standard network interface card (NIC), an instrument that is connected to a PC controller, or a PC integrated with specialized hardware designed specifically to analyze the network under test.

**Handheld test tools.** Handheld test tools typically perform Physical layer tests to verify connectivity at the physical level. Some handheld testers perform limited protocol measurements but do not capture data and perform the full set of protocol analysis measurements. These tools are the most specific in nature of the four categories described, and are useful for troubleshooting specific network problems such as physical cable faults or simple protocol problems.

**24.5 Protocol Analysis**

A *network fault* is any degradation in the expected service of the network. Examples of service degradation include an excessively high level of bit errors on the transmission medium, a single user monopolizing the network bandwidth, a misconfigured device, or a software defect in a device on the network. Regardless of the cause, the network manager's fundamental responsibility is to ensure that such problems are fixed and that the expected level of service is restored.

To accomplish this, the network manager must troubleshoot the network and isolate faults when the inevitable problems occur. Many of the difficult-to-detect, intermittent problems can be recreated only if the network is stressed by sending frames. Many problems are caused by the constantly changing configurations of the devices attached to the network, so the manager must manage the network device configurations. To ensure the bandwidth and performance that users demand, the manager must monitor network performance and plan accordingly for growth. Network managers also are concerned with the security of their networks and use protocol analysis tools to monitor for illegal or unauthorized frames on network segments.

**24.5.1 Typical protocol analysis applications**

Instances where protocol analysis is the appropriate technology include the following:

- Window into the network
- Fault isolation and troubleshooting
- Performance monitoring
- Network baselining
- Security
- Stress testing
- Network mapping
- Connectivity testing
- Conformance testing

**Window into the network.** Protocol analyzers allow users to see the type of traffic on the network. Many problems can be solved quickly by determining what



type of traffic is or is not present and whether protocol errors are occurring. Without a “window into the network,” network managers must rely on indirect measures to determine the network behavior, such as observing the attached nodes and the problems users are experiencing.

**Fault isolation and troubleshooting.** Most network problems, such as downtime, are solved by following a rigorous troubleshooting methodology. This methodology (not unique to network troubleshooting) consists of observing that a problem has occurred, gathering data about the problem, formulating a hypothesis, and then proving or disproving the hypothesis. This process is repeated until the problems are solved. Protocol analysis is used in the troubleshooting process first for observing that a problem has occurred, and next for gathering data (using measurements such as protocol decodes and protocol statistics as described in section 24.6). The user then formulates a hypothesis and uses the protocol analysis measurements to confirm cause of the problem, ultimately leading to a solution. The protocol analyzers then can be used to confirm that the problem has indeed been repaired.

**Performance monitoring.** Determination of the current network utilization, the protocols in use, the errors occurring, the applications executing, and the users logged on is critical to understanding if the network is functioning properly or if problems such as insufficient capacity exist. Performance monitoring can be used over short time periods to troubleshoot problems, or it can be used over long time periods to determine traffic profiles and optimize the configuration and topology of the network.

**Network baselining.** Every network is unique; there are different applications, distinct traffic profiles, products from numerous vendors, and varying topologies. Network managers therefore must determine what is normal operation for their particular networks. A *network baseline* is performed to determine the profile of a particular network over time. A profile is made up of statistical data that includes a network map, the number of users, protocols in use, error information, and traffic levels. This information is recorded on a regular basis (typically daily or weekly) and compared to previously recorded results. The baselining information is used to generate reports describing the network topology, performance, and operation. It is used to evaluate network operation, isolate traffic-related problems, assess the impact of hardware and software changes, and plan for growth.

**Security.** Networks are interconnected on a global scale, meaning it is possible for networks to be illegally accessed. This illegal access can be knowingly performed by someone with criminal intent, or it can be the result of an errored configuration of a device on the network. Protocol analysis tools, with their powerful filtering, triggering, and decoding capabilities, can detect security violations.

**Stress testing.** Many errors on networks are intermittent and can be recreated only by generating traffic to stress network traffic levels or error levels, or by creating specific frame sequences and capturing all of the data. Stress testing a network and observing the results with protocol statistics and decodes allows many difficult problems to be detected.

**Network mapping.** Networks continually grow and change, so one of the big challenges facing network engineers and managers is determining the current topology and configuration. Protocol analysis tools are used to provide automatic node lists of all users connected to the network, as well as graphical maps of the nodes and

the internetworks. This information is used to facilitate the troubleshooting process by allowing users to be located quickly. It also is used as a reference to know the number and location of network users to plan for network growth.

**Connectivity testing.** Many problems are the result of not being able to establish a connection between two devices on the network. A protocol analyzer can become a node on the network and send frames (such as a ping) to a network device and then determine if a response was sent and, if so, the response time. In the case of multiple paths through the network, a more sophisticated test can determine which paths were taken. Connectivity can be verified in many WANs by executing a call placement sequence that establishes a call connection to enable a data transfer.

**Conformance testing.** Conformance testing is used to test data communications devices for conformance to specific standards. These conformance tests consist of a set of test suites (or scenarios) that exercise data communications equipment fully and identify procedural violations that will cause problems. These conformance tests are used by developers of data communications equipment and by carriers to prevent procedural errors before connection to the network is allowed. Conformance tests are based on the applicable protocol standard.

## 24.6 Protocol Analysis Measurements

Protocol analyzer functionality varies depending upon the network technology, such as LAN vs. WAN; the targeted user, such as R&D vs. installation; and the specific application, such as fault isolation vs. performance monitoring. Protocol analysis includes the entire set of measurements that allow a user to analyze the information on a computer communications network. But no single product provides all of the following measurements for all networking technologies:

- Protocol decodes
- Protocol statistics
- Expert analysis
- Traffic generation
- Bit error rate tests
- Simulation
- Programming language

Because protocol analysis requires combining and integrating the results of different measurements, the listed measurements typically are made in combination. Thus, protocol analyzers include different sets or combination of these measurements. A good user interface combines and presents pertinent information for the user in a maximally useful way.

### 24.6.1 Protocol decodes

*Protocol decodes*, also referred to as *packet traces*, interpret the bit streams being transmitted on the physical media. A protocol decode, as the name suggests, actu-

ally decodes the transmitted bit stream, identifying it and breaking it into information fields. The decoded fields are compared to the expected values in the protocol standards, and information is displayed as values, symbols, and text. If unexpected values are encountered, then an error is flagged on the decode display. Protocol decodes follow the individual conversations and point out the significance of the frames on the network by matching replies with requests, monitoring packet sequencing activity, and flagging errors. Protocol decodes let the user analyze data frames in detail by presenting the frames in a variety of formats.

Decodes typically are used for troubleshooting. The ability to troubleshoot a problem is directly related to the ability to capture the frames from the network. Protocol conversations are based on sequences of frames. If a protocol analyzer misses frames, then protocol conversations are not recorded accurately and the troubleshooter cannot tell which is at fault, the protocol analyzer or the network. It is therefore essential to capture all of the frames without dropping or missing any.

Since protocol analysis can be done in one of two ways, either real-time or post-process, most analyzers provide a method for capturing and storing frames for subsequent decoding. Some analyzers also provide the capability to decode frames as they are captured (in real time). This allows the user to display data traffic as it occurs, eliminating the need to stop the testing to find out what is happening on the network.

Network problems are isolated by starting at an overview level and then focusing in on the network traffic in ever-increasing detail. Since a protocol stack is essentially a virtual network operating on the physical medium, and many such virtual networks can coexist on one physical medium, protocol decodes can display all of the network traffic or the protocol decode can be filtered to isolate a single protocol stack or a particular protocol.

**Decode display formats.** Protocol decodes typically are displayed in three formats: a summary decode, a detailed decode, and a data decode (Figures 24.6a, 24.6b, and 24.6c). These examples show a frame captured from the World Wide Web. The HTTP protocol is executing over TCP/IP on a 10Base-T network. The summary decode shows a consolidated view of the frame information. The contents of each frame are summarized into one line on the display. This is used to quickly determine the type of traffic being transmitted on the network and to determine the upper-layer protocol conversations that are occurring.

The summary decode uses an underline cursor to aid in viewing a specific frame, as well as to indicate the frame that is examined in further detail using a “drill-down” user interface. In Figure 24.6a, frame number 4624 is underlined. The summary decode gives a reference number for the frame, a timestamp for the specific frame, the source and destination addresses (in this case the MAC addresses), and the description field of the information contained in the frame, including encapsulated protocols, address and control information, and error information.

The detailed decode (Figure 24.6b) breaks out all of the protocols in the underlined frame and identifies all of the fields in each protocol. It describes each field in the protocol layer and identifies all parameters with a textual description as well as the actual field values. This decode view is used to troubleshoot specific problems within a protocol or in a transaction between devices on a network.

Network Stack Summary Decode					
Control	Config	Actions	Format	Other displays	Print Help
Frame	Time	Source	Destination	Description	
4618	04.171544	Becky's PC	Cisco---01-FE-1D	TCP S=1120 D=8000	
4619	04.250532	HP-----68-38-E4	Broadcast	SAP R HP28699A 0000	
4620	04.260417	HP-----8B-4B-9E	HP-----10-DC-07	NETB Session data	
4621	04.263308	HP-----10-DC-07	HP-----8B-4B-9E	SMB R Write	
4622	04.265490	HP-----8B-4B-9E	HP-----10-DC-07	SMB C F=8162 Write	
4623	04.266372	HP-----8B-4B-9E	HP-----10-DC-07	NETB Session data	
4624	04.308424	Cisco---01-FE-1D	Becky's PC	HTTP R PORT=1122 I	
4625	04.313735	Cisco---01-FE-1D	Becky's PC	HTTP R PORT=1122 I	
4626	04.316341	Cisco---01-FE-1D	Becky's PC	HTTP R PORT=1122 I	
4627	04.317574	Becky's PC	Cisco---01-FE-1D	TCP S=1122 D=8000	
4628	04.317964	Cisco---01-FE-1D	Becky's PC	HTTP R PORT=1122 I	
4629	04.319472	Cisco---01-FE-1D	Becky's PC	HTTP R PORT=1122 I	
4630	04.320265	HP-----0B-2B-79	HP#-----00-00-01	Ethernet TYPE=8000	
4631	04.321837	Becky's PC	Cisco---01-FE-1D	TCP S=1122 D=8000	
4632	04.323591	HP-----10-DC-07	HP-----8B-4B-9E	TCP S=2905 D=4400	
4633	04.324065	HP-----8B-4B-9E	HP-----10-DC-07	NETB Session data	
4634	04.324947	HP-----8B-4B-9E	HP-----10-DC-07	NETB Session data	
4635	04.329378	HP-----10-DC-07	HP-----8B-4B-9E	TCP S=2905 D=4400	
4636	04.329796	HP-----8B-4B-9E	HP-----10-DC-07	NETB Session data	
4637	04.336286	HP-----10-DC-07	HP-----8B-4B-9E	SMB R Write	
4638	04.338457	HP-----8B-4B-9E	HP-----10-DC-07	SMB C F=8162 Write	
4639	04.339338	HP-----8B-4B-9E	HP-----10-DC-07	NETB Session data	
4640	04.340220	HP-----8B-4B-9E	HP-----10-DC-07	NETB Session data	
4641	04.343099	HP-----36-CB-34	Cisco---01-FE-1D	SNMP Trap	
4642	04.346605	HP-----10-DC-07	HP-----8B-4B-9E	TCP S=2905 D=4400	

Figure 24.6a Ethernet summary decode.

Network Stack Detailed Decode					
Control	Config	Actions	Format	Other displays	Print Help
Field	Value	Description			
Frame: 4624 Time: Aug 12@13:50:04.3084244 Length: 232					
HTTP					
Response OK					
HTTP/1.0 200 Document follows					
Date: Mon, 12 Aug 1996 19:52:52 GMT					
Server: NCSA/1.4					
Content-type: image/gif					
Last-modified: Thu, 16 May 1996 19:36:52 GMT					
Content-length: 5028					
TCP					
Source port	8000	HTTP			
Destination port	1122	User prog. port			
Sequence number	1580160001				
Ack number	1434663	Seqnum of next expected data			
Data offset	5	Number of 32-bit words in header			
Reserved	...-0000-00...-...				
Flags:	..01-1000				
Urgent flag	..0.....				
Ack flag	...1....	Ack number field is significant			
Push flag	...1...	Sender requests immediate delivery			
Reset flag	.....0..				
Syn flag	.....0.				
Fin flag	.....0				
Window	57344				

Figure 24.6b Ethernet detailed decode.

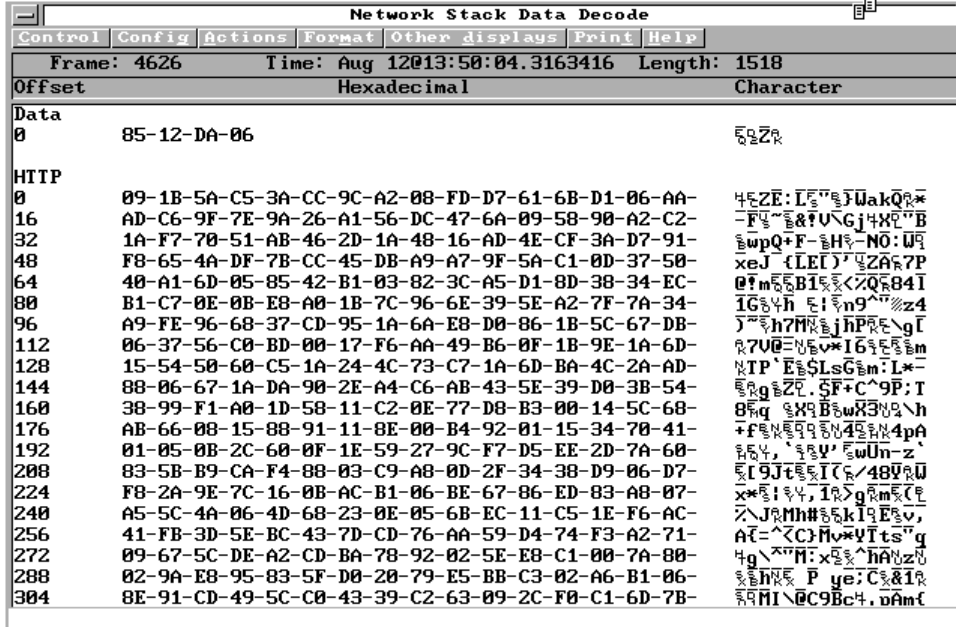


Figure 24.6c Ethernet data decode.

The data decode (Figure 24.6c) shows the values of every byte for frame 4626 in hexadecimal or a data code (such as ASCII or EBCDIC). It also can be used to read strings of text within packets.

**Node list and aliases.** An important feature of protocol analysis is the ability to generate a node list that contains the addresses of all the nodes attached to a network. The node list can be used as a reference to identify new nodes that have shown up, as well as to identify nodes that have gone quiet and possibly no longer exist. The user can create alias names for the obscure network addresses. Instead of referring to the PC in Becky’s office by its MAC address of 08-00-90-2F-3C-7A, for example, an alias can be entered in the node list and the PC will be referenced as “Becky’s PC” in the decodes and other measurements in the protocol analyzer. The alias also will be used in the user interface for configuration screens.

**Timestamps.** Each frame captured by a protocol analyzer from the network under test is provided a timestamp generated by an internal clock in the analyzer. The timestamp is used to determine the particular time at which the frame was received, allowing a user to make critical timing measurements, such as determining if a protocol time-out occurred. Many protocols are based on receiving specific responses within a specified amount of time. In Figure 24.6b, the top of the display shows that frame number 4624 was captured on Monday, August 12, at 13:50:04.3084244. (An abbreviated time stamp is shown in the summary decode.) Typical network troubleshooting involving timing is done at microsecond

( $\mu$ s) or millisecond (ms) resolution. Decodes show time values in one of three methods: absolute, relative, and delta.

- *Absolute time* is the actual value of the timestamp, indicating the real-world time (date, hour, minute, second, microsecond, millisecond, nanosecond) that the frame was received. It is used to determine the exact time that a frame was transmitted on the network.
- *Relative time* shows the frames with times offset from a base set by the user. The user selects a reference frame and the time values are shown as relative to that reference frame. It is used to determine the delay between frames on the network.
- *Delta time* mode shows the time values relative to the preceding time. This is used to determine the timing between subsequent frames on the network.

**Frame numbering.** Captured frames are provided with a sequential number as a convenient user reference. The frame numbers are completely arbitrary, not derived from the network under test nor part of the transmitted information. They do provide a convenient reference that makes it simpler to identify a particular frame from the thousands of frames captured in a single trace capture. For example, in Figure 24.6b the user can see in the upper left corner that he or she is examining frame number 4624 and quickly map that to the specific frame in the summary decode shown in Figure 24.6a.

As a frame is analyzed by a protocol analyzer, the analyzer automatically determines all of the protocols in the frame and invokes the correct methods for decoding the frame. The user does not need to preconfigure the decodes for specific frame types. In addition to the simple parsing of frames, many protocol analyzers will correlate state information from sequences of frames and display it as well. For example, when the final frame is assembled with the preceding frames of a fragmented TCP transmission, the final decode TCP frame will contain an “assembly complete” notification.

### 24.6.2 Protocol statistics

An average high-speed computer network, such as a 10 Mbps Ethernet, handles thousands of frames per second. It takes only seconds to fill the capture buffer of a protocol analyzer. To search for a specific network event, *capture filters* can be used to filter the data going into the buffer; alternatively, *triggers* can stop the capture buffer around a specific event on the network and invoke protocol decodes to troubleshoot the problem. This approach relies on knowing the cause of the trouble, however. In many cases, this is not known. Furthermore, while decodes pinpoint specific problems, they do little to identify trends or provide performance information. Thus the need for protocol statistics exists.

Protocol statistics reduce the volumes of captured data into meaningful information. Protocol statistics provide valuable insight for determining the performance of a network, pinpoint bottlenecks, isolate nodes or stations with errors, and identify stations on the network. Protocol analyzers keep track of hundreds of different statistics. The statistical information is displayed as one or more of:

- Histograms
- Tables

- Line graphs vs. time
- Matrices

There are major classifications of statistical measurements that are applicable to the full range of networks and protocols. Table 24.5a describes the major types of such measurements that can be made. This table is a list of “what” can be measured. Table 24.5b describes “how” it can be measured. The statistical measurement data

**TABLE 24.5a Types of Statistical Measurements.**

Statistical Measurement	Description	Examples
Application	Ascertain the usage distribution of the applications running on the network. Used to determine how a network is being utilized and provide information for optimizing the network.	Data base applications Graphics applications E-mail applications Internet access
Destination Addresses	Determine the type of destination address contained in a frame. Networks with high levels of broadcast traffic sent to all nodes on the network typically experience performance problems.	Unicast address Multicast address Broadcast address Source routing on a Token-Ring network
Error	Detect errors for the entire network, a particular node, or a logical channel. Devices creating errors affect the performance of the entire network and can create problems on other devices.	Transmission errors FCS errors Protocol errors Invalid frames
Frame Size	Calculate the average, minimum, and maximum frame sizes on the network. Used to determine if the network is being utilized with high data throughput (large frames), or if it is being overloaded with high overhead (many small frames).	Frame size distribution Average, minimum, and maximum frame sizes Illegal frame lengths
Node Discovery	Automatically discover the nodes and stations attached to the network and identify their MAC addresses, network address, and alias or “friendly” name.	Create a network node list Identify interconnect devices (e.g., routers) Identify new nodes Identify aged nodes
Number of Nodes	Calculate the number of devices using the network.	Current, average, minimum, and maximum count
Protocol	Ascertain the protocol and protocol stack distribution for the traffic running on the network. Used to determine how a network is being utilized and provide information for optimizing the network. Determine protocol usage, which stations are using bandwidth, and the respective amount of control protocols vs. data transfer protocols.	Distribution of protocols Distribution of protocol stacks Utilization of specific protocols (e.g., Network layer protocols)
Utilization	Calculate the amount of traffic on the network vs. the theoretical bandwidth available on the network. Performance problems and network failures typically increase as the network utilization approaches its theoretical limit.	Utilization (percent vs. time) Current, average, minimum, and maximum throughput

**TABLE 24.5b Methods for Filtering Statistical Data.**

Filter-Criteria	Description	Examples
Connection	Track the performance of conversation pairs of MAC addresses, network addresses, or subnet types. Track errors and bandwidth utilization by connection pair.	Throughput statistics between: File server and a router Client and server Two routers
Logical Channel	Track the performance of the virtual circuits established in a switched network. These statistics are good for isolating problems to a particular device or determining the performance of the switched network.	Throughput statistics for: X.25 Logical Channel Number (LCN) Frame relay Data Link Connection Identifier (DLCI)
Network	Statistics are summed for all of the traffic observed on the network. These statistics give a good overall indication of network performance and can provide a network baseline.	Utilization (percent vs. time) Total Node Count Total Error Count Average throughput Average frame size
Node	In a LAN environment, statistics are sorted on a per-node basis. These statistics are good for isolating problems to a particular node or for identifying the critical nodes such as routers, that have high bandwidth.	Frames received per node Frames transmitted per node Bytes received per node Bytes transmitted per node
Top Talkers	Identify the nodes on the network that are transmitting the most traffic. Network problems are often a function of traffic. Top talkers is also useful for identifying key nodes such as routers and servers.	Nodes generating most frames Nodes generating most IP traffic Nodes generating most file transfers

can be filtered by the criteria listed, with these measurements made by combining the entries listed in the two tables. For example, statistics can be filtered by a particular node, allowing error statistics to be collected for that node, to determine if the node is operating properly.

Figure 24.7 shows a summary statistical display for analyzing the performance of a 10Base-T network. It is typical of the types of measurements made on most networks. The upper left part of the display provides a history of the overall network utilization. In this case, the network utilization has ranged from 0 to 20 percent utilization over the last 3 minutes. Typical statistics measurements can be set up to display the last several minutes of traffic or, with coarser granularity, over one week of data.

This display also shows gauges indicating instantaneous and cumulative counts of collisions and errors. The current node count and bytes/frame are displayed. The percent of multicast vs. unicast vs. broadcast traffic is displayed in the Destination Addresses graph. The percent of different protocols on the network is shown in the Protocols graph. Finally, the average throughput (in average frames/sec) for critical nodes on the network is displayed in the Selected Nodes graph. Each of these measurements has a user-defined threshold so that visual alarms, as well as event log entries, can be made (see section 24.7.5).



### 24.6.3 Expert analysis

Troubleshooting computer networks is made complicated by the wide range of network architectures, protocols, and applications that are simultaneously in use on a typical network. Expert analysis reduces thousands of frames to a handful of significant events by examining the individual frames and the protocol conversations for indications of network problems. It watches continuously for router and bridge misconfigurations, slow file transfers, inefficient window sizes, connection resets, and many other problems. And it does this in real time for each protocol stack running on the network, as the network events occur.

Thus data is transformed into meaningful diagnostic information. The typical network engineer will have expertise in certain protocol stacks and certain applications, but no single person has expertise in all of the applications and protocols. Expert analysis therefore combines the knowledge of many networking experts into a single protocol analyzer. Expert analysis is performed by using the analysis capability (data capture, filters, triggers and actions, statistics, etc.) to monitor the network traffic. Figure 24.8 gives an example of one of the many Expert Analysis screens that simplify network troubleshooting.

Rather than setting filters, collecting frames, decoding the frames, and correlating sequences of frames, the expert analysis screen indicates events occurring on the network, their severity, and appropriate additional information needed to troubleshoot the problem. The events are categorized as either *normal*, *warning*, or *alert*. Normal events, such as the OSPF Router Identified, take note of normal occurrences that may be of particular interest. Warning- and alert-level events indicate an increasing level of severity that may suggest degraded network performance or catastrophic network failures.

The highlighted event in Figure 24.8 shows that two nodes, one with IP address 15.42.144.11 and one with IP address 15.6.74.53, are experiencing excessive TCP

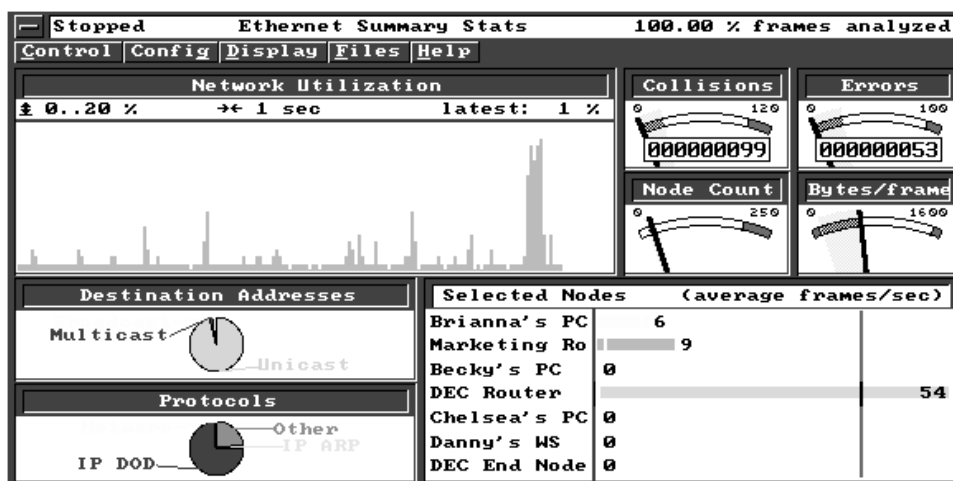


Figure 24.7 Protocol statistics.

```

Network Commentator/Eth.
Control Config Event Filter Print Help
Commentating on: ICMP, TCP/IP
OSPF: Router Identified [Normal] Aug 12@15:39:39.89502
Network Addr: 15.6.72.1
MAC Addr: 00-00-0C-01-FE-1D, Cisco---01-FE-1D
Number of OSPF routers identified: 1
Frame Number(s): 449
TCP: Low Window [Warning] Aug 12@15:40:10.7337211
15.6.76.130 ---> 15.6.72.205
Port: 9100 1955
Window Size: 0
Low Window Duration: 0:00:10.0
Frame Number(s): 2397 - 2397
TCP: Low Window [Warning] Aug 12@15:40:11.7337211
15.6.76.143 ---> 15.6.72.205
Port: 9100 3631
Window Size: 0
Low Window Duration: 0:00:06.9996299
Frame Number(s): 2845 - 2845
TCP: Excessive Retransmissions [Warning] Aug 12@15:40:35.272483
15.42.144.11 ---> 15.6.74.53
Port: 8088 36496
Rtx/Tx: 2 / 10
Frame Number(s): 7353 - 7356

```

Figure 24.8 Expert analysis.

retransmissions. This is a warning-level event, meaning that the network performance is being adversely affected. Most expert analysis applications provide context-sensitive on-line documentation that provides troubleshooting suggestions. They also provide drill-down capability to access additional measurements. In this particular case, frame numbers 7353–7356 in the capture buffer provide more information about the problem.

#### 24.6.4 Traffic generation

Many network problems are difficult to diagnose because they occur intermittently, often showing up only under peak load. Traffic generators provide the capability to simulate network problems by creating a network load that stresses the network or a particular device sending frame sequences on the network. In an ATM network, a traffic generator can be used to create cell loss and cell delay measurements.

Traffic generators also are used to stress a network or a particular device on the network under various loads. They can be used to increase the background traffic level on an in-service network, or they can be used to stress-test an out-of-service network or specific device. In either case, the user can measure the response of a device or an entire network under an increased load. Table 24.6 shows the parameters that can be used to configure a traffic generator.

A powerful feature of traffic generation is *scenario replay*, the ability to capture a buffer of live network data and then use the protocol analyzer to retransmit the data in either its original form or as modified by the user. The user can define the frame, can cut and paste the frame from one already captured, or send a captured file. This

**TABLE 24.6 Traffic Generator Parameters.**

Parameter	Specification	Description
Load	Maintained Traffic Level	Specify a traffic level, in terms of percent utilization, to be maintained under all circumstances. The traffic generator will add a variable load to the existing network traffic to maintain the specified value.
	Loaded Traffic Level	Specify a traffic level, in terms of percent utilization or frames per second, to be transmitted on the network. This load will be transmitted regardless of the existing traffic level on the network.
	Loaded Statistical Distribution	Specify a traffic profile, in terms of a statistical distribution, to be transmitted on the network.
Interframe Spacing	Maximum Interframe Spacing	Specify the maximum amount of time, typically in microseconds, to be allowed before the traffic generator transmits the next frame.
	Minimum Interframe Spacing	Specify the minimum amount of time, typically in microseconds, to be allowed before the traffic generator transmits the next frame.
	Interframe Spacing	Specify the precise amount of time, typically in microseconds, before the traffic generator transmits the next frame.
Payload Specification	Single-frame	Transmit one specific frame.
	Frame Sequence	Transmit a specified sequence of frames.
	Capture Buffer	Transmit the contents (whole or a specified section) of the capture buffer.
	File	Transmit a file of data that was previously captured from the network under test.
Error Generation	Error Injection Rate	Specify the rate at which to inject errors into the payload being transmitted (typically specified in errors per second).
	Protocol Error	Specify the type of protocol errors to be injected.
	Transmission Error	Specify the type of transmission errors to be injected.
Cycle	Single	Transmit the specified payload once.
	Iterative	Transmit the specified payload the number of times specified by the user.
	Continuous	Transmit the specified payload continuously until the traffic generator is stopped by the user.

is used to capture live problems in the field and then to duplicate them in the engineering lab. Scenario replay is used for field service as well as R&D applications.

### 24.6.5 Bit error rate tests

Bit error rate (BER) tests are transmission tests used to determine the error rate of the transmission media or the end-to-end network. While advanced BER measurements reside in the domain of sophisticated transmission test sets, protocol analysis, particularly in a WAN or ATM environment, often requires a verification of the media. BER tests are performed by transmitting a known bit pattern on the network, looping it back at some point on the network, and analyzing the received sequence. The bit error rate is calculated as a percentage of the bits in error compared to the total number of bits received.

### 24.6.6 Stimulus/response testing

While many networking problems can be solved quickly with decodes and statistics, many of the more difficult problems cannot be solved in a nonintrusive manner. In this case it is necessary to communicate actively with the network devices in order to recreate the problem or obtain necessary information to further isolate a problem. The user can actively query or stress the network and observe the results with decodes and statistics.

Observing Ethernet frames with the same IP address and different MAC addresses might indicate a duplicate IP address problem, for example, but it also might just be the result of a complex router topology. To determine the node's identity—router or true duplicate IP address—the user can search through thousands of captured frames with decodes. But by sending an Address Resolution Protocol (ARP) frame to the node on the network and comparing the results to the addresses of the routers on the network, determined with Node Discovery statistics measurements, the user can quickly isolated duplicate IP addresses.

Typical active tests include:

- Connectivity tests
- Installation tests
- Network path analysis
- Adapter status commands
- Address resolution tests
- MIB (Management Information Base) queries

### 24.6.7 Simulation

In the context of protocol analysis, simulation can take two forms: *protocol simulation* and *protocol emulation*.

**Protocol simulation.** Protocol simulation allows the user to send strings of data containing selected protocol headers along with the encapsulated data. In this way, the operation of a network device can be simulated for the purpose of testing a suspected problem or for establishing a link to confirm operation.

**Protocol emulation.** Protocol emulators are software programs that control the operation of the protocol analyzer automatically. In the case of X.25, for exam-

ple, a protocol emulator takes over the functions of bringing up the link at layers one, two, and three and then automatically maintaining the connection. These functions are specified by the user, and the operation is then placed under programmatic control. Typically simulation tests are executed on top of an emulation program that is automatically handling the lower layers of the protocol stack, e.g., emulating a frame relay circuit while simulating a ping.

#### 24.6.8 Programming language

Protocol analyzers offer a wide range of capabilities in programming. Analyzers intended for field service applications often have simple, soft-key-assisted programming languages to control the instrument functions that monitor and simulate network traffic. Those intended for use in development have fully compiled or interpreted language implementations for custom test script generation. The commands in the programming language will give the user access to the full set of functions described in section 24.7. In addition to script generation, programming languages are useful for analyzing acquired data, searching for complicated combinations of events, decoding special protocols, and customizing the data format.

Programming languages usually are communications-enhanced versions of standard programming languages such as C or Visual Basic. They provide the ability to create, modify, debug, and execute user programs. Libraries are provided to implement common networking functions such as calculating an FCS or assembling and transmitting a frame. The programming language environment interfaces to the measurements and analysis capability of the rest of the product.

### 24.7 Protocol Analysis Options

Protocol analysis consists of using measurements like those described in the preceding section to isolate network problems or monitor network performance. The protocol analysis options described in this section are a set of orthogonal capabilities to the measurements described previously. For example, a protocol analyzer can gather statistics on the network traffic, but a more effective troubleshooting approach is created by setting a capture filter and running the statistics so only the data between a file server and a router is analyzed.

#### 24.7.1 Data capture

The most fundamental attribute of protocol analysis is the ability to capture the traffic from a live network and to store this data into a capture buffer. The captured data then can be analyzed and reanalyzed at the user's discretion. The capture buffer is either part of the main memory of a PC, or it is a special RAM buffer dedicated to capturing frames. Information stored in the buffer includes all of the frames on the network, a timestamp indicating when they were captured, and network-specific status information necessary to reproduce the exact state of the captured traffic. Since protocol analysis is used primarily for fault isolation, it is by definition used when the network is either experiencing heavy traffic loads or experiencing errors. It also is necessary to analyze sequences of frames, since

most protocols are state-dependent. It therefore is essential that all traffic on the network is captured regardless of the level of utilization, or whether the frames are errored or damaged.

**Data logging.** Many network troubleshooting sessions can be spread over hours and days. The capture buffer on a typical protocol analyzer is filled in seconds on a high-speed network. Data logging capabilities therefore are crucial for setting up long troubleshooting sessions. The user can specify a filename and a time interval for recording critical information. This information is then regularly written out to disk and can be examined later. Information that is typically stored to disk includes frames matching a user-specified filter, statistics results, or the results of a programmed test.

**Searching.** Once data is captured in the capture buffer, it can be repeatedly examined for problems or events of significance. Search criteria, like filter criteria, are based on address, error, protocol, and bit patterns.

**Marking.** Once a specific frame has been located in the capture buffer, it can be marked for future reference. The marked frame can be returned to with a “go to.” Timing and counting measurements also can be performed on the marked portions of the capture buffer.

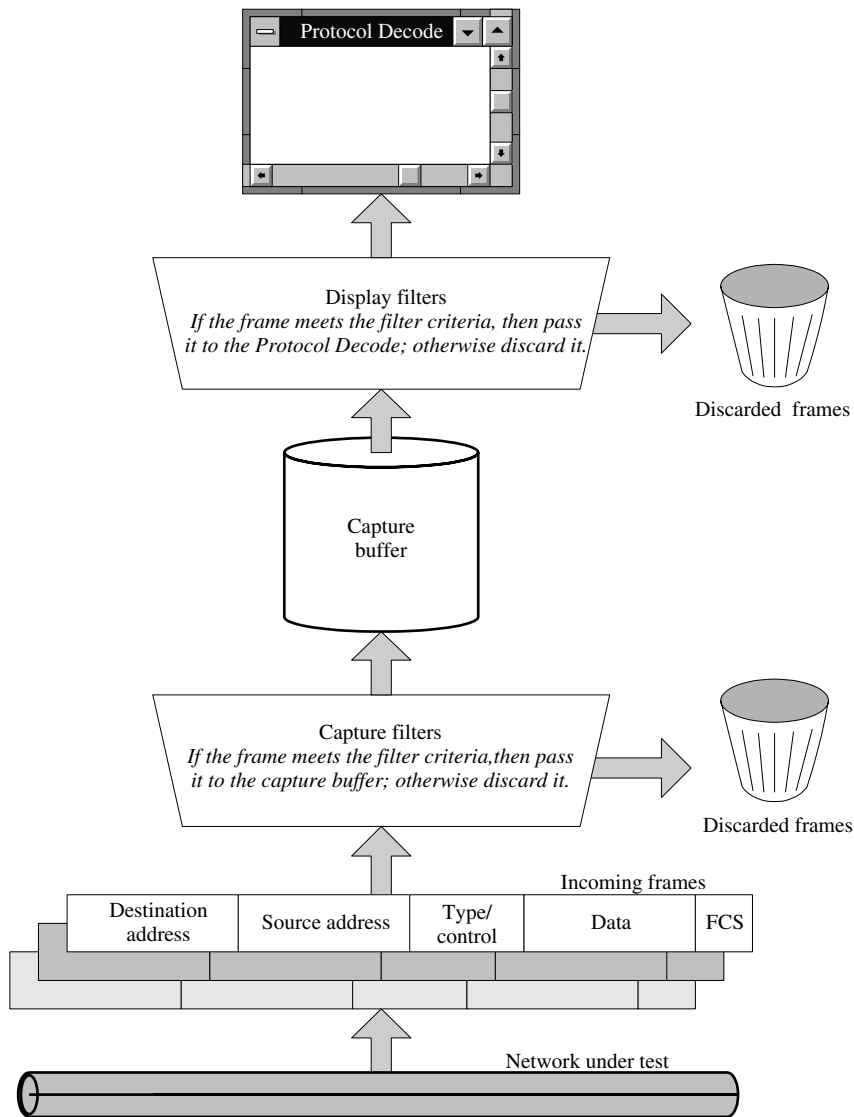
### 24.7.2 Filtering

The key to troubleshooting network problems successfully is based on eliminating the unnecessary information and focusing on the critical information that is essential to solving the problem. Computer networks process thousands of frames per second. A protocol analyzer can quickly fill a 32 MB capture buffer with frames, and a user can sift through protocol decodes searching for the errors. But this is a time-consuming and tedious task. The most powerful function of protocol analysis is the ability to filter the data on the network in order to isolate problems.

The function of a *filter* is very similar to that of a *trigger*; sometimes also called a *trap*. Specific filter patterns are set by the user of a protocol analyzer; these filters then are compared with the data from the network. Filters range from simple bit-pattern matching to sophisticated combinations of address and protocol characteristics. There are two fundamental types of filters, *capture filters* and *display filters* (Figure 24.9).

Capture filters are used to either include or exclude data from being stored in an analyzer's capture buffer. Capture filters make it possible to collect only the frames of interest by eliminating extraneous frames. This effectively increases the usage of the capture buffer. Rather than a 32 MB capture buffer that contains only 6 errored frames out of 40,000 captured frames, the data is filtered so that only errored frames are in the buffer. More frames of interest can be captured and they can be located more quickly. A disadvantage of capture filters is that it is necessary for the user to know what to filter on, that is, to have some idea of what problem to investigate. A second disadvantage is that the frames that were filtered out might contain a sequence of events leading up to the errored frame.

In many situations the source of a network problem is not known, so it is necessary to capture all the frames on the network and use display filters to repeatedly filter the frames. Because all of the frames are stored in the capture buffer, the frames can be played back through the display filters. Display filters act upon the frames



**Figure 24.9 Display and capture filters.** Capture filters compare incoming frames against the user’s filter criteria. Frames are either put in the capture buffer or discarded. Display filters operate on frames already in the capture buffer, displaying them selectively to the user.

once they have been captured. Frames can be selected for display by measurements such as protocol decodes.

Filter conditions can be combined to form more powerful filter criteria. As the troubleshooting process progresses, the user typically discovers more and more information about the network problem. As each new fact is discovered it can be added to the filter criteria, until finally the problem is identified. For example, to isolate a

TABLE 24.7 Filter Criteria.

Filter Class	Filter Type	Filter Criteria
Address	MAC Layer Source Address	Frames with the specific source address of the device sending the frame.
	MAC Layer Destination Address	Frames with the specific destination address of the device specified to receive the frame.
	Network Layer Source Address	Frames with the specific source network address of the system in an end-to-end network that originated the frame.
	Network Layer Destination Address	Frames with the specific destination network address of the system in an end-to-end network that is specified to receive the frame.
	Broadcast	Frames with the destination address equal to a broadcast address (frames to be sent to all devices on the network).
	Multicast	Frames with the destination address equal to a multicast address (frames to be sent to a group of devices on the network).
	LCN, Logical Channel Number	Frames containing a specific LCN. The LCN identifies one of many virtual circuits established in a packet-switched X.25 network.
	DLCI, Data Link Connection Identifier	Frames containing a specific DLCI. The DLCI identifies one of many virtual circuits in a packet-switched frame relay network.
	VPI/VCI, Virtual Path Identifier/Virtual Channel Identifier	Frames containing a specific VPI/VCI pair. The VPI/VCI pair is used to identify a specific virtual circuit in an ATM network.
	Error	FCS (Frame Check Sequence)
All Protocol Errors		Frames experiencing any protocol errors.
Specific Protocol Error		Frames containing a specific protocol error, e.g., an illegally long frame.
All Transmission Errors		Frames experiencing any transmission errors.
Specific Transmission Errors		Frames experiencing a specific transmission error, e.g., an alignment error.
Protocol	Invalid Frames	All invalid frames.
	Specific Protocol	Frames containing a specific protocol type, e.g., IP.
	Specific Protocol Stack	Frames containing the protocols specific to a particular protocol stack, e.g., TCP/IP.
	Specific Protocol Layer	Frames containing protocols specific to a particular layer, e.g., the Network layer.
	Specific Protocol Field	Frames containing a specific protocol field, e.g., a TCP Destination Port of a specified value.
Bit Pattern	Specific Frame Types	Frames of a specific frame type, e.g., an ICMP (Internet Control Message Protocol) Redirect frame.
	Position-dependent bit pattern	Frames containing a hexadecimal pattern at a fixed offset from the beginning of the frame.
	Position-independent bit pattern	Frames containing a hexadecimal pattern anywhere in the frame.

faulty Ethernet network interface card it is necessary to filter on the MAC address of the suspicious node and bad Frame Check Sequence (FCS) simultaneously.

Table 24.7 summarizes the different types of filter criteria. The same criteria usually can be used either by display filters or by capture filters.



### 24.7.3 Triggers and actions

In order to troubleshoot network problems it often is necessary to identify specific frames or fields in frames. Triggers are used to detect events of significance to the user and then initiate some action. Triggers and filters operate the same way in terms of recognizing conditions on the network. The parameters for setting trigger criteria are the same as the filter types specified in Table 24.7. The trigger is a key capability of protocol analyzers, since it allows an automatic search of a data stream for a significant event, resulting in some action to be taken. Possible trigger actions include:

- Visual alarm on the screen
- Audible alarm
- Start capturing data in the capture buffer continuously
- Start capturing data, fill the capture buffer, and stop
- Position the trigger in the capture buffer and stop capturing data
- End the data capture
- Increment a counter
- Start a timer
- Stop a timer
- Make an entry in the event log
- Start a specific measurement
- Send an SNMP trap
- Log data to disk

### 24.7.4 Timers and counters

Timers measure the time interval between two events, usually delineated by triggers. Many protocols are based on responses occurring within a certain elapsed time after a command or request. The protocol implementations keep timers to implement the protocol properly. Protocol analysis tools therefore offer timer services in order to verify the proper operation of the protocol and the network, and can be used to measure the performance of events occurring on the network. Timers also can be used to make latency and performance measurements on network devices.

Counters are used to tally specific events (recognized by triggers) on the network. They can be used to count error conditions or specific sequences of frames. Users can create their own statistical measurements with counters.

### 24.7.5 Event log

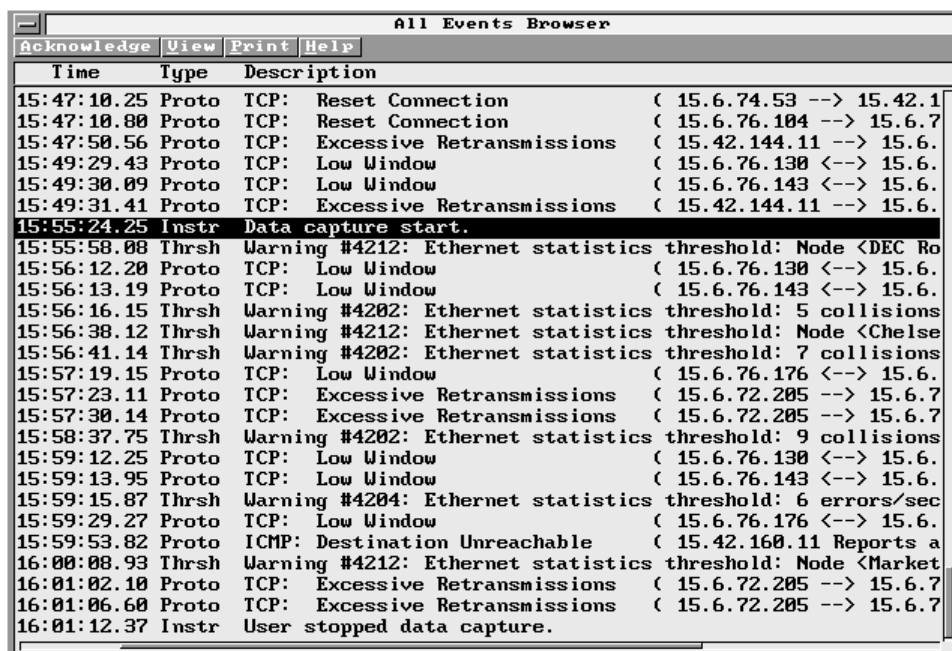
Most tools that include protocol analysis capabilities have event logs to record significant events that have occurred on the network under test. Event logs usually are built as a simple database that records the time of the event, the severity of the event, and a high-level description of what occurred.

Event log entries are categorized by severity, typically normal, warning, and alert levels. Normal events will not impact the network's performance but usually are set up by the user for performance in the monitoring process, e.g., tracking user logins, file transfers, etc. Warning-level events indicate a potential network problem—one that should be investigated by the user—typically including certain statistical thresholds being exceeded, nonresponding routers, and protocol events such as broadcasts or collisions exceeding normal levels.

Alert-level events indicate serious problems that are severely impacting the network. These must be investigated immediately to ensure continued network operation. Examples of alert-level events include high levels of errors or collisions, or high network utilization. Levels used to trigger warning- and alert-level events are user-specified. Typically the levels are selected by performing a network baseline: The performance of the particular network is measured over time and used to determine characteristic normative values.

Figure 24.10 shows the display for a typical event log. Each entry is timestamped when it occurred. A type field indicates if it is a protocol error, a user threshold that was exceeded, or an instrument event such as the start or stop of a data capture. A description field provides a quick summary of the type of event and detailed parameters.

The event log can be used for long-term monitoring of a network without requiring a user to be present, and can be examined periodically.



All Events Browser		
Acknowledge View Print Help		
Time	Type	Description
15:47:10.25	Proto	TCP: Reset Connection ( 15.6.74.53 --> 15.42.1
15:47:10.80	Proto	TCP: Reset Connection ( 15.6.76.104 --> 15.6.7
15:47:50.56	Proto	TCP: Excessive Retransmissions ( 15.42.144.11 --> 15.6.
15:49:29.43	Proto	TCP: Low Window ( 15.6.76.130 <--> 15.6.
15:49:30.09	Proto	TCP: Low Window ( 15.6.76.143 <--> 15.6.
15:49:31.41	Proto	TCP: Excessive Retransmissions ( 15.42.144.11 --> 15.6.
15:55:24.25	Instr	Data capture start.
15:55:58.08	Thrsh	Warning #4212: Ethernet statistics threshold: Node <DEC Ro
15:56:12.20	Proto	TCP: Low Window ( 15.6.76.130 <--> 15.6.
15:56:13.19	Proto	TCP: Low Window ( 15.6.76.143 <--> 15.6.
15:56:16.15	Thrsh	Warning #4202: Ethernet statistics threshold: 5 collisions
15:56:38.12	Thrsh	Warning #4212: Ethernet statistics threshold: Node <Chelse
15:56:41.14	Thrsh	Warning #4202: Ethernet statistics threshold: 7 collisions
15:57:19.15	Proto	TCP: Low Window ( 15.6.76.176 <--> 15.6.
15:57:23.11	Proto	TCP: Excessive Retransmissions ( 15.6.72.205 --> 15.6.7
15:57:30.14	Proto	TCP: Excessive Retransmissions ( 15.6.72.205 --> 15.6.7
15:58:37.75	Thrsh	Warning #4202: Ethernet statistics threshold: 9 collisions
15:59:12.25	Proto	TCP: Low Window ( 15.6.76.130 <--> 15.6.
15:59:13.95	Proto	TCP: Low Window ( 15.6.76.143 <--> 15.6.
15:59:15.87	Thrsh	Warning #4204: Ethernet statistics threshold: 6 errors/sec
15:59:29.27	Proto	TCP: Low Window ( 15.6.76.176 <--> 15.6.
15:59:53.82	Proto	ICMP: Destination Unreachable ( 15.42.160.11 Reports a
16:00:08.93	Thrsh	Warning #4212: Ethernet statistics threshold: Node <Market
16:01:02.10	Proto	TCP: Excessive Retransmissions ( 15.6.72.205 --> 15.6.7
16:01:06.60	Proto	TCP: Excessive Retransmissions ( 15.6.72.205 --> 15.6.7
16:01:12.37	Instr	User stopped data capture.

Figure 24.10 Event log.

### 24.7.6 Remote control

In many networks a problem could occur at a physical location to which the network expert does not have access for connecting a protocol analyzer. The key troubleshooter might be in a different building or a different city. Often one network engineer supports many different networks that are spread over a wide area, so the ability to control a protocol analyzer remotely may become crucial to solving problems. Most protocol analyzers provide such a capability for field service and network management applications. Remote control is implemented by using a PC or a Unix workstation to control the protocol analyzer via a telemetry network implemented with a modem or a LAN connection. Remote control typically includes three levels of control:

- *Virtual terminal remote*, in which a PC or Unix workstation executes an application that displays the screen of the remote protocol analyzer, and provides full keyboard and mouse control. The remote protocol analyzer operates as if it were executing on the local PC or workstation.
- The ability to upload or download captured data files or preconfigured instrument setups.
- The remote control ability to send commands that start and stop tests on the remote protocol analyzer.

Another use for remote control is making a distributed measurement on the entire network. For this application, measurement probes are placed at strategic points throughout a network to monitor activity and report the results to a central control console. With information from many parts of the network, this central console analyzes and displays the overall performance and problems of the entire network.



---

Part

**7**

# Network Test Instrumentation



---

Chapter  
**25**

# Analog Measurement Instrumentation

**Ronald D. Lowe**

*Hewlett Packard Company, Colorado Springs, Colo.*

## 25.1 Introduction

The purpose of a chapter on analog instruments and measurements is to describe what is required to make analog measurements on digital networks. Instruments such as cable testers and Transmission Impairment Measurement Sets (TIMS) make general analog measurements, i.e., voltage, resistance, and frequency, but display results in terms that indicate the health or acceptability of specific network systems. Only those network-specific measurements will be dealt with here. Readers with an interest in how fundamental measurements are made should consult the reference [1] on general test equipment.

Analog parameters exist only at the Physical layer, level 1 (sometimes incorrectly referred to as “level 0”) of the OSI Reference Model for networks. Analog instruments have evolved into three classes: low-cost, general-purpose, and high-performance. The separation of the classes becomes blurred, however, by continuing cost reduction and the appearance of more sophisticated technology in smaller, easier-to-use configurations. Further measurement categorization will be by transport media (copper or fiber optics) and then by LAN vs. WAN within the specified media.

### 25.1.1 Transport media

Transport media have a significant impact on the type of technology required in the instrument. Copper-wire media have LAN and WAN subsets of instruments. Likewise, fiber optic media have similar subsets; LANs using fiber typically use multi-mode, while WANs are made of predominantly single-mode fiber. Test equipment for wireless media tends to be in a special class; see Chapters 17 to 20 of this book for more information. Table 25.1 is a sample list of network media.

TABLE 25.1 Typical Network Transport Media.

Network Type	Defining Standard	Media Type	Transmission Mode	Impedance/Wavelength
10Base2	IEEE 802.3	Coax	Unbalanced	50Ω
10Base5	IEEE 802.3	Coax	Unbalanced	50Ω
10Base-T	IEEE 802.3	UTP	Balanced	100Ω
100Base-ANYLAN	IEEE 802.12	UTP	Balanced	100Ω
Fast Ethernet	IEEE 802.3u	UTP	Balanced	100Ω
FDDI TP-PMD[5]	ANSI X3T9.5	UTP	Balanced	100Ω
WAN voice	IEEE 743-84,96	UTP	Balanced	600, 900, 1200Ω
WAN data	Various	UTP	Balanced	135, 150Ω
Token-Ring	IEEE 802.5	STP	Balanced	150Ω
10Base-F	IEEE 802.3	Fiber optic	Multimode	820 nm (nanometers)
Token-Ring	IEEE 802.5	Fiber optic	Multimode	820 nm (nanometers)
FDDI[5]	ANSI X3.237	Fiber optic	Multimode	1300 nm
WAN data	Various	Fiber optic	Single-mode	1300, 1500 nm
FDDI[5]	ANSI X3.184	Fiber optic	Single-mode	1300 nm

UTP = unshielded twisted pair  
STP = shielded twisted pair

### 25.1.2 LAN or WAN instruments

The distinction between LAN and WAN instruments has its origin in how the instrumentation evolved to support the needs of the service providers and network operators. Other network types, such as MAN (metropolitan area networks) and CAN (campus area networks), are arbitrary designations and have not resulted in additional categories of instrument types or measurements.

Throughout this chapter, references are made to the various standards that exist to define the networks. These references will provide the primary source of the parameters and specifications that are important to both the instrument manufacturer and the network manager. The reader also should be aware that standards are in a constant state of revision. Always request the latest revisions from the responsible publisher.

### 25.1.3 Classes of instruments

The marketplace historically has created demand that has resulted in three classes of analog instruments: limited-function, general-purpose, and high-performance. These classes are becoming arbitrary as manufacturers have continued to decrease costs and increase functionality.

**Limited-function instruments.** The limited-function end of the product spectrum tends to have these characteristics: very portable, handheld, battery-operated,



easy to use, limited or single functionality, go/no-go results, and not necessarily meeting measurement accuracy required by network standards. The population of product in users' hands could equal the number of employees performing the testing. An example is the multipurpose, handheld cable test sets for local area network cable plants.

**General-purpose instruments.** General-purpose analog instruments are referred to in the context of our network testing agenda and should not be confused with general-purpose test equipment. These are larger than low-cost instruments, but are still portable, lightweight, and manageable with one handle. Measurement accuracy meets standards, with some margin. They usually will test only a subset of the measurements specified by the standard. The measurement set is the minimum determined by the users. The population of this product usually is one per work team or department. Cost may be as much as double the handheld class. An example of this class is the portable TIMS instrument used by the operating telephone companies.

**High-performance instruments.** These highly accurate instruments measure all the parameters specified in the standard. They have an internal clock accuracy comparable to secondary time standards. Functionality of measurement is more important than ease of use. The highest-cost instrument usually is more suited for the network engineer as primary operator. They are found most often in engineering laboratories and manufacturers' applications and might not be available in field-portable configurations. They make network-specific measurements that differentiate them from high-end general-purpose test equipment. One example, found at a cable manufacturing facility, is a computer-based automatic cable test system for certifying Category 5 LAN cable.

**Class overlap.** There is a lot of overlap among these classes, both in functionality and marketplace product presence. Handheld units will be found in laboratories. High-performance products will be used by skilled technicians in the field, such as the full-function OTDR (optical time domain reflectometer) used on fiber optic cable. Manufacturers are providing more features constantly, lowering the cost and simplifying the user interface on their products. This has been illustrated most dramatically in the increased capability of the LAN handheld cable testers, which have evolved from simple cable continuity testers to full-function tools based on network wiring certification standards.

## 25.2 LAN Testing on Copper-Wire Media

Cable testers are a class of instruments used for testing copper-wire media on local area networks.

The appropriate standard is the EIA/TIA 568A, TSB-67 Level II, which covers the specification of the copper wire installation.<sup>[2]</sup> This can be coaxial cable or twisted-pair, both shielded and unshielded. IEEE publishes standards for the specifications of the LAN network that will be installed on the media, such as IEEE 802.3 for Ethernet/10Base-T, 802.5 for Token-Ring, etc.<sup>[3]</sup>

**570 Network Test Instrumentation****25.2.1 Automatic test limitations**

Limitations of the autotest feature commonly available in handheld instruments needs special mention. *Autotest* is a generic term for the feature that performs the entire menu of function tests with a single instruction. Manufacturers have had to make a tradeoff between conflicting requirements: time to make a measurement, which includes settling time of the measurement circuits, and the total length of time that the operator must wait to obtain the results. If the test set cycles through the measurements too quickly, some of the tests such as TDR (time domain reflectometry) and NEXT (near end crosstalk) might not have settled enough to give accurate, stable readings. If the test set cycles too slowly, then the operator is required to wait longer for the results.

The “faster” test sets are perceived as being better than those from competitors. Most manufacturers use modern signal processing techniques to help solve this measurement dilemma, but one decision has been made of which the user should be aware: the assumption that design and installation adheres to published standards for the network under test. This is why some users get conflicting results between manual tests and autotest results from the same instrument.

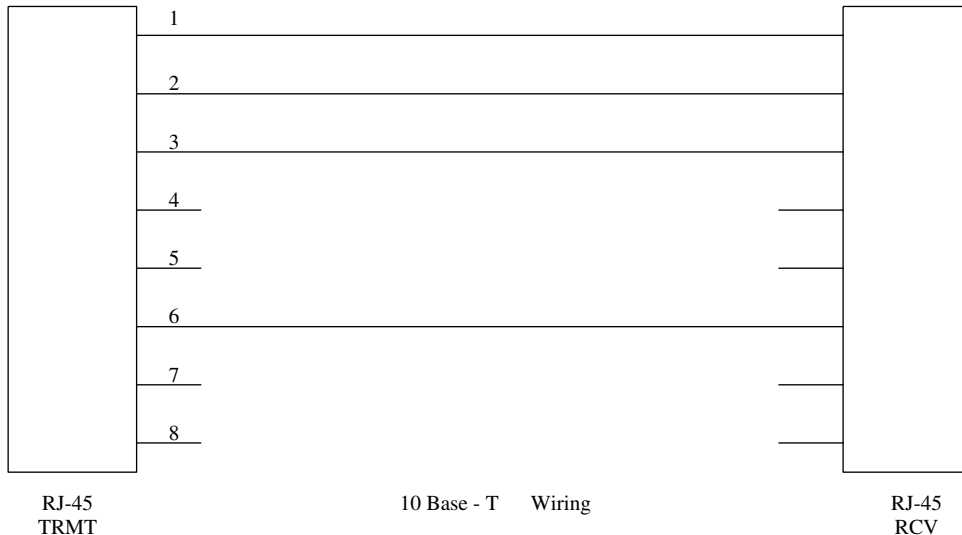
For example, 10Base-T Category 3 cable should not exceed 100 meters in length between nodes. The timing on the measurements in an autotest mode has assumed that fact and has some margin built in. If the operator attempts to use the autotest function on a 1000-meter reel of cable, the settling time for some measurements would be inadequate to get accurate results. Some manufacturers have anticipated this usage and either provide special tests or recommend that manual function testing be performed.

**Cable terminations.** An accessory that is required to support the test set is some form of a terminating unit for the far end of the cable. Analog measurements on LAN cable are intrusive and require the user to take that segment of the network out of service. How this terminating unit is used varies with the manufacturer; its requirements vary with the measurement being made. Its purpose will be mentioned as each measurement is discussed.

**25.2.2 Resistance and wire mapping**

Resistance and wire mapping measurements typically use the same ohms measurement principles. The resistance measurement returns a value, whereas the wire map function detects continuity up to a chosen limit (such as 200 $\Omega$ ) and displays some form of “connection” or “no connection” indication.

This ohmic limit might or might not be explicitly stated by the manufacturer. It must be a resistance value that is greater than the maximum dc resistance specified by the standard for the length and wire gauge of the cable in the specific network segment. The manufacturers have coded the wire resistance based on gauge and length into the instrument’s memory. The user need only select the wiring type to make the measurements. The termination unit may be as simple as looping back the connection to the test set, or it might actually support the measurement for a true end-to-end resistance reading.



**Figure 25.1 Typical display of 10Base-T wire map.** This is one of the several ways to display cable wiring. For 10Base-T cabling the active connector pins are 1, 2, 3, and 6. Some instruments will display only a table of connections with a pass/fail assessment of the connections.

Some manufacturers of test equipment for these applications have integrated the appropriate standards into their test tools. Cable wiring connections are an example of this integration. Historically cable wiring was verified by a tedious manual testing process using an ohmmeter. This verification now has been integrated into a single automated measurement with pass/fail cable integrity results.

For example, the EIA/TIA 568A will define the pairing of a four-pair twisted copper termination at an **RJ-48** connector to be pair 1 at pins 4 and 5, pair 2 at pins 1 and 2, pair 3 at pins 3 and 6, and pair 4 at pins 7 and 8. The IEEE 802.3 specifies that for 10Base-T networks the transmit and receive connection will be made on pairs 2 and 3. A tester’s wire map function must verify that the pairings in the cable are correct, and that for 10Base-T the appropriate pairs have been connected. The tester may provide a go/no-go green or red light indication, or display graphically the actual wiring measured. Figure 25.1 depicts a typical graphical result from a wire map measurement.

**25.2.3 Noise**

Meaningful noise measurements are made by detecting signal energy with a true RMS detector through a specified filter characteristic. LAN standards have not specified noise limits in this classic sense. Some form of noise measurement still can provide useful information if the filter characteristics, e.g., low-pass vs. high-pass, are known. Some manufacturers have provided this measurement using simple filters and average detection of the signal. There are no pass/fail criteria, but the existence of high levels of noise becomes an important datum for understanding a marginal or failed communication link.

### 25.2.4 Time domain reflectometry (TDR)

TDR actually is an analog measurement in the time domain, one that measures length or distance to an imperfection in a cable. It is included here because it is a primary measurement function in almost all network cable testers. There are some pitfalls that stem from its improper usage.


TDR is based on a pulse propagating down a transmission line of a specified characteristic impedance. If the pulse encounters an abrupt change in the cable's characteristic impedance, a reflection of the pulse is created and propagated back to the source. The time between the launching of the original pulse and the receipt of the reflected pulse is measured. If the velocity of propagation  $V_p$  of the cable (as specified by the manufacturer) is known, then the distance from the source and the point where the pulse was reflected can be computed in feet or meters.

The accuracy of the measurement is a function of the test set's ability to detect the return pulse and the accuracy of the cable  $V_p$  figure used in the calculation.  $V_p$  usually is given as a decimal factor representing velocity of a propagating signal compared to the speed of light. For instance, a  $V_p$  of 0.67 means that a signal will propagate on the cable at 67 percent of the speed of light. Some test equipment manufacturers also will provide calibration by allowing the user to select a cable type from a list in the test set's memory.

**Sources of error.** There are three principal sources of error in making TDR cable length measurements: choosing the correct  $V_p$ , determining the cable twist rate, and measuring the time of the reflected pulse. Nominal  $V_p$  is furnished by the cable manufacturer, but once it has been installed the identity of the cable might not be obvious to the user. One way to improve accuracy is to measure a piece of cable physically and then set the  $V_p$  for the tester to read the correct length. Then that  $V_p$  setting will be accurate for measurements on the same type of cable.

Another source of error on Category 5 twisted-pair cable is differing twist rates within the cable. Careful inspection of the bare cable pairs will reveal that the number of turns or twists per foot is different for each of the four pairs in the cable. One should surmise that the pair with the most twists per foot will be measured by the TDR as longer than a pair with fewer twists per foot—because it is, in fact, physically longer.

An error source that the equipment manufacturer can control is the quality of the transmitted pulse shape and how the return pulse is processed. Figure 25.2 shows where error can be introduced by the test equipment's decision-making as to the timing of the return pulse. Some equipment will take several repeated TDR measurements and then average the readings to improve accuracy. A far-end termination unit should not be required for the TDR measurement, since a far-end open circuit provides a perfectly good reflection for the measurement. On some test sets a different length measurement is displayed depending upon whether the termination unit is connected and, if so, whether it is powered on or off. When in doubt, follow the manufacturer's recommendation.




**Figure 25.2 TDR source of error.** This figure represents a transmitted pulse being attenuated as it traverses a long cable. Since the attenuation is higher at higher frequencies, the pulse shape also is degraded. The greater the degradation, the greater the error shown as  $\Delta T$ . Note that this type of error always makes the cable measure longer than its actual length.


**25.2.5 Attenuation**

*Attenuation* is the loss of signal at a specified frequency measured in dB from the source to the end of a terminated cable. The measurement can be made either in the loopback or the end-to-end mode, as shown in Figure 25.3. A loopback measurement simplifies the equipment but can give errors if the cable has poor crosstalk characteristics or isolation between the pairs. (The mathematical roots of this error condition are beyond the scope of this chapter.)

For a qualitative assessment of this error, follow this line of thought: Typical cable lengths are several wavelengths long at the higher frequencies, meaning that the phase of an interfering crosstalk signal either adds to or subtracts from the desired



Attenuation via loop back method  
 $\text{Attn (one way)} = \text{Measured Attn} / 2$   
 (invalid if isolation between pairs is poor)



Attenuation via end to end method

**Figure 25.3 Attenuation measurement.** Two methods of attenuation measurement are shown. The loopback method is the least expensive to implement but is subject to a serious unresolvable error. If there is significant crosstalk between the pairs, then standing waves from the generating signal can result, giving an unpredictable measurement. This is a serious problem at data transmission rates above 5 Mbps. The end-to-end measurement requires some means of calibrating the termination unit to the test set's transmitted level.

signal, and depends both on the frequency and the length of the cable being measured. Therefore the measured signal level will have an error that depends upon the frequency, length of cable, and the crosstalk present. If there also is a mismatch in the cable termination that causes a standing wave phenomenon, then trying to achieve any measurement accuracy becomes an intractable problem for the test set.

The end-to-end measurement avoids all these problems except the termination mismatch. End-to-end measuring equipment adds significantly to the cost of the test set because a precision frequency source must be provided in the termination unit. The manufacturer can control the precision of the termination, but the presumption is that the cable has a known characteristic impedance. Manufacturers' specifica-

tions for characteristic impedance are somewhat loose for making precision measurements in twisted-pair cable. This is one of the serious design challenges that test equipment manufacturers have had to face in providing accurate measurements in the 20–100 MHz range.

It is worth noting that even with these inherent sources of error, these measurements, when used in conjunction with other measurements, give the user useful information on the condition of the cable.

### 25.2.6 Near-end crosstalk (NEXT)

It is important to discuss crosstalk in terms of NEXT, the standard of measurement in LANs, to distinguish it from *far-end crosstalk*, or FEXT. (FEXT will not be covered in this chapter.) Figure 25.4 shows the measurement connections for NEXT.



NEXT is a measure of the quality of isolation of the transmitted signal from the received signal. It should be obvious that poor isolation in a cable will cause the transmitted signal to bleed over into the received signal. The receiver will have difficulty in accurately determining the information content of the received signal. A digital test of this condition would disclose a poor or unacceptable BERT (bit error rate test).

It is important that the far-end termination unit provide a good match to the cable characteristic impedance to minimize any standing wave phenomena. It is also important that the instrument end of the product be designed for good isolation between its transmitter and receiver. This is embodied in the instrument's specification of the best NEXT it can measure accurately.

Unlike attenuation, a plot of NEXT versus frequency is not monotonic (Figure 25.4). This is caused by the interfering signal being in or out of phase with the reference signal. The nulls are very sharp and the peaks are broad. Their position in frequency depends upon the amount of isolation, the cable length, and the accuracy of termination; predicting them is a difficult mathematical problem. It is important that if the transmitter frequency is incremented discretely instead of swept continuously, then the increment must be small enough to identify the peaks of the NEXT. The results are displayed in dB of isolation from the transmitted signal at a specified frequency. Graphical displays are common, with limits specified by the standard incorporated in the display.

Some test equipment will display a computed noise margin as a result of an autotest measurement. The number is a calculation of the difference between the attenuation and the NEXT at a particular frequency. It gives the user an assessment of the quality of signal that can be expected on that specific portion of the network. The measurement is referred to as the *signal-to-crosstalk ratio*.

It is worth noting that NEXT measurement is the primary means of detecting a double-split wiring error. Figure 25.5 depicts a double-split condition. A split is where a wiring error has corrupted the pair integrity of a cable. Instead of correcting the error, the cable was resplit to give end-to-end continuity. End-to-end tests will yield correct wire map, resistance, length, and some degradation of attenuation. A NEXT measurement will show unacceptable results and thus is the primary measurement that will detect this problem.



**Figure 25.4 NEXT measurement.** The near-end crosstalk measurement requires the far end of the cable to be properly terminated. The procedure consists of transmitting on one pair and then measuring any signal level that is coupled to the adjacent pair of wires. This usually is done over a frequency range of interest. The result is a non-monotonic plot of crosstalk in dB. This non-monotonic condition is caused by the reinforcing and interfering effects of the coupled signal's phase. The limit line is specified by the appropriate LAN standard. It is important that if the test frequencies are stepped rather than continuously swept, the step interval must be small enough to detect the peaks in the crosstalk.

**25.2.7 Cable location**

What appears to be obvious and logical on a network diagram can become incomprehensible in the wiring closet. The network manager must be able to trace and lo-



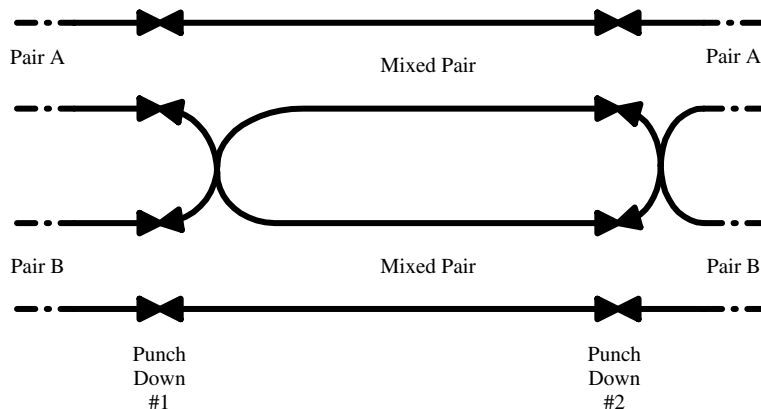
cate the specific cable of interest through all the other cabling present in the facility. Tracing typically involves placing an audible tone on the cable and, using a separate receiving device, detecting and following the cable through the wiring closet and walls. If the cable is routed through a metal conduit, the signal may be lost; applying it directly to the conduit allows the trace to continue.

If the test equipment provides a *locator function*, wiring can be identified from unique terminating units at the far ends of the cables. Numbered passive terminating units can be plugged into different office outlets. By measuring the terminating unit and displaying its unique number, the test equipment can determine where it is connected on the wire frame. These testing techniques can be used together to document an unknown wiring plant.


### 25.3 WAN Testing on Copper-Wire Media

It is appropriate to mention both the similarities and differences between LAN and WAN technologies. WAN has its origins in the basic telephone system, which traditionally has used twisted-pair copper wire at voice-band frequencies, less than 4 kHz. Cable lengths are considered in terms of *kilofeet*, and there are embedded repeaters. With respect to capability and ruggedness, instruments for making continuity and TDR measurements fall into a different class than those for attenuation and crosstalk. One usually does not find all the measurement functions combined into one instrument. Since the principles for making the continuity and TDR measurements are similar to the LAN instruments, they will not be repeated here.

What is unique is that a WAN usually reaches the individual subscriber as an analog channel, whereas a LAN channel is always digital. A type of test instrument called a Transmission Impairment Measuring Set (TIMS) is designed specifically for



**Figure 25.5 Double split pair.** This diagram depicts the double-split condition. A single split is where the pair integrity of the cable is corrupted by interchanging a wire in one pair with a wire in an adjacent pair. This fault would show up on a wire map as an obvious failure. If an attempt to correct this problem is done by creating another single split elsewhere in the cable installation to satisfy the wire map, however, then the double-split condition occurs. It satisfies the wire map test and might be only marginal on the attenuation test. Only a crosstalk measurement will disclose this fault condition.



**Figure 25.6** This is a typical example of a TIMS instrument that measures a subset of the measurements. Such an instrument is usually the instrument of choice for a telephone craft person that has the task of troubleshooting performance problems on analog telephone cabling. A scan of the figure will reveal several features of this type of instrument. There are five different impedance levels for both transmit and receive functions. The noise measurements can be weighted with five different filter characteristics. Timer control is provided for the Impulse Noise measurement. A selection of fixed frequency tones are provided including the ability to change the preset frequency. These test sets are a common measuring tool for the maintenance personnel of the telephone companies.

measuring the transmission characteristics of analog voice frequency circuits (Figure 25.6).

These instruments meet all or part of the requirements defined in IEEE Standard 743-1984 [4] for the North American telephone systems, and ITU-CCITT Vol. 4 Series O [6] for the European telephone systems. (Other regions of the world typically have adopted one of these two dominant standards.) IEEE 743 is being revised at this writing (1996); some measurements will be deleted and others added to support the newer digital subscriber loop requirements. Completion of this new standard will motivate instrument suppliers to make the transition from analog to digital measurement techniques. A digital implementation still will yield analog answers, so a discussion of analog measurements is still a worthwhile venture.

### 25.3.1 TIMS

A general description of measurements is easier if one places them into four categories: basic, transient, incidental modulation, and “other.” This categorization is consistent with that described in the IEEE standard 743-1984. Many products are on the market, with the smaller, less-expensive units implementing only a subset of the standard. Users must determine whether a subset of the measurements will satisfy their needs. The measurements examined here are:

- Basic measurements
  - Level and frequency
  - Noise and signal-to-noise
  - Envelope delay distortion (EDD)

- Transients
  - Impulse noise
  - Phase and gain hits
  - Dropouts
- Incidental modulation
  - Phase jitter
  - Amplitude jitter
- Other
  - Return loss
  - P/AR
  - Intermodulation distortion

### 25.3.2 Instrument parameters common to TMS

TMS instruments contain both transmitter and receiver functions, allowing tests in both loopback and end-to-end modes to be performed. All measurements on twisted-pair cable are made in balanced mode, with signals existing wire-to-wire and isolated from ground or earth potential. They also provide calibrated measurements for a variety of impedances; common terminations are 600, 900, and 1200 $\Omega$  for voice service, and 135 and 150 $\Omega$  for cable in program or data service. The instruments also provide a bridging (high-impedance) termination to facilitate connections to circuits that already are terminated properly (Figure 25.7). Usually there is provision for attaching a telephone handset to allow dialing and holding of the dialed line.

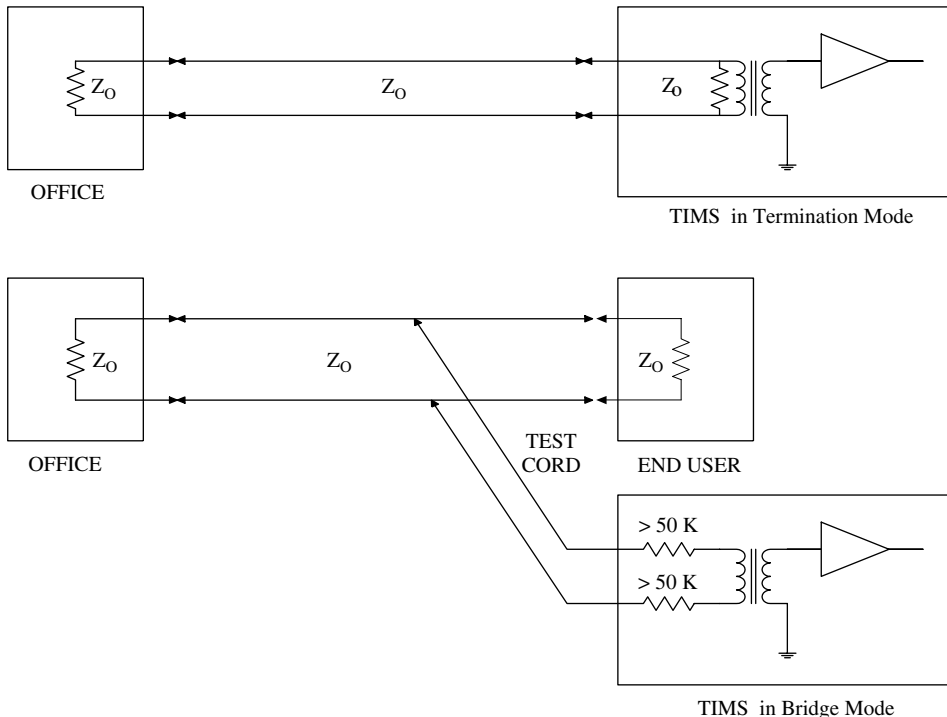
All level readings are in dB and referenced to 1 mW of sinusoidal power into 600 $\Omega$  impedance. The received level is expressed in dBm, equalling  $10 \log_{10}(10^3P)$ , or 0 dBm for 1 mW into 600 ohms. For example, a signal that is 10 dB below the 0 dBm reference would be displayed as –10 dBm. Calibration is provided for level in all possible terminations.

Almost all voice channels in the telephone system are digitized for long-distance transmission using an 8 kHz sampling rate. The standard 1000 Hz test tone was offset to 1004 Hz to avoid being a submultiple of the sampling rate. Other fixed tones are worthy of note. Provision is made to avoid transmitting 2600 Hz (2280 Hz for ITU) if the operator wishes; these tones are a disconnect signal in some telephone single-frequency signaling systems. Fixed tones of 404, 1004, and 2800 Hz allow a quick reading on the frequency response (gain slope) of the circuit under test. (The ITU tones are 300 and 820 Hz, or 1020, 2000, and 3000 Hz.)

### 25.3.3 Basic measurements

As noted in the bullet lists in section 25.3.1, the “basic measurements” discussed here are level and frequency, noise, signal-to-noise, and envelope delay distortion (EDD).

**Level and frequency.** For level measurements the transmitter is a sine wave generator that has a frequency adjustment range from a minimum of 50 Hz to at least 3900 Hz. Most transmitters operate down to 20 Hz and up to 110 kHz. Frequency stability



**Figure 25.7 TIMS termination and bridging.** A TIMS instrument can connect to a telephone cable by terminating the cable for intrusive testing, or by bridging onto an active circuit for monitoring purposes.

and accuracy requirements dictate a crystal-controlled oscillator; most implementations satisfy the requirement by using a frequency synthesizer design.

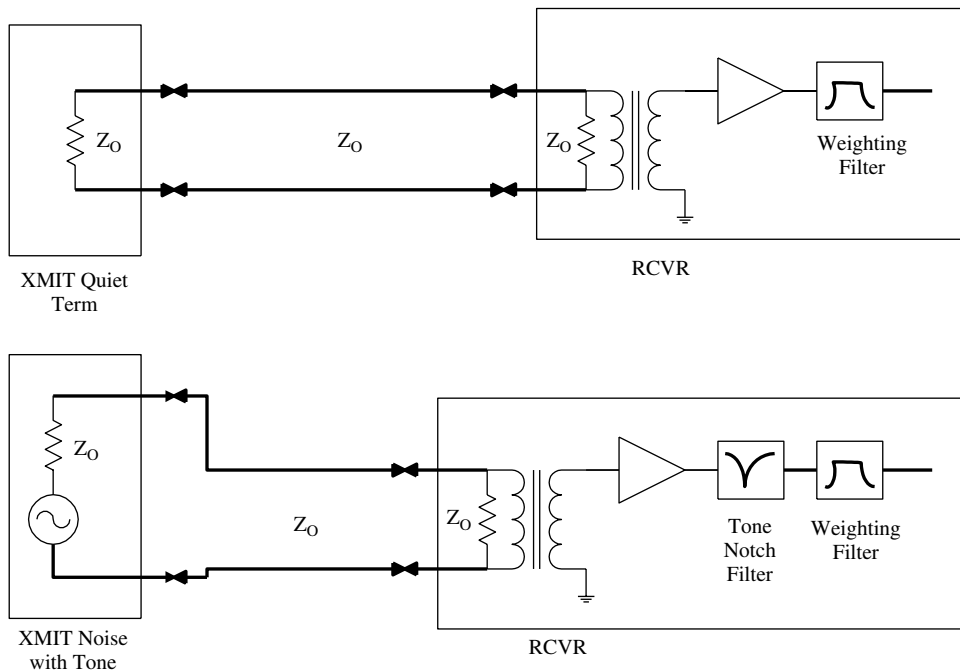
The output level should be adjustable from a maximum of +10 dBm down to a minimum of -40 dBm. The manufacturers tend to exceed these limits as the users demand more performance. Recent interest in higher frequencies (to 200 kHz) has been caused by WAN suppliers needing to convert voice-grade twisted-pair cable to carry digital traffic. The digital system most widely deployed to the individual subscriber is the Integrated Services Digital Network (ISDN). This system has a clock frequency at 192 kHz, with its dominant spectral energy at 96 kHz.

The receiver is simply a voltmeter, which has a bandwidth consistent with the transmitter frequency range, and an average detector. To achieve the dynamic range, some form of autoranging usually is included. The standard requires a sensitivity of at least -50 dBm, but because the same circuitry is used for noise and other measurements, most manufacturers provide a design with -90 dBm sensitivity. The receiver display provides both level in dBm and frequency in Hz or kHz, depending on frequency range. A useful feature is a provision for relative dB level readings. Selecting a Level Zero function sets the level indicator to 0 dB and stores the absolute level internally; level adjustments during that measurement sequence are then displayed in dB readings relative to the stored absolute value.

**Noise.** There are several variations on noise measurements that can be made by TIMS. A noise measurement presumes some sort of conditioning or weighting filter, of which there are several to be considered. To cover the unbalanced cable problem, there also might be an unbalanced or noise-to-ground measurement. Noise measurements require an RMS detector; a QRMS (quasi-RMS) has been developed as a reasonable compromise. The s/n ratio, a computed number based on a sequential level and noise measurement algorithm, is covered in a subsequent section.

The transmitter for noise measurements operates in only two modes. One is an off condition, referred to as a *quiet termination*. It provides the chosen transmitter termination impedance to the line. The second mode generates a fixed tone of 1004 Hz (820 or 1020 Hz for ITU) with a user-adjustable level. If there is a compander (compressor/expander for signal level control) in the circuit, it detects the tone and fixes its gain. Without this tone, the compander would try to compensate for the lack of a signal and wreak havoc with any attempts to measure noise levels. Figure 25.8 depicts the general noise measurement conditions.

The receiver uses much of the same measuring circuitry found in the level measurement. Table 25.2 lists first the weighting filters, followed by brief notes on their purposes.



**Figure 25.8 Noise measurement with and without tone.** Noise measurements always require the selection of a weighting filter. In dedicated telephone lines a tone is not required since the circuit is permanent. On telephone lines used in dial-up service require the application of a holding tone to maintain the circuit. This tone must be filtered out by the TIMS before meaningful noise measurements can be made.

TABLE 25.2 TIMS Noise Filters.

Filter Name	TIMS Version	Description and Purpose
C-message	North American	700–3000 Hz voice characteristics
3 kHz flat	North American	3000 Hz low-pass for low-frequency noise
15 kHz flat	North American	15,000 Hz low-pass for program circuits
Program	North American	700–8000 Hz bandpass for program circuits
50 kilobit	North American	for 56 kbps data service at 135 $\Omega$
Psophometric	ITU (Europe)	Similar to C-Message filter
3 kHz bandpass	ITU (Europe)	275–3000 Hz for impulse noise measurements
Sound unweighted	ITU (Europe)	Program circuits for broadcast industry
Sound weighted	ITU (Europe)	Program circuits, studio-to-transmitter

When the transmitter is sending its 1004 Hz (ITU 820 or 1020 Hz) tone, the receiver invokes a very narrow notch filter to remove the tone prior to the noise measurement. The user must know enough about the circuit under test to determine if a tone is needed during noise measurements. At the time the standard was written, it specified and so named a quasi-RMS detector with characteristics that approximate a true RMS detector. Manufacturers that have implemented circuitry to meet these characteristics label their instruments as measuring QRMS noise.

The North American and the ITU versions differ distinctly in the display of the noise level being measured. The ITU versions display the level in dBm consistently with the level measurement, except they usually provide sensitivity as great as  $-90$  dBm. The North American standard created a new reference at  $-90$  dBm and called it 0 dBm. An instrument measuring a noise signal of level 0 dBm through a C-Message filter would, in a North American version, display 90 dBm. It would be referred to as 90 dBm<sub>C</sub>, with the C meaning the noise level was being made via the C-Message filter.

The standards also provide for a noise-to-ground or unbalanced measurement, which is performed simply by paralleling the input of the twisted-pair cable and reading the level with respect to instrument ground. This gives a measure of the unbalance of the cable.

**Signal-to-noise (s/n) ratio.** An instrument that can perform both the level and the noise measurements has all the circuitry needed to ascertain the signal-to-noise ratio. The measurement uses the 1004 Hz (ITU 820 and 1020 Hz) reference tone for the signal portion of the calculation and stores this level. Then the instrument cycles to a noise measurement with the notch filter invoked and a user-selected noise filter. Internal algorithms then compute and display the s/n ratio in dB. The IEEE 743-1984 standard specifies the maximum time that an instrument can process and display a reading.

**Envelope delay distortion (EDD).** *Phase distortion* is a measure of the nonlinearity of the phase vs. frequency of a voice channel. Because establishing a phase reference

is difficult, the derivative of the phase with respect to the frequency, called the *envelope delay distortion*, is what is measured. Envelope delay is the measure of time to propagate the modulation (the actual information applied to the signal) through the system. Distortion of this envelope while propagating through the system is described by the EDD measurement.


Nonlinearities in phase characteristics go undetected in voice applications but create serious problems with digital data streams. Envelope delay distortion is complex and should be understood before one examines how instruments make the measurement. Figure 25.9 compares a linear to a nonlinear channel. If one modulates the amplitude of a carrier and frequency-sweeps this signal over a nonlinear channel, the phase difference between the lower and upper sidebands will change according to the degree of nonlinearity. This change or relative delay can be measured and constitutes the envelope delay in time vs. frequency of the channel. The presence of changing delay vs. frequency is a measure of the envelope delay distortion of the channel.

The most common measurement configuration is end-to-end, a complication of which is getting the relative phase distortion information back to the transmitter for comparison to the phase of the transmitted signal. The scheme used is to design a transmitter that will operate in two modes. The modulated transmitter signal for the measurement channel, called the “Normal” setting, is capable of sweeping the channel over the voice band of frequencies. The receiver end is configured in a “Repeat” setting, which retransmits the modulated signal back to the sending set on a separate channel but at a fixed carrier frequency. Fixed changes in phase exist throughout this scheme, but the changes in phase vs. frequency are preserved and available for processing at the Normal test set end of the measurement.


The direction of test can be reversed by each set switching roles between Normal and Repeat modes. The North American version uses an 83% Hz modulating frequency on the variable-frequency carrier. The amplitude modulation yields two sidebands and it is their relative phase that is measured. The transmitter in the Repeat mode uses the 83% Hz modulating frequency extracted from the Normal signal, but the return carrier is fixed in frequency at typically 1800 Hz. As a result, all relative phase changes are from the nonlinearities of the Normal channel under test.

The ITU version uses a modulating frequency 41% Hz and also provides a single-channel measurement scheme. This single-channel scheme requires the two test sets to engage in a timesharing mode to alternately transmit test signal, store, and telemeter the data back to the originating test set. (Refer to the ITU standard for more discussion of this scheme.) Test sets also may provide a loopback measurement in which the measured EDD is indicative of the round trip of the circuit under test. These test sets make the assumption that the distortion present is equally distributed on both legs of the channel. The test set then assigns half of the distortion to each channel. (Any asymmetry of the legs will make such an assignment faulty.)


The receiver of the Normal test set receives and demodulates the fixed-carrier Repeat signal. It extracts its own transmitter’s modulating phase information and allows the user to set in a zero delay reading at the voice band center frequency, usually 1800 Hz. Then, as the transmitter carrier is swept across the voice band, the receiver reads out the delay in microseconds vs. frequency. Figure 25.9 depicts a typical EDD curve that results. The receiver in the Repeat mode serves only to de-




(a) Ideal Linear Phase Characteristic




(b) Practical Nonlinear Phase Characteristic




(c) Linear Phase with Superimposed AM Signal Components



(d) Nonlinear Phase with Superimposed AM Signal Components



(e) Envelope Delay Characteristics of (c)



(f) Envelope Delay Characteristics of (d)

**Figure 25.9 Envelope delay distortion (EDD).** This figure is best understood when read top to bottom. The left column depicts three different frequencies with corresponding sidebands imposed upon a circuit with linear phase characteristics. The envelope delay characteristic in the bottom panel (e) is flat to represent no distortion. The right column shows the same frequency spectrum imposed upon a circuit with nonlinear phase characteristics. The resulting curve for the delay (f) is not flat but distorted, hence the term *envelope delay distortion*.

modulate the received swept carrier and remodulates the same signal on the fixed frequency of the Repeat carrier. The Repeat receiver also will display locally the received frequency and level of the swept carrier.



### 25.3.4 Transients

Transients, the second of the four general categories of measurements outlined in section 25.3.1, includes impulse noise, phase and gain hits, and dropouts.

**Impulse noise.** *Impulse noise* is defined as voltage spikes that are much higher than the average background noise levels. The test instrument can measure the thresholds of impulse noise spikes and display a count of occurrences over predetermined lengths of time. Impulse noise in the voice channel can be heard as audible static best described as “pops and clicks.” Impulse noise in data channels is destructive to data packets, usually requiring them to be retransmitted. Amplitude and occurrence frequency are important impulse noise parameters for a communications channel.

The transmitter for impulse noise testing is in the off (quiet termination) mode, or is transmitting the 1004 Hz (or 820 and 1020 Hz) holding tone.

If the holding tone is present, then the receiver must invoke the notch filter before counting the impulses. The standards specify limits on counting rates. An impulse into a bandpass filter can cause ringing, so the counting rate limit is set to avoid counting the ring phenomenon. More sophisticated receivers perform three-level impulse noise measurements. In addition to an adjustable threshold level, they provide three counters that are offset by 3 or 4 dB and allow simultaneous counting at each level.


**Phase hits, gain hits, and dropouts.** Phase hits, gain hits, and dropouts—collectively referred to as *hits and dropouts*—are three additional transient phenomena that have been quantified. Figures 25.10 and 25.11 show representations of these transients.

A *phase hit* can be caused by a change in routing of the channel and can be disruptive of phase-modulated signals. A *gain hit* is an abrupt change in signal level in a channel. When a gain hit causes a signal level reduction greater than 12 dB, it is reclassified as a *dropout*. Companies that provide data channel services document the transient characteristics of a new service, establishing a quality baseline for that circuit. If follow-on maintenance is required later, transient measurements can be made and compared to the original numbers. Possible circuit deterioration thus can be ascertained.

These measurements are all made with the 1004 Hz (ITU 820 and 1020 Hz) tone present. The instrument monitors and counts the occurrence of these transient phenomena in the holding tone. A phase hit would be a sudden shift in phase as small as 5 degrees detected in the received holding tone. In addition, the level is monitored for sudden shifts (gain hits) to higher or lower levels up to 10 dB. A level shift of 12 dB or greater is declared a dropout. All three should be counted simultaneously.

### 25.3.5 Incidental modulation

Incidental modulation could be defined as inadvertent or unwanted change in the information applied to a carrier. In digital circuits, such modulation usually takes the form of *jitter*.



**Figure 25.10 Transient phenomena.** A test signal that is not subject to interfering transient signals would appear as a clean sine wave as shown in the top waveform. Each interfering transient signal is shown as labeled in each succeeding waveform.

<u>Transient</u>	<u>Signal Related</u>	<u>Characteristic</u>	<u>Parameter Duration</u>
Impulse Noise	No	Level & Phase	< 0.1 ms to 4 ms
Phase Hit	Yes	Phase	> 4 ms
Gain Hit	Yes	Level	< 12 dB
Dropout	Yes	Level	> 12 dB

**Figure 25.11** Transient measurements.

**Phase jitter.** Jitter is a descriptive term that implies something is not constant, i.e., it is moving, in this case signal phase. Where “phase hit” describes an abrupt transient, *phase jitter* is about an unwanted phase modulation of the desired signal. Since circuits are maintained with a holding tone, this tone can be analyzed, yielding the inherent phase jitter of the circuit without signal. Excess jitter will have a corrupting influence upon the communication channel. The most common sources of

phase jitter are interference from the 20 Hz ringing current and the 60 Hz commercial power system. A narrowband filter of about 800 Hz centered on the holding tone will allow the test equipment to process and detect this phase modulation.

**Amplitude jitter.** As with phase jitter, the instrument should detect the incidental modulation of the ringing signal and 60 Hz power system on the amplitude of the holding tone.

### 25.3.6 Other measurements

In addition to not fitting into the three preceding categories, some of what follows tends toward obsolescence.

**Return loss.** *Return loss* is a measure of how well a cable's characteristic impedance matches the cable termination. An impedance mismatch gives rise to a reflection of any transmitted signal. The loss is displayed in dB and is calculated by the formula

$$10 \log_{10} \left( \frac{\text{transmitted signal}}{\text{received signal}} \right) \quad (25.1)$$

Telephone system engineers have defined two subsets of return loss that have specific meaning. Test equipment exists for *echo return loss* (ERL) and *singing return loss* (SRL). The ERL and SRL have their origins in the all-analog channel, where echo and singing were readily identifiable complaints in a malfunctioning voice system. The terms *echo* and *singing* tend to be self-descriptive; echo means a reverberative effect on the analog channel, and singing is a high-pitched oscillation found in older analog systems. The spectral content was heuristically determined and test sets are built to specifically measure these phenomena.

**P/AR (peak-to-average ratio).** P/AR is a shorthand for *peak-to-average* ratio. Identified only in the North American standard, it will not appear in the new revised IEEE 743 standard. It is discussed here only because many test sets still have this measurement capability. It is unique in measurement and implementation in that it requires the ability to transmit a special 16-tone spectrum of frequencies and to process the received signal simultaneously with both a peak detector with an average detector. The receiver solves the equation

$$\text{P/AR} = 2 \left( \frac{E_{\text{peak}}}{E_{\text{avg}}} \right) - 1 \quad (25.2)$$

and displays a number of 100 or less P/AR units. A P/AR of 100 units means that the transmitted spectrum has suffered no degradation whatever during transmission.

The measurement tends to lump together the effects of envelope delay distortion, poor return loss, and bandwidth reduction in the channel. It also tends to be insensitive to noise, nonlinear distortion, and transient phenomena. It is a relatively expensive measurement to implement in that it requires special dedicated hardware

and software that cannot be used for other measurements. The P/AR measurement is most useful when the network operator has properly baselined all the circuits in the system. Then any degradation can be compared to the baseline. Without this baseline, the P/AR becomes somewhat subjective.

**Intermodulation distortion.** Intermodulation (IM) distortion occurs when there are nonlinearities in the channel. The measurement described here is the 4-tone method covered in IEEE 743-1984.[4] At this writing a replacement method using 23 tones has been defined and will be published in the revised standard. Manufacturers will be using this newer method in the next generation of test equipment releases.

In the 4-tone system the transmitted signal consists of four tones of equal amplitude. Two are 6 Hz apart and centered at 860 Hz; the other two are 16 Hz apart and centered at 1380 Hz. The receiver has three narrowband filters centered on 1900 Hz for 3rd-order products, and on 520 Hz and 2240 Hz for 2nd-order products. The following equations are solved to provide a display in dB of the 2nd and 3rd intermodulation products:

$$V_S = \sqrt{\frac{V_L^2 + V_H^2}{2}} \quad (25.3)$$

$$\text{2nd IM product} = 20 \log_{10} \left( \frac{V_T}{V_S} \right) \quad (25.4)$$

$$\text{3rd IM product} = 20 \log_{10} \left( \frac{V_T}{V_M} \right) \quad (25.5)$$

where:

$V_T$  = RMS value of 4-tone signal

$V_M$  = RMS value of 1900 Hz filtered signal

$V_L$  = RMS value of 520 Hz filtered signal

$V_H$  = RMS value of 2240 Hz filtered signal

In the 23-tone system, the transmitted signal consists of 23 equally spaced frequencies with known phase relationships. Part of the rationale of choosing a 23-tone system is that it approximates the energy spread of a typical high-speed, echo-canceling, voice-band modem. Other implementation factors include the PCM sampling rate of 8000 Hz and the availability of DSP (digital signal processing) technology for a 512-point DFT (discrete Fourier transform).

The receiver also will be implemented using DSP technology. A 512-point DFT will define 256 measurement bands, each 15.625 Hz wide, in the 0 to 4 kHz voice band. With computer analysis of these bands, one can extract the 2nd, 3rd, and total harmonic distortion of the channel under test. Additionally, one can extract the EDD and the s/n ratio. The revised IEEE 743 standard (reviewed in draft form) extensively documents this new technique. At this writing, however, no commercial products are known to use this technique. It should be presumed that the instruments, when they are made available, will dramatically change how traditional TIMS measurements are made.

## 25.4 Network Testing on Fiber Optic Media

Analog measurements on fiber optic media are limited primarily to power loss and optical return loss at the specified wavelength. Only a cursory overview is presented here because these measurements are not network-specific. They are covered adequately in references on fiber optic measurements.

Power meters are the optical equivalent of electrical power meters, or voltage measurements at a specified impedance. The optoelectric (optical-to-electrical) transducer is a key component and also is the limiting factor on the range of wavelengths that can be measured. The equivalent to the transmitter is called a *source*, either a laser or LED, and is limited to single wavelength or very narrow optical spectrum dispersion. *Loss Test Sets* combine a source and a power meter that can be used to measure loss in dB of a fiber segment at a specified wavelength.

An OTDR (optical time domain reflectometer) can be used to measure return loss and also will give a graphic picture of defects in the fiber. OTDR measurements are covered in Chapter 28.

## 25.4 References

- Coombs, Clyde F., Jr. *Electronic Instrument Handbook, 2nd Ed.* (New York: McGraw-Hill, 1995.) [1]  
*EIA/TIA 568A & TSB-67 Level II, Commercial Building Telecommunications Wiring Standard.* (Washington, D.C.: Electronic Industries Association.) [2]  
*IEEE 802.n, Local Area Networks.* (Piscataway, N.J.: Institute of Electrical and Electronics Engineers.) [3]  
*IEEE 743-1984, Standard Methods and Equipment for Measuring the Transmission Characteristics of Analog Voice Frequency Circuits.* (Piscataway, N.J.: Institute of Electrical and Electronics Engineers, 1984.) [4]  
*ANSI X3T9.5 (FDDI), Standard Fiber Distributed Data Interface.* (Washington, D.C.: American National Standards Institute.) [5]  
*Series O, Vol. 4, International Telecommunications Union.* (Geneva: CCITT.) [6]



# Bit Error Rate Measurements and Error Performance Analysis

Hugh Walker

*Hewlett-Packard Company, South Queensferry, Scotland*

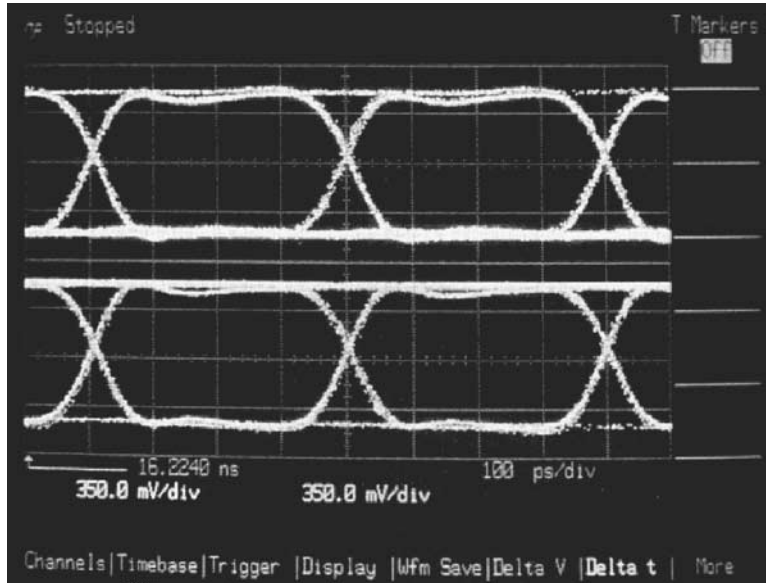
## 26.1 Introduction

The fundamental measure of performance or quality in digital systems is the probability of any stored or transmitted bit being received in error. With the latest equipment, the probabilities are very low, on the order of  $10^{-12}$  or less. It remains necessary, however, to measure the performance of these systems and in particular to analyze the margins of safety available while exploring potential weaknesses that could lead to degraded performance. This is the purpose of digital pattern generators and error detectors, sometimes referred to as *bit error rate test sets*, or BERTS.

As telecommunication systems have become more sophisticated and carry a wide range of voice, data, and video traffic, new quality of service (QoS) standards, referred to as *error performance analysis*, have evolved.

## 26.2 Sources of Errors

Errors in digital systems arise as a result of several distinct practical effects. When viewing a random digital signal on an oscilloscope, the common eye diagram shown in Figure 26.1 is displayed. To obtain the eye diagram display, the sweep is triggered using the data clock source, and the time base is adjusted so that, say, two or four bit periods are displayed. On each successive sweep, the random bit patterns (and transitions) build up on the display, either through the persistence of the screen phosphor or through digital storage. With random data, all possible combinations of bit sequences will be explored, so the eye diagram shows the extent of pulse distortion that might occur. The eye diagram is important to digital circuit designers and testers because it shows at a glance the quality of a system.



**Figure 26.1** The eye diagram of a random digital signal displayed on a digital sampling oscilloscope. Over successive sweeps, the random bit patterns build up a composite picture of all the possible pattern sequences and transitions. The large open area in the center of the pulse is called the *eye opening*, with the 1 value at the top and the 0 value at the bottom. The distance between the top and the bottom at the center of the pulse is called the *eye height*, while the distance between the transitions is called the *eye width*. The eye diagram is a useful qualitative measure of a digital system's performance. An engineer can spot immediately if the eye is degraded by noise, timing jitter, pulse degradation, or intersymbol interference (ISI). As explained in the text, ISI is pattern-dependent, so different test patterns will create different eye diagrams.

The displayed eye diagram exhibits an open area in the center of the pulse separating the 1 level from the 0 level. This is termed the *eye opening*. The wider this gap, the lower the probability of a 1 being confused with a 0 and vice versa. The space between adjacent vertical transitions at the edges of the pulse is termed the *eye width*. The wider this gap, the more tolerant the system will be of the point at which the digital signal is sampled to determine the instantaneous binary value. Errors occur either when the eye opening is partially closed or when the relative sampling instant is displaced by timing jitter, described below. The following subsections discuss sources of errors.

### 26.2.1 Excessive noise

Excessive noise generally arises from thermal noise or crosstalk in sensitive circuits that, by addition to the binary signal, closes the eye opening and creates a “fuzziness” on the eye diagram display. Examples of noise-limited systems are an optical receiver operating near its minimum light level or a digital radio or satellite receiver operating with a low signal attenuated by fading. A characteristic of noise-limited systems is that the errors occur randomly according to a Poisson distribution and are not influenced by the transmitted bit pattern.




### 26.2.2 Intersymbol interference

*Intersymbol interference (ISI)* is caused by short-term storage effects in a digital circuit. At any time, the received signal represents not only the current digital value but also the residues of previous digital values (determined by the system's impulse response). With a random digital signal, these residuals appear on the eye diagram rather like thermal noise. The levels are not random, however, but are determined by the preceding digital pattern. The purpose of the pattern generator is to simulate a wide range of digital sequences and stress patterns to explore the effects of ISI. This sometimes is referred to as the *pattern dependency* of a system.

### 26.2.3 Timing jitter

Timing jitter causes a horizontal displacement of the eye diagram and so reduces the eye width. One of the best definitions of timing jitter has been provided by the ITU-T (International Telecommunication Union-Telecom Standards Committee). "Timing jitter is the short-term variation of the significant instants of a digital signal from their ideal positions in time." The significant instant might be the rising or falling edge of a pulse.

The effect of jitter can be seen in the diagram in Figure 26.2. At certain points in time, the pulse is significantly offset from its correct position. If this offset be-



**Figure 26.2** Timing jitter disturbs the pulse from its ideal position in time, and the perturbations cause a narrowing of the eye area, as shown in (a). Seen in real time (b) at instants  $T_1, T_2, T_3$ , and so on, one can see that the bit pattern is displaced from the ideal positions in time. The instantaneous offsets,  $t_1, t_2, t_3$ , and so on from the ideal positions form the jitter function  $J(t)$ . If jitter becomes excessive, the eye opening will be closed sufficiently to cause errors when data is sampled. Sampling usually is timed to occur at the center of the eye, at the point of greatest eye height.

comes large, then there will be an error when we try to sample and decode the digital signal. The disturbance or offset of the timing instant usually is measured in unit intervals (UIs) peak to peak, where the unit interval is equivalent to one bit period.

Jitter arises from a variety of causes, including superimposed noise and crosstalk affecting the trigger point of logic decision circuits. Another important source of jitter is clock recovery circuits that derive the reference sampling clock from a received data stream. Depending on the pattern density, the clock recovery circuit (usually a tuned tank circuit or phase-locked loop) may drift toward its natural frequency and lose or accumulate phase shift relative to the incoming data. In analyzing the source of errors, it often is necessary to check jitter levels on reclocked data and to check a system's tolerance to specific levels of superimposed jitter. These measurements usually are carried out with a dedicated telecommunications jitter test set, a digital sampling oscilloscope, or a time-interval analyzer. (For more information on jitter, see Chapter 23.)


### 26.3 Error Measurements

To test a digital system, the input is stimulated by a test pattern. Usually this is a *pseudorandom binary sequence* (PRBS), although other specific stress patterns (referred to as *user-defined word patterns*) may be used to explore limits of performance. Typical stress patterns might include long runs of 0s to test clock recovery or patterns with alternate periods of high and low *mark density* (or "1 density") to check storage effects in optical transmitters and receivers that may respond to average mean level in a data signal.

For telecommunications and data communications transmission systems, the object is to simulate the random traffic experienced under normal operating conditions. The problem with a truly random signal is that the error detector will have no means of knowing the actual bit values that were transmitted and therefore no way of detecting errors. Used instead is a pseudorandom signal, which has virtually all the statistical characteristics of a true random signal and appears as such to the item under test. In fact, it is completely deterministic and therefore predictable by the error detector. A range of maximal-length PRBS patterns has been specified for this, as described subsequently. At the error detector (Figure 26.3), the output of the system under test is compared bit by bit with a locally generated, error-free reference pattern.



**Figure 26.3** To test a digital transmission system, a pattern generator is connected to the input of the system under test, and an error detector is connected at the output.



**Figure 26.4** Three methods of computing BER. Method 1, used on early BER testers, simply defines the measurement time by the number of clock periods ( $10^6$ ,  $10^7$ , and so on). The accumulated error count then can be converted easily to BER; however, the measurement period varies according to the bit rate. Method 2 defines the measurement period in seconds, minutes, and hours, and a microprocessor is used to calculate BER from the accumulated error count and clock count in that period. The advantage is that the measurements are consistent with error-performance standards. Method 3 determines the measurement period as that required to accumulate sufficient errors for a statistically significant result, e.g., 100 errors. This may lead to very long measurements at low BER values.

The probability of error in any transmitted bit is a statistical property and has to be treated as such. Any attempt to measure this over a given time period can be expressed in various ways, the most common of which is

$$\text{bit error ratio (BER)} = \frac{\text{number of errors counted in the averaging interval}}{\text{total number of transmitted bits in the averaging interval}} \quad (26.1)$$

Clearly the result will have a statistical variance from the long-term mean error ratio, dependent on the size of the sample taken from the population—in this case the number of errors counted.

Three methods of computing BER are in general use; these are illustrated in Figure 26.4. The first method, common on early test sets, simply counted the number of clock periods to provide a time base or averaging interval. This could be implemented easily using discrete logic decade dividers.

Now that microprocessors are available, more convenient gating periods are used. In the second method, a timed gating period of, say, 1 second, 1 minute, or 1 hour is used, and a calculation of BER is made from the accumulated totals. The advantage of this method is that it provides results that are compatible with the error-performance criteria discussed subsequently.


The third method determines the gating period by counting sufficient errors (typically 100 or more) for statistically reliable results. Again, the processor calculates BER from the accumulated totals. This method can lead to very long gating periods with low BER values. For example, a system running at 100 Mbps with a BER of  $10^{-12}$  would take nearly 12 days to accumulate 100 errors.

The most commonly used method is the second, which calculates BER after a fixed, repetitive gating period. In this case, the variance in the result will continuously change, so it is normal to give some kind of warning if the variance exceeds generally acceptable levels. The most widely accepted level is 10 percent, i.e., an error count of at least 100 errors. In practical digital transmission systems, particularly those using radio propagation, the BER can vary substantially over time. In this case, the long-term mean value provides only part of the story. Communications engineers also are interested in the percentage of time the system under test is unacceptably degraded. This is called *error analysis* or *error performance*, which is discussed in section 26.3.3.

### 26.3.1 Test patterns

As alluded to previously, the choice of test pattern is usually made between a PRBS (to simulate traffic) and specific word patterns (to examine pattern-dependent tendencies or critical timing effects). With a PRBS, the choice of binary sequence and the resulting spectral and run properties are important. These properties may be summarized:


- Sequence length, in bits
- Shift register feedback configuration defining binary run properties
- Spectral line spacing, which depends on bit rate (see Figure 26.5).



**Figure 26.5** This is the ideal spectrum of a binary nonreturn to zero (NRZ) PRBS. It has a line spectrum in which the line spacing is determined by the bit rate ( $f_b$ ) and the sequence length  $2^n - 1$  (see Table 26.1).

**TABLE 26.1 Relationship Between Sequence Length and Bit Rate with Corresponding Spectral Line Spacing for Some Standard Telecommunications Transmission Rates.**

Bit rate ( $f_b$ ), kbps	Sequence Length ( $n$ )	Polynomial	Spectral Line ( $f_b/n$ ), Hz
1,544	$2^{15}-1$ bits	$D_{15} + D_{14} + 1 = 0$	47.1
2,048	$2^{15}-1$ bits	$D_{15} + D_{14} + 1 = 0$	62.5
34,368	$2^{23}-1$ bits	$D_{23} + D_{18} + 1 = 0$	4.1
44,736	$2^{15}-1$ bits	$D_{15} + D_{14} + 1 = 0$	1365.3
139,264	$2^{23}-1$ bits	$D_{23} + D_{18} + 1 = 0$	16.6



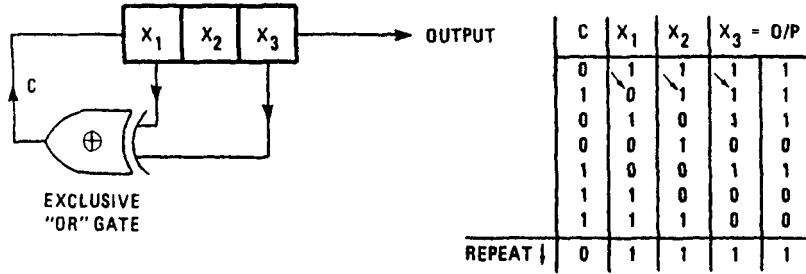
**Figure 26.6** The PRBS pattern generator consists of a shift register with feedback taps connected to an exclusive-OR gate. The output of the shift register may be inverted to create a signal with the maximum run of 0s. The polynomial that defines the PRBS determines where the feedback connections are made. The expression for this diagram would be  $D_n + D_{n-1} + 1 = 0$ .

PRBS patterns have been standardized by the ITU-T for testing digital transmission systems (Recommendations O.151, O.152, and O.153). The most commonly used patterns in digital transmission testing are summarized in Table 26.1.

Table 26.1 shows some examples of sequences specified by the ITU-T for several standard telecommunications bit rates. Note that the longer sequences give closer spectral line spacing; typically, the higher the operating bit rate, the longer is the required sequence to simulate real data traffic. For tests in the Gbps range, some test sets now provide a  $2^{31}-1$  sequence length.

Adequate (i.e., close enough) spectral line spacing is important when testing systems containing relatively narrowband (high-Q) clock timing recovery circuits in order to see the jitter contribution of these and its effect on error performance. The shift register configuration is defined by a polynomial of the type shown in Table 26.1. The letter  $D$  stands for delay; the expression  $D_{15} + D_{14} + 1 = 0$ , for example, means that the outputs of the fifteenth and fourteenth stages of the shift register are connected to an exclusive-OR gate, the output of which drives the first shift register stage, as shown in Figure 26.6.

This basic circuit arrangement generates a sequence with a maximum run of 1s rather than 0s. It is common to invert the output to generate a maximum run of 0s, since this may create more stringent conditions for a clock recovery circuit. The simple three-stage PRBS generator shown in Figure 26.7 with the truth table helps to explain the operation of the feedback shift register. This has a sequence length of  $2^3-1$ , or 7 bits.



**Figure 26.7** The simple three-stage PRBS generator demonstrates how the feedback shift register creates the pseudorandom sequence. The output of the exclusive-OR gate provides the input to the shift register. The truth table shows how the bit pattern steps through the shift register with each clock pulse.



**Figure 26.8** This graph shows the running digital sum for a commonly used PRBS. Notice the initial steep descent of the graph, indicating a predominance of 1s or 0s (depending on inversion) for the initial period of the sequence.

The choice of shift register configuration affects the run properties of the PRBS; an example for the  $D_{15} + D_{14} + 1 = 0$  polynomial is shown in Figure 26.8. This graph shows that there are periods in the sequence with a low number of 1s relative to 0s, which is more stressful to clock recovery circuits. Run properties affect timing jitter in terms of the length of 0 blocks over which phase error is accumulated by the timing recovery circuits. This leads to pattern-dependent jitter and, if not controlled, to errors.

Some test sets allow users to take a standard PRBS and modify it using 0 substitution and variable mark density. A maximal-length ( $m$ -sequence) PRBS usually has a maximum run of 0s equal to one less than the degree of the defining polynomial. This can be extended by deliberately overwriting subsequent bits with 0s to any extent the user requires. Zero substitution is useful for checking the limits of a clock recovery circuit. Variable mark density allows the user to change the density of 1s in the pattern from the regular 50 percent of a PRBS pattern. A test set may allow the ratio to be varied from 12.5 percent to 87.5 percent. The advantage of modifying a standard PRBS is that many of the “random” traffic-simulating characteristics are retained while the user explores the limitations of a system under fairly realistic conditions.

In addition to PRBS patterns, most test equipment provides at least some facilities for the user to program a repetitive-word pattern. At the simplest level, this could be a particular pattern with minimum or maximum density of marks, for example, the 3-in-24 pattern used in North America for testing TI circuits at 1.544 Mbps. This pattern complies with the minimum mark density of 12.5 percent while providing the longest allowed run of 0 bits, namely 15.

With the increasing complexity of telecommunications equipment, however, there also is the need to simulate frame structures and even specific data messages. Indeed, many systems today will not respond to an unstructured, unframed PRBS. Framing is used by telecommunications equipment to identify individual channels that are time division multiplexed together in a serial data stream. To mark the beginning of a sequence of multiplexed channels, the telecommunications equipment transmits a fixed frame word or pattern of bits that is recognized at the receiving end. If the frame word pattern is not present, the receiving equipment will assume a fault and transmit alarms (see Chapter 3).

Pattern generators therefore will require large pattern memories that can be programmed from a computer to simulate a wide range of framed test signals and, of course, to simulate specific malfunctions and errors in those framed test signals. Figure 26.9 shows a typical range of patterns, downloaded from disk, available on a Gbps pattern generator and error detector.


### 26.3.2 Error detection

The basic error-detection process involves comparing the incoming data stream from the item under test with a reference pattern in an exclusive-OR gate. When both data patterns are identical and synchronized, the exclusive-OR gives zero output. When a difference exists, i.e., when an error is present in the received stream, an output is generated as shown in Figure 26.10. This simple circuit will detect all logic errors. An additional feature allows errored 1s and errored 0s to be counted separately. Error mechanisms in some devices, such as laser transmitters, lead to predominance of one polarity over the other.

A further refinement available on some instruments allows the user to identify the exact bit within a long word pattern that is in error. Error generation often is systematic, that is, it is caused by inherent deterministic impairments in the equipment such as intersymbol interference, storage effects, and delay-line reflections. A particular sequence of bits in the pattern creates a high probability of error in a subsequent

		16:22:39 FEB 1, 1993	MENU	
DISC	HP 78842B ERROR DETECTOR (Patterns)			(0,17)
PATT 6				CURRENT PATTERN
	Current Pattern	INACTIVE	Length:	127
DISC	2^7	from Patt 4		
PATT 7	Patt. 1:	SONET STS-48	Length:	8,192
	Patt. 2:	2^11	Length:	2,048
DISC	Patt. 3:	SONET STS-48	Length:	1,152
PATT 8	Patt. 4:	2^7	Length:	127
	Patt. 5:	SONET STS-12	Length:	77,760
DISC	Patt. 6:	SONET STS-48	Length:	311,040
PATT 9	Patt. 7:	CID STM-4	Length:	20,720
	Patt. 8:	CID STM-16	Length:	22,440
DISC	Patt. 9:	SDH STM-4	Length:	77,760
PATT 10	Patt. 10:	SDH STM-16	Length:	311,040
	Patt. 11:	FDDI Jitter	Length:	1,280
DISC	Patt. 12:	FDDI Wander	Length:	90,000
PATT 11				
DISC				
PATT 12				
				CANCEL EDIT

**Figure 26.9** An example of stored word patterns in a BER test set. The list includes some simple PRBS patterns, SONET/SDH frame patterns, and test patterns for FDDI LAN applications.



**Figure 26.10** Error detection using an exclusive-OR gate simply involves comparing bit by bit the received pattern with the reference pattern. When a difference is present, the exclusive-OR gate gives an output.

bit. Errored bit identification helps the user to focus on the portion of pattern that creates problems and then to investigate the causes. Counting bits from the start of the pattern as a reference point, the instrument locates the position of an errored bit and identifies it. On successive repeats of the pattern it calculates the “bit” BER, which, if the cause is systematic, will be far higher than the average BER for the pattern.

Of course, if the error is purely random, then on average the “bit” BER will be similar to the overall average. This capability, in conjunction with a digital sampling oscilloscope, is useful to an engineer designing equipment and trying to locate the often complex sources of errors. For example, reflections in transmission lines often give rise to pulse degradation that is many bits delayed from the launch pulse. By monitoring a frequently errored bit, the engineer can adjust the word pattern to explore the effects of ISI and reflections.

The reference pattern for error detection could be supplied locally from the pattern generator, but normally a separate reference pattern generator is provided in the error detector so that the transmit and receive portions of the test equipment can be separated—an important requirement for end-to-end testing of telecommunications links.



**Synchronization** This raises the question of synchronization of the two patterns to be compared before error detection can commence. In order to synchronize rapidly yet remain synchronized at high error rates or during large error bursts, it's necessary to establish a *sync criterion BER* that has variable gate times over which a test for synchronicity is made.

The sync criterion may be expressed as

$$\text{sync gain} = \text{sync loss} \quad (26.2a)$$


$$x/n = X/N \quad (26.2b)$$

where

$x$  and  $X$  = error counts  
 $n$  and  $N$  = total bit counts  
 $x \ll X$   
 $n \ll N$

(**Note:** To avoid oscillation, it is normal to make the sync loss BER greater than the sync gain BER, while the gate time for sync gain is much less than for sync loss. The longer period for sync loss ensures that the error detector is not thrown out of lock by a burst of errors. For example, one test set has a sync loss criterion of >20,000 errors in 500,000 bits and a sync gain criterion of <4 errors in 100 bits.)

The normal method of achieving synchronization is to open the feedback loop in the reference pattern shift register and feed the input data signal into the register until it is full, close the feedback loop, and test for sync (Figure 26.11). Clearly, two PRBS patterns out of sync have a BER of approximately 0.5, so the sync criterion must be lower than this. Some error detectors allow the user to set the sync gain/loss



**Figure 26.11** To obtain pattern synchronization in a closed-loop error detector, the feedback loop on the reference pattern generator is opened temporarily, and a sample of the input bit pattern is fed into the shift register. The loop is then closed, and a synchronization check is made by checking the output of the error detector. A BER of approximately 0.5 indicates a loss of synchronization. In this case, the process is repeated until synchronization is obtained. Some test sets allow the criteria for sync gain and loss to be adjusted for a particular application.


criteria, for example, by setting the threshold BER in a particular gating period. Since measurements are invalid during resynchronization, it is desirable to minimize these periods.

**In-service measurements** Finally, as mentioned earlier, most communications equipment operates with a standard frame structure. Framing allows multiplexed channels to be identified in a high-speed data stream. Frame structures also carry information on system status and error detection, for example using cyclic redundancy checksum (CRC) or parity. The frame word itself provides a constant repetitive pattern even when operating in live traffic and so provides a simple way of monitoring errors. Some error detectors, specifically designed for maintenance of digital communications systems, lock onto the frame pattern and report on frame and code errors while the system is in service in much the same way as logic errors in a PRBS are counted in an out-of-service test.

### 26.3.3 Error analysis

In practical systems, the BER can vary substantially over time due to propagation effects in radio and satellite systems, electrical interference, and random traffic patterns (Figure 26.12). For example, a system may have a satisfactory long-term BER level while being subjected to short bursts of intense errors interspersed with long periods of error-free operation. If the system is subjected to high BER for periods of more than a few seconds, it is considered to be unusable and is described as “unavailable.” In commercial communications, customers want to know what grade of service a network operator guarantees to provide because this determines the price paid for the service. This is generally classified as the percentage of time a service meets or exceeds expectations.

To meet this measurement requirement, the error detector must measure BER and count errors and classify the results as percentages of the total elapsed mea-

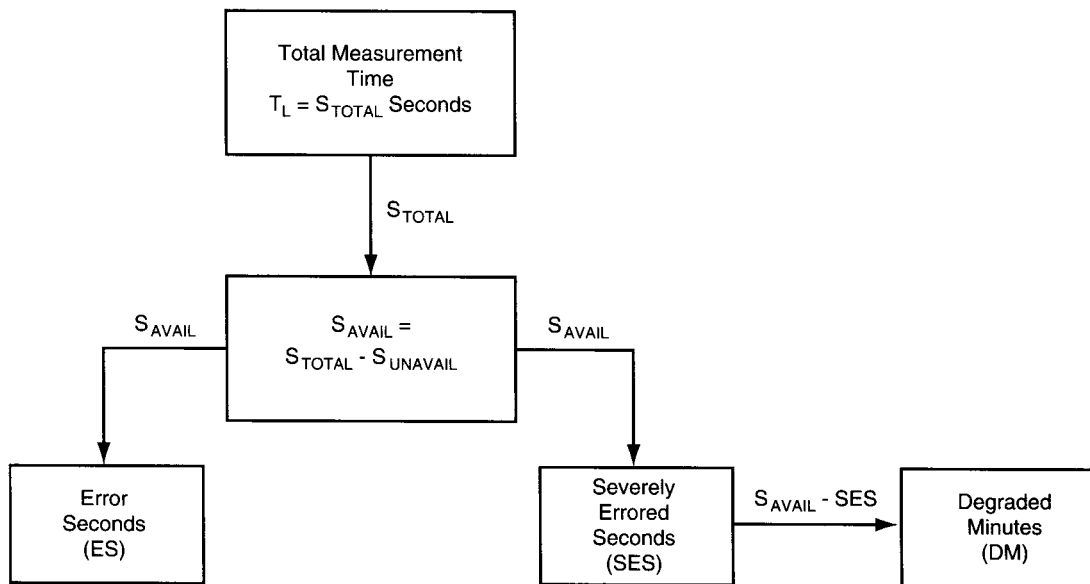


**Figure 26.12** Error performance. In a practical system, particularly those using radio propagation, the BER may vary a great deal over time. One could measure the average long-term BER, but communications engineers find it more useful to categorize the performance into bands and then express the error performance in terms of the percentage of time each threshold is exceeded.

**TABLE 26.2 ITU-T Recommendation G.821: Error Performance of an International Digital Connection Forming Part of an ISDN\***

Performance Classification	Objectives
Degraded minutes (DM)	Fewer than 10 percent of 1-minute intervals to have a bit error ratio worse than $10^{-6}$
Severely errored seconds (SES)	Fewer than 0.2 percent of 1-minute intervals to have a bit error ratio worse than $10^{-3}$
Errored seconds (ES)	Fewer than 8 percent of 1-minute intervals to have any errors (equivalent to 92 percent error-free seconds)

\*Measured over a period  $T_L$  (e.g., 1 month) on a unidirectional 64 kbps channel of the hypothetical reference connection (HRX) of 27,500 km.



**Figure 26.13** This diagram shows how the error detector classifies the error measurements according to the G.821 criteria. From the total measurement time  $T_L$ , the tester subtracts periods of unavailable time (periods of 10 seconds or more when the BER is worse than  $10^{-3}$ ). From the remaining available time ( $S_{avail}$  seconds), errored seconds are accumulated, and simultaneously, any seconds with the BER worse than  $10^{-3}$  are accumulated as severely errored seconds (SES). The remaining nonseverely Errored Second periods are grouped together in 60-second blocks, and any that have an average BER of worse than  $10^{-6}$  are classified as degraded minutes (DM).

surement period for which certain thresholds are exceeded. This is called *error analysis*. The industry standard specification is ITU-T Recommendation G.821, shown in Table 26.2.

G.821 defines how error performance parameters are calculated in accordance with the flow diagram shown in Figure 26.13. The total measurement time is divided into 1-second periods, and unavailable time is subtracted to obtain the available time on which the G.821 parameters are calculated.

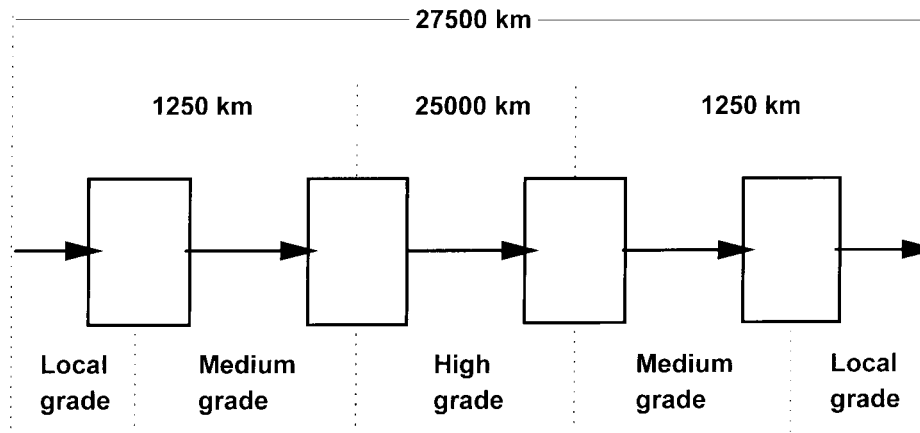
A period of unavailable time begins when the BER in each second is worse than  $10^{-3}$  for a period of 10 consecutive seconds. These 10 seconds are then considered to be unavailable time. A new period of available time begins with the first second of a period of 10 consecutive seconds each of which has a BER better than  $10^{-3}$ .

During available time, any second containing one or more errors is logged as an errored second (ES). Any period of 1 second with a BER exceeding  $10^{-3}$  is classified as a severely errored second (SES) and is subtracted from available time. The remaining seconds are grouped into 60-second periods. Any of these 1-minute periods with a mean BER exceeding  $10^{-6}$  are classified as degraded minutes (DM).

Almost all BER test sets now incorporate the basic G.821 analysis outlined above. When the original G.821 recommendation was formulated, communications engineers principally were concerned with the quality of voice and low-speed data transmission based on channel bit rates.

G.821 refers to the overall performance for the end-to-end 64 kbps connection on a very long (27,500 km) international connection. It's the starting point for error performance measurements, but needs interpretation before being applied to a practical transmission system operating at rates different to 64 kbps or over shorter distances. To handle the distance element, *apportionment* is used to decide how much of the recommended G.821 performance budget can be allocated to different sections of the network.

**Apportionment** The *hypothetical reference connection* (HRX) defined in G.821 is shown in Figure 26.14. The connection consists of a local- and medium-grade section at each end of the link, and a high-grade, long-distance section in the middle. A high-capacity transmission system such as microwave or lightwave would be characterized as a high-grade section. Typically a low-grade section would be the metallic sub-



**Figure 26.14** The Hypothetical Reference Connection (HRX) is a unidirectional 64 kbps international connection of 27,500 km. This is considered to be one of the longest paths in the telecom network and the error performance standards applied to this reference circuit should be exceeded by most practical, shorter connections. The HRX is divided into five sections, a local-grade and medium-grade section at each end, and a long-distance high-grade section in the middle. These definitions are used for apportioning the error performance objectives by network section.

**TABLE 26.3 Allocation of Degraded Minutes and Errored Second Objectives for the Hypothetical Reference Circuit (HRX).**

Circuit classification	Allocation of Degraded Minutes and Errored Seconds objectives
Local-grade (two ends)	15% block allowance to each end
Medium-grade (two ends)	15% block allowance to each end
High-grade	40% (equivalent to conceptual quality of 0.0016%/km)

**TABLE 26.4 Allocation of Severely Errored Seconds in the Hypothetical Reference Circuit (HRX).**

Circuit classification	Allocation of Severely Errored Seconds objectives
Local-grade (two ends)	0.015% block allowance to each end
Medium-grade (two ends)	0.015% block allowance to each end
High-grade	0.04%

scriber loop. Medium-grade would be the connection from, say, a local exchange to a trunk switching center. The three types of section in the network are allocated different portions of the total G.821 specification.

As shown in Tables 26.3 and 26.4, the allocation of degraded minutes and errored seconds is handled in a slightly different way than severely errored seconds. Low- and medium-grade sections are allocated a block allowance of 15 percent of the total G.821 specification at each end (i.e., 1.5 percent DM and 1.2 percent ES) irrespective of length. The longer high-grade section is apportioned on a distance basis so that 40 percent allowance is reduced in the ratio  $L/25000$ .

Thus, for a high-grade section of length  $L$  km, Equations 26.3 and 26.4 give the values for allowable DM and ES.

$$\text{Allowable DM} = \frac{10\% \times 0.4 \times L}{25000} \quad \text{or} \quad 0.00016\%/km \quad (26.3)$$

$$\text{Allowable ES} = \frac{8\% \times 0.4 \times L}{25000} \quad \text{or} \quad 0.000128\%/km \quad (26.4)$$

Severely errored seconds are allocated on a block basis only. Of the total 0.2 percent G.821 specification, 0.1 percent is allocated on a block basis as shown in Table 26.4. The remaining 0.1 percent SES is allocated to medium- and high-grade sections to account for adverse operating conditions such as propagation in microwave radio systems. For example, G.821 recommends that an additional 0.05 percent SES may be allocated to a medium- or high-grade microwave radio section of 2500 km.

**26.3.3.2 Application at higher bit rates** Error performance standards usually refer to measurements of 64 kbps, whereas practical measurements on transmission systems

are invariably made at a higher multiplex rate. The 1988 (Blue Book) version of G.821 (Annex D) gives provisional guidelines for conversion:

1. Percent DM converted directly
2. Percent SES converted directly with the addition of percent time with loss of frame alignment

Error second estimation is given by Equation 26.5, where  $n$  is the number of errors in the  $i$ th second,  $N$  is the higher bit rate divided by 64 kbps, and  $j$  is the total number of seconds.

$$ES_{64 \text{ kbps}} = \frac{1}{j} \sum_{i=1}^{i=j} \left( \frac{n}{N} \right)_i \times 100\% \quad (26.5)$$

Y percent DM measured at the line rate can be compared directly to Y percent DM at 64 kbps. Y percent SES measured at the line rate can be converted directly to Y percent SES at 64 kbps, but if during the test a loss of frame alignment is detected (or a slip), this time as a percentage should be added.

The conversion for errored seconds is more complicated. Since the higher multiplexed bit rate contains many 64 kbps channels, we need to know how many errors are contained in each errored second at the higher rate in order to estimate how many 64 kbps channels have been errored.

Assuming errors are distributed evenly within the frame (the worst-case condition), then Equation 26.5 is representative.

The validity of all these conversion algorithms has been debated in standards committees. Bursts of errors in a high-rate transmission stream may cause loss of synchronization in a subsequent demultiplexer, creating an extended error burst on a tributary output. Some practical measurements seem to support these conclusions.

The reader should refer to the latest version of ITU-T Recommendation G.821 for guidance on conversion algorithms. The more recent G.826 recommendation also should be consulted because it covers the performance of higher-speed digital paths.

### 26.3.4 Related error performance standards

Several other error performance standards exist, but these generally are related to the ITU-T G.821 recommendation.

**ITU-R Recommendation 594-1** Error performance of digital microwave radio systems is characterized using ITU-R (formerly CCIR) recommendations, notably Recommendation 594-1, for 2500 km link (64 kbps unidirectional channel) as follows:

1. BER worse than  $1 \times 10^{-6}$  for less than 0.4 percent of any month.
2. BER worse than  $1 \times 10^{-3}$  for less than 0.054 percent of any month.
3. Errored seconds should not exceed 0.32 percent of any month.
4. Residual BER should not exceed  $5 \times 10^{-9}$  (15-minute integration).

Recommendation 594-1 is compatible with the high-grade portion of G.821. An additional block allowance of 0.05 percent SES has been added for adverse propagation conditions and a specification for residual BER has been added. Both the residual BER (RBER) threshold and the percentage for G.821 parameters should be reduced in proportion for systems less than 2500 km long.

Since these recommendations are compatible with G.821, the same measuring instruments and calculation are used to assess performance. The main consideration is that radio propagation is affected by weather, so results could be misleading unless measured over a reasonable period (such as a month).

**Recommendation G.826** This recommendation (issued at the end of 1993 with subsequent updates) focuses on the error performance parameters and objectives for international constant bit rate (CBR) digital paths at or above the primary rate (i.e., 1.544/2.048 Mbps and above).

It is complementary to G.821; a higher-speed transmission system meeting G.826 should in most cases meet the G.821 requirements for the 64 kbps path carried therein. G.826 applies to PDH and SDH transmission systems and also to ATM transmission paths supported by PDH/SDH framing. The overall end-to-end ATM performance objectives are defined in ITU-T Recommendation I.356.

The measurement definitions in G.826 are applicable to both in-service and out-of-service testing, and are based on the concept of an errored block of data bits such as a frame. The block check for errors could be made by a test set or could be derived from in-service monitoring of CRC or parity errors, described as Error Detection Codes (EDCs).

The error events are defined as follows:

- *Error Block (EB)*: A block in which one or more bits are in error.
- *Errored Seconds (ES)*: A one-second period with one or more block errors.
- *Severely Errored Seconds (SES)*: A one-second period that contains 30 percent errored blocks or at least one Severely Disturbed Period (SDP).

An SDP occurs when, over a period of time equivalent to four contiguous blocks or one ms, whichever is longer, either all contiguous blocks are affected by a high error density  $10^2$ , or a Defect Event (DE) such as a loss of signal, loss of frame synchronization, etc., occurs.

The performance parameters are defined as follows:

- *Errored Second Ratio (ESR)*: The ratio of ES to total seconds of available time during a fixed measurement interval.
- *Severely Errored Second Ratio (SESR)*: The ratio of SES to total available time.
- *Background Block Error Ratio (BBER)*: The ratio of errored blocks to total blocks during a fixed measurement interval, excluding all blocks during SES and unavailable time.

**TABLE 26.5 End-to-End Error Performance Objectives for a 27,500 km Digital Hypothetical Reference Path.**

Rate Mbps	1.5 to 5	>5 to 15	>15 to 55	>55 to 160	>160 to 3500
Bits/block	2000–8000	2000–8000	4000–20,000	6000–20,000	15,000–30,000
ESR	0.04	0.05	0.08	0.16	Note 1
SESR	$2 \times 10^{-3}$	$2 \times 10^{-3}$	$2 \times 10^{-3}$	$2 \times 10^{-3}$	$2 \times 10^{-3}$
BBER	$2 \times 10^{-4}$	$2 \times 10^{-4}$	$2 \times 10^{-4}$	$2 \times 10^{-4}$	$10^{-4}$

Note 1: Under review, proposed 0.16 up to STM-4

**TABLE 26.6 Apportionment for Digital Sections in ITU-T Recommendation G.921.**

Section Quality Classification	HRDS Length (km)	Allocation	To Be Used in Circuit Classification
1	280	0.45%	High-grade
2	280	2%	Medium-grade
3	50	2%	Medium-grade
4	50	5%	Medium-grade

Based on these parameters, the G.826 performance criteria as a function of bit rate are shown in Table 26.5. As with recommendation G.821, the objectives are apportioned to different parts of the end-to-end network.

ITU-T Recommendation G.921 defines the performance of a *Hypothetical Reference Digital Section* (HRDS) and is based on the requirements of G.821. G.921 considers digital sections of 280 km (or multiples of 280 km) and assigns percentage allocation of overall G.821 specifications. Shorter, medium-grade connections also are defined as shown in Table 26.6.

**Recommendation M.2100** ITU-T recommendation M.2100 is titled “Performance limits for bringing-into-service and maintenance of digital paths, sections and transmission sections.” Error, timing, and availability performance are considered. A method for deriving ES and SES from in-service measurements is given for all hierarchical levels. M.2100 defines practical performance criteria for digital circuits, measured over shorter periods than the 1 month defined in G.821. Periods of 15 minutes, 24 hours, and 7 days are recommended. It also recommends margins for aging so that maintenance intervals can be extended.

Furthermore, M.2100 defines Anomaly Events (AE) and Defect Events (DE) for both in-service and out-of-service testing, and indicates the number of events permissible in a measurement period. For in-service testing, it suggests how events should be interpreted in terms of G.821 parameters. This parallels the North American ANSI Ti-231-1993 “Digital Hierarchy – Layer I In-Service Digital Transmission Performance Monitoring” standard.



### 26.3.5 In-service measurements

In-service measurements have become increasingly important as they allow long-term performance monitoring and preventive maintenance without interrupting customer traffic. This is desirable in the deregulated competitive environment because it allows the operator to continuously monitor quality of service and determine whether degraded performance is being caused within the operator's network or by another vendor.

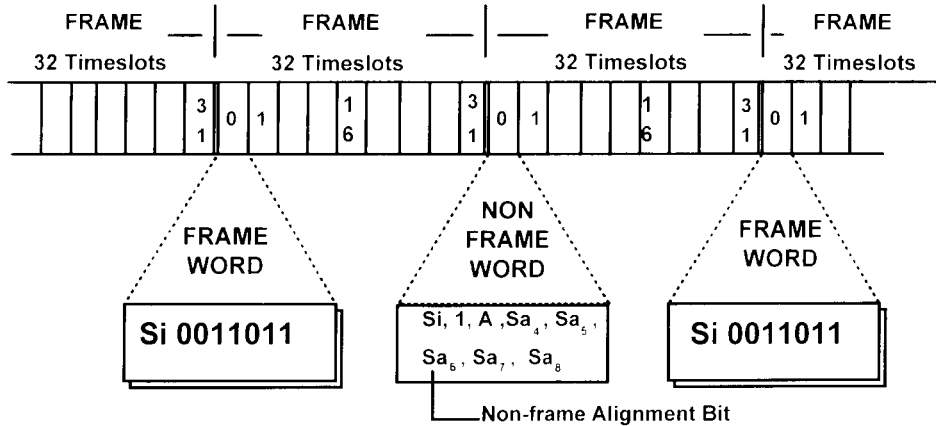
Often the customer's traffic will traverse several different networks between source and destination. The different parts of the route are referred to as *network sections*, while the end-to-end connection is referred to as the *path*. Monitoring overall path performance provides an indication of the service the customer receives, while section measurements are important for troubleshooting and preventing finger-pointing.

In-service measurements cannot rely on a bit-by-bit error check as is possible with an out-of-service PRBS test, because the monitoring equipment has no way of knowing the data content of random customer traffic. Instead, in-service testing has to check for errors in any known fixed patterns such as frame words in the random data stream, or apply error detection codes (EDCs) to blocks of data. The most powerful detection processes are based on computing parity or checksums on blocks of data (including the payload bits) and transmitting the result to the far end for recomputation and comparison. Several available methods for in-service error detection can be built into the network equipment itself, or can be implemented in test equipment or network monitoring.

**Bipolar violations (BPVs) and line code violations.** As explained in Chapter 3, telecommunications systems use standard interface codes to guarantee clock recovery, and line transmission systems also use coding to make efficient use of media. These simple binary and ternary codes have non-allowed states. If one of these illegal states is detected then an error in the data stream also will be likely. This detection process is limited, however, and binary errors easily can slip through as legitimate codes. Also, the detection applies only to a single transmission section as equipment will not retransmit code violations to subsequent sections. For this reason it cannot be used to monitor overall path performance.

**Frame alignment signal (FAS) errors.** These are detected by checking the bits in the repetitive frame alignment word to be found, for example, in timeslot zero (TS0) in alternate frames at the 2 Mbps E1 primary rate (ITU-T G.704) as shown in Figure 26.15. Since the frame alignment word forms only a small part of the overall frame and does not check any of the payload area, it provides only a sample or snapshot of error performance. It can be useful for estimating long-term average BER, assuming that errors are evenly distributed in the frame (Poisson distribution); this assumption may be invalid, however, since some error events occur in clusters.

In PDH systems, checking for errors in the frame alignment signal is the most convenient way of making in-service measurement of error performance. Each level in the hierarchy has a frame structure and associated frame alignment signals as shown in Figure 26.16. These can be checked by a test set or by the operational equipment.



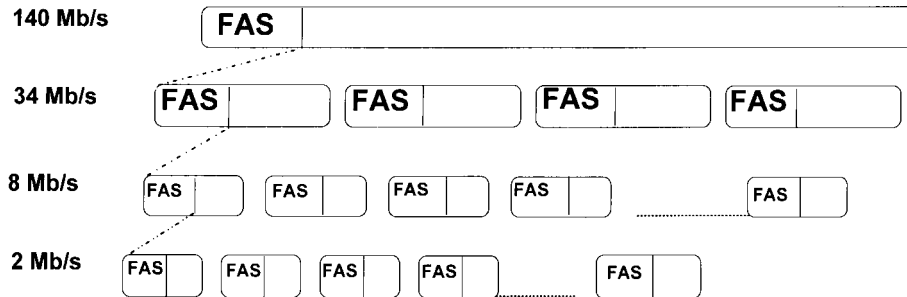
**Si** : Reserved for INTERNATIONAL    **A** : Remote Alarm

**Sa<sub>N</sub>** : Reserved for NATIONAL routes ("1" for INTERNATIONAL routes)

**LOSS of frame** : 3 bad frame words in 4

**Frame RECOVERY**: 2 good frame words + 1 Non-frame word

**Figure 26.15** The frame structure for 2.048 Mbps (E1) primary rate frame according to ITU-T recommendation G.704, showing the frame alignment signal (FAS) in alternate frames in timeslot 0 (TS0).



**Figure 26.16** Frame alignment signals are present at each hierarchical level in PDH. Each can be accessed by a digital transmission test set equipped with full framing and demultiplex capability.

Frame error checking has limitations and cannot guarantee to detect 90 percent of errors as required by the in-service monitoring recommendation of G.826. Because the payload is not checked, errored blocks can pass through undetected. A further problem for in-service monitoring occurs if the 2 Mbps signal is unframed. This is unlikely in the main telephone network, but a characteristic of 2 Mbps service is the ability to provide a clear channel so that any kind of random 2 Mbps signal can be transmitted error-free. For digital leased-line applications this could be a video signal, LAN interconnection, or other packetized data without G.704 framing. The network operator then has no way of checking the quality of the 2 Mbps service.

The limited in-service capability of PDH is one of the reasons for development of next generation SONET and SDH systems. In SDH, the PDH tributary is placed in a *virtual container* (VC) that includes a full set of monitoring, called the *path overhead*, so we no longer need to rely on the PDH frame structure for in-service testing and can check service quality even for unstructured data streams. This is covered in more detail in Chapter 13 and 30 on SONET and SDH.

**Parity errors.** Parity errors are detected by computing odd or even parity for blocks of data and then comparing the transmitted parity bits with the recalculated parity at the receiving end. Any discrepancy indicates one or more errors in the data block. There is a slight chance that the errors will occur in the parity bits themselves rather than in the payload bits; for reasonable block lengths, however, this effect can be neglected.

A bigger problem with simple parity schemes is that they can be fooled by multiple errors in a block that cancel the parity error detection at the receiving end. Parity error detection is reliable for lower error rates (e.g., less than  $10^{-4}$ ). Parity checking is used in the North American DS3 standard (44.736 Mbps), in some transmission line codes in optical and radio systems, and in the SONET/SDH standards where the BIP-8 parity check meets the G.826 requirement of detecting more than 90 percent of errored blocks.

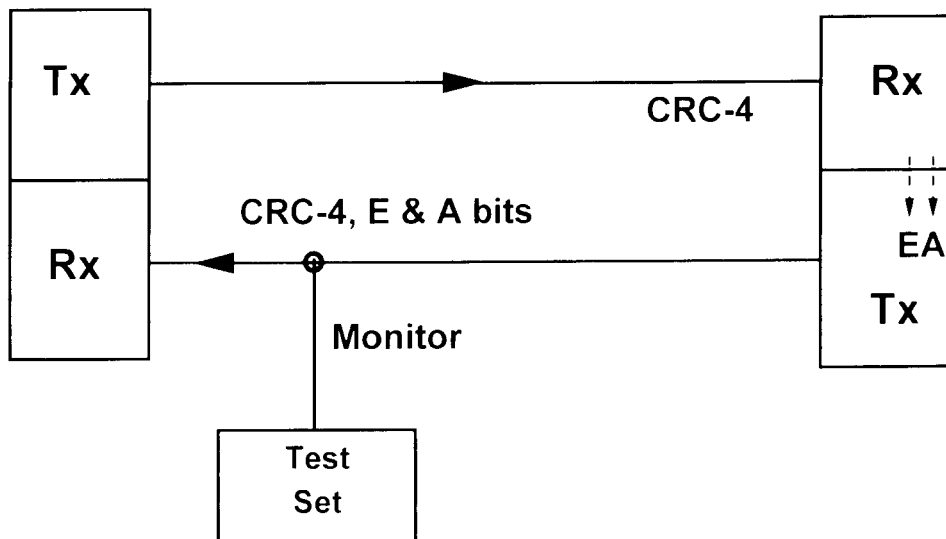
**Cyclic redundancy checksum (CRC).** Cyclic redundancy checksum is computed on blocks of data including payload bits and a CRC remainder sent to the receive end for comparison with the recalculated remainder. A discrepancy indicates one or more errors have occurred in the data block. CRC checking has become standard for G.704 framed primary rate signals at 1.544 Mbps (DS1) and 2.048 Mbps (E1). In North America the earlier non-CRC standard is referred to as *D4 framing* and uses a 12-frame multiframe for establishing frame alignment. The later CRC-6 frame incorporating the error detection is referred to as *Extended Superframe* (ESF) and uses a 24-frame multiframe (4632 bits) for establishing frame alignment.

The international E1 frame structure at 2 Mbps retains the 16-frame multiframe but divides it into two sub-multiframes of 8 frames (2048 bits) for calculating a CRC-4 remainder that is transmitted in the unused bit one position of the TS0 frame alignment word as shown in Figure 26.17. Another facet is the ability to transmit remote error detection back to the sending end. When an errored block is detected at the receiving end, the E bit is set in the return path. This is termed a *Far End Block Error* (FEBE) or *Remote End Block Error* (REBE). By checking CRC-4, E bits (FEBE), and A bits (Alarms), an indication of both go and return paths is possible as shown in Figure 26.18.

CRC error detection is reliable, with around 94 percent of errored blocks being detected by CRC-4 and 98.5 percent with CRC-6, thus meeting the requirements of G.826. CRC error detection also is used in the new generation of ATM cell transmission standards.

	Sub-multiframe (SMF)	Frame number	Bits 1 to 8 of the frame (TS0)								TS 1-31
			1	2	3	4	5	6	7	8	9-256
			Multi-frame	I	0	C1	0	0	1	1	0
1	0	1			A	S	S	0	S	S	
2	C2	0			0	1	1	0	1	1	
3	0	1			A	S	S	S	S	S	
4	C3	0			0	1	1	0	1	1	
5	1	1			A	S	S	S	S	S	
6	C4	0			0	1	1	0	1	1	
7	0	1		A	S	S	S	S	S		
II	8	C1		0	0	1	1	0	1	1	
	9	1		1	A	S	S	S	S	S	
	10	C2		0	0	1	1	0	1	1	
	11	1		1	A	S	S	S	S	S	
	12	C3		0	0	1	1	0	1	1	
	13	E		1	A	S	S	S	S	S	
	14	C4		0	0	1	1	0	1	1	
	15	E	1	A	S	S	S	S	S		

**Figure 26.17** The CRC-4 multiframe structure for 2.048 Mbps (E1) primary rate, showing the contents of successive timeslot 0 (TS0) positions. The main frame alignment signal in alternate frames remains the same as the earlier G.704 standard shown in Figure 26.16, but the reserved first bit now is used to carry the four bits of the CRC-4 remainder for the previous sub-multiframe. The E bits provide the far end block error (FEBE) indication so that the performance of the transmit path can be monitored from the sending end. The A bits provide a remote alarm indication.



**Indicates performance of both go and return paths**

**Figure 26.18** A test set or network management system connected to one direction of a 2 Mbps path operating CRC framing gets an indication of the performance of both directions of transmission by means of CRC-4, E, and A bits.

**TABLE 26.7 In-Service Anomaly Events (ITU-T M.2100) and Performance Primitives (ANSI T1M1.3).**

Interface/Bit Rate	Section/Line	Path
T1, DS1 15.44 Mbps	BPVs (B8ZS, AMI) FAS error Line-code error	CRC-6 error G1–G6 bits (ESF data link)
E1 2.048 Mbps	BPVs (HDB3) FAS error Line-code error	CRC-4 error FEBE (E bits)
DS3 44.736 Mbps	BPVs (B3ZS) FAS error Line-code error M13 parity errors (P bits)	C-bit parity error FEBE bits
E3/E4 34.368/139.264 Mbps	BPVs (HDB3) FAS error Line-code error	
SONET/SDH 51.84 Mbps to 9953.28 Mbps	BIP-8 (Byte B1) (section overhead) BIP-8 (Byte B2) (line overhead)	BIP-8 (Byte B3) (path overhead) FEBE Payload path monitor (e.g., DS1 and DS3)

According to ITU-T recommendation M.2100, system errors are classified as *Anomaly Events* (AE) such as frame errors, parity errors, CRC errors; and *Defect Events* (DE) such as loss of signal, loss of frame synchronization, and so on. Table 26.7 shows the anomaly events detection processes available at the various telecommunication standard bit rates.

### 26.3.6 Data logging

Since error performance measurements might need to run for several hours or days to accumulate statistically significant results, most of the time the test sets will be left unattended. Hence the error detector must have a means to log measurement data and error events/alarms for later analysis. Furthermore, long-term tests may be affected by power supply interruptions. Data logging protects valuable test data that otherwise would be lost. When power is reinstated, the instrument can be designed to recover from the interruption automatically and to recommence the test without operator involvement.

Data logging usually is provided by outputting results and events to a printer as they happen, or by storing information in nonvolatile memory for later analysis and display. An example printout is shown in Figure 26.19.

Long-term tests can produce a large amount of measurement data, which is time-consuming to analyze. Graphic display of results, as shown in the example in Figure 26.20, helps the user quickly identify the periods of interest. Note that the results are timestamped, allowing error/alarm events to be correlated with system malfunctions.

Hewlett Packard nPS111C Instrument Configuration					
<b>RECEIVER</b>					
Receive Signal	: STM-1e				
Payload	: 140 Mb/s				
Payload Type	: UNFRAMED				
Pattern	: 2 <sup>23</sup> -1	Polarity	: INVERTED		
Error Count	BIT (test)	CODE	CRC	REBE	
	23	N/A	N/A	N/A	
Error Ratio	1.161E-09	N/A	N/A	N/A	
<b>Analysis Results :</b>					
<b>G.826 ANALYSIS</b>					
	RS B1	MS B2	MS	PATH	
	BIP	BIP	FEBE	B3 BIP	
Errored Blocks	44011	46378	0	10239	
Errored Sec	7	15	0	17	
Severely Errored Sec	5	15	0	9	
Unavailable Sec	0	0	0	0	
Path Unavailable Sec	N/A	0	0	0	
Background Block Errors	4011	0	0	10236	
Errored Second Ratio	4.930E-02	1.056E-01	0	1.197E-01	
Severely Errored Sec Ratio	3.521E-02	1.056E-01	0	6.338E-02	
Background Block Err Ratio	3.660E-03	0	0	9.620E-03	
<b>G.821 ANALYSIS</b>					
	BIT (test)	FAS 140M	FAS 34M	FAS 8M	FAS 2M
Errored Sec	19	N/A	N/A	N/A	N/A
%Errored Sec	13.38028	N/A	N/A	N/A	N/A
%ES (Annex D)	0.00744	N/A	N/A	N/A	N/A
Error Free Sec	123	N/A	N/A	N/A	N/A
%Error Free Sec	86.61972	N/A	N/A	N/A	N/A
Severely Err Sec	9	N/A	N/A	N/A	N/A
%Severely Err Sec	6.33803	N/A	N/A	N/A	N/A
Degraded Minutes	0	N/A	N/A	N/A	N/A
%Degraded Minutes	0.00000	N/A	N/A	N/A	N/A
Unavailable Sec	0	N/A	N/A	N/A	N/A
%Unavailable Sec	0.00000	N/A	N/A	N/A	N/A
<b>M.2100 ANALYSIS</b>					
	Rx 140Mb/s	Tx	Rx 34Mb/s	Tx	
Errored Seconds	19	N/A	N/A	N/A	
Severely Errored Seconds	9	N/A	N/A	N/A	
Unavailable Seconds	0	N/A	N/A	N/A	

Figure 26.19 The modern BER test set generates a large amount of information on errored events and calculates error performance measurements. This sample printout shows the type of data available.

## 26.4 Bit Error Rate (BER) Instrument Architecture

A bit error rate tester (BERT) consists of a pattern generator and error detector, often combined in a single instrument but sometimes separate. Applications of BER testers can be divided into two broad categories:

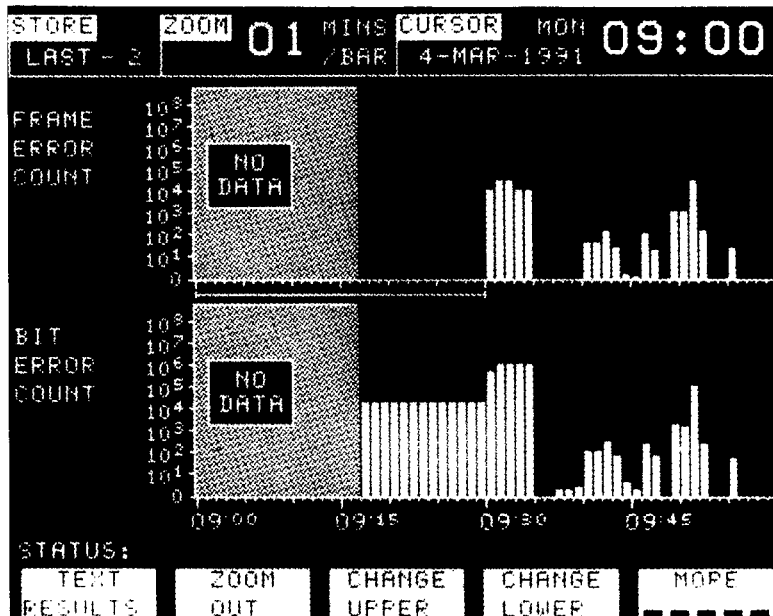
- Telecommunications testing
- General laboratory/production testing of devices and systems

In both cases the measurements and modes of operation are similar; the main difference is in the electrical interfaces required, and the need to provide structured or framed data streams.

If the BERT is used for testing telecommunications equipment and systems either in the field or in the factory, it must have coded interfaces according to the CCITT Recommendation G.703 (or the equivalent North American ANSI T1 standards). As mentioned previously, the G.703 interface specification defines digital codes (such as HDB3, B3ZS, and CMI), which ensure that a minimum number of transitions occur in the data stream even with long runs of 1s or 0s. This characteristic ensures that a timing clock can be recovered reliably in the receiving equipment.

A single output or input connection is all that is required when connecting the test set, and the user need not be concerned about the relative phasing of parallel clock and data signals. G.703 also defines other important interface characteristics such as specific bit rates (see Table 26.1), voltage levels, pulse shape, impedance, frame structure, and so on. Put simply, if the test set complies with the relevant interface standard, it is ready to connect to telecommunications equipment. This also is true for the new generation of synchronous optical network (SONET) equipment for which standards exist up to 2.488 Gbps.

For more general use, conventional clock and binary NRZ (nonreturn to zero) data are necessary. With higher-speed application, clock/data phasing becomes critical, so the test set should have a means of adjusting the clock timing for optimal



**Figure 26.20** An additional feature on some BER test sets is the graphic display of results. In this example, a histogram display shows how bit error count and frame error count are varying with time.

sampling of the eye diagram. To explain some of these concepts, two examples of BER testers follow.

#### 26.4.1 Telecommunications BER tester with coded interfaces

A pattern generator and an error detector for telecommunications applications are shown in the block diagrams in Figures 26.21 and 26.22. This serial implementation is suitable for bit rates up to around 200 Mbps. (Typical maximum rates are 44.736 Mbps (DS3) in North America, and 139.264 Mbps outside North America.)

In Figure 26.21, the PRBS and word generator circuitry is clocked either from a fixed-frequency clock source (to G.703), or from a synthesizer to provide a variable clock rate. Most telecommunications applications require a few specific clock frequencies and the ability to provide small offsets to  $\pm 15$  to  $\pm 50$  ppm. The PRBS and word generator circuit usually provides a trigger pulse to signify repetition of the pattern.

The output from this circuit drives either the binary data output amplifier, which provides DATA and  $\overline{\text{DATA}}$  with an accompanying clock signal, or the coded data output circuitry. This may add framing to the signal (which is a fairly standard requirement in North America but not essential elsewhere, although all SONET/SDH systems require framing). It then adds the appropriate interface code for effective clock recovery. The output amplifier provides a signal conforming to the electrical interface specification, which may require a pseudo-ternary, alternate-mark inversion signal. This means that alternate marks or 1s are encoded as positive and negative pulses in order to provide a near-zero average dc level.

Errors can be added to the pattern by an exclusive-OR gate that is controlled by single-shot or repetitive pulses from the clock generator. The decade divider sets a BER of  $10^{-N}$ .

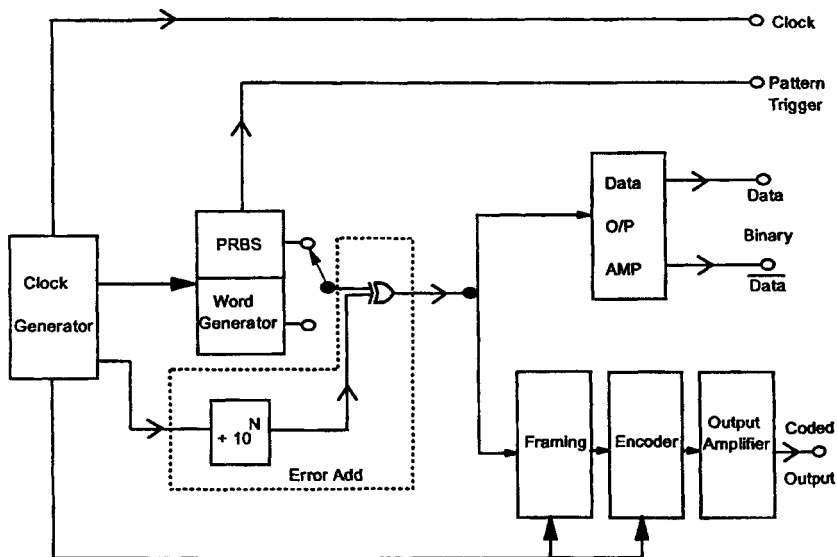


Figure 26.21 Pattern generator block diagram for a telecommunications BER test set.



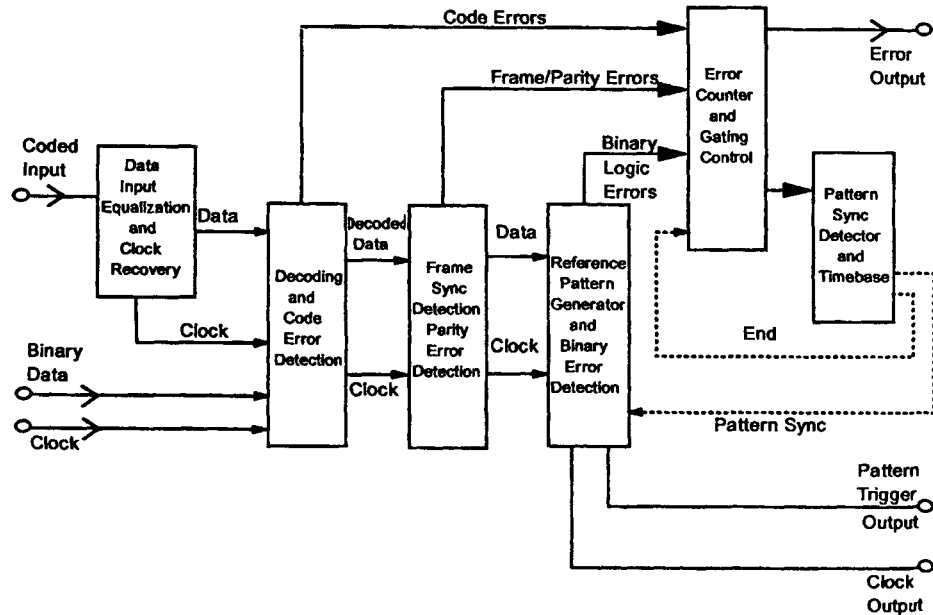


Figure 26.22 Error detector block diagram for a telecommunications BER test set.

The companion error detector, shown in Figure 26.22, receives the standard coded-interface signal, recovers the clock, and strips off the coding to provide binary data and clock signals. In the process, it detects any violations of the interface code algorithm and sends signals to the error counter. This provides the first level of in-service error detection.

For an instrument equipped to deal with framed signals, the receiver then locks onto any framing present, checks for frame errors, and decodes any embedded alarm signals, parity, or CRC bits, thus providing a further level of in-service measurement.

Finally, the binary data and clock are fed to the error detector and reference pattern generator (as described in section 26.3.2), which checks the received pattern bit by bit for logic errors. A time base controls the measurement gating for single-shot, repetitive, and manual gating. The error counts accumulated are processed to provide BER and error performance analysis (see section 26.3.3).

#### 26.4.2 High-speed pattern generator and error detector

Figures 26.23 and 26.24 show the architecture for a 3 Gbps pattern generator and error detector. Because of the high bit rate, it is not practical to implement the PRBS and word generation directly in serial form. Instead, the patterns are generated as parallel 16-bit words at a maximum rate of 200 Mbps, where shift registers and high-capacity memory can be implemented using bipolar technology. (The high-speed circuitry normally would use gallium arsenide ICs.) The high-speed multiplexer (using a pyramid of SPDT switches shown in Figure 26.25) converts the parallel data to a serial stream at rates up to 3 Gbps.

618 Network Test Instrumentation

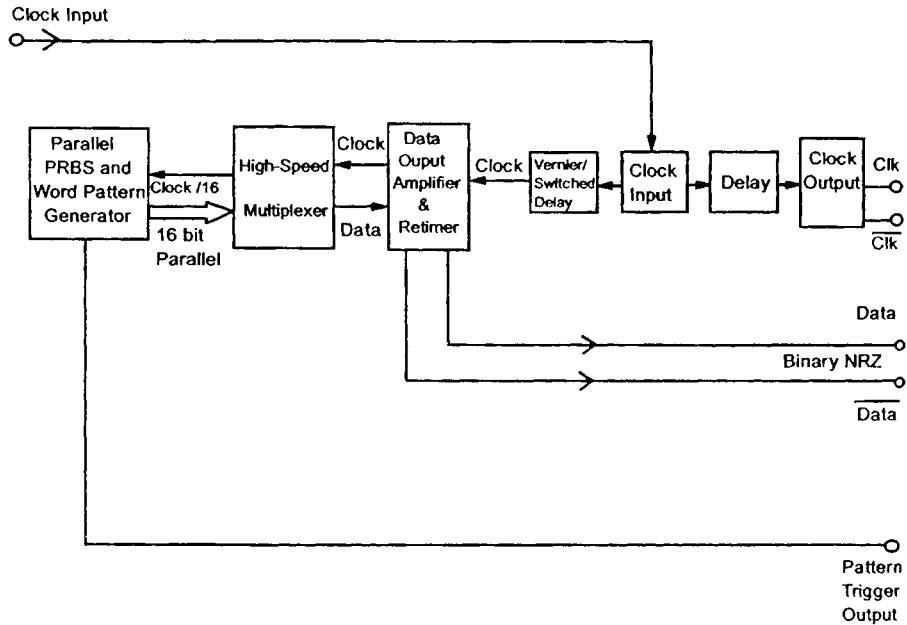


Figure 26.23 High-speed pattern generator block diagram.

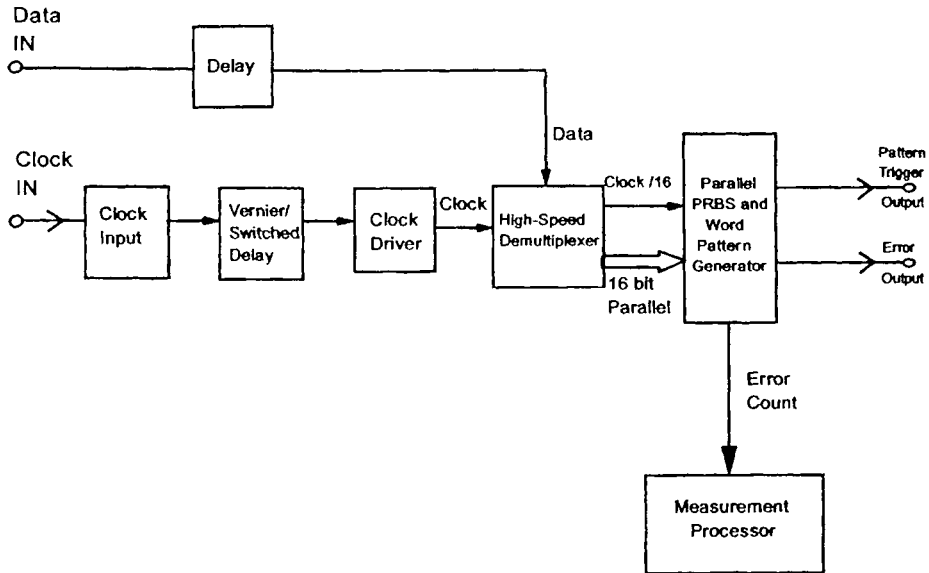
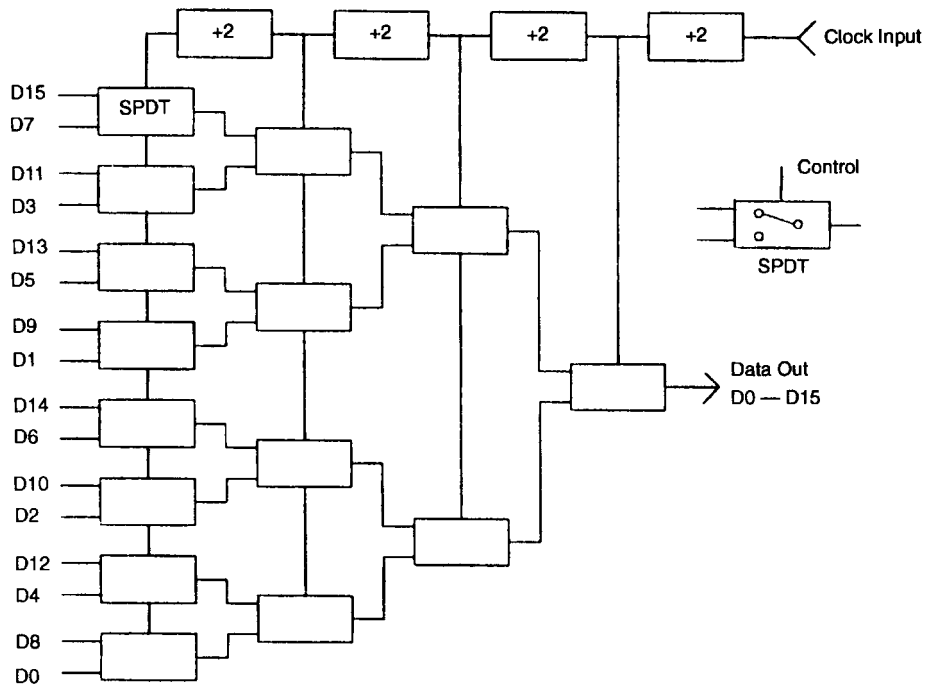


Figure 26.24 High-speed error detector block diagram.



**Figure 26.25** Pyramidal multiplexer configuration. Here, 16-bit data from the parallel gate array are fed to the first eight switches. The output of these switches is fed to the next four switches and so on until a single serial data stream is generated. The control of switch positions for each level is derived by dividing the clock rate by 2 at each stage.

The clock input is generated by a frequency synthesizer. The clock output amplifier is driven through a fixed delay line, and the pattern generation and output amplifier are clocked through switched and vernier delay circuitry so that clock/data phase can be varied both positively and negatively. The switched delays are 250, 500, and 1000 ps, while the vernier provides 0 to 250 ps in 1 ps increments.

The retimer associated with the output amplifier reclocks the data through a D-type flip-flop to maintain minimum jitter. Since this type of test set normally would be used in a laboratory environment, the clock and data output levels and dc offsets can be varied for the particular application.

The companion error detector, shown in Figure 26.24, uses a similar parallel implementation. Clock and data inputs pass through fixed and vernier delays so that any arbitrary clock/data phase can be adjusted for optimal sampling in the error detector. In fact, by adjusting the decision threshold and clock phasing under control of the internal processor, the error detector operating conditions can be optimized automatically. The high-speed demultiplexer converts the serial data stream to 16-bit parallel words along with a divide-by-16 clock signal. The parallel implementation reference pattern generator synchronizes with the incoming data and makes a bit-by-bit comparison. Any errors that occur are counted by one of two counters. One counter counts errors, while the other counter is being read in succession. The measurement processor provides error performance analysis down to 1 ms resolution.

## 26.5 Bit Error Rate Measuring Instrument Specifications

This section describes the major specifications of BER testers and the relevance for different applications, principally telecommunications testing and general-purpose device and system testing.

### 26.5.1 Patterns

The pattern generator and matching error detector will provide a range of PRBS patterns. The range might be quite limited for a dedicated tester aimed at maintenance of 1.544 Mbps (T1) or 2.048 Mbps (E1) links, or might be wide for general-purpose applications.

For international application, the patterns should conform to ITU-T Recommendations O.151, O.152, and O.153, as shown in Table 26.8. For North American applications, you also may find the QRSS (quasi-random signal source), which is a  $2^{20}-1$  PRBS, but with the maximum run of 0s limited to 15. If the tester uses standard PRBS patterns, the user can be almost certain it will interwork with other BER tester models, since the defining polynomials will be the same. When testing telecommunications links, one also can be sure of measuring them in accordance with the international standard in case of a dispute between vendor and customers.

For general laboratory use, a wider range of PRBS patterns may be useful. If not covered by the ITU-T recommendations, it is usual to quote the polynomial used. For example:

$$\begin{aligned}2^{11}-1 \quad D^{31} + D^{28} + 1 &= 0 \\2^{10}-1 \quad D^{10} + D^7 + 1 &= 0 \\2^7-1 \quad D^7 + D^6 + 1 &= 0\end{aligned}$$

As mentioned in section 26.3.1, for general-purpose applications it is useful to be able to modify the PRBS pattern to add arbitrarily long runs of 0s or to vary the mark density, i.e., the ratio of 1s to 0s.

**TABLE 26.8 Test Pattern versus Bit Rate (O.151, O.152, and O.153).**

Bit Rate	Test Pattern
50 bps to 168 kbps (0.153)	$2^9-1, 2^{11}-1, 2^{20}-1$
64 kbps (0.152)	$2^{11}-1$
1.544 Mbps	$2^{15}-1, 2^{20}-1$
2.048 Mbps	$2^{15}-1$
8.448 Mbps	$2^{14}-1$
34.368 Mbps	$2^{23}-1$
44.736 Mbps	$2^{15}-1, 2^{20}-1$
139.264 Mbps	$2^{23}-1$

User-programmable word patterns are becoming increasingly important as a means of simulating extreme data patterns found in live traffic and to simulate and capture frame structures. For maintenance of telecommunications networks, a few standard word patterns are all that is required. In North America it is now common to provide a range of word patterns in a T1 tester for 1.544 Mbps. These might include all 1s/0s, 1:1 pattern density, 1:7 density, 3-in-24, and 55-octet (8-bit) user words.

For general-purpose applications, particularly at high bit rates, very much larger pattern memories are common. For example, a 3 Gbps test set (HP 71603B) has a pattern memory of 4 Mbits so that complete frame structures for SONET/SDH systems and ATM systems can be simulated. With very large pattern memories, you should consider how they will be programmed, because it is impractical to program them by hand using the keyboard and display. Some test sets provide a diskette drive for downloading the word pattern. These disks can be programmed on a PC, and the test equipment manufacturer might well supply a disk with a range of standard patterns already stored.

When reviewing the large pattern memory capability, one should investigate how the memory is arranged. Very large memories will be arranged as parallel bitmaps, which may limit the resolution the pattern can have. In other words, the stored pattern must be an integer number of smallest steps in memory. Some frame structures may require the repetition of several frames in order to make up an integer number of steps. For example, one pattern generator offers 1-bit resolution for pattern lengths from 1 bit to 32 kbits, rising to 128-bit steps for pattern lengths of 2 to 4 Mbits. Using the 2- to 4-Mbit range, your pattern would have to be an exact multiple of 128 bits.

Finally, some pattern generators and error detectors allow switching between two stored patterns in order to investigate how a device under test responds to varying patterns. A gating input usually is provided to control the switching.

### 26.5.2 Clock and data inputs and outputs

The requirements in this area will depend on the intended application of the tester. For testing telecommunications equipment and systems, the main requirement is standard coded data interfaces. (Clock signals generally are not required.) For each of the standard telecommunications bit rates, the interface should conform to a recognized telecommunications standard such as ITU-T G.703, ANSI TI.403, or ANSI TI.102/TI.105.

The error detector also may have high-impedance inputs to allow bridging measurements, and the available sensitivity of the input should be able to handle levels of  $-20$  to  $-30$  dB below standard to cope with protected test points on equipment. For example, the input sensitivity might be quoted as  $+6$  to  $-30$  dB $_{dsx}$ , where  $dsx$  stands for digital crossconnect (the usual test reference point in a telephone exchange or central office). Unless specified otherwise, one can assume that the tester meets the standard jitter specifications at the interface (ITU-T Recommendations G.823, G.824, and G.825), since this is included in G.703.

For general applications, more flexible and detailed specifications are needed. In most applications, both data and clock will be required as generator and error detector.

**622 Network Test Instrumentation**

One should check that the inputs and outputs can interface with standard logic levels such as ICL or TTL. Whereas telecommunications equipment normally operates with 75- $\Omega$  impedances, high-speed general circuit applications will probably require 50- $\Omega$  connections. A general-purpose test set will probably have programmable output levels and offsets, e.g., 0.25 to 2 V p-p with 10 mV resolution. Both data and inverted-data outputs usually are provided.

In high-speed testers, check the jitter specification on the clock and data outputs. This usually is specified in picoseconds (e.g., 10 ps RMS) when transmitting a PRBS. Wave shape also is important and is defined by the transition times and overshoot. A poor wave shape from the pattern generator may limit the testing and hence the quality of development work.

One final point, sometimes overlooked in digital systems, is the return loss or impedance matching at the tester input and output. Compared to analog systems, return loss is not quite so important; anything less than 10 to 15 dB return loss, however, may cause problems with transmission-line reflections and degraded pulse shape. This is particularly important when testing very high-speed systems and components. The G.703 interface specification gives recommended return loss values for telecommunication systems.

**26.5.3 Framing**

This is mainly applicable to telecommunications equipment testing. As with coded-data interfaces, framing is governed by the appropriate standard such as ITU-T Recommendations G.704, G.706, G.742, and G.751. When framing is included in the tester, a wider range of measurements is possible, including in-service monitoring. In some cases with equipment in North America, framing at T1 (1.544 Mbps) and DS3 (44.736 Mbps) is mandatory, as it is with the new generation of SONET/SDH transmission equipment.

General-purpose testers with large pattern memories also can provide frame simulation. The extra flexibility provided by this approach is not required in dedicated applications.

**26.5.4 Error insertion**

This facility provided in the pattern generator allows single errors to be inserted in a PRBS or word pattern or a specified BER to be generated. This is a useful check of the error detector, but of course, errors in a pseudorandom data stream will not be detected by the equipment under test. If, however, the data stream is structured either by using the pattern memory or by built-in frame capability, then errors inserted in the frame will be detected by the equipment under test. The error detector should have appropriate measurements such as code and frame errors in addition to standard logic errors in data.

**26.5.5 Clock rates**

For telecommunications applications, a few standard bit rates are all that is required. Table 26.9 shows the most common rates and the required clock accuracy to meet the interface specifications.

**TABLE 26.9 Clock Tolerance at Hierarchical Interfaces.**

Clock Rate	Tolerance
PDH	
64 kbps	$\pm 100$ ppm
1.544 Mbps (DS1)	$\pm 50$ ppm
2.048 Mbps (E1)	$\pm 50$ ppm
8.448 Mbps	$\pm 30$ ppm
34.368 Mbps	$\pm 20$ ppm
44.736 Mbps (DS3)	$\pm 20$ ppm
139.264 Mbps	$\pm 15$ ppm
SONET/SDH	
51.84 Mbps (STS-1)	$< 4.6$ ppm*
155.52 Mbps (STS-3, STM-1)	$< 4.6$ ppm*
622.08 Mbps (STS-12, STM-3)	$< 4.6$ ppm*

\* Normally synchronized to the system clock to avoid excessive pointer movements.

For general applications, a synthesized clock source is more useful. Sometimes this is a built-in capability; sometimes an external synthesizer is required, in which case a clock input is specified on the pattern generator. The clock input should be compatible with the synthesizer output, which probably will be a sine wave. For best jitter performance, the synthesizer should have low single-sideband phase noise, e.g.,  $-140$  dBc/Hz.

### 26.5.6 Error measurements

All test sets will provide the basic error measurements such as error count and bit error ratio in a timed gating period. If the test set operates only with PRBS or word patterns, then the errors detected will be simply logic errors from bit-by-bit comparison with the reference pattern generator. If the tester has framed data capability, then errors also can be detected in frame words and in coding. The range or error types might include the following:

- Out-of-service measurements
  - Logic errors or bit errors (binary reference pattern comparison)
- In-service measurements
  - Frame errors
  - CRC-4 or CRC-6 code errors
  - Remote end block errors (REBE) or far end block errors (FEBE)
  - Parity errors
  - Interface code errors or bipolar violations

As mentioned in section 26.4.1, an error detector capable of operating with framed signals can derive one or more of the in-service measurements just listed. These might or might not be important in your application, but measurements on operational telecommunications equipment increasingly require this facility. Frame

**624 Network Test Instrumentation**

errors are detected by checking for any errored bits in the periodic frame alignment signal (FAS). Cyclic redundancy checksum (CRC-4 and CRC-6) is calculated on blocks of live data, and the remainder is transmitted as part of the frame structure.

At the receiving end, the error detector decodes these data and compares them with the locally calculated CRC remainder. A discrepancy indicates one or more errors in the data block. This powerful method of in-service error checking is becoming universally accepted internationally (CRC-4) and in North America (CRC-6). In North American DS3 systems and in the new generation of SONET equipment, a similar in-service check is provided by parity error checks on the stat bits in the frame. A further enhancement is REBE or FEBE, whereby block error information detected by CRC or parity is sent back to the transmitting end by means of the frame structure. Lastly, the interface code itself can be checked for errors.

Review the specification of the error detector for how many error types are available and for compliance with the international or North American standards.

Error analysis normally will be included based on ITU-T Recommendation G.821, described in section 26.3.3. This is based on the 1-second error-free interval. Some test sets also provide finer resolution of the interval down to 1 ms, which can be useful in research work and field trials. Some test sets incorporate G.821 Annex D additions and possibly analysis according to the new ITU-T Recommendation M.2100 and G.826 for maintenance of telecommunications links.

Lastly, look at the data-logging capabilities provided. Long-term tests result in a lot of data, so a means for storing them and displaying them graphically can be a valuable facility.



## Protocol Analyzers

**Stephen Witt**

*Hewlett-Packard Co., Colorado Springs, Colorado*

Two broad categories of products are used to implement and manage computer networks, those that test the transmission network and those that test the protocol information transferred over the transmission network. Testing the protocols is referred to as *protocol analysis*. It can be accomplished with several different types of descriptively named products including:

- Network management systems
- Distributed monitoring systems
- Protocol analyzers
- Handheld testers

*Network management systems* are comprehensive, integrated network-wide systems for managing and administrating systems and networks. Protocol analysis is one of many applications network management systems perform. Network troubleshooting is performed by acquiring network data from devices on the network and from instrument probes distributed through the network.

*Distributed monitoring systems* are performance-monitoring and troubleshooting applications that are implemented with instrument probes or protocol analyzers that are distributed throughout the network. The probes and analyzers are controlled with a management application running on a workstation or PC.

*Protocol analyzers* are specialized instruments dedicated to protocol analysis. Protocol analyzers are used to troubleshoot network problems and to monitor network performance.

*Handheld* testers are special-purpose tools that are small, lightweight, and usually battery operated. They perform a variety of measurements such as continuity tests, transmission tests, and simple protocol analysis measurements such as simple statistics and connectivity tests.

## 27.1 Protocol Analyzer Products

The term *protocol analyzer* typically is used to describe a class of instruments that are dedicated to performing protocol analysis. Protocol analyzers are implemented in one of three ways:

- Portable protocol analyzers
- Embedded protocol analyzers
- High-end protocol test sets

Each of these products has unique advantages and disadvantages that make it the appropriate choice for specific applications (Table 27.1).

### 27.1.1 Dispatched portable protocol analyzers

The term protocol analyzer most commonly refers to the portable instruments that are dispatched to trouble sites on the network, i.e., the critical links or segments that are experiencing problems. The emphasis in these products is portability: products that are lightweight, rugged, and include the maximum amount of functionality.

Network troubleshooters require a product that connects to any point in an inter-network regardless of the network's physical implementation. Most popular protocol analyzers therefore are able to accommodate multiple network interface modules. To avoid multiple trips back to the office to get additional equipment, network troubleshooters require a "one-handle solution," a product that integrates as much test capability as possible into one portable product.

Portable protocol analyzers most commonly are used for network troubleshooting in installation and maintenance applications because they focus on installing networks and troubleshooting network problems. Installing networks requires the network engineer to stress the network or network devices using scripts that emulate setting up logical connections, placing calls, and creating high-traffic scenarios. Troubleshooting network problems requires that the network engineer have extensive protocol decodes available to examine whatever frames are present on the network. Protocol statistics and expert analysis are used to identify network errors.


Figure 27.1 shows a typical dispatched portable protocol analyzer that is lightweight, rugged, integrates a standard PC, and will accommodate modules for different network interfaces.

### 27.1.2 Embedded protocol analyzers

The flexibility and portability of a portable protocol analyzer makes it more expensive than protocol analyzer solutions that are implemented as embedded systems. Embedded protocol analyzers are implemented as software applications that run in a computer platform such as a Unix workstation or a desktop PC. The line interface is provided by a special-purpose analysis and acquisition system or a commercially available standard network interface card. Embedded protocol analyzers, like portable protocol analyzers, are used to test established networking technologies.

**TABLE 27.1 Protocol Analyzer Products.**

Product	Description	Application	Advantages	Disadvantages
Dispatched portable protocol analyzers	<ul style="list-style-type: none"> <li>• Integrated, standalone unit</li> <li>• Lightweight, portable, rugged package</li> <li>• Built upon an integrated PC system</li> <li>• Provides “one handle” solution</li> <li>• Supports multiple network interfaces</li> </ul>	<ul style="list-style-type: none"> <li>• Troubleshooting</li> <li>• Installation</li> <li>• Performance monitoring</li> </ul>	<ul style="list-style-type: none"> <li>• Good price performance</li> <li>• Portability</li> <li>• Accommodates many different network interfaces</li> <li>• Turnkey solution—hardware and software configured at the factory</li> <li>• Suitable for simple R&amp;D applications</li> </ul>	<ul style="list-style-type: none"> <li>• Limited multi-port capacity due to size of package</li> <li>• Limited capability for R&amp;D applications</li> </ul>
Embedded protocol analyzers	<ul style="list-style-type: none"> <li>• “Cards and software”: line interface hardware and application software for each network interface</li> <li>• Built on standard computing platform: PC or Unix workstation</li> </ul>	<ul style="list-style-type: none"> <li>• Troubleshooting</li> <li>• Performance monitoring</li> </ul>	<ul style="list-style-type: none"> <li>• Lower relative cost</li> <li>• Takes advantage of an existing PC or workstation</li> <li>• Additional cards and software can be added for additional network interfaces</li> </ul>	<ul style="list-style-type: none"> <li>• Limited to a specific physical site unless built into portable PC</li> <li>• Requires user to set up and configure hardware and software</li> </ul>
High-end protocol test sets	<ul style="list-style-type: none"> <li>• Built on standard computing platform: PC or Unix workstation</li> <li>• Standard card cage allows user to configure the product to the application</li> </ul>	<ul style="list-style-type: none"> <li>• Conformance testing</li> <li>• Stress testing</li> <li>• Traffic generation</li> <li>• R&amp;D</li> </ul>	<ul style="list-style-type: none"> <li>• Sophisticated measurement set</li> <li>• Programmable measurements</li> <li>• Modular to allow upgrade of test capability</li> <li>• Many software applications</li> <li>• Supports new, emerging network technologies</li> </ul>	<ul style="list-style-type: none"> <li>• Higher relative cost</li> <li>• Limited portability</li> <li>• Complex capabilities are typically more difficult to use</li> </ul>



**Figure 27.1** A portable protocol analyzer.

The functionality of an embedded protocol analyzer usually is focused on performance-monitoring applications, making use of the statistical measurement capability.

Because embedded protocol analyzers take advantage of an existing computer platform, they usually are the least expensive alternative for protocol analysis. But lower cost comes with some disadvantages, such as requiring the user to install and configure the hardware and software. They also do not provide a portable, one-handle solution with multiple line interfaces. A network manager usually uses these products to monitor a specific critical segment or link, however, so the fact that these solutions usually are dedicated to a single network interface is an acceptable tradeoff.

Figure 27.2 shows a set of PC cards that are used with a software application to implement an embedded protocol analyzer.

### 27.1.3 High-end test sets


High-end protocol test sets are used in development applications for newly emerging network technologies. R&D engineers developing network equipment, designers of new networks, and engineers who are commissioning new equipment all require their protocol analysis tools to have power, performance, and flexibility. The testing functions in a high-end test set therefore emphasize traffic generation, simulation, emulation, stimulus/response testing, and programming capability.

As new network technologies are deployed, this class of product is used for installation and maintenance until portable or embedded solutions become available. Products in this category are more expensive than portable protocol analyzers and embedded protocol analyzers used for installation and maintenance applications. High-end test sets usually are offered in a card cage, such as VXI, that allows the user to configure the exact network interfaces and measurements required for the application, and to analyze multiple network connections simultaneously (multiport analysis).

Figure 27.3 shows a typical high-end protocol test set offering a Unix-based computing platform with slide-in modules for the line interfaces and analysis and acquisition systems.

## 27.2 Protocol Analyzer Implementations

Because there is a variety of requirements for protocol analysis, there are many different product implementations. Many of the differences between products represent



**Figure 27.2** An embedded protocol analyzer.



**Figure 27.3** A high-end protocol test set.

nothing more than the preferences or innovations offered by individual vendors. But there are some fundamental tradeoffs in the implementation of a protocol analyzer that have significant impact on the feature set and cost of the product. These tradeoffs include:

- Real-time analysis vs. post-processing
- Passive testing vs. intrusive testing
- Dedicated hardware acquisition system vs. software acquisition system
- Single-port vs. multiport

### 27.2.1 Real-time analysis vs. post-processing

Network troubleshooting can be performed in real time as problems occur, or it can be done in a post-process mode by capturing a buffer of data for subsequent analysis.

**Real-time analysis.** The most intuitive way to use a protocol analyzer is in real time, making measurements on the live network as the traffic occurs. Extra processing power is required to keep up with real-time network events. Real-time analysis is performed to troubleshoot network failures as they occur. The network engineer will interact with the protocol analyzer to gather information, query the network, and solve the problem. Real-time analysis also is used to gather performance-monitoring information over long periods of time, which typically would overrun a capture buffer.

**Post-process analysis.** In post-process mode, data is captured and analyzed offline. This allows the user to closely scrutinize the collected data. Post-process typically is used to solve intermittent problems. The most common scenario is to set a specific capture trigger that will cause a capture buffer full of data to be stored. The capture buffer then is studied with protocol decodes, display filters, statistics, or expert analysis measurements to track down the intermittent problem. As a general rule, post-process solutions require less processor bandwidth than do real-time analysis applications because the analysis does not need to keep up with network events as they happen.

### 27.2.2 Passive testing vs. intrusive testing

All protocol analyzers must connect electrically to the network under test. The manner in which they connect is dictated by the specifications of the Physical layer of the network being tested, and this may be a passive or intrusive electrical connection. In terms of the logical or protocol connection, however, the protocol analyzer can be a passive monitor or it can participate actively in the protocol and perform intrusive testing. Figure 27.4 shows the logical flow of the protocol control in a passive and an intrusive situation.

**Passive testing.** Many problems can be discovered by using passive measurements such as statistics and decodes that monitor the network under test. Passive tests do not alter or participate in the protocol on the network. An advantage of passive testing is that it does not contribute to problems that may be occurring, nor is any traffic load added to the network.

**Intrusive testing.** Some troubleshooting scenarios require that the network be stressed or loaded in order to force problems to occur. A protocol analyzer can act as a network node and intrusively stimulate the network under test. Traffic generation, bit error rate (BER) tests, stimulus/response tests, and simulation are intrusive measurements. A protocol analyzer that is designed to implement intrusive testing includes circuitry to maintain a transmit path in addition to a receive path. Unlike a normal node that either transmits or receives data, a protocol analyzer always must maintain its receive path so that it can observe all of the data on the network, including data that which the analyzer itself transmits.

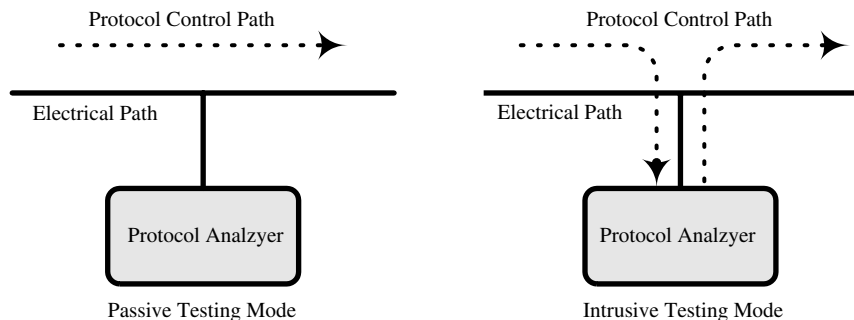


Figure 27.4 Passive vs. intrusive testing.

### 27.2.3 Dedicated vs. software acquisition systems

The distinguishing feature between two broad categories of protocol analyzers is the method of implementation for the analysis and acquisition systems. An analysis and acquisition system can be built as a software application in a PC, or it can be implemented with a dedicated system using a special-purpose microprocessor and specialized hardware. Table 27.2 describes the advantages and disadvantages of the two implementations.

**Software acquisition system.** Low-cost analyzers often are implemented using the same commercial network interface cards commonly used in PC systems or workstations. The acquisition and analysis functionality then is implemented in software as a PC application. The acquisition system therefore is limited in performance by the computing platform, which also must execute the user interface and measurement software. Commercial network interface cards typically are not designed to pass all data and all errored frames to the host, therefore restricting their suitability to protocol analysis.

**Dedicated hardware acquisition system.** In most high-performance protocol analyzers, the data analysis and acquisition system is implemented in dedicated hardware, often with special-purpose processors specially designed to facilitate manipulating and examining data at high speeds. This allows computation-intensive measurements to be executed, and multiple measurements to be executed simultaneously. Simultaneous measurement execution is required for most troubleshooting scenarios. An example of a high-performance measurement is “top talkers” on a 100Base-T LAN (where data is transferred at 100 Mbps). The top talkers measurement displays the network nodes that are transmitting the most traffic. This requires that the protocol analyzer process each frame on the network in real time, examining its source address and keeping a sorted database of the addresses.

### 27.2.4 Single-port vs. multiport

Most network troubleshooting applications require that a protocol analyzer connect only to one point on the network. There are, however, certain situations in which it is necessary to connect a protocol analyzer to more than one point. Typical scenarios usually involve the network interconnect devices such as bridges, routers, hubs, and switches. These devices have multiple network connections and they perform transformations on the protocol information that flows through them. Multiport protocol analyzers typically will support from two to eight simultaneous test ports.

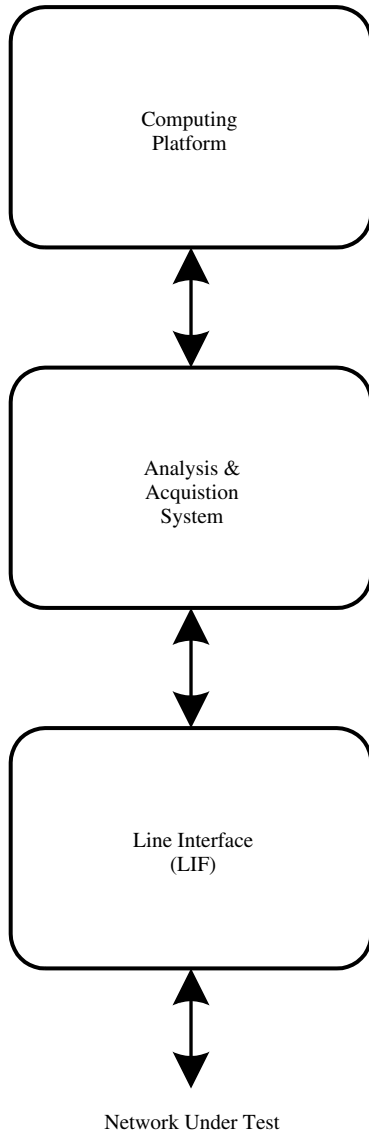
**Single-port.** Single-port analysis is used to describe a protocol analyzer that is using one, and only one, line interface to test the network. Figure 27.5 shows the block diagram for a single-port analyzer. There is a dedicated analysis and acquisition system and a dedicated computing system for the single line interface. Most protocol analysis and network troubleshooting is performed in a single-port configuration.

**Multiport.** Multiport protocol analysis is used to simultaneously test multiple points on the network. The different points could be the same or different physical network technologies. For example, a two-port analyzer could be used to test the latency of frames being processed by a router that routes T1 traffic into a 10Base-T Ethernet network. Multiport protocol analyzers are significantly more complex and more



**TABLE 27.2 Comparison of Analysis and Acquisition Systems. Software analysis and acquisition systems are implemented as applications in the computer platform (PC or Unix). Hardware systems are implemented as separate processors and circuitry dedicated to analysis and acquisition.**

Attribute	Comments	Software Solution	Dedicated Hardware Solution		
Cost	Hardware solutions are more costly due to the material cost of the dedicated processor and circuitry.	★★★★	★★		
Size	Hardware solutions are implemented as computer cards or modules. Software solutions are built as applications on the hard disk of the computing platform that acquires data from the network by using a standard network interface card.	★★★	★★		
Performance	The computing platform shares its processing power with the software solution, while the hardware solution offloads the computing platform.	★	★★★★		
Full-line rate acquisition	Hardware solutions capture traffic at 100 percent network utilization, while software solutions are limited to partial capture on most networks, often discarding frames.	★	★★★★		
Capture all errored frames	The standard network interface cards used with software solutions ignore many errored frames, which are of particular interest in network troubleshooting. Hardware solutions are built with special circuitry to capture these frames.	★	★★★★		
Triggering	The number of triggers and the sophistication of the trigger conditions is far greater in a hardware solution that has special circuitry to implement triggers.	★★	★★★★		
Filtering	The number of filters and the sophistication of the filter conditions is far greater in a hardware solution that has special circuitry to implement filters.	★★	★★★★		
Statistics	Network statistics require examining and sorting network traffic in real time. The processing power of a hardware solution offers more sophisticated real-time statistics.	★★	★★★★		
Expert analysis	Expert analysis requires capturing all the network frames and following complex frame sequences. The capture, filter, and trigger facilities of a hardware solution are superior to those of a software solution.	★	★★★★		
Simple troubleshooting	Simple troubleshooting involves observing network traffic and looking for connections, data transfer, and utilization. This can be performed well with either type of solution.	★★★	★★★★		
In-depth troubleshooting	In-depth troubleshooting requires excellent data capture, triggering, filtering, and the ability to examine errored frames. Hardware solutions provide a much more robust feature set.	★	★★★★		
Performance monitoring	Performance monitoring is a function of the statistical measurement capability of the product and is greatly enhanced by the processing power of a dedicated hardware solution.	★★	★★★★		
Established network technologies	Mature technologies are readily available on standard PC cards and PCMCIA cards for integration into a software solution. These same technologies are also available in custom hardware solutions.	★★★★	★★★★		
Newly emerging technologies	Emerging technologies sometimes are available on standard network interface cards that can be integrated into a software solution. However, the test solutions for these new technologies usually must be implemented as custom hardware solutions.	★★	★★★★		
		★★★★ Excellent	★★★ Good	★★ Fair	★ Poor



**Figure 27.5** Single-port protocol analyzer architecture.

costly than single-port analyzers, with each line interface requiring a dedicated analysis and acquisition system. They most commonly are implemented with additional dedicated hardware that allows the multiple ports to handle all of the data, and to provide time synchronization between the ports so that captured data can be correlated.

Figure 27.6 shows the block diagram for a multiport protocol analyzer that is constructed with a line interface and an analysis and acquisition system dedicated to each port.

## 27.3 Basic Protocol Analyzer Architectures

There are three main components to any protocol analyzer:

- Computing platform
- Analysis and acquisition system
- Line interface

### 27.3.1 Computing platform

The computing platform is a general-purpose processing system, typically a PC or a Unix workstation. The computing platform executes the user interface for the product and controls the measurements that are executed in the analysis and acquisition systems. It is common for other applications to be run in conjunction with a protocol analyzer. These include spreadsheet applications for analyzing data and measurement results, network management software, and terminal emulation applications that log into computer systems on the network. Computing platforms therefore are usually based on an industry-standard system that allows a user to interact openly with the application environment.

### 27.3.2 Analysis and acquisition system

Fundamentally, a protocol analyzer acquires data from the network and then analyzes that data. Thus the analysis and acquisition system is the core of a protocol analyzer. This system is responsible for transferring data from the line interface to the capture buffer, ensuring that all error conditions, the protocol state information, and

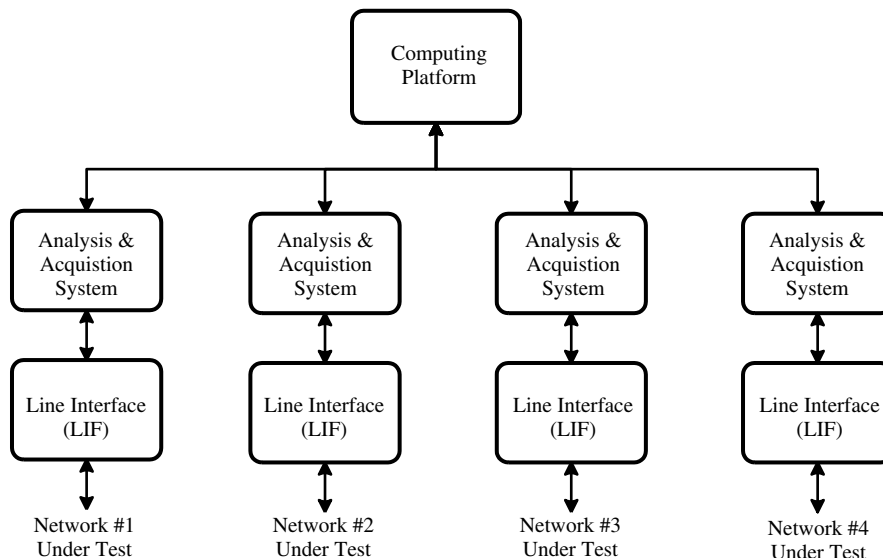


Figure 27.6 Multiport protocol analyzer architecture.

## 636 Network Test Instrumentation

the protocol data are correctly stored and timestamped. During real time and in post-process mode, the triggers and actions, the timers and counters, and the protocol followers are executed in the analysis and acquisition system.

Additionally, the measurements typically are executed in a distributed fashion between the computing platform and the analysis and acquisition system. In low-cost, software-based analyzers, the analysis and acquisition system functions are performed by the computing platform. In high-performance protocol analyzers, a dedicated processor and special-purpose hardware are used to implement the processing required by the analysis and acquisition functions.

### 27.3.3 Line interface

The physical hardware and firmware necessary to attach to the network under test are implemented in the line interface. Additionally, the line interface includes the necessary transmit circuitry to implement simulation functions for intrusive testing. The function of the line interface is to implement the Physical layer of the OSI Reference Model and provide framed data to the analysis and acquisition system. (Figure 24.1 in Chapter 24 provides a detailed description of the OSI network model.)

## 27.4 Detailed Protocol Analyzer Architecture

In order to implement full-featured protocol analysis, it is necessary for an analyzer to include all of the architectural components detailed in Figure 27.7. In order to clarify the functions of each of these blocks, the following discussion will describe how a frame is captured from the network and analyzed. The fundamental objective of the architecture is to accurately collect all of the data from the network under test (including timestamping the data and detecting all errors), and to store the data in the capture buffer so that it can be analyzed by the various protocol analysis measurements.

### 27.4.1 Line interface

This portion comprises two blocks: the network interface block and the simulate block.

**Network interface.** The network interface block provides the electrical interface to the network under test. It converts the analog signals from the network into a digital bit stream that is framed correctly. Any transmission errors encountered in the frame are detected by the network interface. The output of the network interface is a data structure of information formed in bytes that can be stored in the capture buffer and analyzed by the measurements.

**Simulate.** The simulate block provides formatting to form valid frames to be transmitted on the network. The frame type and the content of the individual frames are specified by the specific measurements selected by the user. The network interface block adds the necessary message framing and physically transmits the data onto the network.

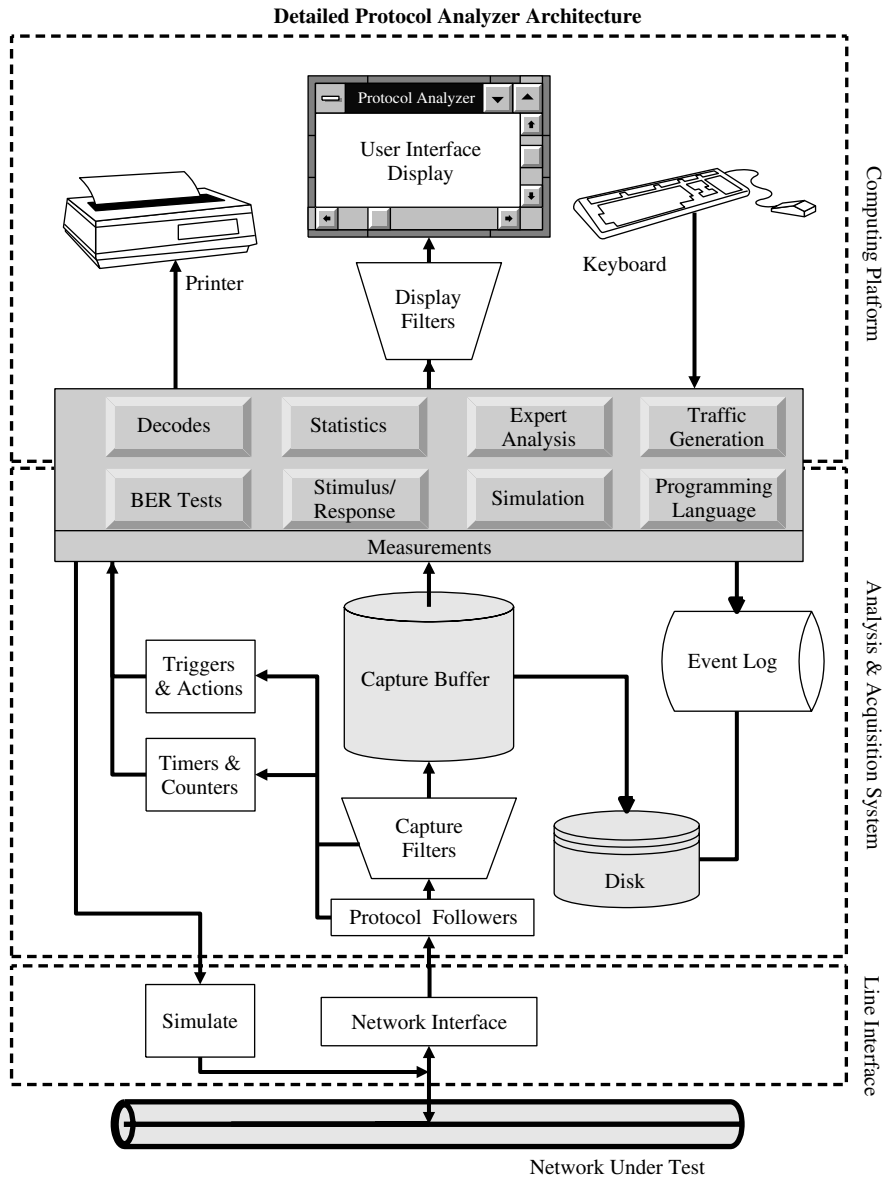


Figure 27.7 Detailed protocol analyzer architecture.

### 27.4.2 Analysis and acquisition system

The principal items in the analysis and acquisition system include protocol followers, capture filters, a capture buffer, a high-capacity fixed disk, the event log, timers and counters, and triggers/actions.

**Protocol followers.** Protocol followers are used to follow the state of the protocol at the Data Link layer. Errors in addressing, error recovery, and frame transmission

are detected and an indication of the error is added to the information stored in the capture buffer.

**Capture filters.** Capture filters compare user-defined specifications (such as address fields, protocol types, error conditions, and logical channel numbers) against the incoming data. If the filter criteria are met, then the frame is passed onto the capture buffer for storage; otherwise it is discarded.

**Capture buffer.** The capture buffer forms the core of a protocol analyzer. All collected data, error conditions, and timing information are stored in the capture buffer. The capture buffer is implemented as a circular queue, with a head pointer to indicate where new incoming data is stored and a tail pointer to indicate where the first frame of information was stored. Capture buffers typically are 2 to 256 MB in size. Data capture can occur in several ways: a circular, continuous, first-in first-out (FIFO) manner; a fill-the-buffer-and-stop mode; or in a triggered mode that stops the capture based on a specified event. In the FIFO mode, the head and tail pointers are next to one another and each new incoming frame overwrites the oldest frame in the buffer.

**Disk.** The disk is part of the computing platform for the storage of the software and data necessary to run the protocol analyzer. In Figure 27.7 it is shown as part of the analysis and acquisition system because it is used as an extension of the capture buffer. The disk can be used for storing capture buffer data as well as measurement results (e.g., statistics result logging).

**Event log.** The event log is implemented with a database on the computing platform's disk. When a measurement such as expert analysis detects a critical network event, or the statistics monitor detects a threshold being exceeded, an entry is made into the event log.

**Triggers and actions.** The triggers and actions form the base for much of the troubleshooting capability of a protocol analyzer. Triggers are implemented with a pattern-matching device (usually implemented in hardware for performance reasons) that matches user-defined patterns against incoming data. Triggers range from simple bit patterns to complicated byte streams that represent specific protocol sequences. When a trigger is matched, the measurement software is immediately interrupted to handle the trigger and invoke an action.

**Timers and counters.** The triggers initiate the timers and counters function. When an incoming frame matches the user-defined trigger condition, a timer is started or stopped, or a counter is incremented or decremented. The results of the timer and counter values are read by measurement applications. The timers usually are implemented in firmware or hardware so that the proper accuracy can be attained.

### 27.4.3 Computing platform

The principal items of interest on the computing platform include the measurement software, the display filters, the user interface display, and assorted peripherals.

**Measurements.** Measurements are software applications that execute on the processors of the computing platform and the analysis and acquisition system. The measurements include:

- Decodes
- Statistics

- Expert analysis
- Traffic generation
- BER tests
- Stimulus/response
- Simulation
- Programming language

These measurements are described in more detail in Chapter 24, section 24.6.

**Display filters.** Display filters are functionally equivalent to capture filters. The difference between the two is that a capture filter limits the incoming data from ever being analyzed or stored in the buffer, while a display filter is used to select what data already in the capture buffer is presented on the screen, typically by the use of a decode. Display filters operate on data that is captured and available for analysis by the measurements built into the product.

**User interface display.** The user interface is the set of screens and input/output devices (keyboard, mouse, printer) that provide the human/machine interface. All of the functionality of the product—the measurement setups, configurations, and measurement results—are accessed through the user interface.

**Peripherals.** The peripherals include standard PC or workstation peripherals such as keyboard, mouse, and printer. The keyboard and mouse are used to navigate the user interface. Measurement results often are printed out for troubleshooting and reporting applications.

## 27.5 Protocol Analyzer Implementation Technologies

The computing platform for many of today's protocol analyzers is a personal computer. The pervasiveness, computing power, and low cost of PC technology make the PC an obvious choice to implement the controller function of a variety of instruments, including protocol analyzers.

**Operating systems.** Most network managers and network engineers are accustomed to operating a variety of computer and test equipment. Therefore a protocol analyzer with a user interface based on an industry-standard operating system, such as Microsoft Windows or Unix, is advantageous. Typical protocol analysis applications have configuration windows, data display windows, and graphical displays. Most of the applications run by a protocol analyzer take advantage of the user interface capabilities offered by these operating systems.

**Network interfaces.** Commercially available network interface chip sets used in network and computing equipment often are incorporated into the line interface designs of protocol analyzers. Because these chip sets are focused on providing communications rather than testing capabilities, it usually is necessary to augment the commercial network chip set with additional hardware to provide the necessary protocol following, acquisition, and error handling.

**High-speed memory systems.** The capture buffer systems usually are implemented with a dedicated high-speed memory system based on DRAM design and

include additional hardware to handle the FIFO functionality necessary to implement a circular buffer.

**RISC (Reduced Instruction Set Computer) processor systems.** RISC processors are especially efficient at sorting, manipulating and analyzing streams of data. This makes them an ideal choice for implementation in the analysis and acquisition systems of a protocol analyzer. The functionality described in Figure 27.7 for the analysis and acquisition system often can be entirely implemented in a RISC processor.

**High-speed digital logic.** The triggering, filtering, and protocol-following functionalities often are implemented in a RISC processor, but in certain situations it is necessary to gain performance benefits by implementing these functions instead in high-speed digital logic, often in programmable logic arrays. The advantage of a programmable logic array is that it can be reconfigured for different network technologies by loading different logic overlays from preconfigured setups on the hard disk. This allows the analysis and acquisition system to be general enough to handle a variety of network technologies, and thus keep the price of the protocol analyzer down.

## 27.6 Protocol Analyzer Selection Criteria

When selecting a protocol analyzer, it is important to recognize that there are a number of products on the market. The user must first determine how the product will be used; the information in Chapter 24 and this chapter describe the functionality that is available. Table 27.3 lists the major categories of product features that can be used to evaluate protocol analyzers. No single product incorporates all of the capabilities outlined, so it is important to look for the combination of criteria that meets your needs.

**TABLE 27.3 Protocol Analyzer Selection Criteria.**

Selection Criteria	Common Attributes	Typical Applications
☑ Package	Portable, < 20 lb (9 kg)	The protocol analyzer is dispatched to the troubleshooting location.
	Transportable, 20-50 lbs (9-22 kg)	The protocol analyzer is moved infrequently or is used on a cart.
	Stationary, > 50 lbs (22 kg)	The protocol analyzer is used in a fixed location such as a lab bench or rack.
☑ Operation	Floor	The protocol analyzer is used in wiring closets and behind racks of equipment.
	Rackmount	The protocol analyzer is used in a test equipment bay or in a rack of network equipment.
	Table	The protocol analyzer is used in an office or lab environment.
☑ Power	AC	Typical operation for most test equipment.
	Battery	Needed for portability or in places where ac power is not available.



**TABLE 27.3 Protocol Analyzer Selection Criteria (Continued).**

Selection Criteria	Common Attributes	Typical Applications
<input checked="" type="checkbox"/> Computing platform	PC with Microsoft Windows  Unix workstation  Vendor-proprietary	Preferred by users familiar with PC systems and used for running other PC applications. Preferred by users familiar with Unix systems and used for running other Unix applications. Typically unique to a specific test application.
<input checked="" type="checkbox"/> Display system	<ul style="list-style-type: none"> <li>• Resolution (VGA, SVGA, etc.)</li> <li>• Color</li> </ul>	Display quality enhances usability but also increases product cost.
<input checked="" type="checkbox"/> User interface	<ul style="list-style-type: none"> <li>• Ease of use</li> <li>• Use of color</li> <li>• Graphical</li> </ul>	Should be intuitive and as easy to use as possible.
<input checked="" type="checkbox"/> I/O	<ul style="list-style-type: none"> <li>• Serial</li> <li>• Parallel</li> <li>• LAN</li> <li>• PCMCIA</li> <li>• HP-IB</li> </ul>	Used for printing measurement results and reports, and also remote control. PCMCIA provides the greatest flexibility because it can handle various interface cards.
<input checked="" type="checkbox"/> Capture performance	<ul style="list-style-type: none"> <li>• All frames</li> <li>• Errored frames</li> </ul>	Acquiring and analyzing all of the frames on the network, regardless of utilization or error levels, is essential to in-depth troubleshooting applications and performance.
<input checked="" type="checkbox"/> Analysis and acquisition system	Software  Dedicated hardware	Used for monitoring and simple troubleshooting applications where a low-cost solution is preferred. Used for performance monitoring, in-depth troubleshooting, and installation where a full-featured, full-capture-rate solution is required.
<input checked="" type="checkbox"/> Number of test ports	Single port  Multiport	Used for most troubleshooting situations where testing is focused on the segment or link. Used to test interconnect devices (bridges, routers, switches, etc.) in internetworking situations.
<input checked="" type="checkbox"/> Network interfaces	<ul style="list-style-type: none"> <li>• Ethernet</li> <li>• Token-passing bus</li> <li>• Token-Ring</li> <li>• 100Base-VG ANYLAN</li> <li>• Fast Ethernet</li> <li>• FDDI</li> <li>• MAN (Metropolitan Area Network)</li> <li>• ITU-T G.804 PDH</li> <li>• T1.105 SONET</li> <li>• G.708 SDH</li> <li>• TAXI</li> <li>• Fiber Channel</li> <li>• IBM Block Encoded</li> <li>• HSSI, High-Speed Serial Interface</li> <li>• V-series</li> </ul>	Ability to connect the analyzer to networks on your site is the most critical purchasing criterion. This is determined by the network interfaces offered by the vendor.
<input checked="" type="checkbox"/> Protocol stacks	<ul style="list-style-type: none"> <li>• AppleTalk</li> <li>• ATM</li> <li>• Banyan VINES</li> <li>• DECnet</li> </ul>	Supported protocol stacks determines if you can test the protocols running on your network. The protocol stacks (including all individual protocols in the stacks) should be supported by all the

## 642 Network Test Instrumentation

TABLE 27.3 Protocol Analyzer Selection Criteria (Continued).

Selection Criteria	Common Attributes	Typical Applications
	<ul style="list-style-type: none"> <li>• Frame relay</li> <li>• IBM/SNA</li> <li>• ISDN</li> <li>• ISO</li> <li>• LAN Manager</li> <li>• MAP</li> <li>• Novell NetWare</li> <li>• Sun</li> <li>• TCP/IP</li> <li>• X Window</li> <li>• X.25</li> <li>• XNS</li> </ul>	measurement capability and options in the analyzer.
<input checked="" type="checkbox"/> File import/export	<ul style="list-style-type: none"> <li>• ASCII</li> <li>• CSV</li> <li>• Vendor-specific</li> </ul>	Data captured by a protocol analyzer is imported into other applications such as databases and spreadsheets. This is accomplished through standard interchange formats.
<input checked="" type="checkbox"/> Measurements	Protocol decodes Statistics  Expert analysis Traffic generation  BER testing Stimulus/response testing Simulation/emulation  Programming language	Used for in-depth troubleshooting. Used for troubleshooting and performance monitoring. Used for in-depth troubleshooting. Used to stress test existing networks and to install new networks. Used to verify the transmission network. Used to query the network under test. Used to create installation tests for new networks and equipment. Used for creating specific network scenarios, usually for development applications or recreating network problems.
<input checked="" type="checkbox"/> Measurement options	Display filtering  Capture filtering  Triggers and actions  Timers and counters	Used in troubleshooting applications to pinpoint specific frames. Used in troubleshooting applications to reduce the amount of data captured by the analyzer. Used in troubleshooting applications to isolate intermittent failures. Used to create simple statistics or to verify protocol timing sequences.
<input checked="" type="checkbox"/> Connectivity	Remote control  SNMP-manageable  Terminal emulation  Network management systems	Used to control a protocol analyzer in a remote location via a LAN or modem connection. Used to query SNMP network devices or to provide SNMP information to queries from other network devices. This typically requires that the protocol analyzer be built with an MIB (Management Information Base). Allows the user to log on to other systems connected to the network. It is used to test connectivity and also to access applications available on other systems. The protocol analyzer can be invoked by Network Management Systems.
<input checked="" type="checkbox"/> Operating specifications	<ul style="list-style-type: none"> <li>• Safety regulations</li> <li>• Electromagnetic compatibility</li> <li>• Temperature</li> </ul>	These specifications vary by country and by the requirements of the application in which the protocol analyzer will be used.

**TABLE 27.3 Protocol Analyzer Selection Criteria (Continued).**

Selection Criteria	Common Attributes	Typical Applications
<input checked="" type="checkbox"/> Reliability	<ul style="list-style-type: none"> <li>• MTBF (Mean Time Between Failures) for the hardware</li> <li>• Software defects and upgrades</li> </ul>	The more critical the network being tested, the higher the reliability required of the protocol analyzer.
<input checked="" type="checkbox"/> Installation	<ul style="list-style-type: none"> <li>• Factory setup</li> <li>• User setup</li> </ul>	Most protocol analyzers are shipped completely configured as turnkey solutions. However, embedded protocol analyzers require the user to install the hardware and software.
<input checked="" type="checkbox"/> Support	<ul style="list-style-type: none"> <li>• Phone (help desk)</li> <li>• Bulletin boards</li> <li>• World Wide Web site</li> <li>• Technical consultants</li> </ul>	The particular type of support required depends on the specific network application, the preference of the user, and the geographic location.
<input checked="" type="checkbox"/> Training	<ul style="list-style-type: none"> <li>• Classes</li> <li>• Computer-based training (CBT)</li> <li>• Tutorials</li> <li>• Application Notes</li> </ul>	The particular type of training required depends on the specific network application, the preference of the user, and the geographic location.
<input checked="" type="checkbox"/> Help systems	<ul style="list-style-type: none"> <li>• On-line help</li> <li>• Documentation</li> </ul>	Using a protocol analyzer requires help systems to describe the operation of the product, the network interfaces, and the protocols supported.



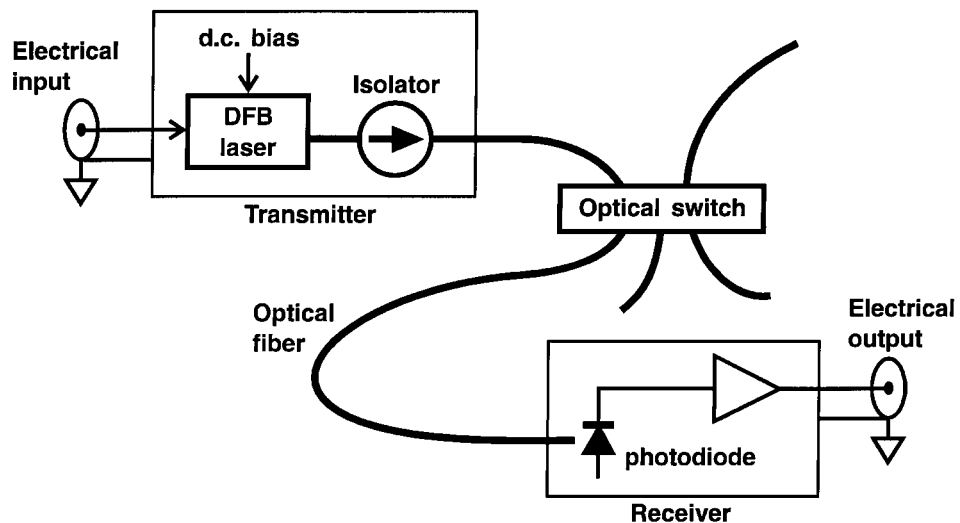
## Optical Element Testers

Brian L. Heffner

*Hewlett-Packard Laboratories, Palo Alto, California*

### 28.1 The Need for Optical Test

The transmission medium of a working communications link is transparent to the user. Whether it's a coaxial cable, a beam of microwave radiation, or an optical fiber that carries information from one location to another is not normally of concern. Accordingly, testing an optical network begins with the same equipment and procedures used to test any network, optical or not. For example, when the optical transmitter of Figure 28.1 is driven by a digital test pattern generator, an eye diagram



**Figure 28.1** A fiber optic communications link. The optical switch is an example of an optical component that may affect transmission of the optical signal.

**TABLE 28.1 Optical Measurements and Their Relation to Optical Network Properties.**

Measurement	Significance of measured quantity
Optical power	Power at receiver affects signal-to-noise ratio
Loss and attenuation	Reduces optical power at receiver
Spectral loss	Differences in performance between channels in wavelength-multiplexed systems
Polarization-dependent loss	Inconsistent loss, loss that drifts over time
Optical time domain reflectometry	Location of fiber breaks, determination of fiber attenuation, determination of splice, connector, and component losses
Multimode dispersion	Pulse distortion in multimode fiber networks
Chromatic dispersion	Pulse distortion in single-mode fiber networks
Polarization-mode dispersion	Pulse distortion in single-mode fiber networks, pulse distortion that drifts over time

(as discussed in Chapters 23 and 26) can be measured at the output of the receiver, allowing estimation of the bit error rate at any transmission rate.

### 28.1.1 Optical measurements

Purely electrical measurements such as these can confirm that a transmission network operates correctly. When such a measurement indicates a failure at some point between the electrical input of the transmitter and the electrical output of the receiver, however, the search for the problem requires an acquaintance with optical test procedures and test equipment.

To test for degradation or failure of an optical transmitter, for example, we would begin by measuring the transmitter output power with an optical power meter. The repeatability and loss of a fiber optic connector can be measured using an optical loss test set. An optical time domain reflectometer makes it possible to find a break in a fiber cable, even when the break and the measurement site are separated by kilometers. When we measure an unexpectedly poor eye diagram at the output of an optical network, we might search for its physical cause by measuring the optical dispersion of the fibers and components in the network.

This chapter describes the equipment and procedures used for these and other optical measurements of optical components and fibers, as described in Table 28.1.

### 28.1.2 Uniquely optical effects

Polarization-dependent loss and polarization-mode dispersion lead to effects that over time can change in networks using more than approximately 100 m of single-mode fiber. In fibers of this length the *birefringence*, i.e., the way the optical polarization is changed by the fiber, is likely to vary significantly in response to changes in temperature, potentially causing the signal strength and pulse distortion at the receiver to slowly change over time.

Birefringence will be described in section 28.5.1, but for now we observe that this effect is inherently optical in origin. Hence, electrical test of the network will not allow an estimate of variations in network performance owing to changes in birefringence; only optical measurement can provide the information needed for such estimates.

## 28.2 Optical Power

Like any traveling electromagnetic wave, light carries energy away from a radiation source. The time rate of change of energy that is radiated is called the *radiant flux*, or more colloquially the *optical power*. Measurement of optical power requires some form of photodetector, i.e., a transducer that converts received radiant flux to an electrical signal. Many types of photodetector commonly are used (Driscoll and Vaughn, 1978), but at the wavelengths used for optical communication the best performance is obtained from the photovoltaic p-n junction or photodiode.

### 28.2.1 Optical power meter

An optical power meter is similar to an optical receiver; both devices employ a photodetector to generate an electrical signal in response to an optical input. An optical power meter is designed to respond to a large range of radiant flux with good linearity. It sometimes is called an *average power meter* because it typically has a small measurement bandwidth, allowing measurement of low levels of optical power but precluding demodulation of information carried by the optical signal. Light propagating through free space can be measured, but measurement of optical power in communications networks implies, in practice, measurement of light guided by fiber optic cables.

A power meter useful for network test consists of a photodiode followed by a precision, gain-ranging current amplifier and display device, usually a numerical display (Figure 28.2). Such a device with basic features can be portable, rugged, and the size

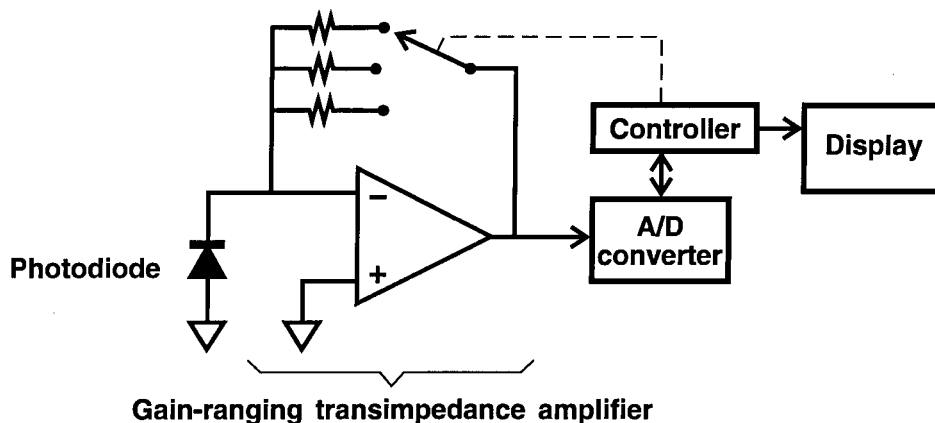


Figure 28.2 Block diagram of an optical power meter.

of a handheld calculator. A more full-featured power meter typically includes a GPIB (general purpose interface bus) interface, data filtering capabilities, and a variety of compatible optical power sensors.

### 28.2.2 Selecting and using a power meter

Selecting and using a power meter requires knowledge of the optical wavelength(s) used for transmission over the network. Different photodiode materials are suitable for use over different ranges of wavelength. Silicon photodiodes cover the approximate range 280–1050 nm, while germanium or indium-gallium arsenide photodiodes typically cover 800–1700 nm. Once the proper photodiode has been selected, it's still necessary to know the wavelength to obtain accurate power measurement because the responsivity (i.e., the conversion factor between photocurrent and optical power) of the photodiode depends on the wavelength.

Power meters require the user to know the wavelength of the light to be measured, and to set manually a wavelength calibration factor via the instrument keypad. Typically, a variety of connector interfaces can be attached to the power sensor to allow connection of various styles of fiber-optic connectors.

Optical power can be expressed in linear units, e.g., milliwatts or nanowatts, or in logarithmic units. The standard logarithmic unit is the decibel referenced to 1 mW, or *dBm*, given by

$$P_{dBm} = 10 \log_{10} \frac{P_{linear}}{1 \text{ mW}} \quad (28.1)$$

where

$P_{dBm}$  is optical power expressed in logarithmic units (dBm)

$P_{linear}$  is optical power expressed in linear units (W)

For example, 500 nW is equivalent to  $10 \log_{10}(500 \text{ nW}/1 \text{ mW}) = -33.0 \text{ dBm}$ .

### 28.3 Loss and Attenuation

Optical loss is perhaps the most fundamental characteristic of a fiber or component. Loss is a term used to describe reduction of optical power, and is a consequence of transmission through any passive device or fiber. For example, loss can occur in fiber connectors when light escapes the fiber core and ultimately is absorbed by the plastic jacket of the fiber. Similarly, an optical switch unavoidably exhibits some loss when the input light is imperfectly coupled into the output fiber.

Excessive loss in an optical network can result in degradation of the bit error rate or signal-to-noise (s/n) ratio, and may cause complete failure of the network. As in electronic systems, a device that causes an optical input power  $P_{in}$  (in linear units) to be reduced to an output power  $P_{out}$  is said to exhibit an optical loss, usually expressed in decibels as:

$$L = 10 \log_{10} \frac{P_{in}}{P_{out}} \quad (28.2)$$



where

$L$  is the loss expressed in dB

$P_{in}$  is the input power expressed in W

$P_{out}$  is the output power expressed in W

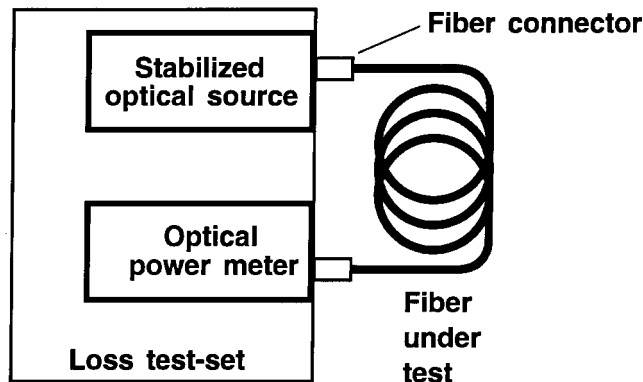
### 28.3.1 Measuring loss

Loss easily can be measured at a particular wavelength of interest by using a stabilized optical source and an optical power meter. Compact loss test sets include both the source and the power meter in a single instrument (Figure 28.3).

Light-emitting diodes and laser diodes are commonly used as optical sources. Laser diodes exhibit a narrower output spectrum and slightly poorer stability than light-emitting diodes, but offer approximately 100 times greater output power, allowing measurement of higher loss and better accuracy for moderate loss. Obviously, measuring the loss of an installed fiber requires a power meter that is physically separate from the optical source at the other end of the fiber. For more convenient and informative measurement of installed fibers, reflectometers (described in section 28.6) usually are used.

### 28.3.2 Test set stability

Test set linearity and stability can result in measurement accuracy of approximately  $\pm 0.005$  dB. When measuring through fiber connectors, however, the user must keep in mind that the repeatability of the connector loss typically is not better than  $\pm 0.10$  dB, since a connector is mated to different samples of a compatible connector, and may be considerably worse depending on the connector design and condition.



**Figure 28.3** Measuring the loss of a spool of fiber at a single wavelength. The output power of the stabilized source is first measured after connecting it to the power meter with a short fiber cable of negligible loss. The difference (in dB) between this initial power level and the power measured through the fiber spool is the loss of the long fiber. The optical source and power meter may be separate units, or may be integrated into a loss test set.

## 28.4 Spectral Loss

The intrinsic attenuation of optical fibers constitutes a large fraction of the total loss of many optical networks. Optical power decays exponentially along a fiber, allowing the attenuation caused by the fiber to be specified in units of dB/km. The attenuation of many modern fibers is limited by Rayleigh scattering, leading to higher attenuation at shorter wavelengths, as shown in Figure 28.4. Rayleigh scattering is proportional to  $1/\lambda^4$ , and is caused by transmission through particles much smaller than the optical wavelength, such as the silica molecules in the fiber core. Consequently, a more general characterization of loss includes the variation of loss as a function of wavelength, or spectral loss. Two approaches to this measurement are practiced, with either a broadband or a tunable source.

### 28.4.1 Spectral loss measurement with broadband source

In the first approach, shown in part (a) of Figure 28.5, a wide-spectrum optical source, such as a light-emitting diode or an incandescent lamp, is coupled into the device or network under test, and the device output is coupled to an optical spectrum analyzer, which is a scanning monochromator integrated with a low-noise power meter. The optical spectrum analyzer then displays the optical power transmitted through the device as a function of wavelength. The advantages of the spectrum analyzer approach include wide available spectral range, no coordination of the measurement conditions between the source and receiver ends of a long fiber, and fast measurement speed.

### 28.4.2 Spectral loss measurement with tunable source

In a second approach, part (b) of Figure 28.5, a tunable laser is coupled through the device into a power meter, which measures the transmitted power as the laser is tuned over the desired measurement range. The tunable laser approach accommodates devices with higher loss and can better resolve features occurring at closely spaced wavelengths than can the spectrum analyzer method.





Figure 28.4 Attenuation of a low-loss fiber as a function of optical wavelength.



**Figure 28.5** Measuring the spectral attenuation of a fiber: (a) Spectrum analyzer technique, where the spectral selectivity is achieved at the receiver end, for example with a grating-based optical spectrum analyzer. (b) Tunable source technique, where spectral selectivity is achieved at the source end. This technique requires coordination between the tunable source and the power meter.

### 28.4.3 Calibration

In both cases, coupling a uniform optical power into the device at all wavelengths is unnecessary. A calibration sweep, in which the power transmitted through a short fiber cable is measured and recorded as a function of wavelength, can be made. The results of a second sweep with the device in place can then be subtracted (expressed in dB) from the calibration sweep to obtain the spectral loss.


## 28.5 Polarization, Birefringence, and Polarization Dependent Loss (PDL)

Just as unexpected loss in an optical network can result in poor transmission quality, variations in loss can lead to changes in a network's performance over time. One mechanism that can lead to loss variation is the interaction of fiber birefringence (described in section 28.5.2) and the *polarization-dependent loss* (PDL) of an optical component. Whereas the loss of a device describes the reduction in transmitted optical power, PDL describes how much the loss changes as the polarization is varied over all possibilities.

Even when components of a network are completely stable from a mechanical standpoint, the birefringence of fiber cables can be expected to drift over time, leading to performance variations because of the effects of PDL. Understanding this phenomenon requires some acquaintance with optical polarization. A brief introduction is provided here; the interested reader is directed to other texts, cited in the chapter references, in which the topic is treated in greater depth (Kliger et al., 1990; Collet, 1993).

### 28.5.1 Optical polarization

Light is a transverse electromagnetic wave, and oscillation of the optical field in the two transverse dimensions allows for a variety of optical polarizations. Linear polarization may be the simplest and most familiar example of polarized light, but the electric field of polarized light does not necessarily oscillate in a single transverse direction.



**Figure 28.6** Electric field oscillations of polarized light: (a) linear polarization, and (b) elliptical polarization. Magnetic fields also are present in the traveling electromagnetic waves, but are not shown. Shown to the right are the electric field oscillations at a fixed  $x$ - $y$  plane.

More generally the polarization is elliptical. Elliptical polarization results when the electric field oscillation has two orthogonal linear components linked by a constant phase relation, as shown in Figure 28.6. The loss of a device exhibiting PDL may take maximum and minimum values when the incident light is elliptically polarized.

### 28.5.2 Birefringence

The birefringence of optical fiber cables can lead to time variation in the performance of optical networks. *Birefringence* is a term used to describe the ability of a crystal, device, or optical fiber to change the polarization of incident light without incurring PDL. For example, propagation of polarized light through a window of crystalline quartz will change the polarization of the light (birefringence), but the light will suffer the same loss regardless of the polarization (no PDL).

A real device can exhibit both significant birefringence and PDL. Practical fiber cables, however, exhibit negligible PDL and contribute only to loss and birefringence. In other words, fiber cables exhibit loss that is largely independent of the transmitted polarization, and also exhibit birefringence, transforming a given input polarization into a different output polarization. At a particular transmission wavelength, the loss of a fiber will remain constant while the birefringence typically will vary in response to changing temperature and aging of the cabling material.

Even when an optical transmitter operates at a fixed optical polarization, the changing birefringence of a fiber cable will result in a changing polarization at the cable output. If this signal then is transmitted through a device exhibiting PDL, or is incident on a receiver in which the responsivity depends on polarization, the effective optical power at the receiver will change over time, resulting in a variable bit error

rate. Fiber is manufactured under conditions carefully controlled to reduce the intrinsic birefringence to very low levels. In spite of the best manufacturing techniques, mechanical stress on the fiber, caused by the cabling process and by bends in the deployed fiber, will induce birefringence. Temperature indirectly changes birefringence primarily by changing the mechanical strain of a fiber.

### 28.5.3 Polarization-dependent loss (PDL)

A linear polarizer, which exhibits a loss that depends on the incident polarization, is an example of a polarization filter. Polarized sunglasses, for example, are designed to block transmission of light in which the electric field oscillates horizontally. The filtering action is not perfect: perhaps 20 percent of the incident vertically polarized light is transmitted, while only 0.2 percent of the horizontally polarized incident light is transmitted. *Polarization-dependent loss* is a term used to describe the polarization selectivity of a partial polarizer, and is defined as the ratio of the transmission of the most favored polarization to that of the least favored polarization, usually expressed in decibels. In the example of the sunglasses, the PDL is  $10 \log_{10} (0.2/0.002)$ , or 20 dB.

#### 28.5.3.1 Polarization adjustment

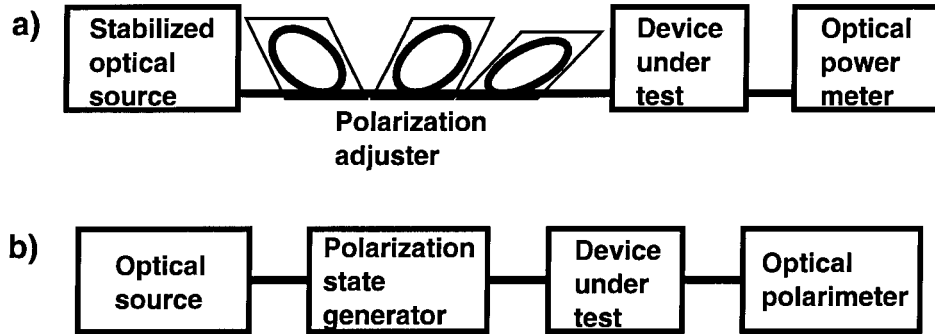
A useful way to measure PDL is simply to measure the loss through a device as the input polarization is varied over all possible values. To control the input polarization, we can take advantage of the strain-induced birefringence in a loop of optical fiber. Bending the fiber induces a uniform strain throughout the loop, resulting in the loop acting as a linear birefringence or waveplate. By rotating the fiber loop to change the orientation of the birefringence, the output polarization can be varied.

An arrangement of three successive fiber loops is called a *polarization adjuster* because it allows any input polarization to be transformed into any desired output polarization according to the orientations of the three loops. In theory, two loops would suffice if they were of precisely the correct diameters for the appropriate wavelength; in practice, three loops are used to allow for variations in diameter, fiber construction, and wavelength.

When the fiber cable diameter is 1 mm or less and at least a meter of fiber can be handled free of any enclosures, a useful polarization adjuster can be improvised by holding two or three loops of fiber in a coil approximately 5 cm in diameter. By rotating the orientation of this coil in the hand, the transmitted state of polarization can be crudely adjusted. Slightly more sophisticated polarization adjusters are commonly made up of three rotatable paddles, usually held in place by friction and adjustable by hand. When a coil of fiber is attached to each paddle, the three paddles can be set repeatably to cause the transmitted signal to match any desired polarization.

#### 28.5.3.2 Measuring PDL

Types of PDL measurement techniques are shown in part (a) of Figure 28.7. The *polarization search* technique requires a stabilized optical source and a power meter,



**Figure 28.7** Techniques for measurement of polarization-dependent loss: (a) polarization search technique, and (b) matrix technique.

both of which can be incorporated into a loss test set. The transmitted optical power is monitored as the polarization adjuster generates all possible polarizations. The ratio of the maximum to minimum transmitted power is the PDL, usually expressed in dB.

The polarization adjuster can be varied manually or driven by motors. Motor-driven polarization adjusters are designed specifically for PDL measurement, and can provide good polarization coverage in a random manner over a period of approximately one minute. By monitoring the optical power over this period of time, the maximum and minimum values can be found, with the ratio of these values determining the PDL. Measurement accuracy of  $\pm 0.004$  dB is attained by commercially available systems.

When polarization is being manually adjusted, each loop is rotated in a direction to maximize the power meter reading. After several iterations of adjusting all loops, a local maximum will have been found. Usually this also will be a global maximum, but this must be confirmed by noting the power level and coil angles, and then repeating the search from a different set of starting angles. If the same maximum power level is found at a different set of optimized coil angles, the measured maximum is taken to be a global maximum. A similar procedure will find the global minimum of power transmission.

Two weaknesses of the polarization-search method of PDL measurement include polarization-dependence of the power meter and coupling of the polarization adjuster loss to the position of the adjuster. The latter problem can occur when the fiber is bent too sharply, typically at the extremes of the coil rotations, causing macrobending losses. This problem is avoided by proper design of the coil rotation mechanism. Most photodiode power detectors exhibit a responsivity which depends slightly on the incident polarization. This dependence might be only  $\pm 0.01$  dB, but it adds directly to the PDL measurement uncertainty. Careful selection of the optical detector is therefore essential. Under some conditions the polarization dependence of a detector may be reduced by incorporation of a depolarizing filter based on multiple scattering or absorption and re-emission of light (Nyman and Wolter, 1993).

A third inherent weakness of the polarization-search method is the search itself, which is not deterministic and can lead to long measurement times. Systems are available that avoid a polarization search by measuring a Jones matrix or a Mueller

matrix that describes the polarization-transforming properties of the device under test. The matrix is measured by stimulating the device with three (Jones) or four (Mueller) accurately known polarizations, and measuring the response at the device output using a polarimeter, as shown in Figure 28.7 part (b). The PDL can be extracted from either of these matrices. Using either matrix technique, PDL can be measured in approximately two seconds to an accuracy of  $\pm 0.05$  dB.

## 28.6 Reflectometry

A *reflectometer* is a tool allowing measurement of reflections along the length of an optical fiber. An optical time domain reflectometer (OTDR) trace displays the strength of the measured reflections as a function of distance from the reflectometer.

Reflections can be grouped into two types, discrete and continuous. A *discrete reflection* occurs at a distinct point along the transmission path, and might be caused by a fiber connector, a receiver, or by an unintended break in the fiber. In contrast, a *continuous reflection* is caused by light scattering, and usually results in a low level of reflection that changes slowly along the length of a fiber. While measurement of discrete reflection sites can be informative, even more information can be deduced from measurement of backscattered light as a function of position along the fiber.

### 28.6.1 Rayleigh backscatter


No optical fiber is completely lossless. Much of the light lost in transmission through a low-loss fiber is scattered, or redirected, by the silica molecules of the fiber core. The strength of this type of scattering, called *Rayleigh scattering*, is inversely proportional to the fourth power of the optical wavelength, leading to stronger scattering for shorter wavelengths.

Most of the scattered light is directed toward the sides of the fiber, and is absorbed by the plastic jacket of the fiber. A small fraction of the scattered light, however, is directed back toward the light source and is guided by the fiber core. This backscattered light, after detection and analysis by a reflectometer, allows measurement of not only the loss of a fiber network, but where that loss occurs. Moreover, the measurement is performed with access to only one end of the fiber, avoiding the problems of coordinating measurement equipment at geographically separate locations.

### 28.6.2 Basic reflectometer measurements

Several basic measurements will be explained with the aid of the schematic reflectometer trace of Figure 28.8. Reflected power is displayed on a dB scale as a function of distance. Peaks and discontinuities in the trace indicate network components or fiber splices, and the sloped lines indicate Rayleigh backscatter in continuous sections of fiber.

**Location.** The reflective features of a fiber network are directly plotted as a function of distance from the reflectometer. The peak at (a) indicates a discrete reflection, for example, from a mechanical splice between sections of fiber. Distances between splices or other features can be directly measured.



**Figure 28.8** A reflectometer trace showing optical power reflected from a fiber as a function of position of the reflection. Diagnostic information exhibited by the trace includes (a) position, (b) reflectivity, (c) fiber attenuation, and (d) discrete loss.

**Reflectivity.** The height of the peak (b) yields the strength of the reflection in relation to other reflections along the fiber.

**Attenuation.** Assuming the level of Rayleigh backscatter is consistent at all points along the fiber, the slope (c) of the backscatter directly yields the attenuation of the fiber in dB/km. The assumption of consistent backscatter at all points is justified for all fibers except those intentionally manufactured with a position-dependent parameter. For example, a fiber made with a graded chromatic dispersion may exhibit a slight gradation in backscatter.

**Discrete loss.** The loss in a splice between fiber sections is given by the reduction (d) in the trace at the splice, assuming the level of Rayleigh backscatter is the same in both fiber sections. The assumption of identical backscatter in different sections of fiber is justified when the sections are manufactured under very similar conditions, such as when a number of spools of fiber of the same model number are purchased from the same manufacturer.

### 28.6.3 Direct-detection OTDR

A direct-detection OTDR acquires the reflection response of a fiber or network by measuring reflected optical power as a function of time in response to a transmitted optical pulse. This is the simplest type of OTDR, as can be seen from the block diagram shown in Figure 28.9. The 3 dB fiber coupler is a device that splits optical pulses from the laser into the two fibers on the right-hand side of the coupler. One of the right-hand fibers is terminated to suppress any reflected signal, while the other is connected to the fiber or network under test. Reflections from the fiber or network under test are split between the laser, (where the signal is lost), and the re-






Figure 28.9 Reflectometer block diagram.

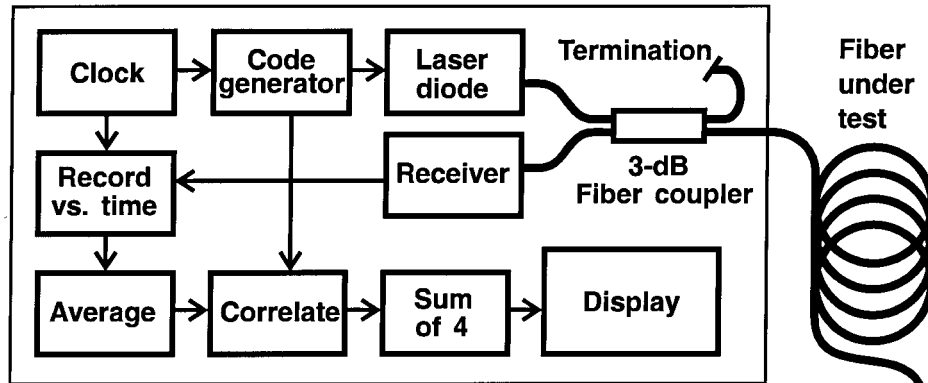


Figure 28.10 Correlation reflectometer block diagram.

ceiver. The received signal is measured as a function of time and displayed as a function of distance using an assumed value for the group velocity of light in a fiber.

After the first trace is measured, a second pulse is transmitted into the fiber, the reflection is detected and added to the first trace, and the average of the two traces is displayed. As the number of traces averaged gradually increases, the displayed noise floor can be observed to fall. Since the noise floor falls by 1.5 dB for every doubling of the number of traces averaged, the rate of improvement of the noise floor diminishes as the averaging process progresses.

#### 28.6.4 Correlation OTDR

The trade-off between total measurement time and noise performance can be greatly improved by use of a *correlation OTDR* (Figure 28.10). Using the terminology of linear, time-invariant systems, the output of an OTDR receiver is the

convolution (denoted by  $*$ ) of the optical pulse shape, the reflective impulse response of the fiber, and the impulse response of the receiver:

$$y(t) = p(t) * f(t) * r(t) \quad (28.3)$$

where

$y(t)$  is the receiver output

$p(t)$  is the optical pulse shape

$f(t)$  is the impulse response of the fiber on reflection

$r(t)$  is the receiver impulse response

By properly shaping the optical pulse and then correlating the receiver output with the known optical pulse shape, we can enjoy the noise advantages of a wide, energetic pulse and still obtain the high resolution of a narrow pulse. The result of the correlation (denoted by  $\star$ ) is

$$p(t) \star y(t) = [p(t) \star p(t)] * f(t) * r(t) \quad (28.4)$$

where

$y(t)$  is the receiver output

$p(t)$  is the optical pulse shape

$f(t)$  is the impulse response of the fiber on reflection

$r(t)$  is the receiver impulse response

When the autocorrelation of the optical pulse approximates a delta function, we obtain the same measurement as if we had used a narrow optical pulse (Nazarathy et al., 1989).

The best solution applicable within the constraints of reflectometry calls for using a family of four pulse shapes. The envelope of each pulse is a binary code designed so that the sum of the autocorrelations of the four pulses is a delta function. To obtain a reflectometer trace, reflections are measured in response to each of the four coded pulses. Each reflection is correlated with the corresponding code, and the four resulting correlations are added together to obtain the reflected signal versus time or distance. The correlation technique allows a given noise level to be achieved with fewer measured pulses, hence more quickly, than without correlation.

### 28.6.5 Continuously-modulated optical carrier reflectometer

Instead of measuring the reflected response to optical pulses as in the preceding two methods, a reflectometer trace can be calculated from the reflected response to a continuous optical signal that carries a modulation covering a range of frequencies. A continuously modulated optical carrier reflectometer (Figure 28.11) consists simply of an electrical network analyzer coupled to an optical source and optical receiver, with a fiber coupler providing access to the fiber under test.

The electrical source generates a swept sine wave, which drives an optical modulator. The resulting optical signal carries a sinusoidal amplitude modulation. A phase-sensitive receiver detects both the amplitude and phase (relative to the source) of the

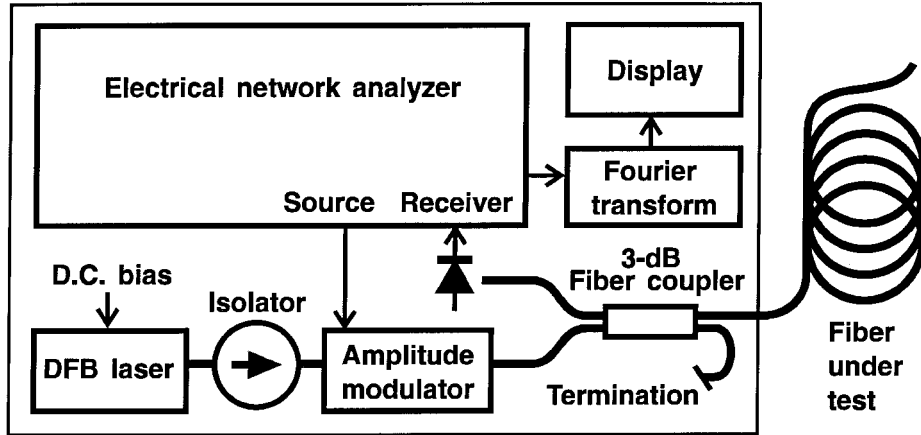


Figure 28.11 Continuously modulated optical carrier reflectometer block diagram.

photocurrent generated from the optical signal reflected by the fiber under test. The amplitude and phase, measured over a range of modulation frequencies, comprise the complex frequency input to an inverse Fourier transform algorithm. The magnitude of the output of the transform (a function of time) yields the reflectivity of the fiber under test.

Measurement over a frequency range of 40 GHz is possible using this method, leading to a two-point resolution of only a few millimeters. The primary limitation is the inability to measure the Rayleigh backscatter signal of most fiber. Range ambiguity also is a possible drawback, owing to the aliasing inherent in the discrete Fourier transform.


### 28.6.6 OTDR specifications

The different types of reflectometers have in common some fundamental specifications that characterize their performance.

- Dynamic range
- Resolution
- Amplitude linearity
- Dead zone

**Dynamic range.** A reflectometer is limited by electronic noise in the optical receiver. The *noise equivalent power*, or NEP, of the reflectometer is the reflected optical power that would cause a hypothetical noiseless receiver to generate an electronic output equivalent to the actual electronic noise. The one-way dynamic range (a) in Figure 28.12 is the displayed difference, in dB, between the initial level  $P_0$  of Rayleigh backscatter and the NEP. The backscatter can be measured over the entire length of a fiber only if the total loss of the fiber is less than the one-way dynamic range.

The dynamic range of a reflectometer is affected by several measurement conditions. A reflectometer trace typically is constructed as an average of many consecutive



**Figure 28.12** OTDR specifications illustrated on a measurement trace: (a) one-way dynamic range, (b) two-point resolution, and (c) dead zone. The broken line at (c) illustrates the linear receiver response in the absence of saturation.

traces in order to reduce the effects of noise. As a result, the effective NEP of an averaged trace is inversely proportional to the square root of the number of traces averaged. The initial backscatter level  $P_0$  is proportional to the optical pulse length, and also depends on the physical characteristics of the fiber under test. Dynamic range therefore depends on the number of traces averaged, the pulse length, and the fiber under test. The dynamic range of a commercial reflectometer is specified for a given pulse width and a given total time for the averaging process. For example, the one-way dynamic range of an OTDR might be specified as 22 dB for a 1  $\mu$ s pulse after 3 minutes of averaging.

**Resolution.** Each feature of a reflectometer trace has a finite width, usually set by the measurement pulse width. Two features separated by less than this width cannot be resolved individually. The two-point resolution is the minimum separation between features that can be resolved, (b) in Figure 28.12. Two-point resolution can be improved by using shorter pulses, but at the expense of dynamic range. A second type of resolution, the display resolution, is the spacing between the data points constituting the trace.

**Amplitude linearity.** An application of reflectometry is the measurement of fiber splice loss and connector loss. Because these losses are estimated from discontinuities in the backscatter trace, amplitude linearity limits the minimum loss that can be measured in this way. Amplitude linearity, specified in dB/dB, expresses the deviation from a straight line that would occur in measurement of the backscatter of a fiber exhibiting perfectly constant loss and constant Rayleigh backscatter along its length.

**Dead zone.** A very strong discrete reflection may drive the optical receiver out of its normal linear range and into saturation. When this happens the receiver will re-

main saturated for a short time even after the optical reflection has fallen to a low level. As a result, events received during this “dead time” will not be measured. The dead zone, (c) in Figure 28.12, is the corresponding distance that is effectively obscured beyond a strong, saturating reflection.

## 28.7 Dispersion (Pulse Distortion)

So far, this chapter has concentrated on locating and quantifying the network elements that cause optical loss. The aim up to this point has been to ensure that the maximum amount of optical power is guided from the transmitter to the receiver. We now address a second important consideration, *pulse distortion*.


A pulse transmitted through a network with absolutely no loss will arrive at the receiver with its initial energy intact, but the shape of the transmitted pulse nonetheless will be changed if the network is dispersive. The detrimental effect of pulse broadening is illustrated in Figure 28.13, where isolated bits are seen to be swamped by the effects of adjacent complementary bits. Without attention to dispersion, we might end up with a network in which plenty of power arrives at the receiver, but which is useless as a communication channel except at unacceptably low bandwidth.

We will consider three different types of dispersion in order of decreasing severity:

1. Multimode dispersion, which affects transmission only through multimode fiber.
2. Chromatic dispersion, which often is the dominant mechanism of dispersion in single-mode fiber, where multimode dispersion is by definition absent.
3. Polarization-mode dispersion, which can become important when chromatic dispersion is nearly eliminated, either by choice of the transmission wavelength or by careful dispersion compensation.

### 28.7.1 Multimode dispersion

An optical fiber consists of a central, light-guiding core surrounded by a cladding material of slightly lower refractive index. Light approaching the core-cladding interface is reflected back into the core by total internal reflection, allowing the fiber to guide light with very low loss by continued multiple reflection. A fiber with a relatively large



**Figure 28.13** Effects of dispersion on pulse transmission: (a) broadening of a single pulse, and (b) reduction in signal fidelity caused by interference of adjacent bits.

core (typically at least 50  $\mu\text{m}$  in diameter), and relatively large difference in refractive index between the core and cladding, is able to transmit light with the power distributed in distinctive spatial patterns across the core and inner cladding. These spatial distributions of power are called *modes* of the fiber (Keiser, 1983).

Multimode fiber is any fiber that supports more than one spatial mode. Multimode fiber typically is used because its large core size leads to easier mechanical alignment of connectors compared to single-mode fiber, and easier coupling of light into the fiber from optical sources such as light-emitting diodes.

The pulse-broadening phenomenon of multimode dispersion occurs because slightly different group velocities are associated with the various spatial modes of a multimode fiber. As a result, a signal coupled simultaneously into many spatial modes at the input of a fiber arrives at the output after a distribution of time delays, and an input pulse with sharp edges arrives at the output both broadened and with its edges smoothed out.

**Fiber length dependence.** The severity of the effect of dispersion depends on the length of fiber traversed. The exact length dependence is strongly affected by how quickly optical power is exchanged between different spatial modes. Power transfer can occur within the fiber, as well as at splices and connectors. The *coupling length* is a measure of the length of fiber traveled between mode-coupling events.

Multimode dispersion is proportional to the square root of the fiber length for lengths much greater than the coupling length. For short lengths of fiber, less than the coupling length, multimode dispersion increases proportionally to fiber length. Because single-mode fiber is usually used for long-distance transmission, and because the coupling length for graded-index multimode fiber may be greater than several kilometers, multimode dispersion usually is assumed to increase proportionally to length, and is characterized by a length-bandwidth product in units of MHz-km. The multimode dispersion of a 62.5  $\mu\text{m}$  core, graded-index fiber might be specified as not greater than 500 MHz-km.

**Testing multimode dispersion.** Multimode dispersion is tested by measuring the frequency response of an amplitude-modulated laser source transmitted through a known length of fiber. The optical spectrum must be narrow enough that chromatic dispersion will not be important, so a distributed-feedback laser typically is used for measurement of graded-index fibers. Using the setup in (a) of Figure 28.14, we measure the frequency  $f_{3dB}$  at which the modulation response is reduced by 3 dB after transmission through a fiber of length  $L$ . This corresponds to a 6 dB reduction in the electrical response because the photodiode generates a *current* proportional to optical *power*. The measured length-bandwidth product is then simply  $f_{3dB}L$ .

### 28.7.2 Chromatic dispersion (CD)

Systems employing single-mode fiber avoid the problems of multimode dispersion entirely. The dispersion of single-mode systems usually is dominated by *chromatic dispersion* (CD), defined as the derivative of the group delay with respect to wavelength for a given length of fiber. Fiber can be designed to allow only one spatial

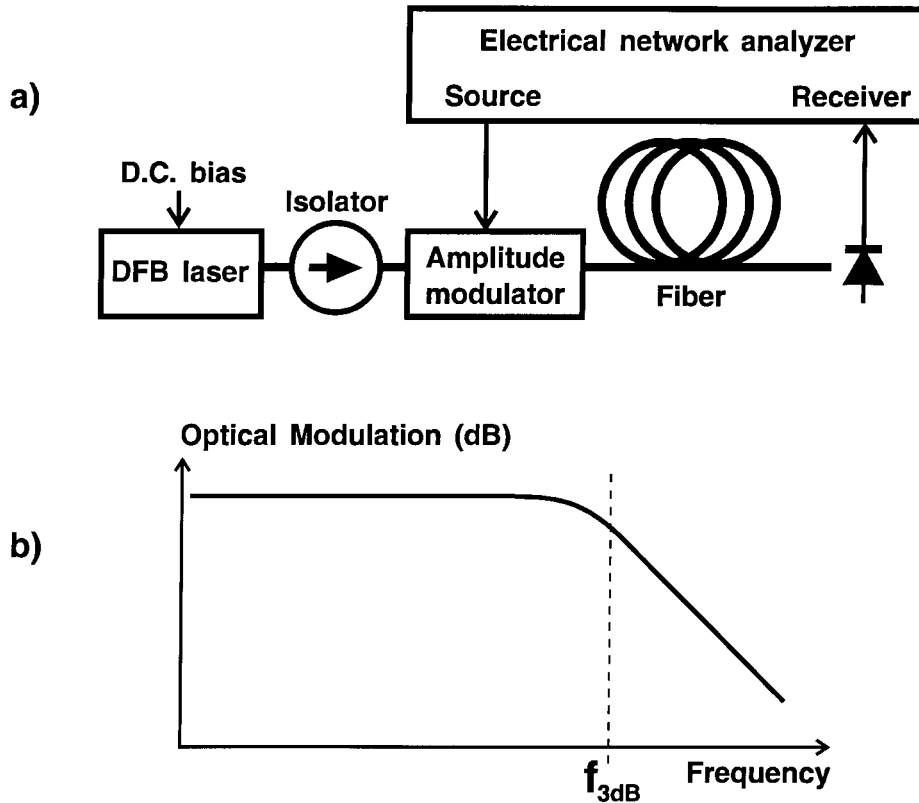



Figure 28.14 (a) Block diagram of multimode dispersion measurement; (b) result of measurement.

mode to propagate by limiting the core size to a small diameter and specifying a small index difference between core and cladding. Doing so assures that only a single spatial mode will be guided by the fiber over a range of optical wavelengths.

For example, a typical fiber designed for long-haul communication will guide a single spatial mode over a wavelength range of approximately 1240–1650 nm. At longer wavelengths the loss at small bends in the fiber increases, while at shorter wavelengths a second spatial mode will be guided in straight sections of the fiber, leading to multimode distortion as described in the preceding section.

The effect of CD arises because different wavelengths of light can travel along a fiber at different speeds. Since a pulse of light at the fiber input is composed of an optical spectrum of finite width, the different propagation delays suffered by the different spectral components result an output pulse with a time evolution different from the input pulse. Usually the pulse is smoothed and broadened by CD, although pulse compression also is possible.

The penalty imposed by this behavior is exacerbated by transmitter “chirp,” i.e., by changes in the optical wavelength within a single transmitted bit. Systems in which CD is a significant limitation employ methods to control chirp of the transmitter, for example by use of distributed-feedback lasers and external modulators. Even



**Figure 28.15** Analysis for chromatic dispersion measurement. Data points are measured relative group delay. The approximately parabolic Sellmeier curve is a least-squares fit to the data. The line indicating dispersion is the derivative of the Sellmeier curve.

when chirp is entirely eliminated, a finite optical spectrum is transmitted owing to the information bandwidth of the signal.

**Common elements of CD measurement.** The CD measurements to be described share some common elements. Delay is measured as a function of wavelength, and the desired CD is the derivative of delay with respect to wavelength. Because we expect a certain functional dependence between delay and wavelength, we fit the measured points with an analytical function. Compared to direct calculation of finite differences, CD derived from the derivative of the fitted analytical function is much more robust in the presence of measurement noise.

Measured delays are shown in Figure 28.15 along with a fitted Sellmeier function and its derivative. A quadratic analytical function usually is used for dispersion-shifted fiber (with a dispersion minimum of approximately 1550 nm), while a 3-term Sellmeier function usually is fitted to measurements of a dispersion-unshifted fiber (with a dispersion minimum of approximately 1310 nm). The analytical functions and the corresponding dispersions are given below:

Quadratic function:

$$\tau(\lambda) = A + B\lambda + C\lambda^2 = \tau_0 + \frac{S_0}{2} (\lambda - \lambda_0)^2 \quad (28.5a)$$

$$D(\lambda) = B + 2C\lambda = S_0 (\lambda - \lambda_0) \quad (28.5b)$$



3-term Sellmeier function:

$$\tau(\lambda) = A\lambda^2 + B + C\lambda^{-2} = \tau_0 + \frac{S_0}{8} \left( \lambda - \frac{\lambda_0^2}{\lambda} \right)^2 \quad (28.6a)$$

$$D(\lambda) = 2A\lambda - 2C\lambda^{-3} = \frac{S_0\lambda}{4} \left( 1 - \frac{\lambda_0^4}{\lambda^4} \right) \quad (28.6b)$$

where

$A, B, C$  are constants determined by a least-squares fit to data

$\tau(\lambda)$  is the group delay as a function of wavelength

$D(\lambda)$  is the chromatic dispersion as a function of wavelength

$\tau_0$  is the group delay at the zero-dispersion wavelength


$S_0$  is the zero-dispersion slope

$\lambda_0$  is the zero-dispersion wavelength

CD usually is specified by the zero-dispersion wavelength  $\lambda_0$  in nm and zero-dispersion slope  $S_0$  in ps/nm-km. The operating wavelength of most fiber optic systems is either in the range 1270–1350 nm where the CD of simple (dispersion-unshifted) silica fiber is minimum, or in the range 1525–1575 nm where the fiber loss is minimum. Accurate CD measurement is most important within these ranges.

Three methods suitable for CD measurement of long fibers are presented here. These methods are suitable for use both in the lab and in the field. CD of short (1 m) lengths of fiber can be measured by use of optical interferometry, but only in a controlled lab environment. The interested reader can consult the chapter references (Cohen, 1985; Pelayo et al., 1988) for details of interferometric CD measurements.

**Pulse delay.** The pulse delay or time-of-flight approach is in many ways the most direct measurement of CD. The relative propagation delay of a short optical pulse is measured at various wavelengths using equipment similar to that shown in the block diagram of Figure 28.16. The optical test source can be a continuous-wave tunable



**Figure 28.16** Pulse-delay measurement of chromatic dispersion. When the ends of the test fiber are close together, the reference path (fixed source, reference fiber, and receiver) can be replaced by an electrical cable indicated by the broken line.


laser followed by a pulsed intensity modulator, or a bank of pulsed distributed-feedback lasers of graduated wavelength, with outputs combined by a fiber coupler, or a wavelength-tunable mode-locked laser. In all cases, optical pulses of width less than approximately 500 ps are necessary.

A second pulsed optical trigger source, of fixed wavelength, is used to send a trigger signal to the oscilloscope at the receiving end. The optical trigger source and the reference fiber are necessary only when the ends of the fiber under test are geographically separate. In a loopback measurement, or before the fiber is deployed, the two ends of the fiber under test can be located close together, and the oscilloscope trigger source can be transmitted directly from the pulse generator over a short length of electrical cable.


As the measurement wavelength is changed, changes in the group delay are directly observed as changes in the delay of the measured pulse. After the relative group delay is measured over a set of wavelengths, an analytical function can be fitted to the data and CD can be derived from the derivative of the function, as was shown in Figure 28.15. Accuracy often can be improved by including a variable attenuator after the test fiber output. At each wavelength the attenuator is adjusted to obtain a constant pulse amplitude at the oscilloscope, avoiding any influence of the pulse amplitude on the measurement of delay.

**Modulation phase shift.** The difficulty of generating very short optical pulses can be avoided by measuring instead the phase shift of sinusoidal intensity modulation carried by light of varying wavelength. The relative phase is measured at various wavelengths using equipment such as that shown in the block diagram of Figure 28.17.

The optical test source is a continuous-wave tunable laser followed by a sinusoidal intensity modulator. A second modulated source, of fixed wavelength, establishes a reference phase at the receiving end to allow measurement of the phase variations of the tunable-source modulation. Exactly as in the case of the pulse-delay method,



**Figure 28.17** Modulation phase shift measurement of chromatic dispersion. When the ends of the test fiber are close together, the reference path (fixed source, intensity modulator, reference fiber, and receiver) can be replaced by an electrical cable indicated by the broken line.



**Figure 28.18** Differential phase shift measurement of chromatic dispersion. When the ends of the test fiber are close together, the reference path (fixed source, intensity modulator, reference fiber, and receiver) can be replaced by an electrical cable indicated by the broken line.

the fixed-wavelength source and the reference fiber are necessary only when the ends of the fiber under test are geographically separate. In a loopback measurement, or before the fiber is deployed, the two ends of the fiber under test can be located close together, and reference phase can be transmitted directly from the sinusoidal oscillator over a short length of electrical cable.

As the measurement wavelength is changed, a relative phase shift  $\phi(\lambda)$  is measured by the phase meter as a function of wavelength. A relative group delay  $\tau(\lambda)$  is then calculated from the phase shift according to Equation 28.7. Finally, an analytical function is fitted to the data and CD is derived from the derivative of the function, as was shown in Figure 28.15.

$$\tau(\lambda) = \frac{\phi(\lambda)}{2\pi f L} \quad (28.7)$$

where

- $\tau(\lambda)$  is the group delay as a function of wavelength
- $\phi(\lambda)$  is the radian phase shift as a function of wavelength
- $f$  is the amplitude modulation frequency
- $L$  is length of the fiber under test

**Differential phase shift.** This method (Figure 28.18) is similar to the modulation phase shift method, but includes in addition a direct modulation of the optical wavelength. A second name for the differential phase shift method is *double demodulation*, because two types of modulation are used together: a high-frequency amplitude

modulation of a broadband optical source such as a light-emitting diode, and a low-frequency dithering of a wavelength-selecting device such as a monochromator.

As in the modulation phase shift method, the phase difference of the amplitude modulation is measured between a test fiber and a reference fiber. Again, the reference fiber and reference receiver can be replaced by an electrical cable when both ends of the fiber under test are close together. After an electrical signal proportional to the phase difference is generated by a phase meter, variations in this signal caused by wavelength dithering are detected by a low-frequency lock-in amplifier. A low-frequency oscillator generates a signal that provides the reference for the lock-in and causes the wavelength-selecting device to dither over a range of  $\pm\Delta\lambda/2$  at a center wavelength of  $\bar{\lambda}$ . The CD is then given by

$$D(\bar{\lambda}) = \frac{\Delta\phi}{2\pi f \Delta\lambda L} \quad (28.8)$$

where

- $D(\bar{\lambda})$  is the CD as a function of wavelength
- $\Delta\phi$  is the radian peak-to-peak modulation phase shift
- $f$  is the amplitude modulation frequency
- $\bar{\lambda}$  is the mean wavelength selected
- $\Delta\lambda$  is the peak-to-peak dither of the wavelength selection
- $L$  is length of the fiber under test


In practice, Equation 28.8 must include a multiplicative correction factor to account for the temporal characteristics of the dither waveform and other factors. The constant correction factor can be determined by measurement of a fiber of known CD. This method determines CD directly, without the need for fitting of an analytical function to measured data.

### 28.7.3 Polarization-mode dispersion (PMD)

In single-mode fibers where chromatic dispersion has been carefully compensated, or when operating at a wavelength near the chromatic dispersion minimum of a fiber, the effects of chromatic dispersion may be small enough that polarization-mode dispersion (PMD) becomes a dominant limitation of the maximum bit rate.

The effect of PMD is simplest in certain components that incorporate birefringent crystals. (Examples include optical isolators and electro-optic modulators.) The group delay through such a component depends on the polarization of the optical signal. A detailed analysis of the situation (Poole and Giles, 1988) shows that only two polarizations called the *principal states of polarization* (PSPs) will experience pure group delays, and that the PSPs are orthogonal in the absence of polarization-dependent loss. The difference between these delays is called the *differential group delay* (DGD).

A pulse coinciding with a PSP propagates through the device without suffering first-order broadening. The energy of a more general input pulse, however, is distributed between the PSPs. At the output, the two PSPs arrive at different times and add together at the photodetector (Figure 28.19) to form an electrical output pulse




**Figure 28.19** Pulse splitting caused by polarization-mode dispersion. The two output pulses are polarized along the principal states of polarization, and separated in time by the differential group delay. A polarization-independent receiver responds to the sum of the two pulses, effectively resulting in pulse broadening.

that is broadened or even split into two. The PMD of a simple component is simply the DGD, usually expressed in picoseconds.

In a long fiber or in a concatenation of many fiber components, the pulse-broadening phenomenon is more complicated. In this case, both the DGD and PSPs depend on the optical wavelength. Moreover, the slow drift in birefringences of different sections of the fiber cause both the DGD and PSPs to vary over time, even when measured at a constant center wavelength. For this reason, the PMD of a long fiber, or of a network incorporating long lengths of fiber, is specified as a statistical quantity.

Both the mean DGD and the RMS DGD, i.e., the square root of the mean squared DGD, are used to characterize the statistical distribution of DGD; in this context, both quantities currently are used to define PMD. Using either definition, fiber PMD is found to increase proportionally to the square root of the fiber's length, so PMD is specified in  $ps/\sqrt{km}$ . The PSPs are assumed to vary uniformly over all states of polarization.

Pulse distortion and broadening are different for a long fiber versus for a simple birefringent component. After transmission through a long fiber, pulses much shorter than the mean or RMS DGD are broken up into a distribution of pulses with an envelope that is roughly Gaussian. As shown in Figure 28.20, broader input pulses result in coarser variation of the output amplitude. When the input pulse width is greater



**Figure 28.20** Pulse broadening caused by a highly mode-coupled long fiber. Input pulses of increasing duration yield output pulses that progress from a distribution of narrow pulses to a cleanly separated pair of output pulses.

than approximately the mean DGD, the output pulses are not broadened, but rather are split into two pulses resolved along the PSPs. As rule of thumb, acceptable transmission in digital networks requires the DGD to be less than one-tenth of a bit period. Under this condition, the pulse-splitting, principal-states model is clearly appropriate.

**Wavelength scanning with extrema counting.** In the most common wavelength scanning setup, broadband polarized light is transmitted through the component or network under test, as shown in Figure 28.21a. Optical power through an analyzer at the output is measured as a function of wavelength using an optical spectrum analyzer, i.e., a tunable optical filter followed by an optical power meter. The measured power is observed to describe a pattern of ripples as a function of wavelength (Figure 28.21b). The measurement principle of this technique, somewhat simplified, is that the mean DGD is proportional to the density of the measured ripples.

As the power spectrum of the source may itself exhibit ripples independent of any PMD, we first measure the source power spectrum so that in later steps we can compensate for its variations. Conceptually, the simplest way to measure the source power spectrum is to replace the component or fiber under test with a short fiber jumper cable with negligible PMD, and to measure its output power spectrum directly without using the output analyzer.

A second, preferred method is to measure two power spectra at the output of the analyzer. The first power spectrum  $H(\lambda)$  is measured with the analyzer oriented horizontally, and the second  $V(\lambda)$  is measured with the analyzer oriented vertically. A transmission function  $A(\lambda)$  is then calculated by normalization:

$$A(\lambda) = H(\lambda)/(H(\lambda)+V(\lambda)) \text{ or } A(\lambda) = V(\lambda)/(H(\lambda)+V(\lambda)) \quad (28.9)$$

where

$A(\lambda)$  is the transmission function

$H(\lambda)$  is the power spectrum through a horizontal polarizer

$V(\lambda)$  is the power spectrum through a vertical polarizer

This normalization technique is able to correct for both source spectral variation and variation in spectral loss of the fiber.

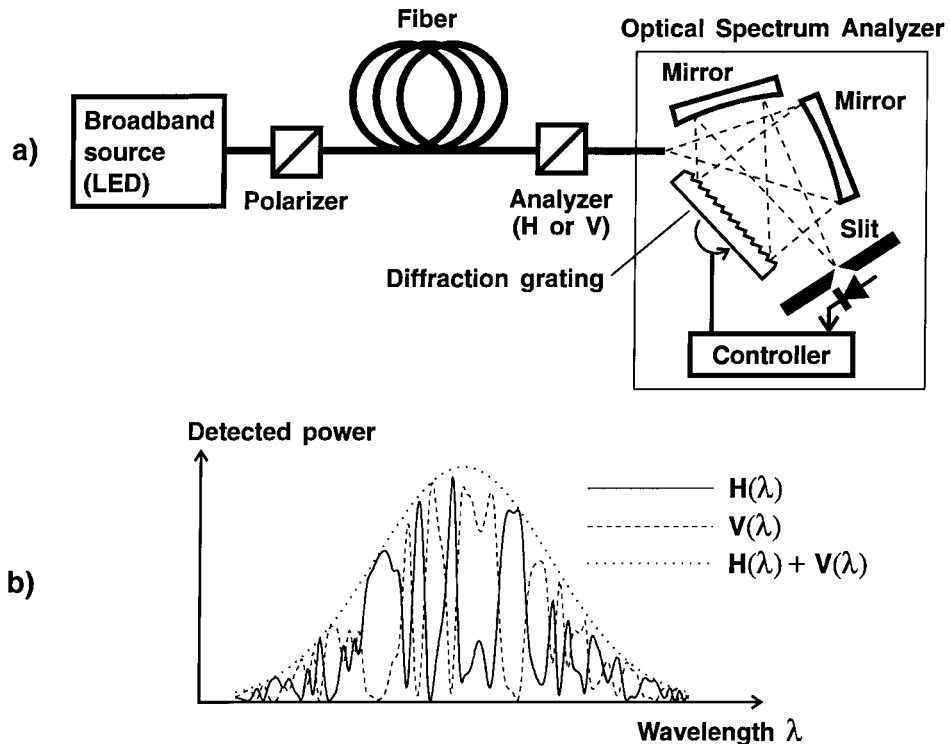
The mean DGD is calculated by counting the number of extrema (minima and maxima)  $N_e$  of the transmission function  $A(\lambda)$  between the wavelengths  $\lambda_1$  and  $\lambda_2$  at the boundaries of the measurement. The mean DGD  $\Delta\tau_{mean}$  is then given by

$$\Delta\tau_{mean} = k \frac{N_e}{2C} \frac{\lambda_1 \lambda_2}{\lambda_2 - \lambda_1} \quad (28.10)$$

where

$\Delta\tau_{mean}$  is the group delay as a function of wavelength

$k$  is a constant depending on polarization mode coupling



**Figure 28.21** Measurement of PMD by wavelength scanning: (a) block diagram, and (b) measured spectra.

## 672 Network Test Instrumentation

- $N_e$  is the number of extrema between  $\lambda_1$  and  $\lambda_2$   
 $\lambda_1$  is the minimum measured wavelength  
 $\lambda_2$  is the maximum measured wavelength

The constant  $k$  depends on the degree of polarization mode coupling in the fiber or component under test:  $k = 1$  for components with no mode coupling, while highly mode-coupled long fibers require  $k = 0.824$ .

**Wavelength scanning with Fourier transformation.** Another way of analyzing the transmission function  $A(\lambda)$  obtained by wavelength scanning is to employ the Fourier transform. This yields a function of time  $t$  from a function of optical frequency  $\nu$ :

$$f(t) = \int_{-\infty}^{\infty} F(\nu) e^{i2\pi\nu t} d\nu \quad (28.11)$$

Typically  $A(\lambda)$  is evaluated at a series of discrete wavelengths, and a function of frequency then is created by calculating a frequency  $\nu_n = c/\lambda_n$  corresponding to each wavelength  $\lambda_n$ , where  $c$  is the speed of light. Since a discrete Fourier transform requires a function sampled at equally spaced frequencies, interpolation of  $A(\nu)$  may be necessary to obtain points at equal frequency intervals.

The transmission function  $A(\nu_n)$  once expressed at constant frequency intervals, is transformed into a time sequence  $a(t_n)$  by the discrete Fourier transform:

$$a(t_n) = \sum_{k=0}^{N-1} A(\nu_k) W(\nu_k) e^{i2\pi\nu_k t_n/N} \quad (28.12)$$

where


- $a(t_n)$  is the desired time sequence  
 $A(\nu_k)$  is the sampled transmission function  
 $W(\nu_k)$  is a weighting function  
 $N$  is the number of samples in the transmission function

$W(\nu_k)$  is typically a Hanning window  $W(\nu_k) = \cos^2(2\pi k/N - \pi)$ , and is used to limit the spurious effects of the sharp boundaries of the transmission function. Because  $A(\nu_k)$  is an arbitrary real function,  $a(t_n)$  generally is complex, but the magnitude  $|a(t_n)|$  is an even, real function centered at  $t = 0$ .  $A(\nu_n)$  and  $|a(t_n)|$  are shown in Figure 28.22 for the case of extensive mode coupling.

A value of PMD is derived from the time series  $|a(t_n)|$  by calculating the square root of its second moment, i.e., its RMS width. This calculation is complicated by the presence of a large spike at  $t = 0$  and by noise at large values of  $t$ . Accordingly, the limits of summation for the second moment calculation are chosen to exclude the regions of  $|a(t_n)|$  where these effects dominate. Typically a noise threshold is chosen equal to 0.05 times the maximum value of  $|a(t_n)|$  excluding the spike at  $t = 0$ . The summation is carried out from the index  $n = 2$  to the index  $n_{thresh}$  at which  $|a(t_n)|$  last exceeds the noise threshold:

$$\Delta\tau_{rms} = c \left( \frac{\sum_{n=2}^{n_{thresh}} |a(t_n)|^2}{\sum_{n=2}^{n_{thresh}} |a(t_n)|} \right)^{1/2} \quad (28.13)$$





**Figure 28.22** A typical wavelength-scanning transmission function and the magnitude of its Fourier transform.

where

$\Delta\tau_{rms}$  is the RMS DGD

$a(t_n)$  is the time sequence calculated by Equation 28.12

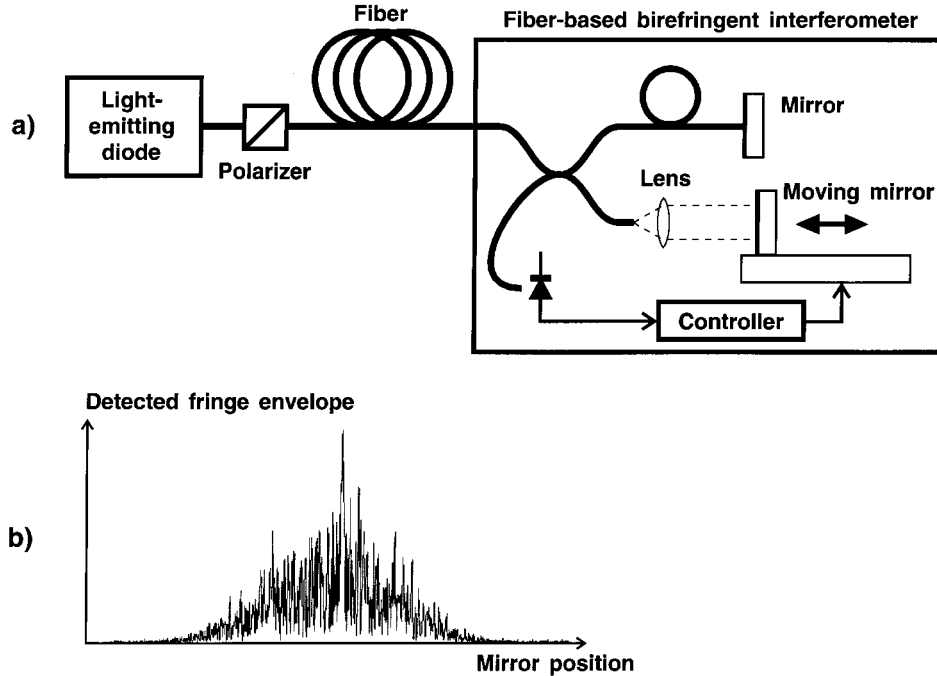
$c$  is a constant depending on polarization mode coupling

$n_{thresh}$  is the index of  $|a(t_n)|$  given by the noise threshold

The constant  $c$  depends on the degree of polarization-mode coupling:  $c = 1$  for a simple birefringence with no mode coupling, while  $c = \sqrt{3/4} \approx 0.866$  for the limiting case of extensive mode coupling in a long fiber.

**Interferometry.** PMD can be measured using an optical interferometer with a variable differential time delay (Figure 28.23). Broadband polarized light is transmitted through the fiber or device under test. At the output the light is analyzed by a variable-delay interferometer, the two arms of which exhibit differing birefringences. A conventional non-birefringent interferometer can also be used if a polarizer is included at the input to the interferometer. A Michelson interferometer is illustrated, but other interferometer configurations can be used.

The principle of measurement is based on observing changes to the coherence function of broadband light. A *coherence function* is a description of the amplitude and phase of interference fringes as a function of the differential delay of an interferometer. A broadband source such as a light-emitting diode has a narrow coherence function centered at zero delay. It typically falls to less than 10 percent of its peak value for delays greater than approximately 50 femtoseconds. This narrow coherence function is changed by transmission through a device exhibiting PMD.



**Figure 28.23** Measurement of PMD by interferometry: (a) block diagram, and (b) result of measurement.

A simple birefringent device without polarization mode coupling leads to a coherence function at the output with peaks on both sides of the zero-delay peak. The delay between either side-peak and the central peak is the differential group delay of the device. When the device is a highly mode-coupled fiber or network, the coherence function amplitude  $\gamma(t_n)$  at the output is an approximately Gaussian function multiplied by noise.

Just as for the method of wavelength scanning with Fourier transformation, a value of PMD is derived from the coherence function amplitude  $\gamma(t_n)$  by calculating the square root of its second moment, i.e., its RMS width. In this measurement method the zero-delay point generally is not known sufficiently accurately beforehand, so calculation of the RMS width must include a term to account for shifts in the center position  $t_{center}$  of  $\gamma(t_n)$ . The center position is first calculated:

$$t_{center} = \frac{\sum_{n=1}^N t_n \gamma(\tau_n)}{\sum_{n=1}^N \gamma(\tau_n)} \quad (28.14)$$

where

- $t_{center}$  is the time at the center of the coherence function
- $t_n$  are the times at which the coherence function is sampled
- $\gamma(t_n)$  is the sampled coherence function amplitude
- $N$  is the total number of samples

It is again necessary to reject the effects of the peak at zero delay and of noise at large delays. Limit indices  $k1, k2, k3,$  and  $k4$  are chosen such that the relative effect of noise is small for  $t_n$  within the interval  $k1 < n < k4$ , and the edges of the central peak are located at  $t_{k2}$  and  $t_{k3}$ . The PMD is then given by

$$\Delta\tau_{rms} = \frac{c}{2} \left( \frac{\sum_{n=k1}^{k2} (t_n - t_{center})^2 \gamma(t_n)}{\sum_{n=k1}^{k2} \gamma(t_n)} \right)^{1/2} + \frac{c}{2} \left( \frac{\sum_{n=k3}^{k4} (t_n - t_{center})^2 \gamma(t_n)}{\sum_{n=k3}^{k4} \gamma(t_n)} \right)^1 \quad (28.15)$$

where

$\Delta\tau_{rms}$  is the RMS DGD

$t_{center}$  is the time at the center of the coherence function

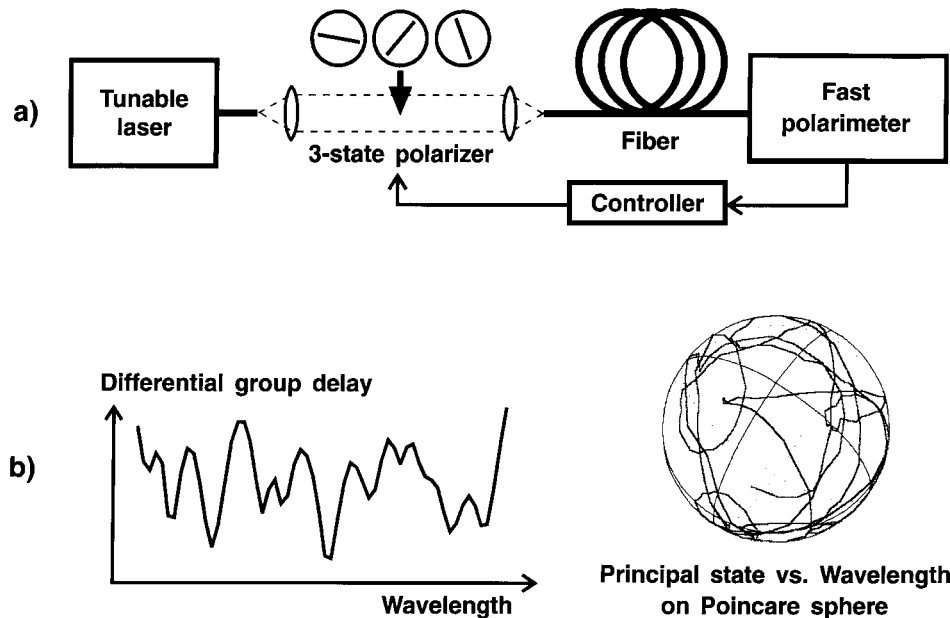
$t_n$  are the times at which the coherence function is sampled

$\gamma(t_n)$  is the sampled coherence function amplitude

$c$  is a constant depending on polarization mode coupling

Here the constant  $c$  accounts for varying degrees of polarization mode coupling:  $c = \sqrt{3}/4 \approx 0.866$  for the case of a highly mode-coupled system, while  $c$  may take values of up to unity for lesser degrees of mode coupling.

**Jones matrix eigenanalysis.** The most complete characterization of PMD is obtained by measuring the Jones matrix of a component or network at a series of optical wavelengths. A block diagram of this measurement (Figure 28.24a) is similar to



**Figure 28.24** Measurement of PMD by Jones matrix eigenanalysis: (a) block diagram, and (b) result of measurements showing variations in both differential group delay and principal states of polarization.

that used for a matrix measurement of polarization-dependent loss (Figure 28.7), with the addition of a means for controlling the measurement wavelength. Often the wavelength is controlled by using a tunable laser as the light source.

The concepts of eigenvalues and eigenvectors of a matrix are integral to this measurement method. An eigenvector  $\mathbf{v}$  of a matrix  $\mathbf{A}$ , and its associated eigenvalue  $\rho$ , satisfy the relation  $\mathbf{A}\mathbf{v} = \rho\mathbf{v}$ . In other words, the linear transformation represented by  $\mathbf{A}$  has the same effect on  $\mathbf{v}$  as simple rescaling of  $\mathbf{v}$  by the factor  $\rho$ . Jones matrices are  $2 \times 2$  complex matrices. The Jones matrices of interest for purposes of PMD measurement have two distinct eigenvalues and two distinct eigenvectors.

It can be shown (Heffner, 1992) that, given a Jones matrix  $\mathbf{T}_1$  representing the polarization transformation caused by transmission through the network at a first wavelength  $\lambda_1$  and a second matrix  $\mathbf{T}_2$  representing transmission at a second wavelength  $\lambda_2$ , that:

1. The eigenvectors of the matrix product  $\mathbf{T}_1\mathbf{T}_2^{-1}$  are Jones vectors representing the output principal states over the wavelength interval  $\lambda_1$  to  $\lambda_2$ .
2. The eigenvectors of the matrix product  $\mathbf{T}_1^{-1}\mathbf{T}_2$  are Jones vectors representing the input principal states over the wavelength interval  $\lambda_1$  to  $\lambda_2$ .
3. The eigenvalues  $\rho_1$  and  $\rho_2$  of the matrix product  $\mathbf{T}_1\mathbf{T}_2^{-1}$  are related to the DGD  $\Delta\tau$  over the wavelength interval  $\lambda_1$  to  $\lambda_2$  by

$$\Delta\tau = \frac{\text{Arg}\left(\frac{\rho_1}{\rho_2}\right)}{\Delta\omega} \quad (28.16)$$

where

$\Delta\tau$  is the DGD

$\rho_1, \rho_2$  are the eigenvalues of  $\mathbf{T}_1\mathbf{T}_2^{-1}$

$\Delta\omega$  is the optical radian frequency interval

$\Delta\omega = 2\pi c(\lambda_2 - \lambda_1)/\lambda_1\lambda_2$  is the difference in optical radian frequency between the measurement wavelengths, and  $\text{Arg}(x)$  is the complex polar angle of  $x$ , i.e.  $b = \text{Arg}(ae^{ib})$ .

In practice, Jones matrices are automatically measured at a series of wavelengths, and eigenanalysis of adjacent pairs of matrices yields the DGD  $\Delta\tau_n$  as a function of wavelength as shown in Figure 28.24b. Once a distribution of  $\Delta\tau_n$  has been measured over a range of wavelengths, mean and RMS values of DGD can be directly calculated:

$$\Delta\tau_{mean} = \frac{1}{N} \sum_{n=1}^N \Delta\tau_n \quad (28.17a)$$

$$\Delta\tau_{rms} = \left[ \frac{1}{N} \sum_{n=1}^N \Delta\tau_n^2 - \left( \frac{1}{N} \sum_{n=1}^N \Delta\tau_n \right)^2 \right]^{1/2} \quad (28.17b)$$

where

$\Delta\tau_{mean}$  is the mean DGD

$\Delta\tau_{rms}$  is the rms DGD

$\Delta\tau_n$  is the DGD at a given wavelength  
 $N$  is the total number of known DGDs

## 28.8 References

- Driscoll, W.G., and Vaughan, W. *Handbook of Optics*. (New York: McGraw-Hill, 1978.) See section 4.
- Kliger, D.S., Lewis, J.W., and Randall, C.E. *Polarized Light in Optics and Spectroscopy*. (New York: Academic Press, 1990.)
- Collet, E. *Polarized Light*. (New York: Marcel Dekker, 1993.)
- Nyman, B.M. and Wolter, G. *IEEE Photonics Technology Letters*, 5 (1993): 817–818.
- Nazarathy, M., Newton, S.A., Giffard, R.P., Moberly, D.S., Sischka, F., Trutna, Jr., W.R., and Foster, S. *IEEE/OSA Journal of Lightwave Technology*, LT-7 (1989): 24–38. Also references therein.
- Keiser, G. *Optical Fiber Communications*. (New York: McGraw-Hill, 1983.) Includes a simple, informative treatment of fiber modes.
- Cohen, L.G. *IEEE/OSA Journal of Lightwave Technology*, LT-3 (1985): 958–966.
- Pelayo, J., Paniello, J., and Villuendas, F. *IEEE/OSA Journal of Lightwave Technology*, LT-6 (1988): 1861–1865.
- Poole, C.D., and Giles, C.R. *Optics Letters*, 13 (1988): 155–157.
- Heffner, B.L. *IEEE Photonics Technology Letters*, 4 (1992): 1066–1069.



---

Chapter  
**29**

## Distributed Network Monitoring

**Colin Wynd**

*Hewlett-Packard, Colorado Springs, Colorado*


### 29.1 Introduction

As a network grows and spreads, the number of segments multiplies. Geographically dispersed segments then are interconnected, making it increasingly difficult for a small staff to manage an enterprise-wide network with only portable analyzers. Continuing competitive pressures mean that corporations must find more efficient methods of monitoring their networks. Distributed network monitoring and analysis products are available to facilitate successful management of medium and large enterprise-wide networks. These products enable network administrators to use the network to manage the network through methodology that provides visibility into remote segments without the need to be physically present.

Network environments have evolved and matured to the point that focus has shifted from technology adoption and implementation issues to management activities needed to support users in the mission-critical environment. As a result, the strategic role that Information Technology (IT) plays in delivering business continuity across the enterprise has been elevated. The elevation of network monitoring as a critical component or function of IT departments can be attributed to three major causes.

**The widening manageability gap.** Figure 29.1 shows that growth in network segments from 1993 to 1997, as projected by 100 managers of large companies, was expected to average 115 percent. At the same time, LAN management budget growth was projected to grow at only 66 percent. This “manageability gap” is forcing IT departments to rethink their network management processes and to use distributed monitoring and analysis tools to allow fewer network staff to monitor and manage a large, growing network.

**The changing face of the business environment.** Corporations are becoming more geographically dispersed. Entities that once were autonomous are working closely together. The network has become a critical method for communicating be-



**Figure 29.1** Infonetics Research, Inc., measured the real costs of management twice in a recent four-year period. The 100 survey respondents also were asked to project 1997 figures for the number of LAN segments they would be managing and for the LAN management budget. Results support the need for cost-effective and time-saving products to help manage all aspects of the enterprise environment.

tween distant business units in a company. Most business applications are client-server instead of being located on a central mainframe. Client-server networking divides the execution of a unit of work between activities initiated by an end user or program (client), and the resource containing data banks or processing power to respond to the activity request (server). A reliable, operational client-server environment is a critical tool for end users who must perform daily tasks vital to the profitability of the organization.

The increased focus on work/life balance has increased the number of work-at-home participants accessing client-server environments. The result is a steadily increasing amount of traffic on the network.

**Growing user expectations.** In the client-server environment, users are starting to expect error-free network connectivity with guaranteed uptime and acceptable response time. Users, to whom the growing complexity of the network is transparent, also expect network services to be delivered regardless of the underlying technology. Business users expect secure data transmission. This increase in reliance on client-server applications as a fundamental part of conducting business means that end users rely on the network as much as (if not more than) they do the telephone. Whether the IT department is ready or not, the network is expected to be as reliable as the phone system.

This chapter introduces the role that distributed network monitoring plays in network administration. Distributed network monitoring will be explained first, followed by the fit with IT management. Finally, the range of functionality that network monitoring brings to the IT manager's arsenal will be covered.



## 29.2 Distributed Network Monitoring Defined

Distributed network monitoring, which encompasses the ability to view and perform monitoring and analysis on a remote network as if it were local, has two components: data collectors and software to process the collected data.

**Data collectors.** Called *agents* or *probes*, data collection devices can be hardware- or software-based. Each segment of a network must be equipped with an agent or probe if network monitoring and analysis application software is to be used to view and manage traffic on the segments. Data collectors, which can be standalone hardware boxes that physically attach to a network segment, are dedicated to monitoring, collecting, and storing communication data. Data collectors also can be embedded in other types of network communication devices, such as routers and hubs. Software-based data collectors can be installed on workstations attached to a network segment, using processing and storage resources of the workstation to monitor, collect, and store data.

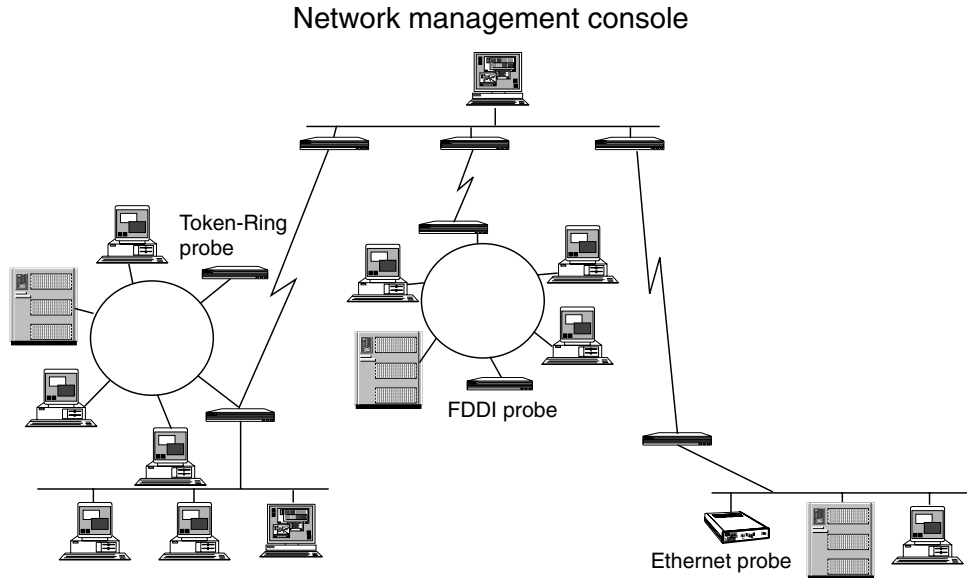
Data collectors can operate by continuously gathering data, storage-intensive activity in which the speed of processing must equal the speed of the communication link. Data collectors also can operate by using statistical sampling methodology, allowing storage capacity to be conserved, and reducing the need for rate of collection to match rate of communication.

Network segments can be of any media type, ranging from local area network (LAN) types such as Ethernet, FDDI, and Token-Ring, to wide area network (WAN) protocols such as frame relay. The segments can be dispersed geographically, but generally they must be interconnected. Distributing data-collection power to the extremities of the network allows network connections to be used to access data pertinent to any segment of a network equipped with a data collector.

Before distributed monitoring, portable devices were carried to remote sites and attached to the network when problems were occurring on a particular segment. Carrying a network monitoring device to a segment only when problems occurred meant that the segment was not being monitored 99 percent of the time. Distributed data collectors placed permanently on mission-critical segments constantly monitor traffic. This enables traffic activity from any one segment to be accessed at any time, and traffic from any combination of segments to be aggregated for a holistic view of network activity.

**Network monitoring and analysis software.** Application software resides at a central processing point—a network management console—where data can be pulled or deposited from the distributed data collectors when needed or as scheduled. The applications analyze and interpret the data, extracting information useful for network traffic visibility, fault management, performance management, capacity planning, network access monitoring, and network availability. Graphical user interfaces (GUI) present the information in easy-to-read charts, maps, tables, and graphs. The network management console communicates with agents using the same network that the agents are monitoring. Out-of-band communication between the manager and agents also is possible via modem link.

Figure 29.2 shows an example of remote monitoring agents installed on a large network and a centralized network management console.



**Figure 29.2** FDDI, Ethernet, and Token-Ring segments are all monitored with data collectors. Data then is sent to, or accessed by, a centralized network management console using network monitoring and analysis application software. This specialized software graphically interprets the data, facilitating decision-making on matters such as balancing traffic loads, resegmenting the network, and adding network resources when needed.

### 29.2.1 Information available from data collectors

Monitoring the network means that information on every single packet on every single segment can be monitored, collected, and stored. Data collectors bank the information according to a set of industry standards; they sort, count, and store various pieces of information that will prove useful for some aspect of network management.

With the limited storage capacity inherent in most data collectors, it's crucial to decide which data is important and should be collected and which data is irrelevant to the task at hand. Corporations with many of segments need to prioritize the pieces of information critical to IT; otherwise they quickly become inundated with unnecessary data. This results in the cost of network analysis potentially exceeding the actual cost of the network. Some of the most important measurements that should be gathered are:

- Utilization
- Protocol distribution
- Top talkers
- Error rates
- Latency measurements (echo tests)

**Utilization.** Given a limited amount of bandwidth on each segment, data collectors monitor use of the bandwidth over time. Then the probes or agents provide

utilization data for the network monitoring and analysis applications to see total usage, breaking out usage into various categories for more detailed analysis. This analysis reveals utilization trends and is useful for network capacity planning and for baselining. It can indicate opportunities for performance improvement, highlighting cycles of activity and identifying traffic sources and destinations.

**Protocol distribution.** At the same time that utilization is being monitored and recorded, components of the traffic are being characterized. Across all seven layers of the OSI model, data on protocols in use on a network segment are gathered and made available for analysis. This information can be used to profile the types of protocols that make up traffic at a point in time or over a period of time. This is useful for examining trends of application use on the network, changing application mixes, monitoring use of new applications, and gauging the effect of new applications on the network.

**Top talkers.** While continuously monitoring the network, data collectors are counting the amount of traffic generated by each node attached to the segment. The counts can be accessed to reveal an ordered list of nodes that are doing the most “talking” on a particular network segment. Visibility into details of each user’s traffic, such as destination node and protocol use, then can be viewed. Identifying these users lets you pinpoint probable causes for performance problems and slow response of the network. The top talkers statistic proves useful for making resegmenting decisions and for load balancing. Top talkers also can indicate potential new applications that are unknown to the network department.

**Error rates.** Just as measures of utilization, protocol breakdown, and traffic generated by each node are gathered, error rates related to network traffic are accumulated. Network monitoring and analysis applications then can view and analyze error information, using the results for many aspects of network management. Total error-rate information serves as a network performance indicator; lots of errors may indicate an impact on response time of the network or on availability of network resources. Baselining the error rate of the network provides measurements for setting thresholds and alarms. Error rates, when correlated with utilization, can indicate potential physical-layer network problems.

**Latency measurements (echo tests).** The ability to test the reachability and response time of network nodes with echo monitoring is a primary indicator of resource availability and network performance.

### 29.2.2 Distributed network information use

Network administrators can use distributed remote monitoring tools to manage all the segments that comprise an enterprise network. Descriptions of some of the functions that the network administrator performs follow.

Network performance management consists of baselining typical network behavior, analyzing application usage by protocol, characterizing client-server activity, and monitoring internetwork traffic and trends. These activities include the ability to set network thresholds to identify anomalies and to create baselines to aid in determining “normal” network performance. This allows network managers to make informed optimization decisions based on many facets of internetwork operation, decisions that can help maximize return on a network infrastructure investment.

**684 Network Test Instrumentation**

Network security monitoring identifies activity on network resources that can be used to secure company-sensitive information. It brings IT the ability to:

- Monitor source and destination of data communications in conjunction with data collectors.
- Identify applications being used.
- Monitor what resources are accessed.
- Capture signatures of potential intruders.
- Trigger packet capture on unauthorized access so conversations can be examined in detail.

These activities can ensure that only authorized users access the network.

*Fault management and availability* means quick responsiveness and timely action to expedite recovery when the network fails. When faced with immediate problems, quick access to network data relevant to the segment, or segments, of the network in distress helps operators understand and deal with the problem. The ability to set thresholds for critical network elements and to identify and view performance indicators is the first step in the tuning or recovery process. In addition, detailed packet interrogation is crucial for a more granular analysis, often needed to solve obstinate or elusive problems.

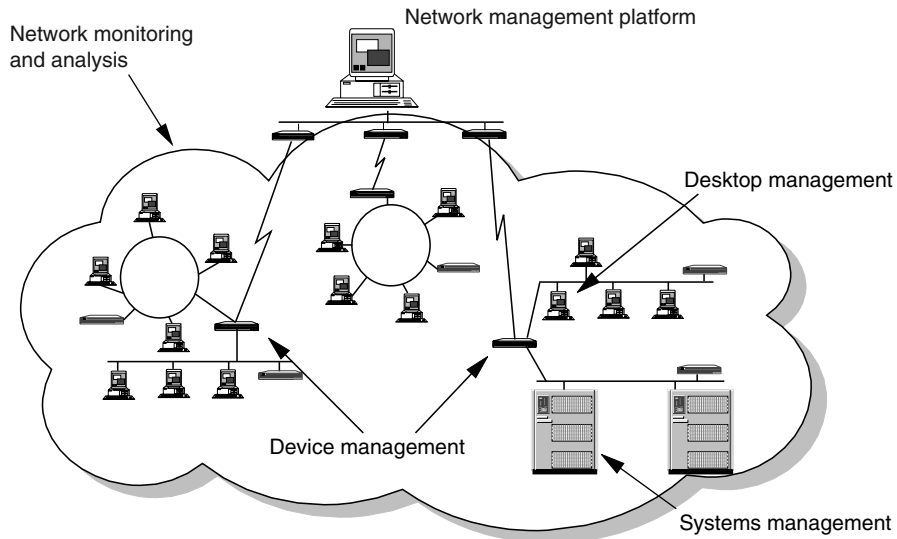
*Capacity planning* relies on the availability of historical information to provide crucial facts about network operation. It also allows quick what-if analysis using traffic profile modeling for balancing or reconfiguring network resources. This information is invaluable for justifying the appropriation of IT funds to maintain and improve delivery of network service for the sustained competitive edge that a network provides a business. Whether purchasing equipment to grow or optimize network needs, or helping justify the expense of moving to a new technology, network monitoring and analysis equips you to make informed decisions.

**29.2.3 Distributed network monitoring and IT management**

The IT management environment, as shown in Figure 29.3, covers the range of devices that reside on the network. It encompasses all issues that surround the business applications that enable business end users to function. The environment breaks down into four components:

- Device management
- Desktop management
- Systems management
- Network monitoring and analysis

*Device management* is concerned with various networking devices such as bridges, routers, and hubs that interconnect the segments of a network. Typical management issues deal with configuration tables, throughput, link states, and port



**Figure 29.3** The network environment that IT must manage is large and diverse. One strategy for managing this environment is to address network devices that make up the infrastructure of the network, desktop devices that use the network, computer systems that provide storage and processing power for network users, and network monitoring and analysis that help keep network communication flowing.

partitioning. A device management application usually shows a picture of the device on the screen, complete with installed cards and indicator lights.

*Desktop management* is concerned with end-user workstations and personal computers, typically the clients in a client-server environment. The management issues are PC configuration files, disk use, and application support.

*Systems management* is concerned with the performance of the computers on the network. The focus of this issue is database performance and disk use on file servers.

*Network monitoring and analysis* primarily is concerned with communication activity on all segments that comprise the network. It looks at the flow of data across the network in an effort to understand network performance issues, to investigate capacity issues, and to resolve problems related to networking protocols. How fast is the server responding to the client's request? Are too many users trying to access one server? What mix of protocols are using a single segment?

#### 29.2.4 Persons responsible for network monitoring

Network monitoring and analysis crosses strategic and tactical boundaries, with many people involved in decision-making and implementation. Some generic descriptions of job titles and functions follow.

The *network manager* usually is responsible for long-term strategic decisions regarding the network, is involved in looking at new technologies such as 100Base-X or ATM, and deciding where and when to modify bandwidth. A network manager looks at network trends and performs forecasting and capacity planning.

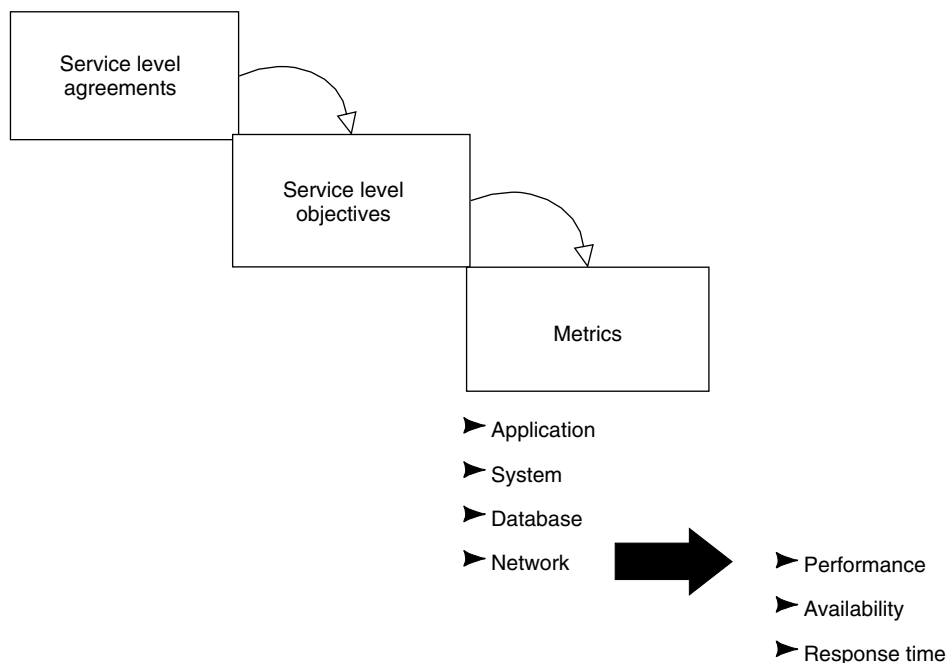
The *network engineer* typically is responsible for day-to-day operation of the network, including upgrading network devices and adding capacity to network devices and interconnects. A network engineer also acts as a second line of support for problems that the operations center engineer cannot resolve.

The *operations center engineer* is the first line of support for network problems and is constantly in a reactive mode. Most large corporations have a centralized network monitoring center staffed with “level-one” engineers who attempt basic troubleshooting on network problems. These engineers monitor for events triggered by servers, workstations, or network devices with built-in ability to alert the operations center of potential problems.

### 29.3 Service Level Management in the Enterprise Environment

*Service level management (SLM)* is the strategy of defining, controlling, and maintaining desired levels of IT service for the business end user. SLM encompasses a wide range of activities: acquiring information to determine appropriate service levels, monitoring for compliance, and providing continuous user feedback regarding IT performance.

Business end users, together with the IT department, define the level of service needed to stay productive based on their level of network reliance. The level of service is turned into a set of objectives that IT can then monitor. Figure 29.4 shows the



**Figure 29.4** Service level agreements between IT and the end user result in metrics being set for service components critical to network success. Some metrics that can be applied to network monitoring and analysis are measures that deal with performance, availability, and response time.

flow of activity that turns service level agreements between IT and end users into objectives and, finally, into metrics that allow IT to monitor compliance.

Within the IT management environment (see Figure 29.3), objectives should be considered for managing all components: device, desktop, system, and network monitoring and analysis. The metrics to measure success in these areas usually concern the applications and databases that contain the information end users require for business transactions. Metrics also relate directly to systems management, measuring the computing power and storage capacity required by applications and databases.

Finally, metrics for network monitoring and analysis allow IT to measure the health of the network and its ability to deliver reliable, responsive business information. Figure 29.4 shows how the service level agreements filter down to deliverables for IT through metrics for network applications, databases, systems, and the network. This chapter focuses on network issues.

### 29.3.1 Distributed monitoring helps deliver on service level agreements

As part of an IT department's *service level agreement* (SLA) with its business end users, IT must maintain a level of network service that satisfies users. To do this, the network must be monitored to ensure error-free connectivity and acceptable responsiveness. If the network were not monitored, it would be impossible for the IT department to guarantee compliance to any agreed-upon level of service.

New client-server applications are springing up in business environments. Some examples of burgeoning applications are the ever-pervasive Internet and electronic mail. If the network is not monitored, the effect of adding even one of these network-intensive applications is unknown. It is entirely possible that electronic mail alone eventually could bring the network to its knees. If the environment is monitored, the utilization of network bandwidth can be determined and traffic trends analyzed to ensure that bandwidth meets or exceeds future growth of resources and applications.

The ability to monitor trends changes IT from being reactive (waiting until something breaks before resolving the problem) to being proactive (resolving potential issues before they cause a break). The IT department should blend into the background, allowing business end users to focus on their own function.

### 29.3.2 How distributed monitoring/analysis supports SLM

Distributed network monitoring and analysis helps position IT to furnish the business and end users with reliable network performance and consistent service. It can help IT focus on delivering dependable and responsive networks with capabilities that support service-level management strategies.

**Network performance management.** Performance management means being able to monitor segment activity as well as to analyze intrasegment traffic. Network managers must be able to examine traffic patterns by source, destination, conversations, protocol/application type, and segment statistics such as utilization and error rates. Network managers must define the performance goals, how notification of

performance problems should be processed, and define the network tolerance levels. Some objectives facing network managers include the following:

- Baselineing and network trending
- Application usage and analysis
- Internetwork perspective
- Data correlation

*Baselineing and network trending* can help determine the true operating envelope for the network by defining certain measurements (segment utilization, error rate, network latency) to check on IT adherence to service level objectives. These tactics also provide the ability to check on out-of-norm conditions that, if unchecked, might have drastic consequences on networked business users' productivity.

*Application usage and analysis* can help network managers answer questions such as what is the overall load of Internet traffic on the corporate network, or what times of day do certain applications load the network. This capability allows network managers to discover important performance information (either real-time or historic) that will help define performance service level objectives for applications in the client-server environment.

*Internetwork perspective* can help discover traffic rates between subnets and reveals which nodes are using WAN links to communicate. This capability is useful when traffic between remote sites and interconnect devices is critical to the normal operation of the business. It also can help define "typical" rates between interconnect devices. Internetwork perspective shows how certain applications use the critical interconnect paths and defines "normal" WAN use for applications.

*Data correlation* allows peak network usage points to be selected throughout the day and discovers:

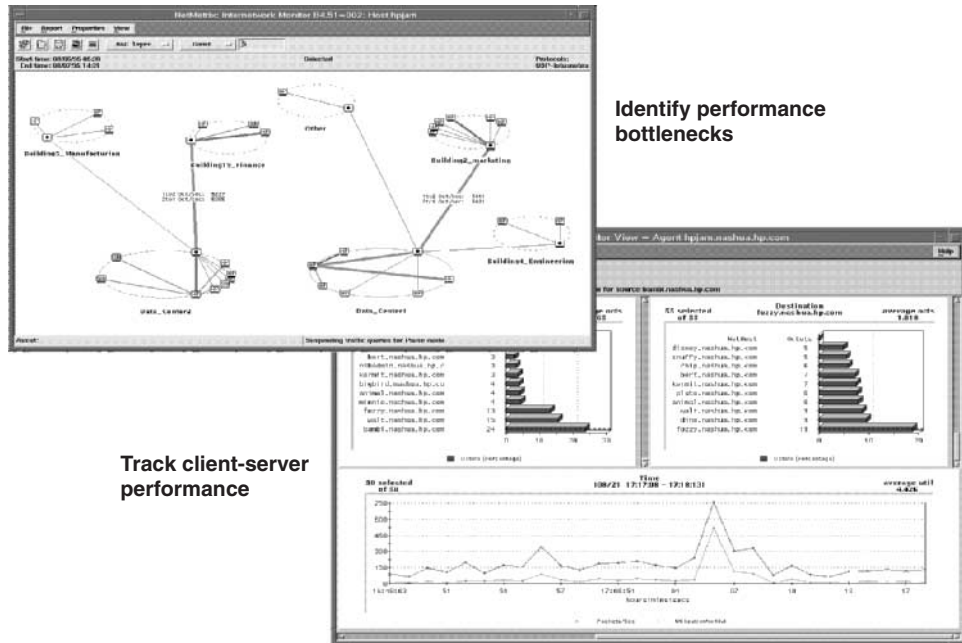
- Which nodes are contributing to network load at that peak point in time.
- To which nodes they were sending traffic.
- Which applications were running between them.

The upper-left screen shot in Figure 29.5 shows an example of the internetwork perspective for performance management by depicting actual traffic flow between several segments of a network. The thickness of the lines indicates the volume of traffic between those two end points. With this information it is easy to identify potential WAN bottlenecks.

The bottom-right screen shot shows how data correlation for clients and servers is associated with a time graph. The ability to determine how much one particular server affects the network can help position that critical network resource.

**Access level (security) monitoring.** Security management encompasses a broad set of network access control policies that span network hosts, network elements, and network access points (firewalls). Consistent policies are the key here; the objective is to support access and connectivity appropriate to the business need, while re-





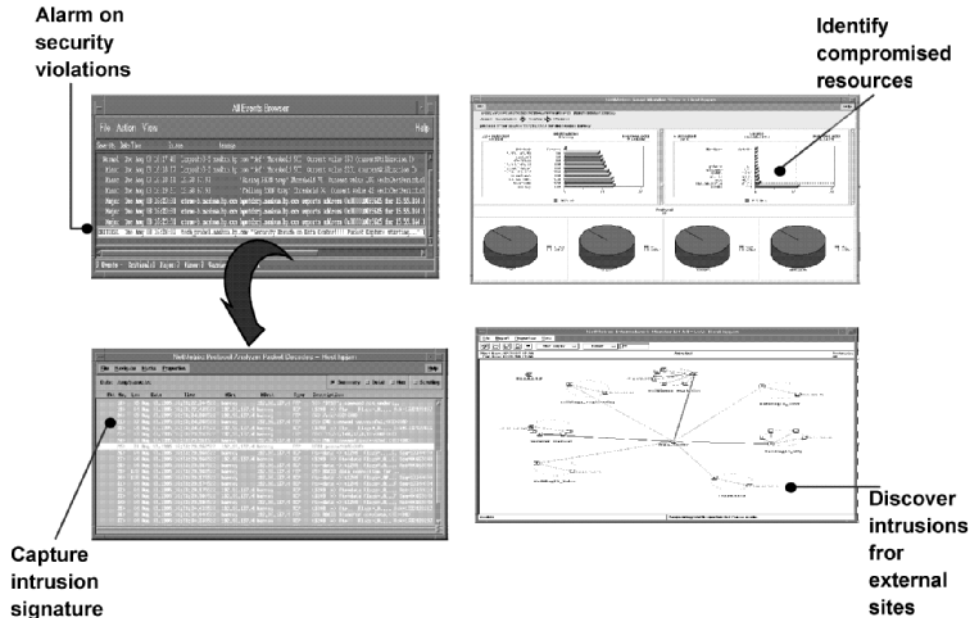
**Figure 29.5** Two aspects of performance management that support service level management strategies are illustrated here. The upper left is an example of internetwork perspective, showing the traffic activity between the various segments of an internetwork. This can indicate opportunities to relocate highly used resources to “quieter” segments. The graph on the lower right is an example of correlating a chosen source and examining all destinations in conversation with that source. Then a graph of utilization over time can be displayed to see the spread of activity between that source and a particular destination over the day. This might provide clues to peak activity times.

stricting clearly inappropriate network-based access. As in other activities, constant monitoring for specific violations is critical, as is a notification mechanism. For certain conditions, immediate automatic action may be required (i.e., “shut down this connection,” or “shut down the firewall”). Monitoring should include both passive and active monitoring (probing).

Access level monitoring ensures that the controls and security in place are performing to expectations. Monitoring the traffic flow to and from a firewall, for instance, ensures that no intruders are internally accessing the network. Access level monitoring polices the “police” and ensures that nothing has been overlooked by security.

Figure 29.6 shows how network access monitoring can help a company protect business-sensitive information. Many books discuss security management in finer detail than is possible here.

**Fault management and availability.** *Fault management* is the continuous monitoring of the network and its elements for the detection of failures within the network environment. When a failure is detected, notification must occur in a timely fashion. The failure also must be qualified with respect to other network problems and prioritized.



**Figure 29.6** The screen on the left shows how a probe has been set to alarm on a security violation (an unauthorized user trying to access network information that is protected). The screen on the bottom left shows that the network signature of the intruder is captured and can be examined in detail for clues as to the source. The screen on the top right shows how examining source and destination traffic volume can pinpoint excessive use of a resource, indicating possible unauthorized access if access exceeds normal limits. The screen on the bottom right shows how easy it is to spot unauthorized access between segments, because the intruder's address is unknown to the server on that segment, returning an error label instead of an address on the graph.

Fault management systems include software bundles to detect and notify a centralized system of failures. The centralized system normally includes some form of discovery and mapping software that allows the network manager to have a graphical view of the network. These notifications must be correlated so that event storms are eliminated. A trouble-ticketing system can be incorporated so that a document trail is kept of the problem, allowing a mechanism for communicating the status of the problem to the end users.

Another aspect of fault management is network resource availability. This includes the monitoring of servers from the end user's perspective to ensure that the machine is available. Tracking and notification of any interruption of client-server access is a critical part of the IT department's function.

The screen shot on the left of Figure 29.7 is a sample packet capture and full seven-layer decode. The screen shot on the right shows a graph of historical statistics that can be used to help troubleshoot performance problems on the network.

**Capacity planning.** Network demand continues to grow at high rates. New applications are encouraging extensive use of the network. Graphics are now sent regularly over the network (either through a corporation's intranet or over the Internet). As

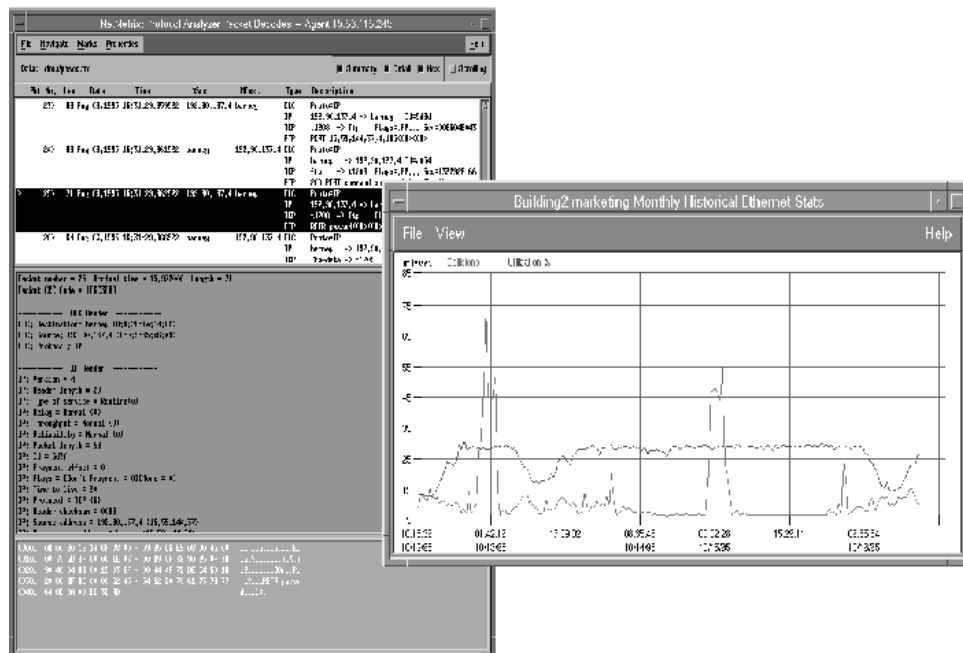
network managers increase bandwidth, new applications for the network (such as voice-over-IP or multimedia) become viable. This causes another spurt of demand for the network.

Capacity planning allows the network manager to anticipate network needs based on past activity. It helps the manager to forecast demand. This enables IT to stay one step ahead of demand.

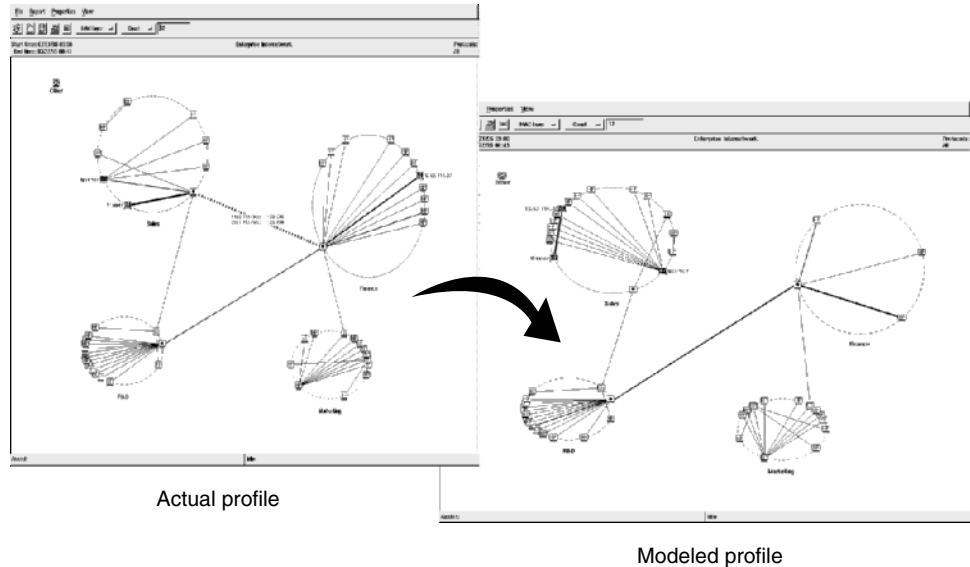
Figure 29.8 shows how data from the network can be collected for several weeks and used to perform what-if scenario planning. This allows IT to see the effect that changing the location of a server will have on intersegment traffic performance.

**Documenting network service level agreements.** Part of IT's function is to demonstrate that they are performing at the level of service agreed upon for end users. Documenting that the network is running within the acceptable envelope prescribed in the service level agreement is an important resource-justification tool. On-demand report generation is a good tool for troubleshooting. Automatic report generation allows reports to be scheduled on a regular basis, desirable when the effort to generate regular reports will be too substantial to make it worthwhile.

An alternative to printing reports is having the option to place them on an intranet web server so that business end users can view the information and print only what



**Figure 29.7** The screen on the left shows a packet decode. Packet information is displayed in three formats. Hexadecimal format is useful for detail investigation of problems. Detail format breaks the information into the component parts of the packet and interprets each part in easy-to-read words. Summary format provides a quick look at conversation flow with summarized information. The screen on the right shows the monthly statistics for utilization and collisions on a segment. The peaks and valleys are noteworthy and can indicate problems (as well as opportunities to tune the network to run more efficiently).



**Figure 29.8** Traffic profile modeling can help a network manager anticipate the user community’s network needs without committing resources of time and money. The effect of proposed network modifications and the impact on traffic patterns can be displayed to assist with making decisions. The actual profile of the network can be compared to the profile after some change in topology has been “virtually” effected, modeled by software, and graphically displayed.

they need, when they need it. This makes the information more accessible and spreads the knowledge exponentially, helping show that the IT department is working for, and with, the business end users.

## 29.4 Standards Overview

Distributed network monitoring has benefited from several standards. The Simple Network Management Protocol (SNMP) provided the foundation for network devices to communicate using the network. It specifies “who” on the network sends information and “who” receives it. It also defines a standard type of information (usually specific to one type of network information) that is passed between sender and receiver.

Within SNMP, the main standard currently defined for network monitoring is the *Remote Monitoring (RMON) Management Information Base (MIB)*, which defines a method of monitoring traffic up to the Data Link layer (layer 2) in the Open Systems Interconnect (OSI) stack. The *Remote Monitoring 2 (RMON2) MIB* standard, currently awaiting ratification by the Internet Engineering Task Force (IETF), the standards-setting body for network communication, defines how to monitor traffic at the Network layer (OSI layer 3) and some portions of the Application layer (layer 7).

Distributed monitoring is a large subject and involves many proprietary protocols. This chapter will discuss only standards-based protocols plus the most widespread proprietary protocols. Topics to be covered include:

- Simple Network Management Protocol (SNMP)
- Simple Network Management Protocol version 2 (SNMPv2)
- SNMP Remote Monitoring (RMON) Management Information Base (MIB)
- Remote Monitoring version 2 (RMON2) MIB

#### 29.4.1 Simple Network Management Protocol (SNMP)

The Simple Network Management Protocol (SNMP) was a draft standard in 1988 and was ratified by the IETF in April, 1989. SNMP is described in full in an IETF Request For Comment (RFC), a detailed technical description that itemizes all the components of SNMP. Each RFC is given a number that distinctly identifies it. SNMP is RFC 1098 of the IETF. SNMP has three basic components; their relationship is shown in Figure 29.9.

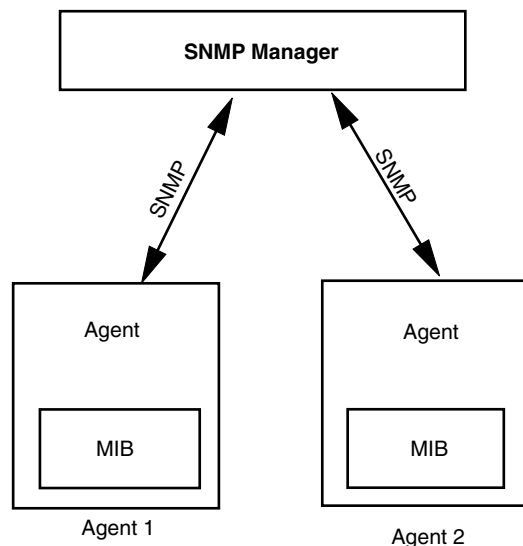
The *agent* (the collector and sender of information) is a software program that resides in a managed *element*, or device, of the network such as a hub, router, or a specialized dedicated data collector. The *manager* (the receiver of information) communicates with the agent by using SNMP commands.

A Management Information Base (MIB) is a database that resides within the agent and holds information in the various data fields that the database is designed to address. RMON is a MIB with a specific set of fields described in more detail subsequently in this section.

Communication between agent and manager takes place with five types of commands:

- Get request
- Get-Next request

- **Manager:** Manages one or more agents. Normally resides on a computer and communicates to one or more agents in the network.
- **Agent :** Collects data. Responds to get/set statements from an SNMP Manager and sends traps to specified destinations.
- **MIB:** Resides on the Agent and stores information for later recall by an SNMP Manager



**Figure 29.9** A manager communicates with agents distributed around the network. The agents collect and store information in databases designed specifically for one purpose, called Management Information Bases (MIBs). SNMP specifies exactly how the communication between manager and agent transpires.

**694 Network Test Instrumentation**

- Set request
- Get response
- Trap

With a *Get* request, a manager requests (from an agent) the value of some variable stored in a field of the database that represents a particular MIB.

A *Get-Next* request is used by a manager to request information on multiple variables, efficiently using the network (reducing traffic) to communicate more than one piece of information at a time. If any one variable is unavailable, no values are returned. This command also is used to retrieve unknown database rows (if available).

With a *Set* request, a manager instructs an agent to set a specific MIB variable (one of the fields in the database) to a desired value.

A *Get response* is sent by an agent as a response to a Set request or Get-Next command; its form depends on what was requested. Examples are:

- An error message
- A response identical to the variable sent in the Set request (showing it was accepted), or
- Actual database values, as requested in a Get-Next.

The manager checks its list of previously sent requests to locate the one that matches this response and, if none is found, the response is discarded; otherwise it is processed further.

A *trap* is one of two unsolicited messages sent from an agent to a manager, often used for out-of-norm event notification when the agent “sees” something that the manager should know about.

**29.4.2 SNMP version 2**

The SNMP standard brings several advantages to the rapid adoption and growth in use of network management applications:

1. The protocol is easy to implement.
2. The protocol requires few resources to operate.
3. The protocol is mature, stable, and well understood.
4. The protocol is widely available (on most computers) and most network devices have some form of embedded agent/MIB.

As networks have grown and the need for network management has become more imperative, several disadvantages of SNMP have become apparent. For example, as a polling-based protocol, each agent is responsible for sending information to the manager as requested. Depending on the number of agents deployed on a network, the amount of traffic this imposes on the network can be burdensome and resource intensive, both in bandwidth utilized to communicate and in processing power to analyze the information.

As more and more mission-critical information is moved to the network, unauthorized access to sensitive information is becoming a high-priority item across all industries. SNMP has limited security provisions.

Currently, table data (blocks of data) cannot be requested to be sent from an agent. This prevents the transfer of multiple variables at once, contributing to unnecessary traffic on the network.

Traps from an agent notify a manager that some event of note has been captured on the agent. SNMP did not provide for an acknowledgment to be sent back to the agent that the event had been received. Because traps are unsolicited, initiated by the occurrence of an event on the agent, it is becoming increasingly important that the manager acknowledge receipt of the trap (typically an out-of-norm condition detected by the agent).

**SNMPv2 improves on SNMP.** SNMP version 2 (SNMPv2) is an addendum to the existing SNMP standard, designed to address practical shortcomings and the evolving needs of emerging technologies.

SNMP defined manager-to-agent communication. SNMPv2 builds on this by defining manager-to-manager communication, facilitating the need to scale network communications to sprawling networks. This deals directly with the problems posed by the network being used to manage the network. The volume of management information flooding the network infrastructure and management console from every segment can, conceivably, predominate the use of network bandwidth.

Manager-to-manager communication will allow an intermediate level of managers to be inserted into large enterprise networks to sit between the network management console and the numerous agents. These mid-level managers each will manage a subset of the agents on the network. This is called a *management domain*. These mid-level managers will then pass relevant information about their domains to the central network management console. This architecture reduces the number of connections to the central manager to one per domain. It enables distributed network monitoring and analysis to be scaled to fit the size of a network while allowing network capacity and processing power to be channeled to business information needs, not network management.

Security issues addressed by SNMPv2 include timestamping collected data so that an indelible audit trail is laid down. A data encryption standard would provide a screen for company-sensitive information traversing networks that are vulnerable to unauthorized listening. The availability of automatic, user-transparent encryption mechanisms could be the single largest stimulus to purchasing via credit card over the Internet. Message authentication would prevent rogue users from masquerading undetected on the network. All of these features are critical to continued use of the network to conduct business.

SNMPv2 will allow whole tables of data to be passed to the manager. A new type of request, a *GetBulkRequest*, should alleviate congestion on the network and promote the use of tables in network monitoring and analysis.

Provisions have been made in SNMPv2 for acknowledgment of traps to be sent back to the agent. The agent will retry notification of the trap until successful, or, if unsuccessful, will notify the user that the manager is not responding. This intro-

duces higher reliability that out-of-norm conditions detected by the agent will be communicated and dealt with in a timely manner.

### 29.4.3 The MIB tree

Management Information Bases are hierarchical in nature and can be viewed logically as a tree structure (Figure 29.10). The branches of the tree are called *variables* or *MIB objects*. Unique identifiers or names are assigned to each branch. MIBs of interest for network monitoring and analysis follow. Like SNMP's RFC number, each MIB has a unique RFC identifier.

- RFC 1213 – MIBII: Basic system information and basic level statistics
- RFC 1757 – RMON: Remote monitoring
- RFC 1513 – RMON: Remote monitoring extension for Token-Ring


### 29.4.4 The Remote Monitoring (RMON) MIB

The RMON MIB is a specific collection of objects used for remote monitoring networks. The RMON MIB is defined by RFC 1757 and RFC 1513. Figure 29.11 shows how the RMON MIB is broken down into ten objects or groups. The first nine define monitoring of Ethernet networks; the tenth defines extensions for Token-Ring. There currently are no standards for monitoring FDDI, 100Base-X, or WAN networks. RMON vendors have added their own proprietary extensions for these additional media types. RMON is limited because it gives visibility only up to the Data Link layer (layer 2) in the OSI stack.

A brief description of the ten RMON MIB groups will provide some insight into the variables that can be collected by monitoring traffic on a network.

**Statistics group.** This group contains segment statistics, collected in 32-bit counters, that include items such as packets, dropped packets, broadcasts, and multicasts. These are just counters and not studies.

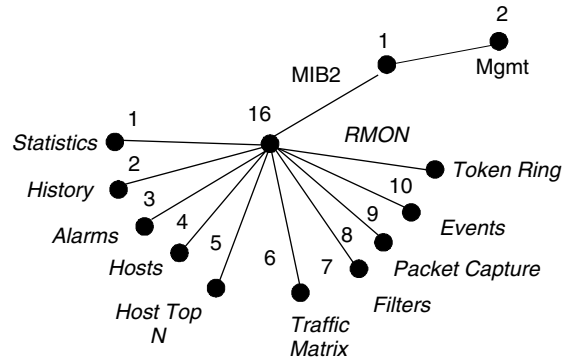
**History group.** This group contains segment history statistics for various counters, such as broadcasts, errors, multicasts, packets, and octets. These numbers are for certain time periods. RMON defines two default time periods, 5 seconds and 1800 seconds.



**Figure 29.10** An MIB tree is like outlining a complex document. By knowing the exact paragraph, it is easy to locate and retrieve data.



- The standard in remote monitoring SNMP MIBs
- A MIB designed to alleviate some critical deficiencies of SNMP not just a poll-based MIB.
- A MIB tuned for SNMP remote monitoring devices
- A MIB rich in functionality and features
- Some groups needed to implement others



SNMP RMON MIB =  
 Simple Network Management Protocol  
 Remote MONitoring  
 Management Information Base  
 (RFC 1757) Ethernet  
 (RFC 1513) Token-Ring

**Figure 29.11** SNMP branches into network-specific MIBs. The MIBs branch out into a collection of objects. Each object may contain several variables where actual values are stored. The RMON MIB of SNMP is specific to remote network monitoring and consists of ten MIB objects.

**Alarms group.** This group covers threshold monitoring (where limits are set for variables and actual values exceed the limits), as well as trap generation when that threshold has been reached. Alarms can be set for various counters and packet match. Traps can start and stop packet capture.

**Hosts group.** This group host table and statistical counters for packets, octets, broadcasts, multicasts, and error packets, plus a timetable of discovery.

**Host Top N.** This group contains studies for a defined time period and a defined number of hosts, listing addresses of the top traffic generators for the study group.

**Traffic matrix group.** This contains a matrix of MAC layer (layer 2) conversations. Information such as errors, packets, and octets are sorted by MAC address.

**Packet capture and filter groups.** These two groups are used together. The packet capture group contains the packets that have been captured. Filters allow capture only of packets that meet certain specified criteria. Multiple instances can be created.

**Events.** An *event* causes an agent to generate an unsolicited notification to a manager when an Alarms group threshold is exceeded. Logs all events with detailed information about the event.

**Token-Ring group.** This group contains information specific to Token-Ring networks, such as ring order, ring station table, and packet size distribution for history studies.

#### 29.4.5 Remote Monitoring version 2 (RMON2) MIB

The RMON standard brought many benefits to the distributed monitoring community, but many features were left out. The RMON2 standard attempts to address

these issues by allowing the monitoring of Network layer (layer 3) information, as well as protocol distribution information up to the Application layer (layer 7). RMON2 provides the ability to see beyond a routed boundary and identify the protocol mix of network activity up through layer 7 of the OSI stack.

In addition, a probe configuration MIB has been added to improve interoperability. This will allow for standard methods to upgrade probes, set network parameters such as IP addresses, and configure SLIP or PPP connections. Finally, a user-defined history mechanism has been added to allow an RMON2 agent to maintain historical studies on arbitrary objects supported by the agent.

#### 29.4.6 Proprietary monitoring protocols

Many proprietary monitoring methods have been developed before and after the RMON/RMON2 standard. The Embedded Advanced Sampling Environment (EASE) protocol is the proprietary protocol in most widespread use.

**Embedded Advanced Sampling Environment (EASE).** The philosophy of EASE, a sampling technology for networks, is that there is no need to deploy expensive RMON technology when low-cost EASE devices perform a similar function. The RMON standard states that every packet must be analyzed. This means that fast hardware must be deployed to cope with the ever-increasing network speed and traffic volume. The potential exists for this to be prohibitively expensive at higher transmission rates (100 Mbps and higher). Currently there are no RMON probes that can handle FDDI utilization above 40 percent because it's just too expensive to produce the hardware.

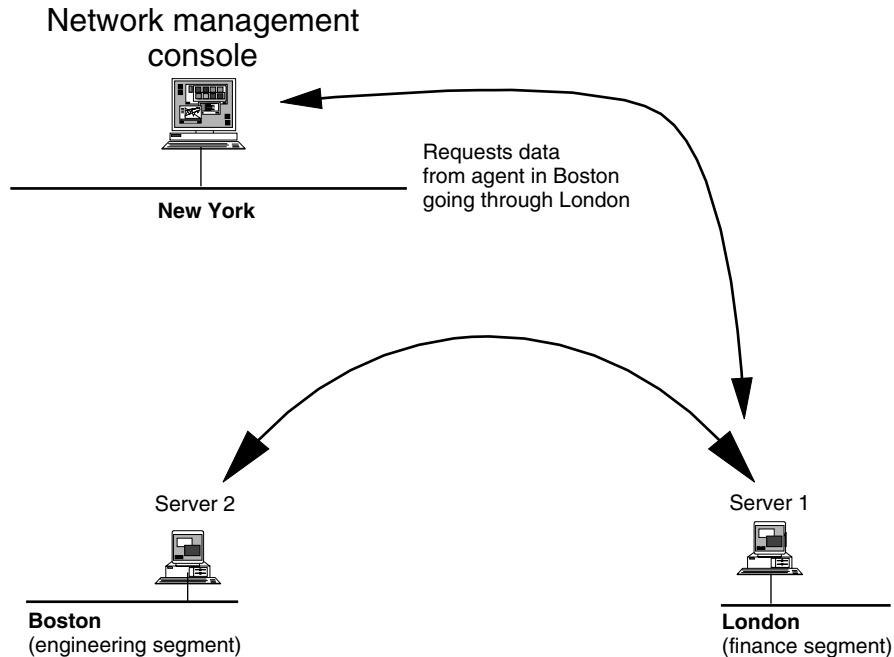
The idea behind EASE is that looking at every packet is not necessary. Sampling already is used in many other areas, such as industrial quality control and financial auditing. By sampling only 2 percent of frames, the cost of the monitoring hardware that needs to operate only at 2 percent of the media speed is greatly reduced.

**Using echo test.** None of the standards heretofore discussed actually allows a network manager to directly monitor the latency in the network. *Latency* is the amount of time it takes a source message sent on the network to reach its destination, and for the source to be notified of receipt. A simple echo test allows network managers to measure the latency in the network.

Figure 29.12 is an example of an echo test between London and Boston. The data from the results can be graphed, revealing potential performance bottlenecks on a network. Echo tests can be configured to work with alarms to notify network staff automatically when latency exceeds acceptable limits.

### 29.5 Limitations of Distributed Network Monitoring

Monitoring the network with the SNMP RMON MIB means that only Data Link layer information is collected. This is not high enough in the OSI stack to provide information about traffic trends of client-server applications.




**Figure 29.12** In this illustration, a network management console initiates a request for an echo test to audit the latency between the console and Boston. The resulting information will give the total latency as well as the component parts of that latency. How much latency exists between New York and London or between London and Boston? This will provide network staff with some clue as to where latency is being introduced.

The RMON2 standard defines a method of monitoring up to layer 7 at certain times. RMON2 does not define continuous monitoring of all layer 7 traffic, nor does RMON2 define any metrics for performance management.

### 29.5.1 Only part of the solution

Network monitoring is only part of a solution that must include the business end users' entire environment. This means that system, database, and application monitoring tools must be deployed in conjunction with network monitoring tools so that the IT department can view the whole environment. *Integrated network and systems management* (INSM) tools are in their infancy. Figure 29.13 includes a screen shot showing database, application, system, and network information correlated together over time.

700 Network Test Instrumentation



- Manage the performance resources from a “single pane of glass”
- Pinpoint problems in system, network, database, or application resources before they affect users.

**Figure 29.13** All aspects of the IT environment can be viewed on a single graph providing a holistic view of a network. These simultaneous graphs are time-correlated and reveal information about how the business is using the network. Is more bandwidth being used by the network than by applications?

# SDH and SONET Analyzers

**Doug Moore**

*Hewlett-Packard Ltd., South Queensferry, Scotland*

## 30.1 Introduction

The increasing bandwidth capacity available in modern optical fiber transmission lines has led to the development of standards for a synchronous digital transport network. The ITU-T countries have defined two standards: SONET (Synchronous Optical Network) for North America, and SDH (Synchronous Digital Hierarchy). Both standards are based on similar frame structures, the only major difference being the primary base rate used to build up their synchronous hierarchy. If you are unfamiliar with SDH/SONET signals, it is recommended that you examine Chapter 13 before continuing with this chapter.

The main service required of today's network operating companies is the ability to provide a faster response to the provisioning of new customer circuits and services. Ultimately their goal is to provide their customers with online control of circuit bandwidth and services. To meet these requirements, network operating companies must improve their ability to manage the bandwidth available in their networks, and they must do this cost-effectively. Consequently, the predominant network requirement has become telecommunications networking, supported by a more advanced approach to network management and maintenance based on computer systems.

Designed for cost-effective, flexible telecommunications networking, the synchronous standards are based on the principles of direct synchronous multiplexing. In essence, this means that individual tributary signals may be multiplexed directly into a higher-rate SDH signal without intermediate stages of multiplexing. Synchronous Network Elements (NEs) then can be interconnected directly, with obvious cost and equipment savings over the existing network.

The synchronous signal structure provides built-in signal capacity for advanced network management and maintenance capabilities required to effectively manage and maintain network flexibility. Approximately 5 percent of the SDH signal structure

**702 Network Test Instrumentation**

is allocated to supporting advanced network management and maintenance procedures and practices.

The synchronous structure provides a flexible signal transportation capability. The signal is capable of transporting all common tributary signals found in the previous asynchronous telecommunication networks. This means that the synchronous network can be deployed as an overlay to the existing network and, where appropriate, can provide enhanced network flexibility by transporting existing signal types. In addition, the standards have the flexibility to accommodate the new types of customer service signals that network operators will need to support in the future.

Synchronous structures can be used in all three traditional telecommunications application areas: long-haul, local network, and loop plant network. This makes it possible for a unified telecommunication network infrastructure to evolve. The fact that the synchronous standards provide a single common standard for this telecommunications network means that equipment supplied by different manufacturers may be interconnected directly.

**30.2 The Basic SDH/SONET Analyzer**

An SDH/SONET analyzer essentially is a dedicated high-speed protocol analyzer (see Chapter 27) with its features tuned to measurement and testing of one or both of the synchronous protocols. It contains hardware to capture and examine the serial data stream, along with circuitry to evaluate the data stream for alarm and error conditions. In addition, most analyzers have the ability to transmit a synchronous data stream, emulating the normal behavior of an SDH/SONET signal. Some analyzers also are able to insert a variety of alarm and stressing conditions into the synchronous signal.

**30.2.1 Uses of the analyzer**

Analyzers are used throughout the life cycle of a network element (NE), from the design of application-specific integrated circuits (ASICs) and PCBs in the R&D labs, and all the way through system integration, conformance testing, and final production. In addition, analyzers are used extensively during the installation and maintenance of the NE and of the overall telecom network. The function of the analyzer varies with the area of use, but in general analyzers are used to monitor signals received from the NEs, and to provide test stimulus to the NE or network. The basic features of the analyzers remain the same across this life cycle, although the relevant importance of each varies with the job required. More details of feature set vs. analyzer function are given in section 30.4.

**30.2.2 Operating modes**

Most analyzers can operate in a variety of modes. The three most important are:

- Terminal mode
- Through mode
- Add-drop multiplex

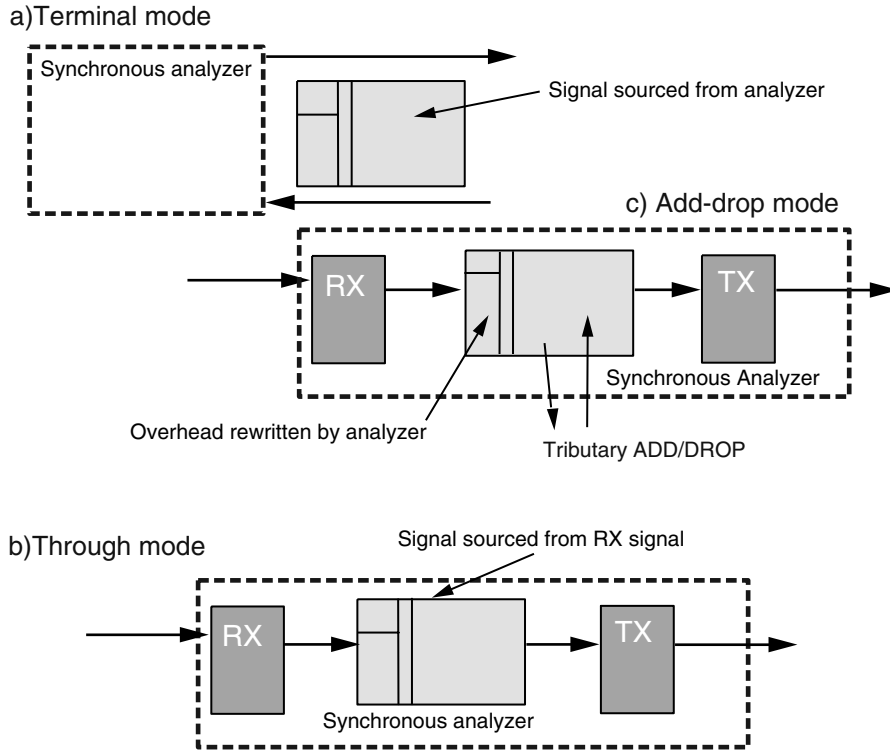


Figure 30.1 Basic SONET/SDH analyzer operating modes

**Terminal mode.** In this mode, the instrument acts as the sink and/or source of the synchronous signal, terminating and examining all parts of the EOC bytes on the sink side, and reporting alarms and errors in the received signal as they occur (Figure 30.1a). It also provides the entire signal on the source side. This type of operation most commonly is used in test situations in which a single NE or component is under test. Typically, this is seen in R&D and manufacturing tests.

**Through mode.** In this mode, the analyzer is positioned between two NEs, and acts as a monitor in the system (Figure 30.1b). This facility is particularly important in network installation and maintenance (I&M) applications. The amount of alteration to the signal available to the user varies considerably from analyzer to analyzer, ranging from a monitor-only capability through to the complete error addition, stressing, and erroring functions outlined in the following sections.

**Add-drop multiplex.** (See Figure 30.1c) The analyzer again is positioned between two NEs, but this time the analyzer acts as a network element in its own right, sinking a part (or all) of the incoming signal, and replacing it with signals sourced internally from the analyzer. This operation most commonly is used when a network or a group of NEs is being tested.

### 30.3 Synchronous Measurements

The features that distinguish synchronous analyzers from the basic protocol analyzers examined in Chapter 24 are related to the measurements required within a synchronous system. This section deals with some of the basic measurements made on synchronous NEs, along with details of the SDH/SONET analyzer features that make the measurements possible. The list is by no means comprehensive, but is intended to give an idea of the type of testing to which NEs are subject.

#### 30.3.1 Functional measurements: static analysis

This type of measurement deals with the purely digital aspects of the synchronous signal. This involves investigating the behavior of an NE under normal operating conditions, as well as measuring its reaction to various alarm and defect conditions defined in the standards. The measurement types can be broken into a number of distinct cases.

**Capturing synchronous data.** This type of signal measurement involves purely passive testing from the analyzer, which must be able to examine the signal correctly with reference to the framing bytes, and report/display the results. The analyzers generally perform static examination of the signal in two ways:

- Capturing synchronous data
- Alarm reporting and trouble scanning

For capturing synchronous data, the analyzer needs to correctly capture the signal with reference to the framing bytes and display the results. The signal can be captured in a number of ways but usually is divided into two areas: The embedded overhead (OH) containing the signaling and alarm /error information, and the actual data being carried by the signal (the payload).

**Overhead capture.** The ability to examine the path and transport overhead of a synchronous signal is an important feature for analyzers. It is accomplished in many ways, but in general there are three main methods:

- Dynamic capture
- Static frame capture
- Selective byte capture

*Dynamic capture* allows visual monitoring of the signal. Most analyzers provide a dynamic display of the overhead (OH) to a screen (Figure 30.2). With a frame rate of 8000 per second, however, it is not possible to give a continuous frame-by-frame update, and the refresh rates typically are on the order of hundreds of milliseconds. This method is a crude and quick way to determine visually whether the OH is functioning correctly, and could allow an operator to ascertain the nature or position of a fault within the signal.

To examine the OH more closely, most analyzers also use some form of *static frame capture*. This method displays a set of consecutive OH frames captured and



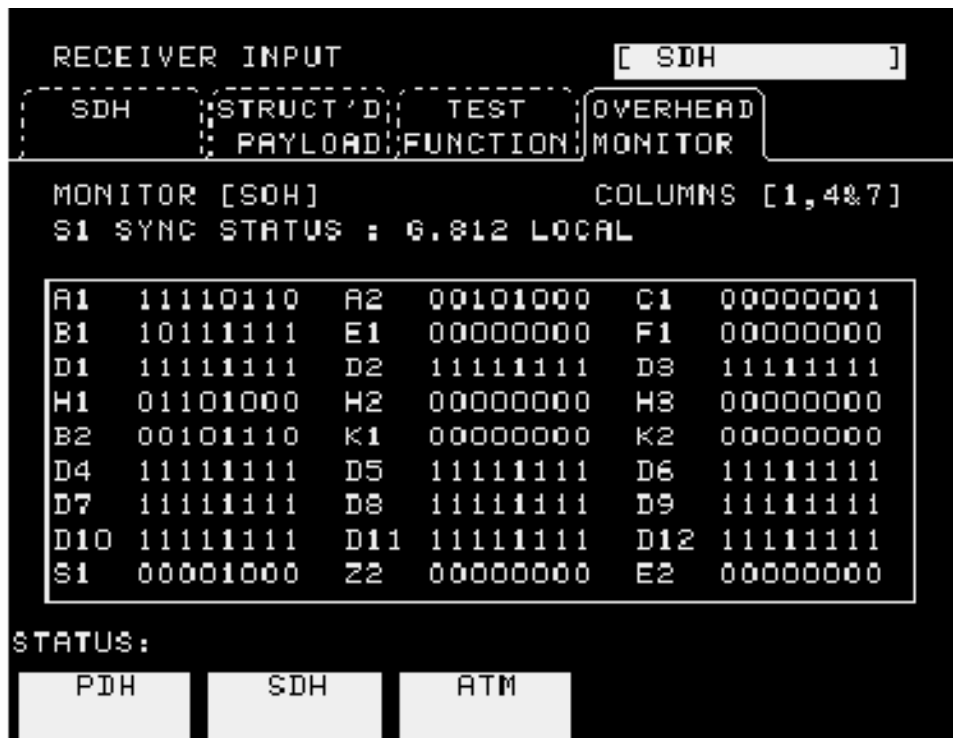


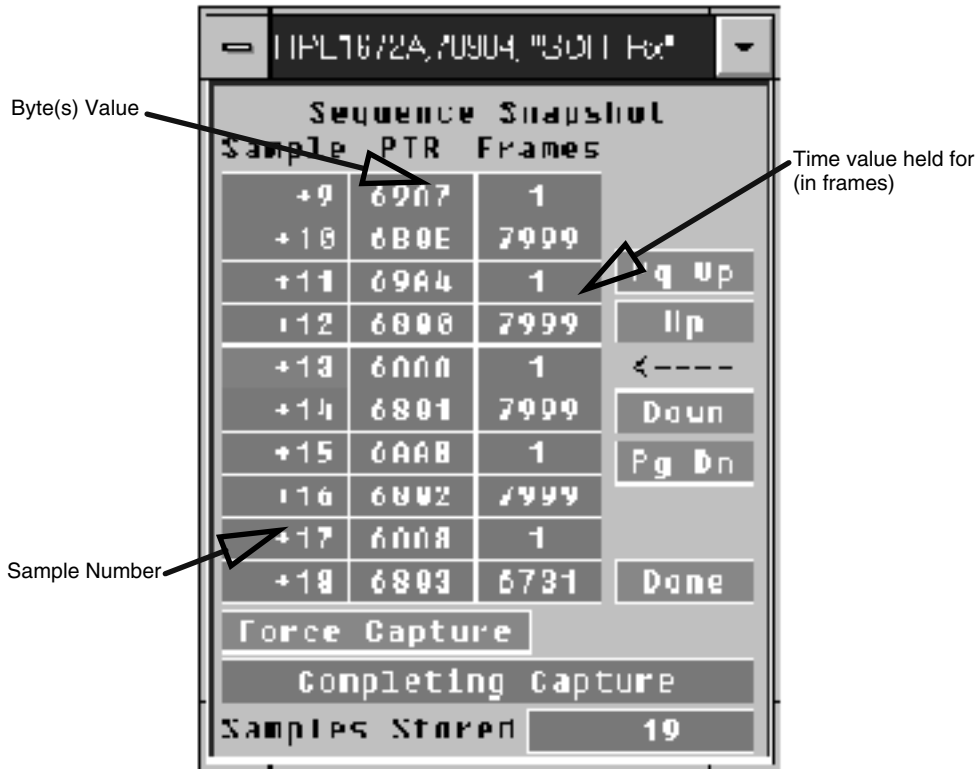
Figure 30.2 Typical analyzer overhead monitoring screen.

held in memory. The frames usually are captured after a trigger event. The range of capture triggers available to the user depends upon the complexity of the analyzer, and can vary from simply the next correct framing bytes (A1A2), to any desired hexadecimal value of any of the overhead bytes. External triggering also might be available.

This type of feature allows a snapshot of the OH's operating behavior over the period covered by the captured frames. The number of frames captured depends upon the depth of memory in each instrument, but ranges of 8 to 200 frames are not uncommon.

Static frame capture is a good technique provided that the time interval to be examined is small (1 ms is 8 frames). Beyond this, the memory requirements for capturing all the bytes start to become prohibitive. Some pointer and protection switching protocol events may occur over a period of several seconds or minutes.

To capture even a second (8000 frames) of the full overhead clearly is impractical, so more sophisticated analyzers use the *selective byte capture* technique. This technique samples a byte (or byte pair, such as the H1H2 pointers) from the OH. The value of the bytes and the number of frames for which they are held are displayed in real time by the instrument, allowing a frame-by-frame analysis of behavior to be performed over an extended period.



**Figure 30.3** Selective byte capture (AU/STS shown here).

An example of such a display is shown in Figure 30.3. In this case the selective capture is focused on the H1H2 pointer value during a pointer movement set to increment at a rate of one per second. The display shows a pointer value held for 7999 frames, an increment flag set for one frame, then a new pointer value held for a further 7999 frames.

The increase in sample time is significant; one analyzer that uses this technique has memory sufficient to capture up to eight frames of a STM-4 signal, a total of 1 ms. In its selective byte capture mode it can hold up to 2048 changes of the monitored data, each holding for upwards of 65,000 frames, giving a total of up to 4.5 hours of continuous monitoring.

**Payload capture.** Transmission of an (unimpaired) payload is the ultimate purpose of a synchronous transmission system, so examination of the payload is an important function of an analyzer. The number of frames of payload captured is limited only by internal memory. As in overhead capture, the methods for capture range from simple manual triggering through sophisticated trigger on errors, byte values, etc. A number of techniques can be readily identified:

- Static capture (entire SPE)
- VT/TU capture
- With overhead (frame capture)

*Static capture* is similar to overhead capture: A specific number of frames of the synchronous payload envelope are captured and displayed. Typical values range from 8 to 64 frames. This capture is useful when the data being carried covers the entire SPE.

Relevant *VT/TU capture* is more informative if a subrate (DS1, E1, etc.) is being carried. Some analyzers will capture the entire SPE and allow individual VTs or TUs to be displayed, while other analyzers simply capture a specified VT or TU position.

*Frame capture* is important for viewing the entire contents of a data stream, to correlate events in the payload with those in the overhead. Many analyzers allow entire frames to be captured and examined, the number of captured frames varying with line rate and available memory. Up to 64 frames is not uncommon.

**Report and trouble scan.** In this mode an analyzer will monitor the received signal and indicate when an error, defect, or alarm has been received. In the mode's crudest form, an event is flagged by front-panel LEDs that indicate the type of alarms received. Most if not all analyzers support this type of indication. More useful is counting, measuring, and processing of the error alarms and defect conditions. Four significant areas of signal monitoring need to be considered:

- Error reporting
- Alarm reporting
- Signal decoding
- Logging

**Error reporting.** For this the analyzer needs to capture the bit interleaved parity (BIP) bytes in the overhead and payload structure and compare the value with the relevant BIP value calculated by the analyzer using the received data stream. Any discrepancy indicates an error in the data transmissions. This error can be reported in a number of ways, the crudest being an LED that flashes whenever a BIP error is received. This is of little practical value other than to give the user a rapid visual indication of a faulty data stream. For more practical use the errors must be stored and displayed. Although this can be done in a number of ways, four methods are commonly used:

- Counts
- Ratios
- Seconds
- G.826

An *error count* display is self-descriptive. It is a count of the number of errors received, usually since the measurement started, or sometimes over a specified time

period. This display type is useful in a situation where a small number of errors is expected. Once the number exceeds a few hundred, however, a more useful definition of the error performance can be obtained using the *error ratio*. This error ratio is obtained by dividing the number of BIP errors received by the number of bits transmitted and covered by the parity byte. It can be seen that this ratio is a function of the transmission rate as well.

Error ratio calculation is best demonstrated by considering as an example a B1 parity byte. Suppose there is one error per minute at the 155.52 Mbps (STM-1 or STS-3) line rate. The B1 byte covers all the bits in the serial data, so in 1 second there are 155,520,000 bits transmitted and checked by the B1 byte. With 60 seconds in a minute, this level of error is represented by an error ratio of

$$\text{error ratio} = \frac{1}{(155,520,000 \times 60)} = 1.07 \times 10^{-7}$$

Under most conditions this would be considered a fairly poor-quality link; common error rates specified for trunk cables (running at 2.4 Gbps and above) are better than  $1 \times 10^{-15}$ , which on a 2.4 Gbps link represents a single BIP error every 5 days!

Another useful diagnostic measure is *errored seconds*. This display counts the number of seconds that contain at least one BIP error. This measure can be useful in determining the type of error being sought. Consider again the foregoing error rate example. At a rate of one error per minute, after an hour of measurement an errored second count of around 60 would indicate a fairly “regular” error mechanism. An errored second count of 2, however, would indicate that the errors were occurring in bursts.

The ability to make *G.826 measurements* is an increasingly important feature (particularly for SDH I&M applications). The ITU-T specification G.826 defines quality of service (QoS) measurements, error performance parameters, and objectives for international, constant bit rates at or above the primary rate.

It is becoming increasingly common to see SDH links specified in terms of severely errored seconds (SES), unavailable seconds (UAS), etc. Any analyzer aimed at the SDH I&M market that doesn't have G.826 capability is at a marked disadvantage.

**Alarm reporting.** Methods used in alarm reporting are similar to those used for error reporting. In this case, however, it is usually only counts and seconds that are useful to the user. It is important that the analyzer's (default) alarm thresholds are consistent with those defined in the relevant standards; particularly for R&D applications, however, it is an advantage to have user-adjustable thresholds available as well.

**Signal decoding.** As well as alarm and error information, the embedded overhead channels (EOC) also contain command and control information necessary to the correct function of a synchronous network. Examples of these types of signal are:

- *J2*: A 16-byte repeating pattern that is part of the path overhead used to route the payload.
- *K1K2*: Two bytes used to transmit protection switching protocol, part of the Line/MS overhead.




Figure 30.4 Typical SDH format logging screen.

- *C2*: One byte, part of the path overhead, used to indicate the type of payload being carried.

This information frequently is in the form of coded patterns that translate to an alphanumeric message. For example, in SDH the all-zero code represents VC-4 *un-equipped* (i.e., not containing any tributary signals); the code 00000001 represents VC-4 *equipped*. A signal decoding feature is available in most analyzers that can decode some or all of these messages and display them, allowing the user to assess the performance of the network element.

**Logging.** All the features just described are useful in increasing an analyzer's ability to examine passively the synchronous signal. For a number of applications, however, long-term testing is necessary, and a function is needed to relate errors/alarms/signals received against time. This feature allows events to be logged versus time so that it is possible to correlate events in the signal with information from the NE's own management system.

Most sophisticated analyzers provide logging of some description, varying in complexity from a simple printout at each event, through to full storage of all events, and access to the information available over remote RS432 or LAN connections. This feature is of greatest importance in the Installation and Maintenance (I&M) phase of deployment, as well as in long-term testing in QA and acceptance testing. Figure 30.4 shows a typical logging screen.

### 30.3.2 Functional measurements: signal stressing

Passive examination of the signal is the simplest form of measurement required by a user. To further test an NE it is necessary for an analyzer to simulate both normal and defect conditions within the synchronous signal. The most common forms of signal stressing are detailed in subsequent paragraphs.

**Alarm stressing.** Most analyzers that can transmit a synchronous bit stream are also capable of stressing the signal to simulate both normal synchronous operations, and defects which can stimulate alarm and near alarm conditions. The type of stressing employed depends upon the type of analyzer, but the most common types are detailed below.

**Alarm soaking.** The basic type of alarm stressing is *soak stressing*, quite simply the ability to turn the alarm condition on and off for a period of time in the data stream. This facility is available in most analyzers, and allows an NE's reaction to alarm conditions to be examined. The type of soak stressing available varies from analyzer to analyzer. In its simplest form it is an on/off button; more complex implementations allow the alarm to be programmed into an On state for a defined number of frames. This "soaking" represents the simplest way to test an NE's reaction to an alarm.

**Alarm sequencing.** Unlike asynchronous signals, however, synchronous signals are detected on a frame-by-frame basis, so a simple "soak"-type stress doesn't provide a rigorous test of the NE's alarm detection capabilities. To test completely an alarm detection algorithm, a (three-stage) alarm sequencer is required (Figure 30.5). From an initial alarm On/Off state, the sequencer sends a single stimulus of

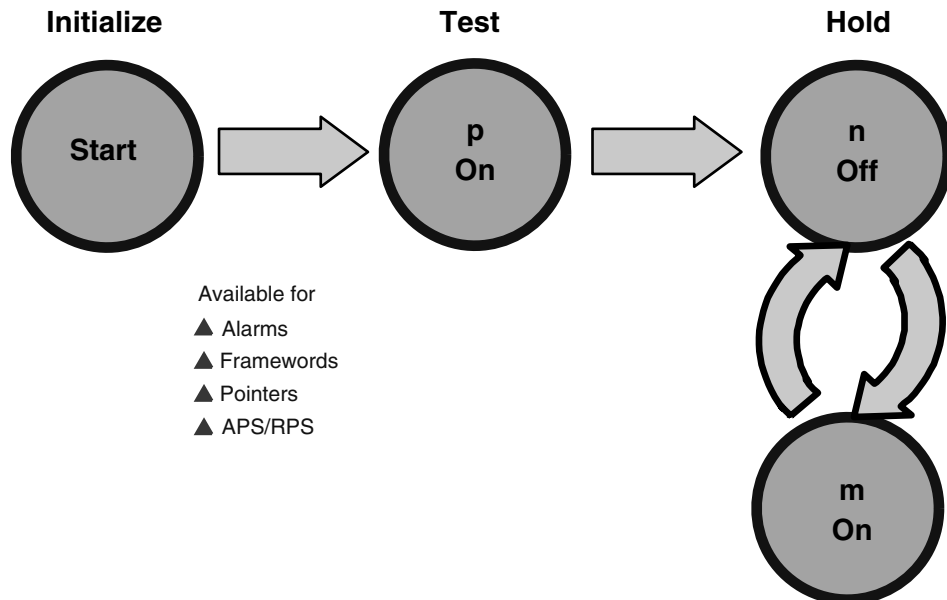


Figure 30.5 P-N-M stressing.

[P] frames of alarm Off/On, followed by a repeating sequence of [M] frames of alarm On/Off and [N] frames of alarm Off/On for as long as desired.

Most sophisticated analyzers can simulate such a sequence; two methods are employed:

- Programmable overhead frames
- Genuine sequencers

Many analyzers have the ability to program all or part of the transport overhead (including the alarm indication bytes) for a number of frames. This allows the P-N-M sequence to be programmed into the overhead. This gives only a partial solution to the three-stage sequence, however. The N-M holding sequence is limited by the maximum number of frames. For most analyzers this maximum is about 64 frames, giving only 8 ms of holding pattern before the P stimulus is repeated. This is a better test than a simple soak but is not the most rigorous.

The most sophisticated analyzers have a three-stage genuine sequencer that can provide exactly the P-N-M sequence. Giving a single [P] stimulus followed by an N-M holding pattern of any desired duration. This is the most rigorous of tests, fully stressing the alarm algorithm. An example follows of how this sequencer could be used.

**A sequencer example.** Consider the SONET Line AIS defect (similar to SDH MSAIS). The entry condition for Line AIS is five frames of all 1s in the K2 bytes; the exit condition is five frames with the K2 byte clear of the AIS indication. This example looks at an AIS exit threshold test, showing both under-threshold (Figure 30.6) and over-threshold (Figure 30.7) conditions. For the under-threshold sequence:

P = 4 frames normal K2  
 M = 1 frame Line AIS K2  
 N = 4 frames normal K2

This sequence gives the maximum number of clear frames without exiting the AIS alarm. Most analyzers with programmable capability can produce this type of under-threshold test. For the over-threshold sequence (Figure 30.7) the following applies:

P = 5 frames normal K2  
 N = 4 frames Line AIS K2  
 M = 1 frame normal K2

This sequence gives one burst of five frames (Line AIS exit threshold), followed by a holding pattern that is as close as possible to the entry threshold without exceeding it. Under these conditions the NE should go clear of the alarm.

These over- and under-threshold sequencers represent the most thorough test of the Line AIS alarm exit criteria. An NE that passes the above criteria is compliant with the SONET standards. It can be seen that this technique can be readily applied to most alarms in the synchronous standards.

**Choice of testing.** A wide variety of test techniques is offered by the synchronous analyzers; which type of testing to use generally is a function of application. In

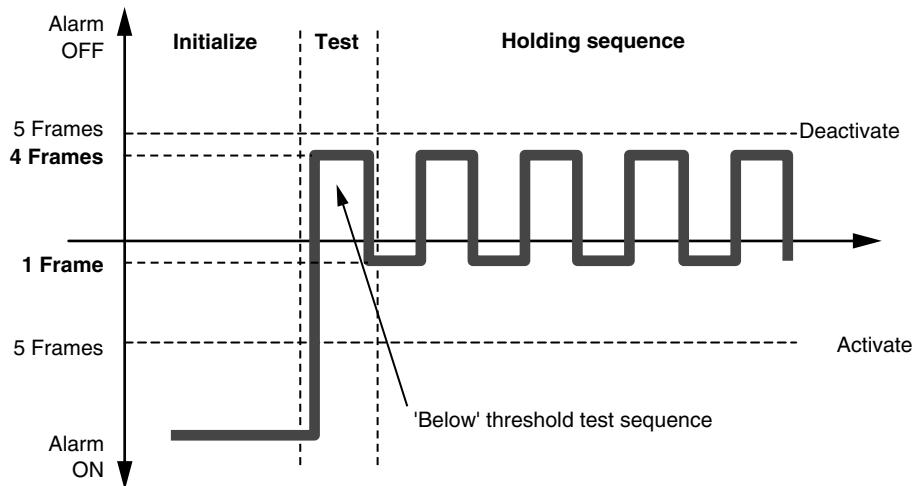


Figure 30.6 P-N-M stressing example: Line AIS exit under-threshold sequence.

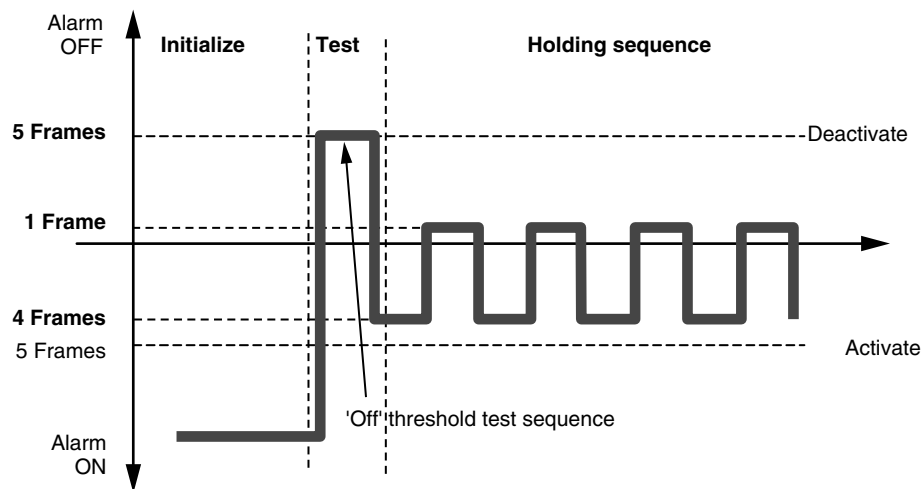


Figure 30.7 P-N-M stressing example: Line AIS exit over-threshold sequence.

general, production testing and I&M applications require quick and simple tests for function; more rigorous testing is usually confined to R&D and QA applications.

**Pointer stressing.** Pointers are vital to the synchronous signal’s ability to carry payload in an efficient and flexible manner. Testing pointer processing algorithms therefore is an important function of most analyzers. Figure 30.8 shows the flow chart of the pointer processing algorithm as defined by the synchronous standards. The pointer can be legally changed at a maximum rate of once every four

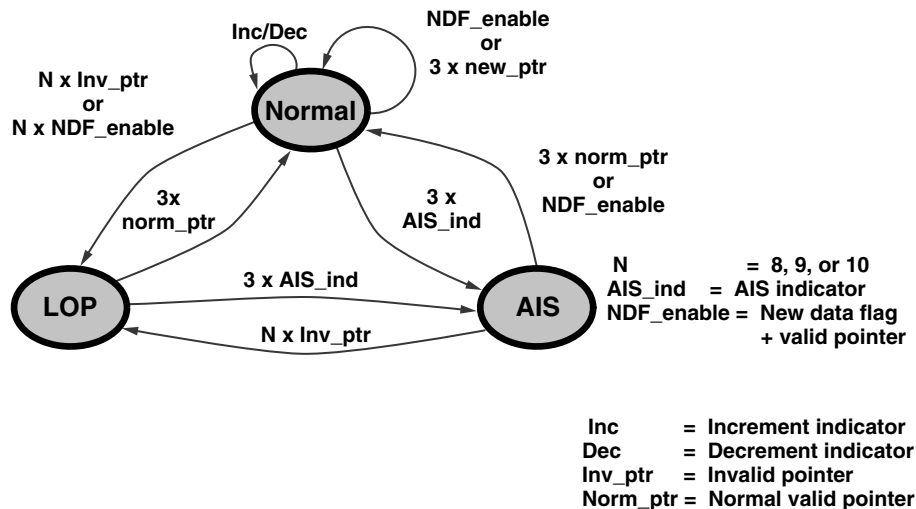


frames. Most analyzers therefore are able to generate the maximum-rate pointer moves. More flexible analyzers also can be programmed to produce a number of rates. The most sophisticated have preprogrammed sequences that conform to test sequences specified in the standards; the 87/3 test specified in GR-253<sup>1</sup> is a common example.

It also must be possible to make pointer changes in the presence of errors in the change indication (New Data Flags NDF) bits. Again, a number of analyzers allow the injection of errors into the NDF bits to stress this feature. The error mechanisms vary from simple, one-shot random error addition, to error rates applied to particular bits of a NDF as defined by a mask. In addition, some analyzers also possess another feature to test pointer algorithms.

**Frequency offsetting.** Pointers allow the synchronous network element to accommodate asynchronous behavior between the reference clocks of differing networks. This asynchrony is limited ( $\pm 20$  ppm maximum), but an NE must be able to operate without error under these conditions. Many analyzers will produce frequency offsets to stimulate pointer movement in an NE.

This frequency offsetting can be done in a number of ways, but the starting point is the same: The network element and the analyzer are locked to the same frequency standard to eliminate any extraneous pointer movements. This is done either by locking both NE and analyzer to an external source or, more practically, by locking the analyzer to the NE's BITS clock. The BITS clock is usually a DS1 for SONET applications or an E1 signal in SDH; most analyzers with frequency offsetting capability have connections for one or both signals.



**Figure 30.8** The pointer algorithm for a synchronous signal.

1. Bellcore GR-253, Dec. 94, Section 5.6.2.3.5.

Once the frequencies are locked, the offsetting can be done. Most analyzers have two main ways in which they will offset frequencies:

- *Line rate offset:* In this case the line rate of the signal emanating from the analyzer is offset by a specified amount. The output from the NE then can be examined to see if it contains the correct pointer adjustments (Figure 30.9).
- *Payload frequency offset:* The line rate of the analyzer is kept the same, but the data in the synchronous payload envelope (SPE) is derived from the offset clock. This will necessitate pointer movements in the signal source from the analyzer. The output from the NE can be examined to confirm correct processing of these movements (Figure 30.10). The most complex analyzers extend these offsets to individual VT/TU to check subrate pointer movements.

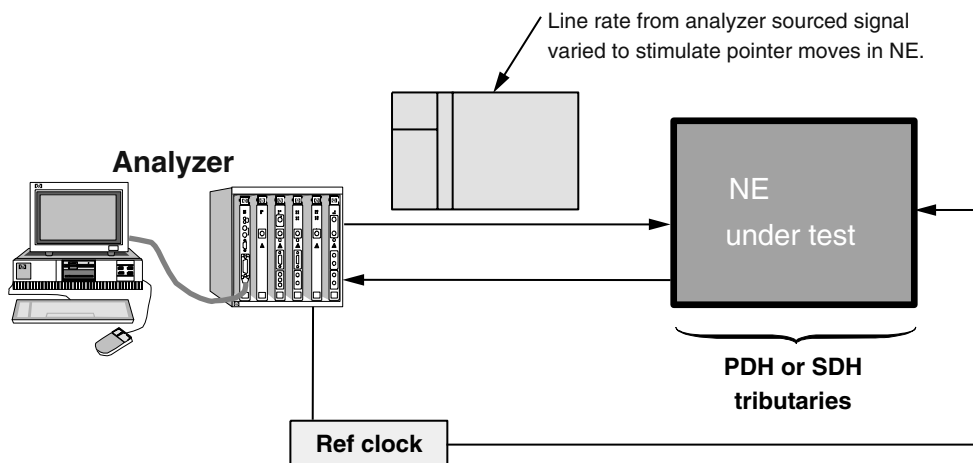


Figure 30.9 Line rate frequency offset testing.




Figure 30.10 Payload offset testing.

**Pointer alarms (LOP).** As Figure 30.8 indicated, the loss of pointer (LOP) alarm can be entered in a number of ways. These entry and exit thresholds can be stressed in exactly the same way as other alarms by use of simple “soak” testing, or, if more rigorous examination is required, a three-state sequencer (described previously) could be employed.

**Framing byte stressing.** The ability to frame up on a synchronous signal, even in the presence of errors in those bytes, is the most fundamental function of a synchronous network element. The means to test this algorithm is also basic to most analyzers.

As with all test capabilities, the type of stressing available will vary from analyzer to analyzer. In the simplest case, *framework stressing* is achieved by setting all the framing bytes (A1A2 in overhead) to an illegal value, usually all 1s. This method will test (crudely) the NE's loss of frame (LOF) and out-of-frame (OOF) alarm processing, and probably is sufficient for quick checks in production and I&M applications.

For other, more rigorous testing (R&D and compliance testing, for example), a closer examination of the framing algorithm is required. Analyzers used in these applications have the ability to error individual bytes in the framing pattern, either with single-shot errors or at specific error rates. A few analyzers can combine this individual bit erroring with the three-stage sequencing (previously described) to allow extremely complex examination of the framing algorithm.

**BIP error addition.** Another important feature of synchronous systems is the ability to identify bit errors in the data stream by means of the bit interleaved parity (BIP) bytes situated in various parts of the overhead. Checking the correct function of these bytes is a basic test function of most analyzers. When stressing for BIP errors, in general it is the BIP bytes themselves that are errored (so that the errors induced can be controlled exactly).

Errors are introduced in a number of ways:

- Single (random)
- Single (via mask)
- Error rates
- Bursts
- G.826 stressing

The single random BIP error is the simplest form of error an analyzer can introduce, and it represents a crude check of an NE's ability to process a BIP error. In addition, it often is used as a continuity check on a system.

Though the single random error is a good quick check, the BIP bytes are not a simple checksum of the bytes they cover. Rather, each bit in the byte is the sum of all the bits in that position in the signal. For example, the B1 byte is the BIP for all the bytes in the preceding frame. Bit 0 represents the checksum for all the bits 0 in that frame, bit 1 represents the checksum for all the bits 1 in the frame, and so on. In order to test this feature of the BIP bytes, most analyzers can inject single errors through a *bit mask*, so that error processing in specific bits or combinations of bits in the parity bytes can be examined.

The self-healing (protection switching) capabilities of the synchronous network are specified to trigger a defined rate of errors in the data stream. In order to test this, analyzers are equipped with the ability to generate specific error rates in the BIP bytes. In simplest form this could be preprogrammed rates set at the switching levels defined in the standards.

For greater flexibility, a number of analyzers have continuously variable error rates, so that both sides of an error threshold can be examined in detail. The method of generating rates also is worth examining. Some analyzers generate the rate on only one bit or a number of bits in a given BIP byte. The best analyzers will distribute the errors evenly over all the bits in the parity byte(s).

Though error rates are given an evenly distributed incidence over a period of time, a network or network element's reaction to a burst of errors also is of interest to some users. Accordingly, some analyzers are equipped to provide bursts of high error rates in a normally error-free data stream.

The quality of service parameters outlined in the ITU-T document G.826 are becoming increasingly important to SDH network operators. Analyzers used to test SDH NEs now are being asked to stimulate specific G.826 errors. The most significant of these is the severely errored seconds (SES) threshold. This is triggered when 30 percent of the frames in a given second contain at least one BIP error. As a result of this need, some analyzers now have the ability to error a given percentage of frames per second.

**Data errors.** Most analyzers, in addition to BIP errors, allow for the introduction of random errors within the data stream. The most popular techniques for adding the data are:

- Single (random)
- Error rates

As with BIP error injection, the best analyzers inject errors into all bytes in the payload and not a selective number.

**Payload stressing.** For confidence, most network equipment manufacturers need to test their equipment under "real" network conditions before handing it on to a customer. This can be achieved in two main ways:

- PRBS loaded signals
- Live traffic addition

*Pseudorandom binary sequences* (PRBS) are explained in detail in the chapter on bit error testing. Broadly speaking, however, they can be considered a predictive bit pattern that has the qualities of a "real" traffic signal. Their predictive nature allows bit errors to be spotted. PRBS loaded signals can be used in a number of ways.

- If a network element doesn't process the SPE, a concatenated SPE containing a PRBS can be used to quickly check for error-free operation under traffic conditions.
- If a full SPE (140 Mbps, DS3, etc.) payload is being carried, a PRBS can be contained within a correctly framed signal within the payload.

- If a subrate tributary (DS1, E1, etc.) is being transmitted using the VT/TU structures, the PRBS can be incorporated into the signal. (More usefully, a separate PRBS type could be carried in the other VT/TUs. This arrangement will allow the VT/TU under test to be readily identified.)

Most analyzers can perform some or all of these techniques of PRBS addition. Which types are applicable depends upon application and the type of testing required.

The ability to add live traffic also is important to some users. This requires the analyzers to act as add/drop multiplexers and add an asynchronous signal to the synchronous data stream.

### 30.3.3 Parametric measurements

Most of these measurements are described in greater detail in other chapters; they will be just here, with emphasis on the important criteria in synchronous work mentioned in passing.

**Jitter (line and tributary).** A fuller explanation of jitter can be found in Chapter 23. In general, the jitter at line rates in synchronous systems offers no surprises, the main difficulty in measurement being caused by the small levels permitted (0.1 p-p) and the high data rates (2.488 Gbps is common).

At the tributary levels, however, where the asynchronous signals are demultiplexed, a unique measuring problem exists. The introduction of pointers as part of the synchronous system has created a type of jitter peculiar to this transmission standard. The pointer slips of 8 or 24 UI at a time (in the line side) introduce high peaks of jitter into asynchronous tributary signals when they are demultiplexed from the synchronous signal. It is worth noting that jitter analyzers designed for measuring conventional asynchronous signals often are not capable of measuring these high instantaneous values. Care therefore must be taken when choosing an analyzer to measure this jitter.

**Wander.** Whereas jitter is defined as short-term variations in frequency, wander is defined as long-term variation. Measurement of this parameter in digital system is dealt with in Chapter 23. Wander is of particular interest in a synchronous system because the synchronizing clock is carried in the data stream, and therefore data stream wander represents variation in a system's synchronizing frequency. A wide template of wander values is defined in the synchronous standards.<sup>2</sup> Analyzers used in this application not only need to generate these levels of wander, but in long-term monitoring applications need to be able to withstand these variations.

**Pulse masks.** As well as the functional behavior, the physical characteristics of synchronous signals are also tightly specified.<sup>3</sup> Some analyzers offer the ability to monitor pulse masks, but more commonly this function is performed by high-speed digital oscilloscopes.

---

2. See ITU-T Recommendation G.707 and Bellcore GR-253, Section 5.4.

3. ITU-T Recommendation G.825.

### 30.4 Choosing an Analyzer

This section deals with general concerns and provides a brief guide to important features and specifications in each of the major areas of analyzer use. Synchronous analyzers are used throughout the life cycle of a network element, from initial design to installation and maintenance within a network. The features that are important in an analyzer vary with its area of use. The following section deals with five main areas of use and the relevant feature sets. Table 30.1 at the end summarizes the chapter.

#### 30.4.1 R&D applications

Research and development (R&D) covers a large area of testing. Testing at this stage is usually manual and extremely rigorous and exhaustive, concentrating as it does on a complete understanding of the component/PCA/NE under test. Broadly speaking, two types of testing occur, *functional* and *parametric*. The parametric tests (pulse mask analysis, bit error rate tests, etc.) usually are confined to the early development stages of a product, i.e. the design of specific circuits, PCBs, and basic ASICs. A synchronous tester required in this area usually is a parametric tester with some (limited) ability to construct and transmit synchronous signals. In addition, some simple alarm generation also might be necessary.

Two types of R&D functional testing can be identified, *leading-edge* and *steady-state*.

The leading-edge tester deals with the standards as they evolve, providing the new services as soon as they are defined. For this type of investigation it is necessary that all the bytes in the frame structure can be accessed, altered, and tested. The ability to perform rigorous and exhaustive tests on all parts of the structure is vital in leading-edge R&D applications, where new capabilities are constantly being addressed. Figure 30.11 shows a typical instrument used for leading-edge design, the HP 75000 Series 90 SDH/SONET analyzer. Based on the VXI standard, the instrument is modular, flexible, and capable of accessing all the EOC bytes. In addition, it has sophisticated stressing and capture capabilities that facilitate the rigorous testing required.

The user of steady-state R&D testing tends to be behind the evolving standards, concentrating on the well-established aspects of the standards, building second-generation equipment that often is cheaper and/or more compact than the leading-edge or first-generation devices. Cost then starts to become a more important factor in the choice of instrument. The range of bytes that might need to be accessible usually is a small, well-defined subset of the entire structure, so a less sophisticated tester probably is called for. Often the steady-state user is designing gear in an area where the standards are well-defined and the behavior of the protocol is understood; the designer is looking for a good source of synchronous signals to do the basic day-to-day testing on an engineer's bench. Typically the steady-state design lab will have access to a "high-end" or leading-edge instrument for final detailed testing.

#### 30.4.2 Quality assurance (QA) and verification testing

The next stage after (or towards the end) of the design cycle is QA and verification testing. This involves comprehensive testing to ensure correct function of an NE to



**Figure 30.11** Typical “leading-edge” R&D analyzer (*courtesy Hewlett-Packard*).

its design specifications, often coupled with rigorous climatic testing. The criteria for instruments used in this application are similar to those for R&D. The complexity of the tests often is lower, however, concentrating on testing the device against a set of limits rather than fully characterizing its performance. In addition, since test development time for QA is often a critical part of many projects, the programmability of a instrument and the ability to operate in an automated environment also are very important.

### 30.4.3 Compliance testing

Performed at the end of the design cycle to ensure that a new NE complies with the relevant synchronous standards, compliance testing sometimes is known as *Type Approval*. It involves the use of detailed and complex tests to exercise fully the NE's error detection, alarm reporting, protection switching, and pointer processing circuits.

Because this is an extremely detailed and complex task, it usually is not performed on every NE, but only on a sample at the end of the design cycle, or at a major revision of the software. Moreover, this testing is not restricted to network equipment manufacturers (NEMs); it also is performed by many network operators before installing a NE into their networks, and by independent testers (such as Bellcore in North America and C.S.E.L.T. in Italy) to allow certificates of conformance to be issued.

The complexity and detail of conformance/TA tests demand that only the most sophisticated testers be used. The functional requirements for a tester in this area are similar to those in the leading-edge R&D applications. Over and above the needs of functional testing, *automation* is another important factor in this area. The tests performed need to be both rigorous and reliable, and, although the first of these requirements can be met by manual testing, it is time-consuming and unreliable. Reliable and repeatable results demand that the tests be automated.

Further, the needs of regression testing mean that the tests must be repeatable at a later date. This regression testing element cannot be underestimated. Synchronous NEs differ significantly from the previous generation of asynchronous elements used for telecom networks; a large portion of the functionality (sometimes upwards of 60 percent) now resides in the software and firmware components of a network element. This functionality is therefore easier to change than before, and major revisions of software might appear every 6 months for a newly developed NE type. From this it is obvious that the need for accurate, repeatable, and quick regression testing is a significant need for those people engaged in compliance testing, and this cannot be achieved without automation.

#### 30.4.4 Production testing

The key objective for production testing is a high throughput of tested network elements. In this phase of the life cycle the network elements usually are measured at a lower testing level than that performed in the QA or verification cycle. In general this testing looks for manufacturing defects, and it usually is sufficient for the tests to exercise all the network element's circuits without specifically testing the total functionality.

The required feature set of the analyzer usually is limited. It often is sufficient to perform simple "soak"-type alarm tests and simple BER tests using PRBS in the synchronous data stream. Though rigor is not as large a factor as with conformance testing, here rapid throughput with consistent results also demands automated testing. A typical automated test stand is shown in Figure 30.12. Note that it includes not only a synchronous analyzer, but switches, oscilloscopes, voltmeters, and a PC controller to drive the automated tests.

Most of the time (hence money) needed to produce an automated test station is spent writing the tests, so the ease with which the analyzer can be programmed is of paramount importance. In addition, now it often is not sufficient that an analyzer can be remotely programmed (via RS232, HPIB, etc.); it is sometimes advantageous that the analyzer have drivers that can be used by a programming language such as HP VEE, or National's Labview.

Form factor also can be important in some cases, particularly where production space is limited. These test stands represent a large investment in both time and money; therefore it is important that the instruments can be upgraded as standards change. This is particularly relevant for the synchronous analyzer, because both SONET and SDH standards are still (slowly) evolving.





Figure 30.12 Typical production test stand.

### 30.4.5 Installation and maintenance (I&M) testing

Many if not most analyzers are used to install and maintain the synchronous networks, so the characteristics for this application include factors that are less important for other classes of testing, but are vital here.

First and foremost, I&M testers must be portable and ruggedized for use outdoors. Because the skill and knowledge level of installing technicians often is lower than that of the design engineers, the user interface's ease-of-use also is of great importance. A typical field analyzer is shown in Figure 30.13: The HP 37717C is portable, rugged, and has an intuitive user interface. This analyzer has been used in large numbers to install and maintain SDH networks. These analyzers are used throughout the network, and a great many may be required to cover a large network. As a result, the cost of an individual unit also is a large factor in analyzer choice.

The feature set required of the analyzer is similar to steady-state R&D features in that a limited and well-defined set of bytes needs to be accessed, and simple function testing also is required. In addition, in-service measurements with the analyzer acting in "through" mode (see section 30.2.2) also might be required. In this mode it might be necessary for the analyzer to be able to log results, and to perform QoS measurements such as those defined in the ITU-T G.826 document.



**Figure 30.13** Typical installation and maintenance analyzer (*courtesy Hewlett-Packard*).

Remote operation of the analyzer is becoming increasingly important for two reasons. First, synchronous NEs are often remotely sited, and it is an advantage to be able to access the monitoring data remotely. Furthermore, highly skilled technicians are a scarce resource, and remote access allows these technicians to be centrally located.

#### 30.4.6 Summary table

Table 30.1 is a summary of the preceding five sections, providing a quick lookup table of the important features in each application area.

### 30.5 Summary

This chapter has been a brief introduction to the synchronous analyzer. As has been shown, these analyzers vary considerably in capabilities and functionality. The choice of analyzer is largely determined by the user's application. Detailed information on analyzers is best obtained from the instrument manufacturers, all of whom produce detailed technical specifications for their products. In addition, some manufacturers also produce detailed application notes that describe the use of analyzers in different test situations. These application notes are often the best source of information for someone looking to build on the basic descriptions in this chapter.

**TABLE 30.1 Synchronous Analyzer Features by Application Area.**

Application Area	Important Features
Research and Development (R&D) <i>Leading-edge</i>  <i>Steady-state</i>	<ul style="list-style-type: none"> <li>• High functionality (including alarm/EOC sequencers, comprehensive capture, etc.)</li> <li>• Access to all bytes in overhead</li> <li>• Cost</li> <li>• Can be upgraded to future standards</li> <li>• Ease of use</li> </ul>
Quality Assurance and Verification	<ul style="list-style-type: none"> <li>• Rigorous testing capability</li> <li>• Easily programmed</li> </ul>
Compliance Test	<ul style="list-style-type: none"> <li>• Rigorous testing capability</li> <li>• Easily programmed</li> <li>• Can be upgraded to future standards</li> </ul>
Production Test	<ul style="list-style-type: none"> <li>• Easily programmed</li> <li>• Cost</li> <li>• Size</li> <li>• Can be upgraded to future standards</li> </ul>
Installation and Maintenance (I&M)	<ul style="list-style-type: none"> <li>• Size and weight</li> <li>• Ruggedized</li> <li>• Ease of use</li> <li>• Cost</li> <li>• Can be upgraded to future standards</li> <li>• Remote access</li> </ul>



## Signaling System 7 (SS7) Testing

Reid Urquhart

*Hewlett-Packard Ltd., South Queensferry, Scotland*

### 31.1 Introduction to Signaling

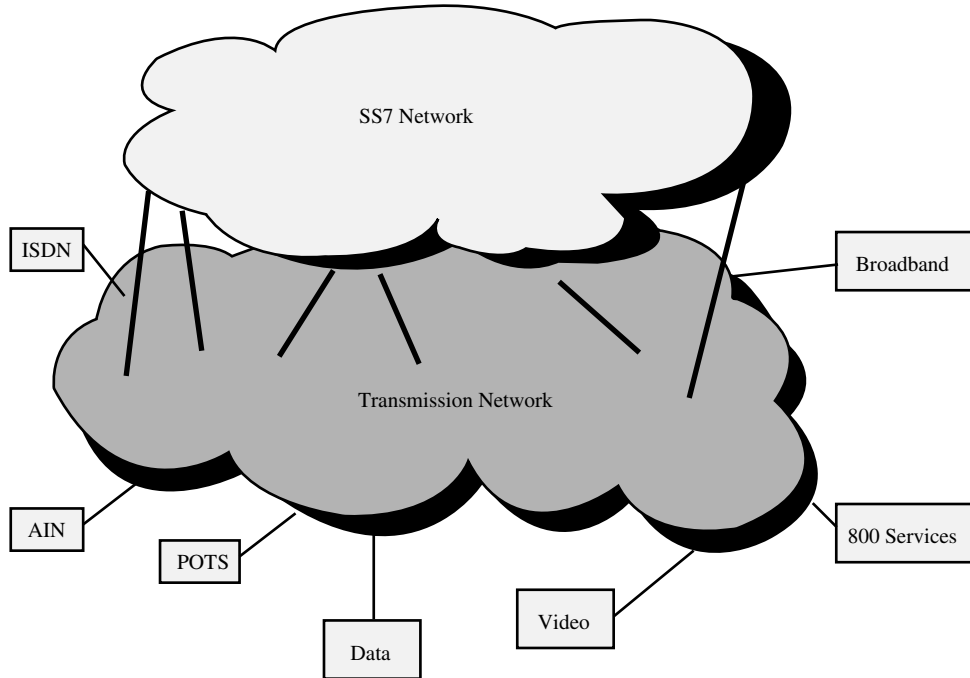
In a telecommunications network, *signaling* is the means by which the user is provided with appropriate network services. Call setup, delivery, teardown, and billing are all controlled by various signaling messages. The first form of signaling, *Channel Associated Signaling* (CAS), was transmitted from point to point on the same circuit used for message transmission. Thus the dialed digits traversed the telecommunications network, set up the voice paths from caller to called party step by step, and ceased when the called party answered. No further signaling messages were possible until the call terminated and the various parts of the voice path were released for use in other calls.

The need to provide a better form of signaling was driven by several factors:

- CAS is inefficient in its use of voice circuits because they are required during both call setup and teardown. They are required during setup even if the called party does not answer.
- Every circuit requires signaling equipment.
- Signaling messages cannot be transmitted during calls.
- New services, such as those where the number dialed is not the actual number called, are difficult to implement with CAS.

### 31.2 Common Channel Signaling

*Common Channel Signaling* (CCS) in effect uses a separate, high-speed data transmission network for all signaling messages. All these messages are transported on a “common” signaling network, which controls the switched network to provide users with the services they require. Figure 31.1 illustrates this concept, showing a



**Figure 31.1** The transmission network providing services under the control of the SS7 network.

separate signal network that sits above the transmission network, with the various connections required to control the transmission network resources.

This illustration is a conceptual one because the high-speed data links required by the signaling network are, of course, part of the transmission network itself. Great care is taken in the allocation of common channel signaling circuits, however, to ensure that transmission network failures do not disrupt the signaling network.

In the late 1970s the first common channel signaling system, CCS System No. 6, was formulated by the ITU-T (formerly CCITT). Despite initial worldwide enthusiasm, it was deployed widely only in North America. This system lacked the flexibility required by the demand for new services and was replaced by a new standard, *Signaling System No. 7 (SS7)*, which has been adopted worldwide.

Although SS7 is an international standard, there are various “flavors” in use throughout the world. While in North America ANSI and Bellcore set the standards, most of the rest of the world follows ITU-T recommendations. Special versions have been set up to support cellular networks, which require added complexity in messages, as well as message sequences to handle cell change-over for roaming users.

Flexibility in the standards definitions allows individual telephone operators or network equipment providers to customize the protocol to the extent that most countries now use their own versions. ITU-T standards are used at international gateways in most cases, however.

Both the North American and ITU-T standards have been continually developed since their inception as new features have been added and new types of service offered.

### 31.3 A Typical North American SS7 Network

The nodes in a signaling network generally are referred to as *Signaling Points*. Figure 31.2 shows the different types of Signaling Points and links found in a typical North American SS7 network. Each of the Signaling Points will be described in turn.

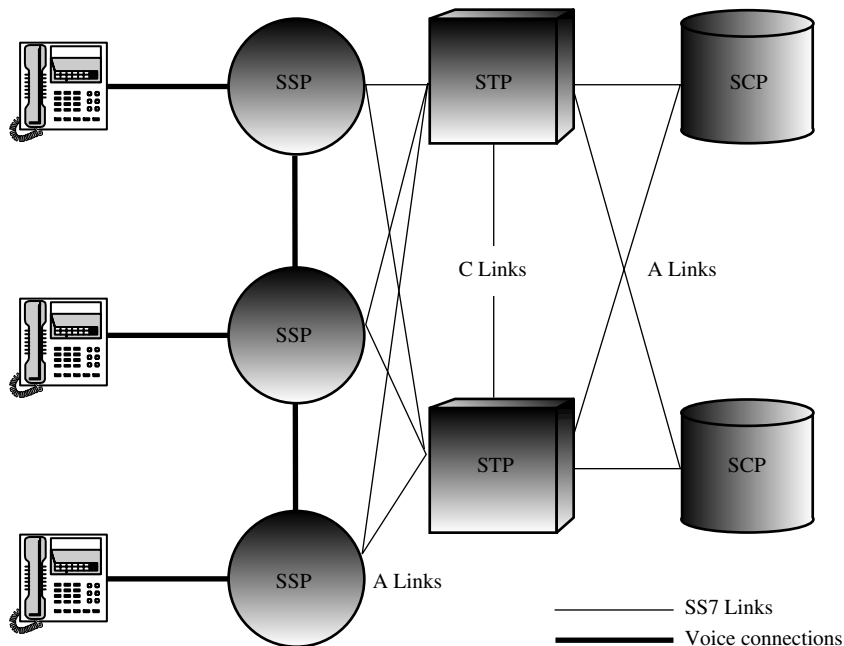
#### 31.3.1 Service Switching Point

The *Service Switching Point* (SSP), at the lowest level in the hierarchy, is the point at which customers connect to the network—the *End Office* or EO. An SSP is a transmission network switch connected to other SSPs or trunk switches as part of the transmission network. The SS7 connections control the switching function. In essence, the SS7 messages command the switch to set up the desired connections for incoming calls.

For outgoing calls, the SSP passes on to the next hierarchical level the SS7 messages required to set up the call, which are generated from the number dialed by the user. Like all other SS7 network elements, the SSP has a unique identifying number, known as a *point code*, which forms part of every SS7 message sent by or addressed to the SSP.

#### 31.3.2 Signaling Transfer Point

The *Signaling Transfer Point* (STP) is a packet switch, which performs the message-routing function in the SS7 network. The STP, unlike the SSP, is only an SS7



**Figure 31.2** The three elements of an SS7 network. The Service Switching point is a network switch, the Switching Transfer Point is a packet switch for SS7 messages and the Switching Control Point is a database.

## 728 Network Test Instrumentation

network element and plays no part in the transmission network. In the North American network, as shown in Figure 31.2, an STP is always one member of a pair. Each SSP is connected, via “A” links, to each member of an STP pair, providing a measure of redundancy in the event of a SS7 link or STP failure. In normal working conditions, the two members of an STP pair share the signaling traffic load equally.

Additional signaling links using different paths through the transmission network provide further diversity. For instance the single “A” links between SSP and STP may in fact be two or even four separate links, normally referred to as a *link set*. As with the STP pair, the members of a link set normally share the traffic equally.

### 31.3.3 Service Control Point

The *Service Control Point* (SCP) is a database that contains the information, required to complete certain kinds of calls. The most obvious example is the 800 call, where the 800 number dialed is not the number actually called. When a user dials an 800 number, a query is automatically sent via the SS7 network to an SCP, which responds by sending back to the originating SSP the actual number to be called. Using this number, the SSP then completes the call in the normal way. This process is fast enough to be invisible to the caller.

SCPs are the key elements in providing all kinds of advanced services. The number returned in response to the originating 800 call need not be fixed but can vary according to time of day, the location of the caller, the number of calls made to the 800 number over a period of time, and various other criteria.

## 31.4 The SS7 Protocol

The SS7 network is a packet-switched network in which individual packets are known as *Signaling Units* or SUs. Three different types of SU are commonly in use:

- LSSU
- FISU
- MSU

**LSSU.** The Signaling Points at either end of a link exchange LSSUs when a link is first being established or when it is being brought back into service following a failure.

**Fill-In Signaling Unit (FISU).** The *Fill-In Signaling Unit*, or FISU, is essentially an empty packet automatically transmitted between Signaling Points at either end of a link when there is no actual signaling traffic. Although FISUs do not contain any signaling data, they do contain the CRC checksums and so provide a continuous check on link quality.

**Message Signaling Unit (MSU).** The *Message Signaling Unit* is the packet that contains the actual signaling data (such as called number) required to control the transmission network.

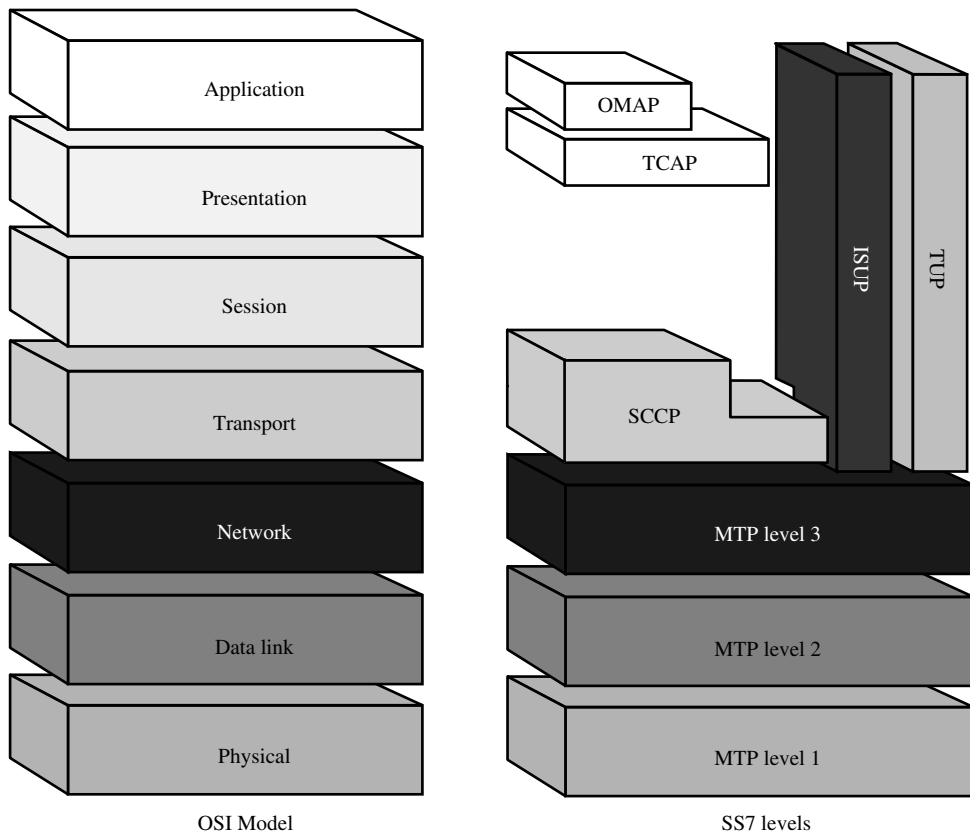


### 31.5 Protocol Levels

The SS7 protocol consists of several layers, known as levels. Since the development of the SS7 protocol predates the well-known OSI model, SS7 does not follow the model completely. Figure 31.3 shows the SS7 levels and compares the two protocol stacks.

**Protocol level one.** This is the physical level. Today this consists of a bidirectional 56 or 64 kbps data channel, although a higher-rate (1.544 Mbps) version currently is under development in North America.

**Protocol level two.** This level provides error detection and correction. Any Signaling Point sending an SU adds to it a cyclic redundancy check code, CRC. The receiving Signaling Point uses the CRC to detect any transmission errors and requests re-transmission of the faulty SU in the event of an error. Level two operates on a link-by-link basis between Signaling Points.



**Figure 31.3** Compares the standard OSI computer communications protocol stack with that used in the SS7 network.

**Protocol level three.** This level provides for the addressing of all SUs. Every MSU contains both an originating point code and a destination point code used by the network Signaling Points to direct messages to their destinations. Level three also contains various network management messages used to inform Signaling Points of network status. Level three controls the rerouting required in the event of congestion or link failures.

Levels one, two, and three together are known as the *Message Transfer Part* (MTP).

**Protocol level four and above.** These levels contain the actual messages required to provide the network users with the services they request. Various parts are defined depending on the kind of service required.

**The Signaling Connection Control Part.** The level three point code routing ensures that messages are routed from origination to destination Signaling Point, but does not contain the information required to address messages to a particular process within a Signaling Point. The SCCP level provides this and effectively supports connectionless message transport for database enquiries such as those required by 800 calls.

**The Telephone User Part.** The TUP controls the setup, management and release of straightforward point-to-point telephone calls. It is most commonly used in telephone networks outside North America.

**The Integrated Services Digital Network (ISDN) User Part.** From its name it can be seen that the original purpose of the ISUP was to set up ISDN connections in the same way as the TUP described above. In North America, however, both straightforward voice calls and ISDN connections are set up by the ISUP; the TUP has never been used.

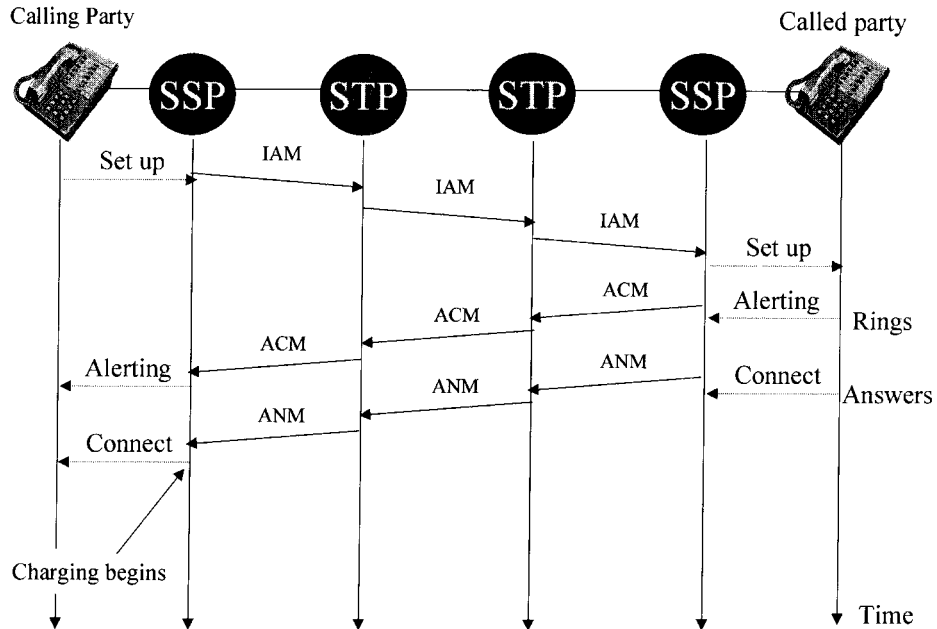
**Transaction Capabilities Application Part.** The TCAP supports intelligent network service by enabling the exchange of information between Signaling Points using the SCCP connectionless service to deliver the messages. For instance, an SSP uses a TCAP query message to get the routing number associated with an 800 call from an SCP. The SCP, in turn, uses TCAP to return the response to the SSP, which then processes the call in the normal way.

### 31.6 Message Sequences

Most SS7 tasks, such as call setup, are performed by sequences of MSUs exchanged between two Signaling Points. These sequences frequently are drawn as *ladder diagrams* in which the ladder uprights represent the Signaling Points exchanging messages and the ladder steps represent MSUs. Figure 31.4 shows one of the simplest but most commonly occurring sequences, a simple call setup.

The illustration shows a call setup in which the called party is available and answers the call. The first six steps are illustrated in the diagram; steps 7 to 10 take place when the call is terminated or busy and are not shown on the ladder diagram.

1. The calling party dials the telephone number on his or her phone.
2. The originating End Office accepts these digits, checks the routing table, and encodes the dialed number into an Initial Address Message (IAM) with the appropriate originating and destination point codes.



**Figure 31.4** The sequence of messages passed between the various Signaling Points in the SS7 network when a simple telephone call is set up.

3. The IAM is transported through the SS7 network via the appropriate SSPs to the terminating End Office.
4. The terminating office responds to the originating office with an Address Complete Message (ACM) and rings the called party's telephone.
5. When the originating End Office receives the ACM, it sends ringing tone to the calling party's telephone.
6. When the called party answers, an Answer Message (ANM) is sent from the terminating office to the originating office, the connection is made and, charging for the call begins.
7. If the calling party hangs up first, the originating office will send a Release (REL) to tear down the call and release the trunk circuit.
8. Upon receiving the REL, the destination office transmits a Release Complete Message (RLC) back to the originating office, finally terminating the billing and releasing all trunks.
9. If the called party hangs up first, the REL will be sent in the opposite direction to tear down the call and release the trunks.
10. If the called party is busy, then the called party End Office will send a REL message with a release cause code of "busy" back to the originating End Office. The originating office will then supply the "busy" tone to the calling party.

## 732 Network Test Instrumentation

### 31.7 Testing SS7

The testing of SS7 networks can be looked at in terms of the tests required at the various protocol levels.

**Testing at level two.** Since level two information is created on a link-by-link basis, level two problems are confined to the link on which they occur and do not propagate through the network. Level two hardware is well established, so failures in a working network are most likely to have a physical cause and can be cured by card replacement.

**Testing at level three.** Level three problems are likely to be more subtle. They can be caused by software defects or by errors in routing tables. Traffic might be misrouted, messages might be lost or sent to the wrong destination, and, in the case of erroneous management messages, working parts of the network might shut down unnecessarily.

Level three test equipment must decode the relevant message octets, search for particular messages, and filter data being captured to concentrate on problem routes. A protocol analyzer that can decode the 1s and 0s of the SS7 bit stream in terms of the protocol definitions and produce the English-language octet descriptions and mnemonics can make a major contribution in this type of testing.

**Testing at level four and above.** These are the most complex parts of the SS7 messages; it is at these levels that messages combine to form sequences and provide services. At level three, sophisticated protocol analysis is essential. At these levels, the capability to trace the sequences of messages, which set up a call or fetch data from a database, is invaluable in tracking down many problems.

Many problems appear only at high load levels; emulation testing, which can produce typical message sequences surrounded by varying load levels, is required to probe networks and Signaling Points to detect any load-related problems.

### 31.8 Continuing Developments

Duplicated signaling links normally are designed to operate with a maximum load of around 40 percent so that each has the capacity to carry all traffic in case of a link failure. Many links are approaching this signaling load level as traffic loads and demand for new services increases.

#### 31.8.1 Local Number Portability (LNP)

The 1996 Telecommunications Act has mandated that network users must be allowed to take their telephone numbers with them when they change suppliers. This commonly is referred to as *Local Number Portability* (LNP), and it will have a profound effect on SS7 networks. LNP means that for normal point-to-point calls, the number dialed (which until now defined the geographic location of the End Office and of the actual telephone addressed) actually can belong to a phone anywhere in the covered region.

### 31.8.2 Database Lookup

All calls ultimately will require a database lookup in an SCP, which contains a table of the telephone numbers and their current location, increasing SS7 traffic dramatically. High-speed SS7 links, which operate at full T1 rates (1.544 Mbps) have been developed and will be brought into service as quickly as possible.

### 31.9 Conclusion

The SS7 protocol, initially developed in the late 1960s to improve telephone call setup, has stood up extremely well to the increasing demands of network expansion, new services, and regulatory changes. It seems set to continue to serve telecom networks in the future.



---

Part

**8**

# Network Management





---

Chapter  
**32**

## Local Area Network Management and Performance Monitoring

**Bill Tomeo**

*Hewlett-Packard Co., Colorado Springs, Colorado*

### 32.1 Introduction

Networks are not self-sustaining. They require the devoted attention of a person or a group of people to implement, maintain, troubleshoot, and grow them. All of these tasks require tools and methodologies that are distinct yet interrelated. For example, the way in which a network is implemented can have a direct impact on both the number and type of problems that occur during network usage. Also, maintenance and troubleshooting problems experienced on a network help influence how a network should be grown, and what is required to accomplish controlled growth.

These issues are not new and unusual to the use of computing resources within organizations. In fact, network management has its roots in data center management. The use of computing power to conduct the day-to-day business of an organization began in the mainframe environment. As the computing power paradigm shifted from centralized processing to the combined power of computers connected by communication lines, the management intentions remained the same.

#### 32.1.1 Functions of network management

Network management encompasses a broad range of functions:

- *User and software administration:* Configure and administer users and applications so that needed functionality is available and accessible to the correct subset of users.
- *User support:* Provide consulting for users on computer and application use and issues.
- *Security:* Monitor information access to ensure data integrity.

**738 Network Management**

- *Fault management:* Fix computer systems and networked components when they fail.
- *Performance management:* Adjust the availability and efficiency of computer resources to accommodate the volume of users and information necessary to conduct the business of the organization.
- *Planning:* Anticipate the need for future resources to keep network resources and components running smoothly.

**32.1.2 Factors driving network complexity**

Coming into the decade of the 1990s, the foundation of network management had been laid around these fundamental and evolving needs. The challenge, however, became two-pronged. The same challenges of the mainframe environment had to be addressed, categorized by management function. The added challenge included:

- The complexity introduced by the number and variety of technologies that converged to make networks viable.
- The variety of equipment needed to implement a network.
- The number of vendors supplying the equipment.
- The geographic distance that separated users from one another.

The luxury and simplicity of single-threaded control in the mainframe environment rapidly unwound as distributed computing dispersed and connected critical computing components across buildings, cities, and countries. As networks spread across physical and geographical boundaries, the list of management issues inherited from the mainframe environment essentially remained the same.

What became problematic, however, was how to apply them to control resources, applications, and users spread across the globe. Networks were evolving, from a medium used to share resources like printers and files, to a strategic business advantage where the flow of mission-critical applications became the lifeblood of an organization. The true measure of network worth to business resided in its reliability. Network management quickly became a necessity, a critical component in the value equation. One accomplishment that facilitated network management was the development and burgeoning deployment of client-server architecture.

**Client-server architecture.** As its name implies, client-server architecture distributes processing between:

- *Clients* as workstations or nodes on the network that request information (pieces of data from a database) or services (access to printers, facsimiles, plotters) from a centralized network resource.
- *Servers* as computers that store data (in centralized databases) and programs (with code that can be accessed by many users from a centralized resource), and provide network-wide services to clients (spooling, peripheral device manage-

ment, file routing). Servers embedded in large internetworks, networks that link several geographically-dispersed sites, run special network operating systems on dedicated server computers. Smaller networks may only need a personal computer running its own operating system with peer-to-peer networking software.

Client-server architecture lends itself to the many different computing work profiles that exist in organizations:

- Local, non-network applications
- Client-based network applications
- Server-based applications
- Split-functionality applications

*Non-network applications* still can reside on client workstations, accommodating software programs not offered in network-compatible versions and allowing for unique uses of software not needed by the vast majority of end users. There also can be *client-based applications* where the application software resides on an individual client, but there is the need and ability to access and share data among users. The prime example of this is spreadsheet applications and data.


*Server-based applications* use the greater computing capacity of servers to move and/or share information between network end users. For example, a large part of the processing for electronic mail takes place on the mail server. It is the central depository for mail that needs to be sent and for mail that has to be delivered. Most electronic mail code resides on the mail server.

*Client-based user interface to server-based searching* allows the processor-intensive storage and retrieval of information to reside on the server, while a “shell” interface runs on the client. (These functional modules sometimes are called the “back end” and “front end,” respectively.) Each client can request speedy retrieval of specific data when it is needed, ensuring data integrity of information that is shared among network users. Database-intensive applications, like order processing or customer service systems, are perfect applications for this profile.

Client-server architecture results in lighter processing loads for client workstations or nodes and increased loads on servers. Servers typically have maximum disk capacity for storage, robust memory, larger caches, and possibly several processors. Servers can be PCs, minicomputers, or mainframes. Servers are mission-critical resources for any organization employing them in any profile. Figure 32.1 depicts typical client-server architecture.

Other environmental factors that had an effect on the rapid adaptation and deployment of networks to conduct all aspects of a business are:

- Economic considerations
- Technical considerations
- Societal considerations
- Political considerations



**Figure 32.1** The advent of client-server architecture concentrated at the server the need for processing power, disk space, and access to widely used applications. Client nodes could pick and choose among the applications available at the server. This resulted in consistent use of applications throughout an organization, and savings in disk space and processing power at each client node. This architecture also made it easier for information technology (IT) technicians to control, maintain, and troubleshoot application problems.

**Economic considerations.** The market demand for networking services created by the massive shift from mainframes and minis to PCs was a by-product of ever-increasing capacity and speed available at prices that steadily dropped. Global competitiveness gave rise to the “PC clone,” capturing a majority of the PC market in the early 1990s. Historically, PCs were purchased because improvements in productivity offset their capital costs. After a decade of installing millions of cheap, high-powered microcomputers, it was time to harness them into company-wide networks for sharing information. In the mid-1990s, over 30 million PCs were wired to LANs in the United States. The true value of the network is its ability to move information efficiently. The value of the information that the network carried quickly exceeded the cost of the components needed to implement the technology.

**Technical considerations.** Standardization in large part accounts for fueling the network management revolution. The beginning of the 1990s saw the definition and rapid adaptation of the Simple Network Management Protocol (SNMP). SNMP was designed specifically to manage network implementations based on Transmission Control Protocol and Internet Protocol (TCP/IP) networking protocols. Simplicity was the strength of SNMP. It was easily implemented and consumed modest processor and network resources. SNMP provided multivendor interoperability between monitoring products and management stations, allowing users to mix and

match network monitors and management stations from different vendors. A more detailed explanation of SNMP is provided later in this chapter.

**Societal considerations.** Many other business forces are forging the re-engineering of traditional information systems to client-server systems. Higher employee productivity and morale, more competitive response to market changes, better customer service, and commitment to total quality also are driving the escalation of client-server systems. Forced to contend with growing networks and shrinking pools of qualified staff, network executives need sophisticated, proactive network management systems that can help keep businesses running smoothly.

**Political considerations.** In the early 1990s, the call for an “information superhighway” initiative brought the concept of networking to the masses. Proof of this is the unprecedented number of connections to network information-access systems like CompuServe and Prodigy. The Internet originated as a cooperative effort between academia and government and was the incubator for enterprise networks. It is the prime example of an interlinked web of networks penetrating over 100 countries with 11,000 separate networks feeding into it, containing up to 1.7 million host computers.

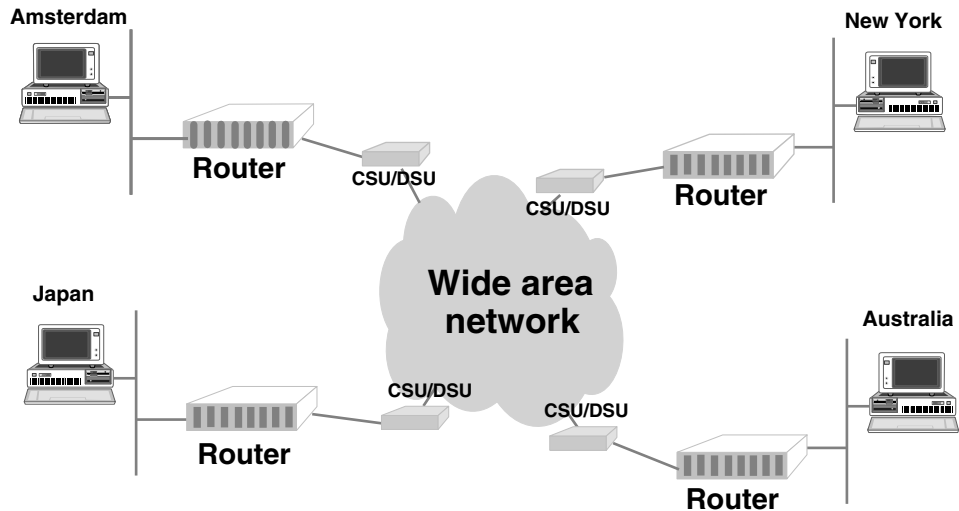
## 32.2 Network Management Boundaries

A network is any number of computers connected by communication channels for the purpose of sharing peripheral devices, centralized software, data, and/or disk space. A local area network (LAN) can be self-contained and small, linking together workers over limited distance such as within a single department in a building or between two buildings. A network also can connect clusters of LANs across geographic distance over permanently installed cables and/or dial-up lines through wide area network (WAN) technology. Figure 32.2 shows a network that consists of clusters of LANs in different cities connected through WAN technology.

### 32.2.1 Network building blocks

Network management embraces all the separate components necessary to build networks; devices, communication technologies, and the enabling communication standards. Network devices are attached to the physical shared medium to establish information access, to control physical access, to extend the reach of the network, and to switch traffic. Basic components include the following:

- **File server:** A computer connected to a network that serves as a central repository for application software, databases, and data. The resources of a file server are accessed by workstations (or clients) attached to the same network. It also serves as the traffic director for workstations seeking access to centralized peripherals.
- **Workstation:** An intelligent terminal, a personal computer, attached through a communication link to a network to achieve access to greater computing resources (file servers), printers, plotters, scanners, or other peripheral devices. (Also referred to as a node or station.)
- **Bridge:** A hardware device used to connect LANs of the same architecture so that they can exchange data. A bridge effectively serves to extend the physical length of a LAN. A bridge reads the address of every packet that flows between the



**CSU - Channel Service Unit** ensures that all signals placed on line are appropriately timed and formed.  
**DSU - Data Service Unit** formats data for transmission on the WAN, ensuring that carrier's requirements are met.

**Figure 32.2** Wide area network (WAN) technologies, like X.25 or frame relay, enable a single organization's system of local area networks to be connected into a global enterprise network.

interconnected LANs. Bridges operate at layers 1 and 2 of the Open Systems Interconnect (OSI) Reference Model, providing connectivity.

- **Hub:** A hardware device that makes it possible to attach additional workstations to a network, either actively through extending cable length by amplifying transmission signals, or passively by splitting the signal to increase number of attachments at the expense of distance. Hubs accommodate multiple LAN architectures over separate ports.
- **Router:** An intelligent hardware device that connects both like and different LAN segment architectures and serves as a central receiving point for traffic (or packets), directing them to the correct destination LAN. Routers operate at layers 1 through 3 of the OSI model to provide connectivity, addressing, and switching.
- **Switch:** An intelligent hub that supports multiple LANs (same or different architectures), multiple media, multiple speeds, and multiple LAN protocols. Switches provide bridging and basic routing capabilities. Switches can provide full LAN bandwidth to multiple, simultaneous communications on a point-to-point basis.
- **Brouter:** A combined bridge and router.
- **Gateway:** A combination of hardware and software that connects LANs to WANs or larger systems where communication protocols differ. A gateway has the intelligence to perform protocol conversion at all seven layers of the OSI model. Slower than bridges and routers, gateways can be the cause of traffic congestion on busy links.

**TABLE 32.1 A summary of characteristics for some of the most commonly deployed LAN networking standards demonstrates the choice and complexity that network managers face. The variety of standards used to communicate on the network allows great flexibility of choice, but introduces complexity in network management.**

	Ethernet	Fast Ethernet	Token-Ring	FDDI	ATM
Standard	IEEE 802.3	IEEE 802.3u	IEEE 802.5	ANSI X3T9.5	ITU-T
Speed	10 Mbps	100 Mbps	4, 16, 20 Mbps	100 Mbps	25 and 155 Mbps
Logical topology	bus	bus	ring	ring	
Physical topology	bus, star	bus, star	ring, star	dual ring, dual star	
Traffic	data	data, voice, video	data	data, voice, video	data, fax, voice, video, image
Transmission	baseband	baseband	baseband	baseband	broadband
Media	coax, UTP, STP	coax, UTP, STP, fiber	coax, UTP, STP	fiber	fiber, UTP, STP
Media access control	CSMA/CD	CSMA/CD	token-passing	token-passing	transmission convergence

### 32.2.2 Communication technologies

The Institute of Electrical and Electronics Engineers (IEEE) is the standards body for LAN communications. They are responsible for defining and approving the Ethernet standard (IEEE 802.3) and the Token-Ring standard (IEEE 802.5). Another widely used LAN standard is Fiber Distributed Data Interface (FDDI), developed by the American National Standards Institute (ANSI), standardized as X.3T9-5.

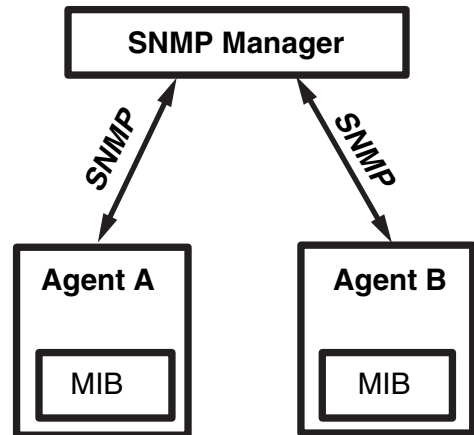
These traditional LAN technologies are giving way to a new breed of fast LAN technologies like “Fast Ethernet” (also known as 100Base-T), 100VG-AnyLAN (in which *VG* means “Voice Grade”), and Asynchronous Transfer Mode (ATM). The steep increase in number of nodes linked to LANs is a surface indication of the traffic growth. Users are demanding higher speed and bandwidth-intensive applications—like simultaneous video, voice, and data—from their networks. Table 32.1 summarizes traditional and some fast LAN technologies.

### 32.2.3 SNMP: A standard communication protocol

The Simple Network Management Protocol (SNMP) provides the foundation for network devices to communicate using the network. It is a standard protocol, part of the TCP/IP suite, used to manage and monitor nodes on a network. It specifies who on the network sends information (an SNMP agent) and who receives it (an SNMP manager). It also defines a standard type of information that is passed between sender and receiver. There are two versions; the first is called simply *SNMP*, and the second is *SNMPv2*.

## 744 Network Management

1. **Manager**— Manages one or more agents. Normally resides on a computer and communicates to one or more agents in the network.
2. **Agent** — Collects data. Responds to get/set statements from an SNMP Manager and sends traps to specified destinations.
3. **MIB**—Resides on the agent and stores information for later recall by an SNMP manager.



**Figure 32.3** A manager communicates with agents distributed around the network. The agents collect and store information in databases designed specifically for one purpose, called *Management Information Bases* (MIBs). SNMP specifies exactly how the communication between manager and agent transpires.

**SNMP.** SNMP has three basic components; their relationship is shown in Figure 32.3. The *agent* (the collector and sender of information) is a software program that resides in a managed *element* (or device) of the network, such as a hub, router, or a specialized dedicated data collector. The *manager* (the receiver of information) communicates with the agent using SNMP commands, usually to query an information repository called a *Management Information Base*, or *MIB*.

MIB refers to a database that resides within an agent and holds information in the various *fields* that the database is designed to address. For example, *Remote Monitoring* (RMON) is a MIB with a specific set of fields that define the data set of information to be collected for monitoring traffic on a LAN segment up to the Data Link layer (layer 2) in the OSI model. MIBs are hierarchical in nature and can be viewed logically as a tree structure. Figure 32.4 is an illustration of the MIB tree.

The “branches” of the MIB tree are called *variables* or *MIB objects*. Unique identifiers or names are assigned to each branch. Each MIB has a unique identifier called an *RFC number*. (The term RFC refers to *Request For Comment*, a type of Internet Engineering Task Force document that defines network protocols and standards.) Some examples of MIBs are MIB-II for basic system information and basic level statistics (RFC1213), and the Remote Monitoring MIB (RFC1757).

Communication between agent and manager takes place with five types of commands:

1. *Get* request
2. *Get-Next* request
3. *Set* request
4. *Get-Response* message
5. *Trap* message

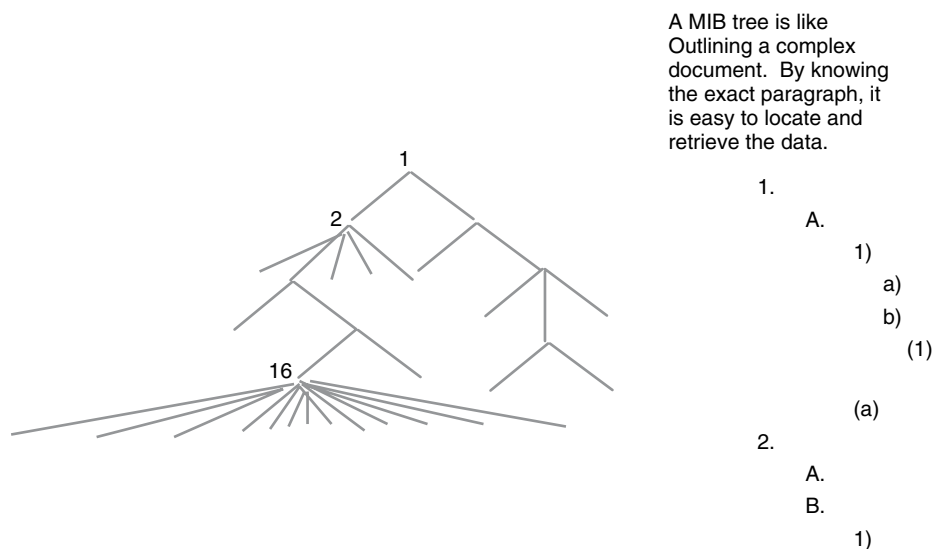


In a *Get* request, a manager requests (from an agent) the value of some variable stored in a field of the database that represent a particular MIB. The *Get-Next* request is used by a manager to request information on multiple variables, efficiently using the network (reducing traffic) to communicate more than one piece of information at a time. If any one variable is unavailable, no values are returned. This command is also used to retrieve unknown rows if available.

In a *Set* request, a manager instructs an agent to set a specific MIB variable (one of the fields in the database) to a desired value.

A *Get-Response* message is sent by an agent as a response to a *Set* request or *Get-Next* command. It can be an error message; if not, it either will be identical to the variable sent in the *Set* request (to show that the request was accepted), or be a *Get-Next* response with actual values, as requested, filled in. The manager checks its list of previously sent requests to locate the one that matches this response and, if none is found, the response is discarded; otherwise it is processed further.

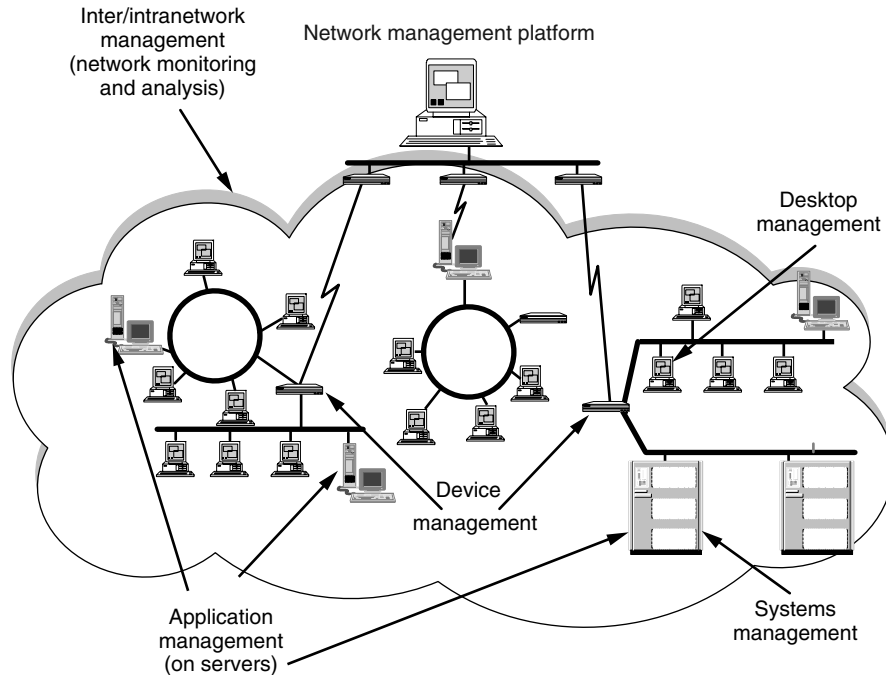
A *trap* is one of two unsolicited messages sent from an agent to a manager, often used for out-of-norm event notification when the agent “sees” something that the manager should know about.



A MIB tree is like Outlining a complex document. By knowing the exact paragraph, it is easy to locate and retrieve the data.

Each paragraph of the tree has an identifier. In SNMP terms, this is called an Object Identifier or system OBJECT IDENTIFIER, such as "1.3.6.1.2.1.1.1".

**Figure 32.4** Given the address of an MIB, the tree structure makes it easy to locate and retrieve data stored there. This makes the SNMP structure flexible enough to accommodate the growth in need for network management information. It also enables proprietary MIBs to be defined and used for unique management needs in the absence of standards-based definitions.



**Figure 32.5** Network management encompasses aspects of both physical topology and network operations. *Desktop management* refers to workstations, such as PCs, micro- and minicomputers, where end users actually access network resources. *Device management* refers to interconnect devices, such as routers and bridges, that enable traffic to be transferred accurately from one place in the network to another. *Application management* refers to the business software that resides on servers throughout the network, enabling end users to perform diverse and numerous business tasks that require computing resources. *Systems management* refers to centralized computing resources, like servers, that provide the horsepower for the network and for applications to be made accessible to all users on a network. *Network management* refers to the management of the traffic on the communication lines.

These commands allow network management applications to request and receive information from all the elements of the network that require management. Figure 32.5 shows the elements of the network that require network management: desktop, device, application, database, systems, and inter/intranetwork management. Aspects of each are an important part of every major network management area: user configuration and change, security, fault management, performance management, and planning.

It is critical that any network management scheme provide visibility, access, and control into the building blocks of the networks as described later in this chapter. SNMP is the enabler. The network management station becomes the SNMP manager; embedded within devices, or connected to LAN segments as separate devices, are the agents that send critical information back to the manager.

**SNMP version 2.** Without the SNMP standard, the adoption and usage growth of network management applications would not have progressed as rapidly as it did. The need for tools to cover all the various network management functions was and is

more critical than the need for tools to conduct business using the network. SNMP facilitated the development of network management applications because the protocol:

- Is easy to implement
- Requires few resources to operate
- Is mature, stable, and well-understood
- Is widely available (on most computers), and most network devices have some form of agent/MIB embedded

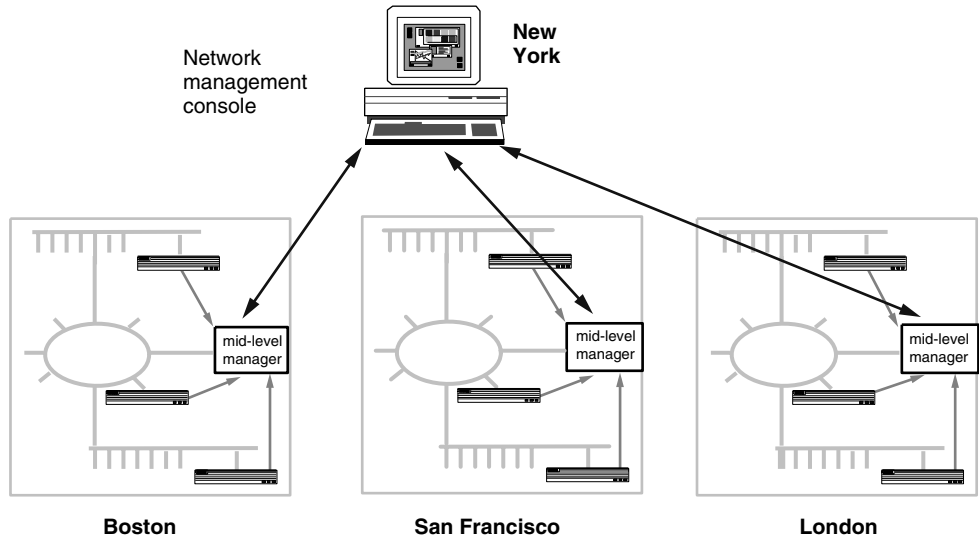
As networks have grown and the need for network management has become more imperative, however, several disadvantages of SNMP have become apparent. Some of these disadvantages are:

1. Because SNMP is a polling-based protocol, each agent is responsible for sending information to the manager as requested. Depending on the number of agents deployed on a network, the amount of traffic this generates can be burdensome and resource-intensive, both in bandwidth utilized and in processing power needed to analyze the information.
2. As more and more mission-critical information is moved to the network, unauthorized access to sensitive information is becoming a high-priority item across all industries. SNMP has limited security provisions.
3. Table data (blocks of data) cannot be requested to be sent from an agent. This prevents the transfer of multiple variables at once, further contributing to unnecessary traffic on the network.
4. Traps from an agent notify a manager that some event of note has been captured on the agent, but SNMP did not provide for an acknowledgment to be sent back to the agent that the event had been received. Because traps are unsolicited, initiated by the occurrence of an event on the agent, it is becoming increasingly important that the manager acknowledge receipt of the trap (typically an out-of-norm condition detected by the agent).

**SNMPv2 improves on SNMP.** SNMP version 2 (SNMPv2) is an addendum to the original SNMP standard, designed to address practical shortcomings and the evolving needs of emerging technologies.

SNMP defined manager-to-agent communication. SNMPv2 builds on this by defining manager-to-manager communication, thus addressing the need to scale network communications to sprawling networks. This deals directly with the problems stemming from using the network to manage the network. The volume of management information flooding the network infrastructure and management console from every segment can, conceivably, appropriate a large part of the network bandwidth for its own use.

Manager-to-manager communication will allow an intermediate level of managers to be inserted into large enterprise networks, sitting between the network management console and the numerous agents. Each of these “middle managers” will manage a subset of the agents on the network. This is called a *management domain*.



**Figure 32.6** A “manager-of-managers” architecture inserts an intermediate level of management between a centralized network management console and numerous agents dispersed throughout a large, enterprise network. These intermediate managers are responsible for collecting, storing, and analyzing SNMP information for different segments of the network, referred to as *domains*. Because each intermediate manager is responsible for a portion of the network, the central management station is relieved of the burden of monitoring each and every device throughout the network. The central manager has a view into the entire network through access to the consolidated information provided by the intermediate managers.

These mid-level managers then will pass relevant information about their domains to the central network management console.

This architecture reduces the number of connections to the central manager to one per domain. It enables distributed network monitoring and analysis to be *scaled* to fit the size of a network while allowing network capacity and processing power to be channeled to business information needs, not network management. Figure 32.6 is an illustration of this “manager of managers” architecture.

Security issues addressed by SNMPv2 include timestamping data collected so that an indelible audit trail is laid down. A data encryption standard would provide a screen for company-sensitive information traversing networks that are vulnerable to unauthorized listening. Resolving this issue might remove the single largest impediment to purchasing with a credit card over the Internet. Message authentication would prevent rogue users from masquerading on the network undetected. All of these features are critical to continued use of the network to conduct business.

SNMPv2 will allow whole tables of data to be passed to the manager. A new type of request, a *GetBulkRequest*, should alleviate congestion on the network and promote the use of tables in network monitoring and analysis.

Provisions have been made in SNMPv2 for acknowledging traps back to the agent. The agent will retry trap notification until successful, or, if unsuccessful, will notify the user that the manager is not responding. This introduces greater confidence that out-of-norm conditions detected by the agent will be communicated and dealt with in a timely manner.

## 32.3 Areas of Network Management

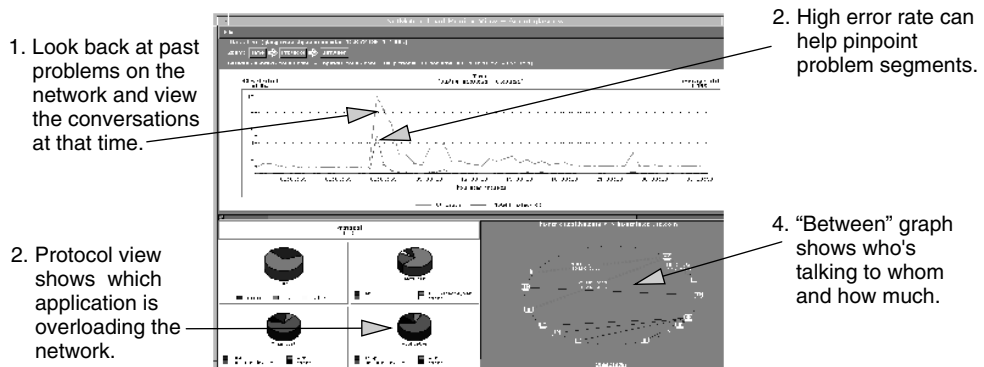
As the primary enabler for network management, SNMP specifies how the network's vital signs will be carried from the extremities to a centralized network management console. At the console, a variety of software will process the data and present it to network staff. Based on the information, network staff must intervene with decision-making and action. They carry out tasks for the same areas of concern that were specified as crucial in the mainframe-based, data center environment. Figure 32.7 summarizes the major areas of network management.

### 32.3.1 Configuration and change management

Adding, maintaining, and deleting users on a network is no small task. The hierarchical nature of networks makes this a continual challenge. At the lowest level, users are formed into workgroups. Within the workgroup, a user may need access to programs and databases on the file server, as well as centralized resources like printers, plotters, or fax machines. As workgroups are interconnected, information is shared locally, regionally, and throughout the enterprise. Proper access and configuration must be applied for users throughout this spreading web of resources. The administration of application software parallels this same model.

Some of the associated efforts that network management must cover include:

- Adding new users and assigning unique network addresses to each user.
- Maintaining existing user information, especially as users switch from one network segment to another.
- Deleting users who no longer have authorized access.
- Maintaining an up-to-date network map showing all workstations and network devices actively connected.



**Figure 32.7** The major areas of network management, each a discipline in its own right, are interrelated. Information deemed necessary in one area invariably proves indispensable in other areas. Correlating data from different areas also can provide insight into one or more of these areas, not possible when distinct areas of management information are maintained.

**750 Network Management**

- Monitoring active network users for immediate visibility into downed nodes.
- Adding new software applications at appropriate places in the network (work-group, campus, or enterprise).
- Maintaining software by applying upgrades and revisions on a timely basis.
- Defining allowed access levels to specific applications for specific users or user classes (workgroups).
- Defining efficient backup procedures to minimize data loss in the event of network failure.
- Defining redundant communication paths in the event of temporary loss of communication access.

**32.3.2 User support**

User support is the critical path to successful and productive use of the network. End users cover the spectrum of computer literacy, experience with application use, perceived satisfaction with network performance, need to access the network, and ability to solve problems. The term *help desk* has pervaded the application software and networking arenas as a way to describe the vast array of tasks that make up user support. It describes a centralized function that exists in organizations where networks have become a strategic enabler for doing business day-to-day. User support encompasses the following issues:

- Access to help staff
- Training of help staff
- Use of support tools

User support means granting access to staff expert in solving surface-level issues and problems that cause users to experience a delay in network use or interruption in productivity. These issues include but are not limited to network connectivity, hardware malfunction, network access configuration, application software or database corruption, inaccessible or unavailable resources, or a lag in response time from network resources.

It means training help desk staff so that they are knowledgeable about networking and application use issues. Training should also include logical problem solving methodology and a foundation in networking standards.

It also means implementing intelligent tools to assist help desk staff with user support, including trouble-ticket systems, management-by-exception tools, and reporting capabilities.

*Trouble-ticket systems* automatically generate a work order or trouble ticket for operations staff responsible for fault and performance management. The ticket is generated based on predefined acceptable operating conditions specific to a particular network. When the defined thresholds are exceeded, the trouble ticket logs location, time, user, node address, and any other data items that might prove useful to operations staff in problem-solving. Optionally, packet capture of network traffic can be triggered by the threshold event, providing even more clues to problem cause.

*Management-by-exception alarm and alert tools* operate on the same premise as trouble ticketing to send an alarm/alert (a notification that a problem has occurred) to operations staff via e-mail, pager, or visual message to the management console. These alarms can be set to occur when symptoms of trouble are present, giving staff the opportunity to address a potential problem before it results in network downtime and lost productivity. Alarms/alerts also can be triggered by occurrence of particular SNMP events.

*Robust reporting capability* allows operations to create reports on a variety of network conditions and resources. Reports can be generated on demand or automatically in conjunction with a management-by-exception scheme. Reports that interpret data into meaningful graphical information are an easy and quick read for operations staff under pressure to solve or avert problems.

### 32.3.3 Security

Who gets access to what resources (hardware and software) and controlling that access is a pervasive and crucial issue for every aspect of networks: file servers, printers, applications, databases, systems, subnets, and nodes. Security issues include:

- Properly identifying appropriate levels of access for all components of the network.
- Determining and defining levels of access for each user according to his or her need to access the appropriate network resources to accomplish a given task.
- Determining the need for and location of *firewalls*, or barriers, to prevent access to company-sensitive computing resources by intruders. These firewalls can be established by routers or by special software running on a dedicated computer that allows only one-way traffic (outward from the protected network).
- Automated monitoring of firewalls as well as access to resources.
- Exception reporting and automated alerts for violations or suspected intrusions.
- Triggering capture of intrusion signatures to trace illegal access to the resource.


### 32.3.4 Fault management

The cost of downtime is rising steadily and is a prime indicator of the growing reliance on the network for conducting day-to-day business, retaining a competitive edge, and using information to its greatest advantage. The ultimate goal of fault management and resource availability is to identify and solve problems that affect a user's ability to access the network to perform necessary business transactions. Because of the steadily increasing complexity of the network, the level of sophistication required of IT to investigate and pinpoint network problems is likewise increasing. Faults can occur in devices, network links, databases, configurations, applications. The list is endless. In addition to a process that defines problem resolution procedures for network staff to follow, fault management must give IT several things.

*Centralized ability to interrogate* is required for devices at the network interface card (NIC) level, for network traffic at the packet level, for memory and storage capacity at the desktop computer level, and for configuration information at the file server level. The key is visibility.

*Ability to correlate threads of information* from devices, systems, applications, databases, and networks is needed on a timely basis. In this way, a holistic view of dynamic interactions as they occurred over time might reveal immediately the cause and effect(s) of difficult-to-detect network problems.

*Ability to magnify the view of a problem* allows examining the components that might be contributing to its cause. This is sometimes called *drill-down* or *zoom* capability. For example, there are network monitoring and analysis tools that will let you examine a problem by interrelating three to four network parameters that, when viewed progressively, will help magnify the problem source. This is a powerful capability that allows a network manager to chase a symptom to the root of the problem. Figure 32.8 is an example of this type of capability.



**Figure 32.8** Network management tools are available that report different pieces of information correlated through time, all relating to the same network issue. By looking at a problem from different perspectives, decision-making or problem-solving can progress quickly and efficiently. In this screen shot, an historical view at the top of the screen shows network utilization plotted beside the total number of errors for a slice of time. Focusing on the spike of utilization and error activity, the bottom left of the screen shows what protocols were responsible for that activity at that point in time. The righthand side of the screen shows the segment of the network that had the greatest amount of traffic (utilization) for the same point in time. This narrows problem investigation and allows other correlations to be performed, leading directly to the problem.



*Immediate notification of faults* is critical, through a system designed to recognize and prioritize problem events that have potential to lead to some form of network failure. The notification could be an on-screen color change, an e-mail message triggered to a distribution list of network staff, a message triggered to a pager, or a report generated.

*Growing intelligence built into the network* grants the ability to self-diagnose and present network staff with a well-defined path to resolution. Network symptoms for particular problems can be characterized and programmed into diagnostic equipment. A combination of symptoms might generate a probable list of problem causes for network staff to pursue. Network staff intervention and decision-making still are needed, however.

*Growing ability for the network to self-heal* provides such things as automatic redundant paths defined for network links that fail.

### 32.3.5 Performance management

Performance management consists of maintaining acceptable levels of response for business users of the network. It is a set of informed decisions based on data that is gathered from specialized collecting devices or embedded in network devices specifically for this purpose. This information is used to optimize the use of network resources and devices. Data that supports decision-making is provided by:

- Baselineing
- Application usage analysis
- Database monitoring
- Traffic monitoring

Baselineing typical network, application, or device behavior establishes a set of “normal” operating parameters. These baselines serve as points of comparison as a network grows and its use increases. They can be used to establish threshold levels indicating out-of-tolerance conditions that require action by network staff. These threshold levels are used in management-by-exception schemes. (This is a term that refers to a proactive methodology for network monitoring.) Network management software watches the network and informs network staff of potential degradation. With the global spreading of many enterprise networks, it is virtually impossible for an operator’s eyes to be everywhere and give equal weight to all aspects of the network. Performance management software makes this possible. It directs attention to network areas that need it most.

Analyzing application usage helps determine how cycles of business activity affect use of the network, in turn helping identify mission-critical applications and impacting decisions about how to segment user communities, what computing and network resources are needed to support various applications, and how to balance load across the entire network so that valuable resources are used to their full potential.

**754 Network Management**

Monitoring database usage is done so that capacity needed does not exceed capacity available. This helps network staff analyze database access to determine peak periods of use and potential bottlenecks that may affect user response.

Monitoring internetwork traffic and trends helps determine if rebalancing or resegmenting will make better use of existing resources.

**32.3.6 Planning**

Planning means constantly archiving and monitoring network information across devices, applications, databases, systems, and network architecture to help make informed decisions about the future of the network, such as:

- Anticipating network needs and modeling the behavior of planned changes without committing resources of time and money.
- Justifying the expense of moving to emerging technologies and standards via cost justifications based on projected costs and savings.
- Justifying the appropriation of IT funds to maintain and improve delivery of network service to end users by using performance management, availability, and response information gathered in other categories of network management.

**32.4 Network Management Methods**

Network management is an integrated set of processes and tasks that may span thousands of miles, different support organizations, and many devices. This typically is referred to as *enterprise management*. It also includes processes designed to handle small groups of users who communicate locally, within the same building or room; this is referred to as *workgroup management*. Most large enterprise networks include the need to manage workgroups as well as bridged network segments that can be globally dispersed.

Within the network, a great variety and number of applications and databases warehouse and process data into meaningful information. The network delivers the information to end users, enabling decision-making, problem-solving, and, in general, increasing the efficiency with which work is accomplished. What is the most efficient way to gather the necessary information about network components to keep the network running smoothly? The *distributed network management model* has emerged as the most efficient tool for accomplishing this enormous task.

**32.4.1 Distributed network management**

Without distributed network management, it would be impractical to monitor tens of thousands of nodes from a single centralized workstation known as a network management station. The load of attendant networking information converging on and emanating from the management station would be impossible to handle.

### 32.4.2 Models for network management

The value of distributed network management can be better understood by comparing it with two other models for network management, *central* and *hierarchical*.

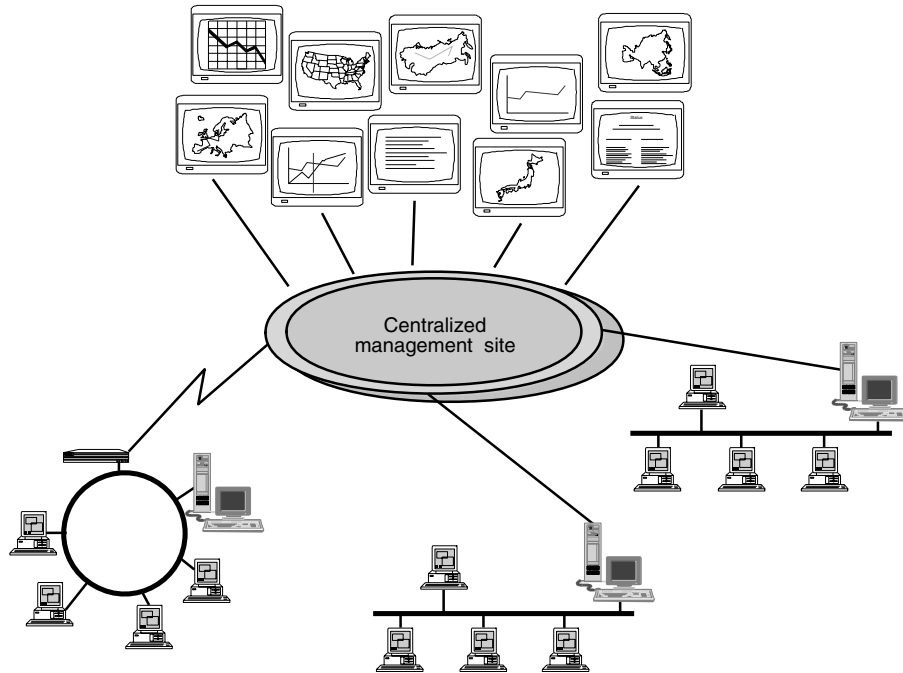
**Central management.** A network typically consists of multiple large sites across the country or around the world. At each site, the number of segments can vary widely. Each segment has a server that centralizes processing for applications, databases, and network “housekeeping” for all the devices attached to that segment. Each site also includes the connecting devices (hubs, routers, bridges, etc.) that link together the various network segments into a cohesive lattice where information smoothly flows into the site from other sites, is delivered to the appropriate device, is processed, and then, as necessary, streams back out to other interconnected sites.

With central management, all of the devices in this web of segments are managed from a central location. Each of several operators at the central site specializes in managing some subset of the network. The subsets can be based on region, device type, management focus (like user administration or security), or some other categorization scheme. Each specialist focuses on his or her view or perspective of the network. Each might have a separate console where disparate, uncorrelated network information can be monitored. When problems occur anywhere on the network, network staff is dispatched by the central site to the problem site to investigate the problem or implement a fix. In this central management model, potentially large amounts of data pass over WAN links between distant sites and the central management hub. Figure 32.9 is a simplified view of a centralized management model.

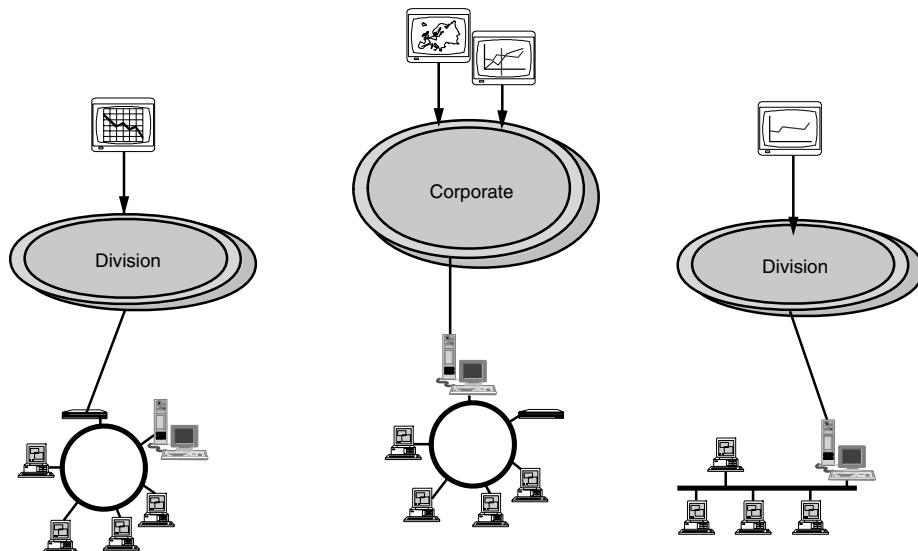
**Hierarchical management.** As with centralized management, this model addresses the needs of a network that consists of multiple large sites, referred to as an *internetwork*. In a hierarchical model, however, there are two or more distinct levels where management is performed from a number (or from all) of these sites within the internetwork. Usually, network staffers at each local site within the internetwork are responsible for the devices on their site, and a group at a central site may oversee and back up the regional efforts and manage the corporate backbone.

The concern is coordination between different, yet important, management efforts spread across the company. Duplication of effort can occur, but it is a small price to pay for a more reliable environment. The hierarchical model is flexible and can accommodate easily the addition of sites at the extremities of the network, or the pooling of sites into a cluster to be managed at an intermediate level (like a region). Figure 32.10 is a simplified view of a hierarchical management model.

**Distributed management.** Distributed management offloads resource-intensive processes (memory, disk space, CPU power) from the traditional network management station or console and places it between the devices that need to be managed and the management stations. These *collection stations* can be structured like the hierarchical model to monitor, collect, and store information on a specific group of network devices. The groups of devices are called *management domains*.



**Figure 32.9** Centralized network management locates all the management resources for a network at a central site, often in one big room. Each distinct function of network management is controlled separately from the others. Each network segment reports directly to the central management location.



**Figure 32.10** Hierarchical network management typically consists of two or more distinct management levels. This could be a corporate site and a number of divisional or branch sites. Each site has its own network management structure and responsibilities for segments within its domain. The corporate site serves as a backup for the divisional or branch sites, and also might maintain responsibility for its own set of segments.

Devices grouped into management domains can be segmented by department function (finance, marketing), geographic location (city, state, or region), or any other scheme of categorization. The collection station consolidates information on the network devices, applications, databases, and communication within its management domain. The collection station then forwards relevant information to one or more management stations upon request or as programmed. Figure 32.11 is a simplified view of a distributed management scheme.

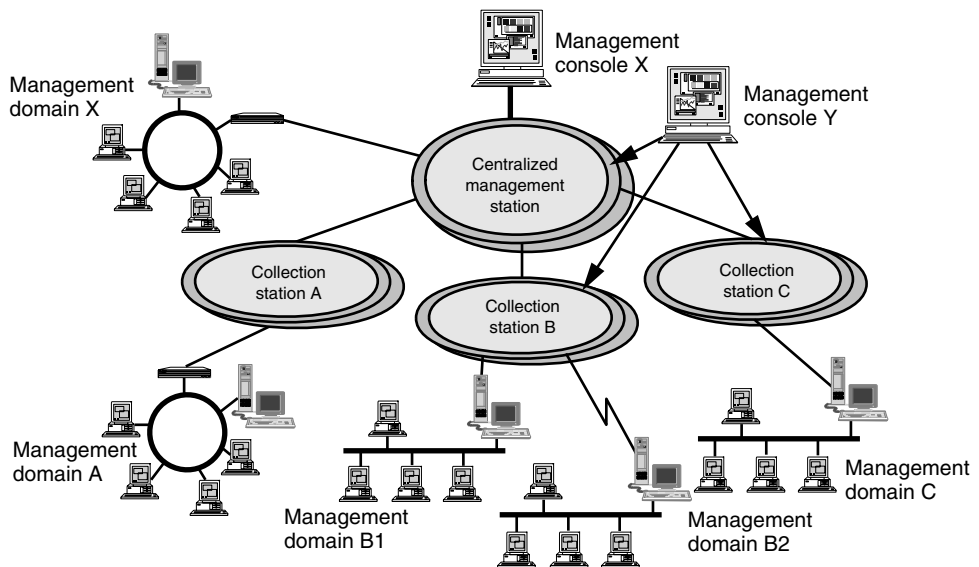
The advantages of distributed management are many, and include the following.

*Reduction in management traffic* frees network resources, like capacity of the communication medium and processing power at the network management console, for business use.

*Reduction in the need to employ WAN links* saves money in moving data to a central management station. The network itself is used to communicate network management information.

*Scalability* means network management can adapt to meet a growing network's needs, partitioning new segments into existing management domains or into their own domains. This scalability allows a network management scheme to accommodate growth with great efficiency. Processing power to handle the added devices or segments resides in the intermediate collection station, sometimes referred to as a mid-level manager. The addition of a segment with 30 nodes, for example, results in the central management station adding just one new domain manager to its monitoring load.

*Freeing-up of processing power and memory* at the management station allows more efficient use of applications that perform network management functions. The



**Figure 32.11** Distributed network management locates processing power at intermediate levels throughout the network. These intermediate managers perform resource-intensive tasks that normally would consume the central management station. The intermediate managers then report only events of interest to the central site, and provide information upon direct query about nodes within their domains.

management station, in addition to keeping track of the topology of a network, also must be the centralized location for application software that performs specifically for any or all of the areas of network management previously identified: configuration and change, user support, security, fault management, performance management, and planning.

The intermediate collection stations offload the processing required in many of these areas. Thresholds can be set at the collection station so that only anomalous events that occur for prescribed network conditions are passed to the central management console. Discovery and mapping of devices within each domain also can be handled by the collection station. With this methodology, an enormous burden is removed from the central management station and is broken down into more manageable, segmented pieces on a domain-by-domain basis.

*Facilitated integration of applications* across network management functions allows the management station to serve as an aggregation point for data from many collection stations.

*A single, holistic view of the network* reveals network interactions between devices, applications, databases, systems, and network communication. Unlike central or hierarchical network management schemes, distributed management grants visibility of an enterprise network at a centralized management console. Views of the network can be shared with other management stations, and control for various aspects of network operations can be divided.

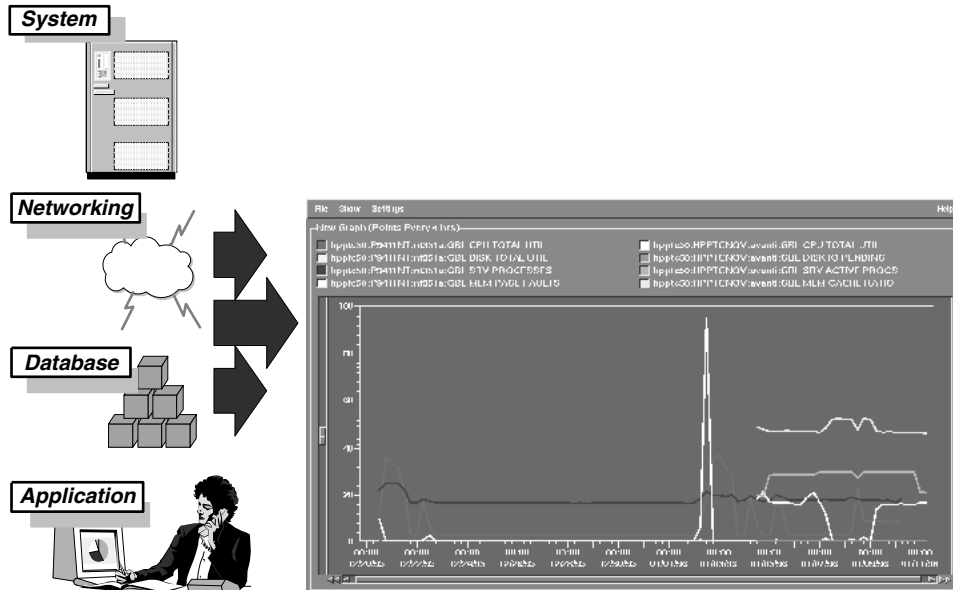
### 32.4.3 Integrated network management

It should be made clear that simply seeding a network with distributed devices (like network test devices or hubs and routers that support SNMP) does not constitute integrated network management. This provides at best a fragmented view of different network components that can be applied in several different network management areas. Correlation of data in this type of scenario is difficult and time-consuming and contributes to inaccurate interpretation of network interactions.

True integrated network management lets interactions within segments and domains, as well as between domains, be observed through a “single pane of glass.” It provides graphical representations of all aspects of the internetwork on a single display. Data correlation is presented side-by-side with consistent formatting, contributing to accurate and timely decision-making. In this way, integrated network management allows problem correlation across a wide range of statistics, alarms, maps, fault location features, graphs, tables, charts, and decodes on a single display. Figure 32.12 is an example of a single display from an integrated network management system.

### 32.4.4 Network management applications

The job of managing networks is no longer just a matter of managing the devices or computers that make up those networks. It involves managing other resources critical to business, such as servers and end-user applications. It is complicated further by the inherent multivendor, heterogeneous nature of networks. A robust network management system addresses all these resources with a choice of options that ex-



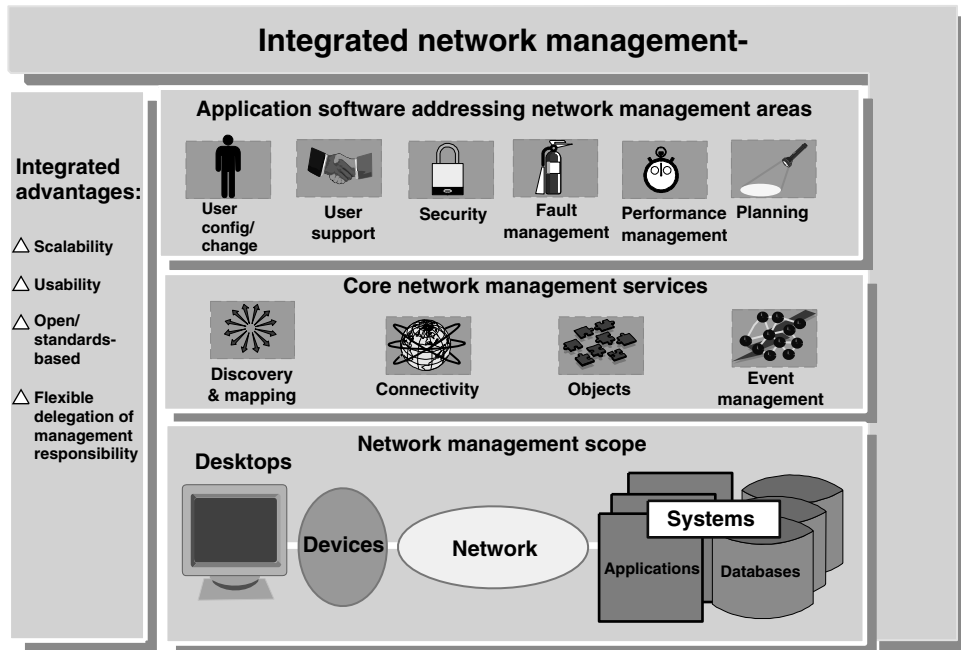
- Manage the performance of resources from a single pane of glass.
- Pinpoint problems in system, network, database, or application resources before they affect users.

**Figure 32.12** Integrated network management enables correlation of network statistics from all network areas. Data are presented side-by-side so that activity from systems, databases, devices, applications, and network traffic can be seen simultaneously. On this screen, total CPU utilization, total disk usage, and total server activity are displayed for an end-to-end network path that may be made up of several different network segments. Also displayed are the same statistics for one segment of that path. This allows network management to monitor the effect on these statistics that one segment may have.

tend the operator's abilities to meet specific network management needs based on unique business needs.

The advantage of integrated network management is that core software provides a foundation for applications addressing all network management areas to share and leverage data where possible, introducing information efficiencies in numerous ways. It serves as a foundation for integrating all the various and necessary network management functions. Using SNMP and standard or proprietary MIBs, it provides a base of access to all the network components through the network communication link. This includes workstations, servers, applications, databases, networking devices (hubs, routers, bridges), and network traffic.

The network management application provides a common set of services into these components, such as discovery, mapping, and event handling (events as defined by SNMP). Network management applications can access these services to take network component data and apply it in useful and creative ways to help staff cope with the complex task of managing network resources. For example, discovery and mapping allow network addresses to be retained so that any device problem can be quickly located and dealt with. Figure 32.13 shows an integrated management scheme.



**Figure 32.13** Integrated network management allows information from all network components to be accessed, consolidated, and correlated at a central management station. This enables all network management functions to examine the information and use it for decision-making and problem-solving as needed. Integrated network management has many advantages over other schemes because it employs scalability to conserve network resources, provides a common user interface that promotes ease of use, provides investment protection because it is open and standards-based for ease of migration to new technologies, and allows network management to be flexibly delegated at intermediate management levels.

MIBs common to all SNMP agents can be used in creative ways to facilitate all aspects of network management. This is the great advantage of standards-based network management. Some agents even allow custom MIBs to be defined so that unique applications particular to the needs of a specific network can be addressed. This broadens the capabilities of SNMP-based management applications to control basic network devices and critical systems and applications.

The MIBs collect the data from network devices and traffic; applications then gather and process the data into meaningful information for presentation at the network management console. In addition to managing devices like routers, bridges, and hubs, they are thus equipped to manage objects such as applications, printers, users, and databases that are central to business success. The ability to control information access to network and system resources, and effortlessly monitor important network components, provides users with unprecedented visibility and control of a network infrastructure.



## SS7 Signaling Monitoring Systems

**George Cooper**

*Hewlett-Packard Ltd., South Queensferry, Scotland*

### 33.1 Introduction

As the world rushes into the information age, telecom customers demand new, faster, and more reliable services. This places increasing demands on digital networks. Expansion of these networks increases the deployment and reliance upon the common-channel signaling system no. 7 (CCS/SS7). SS7 enables faster, more efficient and more reliable routing of communications and creates the opportunity for new revenue-generating services.

With the increased rate of growth, the complexity and diversity of SS7 networks increases. This, along with the shortage of skilled SS7 engineers, makes it increasingly difficult to monitor and manage SS7 networks. An SS7 signaling monitoring system provides the tools to support these SS7 networks. Some important features of a typical system are:

- Network-wide SS7 monitoring and real-time, centralized reporting for early warning of revenue-threatening conditions.
- Graphical display of critical information, for a clear view of key network information.
- Nonintrusive connection to the SS7 network that does not impact SS7 network reliability.
- Nondependence on network elements such as switches, so it can provide network-wide, independent results.
- Alarm management, protocol analysis, and call trace provide the ability to troubleshoot network problems.
- Call detail record collection to facilitate usage billing, fraud detection/control, etc.

### 33.2 Introduction to Out-of-Band Signaling Systems

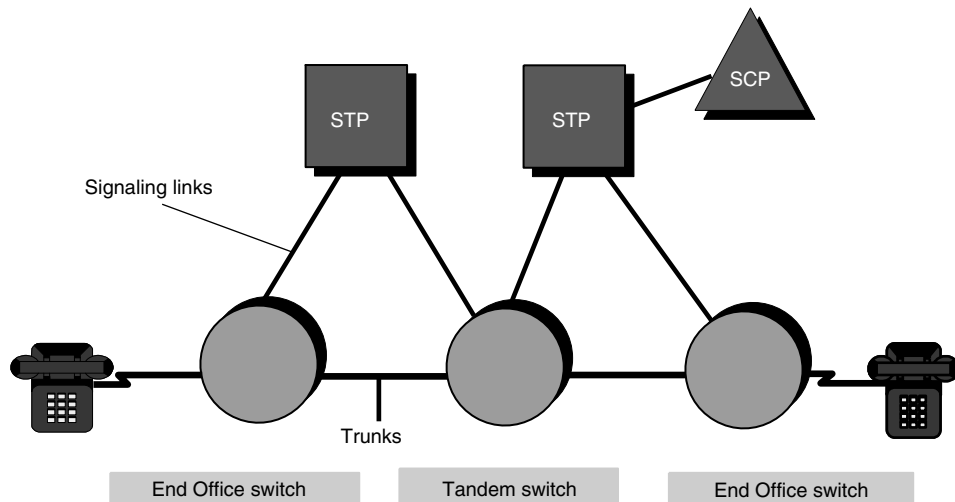
A telephone system is composed of two network types: the *speech path* (trunk) network that carries the voice channels between exchanges and ultimately between the customer's telephones, and the *signaling network* that performs the interexchange control required for call processing. The signaling network carries all the data messages used in setting up calls between computer-controlled telephone switches. These messages conform to a layered network protocol known as *Signaling System No. 7* (or SS7).

When a call is made between two telephones, several steps are required to perform the dialing, ringing (or busy), answer, and hang-up operations. Signaling messages communicate these steps among the various switches involved in processing a call. These signaling messages instruct the switch to generate the appropriate tone (ringing tone, for example), or to divert the call to a prerecorded message such as an announcement that a number no longer is in service.

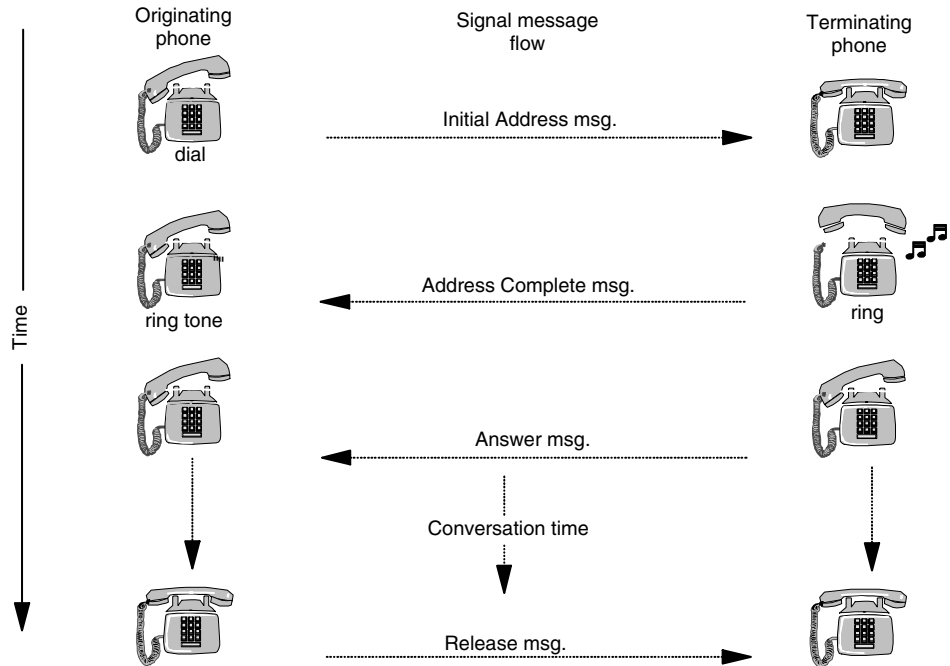
Figure 33.1 shows a simplified view of telephone network infrastructure.

This view of the telephone network shows two telephones connected to End Office switches (local telephone exchanges) and a trunk path between these switches passing through a tandem switch. Above the trunk and switch level there are *signaling links* and *signaling transfer points* or STPs. The STPs act as packet routers, passing signaling messages towards their destination.

The triangular symbol of shown in Figure 33.1 represents a *Service Control Point* or SCP. A common example of a service is an 800 free phone number database, which could be held within an SCP. A switch, receiving an 800 call, sends a



**Figure 33.1** A simplified view of a telephone network infrastructure. Even a simple telephone call requires several steps and end-to-end network access to set up and later release the connection based on information in network databases. Special billing options such as credit card calls and third-party billing require additional database look-ups. All this is done via the SS7 network, a packet-switched network comprising signaling links, switching transfer points (STPs), and switching control points (SCPs). Switching control points are databases filled with network routing and billing information. Signaling transfer points are packet switches.



**Figure 33.2** A simple telephone call requires a sequence of actions to set up the call initially and to release the call when either party hangs up. Upon completion of dialing, a message is sent to the SS7 network, which prepares the end-to-end connection and rings the called telephone. This is confirmed to the caller with an “address complete” message, heard by the caller as ringing tone. Alternatively, a busy signal may be returned. As the called party picks up the phone, an “answer message” instructs the SS7 system to complete the end-to-end connection. At the close of the call, a “release message” instructs the SS7 system to release the intermediate trunks and switches for other calls.

query over the signaling network to the SCP to determine the actual number to which to route the call. The SCP returns the actual number to the switch over the signaling network and the switch sets up the call to that number.

The network carrying the signaling traffic comprises individual signaling links, each link typically running at 56 or 64 kbps. These links pass data to the STPs for routing to destinations based on a message’s *point code address*. When more signaling traffic flows between locations, additional links are grouped into *linksets* to increase the signaling bandwidth. A large telephone company’s SS7 network might contain up to fifty or more STPs and many thousands of SS7 links.

What are the signaling messages? The signaling network carries several forms of signaling messages that serve different needs. Signaling messages are addressed by numeric point code. Each network element, such as a switch or service control point, is allocated a worldwide unique point code; in this regard it is similar in concept to a conventional Ethernet address. One advantage of this addressing scheme is that messages can be identified by their origins or destinations. Among other things, this allows a specific carrier’s network traffic to be identified readily from a mixed stream of messages.

Figure 33.2 shows a simplified view of the SS7 signaling messages exchanged during the setup and hang-up phases of one telephone call. Although the diagram shows

one order of events (an *event* is one or more messages on the signaling network), these could occur in various other sequences, such as if a caller were to hang up before the call was answered.

Note that on any one SS7 link the messages relating to many hundreds or thousands of simultaneous telephone calls are all interleaved. An added complication is that messages relating to any one phone call could transit the SS7 network using different links, possibly diversely routed (that is, via different intermediate STPs).

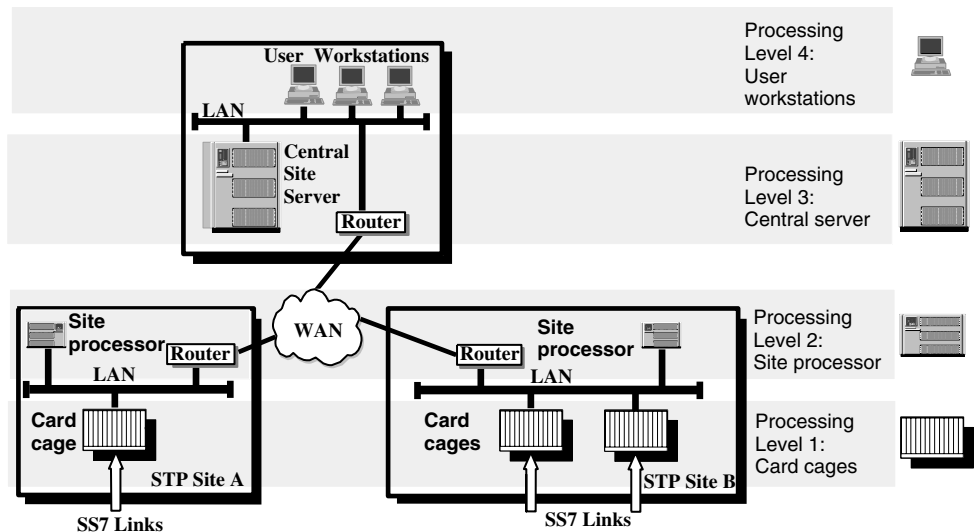
### 33.3 Description of an SS7 Signaling Monitoring System

An SS7 monitoring system is a distributed monitoring system with potentially hundreds of special-purpose data probes connected to signaling links throughout the SS7 network. Data from these probes are gathered, correlated, filtered, and transmitted to a central site, where various applications use the data for network and service management.

#### 33.3.1 System overview

Figure 33.3 shows a basic diagram of an SS7 monitoring system. It consists of four layers of processing. Layers 1 and 2 typically are located at STP system sites that provide a convenient focal point of signaling link sets.

Layer 1 consists of card cages containing SS7 interface cards and processor cards. The interface cards collect the signaling data from the SS7 network via



**Figure 33.3** An SS7 signaling monitoring system is a distributed computing and measurement system with four levels. Level 1 is distributed throughout the geographical region being monitored and includes card cages filled with measurement probes. Level 2 controls the probes and collects and concentrates data from the probes, routing it to the central site server (level 3), where application programs analyze the data for various management, alarm, usage, billing, or other purposes. At level 4, workstations for operator interface connect to the central site server via the LAN.

electrically isolated connections, i.e., without affecting the signaling traffic. A copy of the signaling traffic is fed to the processor cards in the card cage. The processor cards, containing powerful RISC processors, are responsible for time-stamping received SS7 messages, data filtering, and other front-end processing required by the SS7 applications. Each pair of interface and processor cards comprises a protocol analyzer dedicated to the SS7 link to which it is permanently connected.

All card cages at a site are connected to a site processor via a LAN. The site processor is a Unix server. This server is responsible for additional data processing as required by the various SS7 applications.

All sites are connected via a wide area network (WAN) to a central site, where a central site server and user workstations are situated. The central site server is responsible for configuration and management of the system and its applications.

Key aspects of an SS7 monitoring system are:

- SS7 data is available in real time.
- Data reduction should be done by distributed processing at the lowest possible layer. Because of the huge volumes of data in a large SS7 network, transmitting all of it to the central server is not feasible.
- The system should be open, providing published and supported Application Program Interfaces (APIs) for building SS7 monitoring applications.
- The system should support co-residency of multiple applications all running in parallel.

A typical SS7 signaling monitoring system is modular, which facilitates distribution across a geographic region and makes for easy upgrades. The system should be scalable to suit customer network requirements. Typical systems are nonintrusive and operate by collecting and analyzing messages from the SS7 links in a network. It should be independent of network elements and network element vendors. Thus the system can provide a comprehensive, impartial view of the network even during fault conditions. Such a system should be customizable to end-user requirements and designed to integrate into a customer's OSS.

### 33.3.2 Processor hardware

An SS7 signaling monitoring system uses a distributed computing system. Three processor configurations are required in addition to processors embedded in the probes.

**Central site processor.** The central site processor communicates with all the remote site processors over a wide area network (WAN) and provides central control of the monitoring system.

**Remote site processor.** The remote site processor can control a number of card cages. Each remote site processor is responsible for collecting data from the monitored links. Each remote site processor uses a WAN link to the central site

processor and to a peer remote site processor to provide cross-site triggering for call trace and protocol analysis.

**User workstation.** User workstations are clients of the central site server. These provide graphical user interfaces (GUI) for running tools or applications. They can be connected directly to the server LAN or might be located remotely. Computers from various vendors may access applications using the remote X Window protocol.

Additional processing functionality is provided by interface processor cards contained in the measurement card cages. These cards monitor the SS7 links and perform low-level SS7 measurements.

### 33.3.3 Central site hardware configurations

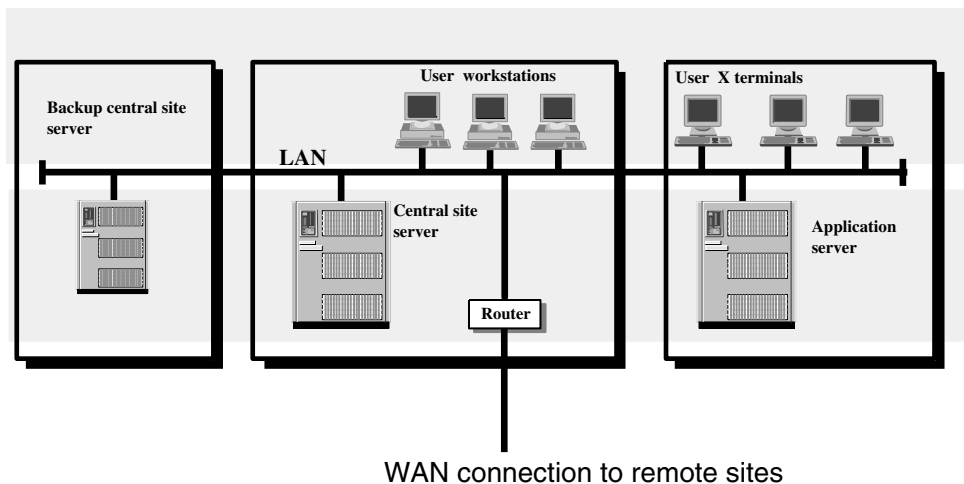
Central site processor configurations depend on the processor model, internal or stand-alone peripherals (such as disks, tape drives, printers), disk configurations (such as size and whether mirrored), and the level of availability/redundancy required. Figure 33.4 illustrates a typical central site server configuration. For a large system, additional processors may be added for backup and to run applications.

### 33.3.4 Remote site hardware configurations

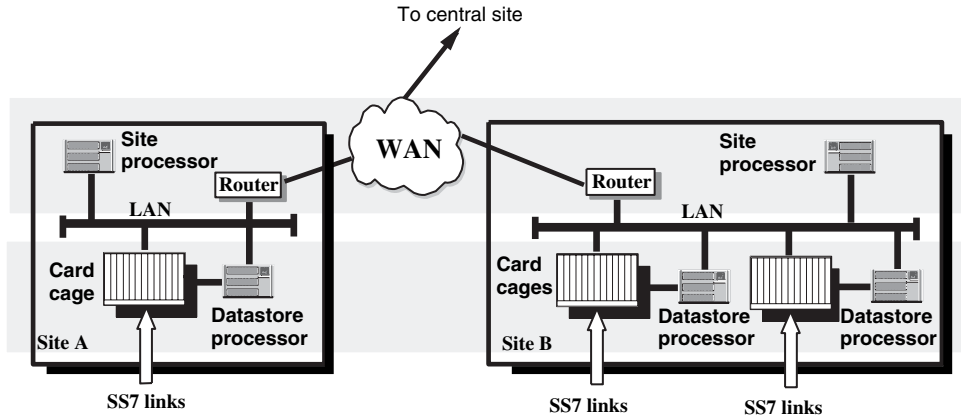
An example remote site processor hardware configuration is shown in Figure 33.5. Remote site processor configurations will depend on the number and type of links to be monitored at the site, and the type of data required by the software applications to be run on the system.

### 33.3.5 Measurement card cage hardware

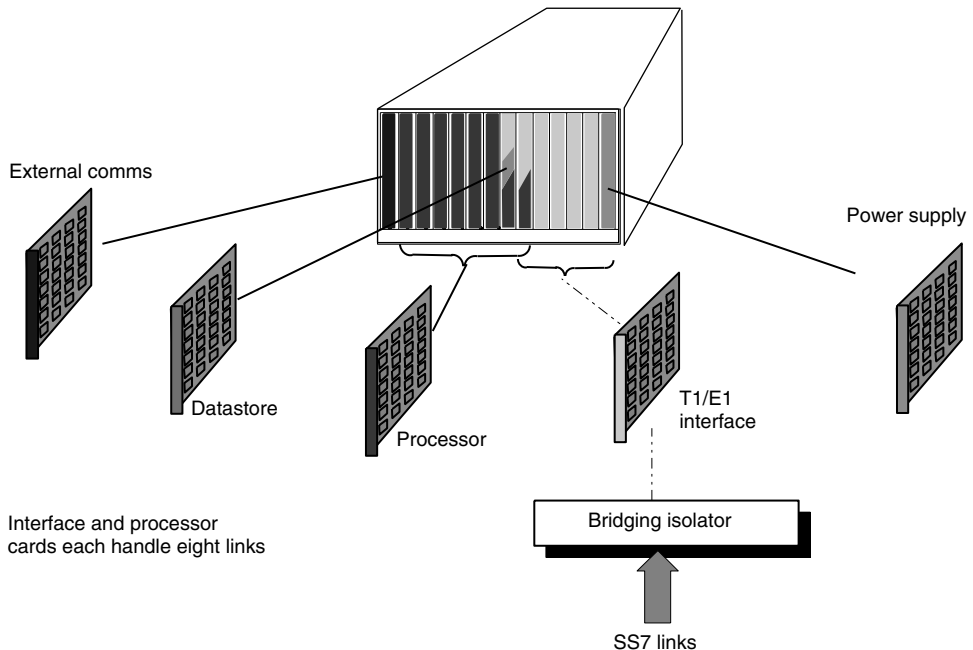
The measurement card cage (Figure 33.6) provides interfacing for the SS7 links, which typically are transported over 1.544 Mbps, 2.048 Mbps, and H64/56 kbps bearers (T1, E1, and DS0A respectively). The interface processor cards are used to



**Figure 33.4** For large installations, the central site hardware may be expanded to include a backup server and a dedicated application server.



**Figure 33.5** The remote site hardware is distributed throughout the geographical region to be monitored. Within a site the components are interconnected through a LAN. Sites are linked together by routers and the wide area network (WAN). The card cages housing the actual monitoring hardware can be scaled from monitoring a few links up to hundreds at a geographical site. The levels of processing power can be matched to the number of applications running by choosing from a range of standard Unix servers. In this configuration, additional processors provide additional local data processing and storage.



**Figure 33.6** The measurement card cage structure allows a number of interface cards and interface processor cards, along with the data storage card, to reside in a North American Building Standard (NEBS)-compliant structure. Any signaling link can be switched to any interface processor card, allowing critical applications to run on dedicated processors.

**768 Network Management**

process data coming from up to eight links per card. A datastore card can be added (optionally) to capture and store all signaling messages at a sustained 100 percent load (1 Erlang) on all links. The data is fed to a datastore processor for storage (Figure 33.5). Additional accessories include bridging isolators and concentrator units for interfacing to V.35 or RS449 links. Together these products form the measurement hardware of SS7 signaling monitoring system.

**33.3.6 Routers**

Any router may be used that supports the appropriate LAN interface and has at least two WAN interfaces that suit the customer's network.

**33.3.7 Communications**

The SS7 signaling monitoring system utilizes an IP (Internet Protocol) internetwork independent from the signal link being monitored (refer to Figure 33.3). Each remote site processor connects to the central site. Connection to the central site is via routers and WAN communication links, which may be T1/E1, or Nx56/64 kbps. Each remote site processor monitoring a Signaling Transfer Point (STP) will have a logical connection to the remote site processor monitoring its paired STP. Remote workstations are supported.

**33.4 Central Server Software and Applications**

The central server is where the data gathered throughout the SS7 network is analyzed for network management applications, billing, fraud control, etc.

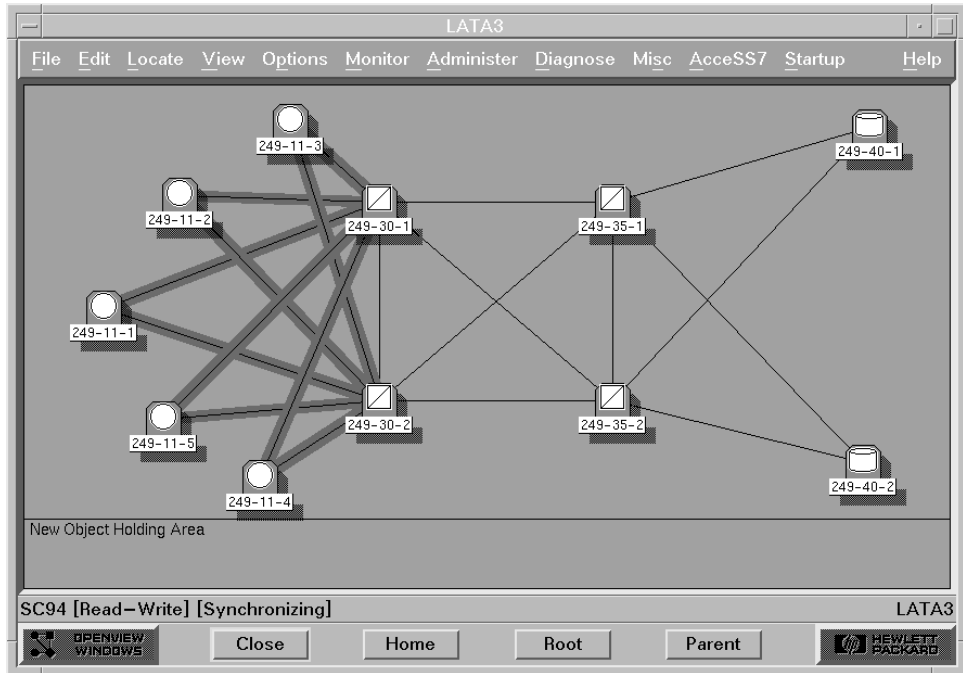
**33.4.1 SS7 Network management application**

Table 33.1 lists a number of network management applications which can be run on data derived from SS7 links.

**TABLE 33.1 SS7 Network Management Application.**

Applications	Features
Link status monitor	Network-wide link status monitoring
Traffic monitor	Network-wide performance monitoring
Protocol analysis	Protocol analysis for selected End Office links
Call trace	Tracing multiple ISUP, 800, and AIN calls
Alarm manager	Reporting and logging alarms
Data capture	Real-time network-wide recording of SS7 data
Fraud management	Real-time detection and reporting of suspicious telephone usage
Statistics	Monitor performance of network resources and measure impact of network changes
Billing usage data	Signaling network usage and voice trunk usage from SS7





**Figure 33.7** In the link status application, the system provides a real-time map showing the interconnection of all the SS7 linksets being monitored. The status of linksets is shown on the GUI network map by color-coding the displayed linksets.

**Link status monitor application.** A real-time graphical user interface (GUI) provides a network-wide map displaying geographical regions being monitored and indicating status of regions via color changes of relevant symbols (Figure 33.7). From the network-wide map, regional maps can be selected. These show the status of all links and switching points in that region via color changes of their relative symbols. A link status display window, which provides link status and statistical information of a selected link, can be requested from these GUI maps. The link status monitor shows the status of linksets on the GUI network map as follows:

- *Red:* All links in the linkset have status oscillating or unavailable.
- *Yellow:* One or more, but not all, links in the linkset have status oscillating or unavailable.
- *Green:* All links in the linkset are available.
- *Blue:* All links in the linkset have a status of unknown or deleted.

The status of each channel on each link is detected from the data carried by the channel. The following six-channel states are defined for each channel:

- Out of Sync
- Out of Service

## 770 Network Management

- Processor Outage
- Busy
- In Service
- Idle

The status of each signaling point is indicated by the status of the links connected to it. The following four signaling point states are defined for each signaling point.

- Isolated from network, all links unavailable/oscillating.
- Isolated from network, all B/C/D links unavailable/oscillating but at least one A link is available (only for signaling points with STP functionality).
- Connected.
- Unknown.

**Link statistics.** The system maintains the following statistics for each link:

- Number of MSU octets over a time interval
- Number of occurrences of each type of LSSU over a time interval
- Number of FIB/BIB inversions over a time interval
- Number of frames in error over a time interval
- Duration of unavailability over a time interval
- Duration of oscillating status over a time interval
- Number of link status changes over a time period when the link is oscillating
- Count of SLTM MSUs over a time interval
- Count of SLTA MSUs over a time interval
- SLC value in last SLTM message
- Current channel state for each channel
- Number of MSUs excluding SLTs over a time period

The link status monitor provides the facility to select linksets and view their status in more detail. For each link in the selected linksets, the system displays the link's current status plus all of the link statistics defined above.

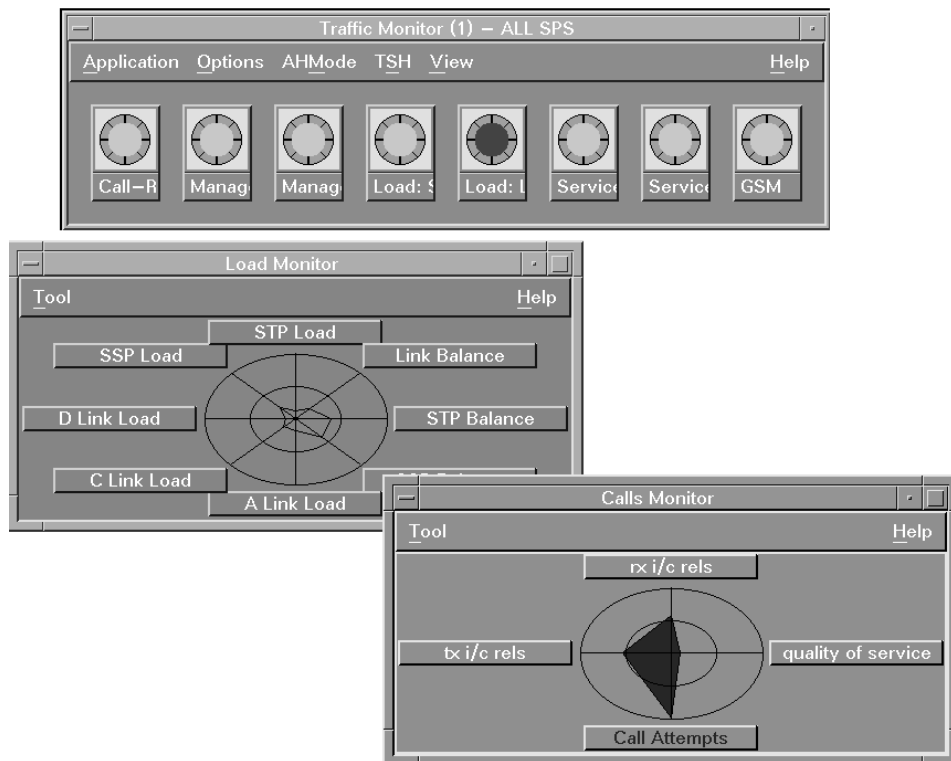
**Traffic monitor application.** The traffic monitor application measures the real-time performance of network resources and the impact of network changes by continually monitoring traffic between SS7 network elements. Up to eight real-time displays show the state of the network.

To ensure that network problems are quickly recognized and resolved, traffic monitor displays measurement information in three levels. At the first level, the network-wide data is correlated and value-sorted, and the worst-case value for each measurement is displayed in a radar diagram. At the second level, up to ten worst-case elements are displayed in an active histogram and updated dynamically when

there are any changes. At the third level, a detailed view of the operation of a specific network element for a particular measurement over a 60-minute period is displayed in a time series histogram. Additionally, these displays are entry points for further investigation of any faults indicated.

Each real-time radar diagram display consists of a number of axes on which are plotted various network parameters. The actual quantity shown on an axis is the current worst case of the parameter represented. Every 15 seconds, information from all SS7 links on the network is analyzed and classified for presentation on the axes. The points on all the axes are joined to form a polygon which is filled with color. The polygons have an inner and an outer circle, which represent the configurable threshold settings for warnings and alarms respectively. If all points are within the inner circle, the polygon is green; if any point exceeds the inner circle, the polygon is yellow; and if any point exceeds the outer circle, it is red. The thresholds on each axis are configurable by the user.

Consider an example of how Traffic Monitor can assist in network management during an unexpected event such as mass calling to a local TV station (Figure 33.8a).




**Figure 33.8a** The traffic monitor application shows SS7 traffic conditions using graphical displays called *radar displays*. The concentric circles represent configurable thresholds and the colored polygons show the labeled traffic parameter value on the appropriate radial line. As the value of a particular parameter moves past its thresholds, the polygon changes color from green to yellow to red. Pointing and clicking on a labeled parameter brings up more detail on that parameter. Shown here are the results of clicking on “Load: Site” (currently green) and “Call-Related” (currently red).

As long as the network is working at normal load, all the icons are green. But suddenly the icon for call-related measurements turns red. This indicates to the network operator that a problem is developing in the network. To get more information, the operator clicks on the red-colored icon. This brings up the radar diagram for call-related measurements.

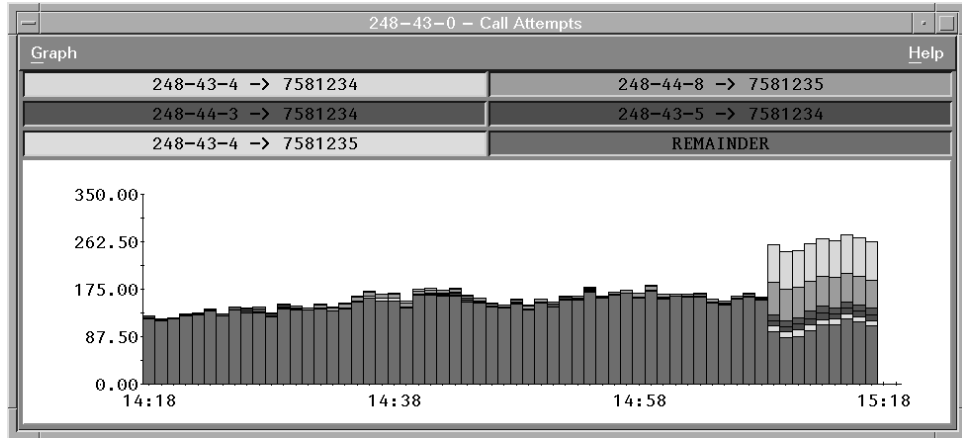
At this top level the problem can be identified as being a high number of call attempts somewhere in the network. This is indicated by the red polygon that exceeds the critical threshold for the measurement Call Attempts. More information is required quickly to resolve the problem. By clicking on the Call Attempts label, which is red as well, the next level of available information, the active histogram (Figure 33.8b), is shown.

The active histogram shows the ten network point codes with the currently highest number of call attempt rates. It is clear that a single location (point code 248-43-0) has received an abnormally high number of calls and is overloaded. This extraordinarily high number of call attempts demands immediate attention to prevent a costly network outage. More information is available by clicking on the appropriate point code button.

The time series histogram (Figure 33.8c) displays the activity at the selected point code during the previous 60 minutes. The extreme right-hand bar on the horizontal axis represents the activity at current time. It clearly illustrates the top five sources (originating point code/phone number pairs) that caused the overload. The network operator can see that for most of the hour the call attempt rate



**Figure 33.8b** An active histogram displays results for up to ten worst-case network elements (point codes) for a particular measurement, in this case call attempts.



**Figure 33.8c** The time series histogram gives a detailed view of the operation of a specific network element: signaling point, signaling link, or signaling linkset, indicated by the point code for a particular measurement. The last 60-minute data period is displayed. This particular display shows the overload resulting from a television station call-in competition, a “mass calling” event.

was normal, but during the last few minutes a pattern of calls to two particular numbers has emerged. Knowing that the called numbers are the numbers of the local television station, the operator deduces that the mass calling event is caused by a TV call-in competition.

The operator now has enough information to advise the maintenance team to take action, such as by introducing call-gapping at the appropriate points in the network. The traffic monitor displays continue to be updated, so that the effect of the corrective action can be verified.

Table 33.2 shows typical groupings of some of the available measurements.

**TABLE 33.2 Typical Measurements in Traffic Monitor Application.**

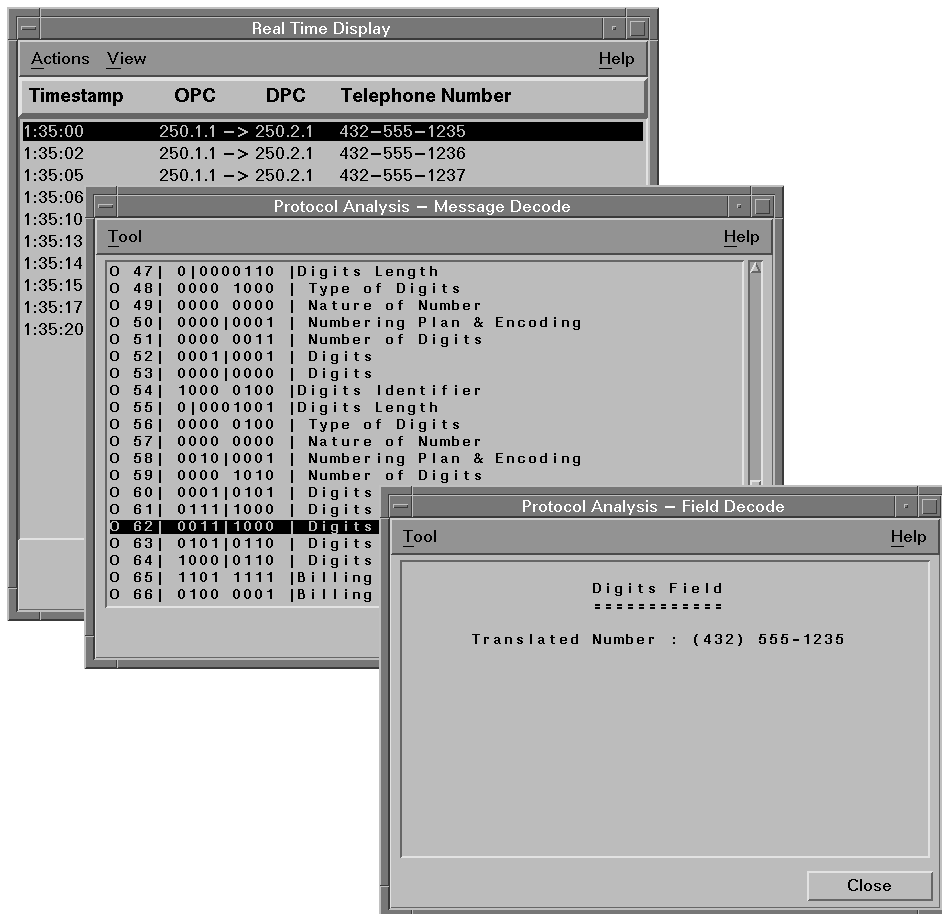
Call measurements		
Axis	Definition	Application
Quality of service	Point codes with highest ratio of abnormal RELs to total IAMs sent.	Shows point codes receiving the worst service from the network.
Call attempts	Point codes receiving the highest number of IAMs.	Shows point codes closest to overload, which might be caused by mass calling event.
Transmitted abnormal RELs	Point codes sending the highest number of abnormal RELs.	Sign of mass calling event or SSP problem.
Received abnormal RELs	Point codes receiving the highest number of abnormal RELs.	Timing problem between SSPs
800 queries	Linksets carrying the highest number of TCAP queries from SSPs.	Shows success of 800 service and indicates where more capacity is required.

TABLE 33.2 Typical Measurements in Traffic Monitor Application (Continued).

<b>Management and maintenance measurements</b>		
Axis	Definition	Application
LSSUs	Links with highest LSSU counts.	Identifies problem links.
Routing	Point codes receiving highest number of route-set-test messages. Clicking on the axis will show the top sources.	Shows point codes to which the network is having problems routing messages.
SP congestion	Point codes receiving highest number of RCT messages. Clicking on the axis will show the top sources.	Shows most congested point codes in the network.
Link changeover	Links most frequently referred to in link changeover messages.	Shows links with high traffic levels.
SLT	Links with greatest SLTM /SLTA traffic.	Indicates link problems.
A-SNM	“A” links with highest ratio of SNM octets to total octets.	Indicates network-level faults.
B-SNM	“B” links with highest ratio of SNM octet to total octets.	Indicates network-level faults.
C-SNM	“C” links with highest ratio of SNM octets to total octets.	Indicates network-level faults.
D-SNM	“D” links with highest ratio of SNM octets to total octets.	Indicates network-level faults.
SIB	Links with highest LSSU counts with busy status indicator.	Shows link congestion.
SIPO	Links with highest number of counts with processor outage status indicator.	Shows signaling point problems.
<b>Load measurements</b>		
Axis	Definition	Application
Quality of service	SSPs with highest ratio of unsuccessful query messages.	Shows SSPs receiving the worst service.
Load	SCPs producing highest number of MSUs.	Shows SCPs closest to overload.
SCMG	SCPs with highest number of subsystem management messages.	Shows SCPs providing the worst service availability.
UDTS	SSPs with the highest number of UDTS & XUDTS messages.	Routing problem between SSPs and SCP providing that service.
Link balance	Linksets with the greatest MSU imbalance between links.	Show routing problems or link faults.
A-link load	“A” links with the greatest MSU load.	Shows links that may be approaching full capacity.
B-link load	“B” links with the greatest MSU load.	Shows links that may be approaching full capacity.
C-link load	“C” links with the greatest MSU load.	Shows links that may be approaching full capacity.
D-link load	“D” links with the greatest MSU load.	Shows links that may be approaching full capacity.

**Protocol analysis application.** The protocol analysis application provides full message and field decodes to analyze traffic on links or groups of links from anywhere in the network. This traffic can be filtered for specific messages. The data source may be either links monitored in real time or a buffer of previously stored data. A full decode of any message captured is available simply by pointing and clicking on the required line.

**Real-time protocol analysis.** Messages of interest on the real-time display window, shown in Figure 33.9, may be individually decoded simply by pointing and clicking on the message. This produces the message decode window. A point-and-click on an individual octet produces the field decode window. Facilities are available to search a buffer of previously captured data for particular messages, to filter messages according to an operator-defined template, and to start and/or stop message storage on receipt of operator-defined trigger conditions. Filters and triggers may be



**Figure 33.9** The upper display shows SS7 messages captured on an SS7 link. Clicking on any message in this display brings up the center display with the message decoded to the octet level. Clicking on an octet brings up the lower display with the octet decoded to fields.

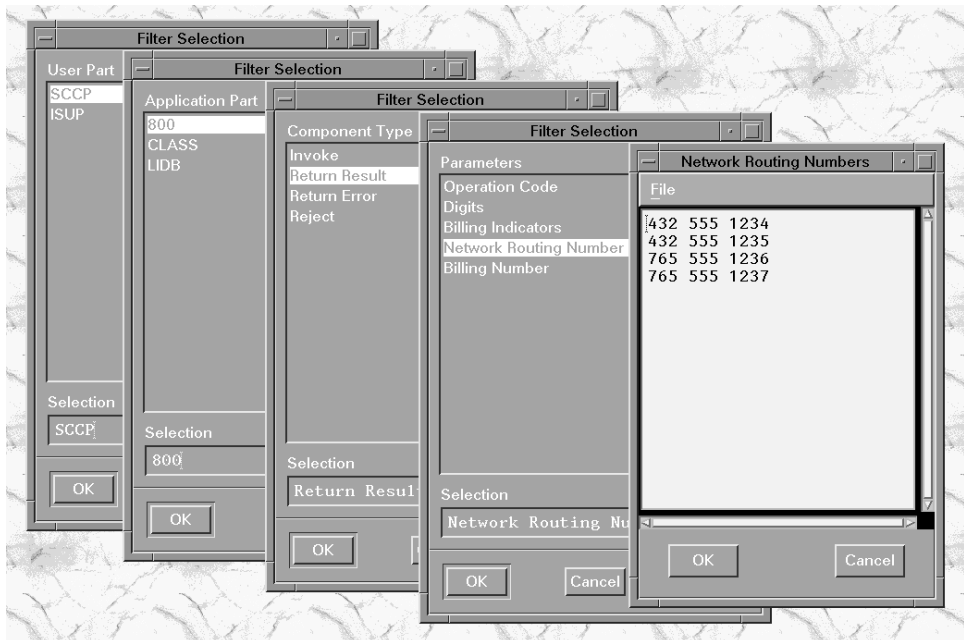
logically combined, applied to any desired link or set of links, and may be used on traffic captured in real time or on previously captured and stored data.

Setting up the protocol analysis application typically consists of the following steps:

- Selecting links
- Setup of triggers and filters
- Starting the measurement
- Real-time display
- Stopping the measurement
- Analysis of captured data

Various filters may be applied to the SS7 traffic being displayed, or stored in the buffer in order to select the required data. In addition, triggers may be set to start and stop information capture. Figure 33.10 shows a particular example of filter selection, with the highlighted line in each window showing the operator's choice. The windows represent a sequence of choices. As each choice is made, the operator is presented with the next level of parameters from which to choose.

Triggers to control the storage of traffic in a buffer are set up using the same process as shown for selecting filters in Figure 33.10. Triggers may be logically combined, and commonly used filter and trigger definitions may be stored and recalled.



**Figure 33.10** Filters for displaying live traffic or captured traffic that has been stored in a buffer are set up in a sequence of choices in the displays shown here.



The user is able to:

- Define triggers and filters
- Modify existing trigger and filter definitions
- Delete existing trigger and filter definitions
- Save a set of trigger and filter definitions to disk
- Restore a set of trigger and filter definitions from disk
- Print trigger and filter definitions in a readable format
- View trigger and filter definitions in a readable format

***Defining triggers and filters.*** Message templates, detailing fields and field values, are used to define system triggers and filters for messages that are of interest to the user. When defining a message template, at the highest level the user can select from a list of available protocols. A series of selection windows is presented, which will lead the user through the protocol, allowing fields to be selected and values for those fields to be specified. The operator can select the fields as required from any part of the message.

Once a new message template has been defined, the user can give the template a name and one-line description for identification. Once identified, the message template may be stored for future use. Stored triggers and filters are displayed as they are defined. Stored triggers and filters may be associated, named, and stored to disk as a configuration file.

***Start triggers.*** A *start trigger* is an event that defines when traffic capture will start. The activation of a start trigger controls data capture simultaneously on all channels defined by the user. The following trigger types are available:

- Start capturing on user request
- Start capturing at specified date and time
- Start capturing on receipt of SU matching a specified message template
- Start capturing on receipt of erroneous SU of type specified by user

The user may define up to five start triggers per channel and define the action performed by the trigger. When a start trigger causes more than one channel to start capturing data, all selected channels are started simultaneously. The system provides the user with the ability to save a start trigger for subsequent recall.

***Stop triggers.*** A *stop trigger* is an event that defines when traffic capture will stop. Six types of stop triggers are available to the user:

- Stop capturing on user request
- Stop capturing at specified date and time
- Stop capturing on receipt of SU matching a specified message template
- Stop capturing on receipt of erroneous SU of type specified by user
- Stop capturing when capture buffer is filled
- Upon receiving any type of erroneous SU

The user may define the action to be taken when a stop trigger occurs. Once one channel has satisfied its stop trigger criteria, data capture stops on all channels configured for protocol analysis by the user. If the stop-on-buffer-full trigger is not selected on a channel, the system treats the capture buffer as a circular buffer.

**Filters.** Filters are used by the system to control which traffic data is written to the capture buffer and display. Filters can be used to include or exclude matching traffic. The user may define the action to include all messages, include or exclude all SUs matching a message template, include or exclude erroneous SUs, or to exclude retransmissions for PCR channels. Where more than one filter is defined for a channel, the system will apply each filter to each message in turn. Only those messages that pass successfully through all filters will be available for further processing. When a filter is applied to a capture buffer, the user has the ability to save the filtered data to a new capture buffer.

**Post-capture protocol analysis.** During a protocol analysis session, traffic is stored in a capture buffer. At the end of the session the user can save this capture buffer for future use. This provides users with the ability to analyze traffic in their own time, and review previous problems and sessions stored by other users. Users are able to:

- Select previously captured traffic from disk
- Display previously captured traffic
- Search for messages in previously captured traffic
- Apply filters to previously captured traffic
- Trace telephone calls in previously captured traffic
- Print previously captured traffic
- Decode selected messages
- Save (subset of) previously captured traffic to disk

**Decoding selected messages.** A “personality” defines the protocols and versions of those protocols that are available. Each personality contains:

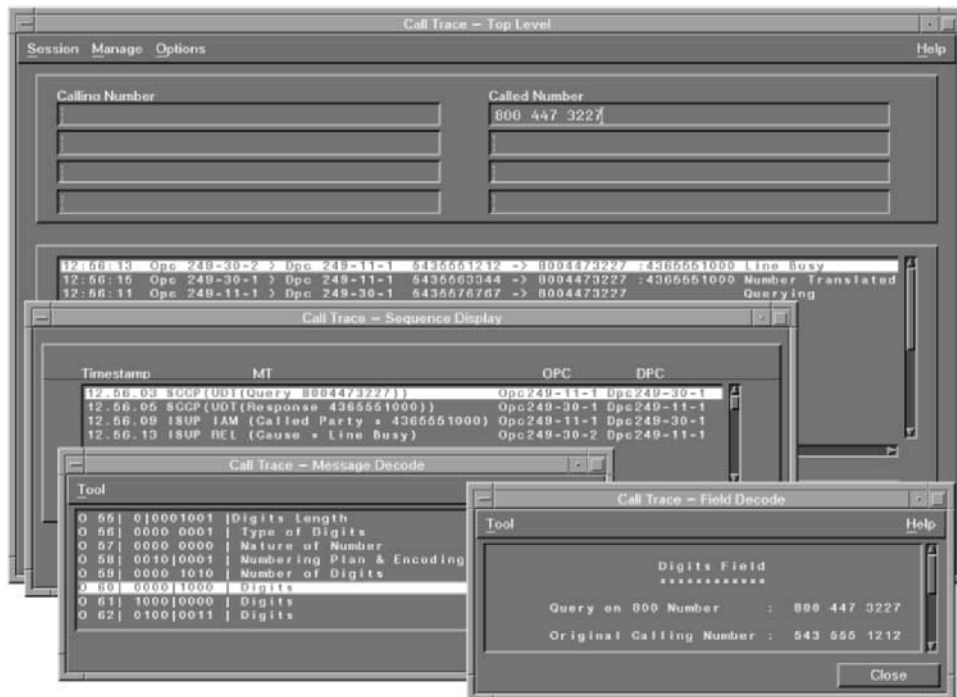
- A unique name for identification purposes
- Point code size and format to be used
- Service indicator mnemonics
- Sub-service field mnemonics
- Message type mnemonics for relevant service indicators
- H0/H1 mnemonics for relevant service indicators
- Level 2 decodes to be used
- MTP decodes to be used
- Service decodes to be used
- User-data decodes to be used

The personality affects how a message is decoded and displayed and also the selections that are available when defining a message template. A number of different personality specifications are available.

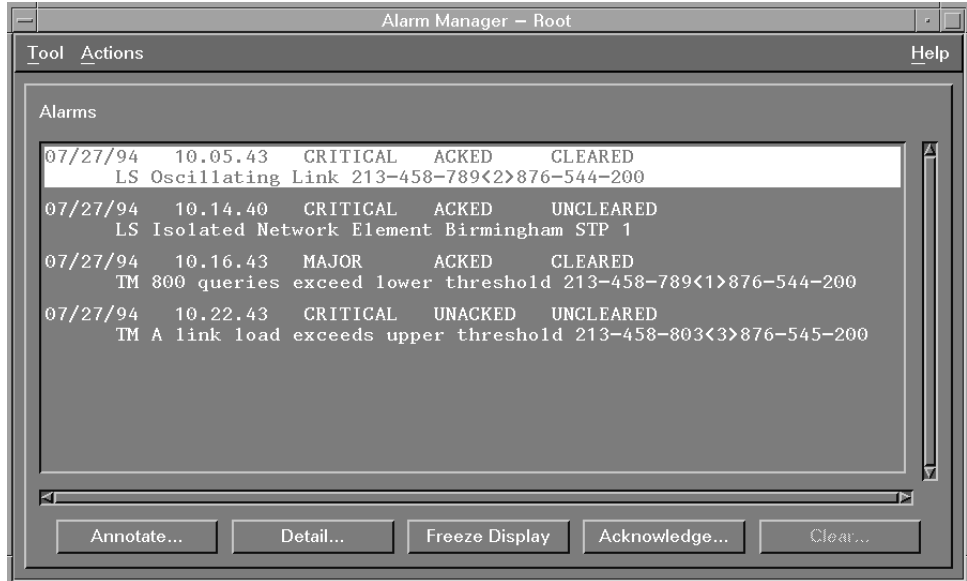
**Decode display.** The user is able to select a message from the display and decode it to the individual fields. On selecting a message to be decoded, a new window is displayed containing a description of each of the fields present in the selected message. The user is allowed to have multiple field decode windows for the same message and multiple message decode windows concurrently on display.

**Call trace application.** The call trace application applies to ISUP and 800 calls. It allows a user to trace one or more calls to or from a number, or between two numbers. The number to be traced can contain up to 26 digits or can be a combination of digits and wildcards. In addition, the user can choose the source of messages to be either a selection of signaling links in real time or from a set of previously recorded capture buffers. In the case of an 800 call, the interaction with the SCP also will be traced. Figure 33.11 shows an example of an 800 number call trace, together with the additional decode information available to the user.

The trace may be applied to one or more monitored linksets. Telephone numbers to be traced can be specified with trailing wild cards. Decoding in call trace follows a



**Figure 33.11** The call trace application begins with the calls listed in the top screen. Sequential pointing and clicking allows the user to drill down on any call to field decode level shown in the far right-hand screen.



**Figure 33.12** The alarm manager application main display provides an overview of all current alarms and is the primary interface to alarm management.

similar format to that used in protocol analysis. By simply pointing and clicking on selected lines, the call can be decoded through message sequence, message decode, right down to field decode.

**Alarm manager application.** The alarm manager logs to the central site server all alarms (threshold violations) generated by the link status monitor and traffic monitor applications, while reporting alarm status to authorized users via a graphical user interface. Figure 33.12 shows an alarm manager main display.

The alarm manager main window shows all current alarms reported from the link status monitor and traffic monitor applications. The main window shows key parameters including:

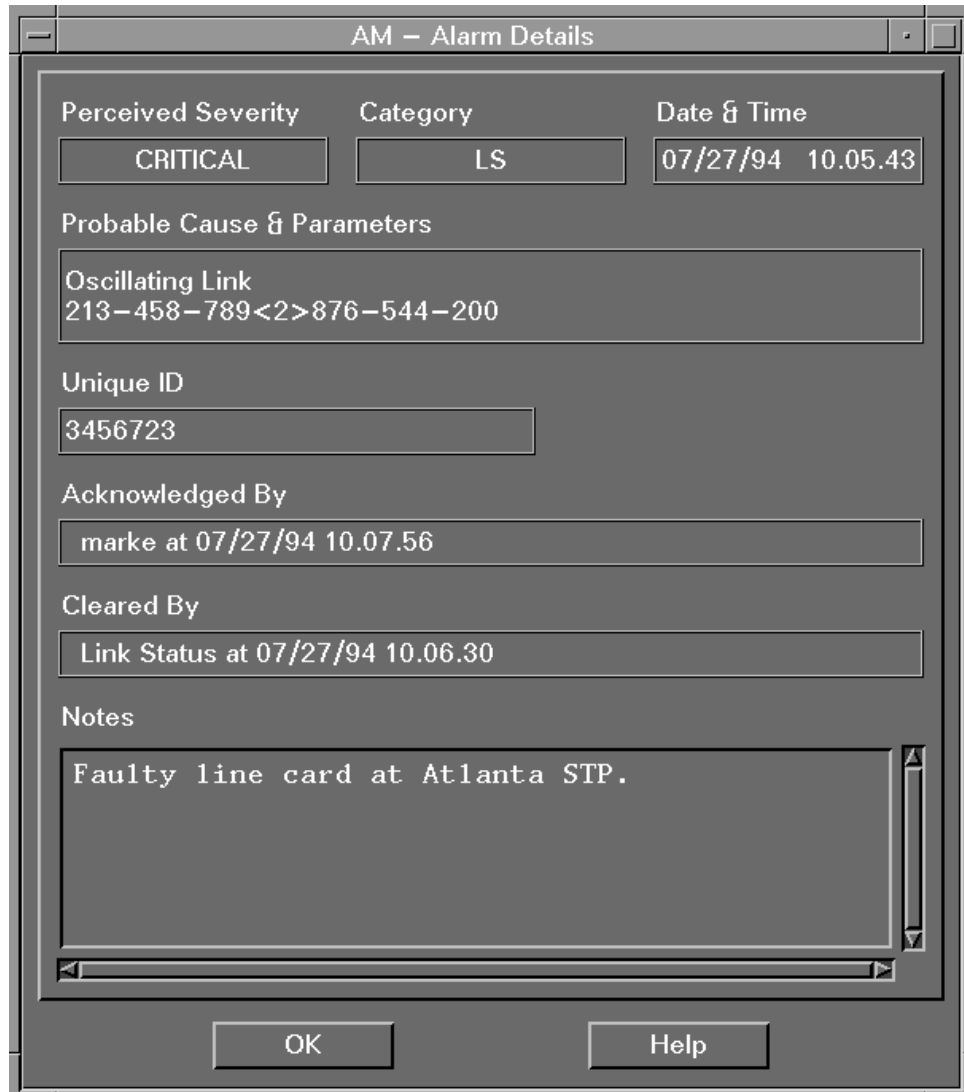
- Time and date of alarm
- Alarm source (link status monitor or traffic monitor)
- Alarm severity
- If an alarm has been acknowledged
- Alarm status (cleared or uncleared)
- Alarm description

The alarm manager main window allows the user to:

- Scroll through alarm log
- Freeze or unfreeze log window

- Acknowledge or clear an alarm
- Select a specific alarm to view more detail or add (annotate) user notes

**Alarm details.** When a user highlights a particular alarm and selects the detail option from the alarm manager main window, the alarm details window is activated (Figure 33.13). The alarm details window shows information common to the alarm manager main window plus a unique alarm ID number and the ID of the user who ac-



**Figure 33.13** Highlighting an alarm shown in the alarm main window and clicking on "Detail" brings up this window, which acts as a trouble log for the alarm.

**782 Network Management**

knowledged the alarm. The alarm details window also allows the user to add a note up to 2000 characters long to an alarm.

Alarm manager features include:

- Logging of all link status monitor and traffic monitor alarms
- Ability of users or applications to clear alarms
- Unique alarm ID number
- Ability to annotate alarms
- Manual or automatic alarm reports
- Manual or automatic alarm archiving
- User access defined by system administrator
- Multiple databases supported
- All alarm details available through a named Unix pipe

**Statistics application.** The statistics application offers great flexibility in terms of measurement setup and result reporting. Each system user can create multiple measurement groups, specify the measurement period and its start and stop times, and choose where to apply these on the network. The statistics application can be used as a surveillance tool, giving an immediate indication of problems, and as a planning tool, gathering data to predict future performance. Each measurement can be checked against a user-defined threshold; a violation will generate an immediate alarm.

Once data has been captured from an SS7 network, the statistics application provides comprehensive reporting tools. At the end of each measurement period, three different reports are generated: a summary view defined on a per-measurement basis, with all the data displayed from the last complete measurement period; an historical view containing data captured over all the measurement periods in the measurement session; and an event log of all threshold violations. These reports can be viewed as tables or graphs, dumped to a printer, or saved in a file for later analysis.

Selecting the parts of the network on which to focus can provide detailed information. For example, selecting a single STP or SCP can provide information on:

- How much traffic is it handling?
- How much traffic is it generating?
- What is the traffic origin?
- What is traffic destination?
- What type of traffic is it (the STP) receiving? Is it SNM, TUP, ISUP, SCCP, TCAP, etc.?
- What services (SCP) are being used?

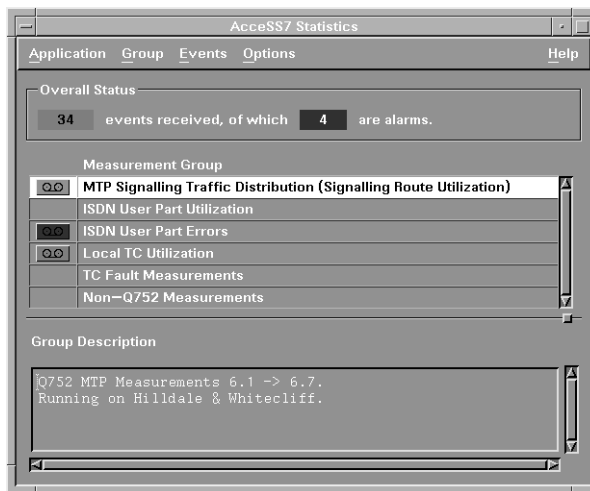
For example, selecting both members of an STP pair can provide information on:

- Are the routing tables correct?
- Is load sharing working correctly?
- Are there any concentrations of traffic to particular destinations?

The application is allocated a defined set of system resources, i.e., bandwidth and processing power. The actual resource consumed depends on the number of users at any one time, as well as the complexity of the measurements they create. The application calculates the resources required for each measurement, and notifies the user of any current resource limitations. This ensures that a user cannot inadvertently dominate the system resources and block other users or other applications.

Figure 33.14 shows the main statistics display and indicates the current set of measurements being run. Figure 33.15 shows the three main display formats for statistics measurement results.

**Data capture application.** The data capture application provides a complete picture of network activity leading up to a significant network event. The stored data can be analyzed to find the cause of the event quickly and easily, and thereby minimize network downtime. The application guarantees the capture and storage of any or all signaling messages, regardless of network loading. It can handle a sustained load of 100 percent signaling messages on all links in the network. This is achieved by allocating dedicated processing and storage capabilities to this application. A data capture card in each card cage pulls all the raw SS7 signaling units from each SS7 link, and passes them directly to an external data capture processor.



**Figure 33.14** The statistics main control window displays the status of each of the current measurement groups and indicates which measurement groups have generated alarms and events.

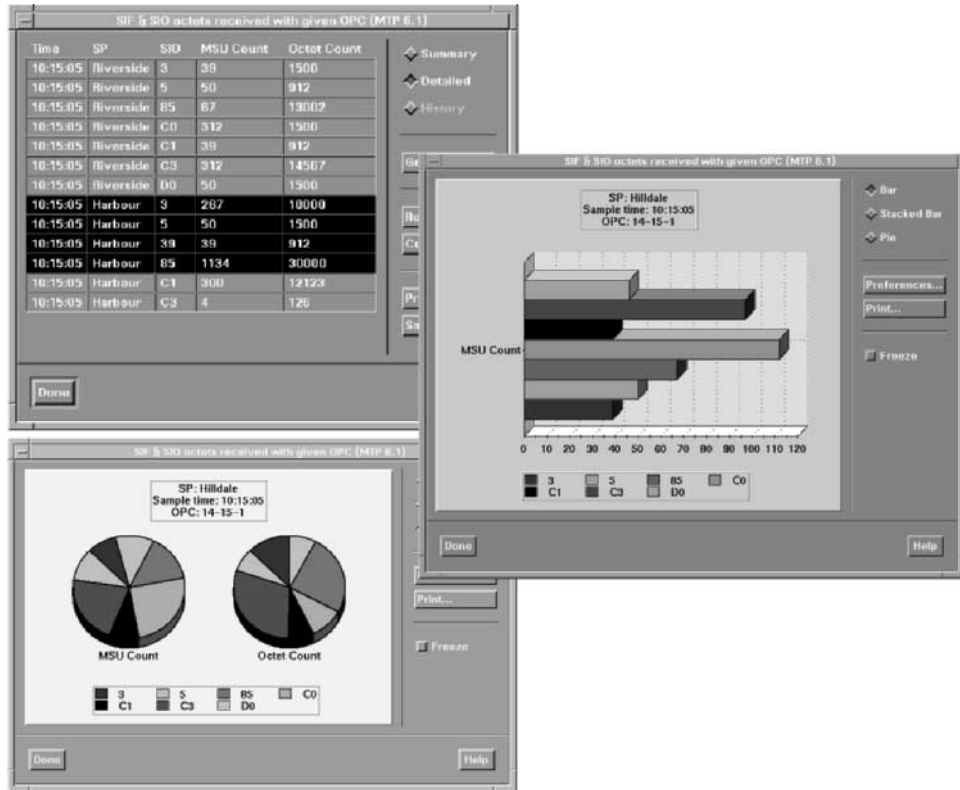


Figure 33.15 Statistics results can be displayed in three formats: tabular, pie chart, and bar chart.

The data capture processor timestamps each signaling unit and then uses filters to determine if the data should be discarded or stored. The default condition is to capture all data on a circular buffer. Once the buffer is full, the oldest data is overwritten. Buffers can be frozen by user-definable triggers. Users graphically control the data capture processors, and set up the capture triggers and filters. The length of storage is a function of network load, and disk storage capacity. The system can be configured to give virtually any desired storage duration. Captured data can be archived to DAT (digital audio tape), or can be analyzed by the data capture processor. This analysis uses the protocol analysis application.

**Call detail recording.** Generation of a *call detail record* (CDR) requires collecting the set of messages that describe the call. Referring back to Figure 33.2, a CDR is generated by extracting information from the SS7 messages flowing through the SS7 network between the originating and terminating switches. Starting with the initial address message (IAM), the system tracks the signaling messages relating to the call and generates a CDR when the call finally completes. These messages together define the call in terms of the called and calling numbers and the call duration.



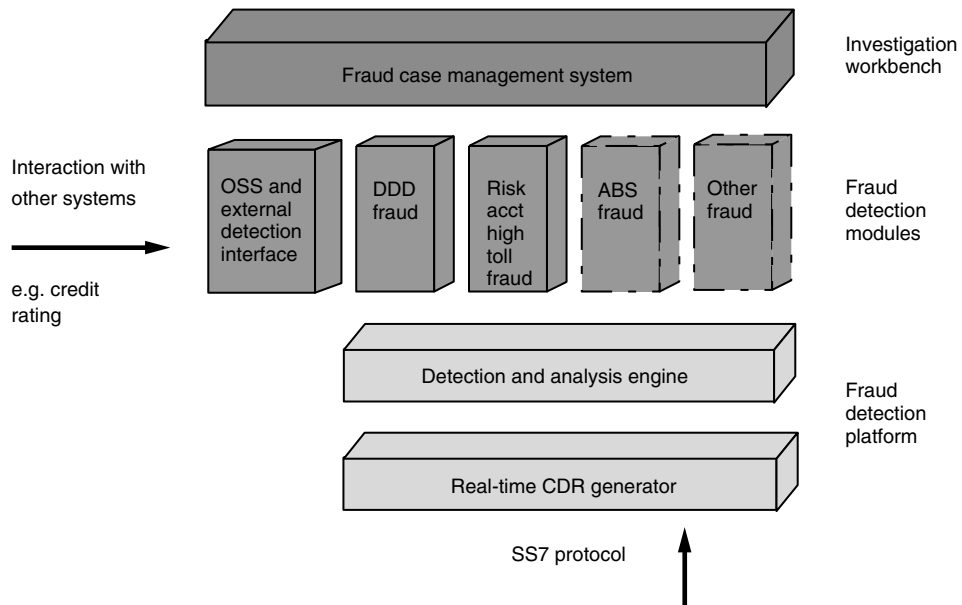
Additional fields within the call detail record are used to convey the carrier and type of call: ISDN, pay phone, etc. Once produced, these call detail records can be fed to a variety of business applications including retail billing, wholesale billing, real-time advice of charge, billing verification, or fraud detection.

A CDR generator collects and correlates CDRs when the messages relating to each call have been diversely routed through the SS7 network (that is, not all present at the same STP monitoring site). By means of sophisticated filtering, cross-triggering, searching, and matching, it can achieve this while still being scalable to networks with hundreds or thousands of call attempts per second.

**Fraud management toolkit.** The fraud management toolkit provides fraud investigators with a source of alerts of suspect calling patterns in real time, as well as a means of pulling these alerts into fraud investigation cases. The architecture of the toolkit is shown in Figure 33.16.

CDRs from the real-time CDR generator are fed to a programmable detection and analysis engine, which matches calls against known suspicious calling patterns and generates fraud alerts when matches are detected. Fraud detection modules control the CDR builder and detection analysis engine to detect selected types of known fraudulent calling patterns and generate appropriate alerts.

The alerts are supplied to a case manager, which gives each alert a severity weighting number and then attaches each alert to a case depending on some key such as the calling party number. Each case will have a severity weighting that is the sum of the



**Figure 33.16** The fraud manager toolkit comprises layered applications that analyze real-time data from the probes distributed throughout the SS7 network. This structure and the open interfaces allow incorporation of other sources of fraud detection information.

severity weightings of the alerts attached to the case. This provides a prioritized list of suspect fraud cases for investigation. Fraud investigators access cases from a workstation screen using the workbench graphical user interface and investigate alerts, add notes, set disposition on cases, and refer cases to other departments for action.

The fraud management toolkit is designed with an open architecture, which permits interaction with other operations and business support systems so that fraud cases can include relevant customer data such as date of installation of service and customer credit records. The interface with these systems is an optional feature of the fraud management toolkit and requires the development of custom software.

The fraud management toolkit offers several significant benefits to fraud investigation departments.

- The workbench provides a common user interface for all fraudulent calling patterns detected, making it easy to investigate new types of fraud.
- The case manager focuses investigator effort on the most serious fraudulent behavior by correlating alerts into cases based on calling party number, presenting cases as a prioritized list with the most serious at the top.
- The toolkit has a modular structure that makes it easy to add functionality to detect new fraudulent calling patterns in future.

A benefit of the open architecture is that it permits interworking with other fraud detection systems, such as billing record-based systems and calling card validation databases, to correlate alert activity detected by these systems with the cases in the case manager. These interworking functions require separate custom software development projects.

### 33.5 Summary

The signaling network (SS7) is in effect the nervous system of the telephone network. It controls setup and release of calls. It enables services beyond simple voice conversations. It provides numerous options for charging for these services.

SS7 network congestion and performance degradation of the signaling network will seriously degrade telephone service. Failure of the signaling network is disastrous, halting all telephone service controlled by the failed SS7 network. A signaling network monitoring system provides the real-time data necessary to manage the SS7 network and keep it operating at top performance.

As important as is maintaining the health of the SS7 network, additional value comes from the data the SS7 signaling monitoring system provides. Call detail records are created in real time to detect and control fraudulent use of telephone services, a serious cause of lost revenue. Cross-carrier billing is made more convenient than ever before. Traffic analysis in real time reveals the effects of mass calling events or new services with call durations well outside the ranges for which the voice network was engineered. With such real-time traffic analyses, carriers can adjust tariffs and re-engineer the network to compensate.