

WORDWARE'S GAME AND GRAPHICS LIBRARY

ADVANCED 3D GAME PROGRAMMING WITH DIRECTX® 10.0

PETER WALSH



Advanced 3D Game Programming with DirectX[®] 10.0

Peter Walsh

Wordware Publishing, Inc.

Library of Congress Cataloging-in-Publication Data

Walsh, Peter, 1980-

Advanced 3D game programming with DirectX 10.0 / by Peter Walsh.

p. cm.

Includes index.

ISBN 10: 1-59822-054-3

ISBN 13: 978-1-59822-054-4

1. Computer games--Programming. 2. DirectX. I. Title.

QA76.76.C672W3823 2007

794.8'1526--dc22

2007041625

© 2008, Wordware Publishing, Inc.

All Rights Reserved

1100 Summit Avenue, Suite 102
Plano, Texas 75074

No part of this book may be reproduced in any form or by
any means without permission in writing from
Wordware Publishing, Inc.

Printed in the United States of America

ISBN 10: 1-59822-054-3

ISBN 13: 978-1-59822-054-4

10 9 8 7 6 5 4 3 2 1

0712

DirectX is a registered trademark of Microsoft Corporation in the United States and/or other countries.

Other brand names and product names mentioned in this book are trademarks or service marks of their respective companies. Any omission or misuse (of any kind) of service marks or trademarks should not be regarded as intent to infringe on the property of others. The publisher recognizes and respects all marks used by companies, manufacturers, and developers as a means to distinguish their products.

This book is sold as is, without warranty of any kind, either express or implied, respecting the contents of this book and any disks or programs that may accompany it, including but not limited to implied warranties for the book's quality, performance, merchantability, or fitness for any particular purpose. Neither Wordware Publishing, Inc. nor its dealers or distributors shall be liable to the purchaser or any other person or entity with respect to any liability, loss, or damage caused or alleged to have been caused directly or indirectly by this book.

All inquiries for volume purchases of this book should be addressed to Wordware Publishing, Inc., at the above address. Telephone inquiries may be made by calling:

(972) 423-0090

To my wife, Lisa

Peter

This page intentionally left blank.

Contents

Acknowledgments	xiii
Introduction	xv
Chapter 1 Windows	1
A Word about Windows	1
Hungarian Notation	3
General Windows Concepts	4
Message Handling in Windows	5
Processing Messages.	6
Hello World—Windows Style	7
Explaining the Code	11
Registering the Application	13
Initializing the Window	13
WndProc—The Message Pump	17
Manipulating Window Geometry	17
Important Window Messages	20
Class Encapsulation	23
COM: The Component Object Model	30
Conclusion.	33
Chapter 2 Getting Started with DirectX 10.	35
What Is DirectX?	35
Installation	36
Setting Up VC++	36
What Happened to DirectDraw?	39
Direct3D	40
2D Graphics 101	41
Textures	44
Complex Textures	46
Describing Textures	46
The ID3D10Texture2D Interface	50
Texture Operations	51
Modifying the Contents of Textures.	51
Creating Textures	52
Implementing Direct3D with cGraphicsLayer.	53
Creating the Graphics Layer.	58

Initializing Direct3D	58
Step 1: Creating a Device and Swap Chain	58
Step 2: Creating a Render Target View	62
Step 3: Putting It All Together	62
Shutting Down Direct3D	64
Sample Application: Direct3D Sample	64
Conclusion.	66
Chapter 3 Input and Sound	67
DirectInput	67
Devices	68
Receiving Device States	70
Cooperative Levels	73
Application Focus and Devices.	73
The DirectInput Object	74
Implementing DirectInput with cInputLayer	74
Additions to cApplication	87
Sound	87
The Essentials of Sound	88
DirectSound Concepts.	89
DirectSound Buffers.	90
Operations on Sound Buffers	93
Loading WAV Files.	96
Implementing DirectSound with cSoundLayer.	103
Creating the DirectSound Object	103
Setting the Cooperative Level	104
Grabbing the Primary Buffer	105
The cSound Class	108
Additions to cApplication	114
Application: DirectSound Sample	114
Conclusion	119
Chapter 4 3D Math Foundations	121
Points.	121
The point3 Structure	124
Basic point3 Functions	125
Assign.	125
Mag and MagSquared	126
Normalize	126
Dist.	126
point3 Operators	127
Addition/Subtraction	127
Vector-Scalar Multiplication/Division.	129
Vector Equality	130
Dot Product	131

Cross Product	134
Polygons	135
Triangles	138
Strips and Fans	139
Planes.	141
Defining Locality with Relation to a Plane	144
Back-face Culling	147
Clipping Lines	148
Clipping Polygons	149
Object Representations	153
Transformations	156
Matrices	156
The matrix4 Structure	166
Translation	168
Basic Rotations	169
Axis-Angle Rotation	170
The LookAt Matrix.	172
Perspective Projection Matrix	174
Inverse of a Matrix	174
Collision Detection with Bounding Spheres	175
Lighting.	178
Representing Color	178
Lighting Models	180
Specular Reflection	182
Light Types	183
Parallel Lights (or Directional Lights)	183
Point Lights	184
Spotlights	185
Shading Models	186
Lambert	186
Gouraud	186
Phong.	187
BSP Trees.	187
BSP Tree Theory	188
BSP Tree Construction	189
BSP Tree Algorithms	194
Sorted Polygon Ordering	194
Testing Locality of a Point	196
Testing Line Segments	196
BSP Tree Code	197
Wrapping It Up.	207
Chapter 5 Artificial Intelligence	209
Starting Point.	210
Locomotion	210

Steering—Basic Algorithms	211
Chasing	211
Evading	211
Pattern-based AI	212
Steering—Advanced Algorithms.	213
Potential Functions	214
The Good	215
The Bad	215
Application: potentialFunc	216
Path Following	218
Groundwork	220
Graph Theory	221
Using Graphs to Find Shortest Paths.	225
Application: Path Planner.	227
Motivation	230
Nondeterministic Finite Automata (NFAs).	230
Genetic Algorithms	233
Rule-Based AI	234
Neural Networks	235
A Basic Neuron.	236
Simple Neural Networks.	238
Training Neural Networks	240
Using Neural Networks in Games	241
Application: NeuralNet.	241
Extending the System	253

Chapter 6 Multiplayer Internet Networking with UDP 255

Terminology	255
Endianness	255
Network Models	257
Protocols	258
Packets	259
Implementation I: MTUDP	260
Design Considerations	260
Things to Watch Out For	260
Mutexes	263
Threads, Monitor, and the Problem of the try/throw/catch Construction	264
MTUDP: The Early Years.	265
MTUDP::Startup() and MTUDP::Cleanup()	266
MTUDP::MTUDP() and MTUDP::~~MTUDP()	267
MTUDP::StartListening().	267
MTUDP::StartSending()	268
MTUDP::ThreadProc().	269
MTUDP::ProcessIncomingData()	270

MTUDP::GetReliableData()	271
Reliable Communications	271
cDataPacket	271
cQueueIn	272
cHost	274
MTUDP::ReliableSendTo()	278
cUnreliableQueueIn	284
cHost::AddACKMessage()/cHost::ProcessIncomingACKs()	285
cNetClock	290
Implementation 2: Smooth Network Play	292
Geographic and Temporal Independence	293
Timing Is Everything	294
Pick and Choose	295
Prediction and Extrapolation	295
Conclusion	297
Chapter 7 Direct3D Fundamentals	299
Introduction to D3D	299
Getting Started with Direct3D	300
Step 1: Creating the ID3D10Device and Swap Chain	300
Step 2: Creating a Depth/Stencil Buffer	302
Bringing It All Together	306
Step 3: Creating a Viewport	308
Step 4: Creating a Default Shader	309
Introduction to Shaders	309
Your First HLSL Shader	311
The Vertex Shader	312
The Pixel Shader	313
The Technique	313
Setting Up the Shader in Code	315
More about Depth Buffers	322
Stencil Buffers	325
Vertex Buffers	325
Lighting with Shaders	328
Application: D3D View	330
The .o3d Format	330
The cModel Class	331
Chapter 8 Advanced 3D Techniques	341
Animation Using Hierarchical Objects	341
Forward Kinematics	343
Inverse Kinematics	346
Application: InvKim	349
Parametric Curves and Surfaces	355
Bezier Curves and Surfaces	355

Bezier Concepts	355
The Math	358
Finding the Basis Matrix	360
Calculating Bezier Curves	361
Forward Differencing	363
The cFwdDiffIterator Class	365
Drawing Curves	367
Drawing Surfaces	367
Application: Teapot	369
B-Spline Curves	376
Application: BSpline	377
Subdivision Surfaces	379
Subdivision Essentials	380
Triangles vs. Quads	382
Interpolating vs. Approximating	382
Uniform vs. Non-Uniform	383
Stationary vs. Non-Stationary	383
Modified Butterfly Method Subdivision Scheme	383
Application: SubDiv	387
Progressive Meshes	399
Progressive Mesh Basics	401
Choosing Our Edges	402
An Edge Selection Algorithm	403
Quadric Error Metrics	403
Implementing a Progressive Mesh Renderer	405
Radiosity	406
Radiosity Foundations	407
Progressive Radiosity	410
The Form Factor	411
Application: Radiosity	412
Conclusion	416
Chapter 9 Advanced Direct3D	417
Alpha Blending	417
The Alpha Blending Equation	418
A Note on Depth Ordering	419
Enabling Alpha Blending	419
Using Alpha Blending from C++	419
Using Alpha Blending from Shaders	422
Texture Mapping 101	423
Fundamentals	424
Affine vs. Perspective Mapping	425
Texture Addressing	426
Wrap	426
Mirror and Mirror Once	427

Clamp	427
Border Color	428
Texture Wrapping	429
Texture Aliasing	430
MIP Maps	432
Filtering	433
Point Sampling	433
Linear Filtering	434
Anisotropic Filtering	435
Textures in Direct3D	436
Texture Loading	437
DDS Format	437
The cTexture Class	438
Activating Textures	440
Creating a Shader View	441
Adding Textures to the Shader	441
Sending the Texture to the Shader	442
Texture Sampling	442
Texture Mapping 202	443
Texture Arrays	443
Effects Using Multiple Textures	443
Light Maps (a.k.a. Dark Maps)	444
Environment Maps	446
Spherical Environment Maps	446
Cubic Environment Maps	450
Specular Maps	452
Detail Maps	452
Application: Detail	455
Glow Maps	463
Gloss Maps	463
Other Effects	464
Application: MultiTex	465
Using the Stencil Buffer	483
Overdraw Counter	485
Dissolves and Wipes	485
Conclusion	485
Chapter 10 Scene Management	487
The Scene Management Problem and Solutions	487
Quadtrees/Octrees	488
Portal Rendering	490
Portal Rendering Concepts	491
Exact Portal Rendering	497
Approximative Portal Rendering	498
Portal Effects	499

- Mirrors 499
- Translocators and Non-Euclidean Movement 502
- Portal Generation 503
- Precalculated Portal Rendering (with PVS). 505
 - Advantages/Disadvantages 506
 - Implementation Details 506
- Application: Mobots Attack! 507
 - Interobject Communication 507
 - Network Communication 511
 - Code Structure 514
- Closing Thoughts 515
- Appendix An STL Primer. 517**
 - Templates 517
 - Containers 518
 - Iterators 519
 - Functors 521
- Index. 523**

Acknowledgments

The DirectX 10 revision of this book would not have been possible without the help of Tim McEvoy at Wordware Publishing. Thanks, Tim. Thanks to Adrian Perez and Dan Royer for paving the way with the first version of the book, *Advanced 3-D Game Programming with DirectX 7.0*. I'd also like to thank my wife, Lisa, for putting up with all the late nights that it took to get this project completed. Also, thanks to all my friends and colleagues at Realtime Worlds, where I nearly went insane and learned a lot more than I thought possible while developing *Crackdown*.

Peter Walsh

This page intentionally left blank.

Introduction

A wise man somewhere, somehow, at some point in history, may have said the best way to start a book is with an anecdote. I would never question the words of a wise man who may or may not have existed, so here we go.

When I was a freshman in high school, I took the required biology class that most kids my age end up having to take. It involved experiments, lab reports, dissecting of various animals, and the like. One of my lab partners was a fellow named Chris V. We were both interested in computers and quickly became friends, to the point where talking about biology in class was second to techno-babble.

One night, in the middle of December, Chris called me up. The lab report that was due the next day required results from the experiment we had done together in class, and he had lost his copy of our experiment results. He wanted to know if I could copy mine and bring them over to his place so he could finish writing up the lab. Of course, this was in those heinous pre-car days, so driving to his house required talking my parents into it, finding his address, and various other hardships. While I was willing to do him the favor, I wasn't willing to do it for free. So I asked him what he could do to reciprocate my kind gesture.

"Well," he said, "I guess I can give you a copy of this game I just got."

"Really? What's it called?" I said.

"*Doom*. By the Wolf 3D guys."

"It's called *Doom*? What kind of name is that??"

After getting the results to his house and the game to mine, I fired up the program on my creaky old 386 DX-20 clone, burning rubber with a whopping 4 MB of RAM. As my space marine took his first tenuous steps down the corridors infested with hellspawn, my life changed. I had done some programming before in school (Logo and Basic), but after I finished playing the first time, I had a clear picture in my head of what I wanted to do with my life: I wanted to write games, something like *Doom*. I popped onto a few local bulletin boards and asked two questions: What language was the game written in, and what compiler was used?

Within a day or so, I purchased Watcom C 10.0 and got my first book on C programming. My first C program was "Hello, World." My second was a slow, crash-happy, non-robust, wireframe spinning cube.

I tip my hat to John Carmack, John Romero, and the rest of the team behind *Doom*; my love for creating games was fully realized via their

masterpiece. It's because of them that I learned everything that I have about this exceptionally interesting and dynamic area of computer acquired programming. The knowledge that I have is what I hope to fill these pages with, so other people can get into graphics and game programming.

I've found that the best way to get a lot of useful information down in a short amount of space is to use the tried-and-true FAQ (frequently asked questions) format. I figured people needed answers to some questions about this book as they stood in their local bookstore trying to decide whether or not to buy it, so read on.

Who are you? What are you doing here?

I am a professional game programmer and have been for quite a few years. I started out like most people these days, getting extremely interested in how games worked after *Doom* came out. After teaching myself programming, I moved on to study for a degree in computer game development at Abertay University in Dundee, Scotland. After that I went on to work for a short while with IC-CAVE, which is a think tank for the next generation of gaming technology. Over the years I've worked on games like *F1 Career Challenge*, *Harry Potter and the Chamber of Secrets*, *SHOX*, and *Medal of Honor: Rising Sun*. Most recently I've worked on *Crackdown* for the Xbox 360. I've developed games for just about every platform you can think of.

I've also read so many programming books that I reckon I have personally wiped out half of the Amazon rainforest. So hopefully all that material will help me write this book in a way that avoids all the pitfalls that other authors have fallen into. I really hope you learn a lot from this book. If you have any questions along the way that you just can't get to the bottom of, please email me at peter_stpwalsh@yahoo.co.uk. Unfortunately, after printing that email address in a previous book it was bombarded by junk mail from spammers and became almost unusable. However, Hotmail has gotten better lately, so hopefully your questions will get through to me!

Why was this book written?

I've learned from many amazingly brilliant people, covered a lot of difficult ground, and asked a lot of dumb questions. One thing that I've found is that the game development industry is all about sharing. If everyone shares, everyone knows more stuff, and the net knowledge of the industry increases. This is a good thing because then we all get to play better games. No one person could discover all the principles behind computer graphics and game programming themselves, and no one can learn in a

vacuum. People took the time to share what they learned with me, and now I'm taking the time to share what I've learned with you.



Note: For the update to DirectX 10 all of the source has also been updated to be Unicode compliant and use secure STL.

Who should read this book?

This book was intended specifically for people who know how to program already but have taken only rudimentary stabs at graphics/game programming or never taken any stab at all, such as programmers in another field or college students looking to embark on some side projects.

Who should not read this book?

This book was not designed for beginners. I'm not trying to sound arrogant or anything; I'm sure a beginner will be able to trudge through this book if he or she feels up to it. However, since I'm so constrained for space, often-times I need to breeze past certain concepts (such as inheritance in C++). If you've never programmed before, you'll have an exceedingly difficult time with this book.

What are the requirements for using the code?

The code was written in C++, using Microsoft Visual C++ 2005 Express Edition, which is available for download for free from Microsoft. The projects and solutions are provided on the downloadable files site (www.wordware.com/files/dx10). This book obviously focuses on DirectX 10, which at the time of writing will only run on Windows Vista or higher, with a DirectX 10 class graphics card. If you don't have these, you'll need to upgrade.

Why use Windows? Why not use Linux?

I chose to use Win32 as the API environment because 90% of computer users currently work on Windows. Win32 is not an easy API to understand, especially after using DOS coding conventions. It isn't terribly elegant either, but I suppose it could be worse. I could choose other platforms to work on, but doing so reduces my target audience by a factor of nine or more.

Why use Direct3D 10?

For those of you who have never used it, OpenGL is another graphics API. Silicon Graphics designed it in the early '90s for use on their high-end graphics workstations. It has been ported to countless platforms and operating systems. Outside of the games industry, in areas like simulation and academic research, OpenGL is the de facto standard for doing computer graphics. It is a simple, elegant, and fast API. Check out www.opengl.org for more information.

But it isn't perfect. First of all, OpenGL has a large amount of functionality in it. Making the interface so simple requires that the implementation take care of a lot of ugly details to make sure everything works correctly. Because of the way drivers are implemented, each company that makes a 3D card has to support the entire OpenGL feature set in order to have a fully compliant OpenGL driver. These drivers are extremely difficult to implement correctly, and the performance on equal hardware can vary wildly based on driver quality. DirectX has the added advantage of being able to move quickly to accommodate new hardware features. DirectX is controlled by Microsoft (which can be a good or bad thing, depending on your view of it), while OpenGL extensions need to be deliberated by committees. And finally, you simply can't get the latest shader support in OpenGL that DirectX provides.

Why use C++? Why not C, .NET, or Java?

I had a few other language choices that I was kicking around when planning this book. Although there are acolytes out there for Delphi, VB, and even C#, the only languages I seriously considered were C++, Java, and C. Java is designed by Sun Microsystems and is an inherently object-oriented language, with some high-level language features like garbage collection. C is about as low level as programming gets without dipping into assembly. It has very few if any high-level constructs and doesn't abstract anything away from the programmer.

C++ is an interesting language because it essentially sits directly between the functionality of the other two languages. C++ supports COM better than C does (this is more thoroughly discussed in Chapter 1). Also, class systems and operator overloading generally make code easier to read (although, of course, any good thing can and will be abused). Java, although very cool, is an interpreted language. Every year this seems to be less important: JIT compilation gets faster and more grunt work is handed off to the APIs. However, I felt C++ would be a better fit for the book. Java is still a very young language and is still going through a lot of change.

Do I need a 3D accelerator?

Yes. And it must be DirectX 10 class running with Windows Vista.

How hardcore is the C++ in this book?

Some people see C++ as a divine blade to smite the wicked. They take control of template classes the likes of which you have never seen. They overload the iostream operators for all of their classes. They see multiple inheritance as a hellspawn of Satan himself. I see C++ as a tool. The more esoteric features of the language (such as the iostream library) I don't use at all. Less esoteric features (like multiple inheritance) I use when it makes sense. Having a coding style you stick to is invaluable. The code for this book was written over an eleven-month period, plus another three for the revision, but I can pick up the code I wrote at the beginning and still read it because I commented it and used some good conventions. If I can understand it, hopefully you can too.

What are the coding conventions used in the source?

One of the greatest books I've ever read on programming was *Code Complete* (Microsoft Press). It's a handbook on how to program well (not just how to program). Nuances like the length of variable names, design of subroutines, and length of files are covered in detail in this book; I strongly encourage anyone who wants to become a great programmer to pick it up. You may notice that some of the conventions I use in this book are similar to the conventions described in *Code Complete*; some of them are borrowed from the great game programmers like John Carmack, and some of them are borrowed from source in DirectX and Win32.

I've tried really hard to make the code in this book accessible to everyone. I comment anything I think is unclear, I strive for good choices in variable names, and I try to make my code look clean while still trying to be fast. Of course, I can't please everyone. Assuredly, there are some C++ coding standards I'm probably not following correctly. There are some pieces of code that would get much faster with a little obfuscation.

If you've never used C++ before or are new to programming, this book is going to be extremely hard to digest. A good discussion on programming essentials and the C++ language is *C++ Primer* (Lippman et al.; Addison-Wesley Publishing).

Class/Structure Names

I prefix my own classes with a lowercase *c* for classes, a lowercase *s* for structs, a lowercase *i* for interfaces, and a lowercase *e* for enumerations (*cButton* or *sButton*).

There is one notable exception. While most classes are intended to hide functionality away and act as components, there are a few classes/structures that are intended to be instantiated as basic primitives. So for basic mathematic primitives like points and matrices, I have no prefix, and I postfix with the dimension of the primitive (2D points are *point2*, 3D points are *point3*, etc.). This is to allow them to have the same look and feel as their closest conceptual neighbor, *float*. For the same reason, all of the mathematic primitives have many overloaded operators to simplify math-laden code.

Variable Names

Semi-long variable names are a good thing. They make your code self-commenting. One needs to be careful though: Make them too long, and they distract from both the code itself and the process of writing it.

I use short variables very sporadically; *int i, j, k* pop up a lot in my code for loops and whatnot, but aside from that I strive to give meaningful names to the variables I use. Usually, this means that they have more than one word in them. The system I use specifies lowercase for the first word and initial cap for each word after that, with no underscores (an example would be *int numObjects*). If the last letter of a word is a capital letter, an underscore is placed to separate it from the next word (example: *class cD3D_App*).

A popular nomenclature for variables is Hungarian notation, which we touch on in Chapter 1. I'm not hardcore about it, but generally my floats are prefixed with "*f*," my ints with "*i*," and my pointers with "*p*" (examples: *float fTimer*; *int iStringSize*; *char *pBuffer*). Note that the prefix counts as the first word, making all words after it start with caps. (I find *pBuffer* much more readable than *pbuffer*.)

I also use prefixes to define special qualities of variables. Global variables are preceded with a "*g_*" (an example would be *int g_hInstance*), static variables are preceded with an "*s_*" (*static float s_fTimer*), and member variables of classes are preceded with an "*m_*" (*int m_iNumElements*).

Companion Files

The companion files can be downloaded from the following web site:

www.wordware.com/files/dx10

These files include the source code discussed in the book along with the game Mobots Attack!. Each chapter (and the game) has its own solution so you can use them independently of each other.

This page intentionally left blank.

Chapter I

Windows

Hello and welcome to the first stage of your journey into the depths of advanced 3D game development with DirectX 10. The first part of your exploration of the world of 3D game programming will start with a look at the foundations of Windows, upon which you will build your knowledge empire. We'll take a quick look through operations like opening and closing a program, handling rudimentary input, and painting basic primitives. If you're familiar with the Windows API, you should breeze through this chapter; otherwise, hold on to your seat!

In this chapter you are going to learn about:

- The theory behind Windows and developing with the Win32 API
- How Win32 game development differs from standard Windows programming
- Messages and how to handle them
- The standard message pump and real-time message pump
- Win32 programming
- COM (the Component Object Model)
- And much more!

A Word about Windows

Windows has had a fairly standard API (application programming interface) for the last 15 years. Anything you wrote for Windows 3.11 in 1992 will, within reason, continue to work in Windows Vista. So what you learn about Windows in this chapter has been the standard way of programming Windows for over a decade. All of the code in this book is designed to work with DirectX 10 and since that only runs on Windows Vista or later, that is the lowest operating system the code in the book is designed to run on. Most of the topics, however, can be modified to work with older versions of Windows and DirectX.

Modern Windows programs are fundamentally different in almost every way from ancient DOS or Win16 applications. In the old days, you had 100% of the processor time and 100% control over all the devices and files in the machine. You also needed an intimate knowledge of all of the devices on a user's machine. (You may remember arcane DOS or Win16 games, which almost always required you to input DMA and IRQ settings for devices such as sound cards.) When a game crashed, it often brought

down the entire system with it, which meant a very frustrating reboot for the end user.

With Windows Vista and DirectX 10, things are thankfully totally different. When your application executes, it is sharing the processor's cores with many other processes, all running concurrently (at the same time). You can't hog control of the sound card, video card, hard disk, or any other system resource for that matter. The input and output is abstracted away.

This is a good thing, apart from the steep learning curve when you first start out.

On one hand, Windows applications have a consistent look and feel. Almost any windowed application you create is automatically familiar to Windows users. They already know how to use menus and toolbars, so if you build your application with the basic Windows constructs, they can pick up the user interface quickly.

On the other hand, you have to put a lot of faith into Windows and other applications. Until DirectX came around, you needed to use the default Windows drawing commands (called the GDI, or graphical device interface). While the GDI can automatically handle any bit depth and work on any monitor, it's exceptionally slow. Windows Vista has a new display driver model that has completely rewritten the user interface code to use DirectX, so it is much faster than previous versions of Windows, but it's still not fast enough for hard-core games. For this reason, many old-school DOS developers swore off ever working in Windows. Pretty much the best you could do with graphics was rendering onto a bitmap that was then drawn into a window, a slow process. You used to have to give up a lot when writing a Windows application.

However, there are a lot of things that Windows can do that would be a nightmare to code in the old world of DOS. You can play sound effects using a single line of code (the `PlaySound` function), query the time stamp counter, use a robust TCP/IP network stack, get access to virtual memory, and the list goes on. Even though you have to take a few speed hits here and there, the advantages of Windows far outweigh the disadvantages.

I'll be using the Windows Vista environment with Visual C++ 2005 Express Edition (available for free from Microsoft) to write all of the applications for this book. We'll be developing our games in C++ using the Win32 API from the Windows platform SDL (Simple DirectMedia Layer). The Win32 API is a set of C functions that an application uses to make a Windows-compliant program. It abstracts away a lot of difficult operations like multitasking and protected memory, as well as providing interfaces to higher-level concepts. Supporting menus, dialog boxes, and multimedia have well-established library functions written for that specific task.

Windows is an extremely broad set of APIs. You can do just about anything, from playing videos to loading web pages. And for every task, there are a slew of different ways to accomplish it. There are some seriously

large books devoted just to the more rudimentary concepts of Windows programming. Subsequently, the discussion here will be limited to what is relevant to allow you to continue with the rest of the book. Instead of covering the tomes of knowledge required to set up dialogs with tree controls, print documents, and read/write keys in the registry, I'm going to deal with the simplest case: creating a window that can draw the world, passing input to the program, and having at least the beginnings of a pleasant relationship with the operating system. If you need any more info, there are many good resources out there on Windows programming. In particular I recommend *Programming Windows* by Charles Petzold.

Hungarian Notation

All of the variable names in Windows land use what is called Hungarian notation. The name came from its inventor, Charles Simonyi, who happened to be Hungarian.

Hungarian notation is the coding convention of just prefixing variables with a few letters to help identify their type. Hungarian notation makes it easier to read other people's code and easy to ensure the correct variables are supplied to functions in the right format. However, it can be really confusing to people who haven't seen it before.

Table 1.1 gives some of the more common prefixes used in most of the Windows and DirectX code that you'll see in this book.

Table 1.1: Some common Hungarian notation prefixes

b	bActive	Variable is a BOOL, a C precursor to the Boolean type found in C++. BOOLS can be true or false.
l	lPitch	Variable is a long integer.
dw	dwWidth	Variable is a DWORD, or unsigned long integer.
w	wSize	Variable is a WORD, or unsigned short integer.
sz	szWindowClass	Variable is a pointer to a string terminated by a zero (a standard C-style string).
p or lp	lpData	Variable is a pointer (lp is a carryover from the far pointers of the 16-bit days; it means long pointer). A pointer-pointer is prefixed by pp or lplp, and so on.
h	hInstance	Variable is a Windows handle.

General Windows Concepts

Notepad.exe is a good example of a simple Windows program. It allows basic text input, lets you do some basic text manipulation like searching and using the clipboard, and also lets you load, save, and print to a file. The program appears in Figure 1.1.

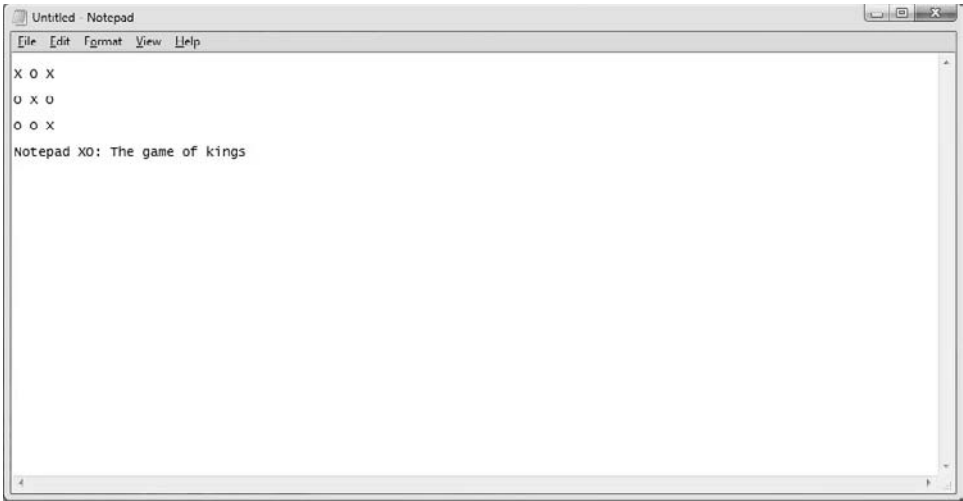


Figure 1.1: Notepad—a very basic window

The windows I show you how to create will be similar to this. A window such as this is partitioned into several distinct areas. Windows manages some of them, but the rest your application manages. The partitioning looks something like Figure 1.2.




Figure 1.2:
GUI window
components

The main parts are:

Title Bar	This area appears in most windows. It gives the name of the window and provides access to the system buttons that allow the user to close, minimize, or maximize an application. The only real control you have over the title bar is via a few flags in the window creation process. You can make it disappear, make it appear without the system icons, or make it thinner.
Menu Bar	The menu is one of the primary forms of interaction in a GUI program. It provides a list of commands the user can execute at any one time. Windows also controls this piece of the puzzle. You create the menu and define the commands, and Windows takes care of everything else.
Resize Bars	Resize bars allow the user to modify the size of the window on screen. You have the option of turning them off during window creation if you don't want to deal with the possibility of the window resizing.
Client Area	The client area is the meat of what you deal with. Windows essentially gives you a sandbox to play with in the client area. This is where you draw your scene. Windows can draw on parts of this region too. When there are scroll bars or toolbars in the application, they are intruding in the client area, so to speak.

Message Handling in Windows

Windows also have something called *focus*. Only one window can have focus at a time. The window that has the focus is the only window that the user can interact with. The rest appear with a different color title bar, and in the background. Because of this, only one application gets to know about the keyboard state.

How does your application know this? How does it know things like when it has focus or when the user clicks on it? How does it know where its window is located on the screen? Well, Windows “tells” the application when certain events happen. Also, you can tell other windows when things happen (in this way, different windows can communicate with each other).

Hold on though... How does Windows “tell” an application anything? This can be a foreign concept to people used to console programming, but it is paramount to the way Windows works. The trick is, Windows (and other applications) share information by sending packets of data called *messages* back and forth. A message is just a structure that contains the message itself, along with some parameters that contain information about the message.

The structure of a Windows message appears below:

```
typedef struct tagMSG {
    HWND hwnd;
    UINT message;
    WPARAM wParam;
    LPARAM lParam;
    DWORD time;
    POINT pt;
} MSG;
```

hwnd	Handle to the window that should receive the message.
message	The identifier of the message. For example, the application receives a msg object when the window is resized, and the message member variable is set to the constant WM_SIZE.
wParam	Information about the message; dependent on the type of message.
lParam	Additional information about the message.
time	Specifies when the message was posted.
pt	Mouse location when the message was posted.

Processing Messages

One of the most important concepts in Windows is that of a HWND, which is basically just an integer, representing a handle to a window. You can think of a HWND as like a barcode, uniquely identifying the window. Each window has its own unique HWND. When a Windows application wants to tell another window to do something, or wants to access a volatile system object like a file on disk, Windows doesn't actually let it change pointers or give it the opportunity to trounce on another application's memory space. Everything is done with handles to objects. It allows the application to send messages to the object, directing it to do things.



Note: The Win32 API predated object-oriented programming, and therefore doesn't take advantage of some newer programming concepts like exception handling. Most functions in Windows instead return an error code (called an HRESULT) that tells the caller how the function did. A non-negative HRESULT means the function succeeded.

If the function returns a negative number, an error occurred. The FAILED() macro returns true if an HRESULT is negative. There are a myriad of different types of errors that can result from a function; two examples are E_FAIL (generic error) and E_NOTIMPL (the function was not implemented).

An annoying side effect of having everything return an error code is that all the calls that retrieve information need to be passed a pointer of data to fill (instead of the more logical choice of just returning the requested data).

Messages can tell a window anything from “Paint yourself” to “You have lost focus” or “User double-clicked at location (x, y).” Each time a message is sent to a window, it is added to a message queue inside Windows. Each window has its own associated local message queue. A message queue ensures that each message gets processed in the order it gets received, even if it arrives while the application is busy processing other messages. In fact, when most Windows applications get stuck in an infinite loop or otherwise stop working, you’ll notice because they’ll stop processing messages, and therefore don’t redraw or process input.

So how does an application process messages? Windows defines a function that all programs must implement called the *window procedure* (or `WndProc` for short). When you create a window, you give Windows your `WndProc()` function in the form of a function pointer. Then, when messages are processed, they are passed as parameters to the function, and the `WndProc()` deals with them. So, for example, when the `WndProc()` function gets passed a message saying “Paint yourself” that is the signal for the window to redraw itself. You would therefore add code to the `WndProc()` to handle the redrawing of your window.

When a Win32 application sends a message, Windows examines the window handle provided, using it to find out where to send the message. The message ID describes the message being sent, and the parameters to the ID are contained in the two other fields in a message, `wParam` and `lParam`. Back in the 16-bit days, `wParam` was a 16-bit (word sized) integer and `lParam` was a 32-bit (long sized) integer, but with Win32 they’re both 32 bits long. The messages wait in a queue until the application receives them.

The window procedure should return 0 for any message it processes. All messages it doesn’t process should be passed to the default message procedure, `DefWindowProc()`. Windows can start behaving erratically if `DefWindowProc()` doesn’t see all of your non-processed messages.

Hello World—Windows Style

To help explain these ideas, let me show you a minimalist Win32 program and analyze what’s going on. This code was modified from the default “Hello, World” code that Visual C++ 2005 Express Edition will automatically generate for you.

```

/*****
*           Advanced 3D Game Programming with DirectX 10.0
*   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *   *
*   Title: HelloWorld.cpp
*   Desc: Simple windows application
*   copyright (c) 2007 by Peter Walsh, Wordware
*****/

#include "stdafx.h"
```

```
#include "Hello World.h"

#define MAX_LOADSTRING 100

// Global Variables:
HINSTANCE hInst;                // current instance
TCHAR szTitle[MAX_LOADSTRING]; // The title bar text
TCHAR szWindowClass[MAX_LOADSTRING]; // the main window class name

// Forward declarations of functions included in this code module:
ATOM                MyRegisterClass(HINSTANCE hInstance);
BOOL                InitInstance(HINSTANCE, int);
LRESULT CALLBACK    WndProc(HWND, UINT, WPARAM, LPARAM);
INT_PTR CALLBACK    About(HWND, UINT, WPARAM, LPARAM);

int APIENTRY _tWinMain(HINSTANCE hInstance,
                      HINSTANCE hPrevInstance,
                      LPTSTR    lpCmdLine,
                      int       nCmdShow)
{
    UNREFERENCED_PARAMETER(hPrevInstance);
    UNREFERENCED_PARAMETER(lpCmdLine);

    // TODO: Place code here.
    MSG msg;
    HACCEL hAccelTable;

    // Initialize global strings
    LoadString(hInstance, IDS_APP_TITLE, szTitle, MAX_LOADSTRING);
    LoadString(hInstance, IDC_HELLOWORLD, szWindowClass, MAX_LOADSTRING);
    MyRegisterClass(hInstance);

    // Perform application initialization:
    if (!InitInstance (hInstance, nCmdShow))
    {
        return FALSE;
    }

    hAccelTable =
        LoadAccelerators(hInstance, MAKEINTRESOURCE(IDC_HELLOWORLD));

    // Main message loop:
    while (GetMessage(&msg, NULL, 0, 0))
    {
        if (!TranslateAccelerator(msg.hwnd, hAccelTable, &msg))
        {
            TranslateMessage(&msg);
            DispatchMessage(&msg);
        }
    }

    return (int) msg.wParam;
}
```

```

//
// FUNCTION: MyRegisterClass()
//
// PURPOSE: Registers the window class.
//
// COMMENTS:
//
// This function and its usage are only necessary if you want this code
// to be compatible with Win32 systems prior to the 'RegisterClassEx'
// function that was added to Windows 95. It is important to call this
// function so that the application will get 'well formed' small icons
// associated with it.
//
ATOM MyRegisterClass(HINSTANCE hInstance)
{
    WNDCLASSEX wcex;

    wcex.cbSize = sizeof(WNDCLASSEX);
    wcex.style      = CS_HREDRAW | CS_VREDRAW;
    wcex.lpfnWndProc = WndProc;
    wcex.cbClsExtra = 0;
    wcex.cbWndExtra = 0;
    wcex.hInstance = hInstance;
    wcex.hIcon      = LoadIcon(hInstance, MAKEINTRESOURCE(IDI_HELLOWORLD));
    wcex.hCursor    = LoadCursor(NULL, IDC_ARROW);
    wcex.hbrBackground = (HBRUSH)(COLOR_WINDOW+1);
    wcex.lpszMenuName = MAKEINTRESOURCE(IDC_HELLOWORLD);
    wcex.lpszClassName = szWindowClass;
    wcex.hIconSm     = LoadIcon(wcex.hInstance, MAKEINTRESOURCE(IDI_SMALL));

    return RegisterClassEx(&wcex);
}

//
// FUNCTION: InitInstance(HINSTANCE, int)
//
// PURPOSE: Saves instance handle and creates main window
//
// COMMENTS:
//
// In this function, we save the instance handle in a global variable
// and create and display the main program window.
//
BOOL InitInstance(HINSTANCE hInstance, int nCmdShow)
{
    HWND hWnd;

    hInst = hInstance; // Store instance handle in our global variable

    hWnd = CreateWindow(szWindowClass, szTitle, WS_OVERLAPPEDWINDOW,
        CW_USEDEFAULT, 0, CW_USEDEFAULT, 0, NULL, NULL, hInstance, NULL);

    if (!hWnd)
    {
        return FALSE;
    }

```



```

    }

    ShowWindow(hWnd, nCmdShow);
    UpdateWindow(hWnd);

    return TRUE;
}

//
// FUNCTION: WndProc(HWND, UINT, WPARAM, LPARAM)
//
// PURPOSE: Processes messages for the main window.
//
// WM_COMMAND   - process the application menu
// WM_PAINT     - paint the main window
// WM_DESTROY   - post a quit message and return
//
//
LRESULT CALLBACK WndProc(HWND hWnd, UINT message, WPARAM wParam, LPARAM lParam)
{
    int wmId, wmEvent;
    PAINTSTRUCT ps;
    HDC hdc;

    TCHAR strOutput[] = L"Hello World!";

    switch (message)
    {
    case WM_COMMAND:
        wmId    = LOWORD(wParam);
        wmEvent = HIWORD(wParam);
        // Parse the menu selections:
        switch (wmId)
        {
        case IDM_ABOUT:
            DialogBox(hInst, MAKEINTRESOURCE(IDD_ABOUTBOX), hWnd, About);
            break;
        case IDM_EXIT:
            DestroyWindow(hWnd);
            break;
        default:
            return DefWindowProc(hWnd, message, wParam, lParam);
        }
        break;
    case WM_PAINT:
        hdc = BeginPaint(hWnd, &ps);

        RECT rt;
        GetClientRect(hWnd, &rt);
        DrawText(hdc, strOutput, (int)wcslen(strOutput), &rt,
            DT_CENTER | DT_VCENTER | DT_SINGLELINE);

        EndPaint(hWnd, &ps);
        break;
    }
}

```

```

        case WM_DESTROY:
            PostQuitMessage(0);
            break;
        default:
            return DefWindowProc(hWnd, message, wParam, lParam);
    }
    return 0;
}

// Message handler for about box.
INT_PTR CALLBACK About(HWND hDlg, UINT message, WPARAM wParam, LPARAM lParam)
{
    UNREFERENCED_PARAMETER(lParam);
    switch (message)
    {
        case WM_INITDIALOG:
            return (INT_PTR)TRUE;

        case WM_COMMAND:
            if (LOWORD(wParam) == IDOK || LOWORD(wParam) == IDCANCEL)
            {
                EndDialog(hDlg, LOWORD(wParam));
                return (INT_PTR)TRUE;
            }
            break;
    }
    return (INT_PTR)FALSE;
}

```

It's easy to get worried when you think about the fact that this is one of the simplest Windows programs you can write, and it's still over 200 lines long. The good thing is that the code above is more or less common to all Windows programs. Most Windows programmers don't remember the exact order everything goes in; they just copy the working Windows initialization code from a previous application and use it like it is their own.

Explaining the Code

Every C/C++ program has its entry point in `main()`, where it is passed control from the operating system. In Windows, things work a little differently. There is some code that the Win32 API runs first, before letting your code run. The actual stub for `main()` lies deep within the Win32 DLLs. However, this application starts at a different point: with a function called `WinMain()`. Windows does its setup work when your application is first run, and then calls `WinMain()`. This is why when you debug a Windows app “WinMain” doesn't appear at the bottom of the call stack; the internal DLL functions that called it are. `WinMain()` is passed the following parameters (in order):

- The instance of the application (another handle, this one representing an instantiation of a running executable). Each process has a separate instance handle that uniquely identifies the process to Windows. This is

different from a window handle, as each application can have many windows under its control. You need to hold on to this instance, as certain Windows API calls need to know what instance is calling them. Think of an instance as just a copy, or even as an image, of the executable in memory. Each executable has a handle so that Windows can tell them apart, manage them, and so on.

- An `HINSTANCE` of another copy of your application currently running. Back in the days before machines had much memory, Windows would have multiple instances of a running program share memory. These days each process is run in its own separate memory space, so this parameter is always `NULL`. It remains this way so that legacy Windows applications still work.
- A pointer to the command line string. When the user drags a file onto an executable in Explorer (not a running copy of the program), Windows runs the program with the first parameter of the command line being the path and filename of the file dragged onto it.
- A set of flags describing how the window should initially be drawn (such as full-screen, minimized, etc.).

The conceptual flow of the function is to do the following:

```
WinMain
    Register the application class with Windows
    Create the main window
    while( not been told to exit )
        Process any messages that Windows has sent
```

`MyRegisterClass()` takes the application instance and tells Windows about the application (registering it, in essence). `InitInstance()` creates the primary window on the screen and starts it drawing. Then the code enters a while loop that remains in execution until the application quits. The function `GetMessage()` looks at the message queue. It always returns 1 unless there is a specific system message in the queue: This is the “Quit now” message and has the message ID `WM_QUIT`. If there is a message in the queue, `GetMessage()` will remove it and fill it into the message structure, which is the “msg” variable above. Inside the while loop, you first take the message and translate it using a function called `TranslateMessage()`.

This is a convenience function. When you receive a message saying a key has been pressed or released, you get the specific key as a virtual key code. The actual values for the IDs are arbitrary, but the namespace is what you care about: When the letter “a” is pressed, one of the message parameters is equivalent to the `#define VK_A`. Since that nomenclature is a pain to deal with if you’re doing something like text input, `TranslateMessage()` does some housekeeping and converts the parameter from `VK_A` to `(char)‘a’`. This makes processing regular text input much easier. Keys without clear ASCII equivalents, such as Page Up and Left Arrow, keep their

virtual key code values (VK_PRIOR and VK_LEFT, respectively). All other messages go through the function and come out unchanged.

The second function, `DispatchMessage()`, is the one that actually processes it. Internally, it looks up which function was registered to process messages (in `MyRegisterClass()`) and sends the message to that function. You'll notice that the code never actually calls the window procedure. That's because Windows does it for you when you ask it to with the `DispatchMessage()` function.

Think of this while loop as the central nervous system for any Windows program. It constantly grabs messages off the queue and processes them as fast as it can. It's so universal it actually has a special name: the *message pump*.

Registering the Application

`MyRegisterClass()` fills a structure that contains the info Windows needs to know about your application before it can create a window, and passes it to the Win32 API. This is where you tell Windows what size to make the icon for the application that appears in the taskbar (`hIcon`, the large version, and `hIconSm`, the smaller version). You can also give it the name of the menu bar if you ever decide to use one. (For now there is none, so it's set to 0.) You need to tell Windows what the application instance is (the one received in the `WinMain()`); this is the `hInstance` parameter. You also tell it which function to call when it processes messages; this is the `lpfnWndProc` parameter. The window class has a name as well, `lpstrClassName`, that is used to reference the class later in the `CreateWindow()` function.



Warning: A window class is completely different from a C++ class. Windows predated the popularity of the C++ language, and therefore some of the names have a tendency to clash.

Initializing the Window

`InitInstance()` creates the window and starts the drawing process. The window is created with a call to `CreateWindow()`, which has the following prototype:

```
HWND CreateWindow(
    LPCTSTR lpClassName,
    LPCTSTR lpWindowName,
    DWORD dwStyle,
    int x,
    int y,
    int nWidth,
    int nHeight,
    HWND hWndParent,
    HMENU hMenu,
```

```
HANDLE hInstance,
LPVOID lpParam
);
```

lpClassName	A string giving the class name for the window class that was registered with RegisterClass(). This defines the basic style of the window, along with which WndProc() will be handling the messages (you can create more than one window class per application).
lpWindowName	The title of the window. This will appear in the title bar of the window and in the taskbar.
dwStyle	A set of flags describing the style for the window (such as having thin borders, being unresizable, and so on). For these discussions, windowed applications will all use WS_OVERLAPPEDWINDOW (this is the standard-looking window, with a resizable edge, a system menu, a title bar, etc.). However, full-screen applications will use the WS_POPUP style (no Windows features at all, not even a border; it's just a client rectangle).
x, y	The x and y location, relative to the top-left corner of the monitor (x increasing right, y increasing down), where the window should be placed.
nWidth, nHeight	The width and height of the window.
hWndParent	A window can have child windows (imagine a paint program like Paint Shop Pro, where each image file exists in its own window). If this is the case and you are creating a child window, pass the HWND of the parent window here.
hMenu	If an application has a menu (yours doesn't), pass the handle to it here.
hInstance	This is the instance of the application that was received in WinMain().
lpParam	Pointer to extra window creation data you can provide in more advanced situations (for now, just pass in NULL).

The width and height of the window that you pass to this function is the width and height for the *entire* window, not just the client area. If you want the client area to be a specific size, say 640 by 480 pixels, you need to adjust the width and height passed to account for the pixels needed for the title bar, resize bars, etc. You can do this with a function called AdjustWindowRect() (discussed later in the chapter). You pass a rectangle structure filled with the desired client rectangle, and the function adjusts the rectangle to reflect the size of the window that will contain the client rectangle, based on the style you pass it (hopefully the same style passed to CreateWindow()). A window created with WS_POPUP has no extra Windows UI features, so the window will go through unchanged. WS_OVERLAPPEDWINDOW has to add space on each side for the resize bar and on the top for the title bar.

If `CreateWindow()` fails (this will happen if there are too many windows or if it receives bad inputs, such as an `hInstance` different from the one provided in `MyRegisterClass()`), you shouldn't try processing any messages for the window (since there is no window!) so return `false`. This is handled in `WinMain()` by exiting the application before entering the message pump. Normally, before exiting, you'd bring up some sort of pop-up alerting the user to the error, instead of just silently quitting. Otherwise, call `ShowWindow()`, which sets the show state of the window just created (the show state was passed to as the last formal parameter in `WinMain()`), and `UpdateWindow()`, which sends a paint message to the window so it can draw itself.



Warning: `CreateWindow()` calls the `WndProc()` function several times before it exits! This can sometimes cause headaches in getting certain Windows programs to work.

Before the function returns and you get the window handle back, `WM_CREATE`, `WM_MOVE`, `WM_SIZE`, and `WM_PAINT` (among others) are sent to the program through the `WndProc()`.

If you're using any components that need the `HWND` of a program to perform work (a good example is a DirectX window, whose surface must be resized whenever it gets a `WM_SIZE` message), you need to tread very carefully so that you don't try to resize the surface before it has been initialized. One way to handle this is to record your window's `HWND` inside `WM_CREATE`, since one of the parameters that gets passed to the `WndProc()` is the window handle to receive the message.

You may wonder how you would alert the user when an event such as an error occurs. In these cases, you have to create dialogs that present information to the user such as why the program failed. Complex dialogs with buttons and edit boxes and whatnot are generally not needed for creating games (usually you create your own interface inside the game); however, there are some basic dialogs that Windows can automatically create, such as the pop-up window presented when exiting any sort of document editing software that says "Save SomeFile.x before exiting?" and has two buttons labeled "Yes" and "No."

The function you use to automate the dialog creation process is called `MessageBox()`. It is one of the most versatile and useful Windows functions. Take a look at its syntax:

```
int MessageBox(
    HWND hWnd,
    LPCTSTR lpText,
    LPCTSTR lpCaption,
    UINT uType
);
```

hWnd	Handle to the owner of the window (this is generally the application's window handle).
lpText	Text for the inside of the message box.
lpCaption	Title of the message box.
uType	A set of flags describing the behavior of the message box. The flags are described in Table 1.2.

The function displays the dialog on the desktop and does not return until the box is closed.

Table 1.2: A set of the common flags used with MessageBox()

MB_OK	The message box has just one button marked OK. This is the default behavior.
MB_ABORTRETRYIGNORE	Three buttons appear—Abort, Retry, and Ignore.
MB_OKCANCEL	Two buttons appear—OK and Cancel.
MB_RETRYCANCEL	Two buttons appear—Retry and Cancel.
MB_YESNO	Two buttons appear—Yes and No.
MB_YESNOCANCEL	Three buttons appear—Yes, No, and Cancel.
MB_ICONEXCLAMATION, MB_ICONWARNING	An exclamation mark icon is displayed.
MB_ICONINFORMATION, MB_ICONASTERISK	An information icon (a lowercase i inscribed in a circle) is displayed.
MB_ICONQUESTION	A question mark icon is displayed.
MB_ICONSTOP, MB_ICONERROR, MB_ICONHAND	A stop sign icon is displayed.

The return value of MessageBox() depends on which button was pressed. Table 1.3 gives the possible return values. Note that this is one of the rare Windows functions that does not return an HRESULT.

Table 1.3: Return values for MessageBox()

IDABORT	The Abort button was pressed.
IDCANCEL	The Cancel button was pressed.
IDIGNORE	The Ignore button was pressed.
IDNO	The No button was pressed.
IDOK	The OK button was pressed.
IDRETRY	The Retry button was pressed.
IDYES	The Yes button was pressed.

WndProc—The Message Pump

WndProc() is the window procedure. This is where everything happens in a Windows application. Since this application is so simple, it will only process two messages (more complex Windows programs will need to process dozens upon dozens of messages). The two messages that most Win32 applications handle are WM_PAINT (sent when Windows would like the window to be redrawn) and WM_DESTROY (sent when the window is being destroyed). An important thing to note is that any message you don't process in the switch statement goes into DefWindowProc(), which defines the default behavior for every Windows message. Anything not processed needs to go into DefWindowProc() for the application to behave correctly.

System messages, such as the message received when the window is being created and destroyed, are sent by Windows internally. You can post messages to your own application (and other applications) with two functions: PostMessage() and SendMessage(). PostMessage() adds the message to the application's message queue to be processed in the message pump. SendMessage() actually calls the WndProc() with the given message itself.

One important point to remember when you're doing Windows programming is that you don't need to memorize any of this. Very few, if any, people know all the parameters to each and every one of the Windows functions; usually it's looked up in MSDN, copied from another place, or filled in for you by a project wizard. So don't worry if you're getting overloaded with new information. One of the most useful investments I ever made was to purchase a second monitor. That way I can program on my main screen with MSDN up on the other, which means I don't have to keep task switching between applications.

One thing you might notice is that for a program that just says "Hello, World!" there sure is a lot of code. Most of it exists in all Windows programs. All applications need to register themselves, they all need to create a window if they want one, and they all need a window procedure. While it may be a bit on the long side, the program does a lot. You can resize it, move it around the screen, have it become occluded by other windows, minimize, maximize, and so on. Windows users automatically take this functionality for granted, but there is a lot of code taking place out of sight.

Manipulating Window Geometry

Since for now the application's use of Windows is so restricted, you only need to concern yourself with two basic Windows structures that are used in geometry functions: POINT and RECT.

In Windows, there are two coordinate spaces. One is the client area coordinate space. The origin (0,0) is the top-left corner of the window (known as *client space*). Coordinates relative to the client area don't need to change when the window is moved around the screen. The other

coordinate space is the desktop coordinate space. This space is absolute, and the origin is the top-left corner of the screen (also known as *screen space*).

Windows uses the POINT structure to represent 2D coordinates. It has two long integers, one for the horizontal component and one for the vertical:

```
typedef struct tagPOINT {
    LONG x;
    LONG y;
} POINT;
```

Since all windows are rectangular, Windows has a structure to represent a rectangle. You'll notice that essentially the structure is two points end to end, the first describing the top-left corner of the rectangle, the other describing the bottom-right corner.

```
typedef struct _RECT {
    LONG left;
    LONG top;
    LONG right;
    LONG bottom;
} RECT;
```

left	Left side of the window.
top	Top of the window.
right	Right side of the window (width is right-left).
bottom	Bottom side of the window (height is bottom-top).

To get the client rectangle of a window you can use the function `GetClientRect()`. The left and top members are always zero, and the right and bottom give you the width and height of the window.

```
BOOL GetClientRect(
    HWND hWnd,
    LPRECT lpRect
);
```

hWnd	Handle to the window you want information about.
lpRect	Pointer to a RECT structure you would like filled with the client rectangle.

Once you have the client rectangle, you often need to know what those points are relative to the desktop coordinate space. `ClientToScreen()`, which has the following syntax, provides this functionality:

```

BOOL ClientToScreen(
    HWND hWnd,
    LPPOINT lpPoint
);

```

<code>hWnd</code>	Handle to the window in which the client point is defined.
<code>lpPoint</code>	Pointer to the client point; this point is changed to screen space.

To change the rectangle you get through `GetClientRect()` to screen space, you can use the `ClientToScreen()` function on the bottom and right members of a rectangle. Slightly inelegant, but it works.

One thing that can mess up window construction is determining the width and height of the window. You could say you want a client rectangle that is 800 pixels by 600 pixels (or some other resolution), but you call `CreateWindow()` giving the dimensions of the *whole* window, including any resize, title bar, and menu bars. Luckily, you can convert a rectangle representing the client rectangle to one representing the window dimensions using `AdjustWindowRect()`. It pushes all of the coordinates out to accommodate the window style `dwStyle`, which should be the same one used in `CreateWindow()` for it to work correctly. For non-pop-up windows, this will make the top and left coordinates negative.

```

BOOL AdjustWindowRect(
    LPRECT lpRect,
    DWORD dwStyle,
    BOOL bMenu
);

```

<code>lpRect</code>	Pointer to the <code>RECT</code> structure to be adjusted.
<code>dwStyle</code>	Style of the intended window, this defines how much to adjust each coordinate. For example, <code>WS_POPUP</code> style windows aren't adjusted at all.
<code>bMenu</code>	Boolean that is true if the window will have a menu. If, like in this case, there is no menu, you can just pass false for this parameter.

Windows has a full-featured graphics library that performs operations on a handle to a graphics device. The package is called the GDI, or graphical device interface. It allows users to draw, among other things, lines, ellipses, bitmaps, and text. (I'll show you its text painting ability in a later chapter.) The sample program uses it to draw the "Hello, World!" text on the screen. I'll show you more of the GDI's functions later in the book.

Important Window Messages

Most of the code in this book uses Windows as a jumping-off point—a way to put a window up on the screen that allows you to draw in it. I’ll only be showing you a small subset of the massive list of window messages in Windows, which is a good thing since they can get pretty mind-numbing after a while. Table 1.4 describes the important messages and their parameters.

Table 1.4: Some important window messages

WM_CREATE	Sent to the application when Windows has completed creating its window but before it is drawn. This is the first time the application will see what the HWND of its window is.
WM_PAINT	<p>Sent to the application when Windows wants the window to draw itself.</p> <p>Parameters:</p> <p>(HDC) wParam</p> <p>A handle to the device context for the window that you can draw in.</p>
WM_ERASEBKGD	<p>Called when the background of a client window should be erased. If you process this message instead of passing it to DefWindowProc(), Windows will let you erase the background of the window (later, I’ll show you why this can be a good thing).</p> <p>Parameters:</p> <p>(HDC) wParam</p> <p>A handle to the device context to draw in.</p>
WM_DESTROY	Sent when the window is being destroyed.
WM_CLOSE	Sent when the window is being asked to close itself. This is where you can, for example, ask for confirmation before closing the window.
WM_SIZE	<p>Sent when the window is resized. When the window is resized, the top-left location stays the same (so when you resize from the top left, both a WM_MOVE and a WM_SIZE message are sent).</p> <p>Parameters:</p> <p>wParam</p> <p>Resizing flag. There are other flags, but the juicy one is SIZE_MINIMIZED; it’s sent when the window is minimized.</p> <p>LOWORD(lParam)</p> <p>New width of the client area (not total window).</p> <p>HIWORD(lParam)</p> <p>New height of the client area (not total window).</p>
WM_MOVE	<p>Sent when the window is moved.</p> <p>Parameters:</p> <p>(int)(short) LOWORD(lParam)</p> <p>New upper-left x-coordinate of client area.</p> <p>(int)(short) HIWORD(lParam)</p> <p>New upper-left y-coordinate of client area.</p>

WM_QUIT	<p>Last message the application gets; upon its receipt the application exits. You never process this message, as it actually never gets through to WndProc(). Instead, it is caught in the message pump in WinMain() and causes that loop to drop out and the application to subsequently exit.</p>
WM_KEYDOWN	<p>Received every time a key is pressed. Also received after a specified time for auto-repeats.</p> <p>Parameters:</p> <p>(int) wParam</p> <p>The virtual key code for the pressed key. If you call TranslateMessage() on the message before processing it, if it is a key with an ASCII code equivalent (letters, numbers, punctuation marks) it will be equivalent to the actual ASCII character.</p>
WM_KEYUP	<p>Received when a key is released.</p> <p>Parameters:</p> <p>(int) wParam</p> <p>The virtual key code for the released key.</p>
WM_MOUSEMOVE	<p>MouseMove is a message that is received almost constantly. Each time the mouse moves in the client area of the window, the application gets notified of the new location of the mouse cursor relative to the origin of the client area.</p> <p>Parameters:</p> <p>LOWORD(1Param)</p> <p>The x-location of the mouse, relative to the upper-left corner of the client area.</p> <p>HIWORD(1Param)</p> <p>The y-location of the mouse, relative to the upper-left corner of the client area.</p> <p>wParam</p> <p>Key flags. This helps you tell the keyboard state for special clicks (such as Alt-left click, for example). Test the key flags to see if certain flags are set. The flags are:</p> <ul style="list-style-type: none"> • MK_CONTROL: Indicates the Control key is down. • MK_LBUTTON: Indicates the left mouse button is down. • MK_MBUTTON: Indicates the middle mouse button is down. • MK_RBUTTON: Indicates the right mouse button is down. • MK_SHIFT: Indicates the Shift key is down.
WM_LBUTTONDOWN	<p>This message is received when the user presses the left mouse button in the client area. You only receive one message when the button is pressed, as opposed to receiving them continually while the button is down.</p> <p>Parameters:</p> <p>LOWORD(1Param)</p> <p>The x-location of the mouse, relative to the upper-left corner of the client area.</p> <p>HIWORD(1Param)</p> <p>The y-location of the mouse, relative to the upper-left corner of the client area.</p>

wParam

Key flags. This helps you tell the keyboard state for special clicks (such as Alt-left click, for example). Test the key flags to see if certain flags are set. The flags are:

- **MK_CONTROL**: Indicates the Control key is down.
- **MK_MBUTTONDOWN**: Indicates the middle mouse button is down.
- **MK_RBUTTONDOWN**: Indicates the right mouse button is down.
- **MK_SHIFT**: Indicates the Shift key is down.

WM_MBUTTONDOWN You receive this message when the user presses the middle mouse button in the client area. You only receive one message when the button is pressed, as opposed to receiving them continually while the button is down.

Parameters:

LOWORD(1Param)

The x-location of the mouse, relative to the upper-left corner of the client area.

HIWORD(1Param)

The y-location of the mouse, relative to the upper-left corner of the client area.

wParam

Key flags. This helps you tell the keyboard state for special clicks (such as Alt-left click, for example). Test the key flags to see if certain flags are set. The flags are:

- **MK_CONTROL**: If set, Control key is down.
- **MK_LBUTTONDOWN**: If set, left mouse button is down.
- **MK_RBUTTONDOWN**: If set, right mouse button is down.
- **MK_SHIFT**: If set, Shift key is down.

WM_RBUTTONDOWN You receive this message when the user presses the right mouse button in the client area. You only receive one message when the button is pressed, as opposed to receiving them continually while the button is down.

Parameters:

LOWORD(1Param)

The x-location of the mouse, relative to the upper-left corner of the client area.

HIWORD(1Param)

The y-location of the mouse, relative to the upper-left corner of the client area.

wParam

Key flags. This helps you tell the keyboard state for special clicks (such as Alt-left click, for example). Test the key flags to see if certain flags are set. The flags are:

- **MK_CONTROL**: Indicates the Control key is down.
 - **MK_LBUTTONDOWN**: Indicates the left mouse button is down.
 - **MK_MBUTTONDOWN**: Indicates the middle mouse button is down.
 - **MK_SHIFT**: Indicates the Shift key is down.
-

WM_LBUTTONDOWN	Received when the user releases the left mouse button in the client area. Parameters: The parameters are the same as for WM_LBUTTONDOWN.
WM_MBUTTONDOWN	Received when the user releases the middle mouse button in the client area. Parameters: The parameters are the same as for WM_MBUTTONDOWN.
WM_RBUTTONDOWN	Received when the user releases the right mouse button in the client area. Parameters: The parameters are the same as for WM_RBUTTONDOWN.
WM_MOUSEWHEEL	Most new mice come equipped with a z-axis control, in the form of a wheel. It can be spun forward and backward and clicked. If it is clicked, it generally sends the middle mouse button messages. However, if it is spun forward or backward, the following parameters are passed. Parameters: (short) HIWORD(wParam) The amount the wheel has spun since the last message. A positive value means the wheel was spun forward (away from the user). A negative value means the wheel was spun backward (toward the user). (short) LOWORD(lParam) The x-location of the mouse, relative to the upper-left corner of the client area. (short) HIWORD(lParam) The y-location of the mouse, relative to the upper-left corner of the client area. LOWORD(wParam) Key flags. This helps you tell the keyboard state for special clicks (such as Alt-left click, for example). Test the key flags to see if certain flags are set. The flags are: <ul style="list-style-type: none"> • MK_CONTROL: Indicates the Control key is down. • MK_LBUTTON: Indicates the left mouse button is down. • MK_MBUTTON: Indicates the middle mouse button is down. • MK_RBUTTON: Indicates the right mouse button is down. • MK_SHIFT: Indicates the Shift key is down.

Class Encapsulation

So, now that you can create a window, I'm going to show you how to design a framework that will sit beneath the Direct3D and other game code and simplify the programming tasks needed in all of the other applications you'll be building in the book.

As a first step, let's look at a list of benefits that could be gained from the encapsulation. In no particular order, it would be good if the application had:

- The ability to control and reimplement the construction and destruction of the application object.
- The ability to automatically create standard system objects (right now just the application window, but later on Direct3D, DirectInput, and so on), and facilities to create your own.
- The ability to add objects that can listen to the stream of window messages arriving to the application and add customized ways to handle them.
- A simple main loop that runs repeatedly until the application exits.

The way I'll do this is with two classes. One of them will abstract the Windows code that needs to be run; it is called `cWindow`. It will be used by a bigger class that is responsible for actually running the application. This class is called `cApplication`. Each new application that you create (with a couple of exceptions) will be subclassed from `cApplication`.

Whenever something goes wrong during the execution that requires the application to exit, the infrastructure is designed so that an error can be thrown. The entire application is wrapped around a try/catch block, so any errors are caught in `WinMain()`, and the application is shut down. A text message describing the error can be passed in the thrown exception, and the string is popped up using a message box before the application exits.

I chose to do this because it can be easier than the alternative of having every single function return an error code, and having each function check the result of each function it calls. Exceptions get thrown so rarely that the added complexity that error codes add seems pretty unnecessary really. With exception handling, the code is nice and clean. The error that almost all of the code in this book throws is called `cGameError`.

```
class cGameError
{
    string m_errorText;
public:
    cGameError( char *errorText )
    {
        DP1("***\n*** [ERROR] cGameError thrown! text: [%s]\n***\n",
            errorText );
        m_errorText = string( errorText );
    }

    const char *GetText()
    {
        return m_errorText.c_str();
    }
};

enum eResult
{
```

```

    resAllGood    = 0,    // function passed with flying colors
    resFalse      = 1,    // function worked and returns 'false'
    resFailed     = -1,   // function failed miserably
    resNotImpl    = -2,   // function has not been implemented
    resForceDWord = 0x7FFFFFFF
};

```

The window abstraction, `cWindow`, is fairly straightforward. `MyRegisterClass()` is replaced with `cWindow::RegisterClass()`, `MyInitInstance()` is now `cWindow::InitInstance()`, and `WndProc()` is now a static function `cWindow::WndProc()`. The function is static because non-static class functions have a hidden first variable passed in (the *this* pointer) that is not compatible with the `WndProc()` function declaration.

The message pump is encapsulated in two functions. `HasMessages()` checks the queue and sees if there are any messages waiting to be processed, returning true if there are any. `Pump()` processes a single message, sending it off to `WndProc()` using `TranslateMessage()/DispatchMessage()`. When `Pump()` receives the `WM_QUIT` message, which again is a notification from Windows that the application should exit, it returns `resFalse`.

Special care needs to be taken to handle thrown exceptions that happen during the window procedure. Between the execution of `DispatchMessage()` and `WndProc()`, the call stack meanders into kernel DLL functions. If a thrown exception flies into them, bad stuff happens (anything from your program crashing to your machine crashing). To handle this, any and all exceptions are caught in the `WndProc()` and saved in a temporary variable. When `Pump()` finishes pumping a message, it checks the temporary variable to see if an error was thrown. If there is an error waiting, `Pump()` rethrows the error and it rises up to `WinMain()`.

```

class cWindow
{
protected:

    int m_width, m_height;
    HWND m_hWnd;
    std::string m_name;
    bool m_bActive;
    static cWindow *m_pGlobalWindow;

public:

    cWindow(
        int width,
        int height,
        const char *name = "Default window name" );
    ~cWindow();

    virtual LRESULT WndProc(
        HWND hWnd,
        UINT uMsg,

```



```

        WPARAM wParam,
        LPARAM lParam );

    virtual void RegisterClass( WNDCLASSEX *pWc = NULL );
    virtual void InitInstance();

    HWND GetHwnd();
    bool IsActive();
    bool HasMessages();
    eResult Pump();
    static cWindow *GetMainWindow();
};

inline cWindow *MainWindow();
    
```

m_width, m_height	Width and height of the client rectangle of the window. This is different from the width and height of the actual window.
m_hwnd	Handle to the window. Use the public function GetHwnd to access it outside the class.
m_name	The name of the window used to construct the window class and window.
m_bActive	Boolean value; true if the window is active (a window is active if it is currently in the foreground).
m_pGlobalWindow	Static variable that points to the single instantiation of a cWindow class for an application. Initially set to NULL.
cWindow()	Constructs a window object. You can only create one instance of this object; this is verified by setting the m_pGlobalWindow object.
~cWindow()	The destructor destroys the window and sets the global window variable to NULL so that it cannot be accessed any longer.
WndProc()	Window procedure for the class. Called by a hidden function inside Window.cpp.
RegisterClass()	Virtual function that registers the window class. This function can be overloaded in child classes to add functionality, such as a menu or different WndProc().
InitInstance()	Virtual function that creates the window. This function can be overloaded in child classes to add functionality, such as changing the window style.
GetHwnd()	Returns the window handle for this window.
IsActive()	Returns true if the application is active and in the foreground.
HasMessages()	True if the window has any messages in its message queue waiting to be processed. Uses PeekMessage() with PM_NOREMOVE.

Pump()	Pumps the first message off the queue and dispatches it to the WndProc(). Returns resAllGood, unless the message gotten off the queue was WM_QUIT, in which case it returns resFalse.
GetMainWindow()	Public function; used by the global function MainWindow() to gain access to the only window object.
MainWindow()	Global function that returns the single instance of the cWindow class for this program. Any piece of code can use this to query information about the window. For example, any code can get the hWnd for the window by calling MainWindow()->GetHWnd().

Finally, there is the cApplication class. Child classes will generally only reimplement SceneInit() and DoFrame(). However, other functions can be reimplemented if added functionality, like the construction of extra system objects, is needed. The game presented in the final chapter will use several other system objects that it will need to construct.

```
class cApplication
{
protected:

    string m_title;
    int m_width;
    int m_height;

    bool m_bActive;

    static cApplication *m_pGlobalApp;

    virtual void InitPrimaryWindow();
    virtual void InitGraphics();
    virtual void InitInput();
    virtual void InitSound();
    virtual void InitExtraSubsystems();

public:

    cApplication();
    virtual ~cApplication();

    virtual void Init();

    virtual void Run();
    virtual void DoFrame( float timeDelta );
    virtual void DoIdleFrame( float timeDelta );
    virtual void ParseCmdLine( char *cmdLine );

    virtual void SceneInit();
    virtual void SceneEnd();

    void Pause();
```

```
void UnPause();

static cApplication *GetApplication();

static void KillApplication();
};

inline cApplication *Application();

HINSTANCE AppInstance();

cApplication *CreateApplication();
```

m_title	Title for the application. Sent to the cWindow when it is constructed.
m_width, m_height	Width and height of the client area of the desired window.
m_bActive	True if the application is active and running. When the application is inactive, input isn't received and the idle frame function is called.
m_pGlobalApp	Static pointer to the single global instance of the application.
InitPrimaryWindow()	Virtual function to initialize the primary window for this application. If bExclusive is true, a pop-up window is created in anticipation of full-screen mode. If it is false, a regular window is made.
InitGraphics()	This function will be discussed later.
InitInput()	This function will be discussed later.
InitSound()	This function will be discussed later.
InitExtraSubsystems()	Virtual function to initialize any additional subsystems the application wants before the scene is initialized.
cApplication()	Constructor; fills in default values for the member variables.
~cApplication()	Shuts down all of the system objects.
Init()	Initializes all of the system objects (which I'll show you in Chapter 3).
Run()	Main part of the application. Displays frames as fast as it can until the WM_QUIT message arrives.
DoFrame()	This function is called every frame by Run(). In it, the subclassing application should perform all game logic and draw the frame. timeDelta is a floating-point value representing how much time has elapsed since the last frame. This is to aid in making applications perform animations at a constant speed independent of the frame rate of the machine.

<code>DoIdleFrame()</code>	This function is called by <code>Run()</code> if the application is currently inactive. Most of the applications that I'll show you won't need this function, but it exists for completeness.
<code>ParseCmdLine()</code>	Virtual function to allow subclasses to view the command line before anything is run.
<code>SceneInit()</code>	Virtual function; overload this to perform scene-specific initialization. Called after the system objects are created.
<code>SceneEnd()</code>	Virtual function; overload to perform scene-specific shutdown code.
<code>Pause()</code>	Pause the application.
<code>UnPause()</code>	Unpause the application.
<code>GetApplication()</code>	Public accessor function to acquire the global application pointer.
<code>KillApplication()</code>	Kills the application and invalidates the global application pointer.
<code>Application()</code>	Global inline function to simplify access to the global application pointer. Equivalent to <code>cApplication::GetApplication()</code> .
<code>AppInstance()</code>	Global inline function to acquire the <code>HINSTANCE</code> of this application.
<code>CreateApplication()</code>	This global function is undefined and must be declared in all further applications. It creates an application object for the code inside <code>GameLib</code> to use. If an application subclasses <code>cApplication</code> with a class <code>cMyApplication</code> , <code>CreateApplication()</code> should simply return <code>(new cMyApplication)</code> .

The `WinMain()` for the application is abstracted away from child applications, hidden inside the `GameLib` code. Just so you don't miss it, the code for it appears below.

```
int APIENTRY WinMain(HINSTANCE hInstance,
                    HINSTANCE hPrevInstance,
                    LPSTR     lpCmdLine,
                    int       nCmdShow)
{
    cApplication *pApp;

    g_hInstance = hInstance;

    try
    {
        pApp = CreateApplication();

        pApp->ParseCmdLine( lpCmdLine );

        pApp->Init();
        pApp->SceneInit();
        pApp->Run();
    }
}
```

```

    }
    catch( cGameError& err )
    {
        /**
         * Knock out the graphics before displaying the dialog,
         * just to be safe.
         */
        if( Graphics() )
        {
            Graphics()->DestroyAll();
        }
        MessageBox(
            NULL,
            err.GetText(),
            "Error!",
            MB_OK|MB_ICONEXCLAMATION );

        // Clean everything up
        delete pApp;
        return 0;
    }

    delete pApp;
    return 0;
}

```

COM: The Component Object Model

Component-based software development is big business. Instead of writing one deeply intertwined piece of software (called *monolithic* software development), a team writes a set of many smaller components that talk to one another. This ends up being an advantage because if the components are modular enough, they can be used in other projects without a lot of headache. Not only that, but the components can be updated and improved independently of each other. As long as the components talk to each other the same way, no problems arise.

To aid in component-based software design, Microsoft created a scheme called the *Component Object Model*, or COM for short. It provides a standard way for objects to communicate with other objects and expose their functionality to other objects that seek it. It is language independent, platform independent, and even machine independent (a COM object can talk to another COM object over a network connection). In this section we cover how COM objects are used in component-based software. As the knowledge required to construct your own COM objects is not necessary for this book, you may want to look in some other books devoted to COM if you need more information. COM is in fact fairly deprecated these days and has been superseded by the .NET architecture. However, it is still used to encapsulate DirectX, so it will be around for some time to come.

A COM object is basically a block of code that implements one or more COM interfaces. (I love circular definitions like this.) A COM interface is just a set of functions. Actually, it's implemented the same way that almost all C++ compilers implement virtual function tables. In C++, COM objects just inherit one or more abstract base classes, which are called COM interfaces. Other classes can get a COM object to do work by calling functions in its interfaces, but that's it. There are no other functions besides the ones in the interfaces, and no access to member variables outside of Get/Set functions existing in the interfaces.

All COM interfaces derive, either directly or indirectly, from a class called IUnknown. In technical terms, this means the first three entries in the vTable of all COM interfaces are the same three functions of IUnknown. The interface is provided in the following:

```
typedef struct interface
interface IUnknown
{
    virtual HRESULT QueryInterface( REFIID iid, void** ppvObject ) = 0;
    virtual ULONG AddRef( void ) = 0;
    virtual ULONG Release( void ) = 0;
};
```

AddRef() and Release() implement reference counting for us. COM objects are created outside of your control. They may be created with new, malloc(), or a completely different memory manager. Because of this you can't simply delete the interface when you're done with it. Reference counting lets the object perform its own memory management. The reference count is the number of other pieces of code that are referencing an object. When you create a COM object, the reference count will most likely be 1, since you're the only one using it. When another piece of code in a different component wants an interface, generally you call AddRef() on the interface to tell the COM object that there is an additional piece of code using it. When a piece of code is done with an interface, it calls Release(), which decrements the reference count. When the reference count reaches 0, it means that no objects are referencing the COM object and it can safely destroy itself.



Warning: If you don't release your COM objects when you're done with them, they won't destroy themselves. This can cause annoying resource leaks in your application.

QueryInterface() is the one function that makes COM work. It allows an object to request another interface from a COM object it has an interface for. You pass QueryInterface() an interface ID, and a pointer to a void pointer to fill with an interface pointer if the requested interface is supported.

As an example, let's consider a car. You create the car object and get an interface pointer to an ICarIgnition interface. If you want to change the radio station, you can ask the owner of the ICarIgnition interface if it also supports the ICarRadio interface.

```
ICarRadio *pRadio = NULL;
HRESULT hr = pIgnition->QueryInterface(
    IID_ICarRadio,
    (VOID**)&pRadio );
if( !pRadio || FAILED( hr ) )
{
    /* handle error */
}

// Now pRadio is ready to use.
```

This is the beauty of COM. The object can be improved without needing to be recompiled. If you decide to add support for a CD player in your car, all a piece of code needs to do is run QueryInterface() for an ICarCDPlayer interface.

Getting COM to work like this forces two restrictions on the design of a system. First up, all interfaces are public. If you poke through the DirectX headers, you'll find the definitions for all of the DirectX interfaces. Any COM program can use any COM object, as long as it has the interface definition and the IDs for the COM interfaces.

A second, bigger restriction is that COM interfaces can never change. Once they are publicly released, they can never be modified in any way (not even fairly harmless modifications, like appending functions to the end of the interface). If this wasn't enforced, applications that used COM objects would need to be recompiled whenever an interface changed, which would defeat COM's whole purpose.

To add functionality to a COM object, you need to add new interfaces. For instance, say you wanted to extend ICarRadio to add bass and treble controls. You can't just add the functions. Instead, you have to put the new functions into a new interface, which would most likely be called ICarRadio2. Any applications that didn't need the new functionality, or ones that predated the addition of the ICarRadio2 interface, wouldn't need to worry about anything and would continue working using the ICarRadio interface. New applications could take advantage of the new functions by simply using QueryInterface() to acquire an ICarRadio2 interface.

The one last big question to address is how COM objects get created. With DirectX, you don't need to worry about a lot of the innards of COM object creation, but I'll give you a cursory overview.

You create a COM object by providing a COM object ID and an interface ID. Internally, the COM creation functions consult the registry, looking for the requested object ID. If the COM object is installed on your system, there will be a registry entry tying an object ID to a DLL. The DLL is loaded by your application, the object is constructed using a DLL-side class factory (returning an IUnknown interface), and then the interface is

QueryInterface()'d for the provided interface ID. If you look up the object ID for Direct3D in the registry, you'll find it sitting there, tied to the Direct 3D DLL.



Note: The registry is a location for Windows to put all sorts of information pertaining to your machine.

So what are these object and interface IDs, and how are they given out? Well, all COM object creators couldn't be in constant communication, making sure the IDs they chose weren't already in use by someone else, so the creators of COM use what are called globally unique identifiers (GUIDs for short). These are 16-byte numbers that are guaranteed to be unique over time and space. (They're made up of an extremely large timestamp in addition to hardware factors like the ID of the network card of the machine that created it.) That way, when an interface is created, a GUID can be generated for it automatically that is guaranteed to not be in use (using a program called GUIDGEN that comes with Visual C++).

Conclusion

So now you know quite a lot about the inner workings of Windows applications. Although we only scratched the surface of Windows, that is thankfully almost all you really need to know to set up the basics of using DirectX.

This page intentionally left blank.

Chapter 2

Getting Started with DirectX 10

Now that you know enough about the basics of Windows, it's time to get down and dirty with DirectX. This chapter shows you everything you need to know to get started with DirectX graphics. In later chapters I'll also show you a little about DirectX Audio and DirectInput; the majority of the book will be about cutting-edge graphics techniques and shaders. So let's get started!

In this chapter I'm going to cover:

- An overview of DirectX
- The components that make up DirectX
- How to initialize Direct3D
- Textures and how to use them
- Sample code for clearing the back buffer, rendering text, and getting output and error messages from DirectX
- And tons more!

What Is DirectX?

Shortly after the release of Windows 95, Microsoft made a push to end DOS's reign as the primary game platform on the PC. Developers weren't convinced by the added abilities that Win32 programming gave (a robust TCP/IP stack, multitasking, access to system information, and so on). They wanted the total control they had in DOS. Besides that, graphics in Windows at that time were done with the Windows GDI, which was obscenely slow, even by the standards of the day.

Microsoft's answer to game developers was *The Game SDK*, which was really the first version of DirectX. At long last, developers could write fast games and still get the advantages of using the Win32 API. With version 2.0, the SDK's name was changed to DirectX. This was because Microsoft realized that game developers weren't the only people who wanted the graphics and audio acceleration provided by SDK; developers of everything from video playback programs to presentation software wanted faster graphics.

Installation

You should install the SDK to its default location in Program Files. For these instructions I'm assuming you have installed to C:\Program Files\Microsoft DirectX SDK (August 2007). If you are using a later version of the SDK or used a different path, keep that in mind for the following instructions.

Installing DirectX is a straightforward process. You'll need to download the DirectX SDK, which is free, from <http://msdn.microsoft.com/DirectX>. Microsoft deploys point releases of DirectX to this web site from time to time with a month and year name attached to it. At the time of writing, the SDK is called the August 2007 SDK and it contains DirectX 9.0 and 10.0 support. Unlike previous version of DirectX, DirectX 10 only runs on Windows Vista, and it requires a DirectX 10 class graphics card.

In the past, DirectX allowed graphics card makers to have a myriad of different types of cards, each with varying feature sets. This created a lot of complexity for programmers, so as a result, with DirectX 10 Microsoft has forced manufacturers to integrate the entire feature set. The only difference is performance. So no matter what kind of DirectX 10 graphics card you buy, you know it will work with your code. The only difference is the speed of the CPU.

Setting Up VC++

All of the code for this book was developed in Visual C++ 2005 Express Edition, which is also free, from <http://msdn.microsoft.com/vstudio/express/visualc/>. You'll also need to download Service Pack 1 and the Windows Platform SDK and follow the somewhat complicated instructions given at <http://msdn.microsoft.com/vstudio/express/visualc/usingpsdk/>.

After you have gotten Visual C++ EE and DirectX installed, you might want to take one of the samples (located in C:\Program Files\Microsoft DirectX SDK (August 2007)\Samples\C++\Direct3D10) and try to compile it. To get it working you'll need to do a couple of things.

Visual C++ needs to know the location of the headers and library files for DirectX so it can correctly compile and link your application. You only need to set this up once, since all projects use the same include and lib directories. To specify the directories, select Options from the Tools menu. In the Options dialog box, select Projects and Solutions > Visual C++ Directories. For include directories, you should enter C:\Program Files\Microsoft DirectX SDK (August 2007)\Include. For lib directories, enter C:\Program Files\Microsoft DirectX SDK (August 2007)\Lib\x86.

The other trick is that the directories need to appear at the top of the list so that they get searched first. When the compiler searches for d3d10.h, for example, it should find C:\Program Files\Microsoft DirectX SDK (August 2007)\Include\D3D10.h (the 10.0 version of the header)

first, and not use any other version. Later versions of the DirectX SDK may have slightly different paths. Also, some versions of Visual Studio include older versions of DirectX, so if you didn't include the new version at the top of the list you would get all kinds of problems.

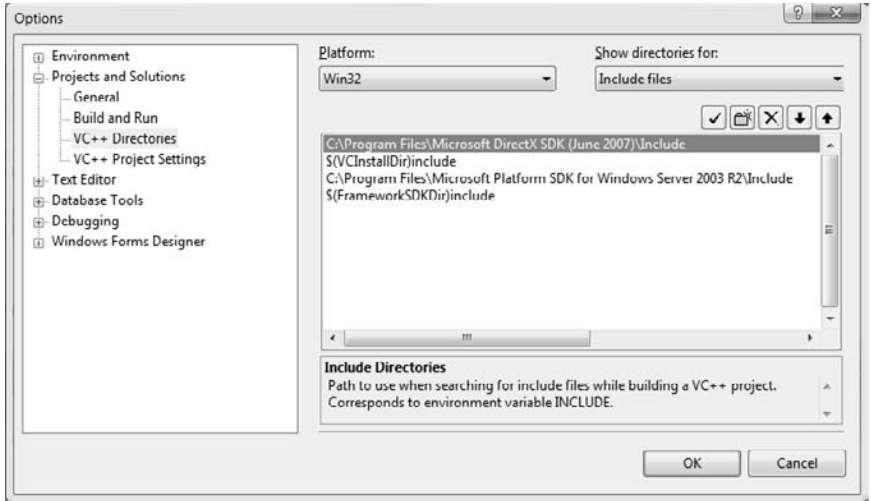


Figure 2.1: The Include directory listing

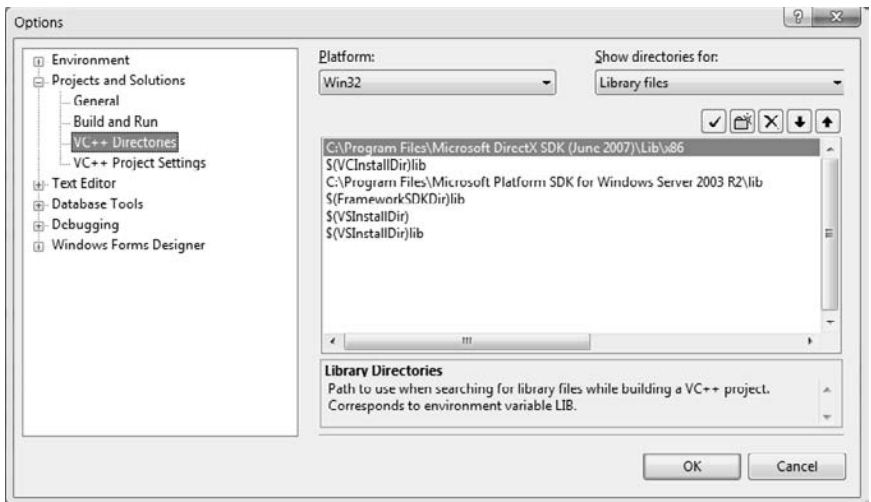


Figure 2.2: The Library directory listing

The other thing you need when building a DirectX application is to have the right libraries listed for the linking step. Most of the applications you write will need the following libraries linked in:

- **winmm.lib**—The Windows multimedia library, which has `timeGetTime`
- **dxguid.lib**—Has the GUIDs for all of the DirectX COM objects
- **d3d10.lib**—Direct3D
- **d3dx10.lib**—Useful D3DX utility extensions
- **dinput8.lib**—DirectInput

Figure 2.3 shows an example of one of the Chapter 9 programs and the first few libraries it links in (it links in several more than can fit in the window).

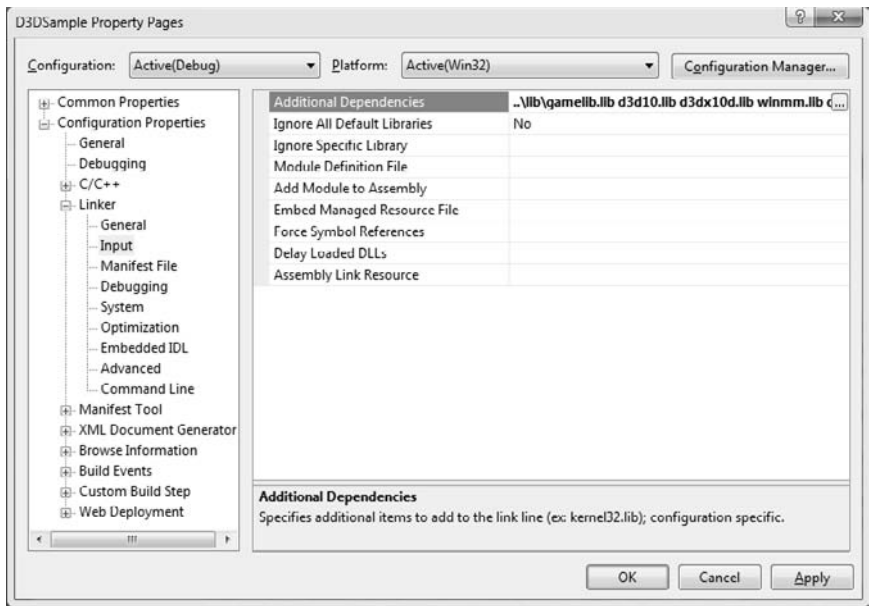


Figure 2.3: Linking in DirectX 10 libraries

If you don't include the libraries correctly, you'll see errors such as these:

```
1>game.lib(GraphicsLayer.obj) : error LNK2019: unresolved external symbol
   _D3D10CreateDeviceAndSwapChain@32 referenced in function "public: void __thiscall
   cGraphicsLayer::InitD3D(int,int,int)" (?InitD3D@cGraphicsLayer@@QAEHHH@Z)
1>.\..\bin\DDSample.exe : fatal error LNK1120: 1 unresolved externals
```

Fixing this would just be a case of linking in `d3d10.lib`.

What Happened to DirectDraw?

If you have any previous experience with DirectX graphics, then you will have probably heard of terms such as DirectDraw, Direct3D, Immediate Mode, and Retained Mode. If not, then don't worry; I'll explain them in a moment. Version 8.0 of DirectX was described by Microsoft as the single most significant upgrade to DirectX since its initial release all those years ago, and version 9.0 continued to build on it. Version 10.0 is to DirectX what Windows XP was to Windows 95. So let me begin with a short introduction into the way things used to be, so that if you come across these terms you will know what is going on.

Graphical output on the PC can be roughly divided into two groups: 2D and 3D, with the latter obviously being far more complex. The implementation of DirectX graphics used to pretty much follow this train of thought. You had *DirectDraw*, which looked after 2D graphics, and *Direct3D*, which looked after 3D. Direct3D was further split into two groups—*Immediate Mode*, which provided a low-level interface to the 3D graphics hardware that was generally considered very complex but fast, and *Retained Mode*, which provided a higher-level, easy-to-use interface to the hardware but was bloated, inflexible, and slow.

As the development of DirectX continued, a number of patterns started to become clear:

- The development of DirectDraw had all but come to an end as of DirectX 5.0. There was just nothing left to do with it, and most resources were being focused on Direct3D.
- The learning curve for DirectDraw was too steep; it was too complicated and required too many tedious steps to set up in code.
- The theoretical split between Direct3D and DirectDraw was becoming a performance bottleneck.
- Direct3D Retained Mode was a complete failure with almost no commercial take-up, and its support was pretty much dropped from DirectX 6.0.
- Direct3D Immediate Mode was too complicated, although it did improve significantly with the release of DirectX 5.0.

To fix these issues, Microsoft took some bold steps in versions 8.0 and 9.0 and completely reorganized Direct3D and DirectDraw. They made the following changes:

- DirectDraw was completely removed as a separate entity and integrated entirely with Direct3D.
- Direct3D Retained Mode was ripped out and was not replaced.
- Direct3D Immediate Mode remains, but is now much more simplified, faster, and just all around more elegant.

- Vertex and pixel shaders were introduced, which allow you to implement advanced visual effects and move away from a fixed function pipeline. In version 9.0 a high-level shader language was introduced, which made shader programming much more intuitive.

Then, when upgrading to DirectX 10, Microsoft made even more drastic changes by rewriting the entire API again, removing the fixed function pipeline, adding in much more advanced HLSL support, and forcing a minimum set of requirements on hardware manufacturers. As you'll see throughout the rest of this book, DirectX 10 is a major leap forward.

Don't forget that although throughout the book I'll be referring to DirectX Graphics as Direct3D, I am not necessarily talking about 3D graphics, since Direct3D now handles the 2D stuff as well. For instance, in the next section I talk heavily about 2D graphics.

OK, so now that you've had your history lesson, let's look at Direct3D in a little more detail.

Direct3D

Direct3D can at first appear confusing; however, once you get used to it, it's really very easy to work with. Although it uses a lot of paradigms you may have never seen before, it will all become clear as you gain experience. It also forces your code to behave nicely with other applications that can be simultaneously using system resources.

Diving into the code that makes Direct3D work will be confusing enough, so to start out I'm just going to talk about the concepts behind the code, which will hopefully make the rocky road ahead a little less painful. If you don't get this stuff immediately, it doesn't mean you're slow and it doesn't mean you're not ready for Windows programming; it means you're normal. DirectX wasn't designed to be easy to use. It was designed to be fast while allowing Windows to maintain some semblance of control over the system. DirectX has gotten much better in recent versions, but it still isn't a trivial thing to pick up.

Direct3D is a set of interfaces and functions that allows a Windows application to talk to the video card(s) in a machine. Only the most basic 2D graphics functions are handled by Direct3D. There are some 2D graphics libraries, such as the GDI, that can do things like draw rounded rectangles, ellipses, lines, thick lines, *n*-sided polygons, and so forth. Direct3D cannot do any of this. Any raster operations need to be developed by you, the game programmer.

What Direct3D *does* do is provide a transparent layer of communication with the hardware on the user's machine. Supported Direct3D functions, like rendering blisteringly fast triangles, are implemented by the video card's super-fast internal hardware.

2D Graphics 101

The way your computer represents images on the screen is as a rectangle of values. Each value is called a *pixel*, short for *picture element*. If the image is m pixels wide and n pixels high, then there are $m*n$ pixels in the image. Each pixel may be anywhere from 1 bit to 4 bytes in size, representing different kinds of color information. The total memory taken up by an image can generally be found as the width of the image times the height of the image times the number of bytes per pixel.

Color on computers is dealt with the same way it is drawn on monitors. CRT computer screens have three cathode ray tubes shining light onto the phosphorous screen dozens of times a second, supplying the color. LCD monitors deliver color with the same basic theory, supplying red, green, and blue pixels, but with quite different technology. By controlling how much red, green, and blue light hits each area of the monitor, the color that results from the phosphors changes. A white pixel, when examined very closely, is actually three smaller pixels: one red, one green, and one blue. You've probably noticed this if you've ever gotten drops of water on your monitor, which magnify what is under them. I wouldn't do this on purpose by the way—try a magnifying glass instead!

There are two main ways that color images are represented on computers. In the first, called *paletted* images, there exists a table of color values (usually with 256 entries) and an image where each pixel is a character indexing into the list. This restricts the image to having 256 unique colors. In the old days, all games used 256-color images for all of their graphics, and before then even fewer colors were used (16 in high-resolution VGA, 4 in EGA and CGA). See Figure 2.4 for a diagram of what this looked like.

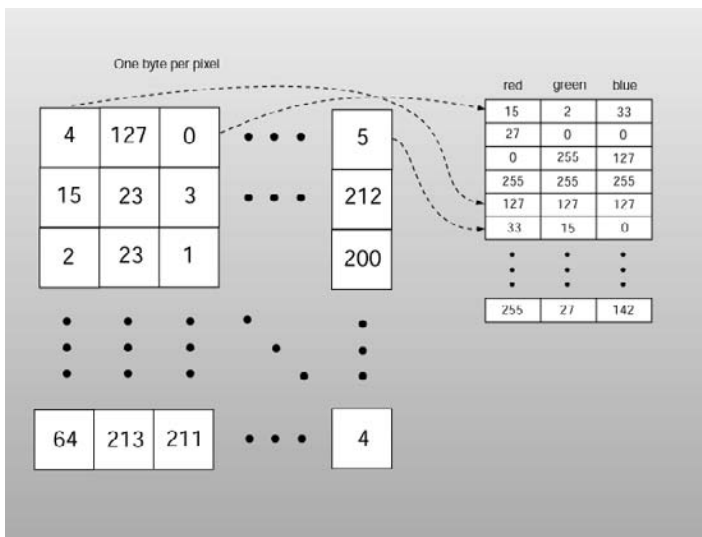


Figure 2.4: How paletted images work

Nowadays every PC you can buy has hardware that can render images with at least 16.7 million individual colors. Rather than have an array with thousands of color entries, the images instead contain explicit color values for each pixel. A 24-bit display, of course, uses 24 bits, or 3 bytes per pixel, for color information. This gives 1 byte, or 256 distinct values each, for red, green, and blue. This is generally called *true color*, because 256^3 (16.7 million) colors is about as much as your eyes can discern, so more color resolution really isn't necessary, at least for computer monitors.

Finally, there is 32-bit color, which is the standard for modern games. The first 24 bytes are for red, green, and blue. The final 8 extra bits per pixel are used to store transparency information, which is generally referred to as the alpha channel, and therefore take up 4 bytes, or 32 bits, of storage per pixel.

Almost universally, all computer images have an origin located at the top-left corner of the image. The top-left corner pixel is referenced with the x,y pair (0,0). The value of x increases to the right; y increases down. This is a departure from the way people usually think of Cartesian coordinates, where the origin is usually in the lower left or center. Figure 2.5 shows the coordinate convention your images will use.




Figure 2.5: Image coordinates

Each horizontal row of pixels is called a *scan line*. The image is stored in memory by attaching each scan line from the top to the bottom end-to-end into a single large one-dimensional array. That way, accessing *pixel* (x,y) on the screen requires you to move across to the correct scan line (the scan line number is y; each scan line is *width* pixels across) and then move across the scan line to the correct pixel.

$$\text{pixel}(x, y) = \text{width} * y + x$$

There's a problem, however. The screen isn't updated instantly. It's a physical device, and as such moves electrons slower than the CPU. The mechanism that lights the screen is internally flying across each scan line of the monitor, reading from the screen's image data and displaying the appropriate colors on the screen. When it reaches the end of a scan line, it moves diagonally down and to the left to the start of the next scan line. When it finishes the last scan line, it moves diagonally up and to the left back to the start of the first scan line. The movement from the bottom-right corner to the top-left corner is called the *vertical blank* or *vertical retrace* (shown in Figure 2.6) and it takes a long time in terms of processor speed. I'm talking years here.




Figure 2.6: The vertical retrace period

Keep this in mind when rendering your images. If you update the screen image at an arbitrary time, the electron gun may be in the middle of the screen. So for that particular frame, the top half of the screen will display the old image, and the bottom half will display the new image. That's assuming you can change the image quickly enough. If you don't, pieces of new image may be smattered all over the screen, creating a horrible, ugly mess. This effect is known as *tearing*.

Because of this, every game under the sun uses a trick called *double buffering*. During rendering, the final image is rasterized into a secondary, off-screen buffer. Then the application waits around for the vertical retrace to begin. When this occurs, it is safe to copy the off-screen image to the on-screen image buffer. The off-screen buffer is generally referred to as the *back buffer*, while the visible image buffer is referred to as the *primary*

surface. You can be fairly sure that the memory copy will finish before the vertical retrace does, so when the electron gun starts drawing again, it's using the new image. While it's drawing the new image, you start rendering the next image into your back buffer, and the cycle continues.



Note: Actually, applications can go a step further and use triple or even quadruple buffering. This is useful to help smooth out jerky frame rates, but requires a lot of video memory (especially at high resolutions).

Textures

2D images in Direct3D are wrapped by objects called *textures*. Internally, a texture is just a structure that manages image data as a contiguous block of memory. The structure keeps track of the vital statistics of the texture, such as its height, width, and format of the pixel. You create textures using the `ID3D10Device` interface, and use the `ID3D10Texture2D` interface to play with them.

One of the features that textures implement is mapping. This is because of the asynchronous (multiple things happening in parallel) nature of many video cards. Instead of having to wait for every operation to finish, you can tell the hardware to do something for you, and it will perform the task in the background while you are attending to other tasks. When multiple things are accessing the same piece of memory at the same time, caution must be taken.

For example, imagine you draw a triangle with a texture. The task gets queued with the other tasks the card is currently doing and will be finished eventually. However, without memory protection, you could quickly copy another image onto the bits of the texture before the render gets executed. When the card got around to performing the render, it would be rendering a different image!

This is a horrible problem. Depending on how much load was on the video card (and whether or not it operates asynchronously; some cards do not), sometimes the texture will be replaced before it is copied, sometimes it won't, sometimes it may even be in the process of being replaced when the card gets to it.

For this reason, you do not have continual access to the raw bits of data that make up your image at all times. The solution DirectX uses is a fairly common concurrency paradigm called a *map*. When you map a texture, you have exclusive access (similar to a lock) to it until you are finished with it. If you request a map on a texture and another piece of code is using it, you won't be able to get it until the other process releases its map on the texture. When you successfully complete a map, you are given a pointer to the raw bits, which you may modify at your leisure, while being confident that no other programs will mess with your memory. In the example of drawing a triangle, Direct3D would map the texture when you requested the render and release it once the render had

completed. If you tried to mangle the bits of the image, your code would not be able to get a pointer to the image data (one of the things you receive when you engage a map) until the map had been released by Direct3D.

Textures, along with having the raw bits of image data, contain a lot of information about the pixels they contain. The width, height, format of the pixel, type of texture, etc., are also stored in the texture. There is another important variable that a texture contains that I should mention, called the *pitch*. Some hardware devices require that image rows begin aligned to 4-pixel boundaries, or 10-pixel boundaries, or any other possible value. If you tried to make an image with an odd width, the card would not be able to handle it. Because of this, Direct3D uses the concept of a pitch in addition to the width.

The pitch of an image is similar to the width; however, it may be a bit bigger to accommodate the way the display adapter works. The address of the pixel directly below the top-left corner of a Direct3D texture is not `texture_width * bytes_per_pixel`. Rather, it is `bytes_per_pixel * texture_pitch`. The texture pitch is *always* measured in bytes; it doesn't vary in relation to the number of bits per pixel. See Figure 2.7.




Figure 2.7: Image width vs. image pitch



Note: This is very important so don't forget it: The pitch of a texture is always measured in bytes and has nothing to do with the number of bits per pixel you are currently working with.

Complex Textures

Textures can be attached to other textures in something called a *texture chain*. This concept is used to represent MIP maps, cubic environment maps (both are discussed in Chapter 9), and swap chains.

Swap chains allow an easier way for the hardware to implement double buffering if two (or more) textures exist in the chain. One of them is actually the screen image, the pixels the electron gun will use to display the image on the monitor. The other is the back buffer, the buffer you render the scene into. When you're done rendering, you can flip the textures, which will actually swap the addresses of the buffers internally on the video card. Then the back buffer becomes the screen image, and the screen image becomes the back buffer. The next time the application starts rendering a frame, it will be rendering into what once was the screen image, while what once was the back buffer is currently being drawn to the screen by the monitor.



Warning: If you're counting on the results from the previous frame when rendering the current frame, be wary. What you get when you're using double buffering with a flipping chain isn't the state of the frame buffer at the end of the previous frame; it's the frame buffer from two frames ago!

Describing Textures

When you create 2D textures or request information about textures, the capabilities and vital statistics for the texture are inscribed in a structure called the *texture description*. The texture description is represented by the D3D10_TEXTURE2D_DESC structure and has the following definition:

```
typedef struct D3D10_TEXTURE2D_DESC {
    UINT Width;
    UINT Height;
    UINT MipLevels;
    UINT ArraySize;
    DXGI_FORMAT Format;
    DXGI_SAMPLE_DESC SampleDesc;
    D3D10_USAGE Usage;
    UINT BindFlags;
    UINT CPUAccessFlags;
    UINT MiscFlags;
} D3D10_TEXTURE2D_DESC;
```

Table 2.1: The D3D10_TEXTURE2D_DESC structure

Width	The width of the texture in pixels.
Height	The height of the texture in pixels.
MipLevels	The number of mip levels in the texture.
ArraySize	The number of textures in the array.

Format	<p>Format of the texture's pixels, which can be any of the following from the Microsoft documentation:</p> <ul style="list-style-type: none"> • <code>DXGI_FORMAT_UNKNOWN</code>—The format is not known. • <code>DXGI_FORMAT_R32G32B32A32_TYPELESS</code>—A four-component, 128-bit typeless format. • <code>DXGI_FORMAT_R32G32B32A32_FLOAT</code>—A four-component, 128-bit floating-point format. • <code>DXGI_FORMAT_R32G32B32A32_UINT</code>—A four-component, 128-bit unsigned-integer format. • <code>DXGI_FORMAT_R32G32B32A32_SINT</code>—A four-component, 128-bit signed-integer format. • <code>DXGI_FORMAT_R32G32B32_TYPELESS</code>—A three-component, 96-bit typeless format. • <code>DXGI_FORMAT_R32G32B32_FLOAT</code>—A three-component, 96-bit floating-point format. • <code>DXGI_FORMAT_R32G32B32_UINT</code>—A three-component, 96-bit unsigned-integer format. • <code>DXGI_FORMAT_R32G32B32_SINT</code>—A three-component, 96-bit signed-integer format. • <code>DXGI_FORMAT_R16G16B16A16_TYPELESS</code>—A four-component, 64-bit typeless format. • <code>DXGI_FORMAT_R16G16B16A16_FLOAT</code>—A four-component, 64-bit floating-point format. • <code>DXGI_FORMAT_R16G16B16A16_UNORM</code>—A four-component, 64-bit unsigned-integer format with values normalized to be between 0 and 1. • <code>DXGI_FORMAT_R16G16B16A16_UINT</code>—A four-component, 64-bit unsigned-integer format. • <code>DXGI_FORMAT_R16G16B16A16_SNORM</code>—A four-component, 64-bit signed-integer format with values normalized to be between -1 and 1. • <code>DXGI_FORMAT_R16G16B16A16_SINT</code>—A four-component, 64-bit signed-integer format. • <code>DXGI_FORMAT_R32G32_TYPELESS</code>—A two-component, 64-bit typeless format. • <code>DXGI_FORMAT_R32G32_FLOAT</code>—A two-component, 64-bit floating-point format. • <code>DXGI_FORMAT_R32G32_UINT</code>—A two-component, 64-bit unsigned-integer format. • <code>DXGI_FORMAT_R32G32_SINT</code>—A two-component, 64-bit signed-integer format. • <code>DXGI_FORMAT_R32G8X24_TYPELESS</code>—A two-component, 64-bit typeless format. • <code>DXGI_FORMAT_D32_FLOAT_S8X24_UINT</code> • <code>DXGI_FORMAT_R32_FLOAT_X8X24_TYPELESS</code> • <code>DXGI_FORMAT_X32_TYPELESS_G8X24_UINT</code> • <code>DXGI_FORMAT_R10G10B10A2_TYPELESS</code>—A four-component, 32-bit typeless format.
--------	---

- `DXGI_FORMAT_R10G10B10A2_UNORM`—A four-component, 32-bit unsigned-integer format with values normalized to be between 0 and 1.
- `DXGI_FORMAT_R10G10B10A2_UINT`—A four-component, 32-bit unsigned-integer format.
- `DXGI_FORMAT_R11G11B10_FLOAT`—A three-component, 32-bit floating-point format.
- `DXGI_FORMAT_R8G8B8A8_TYPELESS`—A four-component, 32-bit typeless format.
- `DXGI_FORMAT_R8G8B8A8_UNORM`—A four-component, 32-bit unsigned-integer format with values normalized to be between 0 and 1.
- `DXGI_FORMAT_R8G8B8A8_UNORM_SRGB`
- `DXGI_FORMAT_R8G8B8A8_UINT`—A four-component, 32-bit unsigned-integer format.
- `DXGI_FORMAT_R8G8B8A8_SNORM`—A four-component, 32-bit signed-integer format with values normalized to be between -1 and 1.
- `DXGI_FORMAT_R8G8B8A8_SINT`—A four-component, 32-bit signed-integer format.
- `DXGI_FORMAT_R16G16_TYPELESS`—A two-component, 32-bit typeless format.
- `DXGI_FORMAT_R16G16_FLOAT`—A two-component, 32-bit floating-point format.
- `DXGI_FORMAT_R16G16_UNORM`—A two-component, 32-bit unsigned-integer format with values normalized to be between 0 and 1.
- `DXGI_FORMAT_R16G16_UINT`—A two-component, 32-bit unsigned-integer format.
- `DXGI_FORMAT_R16G16_SNORM`—A two-component, 32-bit signed-integer format with values normalized to be between -1 and 1.
- `DXGI_FORMAT_R16G16_SINT`—A two-component, 32-bit signed-integer format.
- `DXGI_FORMAT_R32_TYPELESS`—A single-component, 32-bit typeless format.
- `DXGI_FORMAT_D32_FLOAT`
- `DXGI_FORMAT_R32_FLOAT`—A single-component, 32-bit floating-point format.
- `DXGI_FORMAT_R32_UINT`—A single-component, 32-bit unsigned-integer format.
- `DXGI_FORMAT_R32_SINT`—A single-component, 32-bit signed-integer format.
- `DXGI_FORMAT_R24G8_TYPELESS`—A two-component, 32-bit typeless format.
- `DXGI_FORMAT_D24_UNORM_S8_UINT`
- `DXGI_FORMAT_R24_UNORM_X8_TYPELESS`
- `DXGI_FORMAT_X24_TYPELESS_G8_UINT`
- `DXGI_FORMAT_R8G8_TYPELESS`—A two-component, 16-bit typeless format.
- `DXGI_FORMAT_R8G8_UNORM`—A two-component, 16-bit unsigned-integer format with values normalized to be between 0 and 1.

- `DXGI_FORMAT_R8G8_UINT`—A two-component, 16-bit unsigned-integer format.
- `DXGI_FORMAT_R8G8_SNORM`—A two-component, 16-bit signed-integer format with values normalized to be between -1 and 1 .
- `DXGI_FORMAT_R8G8_SINT`—A two-component, 16-bit signed-integer format.
- `DXGI_FORMAT_R16_TYPELESS`—A single-component, 16-bit typeless format.
- `DXGI_FORMAT_R16_FLOAT`—A single-component, 16-bit floating-point format.
- `DXGI_FORMAT_D16_UNORM`
- `DXGI_FORMAT_R16_UNORM`—A single-component, 16-bit unsigned-integer format with values normalized to be between 0 and 1 .
- `DXGI_FORMAT_R16_UINT`—A single-component, 16-bit unsigned-integer format.
- `DXGI_FORMAT_R16_SNORM`—A single-component, 16-bit signed-integer format with values normalized to be between -1 and 1 .
- `DXGI_FORMAT_R16_SINT`—A single-component, 16-bit signed-integer format.
- `DXGI_FORMAT_R8_TYPELESS`—A single-component, 8-bit typeless format.
- `DXGI_FORMAT_R8_UNORM`—A single-component, 8-bit unsigned-integer format with values normalized to be between 0 and 1 .
- `DXGI_FORMAT_R8_UINT`—A single-component, 8-bit unsigned-integer format.
- `DXGI_FORMAT_R8_SNORM`—A single-component, 8-bit signed-integer format with values normalized to be between -1 and 1 .
- `DXGI_FORMAT_R8_SINT`—A single-component, 8-bit signed-integer format.
- `DXGI_FORMAT_A8_UNORM`—A single-component, 8-bit unsigned-integer format with values normalized to be between 0 and 1 .
- `DXGI_FORMAT_R1_UNORM`
- `DXGI_FORMAT_R9G9B9E5_SHAREDEXP`
- `DXGI_FORMAT_R8G8_B8G8_UNORM`
- `DXGI_FORMAT_G8R8_G8B8_UNORM`
- `DXGI_FORMAT_BC1_TYPELESS`—Typeless block-compression format.
- `DXGI_FORMAT_BC1_UNORM`—Block-compression format with values normalized to be between 0 and 1 .
- `DXGI_FORMAT_BC1_UNORM_SRGB`—Block-compression format with values normalized to be between 0 and 1 in sRGB space.
- `DXGI_FORMAT_BC2_TYPELESS`—Typeless block-compression format.
- `DXGI_FORMAT_BC2_UNORM`—Block-compression format with values normalized to be between 0 and 1 .
- `DXGI_FORMAT_BC2_UNORM_SRGB`—Block-compression format with values normalized to be between 0 and 1 in sRGB space.
- `DXGI_FORMAT_BC3_TYPELESS`—Typeless block-compression format.

- `DXGI_FORMAT_BC3_UNORM`—Block-compression format with values normalized to be between 0 and 1.
- `DXGI_FORMAT_BC3_UNORM_SRGB`—Block-compression format with values normalized to be between 0 and 1 in sRGB space.
- `DXGI_FORMAT_BC4_TYPELESS`—Typeless block-compression format.
- `DXGI_FORMAT_BC4_UNORM`—Block-compression format with values normalized to be between 0 and 1.
- `DXGI_FORMAT_BC4_SNORM`—Block-compression format with values normalized to be between -1 and 1.
- `DXGI_FORMAT_BC5_TYPELESS`—Typeless block-compression format.
- `DXGI_FORMAT_BC5_UNORM`—Block-compression format with values normalized to be between 0 and 1.
- `DXGI_FORMAT_BC5_SNORM`—Block-compression format with values normalized to be between -1 and 1.

SampleDesc	Multisampling parameters.
Usage	How the texture is written to and read from; usually set to <code>D3D10_USAGE_DEFAULT</code> .
BindFlags	How the texture is bound to the pipeline; for example, if you wanted to use it as a render target.
CPUAccess-Flags	Flags specifying CPU access to the texture, such as <code>D3D10_CPU_ACCESS_READ</code> .
MiscFlags	Flags specifying uncommon resource options; usually set to 0.

Don't worry, you're not expected to remember all of this yet—just refer back here if you need a refresher. The important thing to do is to keep your head above the mud and keep reading.

Now that you've seen how textures are described, let's look at an actual Direct3D texture.

The ID3D10Texture2D Interface

A Direct3D texture is represented with a COM interface. Since more and more features have been added through the different versions of DirectX, new COM interfaces have been made (remember that you can never change COM interfaces once they're published). The current version of the Direct3D texture interface is called `ID3D10Texture2D`. Unlike the previous version of this interface, `IDirect3DSurface9`, `ID3D10Texture2D` has only three member functions, which are listed in Table 2.2.

Table 2.2: `ID3D10Texture2D` methods

<code>GetDesc()</code>	Fills in a <code>D3D10_TEXTURE2D_DESC</code> structure describing the texture.
<code>Map()</code>	Gets a pointer to the surface data.
<code>Unmap()</code>	Invalidates the pointer acquired through <code>Map()</code> .

Texture Operations

These days DirectX 3D only has a few facilities to help you play with textures directly. One operation you'll need fairly often is copying from one texture to another. To do this you can use the DirectX utility function `D3DX10LoadTextureFromTexture()`, the syntax of which is shown below:

```
HRESULT D3DX10LoadTextureFromTexture(
    ID3D10Resource *pSrcTexture,
    D3DX10_TEXTURE_LOAD_INFO *pLoadInfo,
    ID3D10Resource *pDstTexture
);
```

The term *loading* is the Microsoft lingo for filling a texture with data. As you will see later, you can also load a surface from a file, from a memory resource, or from just about anywhere. Okay, now let's take a moment to look at the function. The parameters are shown in Table 2.3.

Table 2.3: The `D3DX10LoadTextureFromTexture()` parameters

<code>pSrcTexture</code>	Pointer to the source texture.
<code>pLoadInfo</code>	<p>Pointer to the texture loading parameters contained in the <code>D3DX10_TEXTURE_LOAD_INFO</code> structure. This structure looks like this:</p> <pre>typedef struct D3DX10_TEXTURE_LOAD_INFO { D3D10_BOX *pSrcBox; // Source rectangle D3D10_BOX *pDstBox; // Destination rectangle UINT SrcFirstMip; // Index of source mip level UINT DstFirstMip; // Index of destination mip level UINT NumMips; // Number of source mip levels to use UINT SrcFirstElement; // First source array element UINT DstFirstElement; // First destination array element UINT NumElements; // Number of elements to use UINT Filter; // Sampling filter to use UINT MipFilter; // Mip sampling filter to use } D3DX10_TEXTURE_LOAD_INFO;</pre>
<code>pDstTexture</code>	The destination texture.

You will probably use `D3DX10LoadTextureFromTexture()` as your primary tool when copying textures. There is, however, one exception, and that is when you are copying from a back buffer onto the primary surface, which I will cover shortly.

Modifying the Contents of Textures

Applications that wish to perform their own raster functions on textures, such as plotting pixels or drawing triangles with custom rasterizers, must first map the texture before being able to modify the pixels of the image. Mapping allows the CPU to access the texture's raw data.

To map a texture you use the `ID3D10Texture2D::Map()` function, which I have prototyped in the following:

```
HRESULT Map(  
    UINT Subresource,  
    D3D10_MAP MapType,  
    UINT MapFlags,  
    D3D10_MAPPED_TEXTURE2D *pMappedTex2D  
);
```

Table 2.4: Map() parameters

Subresource	The subresource you want to edit. Unless you are working with advanced textures, this is usually 0.
MapType	The CPU permission needed, which is usually D3D10_MAP_READ_WRITE.
MapFlags	Usually set to 0.
pMappedTex2D	A pointer to a D3D10_MAPPED_TEXTURE2D structure, which contains the data you need to write to the texture.

Once you have mapped a texture, the data you need will be contained in the D3D10_MAPPED_TEXTURE2D structure, which looks like this:

```
typedef struct D3D10_MAPPED_TEXTURE2D {  
    void *pData;  
    UINT RowPitch;  
} D3D10_MAPPED_TEXTURE2D;
```

The pData parameter contains a pointer to the raw surface data, and RowPitch contains the pitch of the image in bytes. So if your texture contains standard R8B8G8A8 information, you could cast the data pointer to a D3DCOLOR. Then any (x,y) pixel you want to change would be accessed using the standard pitch * width + x formula.

When you are finished modifying a texture, you must unmap it with the Unmap() function.

Creating Textures

Textures are created by the Direct3D device with the function ID3D10Device::CreateTexture2D(), which is prototyped below:

```
HRESULT CreateTexture2D(  
    const D3D10_TEXTURE2D_DESC *pDesc,  
    const D3D10_SUBRESOURCE_DATA *pInitialData,  
    ID3D10Texture2D **ppTexture2D  
);
```

The first parameter takes the address of a filled out D3D10_TEXTURE2D_DESC structure, which we saw previously. The second parameter can be NULL, and the third takes the address of a pointer to be filled with the address of the newly created texture. When you create a texture, you should set the CPUAccessFlags member of the texture description to be D3D10_CPU_ACCESS_WRITE if you will be mapping and writing directly to the texture's raw data.

Implementing Direct3D with cGraphicsLayer

To implement Direct3D we are going to use a class called cGraphicsLayer, which will only have one instance, cApplication. Its abilities include:

- Initializing windowed Direct3D
- Providing easy access to Direct3D objects if needed
- Being able to initialize Direct3D with the primary display adapter

Notice how in Figure 2.8, since our application is windowed, Direct3D automatically clips our rendering to our window and prevents us from rendering on top of other applications. Let's dive into the code by first having a look at DxHelper.h, which helps simplify some of the programming tasks.




Figure 2.8: Clipped rendering

```

/*****
 *      Advanced 3D Game Programming with DirectX 10.0
 *      * * * * *
 *
 *      See license.txt for modification and distribution information
 *      copyright (c) 2007 by Peter Walsh, Wordware
 *****/

#ifdef _D3DHELPER_H
#define _D3DHELPER_H

#include <memory.h>

```

```

/**
 * This class takes care of the annoying work
 * of having to zero-out and set the size parameters
 * of our Windows and DirectX structures.
 */
template <class T>
struct sAutoZero : public T
{
    sAutoZero()
    {
        memset( this, 0, sizeof(T) );
        dwSize = sizeof(T);
    }
};

/**
 * The right way to release our COM interfaces.
 * If they're still valid, release them, then
 * invalidate them and null them.
 */
template <class T>
inline void SafeRelease( T& iface )
{
    if( iface )
    {
        iface->Release();
        iface = NULL;
    }
}

#endif // _D3DHELPER_H

```

The interface for the graphics layer is in GraphicsLayer.h.

```

/*****
 *          Advanced 3D Game Programming with DirectX 10.0
 *          * * * * *
 *
 *   See license.txt for modification and distribution information
 *   copyright (c) 2007 by Peter Walsh, Wordware
 *****/

#ifndef _GRAPHICSLAYER_H
#define _GRAPHICSLAYER_H

#include <d3d10.h>
#include <d3dx10.h>

#include <list>
#include <string>
using std::wstring;
using std::list;

#include "GameTypes.h"
#include "DxHelper.h"

```

```

//-----
// Manages the insertion point when drawing text
//-----

class CApplication;

class CGraphicsLayer
{

protected:

    HWND                m_hWnd;                // The handle to the window
    ID3D10Device         *m_pDevice;            // The IDirect3DDevice10
                                                // interface
    ID3D10Texture2D     *m_pBackBuffer;        // Pointer to the back buffer
    ID3D10RenderTargetView *m_pRenderTargetView; // Pointer to render target view
    IDXGISwapChain       *m_pSwapChain;        // Pointer to the swap chain
    RECT                 m_rcScreenRect;        // The dimensions of the screen
    CGraphicsLayer(HWND hWnd);                  // Constructor

    // Pointer to the main global gfx object
    static CGraphicsLayer *m_pGlobalGLayer;

    ID3DX10Font          *m_pFont;              // The font used for rendering text

    // Sprites used to hold font characters
    ID3DX10Sprite        *m_pFontSprite;

    // Queue used to hold messages from D3D
    ID3D10InfoQueue      *m_pMessageQueue;

    static const UINT     m_uiMAX_CHARS_PER_FRAME = 512;

public:

    void DestroyAll();
    ~CGraphicsLayer();

    /**
     * Initialization calls.
     */

    void InitD3D(int width, int height, int bpp);

    /**
     * This function uses Direct3DX to write text to the back buffer.
     * It's much faster than using the GDI
     */
    void DrawTextString(int x, int y, D3DXCOLOR color, const TCHAR *strOutput);

    void DumpMessages();
    //===== Accessor functions

    // Gets a pointer to the device
    ID3D10Device *GetDevice()

```

```
{
    return m_pDevice;
}

// Gets a pointer to the back buffer
ID3D10Texture2D *GetBackBuffer()
{
    return m_pBackBuffer;
}

// Gets the screen width
int Width() const
{
    return m_rcScreenRect.right;
}

// Gets the screen height
int Height() const
{
    return m_rcScreenRect.bottom;
}

// Presents the back buffer to the primary surface
void Present();

// Clears the back buffer
void Clear(const float (&colClear)[4]);

// Gets a pointer to the main gfx object
static cGraphicsLayer *GetGraphics()
{
    return m_pGlobalGLayer;
}

// Initializes this object
static void Create(
    HWND hWnd,           // handle to the window
    short width, short height); // Device guid
};

inline cGraphicsLayer *Graphics()
{
    return cGraphicsLayer::GetGraphics();
}

#endif // _GRAPHICSLAYER_H
```

GraphicsLayer.cpp is pretty long, so we'll explain it in sections. First up are a few helper functions that don't directly deal with initialization. The first of these is DrawTextString(), which is used to print strings of text to the screen. It uses a D3DX10 utility object called ID3DX10Font. This object sorts out all the complicated font parameters and renders text to the screen using an array of ID3DX10Sprite objects.

```

void cGraphicsLayer::DrawTextString(int x, int y,
                                   D3DXCOLOR color, const TCHAR *strOutput)
{
    m_pFontSprite->Begin(0);
    RECT rect = {x, y, m_rcScreenRect.right, m_rcScreenRect.bottom};
    m_pFont->DrawText(m_pFontSprite, strOutput, -1, &rect, DT_LEFT, color);
    m_pFontSprite->End();
}

```

The function `Present()` is called once a frame to copy the contents of the back buffer to the primary surface:

```

void cGraphicsLayer::Present()
{
    HRESULT r = S_OK;

    assert(m_pDevice);

    r = m_pSwapChain->Present(0, 0);
    if(FAILED(r))
    {
        OutputDebugString(L"Present Failed!\n");
    }

    DumpMessages();
}

```

The next function, `DumpMessages()`, is used to get any messages from DirectX and dump them to the output window. It uses the `ID3D10InfoQueueObject` to gain access to the messages.

```

void cGraphicsLayer::DumpMessages()
{
    assert(m_pMessageQueue);

    HRESULT r = 0;

    while(1)
    {
        // Get the size of the message
        SIZE_T messageLength = 0;
        r = m_pMessageQueue->GetMessage(0, NULL, &messageLength);
        if(messageLength == 0)
            break;

        // Allocate space and get the message
        D3D10_MESSAGE *pMessage = (D3D10_MESSAGE*)malloc(messageLength);
        r = m_pMessageQueue->GetMessage(0, pMessage, &messageLength);
        if(FAILED(r))
        {
            OutputDebugString(L"Failed to get Direct3D Message");
            break;
        }

        TCHAR strOutput[MAX_PATH];
        swprintf_s(strOutput, MAX_PATH,

```



```

        L"D3DDMSG [Cat[%i] Sev[%i] ID[%i]: %s\n",
        pMessage->Category, pMessage->Severity,
        pMessage->ID, pMessage->pDescription);
    OutputDebugString(strOutput);
}
}

```

The final function we have is called `Clear()`, and it is used to clear the back buffer to a color defined as a parameter.

```

void cGraphicsLayer::Clear(const float (&colClear)[4])
{
    HRESULT r = S_OK;
    assert(m_pDevice);

    m_pDevice->ClearRenderTargetView(m_pRenderTargetView, colClear);
}

```

Creating the Graphics Layer

Before we move on to the code that initializes Direct3D, it's worth looking briefly at how the layer is created in the first place. It is in fact called from `cApplication::InitGraphics`, which calls the static function `cGraphicsLayer::Create()`. This function is shown below:

```

void cGraphicsLayer::Create(HWND hWnd, short width, short height)
{
    new cGraphicsLayer(hWnd); // construct the object.

    // Init Direct3D and the device for full-screen operation
    Graphics()->InitD3D(width, height, 32);
}

```

Now let's check out the `InitD3D()` function, which is where all the action is.

Initializing Direct3D

Step 1: Creating a Device and Swap Chain

Since all our rendering is done with a Direct3D device, the first thing we need to do is construct a device. This also involves constructing a swap chain, consisting of a back buffer and primary surface. Doing this is fairly simple. We fill out a structure called `DXGI_SWAP_CHAIN_DESC` and then call the function `D3D10CreateDeviceAndSwapChain()`. Let's look at the structure first, which has this prototype:

```

typedef struct DXGI_SWAP_CHAIN_DESC {
    DXGI_MODE_DESC BufferDesc;
    DXGI_SAMPLE_DESC SampleDesc;
    DXGI_USAGE BufferUsage;
    UINT BufferCount;
    HWND OutputWindow;
}

```

```

    BOOL Windowed;
    DXGI_SWAP_EFFECT SwapEffect;
    UINT Flags;
} DXGI_SWAP_CHAIN_DESC;

```

Table 2.5: DXGI_SWAP_CHAIN_DESC structure members

BufferDesc	<p>A DXGI_MODE_DESC structure containing information about the format of the primary surface and back buffer. This structure looks like this:</p> <pre> typedef struct DXGI_MODE_DESC { UINT Width; UINT Height; DXGI_RATIONAL RefreshRate; DXGI_FORMAT Format; DXGI_MODE_SCANLINE_ORDER ScanlineOrdering; DXGI_MODE_SCALING Scaling; } DXGI_MODE_DESC, *LPDXGI_MODE_DESC; </pre> <p>Width—The width of the buffer in pixels.</p> <p>Height—The height of the buffer in pixels.</p> <p>RefreshRate—A structure containing a numerator and denominator for the refresh rate, which will usually be 60/1.</p> <p>Format—The format of the back buffer, which will usually be DXGI_FORMAT_R8G8B8A8_UNORM.</p> <p>ScanlineOrdering—Usually set to 0.</p> <p>Scaling—Usually set to 0.</p>
SampleDesc	<p>Defines the multisampling quality settings in a DXGI_SAMPLE_DESC structure, which looks like this:</p> <pre> typedef struct DXGI_SAMPLE_DESC { UINT Count; UINT Quality; } DXGI_SAMPLE_DESC, *LPDXGI_SAMPLE_DESC; </pre> <p>Count—The number of pixels to multisample. We'll set this to 1 for now, which disables multisampling.</p> <p>Quality—The quality of the multisampling, which we'll set to 0 for now.</p>
BufferUsage	Set this to DXGI_USAGE_RENDER_TARGET_OUTPUT.
BufferCount	The number of back buffers to use. Set this to 1.
OutputWindow	The HWND of the window to render to.
Windowed	TRUE for windowed, and FALSE for full-screen.
SwapEffect	The effect of the call to Present(). We usually set this to DXGI_SWAP_EFFECT_DISCARD, which means we don't need the back buffer to remain the same after it is presented to the primary surface.
Flags	Advanced flags—usually set to 0.

Although this structure looks pretty complicated with all its sub-structures, it's actually not so bad since most of the members generally stay the same. Check out the code below, which sets up the structure. Note that since the structure is memset to 0, I don't bother filling out all the values:

```
void cGraphicsLayer::InitD3D(int width, int height, int bpp)
{
    HRESULT r = 0;

    // Structure to hold the creation parameters for the device
    DXGI_SWAP_CHAIN_DESC swapDesc;
    ZeroMemory(&swapDesc, sizeof(swapDesc));

    // Only want one back buffer
    swapDesc.BufferCount = 1;

    // Width and height of the back buffer
    swapDesc.BufferDesc.Width = 640;
    swapDesc.BufferDesc.Height = 480;

    // Standard 32-bit surface type
    swapDesc.BufferDesc.Format = DXGI_FORMAT_R8G8B8A8_UNORM;

    // 60hz refresh rate
    swapDesc.BufferDesc.RefreshRate.Numerator = 60;
    swapDesc.BufferDesc.RefreshRate.Denominator = 1;
    swapDesc.BufferUsage = DXGI_USAGE_RENDER_TARGET_OUTPUT;

    // Connect it to our main window
    swapDesc.OutputWindow = m_hWnd;

    // No multisampling
    swapDesc.SampleDesc.Count = 1;
    swapDesc.SampleDesc.Quality = 0;

    // Windowed mode
    swapDesc.Windowed = TRUE;
```

Now we can create the device and swap chain with a call to `D3D10CreateDeviceAndSwapChain()`, which has the following prototype:

```
HRESULT D3D10CreateDeviceAndSwapChain(
    IDXGIAdapter *pAdapter,
    D3D10_DRIVER_TYPE DriverType,
    HMODULE Software,
    UINT Flags,
    UINT SDKVersion,
    DXGI_SWAP_CHAIN_DESC *pSwapChainDesc,
    IDXGISwapChain **ppSwapChain,
    ID3D10Device **ppDevice
);
```

Table 2.6: D3D10CreateDeviceAndSwapChain() parameters

pAdapter	Pointer to an adapter, or NULL to use the primary display adapter.
DriverType	The type of driver to use. We will usually set this to D3D10_DRIVER_TYPE_HARDWARE.
Software	Handle to a software rendering module, which is almost always NULL.
Flags	Flags for device creation. This can be NULL, but for development it's a good idea to set this to D3D10_CREATE_DEVICE_DEBUG. It's slower but allows you to get helpful debug information out of DirectX to track down problems.
SDKVersion	Just set this to D3D10_SDK_VERSION so that Direct3D knows it's being compiled with the correct version number.
pSwapChainDesc	Address of a filled out DXGI_SWAP_CHAIN_DESC structure we looked at a moment ago.
ppSwapChain	The address of a pointer to an IDXGISwapChain that will be filled with a pointer to the newly created swap chain.
ppDevice	Address of a pointer to an ID3D10Device that will be filled with a pointer to the newly created device.

So to create a standard hardware device with debugging enabled you could use the following code:

```
// Create the device using hardware acceleration
r = D3D10CreateDeviceAndSwapChain(
    NULL,                // Default adapter
    D3D10_DRIVER_TYPE_HARDWARE, // Hardware accelerated device
    NULL,                // Not using a software DLL for rendering
    D3D10_CREATE_DEVICE_DEBUG, // Flag to allow debug output
    D3D10_SDK_VERSION,   // Indicates the SDK version being used
    &swapDesc,
    &m_pSwapChain,
    &m_pDevice);

if(FAILED(r))
{
    throw cGameError(L"Could not create IDirect3DDevice10");
}
```

Step 2: Creating a Render Target View

In previous versions of DirectX, that was all you really needed to do; however, in DirectX 10 we need one final step—to create a view of the render target. This is kind of like converting the render target into a form that the GPU understands. It's very easy to do; here is the code:

```
// Create a render target view
r = m_pDevice->CreateRenderTargetView(
    m_pBackBuffer, NULL, &m_pRenderTargetView);
```

```
if(FAILED(r))
{
    throw cGameError(L"Could not create render target view");
}
```

Step 3: Putting It All Together

Now check out the entire `InitD3D()` function. You'll see it does all the things we talked about and a few others, like saving a pointer to the back buffer and setting up the Direct3D queued debug output system.

```
void cGraphicsLayer::InitD3D(int width, int height, int bpp)
{
    HRESULT r = 0;

    // Structure to hold the creation parameters for the device
    DXGI_SWAP_CHAIN_DESC swapDesc;
    ZeroMemory(&swapDesc, sizeof(swapDesc));

    // Only want one back buffer
    swapDesc.BufferCount = 1;

    // Width and height of the back buffer
    swapDesc.BufferDesc.Width = 640;
    swapDesc.BufferDesc.Height = 480;

    // Standard 32-bit surface type
    swapDesc.BufferDesc.Format = DXGI_FORMAT_R8G8B8A8_UNORM;

    // 60hz refresh rate
    swapDesc.BufferDesc.RefreshRate.Numerator = 60;
    swapDesc.BufferDesc.RefreshRate.Denominator = 1;
    swapDesc.BufferUsage = DXGI_USAGE_RENDER_TARGET_OUTPUT;

    // Connect it to our main window
    swapDesc.OutputWindow = m_hWnd;

    // No multisampling
    swapDesc.SampleDesc.Count = 1;
    swapDesc.SampleDesc.Quality = 0;

    // Windowed mode
    swapDesc.Windowed = TRUE;

    // Create the device using hardware acceleration
    r = D3D10CreateDeviceAndSwapChain(
        NULL,                // Default adapter
        D3D10_DRIVER_TYPE_HARDWARE, // Hardware accelerated device
        NULL,                // Not using a software DLL for rendering
        D3D10_CREATE_DEVICE_DEBUG, // Flag to allow debug output
        D3D10_SDK_VERSION,   // Indicates the SDK version being used
        &swapDesc,
```

```

        &m_pSwapChain,
        &m_pDevice);

if(FAILED(r))
{
    throw cGameError(L"Could not create IDirect3DDevice10");
}

r=m_pDevice->QueryInterface(
    __uuidof(ID3D10InfoQueue), (LPVOID*)&m_pMessageQueue);
if(FAILED(r))
{
    throw cGameError(
        L"Could not create IDirect3DDevice10 message queue");
}
m_pMessageQueue->SetMuteDebugOutput(false);    // No muting
m_pMessageQueue->SetMessageCountLimit(-1);    // Unlimited messages

// Keep a copy of the screen dimensions
m_rcScreenRect.left = m_rcScreenRect.top = 0;
m_rcScreenRect.right = width;
m_rcScreenRect.bottom = height;

// Get a copy of the pointer to the back buffer
r = m_pSwapChain->GetBuffer(0,
    __uuidof(ID3D10Texture2D), (LPVOID*)&m_pBackBuffer);
if(FAILED(r))
{
    throw cGameError(L"Could not get back buffer");
}

// Create a render target view
r = m_pDevice->CreateRenderTargetView(
    m_pBackBuffer, NULL, &m_pRenderTargetView);
if(FAILED(r))
{
    throw cGameError(L"Could not create render target view");
}

// Attach the render target view to the output merger state
m_pDevice->OMSetRenderTargets(1, &m_pRenderTargetView, NULL);

// Create a viewport the same size as the back buffer
D3D10_VIEWPORT vp;
vp.Width = width;
vp.Height = height;
vp.MinDepth = 0.0f;
vp.MaxDepth = 1.0f;
vp.TopLeftX = 0;
vp.TopLeftY = 0;
m_pDevice->RSSetViewports( 1, &vp );

// Create the font for rendering text
D3DX10CreateFont(m_pDevice,
    15, 0,

```

```

        FW_BOLD,
        1,
        FALSE,
        DEFAULT_CHARSET,
        OUT_DEFAULT_PRECIS,
        DEFAULT_QUALITY,
        DEFAULT_PITCH | FF_DONTCARE,
        L"Arial",
        &m_pFont);
assert(m_pFont);

// Create the sprite to use to render letters
D3DX10CreateSprite(m_pDevice, m_uiMAX_CHARS_PER_FRAME, &m_pFontSprite);
}

```

Shutting Down Direct3D

Shutting down Direct3D is very simple. It is handled in `cGraphicsLayer::DestroyAll()`. Each object is released in the opposite order it was created.

```

void cGraphicsLayer::DestroyAll()
{
    SafeRelease(m_pMessageQueue);
    SafeRelease(m_pFont);
    SafeRelease(m_pFontSprite);
    SafeRelease(m_pRenderTargetView);
    SafeRelease(m_pBackBuffer);
    SafeRelease(m_pSwapChain);
    SafeRelease(m_pDevice);

    /**
     * Prevent any further access to the graphics class
     */
    m_pGlobalGLayer = NULL;
}

```

Sample Application: Direct3D Sample

Now let's look at a very simple application to demonstrate how future samples in the book will be put together. This sample just boots up Direct3D and clears the back buffer to random colors. It also renders each text frame to a random location. The source is very short; check it out:

```

/*****
 *           Advanced 3D Game Programming with DirectX 10
 * * * * *
 *
 *   See license.txt for modification and distribution information
 *   copyright (c) 2007 by Peter Walsh, Wordware
 *****/

```

```

#include "stdafx.h"

#include <string>
using namespace std;

class cD3DSampleApp : public cApplication
{
public:

    =    //=====----- cApplication

    virtual void DoFrame( float timeDelta );

    cD3DSampleApp() :
        cApplication()
    {
        m_title = wstring( L"Direct3D Sample" );
    }
};

cApplication *CreateApplication()
{
    return new cD3DSampleApp();
}

void cD3DSampleApp::DoFrame( float timeDelta )
{
    if(!Graphics())
        return;

    // Clear the previous contents of the back buffer
    float colClear[4] = {RandFloat(), RandFloat(), RandFloat(), 1.0f};
    Graphics()->Clear(colClear);

    // Output green text at a random location
    Graphics()->DrawTextString( rand()%640, rand()%480,
        D3DXCOLOR(0, 1, 0, 1), L"Advanced Direct3D 10.0" );

    // Present the back buffer to the primary surface to make it visible
    Graphics()->Present();
}

```





Figure 2.9: Clearing the back buffer and rendering text

As you can see in Figure 2.9, now that all the initialization code is encapsulated into `cGraphicsLayer`, it's very simple to create a Direct3D 10 application. We'll be extending the graphics code later, particularly when we start looking at shaders. However, for now the initialization code is out of the way, which is nice since it always stands in the way of having fun.

Conclusion

This chapter introduced you to setting up DirectX 10 for rendering. We learned about textures and the rendering device, and saw how to create a device, swap chain, and render target. Next, we are going to take a slight detour away from graphics while we set up some basic sound and input.

Chapter 3

Input and Sound

Getting input from the user is probably the most important part of any computer game. Without input, no matter how snazzy the graphics or how great the sound, you are effectively just watching a movie. I am a recent convert to the importance of sound in games. I used to think that the main selling point of any game was the initial shock value of how its graphics looked. However, a while back I attended a Microsoft DirectX Meltdown conference in London, where one of the Microsoft guys showed me a rewritten version of *3D Boids* that included full support for advanced 3D DirectSound and DirectMusic. The results were really spectacular. The original, slightly boring game became incredibly immersive at the press of a key. Since then I have always strived to ensure that audio is a huge part of any project I work on.

In this chapter you will learn all about:

- Acquiring input with DirectInput
- Communicating with the keyboard and mouse
- Cooperative input modes
- Loading wave files and playing them back with DirectSound
- Working with the primary sound buffer
- Incorporating input and sound into a sample application

The core of this book is about 3D graphics, so we won't be covering input and sound at that same level of detail. Instead, this chapter will give you a grounding, allowing you to get started in these exciting areas so you can continue further if you wish. Let's start with input.

DirectInput

A game needs to get information from the keyboard, mouse, or joystick. In the past, programming these different devices into a usable form meant *a lot* of effort and hair pulling, particularly with joysticks. To fix this problem Microsoft created *DirectInput*. Although there have been heavy advances made in DirectX Graphics, nothing has changed with DirectInput in years and it is still called DirectInput 8. In other words, it's pretty much still the same as the version that shipped with DirectX 8 in 2000.

DirectInput was created to provide a direct way to communicate with the input devices that exist on users' systems. It supplies the ability to enumerate all the devices connected to a machine and even enumerate the

capabilities of a particular device. You can take any input device under the sun; as long as it has a DirectInput driver written for it, your application can talk to it. And these days virtually every device has a DirectInput driver.

There are a lot of nifty features in DirectInput like force feedback, but with a limited amount of space to discuss it, this DirectInput discussion will be restricted to just mouse and keyboard usage. However, once you understand the concepts that make DirectInput work, getting more complex things done with it won't be difficult.

The Win32 API has a full set of window messages that can inform you when keys are pressed, when the mouse moves, etc. There is even rudimentary support for joysticks. So what advantages are there to using DirectInput over the standard API calls?

Well, there are several reasons:

- The Win32 API was not designed for games or speed.
- Joystick support under Win32 is flaky at best. Supporting complex joysticks with several axes, 8 to 10 buttons, a point of view hat, etc., just can't be done on Win32.
- The mouse support is limited to three buttons, two axes, and the mouse wheel if one is present. Many mice on the market today have four, five, or even more buttons.
- The keyboard support in Win32 was designed for keyboard entry applications. There is a lot of functionality to handle automatically repeating keys, conversion from key codes to ASCII characters, and so on, that a game just doesn't need, and ends up wasting valuable processor cycles.
- The Win32 keyboard handling code captures some keys for you (like Alt) that require special message processing to handle correctly.
- Message processing isn't the fastest thing in the world. Applications get flooded with mouse message requests, and since you can't render a frame until the message queue is empty, this can slow down the application.

Devices

An `IDirectInputDevice8` represents a physical object that can give input to the computer. The keyboard, the mouse, and any joysticks/joypads are examples of devices. You talk to devices through a COM interface, just like with Direct3D.

Devices are composed of a set of objects, each one defining a button, axis, POV (point of view) hat, etc. A device can enumerate the objects it contains using `IDirectInputDevice8::EnumObjects()`. This is only really useful for joysticks, as keyboards and mice have a standard set of objects.

An object is described by a structure called `DIDeviceObjectInstance`. The set of DirectInput functionality that I'm going to show you doesn't require you to understand the workings of this structure, but I'll give you a peek at it anyway. The structure has, among other things, a GUID that describes the type of object. The current set of object types appears in Table 3.1. More may appear in the future as people create newer and better object types.

Table 3.1: The current set of object type GUIDs

<code>GUID_XAxis</code>	An axis representing movement in the x-axis (for example, left-to-right movement on a mouse).
<code>GUID_YAxis</code>	An axis representing movement in the y-axis (for example, up-to-down movement on a mouse).
<code>GUID_ZAxis</code>	An axis representing movement in the z-axis (for example, the mouse wheel on newer models).
<code>GUID_RxAxis</code>	An axis representing rotation relative to the x-axis.
<code>GUID_RyAxis</code>	An axis representing rotation relative to the y-axis.
<code>GUID_RzAxis</code>	An axis representing rotation relative to the z-axis.
<code>GUID_Slider</code>	A slider axis (for example, the throttle slider that appears on some joysticks).
<code>GUID_Button</code>	A button (on a mouse or joystick).
<code>GUID_Key</code>	A key (on a keyboard).
<code>GUID_POV</code>	A POV hat that appears on some joysticks.
<code>GUID_Unknown</code>	An unknown type of device.

When an application requests the current state of the device, the information needs to be transmitted in some meaningful way. Just getting a list of bytes wouldn't provide enough information, and forcing applications to use a standard communication method wouldn't elegantly solve the problem for all the different types of devices on the market. Because of this, DirectInput lets the application dictate to the device how it wishes to receive its data. If you only want one or two buttons on a joystick, you don't need to request all of the data from the joystick, which may have dozens of buttons. Among other things, the application can decide if any axes on the device should be absolute (centered around a neutral origin, like a joystick axis) or relative (freely moving, like a mouse axis). To tell DirectInput what kind of data you are expecting you must call `IDirectInputDevice8::SetDataFormat()`.

```
HRESULT IDirectInputDevice8::SetDataFormat(
    LPCDIDATAFORMAT lpdf
);
```

lpdf	<p>A pointer to a <code>DIDATAFORMAT</code> structure that defines the format of the data received from the device. There are some defined constants that you can use:</p> <ul style="list-style-type: none"> <code>c_dfDIKeyboard</code>—Standard keyboard structure. An array of 256 characters, one for each key. <code>c_dfDIMouse</code>—Standard mouse structure. Three axes and four buttons. Corresponds to the <code>DIMOUSESTATE</code> structure. <code>c_dfDIMouse2</code>—Extended mouse structure. Three axes and eight buttons. Corresponds to the <code>DIMOUSESTATE2</code> structure. <code>c_dfDIJoystick</code>—Standard joystick. Three positional axes, three rotation axes, two sliders, a POV hat, and 32 buttons. Corresponds to the <code>DIJOYSTATE</code> structure. <code>c_dfDIJoystick2</code>—Extended capability joystick. Refer to the SDK documentation for the truly massive data format definition. Corresponds to the <code>DIJOYSTATE2</code> structure.
------	--

Receiving Device States

There are two ways to receive data from a device: *immediate* data access and *buffered* data access. My code only uses immediate data access, but buffered data access is not without its merits. Buffered data access is useful for when you absolutely need to get every input event that happens. If a key is quickly pressed and released between immediate device state requests, you will miss it since the state changes aren't queued. If the application is running at any reasonable frame rate, however, this won't be a problem. Immediate data access is used to find the current state of the device at some point in time. If buttons were pressed and released between your requests, you don't see them. You ask for the device state using `IDirectInputDevice8::GetDeviceState()`:

```
HRESULT IDirectInputDevice8::GetDeviceState(
    DWORD cbData,
    LPVOID lpvData
);
```

cbData	Size, in bytes, of the data structure being passed in with <code>lpvData</code> .
lpvData	Pointer to a buffer to fill with the device state. The format of the data depends on the format you defined using <code>SetDataFormat()</code> .

For mouse devices, if you set the data format to `c_dfDIMouse`, the parameters to `GetDeviceData()` should be `sizeof(DIMOUSESTATE)` and the address of a valid `DIMOUSESTATE` structure. After the function completes, if it is successful, the structure will be filled with the data from the mouse.

```
typedef struct DIMOUSESTATE {
    LONG lX;
    LONG lY;
```

```

    LONG lZ;
    BYTE rgbButtons[4];
} DIMOUSESTATE, *LPDIMOUSESTATE;

```

lX	X-axis of movement. Relative movement; if the axis hasn't moved since the last time you checked, this will be 0.
lY	Y-axis of movement. Relative movement; if the axis hasn't moved since the last time you checked, this will be 0.
lZ	Z-axis (mouse wheel) movement. Relative movement; if the axis hasn't moved since the last time it was checked, this will be 0.
rgbButtons	A set of bytes, one for each of four mouse buttons. To support a mouse with more buttons, use the DIMOUSESTATE2 structure.

As for the keyboard data, all you do is pass in a 256-element array of characters. Each character represents a certain key. You can index into the array to find a certain key using the `DirectInput` key constants. There is a constant for every possible key on a keyboard. Table 3.2 shows a list of the common ones. Some of the more obscure ones, like the ones for Japanese keyboards and web keyboards, are not included. See the SDK documentation for a complete list at *DirectX 10 C++ Documentation/DirectX Input/DirectInput/Reference/Device Constants/Keyboard Device Constants*.

Table 3.2: The common `DirectInput` keyboard constants

<code>DIK_A ... DIK_Z</code>	A through Z keys
<code>DIK_0 ... DIK_9</code>	0 through 9 keys
<code>DIK_F1 ... DIK_F15</code>	F1 through F15 keys, if they exist
<code>DIK_NUMPAD0 ... DIK_NUMPAD9</code>	Number pad keys. The keys are the same regardless of whether or not Num Lock is on.
<code>DIK_ESCAPE</code>	Esc key
<code>DIK_MINUS</code>	– key on the top row
<code>DIK_EQUALS</code>	= key on the top row
<code>DIK_BACK</code>	Backspace key
<code>DIK_TAB</code>	Tab key
<code>DIK_LBRACKET</code>	[(left bracket) key
<code>DIK_RBRACKET</code>] (right bracket) key
<code>DIK_RETURN</code>	Return key
<code>DIK_LCONTROL</code>	Left-side Ctrl key
<code>DIK_SEMICOLON</code>	; (semicolon) key
<code>DIK_APOSTROPHE</code>	' (apostrophe) key

DIK_GRAVE	` (grave accent) key; usually the same as the tilde (~) key
DIK_LSHIFT	Left-side Shift key
DIK_BACKSLASH	\ (backslash) key
DIK_COMMA	, (comma) key
DIK_PERIOD	. (period) key
DIK_SLASH	/ (forward slash) key
DIK_RSHIFT	Right-side Shift key
DIK_MULTIPLY	* key on numeric pad
DIK_LMENU	Left-side Alt key
DIK_SPACE	Spacebar
DIK_CAPITAL	Caps Lock key
DIK_NUMLOCK	Num Lock key
DIK_SCROLL	Scroll Lock key
DIK_SUBTRACT	– sign on keypad
DIK_ADD	+ sign on keypad
DIK_DECIMAL	. sign on keypad
DIK_NUMPADENTER	Enter on keypad
DIK_RCONTROL	Right-side Ctrl key
DIK_DIVIDE	/ sign on keypad
DIK_SYSRQ	SysRq (same as PrtScrn) key
DIK_RMENU	Right-side Alt key
DIK_PAUSE	Pause key
DIK_HOME	Home key (if there is a set separate from the keypad)
DIK_UP	Up arrow
DIK_PRIOR	PgUp key (if there is a set separate from the keypad)
DIK_LEFT	Left arrow
DIK_RIGHT	Right arrow
DIK_END	End key (if there is a set separate from the keypad)
DIK_DOWN	Down arrow
DIK_NEXT	PgDn key (if there is a set separate from the keypad)
DIK_INSERT	Insert key (if there is a set separate from the keypad)
DIK_DELETE	Delete key (if there is a set separate from the keypad)
DIK_LWIN	Left-side Windows key

DIK_RWIN	Right-side Windows key
DIK_APPS	Application key

Cooperative Levels

DirectInput devices have a concept of a cooperative level, since they are shared by all applications using the system. Setting the cooperative level is the first thing that you should do upon successful creation of an IDirectInputDevice8 interface. The call to set the cooperative level is:

```
HRESULT IDirectInputDevice8::SetCooperativeLevel(
    HWND hwnd,
    DWORD dwFlags
);
```

hwnd	Handle to the window of the application that created the object.
dwFlags	<p>A set of flags describing the cooperative level desired. Can be a combination of the following:</p> <ul style="list-style-type: none"> • DISCL_BACKGROUND—When this flag is set, the application may acquire the device at any time, even if it is not the currently active application. • DISCL_EXCLUSIVE—Application requests exclusive access to the input device. This prevents other applications from simultaneously using the device (for example, Windows itself). If the mouse device is set to exclusive mode, Windows stops sending mouse messages and the cursor disappears. • DISCL_FOREGROUND—When this flag is set, the device is automatically unacquired when the window moves to the background. It can only be reacquired when the application moves to the foreground. • DISCL_NONEXCLUSIVE—Application requests non-exclusive access to the input device. This way it doesn't interfere with the other applications that are simultaneously using the device (for example, Windows itself). • DISCL_NOWINKEY—Disables the use of the Windows key. This prevents the user from accidentally being knocked out of an exclusive application by pressing the Windows key.

All devices must set either **DISCL_FOREGROUND** or **DISCL_BACKGROUND** (but not both), as well as either **DISCL_EXCLUSIVE** or **DISCL_NONEXCLUSIVE** (but not both).

Application Focus and Devices

If you ever can't get the device state from a device, chances are access to it has been lost. For example, when the application doesn't have focus you can't grab the state of the keyboard. The application class will automatically detect when it loses focus and stop the input code from polling the devices until focus is regained. When you get focus, you need to reacquire the device before you can start requesting its state. That is done using the

parameter-free function `IDirectInputDevice8::Acquire()`. You'll see `Acquire()` scattered throughout the input code for both the keyboard and the mouse.

The DirectInput Object

The `DirectInput` object (which has the interface `IDirectInput8`) doesn't have a clear tie to a physical device as the `Direct3D` device object did. However, you need it to enumerate and create available devices.

To create the `DirectInput` object, you use the global function `DirectInput8Create()`, which wraps up all the necessary COM work.

```
HRESULT WINAPI DirectInput8Create(  
    HINSTANCE hinst,  
    DWORD dwVersion,  
    REFIID riidIIF,  
    LPVOID *ppvOut,  
    LPUNKNOWN punkOuter  
);
```

hinst	Handle to the instance of the application that is creating the <code>DirectInput</code> object.
dwVersion	The version number of the <code>DirectInput</code> object that you want to create. You should specify <code>DIRECTINPUT_VERSION</code> for this parameter.
riidIIF	An identifier for the interface you want to create. Specify <code>IID_IDirectInput8</code> for this parameter and you won't go wrong.
ppvOut	Address of a pointer that will receive the address of the newly created interface.
punkOuter	Used for COM aggregation—just specify <code>NULL</code> .

Once you have the `DirectInput` interface, you can use it to enumerate and create devices. Device creation is done using `IDirectInput8::CreateDevice()`.

Implementing DirectInput with cInputLayer

Due to the small subset of the total `DirectInput` functionality I'm showing you, the code to handle `DirectInput` is very simple. Adding support for simple joysticks wouldn't be too much harder, but implementing a robust system that could enumerate device objects and assign tasks to each of them would take considerably more work.

The input layer constructs and holds onto a mouse object and a keyboard object (`cMouse` and `cKeyboard`, respectively). Both the mouse and the keyboard can have *listeners*, or classes that are notified when events happen. To make a class a listener, two things must happen. First, the class must implement the `iKeyboardReceiver` interface (for keyboards) and/or

the `iMouseReceiver` interface (for mouse devices). Second, it must tell the keyboard or mouse to make itself a receiver. This can be done by calling `cKeyboard::SetReceiver()` or `cMouse::SetReceiver()`. Just pass in the address of the class that wishes to become a receiver. Here are the interfaces:

```
/**
 * Any object that implements this interface can receive input
 * from the keyboard.
 */
struct iKeyboardReceiver
{
    virtual void KeyUp( int key ) = 0;
    virtual void KeyDown( int key ) = 0;
};

/**
 * Any object that implements this interface can receive input
 * from the mouse.
 */
struct iMouseReceiver
{
    virtual void MouseMoved( int dx, int dy ) = 0;
    virtual void MouseButtonUp( int button ) = 0;
    virtual void MouseButtonDown( int button ) = 0;
};
```

The input layer is another system object, and like the others can only have one instance. This condition is validated in the constructor.

```

/*****
 *           Advanced 3D Game Programming with DirectX 10.0
 * ****
 *
 * See license.txt for modification and distribution information
 * copyright (c) 2007 by Peter Walsh, Wordware
 *****/
#ifdef _INPUTLAYER_H
#define _INPUTLAYER_H

#include <input.h>
#include "Keyboard.h"
#include "Mouse.h"

class cInputLayer
{
    cKeyboard      *m_pKeyboard;
    cMouse         *m_pMouse;

    // The DI8 object
    LPDIRECTINPUT8  m_pDI;

    static cInputLayer *m_pGlobalILayer;

    cInputLayer(
```

```
        HINSTANCE hInst,
        HWND hWnd,
        bool bExclusive,
        bool bUseKeyboard = true,
        bool bUseMouse = true );

public:

    virtual ~CInputLayer();

    CKeyboard *GetKeyboard()
    {
        return m_pKeyboard;
    }

    CMouse *GetMouse()
    {
        return m_pMouse;
    }

    void UpdateDevices();

    static CInputLayer *GetInput()
    {
        return m_pGlobalILayer;
    }

    LPDIRECTINPUT8 GetDI()
    {
        return m_pDI;
    }

    void SetFocus(); // called when the app gains focus
    void KillFocus(); // called when the app must release focus

    static void Create(
        HINSTANCE hInst,
        HWND hWnd,
        bool bExclusive,
        bool bUseKeyboard = true,
        bool bUseMouse = true )
    {
        // everything is taken care of in the constructor
        new CInputLayer(
            hInst,
            hWnd,
            bExclusive,
            bUseKeyboard,
            bUseMouse );
    }
};

inline CInputLayer *Input()
{
    return CInputLayer::GetInput();
}
```

```

}

#endif // _INPUTLAYER_H

/*****
 *      Advanced 3D Game Programming with DirectX 10.0
 *      * * * * *
 *
 *      See license.txt for modification and distribution information
 *      copyright (c) 2007 by Peter Walsh, Wordware
 *****/
#include "stdafx.h"
#include "InputLayer.h"
#include "Keyboard.h"
#include "Mouse.h"
#include "Application.h"
#include "Window.h"

cInputLayer *cInputLayer::m_pGlobalILayer = NULL;

cInputLayer::cInputLayer(
    HINSTANCE hInst,
    HWND hWnd,
    bool bExclusive,
    bool bUseKeyboard,
    bool bUseMouse )
{

    m_pKeyboard = NULL;
    m_pMouse = NULL;

    if( m_pGlobalILayer )
    {
        throw cGameError("cInputLayer already initialized!\n");
    }
    m_pGlobalILayer = this;

    HRESULT hr;

    /**
     * Create the DI8 object
     */

    hr = DirectInput8Create( hInst, DIRECTINPUT_VERSION,
                           IID_IDirectInput8, (void**)&m_pDI, NULL );
    if( FAILED(hr) )
    {
        throw cGameError("DirectInput8 object could not be created\n");
    }

    try
    {
        if( bUseKeyboard )
        {

```

```
        m_pKeyboard = new cKeyboard( hWnd );
    }
    if( bUseMouse )
    {
        m_pMouse = new cMouse( hWnd, bExclusive );
    }
}
catch( ... )
{
    SafeRelease( m_pDI );
    throw;
}

}

cInputLayer::~cInputLayer()
{
    if( m_pDI )
    {
        if( m_pKeyboard )
        {
            delete m_pKeyboard; // this does all the de-init.
        }

        if( m_pMouse )
        {
            delete m_pMouse; // this does all the de-init.
        }
        SafeRelease( m_pDI );
    }
    m_pGlobalILayer = NULL;
}

void cInputLayer::UpdateDevices()
{
    if( m_pKeyboard )
    {
        m_pKeyboard->Update();
    }
    if( m_pMouse )
    {
        m_pMouse->Update();
    }
}

void cInputLayer::SetFocus()
{
    if( m_pKeyboard )
    {
        m_pKeyboard->ClearTable();
    }
    if( m_pMouse )
    {

```

```

        m_pMouse->Acquire();
    }
}

void cInputLayer::KillFocus()
{
    if( m_pKeyboard )
    {
        m_pKeyboard->ClearTable();
    }
    if( m_pMouse )
    {
        m_pMouse->UnAcquire();
    }
}

```

The keyboard object pretty much wraps around the `IDirectInputDevice8` interface, while providing the listener interface for an easy way for classes to listen to keys that get pressed. If you don't want to use listeners, just call the `Poll()` method on the keyboard object to find the state of a certain key at the time it was last checked.

```

/*****
 *           Advanced 3D Game Programming with DirectX 10.0
 * *****/
*
*   See license.txt for modification and distribution information
*   copyright (c) 2007 by Peter Walsh, Wordware
*****/
#ifdef _KEYBOARD_H
#define _KEYBOARD_H

#include <memory.h>
#include <directinput.h>

class cInputLayer;

/**
 * Any object that implements this interface can receive input
 * from the keyboard.
 */
struct iKeyboardReceiver
{
    virtual void KeyUp( int key ){};
    virtual void KeyDown( int key ){};
};

class cKeyboard
{
    // The DInput device used to encapsulate the keyboard
    LPDIRECTINPUTDEVICE8 m_pDevice;

```

```

char m_keyState[256];

iKeyboardReceiver *m_pTarget;

public:

    void ClearTable()
    {
        memset( m_keyState, 0, sizeof(char)*256 );
    }

    cKeyboard( HWND hWnd );
    ~cKeyboard();

    // Poll to see if a certain key is down
    bool Poll( int key );

    // Use this to establish a KeyboardReceiver as the current input focus
    void SetReceiver( iKeyboardReceiver *pTarget );

    eResult Update();
};

#endif // _KEYBOARD_H

/*****
 *      Advanced 3D Game Programming with DirectX 10.0
 *      * * * * *
 *
 *      See license.txt for modification and distribution information
 *      copyright (c) 2007 by Peter Walsh, Wordware
 *****/
#include "stdafx.h"
#include "InputLayer.h"
#include "window.h"

#include <stack>
using namespace std;

#include "Keyboard.h"

cKeyboard::cKeyboard( HWND hWnd )
{
    m_pTarget = NULL;

    HRESULT hr;

    /**
     * Get the DInput interface pointer
     */
    LPDIRECTINPUT8 pDI = Input()->GetDInput();

    /**
     * Create the keyboard device

```

```

    */
    hr = Input()->GetDInput()->CreateDevice( GUID_SysKeyboard, &m_pDevice,
                                             NULL );

    if( FAILED(hr) )
    {
        throw cGameError("Keyboard could not be created\n");
    }

/**
 * Set the keyboard data format
 */
    hr = m_pDevice->SetDataFormat(&c_dfDIKeyboard);
    if( FAILED(hr) )
    {
        SafeRelease( m_pDevice );
        throw cGameError("Keyboard could not be created\n");
    }

/**
 * Set the cooperative level
 */
    hr = m_pDevice->SetCooperativeLevel(
        hWnd,
        DISCL_FOREGROUND | DISCL_NONEXCLUSIVE);
    if( FAILED(hr) )
    {
        SafeRelease( m_pDevice );
        throw cGameError("Keyboard coop level could not be changed\n");
    }

    memset( m_keyState, 0, 256*sizeof(bool) );
}

cKeyboard::~cKeyboard()
{
    if( m_pDevice )
    {
        m_pDevice->UnAcquire();
        SafeRelease( m_pDevice );
    }
}

void cKeyboard::SetReceiver( iKeyboardReceiver *pTarget )
{
    // Set the new target.
    m_pTarget = pTarget;
}

bool cKeyboard::Poll( int key )
{
    // stuff goes in here.
    if( m_keyState[key] & 0x80 )
        return true;
}

```



```
        return false;
    }

    HRESULT CKeyboard::Update()
    {
        BYTE newState[256];
        HRESULT hr;

        hr = m_pDevice->Poll();
        hr = m_pDevice->GetDeviceState(sizeof(newState), (LPVOID)&newState);

        if( FAILED(hr) )
        {
            hr = m_pDevice->Acquire();
            if( FAILED(hr) )
            {
                return resFailed;
            }

            hr = m_pDevice->Poll();
            hr = m_pDevice->GetDeviceState(sizeof(newState), (LPVOID)&newState);
            if( FAILED(hr) )
            {
                return resFailed;
            }
        }

        if( m_pTarget )
        {
            int i;
            for( i=0; i< 256; i++ )
            {
                if( m_keyState[i] != newState[i] )
                {
                    // Something happened to this key since last checked
                    if( !(newState[i] & 0x80) )
                    {
                        // It was released
                        m_pTarget->KeyUp( i );
                    }
                    else
                    {
                        // Do nothing; it was just pressed, it'll get a keydown
                        // in a bit, and we don't want to send the signal to
                        // the input target twice
                    }
                }
            }

            // copy the state over (we could do a memcpy at the end, but this
            // will have better cache performance)
            m_keyState[i] = newState[i];

            if( Poll( i ) )
```

```

        {
            // It was pressed
            m_pTarget->KeyDown( i );
        }
    }
}
else
{
    // copy the new states over.
    memcpy( m_keyState, newState, 256 );
}

return resAllGood;
}

```

The mouse object is almost identical in function to the keyboard object.

```

/*****
 *      Advanced 3D Game Programming with DirectX 10.0
 *      * * * * *
 *
 *      See license.txt for modification and distribution information
 *      copyright (c) 2007 by Peter Walsh, Wordware
 *****/
#ifndef _MOUSE_H
#define _MOUSE_H

#include <dinput.h>

/**
 * Any object that implements this interface can receive input
 * from the mouse.
 */
struct IMouseReceiver
{
    virtual void MouseMoved( int dx, int dy ){};
    virtual void MouseButtonUp( int button ){};
    virtual void MouseButtonDown( int button ){};
};

class CMouse
{
    LPDIRECTINPUTDEVICE8 m_pDevice;

    DIMOUSESTATE m_lastState;

    IMouseReceiver *m_pTarget;

public:

    CMouse( HWND hWnd, bool bExclusive );
    ~CMouse();

    /**
     * Use this to establish a MouseReceiver as the current

```

```

        * input focus
        */
void SetReceiver( iMouseReceiver *pTarget );

eResult Update();

eResult Acquire();
void UnAcquire();
};

#endif // _MOUSE_H

/*****
 *      Advanced 3D Game Programming with DirectX 10.0
 *      * * * * *
 *
 *      See license.txt for modification and distribution information
 *      copyright (c) 2007 by Peter Walsh, Wordware
 *****/
#include "stdafx.h"
#include "InputLayer.h"
#include "Window.h"

#include "Mouse.h"

cMouse::cMouse( HWND hWnd, bool bExclusive )
{
    m_pTarget = NULL;

    HRESULT hr;

    /**
     * Create the device
     */
    hr = Input()->GetDInput()->CreateDevice( GUID_SysMouse,
                                             &m_pDevice, NULL );

    if( FAILED(hr) )
    {
        throw cGameError("[cMouse::Init]: Couldn't create the device!\n");
    }

    /**
     * Set the data format
     */
    hr = m_pDevice->SetDataFormat(&c_dfDIMouse);
    if( FAILED(hr) )
    {
        SafeRelease( m_pDevice );
        throw cGameError("[cMouse::Init]: SetDataFormat failed\n");
    }

    /**
     * Set the cooperative level
     */
    if( bExclusive )

```

```

    {
        hr = m_pDevice->SetCooperativeLevel( hWnd, DISCL_EXCLUSIVE |
                                                DISCL_NOWINKEY | DISCL_FOREGROUND );
    }
    else
    {
        hr = m_pDevice->SetCooperativeLevel( hWnd, DISCL_NONEXCLUSIVE |
                                                DISCL_FOREGROUND );
    }

    if( FAILED(hr) )
    {
        SafeRelease( m_pDevice );
        throw cGameError("[cMouse::Init]: SetCooperativeLevel failed\n");
    }

    m_lastState.lX = 0;
    m_lastState.lY = 0;
    m_lastState.lZ = 0;
    m_lastState.rgbButtons[0] = 0;
    m_lastState.rgbButtons[1] = 0;
    m_lastState.rgbButtons[2] = 0;
    m_lastState.rgbButtons[3] = 0;
}

cMouse::~cMouse()
{
    if( m_pDevice )
    {
        m_pDevice->UnAcquire();
        SafeRelease( m_pDevice );
    }
}

void cMouse::SetReceiver( iMouseReceiver *pTarget )
{
    m_pTarget = pTarget;
}

eResult cMouse::Update()
{
    DIMOUSESTATE currState;
    HRESULT hr;

    hr = m_pDevice->Poll();
    hr = m_pDevice->GetDeviceState( sizeof(DIMOUSESTATE),
                                    (void*)&currState );

    if( FAILED(hr) )
    {
        hr = m_pDevice->Acquire();
        if( FAILED(hr) )

```

```
{
    return resFailed;
}

hr = m_pDevice->Poll();
hr = m_pDevice->GetDeviceState( sizeof(DIMOUSESTATE),
                                (void*)&currState );

if( FAILED(hr) )
{
    return resFailed;
}

if( m_pTarget )
{
    int dx = currState.lX;
    int dy = currState.lY;
    if( dx || dy )
    {
        m_pTarget->MouseMoved( dx, dy );
    }
    if( currState.rgbButtons[0] & 0x80 )
    {
        // the button got pressed.
        m_pTarget->MouseButtonDown( 0 );
    }
    if( currState.rgbButtons[1] & 0x80 )
    {
        // the button got pressed.
        m_pTarget->MouseButtonDown( 1 );
    }
    if( currState.rgbButtons[2] & 0x80 )
    {
        // the button got pressed.
        m_pTarget->MouseButtonDown( 2 );
    }
    if( !(currState.rgbButtons[0] & 0x80) && (m_lastState.rgbButtons[0]
                                                & 0x80) )
    {
        // the button got released.
        m_pTarget->MouseButtonUp( 0 );
    }
    if( !(currState.rgbButtons[1] & 0x80) && (m_lastState.rgbButtons[1]
                                                & 0x80) )
    {
        // the button got released.
        m_pTarget->MouseButtonUp( 1 );
    }
    if( !(currState.rgbButtons[2] & 0x80) && (m_lastState.rgbButtons[2]
                                                & 0x80) )
    {
        // the button got released.
        m_pTarget->MouseButtonUp( 2 );
    }
}

m_lastState = currState;
```

```

        return resAllGood;
    }

    HRESULT cMouse::Acquire()
    {
        HRESULT hr = m_pDevice->Acquire();
        if( FAILED(hr) )
        {
            return resFailed;
        }
        return resAllGood;
    }

    void cMouse::UnAcquire()
    {
        m_pDevice->UnAcquire();
    }

```

Additions to *cApplication*

The only addition to *cApplication* is the *InitInput()* call. It initializes both the keyboard and the mouse. The method can be overloaded if this behavior isn't what you want. The code is in the following:

```
cInputLayer::Create(AppInstance(), MainWindow()->GetHWnd(), NULL, true, true);
```

The sample for this section is the same as the sound example, so you'll see it toward the end of the chapter.

Once you have created your *DirectInput* code, you can keep it as it is for most of your projects without too much modification. So if, in the future, you ever need to use code for any of your own projects to deal with input, you can just cut and paste what you have learned in this chapter into your code.

Now it's time to move further into the possibilities of *DirectX* with *DirectSound*.

Sound

There was a time, long ago, when computers didn't have sound cards. Sound cards were add-ons that people bought and installed manually. I clearly remember the first time I played the original *Wolfenstein 3D* on a sound card-enabled machine; after that I ran out and bought one. Sound can totally change the experience of electronic entertainment. Instead of just associating visual images with a virtual experience, adding sound to an application makes it still more immersive, especially if the sound effects are well made.

Before the great move to *DirectX*, using sound was a tricky process for programmers. Usually it involved licensing an expensive and complex

third-party sound library that could interface with the different types of sound cards on the market. These libraries could cost hundreds or thousands of dollars. With the advent of DirectSound, the need for these libraries has all but disappeared. DirectSound is an API that can play sound on any Windows-capable sound card (which is basically all of them).

While the Win32 API has some limited sound-playing functionality, it's not something that is practical for most games. Sounds can't be mixed together, signal processing is nonexistent, and it isn't the fastest thing in the world.

As of DirectX 6.1, the DirectX component called DirectMusic allows applications to dynamically improvise music for games. DirectMusic is really a specialty subject that goes too far off track for this book, so I'm going to limit my discussions to DirectSound.

The Essentials of Sound

Sound itself is a wave of kinetic energy caused by the motion of an object. The wave travels through matter at a speed dependent on the type of matter and temperature (very quickly through solids; through air at 24° C (75° F) it moves at about 1240 kph (775 mph)). Sound waves have energy, so they can cause objects to move; when they hit a solid object, some of the sound is transmitted through the object, some is absorbed, and some is reflected back (the reflecting back is known as echo). When the waves hit an object, they make it vibrate. When the vibrating object is your eardrum, the wave is converted into electric signals in your cochlea and then sent to your brain where it enters consciousness.

The waves are sinusoidal in nature, and they have *amplitude* and *frequency*. The amplitude defines how loud the sound is and is usually measured in decibels (dB). The frequency is how many different wave oscillations fit into one second, measured in hertz (Hz). The frequency of a sound defines what its pitch is; lower-pitched sounds resonate less than higher-pitched sounds. The A above middle C has a wave that resonates 440 times a second, so it has a frequency of 440 Hz.

Sound is additive; that is, if two sounds are going through the air together, both apply their energy to the air molecules around them. When the crests of the sound waves match up, their result is a louder sound, while if opposite crests match up, they cancel each other out. The more things there are creating sound in a room, the more sound there generally is in the room.

On a computer, sound is represented as a stream of discrete samples. Each sample is usually an 8-bit or 16-bit integer, representing the amplitude of the sample. With 16 bits the amplitude can be better approximated, since there is a range of about 65,000 values, instead of only 256 found in 8 bits. Successive samples are played, and when enough samples are put together, they approximate the continuous sound curve well enough that the human ear can't tell the difference. In order to

approximate it well, the sampling rate (number of samples every second) is much higher than the frequency of most audible sounds—for CD-quality sound, 44,100 samples per second are used to approximate the waveform. See Figure 3.1 for what this looks like. The figure shows an extremely magnified waveform; the amount of signal shown would probably account for a few hundredths of a second of sound.



Note: By the way, in case you have any experience with previous versions of DirectSound, there are virtually no changes in DirectSound since DirectX 8.0. The changes are pretty obscure and minor but if you're interested, check out *DirectX C++ Documentation/DirectX Audio/DirectSound/What's New In DirectSound*. There are no changes to any methods or interfaces so your old code should work perfectly with DirectX 10.0. One big difference is that Windows Vista no longer supports hardware acceleration for audio, but for this book that won't affect us. If you still want to use hardware accelerated audio, you should check out OpenAL on the net.




Figure 3.1: Continuous waveforms and a digital approximation

DirectSound Concepts

DirectSound centers around a set of interfaces that perform all the work you need to do. The DirectSound interfaces are summed up in Table 3.3. For now, I'm just going to focus on the ability to play sounds, so don't get too caught up with all of those interfaces.

Table 3.3: The main DirectSound interfaces

IDirectSound8	Used in determining the capabilities of the sound card and creating buffers for playback.
IDirectSoundBuffer8	A buffer used to hold onto the data for a playable sound.
IDirectSound3DBuffer8	A buffer used to contain a 3D sound. Has additional information like distance, position, projection cones, and so forth.
IDirectSound3DListener8	An object used to represent a 3D listener. Depending on the location and direction of the listener in 3D space, 3D buffers sound different.
IDirectSoundCapture8	Interface used to create capture buffers.
IDirectSoundCaptureBuffer8	Buffer used to hold sound data recorded from a device such as a microphone.
IDirectSoundNotify8	Object that can be used to notify an application when a sound has finished playing.
IKsPropertySet8	An interface used by sound card manufacturers to add special abilities to their driver without needing to extend the spec. This is for pretty hard-core stuff, and beyond the scope of this book.

DirectSound Buffers

DirectSound buffers are your main tools in DirectSound. They are very similar to the textures used in Direct3D. Just like textures, in order to access their data you need to lock them, and then you unlock them when you're finished. Because the DirectSound driver can operate asynchronously from the user application, care must be taken that no application is reading data when another is reading from it, or vice versa.

There are two kinds of buffers in DirectSound: the *primary buffer* and *secondary buffers*. The primary buffer represents the sound that is currently playing on the card. There is a secondary buffer for each sound effect an application wants to play. Secondary sound buffers are mixed together into the primary buffer and play out through the speakers. By using the *mixer* you get multiple sound effects to play at once. DirectSound has a well-optimized piece of code that can mix a bunch of secondary sound buffers together, and many sound cards can perform this operation in the hardware automatically (except on Vista). Vista has a new sound driver model that is incompatible with DirectSound for the time being, which means all mixing has to be done in software unless you use an intermediate driver like OpenAL.

One key difference between Direct3D textures and DirectSound buffers is that buffers are conceptually circular. When a sound effect is playing, the play marker loops around to the beginning of the buffer when it reaches the end, unless you tell it to do otherwise. The play marker is a

conceptual marker in the buffer that represents where sound data is being retrieved.

Just like textures, buffers are created by filling out a description of what you want in the buffer. The structure used to describe a DirectSound buffer is called `DSBUFFERDESC`:

```
typedef struct DSBUFFERDESC {
    DWORD dwSize;
    DWORD dwFlags;
    DWORD dwBufferBytes;
    DWORD dwReserved;
    LPWAVEFORMATEX lpwfxFormat;
    GUID guid3DAlgorithm;
} DSBUFFERDESC;
```

<code>dwSize</code>	Size of the structure; set this to <code>sizeof(DSBUFFERDESC)</code> .
<code>dwFlags</code>	<p>Flags that describe the capabilities or desired capabilities of the buffer. Can be one or more of the following:</p> <ul style="list-style-type: none"> • <code>DSBCAPS_CTRL3D</code>—The buffer requires 3D control. It may be a primary or secondary buffer. • <code>DSBCAPS_CTRLFREQUENCY</code>—The buffer requires the ability to control its frequency. • <code>DSBCAPS_CTRLPAN</code>—The buffer requires the ability to control panning. • <code>DSBCAPS_CTRLPOSITIONNOTIFY</code>—The buffer requires position notification. • <code>DSBCAPS_CTRLVOLUME</code>—The buffer requires the ability to control its volume. • <code>DSBCAPS_GETCURRENTPOSITION2</code>—Any calls to <code>GetCurrentPosition()</code> should use the new behavior of putting the read position where it is actually reading. The old behavior put it right behind the write position. The old behavior was also only on emulated DirectSound devices. • <code>DSBCAPS_GLOBALFOCUS</code>—Like <code>DSBCAPS_STICKYFOCUS</code>, except the buffer can also be heard when other DirectSound applications have focus. The exception is applications that request exclusive access to the sound cards. All other global sounds will be muted when those applications have focus. • <code>DSBCAPS_LOCDEFER</code>—The buffer can be assigned to either hardware or software playback, depending on the mood of the driver. This flag must be set if the voice management features in version 9.0 are to be used. • <code>DSBCAPS_LOCHARDWARE</code>—Forces the buffer to be mixed in hardware. The application must make sure there is a mixing channel available for the buffer. If there isn't enough memory on the card, or the card doesn't support hardware mixing, calling <code>CreateSoundBuffer()</code> will fail. Using this flag on Windows Vista will always cause the function to fail, since Vista does not support hardware accelerated audio. See <code>OpenAL</code> if you need accelerated audio.

- **DSBCAPS_LOCSOFTWARE**—Forces the buffer to be mixed in software. Required on Vista.
- **DSBCAPS_MUTE3DATMAXDISTANCE**—This flag applies to advanced 3D sound buffers.
- **DSBCAPS_PRIMARYBUFFER**—Indicates that the buffer is the single and only (primary) buffer for the sound card. A secondary buffer is created if this flag is not set.
- **DSBCAPS_STATIC**—Informs the driver that the buffer will be filled once and played many times. This makes the driver more likely to put the buffer in hardware memory.
- **DSBCAPS_STICKYFOCUS**—Changes the focus behavior of a sound buffer. Buffers created with sticky focus aren't muted when the user switches to a non-DirectSound application. This is useful for applications like TV cards, where the user wants to hear what is happening while using another application. However, if the user switches to another DirectSound application, all sound effects are muted.

dwBufferBytes	Size of the buffer, in bytes. When you create the primary surface, this parameter should be set to zero.
dwReserved	Reserved for use by DirectSound; don't use.
lpwfxFormat	Pointer to a WAVEFORMATEX structure describing the format of the wave data in the buffer. This is analogous to the pixel formats describing the format of the pixels in Direct3D textures.
guid3DAlgorithm	GUID that defines the two-speaker virtualization algorithm to be used for software rendering. This GUID is ignored unless the buffer needs 3D control (set by the DSBCAPS_CTRL3D flag). See the documentation for a listing of the available GUIDs for this parameter.

The `lpwfxFormat` parameter of the sound buffer description is a pointer to a `WAVEFORMATEX` structure. The reason why there's no DS prefixing the structure is because it isn't a DirectSound structure, but instead is one used by Windows for its sound playback work.

```
typedef struct WAVEFORMATEX {
    WORD wFormatTag;
    WORD nChannels;
    DWORD nSamplesPerSec;
    DWORD nAvgBytesPerSec;
    WORD nBlockAlign;
    WORD wBitsPerSample;
    WORD cbSize;
} WAVEFORMATEX;
```

wFormatTag	A tag describing the content of the sound data. If the data is compressed, this tag will correspond to the particular method that was used to compress it. For non-compressed data, this will be set to the constant <code>WAVE_FORMAT_PCM</code> .
nChannels	The number of separate audio channels for the sound data. For monaural sound there is one channel; for stereo sound there are two.
nSamplesPerSec	The number of samples per second. For CD-quality audio this is about 44,100; for radio quality it is about 22,050.
nAvgBytesPerSec	The required data throughput to play the sound. This is here so you can deal with compressed sound files.
nBlockAlign	Block alignment in bytes. Essentially this is the amount of data for one sample. If you had two channels of audio and 16 bits (2 bytes) per sample, this would be $2 * 2 = 4$ bytes.
wBitsPerSample	The number of bits for each discrete sample. This is generally either 8 or 16.
cbSize	The size of any extra info that is appended to the structure. This is only used by compressed sound formats.

Operations on Sound Buffers

Once you have created a buffer and filled it with the appropriate data, you would, of course, like to play it. The `Play()` method on the buffer interface plays the sound buffer on the primary buffer. The sound can be stopped by calling the `Stop()` method, which takes no parameters.

```
HRESULT IDirectSoundBuffer8::Play(
    DWORD dwReserved1,
    DWORD dwPriority,
    DWORD dwFlags
);
HRESULT IDirectSoundBuffer8::Stop();
```

dwReserved1	Reserved parameter; must be set to 0.
dwPriority	The priority of the sound. This is used by the sound manager in the event that it needs to evict a playing sound (it evicts the one with the lowest priority). The valid range is anywhere from 0x0 to 0xFFFFFFFF. 0 has the lowest priority. This value shouldn't be used if the buffer wasn't created with the <code>LOCDEFER</code> flag, and should be left as 0.
dwFlags	A set of flags describing the method's behavior. They are: <ul style="list-style-type: none"> • <code>DSBPLAY_LOOPING</code>—Whenever the end of the buffer is reached, <code>DirectSound</code> wraps to the beginning of the buffer and continues playing it. This is useful for sounds like engine hums. The sound effect continues playing until it is explicitly shut off using <code>Stop()</code>.

- **DSBPLAY_LOCHARDWARE**—This flag only affects the buffer created with the **DSBCAPS_LOCDEFER** flag. It forces the buffer to be played in the hardware. If there aren't any voices available and no **TERMINATEBY_*** flags are set, **Play()** will fail. This flag shouldn't be used with **DSBPLAY_LOCSOFTWARE**. Note that it does not work on Windows Vista.
- **DSBPLAY_LOCSOFTWARE**—This flag only affects the buffer created with the **DSBCAPS_LOCDEFER** flag. It forces the buffer to be played in software. If there aren't any voices available and no **TERMINATEBY_*** flags are set, **Play()** will fail. This flag shouldn't be used with **DSBPLAY_LOCHARDWARE**. If neither **LOCSOFTWARE** or **LOCHARDWARE** is specified, the location for playback will be decided by the sound driver, depending on the available resources.
- **DSBPLAY_TERMINATEBY_TIME**—Setting this flag enables the buffer to steal resources from another buffer. The driver is forced to play the buffer in hardware. If no hardware voices are available, the driver chooses a buffer to remove, choosing the buffer that has the least amount of time left to play. The only candidate buffers for removal are ones created with the **DSBCAPS_LOCDEFER** flag.
- **DSBPLAY_TERMINATEBY_DISTANCE**—This flag is only relevant to 3D buffers, which are beyond the scope of this book.
- **DSBPLAY_TERMINATEBY_PRIORITY**—Setting this flag enables the buffer to steal resources from another buffer. The driver is forced to play the buffer in hardware. If no hardware voices are available, the driver chooses a buffer to remove, choosing the buffer that has the lowest priority. The only candidate buffers for removal are ones created with the **DSBCAPS_LOCDEFER** flag.

Unfortunately, there is only one play marker per sound buffer, so you can't play the same sound twice at the same time. However, the code I'll show you can clone the sound effect into a new buffer and play the new effect, so you can have multiple sounds of the same type playing at the same time. To implement this, however, you need to know if the sound buffer is playing at any point in time. You can do this using the **GetStatus()** method on the sound buffer interface:

```
HRESULT IDirectSoundBuffer8::GetStatus(
    LPDWORD lpdwStatus
);
```

lpdwStatus Pointer to a **DWORD** that will be filled with the status of the sound buffer. If the function succeeds, the **DWORD** can check to see if any of the following flags are set:

- **DSBSTATUS_BUFFERLOST**—The sound buffer was lost. Before it can be played or locked, it must be restored using the **Restore()** method on the **DirectSoundBuffer**. **Restore()** takes no parameters and reallocates the required memory for a **DirectSound** buffer.

- **DSBSTATUS_LOOPING**—The buffer is playing and also looping. It won't stop until the `Stop()` method is called on it.
- **DSBSTATUS_PLAYING**—The buffer is currently playing. The buffer is stopped if this flag isn't set.
- **DSBSTATUS_LOCSOFTWARE**—The buffer is playing from system RAM. This flag is only meaningful for buffers that were created with the **DSBCAPS_LOCDEFER** flag.
- **DSBSTATUS_LOCHARDWARE**—The buffer is playing on the sound card's memory. This flag is only meaningful for buffers that were created with the **DSBCAPS_LOCDEFER** flag. This does not work on Windows Vista.
- **DSBSTATUS_TERMINATED**—The buffer was terminated by the sound logic.

To play a buffer with anything meaningful in it, you're going to need to fill it with something. Unfortunately, DirectSound doesn't have the ability to automatically load WAV files, so you have to do that yourself. When you load the file and get the data, you put it into the sound buffer by locking it and getting a pointer to the buffer to write into. This is done using the `Lock()` method on the sound buffer interface.

```
HRESULT IDirectSoundBuffer8::Lock(
    DWORD dwWriteCursor,
    DWORD dwWriteBytes,
    LPVOID lp1pvAudioPtr1,
    LPDWORD lpdwAudioBytes1,
    LPVOID lp1pvAudioPtr2,
    LPDWORD lpdwAudioBytes2,
    DWORD dwFlags
);
```

<code>dwWriteCursor</code>	Offset from the start of the buffer (in bytes) to where the lock should begin.
<code>dwWriteBytes</code>	Number of bytes that should be locked. Remember that sound buffers are circular, conceptually. If more bytes are requested than are left in the file, the lock continues at the beginning of the buffer.
<code>lp1pvAudioPtr1</code>	Pointer to be filled with the requested data pointer of the lock.
<code>lpdwAudioBytes1</code>	Pointer to be filled with the number of bytes of the first data block. This may or may not be the same as <code>dwWriteBytes</code> , depending on whether or not the lock wrapped to the beginning of the sound buffer.
<code>lp1pvAudioPtr2</code>	Pointer to be filled with the secondary data pointer of the lock. This member is only set if the memory requested in the lock wrapped to the beginning of the buffer (it will be set to the beginning of the buffer). If the lock did not require a wrap, this pointer will be set to <code>NULL</code> .

lpdwAudioBytes2	Pointer to be filled with the number of bytes of the second data block. If the lock required a wrap, this will be the number of bytes left over after the wrap around.
dwFlags	A set of flags modifying the behavior of the Lock() method: <ul style="list-style-type: none">• DSBLOCK_FROMWRITECURSOR—Locks from the current write cursor in the buffer.• DSBLOCK_ENTIREBUFFER—Locks the entire sound buffer. The dwWriteBytes parameter is ignored and can be set to zero.

To unlock a sound buffer after filling it, just call the Unlock() method on it. This allows other concurrent tasks on the machine, like the sound hardware, to access the sound buffer's data bits.

```
HRESULT IDirectSoundBuffer8::Unlock(  
    LPVOID lpvAudioPtr1,  
    DWORD dwAudioBytes1,  
    LPVOID lpvAudioPtr2,  
    DWORD dwAudioBytes2  
);
```

lpvAudioPtr1	Pointer to the first block of data to unlock. This must be the same value that was given by Lock().
dwAudioBytes1	Length of the first block of data to unlock. This must be the same value that was given by Lock().
lpvAudioPtr2	Pointer to the second block of data to unlock. This must be the same value that was given by Lock().
dwAudioBytes2	Length of the second block of data to unlock. This must be the same value that was given by Lock().

Loading WAV Files

Call me old fashioned, but I try to avoid reinventing any wheels I can. One I distinctly do not want to reinvent is the WAV-file-loading wheel. The DirectX SDK comes with code to load a WAV file and create a DirectSound buffer, and I'm going to use it verbatim here. For an overview of how it works, here is the source and header file:

```
//-----  
// File: WavRead.h  
//  
// Desc: Support for loading and playing Wave files using DirectSound sound  
//       buffers.  
//  
// Copyright (c) 1999 Microsoft Corp. All rights reserved.  
//-----  
#ifndef WAVE_READ_H
```

```

#define WAVE_READ_H

#include <mmreg.h>
#include <mmsystem.h>

//-----
// Name: class CWaveSoundRead
// Desc: A class to read in sound data from a Wave file
//-----
class CWaveSoundRead
{
public:
    WAVEFORMATEX *m_pwfx;        // Pointer to WAVEFORMATEX structure
    HMMIO          m_hmmioIn;    // MM I/O handle for the WAVE
    MMCKINFO       m_ckIn;       // Multimedia RIFF chunk
    MMCKINFO       m_ckInRiff;   // Use in opening a WAVE file

public:
    CWaveSoundRead();
    ~CWaveSoundRead();

    HRESULT Open( CHAR *strFilename );
    HRESULT Reset();
    HRESULT Read( UINT nSizeToRead, BYTE *pbData, UINT *pnSizeRead );
    HRESULT Close();

};

#endif WAVE_READ_H

//-----
// File: WavRead.cpp
//
// Desc: Wave file support for loading and playing Wave files using DirectSound
//       buffers.
//
// Copyright (c) 1999 Microsoft Corp. All rights reserved.
//-----
#include <windows.h>
#include "WavRead.h"

//-----
// Defines, constants, and global variables
//-----
#define SAFE_DELETE(p) { if(p) { delete (p);    (p)=NULL; } }
#define SAFE_RELEASE(p) { if(p) { (p)->Release(); (p)=NULL; } }

```



```

//-----
// Name: ReadMMIO()
// Desc: Support function for reading from a multimedia I/O stream
//-----
HRESULT ReadMMIO( HMMIO hmmioIn, MMCKINFO *pckInRIFF,
    WAVEFORMATEX** ppwfxInfo )
{
    MMCKINFO      ckIn;          // chunk info for general use.
    PCMWAVEFORMAT pcmWaveFormat; // Temp PCM structure to load in.

    *ppwfxInfo = NULL;

    if( ( 0 != mmioDescend( hmmioIn, pckInRIFF, NULL, 0 ) ) )
        return E_FAIL;

    if( (pckInRIFF->ckid != FOURCC_RIFF) ||
        (pckInRIFF->fccType != mmioFOURCC('W', 'A', 'V', 'E') ) )
        return E_FAIL;

    // Search the input file for the 'fmt' chunk.
    ckIn.ckid = mmioFOURCC('f', 'm', 't', ' ');
    if( 0 != mmioDescend(hmmioIn, &ckIn, pckInRIFF, MMIO_FINDCHUNK) )
        return E_FAIL;

    // Expect the 'fmt' chunk to be at least as large as <PCMWaveFormat>;
    // if there are extra parameters at the end, we'll ignore them
    if( ckIn.cksize < (LONG) sizeof(PCMWaveFormat) )
        return E_FAIL;

    // Read the 'fmt' chunk into <pcmWaveFormat>.
    if( mmioRead( hmmioIn, (HPSTR) &pcmWaveFormat,
        sizeof(pcmWaveFormat)) != sizeof(pcmWaveFormat) )
        return E_FAIL;

    // Allocate the waveformatex, but if it's not pcm format, read the next
    // word, and that's how many extra bytes to allocate.
    if( pcmWaveFormat.wf.wFormatTag == WAVE_FORMAT_PCM )
    {
        if( NULL == ( *ppwfxInfo = new WAVEFORMATEX ) )
            return E_FAIL;

        // Copy the bytes from the pcm structure
        // to the waveformatex structure
        memcpy( *ppwfxInfo, &pcmWaveFormat, sizeof(pcmWaveFormat) );
        (*ppwfxInfo)->cbSize = 0;
    }
    else
    {
        // Read in length of extra bytes.
        WORD cbExtraBytes = 0L;
        if( mmioRead( hmmioIn,
            (CHAR*)&cbExtraBytes, sizeof(WORD)) != sizeof(WORD) )
            return E_FAIL;

        *ppwfxInfo = (WAVEFORMATEX*)

```

```

        new CHAR[ sizeof(WAVEFORMATEX) + cbExtraBytes ];
    if( NULL == *ppwfxInfo )
        return E_FAIL;

    // Copy the bytes from the pcm
    // structure to the waveformatex structure
    memcpy( *ppwfxInfo, &pcmWaveFormat, sizeof(pcmWaveFormat) );
    (*ppwfxInfo)->cbSize = cbExtraBytes;

    // Now, read those extra bytes into the structure
    if( mmioRead( hmmioIn, (CHAR*)
        (((BYTE*)&(*ppwfxInfo)->cbSize))+sizeof(WORD)),
        cbExtraBytes ) != cbExtraBytes )
    {
        delete *ppwfxInfo;
        *ppwfxInfo = NULL;
        return E_FAIL;
    }
}

// Ascend the input file out of the 'fmt ' chunk.
if( 0 != mmioAscend( hmmioIn, &ckIn, 0 ) )
{
    delete *ppwfxInfo;
    *ppwfxInfo = NULL;
    return E_FAIL;
}

return S_OK;
}

//-----
// Name: WaveOpenFile()
// Desc: This function will open a wave input file and prepare it for reading,
//       so the data can be easily read with WaveReadFile. Returns 0 if
//       successful, the error code if not.
//-----
HRESULT WaveOpenFile(TCHAR *strFileName,
                    HMMIO *phmmioIn, WAVEFORMATEX** ppwfxInfo,
                    MMCKINFO *pckInRIFF )
{
    HRESULT hr;
    HMMIO hmmioIn = NULL;

    if( NULL == ( hmmioIn = mmioOpen(
        strFileName, NULL, MMIO_ALLOCBUF|MMIO_READ ) ) )
        return E_FAIL;

    if( FAILED( hr = ReadMMIO( hmmioIn, pckInRIFF, ppwfxInfo ) ) )
    {
        mmioClose( hmmioIn, 0 );
        return hr;
    }
}

```

```

    }

    *phmmioIn = hmmioIn;

    return S_OK;
}

//-----
// Name: WaveStartDataRead()
// Desc: Routine has to be called before WaveReadFile as it searches for the
//       chunk to descend into for reading, that is, the 'data' chunk. For
//       simplicity, this used to be in the open routine, but was taken out and
//       moved to a separate routine so there was more control on the chunks
//       that are before the data chunk, such as 'fact', etc...
//-----
HRESULT WaveStartDataRead( HMMIO *phmmioIn, MMCKINFO *pckIn,
                          MMCKINFO *pckInRIFF )
{
    // Seek to the data
    if( -1 == mmioSeek( *phmmioIn, pckInRIFF->dwDataOffset + sizeof(FOURCC),
                       SEEK_SET ) )
        return E_FAIL;

    // Search the input file for for the 'data' chunk.
    pckIn->ckid = mmioFOURCC('d', 'a', 't', 'a');
    if( 0 != mmioDescend( *phmmioIn, pckIn, pckInRIFF, MMIO_FINDCHUNK ) )
        return E_FAIL;

    return S_OK;
}

//-----
// Name: WaveReadFile()
// Desc: Reads wave data from the wave file. Make sure we're descended into
//       the data chunk before calling this function.
//       hmmioIn    - Handle to mmio.
//       cbRead     - # of bytes to read.
//       pbDest     - Destination buffer to put bytes.
//       cbActualRead - # of bytes actually read.
//-----
HRESULT WaveReadFile( HMMIO hmmioIn, UINT cbRead, BYTE *pbDest,
                     MMCKINFO *pckIn, UINT *cbActualRead )
{
    MMIOINFO mmioinfoIn;    // current status of <hmmioIn>

    *cbActualRead = 0;

```

```

        if( 0 != mmioGetInfo( hmmioIn, &mmioinfoIn, 0 ) )
            return E_FAIL;

        UINT cbDataIn = cbRead;
        if( cbDataIn > pckIn->cksize )
            cbDataIn = pckIn->cksize;

        pckIn->cksize -= cbDataIn;

        for( DWORD cT = 0; cT < cbDataIn; cT++ )
        {
            // Copy the bytes from the io to the buffer.
            if( mmioinfoIn.pchNext == mmioinfoIn.pchEndRead )
            {
                if( 0 != mmioAdvance( hmmioIn, &mmioinfoIn, MMIO_READ ) )
                    return E_FAIL;

                if( mmioinfoIn.pchNext == mmioinfoIn.pchEndRead )
                    return E_FAIL;
            }

            // Actual copy.
            *((BYTE*)pbDest+cT) = *((BYTE*)mmioinfoIn.pchNext);
            mmioinfoIn.pchNext++;
        }

        if( 0 != mmioSetInfo( hmmioIn, &mmioinfoIn, 0 ) )
            return E_FAIL;

        *cbActualRead = cbDataIn;
        return S_OK;
    }

//-----
// Name: CWaveSoundRead()
// Desc: Constructs the class
//-----
CWaveSoundRead::CWaveSoundRead()
{
    m_pwfx = NULL;
}

//-----
// Name: ~CWaveSoundRead()
// Desc: Destructs the class
//-----
CWaveSoundRead::~CWaveSoundRead()
{
    Close();
}

```

```

        SAFE_DELETE( m_pwf );
    }

//-----
// Name: Open()
// Desc: Opens a wave file for reading
//-----
HRESULT CWaveSoundRead::Open( TCHAR *strFilename )
{
    SAFE_DELETE( m_pwf );

    HRESULT hr;

    if( FAILED( hr =
        WaveOpenFile( strFilename, &m_hmmioIn, &m_pwf, &m_ckInRiff ) ) )
        return hr;

    if( FAILED( hr = Reset() ) )
        return hr;

    return hr;
}

//-----
// Name: Reset()
// Desc: Resets the internal m_ckIn pointer so reading starts from the
//       beginning of the file again
//-----
HRESULT CWaveSoundRead::Reset()
{
    return WaveStartDataRead( &m_hmmioIn, &m_ckIn, &m_ckInRiff );
}

//-----
// Name: Read()
// Desc: Reads a wave file into a pointer and returns how much read
//       using m_ckIn to determine where to start reading from
//-----
HRESULT CWaveSoundRead::Read( UINT nSizeToRead, BYTE *pbData, UINT *pnSizeRead )
{
    return WaveReadFile( m_hmmioIn, nSizeToRead, pbData, &m_ckIn, pnSizeRead );
}

```

```
//-----
// Name: Close()
// Desc: Closes an open wave file
//-----
HRESULT CWaveSoundRead::Close()
{
    mmioClose( m_hmmioIn, 0 );
    return S_OK;
}
```

Implementing DirectSound with cSoundLayer

The final system layer I'm going to implement in this chapter is the sound layer. The class is called `cSoundLayer`, and has the same restrictions as the graphics and input layers (note that only one instance of the class may exist in any application).

Creating the sound layer is simple enough. The sound layer has the same interface for creation as the graphics and input layers: a static `Create()` method that takes care of the initialization hassles. The `Create()` method for the sound layer is simple enough, and it appears in the following:

```
static void cSoundLayer::Create( HWND hWnd )
{
    new cSoundLayer( hWnd );
}
```

The code that lies inside the `cSoundLayer` constructor is what I'll dissect next in the step-by-step process of setting up DirectSound.

Creating the DirectSound Object

The first step in initializing DirectSound is to actually acquire the interface pointer to the `IDirectSound8` object. To do this, you call the function `DirectSoundCreate8()`.

```
HRESULT WINAPI DirectSoundCreate8(
    LPCGUID lpcGuid,
    LPDIRECTSOUND8 *ppDS,
    LPUNKNOWN pUnkOuter
);
```

lpcGuid	A pointer to a GUID that describes the device you wish to create. While you can enumerate all of the sound devices with <code>DirectSoundEnumerate()</code> , generally there is only one sound card on a machine. To get the default device (which is what you usually want), set this to <code>NULL</code> .
ppDS	A pointer to an <code>LPDIRECTSOUND8</code> interface pointer that will be filled with a valid interface pointer if the function succeeds.
pUnkOuter	Used for COM aggregation; leave this as <code>NULL</code> .

Sample code to create the sound interface appears in the following:

```
LPDIRECTSOUND8 m_pDSound = 0;

// Create IDirectSound using the primary sound device
hr = DirectSoundCreate8( NULL, &m_pDSound, NULL );
if( FAILED(hr) )
{
    // Handle critical error
}
```

Setting the Cooperative Level

After you acquire the interface pointer, the next step is to declare how cooperative you intend on being. Just like DirectInput, this is done using the SetCooperativeLevel() function.

```
HRESULT IDirectSound8::SetCooperativeLevel(
    HWND hwnd,
    DWORD dwLevel
);
```

hwnd	Handle to the window to be associated with the DirectSound object. This should be the primary window.
dwLevel	One of the following flags, describing the desired cooperative level: <ul style="list-style-type: none">• DSSCL_EXCLUSIVE—Grab exclusive control of the sound device. When the application has focus, it is the only audible application.• DSSCL_NORMAL—Smoothest, yet most restrictive cooperative level. The primary format cannot be changed. This is the cooperative level the sound layer uses.• DSSCL_PRIORITY—Like DSSCL_NORMAL except the primary format may be changed.• DSSCL_WRITEPRIMARY—This is the highest possible priority for an application to have. It can't play any secondary buffers, and it has the ability to manually mangle the bits of the primary buffer. Only for the extremely hardcore!

This code will be changing the primary format of the sound buffer, so I'll go ahead and set this to DSSCL_PRIORITY. Sample code to do this appears in the following:

```
// pDSound is a valid LPDIRECTSOUND8 object.
HRESULT hr = pDSound->SetCooperativeLevel( hWnd, DSSCL_PRIORITY );
if( FAILED(hr) )
{
    /* handle error */
}
```

Grabbing the Primary Buffer

Since the sound layer sets the cooperative level's priority, it can do some crazy things like change the format of the primary buffer. Generally it's best to set the primary buffer to the same format that all of your secondary buffers will be in; this makes the mixer's job easier, as it doesn't have to resample any sound effects to be able to mix them into the primary buffer. You can imagine what would happen if you tried to play a 22 kHz sound effect in a 44 kHz buffer without resampling: You would run out of samples twice as soon as you would expect, and the sound effect would have sort of an inhaled-helium quality to it.

To change the format of the primary buffer, you just need to grab it using `CreateSoundBuffer()`, fill out a new format description, and set it using the `SetFormat()` method on the primary buffer. This code sets the primary format to 22 kHz, 16-bit stereo.

```
// pDSound is a valid LPDIRECTSOUND object.
LPDIRECTSOUNDBUFFER pDSBPrimary = NULL;

sAutoZero<DSBUFFERDESC> dsbd;
dsbd.dwFlags          = DSBCAPS_PRIMARYBUFFER;
dsbd.dwBufferBytes    = 0;
dsbd.lpwfxFormat      = NULL;

HRESULT hr = pDSound->CreateSoundBuffer( &dsbd, &pDSBPrimary, NULL );
if( FAILED(hr) )
{
    /* handle error */
}

// Set primary buffer format to 22 kHz and 16-bit output.
WAVEFORMATEX wfx;
ZeroMemory( &wfx, sizeof(WAVEFORMATEX) );
wfx.wFormatTag      = WAVE_FORMAT_PCM;
wfx.nChannels       = 2;
wfx.nSamplesPerSec  = 22050;
wfx.wBitsPerSample  = 16;
wfx.nBlockAlign     = wfx.wBitsPerSample / 8 * wfx.nChannels;
wfx.nAvgBytesPerSec = wfx.nSamplesPerSec * wfx.nBlockAlign;

HRESULT hr = hr = pDSBPrimary->SetFormat(&wfx)
if( FAILED( ) )
{
    throw cGameError( "SetFormat (DS) failed!" );
}

SafeRelease( pDSBPrimary );
```


With all the code in place, you can actually write the sound layer class.

```

/*****
 *      Advanced 3D Game Programming with DirectX 10.0
 *      * * * * *
 *
 *      See license.txt for modification and distribution information
 *      copyright (c) 2007 by Peter Walsh, Wordware
 *****/
#ifdef _SOUNDLAYER_H
#define _SOUNDLAYER_H

#include <dsound.h>
#include "GameErrors.h"          // Added by ClassView

class cSound;

class cSoundLayer
{
    LPDIRECTSOUND8          m_pDSound;
    LPDIRECTSOUNDBUFFER8    m_pPrimary;    // primary mixer

    static cSoundLayer      *m_pGlobalSLayer;

    cSoundLayer( HWND hWnd );

public:
    virtual ~cSoundLayer();

    static cSoundLayer *GetSound()
    {
        return m_pGlobalSLayer;
    }

    LPDIRECTSOUND8 GetDSound()
    {
        return m_pDSound;
    }

    static void Create( HWND hWnd )
    {
        new cSoundLayer( hWnd );
    }
};

inline cSoundLayer *Sound()
{
    return cSoundLayer::GetSound();
}

#endif // _SOUNDLAYER_H

```

```

/*****
 *      Advanced 3D Game Programming with DirectX 10.0
 *      * * * * *
 *
 *      See license.txt for modification and distribution information
 *      copyright (c) 2007 by Peter Walsh, Wordware
 *****/
#include "stdafx.h"

#include "SoundLayer.h"
#include "Sound.h"

cSoundLayer *cSoundLayer::m_pGlobalSLayer = NULL;

cSoundLayer::cSoundLayer( HWND hWnd )
{
    m_pDSound = NULL;
    m_pPrimary = NULL;

    if( m_pGlobalSLayer )
    {
        throw cGameError( "cSoundLayer already initialized!" );
    }
    m_pGlobalSLayer = this;

    HRESULT          hr;
    LPDIRECTSOUNDBUFFER pDSBPrimary = NULL;

    // Create IDirectSound using the primary sound device
    hr = DirectSoundCreate8( NULL, &m_pDSound, NULL );
    if( FAILED(hr) )
    {
        throw cGameError( "DirectSoundCreate failed!" );
    }

    // Set coop level to DSSCL_PRIORITY
    hr = m_pDSound->SetCooperativeLevel( hWnd, DSSCL_PRIORITY );
    if( FAILED(hr) )
    {
        throw cGameError( "SetCooperativeLevel (DS) failed!" );
    }

    // Get the primary buffer
    sAutoZero<DSBUFFERDESC> dsbd;
    dsbd.dwFlags          = DSBCAPS_PRIMARYBUFFER;
    dsbd.dwBufferBytes = 0;
    dsbd.lpwfxFormat      = NULL;

    hr = m_pDSound->CreateSoundBuffer( &dsbd, &pDSBPrimary, NULL );
    if( FAILED(hr) )
    {
        throw cGameError( "CreateSoundBuffer (DS) failed!" );
    }

    // Set primary buffer format to 22 kHz and 16-bit output.

```

```

    WAVEFORMATEX wfx;
    ZeroMemory( &wfx, sizeof(WAVEFORMATEX) );
    wfx.wFormatTag      = WAVE_FORMAT_PCM;
    wfx.nChannels       = 2;
    wfx.nSamplesPerSec  = 22050;
    wfx.wBitsPerSample  = 16;
    wfx.nBlockAlign     = wfx.wBitsPerSample / 8 * wfx.nChannels;
    wfx.nAvgBytesPerSec = wfx.nSamplesPerSec * wfx.nBlockAlign;

    if( FAILED( hr = pDSBPrimary->SetFormat(&wfx) ) )
    {
        throw cGameError( "SetFormat (DS) failed!" );
    }

    SafeRelease( pDSBPrimary );
}

cSoundLayer::~cSoundLayer()
{
    SafeRelease( m_pPrimary );
    SafeRelease( m_pDSound );
    m_pGlobalSLayer = NULL;
}

```

The cSound Class

To help facilitate the creation and playback of secondary buffers, I constructed an encapsulation class called `cSound`. A `cSound` object can be constructed either from a filename or from another `cSound` object. The copy constructor uses a ref-counting map so that all `cSounds` based on the same WAV file use the same `CWaveSoundRead` object. The overhead of the map could have been avoided if the `CWaveSoundRead` code was changed to accommodate the needed functionality, but I felt it was better to leave the code unchanged from the DirectX SDK.

Without any further ado, let's just dive into the code. The details of how this code works isn't terribly interesting, but have a look through it anyway to get accustomed to it.

```

/*****
 *           Advanced 3D Game Programming with DirectX 10.0
 * *****/
 *
 * See license.txt for modification and distribution information
 *      copyright (c) 2007 by Peter Walsh, Wordware
 *****/
#ifdef _SOUND_H
#define _SOUND_H

#include <map>

#include "SoundLayer.h"
#include "WavRead.h"

```

```

class cSound
{
    CWaveSoundRead *m_pWaveSoundRead;
    LPDIRECTSOUNDBUFFER8 m_pBuffer;
    int m_bufferSize;

    /**
     * Multiple sounds that use the same
     * file shouldn't reread it, they should
     * share the CWSR object. This map
     * implements rudimentary reference counting.
     * I would have just changed CWaveSoundRead,
     * but I wanted to keep it unchanged from the
     * samples.
     */
    static std::map< CWaveSoundRead*, int > m_waveMap;

    void Init();

public:
    cSound( char *filename );
    cSound( cSound& in );
    cSound& operator=( const cSound &in );

    virtual ~cSound();

    void Restore();
    void Fill();
    void Play( bool bLoop = false );

    bool IsPlaying();

};

#endif // _SOUND_H

/*****
 *      Advanced 3D Game Programming with DirectX 10.0
 *      * * * * *
 *
 *      See license.txt for modification and distribution information
 *      copyright (c) 2007 by Peter Walsh, Wordware
 *****/
#include "stdafx.h"

#include "WavRead.h"
#include "Sound.h"

using std::map;

map< CWaveSoundRead*, int > cSound::m_waveMap;

cSound::cSound( char *filename )
{
    m_pWaveSoundRead = NULL;

```

```

    m_pBuffer = NULL;

    // Create a new wave file class
    m_pWaveSoundRead = new CWaveSoundRead();
    m_waveMap[ m_pWaveSoundRead ] = 1;

    // Load the wave file
    if( FAILED( m_pWaveSoundRead->Open( filename ) ) )
    {
        throw cGameError("couldn't open file!");
    }

    Init();
    Fill();
}

cSound::cSound( cSound& in )
{
    m_pWaveSoundRead = in.m_pWaveSoundRead;
    m_waveMap[ m_pWaveSoundRead ]++;
    Init();
    Fill();
}

cSound& cSound::operator=( const cSound &in )
{
    /**
     * Destroy the old object
     */
    int count = --m_waveMap[ m_pWaveSoundRead ];
    if( !count )
    {
        delete m_pWaveSoundRead;
    }
    SafeRelease( m_pBuffer );

    /**
     * Clone the incoming one
     */
    m_pWaveSoundRead = in.m_pWaveSoundRead;
    m_waveMap[ m_pWaveSoundRead ]++;

    Init();
    Fill();

    return *this;
}

cSound::~cSound()
{
    int count = m_waveMap[ m_pWaveSoundRead ];
    if( count == 1 )
    {
        delete m_pWaveSoundRead;
    }
}

```

```

    }
    else
    {
        m_waveMap[ m_pWaveSoundRead ] = count - 1;
    }

    SafeRelease( m_pBuffer );
}

void cSound::Init()
{
    /**
     * Set up the DirectSound surface. The size of the sound file
     * and the format of the data can be retrieved from the wave
     * sound object. Besides that, we only set the STATIC flag,
     * so that the driver isn't restricted in setting up the
     * buffer.
     */
    sAutoZero<DSBUFFERDESC> dsbd;
    dsbd.dwFlags      = DSBCAPS_STATIC;
    dsbd.dwBufferBytes = m_pWaveSoundRead->m_ckIn.cksize;
    dsbd.lpwfxFormat  = m_pWaveSoundRead->m_pwfx;

    HRESULT hr;

    // Temporary pointer to old DirectSound interface
    LPDIRECTSOUNDBUFFER pTempBuffer = 0;

    // Create the sound buffer
    hr = Sound()->GetDSound()->CreateSoundBuffer(
        &dsbd, &pTempBuffer, NULL );
    if( FAILED(hr) )
    {
        throw cGameError("CreateSoundBuffer failed!");
    }

    // Upgrade the sound buffer to version 8
    pTempBuffer->QueryInterface( IID_IDirectSoundBuffer8, (void**)&m_pBuffer );
    if( FAILED(hr) )
    {
        throw cGameError("SoundBuffer query to 8 failed!");
    }

    // Release the temporary old buffer
    pTempBuffer->Release();

    /**
     * Remember how big the buffer is
     */
    m_bufferSize = dsbd.dwBufferBytes;
}

void cSound::Restore()
{

```

```

        HRESULT hr;

        if( NULL == m_pBuffer )
        {
            return;
        }

        DWORD dwStatus;
        if( FAILED( hr = m_pBuffer->GetStatus( &dwStatus ) ) )
        {
            throw cGameError( "couldn't get buffer status" );
        }

        if( dwStatus & DSBSTATUS_BUFFERLOST )
        {
            /**
             * Chances are, we got here because the app /just/
             * started, and DirectSound hasn't given us any
             * control yet. Just spin until we can restore
             * the buffer
             */
            do
            {
                hr = m_pBuffer->Restore();
                if( hr == DSERR_BUFFERLOST )
                    Sleep( 10 );
            }
            while( hr = m_pBuffer->Restore() );

            /**
             * The buffer was restored. Fill 'er up.
             */
            Fill();
        }
    }

void cSound::Fill()
{
    HRESULT hr;
    uchar *pbWavData; // Pointer to actual wav data
    uint cbWavSize; // Size of data
    void *pbData = NULL;
    void *pbData2 = NULL;
    ulong dwLength;
    ulong dwLength2;

    /**
     * How big the wav file is
     */
    uint nWaveFileSize = m_pWaveSoundRead->m_ckIn.cksize;

    /**
     * Allocate enough data to hold the wav file data
     */

```

```

        pbWavData = new uchar[ nWaveFileSize ];
        if( NULL == pbWavData )
        {
            delete [] pbWavData;
            throw cGameError("Out of memory!");
        }

        hr = m_pWaveSoundRead->Read(
            nWaveFileSize,
            pbWavData,
            &cbWavSize );
        if( FAILED(hr) )
        {
            delete [] pbWavData;
            throw cGameError("m_pWaveSoundRead->Read failed");
        }

        /**
        * Reset the file to the beginning
        */
        m_pWaveSoundRead->Reset();

        /**
        * Lock the buffer so we can copy the data over
        */
        hr = m_pBuffer->Lock(
            0, m_bufferSize, &pbData, &dwLength,
            &pbData2, &dwLength2, 0L );
        if( FAILED(hr) )
        {
            delete [] pbWavData;
            throw cGameError("m_pBuffer->Lock failed");
        }

        /**
        * Copy said data over, unlocking afterwards
        */
        memcpy( pbData, pbWavData, m_bufferSize );
        m_pBuffer->Unlock( pbData, m_bufferSize, NULL, 0 );

        /**
        * We're done with the wav data memory.
        */
        delete [] pbWavData;
    }

bool cSound::IsPlaying()
{
    DWORD dwStatus = 0;

    m_pBuffer->GetStatus( &dwStatus );

    if( dwStatus & DSBSTATUS_PLAYING )
        return true;
}

```



```

        else
            return false;
    }

void cSound::Play( bool bLoop )
{
    HRESULT hr;
    if( NULL == m_pBuffer )
        return;

    // Restore the buffers if they are lost
    Restore();

    // Play buffer
    DWORD dwLooped = bLoop ? DSBPLAY_LOOPING : 0L;
    if( FAILED( hr = m_pBuffer->Play( 0, 0, dwLooped ) ) )
    {
        throw cGameError("m_pBuffer->Play failed");
    }
}

```

Additions to *cApplication*

The only addition to *cApplication* is the *InitSound()* call, which initializes the sound layer. After the call completes you can freely create *cSound* objects and play them. If this is not the behavior you would like in your application, the function is overloadable. The code is in the following:

```

void cApplication::InitSound()
{
    cSoundLayer::Create( MainWindow()->GetHWnd() );
}

```

Application: DirectSound Sample

Adrian, the lead author of the original version of this book, has a few interesting hobbies that have nothing to do with programming. One of them is a beatbox group for which he sings bass. One of his jobs in the beatbox group is to take care of some of the vocal percussion.

Beatbox music is percussive sounds made with the human voice. This has spawned an entire subculture of vocal percussionists, each trying to make that perfect snare sound or cymbal crash using only their mouths. The *DirectSound* sample for this chapter was created using Adrian's vocal abilities!

When you load the file *DSSAMPLE* from the companion files, you're presented with a small window that lists six different vocal percussion sounds. The keys 1 through 6 play each of the sounds, and you can press multiple keys simultaneously to play multiple sounds.

Recall that I didn't show you a `DirectInput` sample earlier in this chapter because I figured it would be better to roll `DirectSound` and `DirectInput` into one sample. Well, here it is. `DirectInput` is used to capture the key-strokes. With some practice you can get a pretty swank beat going.

```

/*****
 *      Advanced 3D Game Programming with DirectX 10.0
 *      * * * * *
 *
 *      See license.txt for modification and distribution information
 *      copyright (c) 2007 by Peter Walsh, Wordware
 *****/
#include "stdafx.h"

#include <vector>
#include <string>
using namespace std;

class cDSSampleApp : public cApplication, public iKeyboardReceiver
{
    vector< cSound* > m_sounds[6];
    string m_names[6];

    int m_states[6]; // states of the keys 1-6

public:

    void PlaySound( int num );

    //=====----- cApplication

    virtual void DoFrame( float timeDelta );
    virtual void SceneInit();

    cDSSampleApp() :
        cApplication()
    {
        m_title = string( "DirectSound Sample" );
        m_width = 320;
        m_height = 200;

        for( int i=0; i<6; i++ ) m_states[i] = 0;
    }

    ~cDSSampleApp()
    {
        for( int i=0; i<6; i++ )
        {
            for( int i2=0; i2< m_sounds[i].size(); i2++ )
            {
                delete m_sounds[i][i2];
            }
        }
    }
}

```

```
        virtual void KeyUp( int key );
        virtual void KeyDown( int key );

};

cApplication *CreateApplication()
{
    return new cDSSampleApp();
}

void DestroyApplication( cApplication *pApp )
{
    delete pApp;
}

void cDSSampleApp::SceneInit()
{
    m_names[0] = string("media\\keg.wav");
    m_names[1] = string("media\\crash1.wav");
    m_names[2] = string("media\\crash2.wav");
    m_names[3] = string("media\\bass.wav");
    m_names[4] = string("media\\snare.wav");
    m_names[5] = string("media\\hihat.wav");

    Input()->GetKeyboard()->SetReceiver( this );

    for( int i=0; i<6; i++ )
    {
        m_sounds[i].push_back( new cSound( (char*)m_names[i].c_str() ) );
    }
}

void cDSSampleApp::PlaySound( int num )
{
    /**
     * iterate through the vector, looking
     * for a sound that isn't currently playing.
     */
    vector<cSound*>::iterator iter;
    for( iter = m_sounds[num].begin(); iter != m_sounds[num].end(); iter++ )
    {
        if( !(*iter)->IsPlaying() )
        {
            (*iter)->Play();
            return;
        }
    }

    /**
     * A sound wasn't found. Create a new one.
     */
}
```

```

        DP("spawning a new sound\n");

        cSound *pNew = new cSound( *m_sounds[num][0] );
        m_sounds[num].push_back( pNew );
        m_sounds[num][ m_sounds[num].size() - 1 ]->Play();
    }

void cDSSampleApp::DoFrame( float timeDelta )
{
    // Clear the previous contents of the back buffer
    Graphics()->GetDevice()->Clear( 0, 0, D3DCLEAR_TARGET | D3DCLEAR_ZBUFFER,
                                   D3DCOLOR_XRGB( 0,0,200), 1.0f, 0 );

    // Set up the strings
    string help;
    help += "DirectSound Sample application\n";
    help += "Vocal Percussion with Adrian Perez\n";
    help += " [1]: Keg drum\n";
    help += " [2]: Crash 1\n";
    help += " [3]: Crash 2\n";
    help += " [4]: Bass drum\n";
    help += " [5]: Snare drum\n";
    help += " [6]: Hi-Hat\n";

    // Tell Direct3D we are about to start rendering
    Graphics()->GetDevice()->BeginScene();

    // Output the text
    Graphics()->DrawTextString( 1, 1, D3DCOLOR_XRGB( 0, 255, 0), help.c_str() );

    // Tell Direct3D we are done rendering
    Graphics()->GetDevice()->EndScene();

    // Present the back buffer to the primary surface
    Graphics()->Flip();
}

void cDSSampleApp::KeyDown( int key )
{
    switch( key )
    {
    case DIK_1:
        if( !m_states[0] )
        {
            m_states[0] = 1;
            PlaySound(0);
        }
        break;
    case DIK_2:
        if( !m_states[1] )
        {
            m_states[1] = 1;
            PlaySound(1);
        }
    }
}

```

```
        }
        break;
    case DIK_3:
        if( !m_states[2] )
        {
            m_states[2] = 1;
            PlaySound(2);
        }
        break;
    case DIK_4:
        if( !m_states[3] )
        {
            m_states[3] = 1;
            PlaySound(3);
        }
        break;
    case DIK_5:
        if( !m_states[4] )
        {
            m_states[4] = 1;
            PlaySound(4);
        }
        break;
    case DIK_6:
        if( !m_states[5] )
        {
            m_states[5] = 1;
            PlaySound(5);
        }
        break;
    }
}

void cDSSampleApp::KeyUp( int key )
{
    switch( key )
    {
    case DIK_1:
        m_states[0] = 0;
        break;
    case DIK_2:
        m_states[1] = 0;
        break;
    case DIK_3:
        m_states[2] = 0;
        break;
    case DIK_4:
        m_states[3] = 0;
        break;
    case DIK_5:
        m_states[4] = 0;
        break;
    case DIK_6:
        m_states[5] = 0;
        break;
    }
```

```
}  
}
```

Conclusion

And that is DirectInput and DirectSound. You should be able to build on top of the code I have shown you in this chapter to create the perfect acoustic accompaniment to your own projects without too much difficulty.




Figure 3.2: The sound and input sample

Now it's time to focus on the 3D mathematics that you need to start exploring the next dimension with Direct3D.

This page intentionally left blank.

Chapter 4

3D Math Foundations

When you really get down to it, using 3D graphics is an exercise in math. Some of the math can be interesting; some of it can be seriously dull. It all depends on the eye of the beholder. Love it or hate it, however, you still have to learn it. A solid foundation in math is a requirement if you want to be a successful 3D coder. Don't worry, though; I'll try to keep this chapter as interesting as possible.

Points

Let's start with the most basic of basic primitives: the *3D point*, or *vector*. Points are paramount in 3D graphics. The vertices of objects, the objects' locations and directions, and their velocities/forces are all described with 3D points. Three-dimensional objects have width, height, and depth, which are represented with the shorthand components x (width), y (height), and z (depth). Points and vectors, when used in equations, are referred to as *vector quantities*, while regular numbers are referred to as *scalar quantities*. The three components are written separated by commas like this: $\langle x, y, z \rangle$. They are also represented as a single row matrix (or, equivalently, a transposed single column matrix). At the right is an example of how points are represented using matrix notation:

$$\mathbf{v} = [\mathbf{v}_x \quad \mathbf{v}_y \quad \mathbf{v}_z]$$

and

$$\mathbf{v} = \begin{bmatrix} \mathbf{v}_x \\ \mathbf{v}_y \\ \mathbf{v}_z \end{bmatrix}^T$$



Note: In the book I use the terms “point” and “vector” interchangeably. They loosely mean the same thing. A point is a location in 3D space, and a vector is a line that goes from the origin to a location in 3D space. For all the math that I'm discussing in this book, they can be used interchangeably.

Here are a few examples of three-dimensional points.




Figure 4.1: Examples of 3D vectors

3D points are graphed in a way analogous to the 2D Cartesian coordinate system. There are three principal axes stretching off into infinity in three directions. These are the x-, y-, and z-axes. They meet at the *origin*, a specific point that represents the center of the current coordinate system (typically you have several coordinate systems to worry about, but I'll leave this until later). The coordinates of the origin are, of course, $\langle 0,0,0 \rangle$.

Which way do the axes point? In some systems (for example, some 3D modelers like 3D Studio Max), x increases to the right, y increases forward (into the screen), and z increases up. These directions are all dependent on the orientation of the viewer of the scene. My choice of axes direction is the one used in most 3D games: x increases to the right, y increases up, and z increases forward, into the monitor (or away from you if that makes it clearer).



Note: This book uses a left-handed coordinate space, where x increases to the right, y increases up, and z increases forward (into the screen). In right-handed coordinate systems, z increases coming out of the screen.

A point always exists some distance away from the origin of the coordinate space; this quantity is called the *magnitude* of the vector (or, more intuitively, the *length* of the vector). To compute the magnitude of vectors in 2D, you use the Pythagorean theorem:

$$\text{magnitude} = \sqrt{x^2 + y^2}$$

Luckily, the Pythagorean theorem extends into 3D to measure the length of 3D vectors by simply adding the extra z component into the equation. You can see that the 2D Pythagorean equation is simply a special case of the 3D equation where the z-distance from the origin is zero.

There is a shorthand notation used to denote the magnitude of a vector when used in more complex equations. The notation is the vector surrounded on both sides by double vertical lines. The equation for vector length, given a vector v with components x , y , and z is:

$$\|v\| = \sqrt{x^2 + y^2 + z^2}$$

A special type of vector is one that has a length of 1. This type of vector is called a *unit vector*. Each unit vector touches a point on what is called the *unit sphere*, a conceptual sphere with radius 1, situated at the origin.

It's often the case that you want a unit-length version of a given vector. For example, the unit-length version n of a given vector m would be:

$$n = \frac{m}{\|m\|}$$

For simplicity's sake, however, I'll introduce some shorthand notation. The same equation can be represented by putting a bar over m to signify the unit-length version:

$$n = \overline{m}$$

There are three specific unit vectors that represent the directions along the three primary axes: i $\langle 1,0,0 \rangle$, j $\langle 0,1,0 \rangle$, and k $\langle 0,0,1 \rangle$.




Figure 4.2: The i , j , and k vectors

Many physics texts use the *i*, *j*, and *k* vectors as primitives to describe other 3D vectors, so it is worth mentioning it here. Any point in 3D can be represented as a linear combination of the *i*, *j*, and *k* vectors. You can define any vector as a sum of the scalar components with the three principal vectors. For example, if you had the 3D vector $\mathbf{a} = \langle 3, 5, 2 \rangle$, you could represent it like this:

$$\mathbf{a} = 3\mathbf{i} + 5\mathbf{j} + 2\mathbf{k}$$

This trait will become more important later on, when I discuss matrices and the spaces they represent.



Note: While it isn't really relevant to the level of skill you need to reach in this book, the concept of a linear combination is important when talking about spaces and transformations.

Given n vectors $\mathbf{b}_0 \dots \mathbf{b}_{n-1}$, any vector \mathbf{v} is a linear combination of the set of the vectors if the following equation can be satisfied:

$$\mathbf{v} = k_0\mathbf{b}_0 + k_1\mathbf{b}_1 + \dots + k_{n-1}\mathbf{b}_{n-1}$$

where $k_0 \dots k_{n-1}$ are scalars.

That is, if you want to get to \mathbf{v} , you can start at the origin and walk along any or all of the vectors some amount and reach \mathbf{v} .

You can say the set of \mathbf{b} vectors is linearly independent if no single \mathbf{b} vector is a linear combination of the others.

The point3 Structure

It is always useful to design a class to encapsulate a generic 3D point. The class name I use is `point3`. Unlike most of the other classes you have seen so far, the intent of the `point3` structure is to act as a mathematical primitive like `float` or `int`. The 3 suffix denotes the dimension of the point. I'll also define 2D and 4D versions of points, which are named `point2` and `point4`, respectively.

```
struct point3
{
    union
    {
        struct
        {
            float x,y,z;    // 3 real components of the vector
        };
        float v[3];        // Array access useful in for loops
    };

    // Default constructor
    point3({})

    // Construct a point with 3 given inputs
```

```

point3( float X, float Y, float Z ) :
    x(X), y(Y), z(Z)
{
}
// ... more will go in here.
};

```

This class uses a union in conjunction with a nameless struct. If you've never encountered unions before, a *union* is used to name components that share memory. So, in the above code, the `y` variable and the `v[1]` variable represent the same piece of memory; when one of them changes, both of them change. A *nameless struct* is used to let you define the `x`, `y`, and `z` components as one atomic unit (since I don't want them to each be referring to the same piece of memory). This way you can use the familiar `x,y,z` notation for most of the code, but maintain the ability to index into an array for iteration.



Note: The non-default constructor uses initialization lists. C++ classes should use these whenever possible. They clarify the intent of the code to the compiler, which lets it do its job better (it has a better chance to inline the code, and the code will end up being considerably more efficient, especially for complex structures).

Finally, you may wonder why I'm choosing floats (32 bits/4 bytes) instead of doubles (64 bits/8 bytes) or long doubles (80 bits/10 bytes). Well, I *could* just implement the point as a template class, but there are too many other interactions with other classes to complicate the code that much. Using it as a template in a way defeats the concept of using the point as a generic primitive.

Doubles and long doubles are slower than floats, about twice as slow for things like divides (19 versus 39 cycles), and on top of that they require twice the space. The added precision really isn't important unless you really need a wide range of precision. Within a few years worlds will be big enough and model resolution will be fine enough that you may need to employ larger floating-point resolutions to get the job done. Until then I'd suggest sticking with traditional floats.

Basic point3 Functions

The `point3` structure is pretty basic right now. To spice it up, I'll add some member functions to help perform some basic operations on 3D points, and explain what they are used for.

Assign

Setting a point to a certain value is a common task. It could be done explicitly in three lines, setting the `x`, `y`, and `z` values separately. However, for simplicity's sake, it's easier to set them all at once, with a single

function call. This is also better than just creating a new variable on the stack with a point3 constructor; it's more efficient to reuse stack variables whenever possible.

```
// Reassign a point without making a temporary structure
inline void point3::Assign( float X, float Y, float Z )
{
    x=X;
    y=Y;
    z=Z;
}
```

Mag and MagSquared

The function Mag uses the 3D version of the Pythagorean theorem mentioned previously to calculate the length of the point structure (the distance from the point to the origin).

```
inline float point3::Mag() const
{
    return (float)sqrt( x*x + y*y + z*z );
}
```

Sometimes you want the squared distance (for example, when calculating the attenuation factor for a point-source light). Rather than using the computationally expensive square root and squaring it, you can avoid the cost and simply make an extra function to do it for you, as shown below.

```
inline float point3::MagSquared() const
{
    return ( x*x + y*y + z*z );
}
```

Normalize

Normalize takes a point structure and makes it a unit-length vector pointing in the same direction.

```
inline void point3::Normalize()
{
    float InvertedMagnitude=1/Mag();
    x*=InvertedMagnitude;
    y*=InvertedMagnitude;
    z*=InvertedMagnitude;
}
```

Dist

Dist is a static function that calculates the distance between two point structures. Conceptually, it finds the vector that connects them (which is the vector $b-a$) and computes its length.

```
inline static float point3::Dist( const point3 &a, const point3 &b )
{
    point3 distVec( b.x - a.x, b.y - a.y, b.z - a.z );
    return distVec.Mag();
}
```

point3 Operators

Now that there is a basic primitive I can use, like other primitives (e.g., int or float), I need some way to operate on the data. Since vectors can be added, subtracted, and multiplied (sort of), just like scalars, it would be cool to have an easy way to perform these operations. Operator overloading to the rescue! C++ lets you modify/define the behavior of operators on classes.

Addition/Subtraction

Vector addition and subtraction are useful in moving points around in 3D. Conceptually, adding a vector to another moves the location of the first vector in the direction of the second. Figure 4.3 shows what the result of vector addition looks like, and Figure 4.4 shows the result of vector subtraction.




Figure 4.3: Vector addition

In many respects, vector addition/subtraction is incredibly similar to the normal scalar addition that I'm sure you know and love. For example, if you wanted to find the average location of a set of vectors, you simply add

them together and divide the result by the number of vectors added, which is, of course, the same averaging formula used for scalars.

The code for adding/subtracting vectors is equally similar to their scalar cousins: Simply add (or subtract) each component together separately. I'll give the $+$ and $-$ operators; in the book's sample code you'll find the $+=$ and $-=$ operators.

```
inline point3 operator+(point3 const &a, point3 const &b)
{
    return point3
    (
        a.x+b.x,
        a.y+b.y,
        a.z+b.z
    );
};

inline point3 operator-(point3 const &a, point3 const &b)
{
    return point3
    (
        a.x-b.x,
        a.y-b.y,
        a.z-b.z
    );
};
```




Figure 4.4: Vector subtraction

Vector-Scalar Multiplication/Division

Often, you may want to increase or decrease the length of a vector, while making sure it still points in the same direction. Basically, you want to scale the vector by a scalar. Figure 4.5 shows what scaling vectors looks like.




Figure 4.5: Scaling vectors

Doing this in code is easy enough; just multiply (or divide) each component in the vector by the provided scalar. Below, we use the `*` and `/` operators; the `*=` and `/=` operators are defined in the header. Note that I defined two multiplicative operators; one for `vector * scalar` and another for `scalar * vector`.

```
inline point3 operator*(point3 const &a, float const &b)
{
    return point3
    (
        a.x*b,
        a.y*b,
        a.z*b
    );
};

inline point3 operator*(float const &a, point3 const &b)
{
    return point3
```



```

    (
        a*b.x,
        a*b.y,
        a*b.z
    );
};

inline point3 operator/(point3 const &a, float const &b)
{
    float inv = 1.f / b; // Cache the division.
    return point3
    (
        a.x*inv,
        a.y*inv,
        a.z*inv
    );
};

```

Vector Equality

You'll often need to know if two points represent the same location in 3D; for example, to see if two polygons share an edge, or if two triangles are the same. We overload the equality operator (`==`) to do this.

This one, at first glance, would be a no-brainer; just compare to see if the x, y, and z values match up. However, the answer is not as simple as that. This is one of many points where an important line in the sand must be drawn, a line between the wonderful land of Theory and the painful land of Reality.

In Theory, there is infinite precision for scalar values. The decimal value of $1/3$ has a string of 3s that never ends. When you multiply two scalar numbers together, the solution is exactly the correct one. When comparing two floating-point numbers, you *know* (with infinite precision) whether or not the numbers are equal. When multiplying an extremely large number by an extremely small one, the result is exactly what was expected. Everything is nice, in Theory. However, right now real estate is pretty expensive over there, so you and I are going to have to stay here, in Reality.

In Reality, floating pointers do not have infinite precision. Floats (32 bits) and doubles (64 bits) can only encode so much precision (around 5 and 15 base 10 places, respectively). They do not multiply nicely. If you multiply an extremely large number and an extremely small number, the solution is not the solution you might expect, due to the lack of precision. Finally, they do not handle equality too well. Due to all the imprecision floating around, two different paths of calculation that should result in the same answer can yield subtly different, although technically equal, numbers.



Note: Look on the web for programming horror stories, and you'll see that countless problems in computer programming come from assuming that floating- or fixed-point numbers will have enough precision for what you need.

So, how is this little issue fixed? In practice, the problem really can't be fixed, but it is possible to hack out a solution that works well enough. Epsilon values provide the answer. (You'll see them covered throughout the book; epsilons are used all over the place to correct for numerical imprecision.) What you do is make each point have a certain, extremely small mass, going out in each direction a value epsilon. For instance, I will be using an epsilon of 10^{-3} , or 0.001. That way, in order to see if two points are equal (or, in this case, equal enough), you test to see if the difference between them is less than or equal to epsilon. If this case is satisfied for all three coordinates, then it can safely be said the two points are equal.



Note: In case you haven't picked it up yet, getting a solution to a problem that is not necessarily correct but good enough is one of the mantras of graphics programming. There is never enough time to calculate everything the hard way; the more corners you can cut without people being able to tell, the more of an edge you'll have over your competition.

In code, this becomes:

```
// above somewhere: #define EPSILON 0.001
inline bool operator==(point3 const &a, point3 const &b)
{
    if( fabs(a.x-b.x)<EPSILON )
    {
        if( fabs(a.y-b.y)<EPSILON )
        {
            if( fabs(a.z-b.z)<EPSILON )
            {
                return true; // We passed
            }
        }
    }
    return false; // The points were not equal enough
};
```

Dot Product

The dot product (mathematically represented with the symbol \bullet) is one of the most important operations in 3D graphics. It is used everywhere. Everything from transformation to clipping to BSP tree traversal uses the dot product.

The mathematical definition for the dot product is this:

$$\mathbf{u} \bullet \mathbf{v} = \|\mathbf{u}\| \times \|\mathbf{v}\| \times \cos(\theta)$$

In this equation, \mathbf{u} and \mathbf{v} represent two vectors in 3D. The $\|\mathbf{u}\|$ and $\|\mathbf{v}\|$ represent the lengths of the vectors, and θ represents the angle between the vectors. As you can see from the equation, the result of the dot product equation is a scalar, not a vector.

Conceptually, the dot product describes the relation between two vectors in scalar form. If one of the vectors is normalized, the dot product represents the length of the shadow that the other vector would cast, as shown in Figure 4.6.




Figure 4.6: The conceptual “shadow” of a dot product

This particular trait is used in clipping.

Using the equation given above, you can rearrange the terms to provide a way to find the angle θ between two vectors:

$$\cos(\theta) = \frac{\mathbf{u} \bullet \mathbf{v}}{\|\mathbf{u}\| \times \|\mathbf{v}\|}$$

This works out very conveniently if both vectors are unit-length; two square roots (to find the vector lengths) and a division drop out of the equation and you get:

$$\theta = \cos^{-1}(\mathbf{u} \bullet \mathbf{v}) \quad (\text{if } \mathbf{u} \text{ and } \mathbf{v} \text{ are unit-length})$$

How does this work? This seems like a rather arbitrary trait for the dot product to have. Well, for some insight, think back to your trigonometry days. My trigonometry professor had a little mnemonic device to help remember the basic rules of trigonometry called “SOHCAHTOA.” The middle three letters say that cosine is equal to the adjacent edge divided by the hypotenuse of a right triangle or:

$$\cos(\theta) = \frac{\text{adjacent}}{\text{hypotenuse}}$$

Now, on a unit circle, the hypotenuse will be length 1, so that drops out of the equation. You’re left with $\cos(\theta) = \text{adjacent edge}$. Think of the adjacent edge as the shadow of the hypotenuse onto the x-axis, as shown in Figure 4.7.




Figure 4.7: Cosine terms




Figure 4.8: Vector cosine analog

So if, for the sake of this example, v is a unit vector going out along the x -axis, and u is a unit vector of the hypotenuse, then $u \bullet v$ will give the length of the shadow of u onto v , which is equivalent to the adjacent edge in the right triangle, and therefore $\cos(\theta)$.

The actual code behind the dot product is much simpler than the equations above, devoid of square roots, divides, and cosines (which is great, since the dot product is computed so often!). The dot product is achieved by summing the piecewise multiplication of each of the components. To implement dot products, I'm going to overload the multiplication operator (*). It seems almost mysterious how three multiplications and two additions can make the same result that you get from the complex equation above, but don't look a gift horse in the mouth, as they say.

```
inline float operator*(point3 const &a, point3 const &b)
{
    return a.x*b.x + a.y*b.y + a.z*b.z;
}
```

Cross Product

Another operation that can be performed between two vectors is called the cross product. It's represented mathematically with the symbol \times . The formula for computing the cross product is shown at the right:

$$\mathbf{a} \times \mathbf{b} = \begin{bmatrix} \mathbf{a}_y \mathbf{b}_z - \mathbf{a}_x \mathbf{b}_y \\ \mathbf{a}_z \mathbf{b}_x - \mathbf{a}_x \mathbf{b}_y \\ \mathbf{a}_x \mathbf{b}_y - \mathbf{a}_x \mathbf{b}_y \end{bmatrix}^T$$

The operation returns a vector, not a scalar like the dot product. The vector it returns is mutually orthogonal to both input vectors. A vector mutually orthogonal to two others means that it is perpendicular to both of them. The resultant vector from a cross product is perpendicular (or orthogonal) to both of the input vectors. Once I start discussing planes you'll see how useful they can be. For most applications of the cross product, you want the result to be unit-length.



Warning: If the two input vectors in a cross-product operation are parallel, the result of the operation is undefined (as there are an infinite number of vectors that are perpendicular to one vector).

An important note is that the cross product operation is not commutative. That is, $a \times b$ is *not* the same as $b \times a$. They are very similar, however, as one points in the opposite direction of the other.

Now it's time for some implementation. Since the `*` operator is already used for the dot product, something else needs to be picked instead. The choice of operator is fairly arbitrary, but following the example of a former professor of mine named David Baraff (who's now working at Pixar), I use the XOR operator (`^`). The code, while not the most intuitive thing in the world, follows the equations stated above.

```
inline point3 operator^(point3 const &a, point3 const &b)
{
    return point3
    (
        (a.y*b.z-a.z*b.y),
        (a.z*b.x-a.x*b.z),
        (a.x*b.y-a.y*b.x)
    );
}
```

The full code that defines all the behavior for points in 3D is found in the downloadable files in the code directory for this chapter in `point3.h` and `point3.cpp`. Also included are `point4.h` and `point4.cpp`, which define four-dimensional points (you'll see these later for quaternion rotations and parametric surfaces).

Polygons

The polygon is the bread and butter of computer graphics. Rendered images would be pretty bland if you didn't have the polygon. While there are other primitives used in computer graphics (implicit surfaces, for example), just about every personal computer on the market today has hardware in it to accelerate the drawing of polygons (well, triangles actually...but same difference), so the polygon is king.

All of the polygons dealt with here are convex. Convex polygons have no dents or pits, i.e., no internal obtuse angles. A convex polygon is much

easier to rasterize, easier to clip, easier to cull, and the list goes on. While you could deal with concave polygons, the code to manage them and draw them is harder than in the convex case. It's much easier to represent concave polygons with two or more convex polygons.

Polygons (or triangles, which I'll discuss next) describe the boundary representation of an object (academic and CAD texts often use the term *b-rep* to mean this). A *b-rep* is simply a set of polygons that exactly define the boundary surface of an object. If the object is a cube, the *b-rep* is the six square polygons that make up each face.

Figure 4.9 shows four examples of polygons. Two are convex, and two are not.




Figure 4.9: Examples of polygons

For now, I'll implement polygons with a template class. This is because I don't want to have to reimplement the class for holding indices or point data (I'll discuss what this means later). While polygons could be implemented with a fully dynamic array, like an STL vector, I chose to limit the functionality for the sake of speed. Each polygon is created with a maximum number of possible elements it can contain. For most applications, a number like 8 or 10 will suffice. In addition to this, there is the number of actual elements in the polygon. This number should never be greater than the maximum number of elements. Let's take a look at the polygon class:

```
template <class type>
struct polygon
{
    int nElem; // number of elements in the polygon
```

```

int maxElem;

type *pList;

polygon()
{
    nElem = 0;
    maxElem = 0;
    pList = NULL;
}

polygon( int maxSize )
{
    maxElem = maxSize;
    pList = new type[maxSize];
}

polygon( const polygon &in )
{
    CloneData( in );
}

~polygon()
{
    DestroyData();
}

void CloneData( const polygon &in )
{
    if( !in.pList )
        return;

    pList = new type[in.maxElem];
    maxElem = in.maxElem;
    nElem = in.nElem;
    for( int i=0; i<in.nElem; i++ )
    {
        pList[i] = in.pList[i];
    }
}

void DestroyData( )
{
    delete[] pList;
    pList = NULL;
}

polygon& operator=( const polygon<type> &in )
{
    if( &in != this )
    {
        DestroyData();
    }
}

```



```
        CloneData( in );  
    }  
  
    return *this;  
}  
};
```

Triangles

Triangles are to 3D graphics what pixels are to 2D graphics. Every PC hardware accelerator under the sun uses triangles as the fundamental drawing primitive (well...scan line aligned trapezoids actually, but that's a hardware implementation issue). When you draw a polygon, hardware devices really draw a fan of triangles. Triangles “flesh out” a 3D object, connecting them together to form a skin or mesh that defines the boundary surface of an object. Triangles, like polygons, generally have an orientation associated with them, to help in normal calculations. All of the code in this book uses the convention that you are located in front of a triangle if the ordering of the vertices goes *clockwise* around the triangle. Figure 4.10 shows what a clockwise ordered triangle would look like.




Figure 4.10: A triangle from three points

When defining a mesh of triangles that define the boundary of a solid, you set it up so that all of the triangles along the skin are ordered clockwise when viewed from the outside.

It is impossible to see triangles that face away from you. (You can find this out by computing the triangle's plane normal and performing a dot product with a vector from the camera location to a location on the plane.)

Now let's move on to the code. To help facilitate using the multiple types, I'll implement triangles using templates. The code is fairly simple; it uses triangles as a container class, so I only define constructors and keep the access public so accessors are not needed.

```
template <class type>
struct tri
{
    type v[3]; // Array access useful for loops

    tri()
    {
        // nothing
    }

    tri( type v0, type v1, type v2 )
    {
        v[0] = v0;
        v[1] = v1;
        v[2] = v2;
    }
};
```

Strips and Fans

Lists of triangles are generally represented in one of three ways. The first is an explicit list or array of triangles, where every three elements represent a new triangle. However, there are two additional representations, designed to save bandwidth while sending triangles to dedicated hardware to draw them. They are called *triangle strips* and *triangle fans*.



Warning: DirectX no longer supports triangle fans in version 10.

Triangle fans, conceptually, look like the folding fans you see in Asian souvenir shops. They are a list of triangles that all share a common point. The first three elements indicate the first triangle. Then each new element is combined with the first element and the current last element to form a new triangle. Note that an n -sided polygon can be represented efficiently using a triangle fan. Figure 4.11 illustrates what I'm talking about.

Triangles in a triangle strip, instead of sharing a common element with all other triangles like a fan, only share elements with the triangle immediately preceding them. The first three elements define the first triangle. Then each subsequent element is combined with the two elements before it, in clockwise order, to create a new triangle. See Figure 4.12 for an explanation of strips.




Figure 4.11: Triangle fan




Figure 4.12: Triangle strip

Planes

The next primitive to discuss is the plane. Planes are to 3D what lines are in 2D; they're $n-1$ dimensional hyperplanes that can help you accomplish various tasks. Planes are defined as infinitely large, infinitely thin slices of space, like big pieces of paper. Triangles that make up your model each exist in their own plane. When you have a plane that represents a slice of 3D space, you can perform operations like classification of points and polygons and clipping.

So how do you represent planes? Well, it is best to build a structure from the equation that defines a plane in 3D. The implicit equation for a plane is:

$$ax + by + cz + d = 0$$

What do these numbers represent? The triplet $\langle a, b, c \rangle$ represents what is called the normal of the plane. A *normal* is a unit vector that, conceptually speaking, sticks directly out of a plane. A stronger mathematical definition would be that the normal is a vector that is perpendicular to all of the points that lie in the plane.

The d component in the equation represents the distance from the plane to the origin. The distance is computed by tracing a line toward the plane until you hit it. Finally, the triplet $\langle x, y, z \rangle$ is any point that satisfies the equation. The set of all points $\langle x, y, z \rangle$ that solve the equation is exactly all the points that lie in the plane.

All of the pictures I'm showing you will be of the top-down variety, and the 3D planes will be on edge, appearing as 2D lines. This makes figure drawing much easier; if there is an easy way to represent infinite 3D planes in 2D, I sure don't know it.

Following are two examples of planes. The first has the normal pointing away from the origin, which causes d to be negative (try some sample values for yourself if this doesn't make sense). The second has the normal pointing toward the origin, so d is positive. Of course, if the plane goes through the origin, d is zero (the distance from the plane to the origin is zero). Figures 4.13 and 4.14 provide some insight into this relationship.




Figure 4.13: d is positive when it faces away from the origin




Figure 4.14: d is negative when the normal faces the origin

It's important to notice that technically the normal $\langle a, b, c \rangle$ does not have to be unit-length for it to have a valid plane equation. But since things end up nicer if the normal is unit-length, all of the normals in this book are unit-length. Here is the code for the plane class:

```

struct plane3 {

    point3 n; // Normal of the plane
    float d; // Distance along the normal to the origin

    plane3( float nX, float nY, float nZ, float D ) :
        n( nX, nY, nZ ), d( D )
    {
        // All done.
    }
    plane3( const point3& N, float D ) :
        n( N ), d( D )
    {
        // All done.
    }

    // Construct a plane from three 3D points
    plane3( const point3& a, const point3& b, const point3& c);

    // Construct a plane from a normal direction and
    // a point on the plane
    plane3( const point3& norm, const point3& loc);

    // Construct a plane from a polygon
    plane3( const polygon<point3>& poly );

    plane3()
    {
        // Do nothing
    }

    // Flip the orientation of the plane
    void Flip();
};

```

Constructing a plane given three points that lie in the plane is a simple task, as shown below. You just perform a cross product between the two vectors made up by the three points ($\langle \text{point}_2 - \text{point}_0 \rangle$ and $\langle \text{point}_1 - \text{point}_0 \rangle$) to find a normal for the plane. After generating the normal and making it unit-length, finding the d value for the plane is just a matter of storing the negative dot product of the normal with any of the points. This holds because it essentially solves the plane equation above for d . Of course plugging a point into the plane equation will make it equal 0, and this constructor has three of them.

```

inline plane3::plane3(
    const point3& a,
    const point3& b,
    const point3& c )
{
    n = (b-a)^(c-a);
    n.Normalize();
}

```

```

    d = -(n*a);
}

```

If you already have a normal and also have a point on the plane, the first step can be skipped.

```

inline plane3::plane3( const point3& norm, const point3& loc ) :
    n( norm ), d( -(norm*loc) )
{
    // all done
}

```

Finally, constructing a plane given a polygon of point3 elements is just a matter of taking three of the points and using the constructor given above.

```

inline plane3::plane3( const polygon<point3>& poly )
{
    point3 a = poly.pList[0];
    point3 b = poly.pList[1];
    point3 c = poly.pList[2];

    n = (b-a)^(c-a);
    n.Normalize();
    d = -(n*a);
}

```

This brings up an important point. If you have an n -sided polygon, nothing discussed up to this point is forcing all of the points to be coplanar. However, problems can crop up if some of the points in the polygon aren't coplanar. For example, when I discuss back-face culling in a moment, you may misidentify what is actually behind the polygon, since there won't be a plane that clearly defines what is in front of and what is behind the plane. That is one of the advantages of using triangles to represent geometry—three points define a plane exactly.

Defining Locality with Relation to a Plane

One of the most important operations planes let you perform is defining the location of a point with respect to a plane. If you drop a point into the equation, it can be classified into three cases: in front of the plane, in back of the plane, or coplanar with the plane. Front is defined as the side of the plane the normal sticks out of.

Here, once again, precision will rear its ugly head. Instead of doing things the theoretical way, having the planes infinitely thin, I'm going to give them a certain thickness of (you guessed it) epsilon.

How do you orient a point in relation to a plane? Well, simply plug x , y , and z into the equation, and see what you get on the right side. If you get zero (or a number close enough to zero by plus or minus epsilon), then the point satisfied the equation and lies on the plane. Points like this can be called coplanar. If the number is greater than zero, then you know that you would have to travel farther along the origin following the path of the normal than you would need to go to reach the plane, so the point must be

in front of the plane. If the number is negative, it must be behind the plane. Note that the first three terms of the equation simplify to the dot product of the input vector and the plane normal. Figure 4.15 shows a visual representation of this operation.




Figure 4.15: Classifying points with a plane

```
// Defines the three possible locations of a point in
// relation to a plane
enum ePointLoc
{
    ptFront,
    ptBack,
    ptCoplanar
};

// we're inlining this because we do it constantly
inline ePointLoc plane3::TestPoint( point3 const &point ) const
{
    float dp = (point * n) + d;

    if(dp > EPSILON)
    {
        return ptFront;
    }
    if(dp < -EPSILON )
    {
        return ptBack;
    }
    return ptCoplanar; // it was between EP and -EP
}
```


Once you have code to classify a point, classifying other primitives, like polygons, becomes pretty trivial, as shown in the following code. The one issue is there are now four possible definition states when the element being tested isn't infinitesimally small. The element may be entirely in front of the plane, entirely in back, or perfectly coplanar. It may also be partially in front and partially in back. I'll refer to this state as *splitting* the plane. It's just a term; the element isn't actually splitting anything.

```
// Defines the four possible locations of a point list in
// relation to a plane. A point list is a more general
// example of a polygon.
enum ePListLoc
{
    plistFront,
    plistBack,
    plistSplit,
    plistCoplanar
};

ePListLoc plane3::TestPList( point3 *list, int num ) const
{
    bool allfront=true, allback=true;

    ePointLoc res;

    for( int i=0; i<num; i++ )
    {
        res = TestPoint( list[i] );

        if( res == ptBack )
        {
            allfront = false;
        }
        else if( res == ptFront )
        {
            allback = false;
        }
    }
    if( allfront && !allback )
    {
        // All the points were either in front or coplanar
        return plistFront;
    }
    else if( !allfront && allback )
    {
        // All the points were either in back or coplanar
        return plistBack;
    }
    else if( !allfront && !allback )
    {
        // Some were in front, some were in back
        return plistSplit;
    }
}
```

```
// All were coplanar
return plistCoplanar;
}
```

Back-face Culling

Now that you know how to define a point with respect to a plane, you can perform back-face culling, one of the most fundamental optimization techniques of 3D graphics.

Let's suppose you have a triangle whose elements are ordered in such a fashion that when viewing the triangle from the front, the elements appear in clockwise order. Back-face culling allows you to take triangles defined with this method and use the plane equation to discard triangles that are facing away. Conceptually, any closed mesh, a cube for example, will have some triangles facing you and some facing away. You know for a fact that you'll never be able to see a polygon that faces away from you; it is always hidden by triangles facing toward you. This, of course, doesn't hold if you're allowed to view the cube from its inside, but this shouldn't be allowed to happen if you want to really optimize your engine.

Rather than perform the work necessary to draw all of the triangles on the screen, you can use the plane equation to find out if a triangle is facing toward the camera, and discard it if it is not. How is this achieved? Given the three points of the triangle, you can define a plane that the triangle sits in. Since you know the elements of the triangle are listed in clockwise order, you also know that if you pass the elements in order to the plane constructor, the normal to the plane will be on the front side of the triangle. If you then think of the location of the camera as a point, all you need to do is perform a point-plane test. If the point of the camera is in front of the plane, then the triangle is visible and should be drawn.

There's an optimization to be had. Since you know three points that lie in the plane (the three points of the triangle) you only need to hold onto the normal of the plane, not the entire plane equation. To perform the back-face cull, just subtract one of the triangle's points from the camera location and perform a dot product with the resultant vector and the normal. If the result of the dot product is greater than zero, then the viewpoint was in front of the triangle. Figure 4.16 can help explain this.




Figure 4.16: A visual example of back-face culling

In practice, 3D accelerators actually perform back-face culling by themselves, so as the triangle rates of cards increase, the amount of manual back-face culling that is performed has steadily decreased. However, the information is useful for custom 3D engines that don't plan on using the facilities of Direct3D.

Clipping Lines

One thing that you'll need is the ability to take two points (a and b) that are on different sides of a plane defining a line segment, and find the point making the intersection of the line with the plane.

This is easy enough to do. Think of this parametrically. Point a can be thought of as the point at time 0 and point b as the point at time 1, and the point of intersection you want to find is somewhere between those two.

Take the dot product of a and b. Using them and the inverse of the plane's d parameter, you can find the scale value (which is a value between 0 and 1 that defines the parametric location of the particle when it intersects the plane). Armed with that, you just use the scale value, plugging it into the linear parametric equation to find the intersection location. Figure 4.17 shows this visually.




Figure 4.17: Finding the intersection of a plane and a line

```
inline const point3 plane3::Split( const point3 &a, const point3 &b ) const
{
    float aDot = (a * n);
    float bDot = (b * n);

    float scale = ( -d - aDot ) / ( bDot - aDot );

    return a + (scale * (b - a));
}
```

Clipping Polygons

Along with the ability to clip lines, you can now also clip polygons. Clipping polygons against planes is a common operation. You take a plane and a polygon and want to get a polygon in return that represents only the part of the input polygon that sits in front of the plane. Conceptually, you can think of the plane slicing off the part of the polygon that is behind it.

Clipping polygons are used principally in screen space clipping. If a polygon is sitting in a position such that when it was drawn it would be partially on screen and partially off screen, you want to clip the polygon such that you only draw the part of the polygon that would be sitting on the screen. Trying to draw primitives that aren't in the view can wreak havoc in many programs. Figure 4.18 shows the dilemma.




Figure 4.18: A polygon that needs to be clipped

To implement polygon clipping, I'll use the Sutherland-Hodgeman polygon clipping algorithm, discussed in section 3.14.1 of *Computer Graphics: Principles and Practice in C (2nd Ed.)* by James Foley, et al.

The algorithm is fairly straightforward. In a clockwise fashion, you wind all the way around the polygon, considering each adjacent pair of points. If the first point is on the front side of the plane (found using a plane to point classification call), you add it to the end of the outgoing polygon (it starts out empty). If the first and second vertices are on different sides, find the split point and add that to the list. While it may not intuitively seem obvious, the algorithm does work. The visual steps of it working appear in Figure 4.19. The function returns true if the clipped polygon is not degenerate (has three or more vertices).




Figure 4.19: Clipping using the Sutherland-Hodgeman algorithm

```
bool plane3::Clip( const polygon<point3> &in, polygon<point3> *out ) const
{
    // Make sure our pointer to the out polygon is valid
    assert( out );
    // Make sure we're not passed a degenerate polygon
    assert( in.nElem > 2 );

    int thisInd=in.nElem-1;
    int nextInd=0;

    ePointLoc thisRes = TestPoint( in.pList[thisInd] );
    ePointLoc nextRes;

    out->nElem = 0;

    for( nextInd=0; nextInd<in.nElem; nextInd++ )
    {
        nextRes = TestPoint( in.pList[nextInd] );

        if( thisRes == ptFront || thisRes == ptCoplanar )
        {
            // Add the point
            out->pList[out->nElem++] = in.pList[thisInd];
        }

        if( ( thisRes == ptBack && nextRes == ptFront ) ||
            ( thisRes == ptFront && nextRes == ptBack ) )
        {
            // Add the split point

```

```

        out->pList[out->nElem++] = Split(
            in.pList[thisInd],
            in.pList[nextInd] );
    }

    thisInd = nextInd;

    thisRes = nextRes;
}
if( out->nElem >= 3 )
{
    return true;
}
return false;
}

```

If you have code to take a polygon and clip off the area behind a plane, then creating a function to save the area behind the plane into an additional polygon isn't too hard. The operation takes a polygon that has elements lying on both sides of a plane and splits it into two distinct pieces, one completely in front of and one completely behind the plane. The BSP code at the end of this chapter uses polygon splitting. The algorithm to do this follows directly from the clipping code, and the code is very similar.

```

bool plane3::Split( polygon<point3> const &in, polygon<point3> *pFront,
                    polygon<point3> *pBack ) const
{
    // Make sure our pointer to the out polygon is valid
    assert( pFront );
    // Make sure our pointer to the out polygon is valid
    assert( pBack );
    // Make sure we're not passed a degenerate polygon
    assert( in.nElem > 2 );

    // Start with curr as the last vertex and next as 0.
    pFront->nElem = 0;
    pBack->nElem = 0;

    int thisInd=in.nElem-1;
    int nextInd=0;

    ePointLoc thisRes = TestPoint( in.pList[thisInd] );
    ePointLoc nextRes;

    for( nextInd=0; nextInd<in.nElem; nextInd++ ) {

        nextRes = TestPoint( in.pList[nextInd] );

        if( thisRes == ptFront )
        {
            // Add the point to the front

```

```

        pFront->pList[pFront->nElem++] = in.pList[thisInd];
    }

    if( thisRes == ptBack )
    {
        // Add the point to the back
        pBack->pList[pBack->nElem++] = in.pList[thisInd];
    }

    if( thisRes == ptCoplanar )
    {
        // Add the point to both
        pFront->pList[pFront->nElem++] = in.pList[thisInd];
        pBack->pList[pBack->nElem++] = in.pList[thisInd];
    }

    if( ( thisRes == ptBack && nextRes == ptFront ) ||
        ( thisRes == ptFront && nextRes == ptBack ) )
    {
        // Add the split point to both
        point3 split = Split(
            in.pList[thisInd],
            in.pList[nextInd] );
        pFront->pList[pFront->nElem++] = split;
        pBack->pList[pBack->nElem++] = split;
    }

    thisInd = nextInd;
    thisRes = nextRes;
}
if( pFront->nElem > 2 && pBack->nElem > 2 )
{
    // Nothing ended up degenerate
    return true;
}
return false;
}

```

Object Representations

Now that you have polygons and triangles, you can build objects. An *object* is just a boundary representation (with a few other traits, like materials, textures, and transformation matrices). Representing the boundary representations of objects is one of the ways that differentiate the myriad of 3D engines out there. There are many different ways to represent polygon data, each with its own advantages and disadvantages.

A big concern is that triangles and polygons need more information than just position if anything interesting is going to be drawn. Typically, the points that make up the triangle faces of an object are called *vertices*, to differentiate them from points or vectors. Vertices can have many different types of data in them besides position, from normal information (for smooth shading), to texture coordinates for texture mapping, to diffuse

and specular color information. I'll visit this point in Chapter 7 when I start showing you how to make 3D objects, but for right now keep in mind that the models will be more complex than just a list of points connecting a bunch of triangles.

An unsophisticated first approach to representing an object would be to explicitly list each triangle as a triplet of vertices. This method is bad for several reasons. The main reason is that generally the objects are made up of a closed mesh of triangles. They meet up and touch each other; each vertex is the meeting point of two or more triangles. While a cube actually has only eight vertices, this method would need three distinct vertices for each of the 12 triangles, a total of 36 vertices. Any amount of work to do per-vertex would have to be done four times more than if you had a representation with only eight vertices. Because of this downfall, this method isn't used much.




Figure 4.20: An object made of distinct triangles

However, it isn't without its advantages. For example, if the triangles are all distinct entities, you can do some neat effects, such as having the triangles fly off in separate directions when the object explodes (the game *MDK* did a good job with this; at the end of each level the world broke up into its component triangles and flew up into the sky).

Another big advantage that this method has is it allows triangles that share vertex locations to have different color, texture, and normal information. For example, if you have the eight-vertex cube, where each vertex had a position and a color, all the triangles that share each corner have the same color information for that corner. If you want each face of the cube to have a different color, you can use explicit vertices for each triangle.



Note: A better way to do this would be to only have explicit copies of the color information and just use one vector. However, this style of object representation doesn't work well with Direct3D.

If you don't need distinct information for each triangle, there is a much better way to represent the objects: with two lists. One is a list of vertices representing all of the vertices in the object, and one is a list of triangles, where each triangle is a triplet of integers, not points. The integers represent indices into the vertex list.




Figure 4.21: An object using shared triangles with indices

This is the method used by many 3D applications, and the method most preferred by Direct3D. Later in the book I'll create a format to represent objects of this type, and provide code both to load objects from disk and draw them.

These aren't the only two horses in town. Later I'll talk about objects where vertices need to know adjacency information (that is, which other vertices are connected to it by triangle edges). There are even more esoteric systems, like the quad-edge data structure, whose data structures barely resemble objects at all, essentially being a pure graph of nodes and edges (nodes represent vertices; triangles are represented by loops in the graph).

Transformations

Now that there are objects in the world, it would be good to be able to move them around the scene: animate them, spin them, and so forth. To do this you need to define a set of transformations that act upon the points in the objects. I'll start out with the simplest transformation: translation.

To move an object by a given vector p , all you need to do is add p to each of the points in the object. The translation transformation can be defined by a vector p as $T(p)$. The translation transformation is inverted easily. The transformation $T^{-1}(p)$ that undoes $T(p)$ is just $T(-p)$, in essence subtracting p from each point in the object.

Unfortunately, translation isn't terribly interesting on its own. It is also important to be able to rotate the objects around arbitrary points and axes as well as translating them. The next thing to do is add rotation transformations. Before doing that, however, I need to talk a little about matrices.

Matrices

A matrix is really just a shorthand way to write a set of simultaneous equations. For example, let's say you're trying to solve x , y , and z that satisfy the following three equations:

$$3x - 8y + 12z = 0$$

$$15x + 14y - 2z = 0$$

$$32x + 0.5y - z = 0$$

First, put all the coefficients of the equations into an n by m box called a *matrix*, where n (the vertical dimension) is the number of equations and m (the horizontal dimension) is the number of coefficients:

$$\begin{array}{l} 3x - 8y + 12z \\ 15x + 14y - 2z \\ 32x + 0.5y - z \end{array} \Rightarrow \begin{bmatrix} 3 & -8 & 12 \\ 15 & 14 & -2 \\ 32 & 0.5 & -1 \end{bmatrix}$$

Here's a 3x4 matrix:

$$\begin{bmatrix} a_{11} & a_{12} & a_{13} & a_{14} \\ a_{21} & a_{22} & a_{23} & a_{24} \\ a_{31} & a_{32} & a_{33} & a_{34} \end{bmatrix}$$

The subscript notation used above is how we reference individual elements of a matrix. The first component is the row number, and the second component is the column number.

Matrices can be added together simply by adding each component. However, the matrices must be the same size to be able to add them (you couldn't, for example, add a 3x3 and a 2x2 matrix together).

$$\begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} + \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix} = \begin{bmatrix} a_{11} + b_{11} & a_{12} + b_{12} \\ a_{21} + b_{21} & a_{22} + b_{22} \end{bmatrix}$$

Multiplying matrices together is a bit more involved. To find $AB=C$, each component c_{ij} of the resultant matrix is found by computing the dot product of the i^{th} row of A with the j^{th} column of B. The rules for matrix sizes are different from those in addition. If A is m by n and B is o by p , the multiplication is only valid if $n = o$, and the dimension of the resultant matrix is m by p . Note that multiplication only is valid if the row length of matrix A is the same as the column length of matrix B.

$$\begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix} = \begin{bmatrix} a_{11}b_{11} + a_{12}b_{21} & a_{11}b_{12} + a_{12}b_{22} \\ a_{21}b_{11} + a_{22}b_{21} & a_{21}b_{12} + a_{22}b_{22} \end{bmatrix}$$

Another example (3x3 times 3x1 yields 3x1):

$$\begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix} \begin{bmatrix} b_{11} \\ b_{21} \\ b_{31} \end{bmatrix} = \begin{bmatrix} a_{11}b_{11} + a_{12}b_{21} + a_{13}b_{31} \\ a_{21}b_{11} + a_{22}b_{21} + a_{23}b_{31} \\ a_{31}b_{11} + a_{32}b_{21} + a_{33}b_{31} \end{bmatrix}$$

This way it is easy to represent the problem above of trying to solve a matrix equation. If you multiply out the matrices below into three simultaneous equations, you'll get the same three above.

$$\begin{bmatrix} 3 & -8 & 12 \\ 15 & 14 & -2 \\ 32 & 0.5 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}$$



Warning: Note that multiplication is not commutative. That is, AB is not the same as BA .

Matrix multiplication has an identity value, just like scalar multiplication (which has an identity of 1). The identity is only defined for square matrices, however. It is defined as a zeroed-out matrix with ones running down the diagonal. Here is the 3x3 identity matrix I_3 :

$$I_3 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

Matrix multiplication also has the law of associativity going for it. That means that as long as you preserve left-to-right order, you can multiply matrix pairs together in any order:

$$ABCD = A(BC)D = (AB)(CD) = (((AB)C)D)$$

This will come into play later; right now just keep it in the back of your head.

What does all this have to do with anything? Very good question. Matrices can be used to represent transformations, specifically rotations. You can represent rotations with 3x3 matrices and points as 1x3 matrices, multiplying them together to get transformed vertices.

$$vA = v'$$

$$\begin{bmatrix} x & y & z \end{bmatrix} \begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{bmatrix} = \begin{bmatrix} x' & y' & z' \end{bmatrix}$$

There are three standard matrices to facilitate rotations about the x-, y-, and z-axes by some angle theta. They are:

$$R_x(\theta) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(\theta) & \sin(\theta) \\ 0 & -\sin(\theta) & \cos(\theta) \end{bmatrix}$$

$$R_y(\theta) = \begin{bmatrix} \cos(\theta) & 0 & -\sin(\theta) \\ 0 & 1 & 0 \\ \sin(\theta) & 0 & \cos(\theta) \end{bmatrix}$$

$$R_z(\theta) = \begin{bmatrix} \cos(\theta) & \sin(\theta) & 0 \\ -\sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

To show this happening, let's manually rotate the point $\langle 2, 0, 0 \rangle$ 45 degrees clockwise about the z-axis.

$$v' = R_z(45)v$$

$$v' = \begin{bmatrix} 0.707 & 0.707 & 0 \\ -0.707 & 0.707 & 0 \\ 0 & 0 & 1 \end{bmatrix} \times \begin{bmatrix} 2 \\ 0 \\ 0 \end{bmatrix}$$

$$v' = \begin{bmatrix} 2 \times 0.707 + 0 \times 0.707 + 0 \times 0 \\ 2 \times -0.707 + 0 \times 0.707 + 0 \times 0 \\ 2 \times 0 + 0 \times 0 + 0 \times 0 \end{bmatrix}$$

$$v' = \begin{bmatrix} 1.414 \\ -1.414 \\ 0 \end{bmatrix}$$

Now you can take an object and apply a sequence of transformations to it to make it do whatever you want. All you need to do is figure out the sequence of transformations needed and then apply the sequence to each of the points in the model.

As an example, let's say you want to rotate an object sitting at a certain point p around its z -axis. You would perform the following sequence of transformations to achieve this:

$$\mathbf{v} = \mathbf{vT}(-\mathbf{p})$$

$$\mathbf{v} = \mathbf{vR}_z\left(\frac{\pi}{2}\right)$$

$$\mathbf{v} = \mathbf{vT}(\mathbf{p})$$

The first transformation moves a point such that it is situated about the world origin instead of being situated about the point p . The next one rotates it (remember, you can only rotate about the origin, not arbitrary points in space). Finally, after the point is rotated, you want to move it back so that it is situated about p . The final translation accomplishes this.

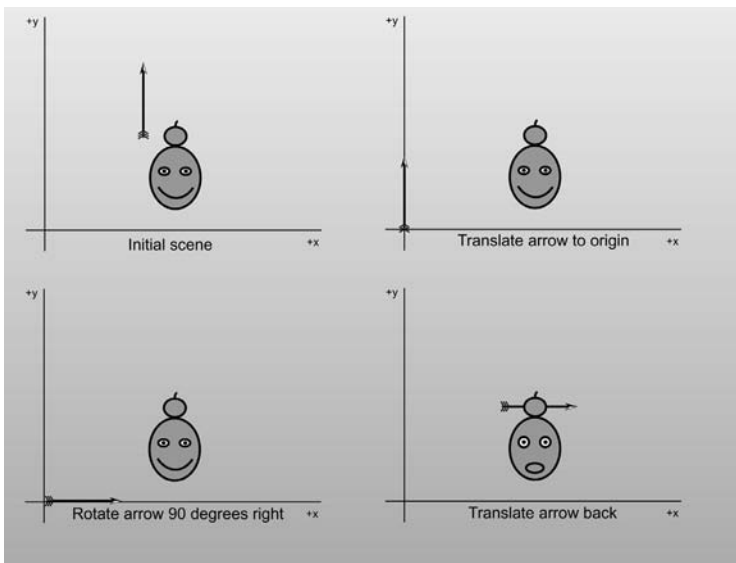


Figure 4.22: Compound transformations

Notice the difference between a rotation followed by a translation and a translation followed by a rotation.

You would be set now, except for one small problem: Doing things this way is kind of slow. There may be dozens of transformations to perform on an object, and if the object has thousands of points, that is dozens of thousands of transformations that need to be trudged through.

The nice thing about matrices is that they can be concatenated together before they are multiplied by points. If there are two rotations, A and B, you know from the associativity law:

$$\mathbf{v}' = (\mathbf{vA})\mathbf{B} \Rightarrow \mathbf{v}' = \mathbf{v}(\mathbf{AB})$$

So before multiplying each of the points by both rotation transformations, you multiply them together into one matrix that represents both rotations, and just multiply the points by the new matrix. If you could also represent translations as matrices, you could concatenate the entire string of matrices together into one big matrix, cutting down on the transformation work quite a bit.

There's a problem: 3x3 matrices can't encode translation. A translation is just an addition by another vector, and because of the semantics of matrix multiplication, you just can't make a 3x3 matrix that adds a vector to an input one.

The way the graphics, robotics, mathematics, and physics communities have solved this problem is to introduce a fourth component to the vectors and an added dimension to the matrices, making them 4x4.

The fourth coordinate is called the *homogenous coordinate*, and is represented with the letter *w*. There are an infinite number of 4D homogenous coordinates for any 3D Cartesian coordinate you can supply. The space of homogenous coordinates given a Cartesian coordinate is defined as this:

$$\begin{bmatrix} x & y & z \end{bmatrix} \Rightarrow \begin{bmatrix} bx & by & bz & b \end{bmatrix} \quad (\text{for all } b \neq 0)$$

To reclaim a Cartesian coordinate from a homogenous coordinate, just make sure the *w* component is 1, and then get the *x*, *y*, and *z* values. If *w* isn't 1, then divide all four components by *w* (removing the *b* from the equation).

Now you can change the translation transformation to a 4x4 matrix:

$$T(\mathbf{p}) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ p_x & p_y & p_z & 1 \end{bmatrix}$$

Note that multiplication by this matrix has the desired behavior:

$$\begin{bmatrix} x & y & z & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ p_x & p_y & p_z & 1 \end{bmatrix} = \begin{bmatrix} x + p_x & y + p_y & z + p_z & 1 \end{bmatrix}$$

The identity and rotation matrices change too, to reflect the added dimension:

$$\mathbf{I}_4 = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

$$\mathbf{R}_x(\theta) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta) & \sin(\theta) & 0 \\ 0 & -\sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

$$\mathbf{R}_y(\theta) = \begin{bmatrix} \cos(\theta) & 0 & -\sin(\theta) & 0 \\ 0 & 1 & 0 & 0 \\ \sin(\theta) & 0 & \cos(\theta) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

$$\mathbf{R}_z(\theta) = \begin{bmatrix} \cos(\theta) & \sin(\theta) & 0 & 0 \\ -\sin(\theta) & \cos(\theta) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

Now that you know how to represent all of the transformations with matrices, you can concatenate them together, saving a load of time and space. This also changes the way you might think about transformations. Each object defines all of its points with respect to a local coordinate system, with the origin representing the center of rotation for the object. Each object also has a matrix, which transforms the points from the local origin to some location in the world. When the object is moved, the matrix can be manipulated to move the points to a different location in the world.

To understand what is going on here, you need to modify the way you perceive matrix transformations. Rather than translate or rotate, they actually become maps from one coordinate space to another. The object is defined in one coordinate space (which is generally called the object's local coordinate space), and the object's matrix maps all of the points to a new location in another coordinate space, which is generally the coordinate space for the entire world (generally called the world coordinate space).

A nice feature of matrices is that it's easy to see where the matrix that transforms from object space to world space is sitting in the world. If you look at the data the right way, you can actually see where the object axes get mapped into the world space.

Consider four vectors, called \mathbf{n} , \mathbf{o} , \mathbf{a} , and \mathbf{p} . The \mathbf{p} vector represents the location of the object coordinate space with relation to the world origin. The \mathbf{n} , \mathbf{o} , and \mathbf{a} vectors represent the orientation of the \mathbf{i} , \mathbf{j} , and \mathbf{k} vectors, respectively.




Figure 4.23: The n , o , a , and p vectors for a transformation

You can get and set these vectors right in the matrix, as they are sitting there in plain sight:

$$\begin{bmatrix} n_x & n_y & n_z & 0 \\ o_x & o_y & o_z & 0 \\ a_x & a_y & a_z & 0 \\ p_x & p_y & p_z & 1 \end{bmatrix}$$

This system of matrix concatenations is how almost all 3D applications perform their transformations. There are four spaces that points can live in: object space, world space, and two new spaces: view space and screen space.

View space defines how images on the screen are displayed. Think of it as a camera. If you move the camera around the scene, the view will change. You see what is in front of the camera (in front is defined as positive z).




Figure 4.24: Mapping from world space to view space

The transformation here is different from the one used to move from object space to world space. Now, while the camera is defined with the same *n*, *o*, *a*, and *p* vectors as defined with the other transforms, the matrix itself is different.

In fact, the view matrix is the inversion of what the object matrix for that position and orientation would be. This is because you're performing a backward transformation: taking points once they're in world space and putting them into a local coordinate space.

As long as you compose the transformations of just rotations and translations (and reflections, by the way, but that comes into play much later in the book), computing the inverse of a transformation is easy. Otherwise, computing an inverse is considerably more difficult and may not even be possible. The inverse of a transformation matrix is given below.

$$\begin{bmatrix} \mathbf{n}_x & \mathbf{n}_y & \mathbf{n}_z & 0 \\ \mathbf{o}_x & \mathbf{o}_y & \mathbf{o}_z & 0 \\ \mathbf{a}_x & \mathbf{a}_y & \mathbf{a}_z & 0 \\ \mathbf{p}_x & \mathbf{p}_y & \mathbf{p}_z & 1 \end{bmatrix}^{-1} = \begin{bmatrix} \mathbf{n}_x & \mathbf{o}_x & \mathbf{a}_x & 0 \\ \mathbf{n}_y & \mathbf{o}_y & \mathbf{a}_y & 0 \\ \mathbf{n}_z & \mathbf{o}_z & \mathbf{a}_z & 0 \\ -(\mathbf{p} \cdot \mathbf{n}) & -(\mathbf{p} \cdot \mathbf{o}) & -(\mathbf{p} \cdot \mathbf{a}) & 1 \end{bmatrix}$$



Warning: This formula for inversion is not universal for all matrices. In fact, the only matrices that can be inverted this way are ones composed exclusively of rotations, reflections, and translations.

There is a final transformation that the points must go through in the transformation process. This transformation maps 3D points defined with respect to the view origin (in view space) and turns them into 2D points that can be drawn on the display. After transforming and clipping the polygons that make up the scene such that they are visible on the screen, the final step is to move them into 2D coordinates, since in order to actually draw things on the screen you need to have absolute x,y coordinates on the screen to draw.

The way this used to be done was without matrices, just as an explicit projection calculation. The point $\langle x,y,z \rangle$ would be mapped to $\langle x',y' \rangle$ using the following equations:

$$x' = \text{scale} \frac{x}{z} + x_{\text{Center}}$$

$$y' = \text{height} - \left(\text{scale} \frac{y}{z} + y_{\text{Center}} \right)$$

where x_{Center} and y_{Center} were half of the width and height of the screen, respectively. These days more complex equations are used, especially since there is now the need to make provisions for z-buffering. While you want x and y to still behave the same way, you don't want to use a value as arbitrary as scale.

Instead, a better value to use in the calculation of the projection matrix is the horizontal field of view (fov). The horizontal fov will be hardcoded, and the code chooses a vertical field of view that will keep the aspect ratio of the screen. This makes sense: You couldn't get away with using the same field of view for both horizontal and vertical directions unless the screen was square; it would end up looking vertically squished.

Finally, you also want to scale the z values appropriately. We'll see more on z-buffers in Chapters 7 and 9, but for right now just make note of an important feature: They let you clip out certain values of z-range. Given the two variables z_{near} and z_{far} , nothing in front of z_{near} will be drawn, nor will anything behind z_{far} . To make the z-buffer work swimmingly on all ranges of z_{near} and z_{far} you need to scale the valid z values to the range of 0.0 to 1.0.

For purposes of continuity, I'll use the same projection matrix definition that Direct3D recommends in the documentation. First, let's define some values. You initially start with the width and height of the viewport and the horizontal field of view.

$$\begin{aligned} \text{aspect} &= \frac{\text{height}}{\text{width}} \\ w &= \text{aspect} \frac{\cos(\text{fov})}{\sin(\text{fov})} \\ h &= \frac{\cos(\text{fov})}{\sin(\text{fov})} \\ q &= \frac{z_{\text{far}}}{z_{\text{far}} - z_{\text{near}}} \end{aligned}$$

With these parameters, the following projection matrix can be made:

$$\begin{bmatrix} w & 0 & 0 & 0 \\ 0 & h & 0 & 0 \\ 0 & 0 & q & 1 \\ 0 & 0 & -q(z_{\text{near}}) & 0 \end{bmatrix}$$

Just for a sanity check, check out the result of this matrix multiplication:

$$\begin{bmatrix} x & y & z & 1 \end{bmatrix} \begin{bmatrix} w & 0 & 0 & 0 \\ 0 & h & 0 & 0 \\ 0 & 0 & q & 1 \\ 0 & 0 & -q(z_{\text{near}}) & 0 \end{bmatrix} = \begin{bmatrix} wx & hy & qz - q(z_{\text{near}}) & z \end{bmatrix}$$

Hmm... this is almost the result wanted, but there is more work to be done. Remember that in order to extract the Cartesian (x,y,z) coordinates from the vector, the homogenous w component must be 1.0. Since, after the multiplication, it's set to z (which can be any value), all four components need to be divided by w to normalize it. This gives the following Cartesian coordinate:

$$\begin{bmatrix} \frac{wx}{z} & \frac{hy}{z} & q\left(1 - \frac{(z_{\text{near}})}{z}\right) & 1 \end{bmatrix}$$

As you can see, this is exactly what was wanted. The width and height are still scaled by values as in the above equation and they are still divided by z . The visible x and y pixels are mapped to $[-1,1]$, so before rasterization Direct3D multiplies and adds the number by x_{Center} or y_{Center} . This, in essence, maps the coordinates from $[-1,1]$ to $[0,\text{width}]$ and $[0,\text{height}]$.

With this last piece of the puzzle, it is now possible to create the entire transformation pipeline. When you want to render a scene, you set up a world matrix (to transform an object's local coordinate points into world space), a view matrix (to transform world coordinate points into a space relative to the viewer), and a projection matrix (to take those viewer-relative points and project them onto a 2D surface so that they can be drawn on the screen). You then multiply the world, view, and projection matrices together (in that order) to get a total matrix that transforms points from object space to screen space.

$$\begin{aligned} \mathbf{v}_{\text{world}} &= \mathbf{v}_{\text{local}} \mathbf{M}_{\text{world}} \\ \mathbf{v}_{\text{view}} &= \mathbf{v}_{\text{world}} \mathbf{M}_{\text{view}} \\ \mathbf{v}_{\text{screen}} &= \mathbf{v}_{\text{view}} \mathbf{M}_{\text{projection}} \\ \mathbf{v}_{\text{screen}} &= \mathbf{v}_{\text{local}} (\mathbf{M}_{\text{world}} \mathbf{M}_{\text{view}} \mathbf{M}_{\text{projection}}) \end{aligned}$$



Warning: OpenGL uses a different matrix convention (where vectors are column vectors, not row vectors, and all matrices are transposed). If you're used to OpenGL, the equation above will seem backward. This is the convention that Direct3D uses, so to avoid confusion, it's what is used here.

To draw a triangle, for example, you would take its local space points defining its three corners and multiply them by the transformation matrix. Then you have to remember to divide through by the w component and voilà! The points are now in screen space and can be filled in using a 2D raster algorithm. Drawing multiple objects is a snap, too. For each object in the scene all you need to do is change the world matrix and reconstruct the total transformation matrix.

The matrix4 Structure

Now that all the groundwork has been laid out to handle transformations, let's actually write some code. The struct is called `matrix4`, because it represents 4D homogenous transformations. Hypothetically, if you wanted to just create rotation matrices, you could do so with a class called `matrix3`.

```
struct matrix4
{
    /**
     * we're using m[y][x] as our notation.
     */
    union
    {
        struct
        {
            float  _11, _12, _13, _14;
            float  _21, _22, _23, _24;
            float  _31, _32, _33, _34;
            float  _41, _42, _43, _44;
        };
        float  m[4][4];
    };

    // justification for a function this ugly:
    // provides an easy way to initialize static matrix variables
    // like base matrices for bezier curves and the identity
    matrix4(float IN_11, float IN_12, float IN_13, float IN_14,
            float IN_21, float IN_22, float IN_23, float IN_24,
            float IN_31, float IN_32, float IN_33, float IN_34,
            float IN_41, float IN_42, float IN_43, float IN_44)
    {
        _11 = IN_11; _12 = IN_12; _13 = IN_13; _14 = IN_14;
        _21 = IN_21; _22 = IN_22; _23 = IN_23; _24 = IN_24;
        _31 = IN_31; _32 = IN_32; _33 = IN_33; _34 = IN_34;
        _41 = IN_41; _42 = IN_42; _43 = IN_43; _44 = IN_44;
    }
};
```

```

    }

    matrix4()
    {
        // Do nothing.
    }

    static const matrix4 Identity;
};

```

The code below contains three main ways to multiply matrices. Two 4x4 matrices can be multiplied together; this is useful for concatenating matrices. A point4 structure can be multiplied by a matrix4 structure; the result is the application of the transformation to the 4D point. Finally, a specialization for multiplying point3 structures and matrix4 structures exists to apply a non-projection transformation to a point3 structure. The matrix4*matrix4 operator creates a temporary structure to hold the result, and isn't terribly fast. Matrix multiplications aren't performed often enough for this to be much of a concern, however.



Warning: If you plan on doing a lot of matrix multiplications per object or even per triangle, you won't want to use the operator. Use the provided MatMult function; it's faster.

```

matrix4 operator*(matrix4 const &a, matrix4 const &b)
{
    matrix4 out;
    for (int j = 0; j < 4; j++) // temporary matrix4 for storing result
        for (int i = 0; i < 4; i++) // transform by columns first
            // then by rows
                out.m[i][j] = a.m[i][0] * b.m[0][j] +
                    a.m[i][1] * b.m[1][j] +
                    a.m[i][2] * b.m[2][j] +
                    a.m[i][3] * b.m[3][j];

    return out;
};

inline const point4 operator*( const matrix4 &a, const point4 &b)
{
    return point4(
        b.x*a._11 + b.y*a._21 + b.z*a._31 + b.w*a._41,
        b.x*a._12 + b.y*a._22 + b.z*a._32 + b.w*a._42,
        b.x*a._13 + b.y*a._23 + b.z*a._33 + b.w*a._43,
        b.x*a._14 + b.y*a._24 + b.z*a._34 + b.w*a._44
    );
};

inline const point4 operator*( const point4 &a, const matrix4 &b)
{
    return b*a;
};

inline const point3 operator*( const matrix4 &a, const point3 &b)

```

```

{
    return point3(
        b.x*a._11 + b.y*a._21 + b.z*a._31 + a._41,
        b.x*a._12 + b.y*a._22 + b.z*a._32 + a._42,
        b.x*a._13 + b.y*a._23 + b.z*a._33 + a._43
    );
};

inline const point3 operator*( const point3 &a, const matrix4 &b)
{
    return b*a;
};

```

There are two ways to create each type of matrix transformation. One performs on an existing matrix4 structure (it doesn't create a temporary matrix4 structure, which is slow). The function for a transformation x is `void matrix4::Tox`. The other is a static function designed to help write cleaner looking code, not for speed. The format for these functions is `static matrix4 matrix4::x`.

Translation

Here again is the matrix for the translation transformation by a given point p :

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ p_x & p_y & p_z & 1 \end{bmatrix}$$

The code to create this type of transformation matrix follows.

```

void matrix4::ToTranslation( const point3& p )
{
    MakeIdent();
    _41 = p.x;
    _42 = p.y;
    _43 = p.z;
}

matrix4 matrix4::Translation( const point3& p )
{
    matrix4 out;
    out.ToTranslation( p );
    return out;
}

```

Basic Rotations

The matrices used to rotate around the three principal axes, again, are:

$$R_x(\theta) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta) & \sin(\theta) & 0 \\ 0 & -\sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

$$R_y(\theta) = \begin{bmatrix} \cos(\theta) & 0 & -\sin(\theta) & 0 \\ 0 & 1 & 0 & 0 \\ \sin(\theta) & 0 & \cos(\theta) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

$$R_z(\theta) = \begin{bmatrix} \cos(\theta) & \sin(\theta) & 0 & 0 \\ -\sin(\theta) & \cos(\theta) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

The code to set up Euler rotation matrices follows.

```
void matrix4::ToXRot( float theta )
{
    float c = (float) cos(theta);
    float s = (float) sin(theta);
    MakeIdent();
    _22 = c;
    _23 = s;
    _32 = -s;
    _33 = c;
}

matrix4 matrix4::XRot( float theta )
{
    matrix4 out;
    out.ToXRot( theta );
    return out;
}

//=====-----

void matrix4::ToYRot( float theta )
{
    float c = (float) cos(theta);
    float s = (float) sin(theta);
    MakeIdent();
    _11 = c;
    _13 = -s;
    _31 = s;
    _33 = c;
}

matrix4 matrix4::YRot( float theta )
```



```

{
    matrix4 out;
    out.ToYRot( theta );
    return out;
}

//=====-----

void matrix4::ToZRot( float theta )
{
    float c = (float) cos(theta);
    float s = (float) sin(theta);
    MakeIdent();
    _11 = c;
    _12 = s;
    _21 = -s;
    _22 = c;
}

matrix4 matrix4::ZRot( float theta )
{
    matrix4 out;
    out.ToZRot( theta );
    return out;
}

```

Axis-Angle Rotation

While there isn't enough space to provide a derivation of the axis-angle rotation matrix, that doesn't stop it from being cool. Axis-angle rotations are the most useful matrix-based rotation. (I say matrix-based because quaternions are faster and more flexible than matrix rotations; see *Real-Time Rendering* by Tomas Möller and Eric Haines for a good discussion on them.)

There are a few problems with using just Euler rotation matrices (the x-rotation, y-rotation, z-rotation matrices you've seen thus far). For starters, there really is no standard way to combine them together.

Imagine that you want to rotate an object around all three axes by three angles. In which order should the matrices be multiplied together? Should the x-rotation come first? The z-rotation? Since no answer is technically correct, usually people pick the one convention that works best for them and stick with it.

A worse problem is that of *gimbal lock*. To explain, look at how rotation matrices are put together. There are really two ways to use rotation matrices. Method 1 is to keep track of the current yaw, pitch, and roll rotations, and build a rotation matrix every frame. Method 2 uses the rotation matrix from the last frame, by just rotating it a small amount to represent any rotation that happened since the last frame.

The second method, while it doesn't suffer from gimbal lock, suffers from other things, namely the fact that all that matrix multiplication brings up some numerical imprecision issues. The *i*, *j*, and *k* vectors of your

matrix *gradually* become non-unit-length and not mutually perpendicular. This is a bad thing. However, there are ways to fix it that are pretty standard, such as renormalizing the vectors, using cross products to assure *orthogonality*.

Gimbal lock pops up when you're using the first method detailed above. Imagine that you perform a yaw rotation first, then pitch, then roll. Also, say that the yaw and pitch rotations are both a quarter-turn (this could come up quite easily in a game like *Descent*). So imagine you perform the first rotation, which takes you from pointing forward to pointing up. The second rotation spins you around the y-axis 90 degrees, so you're still facing up but your up direction is now to the right, not backward.

Now comes the lock. When you go to do the roll rotation, which way will it turn you? About the z-axis, of course. However, given any roll value, you can reach the same final rotation just by changing yaw or pitch. So essentially, you have lost a degree of freedom. This, as you would expect, is bad.

Axis-angle rotations fix both of these problems by doing rotations much more intuitively. You provide an axis that you want to rotate around and an angle amount to rotate around that axis. Simple. The actual matrix to do it, which appears below, isn't quite as simple, unfortunately. For sanity's sake, just treat it as a black box. See *Real-Time Rendering* (Möller and Haines) for a derivation of how this matrix is constructed.

$$\begin{bmatrix} xx(1 - \cos(\theta)) + \cos(\theta) & yx(1 - \cos(\theta)) + z \sin(\theta) & xz(1 - \cos(\theta)) - y \sin(\theta) & 0 \\ xy(1 - \cos(\theta)) - z \sin(\theta) & yy(1 - \cos(\theta)) + \cos(\theta) & yz(1 - \cos(\theta)) + x \sin(\theta) & 0 \\ xz(1 - \cos(\theta)) + y \sin(\theta) & yz(1 - \cos(\theta)) - x \sin(\theta) & zz(1 - \cos(\theta)) + \cos(\theta) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

The code to create an axis-angle matrix follows.

```
void matrix4::ToAxisAngle( const point3& inAxis, float angle )
{
    point3 axis = inAxis.Normalized();
    float s = (float)sin( angle );
    float c = (float)cos( angle );
    float x = axis.x, y = axis.y, z = axis.z;

    _11 = x*x*(1-c)+c;
    _21 = x*y*(1-c)-(z*s);
    _31 = x*z*(1-c)+(y*s);
    _41 = 0;
    _12 = y*x*(1-c)+(z*s);
    _22 = y*y*(1-c)+c;
    _32 = y*z*(1-c)-(x*s);
    _42 = 0;
    _13 = z*x*(1-c)-(y*s);
    _23 = z*y*(1-c)+(x*s);
    _33 = z*z*(1-c)+c;
    _43 = 0;
    _14 = 0;
    _24 = 0;
```

```

    _34 = 0;
    _44 = 1;

}

matrix4 matrix4::AxisAngle( const point3& axis, float angle )
{
    matrix4 out;
    out.ToAxisAngle( axis, angle );
    return out;
}

```

The LookAt Matrix

I discussed earlier that the first three components of the first three rows (the *n*, *o*, and *a* vectors) make up the three principal axes (*i*, *j*, and *k*) of the coordinate space that the matrix represents. I am going to use this to make a matrix that represents a transformation of an object looking a particular direction. This is useful in many cases and is most often used in controlling the camera. Usually, there is a place where the camera is and a place you want the camera to focus on. You can accomplish this using an inverted LookAt matrix (you need to invert it because the camera transformation brings points from world space to view space, not the other way around, like object matrices).

There is one restriction regarding the LookAt matrix: It always assumes that there is a constant up vector, and the camera orients itself to that, so there is no tilt. For the code to work, the camera cannot be looking in the same direction that the up vector points. This is because a cross product is performed with the view vector and the up vector, and if they're the same thing the behavior of the cross product is undefined. In games like *Quake III: Arena*, you can look almost straight up, but there is some infinitesimally small epsilon that prevents you from looking in the exact direction.

Three vectors are passed into the function: a location for the matrix to be, a target to look at, and the up vector (the third parameter will default to *j* <0,1,0> so you don't need to always enter it). The transformation vector for the matrix is simply the location. The *a* vector is the normalized vector representing the target minus the location (or a vector that is the direction you want the object to look in). To find the *n* vector, simply take the normalized cross product of the up vector and the direction vector. (This is why they can't be the same vector; the cross product would return garbage.) Finally, you can get the *o* vector by taking the cross product of the *n* and *a* vectors already found.

I'll show you two versions of this transformation, one to compute the matrix for an object to world transformation, and one that computes the inverse automatically. Use `ObjectLookAt` to make object matrices that look in certain directions, and `CameraLookAt` to make cameras that look in certain directions.

```

void matrix4::ToObjectLookAt(
    const point3& loc,
    const point3& lookAt,
    const point3& inUp )
{

    point3 viewVec = lookAt - loc;
    float mag = viewVec.Mag();
    viewVec /= mag;

    float fDot = inUp * viewVec;
    point3 upVec = inUp - fDot * viewVec;
    upVec.Normalize();

    point3 rightVec = upVec ^ viewVec;

    // The first three rows contain the basis
    // vectors used to rotate the view to point at the lookat point
    _11 = rightVec.x;    _21 = upVec.x;    _31 = viewVec.x;
    _12 = rightVec.y;    _22 = upVec.y;    _32 = viewVec.y;
    _13 = rightVec.z;    _23 = upVec.z;    _33 = viewVec.z;

    // Do the translation values
    _41 = loc.x;
    _42 = loc.y;
    _43 = loc.z;

    _14 = 0;
    _24 = 0;
    _34 = 0;
    _44 = 1;

}

matrix4 matrix4::ObjectLookAt(
    const point3& loc,
    const point3& lookAt,
    const point3& inUp )
{
    matrix4 out;
    out.ToObjectLookAt( loc, lookAt, inUp );
    return out;
}

//=====-----

void matrix4::ToCameraLookAt(
    const point3& loc,
    const point3& lookAt,
    const point3& inUp )
{
    point3 viewVec = lookAt - loc;
    float mag = viewVec.Mag();
    viewVec /= mag;

```

```

float fDot = inUp * viewVec;
point3 upVec = inUp - fDot * viewVec;
upVec.Normalize();

point3 rightVec = upVec ^ viewVec;

// The first three columns contain the basis
// vectors used to rotate the view to point
// at the lookat point
_11 = rightVec.x;   _12 = upVec.x;   _13 = viewVec.x;
_21 = rightVec.y;   _22 = upVec.y;   _23 = viewVec.y;
_31 = rightVec.z;   _32 = upVec.z;   _33 = viewVec.z;

// Do the translation values
_41 = - (loc * rightVec);
_42 = - (loc * upVec);
_43 = - (loc * viewVec);

_14 = 0;
_24 = 0;
_34 = 0;
_44 = 1;
}

matrix4 matrix4::CameraLookAt(
    const point3& loc,
    const point3& lookAt,
    const point3& inUp )
{
    matrix4 out;
    out.ToCameraLookAt( loc, lookAt, inUp );
    return out;
}

```

Perspective Projection Matrix

Creating a perspective projection matrix will be handled by the graphics layer when I add Direct3D to it in Chapter 7, using the matrix discussed earlier in that chapter.

Inverse of a Matrix

Again, the inverse of a matrix composed solely of translations, rotations, and reflections (scales such as $\langle 1, 1, -1 \rangle$ that flip sign but don't change the length) can be computed easily. The inverse matrix looks like this:

$$\begin{bmatrix} \mathbf{n}_x & \mathbf{n}_y & \mathbf{n}_z & 0 \\ \mathbf{o}_x & \mathbf{o}_y & \mathbf{o}_z & 0 \\ \mathbf{a}_x & \mathbf{a}_y & \mathbf{a}_z & 0 \\ \mathbf{p}_x & \mathbf{p}_y & \mathbf{p}_z & 1 \end{bmatrix}^{-1} = \begin{bmatrix} \mathbf{n}_x & \mathbf{o}_x & \mathbf{a}_x & 0 \\ \mathbf{n}_y & \mathbf{o}_y & \mathbf{a}_y & 0 \\ \mathbf{n}_z & \mathbf{o}_z & \mathbf{a}_z & 0 \\ -(\mathbf{p} \bullet \mathbf{n}) & -(\mathbf{p} \bullet \mathbf{o}) & -(\mathbf{p} \bullet \mathbf{a}) & 1 \end{bmatrix}$$

Here is the code to perform the inversion:

```
void matrix4::ToInverse( const matrix4& in )
{
    // first transpose the rotation matrix
    _11 = in._11;
    _12 = in._21;
    _13 = in._31;
    _21 = in._12;
    _22 = in._22;
    _23 = in._32;
    _31 = in._13;
    _32 = in._23;
    _33 = in._33;

    // fix right column
    _14 = 0;
    _24 = 0;
    _34 = 0;
    _44 = 1;

    // now get the new translation vector
    point3 temp = in.GetLoc();

    _41 = -(temp.x * in._11 + temp.y * in._12 + temp.z * in._13);
    _42 = -(temp.x * in._21 + temp.y * in._22 + temp.z * in._23);
    _43 = -(temp.x * in._31 + temp.y * in._32 + temp.z * in._33);
}

matrix4 matrix4::Inverse( const matrix4& in )
{
    matrix4 out;
    out.ToInverse( in );
    return out;
}
```

Collision Detection with Bounding Spheres

Up until now, when I talked about moving 3D objects around, I did so completely oblivious to wherever they may be moving. But suppose there is a sphere slowly moving through the scene. During its journey it collides with another object (for the sake of simplicity, say another sphere). You generally want the reaction that results from the collision to be at least somewhat similar to what happens in the real world.

In the real world, depending on the mass of the spheres, the amount of force they absorb, the air resistance in the scene, and a slew of other factors, they will physically react to each other the moment they collide. If they were rubber balls, they may bounce off of each other. If the spheres were instead made of crazy glue, they would not bounce at all, but would become inextricably attached to each other. Physics simulation aside, you

most certainly do not want to allow any object to blindly fly through another object (unless, of course, that is the effect you're trying to achieve, such as an apparition object like the ghosts in *Super Mario Brothers* games).

There are a million and one ways to handle collisions and the method you use will be very implementation dependent. So for now, all I'm going to discuss here is just getting a rough idea of when a collision has occurred. Most of the time, games only have the horsepower to do very quick and dirty collision detection. Games generally use *bounding boxes* or *bounding spheres* to accomplish this; I'm going to talk about bounding spheres. They try to simplify complex graphics tasks like occlusion and collision detection.

The general idea is that instead of performing tests against possibly thousands of polygons in an object, you can simply hold on to a sphere that approximates the object, and just test against that. Testing a plane or point against a bounding sphere is a simple process, requiring only a subtraction and a vector comparison. When the results you need are approximate, using bounding objects can speed things up nicely. This gives up the ability to get exact results, however. Fire up just about any game and try to just miss an object with a shot. Chances are (if you're not playing something with great collision detection like *MDK*, *Goldeneye*, or *House of the Dead*) you'll hit your target anyway. Most of the time you don't even notice, so giving up exact results isn't a tremendous loss.

Even if you do need exact results, you can still use bounding objects. They allow you to perform trivial rejection. An example is in collision detection. Typically, calculating collision detection exactly is an expensive process (it can be as bad as $O(mn)$, where m and n are the number of polygons in each object). If you have multiple objects in the scene, you need to perform collision tests between all of them, a total of $O(n^2)$ operations where n is the number of objects. This is prohibitive with a large number of complex objects. Bounding object tests are much more manageable, typically being $O(1)$ per test.

To implement bounding spheres, I'll create a structure called `bSphere3`. It can be constructed from a location and a list of points (the location of the object, the object's points) or from an explicit location and radius check. Checking if two spheres intersect is a matter of calling `bSphere3::Intersect()` with both spheres. It returns true if they intersect each other. This is only a baby step that can be taken toward good physics, mind you, but baby steps beat doing nothing!

```
struct bSphere3
{
    float m_radius;
    point3 m_loc;

    bSphere3(){}
}
```

```

bSphere3( float radius, point3 loc ) :
    m_radius( radius ), m_loc( loc )
{
}

bSphere3( point3 loc, int nVerts, point3 *pList )
{
    m_loc = loc;
    m_radius = 0.f;
    float currRad;
    for( int i=0; i< nVerts; i++ )
    {
        currRad = pList[i].Mag();
        if( currRad > m_radius )
        {
            m_radius = currRad;
        }
    }
}

template< class iter >
bSphere3( point3 loc, iter& begin, iter& end )
{
    iter i = begin;
    m_loc = loc;
    m_radius = 0.f;
    float currRad;
    while( i != end )
    {
        currRad = (*i).Mag();
        if( currRad > m_radius )
        {
            m_radius = currRad;
        }
        i++;
    }
}

static bool Intersect( bSphere3& a, bSphere3& b )
{
    // avoid a square root by squaring both sides of the equation
    float magSqrd =
        (a.m_radius + b.m_radius) *
        (a.m_radius + b.m_radius);
    if( (b.m_loc - a.m_loc).MagSquared() > magSqrd )
    {
        return false;
    }
    return true;
}
};

```


Some additional operators are defined in `bSphere3.h`, and plane-sphere classification code is in `plane3.h` as well. See the downloadable files for more detail.

Lighting

Lighting your scenes is a prerequisite if you want them to look realistic. Lighting is a fairly slow and complex system, especially when modeling light correctly (this doesn't happen too often). Later in the book I'll discuss some advanced lighting schemes, including radiosity. I'll discuss two points in this section: how to acquire the amount of light hitting a point in 3D, and how to shade a triangle with those three points. We'll look at much more advanced lighting using HLSL later also.

Representing Color

Before you can go about giving color to anything in a scene, you need to know how to represent color! Usually you use the same red, green, and blue channels discussed in Chapter 2, but for lighting purposes there will also be a fourth component called *alpha*. The alpha component stores transparency information about a texture. It's discussed more in detail in Chapter 9, but for right now let's plan ahead. There will be two structures to ease the color duties: `color3` and `color4`. They both use floating-point values for their components; `color3` has red, green, and blue, while `color4` has the additional fourth component of alpha.

Colors aren't like points—they have a fixed range. Each component can be anything from 0.0 to 1.0 (zero contribution of the channel or complete contribution). If performing operations on colors, such as adding them together, the components may rise above 1.0 or below 0.0. Before trying to use a color, for example feeding it to `Direct3D`, it needs to be saturated. That is what the `Sat()` function does. The conversions to unsigned longs will be used in Chapter 7, when the colors start to get plugged into `Direct3D`.

The code for `color4` follows. I've left out a few routine bits to keep the listing focused.

```
struct color4
{
    union {
        struct
        {
            float r, g, b, a; // Red, Green, Blue, and Alpha color data
        };
        float c[4];
    };

    color4(){}
}
```

```

color4( float inR, float inG, float inB, float inA ) :
    r( inR ), g( inG ), b( inB ), a( inA )
{
}

color4( const color3& in, float alpha = 1.f )
{
    r = in.r;
    g = in.g;
    b = in.b;
    a = alpha;
}

color4( unsigned long color )
{
    b = (float)(color&255) / 255.f;
    color >>= 8;
    g = (float)(color&255) / 255.f;
    color >>= 8;
    r = (float)(color&255) / 255.f;
    color >>= 8;
    a = (float)(color&255) / 255.f;
}

void Assign( float inR, float inG, float inB, float inA )
{
    r = inR;
    g = inG;
    b = inB;
    a = inA;
}

unsigned long MakeDWord()
{
    unsigned long iA = (int)(a * 255.f) << 24;
    unsigned long iR = (int)(r * 255.f) << 16;
    unsigned long iG = (int)(g * 255.f) << 8;
    unsigned long iB = (int)(b * 255.f);
    return iA | iR | iG | iB;
}

unsigned long MakeDWordSafe()
{
    color4 temp = *this;
    temp.Sat();
    return temp.MakeDWord();
}

// if any of the values are >1, cap them.
void Sat()
{

```

```

        if( r > 1 )
            r = 1.f;
        if( g > 1 )
            g = 1.f;
        if( b > 1 )
            b = 1.f;
        if( a > 1 )
            a = 1.f;
        if( r < 0.f )
            r = 0.f;
        if( g < 0.f )
            g = 0.f;
        if( b < 0.f )
            b = 0.f;
        if( a < 0.f )
            a = 0.f;
    }

    color4& operator += ( const color4& in );
    color4& operator -= ( const color4& in );
    color4& operator *= ( const color4& in );
    color4& operator /= ( const color4& in );
    color4& operator *= ( const float& in );
    color4& operator /= ( const float& in );

    // some basic colors.
    static const color4 Black;
    static const color4 Gray;
    static const color4 White;
    static const color4 Red;
    static const color4 Green;
    static const color4 Blue;
    static const color4 Magenta;
    static const color4 Cyan;
    static const color4 Yellow;

};

```

Lighting Models

Lighting an object correctly is an extremely difficult process. Even today, it's still an area of research in academia. There are applications on the market that cost tens of thousands of dollars to perform renderings of scenes that have extremely accurate lighting. These renderings can take inordinate amounts of time to compute, sometimes on the order of several hours or even days for extremely complex images. Think of some of the computer-generated imagery in movies like *Final Fantasy*, *Shrek 3*, and *Ice Age 2*.

In order to avoid those issues, Direct3D and OpenGL graphics programmers use approximations of correct lighting models to get fast but good looking lighting models. While the images invariably end up looking

computer generated, they can be done in real time. True photo realism needs to have incredibly accurate lighting, as human eyes are very sensitive to lighting in a scene. All the kinds of light are cubbyholed into four essential types:

- **Ambient light**—Ambient light can be thought of as the average light in a scene. It is light that is equally transmitted to all points on all surfaces the same amount. Ambient lighting is a horrible hack—an attempt to impersonate the diffuse reflection that is better approximated by radiosity (covered in Chapter 8), but it works well enough for many applications. The difference between ambient light and ambient reflection is that ambient reflection is how much a surface reflects ambient light.
- **Diffuse light**—Diffuse light is light that hits a surface and reflects off equally in all directions. Surfaces that only reflect diffuse light appear lit the same amount, no matter how the camera views it. If modeling chalk or velvet, for example, only diffuse light would be reflected.
- **Specular light**—Specular light is light that only reflects off a surface in a particular direction. This causes a shiny spot on the surface, which is called a *specular highlight*. The highlight is dependent on both the location of the light and the location of the viewer. For example, imagine picking up an apple. The shiny spot on the apple is a good example of a specular highlight. As you move your head, the highlight moves around the surface (which is an indication that it's dependent on the viewing angle).
- **Emissive light**—Emissive light is energy that actually comes off of a surface. A light bulb, for example, looks very bright because it has emissive light. Emissive light does not contribute to other objects in the scene. It is not a light itself; it just modifies the appearance of the surface.

Ambient and diffuse lights have easier equations, so I'll give those first. If the model doesn't reflect specular light at all, you can use the following equation to light each vertex of the object. This is the same diffuse and ambient lighting equation that Direct3D uses (given in the Microsoft DirectX 10.0 SDK documentation). The equation sums all of the lights in the scene.

$$D_v = I_a S_a + S_e + \sum_i A_i (R_{di} S_d L_{di} + S_a L_{ai})$$

Table 4.1: Terms in the ambient/diffuse/emissive lighting equation for a surface

D_v	Final color for the surface.
I_a	Ambient light for the entire scene.
S_a	Ambient color for the surface.

S_e	Emitted color of the surface.
A_i	Attenuation for light i . This value depends on the kind of light you have, but essentially means how much of the total energy from the light hits an object.
R_{di}	Diffuse reflection factor for light i . This is usually the inverse of the dot product between the vertex normal and the direction in which the light travels. That way, normals that are directly facing the light receive more than normals that are turned away from it (of course, if the reflectance factor is less than zero, no diffuse light hits the object). Figure 4.25 shows the calculation visually.
S_d	Diffuse color for the surface.
L_{di}	Diffuse light emitted by light i .
L_{ai}	Ambient light emitted by light i .

The surfaces in the previous equation will end up being vertices of the 3D models once D3D is up and running. The surface reflectance components are usually defined with material structures.




Figure 4.25: Diffuse reflection factor

Specular Reflection

Specular reflections are more complex than ambient, emissive, or diffuse reflections, requiring more computation to use. Many old applications don't use specular reflections because of the overhead involved, or they'll do something like approximate them with an environment map. However, as accelerators are getting faster (especially since DirectX 10 accelerators force you to perform lighting in hardware), specular lighting is increasingly being used to add more realism to scenes.

To find the amount of specular color to attribute to a given vector with a given light, you use the adjacent equations (taken from the Microsoft DirectX 10.0 SDK documentation):

The meanings of the variables are given in Table 4.2.

$$\mathbf{v} = \overline{\mathbf{p}_c - \mathbf{p}_v}$$

$$\mathbf{h} = \overline{\mathbf{v} - \mathbf{l}_d}$$

$$R_s = (\mathbf{n} \cdot \mathbf{h})^p$$

$$S_s = C_s A R_s L_s$$

Table 4.2: Meanings of the specular reflection variables

p_c	Location of the camera.
p_v	Location of the surface.
\mathbf{l}_d	Direction of the light.
\mathbf{h}	The “halfway” vector. Think of this as the vector bisecting the angle made by the light direction and the viewer direction. The closer this is to the normal, the brighter the surface should be. The normal-halfway angle relation is handled by the dot product.
\mathbf{n}	The normal of the surface.
R_s	Specular reflectance. This is, in essence, the intensity of the specular reflection. When the point you’re computing lies directly on a highlight, it will be 1.0; when it isn’t in a highlight at all, it’ll be 0.
p	The “power” of the surface. The higher this number, the sharper the specular highlight. A value of 1 doesn’t look much different from diffuse lighting, but using a value of 15 or 20 gives a nice sharp highlight.
S_s	The color being computed (this is what you want).
C_s	Specular color of the surface. That is, if white specular light were hitting the surface, this is the specular color you would see.
A	Attenuation of the light (how much of the total energy leaving the light actually hits the surface).
L_s	Specular color of the light.

Note that this only solves for one light; you need to solve the same equation for each light, summing up the results as you go.

Light Types

Now that you have a way to find the light hitting a surface, you’re going to need some lights! There are typically three types of lights I am going to discuss, although we’ll see more advanced lighting models when we look at HLSL shaders.

Parallel Lights (or Directional Lights)

Parallel lights cheat a little bit. They represent light that comes from an infinitely far away light source. Because of this, all of the light rays that reach the object are parallel (hence the name). The standard use of a parallel light is to simulate the sun. While it’s not infinitely far away, 93 million miles is good enough!




Figure 4.26: Parallel light sources

The great thing about parallel lights is that a lot of the ugly math goes away. The attenuation factor is always 1 (for point/spotlights, it generally involves divisions if not square roots). The incoming light vector for calculation of the diffuse reflection factor is the same for all considered points, whereas point lights and spotlights involve vector subtractions and a normalization per vertex.

Typically, lighting is the kind of effect that is sacrificed for processing speed. Parallel light sources are the easiest and therefore fastest to process. If you can't afford to do the nicer point lights or spotlights, falling back to parallel lights can keep your frame rates at reasonable levels.

Point Lights

One step better than directional lights are point lights. They represent infinitesimally small points that emit light. Light scatters out equally in all directions. Depending on how much effort you're willing to expend on the light, you can have the intensity falloff based on the inverse squared distance from the light, which is how real lights work.

$$\text{attenuation_factor} = \frac{k}{|\text{surface_location} - \text{light_location}|^2}$$

The light direction is different for each surface location (otherwise the point light would look just like a directional light). The equation for it is:

$$\text{light_direction} = \frac{\text{surface_location} - \text{light_location}}{|\text{surface_location} - \text{light_location}|}$$




Figure 4.27:
Point light source

Spotlights

Spotlights are the most expensive type of light. They model a spotlight not unlike the type you would see in a theatrical production. They are point lights, but light only leaves the point in a particular direction, spreading out based on the aperture of the light.

Spotlights have two angles associated with them. One is the internal cone whose angle is generally referred to as theta (θ). Points within the internal cone receive all of the light of the spotlight; the attenuation is the same as it would be if point lights were used. There is also an angle that defines the outer cone; the angle is referred to as phi (ϕ). Points outside the outer cone receive no light. Points outside the inner cone but inside the outer cone receive light, usually a linear falloff based on how close the point is to the inner cone.




Figure 4.28: Spotlight source

If you think all of this sounds mathematically expensive, you're right. Spotlights can slow down your application a great deal. Then again, they do provide an incredible amount of atmosphere when used correctly, so you will have to figure out a line between performance and aesthetics.

Shading Models

Once you've determined and set up the lighting information, you need to know how to draw the triangles with the supplied information. With DirectX 10's new HLSL (high level shader language) you have unlimited ways to render lighting; however, we'll discuss three common methods: Lambert, Gouraud, and Phong. Figure 4.29 shows a polygon mesh of a sphere, which I'll use to explain the shading models.




Figure 4.29: Wireframe mesh

Lambert

Triangles that use Lambertian shading are painted with one solid color instead of gradients. Typically, each triangle is lit using that triangle's normal. The resulting object looks very angular and sharp. Lambertian shading was used mostly back when computers weren't fast enough to do modern shading in real time. To light a triangle, you compute the lighting equations using the triangle's normal and any of the three vertices of the triangle.



Figure 4.30: Lambert shaded polygon

Gouraud

Gouraud (pronounced *garrow*) shading used to be the de facto shading standard in accelerated 3D hardware, although more advanced models are now available. Instead of specifying one color to use for the entire triangle, each vertex has its own separate color. The color values are linearly interpolated across the triangle, creating a smooth transition between the vertex color values. To calculate the lighting for a vertex, you use the position of the vertex and a vertex normal.

Of course, it's a little hard to correctly define a normal for a vertex. What people do instead is average the normals of all the polygons that share a certain vertex, using that as the vertex normal. When the object is drawn, the lighting color is found for each vertex (rather than each polygon), and then the colors



Figure 4.31: Gouraud shaded polygon

are linearly interpolated across the object. This creates a smoother look, like the one in Figure 4.31.

One problem with Gouraud shading is that the triangles' intensities can never be greater than the intensities at the edges. So if there is a spotlight shining directly into the center of a large triangle, Gouraud shading will interpolate the intensities at the three dark corners, resulting in an incorrectly dark triangle.



Note: The internal highlighting problem usually isn't that bad. If there are enough triangles in the model, the interpolation done by Gouraud shading is usually good enough. If you really want internal highlights but only have Gouraud shading, you can subdivide the triangle into smaller pieces.

Phong

Phong shading is the most realistic shading model I'm going to talk about, and also the most computationally expensive. It tries to solve several problems that arise when you use Gouraud shading. Later in the book we'll use pixel shaders to implement per-pixel lighting.

First of all, Gouraud shading uses a linear gradient. Many objects in real life have sharp highlights, such as the shiny spot on an apple. This is difficult to handle with pure Gouraud shading. The way Phong does this is by interpolating the normal across the triangle face, not the color value, and the lighting equation is solved individually for each pixel.

You'll see how to program your own per-pixel rendering engine using shaders later in the book.




Figure 4.32: Phong shaded polygon

BSP Trees

If all you want to do is just draw lists of polygons and be done with it, then you now have enough knowledge at your disposal to do that. However, there is a lot more to 3D game programming that you must concern yourself with. Hard problems abound, and finding an elegant way to solve the problems is half the challenge of graphics programming (actually *implementing* the solution is the other half).

A lot of the hard graphics problems, such as precise collision detection or ray-object intersection, boil down to a question of spatial relationship. You need to know where objects (defined with a boundary representation of polygons) exist in relation to the other objects around them.

You can, of course, find this explicitly if you'd like, but this leads to a lot of complex and slow algorithms. For example, say you're trying to see if

a ray going through space is hitting any of a list of polygons. The slow way to do it would be to explicitly test each and every polygon against the ray. Polygon-ray intersection is not a trivial operation, so if there are a few thousand polygons, the algorithm can grind to a halt.

A spatial relationship of polygons can help a lot. If you were able to say, “The ray didn’t hit this polygon, but the entire ray is completely in front of the plane the polygon lies in,” then you wouldn’t need to test anything that sat behind the first polygon. BSP trees, as you shall soon see, are one of the most useful ways to partition space.



Note: I implemented a ray-tracer a while ago using two algorithms. One was a brute-force, test-every-polygon-against-every-ray nightmare; the other used BSP trees. The first algorithm took about 90 minutes to render a single frame with about 15K triangles in it. With BSP trees, the rendering time went down to about 45 seconds. Saying BSP trees make a big difference is a major understatement.

It all started when Henry Fuchs and Zvi Kedem, both professors at the University of Texas at Dallas, found a bright young recent grad working at Texas Instruments named Bruce Naylor. They talked him into becoming a graduate student in their graphics department, and he started doing work on computational geometry. Fuchs was a sort of bright, energetic type, and Kedem contained that spark that few other theoretical computer scientists have: He was able to take theory and apply it to practical problems. Out of this triumvirate came two SIGGRAPH papers, and Naylor’s Ph.D. thesis, which gave birth to BSP trees, which were subsequently made famous by ID Software in *Doom* and *Quake* in the ’90s.

BSP Tree Theory

BSP trees are a specialization of binary trees, one of the most basic constructs of computer science. A *BSP tree* represents a region of space (the tree can have any number of nodes, including just one). The tree is made up of *nodes* (having exactly two children) and *leaves* (having exactly zero children). A node represents a partitioning of the space that the tree it is a part of represents. The partitioning creates two new spaces, one in front of the node and one in back of the node.

In 2D applications, such as *Doom* (which used BSP trees to represent the worlds the fearless space marine navigated), the top-level tree represents the entire 2D world. The root node, which contains in it a line equation, defines a partitioning of the world into two pieces, one in front of the line and one in back of it. Each of these pieces is represented by a subtree, itself a BSP tree. The node also contains a line segment that is part of the line equation used in the partitioning. The line segment, with other information like the height and texture ID, becomes a wall in the world. Subtrees are leaves if and only if the space that they represent has no other walls in it. If it does, the wall is used to partition the space yet again.

In 3D applications, things are pretty much the same. The space is partitioned with 3D planes, which contain polygons within them. The plane at each node slices its region into two hunks, one that is further subdivided by its front child node, and the other further subdivided by its back child node.

The recursion downward stops when a space cannot be partitioned any further. For now, this happens when there are no polygons inside of it. At this point, a leaf is created, and it represents a uniform, convex region of space.

There are two primary ways to construct BSP trees. In the first method (called *node-based BSP trees*), nodes contain both polygons and the planes used to partition. Leaves are empty. In the other method (called *leaf-based* or *leafy BSP trees*), nodes only contain planes. Leaves contain all of the polygons that form the boundary of that convex space. I'm only going to talk about node-based BSP trees, but leaf-based BSPs are useful, for example in computing the potentially visible set (PVS) of a scene.

BSP trees are most useful when the set of polygons used to construct the tree represents the boundary representation of an object. The object has a conceptual inside made of solid matter, an outside of empty space surrounding it, and polygons that meet in the middle. Luckily this is how I am representing the objects anyway. When the tree is complete, each leaf represents either solid or empty space. This will prove to be extremely useful, as you shall see in a moment.

BSP Tree Construction

The algorithm to create a node-based BSP tree is simple and recursive, as shown in the pseudocode below. It is fairly time consuming, however, enough so that generally the set of polygons used to construct the BSP tree remains static. This is the case for most of the worlds that players navigate in 3D games, so games such as *Quake III: Arena* consist of a static BSP tree (representing a world) and a set of objects (health boxes, ammo boxes, players, enemies, doors, etc.) that can move around in the world.

I'll go through the tree construction process step by step.

```
struct node
    polygon poly
    plane    part_plane
    ptr      front
    ptr      back
    vector< polygon > coplanar_polygons

struct leaf
    bool solid

leaf Make_Leaf( bool solid )
    leaf out = new leaf
    out.solid = solid
    return out
```

```

polygon Get_Splitting_Polygon( vector< polygon > input_list )
    polygon out = polygon that satisfies some heuristic
    remove out from input_list
    return out

node Make_Node( vector< polygon > input_list )
    vector< polygon > front_list, back_list
    node out = new node
    chosen_polygon = Get_Splitting_Polygon( input_list )
    out.part_plane = Make_Plane_From_Polygon( chosen_polygon )
    out.poly = chosen_polygon
    for( each polygon curr in input_list )
        switch( out.part_plane.Classify_Polygon( curr ) )
            case front
                add curr to front_list
            case back
                add curr to back_list
            case coplanar
                add curr to node.coplanar_polygons
            case split
                split curr into front and back polygons
                add front to front_list
                add back to back_list
    if( front_list is empty )
        out.front = Make_Leaf( false )
    else
        out.front = Make_Node( front_list )
    if( back_list is empty )
        out.back = Make_Leaf( true )
    else
        out.back = Make_Node( back_list )
    return out

node Make_BSP_Tree( vector< polygon > input_list )
    return Make_Node( input_list )

```

Let's step through a sample 2D tree to show what is going on. The initial case will be a relatively small data set with four edges defining a closed region surrounded by empty space. Figure 4.33 shows the initial case, with the polygons on the left and a list on the right that will be processed. Each of the segments also has its plane normal visible; note that they all point out of the solid region.

To create the root node, segment A is used. Segments B, C, and D are all behind segment A, so they all go in the back list. The front list is empty, so the front child is made a leaf representing empty space corresponding to the entire subspace in front of segment A. The back list isn't empty, so it must be recursed, processing the subspace behind segment A. The result of the first partition appears in Figure 4.34.




Figure 4.33: Initial case of the BSP construction




Figure 4.34: Result after the first partitioning

Once recursion into the root node's back child is complete, a polygon to partition with must once again be selected. While real-world applications probably wouldn't choose it, to diversify the example I'm going to use segment B. Segment C is completely in front, but segment D is partially in front and partially in back. It is split into two pieces, one completely in front (which I'll call D_F) and one completely in back (called D_B). After the classification, both the front list and the back list have polygons in them, so they must be recursed with each. Figure 4.35 shows the progress up to this point.




Figure 4.35: Result after the second partition



Note: Notice the dashed line for segment B. It doesn't intrude into the space in front of segment A, because it is only partitioning the subspace *behind* A. This is a very important point you'll need to assimilate if you want to understand BSP trees fully.

I'll partition the front side of the node, the one with a list of D_F and C. I'll use D_F as the partitioning polygon. C is the only polygon to classify, and it's completely behind D_F . The front list is empty, so I create an empty space leaf. This brings the progress up to Figure 4.36.




Figure 4.36: Result after the third partition

Now there are two nodes left to process, C and D_B . I'll consolidate them into one step. They both have no other polygons to classify once the only polygon in each list is selected to be the partitioner. This creates two child leaf nodes, one in back of the polygon representing solid space (represented with a plus sign) and one in front representing empty space (represented with a minus sign). This results in the final BSP tree, which appears in Figure 4.37. I put small dashed lines from each of the leaf nodes to the subspace they represent.




Figure 4.37: The final BSP tree

One piece that's left out of the equation is how you take the list of polygons during each step and choose the polygon to use as the partitioner. There are two heuristics you can try to satisfy: Choose the polygon that causes the least amount of splits, or choose the polygon that most evenly divides the set. One problem, however, is that you can have a ton of polygons in the data set, especially at the top levels of the tree. In Foley's *Computer Graphics* it mentions that after you check about 10% of the polygons, the best candidate found thus far is so similar to the ideal one that it's not worth checking any more. This code will use a strict least-split heuristic, checking the first 10% of the polygon data (or the whole set if it's below some threshold).

BSP Tree Algorithms

Now that you've covered enough ground to create a BSP tree, hopefully a few algorithms will be at your disposal to perform operations on them. A few of the algorithms work in all polygon configurations, but generally they're suited for BSP trees that represent the boundary representation of an object.

Sorted Polygon Ordering

One of the first uses of BSP trees was to get a list of polygons sorted by distance from a given viewpoint. This was used back before hardware z-buffers, when polygons needed to be drawn in back-to-front order to be rendered correctly. It's still useful; however, z-buffer rendering goes faster

if you reject early (so rendering front-to-back can be an advantage), and alpha-blended polygons need to be rendered back-to-front to be rendered correctly.

The fundamental concept behind the algorithm is that if you have a certain plane in the scene dividing it into two pieces, and you are on one side of the plane, then nothing behind the plane can occlude anything in front of the plane. Armed with this rule, all you need to do is traverse the tree. At each node, you check to see which side the camera is on. If it's in front, then you add all the polygons behind the plane (by traversing into the back node), then all the polygons in the plane (the partitioning polygon and any coplanar polygons), and finally the polygons in front of the plane (by traversing into the front node). The opposite applies if it is in back. Leaves just return automatically.



Note: If you don't want to draw polygons facing away, you can automatically discard the node polygon if the camera point is behind the node plane. The coplanar polygons you can check, unless you keep two lists of coplanar polygons—one facing in the same direction as the partitioning polygons and one facing the opposite way.

The algorithm to do this is fairly simple and recursive, as shown in this pseudocode.

```
void node::GetSortedPolyList(
    list< polygon3 > *pList,
    point& camera )
{
    switch( node.part_plane.Classify_Point( camera ) )
    {
        case front
            back. GetSortedPolyList( pList, camera );
            add all node polygons to pList
            front. GetSortedPolyList( pList, camera );
        case back
            front. GetSortedPolyList( pList, camera );
            add all node polygons to pList
            back. GetSortedPolyList( pList, camera );
        case coplanar
            // order doesn't matter
            front. GetSortedPolyList( pList, camera );
            back. GetSortedPolyList( pList, camera );
    }
}

void leaf:: GetSortedPolyList(
    list< polygon3 > *pList,
    point& camera )
{
    return;
}
```

Testing Locality of a Point

A really great use of BSP trees is testing the locality of points. Given a point and a tree, you can tell whether or not the point is sitting in a solid leaf. This is useful for collision detection, among other things.

The algorithm to do it is amazingly simple, as shown in the following pseudocode. At each branch of the tree, you test the point against the plane. If it's in front, you drop it down the front branch; if it's in back, you drop it down the back branch. If the point is coplanar with the polygon, you can pick either one. Whenever you land in a leaf, you have found the region of space that the point is sitting in. If the leaf is tagged as being solid, then the point is sitting in solid space; otherwise it's not.

```
bool node::TestPoint(
    point& pt )
{
    switch( node.part_plane.Classify_Point( pt ) )
    {
        case front
            return front.TestPoint( pt );
        case back
            return back.TestPoint( pt );
        case coplanar
            // Let's drop down the back tree
            return back.TestPoint( pt );
    }
}

bool leaf::TestPoint(
    point& pt )
{
    if( solid )
        return true;
    return false;
}
```

Testing Line Segments

While there are many other algorithms for use with BSP trees, the last one I'll discuss lets you test a line segment against a tree. The algorithm returns true if there is a clear line of sight between both endpoints of the line, and false otherwise. Another way to think of it is to say that the algorithm returns true if and only if the line segment only sits in non-solid leaves.

Like all the other algorithms I've discussed, this is a conceptually simple and elegant algorithm. Starting at the root, you compare the line segment to the plane at a node. If the line segment is completely in front, you drop it down the front side. If it's completely in back, you drop it down the back side. If the line segment is on both sides of the plane, you divide it into two pieces (one in front of the plane and one in back) and recurse

with both of them. If any piece of segment ever lands in a solid cell, then you know there is no line of sight, and you return false.

Source code to do this appears in the following section.

BSP Tree Code

Below is the source code for the BSP class. I'll be using it later to find the form factor in the radiosity simulator. The main difference between this code and the pseudocode given above is this code uses the same node structure to represent both nodes and leaves. This made the code simpler but is an inefficient use of space (leaves only need a single word defining them as solid; here a lot more than that is used).

```

/*****
 *      Advanced 3D Game Programming with DirectX 10.0
 *      * * * * *
 *
 *      See license.txt for modification and distribution information
 *      copyright (c) 2007 by Peter Walsh, Wordware
 *****/
#ifdef _BSPTREE_H
#define _BSPTREE_H

#include <point3.h>
#include <polygon.h>
#include <plane3.h>

#include <vector>
using std::vector;

const float percentageToCheck = .1f; // 10%

/**
 * This code expects the set of polygons we're giving it to be
 * closed, forming a continuous skin. If it's not, weird things
 * may happen.
 */
class cBspTree
{
public:

    // construction/destruction
    cBspTree();
    ~cBspTree();

    // we need to handle copying
    cBspTree( const cBspTree &in );
    cBspTree& operator=( const cBspTree &in );

    // add a polygon to the tree
    void AddPolygon( const polygon<point3>& in );
    void AddPolygonList( vector< polygon<point3> >& in );

    void TraverseTree(

```

```

        vector< polygon<point3>* >* polyList,
        const point3& loc );

    bool LineOfSight( const point3& a, const point3& b );

protected:

private:

class cNode
{
    cNode      *m_pFront;          // pointer to front subtree
    cNode      *m_pBack;           // pointer to back subtree

    polygon<point3> m_poly;
    plane3      m_plane;
    bool        m_bIsLeaf;
    bool        m_bIsSolid;

    vector< polygon<point3> >    m_coplanarList;

    static int BestIndex( vector< polygon<point3> >& polyList );

public:
    cNode( bool bIsSolid );          // leaf constructor
    cNode( const polygon<point3>& in ); // node constructor
    cNode( vector< polygon<point3> >& in ); // node constructor
    ~cNode();

    // we need to handle copying
    cNode( const cNode &in );
    cNode& operator=( const cNode &in );

    void AddPolygon( const polygon<point3>& in );

    void TraverseTree(
        vector< polygon<point3>* >* polyList,
        const point3& loc );

    bool IsLeaf()
    {
        return m_bIsLeaf;
    }

    bool LineOfSight( const point3& a, const point3& b );
};

cNode *m_pHead;                // root node of the tree
};

inline cBspTree::cBspTree( const cBspTree &in )

```

```

{
    // clone the tree
    if( in.m_pHead )
        m_pHead = new cNode( *in.m_pHead );
    else
        m_pHead = NULL;
}

inline cBspTree& cBspTree::operator=( const cBspTree &in )
{
    if( &in != this )
    {
        // delete the tree if we have one already
        if( m_pHead )
            delete m_pHead;

        // clone the tree
        if( in.m_pHead )
            m_pHead = new cNode( *in.m_pHead );
        else
            m_pHead = NULL;
    }

    return *this;
}

inline cBspTree::cNode::cNode( const cNode &in )
{
    m_poly = in.m_poly;
    m_plane = in.m_plane;
    m_bIsLeaf = in.m_bIsLeaf;
    m_bIsSolid = in.m_bIsSolid;

    // clone the trees
    m_pFront = NULL;
    if( in.m_pFront )
        m_pFront = new cNode( *in.m_pFront );

    m_pBack = NULL;
    if( in.m_pBack )
        m_pBack = new cNode( *in.m_pBack );
}

inline cBspTree::cNode& cBspTree::cNode::operator=( const cNode &in )
{
    if( &in != this )
    {
        // delete the subtrees if we have them already
        if( m_pFront )
            delete m_pFront;
        if( m_pBack )
            delete m_pBack;
    }
}

```

```

        // copy all the data over
        m_poly = in.m_poly;
        m_plane = in.m_plane;
        m_bIsLeaf = in.m_bIsLeaf;
        m_bIsSolid = in.m_bIsSolid;

        // clone the trees
        m_pFront = NULL;
        if( in.m_pFront )
            m_pFront = new cNode( *in.m_pFront );

        m_pBack = NULL;
        if( in.m_pBack )
            m_pBack = new cNode( *in.m_pBack );
    }
    return *this;
}

#endif // _BSPTREE_H

/*****
 *      Advanced 3D Game Programming with DirectX 10.0
 *      * * * * *
 *
 *      See license.txt for modification and distribution information
 *      copyright (c) 2007 by Peter Walsh, Wordware
 *****/
#include <template.h>
#include <BspTree.h>

cBspTree::cBspTree()
{
}

cBspTree::~cBspTree()
{
    // destroy the tree
}

void cBspTree::AddPolygon( const polygon<point3>& in )
{
    if( !m_pHead )
    {
        // if there's no tree, make a new one
        m_pHead = new cNode( in );
    }
    else
    {
        // otherwise add it to the tree
        m_pHead->AddPolygon( in );
    }
}

```

```

void cBspTree::AddPolygonList( vector< polygon<point3> >& in )
{
    if( !m_pHead )
    {
        // if there's no tree, make a new one
        m_pHead = new cNode( in );
    }
    else
    {
        /**
         * Adding a list of polygons to
         * an existing tree is unimplemented
         * (exercise to the reader)
         */
        assert( false );
    }
}

void cBspTree::TraverseTree(
    vector<polygon<point3>*>*& polyList,
    const point3& loc )
{
    if( m_pHead )
    {
        // drop it down
        m_pHead->TraverseTree( polyList, loc );
    }
}

bool cBspTree::LineOfSight( const point3& a, const point3& b )
{
    assert( m_pHead ); // make sure there is a tree to test against

    return m_pHead->LineOfSight( a, b );
}

cBspTree::~cNode::~cNode()
{
    delete m_pFront;
    delete m_pBack;
}

cBspTree::cNode::cNode( bool bIsSolid )
: m_bIsLeaf( true )
, m_bIsSolid( bIsSolid )
, m_pFront( NULL )
, m_pBack( NULL )
{
    // all done.
}

```



```
cBspTree::cNode::cNode( const polygon<point3>& in )
: m_bIsLeaf( false )
, m_poly( in )
, m_plane( in )
, m_pFront( new cNode( false ) )
, m_pBack( new cNode( true ) )
{
    // all done.
}

cBspTree::cNode::cNode( vector< polygon<point3> >& in )
: m_bIsLeaf( false )
{
    // if the list is empty, we're bombing out.
    assert( in.size() );

    // get the best index to use as a splitting plane
    int bestIndex = BestIndex( in );

    // we could remove the index from the vector, but that's slow.
    // instead we'll just kind of ignore it during the next phase.
    // remove the best index
    polygon<point3> splitPoly = in[bestIndex];

    m_plane = plane3( splitPoly );
    m_poly = splitPoly;

    // take the rest of the polygons and divide them.
    vector< polygon<point3> > frontList, backList;

    int i;
    for( i=0; i<in.size(); i++ )
    {
        // ignore the polygon if it's the one
        // we're using as the splitting plane
        if( i == bestIndex ) continue;

        // test the polygon against this node.
        pListLoc res = m_plane.TestPoly( in[i] );

        polygon<point3> front, back; // used in PLIST_SPLIT

        switch( res )
        {
            case PLIST_FRONT:
                // drop down the front
                frontList.push_back( in[i] );
                break;
            case PLIST_BACK:
                // drop down the back
                backList.push_back( in[i] );
                break;
            case PLIST_SPLIT:
                // split the polygon, drop the halves down.
                m_plane.Split( in[i], &front, &back );
```

```

        frontList.push_back( front );
        backList.push_back( back );
        break;
    case PLIST_COPLANAR:
        // add the polygon to this node's list
        m_coplanarList.push_back( in[i] );
        break;
    }
}

// we're done processing the polygon list. Deal with them.
if( frontList.size() )
{
    m_pFront = new cNode( frontList );
}
else
{
    m_pFront = new cNode( false );
}
if( backList.size() )
{
    m_pBack = new cNode( backList );
}
else
{
    m_pBack = new cNode( true );
}
}

void cBspTree::cNode::AddPolygon( const polygon<point3>& in )
{
    if( m_bIsLeaf )
    {
        // reinitialize ourselves as a node
        *this = cNode( in );
    }
    else
    {
        // test the polygon against this node.
        plistLoc res = this->m_plane.TestPoly( in );

        polygon<point3> front, back; // used in PLIST_SPLIT
        switch( res )
        {
            case PLIST_FRONT:
                // drop down the front
                m_pFront->AddPolygon( in );
                break;
            case PLIST_BACK:
                // drop down the back
                m_pBack->AddPolygon( in );
                break;
            case PLIST_SPLIT:
                // split the polygon, drop the halves down.

```

```

        m_plane.Split( in, &front, &back );
        m_pFront->AddPolygon( front );
        m_pBack->AddPolygon( back );
        break;
    case PLIST_COPLANAR:
        // add the polygon to this node's list
        m_coplanarList.push_back( in );
        break;
    }
}

}

void cBspTree::cNode::TraverseTree( vector< polygon<point3>* >* polyList,
    const point3& loc )
{
    if( m_bIsLeaf )
    {
        // do nothing.
    }
    else
    {
        // test the loc against the current node
        pointLoc res = m_plane.TestPoint( loc );

        int i;
        switch( res )
        {
            case POINT_FRONT:
                // get back, us, front
                m_pBack->TraverseTree( polyList, loc );
                polyList->push_back( &m_poly ); // the poly at this node
                for( i=0; i<m_coplanarList.size(); i++ )
                {
                    polyList->push_back( &m_coplanarList[i] );
                }
                m_pFront->TraverseTree( polyList, loc );
                break;

            case POINT_BACK:
                // get front, us, back
                m_pFront->TraverseTree( polyList, loc );
                polyList->push_back( &m_poly ); // the poly at this node
                for( i=0; i<m_coplanarList.size(); i++ )
                {
                    polyList->push_back( &m_coplanarList[i] );
                }
                m_pBack->TraverseTree( polyList, loc );
                break;

            case POINT_COPLANAR:
                // get front, back, us
                m_pFront->TraverseTree( polyList, loc );
                m_pBack->TraverseTree( polyList, loc );
                polyList->push_back( &m_poly ); // the poly at this node
                for( i=0; i<m_coplanarList.size(); i++ )

```

```

        {
            polyList->push_back( &m_coplanarList[i] );
        }
        break;
    }
}

int cBspTree::cNode::BestIndex( vector< polygon<point3> >& polyList )
{
    /**
     * The current heuristic is blind least-split
     */
    // run through the list, searching for the best one.
    // the highest polygon we'll bother testing (10% of total)
    int maxCheck;
    maxCheck = (int)(polyList.size() * percentageToCheck);
    if( !maxCheck ) maxCheck = 1;

    int i, i2;
    int bestSplits = 100000;
    int bestIndex = -1;
    int currSplits;
    plane3 currPlane;
    for( i=0; i<maxCheck; i++ )
    {
        currSplits = 0;
        currPlane = plane3( polyList[i] );
        pListLoc res;

        for( i2=0; i2< polyList.size(); i2++ )
        {
            if( i == i2 ) continue;

            res = currPlane.TestPoly( polyList[i2] );
            if( res == PLIST_SPLIT )
                currSplits++;
        }
        if( currSplits < bestSplits )
        {
            bestSplits = currSplits;
            bestIndex = i;
        }
    }
    assert( bestIndex >= 0 );
    return bestIndex;
}

bool cBspTree::cNode::LineOfSight( const point3& a, const point3& b )
{
    if( m_bIsLeaf )
    {
        // if we land in a solid node, then there is no line of sight
    }
}

```

```
        return !m_bIsSolid;
    }

    pointLoc aLoc = m_plane.TestPoint( a );
    pointLoc bLoc = m_plane.TestPoint( b );

    point3 split;

    if( aLoc == POINT_COPLANAR && bLoc == POINT_COPLANAR )
    {
        // for sake of something better to do, be conservative
        //return false;
        return m_pFront->LineOfSight( a, b );
    }

    if( aLoc == POINT_FRONT && bLoc == POINT_BACK )
    {
        //split, then return the logical 'or' of both sides
        split = m_plane.Split( a, b );

        return m_pFront->LineOfSight( a, split )
            && m_pBack->LineOfSight( b, split );
    }

    if( aLoc == POINT_BACK && bLoc == POINT_FRONT )
    {
        // split, then return the logical 'or' of both sides
        split = m_plane.Split( a, b );

        return m_pFront->LineOfSight( b, split )
            && m_pBack->LineOfSight( a, split );
    }

    // the other == POINT_COPLANAR or POINT_FRONT
    if( aLoc == POINT_FRONT || bLoc == POINT_FRONT )
    {
        // drop down the front
        return m_pFront->LineOfSight( a, b );
    }

    else // they're both on the back side
    {
        // drop down the front
        return m_pBack->LineOfSight( a, b );
    }

    return true;
}
```

Wrapping It Up

Most of the code discussed in this chapter is available from the downloadable files in one library, called `math3d.lib`. The rest of it (most notably the lighting pipeline) won't be implemented by you; that's being left wholly to Direct3D HLSL shaders. There aren't any sample applications for this chapter because you won't be able to draw any primitives until Chapter 7. The downloadable source for this chapter contains more complete, more commented versions of the code discussed in this chapter. So feel free to take a look at that.

This page intentionally left blank.

Artificial Intelligence

From a very young age I was fascinated with computers, despite the fact that we only had two games for our home machine: a game where a donkey ran down the road avoiding cars, and an app that used the PC speaker to crudely simulate a piano.

The title of the first AI application I saw escapes me (I believe it may have been just *Animal*), but the premise was simple enough. The object of the game was for the computer to guess an animal you were thinking about. It would ask a series of yes/no questions that would narrow down the possible choices (examples would be “does your animal fly?” or “does your animal have four legs?”), and when it was sure, it would tell you what it thought your animal was. The neat thing was, if it didn’t guess your animal, it would ask you a question that differentiated the two animals, something your animal had that the other didn’t. From then on, the program would be able to guess your animal! It could learn!

This impressed my young mind to no end. After some formal training in programming, I’ve come to accept that it’s a fairly trivial program: The application keeps an internal binary tree with a question at each branch and an animal at each leaf. It descends down the tree asking the question at each branch and taking the appropriate direction. If it reaches a leaf and the animal stored there isn’t yours, it creates a new branch, adds your question, and puts your animal and the animal previously in the leaf in two new leaves.

How the program worked, however, really isn’t that important. The trick is, it *seemed* intelligent to me. Game programmers need to aim for this. While academia argues for the next 50 years over whether or not human-level intelligence is possible with computers, game developers need only be concerned with tricking humans into thinking what they’re playing against is intelligent. And luckily (for both developers and academia), humans aren’t that smart.

This is, of course, not as easy as it sounds. Video games are rife with pretty stupid computer opponents. Early first-person shooters had enemies that walked toward the player in a zigzag pattern, never walking directly toward their target, shooting intermittently. Bad guys in other games would sometimes walk into a corner looking for you, determined that they would eventually find you even though you were several rooms away. Fighting games are even worse. The AI governing computer opponents can become extremely repetitive (so that every time you jump toward the opponent, they execute the same move). I can’t promise to teach you everything you need to know to make the next *Reaper Bot*; that would be

the topic of an entire book all its own. By the end of this chapter, however, you should be able to write an AI that can at least challenge you and maybe even surprise you!

Starting Point

Most AI problems that programmers face fall into three groups. At the lowest level is the problem of physical movement—how to move the unit, how to turn, how to walk, etc. This group is sometimes called *locomotion*, or *motor skills*. Moving up one level is a higher-level view of unit movement, where the unit has to decide how to get from point A to point B, avoiding obstacles and/or other units. This group is called *steering*, or *task generation*. Finally, at the highest level, the meatiest part of AI, is the actual thinking. Any cockroach can turn in a circle and do its best to avoid basic obstacles (like a human's foot). That does not make the cockroach intelligent. The third and highest stage is where the unit decides what to do and uses its ability to move around and plan directions to carry out its wishes. This highest level is called *motivation*, or *action steering*.

Locomotion

Locomotion, depending on how you look at it, is either trivial or trivially complex. An animation-based system can handle locomotion pretty easily, move forward one unit, and use the next frame of animation in the walk cycle. Every game on the market uses something similar to this to handle AI locomotion.

However, that isn't the whole story. When you walk up stairs, you need a stair walking animation; when you descend down a hill, you naturally lean back to retain your balance. The angle you lean back is dependent on the angle of the hill. The amount you dig your feet into the ice is dependent on how slippery the ice is and how sure your footing needs to be before you proceed. Animation systems robust enough to handle cases like this require a lot of special casing and scripting; most animation systems use the same walk animation for all cases.

A branch of control theory attempts to solve this with *physical controllers*. You can actually teach an AI creature how to stand and tell it how to retain its balance, how to walk around, jump, anything. This gives the AI incredible control, as the algorithms can handle any realistic terrain and any conditions. Many people agree that the future of locomotion in games is physical controllers.

However, physical controllers aren't easy. At all. For these purposes, it's total overkill. As Moore's law inevitably marches forward, there will eventually be enough processing power to devote the cycles to letting each creature figure out how to run toward its target. When this happens, games will be one huge step closer to looking like real life.

Steering—Basic Algorithms

Even games with little or no AI at all need to implement some form of steering. Steering allows entities to navigate around the world they exist in. Without it, enemies would just sit there with a rather blank look in their eyes. There are a slew of extremely basic steering algorithms that I'll touch upon, and a couple of slightly more advanced ones that I'll dig into a little deeper.

Chasing

The first AI that most people implement is the ruthless, unthinking, unrelenting Terminator AI. The creature never thinks about rest, about getting health or ammo, about attacking other targets, or even walking around obstacles: It just picks a target and moves toward it each frame relentlessly. The code to handle this sort of AI is trivial. Each frame, the creature takes the position of its target, generates a vector to it, and moves along the vector a fixed amount (the amount is the speed of the creature), as shown in the following pseudocode:

```
void cCreature::Chase( cObject *target )
{
    // Find the locations of the creature and its target.
    point3 creatureLoc = m_loc;
    point3 targetLoc = target->GetLoc();

    // Generate a direction vector between the objects
    point3 direction = targetLoc - creatureLoc;

    // Normalize the direction (make it unit-length)
    direction.Normalize();

    // move our object along the direction vector some fixed amount
    m_loc += direction * m_speed;
}
```

Evading

The inverse of a chasing algorithm is what I could probably get away with calling rabbit AI, but I'll leave it at evading. Each frame, you move directly away from a target as fast as you can (although in this case the target would most likely be a predator).

```
void cCreature::Evade( cObject* target )
{
    // Find the locations of the creature and its target.
    point3 creatureLoc = m_loc;
    point3 targetLoc = target->GetLoc();

    // Generate a direction vector between the objects
```

```

    point3 direction = targetLoc - creatureLoc;

    // Normalize the direction (make it unit-length)
    direction.Normalize();

    // move our object away from the target by multiplying
    // by a negative speed
    m_loc += direction * -m_speed;
}

```

Pattern-based AI

Another fairly simple AI algorithm I'm going to discuss is pattern-based AI. If you have ever played the classic *Space Invaders*, you're familiar with this AI algorithm. Aliens take turns dive-bombing the player, with every type of alien attacking in one uniform way. The way it attacks is called a *pattern*. At each point in time, each creature in the simulation is following some sort of pattern.

The motivation engine (in this case, usually a random number generator) chooses a pattern to perform from a fixed set of patterns. Each pattern encodes a series of movements to be carried out each frame. Following the *Space Invaders* theme, examples of pattern-based AI would be moving back and forth, diving, and diving while shooting. Anyone who has played the game has noticed that each unit type dives toward the player the same way, oblivious to the player's location. When the baddies aren't diving, they all slide back and forth in the same exact way. They're all following the same set of patterns.

The algorithm to run a pattern-based AI creature is straightforward. I'll define a pattern to be an array of points that define direction vectors for each frame of the pattern. Since the arrays can be of any length, I also keep track of the length of the array. During the AI simulation step the creature moves itself by the amount in the current index of the array. When it reaches the end of an array, it randomly selects a new pattern. Let's examine some pseudocode to handle pattern-based AI.

```

struct sPattern
{
    int patternLength;
    point3 *pattern; // array of length patternLength
};

sPattern g_patterns[ NUM_PATTERNS ];

void cCreature::FollowPattern()
{
    // pattFrame is the current frame of the pattern
    // pattNum is the current pattern we're using.

    if( pattFrame >= g_patterns[ pattNum ].patternLength )
    {

```

```

    // new pattern
    pattNum = rand()%NUM_PATTERNS;
    pattFrame = 0;
}

// follow our current pattern.
m_loc += g_patterns[pattNum].pattern[pattFrame++];
}

```

Pattern-based AI can be specialized into what is known as *scripted* AI. When a certain state is reached, the motivation engine can run a certain scripted steering pattern. For example, an important state would be your player entering a room with a bad guy in it. This could cause the creature to follow a specialized animation just for that one game situation. The bad guy could run and trip an alarm, dive out of bed toward his bludgeoning weapon of choice, or anything else you can dream up.

Steering—Advanced Algorithms

In case you haven't figured it out yet, the basic steering algorithms provided so far are *terrible*! They merely provide the creature with an ability to move. Moving in a pattern, moving directly toward an opponent, or fleeing directly away from it is only slightly more believable than picking a random direction to move every frame! No one will ever mistake your basic cretins for intelligent creatures. Real creatures don't follow patterns. Moving directly toward or directly away from you makes them an easy target. A prime example of how this would fail is illustrated in Figure 5.1.




Figure 5.1: Chasing directly toward a target is not always the smartest option.

How can you make the creatures appear more intelligent? It would be cool to give them the ability to navigate through an environment, avoiding

obstacles. If the top-level motivation AI decides that it needs health or ammo, or needs to head to a certain location, the task of getting there intelligently should be handed off to the steering engine. Two general algorithms for achieving this are what I'll discuss next.

Potential Functions

Luckily for game programmers, a lot of the hard problems in steering and autonomous navigation for AI creatures have already been solved by the robotics community. Getting an autonomous unit, like a Mars rover, to plan a path and execute it in a foreign environment is a problem that countless researchers and academics have spent years trying to solve. One of the ways they have come up with to let a robot (or a creature in the game) wander through an unknown scene is to use what are called *potential functions*.

Imagine a scene filled with obstacles, say tree trunks in a forest. There is a path from the start to the goal and no tricky situations (like a U-shaped wall of trees, which ends up being a real problem as you'll see in a moment). The unit should be able to reach the goal; all it needs to do is not run into any trees. Anytime it gets close to a tree, logically, it should adjust its direction vector so it moves away from the tree. The amount it wants to move away from the tree should be a function based on the distance from the obstacle; that is, if it is right next to the obstacle, it will want to avoid it more than if it is half a mile away from it. Figure 5.2 shows an example of this. It will obviously want to try to avoid obstacle 1 more than it tries to avoid obstacle 2, since obstacle 2 is so much farther away.




Figure 5.2: Some obstacles should be avoided more than others.

This statement can be turned into an equation. Initially the direction is the normalized vector leading to the goal (or the goal location minus the current location). Then, for each obstacle, you find the normalized vector that moves directly away from it. Then multiply it by a constant, and divide it by the squared distance from the obstacle. When finished, you have a vector that the object should use as a direction vector (it should be normalized, however).

$$\text{direction} = \frac{\text{goal}_{\text{loc}} - \text{curr}_{\text{loc}}}{\|\text{goal}_{\text{loc}} - \text{curr}_{\text{loc}}\|} + \sum_n \left(\frac{\text{curr}_{\text{loc}} - \text{obst}^n_{\text{loc}}}{\|\text{curr}_{\text{loc}} - \text{obst}^n_{\text{loc}}\|^2} \times \frac{k}{\|\text{curr}_{\text{loc}} - \text{obst}^n_{\text{loc}}\|} \right)$$

Generally the obstacles (and the object navigating) have some radius associated with them, so the last term in the equation can be adjusted to use the distance between the spheres instead of the distance between the spheres' centers.

The Good

Potential functions work great in areas sparsely populated with physical obstacles, particularly outdoor scenes. You can reach a goal on the other side of a map avoiding trees, rocks, and houses beautifully and with little computational effort. A quick and easy optimization is to only consider the objects that are within some reasonable distance from you, say 50 feet; that way you don't need to test against every object in the world (which could have hundreds or thousands of objects in it).

One great advantage of potential functions is in situations where the obstacles themselves can move around. You can use potential functions to model the movement of a squad of units across a map, for example. They can avoid obstacles in the scene (using potential functions or more complex path finding algorithms like A*) and then avoid each other using potential functions.

The Bad

Potential functions are not a silver bullet for all problems, however. As intelligent as units can look being autonomously navigated with this method, they're still incredibly stupid. Mathematically, think of the potential functions as descending into a valley toward a goal, with the obstacles appearing as hills that roll past. If these hill obstacles are grouped in such a way as to create a secondary valley, there is no way of getting out of it, since the only way to move is downward. Similarly, if two obstacles are too close to each other, you won't be able to pass between them, and if obstacles are organized to form a barrier, like a wall or specifically a U-shaped obstacle, you're totally screwed. Figure 5.3 gives an example of this.




Figure 5.3: Potential functions alone can-not get to the goal in this configuration.

Application: potentialFunc

To help explain the ideas described above, I wrote a small test app to show off potential functions. You can use the z, x, and c keys to make large, medium, and small obstacles under the mouse, respectively. The Space key releases a creature under the mouse that heads to the goal, which appears as a green circle. Since the GDI is so slow, I decided against clearing the window every frame, so the creatures leave trails as they move around. For most cases, the creatures (more than one can be created at once) reach their goal well, as evidenced in the following figure:




Figure 5.4: Potential functions doing their job

However, they don't work all the time, as evidenced by this figure:




Figure 5.5:
Potential func-
tions failing
spectacularly

You'll notice that inside the code I sum the potential forces for the objects and move a bit ten times each frame. If I only moved once, I would need to move a fairly large amount to have the speed of the creature be anything interesting. However, when the deltas are that large, the result is some ugly numerical stability problems (characterized by a jittering when the creature gets very close to an obstacle). Sampling multiple times each frame fixes the problem.

The GDI isn't useful for writing 3D games, so I'm covering it very minimally in this book. However, for doing something like a potential function application it turns out to be quite useful. While I'm providing no explanations of how GDI works, armed with this code and the Win32 SDK documentation, figuring it out on your own won't be terribly difficult.

The code uses two main classes, `cCreature` (an entity that tries to reach the goal) and `cObstacle` (something obstructing the path to the goal). The code keeps vectors of all of the creatures and objects in the scene. Each frame, each member of the creature vector gets processed, during which it examines the list of obstacles. A nice extension to the program would be for creatures to also avoid other creatures; currently they blindly run all over each other.

The code for this application is mostly GUI and drawing code, and the important function is `cCreature::Process()`. It is called every frame, and it performs the potential function equation given earlier to find the new location. After each creature gets processed, the entire scene gets drawn. Rather than list all of the code for the program, I'll just give this one function.


```

bool cCreature::Process()
{
    point3 goalVec = g_goalLoc - m_loc;

    if( goalVec.Length() < g_creatureSpeed )
        return false;    // we reached the goal, destroy ourselves

    point3 dirVec = goalVec / goalVec.Length();

    float k = .1f;

    // for each obstacle
    for( int i=0; i<g_obstacles.size(); i++ )
    {
        // find the vector between the creature and the obstacle
        point3 obstacleVec = m_loc - g_obstacles[i].m_loc;

        // compute the length, subtracting object radii to find
        // the distance between the spheres,
        // not the sphere centers
        float dist = obstacleVec.Length() -
            g_obstacles[i].m_rad - m_rad;

        // this is the vector pointing away from the obstacle
        obstacleVec.Normalize();

        dirVec += obstacleVec * ( k / (dist * dist) );
    }
    dirVec.Normalize();

    m_loc += g_creatureSpeed * dirVec;
    return true;    // we should continue processing
}

```

Path Following

Path following is the process of making an agent look intelligent by having it proceed to its destination using a logical path. The term “path following” is really only half of the picture. Following a path once you’re given it is fairly easy. The tricky part is generating a logical path to a target. This is called *path planning*.

Before it is possible to create a logical path, it must be defined. For example, if a creature’s desired destination (handed to it from the motivation code) is on the other side of a steep ravine, a logical path would probably be to walk to the nearest bridge, cross the ravine, then walk to the target. If there were a steep mountain separating it from its target, the most logical path would be to walk around the mountain, instead of whipping out climbing gear.

A slightly more precise definition of a logical path is *the path of least resistance*. Resistance can be defined as one of a million possible things, from a lava pit to a strong enemy to a brick wall. In an example of a world

with no environmental hazards, enemies, cliffs, or whatnot, the path of least resistance is the shortest one, as shown in Figure 5.6.




Figure 5.6:
Choosing paths
based on length
alone

Other worlds are not so constant. Resistance factors can be worked into algorithms to account for something like a room that has the chance of being filled with lava (like the main area of DM2 in *Quake*). Even if traveling through the lava room is the shortest of all possible paths using sheer distance, the most logical path is to avoid the lava room if it made sense. Luckily, once the path finding algorithm is set up, modifying it to support other kinds of cost besides distance is a fairly trivial task. If other factors are taken into account, the chosen path may be different. See Figure 5.7.




Figure 5.7:
Choosing paths
based on other
criterion

Groundwork

While there are algorithms for path planning in just about every sort of environment, I'm going to focus on path planning in networked convex polyhedral cells. Path planning for something like a 2D map (like those seen in *Starcraft*) is better planned with algorithms like A*.

A convex cell will be defined as a region of passable space that a creature can wander through, such as a room or hallway. Convex polyhedrons follow the same rules for convexity as the polygons. For a polygon (2D) or a polyhedron (3D) to be convex, any ray that is traced between any two points in the cell cannot leave the cell. Intuitively, the cell cannot have any dents or depressions in it; that is, no part of the cell sticks inward. Concavity is a very important trait for what is being done here. At any point inside the polyhedron, exiting the polyhedron at any location is possible and there is no need to worry about bumping into walls. Terminator logic can be used from before until the edge of the polyhedron is reached.

The polyhedrons, when all laid out, become the world. They do not intersect each other. They meet up such that there is exactly one convex polygon joining any two cells. This invisible boundary polygon is a special type of polygon called a *portal*. Portals are the doorways connecting rooms and are passable regions themselves. If you enter and exit cells from portals, and you know a cell is convex, then you also know that any ray traveling between two portals will not be obstructed by the walls of the cell (although it may run against a wall). Until objects are introduced into the world, if the paths are followed exactly, there is no need to perform collision tests.




Figure 5.8:
Cells and the
portals connect-
ing them

I'll touch upon this spatial definition later in the book when I discuss hidden surface removal algorithms; portal rendering uses this same paradigm to accelerate hidden surface removal tasks.

The big question that remains is how do you move around this map? To accomplish finding the shortest path between two arbitrary locations on the map (the location of the creature and a location the user chooses), I'm going to build a directed, weighted graph and use Dijkstra's algorithm to find the shortest edge traversal of the graph.

If that last sentence didn't make a whole lot of sense, don't worry, just keep reading!

Graph Theory

The need to find the shortest path in graphs shows up everywhere in computer programming. Graphs can be used to solve a large variety of problems, from finding a good path to send packets through on a network of computers, to planning airline trips, to generating door-to-door directions using map software.

A *weighted, directed graph* is a set of *nodes* connected to each other by a set of *edges*. Nodes contain locations, states you would like to reach, machines, anything of interest. Edges are bridges from one node to another. (The two nodes being connected can be the same node, although for these purposes that isn't terribly useful.) Each edge has a value that describes the cost to travel across the edge, and is unidirectional. To travel from one node to another and back, two edges are needed: one to take you from the first node to the second and one that goes from the second node to the first.

Dijkstra's algorithm allows you to take a graph with positive weights on each edge and a starting location and find the shortest path to all of the other nodes (if they are reachable at all). In this algorithm each node has two pieces of data associated with it: a "parent" node and a "best cost" value. Initially, all of the parent values for all of the nodes are set to invalid values, and the best cost values are set to infinity. The start node's best cost is set to zero, and all of the nodes are put into a priority queue that always removes the element with the lowest cost. Figure 5.9 shows the initial case.




Figure 5.9:
The initial case
for shortest path
computation



Note: Notice that the example graphs I’m using seem to have bidirectional edges (edges with arrows on both sides). These are just meant as shorthand for two unidirectional edges with the same cost in both directions. In the following images, gray circles are visited nodes and dashed lines are parent links.

Iteratively remove the node with the lowest best cost from the queue. Then look at each of its edges. If the current best cost for the destination node for any of the edges is greater than the current node’s cost plus the edges’ cost, then there is a better path to the destination node. Then update the cost of the destination node and the parent node information, pointing them to the current node. The pseudocode is shown below.

```

struct node
    vector< edge > edges
    node parent
    real cost

struct edge
    node dest
    real cost

while( priority_queue is not empty )
    node curr = priority_queue.pop
    for( all edges leaving curr )
        if( edge.dest.cost > curr.cost + edge.cost )
            edge.dest.cost = curr.cost + edge.cost
            edge.dest.parent = curr
    
```

Let me step through the algorithm so I can show you what happens. In the first iteration, I take the starting node off the priority queue (since its best

cost is zero and the rest are all set to infinity). All of the destination nodes are currently at infinity, so they get updated, as shown in Figure 5.10.




Figure 5.10:
Aftermath of the
first step of
Dijkstra's
algorithm

Then it all has to be done again. The new node you pull off the priority queue is the top-left node, with a best cost of 8. It updates the top-right node and the center node, as shown in Figure 5.11.




Figure 5.11:
Step 2

The next node to come off the queue is the bottom-right one, with a value of 10. Its only destination node, the top-right one, already has a best cost of 13, which is less than 15 ($10 + \text{the cost of the edge} - 15$). Thus, the top-right node doesn't get updated, as shown in Figure 5.12.




Figure 5.12:
Step 3

Next is the top-right node. It updates the center node, giving it a new best cost of 14, producing Figure 5.13.




Figure 5.13:
Step 4

Finally, the center node is visited. It doesn't update anything. This empties the priority queue, giving the final graph, which appears in Figure 5.14.




Figure 5.14:
Step 5

Using Graphs to Find Shortest Paths

Now, armed with Dijkstra's algorithm, you can take a point and find the shortest path and shortest distance to all other visitable nodes on the graph. But one question remains: How is the graph to traverse generated? As it turns out, this is a simple automatic process, thanks to the spatial data structure.

First, the kind of behavior that you wish the creature to have needs to be established. When a creature's target exists in the same convex cell the creature is in, the path is simple: Go directly toward the object using something like the Terminator AI I discussed at the beginning of the chapter. There is no need to worry about colliding with walls since the definition of convexity assures that it is possible to just march directly toward the target.



Warning: I'm ignoring the fact that the objects take up a certain amount of space, so the total set of the creature's visitable points is slightly smaller than the total set of points in the convex cell. For the purposes of what I'm doing here, this is a tolerable problem, but a more robust application would need to take this fact into account.

So first there needs to be a way to tell in which cell an object is located. Luckily, this is easy to do. Each polygon in a cell has a plane associated with it. All of the planes are defined such that the normal points into the cell. Simply controlling the winding order of the polygons created does this. Also known is that each point can be classified as either in front of or in back of a plane. For a point to be inside a cell, it must be in front of all of the planes that make up the boundary of the cell.

It may seem mildly counterintuitive to have the normals sticking in toward the center of the object rather than outward, but remember that they're never going to be considered for drawing from the outside. The cells are areas of empty space surrounded by solid matter. You draw from the inside, and the normals point toward you when the polygons are visible, so the normals should point inside.

Now you can easily find out the cell in which both the source and destination are located. If they are in the same cell, you're done (marching toward the target). If not, more work needs to be done. You need to generate a path that goes from the source cell to the destination cell. To do this, you put nodes inside each portal, and throw edges back and forth between all the portals in a cell. An implementation detail is that a node in a portal is actually held by both of the cells on either side of the portal. Once the network of nodes is set up, building the edges is fairly easy. Add two edges (one each way) between each of the nodes in each cell. You have to be careful, as really intricate worlds with lots of portals and lots of nodes have to be carefully constructed so as not to overload the graph. (Naturally, the more edges in the graph, the longer Dijkstra's algorithm will take to finish its task.)

You may be wondering why I'm bothering with directed edges. The effect of having two directed edges going in opposite directions would be the same as having one bi-directed edge, and you would only have half the edges in the graph. In this 2D example there is little reason to have unidirectional edges. But in 3D everything changes. If, for example, the cell on the other side of the portal has a floor 20 feet below the other cell, you can't use the same behavior you use in the 2D example, especially when incorporating physical properties like gravity. In this case, you would want to let the creature walk off the ledge and fall 20 feet, but since the creature wouldn't be able to turn around and miraculously leap 20 feet into the air into the cell above, you don't want an edge that would tell you to do so.

Here is where you can start to see a very important fact about AI. Although a creature seems intelligent now (well...more intelligent than the basic algorithms at the beginning of the chapter would allow), it's following a very standard algorithm to pursue its target. It has no idea what gravity is, and it has no idea that it can't leap 20 feet. The intelligence in this example doesn't come from the algorithm itself, but rather it comes from the implementation, specifically the way the graph is laid out. If it is done poorly (for example, putting in an edge that told the creature to move forward even though the door was 20 feet above it), the creature will follow the same algorithm it always does but will look much less intelligent (walking against a wall repeatedly, hoping to magically cross through the doorway 20 feet above it).

Application: Path Planner

The second application for this chapter is a fully functioning path planner and executor. The code loads a world description off the disk, and builds an internal graph to navigate with. When the user clicks somewhere in the map, the little creature internally finds the shortest path to that location and then moves there.

Parsing the world isn't terribly hard; the data is listed in ASCII format (and was entered manually, yuck!). The first line of the file has one number, providing the number of cells. Following, separated by blank lines, are that many cells. Each cell has one line of header (containing the number of vertices, edges, portals, and items). Items were never implemented for this demo, but they wouldn't be too hard to add. It would be nice to be able to put health in the world and tell the creature "go get health!" and have it go get it.

Points are described with two floating-point coordinates, edges with two indices, and portals with two indices and a third index corresponding to the cell on the other side of the doorway.

The following section of code is from the cell description file.

```
17

6 5 1 0
-8.0 8.0
-4.0 8.0
-4.0 4.0
-5.5 4.0
-6.5 4.0
-8.0 4.0
0 1
1 2
2 3
4 5
5 0
3 4 8

... more cells
```

Building the graph is a little trickier. The way it works is that each pair of doorways (remember, each conceptual doorway has a doorway structure leading out of both of the cells touching it) holds onto a node situated in the center of the doorway. Each cell connects all of its doorway nodes together with dual edges—one going in each direction.

When the user clicks on a location, first the code makes sure that the user clicked inside the boundary of one of the cells. If not, the click is ignored. Only approximate boundary testing is used (using two-dimensional bounding boxes); more work would need to be done to enable exact hit testing (this is left as an exercise for the reader).

When the user clicks inside a cell, then the fun starts. Barring the trivial case (the creature and clicked location are in the same cell), a node is

created inside the cell and edges are thrown out to all of the doorway nodes. Then Dijkstra's algorithm is used to find the shortest path to the node. The shortest path is inserted into a structure called *sPath*, which is essentially just a stack of nodes. While the creature is following a path, it peeks at the top of the stack. If it is close enough to it within some epsilon, the node is popped off the stack and the next one is chosen. When the stack is empty, the creature has reached its destination.

The application uses the GDI for all the graphics, making it fairly slow. Also, the graph searching algorithm uses linear searches to find the cheapest node while it's constructing the shortest path. What fun would it be if I did all the work for you? A screenshot from the path planner appears in Figure 5.15. The creature appears as a small circle.

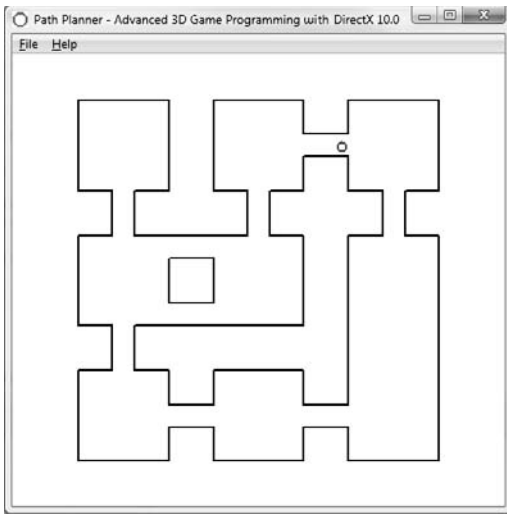


Figure 5.15:
Screenshot from
the path planner

There is plenty of other source code to wander through in this project, but this section that finds the shortest path in the graph seemed like the most interesting part.

```
cNode *cWorld::FindCheapestNode()
{
    // ideally, we would implement a slightly more advanced
    // data structure to hold the nodes, like a heap.
    // since our levels are so simple, we can deal with a
    // linear algorithm.

    float fBestCost = REALLY_BIG;
    cNode *pOut = NULL;
    for( int i=0; i<m_nodeList.size(); i++ )
    {
        if( !m_nodeList[i]->m_bVisited )
        {
            if( m_nodeList[i]->m_fCost < fBestCost )
            {
```

```

        // new cheapest node
        fBestCost = m_nodeList[i]->m_fCost;
        pOut = m_nodeList[i];
    }

}

// if we haven't found a node yet, something is
// wrong with the graph.
assert( pOut );

return pOut;
}

void cNode::Relax()
{
    this->m_bVisited = true;

    for( int i=0; i<m_edgeList.size(); i++ )
    {
        cEdge *pCurr = m_edgeList[i];
        if( pCurr->m_fWeight + this->m_fCost < pCurr->m_pTo->m_fCost )
        {
            // relax the 'to' node
            pCurr->m_pTo->m_pPrev = this;
            pCurr->m_pTo->m_fCost = pCurr->m_fWeight + this->m_fCost;
        }
    }
}

void cWorld::ShortestPath( sPath *pPath, cNode *pTo, cNode *pFrom )
{
    // easy out.
    if( pTo == pFrom ) return;

    InitShortestPath();

    pFrom->m_fCost = 0.f;

    bool bDone = false;
    cNode *pCurr;
    while( 1 )
    {
        pCurr = FindCheapestNode();
        if( !pCurr )
            return; // no path can be found.
        if( pCurr == pTo )
            break; // We found the shortest path

        pCurr->Relax(); // relax this node
    }

    // now we construct the path.

```

```

// empty the path first.
while( !pPath->m_nodeStack.empty() ) pPath->m_nodeStack.pop();

pCurr = pTo;
while( pCurr != pFrom )
{
    pPath->m_nodeStack.push( pCurr );
    pCurr = pCurr->m_pPrev;
}
}

```

Motivation

The final area of AI I'll be discussing is the motivation of a creature. I feel it's the most interesting facet of AI. The job of the motivation engine is to decide, at a very high level, what the creature should be doing. Examples of high-level states would be "get health" or "attack nearest player." Once you have decided on a behavior, you create a set of tasks for the steering engine to accomplish. Using the "get health" example, the motivation engine would look through an internal map of the world for the closest health and then direct the locomotion engine to find the shortest path to it and execute the path. I'll show you a few high-level motivation concepts.

Nondeterministic Finite Automata (NFAs)

NFAs are popular in simpler artificial intelligence systems (and not just in AI; NFAs are used everywhere). If, for example, you've ever used a search program like `grep` (a UNIX searching command), you've used NFAs. They're a classic piece of theoretic computer science, an extension of deterministic finite automata (DFAs).

How do they work? In the classic sense, you have a set of nodes connected with edges. One node (or more) is the start node and one (or more) is the end node. At any point in time, there is a set of active nodes. You send a string of data into an NFA. Each piece is processed individually.

The processing goes as follows: Each active node receives the current piece of data. It makes itself inactive and compares the data to each of its edges. If any of its outgoing edges match the input data, they turn their destination node on. There is a special type of edge called an epsilon edge that turns its destination on regardless of the input.

When all of the data has been processed, you look at the list of active nodes. If any of the end nodes are active, then that means the string of data passed. You construct the NFA to accept certain types of strings and can quickly run a string through an NFA to test it.

Here are a few examples to help make the definition more concrete. Both of the examples are fairly simple NFAs just to show the concepts being explained. Let's say there is an alphabet with exactly two values, A and B. The first example, shown in Figure 5.16, is an NFA that accepts only the string ABB and nothing else.

The second example, shown in Figure 5.17, is an NFA that accepts the string A^*B , where A^* means any number of As, including zero.




Figure 5.16:
NFA that accepts
the string ABB




Figure 5.17:
NFA that accepts
the string A^*B

How is this useful for game programming? If you encode the environment in which the creature exists into a string that you feed into an NFA, you can allow it to process its scene and decide what to do. You could have one goal state for each of the possible behaviors (that is, one for “attack enemy,” one for “get health,” and additional ones for other high-level behaviors). As an example, one of the entries in the array of NFA data could represent how much ammo the character has. Let’s say there are three possible states: {Plenty of ammo, Ammo, Little or no ammo}. The edge that corresponded to “Plenty of ammo” would lead to a section of the

NFA that would contain aggressive end states, while the “Little or no ammo” edges would lead to a section of the NFA that would most likely have the creature decide that it needed to get some ammo. The next piece of data would describe a different aspect of the universe the creature existed in, and the NFA would have branches ready to accept it.

Table 5.1 contains some examples of states that could be encoded in the string of data for the NFA.

Table 5.1: Some example states that could be encoded into an NFA

Proximity to nearest opponent	Very near; Average distance; Very far. If the nearest opponent is very far, the edge could lead to states that encourage the collection of items.
Health	Plenty of health; Adequate health; Dangerous health. If the creature has dangerously low health and the opponent was very near, a kamikaze attack would probably be in order. If the nearest enemy was very far away, it should consider getting some health.
Environment	Tight and close; Medium; Expansive. A state like this would determine which weapon to use. For example, an explosive weapon like a rocket launcher shouldn't be used in tight and close areas.
Enemy health	Plenty of health; Adequate health; Dangerous health. The health of the nearest enemy determines the attacking pattern of the creature. Even if the creature has moderate to low health, it should try for the kill if the enemy has dangerous health.
Enemy altitude	Above; Equal; Below. It's advantageous in most games to be on higher ground than your opponent, especially in games with rocket launcher splash damage. If the creature is below its nearest opponent and the opponent is nearby, it might consider retreating to higher ground before attacking.

One way to implement NFAs would be to have a function pointer in each end state that got executed after the NFA was processed if the end state succeeded.

The only problem with NFAs is that it's extremely difficult to encode fuzzy decisions. For example, it would be better if the creature's health was represented with a floating-point value so there would be a nearly continuous range of responses based on health. I'll show you how to use neural networks to do this. However, NFA-based AI can be more than adequate for many games. If your NFA's behavior is too simple, you generally only need to extend the NFA, adding more behaviors and more states.

Genetic Algorithms

While not directly a motivation concept, genetic algorithms (or GAs) can be used to tweak other motivation engines. They try to imitate nature to solve problems. Typically, when you're trying to solve a problem that has a fuzzy solution (like, for example, the skill of an AI opponent), it's very hard to tweak the numbers to get the best answer.

One way to solve a problem like this is to attack it the way nature does. In nature (according to Darwin, anyway) animals do everything they can to survive long enough to produce offspring. Typically, the only members of a species that survive long enough to procreate are the most superior of their immediate peers. In a pride of lions, only one male impregnates all of the females. Thus, all of the male lions vie for control of the pride so that their genes get carried on.

Added to this system, occasionally, is a bit of mutation. An offspring is the combination of the genes of the two parents, but it may be different from either of the parents by themselves. Occasionally, an animal will be born with bigger teeth, sharper claws, longer legs, or in Simpsonian cases, a third eye. The change might give that particular offspring an advantage over its peers. If it does, that offspring is more likely than the other animals to carry on its genes, and thus, over time, the species improves.

That's nice and all, but what does that have to do with software development? A lot, frankly. What if you could codify the parameters of a problem into genes? You could randomly create a set of animals, each with its own genes. They are set loose, they wreak havoc, and a superior pair of genes is found. Then you combine these two genes, sprinkle some random perturbations in, and repeat the process with the new offspring and another bunch of random creatures.

For example, you could define the behavior of all the creatures in terms of a set of scalar values that define how timid a creature is when it's damaged, how prone it is to change its current goal, how accurate its shots are when it is moving backward, and so forth. Correctly determining the best set of parameters for each of the creatures can prove difficult. Things get worse when you consider other types of variables, like the weapon the creature is using and the type of enemy it's up against.

Genetic algorithms to the rescue! Initially, you create a slew of creatures with a bunch of random values for each of the parameters and put them into a virtual battleground, having them duke it out until only two creatures remain. Those two creatures mate, combining their genes and sprinkling in a bit of mutation to create a whole new set of creatures, and the cycle repeats.

The behavior that genetic algorithms exhibit is called *hill climbing*. You can think of a creature's idealness as a function of n variables. The graph for this function would have many relative maximums and one absolute maximum. In the case where there were only two variables, you would see a graph with a bunch of hills (where the two parameters made a formidable opponent), a bunch of valleys (where the parameters made a bad

opponent), and an absolute maximum (the top of the tallest mountain: the best possible creature).

For each iteration, the creature that survives will hopefully be the one that was the highest on the graph. Then the iteration continues, with a small mutation (you can think of this as sampling the area immediately around the creature). The winner of the next round will be a little bit better than its parent as it climbs the hill. When the children stop getting better, you know you have reached the top of a hill, a relative maximum.

How do you know if you reached the absolute maximum, the tallest hill on the graph? It's extremely hard to do. If you increase the amount of mutation, you increase the area you sample around the creature, so you're more likely to happen to hit a point along the slope of the tallest mountain. However, the more you increase the sampling area, the less likely you are to birth a creature further up the mountain, so the function takes much longer to converge.

Rule-Based AI

The world of reality is governed by a set of rules, rules that control everything from the rising and setting of the sun to the way cars work. The AI algorithms discussed up to this point aren't aware of any rules, so they would have a lot of difficulty knowing how to start a car, for example.

Rule-based AI can help alleviate this problem. You define a set of rules that govern how things work in the world. The creature can analyze the set of rules to decide what to do. For example, let's say that a creature needs health. It knows that there is health in a certain room, but to get into the room the creature must open the door, which can only be done from a security station console. One way to implement this would be to hardcode the knowledge into the creature. It would run to the security station, open the door, run through it, and grab the health.

However, a generic solution has a lot of advantages. The behavior it can exhibit isn't limited to just opening security doors. Anything you can describe with a set of rules is something it can figure out.

```

IF [Health_Room == Visitable]
    THEN [Health == Gettable]
IF [Security_Door == Door_Open]
    THEN [Health_Room == Visitable]
IF [Today == Sunday]
    THEN [Tacos == 0.49]
IF [Creature_Health < 0.25]
    THEN [Creature_State = FindGettableHealth]
IF [Creature_Position NEAR Security_Console]
    THEN [Security_Console_Usable]
IF [Security_Console_Usable] AND [Security_Door != Door_Open]
    THEN [Creature_Use(Security_Console)]
IF [Security_Console_Used]
    THEN [Security_Door == Door_Open]
IF [Creature_Move_To(Security_Console)]
    THEN [Creature_Position NEAR Security_Console]

```

Half the challenge in setting up rule-based systems is to come up with an efficient way to encode the rules. The other half is actually creating the rules. Luckily a lot of the rules, like the `Creature_Move_To` rule at the end of the list, can be automatically generated.

How does the creature figure out what to do, given these rules? It has a goal in mind: getting health. It looks in the rules and finds the goal it wants: `[Health == Gettable]`. It then needs to satisfy the condition for that goal to be true, which is `[Health_Room == Visitable]`. The creature can query the game engine and ask it if the health room is visitable. When the creature finds out that it is not, it has a new goal: making the health room visitable.

Searching the rules again, it finds that `[Health_Room == Visitable]` if `[Security_Door == Door_Open]`. Once again, it sees that the security door is not open, so it analyzes the rule set again, looking for a way to satisfy the condition.

This process continues until the creature reaches the rule saying that if it moves to the security console, it will be near the security console. Finally, a command that it can do! It then uses path planning to get to the security console, presses the button to open the security door, moves to the health room, and picks up the health.

AI like this can be amazingly neat. Nowhere do you tell how to get the health. It actually figured out how to do it all by itself. If you could encode all the rules necessary to do anything in a particular world, then the AI would be able to figure out how to accomplish whatever goals it wanted. The only tricky thing is encoding this information in an efficient way. And if you think that's tricky, try getting the creature to develop its own rules as it goes along. If you can get that, your AI will always be learning, always improving.

Neural Networks

One of the huge areas of research in AI is in neural networks (NNs). They take a very fundamental approach to the problem of artificial intelligence by trying to closely simulate intelligence, in the physical sense.

Years of research have gone into studying how the brain actually works (it's mystifying that evolution managed to design an intelligence capable of analyzing itself). Researchers have discovered the basic building blocks of the brain and have found that, at a biological level, it is just a really, really (REALLY) dense graph. On the order of billions or trillions of nodes, and each node is connected to thousands of others.

The difference between the brain and other types of graphs is that the brain is extremely connected. Thinking of several concepts brings up several other concepts, simply through the fact that the nodes are connected. As an example, think for a moment about an object that is leafy, green, and crunchy. You most likely thought about several things, maybe celery or some other vegetable. That's because there is a strong connection between the leafy part of your brain and things that are leafy. When the leafy

neuron fires, it sends its signal to all the nodes it's connected to. The same goes for green and crunchy. Since, when you think of those things, they all fire and all send signals to nodes, some nodes receive enough energy to fire themselves, such as the celery node.

Now, I'm not going to attempt to model the brain itself, but you can learn from it and build your own network of electronic neurons. Graphs that simulate brain activity in this way are generally called *neural networks*.

Neural networks are still a very active area of research. In the last year or so, a team was able to use a new type of neural network to understand garbled human speech better than humans can! One of the big advantages of neural networks is that they can be trained to remember their past actions. You can teach them, giving them an input and then telling them the correct output. Do this enough times and the network can learn what the correct answer is.

However, that is a big piece of pie to bite down on. Instead, I'm going to delve into a higher-level discussion of neural networks, by explaining how they work sans training and providing code for you to play with.

A Basic Neuron

Think of a generic neuron in your brain as consisting of three biological parts: an *axon*, *dendrites*, and a *soma*. The processing unit is the soma. It takes input coming from the dendrites and outputs to the axon. The axon, in turn, is connected to the dendrites of other neurons, passing the signals on. These processes are all handled with chemicals in real neurons; a soma that fires is, in essence, sending a chemical down its axon that will meet up with other dendrites, sending the fired message to other neurons. Figure 5.18 shows what a real neuron looks like.




Figure 5.18:
A biological
neuron

The digital version is very similar. There is a network of nodes connected by edges. When a node is processed, it takes all of the signals on the incoming edges and adds them together. One of these edges is a special bias or *memory edge*, which is just an edge that is always on. This value can change to modify the behavior of the network (the higher the bias value, the more likely the neuron is to fire). If the summation of the inputting nodes is above the threshold (usually 1.0), then the node sends a fire signal to each of its outgoing edges. The fire signal is not the result of the addition, as that may be much more than 1.0. It is always 1.0. Each edge also has a bias that can scale the signal being passed it higher or lower. Because of this, the input that arrives at a neuron can be just about any value, not just 1.0 (firing neurons) or 0 (non-firing neurons). They may be anywhere; if the edge bias was 5.0, for example, the neuron would receive 5.0 or 0, depending on whether or not the neuron attached to it fired. Using a bias on the edges can also make a fired neuron have a dampening effect on other neurons.

The equation for the output of a neuron can be formalized as follows:

$$x = b(B) + \sum_n \text{bias}_n(\text{node}_n)$$

$$\text{out} = \begin{cases} 1.0 & \text{if } x > 1.0 \\ 0.0 & \text{otherwise} \end{cases}$$

where you sum over the inputs n (the bias of the edge, multiplied by the output of the neuron attached to it) plus the weight of the bias node times the bias edge weight.

Other types of responses to the inputs are possible; some systems use a sigmoid exponential function like the one below. A continuous function such as this makes it easier to train certain types of networks (back propagation networks, for example), but for these purposes the all-or-nothing response will do the job.

$$x = b(B) + \sum_n \text{bias}_n(\text{node}_n)$$

$$\text{out} = \frac{1.0}{1.0 + e^{-x}}$$

One of the capabilities of the brain is the ability to imagine things given a few inputs. Imagine you hear the phrases “vegetable,” “orange,” and “eaten by rabbits.” Your mind’s eye conjures up an image of carrots. Imagine your neural network’s inputs are these words and your outputs are names of different objects. When you hear the word “orange,” somewhere in your network (and your brain) an “orange” neuron fires. It sends a fire signal to objects you have associated with the word “orange” (for example: carrots, oranges, orange crayons, an orange shirt). That signal alone probably won’t be enough for any particular one of those other neurons to fire; they need other signals to help bring the total over the threshold. If you

then hear another phrase, such as “eaten by rabbits,” the “eaten by rabbits” neuron will fire off a signal to all the nodes associated with that word (for example: carrots, lettuce, boisterous English crusaders). Those two signals may be enough to have the neuron fire, sending an output of carrots. Figure 5.19 abstractly shows what is happening.

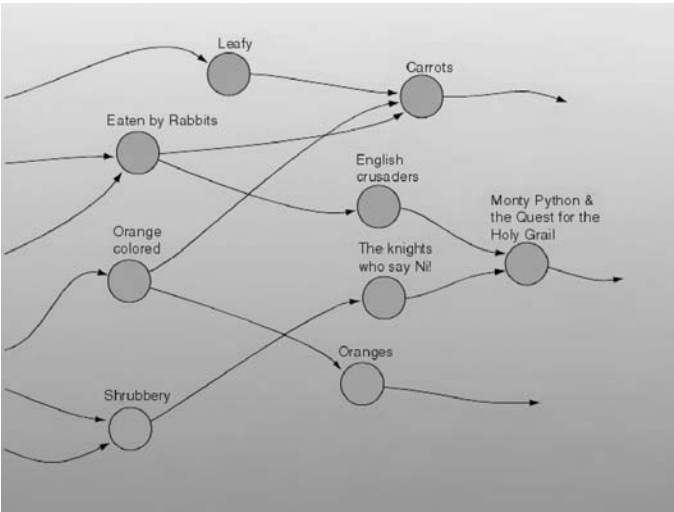


Figure 5.19:
A subsection of a
hypothetical
neural network

Simple Neural Networks

Neural networks are Turing-complete; that is, they can be used to perform any calculation that computers can do, given enough nodes and enough edges. Given that you can construct any processor using nothing but NAND gates, this doesn’t seem like too ridiculous a conjecture. Let’s look at some simpler neural networks before trying to tackle anything more complex.

AND

Binary logic seems like a good place to start. As a first stab at a neural net, let’s try to design a neural net that can perform a binary AND. The network appears in Figure 5.20.




Figure 5.20:
A neural network
that can perform
a binary AND
function

Note that the input nodes have a bias of 0.1. This is to help fuzzify the numbers a bit. You could make the network strict if you'd like (setting the bias to 0.0), but for many applications 0.9 is close enough to 1.0 to count as being 1.0.

OR

Binary OR is similar to AND; the middle edges just have a higher weight so that either one of them can activate the output node. The net appears in Figure 5.21.




Figure 5.21:
A neural network
that can perform
a binary OR
function

XOR

Handling XOR requires a bit more thought. Three nodes alone can't possibly handle XOR; you need to make another layer to the network. Here's a semi-intuitive reasoning behind the workings of Figure 5.22: The top internal node will only be activated if both input nodes fire. The bottom one will fire if either of the input nodes fires. If both internal nodes fire, that means that both input nodes fired (a case you should not accept), which is correctly handled by having a large negative weight for the edge leading from the top internal node to the output node.




Figure 5.22:
A neural network
that can perform
a binary XOR
function

Training Neural Networks

While it's outside the scope of this book, it's important to know one of the most important and interesting features about neural nets: They can be trained. Suppose you create a neural net to solve a certain problem (or put another way, to give a certain output given a set of inputs). You can initially seed the network with random values for all of the edge biases and then have the network learn. Neural nets can be trained or can learn autonomously. An autonomously learning neural net would be, for example, an AI that was trying to escape from a maze. As it moves, it learns more information, but it has no way to check its answer as it goes along. These types of networks learn much slower than trained networks. Trained neural networks, on the other hand, have a cheat sheet; that is, they know the solution to each problem. They run an input and check their output against the correct answer. If it is wrong, the network modifies some of the weights so that it gets the correct answer the next time.

Using Neural Networks in Games

Using a neural network to decide the high-level action to perform in lieu of NFAs has a lot of advantages. For example, the solutions are often much fuzzier. Reaching a certain state isn't as black and white as achieving a certain value in the string of inputs; it's the sum of a set of factors that all contribute to the behavior.

As an example, let's say that you have a state that, when reached, causes your creature to flee its current location in search of health. You may want to do this in many cases. One example would be if there was a strong enemy nearby. Another would be if there was a mildly strong enemy nearby and the main character is low on health. You can probably conjure up a dozen other cases that would justify turning tail and fleeing.

While it's possible to codify all of these cases separately into an NFA, it's rather tedious. It's better to have all of the input states (proximity of nearest enemy, strength of nearest enemy, health, ammo, etc.) become inputs into the neural network. Then you could just have an output node that, when fired, caused the creature to run for health. This way, the behavior emerges from the millions of different combinations for inputs. If enough factors contribute to the turn-and-flee state to make it fire, it will sum over the threshold and fire.

A neural network that does this is exactly what I'm going to show you how to write.

Application: NeuralNet

The NeuralNet sample application is a command-line application to show off a neural network simulator. The network is loaded off disk from a description file; input values for the network are requested from the user, then the network is run and the output appears on the console. I'll also build a sample network that simulates a simple creature AI. In this example, the creature has low health, plenty of ammo, and an enemy nearby. The network decides to select the state [Flee_Enemy_Towards_Health]. If this code were to be used in a game, state-setting functions would be called in lieu of printing out the names of the output states.

```
Advanced 3D Game Programming with DirectX 10.0
```

```
-----  
Neural Net Simulator
```

```
Using nn description file [creature.nn]
```

```
Neural Net Inputs:
```

```
-----  
Ammo (0..1)  
1 - Ammo (0..1)  
Proximity to enemy (0..1)  
1 - Proximity to enemy (0..1)
```



```
Health (0..1)
1 - Health (0..1)

Enter Inputs:
-----
Enter floating point input for [Ammo (0..1)]
1.0
```

The NeuralNet description file (*.nn) details the network that the application will run. Each line that isn't a comment starts with a keyword describing the data contained in the line. The keywords appear in Table 5.2.

Table 5.2: Neural net description keywords

NN_BEGIN	Defines the beginning of the neural network. Always the first line of the file. First token is the number of layers in the neural network. The input layer counts as one, and so does the output layer.
NN_END	Defines the ending of the neural network description.
NEURON	Declares a neuron. The first token is the name of the neuron, and the second is the bias of the neuron.
INPUT	Declares an input. The first token is the name of the neuron to receive the input, and the second token (enclosed in quotes) is the user-friendly name for the input. The list of inputs is iterated for the user prior to running the simulation.
DEFAULTOUT	The default output of the neural network. The only token is the text of the default output.
OUTPUT	Declares an output. The first token is the name of the neuron, the second is the text to print if the neuron fires, and the third is the bias of the neuron.
EDGE	Declares an edge. The first token is the name of the source node, the second token is the name of the destination node, and the third token is the floating-point weight of the edge.

The order in which the neurons appear in the file is pivotally important. They are appended to an STL vector as they are loaded in, and the vector is traversed when the network is run. Therefore, they should appear ordered in the file as they would appear left to right in the diagrams presented thus far (the input nodes at the beginning, the internal nodes in the middle, the output nodes at the end).

The following is a simplistic creature AI that can attack, flee, and find items it needs. The network is simple enough that it's easy to see that adding more states wouldn't be too hard a task. It's important to note that this network is designed to have its inputs range from -1 to 1 (so having health input as 0 means the creature has about 50% health).

```
# First line starts the NN loading and gives the # of layers.
NN_BEGIN 2
#
```

```

# NEURON x y z
# x = layer number
# y = node name
# z = node bias
NEURON 0 health 0.0
NEURON 0 healthInv 0.0
NEURON 0 ammo 0.0
NEURON 0 ammoInv 0.0
NEURON 0 enemy 0.0
NEURON 0 enemyInv 0.0
NEURON 1 findHealth 0.2
NEURON 1 findAmmo 0.2
NEURON 1 attackEnemy 0.5
NEURON 1 fleeToHealth 0.5
NEURON 1 fleeToAmmo 0.5
#
# DEFAULTOUT "string"
# string = the default output
DEFAULTOUT "Chill out"
#
# EDGE x y z
# x = source neuron
# y = dest neuron
# z = edge weight
#
EDGE health attackEnemy 0.5
EDGE ammo attackEnemy 0.5
EDGE enemy attackEnemy 0.5
EDGE healthInv attackEnemy -0.5
EDGE ammoInv attackEnemy -0.5
EDGE enemyInv attackEnemy -0.6
#
EDGE healthInv findHealth 0.6
EDGE enemyInv findHealth 0.6
#
EDGE ammoInv findAmmo 0.6
EDGE enemyInv findAmmo 0.6
#
EDGE healthInv fleeToHealth 0.8
EDGE enemy fleeToHealth 0.5
#
EDGE ammoInv fleeToAmmo 0.8
EDGE enemy fleeToAmmo 0.5
#
# INPUT/OUTPUT x "y"
# x = node for input/output
# y = fancy name for the input/output
INPUT health "Health (0..1)"
INPUT healthInv "1 - Health (0..1)"
INPUT ammo "Ammo (0..1)"
INPUT ammoInv "1 - Ammo (0..1)"
INPUT enemy "Proximity to enemy (0..1)"
INPUT enemyInv "1 - Proximity to enemy (0..1)"
OUTPUT findHealth "Find Health"
OUTPUT findAmmo "Find Ammo"

```

```

OUTPUT attackEnemy "Attack Nearest Enemy"
OUTPUT fleeToHealth "Flee Enemy Towards Health"
OUTPUT fleeToAmmo "Flee Enemy Towards Ammo"
#
NN_END

```

The source code for the neural network simulator appears next.

```

/*****
 *           Advanced 3D Game Programming with DirectX 10.0
 * * * * *
 *
 * See license.txt for modification and distribution information
 * copyright (c) 2007 by Peter Walsh, Wordware
 *****/

#ifndef _NEURALNET_H
#define _NEURALNET_H

#include <string>
#include <vector>
#include <map>

using namespace std;

#include "file.h"

class cNeuralNet
{
protected:
    class cNode;
    class cEdge;

public:
    wstring GetOutput();

    void SendInput( const TCHAR *inputName, float amt );

    void Load( cFile& file );
    void Run();
    void Clear();
    cNeuralNet();
    virtual ~cNeuralNet();

    void ListInputs();
    void GetInputs();

protected:

    cNode *FindNode( const TCHAR *name );

    class cNode
    {
    public:

```

```

void Init( const TCHAR *name, float weight );

void Clear();
virtual void Run();

void AddOutEdge( cNode *target, float edgeWeight );
void SendInput( float in );

const TCHAR *GetName() const;
float GetTotal() const;
protected:

    // Computes the output function given the total.
    virtual float CalcOutput();

    wstring m_name;
    float m_weight; // initial bias in either direction
    float m_total; // total of the summed inputs up to this point

    vector< cEdge > m_outEdges;

};

class cEdge
{
    cNode *m_pSrc;
    cNode *m_pDest;
    float m_weight;
public:

    cEdge( cNode *pSrc, cNode *pDst, float weight );

    void Fire( float amount );
};

friend class cNode;

vector< vector< cNode* > > m_nodes;

// maps the names of output nodes to output strings.
map< wstring, wstring > m_inputs;
map< wstring, wstring > m_outputs;

wstring m_defaultOutput;
};

inline const TCHAR *cNeuralNet::cNode::GetName() const
{
    return m_name.c_str();
}

inline float cNeuralNet::cNode::GetTotal() const
{
    return m_total;
}

```

```

}

#endif // _NEURALNET_H

/*****
 *      Advanced 3D Game Programming with DirectX 10.0
 *      * * * * *
 *
 *      See license.txt for modification and distribution information
 *      copyright (c) 2007 by Peter Walsh, Wordware
 *****/

using namespace std;

#include "stdafx.h"
#include <math.h> // for atof, exp
#include <stdio.h>
#include <stdlib.h> // for atof
#include "NeuralNet.h"

int wmain(int argc, TCHAR *argv[])
{
    // Sorry, I don't do cout.
    wprintf_s( L"Advanced 3D Game Programming with DirectX 10.0\n" );
    wprintf_s( L"-----\n\n" );
    wprintf_s( L"Neural Net Simulator\n\n" );

    if( argc != 2 )
    {
        wprintf_s(L"Usage: neuralnet filename.nn\n");
        return 0;
    }

    wprintf_s(L"Using nn description file [%s]\n\n", argv[1] );

    cNeuralNet nn;
    cFile nnFile;
    nnFile.Open( (TCHAR*)argv[1] );
    nn.Load( nnFile );
    nnFile.Close();

    int done = 0;
    while( !done )
    {
        // Clear the totals
        nn.Clear();

        // List the inputs for the net from the user
        nn.ListInputs();

        // Get the inputs for the net from the user

```

```

        nn.GetInputs();

        // Run the net
        nn.Run();

        // Get the net's output.
        wstring output = nn.GetOutput();

        wprintf_s(L"\nNeural Net output was [%s]\n", output.c_str() );
        wprintf_s(L"\nRun Again? (y/n)\n");
        TCHAR buff[80];
        _getws_s( buff, 80 );

        if( !(buff[0] == 'y' || buff[0] == 'Y') )
        {
            done = 1;
        }
    }
    return 0;
}

cNeuralNet::cNeuralNet()
{
    // no work needs to be done.
}

cNeuralNet::~cNeuralNet()
{
    // delete all of the nodes. (each node will get its outgoing edges)
    int numLayers = m_nodes.size();
    for( int i=0; i<numLayers; i++ )
    {
        int layerSize = m_nodes[i].size();
        for( int j=0; j<layerSize; j++ )
        {
            delete m_nodes[i][j];
        }
    }
}

cNeuralNet::cNode *cNeuralNet::FindNode( const TCHAR *name)
{
    cNode *pCurr;

    // Search for the node.
    int numLayers = m_nodes.size();
    for( int i=0; i<numLayers; i++ )
    {
        int layerSize = m_nodes[i].size();
        for( int j=0; j<layerSize; j++ )
        {
            pCurr = m_nodes[i][j];
            if( 0 == wcsncmp( pCurr->GetName(), name ) )

```

```
        return pCurr;
    }
}

// didn't contain the node (bad)
wprintf_s( L"ERROR IN NEURAL NET FILE!\n");
wprintf_s( L"Tried to look for node named [%s]\n", name );
wprintf_s( L"but couldn't find it!\n");
exit(0);
return NULL;
}

void cNeuralNet::Clear()
{
    // Call clear on each of the networks.
    cNode *pCurr;

    int numLayers = m_nodes.size();
    for( int i=0; i<numLayers; i++ )
    {
        int layerSize = m_nodes[i].size();
        for( int j=0; j<layerSize; j++ )
        {
            pCurr = m_nodes[i][j];
            pCurr->Clear();
        }
    }
}

void cNeuralNet::Run()
{
    // Run each layer, running each node in each layer.
    int numLayers = m_nodes.size();
    for( int i=0; i<numLayers; i++ )
    {
        int layerSize = m_nodes[i].size();
        for( int j=0; j<layerSize; j++ )
        {
            m_nodes[i][j]->Run();
        }
    }
}

void cNeuralNet::SendInput( const TCHAR *inputTarget, float amt)
{
    // Find the node that we're sending the input to, and send it.
    FindNode( inputTarget )->SendInput( amt );
}

void cNeuralNet::cNode::Clear()
{
    // initial total is set to the bias
    m_total = m_weight;
}
```

```

void cNeuralNet::cNode::Run()
{
    // Compute the transfer function
    float output = CalcOutput();

    // Send it to each of our children
    int size = m_outEdges.size();
    for( int i=0; i< size; i++ )
    {
        m_outEdges[i].Fire( output );
    }
}

void cNeuralNet::cNode::Init( const TCHAR *name, float weight)
{
    m_name = wstring( name );
    m_weight = weight;
}

float cNeuralNet::cNode::CalcOutput()
{
    // This can use an exponential-type function
    // but for simplicity's sake we're just doing
    // flat yes/no.
    if( m_total >= 1.0f )
        return 1.0f;
    else
        return 0.f;
}

void cNeuralNet::cNode::SendInput(float in)
{
    // just add the input to the total for the network.
    m_total += in;
}

void cNeuralNet::cNode::AddOutEdge(cNode *target, float edgeWeight)
{
    // Create an edge structure
    m_outEdges.push_back( cEdge( this, target, edgeWeight) );
}

void cNeuralNet::Load(cFile &file)
{
    TCHAR buff[1024];
    TCHAR *pCurr;
    TCHAR delim[] = L" \t\n\r";

    TCHAR *context;

    while( 1 )
    {
        file.ReadNonCommentedLine( buff, '#' );

```



```

pCurr = wcstok_s( buff, delim, &context );
if( 0 == wscmp( pCurr, L"NN_BEGIN" ) )
{
    // Read in the # of layers
    int nLayers = _wtoi( wcstok_s( NULL, delim, &context ) );
    for( int i=0; i<nLayers; i++ )
    {
        // populate the vector-of-vectors with vectors.
        vector< cNeuralNet::cNode *> garbage;
        m_nodes.push_back( garbage );
    }
}
else if( 0 == wscmp( pCurr, L"NN_END" ) )
{
    // We're done loading the network. Break from the loop
    break;
}
else if( 0 == wscmp( pCurr, L"NEURON" ) )
{
    int layerNum = _wtoi( wcstok_s( NULL, delim, &context ) );
    cNode *pNew = new cNode();
    TCHAR *name;
    name = wcstok_s( NULL, delim, &context );
    pCurr = wcstok_s( NULL, delim, &context );
    float val;
    swscanf_s(pCurr, L"%f", &val );
    pNew->Init( name, val );

    // done initializing. Add it to the list
    m_nodes[layerNum].push_back( pNew );
}
else if( 0 == wscmp( pCurr, L"INPUT" ) )
{
    wstring nodeName = wstring( wcstok_s( NULL, delim, &context ) );

    TCHAR *pRestOfString = wcstok_s( NULL, L"\n\r", &context );

    TCHAR *pStrStart = wcschr( pRestOfString, '\"' );
    TCHAR *pStrEnd = wcsrchr( pRestOfString, '\"' );

    assert( pStrStart && pStrEnd && (pStrStart!=pStrEnd) );

    m_inputs[ nodeName ] = wstring( pStrStart+1, pStrEnd );
}
else if( 0 == wscmp( pCurr, L"OUTPUT" ) )
{
    wstring nodeName = wstring( wcstok_s( NULL, delim, &context ) );

    TCHAR *pRestOfString = wcstok_s( NULL, L"\n\r", &context );

    TCHAR *pStrStart = wcschr( pRestOfString, '\"' );
    TCHAR *pStrEnd = wcsrchr( pRestOfString, '\"' );

```

```

        assert( pStrStart && pStrEnd && (pStrStart!=pStrEnd) );

        m_outputs[ nodeName ] = wstring( pStrStart+1, pStrEnd );
    }
    else if( 0 == wcsncmp( pCurr, L"DEFAULTOUT" ) )
    {
        TCHAR *pRestOfString = wcstok_s( NULL, L"\n\r", &context );

        TCHAR *pStrStart = wcschr( pRestOfString, '"' );
        TCHAR *pStrEnd = wcsrchr( pRestOfString, '"' );

        assert( pStrStart && pStrEnd && (pStrStart!=pStrEnd) );

        m_defaultOutput = wstring( pStrStart+1, pStrEnd );
    }
    else if( 0 == wcsncmp( pCurr, L"EDGE" ) )
    {
        TCHAR *src;
        TCHAR *dst;
        src = wcstok_s( NULL, L" \t\n\r", &context );
        dst = wcstok_s( NULL, L" \t\n\r", &context );

        pCurr = wcstok_s( NULL, L" \t\n\r", &context );
        float val;
        swscanf_s(pCurr, L"%f", &val );

        FindNode( src )->AddOutEdge( FindNode( dst ), val );
    }

    }

}

cNeuralNet::cEdge::cEdge( cNode *pSrc, cNode *pDest, float weight)
: m_pSrc( pSrc )
, m_pDest( pDest )
, m_weight( weight)
{
    // all done.
}

void cNeuralNet::cEdge::Fire( float amount )
{
    // Send the signal, multiplied by the weight,
    // to the destination node.
    m_pDest->SendInput( amount * m_weight );
}

//=====-----

void cNeuralNet::ListInputs()
{
    wprintf_s(L"\n");
    wprintf_s(L"Neural Net Inputs:\n");

```

```
wprintf_s(L"-----\n");

map<wstring, wstring>::iterator iter;
for( iter = m_inputs.begin(); iter != m_inputs.end(); iter++ )
{
    wprintf_s(L"%s\n", (*iter).second.c_str() );
}

void cNeuralNet::GetInputs()
{
    wprintf_s(L"\n");
    wprintf_s(L"Enter Inputs:\n");
    wprintf_s(L"-----\n");

    map<wstring, wstring>::iterator iter;
    for( iter = m_inputs.begin(); iter != m_inputs.end(); iter++ )
    {
        wprintf_s(L"Enter floating point input for [%s]\n", (*iter).second.c_str()
    );
        TCHAR buff[80];
        _getws_s( buff, 80 );
        float value = (float)_wtof( buff );

        cNode *pNode = FindNode( (*iter).first.c_str() );

        pNode->SendInput( value );
    }
}

wstring cNeuralNet::GetOutput()
{
    map<wstring, wstring>::iterator iter;

    float fBest = 1.f;
    map<wstring, wstring>::iterator best = m_outputs.end();

    for( iter = m_outputs.begin(); iter != m_outputs.end(); iter++ )
    {
        // check the output. Is it the best?
        cNode *pNode = FindNode( (*iter).first.c_str() );
        if( pNode->GetTotal() > fBest )
        {
            // new best.
            fBest = pNode->GetTotal();
            best = iter;
            wprintf_s(L"new total: %f\n", fBest );
        }
    }
    if( best == m_outputs.end() )
    {
        return m_defaultOutput;
    }
}
```

```
    else
    {
        return (*best).second.c_str();
    }
}
```

Extending the System

Creating a successful AI engine needs the combined effort of a bunch of different concepts and ideas. Neural networks alone can't do much, but combine them with path planning and you can create a formidable opponent. Use genetic algorithms to evolve the neural networks (well, actually just the bias weights for the neural networks) and you can breed an army of formidable opponents, each one different in its own way. It's a truly exciting facet of game programming, and you're only cheating yourself if you don't investigate the topic further!

This page intentionally left blank.

Chapter 6

Multiplayer Internet Networking with UDP

Doom. Quake IV. Crackdown. Unreal. It seems like every game released these days is written to be played on the Internet. It's the wave of the future—the world is becoming a global village and there's a dire need to kill everyone in town. But writing a game and writing a game that can be played over the Internet are two very different things. Far too many games have died in their infancy because programmers assumed it would be a simple matter of adding in network code when everything else was done. Nothing could be further from the truth. In this chapter, I'm going to show you the basics of setting up a network game and reducing the amount of lag, and then investigate some possible optimizations.

Terminology

First, you need to know some basic terminology.

Endianness

There are a few major problems with the Internet. First, it's completely unorganized; second, data sent over the Internet has a good chance of never reaching its destination. It's important to understand some of what is going on inside the Internet in order to overcome these problems.

The reason the Internet is so unorganized is that it is still evolving. Different operating systems, different hardware—it can be a real headache. By far one of the furthest reaching differences is that of *endianness*. When a computer CPU has to store information that takes up more than 1 byte, most types of CPUs will store the bytes in order from largest to smallest. This is known as *little endian*. However, other machines (namely Apples) do it a little differently. Suppose you have a whole number that takes up 2 bytes. In a little endian system, the bytes are stored with bits 0-7 representing the values 20-27 and bits 8-15 representing values 28-35. But the CPU in a *big endian* system stores the same value the other way around, with bits 0-7 representing values 28-35 and bits 8-15 representing values 20-27.




Figure 6.1: Big-endian vs. little-endian



Note: By the way, if you’re wondering where this “endian” terminology comes from, this is the note for you. It originates from the book *Gulliver’s Travels*. Somewhere in the book there are people who are split into two groups. The first group, the little endians, ate their hard-boiled eggs from the small side first. The other group, the big endians, ate their eggs from the big side. So it’s not really a technical term!

This means that when you want to send data over the Internet you have to make sure that your numbers are in the default endianness of the Internet, and when you receive information you have to be sure to switch back to your computer’s default endianness.

Another effect of such an unorganized system is that a system of addresses had to be created so as to tell one machine from another. Known as *host addresses* or *IP addresses*, they take the form *nnn.nnn.nnn.nnn*, where *nnn* is a whole number from 0 to 255. This would imply that there can’t be more than 4 billion machines on the Internet simultaneously, but in practice the limit is quite a bit lower. The reason is that there are three types of networks: class A, class B, and class C. Each class uses a certain range of possible IP addresses, severely limiting the total possible combinations. For example, class C networks use IP addresses in the range 192.0.0.x through 223.255.255.x. Class C networks use the first three bits to identify the network as class C, which leaves 21 bits for identifying a computer on the network. The value 2^{21} is a total of 2,097,152 possible addresses; as of the beginning of 2007, 75% of those have been assigned to computers that are always online. But it gets worse—class B networks have a total of 16,384 combinations, and class A networks only have 128. Fortunately, new legislation is changing all that. For more information, check out the American Registry for Internet Numbers (www.arin.net).

There are two kinds of addresses: For those machines that are always on the Internet, there are static addresses that never change. For those computers that have dial-up connections or that aren’t regularly on the Internet, there are dynamic addresses that are different each time the computer connects.

With all the phone numbers, bank accounts, combination locks, secret passwords, and shoe sizes, there isn't much room left over in most people's memories for a collection of IP addresses. So in order to make things a little more user friendly, *host names* were introduced. A host name such as `www.realtimeworlds.com` or `www.gamedev.net` represents the four-number address. If the IP address changes, the host name keeps working. A host name has to be resolved back into the IP address before it can be used to make a connection attempt. In order to resolve a host name, the computer trying to resolve must already know the address of a *domain name server* (DNS). It contacts the DNS and sends the host name. The DNS server responds by returning the numeric IP address of the host.

With so many programs running on so many different computers around the globe, there has to be a way to separate communication into different "channels," much like separate phone lines or TV stations. Inside any Winsock-compliant computer are 65,534 imaginary *ports* to which data can be sent. Some recognized protocols have default ports—HTTP uses port 80 and FTP uses port 21 (more on protocols in a moment). Any program can send data to any port, but if there's no one listening the data will be ignored, and if the listening program doesn't understand the data then things could get ugly. In order to listen or transmit data to a port, both machines must begin by initializing *Winsock*, a standard library of methods for accessing network firmware/hardware. Winsock is based in part on UNIX sockets, so most methods can be used on either type of operating system without the need for rewriting. Once Winsock has been initialized, the two machines must each create a socket handle and associate that socket with a port. Having successfully completed socket association, all that remains to do is transfer data and then clean up when you're done. But your problems are just beginning. Once you've finished all the fundamental communication code, there should only be one thing on your mind: speed, speed, and more speed.

Network Models

In order to make games run smoothly some kind of order has to be imposed on the Internet; some clearly defined way of making sure that every player in the game sees as close to the same thing as possible. The first thought that leaps to mind is "connect every machine to every other machine!" This is known as a *peer-to-peer* configuration, and it sounds like a good configuration. In fact it was used in some of the first networked games. However, as the number of players rise, this peer-to-peer model quickly becomes impractical. Consider a game with four players. Each player must have three connections to other players for a total of six connections. Each player also has to send the same data three times. Hmm. Dubious. Now consider the same game with six players. Each player has to send the same data out five times and there are a total of 15 connections.

In an eight-player game there are 28 connections. Try it yourself—the equation is

$$\frac{P \times (P - 1)}{2}$$

where P is the number of players.

Another method might be to arrange all the players in a *ring*, with each player connected to two other machines. This sounds a bit better because there are only (P + 1) connections and each player only has to send data once, clockwise around the ring. My computer tells your computer, your computer tells her computer, and so on around the ring until it comes back to me, at which point I do nothing. But consider the amount of time it takes to send information from one computer to another. Even if a computer sends data in both directions at once it will still take too long for data to travel halfway around the ring. Things can become pretty complicated if one of the player’s computers suddenly crashes or leaves the game—all the data that it had received but not yet transmitted to the next machine in the ring suddenly vanishes, leaving some machines with one version of the game and some with another.




Figure 6.2:
Peer-to-peer, ring, and
client/server network
configurations

The most popular design is a *client/server* configuration, which might look like a star because every player is connected to a powerful central computer. This central computer is the server and it makes sure everyone is synchronized and experiencing the same thing. There are, at most, P connections in a client/server configuration and the clients only have to send data once. The server does the bulk of the work, sending out the same data, at most, P times. This method ensures the smallest possible time difference between any two clients but can be quite taxing on the server, so much so that some computers are *dedicated servers* that do nothing but keep games running for other people.

Protocols

Once you’ve decided on which model you’re going to use (I reach out with my mind powers and see that you have chosen client/server...), there comes the decision of what protocol. *Protocols* are accepted standard languages that computers use to transmit data back and forth. Most protocol information does not have to be set up by the programmer (phew) but is required for making sure that data reaches its intended destination. At the core, most protocols consist of basically the same thing. The following is a

list of some of the more commonly used protocols and a brief description of each.

- *Internet Protocol (IP)* is one of the simplest protocols available. Programmers use protocols built on top of IP in order to transmit data.
- *User Datagram Protocol (UDP)* adds the barest minimum of features to IP. It's just enough to transmit data and it does the job very fast. However, UDP data is unreliable, meaning that if you send data through UDP, some of it may not reach the destination machine, and even if it does it may not arrive in the order that it was sent. However, for real-time games this is the protocol of choice and the one I will be covering in this chapter.
- *Transmission Control Protocol (TCP)* is also built on top of IP and adds a lot of stability to network data transmission at the expense of speed. If you want reliable data transmission, this is the protocol for you. It is best suited for turn-based games that don't have to worry about speed so much as making sure the right information reaches its destination. This is not to say it can't be used for real-time games—*NetStorm* (an Activision title) uses TCP/IP. However, it is my considered opinion that the amount of data being transmitted in *NetStorm* is far lower than in, say, *Unreal Tournament*.
- *Internet Control Message Protocol (ICMP)* is built on top of IP and provides a way to return error messages. ICMP is at the heart of every ping utility. Its features are duplicated in so many other protocols that sometimes its features are mistaken as being a part of IP

Packets

Any data you transmit is broken into *packets*, blocks of data of a maximum size. Each packet is then prefixed with a header block containing information such as host address, host port, the amount of time the packet has been traveling through the Internet, and whatever else the selected protocol requires. A UDP packet's maximum size is 4096 bytes. If it doesn't seem like much, you're right. But remember, most dial-up connections have trouble reaching 10 kilobytes per second, and you're going to be transmitting a lot of information both ways. When a packet is sent out into the Internet, it must first travel to one of the servers that forms the backbone of the Internet. Up until that point the packet is traveling in a straight line and there's no confusion. However, the moment the packet reaches the first server on the backbone it starts to follow a new set of rules. Each computer with more than one connection has a series of weights associated with each connection to another computer. These weights determine which computer should receive the most traffic. If these weights should change, the packets from your computer could take a much longer route than necessary and in some cases never even reach their destination. The real drawback is that it means packets may not arrive at their destination in

any particular order. It could also be that every copy of the packet takes a scenic route, dies of old age, and the machine you were trying to send to never gets the packet at all.

In this case I'm willing to sacrifice a little reliability in exchange for the increased speed, but sometimes there are messages being sent to the server that must get through. For that reason I'll show you how to set up a reliable, ordered communication system for the few, the proud, the brave: the UDP packets.

Implementation I: MTUDP

Design Considerations

Since this is a tutorial, I'm only going to develop for the Windows platform. This means that I don't have to be very careful about endianness. I've also chosen to use UDP because it's the fastest way to communicate and makes the most sense in a real-time game. I'm also going to design with the client/server (star) configuration in mind, because it is the most scalable and the most robust.

Things to Watch Out For

In all the online tutorials that I've read about creating multiplayer networked games, there's always one detail that's left out, and that detail is about the same size and level of danger as an out-of-control 18-wheel truck. The problem is multithreading.

Consider for a moment a simple single-thread game: In your main loop you read in from the keyboard, move things in the world, and then draw to the screen. So it would seem reasonable that in a network game you would read in from the keyboard, read in from the Internet, move things, send messages back out to the Internet, and then draw to the screen. Sadly, this is not the case. Oh sure, you can write a game to work like this (I learned the hard way), but it won't work the way you expect it to. Why? Let's say your computer can draw 20 frames per second. That means that most of 50ms is being spent drawing, which means that nearly 50ms go by between any two checks for new data from the Internet. So what? Anyone who's ever played a network game will tell you that 50ms can mean the difference between life and death. So when you send a message to another computer, that message could be waiting to be read for nearly 50ms and the reply could be waiting in your machine's hardware for an extra 50ms for an extra round-trip time of 100ms!

Worse still is the realization that if you stay with a single-threaded app, there's nothing you can do to solve the problem; nothing will make that delay go away. Yes, it would be shorter if the frame rate were higher. But just try to tell people they can't play unless their frame rate is high enough—I bet good money they tar and feather you.

The solution is, of course, to write a multithreaded app. I'll admit the first time I had to write one I was pretty spooked. I thought it was going to be a huge pain. Please believe me when I say that as long as you write clean, careful code, you can get it right the first time and you won't have to debug a multithreaded app. And since everything you do from here on in will depend on that multithreading, let's start the MTUDP (multithreaded UDP) class there. First of all, be sure that you tell the compiler you're not designing a single-threaded app. In Microsoft Visual C++ 2005 Express Edition, the option to change is in Project|Settings|Configuration Properties|C/C++|Code Generation|Use Runtime Library.

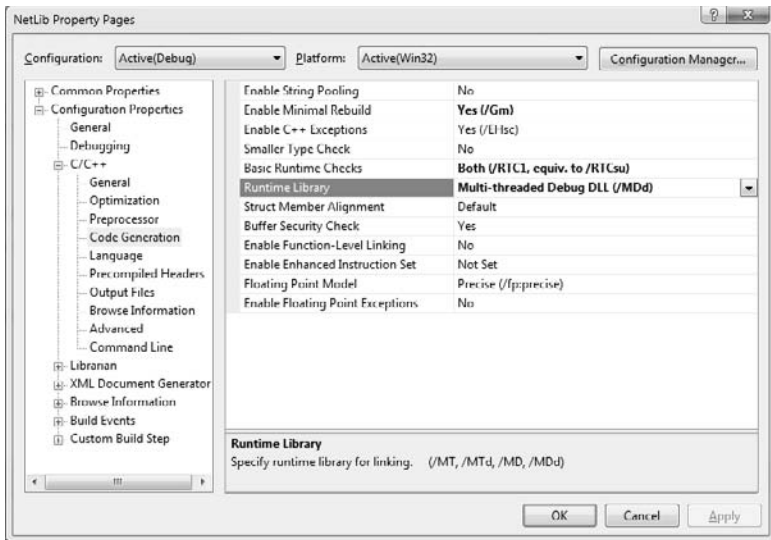


Figure 6.3: Setting up multithreading

Windows multithreading is, from what I hear, completely insane in its design. Fortunately, it's also really easy to get started. `CreateThread()` is a standard Windows function and, while I won't go into much detail about its inner workings here (you have MSDN; look it up!), I will say that I call it as follows:

```
void cThread::Begin()
{
    d_threadHandle = CreateThread( NULL,
                                   0,
                                   (LPTHREAD_START_ROUTINE)gsThreadProc,
                                   this, 0, (LPDWORD)&d_threadID );

    if( d_threadHandle == NULL )
        throw cError( "cThread() - Thread creation failed." );
    d_bIsRunning = true;
}
```

As you can tell, I've encapsulated all my thread stuff into a single class. This gives me a nice place to store `d_threadHandle` so I can kill the thread later, and it means I can use `cThread` as a base class for, oh, say, `MTUDP`, and I can reuse the `cThread` class in all my other applications without making any changes to it.

`CreateThread()` takes six parameters, the most important of which are the third and fourth parameters, `gsThreadProc` and `this`. `this` is a pointer to the instance of `cThread` and will be sent to `gsThreadProc`. This is crucial because `gsThreadProc` *cannot* be a class function because Windows doesn't like that. Instead, `gsThreadProc` is defined at the very beginning of `cThread.cpp` as follows:

```
static DWORD WINAPI cThreadProc( cThread *pThis )
{
    return pThis->ThreadProc();
}
```

I don't know about you, but I think that's pretty sneaky. It also happens to work! Back in the `cThread` class, `ThreadProc()` is a virtual function that returns zero immediately. `ThreadProc()` can return anything you like, but I've always liked to return zero when there is no problem and use all the other numbers as error codes.

Sooner or later you're going to want to stop the thread. Again, this is pretty straightforward.

```
void cThread::End()
{
    if( d_threadHandle != NULL )
    {
        d_bIsRunning = false;
        WaitForSingleObject( d_threadHandle, INFINITE );
        CloseHandle( d_threadHandle );
        d_threadHandle = NULL;
    }
}
```

The function `cThread::End()` is set up in such a way that you can't stop a thread more than once, but the real beauty is hidden. Notice `d_bIsRunning`? Well, you can use it for more than just telling the other threads that you're still working. Let's look at a simple version of a derived class's `ThreadProc()`.

```
DWORD MTUDP::ThreadProc()
{
    while( d_bIsRunning == true )
    {
        // Read and process network data here.
    }
    return 0;
}
```

This means that the moment `d_bIsRunning` is set to false, the thread will quit. Of course, we could get the thread to quit at any time—if it detected

an error, for example. This is an easy way for one thread to have start/stop control on another thread. In fact, if you didn't set `d_bIsRunning` to false, the first thread would stop running forever while it waited for `WaitForSingleObject(d_threadHandle, INFINITE)`. This is because `d_threadHandle` functions like a mutex.

Mutexes

Mutexes are crucial in multithreading because they protect data that is in a critical section. For example, let's say you have a linked list of information. One thread is adding to the linked list and the other thread is removing. What would happen if the two tried to access the linked list at the same time? There's a chance that one thread could walk through the linked list and then step off into "funny RAM" (some undefined location in RAM that is potentially dangerous to modify) because the other thread hadn't finished working with the linked list pointers.

Fortunately, C++ lets you set up a really nice little class to monitor these critical sections and make sure every thread plays nice.

```
class cMonitor
{
protected:

    HANDLE d_mutex;

public:
    cMonitor();
    virtual ~cMonitor();

    void MutexOn() const;
    void MutexOff() const;
};
```

Again, this class is used as a base class for every class that has a critical section. In fact, I defined `cThread` as `class cThread : public cMonitor`. The four `cMonitor` functions are also very interesting.

```
cMonitor::cMonitor()
{
    // This mutex will help the two threads share their toys.
    d_mutex = CreateMutex( NULL, false, NULL );
    if( d_mutex == NULL )
        throw cError( "cMonitor() - Mutex creation failed." );
}

cMonitor::~cMonitor()
{
    if( d_mutex != NULL )
    {
        CloseHandle( d_mutex );
        d_mutex = NULL;
    }
}
```

cMonitor will create a new mutex and clean up after itself.

```
void cMonitor::MutexOn() const
{
    WaitForSingleObject( d_mutex, INFINITE );
}

void cMonitor::MutexOff() const
{
    ReleaseMutex( d_mutex ); // To be safe...
}
```

Once again you see that `WaitForSingleObject()` will stall a thread forever if necessary. The big difference between this and `d_threadHandle` is that `d_threadHandle` was released by Windows. Here, control is left up to a thread. If `WaitForSingleObject()` is called, the thread will gain control of a mutex and all other threads will have to wait until that same thread calls `ReleaseMutex()` before they get a turn, and it's first come, first serve. This means you have to be very careful with how you handle your mutexes—if you don't match every `WaitFor...` with a `ReleaseMutex()`, threads will hang forever and you will soon find yourself turning your computer off to reboot it. I suppose I could have written a version of `MutexOn()` that would wait *n* milliseconds and return an error code, but I haven't found a need for it yet.

Threads, Monitor, and the Problem of the try/throw/catch Construction

Try/throw/catch is a wonderful construction that can simplify your debugging. Unfortunately, it doesn't work very well inside other threads. Actually, it works, but it might surprise you. The following would work, but it would not catch anything thrown by the other thread.

```
// somewhere inside thread #1
try
{
    cThreadDerived myNewThread;

    mNewThread.Begin();

    // Do other stuff.
}
catch( cError &err )
{
    // No error would be reported.
}

// somewhere inside cThreadDerived::ThreadProc()
throw cError( "Gack!" );
```

The solution is to put an extra try/catch inside the ThreadProc() of cThreadDerived and then store the error somewhere for the other thread to read it or process it right there and then.

MTUDP: The Early Years

You've seen the multithreading class and you've got a way to protect the critical sections. So here's how it's going to work: The main thread can read and send data with MTUDP whenever it can get the mutex. The rest of the time it can render, check the keyboard, play music, etc. MTUDP, meanwhile, will be constantly rechecking the network to see if there is data to be read in from the Internet and processing any that arrives.

Now you can start getting down to business!

```
class MTUDP : public cThread
{
protected:
    SOCKET          d_listenSocket,
                   d_sendSocket;
    unsigned short  d_localListenPort,
                   d_foreignListenPort;
    bool            d_bStarted,
                   d_bListening,
                   d_bSending;
    // A list of all the data packets that have arrived.

public:
    MTUDP();
    virtual ~MTUDP();

    virtual ThreadProc();

    void          Startup( unsigned short localListenPort,
                           unsigned short ForeignListenPort );
    void          Cleanup();
    void          StartListening();
    void          StartSending();
    void          StopListening();
    void          StopSending();
    unsigned short GetReliableData( char *const pBuffer,
                                   unsigned short maxLen );
    void          ReliableSendTo( const char *const pStr, unsigned short len );
};
```

Startup() and Cleanup() are the bookends of the class and are required to initialize and tidy up. StartListening() and StartSending() will create the d_listenSocket and d_sendSocket, respectively. One or more of these has to be called before ReliableSendTo() or GetReliableData() will do anything.

MTUDP::Startup() and MTUDP::Cleanup()

```

void MTUDP::Startup( unsigned short localListenPort,
                    unsigned short foreignListenPort )
{
    Cleanup(); // just in case somebody messed up out there...

    WSADATA wsaData;
    int     error;

    error = WSASStartup( MAKEWORD( 2, 2 ), &wsaData );
    if( error == SOCKET_ERROR )
    {
        char errorBuffer[ 100 ];

        error = WSAGetLastError();
        if( error == WSAVERNOTSUPPORTED )
        {
            sprintf( errorBuffer,
"MTUDP::Startup() - WSASStartup() error.\nRequested v2.2, found only v%d.%d.",
LOBYTE( wsaData.wVersion ), HIBYTE( wsaData.wVersion ) );
            WSACleanup();
        }
        else
            sprintf( errorBuffer, "MTUDP::Startup() - WSASStartup() error %d",
WSAGetLastError() );

        throw cError( errorBuffer );
    }

    d_localListenPort = localListenPort;
    d_foreignListenPort = foreignListenPort;
    d_bytesTransferred = 0;
    d_bStarted = true;
}

```

Really the only mystery here is `WSASStartup()`. It takes two parameters: a word that describes what version of Winsock you'd like to use and a pointer to an instance of `WSADATA`, which will contain all kinds of useful information regarding this machine's Winsock capabilities. I admire the way the Winsock programmers handle errors—just about everything will return `SOCKET_ERROR`, at which point you can call `WSAGetLastError()` to find out more information. The two variables passed to `Startup` (`d_localListenPort` and `d_foreignListenPort`) will be used a little later.

```

void MTUDP::Cleanup()
{
    if( d_bStarted == false )
        return;

    d_bStarted = false;

    StopListening();
    StopSending();
}

```

```
// Clean up all data waiting to be read

WSACleanup();
}
```

An important note: `WSAStartup()` causes a DLL to be loaded, so be sure to match every call to `WSAStartup()` with exactly one call to `WSACleanup()`.

MTUDP::MTUDP() and MTUDP::~~MTUDP()

All programming (and, as I've learned, all attempts to explain things to people) should follow the "method of least surprise." `MTUDP`'s creation and destruction methods prove we've been able to stick to that rule.

At this point all the creation method does is initialize `d_bStarted` to false and the destruction method calls `Cleanup()`.

MTUDP::StartListening()

Now we get to put `d_localListenPort` to use.

```
void MTUDP::StartListening()
{
    if( d_bListening == true ||
        d_bStarted == false )
        return;

    d_bListening = true;
    //Nothing special yet; this just prevents you from
    //calling StartListening() twice.

    d_listenSocket = socket( AF_INET, SOCK_DGRAM, 0 );
    if( d_listenSocket == INVALID_SOCKET )
        // throw an error here.
```

The Winsock method `socket()` creates a socket. The three parameters are the address family, the socket type, and the protocol type (which modifies the socket type). The only parameter here you should ever mess with is `SOCK_DGRAM`, which could be changed to `SOCK_STREAM` if you wanted to work in TCP/IP

```
SOCKADDR_IN localAddr;
int          result;

memset( &localAddr, 0, sizeof( SOCKADDR_IN ) );
localAddr.sin_family = AF_INET;
localAddr.sin_addr.s_addr = htonl( INADDR_ANY );
localAddr.sin_port = htons( d_localListenPort );

result = bind( d_listenSocket,
               (sockaddr *)&localAddr,
               sizeof( SOCKADDR_IN ) );
if( result == SOCKET_ERROR )
{
    closesocket( d_listenSocket );
```

```
// throw another error.
}
```

The `bind()` method takes three parameters—the port number on which to open the new listening socket, some information about the type of socket (in the form of a `SOCKADDR` or `SOCKADDR_IN` structure), and the size of the second parameter. Every time you `bind()` a socket you have to make `sin_family` equal to the same thing as the socket's address family. Since this is a listening socket you want it to be on port `d_localListenPort`, so that's what `sin_port` is set to. The parameter `sin_addr.s_addr` is the address you would be sending data to. The listen socket will never send any data, so set it to `INADDR_ANY`. Lastly, if the `bind()` fails, be sure to close the socket. There's only one step left!

```
// We're go for go!
cThread::Begin();
}
```

MTUDP::StartSending()

The start of `StartSending()` is the same old deal—check that a send socket has not been opened (`d_bSending == false`) and create the send socket (which looks exactly the same as it did in `StartListening()`). The only significant change comes in the call to `bind()`.

```
SOCKADDR_IN localAddr;
int          result;

memset( &localAddr, 0, sizeof( SOCKADDR_IN ) );
localAddr.sin_family = AF_INET;
localAddr.sin_addr.s_addr = htonl( INADDR_ANY );
localAddr.sin_port = htons( 0 );

result = bind( d_sendSocket, (sockaddr *)&localAddr, sizeof( SOCKADDR_IN ) );
if( result == SOCKET_ERROR )
    // close the socket and throw an error.
```

I don't care what port the send socket is bound to, so `sin_port` is set to zero. Even though data is being sent, because UDP is being used, the `sin_addr.s_addr` is once again set to `INADDR_ANY`. This would have to change if you wanted to use TCP/IP, because once you open a TCP/IP socket it can only send to one address until it is closed or forced to change.

At the end of `StartSending()`, you do not call `cThread::Begin()`. Thanks to the `cThread` class it wouldn't have an effect, so make sure to call `StartListening()` before `StartSending()`. Another good reason to call `StartListening()` first is because there's a very small chance that the random port Winsock binds your send socket to is the same port you want to use for listening.

MTUDP::ThreadProc()

Now to the real meat and potatoes. I'll explain the whole thing at the end.

```

DWORD MTUDP::ThreadProc()
{
    if( d_bListening == false )
        return 0; // Quit already?! It can happen...

    char            inBuffer[ MAX_UDPBUFFERSIZE ];
    timeval         waitTimeStr;
    SOCKADDR_IN     fromAddr;
    int             fromLen;
    unsigned short  result;
    FD_SET          set;

    try
    {
        while( d_bListening == true )
        {
            // Listen to see if there is data waiting to be read.
            FD_ZERO( &set );
            FD_SET( d_listenSocket, &set );

            waitTimeStr.tv_sec = 0; // Wait 0 seconds
            waitTimeStr.tv_usec = 0; // Wait 0 microseconds (1/(1*10^6) seconds)

            // select() tells us if there is data to be read.
            result = select( FD_SETSIZE, &set, NULL, NULL, &waitTimeStr );
            if( result == 0 )
                continue;
            if( result == SOCKET_ERROR )
                // throw an error.

            // recvfrom() gets the data and puts it in inBuffer.
            fromLen = sizeof( SOCKADDR );
            result = recvfrom( d_listenSocket,
                              inBuffer,
                              MAX_UDPBUFFERSIZE,
                              0,
                              (SOCKADDR *)&fromAddr,
                              &fromLen );

            if( result == 0 )
                continue;
            if( result == SOCKET_ERROR )
                // throw an error.

            // Put the received data in a mutex-protected queue here.
            ProcessIncomingData( inBuffer,
                                result,
                                ntohl( fromAddr.sin_addr.s_addr ),
                                GetTickCount() );

        } // while
    } // try
    catch( cError &err )
    {

```

```

        // do something with err.d_text so that the
        // other thread knows this thread borked.
    }

    // Returns 1 if the close was not graceful.

    return d_bListening == true;
}

```

It may seem a little weird to put a check for `d_bListening` at the start of the thread proc. I added it because there is a short delay between when you call `Begin()` and when `ThreadProc()` is actually called, and even if you clean up properly when you're going to quit, it can make your debug output look a little funny.

`MAX_UDPBUFFERSIZE` is equal to the maximum size of a UDP packet—4096 bytes. I seriously doubt you will ever send a UDP block this big, but it never hurts to play it safe. As you can see, `try/catch/throw` is here, just like I said. The next step is the while loop, which begins with a call to `select()`. `select()` will check any number of sockets to see if there is data waiting to be read, check if one of the sockets can send data, and/or check if an error occurred on one of the sockets. `select()` can be made to wait for a state change as long as you want, but I set `waitTimeStr` to 0 milliseconds so that it would poll the sockets and return immediately. That way it's a little more thread friendly.

Some of you may have some experience with Winsock and are probably wondering why I didn't use something called "asynchronous event notification." Two reasons: First, it takes a lot of effort to get set up and then clean up again. Second, it makes MTUDP dependent on a window handle, which makes it dependent on the speed at which `WndProc()` messages can be parsed, and it would make MTUDP even more dependent on Windows functions, something we'd like to avoid if possible.

The next steps only happen if there is data to be read. `recvfrom()` will read in data from a given socket and return the number of bytes read, but no more than the `MAX_UDPBUFFERSIZE` limit. `recvfrom()` will also supply some information on where the data came from in the `fromAddr` structure.

If some data was successfully read into `inBuffer`, the final step in the while loop is called. This is a new MTUDP function called `ProcessIncomingData()`.

MTUDP::ProcessIncomingData()

Well, I'm sorry to say that, for now, `ProcessIncomingData()` is virtually empty. However, it is the first opportunity to see mutexes in action.

```

void MTUDP::ProcessIncomingData( char *const pData,
                                unsigned short length,
                                DWORD address,
                                DWORD receiveTime )
{
    cMonitor::MutexOn();
}

```

```

    // Add the data to our list of received packets.
    cMonitor::MutexOff();
}

```

MTUDP::GetReliableData()

GetReliableData() is one of the few methods that can be called by another thread. Because it also messes with the list of received packets, mutexes have to be used again.

```

unsigned short MTUDP::GetReliableData( char *const pBuffer,
                                       unsigned short maxLen )
{
    if( pBuffer == NULL )
        throw cError( "MTUDP::GetReliableData() - Invalid parameters." );

    if( maxLen == 0 )
        return 0;

    cMonitor::MutexOn();
    // take one of the received packets off the list.
    cMonitor::MutexOff();
    // fill pBuffer with the contents of the packet.
    // return the size of the packet we just read in.
}

```

You've now got everything required to asynchronously read data from the Internet while the other thread renders, reads input, picks its nose, gives your hard drive a wedgie, you name it; it's coded. Of course, it doesn't really tell you who sent the information, and it's a long way from being reliable.

Reliable Communications

It's a good thing that I left this for the end because some of the code to get ReliableSendTo() working will help with reliable communications. In music circles this next bit would be called a bridge—the melody changes, maybe even enters a new key, but it gets you where you need to go.

cDataPacket

You've probably had all sorts of ideas on how to store the incoming data packets. I'm going to describe my data packet format, which may be a little puzzling at first. Trust me—by the end it will all make perfect sense.

```

// this file is eventually inherited everywhere else, so this seemed
// like a good place to define it.
#define MAX_UDPBUFFERSIZE 4096

class cDataPacket
{
public:
    char          d_data[ MAX_UDPBUFFERSIZE ];

```

```

    unsigned short  d_length,
                   d_timesSent;
    DWORD           d_id,
                   d_firstTime,
                   d_lastTime;

    cDataPacket();
    virtual ~cDataPacket();

    void Init( DWORD time,
              DWORD id,
              unsigned short len,
              const char *const pData );

    cDataPacket &operator=( const cDataPacket &otherPacket );
};

```

As always, it follows the K.I.S.S. (keep it simple, stupid) principle. `Init()` sets `d_firstTime` and `d_lastTime` to `time`, `d_id` to `id`, and `d_length` to `len`, and copies `len` bytes from `pData` into `d_data`. The `=` operator copies one packet into another.

cQueueIn

`cQueueIn` stores all the data packets in a nice, neat, orderly manner. In fact, it keeps two lists—one for data packets that are in order and one for the rest (which are as ordered as can be, given that some may be missing from the list).

```

class cQueueIn : public cMonitor
{
protected:
    list<cDataPacket *> d_packetsOrdered;
    list<cDataPacket *> d_packetsUnordered;
    DWORD              d_currentPacketID,
                      d_count; // number of packets added to this queue.

public:
    cQueueIn();
    virtual ~cQueueIn();

    void          Clear();
    void          AddPacket( DWORD packetID,
                           const char *const pData,
                           unsigned short len,
                           DWORD receiveTime );

    cDataPacket   *GetPacket();
    bool          UnorderedPacketIsQueued( DWORD packetID );
    DWORD         GetHighestID();
    inline DWORD  GetCurrentID(); // returns d_currentPacketID.
    inline DWORD  GetCount();     // returns d_count.
};

```

`d_currentPacketID` is equal to the highest ordered packet ID plus 1. `Clear()` removes all packets from all lists. `GetPacket()` removes the first packet in the `d_packetsOrdered` list (if any) and returns it. `UnorderedPacketIsQueued()` informs the caller if the packet is in the `d_packetsUnordered` list and returns true if `packetID < d_currentPacketID`. `GetHighestID()` returns the highest unordered packet ID plus 1 (or `d_currentPacketID` if `d_packetsUnordered` is empty). In fact, the only tricky part in this whole class is `AddPacket()`.

```
void cQueueIn::AddPacket( DWORD packetID,
                        const char *const pData,
                        unsigned short len,
                        DWORD receiveTime )
{
    if( pData == NULL ||
        len == 0 ||
        d_currentPacketID > packetID )
        return;

    // Create the packet.
    cDataPacket *pPacket;

    pPacket = new cDataPacket;
    if( pPacket == NULL )
        throw cError( "cQueueIn::AddPacket() - insufficient memory." );

    pPacket->Init( receiveTime, packetID, len, pData );

    // Add the packet to the queues.
    cMonitor::MutexOn();

    if( d_currentPacketID == pPacket->d_id )
    {
        // This packet is the next ordered packet. Add it to the ordered list
        // and then move all unordered that can be moved to the ordered list.
        d_packetsOrdered.push_back( pPacket );
        d_currentPacketID++;
        d_count++;

        pPacket = *d_packetsUnordered.begin();
        while( d_packetsUnordered.empty() == false &&
            d_currentPacketID == pPacket->d_id )
        {
            d_packetsUnordered.pop_front();
            d_packetsOrdered.push_back( pPacket );
            d_currentPacketID++;
            pPacket = *d_packetsUnordered.begin();
        }
    }
    else // d_currentPacketID < pPacket->d_id
    {
        // Out of order. Sort into the list.
        list<cDataPacket *>::iterator iPacket;
```



```

bool bExists;

bExists = false;

for( iPacket = d_packetsUnordered.begin();
    iPacket != d_packetsUnordered.end(); ++iPacket )
{
    // Already in list - get out now!
    if( (*iPacket)->d_id == pPacket->d_id )
    {
        bExists = true;
        break;
    }
    if( (*iPacket)->d_id > pPacket->d_id )
        break;
}

if( bExists == true )
    delete pPacket;
else
{
    // We've gone 1 past the spot where pPacket belongs. Back up and insert.
    d_packetsUnordered.insert( iPacket, pPacket );
    d_count++;
}
}

cMonitor::MutexOff();
}

```

Now I could stop right here, add an instance of `cQueueIn` to `MTUDP`, and there would be almost everything needed for reliable communications, but that's not why I went off on this tangent. There is still no way of sending data to another computer and also no way of telling who the data came from.

cHost

Yes, this is another new class. Don't worry, there's only four more, but they won't be mentioned for quite some time. The `cHost` class doesn't contain much yet, but it will be expanded later.

[illegible]

```

void          SetPort( unsigned short port );
bool          SetAddress( const char *const pAddress );
bool          SetAddress( DWORD address );
DWORD        GetAddress(); // returns d_address.
unsigned short GetPort();   // returns d_port.

cQueueIn      &GetInQueue(); // returns d_inQueue.
};

```

There are only two big mysteries here: `SetAddress()` and `ProcessIncomingReliable()`.

```

bool cHost::SetAddress( const char *const pAddress )
{
    if( pAddress == NULL )
        return true;

    IN_ADDR *pAddr;
    HOSTENT *pHe;

    pHe = gethostbyname( pAddress );
    if( pHe == NULL )
        return true;

    pAddr = (in_addr *)pHe->h_addr_list[ 0 ];
    d_address = ntohl( pAddr->s_addr );

    return false;
}

```

The other `SetAddress()` assumes you've already done the work, so it just sets `d_address` equal to `address` and returns.

As I said before, the `cHost` you're working with is a really simple version of the full `cHost` class. Even `ProcessIncomingReliable()`, which I'm about to show, is a simple version of the full `ProcessIncomingReliable()`.

```

unsigned short cHost::ProcessIncomingReliable( char *const pBuffer,
                                              unsigned short maxLen,
                                              DWORD receiveTime )
{
    DWORD        packetID;
    char         *readPtr;
    unsigned short length;

    readPtr = pBuffer;
    memcpy( &packetID, readPtr, sizeof( DWORD ) );
    readPtr += sizeof( DWORD );
    memcpy( &length, readPtr, sizeof( unsigned short ) );
    readPtr += sizeof( unsigned short );

    // If this message is a packet, queue the data
    // to be dealt with by the application later.
    d_inQueue.AddPacket( packetID, (char *)readPtr, length, receiveTime );
    readPtr += length;
}

```

```
// d_inQueue::d_count will be used here at a much much later date.

return readPtr - pBuffer;
}
```

This might seem like overkill, but it will make the program a lot more robust and net-friendly in the near future.

Things are now going to start building on the layers that came before. To start with, MTUDP is going to store a list<> containing all the instances of cHost, so the definition of MTUDP has to be expanded.

```
// Used by classes that call MTUDP, rather than have MTUDP return a pointer.
typedef DWORD   HOSTHANDLE;

class MTUDP : public cThread
{
private:
    // purely internal shortcuts.
    typedef map<HOSTHANDLE, cHost *> HOSTMAP;
    typedef list<cHost *>             HOSTLIST;

protected:
    HOSTLIST      d_hosts;
    HOSTMAP       d_hostMap;
    HOSTHANDLE    d_lastHandleID;

public:
    HOSTHANDLE    HostCreate( const char *const pAddress,
                             unsigned short port );
    HOSTHANDLE    HostCreate( DWORD address, unsigned short port );
    void          HostDestroy( HOSTHANDLE hostID );
    unsigned short HostGetPort( HOSTHANDLE hostID );
    DWORD         HostGetAddress( HOSTHANDLE hostID );
```

So what exactly did I do here? Well, MTUDP returns a unique HOSTHANDLE for each host so that no one can do anything silly (like try to delete a host). It also means that because MTUDP has to be called for everything involving hosts, MTUDP can protect d_hostMap and d_hosts with the cThread::cMonitor.

Now, it may surprise you to know that MTUDP creates hosts at times other than when some outside class calls HostCreate(). In fact, this is a perfect time to also show you just what's going to happen to cHost::QueueIn() by revisiting MTUDP::ProcessIncomingData().

```
void MTUDP::ProcessIncomingData( char *const pData, unsigned short length,
                                DWORD address, DWORD receiveTime )
{
    // Find the host that sent this data.
    cHost      *pHost;
    HOSTLIST::iterator iHost;

    cMonitor::MutexOn();
    // search d_hosts to find a host with the same address.
    if( iHost == d_hosts.end() )
```

```

{
    // Host not found! Must be someone new sending data to this computer.
    DWORD hostID;

    hostID = HostCreate( address, d_foreignListenPort );
    if( hostID == 0 )
        // turn mutex off and throw an error, the host creation failed.
        pHost = d_hostMap[ hostID ];
}
else
    pHost = *iHost;

assert( pHost != NULL );

// This next part will get more complicated later.
pHost->ProcessIncomingReliable( pData, length, receiveTime );
}

```

Of course, that means you now have a list of hosts. Each host might contain some new data that arrived from the Internet, so you're going to have to tell the other thread about it somehow. That means you're going to have to make changes to `MTUDP::GetReliableData()`.

```

unsigned short MTUDP::GetReliableData( char *const pBuffer,
                                       unsigned short maxLen,
                                       HOSTHANDLE *const pHostID )
{
    if( pBuffer == NULL ||
        pHostID == NULL )
        throw cError( "MTUDP::GetReliableData() - Invalid parameters." );

    if( maxLen == 0 )
        return 0;

    cDataPacket      *pPacket;
    HOSTLIST::iterator iHost;

    pPacket = NULL;

    cMonitor::MutexOn();

    // Is there any queued, ordered data?
    for( iHost = d_hosts.begin(); iHost != d_hosts.end(); ++iHost )
    {
        pPacket = (*iHost)->GetInQueue().GetPacket();
        if( pPacket != NULL )
            break;
    }

    cMonitor::MutexOff();

    unsigned short length;

    length = 0;

```

```

if( pPacket != NULL )
{
    length = pPacket->d_length > maxLen ? maxLen : pPacket->d_length;
    memcpy( pBuffer, pPacket->d_data, length );

    delete pPacket;

    *pHostID = (*iHost)->GetAddress();
}

return length;
}

```

See how I deal with pPacket copying into pBuffer after I release the mutex? This is an opportunity to reinforce a very important point: *Hold on to a mutex for as little time as possible*. A perfect example: Before I had a monitor class my network class had one mutex. Naturally, it was being held by one thread or another for vast periods of time (20ms!), and it was creating the same delay effect as when I was only using one thread. Boy, was my face black and blue (mostly from hitting it against my desk in frustration).

MTUDP::ReliableSendTo()

Finally! Code first, explanation later.

```

void MTUDP::ReliableSendTo( const char *const pStr, unsigned short length,
                           HOSTHANDLE hostID )
{
    if( d_bSending == false )
        throw cError( "MTUDP::ReliableSendTo() - Sending not initialized!" );

    cHost *pHost;

    cMonitor::MutexOn();

    pHost = d_hostMap[ hostID ];
    if( pHost == NULL )
        throw cError( "MTUDP::ReliableSendTo() - Invalid parameters." );

    char            outBuffer[ MAX_UDPBUFFERSIZE ];
    unsigned short  count;
    DWORD           packetID;

    count = 0;
    memset( outBuffer, 0, MAX_UDPBUFFERSIZE );

    // Attach the message data.
    packetID = pHost->GetOutQueue().GetCurrentID();
    if( pStr )
    {
        // Flag indicating this block is a message.
        outBuffer[ count ] = MTUDPMSGTYPE_RELIABLE;
    }
}

```

```

count++;

memcpy( &outBuffer[ count ], &packetID, sizeof( DWORD ) );
count += sizeof( DWORD );
memcpy( &outBuffer[ count ], &length, sizeof( unsigned short ) );
count += sizeof( unsigned short );
memcpy( &outBuffer[ count ], pStr, length );
count += length;
}

// Attach the previous message, just to ensure that it gets there.
cDataPacket secondPacket;

if( pHost->GetOutQueue().GetPreviousPacket( packetID, &secondPacket )
    == true )
{
    // Flag indicating this block is a message.
    outBuffer[ count ] = MTUDPMSGTYPE_RELIABLE;
    count++;

    // Append the message
    memcpy( &outBuffer[ count ], &secondPacket.d_id, sizeof( DWORD ) );
    count += sizeof( DWORD );
    memcpy( &outBuffer[ count ],
            &secondPacket.d_length,
            sizeof( unsigned short ) );
    count += sizeof( unsigned short );
    memcpy( &outBuffer[ count ],
            secondPacket.d_data, secondPacket.d_length );
    count += secondPacket.d_length;
}

#ifdef _DEBUG_DROPTTEST
    if( rand() % _DEBUG_DROPTTEST != _DEBUG_DROPTTEST - 1 )
    {
#endif
        // Send
        SOCKADDR_IN remoteAddr;
        unsigned short result;

        memset( &remoteAddr, 0, sizeof( SOCKADDR_IN ) );
        remoteAddr.sin_family = AF_INET;
        remoteAddr.sin_addr.s_addr = htonl( pHost->GetAddress() );
        remoteAddr.sin_port = htons( pHost->GetPort() );

        // Send the data.
        result = sendto( d_sendSocket,
                        outBuffer,
                        count,
                        0,
                        (SOCKADDR *)&remoteAddr,
                        sizeof( SOCKADDR ) );

        if( result < count )

```

```

        // turn off the mutex and throw an error – could not send all data.
        if( result == SOCKET_ERROR )
            // turn off the mutex and throw an error – sendto() failed.

#ifdef _DEBUG_DROPTEST
    }
#endif

    if( pStr )
        pHost->GetOutQueue().AddPacket( pStr, length );

    cMonitor::MutexOff();
}

```

Since I've covered most of this before, there are only four new and interesting things.

The first is `_DEBUG_DROPTEST`. This function will cause a random packet to not be sent, which is equivalent to playing on a really bad network. If your game can still play on a LAN with a `_DEBUG_DROPTEST` as high as four, then you have done a really good job, because that's more than you would ever see in a real game.

The second new thing is `sendto()`. I think any logically minded person can look at the `bind()` code, look at the clearly named variables, and understand how `sendto()` works.

It may surprise you to see that the mutex is held for so long, directly contradicting what I said earlier. As you can see, `pHost` is still being used on the next-to-last line of the program, so the mutex has to be held in case the other thread calls `MTUDP::HostDestroy()`. Of course, the only reason it has to be held so long is because of `HostDestroy()`.

The third new thing is `MTUDPMSGTYPE_RELIABLE`. I'll get to that a little later.

The last and most important new item is `cHost::GetOutQueue()`. Just like its counterpart, `GetOutQueue()` provides access to an instance of `cQueueOut`, which is remarkably similar (but not identical) to `cQueueIn`.

```

class cQueueOut : public cMonitor
{
protected:
    list<cDataPacket *> d_packets;
    DWORD              d_currentPacketID,
                      d_count; // number of packets added to this queue.

public:
    cQueueOut();
    virtual ~cQueueOut();

    void        Clear();
    void        AddPacket( const char *const pData, unsigned short len );
    void        RemovePacket( DWORD packetID );
    bool        GetPacketForResend( DWORD waitTime, cDataPacket *pPacket );
    bool        GetPreviousPacket( DWORD packetID, cDataPacket *pPacket );
    cDataPacket *BorrowPacket( DWORD packetID );

```

```

void        ReturnPacket();
DWORD       GetLowestID();
bool        IsEmpty();

inline DWORD GetCurrentID(); // returns d_currentPacketID.
inline DWORD GetCount();    // returns d_count.
};

```

There are several crucial differences between `cQueueIn` and `cQueueOut`: `d_currentPacketID` is the ID of the last packet sent/added to the queue; `GetLowestID()` returns the ID of the first packet in the list (which, incidentally, would also be the packet that has been in the list the longest); `AddPacket()` just adds a packet to the far end of the list and assigns it the next `d_currentPacketID`; and `RemovePacket()` removes the packet with `d_id == packetID`.

The four new functions are `GetPacketForResend()`, `GetPreviousPacket()`, `BorrowPacket()`, and `ReturnPacket()`, of which the first two require a brief overview and the last two require a big warning. `GetPacketForResend()` checks if there are any packets that were last sent more than `waitTime` milliseconds ago. If there are, it copies that packet to `pPacket` and updates the original packet's `d_lastTime`. This way, if you know the ping to some other computer, then you know how long to wait before you can assume the packet was dropped. `GetPreviousPacket()` is far simpler; it returns the packet that was sent just before the packet with `d_id == packetID`. This is used by `ReliableSendTo()` to “piggyback” an old packet with a new one in the hopes that it will reduce the number of resends caused by packet drops.

`BorrowPacket()` and `ReturnPacket()` are evil incarnate. I say this because they really, really bend the unwritten mutex rule: Lock and release a mutex in the same function. I know I should have gotten rid of them, but when you see how they are used in the code (later), I hope you'll agree it was the most straightforward implementation. I put it to you as a challenge to remove them. Nevermore shall I mention the functions-that-cannot-be-named().

The point of `MTUDPMSGTYPE_RELIABLE` is that it is an identifier that would be read by `ProcessIncomingData()`. When `ProcessIncomingData()` sees `MTUDPMSGTYPE_RELIABLE`, it would call `pHost->ProcessIncomingReliable()`. The benefit of doing things this way is that it means I can send other stuff in the same message and piggyback it just like I did with the old messages and `GetPreviousPacket()`. In fact, I could send a message that had all kinds of data and no `MTUDPMSGTYPE_RELIABLE` (madness! utter madness!). Of course, in order to be able to process these different message types I'd better make some improvements, the first of which is to define all the different types.

```

enum eMTUDPMsgType
{
    MTUDPMSGTYPE_ACKS           = 0,
    MTUDPMSGTYPE_RELIABLE      = 1,

```



```

    MTUDPMSGTYPE_UNRELIABLE = 2,
    MTUDPMSGTYPE_CLOCK      = 3,
    MTUDPMSGTYPE_NUMMESSAGES = 4,
};

```

I defined this enum in `MTUDP.cpp` because it's a completely internal matter that no other class should be messing with.

Although you're not going to work with most of these types (just yet), here's a brief overview of what they're for:

- `MTUDPMSGTYPE_CLOCK` is for a really cool clock I'm going to add later. It is pretty neat when you consider that the clock will read almost exactly the same value on all clients and the server. This is a critical feature of real-time games because it makes sure that you can say "this thing happened at this time" and everyone can correctly duplicate the effect.
- `MTUDPMSGTYPE_UNRELIABLE` is an unreliable message. When a computer sends an unreliable message it doesn't expect any kind of confirmation because it isn't very concerned if the message doesn't reach the intended destination. A good example of this would be the update messages in a game—if you're sending 20 messages a second, a packet drop here and a packet drop there is no reason to have a nervous breakdown. That's part of the reason we made `_DEBUG_DROPTEST` in the first place!
- `MTUDPMSGTYPE_ACKS` is vital to reliable message transmission. If my computer sends a reliable message to your computer, I need to get a message back saying "yes, I got that message!" If I don't get that message, then I have to resend it after a certain amount of time (hence `GetPacketForResend()`).

Now, before I start implementing the stuff associated with `eMTUDPMsgType`, let me go back and improve `MTUDP::ProcessIncomingData()`.

```

assert( pHost != NULL );

// Process the header for this packet.
bool      bMessageArrived;
unsigned char code;
char      *ptr;

bMessageArrived = false;
ptr = pData;

while( ptr < pData + length )
{
    code = *ptr;
    ptr++;

    switch( code )
    {
        case MTUDPMSGTYPE_ACKS:

```

```

        // Process any ACKs in the packet.
        ptr += pHost->ProcessIncomingACKs( ptr,
                                           pData + length - ptr,
                                           receiveTime );

        break;
    case MTUDPMSGTYPE_RELIABLE:
        bMessageArrived = true;
        // Process reliable message in the packet.
        ptr += pHost->ProcessIncomingReliable( ptr,
                                              pData + length - ptr,
                                              receiveTime );

        break;
    case MTUDPMSGTYPE_UNRELIABLE:
        // Process UNReliable message in the packet.
        ptr += pHost->ProcessIncomingUnreliable( ptr,
                                              pData + length - ptr,
                                              receiveTime );

        break;
    case MTUDPMSGTYPE_CLOCK:
        ptr += ProcessIncomingClockData( ptr,
                                         pData + length - ptr,
                                         pHost,
                                         receiveTime );

        break;
    default:
        // Turn mutex off, throw an error. Something VERY BAD has happened,
        // probably a write to bad memory (such as to an uninitialized
        // pointer).
        break;
    }
}

cMonitor::MutexOff();

if( bMessageArrived == true )
{
    // Send an ACK immediately. If this machine is the
    // server, also send a timestamp of the server clock.
    ReliableSendTo( NULL, 0, pHost->GetAddress() );
}
}

```

So `ProcessIncomingData()` reads in the message type, then sends the remaining data off to be processed. It repeats this until there's no data left to be processed. At the end, if a new message arrived, it calls `ReliableSendTo()` again. Why? Because I'm going to make more improvements to it!

```

// some code we've seen before
memset( outBuffer, 0, MAX_UDPBUFFERSIZE );

// Attach the ACKs.
if( pHost->GetInQueue().GetCount() != 0 )

```

```

{
    // Flag indicating this block is a set of ACKs.
    outBuffer[ count ] = MTUDPMSGTYPE_ACKS;
    count++;

    count += pHost->AddACKMessage( &outBuffer[ count ], MAX_UDPBUFFERSIZE );
}

count += AddClockData( &outBuffer[ count ],
                      MAX_UDPBUFFERSIZE - count,
                      pHost );

// some code we've seen before.

```

So now it is sending clock data, ACK messages, and as many as two reliable packets in every message sent out. Unfortunately, there are now a number of outstanding issues:

- `ProcessIncomingUnreliable()` is all well and good, but how do you send unreliable data?
- How do `cHost::AddACKMessage()` and `cHost::ProcessingIncomingACKs()` work?
- Okay, so I ACK the messages. But you said I should only resend packets if I haven't received an ACK within a few milliseconds of the ping to that computer. So how do I calculate ping?
- How do `AddClockData()` and `ProcessIncomingClockData()` work?

Unfortunately, most of those questions have answers that overlap, so I apologize in advance if things get a little confusing.

Remember how I said there were four more classes to be defined? The class `cQueueOut` was one and here come two more.

cUnreliableQueueIn

```

class cUnreliableQueueIn : public cMonitor
{
    list<cDataPacket *> d_packets;
    DWORD               d_currentPacketID;

public:
    cUnreliableQueueIn();
    virtual ~cUnreliableQueueIn();

    void          Clear();
    void          AddPacket( DWORD packetID,
                           const char *const pData,
                           unsigned short len,
                           DWORD receiveTime );

    cDataPacket *GetPacket();
};
cUnreliableQueueOut

```

```

class cUnreliableQueueOut : public cMonitor
{
    list<cDataPacket *> d_packets;
    DWORD              d_currentPacketID;
    unsigned char      d_maxPackets,
                      d_numPackets;

public:
    cUnreliableQueueOut();
    virtual ~cUnreliableQueueOut();

    void Clear();
    void AddPacket( const char *const pData, unsigned short len );
    bool GetPreviousPacket( DWORD packetID, cDataPacket *pPacket );
    void SetMaxPackets( unsigned char maxPackets );

    inline DWORD GetCurrentID(); // returns d_currentPacketID.
};

```

They certainly share a lot of traits with their reliable counterparts. The two differences are that I don't want to hang on to a huge number of outgoing packets, and I only have to sort incoming packets into one list. In fact, my unreliable packet sorting is really lazy—if the packets don't arrive in the right order, the packet with the lower ID gets deleted. As you can see, `cQueueOut` has a function called `SetMaxPackets()` so you can control how many packets are queued. Frankly, you'd only ever set it to 0, 1, or 2.

Now that that's been explained, let's look at `MTUDP::UnreliableSendTo()`. `UnreliableSendTo()` is almost identical to `ReliableSendTo()`. The only two differences are that unreliable queues are used instead of the reliable ones and the previous packet (if any) is put into the `outBuffer` first, followed by the new packet. This is done so that if packet n is dropped, when packet n arrives with packet $n + 1$, my lazy packet queuing won't destroy packet n .

`cHost::AddACKMessage()/cHost::ProcessIncomingACKs()`

Aside from these two functions, there's a few other things that have to be added to `cHost` with regard to ACKs.

```

#define ACK_MAXPERMSG      256
#define ACK_BUFFERLENGTH  48

class cHost : public cMonitor
{
protected:
    // A buffer of the latest ACK message for this host
    char      d_ackBuffer[ ACK_BUFFERLENGTH ];

    unsigned short d_ackLength; // amount of the buffer actually used.

    void ACKPacket( DWORD packetID, DWORD receiveTime );

public:

```

```

unsigned short ProcessIncomingACKs( char *const pBuffer,
                                   unsigned short len,
                                   DWORD receiveTime );
unsigned short AddACKMessage( char *const pBuffer, unsigned short
                             maxLen );
}

```

The idea here is that I'll probably be sending more ACKs than receiving packets, so it only makes sense to save time by generating the ACK message when required and then using a cut and paste. In fact, that's what `AddACKMessage()` does—it copies `d_ackLength` bytes of `d_ackBuffer` into `pBuffer`. The actual ACK message is generated at the end of `cHost::ProcessIncomingReliable()`. Now you'll finally learn what `cQueueIn::GetCount()`, `cQueueIn::GetHighestID()`, `cQueueIn::GetCurrentID()`, and `cQueueIn::UnorderedPacketIsQueued()` are for.

```

// some code we've seen before.
d_inQueue.AddPacket( packetID, (char *)readPtr, length, receiveTime );
readPtr += length;

// Should we build an ACK message?
if( d_inQueue.GetCount() == 0 )
    return ( readPtr - pBuffer );

// Build the new ACK message.
DWORD lowest, highest, ackID;
unsigned char mask, *ptr;

lowest = d_inQueue.GetCurrentID();
highest = d_inQueue.GetHighestID();

// Cap the highest so as not to overflow the ACK buffer
// (or spend too much time building ACK messages).
if( highest > lowest + ACK_MAXPERMSG )
    highest = lowest + ACK_MAXPERMSG;

ptr = (unsigned char *)d_ackBuffer;
// Send the base packet ID, which is the
// ID of the last ordered packet received.
memcpy( ptr, &lowest, sizeof( DWORD ) );
ptr += sizeof( DWORD );
// Add the number of additional ACKs.
*ptr = highest - lowest;
ptr++;

ackID = lowest;
mask = 0x80;

while( ackID < highest )
{
    if( mask == 0 )
    {
        mask = 0x80;
    }
}

```

```

        ptr++;
    }

    // Is there a packet with id 'i' ?
    if( d_inQueue.UnorderedPacketIsQueued( ackID ) == true )
        *ptr |= mask;    // There is
    else
        *ptr &= ~mask;    // There isn't

    mask >>= 1;
    ackID++;
}

// Record the amount of the ackBuffer used.
d_ackLength = ( ptr - (unsigned char *)d_ackBuffer ) + ( mask != 0 );

// return the number of bytes read from
return readPtr - pBuffer;
}

```

For those of you who don't dream in binary, here's how it works. First of all, you know the number of reliable packets that have arrived in the correct order. So telling the other computer about all the packets that have arrived since last time that are below that number is just a waste of bandwidth. For the rest of the packets, I could have sent the IDs of every packet that has been received (or not received), but think about it: Each ID requires 4 bytes, so storing, say, 64 IDs would take 256 bytes! Fortunately, I can show you a handy trick:

```

// pretend ackBuffer is actually 48 * 8 BITS long instead of 48 BYTES.
for( j = 0; j < highest - lowest; j++ )
{
    if( d_inQueue.UnorderedPacketIsQueued( j + lowest ) == true )
        ackBuffer[ j ] == 1;
    else
        ackBuffer[ j ] == 0;
}

```

Even if you used a whole character to store a 1 or a 0, you'd still be using one-fourth the amount of space. As it is, you could store those original 64 IDs in 8 bytes, eight times less than originally planned.

The next important step is `cHost::ProcessIncomingACKs()`. I think you get the idea—read in the first `DWORD` and ACK every packet with a lower ID that's still in `d_queueOut`. Then go one bit at a time through the rest of the ACKs (if any) and if a bit is 1, ACK the corresponding packet. So I guess the only thing left to show is how to calculate the ping using the ACK information.

```

void cHost::ACKPacket( DWORD packetID, DWORD receiveTime )
{
    cDataPacket *pPacket;

    pPacket = d_outQueue.BorrowPacket( packetID );
    if( pPacket == NULL )

```

```

        return; // the mutex was not locked.

    DWORD time;

    time = receiveTime - pPacket->d_firstTime;
    d_outQueue.ReturnPacket();

    unsigned int i;

    if( pPacket->d_timesSent == 1 )
    {
        for( i = 0; i < PING_RECORDLENGTH - 1; i++ )
            d_pingLink[ i ] = d_pingLink[ i + 1 ];
        d_pingLink[ i ] = time;
    }
    for( i = 0; i < PING_RECORDLENGTH - 1; i++ )
        d_pingTrans[ i ] = d_pingTrans[ i + 1 ];
    d_pingTrans[ i ] = time;

    d_outQueue.RemovePacket( packetID );
}

```

If you take a good look at `cHost::ACKPacket()` you'll notice the only line that actually does anything to ACK the packet is the last one! Everything else helps with the next outstanding issue: ping calculation.

There are two kinds of ping: link ping and transmission latency ping. *Link ping* is the shortest possible time it takes a message to go from one computer and back, the kind of ping you would get from using a ping utility (open a DOS box, type “ping [some address]” and see for yourself). *Transmission latency ping* is the time it takes two programs to respond to each other. In this case, it's the average time that it takes a reliably sent packet to be ACKed, including all the attempts to resend it.

In order to calculate ping for each `cHost`, the following has to be added:

```

#define PING_RECORDLENGTH      64
#define PING_DEFAULTVALLINK    150
#define PING_DEFAULTVALTRANS   200

class cHost : public cMonitor
{
protected:
    // Ping records
    DWORD d_pingLink[ PING_RECORDLENGTH ],
          d_pingTrans[ PING_RECORDLENGTH ];

public:
    float GetAverageLinkPing( float percent );
    float GetAverageTransPing( float percent );
}

```

As packets come in and are ACKed, their round-trip time is calculated and stored in the appropriate ping record (as previously described). Of course,

the two ping records need to be initialized and that's what PING_DEFAULTVALLINK and PING_DEFAULTVALTRANS are for. This is done only once, when cHost is created. Picking good initial values is important for those first few seconds before a lot of messages have been transmitted back and forth. Too high or too low and GetAverage...Ping() will be wrong, which could temporarily mess things up.

Since both average ping calculators are the same (only using different lists), I'll only show the first, GetAverageLinkPing(). Remember how in the cThread class I showed you a little cheat with cThreadProc()? I'm going to do something like that again.

```
// This is defined at the start of cHost.cpp for qsort.
static int sSortPing( const void *arg1, const void *arg2 )
{
    if( *(DWORD *)arg1 < *(DWORD *)arg2 )
        return -1;
    if( *(DWORD *)arg1 > *(DWORD *)arg2 )
        return 1;
    return 0;
}

float cHost::GetAverageLinkPing( float bestPercentage )
{
    if( bestPercentage <= 0.0f ||
        bestPercentage > 100.0f )
        bestPercentage = 100.0f;

    DWORD pings[ PING_RECORDLENGTH ];
    float sum, worstFloat;
    int    worst, i;

    // Recalculate the ping list
    memcpy( pings, &d_pingLink, PING_RECORDLENGTH * sizeof( DWORD ) );
    qsort( pings, PING_RECORDLENGTH, sizeof( DWORD ), sSortPing );

    // Average the first bestPercentage / 100.
    worstFloat = (float)PING_RECORDLENGTH * bestPercentage / 100.0f;
    worst = (int)worstFloat + ( ( worstFloat - (int)worstFloat ) != 0 );
    sum = 0.0f;
    for( i = 0; i < worst; i++ )
        sum += pings[ i ];

    return sum / (float)worst;
}
```

The beauty of this seemingly overcomplicated system is that you can get an average of the best n percent of the pings. Want an average ping that ignores the three or four worst cases? Get the best 80%. Want super accurate best times? Get 30% or less. In fact, those super accurate link ping times will be vital when I answer the fourth question: How do AddClockData() and ProcessIncomingClockData() work?

cNetClock

There's only one class left to define and here it is.

```
class cNetClock : public cMonitor
{
protected:
    struct cTimePair
    {
    public:
        DWORD d_actual, // The actual time as reported by GetTickCount()
              d_clock;   // The clock time as determined by the server.
    };

    cTimePair d_start,      // The first time set by the server.
              d_lastUpdate; // the last updated time set by the server.
    bool      d_bInitialized; // first time has been received.

public:
    cNetClock();
    virtual ~cNetClock();

    void Init();
    void Synchronize( DWORD serverTime,
                     DWORD packetSendTime,
                     DWORD packetACKTime,
                     float ping );

    DWORD GetTime() const;
    DWORD TranslateTime( DWORD time ) const;
};
```

The class `cTimePair` consists of two values: `d_actual` (which is the time returned by the local clock) and `d_clock` (which is the estimated server clock time). The value `d_start` is the clock value the first time it is calculated and `d_lastUpdate` is the most recent clock value. Why keep both? Although I haven't written it here in the book, I was running an experiment to see if you could determine the rate at which the local clock and the server clock would drift apart and then compensate for that drift.

Anyhow, about the other methods. `GetTime()` returns the current server clock time. `TranslateTime()` will take a local time value and convert it to server clock time. `Init()` will set up the initial values. That just leaves `Synchronize()`.

```
void cNetClock::Synchronize( DWORD serverTime,
                             DWORD packetSendTime,
                             DWORD packetACKTime,
                             float ping )
{
    cMonitor::MutexOn();

    DWORD dt;

    dt = packetACKTime - packetSendTime;
```

```

if( dt > 10000 )
    // This synch attempt is too old. Release mutex and return now.

if( d_bInitialized == true )
{
    // If the packet ACK time was too long OR the clock is close enough
    // then do not update the clock.
    if( abs( serverTime + ( dt / 2 ) - GetTime() ) <= 5 )
        // The clock is already very synched. Release mutex and return now.

    d_lastUpdate.d_actual = packetACKTime;
    d_lastUpdate.d_clock = serverTime + (DWORD)( ping / 2 );
    d_ratio = (double)( d_lastUpdate.d_clock - d_start.d_clock ) /
        (double)( d_lastUpdate.d_actual - d_start.d_actual );
}
else // d_bInitialized == false
{
    d_lastUpdate.d_actual = packetACKTime;
    d_lastUpdate.d_clock = serverTime + ( dt / 2 );
    d_start.d_actual = d_lastUpdate.d_actual;
    d_start.d_clock = d_lastUpdate.d_clock;
    d_bInitialized = true;
}

cMonitor::MutexOff();
}

```

As you can see, Synchronize() requires three values: serverTime, packetSendTime, and packetACKTime. Two of the values seem to make good sense—the time a packet was sent out and the time that packet was ACKed. But how does serverTime fit into the picture? For that, I have to add more code to MTUDP:

```

class MTUDP : public cThread
{
protected:
    bool    d_bIsServerOn,
            d_bIsClientOn;
    cNetClock d_clock;

    unsigned short  AddClockData( char *const pData,
                                   unsigned short maxlen,
                                   cHost *const pHost );
    unsigned short  ProcessIncomingClockData( char *const pData,
                                                unsigned short len,
                                                cHost *const pHost,
                                                DWORD receiveTime );

public:
    void  StartServer();
    void  StopServer();
    void  StartClient();
    void  StopClient();

    // GetClock returns d_clock and returns a const ptr so

```

```

        // that no one can call Synchronize and screw things up.
        inline const cNetClock &GetClock();
    }

```

All the client/server stuff you see here is required for the clock and only for the clock. In essence, what it does is tell MTUDP who is in charge and has the final say about what the clock should read. When a client calls `AddClockData()` it sends the current time local to that client, not the server time according to the client. When the server receives a clock time from a client, it stores that time in `cHost`. When a message is going to be sent back to the client, the server sends the last clock time it got from the client and the current server time. When the client gets a clock update from the server it now has three values: the time the message was originally sent (`packetSendTime`), the server time when a response was given (`serverTime`), and the current local time (`packetACKTime`). Based on these three values, the current server time should be approximately `cNetClock::d_lastUpdate.d_clock = serverTime + (packetACKTime - packetSendTime) / 2`.

Of course, you'd only do this if the total round-trip was extremely close to the actual ping time because it's the only way to minimize the difference between client net clock time and server net clock time.

As I said, the last client time has to be stored in `cHost`. That means one final addition to `cHost`.

```

class cHost : public cMonitor
{
protected:
    // For clock synchronization
    DWORD d_lastClockTime;
    bool d_bClockTimeSet;
public:
    DWORD GetLastClockTime();           // self-explanatory.
    void SetLastClockTime( DWORD time ); // self-explanatory.

    inline bool WasClockTimeSet();       // returns d_bClockTimeSet.
}

```

And that appears to be that. In just under 35 pages I've shown you how to set up all the harder parts of network game programming. In the next section I'll show you how to use the MTUDP class to achieve first-rate, super-smooth game play.

Implementation 2: Smooth Network Play

Fortunately, this section is a lot shorter. Unfortunately, this section has no code because the solution for any one game probably wouldn't work for another game.

Geographic and Temporal Independence

Although in this book I am going to write a real-time, networked game, it is important to note the other types of network games and how they affect the inner workings. The major differences can be categorized in two ways: the time separation and the player separation, more formally referred to as geographic independence and temporal independence.

Geographic independence means separation between players. A best-case example would be a two-player *Tetris* game where the players' game boards are displayed side by side. There doesn't have to be a lot of accuracy because the two will never interact. A worst-case example would be a crowded room in *Quake*—everybody's shooting, everybody's moving, and it's very hard to keep everybody nicely synched. This is why in a heavy firefight the latency climbs; the server has to send out a lot more information to a lot more people.

Temporal independence is the separation between events. A best-case example would be a turn-based game such as chess. I can't move a piece until you've moved a piece and I can take as long as I want to think about the next move, so there's plenty of time to make sure that each player sees exactly the same thing. Again, the worst-case scenario is *Quake*—everybody's moving as fast as they can, and if you don't keep up you lag and die.

It's important when designing your game to take the types of independence into consideration because it can greatly alter the way you code the inner workings. In a chess game I would only use `MTUDP::ReliableSendTo()`, because every move has to be told to the other player and it doesn't matter how long it takes until he gets the packet; he'll believe I'm still thinking about my move. In a *Tetris* game I might use `ReliableSendTo()` to tell the other player what new piece has appeared at the top of the wall, where the pieces land, and other important messages like "the other player has lost." The in-between part while the player is twisting and turning isn't really all that important, so maybe I would send that information using `MTUDP::UnreliableSendTo()`. That way they look like they're doing something and I can still guarantee that the final version of each player's wall is correctly imitated on the other player's computer.

Real-time games, however, are a far more complicated story. The login and logout are, of course, sent with `Reliable...()`. But so are any name, model, team, color, shoe size, decal changes, votes, chat messages—the list goes on and on. In a game, however, updates about the player's position are sent 20 times a second and they are sent unreliably. Why? At 20 times a second a player can do a lot of fancy dancin' and it will be (reasonably) duplicated on the other computers. But because there are so many updates being sent, you don't really care if one or two get lost—it's no reason to throw yourself off a bridge. If, however, you were sending all the updates with `Reliable...()`, the slightest hiccup in the network would start a chain

reaction of backlogged reliable messages that would very quickly ruin the game.

While all these updates are being sent unreliably, important events like shooting a rocket, colliding with another player, opening a door, or a player death are all sent reliably. The reason for this is because a rocket blast could kill somebody, and if you don't get the message, you would still see the player standing there. Another possibility is that you don't know the rocket was fired, so you'd be walking along and suddenly ("argh!") you'd die for no reason.

Timing Is Everything

The next challenge you'll face is a simple problem with a complicated solution. The client and the server are sending messages to each other at roughly 50 millisecond intervals. Unfortunately, tests will show that over most connections the receiver will get a "burst" of packets followed by a period of silence followed by another burst. This means you definitely cannot assume that packets arrive exactly 50ms apart—you can't even begin to assume when they were first sent. (If you were trying, cut it out!)

The solution comes from our synchronized network clock.

```
cGame::SendUpdate()
{
    if( time to send another update )
    {
        update.SetTime( d_MTUDPInstance.GetClock().GetTime() );
        update.SetPlayerData( pPlayer->ID(), pPlayer->Pos(), pPlayer->Vel() );
        d_MTUDPInstance.UnreliableSendTo( update.Buffer(),
                                          update.BufferLength(),
                                          someHostID );
    }
}

cGame::ProcessIncomingUpdate( anUpdate )
{
    currentTime = d_MTUDPInstance.GetClock().GetTime();
    eventTime = anUpdate.GetTime();

    updatePos = anUpdate.GetPos();
    updateVel = anUpdate.GetVelocity();

    newPos = updatePos + updateVel * ( currentTime - eventTime );
    pPlayer[ playerId ].SetPos( newPos );
}
```

The above case would only work if people moved in a straight line. Since most games don't, you also have to take into account their turning speed, physics, whether they are jumping, etc.

In case it wasn't clear yet, let me make it perfectly crystal: *Latency is public enemy #1*. Of course, getting players to appear isn't the only problem.

Pick and Choose

Reducing the amount of data is another important aspect of network programming. The question to keep in mind when determining what to send is: “What is the bare minimum I have to send to keep the other computer(s) up to date?” For example, in a game like *Quake* there are a lot of ambient noises. Water flowing, lava burbling, moaning voices, wind, and so on. Not one of these effects is an instruction from the server. Why? Because none of these sounds are critical to keeping the game going. In fact, none of the sounds are. Not that it makes any difference, because you can get all your “play this sound” type messages for free.

Every time a sound is played, it’s because something happened. When something happens, it has to be duplicated on every computer. This means that every sound event is implicit in some other kind of event. If your computer gets a message saying “a door opened,” then your machine knows it has to open the door and play the door open sound.

Another good question to keep in mind is “how can I send the same information with less data?” A perfect example is the ACK system. Remember how I used 1 bit per packet and ended up using one-eighth the amount of data? Then consider what happens if, instead of saying “player x is turning left and moving forward” you use 1-bit flags. It only takes 2 bits to indicate left, right, or no turning and the same goes for walking forward/back or left/right. A few more 1-bit flags that mean things like “I am shooting,” “I am reloading,” or “I am shaving my bikini zone,” and you’ve got everything you need to duplicate the events of one computer on another. Another good example of reducing data comes in the form of a parametric movement. Take a rocket, for example. It flies in a nice straight line, so you only have to send the message “a rocket has been fired from position X with velocity Y at time Z” and the other computer can calculate its trajectory from there.

Prediction and Extrapolation

Of course, it’s not just as simple as processing the messages as they arrive. The game has to keep moving things around whether or not it’s getting messages from the other computer(s) for as long as it can. That means that everything in the game has to be predictable: All players of type Y carrying gun X move at speed Z. Without constants like that, the game on one machine would quickly become different from that on other machines and everything would get very annoying. But there’s more to it, and that “more” is a latency related problem.



Note: This is one of the few places where things start to differ between the client and server, so please bear with me.

The server isn't just the final authority on the clock time, it's also the final authority on every single player movement or world event (such as doors and elevators). That means it also has to shoulder a big burden. Imagine that there's a latency of 100 milliseconds between client and server. On the server, a player gets hit with a rocket and dies. The server builds a message and sends it to the client. From the time the server sends the message until the client gets the message the two games are not synchronized. It may not sound like much but it's the culmination of all these little things that make a great game terrible—or fantastic, if they're solved. In this case, the server could try predicting to see where everyone and everything will be n milliseconds from now and send messages that say things like “if this player gets hit by that rocket he'll die.” The client will get the message just in time and no one will be the wiser. In order to predict where everyone will be n milliseconds from now, the server must first extrapolate the players' current position based on the last update sent from the clients. In other words, the server uses the last update from a client and moves the player based on that information every frame. It then uses this new position to predict where the player is *going to be* and then it can tell clients “player X will be at position Y at time Z.” In order to make the game run its smoothest for all clients, the amount of time to predict ahead should be equal to half the client's transmission ping. Of course, this means recalculating the predictions for every player, but it's a small price to pay for super-smooth game play.

The clients, on the other hand, should be receiving “player X will be at position Y at time Z” just about the same moment the clock reaches time Z. You would think that the client could just start extrapolating based on that info, right? Wrong. Although both the clients and the server are showing almost exactly the same thing, the clients have one small problem, illustrated in this example: If a client shoots at a moving target, that target will not be there by the time the message gets to the server. Woe! Sufferance! What to do? Well, the answer is to *predict* where everything will be n milliseconds from now. What is n ? If you guessed half the transmission ping, you guessed right.

You're probably wondering why one is called prediction and the other is extrapolation. When the server is extrapolating, it's using old data to find the current player positions. When a client is predicting, it's using current data to extrapolate future player positions.

Using `cHost::GetAverageTransPing(50.0f)` to get half the transmission ping is not the answer. Using `cHost::GetAverageTransPing(80.0f)/2` would work a lot better. Why? By taking 80% of the transmission pings you can ignore a few of the worst cases where a packet was dropped (maybe even dropped twice!), and since ping is the round-trip time you have to divide it by two.

Although predicting helps to get the messages to the server on time, it doesn't help to solve the last problem—what happens if a prediction is wrong? The players on screen would “teleport” to new locations without

crossing the intermediate distance. It could also mean that a client thinks someone got hit by a rocket when in fact on the server he dodged at just the last second.

The rocket-dodging problem is the easier problem to solve, so I'll tackle it first. Because the server has the final say in everything, the client should perform collision detection as it always would: Let the rocket blow up, spill some blood pixels around the room, and then do nothing to the player until it got a message from the server saying "player X has definitely been hit and player X's new health is Y." Until that message is received, all the animations performed around/with the player should be as non-interfering and superficial as a sound effect. All of which raises an important point: Both the client and the server perform collision detection, but it's the server that decides who lives and who dies.

As for the teleport issue, well, it's a bit trickier. Let's say you are watching somebody whose predicted position is (0,0) and they are running (1,0). Suddenly your client gets an update that says the player's new predicted position is (2,0) running (0,1). Instead of teleporting that player and suddenly turning him, why not interpolate the difference? By that I mean the player would very (very) quickly move from (0,0) to somewhere around (2,0.1) and make a fast turn to the left. Naturally, this can only be done if the updates come within, say, 75 milliseconds of each other. Anything more and you'd have to teleport the players or they might start clipping through walls.

And last but not least, there are times when a real network can suddenly go nuts and lag for as much as 30 seconds. In cases where the last message from a computer was more than two seconds ago, I would freeze all motion and try to get the other machine talking again. If the computer does eventually respond, the best solution for the server would be to send a special update saying where everything is in the game right now and let the client start predicting from scratch. If there's still no response after 15 seconds I would disconnect that other computer from the game (or disconnect myself, if I'm a client).

Conclusion

In this chapter I've divulged almost everything I know about multithreading and network game programming. Well, except for my biggest secrets! There's only two things left to make note of.

First, if `MTUDP::ProcessIncomingData()` is screaming its head off because there's an invalid message type (i.e., the byte read does not equal one of the `eMTUDPMsgType`), then it means that somewhere in the rest of your program you are writing to funny memory, such as writing beyond the bounds of an array or trying to do something funny with an uninitialized pointer.

Second, do not try to add network support to a game that has already been written because it will drive you insane. Try it this way—when most

people start writing an engine, they begin with some graphics, then add keyboard or mouse support because graphics are more important and without graphics, the keyboard and mouse are useless. The network controls a lot of things about how the graphics will appear, which means that the network is more important than the graphics, and thus creating the network should be a priority.

I am sure you will have endless fun with the network topics I have discussed here as long as you incorporate them from the beginning.

Chapter 7

Direct3D Fundamentals

I remember when I was young and went through the rite of passage of learning to ride a bicycle. It wasn't pretty. At first, I was simply terrified of getting near the thing. I figured my own two feet were good enough. Personally, I felt the added speed and features of a bike weren't worth the learning curve. I would attempt to sit on my bike, only to have it violently buck me over its shoulders like some vicious bull at a rodeo. The balance I needed, the speed control, the turning-while-braking—it was almost too much. Every ten minutes, I would burst into my house, looking for my mom so she could bandage up my newly skinned knees. It took a while, but eventually the vicious spirit of the bike was broken and I was able to ride around. Once I got used to it, I wondered why it took me so long to get the hang of it. Once I got over the hump of the learning curve, the rest was smooth sailing.

And with that, I delve into something quite similar to learning to ride a bicycle. Something that initially is hard to grasp, something that may scrape your knees a few times (maybe as deep as the arteries), but something that is worth learning and, once you get used to it, pretty painless: Direct3D programming.

In this chapter we're going to cover:

- The Direct3D 10 device
- How to create a viewport
- Introduction to using shaders and creating a default HLSL shader
- Depth buffers
- Stencil buffers
- Lighting techniques
- Loading 3D models

Introduction to D3D

There is only one major interface that is all-important in Direct3D: the Direct3D device. You came across this peripherally in Chapter 2. Note that if you have experience with previous versions of Direct3D, in version 10 the Direct3D object is no longer used.

The Direct3D device will become the center of your 3D universe. Just about all of the work you do in Direct3D goes through the device. Depending on how you want to work, you can create a specific device type, such

as a normal hardware device that takes full control of your Direct3D 10 class accelerator, or a software reference device that is slow but easier to debug.



Note: This is the first time I've had to really worry about the concept of rasterization, so it makes sense to at least define the term. *Rasterization* is the process of taking a graphic primitive (such as a triangle) and actually rendering it pixel by pixel to the screen. It's an extremely complex (and interesting) facet of computer graphics programming; you're missing out if you've never tried to write your own texture mapper from scratch!

You'll use the device for everything: setting textures, setting render targets, drawing triangles, setting up shaders, and so on. It is your mode of communication with the hardware on the user's machine and you'll use it constantly. Learn the interface, and love it.

Many of the concepts I talked about in Chapter 2 will be put into use here. It's no coincidence that the same types of lights I discussed are the same ones Direct3D supports. In order to grasp the practical concepts of Direct3D, I needed to first show you the essentials of 3D programming. With that in your back pocket you can start exploring the concepts that drive Direct3D programming.

Getting Started with Direct3D

We took a brief look at how to set up Direct3D in Chapter 2. However, for full 3D work we need to go through quite a few more steps to get it all set up. This can get pretty confusing, so first we'll look at an overview of the process and then go through the code step by step. I've broken the initialization down into four steps, which are:

1. Creating the device and swap chain
2. Creating a depth and stencil buffer
3. Creating a viewport
4. Creating a default shader

Step 1: Creating the ID3D10Device and Swap Chain

This is what we looked at in Chapter 2, so I'll not spend too long talking about it except to recap that we filled out a `DXGI_SWAP_CHAIN_DESC` structure and passed it to the function `D3D10CreateDeviceAndSwapChain()`. This creates a device, render target, and back buffer. Here is the code now in its own function:

```
void cGraphicsLayer::CreateDeviceAndSwapChain()
{
    HRESULT r = 0;

    // Structure to hold the creation parameters for the device
    DXGI_SWAP_CHAIN_DESC descSwap;
```

```

ZeroMemory(&descSwap, sizeof(descSwap));

// Only want one back buffer
descSwap.BufferCount = 1;

// Width and height of the back buffer
descSwap.BufferDesc.Width = m_rcScreenRect.right;
descSwap.BufferDesc.Height = m_rcScreenRect.bottom;

// Standard 32-bit surface type
descSwap.BufferDesc.Format = DXGI_FORMAT_R8G8B8A8_UNORM;

// 60hz refresh rate
descSwap.BufferDesc.RefreshRate.Numerator = 60;
descSwap.BufferDesc.RefreshRate.Denominator = 1;
descSwap.BufferUsage = DXGI_USAGE_RENDER_TARGET_OUTPUT;

// Connect it to our main window
descSwap.OutputWindow = m_hWnd;

// No multisampling
descSwap.SampleDesc.Count = 1;
descSwap.SampleDesc.Quality = 0;

// Windowed mode
descSwap.Windowed = TRUE;

// Create the device using hardware acceleration
r = D3D10CreateDeviceAndSwapChain(
    NULL, // Default adapter
    D3D10_DRIVER_TYPE_HARDWARE, // Hardware accelerated device
    NULL, // Not using a software DLL for rendering
    D3D10_CREATE_DEVICE_DEBUG, // Flag to allow debug output
    D3D10_SDK_VERSION, // Indicates the SDK version being used
    &descSwap,
    &m_pSwapChain,
    &m_pDevice);

if(FAILED(r))
{
    throw cGameError(L"Could not create IDirect3DDevice10");
}

// Get a copy of the pointer to the back buffer
r = m_pSwapChain->GetBuffer(0,
    __uuidof(ID3D10Texture2D), (LPVOID*)&m_pBackBuffer);
if(FAILED(r))
{
    throw cGameError(L"Could not get back buffer");
}

// Create a render target view
r = m_pDevice->CreateRenderTargetView(
    m_pBackBuffer, NULL, &m_pRenderTargetView);

```

```

    if(FAILED(r))
    {
        throw cGameError(L"Could not create render target view");
    }

    r = m_pDevice->QueryInterface(
        __uuidof(ID3D10InfoQueue), (LPVOID*)&m_pMessageQueue);
    if(FAILED(r))
    {
        throw cGameError(
            L"Could not create IDirect3DDevice10 message queue");
    }
    m_pMessageQueue->SetMuteDebugOutput(false);    // No muting
    m_pMessageQueue->SetMessageCountLimit(-1);    // Unlimited messages
}

```

The `cGraphicsLayer::CreateDeviceAndSwapChain()` function is called from `cGraphicsLayer::InitD3D()`.

Step 2: Creating a Depth/Stencil Buffer

We'll look at how to use depth and stencil buffers a little later. For now, though, we need to set one up in order to be able to render properly. What is important to know is that a depth buffer is used to provide pixel accurate depth testing so that when you render an object in front of another object they don't come out all mangled up. Stencil buffers are used for advanced effects like volume shadowing.

To create a depth/stencil buffer we start by creating a 2D texture with the same resolution as our back buffer. We do this by filling out a `D3D10_TEXTURE2D_DESC` structure, which we saw in Chapter 2. This time, however, we will fill it out like this:

```

D3D10_TEXTURE2D_DESC descDepth;
ZeroMemory(&descDepth, sizeof(descDepth));
descDepth.Width = m_rcScreenRect.right;
descDepth.Height = m_rcScreenRect.bottom;
descDepth.MipLevels = 1;
descDepth.ArraySize = 1;
descDepth.Format = DXGI_FORMAT_D24_UNORM_S8_UINT;
descDepth.SampleDesc.Count = 1;
descDepth.SampleDesc.Quality = 0;
descDepth.Usage = D3D10_USAGE_DEFAULT;
descDepth.BindFlags = D3D10_BIND_DEPTH_STENCIL;
descDepth.CPUAccessFlags = 0;
descDepth.MiscFlags = 0;

```

The main points to note are that the width and height are exactly the same size as the back buffer. Also notice the format is set to `DXGI_FORMAT_D24_UNORM_S8_UINT`, which in English means a 32-bit buffer, with 24 bits allocated to the depth buffer and 8 bits allocated to the stencil buffer. The buffer holds unsigned integer data. When the structure is all filled out, we can pass it as a parameter to `ID3D10Device::CreateTexture2D()`, which has the following prototype:

```
HRESULT CreateTexture2D(
    const D3D10_TEXTURE2D_DESC *pDesc,
    const D3D10_SUBRESOURCE_DATA *pInitialData,
    ID3D10Texture2D **ppTexture2D
);
```

The first parameter takes the address of our `D3D10_TEXTURE2D_DESC` structure that we just filled out. The second parameter takes initial data to load the texture with, which we are not interested in and can therefore set to `NULL`. The third parameter takes the address of a pointer to a texture, which will be filled in by Direct3D when the texture is created. Here is the call:

```
r = m_pDevice->CreateTexture2D(&descDepth, NULL, &m_pDepthStencilBuffer);
```

So now we have a texture set up to use as our depth/stencil buffer, but Direct3D doesn't know anything about it yet. The next step is to fill in a `D3D10_DEPTH_STENCIL_DESC` structure, which looks like this:

```
typedef struct D3D10_DEPTH_STENCIL_DESC {
    BOOL DepthEnable;
    D3D10_DEPTH_WRITE_MASK DepthWriteMask;
    D3D10_COMPARISON_FUNC DepthFunc;
    BOOL StencilEnable;
    UINT8 StencilReadMask;
    UINT8 StencilWriteMask;
    D3D10_DEPTH_STENCIL_OP_DESC FrontFace;
    D3D10_DEPTH_STENCIL_OP_DESC BackFace;
} D3D10_DEPTH_STENCIL_DESC;
```

DepthEnable	Flags whether the depth part of the buffer is enabled. We'll set this to true.
DepthWriteMask	This can be set to <code>D3D10_DEPTH_WRITE_MASK_ZERO</code> or <code>D3D10_DEPTH_WRITE_MASK_ALL</code> . We'll be setting it to <code>D3D10_DEPTH_WRITE_MASK_ALL</code> , which means all of the depth buffer is available for writing.
DepthFunc	<p>A function for comparing depth data. This can be set to any one of the <code>D3D10_COMPARISON_FUNC</code> enumerated values:</p> <ul style="list-style-type: none"> • <code>D3D10_COMPARISON_NEVER</code>—Test always fails • <code>D3D10_COMPARISON_LESS</code>—Pass if source data is less than destination data • <code>D3D10_COMPARISON_EQUAL</code>—Pass if source data is equal to destination data • <code>D3D10_COMPARISON_LESS_EQUAL</code>—Pass if source data is less than or equal to destination data • <code>D3D10_COMPARISON_GREATER</code>—Pass if the source data is greater than destination data • <code>D3D10_COMPARISON_NOT_EQUAL</code>—Pass if the source data is not equal to destination data

- `D3D10_COMPARISON_GREATER_EQUAL`—Pass if the source data is greater than or equal to destination data
 - `D3D10_COMPARISON_ALWAYS`—Always pass the test
- We'll be using `D3D10_COMPARISON_LESS`.

StencilEnable	Flags whether the stencil buffer portion of the depth stencil buffer is enabled, which we'll set to true.
StencilReadMask	Sets the portion of the stencil buffer for reading, which we'll set to all, or <code>0xFFFFFFFF</code> .
StencilWriteMask	Sets the portion of the stencil buffer for writing, which we'll set to all, or <code>0xFFFFFFFF</code> .
FrontFace	Sets how the results of the depth test should be used for triangles facing the camera in the stencil buffer. There are three <code>D3D10_STENCIL_OP</code> members, which can be set to test for when the stencil fails, when the stencil passes and depth testing fails, and for when they both pass. Finally, there is a <code>D3D10_COMPARISON_FUNC</code> , which is set to a test to compare data. For now we'll set these to <code>D3D10_STENCIL_OP_KEEP</code> , <code>D3D10_STENCIL_OP_INCR</code> , <code>D3D10_STENCIL_OP_KEEP</code> , and <code>D3D10_COMPARISON_ALWAYS</code> , respectively.
BackFace	This member is the same as <code>FrontFace</code> except it defines what to do for triangles that are not facing the camera. For now we'll set this to <code>D3D10_STENCIL_OP_KEEP</code> , <code>D3D10_STENCIL_OP_DECR</code> , <code>D3D10_STENCIL_OP_KEEP</code> , and <code>D3D10_COMPARISON_ALWAYS</code> , respectively.

Here is the code to fill out the depth stencil buffer:

```
D3D10_DEPTH_STENCIL_DESC descDS;
ZeroMemory(&descDS, sizeof(descDS));
descDS.DepthEnable = true;
descDS.DepthWriteMask = D3D10_DEPTH_WRITE_MASK_ALL;
descDS.DepthFunc = D3D10_COMPARISON_LESS;

// Stencil test values
descDS.StencilEnable = true;
descDS.StencilReadMask = (UINT8)0xFFFFFFFF;
descDS.StencilWriteMask = (UINT8)0xFFFFFFFF;

// Stencil op if pixel is front
descDS.FrontFace.StencilFailOp = D3D10_STENCIL_OP_KEEP;
descDS.FrontFace.StencilDepthFailOp = D3D10_STENCIL_OP_INCR;
descDS.FrontFace.StencilPassOp = D3D10_STENCIL_OP_KEEP;
descDS.FrontFace.StencilFunc = D3D10_COMPARISON_ALWAYS;

// Stencil op if pixel is back
descDS.BackFace.StencilFailOp = D3D10_STENCIL_OP_KEEP;
descDS.BackFace.StencilDepthFailOp = D3D10_STENCIL_OP_DECR;
descDS.BackFace.StencilPassOp = D3D10_STENCIL_OP_KEEP;
descDS.BackFace.StencilFunc = D3D10_COMPARISON_ALWAYS;
```

Now that the structure is all filled out we can use it to create a depth stencil state. For each mode you want the depth stencil buffer to be in you can create a state. For example, you could have one state for normal rendering and another state for volumetric shadowing. Once the states are created, it's a single function call to change the states. This is a huge improvement over Direct3D 9, where you would have to manually set all the states each time. To create a state we call `ID3D10Device::CreateDepthStencilState()`, which has the following prototype:

```
HRESULT CreateDepthStencilState(
    const D3D10_DEPTH_STENCIL_DESC *pDepthStencilDesc,
    ID3D10DepthStencilState **ppDepthStencilState
);
```

The first parameter takes the address of the `D3D10_DEPTH_STENCIL_DESC` structure that we just filled in. The second parameter takes the address of a pointer to an `ID3D10DepthStencilState` interface, which will be filled in by Direct3D. Here is the call I used, which stores the default depth stencil state pointer in the `cGraphicsLayer` class:

```
r = m_pDevice->CreateDepthStencilState(&descDS, &m_pDepthStencilState);
if(FAILED(r))
{
    throw cGameError(L"Could not create depth/stencil state");
}
```

Once the state is created, we need to set it as the current state. To do this, we call the function `ID3D10Device::OMSetDepthStencilState()`. The OM stands for output merger. The function has this prototype:

```
void OMSetDepthStencilState(
    ID3D10DepthStencilState *pDepthStencilState,
    UINT StencilRef
);
```

The first parameter takes a pointer to a state to set, and the second takes a reference value to use for the depth stencil comparison functions. I called the function like this:

```
m_pDevice->OMSetDepthStencilState(m_pDepthStencilState, 1);
```

The final step in the process is to create a view of the depth stencil buffer so that Direct3D knows to render depth and stencil information into our buffer. To do this, we fill out yet another structure called `D3D10_DEPTH_STENCIL_VIEW_DESC`, which looks like this:

```
typedef struct D3D10_DEPTH_STENCIL_VIEW_DESC {
    DXGI_FORMAT Format;
    D3D10_DSV_DIMENSION ViewDimension;
    union {
        D3D10_TEX1D_DSV Texture1D;
        D3D10_TEX1D_ARRAY_DSV Texture1DArray;
        D3D10_TEX2D_DSV Texture2D;
    };
};
```



```

        D3D10_TEX2D_ARRAY_DSV Texture2DArray;
        D3D10_TEX2DMS_DSV Texture2DMS;
        D3D10_TEX2DMS_ARRAY_DSV Texture2DMSArray;
    };
} D3D10_DEPTH_STENCIL_VIEW_DESC;

```

Format	This takes the format of the depth stencil buffer, which must be the same as we set in the call to <code>CreateTexture2D()</code> . Here we use <code>DXGI_FORMAT_D24_UNORM_S8_UINT</code> .
ViewDimension	This sets the dimensions of the texture, or in other words if it has any array slices. It doesn't, so we set it to <code>D3D10_DSV_DIMENSION_TEXTURE2D</code> .
Texture2D	This has one member—the number of mip levels. We don't have any so it's set to 0.

Here is the code to fill it in:

```

D3D10_DEPTH_STENCIL_VIEW_DESC descDSView;
ZeroMemory(&descDSView, sizeof(descDSView));
descDSView.Format = DXGI_FORMAT_D24_UNORM_S8_UINT;
descDSView.ViewDimension = D3D10_DSV_DIMENSION_TEXTURE2D;
descDSView.Texture2D.MipSlice = 0;

```

Finally, we need to call the function `ID3D10Device::CreateDepthStencilView()`, which has the following prototype:

```

HRESULT CreateDepthStencilView(
    ID3D10Resource *pResource,
    const D3D10_DEPTH_STENCIL_VIEW_DESC *pDesc,
    ID3D10DepthStencilView **ppDepthStencilView
);

```

The first parameter takes a pointer to our depth stencil buffer texture. The second parameter takes the address of the `D3D10_DEPTH_STENCIL_VIEW_DESC` structure we just filled in. The third parameter takes the address of a pointer that will be filled in by Direct3D.

Bringing It All Together

And that's all that is required to create a depth stencil buffer. I put all this code together into a single function called `cGraphicsLayer::CreateDepthStencilBuffer()`, which is listed here:

```

void cGraphicsLayer::CreateDepthStencilBuffer()
{
    HRESULT r = 0;
    // Create the depth buffer
    D3D10_TEXTURE2D_DESC descDepth;
    ZeroMemory(&descDepth, sizeof(descDepth));
    descDepth.Width = m_rcScreenRect.right;
    descDepth.Height = m_rcScreenRect.bottom;
    descDepth.MipLevels = 1;
    descDepth.ArraySize = 1;

```

```

descDepth.Format = DXGI_FORMAT_D24_UNORM_S8_UINT;
descDepth.SampleDesc.Count = 1;
descDepth.SampleDesc.Quality = 0;
descDepth.Usage = D3D10_USAGE_DEFAULT;
descDepth.BindFlags = D3D10_BIND_DEPTH_STENCIL;
descDepth.CPUAccessFlags = 0;
descDepth.MiscFlags = 0;
r = m_pDevice->CreateTexture2D(&descDepth, NULL, &m_pDepthStencilBuffer);
if(FAILED(r))
{
    throw cGameError(L"Unable to create depth buffer");
}

D3D10_DEPTH_STENCIL_DESC descDS;
ZeroMemory(&descDS, sizeof(descDS));
descDS.DepthEnable = true;
descDS.DepthWriteMask = D3D10_DEPTH_WRITE_MASK_ALL;
descDS.DepthFunc = D3D10_COMPARISON_LESS;

// Stencil test values
descDS.StencilEnable = true;
descDS.StencilReadMask = (UINT8)0xFFFFFFFF;
descDS.StencilWriteMask = (UINT8)0xFFFFFFFF;

// Stencil op if pixel is front
descDS.FrontFace.StencilFailOp = D3D10_STENCIL_OP_KEEP;
descDS.FrontFace.StencilDepthFailOp = D3D10_STENCIL_OP_INCR;
descDS.FrontFace.StencilPassOp = D3D10_STENCIL_OP_KEEP;
descDS.FrontFace.StencilFunc = D3D10_COMPARISON_ALWAYS;

// Stencil op if pixel is back
descDS.BackFace.StencilFailOp = D3D10_STENCIL_OP_KEEP;
descDS.BackFace.StencilDepthFailOp = D3D10_STENCIL_OP_DECR;
descDS.BackFace.StencilPassOp = D3D10_STENCIL_OP_KEEP;
descDS.BackFace.StencilFunc = D3D10_COMPARISON_ALWAYS;

r = m_pDevice->CreateDepthStencilState(&descDS, &m_pDepthStencilState);
if(FAILED(r))
{
    throw cGameError(L"Could not create depth/stencil state");
}
m_pDevice->OMSetDepthStencilState(m_pDepthStencilState, 1);

D3D10_DEPTH_STENCIL_VIEW_DESC descDSView;
ZeroMemory(&descDSView, sizeof(descDSView));
descDSView.Format = DXGI_FORMAT_D24_UNORM_S8_UINT;
descDSView.ViewDimension = D3D10_DSV_DIMENSION_TEXTURE2D;
descDSView.Texture2D.MipSlice = 0;

r = m_pDevice->CreateDepthStencilView(
    m_pDepthStencilBuffer, &descDSView, &m_pDepthStencilView);
if(FAILED(r))
{
    throw cGameError(L"Could not create depth/stencil view");
}

```

```
    }  
}
```

As with `CreateDeviceAndSwapChain()`, this function is also called from `InitD3D()`. Once the render target and depth stencil buffer are created we can assign them to Direct3D with the function we saw in Chapter 2, `ID3D10Device::OMSetRenderTargets()`. This time the third parameter is set to be the depth stencil buffer.

```
m_pDevice->OMSetRenderTargets(1, &m_pRenderTargetView, m_pDepthStencilView);
```

Now let’s check out the viewport.

Step 3: Creating a Viewport

A viewport defines which area of your back buffer is rendered to. We want to render to the entire buffer; however, you could easily change these settings to render to a different portion of the buffer. The viewport also defines the minimum and maximum depth that will be used for your depth buffer. Setting up the viewport is much easier than creating the depth stencil buffer. The first step is to fill in a `D3D10_VIEWPORT` structure, which looks like this:

```
typedef struct D3D10_VIEWPORT {  
    INT TopLeftX;  
    INT TopLeftY;  
    UINT Width;  
    UINT Height;  
    FLOAT MinDepth;  
    FLOAT MaxDepth;  
} D3D10_VIEWPORT;
```

TopLeftX	The top-left X coordinate at which to start rendering.
TopLeftY	The top-left Y coordinate at which to start rendering.
Width	The width in pixels of the area you want to render to.
Height	The height in pixels of the area you want to render to.
MinDepth	The minimum value for the depth buffer.
MaxDepth	The maximum value for the depth buffer.

Here is the code I used to initialize the structure:

```
D3D10_VIEWPORT vp;  
ZeroMemory(&vp, sizeof(vp));  
vp.Width = m_rcScreenRect.right;  
vp.Height = m_rcScreenRect.bottom;  
vp.MinDepth = 0.0f;  
vp.MaxDepth = 1.0f;  
vp.TopLeftX = 0;  
vp.TopLeftY = 0;
```

Once that is filled in we can tell Direct3D about the viewport using the function `ID3D10Device::RSSetViewports()`. The RS prefix stands for rasterizer stage. The function has the following prototype:

```
void RSSetViewports(
    UINT NumViewports,
    const D3D10_VIEWPORT *pViewports
);
```

The first parameter takes the number of viewports to set. The second parameter takes an array of `D3D10_VIEWPORT` structures. Here is the code I used:

```
m_pDevice->RSSetViewports( 1, &vp );
```

Altogether this code is wrapped in the function `cGraphicsLayer::CreateViewport()`, shown below, which again is called from `InitD3D()`:

```
void cGraphicsLayer::CreateViewport()
{
    // Create a viewport the same size as the back buffer
    D3D10_VIEWPORT vp;
    ZeroMemory(&vp, sizeof(vp));
    vp.Width = m_rcScreenRect.right;
    vp.Height = m_rcScreenRect.bottom;
    vp.MinDepth = 0.0f;
    vp.MaxDepth = 1.0f;
    vp.TopLeftX = 0;
    vp.TopLeftY = 0;
    m_pDevice->RSSetViewports( 1, &vp );
}
```

Step 4: Creating a Default Shader

When you reached a point similar to this with DirectX 9.0, you could just go off and start rendering triangles, etc. That is because DirectX 9.0 contains a fixed function pipeline, which means you can set up almost everything you need to do to render by setting a vast range of flags called render states. Direct3D 10 contains no fixed function pipeline and no render states. Absolutely everything you want to render must be done with HLSL shaders.

Introduction to Shaders

A *shader* is just a name for a program that executes on the GPU instead of the CPU. Shaders are written in a language called HLSL, which is very similar to C++. Shaders were originally written in Assembly, which was very complicated and hard to maintain. These days Assembly coding is not allowed and all shaders must be written in HLSL.

The word *shader* originated from the fact that shaders used to be used only to fill in the final surface properties of triangles such as color, reflection, and so on. These days shaders can contain any code, even physics or

audio calculations that have nothing to do with graphics! Each version of DirectX has supported more and more advanced shaders. The current version, 4.0, is the most advanced at the time of writing and is only supported by DirectX 10.

Shaders are written in HLSL just like any other source file in Visual C++ or any other text editor. They are usually saved in FX (effect) files, with the file extension .fx. Once written, the effect files are loaded and compiled into binary format by Direct3D. Usually you have many different shaders for each object you are rendering in your game. For example, you would have a particular shader for rendering water, another for the sky, and so on. Usually you will also have a default shader that runs when no specialized shaders are needed, which is what we'll be creating in this chapter. The process normally goes like this pseudocode:

```
For each object to render
{
    Tell Direct3D to select the correct shader you've written.
    Pass any data to the shader you need to.
    Render the object using the shader.
}
```

Shaders used to come in two varieties called *vertex* shaders and *pixel* shaders. DirectX 10 adds a third type of shader called a *geometry* shader. Each type of shader runs at a different time and is used for a different purpose. Each type of shader is contained in its own function in the FX file. There can be other functions that the shader calls, but at a minimum each shader will be contained in one function.

A geometry shader executes once for every primitive type you render. A *primitive* is the most basic object you can render and can be a point, line, or triangle. So, for example, if your model contains 1,000 triangles, the geometry shader will execute on the graphics card 1,000 times for that object. Geometry shaders are used for advanced features like fur rendering, particle systems, and so on. They are completely optional and we won't be using them in this chapter.

A vertex shader, as you might guess, executes once for every vertex in your model. So, for example, if we take the same 1,000-triangle object and assume each triangle has three vertices, then the vertex shader would run 3,000 times for the object. Vertex shaders are usually used for positioning and high-level lighting calculations. For example, you transform the position of each vertex from local to world space, and then apply the view and projection transformations. It's very common to compute lighting calculations for each vertex in the vertex shader, which we'll see shortly.

Pixel shaders execute once for every pixel rendered to the screen, no matter how many triangles are involved. So if you have a quad made of two triangles that takes up the entire screen on a 640x480 resolution buffer, then the pixel shader would execute 307,200 times. As pixel shaders execute so often it's important to make them as efficient as possible and keep anything that can execute in higher level shaders (i.e., vertex

or geometry) if possible. The great thing about pixel shaders is that if you work out something in a vertex shader and pass it as a parameter to a pixel shader, the data is automatically interpolated for each pixel the shader operates on. So if in a vertex shader you set the three corners of a triangle to be red, green, and blue and spit out the colors in the pixel shader, you will see a smooth blend from each corner thanks to interpolation. The real benefit to this of course is that you can work out per-vertex lighting in the vertex shader and perform the final calculations in the pixel shader to get per-pixel lighting very efficiently.

Your First HLSL Shader

As I mentioned previously we will not be looking at geometry shaders in this chapter, so this example will only contain vertex and pixel shaders. Shaders are held in an FX file, which must contain answers to the following questions:

- What data will the shaders need from the CPU before they execute?
- What kind of data will the vertex shader accept as input and produce as output?
- What data will the pixel shader be expecting?
- If there is more than one shader in the file, which ones should Direct3D use?

All of these questions are very easy to answer. First, let's think about what data the shaders will need. The vertex shader will need per-vertex data such as position, normal, color, etc. However, it will also need some general data, such as the world->view->projection transformation. So the first thing we need to do is to define some global variables to hold the transformation matrices:

```
matrix g_mtxWorld;
matrix g_mtxView;
matrix g_mtxProj;
```

HLSL supports the built-in type `matrix`, which corresponds to the same format as a 4x4 matrix such as `D3DXMATRIX`.

Next, we need to write a structure to define what each vertex will look like. We'll see shortly how to write a corresponding structure in C++ to match this one. Each vertex in a model usually contains a local space position, a normal for lighting calculations, a color, and 2D texture coordinates if the object is textured. We can put all of this information into a structure like this:

```
struct VS_INPUT
{
    float3 vPosition    : POSITION;
    float3 vNormal      : NORMAL;
    float4 vColor       : COLOR;
```

```
float2 vTexCoords    : TEXCOORD;
};
```

This is just like a normal C++ structure except it has a little extra information to let the GPU know how to handle the data. You can see that each structure member uses a built-in HLSL type such as float4, float3, etc. The trailing integer identifies the number of floating-point components of the variable. So `vPosition` is a float3 and has three members, X, Y, and Z. Any of the built-in types can be accessed using either X, Y, Z, W in the case of position information, or R, G, B, A values for colors. I called the structure `VS_INPUT`, but you can call it anything you like.

The next step is to define what format of data the vertex shader will output to pass to the pixel shader. For now we'll output the fully transformed position, normal, and color for lighting in the pixel shader. The output function looks like this:

```
struct VS_OUTPUT
{
    float4 vPosition    : SV_POSITION;
    float3 vNormal      : NORMAL;
    float4 vColor       : COLOR0;
};
```

The SV prefix stands for system value, and it means the variable is attached to a system value that the GPU will use.

The Vertex Shader

The vertex shader is set up just like a normal C++ function. Here is the code:

```
VS_OUTPUT DefaultVS(VS_INPUT dataIn)
{
    VS_OUTPUT result;

    float4 vPos = float4(dataIn.vPosition, 1.0f);

    float4 vFinalPos = mul(vPos, g_mtxWorld);
    vFinalPos = mul(vFinalPos, g_mtxView);
    vFinalPos = mul(vFinalPos, g_mtxProj);

    result.vPosition = vFinalPos;
    result.vNormal = float3(0,0,0);
    result.vColor = float4(1,0,0,0);

    return result;
}
```

The first line, just like in C++, defines the function as being called `DefaultVS`, which takes a `VS_INPUT` structure as input and returns a `VS_OUTPUT`. The next line declares a new `VS_OUTPUT` structure called `result`, which we will fill in and return.

The next line creates a float4 vPos variable and initializes it constructor-like with the X,Y,Z position from the incoming dataIn variable and puts a 1.0f in the W component. The variable is then multiplied by the world, view, and transformation matrices and its final position is stored in vFinalPos.

The next three lines store the final position in the output result structure. Notice the color is set to red (1, 0, 0, 0) and the normal to 0,0,0. We'll see how to use these later for lighting. Finally, the filled-in structure is returned. Once the vertex shader has run for all the vertices in the object, it's time to run the pixel shader.

The Pixel Shader

The pixel shader is another function that looks just like the vertex shader. In fact, it takes the output from the vertex shader as its input, although it doesn't have to:

```
float4 DefaultPS(VS_OUTPUT dataIn) : SV_Target
{
    return dataIn.vColor;
}
```

Notice how the pixel shader returns a float4; this is because the job of the pixel shader is to pick a final color for the pixel. The shader is also bound to SV_Target, which means its output will be sent to the render target rather than anywhere else like the depth/stencil buffer. The SV stands for system value.

The Technique

Since an FX file can contain any number of vertex shaders, we need a way to tell Direct3D which one to use at a given time. To do this, we bundle a vertex shader, pixel shader, and geometry shader together into a *technique*. We can define a technique to render in multiple passes, to use just a vertex shader, and any number of other combinations. Our simple technique is listed below:

```
technique10 DefaultTechnique
{
    pass Pass0
    {
        SetGeometryShader(NULL);
        SetVertexShader(CompileShader(vs_4_0, VS()));
        SetPixelShader(CompileShader(ps_4_0, PS()));
    }
}
```

As you can see, it is called DefaultTechnique and contains a single rendering pass called Pass0. This pass disables use of a geometry shader and sets up the vertex and pixel shaders. Notice the use of the tags vs_4_0 and ps_4_0, which tell Direct3D which version of the pixel shader we are

compiling for. You should set all your DirectX 10 shaders to use version 4 as they are far more advanced than any previous versions.

Now let's look at the entire shader all together:

```

/*****
 *      Advanced 3D Game Programming with DirectX 10.0
 *      *      *      *      *      *      *      *      *      *      *
 *
 *      See license.txt for modification and distribution information
 *      copyright (c) 2007 by Peter Walsh, Wordware
 *****/

matrix g_mtxWorld;
matrix g_mtxView;
matrix g_mtxProj;

#define MAX_LIGHTS 10    // Ensure this is the same as the C++ value

int4 g_vLightStatus;
float4 g_vLightColors[MAX_LIGHTS];
float4 g_vLightDirections[MAX_LIGHTS];

////////////////////////////////////
// Default Vertex Shader
struct VS_INPUT
{
    float3 vPosition    : POSITION;
    float3 vNormal       : NORMAL;
    float4 vColor        : COLOR0;
    float2 vTexCoords    : TEXCOORD;
};

struct VS_OUTPUT
{
    float4 vPosition     : SV_POSITION;
    float3 vNormal        : NORMAL;
    float4 vColor         : COLOR0;
};

VS_OUTPUT DefaultVS(VS_INPUT dataIn)
{
    VS_OUTPUT result;

    float4 vPos = float4(dataIn.vPosition, 1.0f);

    float4 vFinalPos = mul(vPos, g_mtxWorld);
    vFinalPos = mul(vFinalPos, g_mtxView);
    vFinalPos = mul(vFinalPos, g_mtxProj);

    result.vPosition = vFinalPos;
    result.vNormal = float3(0,0,0);
    result.vColor = float4(1,0,0,0);

    return result;
}

```

```

/////////////////////////////////////////////////////////////////
// Default Pixel Shader
float4 DefaultPS(VS_OUTPUT dataIn) : SV_Target
{
    return dataIn.vColor;
}

/////////////////////////////////////////////////////////////////
// Default Technique
technique10 DefaultTechnique
{
    pass Pass0
    {
        SetGeometryShader(NULL);
        SetVertexShader(CompileShader(vs_4_0, VS()));
        SetPixelShader(CompileShader(ps_4_0, PS()));
    }
}

```

Setting Up the Shader in Code

Now that we have our shader written, we need to set up a few more things on the C++ side so that Direct3D knows to render with it. The steps involved are:

1. Load up and compile the FX file.
2. Check for any compile errors.
3. Get a pointer to the technique.
4. Set up the input layout.
5. Create the input layout.
6. Activate the input layout.
7. Get pointers to the global variables in the shader.
8. Set the default values for the globals.

After we complete these steps we are nearly ready to render. So let's start with loading and compiling the FX file. To do that, we use the function `D3DX10CreateEffectFromFile()`, which has the following prototype:

```

HRESULT D3DX10CreateEffectFromFile(
    LPCWSTR pFileName,
    CONST D3D10_SHADER_MACRO *pDefines,
    ID3D10Include *pInclude,
    LPCSTR pProfile,
    UINT HLSLFlags,
    UINT FXFlags,
    ID3D10Device *pDevice,
    ID3D10EffectPool *pEffectPool,
    ID3DX10ThreadPump *pPump,
    ID3D10Effect **ppEffect,
    ID3D10Blob **ppErrors,
    HRESULT *pHResult
);

```

pFileName	String containing the pathname of the FX file to load and compile.
pDefines	Set any preprocessor definitions here with an array of D3D10_SHADER_MACROS. NULL terminate with 0,0 macro.
pInclude	Used for custom file access code. Set to NULL.
pProfile	The shader model to compile to, which for us is version 4, or fx_4_0.
HLSLFlags	Compile options. I use D3D10_SHADER_ENABLE_STRICTNESS to ensure no legacy pre-4.0 code is being used and D3D10_SHADER_DEBUG during debug compiles.
FXFlags	Advanced FX compilation settings, which can be set to NULL.
pDevice	Pointer to the rendering ID3D10Device.
pEffectPool	You can create a pool of effects that share variables. Set to NULL for now.
pPump	You can execute this function in its own thread by specifying a pump here. Set to NULL for now.
ppEffect	Address of the effect pointer, which will be filled by Direct3D.
ppErrors	Pointer to an ID3D10Blob buffer, which contains any compilation errors in multibyte format. You'll need to use mbstowcs() to convert the string to Unicode before you can use it.
pHResult	Set to NULL.

Here is the code to create and compile the effect file and check for errors:

```
ID3D10Blob *pErrors = 0;
// Create the default rendering effect
r = D3DX10CreateEffectFromFile(L"..\\GameLib\\DefaultShader.fx", NULL, NULL,
"fx_4_0", shaderFlags, 0, m_pDevice, NULL, NULL, &m_pDefaultEffect,
&pErrors, NULL);
if(pErrors)
{
    char *pCompileErrors = static_cast<char*>(pErrors->GetBufferPointer());
    TCHAR wcsErrors[MAX_PATH];
    mbstowcs(wcsErrors, pCompileErrors, MAX_PATH);
    OutputDebugString(wcsErrors);
}
```

The next step is to get a pointer to the technique, which is done using the function ID3D10Effect::GetTechniqueByName(). This function takes a string with the name of the technique, like this:

```
m_pDefaultTechnique = m_pDefaultEffect->
GetTechniqueByName("DefaultTechnique");
if(!m_pDefaultTechnique)
{
    throw cGameError(L"Could not find default technique in DefaultShader.fx");
}
```

Now that we have the technique, we need to set up the input layout in C++ so that Direct3D knows how to send data to the shader. An *input*

layout, known as the flexible vertex format in previous versions of DirectX, defines how the vertex data is stored for the vertex shader. Since our shader vertex contains a position, normal, color, and texture coordinates, we need to tell this to Direct3D along with how much space they take up. This is achieved by creating an array of `D3D10_INPUT_ELEMENT_DESC` structures for each entry in a vertex. The structure looks like this:

```
typedef struct D3D10_INPUT_ELEMENT_DESC {
    LPCSTR SemanticName;
    UINT SemanticIndex;
    DXGI_FORMAT Format;
    UINT InputSlot;
    UINT AlignedByteOffset;
    D3D10_INPUT_CLASSIFICATION InputSlotClass;
    UINT InstanceDataStepRate;
} D3D10_INPUT_ELEMENT_DESC;
```

SemanticName	The name of the semantic, such as "POSITION." You can find a full list of all semantic names in the DirectX 10 C++ documentation.
SemanticIndex	If there is more than one vertex entry of the same type this is used to differentiate them. For example, if you had two positions you would set this entry to 0 and 1.
Format	A <code>DXGI_FORMAT</code> describing the format of the data.
InputSlot	The input assembler index that will supply this data.
AlignedByteOffset	You can type the offset of the member in the vertex here or use <code>D3D10_APPEND_ALIGNED_ELEMENT</code> to automatically calculate it if there is no missing space.
InputSlotClass	Set this to <code>D3D10_INPUT_PER_VERTEX_DATA</code> .
InstanceDataStepRate	Used for advanced instancing. Set this to 0 for now.

Here is the code I used to create an input layout:

```
D3D10_INPUT_ELEMENT_DESC defaultLayout[] =
{
    {"POSITION", 0, DXGI_FORMAT_R32G32B32_FLOAT, 0, 0, D3D10_INPUT_PER_VERTEX_DATA, 0},
    {"NORMAL", 0, DXGI_FORMAT_R32G32B32_FLOAT, 0, D3D10_APPEND_ALIGNED_ELEMENT, D3D10_INPUT_PER_VERTEX_DATA, 0},
    {"COLOR", 0, DXGI_FORMAT_R32G32B32A32_FLOAT, 0, D3D10_APPEND_ALIGNED_ELEMENT, D3D10_INPUT_PER_VERTEX_DATA, 0},
    {"TEXCOORD", 0, DXGI_FORMAT_R32G32_FLOAT, 0, D3D10_APPEND_ALIGNED_ELEMENT, D3D10_INPUT_PER_VERTEX_DATA, 0},
};
```

Notice how this layout exactly matches the format of the vertex in the shader. Once the layout is defined we need to use it to create an

ID3D10InputLayout using the function ID3D10Device::CreateInputLayout(). This function has the following prototype:

```
HRESULT CreateInputLayout(  
    const D3D10_INPUT_ELEMENT_DESC *pInputElementDescs,  
    UINT NumElements,  
    const void *pShaderBytecodeWithInputSignature,  
    SIZE_T BytecodeLength,  
    ID3D10InputLayout **ppInputLayout  
);
```

pInputElementDescs	Pointer to the array of D3D10_INPUT_ELEMENT_DESCs we just filled in.
NumElements	The number of elements in the structure.
pShaderBytecodeWithInputSignature	The byte code signature of the shader, which is acquired by looking at the technique description.
BytecodeLength	The size of the compiled shader.
ppInputLayout	Pointer to the address of the interface, which will be filled in by Direct3D.

This is the code to create our default input layout:

```
UINT uiNumElements = sizeof(defaultLayout)/sizeof(defaultLayout[0]);  
D3D10_PASS_DESC descPass;  
m_pDefaultTechnique->GetPassByIndex(0)->GetDesc(&descPass);  
  
r = m_pDevice->CreateInputLayout(defaultLayout, uiNumElements,  
    descPass.pIAInputSignature, descPass.IAInputSignatureSize,  
    &m_pDefaultInputLayout);  
if(FAILED(r))  
{  
    throw cGameError(L"Could not create default layout");  
}
```

The final step is to set the input layout we created as the active input layout with Direct3D using the function ID3D10Device::IASetInputLayout(). The IA stands for input assembler. It has the following prototype:

```
void IASetInputLayout(  
    ID3D10InputLayout *pInputLayout  
);
```

As you can see it takes a single parameter, which is a pointer to a valid input layout. I called it like this:

```
m_pDevice->IASetInputLayout(m_pDefaultInputLayout);
```

At this point the shader is now active and the correct input layout has been selected. Now we need a way to define how to hold our vertices in C++. To do this, I added a class to cGraphicsLayer called cDefaultVertex, which looks like the following:

```

class cDefaultVertex
{
public:
    D3DXVECTOR3 m_vPosition;
    D3DXVECTOR3 m_vNormal;
    D3DXCOLOR m_vColor;
    D3DXVECTOR2 m_TexCoords;
};

```

Again, notice how it matches both the input layout we created and the structure created in the shader. If these three items do not match, then very freaky things will get rendered to your screen.

Now that the shader is up and running, we need a way to communicate with it. For example, if we are going to use it to render each object we need a way to update the transformation matrices. To do this, we get a pointer to each variable we want to update using the function `ID3D10Effect::GetVariableByName()` like this:

```

m_pmtxWorldVar =
m_pDefaultEffect->GetVariableByName("g_mtxWorld")->AsMatrix();
m_pmtxViewVar =
m_pDefaultEffect->GetVariableByName("g_mtxView")->AsMatrix();
m_pmtxProjVar =
m_pDefaultEffect->GetVariableByName("g_mtxProj")->AsMatrix();

```

Notice how the string passed to the function exactly matches the string name in the effect file. We can now use these pointers to update the transformation matrices. I added a function called `UpdateMatrices()` that does this for us:

```

void cGraphicsLayer::UpdateMatrices()
{
    m_pmtxWorldVar->SetMatrix((float*)&m_mtxWorld);
    m_pmtxViewVar->SetMatrix((float*)&m_mtxView);
    m_pmtxProjVar->SetMatrix((float*)&m_mtxProj);
}

```

This function updates the shader transformation matrices with data from the matrices stored in `cGraphicsLayer`. Now let's check out all the shader code together, which I've put into a single function called `CreateDefaultShader()`:

```

void cGraphicsLayer::CreateDefaultShader()
{
    HRESULT r = 0;

    DWORD shaderFlags = D3D10_SHADER_ENABLE_STRICTNESS;
    #if defined( DEBUG ) || defined( _DEBUG )
        // Turn on extra debug info when in debug config
        shaderFlags |= D3D10_SHADER_DEBUG;
    #endif

    ID3D10Blob *pErrors = 0;
    // Create the default rendering effect

```

```

r = D3DX10CreateEffectFromFile(L"..\\GameLib\\DefaultShader.fx",
    NULL, NULL, "fx_4_0", shaderFlags, 0,
    m_pDevice, NULL, NULL, &m_pDefaultEffect, &pErrors, NULL);
if(pErrors)
{
    char *pCompileErrors = static_cast<char*>(
        pErrors->GetBufferPointer());
    TCHAR wcsErrors[MAX_PATH];
    mbstowcs(wcsErrors, pCompileErrors, MAX_PATH);
    OutputDebugString(wcsErrors);
}

if(FAILED(r))
{
    throw cGameError(
        L"Could not create default shader - DefaultShader.fx");
}

m_pDefaultTechnique =
    m_pDefaultEffect->GetTechniqueByName("DefaultTechnique");
if(!m_pDefaultTechnique)
{
    throw cGameError(
        L"Could not find default technique in DefaultShader.fx");
}

// Set up the input layout
D3D10_INPUT_ELEMENT_DESC defaultLayout[] =
{
    {"POSITION", 0, DXGI_FORMAT_R32G32B32_FLOAT, 0, 0,
     D3D10_INPUT_PER_VERTEX_DATA, 0 },
    {"NORMAL", 0, DXGI_FORMAT_R32G32B32_FLOAT, 0,
     D3D10_APPEND_ALIGNED_ELEMENT,
     D3D10_INPUT_PER_VERTEX_DATA, 0},
    {"COLOR", 0, DXGI_FORMAT_R32G32B32A32_FLOAT, 0,
     D3D10_APPEND_ALIGNED_ELEMENT,
     D3D10_INPUT_PER_VERTEX_DATA, 0},
    {"TEXCOORD", 0, DXGI_FORMAT_R32G32_FLOAT, 0,
     D3D10_APPEND_ALIGNED_ELEMENT,
     D3D10_INPUT_PER_VERTEX_DATA, 0},
};

UINT uiNumElements = sizeof(defaultLayout)/sizeof(defaultLayout[0]);
D3D10_PASS_DESC descPass;
m_pDefaultTechnique->GetPassByIndex(0)->GetDesc(&descPass);
r = m_pDevice->CreateInputLayout(defaultLayout, uiNumElements,
    descPass.pIAInputSignature,
    descPass.IAInputSignatureSize, &m_pDefaultInputLayout);
if(FAILED(r))
{
    throw cGameError(L"Could not create default layout");
}

m_pDevice->IASetInputLayout(m_pDefaultInputLayout);

m_pmtxWorldVar = m_pDefaultEffect->GetVariableByName(

```

```

        "g_mtxWorld")->AsMatrix();
    m_pmtxViewVar = m_pDefaultEffect->GetVariableByName(
        "g_mtxView")->AsMatrix();
    m_pmtxProjVar = m_pDefaultEffect->GetVariableByName(
        "g_mtxProj")->AsMatrix();

    D3DXMATRIX mtxWorld;
    D3DXMatrixIdentity(&mtxWorld);
    SetWorldMtx(mtxWorld);

    D3DXMATRIX mtxView;
    D3DXVECTOR3 vCamPos(0.0f, 1.0f, -3.0f);
    D3DXVECTOR3 vCamAt(0.0f, 1.0f, 0.0f);
    D3DXVECTOR3 vCamUp(0.0f, 1.0f, 0.0f);
    D3DXMatrixLookAtLH(&mtxView, &vCamPos, &vCamAt, &vCamUp);
    SetViewMtx(mtxView);

    D3DXMATRIX mtxProj;
    D3DXMatrixPerspectiveFovLH(&mtxProj, (float)D3DX_PI * 0.5f,
        m_rcScreenRect.right/(float)m_rcScreenRect.bottom, 0.1f, 100.0f);
    SetProjMtx(mtxProj);

    UpdateMatrices();

    m_pLightDirVar = m_pDefaultEffect->GetVariableByName(
        "g_vLightDirections" )->AsVector();
    m_pLightColorVar = m_pDefaultEffect->GetVariableByName(
        "g_vLightColors" )->AsVector();
    m_pNumLightsVar = m_pDefaultEffect->GetVariableByName(
        "g_vLightStatus" )->AsVector();
}

```

All of these functions are called from `InitD3D()`, which now looks like this:

```

void cGraphicsLayer::InitD3D(int width, int height)
{
    HRESULT r = 0;

    // Keep a copy of the screen dimensions
    m_rcScreenRect.left = m_rcScreenRect.top = 0;
    m_rcScreenRect.right = width;
    m_rcScreenRect.bottom = height;

    CreateDeviceAndSwapChain();
    CreateViewport();
    CreateDepthStencilBuffer();
    CreateDebugText();

    // Attach the render target view to the output merger state
    m_pDevice->OMSetRenderTargets(1, &m_pRenderTargetView, m_pDepthStencilView);

    CreateDefaultShader();
}

```

Now we're ready to render! But before we do that, let's take another look at buffers and find out how to load up 3D models to render.

More about Depth Buffers

Often in computer graphics you run into the problem of determining which pixels of each triangle are visible to the viewer. A drawing algorithm typically acts in the same way as a painter. When you draw a triangle on the screen, the device draws it right over everything else that's there, like painting on a canvas. This presents an immediate problem: The image can appear incorrect if you draw polygons out of order. Imagine what a picture would look like if an artist placed birds and clouds on the canvas first, then painted the blue sky on top of it, covering everything he had already drawn! Figure 7.1 shows what I am talking about.




Figure 7.1: The depth problem

The old way to solve this problem, before there was readily available hardware to solve the problem for you, was to implement the painter's algorithm. In it, you draw the world the same way a painter would: Draw the farthest things first, the nearest things last. This way, your image ends up being drawn correctly. If it doesn't seem intuitive, just think of how painters create paintings. First, they draw the farthest things away (sky, mountains, whatnot). As the paint dries, they paint on top of what is already there, adding elements in the foreground.

There are a few problems with this algorithm. First of all, it doesn't always work. You have to sort your polygons based on depth, but unless the polygons are parallel with the view plane, how do you determine the depth of the entire polygon? You could use the nearest vertex, the farthest, or the average of all the vertices, but these all have cases that won't work. There are some other cases, like triangles that intersect each other, that

cannot possibly be drawn correctly with the painter's algorithm. Some triangle configurations are also unrenderable using the painter's algorithm (see Figure 7.2). Finally, you need to actually have an ordered list of polygons to draw. That involves a lot of sorting, which can become prohibitive as the triangle count increases. Most naïve sorting algorithms are $O(n^2)$, and while the fastest ones approach $O(n \lg n)$, this still will kill you if you have thousands of triangles visible on the screen at once.




Figure 7.2: Unworkable sorting problems

Isn't there a better way to handle finding the nearest triangle at each pixel? As usual in computer science, there's an extremely simple but inefficient brute-force way to attack the problem. The brute-force way ends up being the one that most cards use. All modern DirectX 10 cards support the method called *z-buffering*, or *depth buffering*.

The z-buffer is a second image buffer you keep in addition to the frame buffer. The z-buffer holds a single number that represents the distance at every pixel (measured with the z component, since you're looking down the z-axis with the coordinate system). Each pixel in the z-buffer holds a z value of the closest pixel drawn up to that point. Note that you don't use the Pythagorean distance in the z-buffer, just the raw z value of the pixels.

Before drawing, you initialize the buffer to an extremely far away value. When you draw a triangle, you iterate not only color information, but also the depth (distance along the z-axis from the camera) of the current pixel you want to draw. When you go to draw, you check the iterated depth against the value currently in the depth buffer. If the current depth is closer than the buffer depth, it means the pixel is in front of the pixel

already in the frame buffer. So you can update the frame buffer with the color value and update the depth buffer with the new depth. If the iterated depth is farther away than the z-buffer depth, that means the pixel already in the frame buffer is closer than the one you're trying to draw, so you do nothing. That way, you never draw a pixel that is obscured by something already in the scene. Figure 7.3 shows how the frame buffer changes as you rasterize a triangle.




Figure 7.3: Depth buffers in action

Z-buffering is not a perfect solution. First of all, you don't have infinite resolution to represent depth (again, numerical imprecision comes up to haunt us). The problem is that the precision of z-buffers doesn't vary linearly with z (because z doesn't vary linearly in screen space). Because of this, a lot of the precision (like 90% of it) is used up in the front 10% of the scene. The end result is artifacts tend to show up, especially in primitives drawn far away from the viewer. Most depth buffers are 24 bits. Going any higher than 24 bits is really a waste; going to 32 bits means there are 4 billion possible depth values between your near and far plane, which is way too much. This is why the top 8 bits of a 32-bit depth buffer are usually used for stencil information.

Another problem with z-buffering is the speed of the algorithm. You need to perform a comparison per-pixel per-triangle. This would be prohibitive, but thankfully the card does this automatically for you, at no speed loss. Thanks to the silicon, you don't have to worry about the rendering order of the triangles (until, of course, you start doing stuff like alpha blending, but that comes later...). Actually, the main reason anyone

uses the z-buffer algorithm is that brute-force algorithms tend to be extremely easy to implement in hardware.

Stencil Buffers

Stencil buffers rose to fame a few years ago, and they are still extremely useful. Originally they existed solely on high-end SGI machines costing barrels of money, but now they are common in the consumer market.

Stencils are used all over the place in the “real world.” For example, when people paint arrows and letters on the street, they lay large pieces of metal on the ground, then spray paint on them. The stencil constrains where the paint can go, so that you get a nice arrow on the ground without having to painstakingly paint each edge.

Stencil buffers work in a similar way. You can use them to constrain drawing to a particular region and to record how many times a particular region has been written to. They can even be used to render dynamic shadows really quickly. We’ll discuss them more later, but for right now just keep them in mind. Know that typically they come in 1-, 4-, and 8-bit varieties, and that they share bits with the depth buffer (in a 32-bit z-buffer, 24 bits are for z, 8 are for stencil). We of course are using 8 bits for our stencil buffer.

Vertex Buffers

DirectX 6.0 was the first version of the SDK to include vertex buffers. The circular definition is that they are buffers filled with vertices, but this is actually about as good a definition as you’ll need. Instead of creating a buffer with image data or sound data, this is a buffer with vertex data. Like any other surface, it must be locked to gain access and unlocked when you relinquish said access. These vertices may be at any stage of the vertex pipeline (transformed, untransformed, lit, unlit). You can draw them using special draw primitive commands that specifically make use of the vertex buffer. In DirectX 10 a vertex buffer, just like any other kind of buffer, is an `ID3D10Buffer`. In past versions of DirectX there were special buffers for each kind of buffer. So when I talk of vertex buffers here I really just mean an `ID3D10Buffer` that happens to be used to hold vertices.

Vertex buffers have two primary uses. First, they accelerate the rendering of static geometry: You create a vertex buffer to hold the geometry for an object, fill it once, and draw the set of vertices every frame. You can use the vertex buffer to optimize the vertex data for the particular device so it can draw it as fast as possible. The other use for vertex buffers is to provide a high-bandwidth stream of data so you can feed the graphics card primitives as fast as you can.

Vertex buffers also make it easier for the hardware to reuse vertex data. For example, let’s say you’re drawing an object with multiple textures on it, or with multiple states for one section or another. Instead of having to separately transform, light, and clip all of the vertices in the object for

each texture, you can run the entire geometry pipeline on the vertex buffer once, then draw groups of vertices as you like. The subdivision surface sample application makes use of vertex buffers, if you want to see a piece of code that takes advantage of this nifty feature.

Creating vertex buffers is very simple; it is done with a call to `ID3D10Device::CreateBuffer()`, which has the following prototype:

```
HRESULT CreateBuffer(
    const D3D10_BUFFER_DESC *pDesc,
    const D3D10_SUBRESOURCE_DATA *pInitialData,
    ID3D10Buffer **ppBuffer
);
```

The first parameter is the address of a `D3D10_BUFFER_DESC` structure, which is shown below. The second parameter takes the address of a `D3D10_SUBRESOURCE_DATA` structure to fill the vertex buffer with initial data, which I'll show you in a bit. The final parameter is where the pointer is set to point at a valid buffer.

The `D3D10_BUFFER_DESC` structure looks like this:

```
typedef struct D3D10_BUFFER_DESC {
    UINT ByteWidth;
    D3D10_USAGE Usage;
    UINT BindFlags;
    UINT CPUAccessFlags;
    UINT MiscFlags;
} D3D10_BUFFER_DESC;
```

ByteWidth	The size of the buffer in bytes.
Usage	Usually set to <code>D3D10_USAGE_DEFAULT</code> .
BindFlags	How the buffer is bound to the rendering pipeline. Since this is a vertex buffer, we set it to <code>D3D10_BIND_VERTEX_BUFFER</code> . If you were creating an index buffer, you would use <code>D3D10_BIND_INDEX_BUFFER</code> .
CPUAccessFlags	Allows you to set whether the CPU can access the buffer, or set to <code>NULL</code> for better performance. If you are not setting initial data, you must give the CPU access to fill the buffer with data later.
MiscFlags	Generally set to <code>NULL</code> .

The second type of structure to fill out is `D3D10_SUBRESOURCE_DATA`, which looks like this:

```
typedef struct D3D10_SUBRESOURCE_DATA {
    const void *pSysMem;
    UINT SysMemPitch;
    UINT SysMemSlicePitch;
} D3D10_SUBRESOURCE_DATA;
```

pSysMem	Pointer in system memory of data to use to fill the vertex buffer with.
SysMemPitch	Always set to 0.
SysMemSlicePitch	Always set to 0.

Here is example code of a vertex buffer being created:

```
D3D10_BUFFER_DESC descBuffer;
memset(&descBuffer, 0, sizeof(descBuffer));
descBuffer.Usage = D3D10_USAGE_DEFAULT;
descBuffer.ByteWidth = sizeof(cGraphicsLayer::cDefaultVertex) * NumVerts();
descBuffer.BindFlags = D3D10_BIND_VERTEX_BUFFER;
descBuffer.CPUAccessFlags = 0;
descBuffer.MiscFlags = 0;

D3D10_SUBRESOURCE_DATA resData;
memset(&resData, 0, sizeof(resData));
resData.pSysMem = &m_verts[0];
Graphics()->GetDevice()->CreateBuffer(&descBuffer, &resData, &m_pVertexBuffer);
```

You must use vertex buffers to draw primitives such as triangles, lines, and so on. In the past you could just draw primitives with calls to certain functions, but these days you *must* package your rendering data into vertex buffers before rendering them with Draw() or DrawIndexed().

To render with vertex buffers you have to set them as active in the input assembler and set the topology. You set them as active with ID3D10Device::IASetVertexBuffers(), which has the following prototype:

```
void IASetVertexBuffers(
    UINT StartSlot,
    UINT NumBuffers,
    ID3D10Buffer *const *ppVertexBuffers,
    const UINT *pStrides,
    const UINT *pOffsets
);
```

StartSlot	The input slot to bind the buffer to.
NumBuffers	The number of buffers to be activated, usually 1.
ppVertexBuffers	Array of vertex buffers to activate.
pStrides	Array of stride values, which is the width in bytes of a single vertex.
pOffsets	Array of offset values, which contain an offset into the vertex buffer to start at.

The *topology* is the new DirectX 10 term for the primitive type, such as a triangle list, etc. You set the topology with the function ID3D10Device::IASetPrimitiveTopology(). Here is some sample code that shows a vertex buffer being made ready for rendering:

```

UINT uiStride = sizeof(cGraphicsLayer::cDefaultVertex);
UINT uiOffset = 0;

Graphics()->GetDevice()->IASetVertexBuffers(
    0, 1, &m_pVertexBuffer, &uiStride, &uiOffset);
Graphics()->GetDevice()->IASetPrimitiveTopology(
    D3D10_PRIMITIVE_TOPOLOGY_TRIANGLELIST);

```

Lighting with Shaders

Now that we have our shaders set up, it's time to add some lights. For this example, we'll program a directional light shader. In later sections of the book, you'll learn about more advanced techniques. A directional light can be represented with just a color and a direction. I've encapsulated these into a simple class called `cLight`, which you can see here:

```

class cLight
{
public:
    cLight()
    {
    }

    D3DXCOLOR m_vColor;
    D3DXVECTOR3 m_vDirection;
};

```

What we want to do is create an array of active lights. For now I've set the maximum number of active lights to 10 using the preprocessor definition `MAX_LIGHTS`. Since each light must be calculated for each pixel, the more lights you have the more computationally intensive your scene becomes to render. You can have any number of lights in your game, but you'll need a system to prioritize only the 10 or so lights nearest to the player to be active at any one time. This is the array of lights I've added to `cGraphicsLayer`:

```

cLight m_aLights[MAX_LIGHTS];           // Light array
int m_iNumLights;                        // Number of active lights

```

Each frame, we need to transfer the data from this array into the GPU of the graphics card so that the shaders can access it. For this we need to add an identical array to the shader:

```

int4 g_vLightStatus;
float4 g_vLightColors[MAX_LIGHTS];
float4 g_vLightDirections[MAX_LIGHTS];

```

The `g_vLightStatus` variable is used to figure out how many lights are active. I also added access variables to `cGraphicsLayer` to allow access to these shader variables:

```

ID3D10EffectVectorVariable *m_pLightDirVar;
ID3D10EffectVectorVariable *m_pLightColorVar;

```

```
ID3D10EffectVectorVariable *m_pNumLightsVar;
```

These variables are set up in the `cGraphicsLayer::CreateDefaultShader()` function like this:

```
m_pLightDirVar = m_pDefaultEffect->GetVariableByName( "g_vLightDirections"
)->AsVector();
m_pLightColorVar = m_pDefaultEffect->GetVariableByName( "g_vLightColors"
)->AsVector();
m_pNumLightsVar = m_pDefaultEffect->GetVariableByName( "g_vLightStatus"
)->AsVector();
```

Now that we have access, we need a way of updating the values each frame. For this I added a function called `cGraphicsLayer::UpdateLights()`, which is called once a frame.

```
void cGraphicsLayer::UpdateLights()
{
    int iLightData[4] = {m_iNumLights, 0, 0, 0};

    m_pNumLightsVar->SetIntVector(iLightData);

    for(int iCurLight = 0 ; iCurLight < m_iNumLights ; iCurLight++)
    {
        m_pLightDirVar->SetFloatVectorArray(
            (float*)m_aLights[iCurLight].m_vDirection, iCurLight, 1);
        m_pLightColorVar->SetFloatVectorArray(
            (float*)m_aLights[iCurLight].m_vColor, iCurLight, 1);
    }
}
```

Nearly there now. We have light data stored in the graphics layer, and corresponding arrays set up in the shader. Each frame the light data is copied from the graphics layer into the shader. Let's add some lights and then see how to calculate the lighting value. To add a red light that is pointing up to the scene you would just add this to your code:

```
Graphics()->AddLight(
    D3DXCOLOR(1.0f, 0.0f, 0.0f, 1.0f), D3DXVECTOR3(0.0f, 1.0f, 0.0f));
```

Now let's look at how the pixel shader calculates the lighting:

```
////////////////////////////////////
// Default Pixel Shader
float4 DefaultPS(VS_OUTPUT dataIn) : SV_Target
{
    float4 finalColor = 0;

    for(int iCurLight = 0 ; iCurLight < g_vLightStatus.x ; iCurLight++)
    {
        finalColor += saturate(dot(g_vLightDirections[iCurLight],
            dataIn.vNormal) * g_vLightColors[iCurLight]);
    }

    return finalColor;
}
```


This code uses a simple lighting calculation by taking the dot product of the face normal, which is passed in from the vertex shader, with the vector to the light. The result is used to calculate the intensity of the light. Note how it loops for each active light and adds all the results together to compute the final color. The screenshot in Figure 7.4 in the upcoming sample application shows the results.

Application: D3D View

The sample application for this chapter is an object viewer. It loads an object file from disk and displays the object spinning around the scene. Before you can draw the spinning object, you of course need a way to load it.

There are a myriad of different object formats out there. OBJ, 3DS, DXF, ASC, and PLG files are available on the net or can be easily constructed. However, they're all either extremely hard to parse or not fully featured enough. Rather than trudge through a parser for one of these data types, I'm going to circumvent a lot of headache and create our own format. The web is rife with parsers for any of these other formats, so if you want to parse it you won't have to reinvent the wheel.

The .o3d Format

The name for the object format will be .o3d (object 3D format). It's a Unicode text file, which makes it easy to edit manually if the need arises. The object is designed for regular objects that have no color information but may have normal or texture information.

```
Tetrahedron 3 1 4 4
-1.0 -1.0 -1.0
1.0 1.0 -1.0
-1.0 1.0 1.0
1.0 -1.0 1.0
2 3 4
1 4 3
1 3 2
1 2 4
```

The first line of the file is the header. It has five fields, separated by spaces. They are, in order:

- The name for the object (spaces within the name are not allowed).
- The number of fields per vertex. This can be three (just position), five (three position and two texture), six (three position and three normal), or eight (three position, three normal, and two texture).
- The offset for the indices. Some index lists are 0-based, while some are 1-based. This offset is subtracted from each of the indices on load.

Since the indices in the tetrahedron list start at 1, the offset is 1 (since index 1 will actually be element 0 internally).

- The number of vertices in the model.
- The number of triangles in the model.

After the header line, there is one line for each of the vertices. Each line has n fields separated by spaces (where n is the number of fields per vertex). The first three fields are always position.

After the list of vertices, there is a list of triangles. Each triangle is defined with three indices separated by spaces. Each index has the offset (defined in the header) subtracted from it.

The *cModel* Class

To load .o3d models, I'm going to create a class that represents a model. It has one constructor that takes a filename on disk. The constructor opens the file, parses it, and extracts the vertex and triangle information. It takes the information and fills up two vectors. If the file it loads does not have normal information defined for it, the class uses face averaging to automatically generate normals for the object.

Face averaging is used often to find normals for vertices that make a model appear rounded when Gouraud shading is used on it. The normals for each of the faces are computed, and the normal is added to each of the face's vertices. When all of the faces have contributed their normals, the vertex normals are normalized. This, in essence, makes each vertex normal the average of the normals of the faces around it. This gives the model a smooth look.

The *cModel* class can automatically draw an object at any world position. It uses `DrawIndexed()` to draw the entire model in one fell swoop.

```

/*****
 *           Advanced 3D Game Programming with DirectX 10.0
 * ****
 *
 *   See license.txt for modification and distribution information
 *   copyright (c) 2007 by Peter Walsh, Wordware
 *****/

#ifdef _MODEL_H
#define _MODEL_H

#include <vector>
#include <string>
#include "..\math3d\tri.h"
#include "..\math3d\mathD3D.h"
#include "GraphicsLayer.h"

class cModel
{

```

```

typedef tri<WORD> sTri;

std::vector<sTri>      m_tris;
std::vector<cGraphicsLayer::cDefaultVertex> m_verts;

std::wstring          m_name;

ID3D10Buffer          *m_pVertexBuffer;
ID3D10Buffer          *m_pIndexBuffer;
public:

    cModel( const TCHAR *filename );
    cModel( const TCHAR *name, int nVerts, int nTris );

    ~cModel();

    float GenRadius();
    void Scale( float amt );

    void Draw();

    //----- Access functions.

    int NumVerts(){ return m_verts.size(); }
    int NumTris(){ return m_tris.size(); }
    const TCHAR *Name(){ return m_name.c_str(); }

    /**
     * Some other classes may end up using cModel
     * to assist in their file parsing. Because of this
     * give them a way to get at the vertex and triangle
     * data.
     */
    cGraphicsLayer::cDefaultVertex *VertData(){ return &m_verts[0]; }
    sTri *TriData(){ return &m_tris[0]; }

};

#endif // _MODEL_H

```

And here is the implementation code:

```

/*****
 *      Advanced 3D Game Programming with DirectX 10.0
 *      * * * * *
 *
 *      See license.txt for modification and distribution information
 *      copyright (c) 2007 by Peter Walsh, Wordware
 *****/

#include "stdafx.h"
#include "../math3d/point3.h"
#include "../math3d/matrix4.h"
#include "Model.h"
#include "file.h"

```

```

using namespace std;

cModel::cModel(const TCHAR *name, int nVerts, int nTris) :
    m_name(name)
{
    int i;

    m_verts.reserve(nVerts);
    cGraphicsLayer::cDefaultVertex vert;
    for(i=0; i<nVerts; i++)
    {
        m_verts.push_back(vert);
    }

    m_tris.reserve(nTris);
    sTri tri;
    for(i=0; i<nTris; i++)
    {
        m_tris.push_back(tri);
    }

    m_pVertexBuffer = NULL;
    m_pIndexBuffer = NULL;
}

cModel::~cModel()
{
    SafeRelease(m_pVertexBuffer);
    SafeRelease(m_pIndexBuffer);
}

cModel::cModel(const TCHAR *filename)
{
    int i;

    cFile file;
    file.Open(filename);

    queue<wstring> m_tokens;

    file.TokenizeNextNCLine(&m_tokens, '#');

    // first token is the name.
    m_name = m_tokens.front();
    m_tokens.pop();

    // next is the # of fields in the vertex info
    int nVertexFields = _wtoi(m_tokens.front().c_str());
    m_tokens.pop();

    // next is the triangle offset
    int offset = _wtoi(m_tokens.front().c_str());
    m_tokens.pop();

    // next is the # of vertices

```

```

int nVerts = _wtoi(m_tokens.front().c_str());
m_tokens.pop();

// next is the # of triangles
int nTris = _wtoi(m_tokens.front().c_str());
m_tokens.pop();

if(!m_tokens.empty())
{
    // additional info in the header
}

// Reserve space in the vector for all the verts.
// This will speed up all the additions, since only
// one resize will be done.
m_verts.reserve(nVerts);
for(i=0; i<nVerts; i++)
{
    //while(!m_tokens.empty()) m_tokens.pop();
    file.TokenizeNextNCLine(&m_tokens, '#');

    cGraphicsLayer::cDefaultVertex curr;

    // Vertex data is guaranteed
    curr.m_vPosition.x = _wtof(m_tokens.front().c_str());
    m_tokens.pop();
    curr.m_vPosition.y = _wtof(m_tokens.front().c_str());
    m_tokens.pop();
    curr.m_vPosition.z = _wtof(m_tokens.front().c_str());
    m_tokens.pop();

    // hack to load color
    if(nVertexFields == 4)
    {
        curr.m_vColor = D3DXCOLOR(_wtoi( m_tokens.front().c_str() ));
        m_tokens.pop();
    }

    // Load normal data if nfields is 6 or 8
    if(nVertexFields == 6 || nVertexFields == 8)
    {
        curr.m_vNormal.x = _wtof(m_tokens.front().c_str());
        m_tokens.pop();
        curr.m_vNormal.y = _wtof(m_tokens.front().c_str());
        m_tokens.pop();
        curr.m_vNormal.z = _wtof(m_tokens.front().c_str());
        m_tokens.pop();
    }
    else
    {
        curr.m_vNormal = D3DXVECTOR3(0, 0, 0);
    }

    // Load texture data if nfields is 5 or 8
    if(nVertexFields == 5 || nVertexFields == 8)

```

```

    {
        curr.m_TexCoords.x = _wtof(m_tokens.front().c_str());
        m_tokens.pop();
        curr.m_TexCoords.y = _wtof(m_tokens.front().c_str());
        m_tokens.pop();
    }
    else
    {
        curr.m_TexCoords.x = 0.f;
        curr.m_TexCoords.y = 0.f;
    }

    m_verts.push_back(curr);
}

// Reserve space in the vector for all the verts.
// This will speed up all the additions, since only
// one resize will be done.
m_tris.reserve(nTris);
for(i=0; i<nTris; i++)
{
    m_tokens.empty();
    file.TokenizeNextNCLine(&m_tokens, '#');

    sTri tri;

    // vertex data is guaranteed
    tri.v[0] = _wtoi(m_tokens.front().c_str()) - offset;
    m_tokens.pop();
    tri.v[1] = _wtoi(m_tokens.front().c_str()) - offset;
    m_tokens.pop();
    tri.v[2] = _wtoi(m_tokens.front().c_str()) - offset;
    m_tokens.pop();

    m_tris.push_back(tri);
}

if(nVertexFields == 3 || nVertexFields == 4 || nVertexFields == 5)
{
    // Normals weren't provided. Generate our own.

    // First set all the normals to zero.
    for(i=0; i<nVerts; i++)
    {
        m_verts[i].m_vNormal = D3DXVECTOR3(0,0,0);
    }

    // Then go through and add each triangle's normal
    // to each of its verts.
    for(i=0; i<nTris; i++)
    {
        plane3 plane(
            m_verts[ m_tris[i].v[0] ].m_vPosition,
            m_verts[ m_tris[i].v[1] ].m_vPosition,

```

```

        m_verts[ m_tris[i].v[2] ].m_vPosition);

    m_verts[ m_tris[i].v[0] ].m_vNormal +=
        D3DXVECTOR3(plane.n.x, plane.n.y, plane.n.z);
    m_verts[ m_tris[i].v[1] ].m_vNormal +=
        D3DXVECTOR3(plane.n.x, plane.n.y, plane.n.z);
    m_verts[ m_tris[i].v[2] ].m_vNormal +=
        D3DXVECTOR3(plane.n.x, plane.n.y, plane.n.z);
}

// Finally normalize all of the normals
for(i=0; i<nVerts; i++)
{
    D3DXVec3Normalize(&m_verts[i].m_vNormal, &m_verts[i].m_vNormal);
}

// Set up vertex and index buffers
m_pVertexBuffer = NULL;
m_pIndexBuffer = NULL;

D3D10_BUFFER_DESC descBuffer;
memset(&descBuffer, 0, sizeof(descBuffer));
descBuffer.Usage = D3D10_USAGE_DEFAULT;
descBuffer.ByteWidth =
    sizeof(cGraphicsLayer::cDefaultVertex) * NumVerts();
descBuffer.BindFlags = D3D10_BIND_VERTEX_BUFFER;
descBuffer.CPUAccessFlags = 0;
descBuffer.MiscFlags = 0;

D3D10_SUBRESOURCE_DATA resData;
memset(&resData, 0, sizeof(resData));
resData.pSysMem = &m_verts[0];
Graphics()->GetDevice()->CreateBuffer(
    &descBuffer, &resData, &m_pVertexBuffer);

descBuffer.Usage = D3D10_USAGE_DEFAULT;
descBuffer.ByteWidth = sizeof(WORD) * NumTris() * 3;
descBuffer.BindFlags = D3D10_BIND_INDEX_BUFFER;
descBuffer.CPUAccessFlags = 0;
descBuffer.MiscFlags = 0;
resData.pSysMem = &m_tris[0];
Graphics()->GetDevice()->CreateBuffer(
    &descBuffer, &resData, &m_pIndexBuffer);
}

void cModel::Scale(float amt)
{
    int size = m_verts.size();
    for(int i=0; i<size; i++)
    {
        D3DXVec3Scale(&m_verts[i].m_vPosition, &m_verts[i].m_vPosition, amt);
    }
}

```

```

void cModel::Draw()
{
    UINT uiStride = sizeof(cGraphicsLayer::cDefaultVertex);
    UINT uiOffset = 0;

    Graphics()->GetDevice()->IASetVertexBuffers(
        0, 1, &m_pVertexBuffer, &uiStride, &uiOffset);
    Graphics()->GetDevice()->IASetIndexBuffer(
        m_pIndexBuffer, DXGI_FORMAT_R16_UINT, 0);
    Graphics()->GetDevice()->IASetPrimitiveTopology(
        D3D10_PRIMITIVE_TOPOLOGY_TRIANGLELIST);

    Graphics()->GetDevice()->DrawIndexed(m_tris.size() * 3, 0, 0);
}

float cModel::GenRadius()
{
    float best = 0.f;
    int size = m_verts.size();
    for(int i=0; i<size; i++)
    {
        float curr = D3DXVec3Length(&m_verts[i].m_vPosition);
        if(curr > best)
            best = curr;
    }
    return best;
}

```

Now that you have a way to load models, a program just needs to be wrapped around it. That is what the D3DSample program does. It takes a filename in the constructor, loads it, creates three colored directional lights, and spins the object around in front of the camera. There is no user input for this program; it's just there to look pretty. See Figure 7.4 for a screenshot of D3DSampleApp in action.

There are a few models in the downloadable files, so you can mess around with them if you want to see what other models look like.

```

/*****
 *      Advanced 3D Game Programming with DirectX 10.0
 *      * * * * *
 *
 *      See license.txt for modification and distribution information
 *      copyright (c) 2007 by Peter Walsh, Wordware
 *****/
#include "stdafx.h"

class cD3DSampleApp : public cApplication
{
public:
    wstring m_filename;
    cModel *m_pModel;

```



```

void InitLights();

//----- cApplication

virtual void DoFrame( float timeDelta );
virtual void SceneInit();

virtual void SceneEnd()
{
    delete m_pModel;
}

cD3DSampleApp() :
    cApplication()
{
    m_title = wstring( L"D3D 10 Sample - Spinning Objects" );
    m_pModel = NULL;
    m_filename = L"..\\BIN\\Media\\rabbit.o3d";
    m_pVertexBuffer = NULL;
}

protected:
    ID3D10Buffer *m_pVertexBuffer;
};

cApplication *CreateApplication()
{
    return new cD3DSampleApp();
}

void DestroyApplication( cApplication *pApp )
{
    delete pApp;
}

void cD3DSampleApp::SceneInit()
{
    /**
     * Create a model with the given filename,
     * and resize it so it fits inside a unit sphere.
     */
    m_pModel = new cModel( m_filename.c_str() );
    m_pModel->Scale(1.0f / m_pModel->GenRadius() );

    InitLights();
}

void cD3DSampleApp::InitLights()
{
    Graphics()->AddLight(D3DXCOLOR(
        1.0f, 0.0f, 0.0f, 1.0f), D3DXVECTOR3(0.0f, 1.0f, 0.0f));
    Graphics()->AddLight(D3DXCOLOR(
        1.0f, 1.0f, 0.0f, 1.0f), D3DXVECTOR3(0.0f, -1.0f, 0.0f));
    Graphics()->AddLight(D3DXCOLOR(
        0.0f, 0.0f, 1.0f, 1.0f), D3DXVECTOR3(0.0f, 1.0f, 1.0f));
}

```

```

Graphics()->AddLight(D3DXCOLOR(
    0.0f, 1.0f, 0.0f, 1.0f), D3DXVECTOR3(0.0f, 1.0f,-1.0f));
}

void cD3DSampleApp::DoFrame( float timeDelta )
{
    if(!Graphics())
        return;

    Graphics()->UpdateMatrices();
    Graphics()->UpdateLights();

    // Clear the previous contents of the back buffer
    float colClear[4] = {0.1f, 0.1f, 0.1f, 1.0f};
    Graphics()->Clear(colClear);
    Graphics()->ClearDepthStencil(1.0f, 0);

    D3D10_TECHNIQUE_DESC descTechnique;
    Graphics()->GetDefaultTechnique()->GetDesc(&descTechnique);
    for(UINT uiCurPass = 0; uiCurPass < descTechnique.Passes; uiCurPass++)
    {
        Graphics()->GetDefaultTechnique()->
            GetPassByIndex(uiCurPass)->Apply(0);
        m_pModel->Draw();
    }

    // Present the back buffer to the primary surface to make it visible
    Graphics()->Present();
}

```

Check out the results in Figure 7.4.




Figure 7.4:
Per-pixel lit 3D
model

Now we'll move on to the next chapter, where we'll build on your foundation of DirectX 10 knowledge.

This page intentionally left blank.

Chapter 8

Advanced 3D Techniques

This is my favorite chapter in the book. Nothing but pure uncut 3D graphics. We're going to take a whirlwind tour of some more advanced topics in 3D programming. Among other things, we'll cover inverse kinematics, subdivision surfaces, and radiosity lighting. This is the most interesting and exciting part of graphics programming—experimenting with cool technology and trying to get it to work well enough to make it into a project.

In this chapter we will look at:

- Hierarchical object animation
- Kinematics
- Parametric curves and surfaces
- Bezier curves and surfaces
- Subdivision
- Progressive meshes
- Radiosity lighting

Animation Using Hierarchical Objects

I wish there were more space to devote to animating our objects, but unfortunately there isn't. Animation is a rich topic, from keyframe animation to motion capture to rotoscoping. I'll just be able to give a sweeping discussion about a few techniques used in animation, then talk about hierarchical objects.

Back in the 2D days, animation was done using *sprites*, which are just bunches of pixels that represent images on the screen. A set of animation frames would be shown in rapid succession to give the illusion of motion. The same technique is used in animated films to give life to the characters.

In 3D, the landscape is much more varied. Some systems use simple extensions from their 2D counterparts. Some games have a complete set of vertex positions for each frame of each animation, which is very similar to 2D games, just replacing pixels with vertices. Newer games move a step further, using interpolation to smoothly morph between frames. This way the playback speed looks good independent of the recording speed; an animation recorded at 10 fps still looks smooth on a 60 fps display.

While systems like this can be very fast (you have to compute, at most, a linear interpolation per vertex), they have a slew of disadvantages. The primary disadvantage is that you must explicitly store each frame of animation in memory. If you have a model with 500 vertices, at 24 bytes (3 floats) per vertex, that's 12 kilobytes of memory needed per frame. If you have several hundred frames of animation, suddenly you're faced with around a megabyte of storage per animated object. In practice, if you have many different types of objects in the scene, the memory requirements become prohibitive.



Note: The memory requirements for each character model in *Quake III: Arena* were so high that the game almost had an eleventh-hour switch over to hierarchical models.

Explicitly placing each vertex in a model each frame is only one solution, and it is lathered in redundancy, as the topology of the model remains about the same. For example, outside of the bending and flexing that occurs at model joints, the relative locations of the vertices in relation to each other stay pretty similar from one frame to the next.

The way humans and other animals move isn't defined by the skin moving around. Your bones are rigid bodies connected by joints that can only bend in certain directions. The muscles in your body are connected to the bones through tendons and ligaments, and the skin sits on top of the muscles. Therefore, the position of your skin is a function of the position of your bones.

This structural paradigm is emulated by bone-based animation. A model is defined once in a neutral position, with a set of bones underlying the structure of the model. All of the vertices in the forearm region of the model are conceptually bound to the forearm bone, and so forth. Instead of explicitly listing a set of vertices per frame for your animation, all this system needs is the orientation of each bone in relation to its parent bone. Typically, the root node is the hip of the model, so the world matrix for the object corresponds to the position of the hip, and the world transformations for each of the other joints are derived from it.




Figure 8.1: Building a hierarchy of rigid objects to make a humanoid

With these orientations you can figure out the layout of each bone of the model, and you use the same transformation matrices. Figuring out the positions of bones, given the angles of the joints, is called *forward kinematics*.

Forward Kinematics

Understanding the way transformations are concatenated is pivotal to understanding forward kinematics. See Chapter 4 for a discussion on this topic if you're rusty on it.

Let's say we're dealing with the simple case of a 2D two-linkage system—an upper arm and a lower arm, with shoulder and elbow joints. We'll define the vertices of the upper arm with a local origin of the shoulder and the vertices sticking out along the x-axis. The lower arm is defined in the same manner, just using the elbow as the local origin. There is a special point in the upper arm that defines where the elbow joint is situated. There is also a point that defines where the shoulder joint is situated relative to the world origin.




Figure 8.2: The untransformed upper- and lower-arm segments

The first task is to transform the points of the upper arm. What you want to do is rotate each of the points about the shoulder axis by the shoulder angle θ_1 , and then translate them so that they are situated relative to the origin. So the transformation becomes:

$$\text{upper-arm transformation} = R_z(\theta_1)T(\text{shoulder_location})$$




Figure 8.3: The result of transforming just the shoulder

Transforming the elbow points is more difficult. Not only are they dependent on the elbow angle θ_2 , but they also depend on the position and angle of their parent: the upper arm and shoulder joint.

We can subdivide the problem to make it easier. If we can transform the elbow points to orient them relative to the origin of the shoulder, then

we can just add to that the transformation for the shoulder to take them to world space. This transformation becomes:

$$\begin{aligned} \text{lower-arm transformation} &= R_Z(\theta_2)T(\text{elbow_location}) \\ &R_Z(\theta_1)T(\text{shoulder_location}) \end{aligned}$$

or

$$\begin{aligned} \text{lower-arm transformation} &= R_Z(\theta_2)T(\text{elbow_location}) \\ \text{upper-arm transformation} & \end{aligned}$$




Figure 8.4: The fully transformed arm

This system makes at least some intuitive sense. Imagine we have some point on the lower arm, initially in object space. The rotation by θ_2 rotates the point about the elbow joint, and the translation moves the point such that the elbow is sticking out to the right of the origin. At this point we have the shoulder joint at the origin, the upper arm sticking out to the right, a jointed elbow, and then our point somewhere on the lower arm. The next transformation we apply is the rotation by θ_1 , which rotates everything up to this point (the lower arm and upper arm) by the shoulder angle. Finally, we apply the transformation to place the shoulder somewhere in the world a little more meaningful than the world space origin.

This system fits into a clean recursive algorithm very well. At each stage of the hierarchy, we compute the transformation that transforms the current joint to the space of the joint above it in the hierarchy, appending it to the front of the current world matrix, and recursing with each of the children.


```

struct hierNode
{
    vert_and_triangle_Data m_data;
    vector< hierNode *children > m_children;
    matrix4 m_matrix;

    void Draw( matrix4 parentMatrix )
    {
        matrix4 curr = m_matrix *parentMatrix;

        // draws the triangles of this node using the provided matrix
        m_data->Draw( curr );
        for( int i=0; i<m_children.size(); i++ )
        {
            m_children[i]->Draw( curr );
        }
    }
};

```

Inverse Kinematics

Forward kinematics takes a set of joint angles and finds the position of the end effector. The inverse of the problem, finding the set of joint angles required to place the end effector in a desired position, is called *inverse kinematics*.

IK is useful in a lot of applications. An example would be having autonomous agents helping the player in a game. During the course of the game, the situation may arise that the autonomous helper needs to press a button, pull a lever, or perform some other action. When this is done without IK, each type of button must be hand-animated by an artist so the agent hits the button accurately. With IK this becomes much easier. The agent just needs to move close enough to it, and find the angles for the shoulder, elbow, and hand to put the pointer finger at the location of the button.

Inverse kinematics is a hard problem, especially when you start solving more complicated cases. It all boils down to degrees of freedom. In all but the simplest case (being a single angular joint and a singular prismatic joint) there are multiple possible solutions for an inverse kinematics system. Take, for example, a shoulder-elbow linkage: two links with two angular joints (shoulder and elbow) and an end effector at the hand. If there is any bend in the arm at all, then there are two possible solutions for the linkage, as evidenced by Figure 8.5.




Figure 8.5: The two joint solutions for a given end effector

These two possible solutions are commonly referred to as elbow up and elbow down. While for this case it's fairly easy to determine the two elbow configurations, it only gets worse. If you had a three-segment linkage, for example, there are potentially an infinite number of solutions to the problem.

There are two ways to go about solving an IK problem. One way is to do it algebraically: The forward kinematics equation gets inverted, the system is solved, and the solution is found. The other way is geometrically: Trigonometric identities and other geometric theorems are used to solve the problem. For more complex IK systems, often a combination of both methods needs to be used. Algebraic manipulation will get you so far toward the solution, then you take what you've gotten thus far and feed it into a geometric solution to get a little further, and so on.

To introduce you to IK, let's solve a simple system: two segments, each with a pivot joint with one degree of freedom. This corresponds closely to a human arm. The base joint is the shoulder, the second joint is the elbow, and the end effector is the wrist. It's a 2D problem, but applying the solution in 3D isn't hard. Ian Davis, a CMU alum currently at Activision, used this type of IK problem to implement autonomous agents in a game. The agents could wander around and help the player. When they wanted to press a button, they moved to the button such that a plane was formed with the arm and button, and then the 2D IK solution was found in the plane.

Being able to solve the two-joint system is also useful in solving slightly more complex systems. If we want to have a third segment (a hand, pivoting at the wrist), there are an infinite number of solutions for most positions that the pointer finger can be in. However, if we force the

hand to be at a particular angle, the problem decomposes into solving a two-segment problem (given the length of the hand and the angle it should be in, the position of the wrist can be found relative to the end effector, and then the wrist, elbow, and shoulder form a solvable two-segment problem).

The two things we'll need to solve the IK problem are two laws of geometry: the law of cosines and the law of sines. They are given in Figure 8.6.




Figure 8.6:
The laws of sines
and cosines

To formally state the problem, we are given as input the lengths of two arm segments L_1 and L_2 , and the desired x,y position of the end effector. We wish to find a valid set of theta angles for the shoulder and elbow joints. The problem configuration appears in Figure 8.7.




Figure 8.7:
The IK problem

We'll be using a bunch of variables to solve the IK problem. They are given in Figure 8.8.




Figure 8.8:
The IK problem
with the variables
we'll use to solve
it

Here is what we do to solve the IK problem, step by step:

1. Find $dist$, using the Pythagorean theorem.
2. Find θ , using the arc-tangent ($\theta = \tan^{-1}(y/x)$).
3. Find θ_{hat} using the law of cosines ($A=dist$, $B=L_1$, $C=L_2$).
4. We can now find the shoulder angle θ_1 by subtracting θ_{hat} from θ .
5. θ_{arm} can be found using the law of cosines as well ($A=L_2$, $B=L_1$, $C=dist$).
6. The elbow angle, θ_2 , is just $\pi - \theta_{\text{arm}}$.

Application: InvKim

To show off inverse kinematics, I wrote a simple application called *InvKim* that solves the two-linkage problem. One of the things that it needs to do is bind the end effector position to the range of possible solutions that can be reached by the arm. The mouse controls a little icon that the end effector always moves toward. You'll notice that the end effector moves at a constant velocity, and the theta angles change to accommodate its movement. When the pointer is in a position that the arm cannot reach, it tries its best to get there, pointing toward its desired goal.




Figure 8.9: The InvKim sample

```
void cIKApp::DrawLinkage()
{
    /**
     * Use the lengths and theta information
     * to compute the forward dynamics of
     * the arm, and draw it.
     */
    cGraphicsLayer::cDefaultVertex box[5];
    cGraphicsLayer::cDefaultVertex joint[20];
    matrix4 rot1, trans1;
    matrix4 rot2, trans2;

    /**
     * create a half circle to give our links rounded edges
     */
    point3 halfCircle[10];
    int i;
    for( i=0; i<10; i++ )
    {
        float theta = (float)i*PI/9.0f;
        halfCircle[i] = point3(
            0.85f * sin( theta ),
            0.85f * cos( theta ),
            0.f );
    }

    rot1.ToZRot( m_theta1 );
    trans1.ToTranslation( point3( m_l1,0, 0 ) );

    rot2.ToZRot( m_theta2 );
```

```

ID3D10Device *pDevice = Graphics()->GetDevice();

/**
 * Make and draw the upper arm
 */
matrix4 shoulderMat = rot1;
for( i=0; i<10; i++ )
{
    point3 temp = halfCircle[i];
    temp.x = -temp.x;

    cGraphicsLayer::cDefaultVertex tempVertex;
    tempVertex.m_vPosition = *(D3DXVECTOR3*)&(shoulderMat * temp);
    tempVertex.m_vNormal = D3DXVECTOR3(0,0,0);
    tempVertex.m_vColor = D3DXCOLOR(1, 0, 0, 1);
    tempVertex.m_TexCoords = D3DXVECTOR2(0,0);

    joint[i] = tempVertex;

    tempVertex.m_vPosition =
        *(D3DXVECTOR3*)&(shoulderMat *
            (halfCircle[i] + point3( m_l1, 0, 0 )));
    tempVertex.m_vColor = D3DXCOLOR(1, 1, 0, 1);

    joint[19-i] = tempVertex;
}

ID3D10Buffer *pJointVertexBuffer = NULL;

D3D10_BUFFER_DESC descBuffer;
memset(&descBuffer, 0, sizeof(descBuffer));
descBuffer.Usage = D3D10_USAGE_DEFAULT;
descBuffer.ByteWidth = sizeof(cGraphicsLayer::cDefaultVertex) * 20;
descBuffer.BindFlags = D3D10_BIND_VERTEX_BUFFER;
descBuffer.CPUAccessFlags = 0;
descBuffer.MiscFlags = 0;

D3D10_SUBRESOURCE_DATA resData;
memset(&resData, 0, sizeof(resData));
resData.pSysMem = joint;

if(descBuffer.ByteWidth == 0)
{
    return;
}

Graphics()->GetDevice()->CreateBuffer(
    &descBuffer, &resData, &pJointVertexBuffer);

if(pJointVertexBuffer)
{
    UINT uiOffset = 0;
    UINT uiStride = sizeof(cGraphicsLayer::cDefaultVertex);

```

```

Graphics()->GetDevice()->IASetVertexBuffers(
    0, 1, &pJointVertexBuffer, &uiStride, &uiOffset);

if(m_pJointVertexBuffer)
{
    m_pJointVertexBuffer->Release();
    m_pJointVertexBuffer = NULL;
}

m_pJointVertexBuffer = pJointVertexBuffer;
Graphics()->GetDevice()->IASetPrimitiveTopology(
    D3D10_PRIMITIVE_TOPOLOGY_LINESTRIP);

D3D10_TECHNIQUE_DESC descTechnique;
Graphics()->GetDefaultTechnique()->GetDesc(&descTechnique);
for(UINT uiCurPass = 0; uiCurPass < descTechnique.Passes; uiCurPass++)
{
    Graphics()->GetDefaultTechnique()->
        GetPassByIndex(uiCurPass)->Apply(0);
    Graphics()->GetDevice()->Draw(20, 0);
}
}

/**
 * Make and draw the lower arm
 */
matrix4 elbowMat = rot2 * transl * rot1;
for( i=0; i<10; i++ )
{
    point3 temp = halfCircle[i];
    temp.x = -temp.x;

    cGraphicsLayer::cDefaultVertex tempVertex;
    tempVertex.m_vPosition = *(D3DXVECTOR3*)&(elbowMat * temp);
    tempVertex.m_vNormal = D3DXVECTOR3(0,0,0);
    tempVertex.m_vColor = D3DXCOLOR(1, 0, 1, 1);
    tempVertex.m_TexCoords = D3DXVECTOR2(0,0);

    joint[i] = tempVertex;

    tempVertex.m_vPosition = *(D3DXVECTOR3*)&
        (elbowMat * (halfCircle[i] + point3( m_l2, 0, 0.f )));
    tempVertex.m_vColor = D3DXCOLOR(0, 0, 1, 1);

    joint[19-i] = tempVertex;
}

ID3D10Buffer *pArmVertexBuffer = NULL;

memset(&descBuffer, 0, sizeof(descBuffer));
descBuffer.Usage = D3D10_USAGE_DEFAULT;
descBuffer.ByteWidth = sizeof(cGraphicsLayer::cDefaultVertex) * 20;
descBuffer.BindFlags = D3D10_BIND_VERTEX_BUFFER;
descBuffer.CPUAccessFlags = 0;

```

```

descBuffer.MiscFlags = 0;

memset(&resData, 0, sizeof(resData));
resData.pSysMem = joint;

if(descBuffer.ByteWidth == 0)
{
    return;
}

Graphics()->GetDevice()->CreateBuffer(
    &descBuffer, &resData, &pArmVertexBuffer);

if(pArmVertexBuffer)
{
    UINT uiOffset = 0;
    UINT uiStride = sizeof(cGraphicsLayer::cDefaultVertex);
    Graphics()->GetDevice()->IASetVertexBuffers(
        0, 1, &pArmVertexBuffer, &uiStride, &uiOffset);

    if(m_pArmVertexBuffer)
    {
        m_pArmVertexBuffer->Release();
        m_pArmVertexBuffer = NULL;
    }

    m_pArmVertexBuffer = pArmVertexBuffer;
    Graphics()->GetDevice()->IASetPrimitiveTopology(
        D3D10_PRIMITIVE_TOPOLOGY_LINESTRIP);

    D3D10_TECHNIQUE_DESC descTechnique;
    Graphics()->GetDefaultTechnique()->GetDesc(&descTechnique);
    for(UINT uiCurPass = 0; uiCurPass < descTechnique.Passes; uiCurPass++)
    {
        Graphics()->GetDefaultTechnique()->
            GetPassByIndex(uiCurPass)->Apply(0);
        Graphics()->GetDevice()->Draw(19, 0);
    }
}

/**
 * Draw a diamond where the mouse is
 */
matrix4 mouseTrans;
mouseTrans.ToTranslation( m_mouse );

cGraphicsLayer::cDefaultVertex tempVertex;

tempVertex.m_vNormal = D3DXVECTOR3(0,0,0);
tempVertex.m_vColor = D3DXCOLOR(0, 1, 1, 1);
tempVertex.m_TexCoords = D3DXVECTOR2(0,0);

tempVertex.m_vPosition =
    *(D3DXVECTOR3*)&(point3(0.5f,0.f,0.f) * mouseTrans);

```



```

box[0] = tempVertex;

tempVertex.m_vPosition =
    *(D3DXVECTOR3*)&(point3(0.f,-0.5f,0.f) * mouseTrans);
box[1] = tempVertex;

tempVertex.m_vPosition =
    *(D3DXVECTOR3*)&(point3(-0.5f,0.f,0.f) * mouseTrans);
box[2] = tempVertex;

tempVertex.m_vPosition =
    *(D3DXVECTOR3*)&(point3(0.f,0.5f,0.f) * mouseTrans);
box[3] = tempVertex;

box[4] = box[0];

ID3D10Buffer *pDiamondVertexBuffer = NULL;

memset(&descBuffer, 0, sizeof(descBuffer));
descBuffer.Usage = D3D10_USAGE_DEFAULT;
descBuffer.ByteWidth = sizeof(cGraphicsLayer::cDefaultVertex) * 5;
descBuffer.BindFlags = D3D10_BIND_VERTEX_BUFFER;
descBuffer.CPUAccessFlags = 0;
descBuffer.MiscFlags = 0;

memset(&resData, 0, sizeof(resData));
resData.pSysMem = box;

if(descBuffer.ByteWidth == 0)
{
    return;
}

Graphics()->GetDevice()->CreateBuffer(
    &descBuffer, &resData, &pDiamondVertexBuffer);

if(pDiamondVertexBuffer)
{
    UINT uiOffset = 0;
    UINT uiStride = sizeof(cGraphicsLayer::cDefaultVertex);
    Graphics()->GetDevice()->IASetVertexBuffers(
        0, 1, &pDiamondVertexBuffer, &uiStride, &uiOffset);

    if(m_pDiamondVertexBuffer)
    {
        m_pDiamondVertexBuffer->Release();
        m_pDiamondVertexBuffer = NULL;
    }

    m_pDiamondVertexBuffer = pDiamondVertexBuffer;
    Graphics()->GetDevice()->IASetPrimitiveTopology(
        D3D10_PRIMITIVE_TOPOLOGY_LINESTRIP);

    D3D10_TECHNIQUE_DESC descTechnique;
    Graphics()->GetDefaultTechnique()->GetDesc(&descTechnique);

```

```

for(UINT uiCurPass = 0; uiCurPass < descTechnique.Passes; uiCurPass++)
{
    Graphics()->GetDefaultTechnique()->
        GetPassByIndex(uiCurPass)->Apply(0);
    Graphics()->GetDevice()->Draw(5, 0);
}
}
}

```

Parametric Curves and Surfaces

Something you may have noticed up to this point is that most of the objects we have been dealing with have been a little on the angular side. We can clearly see the vertices, triangles, and edges that define the boundaries. Objects in the real world, especially organic objects like humans, don't have such sharp definitions. They are curvy to some extent, a trait that is difficult to represent with generic triangle meshes. We can define mathematical entities that allow us to smoothly generate curves (called *splines*) and surfaces (called *patches*). We'll discuss two styles of curves: cubic Bezier and cubic b-spline.



Note: The term spline comes from way, way back, when ships were built from wood. The process of bending planks with weights to form the shape of the hulls of boats is not unlike the math behind curves.

Bezier Curves and Surfaces

A cubic Bezier curve defines a parametric equation that produces a position in space from a given time parameter. Bezier curves can have different degrees, but there are only two that are widely used: quadric and cubic. Quadric curves only use three control points, while cubic curves use four. We'll be covering cubic Bezier curves, so deriving the math for quadric curves won't be difficult once we're through.

Bezier Concepts

Cubic Bezier curves are defined by four points in space. These are called the *control points* of the curve. To avoid confusion, generally lines are drawn between the points to define which way they connect to each other. Figure 8.10 shows an example of four control points that define a cubic Bezier curve.




Figure 8.10:
Four points defining
a control polygon
for a Bezier curve

The actual curve is computed using these four control points by solving an equation with a given t parameter between 0 and 1. At $t=0$, the returned point is sitting at the first control point. At $t=1$, the point is sitting at the last control point. The tangent of the curve (the direction at which the particle moves) at $t=0$ is parallel to the line connecting the first and second control points. For $t=1$, the tangent is parallel to the line connecting the third and fourth control points. During the time in the middle the particle traces a smooth path between these two directions/locations, making a curve that looks like the one in Figure 8.11.




Figure 8.11:
Sample Bezier
curve

With just four points, it's hard to represent too intricate a curve. In order to have anything interesting, we have to combine them to form larger, more complex curves. TrueType fonts are defined this way, as a set of Bezier curves. Actually, they are defined as a special Bezier curve called a quadratic Bezier, which is easier to rasterize as it has only one control point.

But how do we join them? How do we know that the curviness will continue from one curve to the next? This brings up the concept of *continuity*. Bezier curves can meet together in several ways.

The first type of continuity is called C^0 . In this case, the last control point of one curve is in the same position as the first control point of the next curve. Because of this, the particle will go from one curve to the other without a jump. However, remember from before that the positions and distances of the vectors between the first/second and third/fourth control points define the direction of the particle. If the third and fourth control points of one curve are not colinear with the second control point of the next curve, there will be a sharp discontinuity, as shown in Figure 8.12.




Figure 8.12:
Two curves
meeting with C^0
continuity

We can fix this by achieving C^1 continuity. In this case, the second control point of the second curve is colinear with the last two points of the previous curve, but not the same distance from the first control point that the third control point of the previous curve is. Curves with C^1 continuity appear smooth, as shown in Figure 8.13.




Figure 8.13:
Two curves
meeting with C^1
continuity

To make our curves seem totally smooth, we must go for C^2 continuity. To do this, the distance between the third and fourth control points of one curve must be the same direction and same distance apart as the first and second control points of the next one. This puts serious constraints on how

we can model our Bezier surfaces, however. The restrictions we have to impose give us an extremely fair, extremely smooth-looking joint connecting the two curve segments, as shown in Figure 8.14.




Figure 8.14:
Two curves
meeting with C^2
continuity

The Math

Everyone put on your math caps; here comes the fun part. We'll define Bezier curves of degree n parametrically, as a function of t . We can think of t as being the time during the particles' travel. The t variable ranges from 0.0 to 1.0 for each Bezier curve.

$$\mathbf{q}(t) = \sum_{i=0}^n \mathbf{p}_i B_{i,n}(t) \quad 0 \leq t \leq 1$$

where $B_{i,n}(t)$ is the *Bernstein polynomial*:

$$B_{i,n}(t) = \binom{n}{i} (1-t)^{n-i} t^i$$

$$B_{i,n}(t) = \frac{n!}{i!(n-i)!} (1-t)^{n-i} t^i$$

The vector \mathbf{p}_i is control point i .

Let's work out the equations for our cubic ($n=3$) curves. To help with the flow of the derivation we're going to do, we'll expand each equation so it's in the form of $ax^3 + bx^2 + cx + d$.

$$\begin{aligned}
B_{0,3}(t) &= \frac{3!}{0!(3)!} (1-t)^3 t^0 \\
&= (1-t)^3 \\
&= (1-t)(1-2t+t^2) \\
&= (1-2t+t^2) - (t-2t^2+t^3) \\
&= -t^3 + 3t^2 - 3t + 1
\end{aligned}$$

$$\begin{aligned}
B_{1,3}(t) &= \frac{3!}{1!(2)!} (1-t)^2 t^1 \\
&= 3t(1-t)^2 \\
&= 3t(1-2t+t^2) \\
&= 3t^3 - 6t^2 + 3t
\end{aligned}$$

$$\begin{aligned}
B_{2,3}(t) &= \frac{3!}{2!(1)!} (1-t)^1 t^2 \\
&= 3t^2(1-t) \\
&= 3t(1-2t+t^2) \\
&= -3t^3 + 3t^2
\end{aligned}$$

$$\begin{aligned}
B_{3,3}(t) &= \frac{3!}{3!(0)!} (1-t)^0 t^3 \\
&= t^3
\end{aligned}$$

Putting everything together, we get:

$$\begin{aligned}
\mathbf{q}(t) &= \mathbf{p}_0 B_{0,3}(t) + \mathbf{p}_1 B_{1,3}(t) + \mathbf{p}_2 B_{2,3}(t) + \mathbf{p}_3 B_{3,3}(t) \\
\mathbf{q}(t) &= \mathbf{p}_0 (1-t)^3 + \mathbf{p}_1 3t(1-t)^2 + \mathbf{p}_2 3t^2(1-t) + \mathbf{p}_3 t^3
\end{aligned}$$

The equation can be solved using vector mathematics, or we can extract the x, y, and z components of each control point and solve the curve position for that component independently of the other two.

Some insight as to how this equation generates a curve for us comes when we graph the equation. The graph of the four Bernstein blending functions appears in Figure 8.15.




Figure 8.15:
A graph of the
four blending
functions

Note that at $t=0$ we are only influenced by the first control point (the Bernstein for the others evaluates to 0). This agrees with the observation that at $t=0$ our curve is sitting on top of the first control point. The same is true for the last point at $t=1$. Also note that the second and third points never get to contribute completely to the curve (their graphs never reach 1.0), which explains why the middle two control points do not intersect (unless our control points are all colinear, of course).

Finding the Basis Matrix

The equation presented above to find Bezier curve points is a bit clunky, and doesn't fit well into the 3D framework we've set up thus far. Luckily, we can decompose the equation into matrix-vector math, as we'll soon see.

Let us consider each coordinate separately, performing the equations for x , y , and z separately. So when we write p_0 , for example, we're referring to one particular coordinate of the first control vector. If we think about the control points as a vector, we can rewrite the equation as a dot product of two 4D vectors:

$$\mathbf{q}_x(t) = \begin{bmatrix} p_{0x} \\ p_{1x} \\ p_{2x} \\ p_{3x} \end{bmatrix}^T \cdot \begin{bmatrix} (1-t)^3 \\ 3t(1-t)^2 \\ 3t^2(1-t) \\ t^3 \end{bmatrix}^T$$

Note that the equations for y and z are identical, just swapping the corresponding components. We'll exclude the component subscripts for the rest of the equations, but keep them in mind. Now, each term in the second vector is one of the Bernstein terms. Let's fill in their full forms that we figured out above. (I took the liberty of adding a few choice zero terms, to help the logic flow of where we're taking this.)

$$\mathbf{q}(t) = \begin{bmatrix} p_0 \\ p_1 \\ p_2 \\ p_3 \end{bmatrix}^T \cdot \begin{bmatrix} -t^3 + 3t^2 - 3t + 1 \\ 3t^3 - 6t^2 + 3t + 0(1) \\ -3t^3 + 3t^2 + 0t + 0(1) \\ t^3 + 0t^2 + 0t + 0(1) \end{bmatrix}^T$$

Hmmm...well, this is interesting. We have a lot of like terms here. As it turns out, we can represent the right term as the result of the multiplication of a 4×4 matrix and the vector $\langle t^3, t^2, t, 1 \rangle$.

$$\mathbf{q}(t) = \begin{bmatrix} p_0 \\ p_1 \\ p_2 \\ p_3 \end{bmatrix}^T \cdot \left(\begin{bmatrix} t^3 \\ t^2 \\ t \\ 1 \end{bmatrix}^T \times \begin{bmatrix} -1 & 3 & -3 & 1 \\ 3 & -6 & 3 & 0 \\ -3 & 3 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{bmatrix} \right)$$



Note: If you don't follow the jump, go to Chapter 4 to see how we multiply 4×4 and 1×4 matrices together. Try working it out on paper so you see what is happening.

If you've followed along up to this point, pat yourself on the back. We just derived M_B , the basis matrix for Bezier curves:

$$M_B = \begin{bmatrix} -1 & 3 & -3 & 1 \\ 3 & -6 & 3 & 0 \\ -3 & 3 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{bmatrix}$$

Now we're golden: We can find any point $p(t)$ on a Bezier curve. For each component (x, y, and z), we multiply together a vector of those components from the four control points, the vector $\langle t^3, t^2, t, 1 \rangle$, and the basis matrix M_B . We perform the 1D computation for all three axes independently.

Calculating Bezier Curves

So this begs the question of how we render our Bezier curves. Well, the way it's typically done is by stepping across the curve a discrete number of steps, calculating the point along the curve at that point, and then drawing a short line between each pair of points. So our curves are not perfectly curvy (unless we calculate an infinite number of points between $t=0$ and $t=1$, which is a bit on the impossible side). However, we're always bound to some resolution below which we don't really care about. In printing, it's

the dots-per-inch setting of the printer, and in video it's the resolution of the monitor. So if we calculate the curve such that each subline is less than one pixel long, it will appear exactly as the limit curve would, and has much more feasible memory and computational constraints.

Here's the code to do it:

```
matrix4 cBezierPatch::m_basisMatrix = matrix4(
    -1.0f, 3.0f, -3.0f, 1.0f,
    3.0f, -6.0f, 3.0f, 0.0f,
    -3.0f, 3.0f, 0.0f, 0.0f,
    1.0f, 0.0f, 0.0f, 0.0f
);

class cBezierSlowIterator
{
    int     m_i;           // our current step in the iteration
    int     m_nSteps;      // the number of steps
    point4  m_p[3];        // for x, y, and z
    point3  m_cPts[4];

    point3  m_Q;           // Current position
public:
    cBezierSlowIterator(
        int nSteps, point3 p1, point3 p2, point3 p3, point3 p4 )
    {
        m_cPts[0] = p1;
        m_cPts[1] = p2;
        m_cPts[2] = p3;
        m_cPts[3] = p4;

        m_nSteps = nSteps;
        m_p[0].Assign( p1.x, p2.x, p3.x, p4.x );
        m_p[1].Assign( p1.y, p2.y, p3.y, p4.y );
        m_p[2].Assign( p1.z, p2.z, p3.z, p4.z );
    }

    void Start() {
        m_i = 0;
    }

    bool Done() {
        return !(m_i < m_nSteps);
    }

    point3& GetCurr() {
        return m_Q;
    }

    operator point3&() {
        return m_Q;
    }
}
```

```

void CalcNext() {
    float t = (float)m_i / m_nSteps;
    point4 tVec( t*t*t, t*t, t, 1 );
    point4 pVec;

    m_Q.x = m_p[0] * (tVec *cBezierPatch::m_basisMatrix);
    m_Q.y = m_p[1] * (tVec *cBezierPatch::m_basisMatrix);
    m_Q.z = m_p[2] * (tVec *cBezierPatch::m_basisMatrix);
    m_i++;
}
};

```

That code is written for readability; it's terribly slow and it isn't anything anyone would use in production-quality code. Let's write it for speed!

Forward Differencing

When we're computing a linear function (say, color across a polygon) we never try to find the result of the function at each point explicitly like this:

```

int numSteps = 50;
for( int i=0; i<numSteps; i++ )
{
    outColor[i] = FindColor(i);
}

```

Instead, we find the correct value at the first pixel (or point or whatever), and find out how much it will change during each step (this is called a *delta*). Then when we go across, we simply add the delta to the output, like so:

```

int numSteps = 50;
color curr = FindColor(0);
color delta = FindDelta( numSteps );
for( int i=0; i<numSteps; i++ )
{
    outColor[i] = curr;
    curr += delta;
}

```

This can speed up our code a lot because we're replacing the function call with just an addition.

The reason this particular code works is because the function we're interpolating is linear. The graph of the function is a straight line. So if we can compute the slope of the line, we can increment our y by the slope whenever we increment x , and thus we compute $f(x+1)$ in terms of $f(x)$ instead of doing it explicitly.

What about Bezier curves? The delta we would add to the position during each iteration of finding $p(t)$ isn't constant, because it's a cubic function. The first derivative of the curve formed by the Bezier curve isn't a straight line (neither is the second derivative, for that matter). To solve this problem, it is best to use *forward differencing*.

We can define our Bezier equation to be just a regular cubic function (or the dot product of two 4D vectors), like so:

$$\mathbf{q}_x(t) = at^3 + bt^2 + ct + d$$

$$\mathbf{q}_x(t) = \mathbf{t} \bullet \mathbf{c}$$

$$\mathbf{t} = [t^3 \quad t^2 \quad t \quad 1]^T$$

$$\mathbf{c} = [a \quad b \quad c \quad d]^T$$

$$\mathbf{c} = \mathbf{M}_B [\mathbf{p}_{0x} \quad \mathbf{p}_{1x} \quad \mathbf{p}_{2x} \quad \mathbf{p}_{3x}]^T$$

Note that in the above we only define q_x ; q_y and q_z would be essentially identical. For the remainder of the equations, we're going to just abstractly deal with some function $q(t)$ (in code, we'll need to do the work for each component).

Let's define the forward difference as $\Delta q(t)$. The forward difference is defined such that when we add it to $q(t)$, we get the next point in the iteration (the point we get after we increment t by the small inter-step delta value δ). That is...

$$\mathbf{q}(t + \delta) = \mathbf{q}(t) + \Delta \mathbf{q}(t)$$

So now we just need to find $\Delta q(t)$. Don't forget that d is a constant, based on the number we wish to tessellate ($\delta = 1/\text{size}$). Let's get the math down:

$$\Delta \mathbf{q}(t) = \mathbf{q}(t + \delta) - \mathbf{q}(t)$$

$$\Delta \mathbf{q}(t) = a(t + \delta)^3 + b(t + \delta)^2 + c(t + \delta) + d - (at^3 + bt^2 + ct + d)$$

$$\Delta \mathbf{q}(t) = 3at^2\delta + t(3a\delta^2 + 2b\delta) + a\delta^3 + b\delta^2 + c\delta$$

Unfortunately, $\Delta q(t)$ is a function of t , so we would need to calculate it explicitly each iteration. All we've done is add extra work. However, $\Delta q(t)$ is a quadratic equation where $q(t)$ was cubic, so we've improved a bit. Let's calculate the forward difference of $\Delta q(t)$ (that is, $\Delta^2 q(t)$).

$$\Delta^2 \mathbf{q}(t) = \Delta \mathbf{q}(t + \delta) - \Delta \mathbf{q}(t)$$

$$\Delta^2 \mathbf{q}(t) = 3a(t + \delta)^2\delta + (t + \delta)(3a\delta^2 + 2b\delta) + a\delta^3 + b\delta^2 + c\delta -$$

$$(3at^2\delta + t(3a\delta^2 + 2b\delta) + a\delta^3 + b\delta^2 + c\delta)$$

$$\Delta^2 \mathbf{q}(t) = 6at\delta^2 + 6a\delta^3 + 2b\delta^2$$

We're almost there. While $\Delta^2 q(t)$ still is a function of t , this time it's just a linear equation. We just need to do this one more time and calculate

$\Delta^3 q(t)$:

$$\Delta^3 \mathbf{q}(t) = \Delta^2 \mathbf{q}(t + \delta) - \Delta^2 \mathbf{q}(t)$$

$$\Delta^3 \mathbf{q}(t) = 6a(t + \delta)\delta^2 + 6a\delta^3 + 2b\delta^2 - (6at\delta^2 + 6a\delta^3 + 2b\delta^2)$$

$$\Delta^3 \mathbf{q}(t) = 6a\delta^3$$

Eureka! A constant! If you don't share my exuberance, hold on. Let's suppose that at some point along the curve we know $q(t)$, $\Delta q(t)$, $\Delta^2 q(t)$, and $\Delta^3 q(t)$. This will hold true at the initial case when $t=0$: We can explicitly compute all four variables. To arrive at the next step in the iteration, we just do:

$$\begin{aligned} \mathbf{q}(t+\delta) &= \mathbf{q}(t) + \Delta \mathbf{q}(t) \\ \Delta \mathbf{q}(t+\delta) &= \Delta \mathbf{q}(t) + \Delta^2 \mathbf{q}(t) \\ \Delta^2 \mathbf{q}(t+\delta) &= \Delta^2 \mathbf{q}(t) + \Delta^3 \mathbf{q}(t) \end{aligned}$$

As you can see, it's just a bunch of additions. All we need to do is keep track of everything. Suddenly, we only need to do hard work during setup; calculating n points is next to free.

The *cFwdDiffIterator* Class

The *cFwdDiffIterator* class implements the equations listed above to perform forward differencing. Compare and contrast the equations and the code until they make sense.

```
class cFwdDiffIterator
{
    int    m_i;           // our current step in the iteration
    int    m_nSteps;      // the number of steps

    point3 m_p[4];        // The 4 control points

    point3 m_Q;           // the point at the current iteration location
    point3 m_dQ;          // First derivative (initially at zero)
    point3 m_ddQ;         // Second derivative (initially at zero)
    point3 m_dddQ;        // Triple derivative (constant)

public:
    cFwdDiffIterator()
    {
        // Do nothing
    }
    cFwdDiffIterator(
        int nSteps,

        point3 p1,
        point3 p2,
        point3 p3,
        point3 p4 )
    {
        m_nSteps = nSteps;
        m_p[0] = p1;
        m_p[1] = p2;
        m_p[2] = p3;
        m_p[3] = p4;
    }
}
```

```

void Start()
{
    m_i = 0;

    float d = 1.f/(m_nSteps-1);
    float d2 = d*d; // d^2
    float d3 = d*d2; // d^3

    point4 px( m_p[0].x, m_p[1].x, m_p[2].x, m_p[3].x );
    point4 py( m_p[0].y, m_p[1].y, m_p[2].y, m_p[3].y );
    point4 pz( m_p[0].z, m_p[1].z, m_p[2].z, m_p[3].z );

    point4 cVec[3]; // <a, b, c, d> for x, y, and z.
    cVec[0] = px *cBezierPatch::m_basisMatrix;
    cVec[1] = py *cBezierPatch::m_basisMatrix;
    cVec[2] = pz *cBezierPatch::m_basisMatrix;

    m_Q = m_p[0];

    // Do the work for each component
    int i = 3;
    while (i--)
    {
        // remember that t=0 here so many of the terms
        // in the text drop out.
        float a = cVec[i].v[0];
        float b = cVec[i].v[1];
        float c = cVec[i].v[2];
        // luckily d isn't used, which
        // would clash with the other d.

        m_dQ.v[i] = a * d3 + b * d2 + c * d;
        m_ddQ.v[i] = 6 * a * d3 + 2 * b * d2;
        m_dddQ.v[i] = 6 * a * d3;
    }
}

bool Done()
{
    return !(m_i < m_nSteps);
}

point3& GetCurr()
{
    return m_Q;
}

operator point3&()
{
    return m_Q;
}

void CalcNext()
{
    m_Q += m_dQ;
}

```

```

        m_dQ += m_ddQ;
        m_ddQ += m_dddQ;

        m_i++;
    }
};

```

Drawing Curves

Armed with our fast forward difference iterator, drawing curves isn't difficult at all. All we need to do is step across the Bezier curve, sample the curve point at however many locations desired, and draw the data, either as a point list or a line strip. Check out this pseudocode:

```

void DrawCurve(
    const point3& c1,
    const point3& c2,
    const point3& c3,
    const point3& c4 )
{
    // we can tessellate to any level of detail we want, but for the
    // sake of example let's generate 50 points (49 line segments)
    Vertex v[50];
    cFwdDiffIterator iter( 50, c1, c2, c3, c4 );
    int curr = 0;
    for( iter.Start(); !iter.Done(); iter.CalcNext() )
    {
        v[curr++] = Vertex( iter.GetCurr(), 0x00FFFFFF );
    }

    DrawLineStrip(49, v,);
}

```

Drawing Surfaces

While curves are swell and all, what we really want to do is draw curved surfaces. Luckily, we're not far away from being able to do that. Instead of four control points, we're going to have 16. We define a 4×4 grid of points that will form a 3D surface, called a *patch*. A simple patch appears in Figure 8.16.




Figure 8.16:
A control net for
a simple patch

So instead of the function $q(t)$ we had before, now we have a new function $q(s,t)$ that gives the point along the surface for the two inputs $([0,1],[0,1])$. The four corners of our patch are $(0,0)$, $(1,0)$, $(1,1)$, and $(0,1)$. In practice, it would be possible to just iterate across the entire surface with two for loops, calculating the point using the two-dimensional function. However, we can exploit the code written previously for calculating curves.

We can think of the patch as a series of n curves put side to side to give the impression of a surface. If we step evenly along all n curves m times, we will create an $m \times n$ grid of point values. We can use the same forward differencing code we wrote before to step m times along each of these n curves. All we need is the four control points that define each of the n curves.

No problem. We can think of the 16 control points as four sets of control points describing four vertical curves. We simultaneously step n times along each of these four curves. Each of the n iterated points is a control point for a horizontal curve. We take the four iterated points from the four curves and use that to create a horizontal curve, which we iterate across n times. We use our forward differencing code here, too. An example of a Bezier patch appears in Figure 8.17.




Figure 8.17:
The control net
with the tessellated mesh

Application: Teapot

The application for this section is a viewer to display objects composed of Bezier surfaces. A Bezier object (represented by the class `cBezierObject`) holds on to a bunch of separate Bezier patches. To show off the code, we'll use the canonical Bezier patch surface: the Utah Teapot.

The teapot is used in graphics so often it's transcended to the point of being a basic geometric primitive (in both 3ds Max and Direct3D). It's a set of 28 patches that define an object that looks like a teapot. Figure 8.18 shows what the control nets for the patches look like, and Figure 8.19 shows the tessellated Bezier patches.




Figure 8.18:
Control nets for
the Bezier
patches of the
teapot model




Figure 8.19:
Tessellated tea-
pot model

The Bezier patches are loaded from an ASCII `.bez` file. The first line gives the number of patches, and then each patch is listed on four lines (four control points per line). Each patch is separated by a line of white space.

One thing we haven't discussed yet is how to generate normals for our patches. We need to define a vertex normal for each vertex in the grid if we want it to be lit correctly. One way to do it would be to compute them the same way as for regular polygonal surfaces (that is, find normals for each of the triangles, and average the normals of all the triangles that share a vertex, making that the vertex normal). This won't work, however,

at least not completely. The normals of the edges will be biased inward a little (since they don't have the triangles of adjacent patches contributing to the average). This will cause our patches to meet up incorrectly, causing a visual seam where adjacent normals are different.

A better way to calculate Bezier patch normals is to generate them explicitly from the definition of the curve. When we compute the Bezier function using the t -vector $\langle t^3, t^2, t, 1 \rangle$, we compute the position along the curve. If we instead use the first derivative of the Bezier function, we will get the tangent at each point instead of the position. To compute the derivative, we just use a different t -vector, where each component is the derivative of the component in the regular t -vector. This gives us the vector $\langle 3t^2, 2t, 1, 0 \rangle$.

To do this I threw together a quick and dirty iterator class (called `cTangentIterator`) that uses the slow matrix multiplication method to calculate the tangent vectors. Converting the iterator to use forward differencing would not be hard, and is left as an exercise for the reader.

We step across the patch one way and find the position and u -tangent vector. We then step across the perpendicular direction, calculating the v -tangent vectors. Then we cross product the two tangent vectors to get a vector perpendicular to both of them (which is the normal we want). We use the position and normal to build the vertex list. Then when we draw, it's just one `DrawIndexedPrimitive()` call. There's too much code in the project to list here, so I'll just show you the interesting parts.

```
void cBezierPatch::Init( int size )
{
    delete [] m_vertList;
    delete [] m_triList;
    delete [] m_uTangList;
    delete [] m_vTangList;

    m_size = size;

    // allocate our lists
    m_vertList = new cGraphicsLayer::cDefaultVertex[ size * size ];
    m_triList = new sTri[ (size-1) * (size-1) * 2 ];

    m_uTangList = new point3[ size * size ];
    m_vTangList = new point3[ size * size ];

    Tessellate();
}

/**
 * Fill in the grid of values (all the dynamic arrays
 * have been initialized already). The grid is of
 * size mxn where m = n = m_size
 */
void cBezierPatch::Tessellate()
{
    int u, v; // patch-space coordinates.
```

```

point3 p1,p2,p3,p4;

/**
 * These are the four curves that will define the control points for the
 * rest of the curves
 */
cFwdDiffIterator mainCurve1;
cFwdDiffIterator mainCurve2;
cFwdDiffIterator mainCurve3;
cFwdDiffIterator mainCurve4;

int nSteps = m_size;
mainCurve1 = cFwdDiffIterator( nSteps, m_ctrlPoints[0], m_ctrlPoints[4],
    m_ctrlPoints[8], m_ctrlPoints[12] );
mainCurve2 = cFwdDiffIterator( nSteps, m_ctrlPoints[1], m_ctrlPoints[5],
    m_ctrlPoints[9], m_ctrlPoints[13] );
mainCurve3 = cFwdDiffIterator( nSteps, m_ctrlPoints[2], m_ctrlPoints[6],
    m_ctrlPoints[10], m_ctrlPoints[14] );
mainCurve4 = cFwdDiffIterator( nSteps, m_ctrlPoints[3], m_ctrlPoints[7],
    m_ctrlPoints[11], m_ctrlPoints[15] );

mainCurve1.Start();
mainCurve2.Start();
mainCurve3.Start();
mainCurve4.Start();

for(v=0;v<m_size;v++)
{
    /**
     * Generate our four control points for this curve
     */
    p1 = mainCurve1.GetCurr();
    p2 = mainCurve2.GetCurr();
    p3 = mainCurve3.GetCurr();
    p4 = mainCurve4.GetCurr();

    /**
     * Now step along the curve filling in the data
     */
    cTangentIterator tanIter( nSteps, p1, p2, p3, p4 );
    tanIter.Start();
    cFwdDiffIterator iter( nSteps, p1, p2, p3, p4 );
    u = 0;
    for(
        iter.Start(); !iter.Done(); iter.CalcNext(), u++ )
    {
        m_vertList[m_size*v+u].m_vPosition =
            *(D3DXVECTOR3*)&iter.GetCurr();
        m_vertList[m_size*v+u].m_TexCoords = D3DXVECTOR2(0,0);
        m_vertList[m_size*v+u].m_vColor = D3DXCOLOR(0, 0, 0, 1);
        m_vertList[m_size*v+u].m_vNormal = D3DXVECTOR3(0,0,0);

        // We're piggybacking our u-direction tangent vector calculation here.
        m_uTangList[m_size*v+u] = tanIter.GetCurr();
        tanIter.CalcNext();
    }
}

```

```

    }

    mainCurve1.CalcNext();
    mainCurve2.CalcNext();
    mainCurve3.CalcNext();
    mainCurve4.CalcNext();
}

/**
 * Since we can't generate the v-tangents in the same run as the u-tangents
 * (we need to go
 * in the opposite direction), we have to go through the process again, but
 * this time in the
 * perpendicular direction we went the first time
 */
mainCurve1 = cFwdDiffIterator( nSteps, m_ctrlPoints[0], m_ctrlPoints[1],
    m_ctrlPoints[2], m_ctrlPoints[3] );
mainCurve2 = cFwdDiffIterator( nSteps, m_ctrlPoints[4], m_ctrlPoints[5],
    m_ctrlPoints[6], m_ctrlPoints[7] );
mainCurve3 = cFwdDiffIterator( nSteps, m_ctrlPoints[8], m_ctrlPoints[9],
    m_ctrlPoints[10], m_ctrlPoints[11] );
mainCurve4 = cFwdDiffIterator( nSteps, m_ctrlPoints[12], m_ctrlPoints[13],
    m_ctrlPoints[14], m_ctrlPoints[15] );

mainCurve1.Start();
mainCurve2.Start();
mainCurve3.Start();
mainCurve4.Start();

for(v=0;v<m_size;v++)
{
    // create a horizontal Bezier curve by calc'ing points along the 4
    // vertical ones

    p1 = mainCurve1.GetCurr();
    p2 = mainCurve2.GetCurr();
    p3 = mainCurve3.GetCurr();
    p4 = mainCurve4.GetCurr();

    cTangentIterator iter( nSteps, p1, p2, p3, p4 );
    u = 0;
    for( iter.Start(); !iter.Done(); iter.CalcNext(), u++ )
    {
        // We don't get the location because all we
        // want here is the v-tangents
        m_vTangList[m_size*u+v] = iter.GetCurr();
    }

    mainCurve1.CalcNext();
    mainCurve2.CalcNext();
    mainCurve3.CalcNext();
    mainCurve4.CalcNext();
}

int offset;

```

```

for(v=0;v<m_size;v++)
{
    // tessellate across the horizontal Bezier
    for(u=0;u<m_size;u++)
    {
        offset = m_size*v+u;

        point3 norm;
        norm = m_vTangList[offset] ^ m_uTangList[offset];
        norm.Normalize();

        m_vertList[offset].m_vNormal = *(D3DXVECTOR3*)&norm;
        m_vertList[offset].m_TexCoords.x = 0;
        m_vertList[offset].m_TexCoords.y = 0;
    }
}

// use an incremented pointer to the triangle list
sTri *pCurrTri = m_triList;

// build the tri list
for( v=0; v< (m_size-1); v++ )
{
    for( u=0; u< (m_size-1); u++ )
    {
        // tessellating square [u,v]

        // 0, 1, 2
        pCurrTri->v[0] = m_size*(v+0) + (u+0);
        pCurrTri->v[1] = m_size*(v+0) + (u+1);
        pCurrTri->v[2] = m_size*(v+1) + (u+1);
        pCurrTri++;

        // 2, 3, 0
        pCurrTri->v[0] = m_size*(v+1) + (u+1);
        pCurrTri->v[1] = m_size*(v+1) + (u+0);
        pCurrTri->v[2] = m_size*(v+0) + (u+0);
        pCurrTri++;
    }
}

void cBezierPatch::Draw( bool bDrawNet )
{
    static bool bFirst = true;

    // hard code the control mesh lines
    static short netIndices[] = {
        0, 1, 1, 2, 2, 3, 4, 5, 5, 6, 6, 7,
        8, 9, 9, 10, 10, 11, 12, 13, 13, 14, 14, 15,
        0, 4, 4, 8, 8, 12, 1, 5, 5, 9, 9, 13,
        2, 6, 6, 10, 10, 14, 3, 7, 7, 11, 11, 15 };

```

```

static ID3D10Buffer *pNetIndexBuffer = NULL;
static ID3D10Buffer *pNetVertexBuffer = NULL;
if(bFirst)
{
    D3D10_BUFFER_DESC descBuffer;
    memset(&descBuffer, 0, sizeof(descBuffer));
    descBuffer.Usage = D3D10_USAGE_DEFAULT;
    descBuffer.ByteWidth = sizeof(netIndices);
    descBuffer.BindFlags = D3D10_BIND_INDEX_BUFFER;
    descBuffer.CPUAccessFlags = 0;
    descBuffer.MiscFlags = 0;

    D3D10_SUBRESOURCE_DATA resData;
    memset(&resData, 0, sizeof(resData));
    resData.pSysMem = netIndices;

    Graphics()->GetDevice()->CreateBuffer(
        &descBuffer, &resData, &pNetIndexBuffer);

    bFirst = false;
}

if(!pNetIndexBuffer)
    return;

D3D10_BUFFER_DESC descBuffer;
memset(&descBuffer, 0, sizeof(descBuffer));
descBuffer.Usage = D3D10_USAGE_DEFAULT;
descBuffer.ByteWidth = sizeof(cGraphicsLayer::cDefaultVertex) * 16;
descBuffer.BindFlags = D3D10_BIND_VERTEX_BUFFER;
descBuffer.CPUAccessFlags = 0;
descBuffer.MiscFlags = 0;

cGraphicsLayer::cDefaultVertex v[16];
for( int i=0; i<16; i++ )
{
    cGraphicsLayer::cDefaultVertex tempVertex;
    tempVertex.m_TexCoords = D3DXVECTOR2(0,0);
    tempVertex.m_vColor = D3DXCOLOR(0.5f, 0.5f, 0.5f, 1.0f);
    tempVertex.m_vNormal = D3DXVECTOR3(0,0,0);
    tempVertex.m_vPosition = *(D3DXVECTOR3*)&m_ctrlPoints[i];

    v[i] = tempVertex;
}

D3D10_SUBRESOURCE_DATA resData;
memset(&resData, 0, sizeof(resData));
resData.pSysMem = v;

Graphics()->GetDevice()->CreateBuffer(
    &descBuffer, &resData, &pNetVertexBuffer);

```

```

if( bDrawNet )
{
    UINT uiStride = sizeof(cGraphicsLayer::cDefaultVertex);
    UINT uiOffset = 0;
    Graphics()->GetDevice()->IASetVertexBuffers(
        0, 1, &pNetVertexBuffer, &uiStride, &uiOffset);
    Graphics()->GetDevice()->IASetIndexBuffer(
        pNetIndexBuffer, DXGI_FORMAT_R16_UINT, 0);
    Graphics()->GetDevice()->IASetPrimitiveTopology(
        D3D10_PRIMITIVE_TOPOLOGY_LINELIST);

    if(m_pNetVertexBuffer)
        m_pNetVertexBuffer->Release();

    m_pNetVertexBuffer = pNetVertexBuffer;

    D3D10_TECHNIQUE_DESC descTechnique;
    Graphics()->GetDefaultTechnique()->GetDesc(&descTechnique);
    for(UINT uiCurPass = 0; uiCurPass < descTechnique.Passes; uiCurPass++)
    {
        Graphics()->GetDefaultTechnique()->
            GetPassByIndex(uiCurPass)->Apply(0);
        Graphics()->GetDevice()->DrawIndexed(
            sizeof(netIndices) / sizeof(short), 0, 0);
    }
}

ID3D10Buffer *pMainIndexBuffer = NULL;
ID3D10Buffer *pMainVertexBuffer = NULL;

memset(&descBuffer, 0, sizeof(descBuffer));
descBuffer.Usage = D3D10_USAGE_DEFAULT;
descBuffer.ByteWidth = sizeof(WORD) * 2 * (m_size-1) * (m_size-1) * 3;
descBuffer.BindFlags = D3D10_BIND_INDEX_BUFFER;
descBuffer.CPUAccessFlags = 0;
descBuffer.MiscFlags = 0;

memset(&resData, 0, sizeof(resData));
resData.pSysMem = m_triList;

Graphics()->GetDevice()->CreateBuffer(
    &descBuffer, &resData, &pMainIndexBuffer);

memset(&descBuffer, 0, sizeof(descBuffer));
descBuffer.Usage = D3D10_USAGE_DEFAULT;
descBuffer.ByteWidth =
    sizeof(cGraphicsLayer::cDefaultVertex) * m_size * m_size;
descBuffer.BindFlags = D3D10_BIND_VERTEX_BUFFER;
descBuffer.CPUAccessFlags = 0;
descBuffer.MiscFlags = 0;

memset(&resData, 0, sizeof(resData));
resData.pSysMem = m_vertList;

```

```

Graphics()->GetDevice()->CreateBuffer(
    &descBuffer, &resData, &pMainVertexBuffer);

UINT uiStride = sizeof(cGraphicsLayer::cDefaultVertex);
UINT uiOffset = 0;
Graphics()->GetDevice()->IASetVertexBuffers(
    0, 1, &pMainVertexBuffer, &uiStride, &uiOffset);
Graphics()->GetDevice()->IASetIndexBuffer(
    pMainIndexBuffer, DXGI_FORMAT_R16_UINT, 0);
Graphics()->GetDevice()->IASetPrimitiveTopology(
    D3D10_PRIMITIVE_TOPOLOGY_TRIANGLELIST);

if(m_pMainVertexBuffer)
    m_pMainVertexBuffer->Release();
if(m_pMainIndexBuffer)
    m_pMainIndexBuffer->Release();

m_pMainVertexBuffer = pMainVertexBuffer;
m_pMainIndexBuffer = pMainIndexBuffer;

D3D10_TECHNIQUE_DESC descTechnique;
Graphics()->GetDefaultTechnique()->GetDesc(&descTechnique);
for(UINT uiCurPass = 0; uiCurPass < descTechnique.Passes; uiCurPass++)
{
    Graphics()->GetDefaultTechnique()->
        GetPassByIndex(uiCurPass)->Apply(0);
    Graphics()->GetDevice()->DrawIndexed(
        2 * (m_size-1) * (m_size-1) * 3, 0, 0);
}
}

```

B-Spline Curves

Because there are myriad other types of parametric curves and surfaces, we could not hope to cover them all. However, we'll quickly cover one more type of curve, b-spline, before moving on to subdivision surfaces.

Uniform, rational b-splines are quite different from Bezier curves. Rather than have a set of distinct curves, each one made up of four control points, a b-spline is made up of any number of control points (well... any number greater than four). They are C^2 continuous, but they are not interpolative (they don't pass through their control points).

Given a particular control point p_i , we iterate from $t=0$ to $t=1$. The iteration uses the four control points $(p_i, p_{i+1}, p_{i+2}, p_{i+3})$. The curve it steps out sits between p_{i+1} and p_{i+2} , but note that the curve itself probably won't actually go through those points. Figure 8.20 may help you understand this.




Figure 8.20:
Sample b-spline

Each section of the curve (denoted by s_0 , s_1 , etc.) is traced out by the four control points around it. Segment s_0 is traced out by p_0 – p_3 , segment s_1 by p_1 – p_4 , and so on. To compute a point along a b-spline, we use the following equation:

$$\mathbf{q}(t) = \begin{bmatrix} t^3 \\ t^2 \\ t \\ 1 \end{bmatrix}^T \frac{1}{6} \begin{bmatrix} -1 & 3 & -3 & 1 \\ 3 & -6 & 3 & 0 \\ -3 & 0 & 3 & 0 \\ 1 & 4 & 1 & 0 \end{bmatrix} \begin{bmatrix} \mathbf{p}_i \\ \mathbf{p}_{i+1} \\ \mathbf{p}_{i+2} \\ \mathbf{p}_{i+3} \end{bmatrix}$$

The main reason I'm including b-spline curves in this chapter is just to show you that once you've learned one style of parametric curve, you've pretty much learned them all. Almost all use the same style of equation; it's just a matter of choosing the kind of curve you want and plugging it into your code.

Application: BSpline

Just for fun, I threw together a simple application to show off b-splines. It draws a set of six splines whose tails fade off to blackness, spinning around in space. The code running the splines is pretty rudimentary; it's just there to hopefully spark an idea in your head to use them for something more complex. As simple as the code is, it can be pretty mesmerizing, and I feel it's one of the more visually pleasing sample applications in this book.

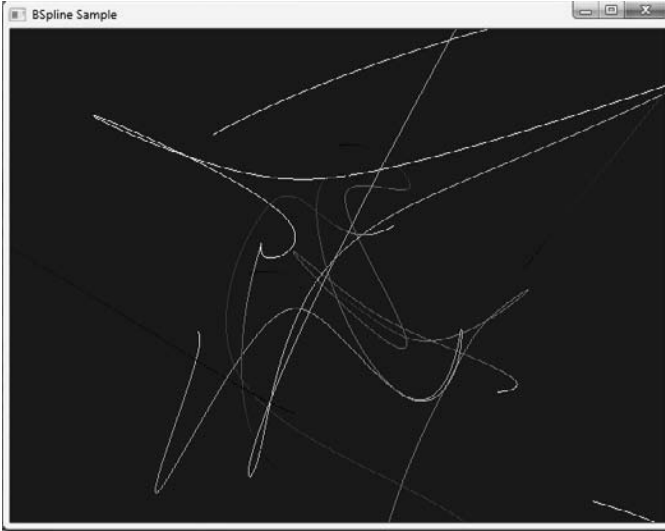


Figure 8.21: The BSpline sample

Check out the code below, which calculates the points along the curve. You can see the full code listing in the downloadable files.

```
/**
 * The b-spline basis matrix
 */
matrix4 cBSpline::m_baseMatrix = matrix4(
    -1, 3, -3, 1,
    3, -6, 3, 0,
    -3, 0, 3, 0,
    1, 4, 1, 0);

point3 cBSpline::Calc( float t, int i0 )
{
    assert(i0+3 < m_ctrlPoints.size() );
    assert(t>=0.f && t<=1.f );
    point4 tVec( t*t*t, t*t, t, 1 );

    point4 xVec(
        m_ctrlPoints[i0].x,
        m_ctrlPoints[i0+1].x,
        m_ctrlPoints[i0+2].x,
        m_ctrlPoints[i0+3].x );
    point4 yVec(
        m_ctrlPoints[i0].y,
        m_ctrlPoints[i0+1].y,
        m_ctrlPoints[i0+2].y,
        m_ctrlPoints[i0+3].y );
    point4 zVec(
        m_ctrlPoints[i0].z,
        m_ctrlPoints[i0+1].z,
```

```

        m_ctrlPoints[i0+2].z,
        m_ctrlPoints[i0+3].z );

    return point3(
        tVec * (1.f/6) * m_baseMatrix * xVec,
        tVec * (1.f/6) * m_baseMatrix * yVec,
        tVec * (1.f/6) * m_baseMatrix * zVec );
}

point3 cBSpline::CalcAbs( float t )
{
    // the T we get isn't right, fix it.
    t *= m_ctrlPoints.size() - 3;
    int vert = (int)(floor(t));
    t -= (float)floor(t);
    return Calc( t, vert );
}

```

Subdivision Surfaces

Parametric surfaces, while really cool, are not without their problems. The main problem is in order to have smoothness, it's usually necessary to keep the valence at patch corners equal to 2 or 4. (That is, at any patch corner there is either one more or three more patches also touching that corner.) Otherwise, the patches don't meet up correctly and there's a seam in the surface. This can be fixed by using *degenerate patches* (patches that are really triangles); however, getting some things to look right (like the meeting point of a handle and a mug) can prove downright maddening.

Subdivision surfaces try to get around this restriction by attacking the problem of creating smooth surfaces a different way. They use a discrete operation that takes a given mesh and subdivides it. If the resultant mesh is subdivided again and again, eventually the surface reaches the *limit surface*. Most subdivision schemes have a limit surface that has C^1 continuity, which is generally all we need for games. You don't need to go all the way to the limit surface, however; each time you subdivide your surface looks smoother and smoother.

Subdivision surfaces have gotten a lot of press in the computer graphics community. Mostly this is because they're fairly straightforward to code, easy to use by artists, and very *very* cool looking. The first mainstream media to use subdivision surfaces was *Geri's Game*, a short by Pixar. The piece won, among other things, an Academy Award for Best Animated Short.

Subdivision Essentials

To begin the discussion of subdivision curves and surfaces, we'll consider a simple 2D case: subdividing a curve. Once we learn how to subdivide that, we can start experimenting with surfaces. Our lowest resolution curve, the control curve, appears in Figure 8.22.




Figure 8.22:
A sample
five-segment
loop

Luckily for us, this is a closed loop, so for our first baby steps we don't need to trip over boundary cases. Let's define an operation that we can perform on our curve and call it an edge split. It takes some particular edge from p_n to p_{n+1} . The edge is subdivided into two new edges. The location of the new internal point (we'll call it $p_{n+0.5}$) depends on the neighborhood of points around it. We want to position the new internal point such that it fits on the curve defined by the points around it.

The formula we'll define to calculate $p_{n+0.5}$ is the following:

$$p_{n+0.5} = -\frac{1}{16}p_{n-1} + \frac{9}{16}p_n + \frac{9}{16}p_{n+1} - \frac{1}{16}p_{n+2}$$

This equation, after some reflection, seems pretty intuitive. Most of the position of the new point is an average of the two points adjacent to it. Then, to perturb it a little, we move it away from the points one hop from it on either side by a small amount. Note that the set of constant values (called the *mask*) all add up to 1.

If we apply the equation to each of the edges in the control curve, we get a new curve. We'll call this curve the level 1 curve. It rears its not-so-ugly head in Figure 8.23.




Figure 8.23:
The loop after
one subdivision

You'll notice that after the subdivision step, we doubled the number of edges in our loop, and our loop got a tiny bit smoother. If we apply it again (shown in Figure 8.24), it gets still smoother.




Figure 8.24:
The loop after
two subdivisions

It's fairly easy to see that eventually this little fella will be about as smooth as we can possibly deal with. How smooth we go depends on the application. If we were Pixar, and we were making the new animated short *Geri's Curve*, we could afford to subdivide it such that all of our line segments are half a pixel wide. Any smoother than that is truly excessive, and even taking it to that level is infeasible for current generation real-time 3D graphics.

Handling surfaces is just as easy. You start out with a control mesh (in some cases this is just a regular triangular model) and each subdivision creates a more tessellated mesh. The beauty is you can decide how much to subdivide based on how much hardware is available to do the work. If someone picks up your game eight years from now, your code could automatically take advantage of the multi-quadrillion triangle rate and subdivide your curves and surfaces from here to kingdom come.

This is the driving idea behind all subdivision surface schemes: They all derive their identity from small little differences. Let's take a look at some of the differences before we decide upon a method to implement.

Triangles vs. Quads

One of the most obvious differences between subdivision schemes is the type of primitive they operate on. Some schemes, such as Catmull-Clark subdivision, operate with control meshes of quadrilaterals. Others, like butterfly subdivision, instead work with triangle meshes.

Using a subdivision mesh based on triangles has a lot of advantages over quadrilateral methods. First of all, most modeling programs can easily create meshes built out of triangles. Making one exclusively out of quadrilaterals can be considerably more difficult, and has a lot of the same problems that arise from attempting to build complex objects out of Bezier patches. Also, being able to use triangle meshes is a big plus because you may be adding subdivision surfaces to an existing project that uses regular triangle models; you won't need to do any work converting your existing media over to a subdividing system.

Interpolating vs. Approximating

After we've decided what primitive our subdivision meshes should be based on, we need to decide if the method we want to implement should be *interpolating* or *approximating*. They define how the new control mesh is reached from the original.

With approximating subdivision, the limit mesh is actually never reached by the vertices, unless the surface is subdivided an infinite number of times. Each time a subdivision is performed, the old mesh is completely thrown away and a new mesh is created that is a bit closer to the limit curve. As subdivisions continue, the surface moves closer and closer to the limit surface, looking more and more like it. This has a few side effects. The primary one is that the initial control mesh tends not to look much like the limit surface at all. Modifying the initial mesh to get the desired result in the limit mesh isn't easy. However, for giving up a bit of intuitive control, you generally get a much nicer-looking mesh and have fewer strange-looking subdivided areas.

Interpolating subdivision, on the other hand, always adds vertices right on the limit surface. The initial control mesh is on the limit surface, each new batch of vertices and triangles we add is on the limit surface, and so on. Essentially the subdivision just interpolates new points on the limit surface, making the surface look smoother and smoother but not too different. You can anticipate what the limit curve will look like when you're examining an interpolating subdivision scheme.

Uniform vs. Non-Uniform

Uniform schemes define a single unified way to divide an edge. No matter what type of edge you have or whatever valence the endpoints have, the same scheme is used to subdivide it. Non-uniform methods tailor themselves to different cases, oftentimes specializing to take care of irregularities in the surface. For example, the modified butterfly scheme (which we'll discuss at length shortly) is non-uniform, since it uses three different ways to subdivide edges based on the types of vertices at the endpoints.

Stationary vs. Non-Stationary

This consideration is similar to the uniform/non-uniform one. When a scheme is stationary, the same scheme is used at each subdivision level. Non-stationary methods may use one method for the first subdivision, then switch to another once the surface is looking moderately smooth.

Modified Butterfly Method Subdivision Scheme

The butterfly subdivision scheme was first birthed in 1990 by Dyn, Gregory, and Levin. It handled certain types of surfaces beautifully, but it had a lot of visual discontinuities in certain situations that made it somewhat undesirable. In 1996, Zorin, Schröder, and Sweldens extended the butterfly subdivision scheme to better handle irregular cases, creating the modified butterfly method subdivision scheme. This is the method we're going to focus on for several reasons. First, it's interpolative, so our limit mesh looks a lot like our initial mesh. Second, it works on triangle meshes, which means we can take existing code and drop in subdivision surfaces pretty easily. Finally, it's visually pleasing and easy to code. What more could you want?

To subdivide our mesh, we take each edge and subdivide it into two pieces, forming four new triangles from each original triangle. This is preferred because our subdivided triangles will have a shape similar to their parent triangle (unlike, for example, creating a split location in the center of the triangle and throwing edges to the corners of the triangle). Figure 8.25 shows what a subdivision step looks like.




Figure 8.25:
Subdividing edges
to add triangles

The equation we use to subdivide an edge depends on the valence of its endpoints. The *valence* of a vertex in this context is defined as the number of other vertices the vertex is adjacent to. There are three possible cases that we have to handle.

The first case is when both vertices of a particular edge have a valence=6. We use a mask on the neighborhood of vertices around the edge. This mask is where the modified butterfly scheme gets its name, because it looks sort of like a butterfly. It appears in Figure 8.26.




Figure 8.26:
The butterfly
mask

The modified butterfly scheme added two points and a tension parameter that lets you control the sharpness of the limit surface. Since this scheme complicates the code, I chose to go with a universal w value of 0.0 instead (which resolves to the above Figure 8.26). The modified butterfly mask appears in Figure 8.27.




Figure 8.27:
The modified
butterfly mask

To compute the location of the subdivided edge vertex (the white circle in both images), we step around the neighborhood of vertices and sum them (multiplying each vector by the weight dictated by the mask). You'll notice that all the weights sum up to 1.0. This is good; it means our subdivided point will be in the right neighborhood compared to the rest of the vertices. You can imagine if the sum was much larger the subdivided vertex would be much farther away from the origin than any of the vertices used to create it, which would be incorrect.

When only one of our vertices is regular (i.e., has a valence=6), we compute the subdivided location using the irregular vertex, otherwise known as a k-vertex. This is where the modified butterfly algorithm shines over the original butterfly algorithm (which handled k-vertices very poorly). An example appears in Figure 8.28. The right vertex has a valence of 6, and the left vertex has a valence of 9, so we use the left vertex to compute the location for the new vertex (indicated by the white circle).




Figure 8.28:
Example of a
k-vertex

The general case for a k -vertex has us step around the vertex, weighting the neighbors using a mask determined by the valence of the k -vertex. Figure 8.29 shows the generic k -vertex and how we name the vertices. Note that the k -vertex itself has a weight of $\frac{3}{4}$, in all cases.




Figure 8.29:
Generic k -vertex

There are three cases to deal with: $k=3$, $k=4$, and $k=5$. The masks for each of them are:

$$s_0 = \frac{5}{12}, s_{1,2} = -\frac{1}{12} \quad \text{for } k = 3$$

$$s_0 = \frac{3}{8}, s_2 = -\frac{1}{8}, s_{1,3} = 0 \quad \text{for } k = 4$$

$$s_i = \frac{1}{k} \left(\frac{1}{4} + \cos \frac{2i\pi}{k} + \frac{1}{2} \cos \frac{4i\pi}{k} \right) \quad \text{for } k \geq 5 \text{ (} k \neq 6 \text{)}$$

The final case we need to worry about is when both endpoints of the current edge are k -vertices. When this occurs we compute the k -vertex for both endpoints using the above weights, and average the results together.

Note that we are assuming that our input triangle mesh is a closed boundary representation (doesn't have any holes in it). The paper describing the modified butterfly scheme discusses ways to handle holes in the model (with excellent results), but the code we'll write next won't be able to handle holes in the model so we won't discuss it.

Using these schema for computing our subdivided locations results in an extremely fair-looking surface. Figure 8.30 shows how an octahedron looks as it is repeatedly subdivided. The application we will discuss next was used to create this image.




Figure 8.30:
A subdivided
octagon model

Application: SubDiv

The SubDiv application implements the modified butterfly subdivision scheme we just discussed. It loads an .o3d file and displays it interactively, giving the users the option of subdividing the model whenever they wish.




Figure 8.31: The SubDiv sample




Figure 8.32: The SubDiv sample after division

The model data is represented with an adjacency graph. Each triangle structure holds pointers to the three vertices it is composed of. Each vertex structure has STL vectors that contain pointers to edge structures (one edge for each vertex it's connected to) and triangle structures. The lists are unsorted (which requires linear searching; fixing this to order the edges in clockwise winding order, for example, is left as an exercise for the reader).

The following code gives the header definitions (and many of the functions) for the vertex, edge, and triangle structures. These classes are all defined inside the subdivision surface class (cSubDivSurf).

```
/**
 * Subdivision Surface vertex
 */
struct sVert
{
    /**
     * These two arrays describe the adjacency information
     * for a vertex. Each vertex knows who all of its neighboring
     * edges and triangles are. An important note is that these
     * lists aren't sorted. We need to search through the list
     * when we need to get a specific adjacent triangle.
     * This is, of course, inefficient. Consider sorted insertion
     * an exercise to the reader.
     */
    std::vector< sTriangle* >    m_triList;
    std::vector< sEdge* >       m_edgeList;

    /**
     * position/normal information for the vertex
     */
}
```

```

cGraphicsLayer::cDefaultVertex    m_vert;

/**
 * Each vertex knows its position in the array it lies in.
 * This helps when we're constructing the arrays of subdivided data.
 */
int            m_index;

void AddEdge( sEdge *pEdge )
{
    assert( 0 == std::count(
        m_edgeList.begin(),
        m_edgeList.end(),
        pEdge ) );
    m_edgeList.push_back( pEdge );
}

void AddTri( sTriangle *pTri )
{
    assert( 0 == std::count(
        m_triList.begin(),
        m_triList.end(),
        pTri ) );
    m_triList.push_back( pTri );
}

/**
 * Valence == How many other vertices are connected to this one
 * which said another way is how many edges the vert has.
 */
int Valence()
{
    return m_edgeList.size();
}

sVert() :
    m_triList( 0 ),
    m_edgeList( 0 )
{
}

/**
 * Given a vertex that we know we are attached to, this function
 * searches the list of adjacent edges looking for the one that
 * contains the input vertex. Asserts if there is no edge for
 * that vertex.
 */
sEdge    *GetEdge( sVert *pOther )
{
    for( int i=0; i<m_edgeList.size(); i++ )
    {
        if( m_edgeList[i]->Contains( pOther ) )
            return m_edgeList[i];
    }
    assert(false); // didn't have it!
}

```

```

        return NULL;
    }
};
/**
 * Edge structure that connects two vertices in a SubSurf
 */
struct sEdge
{
    sVert    *m_v[2];

    /**
     * When we perform the subdivision calculations on all the edges
     * the result is held in this newVLoc structure. Never has any
     * connectivity information, just location and color.
     */
    sVert      m_newVLoc;

    /**
     * true == one of the edges' vertices is the input vertex
     */
    bool Contains( sVert *pVert )
    {
        return (m_v[0] == pVert) || m_v[1] == pVert;
    }

    /**
     * retval = the other vertex than the input one
     */
    sVert *Other( sVert *pVert )
    {
        return (m_v[0] == pVert) ? m_v[1] : m_v[0];
    }

    void Init( sVert *v0, sVert *v1 )
    {
        m_v[0] = v0;
        m_v[1] = v1;

        /**
         * Note that the edge notifies both of its vertices that it's
         * connected to them.
         */
        m_v[0]->AddEdge( this );
        m_v[1]->AddEdge( this );
    }

    /**
     * This function takes into consideration the two triangles that
     * share this edge. It returns the third vertex of the first
     * triangle it finds that is not equal to 'notThisOne'. So if we
     * want one, notThisOne is passed as NULL. If we want the other
     * one, we pass the result of the first execution.
     */
    sVert *GetOtherVert( sVert *v0, sVert *v1, sVert *notThisOne )
    {

```

```

    sTriangle *pTri;
    for( int i=0; i<v0->m_triList.size(); i++ )
    {
        pTri = v0->m_triList[i];
        if( pTri->Contains( v0 ) && pTri->Contains( v1 ) )
        {
            if( pTri->Other( v0, v1 ) != notThisOne )
                return pTri->Other( v0, v1 );
        }
    }
    // when we support boundary edges, we shouldn't assert
    assert(false);
    return NULL;
}

/**
 * Calculate the K-Vertex location of 'prim' vertex. For triangles
 * of valence !=6
 */
point3 CalcKVert( int prim, int sec );

/**
 * Calculate the location of the subdivided point using the
 * butterfly method.
 * for edges with both vertices of valence == 6
 */
point3 CalcButterfly();
};

/**
 * Subdivision surface triangle
 */
struct sTriangle
{
    /**
     * The three vertices of this triangle
     */
    sVert    *m_v[3];
    point3    m_normal;

    void Init( sVert *v0, sVert *v1, sVert *v2 )
    {
        m_v[0] = v0;
        m_v[1] = v1;
        m_v[2] = v2;

        /**
         * Note that the triangle notifies all 3 of its vertices
         * that it's connected to them.
         */
        m_v[0]->AddTri( this );
        m_v[1]->AddTri( this );
        m_v[2]->AddTri( this );
    }

    /**
     * true == the triangle contains the input vertex

```

```

    */
    bool Contains( sVert *pVert )
    {
        return pVert == m_v[0] || pVert == m_v[1] || pVert == m_v[2];
    }

    /**
     * retval = the third vertex (first and second are input).
     * asserts out if input values aren't part of the triangle
     */
    sVert *Other( sVert *v1, sVert *v2 )
    {
        assert( Contains( v1 ) && Contains( v2 ) );
        for( int i=0; i<3; i++ )
        {
            if( m_v[i] != v1 && m_v[i] != v2 )
                return m_v[i];
        }
        assert(false); // something bad happened;
        return NULL;
    }
};

```

The interesting part of the application is when the model is subdivided. Since we used vertex buffers to hold the subdivided data, we have an upper bound of 2^{16} , or 65,536, vertices. This code gets called when the user subdivides the model:

```

eResult cSubDivSurf::Subdivide()
{
    /**
     * We know how many components our subdivided model will have, calc them
     */
    int nNewEdges = 2*m_nEdges + 3*m_nTris;
    int nNewVerts = m_nVerts + m_nEdges;
    int nNewTris = 4*m_nTris;

    /**
     * If the model will have too many triangles
     * (d3d can only handle  $2^{16}$ ), return
     */
    if( nNewVerts >= 65536 || nNewTris >= 65536 )
    {
        return resFalse;
    }

    /**
     * Find the location of the new vertices. Most of the hard work
     * is done here.
     */
    GenNewVertLocs();

    int i;

```

```

sVert *inner[3]; // the vertices on the 3 edges (order: 0..1, 1..2, 2..0)

// Allocate space for the subdivided data
sVert *pNewVerts = new sVert[ nNewVerts ];
sEdge *pNewEdges = new sEdge[ nNewEdges ];
sTriangle *pNewTris = new sTriangle[ nNewTris ];

//----- Step 1: Fill the vertex list
// First batch - the original vertices
for( i=0; i<m_nVerts; i++ )
{
    pNewVerts[i].m_index = i;
    pNewVerts[i].m_vert = m_pVList[i].m_vert;
}
// Second batch - vertices from each edge
for( i=0; i<m_nEdges; i++ )
{
    pNewVerts[m_nVerts + i].m_index = m_nVerts + i;
    pNewVerts[m_nVerts + i].m_vert = m_pEList[i].m_newVLoc.m_vert;
}

//----- Step 2: Fill in the edge list
int currEdge = 0;
// First batch - the 2 edges that are spawned by each original edge
for( i=0; i<m_nEdges; i++ )
{
    pNewEdges[currEdge++].Init(
        &pNewVerts[m_pEList[i].m_v[0]->m_index],
        &pNewVerts[m_pEList[i].m_newVLoc.m_index] );
    pNewEdges[currEdge++].Init(
        &pNewVerts[m_pEList[i].m_v[1]->m_index],
        &pNewVerts[m_pEList[i].m_newVLoc.m_index] );
}
// Second batch - the 3 inner edges spawned by each original triangle
for( i=0; i<m_nTris; i++ )
{
    // find the inner 3 vertices of this triangle
    // ( the new vertex of each of the triangles' edges )
    inner[0] = &m_pTList[i].m_v[0]->GetEdge(
        m_pTList[i].m_v[1] )->m_newVLoc;
    inner[1] = &m_pTList[i].m_v[1]->GetEdge(
        m_pTList[i].m_v[2] )->m_newVLoc;
    inner[2] = &m_pTList[i].m_v[2]->GetEdge(
        m_pTList[i].m_v[0] )->m_newVLoc;

    pNewEdges[currEdge++].Init(
        &pNewVerts[inner[0]->m_index],
        &pNewVerts[inner[1]->m_index] );
    pNewEdges[currEdge++].Init(
        &pNewVerts[inner[1]->m_index],
        &pNewVerts[inner[2]->m_index] );
    pNewEdges[currEdge++].Init(
        &pNewVerts[inner[2]->m_index],
        &pNewVerts[inner[0]->m_index] );
}

```



```
//----- Step 3: Fill in the triangle list
int currTri = 0;
for( i=0; i<m_nTris; i++ )
{
    // find the inner vertices
    inner[0] = &m_pTList[i].m_v[0]->GetEdge(
        m_pTList[i].m_v[1] )->m_newVLoc;
    inner[1] = &m_pTList[i].m_v[1]->GetEdge(
        m_pTList[i].m_v[2] )->m_newVLoc;
    inner[2] = &m_pTList[i].m_v[2]->GetEdge(
        m_pTList[i].m_v[0] )->m_newVLoc;

    // 0, inner0, inner2
    pNewTris[currTri++].Init(
        &pNewVerts[m_pTList[i].m_v[0]->m_index],
        &pNewVerts[inner[0]->m_index],
        &pNewVerts[inner[2]->m_index] );

    // 1, inner1, inner0
    pNewTris[currTri++].Init(
        &pNewVerts[m_pTList[i].m_v[1]->m_index],
        &pNewVerts[inner[1]->m_index],
        &pNewVerts[inner[0]->m_index] );

    // 2, inner2, inner1
    pNewTris[currTri++].Init(
        &pNewVerts[m_pTList[i].m_v[2]->m_index],
        &pNewVerts[inner[2]->m_index],
        &pNewVerts[inner[1]->m_index] );

    // inner0, inner1, inner2
    pNewTris[currTri++].Init(
        &pNewVerts[inner[0]->m_index],
        &pNewVerts[inner[1]->m_index],
        &pNewVerts[inner[2]->m_index] );
}

//----- Step 4: Housekeeping

// Swap out the old data sets for the new ones.

delete [] m_pVList;
delete [] m_pEList;
delete [] m_pTList;

m_nVerts = nNewVerts;
m_nEdges = nNewEdges;
m_nTris = nNewTris;

m_pVList = pNewVerts;
m_pEList = pNewEdges;
m_pTList = pNewTris;

// Calculate the vertex normals of the
// new mesh using face normal averaging
```

```

    CalcNormals();

    //-----
    // Step 5: Make arrays so we can send the triangles in one batch

    delete [] m_d3dTriList;
    if( m_pVertexBuffer )
        m_pVertexBuffer->Release();
    m_pVertexBuffer = NULL;

    GenD3DData();

    return resAllGood;
}

/**
 * This is where the meat of the subdivision work is done.
 * Depending on the valence of the two endpoints of each edge,
 * the code will generate the new edge value.
 */
void cSubDivSurf::GenNewVertLocs()
{
    for( int i=0; i<m_nEdges; i++ )
    {
        int val0 = m_pEList[i].m_v[0]->Valence();
        int val1 = m_pEList[i].m_v[1]->Valence();

        point3 loc;

        /**
         * CASE 1: both vertices are of valence == 6
         * Use the butterfly scheme
         */
        if( val0 == 6 && val1 == 6 )
        {
            loc = m_pEList[i].CalcButterfly();
        }

        /**
         * CASE 2: one of the vertices are of valence == 6
         * Calculate the k-vertex for the non-6 vertex
         */
        else if( val0 == 6 && val1 != 6 )
        {
            loc = m_pEList[i].CalcKVert(1,0);
        }

        else if( val0 != 6 && val1 == 6 )
        {
            loc = m_pEList[i].CalcKVert(0,1);
        }

        /**
         * CASE 3: neither of the vertices are of valence == 6
         * Calculate the k-vertex for each of them, and average

```

```

        * the result
        */
    else
    {
        loc = ( m_pEList[i].CalcKVert(1,0) +
                m_pEList[i].CalcKVert(0,1) ) / 2.f;
    }

    cGraphicsLayer::cDefaultVertex tempVertex;
    tempVertex.m_vPosition = *(D3DXVECTOR3*)&loc;
    tempVertex.m_vNormal = D3DXVECTOR3(0,0,0);
    tempVertex.m_vColor = D3DXCOLOR(0,0,0,0);
    tempVertex.m_TexCoords = D3DXVECTOR2(0,0);

    m_pEList[i].m_newWLoc.m_vert = tempVertex;

    /**
     * Assign the new vertex an index. (This is useful later,
     * when we start throwing vertex pointers around. We
     * could have implemented everything with indices, but
     * the code would be much harder to read. An extra dword
     * per vertex is a small price to pay.)
     */
    m_pEList[i].m_newWLoc.m_index = i + m_nVerts;
}

}

point3 cSubDivSurf::sEdge::CalcButterfly()
{
    point3 out = point3::Zero;

    sVert *other[2];
    other[0] = GetOtherVert( m_v[0], m_v[1], NULL );
    other[1] = GetOtherVert( m_v[0], m_v[1], other[0] );

    // two main ones
    out += (1.f/2.f) * (*(point3*)&(m_v[0]->m_vert.m_vPosition));
    out += (1.f/2.f) * (*(point3*)&(m_v[1]->m_vert.m_vPosition));

    // top/bottom ones
    out += (1.f/8.f) * (*(point3*)&(other[0]->m_vert.m_vPosition));
    out += (1.f/8.f) * (*(point3*)&(other[1]->m_vert.m_vPosition));

    // outside 4 verts
    out += (-1.f/16.f) * (*(point3*)&
        (GetOtherVert( other[0], m_v[0], m_v[1] )->m_vert.m_vPosition));
    out += (-1.f/16.f) * (*(point3*)&
        (GetOtherVert( other[0], m_v[1], m_v[0] )->m_vert.m_vPosition));
    out += (-1.f/16.f) * (*(point3*)&
        (GetOtherVert( other[1], m_v[0], m_v[1] )->m_vert.m_vPosition));
    out += (-1.f/16.f) * (*(point3*)&
        (GetOtherVert( other[1], m_v[1], m_v[0] )->m_vert.m_vPosition));

    return out;
}

```

```

point3 cSubDivSurf::sEdge::CalcKVert(int prim, int sec)
{
    int valence = m_v[prim]->Valence();

    point3 out = point3::Zero;

    out += (3.f / 4.f) * *(point3*)&m_v[prim]->m_vert.m_vPosition;

    if( valence < 3 )
        assert( false );

    else if( valence == 3 )
    {
        for( int i=0; i<m_v[prim]->m_edgeList.size(); i++ )
        {
            sVert *pOther = m_v[prim]->m_edgeList[i]->Other( m_v[prim] );
            if( pOther == m_v[sec] )
                out += (5.f/12.f) * *(point3*)&pOther->m_vert.m_vPosition;
            else
                out += (-1.f/12.f) * *(point3*)&pOther->m_vert.m_vPosition;
        }
    }

    else if( valence == 4 )
    {
        out += (3.f/8.f) * *(point3*)&m_v[sec]->m_vert.m_vPosition;

        sVert *pTemp = GetOtherVert( m_v[0], m_v[1], NULL );
        // get the one after it
        sVert *pOther = GetOtherVert( m_v[prim], pTemp, m_v[sec] );

        out += (-1.f/8.f) * *(point3*)&pOther->m_vert.m_vPosition;
    }

    else // valence >= 5
    {
        sVert *pCurr = m_v[sec];
        sVert *pLast = NULL;
        sVert *pTemp;
        for( int i=0; i< valence; i++ )
        {
            float weight =
                ((1.f/4.f) + (float)cos( 2 * PI *
                    (float)i / (float)valence ) +
                (1.f/2.f) * (float)cos(4*PI*(float)i/
                    (float)valence)) / (float)valence;

            out += weight * *(point3*)&pCurr->m_vert.m_vPosition;

            pTemp = GetOtherVert( m_v[prim], pCurr, pLast );
            pLast = pCurr;
            pCurr = pTemp;
        }
    }

    return out;
}

```

```

    }

    void cSubDivSurf::CalcNormals()
    {
        int i;

        // reset all vertex normals
        for( i=0; i<m_nVerts; i++ )
        {
            m_pVList[i].m_vert.m_vNormal = D3DXVECTOR3(0,0,0);
        }

        // find all triangle normals
        for( i=0; i<m_nTris; i++ )
        {
            m_pTList[i].m_normal = plane3(
                *(point3*)&m_pTList[i].m_v[0]->m_vert.m_vPosition),
                *(point3*)&m_pTList[i].m_v[1]->m_vert.m_vPosition),
                *(point3*)&m_pTList[i].m_v[2]->m_vert.m_vPosition)).n;

            // add the normal to each vertex
            m_pTList[i].m_v[0]->m_vert.m_vNormal+=
                *(D3DXVECTOR3*)&m_pTList[i].m_normal);
            m_pTList[i].m_v[1]->m_vert.m_vNormal +=
                *(D3DXVECTOR3*)&m_pTList[i].m_normal);
            m_pTList[i].m_v[2]->m_vert.m_vNormal +=
                *(D3DXVECTOR3*)&m_pTList[i].m_normal);
        }

        // reset all vertex normals
        for( i=0; i<m_nVerts; i++ )
        {
            D3DXVec3Normalize(
                &m_pVList[i].m_vert.m_vNormal, &m_pVList[i].m_vert.m_vNormal);
        }
    }

    void cSubDivSurf::GenD3DData()
    {
        /**
         * Create a vertex buffer
         */

        HRESULT hr;

        D3D10_BUFFER_DESC descBuffer;
        memset(&descBuffer, 0, sizeof(descBuffer));
        descBuffer.Usage = D3D10_USAGE_DYNAMIC;
        descBuffer.ByteWidth = sizeof(cGraphicsLayer::cDefaultVertex) *m_nVerts;
        descBuffer.BindFlags = D3D10_BIND_VERTEX_BUFFER;
        descBuffer.CPUAccessFlags = D3D10_CPU_ACCESS_WRITE;
        descBuffer.MiscFlags = 0;
    }

```

```

Graphics()->GetDevice()->CreateBuffer(
    &descBuffer, NULL, &m_pVertexBuffer);

if(!m_pVertexBuffer)
{
    throw cGameError(L"Vertex Buffer creation failed!\n");
}

m_d3dTriList = new sTri[ m_nTris ];

cGraphicsLayer::cDefaultVertex *pVert;

m_pVertexBuffer->Map(D3D10_MAP_WRITE_DISCARD, 0, (void**)&pVert);

if(!pVert)
{
    throw cGameError(L"VB map failed\n");
}

int i;

// Copy data into the buffer
for( i=0; i<m_nVerts; i++ )
{
    *pVert++ = m_pVList[i].m_vert;
}
m_pVertexBuffer->Unmap();

for( i=0; i<m_nTris; i++ )
{
    m_d3dTriList[i].v[0] = m_pTList[i].m_v[0]->m_index;
    m_d3dTriList[i].v[1] = m_pTList[i].m_v[1]->m_index;
    m_d3dTriList[i].v[2] = m_pTList[i].m_v[2]->m_index;
}
}

```

Progressive Meshes

The final multiresolution system we are going to discuss is progressive meshes. They're rapidly gaining favor in the game community; many games use them as a way to keep scene detail at a constant level.

Oftentimes when we're playing a 3D game, many of our objects will appear off in the distance. For example, if we're building a combat flight simulator, bogies will appear miles away before we engage them. When an object is this far away, it will appear to be only a few pixels on the screen.

We could simply opt not to draw an object if it is this far away. However, this can lead to a discontinuity of experience for the user. He or she will suddenly remember they're playing a video game, and that should be avoided at all costs. If we have a model with thousands of triangles in it to represent our enemy aircraft, we're going to waste a lot of time transforming and lighting vertices when we'll end up with just a blob of a few pixels.

Drawing several incoming bogie blobs may max out our triangle budget for the frame, and our frame rate will drop. This will hurt the user experience just as much if not more than not drawing the object in the first place.

Even when the object is moderately close, if most of the triangles are smaller than one pixel, we're wasting effort on drawing our models. If we used, instead, a lower-resolution version of the mesh for farther distances, the visual output would be about the same, but we would save a lot of time in model processing.

This is the problem progressive meshes try to solve. They allow us to arbitrarily scale the polygon resolution of a mesh from its max all the way down to two triangles. When our model is extremely far away, we draw the lowest resolution model we can. Then, as it approaches the camera, we slowly add detail polygon by polygon, so the user always will be seeing a nearly ideal image at a much faster frame rate. Moving between detail levels on a triangle-by-triangle basis is much less noticeable than switching between a handful of models at different resolutions. We can even morph our triangle-by-triangle transitions using what are called geomorphs, making them even less noticeable.

Progressive meshes can also help us when we have multiple close objects on the screen. If we used just the distance criterion discussed above to set polygon resolution, we could easily have the case where there are multiple dense objects close to the camera. We would have to draw them all at a high resolution, and we would hit our polygon budget and our frame rate would drop out. In this extreme situation, we can suffer some visual quality loss and turn down the polygon count of our objects. In general, when a user is playing an intense game, he or she won't notice that the meshes are lower resolution. Users will, however, immediately notice a frame rate reduction.

One thing progressive meshes can't do is add detail to a model. Unlike the other two multiresolution surface methods we have discussed, progressive meshes can only vary the detail in a model from its original polygon count down to two polygons.

Progressive meshes were originally described in a 1996 SIGGRAPH paper by Hugues Hoppe. Since then a lot of neat things have happened with them. Hoppe has applied them to view-dependent level-of-detail and terrain rendering. They were added to Direct3D Retained Mode (which is no longer supported). Recently, Hoppe extended research done by Michael Garland and Paul Heckbert, using quadric error metrics to encode normal, color, and texture information. We'll be covering some of the basics of quadric error metrics. Hoppe's web site (<http://www.research.micro-soft.com/~hoppe>) has downloadable versions of all his papers.

Progressive Mesh Basics

How do progressive meshes work? They center around an operation called an *edge collapse*. Conceptually, it takes two vertices that share an edge and merges them. This destroys the edge that was shared and the two triangles that shared the edge.

The cool thing about edge collapse is that it only affects a small neighborhood of vertices, edges, and triangles. We can save the state of those entities in a way such that we can reverse the effect of the edge collapse, splitting a vertex into two, adding an edge, and adding two triangles. This operation, the inverse of the edge collapse, is called a *vertex split*. Figure 8.33 shows how the edge collapse and vertex split work.




Figure 8.33:
The edge collapse
and vertex split
operations

To construct a progressive mesh, we take our initial mesh and iteratively remove edges using edge collapses. Each time we remove an edge, the model loses two triangles. We then save the edge collapse we performed into a stack, and continue with the new model. Eventually, we reach a point where we can no longer remove any edges. At this point we have our lowest resolution mesh and a stack of structures representing each edge that was collapsed. If we want to have a particular number of triangles for our model, all we do is apply vertex splits or edge collapses to get to the required number (plus or minus one, though, since we can only change the count by two).

During run time, most systems have three main areas of data: a stack of edge collapses, a stack of vertex splits, and the model. To apply a vertex split, we pop one off the stack, perform the requisite operations on the mesh, construct an edge collapse to invert the process, and push the newly created edge collapse onto the edge collapse stack. The reverse process applies to edge collapses.

There are a lot of cool side effects that arise from progressive meshes. For starters, they can be stored on disk efficiently. If an application is smart about how it represents vertex splits, storing the lowest resolution mesh and the sequence of vertex splits to bring it back to the highest resolution

model doesn't take much more space than storing the high-resolution mesh on its own.

Also, the entire mesh doesn't need to be loaded all at once. A game could load the first 400 or so triangles of each model at startup and then load more vertex splits as needed. This can save some time if the game is being loaded from disk, and a lot of time if the game is being loaded over the Internet.

Another thing to consider is that since the edge collapses happen in such a small region, many of them can be combined, getting quick jumps from one resolution to another. Each edge collapse/vertex split can even be morphed, smoothly moving the vertices together or apart. This alleviates some of the popping effects that can occur when progressive meshes are used without any morphing. Hoppe calls these transitions *geomorphs*.

Choosing Our Edges

The secret to making a good progressive mesh is choosing the right edge to collapse during each iteration. The sequence is extremely important. If we choose our edges unwisely, our low-resolution mesh won't look anything like our high-resolution mesh.

As an extreme example, imagine we chose our edges completely at random. This can have extremely adverse effects on the way our model looks even after a few edge collapses.



Warning: Obviously, we should not choose vertices completely at random. We have to take other factors into account when choosing an edge. Specifically, we have to maintain the topology of a model. We shouldn't select edges that will cause seams in our mesh (places where more than two triangles meet an edge).

Another naïve method of selecting edges would be to choose the shortest edge at each point in time. This uses the well-founded idea that smaller edges won't be as visible to the user from faraway distances, so they should be destroyed first. However, this method overlooks an important factor that must be considered in our final selection algorithm. Specifically, small details, such as the nose of a human face or the horns of a cow, must be preserved as long as possible if a good low-polygon representation of the model is to be created. We must not only take into account the length of the edge, but also how much the model will change if we remove it. Ideally, we want to pick the edge that changes the visual look of the model the least. Since this is a very fuzzy heuristic, we end up approximating it.

The opposite extreme would be to rigorously try to approximate the least-visual-change heuristic, and spend an awfully long time doing it. While this will give us the best visual model, it is less than ideal. If we can spend something like 5% of the processing time and get a model that looks 95% as good as an ultra-slow ideal method, we should use that one. Let's discuss two different edge selection algorithms.

An Edge Selection Algorithm

Stan Melax wrote an article for *Game Developer* magazine way back in November 1998 that detailed a simple and fast function to compute the relative cost of contracting a vertex v into a vertex u . Since they are different operations, $\text{cost}(u,v)$ will generally be different from $\text{cost}(v,u)$. The algorithm's only shortcoming lies in the fact that it can only collapse one vertex onto another; it cannot take an edge and reposition the final vertex in a location to minimize the total error (as quadric error metrics can do). The cost function is:

$$\text{cost}(\mathbf{u}, \mathbf{v}) = \|\mathbf{u} - \mathbf{v}\| \times \max_{f \in Tu} \left(\min_{n \in Tuv} \{ \|\mathbf{f} \cdot \mathbf{n}\| \div 2 \} \right)$$

where Tu is the set of triangles that share vertex u , and Tuv is the set of triangles that share both vertex u and v .

Quadric Error Metrics

Michael Garland and Paul Heckbert devised an edge selection algorithm in 1997 that was based on quadric error metrics (published as “Surface Simplification Using Quadric Error Metrics” in *Computer Graphics*). The algorithm is not only extremely fast, its output looks very nice. I don't have the space to explain all the math needed to get this algorithm working (specifically, generic matrix inversion code), but we can go over enough to get your feet wet.

Given a particular vertex v and a new vertex v' , we want to be able to find out how much error would be introduced into the model by replacing v with v' . If we think of each vertex as being the intersection point of several planes (in particular, the planes belonging to the set of triangles that share the vertex), then we can define the error as how far the new vertex is from each plane.

This algorithm uses the squared distance. This way we can define an error function for a vertex v given the set of planes p that share the vertex as:

$$\begin{aligned} \Delta(\mathbf{v}) &= \sum_{p \in \text{planes}(\mathbf{v})} (\mathbf{p}^T \mathbf{v})^2 \\ \Delta(\mathbf{v}) &= \sum_{p \in \text{planes}(\mathbf{v})} (\mathbf{v}^T \mathbf{p})(\mathbf{p}^T \mathbf{v}) \\ \Delta(\mathbf{v}) &= \sum_{p \in \text{planes}(\mathbf{v})} \mathbf{v}^T (\mathbf{p} \mathbf{p}^T) \mathbf{v} \\ \Delta(\mathbf{v}) &= \mathbf{v}^T \left(\sum_{p \in \text{planes}(\mathbf{v})} \mathbf{K}_p \right) \mathbf{v} \end{aligned}$$

The matrix K_p represents the coefficients of the plane equation $\langle a, b, c, d \rangle$ for a particular plane p multiplied with its transpose to form a 4x4 matrix. Expanded, the multiplication becomes:

$$K_p = \begin{bmatrix} a^2 & ba & ca & da \\ ab & b^2 & cb & db \\ ac & bc & c^2 & dc \\ ad & bd & cd & d^2 \end{bmatrix}$$

K_p is used to find the squared distance error of a vertex to the plane it represents. We sum the matrices for each plane to form the matrix Q :

$$Q = \sum_{p \in \text{planes}(\mathbf{v})} K_p$$

which makes the error equation:

$$\Delta(\mathbf{v}) = \mathbf{v}^T Q \mathbf{v}$$

Given the matrix Q for each of the vertices in the model, we can find the error for taking out any particular edge in the model. Given an edge between two vertices v_1 and v_2 , we find the ideal vertex v' by minimizing the function:

$$\mathbf{v}'^T (Q_1 + Q_2) \mathbf{v}'$$

where Q_1 and Q_2 are the Q matrices for v_1 and v_2 .

Finding v' is the hard part of this algorithm. If we want to try solving it exactly, we just want to solve the equation:

$$\begin{bmatrix} q_{11} & q_{12} & q_{13} & q_{14} \\ q_{21} & q_{22} & q_{23} & q_{24} \\ q_{31} & q_{32} & q_{33} & q_{34} \\ 0 & 0 & 0 & 1 \end{bmatrix} \mathbf{v}' = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \end{bmatrix}$$

where the 4x4 matrix above is $(Q_1 + Q_2)$ with the bottom row changed around. If the matrix above is invertible, then the ideal v' (the one that has zero error) is just:

$$\mathbf{v}' = \begin{bmatrix} q_{11} & q_{12} & q_{13} & q_{14} \\ q_{21} & q_{22} & q_{23} & q_{24} \\ q_{31} & q_{32} & q_{33} & q_{34} \\ 0 & 0 & 0 & 1 \end{bmatrix}^{-1} \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \end{bmatrix}$$

If the matrix isn't invertible, then the easiest thing to do, short of solving the minimization problem, would be to just choose the vertex causing the

least error out of the set $(v_1, v_2, (v_1 + v_2)/2)$. Finding out if the matrix is invertible, and inverting it, is the ugly part that I don't have space to explain fully. It isn't a terribly hard problem, given a solid background in linear algebra.

We compute the ideal vertex (the one that minimizes the error caused by contracting an edge) and store the error associated with that ideal vertex (since it may not be zero). When we've done this for each of the edges, the best edge to remove is the one with the least amount of error. After we collapse the cheapest edge, we recompute the Q matrices and the ideal vertices for each of the vertices in the immediate neighborhood of the removed edge (since the planes have changed) and continue.

Implementing a Progressive Mesh Renderer

Due to space and time constraints, code to implement progressive meshes is not included in this book. That shouldn't scare you off, however; they're not too hard to implement. The only real trick is making them efficient.

How you implement progressive meshes depends on whether you calculate the mesh as a preprocessing step or at run time. A lot of extra information needs to be kept around during the mesh construction to make it even moderately efficient, so it might be best to write two applications. The first one would take an object, build a progressive mesh out of it, and write the progressive mesh to disk. A separate application would actually load the progressive mesh off the disk and display it. This would have a lot of advantages; most notably you could make both algorithms (construction and display) efficient in their own ways without having to make them sacrifice things for each other.

To implement a progressive mesh constructor efficiently, you'll most likely want something along the lines of the code used in the subdivision surface renderer, where each vertex knows about all the vertices around it. As edges were removed, the adjacency information would be updated to reflect the new topology of the model. This way it would be easy to find the set of vertices and triangles that would be modified when an edge is removed.

Storing the vertex splits and edge collapses can be done in several ways. One way would be to make a structure like the one here:

```
// can double as sVSplit
struct sECol
{
    // the 2 former locations of the vertices
    point3 locs[2];

    // where the collapsed vertex goes
    point3 newLoc;

    // Indices of the two vertices
    int verts[2];
};
```

```

// Indices of the two triangles
int tris[2];

// The indices of triangles that need to
// have vertex indices swapped
vector<int> modTris;
};

```

When it came time to perform a vertex split, you would perform the following steps:

1. Activate (via an active flag) `verts[1]`, `tris[0]`, and `tris[1]` (`verts[0]` is the collapsed vertex, so it's already active).
2. Move `verts[0]` and `verts[1]` to `locs[0]` and `locs[1]`.
3. For each of the triangles in `modTris`, change any indices that point to `verts[0]` and change them to `verts[1]`. You can think of the `modTris` as being the set of triangles below the collapsed triangles in Figure 8.33.

Performing an edge collapse would be a similar process, just reversing everything.

Radiosity

In this section we're going to discuss a way to do lighting that is very accurate, but only handles diffuse light: radiosity lighting.

The wave/particle duality aside, light acts much like any other type of energy. It leaves a source in a particular direction; as it hits objects, some of the energy is absorbed and some is reflected back into the scene. The direction in which it reflects depends on the microscopic structure of the surface. Surfaces that appear smooth at a macroscopic level, like chalk, actually have a really rough microstructure when seen under a microscope.

The light that leaves an object may bounce off of a thousand other points in the scene before it eventually reaches our eye. In fact, only a tiny amount (generally less than a tenth of 1%) of all the energy that leaves a light ever reaches our eye. Because of this, the light that reflects off of other objects affects the total lighting of the scene.

An example: When you're watching a movie at a movie theater, there is generally only one light in the scene (sans exit lights, aisle lights, cell phones, etc.), and that is the movie projector. The only object that directly receives light from the movie projector is the movie screen. However, that is not the only object that receives any light. If you've ever gotten up to get popcorn, you're able to see everyone in the theater watching the movie because light is bouncing off the screen, onto and off of peoples' faces, and bouncing into your eyes. The problem with the lighting models we've discussed so far is that they can't handle this. Sure, we could just turn up the ambient color to simulate the light reflecting off the screen into the

theater, but that won't work; since we only want the front sides of people to be lit, it will look horridly wrong.

What we would like is to simulate the real world, and find not only the light that is emitted from light sources that hits surfaces, but also find the light that is reflected from other surfaces. We want to find the interreflection of light in our 3D scene.

This is both good and bad (but not ugly, thankfully). The good is, the light in our scene will behave more like light we see in the real world. Light will bounce off of all the surfaces in our scene. Modeling this interreflection will give us an extremely slick-looking scene. The bad thing is, the math suddenly becomes much harder, because now all of our surfaces are interrelated. The lighting calculation must be done as a precalculating step, since it's far too expensive to do in real time. We save the radiosity results into the data file we use to represent geometry on disk, so any program using the data can take advantage of the time spent calculating the radiosity solution.



Note: Radiosity isn't for everyone. While *Quake II* used it to great effect to light the worlds, *Quake III* did not. The motivation behind not using it for *Quake III* lies partially in the fact that computing the correct radiosity solution for Bezier surfaces is a total pain, and radiosity doesn't give shadows as sharp as non-interreflective lighting schemes. *Quake* had a very certain look and feel because of how its shadows worked. *Quake III* went back to that.

Radiosity Foundations

We'll begin our discussion of radiosity with some basic terms that we'll use in the rest of the equations:

Table 8.1: Some basic terms used in radiosity

Radiance (or intensity)	The light (or power) coming into (or out of) an area in a given direction. Units: power / (area x solid angle)
Radiosity	The light leaving an area. This value can be thought of as color leaving a surface. Units: power / area
Radiant emitted flux density	The unit for light emission. This value can be thought of as the initial color of a surface. Units: power / area

Our initial scene is composed of a set of closed polygons. We subdivide our polygons into a grid of patches. A *patch* is a discrete element with a computable surface area whose radiosity (and color) remains constant across the whole surface.

The amount we subdivide our polygons determines how intricately our polygon can be lit. You can imagine the worst case of a diagonal shadow falling on a surface. If we don't subdivide enough, we'll be able to see a stepping pattern at the borders between intensity levels. Another way to think of this is drawing a scene in 320x200 versus 1600x1200. The more resolution we add, the better the output picture looks. However, the more patches we add, the more patches we need to work with, which makes our algorithm considerably slower.

Radiosity doesn't use traditional lights (like point lights or spotlights). Instead, certain patches actually emit energy (light) into the scene. This could be why a lot of the radiosity images seen in books like Foley's are offices lit by fluorescent ceiling panel lights (which are quite easy to approximate with a polygon).

Let's consider a particular patch i in our scene. We want to find the radiosity leaving our surface. (This can be a source of confusion: Radiosity is both an algorithm and a unit!) Essentially, the radiosity leaving our surface is the color of the surface when we end up drawing it. For example, the more red energy leaving the surface, the more red light will enter our virtual eye looking at the surface, making the surface appear more red. For all of the following equations, power is equivalent to light.

$$\begin{pmatrix} \text{outgoing} \\ \text{power of} \\ \text{element } i \end{pmatrix} = \begin{pmatrix} \text{power} \\ \text{emitted by} \\ \text{element } i \end{pmatrix} + \begin{pmatrix} \text{power} \\ \text{reflected by} \\ \text{element } i \end{pmatrix}$$

We know how much power each of our surfaces emit. All the surfaces we want to use as lights emit some light; the rest of the surfaces don't emit any. All we need to know is how much is reflected by a surface. This ends up being the amount of energy the surface receives from the other surfaces, multiplied by the reflectance of the surface. Expanding the right side gives:

$$\begin{pmatrix} \text{outgoing} \\ \text{power of} \\ \text{elem. } i \end{pmatrix} = \begin{pmatrix} \text{power} \\ \text{emitted by} \\ \text{elem. } i \end{pmatrix} + \begin{pmatrix} \text{reflectance} \\ \text{of elem. } i \end{pmatrix} \times \sum_{\forall j \neq i} \begin{pmatrix} \text{outgoing} \\ \text{power} \\ \text{of elem. } j \end{pmatrix} \times \begin{pmatrix} \text{fraction of power} \\ \text{leaving elem. } j \text{ that} \\ \text{arrives at elem. } i \end{pmatrix}$$

So this equation says that the energy reflected by element i is equal to the incoming energy times a reflectance term that says how much of the incoming energy is reflected back into the scene. To find the energy incoming to our surface, we take every other surface j in our scene, find out how much of the outgoing power of j hits i , and sum all of the energy terms together. You may have noticed that in order to find the outgoing power of element i we need the outgoing power of element j , and in order to find the outgoing power of element j we need the outgoing power of element i . We'll cover this soon.

Let's define some variables to represent the terms above and flesh out a mathematical equation:

Table 8.2: Variables for our radiosity equations

A_i	Area of patch i . (This is pretty easy to compute for quads.)
e_i	Radiant emitted flux density of patch i . (We are given this. Our luminous surfaces get to emit light of a certain color.)
ρ_i	Reflectance of patch i . (We're given this too. It's how much the patch reflects each color component. Essentially, this is the color of the patch when seen under bright white light.)
b_i	Radiosity of patch i . (This is what we want to find.)
$F_{j \rightarrow i}$	Form factor from patch j to patch i (the fraction of the total radiosity leaving j that directly hits i , which we will compute later).

So if we simply rewrite the equation we have above with our defined variables, we get the following radiosity equation:

$$b_i = e_i + \rho_i \sum_{j=1}^n b_j F_{j \rightarrow i} \frac{A_j}{A_i}$$

We're going to go into the computation of the form factor later. For right now we'll just present a particular trait of the form factor called the Reciprocity Law:

$$A_i F_{ij} = A_j F_{ji}$$

This states that the form factors between subpatches are related to the areas of each of the subpatches. With this law we can simplify and rearrange our equation to get the following:

$$b_i = e_i + \rho_i \sum_{j=1}^n b_j F_{i \rightarrow j} \quad b_i - \rho_i \sum_{j=1}^n b_j F_{i \rightarrow j} = e_i$$

By now you've probably noticed an icky problem: To find the radiosity of some surface i we need to know the radiosity of all of the other surfaces, presenting a circular dependency. To get around this we need to solve all of the radiosity equations simultaneously.

The way this is generally done is to take all n patches in our scene and compose a humongous $n \times n$ matrix, turning all of the equations above into one matrix equation.

$$\begin{bmatrix} 1 - \rho_1 F_{1 \rightarrow 1} & -\rho_1 F_{1 \rightarrow 2} & \cdots & -\rho_1 F_{1 \rightarrow n} \\ -\rho_2 F_{2 \rightarrow 1} & 1 - \rho_2 F_{2 \rightarrow 2} & \cdots & -\rho_2 F_{2 \rightarrow n} \\ \vdots & \vdots & \ddots & \vdots \\ -\rho_n F_{n \rightarrow 1} & -\rho_n F_{n \rightarrow 2} & \cdots & 1 - \rho_n F_{n \rightarrow n} \end{bmatrix} \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{bmatrix} = \begin{bmatrix} e_1 \\ e_2 \\ \vdots \\ e_n \end{bmatrix}$$

I could try to explain how to solve this monstrosity, but hopefully we're all getting the idea that this is the wrong way to go. Getting a good radiosity solution can require several thousand patches for even simple scenes, which will cost us tens of megabytes of memory for the $n \times n$ matrix, and forget about the processing cost of trying to solve said multimegabyte matrix equation.

Unless we can figure out some way around this, we're up a creek. Luckily, there is a way around. In most situations, a lot of the values in the matrix will be either zero or arbitrarily small. This is called a *sparse* matrix. The amount of outgoing energy for most of these patches is really small, and will only contribute to a small subset of the surfaces. Rather than explicitly solve this large sparse matrix, we can solve it progressively, saving us a ton of memory and a ton of time.

Progressive Radiosity

The big conceptual difference between progressive radiosity and matrix radiosity is that in progressive radiosity we shoot light out from patches, instead of receiving it. Each patch has a value that represents how much energy it has to give out (Δ Radiosity, or deltaRad) that is initially set to how much energy the surface emits. Each iteration, we choose the patch that has the most energy to give out (deltaRad * the area of the patch). We then send its energy out into the scene, finding how much of it hits each surface. We add the incoming energy to the radiosity and deltaRad of each other patch. Finally, we set the deltaRad of our source patch to zero (since, at this point, it has released all of its energy) and repeat. Whenever the patch with the most energy has its energy value below a certain threshold, we stop.

Here's pseudocode for the algorithm:

```
For( each patch 'curr' )
    curr.radiosity = curr.emitted
    curr.deltaRad = curr.emitted
while( not done )
    source = patch with max. outgoing energy (deltaRad * area)
    if( source.deltaRad < threshold )
        done = true
    For( each patch 'dest' != source )
        deltaRad = dest.reflectiveness *
            FormFactor( dest, source )
        dest.radiosity += deltaRad
        dest.deltaRad += deltaRad
    source.deltaRad = 0
Draw scene (if desired)
```

The Form Factor

The final piece of the puzzle is the calculation of this mysterious form factor. Again, it represents the amount of energy that leaves a subpatch i that reaches a subpatch j . The initial equation is not as scary as it looks. The definition of the form factor between two subpatches i and j is:

$$F_{i \rightarrow j} = \frac{v_{ij}}{A_i} \int \int \frac{\cos \theta_i \cos \theta_j}{\pi r^2} v_{i \rightarrow j} dA_j dA_i$$

Table 8.3: Variable meanings for the form factor equation

v_{ij}	Visibility relationship between i and j ; 1 if there is a line of sight between the two elements, 0 otherwise.
dA_i, dA_j	Infinitesimally small pieces of the elements i and j .
r	The length of the ray separating i and j .
θ_i, θ_j	The angle between the ray separating i and j and the normals of i and j , respectively (see Figure 8.34).

Figure 8.34 may help you visualize the relationship between some of the variables in the form factor equation.




Figure 8.34:
The θ and r
variables
visualized

Maybe you enjoy working with troublesome double integrals. I don't. I look at equations like this, think about having to write code to handle it, and run for the hills. Luckily we can give up a little bit in accuracy and get rid of those nasty integrals.

What do the integrals mean? Essentially, we want to compute the bulk of the equation an infinite number of times for a set of infinitely small pieces of the patches, and sum them all together. Of course, doing it an

infinite number of times is unreasonable. We can do it enough so that our solution is close enough to what we would get if we had computed the integral properly. We're not even going to do it enough, though; we're just going to do it once.

What is the justification for this? Our patches are generally going to be pretty small, small to the point that the radiosity for each of the subpatches is going to be pretty much the same. We can't get much variance in the amount of light hitting a surface when the surface is only a few inches square. Of course, there are cases where it could fail, but they most likely won't come up, and if they do that's what we get for approximating.

Instead of computing the form factor equation for a bunch of small subpatches, we're just going to compute it once for both patches. The delta areas become the regular areas, and we compute the line of sight only once, using the centers of the patches. This makes our equation much nicer looking:

$$F_{i-j} = \frac{1}{A_i} \frac{\cos\theta_i \cos\theta_j}{\pi r^2} v_{i-j} A_j A_i$$

$$F_{i-j} = \frac{\cos\theta_i \cos\theta_j}{\pi r^2} v_{i-j} A_j$$

This isn't painful at all. To compute the line of sight, we'll just use the BSP tree code we developed in Chapter 4. Testing is quick (anywhere from $O(\lg n)$ to $O(n)$ worst case, where n is the number of polygons), and it's not dependent on the number of patches, just the number of polygons.

Application: Radiosity

With all the pieces in place, we can finally make a stab at implementing a radiosity simulator, which I have taken the liberty of doing. It loads a scene description file off the disk and progressively adds radiosity to the scene. For each frame, it processes the brightest patch and then renders it. That way, as the program is running, light slowly fills the room.

The first non-commented line of the file contains the number of polygons. The following listing shows the header and the first polygon of the provided data file. The first line of the polygon has four floating-point values, the first three of which describe the energy of the surface. Most of the polygons have the energy set to black, but there are three lit polygons in the room to add light to it. The fourth component is the reflectance of the polygon. This should be an RGB triplet as well; making it just a float restricts all the surfaces to be varying shades of gray. After the polygon header there are four lines with three floats each, defining the four corners of the polygon. When the polygon is loaded, it is subdivided into a bunch of subpatches until their area is below a constant threshold.

```
# this is a more complex data set
26
# top of the room, very reflective
0.0 0.0 0.0 0.76 ##
-10.0 10.0 -10.0
10.0 10.0 -10.0
10.0 10.0 -8.0
-10.0 10.0 -8.0
...
```

This code can only correctly deal with square polygons. Adding support for other types of polygons wouldn't be hard, but I didn't want to over-complicate the code for this program. Also, for the sake of simplicity, patches are flat shaded. Computing the right color for the patch corners is harder than you would think. The naïve solution would be to just compute the radiosity equations using the vertices instead of the centers of the patches. The problem occurs at corners. Since the point you're computing is right against the polygon next to it, it won't receive any light, and you'll get an almost black line running around the borders of all your polygons—an unacceptable artifact. There is a nifty algorithm in Foley's *Computer Graphics* in the radiosity section to compute vertex colors from patch colors; implementing it is left as an exercise for the reader.

A screenshot from the radiosity application after it has run its course (it can take a while on older systems—for comparison: five minutes on a 1997 Celeron 366, or 5 seconds on a P4 1.5Ghz, or a few milliseconds on my Core2 Duo 2.4Ghz) appears in Figure 8.35.




Figure 8.35: Screenshot from radiosity calculator

Here is a portion of the code:

```
bool cRadiosityCalc::LineOfSight( sPatch *a, sPatch *b )
{
    // Early-out 1: they're sitting on the same spot
    if( a->m_plane == b->m_plane )
        return false;

    // Early-out 2: b is behind a
    if( a->m_plane.TestPoint( b->m_center ) == ptBack )
        return false;

    // Early-out 3: a is behind b
    if( b->m_plane.TestPoint( a->m_center ) == ptBack )
        return false;

    // Compute the slow
    return m_tree.LineOfSight( a->m_center, b->m_center );
}

float cRadiosityCalc::FormFactor( sPatch *pSrc, sPatch *pDest )
{
    float   angle1, angle2, dist, factor;
    point3  vec;

    // find vij first.  If it's 0, we can early-out.
    if( !LineOfSight( pSrc, pDest ) )
        return 0.f;

    point3 srcLoc = pSrc->m_center;
    point3 destLoc = pDest->m_center;

    vec = destLoc - srcLoc;
    dist = vec.Mag();
    vec /= dist;

    angle1 = vec * pSrc->m_plane.n;
    angle2 = -( vec * pDest->m_plane.n );

    factor = angle1 * angle2 * pDest->m_area;
    factor /= PI * dist * dist;

    return factor;
}

cRadiosityCalc::sPatch *cRadiosityCalc::FindBrightest()
{
    sPatch *pBrightest = NULL;
    float brightest = 0.05f;

    float currIntensity;

    list<sPatch*>::iterator iter;

    // Blech. Linear search
```

```

sPatch *pCurr;
for(
    iter = m_patchList.begin();
    iter != m_patchList.end();
    iter++ )
{
    pCurr = *iter;

    currIntensity = pCurr->m_intensity;

    if( currIntensity > brightest )
    {
        brightest = currIntensity;
        pBrightest = pCurr;
    }
}

// This will be NULL if nothing was bright enough
return pBrightest;
}

bool cRadiosityCalc::CalcNextIteration()
{
    // Find the next patch that we need to
    sPatch *pSrc = FindBrightest();

    // If there was no patch, we're done.
    if( !pSrc )
    {
        DWORD diff = timeGetTime() - m_startTime;
        float time = (float)diff/1000;

        char buff[255];
        sprintf(
            buff,
            "Radiosity : Done - took %f seconds to render",
            time );
        SetWindowText( MainWindow()->GetHwnd(), buff );
        return false;    // no more to calculate
    }

    sPatch *pDest;
    list<sPatch*>::iterator iter;

    float formFactor;    // form factor Fi-j
    color3 deltaRad;    // Incremental radiosity shot from src to dest

    for(
        iter = m_patchList.begin();
        iter != m_patchList.end();
        iter++ )
    {
        pDest = *iter;

        // Skip sending energy to ourself

```

```
        if( pDest == pSrc )
            continue;

        // Compute the form factor
        formFactor = FormFactor( pDest, pSrc );

        // Early out if the form factor was 0.
        if( formFactor == 0.f )
            continue;

        // Compute the energy being sent from src to dest
        deltaRad = pDest->m_reflect * pSrc->m_deltaRad * formFactor;

        // Send said energy
        pDest->m_radiosity += deltaRad;
        pDest->m_deltaRad += deltaRad;

        // Cache the new intensity.
        pDest->m_intensity =
            pDest->m_area *
            (pDest->m_deltaRad.r +
             pDest->m_deltaRad.g +
             pDest->m_deltaRad.b );
    }
    // this patch has shot out all of its energy.
    pSrc->m_deltaRad = color3::Black;
    pSrc->m_intensity = 0.f;

    return true;
}
```

Conclusion

In this chapter we covered a lot of complex material. We started with parametric curves and subdivision before moving on through progressive meshes and radiosity lighting. Using the samples in this chapter as a base should allow you to create some really great effects in your games. Most of the code is fairly modular for you to plug into your own engines if you need to.

Now let's move on to more advanced topics like texturing.

Chapter 9

Advanced Direct3D

While I covered a lot of ground in Chapter 8, I really only scratched the surface of Direct3D's total set of functionality. By the end of this chapter, I'll have discussed everything you could ever want to know about texture mapping, along with alpha blending, multitexture effects, and the stencil buffer.

With Direct3D, there eventually comes a crest in the learning curve. At some point you know enough about the API that figuring out the rest is easy. For example, there comes a point when you've been bitten enough by setting the vertex shader parameters and zeroing out structures that you automatically do it. Hopefully, after learning the material in this chapter, you'll be over the hump. When you get there, learning the rest of the API is a breeze. Among the topics discussed in this chapter are:

- Alpha blending
- Texture mapping
- Pixel shaders
- Environment mapping
- Stencil buffers

Alpha Blending

Up to this point, I've been fairly dismissive of the mysterious alpha component that rides along in all of the D3DCOLOR structures. Now you may finally learn its dark secrets. A lot of power is hidden away inside the alpha component.

Loosely, the alpha component of the RGBA quad represents the opacity of a surface. An alpha value of 255 (or 1.0f for floating point) means the color is completely opaque, and an alpha value of 0 (or 0.0f) means the color is completely transparent. Of course, the value of the alpha component is fairly meaningless unless you actually activate the alpha blending step. If you want, you can set things up a different way, such as having 0 mean that the color is completely opaque. The meaning of alpha is dependent on how you set up the alpha blending step.

The alpha blending step is one of the last in the D3D output merger pipeline. As you rasterize primitives, each pixel that you wish to change in the frame buffer gets sent through the alpha blending step. That pixel is combined using blending factors to the pixel that is currently in the frame

buffer. You can add the two pixels together, multiply them together, linearly combine them using the alpha component, and so forth.

The Alpha Blending Equation

The equation that governs the behavior of the blending performed in Direct3D is defined as follows:

$$\text{final color} = \text{source} \times \text{source blend factor} + \text{destination} \times \text{destination blend factor}$$

Final color is the color that goes to the frame buffer after the blending operation. *Source* is the pixel you are attempting to draw to the frame buffer, generally one of the many pixels in a triangle you have told D3D to draw for you. *Destination* is the pixel that already exists in the frame buffer before you attempt to draw a new one. The *source blend factor* and *destination blend factor* are variables that modify how the colors are combined. The blend factors are the components you have control over in the equation; you cannot modify the positions of any of the terms or modify the operations performed on them.

For example, say you want an alpha blending equation to do nothing—to just draw the pixel from the triangle and not consider what was already there at all (this is the default behavior of the Direct3D output merger). An equation that would accomplish this would be:

$$\text{final color} = \text{source} \times 1.0 + \text{destination} \times 0.0$$

As you can see, the destination blending factor is 0 and the source blending factor is 1. This reduces the equation to:

$$\text{final color} = \text{source}$$

A second example would be if you wanted to multiply the source and destination components together before writing them to the frame buffer. This initially would seem difficult, as in the above equation they are only added together. However, the blending factors defined need not be constants; they can in fact be actual color components (or inverses thereof). The equation setup would be:

$$\text{final color} = \text{source} \times 0.0 + \text{destination} \times \text{source}$$

In this equation, the destination blend factor is set to the source color itself. Also, since the source blend factor is set to zero, the left-hand side of the equation drops away and you are left with:

$$\text{final color} = \text{destination} \times \text{source}$$

A Note on Depth Ordering

Usually if you are using a blending step that changes the color already in the depth buffer, you are attempting to use a semi-transparent surface, such as a puff of smoke or a fading particle effect. For the particle to appear correctly, the value already in the depth buffer must be what you would naturally see behind the specified primitive. For this to work correctly, you need to manually sort all of the alpha-blended primitives into a back-to-front list, drawing them after you draw the rest of the scene polygons. Using `qsort`, the STL generic sort algorithm, or something similar using the view space *z* value of the first vertex of each primitive as the sorting key will generally do the trick.

Enabling Alpha Blending

There are two ways to turn on alpha blending. The first, and best way usually, is in the technique in your shader. However, I'll show you the C++ way first so that the shader method makes more sense. Similar to other settings in Direct3D, such as the depth/stencil buffer, you need to create a *blend state*. The blend state maintains the current settings for how the output merger should use alpha blending as you render your primitives.

You could set up one state for alpha blending turned off, and another with it turned on with advanced effects. Once you've created your states you just have to make a single call to switch between them. This makes it very efficient for Direct3D 10 to change states as it massively reduces the amount of communication required with the graphics card.

Using Alpha Blending from C++

Each blend state you create must be described using a `D3D10_BLEND_DESC` structure, which has the following format:

```
typedef struct D3D10_BLEND_DESC {
    BOOL AlphaToCoverageEnable;
    BOOL BlendEnable[8];
    D3D10_BLEND SrcBlend;
    D3D10_BLEND DestBlend;
    D3D10_BLEND_OP BlendOp;
    D3D10_BLEND SrcBlendAlpha;
    D3D10_BLEND DestBlendAlpha;
    D3D10_BLEND_OP BlendOpAlpha;
    UINT8 RenderTargetWriteMask[8];
} D3D10_BLEND_DESC;
```

Table 9.1: D3D10_BLEND_DESC members

AlphaToCoverageEnable	Used for advanced blending techniques for things like foliage.
BlendEnable	A flag indicating whether or not blending should be enabled. There are eight settings, as Direct3D can render to up to eight render targets at once. Usually you'll only be interested in the first item.
SrcBlend	The source blend option, which defines how the output merger should treat the output from your pixel shader.
DestBlend	The destination blend option, which defines how the output merger should treat the pixel at the destination already in the render target.
BlendOp	The operation that should be performed to combine the source and destination pixels.
SrcBlendAlpha	The blend option defining how to treat the source alpha.
DestBlendAlpha	The blend option defining how to treat the destination alpha.
BlendOpAlpha	The operation defining how to combine the source and destination alphas.
RenderTargetWriteMask	A per-pixel mask that specifies how the destination buffer can be written to. Usually set to D3D10_COLOR_WRITE_ENABLE_ALL.

While this may initially look quite complicated, it's actually pretty easy, as you'll see soon.

Before we do that, let's take a quick look at the values this structure can be filled with. First up are the blend options, which can be any of the following:

Table 9.2: Blend options

D3D10_BLEND_ZERO	The data will be treated as (0,0,0,0).
D3D10_BLEND_ONE	The data will be treated as (1,1,1,1).
D3D10_BLEND_SRC_COLOR	The data will be the source color data from the pixel shader with no pre-blending.
D3D10_BLEND_INV_SRC_COLOR	The data will be the source color data from the pixel shader with an invert pre-blend operation.
D3D10_BLEND_SRC_ALPHA	The data is the alpha from the pixel shader with no pre-blending.
D3D10_BLEND_INV_SRC_ALPHA	The data is the alpha from the pixel shader with an invert pre-blend.
D3D10_BLEND_DEST_ALPHA	The data is the alpha from the destination target with no pre-blend operation.
D3D10_BLEND_INV_DEST_ALPHA	The data is the alpha from the destination target with an invert pre-blend operation.

D3D10_BLEND_DEST_COLOR	The data is the color from the render target with no pre-blend.
D3D10_BLEND_INV_DEST_COLOR	The data is the color from the render target with an invert pre-blend.
D3D10_BLEND_SRC_ALPHA_SAT	The data is the alpha from the pixel shader clamped to 1.0f or less.
D3D10_BLEND_BLEND_FACTOR	The data is taken from the function ID3D10Device::OMSetBlendState(), with no pre-blend.
D3D10_BLEND_INV_BLEND_FACTOR	The data is taken from the function ID3D10Device::OMSetBlendState(), with an invert pre-blend.
D3D10_BLEND_SRC1_COLOR	Used for dual-source rendering, which is beyond the scope of this chapter.
D3D10_BLEND_INV_SRC1_COLOR	Used for dual-source rendering, which is beyond the scope of this chapter.
D3D10_BLEND_SRC1_ALPHA	Used for dual-source rendering, which is beyond the scope of this chapter.
D3D10_BLEND_INV_SRC1_ALPHA	Used for dual-source rendering, which is beyond the scope of this chapter.

Most the time you'll only be using the first two options and D3D10_BLEND_SRC_ALPHA and D3D10_BLEND_INV_SRC_ALPHA. Now let's look at the operations that are used to blend with these options:

Table 9.3: Blend operations

D3D10_BLEND_OP_ADD	Add source 1 and source 2.
D3D10_BLEND_OP_SUBTRACT	Subtract source 1 from source 2.
D3D10_BLEND_OP_REV_SUBTRACT	Subtract source 2 from source 1.
D3D10_BLEND_OP_MIN	Use whichever of the two sources is smaller.
D3D10_BLEND_OP_MAX	Use whichever of the two sources is larger.

Once you fill out a D3D10_BLEND_DESC structure you can create a blend state with the function ID3D10Device::CreateBlendState(), which has the following prototype:

```
HRESULT CreateBlendState(
    const D3D10_BLEND_DESC *pBlendStateDesc,
    ID3D10BlendState **ppBlendState
);
```

The first parameter is the address of the blend structure you filled in, and the second fills in the address of an ID3D10BlendState pointer for you. When a blend state is created, you can set it as active using the function ID3D10Device::OMSetBlendState(), which looks like this:

```
void OMSetBlendState(
    ID3D10BlendState *pBlendState,
    const FLOAT BlendFactor[4],
    UINT SampleMask
);
```

The first parameter takes the pointer to the blend state you created with `CreateBlendState()`. The second parameter takes an array of up to four blend factors, one for each RGBA component. The final parameter takes an optional sample mask, which you can default to `0xFFFFFFFF`. After that, you are ready to rock. So now let's take a look at some minimal code to turn on generic alpha blending to render an object with transparency.

```
D3D10_BLEND_DESC descBlend;
memset(&descBlend, 0, sizeof(D3D10_BLEND_DESC));
descBlend.AlphaToCoverageEnable = false;
descBlend.BlendEnable[0] = true;
descBlend.SrcBlend = D3D10_BLEND_SRC_ALPHA;
descBlend.DestBlend = D3D10_BLEND_INV_SRC_ALPHA;
descBlend.SrcBlendAlpha = D3D10_BLEND_SRC_ALPHA;
descBlend.DestBlendAlpha = D3D10_BLEND_INV_SRC_ALPHA;
descBlend.BlendOp = D3D10_BLEND_OP_ADD;
descBlend.BlendOpAlpha = D3D10_BLEND_OP_ADD;
descBlend.RenderTargetWriteMask[0] = D3D10_COLOR_WRITE_ENABLE_ALL;

ID3D10BlendState *pBlendState = NULL;
Graphics()->GetDevice()->CreateBlendState(&descBlend, &pBlendState);
Graphics()->GetDevice()->OMSetBlendState(pBlendState, NULL, 0xFFFFFFFF);
```

Using Alpha Blending from Shaders

It's also very useful to know how to turn on and use alpha blending states from shaders because it is quicker and easier to do. The first thing you need to do in your effect file is define a shader state, which works almost the same way as in C++:

```
BlendState AlphaBlendingStateOn
{
    BlendEnable[0] = TRUE;
    SrcBlend = SRC_ALPHA;
    DestBlend = INV_SRC_ALPHA;
    BlendOp = ADD;
    SrcBlendAlpha = SRC_ALPHA;
    DestBlendAlpha = INV_SRC_ALPHA;
    BlendOpAlpha = ADD;
    RenderTargetWriteMask[0] = 0x0F;
};
```

Then in the technique before you render, you can set it as the current blend state.

```

////////////////////////////////////
// Default Technique
technique10 DefaultTechnique
{
    pass Pass0
    {
        SetBlendState( AlphaBlendingStateOn, float4( 0.0f, 0.0f, 0.0f, 0.0f ),
0xFFFFFFFF );

        SetGeometryShader(NULL);
        SetVertexShader(CompileShader(vs_4_0, DefaultVS()));
        SetPixelShader(CompileShader(ps_4_0, DefaultPS()));
    }
}

```

Check this out in the following figure, which shows a 50% transparent rabbit over a red surface. It creates a kind of weird orange result, although it appears to be a nice shade of gray in the book! I achieved this by forcing the alpha value from the pixel shader to be 0.5.




Figure 9.1: Alpha blending

Texture Mapping 101

It's kind of hard to think of texture mapping qualifying as advanced Direct3D material. Just about every 3D game that has come out in the last few years has used it, so it can't be terribly complex. When drawing your 3D objects with only a solid color (or even a solid color per vertex that is Gouraud shaded across the triangle pixels), they look rather bland and uninteresting. Objects in the real world have detail all over them, from the rings in woodgrain to the red and white pattern on a brick wall.

You could simulate these types of surfaces by increasing the triangle count a few orders of magnitude, and color each triangle so it could simulate things like woodgrain, bricks, or even the letters that appear on the keys of a keyboard. This, of course, is a terrible idea! You are almost always limited by the number of triangles you can feed the card per frame, so you cannot add the number of triangles you need to simulate that kind of detail. There must be a better way to solve the problem.

Really, what it comes down to is that you generally have the polygon budget to represent something like a brick wall with a handful of triangles. Instead of assigning a color to the vertices, you want to paint the *picture* of a brick wall onto the mesh. Then, at least from an appreciable distance (far enough that you can't notice the bumps and cracks in the wall), the polygons will look a lot like a brick wall.

Welcome to the world of texture mapping. A *texture* is just a regular image with some restrictions on it (such as having a power-of-two width and height). The name “texture” is kind of a misnomer; it does not represent what the uninitiated think of when they hear the word texture. Instead of meaning the physical feeling of a surface (being rough, smooth, etc.), texture in this context just means a special kind of image that you can map onto a polygon.

Fundamentals

Every texture-mapped polygon in our 3D space has a corresponding 2D polygon in texture space. Usually the coordinates for texture space are *u* (the horizontal direction) and *v* (the vertical direction). The upper-left corner of the texture is $\langle 0,0 \rangle$ and the bottom-right corner is $\langle 1,1 \rangle$, regardless of the actual size of the texture; even if the texture is wider than it is tall.

A Direct3D vertex shader is provided with the texture coordinates for the vertices of the triangles. It then interpolates across each pixel in the pixel shader in the triangle, finding the appropriate *u,v* pair, and then fetches that texture coordinate (or *texel*) and uses it as the color for that pixel. Figure 9.2 shows a visual representation of what happens when you texture a primitive.




Figure 9.2:
Texture coordinates
and rendering

Other factors can come into play, such as the diffuse color, multiple textures, and so forth, but I'll get to those later.

Affine vs. Perspective Mapping

To draw a primitive with a texture map, all you need to do is specify texture coordinates for each of the vertices of the primitive. The per-pixel texture coordinates can be found in one of two ways, called *affine mapping* and *perspective mapping*. Affine mapping is considered old technology and was used before there was the computing horsepower available to handle perspective mapping.

Affine mapping interpolates texture coordinates across each scan line of a triangle linearly. The u and v are interpolated the same way that r , g , and b are for Gouraud shading. Because of the simplicity (finding the delta- u and delta- v for a scan line, and then two adds per pixel to find the new u and v), affine mapping was very big in the days predating hardware acceleration.

However, betting that u and v vary linearly across a polygon is grossly incorrect. If the polygon is facing directly toward the viewer, then yes, u and v will vary linearly in relation to the pixels. However, if the polygon is on an angle, there is perspective distortion that prevents this from being true. The PlayStation 1 renders its triangles using affine rendering, and this can be really visible, especially when $1/z$ varies a lot over the space of a triangle. A good example is angled triangles near the camera, such as the ones at the bottom of the screen in racing games.

Perspective mapping, otherwise known as *perspective-correct* mapping, varies u and v correctly across the polygon, correcting for perspective distortion. The short mathematical answer is that while u and v do not vary linearly across a scan line, u/z , v/z , and $1/z$ do. If you interpolate all three of those values, you can find u and v by dividing u/z and v/z by $1/z$. A

division-per-pixel with a software renderer was impossible to do in real time in the old days, so most games found some way around it. *Quake*, for example, did the division every 16 pixels and did a linear interpolation between, which made the texture mapping look perfect in anything but extremely off-center polygons. With modern DirectX 10 acceleration, there is no need to worry; perspective mapping is just as fast as affine mapping on all modern cards.

Texture Addressing

The behavior for choosing texels at the vertices between 0..1 is pretty well defined, but what happens if you choose texels outside that range? How should Direct3D deal with it? This is a texture addressing problem. There are five different ways that Direct3D 10 can do texture addressing: wrap, mirror, mirror once, clamp, and border color. Each mode is described below.

Wrap

In wrap addressing mode, when a texel is selected past the end of the texture, it is wrapped around to the other side of the texture. The texel (1.3, -0.4) would be mapped to (0.3, 0.6). This makes the texture repeat itself like the posters on the walls of a construction site, as shown in Figure 9.3.




Figure 9.3:
Wrap addressing
mode

Care must be taken to make sure textures tile correctly when this addressing mode is used. If not, visible seams between copies of the texture will be visible, as you can see in Figure 9.3.

Mirror and Mirror Once

Mirror addressing mode flips texels outside of the (0..1) region so that it looks as if the texture is mirrored along each axis. This addressing mode can be useful for drawing multiple copies of a texture across a surface, even if the texture was not designed to wrap cleanly. The texel (1.3, -0.4) would be mapped to (0.7, 0.4), as shown in Figure 9.4. Mirror Once is similar except it only performs the mirror operation one time.




Figure 9.4:
Mirror addressing
mode

Clamp

Clamp mode is useful when you only want one copy of the texture map to appear on a polygon. All texture coordinates outside the (0..1) boundary are snapped to the nearest edge so they fall within (0..1). The texel (1.3, -0.4) would be mapped to (1.0, 0.0), as shown in Figure 9.5.




Figure 9.5:
Clamp addressing
mode

Unless the texture is created with a one-pixel boundary of a solid color around the edges, noticeable artifacts can occur (such as the streaks in the image).

Border Color

Border Color mode actually has two stage states to worry about: one to change the state to the addressing mode and one to choose a border color. In this addressing mode, all texture coordinates outside of the (0..1) region become the border color. See Figure 9.6.



Figure 9.6:
Border Color
addressing mode

Texture Wrapping

Texture wrapping is different from the texture addressing problem described above. Instead of deciding how texel coordinates outside the boundary of (0..1) should be mapped to the (0..1) area, it decides how to interpolate between texture coordinates. Usually, when the rasterizer needs to interpolate between two *u* coordinates (say, 0.1 and 0.8), it interpolates horizontally across the texture map, finding a midpoint of 0.45. When wrapping is enabled, it instead interpolates in the shortest direction. This would be to actually move from 0.1 to the left, wrap past 0.0 to 1.0, and then keep moving left to 0.8. The midpoint here would be 0.95.

To enable texture wrapping in your shader you create a sampler state, like the one below, and set the *AddressU* and *AddressV* members to *Wrap*:

```
SamplerState SamplerStateWrap
{
    Filter = MIN_MAG_MIP_LINEAR;
    AddressU = Wrap;
    AddressV = Wrap;
};
```

The *Filter* member is set to linear filtering, which you'll learn about later, along with how to use a sampler state.




Figure 9.7:
Examples of texture wrapping

Texture Aliasing

One of the biggest problems applications that use texture mapping have to deal with is *texture aliasing*. Texture aliasing is a smaller part of aliasing, which is another real problem in computer graphics. *Aliasing*, essentially, is when your image doesn't look the way you would expect; it looks like it was generated with a computer. Texture aliasing can take many forms. If you've ever heard of moiré effects, jaggies, blockies, blurries, texel swimming, or shimmering, you've heard about texture aliasing.

Why does texture aliasing occur? The short answer is because you're trying to discretely sample a signal (i.e., the texture on a polygon displayed on a set of pixels) that we would actually see as continuous (or as continuous as the resolution of our eyes can tell). Take the example of a texture that just had a horizontal repeating sinusoidal color variation on it. If you graphed the intensity as it related to the horizontal position on the screen, you would get something like Figure 9.8.

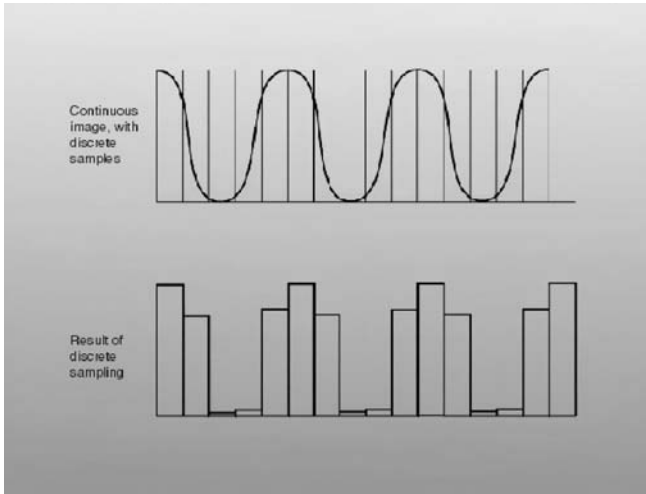


Figure 9.8:
A good result from discrete sampling

Notice that even though it is being sampled discretely, the sample points follow together well, and you can fairly closely approximate the continuous signal you're seeing. Problems start occurring when the signal changes faster than the discrete samples can keep up with. Take Figure 9.9, for example.




Figure 9.9:
A very poor result from discrete sampling

In this graph, the discrete samples don't approximate the continuous signal correctly, and you get a different signal altogether. As the frequency of the sine wave changes, the discrete signal you get varies widely, producing some really ugly effects. The ugly effects become even worse because the texture isn't actually a continuous signal; it's a discrete signal being

sampled at a different frequency. It's easy to imagine the sine function becoming the tiniest bit wider, so that each discrete sample met up with the crest of the sine wave. This tiny difference could happen over a couple of frames of a simulation (imagine the texture slowly moving toward the camera), and the resultant image would change from a wide variation of color to solid white!

If none of this is making sense, fire up an old texture-mapped game, such as *Doom* or *Duke Nukem 3D*. Watch the floors in the distance as you run around. You'll notice that the textures kind of swim around and you see lots of ugly artifacts. That's bad. That effect is what I'm talking about here.

MIP Maps

MIP mapping is a way for Direct3D to alleviate some of the aliasing that can occur by limiting the ratio of pixel size to texel size. The closer the ratio is to 1, the less texture aliasing occurs (because you're taking enough samples to approximate the signal of the texture).



Note: MIP is short for “multum in parvo,” which is Latin for “many things in a small place.”

Instead of keeping just one version of a texture in memory, we keep a chain of MIP maps. The top one is the original texture. Each successive one is half the size in each direction of the previous one (if the top level is 256x256 texels, the first MIP level is 128x128, the next is 64x64, and so on, down to 1x1).

MIP map surfaces can be created automatically using `ID3D10Device::CreateTexture2D()`. Just create a texture with the `MipLevels` member of the `D3D10_TEXTURE2D_DESC` structure set above 1. If you set `MipLevels` to 0, then a full chain down to the smallest size is created for you.

MIP levels can be generated in several ways. The most common is to simply sample each 4x4 pixel square in one MIP level into one pixel of the MIP level below it, averaging the four color values. Luckily, since you're loading DDS (Direct3D Surface) texture files, this is done automatically.




Figure 9.10: MIP maps in action

Filtering

Filtering, or the way in which you get texels from the texture map given a u,v coordinate pair, can dramatically affect the way the final image turns out. The filtering problem is divided into two separate issues: magnification and minification.

Magnification occurs when you try to map a single texel in a texture map to many pixels in the frame buffer. For example, if you were drawing a 64×64 texture onto a 400×400 pixel polygon, the image would suffer the torment of magnification artifacts. Linear filtering helps get rid of these artifacts.

Minification (I didn't make up this word) is the opposite problem—when multiple texels need to be mapped to a single pixel. If you were instead drawing that 64×64 texture onto a 10×10 pixel polygon, our image would instead be feeling the pain from minification. Swimming pixels, such as the type discussed in the texture aliasing discussion above, are tell-tale symptoms.

Direct3D 10 hardware can use three different varieties of filtering to alleviate magnification and minification artifacts: point sampling, linear filtering, and anisotropic filtering. Let's take a look at each of them.


Point Sampling

Point sampling is the simplest kind of filter. In fact, it's hard to think of it as being a filter at all. Given a floating-point (or fixed-point) u,v coordinate, the coordinates are snapped to the nearest integer and the texel at that coordinate pair is used as the final color.

Point sampling suffers from the most aliasing artifacts. If MIP mapping is used, these artifacts can be alleviated somewhat. The PlayStation console uses point sampling for its texture mapping, as did the first generation of 3D games (*Descent* and *Quake*). *Quake* got past some of the visual

artifacts of point sampling by using MIP maps, selecting the MIP map based on distance, and point sampling out of that. However, since no filtering is done between MIP map levels, if you run toward a wall from a far-off distance, you can actually see the MIP level switch as the distance decreases. Figure 9.11 shows worst-case point sampling, a checkerboard pattern with no MIP mapping.

The artifacts caused by point sampling are readily visible in the distance. As the ratio between the signal and the discrete sampling changes, the output signal changes completely, giving rise to the visible banding artifacts.



Courtesy of Paul Heckbert


Figure 9.11:
Worst-case point sampling

Linear Filtering

One step up from point sampling is linear filtering. Instead of snapping to the nearest integer coordinate, the four nearest texels are averaged together based on the relative distances from the sampling point. The closer the ideal coordinate is to an integer coordinate, the more you weight it. For example, if you wanted a texel for the coordinate (8.15,2.75), the result would be:

$$\begin{aligned} \text{Result pixel} = & (1 - 0.15) \times (1 - 0.75) \times \text{Texel}(8,2) + \\ & 0.15 \times (1 - 0.75) \times \text{Texel}(9,2) + \\ & (1 - 0.15) \times 0.75 \times \text{Texel}(8,3) + \\ & 0.15 \times 0.75 \times \text{Texel}(9,3) \end{aligned}$$

Linear filtering can improve image quality a lot, especially if it is combined with MIP mapping. All DirectX 10 hardware can handle linear filtering with MIP maps, so it is used the most often. If MIP maps aren't used, however, it only looks marginally better than point sampling, as evidenced by Figure 9.12.



Courtesy of Paul Heckbert

Figure 9.12:
Linear filtering

Anisotropic Filtering

A problem that arises in linear filtering is that texels are sampled using square sampling. This works well if the polygon is facing directly toward the viewer, but doesn't work properly if the polygons are angled sharply away from the viewer. For example, think of the point of view of a chicken crossing a road. If you could imagine each pixel as a tall, thin pyramid being shot out of the chicken's eye, when the pyramid intersected the road it wouldn't be a square at all. Figure 9.13 illustrates this concept using a circular sample region.

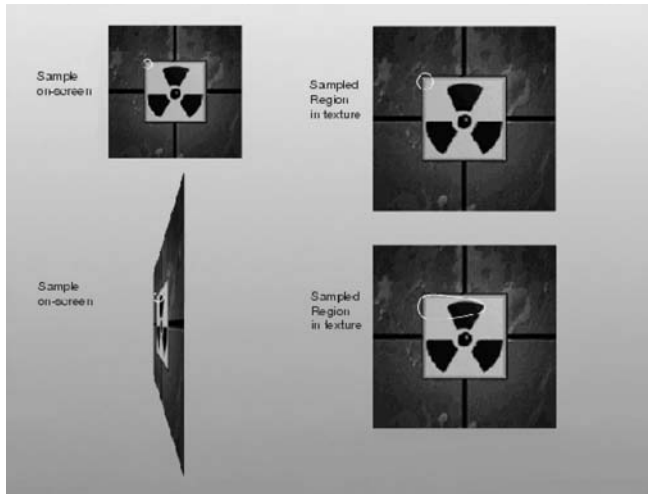


Figure 9.13:
Ideal texture
sampling

Anisotropic filtering looks about as good as it can get, as evidenced in Figure 9.14. Just about all of the banding artifacts or blurring artifacts are gone, and the image looks more and more like a real photograph taken of an infinite checkerboard.




Figure 9.14:
Anisotropic filtering

Textures in Direct3D

Generally, most scenes have many more textures than are visible at any one time. Also, all of the textures in the scene usually can't fit in the video card's texture memory. This, of course, presents a dilemma: How should an application make sure that the textures it currently needs for drawing are loaded into the video card?

To solve this problem you need a subsystem to handle texture management. Whenever you want to draw a primitive with a certain texture, the texture management subsystem makes sure that the texture is available and in memory. If it isn't, the texture is uploaded to the card and another texture is evicted from memory to make space for it.

Which texture do you evict? You can't be haphazard about it; uploading textures is expensive and should be done as little as possible. If you know for certain that a texture won't be used for a while (say it's a special texture used in an area that you're far away from), then you can evict it, but generally such information isn't available. Instead, usually the texture that hasn't been used in the longest time, or the least recently used texture (LRU), is evicted. This system is used almost everywhere where more data is needed than exists places to store it (disk caches, for example). Of course, textures continue to be evicted until there is enough space to store the desired texture.

Direct3D, thankfully, has an automatic texture management system. While it's not as fast as a more specific texture management system would be, for these purposes it's all that's needed.

You can't be too reckless with your newfound power, however. If you tried to draw a scene with 512 MB of textures simultaneously visible on a card with only 256 MB of texture RAM, that would mean that the texture management engine would need to download 256 MB of textures to the graphics card every frame. This can be an expensive operation, and can just murder your frame rate. In situations like this, it is a better bet to

lower the resolutions of your textures so everything can fit on the card at once.

Direct3D lets you control the way in which managed textures are evicted. Direct3D holds onto a time stamp of when each managed texture was used, as well as a priority number for each. When Direct3D needs to remove a texture, it removes the one with the lowest priority, removing the least recently used if there is more than one candidate. You can set the priority of a texture using `SetEvictionPriority()`. There is also a corresponding call to get the priority called `GetEvictionPriority()`.

```
void SetEvictionPriority(
    UINT EvictionPriority
);
```

EvictionPriority	This can be set to any of the following:
	<ul style="list-style-type: none"> • <code>DXGI_RESOURCE_PRIORITY_MINIMUM</code> • <code>DXGI_RESOURCE_PRIORITY_LOW</code> • <code>DXGI_RESOURCE_PRIORITY_NORMAL</code> • <code>DXGI_RESOURCE_PRIORITY_HIGH</code> • <code>DXGI_RESOURCE_PRIORITY_MAXIMUM</code>

Being able to use texture management makes our lives much easier. At startup we just need to load all the textures we want to use, and Direct3D will take care of everything. If there are too many, the extra ones will invisibly sit in system RAM surfaces until they are used, and then they'll get uploaded to the card.

Texture Loading

Just like playing sound effects, in order to do anything interesting with your newly learned texture mapping skills, you need to actually have some data to work with. Be aware that loading certain types of textures can be avoided by generating them algorithmically.

DDS Format

Loading the DDS (or Direct3D Surface) format couldn't be easier. The first four bytes contain a magic number that describes it as a DDS surface (0x20534444, or "DDS" in ASCII). Next is a `DDSURFACEDESC2` structure describing the surface. After that comes the raw data of the surface. If there are any attached surfaces (MIP maps, for example), the data for them is provided after the main surface.

DDS surfaces can be created using the DirectX Texture Tool, which comes with the SDK. It can load a BMP file (and load another into the alpha channel, if desired), generate MIP maps, and save the result as a DDS texture. You can also create compressed DDS texture files. However, you don't need to worry about the intricacies of the surface format because

since version 8.0 it is easy to load DDS files with the `D3DXCreateTextureFromFile()` utility function.




Figure 9.15: The DirectX Texture Tool

The *cTexture* Class

To facilitate the loading of Direct3D textures, I'm going to create a class that wraps around a Direct3D texture object, specifically for use as a texture. The class is fairly simple.

```

/*****
 *      Advanced 3D Game Programming with DirectX 10.0
 *      * * * * *
 *
 *      See license.txt for modification and distribution information
 *      copyright (c) 2007 by Peter Walsh, Wordware
 *****/

#ifndef _TEXTURE_H
#define _TEXTURE_H

#include <string>

class cGraphicsLayer;

class cTexture
{
protected:

    void ReadDDSTexture( ID3D10Texture2D **pTexture );

    ID3D10Texture2D    *m_pTexture;

```

```

        ID3D10ShaderResourceView *m_pShaderResourceView;

        wstring      m_name;

        // The stage for this particular texture.
        DWORD        m_stage;

public:
    cTexture( const TCHAR *filename, DWORD stage = 0 );
    virtual ~cTexture();

    ID3D10Texture2D *GetTexture();

    ID3D10ShaderResourceView *GetShaderView();

};

#endif // _TEXTURE_H

/*****
 *      Advanced 3D Game Programming with DirectX 10.0
 *      * * * * *
 *
 *      See license.txt for modification and distribution information
 *      copyright (c) 2007 by Peter Walsh, Wordware
 *****/

#include "stdafx.h"
#include "Texture.h"
#include "GraphicsLayer.h"

using std::vector;

cTexture::cTexture( const TCHAR *filename, DWORD stage )
{
    m_pShaderResourceView = NULL;
    ID3D10Texture2D *pTempTex = 0;

    m_pTexture = 0;

    m_name = wstring( filename );
    m_stage = stage;

    ReadDDSTexture( &m_pTexture );
}

cTexture::~cTexture()
{
    SafeRelease( m_pTexture );
}

void cTexture::ReadDDSTexture( ID3D10Texture2D **pTexture )
{

```

```

HRESULT r = 0;

r = D3DX10CreateTextureFromFile(
    Graphics()->GetDevice(),
    m_name.c_str(),
    NULL,
    NULL,
    (ID3D10Resource**)pTexture,
    NULL);

if( FAILED(r) )
{
    throw cGameError( L"Bad DDS file\n");
}

ID3D10Texture2D *cTexture::GetTexture()
{
    return m_pTexture;
}

ID3D10ShaderResourceView *cTexture::GetShaderView()
{
    if(!m_pShaderResourceView)
    {
        D3D10_TEXTURE2D_DESC descTexture;
        m_pTexture->GetDesc(&descTexture);

        D3D10_SHADER_RESOURCE_VIEW_DESC descShaderView;
        memset(&descShaderView, 0, sizeof(descShaderView));
        descShaderView.Format = descTexture.Format;
        descShaderView.ViewDimension = D3D10_SRV_DIMENSION_TEXTURE2D;
        descShaderView.Texture2D.MipLevels = descTexture.MipLevels;

        Graphics()->GetDevice()->CreateShaderResourceView(
            m_pTexture, &descShaderView, &m_pShaderResourceView);
    }

    return m_pShaderResourceView;
}

```

I'll discuss the shader view code in a moment.

Activating Textures

With all this talk of texture management and texture addressing, filtering, and wrapping, I still haven't described how to activate a texture for use! After creating a texture surface, filling it with image data, setting all the rendering states we want, and so forth, to actually activate it you need to follow a few steps.

First we need to create a *shader view* of the texture so the shader knows how to deal with it. Then we need to add a texture variable to our

shader and set up that variable to be linked to the texture we loaded. When that is done, we will be able to sample it for rendering.

Creating a Shader View

A shader view of a texture allows your texture to be accessed from a shader. To create the shader view in the previous code listing I added the function `GetShaderView()` to `cTexture`. The first time it is called it creates a shader view, and every subsequent time it just returns a saved pointer to the view. To create a view you fill out a `D3D10_SHADER_RESOURCE_VIEW_DESC` structure, which looks like this:

```
typedef struct D3D10_SHADER_RESOURCE_VIEW_DESC {
    DXGI_FORMAT Format;
    D3D10_SRV_DIMENSION ViewDimension;
    union {
        D3D10_BUFFER_SRV Buffer;
        D3D10_TEX1D_SRV Texture1D;
        D3D10_TEX1D_ARRAY_SRV Texture1DArray;
        D3D10_TEX2D_SRV Texture2D;
        D3D10_TEX2D_ARRAY_SRV Texture2DArray;
        D3D10_TEX2DMS_SRV Texture2DMS;
        D3D10_TEX2DMS_ARRAY_SRV Texture2DMSArray;
        D3D10_TEX3D_SRV Texture3D;
        D3D10_TEXCUBE_SRV TextureCube;
    };
} D3D10_SHADER_RESOURCE_VIEW_DESC;
```

As you can see from the listing, you can get all the information to fill out this structure from the description of the texture. When it is filled in you call `ID3D10Device::CreateShaderResourceView()`, which looks like this:

```
HRESULT CreateShaderResourceView(
    ID3D10Resource *pResource,
    const D3D10_SHADER_RESOURCE_VIEW_DESC *pDesc,
    ID3D10ShaderResourceView **ppSRView
);
```

The first parameter takes a pointer to the resource that you want to create a view of, which in this case is the pointer to the texture. The second parameter takes a pointer to the view description created from the texture description, and the final parameter returns a pointer to the created resource view, which you'll need to hang on to in order to set it up in the shader.

Adding Textures to the Shader

Now that we have a texture ready to use on the C++ side we need to add a texture variable to our shader to hold the texture. You can do this in HLSL with the `Texture2D` variable type. Here I declare a single texture variable called `g_Textures`:


```
Texture2D g_Textures;
```

Sending the Texture to the Shader

Now we need to add some code to `cGraphicsLayer` that will send a texture to the shader. This is just like the variables we added previously that allowed us to update the world and view matrices, except this time we will create a pointer of type `ID3D10EffectShaderResourceVariable`. Here is the one I added to `cGraphicsLayer`:

```
ID3D10EffectShaderResourceVariable *m_pTexturesVar; // Pointer to the texture
                                                    // variable
```

This is created in the function `cGraphicsLayer::CreateDefaultShader()` with this code:

```
m_pTexturesVar = m_pDefaultEffect->GetVariableByName( "g_Textures"
)->AsShaderResource();
```

Now we can use this variable to send textures to the shader with this utility function:

```
void cGraphicsLayer::SetTexture(ID3D10ShaderResourceView *pTexShaderView)
{
    if(m_pTexturesVar)
    {
        m_pTexturesVar->SetResource(pTexShaderView);
    }
}
```

Texture Sampling

When you have all your texture data and texture coordinates set up and your pixel shader runs, it is very easy to sample the color from the texture. You just use the HLSL command `sample()`, which takes two parameters. The first parameter is the sampler to use, which we looked at earlier to define the wrapping mode and filtering method. The second parameter takes the coordinates of the texture to sample. For example, the following code samples using a sampler called `SamplerStateWrap` and the coordinates sent through from the vertex shader:

```
float4 DefaultPS(VS_OUTPUT dataIn) : SV_Target
{
    float4 texColor = (g_Texture.Sample(SamplerStateWrap, dataIn.vTexCoords1);

    return texColor;
}
```

Texture Mapping 202

In this day and age, plain-vanilla texture mapping is not enough. There are many neat effects that you can create using multiple passes, and in this section I'll discuss how to use effects to make objects that look really cool by combining multiple textures together in the shader.

In previous versions of DirectX, this was one of the most complicated tasks to do. You had to set up hundreds of render states, texture transform states, sampler states, and so on until you were blue in the face. These days it's much easier with the power of HLSL. You just send all your textures to the pixel shader and then combine them any way you want.

Texture Arrays

If you are using multiple textures, it can be easier and more useful to set up texture arrays in your shader. The process is very similar to using a single texture. Instead of declaring a single texture like before, you can define any number of textures like this:

```
Texture2D g_Textures[8];
```

You set up the variable in C++ in the same way, except instead of calling `ID3D10EffectShaderResourceVariable::SetResource()`, you call `ID3D10EffectShaderResourceVariable::SetResourceArray()`, which looks like this:

```
HRESULT SetResourceArray(
    ID3D10ShaderResourceView **ppResources,
    UINT Offset,
    UINT Count
);
```

The first parameter takes an array of pointers to your textures. The second parameter takes the offset into the array that you want to set. The final parameter takes the number of items in the array that you are setting.

Effects Using Multiple Textures

Most modern games now use multiple textures per primitive for a variety of effects. While there are many more possible kinds of effects than can be described here, I'm going to run through the most common ones and show how to implement them using both multiple textures per pass and multiple passes.

The way you combine textures and the way you make the textures defines the kind of effect you end up with. Using multitexturing is preferred. Since you only draw the primitive once, it ends up being faster than multipass. Multipass involves drawing each of the separate phases of the effect one at a time. Generally, you change the texture, change the alpha blending effects, and then redraw the primitive. The new pass will be

combined with the previous pass pixel-by-pixel. Figure 9.16 helps explain the kinds of things I'm trying to do. Using multitexture, you would set the first stage to texture A and the second stage to texture B, and then set the operation in texture B to either add, multiply, or subtract the pixels. Using multipass, you would draw texture A first, then change the alpha blending steps to add or multiply the pixels together, and then draw the polygon again using texture B. With DirectX 10 class hardware it is rare to use multipass rendering unless it is for extremely complex effects.

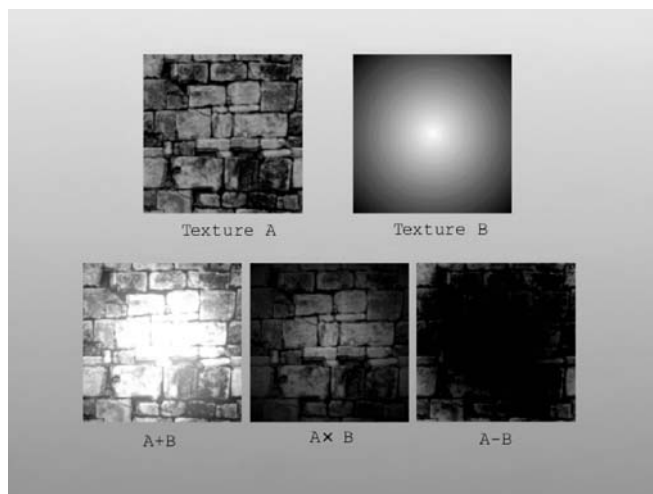


Figure 9.16:
Combining textures
for more advanced
effects

Light Maps (a.k.a. Dark Maps)

Light mapping is practically a standard feature for first-person shooters these days. It allows the diffuse color of a polygon to change non-linearly across the face of a polygon. This is used to create effects like colored lights and shadows.

Using a light-map creation system (usually something like a radiosity calculator, which I created in Chapter 8), texture maps that contain just lighting information are calculated for all of the surfaces in the scene. Since usually the light map doesn't change per pixel nearly as much as the texture map, a lower-resolution texture is used for the light map. *Quake*-style games use about 16^2 texels of light map for each texel of texture map. The base map is just the picture that would appear on the wall if everything were fully and evenly lit, like wallpaper. The light map is modulated (multiplied) with the base map. That way, areas that get a lot of light (which appear white in the light map) appear as they would in the fully lit world (since the base map pixel times white (1) resolves to the base map). As the light map gets darker, the result appears darker. Since a light map can only darken the base map, not lighten it, sometimes the effect is referred to as *dark mapping*.

When you go to draw the polygon, you can do it in several ways. Here is an example in a pixel shader:

```
float4 MT_PS(VS_OUTPUT dataIn) : SV_Target
{
    // Diffuse lighting
    float4 diffuseColor = saturate(dot(vLightDir, dataIn.vNormal) * vSunColor);

    // Texture and light map color
    float4 textureColor = texture.Sample(SamplerStateWrap, dataIn.vTexCoords1);
    float4 lightMapColor = lightmapTex.Sample(SamplerStateWrap,
        dataIn.vTexCoords2);

    // Combine the colors
    float4 finalColor = diffuseColor * textureColor * lightMapColor;

    return finalColor;
}
```

The visual flair that you get from light mapping is amazing. Following is a prime example from *Quake III: Arena*. Figure 9.17 is rendered without light maps. The image looks bland and uninteresting. Figure 9.18 shows the same scene with light mapping enabled. The difference, I'm sure you'll agree, is amazing.



Courtesy of id Software

Figure 9.17:
Quake III: Arena without light maps




Figure 9.18:
Quake III: Arena with light maps


Environment Maps

Environment mapping was one of the first cool effects people used texture maps with. The concept is quite simple: You want a polygon to be able to reflect the scene, as if it were a mirror or shiny surface like chrome. There are two primary ways to do it that Direct3D supports: spherical environment maps and cubic environment maps.

Spherical Environment Maps

Spherical environment maps are one of those classic horrible hacks that happens to look really good in practice. It isn't a perfect effect, but it's more than good enough for most purposes.

The environment mapping maps each vertex into a *u,v* pair in the spherical environment map. Once you have the locations in the sphere map for each vertex, you texture map as normal. The sphere map is called that because the actual picture looks like the scene pictured on a sphere. Real photos are taken with a 180-degree field of view camera lens, or using a ray-tracer to prerender the sphere map. Rendering a texture like this is complex enough that it is infeasible to try to do it in real time; it must be done as a preprocessing step. An example of a sphere map texture appears in Figure 9.19.



Courtesy of nVidia Corporation

Figure 9.19:
A spherical texture

The region outside of the circle in the above image is black, but it can be any color; you're never actually going to be addressing from those coordinates, as you'll see in a moment.



Note: You can use any texture as a sphere map. Just load it up into Photoshop and run it through the Polar Coordinates filter.

Once you have the spherical texture map, the only task left to do is generate the texture coordinates for each vertex. Here comes the trick that runs the algorithm:

The normal for each vertex, when transformed to view space, will vary along each direction from -1 to 1 . What if you took just the x and y components and mapped them to $(0,1)$? You could use the following equation:

$$u = \frac{n_x + 1}{2}$$

$$v = \frac{n_y + 1}{2}$$

You know that the radius of the 2D vector $\langle n_x, n_y \rangle$ will vary between 0 (when z is 1.0 and the normal is facing directly toward the viewer) and 1 (when z is 0.0 and the normal is perpendicular to the viewer). When n_x and n_y are 0 , you'll get a u,v pair of $\langle 0.5, 0.5 \rangle$. This is exactly what was

wanted: The vertex whose normal is pointing directly toward us should reflect the point directly behind us (the point in the center of the sphere map). The vertices along the edges (with radius 1.0) should reflect the regions on the edge of the sphere map. This is exactly what happens.

As you can see in Figure 9.20, this environment mapping method can give really nice-looking results.




Courtesy of nVidia Corporation

Figure 9.20: In a lot of cases spherical environment mapping looks great.

One caveat of this rendering method is that the sphere map must remain the same, even if the camera moves. Because of this, it often isn't useful to reflect certain types of scenes; it's best suited for bland scenery like starscapes.

There are some mechanisms used to attempt to interpolate correct positions for the spherical environment map while the camera is moving, but they are far from perfect. They suffer from precision issues; while texels in the center of the sphere map correspond to relatively small changes in normal direction, near the edges there are big changes, and an infinite change when you reach the edge of the circle. This causes some noticeable artifacts, as evidenced in Figure 9.21. Again, these artifacts only pop up if you try to find the sphere map location while the camera is moving. If your sphere map setup does not change, none of this happens.



Courtesy of nVidia Corporation

Figure 9.21:
Warping artifacts

Calculating Spherical Coordinates in HLSL

Check out this code example, which calculates the spherical coordinates and saves them in `vTexCoords2`.

```
VS_OUTPUT MT_VS(VS_INPUT dataIn)
{
    VS_OUTPUT result;

    float4 vPos = float4(dataIn.vPosition, 1.0f);

    float4 vWorldPos = mul(vPos, g_mtxWorld);
    float4 vViewPos = mul(vWorldPos, g_mtxView);
    float4 vScreenPos = mul(vViewPos, g_mtxProj);

    result.vPosition = vScreenPos;
    result.vColor = dataIn.vColor;
    result.vTexCoords1 = dataIn.vTexCoords1;

    float3 vViewNormal =
        normalize(mul(mul(dataIn.vNormal, g_mtxWorld), g_mtxView));

    result.vNormal = vViewNormal;

    float3 vReflection = normalize(reflect(vViewPos, vViewNormal));
```



```
float fTemp = 2.0f * sqrt(
    vReflection.x * vReflection.x +
    vReflection.y * vReflection.y +
    (vReflection.z+1.0f) * (vReflection.z+1.0f)
);

result.vTexCoords2.x = vReflection.x / fTemp + 0.5f;
result.vTexCoords2.y = vReflection.y / fTemp + 0.5f;

return result;
}
```

Cubic Environment Maps

Cubic environment maps have been used in high-end graphics workstations for some time, and they have a lot of advantages over spherical environment maps.

The biggest advantage is that cubic environment maps don't suffer from the warping artifacts that plague spherical environment maps. You can move around an object, and it will correctly reflect the correct portion of the scene. Also, they're much easier to make, and in fact can be made in real time (producing accurate real-time reflections). You can even use the render target as an environment map showing the reflections (such as off a shiny car) to accurately reflect what the surrounding environment looks like.

A cubic environment map is actually a complex Direct3D texture with six different square textures, one facing in each direction. They are:

- Map 0: +X direction (+Y up, -Z right)
- Map 1: -X direction (+Y up, +Z right)
- Map 2: +Y direction (-Z up, -X right)
- Map 3: -Y direction (+Z up, -X right)
- Map 4: +Z direction (+Y up, +X right)
- Map 5: -Z direction (+Y up, +X right)

The six environment maps that are used in the images for this section appear in Figure 9.22.




Figure 9.22:
The six components of a cubic environment map

How do you actually use this environment map to get texture coordinates for each of the vertices? The first step is to find the reflection vector for each vertex. You can think of a particle flying out of the camera and hitting the vertex. The surface at the vertex has a normal provided by the vertex normal, and the particle bounces off of the vertex back into the scene. The direction in which it bounces off is the reflection vector, and it's a function of the camera to vertex direction and vertex normal. The equation to find the reflection vector r is:

$$d = \frac{v - c}{\|v - c\|}$$


$$r = 2 \times (d \cdot n)n - d$$

where r is the desired reflection vector, v is the vertex location, c is the camera location, and n is the vertex normal. The d vector is the normalized direction vector pointing from the camera to the vertex.

Given the reflection vector, finding the right texel in the cubic environment map isn't that hard. First, you find which component of the three has the greatest magnitude (let's assume it's the x component). This determines which environment map you want to use. So if the absolute value of the x component was the greatest and the x component was also negative, you would want to use the $-X$ direction cubic map (map 1). The other two components, y and z in this example, are used to index into the map. We scale them from the $[-1, 1]$ range to the $[0, 1]$ range. Finally, you use z to choose the u value and y to choose the v value.

To implement cubic environment mapping you just need to implement this formula in HLSL, similar to how I did the spherical mapping in the previous section. There is some interesting literature out on the web for advanced mapping effects.

The sphere you saw being spherically environment mapped earlier appears with cubic environment mapping in Figure 9.23. Notice that all of the artifacts are gone and the sphere looks pretty much perfect.



Courtesy of nVidia Corporation

Figure 9.23:
Cubic environment mapping

Note that to use cubic mapping you would use a `Texture3D` HLSL variable in place of `Texture2D` and sample with 3D texture coordinates.

Specular Maps

The lighting you can approximate with multitexture isn't limited to diffuse color. Specular highlights can also be done using multitexture. It can do neat things that specular highlights done per-vertex cannot, like having highlights in the middle of a polygon.

A specular map is usually an environment map like the kind used in spherical environment mapping that approximates the reflective view of the lights in our scene from the viewpoint of an object's location. Then you just perform normal spherical (or cubic) environment mapping to get the specular highlights.

The added advantage of doing things this way is that some special processing can be done on the specular map to achieve some neat effects. For example, after creating the environment map, you could perform a blur filter on it to make the highlights a little softer. This would approximate a slightly matte specular surface.

Detail Maps

A problem that arises with many textures is that the camera generally is allowed to get too close to them. Take, for example, the texture in Figure

9.24. From a standard viewing distance (15 or 20 feet away), this texture would look perfectly normal on an 8- to 10-foot tall wall.



Figure 9.24:
An example wall texture

However, a free-moving camera can move anywhere it likes. If you position the camera to be only a few inches away from the wall, you get something that looks like Figure 9.25. With point sampling, we get large, ugly, blocky texels. With linear filtering the problem is even worse: You get a blurry mess.




Figure 9.25:
Getting too close

This problem gets really bad in things like flight simulators. The source art for the ground is designed to be viewed from a distance of 30,000 feet above. When the plane is dipping close to the ground, it's almost impossible to correctly gauge distance; there isn't any detail to help you judge how far off the ground it is, resulting in a poor visual experience.

You can use larger textures in the scene, but then you need to page to system RAM more, load times are longer, etc. For this entire headache all you get is improved visual experience for an anomalous occurrence anyway; most of the user's time won't be spent six inches away from a wall or flying three feet above the ground.

What this problem boils down to is the designed signal of an image. Most textures are designed to encode low-frequency signals, the kind that change over several inches. The general color and shape of an image are examples of low-frequency signals.

The real world, however, has high-frequency signals in addition to these low-frequency signals. These are the little details that you notice when you look closely at a surface, the kind that change over fractions of an inch. The bumps and cracks in asphalt, the grit in granite, and the tiny variations of grain in a piece of wood are all good examples of high-frequency signals.

While you could hypothetically make all of the textures 4096 texels on a side and record all of the high-frequency data, you don't need to. The high-frequency image data is generally quite repetitive. If you make it tile correctly, all you need to do is repeat it across the surface. It should be combined with the base map, adding detail to it (making areas darker or lighter).

Figure 9.26 shows the detail map that you'll use in the application coming up in a little bit. The histogram of the image is tightly centered around solid gray (127,127,127). You'll see why in a moment. Also, it's




Figure 9.26:
The detail map

designed without lots of sharp visual distinctions across the surface, so any details quickly fade away with MIP level increases.

If you tile the high-frequency detail map across the low-frequency base map, you can eliminate the blurry artifacts encountered before. As an added bonus, after you get far enough away from a surface, the MIP level for the detail map will be solid gray so you can actually turn it off, if you'd like, for faraway surfaces. Doing this improves the performance penalty. Figure 9.27 shows the base map with the detail map applied.

This essentially does an addition, with one of the textures having signed color values ($-127..128$) instead of the unsigned values ($0..255$) that you're used to. Black corresponds to -127 , white corresponds to 128 , and solid gray corresponds to 0 . If the second texture map is a solid gray image (like the detail map at a low MIP map level), the result of the blend is just the other texture.




Figure 9.27:
The base map combined
with the detail map

Implementing detail maps is as simple as sampling from two textures in the pixel shader and combining them, although the detail map usually has a reduction factor applied so it doesn't overpower the main image. The texture coordinates set up in the vertices take care of the repetition.

Application: Detail

To show off the texture loading code I set up earlier in the chapter and explain detail textures, I threw together a simple application that shows the base map/detail map combo used throughout this section. The application uses the new version of the GameLib library (which has texture support). A screenshot from the application appears in Figure 9.28. Note

that unlike previous versions of DirectX, with DirectX 10 there is no need to check the caps as all DX10 hardware supports the full minimum feature set, apart from a few obscure items.



Figure 9.28:
Alpha blending
for detail

Here is the code:

```

/*****
 *      Advanced 3D Game Programming with DirectX 10.0
 *      *      *      *      *      *      *      *      *      *      *
 *
 *      See license.txt for modification and distribution information
 *      copyright (c) 2007 by Peter Walsh, Wordware
 *****/

#include "stdafx.h"

//Detail Texture example app
//Advanced 3D Game Programming with DirectX 10.0
//
// Controls:
//      [D] : Turn detail textures on/off

class cDetailApp : public cApplication,
                  public iKeyboardReceiver
{
    // Data for the square
    cGraphicsLayer::cDefaultVertex m_square[6];
    void LoadTextures();

    void SetStage1();
    void SetStage2();
}

```

```

    cTexture    *m_pTextures[2];

    // If this is true, draw some help on the screen
    bool m_drawHelp;

    void InitSquare();
    void DrawSquare();

    bool m_bCanDoMultitexture;
    bool m_bCanDoAddSigned;
    bool m_bDrawDetail;

public:

    //=====----- cApplication

    virtual void DoFrame( float timeDelta );
    virtual void SceneInit();

    cDetailApp() :
        cApplication()
    {
        m_pTextures[0] = NULL;
        m_pTextures[1] = NULL;
        m_pTextures[2] = NULL;
        m_title = wstring( L"Detail Textures" );
    }

    ~cDetailApp()
    {
        delete m_pTextures[0];
        delete m_pTextures[1];
    }

    //=====----- iKeyboardReceiver

    virtual void KeyUp( int key );
    virtual void KeyDown( int key );

    //=====----- cDetailApp
};

cApplication *CreateApplication()
{
    return new cDetailApp();
}

void DestroyApplication( cApplication *pApp )
{
    delete pApp;
}

void cDetailApp::SceneInit()
{

```



```

    Input()->GetKeyboard()->SetReceiver( this );

    /// initialize the camera
    Graphics()->SetViewMtx( *(D3DXMATRIX*)&matrix4::Identity );

    // do draw the detail
    m_bDrawDetail = true;

    /// init our geometry
    InitSquare();

    /// load our textures
    LoadTextures();
}

void cDetailApp::DoFrame( float timeDelta )
{
    float colClear[4] = {1.0f, 1.0f, 1.0f, 1.0f};
    Graphics()->Clear(colClear);
    Graphics()->ClearDepthStencil(1.0f, 0);

    static float total =0.f;
    total += timeDelta;
    //
    float dist = 1.1f + 40.f*(1.f + sin(total));

    if(m_bDrawDetail)
        Graphics()->SetNumActiveTexture(1, 0);
    else
        Graphics()->SetNumActiveTexture(2, 0);

    matrix4 mat;
    mat.ToTranslation( point3(0.0f, 0.0f, dist) );
    Graphics()->SetWorldMtx( *(D3DXMATRIX*)&mat );

    Graphics()->UpdateMatrices();

    /// Draw the square
    ID3D10Device *pDevice = Graphics()->GetDevice();

    DrawSquare();

    /// flip the buffer
    Graphics()->Present();
}

void cDetailApp::KeyUp( int key )
{
    if( key == DIK_D )
    {
        m_bDrawDetail = !m_bDrawDetail;
    }
}

void cDetailApp::KeyDown( int key )

```

```

{
}

void cDetailApp::InitSquare()
{
    point3 v[4];
    v[0] = point3(-10.f, 10.f, 0.f);
    v[1] = point3( 10.f, 10.f, 0.f);
    v[2] = point3( 10.f, -10.f, 0.f);
    v[3] = point3(-10.f, -10.f, 0.f);

    m_square[0].m_vPosition = *(D3DXVECTOR3*)&v[0];
    m_square[0].m_vColor = D3DCOLOR(1,1,1,1);
    m_square[0].m_vNormal = D3DXVECTOR3(0,0,0);
    m_square[0].m_TexCoords[0] = D3DXVECTOR2(0.0f, 0.0f);
    m_square[0].m_TexCoords[1] = D3DXVECTOR2(0.0f, 0.0f);

    m_square[1].m_vPosition = *(D3DXVECTOR3*)&v[1];
    m_square[1].m_vColor = D3DCOLOR(1,1,1,1);
    m_square[1].m_vNormal = D3DXVECTOR3(0,0,0);
    m_square[1].m_TexCoords[0] = D3DXVECTOR2(1.0f, 0.0f);
    m_square[1].m_TexCoords[1] = D3DXVECTOR2(1.0f, 0.0f);

    m_square[2].m_vPosition = *(D3DXVECTOR3*)&v[2];
    m_square[2].m_vColor = D3DCOLOR(1,1,1,1);
    m_square[2].m_vNormal = D3DXVECTOR3(0,0,0);
    m_square[2].m_TexCoords[0] = D3DXVECTOR2(1.0f, 1.0f);
    m_square[2].m_TexCoords[1] = D3DXVECTOR2(1.0f, 1.0f);

    m_square[3].m_vPosition = *(D3DXVECTOR3*)&v[0];
    m_square[3].m_vColor = D3DCOLOR(1,1,1,1);
    m_square[3].m_vNormal = D3DXVECTOR3(0,0,0);
    m_square[3].m_TexCoords[0] = D3DXVECTOR2(0.0f, 0.0f);
    m_square[3].m_TexCoords[1] = D3DXVECTOR2(0.0f, 0.0f);

    m_square[4].m_vPosition = *(D3DXVECTOR3*)&v[2];
    m_square[4].m_vColor = D3DCOLOR(1,1,1,1);
    m_square[4].m_vNormal = D3DXVECTOR3(0,0,0);
    m_square[4].m_TexCoords[0] = D3DXVECTOR2(1.0f, 1.0f);
    m_square[4].m_TexCoords[1] = D3DXVECTOR2(1.0f, 1.0f);

    m_square[5].m_vPosition = *(D3DXVECTOR3*)&v[3];
    m_square[5].m_vColor = D3DCOLOR(1,1,1,1);
    m_square[5].m_vNormal = D3DXVECTOR3(0,0,0);
    m_square[5].m_TexCoords[0] = D3DXVECTOR2(0.0f, 1.0f);
    m_square[5].m_TexCoords[1] = D3DXVECTOR2(0.0f, 1.0f);

}

void cDetailApp::DrawSquare()
{
    ID3D10Device *pDevice = Graphics()->GetDevice();

    static ID3D10Buffer *pVertexBuffer = NULL;

```

```

static bool bFirst = true;

if(bFirst)
{
    D3D10_BUFFER_DESC descBuffer;
    memset(&descBuffer, 0, sizeof(descBuffer));
    descBuffer.Usage = D3D10_USAGE_DEFAULT;
    descBuffer.ByteWidth = sizeof(cGraphicsLayer::cDefaultVertex) * 6;
    descBuffer.BindFlags = D3D10_BIND_VERTEX_BUFFER;
    descBuffer.CPUAccessFlags = 0;
    descBuffer.MiscFlags = 0;

    D3D10_SUBRESOURCE_DATA resData;
    memset(&resData, 0, sizeof(resData));
    resData.pSysMem = &m_square;
    Graphics()->GetDevice()->CreateBuffer(
        &descBuffer, &resData, &pVertexBuffer);

    bFirst = false;
}

if(pVertexBuffer)
{
    UINT uiStride = sizeof(cGraphicsLayer::cDefaultVertex);
    UINT uiOffset = 0;

    Graphics()->GetDevice()->IASetVertexBuffers(
        0, 1, &pVertexBuffer, &uiStride, &uiOffset);
    Graphics()->GetDevice()->IASetPrimitiveTopology(
        D3D10_PRIMITIVE_TOPOLOGY_TRIANGLELIST);

    D3D10_TECHNIQUE_DESC descTechnique;
    Graphics()->GetDefaultTechnique()->GetDesc(&descTechnique);
    for(UINT uiCurPass = 0; uiCurPass < descTechnique.Passes; uiCurPass++)
    {
        Graphics()->GetDefaultTechnique()->
            GetPassByIndex(uiCurPass)->Apply(0);
        pDevice->Draw(6, 0);
    }
}

void cDetailApp::LoadTextures()
{
    ID3D10Device *pDevice = Graphics()->GetDevice();

    /**
     * base map in texture stage 0
     */
    m_pTextures[0] = new cTexture( L"..\\bin\\media\\base.dds", 0 );
    m_pTextures[1] = new cTexture( L"..\\bin\\media\\detail3.dds", 0 );

    ID3D10ShaderResourceView *pViews[2];

    pViews[0] = m_pTextures[0]->GetShaderView();

```

```

pViews[1] = m_pTextures[1]->GetShaderView();

Graphics()->SetTextureArray(pViews, 2);
}

```

The sample renders a quad that moves toward and away from the viewer. You can turn the detail map on and off using the D key. Here is the shader that goes with it:

```

/*****
 *           Advanced 3D Game Programming with DirectX 10.0
 * * * * *
 *
 *   See license.txt for modification and distribution information
 *   copyright (c) 2007 by Peter Walsh, Wordware
 *****/

matrix g_mtxWorld;
matrix g_mtxView;
matrix g_mtxProj;

#define MAX_LIGHTS 10    // Ensure this is the same as the C++ value

int4 g_viLightStatus;
float4 g_vLightColors[MAX_LIGHTS];
float4 g_vLightDirections[MAX_LIGHTS];

int4 g_viTextureStatus;
Texture2D g_Textures[8];

//////////////////////////////////////
// Default Vertex Shader
struct VS_INPUT
{
    float3 vPosition    : POSITION;
    float3 vNormal      : NORMAL;
    float4 vColor       : COLOR0;
    float2 vTexCoords1  : TEXCOORD;
    float2 vTexCoords2  : TEXCOORD;
};

struct VS_OUTPUT
{
    float4 vPosition    : SV_POSITION;
    float3 vNormal      : NORMAL;
    float4 vColor       : COLOR0;
    float2 vTexCoords1  : TEXCOORD0;
    float2 vTexCoords2  : TEXCOORD1;
};

VS_OUTPUT DefaultVS(VS_INPUT dataIn)
{
    VS_OUTPUT result;

```

```

float4 vPos = float4(dataIn.vPosition, 1.0f);

float4 vFinalPos = mul(vPos, g_mtxWorld);
vFinalPos = mul(vFinalPos, g_mtxView);
vFinalPos = mul(vFinalPos, g_mtxProj);

result.vPosition = vFinalPos;
result.vNormal = dataIn.vNormal;
result.vColor = dataIn.vColor;
result.vTexCoords1 = dataIn.vTexCoords1;
result.vTexCoords2 = dataIn.vTexCoords2;

return result;
}

////////////////////////////////////////////////////////////////////////////////////////////////////////////////////////////////
// Default Pixel Shader

SamplerState SamplerStateWrap
{
    Filter = MIN_MAG_MIP_LINEAR;
    AddressU = Wrap;
    AddressV = Wrap;
};

float4 DefaultPS(VS_OUTPUT dataIn) : SV_Target
{
    float4 finalColor = dataIn.vColor;

    if(g_viLightStatus.x > 0)
    {
        for(int iCurLight = 0 ; iCurLight < g_viLightStatus.x ; iCurLight++)
        {
            finalColor += saturate(dot(g_vLightDirections[iCurLight],
                dataIn.vNormal) * g_vLightColors[iCurLight]);
        }
    }

    float4 texColor = finalColor;

    if(g_viTextureStatus.x > 0)
        texColor *= g_Textures[0].Sample(
            SamplerStateWrap, dataIn.vTexCoords1);
    if(g_viTextureStatus.x > 1)
        texColor *= (g_Textures[1].Sample(
            SamplerStateWrap, dataIn.vTexCoords2));

    return texColor;
}

////////////////////////////////////////////////////////////////////////////////////////////////////////////////////////////////
// Default Technique
technique10 DefaultTechnique
{
    pass Pass0

```

```

    {
        SetGeometryShader(NULL);
        SetVertexShader(CompileShader(vs_4_0, DefaultVS()));
        SetPixelShader(CompileShader(ps_4_0, DefaultPS()));
    }
}

```

Glow Maps

Glow maps are useful for creating objects that have parts that glow independently of the base map. Examples of this are things like LEDs on a tactical unit, buttons on a weapon or other unit, and the lights on a building or spaceship. The same scenery during the daytime could look completely different at night with the addition of a few glow maps.

To implement it you use a texture map that is mostly black, with lighter areas representing things that will glow on the final image. What you want is the glow map to have no effect on the base map except in glowing areas. Then you just sample the glow texture and add it to your final color instead of modulating it.

The one danger of using glow maps is you can easily saturate the image if you use a bad choice of texture. Saturation occurs when the result of the blending step is greater than 1.0. The value is clamped down to 1.0 of course, but this causes the image to look too bright. If the base map is too bright, consider using darker shades of color for glowing areas of the glow map.

Gloss Maps

Gloss maps are one of the cooler effects that can be done with multitexture. Any other effect you can do (like environment maps or specular maps) can look cooler if you also use gloss maps.

Gloss maps themselves don't do much; they are combined with another multitexture operation. The gloss map controls how much another effect shows through on a surface. For example, let's suppose you're designing a racing car game. The texture map for the car includes everything except the wheels (which are different objects, connected to the parent object). When you go to draw the car, you put an environment map on it, showing some sort of city scene or lighting that is rushing by.

One small issue that can crop up using this method is the fact that the entire surface of the car (windshield, hood, bumper, etc.) reflects the environment map the same amount. *San Francisco Rush* got around this by using a different map for the windshield, but there is another way to go about it: by using a gloss map on the car. The gloss map is brighter in areas that should reflect the environment map more, and darker in areas where it should reflect it less. So, in this example, the area that would cover the windshield would be fairly bright, almost white. The body of the car would be a lighter gray, and the non-reflective bumpers would be dark, almost

black. Figure 9.29 shows how you combine the base map, gloss map, and specular/environment map to make a gloss-mapped image.




Figure 9.29:
The separate
pieces of a gloss
map

You can do some amazing effects with this. For example, let's say you're driving through a mud puddle and mud splatters up on the car. You could use a special mud texture and blit some streaks of mud on top of the base car texture map around the wheels to show that it had just gone through mud. You could also blit the same mud effects to the gloss map, painting black texels instead of mud-colored texels. That way, whatever regions of the car had mud on them would not reflect the environment map, which is exactly the kind of effect wanted.

The way it works is you perform a first pass with the base map modulated by the diffuse color to get the plain-vanilla texture-mapped model we all know and love. Then you perform a second pass that has two textures (the environment map and the gloss map) modulated together. The modulation allows the light areas of the gloss map to cause the environment map to come out more than in dark areas. The result of the modulation is blended with the frame buffer destination color using an addition blend (source factor = 1, dest factor = 1).

Other Effects

There are innumerable other multitexture effects that are possible given the set of blending operations provided by Direct3D. To create a certain effect, all you need to do is dream it up; with HLSL and DirectX 10 there are no limits. To showcase that concept, I'm going to throw together the blending modes I've discussed thus far into a menagerie of multitexture!

Application: MultiTex

To show off some multipass and multitexture effects, I threw together an application that shows an object resembling the Earth with six total passes that can each be toggled on and off.




Figure 9.30: The base texture map



Figure 9.31:
A screenshot of
the base map
with diffuse light



Figure 9.32:
The detail map




Figure 9.33:
The base pass,
with diffuse light
and detail map

The magic of this application occurs in the shader, where all the different textures are combined together to create the final image. Here is the code, followed by the shader:

```

/*****
 *      Advanced 3D Game Programming with DirectX 10.0
 *      * * * * *
 *
 *      See license.txt for modification and distribution information
 *      copyright (c) 2007 by Peter Walsh, Wordware
 *****/

#include <algorithm>

#include "stdafx.h"
#include "GraphicsLayer.h"
#include "Application.h"
#include "GameErrors.h"
#include "DxHelper.h"
#include "Window.h"

using namespace std;

/**
 * This is the static member variable that points to the one
 * (and only one) graphics layer allowed in each application.
 */
cGraphicsLayer *cGraphicsLayer::m_pGlobalGLayer = NULL;

cGraphicsLayer::cGraphicsLayer(HWND hWnd)
{
    if(m_pGlobalGLayer)
        throw cGameError(
            L"GraphicsLayer object already instantiated\n");
}

```

```

    m_pGlobalGLayer = this;

    m_hWnd = hWnd;
    m_pDevice = NULL;
    m_pBackBuffer = NULL;
    m_pSwapChain = NULL;
    m_pRenderTargetView = NULL;
    m_pFont = NULL;
    m_pFontSprite = NULL;
    m_pMessageQueue = NULL;
    m_pDefaultEffect = NULL;
    m_pDefaultTechnique = NULL;
    m_pDefaultInputLayout = NULL;
    m_pmtxWorldVar = NULL;
    m_pmtxViewVar = NULL;
    m_pmtxProjVar = NULL;
    m_pLightDirVar = NULL;
    m_pLightColorVar = NULL;
    m_pDepthStencilBuffer = NULL;
    m_pDepthStencilState = NULL;
    m_pDepthStencilView = NULL;
    m_pNumLightsVar = NULL;

    m_pTexturesVar = NULL;

    // Default to identity
    D3DXMatrixIdentity(&m_mtxWorld);
    D3DXMatrixIdentity(&m_mtxView);
    D3DXMatrixIdentity(&m_mtxProj);

    m_iNumLights = 0;
}

void CGraphicsLayer::DestroyAll()
{
    if(m_pDevice)
        m_pDevice->ClearState();

    // Release in opposite order to creation usually
    SafeRelease(m_pDefaultInputLayout);
    SafeRelease(m_pDefaultEffect);
    SafeRelease(m_pMessageQueue);
    SafeRelease(m_pFont);
    SafeRelease(m_pFontSprite);
    SafeRelease(m_pDepthStencilView);
    SafeRelease(m_pDepthStencilState);
    SafeRelease(m_pRenderTargetView);
    SafeRelease(m_pBackBuffer);
    SafeRelease(m_pDepthStencilBuffer);
    SafeRelease(m_pSwapChain);
    SafeRelease(m_pDevice);

    /**
     * Prevent any further access to the graphics class
     */
}

```

```
m_pGlobalGLayer = NULL;
}

cGraphicsLayer::~cGraphicsLayer()
{
    DestroyAll();
}

void cGraphicsLayer::Present()
{
    HRESULT r = S_OK;

    assert(m_pDevice);

    r = m_pSwapChain->Present(0, 0);
    if(FAILED(r))
    {
        OutputDebugString(L"Present Failed!\n");
    }

    DumpMessages();
}

void cGraphicsLayer::DumpMessages()
{
    assert(m_pMessageQueue);

    HRESULT r = 0;

    int iCount = 0;
    while(1)
    {
        iCount++;
        if(iCount > 10)
            break;

        // Get the size of the message
        SIZE_T messageLength = 0;
        r = m_pMessageQueue->GetMessage(0, NULL, &messageLength);
        if(messageLength == 0)
            break;

        // Allocate space and get the message
        D3D10_MESSAGE * pMessage = (D3D10_MESSAGE*)malloc(messageLength);
        r = m_pMessageQueue->GetMessage(0, pMessage, &messageLength);
        if(FAILED(r))
        {
            OutputDebugString(L"Failed to get Direct3D Message");
            break;
        }

        TCHAR strOutput[2048];

        const char *pCompileErrors =
            static_cast<const char*>(pMessage->pDescription);
```

```

        TCHAR wcsErrors[2048];
        mbstowcs(wcsErrors, pCompileErrors, 2048);

        swprintf_s(strOutput, 2048, L"D3DMSG [Cat[%i] Sev[%i] ID[%i]: %s\n",
            pMessage->Category, pMessage->Severity, pMessage->ID, wcsErrors);
        OutputDebugString(strOutput);
    }
}

void CGraphicsLayer::Clear(const float (&colClear)[4])
{
    assert(m_pDevice);
    m_pDevice->ClearRenderTargetView(m_pRenderTargetView, colClear);
}

void CGraphicsLayer::ClearDepthStencil(const float fDepth, const UINT8 uiStencil)
{
    assert(m_pDevice);
    m_pDevice->ClearDepthStencilView(m_pDepthStencilView,
        D3D10_CLEAR_DEPTH | D3D10_CLEAR_STENCIL, fDepth, uiStencil);
}

void CGraphicsLayer::UpdateLights()
{
    int iLightData[4] = {m_iNumLights, 0, 0, 0};

    m_pNumLightsVar->SetIntVector(iLightData);

    for(int iCurLight = 0 ; iCurLight < m_iNumLights ; iCurLight++)
    {
        m_pLightDirVar->SetFloatVectorArray(
            (float*)m_aLights[iCurLight].m_vDirection, iCurLight, 1);
        m_pLightColorVar->SetFloatVectorArray(
            (float*)m_aLights[iCurLight].m_vColor, iCurLight, 1);
    }
}

void CGraphicsLayer::InitD3D(int width, int height)
{
    HRESULT r = 0;

    // Keep a copy of the screen dimensions
    m_rcScreenRect.left = m_rcScreenRect.top = 0;
    m_rcScreenRect.right = width;
    m_rcScreenRect.bottom = height;

    CreateDeviceAndSwapChain();
    CreateViewport();
    CreateDepthStencilBuffer();
    CreateDebugText();

    // Attach the render target view to the output merger state
    m_pDevice->OMSetRenderTargets(
        1, &m_pRenderTargetView, m_pDepthStencilView);
}

```

```

        CreateDefaultShader();
    }

void CGraphicsLayer::CreateDeviceAndSwapChain()
{
    HRESULT r = 0;

    // Structure to hold the creation parameters for the device
    DXGI_SWAP_CHAIN_DESC descSwap;
    ZeroMemory(&descSwap, sizeof(descSwap));

    // Only want one back buffer
    descSwap.BufferCount = 1;

    // Width and height of the back buffer
    descSwap.BufferDesc.Width = m_rcScreenRect.right;
    descSwap.BufferDesc.Height = m_rcScreenRect.bottom;

    // Standard 32-bit surface type
    descSwap.BufferDesc.Format = DXGI_FORMAT_R8G8B8A8_UNORM;

    // 60hz refresh rate
    descSwap.BufferDesc.RefreshRate.Numerator = 60;
    descSwap.BufferDesc.RefreshRate.Denominator = 1;
    descSwap.BufferUsage = DXGI_USAGE_RENDER_TARGET_OUTPUT;

    // Connect it to our main window
    descSwap.OutputWindow = m_hWnd;

    // No multisampling
    descSwap.SampleDesc.Count = 1;
    descSwap.SampleDesc.Quality = 0;

    // Windowed mode
    descSwap.Windowed = TRUE;

    // Create the device using hardware acceleration
    r = D3D10CreateDeviceAndSwapChain(
        NULL, // Default adapter
        D3D10_DRIVER_TYPE_HARDWARE, // Hardware accelerated device
        NULL, // Not using a software DLL for rendering
        D3D10_CREATE_DEVICE_DEBUG, // Flag to allow debug output
        D3D10_SDK_VERSION, // Indicates the SDK version being used
        &descSwap,
        &m_pSwapChain,
        &m_pDevice);

    if(FAILED(r))
    {
        throw cGameError(L"Could not create IDirect3DDevice10");
    }

    // Get a copy of the pointer to the back buffer
    r = m_pSwapChain->GetBuffer(
        0, __uuidof(ID3D10Texture2D), (LPVOID*)&m_pBackBuffer);
}

```

```

    if(FAILED(r))
    {
        throw cGameError(L"Could not get back buffer");
    }

    // Create a render target view
    r = m_pDevice->CreateRenderTargetView(
        m_pBackBuffer, NULL, &m_pRenderTargetView);
    if(FAILED(r))
    {
        throw cGameError(L"Could not create render target view");
    }

    r = m_pDevice->QueryInterface(
        __uuidof(ID3D10InfoQueue), (LPVOID*)&m_pMessageQueue);
    if(FAILED(r))
    {
        throw cGameError(L"Could not create IDirect3DDevice10 message queue");
    }
    m_pMessageQueue->SetMuteDebugOutput(false);    // No muting
    m_pMessageQueue->SetMessageCountLimit(-1);    // Unlimited messages
}

void cGraphicsLayer::CreateViewport()
{
    // Create a viewport the same size as the back buffer
    D3D10_VIEWPORT vp;
    ZeroMemory(&vp, sizeof(vp));
    vp.Width = m_rcScreenRect.right;
    vp.Height = m_rcScreenRect.bottom;
    vp.MinDepth = 0.0f;
    vp.MaxDepth = 1.0f;
    vp.TopLeftX = 0;
    vp.TopLeftY = 0;
    m_pDevice->RSSetViewports( 1, &vp );
}

void cGraphicsLayer::CreateDebugText()
{
    // Create the font for rendering text
    D3DX10CreateFont(m_pDevice,
        15, 0,
        FW_BOLD,
        1,
        FALSE,
        DEFAULT_CHARSET,
        OUT_DEFAULT_PRECIS,
        DEFAULT_QUALITY,
        DEFAULT_PITCH | FF_DONTCARE,
        L"Arial",
        &m_pFont);
    assert(m_pFont);

    // Create the sprite to use to render letters
    D3DX10CreateSprite(m_pDevice, m_uiMAX_CHARS_PER_FRAME, &m_pFontSprite);

```

```

}

void CGraphicsLayer::CreateDepthStencilBuffer()
{
    HRESULT r = 0;
    // Create the depth buffer
    D3D10_TEXTURE2D_DESC descDepth;
    ZeroMemory(&descDepth, sizeof(descDepth));
    descDepth.Width = m_rcScreenRect.right;
    descDepth.Height = m_rcScreenRect.bottom;
    descDepth.MipLevels = 1;
    descDepth.ArraySize = 1;
    descDepth.Format = DXGI_FORMAT_D24_UNORM_S8_UINT;
    descDepth.SampleDesc.Count = 1;
    descDepth.SampleDesc.Quality = 0;
    descDepth.Usage = D3D10_USAGE_DEFAULT;
    descDepth.BindFlags = D3D10_BIND_DEPTH_STENCIL;
    descDepth.CPUAccessFlags = 0;
    descDepth.MiscFlags = 0;
    r = m_pDevice->CreateTexture2D(&descDepth, NULL, &m_pDepthStencilBuffer);
    if(FAILED(r))
    {
        throw cGameError(L"Unable to create depth buffer");
    }

    D3D10_DEPTH_STENCIL_DESC descDS;
    ZeroMemory(&descDS, sizeof(descDS));
    descDS.DepthEnable = true;
    descDS.DepthWriteMask = D3D10_DEPTH_WRITE_MASK_ALL;
    descDS.DepthFunc = D3D10_COMPARISON_LESS;

    // Stencil test values
    descDS.StencilEnable = true;
    descDS.StencilReadMask = (UINT8)0xFFFFFFFF;
    descDS.StencilWriteMask = (UINT8)0xFFFFFFFF;

    // Stencil op if pixel is front
    descDS.FrontFace.StencilFailOp = D3D10_STENCIL_OP_KEEP;
    descDS.FrontFace.StencilDepthFailOp = D3D10_STENCIL_OP_INCR;
    descDS.FrontFace.StencilPassOp = D3D10_STENCIL_OP_KEEP;
    descDS.FrontFace.StencilFunc = D3D10_COMPARISON_ALWAYS;

    // Stencil op if pixel is back
    descDS.BackFace.StencilFailOp = D3D10_STENCIL_OP_KEEP;
    descDS.BackFace.StencilDepthFailOp = D3D10_STENCIL_OP_DECR;
    descDS.BackFace.StencilPassOp = D3D10_STENCIL_OP_KEEP;
    descDS.BackFace.StencilFunc = D3D10_COMPARISON_ALWAYS;

    r = m_pDevice->CreateDepthStencilState(&descDS, &m_pDepthStencilState);
    if(FAILED(r))
    {
        throw cGameError(L"Could not create depth/stencil state");
    }
    m_pDevice->OMSetDepthStencilState(m_pDepthStencilState, 1);
}

```

```

D3D10_DEPTH_STENCIL_VIEW_DESC descDSView;
ZeroMemory(&descDSView, sizeof(descDSView));
descDSView.Format = DXGI_FORMAT_D24_UNORM_S8_UINT;
descDSView.ViewDimension = D3D10_DSV_DIMENSION_TEXTURE2D;
descDSView.Texture2D.MipSlice = 0;

r = m_pDevice->CreateDepthStencilView(
    m_pDepthStencilBuffer, &descDSView, &m_pDepthStencilView);
if(FAILED(r))
{
    throw cGameError(L"Could not create depth/stencil view");
}
}

void cGraphicsLayer::CreateDefaultShader()
{
    HRESULT r = 0;

    DWORD shaderFlags = D3D10_SHADER_ENABLE_STRICTNESS;
    #if defined( DEBUG ) || defined( _DEBUG )
        // Turn on extra debug info when in debug config
        shaderFlags |= D3D10_SHADER_DEBUG;
    #endif

    ID3D10Blob *pErrors = 0;
    // Create the default rendering effect
    r = D3DX10CreateEffectFromFile(
        L"..\\GameLib\\DefaultShader.fx", NULL, NULL, "fx_4_0", shaderFlags, 0,
        m_pDevice, NULL, NULL, &m_pDefaultEffect, &pErrors, NULL);
    if(pErrors)
    {
        char *pCompileErrors = static_cast<char*>(
            pErrors->GetBufferPointer());
        TCHAR wcsErrors[MAX_PATH];
        mbstowcs(wcsErrors, pCompileErrors, MAX_PATH);
        OutputDebugString(wcsErrors);
    }

    if(FAILED(r))
    {
        throw cGameError(
            L"Could not create default shader - DefaultShader.fx");
    }

    m_pDefaultTechnique =
        m_pDefaultEffect->GetTechniqueByName("DefaultTechnique");
    if(!m_pDefaultTechnique)
    {
        throw cGameError(
            L"Could not find default technique in DefaultShader.fx");
    }

    // Set up the input layout
    D3D10_INPUT_ELEMENT_DESC defaultLayout[] =
    {

```



```

        {"POSITION", 0, DXGI_FORMAT_R32G32B32_FLOAT, 0, 0,
         D3D10_INPUT_PER_VERTEX_DATA, 0},
        {"NORMAL", 0, DXGI_FORMAT_R32G32B32_FLOAT, 0,
         D3D10_APPEND_ALIGNED_ELEMENT, D3D10_INPUT_PER_VERTEX_DATA, 0},
        {"COLOR", 0, DXGI_FORMAT_R32G32B32A32_FLOAT, 0,
         D3D10_APPEND_ALIGNED_ELEMENT, D3D10_INPUT_PER_VERTEX_DATA, 0},
        {"TEXCOORD", 0, DXGI_FORMAT_R32G32_FLOAT, 0, D3D10_APPEND_ALIGNED_ELEMENT,
         D3D10_INPUT_PER_VERTEX_DATA, 0},
        {"TEXCOORD", 1, DXGI_FORMAT_R32G32_FLOAT, 0, D3D10_APPEND_ALIGNED_ELEMENT,
         D3D10_INPUT_PER_VERTEX_DATA, 0},
    };

    UINT uiNumElements = sizeof(defaultLayout)/sizeof(defaultLayout[0]);
    D3D10_PASS_DESC descPass;
    m_pDefaultTechnique->GetPassByIndex(0)->GetDesc(&descPass);
    r = m_pDevice->CreateInputLayout(defaultLayout, uiNumElements,
        descPass.pIAInputSignature,
        descPass.pIAInputSignatureSize, &m_pDefaultInputLayout);
    if(FAILED(r))
    {
        throw cGameError(L"Could not create default layout");
    }
    m_pDevice->IASetInputLayout(m_pDefaultInputLayout);

    m_pmtxWorldVar = m_pDefaultEffect->
        GetVariableByName("g_mtxWorld")->AsMatrix();
    m_pmtxViewVar = m_pDefaultEffect->
        GetVariableByName("g_mtxView")->AsMatrix();
    m_pmtxProjVar = m_pDefaultEffect->
        GetVariableByName("g_mtxProj")->AsMatrix();

    m_pTexturesVar = m_pDefaultEffect->
        GetVariableByName("g_Textures")->AsShaderResource();

    D3DXMATRIX mtxWorld;
    D3DXMatrixIdentity(&mtxWorld);
    SetWorldMtx(mtxWorld);

    D3DXMATRIX mtxView;
    D3DXVECTOR3 vCamPos(0.0f, 1.0f, -12.0f);
    D3DXVECTOR3 vCamAt(0.0f, 1.0f, 0.0f);
    D3DXVECTOR3 vCamUp(0.0f, 1.0f, 0.0f);
    D3DXMatrixLookAtLH(&mtxView, &vCamPos, &vCamAt, &vCamUp);
    SetViewMtx(mtxView);

    D3DXMATRIX mtxProj;
    D3DXMatrixPerspectiveFovLH(&mtxProj, (float)D3DX_PI * 0.5f,
        m_rcScreenRect.right/(float)m_rcScreenRect.bottom, 0.1f, 100.0f);
    SetProjMtx(mtxProj);

    UpdateMatrices();

    m_pLightDirVar = m_pDefaultEffect->GetVariableByName(
        "g_vLightDirections")->AsVector();

```

```

    m_pLightColorVar = m_pDefaultEffect->GetVariableByName(
        "g_vLightColors" )->AsVector();
    m_pNumLightsVar = m_pDefaultEffect->GetVariableByName(
        "g_viLightStatus" )->AsVector();

    m_pTextureDataVar = m_pDefaultEffect->GetVariableByName(
        "g_viTextureStatus" )->AsVector();
}

void CGraphicsLayer::DrawTextString(int x, int y,
                                   D3DXCOLOR color, const TCHAR *strOutput)
{
    m_pFontSprite->Begin(0);
    RECT rect = {x, y, m_rcScreenRect.right, m_rcScreenRect.bottom};
    m_pFont->DrawText(m_pFontSprite, strOutput, -1, &rect, DT_LEFT, color);
    m_pFontSprite->End();
}

void CGraphicsLayer::Create(HWND hWnd, short width, short height)
{
    new CGraphicsLayer(hWnd); // construct the object

    // Init Direct3D and the device for fullscreen operation
    Graphics()->InitD3D(width, height);
}

void CGraphicsLayer::SetNumActiveTexture(int texCount, int iCurrent)
{
    int data[4] = {texCount, iCurrent, 0, 0};
    m_pTextureDataVar->SetIntVector(data);
}

void CGraphicsLayer::UpdateMatrices()
{
    m_pmtxWorldVar->SetMatrix((float*)&m_mtxWorld);
    m_pmtxViewVar->SetMatrix((float*)&m_mtxView);
    m_pmtxProjVar->SetMatrix((float*)&m_mtxProj);
}

void CGraphicsLayer::SetTexture(int iIdx, ID3D10ShaderResourceView
*pTexShaderView)
{
    if(m_pTexturesVar)
    {
        m_pTexturesVar->SetResource(pTexShaderView);
    }
}

```




Figure 9.34: The glow map



Figure 9.35:
Another detail map




Figure 9.36:
A screenshot of
the base map,
diffuse light, and
glow map

```

/*****
 *           Advanced 3D Game Programming with DirectX 10.0
 * ****
 *
 *   See license.txt for modification and distribution information
 *   copyright (c) 2007 by Peter Walsh, Wordware
 *****/

matrix g_mtxWorld;
matrix g_mtxView;
matrix g_mtxProj;

#define MAX_LIGHTS 10    // Ensure this is the same as the C++ value

int4 g_vLightStatus;
float4 g_vLightColors[MAX_LIGHTS];
float4 g_vLightDirections[MAX_LIGHTS];

int4 g_viTextureStatus;
Texture2D g_Textures[8];

////////////////////////////////////
// Default Vertex Shader
struct VS_INPUT
{
    float3 vPosition    : POSITION;
    float3 vNormal      : NORMAL;
    float4 vColor       : COLOR0;
    float2 vTexCoords1  : TEXCOORD;
    float2 vTexCoords2  : TEXCOORD;
};

struct VS_OUTPUT
{
    float4 vPosition    : SV_POSITION;
    float3 vNormal      : NORMAL;
    float4 vColor       : COLOR0;
    float2 vTexCoords1  : TEXCOORD0;
    float2 vTexCoords2  : TEXCOORD1;
};

VS_OUTPUT DefaultVS(VS_INPUT dataIn)
{
    VS_OUTPUT result;

    float4 vPos = float4(dataIn.vPosition, 1.0f);

    float4 vFinalPos = mul(vPos, g_mtxWorld);
    vFinalPos = mul(vFinalPos, g_mtxView);
    vFinalPos = mul(vFinalPos, g_mtxProj);

    result.vPosition = vFinalPos;
    result.vNormal = dataIn.vNormal;
    result.vColor = dataIn.vColor;
    result.vTexCoords1 = dataIn.vTexCoords1;

```

```

        result.vTexCoords2 = dataIn.vTexCoords2;

        return result;
    }

    //////////////////////////////////////
    // Default Pixel Shader

    SamplerState SamplerStateWrap
    {
        Filter = MIN_MAG_MIP_LINEAR;
        AddressU = Wrap;
        AddressV = Wrap;
    };

    float4 DefaultPS(VS_OUTPUT dataIn) : SV_Target
    {
        float4 finalColor = dataIn.vColor;

        if(g_viLightStatus.x > 0)
        {
            for(int iCurLight = 0 ; iCurLight < g_viLightStatus.x ; iCurLight++)
            {
                finalColor += saturate(dot(g_vLightDirections[iCurLight],
                    dataIn.vNormal) * g_vLightColors[iCurLight]);
            }
        }

        float4 texColor = finalColor;

        if(g_viTextureStatus.x > 0)
            texColor *= g_Textures[0].Sample(
                SamplerStateWrap, dataIn.vTexCoords1);
        if(g_viTextureStatus.x > 1)
            texColor *= (g_Textures[1].Sample(
                SamplerStateWrap, dataIn.vTexCoords2));

        return texColor;
    }

    //////////////////////////////////////
    // Default Technique
    technique10 DefaultTechnique
    {
        pass Pass0
        {
            SetGeometryShader(NULL);
            SetVertexShader(CompileShader(vs_4_0, DefaultVS()));
            SetPixelShader(CompileShader(ps_4_0, DefaultPS()));
        }
    }

    VS_OUTPUT MT_VS(VS_INPUT dataIn)
    {

```

```

VS_OUTPUT result;

float4 vPos = float4(dataIn.vPosition, 1.0f);

float4 vWorldPos = mul(vPos, g_mtxWorld);
float4 vViewPos = mul(vWorldPos, g_mtxView);
float4 vScreenPos = mul(vViewPos, g_mtxProj);

result.vPosition = vScreenPos;
result.vColor = dataIn.vColor;
result.vTexCoords1 = dataIn.vTexCoords1;

float3 vViewNormal = normalize(
    mul(mul(dataIn.vNormal, g_mtxWorld), g_mtxView));

result.vNormal = vViewNormal;

float3 vReflection = normalize(reflect(vViewPos, vViewNormal));

float fTemp = 2.0f * sqrt(
    vReflection.x * vReflection.x +
    vReflection.y * vReflection.y +
    (vReflection.z+1.0f) * (vReflection.z+1.0f)
);

result.vTexCoords2.x = vReflection.x / fTemp + 0.5f;
result.vTexCoords2.y = vReflection.y / fTemp + 0.5f;

return result;
}

float4 MT_PS(VS_OUTPUT dataIn) : SV_Target
{
    float4 diffuseColor = 0;

    float3 vLightDir = float3(-1,0.0,-0.2);
    float4 vSunColor = float4(1,1,0.3,0);

    diffuseColor = saturate(dot(vLightDir, dataIn.vNormal) * vSunColor);

    float4 baseColor = g_Textures[0].Sample(
        SamplerStateWrap, dataIn.vTexCoords1);
    float4 detailColor = g_Textures[2].Sample(
        SamplerStateWrap, dataIn.vTexCoords1);
    float4 envColor = g_Textures[5].Sample(
        SamplerStateWrap, dataIn.vTexCoords2);
    float4 glowColor = g_Textures[1].Sample(
        SamplerStateWrap, dataIn.vTexCoords1);
    float4 maskColor = g_Textures[6].Sample(
        SamplerStateWrap, dataIn.vTexCoords1);
    float4 cloudColor = g_Textures[7].Sample(
        SamplerStateWrap, dataIn.vTexCoords1);

```

```

    baseColor = (baseColor * detailColor + diffuseColor);
    envColor = envColor * (1-maskColor) * 0.3f;
    glowColor = pow(glowColor * (1-diffuseColor), 4);

    float4 finalColor = baseColor + envColor + glowColor;

    return finalColor;
}

technique10 MultiTexTechnique
{
    pass PassBackground
    {
        SetGeometryShader(NULL);
        SetVertexShader(CompileShader(vs_4_0, MT_VS()));
        SetPixelShader(CompileShader(ps_4_0, MT_PS()));
    }
}

```




Figure 9.37: The base pass, with diffuse light and the environment map




Figure 9.38: The gloss map




Figure 9.39: The base pass, diffuse lighting, and gloss pass




Figure 9.40: The clouds texture




Figure 9.41: The base pass with diffuse lighting and the clouds




Figure 9.42: All passes turned on

Using the Stencil Buffer

I said previously that I would talk about stencil buffers more, and here we are. They are now so ubiquitous across hardware it's well worth understanding how they work. Stencil buffers allow you to easily perform a lot of nifty effects that otherwise would be extremely difficult, if not impossible.

The stencil buffer is yet another buffer for your application (you already have the frame buffer or render target, back buffer, and the z-buffer). It's never its own buffer; rather, it always piggybacks a few bits of the z-buffer. Generally, when stenciling is desired, you set up a 16-bit z-buffer (15 bits of depth, one bit of stencil) or a 32-bit z-buffer (24 bits of depth, 8 bits of stencil). You can clear it to a default value using `ID3D10Device::ClearDepthStencilView()` just like you did with the back buffer and the z-buffer.

Here's the way it works: Before a pixel is tested against the z-buffer, it's tested against the stencil buffer. The stencil for a pixel is defined by the stencil reference value (which is set with one of the depth/stencil states). They are compared using a comparison function just like the z-buffer. It should be noted that before the reference value and the stencil buffer value are compared, they are both ANDed by the stencil mask. The equation for the stencil test step is:

```
(StencilRef & StencilMask) CompFunc (StencilBufferValue & StencilMask)
```

What happens as a result of the comparison is defined by a bunch of other depth/stencil states. There are three possible cases that can occur, and you can modify what happens to the stencil buffer for each of them. The three cases are:

- The stencil test fails.
- The stencil test succeeds, but then the z test fails.
- Both the stencil test and the z test succeed.

What happens when you get to one of these three cases is defined by setting states to a member of the `D3D10_STENCIL_OP` enumeration, the values of which are listed in Table 9.4.

Table 9.4: Values for the `D3D10_STENCIL_OP` enumeration

<code>D3D10_STENCIL_OP_KEEP</code>	Do not change the value in the stencil buffer.
<code>D3D10_STENCIL_OP_ZERO</code>	Set the entry in the stencil buffer to 0.
<code>D3D10_STENCIL_OP_REPLACE</code>	Set the entry in the stencil buffer to the reference value.
<code>D3D10_STENCIL_OP_INCRSAT</code>	Increment the stencil buffer entry, clamping it to the maximum value—8-bit stencil buffers have a maximum value of 255; 1-bit stencil buffers have a maximum value of 1.
<code>D3D10_STENCIL_OP_DECRSAT</code>	Decrement the stencil buffer entry, clamping it to 0.
<code>D3D10_STENCIL_OP_INVERT</code>	Invert the bits in the stencil buffer entry.
<code>D3D10_STENCIL_OP_INCR</code>	Increment the stencil buffer entry, wrapping to 0 if it goes past the maximum value.
<code>D3D10_STENCIL_OP_DECR</code>	Decrement the stencil buffer entry, wrapping to the maximum value if it goes past 0.

Table 9.5 lists the `D3DCMPFUNC` enumeration, which holds the possible comparison functions we can set `D3D10_STENCIL_OP` to.

Table 9.5: Values for the `D3DCMPFUNC` enumeration

<code>D3DCMP_NEVER</code>	Always fails the test.
<code>D3DCMP_LESS</code>	Passes if the tested pixel is less than the current pixel.
<code>D3DCMP_EQUAL</code>	Passes if the tested pixel is equal to the current pixel.
<code>D3DCMP_LESSEQUAL</code>	Passes if the tested pixel is less than or equal to the current pixel.
<code>D3DCMP_GREATER</code>	Passes if the tested pixel is greater than the current pixel.
<code>D3DCMP_NOTEQUAL</code>	Passes if the tested pixel is not equal to the current pixel.
<code>D3DCMP_GREATEREQUAL</code>	Passes if the tested pixel is greater than or equal to the current pixel.
<code>D3DCMP_ALWAYS</code>	Always passes the test.

Overdraw Counter

One simple use of stencil buffers is to implement an overdraw counter. The bounding factor in most graphics applications (especially games like *Quake III: Arena* or *Unreal Tournament*) is the fill rate of the card. Very few games implement anywhere near exact visibility for their scenes, so many pixels on the screen will be drawn two, three, five, or more times. If you draw every pixel five times and you're running at a high resolution, the application will be completely bound by how fast the card can draw the pixels to the frame buffer.

You can use stencil buffers to help you figure out how much overdraw you're doing when rendering a given scene. Initially you clear the stencil buffer to zero. The stencil comparison function is set to always accept pixels. Then the stencil operations for both PASS and ZFAIL are set to increment the value in the stencil buffer.

Then, after you render your frame, you lock the z-buffer and average together all of the stencil values in the buffer. Of course, this is going to be a really slow operation; calculating the overdraw level is something to be done during development only!

Dissolves and Wipes

Another use for stencil buffers is to block out certain regions of an image. You can use this to do film-like transitions between scenes, such as wiping left to right from one image to another.

To implement it, you have a polygon that grows frame to frame, eventually enveloping the entire screen. You initially clear out the stencil buffer to zero. The wipe polygon doesn't draw anything to the frame buffer, but it sets the stencil pixels it covers to 1. The old scene should be rendered on all the stencil pixels marked with a 0 (the ones that weren't covered by the wipe polygon) and the new scene should be rendered on all the pixels with a stencil value of 1. When the wipe polygon grows to the point that it's completely covering the frame buffer, you can stop rendering the first scene (since it's completely invisible now).

Conclusion

Now you should have an understanding of some of the more advanced texturing effects you can achieve in Direct3D 10. We looked at a number of Direct3D 10 techniques including alpha blending, texture mapping, pixel shaders, and environment mapping. In the next chapter we'll look at more advanced topics like scene management.

This page intentionally left blank.

Scene Management

Sometimes I wish I was in an industry that moved a little slower. Imagine the car industry—if car manufacturers had to move at the speed that computers have to move at, cars would be traveling at supersonic speeds, flying, and driving themselves. Luckily for them this isn't the case. Large strides in most industries happen over years, not days.

The computer industry is an entirely different matter. Users are always clamoring for more—bigger bad guys, more complex physics and AI, higher-resolution textures, and so forth. If a game doesn't provide what the users want, it won't be what users buy.



Aside: In fact, it's quickly becoming the case that the programmers aren't the ones running the fastest to keep up. For many artists, creating a texture twice as detailed takes more than twice as long to make. The same thing goes for models and worlds. Content creation, within the next few years, will become the bottleneck for games.

The Scene Management Problem and Solutions

An extremely large problem that every game has to deal with is managing its world on a per-frame basis. The problem as a whole is called *scene management*. It has many different facets: managing the per-frame polygon count, keeping the $O(n^2)$ physics and AI algorithms in check, and managing the network throughput, among others.

As an example, suppose you're writing a first-person style game, where the world is a large research facility with a scientist you must find. The research facility might be tremendously large, with a hundred rooms or more, and hallways connecting them. Each room has dozens of objects, most with hundreds of triangles. All in all, the entire world could have upward of two or three million triangles in it.

There are a number of issues that you need to deal with to handle this world. For starters, how do you draw it? Early naïve systems would have no structure to the world at all, just a list of two million triangles. There is no choice but to draw the entire two million triangle list. This is pretty ridiculous, as almost all of the polygons drawn won't end up contributing pixels to the final image. In most cases, we'll be standing in one room. If the door is closed, all of the visible polygons belong to the room you're in, and you'll end up drawing less than 1% of the total polygon count. The only way you'll be able to draw the world at interactive frame rates is to somehow chop away polygons that you know won't be visible. Drawing

1% versus drawing 100% of the polygons can mean the difference between 30 frames per second and three seconds per frame!

There's an even worse example: collision detection. Whenever an object (such as our fearless player) moves in the world, you must make sure it hasn't hit any other objects. The brute-force algorithm involves us taking each object and checking it for a collision. True, for most of the several thousand objects in the scene, you will quickly reject them with a trivial bounding box or bounding sphere test. However, you still need to perform all multi-thousand tests, however quick and trivial they may be, for each object that moves! The time complexity of this algorithm, $O(n^2)$, will completely devour any processor power you have, and the frame rate will slow to a crawl.

The idea of doing this is completely ridiculous. If an object is situated in a room, you should test against the dozens of objects in the room, not the thousands of objects in the whole world! As you reduce n (the number of objects each object must test against for collisions), the physics code speed increases quadratically!

The same issues that plague collision detection also attack the networking code. As characters move around the scene, their movements must be broadcast to all the other clients so they can keep an accurate picture of the world. However, each client couldn't care less about where things that they can't see are moving. Rather than knowing where each of the thousands of scene objects are, it just wants to know about the ones that are relevant to it.

There needs to be a system of scene management. You have to be smart about which parts of the scene you choose to render, which sets of objects to perform tests against, and so forth, to keep the running time of the application in check. Ideally, the size of the total scene shouldn't matter; as long as there is RAM to hold it all, you can increase the size of the world without bounds while keeping all of the ugly algorithms running at about the same speed.

In the following sections, I'll go over a few different systems that can be used to manage different types of scenes. I'll end up using one of them (portal rendering) to write a game at the end of this chapter, so I'll obviously be going in-depth on that one the most. The other ones we'll hopefully cover enough that you'll be able to implement them on your own.

Quadtrees/Octrees

Quadtrees are a classic form of spatial partitioning, and are used in scene management algorithms for many systems where the spatial data can be usefully boiled down to two dimensions. The most prevalent example of this is games that take place over terrain, like *Myth* or *Tribes*. While objects have a particular height value, the space of possible height values is much less than either of the lateral dimensions. You can take advantage of this and use the two lateral dimensions to classify the relations of the objects.

The scene is initially bounded on all sides by a bounding square. The recursive algorithm proceeds as follows: If there is more than one object in the bounding square, the square is divided into four smaller squares, or subnodes. The subnodes are associated with their parent node in the form of a tree with either four or zero children at each node (hence the name quadtree). When an object is bridging a boundary between two nodes, both nodes contain a handle to the object. This continues until all the nodes have either one or zero objects in them or some maximum depth level is reached. Figure 10.1 shows how you would turn a scene into a quadtree.




Figure 10.1:
Sample quadtree

The advantage that quadtrees give lies in hierarchical culling. To illustrate, consider the case of drawing the world. Imagine that the total scene is very large, covering many square miles—much more than the viewer could ever see. When you go to draw the scene, you test the four subnodes of the quadtree root against the visible region from the current viewpoint. If one of the subnodes does not sit in the visible region, you know that none of its children do either. You can trivially cull that subnode and everything below it. If any part of the subnode is in the visible region, you recurse on the four children of that subnode. When we reach a leaf, you draw all of the objects inside it. This gives you a pleasingly quick way to cull out large portions of the database during rendering.

As objects move around, they may exit the current node they are in. At this point the quadtree should be recomputed. This way the tree always has the least amount of nodes necessary to represent the data. Alternatively, the tree can be constructed independently of the objects, blindly subdivided to a certain depth. Then as objects move, the code finds out the set of leaves in which the bounding sphere for the object sits.

Interobject collision detection can be done really quickly using a tree such as this. The only objects that some particular object could possibly

intersect with must also be in one of the same leaves that said particular object is in. To do an object intersection test, you get the list of leaves an object is sitting in and then get the list of objects that are sitting in each of those leaves. That set of objects is the space of possible collision candidates. The space of candidates you end up with will be considerably smaller than the total number of objects. Also, if you keep the size of the leaves constant you can increase the size of the total scene almost without bound. Each of these algorithms will take about the same time to run as long as the relative proximity of objects remains the same relative to the area of the leaves.

Octrees are very similar to quadtrees, except you deal with cubes instead of squares and divide the space along all three axes, making eight subregions. Each node has eight subnodes instead of the four that we see in quadtrees. Octrees perform very well in games like 3D space sims such as *Homeworld*, where there are many objects in the scene, but very few of them are likely candidates for collisions.

Portal Rendering

Portal rendering is a really effective way to handle scene management for indoor environments. It's the method I'll use in the game I write at the end of this chapter. It's also used in many games and public domain engines (like *Crystal Space*). Besides being effective for what it tries to do, portal rendering is intuitive and easy to implement.

Imagine you are standing in a room. This room turns out to be a bit on the sparse side; in fact there is absolutely nothing in it. The only thing in the room for us to look at is a doorway leading into another equally empty and uninteresting room. The doorway is jumping on the minimalist bandwagon, so it happens to not have a door in it. There is an unobstructed view into the next room. Also, for the sake of discussion, assume that the walls between the rooms are made of a new space-age construction material that is infinitely thin.

Now draw the room you're standing in. You have some sort of data structure that represents the room, with a list of the polygons that define the walls, floors, and ceiling. Eventually this data structure will be called a *cell*. You also have a special invisible polygon that covers the doorway. This special polygon is a *portal*.

After you have drawn the room you are in, the entire frame buffer will be filled in except for the doorway. You know for a fact that the other room, the one you can see through the doorway, is entirely behind the room you are standing in. The only pixels left to even consider when drawing the next room are the ones seen through the doorway.

Instead of blindly drawing the next room and letting the z-buffer take care of it, you can constrain the renderer, telling it to only draw the pixels that haven't been touched yet. I'll discuss the constraining part in a moment, but for right now I intuitively know that the pixels that haven't been drawn yet are the pixels in the doorway. That way, you don't waste

time drawing triangles that would be obstructed by the z-buffer and not drawn anyway; you only draw pixels that will end up making it onto the screen. The bound for many applications is the fill rate of the card: how fast it can draw pixels on the screen. The fewer pixels you have it draw, the faster the application has the potential to be. An extra advantage comes from the fact that the fewer triangles you process, the less strain we put on the transformation and lighting pipeline.

This algorithm is recursive. For example, say that the second room has another doorway on the opposite wall, looking into a third room. When you finish drawing the second room, the only pixels left will be the ones that lay inside the next doorway. Once again constrain the drawing to only take place inside that doorway and then draw the next room.

Portal Rendering Concepts

A scene that uses portal rendering must be put together in a certain way. In essence, you compose your scene out of a set of rooms that are all connected by doorways. The rooms, which I'll call cells, can have any number of polygons, as long as the cells themselves remain convex. Each cell represents empty space, and all of the polygons in the cell face inward. They can be connected to any number of other cells, and the boundary locations become the portals. You can think of two adjacent cells (call them A and B) having two portals that connect them. One belongs to A and points to B as the neighbor, and the other belongs to B and points to A as the neighbor. The two portals must be exactly the same except for vertex ordering and normal orientation. If this is not the case, portal rendering will not work correctly. To help illustrate the point, consider a standard scene of polygons, as shown in Figure 10.2.




Figure 10.2:
A regular scene of
polygons

This scene is of a few angular rooms, seen from a top-down view. The white area shows the region in which you can move around. As you have it set up now, there is no spatial relationship set up for this scene. You can draw the world using the z-buffer and not have to worry about anything, but you'll have to suffer through all the scene management problems detailed above.

Instead, how about turning the scene into something with which you can portal render? I'll discuss later how you can take an arbitrary polygonal scene and decompose it into a bunch of convex cells, but for right now assume that I have a black box that can do it for you. It might come up with a composition like the one that appears in Figure 10.3.




Figure 10.3:
The scene divided
into eight cells

Now you have divided the scene into eight distinct convex rooms, all connected together with portals. Even without the rendering efficiency of having zero overdraw, this data representation is useful in many ways. Since the cells are convex, you can quickly perform a test between the bounding sphere of an object and the cells in the world to find out which cell(s) an object is touching (it may be situated in a portal such that it is sitting in more than one cell). All you do is perform a plane-sphere test with each polygon and portal of the cell. If the sphere is completely behind any of the planes (remember that the normals all point into the cell), then you know that the sphere isn't touching the cell at all.

If you know the space of cells in which an object exists, then suddenly the scene becomes much more manageable. When you want to do any processing on an object that needs to be tested against other objects (for example, checking for collisions), you don't need to check all the objects in the scene; you just need to check the objects that are in each of the cells that the target object is in. Even if the world has a hundred thousand cells and a hundred thousand objects wandering around those cells, the hard algorithms will only need to be run with an extremely small subset of that

total set; there might only be ten or so other objects in the cell(s) a target object is currently in. You can imagine how much faster this makes things.

The extra bonus that you get with cell-based worlds lies in portal rendering, which allows you to efficiently find the exact set of cells that are visible from that viewpoint. Even better, you can find the exact set of polygons visible from that viewpoint.

To generate this visibility data, I use what I'll call a *viewing cone*. A viewing cone is an n -sided pyramid that extends infinitely forward, with all the sides meeting together in world space at the location of the viewer. The sides bound the visible region of space that can be seen from the camera. Before you do any portal rendering, the sides of the cone represent the sides of the screen; anything outside the cone will be outside the screen and shouldn't be drawn. Note that when you're rendering, the viewing cone has two extra conceptual polygons that lop off the tip of the cone (everything in front of the near z -plane is not drawn) and the bottom of the cone (everything behind the far z -plane is not drawn). You must ignore these two planes for right now; portal rendering won't work correctly if you don't.



Aside: A pyramid with the top and bottom chopped off is called a *frustum*.

Given the viewing cone, clipping a polygon against it is fairly easy. You perform a polygon-plane clipping operation with the polygon and each of the planes of the cone. If at any time we completely cull the polygon, we know that it is invisible and we can stop processing it.

To use this functionality, we'll make a class called `cViewCone`. It can be constructed from a viewer location and a polygon (which can be extracted from a portal in our scene) or from a viewer location and the projection information (width, height, field of view). It clips polygons with `Clip()`, returning true if some part of the polygon was inside the viewing cone (and false otherwise). Also, an output polygon is filled with the inside fragment, if one was there.

```

/*****
 *      Advanced 3D Game Programming with DirectX 10.0
 *      * * * * *
 *
 *      See license.txt for modification and distribution information
 *      copyright (c) 2007 by Peter Walsh, Wordware
 *****/

#ifndef _FRUSTUM_H
#define _FRUSTUM_H

#include "..\math3d\point3.h"
#include "..\math3d\plane3.h"

#define MAX_PLANES 32

```

```

class cViewCone
{
    plane3    m_planes[MAX_PLANES];
    int       m_nPlanes;
    point3    m_camLoc;

    /**
     * We need this functionality twice, encapsulate it
     */
    void GenFromPoly( const point3& camLoc, const polygon< point3 >& in );

public:

    /**
     * Default constructor/destructor
     */
    cViewCone();

    /**
     * Construct a frustum from an input polygon. The polygon
     * is assumed to wind clockwise from the point of view of the
     * camera
     */
    cViewCone( const point3& camLoc, const polygon< point3 >& in );

    /**
     * Construct a frustum from the viewport data. Uses the
     * data to construct a cameraspace polygon,
     * back-transforms it to worldspace, then constructs a
     * frustum out of it.
     */
    cViewCone( float fov, int width, int height, matrix4& viewMat );

    /**
     * Clip a polygon to the frustum.
     * true if there was anything left
     */
    bool Clip(
        const polygon<point3>& in,
        polygon<point3>* out );

    /**
     * Get the center point of a frustum.
     * This is needed when we create frustums
     * from other frustums
     */
    const point3& GetLoc()
    {
        return m_camLoc;
    }
};

#endif // _FRUSTUM_H

```

```

/*****
 *      Advanced 3D Game Programming with DirectX 10.0
 *      * * * * *
 *
 *      See license.txt for modification and distribution information
 *      copyright (c) 2007 by Peter Walsh, Wordware
 *****/

#include "stdafx.h"
#include <assert.h>

#include <algorithm> // for swap()

#include "ViewCone.h"

using namespace std;

cViewCone::cViewCone()
: m_nPlanes( 0 )
, m_camLoc( point3::Zero )
{
    // Do nothing
}

cViewCone::cViewCone( const point3& camLoc, const polygon< point3 >& in )
{
    assert( in.nElem );
    assert( in.pList );
    GenFromPoly( camLoc, in );
}

cViewCone::cViewCone( float fov, int width, int height, matrix4& viewMat )
{
    /**
     * This function is kind of a magic trick, as it tries to
     * invert the projection matrix. If you stare at the way
     * we make projection matrices for long enough this should
     * make sense.
     */
    float aspect = ((float)height) / width;

    float z = 10;

    float w = aspect * (float)( cos(fov/2)/sin(fov/2) );
    float h = 1.0f * (float)( cos(fov/2)/sin(fov/2) );

    float x0 = -z/w;
    float x1 = z/w;
    float y0 = z/h;
    float y1 = -z/h;

    /**
     * Construct a clockwise camera-space polygon
     */
    polygon<point3> poly(4);

```

```

    poly.nElem = 4;
    poly.pList[0] = point3( x0, y0,z); // top-left
    poly.pList[1] = point3( x1, y0,z); // top-right
    poly.pList[2] = point3( x1, y1,z); // bottom-right
    poly.pList[3] = point3( x0, y1,z); // bottom-left

    /**
     * Create a camspace->worldspace transform
     */
    matrix4 camMatInv = matrix4::Inverse( viewMat );

    /**
     * Convert it to worldspace
     */
    poly.pList[0] = poly.pList[0] * camMatInv;
    poly.pList[1] = poly.pList[1] * camMatInv;
    poly.pList[2] = poly.pList[2] * camMatInv;
    poly.pList[3] = poly.pList[3] * camMatInv;

    /**
     * Generate the frustum
     */
    GenFromPoly( camMatInv.GetLoc(), poly );
}

void cViewCone::GenFromPoly( const point3& camLoc, const polygon< point3 >& in )
{
    int i;
    m_camLoc = camLoc;
    m_nPlanes = 0;
    for( i=0; i< in.nElem; i++ )
    {
        /**
         * Plane 'i' contains the camera location and the 'ith'
         * edge around the polygon
         */
        m_planes[ m_nPlanes++ ] = plane3(
            camLoc,
            in.pList[(i+1)%in.nElem],
            in.pList[i] );
    }
}

bool cViewCone::Clip( const polygon<point3>& in, polygon<point3>*& out )
{
    /**
     * Temporary polygons. This isn't thread safe
     */
    static polygon<point3> a(32), b(32);
    polygon<point3>* pSrc = &a;
    polygon<point3>* pDest = &b;

    int i;

```

```

/**
 * Copy the input polygon to a.
 */
a.nElem = in.nElem;
for( i=0; i<a.nElem; i++ )
{
    a.pList[i] = in.pList[i];
}

/**
 * Iteratively clip the polygons
 */
for( i=0; i<m_nPlanes; i++ )
{
    if( !m_planes[i].Clip( *pSrc, pDest ) )
    {
        /**
         * Failure
         */
        return false;
    }
    std::swap( pSrc, pDest );
}

/**
 * If we make it here, we have a polygon that survived.
 * Copy it to out.
 */
out->nElem = pSrc->nElem;
for( i=0; i<pSrc->nElem; i++ )
{
    out->pList[i] = pSrc->pList[i];
}

/**
 * Success
 */
return true;
}

```

You can perform portal rendering in one of two ways, depending on the fill rate of the hardware you're running on and the speed of the host processor. The two methods are *exact portal rendering* and *approximative portal rendering*.

Exact Portal Rendering

To render a portal scene using exact portal rendering, you use a simple recursive algorithm. Each cell has a list of polygons, a list of portals, and a visited bit. Each portal has a pointer to the cell adjacent to it. You start the algorithm knowing where the camera is situated, where it's pointing, and in which cell it is sitting. From this, along with other information like the height, width, and field of view of the camera, you can determine the

initial viewing cone that represents the entire viewable area on the screen. Also, you clear the valid bit for all the cells in the scene.

You draw all of the visible regions of the cell's polygons (the visible regions are found by clipping the polygons against the current viewing cone). Also, you set the visited bit to true. Then you walk the list of portals for the cell. If the cell on the other side hasn't been visited, you try to clip the portal against the viewing cone. If a valid portal fragment results from the operation, you have the area of the portal that was visible from the current viewing cone. Take the resulting portal fragment and use it to generate a new viewing cone. Finally, you recurse into the cell adjacent to the portal in question using the new viewing cone. You repeat this process until there are no new cells to traverse into.

Check out the following pseudocode as an example.

```
void DrawSceneExact
    for( all cells )
        cell.visited = false
    currCell = cell camera is in
    currCone = viewing cone of camera
    currCell.visited = true
    VisitCell( currCell, currCone )

void VisitCell( cell, viewCone )
    for( each polygon in cell )
        polygon fragment = viewCone.clip( current polygon )
        if( polygon fragment is valid )
            draw( polygon fragment )
    for( each portal )
        portal fragment = viewCone.clip( current portal )
        if( portal fragment is valid )
            if( !portal.otherCell.visited )
                portal.otherCell.visited = true
                newCone = viewing cone of portal fragment
                VisitCell( portal.otherCell, newCone )
```

I haven't talked about how to handle rendering objects (such as enemies, players, ammo boxes, and so forth) that would be sitting in these cells. It's almost impossible to guarantee zero overdraw if you have to draw objects that are in cells. Luckily, there is the z-buffer, so you don't need to worry; you just draw the objects for a particular cell when you recurse into it. Handling objects without a depth buffer can get hairy pretty quickly; be happy you have it.

Approximative Portal Rendering

As the fill rate of cards keeps increasing, it's becoming less and less troublesome to just throw up your hands and draw some triangles that won't be seen. The situation is definitely much better than it was a few years ago, when software rasterizers were so slow that you wouldn't even think of wasting time drawing pixels you would never see. Also, since the

triangle rate is increasing so rapidly, it's quickly getting to the point where it takes longer to clip off invisible regions of a triangle than it would to just draw the triangle and let the hardware sort any problems out.

In approximative portal rendering, you only spend time clipping portals. Objects in the cells and the triangles making up the cell boundaries are either trivially rejected or drawn. When you want to draw an object, you test the bounding sphere against the frustum. If the sphere is completely outside the frustum, you know that it's completely obscured by the cells you've already drawn, so you don't draw the object. If any part of it is visible, you just draw the entire object, no questions asked. While you do spend time drawing invisible triangles (since part of the object may be obscured), you make up for it since you can draw the object without any special processing using one big `DrawIndexedPrimitive()` or something similar. The same is true for portal polygons. You can try to trivially reject polygons in the cell and save some rendering time, or just blindly draw all of them when you enter the cell.

Another plus when you go with an approximative portal rendering scheme is that the cells don't need to be strictly convex; they can have any number of concavities in them and still render correctly if a z-buffer is used. Remember, however, that things like containment tests become untrivial when you go with concave cells; you can generally use something like a BSP tree for each cell to get around these problems.

Portal Effects

Assuming that all of the portals and cells are in a fixed location in 3D, there isn't anything terribly interesting that you do with portal rendering. However, that's a restriction you don't necessarily need to put on yourself. There are a few nifty effects that can be done almost for free with a portal rendering engine, two of which I'll cover here: *mirrors* and *teleporters*.

Mirrors

Portals can be used to create mirrors that reflect the scene back to you. Using them is much easier when you're using exact portal rendering (clipping all drawn polygons to the boundaries of the viewing cone for the cell the polygons are in); when they're used with approximative portal rendering, a little more work needs to be done.

Mirrors can be implemented with a special portal that contains a transformation matrix and a pointer back to the parent cell. When this portal is reached, the viewing cone is transformed by the portal's transformation matrix. You then continue the recursive portal algorithm, drawing the cell we're in again with the new transformation matrix that will make it seem as if we are looking in a mirror.



Warning: You should be careful when using multiple mirrors in a scene. If two mirrors can see each other, it is possible to infinitely recurse between both portals until the stack overflows. This can be avoided by keeping track of how many times you have recursed into a mirror portal and stopping after some number of iterations.

To implement mirrors you need the answers to two questions: How do you create the mirror transformation matrix? and How do you transform the viewing cone by that matrix? I'll answer each of these questions separately.

Before you can try to make the mirror transformation matrix, you need an intuitive understanding of what the transformation should do. When you transform the viewing cone by the matrix, you will essentially be flipping it over the mirror such that it is sitting in world space exactly opposite where it was before. Figure 10.4 shows what is happening.




Figure 10.4:
2D example of view
cone reflection

For comprehension's sake, let's give the mirror its own local coordinate space. To define it, you need the n , o , a , and p vectors to put the matrix together (see Chapter 4). The p vector is any point on the mirror; you can just use the first vertex of the portal polygon. The a vector is the normal of the portal polygon (so in the local coordinate space, the mirror is situated at the origin in the x - y plane). The n vector is found by crossing a with any vector that isn't parallel to it (let's just use the up direction, $\langle 0,1,0 \rangle$) and normalizing the result. Given n and a , o is just the normalized cross product of the two. Altogether this becomes:

$$\begin{aligned} a &= \text{mirror}_{\text{normal}} \\ n &= \overline{a \times \langle 0,1,0 \rangle} \\ o &= \overline{a \times n} \\ T_{\text{mirror}} &= \begin{bmatrix} -n & -o & 0 \\ -o & -n & 0 \\ -a & -a & 0 \\ -p & -p & 1 \end{bmatrix} \end{aligned}$$



Warning: The cross product is undefined when the two vectors are parallel, so if the mirror is on the floor or ceiling you should use a different vector rather than $\langle 0, 1, 0 \rangle$. $\langle 1, 0, 0 \rangle$ will suffice.

However, a transformation matrix that converts points local to the mirror to world space isn't terribly useful by itself. To actually make the mirror transformation matrix you need to do a bit more work. The final transformation needs to perform the following steps:

- Transform world space vertices to the mirror's local coordinate space. This can be accomplished by multiplying the vertices by T_{mirror}^{-1} .
- Flip the local space vertices over the x-y plane. This can be accomplished by using a scaling transformation that scales by 1 in the x and y directions and -1 in the z direction (see Chapter 4). We'll call this transformation T_{reflect} .
- Finally, transform the reflected local space vertices back to world space. This can be accomplished by multiplying the vertices by T_{mirror} .

Given these three steps you can compose the final transformation matrix, M_{mirror}

$$M_{\text{mirror}} = T_{\text{mirror}}^{-1} T_{\text{reflect}} T_{\text{mirror}}$$

Given M_{mirror} , how do you apply the transformation to the viewing cone, which is just a single point and a set of planes? I haven't discussed how to apply transformations to planes yet, but now seems like a great time. There is a real way to do it, given the plane defined as a 1x4 matrix:

$$\mathbf{n} = [a \quad b \quad c \quad d]$$

$$\mathbf{n}' = \mathbf{n}(\mathbf{M}^{-1})^T$$

If you don't like that, there's a slightly more intuitive way that requires you to do a tiny bit more work. The problem with transforming normals by a transformation matrix is that you don't want them to be translated, just rotated. If you translated them, they wouldn't be normal-length anymore and wouldn't correctly represent a normal for anything. If you just zero-out the translation component of M_{mirror} (M_{14} , M_{24} , and M_{34}) and multiply it by the normal component of the plane, it will be correctly transformed. Alternatively, you can just do a 1x4 times 4x4 operation, making the first vector $[a, b, c, 0]$.



Warning: This trick only works for rigid-body transforms (ones composed solely of rotations, translations, and reflections).

So you create two transformation matrices, one for transforming regular vectors and one for transforming normals. You multiply the view cone location by the vector transformation matrix and multiply each of the

normals in the view cone planes by the normal transformation matrix. Finally, recompute the d components for each of the planes by taking the negative dot product of the transformed normal and the transformed view cone location (since the location is sitting on each of the planes in the view cone).

You should postpone rendering through a mirror portal until you have finished with all of the regular portals. When you go to draw a mirror portal, you clone the viewing cone and transform it by M_{mirror} . Then you reset all of the visited bits and continue the algorithm in the cell that owned the portal. This is done for all of the mirrors visited. Each time you find one, you add it to a mirror queue of mirror portals left to process.

You must be careful if you are using approximative portal rendering and you try to use mirrors. If you draw cells behind the portal, the polygons will interfere with each other because of z-buffer issues. Technically, what you see in a mirror is a flat image, and the mirror should always occlude things it is in front of. The way you are rendering a mirror (as a regular portal walk) it has depth, and faraway things in the mirror may not occlude near things that should technically be behind it. To fix this, before you render through the mirror portal, change the z-buffer comparison function to `D3DCMP_ALWAYS` and draw a screen space polygon over the portal polygon with the depth set to the maximum depth value. This essentially resets the z-buffer of the portal region so that everything drawn through the mirror portal will occlude anything drawn behind it. I recommend you use exact portal rendering if you want to do mirrors or translocators, which I'll discuss next.

Translocators and Non-Euclidean Movement

One of the coolest effects you can do with portal rendering is create non-Euclidean spaces to explore. One effect is having a doorway floating in the middle of a room that leads to a different area; you can see the different area through the door as you move around it. Another effect is having a small structure with a door, and upon entering the structure you realize there is much more space inside of it than could be possible given the dimensions of the structure from the outside. Imagine a small cube with a small door that opens into a giant amphitheater. Neither of these effects is possible in the real world, making them all the neater to have in a game.

You perform this trick in a way similar to the way you did mirrors, with a special transformation matrix you apply to the viewing cone when you descend through the portal. Instead of a mirror portal that points back to the cell it belongs to, a translocator portal points to a cell that can be anywhere in the scene. There are two portals that are the same size (but not necessarily the same orientation): a source portal and a destination portal. When you look through the source portal, the view is seen as if you were looking through the destination portal. Figure 10.5 may help explain this.




Figure 10.5:
2D representation
of the translocator
transformation

To create the transformation matrix to transform the view cone so that it appears to be looking through the destination portal, you compute local coordinate space matrices for both portals using the same n , o , a , and p vectors we used in the mirrors section. This gives you two matrices: T_{source} and T_{dest} . Then to compute $M_{translocator}$ you do the following steps:

- Transform the vectors from world space to the local coordinate space of the source matrix (multiply them by T_{source}^{-1}).
- Take the local space vectors and transform them back into world space, but use the destination transformation matrix (T_{dest}).

Given these steps you can compose the final transformation matrix:

$$M_{translocator} = T_{source}^{-1} T_{destination}$$

The rendering process for translocators is identical to rendering mirrors and has the same caveats when approximative portal rendering is used.

Portal Generation

Portal generation, or finding the set of convex cells and interconnecting portals given an arbitrary set of polygons, is a fairly difficult problem. The algorithm I'm going to describe is too complex to fully explain here; it would take much more space than can be allotted. However, it should lead you in the right general direction if you wish to implement it. David Black originally introduced me to this algorithm.

The first step is to create a leafy BSP of the data set. Leafy BSPs are built differently than node BSPs (the kind discussed in Chapter 4). Instead of storing polygons and planes at the nodes, only planes are stored. Leaves contain lists of polygons. During construction, you take the array of polygons and attempt to find a plane from the set of polygon planes that

divides the set into two non-zero sets. Coplanar polygons are put into the side that they face, so if the normal to the polygon is the same as the plane normal, it is considered in front of the plane. Trying to find a splitter will fail if and only if the set of polygons forms a convex cell. If this happens, the set of polygons becomes a leaf; otherwise, the plane is used to divide the set into two pieces, and the algorithm recurses on both pieces. An example of tree construction on a simple 12-polygon 2D data set appears in Figure 10.6.




Figure 10.6:
Constructing the
leaves of a BSP tree

The leaves of the tree will become the cells of the data set, and the nodes will become the portals. To find the portal polygon given the plane at a node, you first build a polygon that lies in the plane but extends out in all directions past the boundaries of the data set.

This isn't hideously difficult. You keep track of a *universe box*, a cube that is big enough to enclose the entire data set. You look at the plane normal to find the polygon in the universe box that is the most parallel to it. Each of the four vertices of that universe box polygon are projected into the plane. You then drop that polygon through the tree, clipping it against the cells that it sits in. After some careful clipping work (you need to clip against other polygons in the same plane, polygons in adjacent cells, etc.), you get a polygon that isn't obscured by any of the geometry polygons. This becomes a *portal polygon*.

After you do this for each of the splitting planes, you have a set of cells and a set of portal polygons but no association between them. Generating the associations between cells and portals is fairly involved, unfortunately. The sides of a cell may be defined by planes far away, so it's difficult to match up a portal polygon with a cell that it is abutting. Making the problem worse is the fact that some portal polygons may be too big, spanning several adjacent cells. In this case you would need to split up the cell.

On top of all that, once you get through this mess and are left with the set of cells and portals, you'll almost definitely have way too many cells and way too many portals. Combining cells isn't easy. You could just merge cells only if the new cell they formed was convex, but this will also give you a less-than-ideal solution: You may need to merge together three or more cells to get a nice big convex cell, but you wouldn't be able to reach that cell if you couldn't find pairs of cells out of the set that formed convex cells.

Because of problems like this, many engines just leave the process of portal cell generation up to the artists. If you're using approximative portal rendering, the artists can place portals fairly judiciously and end up with concave cells, leaving them just in things like doorways between rooms and whatnot. *Quake II* used something like this to help with culling scenes behind closed doors; area portals would be covering doors, and scenes behind them would only be traversed if the doors weren't closed.

Precalculated Portal Rendering (with PVS)

Up to this point I have discussed the usage of portal rendering to find the set of visible cells from a certain point in space. This way you can dynamically find the exact set of visible cells from a certain viewpoint. However, you shouldn't forget one of the fundamental optimization concepts in computer programming: Why generate something dynamically if you can precalculate it?

How do you precalculate the set of visible cells from a given viewpoint? The scene has a nearly infinite number of possible viewpoints, and calculating the set of visible cells for each of them would be a nightmare. If you want to be able to precalculate anything, you need to cut down the space of entries or cut down the number of positions for which you need to precalculate.

What if you just considered each cell as a whole? If you found the set of all the cells that were visible from *any* point in the cell, you could just save that. Each of the n cells would have a bit vector with n entries. If bit i in the bit vector is true, then cell i is visible from the current cell.

This technique of precalculating the set of visible cells for each cell was pioneered by Seth Teller in his 1992 thesis. The data associated with each cell is called the Potentially Visible Set, or PVS for short. It has since been used in *Quake*, *Quake II*, and just about every other first-person shooter under the sun.

Doing this, of course, forces you to give up exact visibility. The set of visible cells from all points inside a cell will almost definitely be more than the set of visible cells from one particular point inside the cell, so you may end up drawing some cells that are totally obscured from the camera. However, what you lose in fill rate, you gain in processing time. You don't need to do any expensive frustum generation or cell traversal; you simply step through the bit vector of the particular cell and draw all the cells whose bits are set.

Advantages/Disadvantages

The big reason this system is a win is because it offloads work from the processor to the hardware. True, you'll end up drawing more polygons than you have to, but it won't be that much more. The extra cost in triangle processing and fill rate is more than made up for since you don't need to do any frustum generation or polygon clipping.

However, using this system forces you to give up some freedom. The time it takes to compute the PVS is fairly substantial, due to the complexity of the algorithm. This prevents you from having your cells move around; they must remain static. This, however, is forgivable in most cases; the geometry that defines walls and floors shouldn't be moving around anyway.

Implementation Details

I can't possibly hope to cover the material required to implement PVS rendering; Seth Teller spends 150 pages doing it in his thesis. However, I can give a sweeping overview of the pieces of code involved.

The first step is to generate a cell and portal data set, using something like the algorithm discussed earlier. It's especially important to keep your cell count down, since you have an n^2 memory cost to hold the PVS data (where n is the number of cells). Because of this, most systems use the concept of *detail polygons* when computing the cells. Detail polygons are things like torches or computer terminals—things that don't really define the structural boundaries of a scene but just introduce concavities. Those polygons generally are not considered until the PVS table is calculated. Then they are just added to the cells they belong to. This causes the cells to be concave, but the visibility information will still remain the same, so we're all good.

Once you have the set of portals and cells, you iteratively step through each cell and find the set of visible cells from it. To do this, you do something similar to the frustum generation we did earlier in the chapter, but instead of a viewing cone coming out of a point, you generate a solid that represents what is viewable from all points inside the solid. An algorithm to do this (called *portal stabbing*) is given in Seth Teller's thesis. Also, the source code to QV (the application that performs this operation for the *Quake* engine) is available online.

When finished, and you have the PVS vector for each of the cells, rendering is easy. You can easily find out which cell the viewer is in (since each of the cells is convex). Given that cell, you step through the bit vector for that cell. If bit i is set, you draw cell i and let the z-buffer sort it out.

Application: Mobots Attack!

The intent of *Mobots Attack!* was to make an extremely simple client/server game that would provide a starting point for your own 3D game project. As such, it is severely lacking in some areas but fairly functional in others. There is only one level and it was crafted entirely by hand. Physics support is extremely lacking, as is the user interface. However, it has a fairly robust networking model that allows players to connect to a server, wander about, and shoot rockets at each other.

The objective of the game wasn't to make something glitzy. It doesn't use radiosity, AI opponents, multitexture, or any of the multiresolution modeling techniques we discussed previously. However, adding any of these things wouldn't be terribly difficult. Hopefully, adding cool features to an existing project will prove more fruitful for you than trying to write the entire project yourself. Making a project that was easy to add to was the goal of this game. I'll quickly cover some of the concepts that make this project work.

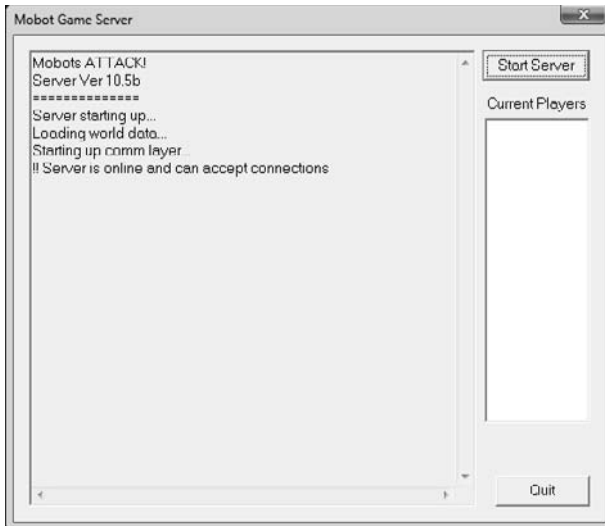


Figure 10.7:
The Mobots server

Interobject Communication

One of the biggest problems in getting a project of this size to work in any sort of reasonable way is interobject communication. For example, when an object hits a wall, some amount of communication needs to go on between the object and the wall so that the object stops moving. When a rocket hits an object, the rocket needs to inform the object that it must lose some of its hit points. When a piece of code wants to print debugging info, it needs to tell the application object to handle it.

Things get even worse. When the client moves, it needs some way to tell the server that its object has moved. But how would it do that? It's not

like it can just dereference a pointer and change the position manually; the server could be on a completely different continent.

To take care of this, a messaging system for objects to communicate with each other was implemented. Every object that wants to communicate needs to implement an interface called `iGameObject`.

```
typedef uint msgRet;

interface iGameObject
{
public:
    virtual objID GetID() = 0;
    virtual void SetID( objID id ) = 0;

    virtual msgRet ProcMsg( const sMsg& msg ) = 0;
};
```

An `objID` is an `int` masquerading as two shorts. The high short defines the class of object that the ID corresponds to, and the low short is the individual instance of that object. Each object in the game has a different `objID`, and that ID is the same across all the machines playing a game (the server and each of the clients). The following code manages the `objID`:

```
typedef uint objID;

inline objID MakeID( ushort segment, ushort offset )
{
    return (((uint)segment)<<16) | ((uint)offset);
}

inline ushort GetIDSegment( objID id )
{
    return (ushort)(id>>16);
}

inline ushort GetIDOffset( objID id )
{
    return (ushort)(id & 0xFFFF);
}

/**
 * These segments define the types of objects
 */
const ushort c_sysSegment = 0;    // System object
const ushort c_cellSegment = 1;   // Cell object
const ushort c_playerSegment = 2; // Player object
const ushort c_spawnSegment = 3;  // Spawning object
const ushort c_projSegment = 4;   // Projectile object
const ushort c_paraSegment = 5;   // Parametric object
const ushort c_tempSegment = 6;   // Temp object
```

All object communication is done by passing messages around. In the same way you would send a message to a window to have it change its screen position in Windows, you send a message to an object to have it perform a

certain task. The message structure holds onto the destination object (an objID), the type of the message (which is a member of the eMsgType enumeration), and then some extra data that has a different meaning for each of the messages. The following shows the sMsg structure.

```

struct sMsg
{
    eMsgType    m_type;
    objID       m_dest;
    union
    {
        struct
        {
            point3 m_pt;
        };
        struct
        {
            plane3 m_plane;
        };
        struct
        {
            color3 m_col;
        };
        struct
        {
            int m_i[4];
        };
        struct
        {
            float m_f[4];
        };
        struct
        {
            void *m_pData;
        };
    };
};

sMsg( eMsgType type = msgForceDword, objID dest = 0 )
: m_type( type )
, m_dest( dest )
{
}

sMsg( eMsgType type, objID dest, float f )
: m_type( type )
, m_dest( dest )
{
    m_f[0] = f;
}

sMsg( eMsgType type, objID dest, int i )
: m_type( type )
, m_dest( dest )
{
    m_i[0] = i;
}

```

```

    }

    sMsg( eMsgType type, objID dest, const point3& pt )
    : m_type( type )
    , m_dest( dest )
    , m_pt(pt)
    {
    }

    sMsg( eMsgType type, objID dest, const plane3& plane )
    : m_type( type )
    , m_dest( dest )
    , m_plane(plane)
    {
    }

    sMsg( eMsgType type, objID dest, void *pData )
    : m_type( type )
    , m_dest( dest )
    , m_pData( pData )
    {
    }
};

```

When an object is created, it registers itself with a singleton object called the message daemon (cMsgDaemon). The registering process simply adds an entry into a map that associates a particular ID with a pointer to an object. Typically what happens is when an object is created, a message will be broadcast to the other connected machines telling them to make the object as well and providing it with the ID to use in the object creation.

```

class cMsgDaemon
{
    map< objID, iGameObject* > m_objectMap;
    static cMsgDaemon *m_pGlobalMsgDaemon;

public:
    cMsgDaemon();
    ~cMsgDaemon();

    static cMsgDaemon *GetMsgDaemon()
    {
        // Accessor to the singleton
        if( !m_pGlobalMsgDaemon )
        {
            m_pGlobalMsgDaemon = new cMsgDaemon;
        }
        return m_pGlobalMsgDaemon;
    }

    void RegObject( objID id, iGameObject *pObj );
    void UnRegObject( objID id );
};

```

```

iGameObject *Get( int id )
{
    return m_objectMap[id];
}

/**
 * Deliver this message to the destination
 * marked in msg.m_dest. Throws an exception
 * if no such object exists.
 */
uint DeliverMessage( const sMsg& msg );
};

```

When one object wants to send a message to another object, it just needs to fill out an sMsg structure and then call cMsgDaemon::DeliverMessage() (or for a nicer-looking wrapper, use the SendMessage() function). In some areas of code, rather than ferry a slew of messages back and forth, a local-scope pointer to an object corresponding to an ID can be acquired with cMsgDaemon::Get() and then member functions can be called.

Network Communication

The networking model for this game is remarkably simple. There is no client-side prediction and no extrapolation. While this makes for choppy game play, hopefully it should make it easier to understand. The messaging model implemented here was strongly based on an article written by Mason McCuskey for GameDev.net called “Why pluggable factories rock my multiplayer world.”

Here’s the essential problem pluggable factories try to solve: Messages arrive to you as datagrams, essentially just buffers full of bits. Those bits represent a message that was sent to you from another client. The first byte (or short, if there are a whole lot of messages) is an ID tag that describes what the message is (a tag of 0x07, for example, may be the tag for a message describing the new position of an object that moved). Using the ID tag, you can figure out what the rest of the data is.

How do you go about figuring out what the rest of the data is? One way would be to just have a massive switch statement with a case label for each message tag that will take the rest of the data and construct a useful message. While that would work, it isn’t the right thing to do, OOP-wise. Higher-level code (that is, the code that constructs the network messages) needs to know details about lower-level code (that is, each of the message IDs and to what each of them correspond).




Figure 10.8:
Portal rendering in
the client

Pluggable factories allow you to get around this. Each message has a class that describes it. Every message derives from a common base class called `cNetMessage`.

```
/**
 * Generic Message
 * Every message class derives from this one.
 */
class cNetMessage
{
public:
    cNetMessage()
    {
    }
    ~cNetMessage()
    {
    }

    /**
     * Write out a bitstream to be sent over the wire
     * that encapsulates the data of the message.
     */
    virtual int SerializeTo( uchar *pOutput )
    {
        return 0;
    }

    /**
     * Take a bitstream as input (coming in over
     * the wire) and convert it into a message
     */
    virtual void SerializeFrom( uchar *pFromData, int datasize )
    {
    }
}
```

```

/**
 * This is called on a newly constructed message.
 * The message in essence executes itself. This
 * works because of the objID system; the message
 * object can communicate its desired changes to
 * the other objects in the system.
 */
virtual void Exec() = 0;

netID GetFrom()
{
    return m_from;
}
netID GetTo()
{
    return m_to;
}

void SetFrom( netID id )
{
    m_from = id;
}

void SetTo( netID id )
{
    m_to = id;
}

protected:

    netID m_from;
    netID m_to;
};

```

Every derived NetMessage class has a sister class that is the maker for that particular class type. For example, the login request message class cNM_LoginRequest has a sister maker class called cNM_LoginRequest-Maker. The maker class's responsibility is to create instances of its class type. The maker registers itself with a map in the maker parent class. The map associates those first-byte IDs with a pointer to a maker object. When a message comes off the wire, a piece of code looks up the ID in the map, gets the maker pointer, and tells the maker to create a message object. The maker creates a new instance of its sister net message class, calls `SerializeFrom()` on it with the incoming data, and returns the instance of the class.

Once a message is created, its `Exec()` method is called. This is where the message does any work it needs to do. For example, when the cNM_LoginRequest is executed (this happens on the server when a client attempts to connect), the message tells the server (using the interobject messaging system discussed previously) to create the player with the given name that was supplied. This will in turn create new messages, like an acknowledgment message notifying the client that it has logged in.




Figure 10.9:
Watch out, mobots!

Code Structure

There are six projects in the game workspace. Three of them you've seen before: `math3D`, `netLib`, and `gameLib`. The other three are `gameServer`, `gameClient`, and `gameCommon`. I made `gameCommon` just to ease the compile times; it has all the code that is common to both the client and the server.

The server is a Win32 dialog app. It doesn't link any of the DirectX headers in, so it should be able to run on any machine with a network card. All of the render code is pretty much divorced from everything else and put into the client library. The `gameClient` derives from `cApplication` just like every other sample app in the book.

The downloadable files contain documentation to help you get the game up and running on your machine; the client can connect to the local host, so a server and a client can both run on the same machine.

Closing Thoughts

I've covered a lot of ground in this book. Hopefully, it has all been lucid and the steps taken haven't been too big. If you've made it to this point, you should have enough knowledge to be able to implement a fairly complex game.

More importantly, you hopefully have acquired enough knowledge about 3D graphics and game programming that learning new things will come easily. Once you make it over the big hump, you start to see all the fundamental concepts that interconnect just about all of the driving concepts and algorithms.

Good luck with all of your endeavors!

Peter Walsh

This page intentionally left blank.

Appendix

An STL Primer

The world has two kinds of people in it: people who love the STL and use it every day, and people who have never learned the STL. If you're one of the latter, this appendix will hopefully help you get started.

The Standard Template Library is a set of classes and functions that help coders use basic containers (like linked lists and dynamic arrays) and basic algorithms (like sorting). It was officially introduced into the C++ library by the ANSI/ISO C++ Standards Committee in July 1994. Almost all C++ compilers (including Visual C++ 2005 Express Edition) implement the STL fairly well.

Almost all of the classes in the STL are template classes. This makes them usable with any type of object or class, and they are also compiled entirely as inline, making them extremely fast.

Templates

A quick explanation of templates: They allow you to define a generic piece of code that will work with any type of data, be it ints, floats, or classes.

The canonical example is Swap. Normally, if you want to swap integers in one place and swap floats in another, you write something like this:

```
void SwapInt( int &a, int &b )
{
    int temp = a;
    a = b;
    b = temp;
}

void SwapFloat( float &a, float &b )
{
    float temp = a;
    a = b;
    b = temp;
}
```

This is tolerable as long as you're only swapping around these two types, but what if you start swapping other things? You would end up with 10 or 15 different Swap functions in some file. The worst part is they're all exactly the same, except for the three tokens that declare the type. Let's make Swap() a template function.

```
template < class swapType >
void Swap( swapType &a, swapType &b )
{
    swapType temp = a;
```

```
    a = b;
    b = temp;
}
```

Here’s how it works: You use the templated `Swap()` function like you would any other. When the compiler encounters a place where you use the function, it checks the types that you’re using and makes sure they’re valid (both the same, since you use `T` for both `a` and `b`). Then it makes a custom piece of code specifically for the two types you’re using and compiles it inline. A way to think of it is the compiler does a find-replace, switching all instances of `swapType` (or whatever you name your template types; most people use `T`) to the types of the two variables you pass into `Swap()`. Because of this, the only penalty for using templates is during compilation; using them at run time is just as fast as using custom functions. There’s also a small penalty since using everything inline can increase your code size. However, for a large part this point is moot—most STL functions are short enough that the code actually ends up being smaller. Inlining the code for small functions takes less space than saving/restoring the stack frame.

Of course, even writing your own templated `Swap()` function is kind of dumb, as the STL library has its own function (`swap()`), but it serves as a good example. Templated classes are syntactically a little different, but we’ll get to those in a moment.

Containers

STL implements a set of basic containers to simplify most programming tasks; I used them everywhere in the text. While there are several more, Table A.1 lists the most popular ones.

Table A.1: The basic container classes

vector	Dynamic array class. You append entries on the end (using <code>push_back()</code>) and then can access them using standard array notation (via an overloaded <code>[]</code> operator). When the array needs more space, it internally allocates a bigger block of memory, copies the data over (explicitly, not bitwise), and releases the old one. Inserting data anywhere but the back is slow, as all the other entries need to be moved back one slot in memory.
deque	DeQueue class. Essentially a dynamic array of dynamic arrays. The data doesn’t sit in memory linearly, but you can get array-style lookups really quickly, and can append to the front or the back quickly.
list	Doubly linked list class. Inserting and removing anywhere is cheap, but you can’t randomly access things; you can only iterate forward or backward.
slist	Singly linked list class. Inserting to the front is quick, while inserting to the back is extremely slow. You shouldn’t need to use this since <code>list</code> is sufficiently fast for most code that would be using a linked list anyway.

map	This is used in a few places in the code; it is an associative container that lets you look up entries given a key. An example would be telephone numbers. You would make a map like so: map<string, int> numMap; and be able to say things like: numMap["joe"] = 5553298;
stack	A simple stack class.
queue	A simple queue class.
string	A vector of characters, with a lot of useful string operations implemented.

Let's look at some sample code. The following listing creates a vector template class of integers, adds some elements, and then asserts both.

```
#include <list>
#include <vector>
#include <string>

using namespace std;

void main()
{
    // Create a vector and add some numbers
    vector<int> intVec;
    intVec.push_back(5);
    intVec.push_back(10);
    assert( intVec[0] == 5 );
    assert( intVec.size() == 2 );
}
```

Notice two things: The headers for STL aren't post-fixed by .h, and the code uses the *using* keyword, which you may not have seen before. Namespaces essentially are blocks of functions, classes, and variables that sit in their own namespace (in this case, the namespace *std*). That way all of STL doesn't cloud the global namespace with all of its types (you may want to define your own class called *string*, for example). Putting the *using* keyword at the top of a .cpp file declares that we want the entire *std* namespace to be introduced into the global namespace so we can just say *vector<int>*. If we don't do that, we would need to specify the namespace we were referring to, so we would put *std::vector<int>*.

Iterators

Accessing individual elements of a vector is pretty straightforward; it's, in essence, an array (just a dynamic one) so we can use the same bracket-style syntax we use to access regular arrays. What about lists? Random access in a list is extremely inefficient, so it would be bad to allow the bracket operator to be used to access random elements. Accessing elements in other containers, like maps, makes even less intuitive sense. To

remedy these problems, STL uses an iterator interface to access elements in all the containers the same way.

Iterators are classes that each container defines that represent elements in the container. Iterators have two important methods: dereference and increment. Dereference accesses the element to which the iterator is currently pointing. Incrementing an iterator just moves it such that it points to the next element in the container.

For vectors of a type *T*, the iterator is just an alias to a pointer to a *T*. Incrementing a pointer will move to the next element in the array, and dereferencing will access the data. Linked lists use an actual class, where increment does something like (`currNode = currNode->pNext`) and dereference does something like (`return currNode->data`).

In order to make it work swimmingly, containers define two important iterators: `begin` and `end`. The `begin` iterator points to the first element in the container (`vec[0]` for vectors, `head.pNext` for lists). The `end` iterator points to one-past-the-last element in the container—the first non-valid element (`vec[size]` for vectors, `tail` for lists). In other words, when our iterator is pointing to `end`, we're done.

```
#include <list>
#include <vector>
#include <string>

using namespace std;

class CFoo
{
    ...
public:
    void DoSomething();
}

void main()
{
    vector< CFoo > fooVec;

    // Fill fooVec with some stuff
    ...

    // Create an iterator
    vector<CFoo>::iterator iter;

    // Iterate over all the elements in fooVec.
    for( iter = fooVec.begin();
        iter != fooVec.end();
        iter++ )
    {
        (*iter).DoSomething();
    }
}
```

You should realize that the picture is much richer than this. There are actually several different types of iterators (forward only iterators, random iterators, bidirectional iterators). I'm just trying to provide enough information to get your feet wet.

Why are iterators so cool? They provide a standard way to access the elements in a container. This is used extensively by the STL generic algorithms. As a first example, consider the generic algorithm `for_each`. It accepts three inputs: an iterator pointing to the first element we want to touch, an iterator pointing to the one-after-the-last element, and a functor. The functor is, as far as we care right now, a function called on each element in the container.

```
// for_each. Apply a function to every element of a range.
template <class iterator, class functor >
functor for_each( iterator curr, iterator last, functor f)
{
    for ( ; curr != last; ++curr)
    {
        f(*curr);
    }
    return f;
}
```

This code will work with any kind of iterator, be it a list iterator or a vector iterator. So you can run the generic function (or any of the other several dozen generic functions and algorithms) on any container. Pretty sweet, huh?

Functors

The last thing we'll talk about in this short run through the STL is functors. Functors are used by many of the generic algorithms and functions (like `for_each`, discussed above). They are classes that implement the parentheses operator. This allows them to mimic the behavior of a regular function, but they can do neat things like save function state (via member variables).

This page intentionally left blank.

Index

- .o3d format, 330-331
- 2D Cartesian coordinate system, 122
- 2D graphics, 41-52
- 32-bit color, 42
- 3D points, 121-122

A

- ACKPacket() function, 287-288
- action steering, 210
- AddACKMessage() function, 285-286
- addition,
 - matrix, 156-157
 - vector, 127-128
- AddPacket() function, 273-274
- AddRef() function, 31
- AdjustWindowRect() function, 19
- affine mapping, 425
- AI, 209-210
 - pattern-based, 212-213
 - rule-based, 234-235
 - scripted, 213
- aliasing, 430 *see also* texture aliasing
- alpha blending, 417-418
 - enabling, 419-423
 - equation, 418
- alpha component, 42, 178, 417
- ambient light, 181
- amplitude, 88
- AND, 238-239
- animation, 341-343
- anisotropic filtering, 435-436
- application, registering, 13
- approximating, 382
- approximative portal rendering, 498-499
- artificial intelligence, *see* AI
- Assign() function, 125-126
- axis-angle rotation matrix, 170-172

B

- back buffer, 43
- back-face culling, 147-148
- Begin() function, 261
- Bezier curve, 355
 - calculating, 361-363
 - defining, 358-360
 - drawing, 355-358
 - using with matrix, 360-361

- big endian, 255-256
- binary tree, *see* BSP tree
- bind() function, 267-268
- blend operations, 421
- blend options, 420-421
- blend state, 419
- border color addressing, 428-429
- BorrowPacket() function, 281
- bounding boxes, 176
- bounding spheres, 176
 - implementing, 176-178
- b-rep, 136
- BSP tree, 187-189
 - algorithms, 194-197
 - creating, 189-194, 197-206
- bSphere3 structure, 176-177
- b-spline curves, 376-377
- b-spline example application, 377-379
- buffers, 325-328
 - back, 43
 - changing format of primary, 105
 - in DirectSound, 90-91
 - primary, 43-44, 90
 - secondary, 90
- butterfly subdivision scheme, 383

C

- C++ , using alpha blending in, 419-422
- C++ class, 13
- cApplication, 24, 27-29
 - modifying, 87, 114
- cDataPacket, 271-272
- cell, 490
- cFwdDiffIterator, 365-367
- cGameError, 24-25
- cGraphicsLayer, 55-56
- chasing algorithm, 211
- cHost, 274-275
- cInputLayer, 75-79
- cKeyboard, 79-83
- clamp addressing, 427-428
- class encapsulation, 23-25
- Cleanup() function, 266-267
- Clear() function, 58
- client area, 4-5
- client space, 17

client/server configuration, 258
 ClientToScreen() function, 18-19
 cLight, 328
 Clip() function, 151-152
 clipping, 148-153
 cModel, 331-332
 cMonitor, 263
 cMouse, 83-87
 cNetClock, 290-292
 collision detection, 175-176
 color, representing, 41-42, 178
 color4 structure, 178-180
 COM, 30, 31-33
 interface, 31
 object, 31
 communication, implementing, 507-514
 Component Object Model, *see* COM
 concave polygons, 135-136
 continuity, 356-358
 control points, 355-356
 convex polygons, 135-136
 cooperative levels, 73
 setting, 104
 coordinate space, 17-18
 left-handed, 122
 coordinate system, 122
 cQueueIn, 272-273
 cQueueOut, 280-281
 Create() function, 58
 CreateBlendState() function, 421
 CreateBuffer() function, 326
 CreateDefaultShader() function, 319-321
 CreateDepthStencilBuffer() function, 306-308
 CreateDepthStencilState() function, 305
 CreateDepthStencilView() function, 306
 CreateDeviceAndSwapChain() function, 300-302
 CreateInputLayout() function, 318
 CreateShaderResourceView() function, 441
 CreateSoundBuffer() function, 105
 CreateTexture2D() function, 52, 303
 CreateThread() function, 261-262
 CreateViewport() function, 309
 CreateWindow() function, 13-15
 cross product, 134-135
 cSlowBezierIterator, 362-363
 cSound, 108-114
 cSoundLayer, 106-108
 cTexture, 438-440
 cThread, 265
 cubic curves, 355
 cubic environment maps, 450-452
 culling, *see* back-face culling
 cUnreliableQueueIn, 284-285
 curves,
 drawing, 367
 subdividing, 380-382
 cViewCone, 493-497
 cWindow, 24-27

D

D3D10_BLEND_DESC structure, 419-420, 422
 D3D10_BUFFER_DESC structure, 326
 D3D10_DEPTH_STENCIL_DESC structure, 303-304
 D3D10_DEPTH_STENCIL_VIEW_DESC structure, 305-306
 D3D10_INPUT_ELEMENT_DESC structure, 317
 D3D10_MAPPED_TEXTURE2D structure, 52
 D3D10_SHADER_RESOURCE_VIEW_DESC structure, 441
 D3D10_STENCIL_OP enumeration, 484
 D3D10_SUBRESOURCE_DATA structure, 326-327
 D3D10_TEXTURE2D_DESC structure, 46-50, 302
 D3D10_VIEWPORT structure, 308
 D3D10CreateDeviceAndSwapChain() function, 60-61
 D3DCMPFUNC enumeration, 484
 D3DX10CreateEffectFromFile() function, 315-316
 D3DX10LoadTextureFromTexture() function, 51
 dark mapping, *see* light mapping
 data, reducing, 295
 data access, 70
 DDS format, 437-438
 DefWindowProc() function, 17
 degenerate patches, 379
 depth buffer, creating, 302-308
 depth buffering, 323-325
 depth ordering, 419
 depth problem, 322-323
 DestroyAll() function, 64
 detail mapping example application, 455-463
 detail maps, 452-455
 detail polygons, 506
 devices, 68-69
 receiving state of, 70
 setting cooperative levels for, 73
 diffuse light, 181
 Dijkstra's algorithm, 221-225
 DIMOUSESTATE structure, 70-71
 Direct3D, 340
 example application, 64-66
 implementing, 53-58
 initializing, 58-64
 shutting down, 64
 textures in, 436-437
 Direct3D device, 299-300
 creating, 58-61
 Direct3D Surface format, *see* DDS format
 DirectDraw, 39-40
 DirectInput, 67-68
 devices, 68
 implementing, 74-87
 keyboard constants, 71-73
 object, 74
 vs. Win32 API, 68

DirectInput8Create() function, 74
 directional lights, 183-184
 creating shader for, 328-330
 DirectSound, 88
 buffers, 90-91
 example application, 114-119
 implementing, 103-114
 interfaces, 89-90
 DirectSoundCreate8() function, 103-104
 DirectX, 35
 DirectX SDK, installing, 36
 DispatchMessage() function, 13
 dissolves, 485
 Dist() function, 126-127
 division, vector, 129-130
 DNS, 257
 domain name server, *see* DNS
 DOS programming vs. Windows programming, 1-2
 dot product, 131-134
 double buffering, 43
 DrawTextString() function, 56-57
 DSBUFFERDESC structure, 91-92
 DumpMessages() function, 57-58
 DXGI_SWAP_CHAIN_DESC structure, 58-59
 dynamic address, 256

E

edge collapse, 401-402
 edge selection, 402
 algorithms, 403-405
 edge split, 380
 emissive light, 181
 encapsulation, 23-25
 End() function, 262
 endianness, 255-256
 environment mapping, 446-452
 equality, vector, 130-131
 error codes, 6
 evading algorithm, 211-212
 exact portal rendering, 497-498
 extrapolation, 295-297

F

filtering, 433-436
 focus, 5, 73-74
 form factor, 411-412
 forward differencing, 363-365
 forward kinematics, 343-346
 frequency, 88
 frustum, 493

G

game example application, 507-514
 genetic algorithms, 233-234
 geographic independence, 293
 geometry shader, 310
 geomorphs, 402
 GetAverageLinkPing() function, 289
 GetClientRect() function, 18

GetDeviceState() function, 70
 GetMessage() function, 12
 GetPacketForResend() function, 281
 GetPreviousPacket() function, 281
 GetReliableData() function, 271, 277-278
 GetStatus() function, 94-95
 gimbal lock, 170-171
 globally unique identifier, *see* GUID
 gloss maps, 463-464
 glow maps, 463
 Gouraud shading, 186-187
 GUID, 33

H

handle, 6
 HasMessages() function, 25
 HelloWorld example program, 7-11
 hierarchical animation model, 342-343
 hill climbing, 233-234
 HLSL shader, creating, 311-312
 homogenous coordinate, 160
 host address, 256
 host names, 257
 HRESULT, 6
 Hungarian notation, 3
 HWND, 6

I

IASetInputLayout() function, 318
 IASetVertexBuffers() function, 327
 ICMP, 259
 ID3D10Device, creating, 300-302
 ID3D10Texture2D interface, 50
 identity matrix, 157
 IDirectInputDevice8 interface, 68
 IK, *see* inverse kinematics
 InitD3D() function, 60, 62-64, 321
 initialization, 13-16
 InitInstance() function, 12-13
 input layout, 316-317
 intensity, 407
 interfaces in DirectSound, 89-90
 Internet Control Message Protocol, *see* ICMP
 Internet Protocol, *see* IP
 interobject communication, 507-511
 interpolating, 382
 inverse kinematics, 346-349
 example application, 349-355
 inverse transformation matrix, 163
 inversion, 174-175
 IP, 259
 address, 256
 IUnknown interface, 31

L

Lambert shading, 186
 leaf-based BSP trees, 189
 leafy BSP trees, 189
 leaves, 188
 left-handed coordinate space, 122

length, 122-123
 libraries, 38
 light mapping, 444-446
 light types, 183-186
 lighting, 178
 models, 180-181
 limit surface, 379
 line segments, comparing, 196-197
 linear filtering, 434-435
 lines, clipping, 148-149
 link ping, 288
 listeners, 74-75
 little endian, 255-256
 local coordinate space, 161
 locality of points, testing, 196
 Lock() function, 95-96
 locomotion, 210
 LookAt matrix, 172-174

M

Mag() function, 126
 magnification, 433
 magnitude, 122-123
 MagSquared() function, 126
 map, 44-45
 Map() function, 52
 mapping, 44
 affine, 425
 perspective, 425-426
 perspective-correct, 425-426
 texture, 423-425
 mask, 380
 mathematical operators, vector, 127-135
 matrices, *see* matrix
 matrix, 156
 addition, 156-157
 inverse of, 174-175
 multiplication, 157-158, 167
 operations, 156-158
 projection, 164-165
 using with Bezier curves, 360-361
 matrix4 structure, 166-167
 menu bar, 4-5
 message pump, 13
 MessageBox() function, 15-16
 messages, 5
 handling in Windows, 5-6
 processing, 6-7
 system, 17
 window, 20-23
 minification, 433
 MIP map, 432-433
 mirror addressing, 427
 mirrors, implementing, 499-502
 Mobots Attack! example game application, 507-514
 modified butterfly subdivision scheme, 383-387
 example application, 387-399
 motivation, 210, 230

motor skills, 210
 MSG structure, 6
 MTUDP classes, 265, 271-278, 284-285, 290-292
 multipass, 443-444
 multiplication,
 matrix, 157-158, 167
 vector, 129-130
 multitexturing, 443-444
 multithreading, 260-261
 mutex, 263-264
 MutexOn() function, 264
 MyRegisterClass() function, 12-13

N

nameless struct, 125
 network communication, 511-514
 network models, 257-258
 network play, implementing, 292-297
 neural networks, 235-236, 238-240
 AND function, 238-239
 example application, 241-253
 OR function, 239
 training, 240
 using in games, 241
 XOR function, 240
 neuron, 236-238
 NFA, 230-232
 node-based BSP trees, 189
 nodes, 188
 non-deterministic finite automata, *see* NFA
 normal, 141
 Normalize() function, 126

O

object viewer example application, 331-339
 objects, 68-69, 153
 GUIDS for, 69
 representing, 154-155
 octrees, 490
 OMSetBlendState() function, 422
 OMSetDepthStencilState() function, 305
 OR, 239
 origin, 122
 overdraw counter, 485

P

packets, 259-260
 painter's algorithm, 322
 paletted color, 41-42
 parallel lights, 183-184
 patch, 355, 367-368, 407
 path, finding shortest, 225-226
 path following, 218-219
 path planning, 218-226
 example application, 227-230
 pattern, 212
 pattern-based AI, 212-213
 peer-to-peer configuration, 257-258
 perspective mapping, 425-426

- perspective projection matrix, 174
- perspective-correct mapping, 425-426
- Phong shading, 187
- physical controllers, 210
- ping, 288
 - calculating, 287-288
- pitch, 45
- pixel, 41
- pixel shader, 310-311
 - creating, 313
- plane, 141-142
 - constructing, 143-144
 - orienting point in, 144-145
 - splitting, 146-147
- plane3 structure, 143-144
- Play() function, 93-94
- point lights, 184-185
- point sampling, 433-434
- POINT structure, 18
- point3 structure, 124-125
 - functions, 125-127
 - operators, 127-135
- points, 121-124
 - 3D, 121-122
 - orienting in plane, 144-145
 - testing locality of, 196
- polygon template class, 136-138
- polygons, 135-138
 - clipping, 149-153
 - ordering, 194-195
- portal, 220, 490
- portal effects, 499-503
- portal generation, 503-505
- portal polygon, 504
- portal rendering, 490-493
 - approximative, 498-499
 - exact, 497-498
 - precalculated, 505-506
- ports, 257
- PostMessage() function, 17
- potential functions, 214-216
 - example application, 216-218
- Potentially Visible Set, *see* PVS
- precalculated portal rendering, 505-506
- prediction, 295-297
- Present() function, 57
- primary buffer, 43-44, 90
 - changing format of, 105
- primitive, 310
- Process() function, 218
- ProcessIncomingACKs() function, 285-286
- ProcessIncomingData() function, 270-271, 276-277, 282-283
- ProcessIncomingReliable() function, 275-276
- progressive meshes, 399-400
 - implementing, 405-406
- progressive radiosity, 410
- projection matrix, 164-165
- protocols, 258-259
- Pump() function, 25
- PVS, 505-506
 - implementing, 506
- Q**
- quadratic curves, 355
- quadtrees, 488-490
- QueryInterface() function, 31
- R**
- radiance, 407
- radiant emitted flux density, 407
- radiosity, 407
 - calculating, 407-410
 - example application, 412-416
 - lighting, 406-407
 - progressive, 410
- rasterization, 300
- RECT structure, 18
- reference counting, 31
- Release() function, 31
- ReliableSendTo() function, 278-280, 283-284
- render target view, creating, 61-62
- resize bars, 4-5
- ReturnPacket() function, 281
- ring configuration, 258
- rotation transformations, 158-159, 169-170
- RSSetViewports() function, 309
- rule-based AI, 234-235
- S**
- scalar quantities, 121
- scan line, 42-43
- scene management, 487-488
- screen space, 17-18
- scripted AI, 213
- secondary buffer, 90
- SendMessage() function, 17
- SendUpdate() function, 294
- SetAddress() function, 275
- SetCooperativeLevel() function, 73, 104
- SetDataFormat() function, 69-70
- SetEvictionPriority() function, 437
- SetResourceArray() function, 443
- shader view, 440-441
 - creating, 441
- shaders, 309-311
 - adding textures to, 441-442
 - creating default, 309-321
 - creating for directional light, 328-330
 - defining technique for, 313-314
 - example application, 314-315
 - sending textures to, 442
 - setting up, 315-321
 - using alpha blending in, 422-423
- shading models, 186-187
- socket() function, 267
- sorted polygon ordering, 194-195
- sound, 87-89
- sound buffers, working with, 93-96

sparse matrix, 410
 specular highlight, 181
 specular light, 181
 specular maps, 452
 specular reflection, 182-183
 spherical coordinates, calculating, 449-450
 spherical environment maps, 446-449
 splines, 355
 Split() function, 149, 152-153
 spotlights, 185-186
 sprites, 341
 Standard Template Library, *see* STL
 StartListening() function, 267
 StartSending() function, 268
 Startup() function, 266
 static address, 256
 steering, 210

- algorithms, 211-230

 stencil buffer, 325, 483

- creating, 302-308
- using, 483-484

 STL, 517

- containers, 518-519
- functors, 521
- iterators, 519-521
- templates, 517-518

 structures,

- bSphere3, 176-177
- color3, 178-180
- D3D10_BLEND_DESC, 419-420, 422
- D3D10_BUFFER_DESC, 326
- D3D10_DEPTH_STENCIL_DESC, 303-304
- D3D10_DEPTH_STENCIL_VIEW_DESC, 305-306
- D3D10_INPUT_ELEMENT_DESC, 317
- D3D10_MAPPED_TEXTURE2D, 52
- D3D10_SHADER_RESOURCE_VIEW_DESC, 441
- D3D10_SUBRESOURCE_DATA, 326-327
- D3D10_TEXTURE2D_DESC, 46-50, 302
- D3D10_VIEWPORT, 308
- DIMOUSESTATE, 70-71
- DSBUFFERDESC, 91-92
- DXGI_SWAP_CHAIN_DESC, 58-59
- matrix4, 166-167
- MSG, 6
- plane3, 143-144
- POINT, 18
- point3, 124-125
- RECT, 18
- WAVEFORMATEX, 92-93

 subdivision schemes, 380-383

- butterfly, 383
- example application, 387-399
- interpolating vs. approximating, 382
- modified butterfly, 383-387
- stationary vs. non-stationary, 383
- triangles vs. quads, 382
- uniform vs. non-uniform, 383

subdivision surfaces, 379
 subtraction, vector, 127-128
 surfaces, drawing, 367-368
 Sutherland-Hodgeman algorithm, 150-151
 swap chains, 46

- creating, 58-61, 300-302

 Synchronize() function, 290-291
 system messages, 17

T

task generation, 210
 TCP, 259
 teapot example application, 369-376
 tearing, 43
 temporal independence, 293
 texel, 424
 texture addressing, 426-429
 texture aliasing, 430-432
 texture arrays, 443
 texture chain, 46
 texture coordinates, 424
 texture description, 46
 texture management, 436-437
 texture mapping, 423-425, 443-483
 texture wrapping, 429-430
 textures, 44-45, 424

- activating, 440-442
- adding to shader, 441-442
- complex, 46
- creating, 52
- describing, 46
- in Direct3D, 436-437
- loading, 51, 437-440
- sampling, 442
- sending to shader, 442
- working with, 51-52

 texturing example application, 465-483
 ThreadProc() function, 262-263, 269-270
 timing, 294
 title bar, 4-5
 ToInverse() function, 175
 topology, 327
 transformation matrix inverse, 163
 transformations, 156

- rotation, 158-159, 169-170
- translation, 156, 168

 TranslateMessage() function, 12-13
 translation transformation, 156, 168
 translocators, 502-503
 Transmission Control Protocol, *see* TCP
 transmission latency ping, 288
 tri template class, 139
 triangle fans, 139-140
 triangle strips, 139-140
 triangles, 138-139
 true color, 42

U

UDP, 259

- implementing, 260-292

unit sphere, 123
 unit vector, 123
 universe box, 504
 Unlock() function, 96
 UpdateLights() function, 329
 UpdateMatrices() function, 319
 User Datagram Protocol, *see* UDP

V

vector, 121
 equality, 130-131
 length of, 123
 magnitude of, 123
 mathematical operators, 127-135
 quantities, 121
 vertex buffers, 325-326
 using, 326-328
 vertex shader, 310
 creating, 312-313
 vertex split, 401-402
 vertical blank, 43
 vertical retrace, 43
 vertices, 153
 view space, 162-163
 viewing cone, 493
 viewport, 308
 creating, 308-309
 Vista, *see* Windows Vista
 Visual C++, setting up, 36-38

W

WAV files, loading, 96-103
 WAVEFORMATEX structure, 92-93
 Win32 API, 2
 vs. DirectInput, 68
 Win32 example program, 7-11
 window, 4
 class, 13
 handle, 6

 initializing, 13-16
 messages, 20-23
 procedure, 7
 Windows, message handling in, 5-6
 Windows programming, 1-3
 vs. DOS programming, 1-2
 Windows Vista, 1-2
 WinMain() function, 11-12, 29-30
 Winsock, 257
 wipes, 485
 WM_CLOSE message, 20
 WM_CREATE message, 20
 WM_DESTROY message, 17, 20
 WM_ERASEBKGD message, 20
 WM_KEYDOWN message, 21
 WM_KEYUP message, 21
 WM_LBUTTONDOWN message, 21-22
 WM_LBUTTONUP message, 23
 WM_MBUTTONDOWN message, 22
 WM_MBUTTONUP message, 23
 WM_MOUSEMOVE message, 21
 WM_MOUSEWHEEL message, 23
 WM_MOVE message, 20
 WM_PAINT message, 17, 20
 WM_QUIT message, 21
 WM_RBUTTONDOWN message, 22
 WM_RBUTTONUP message, 23
 WM_SIZE message, 20
 WndProc() function, 7, 15, 17
 world coordinate space, 161
 world space, 163
 wrap addressing, 426-427

X

XOR, 240

Z

z-buffering, 323-325