Quick answers to common problems

Unity 5.x Shaders and Effects Cookbook

Master the art of Shader programming to bring life to your Unity projects

Alan Zucconi   Kenneth Lammers

[PACKT]
PUBLISHING

Unity 5.x Shaders and Effects Cookbook

Master the art of Shader programming to bring life to your Unity projects

Alan Zucconi

Kenneth Lammers

[PACKT]
PUBLISHING

BIRMINGHAM - MUMBAI

Unity 5.x Shaders and Effects Cookbook

Credits

Authors
Alan Zucconi

Kenneth Lammers

Reviewer
Kenneth Lammers

Commissioning Editor
Priya Singh

Acquisition Editors
Rahul Nair

Erol Staveley

Content Development Editor
Mehvash Fatima

Technical Editors
Pranil Pathare

Danish Shaikh

Copy Editor
Tasneem Fatehi

Project Coordinator
Kinjal Bari

Proofreader
Safis Editing

Indexer
Monica Ajmera Mehta

Graphics
Kirk D'Penha

Disha Haria

Production Coordinator
Nilesh Mohite

Cover Work
Nilesh Mohite

About the Authors

Alan Zucconi is a passionate developer, author, and motivational speaker, recognized as one of Develop's "30 under 30." His expertise has been built over the past 10 years, while he dedicated his time to academia and the gaming industry. He started his independent career to fully explore his creativity, tearing down the wall between art and gaming. Prior to that, he worked at Imperial College London, where he discovered his passion for teaching and writing. His titles include the gravity puzzle, *ORBITALIS*, and the upcoming time travel platformer, *Still Time*.

Kenneth Lammers has over 15 years of experience in the gaming industry, working as a character artist, technical artist, technical art director, and programmer. Throughout his career, he has worked on titles such as Call of Duty 3, Crackdown 2, Alan Wake, and Kinect Star Wars. He currently owns and operates Ozone Interactive along with his business partner, Noah Kaarbo. Together, they have worked with clients such as Amazon, Eline Media, IGT, and Microsoft.

Kenny has worked for Microsoft Games Studios, Activision, and Surreal, and has recently gone out on his own, operating CreativeTD and Ozone Interactive.

Kenny authored the first version of *Unity Shaders and Effects Cookbook* by *Packt Publishing*, and was very happy to be a part of the writing, updating and reviewing of this book.

www.PacktPub.com

eBooks, discount offers, and more

Did you know that Packt offers eBook versions of every book published, with PDF and ePub files available? You can upgrade to the eBook version at `www.PacktPub.com` and as a print book customer, you are entitled to a discount on the eBook copy. Get in touch with us at `customercare@packtpub.com` for more details.

At `www.PacktPub.com`, you can also read a collection of free technical articles, sign up for a range of free newsletters and receive exclusive discounts and offers on Packt books and eBooks.

`https://www2.packtpub.com/books/subscription/packtlib`

Do you need instant solutions to your IT questions? PacktLib is Packt's online digital book library. Here, you can search, access, and read Packt's entire library of books.

Why subscribe?

- ▶ Fully searchable across every book published by Packt
- ▶ Copy and paste, print, and bookmark content
- ▶ On demand and accessible via a web browser

Table of Contents

Preface

Unity 5.x Shaders and Effects Cookbook is your guide to becoming familiar with the creation of shaders and post effects in Unity 5. You will start your journey at the beginning, creating the most basic shaders and learning how the shader code is structured. This foundational knowledge will arm you with the means to progress further through each chapter, learning advanced techniques such as volumetric explosions and fur shading. This edition of the book is written specifically for Unity 5 and will help you to master physically-based rendering and global illumination to get as close to photorealism as possible.

By the end of each chapter, you will have gained new skill sets that will increase the quality of your shaders and even make your shader writing process more efficient. These chapters have been tailored so that you can jump into each section and learn a specific skill from beginner to expert. For those who are new to shader writing in Unity, you can progress through each chapter, one at a time, to build on your knowledge. Either way, you will learn the techniques that make modern games look the way they do.

Once you have completed this book, you will have a set of shaders that you can use in your Unity 3D games as well as the understanding to add to them, accomplish new effects, and address performance needs. So let's get started!

What this book covers

Chapter 1, *Creating Your First Shader*, introduces you to the world of shader coding in Unity 4 and 5.

Chapter 2, *Surface Shaders and Texture Mapping*, covers the most common and useful techniques that you can implement with Surface Shaders, including how to use textures and normal maps for your models.

Chapter 3, *Understanding Lighting Models*, gives you an in-depth explanation of how shaders model the behavior of light. The chapter teaches you how to create custom lighting models used to simulate special effects such as toon shading.

Chapter 4, *Physically Based Rendering in Unity 5*, shows you that physically-based rendering is the standard technology used by Unity 5 to bring realism to your games. This chapter explains how to make the most out of it, mastering transparencies, reflective surfaces, and global illumination.

Chapter 5, *Vertex Functions*, teaches you how shaders can be used to alter the geometry of an object; this chapter introduces vertex modifiers and uses them to bring volumetric explosions, snow shaders, and other effects to life.

Chapter 6, *Fragment Shaders and Grab Passes*, explains how to use grab passes to make materials that emulate the deformations generated by these semi-transparent materials.

Chapter 7, *Mobile Shader Adjustment*, helps you optimize your shaders to get the most out of any devices.

Chapter 8, *Screen Effects with Unity Render Textures*, shows you how to create special effects and visuals that would otherwise be impossible to achieve.

Chapter 9, *Gameplay and Screen Effects*, tells you how post-processing effects can be used to complement your gameplay, simulating, for instance, a night vision effect.

Chapter 10, *Advanced Shading Techniques*, introduces the most advanced techniques in this book, such as fur shading and heatmap rendering.

What you need for this book

The following is a list of the required and optional software to complete the recipes in this book:

 ▸ Unity 5
 ▸ A 3D application such as Maya, Max, or Blender (optional)
 ▸ A 2D image editing application such as Photoshop or Gimp (optional)

Who this book is for

This book is written for developers who want to create their first shaders in Unity 5 or wish to take their game to a whole new level by adding professional post-processing effects. A solid understanding of Unity is required.

Sections

In this book, you will find several headings that appear frequently (Getting ready, How to do it, How it works, There's more, and See also).

To give clear instructions on how to complete a recipe, we use these sections as follows:

Getting ready

This section tells you what to expect in the recipe, and describes how to set up any software or any preliminary settings required for the recipe.

How to do it...

This section contains the steps required to follow the recipe.

How it works...

This section usually consists of a detailed explanation of what happened in the previous section.

There's more...

This section consists of additional information about the recipe in order to make the reader more knowledgeable about the recipe.

See also

This section provides helpful links to other useful information for the recipe.

Conventions

In this book, you will find a number of text styles that distinguish between different kinds of information. Here are some examples of these styles and an explanation of their meaning.

Code words in text, database table names, folder names, filenames, file extensions, pathnames, dummy URLs, user input, and Twitter handles are shown as follows: "Enter the following code into the `Properties` block of your shader."

A block of code is set as follows:

```
void surf (Input IN, inout SurfaceOutput o)
{
  float4 c;
  c = pow((_EmissiveColor + _AmbientColor), _MySliderValue);
  o.Albedo = c.rgb;
  o.Alpha = c.a;
}
```

When we wish to draw your attention to a particular part of a code block, the relevant lines or items are set in bold:

```
void surf (Input IN, inout SurfaceOutputStandard o) {
  fixed4 c = pow((_Color + _AmbientColor), _MySliderValue);
  o.Albedo = c.rgb;
  o.Metallic = _Metallic;
  o.Smoothness = _Glossiness;
  o.Alpha = c.a;
}
```

New terms and **important words** are shown in bold. Words that you see on the screen, for example, in menus or dialog boxes, appear in the text like this: "In the **Project** tab in your Unity editor, right-click on the **Assets** folder and select **Create** | **Folder**."

> Warnings or important notes appear in a box like this.

> Tips and tricks appear like this.

Reader feedback

Feedback from our readers is always welcome. Let us know what you think about this book—what you liked or disliked. Reader feedback is important for us as it helps us develop titles that you will really get the most out of.

To send us general feedback, simply e-mail `feedback@packtpub.com`, and mention the book's title in the subject of your message.

If there is a topic that you have expertise in and you are interested in either writing or contributing to a book, see our author guide at `www.packtpub.com/authors`.

Customer support

Now that you are the proud owner of a Packt book, we have a number of things to help you to get the most from your purchase.

Downloading the example code

You can download the example code files for this book from your account at
`http://www.packtpub.com`. If you purchased this book elsewhere, you can visit
`http://www.packtpub.com/support` and register to have the files e-mailed directly
to you.

You can download the code files by following these steps:

1. Log in or register to our website using your e-mail address and password.
2. Hover the mouse pointer on the **SUPPORT** tab at the top.
3. Click on **Code Downloads & Errata**.
4. Enter the name of the book in the **Search** box.
5. Select the book for which you're looking to download the code files.
6. Choose from the drop-down menu where you purchased this book from.
7. Click on **Code Download**.

Once the file is downloaded, please make sure that you unzip or extract the folder using the
latest version of:

- WinRAR / 7-Zip for Windows
- Zipeg / iZip / UnRarX for Mac
- 7-Zip / PeaZip for Linux

Downloading the color images of this book

We also provide you with a PDF file that has color images of the screenshots/diagrams used
in this book. The color images will help you better understand the changes in the output. You
can download this file from `https://www.packtpub.com/sites/default/files/
downloads/Unity5xShadersAndEffectsCookbook_SecondEdition_Graphics.pdf`.

Errata

Although we have taken every care to ensure the accuracy of our content, mistakes do happen.
If you find a mistake in one of our books—maybe a mistake in the text or the code—we would be
grateful if you could report this to us. By doing so, you can save other readers from frustration
and help us improve subsequent versions of this book. If you find any errata, please report them
by visiting `http://www.packtpub.com/submit-errata`, selecting your book, clicking on
the **Errata Submission Form** link, and entering the details of your errata. Once your errata are
verified, your submission will be accepted and the errata will be uploaded to our website or
added to any list of existing errata under the Errata section of that title.

To view the previously submitted errata, go to `https://www.packtpub.com/books/content/support` and enter the name of the book in the search field. The required information will appear under the **Errata** section.

Piracy

Piracy of copyrighted material on the Internet is an ongoing problem across all media. At Packt, we take the protection of our copyright and licenses very seriously. If you come across any illegal copies of our works in any form on the Internet, please provide us with the location address or website name immediately so that we can pursue a remedy.

Please contact us at `copyright@packtpub.com` with a link to the suspected pirated material.

We appreciate your help in protecting our authors and our ability to bring you valuable content.

Questions

If you have a problem with any aspect of this book, you can contact us at `questions@packtpub.com`, and we will do our best to address the problem.

1

Creating Your First Shader

This chapter will cover some of the more common diffuse techniques found in today's Game Development Shading Pipelines. In this chapter, you will learn about the following recipes:

- ▶ Creating a basic Standard Shader
- ▶ Migrating Legacy Shaders from Unity 4 to Unity 5
- ▶ Adding properties to a shader
- ▶ Using properties in a Surface Shader

Introduction

Let's imagine a cube that has been painted white uniformly. Even if the color used is the same on each face, they will all have different shades of white depending on the direction that the light is coming from and the angle that we are looking at it. This extra level of realism is achieved in 3D graphics by **shaders**, special programs that are mostly used to simulate how light works. A wooden cube and a metal one may share the same 3D model, but what makes them look different is the shader that they use. Recipe after recipe, this first chapter will introduce you to shader coding in Unity. If you have little to no previous experience with shaders, this chapter is what you need to understand what shaders are, how they work, and how to customize them.

By the end of this chapter, you will have learned how to build basic shaders that perform basic operations. Armed with this knowledge, you will be able to create just about any Surface Shader.

Creating a basic Standard Shader

Every Unity game developer should be familiar with the concept of **components**. All the objects that are part of a game contain a number of components that affect their look and behavior. While **scripts** determine how objects should behave, **renderers** decide how they should appear on the screen. Unity comes with several renderers, depending on the type of object that we are trying to visualise; every 3D model typically has `MeshRenderer`. An object should have only one renderer, but the renderer itself can contain several **materials**. Each material is a wrapper for a single shader, the final ring in the food chain of 3D graphics. The relationships between these components can be seen in the following diagram:
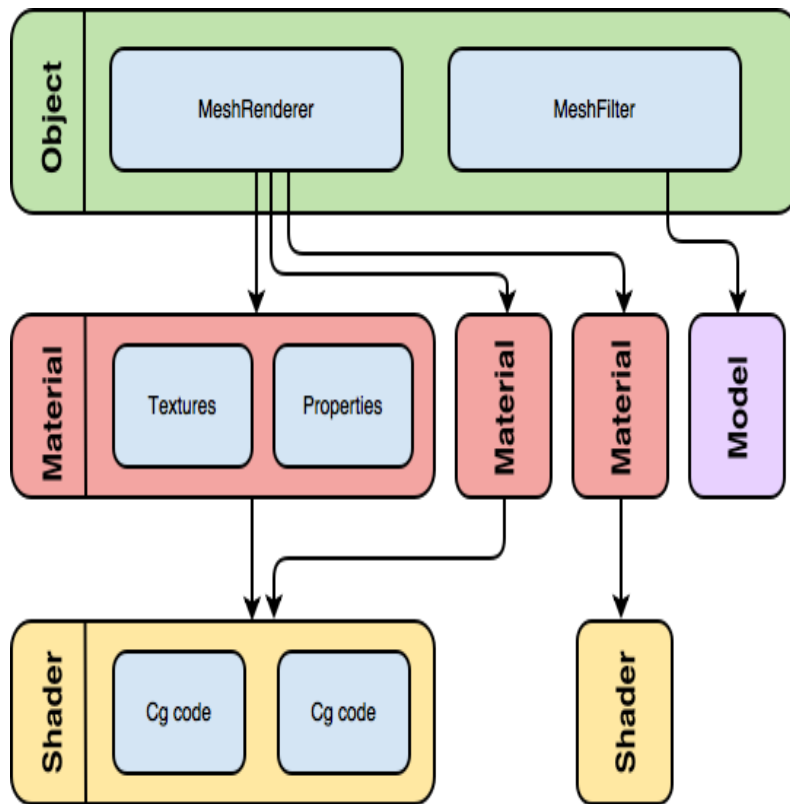
Understanding the difference between these components is essential to understand how shaders work.

Getting ready

To get started with this recipe, you will need to have Unity 5 running and must have created a new project. There will also be a Unity project included with this cookbook, so you can use that one as well and simply add your own custom shaders to it as you step through each recipe. With this completed, you are now ready to step into the wonderful world of real-time shading!

How to do it...

Before getting into our first shader, let's create a small scene for us to work with. This can be done by navigating to **GameObject | Create Empty** in the Unity editor. From here, you can create a plane to act as a ground, a couple of spheres to which we will apply our shader, and a directional light to give the scene some light. With our scene generated, we can move on to the shader writing steps:

1. In the **Project** tab in your Unity editor, right-click on the **Assets** folder and select **Create | Folder**.

    > If you are using the Unity project that came with the cookbook, you can skip to step 4.

2. Rename the folder that you created to `Shaders` by right-clicking on it and selecting **Rename** from the drop-down list or selecting the folder and hitting *F2* on the keyboard.

3. Create another folder and rename it to `Materials`.

4. Right-click on the `Shaders` folder and select **Create | Shader**. Then right-click on the `Materials` folder and select **Create | Material**.

5. Rename both the shader and material to `StandardDiffuse`.

6. Launch the `StandardDiffuse` shader in **MonoDevelop** (the default script editor for Unity) by double-clicking on it. This will automatically launch the editor for you and display the shader code.

    > You will see that Unity has already populated our shader with some basic code. This, by default, will get you a basic Diffuse shader that accepts one texture. We will be modifying this base code so that you can learn how to quickly start developing your own custom shaders.

7. Now let's give our shader a custom folder from which it's selected. The very first line of code in the shader is the custom description that we have to give the shader so that Unity can make it available in the shader drop-down list when assigning to materials. We have renamed our path to `Shader "CookbookShaders/ StandardDiffuse"`, but you can name it to whatever you want and rename it at any time. So don't worry about any dependencies at this point. Save the shader in MonoDevelop and return to the Unity editor. Unity will automatically compile the shader when it recognizes that the file has been updated. This is what your shader should look like at this point:

```
Shader "CookbookShaders/StandardDiffuse" {
  Properties {
    _Color ("Color", Color) = (1,1,1,1)
    _MainTex ("Albedo (RGB)", 2D) = "white" {}
    _Glossiness ("Smoothness", Range(0,1)) = 0.5
    _Metallic ("Metallic", Range(0,1)) = 0.0
  }
  SubShader {
    Tags { "RenderType"="Opaque" }
    LOD 200

    CGPROGRAM
    // Physically based Standard lighting model, and enable
      shadows on all light types
    #pragma surface surf Standard fullforwardshadows

    // Use shader model 3.0 target, to get nicer looking
      lighting
    #pragma target 3.0

    sampler2D _MainTex;

    struct Input {
      float2 uv_MainTex;
    };

    half _Glossiness;
    half _Metallic;
    fixed4 _Color;

    void surf (Input IN, inout SurfaceOutputStandard o) {
      // Albedo comes from a texture tinted by color
      fixed4 c = tex2D (_MainTex, IN.uv_MainTex) * _Color;
      o.Albedo = c.rgb;
      // Metallic and smoothness come from slider variables
      o.Metallic = _Metallic;
      o.Smoothness = _Glossiness;
```

```
        o.Alpha = c.a;
      }
      ENDCG
  }
  FallBack "Diffuse"
}
```

8. Technically speaking, this is a **Surface Shader** based on **physically-based rendering**, which Unity 5 has adopted as its new standard. As the name suggests, this type of shader achieves realism by simulating how light physically behaves when hitting objects. If you are using a previous version of Unity (such as Unity 4), your code will look very different. Prior to the introduction of physically-based shaders, Unity 4 used less sophisticated techniques. All these different types of shader will be further explored in the next chapters of this book.

9. After your shader is created, we need to connect it to a material. Select the material called `StandardDiffuse` that we created in step 4 and look at the **Inspector** tab. From the **Shader** drop-down list, select **CookbookShaders | StandardDiffuse**. (Your shader path might be different if you chose to use a different path name.) This will assign your shader to your material and make it ready for you to assign to an object.

> To assign a material to an object, you can simply click and drag your material from the **Project** tab to the object in your scene. You can also drag a material on to the **Inspector** tab of an object in the Unity editor to assign a material.

The screenshot of an example is as follows:

Not much to look at at this point, but our shader development environment is set up and we can now start to modify the shader to suit our needs.

How it works...

Unity has made the task of getting your shader environment up and running, which is very easy for you. It is simply a matter of a few clicks and you are good to go. There are a lot of elements working in the background with regard to the Surface Shader itself. Unity has taken the Cg shader language and made it more efficient to write by doing a lot of the heavy Cg code lifting for you. The Surface Shader language is a more component-based way of writing shaders. Tasks such as processing your own texture coordinates and transformation matrices have already been done for you, so you don't have to start from scratch any more. In the past, we would have to start a new shader and rewrite a lot of code over and over again. As you gain more experience with Surface Shaders, you will naturally want to explore more of the underlying functions of the Cg language and how Unity is processing all of the low-level **graphics processing unit** (**GPU**) tasks for you.

> All the files in a Unity project are referenced independently from the folder that they are in. We can move shaders and materials from within the editor without the risk of breaking any connection. Files, however, should never be moved from outside the editor as Unity will not be able to update their references.

So, by simply changing the shader's path name to a name of our choice, we have got our basic Diffuse shader working in the Unity environment, with lights and shadows and all that by just changing one line of code!

See also

The source code of the built-in shaders is typically hidden in Unity 5. You cannot open this from the editor like you do with your own shaders.

For more information on where to find a large portion of the built-in Cg functions for Unity, go to your Unity install directory and navigate to `Unity45\Editor\Data\CGIncludes`. In this folder, you can find the source code of the shaders shipped with Unity. Over time, they have changed a lot; **UNITY DOWNLOAD ARCHIVE** (`https://unity3d.com/get-unity/download/archive`) is the place to go if you need to access the source codes of a shader used in a different version of Unity. After choosing the right version, select **Built in shaders** from the drop-down list, as shown in the following image. There are three files that are of note at this point—`UnityCG.cginc`, `Lighting.cginc`, and `UnityShaderVariables.cginc`. Our current shader is making use of all these files at the moment:
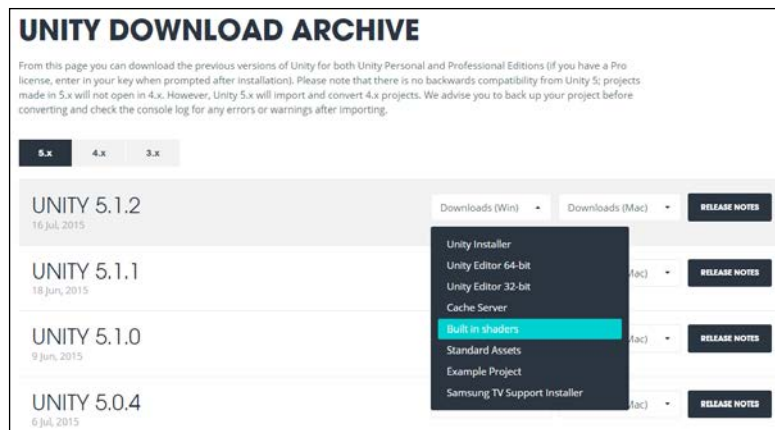
Chapter 10, *Advanced Shading Techniques*, will explore in-depth how to use GcInclude for a modular approach to shader coding.

Migrating Legacy Shaders from Unity 4 to Unity 5

It is undeniable that graphics in videogames have changed massively over the last 10 years. Every new game comes with cutting-edge techniques that are getting us closer to achieving real-time photorealism. We should not be surprised by the fact that shaders themselves have changed massively throughout the lifetime of Unity. This is one of the major sources of confusion when approaching shaders for the first time. Prior to Unity 5, mainly two different shaders were adopted: **Diffuse** and **Specular**. As the names suggest, they were used for matte and shiny materials, respectively. If you are already using Unity 5, you can skip this recipe. This recipe will explain how to replicate these effects using Unity 5.

Getting ready

The starting point of this recipe is having a workspace made in Unity 4, which uses some of the built-in shaders that were originally provided. If you are to start a new game, there is no doubt that you should use the latest version of Unity. However, if your project is already in the later stages of development with an older version, you should be very careful before migrating. Many things have changed behind the curtains of the engine, and even if your built-in shaders will most likely work without any problem, your scripts might not. If you are to migrate your entire workspace, the first thing that you should do is take backup. It is important to remember that saving assets and scenes is not enough as most of the configuration in Unity is stored in its metadata files. The safest option to survive a migration is to duplicate the entire folder that contains your project. The best way of doing this is by physically copying the folder from File Explorer (Windows) or Finder (Mac).

How to do it...

There are two main options if you want to migrate your built-in shaders: upgrading your project automatically or switching to Standard Shaders instead.

Upgrading automatically

This option is the easiest one. Unity 5 can import a project made with an earlier version and upgrade it. You should notice that once the conversion is done, you will not be able to use Unity 4; even if none of your assets may have changed directly, Unity metadata has been converted. To proceed with this, open Unity 5 and click on **OPEN OTHER** to select the folder of your old project. You will be asked if you want to convert it; click on **Upgrade** to proceed. Unity will reimport all of your assets and recompile all of your scripts. The process might last for several hours if your project is big. Once the conversion is done, your built-in shaders from Unity 4 should have been replaced with their legacy equivalent. You can check this from the inspector of your materials that should have changed (for instance) from **Bumped Diffuse** to **Legacy Shader/Bumped Diffuse**.

> Even if Diffuse, Specular, and the other built-in shaders from Unity 4 are now deprecated, Unity 5 keeps them for backward compatibility. They can be found in the drop-down menu of a material under the **Legacy Shaders** folder.

Using Standard Shaders

Instead of using the Legacy Shaders, you might decide to replace them with the new Standard Shaders from Unity 5. Before doing this, you should keep in mind that as they are based on a different lighting model, your materials will most likely look different. Unity 4 came with more than eighty different built-in shaders divided in six different families (Normal, Transparent, Transparent Cutout, Self-Illuminated, and Reflective). In Unity 5, they are all replaced by the Standard Shader introduced in the previous recipe. Unfortunately, there is no magic recipe to convert your shaders directly. However, you can use the following table as a starting point to understand how the Standard Shader can be configured to simulate Unity 4 Legacy Shaders:

Shader	Unity 4	Unity 4 (Legacy)	Unity 5
Diffuse	Diffuse	Legacy Shader/Diffuse	Standard
	Lambert	Lambert	Physically-based rendering: Metallic Workflow
Specular	Specular	Legacy Shader/Specular	Standard (Specular setup)
	Blinn-Phong	Blinn-Phong	Physically-based rendering: Specular Workflow

Shader	Unity 4	Unity 4 (Legacy)	Unity 5
Transparent	Transparent Vertex-Lit	Legacy Shader/Transparent Vertex-Lit	Standard Rendering Mode: Transparent
	Transparent Cutout Vertex-Lit	Legacy Shader/Transparent Cutout Vertex-Lit	Standard Rendering Mode: Cutout

You can change the shader used by your old material using the **Shader** drop-down menu in **Inspector**. All you need to do is simply select the appropriate Standard Shader. If your old shader used textures, colours, and normal maps, they will be automatically used in the new Standard Shader. You might still have to configure the parameters of the Standard Shader to get as close to your original lighting model as possible. The following picture shows the ubiquitous *Stanford bunny* rendered with a Legacy Diffuse Shader (right), converted Standard Shader (left), and Standard Shader with **Smoothness** set to zero (middle):

Migrating custom shaders

If you have written custom shaders in Unity 4, chances are that this will work straightaway in Unity 5. Despite this, Unity has made some minor changes in the way shaders work, which can cause both errors and inconsistencies. The most relevant and important one is the intensity of the light. Lights in Unity 5 are twice as bright. All the Legacy Shaders have been rewritten to take this into account; if you have upgraded your shaders or switched to Standard Shaders, you will not notice any difference. If you have written your own lighting model, you will have to be sure that the intensity of the light is not multiplied by two any more. The following code is used to ensure this:

```
// Unity 4
c.rgb = s.Albedo * _LightColor0.rgb * (diff * atten * 2);
// Unity 5
c.rgb = s.Albedo * _LightColor0.rgb * (diff * atten);
```

If you haven't written a shader yet, don't panic: lighting models will be extensively explained in *Chapter 3*, *Understanding Lighting Models*.

> There are several other changes in the way Unity 5 handles shaders compared to Unity 4. You can see all of them in **Shaders in Unity 5.0** at `http://docs.unity3d.com/Manual/UpgradeGuide5-Shaders.html`.

How it works...

Writing shaders is always a trade-off between realism and efficiency; realistic shaders require intensive computation, potentially introducing a significant lag. It's important to use only those effects that are strictly required: if a material does not need specular reflections, then there is no need to use a shader that calculates them. This has been the main reason why Unity 4 has been shipped with so many different shaders. The new Standard Shader of Unity 5 can potentially replace all of the previous shaders as it incorporates normal mapping, transparency, and reflection. However, it has been cleverly optimized so that only the effects that are really necessary are calculated. If your standard material does not have reflections, they will not be calculated.

Despite this, the Standard Shader is mainly designed for realistic materials. The Legacy Diffuse and Specular shaders, in comparison, were not really designed for realistic materials. This is the reason switching from Legacy to Standard Shaders will mostly introduce slight changes in the way your objects are rendered.

See also

- *Chapter 3*, *Understanding Lighting Models*, explores in-depth how the Diffuse and Specular shaders work. Even if deprecated in Unity 5, understanding them is essential if you want to design new lighting models.
- *Chapter 4*, *Physically Based Rendering in Unity 5*, will show you how to unlock the potential of the Standard Shader in Unity 5.

Adding properties to a shader

Properties of a shader are very important for the shader pipeline as they are the method that you use to let the artist or user of the shader assign textures and tweak your shader values. Properties allow you to expose GUI elements in a material's **Inspector** tab without you having to use a separate editor, which provides visual ways to tweak a shader.

With your shader opened in MonoDevelop, look at the block of lines 2 through 7. This is called the `Properties` block. Currently, it will have one property in it called `_MainTex`. If you look at your material that has this shader applied to it, you will notice that there is one **texture** GUI element in the **Inspector** tab. These lines of code in our shader are creating this GUI element for us.

Again, Unity has made this process very efficient in terms of coding and the amount of time it takes to iterate through changing your properties.

Getting ready

Let's see how this works in our current shader called `StandardDiffuse`, by creating our own properties and learning more about the syntax involved. For this example, we will refit the shader previously created. Instead of using a texture, it will only use its color and some other properties that we will be able to change directly from the **Inspector** tab. Start by duplicating the `StandardDiffuse` shader. You can do this by selecting it in the **Inspector** tab and pressing *Ctrl + D*. This will create a copy called `StandardDiffuse2`.

> You can give a friendlier name to your shader in its first line. For instance, `Shader "CookbookShaders/StandardDiffuse"` tells Unity to call this `StandardDiffuse` shader and move it to a group called `CookbookShaders`. If you duplicate a shader using *Ctrl + D*, your new file will share the same name. To avoid confusion, make sure to change the first line of each new shader so that it uses a unique alias.

How to do it...

Once the `StandardDiffuse2` shader is ready, we can start changing its properties:

1. In our `Properties` block of our shader, remove the current property by deleting the following code from our current shader:

   ```
   _MainTex ("Albedo (RGB)", 2D) = "white" {}
   ```

2. As we have removed an essential property, this shader will not compile until the other references to `_MainTex` are removed. Let's remove this other line:

   ```
   sampler2D _MainTex;
   ```

3. The original shader used `_MainTex` to color the model. Let's change this by replacing the first line of code of the `surf()` function with this:

   ```
   fixed4 c = _Color;
   ```

4. When you save and return to Unity, the shader will compile, and you will see that now our material's **Inspector** tab doesn't have a texture swatch anymore. To complete the refit of this shader, let's add one more property and see what happens. Enter the following code:

```
_AmbientColor ("Ambient Color", Color) = (1,1,1,1)
```

5. We have added another color swatch to the material's **Inspector** tab. Now let's add one more to get a feel for other kinds of properties that we can create. Add the following code to the `Properties` block:

```
_MySliderValue ("This is a Slider", Range(0,10)) = 2.5
```

6. We have now created another GUI element that allows us to visually interact with our shader. This time, we created a slider with the name **This is a Slider**, as shown in the following screenshot:

Properties allow you to create a visual way to tweak shaders without having to change values in the shader code itself. The next recipe will show you how these properties can actually be used to create a more interesting shader.

> While properties belong to shaders, the values associated with them are stored in materials. The same shader can be safely shared between many different materials. On the other hand, changing the property of a material will affect the look of all the objects that are currently using it.

How it works...

Every Unity shader has a built-in structure that it is looking for in its code. The `Properties` block is one of those functions that are expected by Unity. The reason behind this is to give you, the shader programmer, a means of quickly creating GUI elements that tie directly into your shader code. These properties that you declare in the `Properties` block can then be used in your shader code to change values, colors, and textures. The syntax to define a property is as follows:

Let's take a look at what is going on under the hood here. When you first start writing a new property, you will need to give it a **Variable Name**. The variable name is going to be the name that your shader code is going to use in order to get the value from the GUI element. This saves us a lot of time because we don't have to set up this system ourselves.

The next elements of a property are the **Inspector GUI Name** and **Type** of the property, which is contained within parentheses. The **Inspector GUI Name** is the name that is going to appear in the material's **Inspector** tab when the user is interacting with and tweaking the shader. The **Type** is the type of data that this property is going to control. There are many types that we can define for properties inside of Unity shaders. The following table describes the types of variables that we can have in our shaders:

Surface Shader property types	
`Range (min, max)`	This creates a float property as a slider from the minimum value to the maximum value
`Color`	This creates a color swatch in the **Inspector** tab that opens up a `color picker = (float,float,float,float)`
`2D`	This creates a texture swatch that allows a user to drag a texture in the shader
`Rect`	This creates a non-power-of-2 texture swatch and functions the same as the 2D GUI element
`Cube`	This creates a cube map swatch in the **Inspector** tab and allows a user to drag and drop a cube map in the shader
`Float`	This creates a float value in the **Inspector** tab but without a slider
`Vector`	This creates a four-float property that allows you to create directions or colors

Finally, there is the **Default Value**. This simply sets the value of this property to the value that you place in the code. So, in the previous example image, the default value for the property named `_AmbientColor`, which is of the `Color` type, is set to a value of `1,1,1,1`. As this is a color property expecting a color that is `RGBA` or `float4` or `r, g, b, a = x, y, z, w,` this color property, when first created, is set to white.

See also

The properties are documented in the Unity manual at `http://docs.unity3d.com/Documentation/Components/SL-Properties.html`.

Using properties in a Surface Shader

Now that we have created some properties, let's actually hook them up to the shader so that we can use them as tweaks to our shader and make the material process much more interactive.

We can use the properties' values from the material's **Inspector** tab because we have attached a variable name to the property itself, but in the shader code, you have to set up a couple of things before you can start calling the value by its variable name.

How to do it...

The following steps show you how to use the properties in a Surface Shader:

1. To begin, let's remove the following lines of code, as we deleted the property called `_MainTex` in the *Creating a basic Standard Shader* recipe of this chapter:

   ```
   _MainTex ("Albedo (RGB)", 2D) = "white" {}
   sampler2D _MainTex;
   fixed4 c = tex2D (_MainTex, IN.uv_MainTex) * _Color;
   ```

2. Next, add the following lines of code to the shader, below the `CGPROGRAM` line:

   ```
   float4 _AmbientColor;
   float _MySliderValue;
   ```

3. With step 2 complete, we can now use the values from the properties in our shader. Let's do this by adding the value from the `_Color` property to the `_AmbientColor` property and giving the result of this to the `o.Albedo` line of code. So, let's add the following code to the shader in the `surf()` function:

   ```
   void surf (Input IN, inout SurfaceOutputStandard o) {
     fixed4 c = pow((_Color + _AmbientColor),
       _MySliderValue);
     o.Albedo = c.rgb;
     o.Metallic = _Metallic;
   ```

```
    o.Smoothness = _Glossiness;
    o.Alpha = c.a;
}
```

4. Finally, your shader should look like the following shader code. If you save your shader in MonoDevelop and re-enter Unity, your shader will compile. If there were no errors, you will now have the ability to change the ambient and emissive colors of the material as well as increase the saturation of the final color using the slider value. Pretty neat, huh!

```
Shader "CookbookShaders/StandardDiffuse3" {
  // We define Properties in the properties block
  Properties {
    _Color ("Color", Color) = (1,1,1,1)
    _AmbientColor("Ambient Color", Color) = (1,1,1,1)
    _MySliderValue("This is a Slider", Range(0,10)) = 2.5
  }
  SubShader {
    Tags { "RenderType"="Opaque" }
    LOD 200

    // We need to declare the properties variable type
      inside of the
    // CGPROGRAM so we can access its value from the
      properties block.
    CGPROGRAM
    #pragma surface surf Standard fullforwardshadows
    #pragma target 3.0

    struct Input {
      float2 uv_MainTex;
    };


    fixed4 _Color;
    float4 _AmbientColor;
    float _MySliderValue;

    void surf (Input IN, inout SurfaceOutputStandard o) {
      // We can then use the properties values in our
        shader
      fixed4 c = pow((_Color + _AmbientColor),
        _MySliderValue);
      o.Albedo = c.rgb;
      o.Alpha = c.a;
    }
```

```
      ENDCG
   }
   FallBack "Diffuse"
}
```

Downloading the example code

You can download the example code files for all Packt books you have purchased from your account at `http://www.packtpub.com`. If you purchased this book elsewhere, you can visit `http://www.packtpub.com/support` and register to have the files e-mailed directly to you.

The `pow(arg1, arg2)` function is a built-in function that will perform the equivalent math function of power. So, argument 1 is the value that we want to raise to a power and argument 2 is the power that we want to raise it to.

To find out more about the `pow()` function, look at the Cg tutorial. It is a great free resource that you can use to learn more about shading and get a glossary of all the functions available to you in the Cg shading language:

`http://http.developer.nvidia.com/CgTutorial/cg_tutorial_appendix_e.html`

The following screenshot demonstrates the result obtained using our properties to control our material's colors and saturation from within the material's **Inspector** tab:

How it works...

When you declare a new property in the `Properties` block, you are providing a way for the shader to retrieve the tweaked value from the material's **Inspector** tab. This value is stored in the variable name portion of the property. In this case, `_AmbientColor`, `_Color`, and `_MySliderValue` are the variables in which we are storing the tweaked values. In order for you to be able to use the value in the `SubShader{}` block, you need to create three new variables with the same names as the property's variable name. This automatically sets up a link between these two so that they know they have to work with the same data. Additionally, it declares the type of data that we want to store in our subshader variables, which will come in handy when we look at optimizing shaders in a later chapter.

Once you have created the subshader variables, you can then use the values in the `surf()` function. In this case, we want to add the `_Color` and `_AmbientColor` variables together and take it to a power of whatever the `_MySliderValue` variable is equal to in the material's **Inspector** tab.

The vast majority of shaders start out as Standard Shaders and get modified until they match the desired look. We have now created the foundation for any Surface Shader you will create that requires a diffuse component.

> Materials are **assets**. This means that any change made to them while your game is running in the editor are permanent. If you have changed the value of a property by mistake, you can undo it using *Ctrl + Z*.

There's more...

Like any other programming language, Cg does not allow mistakes. As such, your shader will not work if you have a typo in your code. When this happens, your materials are rendered in unshaded magenta:

When a script does not compile, Unity prevents your game from being exported or even executed. Conversely, errors in shaders do not stop your game from being executed.

If one of your shaders appears as magenta, it is time to investigate where the problem is. If you select the incriminated shader, you will see a list of errors displayed in its **Inspector** tab:
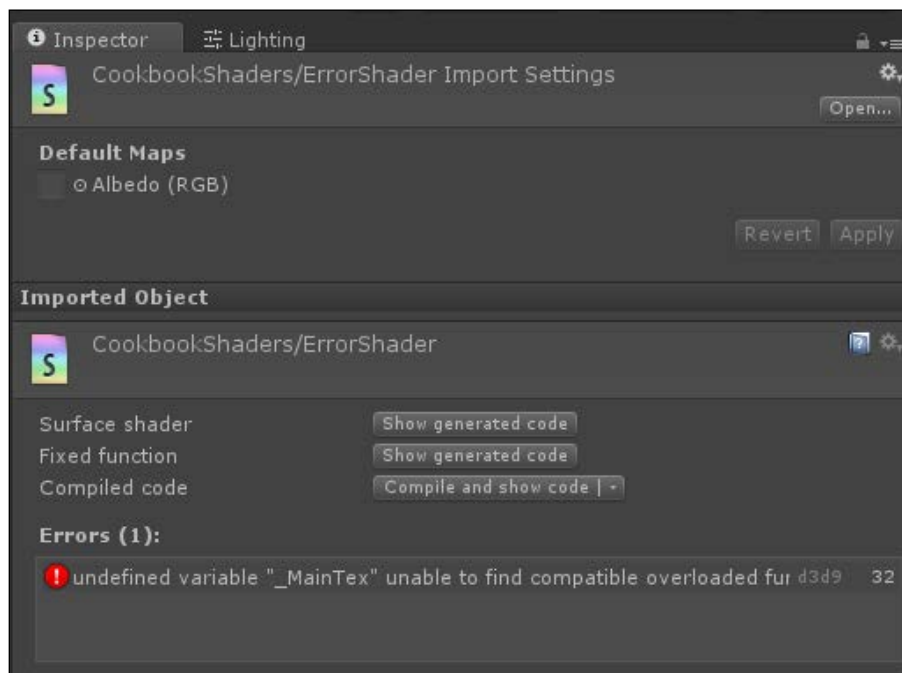
Despite showing the line that raised the error, it rarely means that this is the line that has to be fixed. The error message shown in the previous image is generated by deleting the `sampler2D _MainTex` variable from the `SubShader{}` block. However, the error is raised by the first line that tries to access such a variable.

Finding and fixing what's wrong with code is a process called **debugging**. The most common mistakes that you should check for are as follows:

- ▶ A missing bracket. If you forgot to add a curly bracket to close a section, the compiler is likely to raise errors at the end of the document, at the beginning, or a new section.
- ▶ A missing semicolon. One of the most common mistakes but luckily one of the easiest to spot and fix. Errors are often raised by the following line.
- ▶ A property that has been defined in the `Properties` section but has not been coupled with a variable in the `SubShader{}` block.
- ▶ Conversely to what you might be used to in C# scripts, floating point values in Cg do not need to the followed by an `f`: it's `1.0`, not `1.0f`.

The error messages raised by shaders can be very misleading, especially due to their strict syntactic constraints. If you are in doubt about their meaning, it's best to search the Internet. Unity forums are filled with other developers who are likely to have encountered (and fixed) your problem before.

See also

More information on how to master Surface Shaders and their properties can be found in *Chapter 2*, *Surface Shaders and Texture Mapping*. If you are curious to see what shaders can actually do when used at their full potential, have a look at *Chapter 10*, *Advanced Shading Techniques*, for some of the most advanced techniques covered in this book.

2
Surface Shaders and Texture Mapping

In this chapter, we will explore Surface Shaders. We will start from a very simple matte material and end with holographic projections and advanced terrains blending. We can also use textures to animate, blend, and drive any other property that we want. In this chapter, you will learn about the following methods:

- ▶ Diffuse shading
- ▶ Using packed arrays
- ▶ Adding a texture to a shader
- ▶ Scrolling textures by modifying UV values
- ▶ Normal mapping
- ▶ Creating a transparent material
- ▶ Creating a Holographic Shader
- ▶ Packing and blending textures
- ▶ Creating a circle around your terrain

Introduction

Surface Shaders have been introduced in *Chapter 1*, *Creating Your First Shader*, as the main type of shader used in Unity. This chapter will show in detail what these actually are and how they work. Generally speaking, there are two essential steps in every Surface Shader. First, you have to specify certain physical properties of the material that you want to describe, such as its diffuse color, smoothness, and transparency. These properties are initialized in a function called **surface function** and stored in a structure called **surface output**. Secondly, the surface output is passed to a **lighting model**. This is a special function that will also take information about the nearby lights in the scene. Both these parameters are then used to calculate the final color for each pixel of your model. The lighting function is where the real calculations of a shader take place as it's the piece of code that determines how light should behave when it touches a material.

The following diagram loosely summarizes how a Surface Shader works. Custom lighting models will be explored in *Chapter 3*, *Understanding Lighting Models*, while *Chapter 5*, *Vertex Functions*, will focus on vertex modifiers:
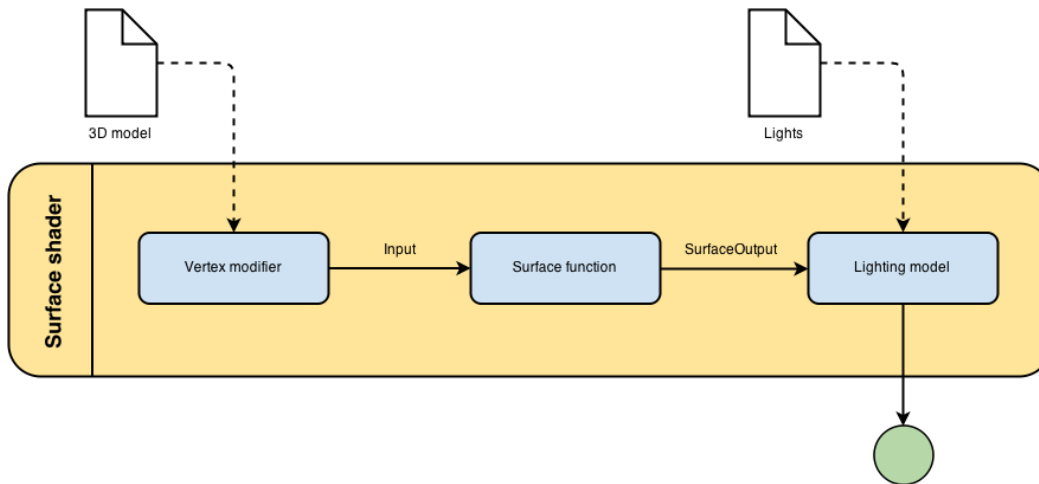
Diffuse shading

Before starting our journey into texture mapping, it is important to understand how diffuse materials work. Certain objects might have a uniform color and smooth surface, but not smooth enough to shine on reflected light. These matte materials are best represented with a Diffuse shader. While in the real world, pure diffuse materials do not exist; Diffuse shaders are relatively cheap to implement and find a large application in games with low-poly aesthetics.

Getting ready

There are several ways in which you can create your own Diffuse shader. A quick way is to start with the Standard Shader in Unity 5 and edit it to remove any texture, similarly to what was previously done in *Chapter 1*, *Creating Your First Shader*.

How to do it...

Let's start with our Standard Shader, and apply the following changes:

1. Remove all the properties except `_Color`:

   ```
   _Color ("Color", Color) = (1,1,1,1)
   ```

2. From the `SubShader{}` section, remove the `_MainTex`, `_Glossiness`, and `_Metallic` variables. You should not remove the reference to `uv_MainTex` as Cg does not allow the `Input` struct to be empty. The value will be simply ignored.

3. Remove the content of the `surf()` function and replace it with the following:

   ```
   o.Albedo = _Color.rgb;
   ```

4. Your shader should look as follows:

   ```
   Shader "CookbookShaders/Diffuse" {
     Properties {
       _Color ("Color", Color) = (1,1,1,1)
     }
     SubShader {
       Tags { "RenderType"="Opaque" }
       LOD 200

       CGPROGRAM
       #pragma surface surf Standard fullforwardshadows
       #pragma target 3.0
       struct Input {
         float2 uv_MainTex;
       };
   ```

```
        fixed4 _Color;

        void surf (Input IN, inout SurfaceOutputStandard o) {
          o.Albedo = _Color.rgb;
        }
        ENDCG
    }
    FallBack "Diffuse"
}
```

As this shader has been refitted from a Standard Shader, it will use physically-based rendering to simulate how light behaves on your models. If you are trying to achieve a non-photorealistic look, you can change the first `#pragma` directive so that it uses `Lambert` rather than `Standard`. If you do so, you should also replace `SurfaceOutputStandard` with `SurfaceOutput`.

How it works...

The way shaders allow you to communicate the rendering properties of your material to their lighting model is via a surface output. It is basically a wrapper around all the parameters that the current lighting model needs. It should not surprise you that different lighting models have different surface output structs. The following table shows the three main output structs used in Unity 5 and how they can be used:

Type of shaders	Unity 4	Unity 5
Diffuse	Any Surface Shader	**Standard**
	`SurfaceOutput`	`SurfaceOutputStandard`
Specular	Any Surface Shader	**Standard (Specular setup)**
	`SurfaceOutput`	`SurfaceOutputStandardSpecular`

The `SurfaceOutput` struct has the following properties:

▸ `fixed3 Albedo;`: This is the diffuse color of the material

▸ `fixed3 Normal;`: This is the tangent space normal, if written

▸ `fixed3 Emission;`: This is the color of the light emitted by the material (this property is declared as `half3` in the Standard Shaders)

▸ `fixed Alpha;`: This is the transparency of the material

▸ `half Specular;`: This is the specular power from `0` to `1`

▸ `fixed Gloss;`: This is the specular intensity

The `SurfaceOutputStandard` struct has the following properties:

- `fixed3 Albedo;`: This is the base color of the material (whether it's diffuse or specular)
- `fixed3 Normal;`
- `half3 Emission;`: This property is declared as `half3`, while it was defined as `fixed3` in `SurfaceOutput`
- `fixed Alpha;`
- `half Occlusion;`: This is the occlusion (default `1`)
- `half Smoothness;`: This is the smoothness (`0` = rough, `1` = smooth)
- `half Metallic;`: `0` = non-metal, `1`= metal

The `SurfaceOutputStandardSpecular` struct has the following properties:

- `fixed3 Albedo;`.
- `fixed3 Normal;`.
- `half3 Emission;`.
- `fixed Alpha;`.
- `half Occlusion;`.
- `half Smoothness;`.
- `fixed3 Specular;`: This is the specular color. This is very different from the `Specular` property in `SurfaceOutput` as it allows specifying a color rather than a single value.

Using a Surface Shader correctly is a matter of initializing the surface output with the correct values.

Using packed arrays

Loosely speaking, the code inside a shader has to be executed for at least every pixel in your screen. This is the reason why GPUs are highly optimized for parallel computing. This philosophy is also evident in the standard type of variables and operators available in Cg. Understanding them is essential not just to use shaders correctly, but also to write highly optimized ones.

How to do it...

There are two types of variables in Cg: single values and packed arrays. The latter can be identified because their type ends with a number such as `float3` or `int4`. As their names suggest, these types of variables are similar to structs, which means that they each contain several single values. Cg calls them **packed arrays**, though they are not exactly *arrays* in the traditional sense.

The elements of a packed array can be accessed as a normal struct. They are typically called `x`, `y`, `z`, and `w`. However, Cg also provides you with another alias for them, that is, `r`, `g`, `b`, and `a`. Despite there being no difference between using `x` or `r`, it can make a huge difference for the readers. Shader coding, in fact, often involves calculation with positions and colors. You might have seen this in the Standard Shaders:

```
o.Alpha = _Color.a;
```

Here, `o` was a struct and `_Color` was a packed array. This is also why Cg prohibits the mixed usage of these two syntaxes: you cannot use `_Color.xgz`.

There is also another important feature of packed arrays that has no equivalent in C#: **swizzling**. Cg allows addressing and reordering elements within packed arrays in just a single line. Once again, this appears in the Standard Shader:

```
o.Albedo = _Color.rgb;
```

`Albedo` is `fixed3`, which means that it contains three values of the `fixed` type. However, `_Color` is defined as `fixed4`. A direct assignment would result in a compiler error as `_Color` is bigger than `Albedo`. The C# way of doing this would be as follows:

```
o.Albedo.r = _Color.r;
o.Albedo.g = _Color.g;
o.Albedo.b = _Color.b;
```

However, it can be compressed in Cg:

```
o.Albedo = _Color.rgb;
```

Cg also allows reordering elements, for instance, using `_Color.bgr` to swap the red and blue channels.

Lastly, when a single value is assigned to a packed array, it is copied to all of its fields:

```
o.Albedo = 0; // Black =(0,0,0)
o.Albedo = 1; // White =(1,1,1)
```

This is referred to as **smearing**.

Swizzling can also be used on the left-hand side of an expression, allowing only certain components of a packed array to be overwritten:

```
o.Albedo.rg = _Color.rg;
```

In which case, it is called **masking**.

Packed matrices

Where swizzling really shows its full potential is when applied to packed matrices. Cg allows types such as `float4x4`, which represents a matrix of floats with four rows and four columns. You can access a single element of the matrix using the *_mRC* notation, where *R* is the row and *C* is the column:

```
float4x4 matrix;
// ...
float first = matrix._m00;
float last = matrix._m33;
```

The *_mRC* notation can also be chained:

```
float4 diagonal = matrix._m00_m11_m22_m33;
```

An entire row can be selected using squared brackets:

```
float4 firstRow = matrix[0];
// Equivalent to
float4 firstRow = matrix._m00_m01_m02_m03;
```

See also

Packed arrays are one of the nicest features of Cg. You can discover more about them here:

```
http://http.developer.nvidia.com/CgTutorial/cg_tutorial_chapter02.html
```

Adding a texture to a shader

Textures can bring our shaders to life very quickly in terms of achieving very realistic effects. In order to effectively use textures, we need to understand how a 2D image is mapped to a 3D model. This process is called **texture mapping**, and it requires some work to be done on the shader and 3D model that we want to use. Models, in fact, are made out of triangles; each vertex can store data that shaders can access. One of the most important information stored in vertices is the **UV data**. It consists of two coordinates, *U* and *V*, ranging from 0 to 1. They represent the *XY* position of the pixel in the 2D image that will be mapped to the vertices. UV data is present only for vertices; when the inner points of a triangle have to be texture-mapped, the GPU interpolates the closest UV values to find the right pixel in the texture to be used. The following image shows you how a 2D texture is mapped to a triangle from a 3D model:

The UV data is stored in the 3D model and requires a modeling software to be edited. Some models lack the UV component, hence they cannot support texture mapping. The Stanford bunny, for example, was not originally provided with one.

Getting ready

For this recipe, you'll need a 3D model with UV data and its texture. They both need to be imported to Unity before starting. You can do this simply by dragging them to the editor. As the Standard Shader supports texture mapping by default, we'll use this and then explain in detail how it works.

How to do it...

Adding a texture to your model using the Standard Shader is incredibly simple, as follows:

1. Create a new Standard Shader called `TexturedShader`.
2. Create a new material called `TexturedMaterial`.
3. Assign the shader to the material by dragging over it.
4. After selecting the material, drag your texture to the empty rectangle called **Albedo (RGB)**. If you have followed all these steps correctly, your material **Inspector** tab should look like this:
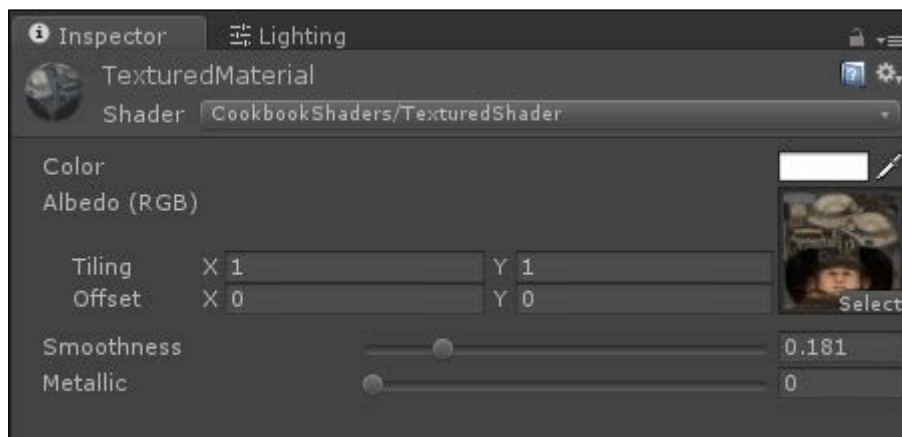
The Standard Shader knows how to map a 2D image to a 3D model using its UV data.

How it works...

When the Standard Shader is used from the inspector of a material, the process behind texture mapping is completely transparent to developers. If we want to understand how it works, it's necessary to take a closer look at `TexturedShader`. From the `Properties` section, we can see that the `Albedo (RGB)` texture is actually referred to in the code as `_MainTex`:

```
_MainTex ("Albedo (RGB)", 2D) = "white" {}
```

In the `CGPROGRAM` section, this texture is defined as `sampler2D`, the standard type for 2D textures:

```
sampler2D _MainTex;
```

The next line shows a struct called `Input`. This is the input parameter for the surface function and contains a packed array called `uv_MainTex`:

```
struct Input {
  float2 uv_MainTex;
};
```

Every time the `surf()` surface function is called, the `Input` structure will contain the UV of `_MainTex` for the specific point of the 3D model that needs to be rendered. The Standard Shader recognizes that the name `uv_MainTex` refers to `_MainTex` and initializes it automatically. If you are interested in understanding how the UV is actually mapped from a 3D space to a 2D texture, you can check *Chapter 3*, *Understanding Lighting Models*.

Finally, the UV data is used to sample the texture in the first line of the surface function:

```
fixed4 c = tex2D (_MainTex, IN.uv_MainTex) * _Color;
```

This is done using the `tex2D()` function of Cg; it takes a texture and UV and returns the color of the pixel at that position.

> The *U* and *V* coordinates go from 0 to 1, where (0,0) and (1,1) correspond to two opposite corners. Different implementations associate UV with different corners; if your texture happens to appear reversed, try inverting the V component.

There's more...

When you import a texture to Unity, you are setting up some of the properties that `sampler2D` will use. The most important is the **Filter** mode, which determines how colors are interpolated when the texture is sampled. It is very unlikely that the UV data will point exactly to the center of a pixel; in all the other cases, you might want to interpolate between the closest pixels to get a more uniform color. The following is the screenshot of the **Inspector** tab of an example texture:
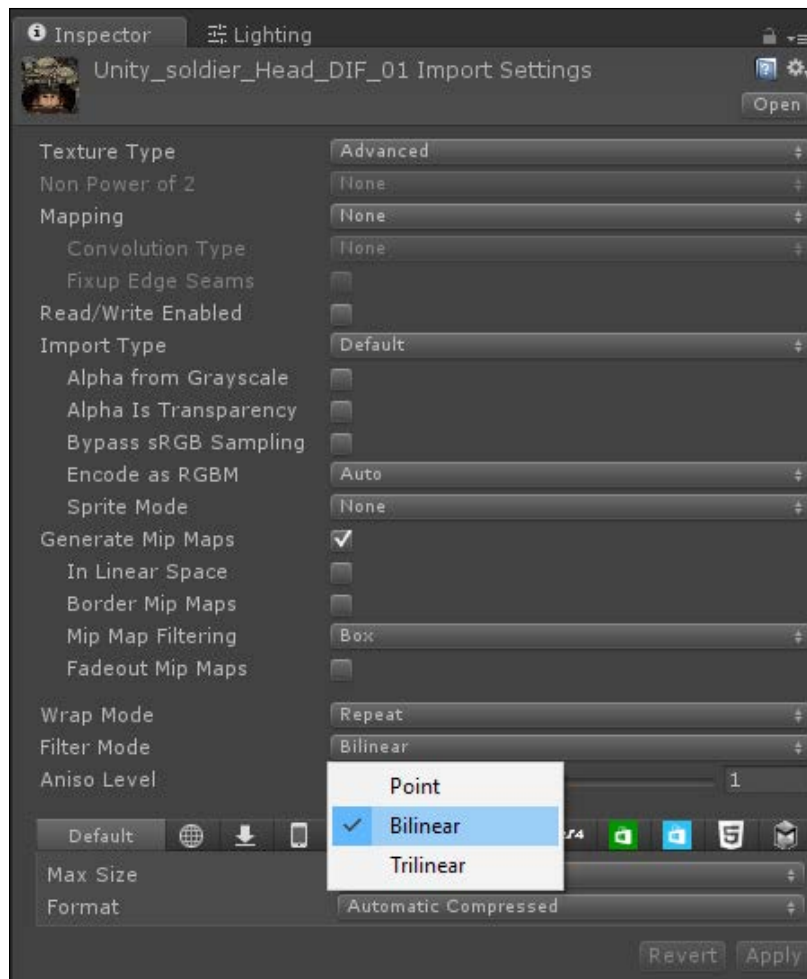
For most applications, **Bilinear** provides an inexpensive yet effective way to smooth the texture. If you are creating a 2D game, however, **Bilinear** might produce blurred tiles. In this case, you can use **Point** to remove any interpolation from the texture sampling.

When a texture is seen from a steep angle, texture sampling is likely to produce visually unpleasant artifacts. You can reduce them by setting **Aniso Level** to a higher value. This is particular useful for floor and ceiling textures, where glitches can break the illusion of continuity.

See also

If you would like to know more about the inner working of how textures are mapped to a 3D surface, you can read the information available at `http://http.developer.nvidia.com/CgTutorial/cg_tutorial_chapter03.html`.

For a complete list of the options available when importing a 2D texture, you can refer to the following website:

```
http://docs.unity3d.com/Manual/class-TextureImporter.html
```

Scrolling textures by modifying UV values

One of the most common texture techniques used in today's game industry is the process of allowing you to scroll the textures over the surface of an object. This allows you to create effects such as waterfalls, rivers, lava flows, and so on. It's also a technique that is the basis to create animated sprite effects, but we will cover this in a subsequent recipe of this chapter. Let's first see how we will create a simple scrolling effect in a Surface Shader.

Getting ready

To begin this recipe, you will need to create a new shader file and material. This will set us up with a nice clean shader that we can use to study the scrolling effect by itself.

How to do it...

To begin with, we will launch our new shader file that we just created and enter the code mentioned in the following steps:

1. The shader will need two new properties that will allow us to control the speed of the texture scrolling. So, let's add a speed property for the *X* direction and a speed property for the *Y* direction, as shown in the following code:

```
Properties {
  _MainTint ("Diffuse Tint", Color) = (1,1,1,1)
  _MainTex ("Base (RGB)", 2D) = "white" {}
  _ScrollXSpeed ("X Scroll Speed", "Range(0,10)) = 2
  _ScrollYSpeed ("Y Scroll Speed", "Range(0,10)) = 2
}
```

2. Modify the Cg properties in the `CGPROGRAM` section and create new variables so that we can access the values from our properties:

```
fixed4 _MainTint;
fixed _ScrollXSpeed;
fixed _ScrollYSpeed;
sampler2D _MainTex;
```

3. Modify the surface function to change the UVs given to the `tex2D()` function. Then, use the built-in `_Time` variable to animate the UVs over time when the play button is pressed in the editor:

```
void surf (Input IN, inout SurfaceOutput o)
```

```
  {
    // Create a separate variable to store our UVs
    // before we pass them to the tex2D() function
    fixed2 scrolledUV = IN.uv_MainTex;

    // Create variables that store the individual x and y
    // components for the UV's scaled by time
    fixed xScrollValue = _ScrollXSpeed * _Time;
    fixed yScrollValue = _ScrollYSpeed * _Time;

    // Apply the final UV offset
    scrolledUV += fixed2(xScrollValue, yScrollValue);

    // Apply textures and tint
    half4 c = tex2D (_MainTex, scrolledUV);
    o.Albedo = c.rgb * _MainTint;
    o.Alpha = c.a;
  }
```

The following image demonstrates the result of utilizing the scrolling UV system to create a simple river motion for your environments. You can notice this effect in the scene called `ScrollingUVs` from the code files provided with this book:

How it works...

The scrolling system starts with the declaration of a couple of properties, which will allow the user of this shader to increase or decrease the speed of the scrolling effect itself. At their core, they are float values being passed from the material's **Inspector** tab to the surface function of the shader. For more information on shader properties, see *Chapter 1*, *Creating Your First Shader*.

Once we have these float values from the material's **Inspector** tab, we can use them to offset our UV values in the shader.

To begin this process, we first store the UVs in a separate variable called `scrolledUV`. This variable has to be `float2/fixed2` because the UV values are being passed to us from the `Input` structure:

```
struct Input
{
    float2 uv_MainTex;
}
```

Once we have access to the mesh's UVs, we can offset them using our scroll speed variables and built-in `_Time` variable. This built-in variable returns a variable of the `float4` type, meaning that each component of this variable contains different values of time as it pertains to game time.

A complete description of these individual time values are described at the following link: `http://docs.unity3d.com/Manual/SL-UnityShaderVariables.html`

This `_Time` variable will give us an incremented float value based on Unity's game time clock. So, we can use this value to move our UVs in a UV direction and scale that time with our scroll speed variables:

```
// Create variables that store the individual x and y
// components for the uv's scaled by time
fixed xScrollValue = _ScrollXSpeed * _Time;
fixed yScrollValue = _ScrollYSpeed * _Time;
```

With the correct offset being calculated by time, we can add the new offset value back to the original UV position. This is why we are using the `+=` operator in the next line. We want to take the original UV position, add the new offset value, and then pass this to the `tex2D()` function as the texture's new UVs. This creates the effect of the texture moving on the surface. We are really manipulating the UVs, so we are faking the effect of the texture moving:

```
scrolledUV += fixed2(xScrollValue, yScrollValue);
half4 c = tex2D (_MainTex, scrolledUV);
```

Normal mapping

Every triangle of a 3D model has a *facing direction*, which is the direction that it is pointing toward. It is often represented with an arrow placed in the center of the triangle and orthogonal to the surface. The facing direction plays an important role in the way light reflects on a surface. If two adjacent triangles face different directions, they will reflect lights at different angles, hence they'll be shaded differently. For curved objects, this is a problem: it is obvious that the geometry is made out of flat triangles.

To avoid this problem, the way the light reflects on a triangle doesn't take into account its facing direction, but its *normal direction* instead. As stated in *Adding a texture to a shader* recipe, vertices can store data; the normal direction is the most used information after the UV data. This is a vector of unit length that indicates the direction faced by the vertex. Regardless of the facing direction, every point within a triangle has its own normal direction that is a linear interpolation of the ones stored in its vertices. This gives us the ability to fake the effect of high-resolution geometry on a low-resolution model. The following image shows the same geometric shape rendered with different per-vertex normals. In the image on the left, normals are orthogonal to the face represented by its vertices; this indicates that there is a clear separation between each face. On the right, normals are interpolated along the surface, indicating that even if the surface is rough, light should reflect as if it's smooth. It's easy to see that even if the three objects in the following image share the same geometry, they reflect light differently. Despite being made out of flat triangles, the object on the right reflects light as if its surface was actually curved:

Smooth objects with rough edges are a clear indication that per-vertex normals have been interpolated. This can be seen if we draw the direction of the normal stored in every vertex, as shown in the following image. You should note that every triangle has only three normals, but as multiple triangles can share the same vertex, more than one line can come out of it:

Calculating the normals from the 3D model is a technique that has rapidly declined in favor of a more advanced one—normal mapping. Similar to what happens with texture mapping, the normal directions can be provided using an additional texture, usually called normal map or bump map. Normal maps are usually RGB images, where the RGB components are used to indicate the X, Y, and Z components of the normal direction. There are many ways to create normal maps these days. Some applications such as **CrazyBump** (`http://www.crazybump.com/`) and **NDO Painter** (`http://quixel.se/ndo/`) will take in 2D data and convert it to normal data for you. Other applications such as **Zbrush 4R7** (`http://www.pixologic.com/`) and **AUTODESK** (`http://usa.autodesk.com`) will take 3D sculpted data and create normal maps for you. The actual process of creating normal maps is definitely out of the scope of this book, but the links in the previous text should help you get started.

Unity makes the process of adding normals to your shaders quite an easy process in the Surface Shader realm using the `UnpackNormals()` function. Let's see how this is done.

Getting ready

Create a new material and shader and set them up on a new object in the **Scene** view. This will give us a clean workspace in which we can look at just the normal mapping technique.

You will need a normal map for this recipe, but there is also one in the Unity project included with this book.

An example normal map included with this book's contents is shown here:

How to do it...

The following are the steps to create a normal map shader:

1. Let's get the `Properties` block set up in order to have a color tint and texture:

```
Properties
{
  _MainTint ("Diffuse Tint", Color) = (1,1,1,1)
  _NormalTex ("Normal Map", 2D) = "bump" {}
}
```

> By initializing the texture as `bump`, we are telling Unity that `_NormalTex` will contain a normal map. If the texture is not set, it will be replaced by a grey texture. The color used `(0.5,0.5,0.5,1)` indicates no bump at all.

2. Link the properties to the Cg program by declaring them in `SubShader{}` below the `CGPROGRAM` statement:

```
CPROGRAM
#pragma surface surf Lambert

// Link the property to the CG program
sampler2D _NormalTex;
float4 _MainTint;
```

3. We need to make sure that we update the `Input` struct with the proper variable name so that we can use the model's UVs for the normal map texture:

```
// Make sure you get the UVs for the texture in the struct
struct Input
{
  float2 uv_NormalTex;
}
```

4. Finally, we extract the normal information from the normal map texture using the built-in `UnpackNormal()` function. Then, you only have to apply these new normals to the output of the Surface Shader:

```
// Get the normal data out of the normal map texture
// using the UnpackNormal function
float3 normalMap = UnpackNormal(tex2D(_NormalTex,
  IN.uv_NormalTex));

// Apply the new normal to the lighting model
o.Normal = normalMap.rgb;
```

The following image demonstrates the result of our normal map shader:

> Shaders can have both a texture map and normal map. It is not uncommon to use the same UV data to address both. However, it is possible to provide a secondary set of UVs in the vertex data (UV2) specifically used for the normal map.

How it works...

The actual math to perform the normal mapping effect is definitely beyond the scope of this chapter, but Unity has done it all for us already. It has created the functions for us so that we don't have to keep doing it over and over again. This is another reason why Surface Shaders are a really efficient way to write shaders.

If you look in the `UnityCG.cginc` file found in the `Data` folder in your Unity installation directory, you will find the definitions for the `UnpackNormal()` function. When you declare this function in your Surface Shader, Unity takes the provided normal map and processes it for you and gives you the correct type of data so that you can use it in your per-pixel lighting function. It's a huge time-saver! When sampling a texture, you get RGB values from 0 to 1; however, the directions of a normal vector range from -1 to +1. `UnpackNormal()` brings these components in the right range.

Once you have processed the normal map with the `UnpackNormal()` function, you send it back to your `SurfaceOutput` struct so that it can be used in the lighting function. This is done by `o.Normal = normalMap.rgb;`. We will see how the normal is actually used to calculate the final color of each pixel in *Chapter 3*, *Understanding Lighting Models*.

There's more...

You can also add some controls to your normal map shader that lets a user adjust the intensity of the normal map. This is easily done by modifying the x and y components of the normal map variable and then adding it all back together. Add another property to the `Properties` block and name it `_NormalMapIntensity`:

```
_NormalMapIntensity("Normal intensity", Range(0,1)) = 1
```

Multiply the x and y components of the unpacked normal map and reapply this value to the normal map variable:

```
fixed3 n = UnpackNormal(tex2D(_BumpTex, IN.uv_ uv_MainTex)).rgb;
n.x *= _NormalMapIntensity;
n.y *= _NormalMapIntensity;
o.Normal = normalize(n);
```

> Normal vectors are supposed to have lengths equal to one. Multiplying them for `_NormalMapIntensity` changes their length, making normalization necessary.

Now, you can let a user adjust the intensity of the normal map in the material's **Inspector** tab. The following image shows the result of modifying the normal map with our scalar values:

Creating a transparent material

All the shaders seen so far have something in common—they are used for solid materials. If you want to improve the look of your game, transparent materials are often a good way to start. They can be used for anything from a fire effect to a window glass. Working with them, unfortunately, is slightly more complicated. Before rendering solid models, Unity orders them according to the distance from the camera (**Z ordering**) and skips all the triangles that are facing away from the camera (**culling**). When rendering transparent geometries, there are instances in which these two aspects can cause problems. This recipe will show you how to solve some of these issues when it comes to creating a transparent Surface Shader. This topic will be heavily revisited in *Chapter 6*, *Fragment Shaders and Grab Passes*, where realistic glass and water shaders will be provided.

Getting ready

This recipe requires a new shader, which we'll be calling `Transparent`, and a new material so that it can be attached to an object. As this is going to be a transparent glass window, a quad or plane is perfect. We will also need several other non-transparent objects to test the effect. In this example, we will use a PNG for the glass texture. The alpha channel of the image will be used to determine the transparency of the glass. The process of creating such an image depends on the software that you are using. However, these are the main steps that you will need to follow:

1. Find the image of the glass you want for your windows.
2. Open it with a photoediting software, such as **GIMP** or **Photoshop**.
3. Select the parts of the image that you want to be semi-transparent.
4. Create a white (full opacity) layer mask on your image.
5. Use the selection previously made to fill the layer mask with a darker color.
6. Save the image and import it to Unity.

The toy image used in this recipe is a picture of a stained glass from the *Meaux Cathedral* in *France* (`https://en.wikipedia.org/wiki/Stained_glass`). If you have followed all the steps, your image should look like this (**RGB** channels on the left, and **A** channel on the right):

How to do it...

As mentioned previously, there are a few aspects that we need to take care of while using a Transparent Shader:

1. In the `SubShader{}` section of the shader, add the following tags that signal the shader is transparent:

```
Tags
{
  "Queue" = "Transparent"
  "IgnoreProjector" = "True"
  "RenderType" = "Transparent"
}
```

2. As this shader is designed for 2D materials, make sure that the back geometry of your model is not drawn by adding the following:

```
Cull Back
```

3. Tell the shader that this material is transparent and needs to be blended with what was drawn on the screen before:

```
#pragma surface surf Standard alpha:fade
```

4. Use this Surface Shader to determine the final color and transparency of the glass:

```
void surf(Input IN, inout SurfaceOutputStandard o)
{
    float4 c = tex2D(_MainTex, IN.uv_MainTex) * _Color;
    o.Albedo = c.rgb;
    o.Alpha = c.a;
}
```

How it works...

This shader introduces several new concepts. First of all, `Tags` are used to add information about how the object is going to be rendered. The really interesting one here is `Queue`. Unity, by default, will sort your objects for you based on the distance from the camera. So, as an object gets nearer to the camera, it is going to be drawn over all the objects that are further away from the camera. For most cases, this works out just fine for games, but you will find certain situations where you will want to have more control over the sorting of your objects in your scene. Unity has provided us with some default render queues, each with a unique value that directs Unity when to draw the object to the screen. These built-in render queues are called `Background`, `Geometry`, `AlphaTest`, `Transparent`, and `Overlay`. These queues weren't just created arbitrarily; they actually serve a purpose to make our lives easier when writing shaders and interacting with the real-time renderer. Refer to the following table for descriptions on the usage of each of these individual render queues:

Render queue	Render queue description	Render queue value
Background	This render queue is rendered first. It is used for skyboxes and so on.	1000
Geometry	This is the default render queue. This is used for most objects. Opaque geometry uses this queue.	2000
AlphaTest	Alpha-tested geometry uses this queue. It's different from the `Geometry` queue as it's more efficient to render alpha-tested objects after all the solid objects are drawn.	2450
Transparent	This render queue is rendered after `Geometry` and `AlphaTest` queues in back-to-front order. Anything alpha-blended (that is, shaders that don't write to the depth buffer) should go here, for example, glass and particle effects.	3000
Overlay	This render queue is meant for overlay effects. Anything rendered last should go here, for example, lens flares.	4000

So, once you know which render queue your object belongs to, you can assign its built-in render queue tag. Our shader used the `Transparent` queue, so we wrote `Tags{"Queue"="Trasparent"}`.

> The fact that the `Transparent` queue is rendered after `Geometry` does *not* mean that our glass will appear on top of all the other solid objects. Unity will draw the glass last, but it will not render pixels that belong to pieces of geometry hidden behind something else. This control is done using a technique called **ZBuffering**. More information on how models are rendered can be found at `http://docs.unity3d.com/Manual/SL-CullAndDepth.html`.

The `IgnoreProjector` tag makes this object unaffected by Unity's projectors. Lastly, `RenderType` plays a role in **shader replacement**, a topic that will be covered briefly in *Chapter 9*, *Gameplay and Screen Effects*.

The last concept introduced is `alpha:fade`. This indicates that all the pixels from this material have to be blended with what was on the screen before according to their alpha values. Without this directive, the pixels will be drawn in the correct order, but they won't have any transparency.

Creating a Holographic Shader

More and more space-themed games are being released every year. An important part of a good sci-fi game is the way futuristic technology is presented and integrated in the gameplay. There's nothing that screams futuristic more than holograms. Despite being present in many flavors, holograms are often represented as semi-transparent, thin projections of an object. This recipe shows you how to create a shader that simulates such effects. Take this as a starting point: you can add noise, animated scanlines, and vibrations to create a truly outstanding holographic effect. The following image shows an example of a holographic effect:

Getting ready

As the holographic effects shows only the outlines of an object, we'll call this shader `Silhouette`. Attach it to a material and assign it to your 3D model.

How to do it...

The following changes will modify our existing shader into a holographic one:

1. Add the following property to the shader:

   ```
   _DotProduct("Rim effect", Range(-1,1)) = 0.25
   ```

2. Add its respective variable to the `CGPROGRAM` section:

   ```
   float _DotProduct;
   ```

3. As this material is transparent, add the following tags:

   ```
   Tags
   {
     "Queue" = "Transparent"
     "IgnoreProjector" = "True"
     "RenderType" = "Transparent"
   }
   ```

   > According to the type of object that you will use, you might want its backside to appear. If this is the case, add `Cull Off` so that the back of the model won't be removed (*culled*).

4. This shader is not trying to simulate a realistic material, so there is no need to use the PBR lighting model. The **Lambertian reflectance**, which is very cheap, is used instead. Additionally, we should disable any lighting with `nolighting` and signal to Cg that this is a Transparent Shader using `alpha:fade`:

   ```
   #pragma surface surf Lambert alpha:fade nolighting
   ```

5. Change the `Input` structure so that Unity will fill it with the current view direction and world normal direction:

   ```
   struct Input
   {
     float2 uv_MainTex;
     float3 worldNormal;
     float3 viewDir;
   };
   ```

6. Use the following surface function. Remember that as this shader is using the Lambertian reflectance as its lighting function, the name of the surface output structure should be changed accordingly to `SurfaeOutput` instead of `SurfaceOutputStandard`:

```
void surf(Input IN, inout SurfaceOutput o)
{
    float4 c = tex2D(_MainTex, IN.uv_MainTex) * _Color;
    o.Albedo = c.rgb;

    float border = 1 - (abs(dot(IN.viewDir,
        IN.worldNormal)));
    float alpha = (border * (1 - _DotProduct) + _DotProduct);
    o.Alpha = c.a * alpha;
}
```

You can now use the **Rim effect** slider to choose the strength of the holographic effect.

How it works...

As mentioned before, this shader works by showing only the silhouette of an object. If we look at the object from another angle, its outline will change. Geometrically speaking, the edges of a model are all those triangles whose *normal direction* is orthogonal (90 degrees) to the current *view direction*. The `Input` structure declares these parameters, `worldNormal` and `viewDir`, respectively.

The problem of understanding when two vectors are orthogonal can be solved using the **dot product**. It's an operator that takes two vectors and returns zero if they are orthogonal. We use `_DotProduct` to determine how close to zero the dot product has to be for the triangle to fade completely.

The second aspect that is used in this shader is the gentle fading between the edge of the model (fully visible) and the angle determined by `_DotProduct` (invisible). This linear interpolation is done as follows:

```
float alpha = (border * (1 - _DotProduct) + _DotProduct);
```

Finally, the original alpha from the texture is multiplied with the newly calculated coefficient to achieve the final look.

There's more...

This technique is very simple and relatively inexpensive. Yet, it can be used for a large variety of effects, such as the following:

- ▸ The slightly colored atmosphere of a planet in sci-fi games
- ▸ The edge of an object that has been selected or is currently under the mouse
- ▸ A ghost or specter
- ▸ Smoke coming out of an engine
- ▸ The shockwave of an explosion
- ▸ The bubble shield of a spaceship under attack

See also

The dot product plays an important role in the way reflections are calculated. *Chapter 3*, *Understanding Lighting Models*, will explain in detail how it works and why it is widely used in so many shaders.

Packing and blending textures

Textures are useful to store not only loads of data, not just pixel colors as we generally tend to think of them, but also for multiple sets of pixels in both the *x* and *y* directions and RGBA channels. We can actually pack multiple images into one single RGBA texture and use each of the R, G, B, and A components as individual textures themselves by extracting each of these components in the shader code.

The result of packing individual grayscale images into a single RGBA texture can be seen in the following image:

| R Channel | G Channel | B Channel | Alpha Channel | Packed Result |

Why is this helpful? Well, in terms of the amount of actual memory that your application takes up, textures are a large portion of your application's size. So, to begin reducing the size of your application, we can look at all of the images that we are using in our shader and see if we can merge these textures into a single texture.

Any texture that is grayscale can be packed into one of the RGBA channels of another texture. This might sound a bit odd at first, but this recipe is going to demonstrate one of the uses of packing a texture and using these packed textures in a shader.

One example of using these packed textures is when you want to blend a set of textures together onto a single surface. You see this most often in terrain type shaders, where you need to blend into another texture nicely using some sort of control texture or the packed texture, in this case. This recipe covers this technique and shows you how you can construct the beginnings of a nice four-texture blended terrain shader.

Getting ready

Let's create a new shader file in your `Shaders` folder and then create a new material for this shader. The naming convention is entirely up to you for your shader and material files, so try your best to keep them organized and easy to reference later on.

Once you have your shader and material ready, create a new scene in which we can test our shader.

You will also need to gather up four textures that you would want to blend together. These can be anything, but for a nice terrain shader, you will want grass, dirt, rocky dirt, and rock textures.

These are the color textures that we will be using for this recipe, which are included with this book.

Finally, we will also need a blending texture that is packed with grayscale images. This will give us the four blending textures that we can use to direct how the color textures will be placed on the object surface.

We can use very intricate blending textures to create a very realistic distribution of terrain textures over a terrain mesh, as seen in the following image:

R Channel   G Channel   B Channel   Alpha Channel   Packed Result

How to do it...

Let's learn how to use packed textures by entering the code shown in the following steps:

1. We need to add a few properties to our `Properties` block. We will need five `sampler2D` objects, or textures, and two color properties:

```
Properties
{
_MainTint ("Diffuse Tint", Color) = (1,1,1,1)

//Add the properties below so we can input all of our
  textures
  _ColorA ("Terrain Color A", Color) = (1,1,1,1)
  _ColorB ("Terrain Color B", Color) = (1,1,1,1)
  _RTexture ("Red Channel Texture", 2D) = ""{}
  _GTexture ("Green Channel Texture", 2D) = ""{}
  _BTexture ("Blue Channel Texture", 2D) = ""{}
  _ATexture ("Alpha Channel Texture", 2D) = ""{}
  _BlendTex ("Blend Texture", 2D) = ""{}
}
```

2. We then need to create the `SubShader{}` section variables that will be our link to the data in the `Properties` block:

```
CGPROGRAM
#pragma surface surf Lambert

float4 _MainTint;
float4 _ColorA;
float4 _ColorB;
sampler2D _RTexture;
sampler2D _GTexture;
sampler2D _BTexture;
sampler2D _BlendTex;
sampler2D _ATexture;
```

3. So, now we have our texture properties and we are passing them to our `SubShader{}` function. In order to allow the user to change the tiling rates on a per-texture basis, we will need to modify our `Input` struct. This will allow us to use the tiling and offset parameters on each texture:

```
struct Input
{
  float2 uv_RTexture;
  float2 uv_GTexture;
  float2 uv_BTexture;
```

```
    float2 uv_ATexture;
    float2 uv_BlendTex;
};
```

4. In the `surf()` function, get the texture information and store them in their own variables so that we can work with the data in a clean, easy-to-understand way:

```
//Get the pixel data from the blend texture
//we need a float 4 here because the texture
//will return R,G,B,and A or X,Y,Z, and W
float4 blendData = tex2D(_BlendTex, IN.uv_BlendTex);

//Get the data from the textures we want to blend
float4 rTexData = tex2D(_RTexture, IN.uv_RTexture);
float4 gTexData = tex2D(_GTexture, IN.uv_GTexture);
float4 bTexData = tex2D(_BTexture, IN.uv_BTexture);
float4 aTexData = tex2D(_ATexture, IN.uv_ATexture);
```

5. Let's blend each of our textures together using the `lerp()` function. It takes three arguments, `lerp(value : a, value : b, blend: c)`. The `lerp()` function takes in two textures and blends them with the float value given in the last argument:

```
//No we need to contruct a new RGBA value and add all
//the different blended texture back together
float4 finalColor;
finalColor = lerp(rTexData, gTexData, blendData.g);
finalColor = lerp(finalColor, bTexData, blendData.b);
finalColor = lerp(finalColor, aTexData,
  blendData.a);finalColor.a = 1.0;
```

6. Finally, we multiply our blended textures with the color tint values and use the red channel to determine where the two different terrain tint colors go:

```
//Add on our terrain tinting colors
float4 terrainLayers = lerp(_ColorA, _ColorB, blendData.r);
finalColor *= terrainLayers;
finalColor = saturate(finalColor);

o.Albedo = finalColor.rgb * _MainTint.rgb;
o.Alpha = finalColor.a;
```

The result of blending together four terrain textures and creating a terrain tinting technique can be seen in the following image:

How it works...

This might seem like quite a few lines of code, but the concept behind blending is actually quite simple. For the technique to work, we have to employ the built-in `lerp()` function from the CgFX standard library. This function allows us to pick a value between argument one and argument two using argument three as the blend amount:

Function	Description
`lerp( a , b, f )`	This involves linear interpolation: `(1 – f )* a + b * f` Here, `a` and `b` are matching vector or scalar types. The `f` parameter can be either a scalar or vector of the same type as `a` and `b`.

So, for example, if we wanted to find the mid-value between 1 and 2, we could feed the value 0.5 as the third argument to the `lerp()` function and it would return the value 1.5. This works perfectly for our blending needs as the values of an individual channel in an RGBA texture are single float values, usually in the range of 0 to 1.

In the shader, we simply take one of the channels from our blend texture and use it to drive the color that is picked in a `lerp()` function for each pixel. For instance, we take our grass texture and dirt texture, use the red channel from our blending texture, and feed this to a `lerp()` function. This will give us the correct blended color result for each pixel on the surface.

A more visual representation of what is happening when using the `lerp()` function is shown in the following image:

The shader code simply uses the four channels of the blend texture and all the color textures to create a final blended texture. This final texture then becomes our color that we can multiply with our diffuse lighting.

Creating a circle around your terrain

Many RTS games display distances (range attack, moving distance, sight, and so on) by drawing a circle around the selected unit. If the terrain is flat, this can be done simply by stretching a quad with the texture of a circle. If that's not the case, the quad will most likely be clipped behind a hill or another piece of geometry. This recipe will show you how to create a shader that allows you to draw circles around an object of arbitrary complexity. If you want to be able to move or animate your circle, we will need both a shader and C# script. The following image shows an example of drawing a circle in a hilly region using a shader:

Getting ready

Despite working with every piece of geometry, this technique is oriented to terrains. Hence, the first step is setting up a terrain in Unity.

1. Let's start by creating a new shader called `RadiusShader` and the respective material, `Radius`.

2. Have the character for your object ready; we will draw a circle around it.

3. From the menu, navigate to **GameObject** | **3D Object** | **Terrain** to create a new terrain.

4. Create the geometry for your terrain. You can either import an existing one or draw your own using the tools available (**Raise/Lower Terrain**, **Paint Height**, **Smooth Height**).

5. Terrains are special objects in Unity, and the way texture mapping works on them is different from traditional 3D models. You cannot provide `_MainTex` from a shader as it needs to be provided directly from the terrain itself. To do this, select **Paint Texture** and then click on **Add Texture...**:
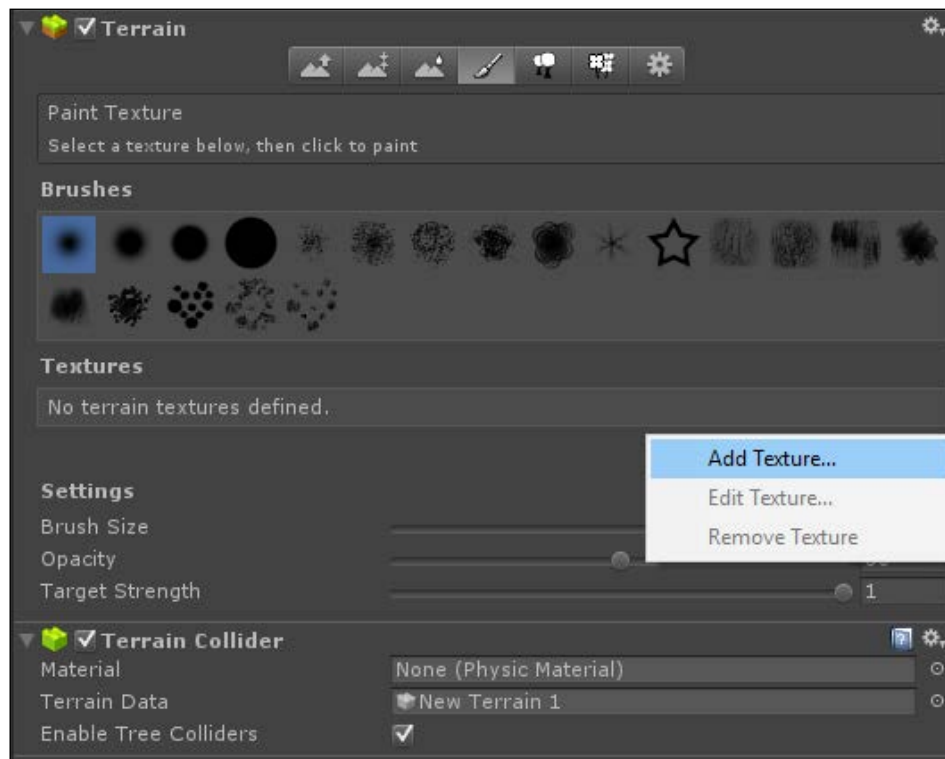
6. Now that the texture is set, you have to change the material of the terrain so that a custom shader can be provided. From **Terrain Settings**, change the **Material** property to **Custom**, and then drag the **Radius** material to the **Custom Material** box.

You are now ready to create your shader.

How to do it...

Let's start by editing the `RadiusShader` file:

1. In the new shader, add these four properties:

```
_Center("Center", Vector) = (0,0,0,0)
_Radius("Radius", Float) = 0.5
_RadiusColor("Radius Color", Color) = (1,0,0,1)
_RadiusWidth("Radius Width", Float) = 2
```

2. Add their respective variables to the `CGPROGRAM` section:

```
float3 _Center;
float _Radius;
fixed4 _RadiusColor;
float _RadiusWidth;
```

3. `Input` to our surface function requires not only the UV of the texture, but also the position (in world coordinates) of every point of the terrain. We can retrieve this parameter by changing the `Input` struct as follows:

```
struct Input
{
  float2 uv_MainTex; // The UV of the terrain texture
  float3 worldPos;   // The in-world position
};
```

4. Lastly, we use this surface function:

```
void surf(Input IN, inout SurfaceOutputStandard o)
{
  float d = distance(_Center, IN.worldPos);
  if (d > _Radius && d < _Radius + _RadiusWidth)
    o.Albedo = _RadiusColor;
  else
    o.Albedo = tex2D(_MainTex, IN.uv_MainTex).rgb;
}
```

These steps are all it takes to draw a circle on your terrain. You can use the material's **Inspector** tab to change the position, radius, and color of the circle.

Moving the circle

If you want the circle to follow your character, other steps are necessary:

1. Create a new C# script called `Radius`.

2. Add these properties to the script:

   ```
   public Material radiusMaterial;
   public float radius = 1;
   public Color color = Color.white;
   ```

3. In the `Update()` method, add these lines of code:

   ```
   radiusMaterial.SetVector("_Center", transform.position);
   radiusMaterial.SetFloat("_Radius", radius);
   radiusMaterial.SetColor("_RadiusColor", color);
   ```

4. Attach the script to your character.

5. Finally, drag the `Radius` material to the **Radius Material** slot of the script.

You can now move your character around and this will create a nice circle around it. Changing the properties of the `Radius` script will change the radius as well.

How it works...

The relevant parameters to draw a circle are its center, radius, and color. They are all available in the shader with the names `_Center`, `_Radius`, and `_RadiusColor`. By adding the `worldPos` variable to the `Input` structure, we are asking Unity to provide us with the position of the pixel that we are drawing expressed in world coordinates. This is the actual position of an object in the editor.

The `surf()` function is where the circle is actually drawn. It calculates the distance from the point being drawn and center of the radius, then it checks whether it is between `_Radius` and `_Radius + _RadiusWidth`; if this is the case, it uses the chosen color. In the other case, it just samples the texture map like all the other shaders seen so far.

3
Understanding Lighting Models

In the previous chapters, we introduced Surface Shaders and explained how we can change physical properties (such as Albedo and Specular) to simulate different materials. How does this really work? At the heart of every Surface Shader, there is its **lighting model**. It's the function that takes these properties and calculates the final shade of each pixel. Unity usually hides this from the developers because in order to write a lighting model, you have to understand how light reflects and refracts onto surfaces. This chapter will finally show you how lighting models work and give you the basics to create your own.

In this chapter, you will learn the following recipes:

- ▶ Creating a custom diffuse lighting model
- ▶ Creating a Toon Shader
- ▶ Creating a Phong Specular type
- ▶ Creating a BlinnPhong Specular type
- ▶ Creating an Anisotropic Specular type

Introduction

Simulating the way light works is a very challenging and resource-consuming task. For many years, video games have used very simple lighting models that, despite lacking realism, were very believable. Even if most 3D engines are now using physically-based renderers, it is worth exploring some simpler techniques. The ones presented in this chapter are reasonably realistic and widely adopted on devices with low resources such as mobile phones. Understanding these simple lighting models is also essential if you want to create your own one.

Creating a custom diffuse lighting model

If you are familiar with Unity 4, you may know that the default shader it provided was based on a lighting model called Lambertian reflectance. This recipe will show you how it is possible to create a shader with a custom lighting model and explain the mathematics and implementation behind it. The following image shows the same geometry rendered with a Standard Shader (right) and diffuse Lambert one (left):

Shaders based on the Lambertian reflectance are classified as non-photorealistic; no object in the real world really looks like this. However, Lambert Shaders are still often used in low poly games as they produce a neat contrast between the faces of complex geometries. The lighting model used to calculate the Lambertian reflectance is also very efficient, making it perfect for mobile games.

Unity has already provided us with a lighting function that we can use for our shaders. It is called the Lambertian lighting model. It is one of the more basic and efficient forms of reflectance, which you can find in a lot of games even today. As it is already built in the Unity Surface Shader language, we thought it is best to start with this first and build on it. You can also find an example in the Unity reference manual, but we will go into more depth with it and explain where the data is coming from and why it is working the way it is. This will help you get a nice grounding in setting up custom lighting models so that we can build on this knowledge in the future recipes in this chapter.

Getting ready

Let's start by carrying out the following steps:

1. Create a new shader and give it a name.
2. Create a new material, give it a name, and assign the new shader to its shader property.
3. Then, create a sphere object and place it roughly in the center of the scene.
4. Finally, let's create a directional light to cast some light on our object.

When your assets have been set up in Unity, you should have a scene that resembles the following screenshot:

How to do it...

The Lambertian reflectance can be achieved with the following changes to the shader:

1. Begin by adding the following properties to the shader's `Properties` block:

    ```
    _MainTex("Texture", 2D) = "white"
    ```

2. Change the `#pragma` directive of the shader so that, instead of `Standard`, it uses our custom lighting model:

    ```
    #pragma surface surf SimpleLambert
    ```

3. Use a very simple surface function, which just samples the texture according to its UV data:

    ```
    void surf(Input IN, inout SurfaceOutput o) {
      o.Albedo = tex2D(_MainTex, IN.uv_MainTex).rgb;
    }
    ```

4. Add a function called `LightingSimpleLambert()` that will contain the following code for the Lambertian reflectance:

    ```
    half4 LightingSimpleLambert (SurfaceOutput s, half3
      lightDir, half atten) {
      half NdotL = dot (s.Normal, lightDir);
      half4 c;
      c.rgb = s.Albedo * _LightColor0.rgb * (NdotL * atten *
        1);
      c.a = s.Alpha;
      return c;
    }
    ```

How it works...

As previously seen in *Chapter 1*, *Creating Your First Shader*, the `#pragma` directive is used to specify which surface function to use. Choosing a different lighting model works in a similar fashion: `SimpleLambert` forces Cg to look for a function called `LightingSimpleLambert()`. Note `Lighting` at the beginning, which is omitted in the directive.

The lighting function takes three parameters: the *surface output* (which contains the physical properties such as the albedo and transparency), the *direction* the light is coming from, and its *attenuation*.

According to the Lambertian reflectance, the amount of light a surface reflects depends on the angle between the incident light and surface normal. If you have played pool billiards, you are surely familiar with this concept; the direction of a ball depends on its incident angle against the wall. If you hit a wall at a 90 degree angle, the ball will come back at you; if you hit it with a very low angle, its direction will be mostly unchanged. The Lambertian model makes the same assumption; if the light hits a triangle with a 90 degree angle, all the light gets reflected back. The lower the angle, the less light is reflected back to you. This concept is shown in the following image:


This simple concept has to be translated into a mathematical form. In vector algebra, the angle between two unit vectors can be calculated via an operator called **dot product**. When the dot product is equal to zero, two vectors are orthogonal, which means that they make a 90 degree angle. When it is equal to one (or minus one), they are parallel to each other. Cg has a function called `dot()`, which implements the dot product extremely efficiently.

The following picture shows a light source (sun) shining on a complex surface. **L** indicates the light direction (called `lightDir` in the shader) and **N** is the normal to the surface. The light is reflected with the same angle that it hits the surface:

The Lambertian reflectance simply uses the `NdotL` dot product as a multiplicative coefficient for the intensity of light:

$$I = N \cdot L$$

When *N* and *L* are parallel, all the light is reflected back to the source, causing the geometry to appear brighter. The `_LightColor0` variable contains the color of the light that is calculated.

> Prior to Unity 5, the intensity of the lights were different. If you are using an old Diffuse shader based on the Lambertian model, you may notice that `NdotL` was multiplied by two: `(NdotL * atten * 2)` rather than `(NdotL * atten)`. If you are importing a custom shader from Unity 4, you will need to correct this manually. Legacy Shaders, however, have already been designed taking this aspect into account.

When the dot product is negative, the light is coming from the opposite side of the triangle. This is not a problem for opaque geometries as triangles that are not facing the camera frontally are *culled* (discarded) and not rendered.

This basic Lambert is a great starting point when you are prototyping your shaders as you can get a lot accomplished in terms of writing the core functionality of the shader while not having to worry about the basic lighting functions.

Unity has provided us with a lighting model that has already taken the task of creating a Lambert lighting for you. If you look at the `UnityCG.cginc` file found in your Unity's installation directory under the `Data` folder, you will notice that you have Lambert and BlinnPhong lighting models available for you to use. The moment you compile your shader with `#pragma surface surf Lambert`, you are telling the shader to utilize Unity's implementation of the Lambert lighting function in the `UnityCG.cginc` file so that we don't have to write that code over and over again. We will explore how the BlinnPhong model works later in this chapter.

Creating a Toon Shader

One of the most used effects in games is the **toon shading**, which is also known as **cel shading** (short for *celluloid*). It is a non-photorealistic rendering technique that makes 3D models appear flat. Many games use it to give the illusion that the graphics are being hand-drawn rather than being 3D-modeled. You can see, in the following picture, a sphere rendered with a Standard Shader (right) and Toon Shader (left):

Achieving this effect using just surface functions is not impossible, but it would be extremely expensive and time-consuming. The surface function, in fact, only works on the properties of the material, not its actual lighting condition. As toon shading requires to change the way light reflects, we need to create our custom lighting model instead.

Getting ready

Let's start this recipe by creating a shader and its material and importing a special texture, as follows:

1. Start by creating a new shader; in this example, we will extend the one made in the previous recipe.

2. Create a new material for the shader and attach it to a 3D model. The toon shading works best on curved surfaces.

3. This recipe requires an additional texture called **ramp map**. It is important that you change its **Wrap Mode** to **Clamp**. If you want the edges between the colors to be sharp, the **Filter Mode** should also be set to **Point**:

How to do it...

The toon aesthetic can be achieved with the following changes to the shader:

1. Add a new property for a texture called `_RampTex`:

   ```
   _RampTex ("Ramp", 2D) = "white" {}
   ```

2. Add its relative variable in the `CGPROGRAM` section:

   ```
   sampler2D _RampTex;
   ```

3. Change the `#pragma` directive so that it points to a function called `LightingToon()`:

    ```
    #pragma surface surf Toon
    ```

4. Use this lighting model:

    ```
    fixed4 LightingToon (SurfaceOutput s, fixed3 lightDir,
      fixed atten)
    {
      half NdotL = dot(s.Normal, lightDir);
      NdotL = tex2D(_RampTex, fixed2(NdotL, 0.5));

      fixed4 c;
      c.rgb = s.Albedo * _LightColor0.rgb * NdotL * atten;
      c.a = s.Alpha;

      return c;
    }
    ```

How it works...

The main characteristic of the toon shading is the way the light is rendered; surfaces are not shaded uniformly. To achieve this effect, we need a ramp map. Its purpose is to remap the Lambertian light intensity `NdotL` to another value. Using a ramp map without a gradient, we can force the lighting to be rendered in steps. The following image shows how the ramp map is used to correct the light intensity:

There's more...

There are many different ways one can achieve a toon shading effect. Using different ramps can produce dramatic changes in the way your models look, so you should experiment in order to find the best one.

An alternative to ramp textures is to *snap* the light intensity `NdotL` so that it can only assume a certain number of values equidistantly sampled from 0 to 1:

```
half4 LightingCustomLambert (SurfaceOutput s, half3 lightDir,
  half3 viewDir, half atten) {
half NdotL = dot (s.Normal, lightDir);
half cel = floor(NdotL * _CelShadingLevels) / (_CelShadingLevels
  -0.5); // Snap

half4 c;
  c.rgb = s.Albedo * _LightColor0.rgb * cel * atten;
  c.a = s.Alpha;

  return c;
}
```

The snapping code multiplies `NdotL` times `_CelShadingLevels`, rounds it to an integer, and then divides it back. By doing this, the `cel` quantity is forced to assume one of the `_CelShadingLevels` equidistant values from 0 to 1. This removes the need for a ramp texture and makes all the color steps of the same size. If you are going for this implementation, remember to add a property called `_CelShadingLevels` to your shader.

Creating a Phong Specular type

The specularity of an object surface simply describes how shiny it is. These types of effects are often referred to as view-dependent effects in the shader world. This is because in order to achieve a realistic Specular effect in your shaders, you need to include the direction that the camera or user is facing the object's surface. The most basic and performance-friendly Specular type is the **Phong Specular** effect. It is the calculation of the light direction reflecting off of the surface compared to the user's view direction. It is a very common Specular model used in many applications, from games to movies. While it isn't the most realistic in terms of accurately modeling the reflected Specular, it gives a great approximation that performs well in most situations. Additionally, if your object is further away from the camera and there is no need for a very accurate Specular, this is a great way to provide a Specular effect to your shaders.

In this recipe, we will be covering how to implement the per-vertex version of the shader and also the per-pixel version using some new parameters in the Surface Shader's `Input` struct. We will see the difference and discuss when and why to use these two different implementations for different situations.

Getting ready

To start with this recipe, perform the following steps:

1. Create a new shader, material, and object, and give them appropriate names so that you can find them later.

2. Attach the shader to the material and the material to the object. To finish off your new scene, create a new directional light so that we can see our Specular effect as we code it.

How to do it...

Follow the following steps to create a Phong lighting model:

1. You might be seeing a pattern at this point, but we always like to start out with our most basic part of the shader writing process: the creation of properties. So, let's add the following properties to the shader:

```
Properties
{
  _MainTint ("Diffuse Tint", Color) = (1,1,1,1)
  _MainTex ("Base (RGB)", 2D) = "white" {}
  _SpecularColor ("Specular Color", Color) = (1,1,1,1)
  _SpecPower ("Specular Power", Range(0,30)) = 1
}
```

2. We then have to make sure to add the corresponding variables to our CGPROGRAM block in our SubShader{} block:

```
float4 _SpecularColor;
sampler2D _MainTex;
float4 _MainTint;
float _SpecPower;
```

3. Now we have to add our custom lighting model so that we can compute our own Phong Specular. Don't worry if it doesn't make sense at this point; we will cover each line of code in the *How it works...* section. Add the following code to the shader's SubShader{} function:

```
fixed4 LightingPhong (SurfaceOutput s, fixed3 lightDir,
  half3 viewDir, fixed atten)
{
  // Reflection
  float NdotL = dot(s.Normal, lightDir);
  float3 reflectionVector = normalize(2.0 * s.Normal *
    NdotL - lightDir);
```

```
    // Specular
    float spec = pow(max(0, dot(reflectionVector, viewDir)),
      _SpecPower);
    float3 finalSpec = _SpecularColor.rgb * spec;

    // Final effect
    fixed4 c;
    c.rgb = (s.Albedo * _LightColor0.rgb * max(0,NdotL) *
      atten) + (_LightColor0.rgb * finalSpec);
    c.a = s.Alpha;
    return c;
}
```

4. Finally, we have to tell the `CGPROGRAM` block that it needs to use our custom lighting function instead of one of the built-in ones. We do this by changing the `#pragma` statement to the following:

```
CPROGRAM
#pragma surface surf Phong
```

The following screenshot demonstrates the result of our custom Phong lighting model using our own custom reflection vector:

How it works...

Let's break down the lighting function by itself, as the rest of the shader should be pretty familiar to you at this point.

In the previous recipes, we have used a lighting function that provided only the light direction, `lightDir`. Unity comes with a set of lighting functions that you can use, including one that provides the view direction, `viewDir`. Refer to the following table or go to `http://docs.unity3d.com/Documentation/Components/SL-SurfaceShaderLighting.html`:

Not view-dependent	`half4 Lighting Name You choose (SurfaceOutput s, half3 lightDir, half atten);`
View-dependent	`half4 Lighting Name You choose (SurfaceOutput s, half3 lightDir, half3 viewDir, half atten);`

In our case, we are doing a Specular shader, so we need to have the view-dependent lighting function structure. So, we have to write the following:

```
CPROGRAM
#pragma surface surf Pong
fixed4 LightingPhong (SurfaceOutput s, fixed3 lightDir, half3 viewDir,
fixed atten)
{
  // ...
}
```

This will tell the shader that we want to create our own view-dependent shader. Always make sure that your lighting function name is the same in your lighting function declaration and the `#pragma` statement, or Unity will not be able to find your lighting model.

The components that play a role in the Phong model are described in the following image. We have the light direction **L** (coupled with its perfect reflection **R**) and normal direction **N**. They have all been encountered before in the Lambertian model, with the exception of **V**, which is the *view direction*:
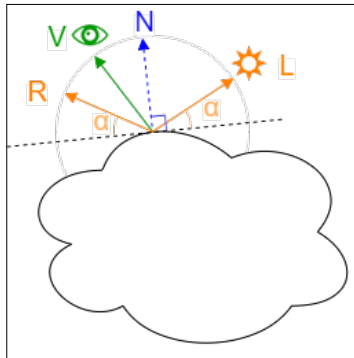
The Phong model assumes that the final light intensity of a reflective surface is given by two components: its diffuse color and Specular value, as follows:

$$I = D + S$$

The diffuse component *D* remains unchanged from the Lambertian model:

$$D = N \cdot L$$

The Specular component *S* is defined as follows:

$$S = (R \cdot V)^p$$

Here, *p* is the Specular power defined as `_SpecPower` in the shader. The only unknown parameter is *R*, which is the reflection of *L* according to *N*. In vector algebra, this can be calculated as follows:

$$R = 2N \cdot (N \cdot L) - L$$

This is exactly what is calculated in the following:

```
float3 reflectionVector = normalize(2.0 * s.Normal * NdotL -
    lightDir);
```

This has the effect of bending the normal towards the light; as a vertex normal is pointing away from the light, it is forced to look at the light. Refer to the following screenshot for a more visual representation. The script that produces this debug effect is included in the book's support page at `https://www.packtpub.com/books/content/support`:

The following screenshot displays the final result of our Phong Specular calculation isolated in the shader:

Creating a BlinnPhong Specular type

Blinn is another more efficient way of calculating and estimating specularity. It is done by getting the half vector from the view direction and light direction. It was brought into the world of Cg by Jim Blinn. He found that it was much more efficient to just get the half vector instead of calculating our own reflection vectors. It cut down on the code and processing time. If you actually look at the built-in BlinnPhong lighting model included in the `UnityCG.cginc` file, you will notice that it is using the half vector as well, hence it is named **BlinnPhong**. It is just a simpler version of the full Phong calculation.

Getting ready

To start with this recipe, perform the following steps:

1. This time, instead of creating a whole new scene, let's just use the objects and scene that we have, and create a new shader and material and name them `BlinnPhong`.

2. Once you have a new shader, double-click on it to launch MonoDevelop so that we can start editing our shader.

How to do it...

Perform the following steps to create a BlinnPhong lighting model:

1. First, we need to add our own properties to the `Properties` block so that we can control the look of the Specular highlight:

```
Properties
{
  _MainTint ("Diffuse Tint", Color) = (1,1,1,1)
  _MainTex ("Base (RGB)", 2D) = "white" {}
  _SpecularColor ("Specular Color", Color) = (1,1,1,1)
  _SpecPower ("Specular Power", Range(0.1,60)) = 3
}
```

2. Then, we need to make sure that we have created the corresponding variables in our `CGPROGRAM` block so that we can access the data from our `Properties` block, in our subshader:

```
sampler2D _MainTex;
float4 _MainTint;
float4 _SpecularColor;
float _SpecPower;
```

3. Now it's time to create our custom lighting model that will process our Diffuse and Specular calculations. The code is as follows:

```
fixed4 LightingCustomBlinnPhong (SurfaceOutput s, fixed3
  lightDir, half3 viewDir, fixed atten)
{
  float NdotL = max(0,dot(s.Normal, lightDir));

  float3 halfVector = normalize(lightDir + viewDir);
  float NdotH = max(0, dot(s.Normal, halfVector));
  float spec = pow(NdotH, _SpecPower) * _SpecularColor;

  float4 c;
```

```
      c.rgb = (s.Albedo * _LightColor0.rgb * NdotL) +
        (_LightColor0.rgb * _SpecularColor.rgb * spec) * atten;
      c.a = s.Alpha;
      return c;
    }
```

4. To complete our shader, we will need to tell our CGPROGRAM block to use our custom lighting model rather than a built-in one by modifying the #pragma statement with the following code:

```
CPROGRAM
#pragma surface surf CustomBlinnPhong
```

The following screenshot demonstrates the results of our BlinnPhong lighting model:

How it works...

The BlinnPhong Specular is almost exactly like the Phong Specular, except that it is more efficient because it uses less code to achieve almost the same effect. Before the introduction of physically-based rendering, this approach was the default choice for Specular reflection in Unity 4.

Calculating the reflection vector **R** is generally expensive. The BlinnPhong Specular replaces it with the half vector **H** between the view direction **V** and light direction **L**:

Instead of calculating our own reflection vector, we are simply going to get the vector halfway between the view direction and light direction, basically simulating the reflection vector. It has actually been found that this approach is more physically accurate than the last approach, but we thought it necessary to show you all the possibilities:

$$S_{Phong} = (R \cdot V)^p, \ S_{BlinnPhong} = (N \cdot H)^p$$

According to vector algebra, the half vector can be calculated as follows:

$$H \frac{V+L}{|V+L|}$$

Here, $|V+L|$ is the length of the vector $V+L$. In Cg, we simply need to add the view direction and light direction together and then normalize the result to a unity vector:

```
float3 halfVector = normalize(lightDir + viewDir);
```

Then, we simply need to dot the vertex normal with this new half vector to get our main Specular value. After this, we just take it to a power of `_SpecPower` and multiply it by the Specular color variable. It's much lighter on the code and math, but still gives us a nice Specular highlight that will work for a lot of real-time situations.

See also

The light models seen in this chapter are extremely simple; no real material is perfectly matte or perfectly specular. Moreover, it is not uncommon for complex materials such as clothing, wood, and skin to require knowledge of how light scatters in the layers beneath the surface.

Use the following table to recap the different lighting models encountered so far:

Technique	Type	Unity 5 shader	Light Intensity (I)
Lambertian	Diffuse	**Legacy Shaders | Diffuse**	$I = N \cdot L$
Phong	Specular		$I = N \cdot L + \left(R \cdot V \right)^p$ $R = 2N \cdot \left(N \cdot L \right) - L$
BlinnPhong	Specular	**Legacy Shaders | Specular**	$I = N \cdot L + \left(N \cdot H \right)^p$ $H = \dfrac{V + L}{\left| V + L \right|}$

There are other interesting models such as the Oren-Nayar lighting model for rough surfaces: `https://en.wikipedia.org/wiki/Oren%E2%80%93Nayar_reflectance_model`

Creating an Anisotropic Specular type

Anisotropic is a type of Specular or reflection that simulates the directionality of grooves in a surface and modifies/stretches the Specular in the perpendicular direction. It is very useful when you want to simulate brushed metals, not a metal with a clear, smooth, and polished surface. Imagine the Specular that you see when you look at the data side of a CD or DVD or the way Specular is shaped at the bottom of a pot or pan. You will notice that if you carefully examine the surface, you will see that there is a direction to the grooves in the surface, usually in the way the metal was brushed. When you apply a Specular to this surface, you get a Specular stretched in the perpendicular direction.

This recipe will introduce you to the concept of augmenting your Specular highlights to achieve different types of brushed surfaces. In future recipes, we will look at ways in which we can use the concepts of this recipe to achieve other effects such as stretched reflections and hair, but here, you are going to learn the fundamentals of the technique first. We will be using this shader as a reference for our own custom Anisotropic Shader:

`http://wiki.unity3d.com/index.php?title=Anisotropic_Highlight_Shader`

The following screenshot shows examples of different types of Specular effects one can achieve using Anisotropic Shaders in Unity:

Getting ready

Let's start this recipe by creating a shader, its material, and some lights for our scene:

1. Create a new scene with some objects and lights so that we can visually debug our shader.

2. Then create a new shader and material, and hook them up to our objects.

3. Lastly, we will need some sort of normal map that will indicate the directionality of our Anisotropic Specular highlight.

The following screenshot shows the Anisotropy normal map we will be using for this recipe. It is available from the book's support page at `https://www.packtpub.com/books/content/support`:

How to do it...

To create an Anisotropic effect, we need to make the following changes to the shader previously created:

1. We first need to add the properties that we are going to need for our shader. These will allow a lot of artistic control over the final appearance of the surface:

```
Properties
{
  _MainTint ("Diffuse Tint", Color) = (1,1,1,1)
  _MainTex ("Base (RGB)", 2D) = "white" {}
  _SpecularColor ("specular Color", Color) = (1,1,1,1)
  _Specular ("Specular Amount", Range(0,1)) = 0.5
  _SpecPower ("Specular Power", Range(0,1)) = 0.5
  _AnisoDir ("Anisotropic Direction", 2D) = "" {}
  _AnisoOffset ("Anisotropic Offset", Range(-1,1)) = -0.2
}
```

2. We then need to make the connection between our `Properties` block and our `SubShader{}` block so that we can use the data being provided by the `Properties` block:

```
sampler2D _MainTex;
sampler2D _AnisoDir;
float4 _MainTint;
float4 _SpecularColor;
float _AnisoOffset;
float _Specular;
float _SpecPower;
```

3. Now we can create our lighting function that will produce the correct Anisotropic effect on our surface. We will use the following code for this:

```
fixed4 LightingAnisotropic(SurfaceAnisoOutput s, fixed3
  lightDir, half3 viewDir, fixed atten)
{
  fixed3 halfVector = normalize(normalize(lightDir) +
    normalize(viewDir));
  float NdotL = saturate(dot(s.Normal, lightDir));

  fixed HdotA = dot(normalize(s.Normal + s.AnisoDirection),
    halfVector);
  float aniso = max(0, sin(radians((HdotA + _AnisoOffset) *
    180)));
```

```
float spec = saturate(pow(aniso, s.Gloss * 128) *
  s.Specular);

fixed4 c;
c.rgb = ((s.Albedo * _LightColor0.rgb * NdotL) +
  (_LightColor0.rgb * _SpecularColor.rgb * spec)) *
  atten;
c.a = s.Alpha;
return c;
}
```

4. In order to use this new lighting function, we need to tell the subshader's `#pragma` statement to look for it instead of using one of the built-in lighting functions. We are also telling the shader to target shader model 3.0 so that we can have more space for textures in our program:

```
CGPROGRAM
#pragma surface surf Anisotropic
#pragma target 3.0
```

5. We have also given the Anisotropic normal map its own UVs by declaring the following code in the `Input` struct. This isn't entirely necessary as we could just use the UVs from the main texture, but this gives us independent control over the tiling of our brushed metal effect so that we can scale it to any size we want:

```
struct Input
{
  float2 uv_MainTex;
  float2 uv_AnisoDir;
};
```

6. We also need to add the `SurfaceAnisoOutput` struct:

```
struct SurfaceAnisoOutput
{
  fixed3 Albedo;
  fixed3 Normal;
  fixed3 Emission;
  fixed3 AnisoDirection;
  half Specular;
  fixed Gloss;
  fixed Alpha;
};
```

7. Finally, we need to use the `surf()` function to pass the correct data to our lighting function. So, we will get the per-pixel information from our Anisotropic normal map and set our Specular parameters as follows:

```
void surf(Input IN, inout SurfaceAnisoOutput o)
{
```

```
    half4 c = tex2D(_MainTex, IN.uv_MainTex) * _MainTint;
    float3 anisoTex = UnpackNormal(tex2D(_AnisoDir,
      IN.uv_AnisoDir));

    o.AnisoDirection = anisoTex;
    o.Specular = _Specular;
    o.Gloss = _SpecPower;
    o.Albedo = c.rgb;
    o.Alpha = c.a;
}
```

The Anisotropic normal map allows us to give the surface direction and helps us disperse the Specular highlight around the surface. The following screenshot demonstrates the result of our Anisotropic Shader:

How it works...

Let's break down this shader into its core components and explain why we are getting the effect. We will mostly be covering the custom lighting function here, as the rest of the shader should be pretty self-explanatory at this point.

We first start by declaring our own `SurfaceAnisoOutput` struct. We need to do this in order to get the per-pixel information from the Anisotropic normal map, and the only way we can do this in a Surface Shader is to use a `tex2D()` function in the `surf()` function. The following code shows the custom surface output structure used in our shader:

```
struct SurfaceAnisoOutput
{
  fixed3 Albedo;
  fixed3 Normal;
  fixed3 Emission;
  fixed3 AnisoDirection;
  half Specular;
  fixed Gloss;
  fixed Alpha;
};
```

We can use the `SurfaceAnisoOutput` struct as a way of interacting between the lighting function and surface function. In our case, we are storing the per-pixel texture information in the variable called `anisoTex` in our `surf()` function and then passing this data to the `SurfaceAnisoOutput` struct by storing it in the `AnisoDirection` variable. Once we have this, we can use the per-pixel information in the lighting function using `s.AnisoDirection`.

With this data connection set up, we can move on to our actual lighting calculations. This begins by getting the usual out of the way, the half vector, so that we don't have to do the full reflection calculation and diffuse lighting, which is the vertex normal dotted with the light vector or direction. This is done in Cg with the following lines:

```
fixed3 halfVector = normalize(normalize(lightDir) +
  normalize(viewDir));
float NdotL = saturate(dot(s.Normal, lightDir));
```

Then, we start the actual modification to the Specular to get the right look. We first dot the normalized sum of the vertex normal and per-pixel vectors from our Anisotropic normal map with `halfVector` calculated in the previous step. This gives us a float value that gives a value of `1` as the surface normal, which is modified by the Anisotropic normal map as it becomes parallel with `halfVector` and `0` as it is perpendicular. Finally, we modify this value with a `sin()` function so that we can basically get a darker middle highlight and ultimately a ring effect based off of `halfVector`. All the previously mentioned operations are summarized in the following two lines of Cg code:

```
fixed HdotA = dot(normalize(s.Normal + s.AnisoDirection),
  halfVector);
float aniso = max(0, sin(radians((HdotA + _AnisoOffset) * 180)));
```

Finally, we scale the effect of the `aniso` value by taking it to a power of `s.Gloss`, and then globally decrease its strength by multiplying it by `s.Specular`:

```
float spec = saturate(pow(aniso, s.Gloss * 128) * s.Specular);
```

This effect is great to create more advanced metal type surfaces, especially the ones that are brushed and seem to have directionality to them. It also works well for hair or any sort of soft surface with directionality to it. The following screenshot shows the result of displaying the final Anisotropic lighting calculation:

4
Physically Based Rendering in Unity 5

One of the biggest changes introduced in Unity 5 is **physically-based rendering**, which is also known as **PBR**. Previous chapters have repeatedly mentioned it without revealing too much about it. If you want to understand not only how PBR works, but how to make the most out of it, this is the chapter you should read.

In this chapter, you will learn the following recipes:

- ► Understanding the metallic setup
- ► Adding transparency to PBR
- ► Creating mirrors and reflective surfaces
- ► Baking lights in your scene

Introduction

All the lighting models encountered in *Chapter 3*, *Understanding Lighting Models*, were very primitive descriptions of how light behaves. The most important aspect during their making was *efficiency*. Real-time shading is expensive, and techniques such as Lambertian or BlinnPhong are a compromise between computational cost and realism. Having a more powerful **graphics processing unit** (**GPU**) has allowed us to write progressively more sophisticated lighting models and rendering engines, with the aim of simulating how light actually behaves. This is, in a nutshell, the philosophy behind PBR. As the name suggests, it tries to get as close as possible to the physics behind the processes that give a unique look to each material. Despite this, the term PBR has been widely used in marketing campaigns and is more of a synonym for *state-of-the-art rendering* rather than a well-defined technique. Unity 5 implements PBR by introducing two important changes. The first is a completely new lighting model (called **Standard**). Surface Shaders allow developers to specify the physical properties of a material, but they do not impose actual physical constraints on them. PBR fills this gap using a lighting model that enforces principles of physics such as *energy conservation* (an object cannot reflect more light than the amount it receives), *microsurface scattering* (rough surfaces reflect light more erratically compared to smooth ones), *Fresnel reflectance* (specular reflections appear at grazing angles), and *surface occlusion* (the darkening of corners and other geometries that are hard to light). All these aspects, and many others, are used to calculate the Standard lighting model. The second aspect that makes PBR so realistic is called **Global Illumination** (**GI**) and is the simulation of physically-based light transport. It means that objects are not drawn in the scene as if they were separate entities. They all contribute to the final rendering as light can reflect on them before hitting something else. This aspect is not captured in the shaders themselves but is an essential part of how the rendering engine works. Unfortunately, accurately simulating how light rays actually bounce over surfaces in real time is beyond the capabilities of modern GPUs. Unity 5 makes some clever optimizations that allow retaining visual fidelity without sacrificing performance. Some of the most advanced techniques (such as reflections), however, require the user input. All of these aspects will be covered in this chapter. It is important to remember that PBR and GI do not automatically guarantee that your game will be photorealistic. Achieving photorealism is a very challenging task and, like every art, it requires great expertize and exceptional skills.

Understanding the metallic setup

Unity 5 provides two different types of PBR shaders; they are referred to in the drop-down menu of the material's **Inspector** tab as **Standard** and **Standard (Specular setup)**. The main difference is that the former exposes the **Metallic** property, while the latter replaces it with **Specular**. Both these metallic and specular setups represent different ways in which one can initialize PBR materials. One of the concepts that has driven PBR is the ability to provide meaningful, physically-related properties that artists and developers can tweak and play with. The properties of some materials are easier to represent indicating how *metallic* they are, while for some, the other is more important in order to define how they reflect lights directly. If you have used Unity 4 in the past, **Standard (Specular setup)** might look more familiar to you. This recipe will show you how to use the **metallic setup** effectively. It's important to remember that the metallic workflow is not just for metallic materials; it is a way to define how materials will look according to how metallic or non-metallic their surface is. Despite being presented as two different types of shaders, both **Metallic** and **Specular** setups are generally equally expressive. As shown in the Unity documentation at `http://docs.unity3d.com/Manual/StandardShaderMetallicVsSpecular.html`, the same materials can usually be recreated with both setups (see the following image):

Getting ready

This recipe will use the Standard Shader provided in Unity 5, so there is no need to create a new one. The steps to start the recipe are as follows:

1. Create a new material.

2. From its **Inspector**, make sure that **Standard** is selected from its **Shader** drop-down menu.

You will also need a textured 3D model.

How to do it...

There are two main textures that need to be configured in the Standard Shader: **Albedo** and **Metallic**. To use the metallic workflow effectively, we need to initialize these maps correctly:

1. The **Albedo** map should be initialized with the unlit texture of the 3D model.

2. To create the **Metallic** map, start by duplicating the file for your **Albedo** map. You can do this by selecting the map from the **Project** tab and pressing *Ctrl + D*.

3. Use white (#ffffff) to color the regions of the map that correspond to materials that are made of pure metal. Use black (#000000) for all the other colors. Shades of grey should be used for dusty, weathered, or worn out metal surfaces, rust, scratched paint, and so on. As a matter of fact, Unity uses only the red channel to store the metallic value; the green and blue ones are ignored.

4. Use the alpha channel of the image to provide information about the **Smoothness** of the material.

5. Assign the **Metallic** map to the material. Both **Metallic** and **Smoothness** sliders will disappear as these two properties are now controlled by the map.

How it works...

Legacy Shaders allow artists to create materials that easily break the illusion of photorealism by having lighting conditions that are impossible in reality. This happens because all the properties of a material exposed in a Legacy Surface Shader are uncorrelated. By introducing the metallic workflow, Unity 5 imposes more constraints on the way objects look, making it harder for artists to create *illogical* materials.

Metals are known for the conducting of electricity; light is in the form of electromagnetic waves, meaning that almost all metals behave in a similar way compared to non-conductors (often referred as **insulators**). Conductors tend to reflect most photons (70-100%), resulting in high reflectance. The remaining light is absorbed, rather than diffused, suggesting that conductors have a very dark diffuse component. Insulators, conversely, have a low reflectance (4%); the rest of the light is scattered on the surface, contributing to their diffused looks.

In the Standard Shader, purely metallic materials have dark diffuse components and the color of their specular reflections is determined by the **Albedo** map. Conversely, the diffuse component of purely non-metallic materials is determined by the **Albedo** map; the color of their specular highlights is determined by the color of the incoming light. Following these principles allows the metallic workflow to combine the albedo and specular into the **Albedo** map, enforcing physically-accurate behaviors. This also allows saving more space, resulting in a significant speed up at the expenses of reduced control over the look of your materials.

See also

For more information about the metallic setup, you can refer to these links:

- **Calibration chart**: How to calibrate a metallic material (`http://blogs.unity3d.com/wp-content/uploads/2014/11/UnityMetallicChart.png`)
- **Material chart**: How to initialize the Standard Shader parameters for common materials (`http://docs.unity3d.com/Manual/StandardShaderMaterialCharts.html`)
- **Quixel MEGASCANS**: A vast library of materials, including textures and PBR parameters (`http://quixel.se/megascans`)
- **PBR Texture Conversion**: How traditional shaders can be converted to PBR shaders (`http://www.marmoset.co/toolbag/learn/pbr-conversion`)
- **Substance Designer**: A node-based software to work with PBR (`https://www.allegorithmic.com/products/substance-designer`)
- **The Theory of Physically-based Rendering**: A complete guide about PBR (`https://www.allegorithmic.com/pbr-guide`)

Adding transparency to PBR

Transparency is such an important aspect in games that the Standard Shader supports three different ways of doing it. This recipe is useful if you need to have realistic materials with transparent or semi-transparent properties. Glasses, bottles, windows, and crystals are good candidates for PBR transparent shaders. This is because you can still have all the realism introduced by PBR with the addition of a transparent or translucent effect. If you need transparency for something different such as UI elements or pixel art, there are more efficient alternatives that are explored in the *Creating a transparent material* recipe in *Chapter 2, Surface Shaders and Texture Mapping*.

> In order to have a transparent Standard material, changing the alpha channel of its **Albedo** color property is not enough. Unless you properly set its **Rendering Mode**, your material will not appear transparent.

Getting ready

This recipe will use the Standard Shader, so there is no need to create a new one:

1. Create a new material.

2. Make sure that the **Shader** property is set to either **Standard** or **Standard (Specular setup)** from the material's **Inspector** tab.

3. Assign the newly created material to the 3D object that you want to be transparent.

How to do it...

The Standard Shader provides three different types of transparencies. Despite being very similar, they have subtle differences and fit different contexts.

Semi-transparent materials

Some materials such as clear plastics, crystal, and glass are semi-transparent. This means that they both require all the realistic effects of PBR (such as specular highlights and Fresnel refraction and reflection) but allow the geometry behind to be seen. If this is what you need, perform the following steps:

1. From the material's **Inspector** tab, set **Rendering Mode** to **Transparent**.

2. The amount of transparency is determined by the alpha channel of the **Albedo** color or the **Albedo** map (if any).

The following picture shows the Unity 5 calibration scene with four different highly polished plastic spheres. From left to right, their transparency is increased. The last sphere is fully transparent, but retains all the added effects of PBR:

The **Transparent** rendering mode is perfect for windows, bottles, gems, and headsets.

> You should notice that many transparent materials don't usually project shadows. On top of this, the **Metallic** and **Smoothness** properties of a material can interfere with the transparency effect. A mirror-like surface can have the alpha set to zero, but if it reflects all the incoming light, it won't appear transparent.

Fading objects

Sometimes, you want an object to fully disappear with a fading effect. In this case, specular reflections and Fresnel refraction and reflection should disappear as well. When a fading object is fully transparent, it should also be invisible. To do this, perform the following steps:

1. From the material's **Inspector** tab, set **Rendering Mode** to **Fade**.

2. As before, use the alpha channel of the **Albedo** color or map to determine the final transparency.

The following picture shows fading spheres. It is clear from the picture that the PBR effects fade with the sphere as well. As you can see in the following image, the last one on the right is almost invisible:

This rendering mode works best for non-realistic objects, such as holograms, laser rays, faux lights, ghosts, and particle effects.

Solid geometries with holes

Most of the materials encountered in a game are solid, meaning that they don't allow light to pass through them. At the same time, many objects have a very complex (yet flat) geometry. Modeling leaves and grass with 3D objects is often overkill. A more efficient approach is to use a quad (rectangle) with a leaf texture. While the leaf itself is solid, the rest of the texture should be fully transparent. If this is what you want, then perform the following steps:

1. From the material's **Inspector** tab, set **Rendering Mode** to **Cutout**.
2. Use the **Alpha Cutoff** slider to determine the cutoff threshold. All the pixels in the **Albedo** map with an alpha value equal to or less than **Alpha Cutoff** will be hidden.

The following image, taken from the **Unity Official Tutorials** on PBR (`https://www.youtube.com/watch?v=fD_ho_ofY6A`), shows you how the effect of the **Cutout** rendering mode can be used to create a hole in the geometry:

It's worth noticing that **Cutout** does not allow the back of the geometry to be seen. In the previous example, you could not see the inner volume of the sphere. If you require such an effect, you need to create your own shader and make sure that the back geometry is not culled.

See also

▶ The examples in these recipe have been created using the Unity 5 **Shader Calibration Scene**, which is freely available in the **Asset Store** at `https://www. assetstore.unity3d.com/en/#!/content/25422`.

▶ More information about albedo and transparency can be found at `http://docs.unity3d.com/Manual/ StandardShaderMaterialParameterAlbedoColor.html`.

Creating mirrors and reflective surfaces

Specular materials reflect lights when objects are viewed from certain angles. Unfortunately, even the Fresnel reflection, which is one of the most accurate models, does not correctly reflect lights from nearby objects. The lighting models examined in the previous chapters took into account only light sources, but ignored light that is reflected from other surfaces. With what you've learned about shaders so far, making a mirror is simply not possible. Global illumination makes this possible by providing PBR shaders with information about their surroundings. This allows objects to have not just specular highlights, but also real reflections, which depend on the other objects around them. Real-time reflections are very costly and require manual setting up and tweaking in order to work. When done properly, they can be used to create mirror-like surfaces, as seen in the following picture:

Getting ready

This recipe will not feature any new shader. Quite the opposite; most of the work is done directly in the editor. Perform the following steps:

1. Create a new scene.
2. Create a quad, which will serve as a mirror.
3. Create a new material and attach it to the mirror.
4. Place the quad in a scene with other objects.
5. Create a new **reflection probe** from **GameObject** | **Light** | **Reflection Probe** and place it in front of the quad.

How to do it...

If the preceding steps have been followed correctly, you should have a quad in the middle of your scene, close to a reflection probe. In order to make it in a mirror, some changes need to be made:

1. Change the shader of the material to **Standard** and its **Rendering Mode** to **Opaque**.
2. Change its **Metallic** and **Smoothness** properties to one. You should see the material reflecting the sky more clearly.
3. Select the reflection probe and change its **Size** and **Probe Origin** until it is in front of the quad and it encloses all the objects that you want to reflect.
4. Finally, change its **Type** to **Realtime**. Make sure that **Culling Mask** is set to **Everything**.

Your reflection probe should be configured, as shown in the following image:
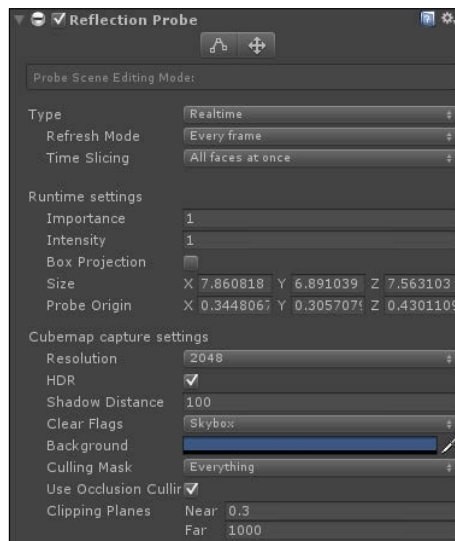
If your probe is used for a real mirror, you should check the **Box Projection** flag. If it is used for other reflective surfaces, such as shiny pieces of metal or glass tables, you can uncheck it.

How it works...

When a shader wants information about its surroundings, it is usually provided in a structure called **cube maps**. They have been briefly mentioned in *Chapter 1, Creating Your First Shader*, as one of the shader property types, among `Color`, `2D`, `Float`, and `Vector`. Loosely speaking, cube maps are the 3D equivalent of 2D textures; they represent a 360-degree view of the world, as seen from a center point. Unity 5 previews cube maps with a spherical projection, as seen in the following picture:

When cube maps are attached with a camera, they are referred to as **skyboxes** as they are used to provide a way to reflect the sky. They can be used to reflect geometries that are not in the actual scene, such as nebulae, clouds, stars, and so on.

The reason why they are called cube maps is because of the way they are created: a cube map is made up of six different textures, each one attached to the face of a cube. You can create a cube map manually or delegate it to a **reflection probe**. You can imagine a reflection probe as a collection of six cameras, creating a 360 mapping of the surrounding area. This also gives you an idea why probes are so expensive. By creating one in our scene, we allow Unity to know which objects are around the mirror. If you need more reflective surfaces, you can add multiple probes. You need no further action for the reflection probes to work. The Standard Shaders will use them automatically.

You should notice that when they are set to **Realtime**, they render their cube map at the beginning of every frame. There is a trick to make this faster; if you know that part of the geometry that you want to reflect does not move, you can bake the reflection. This means that Unity can calculate the reflection before starting the game, allowing more precise (and computationally expensive) calculations. In order to do this, your reflection probe must be set to **Baked** and will work only for objects that are flagged as **Static**. Static objects cannot move or change, which makes them perfect for terrains, buildings, and props. Every time a static object is moved, Unity will regenerate the cube maps for its baked reflection probes. This might take a few minutes to several hours.

You can mix **Realtime** and **Baked** probes to increase the realism of your game. Baked probes will provide very high-quality reflections environmental reflections, while the real-time ones can be used to move objects such as cars or mirrors. The next *Baking lights in your scene* recipe will explain in detail how light baking works.

See also

If you are interested in learning more about reflection probes, you should check these links:

- Unity 5 manual about Reflection Probe: `http://docs.unity3d.com/Manual/class-ReflectionProbe.html`

Baking lights in your scene

Rendering lighting is a very expensive process. Even with state-of-the-art GPUs, accurately calculating the **light transport** (which is how light bounces between surfaces) can take hours. In order to make this process feasible for games, real-time rendering is essential. Modern engines compromise between realism and efficiency; most of the computation is done beforehand in a process called **light baking**. This recipe will explain how light baking works and how you can get the most out of it.

Getting ready

Light baking requires you to have a scene ready. It should have geometries and, obviously, lights. For this recipe, we will rely on Unity's standard features so there is no need to create additional shaders or materials. For a better control, you might want to access the **Lighting** window. If you don't see it, select **Window** | **Lighting** from the menu and dock it where it is more convenient for you.

How to do it...

Light baking requires some manual configuration. There are three essential, yet independent, steps that you need to take.

Configuring the static geometry

These steps must be followed for the configuration:

1. Identify all the objects in your scene that do not change position, size, and material. Possible candidates are buildings, walls, terrains, props, trees, and others.

2. Select these objects, and check the **Static** box from the **Inspector** tab, as shown in the following image. If any of the selected objects has children, Unity will ask if you want them to be considered static as well. If they meet the requirements (fixed position, size, and material), select **Yes, change children** in the pop-up box:

3. If a light qualifies as a static object but illuminates non-static geometry, make sure that its **Baking** property is set to **Mixed**. If it will affect only static objects, set it to **Baked**.

Configuring the light probes

There are objects in your game that will move, such as the main character, enemies, and the other **non-playable characters** (**NPCs**). If they enter a static region that is illuminated, you might want to surround it with light probes. To do this, follow the given steps:

1. From the menu, navigate to **GameObject | Light | Light Probe Group**. A new object called **Light Probe Group** will appear in **Hierarchy**.

2. Once selected, four interconnected spheres will appear. Click and move them around the scene so that they enclose the static region in which your characters can enter. The following picture shows an example of how light probes can be used to enclose the volume of a static office space:

3. Select the moving objects that will enter the light probe region.

4. From their **Inspector**, expand their **renderer component** (usually **Mesh Renderer**) and make sure that **Use Light Probes** is checked (see the following image):
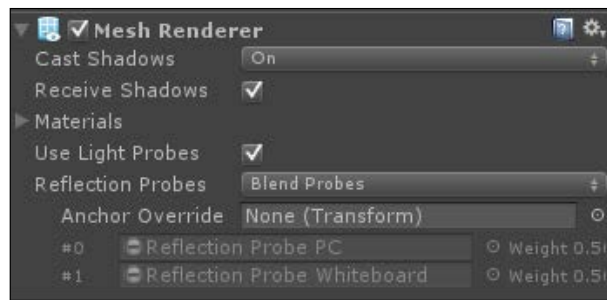
Deciding where and when to use light probes is a critical problem; more information about this can be found in the *How it works...*section.

Baking the lights

To bake the lights, follow the given steps:

1. To finally bake the lights, open the **Lighting** window and select its **Lightmaps** tab.

2. If the **Auto** checkbox is enabled, Unity will automatically execute the baking process in the background. If not, click on **Build**.

> Light baking can take several hours even for a relatively small scene. If you are constantly moving static objects or lights, Unity will restart the process from scratch causing a severe slowdown in the editor. You can uncheck the **Auto** checkbox from the **Lighting** | **Lightmaps** tab to prevent this so that you can decide when to start the process manually.

How it works...

The most complicated part of the rendering is the light transport. During this phase, the GPU calculates how the rays of light bounce between objects. If an object and its lights don't move, this calculation can be done only once as it will never change during the game. Flagging an object as **Static** is how you are telling Unity that such an optimization can be made.

Loosely speaking, light baking refers to the process of calculating the global illumination of a static object and saving it in what is called a **lightmap**. Once baking is completed, lightmaps can be seen in the **Lightmaps** tab of the **Lighting** window:
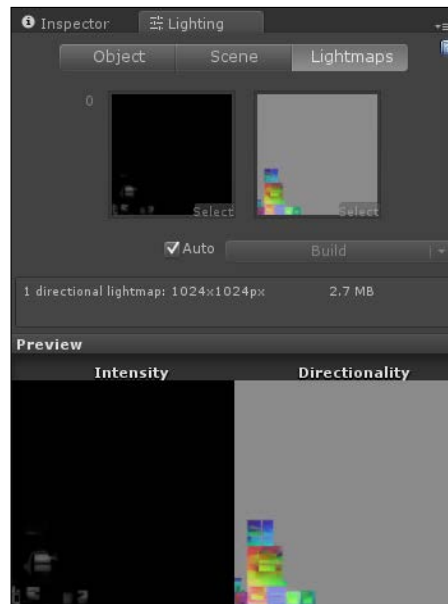
Light baking comes at a great expense: memory. Every static surface is, in fact, retextured so that it already includes its lighting condition. Let's imagine that you have a forest of trees, all sharing the same texture. Once they are made static, each tree will have its very own texture. Light baking not only increases the size of your game, but can take a lot of texture memory if used indiscriminately.

The second aspect introduced in this recipe is **light probing**. Light baking produces extremely high-quality results for static geometries but does not work on moving objects. If your character is entering in a static region, it can look somehow *detached* from the environment. Its shading will not match the surrounding, resulting in an aesthetically unpleasant result. Other objects, such as **skinned mesh renderers**, will not receive global illumination even if made static. Baking lights in real time is not possible, although light probes offer an effective alternative. Every light probe samples the global illumination at a specific point in space. A light probe group can sample several points in space, allowing to interpolate global illumination within a specific volume. This allows us to cast a better light on moving objects, even despite the fact that global illumination has been calculated only for a few points. It is important to remember that light probes need to enclose a volume in order to work. It is best to place light probes in regions where there is a sudden change in the light condition. Similar to lightmaps, probes consume memory and should be placed wisely; remember that they exist only for non-static geometry.

Even while using light probes, there are a few aspects that Unity's global illumination cannot capture. Non-static objects, for instance, cannot reflect light on other objects.

See also

You can read more about light probes at `http://docs.unity3d.com/Manual/LightProbes.html`.

5

Vertex Functions

The term **shader** originates from the fact that Cg has been used mainly to simulate realistic lighting conditions (shadows) on 3D models. Despite this, shaders are now much more than that. They not only define the way the objects are going to look, they can also redefine their shapes entirely. If you want to learn how to manipulate the geometry of a 3D object via shaders, this is the chapter for you.

In this chapter, you will learn the following recipes:

- ▸ Accessing a vertex color in a Surface Shader
- ▸ Animating vertices in a Surface Shader
- ▸ Extruding your models
- ▸ Implementing a snow shader
- ▸ Implementing a volumetric explosion

Introduction

In *Chapter 1*, *Creating Your First Shader*, we explained that 3D models are not just a collection of triangles. Each vertex can contain data that is essential to render the model itself correctly. This chapter will explore how to access this information in order to use it in a shader. We will also explore in detail how the geometry of an object can be deformed simply using Cg code.

Accessing a vertex color in a Surface Shader

Let's begin this chapter by taking a look at how we can access the information of a model's vertex using the vertex function in a Surface Shader. This will arm us with the knowledge to start utilizing the elements contained within a model's vertex to create really useful and visually appealing effects.

A vertex in a vertex function can return information about itself that we need to be aware of. You can actually retrieve the vertices' normal directions as a `float3` value, the position of the vertex as `float3`, and you can even store color values in each vertex and return that color as `float4`. This is what we will take a look at in this recipe. We need to see how to store color information and retrieve this stored color information inside each vertex of a Surface Shader.

Getting ready

In order to write this shader, we need to prepare a few assets. The following steps will set us up to create this Vertex Shader:

1. In order to view the colors of a vertex, we need to have a model that has had color applied to its vertices. While you could use Unity to apply colors, you would have to write a tool to allow an individual to apply the colors or write some scripts to achieve the color application. In the case of this recipe, we simply utilized **Maya** to apply the colors to our model. This model is available on the book's **Support** page at `https://www.packtpub.com/books/content/support`.

2. Create a new scene and place the imported model in the scene.

3. Create a new shader and material. When completed, assign the shader to the material and then the material to the imported model.

Your scene should now look similar to the following screenshot:

How to do it...

With our scene, shader, and material created and ready to go, we can begin to write the code for our shader. Launch the shader by double-clicking on it in the **Project** tab in the Unity editor. Perform the following steps:

1. As we are creating a very simple shader, we will not need to include any properties in our `Properties` block. We will still include a global tint color, just to stay consistent with the other shaders in this book. Enter the following code in the `Properties` block of your shader:

```
Properties
{
  _MainTint("Global Color Tint", Color) = (1,1,1,1)
}
```

2. This next step tells Unity that we will be including a vertex function in our shader:

```
CGPROGRAM
#pragma surface surf Lambert vertex:vert
```

3. As usual, if we have included properties in our `Properties` block, we must make sure to create a corresponding variable in our `CGPROGRAM` statement. Enter the following code just below the `#pragma` statement:

```
float4 _MainTint;
```

4. We now turn our attention to the `Input` struct. We need to add a new variable in order for our `surf()` function to access the data given to us by our `vert()` function:

```
struct Input
{
    float2 uv_MainTex;
    float4 vertColor;
};
```

5. Now, we can write our simple `vert()` function to gain access to the colors stored in each vertex of our mesh:

```
void vert(inout appdata_full v, out Input o)
{
    o.vertColor = v.color;
}
```

6. Finally, we can use the vertex color data from our `Input` struct to be assigned to the `o.Albedo` parameters in the built-in `SurfaceOutput` struct:

```
void surf (Input IN, inout SurfaceOutput o)
{
    o.Albedo = IN.vertColor.rgb * _MainTint.rgb;
}
```

7. With our code completed, we can now re-enter the Unity editor and let the shader compile. If all goes well, you should see something similar to the following screenshot:

How it works...

Unity provides us with a way to access the vertex information of the model to which a shader is attached. This gives us the power to modify things such as the vertices' position and color. With this recipe, we have imported a mesh from Maya (though just about any 3D software application can be used), where vertex colors were added to `Verts`. You'll notice that by importing the model, the default material will not display the vertex colors. We actually have to write a shader to extract the vertex color and display it on the surface of the model. Unity provides us with a lot of built-in functionality when using Surface Shaders, which make the process of extracting this vertex information quick and efficient.

Our first task is to tell Unity that we will be using a vertex function when creating our shader. We do this by adding the `vertex:vert` parameter to the `#pragma` statement of `CGPROGRAM`. This automatically makes Unity look for a vertex function named `vert()` when it goes to compile the shader. If it doesn't find one, Unity will throw a compiling error and ask you to add a `vert()` function to your shader.

This brings us to our next step. We have to actually code the `vert()` function, as seen in step 5. By having this function, we can access the built-in data struct called `appdata_full`. This built-in struct is where the vertex information is stored. So, we then extract the vertex color information by passing it to our `Input` struct by adding the code, `o.vertColor = v.color`.

The `o` variable represents our `Input` struct and the `v` variable is our `appdata_full` vertex data. In this case, we are simply taking the color information from the `appdata_full` struct and putting it in our `Input` struct. Once the vertex color is in our `Input` struct, we can use it in our `surf()` function. In the case of this recipe, we simply apply the color to the `o.Albedo` parameter to the built-in `SurfaceOutput` struct.

There's more...

One can also access a fourth component from the `vert` color data. If you notice, the `vertColor` variable we declared in the `Input` struct is of the `float4` type. This means that we are also passing the alpha value of the vertex colors. Knowing this, you can use it to your advantage for the purpose of storing a fourth vertex color to perform effects such as transparency or giving yourself one more mask to blend two textures. It's really up to you and your production to determine if you really need to use the fourth component, but it is worth mentioning here.

With Unity 5, we now have the ability to target shaders to DirectX 11. This is great, but it means that the compiling process for the shaders is now a bit pickier. This means that we need to include one more line of code to our shader to initialize the output of the vertex information properly. The following code shows what the vertex function code looks like, if you are using DirectX 11 in your shader:

```
void vert(inout appdata_full v, out Input o)
{
  UNITY_INITIALIZE_OUTPUT(Input, o);
  o.vertColor = v.color;
}
```

By including this line of code, your Vertex Shader will not throw any warnings, which say that it won't compile to DirectX 11 appropriately.

Animating vertices in a Surface Shader

Now that we know how to access data on a per-vertex basis, let's expand our knowledge set to include other types of data and position of a vertex.

Using a vertex function, we can access the position of each vertex in a mesh. This allows us to actually modify each individual vertex while the shader does the processing.

In this recipe, we will create a shader that will allow us to modify the positions of each vertex on a mesh with a sine wave. This technique can be used to create animations for objects such as flags or waves on an ocean.

Getting ready

Let's gather our assets together so that we can create the code for our Vertex Shader:

1. Create a new scene and place a plane mesh in the center of the scene.
2. Then create a new shader and material.
3. Finally, assign the shader to the material and the material to the plane mesh.

Your scene should look similar to the following screenshot:

How to do it...

With our scene ready to go, let's double-click on our newly created shader to open it in MonoDevelop:

1. Let's begin with our shader by populating the `Properties` block:

   ```
   Properties
   {
     _MainTex ("Base (RGB)", 2D) = "white" {}
     _tintAmount ("Tint Amount", Range(0,1)) = 0.5
     _ColorA ("Color A", Color) = (1,1,1,1)
     _ColorB ("Color B", Color) = (1,1,1,1)
     _Speed ("Wave Speed", Range(0.1, 80)) = 5
     _Frequency ("Wave Frequency", Range(0, 5)) = 2
     _Amplitude ("Wave Amplitude", Range(-1, 1)) = 1
   }
   ```

2. We now need to tell Unity that we are going to be using a vertex function by adding the following to the `#pragma` statement:

   ```
   CGPROGRAM
   #pragma surface surf Lambert vertex:vert
   ```

3. In order to access the values that have been given to us by our properties, we need to declare a corresponding variable in our `CGPROGRAM` block:

```
sampler2D _MainTex;
float4 _ColorA;
float4 _ColorB;
float _tintAmount;
float _Speed;
float _Frequency;
float _Amplitude;
float _OffsetVal;
```

4. We will be using the vertex position modification as a `vert` color as well. This will allow us to tint our object:

```
struct Input
{
  float2 uv_MainTex;
  float3 vertColor;
}
```

5. At this point, we can perform our vertex modification using a sine wave and vertex function. Enter the following code after the `Input` struct:

```
void vert(inout appdata_full v, out Input o)
{
  float time = _Time * _Speed;
  float waveValueA = sin(time + v.vertex.x * _Frequency) *
    _Amplitude;

  v.vertex.xyz = float3(v.vertex.x, v.vertex.y + waveValueA,
    v.vertex.z);
  v.normal = normalize(float3(v.normal.x + waveValueA,
    v.normal.y, v.normal.z));
  o.vertColor = float3(waveValueA,waveValueA,waveValueA);
}
```

6. Finally, we complete our shader by performing a `lerp()` function between two colors so that we can tint the peaks and valleys of our new mesh, modified by our vertex function:

```
void surf (Input IN, inout SurfaceOutput o)
{
  half4 c = tex2D (_MainTex, IN.uv_MainTex);
  float3 tintColor = lerp(_ColorA, _ColorB, IN.vertColor).rgb;

  o.Albedo = c.rgb * (tintColor * _tintAmount);
  o.Alpha = c.a;
}
```

After completing the code for your shader, switch back to Unity and let the shader compile. Once compiled, you should see something similar to the following screenshot:

How it works...

This particular shader uses the same concept from the last recipe, except that this time, we are modifying the positions of the vertices in the mesh. This is really useful if you don't want to rig up simple objects, such as a flag, and then animate them using a skeleton structure or hierarchy of transforms.

We simply create a sine wave value using the `sin()` function that is built into the Cg language. After calculating this value, we add it to the `y` value of each vertex position, creating a wave-like effect.

We also did a little bit of modification to the normal on the mesh just to give it a more realistic shading based on the sine wave value.

You will see how easy it is to perform more complex vertex effects by utilizing the built-in vertex parameters that Surface Shaders give us.

Extruding your models

One of the biggest problems in games is repetitions. Creating new content is a time-consuming task, and when you have to face thousands of enemies, chances are that they will all look the same. A relatively cheap technique to add variations to your models is using a shader that alters its basic geometry. This recipe will show you a technique called **normal extrusion**, which can be used to create a chubbier or skinnier version of a model, as shown in the following picture with the soldier from the Unity camp demo:

Getting ready

For this recipe, we need to have access to the shader used by the model that you want to alter. Once you have it, we duplicate it so that we can edit it safely. It can be done as follows:

1. Find the shader your model is using and once selected, duplicate it by pressing *Ctrl + D*.

2. Duplicate the original material of the model and assign the cloned shader to it.

3. Assign the new material to your model, and start editing it.

In order for this effect to work, your model should have **normals**.

How to do it...

To create this effect, start by modifying the duplicated shader:

1. Let's start by adding a property to our shader, which will be used to modulate its extrusion. The range presented here goes from `-1` to `+1`, but you might have to adjust this according to your own needs:

   ```
   _Amount ("Extrusion Amount", Range(-1,+1)) = 0
   ```

2. Couple the property with its respective variable:

   ```
   float _Amount;
   ```

3. Change the `#pragma` directive so that it now uses a vertex modifier. You can do this by adding `vertex:function_name` at the end of it. In our case, we have called the function, `vert`:

```
#pragma surface surf Lambert vertex:vert
```

4. Add the following vertex modifier:

```
void vert (inout appdata_full v) {
  v.vertex.xyz += v.normal * _Amount;
}
```

5. The shader is now ready; you can use the **Extrusion Amount** slider in the material's **Inspector** tab to make your model skinnier or chubbier.

How it works...

Surface Shaders works in two steps. In all the previous chapters, we only explored its last one: the surface function. There is another function that can be used: the **vertex modifier**. It takes the data structure of a vertex (which is usually called `appdata_full`) and applies a transformation to it. This gives us the freedom to do virtually everything with the geometry of our model. We signal the **graphics processing unit** (**GPU**) that such a function exists by adding `vertex:vert` to the #pragma directive of the Surface Shader. You can refer to *Chapter 6*, *Fragment Shaders and Grab Passes*, to learn how vertex modifiers can be defined in a Vertex and Fragment Shader instead.

One of the most simple, yet effective, techniques that can be used to alter the geometry of a model is called normal extrusion. It works by projecting a vertex along its normal direction. This is done by the following line of code:

```
v.vertex.xyz += v.normal * _Amount;
```

The position of a vertex is displaced by `_Amount` units toward the vertex normal. If `_Amount` gets too high, the results can be quite unpleasant. With smaller values, however, you can add a lot of variations to your models.

There's more...

If you have multiple enemies and want each one to have its own *weight*, you have to create a different material for each one of them. This is necessary as materials are normally shared between models and changing one will change all of them. There are several ways in which you can do this; the quickest one is to create a script that automatically does it for you. The following script, once attached to an object with a `Renderer`, will duplicate its first material and set the `_Amount` property automatically:

```
using UnityEngine;
public class NormalExtruder : MonoBehaviour {
```

```
    [Range(-0.0001f, 0.0001f)]
    public float amount = 0;

    // Use this for initialization
    void Start () {
      Material material = GetComponent<Renderer>().sharedMaterial;
      Material newMaterial = new Material(material);
      newMaterial.SetFloat("_Amount", amount);
      GetComponent<Renderer>().material = newMaterial;
    }
}
```

Adding extrusion maps

This technique can actually be improved even further. We can add an extra texture (or use the alpha channel of the main one) to indicate the amount of the extrusion. This allows a much better control over which parts are raised or lowered. The following code shows you how it is possible to achieve such an effect:

```
sampler2D _ExtrusionTex;
void vert(inout appdata_full v) {
float4 tex = tex2Dlod (_ExtrusionTex, float4(v.texcoord.xy,0,0));
  float extrusion = tex.r * 2 - 1;
  v.vertex.xyz += v.normal * _Amount * extrusion;
}
```

The red channel of `_ExtrusionTex` is used as a multiplying coefficient for normal extrusion. A value of 0.5 leaves the model unaffected; darker or lighter shades are used to extrude vertices inward or outward, respectively. You should notice that to sample a texture within a vertex modifier, `tex2Dlod` should be used instead of `tex2D`.

> In shaders, color channels go from 0 to 1, although sometimes you need to represent negative values as well (such as inward extrusion). When this is the case, treat 0.5 as zero, having smaller values considered as negative and higher values as positive. This is exactly what happens with normals, which are usually encoded in RGB textures. The `UnpackNormal()` function is used to map a value in the range (0,1) on the range (-1,+1). Mathematically speaking, this is equivalent to `tex.r * 2 -1`.

Extrusion maps are perfect to zombify characters by shrinking the skin to highlight the shape of the bones underneath. The following picture shows you how a *healthy* soldier can be transformed into a corpse using just a shader and extrusion map. Compared to the previous example, you can notice how the clothing is unaffected. The shader used in the following picture also darkens the extruded regions to give an even more emaciated look to the soldier:

Implementing a snow shader

The simulation of snow has always been a challenge in games. The vast majority of games simply includes snow directly in the model's textures so that their tops look white. However, what if one of these objects starts rotating? Snow is not just a lick of paint on a surface; it is a proper accumulation of material and should be treated as such. This recipe shows you how to give a snowy look to your models using just a shader.

This effect is achieved in two steps. First, a white color is used for all the triangles facing the sky. Second, their vertices are extruded to simulate the effect of snow accumulation. You can see the result in the following picture:

> Keep in mind that this recipe does not aim to create a photorealistic snow effect. It provides a good starting point, but it is up to an artist to create the right textures and find the right parameters to make it fit your game.

Getting ready

This effect is purely based on shaders. We will need the following:

1. Create a new shader for the snow effect.
2. Create a new material for the shader.
3. Assign the newly created material to the object that you want to be snowy.

How to do it...

To create a snowy effect, open your shader and make the following changes:

1. Replace the properties of the shader with the following ones:

    ```
    _MainColor("Main Color", Color) = (1.0,1.0,1.0,1.0)
    _MainTex("Base (RGB)", 2D) = "white" {}
    _Bump("Bump", 2D) = "bump" {}
    _Snow("Level of snow", Range(1, -1)) = 1
    _SnowColor("Color of snow", Color) = (1.0,1.0,1.0,1.0)
    _SnowDirection("Direction of snow", Vector) = (0,1,0)
    _SnowDepth("Depth of snow", Range(0,1)) = 0
    ```

2. Complete them with their relative variables:

    ```
    sampler2D _MainTex;
    sampler2D _Bump;
    float _Snow;
    float4 _SnowColor;
    float4 _MainColor;
    float4 _SnowDirection;
    float _SnowDepth;
    ```

3. Replace the `Input` structure with the following one:

    ```
    struct Input {
      float2 uv_MainTex;
      float2 uv_Bump;
      float3 worldNormal;
      INTERNAL_DATA
    };
    ```

4. Replace the surface function with the following one. It will color the snowy parts of the model white:

```
void surf(Input IN, inout SurfaceOutputStandard o) {
  half4 c = tex2D(_MainTex, IN.uv_MainTex);
  o.Normal = UnpackNormal(tex2D(_Bump, IN.uv_Bump));

  if (dot(WorldNormalVector(IN, o.Normal),
    _SnowDirection.xyz) >= _Snow)
    o.Albedo = _SnowColor.rgb;
  else
    o.Albedo = c.rgb * _MainColor;
  o.Alpha = 1;
}
```

5. Configure the `#pragma` directive so that it uses vertex modifiers:

```
#pragma surface surf Standard vertex:vert
```

6. Add the following vertex modifiers, which extrude the vertices covered in snow:

```
void vert(inout appdata_full v) {
  float4 sn = mul(UNITY_MATRIX_IT_MV, _SnowDirection);
  if (dot(v.normal, sn.xyz) >= _Snow)
    v.vertex.xyz += (sn.xyz + v.normal) * _SnowDepth * _Snow;
}
```

You can now use the material's **Inspector** tab to select how much of your model is going to be covered and how thick the snow should be.

How it works...

This shader works in two steps.

Coloring the surface

The first one alters the color of the triangles that are facing the sky. It affects all the triangles with a normal direction similar to `_SnowDirection`. As seen before in *Chapter 3, Understanding Lighting Models*, comparing unit vectors can be done using the **dot product**. When two vectors are orthogonal, their dot product is zero; it is one (or minus one) when they are parallel to each other. The `_Snow` property is used to decide how aligned they should be to be considered facing the sky.

If you look closely at the surface function, you can see that we are not dotting the normal and snow direction directly. This is because they are usually defined in a different space. The snow direction is expressed in world coordinates, while the object normals are usually relative to the model itself. If we rotate the model, its normals will not change, which is not what we want. To fix this, we need to convert the normals from their object coordinates to world coordinates. This is done with the `WorldNormalVector()` function, as seen in the following code:

```
if (dot(WorldNormalVector(IN, o.Normal), _SnowDirection.xyz) >=
  _Snow)
  o.Albedo = _SnowColor.rgb;
else
  o.Albedo = c.rgb * _MainColor;
```

This shader simply colors the model white; a more advanced one should initialize the `SurfaceOutputStandard` structure with textures and parameters from a realistic snow material.

Altering the geometry

The second effect of this shader alters the geometry to simulate the accumulation of snow. Firstly, we identify which triangles have been colored white by testing the same condition used in the surface function. This time, unfortunately, we cannot rely on `WorldNormalVector()` as the `SurfaceOutputStandard` structure is not yet initialized in the vertex modifier. We use this other method instead, which converts `_SnowDirection` to object coordinates:

```
float4 sn = mul(UNITY_MATRIX_IT_MV, _SnowDirection);
```

Then, we can extrude the geometry to simulate the accumulation of snow:

```
if (dot(v.normal, sn.xyz) >= _Snow)
v.vertex.xyz += (sn.xyz + v.normal) * _SnowDepth * _Snow;
```

Once again, this is a very basic effect. One could use a texture map to control the accumulation of snow more precisely or give a peculiar, uneven look.

See also

If you need high-quality snow effects and props for your game, you can also check these resources on the Unity **Asset Store**:

> ▸ **Winter Suite ($30)**: A much more sophisticated version of the snow shader presented in this recipe can be found at `https://www.assetstore.unity3d.com/en/#!/content/13927`

> ▸ **Winter Pack ($60)**: A very realistic set of props and materials for snowy environments can be found at `https://www.assetstore.unity3d.com/en/#!/content/13316`

Implementing a volumetric explosion

The art of game development is a clever trade-off between realism and efficiency. This is particularly true for explosions; they are at the heart of many games, yet the physics behind them is often beyond the computational power of modern machines. Explosions are, essentially, nothing more than very hot balls of gas; hence, the only way to correctly simulate them is by integrating a fluid simulation into your game. As you can imagine, this is infeasible for runtime application, and many games simulate them simply with particles. When an object explodes, it is common to simply instantiate many fire, smoke, and debris particles that, together, can achieve believable results. This approach, unfortunately, is not very realistic and is easy to spot. There is an intermediate technique that can be used to achieve a much more realistic effect: **volumetric explosions**. The idea behind this concept is that explosions are not treated any more like a bunch of particles; they are evolving 3D objects, not just flat 2D textures.

Getting ready

Start this recipe with the following steps:

1. Create a new shader for this effect.

2. Create a new material to host the shader.

3. Attach the material to a sphere. You can create one directly from the editor, navigating to **GameObject** | **3D Object** | **Sphere**.

> This recipe works well with the standard Unity Sphere, but if you need big explosions, you might need to use a more high-poly sphere. In fact, a vertex function can only modify the vertices of a mesh. All the other points will be interpolated using the positions of the nearby vertices. Fewer vertices mean a lower resolution for your explosions.

4. For this recipe, you will also need a ramp texture that has, in a gradient, all the colors your explosions will have. You can create a texture like the following image using **GIMP** or **Photoshop**:

5. Once you have the picture, import it to Unity. Then, from its **Inspector**, make sure that **Filter Mode** is set to **Bilinear** and **Wrap Mode** to **Clamp**. These two settings make sure that the ramp texture is sampled smoothly.

6. Lastly, you will need a noisy texture. You can search on the Internet for freely available **noise textures**. The most commonly used ones are generated using **Perlin noise**.

How to do it...

This effect works in two steps: a vertex function to change the geometry and surface function to give it the right color. The steps are as follows:

1. Add the following properties to the shader:

```
_RampTex("Color Ramp", 2D) = "white" {}
_RampOffset("Ramp offset", Range(-0.5,0.5))= 0

_NoiseTex("Noise tex", 2D) = "gray" {}
_Period("Period", Range(0,1)) = 0.5

_Amount("_Amount", Range(0, 1.0)) = 0.1
_ClipRange("ClipRange", Range(0,1)) = 1
```

2. Add their relative variables so that the Cg code of the shader can actually access them:

```
sampler2D _RampTex;
half _RampOffset;

sampler2D _NoiseTex;
float _Period;

half _Amount;
half _ClipRange;
```

3. Change the `Input` structure so that it receives the UV data of the ramp texture:

```
struct Input {
  float2 uv_NoiseTex;
};
```

4. Add the following vertex function:

```
void vert(inout appdata_full v) {
  float3 disp = tex2Dlod(_NoiseTex,
    float4(v.texcoord.xy,0,0));
  float time = sin(_Time[3] *_Period + disp.r*10);
  v.vertex.xyz += v.normal * disp.r * _Amount * time;
}
```

5. Add the following surface function:

```
void surf(Input IN, inout SurfaceOutput o) {
  float3 noise = tex2D(_NoiseTex, IN.uv_NoiseTex);
  float n = saturate(noise.r + _RampOffset);
  clip(_ClipRange - n);
  half4 c = tex2D(_RampTex, float2(n,0.5));
```

```
    o.Albedo = c.rgb;
    o.Emission = c.rgb*c.a;
}
```

6. We specify the vertex function in the `#pragma` directive, adding the `nolightmap` parameter to prevent Unity from adding realistic lightings to our explosion:

    #pragma surface surf Lambert vertex:vert nolightmap

7. The last step is selecting the material, and from its **Inspector**, attaching the two textures in the relative slots. This is an animated material, meaning that it evolves over time. You can watch the material changing in the editor by clicking on **Animated Materials** from the **Scene** window:

How it works...

If you are reading this recipe, you should already be familiar with how Surface Shaders and vertex modifiers work. The main idea behind this effect is to alter the geometry of the sphere in a seemingly chaotic way, exactly like it happens in a real explosion. The following picture shows you how such an explosion will look inside the editor. You can see that the original mesh has been heavily deformed:

The vertex function is a variant of the technique called **normal extrusion** introduced in the *Extruding your models* recipe of this chapter. The difference here is that the amount of the extrusion is determined both by the time and noise texture.

> When you need a random number in Unity, you can rely on the `Random.Range()` function. There is no standard way to get random numbers in a shader, so the easiest way is to sample a noise texture.

There is no standard way to do this, so take this as an example only:

```
float time = sin(_Time[3] *_Period + disp.r*10);
```

The built-in `_Time[3]` variable is used to get the current time from within the shader, and the red channel of the noise texture `disp.r` is used to make sure that each vertex moves independently. The `sin()` function makes the vertices go up and down, simulating the chaotic behavior of an explosion. Then, the normal extrusion takes place:

```
v.vertex.xyz += v.normal * disp.r * _Amount * time;
```

You should play with these numbers and variables until you find a pattern of movement that you are happy with.

The last part of the effect is achieved by the surface function. Here, the noise texture is used to sample a random color from the ramp texture. However, there are two more aspects that are worth noticing. The first one is the introduction of `_RampOffset`. Its usage forces the explosion to sample colors from the left or right side of the texture. With positive values, the surface of the explosion tends to show more grey tones; exactly what happens when it is dissolving. You can use `_RampOffset` to determine how much fire or smoke there should be in your explosion. The second aspect introduced in the surface function is the usage of `clip()`. What `clip()` does is it clips (removes) pixels from the rendering pipeline. When invoked with a negative value, the current pixel is not drawn. This effect is controlled by `_ClipRange`, which determines which pixels of the volumetric explosions are going to be transparent.

By controlling both `_RampOffset` and `_ClipRange`, you have full control to determine how the explosion behaves and dissolves.

There's more...

The shader presented in this recipe makes a sphere look like an explosion. If you really want to use it, you should couple it with some scripts in order to get the most out of it. The best thing to do is to create an explosion object and make it into a prefab so that you can reuse it every time you need. You can do this by dragging the sphere back into the **Project** window. Once it is done, you can create as many explosions as you want using the `Instantiate()` function.

It is worth noticing, however, that all the objects with the same material share the same look. If you have multiple explosions at the same time, they should not use the same material. When you are instantiating a new explosion, you should also duplicate its material. You can do this easily with this piece of code:

```
GameObject explosion = Instantiate(explosionPrefab) as GameObject;
Renderer renderer = explosion.GetComponent<Renderer>();
Material material = new Material(renderer.sharedMaterial);
renderer.material = material;
```

Lastly, if you are going to use this shader in a realistic way, you should attach a script to it that changes its size, _RampOffset, and _ClipRange according to the type of explosion that you want to recreate.

See also

Much more can be done to make explosions realistic. The approach presented in this recipe only creates an empty shell; inside it, the explosion is actually empty. An easy trick to improve this is to create particles inside it. However, you can only go that far with this. The short movie, **The Butterfly Effect** (`http://unity3d.com/pages/butterfly`) created by **Unity Technologies** in collaboration with **Passion Pictures** and **Nvidia**, is the perfect example. It is based on the same concept of altering the geometry of a sphere, but it renders it with a technique called **volume ray casting**. In a nutshell, it renders the geometry as if it's full. You can see an example in the following picture:

If you are looking for high-quality explosions, check out **Pyro Technix** (`https://www.assetstore.unity3d.com/en/#!/content/16925`) on the Asset Store. It includes volumetric explosions and couples them with realistic shockwaves.

6

Fragment Shaders and Grab Passes

So far, we have relied on Surface Shaders. They have been designed to simplify the way shader coding works, providing meaningful tools for artists. If we want to push our knowledge of shaders further, we need to venture into the territory of Vertex and Fragment Shaders.

In this chapter, you will learn the following recipes:

- ▸ Understanding Vertex and Fragment Shaders
- ▸ Using grab pass
- ▸ Implementing a Glass Shader
- ▸ Implementing a Water Shader for 2D games

Introduction

Compared to Surface Shaders, Vertex and Fragment Shaders come with little to no information about the physical properties that determine how light reflects on surfaces. What they lack in expressivity, they compensate with power: Vertex and Fragment Shaders are not limited by physical constraints and are perfect for non-photorealistic effects. This chapter will focus on a technique called **grab pass**, which allows these shaders to simulate deformations.

Understanding Vertex and Fragment Shaders

The best way to understand how Vertex and Fragment Shaders work is by creating one yourself. This recipe will show you how to write one of these shaders, which will simply apply a texture to a model and multiply it by a given color, as shown in the following image:

The shader presented here is very simple, and it will be used as a starting base for all the other Vertex and Fragment Shaders.

Getting ready

For this recipe, we will need a new shader. Follow these steps:

1. Create a new shader.
2. Create a new material and assign the shader to it.

How to do it...

In all the previous chapters, we have always been able to refit Surface Shaders. This is not the case anymore as Surface and Fragment Shaders are structurally different. We will need the following changes:

1. Delete all the properties of the shader, replacing them with the following:

```
_Color ("Color", Color) = (1,0,0,1) // Red
_MainTex ("Base texture", 2D) = "white" {}
```

2. Delete all the code in the `SubShader` block and replace it with this one:

```
Pass {
    CGPROGRAM

    #pragma vertex vert
    #pragma fragment frag

    half4 _Color;
sampler2D _MainTex;

    struct vertInput {
        float4 pos : POSITION;
float2 texcoord : TEXCOORD0;
    };

    struct vertOutput {
        float4 pos : SV_POSITION;
float2 texcoord : TEXCOORD0;

    };

    vertOutput vert(vertInput input) {
        vertOutput o;
        o.pos = mul(UNITY_MATRIX_MVP, input.pos);
        o.texcoord = input.texcoord;
        return o;
    }

    half4 frag(vertOutput output) : COLOR{
        half4 mainColour = tex2D(_MainTex, output.texcoord);
        return mainColour * _Color;
    }

    ENDCG
    }
```

This will also be the base for all future Vertex and Fragment Shaders.

How it works...

As the name suggests, Vertex and Fragment Shaders work in two steps. The model is first passed through a **vertex function**; the result is then inputted to a **fragment function**. Both these functions are assigned using `pragma` directives:

```
#pragma vertex vert
#pragma fragment frag
```

In this case, they are simply called `vert` and `frag`.

Conceptually speaking, fragments are closely related to pixels; the term **fragment** is often used to refer to the collection of data necessary to draw a pixel. This is also why Vertex and Fragment Shaders are often called **Pixel Shaders**.

The vertex function takes the input data in a structure that is defined as `vertInput` in the shader:

```
struct vertInput {
    float4 pos : POSITION;
float2 texcoord : TEXCOORD0;
};
```

Its name is totally arbitrary, but its content is not. Each field of `struct` must be decorated with a **binding semantic**. This is a feature of Cg that allows us to mark variables so that they will be initialized with certain data, such as normal vectors and vertex position. The binding semantic `POSITION` indicates that when `vertInput` is inputted to the vertex function, `pos` will contain the position of the current vertex. This is similar to the `vertex` field of the `appdata_full` structure in a Surface Shader. The main difference is that `pos` is represented in model coordinates (relative to the 3D object), which we need to convert to view coordinates manually (relative to the position on the screen).

> The vertex function in a Surface Shader is used to alter the geometry of the model only. In a Vertex and Fragment Shader, instead, the vertex function is necessary to project the coordinates of the model to the screen.

The mathematics behind this conversion is beyond the scope of this chapter. However, this transformation can be achieved by multiplying `pos` by a special matrix provided by Unity: `UNITY_MATRIX_MVP`. It is often referred to as the **model-view-projection matrix**, and it is essential to find the position of a vertex on the screen:

```
vertOutput o;
o.pos = mul(UNITY_MATRIX_MVP, input.pos);
```

The other piece of information initialized is `textcoord`, which uses the `TEXCOORD0` binding semantics to get the UV data of the first texture. No further processing is required and this value can be passed directly to the fragment function:

```
o.texcoord = input.texcoord;
```

While Unity will initialise `vertInput` for us, we are responsible for the initialization of `vertOutput`. Despite this, its fields still need to be decorated with binding semantics:

```
struct vertOutput {
    float4 pos : SV_POSITION;
float2 texcoord : TEXCOORD0;
};
```

Once the vertex function has initialised `vertOutput`, the structure is passed to the fragment function. This samples the main texture of the model and multiplies it by the color provided.

As you can see, the Vertex and Fragment Shader has no knowledge of the physical properties of the material; compared to a Surface Shader, it works closer to the architecture of the **graphics processing unit** (**GPU**).

There's more...

One of the most confusing aspects of Vertex and Fragment Shaders is binding semantics. There are many others that you can use and their meaning depends on the context.

Input semantics

The binding semantics in the following table can be used in `vertInput`, which is the structure that Unity provides to the vertex function. The fields decorated with these semantics will be initialized automatically:

Binding semantics	Description
POSITION, SV_POSITION	The position of a vertex in world coordinates (object space)
NORMAL	The normal of a vertex, relative to the world (not to the camera)
COLOR, COLOR0, DIFFUSE, SV_TARGET	The color information stored in the vertex
COLOR1, SPECULAR	The secondary color information stored in the vertex (usually the specular)
TEXCOORD0, TEXCOORD1, …, TEXCOORDi	The i-th UV data stored in the vertex

Output semantics

When binding, semantics are used in `vertOutput`; they do not automatically guarantee that fields will be initialized. Quite the opposite; it's our responsibility to do so. The compiler will make its best to ensure that the fields are initialized with the right data:

Binding semantics	Description
`POSITION, SV_POSITION, HPOS`	The position of a vertex in camera coordinates (clip space, from zero to one for each dimension)
`COLOR, COLOR0, COL0, COL, SV_TARGET`	The front primary color
`COLOR1, COL1`	The front secondary color
`TEXCOORD0, TEXCOORD1, …, TEXCOORDi, TEXi`	The i-th UV data stored in the vertex
`WPOS`	The position, in pixels, in the window (origin in the lower left corner)

If, for any reason, you need a field that will contain a different type of data, you can decorate it with one of the many `TEXCOORD` data available. The compiler will not allow fields to be left undecorated.

See also

You can refer to the *NVIDIA Reference Manual* to check the other binding semantics that are available in Cg:

```
http://developer.download.nvidia.com/cg/Cg_3.1/Cg-3.1_April2012_
ReferenceManual.pdf
```

Using grab pass

In the *Adding transparency to PBR* recipe of *Chapter 4, Creating Test Cases and Writing Scenarios for Behavior Driven Development in Symfony*, we have seen how a material can be made transparent. Even if a transparent material can draw over a scene, it cannot change what has been drawn underneath it. This means that those Transparent Shaders cannot create distortions such as the ones typically seen in glass or water. In order to simulate them, we need to introduce another technique called **grab pass**. This allows us to access what has been drawn on screen so far so that a shader can use it (or alter it) with no restrictions. To learn how to use grab passes, we will create a material that grabs what's rendered behind it and draws it again on the screen. It's a shader that, paradoxically, uses several operations to show no changes at all.

Getting ready

This recipe requires the following operations:

1. Create a shader that we will initialize later.

2. Create a material to host the shader.

3. Attach the material to a flat piece of geometry, such as a quad. Place it in front of some other object so that you cannot see through it. The quad will appear transparent as soon as the shader is complete.

How to do it...

To use grab pass, you need to follow these steps:

1. Remove the `Properties` section; this shader will not use any of them.

2. In the `SubShader` section, add grab pass:

   ```
   GrabPass{   }
   ```

3. After the grab pass, we will need to add this extra pass:

   ```
   Pass {
       CGPROGRAM
   #pragma vertex vert
   #pragma fragment frag
   #include "UnityCG.cginc"

       sampler2D _GrabTexture;

       struct vertInput {
           float4 vertex : POSITION;
       };

       struct vertOutput {
           float4 vertex : POSITION;
           float4 uvgrab : TEXCOORD1;
       };

       // Vertex function
       vertInput vert(vertexInput v) {
           vertexOutput o;
           o.vertex = mul(UNITY_MATRIX_MVP, v.vertex);
           o.uvgrab = ComputeGrabScreenPos(o.vertex);
           return o;
       }
   ```

```
        // Fragment function
        half4 frag(vertexOutput i) : COLOR {
            fixed4 col = tex2Dproj(_GrabTexture, UNITY_PROJ_COORD(i.
uvgrab));
            return col + half4(0.5,0,0,0);
        }
    ENDCG
}
```

How it works...

This recipe not only introduces grab passes but also Vertex and Fragment Shaders; for this reason, we have to analyze the shader in detail.

So far, all the code has always been placed directly in the `SubShader` section. This is because our previous shaders required only a single pass. This time, two passes are required. The first one is the grab pass, which is defined simply by `GrabPass{}`. The rest of the code is placed in the second pass, which is contained in a `Pass` block.

The second pass is not structurally different from the shader shown in the first recipe of this chapter; we use the vertex function `vert` to get the position of the vertex and then we give it a color in the fragment function `frag`. The difference is that `vert` calculates another important detail: the UV data for the grab pass. The grab pass automatically creates a texture that can be referred to as follows:

```
sampler2D _GrabTexture;
```

In order to sample this texture, we need its UV data. The `ComputeGrabScreenPos` function returns data that we can use later to sample the grab texture correctly. This is done in the Fragment Shader using the following line:

```
fixed4 col = tex2Dproj(_GrabTexture, UNITY_PROJ_COORD(i.uvgrab));
```

This is the standard way in which a texture is grabbed and applied to the screen in its correct position. If everything has been done correctly, this shader will simply clone what has been rendered behind the geometry. We will see in the following recipes how this technique can be used to create materials such as water and glass.

There's more...

Every time you use a material with `GrabPass {}`, Unity will have to render the screen to a texture. This operation is very expensive and limits the number of grab passes that you can use in a game. Cg offers a slightly different variation:

```
GrabPass {"TextureName"}
```

This line not only allows you to give a name to the texture, but it also shares the texture with all the materials that have a grab pass called `TextureName`. This means that if you have ten materials, Unity will only do a single grab pass and share the texture to all of them. The main problem of this technique is that it doesn't allow effects that can be stacked. If you are creating a glass with this technique, you won't be able to have two glasses one after the other.

Implementing a Glass Shader

Glass is a very complicated material; it should not be a surprise that other chapters have already created shaders to simulate it in the *Adding transparency to PBR* recipe of *Chapter 4, Creating Test Cases and Writing Scenarios for Behavior Driven Development in Symfony*. However, there is an effect that transparency cannot reproduce deformations. Most glasses are not perfect, hence they create distortions when we look through them. This recipe will teach you how to do this. The idea behind this effect is to use a Vertex and Fragment Shader with a grab pass, and then sample the grab texture with a little change to its UV data to create a distortion. You can see the effect in the following image, using the glass-stained textures from the *Unity Standard Assets*:

Getting ready

The setup for this recipe is similar to the one presented in the previous chapter:

1. Create new Vertex and Fragment Shaders. You can start by copying the one used in the previous recipe, *Using grab pass*, as a base.
2. Create a material that will use the shader.
3. Assign the material to a quad or another flat geometry that will simulate your glass.
4. Place some objects behind it so that you can see the distortion effect.

How to do it...

Let's start by editing the Vertex and Fragment Shaders:

1. Add these two properties to the `Properties` block:

```
_MainTex("Base (RGB) Trans (A)", 2D) = "white" {}
_BumpMap("Noise text", 2D) = "bump" {}
_Magnitude("Magnitude", Range(0,1)) = 0.05
```

2. Add their variables in the second pass:

```
sampler2D _MainTex;
sampler2D _BumpMap;
float   _Magnitude;
```

3. Add the texture information in the input and output structures:

```
float2 texcoord : TEXCOORD0;
```

4. Transfer the UV data from the input to the output structure:

```
o.texcoord = v.texcoord;
```

5. Use the following fragment function:

```
half4 frag(vertOutput i) : COLOR {
        half4 mainColour = tex2D(_MainTex, i.texcoord);

    half4 bump = tex2D(_BumpMap, i.texcoord);
    half2 distortion = UnpackNormal(bump).rg;

    i.uvgrab.xy += distortion * _Magnitude;

    fixed4 col = tex2Dproj(_GrabTexture, UNITY_PROJ_COORD(i.
uvgrab));
    return col * mainColour * _Colour;
}
```

6. This material is transparent so it changes its tags in the `SubShader` block:

```
Tags{ "Queue" = "Transparent" "IgnoreProjector" = "True"
"RenderType" = "Opaque" }
```

7. What's left now is to set the texture for the glass and a normal map to displace the grab texture.

How it works...

The core that this shader uses is a grab pass to take what has already been rendered on the screen. The part where the distortion takes place is in the fragment function. Here, a normal map is unpacked and used to offset the UV data of the grab texture:

```
half4 bump = tex2D(_BumpMap, i.texcoord);
half2 distortion = UnpackNormal(bump).rg;

i.uvgrab.xy += distortion * _Magnitude;
```

The `_Magnitude` slide is used to determine how strong the effect is.

There's more...

This effect is very generic; it grabs the screen and creates a distortion based on a normal map. There is no reason why it shouldn't be used to simulate more interesting things. Many games use distortions around explosions or other sci-fi devices. This material can be applied to a sphere and, with a different normal map, it would simulate the heat wave of an explosion perfectly.

# Implementing a Water Shader for 2D games

The Glass Shader introduced in the previous recipe is static; its distortion never changes. It takes just a few changes to convert it to an animated material, making it perfect for 2D games, which feature water. This recipe uses a similar technique to the one shown in *Chapter 5, Animating Vertices in a Surface Shader*:

Getting ready

This recipe is based on the Vertex and Fragment Shaders described in the *Using grab pass* recipe as it will rely heavily on grab pass.

1. Create a new grab pass shader; you can write your own or start with the one presented in the *Using grab pass* recipe.

2. Create a new material for your shader.

3. Assign the material to a flat geometry that will represent your 2D water. In order for this effect to work, you should have something rendered behind it so that you can see the water-like displacement.

4. This recipe requires a noise texture, which is used to get pseudo-random values. It is important that you choose a seamless noise texture, such as the ones generated by tileable 2D Perlin noise, as shown in the following image. This ensures that when the material is applied to a large object, you will not see any discontinuity. In order for this effect to work, the texture has to be imported in the **Repeat** mode. If you want a smooth and continuous look for your water, you should also set it to **Bilinear** from **Inspector**. These settings ensure that the texture is sampled correctly from the shader:

How to do it...

To create this animated effect, you can start by refitting the shader. Follow these steps:

1. Add the following properties:

```
_NoiseTex("Noise text", 2D) = "white" {}
_Colour ("Colour", Color) = (1,1,1,1)

_Period ("Period", Range(0,50)) = 1
_Magnitude ("Magnitude", Range(0,0.5)) = 0.05
_Scale ("Scale", Range(0,10)) = 1
```

2. Add their respective variables to the second pass of the shader:

```
sampler2D _NoiseTex;
fixed4 _Colour;

float _Period;
float _Magnitude;
float _Scale;
```

3. Define the following output structure for the vertex function:

```
struct vertInput {
    float4 vertex : POSITION;
    fixed4 color : COLOR;
    float2 texcoord : TEXCOORD0;

    float4 worldPos : TEXCOORD1;
    float4 uvgrab : TEXCOORD2;
};
```

4. This shader needs to know the exact position of the space of every fragment. To do this, add the following line to the vertex function:

```
o.worldPos = mul(_Object2World, v.vertex);
```

5. Use the following fragment function:

```
fixed4 frag (vertInput i) : COLOR {
        float sinT = sin(_Time.w / _Period);
    float2 distortion = float2(
 tex2D(_NoiseTex, i.worldPos.xy / _Scale + float2(sinT, 0) ).r -
0.5,
     tex2D(_NoiseTex, i.worldPos.xy / _Scale + float2(0, sinT) ).r
- 0.5
    );
```

```
        i.uvgrab.xy += distortion * _Magnitude;

        fixed4 col = tex2Dproj( _GrabTexture, UNITY_PROJ_COORD(i.
    uvgrab));
        return col * _Colour;
    }
```

How it works...

This shader is very similar to the one introduced in the *Implementing a Glass Shader* recipe. The major difference is that this is an animated material; the displacement is not generated from a normal map but takes into account the current time in order to create a constant animation. The code that displaces the UV data of the grab texture seems quite complicated; let's try to understand how it has been generated. The idea behind it is that a sinusoid function is used to make the water oscillate. This effect needs to evolve over time; to achieve this effect, the distortion generated by the shader depends on the current time that is retrieved with the built-in variable, `_Time`. The `_Period` variable determines the period of the sinusoid, which means how fast the waves appear:

```
float2 distortion = float2( sin(_Time.w/_Period), sin(_Time.w/_Period)
) - 0.5;
```

The problem with this code is that you have the same displacement on the *X* and *Y* axes; as a result, the entire grab texture will rotate in a circular motion, which looks nothing like water. We obviously need to add some randomness to this.

The most common way to add random behaviors to shaders is by including a noise texture. The problem now is to find a way to sample the texture at seemingly random positions. The best way to avoid seeing an obvious sinusoid pattern is to use the sine waves as an offset in the UV data of the noise texture:

```
float sinT = sin(_Time.w / _Period);
float2 distortion = float2
(    tex2D(_NoiseTex, i.texcoord / _Scale + float2(sinT, 0) ).r - 0.5,
    tex2D(_NoiseTex, i.texcoord / _Scale + float2(0, sinT) ).r - 0.5
);
```

The `_Scale` variable determines the size of the waves. This solution is closer to the final version, but has a severe issue—if the water quad moves, the UV data follows it and you can see the water waves following the material rather than being anchored to the background. To solve this, we need to use the world position of the current fragment as the initial position for the UV data:

```
float sinT = sin(_Time.w / _Period);
float2 distortion = float2
(    tex2D(_NoiseTex, i.worldPos.xy / _Scale + float2(sinT, 0) ).r -
0.5,
    tex2D(_NoiseTex, i.worldPos.xy / _Scale + float2(0, sinT) ).r -
0.5
);
i.uvgrab.xy += distortion * _Magnitude;
```

The result is a pleasant, seamless distortion, which doesn't move in any clear direction.

> As it happens with all these special effects, there is no perfect solution. This recipe shows you a technique to create water-like distortion, but you are encouraged to play with it until you find an effect that fits the aesthetics of your game.

7
Mobile Shader Adjustment

In the next two chapters, we are going to take a look at making the shaders that we write performance-friendly for different platforms. We won't be talking about any one platform specifically, but we are going to break down the elements of shaders we can adjust to make them more optimized for mobiles and efficient on any platform in general. These techniques range from understanding what Unity offers in terms of built-in variables that reduce the overhead of the shaders memory to learning about ways in which we can make our own shader code more efficient. This chapter will cover the following recipes:

- ▶ What is a cheap shader
- ▶ Profiling your shaders
- ▶ Modifying our shaders for mobile

Introduction

Learning the art of optimizing your shaders will come up in just about any game project that you work on. There will always come a point in any production where a shader needs to be optimized, or maybe it needs to use less textures but produce the same effect. As a technical artist or shader programmer, you have to understand these core fundamentals to optimize your shaders so that you can increase the performance of your game while still achieving the same visual fidelity. Having this knowledge can also help in setting the way in which you write your shader from the start. For instance, by knowing that the game built using your shader will be played on a mobile device, we can automatically set all our lighting functions to use a half vector as the view direction or set all of our float variable types to fixed or half. These, and many other techniques, all contribute to your shaders running efficiently on your target hardware. Let's begin our journey and start learning how to optimize our shaders.

What is a cheap shader?

When first asked the question, what is a cheap shader, it might be a little tough to answer as there are many elements that go into making a more efficient shader. It could be the amount of memory used up by your variables. It could be the amount of textures the shader is using. It could also be that our shader is working fine, but we can actually produce the same visual effect with half the amount of data by reducing the amount of code we are using or data we are creating. We are going to explore a few of these techniques in this recipe and show how they can be combined to make your shader fast and efficient but still produce the high-quality visuals everyone expects from games today, whether on a mobile or PC.

Getting ready

In order to get this recipe started, we need to gather a few resources together. So let's perform the following tasks:

1. Create a new scene and fill it with a simple sphere object and single directional light.

2. Create a new shader and material and assign the shader to the material.

3. We then need to assign the material we just created to our sphere object in our new scene.

4. Finally, modify the shader so that it uses a diffuse texture and normal map and includes your own custom lighting function. The following image shows the result of modifying our default shader that we created in step 1:

```
Shader "Cookbook/Chapter08/OptimizedShader001"
{
    Properties
    {
        _MainTex ("Base (RGB)", 2D) = "white" {}
        _NormalMap ("Normal Map", 2D) = "bump" {}
    }

    SubShader
    {
        Tags { "RenderType"="Opaque" }
        LOD 200

        CGPROGRAM
        #pragma surface surf SimpleLambert

        sampler2D _MainTex;
        sampler2D _NormalMap;

        struct Input
        {
            float2 uv_MainTex;
            float2 uv_NormalMap;
        };

        inline float4 LightingSimpleLambert (SurfaceOutput s, float3 lightDir, float atten)
        {
            float diff = max (0, dot (s.Normal, lightDir));

            float4 c;
            c.rgb = s.Albedo * _LightColor0.rgb * (diff * atten * 2);
            c.a = s.Alpha;
            return c;
        }

        void surf (Input IN, inout SurfaceOutput o)
        {
            float4 c = tex2D (_MainTex, IN.uv_MainTex);

            o.Albedo = c.rgb;
            o.Alpha = c.a;
            o.Normal = UnpackNormal(tex2D(_NormalMap, IN.uv_NormalMap));
        }
        ENDCG
    }
    FallBack "Diffuse"
}
```

You should now have a setup similar to the following image. This setup will allow us to take a look at some of the basic concepts that go into optimizing shaders using Surface Shaders in Unity:

How to do it...

We are going to build a simple Diffuse shader to take a look at a few ways in which you can optimize your shaders in general.

First, we'll optimize our variable types so that they use less memory when they are processing data:

1. Let's begin with the `struct Input` in our shader. Currently, our UVs are being stored in a variable of the `float2` type. We need to change this to use `half2` instead:

    ```
    struct Input
    {
        half2 uv_MainTex;
        half2 uv_NormalMap;
    };
    ```

2. We can then move to our lighting function and reduce the variable's memory footprint by changing their types to the following:

```
inline fixed4 LightingSimpleLambert (SurfaceOutput s, fixed3 lightDir, fixed atten)
{
    fixed diff = max (0, dot (s.Normal, lightDir));

    fixed4 c;
    c.rgb = s.Albedo * _LightColor0.rgb * (diff * atten * 2);
    c.a = s.Alpha;
    return c;
}
```

3. Finally, we can complete this optimization pass by updating the variables in our `surf()` function:

```
void surf (Input IN, inout SurfaceOutput o)
{
    fixed4 c = tex2D (_MainTex, IN.uv_MainTex);

    o.Albedo = c.rgb;
    o.Alpha = c.a;
    o.Normal = UnpackNormal(tex2D(_NormalMap, IN.uv_NormalMap));
}
```

Now that we have our variables optimized, we are going to take advantage of a built-in lighting function variable so that we can control how lights are processed by this shader. By doing this, we can greatly reduce the amount of lights the shader processes. Modify the `#pragma` statement in your shader with the following code:

```
CGPROGRAM
#pragma surface surf SimpleLambert noforwardadd
```

We can optimize this further by sharing UVs between the normal map and diffuse texture. To do this, we simply change the UV lookup in our `UnpackNormal()` function to use `_MainTex` UVs instead of the UVs of `_NormalMap`:

```
void surf (Input IN, inout SurfaceOutput o)
{
    fixed4 c = tex2D (_MainTex, IN.uv_MainTex);

    o.Albedo = c.rgb;
    o.Alpha = c.a;
    o.Normal = UnpackNormal(tex2D(_NormalMap, IN.uv_MainTex));
}
```

4. As we have removed the need for the normal map UVs, we need to make sure that we remove the normal map UV code from the `Input` struct:

```
struct Input
{
    half2 uv_Diffuse;
};
```

5. Finally, we can further optimize this shader by telling the shader that it only works with certain renderers:

```
CGPROGRAM
#pragma surface surf SimpleLambert exclude_path:prepass noforwardadd
```

The result of our optimization passes show us that we really don't notice a difference in the visual quality, but we have reduced the amount of time it takes for this shader to be drawn to the screen. You will learn about finding out how much time it takes for a shader to render in the next recipe, but the idea to focus on here is that we achieve the same result with less data. So keep this in mind when creating your shaders. The following image shows us the final result of our shader:

How it works...

Now that we have seen the ways in which we can optimize our shaders, let's dive in a bit deeper and really understand why all of these techniques are working and look at a couple of other techniques that you can try for yourself.

Let's first focus our attention on the size of the data each of our variables is storing when we declare them. If you are familiar with programming, then you will understand that you can declare values or variables with different sizes of types. This means that a float actually has a maximum size in memory. The following description will describe these variable types in much more detail:

- **Float**: A float is a full 32-bit precision value and is the slowest of the three different types we see here. It also has its corresponding values of `float2`, `float3`, and `float4`.

- **Half**: The half variable type is a reduced 16-bit floating point value and is suitable to store UV values and color values and is much faster than using a float value. It has its corresponding values like the float type, which are `half2`, `half3`, and `half4`.

- **Fixed**: A fixed value is the smallest in size of the three types, but can be used for lighting calculations and colors and has the corresponding values of `fixed2`, `fixed3`, and `fixed4`.

Our second phase of optimizing our simple shader was to declare the `noforwardadd` value to our `#pragma` statement. This is basically a switch that automatically tells Unity that any object with this particular shader receives only per-pixel light from a single directional light. Any other lights that are calculated by this shader will be forced to be processed as per-vertex lights using Spherical Harmonic values produced internally by Unity. This is especially obvious when we place another light in the scene to light our sphere object because our shader is doing a per-pixel operation using the normal map.

This is great, but what if you wanted to have a bunch of directional lights in the scene and control over which of these lights is used for the main per-pixel light? Well, if you notice, each light has a **Render Mode** drop-down. If you click on this drop-down, you will see a couple of flags that can be set. These are **Auto**, **Important**, and **Not Important**. By selecting a light, you can tell Unity that a light should be considered more as a per-pixel light than a per-vertex light, by setting its render mode to **Important** and vice versa. If you leave a light set to **Auto**, then you will let Unity decide the best course of action.

Place another light in your scene and remove the texture that is currently in the main texture for our shader. You will notice that the second point light does not react with the normal map, only the directional light that we created first. The concept here is that you save on per-pixel operations by just calculating all extra lights as vertex lights, and save performance by just calculating the main directional light as a per-pixel light. The following image visually demonstrates this concept as the point light is not reacting with the normal map:

Effect of using noforwardadd in a shader's #pragam statement

Finally, we did a bit of cleaning up and simply told the normal map texture to use the main texture's UV values, and we got rid of the line of code that pulled in a separate set of UV values specifically for the normal map. This is always a nice way to simplify your code and clean up any unwanted data.

We also declared `exclude_pass: prepass` in our `#pragma` statement so that this shader wouldn't accept any custom lighting from the deferred renderer. This means that we can really use this shader effectively in the forward renderer only, which is set in the main camera's settings.

By taking a bit of time, you will be amazed at how much a shader can be optimized. You have seen how we can pack grayscale textures into a single RGBA texture as well as use lookup textures to fake lighting. There are many ways in which a shader can be optimized, which is why it is always an ambiguous question to ask in the first place, but knowing these different optimization techniques, you can cater your shaders to your game and target platform, ultimately resulting in very streamlined shaders and a nice steady framerate.

Profiling your shaders

Now that we know how we can reduce the overhead that our shaders might take, let's take a look at how to find problematic shaders in a scene where you might have a lot of shaders or a ton of objects, shaders, and scripts, all running at the same time. To find a single object or shader among a whole game can be quite daunting, but Unity provides us with its built-in Profiler. This allows us to actually see, on a frame-by-frame basis, what is happening in the game and each item being used by the GPU and CPU.

Using the Profiler, we can isolate items such as shaders, geometry, and general rendering items using its interface to create blocks of profiling jobs. We can filter out items till we are looking at the performance of just a single object. This then lets us see the effects on the CPU and GPU that the object has while it is performing its functions at runtime.

Let's take a look through the different sections of the Profiler and learn how to debug our scenes and, most importantly, our shaders.

Getting ready

Let's use our Profiler by getting a few assets ready and launching the Profiler window:

1. Let's use the scene from the last recipe and launch the Unity Profiler from **Window|Profiler** or *Ctrl + 7*.

2. Let's also duplicate our sphere a couple more times to see how that affects our rendering.

You should see something similar to the following image:
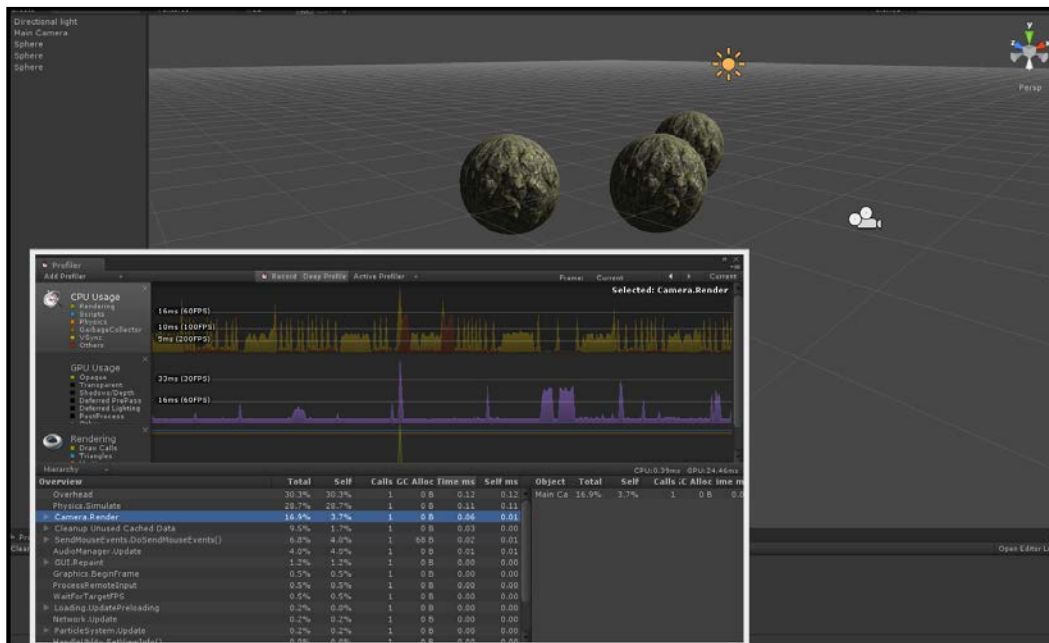
How to do it...

To use the Profiler, we will take a look at some of the UI elements of this window. Before we hit play, let's take a look at how to get the information we need from the profiler:

1. First, click on the larger blocks in the **Profiler** window called **GPU Usage**, **CPU Usage**, and **Rendering**. You will find these blocks on the left-hand side of the upper window:
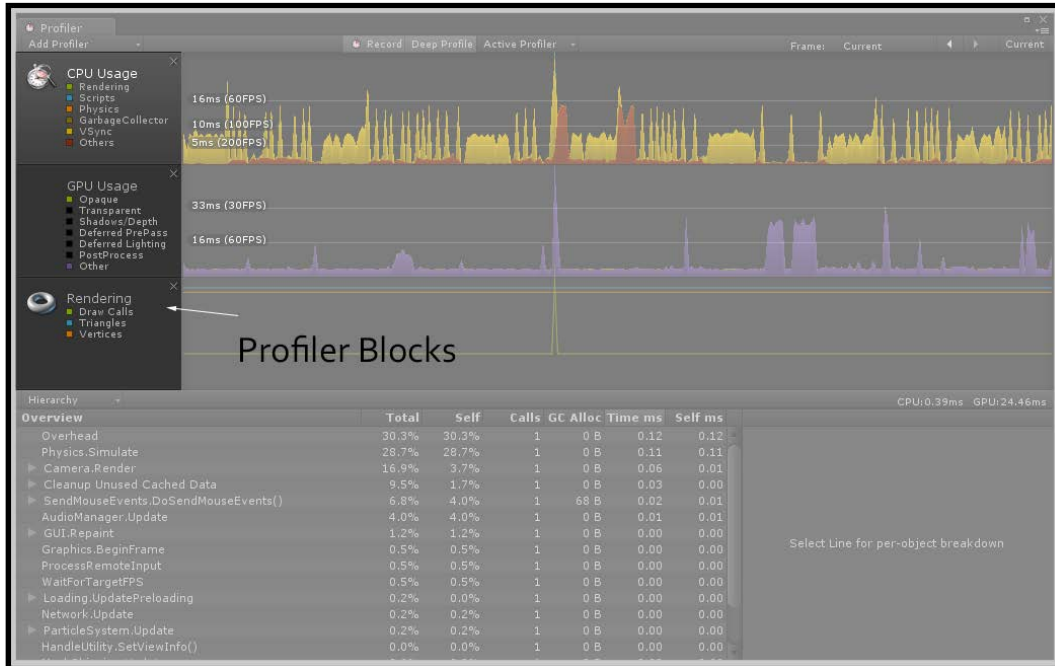

Using these blocks, we can see different data specific to those major functions of our game. The **CPU Usage** is showing us what most of our scripts are doing as well as physics and overall rendering. The **GPU Usage** block is giving us detailed information about the elements that are specific to our lighting, shadows, and render queues. Finally, the **Rendering** block is giving us information about the drawcalls and amount of geometry we have in our scene at any one frame.

By clicking on each of these blocks, we can isolate the type of data we see during our profiling session.

2. Now, click on the tiny colored blocks in one of these **Profile** blocks and hit play or *Ctrl + P* to run the scene.

This lets us dive down even deeper into our profiling session so that we can filter out what is being reported back for us. While the scene is running, uncheck all of the boxes, except for **Opaque** in the **GPU Usage** block. Notice that we can now see just how much time is being used to render the objects that are set to the Render Queue of Opaque:

3. Another great function of the **Profiler** window is the action of clicking and dragging in the graph view. This will automatically pause your game so that you can further analyze a certain spike in the graph to find out exactly which item is causing the performance problem. Click and drag around in the graph view to pause the game and see the effect of using this functionality:

4. Turning our attention now towards the lower half of the **Profiler** window, you will notice that there is a drop-down item available when we have the GPU Block selected. We can expand this to get even more detailed information about the current active profiling session and, in this case, more information about what the camera is currently rendering and how much time it is taking up:
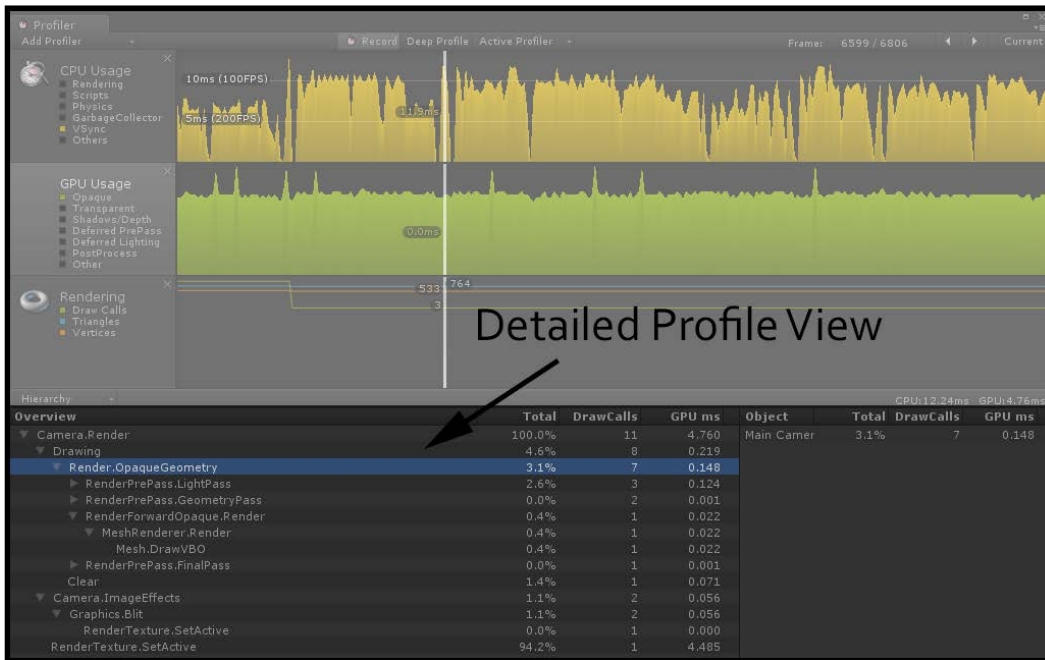
This gives us a complete look at the inner workings of what Unity is processing in this particular frame. In this case, we can see that our three spheres with our optimized shader are taking roughly 0.14 milliseconds to draw to the screen, they are taking up seven drawcalls, and this process is taking 3.1 percent of the GPU's time in every frame. It's this type of information we can use to diagnose and solve performance issues with regard to shaders. Let's conduct a test to see the effects of adding one more texture to our shader and blending two diffuse textures together using a `lerp` function. You will see the effects in the profiler pretty clearly.

5. Modify the `Properties` block of your shader with the following code to give us another texture to use:

```
Properties
{
    _MainTex ("Base (RGB)", 2D) = "white" {}
    _BlendTex ("Blend Texture", 2D) = "white" {}
    _NormalMap ("Normal Map", 2D) = "bump" {}
}
```

6. Then let's feed our texture to `CGPROGRAM`:

```
sampler2D _MainTex;
sampler2D _BlendTex;
sampler2D _NormalMap;
```

7. Now it's time to update our `surf()` function accordingly so that we blend our texture diffuse textures together:

```
void surf (Input IN, inout SurfaceOutput o)
{
    fixed4 c = tex2D (_MainTex, IN.uv_MainTex);
    fixed4 blendTex = tex2D (_BlendTex, IN.uv_MainTex);

    c = lerp(c, blendTex, blendTex.r);

    o.Albedo = c.rgb;
    o.Alpha = c.a;
    o.Normal = UnpackNormal(tex2D(_NormalMap, IN.uv_MainTex));
}
```

Once you save your modifications in your shader and return to Unity's editor, we can run our game and see the increase in milliseconds of our new shader. Press play once you have returned to Unity and let's take a look at the results in our **profiler**:
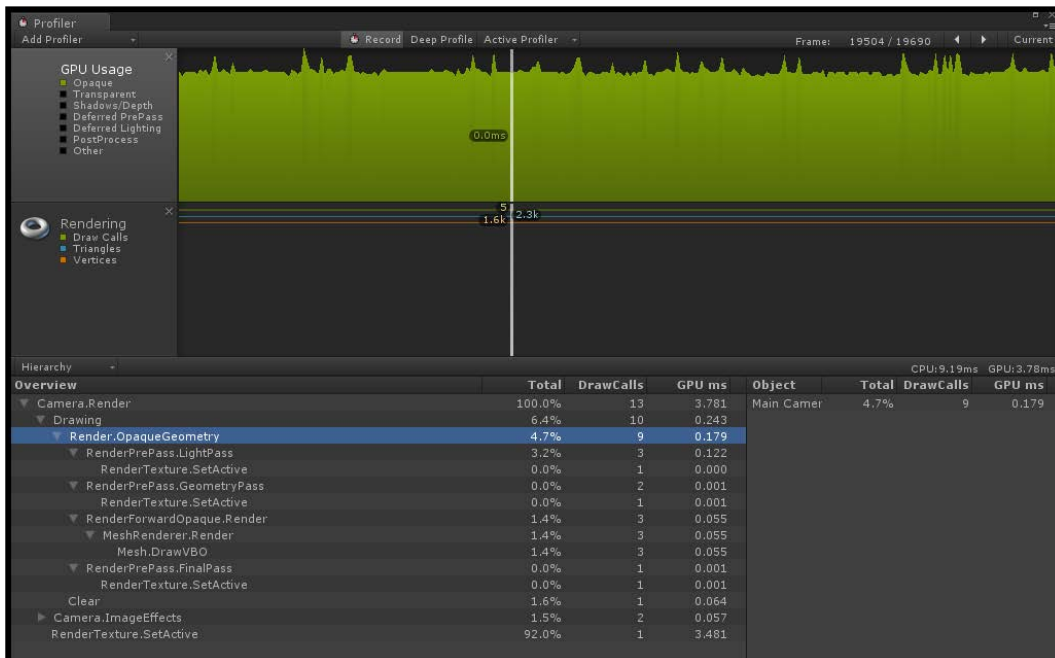
You can see now that the amount of time to render our Opaque Shaders in this scene is taking **0.179** milliseconds, up from 0.140 milliseconds. By adding another texture and using the `lerp()` function, we increased the render time for our spheres. While it's a small change, imagine having 20 shaders all working in different ways on different objects.

Using the information given here, you can pinpoint areas that are causing performance decreases more quickly and solve these issues using the techniques from the previous recipe.

How it works...

While it's completely out of scope of this book to describe how this tool actually works internally, we can surmise that Unity has given us a way to view the computer's performance while our game is running. Basically, this window is tied very tightly to the CPU and GPU to give us real-time feedback of how much time is being taken for each of our scripts, objects, and render queues. Using this information, we have seen that we can track the efficiency of our shader writing to eliminate problematic areas and code.

There's more...

It is also possible to profile specifically for mobile platforms. Unity provides us with a couple of extra features when the Android or IOS build target is set in the Build Settings. We can actually get real-time information from our mobile devices while the game is running. This becomes very useful because you are able to profile directly on the device itself instead of profiling directly in your editor. To find out more about this process, refer to Unity's documentation at the following link:

`http://docs.unity3d.com/Documentation/Manual/MobileProfiling.html`

# Modifying our shaders for mobile

Now that we have seen quite a broad set of techniques to make really optimized shaders, let's take a look at writing a nice, high-quality shader targeted for a mobile device. It is actually quite easy to make a few adjustments to the shaders we have written so that they run faster on a mobile device. This includes elements such as using the `approxview` or `halfasview` lighting function variables. We can also reduce the amount of textures we need and even apply better compression for the textures we are using. By the end of this recipe, we will have a nicely optimized normal-mapped, Specular shader for use in our mobile games.

Getting ready

Before we begin, let's get a fresh new scene and fill it with some objects to apply our Mobile shader:

1. Create a new scene and fill it with a default sphere and single directional light.
2. Create a new material and shader, and assign the shader to the material.
3. Finally, assign the material to our sphere object in our scene.

When completed, you should have a scene similar to the one in the following image:
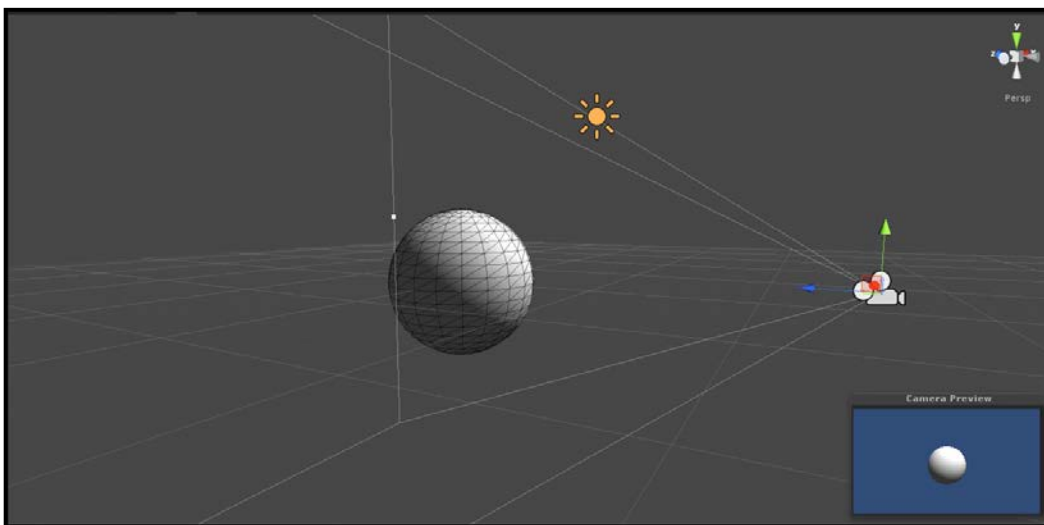
How to do it...

For this recipe, we will write a mobile-friendly shader from scratch and discuss the elements that make it more mobile-friendly:

1. Let's first populate our `Properties` block with the needed textures. In this case, we are going to use a single `Diffuse` texture with the gloss map in its alpha channel, normal map, and slider for specular intensity:

```
Properties
{
    _Diffuse ("Base (RGB) Specular Amount (A)", 2D) = "white" {}
    _SpecIntensity ("Specular Width", Range(0.01, 1)) = 0.5
    _NormalMap ("Normal Map", 2D) = "bump"{}
}
```

2. Our next task is to set up our `#pragma` declarations. This will simply turn certain features of the Surface Shader on and off, ultimately making the shader cheaper or more expensive:

```
CGPROGRAM
#pragma surface surf MobileBlinnPhong exclude_path:prepass nolightmap noforwardadd halfasview
```

3. We then need to make the connection between our `Properties` block and `CGPROGRAM`. This time, we are going to use the fixed variable type for our specular intensity slider to reduce its memory usage:

```
sampler2D _Diffuse;
sampler2D _NormalMap;
fixed _SpecIntensity;
```

4. In order for us to map our textures to the surface of our object, we need to get some UVs. In this case, we are going to get only one set of UVs to keep the amount of data in our shader down to a minimum:

```
struct Input
{
    half2 uv_Diffuse;
};
```

5. The next step is to fill in our lighting function using a few new input variables that are available to us using the new `#pragma` declarations:

```
inline fixed4 LightingMobileBlinnPhong (SurfaceOutput s, fixed3 lightDir, fixed3 halfDir, fixed atten)
{
    fixed diff = max (0, dot (s.Normal, lightDir));
    fixed nh = max (0, dot (s.Normal, halfDir));
    fixed spec = pow (nh, s.Specular*128) * s.Gloss;

    fixed4 c;
    c.rgb = (s.Albedo * _LightColor0.rgb * diff + _LightColor0.rgb * spec) * (atten*2);
    c.a = 0.0;
    return c;
}
```

6. Finally, we complete the shader by creating the `surf()` function and processing the final color of our surface:

```
void surf (Input IN, inout SurfaceOutput o)
{
    fixed4 diffuseTex = tex2D (_Diffuse, IN.uv_Diffuse);
    o.Albedo = diffuseTex.rgb;
    o.Gloss = diffuseTex.a;
    o.Alpha = 0.0;
    o.Specular = _SpecIntensity;
    o.Normal = UnpackNormal(tex2D(_NormalMap, IN.uv_Diffuse));
}
```

When completed with the code portion of this recipe, save your shader and return to the Unity editor to let the shader compile. If no errors occurred, you should see a result similar to the following image:

How it works...

So, let's begin the description of this shader by explaining what it does and doesn't do. First, it excludes the deferred lighting pass. This means that if you created a lighting function that was connected to the deferred renderer's prepass, it wouldn't use that particular lighting function and would look for the default lighting function like the ones that we have been creating thus far in this book.

This particular shader does not support **Lightmapping** by Unity's internal light-mapping system. This just keeps the shader from trying to find light maps for the object that the shader is attached to, making the shader more performance friendly because it is not having to perform the lightmapping check.

We included the `noforwardadd` declaration so that we process only per-pixel textures with a single directional light. All other lights are forced to become per-vertex lights and will not be included in any per-pixel operations you might do in the `surf()` function.

Finally, we are using the `halfasview` declaration to tell Unity that we aren't going to use the main `viewDir` parameter found in a normal lighting function. Instead, we are going to use the half vector as the view direction and process our specular with this. This becomes much faster for the shader to process as it will be done on a per-vertex basis. It isn't completely accurate when it comes to simulating specular in the real world, but visually on a mobile device, it looks just fine and the shader is more optimized.

Its techniques like these that make a shader more efficient and cleaner, codewise. Always make sure that you are using only the data you need while weighing this against your target hardware and the visual quality that the game requires. In the end, it becomes a cocktail of these techniques that ultimately make up your shaders for your games.

8
Screen Effects with Unity Render Textures

In this chapter, you will learn the following recipes:

- ▶ Setting up the screen effects script system
- ▶ Using brightness, saturation, and contrast with screen effects
- ▶ Using basic Photoshop-like Blend modes with screen effects
- ▶ Using the Overlay Blend mode with screen effects

Introduction

One of the most impressive aspects of learning to write shaders is the process of creating your own screen effects, also known as post effects. With these screen effects, we can create stunning real-time images with Bloom, Motion Blur, HDR effects, and so on. Most modern games out in the market today make heavy use of these Screen effects for their depth of field effects, bloom effects, and even color correction effects.

Throughout this chapter, you will learn how to build up the script system that gives us the control to create these screen effects. We will cover Render Textures, what the depth buffer is, and how to create effects that give you Photoshop-like control over the final rendered image of your game. By utilizing screen effects for your games, you not only round out your shader writing knowledge, but you will also have the power to create your own incredible real-time renders with Unity.

Setting up the screen effects script system

The process of creating screen effects is one in which we grab a fullscreen image (or texture), use a shader to process its pixels on the GPU, and then send it back to Unity's renderer to apply it to the whole rendered image of the game. This allows us to perform per-pixel operations on the rendered image of the game in real time, giving us a more global artistic control.

Imagine if you had to go through and adjust each material on each object in your game to just adjust the contrast of the final look of your game. While not impossible, this would take a bit of labor to perform. By utilizing a screen effect, we can adjust the screen's final look as a whole, thereby giving us a more Photoshop-like control over our game's final appearance.

In order to get a Screen effect system up and running, we have to set up a single script to act as the courier of the game's current rendered image or, what Unity calls, the Render Texture. By utilizing this script to pass the Render Texture to a shader, we can create a flexible system to create screen effects. For our first screen effect, we are going to create a very simple grayscale effect, where we can make our game look black and white. Let's take a look at how this is done.

Getting ready

In order to get our Screen Effects system up and running, we need to create a few assets for our current Unity project. By doing this, we will set ourselves up for the steps in the following sections:

1. In the current project, we need to create a new C# script and call it `TestRenderImage.cs`.

2. Create a new shader and call it `ImageEffect.shader`.

3. Create a simple sphere in the scene and assign it a new material. This new material can be anything, but for our example, we will make a simple red, specular material.

4. Finally, create a new directional light and save the scene.

With all of our assets ready, you should have a simple scene setup, which looks similar to the following image:

How to do it...

In order to make our grayscale screen effect work, we need a script and shader. So, we will complete these two new items here and fill them in with the appropriate code to produce our first screen effect. Our first task is to complete the C# script. This will get the whole system running. After this, we will complete the shader and see the results of our Screen Effect. Let's complete our script and shader with the following steps:

1. Open the `TestRenderImage.cs` C# script and let's begin by entering a few variables that we will need to store important objects and data. Enter the following code at the very top of the `TestRenderImage` class:

    ```
    public class TestRenderImage : MonoBehaviour
    {
        #region Variables
        public Shader curShader;
        public float grayScaleAmount = 1.0f;
        private Material curMaterial;
        #endregion
    ```

2. In order for us to edit the Screen Effect in real time, when the Unity editor isn't playing, we need to enter the following line of code just above the declaration of the `TestRenderImage` class:

```
[ExecuteInEditMode]
public class TestRenderImage : MonoBehaviour
{
```

3. As our Screen Effect is using a shader to perform the pixel operations on our Screen image, we have to create a material to run the shader. Without this, we can't access the properties of the shader. For this, we will create a C# property to check for a material, and create one if it doesn't find one. Enter the following code just after the declaration of the variables from step 1:

```
#region Properties
Material material
{
    get
    {
        if(curMaterial == null)
        {
            curMaterial = new Material(curShader);
            curMaterial.hideFlags = HideFlags.HideAndDontSave;
        }
        return curMaterial;
    }
}
#endregion
```

4. We now want to set up some checks in our script to see if the current target platform that we are building the Unity game on actually supports image effects. If it doesn't find anything at the start of this script, then the script will disable itself:

```
void Start()
{
    if(!SystemInfo.supportsImageEffects)
    {
        enabled = false;
        return;
    }

    if(!curShader && !curShader.isSupported)
    {
        enabled = false;
    }
}
```

5. To actually grab the Rendered Image from the Unity Renderer, we need to make use of the following built-in function that Unity provides us, called `OnRenderImage()`. Enter the following code so that we can have access to the current Render Texture:

```
void OnRenderImage(RenderTexture sourceTexture, RenderTexture destTexture)
{
    if(curShader != null)
    {
        material.SetFloat("_LuminosityAmount", grayScaleAmount);
        Graphics.Blit(sourceTexture, destTexture, material);
    }
    else
    {
        Graphics.Blit(sourceTexture, destTexture);
    }
}
```

6. Our Screen effect has a variable called `grayScaleAmount` with which we can control how much grayscale we want for our final Screen Effect. So, in this case, we need to make the value go from 0 – 1, where 0 is no grayscale effect and 1 is full grayscale effect. We will perform this operation in the `Update()` function so that it sets every frame this script is running:

```
void Update()
{
    grayScaleAmount = Mathf.Clamp(grayScaleAmount, 0.0f, 1.0f);
}
```

7. Finally, we complete our script by doing a little bit of clean up on objects we created when the script started:

```
void OnDisable()
{
    if(curMaterial)
    {
        DestroyImmediate(curMaterial);
    }
}
```

At this point, we can now apply this script to the camera, if it compiled without errors, in Unity. Let's apply the `TestRenderImage.cs` script to our main camera in our scene. You should see the `grayScaleAmount` value and a field for a shader, but the script throws an error to the console window. It says that it is missing an instance to an object and so won't process appropriately. If you recall from step 4, we are doing some checks to see whether we have a shader and the current platform supports the shader. As we haven't given the Screen Effect script a shader to work with, then the `curShader` variable is just null, which throws the error. Let's continue our Screen Effects system by completing the shader.

8. To begin our shader, we will populate our properties with some variables so that we can send data to this shader:

```
Properties
{
    _MainTex ("Base (RGB)", 2D) = "white" {}
    _LuminosityAmount ("GrayScale Amount", Range(0.0, 1)) = 1.0
}
```

9. Our shader is now going to utilize pure CG shader code instead of utilizing Unity's built-in Surface Shader code. This will make our Screen Effect more optimized as we need to work only with the pixels of the Render Texture. So, we will create a new `Pass` block in our shader and fill it with some new `#pragma` statements that we haven't seen before:

```
SubShader
{
    Pass
    {
        CGPROGRAM
        #pragma vertex vert_img
        #pragma fragment frag
        #pragma fragmentoption ARB_precision_hint_fastest
        #include "UnityCG.cginc"
```

10. In order to access the data being sent to the shader from the Unity editor, we need to create the corresponding variables in our `CGPROGRAM`:

```
        uniform sampler2D _MainTex;
        fixed _LuminosityAmount;
```

11. Finally, all we need to do is set up our pixel function, in this case, called `frag()`. This is where the meat of the Screen Effect is. This function will process each pixel of the Render Texture and return a new image to our `TestRenderImage.cs` script:

```
fixed4 frag(v2f_img i) : COLOR
{
    //Get the colors from the RenderTexture and the uv's
    //from the v2f_img struct
    fixed4 renderTex = tex2D(_MainTex, i.uv);

    //Apply the Luminosity values to our render texture
    float luminosity = 0.299 * renderTex.r + 0.587 * renderTex.g + 0.114 * renderTex.b;
    fixed4 finalColor = lerp(renderTex, luminosity, _LuminosityAmount);


    return finalColor;
}
```

Once the shader is complete, return to Unity and let it compile to see if any errors occurred. If not, assign the new shader to the `TestRenderImage.cs` script and change the value of the grayscale amount variable. You should see the game view go from a colored version of the game to a grayscale version of the game. The following image demonstrates this Screen Effect:

Grayscale Value = 0.0   Grayscale Value = 0.5   Grayscale Value = 1.0

With this complete, we now have an easy way to test out new Screen Effect shaders without having to write our whole Screen Effect system over and over again. Let's dive in a little deeper and learn about what's going on with the Render Texture and how it is processed throughout its existence.

How it works...

To get a screen effect up and running inside of Unity, we need to create a script and shader. The script drives the real-time update in the editor and is also responsible for capturing the Render Texture from the main camera and passing it to the shader. Once the render texture gets to the shader, we can use the shader to perform per-pixel operations.

At the start of the script, we perform a few checks to make sure that the current selected build platform actually supports screen effects and the shader itself. There are instances where a current platform will not support Screen Effects or the shader that we are using. So the checks that we do in the `Start()` function make sure we don't get any errors if the platform doesn't support the screen system.

Once the script passes these checks, we initiate the Screen Effects system by calling the built-in function, `OnRenderImage()`. This function is responsible for grabbing the renderTexture, giving it to the shader using the `Graphics.Blit()` function, and returning the processed image to the Unity renderer. You can find more information on these two functions at the following URLs:

- **OnRenderImage:** `http://docs.unity3d.com/Documentation/ScriptReference/MonoBehaviour.OnRenderImage.html`
- **Graphics.Blit:** `http://docs.unity3d.com/Documentation/ScriptReference/Graphics.Blit.html`

Once the current render texture reaches the shader, the shader takes it, processes it through the `frag()` function, and returns the final color for each pixel.

You can see how powerful this becomes as it gives us Photoshop-like control over the final rendered image of our game. These screen effects work sequentially like Photoshop layers in the camera. When you place these screen effects one after the other, they will be processed in that order. These are just the bare bones steps to get a screen effect working, but it is the core of how the screen effects system works.

There's more...

Now that we have our simple Screen Effect system up and running, let's take a look at some of the other useful information we can obtain from Unity's renderer:

Scene Depth

We can actually get the depth of everything in our current game by turning on Unity's built-in Depth mode. Once this is turned on, we can use the depth information for a ton of different effects. Let's take a look at how this is done:

1. Create a new shader and call it `SceneDepth_Effect`. Then double-click on this shader to open it in the **MonoDevelop** editor.

2. We will create the Main Texture property and a property to control the power of the scene depth effect. Enter the following code in your shader:

```
Properties
{
    _MainTex ("Base (RGB)", 2D) = "white" {}
    _DepthPower ("Depth Power", Range(1, 5)) = 1
}
```

3. Now we need to create the corresponding variables in our CGPROGRAM. We are going to add one more variable called _CameraDepthTexture. This is a built-in variable that Unity has provided us with through the use of the UnityCG cginclude file. It gives us the depth information from the camera:

```
Pass
{
    CGPROGRAM
    #pragma vertex vert_img
    #pragma fragment frag
    #pragma fragmentoption ARB_precision_hint_fastest
    #include "UnityCG.cginc"

    uniform sampler2D _MainTex;
    fixed _DepthPower;
    sampler2D _CameraDepthTexture;
```

4. We will complete our depth shader by utilizing a couple of built-in functions that Unity provides us with, the UNITY_SAMPLE_DEPTH() and linear01Depth() functions. The first function actually gets the depth information from our _CameraDepthTexture and produces a single float value for each pixel. The Linear01Depth() function then makes sure that the values are within the 0-1 range by taking this final depth value to a power we can control, where the mid-value on the 0-1 range sits in the scene based off of the camera position:

```
fixed4 frag(v2f_img i) : COLOR
{
    //Get the colors from the RenderTexture and the uv's
    //from the v2f_img struct
    float d = UNITY_SAMPLE_DEPTH( tex2D(_CameraDepthTexture, i.uv.xy) );
    d = pow(Linear01Depth(d), _DepthPower);

    return d;
}
```

5. With our shader complete, let's turn our attention to our Screen Effects script. We need to add the depthPower variable to the script so that we can let users change the value in the editor:

```
#region Variables
public Shader curShader;
private Material curMaterial;

public float depthPower = 1.0f;
#endregion
```

6. Our `OnRenderImage()` function then needs to be updated so that it is passing the right value to our shader:

```
void OnRenderImage(RenderTexture sourceTexture, RenderTexture destTexture)
{
    if(curShader != null)
    {
        material.SetFloat("_DepthPower", depthPower);
        Graphics.Blit(sourceTexture, destTexture, material);
    }
    else
    {
        Graphics.Blit(sourceTexture, destTexture);
    }
}
```

7. To complete our depth Screen effect, we need to tell Unity to turn on the depth rendering in the current camera. This is done by simply setting the main camera's `depthTextureMode`:

```
void Update()
{
    Camera.main.depthTextureMode = DepthTextureMode.Depth;
    depthPower = Mathf.Clamp(depthPower, 0, 5);
}
```

With all the code set up, save your script and shader and return to Unity to let them both compile. If no errors are encountered, you should see a result similar to the following image:

Using brightness, saturation, and contrast with screen effects

Now that we have our screen effects system up and running, we can explore how to create more involved pixel operations to perform some of the more common Screen Effects found in games today.

To begin, using a screen effect to adjust the overall final colors of your game is crucial in giving artists a global control over the final look of the game. Techniques such as color adjustment sliders to adjust the intensity for the reds, blues, and greens of the final rendered game or techniques like putting a certain tone of color over the whole screen as seen in something like a sepia film effect.

For this particular recipe, we are going to cover some of the more core color adjustment operations we can perform on an image. These are brightness, saturation, and contrast. Learning how to code these color adjustments gives us a nice base to learn the art of screen effects.

Getting ready

We will need to create a couple of new assets. We can utilize the same scene as our test scene, but we will need a new script and shader:

1. Create a new script and call it `BSC_ImageEffect`.

2. Create a new shader called `BSC_Effect`.

3. Now we simply need to copy the code we had from the C# script in the previous recipe to our new C# script. This will allow us to focus just on the mathematics for the brightness, saturation, and contrast effect.

4. Copy the code from the shader in the previous recipe to our new shader.

5. Create a couple of new objects in the scene, set up some different colored diffuse materials, and randomly assign them to the new objects in the scene. This will give us a good range of colors to test with our new screen effect.

When completed, you should have a scene similar to the following image:
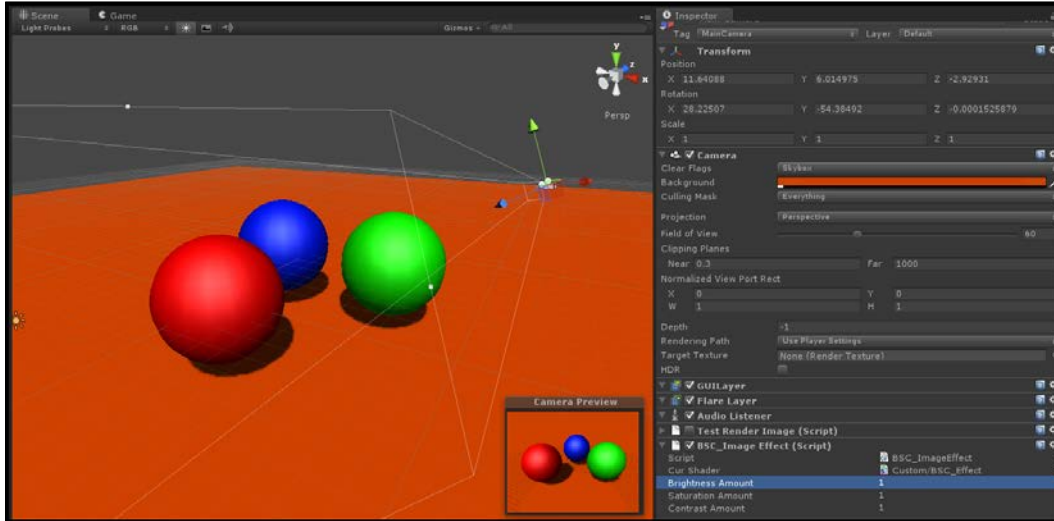
How to do it...

Now that we have completed our scene setup and created our new script and shader, we can begin to fill in the code necessary to achieve the brightness, saturation, and contrast Screen Effect. We will be focusing on just the pixel operation and variable setup for our script and shader, as getting a Screen Effect system up and running is described in the *Setting up the screen effects script system* recipe:

1. Let's begin by launching our new shader and script in MonoDevelop. Simply double-click on the two files in the project view to perform this action.

2. Editing the shader first makes more sense so that we know what kind of variables we will need for our C# script. Let's begin this by entering the appropriate properties for our brightness, saturation, and contrast effect. Remember, we need to keep the `_MainTex` property in our shader as this is the property that the RenderTexture targets when creating Screen Effects:

```
Properties
{
    _MainTex ("Base (RGB)", 2D) = "white" {}
    _BrightnessAmount ("Brightness Amount", Range(0.0, 1)) = 1.0
    _satAmount ("Saturation Amount", Range(0.0, 1)) = 1.0
    _conAmount ("Contrast Amount", Range(0.0, 1)) = 1.0
}
```

3. As usual, in order for us to access the data coming in from our properties in our `CGPROGRAM`, we need to create the corresponding variables in the `CGPROGRAM`:

```
Pass
{
    CGPROGRAM
    #pragma vertex vert_img
    #pragma fragment frag
    #pragma fragmentoption ARB_precision_hint_fastest
    #include "UnityCG.cginc"

    uniform sampler2D _MainTex;
    fixed _BrightnessAmount;
    fixed _satAmount;
    fixed _conAmount;
```

4. Now we need to create the operations that will perform the brightness, saturation, and contrast effects. Enter the following new function in our shader, just above the `frag()` function. Don't worry if it doesn't make sense just yet; all the code will be explained in the next recipe:

```
float3 ContrastSaturationBrightness(float3 color, float brt, float sat, float con)
{
    // Increase or decrease theese values to
    //adjust r, g and b color channels seperately
    float AvgLumR = 0.5;
    float AvgLumG = 0.5;
    float AvgLumB = 0.5;

    //Luminance coefficients for getting lumoinance from the image
    float3 LuminanceCoeff = float3(0.2125, 0.7154, 0.0721);

    //Operation for brightness
    float3 AvgLumin = float3(AvgLumR, AvgLumG, AvgLumB);
    float3 brtColor = color * brt;
    float intensityf = dot(brtColor, LuminanceCoeff);
    float3 intensity = float3(intensityf, intensityf, intensityf);

    //Operation for Saturation
    float3 satColor = lerp(intensity, brtColor, sat);

    //Operation for Contrast
    float3 conColor = lerp(AvgLumin, satColor, con);
    return conColor;
}
```

5. Finally, we just need to update our `frag()` function to actually use the `ContrastSaturationBrightness()` function. This will process all the pixels of our Render Texture and pass it back to our script:

```
fixed4 frag(v2f_img i) : COLOR
{
    //Get the colors from the RenderTexture and the uv's
    //from the v2f_img struct
    fixed4 renderTex = tex2D(_MainTex, i.uv);

    //Apply the Brughtness, saturation, contrast operations
    renderTex.rgb = ContrastSaturationBrightness(renderTex.rgb,
                                                  _BrightnessAmount,
                                                  _satAmount,
                                                  _conAmount);

    return renderTex;
}
```

With the code entered in the shader, return to the Unity editor to let the new shader compile. If there are no errors, we can return to MonoDevelop to work on our script. Let's begin this by creating a couple of new lines of code that will send the proper data to our shader:

1. Our first step in modifying our script is to add the proper variables that will drive the values of our Screen Effect. In this case, we will need a slider for brightness, a slider for saturation, and a slider for contrast:

```
#region Variables
public Shader curShader;
public float brightnessAmount = 1.0f;
public float saturationAmount = 1.0f;
public float contrastAmount = 1.0f;
private Material curMaterial;
#endregion
```

2. With our variables set up, we now need to tell the script to pass their data to the shader. We do this in the `OnRenderImage()` function:

```
void OnRenderImage(RenderTexture sourceTexture, RenderTexture destTexture)
{
    if(curShader != null)
    {
        material.SetFloat("_BrightnessAmount", brightnessAmount);
        material.SetFloat("_satAmount", saturationAmount);
        material.SetFloat("_conAmount", contrastAmount);

        Graphics.Blit(sourceTexture, destTexture, material);
    }
    else
    {
        Graphics.Blit(sourceTexture, destTexture);
    }
}
```

3. Finally, all we need to do is clamp the values of the variables within a range that is reasonable. These clamp values are entirely preferential, so you can use whichever values you see fit:

```
void Update()
{
    brightnessAmount = Mathf.Clamp(brightnessAmount, 0.0f, 2.0f);
    saturationAmount = Mathf.Clamp(saturationAmount, 0.0f, 2.0f);
    contrastAmount = Mathf.Clamp(contrastAmount, 0.0f, 3.0f);
}
```

With the script completed and shader finished, we simply assign our script to our main camera and our shader to the script, and you should see the effects of brightness, saturation, and contrast by manipulating the slider values. The following image shows a result you can achieve with this screen effect:
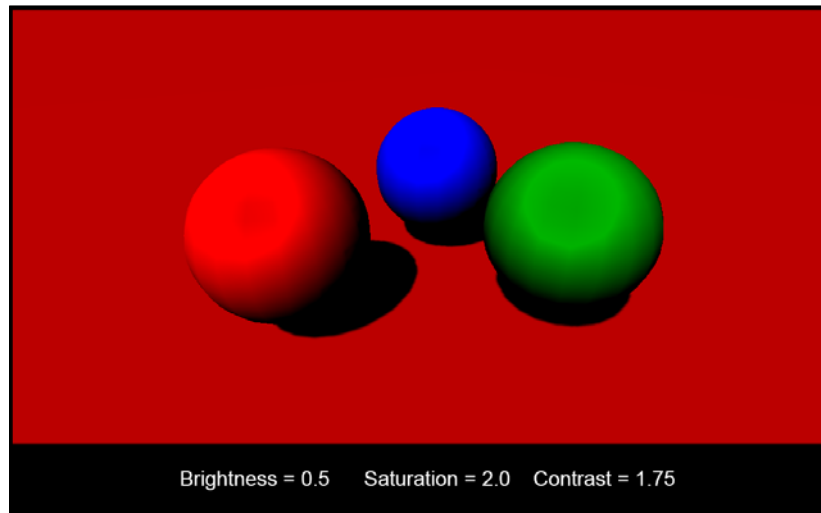
The following image shows another example of what can be done by adjusting the colors of the render image:

How it works...

Since we now know how the basic Screen Effects system works, let's just cover the per-pixel operations we created in the `ContrastSaturationBrightness()` function.

The function starts by taking a few arguments. The first and most important is the current render texture. The other arguments simply adjust the overall effect of the screen effect and are represented by sliders in the screen effects' **Inspector** tab. Once the function receives the render texture and the adjustment values, it declares a few constant values that we use to modify and compare against the original render texture.

The `luminanceCoeff` variable stores the values that will give us the overall brightness of the current image. These coefficients are based on the CIE color matching functions and are pretty standard throughout the industry. We can find the overall brightness of the image by getting the dot product of the current image dotted with these luminance coefficients. Once we have the brightness, we simply use a couple of `lerp` functions to blend from the grayscale version of the brightness operation and the original image multiplied by the brightness value, being passed into the function.

The screen effects, like this one, are crucial to achieve high-quality graphics for your games as it lets you tweak the final look of your game without having to edit each material in your current game scene.

Using basic Photoshop-like Blend modes with screen effects

The screen effects aren't just limited to adjusting the colors of a rendered image from our game. We can also use them to combine other images with our Render Texture. This technique is no different than creating a new layer in Photoshop and choosing a blend mode to blend two images together or, in our case, a texture with a Render Texture. This becomes a very powerful technique as it gives the artists in a production environment a way to simulate their blending modes in the game rather than just in Photoshop.

For this particular recipe, we are going to take a look at some of the more common blend modes, such as **Multiply**, **Add**, and **Overlay**. You will see how simple it is to have the power of Photoshop Blend modes in your game.

Getting ready

To begin, we have to get our assets ready. So let's follow the next few steps to get our screen effects system up and running for our new Blend mode screen effect:

1. Create a new script and call it `BlendMode_ImageEffect`.
2. Create a new shader called `BlendMode_Effect`.
3. Now we simply need to copy the code we had from the C# script in the first recipe of this chapter to our new C# script. This will allow us to focus on just the mathematics for the brightness, saturation, and contrast effect.
4. Copy the code from the shader in the first recipe in this chapter to our new shader.
5. Finally, we will need another texture to perform our blend mode effect. In this recipe, we will use a grunge type texture. This will make the effect very obvious when we are testing it out.

The following image is the grunge map used in the making of this effect. Finding a texture with enough detail and a nice range of grayscale values will make for a nice texture to test our new effect:

How to do it...

Our first blend mode that we will implement is the **Multiply** blend mode as seen in Photoshop. Let's begin by modifying the code in our shader first.

1. Launch the shader in **MonoDevelop** by double-clicking on it in Unity's project view.

2. We need to add some new properties so that we have a texture to blend with and a slider for an opacity value. Enter the following code in your new shader:

```
Properties
{
    _MainTex ("Base (RGB)", 2D) = "white" {}
    _BlendTex ("Blend Texture", 2D) = "white"{}
    _Opacity ("Blend Opacity", Range(0,1)) = 1
}
```

3. Enter the corresponding variables in our CGPROGRAM so that we can access the data from our Properties block:

```
Pass
{
    CGPROGRAM
    #pragma vertex vert_img
    #pragma fragment frag
    #pragma fragmentoption ARB_precision_hint_fastest
    #include "UnityCG.cginc"

    uniform sampler2D _MainTex;
    uniform sampler2D _BlendTex;
    fixed _Opacity;
```

4. Finally, we modify our frag() function so that it performs the multiply operation on our two textures:

```
fixed4 frag(v2f_img i) : COLOR
{
    //Get the colors from the RenderTexture and the uv's
    //from the v2f_img struct
    fixed4 renderTex = tex2D(_MainTex, i.uv);
    fixed4 blendTex = tex2D(_BlendTex, i.uv);

    //Perform a multiply Blend mode
    fixed4 blendedMultiply = renderTex * blendTex;

    //Adjust amount of Blend Mode with a lerp
    renderTex = lerp(renderTex, blendedMultiply, _Opacity);

    return renderTex;
}
```

5. Save the shader and return to the Unity editor to let the new shader code compile and check for errors. If no errors occurred, then double-click on the C# script file to launch it in the MonoDevelop editor.

6. In our script file as well, we need to create the corresponding variables. So we will need a texture so that we can assign one to the shader and a slider to adjust the final amount of the blend mode we want to use:

```
#region Variables
public Shader curShader;
public Texture2D blendTexture;
public float blendOpacity = 1.0f;
private Material curMaterial;
#endregion
```

7. We then need to send our variable data to the shader through the `OnRenderImage()` function:

```
void OnRenderImage(RenderTexture sourceTexture, RenderTexture destTexture)
{
    if(curShader != null)
    {
        material.SetTexture("_BlendTex", blendTexture);
        material.SetFloat("_Opacity", blendOpacity);

        Graphics.Blit(sourceTexture, destTexture, material);
    }
    else
    {
        Graphics.Blit(sourceTexture, destTexture);
    }
}
```

8. To complete the script, we simply fill in our `Update()` function so that we can clamp the value of the `blendOpacity` variable between a value of `0.0` and `1.0`:

```
void Update()
{
    blendOpacity = Mathf.Clamp(blendOpacity, 0.0f, 1.0f);
}
```

With this complete, we assign the screen effect script to our main camera and our screen effect shader to our script so that it has a shader to use for the per-pixel operations. Finally, in order for the effect to be fully functional, the script and shader is looking for a texture. You can assign any texture to the texture field in the **Inspector** for the screen effect script. Once this texture is in place, you will see the effect of multiplying this texture over the game's rendered image. The following image demonstrates the screen effect:
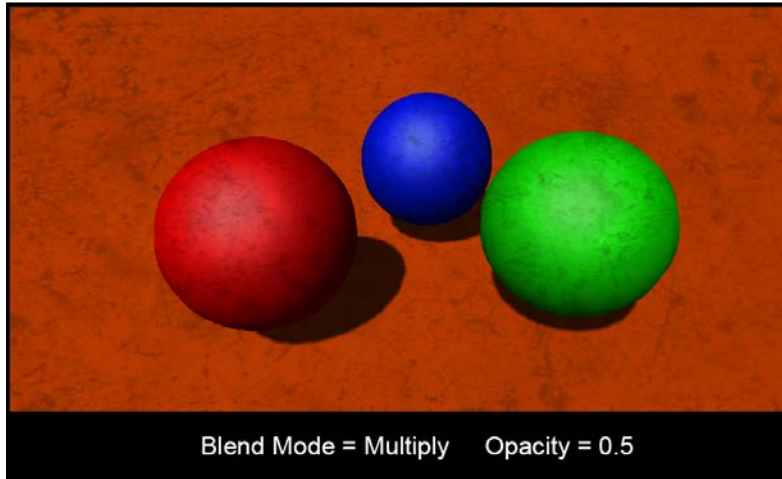
The following image demonstrates a higher intensity of opacity, making the multiplied image much more apparent over our render image:

With our first blend mode set up, we can begin to add a couple of simpler blend modes to get a better understanding of how easy it is to add more effects and really fine-tune the final result in your game. However, first let's break down what is happening here.

How it works...

Now we are starting to gain a ton of power and flexibility in our Screen Effects programming. I am sure that you are now starting to understand how much one can do with this simple system in Unity. We can literally replicate the effects of Photoshop layer blending modes in our game to give artists the flexibility they need to achieve high-quality graphics in a short amount of time.

With this particular recipe, we looked at how to multiply two images together, add two images together, and perform a screen blending mode, using just a little bit of mathematics. When working with blend modes, one has to think on a per-pixel level. For instance, when we are using a multiply blend mode, we literally take each pixel from the original render texture and multiply them with each pixel of the blend texture. The same goes for the add blend mode. It is just a simple mathematical operation of adding each pixel from the source texture, or render texture, to the blend texture.

The screen blend mode is definitely a bit more involved, but it is actually doing the same thing. It takes each image, render texture, and blend texture, inverts them, then multiplies them together, and inverts them again to achieve the final look. Just like Photoshop blends its textures together using blend modes, we can do the same with screen effects.

There's more...

Let's continue this recipe by adding a couple of more blend modes to our screen effect.

In the screen effect shader, let's add the following code to our `frag()` function and change the value we are returning to our script. We will also need to comment out the multiply blend so that we don't return that as well:

```
fixed4 frag(v2f_img i) : COLOR
{
    //Get the colors from the RenderTexture and the uv's
    //from the v2f_img struct
    fixed4 renderTex = tex2D(_MainTex, i.uv);
    fixed4 blendTex = tex2D(_BlendTex, i.uv);

    //Perform a multiply Blend mode
    //fixed4 blendedMultiply = renderTex * blendTex;
    fixed4 blendedMultiply = renderTex + blendTex;

    //Adjust amount of Blend Mode with a lerp
    renderTex = lerp(renderTex, blendedMultiply, _Opacity);

    return renderTex;
}
```

1. Save the shader file in MonoDevelop and return to the Unity editor to let the shader compile. If no errors occurred, you should see a result similar to the following image. This is a simple add blending mode:

Add Blending Mode

As you can see, this has the opposite effect of multiply because we are adding the two images together.

2. Finally, let's add one more blend mode called a Screen Blend. This one is a little bit more involved, from a mathematical standpoint, but still simple to implement. Enter the following code in the `frag()` function of our shader:

```
fixed4 frag(v2f_img i) : COLOR
{
    //Get the colors from the RenderTexture and the uv's
    //from the v2f_img struct
    fixed4 renderTex = tex2D(_MainTex, i.uv);
    fixed4 blendTex = tex2D(_BlendTex, i.uv);

    //Perform a multiply Blend mode
    //fixed4 blendedMultiply = renderTex * blendTex;
    //fixed4 blendedAdd = renderTex + blendTex;
    fixed4 blendedScreen = (1.0 - ((1.0 - renderTex) * (1.0 - blendTex)));

    //Adjust amount of Blend Mode with a lerp
    renderTex = lerp(renderTex, blendedScreen, _Opacity);

    return renderTex;
}
```

The following image demonstrates the results of using a Screen type blend mode to blend two images together in a screen effect:

Using the Overlay Blend mode with screen effects

For our final recipe, we are going to take a look at another type of blend mode, the Overlay Blend mode. This blending actually makes use of some conditional statements that determine the final color of each pixel in each channel. So, the process of using this type of blend mode needs a bit more coding to work. Let's take a look at how this is done in the next few recipes.

Getting ready

For this last Screen Effect, we will need to set up our two scripts as we have in the previous recipes in this chapter. For this recipe, we will be using the same scene we have been using, so we don't have to create a new one:

1. Create a new script file called `Overlay_ImageEffect` and shader file called `Overlay_Effect`.

2. Copy the code from the previous C# script file to our new script file.

3. Copy the code from the previous shader file to our new shader file.

4. Assign the `Overlay_ImageEffect` script to the main camera and `Overlay_Effect` to the script component in the **Inspector**.

5. Finally, double-click on the script and shader files to open them in the MonoDevelop editor.

How to do it...

To begin our Overlay Screen Effect, we will need to get the code of our shader up and running without errors. We can then modify our script file to feed the correct data to the shader.

1. We first need to set up our properties in our `Properties` block. We will use the same properties from the previous few recipes in this chapter:

```
Properties
{
    _MainTex ("Base (RGB)", 2D) = "white" {}
    _BlendTex ("Blend Texture", 2D) = "white"{}
    _Opacity ("Blend Opacity", Range(0,1)) = 1
}
```

2. We then need to create the corresponding variables in our `CGPROGRAM`:

```
Pass
{
    CGPROGRAM
    #pragma vertex vert_img
    #pragma fragment frag
    #pragma fragmentoption ARB_precision_hint_fastest
    #include "UnityCG.cginc"

    uniform sampler2D _MainTex;
    uniform sampler2D _BlendTex;
    fixed _Opacity;
```

3. In order for the Overlay Blend effect to work, we will have to process each pixel from each channel individually. To do this in a shader, we have to write a custom function that will take in a single channel, for instance, the red channel, and perform the Overlay operation. Enter the following code in the shader just below the variable declarations:

```
fixed OverlayBlendMode(fixed basePixel, fixed blendPixel)
{
    if(basePixel < 0.5)
    {
        return (2.0 * basePixel * blendPixel);
    }
    else
    {
        return (1.0 - 2.0 * (1.0 - basePixel) * (1.0 - blendPixel));
    }
}
```

4. Finally, we need to update our `frag()` function to process each channel of our textures to perform the blending:

```
fixed4 frag(v2f_img i) : COLOR
{
    //Get the colors from the RenderTexture and the uv's
    //from the v2f_img struct
    fixed4 renderTex = tex2D(_MainTex, i.uv);
    fixed4 blendTex = tex2D(_BlendTex, i.uv);

    fixed4 blendedImage = renderTex;

    blendedImage.r = OverlayBlendMode(renderTex.r, blendTex.r);
    blendedImage.g = OverlayBlendMode(renderTex.g, blendTex.g);
    blendedImage.b = OverlayBlendMode(renderTex.b, blendTex.b);

    //Adjust amount of Blend Mode with a lerp
    renderTex = lerp(renderTex, blendedImage, _Opacity);

    return renderTex;
}
```

5. With the code completed in the shader, our effect should be working. Save the shader and return to the Unity editor to let the shader compile. Our script is already set up, so we don't have to modify it any further. Once the shader compiles, you should see a result similar to the following image:

Overlay Blend Mode

How it works...

Our Overlay blend mode is definitely a lot more involved, but if you really break down the function, you will notice that it is simply a multiply blend mode and screen blend mode. It's just that, in this case, we are doing a conditional check to apply one or the other blend mode to a pixel.

With this particular Screen Effect, when the Overlay function receives a pixel, it checks to see whether it is less than 0.5. If it is, then we apply a modified multiply blend mode to that pixel; if it's not, then we apply a modified screen blend mode to the pixel. We do this for each pixel for each channel, giving us the final RGB pixel values for our Screen effect.

As you can see, there are many things that can be done with screen effects. It really just depends on the platform and amount of memory you have allocated for screen effects. Usually, this is determined throughout the course of a game project, so have fun and get creative with your screen effects.

9

Gameplay and Screen Effects

When it comes to creating believable and immersive games, material design is not the only aspect that we need to take into account. The overall feeling can be altered using screen effects. This is very common in movies, for instance, when colors are corrected in the post-production phase. You can implement these techniques in your games too, using the knowledge from *Chapter 8, Screen Effects with Unity Render Texture*. Two interesting effects are presented in this chapter; you can, however, adapt them to fit your needs and create your very own screen effect.

In this chapter, you will learn the following recipes:

▶   Creating an old movie screen effect

▶   Creating a night vision screen effect

Introduction

If you are reading this book, you are most likely a person who has played a game or two in your time. One of the aspects of real-time games is the effect of immersing a player into a world to make it feel as if they were actually playing in the real world. The more modern games make heavy use of screen effects to achieve this immersion.

With screen effects, we can turn the mood of a certain environment from calm to scary, just by changing the look of the screen. Imagine walking into a room that is contained within a level, then the game takes over and goes into a cinematic moment. Many modern games will turn on different screen effects to change the mood of the current moment. Understanding how to create effects triggered by gameplay is next in our journey of shader writing.

In this chapter, we are going to take a look at some of the more common gameplay screen effects. You are going to learn how to change the look of the game from normal to an old movie effect, and we are going to take a look at how many first-person shooter games apply their night vision effects to the screen. With each of these recipes, we are going to look at how to hook these up to game events so that they are turned on and off as the game's current presentation needs.

Creating an old movie screen effect

Many games are set in different times. Some take place in fantasy worlds or future sci-fi worlds, and some even take place in the old west, where film cameras were just being developed and the movies that people watched were black and white or sometimes tinted with what is called a sepia effect. The look is very distinct, and we are going to replicate this look using a screen effect in Unity.

There are a few steps to achieve this look, and just to make the whole screen black and white or grayscale, we need to break down this effect into its component parts. If we analyze some reference footage of an old movie, we can begin to do this. Let's take a look at the following image and break down the elements that make up the old movie look:

We constructed this image using a few reference images found online. It is always a good idea to try and utilize Photoshop to construct images like this to aid you in creating a plan for your new screen effect. Performing this process not only tells us the elements we will have to code in, but it also gives us a quick way to see which blending modes work and how we will construct the layers of our screen effect. The Photoshop file we created for this recipe is included in this book's support page at www.packtpub.com/support and is called OldFilmEffect_Research_Layout.psd.

Getting ready

Now that we know what we have to make, let's take a look at how each of the layers is combined to create the final effect and gather some resources for our shader and screen effect script.

> ▶ **Sepia tone**: This is a relatively simple effect to achieve, as we just need to bring all the pixel colors of the original render texture to a single color range. This is easily achieved using the luminance of the original image and adding a constant color. Our first layer will look like the following image:

> ▶ **Vignette effect**: We can always see some sort of soft border around old films when they are being projected with an old movie projector. This is caused because the bulb being used for the movie projector has more brightness in the middle than it does at the edges of the film. This effect is generally called the vignette effect and is our second layer in our screen effect. We can achieve this with an overlaid texture over the whole screen. The following image demonstrates what this layer looks like, isolated as a texture:

► **Dust and scratches**: The third and final layer in our old movie screen effect is dust and scratches. This layer will utilize two different tiled textures, one for scratches and one for dust. The reason is that we will want to animate these two textures over time at different tiling rates. This will give the effect that the film is moving along and there are small scratches and dust on each frame of the old film. The following image demonstrates this effect isolated to its own texture:

Let's get our screen effect system ready with the preceding textures. Perform the following steps:

1. Gather up a vignette texture and dust and scratches texture, like the ones we just saw.

2. Create a new script called `OldFilmEffect.cs` and a new shader called `OldFilmEffectShader.shader`.

3. With our new files created, fill in the code necessary to get the screen effect system up and running. For references on how to do this, see *Chapter 8*, *Screen Effects with Unity Render Textures*.

Finally, with our screen effect system up and running and our textures gathered, we can begin the process of recreating this old film effect.

How to do it...

Our individual layers for our old film screen effect are quite simple, but when combined, we get some very visually stunning effects. Let's run through how to construct the code for our script and shader, then we can step through each line of code and learn why things are working the way they are. At this point, you should have the screen effects system up and running, as we will not be covering how to set this up in this recipe.

1. We will begin by entering the code in our script. Our first block of code that we will enter will define our variable that we want to expose to **Inspector** in order to let the user of this effect adjust it as they see fit. We can also use our mocked-up Photoshop file as a reference when deciding what we will need to expose to the **Inspector** of this effect. Enter the following code in your effect script:

```
#region Variables
public Shader oldFilmShader;

public float OldFilmEffectAmount = 1.0f;

public Color sepiaColor = Color.white;
public Texture2D vignetteTexture;
public float vignetteAmount = 1.0f;

public Texture2D scratchesTexture;
public float scratchesYSpeed = 10.0f;
public float scratchesXSpeed = 10.0f;

public  Texture2D dustTexture;
public float dustYSpeed = 10.0f;
public float dustXSpeed = 10.0f;

private Material curMaterial;
private float randomValue;
#endregion
```

2. Next, we need to fill in the contents of our `OnRenderImage()` function. Here, we will be passing the data from our variables to our shader so that the shader can then use this data in the processing of the render texture:

```
void OnRenderImage(RenderTexture sourceTexture, RenderTexture destTexture)
    {
        if(oldFilmShader != null)
        {
            material.SetColor("_SepiaColor", sepiaColor);
            material.SetFloat("_VignetteAmount", vignetteAmount);
```

```
            material.SetFloat("_EffectAmount",
OldFilmEffectAmount);

            if(vignetteTexture)
            {
                material.SetTexture("_VignetteTex",
vignetteTexture);
            }

            if(scratchesTexture)
            {
                material.SetTexture("_ScratchesTex",
scratchesTexture);
                material.SetFloat("_ScratchesYSpeed",
scratchesYSpeed);
                material.SetFloat("_ScratchesXSpeed",
scratchesXSpeed);
            }

            if(dustTexture)
            {
                material.SetTexture("_DustTex", dustTexture);
                material.SetFloat("_dustYSpeed", dustYSpeed);
                material.SetFloat("_dustXSpeed", dustXSpeed);
                material.SetFloat("_RandomValue", randomValue);
            }

            Graphics.Blit(sourceTexture, destTexture, material);
        }
        else
        {
            Graphics.Blit(sourceTexture, destTexture);
        }
    }
```

3. To complete the script portion of this effect, we simply need to make sure that we clamp the values of the variables that need to have a clamped range instead of being any value:

```
void Update()
    {
        vignetteAmount = Mathf.Clamp01(vignetteAmount);
        OldFilmEffectAmount = Mathf.Clamp(OldFilmEffectAmount, 0f,
1.5f);
        randomValue = Random.Range(-1f,1f);
    }
```

4. With our script complete, let's turn our attention to our shader file. We need to create the corresponding variables, which we created in our script in our shader. This will allow the script and shader to communicate with one another. Enter the following code in the `Properties` block of the shader:

```
Properties
{
    _MainTex ("Base (RGB)", 2D) = "white" {}
    _VignetteTex ("Vignette Texture", 2D) = "white"{}
    _ScratchesTex ("Scartches Texture", 2D) = "white"{}
    _DustTex ("Dust Texture", 2D) = "white"{}
    _SepiaColor ("Sepia Color", Color) = (1,1,1,1)
    _EffectAmount ("Old Film Effect Amount", Range(0,1)) = 1.0
    _VignetteAmount ("Vignette Opacity", Range(0,1)) = 1.0
    _ScratchesYSpeed ("Scratches Y Speed", Float) = 10.0
    _ScratchesXSpeed ("Scratches X Speed", Float) = 10.0
    _dustXSpeed ("Dust X Speed", Float) = 10.0
    _dustYSpeed ("Dust Y Speed", Float) = 10.0
    _RandomValue ("Random Value", Float) = 1.0
    _Contrast ("Contrast", Float) = 3.0
}
```

5. Then, as usual, we need to add these same variable names to our `CGPROGRAM` block so that the `Properties` block can communicate with the `CGPROGRAM` block:

```
CGPROGRAM
#pragma vertex vert_img
#pragma fragment frag
#pragma fragmentoption ARB_precision_hint_fastest
#include "UnityCG.cginc"

uniform sampler2D _MainTex;
uniform sampler2D _VignetteTex;
uniform sampler2D _ScratchesTex;
uniform sampler2D _DustTex;
fixed4 _SepiaColor;
fixed _VignetteAmount;
fixed _ScratchesYSpeed;
fixed _ScratchesXSpeed;
fixed _dustXSpeed;
fixed _dustYSpeed;
fixed _EffectAmount;
fixed _RandomValue;
fixed _Contrast;
```

6. Now, we simply fill in the guts of our `frag()` function so that we process the pixels for our screen effect. To start with, let's get the render texture and vignette texture passed to us by the script:

```
fixed4 frag(v2f_img i) : COLOR
{
    //Get the colors from the RenderTexture and the uv's
    //from the v2f_img struct
    half2 distortedUV = barrelDistortion(i.uv);
    distortedUV = half2(i.uv.x, i.uv.y + (_RandomValue * _
SinTime.z * 0.005));
    fixed4 renderTex = tex2D(_MainTex, i.uv);

    //Get the pixels from the Vignette Texture
    fixed4 vignetteTex = tex2D(_VignetteTex, i.uv);
```

7. We then need to add the process for the dust and scratches by entering the following code:

```
//Process the Scratches UV and pixels
half2 scratchesUV = half2(i.uv.x + (_RandomValue * _SinTime.z *
_ScratchesXSpeed), i.uv.y + (_Time.x * _ScratchesYSpeed));
fixed4 scratchesTex = tex2D(_ScratchesTex, scratchesUV);

//Process the Dust UV and pixels
half2 dustUV = half2(i.uv.x + (_RandomValue * (_SinTime.z
* _dustXSpeed)), i.uv.y + (_RandomValue * (_SinTime.z * _
dustYSpeed)));
fixed4 dustTex = tex2D(_DustTex, dustUV);
```

8. The sepia tone process is next on our list:

```
// get the luminosity values from the render texture using the YIQ
values.
fixed lum = dot (fixed3(0.299, 0.587, 0.114), renderTex.rgb);

//Add the constant color to the lum values
fixed4 finalColor = lum + lerp(_SepiaColor, _SepiaColor +
 fixed4(0.1f,0.1f,0.1f,1.0f), _RandomValue);
finalColor = pow(finalColor, _Contrast);
```

9. Finally, we combine all of our layers and colors and return the final screen effect texture:

```
//Create a constant white color we can use to adjust opacity of
effects
fixed3 constantWhite = fixed3(1,1,1);
```

```
//Composite together the different layers to create finsl Screen
Effect
finalColor = lerp(finalColor, finalColor * vignetteTex, _
VignetteAmount);
finalColor.rgb *= lerp(scratchesTex, constantWhite, (_
RandomValue));
finalColor.rgb *= lerp(dustTex.rgb, constantWhite, (_RandomValue *
_SinTime.z));
finalColor = lerp(renderTex, finalColor, _EffectAmount);

return finalColor;
```

10. With all of our code entered and no errors, you should have a result very similar to the following image. Hit play in the Unity editor to see the effects of the dust and scratches and the slight image shift that we gave the screen effect:

How it works...

Now, let's walk through each of the layers in this screen effect, break down why each of the lines of code is working the way it is, and get more insight as to how we can add more to this screen effect.

Now that our old film screen effect is working, let's step through the lines of code in our `frag()` function as all the other code should be pretty self-explanatory at this point in the book.

Just like our Photoshop layers, our shader is processing each layer and then compositing them together, so while we go through each layer, try to imagine how the layers in Photoshop work. Keeping this concept in mind always helps when developing new screen effects.

Here, we have the first set of lines of code in our `frag()` function:

```
fixed4 frag(v2f_img i) : COLOR
{
    //Get the colors from the RenderTexture and the uv's
    //from the v2f_img struct
    half2 distortedUV = barrelDistortion(i.uv);
    fixed4 renderTex = tex2D(_MainTex, i.uv);
    fixed4 vignetteTex = tex2D(_VignetteTex, i.uv);
```

The first line of code, just after the `frag()` function declaration, is the definition of how the UVs should work for our main render texture or the actual rendered frame of our game. As we are looking to fake the effect of an old film style, we want to adjust the UVs of our render texture, every frame, such that they flicker. This flickering simulates how the winding of the film's projector is just a bit off. This tells us that we need to animate the UVs and this is what this first line of code is doing.

We used the built-in `_SinTime` variable, which Unity provides, to get a value between -1 and 1. We then multiply this by a very small number, in this case, `0.005`, to reduce the intensity of the effect. The final value is then multiplied again by the `_RandomValue` variable, which we generated in the effect script. This value bounces back and forth between -1 and 1 to basically flip the direction of the motion back and forth.

Once our UVs are built and stored in the `renderTexUV` variable, we can sample the render texture using a `tex2D()` function. This operation then gives us our final render texture, which we can use to process further in the rest of the shader.

Moving on to the last line in the previous image, we simply do a straight sample of the vignette texture using the `tex2D()` function. We don't need to use the animated UVs we already created, as the vignette texture will be tied to the motion of the camera itself and not to the flickering of the camera film.

The following code snippet illustrates the second set of lines of code in our `frag()` function:

```
//Process the Scratches UV and pixels
half2 scratchesUV = half2(i.uv.x + (_RandomValue * _SinTime.z * _
ScratchesXSpeed),
        i.uv.y + (_Time.x * _ScratchesYSpeed));
fixed4 scratchesTex = tex2D(_ScratchesTex, scratchesUV);

//Process the Dust UV and pixels
half2 dustUV = half2(i.uv.x + (_RandomValue * (_SinTime.z * _
dustXSpeed)),
        i.uv.y + (_RandomValue * (_SinTime.z * _dustYSpeed)));
fixed4 dustTex = tex2D(_DustTex, dustUV);
```

These lines of code are almost exactly like the previous lines of code in which we need to generate unique animated UV values to modify the position of our screen effect layers. We simply use the built-in `_SinTime` value to get a value between -1 and 1, multiply it by our random value, and then by another multiplier to adjust the overall speed of the animation. Once these UV values are generated, we can then sample our dust and scratches texture using these new animated values.

Our next set of code handles the creation of the colorizing effect for our old film screen effect. The following code snippet demonstrates these lines:

```
// get the luminosity values from the render texture using the YIQ
values.
fixed lum = dot (fixed3(0.299, 0.587, 0.114), renderTex.rgb);

//Add the constant color to the lum values
fixed4 finalColor = lum + lerp(_SepiaColor, _SepiaColor +
    fixed4(0.1f,0.1f,0.1f,1.0f), _RandomValue);
```

With this set of code, we are creating the actual color tinting of the entire render texture. To accomplish this, we first need to turn the render texture into the grayscale version of itself. To do this, we can use the luminosity values given to us by the `YIQ` values. `YIQ` values are the color space used by the NTSC color TV system. Each letter in `YIQ` actually stores color constants that are used by TVs to adjust the color for readability.

While it is not necessary to actually know the reasons for this color scale, it should be known that the `Y` value in `YIQ` is the constant luminance value for any image. So, we can generate a grayscale image of our render texture by taking each pixel of the render texture and dotting it with our luminance values. This is what the first line in this set is doing.

Once we have the luminance values, we can simply add the color we want to tint the image with. This color is passed from our script to our shader, then to our `CGPROGRAM` block, where we can add it to our grayscale render texture. Once completed, we will have a perfectly tinted image.

Finally, we create the blending between each of our layers in our screen effect. The following code snippet shows the set of code we are looking at:

```
//Create a constant white color we can use to adjust opacity of
effects
fixed3 constantWhite = fixed3(1,1,1);

//Composite together the different layers to create finsl Screen
Effect
finalColor = lerp(finalColor, finalColor * vignetteTex, _
VignetteAmount);
finalColor.rgb *= lerp(scratchesTex, constantWhite, (_RandomValue));
finalColor.rgb *= lerp(dustTex.rgb, constantWhite, (_RandomValue *
_SinTime.z));
finalColor = lerp(renderTex, finalColor, _EffectAmount);

return finalColor
```

Our last set of code is relatively simple and doesn't really need a ton of explanation. In short, it is simply multiplying all the layers together to reach our final result. Just like we multiplied our layers together in Photoshop, we multiply them together in our shader. Each layer is processed through a `lerp()` function so that we can adjust the opacity of each layer, which gives more artistic control over the final effect. The more tweaks one can offer, the better when it comes to screen effects.

See also

For more information on the YIQ values, refer to the following links:

▸ http://en.wikipedia.org/wiki/YIQ
▸ http://www.blackice.com/colorspaceYIQ.htm

Creating a night vision screen effect

Our next screen effect is definitely a more popular one. The night vision screen effect is seen in *Call of Duty Modern Warfare*, *Halo*, and just about any first-person shooter out in the market today. It is the effect of brightening the whole image using that very distinct lime green color.

In order to achieve our night vision effect, we need to break down our effect using Photoshop. It is a simple process of finding some reference images online and composing a layered image to see what kind of blending modes you will need or in which order we will need to combine our layers. The following image shows the result of performing just this process in Photoshop:

Let's begin to break down our rough Photoshop composite image into its component parts so that we can better understand the assets we will have to gather. In the next recipe, we will cover the process of doing this.

Getting ready

Let's begin this screen effect by again breaking down our effect into its component layers. Using Photoshop, we can construct a layered image to better illustrate how we can go about capturing the effect of night vision:

▶ **Tinted green**: Our first layer in our screen effect is the iconic green color, found in just about every night vision image. This will give our effect that signature night vision look, as shown in the following image:

▶ **Scan lines**: To increase the effect of this being a new type of display for the player, we include scan lines over the top of our tinted layer. For this, we will use a texture created in Photoshop and let the user tile it so that the scan lines can be bigger or smaller.

▶ **Noise**: Our next layer is a simple noise texture that we tile over the tinted image and scan lines to break up the image and add even more detail to our effect. This layer simply emphasizes that digital read-out look:

▶ **Vignette**: The last layer in our night vision effect is the vignette. If you look at the night vision effect in *Call of Duty Modern Warfare*, you will notice that it uses a vignette that fakes the effect of looking down a scope. We will do that for this screen effect:

Let's create a screen effect system by gathering our textures. Perform the following steps:

1. Gather up a vignette texture, noise texture, and scan line texture, like the ones we just saw.

2. Create a new script called `NightVisionEffect.cs` and a new shader called `NightVisionEffectShader.shader`.

3. With our new files created, fill in the code necessary to get the screen effect system up and running. For instructions on how to do this, refer to *Chapter 8*, *Screen Effects with Unity Render Textures*.

Finally, with our screen effect system up and running and our textures gathered, we can begin the process of recreating this old film effect.

How to do it...

With all of our assets gathered and screen effect system running smoothly, let's begin to add the code necessary to both the script and shader. We will begin our coding with the `NightVisionEffect.cs` script, so double-click on this file now to open it in MonoDevelop.

1. We need to create a few variables that will allow the user of this effect to adjust it in the script's **Inspector**. Enter the following code in the `NightVisionEffect.cs` script:

```
#region Variables
    public Shader nightVisionShader;
```

```
        public float contrast = 2.0f;
        public float brightness = 1.0f;
        public Color nightVisionColor = Color.white;

        public Texture2D vignetteTexture;

        public Texture2D scanLineTexture;
        public float scanLineTileAmount = 4.0f;

        public Texture2D nightVisionNoise;
        public float noiseXSpeed = 100.0f;
        public float noiseYSpeed = 100.0f;

        public float distortion = 0.2f;
        public float scale = 0.8f;

        private float randomValue = 0.0f;
        private Material curMaterial;
        #endregion
```

2. Next, we need to complete our `OnRenderImage()` function so that we are passing the right data to the shader in order for the shader to process the screen effect properly. Complete the `OnRenderImage()` function with the following code:

```
void OnRenderImage(RenderTexture sourceTexture, RenderTexture
destTexture)
    {
        if(nightVisionShader != null)
        {
            material.SetFloat("_Contrast", contrast);
            material.SetFloat("_Brightness", brightness);
            material.SetColor("_NightVisionColor",
nightVisionColor);
            material.SetFloat("_RandomValue", randomValue);
            material.SetFloat("_distortion", distortion);
            material.SetFloat("_scale",scale);

            if(vignetteTexture)
            {
                material.SetTexture("_VignetteTex",
vignetteTexture);
            }

            if(scanLineTexture)
            {
                material.SetTexture("_ScanLineTex",
scanLineTexture);
```

```
                material.SetFloat("_ScanLineTileAmount",
    scanLineTileAmount);
            }

            if(nightVisionNoise)
            {
                material.SetTexture("_NoiseTex",
    nightVisionNoise);
                material.SetFloat("_NoiseXSpeed", noiseXSpeed);
                material.SetFloat("_NoiseYSpeed", noiseYSpeed);
            }

            Graphics.Blit(sourceTexture, destTexture, material);
        }
        else
        {
            Graphics.Blit(sourceTexture, destTexture);
        }
    }
```

3. To complete the `NightVisionEffect.cs` script, we simply need to make sure that we clamp certain variables so that they stay within a range. These ranges are arbitrary and can be changed at a later time. These are just values that worked well:

```
void Update()
    {
        contrast = Mathf.Clamp(contrast, 0f,4f);
        brightness = Mathf.Clamp(brightness, 0f, 2f);
        randomValue = Random.Range(-1f,1f);
        distortion = Mathf.Clamp(distortion, -1f,1f);
        scale = Mathf.Clamp(scale, 0f, 3f);
    }
```

4. We can now turn our attention over to the shader portion of this screen effect. Open the shader, if you haven't already, and begin by entering the following properties in the `Properties` block:

```
Properties
    {
        _MainTex ("Base (RGB)", 2D) = "white" {}
        _VignetteTex ("Vignette Texture", 2D) = "white"{}
        _ScanLineTex ("Scan Line Texture", 2D) = "white"{}
        _NoiseTex ("Noise Texture", 2D) = "white"{}
        _NoiseXSpeed ("Noise X Speed", Float) = 100.0
        _NoiseYSpeed ("Noise Y Speed", Float) = 100.0
        _ScanLineTileAmount ("Scan Line Tile Amount", Float) = 4.0
        _NightVisionColor ("Night Vision Color", Color) =
    (1,1,1,1)
```

```
            _Contrast ("Contrast", Range(0,4)) = 2
            _Brightness ("Brightness", Range(0,2)) = 1
            _RandomValue ("Random Value", Float) = 0
            _distortion ("Distortion", Float) = 0.2
            _scale ("Scale (Zoom)", Float) = 0.8
        }
```

5. To make sure that we are passing the data from our `Properties` block to our `CGPROGRAM` block, we need to make sure to declare them with the same name in the `CGPROGRAM` block:

```
CGPROGRAM
#pragma vertex vert_img
#pragma fragment frag
#pragma fragmentoption ARB_precision_hint_fastest
#include "UnityCG.cginc"

uniform sampler2D _MainTex;
uniform sampler2D _VignetteTex;
uniform sampler2D _ScanLineTex;
uniform sampler2D _NoiseTex;
fixed4 _NightVisionColor;
fixed _Contrast;
fixed _ScanLineTileAmount;
fixed _Brightness;
fixed _RandomValue;
fixed _NoiseXSpeed;
fixed _NoiseYSpeed;
fixed _distortion;
fixed _scale;
```

6. Our effect is also going to include a lens distortion to further convey the effect that we are looking through a lens and the edges of the image are being distorted by the angle of the lens. Enter the following function just after the variable declarations in the `CGPROGRAM` block:

```
float2 barrelDistortion(float2 coord)
    {
        // lens distortion algorithm
        // See http://www.ssontech.com/content/lensalg.htm

        float2 h = coord.xy - float2(0.5, 0.5);
        float r2 = h.x * h.x + h.y * h.y;
        float f = 1.0 + r2 * (_distortion * sqrt(r2));

        return f * _scale * h + 0.5;
    }
```

7. We can now concentrate on the meat of our `NightVisionEffect` shader. Let's start this by entering the code that is necessary to get the render texture and vignette texture. Enter the following code in the `frag()` function of our shader:

```
fixed4 frag(v2f_img i) : COLOR
    {
        //Get the colors from the RenderTexture and the uv's
        //from the v2f_img struct
        half2 distortedUV = barrelDistortion(i.uv);
        fixed4 renderTex = tex2D(_MainTex, distortedUV);
        fixed4 vignetteTex = tex2D(_VignetteTex, i.uv);
```

8. The next step in our `frag()` function is to process the scan lines and noise textures and apply the proper animated UVs to them:

```
//Process scan lines and noise
        half2 scanLinesUV = half2(i.uv.x * _ScanLineTileAmount,
i.uv.y * _ScanLineTileAmount);
        fixed4 scanLineTex = tex2D(_ScanLineTex, scanLinesUV);

        half2 noiseUV = half2(i.uv.x + (_RandomValue * _SinTime.z
* _NoiseXSpeed),
                                        i.uv.y + (_Time.x * _
NoiseYSpeed));
        fixed4 noiseTex = tex2D(_NoiseTex, noiseUV);
```

9. To complete all of our layers in the screen effect, we simply need to process the luminance value of our render texture, and then apply the night vision color to it to achieve that iconic night vision look:

```
// get the luminosity values from the render texture using the YIQ
values.
        fixed lum = dot (fixed3(0.299, 0.587, 0.114), renderTex.
rgb);
        lum += _Brightness;
        fixed4 finalColor = (lum *2) + _NightVisionColor;
```

10. Lastly, we will combine all the layers together and return the final color of our night vision effect:

```
//Final output
        finalColor = pow(finalColor, _Contrast);
        finalColor *= vignetteTex;
        finalColor *= scanLineTex * noiseTex;

        return finalColor;
```

When you have finished entering the code, return to the Unity editor to let the script and shader compile. If there are no errors, hit play in the editor to see the results. You should see something similar to the following image:

How it works...

The night vision effect is actually very similar to the old film screen effect, which shows us just how modular we can make these components. Just by simply swapping the textures that we are using for overlays and changing the speed at which our tiling rates are being calculated, we can achieve very different results using the same code.

The only difference with this effect is the fact that we are including a lens distortion to our screen effect. So let's break this down so that we can get a better understanding of how it works.

The following code snippet illustrates the code used in processing our lens distortion. It is a snippet of code provided to us by the makers of SynthEyes, and the code is freely available to use in your own effects:

```
float2 barrelDistortion(float2 coord)
{
    // lens distortion algorithm
    // See http://www.ssontech.com/content/lensalg.htm
    float2 h = coord.xy - float2(0.5, 0.5);
    float r2 = h.x * h.x + h.y * h.y;
```

```
        float f = 1.0 + r2 * (_distortion * sqrt(r2));

        return f * _scale * h + 0.5;
    }
```

There's more...

It is not uncommon in video games to have the need to highlight certain objects. For instance, a thermal visor should apply a post-processing effect only to people and other sources of heat. Doing this is already possible with the knowledge gathered so far in this book; you can, in fact, change the shader or material of an object by code. However, this is often laborious and has to be replicated on all the objects.

A more effective way is using replaced shaders. Each shader has a tag called `RenderType` that has never been used so far. This property can be used to force a camera to apply a shader only to certain objects. You can do this by attaching the following script to the camera:

```
using UnityEngine;

public class ReplacedShader : MonoBehaviour {

    public Shader shader;
    void Start () {
        GetComponent<Camera>().SetReplacementShader(shader, "Heat");
    }
}
```

After entering the play mode, the camera will query all the objects that it has to render. If they don't have a shader decorated with `RenderType = "Heat"`, they will not be rendered. Objects with such a tag will be rendered with the shader attached to the script.

10
Advanced Shading Techniques

In this chapter, you will learn the following recipes:

- ► Using CgInclude files that are built into Unity
- ► Making your shader world modular with CgInclude
- ► Implementing a Fur Shader
- ► Implementing heatmaps with arrays

Introduction

This final chapter covers some advanced shader techniques that you can use for your game. You should remember that many of the most eye-catching effects you can see in games are made by testing the limit of what shaders can do. This book provides you with the technical basis to modify and create shaders, but you are strongly encouraged to play and experiment with them as much as you can. Making a good game is not a quest for photorealism; you should not approach shaders with the intention of replicating reality because this is unlikely to happen. Instead, you should try to use shaders as a tool to make your game truly unique. With the knowledge of this final chapter, you will be able to create the materials that you want.

Using CgInclude files that are built into Unity

Our first step in writing our own CgInclude files is to understand what Unity is already providing us with for shaders. By writing Surface Shaders, there is a lot happening under the hood, which makes the process of writing Surface Shaders so efficient. We can see this code in the included CgInclude files found in your Unity install folder at `Editor | Data | CGIncludes`. All the files contained within this folder do their part to render our objects with our shaders to the screen. Some of these files take care of shadows and lighting, some take care of helper functions, and some manage platform dependencies. Without them, our shader writing experience would be much more laborious.

You can find a list of information that Unity has provided us with at the following link: `http://docs.unity3d.com/Documentation/Components/SL-BuiltinIncludes.html`

Let's begin the process of understanding these built-in CgInclude files, using some of the built-in helper functions from the `UnityCG.cginc` file.

Getting ready

Before we start diving into the meat of writing the shader, we need to get a few items set up in our scene. Let's create the following and then open the shader in MonoDevelop:

1. Create a new scene and fill it with a simple sphere model.
2. Create a new shader and material.
3. Attach the new shader to the new material and assign the material to the sphere.
4. Then, let's create a directional light and position it above our sphere.
5. Finally, we are going to want to open the `UnityCG.cginc` file from Unity's `CgInclude` folder located in Unity's install directory. This will let us analyze some of the helper function's code so that we can understand better what is happening when we use them.

6. You should have a simple scene set up to work on the shader. Refer to the following screenshot as an example:

How to do it...

With the scene prepared, we can now begin the process of experimenting with some of the built-in helper functions included with the `UnityCG.cginc` file. Double-click on the shader that was created for this scene in order to open it in MonoDevelop and insert the code given in the following steps:

1. Add the following code to the `Properties` block of the new shader file. We will need a single texture and slide for our example shader:

```
Properties
{
    _MainTex ("Base (RGB)", 2D) = "white" {}
    _DesatValue ("Desaturate", Range(0,1)) = 0.5
}
```

2. We then need to make sure that we create the data connection between our `Properties` and `CGPROGRAM` blocks, with the following code placed after the `CGPROGRAM` declaration and `#pragma` directives:

```
sampler2D _MainTex;
fixed _DesatValue;
```

3. Finally, we just have to update our `surf()` function to include the following code. We introduce a new function that we haven't seen yet, which is built into Unity's `UnityCG.cginc` file:

```
void surf (Input IN, inout SurfaceOutput o)
{
    half4 c = tex2D (_MainTex, IV.uv_MainTex);
    c.rgb = lerp(c.rgb, Luminance(r.rgb), _DesatValue);

    o.Albedo = c.rgb;
    o.Alpha = c.a;
}
```

With the shader code modified, you should see something similar to the following screenshot. We have simply used a helper function, built into Unity's CgInclude file, to give us an effect of desaturating the main texture of our shader:

Luminace Value = 0.0   Luminace Value = 0.5   Luminace Value = 1.0

How it works...

Using the built-in helper function named `Luminance()`, we are able to quickly get a desaturation or grayscale effect on our shaders. This is all possible because the `UnityCG.cginc` file is brought automatically to our shader as we are using a Surface shader.

If you search through the `UnityCG.cginc` file, opened in MonoDevelop, you will find the implementation of this function at line 276. The following snippet is taken from the file:

```
inline fixed Luminance (fixed3 c)
{
    return dot(c, fixed3(0.22, 0.707, 0.071));
}
```

As this function is included in the file and Unity automatically compiles with this file, we can use the function in our code as well, thereby reducing the amount of code that we have to write over and over again.

If you notice there is also a `Lighting.cginc` file that Unity comes with. This file houses all the lighting models that we use when we declare something like `#pragma Surface surf Lambert`. Sifting through this file reveals that all the built-in lighting models are defined here for reuse and modularity.

Making your shader world modular with CgInclude

Knowing about the built-in CgInclude files is great, but what if we wanted to build our own CgInclude files to store our own lighting models and helper functions? We can, in fact, create our own CgInclude files, but we need to learn a little more code syntax before we can start using them efficiently in our shader writing pipelines. Let's take a look at the process of creating a new CgInclude file from scratch.

Getting ready

Let's walk through the process of generating a new item for this recipe.

1. Begin by creating a new text file and call it something like `MyCgInclude.txt`.
2. Then change its file extension to `.cginc`. Windows will give you a warning message saying that the file may become unusable, but it will still work.
3. Import this new `.cginc` file to your Unity project and let it compile. If all goes well, you will see that Unity knew to compile it to a CgInclude file.

We are now ready to begin creating our own custom CgInclude code. Simply double-click on the CgInclude file that you created in order to open it in MonoDevelop.

How to do it...

With our CgInclude file open, we can begin to enter the code that will get it working with our Surface Shaders. The following code will get our CgInclude file ready for use within our Surface Shaders and allow us to continually add more code to it as we develop more shaders:

1. We begin our CgInclude file with what is called a preprocessor directive. These are statements such as `#pragma` and `#include`. In this case, we want to define a new set of code that will be executed if our shader includes this file in its compiler directives. Enter the following code at the top of your CgInclude file:

```
#ifndef MY_CG_INCLUDE
#define MY_CG_INCLUDE
```

2. We always need to make sure that we close `#ifndef` or `#ifdef` with `#endif` to close the definition check, just like an `if` statement needs to be closed with two brackets in C#. Enter the following code just after the `#define` directive:

```
#endif
```

3. At this point, we just need to fill in the guts of the CgInclude file. So we finish off our CgInclude file by entering the following code:

```
fixed4 _MyColor;

inline fixed4 LightingHalfLamber (SurfaceOutput s, fixed3
lightDir, fixed atten)
{
    fixed diff = max(0, dot(s.Normal, lightDir));
    diff = (diff + 0.5)*0.5;

    fixed c;
    c.rgb = s.Albedo * _LightColor0.rgb * ((diff * _MyColor.rgb) *
atten);
    c.a = s.Alpha;
    return c;
}
#endif
```

4. With this completed, you now have your very first CgInclude file. With just this little bit of code, we can greatly reduce the amount of code that we have to rewrite, and we can begin to store lighting models that we use all the time here so that we never lose them. Your CgInclude file should look similar to the following code shown:

```
#ifndef MY_CG_INCLUDE
#define MY_CG_INCLUDE

fixed4 _MyColor;
```

```
inline fixed4 LightingHalfLamber (SurfaceOutput s, fixed3
lightDir, fixed atten)
{
    fixed diff = max(0, dot(s.Normal, lightDir));
    diff = (diff + 0.5)*0.5;

    fixed c;
    c.rgb = s.Albedo * _LightColor0.rgb * ((diff * _MyColor.rgb) *
atten);
    c.a = s.Alpha;
    return c;
}
#endif
```

There are a couple more steps that we need to complete before we can fully utilize this CgInclude file. We simply need to tell the current shader we are working with to use this file and its code. To complete the process of creating and using CgInclude files, let's complete the next set of steps:

1. If we turn our attention to our shader, we need to tell our `CGPROGRAM` block to include our new CgInclude file so that we can access the code it contains. Modify the directives of our `CGPROGRAM` block to include the following code:

```
CGPROGRAM
#include "MyCGInclude.cginc"
#pragma surface surf Lambert
```

2. Our current shader is currently using the built-in Lambert lighting model, but we want to use the Half Lambert lighting model that we created in our CgInclude. As we included the code from our CgInclude file, we can use the Half Lambert lighting model with the following code:

```
CGPROGRAM
#include "MyCGInclude.cginc"
#pragma surface surf HalfLambert
```

3. Finally, we have also declared a custom variable in our CgInclude file to show that we can set up default variables for our shaders to use. To see this in action, enter the following code in the `Properties` block of your shader:

```
Properties
{
    _MainTex ("Base (RGB)", 2D) = "white" {}
    _DesatValue ("Desaturate", Range(0,1)) = 0.5
    _MyColor ("My Color", Color) = (1,1,1,1)
}
```

4. When we return to Unity, the shader and CgInclude file will compile, and if you do not see any errors, you will notice that in fact we are using our new Half Lambert lighting model and a new color swatch appears in our material's **Inspector**. The following screenshot shows the result of using our CgInclude file:

How it works...

When using shaders, we can include other sets of code using the `#include` preprocessor directive. This tells Unity that we want to let the current shader use the code from within the included file in the shader; this is the reason why these files are called CgInclude files. We are including snippets of Cg code using the `#include` directive.

Once we declare the `#include` directive and Unity is able to find the file in the project, Unity will then look for code snippets that have been defined. This is where we start to use the `#ifndef` and `#endif` directives. When we declare the `#ifndef` directive, we are simply saying, *if not defined, define something with a name*. In this recipe's case, we said we wanted to `#define MY_CG_INCLUDE`. So if Unity doesn't find a definition called `MY_CG_INCLUDE`, it goes and creates it when the CgInclude file is compiled, thereby giving us access to the code that follows. The `#endif` method simply says that this is the end of this definition, so stop looking for more code.

You can now see how powerful this becomes as we can now store all of our lighting models and custom variables in one file and greatly reduce the amount of code that we have to write. The real power is when you can begin to give your shaders the flexibility by defining multiple states of functions in the CgInclude files.

Implementing a Fur Shader

The look of a material depends on its physical structure. The shaders attempt to simulate them, but in doing so, they oversimplify the way light behaves. Materials with a complex macroscopic structure are particularly hard to render. This is the case for many textile fabrics and animal furs. This recipe will show you how it is possible to simulate fur and other materials (such as grass) that are more than just a flat surface. In order to do this, the same material is drawn multiple times over and over, increasing its size every time. This creates the illusion of fur.

The shader presented here is based on the work of *Jonathan Czeck* and *Aras Pranckevičius*:

Getting ready

In order for this recipe to work, you will need two things. The first one is the texture of the fur as it appears from the outside. The second texture will be used to control the distribution of the fur and is deeply connected to the original one. The following image shows a leopard fur (left) and possible control mask (right):

The white pixels in the control mask will be extruded from the original material, simulating a fur. It is important that the distribution of these white pixels is sparse in order to give an illusion that the material is made out of many small hair strands. A loose way to create such a texture is as follows:

1. Apply a threshold to your original texture to better capture patches where the fur is less dense.

2. Apply a noise filter that pixelates the image. The RGB channels of noise must not be dependent in order to produce a black and white result.

3. For a more realistic look, overlay a Perlin noise filter that adds to the variability of the fur.

4. Finally, apply a threshold filter again to better separate the pixels in your texture.

Like all the other shaders before, you will need to create a new standard shader and material to host it.

How to do it...

For this recipe, we can start modifying a **Standard shader**:

1. Add the following `Properties`:

    ```
    _FurLength ("Fur Length", Range (.0002, 1)) = .25
    _Cutoff ("Alpha cutoff", Range(0,1)) = 0.5
    _CutoffEnd ("Alpha cutoff end", Range(0,1)) = 0.5
    _EdgeFade ("Edge Fade", Range(0,1)) = 0.4
    _Gravity ("Gravity direction", Vector) = (0,0,1,0)
    _GravityStrength ("G strenght", Range(0,1)) = 0.25
    ```

2. This shader requires you to repeat the same pass several times. We will use the technique introduced in the *Making your shader world modular with CgIncludes* section to group all the code necessary from a single pass in an external file. Let's start creating a new CgInclude file called `FurPass.cginc` with the following code:

    ```
    #pragma target 3.0

    fixed4 _Color;
    sampler2D _MainTex;
    half _Glossiness;
    half _Metallic;

    uniform float _FurLength;
    uniform float _Cutoff;
    uniform float _CutoffEnd;
    uniform float _EdgeFade;
    ```

```
uniform fixed3 _Gravity;
uniform fixed _GravityStrength;

void vert (inout appdata_full v)
{
    fixed3 direction = lerp(v.normal, _Gravity * _GravityStrength
+ v.normal * (1-_GravityStrength), FUR_MULTIPLIER);
    v.vertex.xyz += direction * _FurLength * FUR_MULTIPLIER *
v.color.a;
}

struct Input {
    float2 uv_MainTex;
    float3 viewDir;
};

void surf (Input IN, inout SurfaceOutputStandard o) {
    fixed4 c = tex2D (_MainTex, IN.uv_MainTex) * _Color;
    o.Albedo = c.rgb;
    o.Metallic = _Metallic;
    o.Smoothness = _Glossiness;

    //o.Alpha = step(_Cutoff, c.a);
    o.Alpha = step(lerp(_Cutoff,_CutoffEnd,FUR_MULTIPLIER), c.a);

    float alpha = 1 - (FUR_MULTIPLIER * FUR_MULTIPLIER);
    alpha += dot(IN.viewDir, o.Normal) - _EdgeFade;

    o.Alpha *= alpha;
}
```

3. Get back to your original shader and add this extra pass after the ENDCG section:

```
CGPROGRAM
#pragma surface surf Standard fullforwardshadows alpha:blend
vertex:vert
#define FUR_MULTIPLIER 0.05
#include "FurPass.cginc"
ENDCG
```

4. Add more passes, progressively increasing FUR_MULTIPLIER. You can get decent results with 20 passes, from 0.05 to 0.95.

Once the shader is compiled and attached to a material, you can change its appearance from the **Inspector**. The `Fur Length` property determines the space between the fur shells, which will be altering the length of the fur. A longer fur might require more passes to look realistic. `Alpha Cutoff` and `Alpha Cutoff End` are used to control the density of the fur and how it gets progressively thinner. `Edge Fade` determines the final transparency of the fur, resulting in a fuzzier look. Softer materials should have a high `Edge Fade`. Finally, `Gravity Direction` and `Gravity Strength` curve the fur shells to simulate the effect of gravity.

How it works...

The technique presented in this recipe is known as Lengyel's concentric fur shell technique or, simply, shell technique. It works by creating progressively bigger copies of the geometry that needs to be rendered. With the right transparency, it gives the illusion of a continuous thread of hair:

Real fur: solid geometry   Shell fur: several spheres

The shell technique is extremely versatile and relatively easy to implement. Realistic, real fur requires not only extruding the geometry of the model, but also altering its vertices. This is possible with tessellation shaders, which are much more advanced and not covered in this book.

Each pass in this Fur Shader is contained in `FurPass.cginc`. The vertex function creates a slightly bigger version of the model, which is based on the principle of normal extrusion. Additionally, the effect of gravity is taken into account so that it gets more intense the further we are from the centre:

```
void vert (inout appdata_full v)
{
    fixed3 direction = lerp(v.normal, _Gravity * _GravityStrength +
    v.normal * (1-_GravityStrength), FUR_MULTIPLIER);
```

```
    v.vertex.xyz += direction * _FurLength * FUR_MULTIPLIER *
v.color.a;
}
```

In this example, the alpha channel is used to determine the final length of the fur. This allows for a more precise control.

Finally, the surface function reads the control mask from the alpha channel. It uses the cutoff value to determine which pixels to show and which ones to hide. This value changes from the first to the final fur shell to match `Alpha Cutoff` and `Alpha Cutoff End`:

```
o.Alpha = step(lerp(_Cutoff,_CutoffEnd,FUR_MULTIPLIER), c.a);

float alpha = 1 - (FUR_MULTIPLIER * FUR_MULTIPLIER);
alpha += dot(IN.viewDir, o.Normal) - _EdgeFade;

o.Alpha *= alpha;
```

The final alpha value of the fur also depends on its angle from the camera, giving it a softer look.

There's more...

The Fur Shader has been used to simulate fur. However, it can be used for a variety of other materials. It works very well for materials that are naturally made of multiple layers, such as forest canopies, fuzzy clouds, human hair, and even grass.

There are many other improvements that can dramatically increase its realism. You can add a very simple wind animation by changing the direction of the gravity depending on the current time. If calibrated correctly, this can give the impression that the fur is moving because of the wind.

Additionally, you can make your fur move when the character is moving. All these little tweaks contribute to the believability of your fur, giving the illusion that it is not just a static material drawn on the surface. Unfortunately, this shader comes at a price: 20 passes are very heavy to compute. The number of passes roughly determines how believable the material is. You should play with fur length and passes in order to get the effect that works best for you. Given the performance impact of this shader, it is advisable to have several materials with different numbers of passes; you can use them at different distances and save a lot of computation.

Implementing heatmaps with arrays

One characteristic that makes shaders hard to master is the lack of a proper documentation. Most developers learn shaders by messing up with the code, without having a deep knowledge of what's going on. The problem is amplified by the fact that Cg/HLSL makes a lot of assumptions, some of which are not properly advertised. Unity3D allows C# scripts to communicate with shaders using methods such as SetFloat, SetInt, SetVector, and so on. Unfortunately, Unity3D doesn't have a SetArray method, which led many developers to believe that Cg/HLSL doesn't support arrays either. This is not true. This post will show you how it's possible to pass arrays to shaders. Just remember that GPUs are highly optimized for parallel computations, and using for loops in a shader will dramatically drop its performance.

For this recipe, we will implement a heatmap, as shown in the following image:
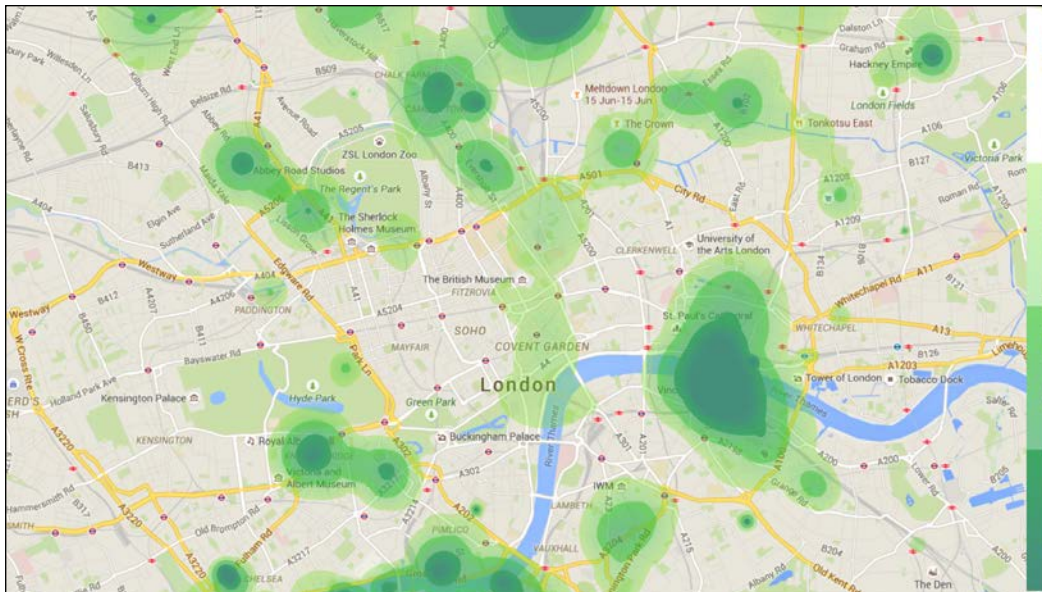
Getting ready

The effect in this recipe creates a heatmap from a set of points. This heatmap can be overlaid on top of another picture, like in the preceding image. The following steps are necessary:

1. Create a quad with the texture that you want to use for the heatmap. In this example, a map of London has been used.

2. Create another quad, and place it on top of the previous one. Our heatmap will appear on this quad.

3. Attach a new material and shader to the second quad.

How to do it...

This shader is quite different from the ones created before, yet it is relatively short. For this reason, the entire code is provided in the following points:

1. Copy this code to the newly created shader:

```
shader " Heatmap" {
    Properties {
        _HeatTex ("Texture", 2D) = "white" {}
    }
    Subshader {
        Tags {"Queue"="Transparent"}
        Blend SrcAlpha OneMinusSrcAlpha // Alpha blend

        Pass {
            CGPROGRAM
            #pragma vertex vert
            #pragma fragment frag

            struct vertInput {
                float4 pos : POSITION;
            };

            struct vertOutput {
                float4 pos : POSITION;
                fixed3 worldPos : TEXCOORD1;
            };

            vertOutput vert(vertInput input) {
                vertOutput o;
                o.pos = mul(UNITY_MATRIX_MVP, input.pos);
                o.worldPos = mul(_Object2World, input.pos).xyz;
                return o;
            }

            uniform int _Points_Length = 0;
            uniform float3 _Points [20];
// (x, y, z) = position
            uniform float2 _Properties [20];
// x = radius, y = intensity

            sampler2D _HeatTex;

            half4 frag(vertOutput output) : COLOR {
                // Loops over all the points
                half h = 0;
```

```
                for (int i = 0; i < _Points_Length; i ++)
                {
                    // Calculates the contribution of each point
                    half di = distance(output.worldPos, _
Points[i].xyz);

                    half ri = _Properties[i].x;
                    half hi = 1 - saturate(di / ri);

                    h += hi * _Properties[i].y;
                }

                // Converts (0-1) according to the heat texture
                h = saturate(h);
                half4 color = tex2D(_HeatTex, fixed2(h, 0.5));
                return color;
            }
            ENDCG
        }
    }
    Fallback "Diffuse"
}
```

2. Once you have attached this script to your material, you should provide a ramp texture for the heatmap. It's important to configure it so that its **Wrap Mode** is set to **Clamp**. The following one has been used for this example:

> If your heatmap is going to be used as an overlay, then make sure that the ramp texture has an alpha channel and the texture is imported with the option, **Alpha is Transparency**.

3. Create a new script called `Heatmaps` using the following code:

```
using UnityEngine;
using System.Collections;

public class Heatmap : MonoBehaviour {

    public Vector3[] positions;
    public float[] radiuses;
    public float[] intensities;
```

```
    public material material;

    void Start ()
    {
        material.SetInt("_Points_Length", positions.Length);
        for (int i = 0; i < positions.Length; i ++)
        {
            material.SetVector("_Points" + i.ToString(),
positions[i]);

            Vector2 properties = new Vector2(radiuses[i],
intensities[i]);
            material.SetVector("_Properties" + i.ToString(),
properties);
        }
    }
}
```

4. Attach the script to an object in your scene, preferably to the quad. Then, drag the material created for this effect to the `material` slot of the script. By doing this, the script will be able to access the material and initialize it.

5. Lastly, expand the positions, radiuses, and intensities fields of your script and fill them with the values of your heatmap. Positions indicate the points (in world coordinates) of your heatmaps, radii indicate their size, and intensities indicate how strongly they affect the surrounding area:
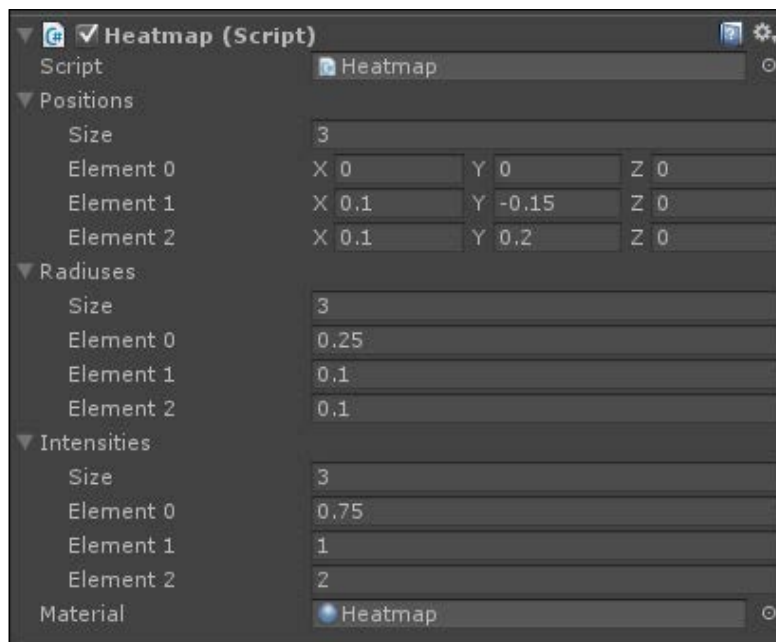
How it works...

This shader relies on things that have never been introduced before in this book; the first one is arrays. Cg allows arrays that can be created with the following syntax:

```
uniform float3 _Points [20];
```

Cg doesn't support arrays with an unknown size: you must preallocate all the space that you need beforehand. The preceding line of code creates an array of 20 elements.

Unity does not expose any method to initialize these arrays directly. However, single elements are accessible using the name of the array (`_Points`) followed by the position, such as `_Points0` or `_Points10`. This currently works only for certain types of arrays, such as `float3` and `float2`. The script attached to the quad initializes the shader's arrays, element by element.

In the fragment function of the shader, there is a similar for loop that, for each pixel of the material, queries all the points to find their contribution to the heatmap:

```
half h = 0;
for (int i = 0; i < _Points_Length; i ++)
{
    // Calculates the contribution of each point
    half di = distance(output.worldPos, _Points[i].xyz);

    half ri = _Properties[i].x;
    half hi = 1 - saturate(di / ri);

    h += hi * _Properties[i].y;
}
```

The `h` variable stores the heat from all the points, given their radii and intensities. It is then used to look up which color to use from the ramp texture.

The shaders and arrays are a winning combination, especially as very few games are using them at their full potential. However, they introduce a significance bottleneck as for each pixel, the shader has to loop through all the points.

Index

Symbols

[PACKT] open source ✳
PUBLISHING   community experience distilled

Thank you for buying
Unity 5.x Shaders and Effects Cookbook

About Packt Publishing

Packt, pronounced 'packed', published its first book, *Mastering phpMyAdmin for Effective MySQL Management*, in April 2004, and subsequently continued to specialize in publishing highly focused books on specific technologies and solutions.

Our books and publications share the experiences of your fellow IT professionals in adapting and customizing today's systems, applications, and frameworks. Our solution-based books give you the knowledge and power to customize the software and technologies you're using to get the job done. Packt books are more specific and less general than the IT books you have seen in the past. Our unique business model allows us to bring you more focused information, giving you more of what you need to know, and less of what you don't.

Packt is a modern yet unique publishing company that focuses on producing quality, cutting-edge books for communities of developers, administrators, and newbies alike. For more information, please visit our website at `www.packtpub.com`.

About Packt Open Source

In 2010, Packt launched two new brands, Packt Open Source and Packt Enterprise, in order to continue its focus on specialization. This book is part of the Packt Open Source brand, home to books published on software built around open source licenses, and offering information to anybody from advanced developers to budding web designers. The Open Source brand also runs Packt's Open Source Royalty Scheme, by which Packt gives a royalty to each open source project about whose software a book is sold.

Writing for Packt

We welcome all inquiries from people who are interested in authoring. Book proposals should be sent to `author@packtpub.com`. If your book idea is still at an early stage and you would like to discuss it first before writing a formal book proposal, then please contact us; one of our commissioning editors will get in touch with you.

We're not just looking for published authors; if you have strong technical skills but no writing experience, our experienced editors can help you develop a writing career, or simply get some additional reward for your expertise.

Unity 5.x Cookbook

ISBN: 978-1-78439-136-2   Paperback: 570 pages

Over 100 recipes exploring the new and exciting features of Unity 5 to spice up your Unity skill set

1. 2D and 3D game development skills covered.

2. The authors are very experienced and give you a plethora of knowledge in this book.

3. This book will be useful for both base level and skilled developers.

Learning Unity iOS Game Development

ISBN: 978-1-78439-980-1   Paperback: 230 pages

Build exciting games with Unity on iOS and publish them on the App Store

1. Coming from an eminent author, this book offers you the fine technicalities of how iTunes and Unity can be used to build a fully interactive and reliable game for the App Store.

2. This book will not only help you learn to build and deploy the App but also give it for review to the App store.

3. The approach of moving from basic to advance skills, will help you catch learn easily.

Please check **www.PacktPub.com** for information on our titles

Unreal Engine Lighting and Rendering Essentials

ISBN: 978-1-78528-906-4   Paperback: 242 pages

Learn the principles of lighting and rendering in Unreal Engine

1. This book will not only help you build and deploy the app, but will also help you maintain the game performance on cross platforms.

2. The author is eminent, so you will learn tips and tricks to build cross-platform mobile and desktop apps from scratch.

3. Easy-to-follow examples with images to create stunning visual effects.

Unreal Development Kit Game Programming with UnrealScript [Video]

ISBN: 978-1-84969-632-6   Duration: 02:23 hrs

Kick-start your career in game development with UnrealScript and the UDK

1. Step-by-step guide on how to set up the UDK and create a game with UnrealScript.

2. Explore core UnrealScript features and configurations.

3. Ideal for newcomers to UnrealScript.

Please check **www.PacktPub.com** for information on our titles