

DRAFT REVISION DRAFT REVISION

KENWRIGHT

INTRODUCTION TO COMPUTER GRAPHICS AND THE VULKAN API

TECHNICAL BOOK



Introduction to Computer Graphics and the Vulkan API

Kenwright

Copyright © 2017 Kenwright
All rights reserved.

No part of this book may be used or reproduced in any manner whatsoever without written permission of the author except in the case of brief quotations embodied in critical articles and reviews.

BOOK TITLE:

Introduction to Computer Graphics and the Vulkan API

ISBN-13: 978-1-548-61617-5

ISBN-10: 1-548-61617-6

The author accepts no responsibility for the accuracy, completeness or quality of the information provided, nor for ensuring that it is up to date. Liability claims against the author relating to material or non-material damages arising from the information provided being used or not being used or from the use of inaccurate and incomplete information are excluded if there was no intentional or gross negligence on the part of the author. The author expressly retains the right to change, add to or delete parts of the book or the whole book without prior notice or to withdraw the information temporarily or permanently.

Revision: 035628082017

First published, July 2017

DRAFT

Contents

1	<i>Introduction & Overview</i>	15
1.1	<i>Getting Started</i>	15
1.2	<i>Computer Graphics</i>	15
1.3	<i>Aim of this Book</i>	16
1.4	<i>Prerequisite (Setting-up Vulkan)</i>	17
1.5	<i>Summary</i>	18
2	<i>Background (OpenGL and Vulkan)</i>	19
2.1	<i>Introduction</i>	19
2.2	<i>History of Vulkan</i>	21
2.3	<i>11 Steps</i>	22
2.4	<i>Naming Convention</i>	24
2.5	<i>Exercises</i>	24
2.5.1	<i>Chapter Questions</i>	24
3	<i>Mathematics</i>	26
3.1	<i>Introduction</i>	26
3.2	<i>Vector</i>	26
3.2.1	<i>What is a Vector?</i>	26
3.2.2	<i>Vectors and Points</i>	27
3.2.3	<i>Vector3</i>	27
3.2.4	<i>Dot Product</i>	28
3.2.5	<i>Cross Product</i>	29
3.2.6	<i>Reconstructing Angles from Positions</i>	30
3.2.7	<i>Plane Equation</i>	30

3.2.8	Support Function	31
3.3	Matrix	31
3.3.1	Why Matrices?	31
3.3.2	Column or Row Major	31
3.3.3	A 4x4 Matrix	32
3.3.4	Creating a Matrix	33
3.3.4.1	Identity Matrix	34
3.3.4.2	Translation Matrix	34
3.3.4.3	Scale Matrix	35
3.3.4.4	Rotation Matrix	35
3.3.5	Matrix-Matrix Multiplication	36
3.3.6	'Pure' Rotation	37
3.3.6.1	Orthogonal Matrices (Useful-Axis)	37
3.3.6.2	Transpose and Inverse	37
3.3.7	Transforming a Vector	37
3.3.7.1	Little Test	38
3.3.8	Matrix Inversion	39
3.4	Quaternion	41
3.4.1	Why Quaternions?	41
3.4.2	Unit-Quaternion (Always)	41
3.4.3	Creating a Quaternion	41
3.4.3.1	Quaternion from Axis-Angle	42
3.4.3.2	Quaternion to Axis-Angle	42
3.4.3.3	Quaternion to Matrix	43
3.4.3.4	Quaternion from Matrix	43
3.4.4	Quaternion-Quaternion Multiplication	45
3.4.5	Quaternion Inverse (Conjugate)	46
3.4.6	Transform a Vector by a Quaternion	46
3.5	Summary	47
3.6	Exercises	47
3.6.1	Chapter Questions	48
4	Graphical Principles	49

4.1	<i>Basic Types</i>	49
4.2	<i>Transforms</i>	50
4.2.1	<i>Homogeneous Coordinates (or Projective Coordinates)</i>	52
4.2.2	<i>Normalized Device Coordinates (NDC)</i>	53
4.2.3	<i>Eye Coordinates</i>	54
4.2.4	<i>Projection</i>	55
4.2.4.1	<i>Orthogonal</i>	56
4.2.4.2	<i>Perspective</i>	56
4.2.5	<i>Camera (LookAt)</i>	59
4.3	<i>Primitives</i>	61
4.3.1	<i>Backface Culling (Clockwise/Counter-Clockwise)</i>	62
4.4	<i>Data/Geometry</i>	63
4.5	<i>Drawing Principles</i>	64
4.6	<i>Programmable Graphics & Shaders</i>	64
4.7	<i>Exercises</i>	67
4.7.1	<i>Chapter Questions</i>	67
5	<i>Shaders</i>	70
5.1	<i>Introduction</i>	70
5.1.1	<i>Anatomy of Shaders</i>	72
5.2	<i>Link between Vulkan and Shaders</i>	74
5.3	<i>Linking data to Uniforms</i>	75
5.3.1	<i>Qualifiers</i>	76
5.3.2	<i>Uniforms</i>	76
5.3.3	<i>Varying</i>	76
5.4	<i>Developing Shaders</i>	77
5.5	<i>Summary</i>	79
5.6	<i>Exercises</i>	79
5.6.1	<i>Chapter Questions</i>	79
6	<i>Programming (11 Steps)</i>	81
6.1	<i>(Step 1 & 2) Initializing Vulkan (Instance Creation)</i>	84
6.2	<i>Debugging</i>	87

6.3	(Step 3) Device(s)	90
6.4	(Step 4) Swap-Chain	92
6.5	(Step 5) FrameBuffer & Render-Pass	95
6.6	(Step 6) Command-Buffers	100
6.7	Buffers	103
6.8	Memory Allocations	104
6.9	(Step 7) Vertex Buffer (Data)	105
6.10	(Step 8) Shaders & Uniforms	109
6.11	Synchronization	114
6.12	(Step 9) Descriptors	115
6.13	(Step 10) Graphics Pipeline	117
6.14	Images And ImageView	122
6.15	(Step 11) Render Loop	123
6.16	Exercises	127
6.16.1	Chapter Questions	127
6.16.2	Practical Exercises	128
7	Texturing	129
7.1	Introduction	129
7.2	Texture Coordinates	130
7.3	Texturing Objects with Vulkan	131
7.4	Texture Blur (Gaussian)	137
7.5	Swirling	138
7.6	Pixelate	139
7.7	Edge-Detection	140
7.8	Black & White	142
7.9	Fish-Eye	143
7.10	Height-Map	144
7.11	Exercises	147
7.11.1	Chapter Questions	148
7.11.2	Practical Exercises	148
8	Lighting	149
8.1	Why is Lighting so Important?	149

8.2	<i>Global vs Local</i>	150
8.3	<i>Lighting Components (Diffuse, Specular and Ambience)</i>	150
8.4	<i>Diffuse (Lambert's Law)</i>	151
8.4.1	<i>Beware of the Sign</i>	152
8.4.2	<i>Normals to 'World' Space</i>	152
8.5	<i>Flat, Gouraud and Phong Shading</i>	155
8.6	<i>Lighting Calculations</i>	157
8.6.1	<i>Phong Shading</i>	157
8.6.2	<i>Gouraud Shading</i>	161
8.7	<i>Single and Multiple Lights</i>	163
8.7.1	<i>Reflection</i>	165
8.7.2	<i>Blinn-Phong</i>	165
8.8	<i>Light Types (Directional, Point, Cone,)</i>	165
8.9	<i>Distance Attenuation</i>	166
8.10	<i>Exercises</i>	166
8.10.1	<i>Chapter Questions</i>	167
8.10.2	<i>Practical Exercises</i>	167
9	<i>Geometry Shader</i>	168
9.1	<i>Pass-Through</i>	168
9.2	<i>Adding Geometry (Shells)</i>	172
9.3	<i>Wireframe Normals</i>	173
9.4	<i>Billboarding</i>	175
9.5	<i>Exercises</i>	178
9.5.1	<i>Chapter Questions</i>	178
9.5.2	<i>Practical Exercises</i>	178
10	<i>Cube Maps, SkyBox & Reflection</i>	179
10.1	<i>Introduction</i>	179
10.1.1	<i>Cube Map Coordinates (Normal)</i>	180
10.2	<i>SkyBox</i>	182
10.2.1	<i>Disable Depth Writing</i>	183
10.3	<i>Reflection (with Cube Maps)</i>	183

10.4 Exercises	184
10.4.1 Chapter Questions	184
10.4.2 Practical Exercises	184
11 Fog & Depth	185
11.1 Introduction	185
11.2 Exercises	188
11.2.1 Practical Exercises	188
12 Bump, Normal, Parallax and Displacement Mapping	189
12.1 Introduction	189
12.2 Coordinate Spaces (Tangent Space)	189
12.3 Normal Mapping (World Space)	190
12.4 Displacement Mapping	192
12.5 Parallax Mapping	192
12.6 Exercises	194
12.6.1 Chapter Questions	194
13 Instancing	196
13.1 Introduction	196
13.2 Instance ID	197
13.3 Instance Uniforms	197
13.4 Exercises	201
13.4.1 Practical Exercises	201
14 Tessellation	202
14.1 Introduction	202
14.2 Pass-Through	203
14.3 gl_TessLevel	207
14.4 Quads	207
14.5 Summary	208
14.6 Exercises	209
14.6.1 Practical Exercises	209
15 Shadows	210

15.1	<i>Introduction</i>	210
15.1.1	<i>Simple and Beautiful (Shadow Mapping)</i>	210
15.2	<i>Theory</i>	211
15.3	<i>3D World Position to 2D and Depth</i>	212
15.3.1	<i>Depth Comparison</i>	213
15.3.2	<i>Linear Depth Buffer</i>	214
15.3.3	<i>Orthographic or Perspective Projection Matrix</i>	214
15.3.4	<i>Visual Artefacts</i>	215
15.4	<i>Discussion</i>	215
15.5	<i>Filtering, Smoothing and Soft-Shadows</i>	216
15.6	<i>Transparency</i>	217
15.7	<i>Debugging</i>	217
15.8	<i>Image Resolution</i>	217
15.8.1	<i>Typical Problems</i>	217
15.8.2	<i>Multiple Shadow Maps (Resolution)</i>	219
15.8.3	<i>Clipping and Soft Fall-Off</i>	219
15.8.4	<i>Scaling and Offsetting the Stored Z-Depth</i>	219
15.9	<i>Implementation</i>	220
15.10	<i>Conclusion</i>	222
15.11	<i>Exercises</i>	222
15.11.1	<i>Chapter Questions</i>	222
15.11.2	<i>Practical Exercises</i>	223
16	Skinning	224
16.1	<i>Introduction</i>	224
16.2	<i>Hierarchies & Transforms</i>	225
16.2.1	<i>Skeletal Animation</i>	226
16.3	<i>Linear Deformation</i>	226
16.3.1	<i>Algorithm</i>	226
16.4	<i>Skinned Tubular Mesh (Hand-Crafted)</i>	227
16.5	<i>Summary</i>	233
16.6	<i>Exercises</i>	234
16.6.1	<i>Chapter Questions</i>	234

16.6.2	<i>Practical Exercises</i>	234
17	<i>Post-Processing & Deferred Rendering</i>	235
17.1	<i>Introduction</i>	235
17.1.1	<i>Post-Processing</i>	236
17.1.2	<i>Deferred Rendering</i>	236
17.2	<i>Multiple Buffers & Phases</i>	237
17.3	<i>Summary</i>	244
17.4	<i>Exercises</i>	244
17.4.1	<i>Chapter Questions</i>	245
17.4.2	<i>Practical Exercises</i>	245
18	<i>Good Practices</i>	246
19	<i>Reading Lists</i>	248
20	<i>Troubleshooting and Q&A</i>	249
21	<i>Appendix</i>	251
21.1	<i>Simple Shader Listing</i>	251
	<i>Bibliography</i>	252
	<i>Index</i>	256

*Dedicated to those who appreciate the
beauty and complexity of computer
graphics*

DRAFT REVISION DRAFT REVISION

DRAFT

1

Introduction & Overview

1.1 Getting Started

This book provides an introductory guide to getting started with computer graphics using the Vulkan API. The book focuses on the practical aspects with details regarding previous and current generation approaches, such as, the shift towards more efficient multi-threaded solutions. The book has been formatted and designed, so whether or not you are currently an expert in computer graphics, actively working with an existing API (OpenGL), or completely in the dark about this mysterious topic, this book has something for you. If you're an experienced developer, you'll find this book a light refresher to the subject, and if you're deciding whether or not to delve into graphics and the Vulkan API, this book may help you make that significant decision. This is an ambitious book, but not unrealistic, and we know that computer graphics is a little bit of an art and involves a variety of skills and abilities. There is so much more to know than this book is able to present - however, it presents the essential facts of the subject with a high-level introduction to the core components and their mechanics. It's not that we necessarily excluded anything critical from this book, but it would be unrealistic to try and cover every possible aspect in a single text. For the sake of practicality, we discuss a variety of important aspects of the Vulkan API, such as, the differences between traditional graphical API paradigms, setting up a Vulkan project, performance factors and real-world applications and examples.

1.2 Computer Graphics

Computer graphics is an exciting and important multi-discipline subject with applications in:

Sample program listings and support material are provided online to complement the book.



Figure 1.1: Designed and maintained by Khronos Group for high performance on rendering and compute [6].

- visualisation solutions,
- video games,
- image and video processing,
- graphical modeling,
- animation,
- augmented and virtual reality,
- production/tool optimisation (CPU/GPU),
- real-time solutions,
- rendering & simulation,
- visual effects,
- user interaction
- robotics
- ...

Computer graphics covers topics from extraction and visualisation to generation and manipulation in both 2-dimensional and 3-dimensional contexts. In this book, you'll focus primarily on 3-dimensional visual solutions. However, you'll still require and apply 2-dimensional principles like texture manipulation and mapping to pixel and screen space effects (e.g., blurring, edge detection and smoothing). You'll discover that computer graphics gives you the power to create worlds of infinite possibilities (e.g., from chocolate cities 'choco-land' to real-world locations like London) or help visualise complex problems (like structural stress in buildings or the workings of internal organs in the human body). The implementations can range in complexity as well - from a simple single triangle with no lighting or texturing requiring a couple of hundred lines of code to a complete renderer engine that's able to display realistic human models accurately down to the hairs on their head (requiring thousand or more lines of code with dozens of different shaders and optimisations). What is more, these solutions may be off-line taking minutes or days to calculate or microseconds for real-time interactive virtual environments (video games).

1.3 Aim of this Book

This book aims to introduce computer graphics programming in a practical context while addressing a number of crucial questions with regard to 'another' graphical application programming interface (API), for example:

- ✓ What exactly is Computer Graphics and the Vulkan API?
- ✓ Why is understanding the 'differences' between the API important?
- ✓ How do you get started programming a graphical application with Vulkan?

Name: 'VULKAN'

The Vulkan API was a ground-up re-design of the popular OpenGL API, previously referred to as the 'Next Generation OpenGL' (GLNext) initiative - however, over time it was decided to rename the API to 'Vulkan' to help emphasise the radical change in thinking, i.e., the aim to provide applications low-level direct control over processor (GPU/APU/CPU) acceleration for maximized performance and predictability.

At the end of this book, you should feel comfortable enough to work with the Vulkan API (i.e., create, customize and generate a variety of simple graphical applications). You should be able to explain the core components of the API, and importantly, why and how they fit together to accomplish the necessary graphical technique [12, 9].




Figure 1.2: Vulkan has a steep learning curve initially - but over time the benefits and freedom provided by the API are rewarded compared to existing solutions (greater optimisations and customisability).

- ✓ Understanding where and why a graphical program 'fails' - e.g., perform worse than current or existing graphical API
- ✓ Dealing with problems, such as, cross-platform, memory leaks, graphical issues, rapid prototyping, versions, ...
- ✓ How to work effectively on complex projects with Vulkan
- ✓ Background introduction to the history of different graphical API
- ✓ Revision on basic graphical principles and techniques (shaders, lighting, transforms, triangles)
- ✓ Managing Vulkan API (structured modular programming)
- ✓ Implement a basic graphical application from the ground up using the native Vulkan API
- ✓ Essential graphical principles and how to implement them with Vulkan
- ✓ How to implement popular graphical effects (e.g., lighting, bump maps, instancing and texturing)

1.4 Prerequisite (Setting-up Vulkan)

Pre-requisites to working with Vulkan The computer graphics samples in this book are build around the Vulkan API - hence, to implement and run the examples you'll need to download and install one of the Vulkan SDK libraries on your machine.

To download and install the necessary Vulkan API drivers and SDK (if you don't already have them installed on your system) is very straightforward. For example, a popular Vulkan API SDK is:

Lunar-G (<http://lunarg.com/>)

In addition, you'll need to have a basic understanding of core programming principles (e.g., functions, pointers, libraries and the ability



Figure 1.3: The LunarG SDK provides the development and runtime components for building, running, and debugging Vulkan applications. This includes the Vulkan loader, Vulkan layers, debugging tools, SPIR-V tools, the Vulkan run time installer, documentation, samples, and demos.

to read simple computer programs written in C, C++ or Java). While basic knowledge of computer graphics concepts would be beneficial (for example, framebuffers and refresh rate), however, it's not required, as you'll be guided through the process of writing a basic graphic applications from the ground-up.

The practical examples and listings in the book are implemented using C/C++.

1.5 *Summary*

These are exciting times for computer graphics. With advancements in technologies and hardware we're seeing breakthroughs in realism and creativity. The material to create amazing effects is freely available (e.g., free open source libraries, online tutorials and free 3-dimensional models). While computer graphics can seem daunting and difficult initially - especially if your mathematics is a bit rusty - the rewards at the end are well worth the time and effort.

2

Background (OpenGL and Vulkan)

2.1 Introduction

Since OpenGL was first released in 1992 by Silicon Graphics Inc., it has been widely adopted across the world by industry as well as academia. The API reduced the engineering complexities and what followed over the coming years was the birth of visually breathtaking solutions that captured the imagination (both visually and inspirationally). The ability to accomplish stunning computer generated images was made possible through further technological advancements. Computer graphics has become increasingly challenging using conventional approaches and expectations have and continue to grow, especially in areas involved with films, games and virtual reality. One specific challenge is the ability to exploit the advancements in rapidly changing technologies. For example, despite the ready availability of multiple high performance graphics cards, the limitations of existing libraries has made it difficult if not impossible to exploit the full potential of the hardware (distributing the workload for processing and rendering high fidelity images in real-time across multiple devices efficiently [4]). While parallel processing paradigms have become an attractive solution in recent years, with multiple cores and threads working together to offering tremendous performance gains, developing parallel applications that exploit these parallel speed-ups efficiently and reliably is a significant challenge.

Vulkan is an exciting multi-platform cross-language graphical and compute interface that exploits the latest 'parallel' hardware architectures. Vulkan provide you and developers with a powerful interface to create stunning visuals for a wide range of applications. Vulkan still follows the same original 'OpenGL' initiatives, i.e., to develop a high quality open source, cross-platform API (Mac, Windows, Linux, Android, Solaris and FreeBSD). OpenGL has come a long way and

Khronos launched the Vulkan 1.0 specification on February 16th, 2016.



Figure 2.1: The Khronos Group is a non-profit, member-funded consortium focused on the creation of royalty-free open standards for parallel computing, graphics and vision processing on a wide variety of platforms and devices. Currently there are 100+ industry-leading company members across the globe.

done amazingly well over the last 25 years (Figure 2.2). Be that as it may, it is time for a major update. As the original OpenGL API follows a state machine architecture this ties the API to a single on-screen context. In addition the OpenGL API is blind to everything the GPU is doing (optimised and managed within the driver - and hidden from the developer). Vulkan takes a different approach - following an object-based API with no global state so all state concepts are localized to a Command-Buffer (you'll learn about Command-Buffers in Section 6.6). What is more Vulkan is more explicit about what the GPU is doing (less hiding what is happening within the driver).

API improvements:

- Explicit Control
- Multi-Threading Friendly
- Direct State Access (DSA)
- Bindless Graphics
- Framebuffer Memory Info
- Texture Barrier
- Acceleration for applications (e.g., Browsers, WebGL, ..)

The principle of explicit control, means you promise to tell the driver every detail. So the driver doesn't have to guess or make assumptions. In return, the driver is more streamlined and efficient (does what you asked for when you asked for it quickly). For instance, memory management in Vulkan gives the control to the application (total memory usage is more visible and simplifies operations, such as as for streaming data). Remember, the application is in charge (so doing it correctly is your responsibility).

While the latest OpenGL graphical API (known as Vulkan) might seem like another iteration, it is well worth learning or even reviewing. At the same time, Vulkan is in its first release (revision 1.0) - and possesses a huge number of changes/improvements compared to any previous update. Importantly, these improvements should not be ignored, as they offer possibilities that were previously not feasible. These key features will help you get more out of GPUs. However, to gain improvements it is important you understand the differences (i.e., applications need to be written differently to utilize these additional features and control - OpenGL ! = Vulkan). As shown in Figure 2.3, you'll notice the shift of power between the driver and the application. Vulkan's abstraction means your application is much closer the hardware compared to traditional APIs. Your application is driving the hardware directly, while leaving just enough abstraction to make things portable. You're not being second-guessed by the driver, while at the same time you're not being first-guessed either. You now have

OpenGL to Vulkan Progression




Figure 2.2: Evolution of the OpenGL API to the most recent incarnation known as 'Vulkan'.

all the control you need to get the best out of your hardware. If it doesn't go fast in Vulkan, it's your fault (of course, remember, with great power comes great responsibility).

A few of the "big tick" items with Vulkan is:

- Explicit control,
- Support for multi-core/threading,
- Predictability,
- Texture formats, memory management, and syncing are client-controlled
- Vulkan drivers do no error checking and
- Bandwidth efficiency.

2.2 History of Vulkan

The Vulkan API was designed and is maintained by the Khronos Group to meet current and future demands for achieving high performance rendering and compute solutions. The Vulkan API achieves this by allowing greater low level control (explicitly) - moving away from 'default' parameters/assumptions set within the driver. The developer has to manage the memory, resource updates, batching, scheduling, ... Hence, the Vulkan API initially seems verbose and complicated due to the large amount of initiation and management (through functions, parameters and structures), yet this is crucial for Vulkan's success. It should also be noted, that DirectX 12 from Microsoft follows a similar design to Vulkan (explicit low level control). For instance, previously,

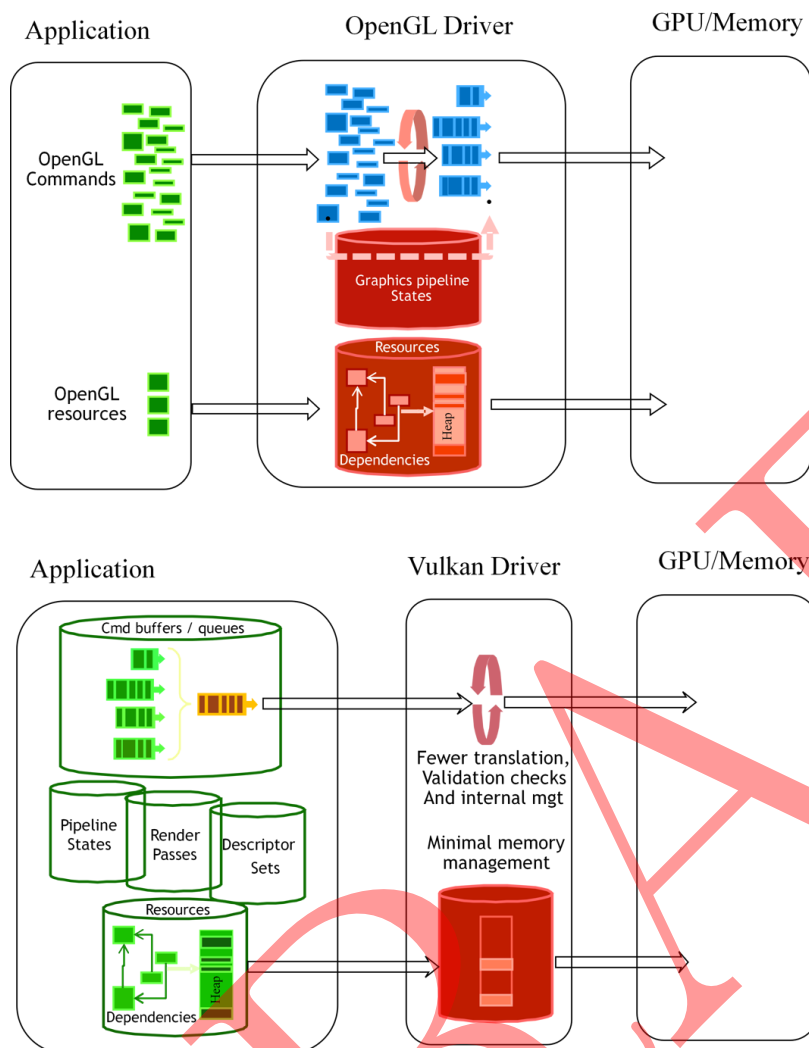


Figure 2.3: High-level view of what has changed between OpenGL and Vulkan. Importantly, the shift in power and work from the driver to the application. The application is now responsible for a number of crucial tasks that were previous hidden to the developer, such as, memory management, resources and command-buffers. This modifications provides a more 'streamline' and efficient solution (bringing the application developer nearer to the hardware).

'OpenGL' did not address multi-threading and was not designed to support the concurrent and parallel paradigm which would be a serious problem in todays multi-core multi-threaded environment. However, the Vulkan API is designed to exploit these multi-threaded environments (and is how it is able to outperforms previous API).

2.3 11 Steps

You'll see an overview of the essential components in most Vulkan graphical applications in Figure 2.4. To complement this programming section, the components have been grouped into 11 distinct steps. From step 1 which initializes the application and creates the window -

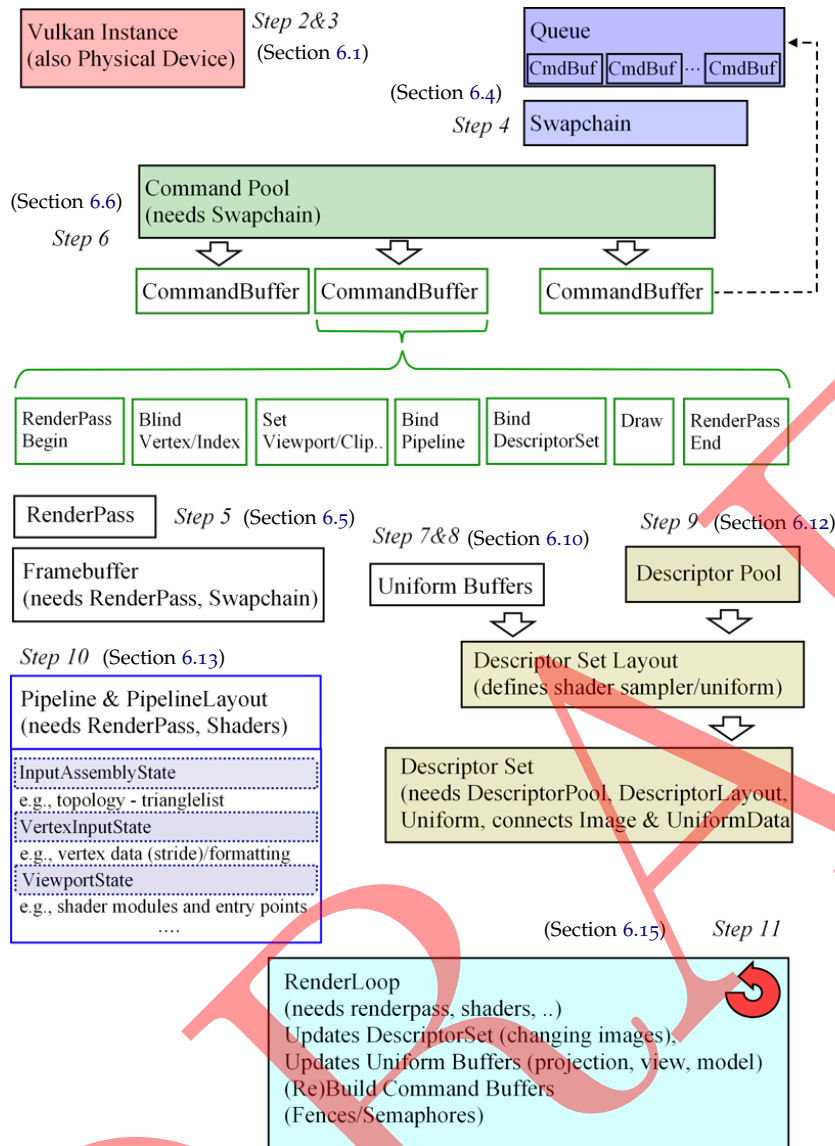


Figure 2.4: Decomposition of the most popular Vulkan components available in common graphical applications. The illustration provides an overview of the elements and how they fit together in a simple application. You'll review and discuss each of the components in the following sections (i.e., from initializing Vulkan and the physical device to building a Command-Buffer and rendering to the screen).

to the final stage 11 which performs the updates and rendering. Each step provides a self-contained set of material which is used to breakup an otherwise complex system. The steps enable you to progress in an orderly manner as you learn the rational behind each of the elements and how they fit together (e.g., swap-chains and command-buffer). Each step builds upon the previous steps and enable you to incrementally build a complete graphical application using the Vulkan API that utilizes all of the features (in a modular manner).

Briefly, the 11 steps are:

1. Initialize application and create a window (operating system specific)
2. Initialize Vulkan (Vulkan Instance)
3. Initialize Device (e.g., GPU)
4. Create Swap-Chain (managing the display output)
5. FrameBuffer & Render-Pass (output image surfaces)
6. Command-Buffer & Command-Pool (essential for graphics - as all draw commands need to be in a command-buffer)
7. Vertex Data (geometry you'll be drawing)
8. Shaders & Uniform Buffers (essential for graphics to have a vertex and fragment shader in addition to any parameters/passing of data to the shaders)
9. Descriptors (glue that holds everything together, such as, the shaders and geometry vertex data)
10. Graphics Pipeline (connecting everything together and enabling features)
11. Render Loop (drawing/syncing)

2.4 Naming Convention

The Vulkan API variables and functions follow a consistent naming convention. While both variables and functions start with the letters 'vk', you need to remember, functions start with a lowercase letter while variables start with an uppercase letter, for example:

Function: **vkCreateInstance(..)**

Variable: **VkResult**

In the example listings that follow in subsequent sections, the Vulkan API functions and structures have been emphasised to help you identify the key elements.

2.5 Exercises

2.5.1 Chapter Questions

Question When was the Vulkan 1.0 specification released?

Question What is the naming convention for Vulkan variables and functions?

Question What is the root methodology behind Vulkan compared to previous graphical API?

Question What is a ray tracing algorithm and how does it compare to a rasterization approach?

3

*Mathematics*3.1 *Introduction*

There are a few fundamental mathematical concepts that are indispensable when working with computer graphics and geometric systems (e.g., vectors and matrices including concepts such as normals and the dot product). The main mathematical tools that you'll review in this chapter are:

- Vectors
 - Dot
 - Cross
- Matrices
 - Transforms
- Quaternions
 - Rotations

Hence, you'll briefly review the workings and implementation details for each mathematical concept. However, in practice you may prefer to use existing pre-written libraries (e.g., glm), but be careful you don't get caught with problems, such as, "handed" convention (i.e., left or right handed differences) or function speed-up hacks, which can cause large numerical errors.

3.2 *Vector*3.2.1 *What is a Vector?*

A vector represents a mathematical or physical direction and length (or magnitude) and is depicted by an arrow (with the arrow symbolizing the direction and the length of the arrow the magnitude). For example, the wind has a direction and speed, as shown on weather maps. You

$$M = \begin{bmatrix} 1 & 5 & 9 & 13 \\ 2 & 6 & 10 & 14 \\ 3 & 7 & 11 & 15 \\ 4 & 8 & 12 & 16 \end{bmatrix}$$

$$= [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16]$$

Figure 3.1: A large majority of computer graphics principles requires you understand common mathematical topics (e.g., matrix and vector mathematics, linear algebra and trigonometry).

can have different dimensions of vectors (i.e., 1D, 2D, 3D, 4D, ...). Note, a 1D vector would just be a scalar float. However, you'll primarily be dealing with 3D vectors composed of an x, y, and z). If you want a 2D vector just remove the z. In code, a vector is nothing more than an array of variables (e.g., `float[3]`). So that you can distinguish the dimensions of your vector, you'll add the number to the end, e.g., "Vector3" and "Vector2". You'll use a class or structure to represent your vector since it makes the code more readable and you'll be able to exploit operator overload.

Let's get this out the way right at the start - what is the difference "in code" between a "Point" and a "Vector"? For example, a Vector3 and a Point3 structure. The answer: Nothing! The code is identical, except for the name of course.

In short, don't make work for yourself. Don't create structures or variables that accomplish the same task but use different names. For example, you might be tempted to use Vectors for direction and Points for position. However, the name of the variable should be sufficient for a detailed description of what the variable is does. For example, Listing 3.1 shown below:

Listing 3.1: Application of Vector3 (e.g., Positions and Directions)

```
1 Vector3 position;  
2 Vector3 direction;  
3 Vector3 velocity;  
4 Vector3 force;
```

3.2.2 *Vectors and Points*

A 3D vector differ from a 3D point tuple (x,y,z) in 3D game mathematics. They are different 'mathematically', while you represent them the same pragmatically. The difference is that a vector is an algebraic object that may or may not be given as a set of coordinates in some space. A point is just a point given by coordinates. Generally, you can conflate the two. An intuitive way to think about the association between a vector and a point is that a vector tells you how to get from the origin (that one point in space to which you assign the coordinates $\langle 0, 0, 0 \rangle$) to its associated point. While in code they may appear the same (e.g., Vector3 for a variable), ensure you know 'mathematically', what that variable stands for (i.e., a 3d position in space or a vector direction with magnitude).

3.2.3 *Vector3*

Listing 3.2: Unsophisticated Vector3 Class Implementation

```

1  class Vector3
2  {
3  public:
4      float x;
5      float y;
6      float z;
7  };

```

Without vectors, basic geometric calculations would be very complex, difficult to read, and time consuming when debugging. Furthermore, once you understand vectors and how to use them, in combination with the various routines (e.g., dot and cross product) you'll be able to tackle daunting geometric problems without even breaking a sweat.

3.2.4 Dot Product

In a nutshell, the dot product is amazing. It's flexible, computationally efficient, and straightforward to use. To summarize, here are the main features the dot product offers:

- Magnitude squared distance of two vectors is the dot product operation
- Sign of the result of the dot product enables us to determine if vectors are facing towards or away from one-another

Word of caution, this operation does not require the vectors to be of unit-length, so you can avoid the cost of normalizing the vectors

- Cosine of the angle between two vectors

Warning, the vectors must be of unit-length, also the 'sign' of direction is not provided (i.e., only provides the shortest path and doesn't tell us the direction)

- Project a vector onto another vector

Note, the vector you are projecting onto should be a unit-vector

- Dot product doesn't involve any complex computational operations (e.g., sqrt, sin) and can be performed using simple multiplication and addition

The dot product can be speeded-up on modern hardware technology since operations such as multiplication can be performed in parallel (e.g., dot product can be done in a single instruction on some processors)

The dot product returns a single scalar value and can easily be implemented, as shown in Listing 3.3.

Listing 3.3: Unsophisticated Vector3 Dot Product Implementation.

```

1  inline
2  float Dot( const Vector3& A, const Vector3& B)
3  {
4      return ( A.x * B.x + A.y * B.y + A.z * B.z );

```



```
5 };
```

3.2.5 Cross Product

While the dot product may come first for usefulness and features the cross product is not far behind for providing a similar list of useful operations. The cross product of two vectors (a and b) is written as $a \times b$ and returns a vector. In three dimensional space, the cross product of two vectors is a vector that is “perpendicular” to both the initial vectors.

The main features of the cross product are:

- Calculates a vector perpendicular to two unit vectors
- Can be combined with the dot product to provide a direction of rotation between two unit vectors (i.e., dot product provides the angle between the two unit vectors but doesn’t provide the direction of rotation)
- Cross product doesn’t involve any complex computational operations (e.g., sqrt, sin) and can be performed using simple multiplication, addition and subtraction

Note, modern hardware can perform the cross product in a single operation due to the parallel nature of the operation

- the area of a parallelogram with sides AB and AC is equal to the magnitude of the cross product of vectors representing two adjacent sides (while the area of a triangle would be half that)

The direction of the resulting vector cross product is given by the “right-hand” convention. With your right hand, if your first finger is vector a , and your second finger is vector b , then your thumb is the cross product result $a \times b$. The implementation details in code are shown in Listing 3.4.

Listing 3.4: Unsophisticated Vector3 Cross Product Implementation.

```
1 inline
2 Vector3 Cross( const Vector3& A, const Vector3& B)
3 {
4     Vector3 vec;
5     vec.x = (A.y*B.z) - (B.y*A.z);
6     vec.y = (A.z*B.x) - (B.z*A.x);
7     vec.z = (A.x*B.y) - (A.y*B.x);
8     return vec;
9 };
```

Be warned that the cross product is “non-commutative”, i.e., $(a \times b)$ does “not” equal $(b \times a)$.

3.2.6 Reconstructing Angles from Positions

Given a set of points, you can reconstruct the link's angle from the positional information as shown in Figure 3.2. This can be valuable when you have a set of animation capture points, and you want to reconstruct the articulated character's bone structure (i.e., rigid bodies and joint angles).




Figure 3.2: **Direction to Angle** - Illustrating how to reconstruct angles from points. You subtract P_3 and P_1 from P_2 to construct two vector directions. Dividing them by their magnitudes normalizes them (i.e., to their unit length values). Finally, the dot product of the two vectors gives us the cosine of the angle between them.

3.2.7 Plane Equation

The plane equation is a mathematical method for representing the valuable concept of a planar surface. The plane equation is probably one of the most useful tools in your algorithm artillery. It boasts the advantage of being uncomplicated and computationally fast. To start with, you can define a plane mathematically by four different methods, but you most commonly represented it as 'a point and a normalized vector'. The normalized vector is perpendicular to the plane, while the known point can be anywhere on the planes surface. As you'll see, the Cartesian form of the plane equation is formally defined as: $Ax + By + Cz + d = 0$, where $\langle A, B, C \rangle$ is the vector normal to the plane, $\langle x, y, z \rangle$ is a point on the plane, and d is the shortest distance from the plane to the origin. The plane equation is used for an assortment of crucial techniques and forms the backbone of a number of fundamental algorithms.

Plane Equation & Dot Product The plane equation can be calculated using the dot product. To define a plane, you need two pieces of information. First, you need a point on the plane, anywhere on the plane; it doesn't matter as long as the point is on the plane. Second, you need the normal of the plane (i.e. the direction the plane is facing).

$$d = \hat{n} \cdot \vec{p} \quad (3.1)$$

where \hat{n} is the plane normal in Cartesian coordinates (unit-length), while the \vec{p} represents the coordinates of a point on the plane, and d

represents the shortest distance from the plane to the origin. Note, the point \vec{p} can be any point on the surface of the plane.

3.2.8 Support Function

Many algorithms make use of a mathematical tool called the **support function**, a.k.a **support mapping**. A support function takes a **direction** and an **array of vertices** as input and returns a point as output. The **output point is the furthest point along the given direction** given all the vertices. Note, there can be multiple points that are valid support function outputs for a particular array of vertices. For instance, the support function of an AABB, given the positive x-axis direction, can return any point on the AABB's face in the positive x-axis direction.

3.3 Matrix

3.3.1 Why Matrices?

Matrices are a compact way of representing and combining transformations (e.g., rotations and translation). Matrices are so common that most computer hardware (e.g., graphical processing units (GPUs) and CPUs) are optimized to perform very efficient matrix operations (i.e., with special instructions and by means of parallelization).

3.3.2 Column or Row Major

A matrix can be ordered using either Column or Row ordering (i.e., depending upon your preference). While DirectX uses Row Major ordering to store the matrix in memory, OpenGL uses Column Major ordering. For this book, you primarily use Column Major ordering.

$$\begin{array}{cc}
 \text{row matrix} & \text{column matrix} \\
 M = \begin{bmatrix} 1 & 2 & 3 & 4 \\ 5 & 6 & 7 & 8 \\ 9 & 10 & 11 & 12 \\ 13 & 14 & 15 & 16 \end{bmatrix} & M = \begin{bmatrix} 1 & 5 & 9 & 13 \\ 2 & 6 & 10 & 14 \\ 3 & 7 & 11 & 15 \\ 4 & 8 & 12 & 16 \end{bmatrix}
 \end{array}$$

Figure 3.3: **Column or Row Major** - Visually illustrating the difference between a column and row matrix organisation.

```

1  /*
2  * Column-major 4x4 matrix
3  *
4  * Layout:
5  *   0 4 8 12
6  *   1 5 9 13
7  *   2 6 10 14
8  *   3 7 11 15
9  */

```

```
10 * 3x3 Rotation Matrix Indices
11 * 0 4 8
12 * 1 5 9
13 * 2 6 10
14 *
15 * 3x1 Translation Indices
16 * 12
17 * 13
18 * 14
19 *
20 */
```

$$M = \begin{bmatrix} 1 & 5 & 9 & 13 \\ 2 & 6 & 10 & 14 \\ 3 & 7 & 11 & 15 \\ 4 & 8 & 12 & 16 \end{bmatrix}$$

`float m[16] = [1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16]`

Figure 3.4: **Matrix Elements** - The OpenGL specification for a 4x4 matrix uses column-major ordering. This means that the values of the matrix are stored by filling the columns with values, and only moves onto the next column once the current column is completely filled. This column-major order is adhered to throughout all the matrix operations within OpenGL. Remember, every fourth consecutive elements in an array represents a column in a 4x4 matrix.

3.3.3 A 4x4 Matrix

A 4x4 matrix (aka a homogeneous transformation matrix) can contain multiple different transformations (e.g., scaling, rotation, and translation), as shown in Figure 3.5. Rather than working with multiple different types of matrix, you will only work with a 4x4 matrix. Note, just in-case you didn't catch-on, a vector3 is technically a 1x3 matrix.




Figure 3.5: **4x4 Homogeneous Transformation Matrix** - Illustrating the different parts of a 4x4 matrix that represent the different transformations.

Figure 3.5 illustrates the decomposition of a 4x4 matrix into a 3x3 rotation matrix and a 1x3 translation matrix. Also, it shows how the diagonal components of the 3x3 rotation affect scaling along the x, y, and z axis, while the global scaling and perspective values are typically fixed.

3.3.4 Creating a Matrix

A matrix is just an array of variables. For example, an uncomplicated 4x4 matrix is shown below in Listing 3.5 is merely an array of 16 floats.

Listing 3.5: Uncomplicated Matrix.

```
1 class Matrix4
2 {
3 public:
4     float M[16];
5 };
```

For C++ and C#, you take advantage of accessor functions and operator overloading to make using variables easier and safer (i.e., sanity checks within the accessors), as shown in Listing 3.6.

Listing 3.6: Basic Matrix4 class for C++.

```
1 class Matrix4
2 {
3 public:
4     // Row - Column Format
5     // e.g., mat.Get(2,4) is row=2, and column=4
6     // or mat(2,4) - using operator overloading
7     float M[16];
8
9     // Accessor with sanity checks (i.e., boundary and
10    // valid number asserts)
11
12    float Get(int row, int col) const
13    {
14        DBG_ASSERT(row>=0 && row<4);
15        DBG_ASSERT(col>=0 && col<4);
16        return M[row*4+col];
17    } // End Get(..)
18
19    // Note - you can't overload operator[] to
20    // accept multiple arguments. Instead -
21    // instead you can overload operator() if you want to access
22    // values using (x,y) syntax
23    float& operator() (int row, int col)
24    {
25        DBG_ASSERT(row>=0 && row<4);
26        DBG_ASSERT(col>=0 && col<4);
27        return M[row*4+col];
28    }
29 };
30
31 // example:
32 Matrix4 mat;
33 mat(0,3) = 2;
34 // and
35 float val = mat.Get(0,3);
```

Note, a matrix can be stored in row-column or column-row form. Make sure you know which is which and be consistent.

3.3.4 Identity Matrix

The identity matrix is analogous to the number 1. If you multiply any matrix by an identity matrix, you will get the original matrix. The format for an identity matrix is all zeros except for the diagonal components, as shown in Equation 3.2.

$$M_{identity} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (3.2)$$

The implementation for creating an identity matrix is shown in Listing 3.7.

Listing 3.7: Creating a 4x4 Identity Matrix.

```
1 // Returns an instance of an identity matrix
2 static
3 Matrix4 Identity()
4 {
5     Matrix4 m;
6     m(0,0)=1; m(0,1)=0; m(0,2)=0; m(0,3)=0;
7     m(1,0)=0; m(1,1)=1; m(1,2)=0; m(1,3)=0;
8     m(2,0)=0; m(2,1)=0; m(2,2)=1; m(2,3)=0;
9     m(3,0)=0; m(3,1)=0; m(3,2)=0; m(3,3)=1;
10    return m;
11 }
```

3.3.4 Translation Matrix

The translation matrix represents a 3D world positions (i.e., an x , y , and z Cartesian point in space).

Essentially, if you start with an identity matrix, which does nothing when multiplied with another matrix. Then the bottom three values describe the translational information, as shown in Equation 3.3.

$$M_{translation} = \begin{bmatrix} 1 & 0 & 0 & tx \\ 0 & 1 & 0 & ty \\ 0 & 0 & 1 & tz \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (3.3)$$

Listing 3.8: Creating 4x4 Translation Matrix.

```
1 static
2 Matrix4 CreateTranslation(float x, float y, float z)
3 {
4     Matrix4 m = Matrix4.Identity;
```

```

5  m(0,3)=x;
6  m(1,3)=y;
7  m(2,3)=z;
8  return m;
9  }

```

3.3.4 Scale Matrix

You'll want to make things smaller and bigger! You can scale objects with a scaling matrix. You can scale the x, y, and z axis by modifying the diagonal elements of the matrix, as shown in Equation 3.4.

$$M_{scale} = \begin{bmatrix} sx & 0 & 0 & 0 \\ 0 & sy & 0 & 0 \\ 0 & 0 & sz & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (3.4)$$

Listing 3.9: Creating 4x4 Scale Matrix.

```

1  Matrix4 CreateScale(float x, float y, float z)
2  {
3      Matrix4 m = Matrix4.Identity;
4      m(0,0)=x;
5      m(1,1)=y;
6      m(2,2)=z;
7      return m;
8  }

```

3.3.4 Rotation Matrix

You need to be able to rotate your objects. You formulate the three main axis rotation matrices, as shown in Equation 3.5.

$$\begin{aligned}
 M_{XRotation} &= \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(-XAngle) & -\sin(-XAngle) & 0 \\ 0 & \sin(-XAngle) & \cos(-XAngle) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \\
 M_{YRotation} &= \begin{bmatrix} \cos(-YAngle) & 0 & \sin(-YAngle) & 0 \\ 0 & 1 & 0 & 0 \\ -\sin(-YAngle) & 0 & \cos(-YAngle) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \\
 M_{ZRotation} &= \begin{bmatrix} \cos(-ZAngle) & -\sin(-ZAngle) & 0 & 0 \\ \sin(-ZAngle) & \cos(-ZAngle) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}
 \end{aligned} \quad (3.5)$$

The order that you multiply the matrices determines the order the rotations will be applied to the point. For example:

$P \times (X \times Y \times Z)$ Rotates in X, Y, then Z

$P \times (Y \times X \times Z)$ Rotates in Y, X, then Z

$P \times (Z \times X \times Y)$ Rotates in Z, X, then Y

where P is the point, and X, Y, and Z represent the matrix-axis rotation.

Listing 3.10: Rotation Matrix Implementation.

```

1  static
2  Matrix4 CreateRotationX(float ax)
3  {
4      Matrix4 m = Matrix4.Identity();
5      m(1,1) = (float)Math.Cos(-ax); m(1,2) = -(float)Math.Sin(-ax);
6      m(2,1) = (float)Math.Sin(-ax); m(2,2) = (float)Math.Cos(-ax);
7      return m;
8  }
9
10 Matrix4 CreateRotationY(float ay)
11 {
12     Matrix4 m = Matrix4.Identity();
13     m(0,0) = (float)Math.Cos(-ay); m(0,2) = (float)Math.Sin(-ay);
14     m(2,0) = -(float)Math.Sin(-ay); m(2,2) = (float)Math.Cos(-ay);
15     return m;
16 }
17
18 Matrix4 CreateRotationZ(float az)
19 {
20     Matrix4 m = Matrix4.Identity();
21     m(0,0) = (float)Math.Cos(-az); m(0,1) = -(float)Math.Sin(-az);
22     m(1,0) = (float)Math.Sin(-az); m(1,1) = (float)Math.Cos(-az);
23     return m;
24 }
```

3.3.5 Matrix-Matrix Multiplication

You can construct matrices that represent different transformations (e.g., scaling, translation, and rotation), which you combine through multiplication.

Always remember matrix multiplication is **NOT commutative**. For example, if you want to rotate the object first then translate its position, you have to be sure you do the multiplication in the correct order; otherwise, you'll end up, translating the object then rotating it.

Listing 3.11: Matrix Multiplication Implementation (result = A * B)

```

1
2  Matrix4 Multiply(const Matrix4& ma, const Matrix4& mb)
3  {
4      Matrix4 result;;
```



```
5  for ( int i = 0; i < 4; ++i )
6      for ( int j = 0; j < 4; ++j )
7          result(i,j) =  ma.Get(i,0) * mb.Get(0,j)
8                      + ma.Get(i,1) * mb.Get(1,j)
9                      + ma.Get(i,2) * mb.Get(2,j)
10                     + ma.Get(i,3) * mb.Get(3,j);
11  return result;
12 }
```

3.3.6 'Pure' Rotation

Matrices that contain only rotation possess special features. For example, they can be easily inverted and converted to and from quaternion or axis-angle format.

3.3.6 Orthogonal Matrices (Useful-Axis)

A matrix that contains only rotational information is termed an 'orthogonal' matrix.

3.3.6 Transpose and Inverse

The inverse of an orthogonal (i.e., 'pure' rotation) matrix is its transpose (i.e., you swap the columns and rows). This is extremely valuable since it is computationally fast, since it requires no complex mathematical operations (e.g., sin and cos), and is straightforward and simple to implement in code, as shown in Listing 3.12.

Listing 3.12: Matrix Transpose Implementation.

```
1  Matrix4 Transpose(const Matrix4& m)
2  {
3      Matrix4 result;
4      for ( int i = 0; i < 4; ++i )
5          for ( int j = 0; j < 4; ++j )
6              result(i,j) = m.Get(j,i);
7      return result;
8  }
```

3.3.7 Transforming a Vector

A vector is basically a matrix with a single row, or column, depending upon your configuration. You can multiply your 4x4 matrix by a 4x1

vector (you'll convert your 3x1 to a 4x1 vector with the last component set to zero). The operation is shown in Equation 3.6.

$$\begin{bmatrix} x & y & z & 1 \end{bmatrix} * \begin{bmatrix} m_{11} & m_{12} & m_{13} & m_{14} \\ m_{21} & m_{22} & m_{23} & m_{24} \\ m_{31} & m_{32} & m_{33} & m_{34} \\ m_{41} & m_{42} & m_{43} & m_{44} \end{bmatrix} = \begin{bmatrix} a & b & c & d \end{bmatrix} \quad (3.6)$$

Listing 3.13: Matrix-Vector Transform Implementation.

```

1 // Row-Vector Convention
2 Vector3 Transform (const Vector3& v, const Matrix4& m)
3 {
4     float result[4];
5     for ( int i = 0; i < 4; ++i )
6     {
7         result[i] = v.X*m(0,i) + v.Y*m(1,i) + v.Z*m(2,i) + m(3,i);
8     }
9     return Vector3( result[0]/result[3],
10                    result[1]/result[3],
11                    result[2]/result[3] );
12 }
```

3.3.7 Little Test

So does your implementation work? You'll do a simple example to demonstrate your matrix and vector are performing the correct calculation. Don't just walk away and 'assume' it works. You should always ask the question, have I tested and am able to 'prove' that the code works - even if it's just modifying a few lines for optimisation reasons - did the optimisation or modification break the original implementation?

Let's create a simple Vector3 (e.g., 0,1,0) pointing straight-up, then you'll create a simple rotation matrix (e.g., rotate $\frac{\pi}{2}$ (i.e., 90 degrees) around z-axis). If you typed the code correctly, you should end up with a Vector3 pointing to the right (e.g., -1,0,0). Listing 3.14 demonstrates a simple implementation example for transforming a vector in code.

Note!!! You "Always" work with radians!! Not degrees, potatoes, or bananas, but "radians". Furthermore, positive rotation is counterclockwise, not clockwise. That is why when you rotate the Vector3(0,1,0), around the z-axis by $\frac{\pi}{2}$, you get Vector3(-1,0,0).

Listing 3.14: Basic Matrix-Vector Transform Sanity Test.

```

1
```

```

2 // Start with <0,1,0>, rotate it, and get <-1,0,0> back
3 Vector3 vy = Vector3(0,1,0);
4 Matrix4 rotZ = Matrix4.CreateRotationZ( (float)Math.PI*0.5f );
5 Vector3 vr = Vector3.Transform( vy, rotZ );
6 // If all went well, vr equals <-1,0,0>; well approximately, e.g., <-0.9999, 0, 0>,
   due to numerical errors and floating point precision ;

```

3.3.8 Matrix Inversion

A matrix is just a rectangle array of numbers or symbols organised into rows and columns. For example:

$$[A] = \begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix} \quad (3.7)$$

$$[A] = \begin{bmatrix} 0.25 & 0.33 \\ 0.125 & 0.66 \end{bmatrix}$$

So given the popular equation $F = ma$, if you know the force and the acceleration, you can work out the mass from $m = \frac{F}{a}$. However, for matrices, the division operation does not exist, hence, you use the ‘inverse’: $m = a^{-1}F$.

Matrix Inverse Properties Given a square matrix $[A]$ (i.e., equal number of rows and columns), then you can say:

$$[A]^{-1}[A] = [A][A]^{-1} = [I] \quad (3.8)$$

where $[I]$ is the identity matrix (i.e., matrix equivalent of 1).

There are two methods for inverting a matrix:

- Analytical
- Numerical

Analytical Matrix Inversion For small matrix problems (e.g., 2x2 or 3x3), the solution can be computed by hand, for example:

$$[A] = \begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix}$$

$$[A]^{-1} = \begin{bmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{bmatrix} \quad (3.9)$$

$$[A]^{-1}[A] = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} = [I]$$

Numerical Matrix Inversion When an analytical solution does not exist, then a numerical solution can be sought. For example:

$$\begin{aligned} [A] &= \begin{bmatrix} 0.25 & 0.33 \\ 0.125 & 0.666 \end{bmatrix} \\ [A]^{-1} &= \begin{bmatrix} 5.31734 & -2.6347 \\ -0.998 & 1.996 \end{bmatrix} \\ [A]^{-1}[A] &= \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} = [I] \end{aligned} \quad (3.10)$$

Techniques for numerically inverting a matrix, include:

- Gaussian Elimination (LU factorization, Gauss-Seidel)
- Singular Value Decomposition (SVD)
- Cholesky Factorization (symmetric defined matrices)

When considering a numerical routine, computational cost and robustness are important factors - for example, you may want the algorithm to converge on a best guess solution for singular matrix problems (i.e, non-convertible matrix - analogous to a divide by zero issue).

Singular Systems If a matrix is 'not' invertible it is said to be singular (it exists on its own). When a matrix is singular, the determinant of a matrix is equal to zero.

Singular systems arise when:

- the equations representing the rows in a matrix are closely inter-related
- data in the matrix contains significant errors which makes it seem as if the rows in the matrix are closely inter-related

Determinant The determinant of a matrix is a single scalar value. Every square matrix has a determinant. For example, to calculate the determinant for a 2×2 matrix:

$$\begin{aligned} \det[A] &= \begin{vmatrix} a & b \\ c & d \end{vmatrix} \\ &= ad - bc \end{aligned} \quad (3.11)$$

When the determinant of a matrix is zero, it is not invertible.

3.4 Quaternion

Quaternions are an efficient, straightforward and robust way of representing rotations. You can represent a rotation using a 3×3 matrix, however, a quaternion only uses 4 variables instead of the 9 variables for a 3×3 matrix. Allows you to easily be interpolated, combine, and re-normalized orientations during drifting (numerical errors).

3.4.1 Why Quaternions?

If quaternions are compared with other types of methods for representing rotation (e.g., Euler's angles, matrices, axis-angle) the quaternion comes out on top. In summary:

- They don't suffer from gimbal lock
- They use the minimum number of variables (i.e., 4-floats) to *uniquely* represent a rotation with no ambiguity
- They are easy to combine (i.e. through multiplication the same as with matrices)
- They can be inverted easily (i.e., unit-quaternion's inverse is its conjugate, which is simply the negative of the vector components)

Hence, you can calculate angular difference between pairs of unit-quaternions easily and fastly

- Interpolating is a breeze
- Drifting due to numerical errors is easier to correct (i.e., re-normalizing the unit-quaternion) compared to matrices

3.4.2 Unit-Quaternion (Always)

In the majority of cases your quaternions will always be unit-quaternions. If they aren't then something has gone wrong. Hence, assert and check that the length of your quaternions is always (approximately) equal to one.

3.4.3 Creating a Quaternion

Essentially, a quaternion is just a 4 vector class, and its implementation is very simple, as shown in Listing 3.15. However, it's all the helper methods that make the quaternion tool invaluable (e.g., multiplication and interpolation methods) that you go into next.

Listing 3.15: Implementation of a Quaternion class.

```
1 class
2 Quaternion
```

```
3 {  
4 public:  
5     float w, x, y, z;  
6 };
```

3.4.3 Quaternion from Axis-Angle

Listing 3.16: Quaternion From Axis-Angle Implementation.

```
1  
2 Quaternion QuaternionFromAxisAngle(const Vector3& axis, float angle)  
3 {  
4     Quaternion q;  
5     q.X = axis.X * (float)Math.Sin(angle/2);  
6     q.Y = axis.Y * (float)Math.Sin(angle/2);  
7     q.Z = axis.Z * (float)Math.Sin(angle/2);  
8     q.W = (float)Math.Cos(angle/2);  
9     return q;  
10 }
```

3.4.3 Quaternion to Axis-Angle

Listing 3.17: Quaternion To Axis-Angle Implementation.

```
1  
2 void QuaternionToAxisAngle(const Quaternion& q,  
3     Vector3& outAxis,  
4     float& outAngle)  
5 {  
6     outAngle = 2 * (float)Math.Acos(q.w);  
7     float s = (float)Math.Sqrt(1-q.w*q.w); // assuming quaternion normalised then w is  
8         // less than 1, so term always positive.  
9     if (s < 0.001)  
10     { // test to avoid divide by zero, s is always positive due to sqrt  
11         // if s close to zero then direction of axis not important  
12         outAxis.x = q.x; // if it is important that axis is normalised then replace with  
13             x=1; y=z=0;  
14         axis.y = q.y;  
15         axis.z = q.z;  
16         return;  
17     }  
18     outAxis.X = q.x / s; // normalize axis  
19     outAxis.Y = q.y / s;  
20     outAxis.Z = q.z / s;  
21 }
```

3.4.3 Quaternion to Matrix

The top left 3x3 part of the rotation matrix is formed with Equation 3.12.

$$\begin{bmatrix} 1 - 2q_y^2 - 2q_z^2 & 2q_xq_y - 2q_zq_w & 2q_xq_z + 2q_yq_w & 0 \\ 2q_xq_y + 2q_zq_w & 1 - 2q_x^2 - 2q_z^2 & 2q_yq_z - 2q_xq_w & 0 \\ 2q_xq_z - 2q_yq_w & 2q_yq_z + 2q_xq_w & 1 - 2q_x^2 - 2q_y^2 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (3.12)$$

Listing 3.18: Quaternion to Matrix Implementation.

```

1
2 Matrix4 QuaternionToMatrix(const Quaternion& q)
3 {
4     float sqw = q.W*q.W;
5     float sqx = q.X*q.X;
6     float sqy = q.Y*q.Y;
7     float sqz = q.Z*q.Z;
8
9     Matrix4 m = Matrix4.Identity();
10
11     // invs (inverse square length) is only required if quaternion is not already
12     // normalised
13     float invs = 1 / (sqx + sqy + sqz + sqw);
14     m(0,0) = ( sqx - sqy - sqz + sqw)*invs; // since sqw + sqx + sqy + sqz = 1/invs*invs
15     m(1,1) = (-sqx + sqy - sqz + sqw)*invs;
16     m(2,2) = (-sqx - sqy + sqz + sqw)*invs;
17
18     float tmp1 = q.X*q.Y;
19     float tmp2 = q.Z*q.W;
20     m(1,0) = 2.0f * (tmp1 + tmp2)*invs;
21     m(0,1) = 2.0f * (tmp1 - tmp2)*invs;
22
23     tmp1 = q.X*q.Z;
24     tmp2 = q.Y*q.W;
25     m(2,0) = 2.0f * (tmp1 - tmp2)*invs;
26     m(0,2) = 2.0f * (tmp1 + tmp2)*invs;
27
28     tmp1 = q.Y*q.Z;
29     tmp2 = q.X*q.W;
30     m(2,1) = 2.0f * (tmp1 + tmp2)*invs;
31     m(1,2) = 2.0f * (tmp1 - tmp2)*invs;
32     return m;
33 }
```

3.4.3 Quaternion from Matrix

As show in Equation 3.12, with the rule that your quaternion is a unit-quaternion (i.e., $q = (q_w, q_x, q_y, q_z)$ where $|q| = 1$)

You need to know how a rotation matrix (i.e., a ‘pure’ 3x3 rotation matrix without scaling) can be compared with the result of a quaternion

(e.g., see Figure 3.5 for the components of a matrix.).

$$\begin{bmatrix} q_w^2 + q_x^2 - q_y^2 - q_z^2 & 2(q_x q_y - q_z q_w) & 2(q_x q_z + q_y q_w) \\ 2(q_x q_y + q_z q_w) & q_w^2 - q_x^2 + q_y^2 - q_z^2 & 2(q_y q_z - q_x q_w) \\ 2(q_x q_z - q_y q_w) & 2(q_y q_z + q_x q_w) & q_w^2 - q_x^2 - q_y^2 + q_z^2 \end{bmatrix} = \begin{bmatrix} r_{11} & r_{12} & r_{13} \\ r_{21} & r_{22} & r_{23} \\ r_{31} & r_{32} & r_{33} \end{bmatrix} \quad (3.13)$$

by remembering that $q_w^2 + q_x^2 + q_y^2 + q_z^2 = 1$, you can rearrange and solve Equation 3.13 to calculate the 3x3 rotation matrix components.

Listing 3.19: Quaternion from Matrix.

```

1
2 static float SIGN(float x) {return (x >= 0.0f) ? +1.0f : -1.0f;}
3 static float NORM(float a, float b, float c, float d) {return sqrt(a * a + b * b + c
   * c + d * d);}
4
5 static
6 Quaternion QuaternionFromMatrix(const Matrix4& m)
7 {
8     /*
9         | 00, 01, 02 |
10        m = | 10, 11, 12 |
11             | 20, 21, 22 |
12
13         q = | qx, qy, qz, qw |
14     */
15     float qx = ( m(0,0) + m(1,1) + m(2,2) + 1.0f) / 4.0f;
16     float qy = ( m(0,0) - m(1,1) - m(2,2) + 1.0f) / 4.0f;
17     float qz = (-m(0,0) + m(1,1) - m(2,2) + 1.0f) / 4.0f;
18     float qw = (-m(0,0) - m(1,1) + m(2,2) + 1.0f) / 4.0f;
19     if (qx < 0.0f) qx = 0.0f;
20     if (qy < 0.0f) qy = 0.0f;
21     if (qz < 0.0f) qz = 0.0f;
22     if (qw < 0.0f) qw = 0.0f;
23     qx = sqrt(qx);
24     qy = sqrt(qy);
25     qz = sqrt(qz);
26     qw = sqrt(qw);
27     if (qx >= qy && qx >= qz && qx >= qw)
28     {
29         qx *= +1.0f;
30         q1 = SIGN(m(2,1) - m(1,2));
31         q2 = SIGN(m(0,2) - m(2,0));
32         q3 = SIGN(m(1,0) - m(0,1));
33     }
34     else if (qy >= qx && qy >= qz && qy >= qw)
35     {
36         qx = SIGN(m(2,1) - m(1,2));
37         qy *= 1.0f;
38         qz = SIGN(m(1,0) + m(0,1));
39         qw = SIGN(m(0,2) + m(2,0));
40     }
41     else if (qz >= qx && qz >= qy && qz >= qw)

```



```

42 {
43     qx *= SIGN(m(0,2) - m(2,0));
44     qy *= SIGN(m(1,0) + m(0,1));
45     qz *= 1.0f;
46     qw *= SIGN(m(2,1) + m(1,2));
47 }
48 else if (qw >= qx && qw >= qy && qw >= qz)
49 {
50     qx *= SIGN(m(1,0) - m(0,1));
51     qy *= SIGN(m(2,0) + m(0,2));
52     qz *= SIGN(m(2,1) + m(1,2));
53     qw *= 1.0f;
54 }
55 else
56 {
57     Debug_c.Assert("**error**\n");
58 }
59 r = NORM(qx, qy, qz, qw);
60 qx /= r;
61 qy /= r;
62 qz /= r;
63 qw /= r;
64 return Quaternion(qx,qy,qz,qw);
65 }

```

3.4.4 Quaternion-Quaternion Multiplication

You multiply quaternions together to concatenate the rotational transforms (i.e., analogous to how you multiply matrices together to combine the individual transforms into a single unified solution). The quaternion multiplication mathematics is easier to digest, if you subdivide the quaternion elements into a ‘scalar’ s and ‘vector’ v component and use the dot and cross product:

$$\begin{aligned}
 (sa, \vec{va})(sb, \vec{vb}) &= (sa)(sb) + (sa)(\vec{vb}) + (sb)(\vec{va}) + ((\vec{va}) \times (\vec{vb})) - ((\vec{va}) \cdot (\vec{vb})) \\
 &\text{group into parts} \\
 &= ((sa)(sb) - ((\vec{va}) \cdot (\vec{vb}))), \text{ scalar part} \\
 &\quad ((sa)(\vec{vb}) + (sb)(\vec{va}) + ((\vec{va}) \times (\vec{vb}))) \text{ vector part} \\
 &\hspace{15em} (3.14)
 \end{aligned}$$

Listing 3.20: Quaternion Quaternion Multiplication.

```

1
2 Quaternion Multiplication(const Quaternion& qa,
3     const Quaternion& qb)
4 {
5     Quaternion qr = Quaternion.Identity;
6     Vector3 va = Vector3(qa.x, qa.y, qa.z);
7     Vector3 vb = Vector3(qb.x, qb.y, qb.z);
8     qr.w = qa.w*qb.w - Vector3.Dot(va,vb);
9     Vector3 vr = Vector3.Cross(va,vb) + qa.w*vb + qb.w*va;
10    qr.x = vr.x;

```

```

11   qr.y = vr.y;
12   qr.z = vr.z;
13   return qr;
14 }

```

3.4.5 Quaternion Inverse (Conjugate)

For a unit-quaternion the conjugate is the same as the inverse. You represent the conjugate by the $*$ symbol, e.g., $\text{conjugate}(q) = q^*$.

The conjugate is useful because it has the following properties:

- $q_a^* q_b^* = (q_b q_a)^*$ In this way you can change the order of the multiplicands.
- $qq^* = a^2 + b^2 + c^2 + d^2 = \text{real number}$. Multiplying a quaternion by its conjugate gives a real number. This makes the conjugate useful for finding the multiplicative inverse. For instance, if you are using a quaternion q to represent a rotation then $\text{conj}(q)$ represents the same rotation in the reverse direction.
- $P_{out} = q P_{in} q^*$ you use this to calculate a rotation transform.

Listing 3.21: Quaternion Conjugate.

```

1  Quaternion Conjugate(const Quaternion& q)
2  {
3      // Note, you invert the vector component
4      Quaternion qr (q.w, -q.x, -q.y, -q.z);
5      return qr;
6  }
7  }

```

3.4.6 Transform a Vector by a Quaternion

As pointed out, you can use the Conjugate to make transforming a `Vector3` a piece of cake. You convert the `Vector3` to a quaternion (i.e., set the scalar W component to 0), then multiply them to get the result. You extract the transformed `Vector3` (i.e., the x , y , and z component of the resulting multiplied quaternions). Simple eh?. The formulation is given by:

$$v_{out} = q v_{in} q^* \quad (3.15)$$

where v_{in} is the original point converted to a quaternion (i.e., w component is set to zero), q and q^* are the quaternion and quaternion conjugate, and v_{out} is the transformed point (i.e., x , y , and z component of the resulting quaternion).

Listing 3.22: Quaternion Vector Transform.

```
1
2 Vector3 Transform(const Vector3& v, const Quaternion& q)
3 {
4     Quaternion qv(0, v.x, v.y, v.z);
5     Quaternion qr = q * qv * Conjugate(q);
6     return new Vector3( qr.x, qr.y, qr.z );
7 }
```

3.5 Summary

Most of the time, you'll use pre-written math libraries (such as, `vmath` or `glm`). If you do write a set of math libraries, you'll probably write them once and never need to worry about writing them again. However, having a solid understanding of how the vector mathematics works can dramatically help you understand the creation, optimization, and debugging of algorithms (both from a theoretical and practical perspective).

Type	Representation
Vector \vec{v}	x, y, z
Unit Vector \hat{v}	$\frac{\vec{v}}{ \vec{v} } \implies \hat{v} = 1$
Position \vec{p}	x, y, z
Rotation \vec{q}	x, y, z, w (quaternion)
Sphere	\vec{p}, r
Plane	\vec{p}, \hat{n}
AABB	\vec{p}, \vec{e}
OBB	$\vec{p}, \vec{q}, \vec{e}$
Line/Segment	\vec{p}_0, \vec{p}_1
Ray	\vec{p}, \hat{n}
Triangle (t)	$\vec{p}_0, \vec{p}_1, \vec{p}_2$
Mesh	$\sum t$
Capsule	\vec{p}_0, \vec{p}_1, r

Table 3.1: Defining primitive objects using a mathematical representation (e.g., a sphere is represented by a centre position p and a radius r).

You use the symbols in Table 3.1, such as, arrows and hats above vectors, to enable us to read mathematical equations at a glance. For instance, you can easily identify a scalar a and a \vec{a} quickly; or a vector \vec{b} and a unit-vector \hat{b} . You also provide simple implementation listings to solidify the your understanding.

3.6 Exercises

After you're familiar with the core mathematical principles, you'll need to constantly practice to strengthen your understanding. The

following example questions provide you this opportunity.

3.6.1 Chapter Questions

Question Given the three matrices A: translation along the vector $v = (4, 0, 2)$, B: rotation 90 degrees around the z-axis and C: a non-uniform scaling with 2 in x, 3 in y and 4 in z.

- Give the (4×4) matrix form of each of A, B and C.
- Calculate the transformed point P' , given the point $P = (1, 2, 3, 1)$. i.e., $P' = CABP$

Question What does it mean if two vectors are orthogonal? How can you determine if two vectors are orthogonal?

Question Give a 3×3 homogeneous matrix to rotate an image clockwise by 90 degrees. Then shift the image to the right by 10 units. Finally scale the image by twice as large. All these transformations are to be done one after the other in sequence

Question What are the basic 2D geometric transformations? Explain each with its matrix representation

Question Show that the composition of two rotations is additive by concatenating the matrix representations for $R(\theta_1)$ and $R(\theta_2)$ to obtain $R(\theta)$. $R(\theta) = R(\theta_1 + \theta_2)$

Question Derive the transformation matrix for rotation about any axis

Question Given a triangle A(0,0), B(1,1) and C(6,2). Write down the transformation matrix to magnify the triangle to twice its size keeping C(6,2) fixed.

Question Explain basic 2D transformations? Give the homogeneous matrix representations for each transformation.

4

Graphical Principles

As you might be new to graphical programming, you might find a number of concepts confusing and alien when discussed in the context of the Vulkan API, such as, shaders and projection transforms. While it would be beyond the scope of a single Chapter to teach a complete graphical syllabus, instead this Chapter aims to review a number of core graphical principles that are fundamental to most graphical solutions. In addition, you're encouraged to read around the subject to complement your understanding of the material (e.g., computer graphics books, introductory graphics/maths articles and online tutorials). In this Chapter, you'll quickly review the following concepts:

- Basic Types
 - Scalars, Vectors, Floats, Colors, ..
- Transforms
 - Coordinate Spaces
 - Camera and Projection
- Primitives
 - Lines, Triangles
- Data/Geometry
 - Vertices & Indices
- Drawing Principles
 - Draw Ordering (Counter)-Clockwise, Texturing, Depth Buffer, Clipping
 - Render output and clipping-cube
- Programmable Graphics
 - Shaders, Pipeline, ..

4.1 Basic Types

As with any standard programming language, you'll have a set of standard data types, such as, floats, doubles and strings. You'll also




Figure 4.1: How you go from 3-dimensional geometry using transforms and rasterization techniques to produce a graphical output on your screen (pixels).

need to create a number of structures to encapsulate data for ease of use and readability. For example, arrays of data for representing your geometry, matrix transforms and color information. You need to be aware of overheads, such as, the sizes of variables in memory, alignment specifics (structure padding) and conversion costs (doubles to floats). For example, see Figure 4.2 for a short list of common Vulkan types and Listing 4.1 for a simple power of two test function.

Vulkan Data Types

Vulkan Data Type	Representation	C-Type
VkFlags	32-bit	uint32_t
VkBool32	32-bit	uint32_t
VkDeviceSize	64-bit	uint64_t
VkSampleMask	32-bit	uint32_t
.....

Figure 4.2: To help with cross-platform development Vulkan defines various data types that map standard C/C++ data types.

Translation $\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & d_x \\ 0 & 1 & d_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$

Scaling $\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} s_x & 0 & 0 \\ 0 & s_y & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$

Rotation $\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$

Figure 4.3: Matrix representation of basic transformations (2-dimensions) - see Chapter 3 for a review of vector/matrix mathematical concepts.

Listing 4.1: Simple example of understanding binary data - testing if a value is a power of two (e.g., 2, 4, 8, 16, 32, 64, 128 - such as for memory allocations and texture dimensions/widths/heights).

```
1 bool IsAlignedPowerOfTwo(uint32_t alignment)
2 {
3     // 2- 10
4     // 4- 100
5     // 8- 1000
6     // ...
7     // Returns true if a power of 2
8     return ((alignment & (alignment-1)) == 0);
9
10    /*
11    e.g.
12    val      = 1000 (8) - power of 2
13    val-1    = 0111
14    val & (val-1) = 0000 (==0) return true
15
16    val      = 0110 (6) - not a power of 2
17    val-1    = 0101
18    val & (val-1) = 0100 (!=0) return false
19
20    */
21 }
```

4.2 Transforms

Mathematics has many applications in computer graphics especially matrices as discussed in Chapter 3. Matrices represent groups of equa-

tions that provide a compact, efficient and systematic way of doing the mathematical operations, such as, rotation, translation, scaling and projection (i.e., the representation of any transformation affine or non-affine). Importantly, the hardware within the computer (like the GPU) is optimised for matrix arithmetic. Of course, one of the most powerful feature that matrices give you is the ability to concatenate several transformations into a single matrix.

Common vector and matrix graphical operations that you'll come across again and again (and should ideally be comfortable with), include:

- Matrix-Vector Transform
- Matrix-Matrix Multiplication
- Vector Cross/Dot Product
- Rotation Matrix (x, y and z axis)
- Scale Matrix
- Translation Matrix
- Projection Matrix
- View Matrix

For a refresher on basic vector and matrix operations see Chapter 3.




Figure 4.4: Simplified graphical overview of the transformation stages between spaces (local space, world space, camera space and projection space) using matrix transforms (model matrix, view matrix, world matrix and a projection matrix).

As shown in Figure 4.4, there are multiple coordinate systems involved in 3-dimensional graphics, such as, Object Space, World Space (aka Model Space), Camera Space (aka Eye Space or View Space), and Screen Space (aka Clip Space). The best thing is, the conversion between the different transform spaces is effortless. You switch between different spaces by multiplying by a transform matrix. For instance, switching from world space to camera space you'd use your 'view matrix'.

While it's important you know how matrix and vector operations work (especially for 3-dimensional graphics), you don't always have to write your own, and a number of free open source libraries are available. For example, one popular mathematics library is:

OpenGL Mathematics (GLM) [1] library for graphics software based on the OpenGL Shading Language (GLSL) specifications. GLM is a header only C++ mathematics library that provides classes and functions designed and implemented with the same naming conventions and functionalities than GLSL.

4.2.1 Homogeneous Coordinates (or Projective Coordinates)

Cartesian coordinate transforms, such as, translation and perspective projection, cannot be expressed through matrix multiplication alone and is one of the core reasons you need to use homogeneous coordinate. Your graphics card takes advantage of homogeneous coordinates to perform transforms efficiently using vector processors with 4-element registers (e.g., programmable shaders and pipeline operations) - making matrix operations highly desirable. Any transformation can be represented as a matrix with each matrix having four columns of four rows due to the homogeneous coordinate system (Figure 4.6). In order to position and align your objects and set up representations of your scene inside your computer, you'll need to be able to transform your objects. As pointed out earlier, there are many transformations available to you (like stretching, twisting and bending), but the three absolutely necessary transforms you need to know are rotation, translation and scaling (Figure 4.5).

<p>Identity</p> $\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$	<p>Translation</p> $\begin{bmatrix} 1 & 0 & 0 & t_x \\ 0 & 1 & 0 & t_y \\ 0 & 0 & 1 & t_z \\ 0 & 0 & 0 & 1 \end{bmatrix}$	<p>Scaling</p> $\begin{bmatrix} s_x & 0 & 0 & 0 \\ 0 & s_y & 0 & 0 \\ 0 & 0 & s_z & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$
<p>Rotation Around X</p> $\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(v) & -\sin(v) & 0 \\ 0 & \sin(v) & \cos(v) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$	<p>Rotation Around Y</p> $\begin{bmatrix} \cos(v) & 0 & \sin(v) & 0 \\ 0 & 1 & 0 & 0 \\ -\sin(v) & 0 & \cos(v) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$	<p>Rotation Around Z</p> $\begin{bmatrix} \cos(v) & -\sin(v) & 0 & 0 \\ \sin(v) & \cos(v) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$

Figure 4.5: Core 3-dimensional homogeneous matrix transforms. Matrices are able to represent a variety of geometric transformations - which are able to be combined with each other by matrix multiplication. As a result, any perspective projection of space can be represented as a single matrix.

Remember, when you transforms your points the result is always made homogeneous. This means that your coordinate values are divided with 'W' (see Figure 4.6).

Point transformed by a translation matrix

$$P' = M \cdot P = \begin{bmatrix} 1 & 0 & 0 & a \\ 0 & 1 & 0 & b \\ 0 & 0 & 1 & c \\ 0 & 0 & 0 & w \end{bmatrix} \cdot \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} = \begin{bmatrix} x+a \\ y+b \\ z+c \\ w \end{bmatrix}$$

and homogeneous

$$\begin{bmatrix} (x+a)/w \\ (y+b)/w \\ (z+c)/w \\ 1 \end{bmatrix}$$

Figure 4.6: Homogeneous coordinates are crucial in computer graphics and 3D computer vision as they allow affine transformations and, in general, projective transformations to be easily represented by a matrix.

4.2.2 Normalized Device Coordinates (NDC)

The Normalize Device Coordinates (NDC) come into action towards the end of the processing (i.e., during the transition from 4D Homogeneous coordinates to screen pixels):

1. Your 4x4 PROJECTION transform takes you from 4D eye coordinates to 4D clip coordinates
2. Then the perspective divide takes you from 4D clip coordinates to 3D NDC coordinates
3. Then the viewport transformation takes those 3D NDC coordinates into 3D window coordinates.

Remember, multiplying by your 'PROJECTION' matrix takes you to 4D CLIP space.

Then the perspective divide gets you to the 3D NDC space.

(EYE-SPACE) →


(PROJECTION TRANSFORM) →

(CLIP-SPACE) →

(PERSPECTIVE DIVIDE) →

NDC-SPACE

Perspective



Orthographic




Figure 4.7: 3-dimensional model to a 2-dimensional image.





Figure 4.8: Mapping the pyramid frustum (eye coordinates) to the cube Normalize Device Coordinates (NDC).

4.2.3 Eye Coordinates

When you transform your geometry by the model and view matrix - this takes you to 'eye coordinates'. In other words, Vulkan defines the camera to be always located at (0, 0, 0) and facing to -Z axis in the eye space coordinates. You transform your vertices (or geometry) from

object space to eye space using your ‘model-view’ matrix which you perform on the GPU in the shader (e.g., vertex shader). The ‘model-view’ matrix is a combination of the ‘Model’ and ‘View’ matrices.



(a)
$$\begin{bmatrix} x_{\text{obj}} \\ y_{\text{obj}} \\ z_{\text{obj}} \\ w_{\text{obj}} \end{bmatrix} \mathbf{M}_{\text{model}} = \begin{bmatrix} x_{\text{world}} \\ y_{\text{world}} \\ z_{\text{world}} \\ w_{\text{world}} \end{bmatrix}$$

(b)
$$\begin{bmatrix} x_{\text{world}} \\ y_{\text{world}} \\ z_{\text{world}} \\ w_{\text{world}} \end{bmatrix} \mathbf{M}_{\text{view}} = \begin{bmatrix} x_{\text{eye}} \\ y_{\text{eye}} \\ z_{\text{eye}} \\ w_{\text{eye}} \end{bmatrix}$$

(c)
$$\begin{bmatrix} x_{\text{eye}} \\ y_{\text{eye}} \\ z_{\text{eye}} \\ w_{\text{eye}} \end{bmatrix} \mathbf{M}_{\text{projection}} = \begin{bmatrix} x_{\text{clip}} \\ y_{\text{clip}} \\ z_{\text{clip}} \\ w_{\text{clip}} \end{bmatrix}$$

(d)
$$\begin{bmatrix} x_{\text{clip}}/w_{\text{clip}} \\ y_{\text{clip}}/w_{\text{clip}} \\ z_{\text{clip}}/w_{\text{clip}} \end{bmatrix} = \begin{bmatrix} x_{\text{ncd}} \\ y_{\text{ncd}} \\ z_{\text{ncd}} \end{bmatrix}$$

(e) Viewport (x, y, w, h)
Depth Range (n, f)
$$\begin{bmatrix} \frac{w}{2} x_{\text{ncd}} + (x + \frac{w}{2}) \\ \frac{h}{2} x_{\text{ncd}} + (y + \frac{h}{2}) \\ \frac{f-n}{2} x_{\text{ncd}} + (\frac{f+n}{2}) \end{bmatrix} = \begin{bmatrix} x_{\text{window}} \\ y_{\text{window}} \\ z_{\text{window}} \end{bmatrix}$$

Figure 4.9: Transforms: (a) World Coordinates, (b) Eye Coordinates, (c) Clip Coordinates, (d) Normalized Device Coordinates (NDC), (e) Window Coordinates (Screen Coordinates).

4.2.4 Projection

The two main categories of projection are (1) perspective and (2) parallel projection as shown in Figure 4.10. For parallel projections, you’ll typically use a basic Orthographic Projection, while for Perspective you’ll use something more fancy to capture the real-world perception of objects getting smaller as they get further away. Applications of the the two projection techniques:

- Parallel projection are used for on screen menus or technical drawings
- Perspective projections are used for full 3-dimensional scenes that mimic the real-world (depth and distance)




Figure 4.10: Broadly categorising sorting projection methods into two main categories.

4.2.4 Orthogonal

A simple orthographic transformation where the original world units would be preserved (the z-coordinate is simply thrown away) is shown below in Equation 4.1:

$$\begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} x' \\ y' \\ z' = 0 \\ 1 \end{bmatrix} \quad (4.1)$$

4.2.4 Perspective

For perspective transforms, this is closer to what you see in the real-world, where objects closer to viewer look larger and parallel lines appear to converge to single point when they go off into the distance (as with train tracks - Figure 4.11). The mathematics is a little more involved for calculating the projection matrix. However, the principles are governed by simple geometric concepts. As shown in Figure 4.12, the projection matrix works by ‘projecting’ the object onto a surface from a pin-point camera location. Due to the importance of the projection matrix in computer graphics the steps for calculating a simple projection matrix follow.

You’ll typically use a one point perspective - however, multi-point perspective projections are possible (e.g., two and three point). These different linear perspective method use a lines to create the illusion




Figure 4.11: Train Track - Linear perspective projection with one vanishing point.

of space on a flat surface. There are three types of linear perspective. One point perspective uses one vanishing point placed on the horizon line. Two point perspective uses two points placed on the horizon line. Three point perspective uses three vanishing points (Figure 4.10).




Figure 4.12: Perspective (3D world to 2D screen window). The horizontal and vertical calculations are done independently. The perspective projection calculation uses basic trigonometric principles (e.g., similar triangles) to derive the perspective matrix. For example, in the diagram, you know all of the values except x' for the projection of the point $P[x, y, z, 1]$ onto the view plane.

As shown in Figure 4.12, the perspective transformation to project the coordinates onto a simple plane is given by Equation 4.2:

$$\begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1/n & 0 \end{bmatrix} = \begin{bmatrix} x' \\ y' \\ z' \\ 1 \end{bmatrix} \quad (4.2)$$

where n is the near viewing plane distance (see Figure 4.12). In the perspective case, you use similar triangles to solve for the intersection point on the plane's surface.

$$\begin{aligned} \frac{x_c}{n} &= \frac{x_e}{z_e} \\ \frac{y_c}{n} &= \frac{y_e}{z_e} \end{aligned} \quad (4.3)$$

therefore:

$$\begin{aligned} x_c &= \frac{x_e}{z_e/n} \\ y_c &= \frac{y_e}{z_e/n} \end{aligned} \quad (4.4)$$

For a real-world projection matrix, you'd have to specify a number of parameters, and the projection surface may not be square (rectangular with an aspect ratio). You'll now go through the steps to creating a more usable final perspective matrix in detail. You'll start with an

empty matrix and add the specifics for each matrix element as you progress through each step.

Step 1 Pass through z_e onto w_c ($w_c = -z_e$):

$$\begin{bmatrix} x_e \\ y_e \\ z_e \\ w_e \end{bmatrix} \begin{bmatrix} ? & ? & ? & ? \\ ? & ? & ? & ? \\ ? & ? & ? & ? \\ 0 & 0 & -1 & 0 \end{bmatrix} \begin{bmatrix} x_c \\ y_c \\ z_c \\ w_c \end{bmatrix} \quad (4.5)$$

Step 2 Map the input coordinates to the NDC coordinates using the relationship $[l, r] \rightarrow [-1, 1]$, $[b, t] \rightarrow [-1, 1]$:

$$\begin{bmatrix} x_e \\ y_e \\ z_e \\ w_e \end{bmatrix} \begin{bmatrix} \frac{2n}{r-l} & 0 & \frac{r+l}{r-l} & 0 \\ 0 & \frac{2n}{t-b} & \frac{t+b}{t-b} & 0 \\ ? & ? & ? & ? \\ 0 & 0 & -1 & 0 \end{bmatrix} \begin{bmatrix} x_c \\ y_c \\ z_c \\ w_c \end{bmatrix} \quad (4.6)$$

Step 3 z_c needs to be modified to include depth information for clipping (e.g., depth test) and is 'not' just the near (n) value. Hence, you need to work out how z_e maps to the near-far. Importantly, the z calculation does not depend on the x or y coordinates. Updating the matrix to include to extra variables 'A' and 'B' and solve them:

$$\begin{bmatrix} x_e \\ y_e \\ z_e \\ w_e \end{bmatrix} \begin{bmatrix} \frac{2n}{r-l} & 0 & \frac{r+l}{r-l} & 0 \\ 0 & \frac{2n}{t-b} & \frac{t+b}{t-b} & 0 \\ 0 & 0 & A & B \\ 0 & 0 & -1 & 0 \end{bmatrix} \begin{bmatrix} x_c \\ y_c \\ z_c \\ w_c \end{bmatrix} \quad (4.7)$$

$$z_n = \frac{z_c}{w_c} = \frac{Az_e + Bw_e}{-z_e} \quad (4.8)$$

You have one equation and two unknowns, so it's impossible to solve unless you add some additional information. Hence, to accomplish this by specifying the value for z_n when the point is on the near (n) and far (f) planes.

$$\begin{aligned} z_n = -1 & \quad \text{when} \quad z_e = -n \\ z_n = 1 & \quad \text{when} \quad z_e = f \end{aligned} \quad (4.9)$$

$$\begin{aligned} \frac{-An + B}{n} = -1 & \quad - > \quad -An + B = -n \\ \frac{-Af + B}{f} = 1 & \quad - > \quad -Af + B = f \end{aligned} \quad (4.10)$$

Hence, you have two equations and two unknowns and should be able to solve for A and B :

$$\begin{aligned} A &= -\frac{f+n}{f-n} \\ B &= -\frac{2fn}{f-n} \end{aligned} \quad (4.11)$$

$$\begin{bmatrix} x_e \\ y_e \\ z_e \\ w_e \end{bmatrix} \begin{bmatrix} \frac{2n}{r-l} & 0 & \frac{r+l}{r-l} & 0 \\ 0 & \frac{2n}{t-b} & \frac{t+b}{t-b} & 0 \\ 0 & 0 & -\frac{f+n}{f-n} & -\frac{2fn}{f-n} \\ 0 & 0 & -1 & 0 \end{bmatrix} \begin{bmatrix} x_c \\ y_c \\ z_c \\ w_c \end{bmatrix} \quad (4.12)$$

Step 4 Simplify the perspective matrix for a general frustum. When the viewing volume is symmetric: $r = -l$ and $t = -b$ it simplifies to:

$$\begin{bmatrix} x_e \\ y_e \\ z_e \\ w_e \end{bmatrix} \begin{bmatrix} \frac{n}{r} & 0 & 0 & 0 \\ 0 & \frac{n}{t} & 0 & 0 \\ 0 & 0 & -\frac{f+n}{f-n} & -\frac{2fn}{f-n} \\ 0 & 0 & -1 & 0 \end{bmatrix} \begin{bmatrix} x_c \\ y_c \\ z_c \\ w_c \end{bmatrix} \quad (4.13)$$

Listing 4.2: Example implementation of a perspective matrix (see Equation 4.13).

```

1  inline
2  Matrix4 Perspective(float fov,
3                     float aspect,
4                     float nearz,
5                     float farz)
6  {
7      float top = tan(fov * 0.00872664625) * nearz; /* 0.00872664625 = PI/360 */
8
9      Matrix4 matrix;
10     memset(matrix, 0, sizeof(GLfloat) * 16);
11
12
13     matrix[0] = nearz / (top * aspect);
14     matrix[5] = nearz / top;
15     matrix[10] = -(farz + nearz) / (farz - nearz);
16     matrix[11] = -1;
17     matrix[14] = -(2 * farz * nearz) / (farz - nearz);
18
19     return matrix;
20 } // End Perspective(..)

```

4.2.5 Camera (LookAt)

In Vulkan, you'll need to explicitly define a camera object for camera transformation. The camera or view matrix is responsible for trans-

forming the entire scene inversely to the origin (0,0,0) and always looking along -Z axis (this space is called eye space).

You construct a view matrix using the LookAt technique. You define the camera location at the eye position, the position you want the camera to look at (or rotating to) the target point target position. You must remember, the eye position and target are defined in 'world space'. The camera LookAt transformation consists of two transformations:

- (M_T) translating the whole scene inversely from the eye position to the origin
- (M_R) rotating the scene with reverse orientation (MR), so the camera is positioned at the origin and facing to the -Z axis

$$M_{view} = M_R \cdot M_T \quad (4.14)$$

The translation part of LookAt transformation is the simplest part to remember as all you need to do is move the camera position to the origin. The translation matrix M_T would be the negation of the eye position.

$$M_T = \begin{bmatrix} 1 & 0 & 0 & -x_e \\ 0 & 1 & 0 & -y_e \\ 0 & 0 & 1 & -z_e \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (4.15)$$

The rotation part of the LookAt transformation requires you to calculate 1st, 2nd and 3rd columns of the rotation matrix.

$$M_R = \begin{bmatrix} l_x & u_x & f_x & 0 \\ l_y & u_y & f_y & 0 \\ l_z & u_z & f_z & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}^{-1} = \begin{bmatrix} l_x & u_x & f_x & 0 \\ l_y & u_y & f_y & 0 \\ l_z & u_z & f_z & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}^T = \begin{bmatrix} l_x & l_y & l_z & 0 \\ u_x & u_y & u_z & 0 \\ f_x & f_y & f_z & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (4.16)$$

Finally, the view matrix for camera's LookAt transform is multiplying M_T and M_R together:

$$M_{view} = M_R M_T = \begin{bmatrix} l_x & l_y & l_z & 0 \\ u_x & u_y & u_z & 0 \\ f_x & f_y & f_z & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & -x_e \\ 0 & 1 & 0 & -y_e \\ 0 & 0 & 1 & -z_e \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} l_x & l_y & l_z & -l_x x_e - l_y y_e - l_z z_e \\ u_x & u_y & u_z & -u_x x_e - u_y y_e - u_z z_e \\ f_x & f_y & f_z & -f_x x_e - f_y y_e - f_z z_e \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (4.17)$$

Listing 4.3: Unsophisticated LookAt Camera View Implementation.

```

1  inline
2  Matrix4 LookAt( const Vector3& eye,
3                 const Vector3& target,
4                 const Vector3& upDir)
5  {
6      // calculate the forward vector from target to eye
7      Vector3 forward = eye - target;
8      forward = Vector3::Normalize(forward); // make unit length
9
10     // calculate the left vector
11     Vector3 left = Vector3::Cross(upDir, forward); // cross product
12     left = Vector3::Normalize(left);
13
14     // recalculate the orthonormal up vector
15     Vector3 up = Vector3::Cross(forward, left); // cross product
16
17     // init 4x4 matrix
18     Matrix4 matrix;
19     matrix = Matrix4::Identity();
20
21     // set rotation part, inverse rotation matrix:  $M^{-1} = M^T$  for Euclidean transform
22     matrix[0] = left.x;
23     matrix[4] = left.y;
24     matrix[8] = left.z;
25     matrix[1] = up.x;
26     matrix[5] = up.y;
27     matrix[9] = up.z;
28     matrix[2] = forward.x;
29     matrix[6] = forward.y;
30     matrix[10] = forward.z;
31
32     // set translation part
33     matrix[12] = -left.x * eye.x - left.y * eye.y - left.z * eye.z;
34     matrix[13] = -up.x * eye.x - up.y * eye.y - up.z * eye.z;
35     matrix[14] = -forward.x * eye.x - forward.y * eye.y - forward.z * eye.z;
36
37     return matrix;
38 } // End LookAt(..)

```

A typical implementation of the LookAt transformation calculation may look something like Listing 4.3.

4.3 Primitives

Primitives are the basic drawing elements (the building blocks for more complex geometry). For example, the most common and simplest primitive is the triangle. However, other simple primitives includes, points, lines, and even squares. In Vulkan, you need to specify the primitive type you'll be using, so the render output knows how to interpret your stream of data (e.g., three points for a triangle or two points for a line).

1. Lines

Lists, Strips, Fans, ..

2. Triangles

Lists Strips, Fans, ..

The primitive topology is described in Vulkan via the `VkPrimitiveTopology` enumerated type as shown below in Listing 4.3:

```
1  typedef enum VkPrimitiveTopology {
2      VK_PRIMITIVE_TOPOLOGY_POINT_LIST          = 0,
3      VK_PRIMITIVE_TOPOLOGY_LINE_LIST          = 1,
4      VK_PRIMITIVE_TOPOLOGY_LINE_STRIP        = 2,
5      VK_PRIMITIVE_TOPOLOGY_TRIANGLE_LIST     = 3,
6      VK_PRIMITIVE_TOPOLOGY_TRIANGLE_STRIP    = 4,
7      VK_PRIMITIVE_TOPOLOGY_TRIANGLE_FAN      = 5,
8      VK_PRIMITIVE_TOPOLOGY_LINE_LIST_WITH_ADJACENCY = 6,
9      VK_PRIMITIVE_TOPOLOGY_LINE_STRIP_WITH_ADJACENCY = 7,
10     VK_PRIMITIVE_TOPOLOGY_TRIANGLE_LIST_WITH_ADJACENCY = 8,
11     VK_PRIMITIVE_TOPOLOGY_TRIANGLE_STRIP_WITH_ADJACENCY = 9,
12     VK_PRIMITIVE_TOPOLOGY_PATCH_LIST         = 10,
13     VK_PRIMITIVE_TOPOLOGY_BEGIN_RANGE = VK_PRIMITIVE_TOPOLOGY_POINT_LIST,
14     VK_PRIMITIVE_TOPOLOGY_END_RANGE   = VK_PRIMITIVE_TOPOLOGY_PATCH_LIST,
15     VK_PRIMITIVE_TOPOLOGY_RANGE_SIZE  = (VK_PRIMITIVE_TOPOLOGY_PATCH_LIST -
16     VK_PRIMITIVE_TOPOLOGY_POINT_LIST + 1),
16     VK_PRIMITIVE_TOPOLOGY_MAX_ENUM    = 0x7FFFFFFF
17 } VkPrimitiveTopology;
```

4.3.1 Backface Culling (Clockwise/Counter-Clockwise)

The draw order enables the graphical API to ‘cull’ unseen triangles (i.e., triangles have two sides - front and back - the triangles facing away from the viewer aren’t drawn). Hence, you need to define the preferred drawing order when you initialize Vulkan (or any other graphical API). The draw order is described in Vulkan via the `VkFrontFace` enumerated type as shown below in Listing 4.3.1:

```
1  typedef enum VkFrontFace {
2      VK_FRONT_FACE_COUNTER_CLOCKWISE = 0,
3      VK_FRONT_FACE_CLOCKWISE        = 1,
4      VK_FRONT_FACE_BEGIN_RANGE      = VK_FRONT_FACE_COUNTER_CLOCKWISE,
5      VK_FRONT_FACE_END_RANGE        = VK_FRONT_FACE_CLOCKWISE,
6      VK_FRONT_FACE_RANGE_SIZE       = (VK_FRONT_FACE_CLOCKWISE -
6      VK_FRONT_FACE_COUNTER_CLOCKWISE + 1),
7      VK_FRONT_FACE_MAX_ENUM         = 0x7FFFFFFF
8  } VkFrontFace;
```

To distinguish between the two sides you use the following convention (see Figure 4.14):

$$\begin{aligned} e0 &= v1 - v0 \\ e1 &= v2 - v0 \\ n &= \frac{e0 \times e1}{||e0 \times e1||} \end{aligned} \quad (4.18)$$




Figure 4.13: Format and configuration of the geometric data needs to be specified.




Figure 4.14: Backface culling removes (doesn’t draw) triangles that are facing away from the viewer. The direction of the triangle (front/back) is determined by the winding order.

where v_0 , v_1 and v_2 are the three corner positions of the triangle, e_0 and e_1 are the edges of the triangle and n is the triangle normal. The side the normal vector emanates from is the front side and the other side is the back side. You say the triangle is front-facing if the viewer (camera) sees the front side of the triangle, while the triangle is back-facing if the viewer sees the back side. Importantly, the front or back facing direction is determined by the ‘ordering’ of the vertices. This is not hard-coded either - as you set the ordering in the Vulkan API (i.e., the way you compute the triangle normal), a triangle ordered clockwise (with respect to that viewer) is front-facing, and a triangle ordered counter-clockwise (with respect to that viewer) is back-facing.

In reality, most 3-dimensional meshes are solid (totally enclosed). The object is constructed so the outside surface of the object has the triangle normals facing outwards. Resulting in the camera seeing the front-facing triangles of a solid object while the back-facing triangles are occluded (culled).

In addition to defining the front facing triangle draw order you also must explicitly define the culling mode. The culling mode is described in Vulkan via the `VkCullModeFlagBits` enumerated type as shown below in Listing 4.3.1:

```
1 typedef enum VkCullModeFlagBits {
2     VK_CULL_MODE_NONE           = 0,
3     VK_CULL_MODE_FRONT_BIT      = 0x00000001,
4     VK_CULL_MODE_BACK_BIT       = 0x00000002,
5     VK_CULL_MODE_FRONT_AND_BACK = 0x00000003,
6     VK_CULL_MODE_FLAG_BITS_MAX_ENUM = 0x7FFFFFFF
7 } VkCullModeFlagBits;
```

You’ll apply these enumeration types in later Sections when you implement your Vulkan graphical application (e.g., when constructing the Vulkan render pipeline in Listing 6.13).

4.4 Data/Geometry

The basic building block of all 3D object (and scenes) is typically a triangle. A triangle can be created by connecting 3 points or vertices to each other (in 2D or 3D). More complex shapes can be created by adding and assembling more triangles. The geometry can be stored in various file formats or generated procedurally. In addition, the triangles may have color information, texture details and even lighting specific data added to them to generate highly realistic outputs.

In this book, you’ll use simple geometry, such as, triangles, planes and cubes to demonstrate various graphical techniques. However, you’ll




Figure 4.15: Simplified illustration of Vertices to Pixels.




Figure 4.16: Wireframe render of a high-poly car with no materials (i.e., basic lighting) to illustrate the decomposition of a model as simpler primitives/triangles.

eventually want to draw more complex models/scenes (e.g., Figure 4.16). Of course, manually typing in the vertex/color information for details meshes would be insane. In computer graphics there are usually lots of complicated and interesting models freely available which are prettier to look at than simple planes and cubes.

While you might want to write your own model loading implementation, a quick solution is to take advantage of popular free open source solutions solution. For example, one such model loading library is:

Open Asset Import Library (short name: Assimp) [2] which is a portable Open Source library to import various well-known 3D model formats in an uniform manner. Assimp is able to import dozens of different model file formats by loading all the model's data into generalized data structures. As soon as Assimp has loaded the model, you can retrieve all the data you need from the data structures and convert them to your Vulkan specified layout. This becomes valuable once you've got your Vulkan application up and running and you want to start adding to its functionality.

4.5 Drawing Principles

In Vulkan (and with other modern graphical API), the viewing frustum is mapped to a cube that extends from -1 to 1 in the x , y and z (see Figure 4.17. Note, you can flip the z -axis to create a left handed coordinate system during projection transformation discussed in previous Sections when you convert from 3D to 2D.

1. Data (triangles) are passed to the renderer
2. Transforms are applied (on the shader) to the vertices (triangles)
3. The 'rasterization' process draws the geometry to the image (back-buffer screen)
4. Various optimisations/enhancements take place:
 - Depth buffer so geometry is drawn in the correct order
 - Culling so only clockwise (or counter-clockwise depending upon the configuration) triangles are drawn (i.e., backface culling)
 - Clipping (the output render frustum is mapped to a -1 to 1 clip space region)

4.6 Programmable Graphics & Shaders

Shaders basically give you the ability to customize your graphics card (akin to programming your CPU). The GPU has different programmable stages that are specifically optimised to perform specific




Figure 4.17: Pass-through (identity-transform) for the graphics renderer will output primitives within the clip space clip cube.

operations (e.g., vertex level or pixel level). For instance, the vertex shader transforms all the vertices positions in virtual space (your 3D model space) to the 2D coordinate which appear on screen (2D screen space). The fragment shader basically gives you the ability to manipulate the pixel information, such as, the pixel color/brightness.

In addition to the vertex and pixel shader there are other shader types. The shader types available in Vulkan are accessed via the `VkShaderStageFlagBits` enumerated type as shown below in Listing 4.6:

```

1  typedef enum VkShaderStageFlagBits {
2      VK_SHADER_STAGE_VERTEX_BIT          = 0x00000001,
3      VK_SHADER_STAGE_TESSELLATION_CONTROL_BIT = 0x00000002,
4      VK_SHADER_STAGE_TESSELLATION_EVALUATION_BIT = 0x00000004,
5      VK_SHADER_STAGE_GEOMETRY_BIT        = 0x00000008,
6      VK_SHADER_STAGE_FRAGMENT_BIT        = 0x00000010,
7      VK_SHADER_STAGE_COMPUTE_BIT         = 0x00000020,
8      VK_SHADER_STAGE_ALL_GRAPHICS        = 0x0000001F,
9      VK_SHADER_STAGE_ALL                  = 0x7FFFFFFF,
10     VK_SHADER_STAGE_FLAG_BITS_MAX_ENUM   = 0x7FFFFFFF
11 } VkShaderStageFlagBits;
```

Multiple shader flags, such as, 'VK_SHADER_STAGE_ALL_GRAPHICS' will become apparent when you start programming the graphical effects with the Vulkan API in later Chapters.

As you might be starting to see, most stages feed their output directly onto the next stage of the pipeline (hence the name 'pipeline'). For instance, the Vertex Shader stage inputs data from the Input Assembler stage, does its own work, and then outputs its results to the Geometry Shader stage (see Figure 4.18).

- **Input Assembler Stage** The start of the pipeline - reads geometric data (vertices and indices) from memory and uses it to assemble geometric primitives (such as, triangles and lines)
- **Vertex Stage** After the primitives have been assembled, the vertices are fed into the vertex shader stage. The vertex shader can be thought of as a function that inputs a vertex and outputs a vertex (one vertex in and one vertex out).
- **Tessellator Stage** As the name indicates, this stage is responsible for tessellation - that is this stage subdivides the triangles of a mesh to add new triangles. These new triangles can then be offset into new positions to create finer mesh detail
- **Geometry Stage** The geometry shader stage is optional. You'll learn about the geometry shader in Chapter 9. The geometry shader inputs entire primitives. For example, if you were drawing triangle lists, then the input to the geometry shader would be the three vertices defining the triangle. Crucially, the geometry shader is able to create and destroy geometry (unlike the vertex stage). For example,

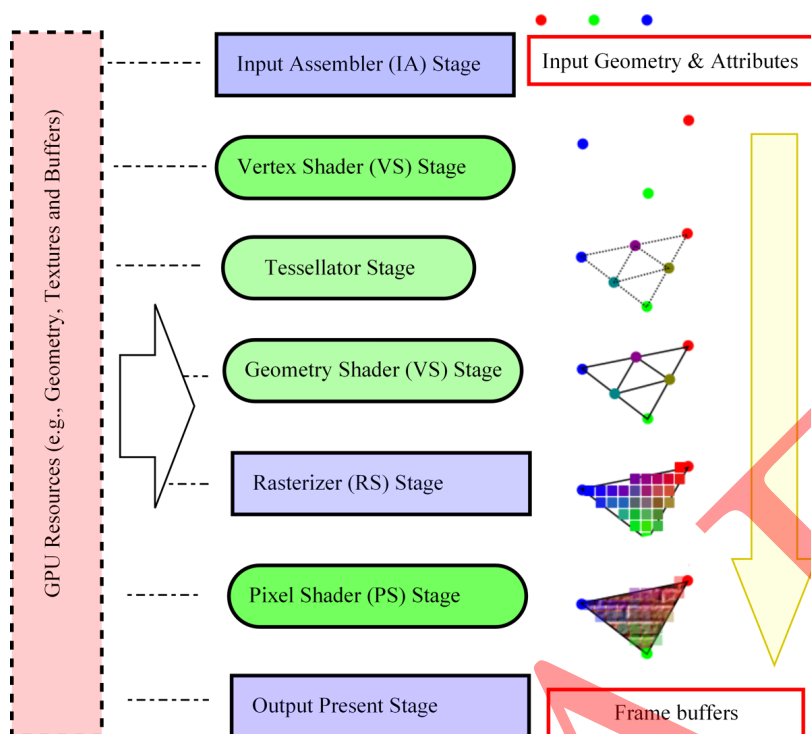


Figure 4.18: Common rendering pipeline stages. The arrow going from the GPU resources to the different stage means different stages can access the resources as input. For instance, the pixel shader stage may need to read data from a texture resource stored in memory in order to do its work. Also the downward arrow indicates the 'pipeline' flow from stage to stage.

the input primitive can be expanded into one or more other primitives, or the geometry shader can choose not to output a primitive based on some condition. You'd be able to pass in a single vertex to the geometry shader and output an entire geometric shape (or no shape at all)

- **Rasterization Stage** The main job of the rasterization stage is to compute pixel colors from the projected 3D triangles
- **Pixel (or Fragment) Stage** A pixel or fragment shader is executed for each pixel fragment and uses the interpolated vertex attributes as input to compute a color. A pixel shader can be as simple as returning a constant color, to doing more complicated things like per-pixel lighting, reflections and shadowing effects
- **Final Output Stage** After pixel fragments have been generated by the pixel (fragment) shader, they move onto the final output stage of the rendering pipeline. In this stage, some pixel fragments may be rejected (e.g., from the depth or stencil buffer tests). Pixel fragments that are not rejected are written to the back buffer. Blending is also done in this stage, where a pixel may be blended with the pixel currently on the back buffer instead of overriding it completely

4.7 Exercises

A number of well written books are available on the principles of computer graphics which complement this text (e.g., mathematics and geometry). Once you've got your Vulkan graphical application up and running you'll be able to extend your simple implementation developed in this book to encapsulate advanced features, such as, ambient occlusion, instancing, tessellation shader and post-processing.

Recommended books specifically focusing on Graphical Principles and Mathematics include:

- 3D math primer for graphics and game development by Dunn, Fletcher and Parberry, Ian [5]
- Computer Graphics: Principles and Practice (3rd Edition) by John Hughes et al. [7]
- Real-Time Rendering by Tomas Akenine-Moller et al. [3]
- 3D Graphics Programming: Games and Beyond by Sergei Savchenko [10]

4.7.1 Chapter Questions

Question What is the clip space?

Question Why are matrices used in graphical transforms?

Question What is a primitive?

Question What is the depth buffer used for?

Question Mention three differences between real-time graphics and off-line (photo-realistic) computer graphics. In this context, also explain why graphics hardware, e.g., graphics cards are useful for computer graphics

Question In the context of the graphics pipeline, describe the responsibility of the vertex shader, rasterizer and pixel shader stage of the graphics pipeline.

Question Mention three coordinate systems (spaces) that you may encounter in a rendering pipeline. Briefly explain the purpose of each system.

Question What is backface culling, why is it useful and where in the graphics pipeline can a backface culling test be executed?

Question What is a viewing frustum?

Question Why is the triangle strip more desirable geometric primitive than a list of triangles?

Question What is the difference between convex and concave objects?

Question For the eye position $e=[0,2,0]$ and a target position $t=[0,-1,0]$ and a view-up vector $up=[1,1,0]$, what is the camera transformation matrix?

Question Write the perspective projection matrix. Multiply it by the given homogeneous point to demonstrate how it generates pixel coordinates that reflect perspective foreshortening:

$$\begin{bmatrix} & & & \\ & & & \\ & & & \\ & & & \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ w \end{bmatrix} = \begin{matrix} \text{Transformed 4D point} \\ [\quad , \quad , \quad , \quad] \end{matrix} \quad \begin{matrix} \text{Pixel value} \\ \text{becomes} \end{matrix} \begin{bmatrix} \quad , \quad \end{bmatrix}$$

Question Explain the differences between 'raster' and 'vector' graphics.

Question Distinguish between window port and viewport.

Question A cube is placed at the origin of a 3D system. Such that all its vertices have positive coordinate values and sides are parallel to the three principal axes. Indicate a convenient position of a viewer at which he can see a 2-point perspective projection.

Question Define vanishing points. Is the location of the vanishing point directly related to the viewing point?

Question What are the various logical graphic input primitives? What are the various input modes?

Question What are the different projection methods? Explain

Question Explain RGB and HSV color modelling?

Question What is homogeneous co-ordinate? Why is a homogeneous co-ordinate system needed in transformation matrix?

Question Derive the transformation matrix for perspective projection.

Question Explain the transformations with examples: (i) Reflection. (ii) Shear.

Question Explain what the Depth-Buffer method is and why we need it?

Question Explain parallel and perspective projections.

Question Discuss the Back-Face surface removal algorithm.

Question Explain depth buffer for visible surface detection in 3D

graphics.

Question What is view volume? How is it specified?

Question Discuss the Back-Face surface removal algorithm.

Question Explain window to viewport transformation.

5

*Shaders*5.1 *Introduction*

Shaders are the chocolate sauce on your ice-cream. They offer truly limitless possibilities. Shader are in almost all recent real-time graphical applications (like video games), not to mention, animated CGI movies. Some popular techniques that use shaders are: parallax-mapping (bump-mapping), phong-shading, cell-shading, bloom and high dynamic range lighting (HDR). So what are shaders? Shaders are small programs developed by 'you' with the ability to customize the graphical pipeline, such as:

- transform data (manipulate your geometry)
- determine colors (principles of light)
- animate and move data
- and much more

Once upon a time, long ago, graphical processing units had static pipelines that were configurable through flags and states using the configurable API. This, of course, stunted the creative juices of developers (prevented any customization). As time progressed OpenGL and DirectX (and now Vulkan) solved this problem by making the pipeline 'programmable' (initially via low-level assembly shader languages and later high level languages like GLSL (OpenGL Shading Language) and HLSL (High-Level Shader Language)).

Currently there are three major shader languages:

- Cg (Nvidia)
- HLSL (Microsoft)
Derived from Cg
- GLSL (OpenGL)

Note you're still able to write shaders in assembly for highly optimised

```
vec3 gamma = vec3(1.0/2.2, 1.0/2.2, 1.0/2.2);
finalColor = vec3(pow(linearColor.r, gamma.r),
                  pow(linearColor.g, gamma.g),
                  pow(linearColor.b, gamma.b));

vec3 gamma = vec3(1.0/2.2);
finalColor = pow(linearColor, gamma);
```

solutions but it's less common due advancements in compiler technologies and computational processing power. The main influences on the development and steering of these shader languages over the years have been the C-language and pre-existing solutions developed in universities and industry with the HLSL coming from Microsoft in 2002 and later GLSL for OpenGL ARB in 2003.

Example applications of vertex shaders (run on per-vertex level) include:

- Color
- Texture
- Position
- Do not change the data type (pass-through)

Example applications of fragment shaders (pixel shaders - run on per-pixel level) include:

- Lighting values
- Output certain color
- Computationally expensive for complex effects due to per-pixel calculation (i.e., every pixel vs every vertex in the vertex shader)

Example applications of the geometry shaders (manipulates graphical primitives to create new primitives - points, lines and triangles) include:

- Shadow volumes
- Cube map (skybox)

The aim of this chapter is to provide concept and language fundamentals essential to high level shader languages common for graphical processing (e.g., GLSL 4+) (i.e., not provide a full in depth programming guide on shader programming). For those of you who are new to shader programming, there are several resources available, such as, books and web coding sandboxes (for instance Shader Toys, GLSL Sandbox and Vertex Shader Art), that are recommended for learning and developing your shader programming skills. Furthermore, shader experience gained through using different programming interfaces, platforms and tool-kits can be easily translated between the different APIs/tools (Vulkan, DirectX and OpenGL) - large amount of overlap and similarity.

5.1.1 Anatomy of Shaders

Shader is a program written in textual form (human readable). You'll find these small (shader) programs have these parts:

- Global variables
- Functions
 - Local variables (also variables passed through functions)
- Means of pass arbitrary data from Application to Shader (e.g., uniforms)
- Data structure definitions

Current shaders are written in C-type languages. Vulkan only accepts the Spir-V shader format, however, GLSL shader files can be compile to Spir-V files using the shader SDK compiler (e.g., glslangValidator.exe). Hence, this chapter will focus on GLSL examples, which can be compiled and implemented with the Vulkan API, and should be straightforward to port to other API (DirectX HLSL or existing OpenGL implementations). The compiled Spir-V files will typically have the extension '.spv'. The options for the glslangValidator compiler are given below in Listing 5.1:

Listing 5.1: glslangValidator command prompt options.

```

1  Usage: glslangValidator [option]... [file]...
2
3  Where: each 'file' ends in <stage>, where <stage> is one of
4  .conf to provide an optional config file that replaces the default configuration
5  (see -c option below for generating a template)
6  .vert for a vertex shader
7  .tesc for a tessellation control shader
8  .tese for a tessellation evaluation shader
9  .geom for a geometry shader
10 .frag for a fragment shader
11 .comp for a compute shader
12
13 Compilation warnings and errors will be printed to stdout.
14
15 To get other information, use one of the following options:
16 Each option must be specified separately.
17 -V create SPIR-V binary, under Vulkan semantics; turns on -l;
18     default file name is <stage>.spv (-o overrides this)
19     (unless -o is specified, which overrides the default file name)
20 -G create SPIR-V binary, under OpenGL semantics; turns on -l;
21     default file name is <stage>.spv (-o overrides this)
22 -H print human readable form of SPIR-V; turns on -V
23 -E print pre-processed GLSL; cannot be used with -l;
24     errors will appear on stderr.
25 -c configuration dump;
26     creates the default configuration file (redirect to a .conf file)
27 -d default to desktop (#version 110) when there is no shader #version
28     (default is ES version 100)
29 -h print this usage message
30 -i intermediate tree (glslang AST) is printed out
31 -l link all input files together to form a single module

```

```

32  -m          memory leak mode
33  -o <file>  save binary into <file>, requires a binary option (e.g., -V)
34  -q          dump reflection query database
35  -r          relaxed semantic error-checking mode
36  -s          silent mode
37  -t          multi-threaded mode
38  -v          print version strings
39  -w          suppress warnings (except as required by #extension : warn)

```

In order to understand how shaders work and how you'll use them to extend the drawing capabilities of graphical processing, it is necessary to have an overview of the key concepts of shader programming, first in general and then from the point of view of the graphical processing unit (GPU).

Initially, there was only two programmable stages (e.g., vertex and fragment stages) but as the thirst for freedom continued to grow - more and more stages and control has been given to developers. In this book, the four main shader stages you'll explore are:

- Vertex Stage (Per-Vertex Processing) - e.g., transforming geometry (vertices) to their final space/position
- Fragment Stage (Per-Fragment Processing) - e.g., providing coloring information to the pixel
- (Optional) Geometry Stage - e.g., extends the Vertex Stage with the added ability to add/remove geometry (also able to know about neighbouring primitives)
- (Optional) Tessellation Stage - e.g., ability to add detail to the geometry (add/remove triangles)

These different stages are not static but programmable. These stages are controlled by programs known as 'shaders'. Importantly, you have to implement the 'vertex' and 'fragment' shader (compulsory shaders required to output to the screen), while the geometry and tessellation stages are optional extras (e.g., you don't need to implement them to generate graphical renders). The shaders responsible for processing the different stages are all compiled using the same 'glslangValidator.exe' (see Listing 5.1).

The extension for the shader programs are:

- .vert** Vertex Shader
- .frag** Fragment Shader (or Pixel Shader)
- .geom** Geometry Shader (Chapter 9)
- .tesc** Tessellation Shader - Control Stage (Chapter 14)
- .tese** Tessellation Shader - Evaluation Stage (Chapter 14)

Even though the graphical pipeline has changed from a static to a programmable paradigm each stage of the pipeline is still responsible for their original tasks (e.g., transforms and lighting).

5.2 Link between Vulkan and Shaders

The Vulkan API has two fronts. The client-side and the server-side. The Vulkan API operates on your application side (client-side), while the shaders operate on the GPU side (server-side). One important responsibility of Vulkan is to link the data to the shaders' (e.g., using layouts and uniforms).

Data in your application is transported to the GPU. Once on the GPU this data your shader will be able to use this data. Your data is linked to your shader through attributes (e.g., specify bindings and locations).

Listing 5.2: Vertex Shader.

```
1  // shader version
2  #version 420
3
4  // 1. input attribute from your program declared as 'inPos'
5  layout( location = 0 ) in vec4 inPos;
6
7  // 2. Uniforms for the model-view-projection transforms
8  layout ( binding = 0 ) uniform buffer
9  {
10     mat4 inProjectionMatrix;
11     mat4 inViewMatrix;
12     mat4 inModelMatrix;
13 };
14
15 // 3. shader program entry point
16 void main()
17 {
18     // combine the matrices
19     mat4 mvp = inModelMatrix * inViewMatrix * inProjectionMatrix;
20
21     // transform position by matrices
22     gl_Position = inPos.xyz * mvp;
23 } // End main(..)
```

The shader version number at the top of each shader file (e.g., #version 420) allows you to know what features/syntax are used by your shader. When no shader version is specified, the default is ES version 100 (#version 110).

Shaders files follow the standard C/C++ commenting syntax (allowing you to make notes/explain your shader code):

- `/* comment */` - starting and ending markers for comments (suitable for multiple line comments)
- `//` - everything after the double slash until a newline is a comment (suitable for single line comments)

5.3 Linking data to Uniforms

Linking data to Uniforms is very similar to linking data to attributes. Uniform variables are those that remain constant for each vertex in the scene. You create a buffer (allocate memory on the GPU for the uniforms). You then specify the location of the uniform in the shader. Once you know the location, you provide data to the uniform (e.g., lock and copy the data across). The model, view and projection matrices for transformations fall in this category, since each vertex in the scene is affected by the same model/view/projection matrices.

```

1  // ** 1 ** Creation of buffer/uniform
2
3  // Create 'uniform' buffer for passing constant
4  // data to the shader (connecting shader with the data)
5
6  // create our uniforms buffers:
7  VkBufferCreateInfo bufferCreateInfo;
8  memset(&bufferCreateInfo, 0, sizeof(bufferCreateInfo));
9  bufferCreateInfo.sType      = VK_STRUCTURE_TYPE_BUFFER_CREATE_INFO;
10 // size in bytes
11 bufferCreateInfo.size       = sizeof(stBuffer);
12 bufferCreateInfo.usage      = VK_BUFFER_USAGE_UNIFORM_BUFFER_BIT;
13 bufferCreateInfo.sharingMode = VK_SHARING_MODE_EXCLUSIVE;
14
15 VkResult result =
16 vkCreateBuffer( device,
17                &bufferCreateInfo,
18                NULL,
19                outBuffer );
20 DBG_ASSERT_VULKAN_MSG( result,
21                        "Failed to create uniforms buffer." );
22
23 // ** 2 ** Allocate memory for buffer:
24 VkMemoryRequirements bufferMemoryRequirements = {};
25 vkGetBufferMemoryRequirements( device,
26                                *outBuffer,
27                                &bufferMemoryRequirements );
28
29 VkMemoryAllocateInfo matrixAllocateInfo = {};
30 matrixAllocateInfo.sType      = VK_STRUCTURE_TYPE_MEMORY_ALLOCATE_INFO;
31 matrixAllocateInfo.allocationSize = bufferMemoryRequirements.size;
32
33 VkPhysicalDeviceMemoryProperties memoryProperties;
34 vkGetPhysicalDeviceMemoryProperties( physicalDevice, &memoryProperties );
35
36 for ( uint32_t i = 0; i < VK_MAX_MEMORY_TYPES; ++i )
37 {
38     VkMemoryType memoryType = memoryProperties.memoryTypes[i];
39     // is this the memory type we are looking for?
40     if ( ( memoryType.propertyFlags & VK_MEMORY_PROPERTY_HOST_VISIBLE_BIT ) )
41     {
42         // save location
43         matrixAllocateInfo.memoryTypeIndex = i;
44         break;
45     }
46 } // End for i

```

```

47
48     result =
49     vkAllocateMemory( device, &matrixAllocateInfo, NULL, outMemory );
50     DBG_ASSERT_VULKAN_MSG( result,
51         "Failed to allocate uniforms buffer memory." );
52
53     result =
54     vkBindBufferMemory( device, *outBuffer, *outMemory, 0 );
55     DBG_ASSERT_VULKAN_MSG( result,
56         "Failed to bind uniforms buffer memory." );
57
58     // ** 3 ** Lock/update the uniforms contentsline
59     uint8_t *pData;
60     DBG_ASSERT_VULKAN(vkMapMemory(device, memory, 0, sizeof(ubo), 0, (void **)&pData));
61     memcpy(pData, &ubo, sizeof(ubo));
62     vkUnmapMemory(device, memory);

```

5.3.1 Qualifiers

Variables in shaders take on different behaviours. Some variables can only receive data, others can only provide data. In fact, some of these variables can only be used in the Vertex Shaders, while other variables can only be used in the Geometry or Fragment Shaders. To differentiate these type of variables there are Qualifier Types. For example:

- in/out
- uniforms
- varying

The in/out shader qualifier defines the receiving/sending of data from buffers and whose value may change frequently.

5.3.2 Uniforms

A Uniform is a shader qualifier whose value may rarely change. You can think of uniforms as global variables which can be seen by all shader types.

5.3.3 Varying

There are times when attribute data needs to be used in different stages of the pipeline (different shaders). In this cases, special type of qualifiers called Varying are used. They take attribute data from the current shader and pass them along to the next shader stage.

5.4 Developing Shaders

You're now in the position to develop your own shaders. You're going to write source code for a vertex and fragment shader. The shader implementations are written in a C-style format. Hence, you can use your favourite text editing program (e.g., notepad or Visual Studio).

The Listing 5.3 below shows the simple.vert file contents. This file will contain your vertex shader source code. Your vertex shader will simply receive vertex data through the input. You'll also receive a Model-View-Projection matrix through a uniform (MVP). You'll then transform the vertex positions by this matrix and set it as the output of the shader. As you'll notice in the vertex shader below, you provide the result to the output of the shader 'gl_Position'. This is a built-in variable for the graphical pipeline (i.e., non-programmable aspects of the pipeline - for instance, determining clipping/calculating specific data for the next stage).

Listing 5.3: Basic Vertex Shader.

```
1 // vertex shader version
2 #version 420
3
4 // input vertex data (i.e., single position)
5 layout (location = 0) in vec4 inPos;
6
7 // single uniform parameter (transforms) - shared by all vertices
8 layout (binding = 0) uniform UBO
9 {
10     mat4 MVP;
11 } ubo;
12
13 void main()
14 {
15     // transformed vertex position for the next stage
16     gl_Position = ubo.MVP * inPos;
17 } // End main(..)
```

Next the file called simple.frag shown below Listing 5.4 is the fragment shader:

Listing 5.4: Basic Fragment Shader.

```
1 // fragment version
2 #version 420
3
4 // output for this fragment (single pixel color)
5 layout (location = 0) out vec4 outFragColor;
6
7 void main()
8 {
9     // constant color - white - all the triangles/primitives would be white
10    outFragColor = vec4(1);
```

```
11  } // End main(..)
```

The shaders above need to be compiled prior to being loaded and used by the Vulkan API (i.e., binary file that is readable by the GPU). When compiling your shader files, ensure you check the output for errors (e.g., typing errors, like spelling mistakes or missing semi-colons). If there are errors in your shader text file your shader compiler will not generate a binary (check the compiler output to ensure it says 'successful').

Common data types:

- int, float, bool, void
- vec2, vec3, vec4
- ivec2, ivec3, ivec4
- mat3, mat4
- sampler2D

For vectors you use the following accessors: 'xyzw' or 'rgba' (including combinations, such as .xy, .xyz) - this makes the shader implementation very flexible and compact.

Examples:

```
1  // mat4 to mat3
2  mat3 viewMatrix = mat3(inViewMatrix);
3
4  // selecting a row from a matrix and convert it to a vector
5  vec3 eye = -inViewMatrix[3].xyz;
6
7  // combine the matrices (multiply matrices together)
8  mat4 mvp = inModelMatrix * inViewMatrix * inProjectionMatrix;
9
10 // converting between types (explicitly)
11 gl_Position = vec4(inPos.xyz, 1.0) * mvp;
```

Common built in functions:

- max
- min
- clamp
- mix
- normalize
- length
- dot
- cross
- texture
- reflect
- pow
- transpose

- inverse
- cos
- sin
- tan
- sqrt

Subsequent chapters you'll implement different shader techniques that enable you to understand the concepts in greater detail, such as, texturing, lighting and geometrical manipulation.

You can create your own structures (i.e., using the 'struct' definition) and your own functions to make your code more manageable and scalable (i.e., you don't need to repeat code but create reusable functions - readability).

5.5 Summary

At the end of this chapter you should be starting to see the incredible power of shaders. With little information shaders are able to create an infinite number of possibilities. In following chapters, you'll be delving much deeper into what you can do with shaders. Everything drawn on the screen has been processed by the appropriate 'shader' running behind the scenes. Modifying shaders incorporates a new set of functions and variables allows you to replace the default techniques with your own. This opens up many exciting possibilities: rendering 3D scenes using more creative and sophisticated solutions and algorithms.

5.6 Exercises

After you're familiar with the shaders, you'll need to constantly practice to strengthen your understanding. The following example questions provide you this opportunity.

5.6.1 Chapter Questions

Question What is the advantage of a programmable pipeline vs a fixed pipeline?

Question In addition to the 'vertex shader' and the 'fragment shader' - name two additional pipeline shaders?

Question What are 'uniforms' for and how and when would you use a uniform?

80 introduction to computer graphics and the vulkan api

Question Write a very basic ‘vertex’ and ‘fragment’ shader (which transforms the vertex coordinates and outputs colored geometry).

Question In your shader how would you convert a ‘vec3’ to a ‘vec4’?

Question In your shader how would you convert a ‘mat4’ to a ‘mat4’?

Question In your shader how would you extract row from a mat4 (i.e., vec4)?

6

Programming (11 Steps)

Welcome to the first coding steps to writing your Vulkan application. This section, you'll learn how to put together the various API elements in context. Essentially, you'll take a difficult and somewhat overwhelming task and develop a set of clear easy to understand functions. The implementation in this book has been broken down into 11 easy steps - making the implementation more manageable. Writing a native Vulkan graphical program can be a bit intimidating initially due to the amount of code and details (1000+ lines). Hence, to make setup/api programming aspect digestible and easier to debug, you'll subdivided your implementation into a set of self-contained functions (see Listing 6.1) as presented by Kenwright [8]. The implementation is 'functional' so variables are passed around, while and returned data/values are stored and used in subsequent methods. This way you avoid globals while learning and analysing the reasons behind the API/graphical concepts.

Figure 6.1: Example listings use the Microsoft Windows API for the platform specific details.

If you're completely new to the Vulkan API - manually typing in the code samples instead of just running the working program will help you absorb and understanding the principles better (more time consuming but aids in deep learning the subject).

Listing 6.1: Steps to initializing and running a basic Vulkan graphical application (split into 11 easy to understand functional stages). As you master each element, you'll expand and customize the implementation to deepen your understanding. Each of the steps is explain in detail in subsequent sections.

```

1 void main(int argc, char** argv)
2 {
3     // Step 1 - Initializing the Window
4     int width = 800;
5     int height = 600;
6     HWND windowHandle = NULL;
7     SetupWindow(width, height, &windowHandle);
8
9     // Step 2 - Initialize Vulkan (Section 6.1)
10    VkInstance instance = NULL;
11    VkSurfaceKHR surface = NULL;
12    SetupVulkanInstance(windowHandle,
13                        &instance,
14                        &surface);
15
16    // Step 3 - Find/Create Device and (Section 6.3)

```

```
17 // Set-up your selected device
18 VkPhysicalDevice physicalDevice = NULL;
19 VkDevice device = NULL;
20 SetupPhysicalDevice(instance,
21                     &physicalDevice,
22                     &device);
23
24
25 // Step 4 - Initialize Swap-Chain (Section 6.4)
26 VkSwapchainKHR swapChain = NULL;
27 VkImage* presentImages = NULL;
28 VkImageView* presentImageViews = NULL;
29 SetupSwapChain( device,
30                physicalDevice,
31                surface,
32                &width,
33                &height,
34                &swapChain,
35                &presentImages,
36                &presentImageViews);
37
38 // Step 5 - Create Render Pass (Section 6.5)
39 VkRenderPass renderPass = NULL;
40 VkFramebuffer* frameBuffers = NULL;
41 SetupRenderPass(device,
42                physicalDevice,
43                width,
44                height,
45                presentImageViews,
46                &renderPass,
47                &frameBuffers);
48
49 // Step 6 - Create Command Pool/Buffer (Section 6.6)
50 VkCommandBuffer commandBuffer = NULL;
51 SetupCommandBuffer(device,
52                physicalDevice,
53                &commandBuffer);
54
55 // Step 7 - Vertex Data/Buffer (Section 6.9)
56 VkBuffer vertexInputBuffer = NULL;
57 int vertexSize = 0;
58 int numberOfTriangles = 0;
59 SetupVertexBuffer(device,
60                physicalDevice,
61                &vertexSize,
62                &numberOfTriangles,
63                &vertexInputBuffer );
64
65 // Step 8 - Load/Setup Shaders (Section 6.10)
66 VkShaderModule vertShaderModule = NULL;
67 VkShaderModule fragShaderModule = NULL;
68 VkBuffer buffer = NULL;
69 VkDeviceMemory memory = NULL;
70 SetupShaderandUniforms(device,
71                physicalDevice,
72                width,
73                height,
74                &vertShaderModule,
75                &fragShaderModule,
76                &buffer,
77                &memory);
```

```

78
79 // Step 9 - Setup Descriptors/Sets (Section 6.12)
80 VkDescriptorSet descriptorSet = NULL;
81 VkDescriptorSetLayout descriptorSetLayout = NULL;
82 SetupDescriptors(device,
83                 buffer,
84                 &descriptorSet,
85                 &descriptorSetLayout);
86
87 // Step 10 - Pipeline (Section 6.13)
88 VkPipeline pipeline = NULL;
89 VkPipelineLayout pipelineLayout = NULL;
90 SetupPipeline( device,
91              width,
92              height,
93              vertexSize,
94              descriptorSetLayout,
95              vertShaderModule,
96              fragShaderModule,
97              renderPass,
98              &pipeline,
99              &pipelineLayout);
100
101 // Step 11 - Render Loop (Section 6.15)
102 MSG msg;
103 while( true )
104 {
105     // Continually force the window to be redrawn as long as no other Win32
106     // messages are pending.
107     PeekMessage( &msg, NULL, NULL, NULL, PM_REMOVE );
108     if( msg.message == WM_QUIT ) break;
109
110     // Your Window's applications is responsible for retrieving and
111     // dispatching input messages to the window GUI in the main message loop
112     TranslateMessage( &msg );
113     DispatchMessage( &msg );
114
115     // Render the screen
116     RenderLoop( device,
117              width,
118              height,
119              numberOfTriangles,
120              swapChain,
121              commandBuffer,
122              presentImages,
123              frameBuffers,
124              renderPass,
125              vertexInputBuffer,
126              descriptorSet,
127              pipelineLayout,
128              pipeline);
129 } // End while(..)
130 return 0;
131
132 } // End WinMain(..)

```

6.1 (Step 1 & 2) Initializing Vulkan (Instance Creation)

The initial step is to setup your window for your operating system. As you have to let Vulkan where you're going to draw to (e.g., screen or off-screen texture). The OS specific parts of the implementation will be done for Windows, however, it should be straightforward to modify these few occurrences for different systems (e.g., Android and Linux).

The first Vulkan specific step you'll need to do after setting up your window is to initialize Vulkan. This is subdivided into two main parts. To begin with, you have to identify the Vulkan driver and characteristics you want to enable (e.g., standard LUNARG driver or the NVidia one, also what layers are available). For example, in the below implementation, you'll use 'vkEnumerateInstanceLayerProperties' and 'vkEnumerateInstanceExtensionProperties' to list all the layers and the instance properties. For the example, in the example implementation below, 'VK_LAYER_NV_optimus' has been hardcoded as the layer. Typically, you'll then also have three extensions, one will be the debug extension ('VK_EXT_debug_report'), which you'll include if you want to initialize the error callback notifications (discuss next). The next two extensions will depend upon your operating system and what you're using Vulkan for (e.g., graphics, compute, ...).

Just to note, in the Listing examples in subsequent sections, you'll often encounter additional 'curly brackets' .. inside the functions. These additional curly brackets have been added to clump together blocks of code so it's easier to read (i.e., self-contained modular components).




Figure 6.2: Uncomplicated window displaying using the Windows API (Step 1 - Initializing the Window).

Listing 6.2: Initializing Vulkan - Vulkan doesn't exist until you create an the Vulkan instance (vkCreateInstance).

```

1  // Step 2 - Initialize Vulkan
2  void SetupVulkanInstance(HWND windowHandle,
3                           VkInstance* outInstance,
4                           VkSurfaceKHR* outSurface)
5  {
6      // Initialize VULKAN
7
8      // Layer properties
9      uint32_t count = 0;
10
11     // Returns the number of layer properties available, If the VkLayerProperties*
12     // is NULL, then the number of layer properties available is returned
13     VkResult result =
14         vkEnumerateInstanceLayerProperties(&count, // uint32_t*
15                                           // pointer to the number of layer properties available
16                                           NULL); // VkLayerProperties*
17     // pointer to an array of VkLayerProperties structures
18
19
20     DBG_ASSERT( VK_SUCCESS == result );

```

A

86 introduction to computer graphics and the vulkan api

```

81 #endif
82
83 {
84     VkApplicationInfo ai    = { };
85     ai.sType                = VK_STRUCTURE_TYPE_APPLICATION_INFO;
86     ai.pApplicationName     = "Hello Vulkan";
87     ai.engineVersion        = 1;
88     ai.apiVersion           = VK_API_VERSION_1_0;
89
90     VkInstanceCreateInfo ici = { };
91     ici.sType                = VK_STRUCTURE_TYPE_INSTANCE_CREATE_INFO;
92     ici.flags                = 0;
93     ici.pNext                = 0;
94     ici.pApplicationInfo     = &ai;
95     ici.enabledLayerCount    = 1;
96     ici.ppEnabledLayerNames  = layers;
97     ici.enabledExtensionCount = 2;
98 #ifdef ENABLE_VULKAN_DEBUG_CALLBACK // access debug callbacks
99     ici.enabledExtensionCount = 3;
100 #endif
101     ici.ppEnabledExtensionNames = extensions;
102
103     // vkCreateInstance verifies that the requested layers exist. If not,
104     // vkCreateInstance will return VK_ERROR_LAYER_NOT_PRESENT
105     VkResult result =
106     vkCreateInstance( &ici,           // const VkInstanceCreateInfo*
107                     // points to an instance of VkInstanceCreateInfo controlling creation
108                     NULL,             // const VkAllocationCallbacks*
109                     // controls host memory allocation
110                     outInstance ); // VkInstance*
111     // pointer to a VkInstance handle for the returning resulting instance
112
113     DBG_ASSERT_VULKAN_MSG( result,
114         "Failed to create vulkan instance." );
115     DBG_ASSERT(*outInstance!=NULL);
116 }
117
118 // Optional - if you want Vulkan to tell you if something is wrong
119 // you must set the callback
120 #ifdef ENABLE_VULKAN_DEBUG_CALLBACK
121 ...
122 #endif
123
124
125 // You need to define what type of surface you'll be
126 // rendering to - this will depend on your computer
127 // and operating system (Win32)
128 HINSTANCE hInst      = GetModuleHandle(NULL);
129
130 // setup parameters for your new windows
131 // surface you'll render into:
132 VkWin32SurfaceCreateInfoKHR sci = { };
133 sci.sType                = VK_STRUCTURE_TYPE_WIN32_SURFACE_CREATE_INFO_KHR;
134 // parameter is NULL, GetModuleHandle returns a handle
135 // to the file used to create the calling process
136 sci.hinstance            = hInst;
137 // Your window handle (HWND)
138 sci.hwnd                 = windowHandle;
139
140 DBG_ASSERT(*outSurface=NULL);
141

```

E

```

142 // To create a VkSurfaceKHR object for a Win32 window, call:
143 VkResult result =
144     vkCreateWin32SurfaceKHR( *outInstance, // VkInstance
145                             // instance to associate with the surface
146                             &sci,          // const VkWin32SurfaceCreateInfoKHR*
147                             // pointer to VkWin32SurfaceCreateInfoKHR structure parameters for the surface object
148                             NULL,          // const VkAllocationCallbacks*
149                             // allocator used for host memory allocated for the surface object
150                             outSurface ); // VkSurfaceKHR*
151                             // VkSurfaceKHR handle in which the created surface object is returned
152
153 DBG_ASSERT_VULKAN_MSG( result,
154     "Could not create surface." );
155 DBG_ASSERT(outSurface!=NULL);
156
157 }// End SetupVulkanInstance(..)

```

The process for setting up the Vulkan instance is:

- A** identify the available layers and extensions (e.g., `vkEnumerateInstanceLayerProperties`)
- E** create the Vulkan instance (completing the structure parameters for all the information, such as, version, name, ..)
- F** create the output surface and connect it with the operating system specific window (Window handle in this case)

While Listing 6.2, focuses on a Microsoft Windows solution, similar functions are available for platform specific Vulkan API (e.g., `vkCreateWin32SurfaceKHR`), such as, Android:

```

1 // To create a VkSurfaceKHR object for an Android native window, you'd call:
2 VkResult vkCreateAndroidSurfaceKHR(
3     VkInstance instance,
4     const VkAndroidSurfaceCreateInfoKHR* pCreateInfo,
5     const VkAllocationCallbacks* pAllocator,
6     VkSurfaceKHR* pSurface);

```

6.2 Debugging

You should start thinking about debugging (and defensive programming) from the start. For instance, a few reasons debugging in Vulkan is challenging:

- May be no obvious relationship between the manifestation(s) of the error and the causes(s)
- Symptoms and cause may be in remote/different parts of the program
- Changes (new features and bug fixes) in the program may mask (or modify) bugs

```

1 // Optional - if you want Vulkan to tell you if something is wrong
2 // you must set the callback
3 #ifdef ENABLE_VULKAN_DEBUG_CALLBACK
4 {
5     // Register your error logging function (defined at the top of the file)
6     VkDebugReportCallbackEXT error_callback = VK_NULL_HANDLE;
7     VkDebugReportCallbackEXT warning_callback = VK_NULL_HANDLE;
8
9     PFN_vkCreateDebugReportCallbackEXT vkCreateDebugReportCallbackEXT = NULL;
10
11     *(void **)&vkCreateDebugReportCallbackEXT =
12         vkGetInstanceProcAddr( *outInstance, "vkCreateDebugReportCallbackEXT" );
13     DBG_ASSERT(vkCreateDebugReportCallbackEXT);
14
15     VkDebugReportCallbackCreateInfoEXT cb_create_info = {};
16     cb_create_info.sType = VK_STRUCTURE_TYPE_DEBUG_REPORT_CREATE_INFO_EXT;
17     cb_create_info.flags = VK_DEBUG_REPORT_ERROR_BIT_EXT;
18     cb_create_info.pfnCallback = &MyDebugReportCallback;
19
20     // Setup error callback function notifications
21     VkResult result =
22         vkCreateDebugReportCallbackEXT( *outInstance,
23                                         // valid VkInstance handle
24                                         &cb_create_info,
25                                         // // pointer to a valid VkDebugReportCallbackCreateInfoEXT structure
26                                         nullptr,
27                                         // If pointer is not NULL then allocator callback manager
28                                         &error_callback;
29                                         // pointer to a VkDebugReportCallbackEXT handle
30
31     DBG_ASSERT_VULKAN_MSG(result, "vkCreateDebugReportCallbackEXT(ERROR) failed");
32
33     // Capture warning as well as errors
34     cb_create_info.flags = VK_DEBUG_REPORT_WARNING_BIT_EXT |
35         VK_DEBUG_REPORT_PERFORMANCE_WARNING_BIT_EXT;
36     cb_create_info.pfnCallback = &MyDebugReportCallback;
37
38     // Setup warning callback function notifications
39     result =
40         vkCreateDebugReportCallbackEXT( *outInstance,
41                                         // valid VkInstance handle
42                                         &cb_create_info,
43                                         // // pointer to a valid VkDebugReportCallbackCreateInfoEXT structure
44                                         nullptr,
45                                         // If pointer is not NULL then allocator callback manager
46                                         &warning_callback;
47                                         // pointer to a VkDebugReportCallbackEXT handle
48
49     // Done
50 }
51 #endif
52
53 // Done
54 }
55
56 // Done
57 }
58
59 // Done
60 }
61
62 // Done
63 }
64
65 // Done
66 }
67
68 // Done
69 }
70
71 // Done
72 }
73
74 // Done
75 }
76
77 // Done
78 }
79
80 // Done
81 }
82
83 // Done
84 }
85
86 // Done
87 }
88
89 // Done
90 }
91
92 // Done
93 }
94
95 // Done
96 }
97
98 // Done
99 }
100
101 // Done
102 }
103
104 // Done
105 }
106
107 // Done
108 }
109
110 // Done
111 }
112
113 // Done
114 }
115
116 // Done
117 }
118
119 // Done
120 }
121
122 // Done
123 }
124
125 // Done
126 }
127
128 // Done
129 }
130
131 // Done
132 }
133
134 // Done
135 }
136
137 // Done
138 }
139
140 // Done
141 }
142
143 // Done
144 }
145
146 // Done
147 }
148
149 // Done
150 }
151
152 // Done
153 }
154
155 // Done
156 }
157
158 // Done
159 }
160
161 // Done
162 }
163
164 // Done
165 }
166
167 // Done
168 }
169
170 // Done
171 }
172
173 // Done
174 }
175
176 // Done
177 }
178
179 // Done
180 }
181
182 // Done
183 }
184
185 // Done
186 }
187
188 // Done
189 }
190
191 // Done
192 }
193
194 // Done
195 }
196
197 // Done
198 }
199
200 // Done
201 }
202
203 // Done
204 }
205
206 // Done
207 }
208
209 // Done
210 }
211
212 // Done
213 }
214
215 // Done
216 }
217
218 // Done
219 }
220
221 // Done
222 }
223
224 // Done
225 }
226
227 // Done
228 }
229
230 // Done
231 }
232
233 // Done
234 }
235
236 // Done
237 }
238
239 // Done
240 }
241
242 // Done
243 }
244
245 // Done
246 }
247
248 // Done
249 }
250
251 // Done
252 }
253
254 // Done
255 }
256
257 // Done
258 }
259
260 // Done
261 }
262
263 // Done
264 }
265
266 // Done
267 }
268
269 // Done
270 }
271
272 // Done
273 }
274
275 // Done
276 }
277
278 // Done
279 }
280
281 // Done
282 }
283
284 // Done
285 }
286
287 // Done
288 }
289
290 // Done
291 }
292
293 // Done
294 }
295
296 // Done
297 }
298
299 // Done
300 }
301
302 // Done
303 }
304
305 // Done
306 }
307
308 // Done
309 }
310
311 // Done
312 }
313
314 // Done
315 }
316
317 // Done
318 }
319
320 // Done
321 }
322
323 // Done
324 }
325
326 // Done
327 }
328
329 // Done
330 }
331
332 // Done
333 }
334
335 // Done
336 }
337
338 // Done
339 }
340
341 // Done
342 }
343
344 // Done
345 }
346
347 // Done
348 }
349
350 // Done
351 }
352
353 // Done
354 }
355
356 // Done
357 }
358
359 // Done
360 }
361
362 // Done
363 }
364
365 // Done
366 }
367
368 // Done
369 }
370
371 // Done
372 }
373
374 // Done
375 }
376
377 // Done
378 }
379
380 // Done
381 }
382
383 // Done
384 }
385
386 // Done
387 }
388
389 // Done
390 }
391
392 // Done
393 }
394
395 // Done
396 }
397
398 // Done
399 }
400
401 // Done
402 }
403
404 // Done
405 }
406
407 // Done
408 }
409
410 // Done
411 }
412
413 // Done
414 }
415
416 // Done
417 }
418
419 // Done
420 }
421
422 // Done
423 }
424
425 // Done
426 }
427
428 // Done
429 }
430
431 // Done
432 }
433
434 // Done
435 }
436
437 // Done
438 }
439
440 // Done
441 }
442
443 // Done
444 }
445
446 // Done
447 }
448
449 // Done
450 }
451
452 // Done
453 }
454
455 // Done
456 }
457
458 // Done
459 }
460
461 // Done
462 }
463
464 // Done
465 }
466
467 // Done
468 }
469
470 // Done
471 }
472
473 // Done
474 }
475
476 // Done
477 }
478
479 // Done
480 }
481
482 // Done
483 }
484
485 // Done
486 }
487
488 // Done
489 }
490
491 // Done
492 }
493
494 // Done
495 }
496
497 // Done
498 }
499
500 // Done
501 }
502
503 // Done
504 }
505
506 // Done
507 }
508
509 // Done
510 }
511
512 // Done
513 }
514
515 // Done
516 }
517
518 // Done
519 }
520
521 // Done
522 }
523
524 // Done
525 }
526
527 // Done
528 }
529
530 // Done
531 }
532
533 // Done
534 }
535
536 // Done
537 }
538
539 // Done
540 }
541
542 // Done
543 }
544
545 // Done
546 }
547
548 // Done
549 }
550
551 // Done
552 }
553
554 // Done
555 }
556
557 // Done
558 }
559
560 // Done
561 }
562
563 // Done
564 }
565
566 // Done
567 }
568
569 // Done
570 }
571
572 // Done
573 }
574
575 // Done
576 }
577
578 // Done
579 }
580
581 // Done
582 }
583
584 // Done
585 }
586
587 // Done
588 }
589
590 // Done
591 }
592
593 // Done
594 }
595
596 // Done
597 }
598
599 // Done
600 }
601
602 // Done
603 }
604
605 // Done
606 }
607
608 // Done
609 }
610
611 // Done
612 }
613
614 // Done
615 }
616
617 // Done
618 }
619
620 // Done
621 }
622
623 // Done
624 }
625
626 // Done
627 }
628
629 // Done
630 }
631
632 // Done
633 }
634
635 // Done
636 }
637
638 // Done
639 }
640
641 // Done
642 }
643
644 // Done
645 }
646
647 // Done
648 }
649
650 // Done
651 }
652
653 // Done
654 }
655
656 // Done
657 }
658
659 // Done
660 }
661
662 // Done
663 }
664
665 // Done
666 }
667
668 // Done
669 }
670
671 // Done
672 }
673
674 // Done
675 }
676
677 // Done
678 }
679
680 // Done
681 }
682
683 // Done
684 }
685
686 // Done
687 }
688
689 // Done
690 }
691
692 // Done
693 }
694
695 // Done
696 }
697
698 // Done
699 }
700
701 // Done
702 }
703
704 // Done
705 }
706
707 // Done
708 }
709
710 // Done
711 }
712
713 // Done
714 }
715
716 // Done
717 }
718
719 // Done
720 }
721
722 // Done
723 }
724
725 // Done
726 }
727
728 // Done
729 }
730
731 // Done
732 }
733
734 // Done
735 }
736
737 // Done
738 }
739
740 // Done
741 }
742
743 // Done
744 }
745
746 // Done
747 }
748
749 // Done
750 }
751
752 // Done
753 }
754
755 // Done
756 }
757
758 // Done
759 }
760
761 // Done
762 }
763
764 // Done
765 }
```

```

44                                     // pointer to a valid VkDebugReportCallbackCreateInfoEXT structure
45                                     nullptr,
46                                     // If pointer is not NULL then allocator callback manager
47                                     &warning_callback);
48                                     // pointer to a VkDebugReportCallbackEXT handle
49
50     DBG_ASSERT_VULKAN_MSG(result, "vkCreateDebugReportCallbackEXT(WARN) failed");
51 }
52 #endif

```

A good habit to get into is using regular sanity checks throughout your implementation. For example, debug asserts (DBG_ASSERT) as shown below. The Vulkan API requires a large number of structures and fields to be setup and configured. For any unknown reason, such as, a typing mistake or some custom detail specific to the hardware, may result in the graphical application failing - importantly, leaving you struggling to work out why. Hence, try and check every return value (e.g., 'VK_SUCCESS') and if a function fails - trigger an assert (don't try and hide the problem) - have the error shout out with the details so you can investigate why it failed and resolve the problem asap. This is also a good habit to get into for helping others, as it makes your code more readable - so other developers are able to understand your code easier and if it fails they're also able to fix the problem quickly as well.

Listing 6.4: Custom asserts to provide an additional layer of checking. Custom asserts also provide flexibility (e.g., write to log files, trigger breakpoints, or disable them easily).

```

1  #if defined(_WIN32)
2  #define DBG_ASSERT(f) { if(!(f)){ __debugbreak(); }; }
3  #else
4  #define DBG_ASSERT(f) { } // Other operating system debug
5  #endif
6
7  // NAN Test
8  #define DBG_VALID(f) { if( (f) != (f) ){ DBG_ASSERT(false); } }
9
10 // Assert with message
11 #define DBG_ASSERT_MSG( val, errmsg ) \
12     if(!(val)){ \
13         DBG_ASSERT( false ) \
14     }
15
16 #define DBG_ASSERT_VULKAN_MSG( val, errmsg ) \
17     if(!( (VK_SUCCESS == val) )){ \
18         DBG_ASSERT( false ) \
19     }

```

While in the long run, you'd incorporate a variety of complex test functions within a structured framework (unit tests), yet asserts provide an effective and efficient debugging tool for identifying issues during the initial phases. You need to use a custom 'macro' instead of the system

assert directly, so you're able to control your asserts - like having the assert trigger a breakpoint at the line causing the validation fault. Furthermore, for release builds, you'd be able to customize the macro so instead of 'triggering' a breakpoint, you may want to write the error to a log file or bring up a dialog error window.

6.3 (Step 3) Device(s)

The system may have multiple devices. Each device may have similar or different capabilities. The physical device is identified in Vulkan using the type 'VkPhysicalDevice'. Provides a handle to query the device about its capabilities, such as, Memory Management Queues Objects Buffers Images and Sync Primitives. For example, the 'Geforce GTX 980' has different capabilities than the 'Tegra X1'.



Figure 6.4: The Vulkan API is designed to support 'multiple' devices with varying capabilities.

Listing 6.5: Determining what devices are on your system and with what capabilities.

```

1 // Step 3 - Find/Create Device
2 void SetupPhysicalDevice(VkInstance instance,
3                          VkPhysicalDevice* outPhysicalDevice,
4                          VkDevice* outDevice)
5 {
6     // Query how many devices are present in the system
7     uint32_t deviceCount = 0;
8     // Enumerates the physical devices accessible to a Vulkan instance
9     // The instance is the handle to a Vulkan instance you previously
10    // created with vkCreateInstance. The VkPhysicalDevice pointer
11    // can be either NULL or a pointer to an array of VkPhysicalDevice handles.
12    VkResult result =
13    vkEnumeratePhysicalDevices(instance, // VkInstance
14                              // handle to a Vulkan instance previously created with vkCreateInstance
15                              &deviceCount, // uint32_t*
16                              // pointer to an integer related to the number of physical devices available or queried
17                              NULL); // VkPhysicalDevice*
18    // either NULL or a pointer to an array of VkPhysicalDevice handles
19
20    DBG_ASSERT_VULKAN_MSG(result,
21                          "Failed to query the number of physical devices present");
22
23    // There has to be at least one device present
24    DBG_ASSERT_MSG(0 != deviceCount,
25                  "Couldn't detect any device present with Vulkan support");
26
27    // Get the physical devices
28    vector<VkPhysicalDevice> physicalDevices(deviceCount);
29
30    // Gets the VkPhysicalDevice handles.
31    result =
32    vkEnumeratePhysicalDevices(instance, // VkInstance
33                              // handle to a Vulkan instance previously created with vkCreateInstance
34                              &deviceCount, // uint32_t*
35                              // pointer to an integer related to the number of physical devices available or queried
36                              &physicalDevices[0]); // VkPhysicalDevice*
37    // either NULL or a pointer to an array of VkPhysicalDevice handles

```

```

38     DBG_ASSERT_VULKAN_MSG(result,
39         "Failed to enumerate physical devices present");
40     DBG_ASSERT(physicalDevices.size() > 0);
41
42     // Use the first available device
43     *outPhysicalDevice = physicalDevices[0];
44
45
46     // Enumerate all physical devices and print out the details
47     for (uint32_t i = 0; i < deviceCount; ++i)
48     {
49         VkPhysicalDeviceProperties deviceProperties;
50         memset(&deviceProperties, 0, sizeof deviceProperties);
51
52         // Gets the properties of a physical device
53         vkGetPhysicalDeviceProperties(physicalDevices[i], // physicalDevice
54                                     // handle to the physical device whose properties will be queried
55                                     &deviceProperties); // pProperties
56                                     // pointer to VkPhysicalDeviceProperties structure, that is filled with information
57
58         dprintf("Driver Version: %d\n", deviceProperties.driverVersion);
59         dprintf("Device Name: %s\n", deviceProperties.deviceName);
60         dprintf("Device Type: %d\n", deviceProperties.deviceType);
61         dprintf("API Version: %d.%d.%d\n",
62             (deviceProperties.apiVersion >> 22) & 0x3FF,
63             (deviceProperties.apiVersion >> 12) & 0x3FF,
64             (deviceProperties.apiVersion & 0xFFF));
65     } //End for i
66
67
68     // Fill up the physical device memory properties:
69     VkPhysicalDeviceMemoryProperties memoryProperties;
70     vkGetPhysicalDeviceMemoryProperties(*outPhysicalDevice,
71                                     // handle to the device to query
72                                     &memoryProperties);
73                                     // pointer to VkPhysicalDeviceMemoryProperties structure returned with properties
74
75     // Here's where you initialize your queues
76     // You'll discuss queues next - however, you need to specify the queue
77     // details for the device creation info
78     VkDeviceQueueCreateInfo queueCreateInfo = {};
79     queueCreateInfo.sType = VK_STRUCTURE_TYPE_DEVICE_QUEUE_CREATE_INFO;
80     // Use the first queue family in the family list
81     queueCreateInfo.queueFamilyIndex = 0;
82     queueCreateInfo.queueCount = 1;
83     float queuePriorities[] = { 1.0f };
84     queueCreateInfo.pQueuePriorities = queuePriorities;
85
86     // Same extension you specified when initializing Vulkan
87     const char *deviceExtensions[] = { "VK_KHR_swapchain" };
88     const char *layers[] = { "VK_LAYER_NV_optimus" };
89
90     VkDeviceCreateInfo dci = {};
91     dci.sType = VK_STRUCTURE_TYPE_DEVICE_CREATE_INFO;
92     // Set queue info on your device
93     dci.queueCreateInfoCount = 1;
94     dci.pQueueCreateInfos = &queueCreateInfo;
95     dci.enabledLayerCount = 0;
96     dci.ppEnabledLayerNames = layers;
97     dci.enabledExtensionCount = 1;

```

```

98   dci.ppEnabledExtensionNames    = deviceExtensions;
99
100   VkPhysicalDeviceFeatures features = {};
101   features.shaderClipDistance    = VK_TRUE;
102   dci.pEnabledFeatures           = &features;
103
104   // Ideally, you'd want to enumerate and find the best
105   // device, however, you just use the first device
106   // 'physicalDevices[0]' for your sample, which you
107   // stored in the previous section
108   result =
109   vkCreateDevice( *outPhysicalDevice, // physicalDevice
110                  // valid handles returned from vkEnumeratePhysicalDevices
111                  &dci,                // pCreateInfo
112                  // pointer to a VkDeviceCreateInfo structure containing device data
113                  NULL,                // pAllocator
114                  // optional control of host memory allocation
115                  outDevice );          // pDevice
116   // pointer to a handle in which the created VkDevice is returned
117
118   DBG_ASSERT_VULKAN_MSG( result, "Failed to create logical device!" );
119 } // End SetupPhysicalDevice(..)

```

With reference to Listing 6.5, the flow of logic to finding and creating the physical device are:

- A** vkEnumeratePhysicalDevices - query how many devices are present in the system
- B** vkEnumeratePhysicalDevices - call again to get the physical devices
- C** vkGetPhysicalDeviceProperties - get properties for each device (help make your decision on which one to choose or on the selected one)
- D** vkGetPhysicalDeviceMemoryProperties - more properties on the chosen device before you go ahead and create the device
- E** vkCreateDevice - finally create your device

6.4 (Step 4) Swap-Chain

There is 'no' default framebuffer in Vulkan. You are able to create an application that displays everything or nothing (total control). Hence, to display something you'll need to create a set of render buffers. These buffers (and their properties) are called the 'swap chain'. As emphasised, you have total control over your swap chain, which means, you can create and use lots of buffers however you want. A few important details when creating your swap chain image buffers:

1. define the surface format
2. create rendering context (connect the swap chain with the presentation output)




Figure 6.5: The swap-chain is a chain of buffers that swap position each time a new frame is rendered.

- Listing 6.6: Managing the screen capabilities and render surfaces.

[illegible]

```

54                                     // resulting swapchain
55     DBG_ASSERT_VULKAN_MSG( result,
56         "Failed to create swapchain." );
57 }
58
59 // Create your images 'double' buffering
60 {
61     uint32_t imageCount = 0;
62     // You'll need to obtain the array of presentable images associated
63     // with the swapchain you created. First, you pass in 'NULL' to
64     // obtain the number of images (i.e., should be 2)
65     vkGetSwapchainImagesKHR( device,          // device
66                             // device associated with swapchain
67                             *outSwapChain,    // swapchain
68                             // swapchain to query
69                             &imageCount,     // pSwapchainImageCount
70                             // pointer to an integer related to the number of format pairs available
71                             NULL );          // pSwapchainImages
72                                     // either NULL or a pointer to an array of VkSwapchainImageKHR structures
73     DBG_ASSERT(imageCount==2);
74
75     // this should be 2 for double-buffering
76     *outPresentImages = new VkImage[imageCount];
77
78     // Obtain the presentable images and link them to
79     // the images in the swapchain
80     VkResult result =
81     vkGetSwapchainImagesKHR( device,          // device
82                             // device associated with swapchain
83                             *outSwapChain,    // swapchain
84                             // swapchain to query
85                             &imageCount,     // pSwapchainImageCount
86                             // pointer to an integer related to the number of format pairs available
87                             *outPresentImages // pSwapchainImages
88                             // either NULL or a pointer to an array of VkSwapchainImageKHR structures
89
90     DBG_ASSERT_VULKAN_MSG( result,
91         "Failed to create swap-chain images" );
92 }
93
94 {
95     // You have 2 for double buffering
96     *outPresentImageViews = new VkImageView[2];
97     for( uint32_t i = 0; i < 2; ++i )
98     {
99         // create VkImageViews for your swap chain
100         // VkImages buffers:
101         VkImageViewCreateInfo ivci = {};
102         ivci.sType                  = VK_STRUCTURE_TYPE_IMAGE_VIEW_CREATE_INFO;
103         ivci.viewType               = VK_IMAGE_VIEW_TYPE_2D;
104         ivci.format                 = VK_FORMAT_B8G8R8A8_UNORM;
105         ivci.components.r           = VK_COMPONENT_SWIZZLE_R;
106         ivci.components.g           = VK_COMPONENT_SWIZZLE_G;
107         ivci.components.b           = VK_COMPONENT_SWIZZLE_B;
108         ivci.components.a           = VK_COMPONENT_SWIZZLE_A;
109         ivci.subresourceRange.aspectMask = VK_IMAGE_ASPECT_COLOR_BIT;
110         ivci.subresourceRange.baseMipLevel = 0;
111         ivci.subresourceRange.levelCount = 1;
112         ivci.subresourceRange.baseArrayLayer = 0;
113

```

```

114   ivci.subresourceRange.layerCount    = 1;
115   ivci.image                          = (*outPresentImages)[i];
116
117   // Create an image view from an existing image
118   VkResult result =
119   vkCreateImageView( device,              // device
120                     // logical device that creates the image view
121                     &ivci,              // pCreateInfo
122                     // pointer to instance of the VkImageViewCreateInfo structure containing parameters for the image view
123                     NULL,              // pAllocator
124                     // optional controls host memory allocation
125                     &(*outPresentImageViews)[i] ); // pView
126                     // pointer to VkImageView handle for returned image view object
127
128   DBG_ASSERT_VULKAN_MSG( result,
129       "Could not create ImageView." );
130   } // End for i
131   }
132 } // End SetupSwapChain(..)

```

Looking at Listing 6.6, you'll see the implementation specifics for configuring and setting up your swapchain:

- A vkGetPhysicalDeviceSurfaceCapabilitiesKHR
- B vkCreateSwapchainKHR
- C vkGetSwapchainImagesKHR is called twice as you'll want to double buffer the swap chain (front and back buffer)
- E vkCreateImageView

6.5 (Step 5) FrameBuffer & Render-Pass

The framebuffer in Vulkan is simpler than previous traditional OpenGL implementations. In Vulkan you have a 'Bag' or 'Repository' of resource views. The render-pass defines the role of framebuffer resources. Importantly, you can have more than one pass with each pass defining which framebuffer resource to use. While the render-pass might seem like additional work, as you start to generate more complex scenes, the render-pass gives you additional screen control. For example, post-processing and deferred rendering (e.g., mapping specific regions or order of processing to different threads/GPUs). The listing below sets a basic fullscreen render-pass (i.e., one display update with no sub-passes). With reference to the command-buffer (in the next section), you can use the command-buffer for several render-passes. You can also use a single command-buffer to draw a whole frames with the multiple passes contributing to techniques like shadow mapping and post-processing (managing these process more efficiently).

```

1  // Step -5-
2  void SetupRenderPass(VkDevice      device,
3                      VkPhysicalDevice physicalDevice,
4                      int            width,
5                      int            height,
6                      VkImageView*   presentImageViews,
7                      VkRenderPass*  outRenderPass,
8                      VkFramebuffer* outFrameBuffers)
9  {
10     // Frame buffer
11     // define your attachment points
12
13     #ifndef DEPTH_BUFFER
14         // Extension (Depth Buffer)
15         VkImage      depthImage      = NULL;
16         VkImageView   depthImageView = NULL;
17
18         {
19             // create a depth image:
20             VkImageCreateInfo imageCreateInfo = {};
21             imageCreateInfo.sType             = VK_STRUCTURE_TYPE_IMAGE_CREATE_INFO;
22             imageCreateInfo.imageType         = VK_IMAGE_TYPE_2D;
23             imageCreateInfo.format            = VK_FORMAT_D16_UNORM;
24             VkExtent3D ef = { width, height, 1 };
25             imageCreateInfo.extent            = ef;
26             imageCreateInfo.mipLevels        = 1;
27             imageCreateInfo.arrayLayers      = 1;
28             imageCreateInfo.samples          = VK_SAMPLE_COUNT_1_BIT;
29             imageCreateInfo.tiling           = VK_IMAGE_TILING_OPTIMAL;
30             imageCreateInfo.usage            = VK_IMAGE_USAGE_DEPTH_STENCIL_ATTACHMENT_BIT;
31             imageCreateInfo.sharingMode      = VK_SHARING_MODE_EXCLUSIVE;
32             imageCreateInfo.queueFamilyIndexCount = 0;
33             imageCreateInfo.pQueueFamilyIndices = NULL;
34             imageCreateInfo.initialLayout     = VK_IMAGE_LAYOUT_UNDEFINED;
35
36             // Create a new image object for your depth buffer
37             VkResult result =
38                 vkCreateImage( device,           // device
39                               // logical device that creates the image
40                               &imageCreateInfo, // pCreateInfo
41                               // pointer to VkImageCreateInfo structure with parameters for the created image
42                               NULL,             // pAllocator
43                               // optional control host memory allocation
44                               &depthImage );    // pImage
45             // pointer to VkImage handle returned image object
46
47             DBG_ASSERT_VULKAN_MSG( result,
48                                   "Failed to create depth image." );
49
50             VkMemoryRequirements memoryRequirements = {};
51             vkGetImageMemoryRequirements( device,           // device
52                                           // logical device that owns the image
53                                           depthImage,       // image
54                                           // image to query
55                                           &memoryRequirements ); // pMemoryRequirements
56             // instance pointer to VkMemoryRequirements structure returned memory requirements
57
58             // memoryRequirements contains memoryTypeBits member which is a bitmask - each one of the
59             // bits is set for every supported memory type for the resource. Bit i is set if and only
60             // if the memory type i in the VkPhysicalDeviceMemoryProperties structure for the physical

```

```

61 // device is supported for the resource.
62
63 // Allocate memory for your depth buffer
64 VkMemoryAllocateInfo imageAllocateInfo = {};
65 imageAllocateInfo.sType = VK_STRUCTURE_TYPE_MEMORY_ALLOCATE_INFO;
66 imageAllocateInfo.allocationSize = memoryRequirements.size;
67
68 // memoryTypeBits is a bitfield where if bit i is set, it means that
69 // the VkMemoryType i of the VkPhysicalDeviceMemoryProperties structure
70 // satisfies the memory requirements:
71 // read the device memory properties
72 VkPhysicalDeviceMemoryProperties memoryProperties;
73 vkGetPhysicalDeviceMemoryProperties( physicalDevice,
74                                     // handle to the device to query.
75                                     &memoryProperties );
76                                     // returned pointer to instance of VkPhysicalDeviceMemoryProperties structure
77
78 uint32_t memoryTypeBits = memoryRequirements.memoryTypeBits;
79 for( uint32_t i = 0; i < VK_MAX_MEMORY_TYPES; ++i )
80 {
81     VkMemoryType memoryType = memoryProperties.memoryTypes[i];
82     if( memoryTypeBits & 1 )
83     {
84         if( ( memoryType.propertyFlags & VK_MEMORY_PROPERTY_DEVICE_LOCAL_BIT ) )
85         {
86             // save index
87             imageAllocateInfo.memoryTypeIndex = i;
88             break;
89         }
90     }
91     memoryTypeBits = memoryTypeBits >> 1;
92 } // End for i
93
94 VkDeviceMemory imageMemory = { 0 };
95 result = vkAllocateMemory( device,
96                            // logical device that owns the memory
97                            &imageAllocateInfo,
98                            // pointer to VkMemoryAllocateInfo structure describing parameters of the allocation
99                            NULL,
100                            // optional control host memory allocation
101                            &imageMemory );
102                            // pointer to returned VkDeviceMemory handle with information about the allocated memory
103
104 DBG_ASSERT_VULKAN_MSG( result, "Failed to allocate device memory." );
105
106 result = vkBindImageMemory( device,
107                             // logical device that owns the image and memory
108                             depthImage,
109                             // image to bind
110                             imageMemory, 0 );
111                             // start offset of the region of memory which is to be bound to the image
112
113 DBG_ASSERT_VULKAN_MSG( result, "Failed to bind image memory." );
114
115 // create the depth image view:
116 VkImageAspectFlags aspectMask = VK_IMAGE_ASPECT_DEPTH_BIT;
117 VkImageViewCreateInfo imageViewCreateInfo = {};
118 imageViewCreateInfo.sType = VK_STRUCTURE_TYPE_IMAGE_VIEW_CREATE_INFO;
119 imageViewCreateInfo.image = depthImage;
120 imageViewCreateInfo.viewType = VK_IMAGE_VIEW_TYPE_2D;

```

```

121 imageViewCreateInfo.format                = imageCreateInfo.format;
122 VkComponentMapping g                      = { VK_COMPONENT_SWIZZLE_IDENTITY,
123                                               VK_COMPONENT_SWIZZLE_IDENTITY,
124                                               VK_COMPONENT_SWIZZLE_IDENTITY,
125                                               VK_COMPONENT_SWIZZLE_IDENTITY };
126 imageViewCreateInfo.components            = g;
127 imageViewCreateInfo.subresourceRange.aspectMask = aspectMask;
128 imageViewCreateInfo.subresourceRange.baseMipLevel = 0;
129 imageViewCreateInfo.subresourceRange.levelCount = 1;
130 imageViewCreateInfo.subresourceRange.baseArrayLayer = 0;
131 imageViewCreateInfo.subresourceRange.layerCount = 1;
132 result =
133     vkCreateImageView( device,
134                       // logical device that creates the image view
135                       &imageViewCreateInfo,
136                       // pointer to instance of the VkImageViewCreateInfo structure containing parameters for created image view
137                       NULL,
138                       // optional control host memory allocation
139                       &depthImageView );
140                       // pointer to returned VkImageView handle object
141
142     DBG_ASSERT_VULKAN_MSG( result,
143                           "Failed to create image view." );
144 }
145 #endif // DEPTH_BUFFER
146
147 // 0 - color screen buffer
148 VkAttachmentDescription pass[2] = { };
149 pass[0].format                = VK_FORMAT_B8G8R8A8_UNORM;
150 pass[0].samples               = VK_SAMPLE_COUNT_1_BIT;
151 pass[0].loadOp                = VK_ATTACHMENT_LOAD_OP_CLEAR;
152 pass[0].storeOp               = VK_ATTACHMENT_STORE_OP_STORE;
153 pass[0].stencilLoadOp         = VK_ATTACHMENT_LOAD_OP_DONT_CARE;
154 pass[0].stencilStoreOp        = VK_ATTACHMENT_STORE_OP_DONT_CARE;
155 pass[0].initialLayout         = VK_IMAGE_LAYOUT_COLOR_ATTACHMENT_OPTIMAL;
156 pass[0].finalLayout           = VK_IMAGE_LAYOUT_COLOR_ATTACHMENT_OPTIMAL;
157
158 VkAttachmentReference car      = {};
159 car.attachment                 = 0;
160 car.layout                     = VK_IMAGE_LAYOUT_COLOR_ATTACHMENT_OPTIMAL;
161
162 // 1 - depth buffer
163 pass[1].format                = VK_FORMAT_D16_UNORM;
164 pass[1].samples               = VK_SAMPLE_COUNT_1_BIT;
165 pass[1].loadOp                = VK_ATTACHMENT_LOAD_OP_CLEAR;
166 pass[1].storeOp               = VK_ATTACHMENT_STORE_OP_DONT_CARE;
167 pass[1].stencilLoadOp         = VK_ATTACHMENT_LOAD_OP_DONT_CARE;
168 pass[1].stencilStoreOp        = VK_ATTACHMENT_STORE_OP_DONT_CARE;
169 pass[1].initialLayout         = VK_IMAGE_LAYOUT_DEPTH_STENCIL_ATTACHMENT_OPTIMAL;
170 pass[1].finalLayout           = VK_IMAGE_LAYOUT_DEPTH_STENCIL_ATTACHMENT_OPTIMAL;
171
172 // create the one main subpass of your renderpass:
173 VkSubpassDescription subpass = {};
174 subpass.pipelineBindPoint     = VK_PIPELINE_BIND_POINT_GRAPHICS;
175 subpass.colorAttachmentCount  = 1;
176 subpass.pColorAttachments     = &car;
177 subpass.pDepthStencilAttachment = NULL;
178
179 #ifdef DEPTH_BUFFER
180 VkAttachmentReference dar      = {};
181 dar.attachment                 = 1;

```



```

242         // pointer to VkFramebufferCreateInfo structure describing framebuffer creation
243         NULL, // pAllocator
244         // optional control of host memory allocation
245         &(*outFrameBuffers)[i]); // pFramebuffer
246         // pointer to returned VkFramebuffer handle for the framebuffer object
247
248         DBG_ASSERT_VULKAN_MSG( result,
249             "Failed to create framebuffer.");
250     } // End for i
251 } // End SetupRenderPass(..)

```

Looking at Listing 6.5, you'll see the implementation specifics for configuring and setting up your framebuffer and render-pass:

- A vkCreateImage
- B vkGetImageMemoryRequirements
- C vkGetPhysicalDeviceMemoryProperties
- D vkAllocateMemory
- E vkBindImageMemory
- F vkCreateImageView
- G vkCreateRenderPass
- H vkCreateFramebuffer

6.6 (Step 6) Command-Buffers

Vulkan Rendering is done through Command-Buffers. The Command-Buffers are allocated from Command-Pools. Typically you have a Command-Pools associated with each thread and only use this thread when you write to the Command-Buffers allocated from its Command-Pool. This is because, it would be inefficient to externally synchronize access between the Command-Buffers and the Command-Pools (i.e., added overhead). Each Command-Buffer can be created either for one shot case or for multiple frames/submissions. Cannot call Command-Buffers from GPU (command-lists can). The API commands for filling the Command-Buffer begin with 'vkCmd'..() and need to be done between a 'Begin' and 'End'. Importantly, the Command-Buffer mechanism is designed to be multi-threading friendly. The 'primary' Command-Buffer can call many secondary Command-Buffers.

Listing 6.7: Command-Buffers are crucial elements for controlling the rendering.

```

1 // Step -6-
2 void SetupCommandBuffer(VkDevice device,
3     VkPhysicalDevice physicalDevice,
4     VkCommandBuffer* outCommandBuffer)
5 {
6     // Give your device some commands (orders)
7     {

```