



SIGGRAPH THINK BEYOND



Computer Graphics

Introduction to the Vulkan Computer Graphics API

Mike Bailey
mjb@cs.oregonstate.edu



This work is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](https://creativecommons.org/licenses/by-nc-nd/4.0/).

<http://cs.oregonstate.edu/~mjb/vulkan>

SIGGRAPH THINK BEYOND

FULL PAGE mjb - July 24, 2020

Course Goals

- Give a sense of how Vulkan is different from OpenGL
- Show how to do basic drawing in Vulkan
- Leave you with working, documented sample code

<http://cs.oregonstate.edu/~mjb/vulkan>


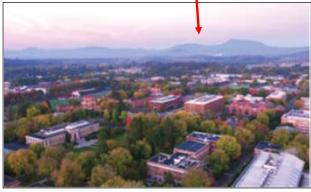
SIGGRAPH THINK BEYOND

mjb - July 24, 2020

Mike Bailey

- Professor of Computer Science, Oregon State University
- Has been in computer graphics for over 30 years
- Has had over 8,000 students in his university classes
- mjb@cs.oregonstate.edu

Welcomel I'm happy to be here. I hope you are too!

<http://cs.oregonstate.edu/~mjb/vulkan>

SIGGRAPH THINK BEYOND

mjb - July 24, 2020

Sections

1. Introduction	13. Swap Chain
2. Sample Code	14. Push Constants
3. Drawing	15. Physical Devices
4. Shaders and SPIR-V	16. Logical Devices
5. Dats Buffers	17. Dynamic State Variables
6. GLFW	18. Getting Information Back
7. GLM	19. Compute Shaders
8. Instancing	20. Specialization Constants
9. Graphics Pipeline Data Structure	21. Synchronization
10. Descriptor Sets	22. Pipeline Barriers
11. Textures	23. Multisampling
12. Queues and Command Buffers	24. Multipass
	25. Ray Tracing

SIGGRAPH THINK BEYOND

mjb - July 24, 2020

My Favorite Vulkan Reference



Graham Sellers, *Vulkan Programming Guide*, Addison-Wesley, 2017.

SIGGRAPH THINK BEYOND

mjb - July 24, 2020



Introduction

Mike Bailey
mjb@cs.oregonstate.edu

<http://cs.oregonstate.edu/~mjb/vulkan>

SIGGRAPH THINK BEYOND

mjb - July 24, 2020

Acknowledgements



First of all, thanks to the inaugural class of 19 students who braved new, unrefined, and just-in-time course materials to take the first Vulkan class at Oregon State University – Winter Quarter, 2018. Thanks for your courage and patience!



Second, thanks to NVIDIA for all of their support!



Third, thanks to the Khronos Group for the great laminated Vulkan Quick Reference Cards! (Look at those happy faces in the photo holding them.)



Ali Alsalehy	Alan Needs
Natasha Anisimova	Raja Petroff
Jianchang Bi	Bei Rong
Christopher Cooper	Lawrence Roy
Richard Cunard	Lily Shellhammer
Braxton Cuneo	Hannah Solorzano
Benjamin Fields	Jian Tang
Trevor Hammock	Glenn Uphagrove
Zach Lerew	Logan Wingard
Victor Li	

SIGGRAPH THE GPU SHOW 7

History of Shaders

- 2004: OpenGL 2.0 / GLSL 1.10 includes Vertex and Fragment Shaders
- 2008: OpenGL 3.0 / GLSL 1.30 adds features left out before
- 2010: OpenGL 3.3 / GLSL 3.30 adds Geometry Shaders
- 2010: OpenGL 4.0 / GLSL 4.00 adds Tessellation Shaders
- 2012: OpenGL 4.3 / GLSL 4.30 adds Compute Shaders
- 2017: OpenGL 4.6 / GLSL 4.60

There is lots more detail at:
https://www.khronos.org/opengl/wiki/History_of_OpenGL

SIGGRAPH THE GPU SHOW 8

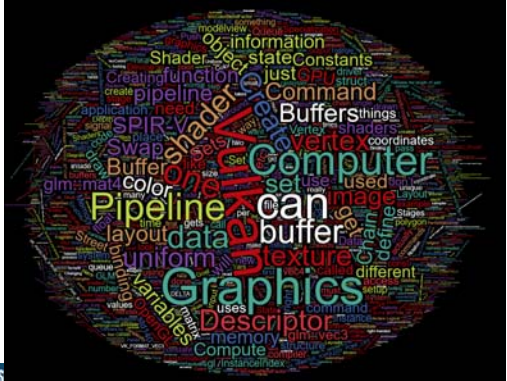
History of Shaders

- 2014: Khronos starts Vulkan effort
- 2016: Vulkan 1.0
- 2016: Vulkan 1.1
- 2020: Vulkan 1.2

There is lots more detail at:
[https://en.wikipedia.org/wiki/Vulkan_\(API\)](https://en.wikipedia.org/wiki/Vulkan_(API))

SIGGRAPH THE GPU SHOW 9

Everything You Need to Know is Right Here ... Somewhere ☺



SIGGRAPH THE GPU SHOW 10

Top Three Reasons that Prompted the Development of Vulkan

1. Performance
2. Performance
3. Performance

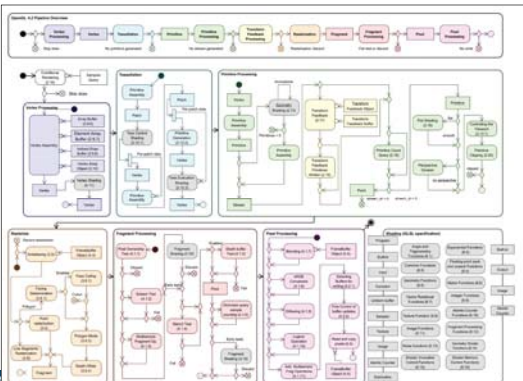
Vulkan is better at keeping the GPU busy than OpenGL is. OpenGL drivers need to do a lot of CPU work before handing work off to the GPU. Vulkan lets you get more power from the GPU card you already have.

This is especially important if you can hide the complexity of Vulkan from your customer base and just let them see the improved performance. Thus, Vulkan has had a lot of support and interest from game engine developers, 3rd party software vendors, etc.

As an aside, the Vulkan development effort was originally called "glNext", which created the false impression that this was a replacement for OpenGL. It's not.

SIGGRAPH THE GPU SHOW 11

OpenGL 4.2 Pipeline Flowchart



SIGGRAPH THE GPU SHOW 12

Why is it so important to keep the GPU Busy?


	Titan V	Titan V100	Titan P100	GTX 1080 S	GTX 1080
GPU	GV100	GV100	GP100 Col. Exec. Pascal	GP104 Pascal	GP104 4th Pascal
Transistor Count	24.1B	24.1B	6.3B	5.3B	7.2B
Full Precision	13.6B/17B	13.6B/17B	16.6B/17B	16.6B/17B	16.6B/17B
CUDA Cores / Tensor Cores	5200/640	5200/640	3684/18	3072/48	2560/0
Tensor	320	320	224	224	160
ROPs	7	7	96/175	64	64
Core Clock	1320MHz	1320MHz	1320MHz	1605MHz	1605MHz
Boost Clock	1400MHz	1375MHz	1400MHz	1605MHz	1605MHz
FP32 TR/GPs	14.1TB/GPs	14.1TB/GPs	10.8TB/GPs	11.4TB/GPs	10.1TB/GPs
Memory Type	HBM2	HBM2	GDDR5X	GDDR5X	GDDR5X
Memory Capacity	12GB	16GB	12GB	6GB	6GB
Memory Clock	1.77Gbps HBM2	1.77Gbps HBM2	710Gbps	500Gbps	GDDR5X
Max Power	357W	400W	400W	175W	175W
Memory Bandwidth	800GB/s	800GB/s	716GB/s	200GB/s	200GB/s
Total Power Budget (TDP)	350W	350W	300W	200W	200W
Power Connectors	1x 8-pin 1x 6-pin	1x 8-pin 1x 6-pin	1x 8-pin 1x 6-pin	1x 8-pin 1x 6-pin	1x 8-pin 1x 6-pin
Release Date	12/07/2017	02/06/2017	1/05	03/27/2013	03/27/2013
Release Price	\$3000	\$10000	\$700	\$199	\$199

The Nvidia Titan V graphics card is not targeted at gamers, but rather at scientific and machine-learning applications. That does not, however, mean that the card is incapable of gaming, but does mean that we can't anticipate future key performance metrics for it. The Titan V is a replacement for the earlier consumer GTX1080 GPU, part of the Pascal architectural card series. The key difference is that the Titan V offers 4x more FP32 TR/GPs, whereas the Titan V100 was available as part of a \$10,000 developer kit. The Titan V100 still offers greater memory capacity to 16GB, 160GB/s memory bandwidth, and three smaller memory profiles, but other card features remain mostly the same. Core clock, for one, is 1320 MHz across each GPU, with 1400 boost across board for the full board. Source: Nvidia.

Who was the original Vulcan?

From Wikipedia:

"Vulcan is the god of fire including the fire of volcanoes, metalworking, and the forge in ancient Roman religion and myth. Vulcan is often depicted with a blacksmith's hammer. The **Vulcanalia** was the annual festival held August 23 in his honor. His Greek counterpart is Hephaestus, the god of fire and smelting. In Etruscan religion, he is identified with Sethlans. Vulcan belongs to the most ancient stage of Roman religion: Varro, the ancient Roman scholar and writer, citing the *Annales Maximi*, records that king Titus Tatius dedicated altars to a series of deities among which Vulcan is mentioned."



[https://en.wikipedia.org/wiki/Vulcan_\(mythology\)](https://en.wikipedia.org/wiki/Vulcan_(mythology))

Why Name it after the God of the Forge?



Who is the Khronos Group?

The Khronos Group, Inc. is a non-profit member-funded industry consortium, focused on the creation of open standard, royalty-free application programming interfaces (APIs) for authoring and accelerated playback of dynamic media on a wide variety of platforms and devices. Khronos members may contribute to the development of Khronos API specifications, vote at various stages before public deployment, and accelerate delivery of their platforms and applications through early access to specification drafts and conformance tests.



Playing "Where's Waldo?" with Khronos Membership




Who's Been Specifically Working on Vulkan?



Vulkan 19


- Originally derived from AMD's *Mantle* API
- Also heavily influenced by Apple's *Metal* API and Microsoft's *DirectX 12*
- Goal: much less driver complexity and overhead than OpenGL has
- Goal: much less user hand-holding
- Goal: higher single-threaded performance than OpenGL can deliver
- Goal: able to do multithreaded graphics
- Goal: able to handle tiled rendering



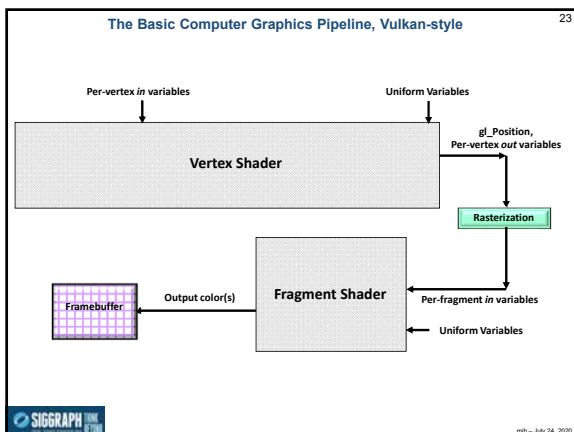
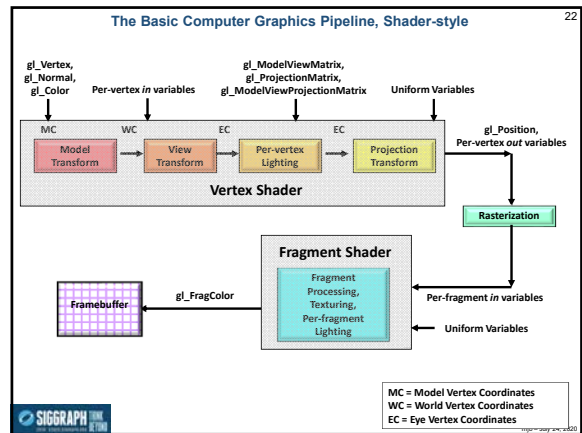
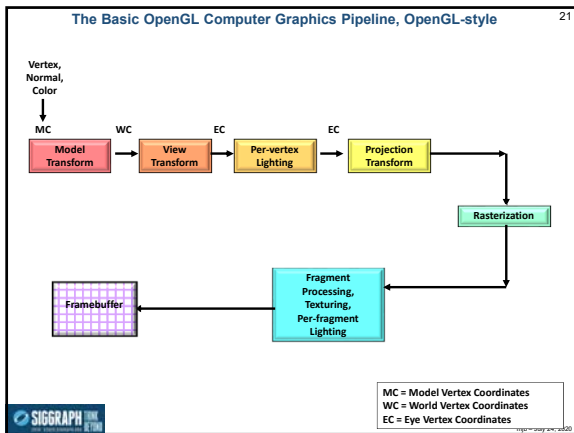
mb - July 24, 2020

Vulkan Differences from OpenGL 20

- More low-level information must be provided (by you!) in the application, rather than the driver
- Screen coordinate system is Y-down
- No "current state", at least not one maintained by the driver
- All of the things that we have talked about being *deprecated* in OpenGL are *really deprecated* in Vulkan: built-in pipeline transformations, begin-end, fixed-function, etc.
- You must manage your own transformations.
- All transformation, color and texture functionality must be done in shaders.
- Shaders are pre-"half-compiled" outside of your application. The compilation process is then finished during the runtime pipeline-building process.




mb - July 24, 2020



Moving part of the driver into the application 24

<p>Complex drivers lead to driver overhead and cross vendor unpredictability</p> <p>Error management is always active</p> <p>Driver processes full shading language source</p> <p>Separate APIs for desktop and mobile markets</p>	<p>Application</p> <p>Traditional graphics drivers include significant context, memory and error management</p> <p>GPU</p>	<p>Vulkan</p> <p>Application responsible for memory allocation and thread management to generate command buffers</p> <p>Direct GPU Control</p> <p>GPU</p> <p style="text-align: center; font-size: x-small;">Khronos Group</p>	<p>Simpler drivers for low-overhead efficiency and cross vendor portability</p> <p>Layered architecture so validation and debug layers can be unloaded when not needed</p> <p>Run-time only has to ingest SPIR-V intermediate language</p> <p>Unified API for mobile, desktop, console and embedded platforms</p>
--	--	--	---



mb - July 24, 2020

Vulkan Highlights: Command Buffers

- Graphics commands are sent to command buffers
- E.g., `vkCmdDoSomething(cmdBuffer, ...);`
- You can have as many simultaneous Command Buffers as you want
- Buffers are flushed to Queues when the application wants them to be flushed
- Each command buffer can be filled from a different thread

CPU Thread [] [] [] [] **Cmd buffer**

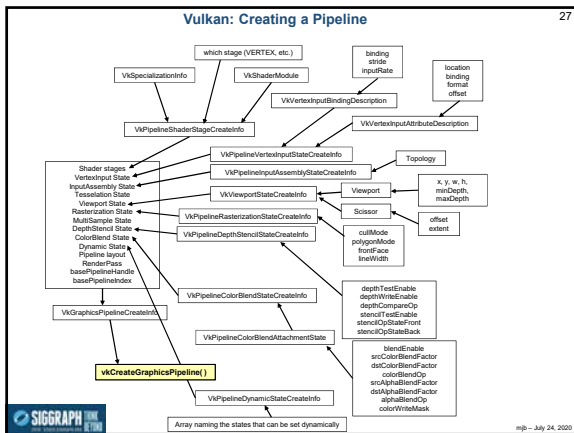
CPU Thread [] [] [] [] **Cmd buffer**

CPU Thread [] [] [] [] **Cmd buffer**

CPU Thread [] [] [] [] **Cmd buffer**

Vulkan Highlights: Pipeline State Objects

- In OpenGL, your "pipeline state" is the combination of whatever your current graphics attributes are: color, transformations, textures, shaders, etc.
- Changing the state on-the-fly one item at-a-time is very expensive
- Vulkan forces you to set all your state variables at once into a "pipeline state object" (PSO) data structure and then invoke the entire PSO *at once* whenever you want to use that state combination
- Think of the pipeline state as being immutable.
- Potentially, you could have thousands of these pre-prepared pipeline state objects



Querying the Number of Something

```
uint32_t count;
result = vkEnumeratePhysicalDevices( Instance, OUT &count, OUT (VkPhysicalDevice *)nullptr );
```

VkPhysicalDevice * physicalDevices = new VkPhysicalDevice[count];
 result = vkEnumeratePhysicalDevices(Instance, OUT &count, OUT physicalDevices);

This way of querying information is a recurring OpenCL and Vulkan pattern (get used to it):

```
result = vkEnumeratePhysicalDevices( Instance, &count, nullptr );
```

How many total there are | Where to put them

```
result = vkEnumeratePhysicalDevices( Instance, &count, physicalDevices );
```

Vulkan Code has a Distinct "Style" of Setting Information in structs and then Passing that Information as a pointer-to-the-struct

```
VkBufferCreateInfo vbci;
vbci.sType = VK_STRUCTURE_TYPE_BUFFER_CREATE_INFO;
vbci.pNext = nullptr;
vbci.flags = 0;
vbci.size = << buffer size in bytes >>
vbci.usage = VK_USAGE_UNIFORM_BUFFER_BIT;
vbci.sharingMode = VK_SHARING_MODE_EXCLUSIVE;
vbci.queueFamilyIndexCount = 0;
vbci.pQueueFamilyIndices = nullptr;

VK_RESULT result = vkCreateBuffer ( LogicalDevice, IN &vbci, PALLOCATOR, OUT &Buffer );

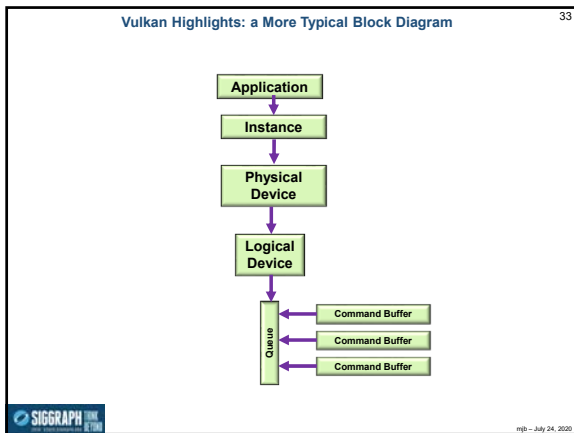
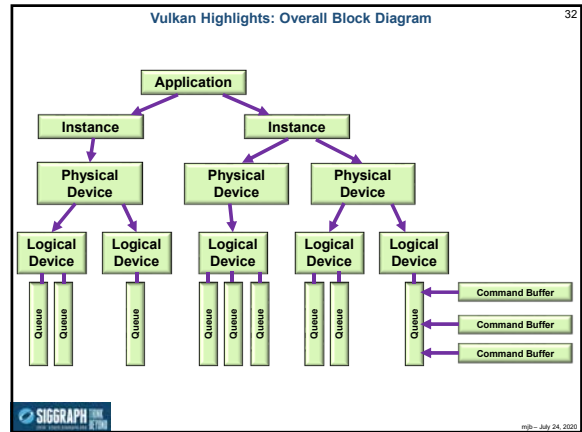
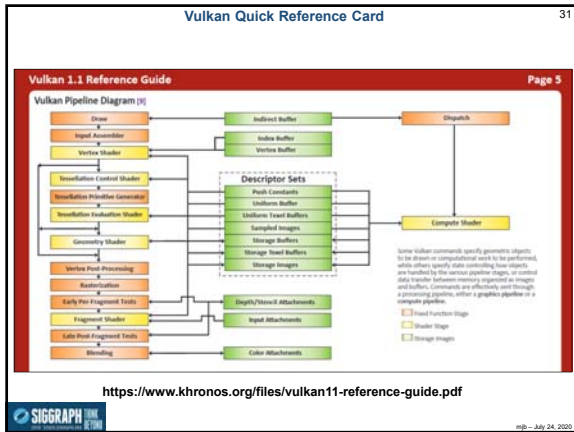
VkMemoryRequirements vmr;
result = vkGetBufferMemoryRequirements( LogicalDevice, Buffer, OUT &vmr ); // fills vmr

VkMemoryAllocateInfo vmai;
vmai.sType = VK_STRUCTURE_TYPE_MEMORY_ALLOCATE_INFO;
vmai.pNext = nullptr;
vmai.flags = 0;
vmai.allocationSize = vmr.size;
vmai.memoryTypeIndex = 0;

result = vkAllocateMemory( LogicalDevice, IN &vmai, PALLOCATOR, OUT &MatrixBufferMemoryHandle );
result = vkBindBufferMemory( LogicalDevice, Buffer, MatrixBufferMemoryHandle, 0 );
```

Vulkan Quick Reference Card – I Recommend you Print This!

<https://www.khronos.org/files/vulkan11-reference-guide.pdf>




- ### Steps in Creating Graphics using Vulkan
1. Create the Vulkan Instance
 2. Setup the Debug Callbacks
 3. Create the Surface
 4. List the Physical Devices
 5. Pick the right Physical Device
 6. Create the Logical Device
 7. Create the Uniform Variable Buffers
 8. Create the Vertex Data Buffers
 9. Create the texture sampler
 10. Create the texture images
 11. Create the Swap Chain
 12. Create the Depth and Stencil Images
 13. Create the RenderPass
 14. Create the Framebuffer(s)
 15. Create the Descriptor Set Pool
 16. Create the Command Buffer Pool
 17. Create the Command Buffer(s)
 18. Read the shaders
 19. Create the Descriptor Set Layouts
 20. Create and populate the Descriptor Sets
 21. Create the Graphics Pipeline(s)
 22. Update-Render-Update-Render- ...
- SIGGRAPH 2019 mb - July 24, 2020

- ### Vulkan GPU Memory
- Your application allocates GPU memory for the objects it needs
 - To write and read that GPU memory, you map that memory to the CPU address space
 - Your application is responsible for making sure that what you put into that memory is actually in the right format, is the right size, has the right alignment, etc.
- SIGGRAPH 2019 mb - July 24, 2020

- ### Vulkan Render Passes
- Drawing is done inside a render pass
 - Each render pass contains what framebuffer attachments to use
 - Each render pass is told what to do when it begins and ends
- SIGGRAPH 2019 mb - July 24, 2020

Vulkan Compute Shaders


- Compute pipelines are allowed, but they are treated as something special (just like OpenGL treats them)
- Compute passes are launched through dispatches
- Compute command buffers can be run asynchronously



mjb - July 24, 2020

Vulkan Synchronization

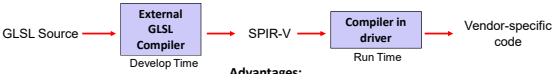
- Synchronization is the responsibility of the application
- Events can be set, polled, and waited for (much like OpenGL)
- Vulkan itself does not ever lock – that’s your application’s job
- Threads can concurrently read from the same object
- Threads can concurrently write to different objects



mjb - July 24, 2020


Vulkan Shaders

- GLSL is the same as before ... almost
- For places it’s not, an implied `#define VULKAN 100` is automatically supplied by the compiler
- You pre-compile your shaders with an external compiler
- Your shaders get turned into an intermediate form known as SPIR-V (Standard Portable Intermediate Representation for Vulkan)
- SPIR-V gets turned into fully-compiled code at runtime
- The SPIR-V spec has been public for years –new shader languages are surely being developed
- OpenCL and OpenGL have adopted SPIR-V as well




Advantages:

1. Software vendors don’t need to ship their shader source
2. Software can launch faster because half of the compilation has already taken place
3. This guarantees a common front-end syntax
4. This allows for other language front-ends




mjb - July 24, 2020


Your Sample2019.zip File Contains This



The “19” refers to the version of Visual Studio, not the year of development.

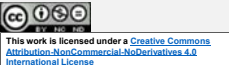


mjb - July 24, 2020




The Vulkan Sample Code Included with These Notes

Mike Bailey
mjb@cs.oregonstate.edu

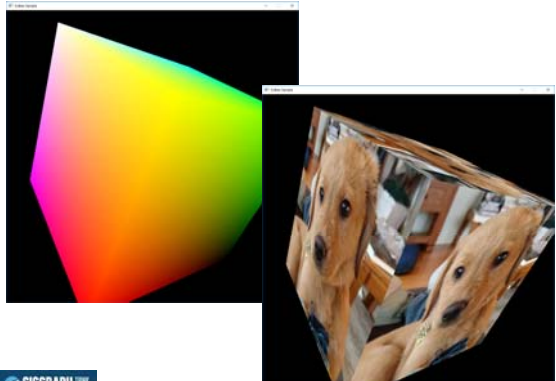



<http://cs.oregonstate.edu/~mjb/vulkan>



mjb - July 24, 2020

Sample Program Output

mjb - July 24, 2020

Sample Program Keyboard Inputs

'l', 'L':	Toggle lighting off and on
'm', 'M':	Toggle display mode (textures vs. colors, for now)
'p', 'P':	Pause the animation
'q', 'Q':	quit the program
Esc:	quit the program
'r', 'R':	Toggle rotation-animation and using the mouse
't', 'T':	Toggle using a vertex buffer only vs. an index buffer (in the index buffer version)
'1', '4', '9'	Set the number of instances (in the instancing version)

Caveats on the Sample Code, I

1. I've written everything out in appalling longhand.
2. Everything is in one .cpp file (except the geometry data). It really should be broken up, but this way you can find everything easily.
3. At times, I could have hidden complexity, but I didn't. At all stages, I have tried to err on the side of showing you *everything*, so that nothing happens in a way that's kept a secret from you.
4. I've setup Vulkan structs every time they are used, even though, in many cases (most?), they could have been setup once and then re-used each time.
5. At times, I've setup things that didn't need to be setup just to show you what could go there.

Caveats on the Sample Code, II

6. There are great uses for C++ classes and methods here to hide some complexity, but I've not done that.
7. I've typedef'ed a couple things to make the Vulkan phraseology more consistent.
8. Even though it is not good software style, I have put persistent information in global variables, rather than a separate data structure
9. At times, I have copied lines from vulkan.h into the code as comments to show you what certain options could be.
10. I've divided functionality up into the pieces that make sense to me. Many other divisions are possible. Feel free to invent your own.

Main Program

```

int main( int argc, char * argv[] )
{
    Width = 800;
    Height = 600;

    armo_err = fopen_err( &FpDebug, DEBUGFILE, "w" );
    if( err != 0 )
    {
        fprintf( stderr, "Cannot open debug print file \"%s\n", DEBUGFILE );
        FpDebug = stderr;
    }
    fprintf( FpDebug, "FpDebug: Width = %d ; Height = %d\n", Width, Height );

    Reset();
    InitGraphics();

    // loop until the user closes the window:
    while( glfwWindowShouldClose( MainWindow ) == 0 )
    {
        glfwPollEvents();
        Time = glfwGetTime(); // elapsed time, in double-precision seconds
        UpdateScene();
        RenderScene();
    }

    fprintf( FpDebug, "Closing the GLFW window\n" );

    vkQueueWaitIdle( Queue );
    vkDeviceWaitIdle( LogicalDevice );
    DestroyAllVulkan();
    glfwDestroyWindow( MainWindow );
    glfwTerminate();
    return 0;
}
    
```

InitGraphics(), I

```

void InitGraphics()
{
    HERE_I_AM( "InitGraphics" );

    VkResult result = VK_SUCCESS;

    Init01Instance();
    InitGLFW();
    Init02CreateDebugCallbacks();

    Init03PhysicalDeviceAndGetQueueFamilyProperties();
    Init04LogicalDeviceAndQueue();

    Init05UniformBuffer( sizeof( Matrices ), &MyMatrixUniformBuffer );
    Fill05DataBuffer( MyMatrixUniformBuffer, (void *) &Matrices );

    Init05UniformBuffer( sizeof( Light ), &MyLightUniformBuffer );
    Fill05DataBuffer( MyLightUniformBuffer, (void *) &Light );

    Init05MyVertexDataBuffer( sizeof( VertexData ), &MyVertexDataBuffer );
    Fill05DataBuffer( MyVertexDataBuffer, (void *) VertexData );

    Init06CommandPool();
    Init06CommandBuffers();
}
    
```

InitGraphics(), II

```

Init07TextureSampler( &MyPuppyTexture.texSampler );
Init07TextureBufferAndFillFromBmpFile( "puppy.bmp", &MyPuppyTexture );

Init08Swapchain();

Init09DepthStencilImage();

Init10RenderPasses();

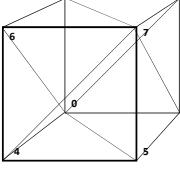
Init11Framebuffers();

Init12SpirvShader( "sample-vert.spv", &ShaderModuleVertex );
Init12SpirvShader( "sample-frag.spv", &ShaderModuleFragment );

Init13DescriptorSetPool();
Init13DescriptorSetLayouts();
Init13DescriptorSets();

Init14GraphicsVertexFragmentPipeline( ShaderModuleVertex, ShaderModuleFragment,
VK_PRIMITIVE_TOPOLOGY_TRIANGLE_LIST, &GraphicsPipeline );
}
    
```


A Colored Cube



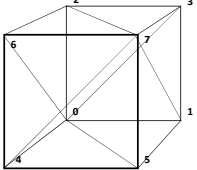
```
static GLfloat CubeColors[ ][3] =
{
    {0.0, 0.0, 0.0},
    {1.0, 0.0, 0.1},
    {0.0, 1.0, 0.1},
    {1.0, 1.0, 0.1},
    {0.0, 0.1},
    {1.0, 0.1, 1.1},
    {0.0, 1.1},
    {1.0, 1.1, 1.1}
};

static GLuint CubeTriangleIndices[ ][3] =
{
    {0, 2, 3},
    {0, 3, 1},
    {4, 5, 7},
    {4, 7, 6},
    {1, 3, 7},
    {1, 7, 5},
    {0, 4, 6},
    {0, 6, 2},
    {2, 6, 7},
    {2, 7, 3},
    {0, 1, 5},
    {0, 5, 4}
};

static GLuint CubeNormals[ ][3] =
{
    {-1, -1, -1},
    {-1, -1, 1},
    {-1, 1, -1},
    {-1, 1, 1},
    {1, -1, -1},
    {1, -1, 1},
    {1, 1, -1},
    {1, 1, 1},
    {1, 1, 1},
    {1, 1, 1},
    {1, 1, 1},
    {1, 1, 1}
};
```

SIGGRAPH 2018 July 24, 2020

A Colored Cube



```
struct vertex
{
    glm::vec3 position;
    glm::vec3 normal;
    glm::vec3 color;
    glm::vec2 texCoord;
};

struct vertex VertexData[ ] =
{
    // triangle 0-2-3:
    // vertex #0:
    {
        {-1, -1, -1},
        {0, 0, -1},
        {0, 0, 0},
        {1, 0}
    },
    // vertex #2:
    {
        {-1, 1, -1},
        {0, 0, -1},
        {0, 1, 0},
        {1, 1}
    },
    // vertex #3:
    {
        {1, 1, -1},
        {0, 0, -1},
        {1, 1, 0},
        {0, 1}
    }
};
```

SIGGRAPH 2018 July 24, 2020

The Vertex Data is in a Separate File

```
#include "SampleVertexData.cpp"

struct vertex
{
    glm::vec3 position;
    glm::vec3 normal;
    glm::vec3 color;
    glm::vec2 texCoord;
};

struct vertex VertexData[ ] =
{
    // triangle 0-2-3:
    // vertex #0:
    {
        {-1, -1, -1},
        {0, 0, -1},
        {0, 0, 0},
        {1, 0}
    },
    // vertex #2:
    {
        {-1, 1, -1},
        {0, 0, -1},
        {0, 1, 0},
        {1, 1}
    },
    ...
};
```

SIGGRAPH 2018 July 24, 2020

What if you don't need all of this information?

```
struct vertex
{
    glm::vec3 position;
    glm::vec3 normal;
    glm::vec3 color;
    glm::vec2 texCoord;
};
```

For example, what if you are not doing texturing in this application? Should you re-do this struct and leave the texCoord element out?

As best as I can tell, the only costs for retaining vertex attributes that you aren't going to use are some GPU memory space and possibly some inefficient uses of the cache, but not gross performance. So, I recommend keeping this struct intact, and, if you don't need texturing, simply don't use the texCoord values in your vertex shader.

SIGGRAPH 2018 July 24, 2020

Vulkan Software Philosophy

Vulkan has lots of typedefs that define C/C++ structs and enums

Vulkan takes a non-C++ object-oriented approach in that those typedef'd structs pass all the necessary information into a function. For example, where we might normally say in C++:

```
result = LogicalDevice->vkGetDeviceQueue ( queueFamilyIndex, queueIndex, OUT &Queue );
```

we would actually say in C:

```
result = vkGetDeviceQueue ( LogicalDevice, queueFamilyIndex, queueIndex, OUT &Queue );
```

SIGGRAPH 2018 July 24, 2020

Vulkan Conventions

VkXxx is a typedef, probably a struct

VkYyy() is a function call

VK_ZZZ is a constant

My Conventions

"Init" in a function call name means that something is being setup that only needs to be setup once

The number after "Init" gives you the ordering

In the source code, after main() comes InitGraphics(), then all of the InitxYYY() functions in numerical order. After that comes the helper functions

"Find" in a function call name means that something is being looked for

"Fill" in a function call name means that some data is being supplied to Vulkan

"IN" and "OUT" ahead of function call arguments are just there to let you know how an argument is going to be used by the function. Otherwise, IN and OUT have no significance. They are actually #define'd to nothing.

SIGGRAPH 2018 July 24, 2020

Querying the Number of Something and Allocating Enough Structures to Hold Them All

```
uint32_t count;
result = vkEnumeratePhysicalDevices( Instance, OUT &count, OUT (VkPhysicalDevice *)nullptr );

VkPhysicalDevice * physicalDevices = new VkPhysicalDevice[ count ];
result = vkEnumeratePhysicalDevices( Instance, OUT &count, OUT &physicalDevices[0] );
```

This way of querying information is a recurring OpenCL and Vulkan pattern (get used to it):

```
result = vkEnumeratePhysicalDevices( Instance, &count, nullptr );
result = vkEnumeratePhysicalDevices( Instance, &count, &physicalDevices[0] );
```

How many total there are Where to put them

mb - July 24, 2020

Your Sample2019.zip File Contains This

Linux shader compiler

Windows shader compiler

Double-click here to launch Visual Studio 2019 with this solution

The "19" refers to the version of Visual Studio, not the year of development.

mb - July 24, 2020

Reporting Error Results, I

```
struct errorcode
{
    VkResult resultCode;
    std::string meaning;
}
ErrorCodes[] =
{
    {VK_NOT_READY, "Not Ready" },
    {VK_TIMEOUT, "Timeout" },
    {VK_EVENT_SET, "Event Set" },
    {VK_EVENT_RESET, "Event Reset" },
    {VK_INCOMPLETE, "Incomplete" },
    {VK_ERROR_OUT_OF_HOST_MEMORY, "Out of Host Memory" },
    {VK_ERROR_OUT_OF_DEVICE_MEMORY, "Out of Device Memory" },
    {VK_ERROR_INITIALIZATION_FAILED, "Initialization Failed" },
    {VK_ERROR_DEVICE_LOST, "Device Lost" },
    {VK_ERROR_MEMORY_MAP_FAILED, "Memory Map Failed" },
    {VK_ERROR_LAYER_NOT_PRESENT, "Layer Not Present" },
    {VK_ERROR_EXTENSION_NOT_PRESENT, "Extension Not Present" },
    {VK_ERROR_FEATURE_NOT_PRESENT, "Feature Not Present" },
    {VK_ERROR_INCOMPATIBLE_DRIVER, "Incompatible Driver" },
    {VK_ERROR_TOO_MANY_OBJECTS, "Too Many Objects" },
    {VK_ERROR_FORMAT_NOT_SUPPORTED, "Format Not Supported" },
    {VK_ERROR_FRAGMENTED_POOL, "Fragmented Pool" },
    {VK_ERROR_SURFACE_LOST_KHR, "Surface Lost" },
    {VK_ERROR_NATIVE_WINDOW_IN_USE_KHR, "Native Window in Use" },
    {VK_SUBOPTIMAL_KHR, "Suboptimal" },
    {VK_ERROR_OUT_OF_DATE_KHR, "Error Out of Date" },
    {VK_ERROR_INCOMPATIBLE_DISPLAY_KHR, "Incompatible Display" },
    {VK_ERROR_VALIDATION_FAILED_EXT, "Validation Failed" },
    {VK_ERROR_INVALID_SHADER_NV, "Invalid Shader" },
    {VK_ERROR_OUT_OF_POOL_MEMORY_KHR, "Out of Pool Memory" },
    {VK_ERROR_INVALID_EXTERNAL_HANDLE, "Invalid External Handle" },
};
```

mb - July 24, 2020

Reporting Error Results, II

```
void PrintVkError( VkResult result, std::string prefix )
{
    if (Verbose && result == VK_SUCCESS)
    {
        fprintf(FpDebug, "%s: %s\n", prefix.c_str(), "Successful");
        fflush(FpDebug);
        return;
    }

    const int numErrorCodes = sizeof( ErrorCodes ) / sizeof( struct errorcode );
    std::string meaning = "";
    for( int i = 0; i < numErrorCodes; i++ )
    {
        if( result == ErrorCodes[i].resultCode )
        {
            meaning = ErrorCodes[i].meaning;
            break;
        }
    }

    fprintf( FpDebug, "%s: %s\n", prefix.c_str(), meaning.c_str() );
    fflush(FpDebug);
}
```

mb - July 24, 2020

Extras in the Code

```
#define REPORT(s) { PrintVkError( result, s ); fflush(FpDebug); }

#define HERE_I_AM(s) if ( Verbose ) { fprintf( FpDebug, "***** %s *****\n", s ); fflush(FpDebug); }

bool Paused;

bool Verbose;

#define DEBUGFILE "VulkanDebug.txt"
errno_t err = fopen_s( &FpDebug, DEBUGFILE, "w" );

const int32_t OFFSET_ZERO = 0;
```

mb - July 24, 2020

Drawing

Mike Bailey
mjb@cs.oregonstate.edu

This work is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](https://creativecommons.org/licenses/by-nc-nd/4.0/).

<http://cs.oregonstate.edu/~mjb/vulkan>

mb - July 24, 2020

Vulkan Topologies

VK_PRIMITIVE_TOPOLOGY_POINT_LIST

VK_PRIMITIVE_TOPOLOGY_TRIANGLE_LIST

VK_PRIMITIVE_TOPOLOGY_LINE_LIST

VK_PRIMITIVE_TOPOLOGY_TRIANGLE_STRIP

VK_PRIMITIVE_TOPOLOGY_LINE_STRIP

VK_PRIMITIVE_TOPOLOGY_TRIANGLE_FAN

SIGGRAPH 2018 | July 24, 2020

Vulkan Topologies

```

typedef enum VkPrimitiveTopology
{
    VK_PRIMITIVE_TOPOLOGY_POINT_LIST
VK_PRIMITIVE_TOPOLOGY_LINE_LIST
VK_PRIMITIVE_TOPOLOGY_LINE_STRIP
VK_PRIMITIVE_TOPOLOGY_TRIANGLE_LIST
VK_PRIMITIVE_TOPOLOGY_TRIANGLE_STRIP
VK_PRIMITIVE_TOPOLOGY_TRIANGLE_FAN
VK_PRIMITIVE_TOPOLOGY_LINE_LIST_WITH_ADJACENCY
VK_PRIMITIVE_TOPOLOGY_LINE_STRIP_WITH_ADJACENCY
VK_PRIMITIVE_TOPOLOGY_TRIANGLE_LIST_WITH_ADJACENCY
VK_PRIMITIVE_TOPOLOGY_TRIANGLE_STRIP_WITH_ADJACENCY
    VK_PRIMITIVE_TOPOLOGY_PATCH_LIST
} VkPrimitiveTopology;
    
```

SIGGRAPH 2018 | July 24, 2020

A Colored Cube Example

```

static const glm::vec3 CubeColors[8] =
{
    { 0.0, 0.0, 1.0 },
    { 1.0, 0.0, 1.0 },
    { 0.0, 1.0, 1.0 },
    { 1.0, 1.0, 1.0 },
    { 0.0, 0.0, 0.0 },
    { 0.0, 0.0, 1.0 },
    { 0.0, 1.0, 0.0 },
    { 1.0, 1.0, 0.0 }
};

static GLuint CubeTriangleIndices[36] =
{
    { 0, 2, 3 },
    { 0, 3, 1 },
    { 4, 5, 7 },
    { 4, 7, 6 },
    { 1, 3, 7 },
    { 1, 7, 5 },
    { 0, 4, 6 },
    { 0, 6, 2 },
    { 2, 6, 7 },
    { 2, 7, 3 },
    { 0, 1, 5 },
    { 0, 5, 4 }
};
    
```

SIGGRAPH 2018 | July 24, 2020

Triangles Represented as an Array of Structures

```

From the file SampleVertexData.cpp:
struct vertex
{
    glm::vec3 position;
    glm::vec3 normal;
    glm::vec3 color;
    glm::vec2 texCoord;
};

struct vertex VertexData[ ] =
{
    // triangle 0-2-3:
    // vertex #0:
    { -1.0, -1.0, 1.0 },
    { 0.0, -1.0, 1.0 },
    { 0.0, 0.0, 1.0 },
    // vertex #1:
    { -1.0, -1.0, 1.0 },
    { 0.0, -1.0, 1.0 },
    { 0.0, 1.0, 1.0 },
    // vertex #2:
    { -1.0, -1.0, 1.0 },
    { 0.0, -1.0, 1.0 },
    { 0.0, 1.0, 1.0 },
    // vertex #3:
    { 1.0, -1.0, 1.0 },
    { 0.0, -1.0, 1.0 },
    { 0.0, 1.0, 1.0 }
};
    
```

Modeled in right-handed coordinates

SIGGRAPH 2018 | July 24, 2020

Non-indexed Buffer Drawing

```

From the file SampleVertexData.cpp:
struct vertex
{
    glm::vec3 position;
    glm::vec3 normal;
    glm::vec3 color;
    glm::vec2 texCoord;
};

struct vertex VertexData[ ] =
{
    // triangle 0-2-3:
    // vertex #0:
    { -1.0, -1.0, 1.0 },
    { 0.0, -1.0, 1.0 },
    { 0.0, 0.0, 1.0 },
    // vertex #1:
    { -1.0, -1.0, 1.0 },
    { 0.0, -1.0, 1.0 },
    { 0.0, 1.0, 1.0 },
    // vertex #2:
    { -1.0, -1.0, 1.0 },
    { 0.0, -1.0, 1.0 },
    { 0.0, 1.0, 1.0 },
    // vertex #3:
    { 1.0, -1.0, 1.0 },
    { 0.0, -1.0, 1.0 },
    { 0.0, 1.0, 1.0 }
};
    
```

SIGGRAPH 2018 | July 24, 2020

Filling the Vertex Buffer

```

struct vertex VertexData[ ] =
{
    ...
};

MyBuffer MyVertexBuffer;

InitOSMyVertexBuffer( sizeof(VertexData), OUT &MyVertexBuffer );
FillOSDataBuffer( MyVertexBuffer, (void *) VertexData );

VkResult
InitOSMyVertexBuffer( IN VkDeviceSize size, OUT MyBuffer * pMyBuffer )
{
    VkResult result;
    result = InitOSDataBuffer( size, VK_BUFFER_USAGE_VERTEX_BUFFER_BIT, pMyBuffer );
    return result;
}
    
```

SIGGRAPH 2018 | July 24, 2020

A Preview of What Init05DataBuffer Does

```

VkResult
Init05DataBuffer( VkDeviceSize size, VkBufferUsageFlags usage, OUT MyBuffer * pMyBuffer )
{
    VkResult result = VK_SUCCESS;
    VkBufferCreateInfo vbc;
    vbc.sType = VK_STRUCTURE_TYPE_BUFFER_CREATE_INFO;
    vbc.pNext = nullptr;
    vbc.flags = 0;
    vbc.size = pMyBuffer->size;
    vbc.usage = usage;
    vbc.sharingMode = VK_SHARING_MODE_EXCLUSIVE;
    vbc.queueFamilyIndexCount = 0;
    vbc.pQueueFamilyIndices = (const uint32_t *)nullptr;
    result = vkCreateBuffer( LogicalDevice, IN &vbc, PALLOCATOR, OUT &pMyBuffer->buffer );

    VkMemoryRequirements vmr;
    vkGetBufferMemoryRequirements( LogicalDevice, IN pMyBuffer->buffer, OUT &vmr ); // fills vmr

    VkMemoryAllocateInfo vmai;
    vmai.sType = VK_STRUCTURE_TYPE_MEMORY_ALLOCATE_INFO;
    vmai.pNext = nullptr;
    vmai.allocationSize = vmr.size;
    vmai.memoryTypeIndex = FindMemoryThatIsHostVisible();

    VkDeviceMemory vdm;
    result = vkAllocateMemory( LogicalDevice, IN &vmai, PALLOCATOR, OUT &vdm );
    pMyBuffer->vdm = vdm;

    result = vkBindBufferMemory( LogicalDevice, pMyBuffer->buffer, IN vdm, 0 ); // 0 is the offset
    return result;
}
    
```

Telling the Pipeline about its Input

We will come to the Pipeline later, but for now, know that a Vulkan pipeline is essentially a very large data structure that holds (what OpenGL would call) the state, including how to parse its input.

C/C++:

```

struct vertex
{
    glm::vec3 position;
    glm::vec3 normal;
    glm::vec3 color;
    glm::vec2 texCoord;
};
        
```

GLSL Shader:

```

layout( location = 0 ) in vec3 aVertex;
layout( location = 1 ) in vec3 aNormal;
layout( location = 2 ) in vec3 aColor;
layout( location = 3 ) in vec2 aTexCoord;
        
```

```

VkVertexInputBindingDescription vbld[1]; // one of these per buffer data buffer
vbld[0].binding = 0; // which binding # this is
vbld[0].stride = sizeof( struct vertex ); // bytes between successive structs
vbld[0].inputRate = VK_VERTEX_INPUT_RATE_VERTEX;
    
```

Telling the Pipeline about its Input

C/C++:

```

struct vertex
{
    glm::vec3 position;
    glm::vec3 normal;
    glm::vec3 color;
    glm::vec2 texCoord;
};
        
```

GLSL Shader:

```

layout( location = 0 ) in vec3 aVertex;
layout( location = 1 ) in vec3 aNormal;
layout( location = 2 ) in vec3 aColor;
layout( location = 3 ) in vec2 aTexCoord;
        
```

```

VkVertexInputAttributeDescription vviad[4]; // array per vertex input attribute
// 4 = vertex, normal, color, texture coord
vviad[0].location = 0; // location in the layout decoration
vviad[0].binding = 0; // which binding description this is part of
vviad[0].format = VK_FORMAT_VEC3; // x, y, z
vviad[0].offset = offsetof( struct vertex, position ); // 0

vviad[1].location = 1;
vviad[1].binding = 0;
vviad[1].format = VK_FORMAT_VEC3; // nx, ny, nz
vviad[1].offset = offsetof( struct vertex, normal ); // 12

vviad[2].location = 2;
vviad[2].binding = 0;
vviad[2].format = VK_FORMAT_VEC3; // r, g, b
vviad[2].offset = offsetof( struct vertex, color ); // 24

vviad[3].location = 3;
vviad[3].binding = 0;
vviad[3].format = VK_FORMAT_VEC2; // s, t
vviad[3].offset = offsetof( struct vertex, texCoord ); // 36
    
```

Always use the C/C++ construct **offsetof**, rather than hardcoding the value!

Telling the Pipeline about its Input

We will come to the Pipeline later, but for now, know that a Vulkan Pipeline is essentially a very large data structure that holds (what OpenGL would call) the state, including how to parse its vertex input.

```

VkPipelineVertexInputStateCreateInfo vpvisci; // used to describe the input vertex attributes
vpvisci.sType = VK_STRUCTURE_TYPE_PIPELINE_VERTEX_INPUT_STATE_CREATE_INFO;
vpvisci.pNext = nullptr;
vpvisci.flags = 0;
vpvisci.vertexBindingDescriptionCount = 1;
vpvisci.pVertexBindingDescriptions = &vbld;
vpvisci.vertexAttributeDescriptionCount = 4;
vpvisci.pVertexAttributeDescriptions = &vviad;

VkPipelineInputAssemblyStateCreateInfo vpiasci;
vpiasci.sType = VK_STRUCTURE_TYPE_PIPELINE_INPUT_ASSEMBLY_STATE_CREATE_INFO;
vpiasci.pNext = nullptr;
vpiasci.flags = 0;
vpiasci.topology = VK_PRIMITIVE_TOPOLOGY_TRIANGLE_LIST;;
    
```

Telling the Pipeline about its Input

We will come to the Pipeline later, but for now, know that a Vulkan Pipeline is essentially a very large data structure that holds (what OpenGL would call) the state, including how to parse its vertex input.

```

VkGraphicsPipelineCreateInfo vgpcci;
vgpcci.sType = VK_STRUCTURE_TYPE_GRAPHICS_PIPELINE_CREATE_INFO;
vgpcci.pNext = nullptr;
vgpcci.flags = 0;
vgpcci.stageCount = 2; // number of shader stages in this pipeline
vgpcci.pStages = vpscsi;
vgpcci.pVertexInputState = &vpvisci;
vgpcci.pInputAssemblyState = &vpiasci;
vgpcci.pTessellationState = (VkPipelineTessellationStateCreateInfo *)nullptr; // &vptsci
vgpcci.pViewportState = &vpvsci;
vgpcci.pRasterizationState = &vrpisci;
vgpcci.pMultisampleState = &vpmsci;
vgpcci.pDepthStencilState = &vpdssci;
vgpcci.pColorBlendState = &vpbcsci;
vgpcci.pDynamicState = &vpdsci;
vgpcci.layout = IN GraphicsPipelineLayout;
vgpcci.renderPass = IN RenderPass;
vgpcci.subpass = 0; // subpass number
vgpcci.basePipelineHandle = (VkPipeline) VK_NULL_HANDLE;
vgpcci.basePipelineIndex = 0;

result = vkCreateGraphicsPipelines( LogicalDevice, VK_NULL_HANDLE, 1, IN &vgpcci, PALLOCATOR, OUT &GraphicsPipeline );
    
```

Telling the Command Buffer what Vertices to Draw

We will come to Command Buffers later, but for now, know that you will specify the vertex buffer that you want drawn.

```

VkBuffer buffers[1] = MyVertexDataBuffer.buffer;

vkCmdBindVertexBuffers( CommandBuffers[nextImageIndex], 0, 1, vertexDataBuffers, offsets );

const uint32_t vertexCount = sizeof( VertexData ) / sizeof( VertexData[0] );
const uint32_t instanceCount = 1;
const uint32_t firstVertex = 0;
const uint32_t firstInstance = 0;

vkCmdDraw( CommandBuffers[nextImageIndex], vertexCount, instanceCount, firstVertex, firstInstance );
    
```

Always use the C/C++ construct **sizeof**, rather than hardcoding a count!

Drawing with an Index Buffer

```

struct JustVertexData[] =
{
    // vertex #0:
    { -1., -1., -1. },
    { 0., 0., -1. },
    { 0., 0., 0. },
    { 1., 0. },
},
    // vertex #1:
    { -1., -1., -1. },
    { 0., 0., -1. },
    { 1., 0., 0. },
    { 0., 0. },
},
    ...
int JustIndexData[] =
{
    0, 2, 3,
    0, 3, 1,
    4, 5, 7,
    4, 7, 6,
    1, 5, 7,
    1, 7, 5,
    0, 4, 6,
    0, 6, 2,
    2, 6, 7,
    2, 7, 3,
    0, 1, 5,
    0, 5, 4,
};
                
```

Stream of Vertices

Vertex 7
Vertex 5
Vertex 4
Vertex 1
Vertex 0
Vertex 3
Vertex 2
Vertex 0

Vertex Lookup

{ -1., -1., -1. }
{ 0., 0., -1. }
{ -1., 1., -1. }
{ 1., 1., -1. }
{ -1., -1., 1. }
{ 1., -1., 1. }
{ -1., 1., 1. }
{ 1., 1., 1. }

Stream of Indices

7
5
4
1
3
0
3
2
0

Triangles

Draw

Drawing with an Index Buffer

```

vkCmdBindVertexBuffers( commandBuffer, firstBinding, bindingCount, vertexDataBuffers, vertexOffsets );
vkCmdBindIndexBuffer( commandBuffer, indexDataBuffer, indexOffset, indexType );
                
```

```

typedef enum VkIndexType
{
    VK_INDEX_TYPE_UINT16 = 0, // 0 - 65,535
    VK_INDEX_TYPE_UINT32 = 1, // 0 - 4,294,967,295
} VkIndexType;
                
```

```

vkCmdDrawIndexed( commandBuffer, indexCount, instanceCount, firstIndex, vertexOffset, firstInstance );
                
```

Drawing with an Index Buffer

```

VKResult
Init05MyIndexDataBuffer( IN VkDeviceSize size, OUT MyBuffer * pMyBuffer )
{
    VkResult result = Init05DataBuffer( size, VK_BUFFER_USAGE_INDEX_BUFFER_BIT, pMyBuffer );
    // fills pMyBuffer
    return result;
}

Init05MyVertexDataBuffer( sizeof( JustVertexData ), IN &MyJustVertexDataBuffer );
Fill05DataBuffer( MyJustVertexDataBuffer, (void *) JustVertexData );

Init05MyIndexDataBuffer( sizeof( JustIndexData ), IN &MyJustIndexDataBuffer );
Fill05DataBuffer( MyJustIndexDataBuffer, (void *) JustIndexData );
                
```

Drawing with an Index Buffer

```

VkBuffer vBuffers[1] = { MyJustVertexDataBuffer.buffer };
VkBuffer iBuffer = { MyJustIndexDataBuffer.buffer };

vkCmdBindVertexBuffers( CommandBuffers[nextImageIndex], 0, 1, vBuffers, offsets );
// 0, 1 = firstBinding, bindingCount
vkCmdBindIndexBuffer( CommandBuffers[nextImageIndex], iBuffer, 0, VK_INDEX_TYPE_UINT32 );

const uint32_t vertexCount = sizeof( JustVertexData ) / sizeof( JustVertexData[0] );
const uint32_t indexCount = sizeof( JustIndexData ) / sizeof( JustIndexData[0] );
const uint32_t instanceCount = 1;
const uint32_t firstVertex = 0;
const uint32_t firstIndex = 0;
const uint32_t firstInstance = 0;
const uint32_t vertexOffset = 0;

vkCmdDrawIndexed( CommandBuffers[nextImageIndex], indexCount, instanceCount, firstIndex,
vertexOffset, firstInstance );
                
```

Indirect Drawing (not to be confused with Indexed)

```

typedef struct
VkDrawIndirectCommand
{
    uint32_t vertexCount;
    uint32_t instanceCount;
    uint32_t firstVertex;
    uint32_t firstInstance;
} VkDrawIndirectCommand;

vkCmdDrawIndirect( CommandBuffers[nextImageIndex], buffer, offset, drawCount, stride );

Compare this with:
vkCmdDraw( CommandBuffers[nextImageIndex], vertexCount, instanceCount, firstVertex, firstInstance );
                
```

Indexed Indirect Drawing (i.e., both Indexed and Indirect)

```

vkCmdDrawIndexedIndirect( commandBuffer, buffer, offset, drawCount, stride );

typedef struct
VkDrawIndexedIndirectCommand
{
    uint32_t indexCount;
    uint32_t instanceCount;
    uint32_t firstIndex;
    int32_t vertexOffset;
    uint32_t firstInstance;
} VkDrawIndexedIndirectCommand;

Compare this with:
vkCmdDrawIndexed( commandBuffer, indexCount, instanceCount, firstIndex, vertexOffset, firstInstance );
                
```

Sometimes the Same Point Needs Multiple Attributes

Sometimes a point that is common to multiple faces has the same attributes, no matter what face it is in. Sometimes it doesn't.

A color-interpolated cube like this actually has both. Point #7 above has the same color, regardless of what face it is in. However, Point #7 has 3 different normal vectors, depending on which face you are defining. Same with its texture coordinates.

Thus, when using indexed buffer drawing, you need to create a new vertex struct if any of (position, normal, color, texCoords) changes from what was previously stored at those coordinates.

SIGGRAPH 2018

Sometimes the Same Point Needs Multiple Attributes

Where values do not match at the corners (texture coordinates)

Where values match at the corners (color)

SIGGRAPH 2018

The OBJ File Format – a triple-indexed way of Drawing

```

v 1.710541 1.283360 -0.040860
v 1.714593 1.273043 -0.041268
v 1.706114 1.279109 -0.040795
v 1.719083 1.272235 -0.041195
v 1.722786 1.267216 -0.041939
v 1.727196 1.271285 -0.041795
v 1.730680 1.261384 -0.042630
v 1.723121 1.280378 -0.037323
v 1.714513 1.286599 -0.037101
v 1.706156 1.293797 -0.037073
v 1.702207 1.290297 -0.040704
v 1.697843 1.285852 -0.040489
v 1.709169 1.295845 -0.029862
v 1.717523 1.289344 -0.029807
...
vn 0.1725 0.2557 -0.9512
vn -0.1979 -0.1899 -0.9616
vn -0.2050 -0.2127 -0.9554
vn 0.1664 0.3020 -0.9387
vn -0.2040 -0.1718 -0.9638
vn 0.1645 0.3203 -0.9329
vn -0.2055 -0.1698 -0.9638
vn 0.4419 0.6436 -0.6249
vn 0.4573 0.5682 -0.6841
vn 0.5160 0.5538 -0.6535
vn 0.1791 0.2082 -0.9616
vn -0.2167 -0.2250 -0.9499
vn 0.6624 0.6871 -0.2987
...
vt 0.816406 0.955536
vt 0.822754 0.959168
vt 0.815918 0.959442
vt 0.823242 0.952952
vt 0.829102 0.958862
vt 0.829590 0.955109
vt 0.835449 0.958618
vt 0.824219 0.951263
vt 0.817383 0.951538
vt 0.810050 0.951385
vt 0.809570 0.955383
vt 0.809082 0.959320
vt 0.811035 0.946381
...
f 173/73/75 65/65/67 66/66/68
66/66/68 74/74/76 73/73/75
174/74/76 66/66/68 67/67/69
167/67/69 75/75/77 74/74/76
75/75/77 67/67/69 69/69/71
69/69/71 76/76/78 75/75/77
171/71/73 72/72/74 71/71/79
172/72/74 78/78/80 71/71/79
78/78/80 72/72/74 73/73/75
173/73/75 79/79/81 78/78/80
179/79/81 73/73/75 74/74/76
174/74/76 80/80/82 79/79/81
180/80/82 74/74/76 75/75/77
175/75/77 81/81/83 80/80/82
    
```

Note: The OBJ file format uses 1-based indexing for faces!

V / T / N

http://cs.oregonstate.edu/~mjb/vulkan

SIGGRAPH 2018

Vulkan. Shaders and SPIR-V

Mike Bailey
mjb@cs.oregonstate.edu

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License

http://cs.oregonstate.edu/~mjb/vulkan

SIGGRAPH 2018

The Shaders' View of the Basic Computer Graphics Pipeline

- In general, you want to have a vertex and fragment shader as a minimum.
- A missing stage is OK. The output from one stage becomes the input of the next stage that is there.
- The last stage before the fragment shader feeds its output variables into the rasterizer. The interpolated values then go to the fragment shaders

Legend:
 = Fixed Function
 = Programmable

SIGGRAPH 2018

Vulkan Shader Stages

Shader stages

```

typedef enum VkPipelineStageFlagsBits {
    VK_PIPELINE_STAGE_TOP_OF_PIPE_BIT = 0x00000001,
    VK_PIPELINE_STAGE_DRAW_INDIRECT_BIT = 0x00000002,
    VK_PIPELINE_STAGE_VERTEX_INPUT_BIT = 0x00000004,
    VK_PIPELINE_STAGE_VERTEX_SHADER_BIT = 0x00000008,
    VK_PIPELINE_STAGE_TESSELLATION_CONTROL_SHADER_BIT = 0x00000010,
    VK_PIPELINE_STAGE_TESSELLATION_EVALUATION_SHADER_BIT = 0x00000020,
    VK_PIPELINE_STAGE_GEOMETRY_SHADER_BIT = 0x00000040,
    VK_PIPELINE_STAGE_FRAGMENT_SHADER_BIT = 0x00000080,
    VK_PIPELINE_STAGE_EARLY_FRAGMENT_TESTS_BIT = 0x00000100,
    VK_PIPELINE_STAGE_LATE_FRAGMENT_TESTS_BIT = 0x00000200,
    VK_PIPELINE_STAGE_COLOR_ATTACHMENT_OUTPUT_BIT = 0x00000400,
    VK_PIPELINE_STAGE_COMPUTE_SHADER_BIT = 0x00000800,
    VK_PIPELINE_STAGE_TRANSFER_BIT = 0x00001000,
    VK_PIPELINE_STAGE_BOTTOM_OF_PIPE_BIT = 0x00002000,
    VK_PIPELINE_STAGE_HOST_BIT = 0x00004000,
    VK_PIPELINE_STAGE_ALL_GRAPHICS_BIT = 0x00008000,
    VK_PIPELINE_STAGE_ALL_COMMANDS_BIT = 0x00010000,
} VkPipelineStageFlagsBits;
    
```

SIGGRAPH 2018

How Vulkan GLSL Differs from OpenGL GLSL

85

Detecting that a GLSL Shader is being used with Vulkan/SPIR-V:

- In the compiler, there is an automatic `#define VULKAN 100`

Vulkan Vertex and Instance indices:

```
gl_VertexIndex
gl_InstanceIndex
```


OpenGL uses:

```
gl_VertexID
gl_InstanceID
```

- Both are 0-based

gl_FragColor:

- In OpenGL, `gl_FragColor` broadcasts to all color attachments
- In Vulkan, it just broadcasts to color attachment location #0
- Best idea: don't use it at all – explicitly declare out variables to have specific location numbers



mb - July 24, 2020

How Vulkan GLSL Differs from OpenGL GLSL

86

Shader combinations of separate texture data and samplers:

```
uniform sampler s;
uniform texture2D t;
vec4 rgba = texture( sampler2D( t, s ), vST );
```

Note: our sample code doesn't use this.

Descriptor Sets:

```
layout( set=0, binding=0 ) ... ;
```

Push Constants:

```
layout( push_constant ) ... ;
```

Specialization Constants:


```
layout( constant_id = 3 ) const int N = 5;
```

- Only for scalars, but a vector's components can be constructed from specialization constants

Specialization Constants for Compute Shaders:

```
layout( local_size_x_id = 8, local_size_y_id = 16 );
```

- This sets `gl_WorkGroupSize.x` and `gl_WorkGroupSize.y`
- `gl_WorkGroupSize.z` is set as a constant



mb - July 24, 2020

Vulkan: Shaders' use of Layouts for Uniform Variables

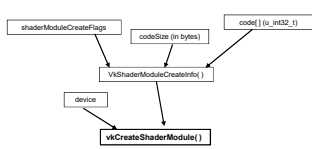

87

```
// non-sampler variables must be in a uniform block:
layout( std140, set = 0, binding = 0 ) uniform matBuf
{
    mat4 uModelMatrix;
    mat4 uViewMatrix;
    mat4 uProjectionMatrix;
    mat3 uNormalMatrix;
} Matrices;

// non-sampler variables must be in a uniform block:
layout( std140, set = 1, binding = 0 ) uniform lightBuf
{
    vec4 uLightPos;
} Light;

layout( set = 2, binding = 0 ) uniform sampler2D uTexUnit;
```

All non-sampler uniform variables must be in block buffers





mb - July 24, 2020

Vulkan Shader Compiling


88

- You half-compile your shaders with an external compiler
- Your shaders get turned into an intermediate form known as SPIR-V, which stands for **Standard Portable Intermediate Representation**.
- SPIR-V gets turned into fully-compiled code at runtime, when the pipeline structure is finally created
- The SPIR-V spec has been public for a few years –new shader languages are surely being developed
- OpenGL and OpenCL have now adopted SPIR-V as well



Advantages:

- Software vendors don't need to ship their shader source
- Syntax errors appear during the SPIR-V step, not during runtime
- Software can launch faster because half of the compilation has already taken place
- This guarantees a common front-end syntax
- This allows for other language front-ends



mb - July 24, 2020

SPIR-V: Standard Portable Intermediate Representation for Vulkan

89

`glslangValidator shaderFile -V [-H] [-I<dir>] [-S <stage>] -o shaderBinaryFile.spv`

Shaderfile extensions:


.vert	Vertex
.tesc	Tessellation Control
.tese	Tessellation Evaluation
.geom	Geometry
.frag	Fragment
.comp	Compute

(Can be overridden by the `-S` option)

Options:

- `-V` Compile for Vulkan
- `-G` Compile for OpenGL
- `-I` Directory(ies) to look in for `#includes`
- `-S` Specify stage rather than get it from shaderfile extension
- `-c` Print out the maximum sizes of various properties

Windows: `glslangValidator.exe`
Linux: `glslangValidator`

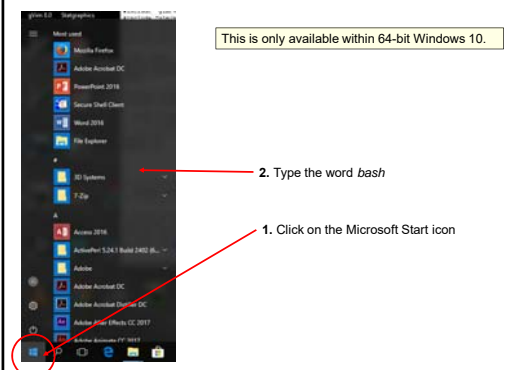


mb - July 24, 2020

You Can Run the SPIR-V Compiler on Windows from a Bash Shell

90

This is only available within 64-bit Windows 10.



- Click on the Microsoft Start icon
- Type the word `bash`



mb - July 24, 2020

You Can Run the SPIR-V Compiler on Windows from a Bash Shell

This is only available within 64-bit Windows 10.

Pick one:

- Can get to your personal folders
- Does not have make
- Can get to your personal folders
- Does have make

SIGGRAPH 2020

Running glslangValidator.exe

```

MINGW64:/y/Vulkan/Sample2017
$ !85
glslangValidator.exe -V sample-vert.vert -o sample-vert.spv
sample-vert.vert
$ !86
glslangValidator.exe -V sample-frag.frag -o sample-frag.spv
sample-frag.frag
MINGW64:/y/Vulkan/Sample2017
$
    
```

SIGGRAPH 2020

Running glslangValidator.exe

```

glslangValidator.exe -V sample-vert.vert -o sample-vert.spv
    
```

Compile for Vulkan ("-G" is compile for OpenGL)

The input file. The compiler determines the shader type by the file extension:

- .vert Vertex shader
- .tccs Tessellation Control Shader
- .tecs Tessellation Evaluation Shader
- .geom Geometry shader
- .frag Fragment shader
- .comp Compute shader

Specify the output file

SIGGRAPH 2020

How do you know if SPIR-V compiled successfully?

Same as C/C++ -- the compiler gives you no nasty messages.

Also, if you care, legal .spv files have a magic number of **0x07230203**

So, if you do an **od -x** on the .spv file, the magic number looks like this:

```

0203 0723 ...
    
```

SIGGRAPH 2020

Reading a SPIR-V File into a Vulkan Shader Module

```

#define SPIRV_MAGIC    0x07230203
...
VkResult
Init12spirvShader( std::string filename, VkShaderModule * pShaderModule )
{
    FILE *fp;
    (void) fopen_s( &fp, filename.c_str(), "rb" );
    if( fp == NULL )
    {
        fprintf( FpDebug, "Cannot open shader file '%s'\n", filename.c_str() );
        return VK_SHOULD_EXIT;
    }
    uint32_t magic;
    fread( &magic, 4, 1, fp );
    if( magic != SPIRV_MAGIC )
    {
        fprintf( FpDebug, "Magic number for spir-v file '%s' is 0x%08x -- should be 0x%08x\n",
            filename.c_str(), magic, SPIRV_MAGIC );
        return VK_SHOULD_EXIT;
    }

    fseek( fp, 0L, SEEK_END );
    int size = ftell( fp );
    rewind( fp );
    unsigned char *code = new unsigned char [size];
    fread( code, size, 1, fp );
    fclose( fp );
}
    
```

SIGGRAPH 2020

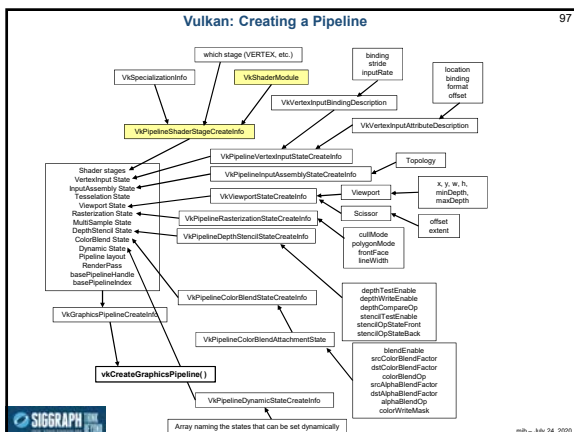
Reading a SPIR-V File into a Shader Module

```

VkShaderModule ShaderModuleVertex;
...
VkShaderModuleCreateInfo vsmci;
vsmci.sType = VK_STRUCTURE_TYPE_SHADER_MODULE_CREATE_INFO;
vsmci.pNext = nullptr;
vsmci.flags = 0;
vsmci.codeSize = size;
vsmci.pCode = (uint32_t *)code;

VkResult result = vkCreateShaderModule( LogicalDevice, &vsmci, PALLOCATOR, OUT & ShaderModuleVertex );
fprintf( FpDebug, "Shader Module '%s' successfully loaded!\n", filename.c_str() );
delete [ ] code;
return result;
}
    
```

SIGGRAPH 2020



You can also take a look at SPIR-V Assembly

gslangValidator.exe -V **H sample-vert.vert -o sample-vert.spv**

This prints out the SPIR-V "assembly" to standard output. Other than nerd interest, there is no graphics-programming reason to look at this. ☺

SIGGRAPH 2020

mb - July 24, 2020

For example, if this is your Shader Source

```
#version 400
#extension GL_ARB_separate_shader_objects : enable
#extension GL_ARB_shading_language_400pack : enable
layout(std140, set = 0, binding = 0) uniform mat4f
{
    mat4 uModelMatrix;
    mat4 uViewMatrix;
    mat4 uProjectionMatrix;
    mat3 uNormalMatrix;
};

// non-opaque must be in a uniform block:
layout(std140, set = 1, binding = 0) uniform lightBuf
{
    vec4 uLightPos;
};

// Light
layout(location = 0) in vec3 aVertex;
layout(location = 1) in vec3 aNormal;
layout(location = 2) in vec3 aColor;
layout(location = 3) in vec3 aTexCoord;

layout(location = 0) out vec3 vNormal;
layout(location = 1) out vec3 vColor;
layout(location = 2) out vec3 vTexCoord;

void main()
{
    mat4 PVM = Matrices * ProjectionMatrix * Matrices * uViewMatrix * Matrices * uModelMatrix;
    gl_Position = PVM * vec4(aVertex, 1.);
    vNormal = Matrices * uNormalMatrix * aNormal;
    vColor = aColor;
    vTexCoord = aTexCoord;
}
```

SIGGRAPH 2020

mb - July 24, 2020

This is the SPIR-V Assembly, Part I

```
#version 400
#extension GL_ARB_separate_shader_objects : enable
#extension GL_ARB_shading_language_400pack : enable
layout(std140, set = 0, binding = 0) uniform mat4f
{
    mat4 uModelMatrix;
    mat4 uViewMatrix;
    mat4 uProjectionMatrix;
    mat3 uNormalMatrix;
};

// non-opaque must be in a uniform block:
layout(std140, set = 1, binding = 0) uniform lightBuf
{
    vec4 uLightPos;
};

// Light
layout(location = 0) in vec3 aVertex;
layout(location = 1) in vec3 aNormal;
layout(location = 2) in vec3 aColor;
layout(location = 3) in vec3 aTexCoord;

layout(location = 0) out vec3 vNormal;
layout(location = 1) out vec3 vColor;
layout(location = 2) out vec3 vTexCoord;

void main()
{
    mat4 PVM = Matrices * ProjectionMatrix * Matrices * uViewMatrix * Matrices * uModelMatrix;
    gl_Position = PVM * vec4(aVertex, 1.);
    vNormal = Matrices * uNormalMatrix * aNormal;
    vColor = aColor;
    vTexCoord = aTexCoord;
}
```

Capability Shader
 ExtendingOp GLSL_std_400
 MemoryModel Logical GLSL400
 EntryPoint Vertices 4, "main" 34, 37, 48, 53, 56, 57, 61, 63
 Source GLSL 400
 SourceExtension GL_ARB_separate_shader_objects
 SourceExtension GL_ARB_shading_language_400pack
 Name 4 "main"
 Name 10 "PVM"
 Name 13 "uModelMatrix"
 MemberName 13(mat4) 0 "uModelMatrix"
 MemberName 13(mat4) 1 "uViewMatrix"
 MemberName 13(mat4) 2 "uProjectionMatrix"
 MemberName 13(mat3) 3 "uNormalMatrix"
 Name 15 "lightBuf"
 MemberName 15(vec4) 0 "gl_Position"
 MemberName 15(vec4) 1 "gl_PointSize"
 MemberName 15(vec4) 2 "gl_ClipDistance"
 Name 34 "main"
 Name 37 "vNormal"
 Name 48 "vColor"
 Name 53 "vTexCoord"
 Name 56 "vColor"
 Name 61 "vTexCoord"
 Name 63 "aVertex"
 Name 65 "aNormal"
 Name 66 "aColor"
 MemberName 66(lightBuf) 0 "uLightPos"
 Name 67 "Light"
 MemberDecorate 13(mat4) 0 ColMajor
 MemberDecorate 13(mat4) 1 ColMajor
 MemberDecorate 13(mat4) 2 ColMajor
 MemberDecorate 13(mat4) 1 Offset 64
 MemberDecorate 13(mat4) 2 Offset 64
 MemberDecorate 13(mat4) 2 ColMajor
 MemberDecorate 13(mat4) 2 Offset 128
 MemberDecorate 13(mat4) 2 MatrixStride 16
 MemberDecorate 13(mat4) 3 ColMajor
 MemberDecorate 13(mat4) 3 Offset 192
 MemberDecorate 13(mat4) 3 MatrixStride 16
 Decorate 15(mat4) Block
 Decorate 15(mat4) DescriptorSet 0

SIGGRAPH 2020

mb - July 24, 2020

This is the SPIR-V Assembly, Part II

```
Decorate 15(mat4) Block  

MemberDecorate 32(gl_PerVertex) 0 BuiltIn Position  

MemberDecorate 32(gl_PerVertex) 1 BuiltIn PointSize  

MemberDecorate 32(gl_PerVertex) 2 BuiltIn ClipDistance  

Decorate 32(gl_PerVertex) Block  

Decorate 32(gl_PerVertex) Location 0  

Decorate 48(vNormal) Location 1  

Decorate 48(vColor) Location 1  

Decorate 61(vTexCoord) Location 2  

Decorate 61(vTexCoord) Location 2  

Decorate 61(vTexCoord) Location 3  

MemberDecorate 65(lightBuf) 0 Offset 0  

Decorate 65(lightBuf) Block  

Decorate 67(Light) DescriptorSet 1  

Decorate 67(Light) Binding 0  

TypeVec 2  

TypeFunction 2  

TypeFloat 32  

TypeVector 6(float) 4  

TypeMatrix 7(float) 4  

TypePointer Function 8  

TypeVector 6(float) 3  

TypeMatrix 11(float) 3  

TypeStruct 11(1vec3) 3  

TypePointer Uniform 13(mat4f)  

15(matrices): 14gl_PerVertex Uniform  

16: TypeLight Uniform 8  

17: 16pptr Constant 2  

18: TypePointer Uniform 8  

21: 16pptr Constant 1  

22: 16pptr Constant 1  

23: 16pptr Constant 1  

29: TypeLight 32 0  

30: 29pptr Constant 1  

31: TypeArray 6(float) 30  

32(gl_PerVertex): TypeStruct 7(float) 6(float) 31  

32(gl_PerVertex) Output 32(gl_PerVertex)  

34: 32pptr Variable Output  

36: TypePointer Input 11(float) 3  

37(pptr): 36pptr Variable Input  

38: 6(float) Constant 106533216  

45: TypePointer Output 7(float) 4  

47: TypePointer Output 11(float) 3  

48(vNormal): 47pptr Variable Output  

49: 16pptr Constant 3
```

SIGGRAPH 2020

mb - July 24, 2020

This is the SPIR-V Assembly, Part III

```
50: 50pptr Variable Uniform 12  

51(hvec3): 50pptr Variable Input  

56(vColor): 47pptr Variable Output  

57(vColor): 32pptr Variable Input  

59: TypeVector 6(float) 2  

60: TypePointer Output 58(float) 2  

61(vTexCoord): 60pptr Variable Output  

62: TypePointer Input 59(float) 2  

63(vTexCoord): 62pptr Variable Input  

65(lightBuf): TypePointer Uniform 65(lightBuf)  

67(Light): 2 Function Note 3  

4pptr) 5  

Label 5  

10pptr) 5pptr Variable Function  

18(pptr) AccessChain 15(matrices) 17  

5: Load 19  

8: Load 22  

8: MatrixTimesMatrix 20, 23  

18(pptr) AccessChain 15(matrices) 25  

8: Load 26  

8: MatrixTimesMatrix 24, 27  

8: Load 19(PVM)  

12: Load 21  

11(hvec3) Load 37(vNormal)  

6(float) CompositeLoad 38, 2  

6(float) CompositeLoad 38, 1  

7(float) CompositeLoad 38, 2  

44(pptr) CompositeConstant 40, 41, 42, 39  

44: 7(float) MatrixTimesVector 35, 43  

44pptr) AccessChain 34, 25  

Store 46, 44  

50pptr) AccessChain 15(matrices) 49  

12: Load 21  

11(hvec3) Load 53(vNormal)  

55: 11(hvec3) MatrixTimesVector 52, 54  

Store 48(vColor) 55  

55: 11(hvec3) Load 57(vColor)  

Store 48(vColor) 58  

64: 58(hvec3) Load 63(vTexCoord)  

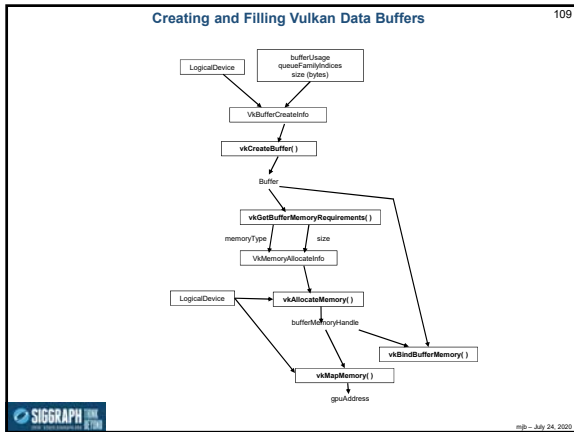
Store 61(vTexCoord) 64  

Return  

FunctionEnd
```

SIGGRAPH 2020

mb - July 24, 2020



Creating a Vulkan Data Buffer 110

```

VkBufferCreateInfo vbc;
vbc.sType = VK_STRUCTURE_TYPE_BUFFER_CREATE_INFO;
vbc.pNext = nullptr;
vbc.flags = 0;
vbc.size = << buffer size in bytes >>
vbc.usage = <<or'ed bits of: >>
    VK_USAGE_TRANSFER_SRC_BIT
    VK_USAGE_TRANSFER_DST_BIT
    VK_USAGE_UNIFORM_TEXEL_BUFFER_BIT
    VK_USAGE_STORAGE_TEXEL_BUFFER_BIT
    VK_USAGE_UNIFORM_BUFFER_BIT
    VK_USAGE_STORAGE_BUFFER_BIT
    VK_USAGE_INDEX_BUFFER_BIT
    VK_USAGE_VERTEX_BUFFER_BIT
    VK_USAGE_INDIRECT_BUFFER_BIT
vbc.sharingMode = << one of: >>
    VK_SHARING_MODE_EXCLUSIVE
    VK_SHARING_MODE_CONCURRENT
vbc.queueFamilyIndexCount = 0;
vbc.pQueueFamilyIndices = (const int32_t) nullptr;

result = vkCreateBuffer( LogicalDevice, IN &vbc, PALLOCATOR, OUT &Buffer );
  
```

SIGGRAPH 2020

Allocating Memory for a Vulkan Data Buffer, Binding a Buffer to Memory, and Writing to the Buffer 111

```

VkMemoryRequirements vmr;
result = vkGetBufferMemoryRequirements( LogicalDevice, Buffer, OUT &vmr );

VkMemoryAllocateInfo vmai;
vmai.sType = VK_STRUCTURE_TYPE_MEMORY_ALLOCATE_INFO;
vmai.pNext = nullptr;
vmai.flags = 0;
vmai.allocationSize = vmr.size;
vmai.memoryTypeIndex = FindMemoryThatIsHostVisible();

...

VkDeviceMemory vdm;
result = vkAllocateMemory( LogicalDevice, IN &vmai, PALLOCATOR, OUT &vdm );

result = vkBindBufferMemory( LogicalDevice, Buffer, IN vdm, 0 ); // 0 is the offset

...

result = vkMapMemory( LogicalDevice, IN vdm, 0, VK_WHOLE_SIZE, 0, &ptr );

<< do the memory copy >>

result = vkUnmapMemory( LogicalDevice, IN vdm );
  
```

SIGGRAPH 2020

Finding the Right Type of Memory 112

```

int FindMemoryThatIsHostVisible()
{
    VkPhysicalDeviceMemoryProperties vpdmp;
    vkGetPhysicalDeviceMemoryProperties( PhysicalDevice, OUT &vpdmp );
    for( unsigned int i = 0; i < vpdmp.memoryTypeCount; i++ )
    {
        VkMemoryType vmt = vpdmp.memoryTypes[i];
        if( ( vmt.propertyFlags & VK_MEMORY_PROPERTY_HOST_VISIBLE_BIT ) != 0 )
        {
            return i;
        }
    }
    return -1;
}
  
```

SIGGRAPH 2020

Finding the Right Type of Memory 113

```

int FindMemoryThatIsDeviceLocal()
{
    VkPhysicalDeviceMemoryProperties vpdmp;
    vkGetPhysicalDeviceMemoryProperties( PhysicalDevice, OUT &vpdmp );
    for( unsigned int i = 0; i < vpdmp.memoryTypeCount; i++ )
    {
        VkMemoryType vmt = vpdmp.memoryTypes[i];
        if( ( vmt.propertyFlags & VK_MEMORY_PROPERTY_DEVICE_LOCAL_BIT ) != 0 )
        {
            return i;
        }
    }
    return -1;
}
  
```

SIGGRAPH 2020

Finding the Right Type of Memory 114

```

VkPhysicalDeviceMemoryProperties vpdmp;
vkGetPhysicalDeviceMemoryProperties( PhysicalDevice, OUT &vpdmp );
  
```

11 Memory Types:

- Memory 0:
- Memory 1:
- Memory 2:
- Memory 3:
- Memory 4:
- Memory 5:
- Memory 6:
- Memory 7: DeviceLocal
- Memory 8: DeviceLocal
- Memory 9: HostVisible HostCoherent
- Memory 10: HostVisible HostCoherent HostCached

2 Memory Heaps:

- Heap 0: size = 0xb7c00000 DeviceLocal
- Heap 1: size = 0xfac00000


SIGGRAPH 2020

Sidebar: The Vulkan Memory Allocator (VMA) 115

The **Vulkan Memory Allocator** is a set of functions to simplify your view of allocating buffer memory. I don't have experience using it (yet), so I'm not in a position to confidently comment on it. But, I am including its github link here and a little sample code in case you want to take a peek.

<https://github.com/GPUOpen-LibrariesAndSDKs/VulkanMemoryAllocator>

This repository includes a smattering of documentation.




mb - July 24, 2020

Sidebar: The Vulkan Memory Allocator (VMA) 116

```
#define VMA_IMPLEMENTATION
#include "vk_mem_alloc.h"
...
VkBufferCreateInfo          vbci;
...
VmaAllocationCreateInfo     vaci;
vac_i.physicalDevice = PhysicalDevice;
vac_i.device = LogicalDevice;
vac_i.usage = VMA_MEMORY_USAGE_GPU_ONLY;

VmaAllocator                var;
vmaCreateAllocator( IN &vac_i, OUT &var );
...
VkBuffer                    Buffer;
VmaAllocation               van;
vmaCreateBuffer( IN var, IN &vbci, IN &vac_i, OUT &Buffer, OUT &van, nullptr );

void *mappedDataAddr;
vmaMapMemory( IN var, IN van, OUT &mappedDataAddr );
memcpy( mappedDataAddr, &MyData, sizeof(MyData) );
vmaUnmapMemory( IN var, IN van );
```



mb - July 24, 2020


Something I've Found Useful 117

I find it handy to encapsulate buffer information in a struct:

```
typedef struct MyBuffer
{
    VkDataBuffer    buffer;
    VkDeviceMemory vdm;
    VkDeviceSize    size;
} MyBuffer;
...
MyBuffer           MyMatrixUniformBuffer;
```

It's the usual object-oriented benefit – you can pass around just one data-item and everyone can access whatever information they need.

It also makes it impossible to accidentally associate the wrong VkDeviceMemory and/or VkDeviceSize with the wrong data buffer.




mb - July 24, 2020

Initializing a Data Buffer 118

It's the usual object-oriented benefit – you can pass around just one data-item and everyone can access whatever information they need.

```
VkResult
Init05DataBuffer( VkDeviceSize size, VkBufferUsageFlags usage, OUT MyBuffer * pMyBuffer )
{
    ...
    vbci.size = pMyBuffer->size = size;
    ...
    result = vkCreateBuffer ( LogicalDevice, IN &vbci, PALLOCATOR, OUT &pMyBuffer->buffer );
    ...
    pMyBuffer->vdm = vdm;
    ...
}
```




mb - July 24, 2020

Here's a C struct used by the Sample Code to hold some uniform variables 119

```
struct matBuf
{
    glm::mat4 uModelMatrix;
    glm::mat4 uViewMatrix;
    glm::mat4 uProjectionMatrix;
    glm::mat3 uNormalMatrix;
} Matrices;
```

Here's the associated GLSL shader code to access those uniform variables

```
layout( std140, set = 0, binding = 0 ) uniform matBuf
{
    mat4 uModelMatrix;
    mat4 uViewMatrix;
    mat4 uProjectionMatrix;
    mat4 uNormalMatrix;
} Matrices;
```



mb - July 24, 2020

Filling those Uniform Variables 120

```
uint32_t          Height, Width;
const double FOV =          glm::radians(60.); // field-of-view angle in radians


glm::vec3 eye(0.,0.,_EYEDIST);
glm::vec3 look(0.,0.,0.);
glm::vec3 up(0.,1.,0.);

Matrices.uModelMatrix    = glm::mat4( 1. ); // identity
Matrices.uViewMatrix     = glm::lookAt( eye, look, up );
Matrices.uProjectionMatrix = glm::perspective( FOV, (double)Width/(double)Height, 0.1, 1000. );
Matrices.uProjectionMatrix[1][1] *= -1.; // account for Vulkan's LH screen coordinate system
Matrices.uNormalMatrix = glm::inverseTranspose( glm::mat3( Matrices.uModelMatrix ) );
```

This code assumes that this line:

```
#define GLM_FORCE_RADIANS
```

is listed before GLM is included!



mb - July 24, 2020

The Parade of Buffer Data

121

MyBuffer MyMatrixUniformBuffer;

The MyBuffer does not hold any actual data itself. It just information about what is in the data buffer

```

VkResult
Init05DataBuffer( VkDeviceSize size, VkBufferUsageFlags usage, OUT MyBuffer * pMyBuffer )
{
    ...
    result = vkCreateBuffer( LogicalDevice, IN &vma, PALLOCATOR, OUT &pMyBuffer->buffer );
    pMyBuffer->vdm = vdm;
    ...
}
  
```

This C struct is holding the original data, written by the application.

```

struct matBuf Matrices;
glm::mat4 eye(0.,0.,EYEDIST);
glm::vec3 look(0.,0.,0.);
glm::vec3 up(0.,1.,0.);
Matrices.uModelMatrix = glm::mat4( eye, look, up );
Matrices.uViewMatrix = glm::lookAt( eye, look, up );
Matrices.uProjectionMatrix = glm::perspective( FOV, (double)Width/(double)Height, 0.1, 1000. );
Matrices.uNormalMatrix = glm::inverseTranspose( glm::mat3( Matrices.uModelMatrix ) );
  
```

Memory-mapped copy operation

The Data Buffer in GPU memory is holding the copied data. It is readable by the shaders

```

uniform matBuf Matrices;
layout( std142, set = 0, binding = 0 ) uniform matBuf
{
    mat4 uModelMatrix;
    mat3 uNormalMatrix;
    mat4 uProjectionMatrix;
    mat4 uViewMatrix;
} Matrices;
  
```

Filling the Data Buffer

122

```

void * pGpuMemory;
vkMapMemory( LogicalDevice, IN myBuffer.vdm, 0, VK_WHOLE_SIZE, 0, OUT &pGpuMemory );
memcopy( pGpuMemory, data, (size_t)myBuffer.size );
vkUnmapMemory( LogicalDevice, IN myBuffer.vdm );
return VK_SUCCESS;
  
```

Init05UniformBuffer(sizeof(Matrices), **OUT &MyMatrixUniformBuffer**);

Fill05DataBuffer(MyMatrixUniformBuffer, **IN (void *) &Matrices**);

```

glm::vec3 eye(0.,0.,EYEDIST);
glm::vec3 look(0.,0.,0.);
glm::vec3 up(0.,1.,0.);
Matrices.uModelMatrix = glm::mat4( ); // identity
Matrices.uViewMatrix = glm::lookAt( eye, look, up );
Matrices.uProjectionMatrix = glm::perspective( FOV, (double)Width/(double)Height, 0.1, 1000. );
Matrices.uProjectionMatrix[1][1] *= -1.;
Matrices.uNormalMatrix = glm::inverseTranspose( glm::mat3( Matrices.uModelMatrix ) );
  
```

Creating and Filling the Data Buffer – the Details

123

```

VkResult
Init05DataBuffer( VkDeviceSize size, VkBufferUsageFlags usage, OUT MyBuffer * pMyBuffer )
{
    VkResult result = VK_SUCCESS;
    VkBufferCreateInfo vbc;
    vbc.sType = VK_STRUCTURE_TYPE_BUFFER_CREATE_INFO;
    vbc.pNext = nullptr;
    vbc.flags = 0;
    vbc.size = pMyBuffer->size;
    vbc.usage = usage;
    vbc.sharingMode = VK_SHARING_MODE_EXCLUSIVE;
    vbc.queueFamilyIndexCount = 0;
    vbc.queueFamilyIndices = (const uint32_t *)nullptr;
    result = vkCreateBuffer( LogicalDevice, IN &vbc, PALLOCATOR, OUT &pMyBuffer->buffer );
    VkMemoryRequirements vmr;
    vkGetBufferMemoryRequirements( LogicalDevice, IN pMyBuffer->buffer, OUT &vmr ); // fills vmr
    VkMemoryAllocateInfo vmai;
    vmai.sType = VK_STRUCTURE_TYPE_MEMORY_ALLOCATE_INFO;
    vmai.pNext = nullptr;
    vmai.allocationSize = vmr.size;
    vmai.memoryTypeIndex = FindMemoryThatIsHostVisible( );
    VkDeviceMemory vdm;
    result = vkAllocateMemory( LogicalDevice, IN &vmai, PALLOCATOR, OUT &vdm );
    pMyBuffer->vdm = vdm;
    result = vkBindBufferMemory( LogicalDevice, pMyBuffer->buffer, IN vdm, OFFSET_ZERO );
    return result;
}
  
```

Creating and Filling the Data Buffer – the Details

124

```

VkResult
Fill05DataBuffer( IN MyBuffer myBuffer, IN void * data )
{
    // the size of the data had better match the size that was used to init the buffer!
    void * pGpuMemory;
    vkMapMemory( LogicalDevice, IN myBuffer.vdm, 0, VK_WHOLE_SIZE, 0, OUT &pGpuMemory );
    memcopy( pGpuMemory, data, (size_t)myBuffer.size ); // 0 and 0 are offset and flags
    vkUnmapMemory( LogicalDevice, IN myBuffer.vdm );
    return VK_SUCCESS;
}
  
```

Remember – to Vulkan and GPU memory, these are just bits. It is up to you to handle their meaning correctly.

Creating and Filling the Data Buffer – the Details

125

```

VkResult
Fill05DataBuffer( IN MyBuffer myBuffer, IN void * data )
{
    // the size of the data had better match the size that was used to init the buffer!
    void * pGpuMemory;
    vkMapMemory( LogicalDevice, IN myBuffer.vdm, 0, VK_WHOLE_SIZE, 0, OUT &pGpuMemory );
    memcopy( pGpuMemory, data, (size_t)myBuffer.size ); // 0 and 0 are offset and flags
    vkUnmapMemory( LogicalDevice, IN myBuffer.vdm );
    return VK_SUCCESS;
}
  
```

Remember – to Vulkan and GPU memory, these are just bits. It is up to you to handle their meaning correctly.

GLFW

Mike Bailey
mjb@cs.oregonstate.edu

<http://cs.oregonstate.edu/~mjb/vulkan>

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License

127

<http://www.glfw.org/>

GLFW is an Open Source, cross-platform library for OpenGL, OpenGL ES and Vulkan development on the desktop. It provides a simple API for creating native windows and surfaces, receiving input and events.

GLFW is written in C and has native support for Windows, macOS and many Unix-like systems using the X Window System, such as Linux and FreeBSD.

GLFW is licensed under the [Zlib/libpng license](#).

- Open your window and OpenGL context with just two function calls.
- Support for OpenGL, OpenGL ES, Vulkan and related options, flags and extensions.
- Support for multiple windows, multiple monitors, high-DPI and gamma control.
- Support for keyboard, mouse, gamepad, touch and window event input, via polling or callbacks.
- Comes with guides, a tutorial, reference documentation, examples and test programs.
- Open Source with an OSI certified license allowing commercial use.
- Access to native objects and complete time options for platform specific features.
- Community maintained bindings for many different languages.

No library can be perfect for everyone. If GLFW isn't what you're looking for, the [web community](#) might help.

SIGGRAPH 2018

mb - July 24, 2020

128

Setting Up GLFW

```
#define GLFW_INCLUDE_VULKAN
#include "glfw3.h"
...
uint32_t      Width, Height;
VkSurfaceKHR  Surface;
...

void
InitGLFW()
{
    glfwInit();
    if( !glfwVulkanSupported() )
    {
        fprintf( stderr, "Vulkan is not supported on this system!\n" );
        exit( 1 );
    }
    glfwWindowHint( GLFW_CLIENT_API, GLFW_NO_API );
    glfwWindowHint( GLFW_RESIZABLE, GLFW_FALSE );
    MainWindow = glfwCreateWindow( Width, Height, "Vulkan Sample", NULL, NULL );
    VkResult result = glfwCreateWindowSurface( Instance, MainWindow, NULL, OUT &Surface );

    glfwSetErrorCallback( GLFWErrorCallback );
    glfwSetKeyCallback( MainWindow, GLFWKeyboard );
    glfwSetCursorPosCallback( MainWindow, GLFWMouseMotion );
    glfwSetMouseButtonCallback( MainWindow, GLFWMouseButton );
}
```

SIGGRAPH 2018

mb - July 24, 2020

129

You Can Also Query What Vulkan Extensions GLFW Requires

```
uint32_t count;
const char ** extensions = glfwGetRequiredInstanceExtensions( &count );

fprintf( FpDebug, "nFound %d GLFW Required Instance Extensions:\n", count );

for( uint32_t i = 0; i < count; i++ )
{
    fprintf( FpDebug, "%t%s\n", extensions[ i ] );
}
```

Found 2 GLFW Required Instance Extensions:
VK_KHR_surface
VK_KHR_win32_surface

SIGGRAPH 2018

mb - July 24, 2020

130

GLFW Keyboard Callback

```
void
GLFWKeyboard( GLFWwindow * window, int key, int scancode, int action, int mods )
{
    if( action == GLFW_PRESS )
    {
        switch( key )
        {
            //case GLFW_KEY_M:
            case 'm':
            case 'M':
                Mode++;
                if( Mode >= 2 )
                    Mode = 0;
                break;

            default:
                fprintf( FpDebug, "Unknown key hit: 0x%04x = %c\n", key );
                fflush( FpDebug );
        }
    }
}
```

SIGGRAPH 2018

mb - July 24, 2020

131

GLFW Mouse Button Callback

```
void
GLFWMouseButton( GLFWwindow *window, int button, int action, int mods )
{
    int b = 0; // LEFT, MIDDLE, or RIGHT

    // get the proper button bit mask:
    switch( button )
    {
        case GLFW_MOUSE_BUTTON_LEFT:
            b = LEFT; break;

        case GLFW_MOUSE_BUTTON_MIDDLE:
            b = MIDDLE; break;

        case GLFW_MOUSE_BUTTON_RIGHT:
            b = RIGHT; break;

        default:
            b = 0;
            fprintf( FpDebug, "Unknown mouse button: %d\n", button );
    }

    // button down sets the bit, up clears the bit:
    if( action == GLFW_PRESS )
    {
        double xpos, ypos;
        glfwGetCursorPos( window, &xpos, &ypos );
        Xmouse = (int)xpos;
        Ymouse = (int)ypos;
        ActiveButton |= b; // set the proper bit
    }
    else
    {
        ActiveButton &= ~b; // clear the proper bit
    }
}
```

SIGGRAPH 2018

mb - July 24, 2020

132

GLFW Mouse Motion Callback

```
void
GLFWMouseMotion( GLFWwindow *window, double xpos, double ypos )
{
    int dx = (int)xpos - Xmouse; // change in mouse coords
    int dy = (int)ypos - Ymouse;

    if( ( ActiveButton & LEFT ) != 0 )
    {
        Xrot += ( ANGFACT * dy );
        Yrot += ( ANGFACT * dx );
    }

    if( ( ActiveButton & MIDDLE ) != 0 )
    {
        Scale += SCLFACT * (float) ( dx - dy );

        // keep object from turning inside-out or disappearing:
        if( Scale < MINSCALE )
            Scale = MINSCALE;
    }

    Xmouse = (int)xpos; // new current position
    Ymouse = (int)ypos;
}
```

SIGGRAPH 2018

mb - July 24, 2020

Looping and Closing GLFW

133

```

while( glfwWindowShouldClose( MainWindow ) == 0 )
{
    glfwPollEvents();
    Time = glfwGetTime(); // elapsed time, in double-precision seconds
    UpdateScene();
    RenderScene();
}

vkQueueWaitIdle( Queue );
vkDeviceWaitIdle( LogicalDevice );
DestroyAllVulkan();
glfwDestroyWindow( MainWindow );
glfwTerminate();
    
```

Does not block – processes any waiting events, then returns

mjb - July 24, 2020

Looping and Closing GLFW

134

If you would like to *block* waiting for events, use:

```
glfwWaitEvents();
```

You can have the blocking wake up after a timeout period with:

```
glfwWaitEventsTimeout( double secs );
```

You can wake up one of these blocks from another thread with:

```
glfwPostEmptyEvent();
```

mjb - July 24, 2020

GLM

Mike Bailey
mjb@cs.oregonstate.edu

This work is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](https://creativecommons.org/licenses/by-nc-nd/4.0/)

<http://cs.oregonstate.edu/~mjb/vulkan>

mjb - July 24, 2020

What is GLM?

136

GLM is a set of C++ classes and functions to fill in the programming gaps in writing the basic vector and matrix mathematics for OpenGL applications. However, even though it was written for OpenGL, it works fine with Vulkan.

Even though GLM looks like a library, it actually isn't – it is all specified in *.hpp header files so that it gets compiled in with your source code.

You can find it at:
<http://glm.g-truc.net/0.9.8.5/>

You invoke GLM like this:

```
#define GLM_FORCE_RADIANS
#include <glm/glm.hpp>
#include <glm/gtc/matrix_transform.hpp>
#include <glm/gtc/matrix_inverse.hpp>
```

OpenGL treats all angles as given in degrees. This line forces GLM to treat all angles as given in radians. I recommend this so that all angles you create in all programming will be in radians.

If GLM is not installed in a system place, put it somewhere you can get access to. Later on, these notes will show you how to use it from there.

mjb - July 24, 2020

Why are we even talking about this?

137

All of the things that we have talked about being *deprecated* in OpenGL are really **deprecated** in Vulkan -- built-in pipeline transformations, begin-end, fixed-function, etc. So, where you might have said in OpenGL:

```

glMatrixMode( GL_MODELVIEW );
glLoadIdentity();
glLookAt( 0.0, 3.0, 0.0, 0.0, 0.0, 1.0 );
glRotatef( (GLfloat)Yrot, 0.0, 1.0, 0.0 );
glRotatef( (GLfloat)Xrot, 1.0, 0.0, 0.0 );
glScalef( (GLfloat)Scale, (GLfloat)Scale, (GLfloat)Scale );
    
```

you would now say:

```

glm::mat4 modelview = glm::mat4( 1. ); // identity
glm::vec3 eye(0.0, 3.0);
glm::vec3 look(0.0, 0.0);
glm::vec3 up(0.1, 1.0);
modelview = glm::lookAt( eye, look, up ); // {x,y,z} = [v]*{x,y,z}
modelview = glm::rotate( modelview, D2R*Yrot, glm::vec3(0.0, 1.0, 0.0) ); // {x,y,z} = [v]*[r]*{x,y,z}
modelview = glm::rotate( modelview, D2R*Xrot, glm::vec3(1.0, 0.0, 0.0) ); // {x,y,z} = [v]*[r]*[x]*{x,y,z}
modelview = glm::scale( modelview, glm::vec3(Scale, Scale, Scale) ); // {x,y,z} = [v]*[r]*[x]*[s]*{x,y,z}
    
```

This is exactly the same concept as OpenGL, but a different expression of it. Read on for details ...

mjb - July 24, 2020

The Most Useful GLM Variables, Operations, and Functions

138

```

// constructor:
glm::mat4( 1. ); // identity matrix
glm::vec4();
glm::vec3();
    
```

GLM recommends that you use the "glm::" syntax and avoid "using namespace" syntax because they have not made any effort to create unique function names

```

// multiplications:
glm::mat4 * glm::mat4
glm::mat4 * glm::vec4
glm::mat4 * glm::vec4( glm::vec3, 1. ) // promote a vec3 to a vec4 via a constructor
    
```

```

// emulating OpenGL transformations with concatenation:
glm::mat4 glm::rotate( glm::mat4 const& m, float angle, glm::vec3 const& axis );
glm::mat4 glm::scale( glm::mat4 const& m, glm::vec3 const& factors );
glm::mat4 glm::translate( glm::mat4 const& m, glm::vec3 const& translation );
    
```

mjb - July 24, 2020

The Most Useful GLM Variables, Operations, and Functions


// viewing volume (assign, not concatenate):

```
glm::mat4 glm::ortho( float left, float right, float bottom, float top, float near, float far );
glm::mat4 glm::ortho( float left, float right, float bottom, float top );
```

```
glm::mat4 glm::frustum( float left, float right, float bottom, float top, float near, float far );
glm::mat4 glm::perspective( float fovy, float aspect, float near, float far );
```

// viewing (assign, not concatenate):

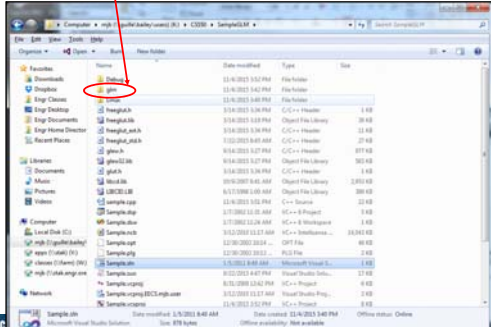

```
glm::mat4 glm::lookAt( glm::vec3 const & eye, glm::vec3 const & look, glm::vec3 const & up );
```



mb - July 24, 2020

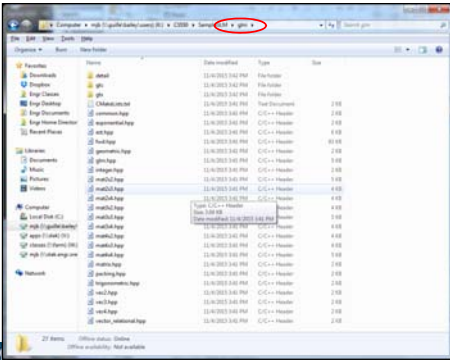

Installing GLM into your own space

I like to just put the whole thing under my Visual Studio project folder so I can zip up a complete project and give it to someone else.

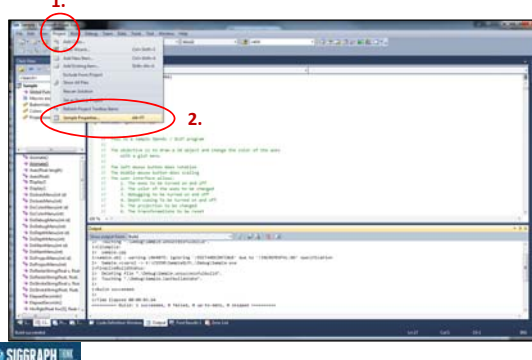

mb - July 24, 2020

Here's what that GLM folder looks like

mb - July 24, 2020

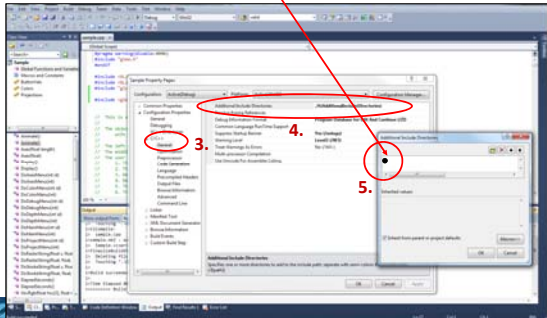

Telling Visual Studio about where the GLM folder is

mb - July 24, 2020

Telling Visual Studio about where the GLM folder is

A *period*, indicating that the **project folder** should also be searched when a **#include <xxx>** is encountered. If you put it somewhere else, enter that full or relative path instead.

mb - July 24, 2020

GLM in the Vulkan sample.cpp Program

```

#if UseMouse
{
    if (Scale < MINSCALE)
        Scale = MINSCALE;
    Matrices.uModelMatrix = glm::mat4( 1. ); // Identity
    Matrices.uModelMatrix = glm::rotate( Matrices.uModelMatrix, Yrot, glm::vec3( 0, 1, 0. ) );
    Matrices.uModelMatrix = glm::rotate( Matrices.uModelMatrix, Xrot, glm::vec3( 1, 0, 0. ) );
    Matrices.uModelMatrix = glm::scale( Matrices.uModelMatrix, glm::vec3(Scale,Scale,Scale) );
    // done this way, the Scale is applied first, then the Xrot, then the Yrot
}
else
{
    if ( ! Paused )
    {
        const glm::vec3 axis = glm::vec3( 0, 1, 0. );
        Matrices.uModelMatrix = glm::rotate( glm::mat4( 1. ), (float)glm::radians( 360*Time*SECONDS_PER_CYCLE ), axis );
    }
}

glm::vec3 eye( 0, 0, _EYEDIST );
glm::vec3 look( 0, 0, 0. );
glm::vec3 up( 0, 1, 0. );
Matrices.uViewMatrix = glm::lookAt( eye, look, up );


Matrices.uProjectionMatrix = glm::perspective( FOV, (double)Width/(double)Height, 0.1f, 1000.f );
Matrices.uProjectionMatrix[1][1] *= -1; // Vulkan's projected Y is inverted from OpenGL

Matrices.uNormalMatrix = glm::inverseTranspose( glm::mat3( Matrices.uModelMatrix ); // note: inverseTransform!

FIOSDataBuffer( MyMatrixUniformBuffer, (void *) &Matrices );

Misc.uTime = (float)Time;
Misc.uMode = Mode;
FIOSDataBuffer( MyMiscUniformBuffer, (void *) &Misc );

```



mb - July 24, 2020

How Does this Matrix Stuff Really Work?

145

$x' = Ax + By + Cz + D$
 $y' = Ex + Fy + Gz + H$
 $z' = Ix + Jy + Kz + L$

This is called a "Linear Transformation" because all of the coordinates are raised to the 1st power, that is, there are no $x^2, x^3,$ etc. terms.

Or, in matrix form:

x constant column
 y constant column
 z constant column
 constant column

x' producing row
 y' producing row
 z' producing row

$$\begin{bmatrix} x' \\ y' \\ z' \\ 1 \end{bmatrix} = \begin{bmatrix} A & B & C & D \\ E & F & G & H \\ I & J & K & L \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix}$$

SIGGRAPH 2000 mb - July 24, 2003

Transformation Matrices

146

Translation

$$\begin{bmatrix} x' \\ y' \\ z' \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & T_x \\ 0 & 1 & 0 & T_y \\ 0 & 0 & 1 & T_z \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix}$$

Rotation about X

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

Scaling

$$\begin{bmatrix} x' \\ y' \\ z' \\ 1 \end{bmatrix} = \begin{bmatrix} S_x & 0 & 0 & 0 \\ 0 & S_y & 0 & 0 \\ 0 & 0 & S_z & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix}$$

Rotation about Y

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

Rotation about Z

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

SIGGRAPH 2000 mb - July 24, 2003

How it Really Works :-)

147

$$\begin{bmatrix} \cos 90^\circ & \sin 90^\circ \\ -\sin 90^\circ & \cos 90^\circ \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

<http://xkcd.com>

SIGGRAPH 2000 mb - July 24, 2003

The Rotation Matrix for an Angle (θ) about an Arbitrary Axis (A_x, A_y, A_z)

148

$$[M] = \begin{bmatrix} A_x A_x + \cos \theta (1 - A_x A_x) & A_x A_y - \cos \theta (A_x A_y) - \sin \theta A_z & A_x A_z - \cos \theta (A_x A_z) + \sin \theta A_y \\ A_x A_y - \cos \theta (A_x A_y) + \sin \theta A_z & A_x A_x + \cos \theta (1 - A_x A_x) & A_x A_z - \cos \theta (A_x A_z) - \sin \theta A_y \\ A_x A_z - \cos \theta (A_x A_z) - \sin \theta A_y & A_x A_z - \cos \theta (A_x A_z) + \sin \theta A_y & A_x A_x + \cos \theta (1 - A_x A_x) \end{bmatrix}$$

For this to be correct, A must be a unit vector

SIGGRAPH 2000 mb - July 24, 2003

Compound Transformations

149

Q: Our rotation matrices only work around the origin? What if we want to rotate about an arbitrary point (A,B)?

A: We create more than one matrix.

Write it

$$\begin{bmatrix} x' \\ y' \\ z' \\ 1 \end{bmatrix} = \begin{bmatrix} 3 \\ 2 \\ 1 \end{bmatrix} \begin{bmatrix} T_{+A,+B} \\ R_\theta \\ T_{-A,-B} \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix}$$

Say it

SIGGRAPH 2000 mb - July 24, 2003

Matrix Multiplication is not Commutative

150


Rotate, then translate

Translate, then rotate

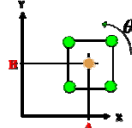
SIGGRAPH 2000 mb - July 24, 2003

Matrix Multiplication is Associative 151

$$\begin{pmatrix} x' \\ y' \\ z' \\ 1 \end{pmatrix} = \begin{pmatrix} T_{+A,+B} \\ \cdot \\ [R_\theta] \\ \cdot \\ [T_{-A,-B}] \end{pmatrix} \cdot \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix}$$

$$\begin{pmatrix} x' \\ y' \\ z' \\ 1 \end{pmatrix} = \underbrace{\begin{pmatrix} T_{+A,+B} \\ \cdot \\ [R_\theta] \\ \cdot \\ [T_{-A,-B}] \end{pmatrix}}_{\text{One matrix - the Current Transformation Matrix, or CTM}} \cdot \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix}$$


One Matrix to Rule Them All 152



$$\begin{pmatrix} x' \\ y' \\ z' \\ 1 \end{pmatrix} = \begin{pmatrix} T_{+A,+B} \\ \cdot \\ [R_\theta] \\ \cdot \\ [T_{-A,-B}] \end{pmatrix} \cdot \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix}$$


```

glm::mat4 Model = glm::mat4( 1. );
Model = glm::translate(Model, glm::vec3( A, B, 0. ) );
Model = glm::rotate(Model, thetaRadians, glm::vec3( Ax, Ay, Az ) );
Model = glm::translate(Model, glm::vec3( -A, -B, 0. ) );

glm::vec3 eye(0.,0.,EYEDIST);
glm::vec3 look(0.,0.,0.); glm::vec3 up(0.,1.,0.);
glm::mat4 View = glm::lookAt( eye, look, up );

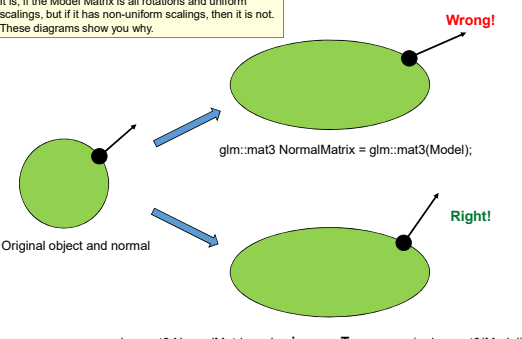
glm::mat4 Projection = glm::perspective( FOV, (double)Width/(double)Height, 0.1, 1000. );
Projection[1][1] *= -1.;

...
glm::mat3 Matrix = Projection * View * Model;
glm::mat3 NormalMatrix = glm::inverseTranspose( glm::mat3(Model) );
    
```



Why Isn't The Normal Matrix exactly the same as the Model Matrix? 153

It is, if the Model Matrix is all rotations and uniform scalings, but if it has non-uniform scalings, then it is not. These diagrams show you why.




Wrong!

glm::mat3 NormalMatrix = glm::mat3(Model);

Right!

glm::mat3 NormalMatrix = glm::inverseTranspose(glm::mat3(Model));




Vulkan.

Instancing

Mike Bailey
mjb@cs.oregonstate.edu

<http://cs.oregonstate.edu/~mjb/vulkan>

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License




Instancing – What and why? 155

- Instancing is the ability to draw the same object multiple times
- It uses all the same vertices and graphics pipeline each time
- It avoids the overhead of the program asking to have the object drawn again, letting the GPU/driver handle all of that

```
vkCmdDraw( CommandBuffers[nextImageIndex], vertexCount, instanceCount, firstVertex, firstInstance );
```

But, this will only get us multiple instances of identical objects drawn on top of each other. How can we make each instance look differently?

BTW, when not using instancing, be sure the **instanceCount** is 1, not 0!



Making each Instance look differently – Approach #1 156

Use the built-in vertex shader variable **gl_InstanceIndex** to define a unique display property, such as position or color.

gl_InstanceIndex starts at 0

In the vertex shader:


```


out vec3 vColor;
const int NUMINSTANCES = 16;
const float DELTA = 3.0;

float xdelta = DELTA * float( gl_InstanceIndex % 4 );
float ydelta = DELTA * float( gl_InstanceIndex / 4 );
vColor = vec3( 1., float( (1.+gl_InstanceIndex) / float( NUMINSTANCES ), 0. );

xdelta -= DELTA * sqrt( float(NUMINSTANCES) ) / 2.;
ydelta -= DELTA * sqrt( float(NUMINSTANCES) ) / 2.;
vec4 vertex = vec4( aVertex.xyz + vec3( xdelta, ydelta, 0. ), 1. );

gl_Position = PVM * vertex; // [p][v]*[m]
    
```





Making each Instance look differently -- Approach #2

Put the unique characteristics in a uniform buffer array and reference them

Still uses `gl_InstanceIndex`

In the vertex shader:

```

layout( std140, set = 3, binding = 0 ) uniform colorBuf
{
    vec3 uColors[1024];
} Colors;

out vec3 vColor;

...

int index = gl_InstanceIndex % 1024; // or "& 1023" - gives 0 - 1023
vColor = Colors.uColors[ index ];

vec4 vertex = ...

gl_Position = PVM * vertex; // [p]*[v]*[m]
    
```

Vulkan.

The Graphics Pipeline Data Structure

Mike Bailey
mjb@cs.oregonstate.edu

<http://cs.oregonstate.edu/~mjb/vulkan>

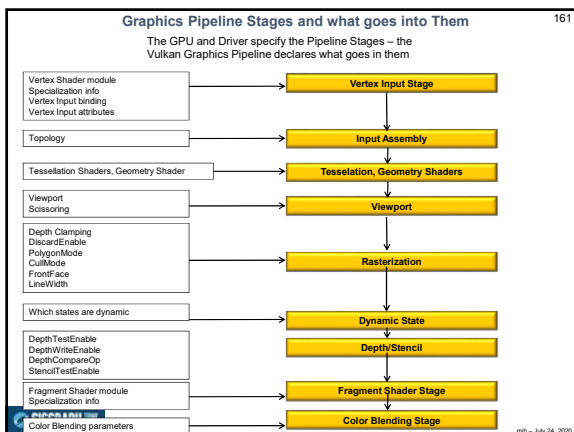
What is the Vulkan Graphics Pipeline?

Don't worry if this is too small to read - a larger version is coming up.

There is also a Vulkan Compute Pipeline Data Structure - we will get to that later.

Here's what you need to know:

- The Vulkan Graphics Pipeline is like what OpenGL would call "The State", or "The Context". It is a **data structure**.
- The Vulkan Graphics Pipeline is **not** the processes that OpenGL would call "the graphics pipeline".
- For the most part, the Vulkan Graphics Pipeline Data Structure is immutable - that is, once this combination of state variables is combined into a Pipeline, that Pipeline never gets changed. To make new combinations of state variables, create a new Graphics Pipeline.
- The shaders get compiled the rest of the way when their Graphics Pipeline gets created.



The First Step: Create the Graphics Pipeline Layout

The Graphics Pipeline Layout is fairly static. Only the layout of the Descriptor Sets and information on the Push Constants need to be supplied.

```

VkResult
Init4GraphicsPipelineLayout()
{
    VkResult result;

    VkPipelineLayoutCreateInfo
    vpclci = { VK_STRUCTURE_TYPE_PIPELINE_LAYOUT_CREATE_INFO,
              vpclci.pNext = nullptr,
              vpclci.flags = 0,
              vpclci.setLayoutCount = 4,
              vpclci.setLayouts = &DescriptorSetLayout[0],
              vpclci.pushConstantRangeCount = 0,
              vpclci.pushConstantRanges = (VkPushConstantRange*)nullptr };

    result = vkCreatePipelineLayout( LogicalDevice, IN &vpclci, ALLOCATOR, OUT &GraphicsPipelineLayout );

    return result;
}
    
```

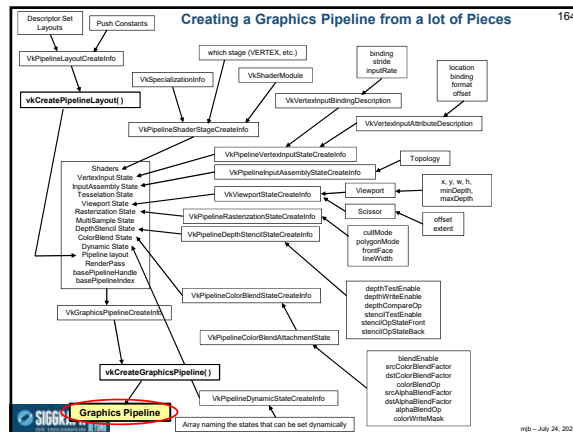
Let the Pipeline Layout know about the Descriptor Set and Push Constant layouts.

Why is this necessary? It is because the Descriptor Sets and Push Constants data structures have different sizes depending on how many of each you have. So, the exact structure of the Pipeline Layout depends on you telling Vulkan about the Descriptor Sets and Push Constants that you will be using.

A Pipeline Data Structure Contains the Following State Items: 163

- Pipeline Layout: Descriptor Sets, Push Constants
- Which Shaders to use
- Per-vertex input attributes: location, binding, format, offset
- Per-vertex input bindings: binding, stride, inputRate
- Assembly: topology
- **Viewport** x, y, w, h, minDepth, maxDepth
- **Scissoring**: x, y, w, h
- Rasterization: cullMode, polygonMode, frontFace, **lineWidth**
- Depth: depthTestEnable, depthWriteEnable, depthCompareOp
- Stencil: stencilTestEnable, stencilOpStateFront, stencilOpStateBack
- Blending: blendEnable, **srcColorBlendFactor**, **dstColorBlendFactor**, **colorBlendOp**, **srcAlphaBlendFactor**, **dstAlphaBlendFactor**, alphaBlendOp, colorWriteMask
- DynamicState: which states can be set dynamically (bound to the command buffer, outside the Pipeline)

Bold/Italics indicates that this state item can also be set with Dynamic State Variables



Creating a Typical Graphics Pipeline 165

```

VkResult
Init4GraphicsVertexFragmentPipeline(VkShaderModule vertexShader, VkShaderModule fragmentShader,
    VkPrimitiveTopology topology, OUT VkPipeline *pGraphicsPipeline )
{
    #ifdef ASSUMPTIONS
        vvb[0].inputRate = VK_VERTEX_INPUT_RATE_VERTEX;
        vprsc.depthClampEnable = VK_FALSE;
        vprsc.rasterizerDiscardEnable = VK_FALSE;
        vprsc.polygonMode = VK_POLYGON_MODE_FILL;
        vprsc.cullMode = VK_CULL_MODE_NONE; // best to do this because of the projectionMatrix[1][1] * -1;
        vprsc.frontFace = VK_FRONT_FACE_COUNTER_CLOCKWISE;
        vprsc.rasterizationSamples = VK_SAMPLE_COUNT_ONE_BIT;
        vpcbas.blendEnable = VK_FALSE;
        vpcbas.logicOpEnable = VK_FALSE;
        vpcssci.depthTestEnable = VK_TRUE;
        vpcssci.depthWriteEnable = VK_TRUE;
        vpcssci.depthCompareOp = VK_COMPARE_OP_LESS;
    #endif

    // These settings seem pretty typical to me. Let's write a simplified
    // Pipeline-creator that accepts Vertex and Fragment shader modules
    // and the topology, and always uses the settings in red above.
}
    
```

The Shaders to Use 166

```

VkPipelineShaderStageCreateInfo
vpsc[0].sType = VK_STRUCTURE_TYPE_PIPELINE_SHADER_STAGE_CREATE_INFO;
vpsc[0].pNext = nullptr;
vpsc[0].flags = 0;
vpsc[0].stage = VK_SHADER_STAGE_VERTEX_BIT;
vpsc[0].module = vertexShader;
vpsc[0].pName = "main";
vpsc[0].pSpecializationInfo = (VkSpecializationInfo *)nullptr;

// Shader BITS
VK_SHADER_STAGE_VERTEX_BIT
VK_SHADER_STAGE_TESSELLATION_CONTROL_BIT
VK_SHADER_STAGE_TESSELLATION_EVALUATION_BIT
VK_SHADER_STAGE_GEOMETRY_BIT
VK_SHADER_STAGE_FRAGMENT_BIT
VK_SHADER_STAGE_COMPUTE_BIT
VK_SHADER_STAGE_ALL_GRAPHICS
VK_SHADER_STAGE_ALL

vpsc[1].sType = VK_STRUCTURE_TYPE_PIPELINE_SHADER_STAGE_CREATE_INFO;
vpsc[1].pNext = nullptr;
vpsc[1].flags = 0;
vpsc[1].stage = VK_SHADER_STAGE_FRAGMENT_BIT;
vpsc[1].module = fragmentShader;
vpsc[1].pName = "main";
vpsc[1].pSpecializationInfo = (VkSpecializationInfo *)nullptr;

VkVertexInputBindingDescription
v vibd[0].binding = 0; // which binding this is
v vibd[0].stride = sizeof(struct vertex); // bytes between successive
v vibd[0].inputRate = VK_VERTEX_INPUT_RATE_VERTEX;

// Shader CHOICES
VK_VERTEX_INPUT_RATE_VERTEX
VK_VERTEX_INPUT_RATE_INSTANCE
    
```

Use one **vpsc[i]** array member per shader module you are using

Use one **vibd[i]** array member per vertex input array-of-structures you are using

Link in the Per-Vertex Attributes 167

```

VkVertexInputAttributeDescription
// A = vertex, normal, color, texture coordinates
vviad[0].location = 0; // location in the layout
vviad[0].binding = 0; // which binding description this is part of
vviad[0].format = VK_FORMAT_VEC3; // x, y, z
vviad[0].offset = offsetOf( struct vertex, position ); // 0

//def EXTRAS_DEFINED_AT_THE_TOP
// these are here for convenience and readability;
#define VK_FORMAT_VEC4 VK_FORMAT_R32G32B32A32_SFLOAT
#define VK_FORMAT_XYZW VK_FORMAT_R32G32B32A32_SFLOAT
#define VK_FORMAT_VEC3 VK_FORMAT_R32G32B32_SFLOAT
#define VK_FORMAT_STP VK_FORMAT_R32G32B32_SFLOAT
#define VK_FORMAT_XYZ VK_FORMAT_R32G32B32_SFLOAT
#define VK_FORMAT_VEC2 VK_FORMAT_R32G32_SFLOAT
#define VK_FORMAT_ST VK_FORMAT_R32G32_SFLOAT
#define VK_FORMAT_XY VK_FORMAT_R32G32_SFLOAT
#define VK_FORMAT_FLOAT VK_FORMAT_R32_SFLOAT
#define VK_FORMAT_S VK_FORMAT_R32_SFLOAT
#define VK_FORMAT_X VK_FORMAT_R32_SFLOAT
#endif

vviad[1].location = 1;
vviad[1].binding = 0;
vviad[1].format = VK_FORMAT_VEC3; // rx, ny, nz
vviad[1].offset = offsetOf( struct vertex, normal ); // 12

vviad[2].location = 2;
vviad[2].binding = 0;
vviad[2].format = VK_FORMAT_VEC3; // r, g, b
vviad[2].offset = offsetOf( struct vertex, color ); // 24

vviad[3].location = 3;
vviad[3].binding = 0;
vviad[3].format = VK_FORMAT_VEC2; // s, t
vviad[3].offset = offsetOf( struct vertex, texCoord ); // 36
    
```

Use one **vviad[i]** array member per element in the struct for the array-of-structures element you are using as vertex input

These are defined at the top of the sample code so that you don't need to use confusing image-looking formats for positions, normals, and tex coords

Link in the Per-Vertex Attributes 168

```

VkPipelineVertexInputStateCreateInfo
vvisci.sType = VK_STRUCTURE_TYPE_PIPELINE_VERTEX_INPUT_STATE_CREATE_INFO;
vvisci.pNext = nullptr;
vvisci.pVertexBindingDescriptions = vvb;
vvisci.pVertexAttributeDescriptions = vviad;

VkPipelineInputAssemblyStateCreateInfo
vpiasci.sType = VK_STRUCTURE_TYPE_PIPELINE_INPUT_ASSEMBLY_STATE_CREATE_INFO;
vpiasci.pNext = nullptr;
vpiasci.flags = 0;
vpiasci.topology = VK_PRIMITIVE_TOPOLOGY_TRIANGLE_LIST;

//def CHOICES
VK_PRIMITIVE_TOPOLOGY_POINT_LIST
VK_PRIMITIVE_TOPOLOGY_LINE_LIST
VK_PRIMITIVE_TOPOLOGY_TRIANGLE_LIST
VK_PRIMITIVE_TOPOLOGY_TRIANGLE_STRIP
VK_PRIMITIVE_TOPOLOGY_TRIANGLE_FAN
VK_PRIMITIVE_TOPOLOGY_LINE_LIST_WITH_ADJACENCY
VK_PRIMITIVE_TOPOLOGY_TRIANGLE_LIST_WITH_ADJACENCY
VK_PRIMITIVE_TOPOLOGY_TRIANGLE_STRIP_WITH_ADJACENCY
#endif

vpiasci.primitiveRestartEnable = VK_FALSE;

VkPipelineTessellationStateCreateInfo
vptsci.sType = VK_STRUCTURE_TYPE_PIPELINE_TESSELLATION_STATE_CREATE_INFO;
vptsci.pNext = nullptr;
vptsci.flags = 0;
vptsci.patchControlPoints = 0; // number of patch control points

VkPipelineGeometryStateCreateInfo
vpgsci.sType = VK_STRUCTURE_TYPE_PIPELINE_GEOMETRY_STATE_CREATE_INFO;
vpgsci.pNext = nullptr;
vpgsci.flags = 0;
    
```

used to describe the input vertex attributes

Declare the binding descriptions and attribute descriptions

Declare the vertex topology

Tessellation Shader info

Geometry Shader info

Options for vpiasci.topology

169

VK_PRIMITIVE_TOPOLOGY_POINT_LIST

VK_PRIMITIVE_TOPOLOGY_TRIANGLE_LIST

VK_PRIMITIVE_TOPOLOGY_LINE_LIST

VK_PRIMITIVE_TOPOLOGY_TRIANGLE_STRIP

VK_PRIMITIVE_TOPOLOGY_LINE_STRIP

VK_PRIMITIVE_TOPOLOGY_TRIANGLE_FAN

SIGGRAPH 2018 July 24, 2020

What is "Primitive Restart Enable"?

170

`vpiasci.primitiveRestartEnable = VK_FALSE;`

"Restart Enable" is used with:

- Indexed drawing.
- Triangle Fan and "Strip topologies"

If `vpiasci.primitiveRestartEnable` is `VK_TRUE`, then a special "index" indicates that the primitive should start over. This is more efficient than explicitly ending the current primitive and explicitly starting a new primitive of the same type.

```
typedef enum VkIndexType
{
    VK_INDEX_TYPE_UINT16 = 0, // 0 - 65,535
    VK_INDEX_TYPE_UINT32 = 1, // 0 - 4,294,967,295
} VkIndexType;
```

If your `VkIndexType` is `VK_INDEX_TYPE_UINT16`, then the special index is `0xffff`.
 If your `VkIndexType` is `VK_INDEX_TYPE_UINT32`, it is `0xffffffff`.

SIGGRAPH 2018 July 24, 2020

One Really Good use of Restart Enable is in Drawing Terrain Surfaces with Triangle Strips

171

Triangle Strip #0:
 Triangle Strip #1:
 Triangle Strip #2:
 ...

SIGGRAPH 2018 July 24, 2020

vkViewport, vkRect2D, vkPipelineViewportStateCreateInfo

172

```
VkViewport
{
    v.x = 0;
    v.y = 0;
    v.width = (float)Width;
    v.height = (float)Height;
    v.minDepth = 0.0f;
    v.maxDepth = 1.0f;
}

VkRect2D
{
    r.offset.x = 0;
    r.offset.y = 0;
    r.extent.width = Width;
    r.extent.height = Height;
}

VkPipelineViewportStateCreateInfo
{
    vpiasciType = VK_STRUCTURE_TYPE_PIPELINE_VIEWPORT_STATE_CREATE_INFO;
    vpiascipNext = nullptr;
    vpiasciFlags = 0;
    vpiasciViewportCount = 1;
    vpiasciViewports = &v;
    vpiasciScissorCount = 1;
    vpiasciScissors = &r;
}
```

vk (circled in red) - Declare the viewport information

vr (circled in red) - Declare the scissoring information

vpiasci (circled in red) - Group the viewport and scissor information together

SIGGRAPH 2018 July 24, 2020

What is the Difference Between Changing the Viewport and Changing the Scissoring?

173

Viewport:
 Viewporting operates on **vertices** and takes place right before the rasterizer. Changing the vertical part of the **viewport** causes the entire scene to get scaled (scrunched) into the viewport area.

Original Image

Scissoring:
 Scissoring operates on **fragments** and takes place right after the rasterizer. Changing the vertical part of the **scissor** causes the entire scene to get clipped where it falls outside the scissor area.

SIGGRAPH 2018 July 24, 2020

Setting the Rasterizer State

174

```
VkPipelineRasterizationStateCreateInfo
{
    vpiasciType = VK_STRUCTURE_TYPE_PIPELINE_RASTERIZATION_STATE_CREATE_INFO;
    vpiasciNext = nullptr;
    vpiasciFlags = 0;
    vpiasciDepthClampEnable = VK_FALSE;
    vpiasciRasterizerDiscardEnable = VK_FALSE;
    vpiasciPolygonMode = VK_POLYGON_MODE_FILL;
}

//def CHOICES
VK_POLYGON_MODE_FILL
VK_POLYGON_MODE_LINE
VK_POLYGON_MODE_POINT
//endif

vpiasciCullMode = VK_CULL_MODE_NONE; // recommend this because of the projMatrix[1][1] != -1;

//def CHOICES
VK_CULL_MODE_NONE
VK_CULL_MODE_FRONT_BIT
VK_CULL_MODE_BACK_BIT
VK_CULL_MODE_FRONT_AND_BACK_BIT
//endif

vpiasciFrontFace = VK_FRONT_FACE_COUNTER_CLOCKWISE;

//def CHOICES
VK_FRONT_FACE_COUNTER_CLOCKWISE
VK_FRONT_FACE_CLOCKWISE
//endif

vpiasciDepthBiasEnable = VK_FALSE;
vpiasciDepthBiasConstantFactor = 0.f;
vpiasciDepthBiasSlopeFactor = 0.f;
vpiasciDepthBiasSlopeFactor = 0.f;
vpiasciLineWidth = 1.f;
```

vpiasci (circled in red) - Declare information about how the rasterization will take place

SIGGRAPH 2018 July 24, 2020

What is "Depth Clamp Enable"?

`vprsci.depthClampEnable = VK_FALSE;`

Depth Clamp Enable causes the fragments that would normally have been discarded because they are closer to the viewer than the near clipping plane to instead get projected to the near clipping plane and displayed.

A good use for this is **Polygon Capping**:

The front of the polygon is clipped, revealing to the viewer that this is really a shell, not a solid

The gray area shows what would happen with depthClampEnable (except it would have been red).

SIGGRAPH 2019 | July 24, 2020

What is "Depth Bias Enable"?

`vprscib.depthBiasEnable = VK_FALSE;`
`vprscib.depthBiasConstantFactor = 0.f;`
`vprscib.depthBiasClamp = 0.f;`
`vprscib.depthBiasSlopeFactor = 0.f;`

Depth Bias Enable allows scaling and translation of the Z-depth values as they come through the rasterizer to avoid Z-fighting.

Z-fighting

SIGGRAPH 2019 | July 24, 2020

MultiSampling State

`VkPipelineMultisampleStateCreateInfo`

```

vprsci.sType = VK_STRUCTURE_TYPE_PIPELINE_MULTISAMPLE_STATE_CREATE_INFO;
vprsci.pNext = nullptr;
vprsci.flags = 0;
vprsci.rasterizationSamples = VK_SAMPLE_COUNT_1_BIT;
vprsci.sampleShadingEnable = VK_FALSE;
vprsci.minSampleShading = 0;
vprsci.pSampleMask = (VkSampleMask*)nullptr;
vprsci.alphaToCoverageEnable = VK_FALSE;
vprsci.alphaToOneEnable = VK_FALSE;
    
```

Declare information about how the multisampling will take place

We will discuss MultiSampling in a separate noteset.

SIGGRAPH 2019 | July 24, 2020

Color Blending State for each Color Attachment *

Create an array with one of these for each color buffer attachment. Each color buffer attachment can use different blending operations.

`VkPipelineColorBlendAttachmentState`

```

vpcbas.blendEnable = VK_FALSE;
vpcbas.srcColorBlendFactor = VK_BLEND_FACTOR_SRC_COLOR;
vpcbas.dstColorBlendFactor = VK_BLEND_FACTOR_ONE_MINUS_SRC_COLOR;
vpcbas.colorBlendOp = VK_BLEND_OP_ADD;
vpcbas.srcAlphaBlendFactor = VK_BLEND_FACTOR_ONE;
vpcbas.dstAlphaBlendFactor = VK_BLEND_FACTOR_ZERO;
vpcbas.alphaBlendOp = VK_BLEND_OP_ADD;
vpcbas.colorWriteMask = VK_COLOR_COMPONENT_R_BIT |
                       VK_COLOR_COMPONENT_G_BIT |
                       VK_COLOR_COMPONENT_B_BIT |
                       VK_COLOR_COMPONENT_A_BIT;
    
```

This controls blending between the output of each color attachment and its image memory:

$$Color_{new} = (1-\alpha) * Color_{existing} + \alpha * Color_{incoming}$$

$$0 \leq \alpha \leq 1.$$

*A "Color Attachment" is a framebuffer to be rendered into. You can have as many of these as you want.

SIGGRAPH 2019 | July 24, 2020

Raster Operations for each Color Attachment

`VkPipelineColorBlendStateCreateInfo`

```

vpcbsci.sType = VK_STRUCTURE_TYPE_PIPELINE_COLOR_BLEND_STATE_CREATE_INFO;
vpcbsci.pNext = nullptr;
vpcbsci.flags = 0;
vpcbsci.logicOpEnable = VK_FALSE;
vpcbsci.logicOp = VK_LOGIC_OP_COPY;

//def CHOICES
VK_LOGIC_OP_CLEAR
VK_LOGIC_OP_AND
VK_LOGIC_OP_AND_REVERSE
VK_LOGIC_OP_COPY
VK_LOGIC_OP_AND_INVERTED
VK_LOGIC_OP_NO_OP
VK_LOGIC_OP_XOR
VK_LOGIC_OP_OR
VK_LOGIC_OP_NOR
VK_LOGIC_OP_EQUIVALENT
VK_LOGIC_OP_INVERT
VK_LOGIC_OP_OR_REVERSE
VK_LOGIC_OP_COPY_INVERTED
VK_LOGIC_OP_OR_INVERTED
VK_LOGIC_OP_NAND
VK_LOGIC_OP_SET
//endif

vpcbsci.attachmentCount = 1;
vpcbsci.pAttachments = &vpcbas;
vpcbsci.blendConstants[0] = 0;
vpcbsci.blendConstants[1] = 0;
vpcbsci.blendConstants[2] = 0;
vpcbsci.blendConstants[3] = 0;
    
```

This controls blending between the output of the fragment shader and the input to the color attachments.

SIGGRAPH 2019 | July 24, 2020

Which Pipeline Variables can be Set Dynamically

Just used as an example in the Sample Code

```

VkDynamicState
//def CHOICES
VK_DYNAMIC_STATE_VIEWPORT -- vkCmdSetViewport()
VK_DYNAMIC_STATE_SCISSOR -- vkCmdSetScissor()
VK_DYNAMIC_STATE_LINE_WIDTH -- vkCmdSetLineWidth()
VK_DYNAMIC_STATE_DEPTH_BIAS -- vkCmdSetDepthBias()
VK_DYNAMIC_STATE_BLEND_CONSTANTS -- vkCmdSetBlendConstants()
VK_DYNAMIC_STATE_DEPTH_BOUNDS -- vkCmdSetDepthBounds()
VK_DYNAMIC_STATE_STENCIL_COMPARE_MASK -- vkCmdSetStencilCompareMask()
VK_DYNAMIC_STATE_STENCIL_WRITE_MASK -- vkCmdSetStencilWriteMask()
VK_DYNAMIC_STATE_STENCIL_REFERENCE -- vkCmdSetStencilReference()
//endif

VkPipelineDynamicStateCreateInfo
vprdsi.sType = VK_STRUCTURE_TYPE_PIPELINE_DYNAMIC_STATE_CREATE_INFO;
vprdsi.pNext = nullptr;
vprdsi.flags = 0;
vprdsi.dynamicStateCount = 0; // leave turned off for now
vprdsi.pDynamicStates = vds;
    
```

SIGGRAPH 2019 | July 24, 2020

The Stencil Buffer

181

Update

Depth

Stencil

Render

Here's how the Stencil Buffer works:

1. While drawing into the Render Buffer, you can write values into the Stencil Buffer at the same time.
2. While drawing into the Render Buffer, you can do arithmetic on values in the Stencil Buffer at the same time.
3. When drawing into the Render Buffer, you can write-protect certain parts of the Render Buffer based on values that are in the Stencil Buffer

SIGGRAPH 2018 mb - July 24, 2020

Using the Stencil Buffer to Create a Magic Lens

182

SIGGRAPH 2018 mb - July 24, 2020

Using the Stencil Buffer to Create a Magic Lens

183

1. Clear the SB = 0
2. Write protect the color buffer
3. Fill a square, setting SB = 1
4. Write-enable the color buffer
5. Draw the solids wherever SB == 0
6. Draw the wireframes wherever SB == 1

SIGGRAPH 2018 mb - July 24, 2020

Using the Stencil Buffer to Perform Polygon Capping

184

SIGGRAPH 2018 mb - July 24, 2020

Using the Stencil Buffer to Perform Polygon Capping

185

1. Clear the SB = 0
2. Draw the polygons, setting SB = ~ SB
3. Draw a large gray polygon across the entire scene wherever SB != 0

SIGGRAPH 2018 mb - July 24, 2020

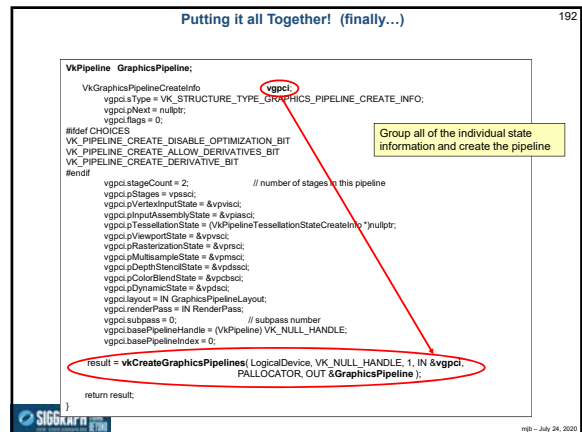
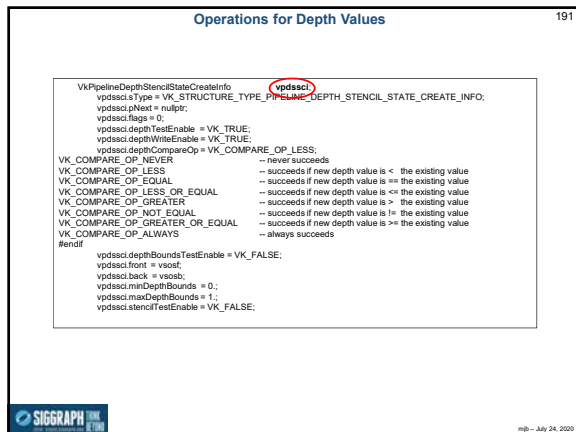
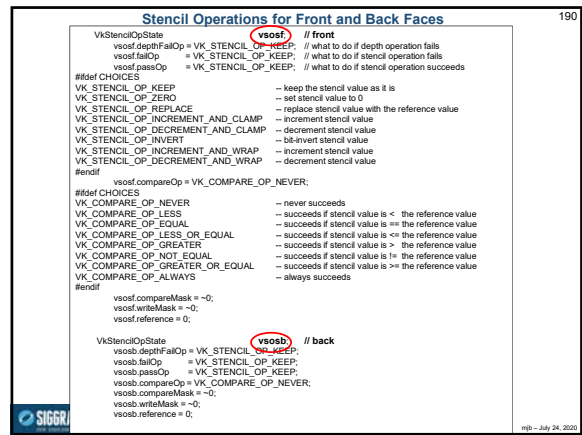
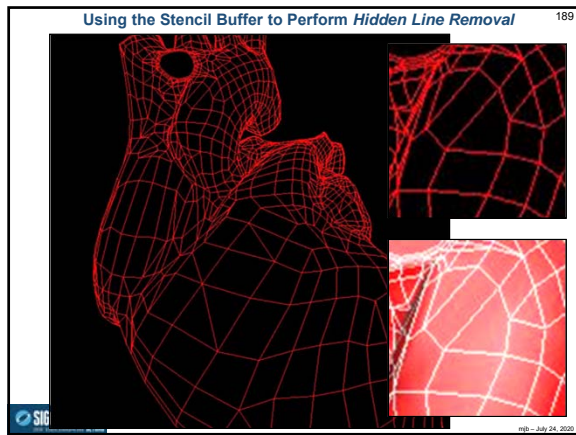
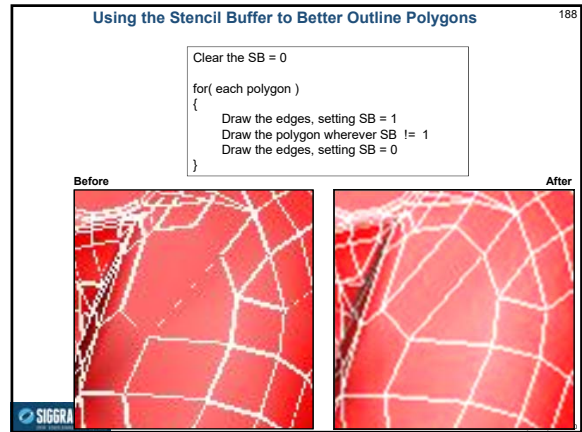
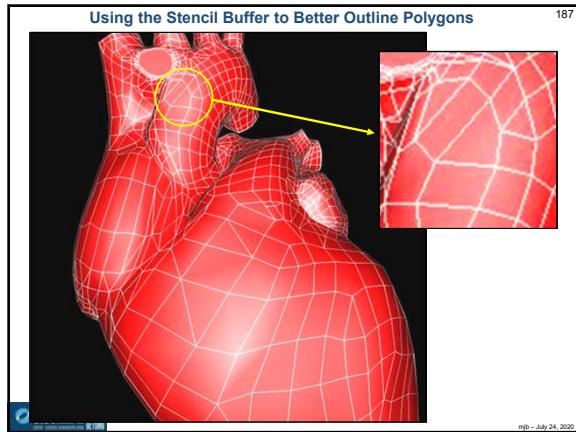
Outlining Polygons the Naïve Way

186

1. Draw the polygons
2. Draw the edges

Z-fighting


SIGGRAPH 2018 mb - July 24, 2020



193

Later on, we will Bind a Specific Graphics Pipeline Data Structure to the Command Buffer when Drawing

```
vkCmdBindPipeline( CommandBuffers[nextImageIndex],
VK_PIPELINE_BIND_POINT_GRAPHICS, GraphicsPipeline );
```



mjb - July 24, 2020

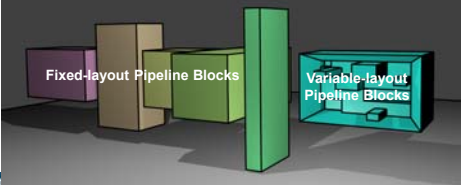

194

Sidebar: What is the Organization of the Pipeline Data Structure?

If you take a close look at the pipeline data structure creation information, you will see that almost all the pieces have a *fixed* size. For example, the viewport only needs 6 pieces of information – ever:

```
VkViewport          vw:
vw.x = 0;
vw.y = 0;
vw.width = (float)Width;
vw.height = (float)Height;
vw.minDepth = 0.0f;
vw.maxDepth = 1.0f;
```

There are two exceptions to this – the Descriptor Sets and the Push Constants. Each of these two can be almost any size, depending on what you allocate for them. So, I think of the Pipeline Data Structure as consisting of some fixed-layout blocks and 2 variable-layout blocks, like this:

mjb - July 24, 2020

195



Descriptor Sets

Mike Bailey
mjb@cs.oregonstate.edu



This work is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](https://creativecommons.org/licenses/by-nc-nd/4.0/)

<http://cs.oregonstate.edu/~mjb/vulkan>



mjb - July 24, 2020

196


In OpenGL

OpenGL puts all uniform data in the same "set", but with different binding numbers, so you can get at each one.

Each uniform variable gets updated one-at-a-time.

Wouldn't it be nice if we could update a collection of related uniform variables all at once, without having to update the uniform variables that are not related to this collection?

```
layout( std140, binding = 0 ) uniform mat4   uModelMatrix;
layout( std140, binding = 1 ) uniform mat4   uViewMatrix;
layout( std140, binding = 2 ) uniform mat4   uProjectionMatrix;
layout( std140, binding = 3 ) uniform mat3   uNormalMatrix;
layout( std140, binding = 4 ) uniform vec4   uLightPos;
layout( std140, binding = 5 ) uniform float  uTime;
layout( std140, binding = 6 ) uniform int    uMode;
layout(          binding = 7 ) uniform sampler2D uSampler;
```



mjb - July 24, 2020


197

What are Descriptor Sets?

Descriptor Sets are an intermediate data structure that tells shaders how to connect information held in GPU memory to groups of related uniform variables and texture sampler declarations in shaders. There are three advantages in doing things this way:

- Related uniform variables can be updated as a group, gaining efficiency.
- Descriptor Sets are activated when the Command Buffer is filled. Different values for the uniform buffer variables can be toggled by just swapping out the Descriptor Set that points to GPU memory, rather than re-writing the GPU memory.
- Values for the shaders' uniform buffer variables can be compartmentalized into what quantities change often and what change seldom (scene-level, model-level, draw-level), so that uniform variables need to be re-written no more often than is necessary.

```
for( each scene )
{
    Bind Descriptor Set #0
    for( each object )
    {
        Bind Descriptor Set #1
        for( each draw )
        {
            Bind Descriptor Set #2
            Do the drawing
        }
    }
}
```



mjb - July 24, 2020

198

Descriptor Sets


Our example will assume the following shader uniform variables:

```
// non-opaque must be in a uniform block:
layout( std140, set = 0, binding = 0 ) uniform matBuf
{
    mat4 uModelMatrix;
    mat4 uViewMatrix;
    mat4 uProjectionMatrix;
    mat3 uNormalMatrix;
} Matrices;

layout( std140, set = 1, binding = 0 ) uniform lightBuf
{
    vec4 uLightPos;
} Light;

layout( std140, set = 2, binding = 0 ) uniform miscBuf
{
    float uTime;
    int uMode;
} Misc;

layout( set = 3, binding = 0 ) uniform sampler2D uSampler;
```



mjb - July 24, 2020

Descriptor Sets

CPU:

Uniform data created in a C++ data structure

- Knows the CPU data structure
- Knows where the data starts
- Knows the data's size

GPU:

Uniform data in a "blob"

- Knows where the data starts
- Knows the data's size
- Doesn't know the CPU or GPU data structure

GPU:

Uniform data used in the shader

- Knows the shader data structure
- Doesn't know where each piece of data starts

```

struct matBuf
{
    glm::mat4 uModelMatrix;
    glm::mat4 uViewMatrix;
    glm::mat4 uProjectionMatrix;
    glm::mat3 uNormalMatrix;
};

struct lightBuf
{
    glm::vec4 uLightPos;
};

struct miscBuf
{
    float uTime;
    int uMode;
};

layout( std140, set = 0, binding = 0 ) uniform matBuf
{
    mat4 uModelMatrix;
    mat4 uViewMatrix;
    mat4 uProjectionMatrix;
    mat3 uNormalMatrix;
} Matrices;

layout( std140, set = 1, binding = 0 ) uniform lightBuf
{
    vec4 uLightPos;
} Light;

layout( std140, set = 2, binding = 0 ) uniform miscBuf
{
    float uTime;
    int uMode;
} Misc;

layout( set = 3, binding = 0 ) uniform sampler2D uSampler;
    
```

* "binary large object"

Step 1: Descriptor Set Pools

You don't allocate Descriptor Sets on the fly – that is too slow. Instead, you allocate a "pool" of Descriptor Sets and then pull from that pool later.

```

VkResult
Init13DescriptorSetPool()
{
    VkResult result;

    VkDescriptorPoolSize
    vdpse[0] type = VK_DESCRIPTOR_TYPE_UNIFORM_BUFFER;
    vdpse[0] descriptorCount = 1;
    vdpse[1] type = VK_DESCRIPTOR_TYPE_UNIFORM_BUFFER;
    vdpse[1] descriptorCount = 1;
    vdpse[2] type = VK_DESCRIPTOR_TYPE_UNIFORM_BUFFER;
    vdpse[2] descriptorCount = 1;
    vdpse[3] type = VK_DESCRIPTOR_TYPE_COMBINED_IMAGE_SAMPLER;
    vdpse[3] descriptorCount = 1;

    #if defined( VK_DESCRIPTOR_TYPE_SAMPLER )
    VK_DESCRIPTOR_TYPE_SAMPLER
    #endif
    #if defined( VK_DESCRIPTOR_TYPE_STORAGE_IMAGE )
    VK_DESCRIPTOR_TYPE_STORAGE_IMAGE
    #endif
    #if defined( VK_DESCRIPTOR_TYPE_UNIFORM_TEXEL_BUFFER )
    VK_DESCRIPTOR_TYPE_UNIFORM_TEXEL_BUFFER
    #endif
    #if defined( VK_DESCRIPTOR_TYPE_STORAGE_TEXEL_BUFFER )
    VK_DESCRIPTOR_TYPE_STORAGE_TEXEL_BUFFER
    #endif
    #if defined( VK_DESCRIPTOR_TYPE_STORAGE_BUFFER )
    VK_DESCRIPTOR_TYPE_STORAGE_BUFFER
    #endif
    #if defined( VK_DESCRIPTOR_TYPE_UNIFORM_BUFFER_DYNAMIC )
    VK_DESCRIPTOR_TYPE_UNIFORM_BUFFER_DYNAMIC
    #endif
    #if defined( VK_DESCRIPTOR_TYPE_INPUT_ATTACHMENT )
    VK_DESCRIPTOR_TYPE_INPUT_ATTACHMENT
    #endif
    #endif

    VkDescriptorPoolCreateInfo
    vdpci type = VK_STRUCTURE_TYPE_DESCRIPTOR_POOL_CREATE_INFO;
    vdpci.pNext = nullptr;
    vdpci.flags = 0;
    vdpci.maxSets = 4;
    vdpci.poolSizeCount = 4;
    vdpci.poolSizes = &vdpse[0];

    result = vkCreateDescriptorPool( LogicalDevice, IN &vdpci, PALLOCATOR_OUT &DescriptorPool);
    return result;
}
    
```

Step 2: Define the Descriptor Set Layouts

I think of Descriptor Set Layouts as a kind of "Rosetta Stone" that allows the Graphics Pipeline data structures to allocate room for the uniform variables and to access them.

```

VkResult
Init13DescriptorSetLayouts()
{
    VkResult result;

    // DS #0:
    VkDescriptorSetLayoutBinding
    MatrixSet[0] binding = 0;
    MatrixSet[0] descriptorType = VK_DESCRIPTOR_TYPE_UNIFORM_BUFFER;
    MatrixSet[0] descriptorCount = 1;
    MatrixSet[0] stageFlags = VK_SHADER_STAGE_VERTEX_BIT;
    MatrixSet[0] pImmutableSamplers = (VkSampler *)nullptr;

    // DS #1:
    VkDescriptorSetLayoutBinding
    LightSet[0] binding = 0;
    LightSet[0] descriptorType = VK_DESCRIPTOR_TYPE_UNIFORM_BUFFER;
    LightSet[0] descriptorCount = 1;
    LightSet[0] stageFlags = VK_SHADER_STAGE_VERTEX_BIT | VK_SHADER_STAGE_FRAGMENT_BIT;
    LightSet[0] pImmutableSamplers = (VkSampler *)nullptr;

    // DS #2:
    VkDescriptorSetLayoutBinding
    MiscSet[0] binding = 0;
    MiscSet[0] descriptorType = VK_DESCRIPTOR_TYPE_UNIFORM_BUFFER;
    MiscSet[0] descriptorCount = 1;
    MiscSet[0] stageFlags = VK_SHADER_STAGE_VERTEX_BIT | VK_SHADER_STAGE_FRAGMENT_BIT;
    MiscSet[0] pImmutableSamplers = (VkSampler *)nullptr;

    // DS #3:
    VkDescriptorSetLayoutBinding
    TexSamplerSet[0] binding = 0;
    TexSamplerSet[0] descriptorType = VK_DESCRIPTOR_TYPE_COMBINED_IMAGE_SAMPLER;
    TexSamplerSet[0] descriptorCount = 1;
    TexSamplerSet[0] stageFlags = VK_SHADER_STAGE_FRAGMENT_BIT;
    TexSamplerSet[0] pImmutableSamplers = (VkSampler *)nullptr;
    uniform sampler2D uSampler;
    vec4 rgba = texture( uSampler, vST );
}
    
```

Step 2: Define the Descriptor Set Layouts

205

```

VkDescriptorSetLayoutCreateInfo vdslo0;
vdslo0.sType = VK_STRUCTURE_TYPE_DESCRIPTOR_SET_LAYOUT_CREATE_INFO;
vdslo0.pNext = nullptr;
vdslo0.flags = 0;
vdslo0.bindingCount = 1;
vdslo0.pBindings = &MatrixSet[0];

VkDescriptorSetLayoutCreateInfo vdslc1;
vdslc1.sType = VK_STRUCTURE_TYPE_DESCRIPTOR_SET_LAYOUT_CREATE_INFO;
vdslc1.pNext = nullptr;
vdslc1.flags = 0;
vdslc1.bindingCount = 1;
vdslc1.pBindings = &LightSet[0];

VkDescriptorSetLayoutCreateInfo vdslc2;
vdslc2.sType = VK_STRUCTURE_TYPE_DESCRIPTOR_SET_LAYOUT_CREATE_INFO;
vdslc2.pNext = nullptr;
vdslc2.flags = 0;
vdslc2.bindingCount = 1;
vdslc2.pBindings = &MiscSet[0];

VkDescriptorSetLayoutCreateInfo vdslc3;
vdslc3.sType = VK_STRUCTURE_TYPE_DESCRIPTOR_SET_LAYOUT_CREATE_INFO;
vdslc3.pNext = nullptr;
vdslc3.flags = 0;
vdslc3.bindingCount = 1;
vdslc3.pBindings = &TexSamplerSet[0];

result = vkCreateDescriptorSetLayout(LogicalDevice, IN &vdslo0, PALLOCATOR, OUT &DescriptorSetLayouts[0]);
result = vkCreateDescriptorSetLayout(LogicalDevice, IN &vdslc1, PALLOCATOR, OUT &DescriptorSetLayouts[1]);
result = vkCreateDescriptorSetLayout(LogicalDevice, IN &vdslc2, PALLOCATOR, OUT &DescriptorSetLayouts[2]);
result = vkCreateDescriptorSetLayout(LogicalDevice, IN &vdslc3, PALLOCATOR, OUT &DescriptorSetLayouts[3]);

return result;
    
```

SIGGRAPH 2018

mb - July 24, 2020

Step 3: Include the Descriptor Set Layouts in a Graphics Pipeline Layout 206

```

VkResult
Init4GraphicsPipelineLayout()
{
    VkResult result;

    VkPipelineLayoutCreateInfo vplci;
    vplci.sType = VK_STRUCTURE_TYPE_PIPELINE_LAYOUT_CREATE_INFO;
    vplci.pNext = nullptr;
    vplci.flags = 0;
    vplci.setLayoutCount = 4;
    vplci.pSetLayouts = &DescriptorSetLayouts[0];
    vplci.pushConstantRangeCount = 0;
    vplci.pushConstantRanges = (VkPushConstantRange *)nullptr;

    result = vkCreatePipelineLayout(LogicalDevice, IN &vplci, PALLOCATOR, OUT &GraphicsPipelineLayout);

    return result;
}
    
```

SIGGRAPH 2018

mb - July 24, 2020

Step 4: Allocating the Memory for Descriptor Sets 207

```

graph TD
    DSP[DescriptorSetPool] --> VDSAI[VKDescriptorSetAllocateInfo]
    DSC[descriptorSetCount] --> VDSAI
    DSL[DescriptorSetLayouts] --> VDSAI
    VDSAI --> VAK[VKAllocateDescriptorSets()]
    VAK --> DS[Descriptor Set]
    
```

SIGGRAPH 2018

mb - July 24, 2020

Step 4: Allocating the Memory for Descriptor Sets 208

```

VkResult
Init3DescriptorSets()
{
    VkResult result;

    VkDescriptorSetAllocateInfo vdsai;
    vdsai.sType = VK_STRUCTURE_TYPE_DESCRIPTOR_SET_ALLOCATE_INFO;
    vdsai.pNext = nullptr;
    vdsai.descriptorPool = DescriptorPool;
    vdsai.descriptorSetCount = 4;
    vdsai.pSetLayouts = DescriptorSetLayouts;

    result = vkAllocateDescriptorSets(LogicalDevice, IN &vdsai, OUT &DescriptorSets[0]);
}
    
```

SIGGRAPH 2018

mb - July 24, 2020

Step 5: Tell the Descriptor Sets where their CPU Data is 209

<pre> VkDescriptorBufferInfo vdbi0; vdbi0.buffer = MyMatrixUniformBuffer.buffer; vdbi0.offset = 0; vdbi0.range = sizeof(Matrices); </pre>	<p>vdbi0;</p> <p>This struct identifies what buffer it owns and how big it is</p>
<pre> VkDescriptorBufferInfo vdbi1; vdbi1.buffer = MyLightUniformBuffer.buffer; vdbi1.offset = 0; vdbi1.range = sizeof(Light); </pre>	<p>vdbi1;</p> <p>This struct identifies what buffer it owns and how big it is</p>
<pre> VkDescriptorBufferInfo vdbi2; vdbi2.buffer = MyMiscUniformBuffer.buffer; vdbi2.offset = 0; vdbi2.range = sizeof(Misc); </pre>	<p>vdbi2;</p> <p>This struct identifies what buffer it owns and how big it is</p>
<pre> VkDescriptorImageInfo vdi0; vdi0.sampler = MyPuppyTexture.texSampler; vdi0.imageView = MyPuppyTexture.texImageView; vdi0.imageLayout = VK_IMAGE_LAYOUT_SHADER_READ_ONLY_OPTIMAL; </pre>	<p>vdi0;</p> <p>This struct identifies what texture sampler and image view it owns</p>

SIGGRAPH 2018

mb - July 24, 2020

Step 5: Tell the Descriptor Sets where their CPU Data is 210

```

VkWriteDescriptorSet vwdso;
vwdso.sType = VK_STRUCTURE_TYPE_WRITE_DESCRIPTOR_SET;
vwdso.pNext = nullptr;
vwdso.descriptorCount = 1;
vwdso.descriptorType = VK_DESCRIPTOR_TYPE_UNIFORM_BUFFER;
vwdso.pBufferInfo = IN &vdbi0;
vwdso.pImageInfo = (VkDescriptorImageInfo *)nullptr;
vwdso.pTexelBufferView = (VkBufferView *)nullptr;

// ds 1:
VkWriteDescriptorSet vwdso1;
vwdso1.sType = VK_STRUCTURE_TYPE_WRITE_DESCRIPTOR_SET;
vwdso1.pNext = nullptr;
vwdso1.descriptorCount = 1;
vwdso1.descriptorType = VK_DESCRIPTOR_TYPE_UNIFORM_BUFFER;
vwdso1.pBufferInfo = IN &vdbi1;
vwdso1.pImageInfo = (VkDescriptorImageInfo *)nullptr;
vwdso1.pTexelBufferView = (VkBufferView *)nullptr;
    
```

This struct links a Descriptor Set to the buffer it is pointing to

This struct links a Descriptor Set to the buffer it is pointing to

SIGGRAPH 2018

mb - July 24, 2020

Step 5: Tell the Descriptor Sets where their data is

```

VkWriteDescriptorSet wvds2;
// ds 2:
wvds2.sType = VK_STRUCTURE_TYPE_WRITE_DESCRIPTOR_SET;
wvds2.pNext = nullptr;
wvds2.dstSet = DescriptorSets[2];
wvds2.dstBinding = 0;
wvds2.dstArrayElement = 0;
wvds2.descriptorCount = 1;
wvds2.descriptorType = VK_DESCRIPTOR_TYPE_UNIFORM_BUFFER;
wvds2.pBufferInfo = IN &wvds2;
wvds2.pImageInfo = (VkDescriptorImageInfo *)nullptr;
wvds2.pTexelBufferView = (VkBufferView *)nullptr;

// ds 3:
VkWriteDescriptorSet wvds3;
wvds3.sType = VK_STRUCTURE_TYPE_WRITE_DESCRIPTOR_SET;
wvds3.pNext = nullptr;
wvds3.dstSet = DescriptorSets[3];
wvds3.dstBinding = 0;
wvds3.dstArrayElement = 0;
wvds3.descriptorCount = 1;
wvds3.descriptorType = VK_DESCRIPTOR_TYPE_COMBINED_IMAGE_SAMPLER;
wvds3.pBufferInfo = (VkDescriptorBufferInfo *)nullptr;
wvds3.pImageInfo = IN &wvds3;
wvds3.pTexelBufferView = (VkBufferView *)nullptr;

uint32_t copyCount = 0;

// this could have been done with one call and an array of VkWriteDescriptorSets:
vkUpdateDescriptorSets(LogicalDevice, 1, IN &wvds0, IN copyCount, (VkCopyDescriptorSet *)nullptr);
vkUpdateDescriptorSets(LogicalDevice, 1, IN &wvds1, IN copyCount, (VkCopyDescriptorSet *)nullptr);
vkUpdateDescriptorSets(LogicalDevice, 1, IN &wvds2, IN copyCount, (VkCopyDescriptorSet *)nullptr);
vkUpdateDescriptorSets(LogicalDevice, 1, IN &wvds3, IN copyCount, (VkCopyDescriptorSet *)nullptr);
    
```

This struct links a Descriptor Set to the buffer it is pointing to

This struct links a Descriptor Set to the image it is pointing to

SIGGRAPH 2019

Step 6: Include the Descriptor Set Layout when Creating a Graphics Pipeline

```

VkGraphicsPipelineCreateInfo vgpcci;
vgpcci.sType = VK_STRUCTURE_TYPE_GRAPHICS_PIPELINE_CREATE_INFO;
vgpcci.pNext = nullptr;
vgpcci.flags = 0;
#ifdef CHOICES
VK_PIPELINE_CREATE_DISABLE_OPTIMIZATION_BIT
VK_PIPELINE_CREATE_ALLOW_DERIVATIVES_BIT
VK_PIPELINE_CREATE_DERIVATIVE_BIT
#endif
vgpcci.stageCount = 2; // number of stages in this pipeline
vgpcci.pStages = vpssci;
vgpcci.pVertexInputState = &vpvsci;
vgpcci.pInputAssemblyState = &vpasci;
vgpcci.pTessellationState = (VkPipelineTessellationStateCreateInfo *)nullptr;
vgpcci.pViewportState = &vpvsci;
vgpcci.pRasterizationState = &vpvsci;
vgpcci.pMultisampleState = &vpmsci;
vgpcci.pDepthStencilState = &vpdssci;
vgpcci.pColorBlendState = &vpcbsci;
vgpcci.pDynamicState = &vpdsci;
vgpcci.layout = &GraphicsPipelineLayout;
vgpcci.renderPass = IN RenderPass; // subpass number
vgpcci.basePipelineHandle = (VkPipeline) VK_NULL_HANDLE;
vgpcci.basePipelineIndex = 0;

result = vkCreateGraphicsPipelines(LogicalDevice, VK_NULL_HANDLE, 1, IN &vgpcci, PALLOCATOR, OUT &GraphicsPipeline);
    
```

SIGGRAPH 2019

Step 7: Bind Descriptor Sets into the Command Buffer when Drawing

```

vkCmdBindDescriptorSets( CommandBuffers[nextImageIndex],
VK_PIPELINE_BIND_POINT_GRAPHICS, GraphicsPipelineLayout,
0, 4, DescriptorSets, 0, (uint32_t *)nullptr );
    
```

So, the Pipeline Layout contains the **structure** of the Descriptor Sets.
Any collection of Descriptor Sets that match that structure can be bound into that pipeline.

SIGGRAPH 2019

Sidebar: The Entire Collection of Descriptor Set Paths

- vkCreateDescriptorPool()** - Create the pool of Descriptor Sets for future use
- vkCreateDescriptorSetLayout()** - Describe a particular Descriptor Set layout and use it in a specific Pipeline layout
- vkCreatePipelineLayout()** - Describe a particular Descriptor Set layout and use it in a specific Pipeline layout
- vkAllocateDescriptorSets()** - Allocate memory for particular Descriptor Sets
- vkUpdateDescriptorSets()** - Tell a particular Descriptor Set where its CPU data is; Re-write CPU data into a particular Descriptor Set
- vkCmdBindDescriptorSets()** - Make a particular Descriptor Set "current" for rendering

SIGGRAPH 2019

Sidebar: Why Do Descriptor Sets Need to Provide Layout Information to the Pipeline Data Structure?

The pieces of the Pipeline Data Structure are fixed in size – with the exception of the Descriptor Sets and the Push Constants. Each of these two can be any size, depending on what you allocate for them. So, the Pipeline Data Structure needs to know how these two are configured before it can set its own total layout.

Think of the DS layout as being a particular-sized hole in the Pipeline Data Structure. Any data you have that matches this hole's shape and size can be plugged in there.

The Pipeline Data Structure


SIGGRAPH 2019

Sidebar: Why Do Descriptor Sets Need to Provide Layout Information to the Pipeline Data Structure?

Any set of data that matches the Descriptor Set Layout can be plugged in there.

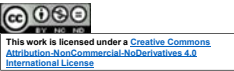
SIGGRAPH 2019

217




Textures

Mike Bailey
mjb@cs.oregonstate.edu



This work is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](https://creativecommons.org/licenses/by-nc-nd/4.0/).

<http://cs.oregonstate.edu/~mjb/vulkan>



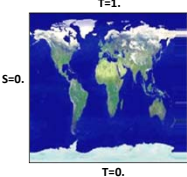
mjb - July 24, 2020


The Basic Idea 218

Texture mapping is a computer graphics operation in which a separate image, referred to as the **texture**, is stretched onto a piece of 3D geometry and follows it however it is transformed. This image is also known as a **texture map**.

Also, to prevent confusion, the texture pixels are not called **pixels**. A pixel is a dot in the final screen image. A dot in the texture image is called a **texture element**, or **texel**.

Similarly, to avoid terminology confusion, a texture's width and height dimensions are not called X and Y. They are called **S** and **T**. A texture map is not generally indexed by its actual resolution coordinates. Instead, it is indexed by a coordinate system that is resolution-independent. The left side is always **S=0**, the right side is **S=1**, the bottom is **T=0**, and the top is **T=1**. Thus, you do not need to be aware of the texture's resolution when you are specifying coordinates that point into it. Think of S and T as a measure of what fraction of the way you are into the texture.

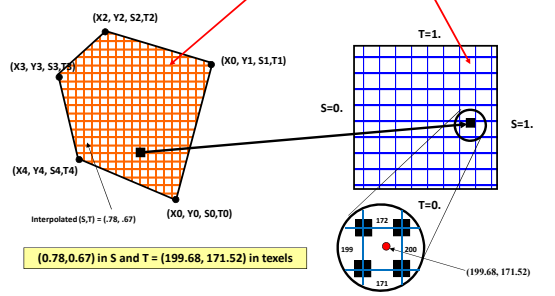





mjb - July 24, 2020

The Basic Idea 219

The mapping between the geometry of the 3D object and the S and T of the texture image works like this:



You specify an (s,t) pair at each vertex, along with the vertex coordinate. At the same time that the rasterizer is interpolating the coordinates, colors, etc. inside the polygon, it is also interpolating the (s,t) coordinates. Then, when it goes to draw each pixel, it uses that pixel's interpolated (s,t) to look up a color in the texture image.



mjb - July 24, 2020

In OpenGL terms: assigning an (s,t) to each vertex 220

Enable texture mapping:

```
glEnable( GL_TEXTURE_2D );
```

Draw your polygons, specifying s and t at each vertex:


```
glBegin( GL_POLYGON );
  glTexCoord2f( s0, t0 );
  glNormal3f( nx0, ny0, nz0 );
  glVertex3f( x0, y0, z0 );

  glTexCoord2f( s1, t1 );
  glNormal3f( nx1, ny1, nz1 );
  glVertex3f( x1, y1, z1 );

  ...
glEnd();
```

Disable texture mapping:

```
glDisable( GL_TEXTURE_2D );
```

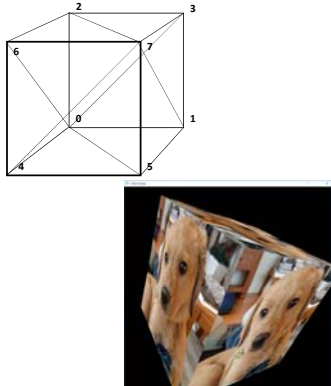



mjb - July 24, 2020

Triangles in an Array of Structures 221

```
struct vertex
{
  glm::vec3  position;
  glm::vec3  normal;
  glm::vec3  color;
  glm::vec2  texCoords;
};

struct vertex VertexData[] =
{
  // triangle 0-2-3:
  // vertex #0:
  { -1, -1, -1,
    { 0, 0, -1,
      { 1, 0 }
  },
  // vertex #2:
  { -1, 1, -1,
    { 0, 0, -1,
      { 1, 1 }
  },
  // vertex #3:
  { 1, 1, -1,
    { 0, 0, -1,
      { 0, 1 }
  },
};
```

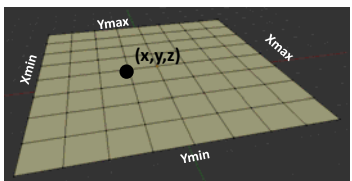





mjb - July 24, 2020

Using a Texture: How do you know what (s,t) to assign to each vertex? 222

The easiest way to figure out what s and t are at a particular vertex is to figure out what *fraction* across the object the vertex is living at. For a plane,



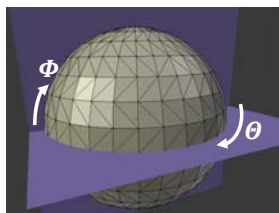
$$s = \frac{x - Xmin}{Xmax - Xmin} \quad t = \frac{y - Ymin}{Ymax - Ymin}$$



mjb - July 24, 2020


Using a Texture: How do you know what (s,t) to assign to each vertex? 223

Or, for a sphere,



$$s = \frac{\theta - (-\pi)}{2\pi} \quad t = \frac{\Phi - (-\frac{\pi}{2})}{\pi}$$

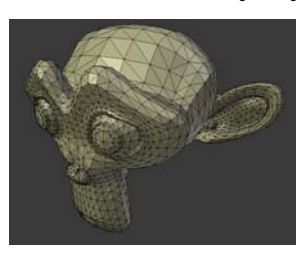
$$s = (\text{lng} + M_PI) / (2 * M_PI);$$

$$t = (\text{lat} + M_PI/2.) / M_PI;$$


SIGGRAPH 2014

Using a Texture: How do you know what (s,t) to assign to each vertex? 224


Uh-oh. Now what? Here's where it gets tougher....



$s = ?$ $t = ?$

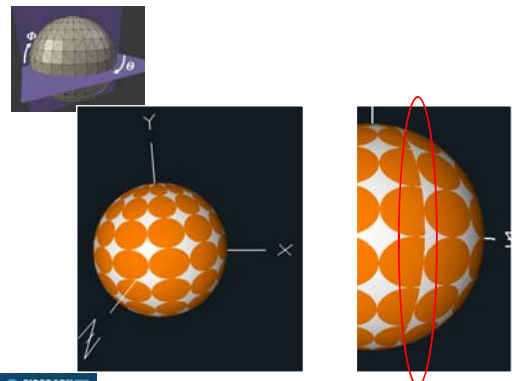
SIGGRAPH 2014

You really are at the mercy of whoever did the modeling... 225

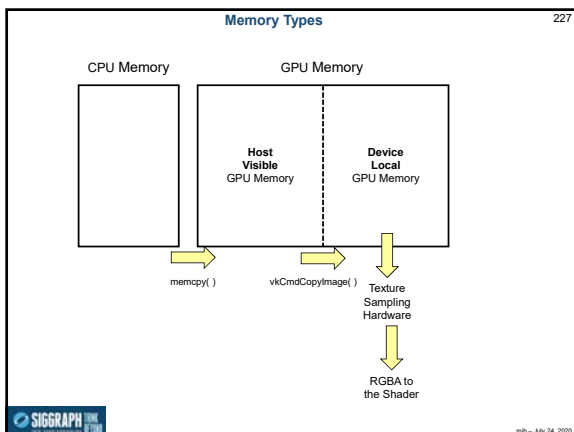


SIGGRAPH 2014

Be careful where s abruptly transitions from 1. back to 0. 226



SIGGRAPH 2014



Memory Types 228

NVIDIA Discrete Graphics:

11 Memory Types:
 Memory 0:
 Memory 1:
 Memory 2:
 Memory 3:
 Memory 4:
 Memory 5:
 Memory 6:
 Memory 7: DeviceLocal
 Memory 8: DeviceLocal
 Memory 9: HostVisible HostCoherent
 Memory 10: HostVisible HostCoherent HostCached

Intel Integrated Graphics:

3 Memory Types:
 Memory 0: DeviceLocal
 Memory 1: DeviceLocal HostVisible HostCoherent
 Memory 2: DeviceLocal HostVisible HostCoherent HostCached

SIGGRAPH 2014


```

235
void * gpuMemory;
vkMemory(LogicalDevice, vdm, 0, VK_WHOLE_SIZE, 0, OUT &gpuMemory);
// 0 and 0 = offset and memory map flags
if (vst.rowPitch == 4 * texWidth)
{
    memcpy(gpuMemory, void * texture, (size_t) textureSize);
}
else
{
    unsigned char *gpuBytes = (unsigned char *) gpuMemory;
    for (unsigned int y = 0; y < texHeight; ++y)
    {
        memcpy(gpuBytes + vst.rowPitch * y, texture + (4 * y * texWidth));
    }
}
vkUnmapMemory(LogicalDevice, vdm);
}

```

```

236
// This second [...] is to create the actual texture image
VkImageCreateInfo vci = { VK_STRUCTURE_TYPE_IMAGE_CREATE_INFO,
    vci.pNext = nullptr,
    vci.flags = 0,
    vci.imageType = VK_IMAGE_TYPE_2D,
    vci.format = VK_FORMAT_RGBA8_UNORM,
    vci.extent.width = texWidth,
    vci.extent.height = texHeight,
    vci.extent.depth = 1,
    vci.mipLevels = 1,
    vci.arrayLayers = 1,
    vci.samples = VK_SAMPLE_COUNT_1_BIT,
    vci.tiling = VK_TILING_OPTIMAL,
    vci.usage = VK_IMAGE_USAGE_TRANSFER_DST_BIT | VK_IMAGE_USAGE_SAMPLED_BIT,
    vci.sharingMode = VK_SHARING_MODE_EXCLUSIVE,
    vci.initialLayout = VK_IMAGE_LAYOUT_PREINITIALIZED,
    vci.queueFamilyIndex = 0,
    vci.queueFamilyIndex = (const uint32_t) "nullptr";
result = vkCreateImage(LogicalDevice, IN &vci, PALLOCATOR, OUT &textureImage); // allocated, but not filled
VkMemoryRequirements vmr;
vkGetImageMemoryRequirements(LogicalDevice, IN textureImage, IN &vmr);
// [VkDevice]
{
    fprintf(FDebug, "Texture vmr size = %ld", vmr.size);
    fprintf(FDebug, "Texture vmr alignment = %ld", vmr.alignment);
    fprintf(FDebug, "Texture vmr memoryTypeBits = %ld", vmr.memoryTypeBits);
    fflush(FDebug);
}
VkMemoryAllocateInfo vmi = { VK_STRUCTURE_TYPE_MEMORY_ALLOCATE_INFO,
    vmi.pNext = nullptr,
    vmi.allocateType = VK_MEMORY_TYPE_DEVICE_LOCAL, // because we want to sample from it
    vmi.memoryTypeIndex = FindMemoryTypeFromDeviceLocal() };
VkDeviceMemory vdm;
result = vkAllocateMemory(LogicalDevice, IN &vmi, PALLOCATOR, OUT &vdm);
result = vkBindImageMemory(LogicalDevice, IN textureImage, IN vdm, 0); // 0 = offset
}

```

```

237
// copy pixels from the staging image to the texture:
VkCommandBufferBeginInfo vcbi = { VK_STRUCTURE_TYPE_COMMAND_BUFFER_BEGIN_INFO,
    vcbi.pNext = nullptr,
    vcbi.flags = VK_COMMAND_BUFFER_USAGE_ONE_TIME_SUBMIT_BIT,
    vcbi.pInheritanceInfo = (VkCommandBufferInheritanceInfo *) "nullptr";
result = vkBeginCommandBuffer(TextureCommandBuffer, IN &vcbi);
// transition the staging buffer layout:
{
    VkImageSubresourceRange vrsr = { VK_IMAGE_ASPECT_COLOR_BIT,
        vrsr.baseMipLevel = 0,
        vrsr.levelCount = 1,
        vrsr.baseArrayLayer = 0,
        vrsr.layerCount = 1 };
    VkImageMemoryBarrier vmb = { VK_STRUCTURE_TYPE_IMAGE_MEMORY_BARRIER,
        vmb.pNext = nullptr,
        vmb.srcAccessMask = VK_IMAGE_LAYOUT_PREINITIALIZED,
        vmb.dstAccessMask = VK_IMAGE_LAYOUT_TRANSFER_DST_OPTIMAL,
        vmb.srcQueueFamilyIndex = VK_QUEUE_FAMILY_IGNORED,
        vmb.dstQueueFamilyIndex = VK_QUEUE_FAMILY_IGNORED,
        vmb.srcSubresourceRange = vrsr,
        vmb.dstSubresourceRange = vrsr,
        vmb.srcAccessMask = VK_ACCESS_HOST_WRITE_BIT,
        vmb.dstAccessMask = 0,
        vmb.subresourceRange = vrsr };
    vkCmdPipelineBarrier(TextureCommandBuffer,
        VK_PIPELINE_STAGE_HOST_BIT, VK_PIPELINE_STAGE_FRAGMENT_SHADER_BIT,
        0, (VkMemoryBarrier *) "nullptr",
        1, IN &vmb);
}
}

```

```

238
// transition the texture buffer layout:
{
    VkImageSubresourceRange vrsr = { VK_IMAGE_ASPECT_COLOR_BIT,
        vrsr.baseMipLevel = 0,
        vrsr.levelCount = 1,
        vrsr.baseArrayLayer = 0,
        vrsr.layerCount = 1 };
    VkImageMemoryBarrier vmb = { VK_STRUCTURE_TYPE_IMAGE_MEMORY_BARRIER,
        vmb.pNext = nullptr,
        vmb.dstAccessMask = VK_IMAGE_LAYOUT_TRANSFER_DST_OPTIMAL,
        vmb.srcAccessMask = VK_ACCESS_SHADER_READ_BIT,
        vmb.srcQueueFamilyIndex = VK_QUEUE_FAMILY_IGNORED,
        vmb.dstQueueFamilyIndex = VK_QUEUE_FAMILY_IGNORED,
        vmb.srcSubresourceRange = vrsr,
        vmb.dstSubresourceRange = vrsr,
        vmb.srcAccessMask = 0,
        vmb.dstAccessMask = VK_ACCESS_SHADER_READ_BIT,
        vmb.subresourceRange = vrsr };
    vkCmdPipelineBarrier(TextureCommandBuffer,
        VK_PIPELINE_STAGE_FRAGMENT_SHADER_BIT, VK_PIPELINE_STAGE_TRANSFER_BIT, 0,
        (VkMemoryBarrier *) "nullptr",
        0, (VkImageMemoryBarrier *) "nullptr",
        1, IN &vmb);
}
// now do the final image transfer:
VkImageSubresourceLayers vsl = { VK_IMAGE_ASPECT_COLOR_BIT,
    vsl.baseMipLevel = 0,
    vsl.layerCount = 1,
    vsl.level = 0 };
VkOffset3D vco = { 0, 0, 0 };
VkExtent3D vce = { texWidth, texHeight, 1 };
}

```

```

239
VkImageCopy vci = { VK_STRUCTURE_TYPE_IMAGE_COPY,
    vci.pNext = nullptr,
    vci.srcSubresource = vrsr,
    vci.dstSubresource = vrsr,
    vci.srcOffset = vco,
    vci.dstOffset = vco,
    vci.extent = vce };
vkCmdCopyImage(TextureCommandBuffer,
    stagingImage, VK_IMAGE_LAYOUT_TRANSFER_SRC_OPTIMAL,
    textureImage, VK_IMAGE_LAYOUT_TRANSFER_DST_OPTIMAL, 1, IN &vci);
}
}

```

```

240
// transition the texture buffer layout a second time:
{
    VkImageSubresourceRange vrsr = { VK_IMAGE_ASPECT_COLOR_BIT,
        vrsr.baseMipLevel = 0,
        vrsr.levelCount = 1,
        vrsr.baseArrayLayer = 0,
        vrsr.layerCount = 1 };
    VkImageMemoryBarrier vmb = { VK_STRUCTURE_TYPE_IMAGE_MEMORY_BARRIER,
        vmb.pNext = nullptr,
        vmb.dstAccessMask = VK_IMAGE_LAYOUT_TRANSFER_DST_OPTIMAL,
        vmb.srcAccessMask = VK_ACCESS_SHADER_READ_BIT,
        vmb.srcQueueFamilyIndex = VK_QUEUE_FAMILY_IGNORED,
        vmb.dstQueueFamilyIndex = VK_QUEUE_FAMILY_IGNORED,
        vmb.srcSubresourceRange = vrsr,
        vmb.dstSubresourceRange = vrsr,
        vmb.srcAccessMask = 0,
        vmb.dstAccessMask = VK_ACCESS_SHADER_READ_BIT,
        vmb.subresourceRange = vrsr };
    vkCmdPipelineBarrier(TextureCommandBuffer,
        VK_PIPELINE_STAGE_FRAGMENT_SHADER_BIT, VK_PIPELINE_STAGE_FRAGMENT_SHADER_BIT, 0,
        (VkMemoryBarrier *) "nullptr",
        0, (VkImageMemoryBarrier *) "nullptr",
        1, IN &vmb);
}
result = vkEndCommandBuffer(TextureCommandBuffer);
VkSubmitInfo vsi = { VK_STRUCTURE_TYPE_SUBMIT_INFO,
    vsi.pNext = nullptr,
    vsi.commandBufferCount = 1,
    vsi.pCommandBuffers = &TextureCommandBuffer,
    vsi.waitSemaphoreCount = 0,
    vsi.pWaitSemaphores = (VkSemaphore *) "nullptr",
    vsi.signalSemaphoreCount = 0,
    vsi.pSignalSemaphores = (VkSemaphore *) "nullptr",
    vsi.pWaitDstStageMask = (VkPipelineStageFlags *) "nullptr";
result = vkQueueSubmit(Queue, 1, IN &vsi, VK_NULL_HANDLE);
result = vkQueueWaitIdle(Queue);
}

```

241

```

// create an image view for the texture image:
// (an "image view" is used to indirectly access an image)

VkImageSubresourceRange
vkr.aspectMask = VK_IMAGE_ASPECT_COLOR_BIT;
vkr.baseMipLevel = 0;
vkr.levelCount = 1;
vkr.baseArrayLayer = 0;
vkr.layerCount = 1;

VkImageViewCreateInfo
vkciv.sType = VK_STRUCTURE_TYPE_IMAGE_VIEW_CREATE_INFO;
vkciv.pNext = nullptr;
vkciv.flags = 0;
vkciv.image = textureImage;
vkciv.viewType = VK_IMAGE_VIEW_TYPE_2D;
vkciv.format = VK_FORMAT_R8G8B8A8_UNORM;
vkciv.components.r = VK_COMPONENT_SWIZZLE_R;
vkciv.components.g = VK_COMPONENT_SWIZZLE_G;
vkciv.components.b = VK_COMPONENT_SWIZZLE_B;
vkciv.components.a = VK_COMPONENT_SWIZZLE_A;
vkciv.subresourceRange = vkr;

result = vkCreateImageView(LogicalDevice, IN &vkciv, PALLOCATOR, OUT &MyTexture->imageView);
return result;
    
```

Note that, at this point, the Staging Buffer is no longer needed, and can be destroyed.

SIGGRAPH 2019

242

Reading in a Texture from a BMP File

```

typedef struct MyTexture
{
    uint32_t width;
    uint32_t height;
    VkImage vkImage;
    VkImageView vkImageView;
    VkSampler vkSampler;
    VkDeviceMemory vdm;
} MyTexture;

MyTexture MyPuppyTexture;
    
```

```

result = Init06TextureBufferAndFillFromBmpFile ("puppy.bmp", &MyTexturePuppy);
Init06TextureSampler( &MyPuppyTexture.texSampler );
    
```

This function can be found in the `sample.cpp` file. The BMP file needs to be created by something that writes uncompressed 24-bit color BMP files, or was converted to the uncompressed BMP format by a tool such as ImageMagick's `convert`, Adobe *Photoshop*, or GNU's *GIMP*.

SIGGRAPH 2019

243

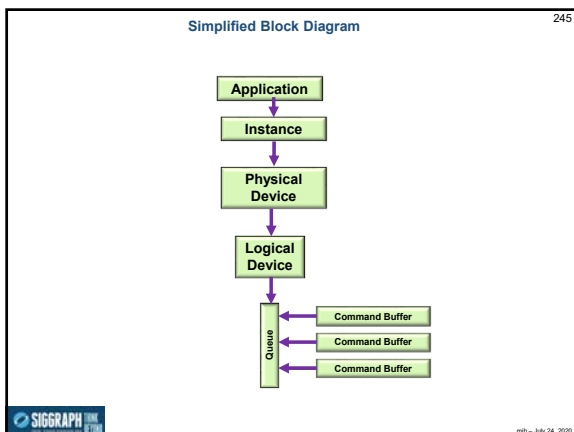
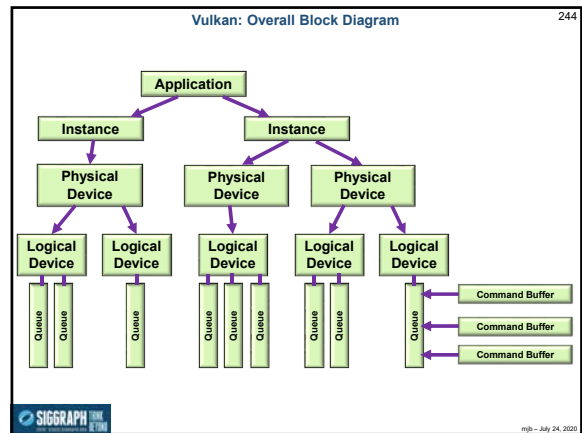
Queues and Command Buffers

Mike Bailey
mjb@cs.oregonstate.edu

<http://cs.oregonstate.edu/~mjb/vulkan>

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License

SIGGRAPH 2019



246

Vulkan Queues and Command Buffers

- Graphics commands are recorded in command buffers, e.g., `vkCmdDoSomething(cmdBuffer, ...)`;
- You can have as many simultaneous Command Buffers as you want
- Each command buffer can be filled from a different thread
- Command Buffers record commands, but no work takes place until a Command Buffer is submitted to a Queue
- We don't create Queues – the Logical Device has them already
- Each Queue belongs to a Queue Family
- We don't create Queue Families – the Physical Device already has them

SIGGRAPH 2019

Querying what Queue Families are Available 247


```

uint32_t count;
VkGetPhysicalDeviceQueueFamilyProperties( IN PhysicalDevice, &count, OUT (VkQueueFamilyProperties *) nullptr );
VkQueueFamilyProperties *vqfp = new VkQueueFamilyProperties[ count ];
VkGetPhysicalDeviceQueueFamilyProperties( PhysicalDevice, &count, OUT &vqfp );

for( unsigned int i = 0; i < count; i++ )
{
    fprintf( FpDebug, "tfid: Queue Family Count = %zd ; ", i, vqfp[i].queueCount );
    if ( ( vqfp[i].queueFlags & VK_QUEUE_GRAPHICS_BIT ) != 0 ) fprintf( FpDebug, " Graphics" );
    if ( ( vqfp[i].queueFlags & VK_QUEUE_COMPUTE_BIT ) != 0 ) fprintf( FpDebug, " Compute" );
    if ( ( vqfp[i].queueFlags & VK_QUEUE_TRANSFER_BIT ) != 0 ) fprintf( FpDebug, " Transfer" );
    fprintf( FpDebug, "\n" );
}
    
```

Found 3 Queue Families:

- 0: Queue Family Count = 16 : Graphics Compute Transfer
- 1: Queue Family Count = 1 : Transfer
- 2: Queue Family Count = 8 : Compute



mb - July 24, 2020


Similarly, we Can Write a Function that Finds the Proper Queue Family 248

```

int
FindQueueFamilyThatDoesGraphics( )
{
    uint32_t count = -1;
    vkGetPhysicalDeviceQueueFamilyProperties( IN PhysicalDevice, OUT &count, OUT (VkQueueFamilyProperties *) nullptr );

    VkQueueFamilyProperties *vqfp = new VkQueueFamilyProperties[ count ];
    vkGetPhysicalDeviceQueueFamilyProperties( IN PhysicalDevice, IN &count, OUT vqfp );

    for( unsigned int i = 0; i < count; i++ )
    {
        if ( ( vqfp[i].queueFlags & VK_QUEUE_GRAPHICS_BIT ) != 0 )
            return i;
    }
    return -1;
}
    
```



mb - July 24, 2020

Creating a Logical Device Needs to Know Queue Family Information 249

```



float *queuePriorities[] =
{
    1. // one entry per queueCount
};

VkDeviceQueueCreateInfo vdcqi[];
vdcqi[0].sType = VK_STRUCTURE_TYPE_QUEUE_CREATE_INFO;
vdcqi[0].pNext = nullptr;
vdcqi[0].flags = 0;
vdcqi[0].queueFamilyIndex = FindQueueFamilyThatDoesGraphics();
vdcqi[0].queueCount = 1;
vdcqi[0].queuePriorities = (float *) queuePriorities;

VkDeviceCreateInfo vdc;
vdc.sType = VK_STRUCTURE_TYPE_DEVICE_CREATE_INFO;
vdc.pNext = nullptr;
vdc.flags = 0;
vdc.queueCreateInfoCount = 1; // # of device queues wanted
vdc.queueCreateInfos = IN &vdcqi[0]; // array of VkDeviceQueueCreateInfo's
vdc.enabledLayerCount = sizeof(myDeviceLayers) / sizeof(char *);
vdc.ppEnabledLayerNames = myDeviceLayers;
vdc.enabledExtensionCount = sizeof(myDeviceExtensions) / sizeof(char *);
vdc.ppEnabledExtensionNames = myDeviceExtensions;
vdc.pEnabledFeatures = IN &PhysicalDeviceFeatures; // already created

result = vkCreateLogicalDevice( PhysicalDevice, IN &vdc, PALLOCATOR, OUT &LogicalDevice );

VkQueue Queue;
uint32_t queueFamilyIndex = FindQueueFamilyThatDoesGraphics();
uint32_t queueIndex = 0;
result = vkCreateDeviceQueue( LogicalDevice, queueFamilyIndex, queueIndex, OUT &Queue );
    
```

mb - July 24, 2020

Creating the Command Pool as part of the Logical Device 250

```


VkResult
Init06CommandPool( )
{
    VkResult result;

    VkCommandPoolCreateInfo vcp;
    vcp.sType = VK_STRUCTURE_TYPE_COMMAND_POOL_CREATE_INFO;
    vcp.pNext = nullptr;
    vcp.flags = VK_COMMAND_POOL_CREATE_RESET_COMMAND_BUFFER_BIT | VK_COMMAND_POOL_CREATE_TRANSIENT_BIT;

    #ifdef CHOICES
    VK_COMMAND_POOL_CREATE_TRANSIENT_BIT
    VK_COMMAND_POOL_CREATE_RESET_COMMAND_BUFFER_BIT
    #endif
    vcp.queueFamilyIndex = FindQueueFamilyThatDoesGraphics();

    result = vkCreateCommandPool( LogicalDevice, IN &vcp, PALLOCATOR, OUT &CommandPool );

    return result;
}
    
```



mb - July 24, 2020

Creating the Command Buffers 251



```

VkResult
Init06CommandBuffers( )
{
    VkResult result;

    // allocate 2 command buffers for the double-buffered rendering;
    {
        VkCommandBufferAllocateInfo vcbal;
        vcbal.sType = VK_STRUCTURE_TYPE_COMMAND_BUFFER_ALLOCATE_INFO;
        vcbal.pNext = nullptr;
        vcbal.commandPool = CommandPool;
        vcbal.level = VK_COMMAND_BUFFER_LEVEL_PRIMARY;
        vcbal.commandBufferCount = 2; // 2, because of double-buffering
        result = vkAllocateCommandBuffers( LogicalDevice, IN &vcbal, OUT &CommandBuffers[0] );
    }

    // allocate 1 command buffer for the transferring pixels from a staging buffer to a texture buffer;
    {
        VkCommandBufferAllocateInfo vcbal;
        vcbal.sType = VK_STRUCTURE_TYPE_COMMAND_BUFFER_ALLOCATE_INFO;
        vcbal.pNext = nullptr;
        vcbal.commandPool = CommandPool;
        vcbal.level = VK_COMMAND_BUFFER_LEVEL_PRIMARY;
        vcbal.commandBufferCount = 1;
        result = vkAllocateCommandBuffers( LogicalDevice, IN &vcbal, OUT &TextureCommandBuffer );
    }

    return result;
}
    
```

mb - July 24, 2020

Beginning a Command Buffer – One per Image 252

```

VkSemaphoreCreateInfo vsc;
vsc.sType = VK_STRUCTURE_TYPE_SEMAPHORE_CREATE_INFO;
vsc.pNext = nullptr;
vsc.flags = 0;

VkSemaphore imageReadySemaphore;
result = vkCreateSemaphore( LogicalDevice, IN &vsc, PALLOCATOR, OUT &imageReadySemaphore );


uint32_t nextImageIndex;
vkAcquireNextImageKHR( LogicalDevice, IN SwapChain, IN UINT64_MAX, IN imageReadySemaphore, IN VK_NULL_HANDLE, OUT &nextImageIndex );

VkCommandBufferBeginInfo vcbbi;
vcbbi.sType = VK_STRUCTURE_TYPE_COMMAND_BUFFER_BEGIN_INFO;
vcbbi.pNext = nullptr;
vcbbi.flags = VK_COMMAND_BUFFER_USAGE_ONE_TIME_SUBMIT_BIT;
vcbbi.pInheritanceInfo = (VkCommandBufferInheritanceInfo *) nullptr;

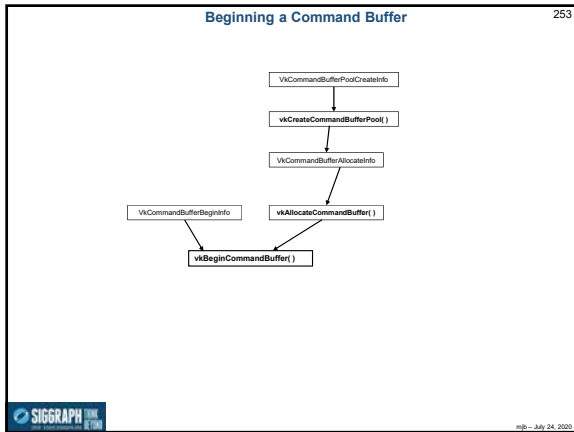
result = vkBeginCommandBuffer( CommandBuffers[nextImageIndex], IN &vcbbi );

...

vkEndCommandBuffer( CommandBuffers[nextImageIndex] );
    
```



mb - July 24, 2020



These are the Commands that could be entered into the Command Buffer, I 254

```

vkCmdBeginQuery( commandBuffer, flags );
vkCmdBeginRenderPass( commandBuffer, const contents );
vkCmdBindDescriptorSets( commandBuffer, pDynamicOffsets );
vkCmdBindIndexBuffer( commandBuffer, indexType );
vkCmdBindPipeline( commandBuffer, pipeline );
vkCmdBindVertexBuffers( commandBuffer, firstBinding, bindingCount, const pOffsets );
vkCmdBindImage( commandBuffer, filter );
vkCmdClearColor( commandBuffer, pRanges );
vkCmdClearDepthStencilImage( commandBuffer, pRanges );
vkCmdCopyBuffer( commandBuffer, pRegions );
vkCmdCopyBufferToImage( commandBuffer, pRegions );
vkCmdCopyImage( commandBuffer, pRegions );
vkCmdCopyImageToBuffer( commandBuffer, pRegions );
vkCmdCopyQueryPoolResults( commandBuffer, flags );
vkCmdDebugMarkerBeginEXT( commandBuffer, pMarkerInfo );
vkCmdDebugMarkerEndEXT( commandBuffer );
vkCmdDispatch( commandBuffer, groupCountX, groupCountY, groupCountZ );
vkCmdDispatchIndirect( commandBuffer, offset );
vkCmdDraw( commandBuffer, vertexCount, instanceCount, firstVertex, firstInstance );
vkCmdDrawIndexed( commandBuffer, indexCount, instanceCount, firstIndex, int32_1_vertexOffset, firstInstance );
vkCmdDrawIndexedIndirect( commandBuffer, stride );
vkCmdDrawIndexedIndirectCountAMD( commandBuffer, stride );
vkCmdDrawIndirect( commandBuffer, stride );
vkCmdDrawIndirectCountAMD( commandBuffer, stride );
vkCmdEndQuery( commandBuffer, query );
vkCmdEndRenderPass( commandBuffer );
vkCmdExecuteCommands( commandBuffer, commandBufferCount, const pCommandBuffers );
  
```

These are the Commands that could be entered into the Command Buffer, II 255

```

vkCmdFillBuffer( commandBuffer, dstBuffer, dstOffset, size, data );
vkCmdNextSubpass( commandBuffer, contents );
vkCmdPipelineBarrier( commandBuffer, srcStageMask, dstStageMask, dependencyFlags, memoryBarrierCount, pMemoryBarriers,
bufferMemoryBarrierCount, pBufferMemoryBarriers, imageMemoryBarrierCount, pImageMemoryBarriers );
vkCmdProcessCommandsNVX( commandBuffer, pProcessCommandIDs );
vkCmdPushConstants( commandBuffer, layout, stageFlags, offset, size, pValues );
vkCmdPushDescriptorSetKHR( commandBuffer, pipelineBindPoint, layout, set, descriptorUpdateTemplate, pDescriptorWrites );
vkCmdPushDescriptorSetWithTemplateKHR( commandBuffer, descriptorUpdateTemplate, layout, set, pData );
vkCmdResetEvent( commandBuffer, pEvent );
vkCmdResetEvent( commandBuffer, event, stageMask );
vkCmdResetQueryPool( commandBuffer, queryPool, firstQuery, queryCount );
vkCmdResetImage( commandBuffer, srcImage, srcImageLayout, dstImage, dstImageLayout, regionCount, pRegions );
vkCmdSetBlendConstants( commandBuffer, blendConstants[4] );
vkCmdSetDepthBias( commandBuffer, depthBiasConstantFactor, depthBiasClamp, depthBiasSlopeFactor );
vkCmdSetDepthBounds( commandBuffer, minDepthBounds, maxDepthBounds );
vkCmdSetDeviceMaskNVX( commandBuffer, deviceMask );
vkCmdSetDiscardRectangleEXT( commandBuffer, firstDiscardRectangle, discardRectangleCount, pDiscardRectangles );
vkCmdSetEvent( commandBuffer, event, stageMask );
vkCmdSetLineWidth( commandBuffer, lineWidth );
vkCmdSetScissor( commandBuffer, firstScissor, scissorCount, pScissors );
vkCmdSetStencilCompareMask( commandBuffer, faceMask, compareMask );
vkCmdSetStencilReference( commandBuffer, faceMask, reference );
vkCmdSetStencilWriteMask( commandBuffer, faceMask, writeMask );
vkCmdSetViewport( commandBuffer, firstViewport, viewportCount, pViewports );
vkCmdSetViewportWScalingNV( commandBuffer, firstViewport, viewportCount, pViewportWScalings );
vkCmdUpdateBuffer( commandBuffer, dstBuffer, dstOffset, dataSize, pData );
vkCmdWaitEvents( commandBuffer, eventCount, pEvents, srcStageMask, dstStageMask, memoryBarrierCount, pMemoryBarriers,
bufferMemoryBarrierCount, pBufferMemoryBarriers, imageMemoryBarrierCount, pImageMemoryBarriers );
vkCmdWriteTimestamp( commandBuffer, pipelineStage, queryPool, query );
  
```

These are the Commands that could be entered into the Command Buffer, III 256

```

VkResult
RenderScene()
{
    VkResult result;
    VkSemaphoreCreateInfo vsci;
    vsci.sType = VK_STRUCTURE_TYPE_SEMAPHORE_CREATE_INFO;
    vsci.pNext = nullptr;
    vsci.flags = 0;

    VkSemaphore imageReadySemaphore;
    result = vkCreateSemaphore( LogicalDevice, IN &vsci, PALLOCATOR, OUT &imageReadySemaphore );

    uint32_t nextImageIndex;
    vkAcquireNextImageKHR( LogicalDevice, IN SwapChain, IN uint64_t MAX_IN_VK_NULL_HANDLE,
    IN VK_NULL_HANDLE, OUT &nextImageIndex );

    VkCommandBufferBeginInfo vcbi;
    vcbi.sType = VK_STRUCTURE_TYPE_COMMAND_BUFFER_BEGIN_INFO;
    vcbi.pNext = nullptr;
    vcbi.flags = VK_COMMAND_BUFFER_USAGE_ONE_TIME_SUBMIT_BIT;
    vcbi.pInheritanceInfo = (VkCommandBufferInheritanceInfo*) nullptr;

    result = vkBeginCommandBuffer( CommandBuffers[nextImageIndex], IN &vcbi );
}
  
```

These are the Commands that could be entered into the Command Buffer, IV 257

```

VkClearColorValue vccv;
vccv.float32[0] = 0.0;
vccv.float32[1] = 0.0;
vccv.float32[2] = 0.0;
vccv.float32[3] = 1.0;

VkClearDepthStencilValue vcdsv;
vcdsv.depth = 1.f;
vcdsv.stencil = 0;

VkClearColorValue vcv[2];
vcv[0].color = vccv;
vcv[1].depthStencil = vcdsv;

VkOffset2D o2d = { 0, 0 };
VkExtent2D e2d = { Width, Height };
VkRect2D r2d = { o2d, e2d };

VkRenderPassBeginInfo vrbpi;
vrbpi.sType = VK_STRUCTURE_TYPE_RENDER_PASS_BEGIN_INFO;
vrbpi.pNext = nullptr;
vrbpi.renderPass = RenderPass;
vrbpi.framebuffer = Framebuffers[nextImageIndex];
vrbpi.renderArea = r2d;
vrbpi.clearValueCount = 2;
vrbpi.pClearValues = vcv; // used for VK_ATTACHMENT_LOAD_OP_CLEAR

vkCmdBeginRenderPass( CommandBuffers[nextImageIndex], IN vrbpi, IN VK_SUBPASS_CONTENTS_INLINE );
  
```

These are the Commands that could be entered into the Command Buffer, V 258

```

VkViewport viewport;
viewport.x = 0; // x
viewport.y = 0; // y
viewport.width = (float)Width;
viewport.height = (float)Height;
viewport.minDepth = 0; // minDepth
viewport.maxDepth = 1; // maxDepth

vkCmdSetViewport( CommandBuffers[nextImageIndex], 0, 1, IN &viewport ); // firstViewport, 1 viewportCount

VkRect2D scissor;
scissor.x = 0;
scissor.y = 0;
scissor.width = Width;
scissor.height = Height;

vkCmdSetScissor( CommandBuffers[nextImageIndex], 0, 1, IN &scissor );

vkCmdBindDescriptorSets( CommandBuffers[nextImageIndex], VK_PIPELINE_BIND_POINT_GRAPHICS,
GraphicsPipelineLayout, 0, 4, DescriptorSets, 0, (uint32_t*) nullptr ); // dynamic offset count, dynamic offsets

vkCmdBindPushConstants( CommandBuffers[nextImageIndex], PipelineLayout, VK_SHADER_STAGE_ALL, offset, size, void values );
VkBuffer buffers[1] = { MyVertexBuffer };
VkDeviceSize offsets[1] = { 0 };

vkCmdBindVertexBuffers( CommandBuffers[nextImageIndex], 0, 1, buffers, offsets ); // 0, 1 = firstBinding, bindingCount


const uint32_t vertexCount = sizeof(VertexData) / sizeof(VertexData[0]);
const uint32_t instanceCount = 1;
const uint32_t firstVertex = 0;
const uint32_t firstInstance = 0;
vkCmdDraw( CommandBuffers[nextImageIndex], vertexCount, instanceCount, firstVertex, firstInstance );

vkCmdEndRenderPass( CommandBuffers[nextImageIndex] );
vkEndCommandBuffer( CommandBuffers[nextImageIndex] );
  
```

Submitting a Command Buffer to a Queue for Execution 259

```

VkSubmitInfo vsi;
vsi.sType = VK_STRUCTURE_TYPE_SUBMIT_INFO;
vsi.pNext = nullptr;
vsi.commandBufferCount = 1;
vsi.pCommandBuffers = &CommandBuffer;
vsi.waitSemaphoreCount = 1;
vsi.pWaitSemaphores = &imageReadySemaphore;
vsi.signalSemaphoreCount = 0;
vsi.pSignalSemaphores = (VkSemaphore *)nullptr;
vsi.pWaitDstStageMask = (VkPipelineStageFlags *)nullptr;
  
```



mb - July 24, 2020

The Entire Submission / Wait / Display Process 260

```

VkFenceCreateInfo vci;
vci.sType = VK_STRUCTURE_TYPE_FENCE_CREATE_INFO;
vci.pNext = nullptr;
vci.flags = 0;

VkFence renderFence;
vkCreateFence(LogicalDevice, IN &vci, PALLOCATOR, OUT &renderFence;
result = VK_SUCCESS;

VkPipelineStageFlags waitABottom = VK_PIPELINE_STAGE_BOTTOM_OF_PIPE_BIT;
VkQueue presentQueue;
vkGetDeviceQueue(LogicalDevice, FindQueueFamilyThatDoesGraphics(), 0, OUT &presentQueue);


VkSubmitInfo vsi;
vsi.sType = VK_STRUCTURE_TYPE_SUBMIT_INFO;
vsi.pNext = nullptr;
vsi.waitSemaphoreCount = 1;
vsi.pWaitSemaphores = &imageReadySemaphore;
vsi.pWaitDstStageMask = &waitABottom;
vsi.commandBufferCount = 1;
vsi.pCommandBuffers = &CommandBuffer(nextImageIndex);
vsi.signalSemaphoreCount = 0;
vsi.pSignalSemaphores = &SemaphoreRenderFinished;

result = vkQueueSubmit(presentQueue, 1, IN &vsi, IN renderFence); // 1 = submitCount
result = vkWaitForFences(LogicalDevice, 1, IN &renderFence, VK_TRUE, UINT64_MAX); // waitAll_timeout

vkDestroyFence(LogicalDevice, renderFence, PALLOCATOR);

VkPresentInfoKHR vpi;
vpi.sType = VK_STRUCTURE_TYPE_PRESENT_INFO_KHR;
vpi.pNext = nullptr;
vpi.waitSemaphoreCount = 0;
vpi.pWaitSemaphores = (VkSemaphore *)nullptr;
vpi.swapchainCount = 1;
vpi.pSwapchains = &SwapChain;
vpi.pImageIndices = &nextImageIndex;
vpi.pResults = (VkResult *)nullptr;

result = vkQueuePresentKHR(presentQueue, IN &vpi);
  
```



mb - July 24, 2020

What Happens After a Queue has Been Submitted? 261


As the Vulkan 1.1 Specification says:

"Command buffer submissions to a single queue respect submission order and other implicit ordering guarantees, but otherwise may overlap or execute out of order. Other types of batches and queue submissions against a single queue (e.g. sparse memory binding) have no implicit ordering constraints with any other queue submission or batch. Additional explicit ordering constraints between queue submissions and individual batches can be expressed with semaphores and fences."

In other words, the Vulkan driver on your system will execute the commands in a single buffer in the order in which they were put there.

But, between different command buffers submitted to different queues, the driver is allowed to execute commands between buffers in-order or out-of-order or overlapped-order, depending on what it thinks it can get away with.

The message here is, I think, always consider using some sort of Vulkan synchronization when one command depends on a previous command reaching a certain state first.



mb - July 24, 2020


Vulkan.

The Swap Chain

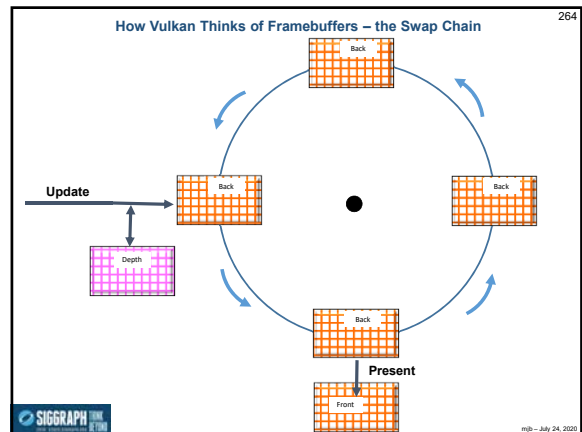
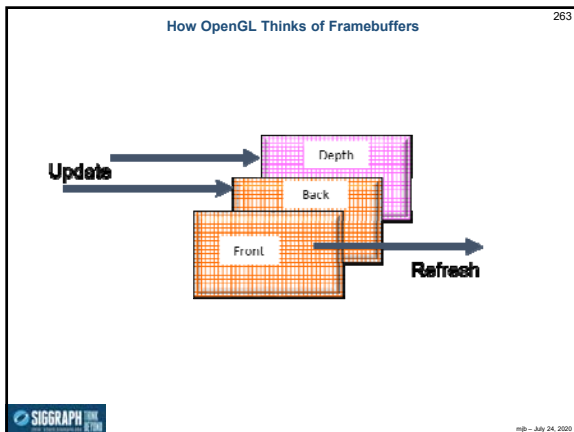
Mike Bailey
mjb@cs.oregonstate.edu

 This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License

<http://cs.oregonstate.edu/~mjb/vulkan>



mb - July 24, 2020



Push Constants

On the shader side, if, for example, you are sending a 4x4 matrix, the use of push constants in the shader looks like this:

```
layout( push_constant ) uniform matrix
{
    mat4 modelMatrix;
} Matrix;
```

On the application side, push constants are pushed at the shaders by binding them to the Vulkan Command Buffer:

```
vkCmdPushConstants( CommandBuffer, PipelineLayout, stageFlags,
    offset, size, pValues );
```

where:
stageFlags are or'ed bits of VK_PIPELINE_STAGE_VERTEX_SHADER_BIT, VK_PIPELINE_STAGE_FRAGMENT_SHADER_BIT, etc.
size is in bytes
pValues is a void * pointer to the data, which, in this 4x4 matrix example, would be of type `glm::mat4`.

Setting up the Push Constants for the Pipeline Structure

Prior to that, however, the pipeline layout needs to be told about the Push Constants:


```
VkPushConstantRange
vpcr[0].stageFlags =
    VK_PIPELINE_STAGE_VERTEX_SHADER_BIT
    | VK_PIPELINE_STAGE_FRAGMENT_SHADER_BIT;
vpcr[0].offset = 0;
vpcr[0].size = sizeof( glm::mat4 );
```

```
VkPipelineLayoutCreateInfo
vplci.sType = VK_STRUCTURE_TYPE_PIPELINE_LAYOUT_CREATE_INFO;
vplci.pNext = nullptr;
vplci.flags = 0;
vplci.setLayoutCount = 4;
vplci.pSetLayouts = DescriptorSetLayouts;
vplci.pushConstantRangeCount = 1;
vplci.pPushConstantRanges = vpcr;
```

```
result = vkCreatePipelineLayout( LogicalDevice, IN &vplci, PALLOCATOR,
    OUT &GraphicsPipelineLayout );
```

An Robotic Example using Push Constants

A robotic animation (i.e., a hierarchical transformation system)



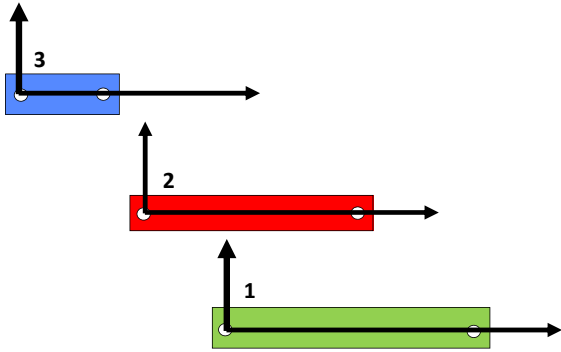
Where each arm is represented by:

```
struct arm
{
    glm::mat4  armMatrix;
    glm::vec3  armColor;
    float      armScale; // scale factor in x
};
```

```
struct armArm1;
struct armArm2;
struct armArm3;
```

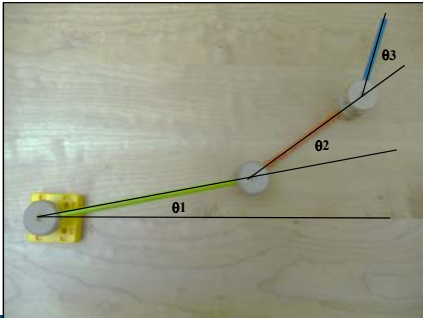
Forward Kinematics:

You Start with Separate Pieces, all Defined in their Own Local Coordinate System



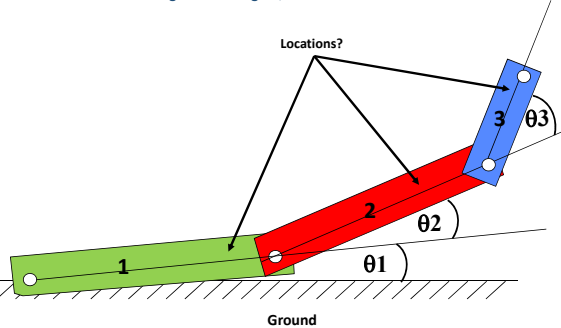
Forward Kinematics:

Hook the Pieces Together, Change Parameters, and Things Move (All Young Children Understand This)



Forward Kinematics:

Given the Lengths and Angles, Where do the Pieces Move To?




Positioning Part #1 With Respect to Ground 283

1. Rotate by Θ_1
2. Translate by $T_{1/G}$

Write it →

$$[M_{1/G}] = [T_{1/G}] * [R_{\theta_1}]$$

← Say it



mb - July 24, 2020

Why Do We Say it Right-to-Left? 284

Write it →

$$[M_{x,y,z}] = [T_{x,y,z}] * [R_{\theta_1}]$$


← Say it

We adopt the convention that the coordinates are multiplied on the right side of the matrix:

$$\begin{bmatrix} x' \\ y' \\ z' \\ 1 \end{bmatrix} = \begin{bmatrix} A & B & C & D \\ E & F & G & H \\ I & J & K & L \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix}$$

$$\begin{bmatrix} x' \\ y' \\ z' \\ 1 \end{bmatrix} = [M_{1/G}] \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} = [T_{1/G}] * [R_{\theta_1}] \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix}$$

So the right-most transformation in the sequence multiplies the (x,y,z,1) first and the left-most transformation multiplies it last



mb - July 24, 2020

Positioning Part #2 With Respect to Ground 285


1. Rotate by Θ_2
2. Translate the length of part 2
3. Rotate by Θ_1
4. Translate by $T_{1/G}$

Write it →

$$[M_{2/G}] = [T_{1/G}] * [R_{\theta_1}] * [T_{2/1}] * [R_{\theta_2}]$$

$$[M_{2/G}] = [M_{1/G}] * [M_{2/1}]$$

← Say it



mb - July 24, 2020

Positioning Part #3 With Respect to Ground 286


1. Rotate by Θ_3
2. Translate the length of part 2
3. Rotate by Θ_2
4. Translate the length of part 1
5. Rotate by Θ_1
6. Translate by $T_{1/G}$

Write it →

$$[M_{3/G}] = [T_{1/G}] * [R_{\theta_1}] * [T_{2/1}] * [R_{\theta_2}] * [T_{3/2}] * [R_{\theta_3}]$$

$$[M_{3/G}] = [M_{1/G}] * [M_{2/1}] * [M_{3/2}]$$

← Say it



mb - July 24, 2020

In the Reset Function 287


```

struct arm      Arm1;
struct arm      Arm2;
struct arm      Arm3;
...
Arm1.armMatrix = glm::mat4( 1. );
Arm1.armColor  = glm::vec3( 0.f, 1.f, 0.f );
Arm1.armScale  = 6.f;

Arm2.armMatrix = glm::mat4( 1. );
Arm2.armColor  = glm::vec3( 1.f, 0.f, 0.f );
Arm2.armScale  = 4.f;

Arm3.armMatrix = glm::mat4( 1. );
Arm3.armColor  = glm::vec3( 0.f, 0.f, 1.f );
Arm3.armScale  = 2.f;
    
```

The constructor **glm::mat4(1.)** produces an identity matrix. The actual transformation matrices will be set in *UpdateScene()*.



mb - July 24, 2020

Setup the Push Constant for the Pipeline Structure 288


```

VkPushConstantRange
vpcr[0].stageFlags =
    VK_PIPELINE_STAGE_VERTEX_SHADER_BIT
    | VK_PIPELINE_STAGE_FRAGMENT_SHADER_BIT;
vpcr[0].offset = 0;
vpcr[0].size = sizeof( struct arm );

VkPipelineLayoutCreateInfo
vpcli.sType = VK_STRUCTURE_TYPE_PIPELINE_LAYOUT_CREATE_INFO;
vpcli.pNext = nullptr;
vpcli.flags = 0;
vpcli.setLayoutCount = 4;
vpcli.pSetLayouts = DescriptorSetLayouts;
vpcli.pushConstantRangeCount = 1;
vpcli.pPushConstantRanges = vpcr;

result = vkCreatePipelineLayout( LogicalDevice, IN &vpcli, PALLOCATOR,
    OUT &GraphicsPipelineLayout );
    
```

Annotations: **vpcr[1]** and **vpcli** are circled in red in the original image.



mb - July 24, 2020

In the UpdateScene Function 289

```

float rot1 = (float)Time;
float rot2 = 2.f * rot1;
float rot3 = 2.f * rot2;

glm::vec3 zaxis = glm::vec3(0., 0., 1.);

glm::mat4 m1g = glm::mat4( 1. ); // identity
m1g = glm::translate(m1g, glm::vec3(0., 0., 0.));
m1g = glm::rotate(m1g, rot1, zaxis); // [T][R]

glm::mat4 m21 = glm::mat4( 1. ); // identity
m21 = glm::translate(m21, glm::vec3(2.*Arm1.armScale, 0., 0.));
m21 = glm::rotate(m21, rot2, zaxis); // [T][R]
m21 = glm::translate(m21, glm::vec3(0., 0., 2.)); // z-offset from previous arm

glm::mat4 m32 = glm::mat4( 1. ); // identity
m32 = glm::translate(m32, glm::vec3(2.*Arm2.armScale, 0., 0.));
m32 = glm::rotate(m32, rot3, zaxis); // [T][R]
m32 = glm::translate(m32, glm::vec3(0., 0., 2.)); // z-offset from previous arm

Arm1.armMatrix = m1g; // m1g
Arm2.armMatrix = m1g * m21; // m2g
Arm3.armMatrix = m1g * m21 * m32; // m3g
    
```

siggraph 2018

In the RenderScene Function 290

```

VkBuffer buffers[1] = { MyVertexDataBuffer.buffer };

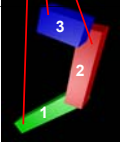
vkCmdBindVertexBuffers( CommandBuffers[nextImageIndex], 0, 1, buffers, offsets );

vkCmdPushConstants( CommandBuffers[nextImageIndex], GraphicsPipelineLayout,
    VK_SHADER_STAGE_ALL, 0, sizeof(struct arm), (void *)&Arm1 );
vkCmdDraw( CommandBuffers[nextImageIndex], vertexCount, instanceCount, firstVertex, firstInstance );

vkCmdPushConstants( CommandBuffers[nextImageIndex], GraphicsPipelineLayout,
    VK_SHADER_STAGE_ALL, 0, sizeof(struct arm), (void *)&Arm2 );
vkCmdDraw( CommandBuffers[nextImageIndex], vertexCount, instanceCount, firstVertex, firstInstance );

vkCmdPushConstants( CommandBuffers[nextImageIndex], GraphicsPipelineLayout,
    VK_SHADER_STAGE_ALL, 0, sizeof(struct arm), (void *)&Arm3 );
vkCmdDraw( CommandBuffers[nextImageIndex], vertexCount, instanceCount, firstVertex, firstInstance );
    
```

The strategy is to draw each link using the same vertex buffer, but modified with a unique color, length, and matrix transformation



siggraph 2018

In the Vertex Shader 291

```

layout( push_constant ) uniform arm
{
    mat4 armMatrix;
    vec3 armColor;
    float armScale; // scale factor in x
} RobotArm;

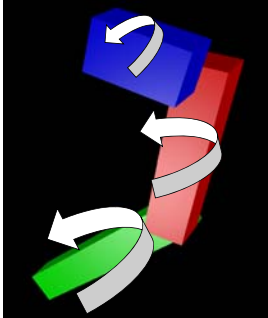
layout( location = 0 ) in vec3 aVertex;
...

vec3 bVertex = aVertex; // arm coordinate system is [-1., 1.] in X
bVertex.x += 1.; // now is [0., 2.]
bVertex.x /= 2.; // now is [0., 1.]
bVertex.x = (RobotArm.armScale); // now is [0., 1.]
bVertex = vec3( RobotArm.armMatrix * vec4( bVertex, 1. ) );
...


gl_Position = PVM * vec4( bVertex, 1. ); // Projection * Viewing * Modeling matrices
    
```

siggraph 2018


292



siggraph 2018

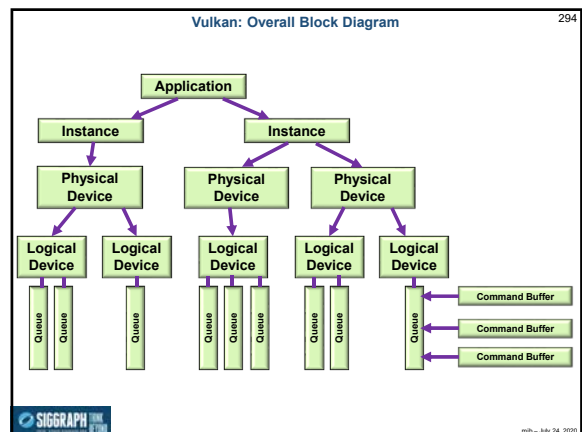

Physical Devices

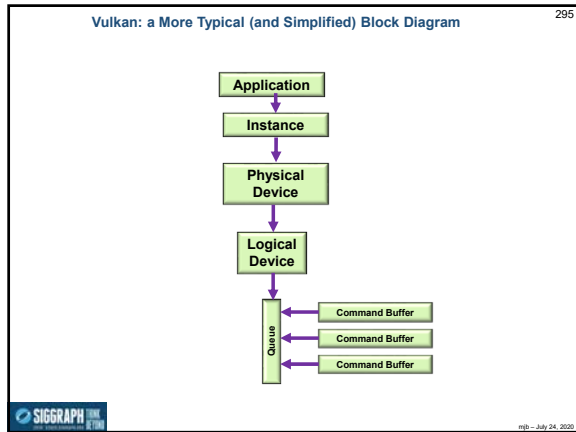
 Mike Bailey
 mjb@cs.oregonstate.edu


 This work is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](https://creativecommons.org/licenses/by-nc-nd/4.0/).

<http://cs.oregonstate.edu/~mjb/vulkan>

siggraph 2018





Querying the Number of Physical Devices

```
uint32_t count;
result = vkEnumeratePhysicalDevices( Instance, OUT &count, OUT (VkPhysicalDevice *)nullptr );

VkPhysicalDevice * physicalDevices = new VkPhysicalDevice[ count ];
result = vkEnumeratePhysicalDevices( Instance, OUT &count, OUT physicalDevices );
```

This way of querying information is a recurring OpenCL and Vulkan pattern (get used to it):

```
result = vkEnumeratePhysicalDevices( Instance, &count, nullptr );
result = vkEnumeratePhysicalDevices( Instance, &count, physicalDevices );
```

How many total there are: **&count**
Where to put them: **nullptr** / **physicalDevices**

296

Vulkan: Identifying the Physical Devices

```
VkResult result = VK_SUCCESS;
result = vkEnumeratePhysicalDevices( Instance, OUT &PhysicalDeviceCount, (VkPhysicalDevice *)nullptr );
if( result != VK_SUCCESS || PhysicalDeviceCount <= 0 )
{
    fprintf( FpDebug, "Could not count the physical devices!\n" );
    return VK_SHOULD_EXIT;
}
fprintf( FpDebug, "%d physical devices found.\n", PhysicalDeviceCount );

VkPhysicalDevice * physicalDevices = new VkPhysicalDevice[ PhysicalDeviceCount ];
result = vkEnumeratePhysicalDevices( Instance, OUT &PhysicalDeviceCount, OUT physicalDevices );
if( result != VK_SUCCESS )
{
    fprintf( FpDebug, "Could not enumerate the %d physical devices!\n", PhysicalDeviceCount );
    return VK_SHOULD_EXIT;
}
```

297

Which Physical Device to Use, I

```
int discreteSelect = -1;
int integratedSelect = -1;
for( unsigned int i = 0; i < PhysicalDeviceCount; i++ )
{
    VkPhysicalDeviceProperties vpdp;
    vkGetPhysicalDeviceProperties( IN physicalDevices[ i ], OUT &vpdp );
    if( result != VK_SUCCESS )
    {
        fprintf( FpDebug, "Could not get the physical device properties of device %d!\n", i );
        return VK_SHOULD_EXIT;
    }

    fprintf( FpDebug, "\nDevice %2d!\n", i );
    fprintf( FpDebug, "API version: %d!\n", vpdp.apiVersion );
    fprintf( FpDebug, "Driver version: %d!\n", vpdp.driverVersion );
    fprintf( FpDebug, "Vendor ID: 0x%04x!\n", vpdp.vendorID );
    fprintf( FpDebug, "Device ID: 0x%04x!\n", vpdp.deviceID );
    fprintf( FpDebug, "Physical Device Type: %d = ", vpdp.deviceType );
    if( vpdp.deviceType == VK_PHYSICAL_DEVICE_TYPE_DISCRETE_GPU ) fprintf( FpDebug, "(Discrete GPU)\n" );
    if( vpdp.deviceType == VK_PHYSICAL_DEVICE_TYPE_INTEGRATED_GPU ) fprintf( FpDebug, "(Integrated GPU)\n" );
    if( vpdp.deviceType == VK_PHYSICAL_DEVICE_TYPE_VIRTUAL_GPU ) fprintf( FpDebug, "(Virtual GPU)\n" );
    if( vpdp.deviceType == VK_PHYSICAL_DEVICE_TYPE_CPU ) fprintf( FpDebug, "(CPU)\n" );
    fprintf( FpDebug, "Device Name: %s!\n", vpdp.deviceName );
    fprintf( FpDebug, "Pipeline Cache Size: %d!\n", vpdp.pipelineCacheUID[0] );
}
```

298

Which Physical Device to Use, II

```
// need some logical here to decide which physical device to select:
if( vpdp.deviceType == VK_PHYSICAL_DEVICE_TYPE_DISCRETE_GPU )
    discreteSelect = i;
if( vpdp.deviceType == VK_PHYSICAL_DEVICE_TYPE_INTEGRATED_GPU )
    integratedSelect = i;
}

int which = -1;
if( discreteSelect >= 0 )
{
    which = discreteSelect;
    PhysicalDevice = physicalDevices[which];
}
else if( integratedSelect >= 0 )
{
    which = integratedSelect;
    PhysicalDevice = physicalDevices[which];
}
else
{
    fprintf( FpDebug, "Could not select a Physical Device!\n" );
    return VK_SHOULD_EXIT;
}
```

299

Asking About the Physical Device's Features

```
VkPhysicalDeviceProperties PhysicalDeviceFeatures;
vkGetPhysicalDeviceFeatures( IN PhysicalDevice, OUT &PhysicalDeviceFeatures );

fprintf( FpDebug, "\nPhysical Device Features!\n" );
fprintf( FpDebug, "geometryShader = %2d!\n", PhysicalDeviceFeatures.geometryShader );
fprintf( FpDebug, "tessellationShader = %2d!\n", PhysicalDeviceFeatures.tessellationShader );
fprintf( FpDebug, "multiDrawIndirect = %2d!\n", PhysicalDeviceFeatures.multiDrawIndirect );
fprintf( FpDebug, "wideLines = %2d!\n", PhysicalDeviceFeatures.wideLines );
fprintf( FpDebug, "largePoints = %2d!\n", PhysicalDeviceFeatures.largePoints );
fprintf( FpDebug, "multiViewport = %2d!\n", PhysicalDeviceFeatures.multiViewport );
fprintf( FpDebug, "occlusionQueryPrecise = %2d!\n", PhysicalDeviceFeatures.occlusionQueryPrecise );
fprintf( FpDebug, "pipelineStatisticsQuery = %2d!\n", PhysicalDeviceFeatures.pipelineStatisticsQuery );
fprintf( FpDebug, "shaderFloat64 = %2d!\n", PhysicalDeviceFeatures.shaderFloat64 );
fprintf( FpDebug, "shaderInt64 = %2d!\n", PhysicalDeviceFeatures.shaderInt64 );
fprintf( FpDebug, "shaderInt16 = %2d!\n", PhysicalDeviceFeatures.shaderInt16 );
```

300

Here's What the NVIDIA RTX 2080 Ti Produced


301

```

vkEnumeratePhysicalDevices:

Device 0:
  API version: 4198499
  Driver version: 4198499
  Vendor ID: 0x10de
  Device ID: 0x1e04
  Physical Device Type: 2 = (Discrete GPU)
  Device Name: RTX 2080 Ti
  Pipeline Cache Size: 206
Device #0 selected ('RTX 2080 Ti')

Physical Device Features:
  geometryShader = 1
  tessellationShader = 1
  multiDrawIndirect = 1
  wideLines = 1
  largePoints = 1
  multiViewport = 1
  occlusionQueryPrecise = 1
  pipelineStatisticsQuery = 1
  shaderFloat64 = 1
  shaderInt64 = 1
  shaderInt16 = 1
    
```



mb - July 24, 2020

Here's What the Intel HD Graphics 520 Produced


302

```

vkEnumeratePhysicalDevices:

Device 0:
  API version: 4194360
  Driver version: 4194360
  Vendor ID: 0x8086
  Device ID: 0x1916
  Physical Device Type: 1 = (Integrated GPU)
  Device Name: Intel(R) HD Graphics 520
  Pipeline Cache Size: 213
Device #0 selected ('Intel(R) HD Graphics 520')

Physical Device Features:
  geometryShader = 1
  tessellationShader = 1
  multiDrawIndirect = 1
  wideLines = 1
  largePoints = 1
  multiViewport = 1
  occlusionQueryPrecise = 1
  pipelineStatisticsQuery = 1
  shaderFloat64 = 1
  shaderInt64 = 1
  shaderInt16 = 1
    
```



mb - July 24, 2020

Asking About the Physical Device's Different Memories


303

```

VkPhysicalDeviceMemoryProperties vpdmp;
vkGetPhysicalDeviceMemoryProperties( PhysicalDevice, OUT &vpdmp );

fprintf( FpDebug, "In%d Memory Types:\n", vpdmp.memoryTypeCount );
for( unsigned int i = 0; i < vpdmp.memoryTypeCount; i++ )
{
  VkMemoryType vmt = vpdmp.memoryTypes[i];
  fprintf( FpDebug, "Memory %2d: ", i );
  if( ( vmt.propertyFlags & VK_MEMORY_PROPERTY_DEVICE_LOCAL_BIT ) != 0 ) fprintf( FpDebug, " DeviceLocal" );
  if( ( vmt.propertyFlags & VK_MEMORY_PROPERTY_HOST_VISIBLE_BIT ) != 0 ) fprintf( FpDebug, " HostVisible" );
  if( ( vmt.propertyFlags & VK_MEMORY_PROPERTY_HOST_COHERENT_BIT ) != 0 ) fprintf( FpDebug, " HostCoherent" );
  if( ( vmt.propertyFlags & VK_MEMORY_PROPERTY_HOST_CACHED_BIT ) != 0 ) fprintf( FpDebug, " HostCached" );
  if( ( vmt.propertyFlags & VK_MEMORY_PROPERTY_LAZILY_ALLOCATED_BIT ) != 0 ) fprintf( FpDebug, " LazilyAllocated" );
  fprintf( FpDebug, "\n" );
}

fprintf( FpDebug, "In%d Memory Heaps:\n", vpdmp.memoryHeapCount );
for( unsigned int i = 0; i < vpdmp.memoryHeapCount; i++ )
{
  fprintf( FpDebug, "Heap %d: ", i );
  VkMemoryHeap vmh = vpdmp.memoryHeaps[i];
  fprintf( FpDebug, " size = %x%08x", (unsigned long) vmh.size );
  if( ( vmh.flags & VK_MEMORY_HEAP_DEVICE_LOCAL_BIT ) != 0 ) fprintf( FpDebug, " DeviceLocal" ); // only one in use
  fprintf( FpDebug, "\n" );
}
    
```



mb - July 24, 2020


Here's What I Got

304

```

11 Memory Types:
Memory 0:
Memory 1:
Memory 2:
Memory 3:
Memory 4:
Memory 5:
Memory 6:
Memory 7: DeviceLocal
Memory 8: DeviceLocal
Memory 9: HostVisible HostCoherent
Memory 10: HostVisible HostCoherent HostCached

2 Memory Heaps:
Heap 0: size = 0xb7c00000 DeviceLocal
Heap 1: size = 0xfac00000
    
```



mb - July 24, 2020


Asking About the Physical Device's Queue Families

305

```

uint32_t count = -1;
vkGetPhysicalDeviceQueueFamilyProperties( IN PhysicalDevice, &count, OUT (VkQueueFamilyProperties *) nullptr );
fprintf( FpDebug, "InFound %d Queue Families:\n", count );

VkQueueFamilyProperties *vqfp = new VkQueueFamilyProperties[ count ];
vkGetPhysicalDeviceQueueFamilyProperties( IN PhysicalDevice, &count, OUT vqfp );
for( unsigned int i = 0; i < count; i++ )
{
  fprintf( FpDebug, "i%d: queueCount = %2d : ", i, vqfp[i].queueCount );
  if( ( vqfp[i].queueFlags & VK_QUEUE_GRAPHICS_BIT ) != 0 ) fprintf( FpDebug, " Graphics" );
  if( ( vqfp[i].queueFlags & VK_QUEUE_COMPUTE_BIT ) != 0 ) fprintf( FpDebug, " Compute" );
  if( ( vqfp[i].queueFlags & VK_QUEUE_TRANSFER_BIT ) != 0 ) fprintf( FpDebug, " Transfer" );
  fprintf( FpDebug, "\n" );
}
    
```




mb - July 24, 2020

Here's What I Got

306


```

Found 3 Queue Families:
0: queueCount = 16 ; Graphics Compute Transfer
1: queueCount = 2 ; Transfer
2: queueCount = 8 ; Compute
    
```




mb - July 24, 2020

307




Logical Devices

Mike Bailey
mjb@cs.oregonstate.edu

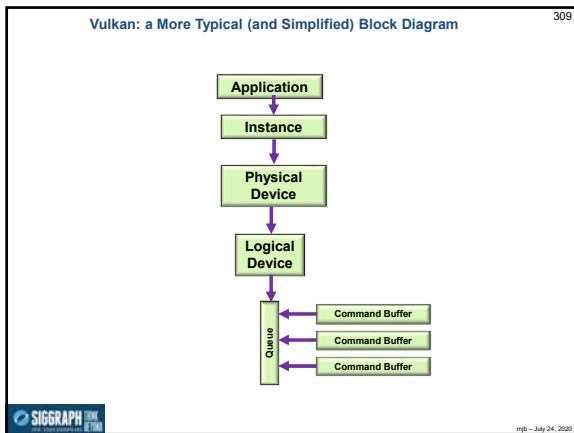
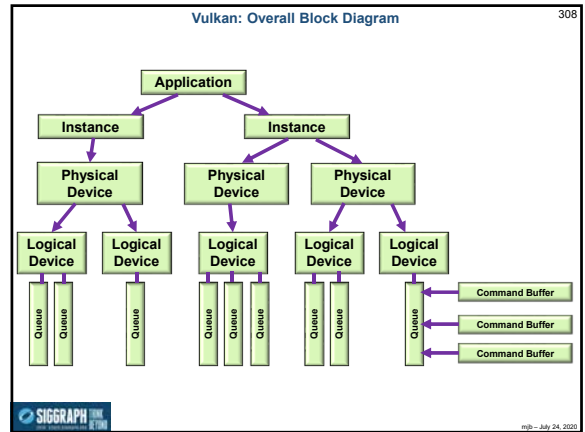


This work is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](https://creativecommons.org/licenses/by-nc-nd/4.0/).

<http://cs.oregonstate.edu/~mjb/vulkan>



mjb - July 24, 2020




Looking to See What Device Layers are Available 310

```

const char * myDeviceLayers[] =
{
    // "VK_LAYER_LUNARG_api_dump",
    // "VK_LAYER_LUNARG_core_validation",
    // "VK_LAYER_LUNARG_image",
    "VK_LAYER_LUNARG_object_tracker",
    "VK_LAYER_LUNARG_parameter_validation",
    // "VK_LAYER_NV_optimus"
};

const char * myDeviceExtensions[] =
{
    "VK_KHR_surface",
    "VK_KHR_win32_surface",
    "VK_EXT_debug_report",
    // "VK_KHR_swapchain"
};

// see what device layers are available:
uint32_t layerCount;
vkEnumerateDeviceLayerProperties(PhysicalDevice, &layerCount, (VkLayerProperties *)nullptr);
VkLayerProperties * deviceLayers = new VkLayerProperties[layerCount];
result = vkEnumerateDeviceLayerProperties(PhysicalDevice, &layerCount, deviceLayers);
    
```




mjb - July 24, 2020

Looking to See What Device Extensions are Available 311

```

// see what device extensions are available:
uint32_t extensionCount;
vkEnumerateDeviceExtensionProperties(PhysicalDevice, deviceLayers[layerName],
    &extensionCount, (VkExtensionProperties *)nullptr);
VkExtensionProperties * deviceExtensions = new VkExtensionProperties[extensionCount];
result = vkEnumerateDeviceExtensionProperties(PhysicalDevice, deviceLayers[layerName],
    &extensionCount, deviceExtensions);
    
```



mjb - July 24, 2020

What Device Layers and Extensions are Available 312


```

4 physical device layers enumerated:
0x00401063 1 "VK_LAYER_NV_optimus" "NVIDIA Optimus layer"
0 device extensions enumerated for "VK_LAYER_NV_optimus":

0x00401072 1 "VK_LAYER_LUNARG_core_validation" "LunarG Validation Layer"
2 device extensions enumerated for "VK_LAYER_LUNARG_core_validation":
0x00000001 "VK_EXT_validation_cache"
0x00000004 "VK_EXT_debug_marker"

0x00401072 1 "VK_LAYER_LUNARG_object_tracker" "LunarG Validation Layer"
2 device extensions enumerated for "VK_LAYER_LUNARG_object_tracker":
0x00000001 "VK_EXT_validation_cache"
0x00000004 "VK_EXT_debug_marker"

0x00401072 1 "VK_LAYER_LUNARG_parameter_validation" "LunarG Validation Layer"
2 device extensions enumerated for "VK_LAYER_LUNARG_parameter_validation":
0x00000001 "VK_EXT_validation_cache"
0x00000004 "VK_EXT_debug_marker"
    
```



mjb - July 24, 2020

Vulkan: Creating a Logical Device

```

float queuePriorities[1]=
{
    1.
};
VkDeviceQueueCreateInfo vdqci;
vdqci.sType = VK_STRUCTURE_TYPE_DEVICE_QUEUE_CREATE_INFO;
vdqci.pNext = nullptr;
vdqci.flags = 0;
vdqci.queueFamilyIndex = 0;
vdqci.queueCount = 1;
vdqci.pQueueProperties = queuePriorities;

VkDeviceCreateInfo vdc;
vdc.sType = VK_STRUCTURE_TYPE_DEVICE_CREATE_INFO;
vdc.pNext = nullptr;
vdc.flags = 0;
vdc.queueCreateInfoCount = 1; // # of device queues
vdc.pQueueCreateInfos = IN vdqci; // array of VkDeviceQueueCreateInfos
vdc.enabledLayerCount = sizeof(myDeviceLayers) / sizeof(char *);
vdc.enabledLayerNames = myDeviceLayers;
vdc.ppEnabledExtensionNames = nullptr; // no extensions
vdc.enabledExtensionCount = sizeof(myDeviceExtensions) / sizeof(char *);
vdc.ppEnabledExtensionNames = myDeviceExtensions;
vdc.pEnabledFeatures = IN &PhysicalDeviceFeatures;

result = vkCreateLogicalDevice( PhysicalDevice, IN &vdc, PALLOCATOR, OUT &LogicalDevice );
    
```

Vulkan: Creating the Logical Device's Queue

```

// get the queue for this logical device:
vkGetDeviceQueue( LogicalDevice, 0, 0, OUT &Queue ); // 0, 0 = queueFamilyIndex, queueIndex
    
```

Vulkan. Dynamic State Variables

Mike Bailey
mjb@cs.oregonstate.edu

<http://cs.oregonstate.edu/~mjb/vulkan>

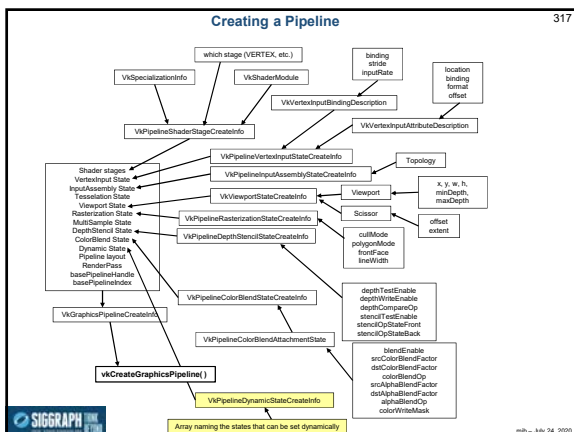
This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License

Creating a Pipeline with Dynamically Changeable State Variables

The graphics pipeline data structure is full of state information, and, as previously-discussed, is largely immutable, that is, the information contained inside it is fixed, and can only be changed by creating a new graphics pipeline data structure with new information.

That isn't quite true. To a certain extent, Vulkan allows you to declare parts of the pipeline state changeable. This allows you to alter pipeline state information on the fly.

This is useful for managing state information that needs to change frequently. This also creates possible optimization opportunities for the Vulkan driver.



Which Pipeline State Variables can be Changed Dynamically

The possible dynamic variables are shown in the `VkDynamicState` enum:

- VK_DYNAMIC_STATE_VIEWPORT
- VK_DYNAMIC_STATE_SCISSOR
- VK_DYNAMIC_STATE_LINE_WIDTH
- VK_DYNAMIC_STATE_DEPTH_BIAS
- VK_DYNAMIC_STATE_BLEND_CONSTANTS
- VK_DYNAMIC_STATE_DEPTH_BOUNDS
- VK_DYNAMIC_STATE_STENCIL_COMPARE_MASK
- VK_DYNAMIC_STATE_STENCIL_WRITE_MASK
- VK_DYNAMIC_STATE_STENCIL_REFERENCE

Creating a Pipeline

319

```

VkDynamicState
{
    VK_DYNAMIC_STATE_VIEWPORT,
    VK_DYNAMIC_STATE_LINE_WIDTH,
};

VkPipelineDynamicStateCreateInfo
vpdsci.sType = VK_STRUCTURE_TYPE_PIPELINE_DYNAMIC_STATE_CREATE_INFO;
vpdsci.pNext = nullptr;
vpdsci.flags = 0;
vpdsci.dynamicStateCount = sizeof(vds) / sizeof(VkDynamicState); // i.e., 2
vpdsci.pDynamicStates = &vds;

VkGraphicsPipelineCreateInfo
...
vgpci.pDynamicState = &vpdsci;
...

vkCreateGraphicsPipelines(LogicalDevice, pipelineCache, 1, &vgpci, PALLOCATOR, &GraphicsPipeline);
    
```

If you declare certain state variables to be dynamic like this, then you **must** fill them in the command buffer! Otherwise, they are **undefined**.

Filling the Dynamic State Variables in the Command Buffer

320


First call:

```
vkCmdBindPipeline( ... );
```

Then, the command buffer-bound function calls to set these dynamic states are:


```

vkCmdSetViewport( commandBuffer, firstViewport, viewportCount, pViewports );
vkCmdSetScissor( commandBuffer, firstScissor, scissorCount, pScissors );
vkCmdSetLineWidth( commandBuffer, lineWidth );
vkCmdSetDepthBias( commandBuffer, depthBiasConstantFactor, depthBiasClamp, depthBiasSlopeFactor );
vkCmdSetBlendConstants( commandBuffer, blendConstants[4] );
vkCmdSetDepthBounds( commandBuffer, minDepthBounds, maxDepthBounds );
vkCmdSetStencilCompareMask( commandBuffer, faceMask, compareMask );
vkCmdSetStencilWriteMask( commandBuffer, faceMask, writeMask );
vkCmdSetStencilReference( commandBuffer, faceMask, reference );
    
```



Getting Information Back from the Graphics System

Mike Bailey
mjb@cs.oregonstate.edu



This work is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](https://creativecommons.org/licenses/by-nc-nd/4.0/).

http://cs.oregonstate.edu/~mjb/vulkan

Setting up Query Pools

322

- There are 3 types of Queries: Occlusion, Pipeline Statistics, and Timestamp
- Vulkan requires you to first setup "Query Pools", one for each specific type
- This indicates that Vulkan thinks that Queries are time-consuming (relatively) to setup, and thus better to set them up in program-setup than in program-runtime

Setting up Query Pools

323

```

VkQueryPoolCreateInfo
vgpci.sType = VK_STRUCTURE_TYPE_QUERY_POOL_CREATE_INFO;
vgpci.pNext = nullptr;
vgpci.flags = 0;
vgpci.queryType = << one of: >>
    VK_QUERY_TYPE_OCCLUSION
    VK_QUERY_TYPE_PIPELINE_STATISTICS
    VK_QUERY_TYPE_TIMESTAMP
vgpci.queryCount = 1;
vgpci.pipelineStatistics = 0; // bitmask of what stats you are querying for if you are doing a pipeline statistics query

VK_QUERY_PIPELINE_STATISTIC_INPUT_ASSEMBLY_PRIMITIVES_BIT
VK_QUERY_PIPELINE_STATISTIC_INPUT_ASSEMBLY_PATCHES_BIT
VK_QUERY_PIPELINE_STATISTIC_VERTEX_SHADER_INVOCATIONS_BIT
VK_QUERY_PIPELINE_STATISTIC_GEOMETRY_SHADER_INVOCATIONS_BIT
VK_QUERY_PIPELINE_STATISTIC_GEOMETRY_SHADER_PRIMITIVES_BIT
VK_QUERY_PIPELINE_STATISTIC_CLIPPING_INVOCATIONS_BIT
VK_QUERY_PIPELINE_STATISTIC_FRAGMENT_SHADER_INVOCATIONS_BIT
VK_QUERY_PIPELINE_STATISTIC_TESSellation_CONTROL_SHADER_PATCHES_BIT
VK_QUERY_PIPELINE_STATISTIC_TESSellation_EVALUATION_SHADER_INVOCATIONS_BIT
VK_QUERY_PIPELINE_STATISTIC_COMPUTE_SHADER_INVOCATIONS_BIT

VkQueryPool occlusionQueryPool;
result = vkCreateQueryPool( LogicalDevice, IN &vgpci, PALLOCATOR, OUT &occlusionQueryPool );

VkQueryPool statisticsQueryPool;
result = vkCreateQueryPool( LogicalDevice, IN &vgpci, PALLOCATOR, OUT &statisticsQueryPool );

VkQueryPool timestampQueryPool;
result = vkCreateQueryPool( LogicalDevice, IN &vgpci, PALLOCATOR, OUT &timestampQueryPool );
    
```

Resetting, Filling, and Examining a Query Pool

324

```

vkCmdResetQueryPool( CommandBuffer, occlusionQueryPool, 0, 1 );

vkCmdBeginQuery( CommandBuffer, occlusionQueryPool, 0, VK_QUERY_CONTROL_PRECISE_BIT );
...
vkCmdEndQuery( CommandBuffer, occlusionQueryPool, 0 );

#define DATASIZE 128
uint32_t data[DATASIZE];

result = vkGetQueryPoolResults( LogicalDevice, occlusionQueryPool, 0, 1, DATASIZE* sizeof( uint32_t ), data, stride, flags );
// or'ed combinations of:
// VK_QUERY_RESULT_64_BIT
// VK_QUERY_RESULT_WAIT_BIT
// VK_QUERY_RESULT_WITH_AVAILABILITY_BIT
// VK_QUERY_RESULT_PARTIAL_BIT
// stride is # of bytes in between each result
    
```

Occlusion Query

325


Occlusion Queries count the number of fragments drawn between the **vkCmdBeginQuery** and the **vkCmdEndQuery** that pass both the Depth and Stencil tests

This is commonly used to see what level-of-detail should be used when drawing a complicated object

Some hints:

- Don't draw the whole scene – just draw the object(s) you are interested in
- Don't draw the whole object – just draw a simple bounding volume at least as big as the object(s)
- Don't draw the whole bounding volume – cull away the back faces (two reasons: time and correctness)
- Don't draw the colors – just draw the depths (especially if the fragment shader is time-consuming)

```
uint32_t fragmentCount;
result = vkGetQueryPoolResults( LogicalDevice, occlusionQueryPool, 0, 1,
                               sizeof(uint32_t), &fragmentCount, 0, VK_QUERY_RESULT_WAIT_BIT );
```



mb - July 24, 2020

Pipeline Statistics Query


326

Pipeline Statistics Queries count how many of various things get done between the **vkCmdBeginQuery** and the **vkCmdEndQuery**

```
uint32_t counts[NUM_STATS];
result = vkGetQueryPoolResults( LogicalDevice, statisticsQueryPool, 0, 1,
                               NUM_STATS*sizeof(uint32_t), counts, 0, VK_QUERY_RESULT_WAIT_BIT );
```

// vqpci.pipelineStatistics = or'ed bits of:

- // VK_QUERY_PIPELINE_STATISTIC_INPUT_ASSEMBLY_VERTICES_BIT
- // VK_QUERY_PIPELINE_STATISTIC_INPUT_ASSEMBLY_PRIMITIVES_BIT
- // VK_QUERY_PIPELINE_STATISTIC_VERTEX_SHADER_INVOCATIONS_BIT
- // VK_QUERY_PIPELINE_STATISTIC_GEOMETRY_SHADER_INVOCATIONS_BIT
- // VK_QUERY_PIPELINE_STATISTIC_GEOMETRY_SHADER_PRIMITIVES_BIT
- // VK_QUERY_PIPELINE_STATISTIC_CLIPPING_INVOCATIONS_BIT
- // VK_QUERY_PIPELINE_STATISTIC_CLIPPING_PRIMITIVES_BIT
- // VK_QUERY_PIPELINE_STATISTIC_FRAGMENT_SHADER_INVOCATIONS_BIT
- // VK_QUERY_PIPELINE_STATISTIC_TESSELLATION_CONTROL_SHADER_PATCHES_BIT
- // VK_QUERY_PIPELINE_STATISTIC_TESSELLATION_EVALUATION_SHADER_INVOCATIONS_BIT
- // VK_QUERY_PIPELINE_STATISTIC_COMPUTE_SHADER_INVOCATIONS_BIT




mb - July 24, 2020

Timestamp Query

327

Timestamp Queries count how many nanoseconds of time elapsed between the **vkCmdBeginQuery** and the **vkCmdEndQuery**.

```
uint64_t nanosecondsCount;
result = vkGetQueryPoolResults( LogicalDevice, timestampQueryPool, 0, 1,
                               sizeof(uint64_t), &nanosecondsCount, 0,
                               VK_QUERY_RESULT_64_BIT | VK_QUERY_RESULT_WAIT_BIT);
```



mb - July 24, 2020

Timestamp Query


328

The **vkCmdWriteTimeStamp()** function produces the time between when this function is called and when the first thing reaches the specified pipeline stage.

Even though the stages are "bits", you are supposed to only specify one of them, not "or" multiple ones together

```
vkCmdWriteTimeStamp( CommandBuffer, pipelineStages, timestampQueryPool, 0 );
```


- // VK_PIPELINE_STAGE_TOP_OF_PIPE_BIT
- // VK_PIPELINE_STAGE_DRAW_INDIRECT_BIT
- // VK_PIPELINE_STAGE_VERTEX_INPUT_BIT
- // VK_PIPELINE_STAGE_VERTEX_SHADER_BIT
- // VK_PIPELINE_STAGE_TESSELLATION_CONTROL_SHADER_BIT
- // VK_PIPELINE_STAGE_TESSELLATION_EVALUATION_SHADER_BIT
- // VK_PIPELINE_STAGE_GEOMETRY_SHADER_BIT
- // VK_PIPELINE_STAGE_FRAGMENT_SHADER_BIT
- // VK_PIPELINE_STAGE_FRAGMENT_SHADER_BIT | VK_PIPELINE_STAGE_EARLY_FRAGMENT_TESTS_BIT
- // VK_PIPELINE_STAGE_LATE_FRAGMENT_TESTS_BIT | VK_PIPELINE_STAGE_COLOR_ATTACHMENT_OUTPUT_BIT
- // VK_PIPELINE_STAGE_COMPUTE_SHADER_BIT
- // VK_PIPELINE_STAGE_TRANSFER_BIT
- // VK_PIPELINE_STAGE_BOTTOM_OF_PIPE_BIT
- // VK_PIPELINE_STAGE_HOST_BIT




mb - July 24, 2020

Vulkan. Compute Shaders

Mike Bailey
mjb@cs.oregonstate.edu



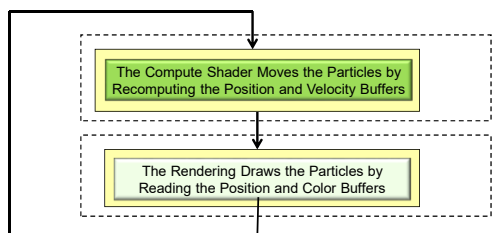
<http://cs.oregonstate.edu/~mjb/vulkan>



mb - July 24, 2020


The Example We Are Going to Use Here is a Particle System

330



```

graph TD
    A[The Compute Shader Moves the Particles by  
Recomputing the Position and Velocity Buffers] --> B[The Rendering Draws the Particles by  
Reading the Position and Color Buffers]
    
```



mb - July 24, 2020

The Data in your C/C++ Program will look like This

This is a Particle System application, so we need Positions, Velocities, and (possibly) Colors

```

#define NUM_PARTICLES (1024*1024) // total number of particles to move
#define NUM_WORK_ITEMS_PER_GROUP 64 // # work-items per work-group
#define NUM_X_WORK_GROUPS (NUM_PARTICLES / NUM_WORK_ITEMS_PER_GROUP)

struct pos
{
    glm::vec4; // positions
};

struct vel
{
    glm::vec4; // velocities
};

struct col
{
    glm::vec4; // colors
};
    
```

Note that .w and .vw are not actually needed. But, by making these structure sizes a multiple of 4 floats, it doesn't matter if they are declared with the std140 or the std430 qualifier. I think this is a good thing.

The Data in your Compute Shader will look like This

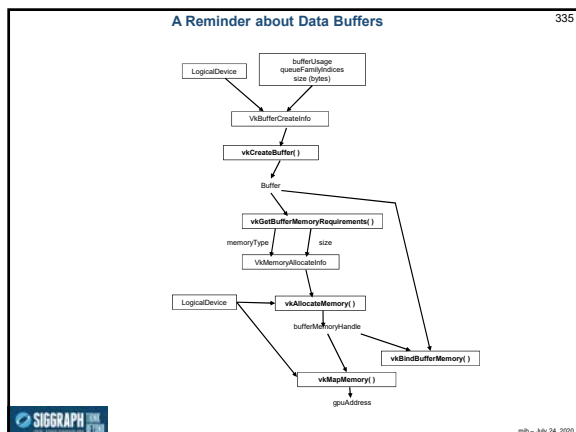
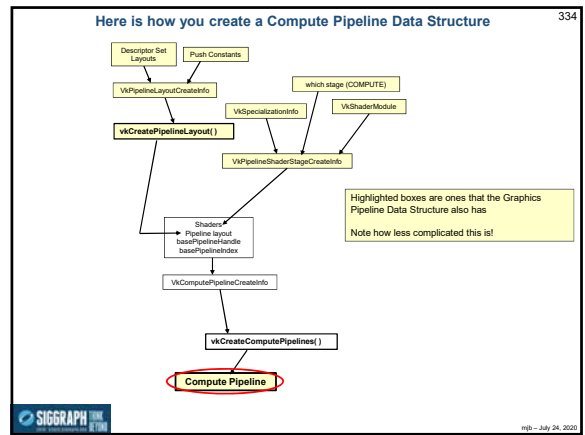
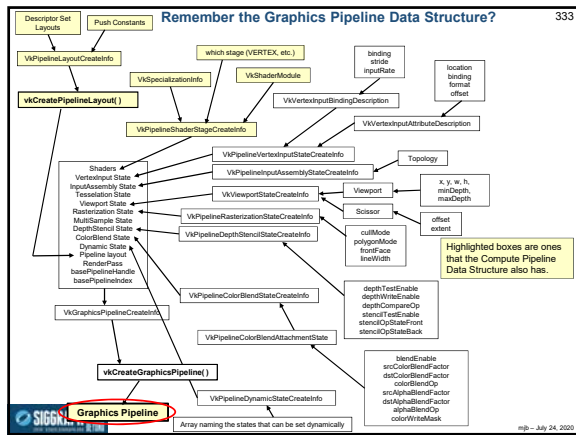
```

layout( std140, set = 0, binding = 0 ) buffer Pos
{
    vec4 Positions[ ]; // array of structures
};

layout( std140, set = 0, binding = 1 ) buffer Vel
{
    vec4 Velocities[ ]; // array of structures
};

layout( std140, set = 0, binding = 2 ) buffer Col
{
    vec4 Colors[ ]; // array of structures
};
    
```

You can use the empty brackets, but only on the last element of the buffer. The actual dimension will be determined for you when Vulkan examines the size of this buffer's data store.



Creating a Shader Storage Buffer

```

VkBuffer PosBuffer;
...
VkBufferCreateInfo vbci;
vbci.sType = VK_STRUCTURE_TYPE_BUFFER_CREATE_INFO;
vbci.pNext = nullptr;
vbci.flags = 0;
vbci.size = NUM_PARTICLES * sizeof( glm::vec4 );
vbci.usage = VK_USAGE_SHADER_STORAGE_BUFFER_BIT;
vbci.sharingMode = VK_SHARING_MODE_EXCLUSIVE;
vbci.queueFamilyIndexCount = 0;
vbci.pQueueFamilyIndices = (const uint32_t*) nullptr;

result = vkCreateBuffer( LogicalDevice, IN &vbci, ALLOCATOR, OUT &PosBuffer );
    
```

Allocating Memory for a Buffer, Binding a Buffer to Memory, and Filling the Buffer

337

```

VkMemoryRequirements
result = vkGetBufferMemoryRequirements( LogicalDevice, PosBuffer, OUT &vmr );

VkMemoryAllocateInfo
vmai.sType = VK_STRUCTURE_TYPE_MEMORY_ALLOCATE_INFO;
vmai.pNext = nullptr;
vmai.flags = 0;
vmai.allocationSize = vmr.size;
vmai.memoryTypeIndex = FindMemoryThatIsHostVisible();
...

VkDeviceMemory
result = vkAllocateMemory( LogicalDevice, IN &vmai, PALLOCATOR, OUT &vdm );
result = vkBindBufferMemory( LogicalDevice, PosBuffer, IN vdm, 0 ); // 0 is the offset
    
```

SIGGRAPH 2020 mb - July 24, 2020

Create the Compute Pipeline Layout

338

```

VkDescriptorSetLayoutBinding
ComputeSet[0].binding = 0;
ComputeSet[0].descriptorType = VK_DESCRIPTOR_TYPE_STORAGE_BUFFER;
ComputeSet[0].descriptorCount = 1;
ComputeSet[0].stageFlags = VK_SHADER_STAGE_COMPUTE_BIT;
ComputeSet[0].pImmutableSamplers = (VkSampler*)nullptr;

ComputeSet[1].binding = 1;
ComputeSet[1].descriptorType = VK_DESCRIPTOR_TYPE_STORAGE_BUFFER;
ComputeSet[1].descriptorCount = 1;
ComputeSet[1].stageFlags = VK_SHADER_STAGE_COMPUTE_BIT;
ComputeSet[1].pImmutableSamplers = (VkSampler*)nullptr;

ComputeSet[2].binding = 2;
ComputeSet[2].descriptorType = VK_DESCRIPTOR_TYPE_STORAGE_BUFFER;
ComputeSet[2].descriptorCount = 1;
ComputeSet[2].stageFlags = VK_SHADER_STAGE_COMPUTE_BIT;
ComputeSet[2].pImmutableSamplers = (VkSampler*)nullptr;

VkDescriptorSetLayoutCreateInfo
vdsi.sType = VK_STRUCTURE_TYPE_DESCRIPTOR_SET_LAYOUT_CREATE_INFO;
vdsi.pNext = nullptr;
vdsi.flags = 0;
vdsi.bindingCount = 3;
vdsi.pBindings = &ComputeSet[0];
    
```

SIGGRAPH 2020 mb - July 24, 2020

Create the Compute Pipeline Layout

339

```

VkPipelineLayout
VkDescriptorSetLayout
...

result = vkCreateDescriptorSetLayout( LogicalDevice, IN &vdsi, PALLOCATOR, OUT &ComputeSetLayout );

VkPipelineLayoutCreateInfo
vpcli.sType = VK_STRUCTURE_TYPE_PIPELINE_LAYOUT_CREATE_INFO;
vpcli.pNext = nullptr;
vpcli.flags = 0;
vpcli.setLayoutCount = 1;
vpcli.pSetLayouts = &ComputeSetLayout;
vpcli.pushConstantRangeCount = 0;
vpcli.pPushConstantRanges = (VkPushConstantRange*)nullptr;

result = vkCreatePipelineLayout( LogicalDevice, IN &vpcli, PALLOCATOR, OUT &ComputePipelineLayout );
    
```

SIGGRAPH 2020 mb - July 24, 2020

Create the Compute Pipeline

340

```

VkPipeline
...

VkPipelineShaderStageCreateInfo
vpssci.sType = VK_STRUCTURE_TYPE_PIPELINE_SHADER_STAGE_CREATE_INFO;
vpssci.pNext = nullptr;
vpssci.flags = 0;
vpssci.stage = VK_SHADER_STAGE_COMPUTE_BIT;
vpssci.module = computeShader;
vpssci.pName = "main";
vpssci.pSpecializationInfo = (VkSpecializationInfo*)nullptr;

VkComputePipelineCreateInfo
vpcpi.sType = VK_STRUCTURE_TYPE_COMPUTE_PIPELINE_CREATE_INFO;
vpcpi.pNext = nullptr;
vpcpi.flags = 0;
vpcpi.stage = vpssci;
vpcpi.layout = ComputePipelineLayout;
vpcpi.basePipelineHandle = VK_NULL_HANDLE;
vpcpi.basePipelineIndex = 0;

result = vkCreateComputePipelines( LogicalDevice, VK_NULL_HANDLE, 1, &vpcpi[0], PALLOCATOR, &ComputePipeline );
    
```

SIGGRAPH 2020 mb - July 24, 2020

Creating a Vulkan Data Buffer

341

```

VkBuffer Buffer;

VkBufferCreateInfo vbci;
vbci.sType = VK_STRUCTURE_TYPE_BUFFER_CREATE_INFO;
vbci.pNext = nullptr;
vbci.flags = 0;
vbci.size = NUM_PARTICLES * sizeof( glm::vec4 );
vbci.usage = VK_USAGE_STORAGE_BUFFER_BIT;
vbci.sharingMode = VK_SHARING_MODE_CONCURRENT;
vbci.queueFamilyIndexCount = 0;
vbci.pQueueFamilyIndices = (const int32_t*)nullptr;

result = vkCreateBuffer ( LogicalDevice, IN &vbci, PALLOCATOR, OUT &posBuffer );
    
```

SIGGRAPH 2020 mb - July 24, 2020

Allocating Memory and Binding the Buffer

342

```

VkMemoryRequirements
result = vkGetBufferMemoryRequirements( LogicalDevice, posBuffer, OUT &vmr );

VkMemoryAllocateInfo
vmai.sType = VK_STRUCTURE_TYPE_MEMORY_ALLOCATE_INFO;
vmai.pNext = nullptr;
vmai.flags = 0;
vmai.allocationSize = vmr.size;
vmai.memoryTypeIndex = FindMemoryThatIsHostVisible();

VkDeviceMemory
result = vkAllocateMemory( LogicalDevice, IN &vmai, PALLOCATOR, OUT &vdm );
result = vkBindBufferMemory( LogicalDevice, posBuffer, IN vdm, 0 ); // 0 is the offset

MyBuffer myPosBuffer;
myPosBuffer.size = vbci.size;
myPosBuffer.buffer = posBuffer;
myPosBuffer.vdm = vdm;
    
```

SIGGRAPH 2020 mb - July 24, 2020

Fill the Buffers

343

```

struct pos * positions;
vkMapMemory(LogicalDevice, IN myPosBuffer.vdm, 0, VK_WHOLE_SIZE, 0, OUT (void *) &positions );
for( int i = 0; i < NUM_PARTICLES; i++)
{
    positions[i].x = Ranf( XMIN, XMAX );
    positions[i].y = Ranf( YMIN, YMAX );
    positions[i].z = Ranf( ZMIN, ZMAX );
    positions[i].w = 1.;
}
vkUnmapMemory(LogicalDevice, IN myPosBuffer.vdm );

struct vel * velocities;
vkMapMemory(LogicalDevice, IN myVelBuffer.vdm, 0, VK_WHOLE_SIZE, 0, OUT (void *) &velocities );
for( int i = 0; i < NUM_PARTICLES; i++)
{
    velocities[i].x = Ranf( VXMIN, VXMAX );
    velocities[i].y = Ranf( VYMIN, VYMAX );
    velocities[i].z = Ranf( VZMIN, VZMAX );
    velocities[i].w = 0.;
}
vkUnmapMemory(LogicalDevice, IN myVelBuffer.vdm );

struct col * colors;
vkMapMemory(LogicalDevice, IN myColBuffer.vdm, 0, VK_WHOLE_SIZE, 0, OUT (void *) &colors );
for( int i = 0; i < NUM_PARTICLES; i++)
{
    colors[i].r = Ranf( .3f, 1. );
    colors[i].g = Ranf( .3f, 1. );
    colors[i].b = Ranf( .3f, 1. );
    colors[i].a = 1.;
}
vkUnmapMemory(LogicalDevice, IN myColBuffer.vdm );
    
```

mb - July 24, 2020

Fill the Buffers

344

```

#include <stdlib.h>

#define TOP 2147483647. // 2^31 - 1

float
Ranf( float low, float high )
{
    long random( ); // returns integer 0 - TOP

    float r = (float)rand( );
    return low + r * ( high - low ) / (float)RAND_MAX ;
}
    
```

SIGGRAPH 2012
mb - July 24, 2020

The Particle System Compute Shader

345

```

layout( std140, set = 0, binding = 0 ) buffer Pos
{
    vec4 Positions[ ]; // array of structures
};



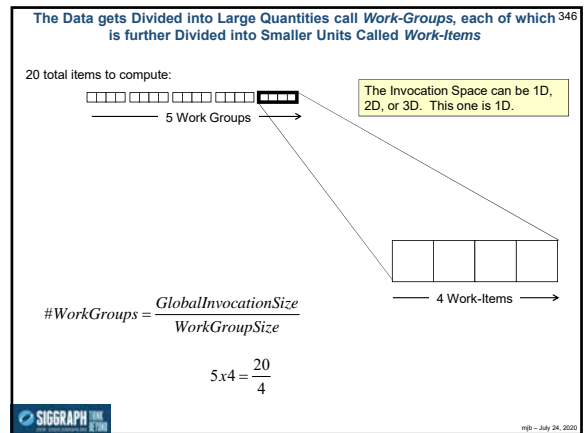
layout( std140, set = 0, binding = 1 ) buffer Vel
{
    vec4 Velocities[ ]; // array of structures
};

layout( std140, set = 0, binding = 2 ) buffer Col
{
    vec4 Colors[ ]; // array of structures
};

layout( local_size_x = 64, local_size_y = 1, local_size_z = 1 ) in;
    
```

This is the number of **work-items per work-group**, set in the compute shader. The number of work-groups is set in the `vkCmdDispatch(commandBuffer, workGroupCountX, workGroupCountY, workGroupCountZ);` function call in the application program.

SIGGRAPH 2012
mb - July 24, 2020



The Particle System Compute Shader – The Physics

349

```

#define POINT          vec3
#define VELOCITY      vec3
#define VECTOR        vec3
#define SPHERE        vec4
#define PLANE         vec4
// xc, yc, zc, r
// a, b, c, d

const VECTOR G       = VECTOR(0., -9.8, 0.);
const float DT       = 0.1;

const SPHERE Sphere = vec4(-100., -800., 0., 600.); // x, y, z, r
...
uint gid = gl_GlobalInvocationID.x; // where I am in the global dataset (6 in this example)
// (as a 1d problem, the .y and .z are both 1)


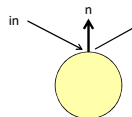
POINT p = Positions[ gid ].xyz;
VELOCITY v = Velocities[ gid ].xyz;

POINT pp = p + v*DT + .5*DT*DT*G;
VELOCITY vp = v + G*DT;

Positions[ gid ].xyz = pp;
Velocities[ gid ].xyz = vp;
    
```

$$p' = p + v \cdot t + \frac{1}{2} G \cdot t^2$$

$$v' = v + G \cdot t$$

SIGGRAPH 2018 July 24, 2020

The Particle System Compute Shader – How About Introducing a Bounce?

350

```

VELOCITY
Bounce( VELOCITY vin, VECTOR n )
{
    VELOCITY vout = reflect( vin, n );
    return vout;
}

// plane equation: Ax + By + Cz + D = 0
// ( it turns out that (A,B,C) is the normal )

VELOCITY
BouncePlane( POINT p, VELOCITY v, PLANE pl )
{
    VECTOR n = normalize( VECTOR( pl.xyz ) );
    return Bounce( v, n );
}

bool
IsUnderPlane( POINT p, PLANE pl )
{
    float r = pl.x*p.x + pl.y*p.y + pl.z*p.z + pl.w;
    return ( r < 0. );
}
    
```

Note: a surface in the x-z plane has the equation:
 $0x + 1y + 0z + 0 = 0$
 and thus its normal vector is (0,1,0)

SIGGRAPH 2018 July 24, 2020

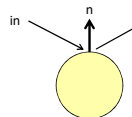
The Particle System Compute Shader – How About Introducing a Bounce?

351

```

VELOCITY
BounceSphere( POINT p, VELOCITY v, SPHERE s )
{
    VECTOR n = normalize( p - s.xyz );
    return Bounce( v, n );
}

bool
IsInsideSphere( POINT p, SPHERE s )
{
    float r = length( p - s.xyz );
    return ( r < s.w );
}
    
```



SIGGRAPH 2018 July 24, 2020

The Particle System Compute Shader – How About Introducing a Bounce?

352

```

uint gid = gl_GlobalInvocationID.x; // the .y and .z are both 1 in this case

POINT p = Positions[ gid ].xyz;
VELOCITY v = Velocities[ gid ].xyz;

POINT pp = p + v*DT + .5*DT*DT*G;
VELOCITY vp = v + G*DT;

if( IsInsideSphere( pp, Sphere ) )
{
    vp = BounceSphere( p, v, S );
    pp = p + vp*DT + .5*DT*DT*G;
}

Positions[ gid ].xyz = pp;
Velocities[ gid ].xyz = vp;
    
```

$$p' = p + v \cdot t + \frac{1}{2} G \cdot t^2$$

$$v' = v + G \cdot t$$

Graphics Trick Alert: Making the bounce happen from the surface of the sphere is time-consuming. Instead, bounce from the previous position in space. If DT is small enough (and it is), nobody will ever know...

SIGGRAPH 2018 July 24, 2020

Dispatching the Compute Shader from the Command Buffer

353

```

#define NUM_PARTICLES (1024*1024)
#define NUM_WORK_ITEMS_PER_GROUP 64
#define NUM_X_WORK_GROUPS ( NUM_PARTICLES / NUM_WORK_ITEMS_PER_GROUP )
...
vkCmdBindPipeline( CommandBuffer, VK_PIPELINE_BIND_POINT_COMPUTE, ComputePipeline );
vkCmdDispatch( CommandBuffer, NUM_X_WORK_GROUPS, 1, 1 );
    
```

This is the number of work-groups, set in the application program.
 The number of work-items per work-group is set in the layout in the compute shader:

$$\text{layout}(\text{local_size_x} = 64, \text{local_size_y} = 1, \text{local_size_z} = 1) \text{ in};$$

SIGGRAPH 2018 July 24, 2020

Displaying the Particles

354

```

VkVertexInputBindingDescription vvbld[3]; // one of these per buffer data buffer
vwbld[0].binding = 0; // which binding # this is
vwbld[0].stride = sizeof( struct pos ); // bytes between successive structs
vwbld[0].inputRate = VK_VERTEX_INPUT_RATE_VERTEX;

vwbld[1].binding = 1;
vwbld[1].stride = sizeof( struct vel );
vwbld[1].inputRate = VK_VERTEX_INPUT_RATE_VERTEX;

vwbld[2].binding = 2;
vwbld[2].stride = sizeof( struct col );
vwbld[2].inputRate = VK_VERTEX_INPUT_RATE_VERTEX;

layout( location = 0 ) in vec4 aPosition;
layout( location = 1 ) in vec4 aVelocity;
layout( location = 2 ) in vec4 aColor;
    
```

SIGGRAPH 2018 July 24, 2020

Displaying the Particles


355

```

VkVertexInputAttributeDescription  viad[3]; // array per vertex input attribute
// 3 = position, velocity, color
viad[0].location = 0; // location in the layout decoration
viad[0].binding = 0; // which binding description this is part of
viad[0].format = VK_FORMAT_VEC4; // x, y, z, w
viad[0].offset = offsetof( struct pos, pos ); // 0

viad[1].location = 1;
viad[1].binding = 0;
viad[1].format = VK_FORMAT_VEC4; // nx, ny, nz
viad[1].offset = offsetof( struct vel, vel ); // 0

viad[2].location = 2;
viad[2].binding = 0;
viad[2].format = VK_FORMAT_VEC4; // r, g, b, a
viad[2].offset = offsetof( struct col, col ); // 0
    
```

 mb - July 24, 2020


Telling the Pipeline about its Input

356

```

VkPipelineVertexInputStateCreateInfo  vpvisci; // used to describe the input vertex attributes
vpvisci.sType = VK_STRUCTURE_TYPE_PIPELINE_VERTEX_INPUT_STATE_CREATE_INFO;
vpvisci.pNext = nullptr;
vpvisci.flags = 0;
vpvisci.vertexBindingDescriptionCount = 3;
vpvisci.pVertexBindingDescriptions = vviad;
vpvisci.vertexAttributeDescriptionCount = 3;
vpvisci.pVertexAttributeDescriptions = vviad;

VkPipelineInputAssemblyStateCreateInfo  vpiasci;
vpiasci.sType = VK_STRUCTURE_TYPE_PIPELINE_INPUT_ASSEMBLY_STATE_CREATE_INFO;
vpiasci.pNext = nullptr;
vpiasci.flags = 0;
vpiasci.topology = VK_PRIMITIVE_TOPOLOGY_POINT_LIST;
    
```

 mb - July 24, 2020

Telling the Pipeline about its Input


357

We will come to the Pipeline later, but for now, know that a Vulkan Pipeline is essentially a very large data structure that holds (what OpenGL would call) the state, including how to parse its vertex input.

```

VkGraphicsPipelineCreateInfo
vgpci.sType = VK_STRUCTURE_TYPE_GRAPHICS_PIPELINE_CREATE_INFO;
vgpci.pNext = nullptr;
vgpci.flags = 0;
vgpci.stageCount = 2; // number of shader stages in this pipeline
vgpci.pStages = vpsci;
vgpci.pVertexInputState = &vpvisci;
vgpci.pInputAssemblyState = &vpiasci;
vgpci.pTessellationState = (VkPipelineTessellationStateCreateInfo *) nullptr; // &vptsci
vgpci.pViewportState = &vpvsci;
vgpci.pRasterizationState = &vprsci;
vgpci.pMultisampleState = &vpmsci;
vgpci.pDepthStencilState = &vpdsci;
vgpci.pColorBlendState = &vpbcsci;
vgpci.pDynamicState = &vpdsci;
vgpci.layout = IN GraphicsPipelineLayout;
vgpci.renderPass = IN RenderPass;
vgpci.subpass = 0; // subpass number
vgpci.basePipelineHandle = (VkPipeline) VK_NULL_HANDLE;
vgpci.basePipelineIndex = 0;

result = vkCreateGraphicsPipelines( LogicalDevice, VK_NULL_HANDLE, 1, IN &vgpci,
PALLOCATOR, OUT &GraphicsPipeline );
    
```

 mb - July 24, 2020


Setting a Pipeline Barrier so the Drawing Waits for the Compute

358

```

VkBufferMemoryBarrier
vbmb.sType = VK_STRUCTURE_TYPE_BUFFER_MEMORY_BARRIER;
vbmb.pNext = nullptr;
vbmb.srcAccessFlags = VK_ACCESS_SHADER_WRITE_BIT;
vbmb.dstAccessFlags = VK_ACCESS_VERTEX_ATTRIBUTE_READ_BIT;
vbmb.srcQueueFamilyIndex = 0;
vbmb.dstQueueFamilyIndex = 0;
vbmb.buffer =
vbmb.offset = 0;
vbmb.size = NUM_PARTICLES * sizeof( glm::vec4 );

const uint32_t bufferMemoryBarrierCount = 1;
VkCmdPipelineBarrier
(
    commandBuffer,
    VK_PIPELINE_STAGE_COMPUTE_SHADER_BIT, VK_PIPELINE_STAGE_VERTEX_INPUT_BIT,
    VK_DEPENDENCY_BY_REGION_BIT, 0, nullptr, bufferMemoryBarrierCount,
    IN &vbmb, 0, nullptr
);
    
```

 mb - July 24, 2020

Drawing

359


```

VkBuffer buffers[ ] = MyPosBuffer.buffer, MyVelBuffer.buffer, MyColBuffer.buffer ;
size_t offsets[ ] = { 0, 0, 0 };

vkCmdBindVertexBuffers( CommandBuffers[nextImageIndex], 0, 3, buffers, offsets );

const uint32_t vertexCount = NUM_PARTICLES;
const uint32_t instanceCount = 1;
const uint32_t firstVertex = 0;
const uint32_t firstInstance = 0;

vkCmdDraw( CommandBuffers[nextImageIndex], NUM_PARTICLES, 1, 0, 0 );
// vertexCount, instanceCount, firstVertex, firstInstance
    
```

 mb - July 24, 2020


Setting a Pipeline Barrier so the Compute Waits for the Drawing

360

```

VkBufferMemoryBarrier
vbmb.sType = VK_STRUCTURE_TYPE_BUFFER_MEMORY_BARRIER;
vbmb.pNext = nullptr;
vbmb.srcAccessFlags = 0;
vbmb.dstAccessFlags = VK_ACCESS_UNIFORM_READ_BIT;
vbmb.srcQueueFamilyIndex = 0;
vbmb.dstQueueFamilyIndex = 0;
vbmb.buffer =
vbmb.offset = 0;
vbmb.size = ??

const uint32_t bufferMemoryBarrierCount = 1;
VkCmdPipelineBarrier
(
    commandBuffer,
    VK_PIPELINE_STAGE_BOTTOM_OF_PIPE_BIT, VK_PIPELINE_STAGE_COMPUTE_SHADER_BIT,
    VK_DEPENDENCY_BY_REGION_BIT, 0, nullptr, bufferMemoryBarrierCount,
    IN &vbmb, 0, nullptr
);
    
```

 mb - July 24, 2020

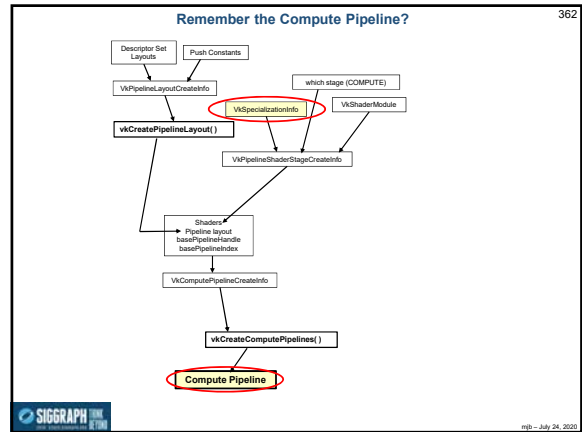
Vulkan.
Specialization Constants

Mike Bailey
mjb@cs.oregonstate.edu

<http://cs.oregonstate.edu/~mjb/vulkan>

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License

SIGGRAPH 2020 July 24, 2020 361



What Are Specialization Constants?

In Vulkan, all shaders get half-way-compiled into SPIR-V and then the rest-of-the-way compiled by the Vulkan driver.

Normally, the half-way compile finalizes all constant values and compiles the code that uses them.

But, it would be nice every so often to have your Vulkan program sneak into the half-way-compiled binary and manipulate some constants at runtime. This is what Specialization Constants are for. A Specialization Constant is a way of injecting an integer, Boolean, uint, float, or double constant into a *halfway-compiled* version of a shader right before the *rest-of-the-way* compilation.

That final compilation happens when you call `vkCreateComputePipelines()`

Without Specialization Constants, you would have to commit to a final value before the SPIR-V compile was done, which could have been a long time ago

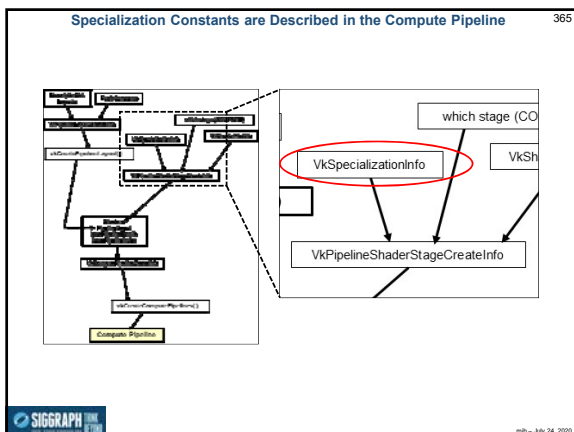
SIGGRAPH 2020 July 24, 2020 363

Why Do We Need Specialization Constants?

Specialization Constants could be used for:

- Setting the work-items per work-group in a compute shader
- Setting a Boolean flag and then eliminating the if-test that used it
- Setting an integer constant and then eliminating the switch-statement that looked for it
- Making a decision to unroll a for-loop because the number of passes through it are small enough
- Collapsing arithmetic expressions into a single value
- Collapsing trivial simplifications, such as adding zero or multiplying by 1

SIGGRAPH 2020 July 24, 2020 364



Specialization Constant Example -- Setting an Array Size

In the compute shader

```
layout( constant_id = 7 ) const int ASIZE = 32;
int array[ASIZE];
```

In the Vulkan C/C++ program:

```
int asize = 64;
VkSpecializationMapEntry vsme[1]; // one array element for each
// Specialization Constant
vsme[0].constantID = 7; // # bytes into the Specialization Constant
vsme[0].offset = 0; // array this one item is
vsme[0].size = sizeof(asize); // size of just this Specialization Constant

VkSpecializationInfo vsi;
vsi.mapEntryCount = 1;
vsi.pMapEntries = &vsme[0];
vsi.dataSize = sizeof(asize); // size of all the Specialization Constants together
vsi.pData = &asize; // array of all the Specialization Constants
```

SIGGRAPH 2020 July 24, 2020 366

Linking the Specialization Constants into the Compute Pipeline

367

```

int asize = 64;

VkSpecializationMapEntry vsme[1];
vsme[0].constantID = 7;
vsme[0].offset = 0;
vsme[0].size = sizeof(asize);

VkSpecializationInfo vsi;
vsi.mapEntryCount = 1;
vsi.pMapEntries = &vsme[0];
vsi.dataSize = sizeof(asize);
vsi.pData = &asize;

VkPipelineShaderStageCreateInfo vpssci;
vpssci.sType = VK_STRUCTURE_TYPE_PIPELINE_SHADER_STAGE_CREATE_INFO;
vpssci.pNext = nullptr;
vpssci.flags = 0;
vpssci.stage = VK_SHADER_STAGE_COMPUTE_BIT;
vpssci.module = computeShader;
vpssci.pName = "main";
vpssci.pSpecializationInfo = &vsi;

VkComputePipelineCreateInfo vcpcci[1];
vcpcci[0].sType = VK_STRUCTURE_TYPE_COMPUTE_PIPELINE_CREATE_INFO;
vcpcci[0].pNext = nullptr;
vcpcci[0].flags = 0;
vcpcci[0].stage = vpssci;
vcpcci[0].layout = ComputePipelineLayout;
vcpcci[0].basePipelineHandle = VK_NULL_HANDLE;
vcpcci[0].basePipelineIndex = 0;

result = vkCreateComputePipelines(LogicalDevice, VK_NULL_HANDLE, 1, &vcpcci[0], ALLOCATOR_OUT &ComputePipeline);
    
```

Specialization Constant Example – Setting Multiple Constants

368

```

In the compute shader
layout( constant_id = 9 ) const int a = 1;
layout( constant_id = 10 ) const int b = 2;
layout( constant_id = 11 ) const float c = 3.14;

In the C/C++ program:
struct abc { int a, int b, float c; } abc;

VkSpecializationMapEntry vsme[3];
vsme[0].constantID = 9;
vsme[0].offset = offsetof( abc, a );
vsme[0].size = sizeof(abc.a);
vsme[1].constantID = 10;
vsme[1].offset = offsetof( abc, b );
vsme[1].size = sizeof(abc.b);
vsme[2].constantID = 11;
vsme[2].offset = offsetof( abc, c );
vsme[2].size = sizeof(abc.c);

VkSpecializationInfo vsi;
vsi.mapEntryCount = 3;
vsi.pMapEntries = &vsme[0];
vsi.dataSize = sizeof(abc);
vsi.pData = &abc;
// size of all the Specialization Constants together
// array of all the Specialization Constants
    
```

It's important to use sizeof() and offsetof() instead of hardcoding numbers!

Specialization Constants – Setting the Number of Work-items Per Work-Group in the Compute Shader

369

```

In the compute shader
layout( local_size_x_id=12 ) in;
layout( local_size_x = 32, local_size_y = 1, local_size_z = 1 ) in;

In the C/C++ program:
int numXworkItems = 64;

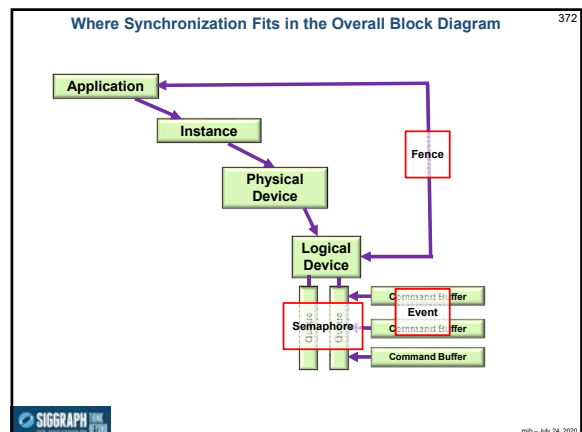
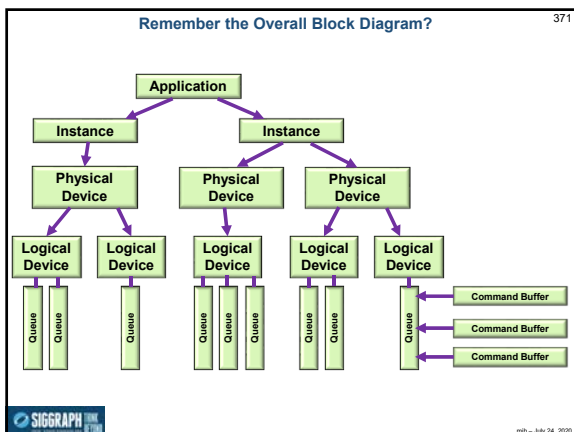
VkSpecializationMapEntry vsme[1];
vsme[0].constantID = 12;
vsme[0].offset = 0;
vsme[0].size = sizeof(int);

VkSpecializationInfo vsi;
vsi.mapEntryCount = 1;
vsi.pMapEntries = &vsme[0];
vsi.dataSize = sizeof(int);
vsi.pData = &numXworkItems;
    
```

Vulkan Synchronization

Mike Bailey
mjb@cs.oregonstate.edu

<http://cs.oregonstate.edu/~mjb/vulkan>



Semaphores

373

- Used to synchronize work executing on different queues within the same logical device
- You create them, and give them to a Vulkan function which sets them. Later on, you tell a Vulkan function to wait on this particular semaphore
- You don't end up setting, resetting, or checking the semaphore yourself
- Semaphores must be initialized ("created") before they can be used

mb - July 24, 2020

Creating a Semaphore

374

```

VkSemaphoreCreateInfo
  vscl.sType = VK_STRUCTURE_TYPE_SEMAPHORE_CREATE_INFO;
  vscl.pNext = nullptr;
  vscl.flags = 0;

VkSemaphore semaphore;
result = vkCreateSemaphore( LogicalDevice, IN &vscl, PALLOCATOR, OUT &semaphore );
    
```

This doesn't actually do anything with the semaphore – it just sets it up

mb - July 24, 2020

Semaphores Example during the Render Loop

375

```

VkSemaphore imageReadySemaphore;

VkSemaphoreCreateInfo
  vscl;
  vscl.sType = VK_STRUCTURE_TYPE_SEMAPHORE_CREATE_INFO;
  vscl.pNext = nullptr;
  vscl.flags = 0;

result = vkCreateSemaphore( LogicalDevice, IN &vscl, PALLOCATOR, OUT &imageReadySemaphore );

uint32_t nextImageIndex;
vkAcquireNextImageKHR( LogicalDevice, IN SwapChain, IN UINT64_MAX,
  IN imageReadySemaphore, IN VK_NULL_HANDLE, OUT &nextImageIndex );
  ...
  Set the semaphore

VkPipelineStageFlags waitAtBottom = VK_PIPELINE_STAGE_BOTTOM_OF_PIPE_BIT;
VkSubmitInfo
  vsi;
  vsi.sType = VK_STRUCTURE_TYPE_SUBMIT_INFO;
  vsi.pNext = nullptr;
  vsi.waitSemaphoreCount = 1;
  vsi.pWaitSemaphores = &imageReadySemaphore;
  vsi.pWaitDstStageMask = &waitAtBottom;
  vsi.commandBufferCount = 1;
  vsi.pCommandBuffers = &CommandBuffers[nextImageIndex];
  vsi.signalSemaphoreCount = 0;
  vsi.pSignalSemaphores = (VkSemaphore) nullptr;
  You do this to wait for an image
  to be ready to be rendered into

result = vkQueueSubmit( presentQueue, 1, IN &vsi, IN renderFence );
    
```

mb - July 24, 2020

Fences

376

- Used when the host needs to wait for the device to complete something big
- Used to synchronize the application with commands submitted to a queue
- Announces that queue-submitted work is finished
- Much finer control than semaphores
- You can un-signal, signal, test or block-while-waiting

mb - July 24, 2020

Fences

377

```

#define VK_FENCE_CREATE_UNSIGNALED_BIT 0

VkFenceCreateInfo
  vfci;
  vfci.sType = VK_STRUCTURE_TYPE_FENCE_CREATE_INFO;
  vfci.pNext = nullptr;
  vfci.flags = VK_FENCE_CREATE_UNSIGNALED_BIT; // = 0
  // VK_FENCE_CREATE_UNSIGNALED_BIT is only other option

VkFence fence;
result = vkCreateFence( LogicalDevice, IN &vfci, PALLOCATOR, OUT &fence );
  ...
  Set the fence

// returns to the host right away:
result = vkGetFenceStatus( LogicalDevice, IN fence );
// result = VK_SUCCESS means it has signaled
// result = VK_NOT_READY means it has not signaled

// blocks the host from executing:
result = vkWaitForFences( LogicalDevice, IN &fence, waitForAll, timeout );
// waitForAll = VK_TRUE: wait for all fences in the list
// waitForAll = VK_FALSE: wait for any one fence in the list
// timeout is a uint64_t timeout in nanoseconds (could be 0, which means to return immediately)
// timeout can be up to UINT64_MAX = 0xffffffffffff ( = 580+ years )
// result = VK_SUCCESS means it returned because a fence (or all fences) signaled
// result = VK_TIMEOUT means it returned because the timeout was exceeded
    
```

mb - July 24, 2020

Fence Example

378

```

VkFence renderFence;
vkCreateFence( LogicalDevice, &vfci, PALLOCATOR, OUT &renderFence );

VkPipelineStageFlags waitAtBottom = VK_PIPELINE_STAGE_BOTTOM_OF_PIPE_BIT;
VkQueue presentQueue;
vkGetDeviceQueue( LogicalDevice, FindQueueFamilyThatDoesGraphics( ), 0, OUT &presentQueue );

VkSubmitInfo
  vsi;
  vsi.sType = VK_STRUCTURE_TYPE_SUBMIT_INFO;
  vsi.pNext = nullptr;
  vsi.waitSemaphoreCount = 1;
  vsi.pWaitSemaphores = &imageReadySemaphore;
  vsi.pWaitDstStageMask = &waitAtBottom;
  vsi.commandBufferCount = 1;
  vsi.pCommandBuffers = &CommandBuffers[nextImageIndex];
  vsi.signalSemaphoreCount = 0;
  vsi.pSignalSemaphores = (VkSemaphore) nullptr;


result = vkQueueSubmit( presentQueue, 1, IN &vsi, IN renderFence );
  ...
  Wait on the fence(s)

result = vkWaitForFences( LogicalDevice, 1, IN &renderFence, VK_TRUE, UINT64_MAX );
  ...
  result = vkQueuePresentKHR( presentQueue, IN &pi );
    
```

mb - July 24, 2020

Events 379

- Events provide even finer-grained synchronization
- Events are a primitive that can be signaled by the host or the device
- Can even signal at one place in the pipeline and wait for it at another place in the pipeline
- Signaling in the pipeline means "signal me as the last piece of this draw command passes that point in the pipeline".
- You can signal, un-signal, or test from a vk function or from a vkCmd function
- Can wait from a vkCmd function

 mb - July 24, 2020

Controlling Events from the Host 380

```

VkEventCreateInfo
veci.sType = VK_STRUCTURE_TYPE_EVENT_CREATE_INFO;
veci.pNext = nullptr;
veci.flags = 0;


VkEvent event;
result = vkCreateEvent( LogicalDevice, IN &veci, PALLOCATOR, OUT &event );

result = vkSetEvent( LogicalDevice, IN event );

result = vkResetEvent( LogicalDevice, IN event );

result = vkGetEventStatus( LogicalDevice, IN event );
// result = VK_EVENT_SET, signaled
// result = VK_EVENT_RESET, not signaled
  
```

Note: the host cannot block waiting for an event, but it can test for it

 mb - July 24, 2020

Controlling Events from the Device 381

```


result = vkCmdSetEvent( CommandBuffer, IN event, pipelineStageBits );
result = vkCmdResetEvent( CommandBuffer, IN event, pipelineStageBits );
result = vkCmdWaitEvents( CommandBuffer, 1, &event,
    srcPipelineStageBits, dstPipelineStageBits,
    memoryBarrierCount, pMemoryBarriers,
    bufferMemoryBarrierCount, pBufferMemoryBarriers,
    imageMemoryBarrierCount, pImageMemoryBarriers
);
  
```


Could be an array of events

Where signaled, where wait for the signal


Memory barriers get executed after events have been signaled

Note: the device cannot test for an event, but it can block


 mb - July 24, 2020


Pipeline Barriers

Mike Bailey
mjb@cs.oregonstate.edu


This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License


<http://cs.oregonstate.edu/~mjb/vulkan>

 mb - July 24, 2020

From the Command Buffer Notes: 383
These are the Commands that can be entered into the Command Buffer, I

```


vkCmdBeginQuery( commandBuffer, flags );
vkCmdBeginRenderPass( commandBuffer, const contents );
vkCmdBindDescriptorSet( commandBuffer, pDynamicOffsets );
vkCmdBindIndexBuffer( commandBuffer, indexType );
vkCmdBindPipeline( commandBuffer, pipeline );
vkCmdBindVertexBuffer( commandBuffer, firstBinding, bindingCount, const pOffsets );
vkCmdBindImage( commandBuffer, filler );
vkCmdClearAttachments( commandBuffer, attachmentCount, const pRects );
vkCmdClearColorImage( commandBuffer, pRanges );
vkCmdClearDepthStencilImage( commandBuffer, pRanges );
vkCmdCopyBuffer( commandBuffer, pRegions );
vkCmdCopyImage( commandBuffer, pRegions );
vkCmdCopyImageToBuffer( commandBuffer, pRegions );
vkCmdCopyQueryPoolResults( commandBuffer, flags );
vkCmdDebugMarkerBeginEXT( commandBuffer, pMarkerInfo );
vkCmdDebugMarkerEndEXT( commandBuffer );
vkCmdDispatch( commandBuffer, groupCountX, groupCountY, groupCountZ );
vkCmdDispatchIndirect( commandBuffer, offset );
vkCmdDraw( commandBuffer, vertexCount, instanceCount, firstVertex, firstInstance );
vkCmdDrawIndexed( commandBuffer, indexCount, instanceCount, firstIndex, inst32, vertexOffset, firstInstance );
vkCmdDrawIndexedIndirect( commandBuffer, stride );
vkCmdDrawIndexedIndirectCountAMD( commandBuffer, stride );
vkCmdDrawIndirect( commandBuffer, stride );
vkCmdDrawIndirectCountAMD( commandBuffer, stride );
vkCmdEndQuery( commandBuffer, query );
vkCmdEndRenderPass( commandBuffer );
vkCmdExecuteCommands( commandBuffer, commandBufferCount, const pCommandBuffers );
  
```

 mb - July 24, 2020

From the Command Buffer Notes: 384
These are the Commands that can be entered into the Command Buffer, II

```

vkCmdFillBuffer( commandBuffer, dstBuffer, dstOffset, size, data );
vkCmdNextSubpass( commandBuffer, contents );
vkCmdPipelineBarrier( commandBuffer, srcStageMask, dstStageMask, dependencyFlags, memoryBarrierCount, pMemoryBarriers,
    bufferMemoryBarrierCount, pBufferMemoryBarriers, imageMemoryBarrierCount, pImageMemoryBarriers );
vkCmdProcessCommandsNVX( commandBuffer, pProcessCommandInfo );
vkCmdPushConstants( commandBuffer, layout, stageFlags, offset, size, pValues );
vkCmdPushDescriptorSetKHR( commandBuffer, pipelineBindPoint, layout, set, descriptorWriteCount, pDescriptorWrites );
vkCmdPushDescriptorSetWithTemplateKHR( commandBuffer, descriptorUpdateTemplate, layout, set, pData );
vkCmdReserveSpaceForCommandsNVX( commandBuffer, pReserveSpaceInfo );
vkCmdResetEvent( commandBuffer, event, stageMask );
vkCmdResetQueryPool( commandBuffer, queryPool, firstQuery, queryCount );
vkCmdResolveImage( commandBuffer, srcImage, srcImageLayout, dstImage, dstImageLayout, regionCount, pRegions );
vkCmdSetBlendConstants( commandBuffer, blendConstants[4] );
vkCmdSetDepthBias( commandBuffer, depthBiasConstantFactor, depthBiasClamp, depthBiasSlopeFactor );
vkCmdSetDepthBounds( commandBuffer, minDepthBounds, maxDepthBounds );
vkCmdSetDiscardRectangleEXT( commandBuffer, firstDiscardRectangle, discardRectangleCount, pDiscardRectangles );
vkCmdSetEvent( commandBuffer, event, stageMask );
vkCmdSetEventWait( commandBuffer, event, stageMask );
vkCmdSetLineWidth( commandBuffer, lineWidth );
vkCmdSetScissor( commandBuffer, firstScissor, scissorCount, pScissors );
vkCmdSetStencilCompareMask( commandBuffer, faceMask, compareMask );
vkCmdSetStencilReference( commandBuffer, faceMask, reference );
vkCmdSetStencilWriteMask( commandBuffer, faceMask, writeMask );
vkCmdSetViewport( commandBuffer, firstViewport, viewportCount, pViewports );
vkCmdSetViewportWScalingNV( commandBuffer, firstViewport, viewportCount, pViewportWScalings );
vkCmdSetVertex( commandBuffer, dstBuffer, dstOffset, dataSize, pData );
vkCmdWaitEvents( commandBuffer, eventCount, pEvents, srcStageMask, dstStageMask, memoryBarrierCount, pMemoryBarriers,
    bufferMemoryBarrierCount, pBufferMemoryBarriers, imageMemoryBarrierCount, pImageMemoryBarriers );
vkCmdWriteTimestamp( commandBuffer, pipelineStage, queryPool, query );
  
```

 mb - July 24, 2020

Potential Memory Race Conditions that Pipeline Barriers can Prevent

1. Write-then-Read (WtR) – the memory write in one operation starts overwriting the memory that another operation's read needs to use
2. Read-then-Write (RtW) – the memory read in one operation hasn't yet finished before another operation starts overwriting that memory
3. Write-then-Write (WtW) – two operations start overwriting the same memory and the end result is non-deterministic

Note: there is no problem with Read-then-Read (RtR) as no data has been changed

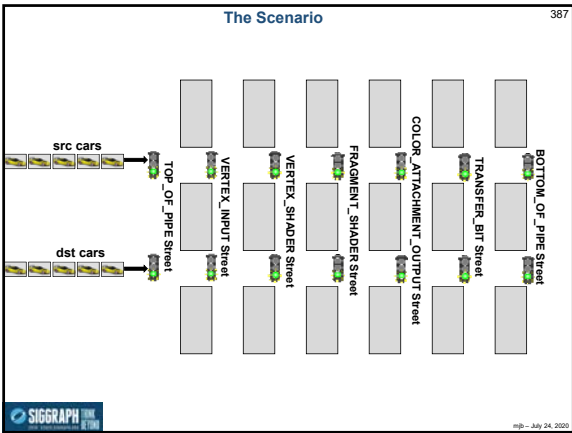
vkCmdPipelineBarrier() Function Call

A Pipeline Barrier is a way to establish a memory dependency between commands that were submitted before the barrier and commands that are submitted after the barrier

```

vkCmdPipelineBarrier( commandBuffer,
    srcStageMask,
    dstStageMask,
    VK_DEPENDENCY_BY_REGION_BIT,
    memoryBarrierCount,    pMemoryBarriers,
    bufferMemoryBarrierCount, pBufferMemoryBarriers,
    imageMemoryBarrierCount, pImageMemoryBarriers
);
    
```

- srcStageMask**: Guarantee that this pipeline stage is completely done being used before ...
- dstStageMask**: ... allowing this pipeline stage to be used
- pMemoryBarriers**, **pBufferMemoryBarriers**, **pImageMemoryBarriers**: Defines what data we will be blocking on or un-blocking on

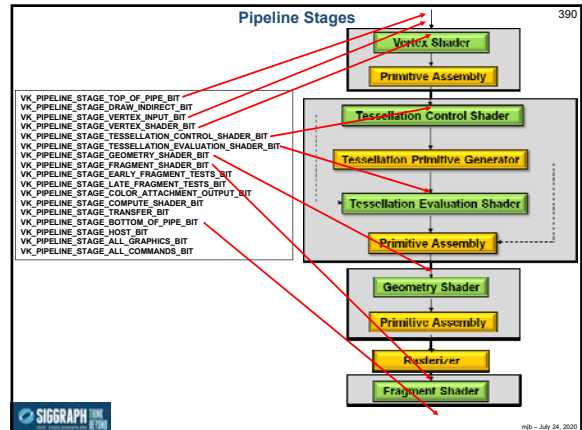


The Scenario

1. The cross-streets are named after pipeline stages
2. All traffic lights start out green
3. There are special sensors at all intersections that will know when **any car in the src group** is in that intersection
4. There are connections from those sensors to the traffic lights so that when **any car in the src group** is in the intersection, the proper **dst** traffic light will be turned red
5. When the **last car in the src group** completely makes it through its intersection, the proper **dst** traffic light is turned back to green
6. The Vulkan command pipeline ordering is this: (1) the **src** cars get released, (2) the pipeline barrier is invoked (which turns some light red), (3) the **dst** cars stop at the red light, (4) the **src** intersection clears, (5) all lights are now green, (6) the **dst** cars continue.

Pipeline Stage Masks – Where in the Pipeline is this Memory Data being Generated or Consumed?

- VK_PIPELINE_STAGE_TOP_OF_PIPE_BIT
- VK_PIPELINE_STAGE_DRAW_INDIRECT_BIT
- VK_PIPELINE_STAGE_VERTEX_INPUT_BIT
- VK_PIPELINE_STAGE_VERTEX_SHADER_BIT
- VK_PIPELINE_STAGE_TESSELLATION_CONTROL_SHADER_BIT
- VK_PIPELINE_STAGE_TESSELLATION_EVALUATION_SHADER_BIT
- VK_PIPELINE_STAGE_GEOMETRY_SHADER_BIT
- VK_PIPELINE_STAGE_FRAGMENT_SHADER_BIT
- VK_PIPELINE_STAGE_EARLY_FRAGMENT_TESTS_BIT
- VK_PIPELINE_STAGE_LATE_FRAGMENT_TESTS_BIT
- VK_PIPELINE_STAGE_COLOR_ATTACHMENT_OUTPUT_BIT
- VK_PIPELINE_STAGE_COMPUTE_SHADER_BIT
- VK_PIPELINE_STAGE_TRANSFER_BIT
- VK_PIPELINE_STAGE_BOTTOM_OF_PIPE_BIT
- VK_PIPELINE_STAGE_HOST_BIT
- VK_PIPELINE_STAGE_ALL_GRAPHICS_BIT
- VK_PIPELINE_STAGE_ALL_COMMANDS_BIT



Access Masks – What are you Interested in Generating or Consuming this Memory for?

VK_ACCESS_INDIRECT_COMMAND_READ_BIT
 VK_ACCESS_INDEX_READ_BIT
 VK_ACCESS_VERTEX_ATTRIBUTE_READ_BIT
 VK_ACCESS_UNIFORM_READ_BIT
 VK_ACCESS_INPUT_ATTACHMENT_READ_BIT
 VK_ACCESS_SHADER_READ_BIT
 VK_ACCESS_SHADER_WRITE_BIT
 VK_ACCESS_COLOR_ATTACHMENT_READ_BIT
 VK_ACCESS_COLOR_ATTACHMENT_WRITE_BIT
 VK_ACCESS_DEPTH_STENCIL_ATTACHMENT_READ_BIT
 VK_ACCESS_DEPTH_STENCIL_ATTACHMENT_WRITE_BIT
 VK_ACCESS_TRANSFER_READ_BIT
 VK_ACCESS_TRANSFER_WRITE_BIT
 VK_ACCESS_HOST_READ_BIT
 VK_ACCESS_HOST_WRITE_BIT
 VK_ACCESS_MEMORY_READ_BIT
 VK_ACCESS_MEMORY_WRITE_BIT

391

Pipeline Stages and what Access Operations are Allowed

Stage	VK_ACCESS_INDIRECT_COMMAND_READ_BIT	VK_ACCESS_INDEX_READ_BIT	VK_ACCESS_VERTEX_ATTRIBUTE_READ_BIT	VK_ACCESS_UNIFORM_READ_BIT	VK_ACCESS_INPUT_ATTACHMENT_READ_BIT	VK_ACCESS_SHADER_READ_BIT	VK_ACCESS_SHADER_WRITE_BIT	VK_ACCESS_COLOR_ATTACHMENT_READ_BIT	VK_ACCESS_COLOR_ATTACHMENT_WRITE_BIT	VK_ACCESS_DEPTH_STENCIL_ATTACHMENT_READ_BIT	VK_ACCESS_DEPTH_STENCIL_ATTACHMENT_WRITE_BIT	VK_ACCESS_TRANSFER_READ_BIT	VK_ACCESS_TRANSFER_WRITE_BIT	VK_ACCESS_HOST_READ_BIT	VK_ACCESS_HOST_WRITE_BIT	VK_ACCESS_MEMORY_READ_BIT	VK_ACCESS_MEMORY_WRITE_BIT
1 VK_PIPELINE_STAGE_TOP_OF_PIPE_BIT																	
2 VK_PIPELINE_STAGE_DRAW_INDIRECT_BIT																	
3 VK_PIPELINE_STAGE_VERTEX_INPUT_BIT																	
4 VK_PIPELINE_STAGE_VERTEX_SHADER_BIT																	
5 VK_PIPELINE_STAGE_TESSELLATION_CONTROL_SHADER_BIT																	
6 VK_PIPELINE_STAGE_TESSELLATION_EVALUATION_SHADER_BIT																	
7 VK_PIPELINE_STAGE_GEOMETRY_SHADER_BIT																	
8 VK_PIPELINE_STAGE_FRAGMENT_SHADER_BIT																	
9 VK_PIPELINE_STAGE_EARLY_FRAGMENT_TESTS_BIT																	
10 VK_PIPELINE_STAGE_LATE_FRAGMENT_TESTS_BIT																	
11 VK_PIPELINE_STAGE_COLOR_ATTACHMENT_OUTPUT_BIT																	
12 VK_PIPELINE_STAGE_BOTTOM_OF_PIPE_BIT																	
VK_PIPELINE_STAGE_COMPUTE_SHADER_BIT																	
VK_PIPELINE_STAGE_HOST_BIT																	
VK_PIPELINE_STAGE_ALL_GRAPHICS_BIT																	
VK_PIPELINE_STAGE_ALL_COMMANDS_BIT																	

392

Access Operations and what Pipeline Stages they can be used In

Access Operation	1	2	3	4	5	6	7	8	9	10	11	12	
VK_ACCESS_INDIRECT_COMMAND_READ_BIT													
VK_ACCESS_INDEX_READ_BIT													
VK_ACCESS_VERTEX_ATTRIBUTE_READ_BIT													
VK_ACCESS_UNIFORM_READ_BIT													
VK_ACCESS_INPUT_ATTACHMENT_READ_BIT													
VK_ACCESS_SHADER_READ_BIT													
VK_ACCESS_SHADER_WRITE_BIT													
VK_ACCESS_COLOR_ATTACHMENT_READ_BIT													
VK_ACCESS_COLOR_ATTACHMENT_WRITE_BIT													
VK_ACCESS_DEPTH_STENCIL_ATTACHMENT_READ_BIT													
VK_ACCESS_DEPTH_STENCIL_ATTACHMENT_WRITE_BIT													
VK_ACCESS_TRANSFER_READ_BIT													
VK_ACCESS_TRANSFER_WRITE_BIT													
VK_ACCESS_HOST_READ_BIT													
VK_ACCESS_HOST_WRITE_BIT													
VK_ACCESS_MEMORY_READ_BIT													
VK_ACCESS_MEMORY_WRITE_BIT													

393

Example: Be sure we are done writing an output image before using it for something else

Stages

VK_PIPELINE_STAGE_TOP_OF_PIPE_BIT
 VK_PIPELINE_STAGE_DRAW_INDIRECT_BIT
 VK_PIPELINE_STAGE_VERTEX_INPUT_BIT
 VK_PIPELINE_STAGE_VERTEX_SHADER_BIT
 VK_PIPELINE_STAGE_TESSELLATION_CONTROL_SHADER_BIT
 VK_PIPELINE_STAGE_TESSELLATION_EVALUATION_SHADER_BIT
 VK_PIPELINE_STAGE_GEOMETRY_SHADER_BIT
 VK_PIPELINE_STAGE_FRAGMENT_SHADER_BIT ← **src**
 VK_PIPELINE_STAGE_EARLY_FRAGMENT_TESTS_BIT
 VK_PIPELINE_STAGE_LATE_FRAGMENT_TESTS_BIT
 VK_PIPELINE_STAGE_COLOR_ATTACHMENT_OUTPUT_BIT
 VK_PIPELINE_STAGE_COMPUTE_SHADER_BIT
 VK_PIPELINE_STAGE_TRANSFER_BIT
 VK_PIPELINE_STAGE_BOTTOM_OF_PIPE_BIT ← **dst**
 VK_PIPELINE_STAGE_ALL_GRAPHICS_BIT
 VK_PIPELINE_STAGE_ALL_COMMANDS_BIT

Access types

VK_ACCESS_INDIRECT_COMMAND_READ_BIT
 VK_ACCESS_INDEX_READ_BIT
 VK_ACCESS_VERTEX_ATTRIBUTE_READ_BIT
 VK_ACCESS_UNIFORM_READ_BIT
 VK_ACCESS_INPUT_ATTACHMENT_READ_BIT
 VK_ACCESS_SHADER_READ_BIT ← **src**
 VK_ACCESS_SHADER_WRITE_BIT
 VK_ACCESS_COLOR_ATTACHMENT_READ_BIT
 VK_ACCESS_COLOR_ATTACHMENT_WRITE_BIT
 VK_ACCESS_DEPTH_STENCIL_ATTACHMENT_READ_BIT
 VK_ACCESS_DEPTH_STENCIL_ATTACHMENT_WRITE_BIT
 VK_ACCESS_TRANSFER_READ_BIT
 VK_ACCESS_TRANSFER_WRITE_BIT
 VK_ACCESS_HOST_READ_BIT
 VK_ACCESS_HOST_WRITE_BIT
 VK_ACCESS_MEMORY_READ_BIT
 VK_ACCESS_MEMORY_WRITE_BIT

dst (no access setting needed)

394

The Scenario

395

Example: Don't read a buffer back to the host until a shader is done writing it

Stages

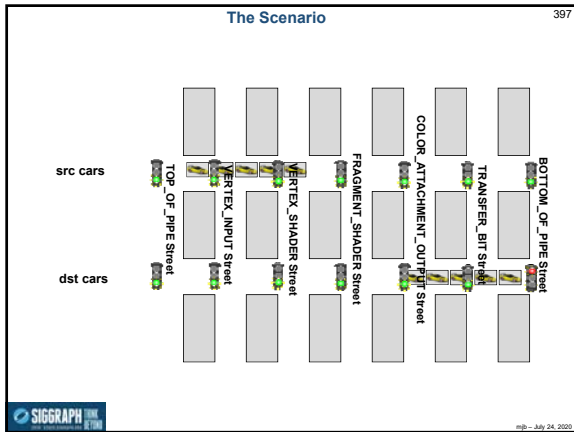
VK_PIPELINE_STAGE_TOP_OF_PIPE_BIT
 VK_PIPELINE_STAGE_DRAW_INDIRECT_BIT
 VK_PIPELINE_STAGE_VERTEX_SHADER_BIT ← **src**
 VK_PIPELINE_STAGE_TESSELLATION_CONTROL_SHADER_BIT
 VK_PIPELINE_STAGE_TESSELLATION_EVALUATION_SHADER_BIT
 VK_PIPELINE_STAGE_GEOMETRY_SHADER_BIT
 VK_PIPELINE_STAGE_FRAGMENT_SHADER_BIT
 VK_PIPELINE_STAGE_EARLY_FRAGMENT_TESTS_BIT
 VK_PIPELINE_STAGE_LATE_FRAGMENT_TESTS_BIT
 VK_PIPELINE_STAGE_COLOR_ATTACHMENT_OUTPUT_BIT
 VK_PIPELINE_STAGE_COMPUTE_SHADER_BIT
 VK_PIPELINE_STAGE_TRANSFER_BIT
 VK_PIPELINE_STAGE_BOTTOM_OF_PIPE_BIT ← **dst**
 VK_PIPELINE_STAGE_ALL_GRAPHICS_BIT
 VK_PIPELINE_STAGE_ALL_COMMANDS_BIT

Access types

VK_ACCESS_INDIRECT_COMMAND_READ_BIT
 VK_ACCESS_INDEX_READ_BIT
 VK_ACCESS_VERTEX_ATTRIBUTE_READ_BIT
 VK_ACCESS_UNIFORM_READ_BIT
 VK_ACCESS_INPUT_ATTACHMENT_READ_BIT
 VK_ACCESS_SHADER_READ_BIT ← **src**
 VK_ACCESS_SHADER_WRITE_BIT
 VK_ACCESS_COLOR_ATTACHMENT_READ_BIT
 VK_ACCESS_COLOR_ATTACHMENT_WRITE_BIT
 VK_ACCESS_DEPTH_STENCIL_ATTACHMENT_READ_BIT
 VK_ACCESS_DEPTH_STENCIL_ATTACHMENT_WRITE_BIT
 VK_ACCESS_TRANSFER_READ_BIT
 VK_ACCESS_TRANSFER_WRITE_BIT
 VK_ACCESS_HOST_READ_BIT
 VK_ACCESS_HOST_WRITE_BIT
 VK_ACCESS_MEMORY_READ_BIT
 VK_ACCESS_MEMORY_WRITE_BIT

dst (no access setting needed)

396



VkImageLayout – How an Image gets Laid Out in Memory depends on how it will be Used

```

VkImageMemoryBarrier vimb;
vimb.sType = VK_STRUCTURE_TYPE_IMAGE_MEMORY_BARRIER;
vimb.pNext = nullptr;
vimb.srcAccessMask = ??;
vimb.dstAccessMask = ??;
vimb.oldLayout = ??;
vimb.newLayout = ??;
vimb.srcQueueFamilyIndex = VK_QUEUE_FAMILY_IGNORED;
vimb.dstQueueFamilyIndex = VK_QUEUE_FAMILY_IGNORED;
vimb.image = ??;
vimb.subresourceRange = visir;
    
```

- VK_IMAGE_LAYOUT_UNDEFINED
- VK_IMAGE_LAYOUT_GENERAL
- VK_IMAGE_LAYOUT_COLOR_ATTACHMENT_OPTIMAL → Used as a color attachment
- VK_IMAGE_LAYOUT_DEPTH_STENCIL_ATTACHMENT_OPTIMAL
- VK_IMAGE_LAYOUT_DEPTH_STENCIL_READ_ONLY_OPTIMAL
- VK_IMAGE_LAYOUT_SHADER_READ_ONLY_OPTIMAL → Read into a shader as a texture
- VK_IMAGE_LAYOUT_TRANSFER_SRC_OPTIMAL → Copy from
- VK_IMAGE_LAYOUT_TRANSFER_DST_OPTIMAL → Copy to
- VK_IMAGE_LAYOUT_PREINITIALIZED
- VK_IMAGE_LAYOUT_PRESENT_SRC_KHR → Show image to viewer
- VK_IMAGE_LAYOUT_SHARED_PRESENT_KHR

Here, the use of vkCmdPipelineBarrier() is to simply change the layout of an image

SIGGRAPH 2020 mb - July 24, 2020

Vulkan.


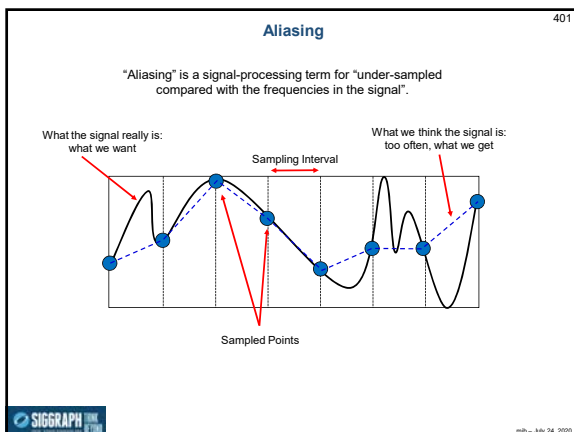

Antialiasing and Multisampling

Mike Bailey
mjb@cs.oregonstate.edu

<http://cs.oregonstate.edu/~mjb/vulkan>

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License

SIGGRAPH 2020 mb - July 24, 2020



The Nyquist Criterion

403

"The Nyquist [sampling] rate is twice the maximum component frequency of the function [i.e., signal] being sampled." -- Wikipedia

SIGGRAPH 2016

MultiSampling

404

Oversampling is a computer graphics technique to improve the quality of your output image by looking inside every pixel to see what the rendering is doing there.

There are two approaches to this:

- Supersampling:** Pick some number of sub-pixels within that pixel that pass the depth and stencil tests. Render the image at each of these sub-pixels.

- Multisampling:** Pick some number of sub-pixels within that pixel that pass the depth and stencil tests. If any of them pass, then perform a single color render for the one pixel and assign that single color to all the sub-pixels that passed the depth and stencil tests.

The final step will be to average those sub-pixels' colors to produce one final color for this whole pixel. This is called **resolving** the pixel.

SIGGRAPH 2016

Vulkan Specification Distribution of Sampling Points within a Pixel

405

SIGGRAPH 2016

Vulkan Specification Distribution of Sampling Points within a Pixel

406

VK_SAMPLE_COUNT_2_BIT	VK_SAMPLE_COUNT_4_BIT	VK_SAMPLE_COUNT_8_BIT	VK_SAMPLE_COUNT_16_BIT
	(0.375, 0.125)	(0.5625, 0.3125)	(0.5625, 0.5625)
		(0.4375, 0.6875)	(0.4375, 0.3125)
			(0.3125, 0.625)
(0.25, 0.25)		(0.8125, 0.5625)	(0.75, 0.4375)
	(0.875, 0.375)	(0.1875, 0.375)	(0.1875, 0.375)
		(0.3125, 0.1875)	(0.625, 0.8125)
		(0.1875, 0.8125)	(0.8125, 0.6875)
			(0.6875, 0.1875)
	(0.125, 0.625)	(0.0625, 0.4375)	(0.375, 0.875)
			(0.5, 0.0625)
(0.75, 0.75)		(0.6875, 0.9375)	(0.25, 0.125)
	(0.625, 0.875)	(0.9375, 0.0625)	(0.125, 0.75)
			(0.0, 0.5)
			(0.9375, 0.25)
			(0.875, 0.9375)
			(0.0625, 0.0)

SIGGRAPH 2016

Consider Two Triangles Who Pass Through the Same Pixel

407

Let's assume (for now) that the two triangles don't overlap – that is, they look this way because they butt up against each other.

SIGGRAPH 2016

Supersampling

408

$$\text{Final Pixel Color} = \frac{\sum_{i=1}^8 \text{Color sample from subpixel}_i}{8}$$

Fragment Shader calls = 8

SIGGRAPH 2016

Multisampling

Final Pixel Color = $\frac{3 \cdot \text{One color sample from A} + 5 \cdot \text{One color sample from B}}{8}$

Fragment Shader calls = 2

Consider Two Triangles Who Pass Through the Same Pixel

Let's assume (for now) that the two triangles don't overlap – that is, they look this way because they butt up against each other.

	Multisampling	Supersampling
Blue fragment shader calls	1	5
Red fragment shader calls	1	3

Consider Two Triangles Who Pass Through the Same Pixel

Q: What if the blue triangle completely filled the pixel when it was drawn, and then the red one, which is closer to the viewer than the blue one, came along and partially filled the pixel?

A: The ideas are all still the same, but the blue one had to deal with 8 sub-pixels (instead of 5 like before). But, the red triangle came along and obsoleted 3 of those blue sub-pixels. Note that the resolved image will still turn out the same as before.

Consider Two Triangles Who Pass Through the Same Pixel

What if the blue triangle completely filled the pixel when it was drawn, and then the red one, which is closer to the viewer than the blue one, came along and partially filled the pixel?

	Multisampling	Supersampling
Blue fragment shader calls	1	8
Red fragment shader calls	1	3

Setting up the Image

```

VkPipelineMultisampleStateCreateInfo
vpmisci.Type = VK_STRUCTURE_TYPE_PIPELINE_MULTISAMPLE_STATE_CREATE_INFO;
vpmisci.pNext = nullptr;
vpmisci.flags = 0;
vpmisci.rasterizationSamples = VK_SAMPLE_COUNT_8_BIT;
vpmisci.sampleShadingEnable = VK_TRUE;
vpmisci.minSampleShading = 0.5f;
vpmisci.pSampleMask = (VkSampleMask*)nullptr;
vpmisci.alphaToCoverageEnable = VK_FALSE;
vpmisci.alphaToOneEnable = VK_FALSE;

VkGraphicsPipelineCreateInfo
vgpci.Type = VK_STRUCTURE_TYPE_GRAPHICS_PIPELINE_CREATE_INFO;
vgpci.pNext = nullptr;
...
vgpci.MultisampleState = &vpmisci;

result = vkCreateGraphicsPipelines(LogicalDevice, VK_NULL_HANDLE, 1, IN &vgpci,
PALLOCATOR, OUT pGraphicsPipeline);
    
```

Annotations:

- vpmisci:** points to the MultisampleStateCreateInfo struct.
- vkSampleCount_8_bit:** points to the rasterizationSamples field. Note: "How dense is the sampling".
- VK_TRUE:** points to the sampleShadingEnable field. Note: "VK_TRUE means to allow some sort of multisampling to take place".
- 0.5f:** points to the minSampleShading field.
- vgpci:** points to the GraphicsPipelineCreateInfo struct.

Setting up the Image

```

VkPipelineMultisampleStateCreateInfo vpmisci;
...
vpmisci.minSampleShading = 0.5;
...
    
```

At least this fraction of samples will get their own fragment shader calls (as long as they pass the depth and stencil tests).

- 0. produces simple multisampling
- (0.,1) produces partial supersampling
- 1. Produces complete supersampling

Setting up the Image

```

VkAttachmentDescription
vad[0].format = VK_FORMAT_B8G8R8A8_SRGB;
vad[0].samples = VK_SAMPLE_COUNT_8_BIT;
vad[0].loadOp = VK_ATTACHMENT_LOAD_OP_CLEAR;
vad[0].storeOp = VK_ATTACHMENT_STORE_OP_STORE;
vad[0].stencilLoadOp = VK_ATTACHMENT_LOAD_OP_DONT_CARE;
vad[0].stencilStoreOp = VK_ATTACHMENT_STORE_OP_DONT_CARE;
vad[0].initialLayout = VK_IMAGE_LAYOUT_UNDEFINED;
vad[0].finalLayout = VK_IMAGE_LAYOUT_PRESENT_SRC_KHR;
vad[0].flags = 0;

vad[1].format = VK_FORMAT_D32_SFLOAT_S8_UINT;
vad[1].samples = VK_SAMPLE_COUNT_8_BIT;
vad[1].loadOp = VK_ATTACHMENT_LOAD_OP_CLEAR;
vad[1].storeOp = VK_ATTACHMENT_STORE_OP_DONT_CARE;
vad[1].stencilLoadOp = VK_ATTACHMENT_LOAD_OP_DONT_CARE;
vad[1].stencilStoreOp = VK_ATTACHMENT_STORE_OP_DONT_CARE;
vad[1].initialLayout = VK_IMAGE_LAYOUT_UNDEFINED;
vad[1].finalLayout = VK_IMAGE_LAYOUT_DEPTH_STENCIL_ATTACHMENT_OPTIMAL;
vad[1].flags = 0;

VkAttachmentReference
colorReference.attachment = 0;
colorReference.layout = VK_IMAGE_LAYOUT_COLOR_ATTACHMENT_OPTIMAL;

VkAttachmentReference
depthReference.attachment = 1;
depthReference.layout = VK_IMAGE_LAYOUT_DEPTH_STENCIL_ATTACHMENT_OPTIMAL;
    
```

to next slide

Setting up the Image

```

VkSubpassDescription
vsd.flags = 0;
vsd.pipelineBindPoint = VK_PIPELINE_BIND_POINT_GRAPHICS;
vsd.inputAttachmentCount = 0;
vsd.pInputAttachments = (VkAttachmentReference *)nullptr;
vsd.colorAttachmentCount = 1;
vsd.pColorAttachments = &colorReference;
vsd.pResolveAttachments = (VkAttachmentReference *)nullptr;
vsd.pDepthStencilAttachment = &depthReference;
vsd.pPreserveAttachmentCount = 0;
vsd.pPreserveAttachments = (uint32_t *)nullptr;

VkRenderPassCreateInfo
vrpci.sType = VK_STRUCTURE_TYPE_RENDER_PASS_CREATE_INFO;
vrpci.pNext = nullptr;
vrpci.flags = 0;
vrpci.attachmentCount = 2; // color and depth/stencil
vrpci.attachments = vad;
vrpci.subpassCount = 1;
vrpci.pSubpasses = IN &vsd;
vrpci.dependencyCount = 0;
vrpci.pDependencies = (VkSubpassDependency *)nullptr;

result = vkCreateRenderPass(LogicalDevice, IN &vrpci, PALLOCATOR, OUT &RenderPass);
    
```

from previous slide

Resolving the Image: Converting the Multisampled Image to a VK_SAMPLE_COUNT_1_BIT image

```

VlOffset3D
vo3.x = 0;
vo3.y = 0;
vo3.z = 0;


VkExtent3D
ve3.width = Width;
ve3.height = Height;
ve3.depth = 1;

VkImageSubresourceLayers
visl.aspectMask = VK_IMAGE_ASPECT_COLOR_BIT;
visl.mipLevel = 0;
visl.baseArrayLayer = 0;
visl.layerCount = 1;

VkImageResolve
vir.srcSubresource = visl;
vir.srcOffset = vo3;
vir.dstSubresource = visl;
vir.dstOffset = vo3;
vir.extent = ve3;

vkCmdResolveImage(cmdBuffer, srcImage, srcImageLayout, dstImage, dstImageLayout, 1, IN &vir);
    
```

For the *imageLayout, use VK_IMAGE_LAYOUT_GENERAL



Multipass Rendering

Mike Bailey
mjb@cs.oregonstate.edu

<http://cs.oregonstate.edu/~mjb/vulkan>

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License

Multipass Rendering uses Attachments -- What is a Vulkan Attachment Anyway?

"[An attachment is] an image associated with a renderpass that can be used as the input or output of one or more of its subpasses."
-- Vulkan Programming Guide

An attachment can be written to, read from, or both.

For example:

```

graph TD
    Attachment1[Attachment] --> Subpass1[Subpass]
    Attachment1 --> Subpass2[Subpass]
    Attachment1 --> Subpass3[Subpass]
    Subpass1 --> Framebuffer[Framebuffer]
    Subpass2 --> Framebuffer
    Subpass3 --> Framebuffer
    
```

What is an Example of Wanting to do This?

There is a process in computer graphics called *Deferred Rendering*. The idea is that a game-quality fragment shader takes a long time (relatively) to execute, but, with all the 3D scene detail, a lot of the rendered fragments are going to get z-buffered away anyhow. So, why did we invoke the fragment shaders so many times when we didn't need to?

Here's the trick:

Let's create a grossly simple fragment shader that writes out (into multiple framebuffers) each fragment's:

- position (x,y,z)
- normal (nx,ny,nz)
- material color (r,g,b)
- texture coordinates (s,t)

As well as:

- the current light source positions and colors
- the current eye position

When we write these out, the final framebuffers will contain just information for the pixels that *can be seen*. We then make a second pass running the expensive lighting model *just* for those pixels. This known as the *G-buffer Algorithm*.

Back in Our Single-pass Days

421

So far, we've only performed single-pass rendering, within a single Vulkan RenderPass.

```

    graph TD
      subgraph "3D Rendering Pass"
        A1[Attachment #1  
Depth Attachment]
        A0[Attachment #0  
Output]
      end
      subgraph "Subpass #0"
        A1
        A0
      end
  
```

Here comes a quick reminder of how we did that.

Afterwards, we will extend it.

SIGGRAPH 2020

Back in Our Single-pass Days, I

422

```

VkAttachmentDescription
vad[0] flags = 0;
vad[0] format = VK_FORMAT_B8G8R8A8_SRGB;
vad[0] samples = VK_SAMPLE_COUNT_1_BIT;
vad[0] loadOp = VK_ATTACHMENT_LOAD_OP_CLEAR;
vad[0] storeOp = VK_ATTACHMENT_STORE_OP_STORE;
vad[0] stencilLoadOp = VK_ATTACHMENT_LOAD_OP_DONT_CARE;
vad[0] stencilStoreOp = VK_ATTACHMENT_STORE_OP_DONT_CARE;
vad[0] initialLayout = VK_IMAGE_LAYOUT_UNDEFINED;
vad[0] finalLayout = VK_IMAGE_LAYOUT_PRESENT_SRC_KHR;

vad[1] flags = 0;
vad[1] format = VK_FORMAT_D32_SFLOAT_S8_UINT;
vad[1] samples = VK_SAMPLE_COUNT_1_BIT;
vad[1] loadOp = VK_ATTACHMENT_LOAD_OP_CLEAR;
vad[1] storeOp = VK_ATTACHMENT_STORE_OP_DONT_CARE;
vad[1] stencilLoadOp = VK_ATTACHMENT_LOAD_OP_DONT_CARE;
vad[1] stencilStoreOp = VK_ATTACHMENT_STORE_OP_DONT_CARE;
vad[1] initialLayout = VK_IMAGE_LAYOUT_UNDEFINED;
vad[1] finalLayout = VK_IMAGE_LAYOUT_DEPTH_STENCIL_ATTACHMENT_OPTIMAL;

VkAttachmentReference
colorReference.attachment = 0;
colorReference.layout = VK_IMAGE_LAYOUT_COLOR_ATTACHMENT_OPTIMAL;

VkAttachmentReference
depthReference.attachment = 1;
depthReference.layout = VK_IMAGE_LAYOUT_DEPTH_STENCIL_ATTACHMENT_OPTIMAL;
  
```

SIGGRAPH 2020

Back in Our Single-pass Days, II

423

```

VkSubpassDescription
vsd flags = 0;
vsd pipelineBindPoint = VK_PIPELINE_BIND_POINT_GRAPHICS;
vsd inputAttachmentCount = 0;
vsd pInputAttachments = (VkAttachmentReference *) nullptr;
vsd colorAttachmentCount = 1;
vsd pColorAttachments = &colorReference;
vsd pResolveAttachments = (VkAttachmentReference *) nullptr;
vsd pDepthStencilAttachment = &depthReference;
vsd preserveAttachmentCount = 0;
vsd pPreserveAttachments = (uint32_t *) nullptr;

VkRenderPassCreateInfo
vrpci.sType = VK_STRUCTURE_TYPE_RENDER_PASS_CREATE_INFO;
vrpci.pNext = nullptr;
vrpci.flags = 0;
vrpci.attachmentCount = 2; // color and depth/stencil
vrpci.pAttachments = vad;
vrpci.subpassCount = 1;
vrpci.pSubpasses = &vsd;
vrpci.dependencyCount = 0;
vrpci.pDependencies = (VkSubpassDependency *) nullptr;

result = vkCreateRenderPass( LogicalDevice, IN &vrpci, ALLOCATOR, OUT &RenderPass );
  
```

SIGGRAPH 2020

Multipass Rendering

424

So far, we've only performed single-pass rendering, but within a single Vulkan RenderPass, we can also have several subpasses, each of which is feeding information to the next subpass or subpasses.

In this case, we will look at following up a 3D rendering with Gbuffer operations.

SIGGRAPH 2020

Multipass, I

425

```

VkAttachmentDescription
vad[0] flags = 0;
vad[0] format = VK_FORMAT_D32_SFLOAT_S8_UINT;
vad[0] samples = VK_SAMPLE_COUNT_1_BIT;
vad[0] loadOp = VK_ATTACHMENT_LOAD_OP_DONT_CARE;
vad[0] storeOp = VK_ATTACHMENT_STORE_OP_DONT_CARE;
vad[0] stencilLoadOp = VK_ATTACHMENT_LOAD_OP_DONT_CARE;
vad[0] stencilStoreOp = VK_ATTACHMENT_STORE_OP_DONT_CARE;
vad[0] initialLayout = VK_IMAGE_LAYOUT_UNDEFINED;
vad[0] finalLayout = VK_IMAGE_LAYOUT_UNDEFINED;

vad[1] flags = 0;
vad[1] format = VK_FORMAT_R32G32B32A32_UINT;
vad[1] samples = VK_SAMPLE_COUNT_1_BIT;
vad[1] loadOp = VK_ATTACHMENT_LOAD_OP_DONT_CARE;
vad[1] storeOp = VK_ATTACHMENT_STORE_OP_STORE;
vad[1] stencilLoadOp = VK_ATTACHMENT_LOAD_OP_DONT_CARE;
vad[1] stencilStoreOp = VK_ATTACHMENT_STORE_OP_DONT_CARE;
vad[1] initialLayout = VK_IMAGE_LAYOUT_UNDEFINED;
vad[1] finalLayout = VK_IMAGE_LAYOUT_UNDEFINED;

vad[2] flags = 0;
vad[2] format = VK_FORMAT_R8G8B8A8_SRGB;
vad[2] samples = VK_SAMPLE_COUNT_1_BIT;
vad[2] loadOp = VK_ATTACHMENT_LOAD_OP_DONT_CARE;
vad[2] storeOp = VK_ATTACHMENT_STORE_OP_STORE;
vad[2] stencilLoadOp = VK_ATTACHMENT_LOAD_OP_DONT_CARE;
vad[2] stencilStoreOp = VK_ATTACHMENT_STORE_OP_DONT_CARE;
vad[2] initialLayout = VK_IMAGE_LAYOUT_UNDEFINED;
vad[2] finalLayout = VK_IMAGE_LAYOUT_PRESENT_SRC;
  
```

SIGGRAPH 2020

Multipass, II

426

```

VkAttachmentReference
depthOutput.attachment = 0;
depthOutput.layout = VK_IMAGE_LAYOUT_DEPTH_STENCIL_ATTACHMENT_OPTIMAL;

VkAttachmentReference
gBufferInput.attachment = 0;
gBufferInput.layout = VK_IMAGE_LAYOUT_COLOR_ATTACHMENT_OPTIMAL;

VkAttachmentReference
gBufferOutput.attachment = 1;
gBufferOutput.layout = VK_IMAGE_LAYOUT_COLOR_ATTACHMENT_OPTIMAL;

VkAttachmentReference
lightingInput[2].attachment = 0;
lightingInput[0].layout = VK_IMAGE_LAYOUT_DEPTH_STENCIL_READ_ONLY_OPTIMAL;
lightingInput[1].attachment = 1;
lightingInput[1].layout = VK_IMAGE_LAYOUT_SHADER_READ_ONLY_OPTIMAL;

VkAttachmentReference
lightingOutput.attachment = 2;
lightingOutput.layout = VK_IMAGE_LAYOUT_COLOR_ATTACHMENT_OPTIMAL;
  
```

SIGGRAPH 2020

Multipass, III 427

```

VkSubpassDescription
vsdp[0].flags = 0;
vsdp[0].pipelineBindPoint = VK_PIPELINE_BIND_POINT_GRAPHICS;
vsdp[0].inputAttachmentCount = 0;
vsdp[0].pInputAttachments = (VkAttachmentReference *) nullptr;
vsdp[0].colorAttachmentCount = 0;
vsdp[0].pColorAttachments = (VkAttachmentReference *) nullptr;
vsdp[0].pResolveAttachments = (VkAttachmentReference *) nullptr;
vsdp[0].pDepthStencilAttachment = &depthOutput;
vsdp[0].preserveAttachmentCount = 0;
vsdp[0].pPreserveAttachments = (uint32_t *) nullptr;

vsdp[1].flags = 0;
vsdp[1].pipelineBindPoint = VK_PIPELINE_BIND_POINT_GRAPHICS;
vsdp[1].inputAttachmentCount = 0;
vsdp[1].pInputAttachments = (VkAttachmentReference *) nullptr;
vsdp[1].colorAttachmentCount = 1;
vsdp[1].pColorAttachments = &gbufferOutput;
vsdp[1].pResolveAttachments = (VkAttachmentReference *) nullptr;
vsdp[1].pDepthStencilAttachment = (VkAttachmentReference *) nullptr;
vsdp[1].preserveAttachmentCount = 0;
vsdp[1].pPreserveAttachments = (uint32_t *) nullptr;

vsdp[2].flags = 0;
vsdp[2].pipelineBindPoint = VK_PIPELINE_BIND_POINT_GRAPHICS;
vsdp[2].inputAttachmentCount = 2;
vsdp[2].pInputAttachments = &lightingInput[0];
vsdp[2].colorAttachmentCount = 1;
vsdp[2].pColorAttachments = &lightingOutput;
vsdp[2].pResolveAttachments = (VkAttachmentReference *) nullptr;
vsdp[2].pDepthStencilAttachment = (VkAttachmentReference *) nullptr;
vsdp[2].preserveAttachmentCount = 0;
vsdp[2].pPreserveAttachments = (uint32_t *) nullptr;

```

SIGGRAPH 2020 mb - July 24, 2020

Multipass, IV 428

```

VkSubpassDependency
vsdp[0].srcSubpass = 0; // depth rendering ->
vsdp[0].dstSubpass = 1; // gbuffer
vsdp[0].srcStageMask = VK_PIPELINE_STAGE_COLOR_ATTACHMENT_OUTPUT_BIT;
vsdp[0].dstStageMask = VK_PIPELINE_STAGE_FRAGMENT_SHADER_BIT;
vsdp[0].srcAccessMask = VK_ACCESS_COLOR_ATTACHMENT_WRITE_BIT;
vsdp[0].dstAccessMask = VK_ACCESS_SHADER_READ_BIT;
vsdp[0].dependencyFlags = VK_DEPENDENCY_BY_REGION_BIT;

vsdp[1].srcSubpass = 1; // gbuffer ->
vsdp[1].dstSubpass = 2; // color output
vsdp[1].srcStageMask = VK_PIPELINE_STAGE_COLOR_ATTACHMENT_OUTPUT_BIT;
vsdp[1].dstStageMask = VK_PIPELINE_STAGE_FRAGMENT_SHADER_BIT;
vsdp[1].srcAccessMask = VK_ACCESS_COLOR_ATTACHMENT_WRITE_BIT;
vsdp[1].dstAccessMask = VK_ACCESS_SHADER_READ_BIT;
vsdp[1].dependencyFlags = VK_DEPENDENCY_BY_REGION_BIT;

```

Notice how similar this is to creating a **Directed Acyclic Graph (DAG)**.

SIGGRAPH 2020 mb - July 24, 2020

Multipass, V 429

```

VkRenderPassCreateInfo
vrpci.sType = VK_STRUCTURE_TYPE_RENDER_PASS_CREATE_INFO;
vrpci.pNext = nullptr;
vrpci.flags = 0;
vrpci.attachmentCount = 3; // depth, gbuffer, output
vrpci.pAttachments = vad;
vrpci.subpassCount = 3;
vrpci.pSubpasses = vsdp;
vrpci.dependencyCount = 2;
vrpci.pDependencies = vsdp;

result = vkCreateRenderPass( LogicalDevice, IN &vrpci, PALLOCATOR, OUT &RenderPass );

```

SIGGRAPH 2020 mb - July 24, 2020

Multipass, VI 430

```

vkCmdBeginRenderPass( CommandBuffers[nextImageIndex], IN &vrpci, IN VK_SUBPASS_CONTENTS_INLINE );
// subpass #0 is automatically started here

vkCmdBindPipeline( CommandBuffers[nextImageIndex], VK_PIPELINE_BIND_POINT_GRAPHICS, GraphicsPipeline );
vkCmdBindDescriptorSets( CommandBuffers[nextImageIndex], VK_PIPELINE_BIND_POINT_GRAPHICS,
    GraphicsPipelineLayout, 0, 4, DescriptorSets, 0, (uint32_t *) nullptr );
vkCmdBindVertexBuffers( CommandBuffers[nextImageIndex], 0, 1, vBuffers, offsets );
vkCmdDraw( CommandBuffers[nextImageIndex], vertexCount, instanceCount, firstVertex, firstInstance );

vkCmdNextSubpass( CommandBuffers[nextImageIndex], VK_SUBPASS_CONTENTS_INLINE );
// subpass #1 is started here

vkCmdNextSubpass( CommandBuffers[nextImageIndex], VK_SUBPASS_CONTENTS_INLINE );
// subpass #2 is started here

vkCmdEndRenderPass( CommandBuffers[nextImageIndex] );

```

SIGGRAPH 2020 mb - July 24, 2020

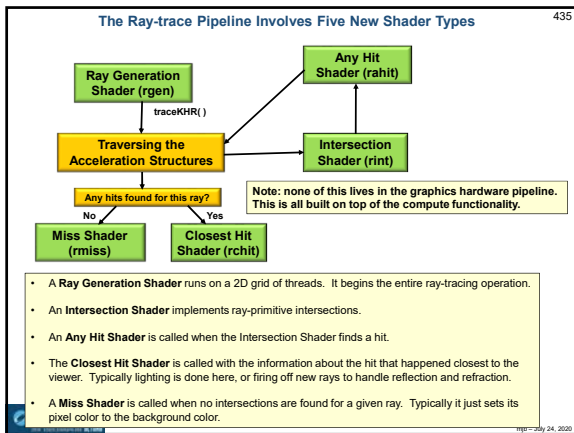
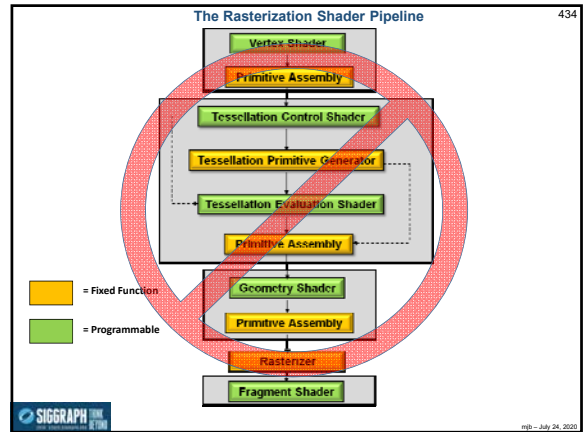

Vulkan.
Vulkan Ray Tracing

This work is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](https://creativecommons.org/licenses/by-nc-nd/4.0/).

SIGGRAPH 2020 mb - July 24, 2020

Analog Ray Tracing Example 432

SIGGRAPH 2020 mb - July 24, 2020



The Ray Intersection Process for a Sphere

1. Sphere equation: $(x-x_0)^2 + (y-y_0)^2 + (z-z_0)^2 = R^2$

2. Ray equation: $(x,y,z) = (x_0,y_0,z_0) + t*(dx,dy,dz)$

Plugging (x,y,z) from the second equation into the first equation and multiplying through and simplifying gives:

$$At^2 + Bt + C = 0$$

Solve for t_1, t_2

If both t_1 and t_2 are complex, then the ray missed the sphere.
 If $t_1 = t_2$, then the ray brushed the sphere at a tangent point.
 If both t_1 and t_2 are real and different, then the ray entered and exited the sphere.

In Vulkan terms:
 $gl_WorldRayOriginKHR = (x_0,y_0,z_0)$
 $gl_HitKHR = t$
 $gl_WorldRayDirectionKHR = (dx,dy,dz)$

436

SIGGRAPH 2018

mb - July 24, 2020

The Ray Intersection Process for a Cube

1. Plane equation: $Ax + By + Cz + D = 0$

2. Ray equation: $(x,y,z) = (x_0,y_0,z_0) + t*(dx,dy,dz)$

Plugging (x,y,z) from the second equation into the first equation and multiplying through and simplifying gives:

$$At + B = 0$$

Solve for t


A cube is actually the intersection of 6 half-space planes (just 4 are shown here). Each of these will produce its own t intersection value. Treat them as pairs: $(t_{x1}, t_{x2}), (t_{y1}, t_{y2}), (t_{z1}, t_{z2})$

The ultimate entry and exit values are:
 $t_{min} = \max(\min(t_{x1}, t_{x2}), \min(t_{y1}, t_{y2}), \min(t_{z1}, t_{z2}))$
 $t_{max} = \min(\max(t_{x1}, t_{x2}), \max(t_{y1}, t_{y2}), \max(t_{z1}, t_{z2}))$

437

SIGGRAPH 2018

mb - July 24, 2020



Acceleration Structures

439

- Bottom-level Acceleration Structure (BLAS) holds the vertex data and is built from vertex and index VkBuffers
- The BLAS can also hold transformations, but it looks like usually the BLAS holds vertices in the original Model Coordinates.
- Top-level Acceleration Structure (TLAS) holds a pointer to elements of the BLAS and a transformation.
- The BLAS is used as a Model Coordinate bounding box.
- The TLAS is used as a World Coordinate bounding box.
- A TLAS can instance multiple BLAS's.

Top Level Acceleration Structure

Bottom Level Acceleration Structure

sig - July 24, 2020

Creating Bottom Level Acceleration Structures

440

```

vkCreateAccelerationStructureKHR BottomLevelAccelerationStructure;
VkAccelerationStructureInfoKHR vasi;
vasi.sType = VK_ACCELERATION_STRUCTURE_TYPE_BOTTOM_LEVEL_KHR;
vasi.flags = 0;
vasi.pNext = nullptr;
vasi.instanceCount = 0;
vasi.geometryCount = << number of vertex buffers >>
vasi.pGeometries = << vertex buffer pointers >>
VkAccelerationStructureCreateInfoKHR vasci;
vasci.sType = VK_STRUCTURE_TYPE_ACCELERATION_STRUCTURE_CREATE_INFO_KHR;
vasci.pNext = nullptr;
vasci.info = &vasi;
vasci.compactedSize = 0;
result = vkCreateAccelerationStructureKHR(LogicalDevice, IN &vasci, PALLOCATOR, OUT &BottomLevelAccelerationStructure);
    
```

Top Level Acceleration Structure

Bottom Level Acceleration Structure

sig - July 24, 2020

Creating Top Level Acceleration Structures

441

```

vkCreateAccelerationStructureKHR TopLevelAccelerationStructure;
VkAccelerationStructureInfoKHR vasi;
vasi.sType = VK_ACCELERATION_STRUCTURE_TYPE_TOP_LEVEL_KHR;
vasi.flags = 0;
vasi.pNext = nullptr;
vasi.instanceCount = << number of bottom level acceleration structure instances >>
vasi.geometryCount = 0;
vasi.pGeometries = VK_NULL_HANDLE;
VkAccelerationStructureCreateInfoKHR vasci;
vasci.sType = VK_STRUCTURE_TYPE_ACCELERATION_STRUCTURE_CREATE_INFO_KHR;
vasci.pNext = nullptr;
vasci.info = &vasi;
vasci.compactedSize = 0;
result = vkCreateAccelerationStructureKHR(LogicalDevice, &vasci, PALLOCATOR, &TopLevelAccelerationStructure);
    
```

Top Level Acceleration Structure

Bottom Level Acceleration Structure

sig - July 24, 2020

Ray Generation Shader

442

Gets all of the rays going and writes the final color to the pixel

```

layout(location = 1) rayPayloadKHR myPayload;
vec4 color;
void main()
{
    traceKHR( topLevel, ... );
    imageStore( framebuffer, gl_GlobalInvocationIDKHR.xy, color );
}
    
```

A "payload" is information that keeps getting passed through the process. Different stages can add to it. It is finally consumed at the very end, in this case by writing color into the pixel being worked on.

sig - July 24, 2020

A New Built-in Function

443

```

void traceKHR
(
    accelerationStructureKHR topLevel,
    uint rayFlags,
    uint cullMask,
    uint sbtRecordOffset,
    uint sbtRecordStride,
    uint missIndex,
    vec3 origin,
    float tmin,
    vec3 direction,
    float tmax,
    int payload
);
    
```

In Vulkan terms:
 gl_WorldRayOriginKHR = (x₀, y₀, z₀)
 gl_HitKHR = t
 gl_WorldRayDirectionKHR = (dx, dy, dz)

sig - July 24, 2020

Intersection Shader

444

Intersect a ray with an arbitrary 3D object. Passes data to the Any Hit shader. There is a built-in ray-triangle Intersection Shader.

```

hitAttributeKHR vec3 attribs
void main()
{
    SpherePrimitive sph = spheres[ gl_PrimitiveID ];
    vec3 orig = gl_WorldRayOriginKHR;
    vec3 dir = normalize( gl_WorldRayDirectionKHR );
    ...
    float discr = b*b - 4.*a*c;
    if( discr < 0. )
        return;
    float tmp = (-b - sqrt(discr)) / (2.*a);
    if( gl_RayTminKHR < tmp && tmp < gl_RayTmaxKHR )
    {
        vec3 p = orig + tmp * dir;
        attribs = p;
        reportIntersectionKHR( tmp, 0 );
        return;
    }
    tmp = (-b + sqrt(discr)) / (2.*a);
    if( gl_RayTminKHR < tmp && tmp < gl_RayTmaxKHR )
    {
        vec3 p = orig + tmp * dir;
        attribs = p;
        reportIntersectionKHR( tmp, 0 );
        return;
    }
}
    
```

sig - July 24, 2020

Miss Shader 445

Handle a ray that doesn't hit any objects

```

rayPayloadKHR myPayload
{
    vec4 color;
};
void main()
{
    color = vec4( 0., 0., 0., 1. );
}
    
```

SIGGRAPH 2020 mb - July 24, 2020

Any Hit Shader 446

Handle a ray that hits anything. Store information on each hit. Can reject a hit.

```

layout( binding = 4, set = 0) buffer outputProperties
{
    float outputValues[];
}
outputData;

layout(location = 0) rayPayloadInKHR uint outputId;
layout(location = 1) rayPayloadInKHR uint hitCounter;
hitAttributeKHR vec 3 attribs;

void main()
{
    outputData.outputValues[ outputId + hitCounter ] = gl_PrimitiveID;
    hitCounter = hitCounter + 1;
}
    
```

SIGGRAPH 2020 mb - July 24, 2020

Closest Hit Shader 447

Handle the intersection closest to the viewer. Collects data from the Any Hit shader. Can spawn more rays.

```

rayPayloadKHR myPayload
{
    vec4 color;
};
void main()
{
    vec3 stp = gl_WorldRayOriginKHR + gl_HitKHR * gl_WorldRayDirectionKHR;
    color = texture( MaterialUnit, stp ); // material properties lookup
}
    
```

In Vulkan terms:
 gl_WorldRayOriginKHR = (x₀, y₀, z₀)
 gl_HitKHR = t
 gl_WorldRayDirectionKHR = (dx, dy, dz)

SIGGRAPH 2020 mb - July 24, 2020

Other New Built-in Functions 448

```

void terminateRayKHR( );
void ignoreIntersectionKHR( );
void reportIntersectionKHR( float hit, uint hitKind );
    
```

Loosely equivalent to "discard"

SIGGRAPH 2020 mb - July 24, 2020

Ray Trace Pipeline Data Structure 449

```

VkPipeline RaytracePipeline;
VkPipelineLayout PipelineLayout;

VkPipelineLayoutCreateInfo vplci;
vplci.sType = VK_STRUCTURE_TYPE_PIPELINE_LAYOUT_CREATE_INFO;
vplci.pNext = nullptr;
vplci.flags = 0;
vplci.setLayoutCount = 1;
vplci.pSetLayouts = &descriptorSetLayout;
vplci.pushConstantRangeCount = 0;
vplci.pushConstantRanges = nullptr;

result = vkCreatePipelineLayout( LogicalDevice, IN &vplci, nullptr, OUT &PipelineLayout );

VkRayTracingPipelineCreateInfoKHR vtrpci;
vtrpci.sType = VK_STRUCTURE_TYPE_RAY_TRACING_PIPELINE_CREATE_INFO_KHR;
vtrpci.pNext = nullptr;
vtrpci.flags = 0;
vtrpci.stageCount = << # of shader stages in the ray-trace pipeline >>
vtrpci.pStages = << what those shader stages are >>
vtrpci.groupCount = << # of shader groups >>
vtrpci.pGroups = << pointer to the groups (a group is a combination of shader programs >>
vtrpci.maxRecursionDepth = << how many recursion layers deep the ray tracing is allowed to go >>
vtrpci.layout = PipelineLayout;
vtrpci.basePipelineHandle = VK_NULL_HANDLE;
vtrpci.basePipelineIndex = 0;

result = vkCreateRayTracingPipelinesKHR( LogicalDevice, PALLOCATOR, 1, IN &vtrpci, nullptr, OUT &RaytracePipeline );
    
```

SIGGRAPH 2020 mb - July 24, 2020

The Trigger comes from the Command Buffer: vkCmdBindPipeline() and vkCmdTraceRaysKHR() 450

```

vkCmdBindPipeline( CommandBuffer, VK_PIPELINE_BIND_POINT_RAYTRACING_KHR, RaytracePipeline );

vkCmdTraceRaysKHR( CommandBuffer,
    raygenShaderBindingTableBuffer, raygenShaderBindingOffset,
    missShaderBindingTableBuffer, missShaderBindingOffset, missShaderBindingStride,
    callableShaderBindingTableBuffer, hitShaderBindingOffset, hitShaderBindingStride,
    callableShaderBindingTableBuffer, callableShaderBindingOffset, callableShaderBindingStride,
    width, height, depth );
    
```

SIGGRAPH 2020 mb - July 24, 2020

451




<https://www.youtube.com/watch?v=QL7sXc2iNJ8>


SIGGRAPH THINK BEYOND

mb - July 24, 2020


452



2019 SIGGRAPH **THINK BEYOND**




Vulkan.



Computer Graphics

Introduction to the Vulkan Computer Graphics API

Mike Bailey
mjb@cs.oregonstate.edu



This work is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](#).

<http://cs.oregonstate.edu/~mjb/vulkan>

SIGGRAPH THINK BEYOND

FULL.ppt

mb - July 24, 2020