

1



Vulkan.

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>

FULL.pptx mjb – July 24, 2020

2

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
2020 SIGGRAPH.org BEYOND

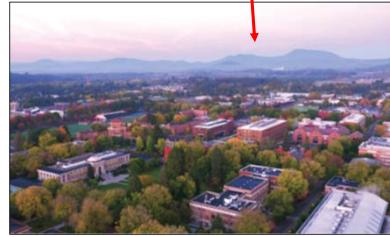
mjb – July 24, 2020

3

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

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

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

FULL.pptx mjb – July 24, 2020

4

Sections

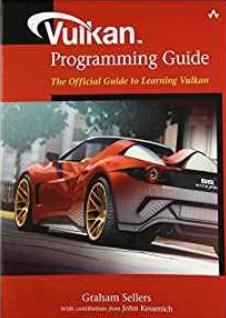
1. Introduction	13. Swap Chain
2. Sample Code	14. Push Constants
3. Drawing	15. Physical Devices
4. Shaders and SPIR-V	16. Logical Devices
5. Dats Buffers	17. Dynamic State Variables
6. GLFW	18. Getting Information Back
7. GLM	19. Compute Shaders
8. Instancing	20. Specialization Constants
9. Graphics 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
2020 SIGGRAPH.org BEYOND

mjb – July 24, 2020

5

My Favorite Vulkan Reference



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

SIGGRAPH THINK
CREATE. DESIGN. INNOVATE. BEYOND

mjb – July 24, 2020

6

Vulkan.

Introduction

Mike Bailey
mjb@cs.oregonstate.edu

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

SIGGRAPH THINK
CREATE. DESIGN. INNOVATE. BEYOND

mjb – July 24, 2020

7

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.)

KHRONOS GROUP

mjb – July 24, 2020

8

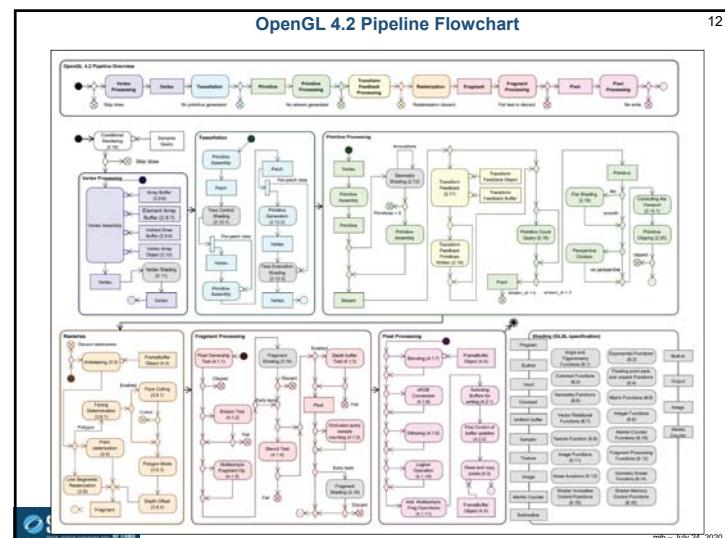
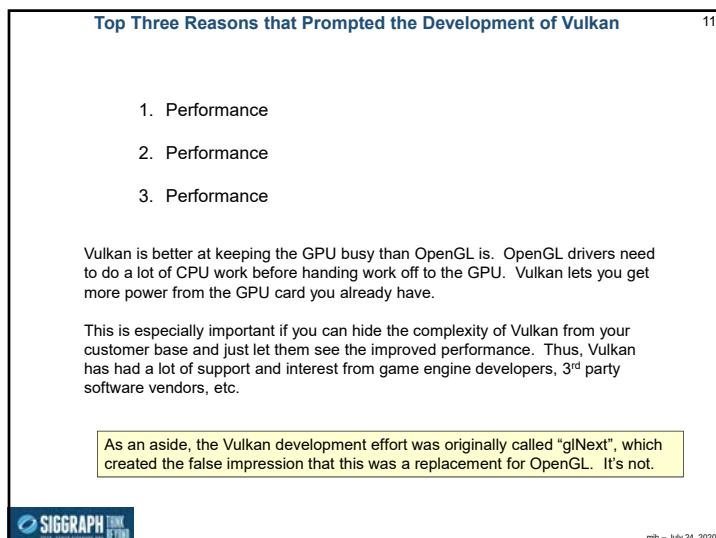
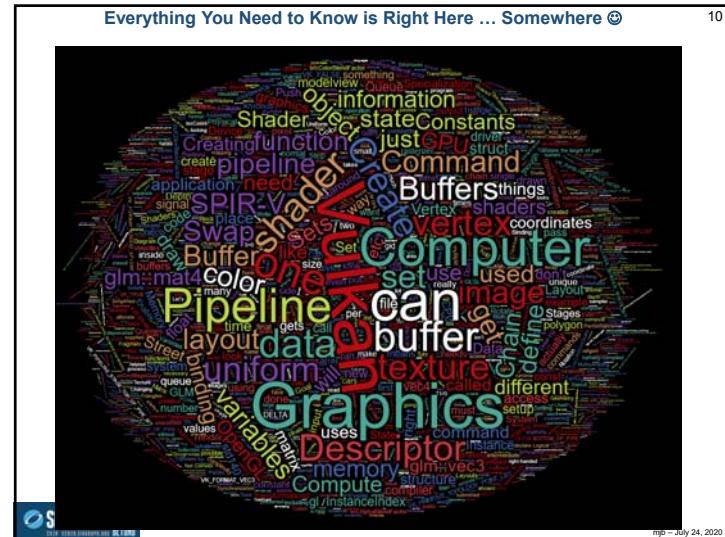
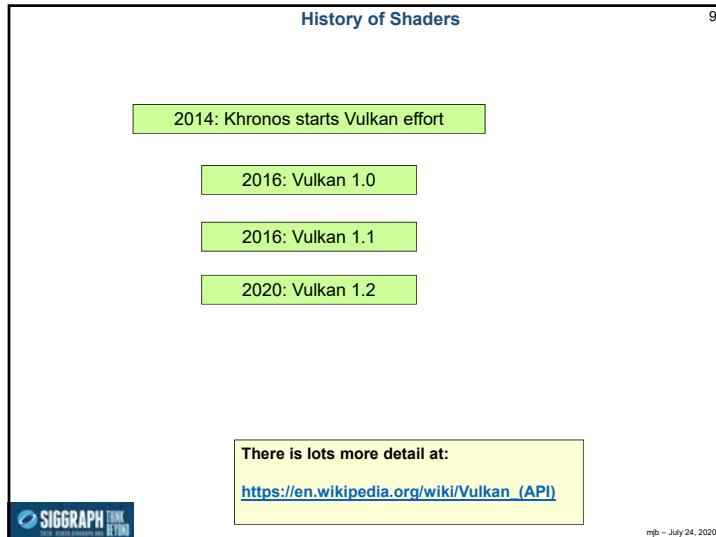
History of Shaders

- 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 THINK
CREATE. DESIGN. INNOVATE. BEYOND

mjb – July 24, 2020



Why is it so important to keep the GPU Busy?

13

NVIDIA Titan V Specs vs. Titan Xp, 1080 Ti					
	Titan V	Tesla V100	Tesla P100	GTX 1080 Ti	GTX 1080
GPU	GV100	GV100	GP100 Cut-Down Pascal	GP102 Pascal	GP104-400 Pascal
Transistor Count	21.1B	21.1B	15.3B	12B	7.2B
Fab Process	12nm FINFET	12nm FINFET	16nm FINFET	16nm FINFET	16nm FINFET
CUDA Cores / Tensor Cores	5120 / 640	5120 / 640	3584 / 0	3584 / 0	2560 / 0
TFLOPs	320	224	224	160	64
Core Clock	1200MHz	1230MHz	1200MHz	1000MHz	1733MHz
Boost Clock	1455MHz	1370MHz	1600MHz	-	-
FP32 TFLOPs	151TFLOPs	141TFLOPs	10.6TFLOPs	-11.4TFLOPs	9TFLOPs
Memory Type	HBM2	HBM2	HBM2	GDDR5X	GDDR5X
Memory Capacity	12GB	16GB	16GB	11GB	8GB
Memory Clock	17Gbps HBM2	17.5Gbps HBM2	7	11Gbps	10Gbps GDDR5X
Memory Interface	3072-bit	4096-bit	4096-bit	32-bit	256-bit
Memory Bandwidth	653.9GB/s	900.0GB/s	7	-484GB/s	370.3GB/s
Total Power Budget ("TDP")	250W	250W	300W	250W	190W
Power Connectors	1x 8-pin 1x 6-pin	?	1x 8-pin 1x 6-pin	1x 8-pin	
Release Date	12/07/2017	4Q16-1Q17	TBD	5/27/2016	
Release Price	\$3000	\$10000	-	\$700	MSRP: \$600 Now: \$500

The NVIDIA Titan V graphics card is not targeted at gaming, but rather scientific and machine/deep learning applications. That does not mean it's not good for gaming, but the performance numbers are not yet available. The Tesla V100 is the latest member of the Tesla V series. The Titan V is a derivative of the earlier enhanced GV100 GPU, part of the Tesla accelerator card series. The key differentiation is that the Titan V ships at \$3000, whereas the Tesla V100 was available as part of a \$10,000 developer kit. The Tesla V100 still offers greater memory capacity by 4GB – 16GB HBM2 versus 12GB HBM2 – and has a wider memory interface, but other core features remain matched or nearly matched. Core count, for one, is 5120 CUDA cores on each GPU, with 640 Tensor cores (used for Tensorflow deep/practice learning workloads) on each GPU.

mb – July 24, 2020



Who was the original Vulcan?

14

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 smithery. 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))



mb – July 24, 2020

Who is the Khronos Group?

16

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.



mb – July 24, 2020

Playing "Where's Waldo" with Khronos Membership

17



mb - July 24, 2020

Who's Been Specifically Working on Vulkan?

18



mb - July 24, 2020

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

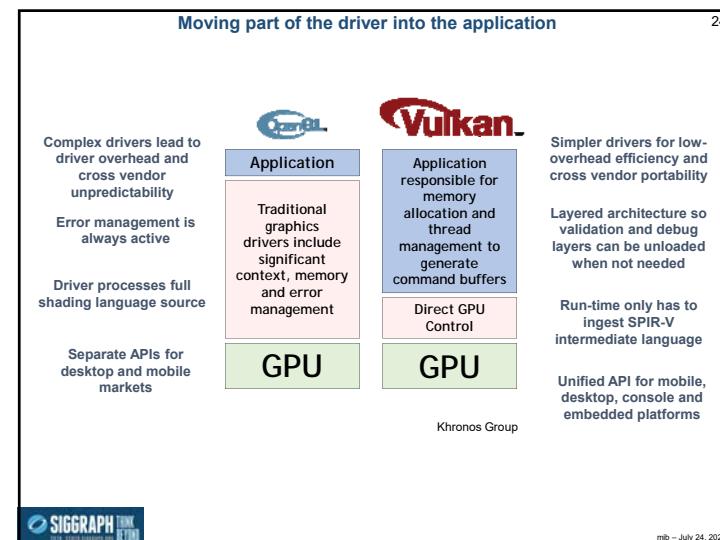
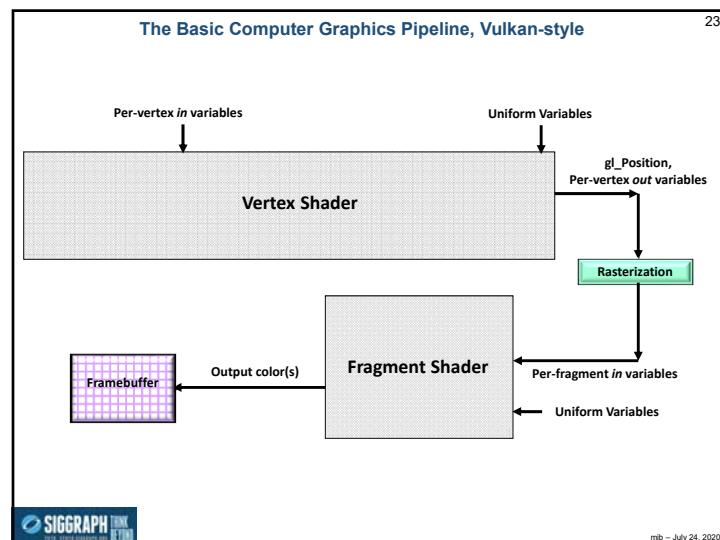
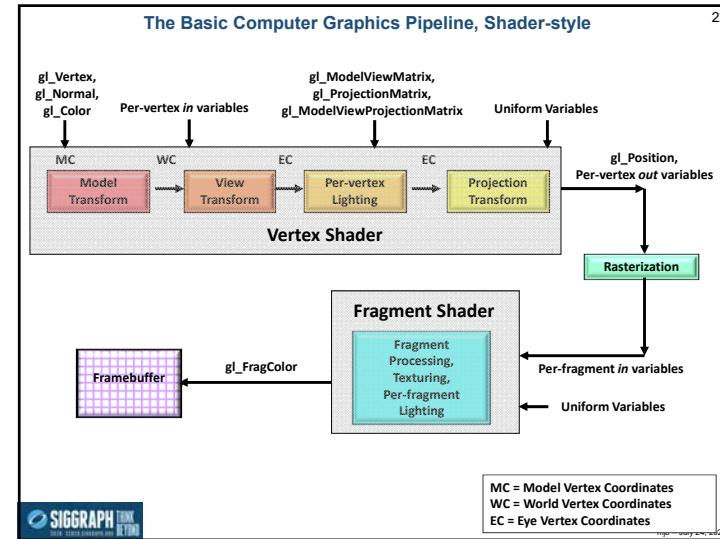
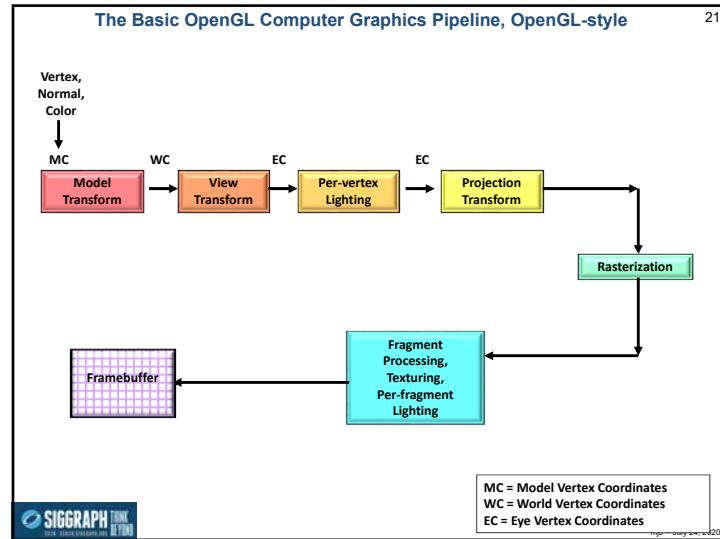
Vulkan Differences from OpenGL

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.



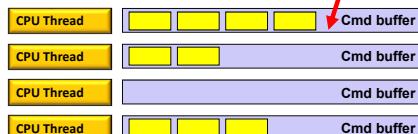
mb - July 24, 2020



Vulkan Highlights: Command Buffers

25

- 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



SIGGRAPH
THINK BEYOND

mjh - July 24, 2020

Vulkan Highlights: Pipeline State Objects

26

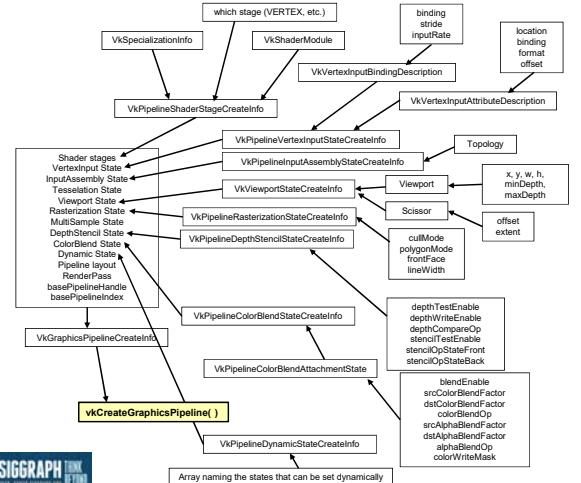
- In OpenGL, your “pipeline state” is the combination of whatever your current graphics attributes are: color, transformations, textures, shaders, etc.
 - Changing the state on-the-fly one item at-a-time is very expensive
 - Vulkan forces you to set all your state variables at once into a “pipeline state 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

SIGGRAPH
THINK BEYOND

min - July 24, 2022

Vulkan: Creating a Pipeline

27



SIGGRAPH
THINK BEYOND

mib - July 31, 2020

Querying the Number of Something

28

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

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

	How many total there are	Where to put them
result = vkEnumeratePhysicalDevices(Instance,	&count,	nullptr);
result = vkEnumeratePhysicalDevices(Instance,	&count,	physicalDevices);

SIGGRAPH

mlb - July 24, 2023

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

29

```

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

VK_RESULT result = vkCreateBuffer( LogicalDevice, IN &vbc, 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 );

```



mb - July 24, 2020

Vulkan Quick Reference Card – I Recommend you Print This!

30

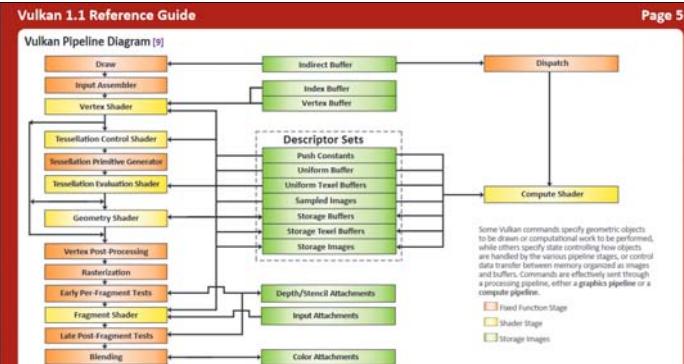


<https://www.khronos.org/files/vulkan11-reference-guide.pdf>

mb - July 24, 2020

Vulkan Quick Reference Card

31



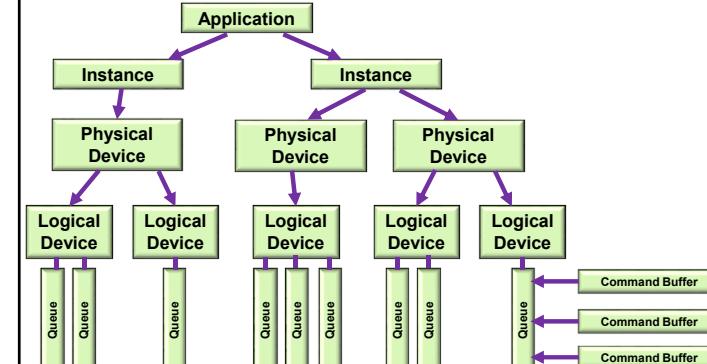
<https://www.khronos.org/files/vulkan11-reference-guide.pdf>



mb - July 24, 2020

Vulkan Highlights: Overall Block Diagram

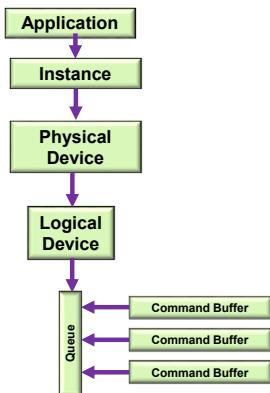
32



mb - July 24, 2020

Vulkan Highlights: a More Typical Block Diagram

33



mb - July 24, 2020

Steps in Creating Graphics using Vulkan

34

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- ...



mb - July 24, 2020

Vulkan GPU Memory

35

- 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.



mb - July 24, 2020

Vulkan Render Passes

36

- 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



mb - July 24, 2020

Vulkan Compute Shaders 37

- 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

 mb - July 24, 2020

Vulkan Synchronization 38

- Synchronization is the responsibility of the application
- Events can be set, polled, and waited for (much like OpenCL)
- 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

 mb - July 24, 2020

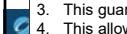
Vulkan Shaders 39

- 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:

- Software vendors don't need to ship their shader source
- 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

Your Sample2019.zip File Contains This 40

Name	Date modified	Type	Size
vs	3/4/2019 2:34 PM	File Folder	709 KB
Debug	3/4/2019 2:49 PM	File Folder	2 KB
glm	3/4/2019 2:34 PM	File Folder	240 KB
glm.glsl	3/4/2019 2:34 PM	File	1,631 KB
glm.0.8.3.a2	3/4/2019 2:34 PM	File	6 KB
oglplus-prj	3/4/2019 2:34 PM	File	1 KB
frog.h	3/10/2018 8:07 AM	SPV File	2 KB
glf.h	3/26/2017 10:48 AM	C/C++ Header	149 KB
glf.h.hdr	3/16/2019 5:06 AM	Object File Library	240 KB
glslangValidator	3/21/2017 5:24 PM	File	1,817 KB
glslangValidator.exe	3/15/2017 12:33 PM	Application	1,631 KB
glslangValidator-help	10/6/2017 2:31 PM	HELP File	6 KB
Makafila	3/31/2018 11:41 AM	File	1 KB
puppy.bmp	3/10/2018 8:13 AM	BMP File	3,072 KB
puppy.jpg	3/10/2018 8:13 AM	JPG File	443 KB
puppyD.bmp	3/1/2018 9:57 AM	BMP File	3,077 KB
puppyD.jpg	3/1/2018 9:58 AM	JPG File	455 KB
sample.cpp	3/4/2019 2:49 PM	C++ Source	158 KB
sample.shader.cpp	3/1/2018 12:46 PM	C++ Source	135 KB
Sample.jn	3/27/2017 9:46 AM	Microsoft Visual Studio Project File	2 KB
Sample.vcxproj	3/4/2019 2:37 PM	VC++ Project	7 KB
Sample.vcxproj.filters	3/26/2018 9:41 AM	VC++ Project Filter	1 KB
Sample.vcxproj.user	1/6/2018 11:28 AM	Per-User Project Data	1 KB
sample.pdf	1/6/2018 11:28 AM	Adobe Acrobat Document	34 KB
sample10.pdf	1/6/2018 11:28 AM	Adobe Acrobat Document	39 KB
sample10.ppt	1/6/2018 11:28 AM	Adobe Acrobat Document	54 KB
sample-compo.comp	2/14/2018 12:25 PM	COMP File	2 KB
sample-compo.spv	2/14/2018 12:25 PM	SPV File	4 KB
sample-frag.frag	2/18/2018 10:52 AM	FRAG File	2 KB

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

 mb - July 24, 2020

41

Vulkan.

The Vulkan Sample Code Included with These Notes

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
CREATE. DESIGN. INNOVATE. BEYOND

mjb – July 24, 2020

42

Sample Program Output

SIGGRAPH THINK
CREATE. DESIGN. INNOVATE. BEYOND

mjb – July 24, 2020

43

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
'i', 'I':	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)

SIGGRAPH THINK
CREATE. DESIGN. INNOVATE. BEYOND

mjb – July 24, 2020

44

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.

SIGGRAPH THINK
CREATE. DESIGN. INNOVATE. BEYOND

mjb – July 24, 2020

Caveats on the Sample Code, II

45

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.



mb - July 24, 2020

Main Program

46

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

    errno_t err = fopen_s( &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");

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

mb - July 24, 2020

InitGraphics(), I

47

```
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();
}
```



mb - July 24, 2020

InitGraphics(), II

48

```
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 );
}
```

mb - July 24, 2020

A Colored Cube

```
static GLfloat CubeColors[ ][3] =
{
    { 0., 0., 0. },
    { 1., 0., 0. },
    { 0., 1., 0. },
    { 1., 1., 0. },
    { 0., 0., 1. },
    { 1., 0., 1. },
    { 0., 1., 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 }
};
```

Hello GLSL! CubeVertices.h

```
#include "SampleVertexData.h"
```

mb - July 24, 2020

**SIGGRAPH THINK
Create. Share. Inspire. REACH**

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. },
        { 0., 1., 0. },
        { 0., 1. }
    }
};
```

**SIGGRAPH THINK
Create. Share. Inspire. REACH**

mb - 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. }
    },
    ...
};
```

mb - July 24, 2020

**SIGGRAPH THINK
Create. Share. Inspire. REACH**

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.

mb - July 24, 2020

**SIGGRAPH THINK
Create. Share. Inspire. REACH**

Vulkan Software Philosophy

53

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

Vulkan takes a non-C++ object-oriented approach in that those typedefed 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 );
```



mb - July 24, 2020

Vulkan Conventions

54

VkXxx is a typedef, probably a struct

vkYyy() is a function call

VK_ZZZ is a constant

My Conventions

"Init" in a function 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 InitxxYYY() functions in numerical order. After that comes the helper functions

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

"Fill" in a function 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.



mb - July 24, 2020

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

55

```
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):

How many total there are	Where to put them
result = vkEnumeratePhysicalDevices(Instance, &count, nullptr);	
result = vkEnumeratePhysicalDevices(Instance, &count, &physicalDevices[0]);	



mb - July 24, 2020

Your Sample2019.zip File Contains This

56

Linux shader compiler

Windows shader compiler

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

Name	Date modified	Type	Size
vs	9/4/2019 2:34 PM	File folder	
Debug	9/4/2019 2:49 PM	File folder	
glm	9/4/2019 2:34 PM	File folder	
glm-0.9.8.5	9/4/2019 2:34 PM	File folder	
glm-0.9.8.5-a2	9/4/2019 2:34 PM	File folder	
EMROSS.pch	9/4/2019 2:34 PM	File	2 KB
glslangValidator	1/10/2019 8:07 AM	MSI File	2 KB
glslang	1/29/2017 10:48 AM	C/C++ Header	540 KB
glslang.3.h	3/16/2016 5:06 AM	Object File Library	240 KB
glslang.3.h.pch	1/21/2017 5:31 PM	File	1,817 KB
glslangValidation	5/15/2017 12:31 PM	Application	1,623 KB
glslangValidation.exe	10/6/2017 2:31 PM	HELP File	6 KB
glslangValidationHelp	1/23/2018 11:47 AM	File	1 KB
Makefile	1/10/2018 8:13 AM	BMP File	3,073 KB
puppy.bmp	1/10/2018 8:13 AM	JPG File	443 KB
puppy0.bmp	1/1/2018 9:57 AM	BMP File	3,073 KB
puppy0.jpg	1/1/2018 9:58 AM	JPG File	453 KB
sample.cpp	9/4/2019 2:49 PM	C++ Source	158 KB
sample.frag	3/1/2018 12:46 PM	C++ Source	155 KB
sample.vert	3/27/2017 9:40 AM	Microsoft Visual Studio...	2 KB
Sample.sln	9/4/2019 2:37 PM	VC++ Project	7 KB
Sample.vcxproj	5/27/2018 11:26 AM	VC++ Project File	1 KB
Sample.vcxproj.filters	6/29/2018 8:49 PM	Per-User Project File	1 KB
Sample.vcxproj.user	1/8/2018 11:26 AM	Adobe Acrobat D...	54 KB
sample0.pdf	1/8/2018 11:26 AM	Adobe Acrobat D...	89 KB
sample1.pdf	1/9/2018 11:26 AM	Adobe Acrobat D...	54 KB
sample10.pdf	3/14/2018 12:23 PM	COMP File	2 KB
sample10.psv	3/14/2018 12:23 PM	SPV File	4 KB
sample.frag.frag	2/18/2018 10:52 AM	FRAG File	2 KB

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



mb - July 24, 2020

Reporting Error Results, I 57

```

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" }
};

```

July 24, 2020

Reporting Error Results, II 58

```

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

```

mb - July 24, 2020

Extras in the Code 59

```

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

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

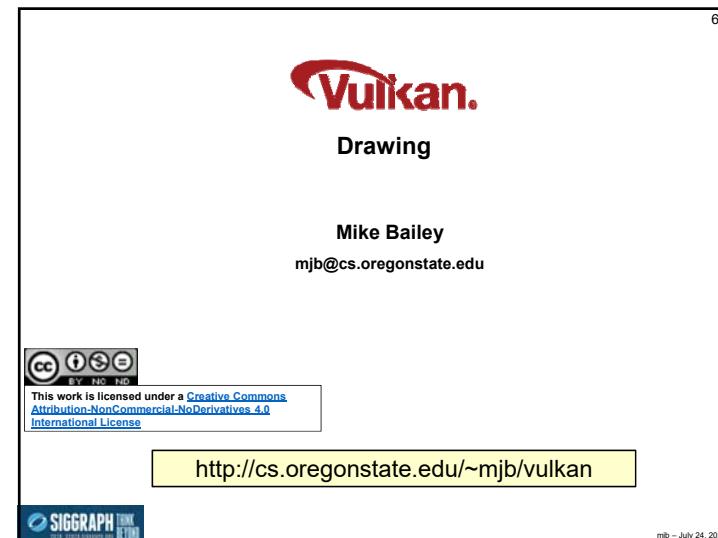
bool Paused;
bool Verbose;

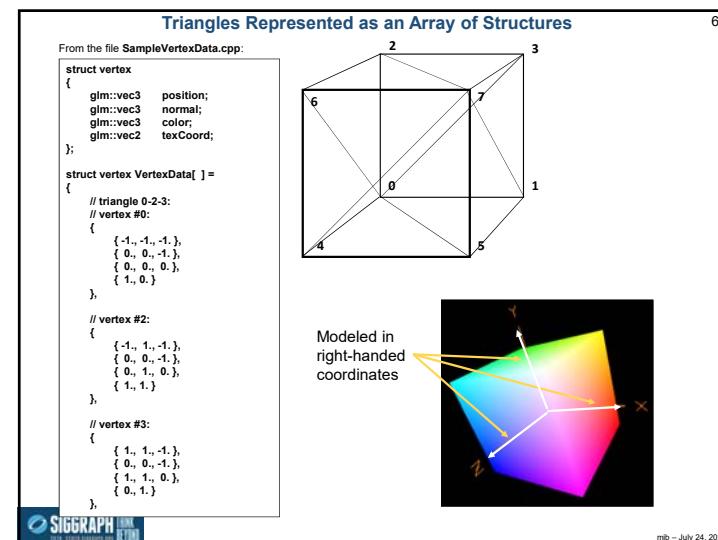
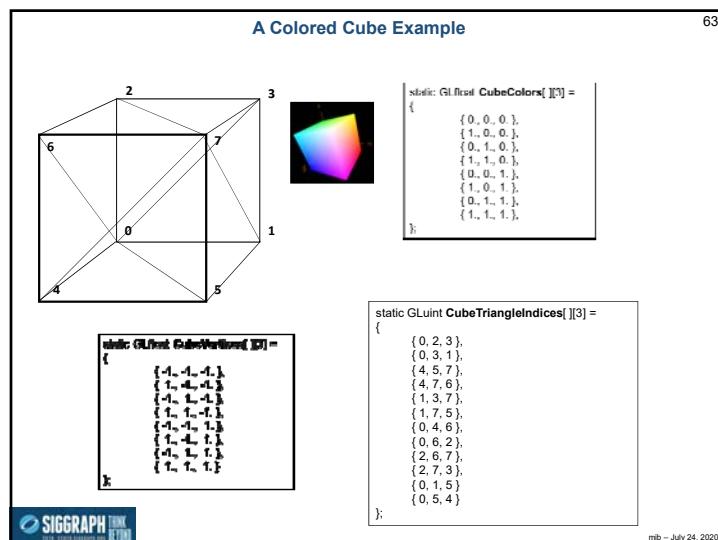
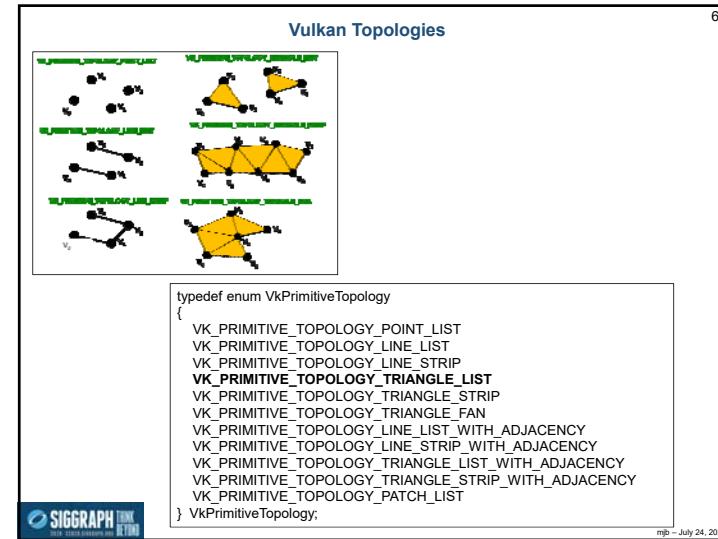
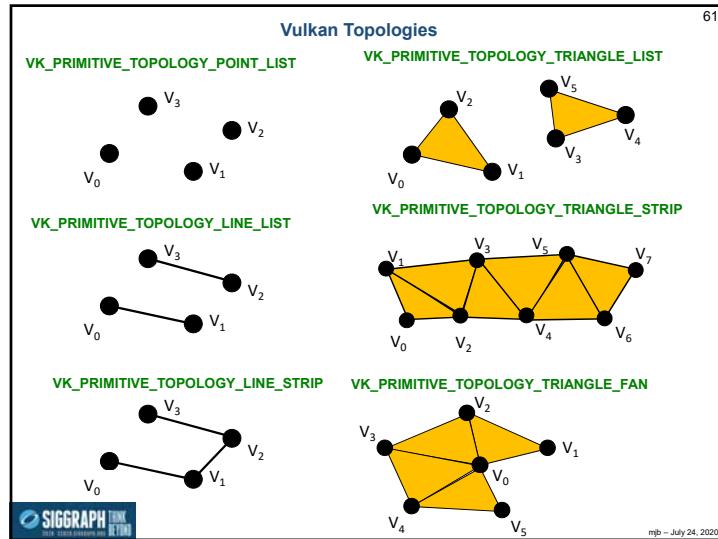
#define DEBUGFILE "VulkanDebug.txt"
errno_t err = fopen_s( &FpDebug, DEBUGFILE, "w" );

const int32_t OFFSET_ZERO = 0;

```

mb - July 24, 2020





Non-indexed Buffer Drawing 65

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., -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. } },
};

```

Stream of Vertices

```

graph TD
    A[Vertex 7] --> B[Vertex 5]
    B --> C[Vertex 4]
    C --> D[Vertex 1]
    D --> E[Vertex 3]
    E --> F[Vertex 0]
    F --> G[Vertex 3]
    G --> H[Vertex 2]
    H --> I[Vertex 0]
    I --> J[Triangles]
    J --> K[Draw]

```

mb – July 24, 2020

Filling the Vertex Buffer 66

```

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

MyBuffer MyVertexDataBuffer;

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

VkResult
Init05MyVertexDataBuffer( IN VkDeviceSize size, OUT MyBuffer * pMyBuffer )
{
    VkResult result;
    result = Init05DataBuffer( size, VK_BUFFER_USAGE_VERTEX_BUFFER_BIT, pMyBuffer );
    return result;
}

```

SIGGRAPH THINK
Create. Connect. Inspire. REINVENT.

mb – July 24, 2020

A Preview of What `Init05DataBuffer` Does 67

```

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 = size;
    vbc.usage = usage;
    vbc.sharingMode = VK_SHARING_MODE_EXCLUSIVE;
    vbc.queueFamilyIndexCount = 0;
    vbc.pQueueFamilyIndices = (const b16* )nullptr;
    result = vkCreateBuffer( LogicalDevice, IN &vbc, PALLOCATOR, OUT &pMyBuffer->buffer );

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

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

    Vkmemory
    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;
}

```

SIGGRAPH THINK
Create. Connect. Inspire. REINVENT.

mb – July 24, 2020

Telling the Pipeline about its Input 68

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

```

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

```

SIGGRAPH THINK
Create. Connect. Inspire. REINVENT.

mb – July 24, 2020

Telling the Pipeline about its Input 69

```

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

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

```

mb - July 24, 2020



Telling the Pipeline about its Input

```

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

```



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

Telling the Pipeline about its Input 70

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 = vviad;
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;;

```

mb - July 24, 2020



Telling the Pipeline about its Input 71

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;
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 = vpssci;
vgpci.pVertexInputState = &vpvisci;
vgpci.pInputAssemblyState = &vpiasci;
vgpci.pTessellationState = (VkPipelineTessellationStateCreateInfo *)nullptr; // &vpptsci
vgpci.pViewportState = &vpvsci;
vgpci.pRasterizationState = &vprsc;
vgpci.pMultisampleState = &vpmsc;
vgpci.pDepthStencilState = &vpdssci;
vgpci.pColorBlendState = &vpcbsci;
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



Telling the Pipeline about its Input

```

vertexCount = sizeof( VertexData ) / sizeof( VertexData[0] );

```



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

Telling the Command Buffer what Vertices to Draw 72

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] = MyVertexBuffer.buffer;
vkCmdBindVertexBuffers( CommandBuffers[nextImageIndex], 0, 1, vertexDataBuffers, offsets );

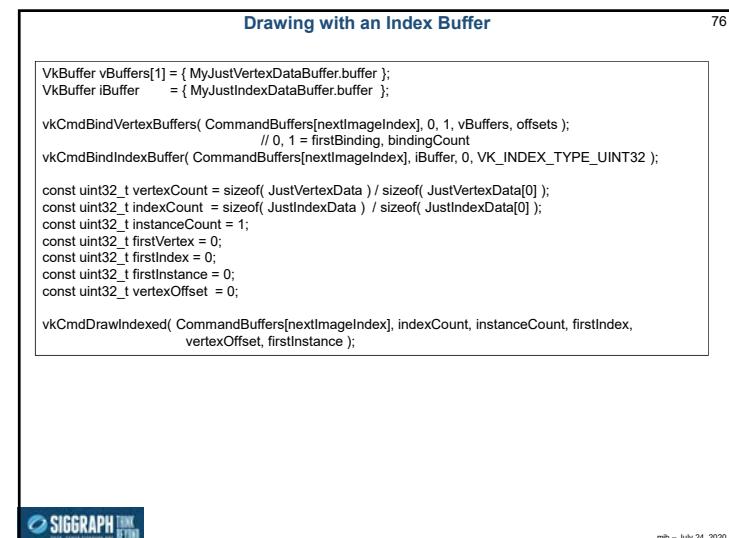
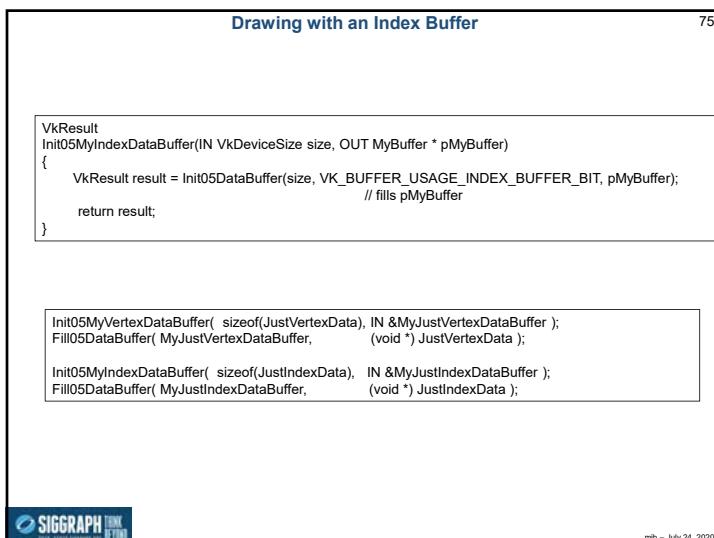
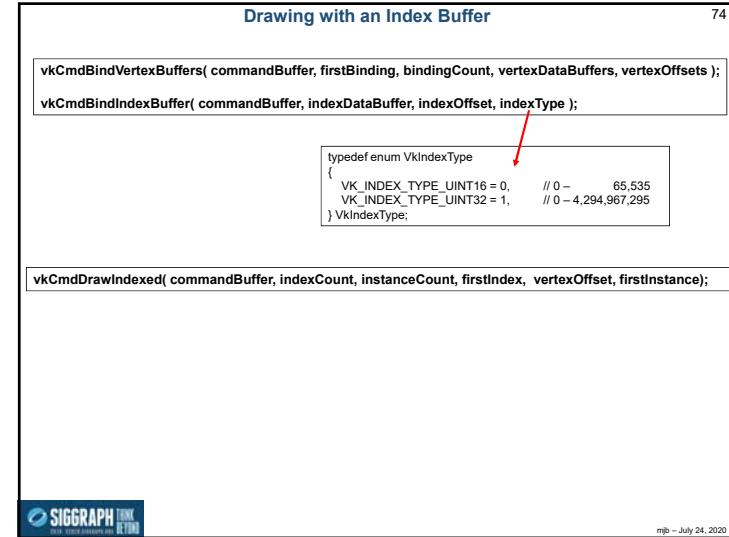
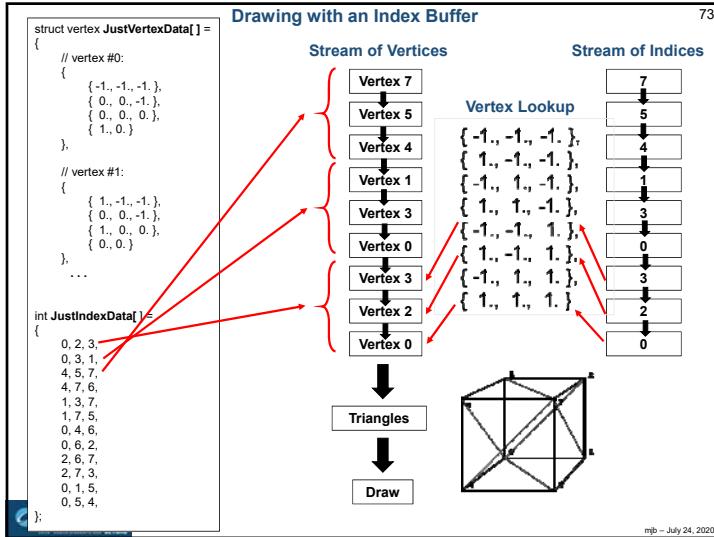
const uint32_t vertexCount = sizeof( VertexData ) / sizeof( VertexData[0] );
const uint32_t instanceCount = 1;
const uint32_t firstVertex = 0;
const uint32_t firstInstance = 0;

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

```

mb - July 24, 2020





Indirect Drawing (not to be confused with Indexed)

77

```
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 );
```



mb - July 24, 2020

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

78

```
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:

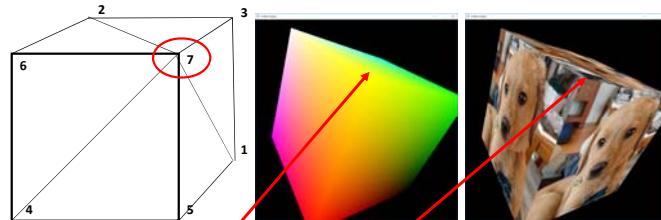
```
vkCmdDrawIndexed( commandBuffer, indexCount, instanceCount, firstIndex, vertexOffset, firstInstance );
```



mb - July 24, 2020

Sometimes the Same Point Needs Multiple Attributes

79



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.

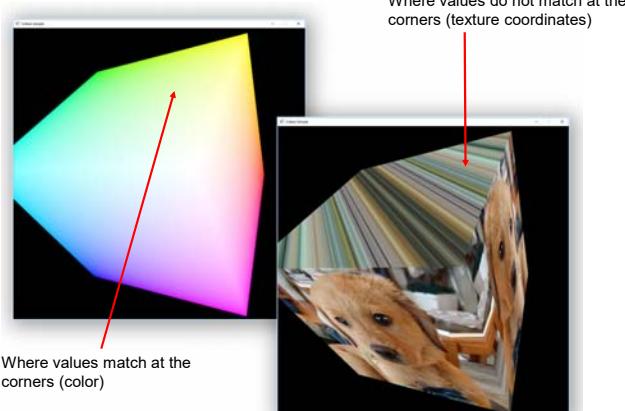
Thus, when using indexed buffer drawing, you need to create a new vertex struct if any of {position, normal, color, texCoords} changes from what was previously stored at those coordinates.



mb - July 24, 2020

Sometimes the Same Point Needs Multiple Attributes

80



Where values match at the corners (color)

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

mb - July 24, 2020

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

81



```
v 1.710541 1.283360 -0.040860
v 1.714593 1.279043 -0.041268
v 1.706114 1.279109 -0.040795
v 1.719083 1.277235 -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 1.286594 -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 1.268344 -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
```

V / T / N

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

mjb – July 24, 2020

82

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://creativecommons.org/licenses/by-nc-nd/4.0/)

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

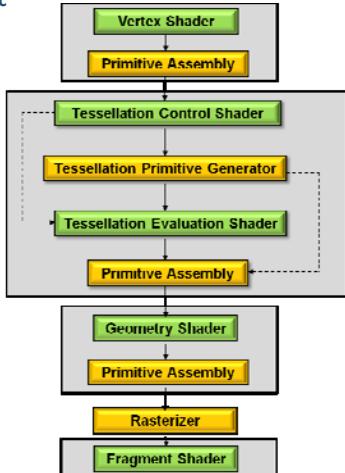


mjb – July 24, 2020

The Shaders' View of the Basic Computer Graphics Pipeline

83

- 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



= Fixed Function
= Programmable



mjb – July 24, 2020

84

Vulkan Shader Stages

Shader stages

```
typedef enum VkPipelineStageFlagBits {
    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,
} VkPipelineStageFlagBits;
```



mjb – July 24, 2020

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
```

- Both are 0-based

OpenGL uses:

```
gl_VertexID  
gl_InstanceID
```

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



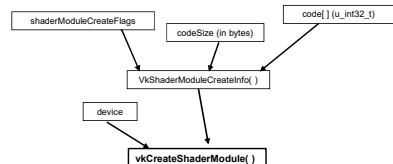
mb - July 24, 2020

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

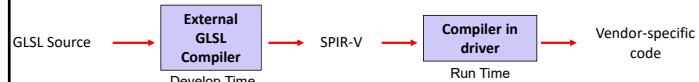


mb - July 24, 2020

Vulkan Shader Compiling

88

- You half-precompile 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

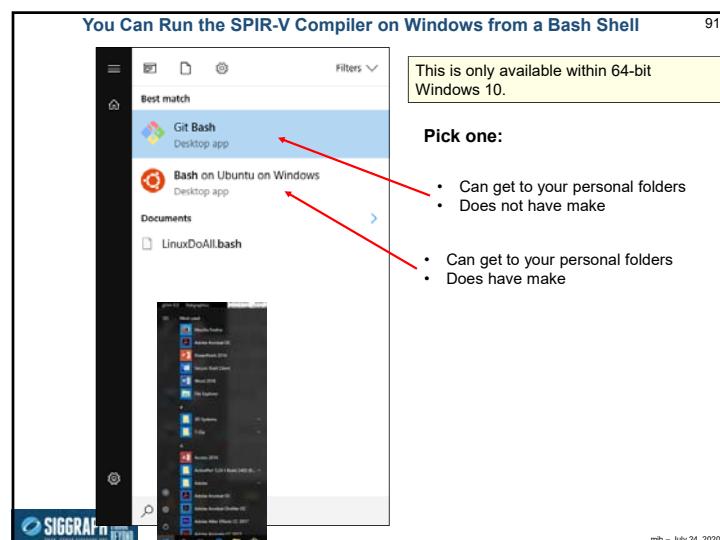
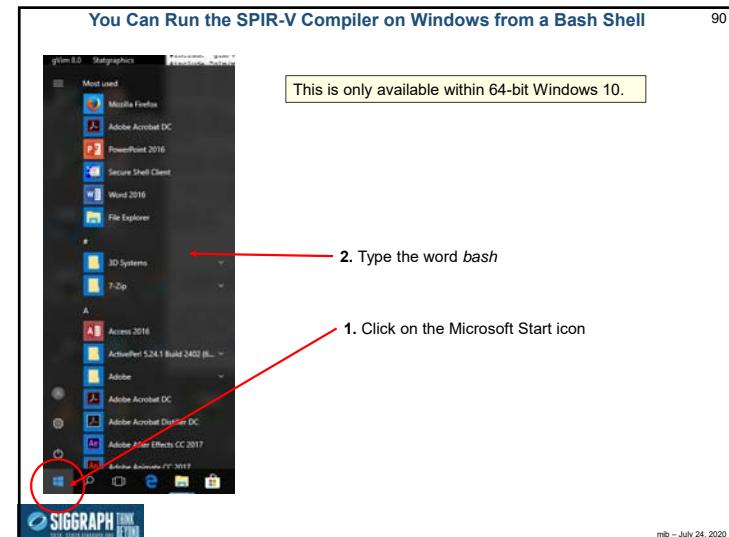
(Can be overridden by the -S option)

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

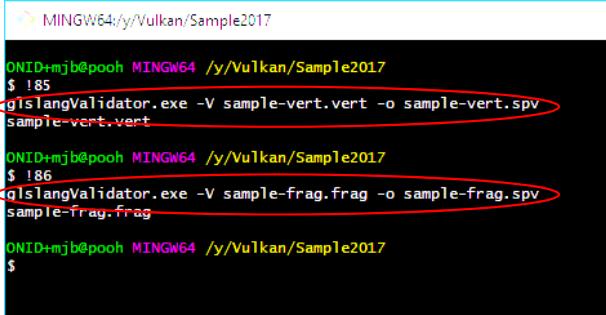
Windows: glslangValidator.exe
Linux: glslangValidator



mb - July 24, 2020



Running glslangValidator.exe 92



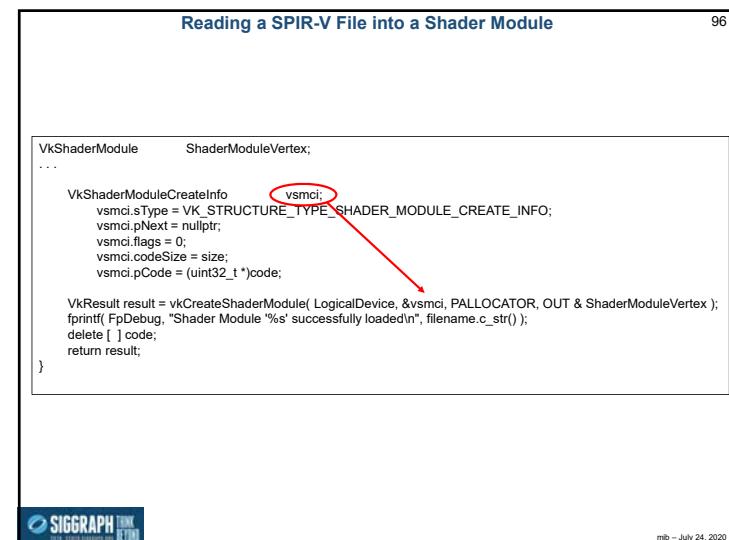
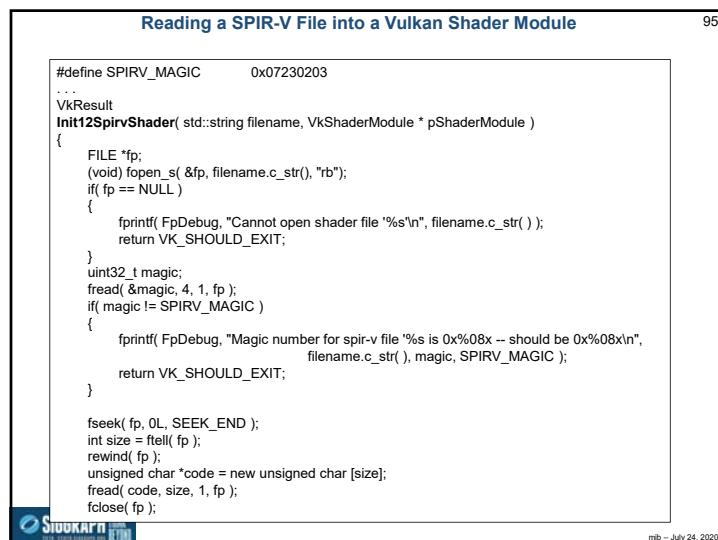
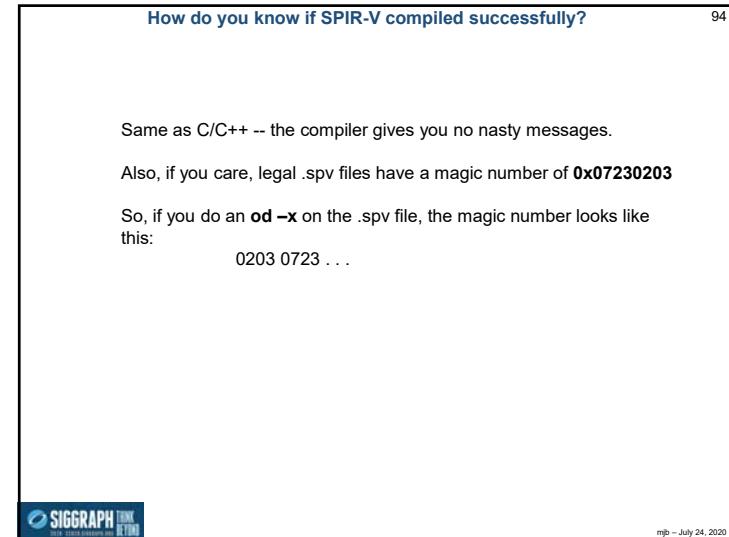
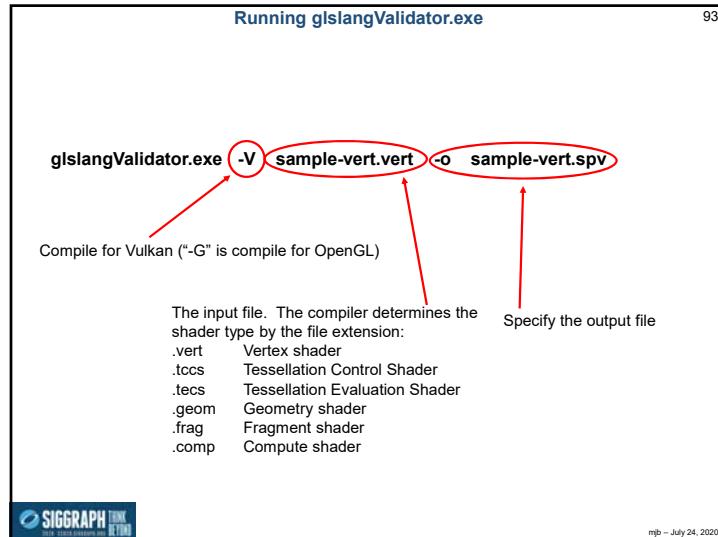
```
MINGW64:/y/Vulkan/Sample2017
$ 185
glslangValidator.exe -V sample-vert.vert -o sample-vert.spv
sample-vert.vert

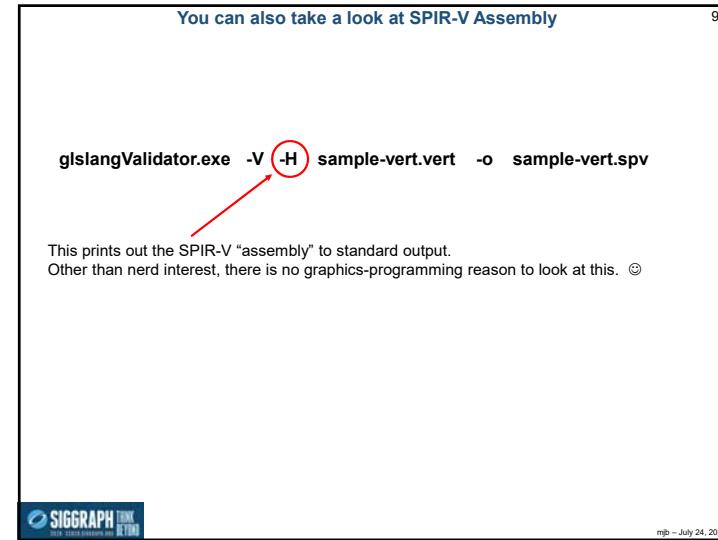
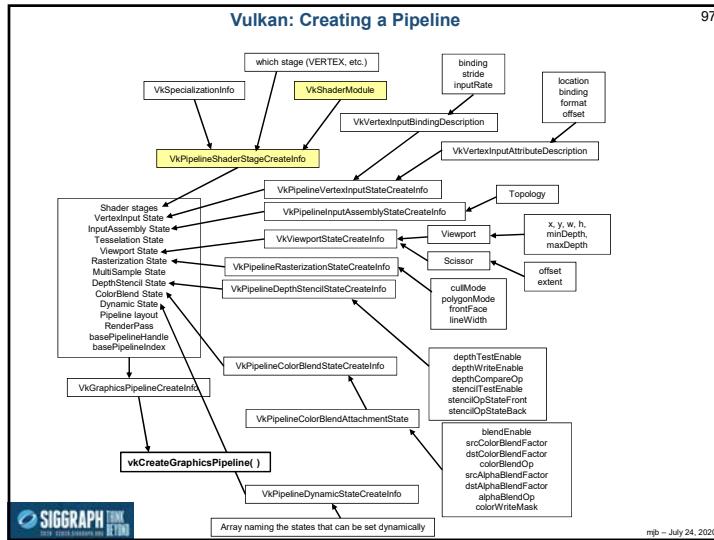
ONID+mjb@pooh MINGW64 /y/Vulkan/Sample2017
$ 186
glslangValidator.exe -V sample-frag.frag -o sample-frag.spv
sample-frag.frag

ONID+mjb@pooh MINGW64 /y/Vulkan/Sample2017
$
```



mb - July 24, 2020





For example, if this is your Shader Source 99

```
#version 400
#extension GL_ARB_separate_shader_objects : enable
#extension GL_ARB_shading_language_420pack : enable
layout( std140, set = 0, binding = 0 ) uniform matBuf
{
    mat4 uModelMatrix;
    mat4 uViewMatrix;
    mat4 uProjectionMatrix;
    mat3 uNormalMatrix;
} Matrices;

// 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 vec2 aTexCoord;

layout( location = 0 ) out vec3 vNormal;
layout( location = 1 ) out vec3 vColor;
layout( location = 2 ) out vec2 vTexCoord;

void main()
{
    mat4 PVM = Matrices.uProjectionMatrix * Matrices.uViewMatrix * Matrices.uModelMatrix;
    gl_Position = PVM * vec4( aVertex, 1. );

    vNormal = Matrices.uNormalMatrix * aNormal;
    vColor = aColor;
    vTexCoord = aTexCoord;
}
```

SIGGRAPH 2018: State-of-the-art in Vulkan

mb - July 24, 2020

This is the SPIR-V Assembly, Part I 100

```
#version 400
#extension GL_ARB_separate_shader_objects : enable
#extension GL_ARB_shading_language_420pack : enable
layout( std140, set = 0, binding = 0 ) uniform matBuf
{
    mat4 uModelMatrix;
    mat4 uViewMatrix;
    mat4 uProjectionMatrix;
    mat3 uNormalMatrix;
} Matrices;

// 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 vec2 aTexCoord;

layout( location = 0 ) out vec3 vNormal;
layout( location = 1 ) out vec3 vColor;
layout( location = 2 ) out vec2 vTexCoord;

void main()
{
    mat4 PVM = Matrices.uProjectionMatrix * Matrices.uViewMatrix * Matrices.uModelMatrix;
    gl_Position = PVM * vec4( aVertex, 1. );

    vNormal = Matrices.uNormalMatrix * aNormal;
    vColor = aColor;
    vTexCoord = aTexCoord;
}
```

```
Capability Shader
ExtImport #include <GLSL.std.450>
MemoryModel Logical GLSL450
EntryPoint Vertex 4 "main" 34 37 48 53 56 57 61 63
Source GLSL 400
SourceExtension "GL_ARB_separate_shader_objects"
SourceExtension "GL_ARB_shading_language_420pack"
Name 4 "main"
Name 10 "PVM"
Name 11 "uViewMatrix"
MemberName 13(matBuf) 0 "uModelMatrix"
MemberName 13(matBuf) 1 "uViewMatrix"
MemberName 13(matBuf) 2 "uProjectionMatrix"
MemberName 13(matBuf) 3 "uNormalMatrix"
Name 15 "Matrices"
Name 32 "g_PerVertex"
MemberName 32(g_PerVertex) 0 "gl_Position"
MemberName 32(g_PerVertex) 1 "gl_PointSize"
MemberName 32(g_PerVertex) 2 "gl_ClipDistance"
Name 34 "Light"
Name 37 "aVertex"
Name 48 "aNormal"
Name 55 "aColor"
Name 56 "vNormal"
Name 57 "vColor"
Name 61 "vTexCoord"
Name 62 "vColor"
Name 63 "vNormal"
Name 65 "lightBuf"
MemberName 65(lightBuf) 0 "uLightPos"
Name 67 "Light"
MemberDecorate 13(matBuf) 0 MatrixStride 16
MemberDecorate 13(matBuf) 0 Offset 0
MemberDecorate 13(matBuf) 0 MatSize 16
MemberDecorate 13(matBuf) 1 ColMajor
MemberDecorate 13(matBuf) 1 RowMajor
MemberDecorate 13(matBuf) 1 MatStride 16
MemberDecorate 13(matBuf) 2 OffSet 16
MemberDecorate 13(matBuf) 2 ColMajor
MemberDecorate 13(matBuf) 2 RowMajor
MemberDecorate 13(matBuf) 2 MatStride 16
MemberDecorate 13(matBuf) 3 OffSet 32
MemberDecorate 13(matBuf) 3 ColMajor
MemberDecorate 13(matBuf) 3 RowMajor
MemberDecorate 13(matBuf) 3 MatStride 16
Decorate 13(matBuf) Block
Decorate 15(Matrices) DescriptorSet 0
```

SIGGRAPH 2018: State-of-the-art in Vulkan

mb - July 24, 2020

This is the SPIR-V Assembly, Part II

10

```

    new(Vector4D) ARGB Caparis Shader object white  

    new(Vector4D) ARGB Caparis Shader object black  

    layout std140, binding = 0 uniform mat4  

        mat4 uMVPMatrix;  

        mat4 uNormalMatrix;  

        mat4 uWorldMatrix;  

    } Matrices  

void draw (Program program block in uniform mat4 lightDef  

    [ ] vec4 lightPos;  

    Light  

    uniform vec3 uLightColor;  

    uniform vec3 uLightIntensity;  

    uniform vec3 uLightPosition;  

    uniform vec3 uLightNormal;  

    uniform vec3 uEyePosition;  

    uniform vec3 uEyeNormal;  

    uniform vec3 uEyeViewDir;  

    uniform vec3 uEyeViewNormal;  

void  

main ()  

{  

    cout << "Molten - " molten version + " Molten - " molten version + " Molten - " molten version  

    + " Molten - " molten version + " Molten - " molten version  

    + " Molten - " molten version + " Molten - " molten version  

    + " Molten - " molten version + " Molten - " molten version  

}

```



January 24, 20

This is the SPIR-V Assembly, Part III

10

```

function G( ATC Attribute shader block, uniform mat4 modelViewProj, uniform vec3 color, uniform vec2 uvCoord, uniform vec2 binding ) {
    mat4 modelView;
    mat4 view;
    mat4 model;
    mat4 projection;
    mat3 normalMatrix;
    mat2 texCoordMatrix;
}

Matrices = {
    modelViewProj: modelViewProj,
    modelView: modelView,
    view: view,
    model: model,
    projection: projection,
    normalMatrix: normalMatrix,
    texCoordMatrix: texCoordMatrix
};

if opaque model bc is a uniform block, then
    modelViewProj.set( modelViewProj, modelView.bind( binding ) );
    modelView.set( modelView, model.bind( binding ) );
    model.set( model, model.bind( binding ) );
    projection.set( projection, projection.bind( binding ) );
    normalMatrix.set( normalMatrix, normalMatrix.bind( binding ) );
    texCoordMatrix.set( texCoordMatrix, texCoordMatrix.bind( binding ) );
}

Light = {
    modelViewProjectionMatrix: modelViewProj,
    modelViewMatrix: modelView,
    viewMatrix: view,
    modelMatrix: model,
    projectionMatrix: projection
};

Normal = {
    modelViewProjectionMatrix: modelViewProj,
    modelViewMatrix: modelView,
    viewMatrix: view,
    modelMatrix: model,
    projectionMatrix: projection
};

Lighting = {
    modelViewProjectionMatrix: modelViewProj,
    modelViewMatrix: modelView,
    viewMatrix: view,
    modelMatrix: model,
    projectionMatrix: projection
};

World = {
    modelViewProjectionMatrix: modelViewProj,
    modelViewMatrix: modelView,
    viewMatrix: view,
    modelMatrix: model,
    projectionMatrix: projection,
    normalMatrix: normalMatrix
};

WorldNormal = {
    modelViewProjectionMatrix: modelViewProj,
    modelViewMatrix: modelView,
    viewMatrix: view,
    modelMatrix: model,
    projectionMatrix: projection,
    normalMatrix: normalMatrix
};

WorldLighting = {
    modelViewProjectionMatrix: modelViewProj,
    modelViewMatrix: modelView,
    viewMatrix: view,
    modelMatrix: model,
    projectionMatrix: projection
};

```



mjb - July 24, 2014

SPIR-V: Printing the Configuration

10

```
MaxLights 32
MaxCapPlanes 6
MaxTextureUnits 32
MaxTextureCoords 32
MaxTextureSize 64
MaxVertexUniformComponents 4096
MaxVaryingFloats 4
MaxVaryingFloats 32
MaxCombinedImageUnits 32
MaxCombinedTextureImageUnits 80
MaxTextureImageUnits 32
MaxProgramUniformComponents 4096
MaxDrawPrimitives 8
MaxVertexUniformVectors 128
MaxVaryingVectors 8
MaxVaryingVectorPatches 16
MaxVertexOutputVectors 16
MaxFragmentInputVectors 15
MaxProgramParameters 128
MaxProgramTextOffset 7
MaxClipDistances 8
MaxComputeGroupCount 65535
MaxComputeWorkGroupCount 65535
MaxComputeWorkGroupSize 65535
MaxComputeWorkGroupSize 1024
MaxComputeWorkGroupSize 64
MaxComputeUniformComponents 1024
MaxComputeSharedMemoryBytes 16
MaxComputeImageUniforms 16
MaxComputeAtomicCounters 4
MaxComputeAtomicOperations 4
MaxComputeSharedMemoryAccess 16
MaxVaryingComponents 60
MaxVertexOutputComponents 64
MaxFragmentOutputComponents 64
MaxProgramOutputComponents 128
MaxFragmentInputComponents 128
MaxCombinedImageUnitsAndFragmentOutput 8
MaxCombinedSharedImageUnitsAndFragmentOutput 8
MaxVertexInDescriptors 0
MaxTextureControlImageUniforms 0
MaxTextureEvaluationImageUniforms 0
MaxTextureImageUniforms 0
MaxFragmentImageUniform 81
```

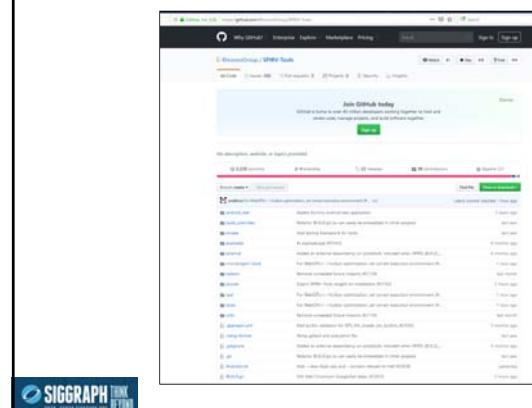


mkb - July 24, 2020

SPIR-V: More Information

10

SPIR-V Tools:



mlb - July 24, 2019

A Google-Wrapped Version of glslangValidator 105

The shadrc project from Google (<https://github.com/google/shadrc>) provides a glslangValidator wrapper program called **glslc** that has a much improved command-line interface. You use, basically, the same way:

```
glslc.exe --target-env=vulkan sample.vert.vert -o sample.vert.spv
```

There are several really nice features. The two I really like are:

1. You can #include files into your shader source
2. You can "#define" definitions on the command line like this:

```
glslc.exe --target-env=vulkan -DNUMPONTS=4 sample.vert.vert -o sample.vert.spv
```

glslc is included in your Sample .zip file

 SIGGRAPH THINK
GEAR UP FOR THE FUTURE

mjb – July 24, 2020

106

Vulkan.

Data Buffers

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
GEAR UP FOR THE FUTURE

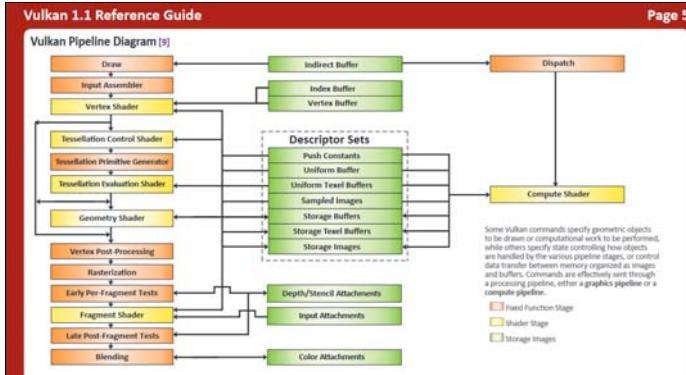
mjb – July 24, 2020

From the Quick Reference Card 107

Vulkan 1.1 Reference Guide

Page 5

Vulkan Pipeline Diagram [9]



Some Vulkan commands specify geometric objects to be drawn or computed. Work to be performed, including state control, new objects, are handled by the various pipeline stages, or control data to be used by the stages. Some commands are sent through buffers. Commands are effectively sent through a processing pipeline, either a graphics pipeline or a compute pipeline.

- Fixed Function Stage
- Shader Stage
- Storage Images

 SIGGRAPH THINK
GEAR UP FOR THE FUTURE

mjb – July 24, 2020

108

Terminology Issues

A Vulkan **Data Buffer** is just a group of contiguous bytes in GPU memory. They have no inherent meaning. The data that is stored there is whatever you want it to be. (This is sometimes called a "Binary Large Object", or "BLOB".)

It is up to you to be sure that the writer and the reader of the Data Buffer are interpreting the bytes in the same way!

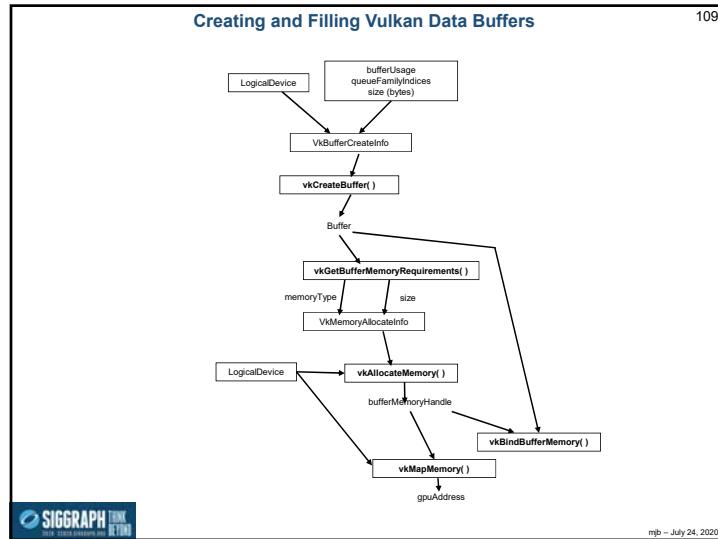
Vulkan calls these things "Buffers". But, Vulkan calls other things "Buffers", too, such as Texture Buffers and Command Buffers. So, I sometimes have taken to calling these things "Data Buffers" and have even gone to far as to override some of Vulkan's own terminology:

```
typedef VkBuffer VkDataBuffer;
```

This is probably a bad idea in the long run.

 SIGGRAPH THINK
GEAR UP FOR THE FUTURE

mjb – July 24, 2020



Creating a Vulkan Data Buffer 110

```

VkBuffer Buffer;
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'd 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 THINK
CREATE. DESIGN. INNOVATE. BEYOND.

mb - July 24, 2020

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

```

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

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

...
VkDeviceMemory
result = vkAllocateMemory( LogicalDevice, IN &vmal, 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 THINK
CREATE. DESIGN. INNOVATE. BEYOND.

mb - July 24, 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 THINK
CREATE. DESIGN. INNOVATE. BEYOND.

mb - July 24, 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.memoryType[ i ];
        if( ( vmt.propertyFlags & VK_MEMORY_PROPERTY_DEVICE_LOCAL_BIT ) != 0 )
        {
            return i;
        }
    }
    return -1;
}
```



mb - July 24, 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



mb - July 24, 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          vbsci;
...
VmaAllocationCreateInfo      vaci;
vac.i.physicalDevice = PhysicalDevice;
vac.i.device = LogicalDevice;
vac.i.usage = VMA_MEMORY_USAGE_GPU_ONLY;
```

```
VmaAllocator          var;
vmaCreateAllocator( IN &vac, OUT &var );
```

```
...
VkBuffer          Buffer;
VmaAllocation      van;
vmaCreateBuffer( IN var, IN &vbsci, IN &vac, 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 )
{
    ...
    vbsci.size = pMyBuffer->size = size;
    ...
    result = vkCreateBuffer( LogicalDevice, IN &vbsci, 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 )
{
    ...
    vmaAllocation = pMyBuffer->vma;
    ...
    result = vmaAllocBuffer( LogicalDevice, IN &size, PALLOCATOR, OUT &pMyBuffer->buffer );
    ...
    pMyBuffer->vdm = vdm;
}
```

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

struct matBuf

```
glm::vec3 eye(0.0,EYEDIST);
glm::vec3 look(0.0,0.0);
glm::vec3 up(0.0,1.0);

Matrices.uModelMatrix = glm::mat4( ); // identity
Matrices.uViewMatrix = glm::lookAt( eye, look, up );
Matrices.uProjectionMatrix = glm::perspective( 1.3f, (double)width/(double)height, 0.1, 1000.0 );
Matrices.uProjectionMatrix[1][1] *= -1;
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( std430, set = 0, binding = 0 ) uniform matBuf
{
    mat4 uModelMatrix;
    mat4 uViewMatrix;
    mat4 uProjectionMatrix;
    mat3 uNormalMatrix;
} Matrices;
```



mb - July 24, 2020

Filling the Data Buffer

122

```
typedef struct MyBuffer
{
    VKBufferHeader header;
    VKDeviceMemory vdm;
    VKAllocation allocation;
} MyBuffer;
```

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

```
glm::vec3 eye(0.0,EYEDIST);
glm::vec3 look(0.0,0.0);
glm::vec3 up(0.0,1.0.);
```

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));

mb - July 24, 2020

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 = 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
    vkmr;
    vkGetBufferMemoryRequirements( LogicalDevice, IN pMyBuffer->buffer, OUT &vkmr ); // fills vkmr

    VKMemoryAllocateInfo
    vmai;
    vmai.sType = VK_STRUCTURE_TYPE_MEMORY_ALLOCATE_INFO;
    vmai.pNext = nullptr;
    vmai.allocationSize = vkmr.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;
}
```



mb - July 24, 2020

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
// 0 and 0 are offset and flags
memcpy( pGpuMemory, data, (size_t)myBuffer.size );
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.



mb - July 24, 2020

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 );
    // 0 and 0 are offset and flags

    memcpy( pGpuMemory, data, (size_t)myBuffer.size );
    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.



mjb - July 24, 2020



GLFW

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>



mjb - July 24, 2020

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

GLFW is an Open Source, multi-platform library for OpenGL, OpenGL ES and Vulkan development on the desktop. It provides a simple API for creating window contexts 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/png license](#).

- Gives you a 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 ramps
- Support for keyboard, mouse, gamepad, timer 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 compile-time options for platform specific features
- Community-maintained bindings for many different languages



mjb - July 24, 2020

Setting Up GLFW

128

```
#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 );
}
```



mjb - July 24, 2020

You Can Also Query What Vulkan Extensions GLFW Requires

129

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

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

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

Found 2 GLFW Required Instance Extensions:
VK_KHR_surface
VK_KHR_win32_surface



mb - July 24, 2020

GLFW Keyboard Callback

130

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



mb - July 24, 2020

GLFW Mouse Button Callback

131

```
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;
            printf( 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;
    }
    else
    {
        ActiveButton &= ~b;   // clear the proper bit
    }
}
```



mb - July 24, 2020

GLFW Mouse Motion Callback

132

```
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;
}
```



mb - July 24, 2020

Looping and Closing GLFW 133

```

while( glfwWindowShouldClose( MainWindow ) == 0 )
{
    glfwPollEvents(); // Does not block – processes any waiting events, then returns
    Time = glfwGetTime(); // elapsed time, in double-precision seconds
    UpdateScene();
    RenderScene();
}

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

```

SIGGRAPH THINK
CREATE. DESIGN. EXPERIMENT. REDEFINE

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();
```

SIGGRAPH THINK
CREATE. DESIGN. EXPERIMENT. REDEFINE

mjb - July 24, 2020

135

Vulkan.

GLM

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
CREATE. DESIGN. EXPERIMENT. REDEFINE

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
```

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.

SIGGRAPH THINK
CREATE. DESIGN. EXPERIMENT. REDEFINE

mjb - July 24, 2020

Why are we even talking about this?

The Most Useful GLM Variables, Operations, and Functions

The Most Useful GLM Variables, Operations, and Functions

139

```
// 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):
```

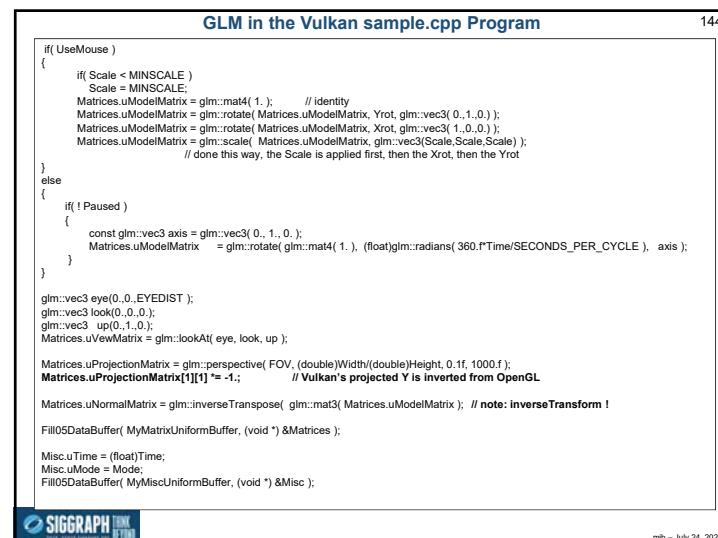
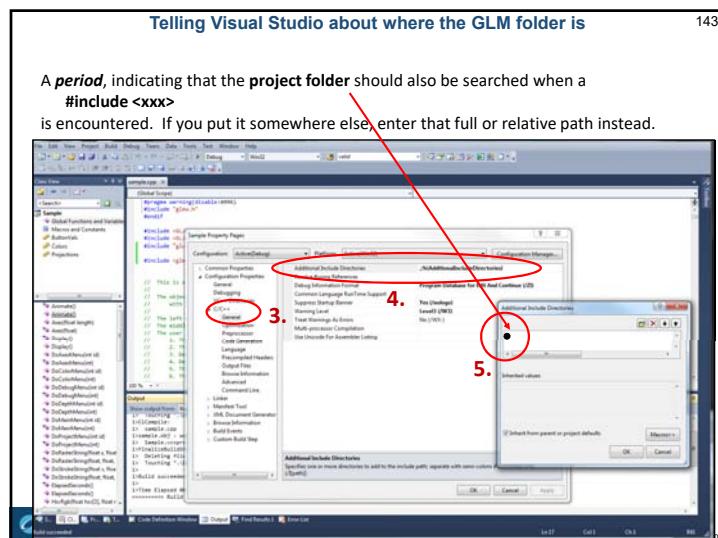
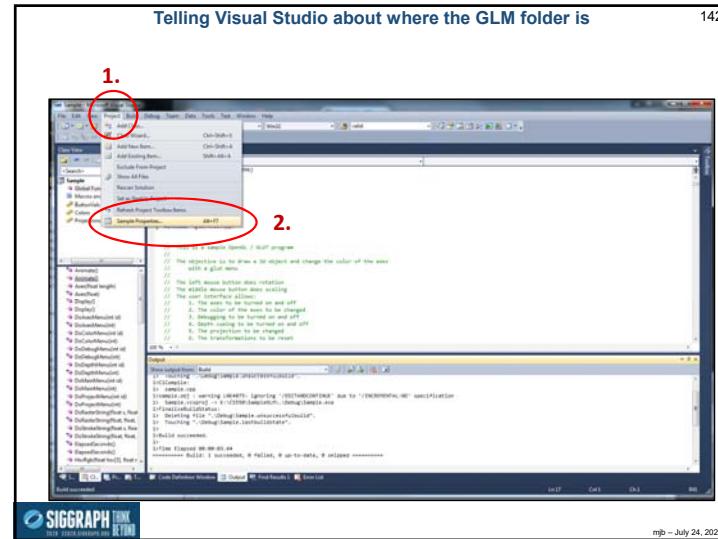
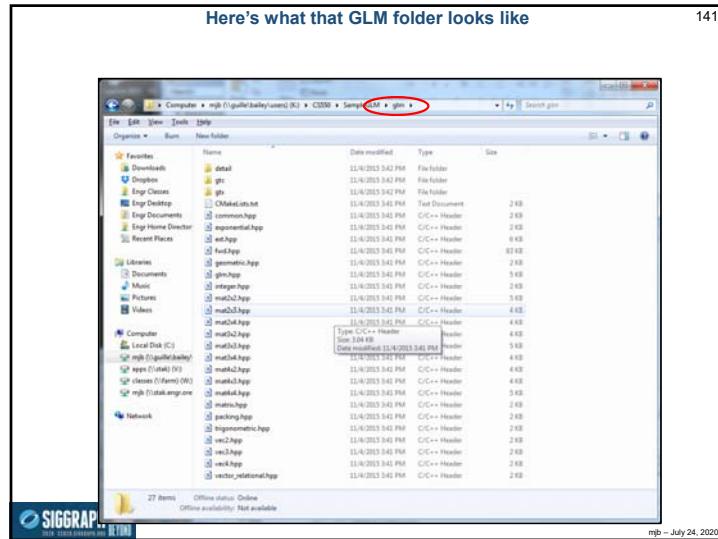
```
glm::mat4 glm::lookAt( glm::vec3 const & eye, glm::vec3 const & look, glm::vec3 const & up );
```

SIGGRAPH 2019

July 24, 2019

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.

Name	Date modified	Type	Size
Debug	11/14/2013 3:52 PM	File folder	
Debug.pdb	11/14/2013 3:49 PM	File folder	
freeglut.h	3/14/2013 3:34 PM	C/C++ Header	1 KB
freeglut.lib	3/14/2013 3:31 PM	Object File Library	39 KB
freeglut_adv.h	3/14/2013 3:34 PM	C/C++ Header	11 KB
freeglut_std.h	7/22/2013 8:45 AM	C/C++ Header	27 KB
glew.h	9/14/2013 3:27 PM	C/C++ Header	977 KB
glew32.lib	9/14/2013 3:27 PM	Object File Library	363 KB
glut.h	3/14/2013 3:34 PM	C/C++ Header	1 KB
include	3/14/2013 3:34 PM	Object File Library	2,000 KB
libglew32.lib	6/21/2006 1:00 AM	Object File Library	369 KB
sample.cpp	13/14/2013 3:51 PM	C/C++ Source	
sample.dsp	1/7/2002 11:17 AM	VC++ 6 Project	5 KB
sample.dsw	1/7/2002 11:26 AM	VC++ 6 Workspace	1 KB
sample.lib	3/12/2013 11:17 AM	VC++ 6 Intellisense	14,043 KB
sample.ncb	3/12/2013 10:54 AM	GIFT File	48 KB
sample.opt	12/30/2002 10:12 PM	PLUG File	2 KB
sample.plg	12/30/2002 10:12 PM	PLUG File	
SampleGLM.sln	1/5/2013 8:49 AM	Microsoft Visual Studio Solution	1 KB
SampleGLM.vcb	8/22/2013 4:47 PM	Visual Studio Sln	17 KB
SampleGLM.vcxproj	8/21/2008 11:54 AM	VC++ Project	6 KB
SampleGLM.vcxproj.filters	8/21/2008 11:54 AM	VC++ Filter Proj	2 KB
SampleGLM.vcxproj.user	8/21/2008 11:54 AM	VC++ Project	8 KB
SampleGLM.vip	8/21/2013 3:51 PM	VC++ Filter	



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:

SIGGRAPH THINK
THINK. DESIGN. INNOVATE. BEYOND.

Transformation Matrices 146

Translation

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

Rotation about X

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

Scaling

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

Rotation about Y

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

Rotation about Z

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

SIGGRAPH THINK
THINK. DESIGN. INNOVATE. BEYOND.

How it Really Works :-) 147

<http://xkcd.com>

SIGGRAPH THINK
THINK. DESIGN. INNOVATE. BEYOND.

The Rotation Matrix for an Angle (θ) about an Arbitrary Axis (A_x, A_y, A_z) 148

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

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

SIGGRAPH THINK
THINK. DESIGN. INNOVATE. BEYOND.

Compound Transformations

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

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

Say it

mb - July 24, 2020

Matrix Multiplication is not Commutative

Rotate, then translate

Translate, then rotate

mb - July 24, 2020

Matrix Multiplication is Associative

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

$$\begin{pmatrix} x' \\ y' \\ z' \\ 1 \end{pmatrix} = \left(\begin{bmatrix} T_{+A,+B} \end{bmatrix} \cdot \begin{bmatrix} R_\theta \end{bmatrix} \cdot \begin{bmatrix} T_{-A,-B} \end{bmatrix} \right) \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix}$$

One matrix – the Current Transformation Matrix, or CTM

mb - July 24, 2020

One Matrix to Rule Them All

$$\begin{pmatrix} x' \\ y' \\ z' \\ 1 \end{pmatrix} = \left(\begin{bmatrix} T_{+A,+B} \end{bmatrix} \cdot \begin{bmatrix} R_\theta \end{bmatrix} \cdot \begin{bmatrix} T_{-A,-B} \end{bmatrix} \right) \begin{pmatrix} x \\ y \\ z \\ 1 \end{pmatrix}$$

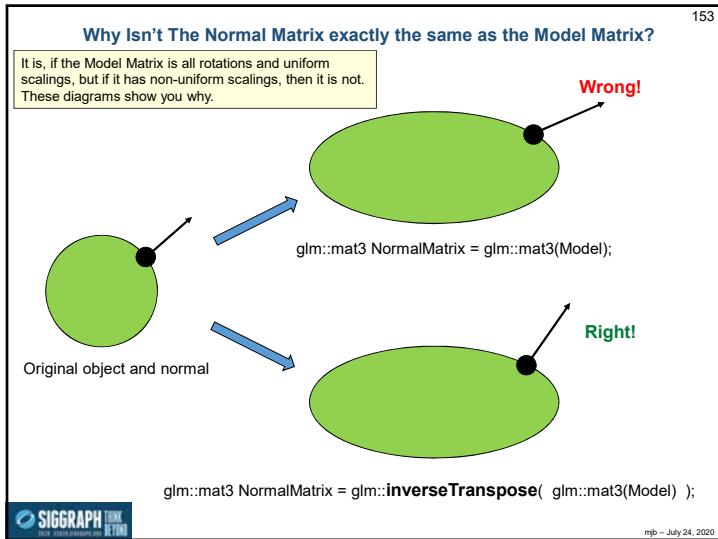
```
glm::mat4 Model = glm::mat4( 1. );
Model = glm::translate(Model, glm::vec3(A, B, 0. ) );
Model = glm::rotate(Model, thetaRadians, glm::vec3(Ax, Ay, Az) );
Model = glm::translate(Model, glm::vec3(-A, -B, 0. ) );

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) );
```

mb - July 24, 2020



154

Vulkan. Instancing

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>

mjb - July 24, 2020

SIGGRAPH THINK
CREATE. DESIGN. INNOVATE. BEYOND.

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 different?

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

mjb - July 24, 2020

SIGGRAPH THINK
CREATE. DESIGN. INNOVATE. BEYOND.

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;

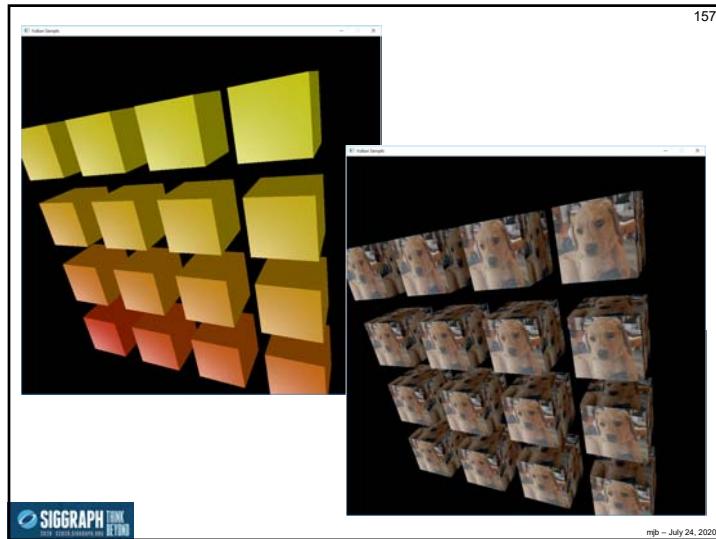
float xdelta = DELTA * float( gl_InstanceIndex % 4 );
float ydelta = DELTA * float( gl_InstanceIndex / 4 );
vColor = vec3( 1., float( (1.+gl_InstanceIndex) ) / float( NUMINSTANCES ), 0. );

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]
```

mjb - July 24, 2020

SIGGRAPH THINK
CREATE. DESIGN. INNOVATE. BEYOND.

**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]
```



158

mjb – July 24, 2020

Vulkan.

The Graphics Pipeline Data Structure

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 logo

159

mjb – July 24, 2020

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.

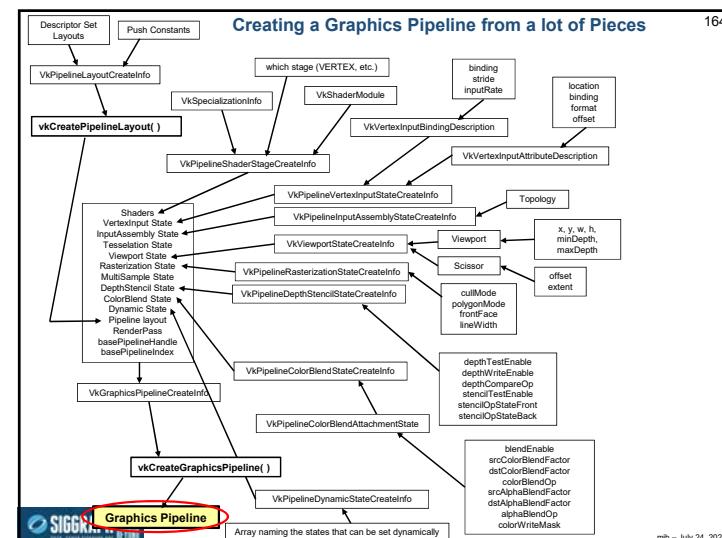
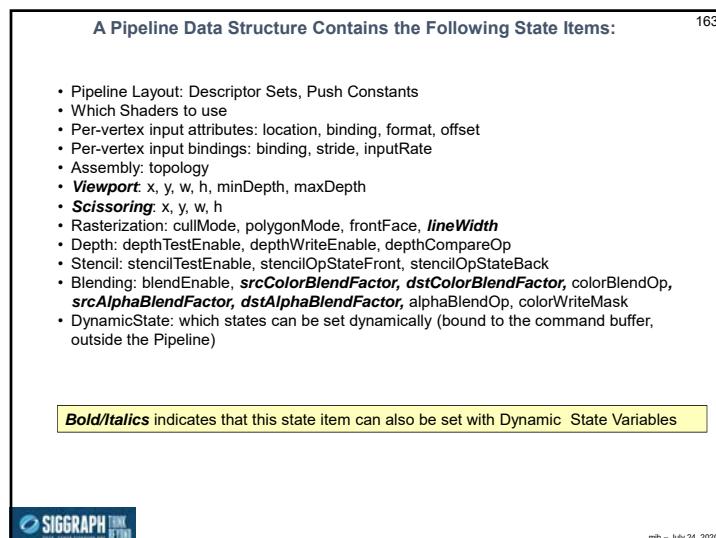
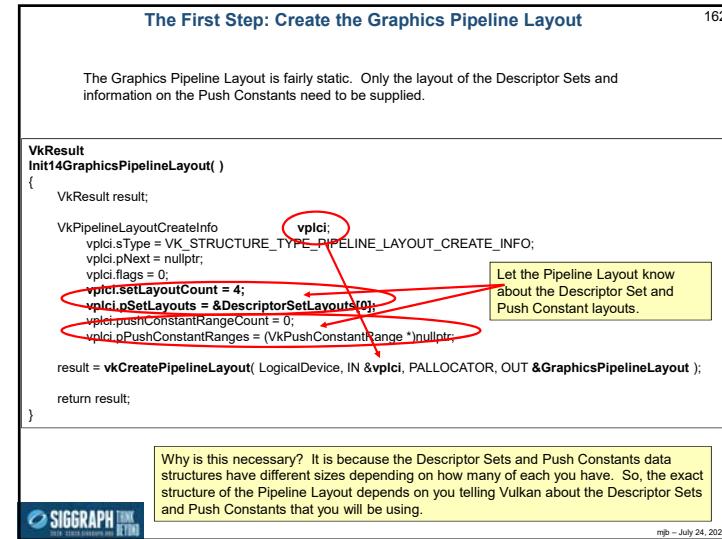
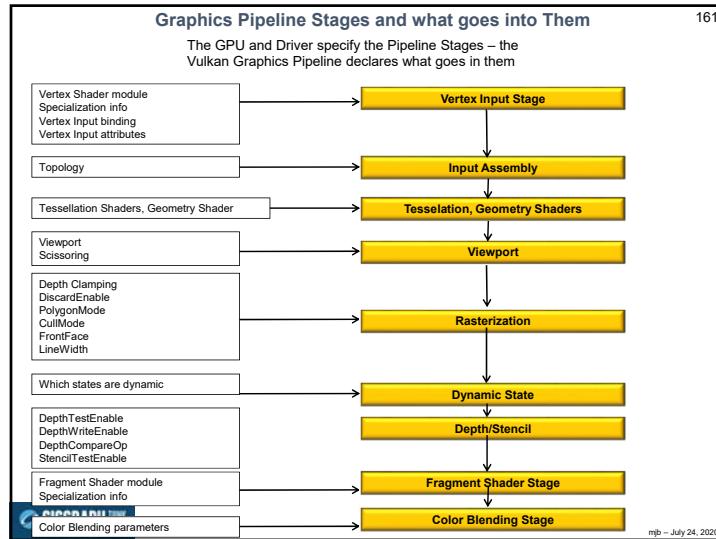
Here's what you need to know:

1. The Vulkan Graphics Pipeline is like what OpenGL would call "The State", or "The Context". It is a **data structure**.
2. The Vulkan Graphics Pipeline is *not* the processes that OpenGL would call "the graphics pipeline".
3. 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.
4. The shaders get compiled the rest of the way when their Graphics Pipeline gets created.

SIGGRAPH logo

160

mjb – July 24, 2020



Creating a Typical Graphics Pipeline 165

```

VkResult
Init14GraphicsVertexFragmentPipeline( VkShaderModule vertexShader, VkShaderModule fragmentShader,
                                     VkPrimitiveTopology topology, OUT VkPipeline *pGraphicsPipeline )
{
    #ifdef ASSUMPTIONS
        vpbsci.inputRate = VK_VERTEX_INPUT_RATE_VERTEX;
        vpbsci.depthClampEnable = VK_FALSE;
        vpbsci.rasterizerDiscardEnable = VK_FALSE;
        vpbsci.polygonMode = VK_POLYGON_MODE_FILL;
        vpbsci.cullMode = VK_CULL_MODE_NONE; // best to do this because of the projectionMatrix[1][1] *= -1;
        vpbsci.frontFace = VK_FRONT_FACE_COUNTER_CLOCKWISE;
        vpbsci.rasterizationSamples = VK_SAMPLE_COUNT_ONE_BIT;
        vpbsci.blendEnable = VK_FALSE;
        vpbsci.logicOpEnable = VK_FALSE;
        vpbsci.depthTestEnable = VK_TRUE;
        vpbsci.depthWriteEnable = VK_TRUE;
        vpssci.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.

```

 mb - July 24, 2020

The Shaders to Use 166

```

VkPipelineShaderStageCreateInfo vpsc[2];
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;

#def 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
#endif

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 vvid[1]; // an array containing one of these per buffer being used
vvid[0].binding = 0; // which binding # this is
vvid[0].stride = sizeof( struct vertex ); // bytes between successive
vvid[0].inputRate = VK_VERTEX_INPUT_RATE_VERTEX;

#def CHOICES
VK_VERTEX_INPUT_RATE_VERTEX
VK_VERTEX_INPUT_RATE_INSTANCE
#endif

```

 mb - July 24, 2020

Link in the Per-Vertex Attributes 167

```

VkVertexInputAttributeDescription vviad[4]; // an array containing one of these per vertex attribute in all bindings
// 4 = vertex, normal, color, texture
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_VEC2 VK_FORMAT_R32G32_SFLOAT
#define VK_FORMAT_VEC2I VK_FORMAT_R32I_SFLOAT
#define VK_FORMAT_XY VK_FORMAT_R32_SFLOAT
#define VK_FORMAT_SFLOAT VK_FORMAT_R32_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; // 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

```

 mb - July 24, 2020

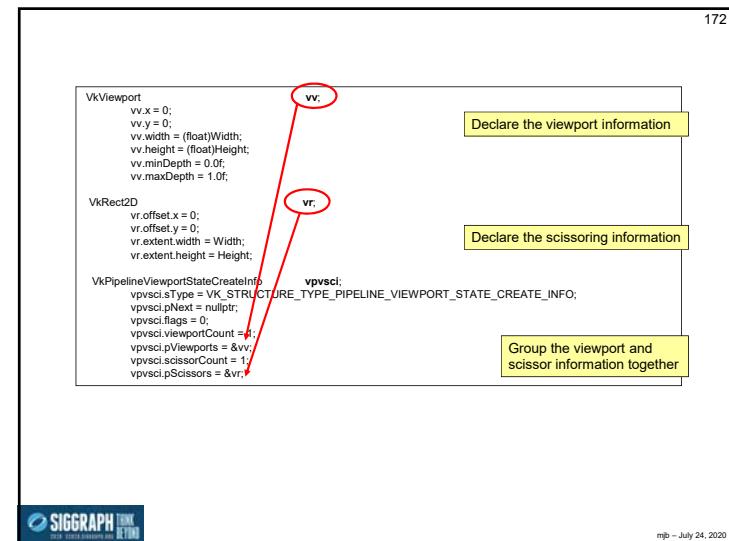
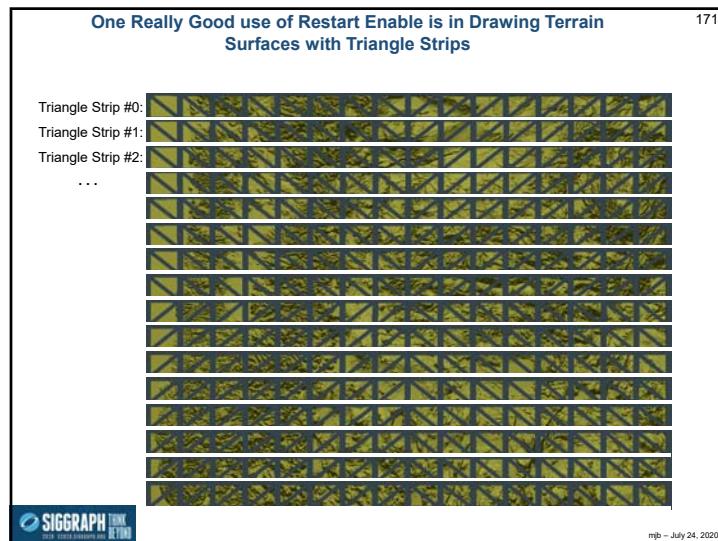
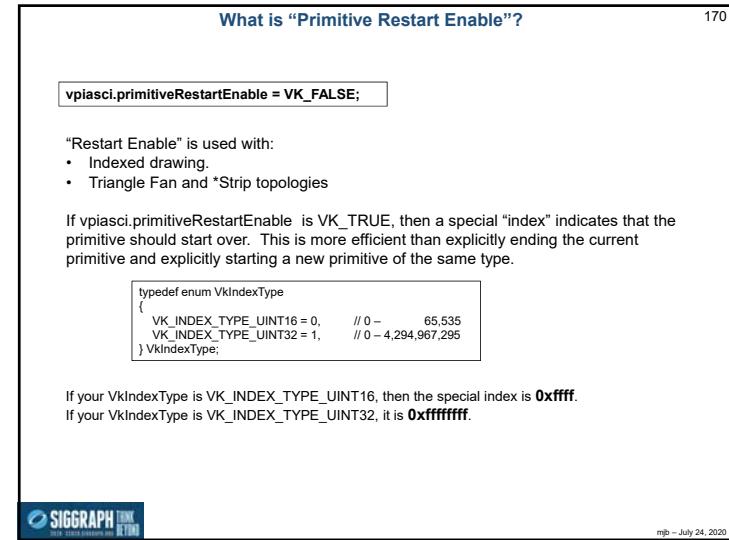
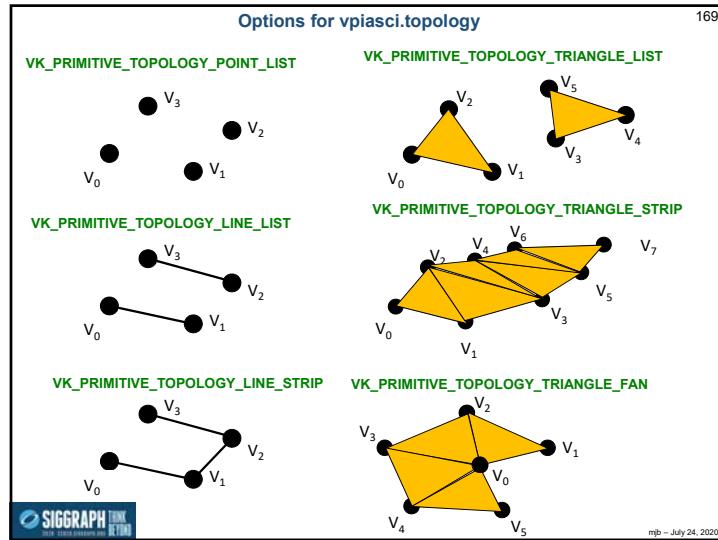
VkPipelineVertexInputStateCreateInfo vpvisci; // used to describe the input vertex attributes
vpvisci.sType = VK_STRUCTURE_TYPE_PIPELINE_VERTEX_INPUT_STATE_CREATE_INFO;
vpvisci.pNext = nullptr;
vpvisci.flags = 0;
vpvisci.vertexBindingDescriptionCount = 1;
vpvisci.pVertexBindingDescriptions = vvid;
vpvisci.vertexAttributeDescriptionCount = 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;
#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_LINE_STRIP_WITH_ADJACENCY
VK_PRIMITIVE_TOPOLOGY_TRIANGLE_LIST_WITH_ADJACENCY
VK_PRIMITIVE_TOPOLOGY_TRIANGLE_STRIP_WITH_ADJACENCY
#endif
vpiasci.primitiveRestartEnable = VK_FALSE;

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

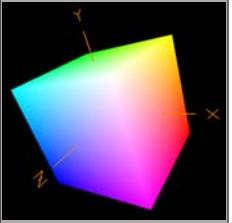
VkPipelineGeometryStateCreateInfo vpgsci;
vpgsci.sType = VK_STRUCTURE_TYPE_PIPELINE_TESSELLATION_STATE_CREATE_INFO;
vpgsci.pNext = nullptr;
vpgsci.flags = 0;

 mb - July 24, 2020

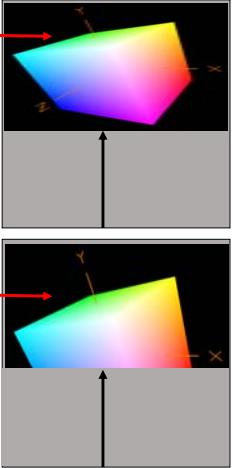


What is the Difference Between Changing the Viewport and Changing the Scissoring?¹⁷³

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 THINK
THINK. DESIGN. INNOVATE. BEYOND

Setting the Rasterizer State¹⁷⁴

```

VkPipelineRasterizationStateCreateInfo vprsc;
vprsc.sType = VK_STRUCTURE_TYPE_PIPELINE_RASTERIZATION_STATE_CREATE_INFO;
vprsc.flags = 0;
vprsc.depthClampEnable = VK_FALSE;
vprsc.rasterizerDiscardEnable = VK_FALSE;
vprsc.polygonMode = VK_POLYGON_MODE_FILL;
#endif
vprsc.cullMode = VK_CULL_MODE_NONE; // recommend this because of the projMatrix[1][1] *= -1;
#ifndef CHOICES
VK_CULL_MODE_NONE
VK_CULL_MODE_FRONT_BIT
VK_CULL_MODE_BACK_BIT
VK_CULL_MODE_FRONT_AND_BACK_BIT
#endif
vprsc.frontFace = VK_FRONT_FACE_COUNTER_CLOCKWISE;
#ifndef CHOICES
VK_FRONT_FACE_COUNTER_CLOCKWISE
VK_FRONT_FACE_CLOCKWISE
#endif
vprsc.depthBiasEnable = VK_FALSE;
vprsc.depthBiasConstantFactor = 0.f;
vprsc.depthBiasClamp = 0.f;
vprsc.depthBiasSlopeFactor = 0.f;
vprsc.lineWidth = 1.f;

```

Declare information about how the rasterization will take place

SIGGRAPH THINK
THINK. DESIGN. INNOVATE. BEYOND

What is “Depth Clamp Enable”?¹⁷⁵

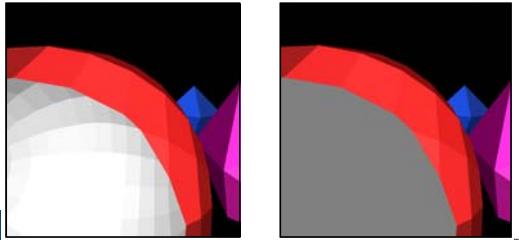
```
vprsc.depthClampEnable = VK_FALSE;
```

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).



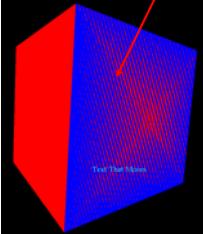
SIGGRAPH THINK
THINK. DESIGN. INNOVATE. BEYOND

What is “Depth Bias Enable”?¹⁷⁶

```
vprsc.depthBiasEnable = VK_FALSE;
vprsc.depthBiasConstantFactor = 0.f;
vprsc.depthBiasClamp = 0.f;
vprsc.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 THINK
THINK. DESIGN. INNOVATE. BEYOND

MultiSampling State 177

```

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

```

Declare information about how the multisampling will take place

We will discuss MultiSampling in a separate noteset.

SIGGRAPH THINK
Create. Share. Inspire. REACH

mjb - July 24, 2020

Color Blending State for each Color Attachment * 178

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

```

vpcbas
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.

$$\text{Color}_{\text{new}} = (1-\alpha) * \text{Color}_{\text{existing}} + \alpha * \text{Color}_{\text{incoming}}$$

$0 \leq \alpha \leq 1$.

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

SIGGRAPH THINK
Create. Share. Inspire. REACH

mjb - July 24, 2020

Raster Operations for each Color Attachment 179

```

vpcbsci
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;
#endif
CHOICES
VK_LOGIC_OP_CLEAR
VK_LOGIC_OP_AND
VK_LOGIC_OP_AND_REVERSE
VK_LOGIC_OP_COPY
VK_LOGIC_OP_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_INVERSE
VK_LOGIC_OP_INVERSE
VK_LOGIC_OP_COPY_INVERTED
VK_LOGIC_OP_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 THINK
Create. Share. Inspire. REACH

mjb - July 24, 2020

Which Pipeline Variables can be Set Dynamically 180

Just used as an example in the Sample Code

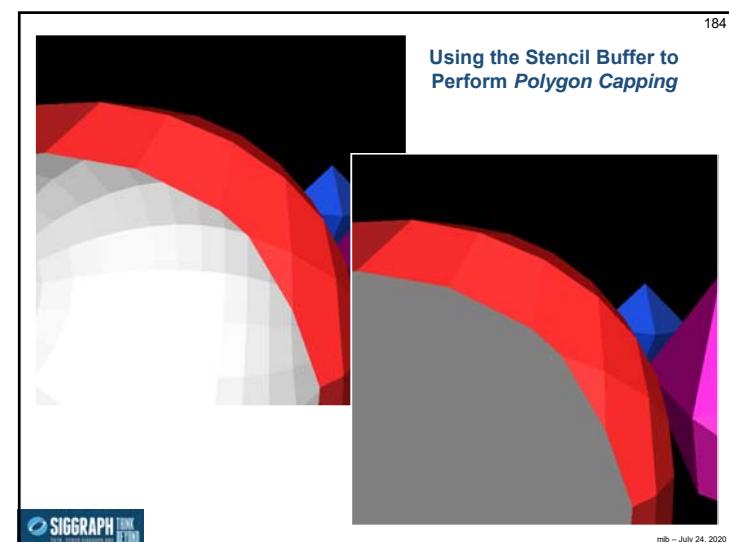
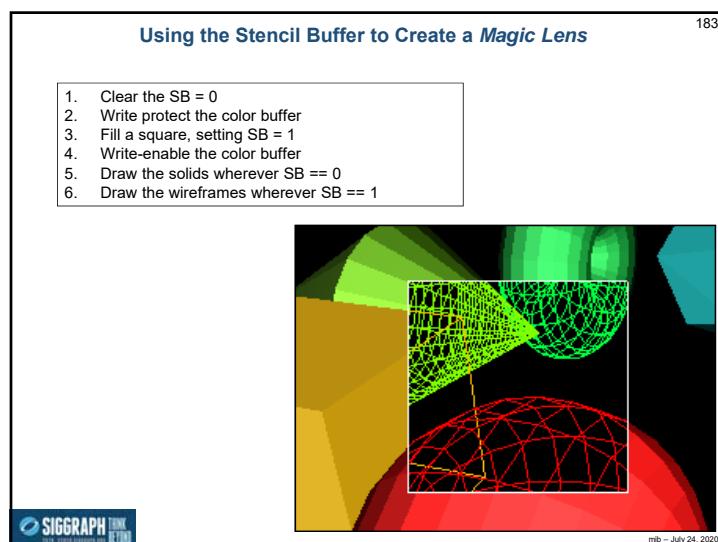
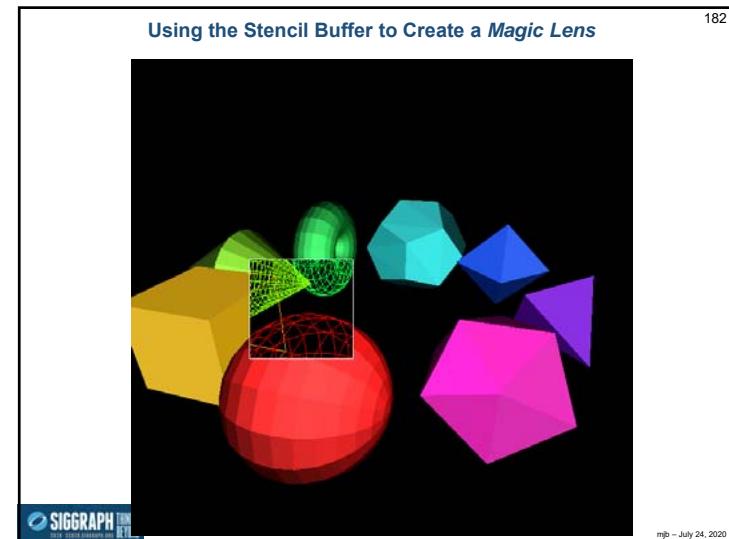
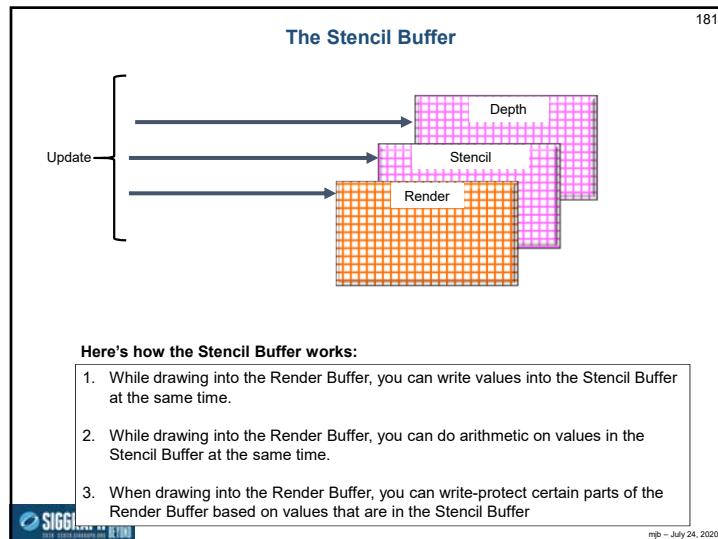
```

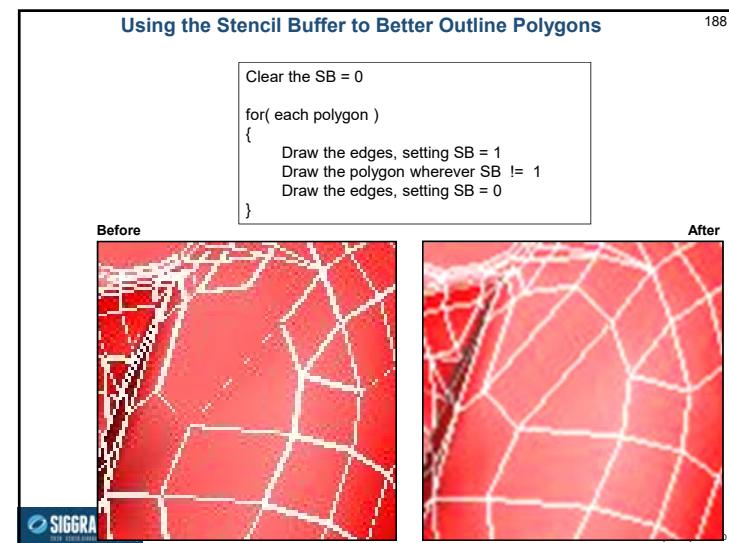
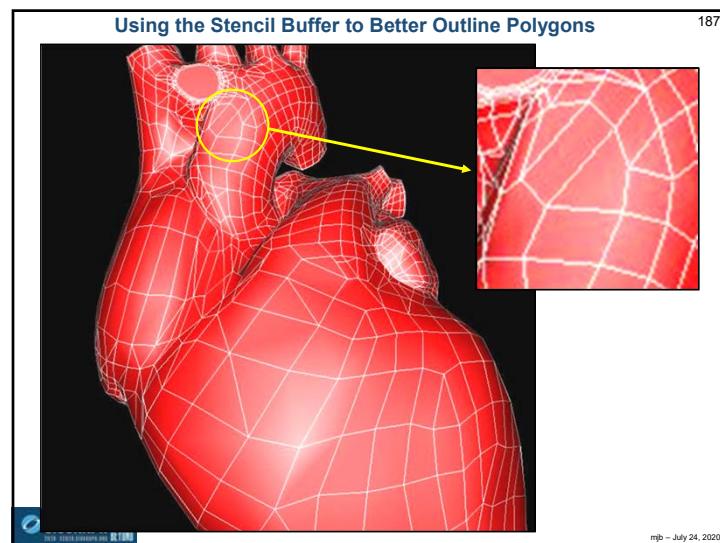
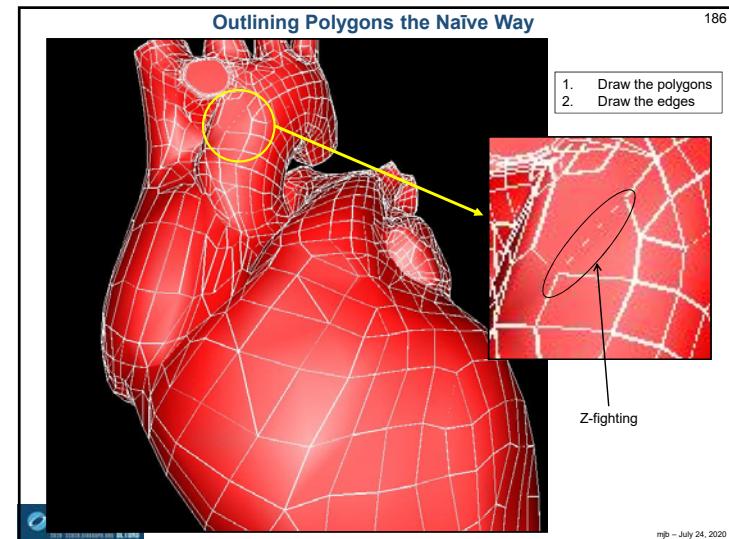
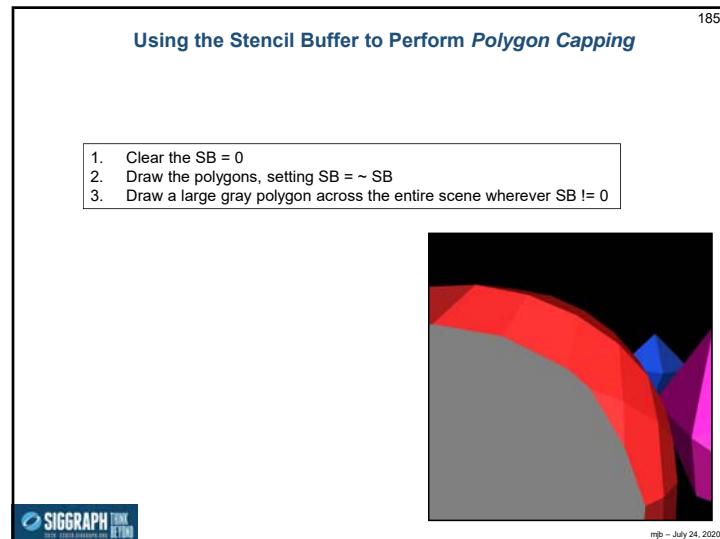
vds |= VK_DYNAMIC_STATE_VIEWPORT, VK_DYNAMIC_STATE_SCISSOR
VkDynamicState
#ifdef CHOICES
VK_DYNAMIC_STATE_VIEWPORT
VK_DYNAMIC_STATE_SCISSOR
VK_DYNAMIC_STATE_LINE_WIDTH
VK_DYNAMIC_STATE_DEPTH_BIAS
VK_DYNAMIC_STATE_BLEND_CONSTANTS
VK_DYNAMIC_STATE_DEPTH_BOUNDS
VK_DYNAMIC_STATE_STENCIL_COMPARE_MASK
VK_DYNAMIC_STATE_STENCIL_WRITE_MASK
VK_DYNAMIC_STATE_STENCIL_REFERENCE
#endif
VkPipelineDynamicStateCreateInfo
vpdsi.sType = VK_STRUCTURE_TYPE_PIPELINE_DYNAMIC_STATE_CREATE_INFO;
vpdsi.pNext = nullptr;
vpdsi.flags = 0;
vpdsi.dynamicStateCount = 0;
vpdsi.pDynamicStates = vds;
// leave turned off for now

```

SIGGRAPH THINK
Create. Share. Inspire. REACH

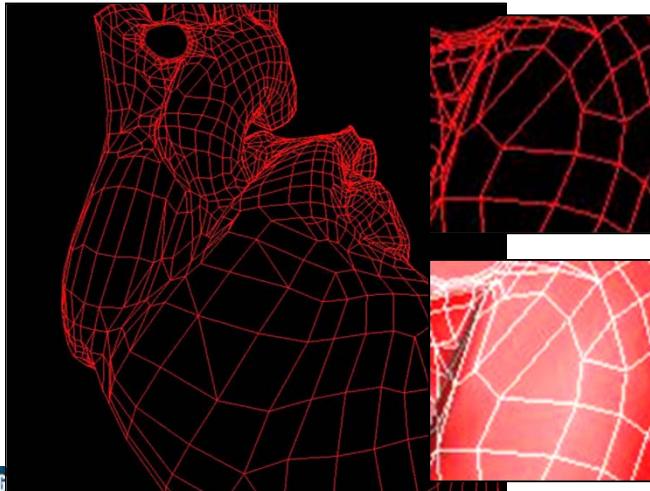
mjb - July 24, 2020





Using the Stencil Buffer to Perform Hidden Line Removal

189



mb - July 24, 2020

Stencil Operations for Front and Back Faces

190

```

VkStencilOpState
    vsosf.depthFailOp = VK_STENCIL_OP_KEEP; // what to do if depth operation fails
    vsosf.failOp = VK_STENCIL_OP_KEEP; // what to do if stencil operation fails
    vsosf.passOp = VK_STENCIL_OP_KEEP; // what to do if stencil operation succeeds

#define CHOICES
VK_STENCIL_OP_KEEP
-- keep the stencil value as it is
VK_STENCIL_OP_ZERO
-- set stencil value to 0
VK_STENCIL_OP_REPLACE
-- replace stencil value with the reference value
VK_STENCIL_OP_INCR_AND_CLAMP
-- increment stencil value
VK_STENCIL_OP_DECREMENT_AND_CLAMP
-- decrement stencil value
VK_STENCIL_OP_INVERT
-- bit-invert stencil value
VK_STENCIL_OP_INCREMENT_AND_WRAP
-- increment stencil value
VK_STENCIL_OP_DECREMENT_AND_WRAP
-- decrement stencil value
#endif

    vsosf.compareOp = VK_COMPARE_OP_NEVER;
#define CHOICES
VK_COMPARE_OP_NEVER
-- never succeeds
VK_COMPARE_OP_LESS
-- succeeds if new depth value is < the existing value
VK_COMPARE_OP_EQUAL
-- succeeds if new depth value is == the existing value
VK_COMPARE_OP_LESS_OR_EQUAL
-- succeeds if new depth value is <= the existing value
VK_COMPARE_OP_GREATER
-- succeeds if new depth value is > the existing value
VK_COMPARE_OP_NOT_EQUAL
-- succeeds if new depth value is != the existing value
VK_COMPARE_OP_GREATER_OR_EQUAL
-- succeeds if new depth value is >= the existing value
VK_COMPARE_OP_ALWAYS
-- always succeeds
#endif

    vsosf.compareMask = -0;
    vsosf.writeMask = -0;
    vsosf.reference = 0;

VkStencilOpState
    vsosb.depthFailOp = VK_STENCIL_OP_KEEP; // what to do if depth operation fails
    vsosb.failOp = VK_STENCIL_OP_KEEP; // what to do if stencil operation fails
    vsosb.passOp = VK_STENCIL_OP_KEEP; // what to do if stencil operation succeeds
    vsosb.compareOp = VK_COMPARE_OP_NEVER;
    vsosb.compareMask = -0;
    vsosb.writeMask = -0;
    vsosb.reference = 0;

```

mb - July 24, 2020

Operations for Depth Values

191

```

VkPipelineDepthStencilCreateInfo
    vpdssci.sType = VK_STRUCTURE_TYPE_PIPELINE_DEPTH_STENCIL_STATE_CREATE_INFO;
    vpdssci.pNext = nullptr;
    vpdssci.flags = 0;
    vpdssci.depthTestEnable = VK_TRUE;
    vpdssci.depthWriteEnable = VK_TRUE;
    vpdssci.depthCompareOp = VK_COMPARE_OP_LESS;
    VK_COMPARE_OP_NEVER
        -- never succeeds
    VK_COMPARE_OP_LESS
        -- succeeds if new depth value is < the existing value
    VK_COMPARE_OP_EQUAL
        -- succeeds if new depth value is == the existing value
    VK_COMPARE_OP_LESS_OR_EQUAL
        -- succeeds if new depth value is <= the existing value
    VK_COMPARE_OP_GREATER
        -- succeeds if new depth value is > the existing value
    VK_COMPARE_OP_NOT_EQUAL
        -- succeeds if new depth value is != the existing value
    VK_COMPARE_OP_GREATER_OR_EQUAL
        -- succeeds if new depth value is >= the existing value
    VK_COMPARE_OP_ALWAYS
        -- always succeeds
#endif

    vpdssci.depthBoundsTestEnable = VK_FALSE;
    vpdssci.front = vsosf;
    vpdssci.back = vsosb;
    vpdssci.minDepthBounds = 0;
    vpdssci.maxDepthBounds = 1;
    vpdssci.stencilTestEnable = VK_FALSE;

```

mb - July 24, 2020

Putting it all Together! (finally...)

192

VkPipeline GraphicsPipeline;

```

VkGraphicsPipelineCreateInfo
    vgpci.sType = VK_STRUCTURE_TYPE_GRAPHICS_PIPELINE_CREATE_INFO;
    vgpci.pNext = nullptr;
    vgpci.flags = 0;
#define CHOICES
VK_PIPELINE_CREATE_DISABLE_OPTIMIZATION_BIT
VK_PIPELINE_CREATE_ALLOW_DERIVATIVES_BIT
VK_PIPELINE_CREATE_DERIVATIVE_BIT
#endif

    vgpci.stageCount = 2; // number of stages in this pipeline
    vgpci.pStages = vpsci;
    vgpci.pVertexInputState = &vpsci;
    vgpci.pInputAssemblyState = &vpsci;
    vgpci.pTessellationState = (VkPipelineTessellationStateCreateInfo *)nullptr;
    vgpci.pViewportsState = &vpsci;
    vgpci.pRasterizationState = &vprsc;
    vgpci.pMultisampleState = &vpmssi;
    vgpci.pDepthStencilState = &vpdssci;
    vgpci.pColorBlendState = &vpcbsci;
    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);

```

Group all of the individual state information and create the pipeline

mb - July 24, 2020

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

193

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



mjb - July 24, 2020

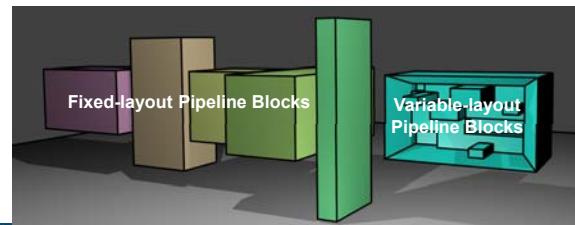
Sidebar: What is the Organization of the Pipeline Data Structure?

194

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
vv.x = 0;
vv.y = 0;
vv.width = (float)Width;
vv.height = (float)Height;
vv.minDepth = 0.0f;
vv.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

Vulkan. Descriptor Sets

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>



mjb - July 24, 2020

In OpenGL

196

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

What are Descriptor Sets? 197

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
        }
    }
}

```

mb - July 24, 2020

Descriptor Sets 198

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;

```

mb - July 24, 2020

Descriptor Sets 199

CPU:	GPU:	GPU:
Uniform data created in a C++ data structure	Uniform data in a "blob"	Uniform data used in the shader
<ul style="list-style-type: none"> • Knows the CPU data structure • Knows where the data starts • Knows the data's size 	<ul style="list-style-type: none"> • Knows where the data starts • Knows the data's size • Doesn't know the CPU or GPU data structure 	<ul style="list-style-type: none"> • 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;
};

```

* "binary large object"

mb - July 24, 2020

Step 1: Descriptor Set Pools 200

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.

```

graph TD
    flags --> VkDescriptorPoolCreateInfo
    maxSets --> VkDescriptorPoolCreateInfo
    podSizeCount --> VkDescriptorPoolCreateInfo
    poolSizes --> VkDescriptorPoolCreateInfo
    device --> VkDescriptorPoolCreateInfo
    VkDescriptorPoolCreateInfo --> vkCreateDescriptorPool()
    vkCreateDescriptorPool() --> DescriptorSetPool

```

mb - July 24, 2020

```

201
VkResult
Init13DescriptorSetPool()
{
    VkResult result;

    VkDescriptorPoolSize
        vdps[4];
    vdps[0].type = VK_DESCRIPTOR_TYPE_UNIFORM_BUFFER;
    vdps[0].descriptorCount = 1;
    vdps[1].type = VK_DESCRIPTOR_TYPE_UNIFORM_BUFFER;
    vdps[1].descriptorCount = 1;
    vdps[2].type = VK_DESCRIPTOR_TYPE_UNIFORM_BUFFER;
    vdps[2].descriptorCount = 1;
    vdps[3].type = VK_DESCRIPTOR_TYPE_COMBINED_IMAGE_SAMPLER;
    vdps[3].descriptorCount = 1;

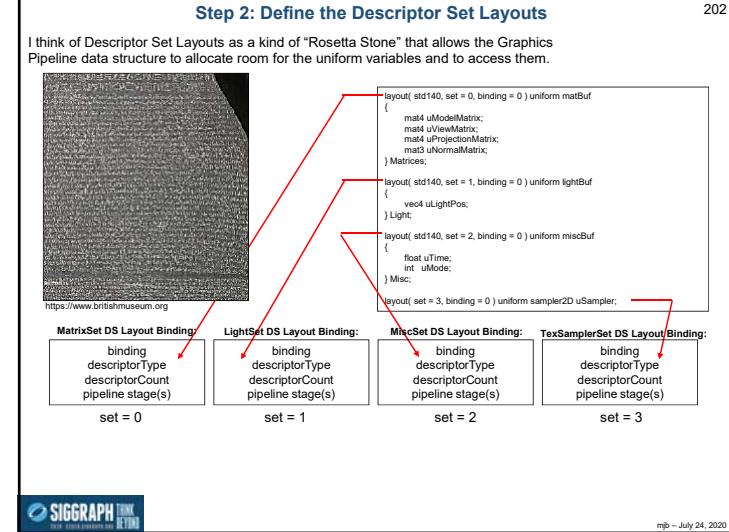
    #ifdef CHOICES
    VK_DESCRIPTOR_TYPE_SAMPLER
    VK_DESCRIPTOR_TYPE_SAMPLED_IMAGE
    VK_DESCRIPTOR_TYPE_STORAGE_IMAGE
    VK_DESCRIPTOR_TYPE_STORAGE_BUFFER
    VK_DESCRIPTOR_TYPE_UNIFORM_TEXEL_BUFFER
    VK_DESCRIPTOR_TYPE_STORAGE_TEXEL_BUFFER
    VK_DESCRIPTOR_TYPE_UNIFORM_BUFFER
    VK_DESCRIPTOR_TYPE_STORAGE_BUFFER
    VK_DESCRIPTOR_TYPE_UNIFORM_TEXEL_BUFFER_DYNAMIC
    VK_DESCRIPTOR_TYPE_STORAGE_TEXEL_BUFFER_DYNAMIC
    VK_DESCRIPTOR_TYPE_INPUT_ATTACHMENT
    #endif

    VkDescriptorPoolCreateInfo
        vdpCreateInfo;
    vdpCreateInfo.sType = VK_STRUCTURE_TYPE_DESCRIPTOR_POOL_CREATE_INFO;
    vdpCreateInfo.pNext = nullptr;
    vdpCreateInfo.flags = 0;
    vdpCreateInfo.maxSets = 4;
    vdpCreateInfo.poolSizeCount = 4;
    vdpCreateInfo.pPoolSizes = &vdps[0];
    vdpCreateInfo.pAllocateMask = nullptr;

    result = vkCreateDescriptorPool( LogicalDevice, IN &vdpCreateInfo, OUT &DescriptorPool );
    return result;
}

SIGGRAPH THINK
mjb - July 24, 2020

```



```

203
VkResult
Init13DescriptorSetLayouts()
{
    VkResult result;

    // DS #0:
    VkDescriptorSetLayoutBinding MatrixSet[1];
    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[1];
    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;

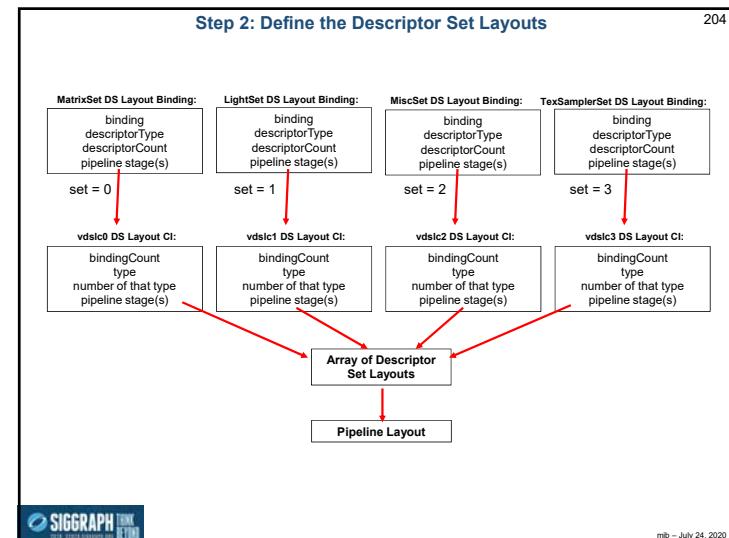
    // DS #2:
    VkDescriptorSetLayoutBinding MiscSet[1];
    MiscSet[0].binding = 0;
    MiscSet[0].descriptorType = VK_DESCRIPTOR_TYPE_UNIFORM_BUFFER;
    MiscSet[0].descriptorCount = 1;
    MiscSet[0].stageFlags = VK_SHADER_STAGE_VERTEX_BIT | VK_SHADER_STAGE_FRAGMENT_BIT;
    MiscSet[0].pImmutableSamplers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

    // DS #3:
    VkDescriptorSetLayoutBinding TexSamplerSet[1];
    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 );

SIGGRAPH THINK
mjb - July 24, 2020

```



205

```

VkDescriptorSetLayoutCreateInfo vdsi0;
vdsi0.sType = VK_STRUCTURE_TYPE_DESCRIPTOR_SET_LAYOUT_CREATE_INFO;
vdsi0.pNext = nullptr;
vdsi0.flags = 0;
vdsi0.bindingCount = 1;
vdsi0.pBindings = &MatrixSet[0];

VkDescriptorSetLayoutCreateInfo vdsi1;
vdsi1.sType = VK_STRUCTURE_TYPE_DESCRIPTOR_SET_LAYOUT_CREATE_INFO;
vdsi1.pNext = nullptr;
vdsi1.flags = 0;
vdsi1.bindingCount = 1;
vdsi1.pBindings = &LightSet[0];

VkDescriptorSetLayoutCreateInfo vdsi2;
vdsi2.sType = VK_STRUCTURE_TYPE_DESCRIPTOR_SET_LAYOUT_CREATE_INFO;
vdsi2.pNext = nullptr;
vdsi2.flags = 0;
vdsi2.bindingCount = 1;
vdsi2.pBindings = &MiscSet[0];

VkDescriptorSetLayoutCreateInfo vdsi3;
vdsi3.sType = VK_STRUCTURE_TYPE_DESCRIPTOR_SET_LAYOUT_CREATE_INFO;
vdsi3.pNext = nullptr;
vdsi3.flags = 0;
vdsi3.bindingCount = 1;
vdsi3.pBindings = &TexSamplerSet[0];

result = vkCreateDescriptorSetLayout( LogicalDevice, IN &vdsi0, PALLOCATOR, OUT &DescriptorSetLayouts[0] );
result = vkCreateDescriptorSetLayout( LogicalDevice, IN &vdsi1, PALLOCATOR, OUT &DescriptorSetLayouts[1] );
result = vkCreateDescriptorSetLayout( LogicalDevice, IN &vdsi2, PALLOCATOR, OUT &DescriptorSetLayouts[2] );
result = vkCreateDescriptorSetLayout( LogicalDevice, IN &vdsi3, PALLOCATOR, OUT &DescriptorSetLayouts[3] );

return result;
}

```

 mb - July 24, 2020

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

```

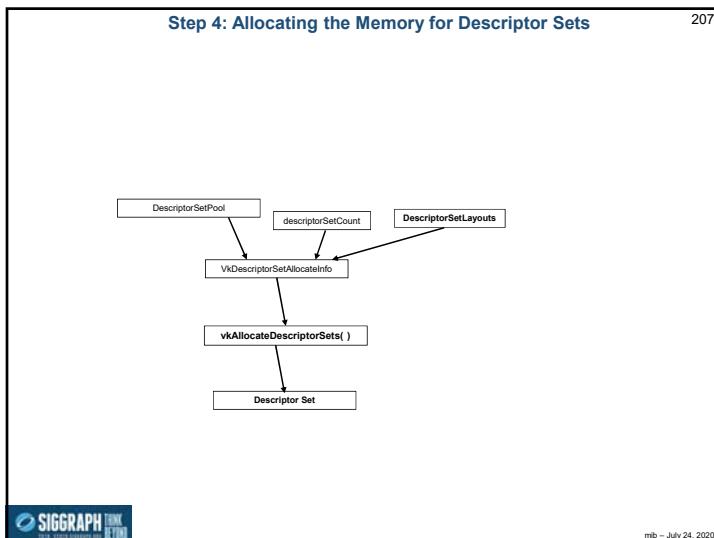
VkResult Init14GraphicsPipelineLayout()
{
    VkResult result;

    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.pPushConstantRanges = (VkPushConstantRange *)nullptr;

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

 mb - July 24, 2020

```



Step 4: Allocating the Memory for Descriptor Sets 208

```

VkResult Init13DescriptorSets()
{
    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] );
}

```

 mb - July 24, 2020

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

209

```

VkDescriptorBufferInfo          vdbi0;
vdbi0.buffer = MyMatrixUniformBuffer.buffer;
vdbi0.offset = 0;
vdbi0.range = sizeof(Matrices);

VkDescriptorBufferInfo          vdbi1;
vdbi1.buffer = MyLightUniformBuffer.buffer;
vdbi1.offset = 0;
vdbi1.range = sizeof(Light);

VkDescriptorBufferInfo          vdbi2;
vdbi2.buffer = MyMiscUniformBuffer.buffer;
vdbi2.offset = 0;
vdbi2.range = sizeof(Misc);

VkDescriptorImageInfo           vdii0;
vdii0.sampler = MyPuppyTexture.texSampler;
vdii0.imageView = MyPuppyTexture.textureView;
vdii0.imageLayout = VK_IMAGE_LAYOUT_SHADER_READ_ONLY_OPTIMAL;

```

This struct identifies what buffer it owns and how big it is

This struct identifies what buffer it owns and how big it is

This struct identifies what buffer it owns and how big it is

This struct identifies what texture sampler and image view it owns



mb - July 24, 2020

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

210

```

VkWriteDescriptorSet             vwdso;
// ds 0:
vwdso.sType = VK_STRUCTURE_TYPE_WRITE_DESCRIPTOR_SET;
vwdso.pNext = nullptr;
vwdso.dsSet = DescriptorSets[0];
vwdso.dsBinding = 0;
vwdso.dsArrayElement = 0;
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.dsSet = DescriptorSets[1];
vwdso1.dsBinding = 0;
vwdso1.dsArrayElement = 0;
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



mb - July 24, 2020

Step 5: Tell the Descriptor Sets where their data is

211

```

VkWriteDescriptorSet             vwdso2;
// ds 2:
vwdso2.sType = VK_STRUCTURE_TYPE_WRITE_DESCRIPTOR_SET;
vwdso2.pNext = nullptr;
vwdso2.dsSet = DescriptorSets[2];
vwdso2.dsBinding = 0;
vwdso2.dsArrayElement = 0;
vwdso2.descriptorCount = 1;
vwdso2.descriptorType = VK_DESCRIPTOR_TYPE_UNIFORM_BUFFER;
vwdso2.pBufferInfo = IN &vdbi2;
vwdso2.pImageInfo = (VkDescriptorImageInfo *)nullptr;
vwdso2.pTexelBufferView = (VkBufferView *)nullptr;

// ds 3:
VkWriteDescriptorSet             vwdso3;
vwdso3.sType = VK_STRUCTURE_TYPE_WRITE_DESCRIPTOR_SET;
vwdso3.pNext = nullptr;
vwdso3.dsSet = DescriptorSets[3];
vwdso3.dsBinding = 0;
vwdso3.dsArrayElement = 0;
vwdso3.descriptorCount = 1;
vwdso3.descriptorType = VK_DESCRIPTOR_TYPE_COMBINED_IMAGE_SAMPLER;
vwdso3.pBufferInfo = (VkDescriptorBufferInfo *)nullptr;
vwdso3.pImageInfo = IN &vdii0;
vwdso3.pTexelBufferView = (VkBufferView *)nullptr;

uint32_t copyCount = 0;

// this could have been done with one call and an array of VkWriteDescriptorSets:

vkUpdateDescriptorSets( LogicalDevice, 1, IN &vwdso, IN copyCount, (VkCopyDescriptorSet *)nullptr );
vkUpdateDescriptorSets( LogicalDevice, 1, IN &vwdso1, IN copyCount, (VkCopyDescriptorSet *)nullptr );
vkUpdateDescriptorSets( LogicalDevice, 1, IN &vwdso2, IN copyCount, (VkCopyDescriptorSet *)nullptr );
vkUpdateDescriptorSets( LogicalDevice, 1, IN &vwdso3, IN copyCount, (VkCopyDescriptorSet *)nullptr );

```

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



mb - July 24, 2020

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

212

```

VkGraphicsPipelineCreateInfo      vgpcl;
vgpcl.sType = VK_STRUCTURE_TYPE_GRAPHICS_PIPELINE_CREATE_INFO;
vgpcl.pNext = nullptr;
vgpcl.flags = 0;
#ifndef CHOICES
VK_PIPELINE_CREATE_DISABLE_OPTIMIZATION_BIT
VK_PIPELINE_CREATE_ALLOW_DERIVATIVES_BIT
VK_PIPELINE_CREATE_DERIVATIVE_BIT
#endif
vgpcl.stageCount = 2; // number of stages in this pipeline
vgpcl.pStages = vpsc;
vgpcl.pVertexInputState = &vpvisci;
vgpcl.pInputAssemblyState = &viaci;
vgpcl.pTessellationState = (VkPipelineTessellationStateCreateInfo *)nullptr;
vgpcl.pViewportState = &vpvsci;
vgpcl.pRasterizationState = &vprsci;
vgpcl.pMultisampleState = &vpmisci;
vgpcl.pDepthStencilState = &vpdssci;
vgpcl.pColorBlendState = &vpcbsci;
vgpcl.pDynamicState = &vdyds;
vgpcl.layout = ON GraphicsPipelineLayout; // subpass number
vgpcl.renderPass = IN RenderPass;
vgpcl.subpass = 0;
vgpcl.basePipelineHandle = (VkPipeline) VK_NULL_HANDLE;
vgpcl.basePipelineIndex = 0;

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

```



mb - July 24, 2020

Step 7: Bind Descriptor Sets into the Command Buffer when Drawing 213

```
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 THINK

mb - July 24, 2020

Sidebar: The Entire Collection of Descriptor Set Paths 214

VkDescriptorPoolCreateInfo vkCreateDescriptorPool()	Create the pool of Descriptor Sets for future use
VkDescriptorSetLayoutBinding VkDescriptorSetLayoutCreateInfo vkCreateDescriptorSetLayout() vkCreatePipelineLayout()	Describe a particular Descriptor Set layout and use it in a specific Pipeline layout
VkDescriptorSetAllocateInfo vkAllocateDescriptorSets()	Allocate memory for particular Descriptor Sets
VkDescriptorBufferInfo VkDescriptorImageInfo VkWriteDescriptorSet 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 THINK

mb - July 24, 2020

Sidebar: Why Do Descriptor Sets Need to Provide Layout Information to the Pipeline Data Structure? 215

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 THINK

mb - July 24, 2020

Sidebar: Why Do Descriptor Sets Need to Provide Layout Information to the Pipeline Data Structure? 216

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

SIGGRAPH THINK

mb - July 24, 2020

217

Vulkan.

Textures

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>



mjb - July 24, 2020

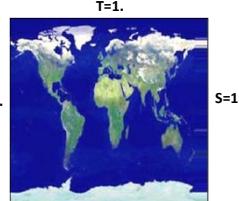
218

The Basic Idea

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.



T=1.
S=0.
S=1.
T=0.

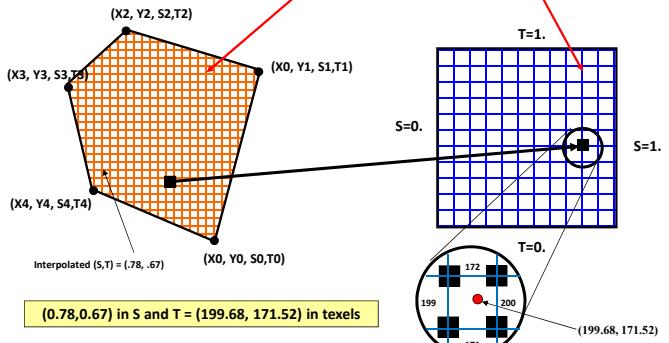


mjb - July 24, 2020

219

The Basic Idea

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



(0.78, 0.67) in S and T = (199.68, 171.52) in texels

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 lookup a color in the texture image.



mjb - July 24, 2020

220

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

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 );

    ...

```

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 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. }
    },
};

```

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,

$$s = \frac{x - X_{min}}{X_{max} - X_{min}} \quad t = \frac{y - Y_{min}}{Y_{max} - Y_{min}}$$

SIGGRAPH THINK
CREATE. DESIGN. ENTERTAIN. BEYOND.

mjb - July 24, 2020

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

Or, for a sphere,

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

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

SIGGRAPH THINK
CREATE. DESIGN. ENTERTAIN. BEYOND.

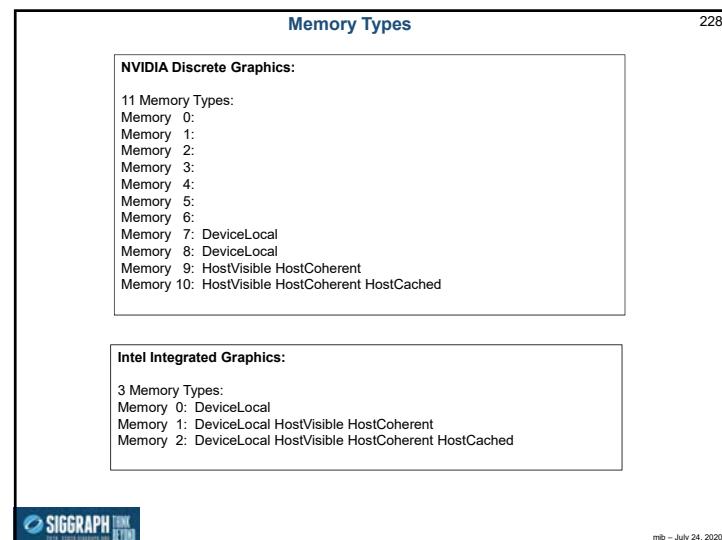
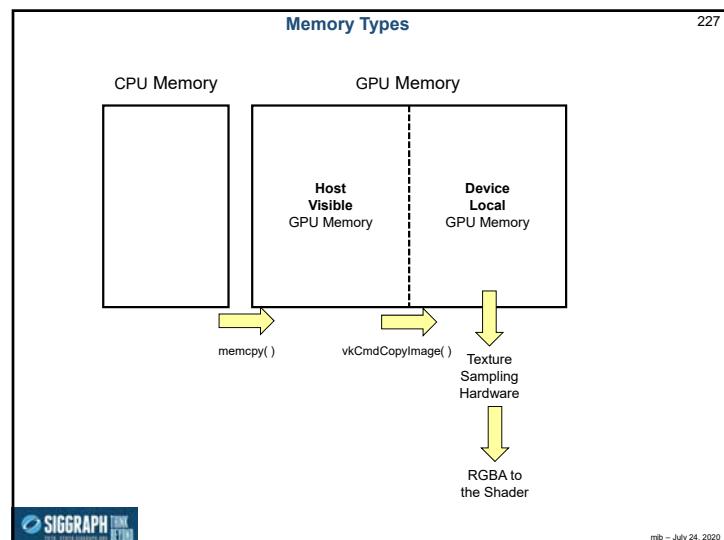
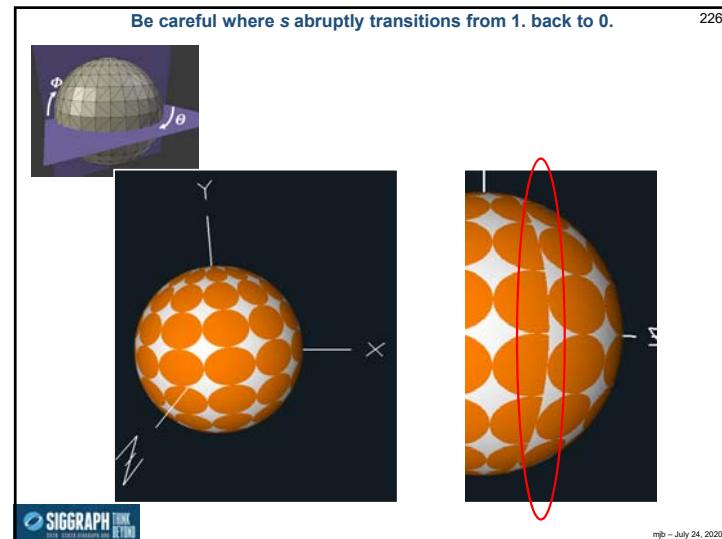
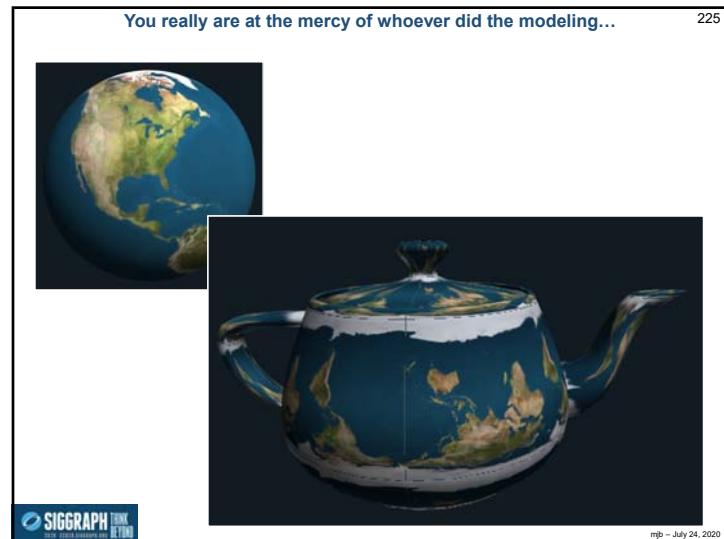
Using a Texture: How do you know what (s,t) to assign to each vertex? 224

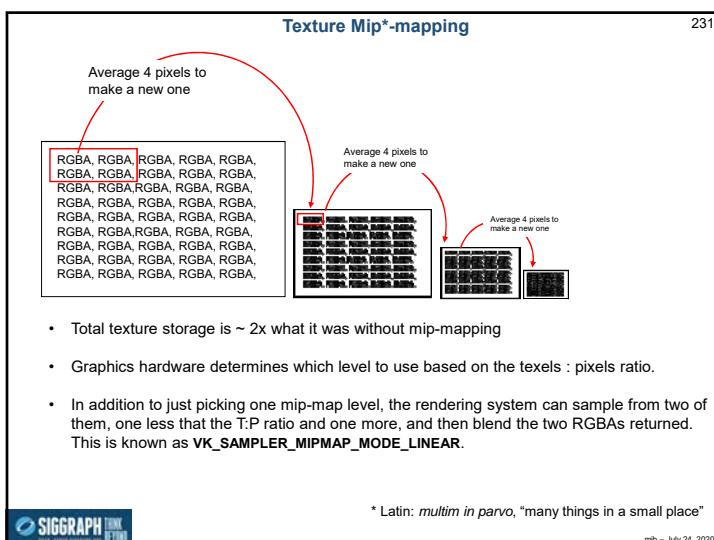
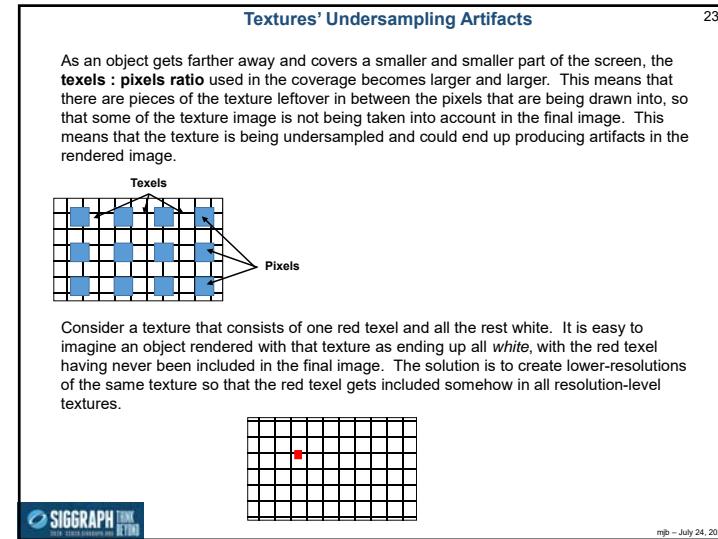
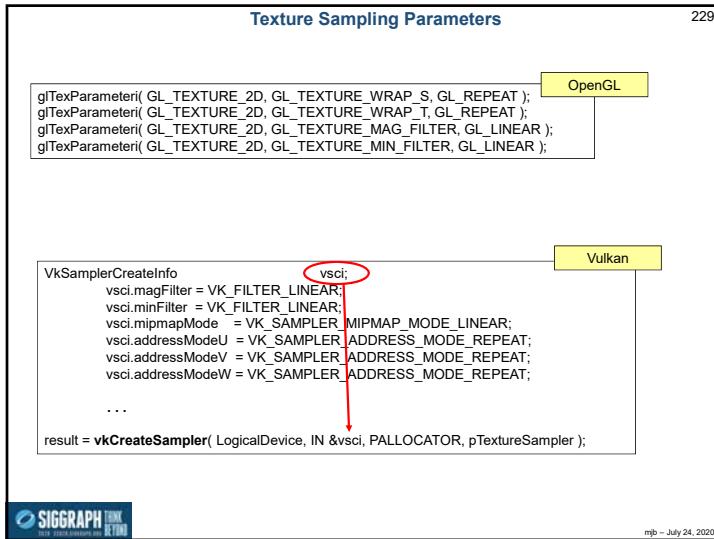
Uh-oh. Now what? Here's where it gets tougher...,

$$s = ? \quad t = ?$$

SIGGRAPH THINK
CREATE. DESIGN. ENTERTAIN. BEYOND.

mjb - July 24, 2020





```
VkResult
Int07TextureSampler(MyTexture *pMyTexture)
{
    VkResult result;
    VkSamplerCreateInfo vsc;
    vsc.sType = VK_STRUCTURE_TYPE_SAMPLER_CREATE_INFO;
    vsc.pNext = nullptr;
    vsc.pFilter = VK_FILTER_LINEAR;
    vsc.minFilter = VK_FILTER_LINEAR;
    vsc.magFilter = VK_FILTER_BILINEAR;
    vsc.addressModeU = VK_SAMPLER_ADDRESS_MODE_REPEAT;
    vsc.addressModeV = VK_SAMPLER_ADDRESS_MODE_REPEAT;
    vsc.addressModeW = VK_SAMPLER_ADDRESS_MODE_REPEAT;

#ifndef CHOICES
    VK_SAMPLER_ADDRESS_MODE_REPEAT
    VK_SAMPLER_ADDRESS_MODE_CLAMP_TO_EDGE
    VK_SAMPLER_ADDRESS_MODE_CLAMP_TO_BORDER
    VK_SAMPLER_ADDRESS_MODE_MIRROR_CLAMP_TO_EDGE
#endif

    vsc.mipmapLodBias = 0;
    vsc.compareMode = VK_COMPARE_MODE_NONE;
    vsc.compareFunc = VK_COMPARE_FUNC_ALWAYS;
    vsc.compareOp = VK_COMPARE_OP_NEVER;

#ifndef CHOICES
    VK_COMPARE_OP_NEVER
    VK_COMPARE_OP_LESS
    VK_COMPARE_OP_EQUAL
    VK_COMPARE_OP_LESS_OR_EQUAL
    VK_COMPARE_OP_GREATER
    VK_COMPARE_OP_GREATER_OR_EQUAL
    VK_COMPARE_OP_GREATER_OR_EQUAL
    VK_COMPARE_OP_ALWAYS
#endif

    vsc.mipmapLod = 0;
    vsc.mipLevelBias = 0;

    vsc.borderColor = VK_BORDER_COLOR_FLOAT_OPAQUE_BLACK;
    vsc.colorSpace = VK_COLOR_SPACE_SRGB_NONLINEAR_KHR;

#ifndef CHOICES
    VK_BORDER_COLOR_FLOAT_TRANSPARENT_BLACK
    VK_BORDER_COLOR_INT_TRANSPARENT_BLACK
    VK_BORDER_COLOR_FLOAT_OPAQUE_BLACK
    VK_BORDER_COLOR_INT_OPAQUE_BLACK
    VK_BORDER_COLOR_FLOAT_OPAQUE_WHITE
    VK_BORDER_COLOR_INT_OPAQUE_WHITE
#endif

    vsc.unnormalizedCoordinates = VK_FALSE; // VK_TRUE means we are use raw levels as the index
    vsc.compareOp = VK_COMPARE_OP_NEVER; // VK_FALSE means we are using the usual 0..1.

    result = vkCreateSampler(LogicalDevice, IN &vsc, PALLOCATOR, OUT 8pMyTexture->texSampler);
}
```

```

VkResult
Init3DTextureBuffer(INOUT MyTexture * pMyTexture)
{
    VkResult result;

    uint32_t texWidth = pMyTexture->width;
    uint32_t texHeight = pMyTexture->height;
    unsigned char *texture = pMyTexture->pixels;
    VkDeviceSize textureSize = texWidth * texHeight * 4; // rgba, 16

    VkImage stagingImage;
    VkImage textureImage;

    // ****
    // this first (...) is to create the staging image:
    // ****
    VkImageCreateInfo info;
    info.sType = VK_STRUCTURE_TYPE_IMAGE_CREATE_INFO;
    info.pNext = NULL;
    info.format = VK_FORMAT_R8G8B8A8_UNORM;
    info.extent.width = texWidth;
    info.extent.height = texHeight;
    info.extent.depth = 1;
    info.mipLevels = 1;
    info.arrayLayers = 1;
    info.samples = VK_SAMPLE_COUNT_1_BIT;
    info.tiling = VK_IMAGE_TILING_LINEAR;

    #ifdef _CHOICES_
    info.usage = VK_IMAGE_USAGE_TRANSFER_SRC_BIT;
    #endif
    #ifdef _CHOICES_
    info.usage |= VK_IMAGE_USAGE_TRANSFER_DST_BIT;
    #endif
    #ifdef _CHOICES_
    info.usage |= VK_IMAGE_USAGE_STORAGE_BIT;
    #endif
    #ifdef _CHOICES_
    info.usage |= VK_IMAGE_USAGE_COLOR_ATTACHMENT_BIT;
    #endif
    #ifdef _CHOICES_
    info.usage |= VK_IMAGE_USAGE_DEPTH_STENCIL_ATTACHMENT_BIT;
    #endif
    #ifdef _CHOICES_
    info.usage |= VK_IMAGE_USAGE_TRANSIENT_ATTACHMENT_BIT;
    #endif
    #ifdef _CHOICES_
    info.usage |= VK_IMAGE_USAGE_INPUT_ATTACHMENT_BIT;
    #endif

    result = vkCreateImage(vk,
                          &info,
                          &pMyTexture->stagingImage);
    if(result != VK_SUCCESS)
        return result;

    vci.sharingMode = VK_SHARING_MODE_EXCLUSIVE;

```

```
void *gpuMemory;
vkMapMemory(LogicalDevice, vdm, 0, VK_WHOLE_SIZE, 0, OUT &gpuMemory);
// 0 and 0 = offset and memory map flags

if (vsl.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[y * vsl.rowPitch], &texture[4 * y * texWidth], (size_t)(4 * texWidth));
    }
}

vkUnmapMemory(LogicalDevice, vdm);
}
//-----
```

mib - July 24, 2020

```

VK_STRUCTURE_TYPE_MEMORY_ALLOCATE_INFO vmaInfo;
vmaInfo.allocationSize = vma.size;
vmaInfo.alignment = vma.alignment;
vmaInfo.memoryTypeIndex = FindMemoryTypeWithVisible(); // because we want to mmap it

VKDeviceMemory vdm;
vkAllocateMemory(LogicalDevice, &vmaInfo, PALLOCATOR, OUT &vdm);
pMyTexture->vdm = vdm;

result = vkBindImageMemory(LogicalDevice, IN stagingImage, IN vdm, 0); // 0 == offset

// we have now created the staging image - fill with the pixel data:

VlImageResource vrs;
vrs.accessMask = VK_IMAGE_ACCESS_TYPE_MEMORY_ALLOCATE_INFO;
vrs.imageLayout = 0;
vrs.arrayLayer = 0;

VlImageSubresource vrsi;
vrsi.aspectMask = VK_IMAGE_ASPECT_COLOR_BIT;
vrsi.offset = 0;
vrsi.size = vdm.size;
vrsi.layer = 0;

VlImageSubresourceLayout vrl;
vrl.aspectMask = VK_IMAGE_ASPECT_COLOR_BIT;
vrl.offset = 0;
vrl.size = vdm.size;
vrl.layer = 0;

vkCreateImage(VL_VK, &vrs, PALLOCATOR, OUT &vrs);
vkGetImageMemoryRequirements(LogicalDevice, IN stagingImage, OUT &vrsi);
vkBindImageMemory(LogicalDevice, IN vrs, IN vrsi, OUT &vrl);

if(Verbose)
{
    viciQueueFamilyIndexCount = 0;
    viciQueueFamilyIndices = (const uint32_t*)NULLptr;
}

result = VcCreateImageLogicalDevice((LogicalDevice, IN vici, PALLOCATOR, OUT &stagingImage)); // allocated, but not filled yet

```

July 24, 2000

```

//=====
// this second (...) is to create the actual texture image:
//=====

{
    VkImageCreateInfo vci;
    vci.sType = VK_STRUCTURE_TYPE_IMAGE_CREATE_INFO;
    vci.pNext = nullptr;
    vci.imageType = VK_IMAGE_TYPE_2D;
    vci.format = VK_FORMAT_R8G8B8A_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_IMAGE_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.queueFamilyIndexCount = 0;
    vci.pQueueFamilyIndices = const_cast<int32_t*>(nullptr);
}

result = vkCreateImage(LogicalDevice, IN &vci, PALLOCATOR, OUT &textureImage); // allocated, but not file

VKMemoryRequirements vmr;
vkGetImageMemoryRequirements( LogicalDevice, IN textureImage, OUT &vmr);

if ( Verbose )
{
    fprintf( fPDebug, "Texture vmr size = %f\n", vmr.size );
    fprintf( fPDebug, "Texture vmr alignment = %f\n", vmr.alignment );
    fprintf( fPDebug, "Texture vmr memoryTypeBits = 0x%08X\n", vmr.memoryTypeBits );
    fflush( fPDebug );
}

VMMemoryAllocateInfo vmai;
vmai.sType = VK_STRUCTURE_TYPE_MEMORY_ALLOCATE_INFO;
vmai.pNext = nullptr;
vmai.allocationSize = vmr.size;
vmai.memoryTypeIndex = FindMemoryThatIsDeviceLocal(); // because we want to sample from it

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

result = vkBindImageMemory( LogicalDevice, IN textureImage, IN &vdm, 0 ); // == offset

```

m12 - July 24, 20

237

```

// copy pixels from the staging image to the texture:
VkCommandBufferBeginInfo vcbi;
vcbi.sType = VK_STRUCTURE_TYPE_COMMAND_BUFFER_BEGIN_INFO;
vcbi.pNext = NULL;
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 visr;
    visr.aspectMask = VK_IMAGE_ASPECT_COLOR_BIT;
    visr.baseMipLevel = 0;
    visr.baseArrayLayer = 0;
    visr.layerCount = 1;

    VkImageMemoryBarrier vimb;
    vimb.sType = VK_STRUCTURE_TYPE_IMAGE_MEMORY_BARRIER;
    vimb.pNext = NULLptr;
    vimb.oldLayout = VK_IMAGE_LAYOUT_PREINITIALIZED;
    vimb.newLayout = VK_IMAGE_LAYOUT_TRANSFER_SRC_OPTIMAL;
    vimb.srcQueueFamilyIndex = VK_QUEUE_FAMILY_IGNORED;
    vimb.dstQueueFamilyIndex = VK_QUEUE_FAMILY_IGNORED;
    vimb.image = stagingImage;
    vimb.sAccessMask = VK_ACCESS_HOST_WRITE_BIT;
    vimb.dAccessMask = 0;
    vimb.subresourceRange = visr;

    vkCmdPipelineBarrier( TextureCommandBuffer,
        VK_PIPELINE_STAGE_HOST_BIT, VK_PIPELINE_STAGE_HOST_BIT, 0,
        0, (VkMemoryBarrier *)NULLptr,
        0, (VkBuffersMemoryBarrier *)NULLptr,
        1, IN &vimb );
}

```

SIGGRAPH THINK BEYOND
2016 SIGGRAPH Asia

mb - July 24, 2020

238

```

// transition the texture buffer layout:
{
    VkImageSubresourceRange visr;
    visr.aspectMask = VK_IMAGE_ASPECT_COLOR_BIT;
    visr.baseMipLevel = 0;
    visr.levelCount = 1;
    visr.baseArrayLayer = 0;
    visr.layerCount = 1;

    VkImageMemoryBarrier vimb;
    vimb.sType = VK_STRUCTURE_TYPE_IMAGE_MEMORY_BARRIER;
    vimb.pNext = NULLptr;
    vimb.oldLayout = VK_IMAGE_LAYOUT_PREINITIALIZED;
    vimb.newLayout = VK_IMAGE_LAYOUT_TRANSFER_DST_OPTIMAL;
    vimb.srcQueueFamilyIndex = VK_QUEUE_FAMILY_IGNORED;
    vimb.dstQueueFamilyIndex = VK_QUEUE_FAMILY_IGNORED;
    vimb.image = textureImage;
    vimb.sAccessMask = 0;
    vimb.dAccessMask = VK_ACCESS_TRANSFER_WRITE_BIT;
    vimb.subresourceRange = visr;

    vkCmdPipelineBarrier( TextureCommandBuffer,
        VK_PIPELINE_STAGE_TOP_OF_PIPE_BIT, VK_PIPELINE_STAGE_TRANSFER_BIT, 0,
        0, (VkBuffersMemoryBarrier *)NULLptr,
        0, (VkBuffersMemoryBarrier *)NULLptr,
        1, IN &vimb );

    // now do the final image transfer:
    VkImageSubresourceLayers visrl;
    visrl.aspectMask = VK_IMAGE_ASPECT_COLOR_BIT;
    visrl.baseArrayLayer = 0;
    visrl.levelCount = 1;
    visrl.layerCount = 1;

    VkExtent3D vo3;
    vo3.x = 0;
    vo3.y = 0;
    vo3.z = 0;

    VkExtent3D ve3;
    ve3.width = texWidth;
    ve3.height = texHeight;
    ve3.depth = 1;
}

```

SIGGRAPH THINK BEYOND
2016 SIGGRAPH Asia

mb - July 24, 2020

239

```

VklImageCopy vic;
vic.srcSubresource = visl;
vic.srcOffset = vo3;
vic.dstSubresource = visl;
vic.dstOffset = ve3;
vic.extent = ve3;

vkCmdCopyImage( TextureCommandBuffer,
    stagingImage, VK_IMAGE_LAYOUT_TRANSFER_SRC_OPTIMAL,
    textureImage, VK_IMAGE_LAYOUT_TRANSFER_DST_OPTIMAL, 1, IN &vic );

```

SIGGRAPH THINK BEYOND
2016 SIGGRAPH Asia

mb - July 24, 2020

240

```

// transition the texture buffer layout a second time:
{
    VkImageSubresourceRange visr;
    visr.aspectMask = VK_IMAGE_ASPECT_COLOR_BIT;
    visr.baseMipLevel = 0;
    visr.levelCount = 1;
    visr.baseArrayLayer = 0;
    visr.layerCount = 1;

    VkImageMemoryBarrier vimb;
    vimb.sType = VK_STRUCTURE_TYPE_IMAGE_MEMORY_BARRIER;
    vimb.pNext = NULLptr;
    vimb.oldLayout = VK_IMAGE_LAYOUT_TRANSFER_DST_OPTIMAL;
    vimb.newLayout = VK_IMAGE_LAYOUT_SHADER_READ_ONLY_OPTIMAL;
    vimb.srcQueueFamilyIndex = VK_QUEUE_FAMILY_IGNORED;
    vimb.dstQueueFamilyIndex = VK_QUEUE_FAMILY_IGNORED;
    vimb.image = textureImage;
    vimb.sAccessMask = 0;
    vimb.dAccessMask = VK_ACCESS_SHADER_READ_BIT;
    vimb.subresourceRange = visr;

    vkCmdPipelineBarrier( TextureCommandBuffer,
        VK_PIPELINE_STAGE_TRANSFER_BIT, VK_PIPELINE_STAGE_FRAGMENT_SHADER_BIT, 0,
        0, (VkBuffersMemoryBarrier *)NULLptr,
        0, (VkBuffersMemoryBarrier *)NULLptr,
        1, IN &vimb );
}

result = vkEndCommandBuffer( TextureCommandBuffer );

VkSubmitInfo vsi;
vsi.sType = VK_STRUCTURE_TYPE_SUBMIT_INFO;
vsi.pNext = NULLptr;
vsi.commandBufferCount = 1;
vsi.pCommandBuffers = &textureCommandBuffer;
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 );

```

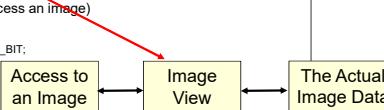
SIGGRAPH THINK BEYOND
2016 SIGGRAPH Asia

mb - July 24, 2020

241

```
// create an image view for the texture image:  
// (an "image view" is used to indirectly access an image)  
  
VkImageSubresourceRange  
visr;  
visr.aspectMask = VK_IMAGE_ASPECT_COLOR_BIT;  
visr.baseMipLevel = 0;  
visr.levelCount = 1;  
visr.baseArrayLayer = 0;  
visr.layerCount = 1;  
  
VkImageViewCreateInfo  
vivci.type = VK_STRUCTURE_TYPE_IMAGE_VIEW_CREATE_INFO;  
vivci.pNext = nullptr;  
vivci.flags = 0;  
vivci.image = textureImage;  
vivci.viewType = VK_IMAGE_VIEW_TYPE_2D;  
vivci.format = VK_FORMAT_R8G8B8A_UNORM; ← 8 bits Red 8 bits Green 8 bits Blue 8 bits Alpha  
vivci.components.r = VK_COMPONENT_SWIZZLE_R;  
vivci.components.g = VK_COMPONENT_SWIZZLE_G;  
vivci.components.b = VK_COMPONENT_SWIZZLE_B;  
vivci.components.a = VK_COMPONENT_SWIZZLE_A;  
vivci.subresourceRange = visr;  
  
result = vkCreateImageView( LogicalDevice, IN &vivci, PALLOCATOR, OUT &pMyTexture->texImageView );  
  
return result;  
}  
  
Note that, at this point, the Staging Buffer is no longer needed, and can be destroyed.
```

 mjb - July 24, 2020



242

Reading in a Texture from a BMP File

```
typedef struct MyTexture  
{  
    uint32_t width;  
    uint32_t height;  
    VkImage texImage;  
    VkImageView texImageView;  
    VkSampler texSampler;  
    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*.

 mjb - July 24, 2020

243

Vulkan.

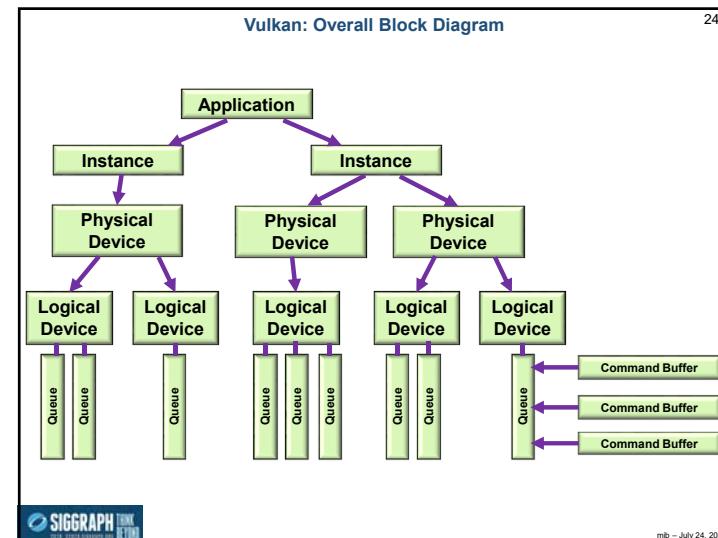
Queues and Command Buffers

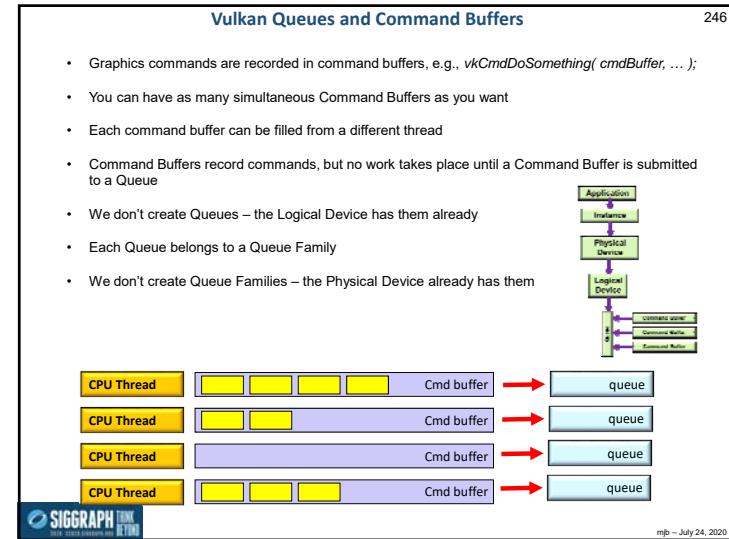
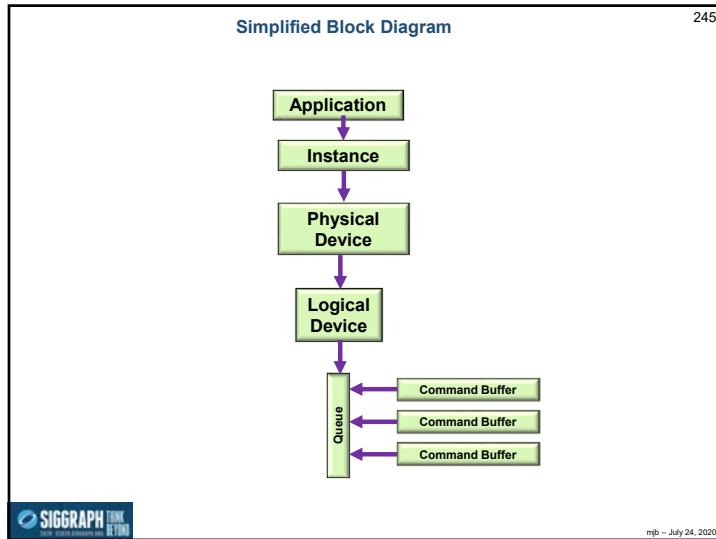
Mike Bailey
mjb@cs.oregonstate.edu


This work is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](#)

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

 mjb - July 24, 2020





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, "\t%ld: Queue Family Count = %d ; ", 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

The code demonstrates how to query the number of queue families available on a physical device and then iterate through them to print their properties. The output shows three queue families: one with 16 graphics, compute, and transfer capabilities, one with 1 transfer capability, and one with 8 compute capabilities.

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;
}
  
```

The code provides a function named `FindQueueFamilyThatDoesGraphics()` that iterates through the queue families of a physical device to find the first one that supports graphics operations. It uses the `vkGetPhysicalDeviceQueueFamilyProperties` function to get the properties of all queue families and then checks the `queueFlags` field for each family to determine if it supports graphics.

Creating a Logical Device Needs to Know Queue Family Information 249

```

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

VkDeviceQueueCreateInfo vdcqi[1];
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.queueCreateCount = 1; // # of device queues wanted
vdc.queueCreateInfo = &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.enabledFeatures = IN &PhysicalDeviceFeatures; // already created

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

VkQueue Queue;
uint32_t queueFamilyIndex = FindQueueFamilyThatDoesGraphics();
uint32_t queueIndex = 0;

result = vkGetDeviceQueue( 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 vcpcl;
    vcpcl.sType = VK_STRUCTURE_TYPE_COMMAND_POOL_CREATE_INFO;
    vcpcl.pNext = nullptr;
    vcpcl.flags = VK_COMMAND_POOL_CREATE_RESET_COMMAND_BUFFER_BIT | VK_COMMAND_POOL_CREATE_TRANSIENT_BIT;
#ifndef CHOICES
VK_COMMAND_POOL_CREATE_TRANSIENT_BIT
VK_COMMAND_POOL_CREATE_RESET_COMMAND_BUFFER_BIT
#endif
    vcpcl.queueFamilyIndex = FindQueueFamilyThatDoesGraphics();

    result = vkCreateCommandPool( LogicalDevice, IN &vcpcl, 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 vcbai;
        vcbai.sType = VK_STRUCTURE_TYPE_COMMAND_BUFFER_ALLOCATE_INFO;
        vcbai.pNext = nullptr;
        vcbai.commandPool = CommandPool;
        vcbai.level = VK_COMMAND_BUFFER_LEVEL_PRIMARY;
        vcbai.commandBufferCount = 2; // 2, because of double-buffering

        result = vkAllocateCommandBuffers( LogicalDevice, IN &vcbai, OUT &CommandBuffers[0] );
    }

    // allocate 1 command buffer for the transferring pixels from a staging buffer to a texture buffer:
    {
        VkCommandBufferAllocateInfo vcbai;
        vcbai.sType = VK_STRUCTURE_TYPE_COMMAND_BUFFER_ALLOCATE_INFO;
        vcbai.pNext = nullptr;
        vcbai.commandPool = CommandPool;
        vcbai.level = VK_COMMAND_BUFFER_LEVEL_PRIMARY;
        vcbai.commandBufferCount = 1;

        result = vkAllocateCommandBuffers( LogicalDevice, IN &vcbai, OUT &TextureCommandBuffer );
    }

    return result;
}

```

mb - July 24, 2020

Beginning a Command Buffer – One per Image 252

```

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 );

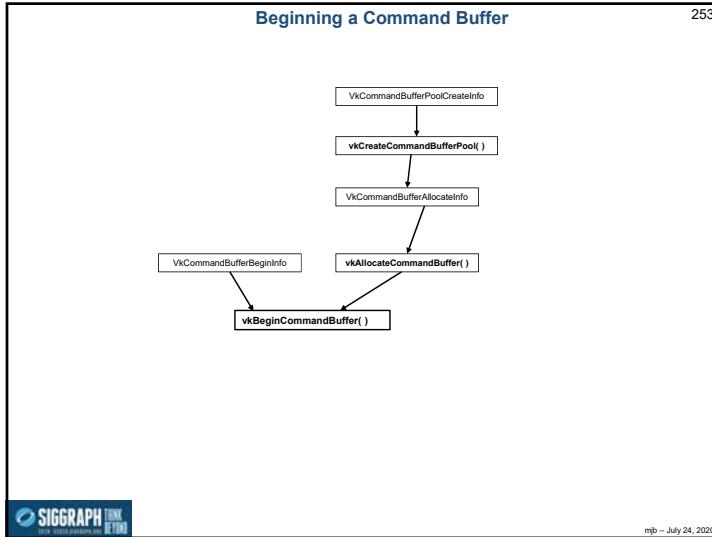
uint32_t nextImageIndex;
vkAcquireNextImageKHR( LogicalDevice, IN SwapChain, IN UINT64_MAX,
                      IN imageReadySemaphore, IN VK_NULL_HANDLE, OUT &nextImageIndex );

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 );
vkCmdBeginDescriptorSets( commandBuffer, pDynamicOffsets );
vkCmdBindImageTextures( commandBuffer, pImageTextures, indexType );
vkCmdBindPipeline( commandBuffer, pipeline );
vkCmdBindVertexBuffers( commandBuffer, firstBinding, bindingCount, const pOffsets );
vkCmdBindImage( commandBuffer, filter );
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, int32_t vertexOffset, firstInstance );
vkCmdDrawIndexedIndirect( commandBuffer, stride );
vkCmdDrawIndexedPrimitivesAMD( commandBuffer, stride );
vkCmdDrawIndirect( commandBuffer, stride );
vkCmdDrawIndirectCountAMD( commandBuffer, stride );
vkCmdEndQuery( commandBuffer, query );
vkCmdEndRenderPass( commandBuffer );
vkCmdExecuteCommands( commandBuffer, commandBufferCount, const pCommandBuffers );
  
```

mb - July 24, 2020

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

```

vkCmdFillBuffer( commandBuffer, dstBuffer, dstOffset, size, data );
vkCmdResetSubpass( commandBuffer, contents );
vkCmdPipelineBarrier( commandBuffer, dstStageMask, dependencyFlags, memoryBarrierCount, VkMemoryBarrier* pMemoryBarriers, bufferMemoryBarrierCount, pBufferMemoryBarriers, imageMemoryBarrierCount, pImageMemoryBarriers );
vkCmdProcessCommandsNVX( commandBuffer, pProcessCommandsInfo );
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, depthBiasScaleFactor, depthBiasClamp );
vkCmdSetDepthBounds( commandBuffer, depthBoundsMin, depthBoundsMax, maxDepthBounds );
vkCmdSetDeviceMaskKHR( commandBuffer, deviceMask );
vkCmdSetDiscardRectangleEXT( commandBuffer, firstDiscardRectangle, discardRectangleCount, pDiscardRectangles );
vkCmdSetEvent( commandBuffer, event, stageMask );
vkCmdSetLineWidth( commandBuffer, lineWidth );
vkCmdSetScissor( commandBuffer, firstScissor, scissorCount, pScissors );
vkCmdSetStencilRef( commandBuffer, stencilRef, stencilMask );
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 );
  
```

mb - July 24, 2020

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_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 );
  
```

mb - July 24, 2020

257

```

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

VkClearDepthStencilValue
vcdsv.depth = 1.0;
vcdsv.stencil = 0;

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

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

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

vkCmdBeginRenderPass( CommandBuffers[nextImageIndex], IN &vrpb, IN VK_SUBPASS_CONTENTS_INLINE );

```



mb - July 24, 2020

258

```

VkViewport viewport =
{
    0.0f, // x
    0.0f, // y
    (float)Width, // width
    (float)Height, // height
    0.0f, // minDepth
    1.0f // maxDepth
};

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

VkRect2D scissor =
{
    0.0f, // x
    0.0f, // y
    Width, // width
    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[] = { MyVertexDataBuffer.buffer };

VkDeviceSize offsets[1] = { 0 }; // 0, 1 = firstBinding, bindingCount

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] );

```

24.2020

259

Submitting a Command Buffer to a Queue for Execution

```

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

260

The Entire Submission / Wait / Display Process

```

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

vkCreateFence( LogicalDevice, IN &vfc, PALLOCATOR, OUT &renderFence );
result = VK_SUCCESS;

VkPipelineStageFlags waitAtBottom = VK_PIPELINE_STAGE_BOTTOM_OF_PIPE_BIT;
VkQueue presentQueue;
vkQueue presentQueue;
vkGetDeviceQueue( LogicalDevice, FindQueueFamilyThatDoesGraphics(), 0, OUT &presentQueue );
// 0 = queueIndex

VkSubmitInfo vsi;
vsi.sType = VK_STRUCTURE_TYPE_SUBMIT_INFO;
vsi.pNext = nullptr;
vsi.waitSemaphoreCount = 1;
vsi.pWaitSemaphores = &imageReadySemaphore;
vsi.signalSemaphoreCount = 1;
vsi.pSignalSemaphores = &SemaphoreRenderFinished;
vsi.commandBufferCount = 1;
vsi.pCommandBuffers = &CommandBuffers[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.plmIndices = &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.

SIGGRAPH THINK

mb - July 24, 2020

262

Vulkan.

The Swap Chain

Mike Bailey
mjb@cs.oregonstate.edu

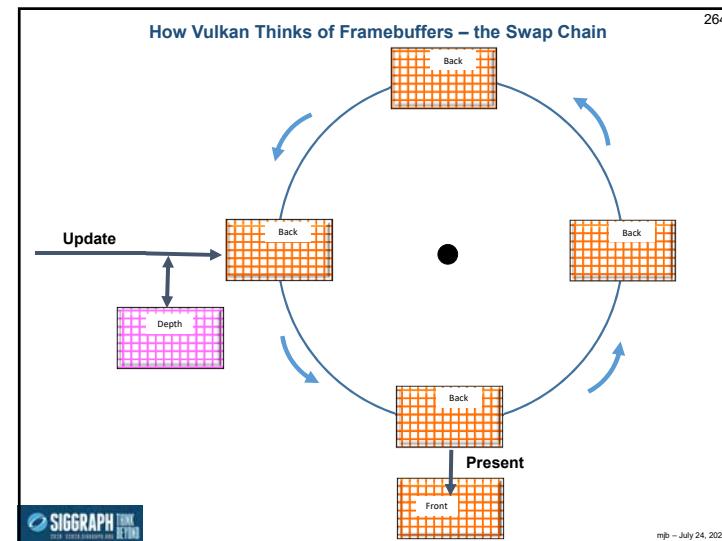
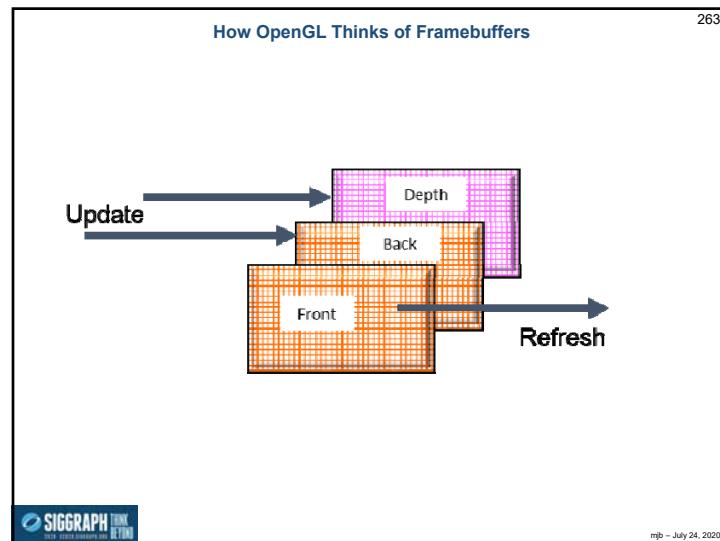


This work is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](http://creativecommons.org/licenses/by-nc-nd/4.0/)

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

SIGGRAPH THINK

mb - July 24, 2020



What is a Swap Chain? 265

Vulkan does not use the idea of a "back buffer". So, we need a place to render into before moving an image into place for viewing. This is called the **Swap Chain**.

In essence, the Swap Chain manages one or more image objects that form a sequence of images that can be drawn into and then given to the Surface to be presented to the user for viewing.

Swap Chains are arranged as a ring buffer 

Swap Chains are tightly coupled to the window system.

After creating the Swap Chain in the first place, the process for using the Swap Chain is:

1. Ask the Swap Chain for an image
2. Render into it via the Command Buffer and a Queue
3. Return the image to the Swap Chain for presentation
4. Present the image to the viewer (copy to "front buffer")



mb - July 24, 2020

We Need to Find Out What our Display Capabilities Are 266

```

VkSurfaceCapabilitiesKHR vsc;
vkGetPhysicalDeviceSurfaceCapabilitiesKHR(PhysicalDevice, Surface, OUT &vsc);
VkExtent2D surfaceRes = vsc.currentExtent;
fprintf( FpDebug, "\nvkGetPhysicalDeviceSurfaceCapabilitiesKHR:\n" );
...
VkBool32 supported;
result = vkGetPhysicalDeviceSurfaceSupportKHR(PhysicalDevice, FindQueueFamilyThatDoesGraphics( ), Surface, &supported);
if( supported == VK_TRUE )
    fprintf( FpDebug, "*** This Surface is supported by the Graphics Queue **\n" );
...
uint32_t formatCount;
vkGetPhysicalDeviceSurfaceFormatsKHR(PhysicalDevice, Surface, &formatCount, (VkSurfaceFormatKHR *) nullptr);
VkSurfaceFormatKHR *surfaceFormats = new VkSurfaceFormatKHR[formatCount];
vkGetPhysicalDeviceSurfaceFormatsKHR(PhysicalDevice, Surface, &formatCount, surfaceFormats);
fprintf( FpDebug, "\nFound %d Surface Formats:\n", formatCount );
...
uint32_t presentModeCount;
vkGetPhysicalDeviceSurfacePresentModesKHR(PhysicalDevice, Surface, &presentModeCount, (VkPresentModeKHR *) nullptr);
VkPresentModeKHR *presentModes = new VkPresentModeKHR[presentModeCount];
vkGetPhysicalDeviceSurfacePresentModesKHR(PhysicalDevice, Surface, &presentModeCount, presentModes);
fprintf( FpDebug, "\nFound %d Present Modes:\n", presentModeCount );
...

```

mb - July 24, 2020

We Need to Find Out What our Display Capabilities Are 267

VulkanDebug.txt output:

```

vkGetPhysicalDeviceSurfaceCapabilitiesKHR:
minImageCount = 2 : maxImageCount = 8
currentExtent = 1024 x 1024
minImageExtent = 1024 x 1024
maxImageExtent = 1024 x 1024
maxImageArrayLayers = 1
supportedTransforms = 0x0001
currentTransform = 0x0001
supportedCompositeAlpha = 0x0001
supportedUsageFlags = 0x009f

** This Surface is supported by the Graphics Queue **

Found 2 Surface Formats:
0: 44 0 ( VK_FORMAT_B8G8R8A8_UNORM, VK_COLOR_SPACE_SRGB_NONLINEAR_KHR )
1: 50 0 ( VK_FORMAT_B8G8R8A8_SRGB, VK_COLOR_SPACE_SRGB_NONLINEAR_KHR )

Found 3 Present Modes:
0: 2 ( VK_PRESENT_MODE_FIFO_KHR )
1: 3 ( VK_PRESENT_MODE_FIFO_RELAXED_KHR )
2: 1 ( VK_PRESENT_MODE_MAILBOX_KHR )

```

mb - July 24, 2020

Creating a Swap Chain 268

```

vkGetDevicePhysicalSurfaceCapabilities()
    |
    +--> VkSurfaceCapabilities
        |
        +--> VkSwapchainCreateInfo
            |
            +--> vkCreateSwapchain()
                |
                +--> vkGetSwapChainImages()
                    |
                    +--> vkCreateImageView()

```

mb - July 24, 2020

Creating a Swap Chain 269

```

VkSurfaceCapabilitiesKHR vsc;
vkGetPhysicalDeviceSurfaceCapabilitiesKHR( PhysicalDevice, Surface, OUT &vsc );
VkExtent2D surfaceRes = vsc.currentExtent;

VkSwapchainCreateInfoKHR vscci;
vscci.sType = VK_STRUCTURE_TYPE_SWAPCHAIN_CREATE_INFO_KHR;
vscci.pNext = nullptr;
vscci.flags = 0;
vscci.surface = Surface;
vscci.minImageCount = 2; // double buffering
vscci.imageFormat = VK_FORMAT_B8G8R8A8_UNORM;
vscci.imageColorSpace = VK_COLORSPACE_SRGB_NONLINEAR_KHR;
vscci.imageExtent.width = surfaceRes.width;
vscci.imageExtent.height = surfaceRes.height;
vscci.imageUsage = VK_IMAGE_USAGE_COLOR_ATTACHMENT_BIT;
vscci.preTransform = VK_SURFACE_TRANSFORM_IDENTITY_BIT_KHR;
vscci.compositeAlpha = VK_COMPOSITE_ALPHA_OPAQUE_BIT_KHR;
vscci.imageArrayLayers = 1;
vscci.imageSharingMode = VK_SHARING_MODE_EXCLUSIVE;
vscci.queueFamilyIndexCount = 0;
vscci.pQueueFamilyIndices = (const uint32_t *)nullptr;
vscci.presentMode = VK_PRESENT_MODE_MAILBOX_KHR;
vscci.oldSwapchain = VK_NULL_HANDLE;
vscci.clipped = VK_TRUE;

result = vkCreateSwapchainKHR( LogicalDevice, IN &vscci, PALLOCATOR, OUT &SwapChain );

```

SIGGRAPH THINK
BEYOND

mb - July 24, 2020

Creating the Swap Chain Images and Image Views 270

```

uint32_t imageCount; // # of display buffers - 2? 3?
result = vkGetSwapchainImagesKHR( LogicalDevice, IN SwapChain, OUT &imageCount, (VkImage *)nullptr );

PresentImages = new VkImage[ imageCount ];
result = vkGetSwapchainImagesKHR( LogicalDevice, SwapChain, OUT &imageCount, PresentImages );

// present views for the double-buffering:
PresentImageViews = new VkImageView[ imageCount ];

for( unsigned int i = 0; i < imageCount; i++ )
{
    VkImageViewCreateInfo vivci;
    vivci.sType = VK_STRUCTURE_TYPE_IMAGE_VIEW_CREATE_INFO;
    vivci.pNext = nullptr;
    vivci.flags = 0;
    vivci.viewType = VK_IMAGE_VIEW_TYPE_2D;
    vivci.format = VK_FORMAT_B8G8R8A8_UNORM;
    vivci.components.r = VK_COMPONENT_SWIZZLE_R;
    vivci.components.g = VK_COMPONENT_SWIZZLE_G;
    vivci.components.b = VK_COMPONENT_SWIZZLE_B;
    vivci.components.a = VK_COMPONENT_SWIZZLE_A;
    vivci.subresourceRange.aspectMask = VK_IMAGE_ASPECT_COLOR_BIT;
    vivci.subresourceRange.baseMipLevel = 0;
    vivci.subresourceRange.levelCount = 1;
    vivci.subresourceRange.baseArrayLayer = 0;
    vivci.subresourceRange.layerCount = 1;
    vivci.image = PresentImages[ i ];

    result = vkCreateImageView( LogicalDevice, IN &vivci, PALLOCATOR, OUT &PresentImageViews[ i ] );
}

```

SIGGRAPH THINK
BEYOND

mb - July 24, 2020

Rendering into the Swap Chain, I 271

```

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;
uint64_t timeout = UINT64_MAX;
vkAcquireNextImageKHR( LogicalDevice, IN SwapChain, IN timeout, IN imageReadySemaphore,
                      IN VK_NULL_HANDLE, OUT &nextImageIndex );
...

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

vkCmdBeginRenderPass( CommandBuffers[ nextImageIndex ], IN &vrpb,
                      IN VK_SUBPASS_CONTENTS_INLINE );
vkCmdBindPipeline( CommandBuffers[ nextImageIndex ], VK_PIPELINE_BIND_POINT_GRAPHICS, GraphicsPipeline );
...

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

```

SIGGRAPH THINK
BEYOND

mb - July 24, 2020

Rendering into the Swap Chain, II 272

```

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

VkFence renderFence;
vkCreateFence( LogicalDevice, &vfc, PALLOCATOR, OUT &renderFence );

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 = &waitForBottom;
vsi.commandBufferCount = 1;
vsi.pCommandBuffers = &CommandBuffers[ nextImageIndex ];
vsi.signalSemaphoreCount = 0;
vsi.pSignalSemaphores = &SemaphoreRenderFinished;

result = vkQueueSubmit( presentQueue, 1, IN &vsi, IN renderFence ); // 1 = submitCount

```

SIGGRAPH THINK
BEYOND

mb - July 24, 2020

Rendering into the Swap Chain, III 273

```

result = vkWaitForFences( LogicalDevice, 1, IN &renderFence, VK_TRUE, UINT64_MAX );

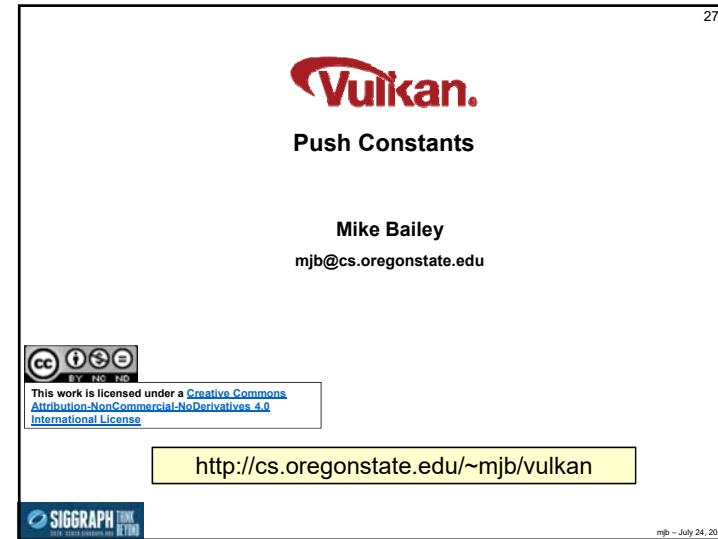
VkPresentInfoKHR
{
    vpi;
    vpi.sType = VK_STRUCTURE_TYPE_PRESENT_INFO_KHR;
    vpi.pNext = nullptr;
    vpi.waitSemaphoreCount = 0;
    vpi.pWaitSemaphores = (VkSemaphore *)nullptr;
    vpi.swapchainCount = 1;
    vpi.pSwapchains = &SwapChain;
    vpi.pImageIndices = &nextImageIndex;
    vpi.pResults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

result = vkQueuePresentKHR( presentQueue, IN &vpi );

```

SIGGRAPH THINK
CREATE. DESIGN. EXPERIMENT. REFINEMENT.

mjb - July 24, 2020



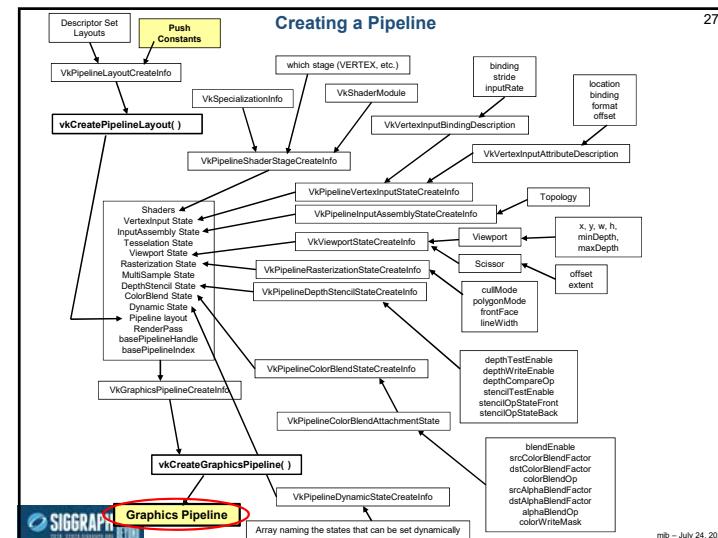
Push Constants 275

In an effort to expand flexibility and retain efficiency, Vulkan provides something called **Push Constants**. Like the name implies, these let you “push” constant values out to the shaders. These are typically used for small, frequently-updated data values. This is good, since Vulkan, at times, makes it cumbersome to send changes to the graphics.

By “small”, Vulkan specifies that these must be at least 128 bytes in size, although they can be larger. For example, the maximum size is 256 bytes on the NVIDIA 1080ti. (You can query this limit by looking at the **maxPushConstantSize** parameter in the **VkPhysicalDeviceLimits** structure.) Unlike uniform buffers and vertex buffers, these are not backed by memory. They are actually part of the Vulkan pipeline.

SIGGRAPH THINK
CREATE. DESIGN. EXPERIMENT. REFINEMENT.

mjb - July 24, 2020



Push Constants 277

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`.

 mjb - July 24, 2020

Setting up the Push Constants for the Pipeline Structure 278

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 );
```

 mjb - July 24, 2020

An Robotic Example using Push Constants 279

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

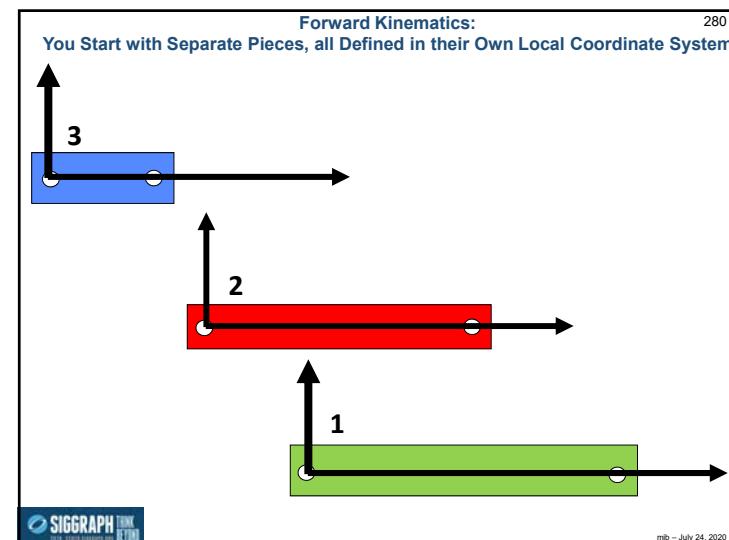


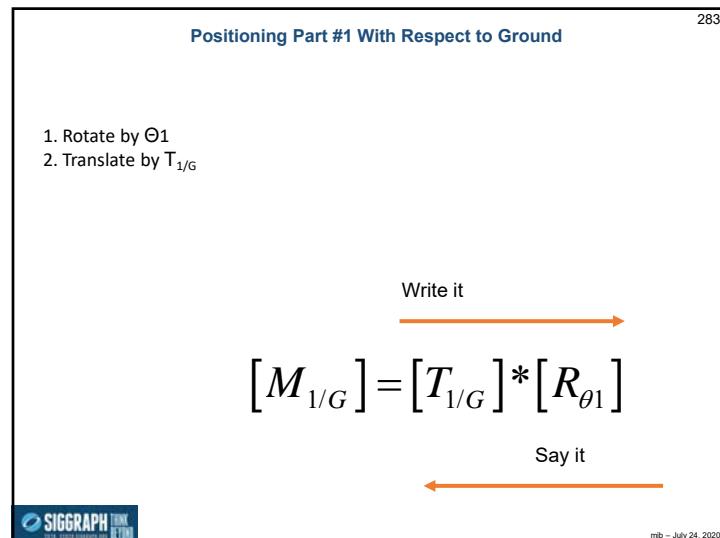
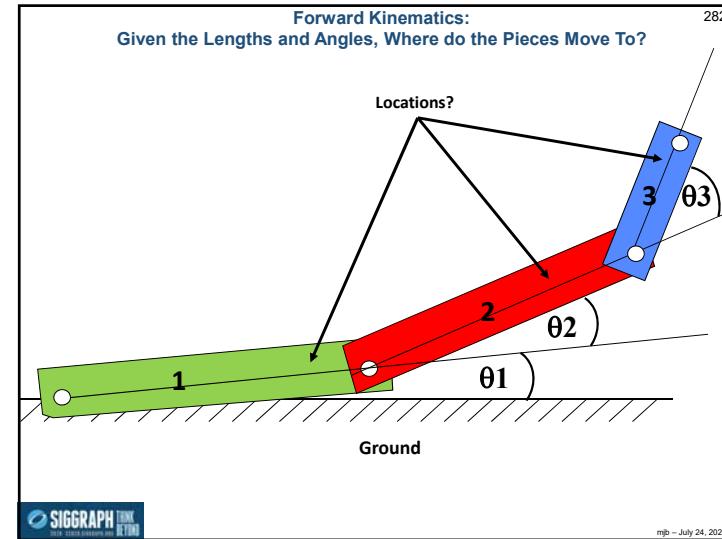
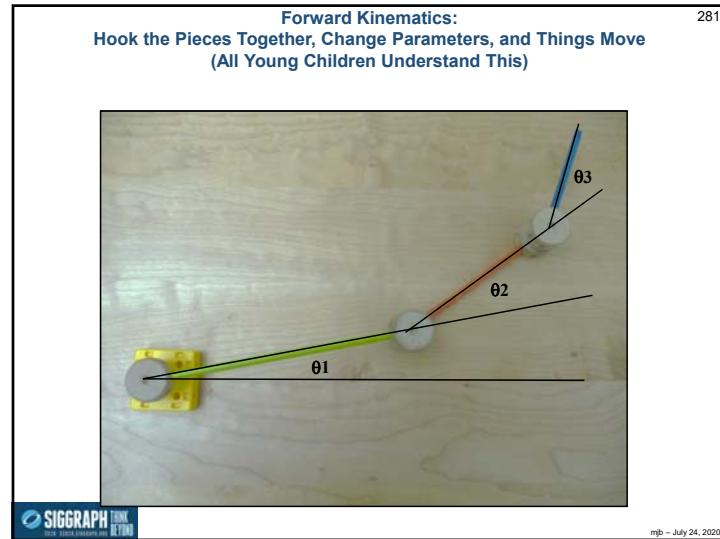
Where each arm is represented by:

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

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

 mjb - July 24, 2020





Why Do We Say it Right-to-Left?

284

$M_{A/B}I = [R_{A/B}I] * [T_{B/I}]$

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

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

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

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

SIGGRAPH THINK

Positioning Part #2 With Respect to Ground 285

1. Rotate by Θ_2
2. Translate the length of part 1
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
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 );

```

 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

```



mb - July 24, 2020

In the *RenderScene* Function

290

```

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

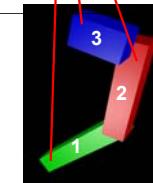
vkCmdBindVertexBuffers( CommandBuffers[nextImageIndex], GraphicsPipelineLayout,
VK_SHADER_STAGE_ALL, 0, sizeof(struct arm), (void *)&Arm1 );
vkCmdDraw( CommandBuffers[nextImageIndex], vertexCount, instanceCount, firstVertex, firstInstance );

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



mb - July 24, 2020

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;
...

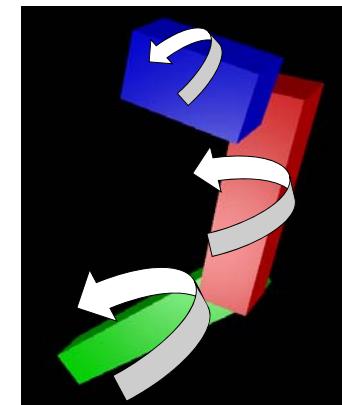
vec3 bVertex = aVertex; // arm coordinate system is [-1., 1.] in X
bVertexx += 1.; // now is [0., 2]
bVertexx /= 2.; // now is [0., 1]
bVertexx *= (RobotArm.armScale); // now is [0., RobotArm.armScale]
bVertex = vec3( RobotArm.armMatrix * vec4( bVertex, 1. ) );

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

```

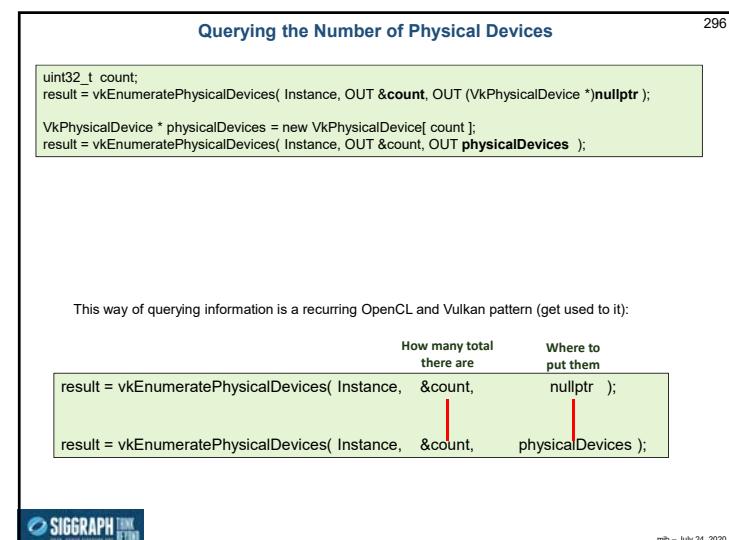
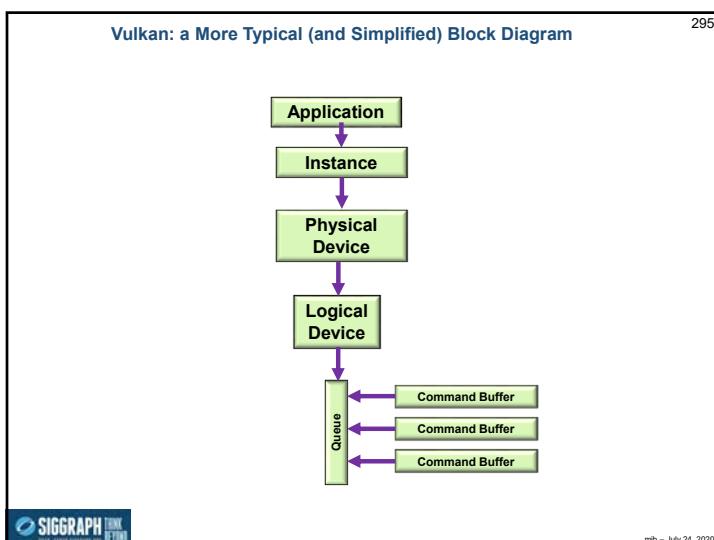
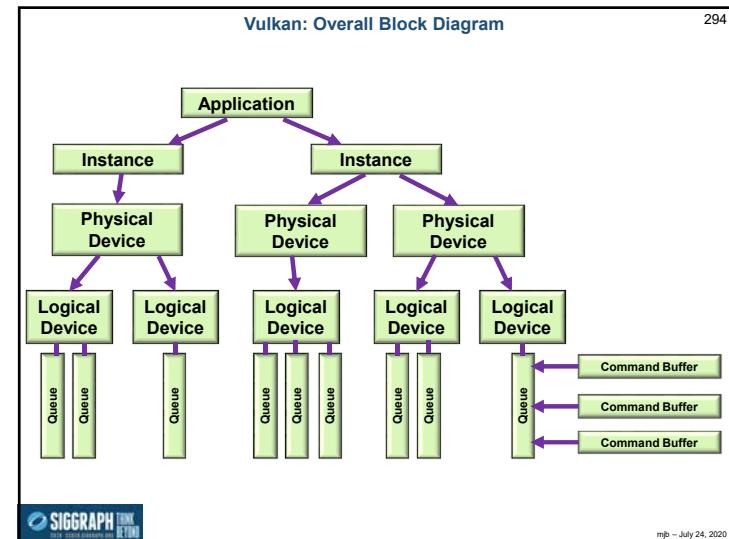
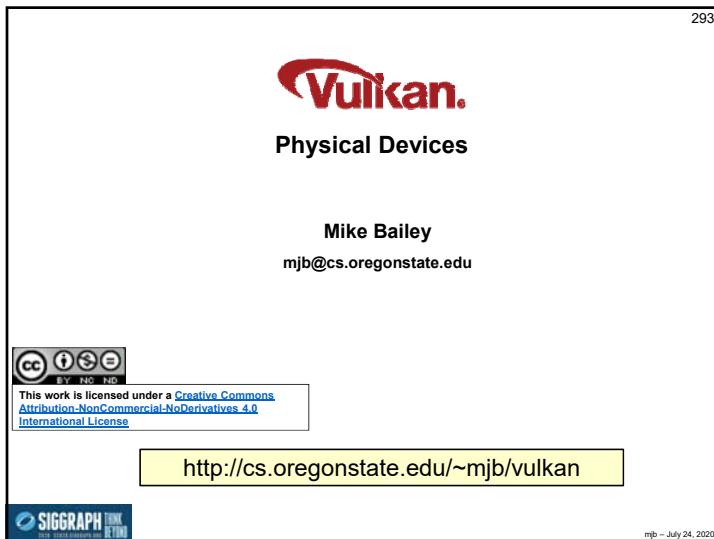


mb - July 24, 2020



292

mb - July 24, 2020



Vulkan: Identifying the Physical Devices

297

```
VkResult result = VK_SUCCESS;

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

printf(FpDebug, "\n%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;
}
```



mb - July 24, 2020

Which Physical Device to Use, I

298

```
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;
    }

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



mb - July 24, 2020

Which Physical Device to Use, II

299

```
// need some logic 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;
}
```



mb - July 24, 2020

Asking About the Physical Device's Features

300

```
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 );
```



mb - July 24, 2020

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 vpdm;
vkGetPhysicalDeviceMemoryProperties(PhysicalDevice, OUT &vpdm);

fprintf(FpDebug, "%n%d Memory Types:\n", vpdm.memoryTypeCount);
for( unsigned int i = 0; i < vpdm.memoryTypeCount; i++ )
{
    VkMemoryType vmt = vpdm.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, "%n%d Memory Heaps:\n", vpdm.memoryHeapCount);
for( unsigned int i = 0; i < vpdm.memoryHeapCount; i++ )
{
    fprintf(FpDebug, "Heap %d: ", i);
    VKMemoryHeap vmh = vpdm.memoryHeaps[i];
    fprintf(FpDebug, "size = 0x%08x", (unsigned long int)vmh.size );
    if( ( vmh.flags & VK_MEMORY_HEAP_DEVICE_LOCAL_BIT )!=0 ) fprintf(FpDebug, " DeviceLocal" );
    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, "Found %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, "%t%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");
}

```



mjb - 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



mjb - July 24, 2020



Logical Devices

Mike Bailey

mjb@cs.oregonstate.edu



This work is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](#)

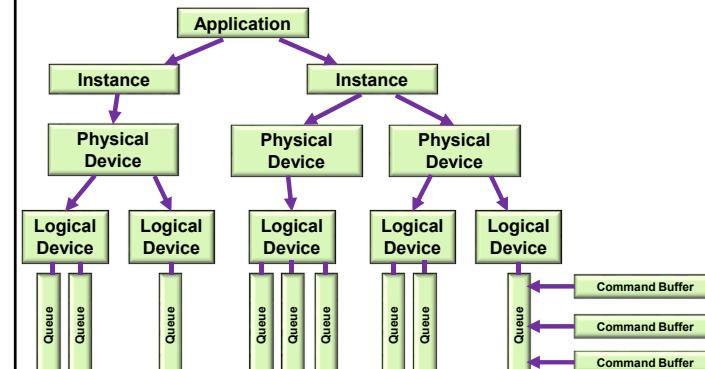
<http://cs.oregonstate.edu/~mjb/vulkan>



mjb - July 24, 2020

Vulkan: Overall Block Diagram

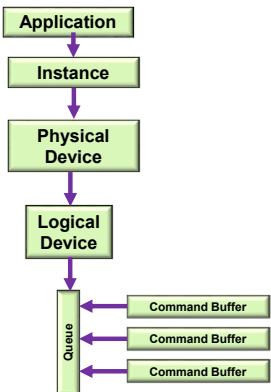
308



mjb - July 24, 2020

Vulkan: a More Typical (and Simplified) Block Diagram

309



mb - 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_swapchains"
};

// see what device layers are available:
uint32_t layerCount;
vkEnumerateDeviceLayerProperties(PhysicalDevice, &layerCount, (VkLayerProperties *)nullptr);

VkLayerProperties * deviceLayers = new VkLayerProperties[layerCount];

result = vkEnumerateDeviceLayerProperties(PhysicalDevice, &layerCount, deviceLayers);
  
```



mb - July 24, 2020

Looking to See What Device Extensions are Available

311

```

// see what device extensions are available:
uint32_t extensionCount;
vkEnumerateDeviceExtensionProperties(PhysicalDevice, deviceLayers[i].layerName,
                                    &extensionCount, (VkExtensionProperties *)nullptr);

VkExtensionProperties * deviceExtensions = new VkExtensionProperties[extensionCount];

result = vkEnumerateDeviceExtensionProperties(PhysicalDevice, deviceLayers[i].layerName,
                                              &extensionCount, deviceExtensions);
  
```



mb - 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'
  
```



mb - July 24, 2020

Vulkan: Creating a Logical Device 313

```

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 vdcii;
vdcii.sType = VK_STRUCTURE_TYPE_DEVICE_CREATE_INFO;
vdcii.pNext = nullptr;
vdcii.flags = 0;
vdcii.queueCreateInfoCount = 1;           // # of device queues
vdcii.pQueueCreateInfos = &vdqci;        // array of VkDeviceQueueCreateInfo's
vdcii.enabledLayerCount = 0;
vdcii.ppEnabledLayerNames = myDeviceLayers;
vdcii.enabledExtensionCount = 0;
vdcii.ppEnabledExtensionNames = (const char**)nullptr;      // no extensions
vdcii.enabledExtensionCount = sizeof(myDeviceExtensions) / sizeof(char *);
vdcii.ppEnabledExtensionNames = myDeviceExtensions;
vdcii.eEnabledFeatures = IN &PhysicalDeviceFeatures;

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

```

mjb – July 24, 2020

Vulkan: Creating the Logical Device's Queue 314

```

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

```

SIGGRAPH 2020 SIGGRAPH.org BEYOND

mjb – July 24, 2020

Vulkan.

Dynamic State Variables

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 2020 SIGGRAPH.org BEYOND

mjb – July 24, 2020

Creating a Pipeline with Dynamically Changeable State Variables 316

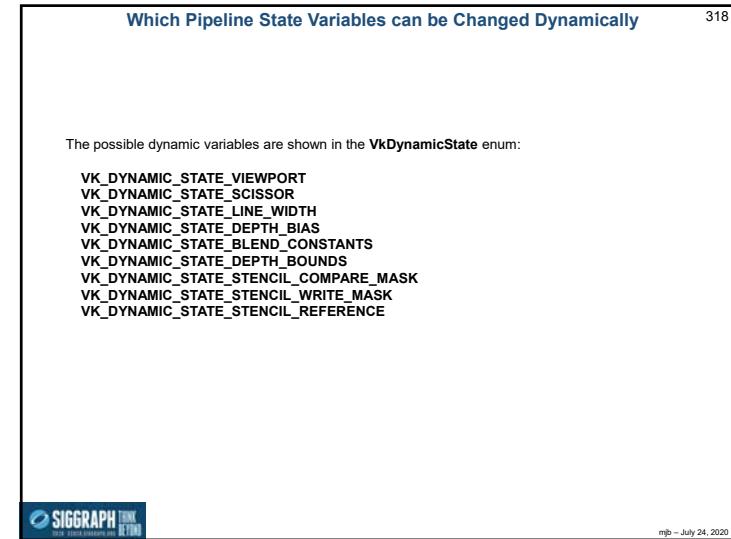
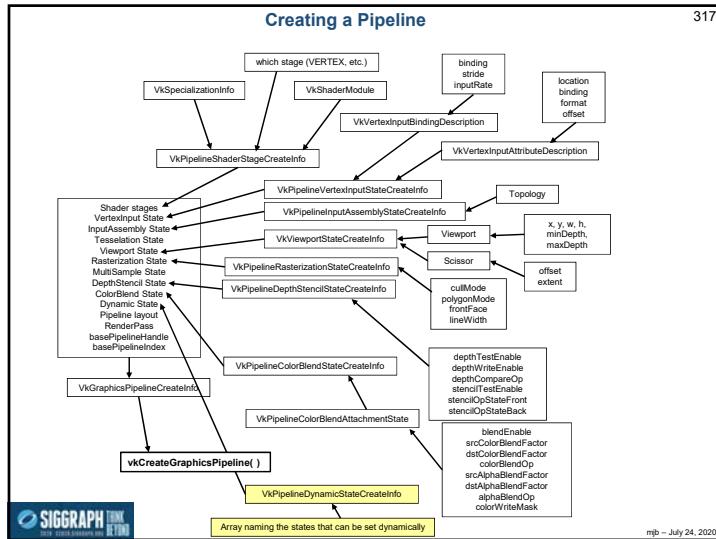
The graphics pipeline data structure is full of state information, and, as previously-discussed, is largely immutable, that is, the information contained inside it is fixed, and can only be changed by creating a new graphics pipeline data structure with new information.

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.

SIGGRAPH 2020 SIGGRAPH.org BEYOND

mjb – July 24, 2020



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**.

SIGGRAPH THINK
CREATE. DESIGN. INNOVATE. REINVENT.

mb - July 24, 2020

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, pScissiors );
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 );

```

SIGGRAPH THINK
CREATE. DESIGN. INNOVATE. REINVENT.

mb - July 24, 2020

321



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](#)

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

SIGGRAPH 2020

mjb - July 24, 2020

322

Setting up Query Pools

- 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

SIGGRAPH 2020

mjb - July 24, 2020

323

Setting up Query Pools

```

VkQueryPoolCreateInfo
    vqpci; // circled
    vqpci.sType = VK_STRUCTURE_TYPE_QUERY_POOL_CREATE_INFO;
    vqpci.pNext = nullptr;
    vqpci.flags = 0;
    vqpci.queryType = << one of: >>
        VK_QUERY_TYPE_OCCULSION
        VK_QUERY_TYPE_PIPELINE_STATISTICS
        VK_QUERY_TYPE_TIMESTAMP
    vqpci.queryCount = 1;
    vqpci.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_PRIMITIVES_BIT
VK_QUERY_PIPELINE_STATISTIC_VERTEX_SHADER_INVOCATIONS_BIT
VK_QUERY_PIPELINE_STATISTIC_GEOMETRY_SHADER_INVOCATIONS_BIT
VK_QUERY_PIPELINE_STATISTIC_FRAGMENT_SHADER_INVOCATIONS_BIT
VK_QUERY_PIPELINE_STATISTIC_CLIPPING_PRIMITIVES_BIT
VK_QUERY_PIPELINE_STATISTIC_CLIPPING_PRIMITIVES_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 &vqpci, PALLOCATOR, OUT &occlusionQueryPool );

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

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

```

mjb - July 24, 2020

324

Resetting, Filling, and Examining a Query Pool

```

vkCmdResetQueryPool( CommandBuffer, occlusionQueryPool, 0, 1 );

vkCmdBeginQuery( CommandBuffer, occlusionQueryPool, 0, VK_QUERY_CONTROL_PRECISE_BIT );
    ...
vkCmdEndQuery( CommandBuffer, occlusionQueryPool, 0 ); // query index number

#define DATASIZE 128
uint32_t data[DATASIZE];

result = vkGetQueryPoolResults( LogicalDevice, occlusionQueryPool, 0, 1, DATASIZE*sizeof(uint32_t), data, stride, flags );
// or'd 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

```

SIGGRAPH 2020

mjb - July 24, 2020

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 );
```

 mjb - 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 );

// vqpc.pipelineStatistics = or'd 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_FRAGMENT_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
```

 mjb - 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 );
```

 mjb - 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_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
```

 mjb - July 24, 2020

329

Vulkan.

Compute Shaders

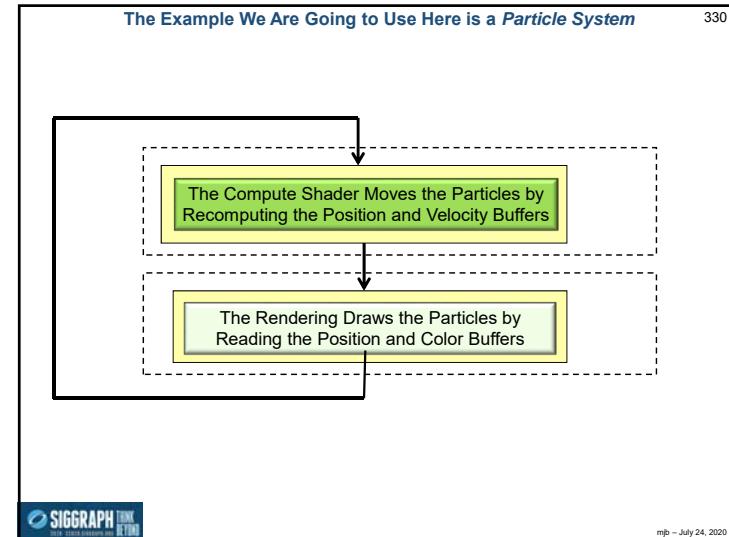
Mike Bailey
mjb@cs.oregonstate.edu

This work is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](#).

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

SIGGRAPH TALK

mjb - July 24, 2020



The Data in your C/C++ Program will look like This 331

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.
```

SIGGRAPH TALK

mjb - July 24, 2020

The Data in your Compute Shader will look like This 332

```
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
};

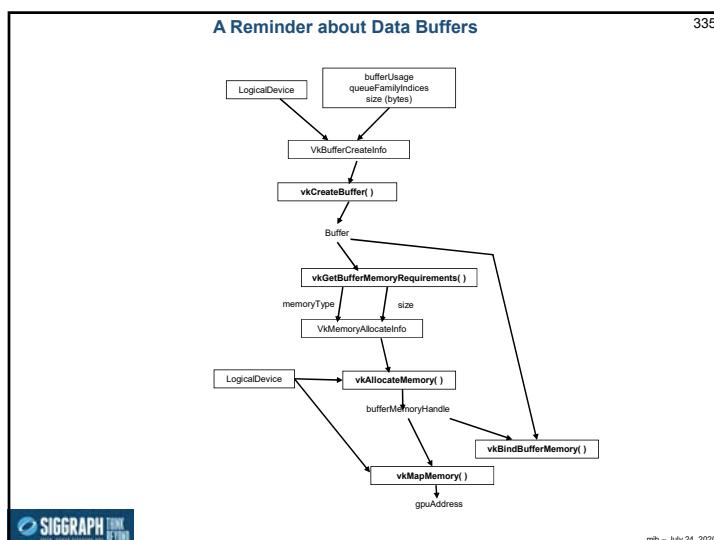
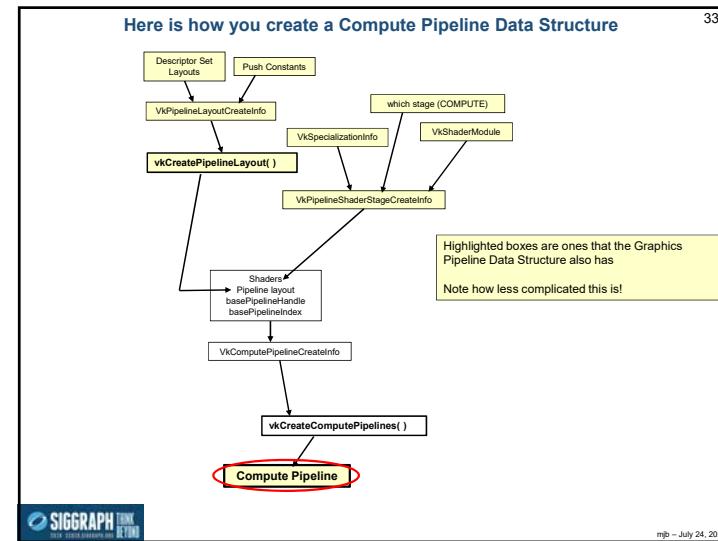
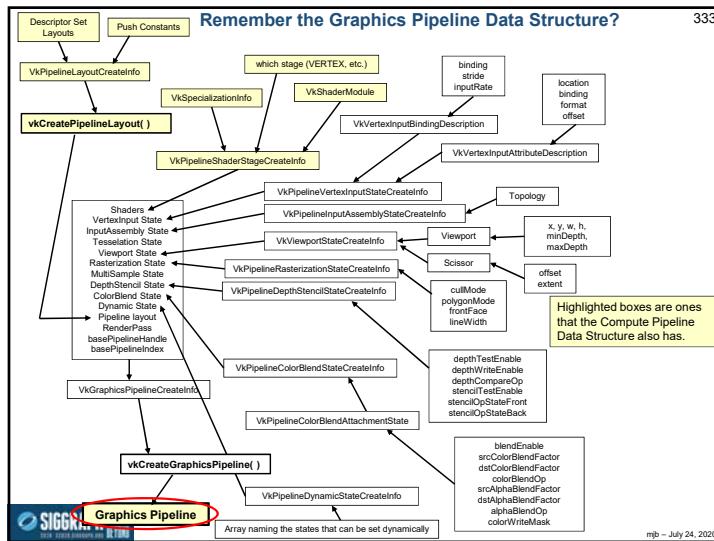
layout( std140, set = 0, binding = 2 ) buffer Col
{
    vec4 Colors[ ]; // array of structures
};
```

1 2 3

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.

SIGGRAPH TALK

mjb - July 24, 2020



```
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_STORAGE_BUFFER_BIT;
vbcI.sharingMode = VK_SHARING_MODE_EXCLUSIVE;
vbcI.queueFamilyIndexCount = 0;
vbcI.pQueueFamilyIndices = (const int32_t*)nullptr;

result = vkCreateBuffer( LogicalDevice, IN &vbcI, PALLOCATOR, OUT &
```

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
```



mb - July 24, 2020

Create the Compute Pipeline Layout

338

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



mb - July 24, 2020

Create the Compute Pipeline Layout

339

```
VkPipelineLayout ComputePipelineLayout;
VkDescriptorSetLayout ComputeSetLayout;
...
result = vkCreateDescriptorSetLayout( LogicalDevice, IN &ComputeSetLayout, OUT &ComputePipelineLayout );
VkPipelineLayoutCreateInfo vplci;
vplci.sType = VK_STRUCTURE_TYPE_PIPELINE_LAYOUT_CREATE_INFO;
vplci.pNext = nullptr;
vplci.flags = 0;
vplci.setLayoutCount = 1;
vplci.pSetLayouts = ComputeSetLayout;
vplci.pushConstantRangeCount = 0;
vplci.pPushConstantRanges = (VkPushConstantRange *)nullptr;
result = vkCreatePipelineLayout( LogicalDevice, IN &vplci, PALLOCATOR, OUT &ComputePipelineLayout );
```



mb - July 24, 2020

Create the Compute Pipeline

340

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

VkComputePipelineCreateInfo vcpcl[1];
vcpcl[0].sType = VK_STRUCTURE_TYPE_COMPUTE_PIPELINE_CREATE_INFO;
vcpcl[0].pNext = nullptr;
vcpcl[0].flags = 0;
vcpcl[0].stage = vpsc;
vcpcl[0].layout = ComputePipelineLayout;
vcpcl[0].basePipelineHandle = VK_NULL_HANDLE;
vcpcl[0].basePipelineIndex = 0;
result = vkCreateComputePipelines( LogicalDevice, VK_NULL_HANDLE, 1, &vcpcl[0], PALLOCATOR, &ComputePipeline );
```



mb - July 24, 2020

Creating a Vulkan Data Buffer 341

```

VkBuffer Buffer;
VkBufferCreateInfo vbc;
vbc.sType = VK_STRUCTURE_TYPE_BUFFER_CREATE_INFO;
vbc.pNext = nullptr;
vbc.flags = 0;
vbc.size = NUM_PARTICLES * sizeof( glm::vec4 );
vbc.usage = VK_USAGE_STORAGE_BUFFER_BIT;
vbc.sharingMode = VK_SHARING_MODE_CONCURRENT;
vbc.queueFamilyIndexCount = 0;
vbc.pQueueFamilyIndices = (const int32_t*) nullptr;
result = vkCreateBuffer( LogicalDevice, &vbc, PALLOCATOR, &posBuffer );

```

 mb - July 24, 2020

Allocating Memory and Binding the Buffer 342

```

VkMemoryRequirements vmr;
result = vkGetBufferMemoryRequirements( LogicalDevice, posBuffer, &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, &vmai, PALLOCATOR, &vdm );
result = vkBindBufferMemory( LogicalDevice, posBuffer, IN vdm, 0 ); // 0 is the offset

MyBuffer myPosBuffer;
myPosBuffer.size = vbc.size;
myPosBuffer.buffer = PosBuffer;
myPosBuffer.vdm = vdm;

```

 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 = 1.0;
}
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( V ymin, V ymax );
    velocities[i].z = Ranf( Vzmin, Vzmax );
    velocities[i].w = 0.0;
}
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, 1.0 );
    colors[i].g = Ranf( .3f, 1.0 );
    colors[i].b = Ranf( .3f, 1.0 );
    colors[i].a = 1.0;
}
vkUnmapMemory( LogicalDevice, IN myColBuffer.vdm );

```

 mb - July 24, 2020

Fill the Buffers 344

```

#include <stdlib.h>

#define TOP 2147483647.0 // 2^31 - 1

float Ranf( float low, float high )
{
    long random(); // returns integer 0 - TOP
    float r = (float)random() / (float)RAND_MAX;
    return low + r * (high - low) / (float)RAND_MAX;
}

```

 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
};

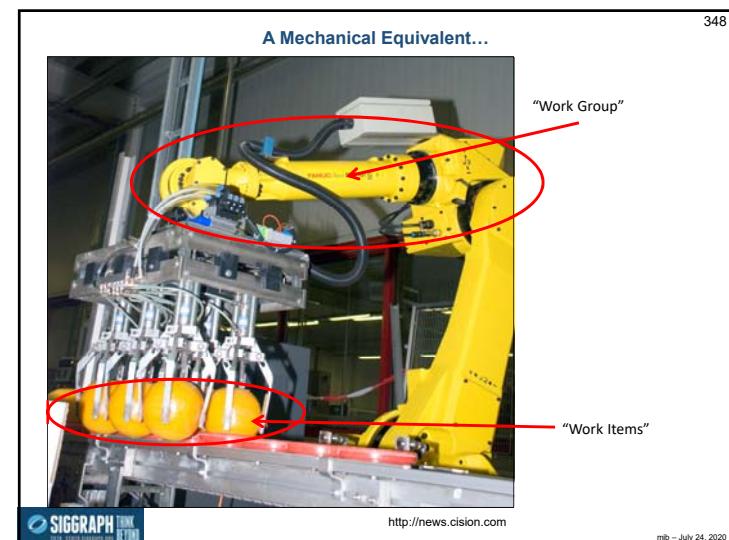
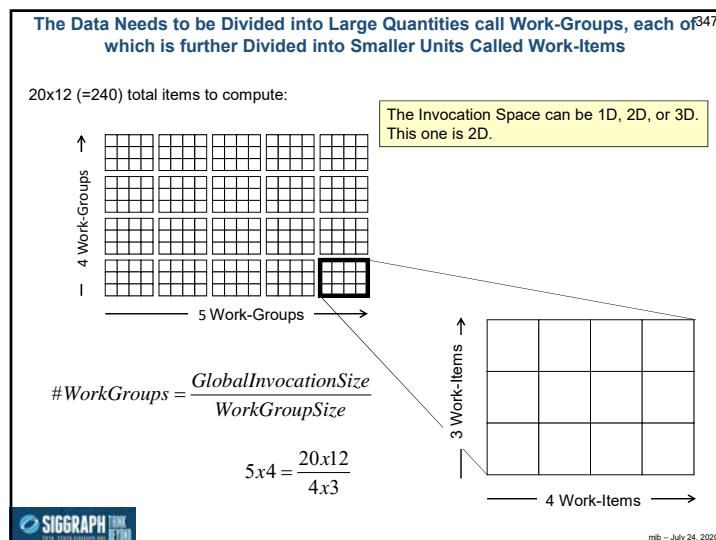
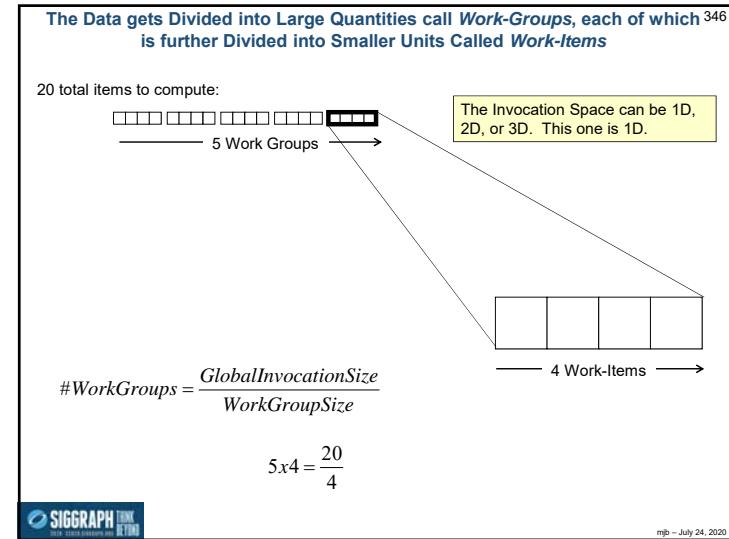
layout( std140, set = 0, binding = 2 ) buffer Col
{
    vec4 Colors[ ];             // array of structures
};

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 THINK
CREATE. DESIGN. INNOVATE. BEYOND



The Particle System Compute Shader – The Physics

349

```
#define POINT      vec3
#define VELOCITY   vec3
#define VECTOR     vec3
#define SPHERE     vec4      // xc, yc, zc, r
#define PLANE      vec4      // 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;
```

$$p' = p + v \cdot t + \frac{1}{2} G \cdot t^2$$

$$v' = v + G \cdot t$$



mb - July 24, 2020

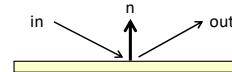
The Particle System Compute Shader – How About Introducing a Bounce?

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 );
}
```



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)



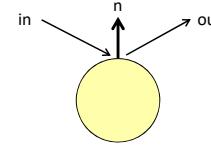
mb - July 24, 2020

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 );
}
```



mb - 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...



mb - 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:

```
layout( local_size_x = 64, local_size_y = 1, local_size_z = 1 ) in;
```



mb - July 24, 2020

Displaying the Particles

354

```
VkVertexInputBindingDescription vvid[3]; // one of these per buffer data buffer
vvid[0].binding = 0; // which binding # this is
vvid[0].stride = sizeof( struct pos ); // bytes between successive structs
vvid[0].inputRate = VK_VERTEX_INPUT_RATE_VERTEX;

vvid[1].binding = 1;
vvid[1].stride = sizeof( struct vel );
vvid[1].inputRate = VK_VERTEX_INPUT_RATE_VERTEX;

vvid[2].binding = 2;
vvid[2].stride = sizeof( struct col );
vvid[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;
```



mb - July 24, 2020

Displaying the Particles

355

```
VkVertexInputAttributeDescription vviad[3]; // array per vertex input attribute
// 3 = position, velocity, color
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_VEC4; // x, y, z, w
vviad[0].offset = offsetof( struct pos, pos ); // 0

vviad[1].location = 1;
vviad[1].binding = 0;
vviad[1].format = VK_FORMAT_VEC4; // nx, ny, nz
vviad[1].offset = offsetof( struct vel, vel ); // 0

vviad[2].location = 2;
vviad[2].binding = 0;
vviad[2].format = VK_FORMAT_VEC4; // r, g, b, a
vviad[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 = vvid;
vpvisci.vertexAttributeDescriptionCount = 3;
vpvisci.pVertexAttributeDescriptions = vviad;

VkPipelineInputAssemblyStateCreateInfo vpiaci;
vpiaci.sType = VK_STRUCTURE_TYPE_PIPELINE_INPUT_ASSEMBLY_STATE_CREATE_INFO;
vpiaci.pNext = nullptr;
vpiaci.flags = 0;
vpiaci.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;
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 = vpssci;
vgpci.pVertexInputState = &vpvisci;
vgpci.pInputAssemblyState = &vpiasci;
vgpci.pTessellationState = (VkPipelineTessellationStateCreateInfo *)nullptr; // &vptsci
vgpci.pViewportState = &vpvsc;
vgpci.pRasterizationState = &vprsc;
vgpci.pMultisampleState = &vpmisci;
vgpci.pDepthStencilState = &vpdssci;
vgpci.pColorBlendState = &vpcbsci;
vgpci.pDynamicState = &vpdscli;
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,
PALLOCATION, OUT &GraphicsPipeline );

```

 mjb - July 24, 2020

Setting a Pipeline Barrier so the Drawing Waits for the Compute 358

```

VkBufferMemoryBarrier vmb;
vmb.sType = VK_STRUCTURE_TYPE_BUFFER_MEMORY_BARRIER;
vmb.pNext = nullptr;
vmb.srcAccessFlags = VK_ACCESS_SHADER_WRITE_BIT;
vmb.dstAccessFlags = VK_ACCESS_VERTEX_ATTRIBUTE_READ_BIT;
vmb.srcQueueFamilyIndex = 0;
vmb.dstQueueFamilyIndex = 0;
vmb.buffer =
vmb.offset = 0;
vmb.size = NUM_PARTICLES * sizeof( glm::vec4 );

const uint32 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 &vmb, 0, nullptr
);

```

 mjb - 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

```

 mjb - July 24, 2020

Setting a Pipeline Barrier so the Compute Waits for the Drawing 360

```

VkBufferMemoryBarrier vmb;
vmb.sType = VK_STRUCTURE_TYPE_BUFFER_MEMORY_BARRIER;
vmb.pNext = nullptr;
vmb.srcAccessFlags = 0;
vmb.dstAccessFlags = VK_ACCESS_UNIFORM_READ_BIT;
vmb.srcQueueFamilyIndex = 0;
vmb.dstQueueFamilyIndex = 0;
vmb.buffer =
vmb.offset = 0;
vmb.size = ???

const uint32 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 &vmb, 0, nullptr
);

```

 mjb - July 24, 2020

361

Vulkan.

Specialization Constants

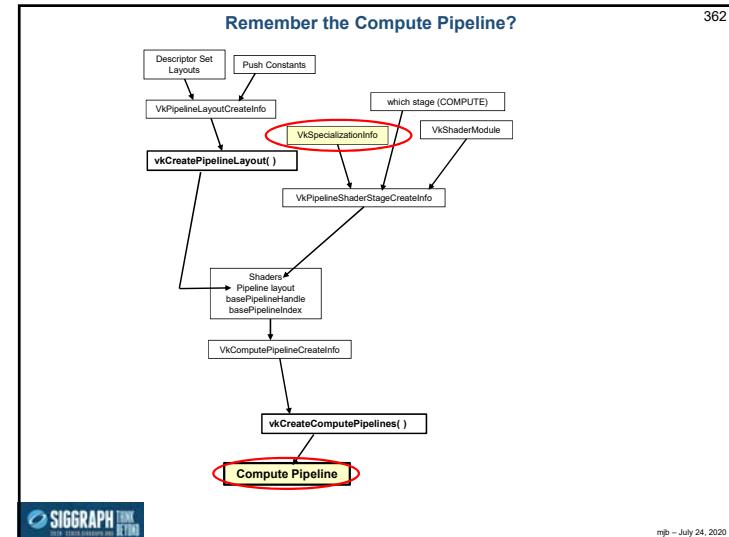
Mike Bailey
mjb@cs.oregonstate.edu

This work is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](#)

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



mjb – July 24, 2020



363

What Are Specialization Constants?

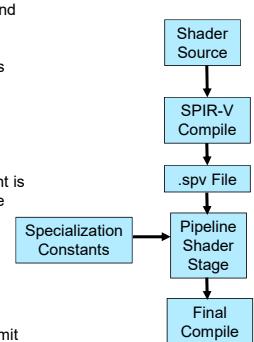
In Vulkan, all shaders get halfway-compiled into SPIR-V and then the rest-of-the-way compiled by the Vulkan driver.

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 halfway-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





mjb – July 24, 2020

364

Why Do We Need Specialization Constants?

Specialization Constants could be used for:

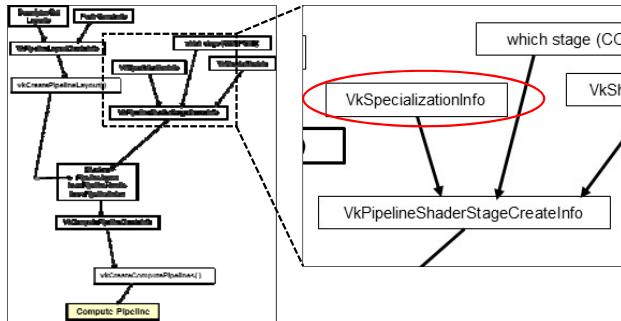
- Setting the work-items per work-group in a compute shader
- Setting a Boolean flag and then eliminating the if-test that used it
- Setting an integer constant and then eliminating the switch-statement that looked for it
- Making a decision to unroll a for-loop because the number of passes through it are small enough
- Collapsing arithmetic expressions into a single value
- Collapsing trivial simplifications, such as adding zero or multiplying by 1



mjb – July 24, 2020

Specialization Constants are Described in the Compute Pipeline

365



mb - July 24, 2020

Specialization Constant Example -- Setting an Array Size

366

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];
vsme[0].constantID = 7;
vsme[0].offset = 0;
vsme[0].size = sizeof(asize);
// one array element for each
// Specialization Constant
vsme[0].pData = &asize;
// # bytes into the Specialization Constant
// array this one item is
// size of just this Specialization Constant

VkSpecializationInfo vsi;
vsi.mapEntryCount = 1;
vsi.pMapEntries = &vsme[0];
vsi.dataSize = sizeof(asize);
vsi.pData = &asize;
// size of all the Specialization Constants together
// array of all the Specialization Constants
```



mb - July 24, 2020

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 vcpcl[1];
vcpcl[0].sType = VK_STRUCTURE_TYPE_COMPUTE_PIPELINE_CREATE_INFO;
vcpcl[0].pNext = nullptr;
vcpcl[0].flags = 0;
vcpcl[0].stage = vpssci;
vcpcl[0].layout = ComputePipelineLayout;
vcpcl[0].basePipelineHandle = VK_NULL_HANDLE;
vcpcl[0].basePipelineIndex = 0;

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

2020

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!



mb - July 24, 2020

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;
```

 mjb – July 24, 2020


Synchronization

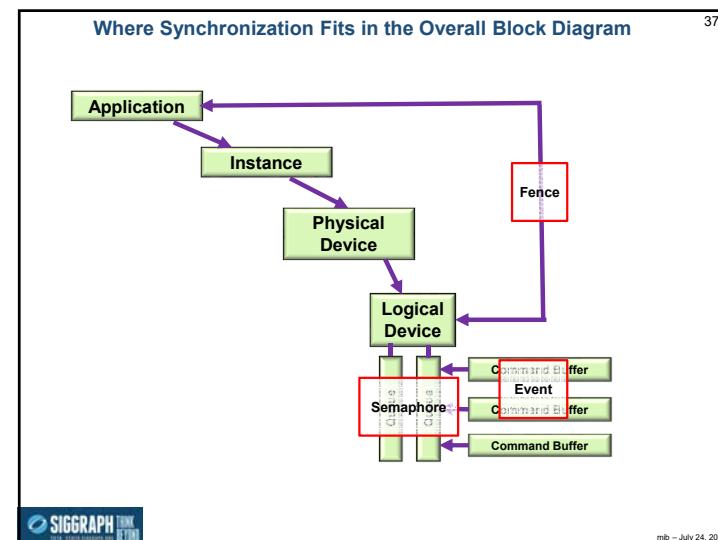
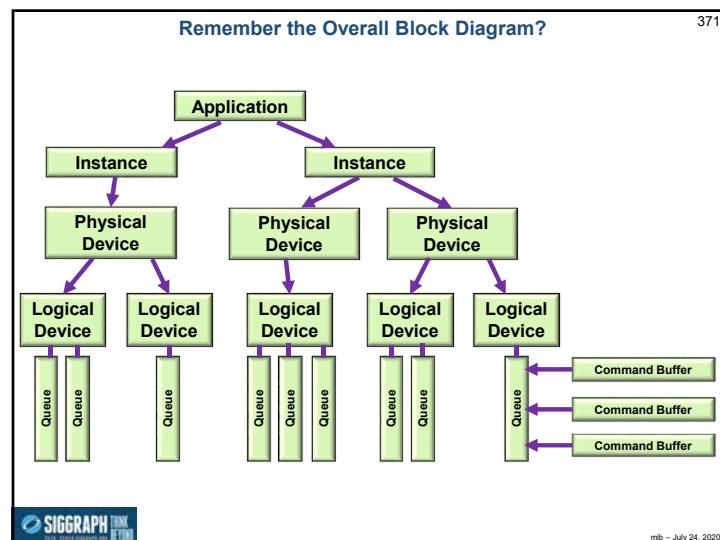
Mike Bailey
mjb@cs.oregonstate.edu



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

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

 mjb – July 24, 2020



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
vsci.sType = VK_STRUCTURE_TYPE_SEMAPHORE_CREATE_INFO;
vsci.pNext = nullptr;
vsci.flags = 0;

VkSemaphore      semaphore;
result = vkCreateSemaphore( LogicalDevice, IN &vsci, 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
vsci.sType = VK_STRUCTURE_TYPE_SEMAPHORE_CREATE_INFO;
vsci.pNext = nullptr;
vsci.flags = 0;
result = vkCreateSemaphore( LogicalDevice, IN &vsci, 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.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 );

```

You do this to wait for an image to be ready to be rendered into

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 vfc;
vfc.sType = VK_STRUCTURE_TYPE_FENCE_CREATE_INFO;
vfc.pNext = nullptr;
vfc.flags = VK_FENCE_CREATE_UNSIGNALED_BIT; // = 0
// VK_FENCE_CREATE_SIGNALLED_BIT is only other option

VkFence fence;
result = vkCreateFence( LogicalDevice, IN &vfc, 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, 1, IN &fence, waitForAll, timeout );
    Wait on the fence(s)
    // 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, &vfc, 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 );
    ...
result = vkWaitForFences( LogicalDevice, 1, IN &renderFence, VK_TRUE, UINT64_MAX );
    ...
result = vkQueuePresentKHR( presentQueue, IN &vpi );

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;
veci.sType = VK_STRUCTURE_TYPE_EVENT_CREATE_INFO;
veci.pNext = nullptr;
veci.flags = 0;

VkEvent event;
result = vkCreateEvent( LogicalDevice, IN &veci, PALLOCATOR, OUT &event );
    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, // Could be an array of events
                        srcPipelineStageBits, dstPipelineStageBits, // Where signaled, where wait for the signal
                        memoryBarrierCount, pMemoryBarriers,
                        bufferMemoryBarrierCount, pBufferMemoryBarriers,
                        imageMemoryBarrierCount, pImageMemoryBarriers );
// Memory barriers get executed after events have been signaled

Note: the device cannot test for an event, but it can block

```

SIGGRAPH THINK

mjb – July 24, 2020

382

Vulkan.

Pipeline Barriers

Mike Bailey
mjb@cs.oregonstate.edu



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

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

SIGGRAPH THINK

mjb – July 24, 2020

From the Command Buffer Notes:

These are the Commands that can be entered into the Command Buffer, I 383

```

vkCmdBeginQuery( commandBuffer, flags );
vkCmdBeginRenderPass( commandBuffer, const contents );
vkCmdBindDescriptorSets( commandBuffer, pDynamicOffsets );
vkCmdBindIndexBuffer( commandBuffer, indexType );
vkCmdBindPipeline( commandBuffer, pipeline );
vkCmdBindVertexBuffers( commandBuffer, firstBinding, bindingCount, const pOffsets );
vkCmdBlitImage( commandBuffer, filter, attachmentCount, const pRects );
vkCmdClearDepthStencilImage( commandBuffer, pRanges );
vkCmdCopyImage( commandBuffer, pRegions );
vkCmdCopyBufferToImage( commandBuffer, pRegion );
vkCmdCopyImage( commandBuffer, pRegions );
vkCmdCopyImageToBuffer( commandBuffer, pRegions );
vkCmdDebugMarkerBeginEXT( commandBuffer, pMarkerInfo );
vkCmdDebugMarkerEndEXT( commandBuffer );
vkCmdDebugMarkerInsertEXT( commandBuffer, pMarkerInfo );
vkCmdDispatch( commandBuffer, groupCountX, groupCountY, groupCountZ );
vkCmdDispatchIndirect( commandBuffer, offset );
vkCmdDraw( commandBuffer, vertexCount, instanceCount, firstVertex, firstInstance );
vkCmdDrawIndexed( commandBuffer, indexCount, instanceCount, firstIndex, int32_t vertexOffset, firstInstance );
vkCmdDrawIndexedIndirect( commandBuffer, stride );
vkCmdDrawIndexedMultiAMD( commandBuffer, stride );
vkCmdDrawIndirectCountAMD( commandBuffer, stride );
vkCmdEndQuery( commandBuffer, query );
vkCmdEndRenderPass( commandBuffer );
vkCmdExecuteCommands( commandBuffer, commandBufferCount, const pCommandBuffers );

```

SIGGRAPH THINK

mjb – July 24, 2020

From the Command Buffer Notes:

These are the Commands that can be entered into the Command Buffer, II 384

```

vkCmdFillBuffer( commandBuffer, dstBuffer, dstOffset, size, data );
vkCmdNextSubpass( commandBuffer, contents );
vkCmdPipelineBarrier( commandBuffer, srcStageMask, dstStageMask, dependencyFlags, memoryBarrierCount, VkMemoryBarrier* pMemoryBarriers,
                     bufferMemoryBarrierCount, pBufferMemoryBarriers, imageMemoryBarrierCount, pImageMemoryBarriers );
vkCmdProcessCommandsNVX( commandBuffer, pProcessCommandsInfo );
vkCmdPushDescriptorSet( commandBuffer, layout, stageFlags, offset, size, pValues );
vkCmdPushDescriptorSetWithTemplate( commandBuffer, pipelineBindPoint, layout, set, descriptorWriteCount, pDescriptorWrites );
vkCmdPushDescriptorSetWithTemplate( commandBuffer, pipelineBindPoint, layout, set, descriptorUpdateTemplate, pData );
vkCmdReserveSpaceForCommandsNVX( commandBuffer, pReserveSpaceInfo );
vkCmdResetEvent( commandBuffer, event, stageMask );
vkCmdResetQueryPool( commandBuffer, queryPool, firstQuery, queryCount );
vkCmdResolveDepthStencilImage( commandBuffer, srcImage, srcImageLayout, dstImage, dstImageLayout, regionCount, pRegions );
vkCmdSetBlendConstants( commandBuffer, blendConstants );
vkCmdSetBlendEquation( commandBuffer, blendEquation );
vkCmdSetBlendEquationSeparate( commandBuffer, minBlendEquation, maxBlendEquation, depthBiasClamp, depthBiasSlopeFactor );
vkCmdSetDepthBounds( commandBuffer, minDepthBounds, maxDepthBounds );
vkCmdSetDeviceMaskKKHX( commandBuffer, deviceMask );
vkCmdSetDiscardRectangleEXT( commandBuffer, firstDiscardRectangle, discardRectangleCount, pDiscardRectangles );
vkCmdSetEvent( commandBuffer, event, stageMask );
vkCmdSetLineWidth( commandBuffer, lineWidth );
vkCmdSetLineWidth( commandBuffer, lineWidth );
vkCmdSetScissor( commandBuffer, firstScissor, scissorCount, pScissors );
vkCmdSetScissorWithSize( commandBuffer, firstScissor, scissorCount, pScissors, faceMask );
vkCmdSetStencilRefrence( commandBuffer, faceMask, reference );
vkCmdSetStencilWriteMask( commandBuffer, faceMask, writeMask );
vkCmdSetViewport( commandBuffer, firstViewport, viewportCount, pViewports );
vkCmdSetViewportScalingNV( commandBuffer, firstViewport, viewportCount, pViewportScalings );
vkCmdUpdateBuffer( commandBuffer, dstBuffer, dstOffset, dataSize, pData );
vkCmdWaitEvents( commandBuffer, eventCount, pEvents, srcStageMask, dstStageMask, memoryBarrierCount, pMemoryBarriers,
                 bufferMemoryBarrierCount, pBufferMemoryBarriers, imageMemoryBarrierCount, pImageMemoryBarriers );
vkCmdWriteTimestamp( commandBuffer, pipelineStage, queryPool, query );

```

SIGGRAPH THINK

mjb – July 24, 2020

Potential Memory Race Conditions that Pipeline Barriers can Prevent 385

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

mb - July 24, 2020

vkCmdPipelineBarrier() Function Call 386

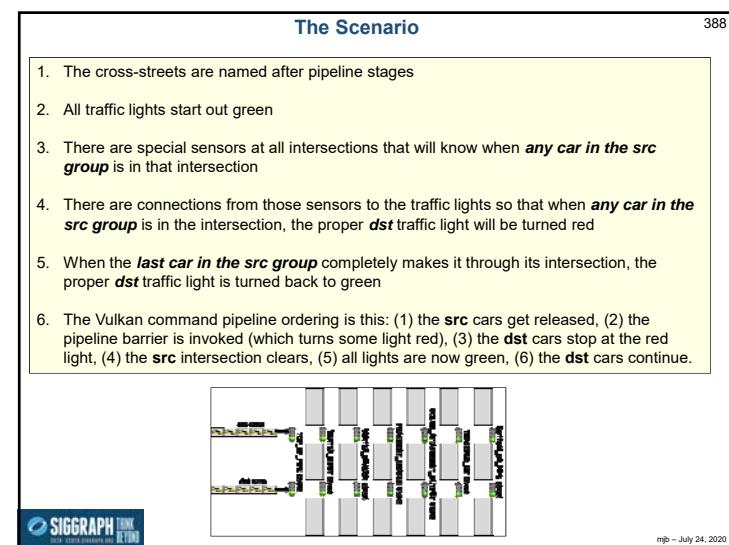
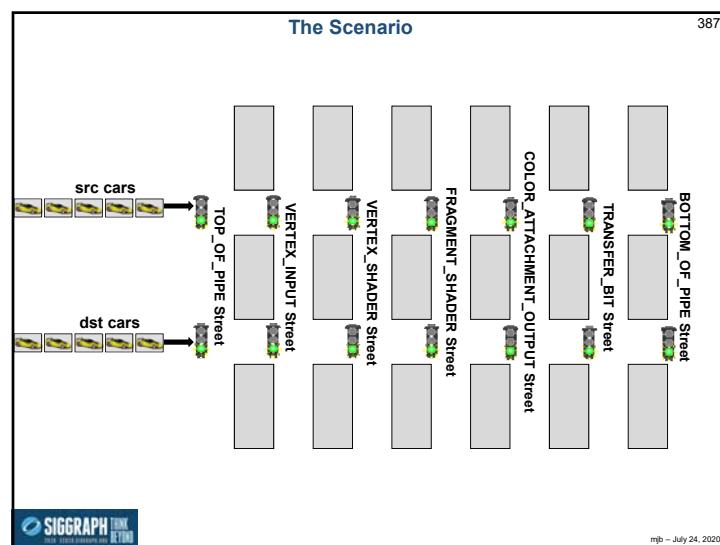
A Pipeline Barrier is a way to establish a memory dependency between commands that were submitted before the barrier and commands that are submitted after the barrier

```

vkCmdPipelineBarrier( commandBuffer,
    srcStageMask,           // Guarantee that this pipeline stage is completely done being used
                           // before ...
    dstStageMask,           // ... allowing this pipeline stage to be used
    VK_DEPENDENCY_BY_REGION_BIT,
    memoryBarrierCount,     // pMemoryBarriers,
    bufferMemoryBarrierCount, // pBufferMemoryBarriers,
    imageMemoryBarrierCount, // pImageMemoryBarriers
);
  
```

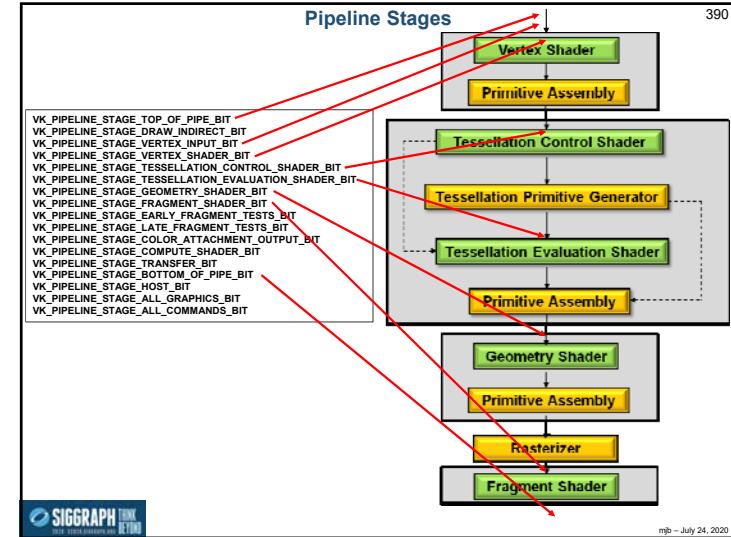
Defines what data we will be blocking on or un-blocking on

mb - July 24, 2020



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	

mb – July 24, 2020

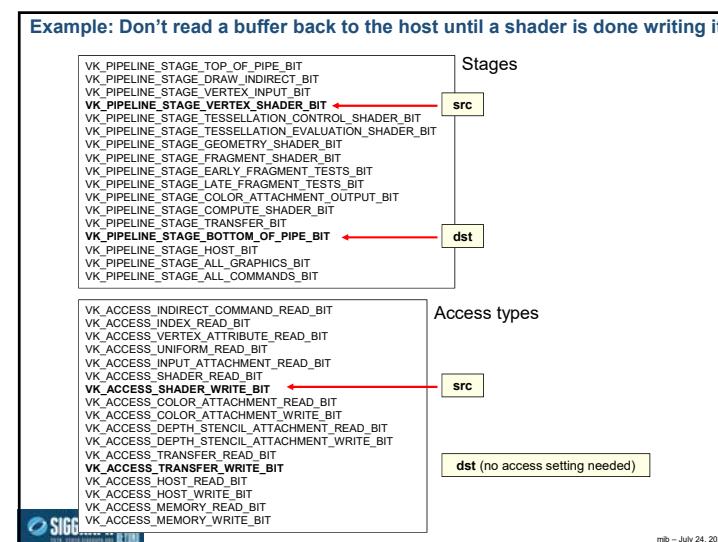
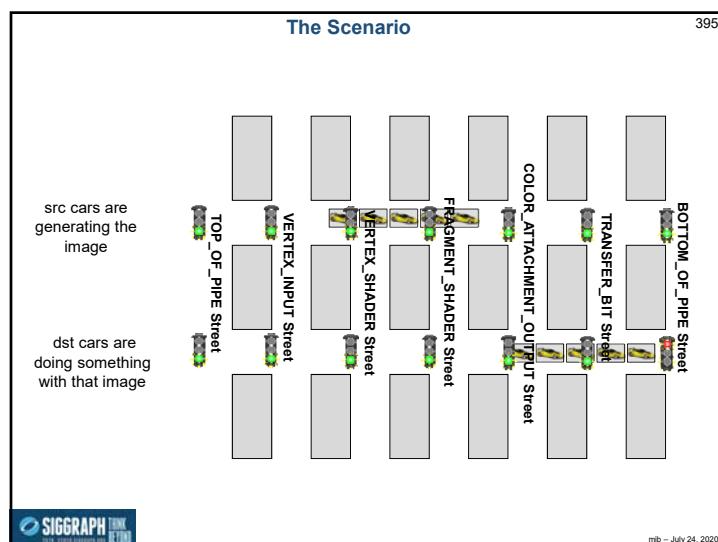
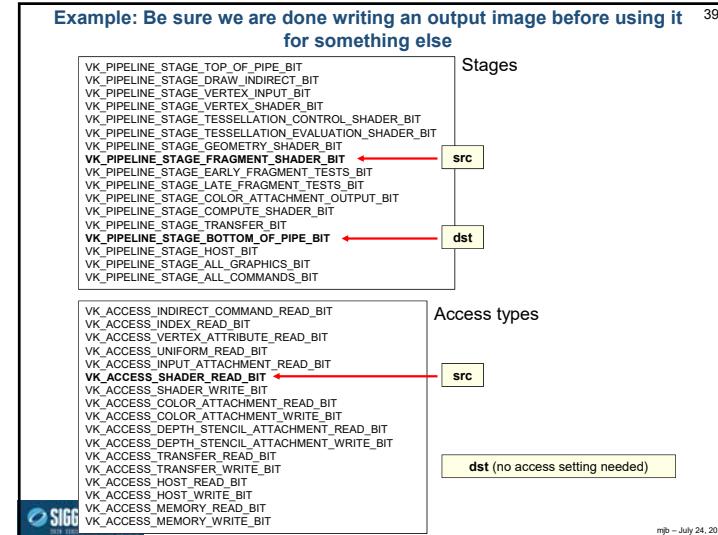
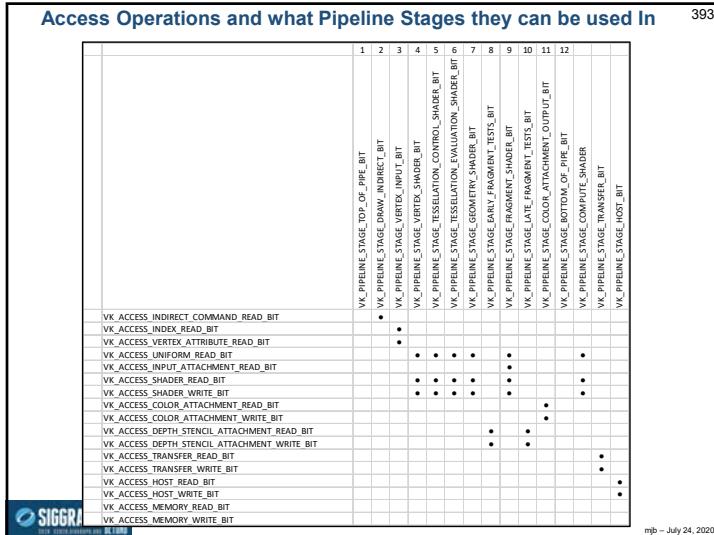


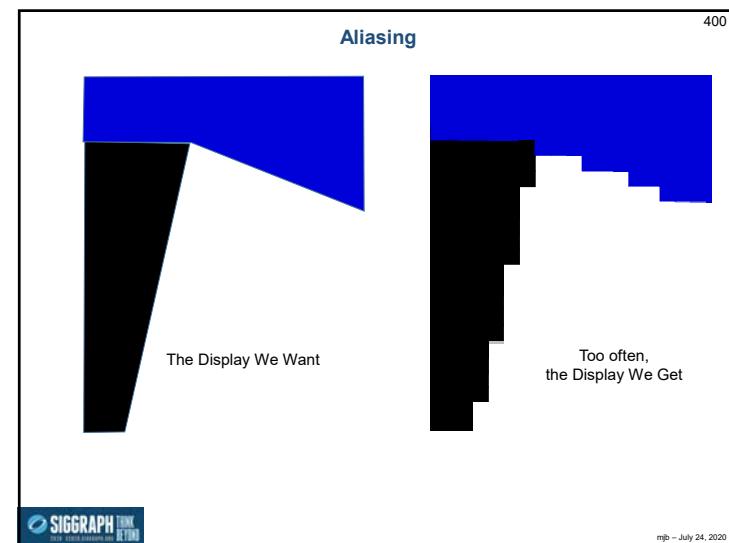
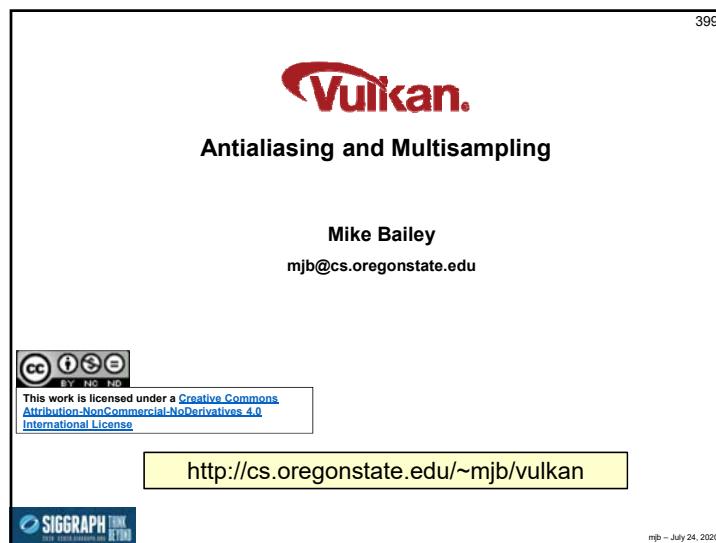
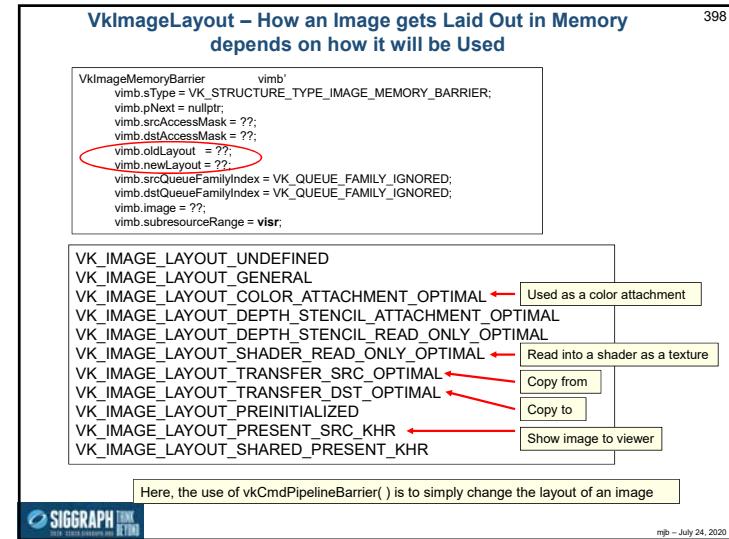
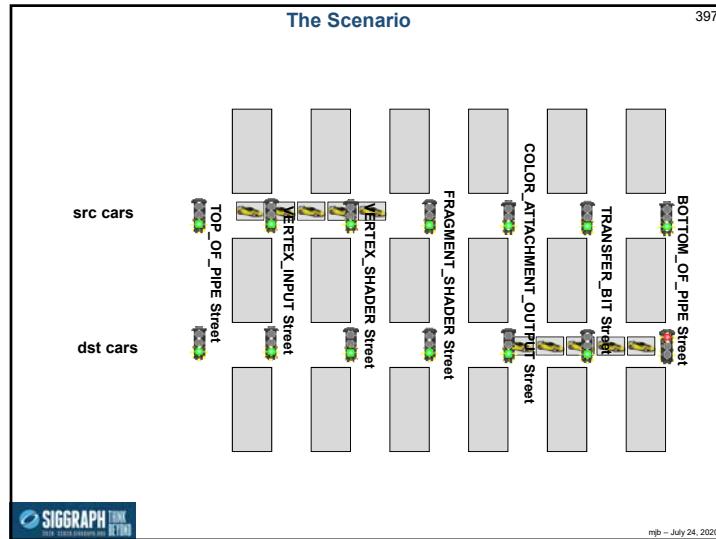
Access Masks –	
What are you Interested in Generating or Consuming this Memory for?	
VK_ACCESS_INDIRECT_COMMAND_READ_BIT	
VK_ACCESS_INDEX_READ_BIT	
VK_ACCESS_VERTEX_ATTRIBUTE_READ_BIT	
VK_ACCESS_UNIFORM_READ_BIT	
VK_ACCESS_INPUT_ATTACHMENT_READ_BIT	
VK_ACCESS_SHADER_READ_BIT	
VK_ACCESS_SHADER_WRITE_BIT	
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	

mb – July 24, 2020

Pipeline Stages and what Access Operations are Allowed	
VK_PIPELINE_STAGE_TOP_OF_PIPE_BIT	• VK_ACCESS_INDIRECT_COMMAND_READ_BIT
VK_PIPELINE_STAGE_DRAW_INDIRECT_BIT	• VK_ACCESS_INDEX_READ_BIT
VK_PIPELINE_STAGE_VERTEX_INPUT_BIT	• VK_ACCESS_VERTEX_ATTRIBUTE_READ_BIT
VK_PIPELINE_STAGE_VERTEX_SHADER_BIT	• VK_ACCESS_UNIFORM_READ_BIT
VK_PIPELINE_STAGE_TESSELLATION_CONTROL_SHADER_BIT	• VK_ACCESS_INPUT_ATTACHMENT_READ_BIT
VK_PIPELINE_STAGE_TESSELLATION_EVALUATION_SHADER_BIT	• VK_ACCESS_SHADER_READ_BIT
VK_PIPELINE_STAGE_GEOMETRY_SHADER_BIT	• VK_ACCESS_COLOR_ATTACHMENT_READ_BIT
VK_PIPELINE_STAGE_FRAGMENT_SHADER_BIT	• VK_ACCESS_DEPTH_STENCIL_ATTACHMENT_READ_BIT
VK_PIPELINE_STAGE_EARLY_FRAGMENT_TESTS_BIT	• VK_ACCESS_SHADER_WRITE_BIT
VK_PIPELINE_STAGE_LATE_FRAGMENT_TESTS_BIT	• VK_ACCESS_COLOR_ATTACHMENT_WRITE_BIT
VK_PIPELINE_STAGE_COLOR_ATTACHMENT_OUTPUT_BIT	• VK_ACCESS_DEPTH_STENCIL_ATTACHMENT_WRITE_BIT
VK_PIPELINE_STAGE_BOTTOM_OF_PIPE_BIT	• VK_ACCESS_TRANSFER_READ_BIT
VK_PIPELINE_STAGE_COMPUTE_SHADER_BIT	• VK_ACCESS_TRANSFER_WRITE_BIT
VK_PIPELINE_STAGE_TRANSFER_BIT	• VK_ACCESS_HOST_READ_BIT
VK_PIPELINE_STAGE_HOST_BIT	• VK_ACCESS_HOST_WRITE_BIT
VK_PIPELINE_STAGE_MEMORY_READ_BIT	• VK_ACCESS_MEMORY_READ_BIT
VK_PIPELINE_STAGE_MEMORY_WRITE_BIT	• VK_ACCESS_MEMORY_WRITE_BIT

mb – July 24, 2020





Aliasing 401

"Aliasing" is a signal-processing term for "under-sampled compared with the frequencies in the signal".

What the signal really is: what we want
Sampling Interval
What we think the signal is: too often, what we get
Sampled Points

SIGGRAPH THINK THINK. DESIGN. INNOVATE. BEYOND.

mb - July 24, 2020

Aliasing 402

SIGGRAPH THINK THINK. DESIGN. INNOVATE. BEYOND.

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 THINK THINK. DESIGN. INNOVATE. BEYOND.

mb - July 24, 2020

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:

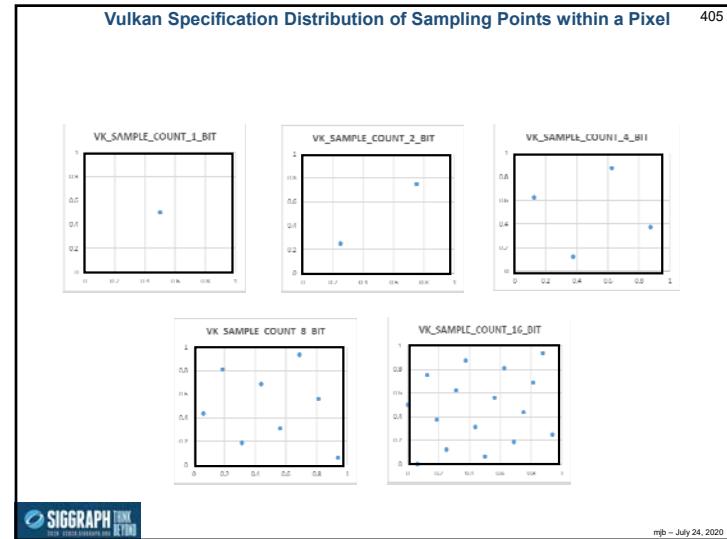
- Supersampling:** Pick some number of sub-pixels within that pixel that pass the depth and stencil tests. Render the image at each of these sub-pixels..

- 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 THINK THINK. DESIGN. INNOVATE. BEYOND.

mb - July 24, 2020

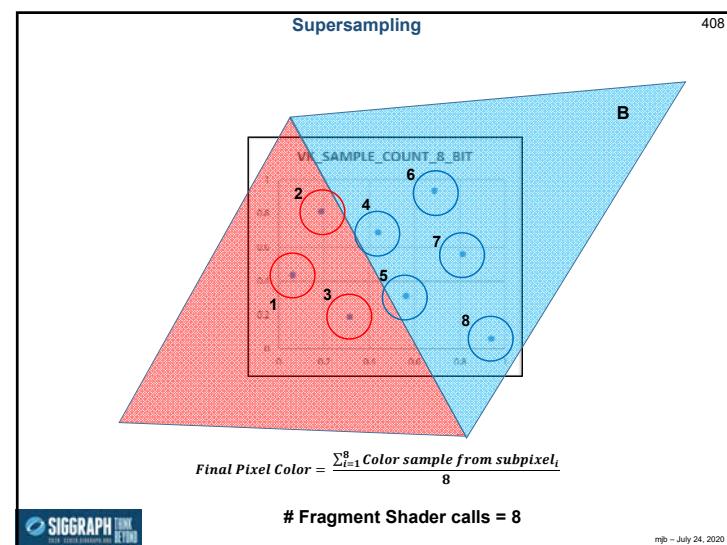
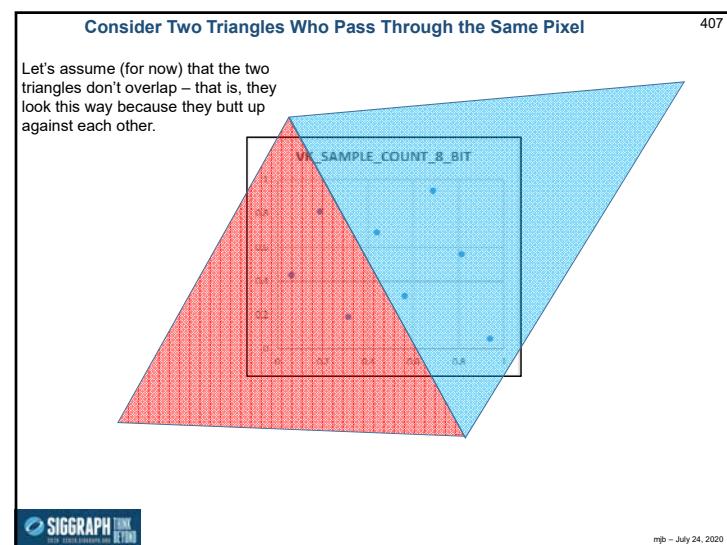


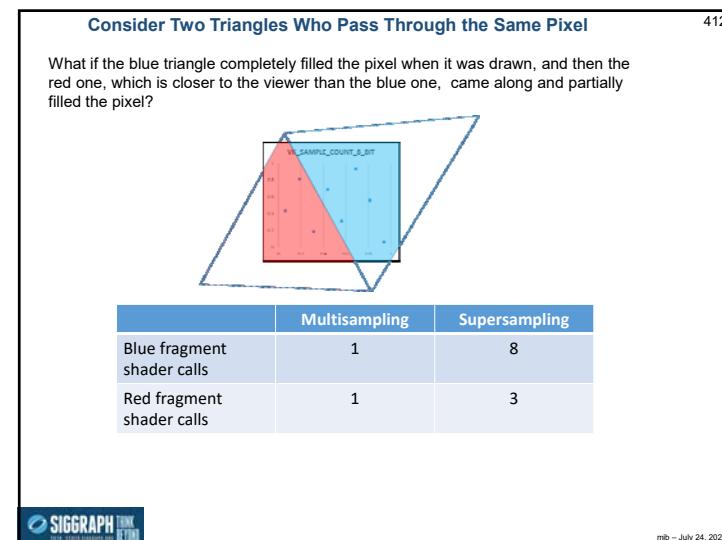
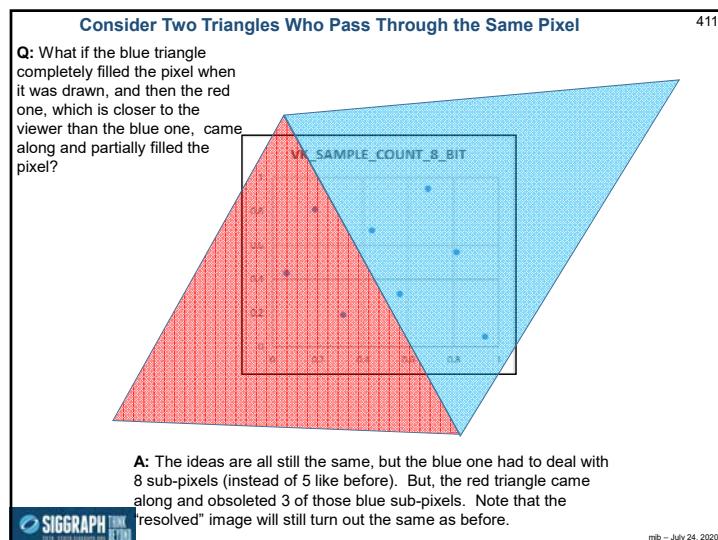
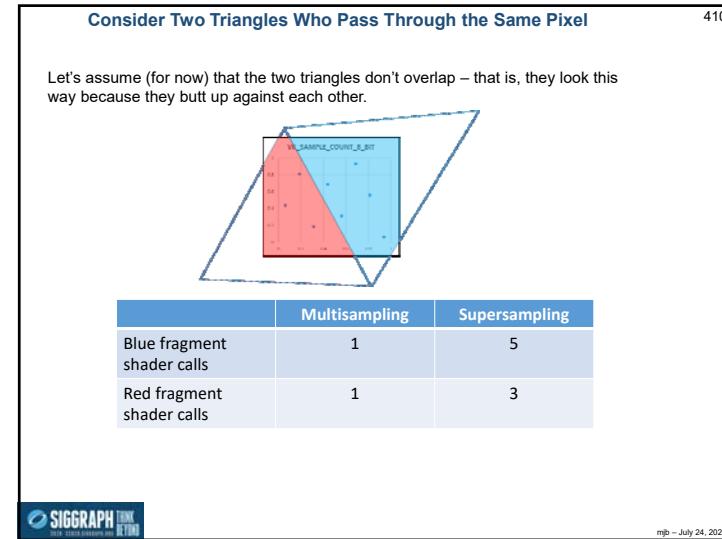
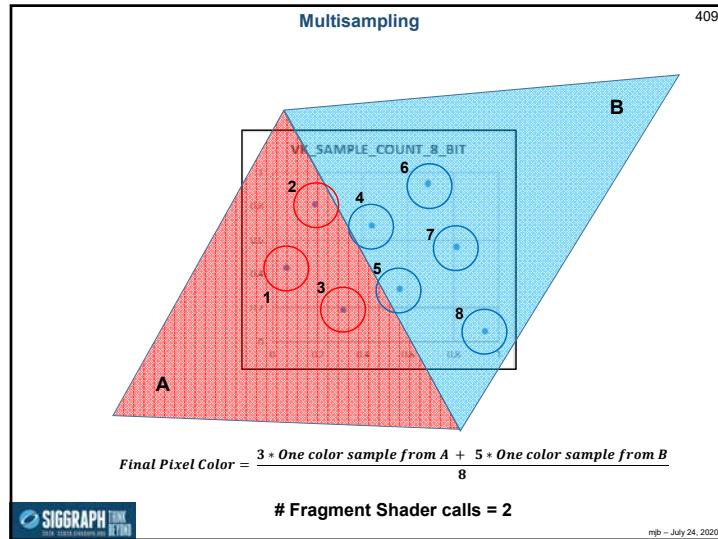
Vulkan Specification Distribution of Sampling Points within a Pixel 406

VK_SAMPLE_COUNT_2_BIT	VK_SAMPLE_COUNT_4_BIT	VK_SAMPLE_COUNT_8_BIT	VK_SAMPLE_COUNT_16_BIT
(0.25, 0.25)	(0.375, 0.125)	(0.4375, 0.6875)	(0.5625, 0.5625)
	(0.875, 0.375)	(0.8125, 0.5625)	(0.4375, 0.3125)
	(0.125, 0.625)	(0.1875, 0.8125)	(0.3125, 0.625)
(0.75, 0.75)		(0.6875, 0.9375)	(0.8125, 0.6875)
		(0.625, 0.875)	(0.5625, 0.1875)
		(0.0625, 0.4375)	(0.375, 0.125)
		(0.0, 0.5)	(0.5, 0.0625)
		(0.9375, 0.0625)	(0.875, 0.9375)
		(0.9375, 0.25)	(0.625, 0.0)

SIGGRAPH THINK
GEAR. DESIGN. INNOVATE. REINVENT.

mb - July 24, 2020





Setting up the Image 413

```

VkPipelineMultisampleStateCreateInfo
vpmci.sType = VK_STRUCTURE_TYPE_PIPELINE_MULTISAMPLE_STATE_CREATE_INFO;
vpmci.pNext = nullptr;
vpmci.flags = 0;
vpmci.rasterizationSamples = VK_SAMPLE_COUNT_8_BIT;
vpmci.sampleShadingEnable = VK_TRUE;
vpmci.minSampleShading = 0.5f;
vpmci.sampleMask = (VkSampleMask *)nullptr;
vpmci.alphaToCoverageEnable = VK_FALSE;
vpmci.alphaToOneEnable = VK_FALSE;

VkGraphicsPipelineCreateInfo
vgci.sType = VK_STRUCTURE_TYPE_GRAPHICS_PIPELINE_CREATE_INFO;
vgci.pNext = nullptr;
...
vgci.pMultisampleState = &vpmci;

result = vkCreateGraphicsPipelines( LogicalDevice, VK_NULL_HANDLE, 1, IN &vgci,
PALLOCATOR, OUT pGraphicsPipeline );

```

vpmci: How dense is the sampling
VK_TRUE means to allow some sort of multisampling to take place

vgci:

SIGGRAPH THINK
CREATE, EXPERIMENT, REFINEMENT

mb - July 24, 2020

Setting up the Image 414

```

VkPipelineMultisampleStateCreateInfo vpmci;
...
vpmci.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

```

VK_SAMPLE_COUNT_8_BIT

SIGGRAPH THINK
CREATE, EXPERIMENT, REFINEMENT

mb - July 24, 2020

Setting up the Image 415

```

VkAttachmentDescription
vad[0].format = VK_FORMAT_B3G6R8A8_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;
colorReference.attachment = 0;
colorReference.layout = VK_IMAGE_LAYOUT_COLOR_ATTACHMENT_OPTIMAL;

VkAttachmentReference depthReference;
depthReference.attachment = 1;
depthReference.layout = VK_IMAGE_LAYOUT_DEPTH_STENCIL_ATTACHMENT_OPTIMAL;

```

vad[2]: to next slide

SIGGRAPH THINK
CREATE, EXPERIMENT, REFINEMENT

mb - July 24, 2020

Setting up the Image 416

```

vsd: from previous slide

VkSubpassDescription vsd;
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
vrpc.i.sType = VK_STRUCTURE_TYPE_RENDER_PASS_CREATE_INFO;
vrpc.i.pNext = nullptr;
vrpc.i.flags = 0;
vrpc.i.attachmentCount = 2;
vrpc.i.pAttachments = vad; // color and depth/stencil
vrpc.i.subpassCount = 1;
vrpc.i.pSubpasses = IN &vsd;
vrpc.i.dependencyCount = 0;
vrpc.i.pDependencies = (VkSubpassDependency *)nullptr;

result = vkCreateRenderPass( LogicalDevice, IN &vrpc.i, PALLOCATOR, OUT &RenderPass );

```

SIGGRAPH THINK
CREATE, EXPERIMENT, REFINEMENT

mb - July 24, 2020

**Resolving the Image:
Converting the Multisampled Image to a VK_SAMPLE_COUNT_1_BIT image**

```

VIOffset3D
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, srclImage, srclImageLayout, dstImage, dstImageLayout, 1, IN &vir );

```

For the *ImageLayout, use VK_IMAGE_LAYOUT_GENERAL

mjb - July 24, 2020

417

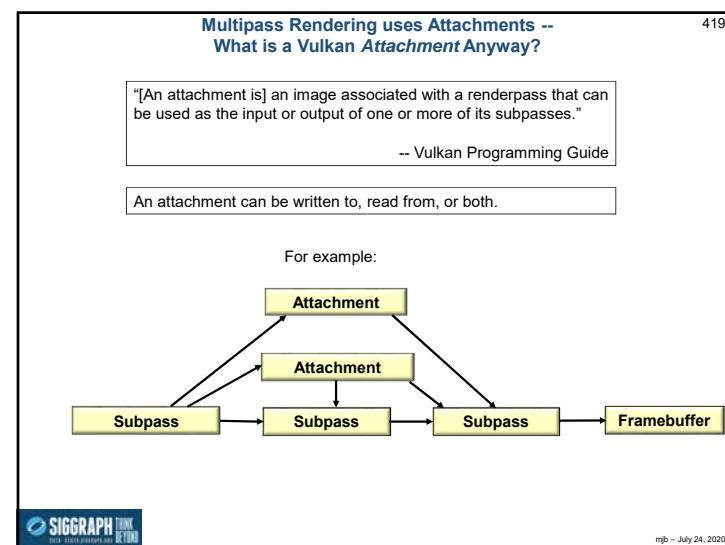
Vulkan.
Multipass Rendering

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>

mjb - July 24, 2020



419

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**.

mjb - July 24, 2020

Back in Our Single-pass Days 421

So far, we've only performed single-pass rendering, within a single Vulkan RenderPass.

```

graph TD
    A[Attachment #1] --> B[Depth Attachment]
    C[Attachment #0] --> D[Output]
    B --> E[3D Rendering Pass]
    D --> E
    E --> F[Output]
    subgraph Subpass_0 [Subpass #0]
        E
    end

```

Here comes a quick reminder of how we did that.

Afterwards, we will extend it.

SIGGRAPH THINK
CREATE. DESIGN. ENTERTAIN. REINVENT.

mb - July 24, 2020

Back in Our Single-pass Days, I 422

```

VkAttachmentDescription vad[2];
vad[0].flags = 0;
vad[0].format = VK_FORMAT_B8GRRB8_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;
colorReference.attachment = 0;
colorReference.layout = VK_IMAGE_LAYOUT_COLOR_ATTACHMENT_OPTIMAL;

VkAttachmentReference depthReference;
depthReference.attachment = 1;
depthReference.layout = VK_IMAGE_LAYOUT_DEPTH_STENCIL_ATTACHMENT_OPTIMAL;

```

SIGGRAPH THINK
CREATE. DESIGN. ENTERTAIN. REINVENT.

mb - July 24, 2020

Back in Our Single-pass Days, II 423

```

VkSubpassDescription vsd;
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;
vrpci.sType = VK_STRUCTURE_TYPE_RENDER_PASS_CREATE_INFO;
vrpci.pNext = nullptr;
vrpci.flags = 0;
vrpci.attachmentCount = 2;
vrpci.pAttachments = &vad;
vrpci.subpassCount = 1;
vrpci.pSubpasses = &vsd;
vrpci.dependencyCount = 0;
vrpci.pDependencies = (VkSubpassDependency *)nullptr;

result = vkCreateRenderPass( LogicalDevice, IN &vrpci, PALLOCATOR, OUT &RenderPass );

```

SIGGRAPH THINK
CREATE. DESIGN. ENTERTAIN. REINVENT.

mb - July 24, 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.

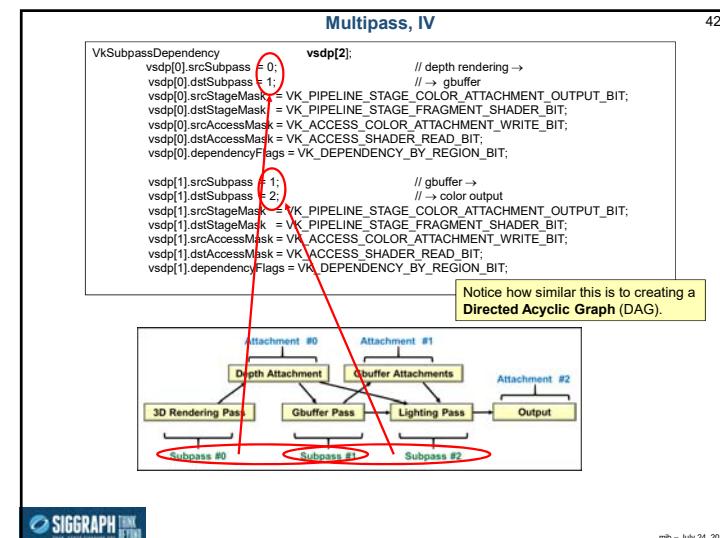
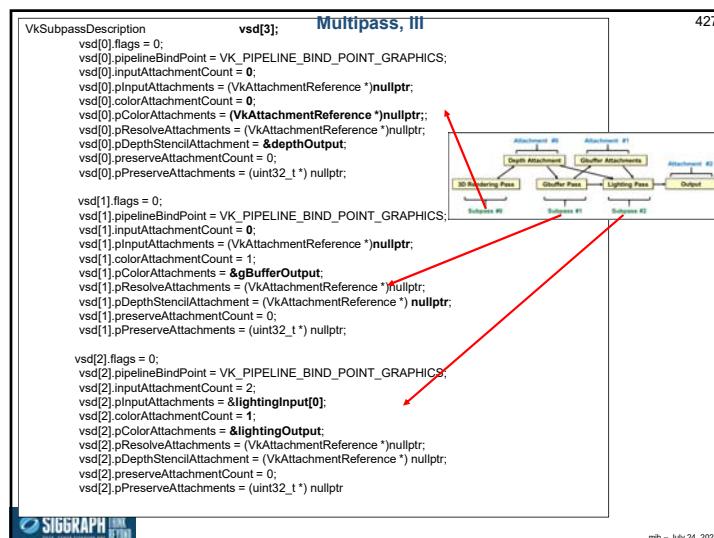
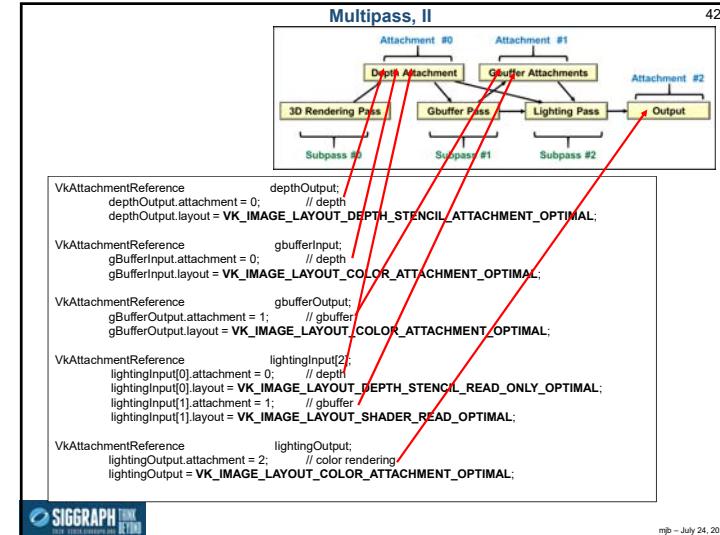
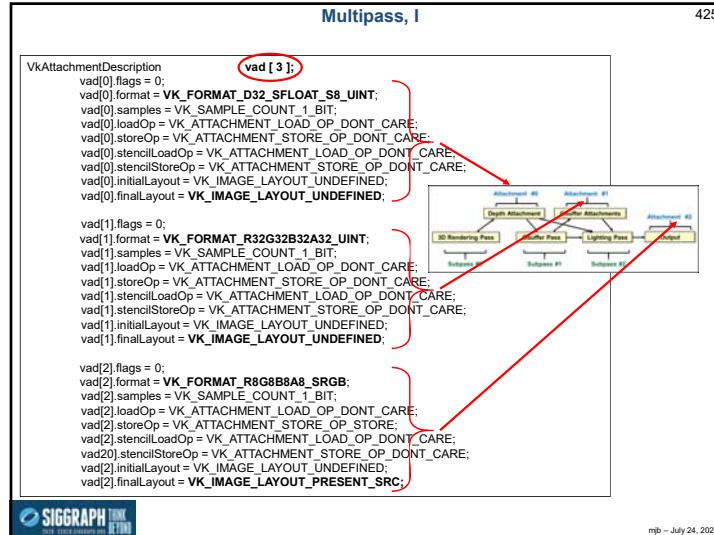
```

graph LR
    A[Attachment #0] --> B[3D Rendering Pass]
    A --> C[Gbuffer Pass]
    D[Attachment #1] --> E[Lighting Pass]
    C --> F[Output]
    C --> E
    E --> F
    subgraph Subpass_0 [Subpass #0]
        B
    end
    subgraph Subpass_1 [Subpass #1]
        C
    end
    subgraph Subpass_2 [Subpass #2]
        E
    end

```

SIGGRAPH THINK
CREATE. DESIGN. ENTERTAIN. REINVENT.

mb - July 24, 2020



Multipass, V 429

```

VkRenderPassCreateInfo vrpcl;
vrpcl.sType = VK_STRUCTURE_TYPE_RENDER_PASS_CREATE_INFO;
vrpcl.pNext = nullptr;
vrpcl.flags = 0;
vrpcl.attachmentCount = 3; // depth, gbuffer, output
vrpcl.pAttachments = vad;
vrpcl.subpassCount = 3;
vrpcl.pSubpasses = vsd;
vrpcl.dependencyCount = 2;
vrpcl.pDependencies = vsdp;

result = vkCreateRenderPass( LogicalDevice, IN &vrpcl, PALLOCATOR, OUT &RenderPass );

```

SIGGRAPH THINK

mb - July 24, 2020

Multipass, VI 430

```

vkCmdBeginRenderPass( CommandBuffers[nextImageIndex], IN &vrpb, IN VK_SUBPASS_CONTENTS_INLINE );
// subpass #0 is automatically started here

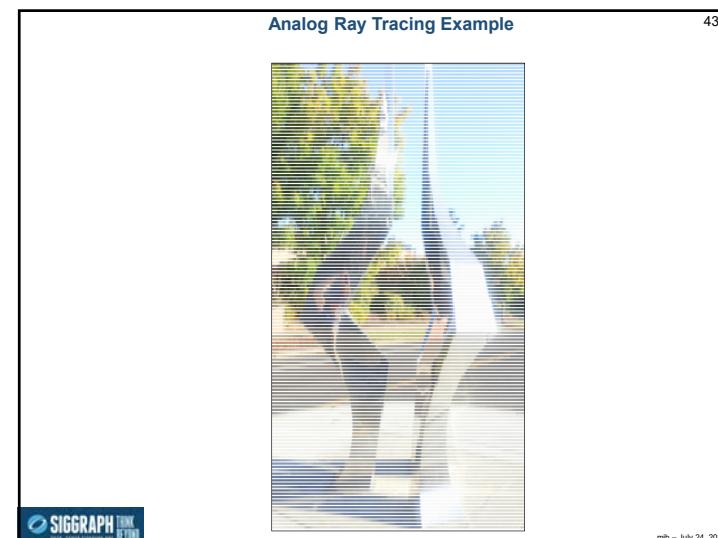
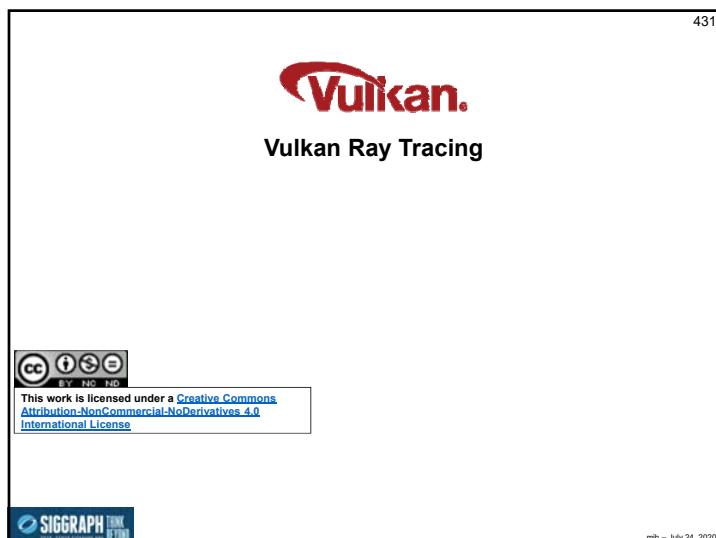
vkCmdBindPipeline( CommandBuffers[nextImageIndex], VK_PIPELINE_BIND_POINT_GRAPHICS, GraphicsPipeline );
vkCmdBindDescriptorSets( CommandBuffers[nextImageIndex], VK_PIPELINE_BIND_POINT_GRAPHICS,
GraphicsPipelineLayout, 0, 4, DescriptorSets, 0, (uint32_t*) nullptr );
vkCmdBindVertexBuffers( CommandBuffers[nextImageIndex], 0, 1, vBuffers, offsets );
vkCmdDraw( CommandBuffers[nextImageIndex], vertexCount, instanceCount, firstVertex, firstInstance );
...
vkCmdNextSubpass( CommandBuffers[nextImageIndex], VK_SUBPASS_CONTENTS_INLINE );
// subpass #1 is started here

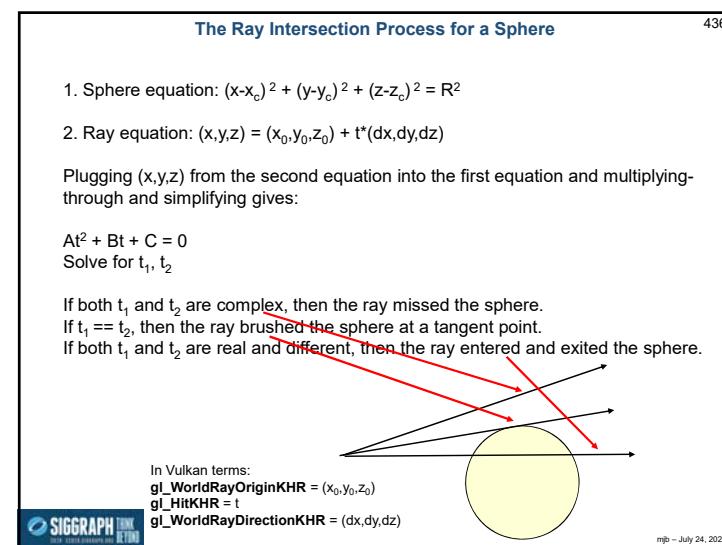
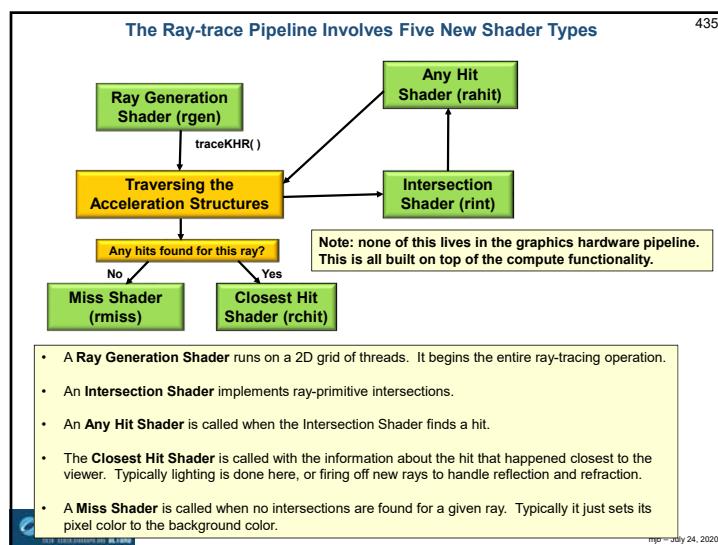
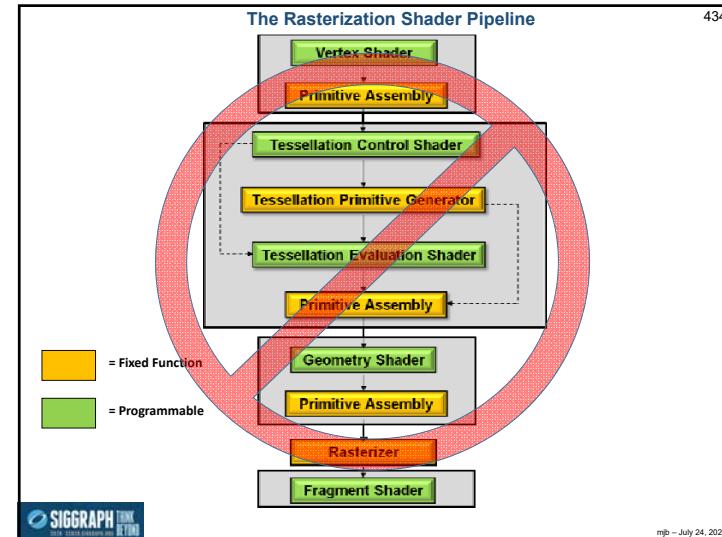
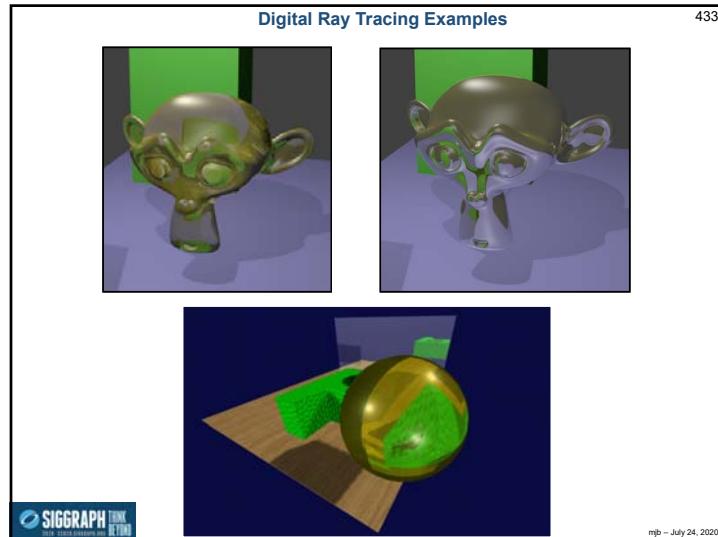
vkCmdNextSubpass( CommandBuffers[nextImageIndex], VK_SUBPASS_CONTENTS_INLINE );
// subpass #2 is started here
...
vkCmdEndRenderPass( CommandBuffers[nextImageIndex] );

```

SIGGRAPH THINK

mb - July 24, 2020





The Ray Intersection Process for a Cube 437

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:

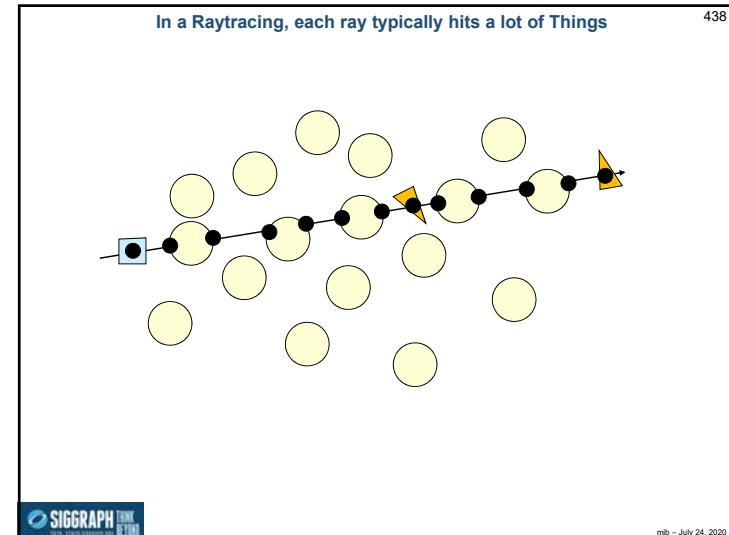
At + B = 0
Solve for t

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:
 $t_{\min} = \max(\min(t_{x1}, t_{x2}), \min(t_{y1}, t_{y2}), \min(t_{z1}, t_{z2}))$
 $t_{\max} = \min(\max(t_{x1}, t_{x2}), \max(t_{y1}, t_{y2}), \max(t_{z1}, t_{z2}))$

SIGGRAPH THINK
SIGGRAPH 2019: www.siggraph.org/2019

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.

SIGGRAPH THINK
SIGGRAPH 2019: www.siggraph.org/2019

mb - July 24, 2020

Creating Bottom Level Acceleration Structures 440

```

vkCreateAccelerationStructureKHR          BottomLevelAccelerationStructure;
{
    VkAccelerationStructureCreateInfoKHR 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 );

```

SIGGRAPH THINK
SIGGRAPH 2019: www.siggraph.org/2019

mb - July 24, 2020

Creating Top Level Acceleration Structures 441

```

vkCreateAccelerationStructureKHR          TopLevelAccelerationStructure;
VkAccelerationStructureCreateInfoKHR      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 );

```

mb - 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.

mb - 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` = (x_0, y_0, z_0)
`gl_HitKHR` = t
`gl_WorldRayDirectionKHR` = (dx, dy, dz)

mb - 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 attrs
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;
        attrs = p;
        reportIntersectionKHR( tmp, 0 );
        return;
    }
    tmp = (-b + sqrt(discr)) / (2.*a);
    if( gl_RayTminKHR < tmp && tmp < gl_RayTmaxKHR )
    {
        vec3 p = orig + tmp * dir;
        attrs = p;
        reportIntersectionKHR( tmp, 0 );
        return;
    }
}

```

mb - 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. );
}
```

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;
}
```

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:
 $gl_WorldRayOriginKHR = (x_0, y_0, z_0)$
 $gl_HitKHR = t$
 $gl_WorldRayDirectionKHR = (dx, dy, dz)$

mb - July 24, 2020

Other New Built-in Functions 448

Loosely equivalent to "discard"

```
void terminateRayKHR();
void ignoreIntersectionKHR();
void reportIntersectionKHR( float hit, uint hitKind );
```

mb - July 24, 2020

Ray Trace Pipeline Data Structure 449

```

VkPipeline RaytracePipeline;
VkPipelineLayout PipelineLayout;

VkPipelineLayoutCreateInfo vrpcl;
vrpcl.sType = VK_STRUCTURE_TYPE_PIPELINE_LAYOUT_CREATE_INFO;
vrpcl.pNext = nullptr;
vrpcl.flags = 0;
vrpcl.setLayoutCount = 1;
vrpcl.pSetLayouts = &descriptorsetLayout;
vrpcl.pushConstantRangeCount = 0;
vrpcl.pPushConstantRanges = nullptr;

result = vkCreatePipelineLayout( LogicalDevice, IN &vrpcl, nullptr, OUT &PipelineLayout );

```

VkRayTracingPipelineCreateInfoKHR

```

vrpcl.sType = VK_STRUCTURE_TYPE_RAY_TRACING_PIPELINE_CREATE_INFO_KHR;
vrpcl.pNext = nullptr;
vrpcl.flags = 0;
vrpcl.stageCount = << # of shader stages in the ray-trace pipeline >>;
vrpcl.pStages = << what those shader stages are >>;
vrpcl.groupCount = << # of shader groups >>;
vrpcl.pGroups = << pointer to the groups (a group is a combination of shader programs) >>;
vrpcl.maxRecursionDepth = << how many recursion layers deep the ray tracing is allowed to go >>;
vrpcl.layout = PipelineLayout;
vrpcl.basePipelineHandle = VK_NULL_HANDLE;
vrpcl.basePipelineIndex = 0;

result = vkCreateRayTracingPipelinesKHR( LogicalDevice, PALLOCATOR, 1, IN &vrpcl, nullptr, OUT &RaytracePipeline );

```

 mjb - 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,
    hitShaderBindingTableBuffer, hitShaderBindingOffset, hitShaderBindingStride,
    callableShaderBindingTableBuffer, callableShaderBindingOffset, callableShaderBindingStride
    width, height, depth );

```

 mjb - July 24, 2020

