

Vulkan.
Introduction to the Vulkan Computer Graphics API
Mike Bailey
mjb@cs.oregonstate.edu
SIGGRAPH 2020 Abridged Version
<http://cs.oregonstate.edu/~mjb/vulkan>

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

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, understandable sample code

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

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>

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. Data 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. Multitpass
	25. Ray Tracing

Section titles that have been greyed-out have not been included in the ABRIDGED noteset, i.e., the one that has been made to fit in SIGGRAPH's reduced time slot. These topics are in the FULL noteset, however, which can be found on the web page:
<http://cs.oregonstate.edu/~mjb/vulkan>

My Favorite Vulkan Reference

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

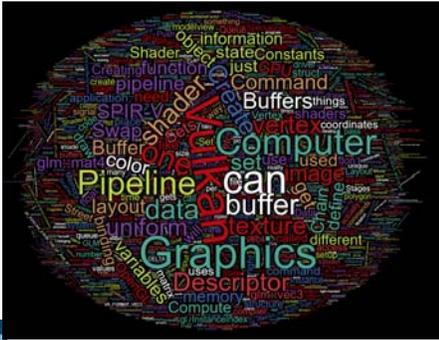
Vulkan.
Introduction

Mike Bailey
mjb@cs.oregonstate.edu

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



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




Top Three Reasons that Prompted the Development of Vulkan

1. Performance
2. Performance
3. Performance

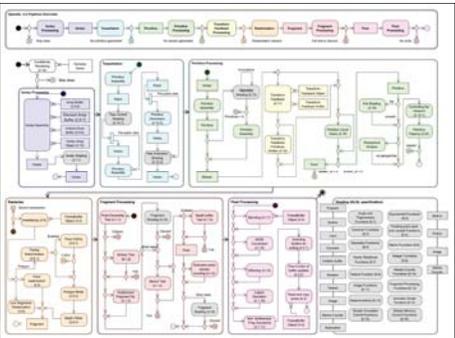
Vulkan is better at keeping the GPU busy than OpenGL is. OpenGL drivers need to do a lot of CPU work before handing work off to the GPU. Vulkan lets you get more power from the GPU card you already have.

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

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



OpenGL 4.2 Pipeline Flowchart




Who is the Khronos Group?

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




Playing "Where's Waldo?" with Khronos Membership




Who's Been Specifically Working on Vulkan?

SIGGRAPH 2018

Vulkan Differences from OpenGL

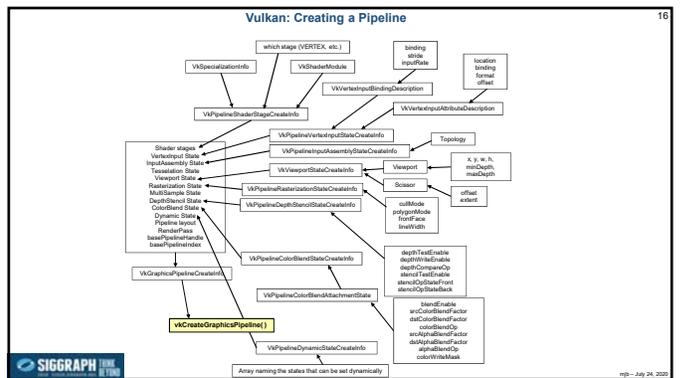
- More low-level information must be provided (by you!) in the application, rather than the driver
- Screen coordinate system is Y-down
- No "current state", at least not one maintained by the driver
- All of the things that we have talked about being *deprecated* in OpenGL are really **deprecated** in Vulkan: built-in pipeline transformations, begin-vertex-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.

SIGGRAPH 2018

Vulkan Highlights: Pipeline State Data Structure

- 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 2018



Querying the Number of Something

```

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

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

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

How many total there are	Where to put them
<code>&count</code>	<code>nullptr</code>
<code>&count</code>	<code>physicalDevices</code>

SIGGRAPH 2018

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

```

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

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

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

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

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

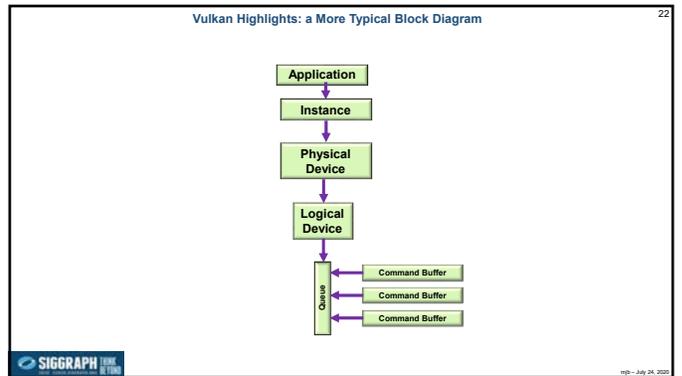
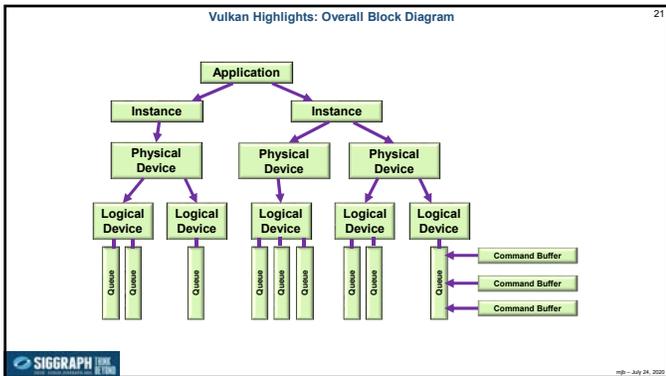
SIGGRAPH 2018

Vulkan Quick Reference Card – I Recommend you Print This!

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

Vulkan Quick Reference Card

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

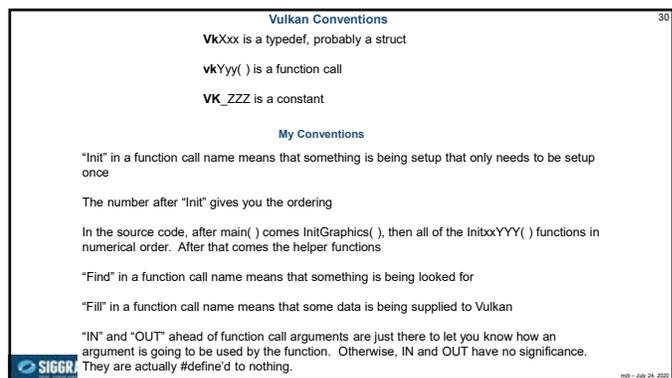
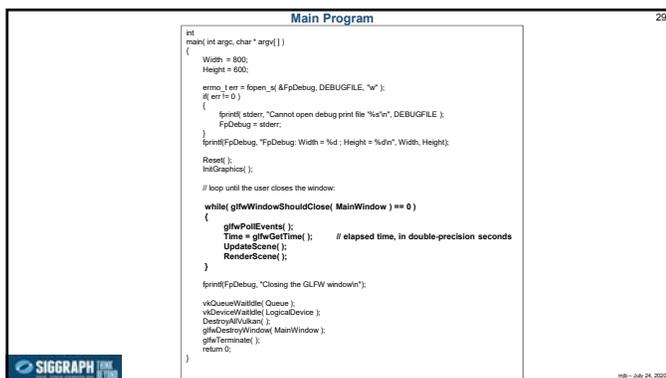
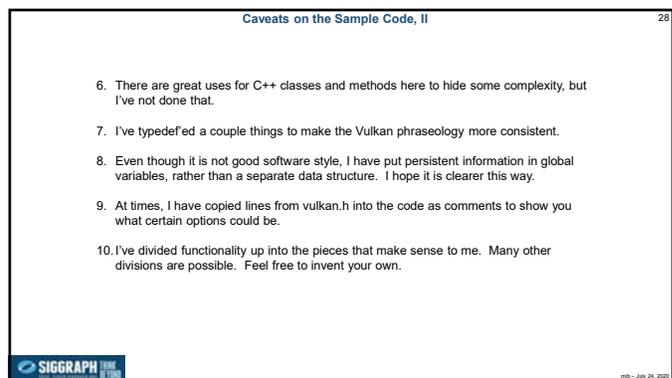
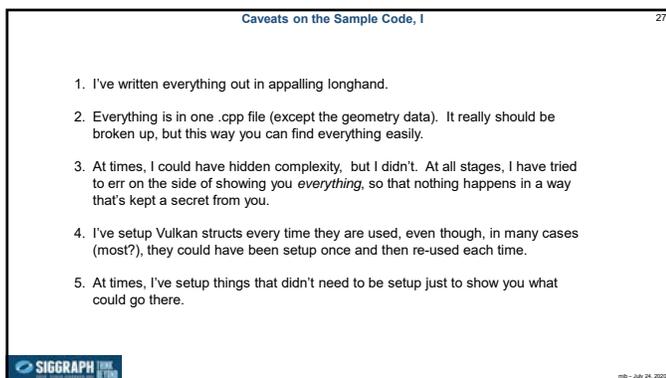
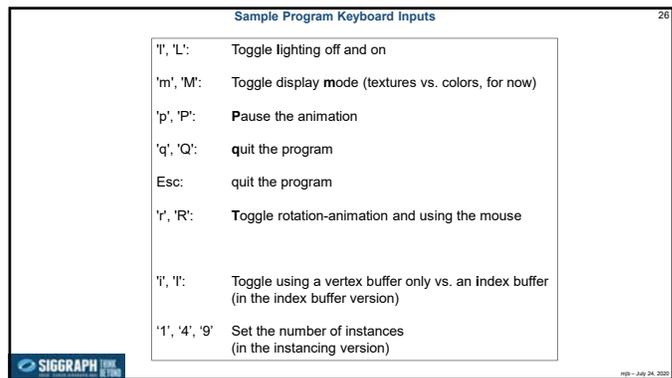
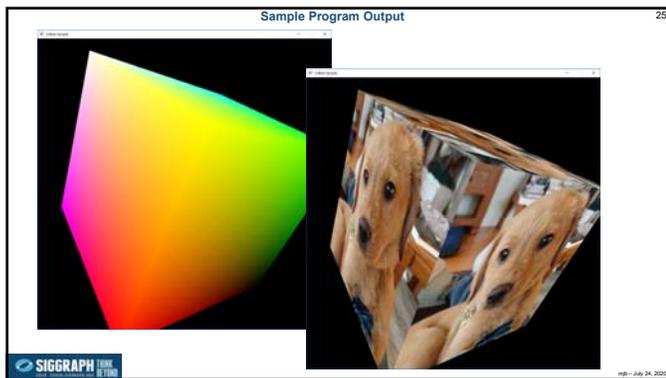


- Steps in Creating Graphics using Vulkan
1. Create the Vulkan Instance
 2. Setup the Debug Callbacks
 3. Create the Surface
 4. List the Physical Devices
 5. Pick the right Physical Device
 6. Create the Logical Device
 7. Create the Uniform Variable Buffers
 8. Create the Vertex Data Buffers
 9. Create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

The Vulkan Sample Code Included with These Notes

Mike Bailey
mjb@cs.oregonstate.edu

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



Your Sample2019.zip File Contains This

Linux shader compiler

Windows shader compiler

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

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

Reporting Error Results, I

```

struct errorcode
{
    VkResult resultCode;
    std::string meaning;
};

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

Reporting Error Results, II

```

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

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

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

Extras in the Code

```

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

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

bool Paused;

bool Verbose;

#define DEBUGFILE "VulkanDebug.txt"

errno_t err = fopen_s( &FpDebug, DEBUGFILE, "w" );

const int32_t OFFSET_ZERO = 0;
    
```

Vulkan.

Drawing

Mike Bailey
mjb@cs.oregonstate.edu

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

Vulkan Topologies

VK_PRIMITIVE_TOPOLOGY_POINT_LIST

VK_PRIMITIVE_TOPOLOGY_TRIANGLE_LIST

VK_PRIMITIVE_TOPOLOGY_LINE_LIST

VK_PRIMITIVE_TOPOLOGY_TRIANGLE_STRIP

VK_PRIMITIVE_TOPOLOGY_LINE_STRIP

VK_PRIMITIVE_TOPOLOGY_TRIANGLE_FAN

Vulkan Topologies

```

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

A Colored Cube Example

```

static float CubeColors[3][8] =
{
    { 0.0, 0.0, 0.0, 1.0, 0.0, 0.0, 0.0, 1.0 },
    { 0.0, 1.0, 0.0, 1.0, 1.0, 0.0, 1.0, 1.0 },
    { 0.0, 0.0, 1.0, 1.0, 0.0, 1.0, 0.0, 1.0 }
};

static GLuint CubeTriangleIndices[3][8] =
{
    { 0, 2, 3, 0, 3, 1, 4, 5, 7 },
    { 1, 2, 7, 1, 7, 5, 1, 7, 5 },
    { 0, 4, 6, 0, 6, 2, 2, 6, 7 },
    { 2, 7, 3, 2, 7, 3, 0, 1, 5 }
};
    
```

Triangles Represented as an Array of Structures

```

From the file SampleVertexData.cpp

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

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

Non-indexed Buffer Drawing

```

From the file SampleVertexData.cpp

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

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

// vertex #0:
{ -1.0, -1.0, 1.0, 0.0, 0.0, 0.0, 1.0, 0.0 },
// vertex #1:
{ 1.0, -1.0, 1.0, 0.0, 1.0, 0.0, 0.0, 1.0 },
// vertex #2:
{ -1.0, -1.0, 1.0, 0.0, 0.0, 1.0, 1.0, 1.0 },
// vertex #3:
{ 1.0, -1.0, 1.0, 0.0, 1.0, 0.0, 0.0, 1.0 },
// vertex #4:
{ 1.0, 1.0, 1.0, 0.0, 0.0, 0.0, 0.0, 1.0 },
// vertex #5:
{ 1.0, 1.0, 1.0, 0.0, 0.0, 1.0, 1.0, 1.0 },
// vertex #6:
{ -1.0, 1.0, 1.0, 0.0, 0.0, 0.0, 0.0, 1.0 },
// vertex #7:
{ -1.0, 1.0, 1.0, 0.0, 0.0, 1.0, 1.0, 1.0 }
};
    
```

Filling the Vertex Buffer

```

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;
}
    
```

A Preview of What Init05DataBuffer Does

```

VkResult
Init05DataBuffer( VkDeviceSize size, VkBufferUsageFlags usage, OUT MyBuffer * pMyBuffer )
{
    VkResult result = VK_SUCCESS;
    VkBufferCreateInfo vbcInfo = {
        .sType = VK_STRUCTURE_TYPE_BUFFER_CREATE_INFO,
        .pNext = nullptr,
        .flags = 0,
        .size = pMyBuffer->size,
        .usage = usage,
        .sharingMode = VK_SHARING_MODE_EXCLUSIVE,
        .queueFamilyIndexCount = 0,
        .queueFamilyIndices = nullptr,
        .result = VK_SUCCESS,
        .logicalDevice = LogicalDevice,
        .allocator = LogicalDevice->allocator
    };

    VkMemoryRequirements memReqs;
    vkGetBufferMemoryRequirements( LogicalDevice, vbcInfo.buffer, &memReqs );

    VkMemoryAllocateInfo vmaInfo = {
        .sType = VK_STRUCTURE_TYPE_MEMORY_ALLOCATE_INFO,
        .pNext = nullptr,
        .allocationSize = memReqs.size,
        .memoryTypeIndex = FindMemoryType( LogicalDevice, memReqs.memoryTypeBits )
    };

    VkDeviceMemory mem;
    result = vkAllocateMemory( LogicalDevice, vmaInfo, 0, &mem );
    pMyBuffer->vdm = mem;

    return result;
}
    
```

Telling the Pipeline about its Input

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

```

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

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

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

Telling the Pipeline about its Input

```

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

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

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

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

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

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

Telling the Pipeline about its Input

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

```

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

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

Telling the Pipeline about its Input

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

```

VkGraphicsPipelineCreateInfo vgpcci;
vgpcci.sType = VK_STRUCTURE_TYPE_GRAPHICS_PIPELINE_CREATE_INFO;
vgpcci.pNext = nullptr;
vgpcci.flags = 0;
vgpcci.stageCount = 2; // number of shader stages in this pipeline
vgpcci.pStages = &vpsscii;
vgpcci.pVertexInputState = &vpvscii;
vgpcci.pInputAssemblyState = &vpiascii;
vgpcci.pTessellationState = (VkPipelineTessellationStateCreateInfo*)nullptr; // &vptscii
vgpcci.pViewportState = &vpvscii;
vgpcci.pRasterizationState = &vprscii;
vgpcci.pMultisampleState = &vpmscii;
vgpcci.pDepthStencilState = &vpdscii;
vgpcci.pColorBlendState = &vpbcscii;
vgpcci.pDynamicState = &vpdscii;
vgpcci.layout = IN GraphicsPipelineLayout;
vgpcci.renderPass = IN RenderPass;
vgpcci.subpass = 0; // subpass number
vgpcci.basePipelineHandle = (VkPipeline)VK_NULL_HANDLE;
vgpcci.basePipelineIndex = 0;

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

Telling the Command Buffer what Vertices to Draw

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

```

VkBuffer buffers[1] = MyVertexDataBuffer.buffer;

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

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

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

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

Drawing with an Index Buffer

```

struct vertex JustVertexData[] =
{
    // vertex #0:
    { -1, -1, -1, 0, 0, -1, 0, 0, 0, 1, 0 },
    // vertex #1:
    { 1, -1, -1, 0, 0, -1, 1, 0, 0, 1, 0 },
    // ...
};

int JustIndexData[] =
{
    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,
};
  
```

Drawing with an Index Buffer 49

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

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

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



#00 - July 24, 2020

Drawing with an Index Buffer 50

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

```
Init05MyVertexBuffer( sizeof(JustVertexData), IN &MyJustVertexDataBuffer );
Fill05DataBuffer( MyJustVertexDataBuffer, (void *) JustVertexData );

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



#00 - July 24, 2020

Drawing with an Index Buffer 51

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

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

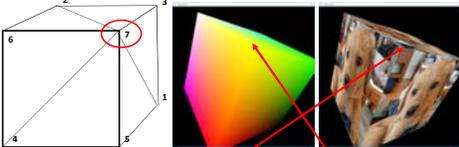
const uint32_t vertexCount = sizeof( JustVertexData ) / sizeof( JustVertexData[0] );
const uint32_t indexCount = sizeof( JustIndexData ) / sizeof( JustIndexData[0] );
const uint32_t instanceCount = 1;
const uint32_t firstVertex = 0;
const uint32_t firstIndex = 0;
const uint32_t firstInstance = 0;
const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

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



#00 - July 24, 2020

Sometimes the Same Point Needs Multiple Attributes 52



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

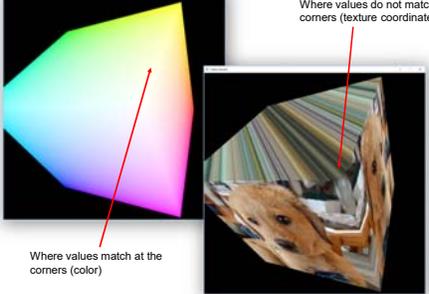
A color-interpolated cube like this actually has both. Point #7 above has the same color, regardless of what face it is in. However, Point #7 has 3 different normal vectors, depending on which face you are defining. Same with its texture coordinates.

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.



#00 - July 24, 2020

Sometimes the Same Point Needs Multiple Attributes 53



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

Where values match at the corners (color)



#00 - July 24, 2020

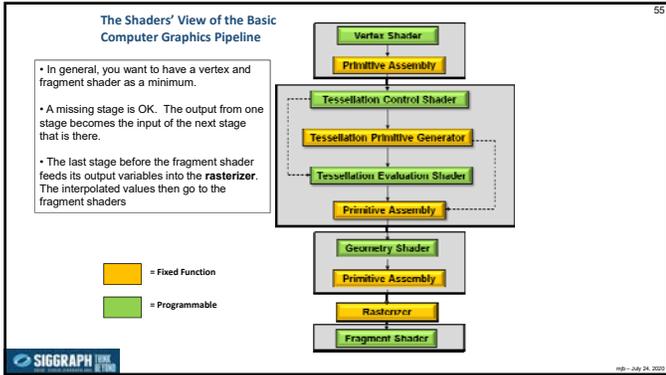
Vulkan.
Shaders and SPIR-V

Mike Bailey
mjb@cs.oregonstate.edu

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



#00 - July 24, 2020



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;
    
```

SIGGRAPH 2016 JULY 24, 2020

How Vulkan GLSL Differs from OpenGL GLSL

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

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

Vulkan Vertex and Instance indices:

```

gl_VertexIndex
gl_InstanceIndex
    
```

OpenGL uses:

```

gl_VertexID
gl_InstanceID
    
```

- Both are 0-based

gl_FragColor:

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

SIGGRAPH 2016 JULY 24, 2020

How Vulkan GLSL Differs from OpenGL GLSL

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

SIGGRAPH 2016 JULY 24, 2020

Vulkan: Shaders' use of Layouts for Uniform Variables

```

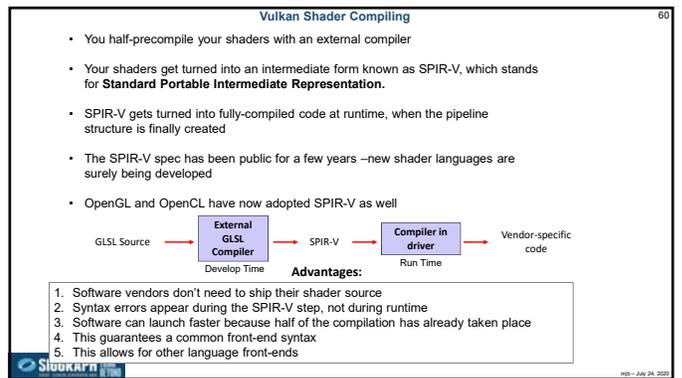
// non-sampler variables must be in a uniform block:
layout( std140, set = 0, binding = 0 ) uniform matBuf
{
    mat4 uModelMatrix;
    mat4 uViewMatrix;
    mat4 uProjectionMatrix;
    mat3 uNormalMatrix;
} Matrices;

// 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

SIGGRAPH 2016 JULY 24, 2020



SPIR-V:
Standard Portable Intermediate Representation for Vulkan

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(ies) to look in for #includes
- S Specify stage rather than get it from shaderfile extension
- c Print out the maximum sizes of various properties

Windows: glslangValidator.exe
Linux: glslangValidator

Running glslangValidator.exe

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

Compile for Vulkan ("-G" is compile for OpenGL) Specify the output file

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

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

Running glslangValidator.exe

```

MINGW64:/y/Vulkan/Sample2017
ONID+mjb@poooh MINGW64 /y/Vulkan/Sample2017
$ 185
glslangValidator.exe -V sample-vert.vert -o sample-vert.spv
sample-vert.vert
ONID+mjb@poooh MINGW64 /y/Vulkan/Sample2017
$ 186
glslangValidator.exe -V sample-frag.frag -o sample-frag.spv
sample-frag.frag
ONID+mjb@poooh MINGW64 /y/Vulkan/Sample2017
$
  
```

How do you know if SPIR-V compiled successfully?

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

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

So, if you do an **od -x** on the .spv file, the magic number looks like this:
0203 0723 ...

Reading a SPIR-V File into a Vulkan Shader Module

```

#define SPIRV_MAGIC 0x07230203
...
VkResult
Init2SPIRVShader( std::string filename, VkShaderModule * pShaderModule )
{
    FILE *fp;
    (void) fopen_s( &fp, filename.c_str(), "rb" );
    if( fp == NULL )
    {
        fprintf( FpDebug, "Cannot open shader file '%s'", filename.c_str() );
        return VK_SHOULD_EXIT;
    }
    uint32_t magic;
    fread( &magic, 4, 1, fp );
    if( magic != SPIRV_MAGIC )
    {
        fprintf( FpDebug, "Magic number for spir-v file '%s' is 0x%08x -- should be 0x%08xn",
                filename.c_str(), magic, SPIRV_MAGIC );
        return VK_SHOULD_EXIT;
    }
    fseek( fp, 0L, SEEK_END );
    int size = ftell( fp );
    rewind( fp );
    unsigned char *code = new unsigned char [size];
    fread( code, size, 1, fp );
    fclose( fp );
  }
  
```

Reading a SPIR-V File into a Shader Module

```

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

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


A Google-Wrapped Version of glslangValidator 73

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

```
glsic.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:


```
glsic.exe -target-env=vulkan -DNUMPONTs=4 sample-vert.vert -o sample-vert.spv
```

glsic is included in your Sample .zip file



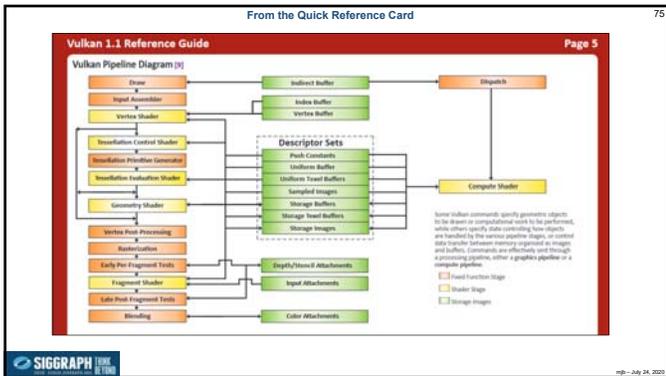

Vulkan.

Data Buffers

Mike Bailey
mjb@cs.oregonstate.edu

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





Terminology Issues 76

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".)

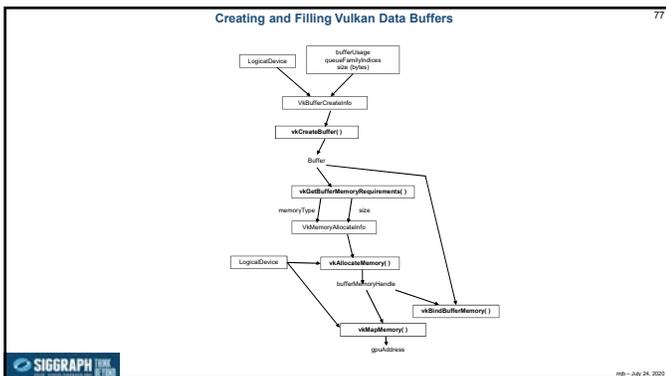
It is up to you to be sure that the writer and the reader of the Data Buffer are interpreting the bytes in the same way!

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.





Creating a Vulkan Data Buffer 78

```
VkBuffer Buffer;
VkBufferCreateInfo vbci;
vbci.sType = VK_STRUCTURE_TYPE_BUFFER_CREATE_INFO;
vbci.pNext = nullptr;
vbci.flags = 0;
vbci.size = << buffer size in bytes >>
vbci.usage = <<or'ed bits of: >>
VK_USAGE_TRANSFER_SRC_BIT
VK_USAGE_TRANSFER_DST_BIT
VK_USAGE_UNIFORM_TEXEL_BUFFER_BIT
VK_USAGE_STORAGE_TEXEL_BUFFER_BIT
VK_USAGE_UNIFORM_BUFFER_BIT
VK_USAGE_STORAGE_BUFFER_BIT
VK_USAGE_INDEX_BUFFER_BIT
VK_USAGE_VERTEX_BUFFER_BIT
VK_USAGE_INDIRECT_BUFFER_BIT
vbci.sharingMode = << one of: >>
VK_SHARING_MODE_EXCLUSIVE
VK_SHARING_MODE_CONCURRENT
vbci.queueFamilyIndexCount = 0;
vbci.queueFamilyIndices = (const int32_t) nullptr;
result = vkCreateBuffer ( LogicalDevice, IN &vbci, PALLOCATOR, OUT &Buffer );
```




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

```

VkMemoryRequirements
result = vkGetBufferMemoryRequirements( LogicalDevice, Buffer, 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, 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 );

```

Annotations: Red circles highlight `vmr`, `vmai`, and `vdm`. Red arrows point from these circles to their respective variables in the code.

Finding the Right Type of Memory

```

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;
}

```

Finding the Right Type of Memory

```

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

```

Finding the Right Type of Memory

```

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

Sidebar: The Vulkan Memory Allocator (VMA)

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.

Sidebar: The Vulkan Memory Allocator (VMA)

```

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

VmaAllocator var;
vmaCreateAllocator( IN &vac_i, OUT &var );
...

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

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

```

Something I've Found Useful

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.

Initializing a Data Buffer

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 )
{
    ...
    vbc.size = pMyBuffer->size = size;
    ...
    result = vkCreateBuffer ( LogicalDevice, IN &vbc, PALLOCATOR, OUT &pMyBuffer->buffer );
    ...
    pMyBuffer->vdm = vdm;
    ...
}
```

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

```
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;
```

Filling those Uniform Variables

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

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

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!

The Parade of Buffer Data

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 )
{
    ...
    vbc.size = pMyBuffer->size = size;
    ...
    result = vkCreateBuffer ( LogicalDevice, IN &vbc, PALLOCATOR, OUT &pMyBuffer->buffer );
    ...
    pMyBuffer->vdm = vdm;
    ...
}
```

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

```
struct matBuf Matrices;
```

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

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

Memory mapped copy operation

Filling the Data Buffer

```
void Init05UniformBuffer(
    const MyMatrixUniformBuffer &matrices,
    MyBuffer &myBuffer,
    ...
)
{
    ...
    Init05DataBuffer( sizeof(matrices), OUT &myBuffer );
    Fill05DataBuffer( MyMatrixUniformBuffer, IN (void *) &matrices );
    ...
}
```

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

Matrices.uModelMatrix = glm::mat4( ); // identity
Matrices.uViewMatrix = glm::lookAt( eye, look, up );

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

Matrices.uNormalMatrix = glm::inverseTranspose( glm::mat3( Matrices.uModelMatrix ) );
```

Creating and Filling the Data Buffer – the Details

```

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

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

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

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

    result = vkBindBufferMemory( LogicalDevice, pMyBuffer->buffer, IN vdm, OFFSET_ZERO );
    return result;
}
    
```

Creating and Filling the Data Buffer – the Details

```

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.



GLFW

Mike Bailey
mjb@cs.oregonstate.edu

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

Setting Up GLFW

```

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

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

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

You Can Also Query What Vulkan Extensions GLFW Requires

```

uint32_t count;
const char ** extensions = glfwGetRequiredInstanceExtensions( &count );
fprintf( FpDebug, "nFound %d GLFW Required Instance Extensions:\n", count );
for( uint32_t i = 0; i < count; i++ )
{
    fprintf( FpDebug, "%s\n", extensions[i] );
}
    
```

Found 2 GLFW Required Instance Extensions:
VK_KHR_surface
VK_KHR_win32_surface

GLFW Keyboard Callback

```

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

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

GLFW Mouse Button Callback

```

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

    // get the proper button bit mask:
    switch( button )
    {
        case GLFW_MOUSE_BUTTON_LEFT:   break;
        case GLFW_MOUSE_BUTTON_MIDDLE:  break;
        case GLFW_MOUSE_BUTTON_RIGHT:   break;

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

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

GLFW Mouse Motion Callback

```

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

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

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

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

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

Looping and Closing GLFW

```

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

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

Does not block – processes any waiting events, then returns

Looping and Closing GLFW

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( );
```



GLM

Mike Bailey
mjb@cs.oregonstate.edu

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

What is GLM?

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.

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

If GLM is not installed in a system place, put it somewhere you can get access to.

Why are we even talking about this?

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

```

glm::mat4x4 GL_MODELVIEW;
glm::loadIdentity();
gluLookAt( 0.0, 3.0, 0.0, 0.0, 0.0, 1.0, 0.0 );
glRotatef( (GLfloat)Yrot, 0.0, 1.0, 0.0 );
glRotatef( (GLfloat)Xrot, 1.0, 0.0, 0.0 );
glScalef( (GLfloat)Scale, (GLfloat)Scale, (GLfloat)Scale );

```

you would now say:

```

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

```

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

The Most Useful GLM Variables, Operations, and Functions

```

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

```

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

```

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

```

```

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

```

The Most Useful GLM Variables, Operations, and Functions

```

// viewing volume (assign, not concatenate):
glm::mat4 glm::ortho( float left, float right, float bottom, float top, float near, float far );
glm::mat4 glm::orthoL( 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 );

```

Installing GLM into your own space

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

Here's what that GLM folder looks like

GLM in the Vulkan sample.cpp Program

```

{ UseMouse }
{ Scale = MINSCALE }
Scale = MINSCALE;
Matrices.uModelMatrix = glm::mat4( 1. ); // identity
Matrices.uModelMatrix = glm::rotate( Matrices.uModelMatrix, Yrot, glm::vec3( 0.1, 0.0, 1.0 ) );
Matrices.uModelMatrix = glm::rotate( Matrices.uModelMatrix, Xrot, glm::vec3( 1.0, 0.0, 0.0 ) );
Matrices.uModelMatrix = glm::scale( Matrices.uModelMatrix, glm::vec3( Scale, Scale, Scale ) );
// Done this way, the Scale is applied first, then the Xrot, then the Yrot
}
else
{
{ I Paused }
const glm::vec3 axis = glm::vec3( 0.0, 1.0, 0.0 ); // (float)glm::radians( 300.F*PI*TIMESECONDS_PER_CYCLE ); axis );
Matrices.uModelMatrix = glm::rotate( glm::mat4( 1. ), (float)glm::radians( 300.F*PI*TIMESECONDS_PER_CYCLE ); axis );
}
}
glm::vec3 eye( 0.0, 0.0, EYEDIST );
glm::vec3 look( 0.0, 0.0, 0.0 );
glm::vec3 up( 0.0, 1.0, 0.0 );
Matrices.uViewMatrix = glm::lookAt( eye, look, up );

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

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

FIDSDatabuffer MyMatrUniformBuffer( void* ) &Matrices;

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

```

109

Sidebar: Why isn't The Normal Matrix exactly the same as the Model Matrix?

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

Original object and normal

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

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

Wrong!

Right!

SIGGRAPH

110

Vulkan.

Instancing

Mike Bailey
mjb@cs.oregonstate.edu

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

SIGGRAPH

111

Instancing – What and why?

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

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

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

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

SIGGRAPH

112

Making each Instance look differently -- Approach #1

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]
    
```

SIGGRAPH

113

SIGGRAPH

114

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]
    
```

SIGGRAPH

Vulkan.

The Graphics Pipeline Data Structure

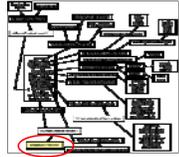
Mike Bailey
mjb@cs.oregonstate.edu

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



115

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.

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.




116

Graphics Pipeline Stages and what goes into Them

The GPU and Driver specify the Pipeline Stages – the Vulkan Graphics Pipeline declares what goes in them

Vertex Shader module Specialization info Vertex input binding Vertex input attributes	↓	Vertex Input Stage
Topology	↓	Input Assembly
Tessellation Shaders, Geometry Shader	↓	Tessellation, Geometry Shaders
Viewport Scissoring	↓	Viewport
Depth Clamping DiscardEnable PolygonMode CullMode FrontFace LineWidth	↓	Rasterization
Which states are dynamic	↓	Dynamic State
DepthTestEnable DepthWriteEnable DepthCompareOp StencilTestEnable	↓	Depth/Stencil
Fragment Shader module Specialization info	↓	Fragment Shader Stage
Color Blending parameters	↓	Color Blending Stage



117

The First Step: Create the Graphics Pipeline Layout

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

```

VkResult
Init4(GraphicsPipelineLayout)
{
    VkResult result;

    VkPipelineLayoutCreateInfo
    vpcli = { VK_STRUCTURE_TYPE_PIPELINE_LAYOUT_CREATE_INFO,
             vpcli.pNext = nullptr,
             vpcli.flags = 0,
             vpcli.setLayoutCount = 4,
             vpcli.pushConstantRanges = (VkPushConstantRange*)0 };

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

    return result;
}
    
```

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

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



118

A Pipeline Data Structure Contains the Following State Items:

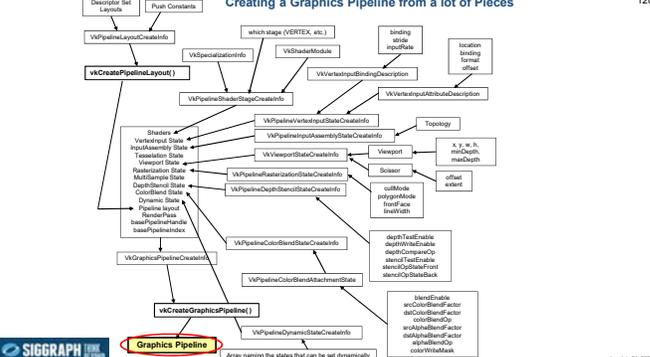
- Pipeline Layout: Descriptor Sets, Push Constants
- Which Shaders to use
- Per-vertex input attributes: location, binding, format, offset
- Per-vertex input bindings: binding, stride, inputRate
- Assembly: topology
- **Viewport**: x, y, w, h, minDepth, maxDepth
- **Scissoring**: x, y, w, h
- Rasterization: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

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



119

Creating a Graphics Pipeline from a lot of Pieces



Array naming the states that can be set dynamically



120

Creating a Typical Graphics Pipeline

```

VkResult
Init14GraphicsVertexFragmentPipeline( VkShaderModule vertexShader, VkShaderModule fragmentShader,
VkPrimitiveTopology topology, OUT VkPipeline *pGraphicsPipeline )
{
    #ifdef ASSUMPTIONS
        vvbld[0].inputRate = VK_VERTEX_INPUT_RATE_VERTEX;
        vpsscI.depthClampEnable = VK_FALSE;
        vpsscI.rasterizerDiscardEnable = VK_FALSE;
        vpsscI.polygonMode = VK_POLYGON_MODE_FILL;
        vpsscI.cullMode = VK_CULL_MODE_NONE; // best to do this because of the projectionMatrix[1][1] != -1;
        vpsscI.frontFace = VK_FRONT_FACE_COUNTER_CLOCKWISE;
        vpsscI.rasterizationSamples = VK_SAMPLE_COUNT_ONE_BIT;
        vpccbas.blendEnable = VK_FALSE;
        vpccbs.logicOpEnable = VK_FALSE;
        vpsscI.depthTestEnable = VK_TRUE;
        vpsscI.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.

The Shaders to Use

```

VkPipelineShaderStageCreateInfo
vpsscI( sType = VK_STRUCTURE_TYPE_PIPELINE_SHADER_STAGE_CREATE_INFO;
pStage = nullptr;
flags = 0;
stage = VK_SHADER_STAGE_VERTEX_BIT;
module = vertexShader;
pName = "main";
pSpecializationInfo = (VkSpecializationInfo *)nullptr;
)

VkShaderStageVertexBit
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

VkPipelineShaderStageCreateInfo
vpsscI( sType = VK_STRUCTURE_TYPE_PIPELINE_SHADER_STAGE_CREATE_INFO;
pStage = nullptr;
flags = 0;
stage = VK_SHADER_STAGE_FRAGMENT_BIT;
module = fragmentShader;
pName = "main";
pSpecializationInfo = (VkSpecializationInfo *)nullptr;
)

VkVertexInputBindingDescription
vbld( binding = 0; // which binding this is
stride = sizeof( struct vertex ); // bytes between successive
inputRate = VK_VERTEX_INPUT_RATE_VERTEX;
)

VkVertexInputAttributeDescription
vat( location = 0; // location in the layout
binding = 0; // which binding description this is part of
format = VK_FORMAT_VEC3; // s, r, g
offset = offsetOf( struct vertex, position ); // 0
)

```

Use one **vpsscI** array member per shader module you are using

Use one **vbld** array member per vertex input array-of-structures you are using

Link in the Per-Vertex Attributes

```

VkVertexInputAttributeDescription
vat( location = 0; // location in the layout
binding = 0; // which binding description this is part of
format = VK_FORMAT_VEC3; // s, r, g
offset = offsetOf( struct vertex, position ); // 0
)

// EXTRA: DEFINED AT THE TOP
// these are here for convenience and readability.
#define VK_FORMAT_VEC2 VK_FORMAT_R32G32B32A32_SFLOAT
#define VK_FORMAT_XYZW VK_FORMAT_R32G32B32A32_SFLOAT
#define VK_FORMAT_VEC3 VK_FORMAT_R32G32B32_SFLOAT
#define VK_FORMAT_STP VK_FORMAT_R32G32B32_SFLOAT
#define VK_FORMAT_XYZ VK_FORMAT_R32G32B32_SFLOAT
#define VK_FORMAT_VEC22 VK_FORMAT_R32G32_SFLOAT
#define VK_FORMAT_ST VK_FORMAT_R32G32_SFLOAT
#define VK_FORMAT_XYZ VK_FORMAT_R32G32B32_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

```

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

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

```

VkPipelineVertexInputStateCreateInfo
vpvsci( sType = VK_STRUCTURE_TYPE_PIPELINE_VERTEX_INPUT_STATE_CREATE_INFO;
pVertexBindingDescriptions = vpvsci;
pVertexAttributeDescriptions = vat;
)

VkPipelineInputAssemblyStateCreateInfo
vpascI( sType = VK_STRUCTURE_TYPE_PIPELINE_ASSEMBLY_STATE_CREATE_INFO;
pTopology = VK_PRIMITIVE_TOPOLOGY_TRIANGLE_LIST;
)

// CHANGES
VK_PRIMITIVE_TOPOLOGY_POINT_LIST
VK_PRIMITIVE_TOPOLOGY_LINE_LIST
VK_PRIMITIVE_TOPOLOGY_TRIANGLE_LIST
VK_PRIMITIVE_TOPOLOGY_TRIANGLE_FAN
VK_PRIMITIVE_TOPOLOGY_TRIANGLE_STRIP
VK_PRIMITIVE_TOPOLOGY_TRIANGLE_STRIP_WITH_ADJACENCY
VK_PRIMITIVE_TOPOLOGY_LINE_STRIP_WITH_ADJACENCY
VK_PRIMITIVE_TOPOLOGY_TRIANGLE_LIST_WITH_ADJACENCY
VK_PRIMITIVE_TOPOLOGY_TRIANGLE_STRIP_WITH_ADJACENCY

```

Declare the binding descriptions and attribute descriptions

Declare the vertex topology

vpsscI Tessellation Shader info

vpsscI Geometry Shader info

Options for vpascI.topology

- VK_PRIMITIVE_TOPOLOGY_POINT_LIST**: Shows four isolated vertices labeled V0, V1, V2, V3.
- VK_PRIMITIVE_TOPOLOGY_TRIANGLE_LIST**: Shows two separate triangles. The first has vertices V0, V1, V2; the second has V3, V4, V5.
- VK_PRIMITIVE_TOPOLOGY_TRIANGLE_STRIP**: Shows a continuous strip of triangles. The first triangle has vertices V0, V1, V2; the second has V1, V2, V3; the third has V2, V3, V4.
- VK_PRIMITIVE_TOPOLOGY_LINE_LIST**: Shows two separate line segments. The first has vertices V0, V1; the second has V2, V3.
- VK_PRIMITIVE_TOPOLOGY_LINE_STRIP**: Shows a continuous strip of line segments. The first has vertices V0, V1; the second has V1, V2; the third has V2, V3.
- VK_PRIMITIVE_TOPOLOGY_TRIANGLE_FAN**: Shows a fan of triangles sharing a common vertex V0. The first triangle has vertices V0, V1, V2; the second has V0, V2, V3; the third has V0, V3, V4.

What is "Primitive Restart Enable"?

vpascI.primitiveRestartEnable = VK_FALSE;

"Restart Enable" is used with:

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

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

```

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

```

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

127

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

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

SIGGRAPH

128

```

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

VkRect2D
  vr.offset.x = 0;
  vr.offset.y = 0;
  vr.extent.width = Width;
  vr.extent.height = Height;

VkPipelineViewportStateCreateInfo
  vpsc.iType = VK_STRUCTURE_TYPE_PIPELINE_VIEWPORT_STATE_CREATE_INFO;
  vpsc.pNext = nullptr;
  vpsc.flags = 0;
  vpsc.viewportCount = 1;
  vpsc.pViewports = &vr;
  vpsc.scissorCount = 1;
  vpsc.pScissors = &vr;
  
```

Declare the viewport information

Declare the scissoring information

Group the viewport and scissor information together

SIGGRAPH

129

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

Viewport
Viewporing 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.

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.

Original Image

SIGGRAPH

130

Setting the Rasterizer State

```

VkPipelineRasterizationStateCreateInfo
  vprsc.iType = VK_STRUCTURE_TYPE_PIPELINE_RASTERIZATION_STATE_CREATE_INFO;
  vprsc.pNext = nullptr;
  vprsc.flags = 0;
  vprsc.depthClampEnable = VK_FALSE;
  vprsc.rasterizerDiscardEnable = VK_FALSE;
  vprsc.polygonMode = VK_POLYGON_MODE_FILL;

  #if defined CHOICES
  VK_POLYGON_MODE_FILL
  VK_POLYGON_MODE_LINE
  VK_POLYGON_MODE_POINT
  #endif

  vprsc.cullMode = VK_CULL_MODE_NONE; // recommend this because of the projMatrix[1][1] * -

  1.;
  #if defined 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;
  #if defined CHOICES
  VK_FRONT_FACE_COUNTER_CLOCKWISE
  VK_FRONT_FACE_CLOCKWISE
  #endif

  vprsc.depthBiasEnable = VK_FALSE;
  vprsc.depthBiasConstantFactor = 0.0;
  vprsc.depthBiasClamp = 0.0;
  vprsc.depthBiasSlopeFactor = 0.0;
  vprsc.lineWidth = 1.0;
  
```

Declare information about how the rasterization will take place

SIGGRAPH

131

What is "Depth Clamp Enable"?

```

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 happen with depthClampEnable (except it would have been red).

SIGGRAPH

132

What is "Depth Bias Enable"?

```

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

MultiSampling State

```

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

vpmisci

Declare information about how the multisampling will take place

We will discuss MultiSampling in a separate noteset.

Color Blending State for each Color Attachment *

```

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;
    
```

vpcbas

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

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

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

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

Raster Operations for each Color Attachment

```

VkPipelineColorBlendStateCreateInfo
vpcbsci.sType = VK_STRUCTURE_TYPE_PIPELINE_COLOR_BLEND_STATE_CREATE_INFO;
vpcbsci.pNext = nullptr;
vpcbsci.flags = 0;
vpcbsci.logicOpEnable = VK_FALSE;
vpcbsci.logicOp = VK_LOGIC_OP_COPY;
//ifdef CHOICES
VK_LOGIC_OP_CLEAR
VK_LOGIC_OP_AND
VK_LOGIC_OP_AND_REVERSE
VK_LOGIC_OP_COPY
VK_LOGIC_OP_AND_INVERTED
VK_LOGIC_OP_NO_OP
VK_LOGIC_OP_XOR
VK_LOGIC_OP_OR
VK_LOGIC_OP_NOR
VK_LOGIC_OP_EQUIVALENT
VK_LOGIC_OP_INVERT
VK_LOGIC_OP_OR_REVERSE
VK_LOGIC_OP_COPY_INVERTED
VK_LOGIC_OP_OR_INVERTED
VK_LOGIC_OP_NAND
VK_LOGIC_OP_SET
//endif
vpcbsci.attachmentCount = 1;
vpcbsci.pAttachments = &vpcbas;
vpcbsci.blendConstants[0] = 0;
vpcbsci.blendConstants[1] = 0;
vpcbsci.blendConstants[2] = 0;
vpcbsci.blendConstants[3] = 0;
    
```

vpcbsci

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

Which Pipeline Variables can be Set Dynamically

```

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
    
```

Just used as an example in the Sample Code

```

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

vpdsci

The Stencil Buffer

Update

Depth

Stencil

Render

Here's how the Stencil Buffer works:

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

Using the Stencil Buffer to Create a Magic Lens

Using the Stencil Buffer to Create a Magic Lens

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

SIGGRAPH 2014

Outlining Polygons the Naïve Way

1. Draw the polygons
2. Draw the edges

Z-fighting

SIGGRAPH 2014

Using the Stencil Buffer to Better Outline Polygons

SIGGRAPH 2014

Using the Stencil Buffer to Better Outline Polygons

```

Clear the SB = 0
for( each polygon )
{
    Draw the edges, setting SB = 1
    Draw the polygon wherever SB != 1
    Draw the edges, setting SB = 0
}
    
```

Before

After

SIGGRAPH 2014

Using the Stencil Buffer to Perform Hidden Line Removal

SIGGRAPH 2014

Stencil Operations for Front and Back Faces

```

VK_STENCIL_OP_STATE // front
vsodp.depthFailOp = VK_STENCIL_OP_KEEP // what to do if depth operation fails
vsodp.failOp = VK_STENCIL_OP_KEEP // what to do if stencil operation fails
vsodp.passOp = VK_STENCIL_OP_KEEP // what to do if stencil operation succeeds

#ifdef 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_INCREMENT_AND_CLAMP // increment stencil value
VK_STENCIL_OP_DECREMENT_AND_CLAMP // decrement stencil value
VK_STENCIL_OP_INVERT // invert stencil value
VK_STENCIL_OP_INCREMENT_AND_WRAP // increment stencil value
VK_STENCIL_OP_DECREMENT_AND_WRAP // decrement stencil value
#endif

vsodp.compareOp = VK_COMPARE_OP_NEVER;

#ifdef CHOICES
VK_COMPARE_OP_NEVER // never succeeds
VK_COMPARE_OP_LESS // succeeds if stencil value is < the reference value
VK_COMPARE_OP_EQUAL // succeeds if stencil value is == the reference value
VK_COMPARE_OP_LESS_OR_EQUAL // succeeds if stencil value is <= the reference value
VK_COMPARE_OP_GREATER // succeeds if stencil value is > the reference value
VK_COMPARE_OP_NOT_EQUAL // succeeds if stencil value is != the reference value
VK_COMPARE_OP_GREATER_OR_EQUAL // succeeds if stencil value is >= the reference value
VK_COMPARE_OP_ALWAYS // always succeeds
#endif

vsodp.compareMask = 0;
vsodp.writeMask = 0;
vsodp.reference = 0;

VK_STENCIL_OP_STATE // back
vsodp.depthFailOp = VK_STENCIL_OP_KEEP // front
vsodp.failOp = VK_STENCIL_OP_KEEP // front
vsodp.passOp = VK_STENCIL_OP_KEEP // front
vsodp.compareOp = VK_COMPARE_OP_NEVER;
vsodp.compareMask = 0;
vsodp.writeMask = 0;
vsodp.reference = 0;
    
```

SIGGRAPH 2014

Operations for Depth Values 145

```

VKPipelineDepthStencilStateCreateInfo
    vpdssciType = VK_STRUCTURE_TYPE_PIPELINE_DEPTH_STENCIL_STATE_CREATE_INFO,
    vpdsscipNext = nullptr,
    vpdssciFlags = 0,
    vpdssciDepthTestEnable = VK_TRUE,
    vpdssciDepthWriteEnable = VK_TRUE,
    vpdssciDepthCompareOp = 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
    vpdssciDepthBoundsTestEnable = VK_FALSE;
    vpdssciFront = vpsci;
    vpdssciBack = vpsci;
    vpdssciMinDepthBounds = 0.;
    vpdssciMaxDepthBounds = 1.;
    vpdssciStencilTestEnable = VK_FALSE;
    #endif
    
```

Putting it all Together! (finally...) 146

```

VkPipeline GraphicsPipeline;
VkGraphicsPipelineCreateInfo
    vgpcliType = VK_STRUCTURE_TYPE_GRAPHICS_PIPELINE_CREATE_INFO,
    vgpclipNext = nullptr,
    vgpcliFlags = 0,
    #if defined(CHOICES)
    VK_PIPELINE_CREATE_DISABLE_OPTIMIZATION_BIT
    VK_PIPELINE_CREATE_ALLOW_DERIVATIVES_BIT
    VK_PIPELINE_CREATE_DERIVATIVE_BIT
    #endif
    vgpcliStageCount = 2; // number of stages in this pipeline
    vgpcliPStages = vpsci;
    vgpcliPVertexInputState = &vpsci;
    vgpcliPInputAssemblyState = &vpsci;
    vgpcliPTessellationState = (VkPipelineTessellationStateCreateInfo*) nullptr;
    vgpcliPViewportState = &vpsci;
    vgpcliPRasterizationState = &vpsci;
    vgpcliPMultisampleState = &vpsci;
    vgpcliPDepthStencilState = &vpdssci;
    vgpcliPColorBlendState = &vpsci;
    vgpcliPDynamicState = &vpsci;
    vgpcliLayout = IN GraphicsPipelineLayout;
    vgpcliRenderPass = IN RenderPass;
    vgpcliSubpass = 0; // subpass number
    vgpcliBasePipelineHandle = (VkPipeline) VK_NULL_HANDLE;
    vgpcliBasePipelineIndex = 0;

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

return result;
    
```

Later on, we will Bind a Specific Graphics Pipeline Data Structure to the Command Buffer when Drawing 147

```

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

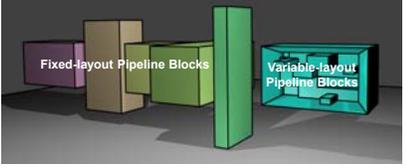
Sidebar: What is the Organization of the Pipeline Data Structure? 148

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;
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:



In OpenGL 150

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;
    
```

Vulkan.

Descriptor Sets

Mike Bailey
mjb@cs.oregonstate.edu

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

In OpenGL 150

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;
    
```

25

What are Descriptor Sets?

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

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

```

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

```

Descriptor Sets

Our example will assume the following shader uniform variables:

```

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

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

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

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

```

Descriptor Sets

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;
};

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

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

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

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

```

* "binary large object"

Step 1: Descriptor Set Pools

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

```

graph TD
  flags --> vkDescriptorPoolCreateInfo
  matSets --> vkDescriptorPoolCreateInfo
  poolSizeCount --> vkDescriptorPoolCreateInfo
  poolSizes --> vkDescriptorPoolCreateInfo
  device --> vkCreateDescriptorPool
  vkDescriptorPoolCreateInfo --> vkCreateDescriptorPool
  vkCreateDescriptorPool --> DescriptorSetPool

```

```

VkResult vkCreateDescriptorPool(
VkDevice device,
const VkDescriptorPoolCreateInfo* pCreateInfo,
VkDescriptorPool* pPool)
{
  VkResult result;

  result = vkDeviceWaitIdle(device);
  if (result != VK_SUCCESS) return result;

  result = vkCreateDescriptorPool(device, pCreateInfo, pPool);
  if (result != VK_SUCCESS) return result;

  return result;
}

```

vkDescriptorPoolCreateInfo

vkDescriptorPool

Step 2: Define the Descriptor Set Layouts

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

```

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;

```

MatrixSet DS Layout Binding: binding descriptorType descriptorCount pipeline stage(s) set = 0

LightSet DS Layout Binding: binding descriptorType descriptorCount pipeline stage(s) set = 1

MiscSet DS Layout Binding: binding descriptorType descriptorCount pipeline stage(s) set = 2

TexSamplerSet DS Layout Binding: binding descriptorType descriptorCount pipeline stage(s) set = 3

```

VkResult
Init13DescriptorSetLayouts()
{
    VkResult result;

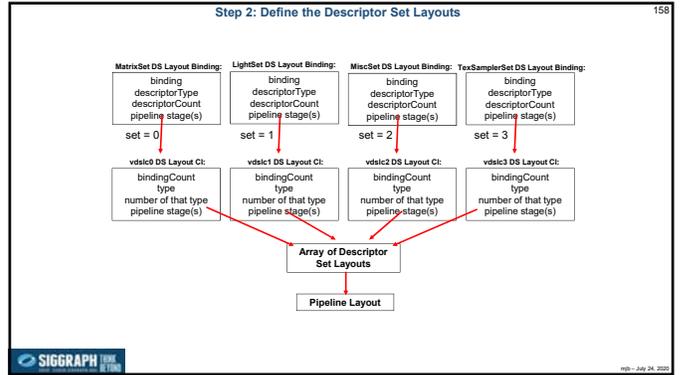
    // DS #0:
    VkDescriptorSetLayoutBinding
    MatrixSet[0] binding = 0;
    MatrixSet[0].descriptorType = VK_DESCRIPTOR_TYPE_UNIFORM_BUFFER;
    MatrixSet[0].descriptorCount = 1;
    MatrixSet[0].stageFlags = VK_SHADER_STAGE_VERTEX_BIT;
    MatrixSet[0].pImmutableSamplers = (VkSampler *)nullptr;

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

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

    // DS #3:
    VkDescriptorSetLayoutBinding
    TexSamplerSet[0] binding = 0;
    TexSamplerSet[0].descriptorType = VK_DESCRIPTOR_TYPE_COMBINED_IMAGE_SAMPLER;
    TexSamplerSet[0].descriptorCount = 1;
    TexSamplerSet[0].stageFlags = VK_SHADER_STAGE_FRAGMENT_BIT;
    TexSamplerSet[0].pImmutableSamplers = (VkSampler *)nullptr;
    uniform sampler2D uSampler;
    vec3 rgb = texture( uSampler, vST );
}

```



```

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

VkDescriptorSetLayoutCreateInfo
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.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;
}

```

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

```

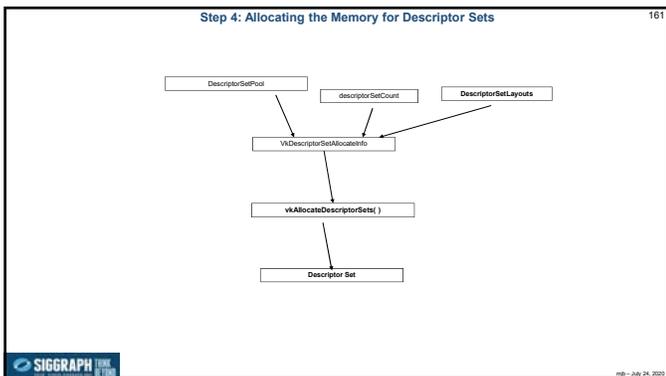
VkResult
Init14GraphicsPipelineLayout()
{
    VkResult result;

    VkPipelineLayoutCreateInfo
    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;
}

```



Step 4: Allocating the Memory for Descriptor Sets

```

VkResult
Init13DescriptorSets()
{
    VkResult result;

    VkDescriptorSetAllocateInfo
    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] );
}

```

Step 5: Tell the Descriptor Sets where their CPU Data is 163

```

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 vdi0;
vdi0.sampler = MyPuppyTexture.texSampler;
vdi0.imageView = MyPuppyTexture.texImageView;
vdi0.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

Step 5: Tell the Descriptor Sets where their CPU Data is 164

```

VkWriteDescriptorSet vwd0;
// ds 0
vwd0.sType = VK_STRUCTURE_TYPE_WRITE_DESCRIPTOR_SET;
vwd0.pNext = nullptr;
vwd0.dstSet = DescriptorSets[0];
vwd0.dstBinding = 0;
vwd0.dstArrayElement = 0;
vwd0.descriptorCount = 1;
vwd0.descriptorType = VK_DESCRIPTOR_TYPE_UNIFORM_BUFFER;
vwd0.pBufferInfo = IN &vdbi0;
vwd0.pImageInfo = (VkDescriptorImageInfo *)nullptr;
vwd0.pTexelBufferView = (VkBufferView *)nullptr;

// ds 1
VkWriteDescriptorSet vwd1;
vwd1.sType = VK_STRUCTURE_TYPE_WRITE_DESCRIPTOR_SET;
vwd1.pNext = nullptr;
vwd1.dstSet = DescriptorSets[1];
vwd1.dstBinding = 0;
vwd1.dstArrayElement = 0;
vwd1.descriptorCount = 1;
vwd1.descriptorType = VK_DESCRIPTOR_TYPE_UNIFORM_BUFFER;
vwd1.pBufferInfo = IN &vdbi1;
vwd1.pImageInfo = (VkDescriptorImageInfo *)nullptr;
vwd1.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

Step 5: Tell the Descriptor Sets where their data is 165

```

VkWriteDescriptorSet vwd2;
// ds 2
vwd2.sType = VK_STRUCTURE_TYPE_WRITE_DESCRIPTOR_SET;
vwd2.pNext = nullptr;
vwd2.dstSet = DescriptorSets[2];
vwd2.dstBinding = 0;
vwd2.dstArrayElement = 0;
vwd2.descriptorCount = 1;
vwd2.descriptorType = VK_DESCRIPTOR_TYPE_UNIFORM_BUFFER;
vwd2.pBufferInfo = IN &vdbi2;
vwd2.pImageInfo = (VkDescriptorImageInfo *)nullptr;
vwd2.pTexelBufferView = (VkBufferView *)nullptr;

// ds 3
VkWriteDescriptorSet vwd3;
vwd3.sType = VK_STRUCTURE_TYPE_WRITE_DESCRIPTOR_SET;
vwd3.pNext = nullptr;
vwd3.dstSet = DescriptorSets[3];
vwd3.dstBinding = 0;
vwd3.dstArrayElement = 0;
vwd3.descriptorCount = 1;
vwd3.descriptorType = VK_DESCRIPTOR_TYPE_COMBINED_IMAGE_SAMPLER;
vwd3.pBufferInfo = (VkDescriptorBufferInfo *)nullptr;
vwd3.pImageInfo = IN &vdi0;
vwd3.pTexelBufferView = (VkBufferView *)nullptr;

uint32_t copyCount = 0;

// this could have been done with one call and an array of VkWriteDescriptorSet:
vkUpdateDescriptorSets(LogicalDevice, 1, IN &vwd0, IN copyCount, (VkCopyDescriptorSet *)nullptr);
vkUpdateDescriptorSets(LogicalDevice, 1, IN &vwd1, IN copyCount, (VkCopyDescriptorSet *)nullptr);
vkUpdateDescriptorSets(LogicalDevice, 1, IN &vwd2, IN copyCount, (VkCopyDescriptorSet *)nullptr);
vkUpdateDescriptorSets(LogicalDevice, 1, IN &vwd3, IN copyCount, (VkCopyDescriptorSet *)nullptr);

```

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

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

Step 6: Include the Descriptor Set Layout when Creating a Graphics Pipeline 166

```

VkGraphicsPipelineCreateInfo vgpcli;
vgpcli.sType = VK_STRUCTURE_TYPE_GRAPHICS_PIPELINE_CREATE_INFO;
vgpcli.pNext = nullptr;
vgpcli.flags = 0;

// shader CHOICES
VK_PIPELINE_CREATE_DISABLE_OPTIMIZATION_BIT
VK_PIPELINE_CREATE_ALLOW_DERIVATIVES_BIT
VK_PIPELINE_CREATE_DERIVATIVE_BIT
uint32_t derivCnt;

vgpcli.stageCount = 2; // number of stages in this pipeline
vgpcli.pStages = vpsci;
vgpcli.pVertexInputState = &vpi;
vgpcli.pInputAssemblyState = &vpasci;
vgpcli.pTessellationState = (VkPipelineTessellationStateCreateInfo *)nullptr;
vgpcli.pViewportState = &vpvsci;
vgpcli.pRasterizationState = &vpasrci;
vgpcli.pMultisampleState = &vpmsci;
vgpcli.pDepthStencilState = &vpdscsi;
vgpcli.pColorBlendState = &vpbcsci;
vgpcli.pDynamicState = &vpdsci;
vgpcli.layout = &vplscli; // GraphicsPipelineLayout
vgpcli.renderPass = IN RenderPass;
vgpcli.subpass = 0; // subpass number
vgpcli.basePipelineHandle = (VkPipeline)VK_NULL_HANDLE;
vgpcli.basePipelineIndex = 0;

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

```

vkCreateGraphicsPipelines

Step 7: Bind Descriptor Sets into the Command Buffer when Drawing 167

```

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.

Sidebar: The Entire Collection of Descriptor Set Paths 168

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

169

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

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

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

The Pipeline Data Structure

SIGGRAPH THINK AT THE EDGE

mjb - July 24, 2020

170

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

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

SIGGRAPH THINK AT THE EDGE

mjb - July 24, 2020

171

Textures

Mike Bailey
mjb@cs.oregonstate.edu

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

SIGGRAPH THINK AT THE EDGE

mjb - July 24, 2020

172

Triangles in an Array of Structures

```

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, 0, 0,
          { 1, 1 }
        }
      },
    // vertex #3:
    { 1, -1, -1,
      { 0, 0, -1,
        { 0, 0, 0,
          { 0, 1 }
        }
      }
    };
    
```

SIGGRAPH THINK AT THE EDGE

mjb - July 24, 2020

173

Memory Types

SIGGRAPH THINK AT THE EDGE

mjb - July 24, 2020

174

Memory Types

NVIDIA Discrete Graphics:

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

Intel Integrated Graphics:

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

SIGGRAPH THINK AT THE EDGE

mjb - July 24, 2020

187

```

// create an image view for the texture image:
// (an "image view" is used to indirectly access an image)
VkImageSubresourceRange
    vkrange{baseMipLevel = 0,
            baseLayerCount = 1,
            baseArrayLayer = 0,
            layerCount = 1;

VkImageCreateInfo
    vici{stType = VK_STRUCTURE_TYPE_IMAGE_VIEW_CREATE_INFO,
        pNext = nullptr,
         flags = 0,
         image = textureImage,
         viewType = VK_IMAGE_VIEW_TYPE_2D,
         format = VK_FORMAT_R8G8B8A8_UNORM,
         components = {VK_COMPONENT_SWIZZLE_R,
                       VK_COMPONENT_SWIZZLE_G,
                       VK_COMPONENT_SWIZZLE_B,
                       VK_COMPONENT_SWIZZLE_A},
         subresourceRange = vkrange;

result = vkCreateImageView( LogicalDevice, IN &vici, PALLOCATOR_OUT &MyTexture->imageView);
return result;
    
```

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

SIGGRAPH 2019 JULY 24, 2020

188

Reading in a Texture from a BMP File

```

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

...
MyTexture MyPuppyTexture;
    
```

```

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

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

SIGGRAPH 2019 JULY 24, 2020

189

Queues and Command Buffers

Mike Bailey
mjb@cs.oregonstate.edu

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

SIGGRAPH 2019 JULY 24, 2020

190

Simplified Block Diagram

```

graph TD
    Application --> Instance
    Instance --> PhysicalDevice
    PhysicalDevice --> LogicalDevice
    LogicalDevice --> Queue
    Queue --> CB1[Command Buffer]
    Queue --> CB2[Command Buffer]
    Queue --> CB3[Command Buffer]
    
```

SIGGRAPH 2019 JULY 24, 2020

191

Vulkan Queues and Command Buffers

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

SIGGRAPH 2019 JULY 24, 2020

192

Querying what Queue Families are Available

```

uint32_t count;
VkGetPhysicalDeviceQueueFamilyProperties( IN PhysicalDevice, &count, OUT (VkQueueFamilyProperties *) nullptr );
VkQueueFamilyProperties *vqfp = new VkQueueFamilyProperties[ count ];
VkGetPhysicalDeviceFamilyProperties( PhysicalDevice, &count, OUT &vqfp );
for( unsigned int i = 0; i < count; i++ )
{
    fprintf( FpDebug, "%d: Queue Family Count = %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" );
}
    
```

Found 3 Queue Families:

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

SIGGRAPH 2019 JULY 24, 2020

Similarly, we Can Write a Function that Finds the Proper Queue Family 193

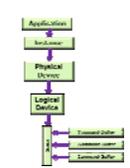
```

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

    VkQueueFamilyProperties vqfp;
    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;
}

```



SIGGRAPH 2020

Creating a Logical Device Needs to Know Queue Family Information 194

```

for( QueuePriorities )
{
    1. // one entry per queueCount

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

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

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

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

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

```

SIGGRAPH 2020

Creating the Command Pool as part of the Logical Device 195

```

VkResult
InitCommandPool()
{
    VkResult result;

    VkCommandPoolCreateInfo vcpci;
    vcpci.sType = VK_STRUCTURE_TYPE_COMMAND_POOL_CREATE_INFO;
    vcpci.pNext = nullptr;
    vcpci.flags = VK_COMMAND_POOL_CREATE_RESET_COMMAND_BUFFER_BIT;

    #ifdef CHOICES
    vcpci.queueFamilyIndex = FindQueueFamilyThatDoesGraphics();
    #endif

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

    return result;
}

```

SIGGRAPH 2020

Creating the Command Buffers 196

```

VkResult
InitCommandBuffers()
{
    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:
    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;
}

```

SIGGRAPH 2020

Beginning a Command Buffer – One per Image 197

```

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] );

```

SIGGRAPH 2020

Beginning a Command Buffer 198

```

graph TD
    A[VKCommandBufferPoolCreateInfo] --> B[vkCreateCommandBufferPool()]
    B --> C[VKCommandBufferAllocateInfo]
    C --> D[vkAllocateCommandBuffers()]
    D --> E[VKCommandBufferBeginInfo]
    E --> F[vkBeginCommandBuffer()]

```

SIGGRAPH 2020

These are the Commands that could be entered into the Command Buffer, I 199

```

VkCmdBeginQuery( commandBuffer, flags );
VkCmdBeginRenderPass( commandBuffer, const contents );
VkCmdBindDescriptorSets( commandBuffer, pDynamicOffsets );
VkCmdBindIndexBuffer( commandBuffer, indexType );
VkCmdBindPipeline( commandBuffer, pipeline );
VkCmdBindVertexBuffers( commandBuffer, firstBinding, bindingCount, const pOffsets );
VkCmdBindImage( commandBuffer, filter );
VkCmdClearAttachments( commandBuffer, attachmentCount, const pRects );
VkCmdClearColorImage( commandBuffer, pRanges );
VkCmdClearDepthStencilImage( commandBuffer, pRanges );
VkCmdCopyBuffer( commandBuffer, pRegions );
VkCmdCopyBufferToImage( commandBuffer, pRegions );
VkCmdCopyImage( commandBuffer, pRegions );
VkCmdCopyImageToBuffer( commandBuffer, pRegions );
VkCmdCopyQueryPoolResults( commandBuffer, flags );
VkCmdDebugMarkerBeginEXT( commandBuffer, pMarkerInfo );
VkCmdDebugMarkerEndEXT( commandBuffer );
VkCmdDebugMarkerInsertEXT( commandBuffer, pMarkerInfo );
VkCmdDispatch( commandBuffer, groupCountX, groupCountY, groupCountZ );
VkCmdDispatchIndirect( commandBuffer, offset );
VkCmdDraw( commandBuffer, vertexCount, instanceCount, firstVertex, firstInstance );
VkCmdDrawIndexed( commandBuffer, indexCount, instanceCount, firstIndex, iH32_1_vertexOffset, firstInstance );
VkCmdDrawIndexedIndirect( commandBuffer, stride );
VkCmdDrawIndexedIndirectCountAMD( commandBuffer, stride );
VkCmdDrawIndirect( commandBuffer, stride );
VkCmdDrawIndirectCountAMD( commandBuffer, stride );
VkCmdEndQuery( commandBuffer, query );
VkCmdEndRenderPass( commandBuffer );
VkCmdExecuteCommands( commandBuffer, commandBufferCount, const pCommandBuffers );

```

7/26/2020

These are the Commands that could be entered into the Command Buffer, II 200

```

VkCmdFillBuffer( commandBuffer, dstBuffer, dstOffset, size, data );
VkCmdNextSubpass( commandBuffer, contents );
VkCmdPipelineBarrier( commandBuffer, srcStageMask, dstStageMask, dependencyFlags, memoryBarrierCount, VMemoryBarrier* pMemoryBarriers, bufferMemoryBarrierCount, pBufferMemoryBarriers, imageMemoryBarrierCount, pImageMemoryBarriers );
VkCmdProcessCommandsNVX( commandBuffer, pProcessCommandInfo );
VkCmdPushConstants( commandBuffer, layout, stageFlags, offset, size, pValues );
VkCmdPushDescriptorSetKHR( commandBuffer, pipelineBindPoint, layout, set, descriptorWriteCount, pDescriptorWrites );
VkCmdPushDescriptorSetWithTemplateKHR( commandBuffer, descriptorUpdateTemplate, layout, set, pData );
VkCmdReserveSpaceForCommandsNVX( commandBuffer, pReserveSpaceInfo );
VkCmdResetEvent( commandBuffer, event, stageMask );
VkCmdResetQueryPool( commandBuffer, queryPool, firstQuery, queryCount );
VkCmdResolveImage( commandBuffer, srcImage, srcImageLayout, dstImage, dstImageLayout, regionCount, pRegions );
VkCmdSetBlendConstants( commandBuffer, blendConstants[4] );
VkCmdSetDepthBias( commandBuffer, depthBiasConstantFactor, depthBiasClamp, depthBiasSlopeFactor );
VkCmdSetDepthBounds( commandBuffer, minDepthBounds, maxDepthBounds );
VkCmdSetDeviceMaskKHV( commandBuffer, deviceMask );
VkCmdSetDiscardRectangleEXT( commandBuffer, firstDiscardRectangle, discardRectangleCount, pDiscardRectangles );
VkCmdSetEvent( commandBuffer, event, stageMask );
VkCmdSetLineWidth( commandBuffer, lineWidth );
VkCmdSetScissor( commandBuffer, firstScissor, scissorCount, pScissors );
VkCmdSetStencilCompareMask( commandBuffer, faceMask, compareMask );
VkCmdSetStencilReference( commandBuffer, faceMask, reference );
VkCmdSetStencilWriteMask( commandBuffer, faceMask, writeMask );
VkCmdSetViewport( commandBuffer, firstViewport, viewportCount, pViewports );
VkCmdSetViewportWScalingNV( commandBuffer, firstViewport, viewportCount, pViewportWScalings );
VkCmdUpdateBuffer( commandBuffer, dstBuffer, dstOffset, dataSize, pData );
VkCmdWaitEvents( commandBuffer, eventCount, pEvents, srcStageMask, dstStageMask, memoryBarrierCount, pBufferMemoryBarriers, bufferMemoryBarrierCount, pBufferMemoryBarriers, imageMemoryBarrierCount, pImageMemoryBarriers );
VkCmdWriteTimestamp( commandBuffer, pipelineStage, queryPool, query );

```

7/26/2020

201

```

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 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 );
}

```

7/26/2020

202

```

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

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

VkClearColor vccv[2];
vccv[1].depthStencil = vcdsv;

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

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

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

```

7/26/2020

203

```

VkViewport viewport;
{
    0, // x
    0, // y
    (float)Width, // xMax
    (float)Height, // yMax
    0, // minDepth
    1, // maxDepth
};

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

VkRect2D scissor;
{
    0, // x
    0, // y
    Width, // xMax
    Height, // yMax
};

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

VkCmdBindDescriptorSets( CommandBuffers[nextImageIndex], VK_PIPELINE_BIND_POINT_GRAPHICS, DispatchIndex, 0, 4, Descriptors, 0, (uint32_t) 1, Trainers );

VkCmdBindPushConstants( CommandBuffers[nextImageIndex], PipelineLayout, VK_SHADER_STAGE_ALL, offset, size, void* values );
VkBuffer buffers[1] = { MyVertexDataBuffer };
VkDeviceSize offset[1] = { 0 };

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

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

VkCmdEndRenderPass( CommandBuffers[nextImageIndex] );
VkCmdEndCommandBuffer( CommandBuffers[nextImageIndex] );

```

7/26/2020

Submitting a Command Buffer to a Queue for Execution 204

```

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

```

7/26/2020

The Entire Submission / Wait / Display Process

```

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

VkFence renderFence;
vkCreateFence(logicalDevice, IN &info, PALLOCATOR, OUT &renderFence);
result = VK_SUCCESS;

VkPipelineStageFlags waitABitton = VK_PIPELINE_STAGE_BOTTOM_OF_PIPE_BIT;
VkQueue presentQueue;
vkGetDeviceQueue(logicalDevice, FindQueueFamilyThatDoesntSupport(VK_QUEUE_FAMILY_UNSUPPORTED), 0, OUT &presentQueue);
// 0 = queue index

VkSubmitInfo info;
info.sType = VK_STRUCTURE_TYPE_SUBMIT_INFO;
info.pNext = nullptr;
info.waitSemaphoreCount = 1;
info.pWaitSemaphores = &renderFence;
info.pWaitDstStageMask = &waitABitton;
info.commandBufferCount = 1;
info.pCommandBuffers = &commandBuffer;
info.signalSemaphoreCount = 0;
info.pSignalSemaphores = &renderFence;

result = vkQueueSubmit(presentQueue, 1, IN &info, IN renderFence); // 1 = submit count
result = vkWaitForFences(logicalDevice, 1, IN &renderFence, VK_TRUE, UINT64_MAX); // wait! timeout

vkDestroyFence(logicalDevice, renderFence, PALLOCATOR);

VkPresentInfoKHR info;
info.sType = VK_STRUCTURE_TYPE_PRESENT_INFO_KHR;
info.pNext = nullptr;
info.waitSemaphoreCount = 0;
info.pWaitSemaphores = (VkSemaphore*) nullptr;
info.swapchainCount = 1;
info.pSwapchains = &swapChain;
info.pImageIndices = &imageIndex;
info.pResults = (VkResult*) nullptr;

result = vkQueuePresentKHR(presentQueue, IN &info);

```

What Happens After a Queue has Been Submitted?

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.

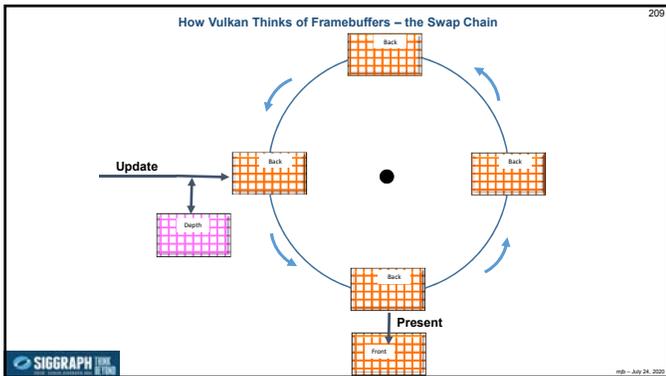
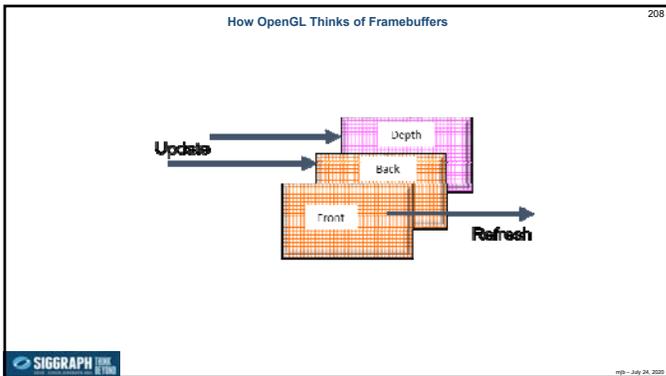
The message here is, I think, always consider using some sort of Vulkan synchronization when one command depends on a previous command reaching a certain state first.

Vulkan.

The Swap Chain

Mike Bailey
mjb@cs.oregonstate.edu

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



What is a Swap Chain?

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”)

We Need to Find Out What our Display Capabilities Are

```

VkSurfaceCapabilitiesKHR vsc;
vkGetPhysicalDeviceSurfaceCapabilitiesKHR(PhysicalDevice, Surface, OUT &vsc);
VkExtent2D surfaceRes = vsc.currentExtent;
printf( FpDebug, "vkGetPhysicalDeviceSurfaceCapabilitiesKHR:\n" );
...
VkBool32 supported;
result = vkGetPhysicalDeviceSurfaceSupportKHR( PhysicalDevice, FindQueueFamilyThatDoesGraphics(), Surface, &supported );
if( supported == VK_TRUE )
    printf( 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 );
printf( FpDebug, "Found %d Surface Formats:\n", formatCount );
...
uint32_t presentModeCount;
vkGetPhysicalDeviceSurfacePresentModesKHR( PhysicalDevice, Surface, &presentModeCount, (VkPresentModeKHR *) nullptr );
VkPresentModeKHR * presentModes = new VkPresentModeKHR[presentModeCount];
vkGetPhysicalDeviceSurfacePresentModesKHR( PhysicalDevice, Surface, &presentModeCount, presentModes );
printf( FpDebug, "Found %d Present Modes:\n", presentModeCount );
...

```

We Need to Find Out What our Display Capabilities Are

VulkanDebug.txt output:

```

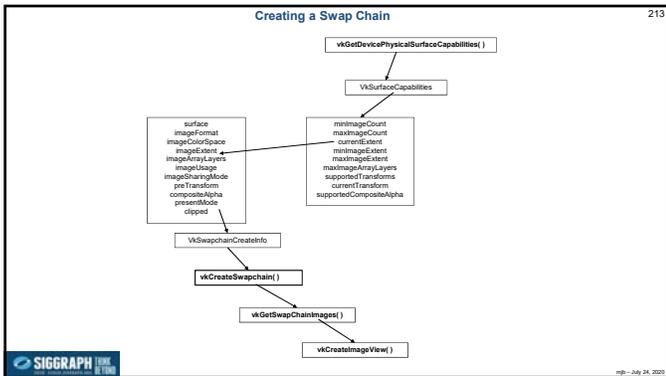
vkGetPhysicalDeviceSurfaceCapabilitiesKHR:
minImageCount = 2; maxImageCount = 8
currentExtent = 1024 x 1024
minImageExtent = 1024 x 1024
maxImageExtent = 1024 x 1024
maxImageArrayLayers = 1
supportedTransforms = 0x0001
currentTransform = 0x0001
supportedCompositeAlpha = 0x0001
supportedUsageFlags = 0x000f

** 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)

```



Creating a Swap Chain

```

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

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

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

```

Creating the Swap Chain Images and Image Views

```

uint32_t imageCount; // # of display buffers - 2?
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] );
}

```

Rendering into the Swap Chain, I

```

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

VkSemaphore ImageReadySemaphore;
result = vkCreateSemaphore( LogicalDevice, IN &vscsi, 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 &vrbpi,
    IN VK_SUBPASS_CONTENTS_INLINE );

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

...

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

```


Setting up the Push Constants for the Pipeline Structure

223

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

```

VkPushConstantRange
vpcr[0].stageFlags = vpcr[1]
                    VK_PIPELINE_STAGE_VERTEX_SHADER_BIT |
                    VK_PIPELINE_STAGE_FRAGMENT_SHADER_BIT;
vpcr[0].offset = 0;
vpcr[0].size = sizeof( glm::mat4 );

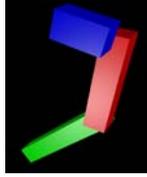
VkPipelineLayoutCreateInfo
vpcli.sType = VK_STRUCTURE_TYPE_PIPELINE_LAYOUT_CREATE_INFO;
vpcli.pNext = nullptr;
vpcli.flags = 0;
vpcli.setLayoutCount = 4;
vpcli.pSetLayouts = DescriptorSetLayouts;
vpcli.pushConstantRangeCount = 1;
vpcli.pPushConstantRanges = vpcr;

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

An Robotic Example using Push Constants

224

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



Where each arm is represented by:

```

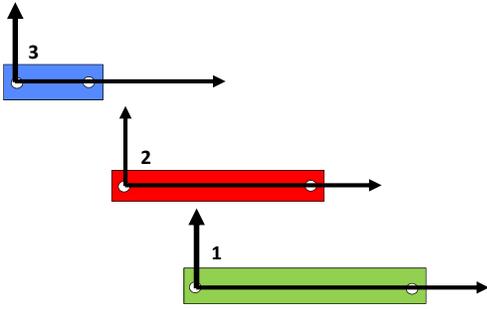
struct arm
{
    glm::mat4  armMatrix;
    glm::vec3  armColor;
    float      armScale; // scale factor in x
};

struct arm  Arm1;
struct arm  Arm2;
struct arm  Arm3;
    
```

Forward Kinematics:

225

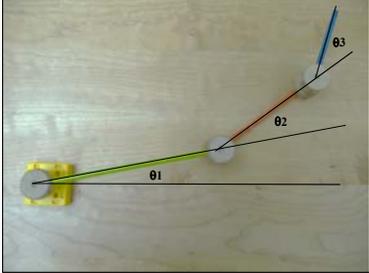
You Start with Separate Pieces, all Defined in their Own Local Coordinate System



Forward Kinematics:

226

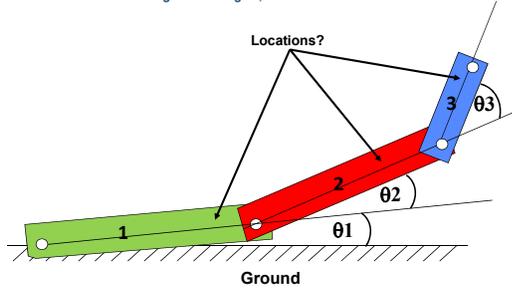
Hook the Pieces Together, Change Parameters, and Things Move (All Young Children Understand This)



Forward Kinematics:

227

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



Positioning Part #1 With Respect to Ground

228

1. Rotate by θ_1
2. Translate by $T_{1/G}$

Write it \rightarrow

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

Say it \leftarrow

Positioning Part #2 With Respect to Ground

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

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

Say it

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

Positioning Part #3 With Respect to Ground

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

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

Say it

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

In the Reset Function

```

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()`.

Setup the Push Constant for the Pipeline Structure

```

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

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

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

In the UpdateScene Function

```

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
    
```

In the RenderScene Function

```

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

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

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

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

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

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

In the Vertex Shader 235

```

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

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

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

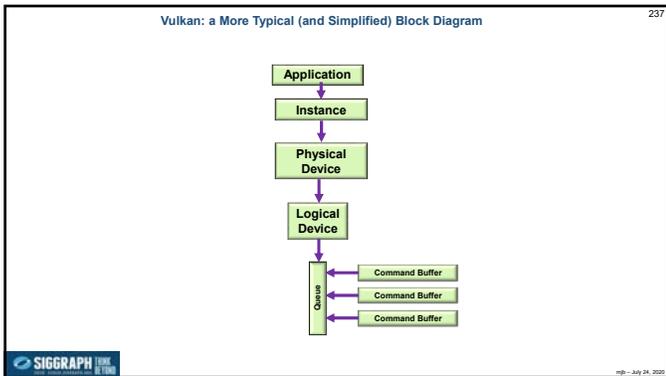
gl_Position = PVM * vec4( bVertex, 1. ); // Projection * Viewing * Modeling matrices
    
```




Physical Devices

 Mike Bailey
 mjb@cs.oregonstate.edu

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

Querying the Number of Physical Devices 238

```

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

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

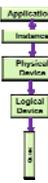
This way of querying information is a recurring OpenCL and Vulkan pattern (get used to it):

```

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



Vulkan: Identifying the Physical Devices 239



```

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



Which Physical Device to Use, I 240

```

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

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



Which Physical Device to Use, II

```

// need some logical here to decide which physical device to select:
if( vpdp.deviceType == VK_PHYSICAL_DEVICE_TYPE_DISCRETE_GPU )
    discreteSelect = 1;

if( vpdp.deviceType == VK_PHYSICAL_DEVICE_TYPE_INTEGRATED_GPU )
    integratedSelect = 1;
}

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;
}
    
```

241

Asking About the Physical Device's Features

```

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

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

242

Here's What the NVIDIA RTX 2080 Ti Produced

```

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
    
```

243

Here's What the Intel HD Graphics 520 Produced

```

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
    
```

244

Asking About the Physical Device's Different Memories

```

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

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

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

245

Here's What I Got

```

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
    
```

246

Asking About the Physical Device's Queue Families 247

```

uint32_t count = -1;
vkGetPhysicalDeviceQueueFamilyProperties( IN PhysicalDevice, &count, OUT (VkQueueFamilyProperties *)nullptr );
fprintf( FpDebug, "nFound %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, "l%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 248

Found 3 Queue Families:

0: queueCount = 16 ; Graphics Compute Transfer

1: queueCount = 2 ; Transfer

2: queueCount = 8 ; Compute

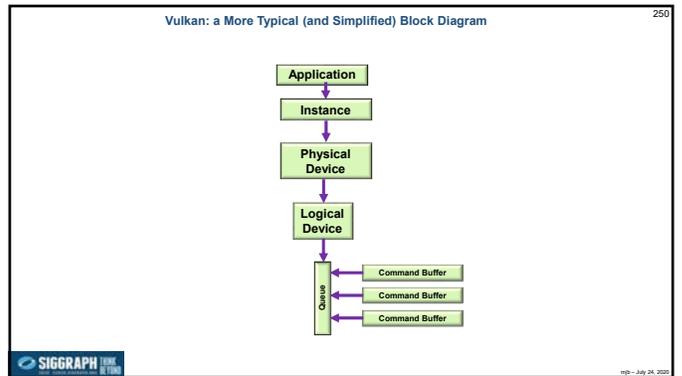
mjb - July 24, 2020

Vulkan.
Logical Devices

Mike Bailey
mjb@cs.oregonstate.edu

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

mjb - July 24, 2020



Looking to See What Device Layers are Available 251

```

const char * myDeviceLayers[] =
{
    // "VK_LAYER_LUNARG_api_dump",
    // "VK_LAYER_LUNARG_core_validation",
    // "VK_LAYER_LUNARG_image",
    "VK_LAYER_LUNARG_objct_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);
    
```

mjb - July 24, 2020

Looking to See What Device Extensions are Available 252

```

// 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);
        
```

mjb - July 24, 2020

What Device Layers and Extensions are Available 253

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'

```

SIGGRAPH THINK BEYOND HO - July 24, 2020

Vulkan: Creating a Logical Device 254

```

float queuePriorities[] =
{
    1.
};
VkDeviceQueueCreateInfo (vdqci)
vdqci.sType = VK_STRUCTURE_TYPE_DEVICE_QUEUE_CREATE_INFO;
vdqci.pNext = nullptr;
vdqci.flags = 0;
vdqci.queueFamilyIndex = 8;
vdqci.queueCount = 1;
vdqci.pQueueProperties = queuePriorities;

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

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

```

SIGGRAPH THINK BEYOND HO - July 24, 2020

Vulkan: Creating the Logical Device's Queue 255

```

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

```

SIGGRAPH THINK BEYOND HO - July 24, 2020

Vulkan.

Introduction to the Vulkan Computer Graphics API

Mike Bailey
mjb@cs.oregonstate.edu

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

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

SIGGRAPH THINK BEYOND HO - July 24, 2020