



Vulkan.


Shaders and SPIR-V



Oregon State University
Mike Bailey
mjb@cs.oregonstate.edu

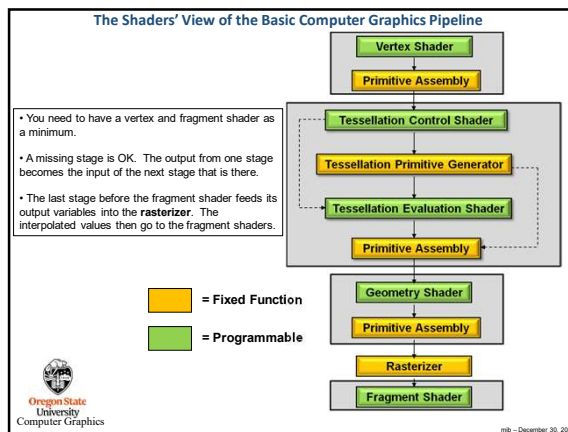


This work is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](https://creativecommons.org/licenses/by-nc-nd/4.0/).



ShadersAndSPIR-V.pptx

mjb - December 30, 2022




Vulkan Shader Stages

Shader stages

```

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



mjb - December 30, 2022

How Vulkan GLSL Differs from OpenGL GLSL

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

```

#define VULKAN 130
or whatever the current version number is.
Typically you use this like:
#ifdef VULKAN
...
#endif
    
```

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:
layout (location = 0) out vec4 fFragColor;



mjb - December 30, 2022

How Vulkan GLSL Differs from OpenGL GLSL

Shader combinations of separate texture data and samplers as an option:

```

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

For example, Specialization Constants can be used with 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



mjb - December 30, 2022

Vulkan: Shaders' use of Layouts for Uniform Variables

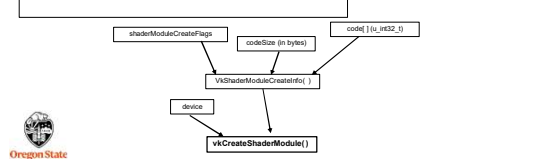
```


layout( std140, set = 0, binding = 0 ) uniform sceneMatBuf
{
    mat4 uProjectionMatrix;
    mat4 uViewMatrix;
    mat4 uSceneMatrix;
} SceneMatrices;

layout( std140, set = 1, binding = 0 ) uniform objectMatBuf
{
    mat4 uModelMatrix;
    mat4 uNormalMatrix;
} ObjectMatrices;

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

All non-sampler uniform variables must be in block buffers





mjb - December 30, 2022

Vulkan Shader Compiling

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

Advantages:

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

Computer Graphics | mjb - December 30, 2022

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

Computer Graphics | mjb - December 30, 2022

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

You can run the glslangValidator program from the Windows Command Prompt, but I have found it easier to run the SPIR-V compiler from Windows-Bash.

To install the bash shell on your own Windows machine, go to this URL:
<https://www.msn.com/en-us/news/technology/how-to-install-and-run-bash-on-windows-11/ar-AA10EoPk>

Or, follow these instructions:

- Head to the **Start** menu search bar, type in 'terminal,' and launch the Windows Terminal as administrator. (On some systems, this is called the **Command Prompt**.)
- Type in the following command in the administrator: **wsl --install**
- Restart your PC once the installation is complete.

As soon as your PC boots up, the installation will begin again. Your PC will start downloading and installing the Ubuntu software. You'll soon get asked to set up a username and password. This can be the same as your system's username and password, but doesn't have to be. The installation will automatically start off from where you left it.

Computer Graphics | mjb - December 30, 2022

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

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

BTW, within bash, if you want to list your files without that sometimes-hard-to-read filename coloring, do this:
ls -l --color=none
 (ell-ess minus-ell minus-minus-color=none)

Computer Graphics | mjb - December 30, 2022

Running glslangValidator.exe in bash

As long as I am on bash, I like using the `make` utility. To do that, put these shader compile lines in a file called `Makefile`:

```
ALLSHADERS: sample-vert.vert sample-frag.frag
glslangValidator.exe -V sample-vert.vert -o sample-vert.spv
glslangValidator.exe -V sample-frag.frag -o sample-frag.spv
```

Then type `make ALLSHADERS`:

```
ljb@PC: /mnt/c/MJB/Vulkan/Sample2019-COLOREDCUBE$ make ALLSHADERS
glslangValidator.exe -V sample-vert.vert -o sample-vert.spv
sample-vert.vert
glslangValidator.exe -V sample-frag.frag -o sample-frag.spv
sample-frag.frag
ljb@PC: /mnt/c/MJB/Vulkan/Sample2019-COLOREDCUBE$
```

Computer Graphics | mjb - December 30, 2022

Running glslangValidator.exe

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

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

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

<code>.vert</code>	Vertex shader
<code>.tccs</code>	Tessellation Control Shader
<code>.tecs</code>	Tessellation Evaluation Shader
<code>.geom</code>	Geometry shader
<code>.frag</code>	Fragment shader
<code>.comp</code>	Compute shader

Computer Graphics | mjb - December 30, 2022

How do you know if SPIR-V compiled successfully?

Same as C/C++ -- the compiler gives you no nasty messages, it just prints the name of the source file you just compiled.

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


So, if you use the Linux command **od -x** on the .spv file, like this:

```
od -x sample-vert.spv
```

the magic number shows up like this:

```
00000000 0203 0723 0000 0001 000a 0008 007e 0000
00000020 0000 0000 0011 0002 0001 0000 000b 0006
...
```

"od" stands for "octal dump", even though it can format the raw bits as most anything: octal, hexadecimal, bytes, characters, etc. "-x" means to format in hexadecimal.

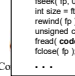


mpb - December 30, 2022

Reading a SPIR-V File into a Vulkan Shader Module

```
#ifndef _WIN32
typedef int errno_t;
int fopen_s(FILE**, const char*, const char*);
#endif

#define SPIRV_MAGIC 0x07230203
...
VkResult Init2SpvShader(std::string filename, VkShaderModule* pShaderModule)
{
    FILE* fp;
    #ifdef _WIN32
    errno_t err = fopen_s(&fp, filename.c_str(), "rb");
    if (err != 0)
    #else
    fp = fopen(filename.c_str(), "rb");
    if (fp == NULL)
    #endif
    {
        fprintf(stderr, "Cannot open shader file '%s'\n", filename.c_str());
        return VK_SHOULD_EXIT;
    }
    uint32_t magic;
    fread(&magic, 4, 1, fp);
    if (magic != SPIRV_MAGIC)
    {
        fprintf(stderr, "Magic number for spir-v file '%s' is 0x%08x -- should be 0x%08x\n", filename.c_str(), magic, SPIRV_MAGIC);
        return VK_SHOULD_EXIT;
    }
    fseek(fp, 0L, SEEK_END);
    int size = ftell(fp);
    rewind(fp);
    unsigned char* code = new unsigned char[size];
    fread(code, size, 1, fp);
    delete fp;
    ...
}
```




mpb - December 30, 2022

Reading a SPIR-V File into a Shader Module

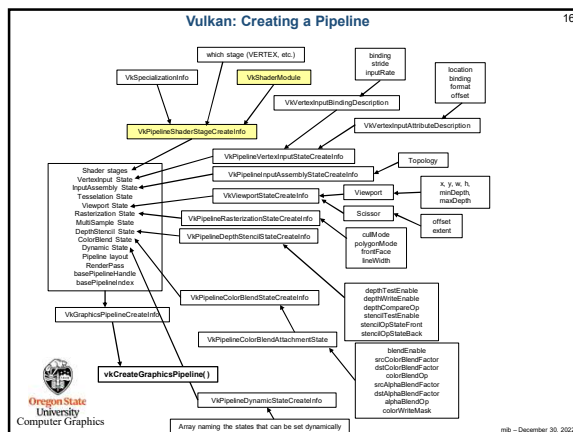
```
...
VkShaderModule ShaderModuleVertex;
...

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

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




mpb - December 30, 2022



You can also take a look at SPIR-V Assembly

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

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



mpb - December 30, 2022

For example, if this is your Shader Source

```
#version 450
#extension GL_ARB_separate_shader_objects : enable
#extension GL_ARB_shading_language_420pack : enable
layout(std140, set = 0, binding = 0) uniform mat4Buf
{
    mat4 uModelMatrix;
    mat4 uViewMatrix;
    mat4 uProjectionMatrix;
    mat3 uNormalMatrix;
} Matrices;

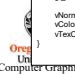
// non-opaque must be in a uniform block:
layout(std140, set = 1, binding = 0) uniform lightBuf
{
    vec4 uLightPos;
} Light;

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

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

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

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



mpb - December 30, 2022

