

Writing an Efficient Vulkan Renderer for Quest Pro and Quest 3

Sergey Kosarevsky, Meta
Alexey Medvedev, Meta



Agenda

- ▶ What is IGL?
- ▶ IGL Shell?
- ▶ Compiling GLSL shaders at run-time
- ▶ Command buffers
- ▶ Fences
- ▶ Staging device
- ▶ Render pipeline states
- ▶ Layout transitions & dependencies
- ▶ Future plans for IGL API

What is IGL?

- ▶ Cross-platform, graphics API agnostic rendering library
- ▶ Used in 20+ Meta's projects (internal and external)
- ▶ **Windows, macOS, iOS, Android, Linux, WebAssembly, OpenXR**
- ▶ **GLES 2, GLES 3+, WebGL 2, GL 4.1+, Metal 2+**
- ▶ **Vulkan 1.1** + optional bindless extensions
VK_KHR_buffer_device_address,
VK_EXT_descriptor_indexing
- ▶ **Open Source MIT project**
<https://github.com/facebook/igl>

IGL



What is IGL?



IGL Shell?

- ▶ Cross-platform, application agnostic framework
- ▶ Used as internal test farmework and external sample app
- ▶ RenderSession — "application"
- ▶ No changes required to run on any platform
- ▶ **Windows, macOS, iOS, Android, Linux, WebAssembly, OpenXR**

IGL



What is IGL?

Modelled after the Metal API and has a very strong OpenGL ES 2 influence:

```
1 class IRenderCommandEncoder
2 {
3 public:
4     ...
5     virtual void bindRenderPipelineState(IRenderPipelineState& pipelineState) = 0;
6     virtual void bindBuffer(int index, IBuffer& buffer, size_t bufferSize) = 0;
7     virtual void bindSamplerState(int index, ISamplerState& samplerState) = 0;
8     virtual void bindTexture(int index, ITexture& texture) = 0;
9     virtual void drawIndexed(PrimitiveType primitiveType, uint32_t indexCount,
10         IndexFormat indexFormat, IBuffer& indexBuffer, size_t indexBufferOffset) = 0;
11     ...
12 };
```

Compiling GLSL shaders at run-time

- ▶ IGL allows GLSL shaders
- ▶ Vulkan accepts only binary SPIR-V
- ▶ Use 'glslang' to compile GLSL at run-time

Compiling GLSL shaders at run-time

```
1 Result compileShader(VkDevice device,
2 VkShaderStageFlagBits stage,
3 const char* code,
4 std::vector<uint32_t>& outSPIRV,
5 const glslang_resource_t* glslLangResource)
6 {
7     const glslang_input_t input = {
8         .language = GLSLANG_SOURCE_GLSL,
9         .stage = getGLSLangShaderStage(stage),
10        .client = GLSLANG_CLIENT_VULKAN,
11        .client_version = GLSLANG_TARGET_VULKAN_1_1,
12        .target_language = GLSLANG_TARGET_SPV,
13        .target_language_version = GLSLANG_TARGET_SPV_1_3,
14        .code = code,
15        .default_version = 100,
16        .default_profile = GLSLANG_NO_PROFILE,
17        .messages = GLSLANG_MSG_DEFAULT_BIT,
18        .resource = glslLangResource,
19    };
20
21    glslang_shader_t* shader = glslang_shader_create(&input);
22    IGL_SCOPE_EXIT { glslang_shader_delete(shader); };
23
24    if (!glslang_shader_preprocess(shader, &input)) {
25        IGL_LOG_ERROR("Preprocessing failed:\n %s\n %s\n",
26            glslang_shader_get_info_log(shader),
27            glslang_shader_get_info_debug_log(shader));
28        logShaderSource(code);
29        return Result(InvalidOperation);
30    }
31
32    if (!glslang_shader_parse(shader, &input)) {
33        IGL_LOG_ERROR("Parsing failed:\n %s\n %s\n",
34            glslang_shader_get_info_log(shader),
35            glslang_shader_get_info_debug_log(shader));
36        logShaderSource(
37            glslang_shader_get_preprocessed_code(shader));
38        return Result(InvalidOperation);
39    }
40
41    glslang_program_t* p = glslang_program_create();
42    glslang_program_add_shader(p, shader);
43    IGL_SCOPE_EXIT { glslang_program_delete(p); };
44    if (!glslang_program_link(p, GLSLANG_MSG_SPV_RULES_BIT |
45        GLSLANG_MSG_VULKAN_RULES_BIT)) {
46        IGL_LOG_ERROR("Linking failed:\n %s\n %s\n",
47            glslang_program_get_info_log(p),
48            glslang_program_get_info_debug_log(p));
49        return Result(InvalidOperation);
50    }
51
52    glslang_spv_options_t options = { .optimize_size = true };
53    glslang_program_SPIRV_generate_with_options(
54        p, input.stage, &options);
55    if (glslang_program_SPIRV_get_messages(p)) {
56        IGL_LOG_ERROR("%s\n",
57            glslang_program_SPIRV_get_messages(p));
58    }
59
60    const uint32_t* ptr = glslang_program_SPIRV_get_ptr(p);
61    outSPIRV = std::vector(ptr,
62        ptr + glslang_program_SPIRV_get_size(p));
63    return Result();
64 }
```

Command buffers

- ▶ Command pools and buffers are encapsulated in **VulkanImmediateCommands**
- ▶ Expose a high-level interface for command buffers

Command buffers

```
1 class VulkanImmediateCommands final {
2 public:
3     struct SubmitHandle {
4         uint32_t bufferIndex_ = 0;
5         uint32_t submitId_ = 0;
6         bool empty() const { return submitId_ == 0; }
7         uint64_t handle() const { return (uint64_t(submitId_) << 32) + bufferIndex_;}
8     };
9     struct CommandBufferWrapper {
10        VkCommandBuffer cmdBuf_ = VK_NULL_HANDLE;
11        VkFence fence_;
12        VkSemaphore semaphore_;
13        SubmitHandle handle_ = {};
14    };
15    const CommandBufferWrapper& acquire();
16    SubmitHandle submit(const CommandBufferWrapper& wrapper);
17
18    void waitSemaphore(VkSemaphore semaphore); // inject a swapchain semaphore etc
19
20    bool isReady(SubmitHandle handle) const;
21    void wait(SubmitHandle handle, uint64_t timeoutNanoseconds = UINT64_MAX);
22    void waitAll();
23    ...
24 private:
25    VkCommandPool commandPool_;
26    std::vector<CommandBufferWrapper> buffers_;
27    ...
28 };
```

Command buffers

```
1 VulkanImmediateCommands::SubmitHandle VulkanImmediateCommands::submit(const CommandBufferWrapper& wrapper) {
2     VK_ASSERT(vkEndCommandBuffer(wrapper.cmdBuf_));
3
4     VkSemaphore waitSemaphores[] = {VK_NULL_HANDLE, VK_NULL_HANDLE};
5     uint32_t numWaitSemaphores = 0;
6     if (waitSemaphore_) waitSemaphores[numWaitSemaphores++] = waitSemaphore_;
7     if (lastSubmitSemaphore_) waitSemaphores[numWaitSemaphores++] = lastSubmitSemaphore_;
8
9     const VkPipelineStageFlags waitStageMasks[] =
10     { VK_PIPELINE_STAGE_ALL_COMMANDS_BIT, VK_PIPELINE_STAGE_ALL_COMMANDS_BIT };
11     const VkSubmitInfo si = {
12         .sType = VK_STRUCTURE_TYPE_SUBMIT_INFO,
13         .waitSemaphoreCount = numWaitSemaphores,
14         .pWaitSemaphores = numWaitSemaphores ? waitSemaphores : nullptr,
15         .pWaitDstStageMask = waitStageMasks,
16         .commandBufferCount = 1,
17         .pCommandBuffers = &wrapper.cmdBuf_,
18         .signalSemaphoreCount = 1,
19         .pSignalSemaphores = &wrapper.semaphore_.vkSemaphore_,
20     };
21     VK_ASSERT(vkQueueSubmit(queue_, 1u, &si, wrapper.fence_.vkFence_));
22
23     lastSubmitSemaphore_ = wrapper.semaphore_.vkSemaphore_;
24     lastSubmitHandle_ = wrapper.handle_;
25     waitSemaphore_ = VK_NULL_HANDLE;
26     submitCounter_++;
27
28     return lastSubmitHandle_;
29 }
```

Fences

- ▶ Vulkan synchronization is too low-level
- ▶ Expose a high-level interface for fences

```
1
2 // GPU Fence Handle – exposed on the IGL API level
3 using SubmitHandle = uint64_t;
4
5 class ICommandQueue {
6 public:
7     virtual SubmitHandle submit(const ICommandBuffer& commandBuffer) = 0;
8 };
```

Fences

```
1
2 bool VulkanImmediateCommands::isReady(const SubmitHandle handle) const {
3     IGL_ASSERT(handle.bufferIndex_ < kMaxCommandBuffers);
4
5     if (handle.empty()) {
6         // a null handle
7         return true;
8     }
9
10    const CommandBufferWrapper& buf = buffers_[handle.bufferIndex_];
11
12    if (buf.cmdBuf_ == VK_NULL_HANDLE) {
13        // already recycled and not yet reused
14        return true;
15    }
16
17    if (buf.handle_.submitId_ != handle.submitId_) {
18        // already recycled and reused by another command buffer
19        return true;
20    }
21
22    return vkWaitForFences(device_, 1, &buf.fence_.vkFence_, VK_TRUE, 0) == VK_SUCCESS;
23 }
```

Staging device

- ▶ Upload texture/buffer data to GPU resources
- ▶ Manages lazily allocated staging buffers
- ▶ Can grow up to 256Mb size based on workload

Staging device

```
1 class VulkanStagingDevice final {
2     struct MemoryRegion {
3         VkDeviceSize offset = 0u;
4         VkDeviceSize size = 0u;
5         VkDeviceSize alignedSize = 0u;
6         VulkanImmediateCommands::SubmitHandle handle;
7     };
8
9     MemoryRegion nextFreeBlock(VkDeviceSize size);
10    void growStagingBuffer(VkDeviceSize minimumSize);
11    ...
12
13 public:
14    void bufferSubData(VulkanBuffer& buffer, size_t dstOffset, size_t size, const void* data);
15    void getBufferSubData(VulkanBuffer& buffer, size_t srcOffset, size_t size, void* data);
16    void imageData(VulkanImage& image,
17                  TextureType type,
18                  const TextureRangeDesc& range,
19                  const TextureFormatProperties& properties,
20                  uint32_t bytesPerRow,
21                  const void* data);
22    ...
23 };
```

Render pipeline states

- ▶ IGL has more dynamic state than what unextended Vulkan 1.1 can provide
- ▶ Don't enable extra extensions
- ▶ Use a hash map to map IGL render states to **VkPipeline** objects
- ▶ Treat Vulkan render passes as a part of dynamic state

Render pipeline states

```
1
2 class alignas(sizeof(uint64_t)) RenderPipelineDynamicState {
3     uint32_t topology_           : 4;
4     uint32_t depthCompareOp_     : 3;
5     uint32_t stencilFrontFailOp_ : 3;
6     uint32_t stencilFrontPassOp_ : 3;
7     uint32_t stencilFrontDepthFailOp_ : 3;
8     uint32_t stencilFrontCompareOp_ : 3;
9     uint32_t stencilBackFailOp_   : 3;
10    uint32_t stencilBackPassOp_    : 3;
11    uint32_t stencilBackDepthFailOp_ : 3;
12    uint32_t stencilBackCompareOp_ : 3;
13    uint32_t renderPassIndex_     : 8;
14    uint32_t depthBiasEnable_     : 1;
15    uint32_t depthWriteEnable_    : 1;
16
17    bool operator==(const RenderPipelineDynamicState& other) const {
18        return *(uint64_t*)this == *(uint64_t*)&other;
19    }
20    struct HashFunction {
21        uint64_t operator()(const RenderPipelineDynamicState& s) const { return *(const uint64_t*)&s; }
22    };
23 };
24
25 class RenderPipelineState final : public IRenderPipelineState {
26     ...
27     unordered_map<RenderPipelineDynamicState, VkPipeline, RenderPipelineDynamicState::HashFunction> pipelines_;
28 };
```

Layout transitions & dependencies

- ▶ Vulkan image layout transitions are too verbose — track internally
- ▶ Expose an API to declare dependencies between graphics and compute work

Layout transitions & dependencies

```
1
2 struct Dependencies {
3     static constexpr uint32_t IGL_MAX_TEXTURE_DEPENDENCIES = 4;
4     static constexpr uint32_t IGL_MAX_BUFFER_DEPENDENCIES = 4;
5     ITexture* textures[IGL_MAX_TEXTURE_DEPENDENCIES] = {};
6     IBuffer* buffers[IGL_MAX_BUFFER_DEPENDENCIES] = {};
7 };
8
9 class ICommandBuffer {
10 public:
11     virtual std::unique_ptr<IRenderCommandEncoder> createRenderCommandEncoder(
12         const RenderPassDesc& renderPass,
13         IFramebuffer& framebuffer,
14         const Dependencies& dependencies,
15         Result* outResult) = 0;
16     virtual std::unique_ptr<IComputeCommandEncoder> createComputeCommandEncoder(const Dependencies& dependencies) = 0;
17     ...
18 };
```

Layout transitions & dependencies

```
1
2 for (auto& tex : dependencies.textures) {
3     transitionToShaderReadOnly(tex);
4 }
5 for (auto& buf : dependencies.buffers) {
6     if (!buf) continue;
7     VkPipelineStageFlags dstStageFlags = VK_PIPELINE_STAGE_VERTEX_SHADER_BIT | VK_PIPELINE_STAGE_FRAGMENT_SHADER_BIT;
8     if ((buf->vkUsageFlags_ & VK_BUFFER_USAGE_INDEX_BUFFER_BIT) ||
9         (buf->vkUsageFlags_ & VK_BUFFER_USAGE_VERTEX_BUFFER_BIT)) {
10        dstStageFlags |= VK_PIPELINE_STAGE_VERTEX_INPUT_BIT;
11    }
12    if (buf->vkUsageFlags_ & VK_BUFFER_USAGE_INDIRECT_BUFFER_BIT) {
13        dstStageFlags |= VK_PIPELINE_STAGE_DRAW_INDIRECT_BIT;
14    }
15    bufferBarrier(buf, VK_PIPELINE_STAGE_COMPUTE_SHADER_BIT, dstStageFlags);
16 }
```

Future plans for IGL API

- ▶ Better dependencies & barriers — a frame graph
- ▶ An API for Vulkan subpasses
- ▶ Replace smart pointers with integer handles

Questions?

Thank you!

Sergey Kosarevsky

 corporateshark
sk@linderdaum.com

Alexey Medvedev

 aleksrudybear
rudybear@gmail.com

<https://github.com/facebook/igl>

<https://github.com/corporateshark/lightweightvk> ← lots of bindless stuff in a fork