

Hello, my name is Graham Wihlidal, and I'm a senior rendering engineer at SEED, and the technical architect for Halcyon.



SEED

SEED is a technical and creative research division of Electronic Arts. We are a cross-disciplinary team with a mission to explore the future of interactive entertainment.

Disruptive Technology?

- Technology that significantly alters the way your business operates.
- Often forces companies to change their business for fear of losing market share or becoming irrelevant.

I'd like to start off with a definition – disruptive technology is technology that significantly alters the way your business operates.

Often this technology forces companies to change their business for fear of losing market share or becoming irrelevant.

Embrace Disruption

- **Maintain existing business**
 - Safe bets, manage risk
 - Also important!
- **Avoid “horse blinders”**
 - Evolving market
 - Evolving users
 - Evolving technology



It's important to embrace disruption, but it is also important to maintain your existing business; through the usual safe bets, and risk management.

After all, this is how you're currently successful, so don't ruin a good thing.

While being successful in the present, it's also very important to avoid “horse blinders” – which is where you ignore a rapidly evolving market, user base, and available technologies.

Some technologies may look like a gimmick or a fad, and many of them are, but some may turn out to be the next big thing, and the entire market pivots towards it.

Embrace Disruption

- Innovation is important!
- **Don't be Kodak!**
 - Invent the digital camera
 - Fear of disruption
 - “That's cute - but don't tell anyone about it.” [Mui 2012]
 - Competition wins

Because of this, innovation is extremely important.

My advice to you is don't be Kodak!

The digital camera was invented at Kodak in 1975 by an engineer named Steve Sasson, and brought to management where they said “that's cute – but don't tell anyone about it”, for fear of disrupting their analog film business.

In the end, the competition won, by successfully getting the digital camera out to market before Kodak knew what hit them. There are also similar misses around photo sharing and photo viewing.

This is all too common in large companies.

Game Development

- Typically, **large engines** have...
 - Complex build systems
 - Inter-connected dependencies
 - **Opinionated and rigid APIs**
 - Specialized systems
 - **Steep learning curve**


Lets frame this to game development in particular.

Typically, large game engines have complex build systems, inter-connected dependencies, opinionated and rigid APIs, specialized systems, and a very steep learning curve.

SEED // Halcyon: Rapid Innovation using Modern Graphics

Why?

- For good reason!
 - Ship great games
- Performance > flexibility



Why would you ever want that? For good reason, to ship great games!

Often, flexibility and extensibility suffer at the hand of performance and scalability.

Opinionated and rigid APIs help bring about performance, and a consistent way to interact with the game engine and its systems.

But...

- Constrains agile prototyping
- Hard to rapidly pivot architecture
 - New rendering engine?
 - No more triangles?
 - Emerging platforms?
 - Cloud, VR, AR, MR, XR, Mobile, Social, ...
 - ...
- High risk, requiring large investment

But, these decisions generally constrain agile prototyping, and make it hard to rapidly pivot the architecture to test out significant changes like a new rendering engine, not using triangles (Dreams PS4 or Claybook, as an example), emerging platforms like cloud, VR/AR/MR/XR, mobile, social, etc..

The engines can of course target these areas on the roadmap, but at a high risk, and often requiring a large investment of limited time and resources.

And what if these ideas fail, after all that effort?

Solution!

- Agile R&D engine
- Prove out technology
 - Scout ahead
 - Share findings
 - Mitigate risk
 - Guide adoption

Our solution, is to have an agile R&D engine, which allows us to quickly prove out technology.

The general premise is we can use this technology platform to scout ahead, share findings back to various teams, mitigate overall risk, and guide adoption of successful technology.

Comparison



Game Engine



R&D Engine

To show this with a couple images, the game engine is like a Formula 1 race car – built for speed, requiring an expert driver, and a large support team behind it to keep it tuned and on the track.

The race crew aren't making drastic changes to the car, for fear of losing the race, or worse. There is a ton of investment in the car, driver, crew, and race outcomes - failure is not an option.

On the other side, we have the R&D engine, which is like a soapbox racer – built for a good time, constantly evolving after crashes or losses, anyone can drive it, the operating cost is negligible, and the race is more about the experience than about the outcome.

Comparison



Game Engine



R&D Engine

For another comparison – a game engine is like a highly specialized artisan knife, built for just a few purposes, but purposes it excels at.

An R&D engine, on the other hand, is like a swiss-army knife. It looks goofy as hell, but it can fulfill a large assortment of purposes, with reasonable enough quality.



So lets talk a bit about our swiss-army knife, called Halcyon.

SEED // Halcyon: Rapid Innovation using Modern Graphics

Halcyon


- **Rapid prototyping** engine
- Windows, Linux, macOS
- **Different purpose than Frostbite**
 - Fast experimentation vs. AAA games



Halcyon is a rapid prototyping engine - currently supported on Windows, Linux, and macOS – and serving a very different purpose than our flagship AAA engine, Frostbite.

“PICA PICA”

- **Exploratory** mini-game & world
[Andersson 2018] [Harmer 2018] [Opara 2018]
- Built with **Halcyon**
- Goals
 - **Hybrid rendering**
 - Clean and consistent visuals
 - Self-learning AI agents
 - Procedural worlds
 - **No precomputation**



One of our recent projects using Halcyon was called PICA PICA, which was showcased at GDC 2018.

We built PICA PICA from the ground up using Halcyon. The main goal of the project was to explore and experiment with hybrid rendering and ray tracing; aiming for clean and consistent visuals.

We also wanted to incorporate procedural world creation, and self-learning AI agents to navigate and operate the environment. And all of these goals were under the banner of no precomputation.

The project was very successful, both inside and outside of Electronic Arts, and provided a lot of cutting edge techniques and findings that were provided back to the organization.

Here is a video we circulated, showing our results from this project.



Halcyon Goals

- **Minimize or eliminate busy-work**
 - Artist “meta-data” meshes
 - Occlusion
 - GI / Lighting
 - Collision
 - Level-of-detail
- **Live reloading of all assets**
 - Insanely fast iteration times

Halcyon allowed us to iterate and move very fast with the development of PICA PICA.

A major goal of Halcyon is to minimize or eliminate busy-work; something I call - artist “meta-data” meshes.

Show me one artist that actually enjoys making these meshes over something more creative, and I guarantee you they are brainwashed, and need an intervention and our caring support.

Another critical goal, is the live reloading of all assets. We don’t want to take a coffee break while we shut down Halcyon, launch a data build, come back, and resume whatever we were doing.

Halcyon Goals

- Only target modern APIs
 - Vulkan 1.1
 - Direct3D 12
 - Metal 2
- Multi-GPU
 - Explicit heterogeneous mGPU
 - No AFR / No linked adapters



One luxury we had by starting from scratch, was choosing our feature set and min spec. We decided to only target modern APIs, so we are not restricted by legacy.

Another interesting goal is to provide easy access to multiple GPUs, without sacrificing API cleanliness or maintainability. To accomplish this, we decided on explicit heterogeneous mGPU, not linked adapters.

We also are avoiding any AFR nonsense for a number of reasons, including problems with temporal techniques.

Halcyon Goals

- **Scalable computation**
 - All cores in a workstation
 - Multiple graphics adapters (mGPU)
 - Local cluster
 - Google Cloud Platform
- Under a **single abstraction**
 - Easy, right?

Since we also want to experiment with techniques that may not map to average present-day hardware, we also decided on supporting scalable computation.

We scale to all cores in a workstation, multiple graphics adapters, local Kubernetes clusters, and also Google Cloud Platform.

We also designed our scaling approaches to work under a single abstraction.

Halcyon Goals


- Local or remote streaming
- Minimal boilerplate code
- Variety of rendering techniques and approaches
 - Rasterization
 - Path and ray tracing
 - Hybrid

We do traditional local rendering, but in addition to scaling computation, we also chose to support scaling rendering remotely, and transmitting the results back to the application (such as buffers, textures, or swap chains).


In order to deliver on our promise of fast experimentation, we needed to ensure a minimal amount of boilerplate code. This includes code to load a shader, set up pipelines and render state, query scene representations, etc.

This was critical in developing and supporting a vast number of rendering techniques and approaches that we have in Halcyon.


Hybrid Rendering



Deferred Shading
(raster)



Direct Shadows
(ray trace or raster)



Direct Lighting
(compute)



Reflections
(ray trace or compute)



Global Illumination
(ray trace)



Ambient Occlusion
(ray trace or compute)



Transparency & Translucency
(ray trace)




Post Processing
(compute)

We have a unique rendering pipeline which is inspired by classic techniques in real-time rendering, with ray tracing sprinkled on top.


We have a deferred renderer with compute-based lighting, and a fairly standard post-processing stack. And then there are a few pluggable components. We can render shadows via ray tracing or cascaded shadow maps.

Reflections can be traced or screen-space marched. Same story for ambient occlusion.

Only our global illumination and translucency actually require ray tracing.



Here is an example of a scene using our hybrid rendering pipeline



And here is the same scene uses our traditional pure rasterization rendering pipeline.

As you can see, most of our visual fidelity holds up between the two, especially with respect to our materials.

Guiding Principles

- No concept of a classic “frame”
- Rendering occurs at variable frequency
- No promise that data is locally resident
 - Design for latency
 - Design for massive data

Some guiding principles behind our architecture is that there is no concept of a classic “frame”, and rendering occurs at variable frequency.

In addition to the frequency, there is also no promise that data is locally resident, requiring us to design our systems to handle latency, and also massive data.

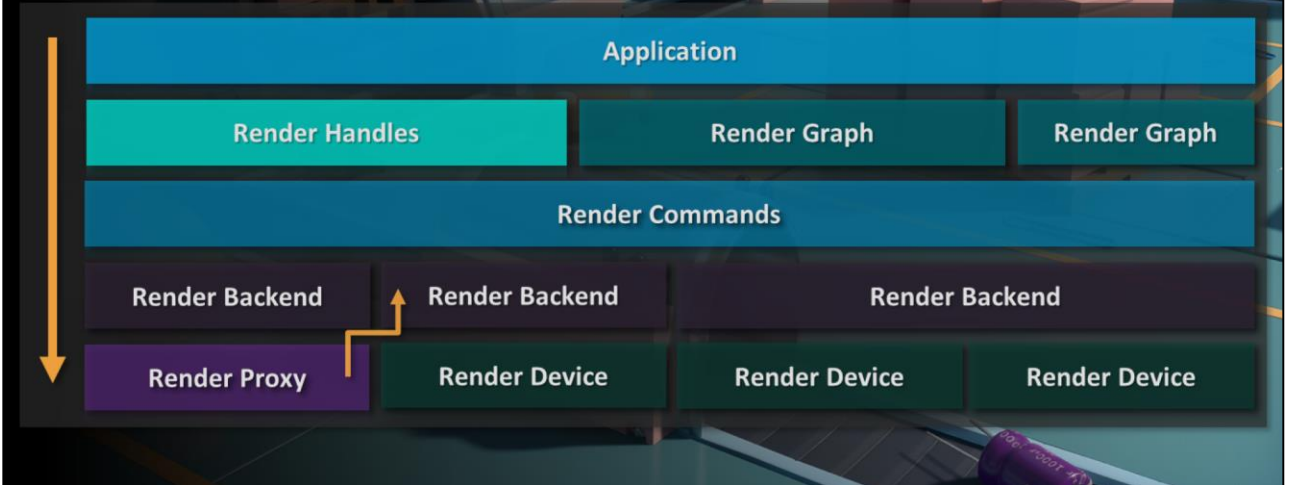
Guiding Principles

- “Separation of Concerns”
 - Design for a single purpose
 - No monolithic systems
- Verbosity hidden with layers
 - Explicit low-level API
 - Convenience APIs on top

I am a big proponent of “separation of concerns”, where systems are designed for a single purpose, and monolithic systems are avoided at all cost.

Another principle is that verbosity is hidden with layers. We provide an explicit low-level API to users, which can be used directly if desired, but then there are convenience APIs on top to avoid the “5000 lines of code to render a triangle” problem.

Rendering Layers



There are a number of rendering layers making up our architecture. The application sits on top (for example, PICA PICA) and provides very high level abstractions on top of Halcyon itself.

Underneath the application are a number of convenience APIs like render graph, and under those are the explicit APIs like render commands going to the various render backend and device implementations.



The first layer I will talk about is the render backend

Render Backend

- **Live-reloadable DLLs**
- Enumerates adapters and capabilities
 - Swap chain support
 - Extensions (i.e. ray tracing, sub groups, ...)
 - **Determine adapter(s) to use**

Render backends are implemented as live-reloadable DLLs. Each backend represents a particular API, and provides enumeration of adapters and capabilities.

In addition, the backends will determine the ideal adapters to use, depending on the desired purpose.

Render Backend

- Provides debugging and profiling
 - RenderDoc integration, validation layers, ...
- Create and destroy render devices

Each render backend also supports a debugging and profiling interface – this provides functionality like RenderDoc integration, CPU and GPU validation layers, etc..

Most importantly, render backends support the creation and destruction of render devices, which I'll cover shortly.

Render Backend

- Direct3D 12
- Vulkan 1.1
- Metal 2
- Proxy
- Mock

We have a variety of render backends implemented



Render Backend

- Direct3D 12
 - Shader Model 6.X
 - **DirectX Ray Tracing**
 - Bindless Resources
 - Explicit Multi-GPU
 - **DirectML**
 - ...

For Direct3D 12, we support the latest variant of shader model 6, DirectX ray tracing, full bindless resources, explicit heterogeneous mGPU, and DirectML.

Render Backend

- Vulkan 1.1
 - Sub-groups
 - Descriptor indexing
 - External memory
 - Multi-draw indirect
 - NV Ray tracing
 - ...


Vulkan is a similar story to Direct3D 12 – it supports basically the same features, but it has the added benefit of being cross platform.

Additionally, it is much easier to work with GPU vendors to get access to hardware features through Vulkan's extension system than it is with Direct3D 12, making it an ideal render backend for high end research and development.

Render Backend

- Metal 2
 - Early development
 - Primarily desktop
 - Argument buffers
 - Machine learning
 - ...

Metal 2 is still in early development, but very promising.



Render Backend

- Proxy
 - Discussed later in the presentation



We have another pretty crazy backend called render proxy, which I discuss later in this presentation.

Render Backend

- Mock
 - Performs resource tracking and validation
 - Command stream is parsed and evaluated
 - No submission to an API
 - Useful for unit tests and debugging

Finally, we have our mock render backend which is for unit testing, debugging, and validation. This backend does all the same work the other backends do, except translation from high level to low level just runs a validation test suite instead of submitting to a graphics API.

Render Device

The next layer to discuss is render device

Render Device

- **Abstraction** of a logical **GPU adapter**
 - e.g. VkDevice, ID3D12Device, ...
- Provides **interface to GPU queues**
- Command list **submission**

Render device is an abstraction of a logical GPU adapter, such as VkDevice or ID3D12Device.

It provides an interface to various GPU queues, and provides API specific command list scheduling and submission functionality

Render Device

- Ownership of GPU resources
 - Create & Destroy
- Lifetime tracking of resources
- Mapping render handles → device resources

Render device has full ownership and lifetime tracking of its GPU resources, and provides create and destroy functionality for all the high level render resource types.

The high level application refers to these resources using handles, so render device also provides an efficient mapping of these handles to internal device resources.



First off, I will describe what our render handles are

Render Handles

- Resources associated by handle
- Lightweight (64 bits)
- Constant-time lookup
- Type safety (i.e. buffer vs texture)
- Can be serialized or transmitted
- Generational for safety
 - e.g. double-delete, usage after delete

Our rendering resources are associated by handle.

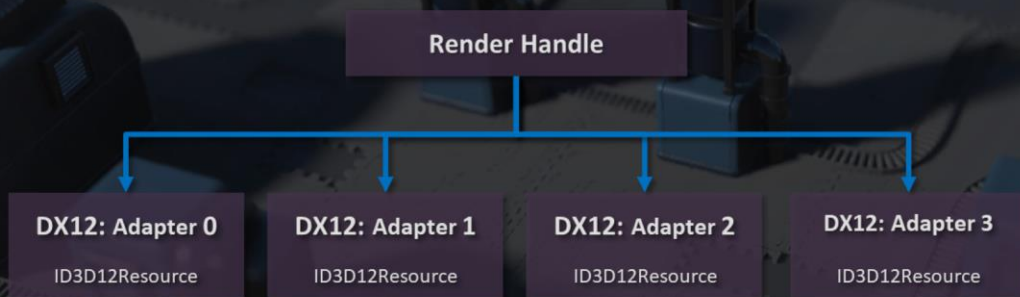
The handles are lightweight (just 64 bits), and the associated resources are fetched with a constant-time lookup. The handles are type safe, protecting against errors like trying to pass a buffer in place of a texture.

The handles can be serialized or transmitted, as long as some contract is in place about what the handle represents.

And the handles are generational for safety, which protects against double-delete, or using a handle after deletion

Render Handles

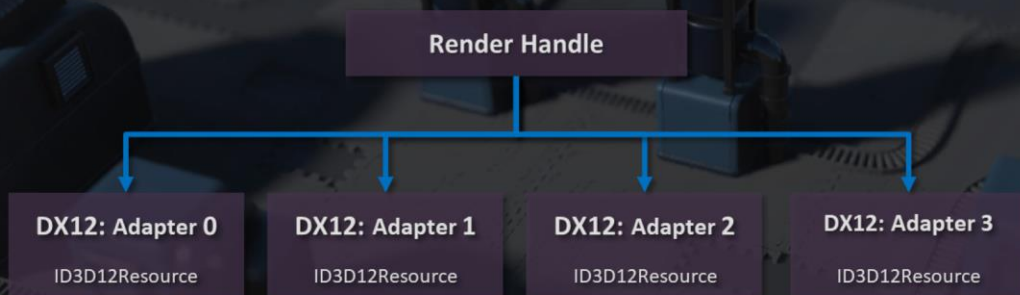
- Handles allow **one-to-many cardinality** *[handle->devices]*
- Each device can have a unique representation of the handle



Handles allow a one-to-many cardinality where a single handle can have a unique representation on each device

Render Handles

- Can query if a device has a handle loaded
- Safely add and remove devices
 - Handle owned by application, representation can reload on device



There is an API to query if a particular render device has a handle loaded or not. This makes it safe to add and remove devices while the application is running.

The handle is owned by the application, so the representation can be loaded or reloaded for any number of devices.

Render Handles

- Shared resources are supported
- Primary device owner, secondaries alias primary



Shared resources are also supported, allowing for copying data between various render devices.

Render Handles

- Can also **mix and match backends** in the **same process!**
 - Makes debugging new backends much easier
 - **DX12 on left half** of screen, **VK on right half** of screen



A crazy feature of this architecture is that we can mix and match different backends in the same process!

This made implementing and debugging new backends like Vulkan much easier, as I could have DX12 rendering on the left half of the screen, while Vulkan was rendering on the right half of the screen.



Another key rendering component is our high level command stream

Render Commands

- Draw
- DrawIndirect
- Dispatch
- DispatchIndirect
- UpdateBuffer
- UpdateTexture
- CopyBuffer
- CopyTexture
- Barriers
- Transitions
- BeginTiming
- EndTiming
- ResolveTimings
- BeginEvent
- EndEvent
- BeginRenderPass
- EndRenderPass
- RayTrace
- UpdateTopLevel
- UpdateBottomLevel
- UpdateShaderTable

We developed an API agnostic high level command stream that allows the application to efficiently express how a scene should be updated and rendered, but allowing each backend to control how this is done, including render state changes and resource barriers or transitions.

Render Commands

- Queue type specified
- Spec validation
 - Allowed to run?
 - *e.g. draws on compute*
- Automatic scheduling
 - Where can it run?
 - Async compute

```
enum class RenderCommandQueueType : uint8
{
    None = 0x00,
    Copy = 0x01,
    Compute = 0x02,
    Graphics = 0x04,
    All = Copy | Compute | Graphics,
};

struct RenderCommand
{
    RenderCommandType type = RenderCommandType::Count;
};

template<RenderCommandType TYPE, RenderCommandQueueType QUEUETYPE>
struct RenderCommandTyped : RenderCommand
{
    static const RenderCommandType Type = TYPE;
    static const RenderCommandQueueType QueueType = QUEUETYPE;

    RenderCommandTyped() { type = Type; }
};
```

Each render command specifies a queue type, which is primarily for spec validation (such as putting graphics work on a compute queue), and also to aid in automatic scheduling, like with async compute.

Render Commands

```
struct RenderCommandDispatch : RenderCommandTyped<RenderCommandType::Dispatch, RenderCommandQueueType::Compute>
{
    RenderResourceHandle pipelineState;
    ShaderArgument shaderArguments[MaxShaderParameters];
    uint32 shaderArgumentsCount = 0;

    uint32 dispatchX = 0;
    uint32 dispatchY = 0;
    uint32 dispatchZ = 0;
};
```

As an example, here is a compute dispatch command, which specifies Compute as the queue type. The underlying backend and scheduler can interpret this accordingly

Render Command List

- Encodes high level commands
- Tracks queue types encountered
 - Queue mask indicating scheduling rules
- Commands are stateless - parallel recording

The high level commands are encoded into a render command list. As each command is encoded, a queue mask is updated which helps indicate the scheduling rules and restrictions for that command list.

The commands are stateless, which allows for fully parallel recording.

Render Compilation

- Render command lists are “compiled”
 - Translation to low level API
 - Can compile once, submit multiple times
- Serial operation (memcpy speed)
 - Perfect redundant state filtering

The render command lists are compiled by each render backend, which means the high level commands are translated to the low level API. Command lists can be compiled once, and submitted multiple times, in appropriate cases.

The low level translation is a serial operation by design, as this is basically running at memcpy speed. Doing this part serial means we get perfect redundant state filtering.

Vulkan Compilation

- Translate **commands**
 - Read command list
 - Write Vulkan API

```
#define COMPILER_PACKET(TYPE_STRUCT)
case TYPE_STRUCT::Type:
    if (!compileCommand(*static_cast<const TYPE_STRUCT*>(command)))
        return false;
    break

for (const auto* command : recorded)
{
    switch (command->type)
    {
        COMPILER_PACKET(RenderCommandDraw);
        COMPILER_PACKET(RenderCommandDrawIndirect);
        COMPILER_PACKET(RenderCommandDispatch);
        COMPILER_PACKET(RenderCommandDispatchIndirect);
        COMPILER_PACKET(RenderCommandUpdateBuffer);
        COMPILER_PACKET(RenderCommandUpdateTexture);
        COMPILER_PACKET(RenderCommandCopyBuffer);
        COMPILER_PACKET(RenderCommandCopyTexture);
        COMPILER_PACKET(RenderCommandBarriers);
        COMPILER_PACKET(RenderCommandTransitions);
        COMPILER_PACKET(RenderCommandBeginTiming);
        COMPILER_PACKET(RenderCommandEndTiming);
        COMPILER_PACKET(RenderCommandResolveTimings);
        COMPILER_PACKET(RenderCommandBeginEvent);
        COMPILER_PACKET(RenderCommandEndEvent);
        COMPILER_PACKET(RenderCommandBeginRenderPass);
        COMPILER_PACKET(RenderCommandEndRenderPass);
        break;
    default:
        HcyFail();
        return false;
    }
}

#undef COMPILER_PACKET
```

Another important aspect of a new render backend is the translation from our high level command stream to low level API calls.

S E E D // Halcyon: Rapid Innovation using Modern Graphics

```
bool RenderCompileContextVulkan::compileCommand(const RenderCommandDispatch& command)
{
    const auto pipelineState = device.getComputePipelineState(command.pipelineState);
    if (!applyPipelineState(pipelineState))
    {
        return false;
    }

    applyShaderArguments(command.pipelineState, command.shaderArguments, command.shaderArgumentsCount);
    applyTransitions();

    vkCmdDispatch(commandBuffer, command.dispatchX, command.dispatchY, command.dispatchZ);
    return true;
}
```

Here is an example of translating a high level compute dispatch command to Vulkan.


```
bool RenderCompileContextVulkan::compileCommand(const RenderCommandBeginTiming& command)
{
    auto timingHeap = const_cast<RenderTimingHeapVulkan*>(device.getTimingHeap(command.timingHeap));
    HcyAssert(timingHeap);
    const uint32 slot = 2 * command.region + 0;
    HcyAssert(slot >= 0 && slot < (timingHeap->desc.regionCount * 2));
    auto& buffer = timingHeap->buffers[timingHeap->currentBuffer];
    vkCmdWriteTimestamp(
        commandBuffer,
        VK_PIPELINE_STAGE_BOTTOM_OF_PIPE_BIT,
        buffer.queryPool,
        slot);
    buffer.regions[command.region].begin = slot;
    return true;
}
```

And an example of translating a high level begin timing command to Vulkan timestamps

```

bool RenderCompileContextVulkan::compileCommand(const RenderCommandResolveTimings& command)
{
    auto timingHeap = const_cast<RenderTimingHeapVulkan*>(device.getTimingHeap(command.timingHeap));
    HcyAssert(timingHeap);
    HcyAssert(command.regionCount > 0);
    HcyAssert(command.regionCount <= timingHeap->desc.regionCount);
    auto& buffer = timingHeap->buffers[timingHeap->currentBuffer];
    buffer.resolveStart = buffer.regions[command.regionStart].begin;
    buffer.resolveCount = (command.regionCount * 2);

    vkCmdCopyQueryPoolResults(
        commandBuffer,
        buffer.queryPool,
        buffer.resolveStart,
        buffer.resolveCount,
        buffer.readBack->buffer(),
        buffer.resolveStart * sizeof(uint64),
        sizeof(uint64),
        VK_QUERY_RESULT_64_BIT | VK_QUERY_RESULT_WAIT_BIT);

    buffer.writeFence->signalGpu(queue.commandQueue());

    vkCmdResetQueryPool(commandBuffer, buffer.queryPool, buffer.resolveStart, buffer.resolveCount);

    timingHeap->previousBuffer = timingHeap->currentBuffer;
    timingHeap->currentBuffer = (device.m_frameIndex % timingHeap->buffers.size());


    return true;
}

```

And here is an example of translating a high level resolve timings command to Vulkan.

Notice that the complexity behind fence tracking or resetting the query pool is completely hidden from the calling code.

Each backend implementation can choose how to handle this efficiently.



Another significant rendering layer is our render graph

Render Graph

- Inspired by FrameGraph [O'Donnell 2017]
- Automatically handle transient resources
- Import explicitly managed resources
- Automatic resource transitions
 - Render target batching
 - DiscardResource
 - Memory aliasing barriers
 - ...

Render Graph is inspired by Frostbite's Frame Graph

Just like Frame Graph, we automatically handle transient resources. Complex rendering uses a number of temporary buffers and textures, that stay alive just for the purpose and duration of an algorithm. These are called transient resources, and the graph automatically manages these lifetimes efficiently.

Automatic resource transitions are also performed

Render Graph

- **Frame Graph** → **Render Graph**: No concept of a “frame”
- Fully **automatic transitions** and **split barriers**
- Single implementation, regardless of backend
 - Translation from high level render command stream
 - API differences hidden from render graph
- **Support for mGPU**
 - Mostly implicit and automatic
 - Can specify a scheduling policy


We call our implementation Render Graph, because we don't have the concept of a “frame”, as mentioned earlier.

Our transitions and split barriers are fully automatic, and there is a single implementation, regardless of backend. This is thanks to our high level render command stream, which hides API differences from render graph.

We also support multi-gpu, which is mostly implicit and automatic, with the exception of a scheduling policy you configure on your pass – this represents what devices a pass runs on, and what transfers are performed going in or out of the pass.

Render Graph

- Composition of multiple graphs at varying frequencies
 - Same GPU: async compute
 - mGPU: graphs per GPU
 - Out-of-core: server cluster, remote streaming



Render graph supports composition of multiple graphs at varying frequencies.

These graphs can run on the same GPU - such as async compute.

They can run with multi-gpu, with a graph on each GPU.

And they can even run out of core, such as in a server cluster, or on another machine streamed remotely.

Render Graph

- Composition of multiple graphs at varying frequencies
 - e.g. translucency, refraction, global illumination



There are a number of techniques that can easily run at a different frequency from the main graph, such as object space translucency and refraction, or our surfel based global illumination.

Render Graph

- Two phases
- **Graph construction**
 - Specify inputs and outputs
 - Serial operation (by design)
- **Graph evaluation**
 - Highly parallelized
 - Record high level render commands
 - Automatic barriers and transitions

Render graph runs in two phases, construction and evaluation.

Construction specifies the input and output dependencies, and this is a serial operation by design.

If you implement a hierarchical gaussian blur as a single pass that gets chained X times, you want to read the size of the input in each step, and generate the proper sized output. Maybe you could split it up into some threads, but tracking dependencies in order to do construction in a parallel fashion might be more costly than just running it serially.

Evaluation is highly parallelized, and this is where high level render commands are recorded, and automatic barriers and transitions are inserted.

```

renderGraph.addPassCallback("Present Pass", [&](RenderGraphBuild& build)
{
    auto& outputTexData = scope.get<RenderGraphOutputTexture>();
    const auto& viewData = scope.get<RenderGraphViewData>();
    const auto& custom = scope.getOptional<RenderGraphCustomFinalTexture>();

    auto finalTexture = build.read(
        custom ? custom->finalTexture
        : scope.get<RenderGraphFinalTexture>().finalTexture, RenderBindFlags::ShaderResource);

    auto outputTexture = outputTexData.outputTexture = build.write(outputTexData.outputTexture, RenderBindFlags::UnorderedAccess);

    return [=](RenderGraphRegistry& registry, RenderCommandList& commandList)
    {
        RenderPoint point = {};
        RenderBox box = {};
        box.w = viewData.viewWidth;
        box.h = viewData.viewHeight;

        commandList.copyTexture(
            registry.getTexture(outputTexture),
            0,
            point,
            registry.getTexture(finalTexture),
            0,
            box);
    };
});

```


← Construction phase

← Evaluation phase

Here is an example render graph pass, with the construction phase up top, and the evaluation phase at the bottom.

Render Graph

- Explicit heterogeneous mGPU
- Parallel fork-join approach
- Resources copied through system memory using copy queue
 - ~1ms for every 15mb transferred
- Minimize PCI-E transfers
 - Immutable data replicated
 - Tightly pack data




As mentioned, we support explicit heterogeneous multi-GPU. We use a parallel fork-join approach, and we copy resources through system memory using the copy queue.

It costs roughly 1ms to transfer 15mb of data, so it's important to minimize how much we transfer.

This is done by redundantly replicating immutable data to all GPUs (such as meshes and textures), and also tightly packing or compressing data to minimize the transfer size.

Render Graph

- Workloads are divided into **partitions**
 - Based on GPU device count
- **Single primary device**
- **Other devices are secondaries**
- Variety of scheduling and transfer patterns are necessary
- **Simple rules engine**



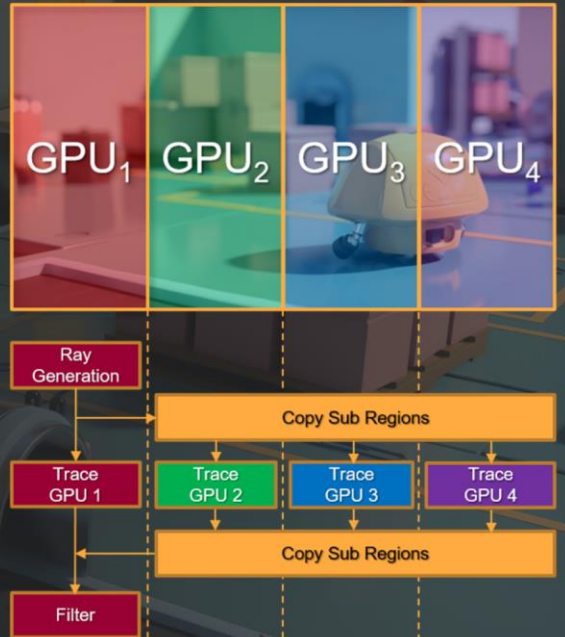
Our workloads are divided into partitions, where the number of partitions is based on the number of available GPUs.

We designate one of the GPUs as the primary device, and the other devices are designated as secondaries.

With complex rendering and computation, there are a variety of scheduling and transfer patterns needed. We built a simple rules engine to drive the logic of our workloads.

Render Graph

- Run ray generation on primary GPU
- Copy results in sub-regions to other GPUs
- Run tracing phases on separate GPUs
- Copy tracing results back to primary GPU
- Run filtering on primary GPU



For our ray tracing workloads, the common pattern is to run the ray generation on the primary GPU. Then we copy slices or sub-regions of this data to the other secondary GPUs, where each GPU performs tracing on its own sub-region or partition.

The tracing results are then copied back to the primary GPU, and then filtering is performed exclusively on the primary GPU.

This approach avoids many problems with temporal techniques, and also means we don't need to transfer as much data, such as the full g-buffer, to the secondary GPUs.

Render Graph

- Only width is divided
- Simplifies textures vs. buffers
- Passes are unaware of GPU count

```
glm::vec4 partitionWindow = scheduleData.mgpuShadows ? registry.partitionIsolated() : registry.partitionAll();
const uint32 dispatchOffset = uint32(desc.width * partitionWindow.x);
const uint32 dispatchWidth = uint32(desc.width * partitionWindow.z);

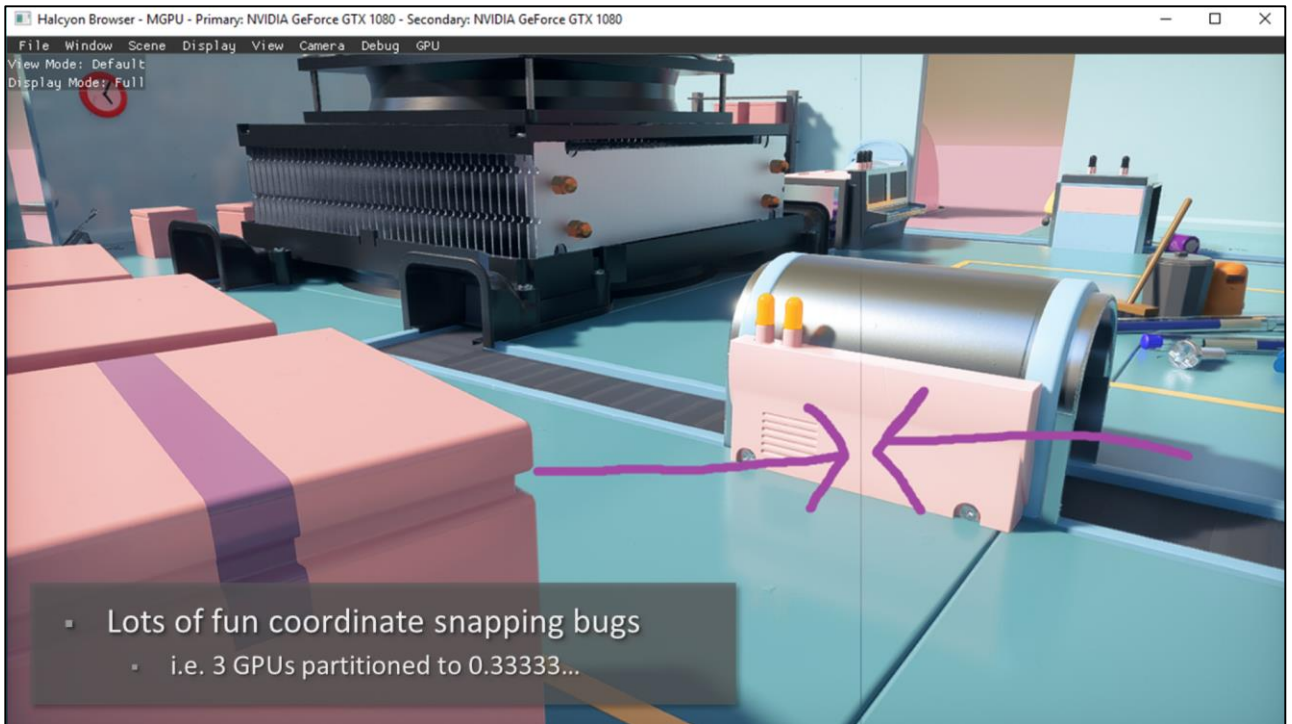
// ...

commandList.dispatch2d(
    pipelineState,
    { ShaderArgument(dynamicConstants.buffer, constantsOffset, registry.createShaderViews(srvs, uavs)) },
    dispatchWidth, desc.height
);
```

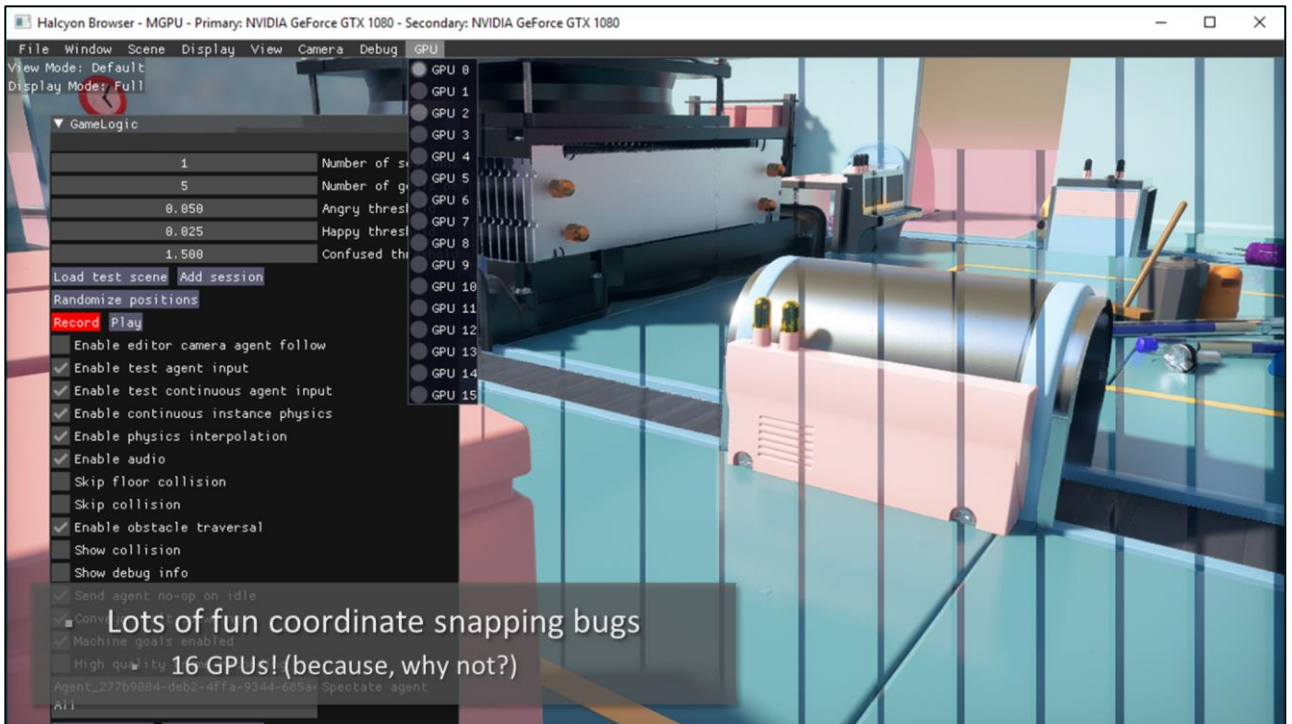
The workloads are only divided into partitions based on the X axis, or width. This simplifies textures vs. buffers, since we can treat all sub-regions as 1D instead of 2D.

Passes are unaware of the GPU count, which keeps the code clean, and avoids any edge cases or bugs with untested mGPU configurations.

This code snippet shows all that is needed for a pass to scale the workloads, regardless of how many GPUs are in use.



The automatic scaling window works quite well for us, but there were lots of fun coordinate snapping bugs to fix up in the beginning, like 3 GPUs partitioning to 0.3 repeating.



And then some interesting bugs when we did crazy configurations, like 16 GPUs, because why not? ☺

Render Graph

- **RenderGraphSchedule**

- **NoDevices** → Pass is disabled
- **AllDevices** → Pass runs on all devices
- **PrimaryDevice** → Pass only runs on primary device
- **SecondaryDevices** → Pass runs on secondaries if count > 1, otherwise primary
- **OnlySecondaryDevices** → Pass only runs on secondary devices, disabled unless mGPU

Requested **Per Pass** →

```
void setSchedule(RenderGraphSchedule passSchedule);  
  
void setTransfer(  
    RenderGraphResource resource,  
    RenderTransferPartition src,  
    RenderTransferPartition dst,  
    RenderTransferFilter dstFilter);
```

Our simple rules engine allows for a pass to specify the scheduling, and also request transfers in or out of the pass.

Render Graph

- **RenderTransferPartition**
 - **PartitionAll** → Select all partitions from device
 - **PartitionIsolated** → Select isolated region from device
- **RenderTransferFilter**
 - **AllDevices** → Transfer completes on all devices
 - **PrimaryDevice** → Transfer completes on the primary device
 - **SecondaryDevices** → Transfer completes on all secondary devices

Requested **Per Pass** →

```
void setSchedule(RenderGraphSchedule passSchedule);  
  
void setTransfer(  
    RenderGraphResource resource,  
    RenderTransferPartition src,  
    RenderTransferPartition dst,  
    RenderTransferFilter dstFilter);
```

The transfers specify what partitions are copied from source to destination.

And the transfers also specify rules for the transfer destination, allowing for basic copies, multi-cast copies, etc..

Render Graph

- **PartitionAll** → **PartitionAll**
 - Copies full resource on one GPU to full resource on all specified GPUs
- **PartitionAll** → **PartitionIsolated**
 - Copies full resource on one GPU to isolated regions on all specified GPUs (**partial copies**)
- **PartitionIsolated** → **PartitionAll**
 - (**Invalid** configuration)
- **PartitionIsolated** → **PartitionIsolated**
 - Copies isolated region on one GPU to isolated regions on all specified GPUs (**partial copies**)

The transfer partition is specified for both the source and the destination GPU, allowing for a variety of transfer patterns.


```

renderGraph.addPassCallback("Shadow Mask Pack", [&](RenderGraphBuild& build)
{
    build.setSchedule(RenderGraphSchedule::AllDevices); // ← Devices this pass will run on

    // ...

    build.setTransfer(
        outputTexture,
        RenderTransferPartition::PartitionIsolated, // ← Schedule transfers in or out
        RenderTransferPartition::PartitionIsolated,
        RenderTransferFilter::PrimaryDevice);

    auto pipelineState = pipelines.pipelineState(ShaderPipelineId::ShadowMaskPack);

    return [=](RenderGraphRegistry& registry, RenderCommandList& commandList)
    {
        glm::vec4 partitionWindow = scheduleData.mgpuShadows ? registry.partitionIsolated() : registry.partitionAll();
        const uint32 dispatchOffset = uint32(desc.width * partitionWindow.x);
        const uint32 dispatchWidth = uint32(desc.width * partitionWindow.z); // ← Scaling work dimensions for each GPU

        // ...

        commandList.dispatch2d(
            pipelineState,
            { ShaderArgument(dynamicConstants.buffer, constantsOffset, registry.createShaderViews(srvs, uavs)) },
            dispatchWidth, desc.height
        );
    };
});

```


Here is an example render graph pass which specifies the mGPU scheduling policy, in this case, to run on all available devices.

This example also schedules a transfer of the resultant data back to the primary device. Isolated to isolated means that each device will only copy the sub-region or partition that it was responsible for computing.

In the bottom half of the code, you can see the evaluation phase which is scaling the work dimensions for the compute dispatch, based on the mGPU configuration.




When developing mGPU support for our render graph passes, some scheduling or transfer bugs were obvious.



Here's another example

Render Graph

- Some bugs were subtle
 - Weird cell shading? ☺
- Incorrect **transfers**?
 - Transfers in (input data)
 - Transfers out (result data)
- Incorrect **scheduling**?
 - Pass not running
 - Pass running when it shouldn't
 - Partition window



And then some bugs were very subtle, like this weird cell shading.

Incorrect transfers would result in input data being incorrect or uninitialized, or result data not being copied back to the main GPU for filtering.

Incorrect scheduling would cause passes to not run, to run when they shouldn't, or have an incorrect partitioning window, interfering with the work of another GPU.

Render Graph

Some of our render graph passes:

- Bloom
- BottomLevelUpdate
- Brdflut
- CocDerive
- DepthPyramid
- DiffuseSh
- Dof
- Final
- GBuffer
- Gtao
- IblReflection
- ImGui
- InstanceTransforms
- Lighting
- MotionBlur
- Present
- RayTracing
- RayTracingAccum
- ReflectionFilter
- ReflectionSample
- ReflectionTrace
- Rtao
- Screenshot
- Segmentation
- ShaderTableUpdate
- ShadowFilter
- ShadowMask
- ShadowCascades
- ShadowTrace
- Skinning
- Ssr
- SurfelGapFill
- SurfelLighting
- SurfelPositions
- SurfelSpawn
- Svgf
- TemporalAa
- TemporalReproject
- TopLevelUpdate
- TranslucencyTrace
- Velocity
- Visibility

This is a list of some of our render graph passes in PICA PICA – we ended up adding mGPU support to reflections, shadows, lighting, global illumination, and ambient occlusion.

Render Graph

- Implicit data flow via explicit **scopes**
 - “Long-distance” extensible parameter passing
- **Scope given to each render pass**
 - Can create **nested** scope for sub-graph
 - Results stored into scope
- Hygiene via nesting and **shadowing**

```
{  
  gbuffer <- render_opaque()  
  gbuffer <- render_decals(gbuffer)  
  
  {  
    gbuffer <- render_opaque()  
    render_lighting(gbuffer)  
  } -> envmap  
  
  apply_envmap(gbuffer, envmap)  
}
```

```
struct RenderGraphAreaLight  
{  
  RenderGraphResource triangleLightList;  
  uint32 triangleCount;  
};
```

The initial implementation of render graph would explicitly chain input and output dependencies between passes, but this lead to messy coupling, and also explicit checks like if a depth pre-pass has already run, import the depth target, and if not, clear the depth target or disable depth.

We improved the implementation to support implicit data flow using explicit scopes. This allows for long-distance extensible parameter passing.

A scope is given to each render pass, which supports nested scopes for sub-graphs, and the results are stored in the scope.

Hygiene is provided via nesting and shadowing.

Render Graph

- Lookup by **type**
 - `scope.get<T>() -> &T`
- Parameters in “plain old data” structs
 - `RenderGraphResource`, `RenderHandle`
 - `float`, `int`, `mat4`, etc.

```
{  
  gbuffer <- render_opaque()  
  gbuffer <- render_decals(gbuffer)  
  
  {  
    gbuffer <- render_opaque()  
    render_lighting(gbuffer)  
  } -> envmap  
  
  apply_envmap(gbuffer, envmap)  
}
```

```
struct RenderGraphAreaLight  
{  
  RenderGraphResource triangleLightList;  
  uint32 triangleCount;  
};
```

The scopes can be looked up by type, such as depth or gbuffer resources. The lookup returns parameters stored in plain old data structs.

Render Graph DSL

- Experimental
- Macro Magic

```
_def_pass(Present) {  
    _use_pipeline(PresentCs) {  
        auto& output = _scope(RenderGraphOutputTexture);  
        const auto viewData = _scope(RgViewData);  
  
        auto inputTexture = rg.read(_scope(RgMainTexture).texture, RenderBindFlags::ShaderResource);  
        output.texture = rg.write(output.texture, RenderBindFlags::UnorderedAccess);  
  
        _when_drawn {  
            _def_srvs(  
                _res_tex2d(inputTexture),  
            );  
  
            _def_uavs(  
                _res_tex2d(output.texture),  
            );  
  
            commandList.dispatch2d(  
                _shader_args,  
                viewData.viewWidth,  
                viewData.viewHeight,  
            );  
        };  
    };  
}
```

We have also been experimenting with various implementations of a DSL for Render Graph – currently it's using hacky macro magic, but we're thinking about writing a code generator using Rust, relying on many amazing features of the language to develop our DSL.

The goal here is to make render graph fully data driven, and even serializable.

SEED // Halcyon: Rapid Innovation using Modern Graphics

Render Graph

- Automatic profiling data
- GPU and CPU counters per-pass
- Works with mGPU
 - Each GPU is profiled



Render graph can collect automatic profiling data, presenting you with GPU and CPU counters per-pass

Additionally, this works with mGPU, where each GPU is shown in the list

Render Graph

- Live debugging overlay
- Evaluated passes in-order of execution
- Input and output dependencies
- Resource version information

There is also a live debugging overlay, using ImGui. This overlay shows the evaluated render graph passes in-order of execution, the input and output dependencies, and the resource version information.

Halcyon Browser - SGPU - NVIDIA TITAN V

File Window Scene Display View Camera Debug GPU

View Mode: Default
Display Mode: Full

Render Graph

Pass: Reflection Filter Pass

Inputs	Outputs
Reflection Filter Uniforms (version:0)	Reflection Filter Uniforms (version:1)
Reflection Temporal 0 (version:0)	Reflection Temporal 0 (version:1)
Reflection Temporal 1 (version:0)	Reflection Meta Temporal 0 (version:1)
Reflection Meta Temporal 0 (version:0)	Reflection Output Texture (version:1)
Reflection Meta Temporal 1 (version:0)	Reflection Meta Output Texture (version:1)
Reflection Hit0 Texture (version:2)	Reflection Variance Output Texture (version:1)
Reflection Hit1 Texture (version:2)	Reflection Velocity Texture (version:1)
Reflection Output Texture (version:0)	

Pass: Reflection Filter 2 Pass

Inputs	Outputs
Reflection Filter Uniforms (version:0)	Reflection Filter Uniforms (version:1)
Reflection Output Texture (version:1)	Reflection Output Texture (version:1)
Reflection Meta Output Texture (version:1)	
Reflection Variance Output Texture (version:1)	
Reflection Output Texture (version:0)	
Gbuffer History 0 (version:2)	
Depth (version:2)	
Default Texture (version:0)	

Pass: Reflection Merge Pass

Inputs	Outputs
Reflection Output Texture (version:1)	Lighting (version:2)
Lighting (version:1)	

Pass: Temporal AA Pass

Inputs	Outputs
Lighting (version:2)	TAA History Texture 1 (version:1)
TAA History Texture 0 (version:0)	TAA Output Texture (version:1)
Temporal Reprojection Texture (version:1)	
TAA History Texture 1 (version:0)	
TAA Output Texture (version:0)	

Pass: Velocity reduce

Inputs	Outputs
Velocity Texture 0 (version:2)	Reduced velocity texture (version:1)
Reduced velocity texture (version:0)	

Pass: Velocity reduce

Inputs	Outputs
Reduced velocity texture (version:1)	Reduced velocity texture (version:1)
Reduced velocity texture (version:0)	

Pass: Velocity dilate

Inputs	Outputs
Reduced velocity texture (version:1)	Dilated velocity texture (version:1)
Dilated velocity texture (version:0)	

Pass: Motion blur

Inputs	Outputs
Depth (version:2)	Motion blurred image (version:1)
Velocity Texture 0 (version:2)	
Dilated velocity texture (version:1)	
TAA Output Texture (version:1)	
Motion blurred image (version:0)	

Pass: Bloom Extract Pass

THE METAVERSE
A Simple Game
Author: TBD



Virtual Multi-GPU

In order to make multi-GPU development easier, we built something we call virtual multi-GPU

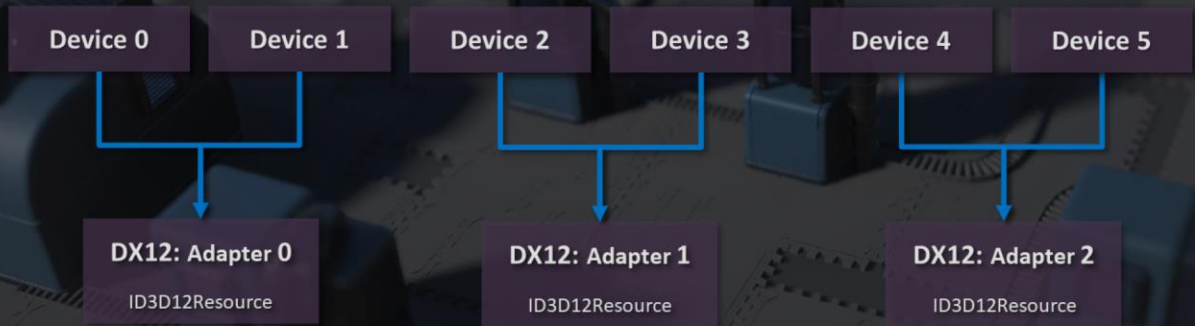
Virtual Multi-GPU

- Most developers have single GPU
- Uncommon for 2 GPU machines
- Rare for 3+ GPU
 - Practical for show floor and cranking up to 11
 - Impractical for regular development 😊

Most developers typically have just a single GPU in their workstations. It is quite uncommon for 2 GPU machines, and it is rare for machines to have more than 2 GPUs. Having a huge machine is practical for the show floor, and cranking settings up to 11, but this is impractical for regular development. Testing these configurations becomes challenging.

Virtual Multi-GPU

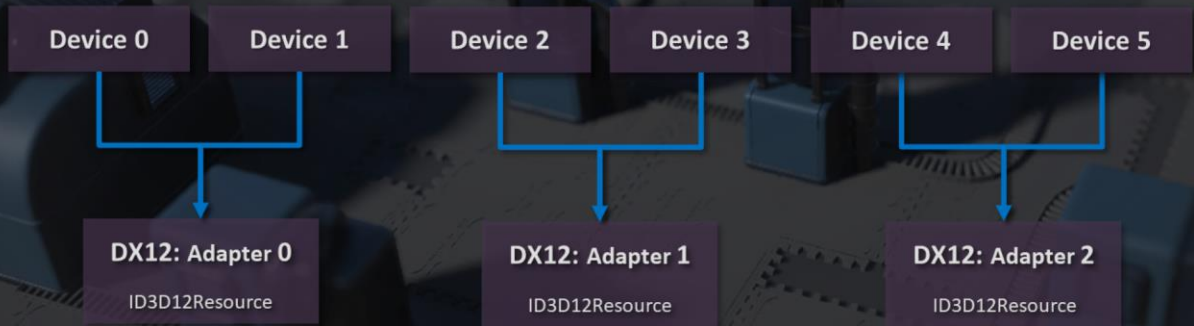
- Build **device indirection table**
- Virtual device index → adapter index



Our solution was to build a device indirection table. We enumerate available device adapters, and based on a setting specifying how many copies of a device are desired, we redirect our own virtual device indices to actual adapter indices.

Virtual Multi-GPU

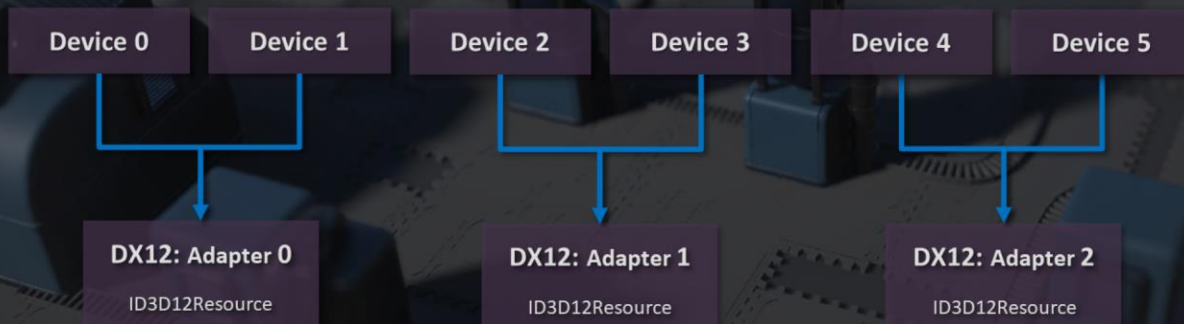
- Create multiple instances of a device
- **Virtual GPUs execute sequentially** (WDDM)



We do this by creating multiple instances of a given device. Similar to multiple applications sharing a GPU, the OS (such as WDDM on Windows 10) will execute multiple instances of a device sequentially.

Virtual Multi-GPU

- Increases overall wall time (don't use for profiling)
- **Amazing for development and testing!**



This approach increases the overall wall time, so don't use it for end to end profiling. However, this is amazing for development and testing, even estimating the possible performance characteristics of the mGPU algorithm in isolation.

Virtual Multi-GPU

- PICA PICA developers all had 1 GPU
- Limited testing with 2 GPUs
- Show floor at GDC 2018 was 4 GPUs
 - Virtual-only testing...
 - Crossed fingers
 - Worked flawlessly!

There is a fun story around this stuff.. For PICA PICA, all developers only had a single GPU, and there was very limited testing with 2 GPUs.

The show floor at GDC 2018 was a machine with 4 GPUs, a configuration we had only ever tested using virtual mGPU. We crossed our fingers, and hoped we didn't make a horrible mistake by relying exclusively on virtual mGPU for testing.

The demo worked flawlessly on the first try with the actual hardware, and gave us the wall time improvements we had estimated.

Virtual Multi-GPU

- Develop and debug **multi-GPU with only a single GPU**
- Virtual mGPU reliably reproduces most bugs!
- Entire **features developed without physical mGPU**
 - i.e. Surfel GI (the night before GDC.. 😊)

Virtual mGPU allows us to develop multi-GPU with only a single GPU

Virtual mGPU even reproduces most mGPU bugs in the same way, which is incredibly useful.

Some developers never tested on physical mGPU, and developed entire features like our surfel GI using only virtual mGPU. (The night before GDC, I might add..)



Another interesting capability of our architecture is our render proxy

Render Proxy

- Remote render backend
- Any API / Any OS



Render proxy is another render backend implementation, like DirectX12 or Vulkan, except this allows us to mix any API with any OS.

Render Proxy

- **Render API calls are routed remotely**
- Uses gRPC (high performance RPC framework)
- Use an API on an incompatible OS
 - e.g. Direct3D 12 on macOS or Linux

The way it works is our render API calls get routed remotely using gRPC, which is a high performance RPC framework developed by Google.

This lets us use an API on an incompatible OS, such as using Direct3D 12 on macOS or Linux.

This is especially nice for bringing up new platforms, as we can get base memory, IO, and core functionality running, without initially worrying about the graphics system.

Render Proxy

- **Scale large workloads** with a GPU cluster
 - Some API as render graph mGPU
- Only rendering is routed, **scene state is local**
- **Work from the couch!**
 - i.e. NV ray tracing on a MacBook 😊

Render proxy allows us to scale large workloads with a GPU cluster. Render graph passes that support mGPU can automatically scale to the cluster without any additional code, which is a really powerful design.

Another interesting property is that only the rendering is routed, the scene state is local. This means that the rendering is based on the actual local version of code and data, which allows for fast iteration and development.

This lets us work from the couch! Such as using Nvidia ray tracing with a Turing GPU on a MacBook

SEED // Halcyon: Rapid Innovation using Modern Graphics

Render Proxy

- The possibilities are **endless!**




With our render proxy architecture, the possibilities are truly endless!



I also wanted to briefly cover some of our machine learning support in Halcyon

Machine Learning

- Deep reinforcement learning
- **Rendering 36 semantic views**
- Training with TensorFlow
 - On-premise GPU cluster
- Inference with TensorFlow
 - CPU AVX2
 - **In-process**



PICA PICA uses deep reinforcement learning, trained with Halcyon rendering 36 semantic views.

The training is performed with TensorFlow and our on-premise GPU cluster.

The PICA PICA demo did in-process inference with TensorFlow using AVX2 on the CPU

Machine Learning

- Adding inferencing with **DirectML**
 - **Hardware accelerated** inferencing operators
 - Resource management
 - Schedule ML work explicitly
 - **Interleave** ML work **with other GPU workloads**
- Fall back for other APIs

We are adding inferencing support with DirectML to Halcyon. This will provide hardware accelerated inferencing operators.

DirectML allows us to do our own resource management, schedule ML work explicitly, and also interleave ML work with other GPU workloads.

We are investigating a fall back for other APIs if DirectML is not available.

Machine Learning

- Treat trained ML models like any other 3D asset
- **Render Graph abstractions**
 - Reference the same render resources
 - Similar to chaining compute passes
- Record “meta” render commands
 - **Backends can fuse or transform**, if desired

We can treat the trained ML models like any other 3D assets.

We will also expose DirectML with render graph abstractions. We want to reference the same render resources, and provide a similar approach to chaining compute passes.

Similar to our other high level commands bracketed within a render pass, we will record some form of “meta” render commands for ML, bracketed together, and this will allow the various backends to fuse or transform these commands for performance, if desired.



I wanted to mention some interesting aspects of our asset pipelines

Asset Pipelines

- Geometry
- Animations
- Shaders
- Sounds
- Music
- Textures
- Scenes
- etc.



100% Rust!



Just like any other engine or framework, we have a variety of asset types that we load from a source or intermediate representation, process with a pipeline, and produce a more optimal representation for runtime.

An exciting spin is that for geometry and animations, we've ditched traditional FBX in favor of glTF.

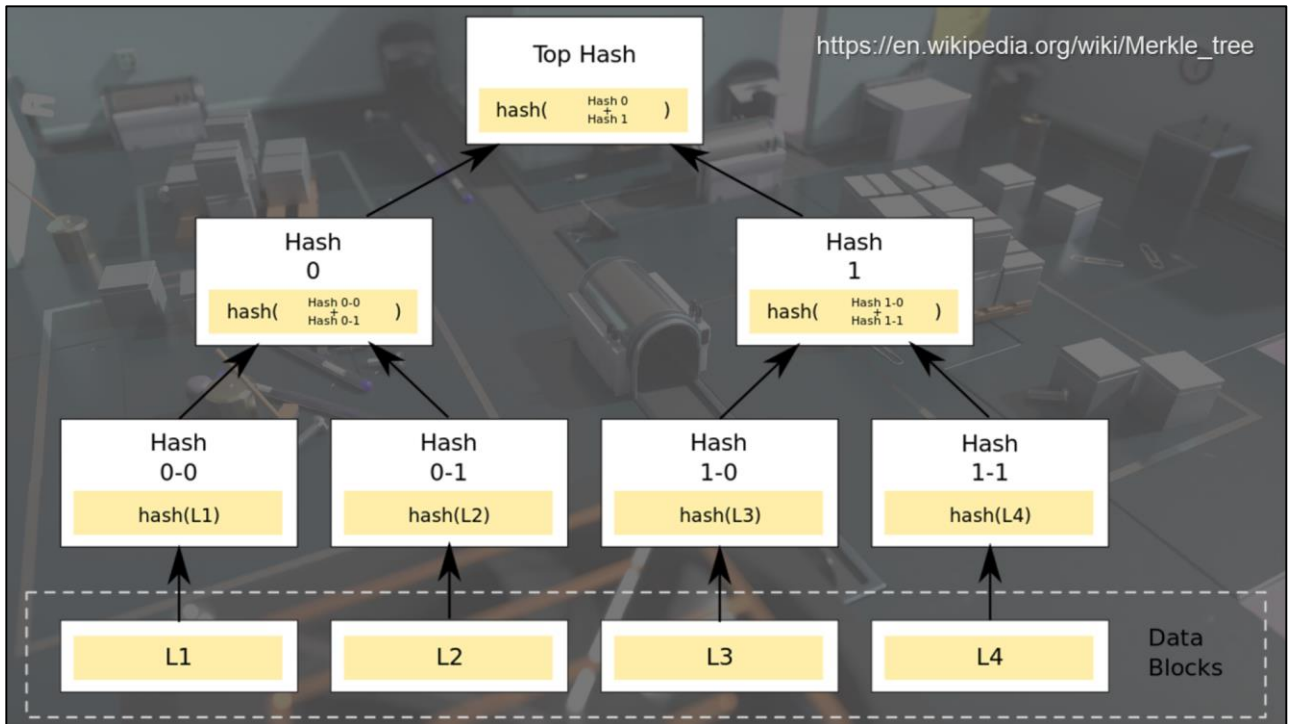
We've also written all our asset pipelines in the Rust programming language. Though, I'll refrain from getting on my usual soapbox and shouting about how awesome Rust is.

Asset Pipelines

- Everything is **content addressable**
 - Hash of data is the identity
 - Sha256
- **Merkle trees!**
 - Dependency evaluation

The cool thing is that all our assets are content addressable - we don't reference resources by virtual file system path, instead we only reference by a hash like sha256.

We use Merkle trees for efficient dependency evaluation



The general idea with Merkle trees is that the leaf nodes are data blocks, which are hashed, and nodes further up the tree are hashes of their respective children.

You can test if a leaf node is part of a given tree by computing a number of hashes proportional to the logarithm of the number of leaf nodes in the tree. This allows for efficient dependency evaluation, avoiding redundant network transmission of data, and secure verification of data contents.

Asset Pipelines

- Containerized, running on **Kubernetes**
 - Google Cloud Platform
 - On-Premise Cluster
 - AMD 1950X TR
 - NV Titan V
- Communication using **gRPC and Protobuf**
- Google Cloud Storage



kubernetes

gRPC

Our pipelines are all containerized in Linux Docker, running on Kubernetes. We can scale our pipelines on Google Cloud Platform, and also in our on-premise cluster, comprised of AMD Threadrippers and Nvidia Titan Vs.

Communication between Halcyon and the pipelines is done with gRPC and Protobuf, and our content addressable data is all backed by Google Cloud Storage.

Asset Pipelines

- **Analytics with Prometheus and Grafana**
 - Publish custom metrics
 - Scraped into rich UI
 - **Collecting data is important!**



Our pipelines support detailed analytics with Prometheus and Grafana. We publish custom metrics to an HTTP endpoint, and these metrics get scraped into a rich UI.

We can't stress enough how important collecting data is – for profiling how long pipelines are taking including bottlenecks to improve, and also for tracking down errors and bugs.



This is what our on-premise GPU cluster metrics look like when displayed with Grafana. This is a very rich and responsive Web UI, and it is very easy for us to add in all sorts of interesting metrics and dashboards for whatever we want to track.



The last thing I'll discuss today, is how we handle shaders

Shaders

- **Complex materials**
 - Multiple microfacet layers
 - [Stachowiak 2018]
- **Energy conserving**
 - Automatic Fresnel between layers
- **All lighting & rendering modes**
 - Raster, path-traced reference, hybrid
- **Iterate with different looks**
 - Bake down permutations for production



Objects with Multi-Layered Materials

We implemented a system which supports multiple microfacet layers arranged into stacks. The stacks could also be probabilistically blended to support prefiltering of material mixtures.

This allowed us to rapidly experiment with the look of PICA PICA, while at the same time enforcing energy conservation, and avoiding common game dev hacks like “metalness”.

An in-engine material editor meant that we could quickly jump in and change the look of everything.

The material system works with all our render modes, be that pure rasterization, path tracing, or hybrid.

And for performance, we can bake down certain permutations for production.

Shaders

- Exclusively HLSL
 - Shader Model 6.X
- Majority are compute shaders
- Performance is critical
 - Group shared memory
 - Wave-ops / Sub-groups

We exclusively use HLSL 6 as a source language, and the majority of our shaders are compute

Performance is critical, so we rely heavily on group shared memory, and wave operations

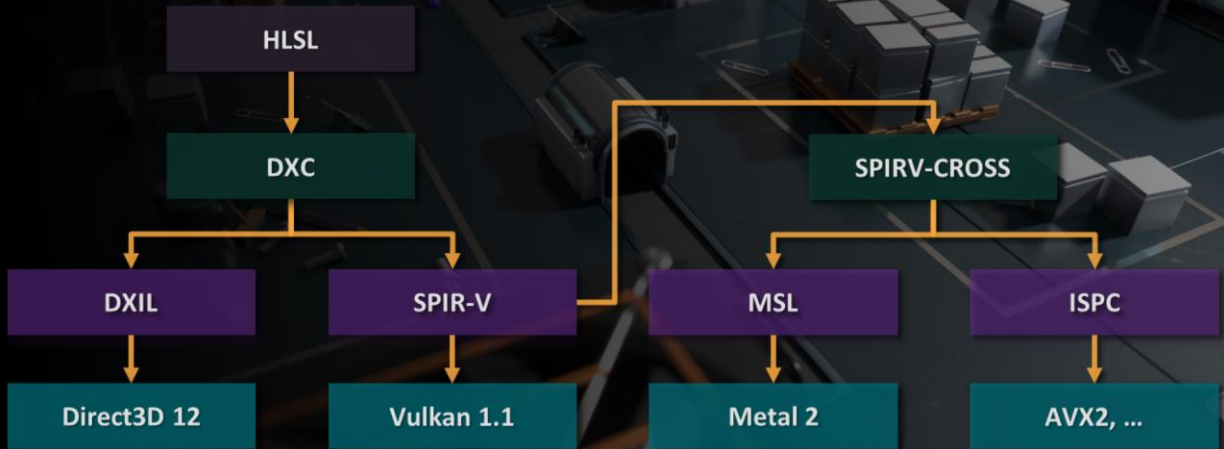
Shaders

- No reflection
 - Avoid costly lookups
 - Only explicit bindings
 - ... except for validation
- Extensive use of HLSL spaces
 - Updates at varying frequency
- Bindless

We don't rely on any reflection information, outside of validation. We avoid costly lookups by requiring explicit bindings.

We also make extensive use of HLSL spaces, where the spaces can be updated at varying frequency.

Shaders



Here is a flow graph of our shader compilation. We always start with HLSL compiled with Microsoft's DXC shader compiler. The DXIL is used by Direct3D 12, and the SPIR-V is used by Vulkan 1.1. We also have support for taking the SPIR-V, running it through SPIRV-CROSS, and generating MSL for Metal, or ISPC to run on the CPU.

Shader Arguments

- Commands refer to resources with “Shader Arguments”
 - Each argument represents an HLSL space
 - MaxShaderParameters → 4 [Configurable]
 - # of spaces, not # of resources

```
struct RenderCommandDispatch : RenderCommandTyped<RenderCommandType::Dispatch, RenderCommandQueueType::Compute>
{
    RenderResourceHandle pipelineState;
    ShaderArgument shaderArguments[MaxShaderParameters]; ←
    uint32 shaderArgumentsCount = 0;

    uint32 dispatchX = 0;
    uint32 dispatchY = 0;
    uint32 dispatchZ = 0;
};
```

We have a concept in Halcyon called shader arguments, where each argument represents an HLSL space. We limit the maximum number of arguments to 4 for efficiency, but this can be configured. It is important to note that this limit represents the maximum number of spaces, not the maximum number of resources.

Shader Arguments

- Each argument contains:
 - “ShaderViews” handle
 - Constant buffer handle and offset
- “ShaderViews”
 - Collection of SRV and UAV handles

Each shader argument contains a “ShaderViews” handle, which refers to a collection of SRV and UAV handles. Additionally, each shader argument also contains a constant buffer handle and offset into the buffer.

Shader Arguments

- Constant buffers are **all dynamic**
 - Avoid temporary descriptors
 - Just a few large buffers, offsets change frequently
 - `VK_DESCRIPTOR_TYPE_UNIFORM_BUFFER_DYNAMIC`
 - *DX12 Root Descriptors* (pass in GPU VA)
- **All descriptor sets are only written once**
 - Persisted / cached

Our constant buffers are all dynamic, and we avoid having temporary descriptors. We have just a few large buffers, and offsets into these buffers change frequently.

For Vulkan, we use the uniform buffer dynamic descriptor type, and on Dx12 we use root descriptors, just passing in a GPU virtual address.

All our descriptor sets are only written once, then persisted or cached.


```

Texture2D<float4> g_albedo           : register(t0, space0);
Texture2D<float4> g_normal           : register(t1, space0);
Texture2D<float4> g_roughness        : register(t2, space0);
Texture2D<float4> g_metalness        : register(t3, space0);
Texture2D<float4> g_ao               : register(t4, space0);
Texture2D<float4> g_emissive         : register(t5, space0);
Texture2DArray<float4> g_translucency : register(t6, space0);

StructuredBuffer<uint3> g_indexBuffer : register(t0, space1);
StructuredBuffer<VtxInputPosition> g_vbPositions : register(t1, space1);
StructuredBuffer<VtxInputTangentSpace> g_vbTangentSpace : register(t2, space1);
StructuredBuffer<VtxInputTexCoordAndColor> g_vbTexAndColor : register(t3, space1);

ConstantBuffer<GeometryConstants> g_geometry : register(b0, space1);
ConstantBuffer<MaterialConstants> g_material : register(b0, space0);

```

```

{
    // space0
    ShaderArgument(materialConstants.buffer, materialConstantsOffset, materialShaderViews),

    // space1
    ShaderArgument(geometryConstants.buffer, geometryConstantsOffset, geometryShaderViews)
},

```

Here is an example HLSL snippet showing one usage of spaces. Each space contains a collection of SRVs and a constant buffer. The shader argument configuration on the CPU is shown below.

SPIR-V Patching

- Shader compilation (**HLSL** → **SPIR-V**)
 - Patch SPIR-V to match DX12
 - Using **spirv-reflect** from Hai and Cort
- `spvReflectCreateShaderModule`
- `spvReflectEnumerateDescriptorSets`
- `spvReflectChangeDescriptorBindingNumbers`
- `spvReflectGetCodeSize` / `spvReflectGetCode`
- `spvReflectDestroyShaderModule`

An important part of our Vulkan backend, was consuming HLSL as a source language, and fixing up the SPIR-V to behave the same as our Dx12 resource binding model.

We decided to use the spirv-reflect library from Hai Nguyen and Cort Stratton; it does a great job at providing SPIR-V reflection using DX12 terminology, but we use it exclusively to patch up our descriptor binding numbers.

SPIR-V Patching

- **SPV_REFLECT_RESOURCE_FLAG_SRV**
 - Offset += 1000
- **SPV_REFLECT_RESOURCE_FLAG_SAMPLER**
 - Offset += 2000
- **SPV_REFLECT_RESOURCE_FLAG_UAV**
 - Offset += 3000

SRVs, Samplers, and UAVs are simple. These types are uniquely name spaced in DX12, so t0 and s0 wouldn't collide. This is not the case in SPIR-V, so we apply a simple offset to each type to emulate this behavior.

SPIR-V Patching

- **SPV_REFLECT_RESOURCE_FLAG_CBV**
 - Offset Unchanged: 0
 - Descriptor Set += **MAX_SHADER_ARGUMENTS**
- CBVs move to their own descriptor sets
 - ShaderViews become persistent and immutable

Constant or uniform buffers are a bit more interesting. We want to move CBVs to their own descriptor sets, in order to make our ShaderViews representing the other resource types persistent and immutable.

To do so, we don't adjust the offset, as we'll have a single CBV per descriptor set. However, we do shift the descriptor set number by the max number of shader arguments.

This means if descriptor set 0 contained a constant buffer, that constant buffer would move to descriptor set 5 (if max shader arguments is 4).

SPIR-V Patching

- If 2 of 4 HLSL spaces in use:

Set 0	SRVs (≥ 1000)	Samplers (≥ 2000)	UAVs (≥ 3000)
Set 1	SRVs (≥ 1000)	Samplers (≥ 2000)	UAVs (≥ 3000)
Set 2	Unbound		
Set 3	Unbound		
Set 4	Dynamic Constant Buffer (Offset: 0)		
Set 5	Dynamic Constant Buffer (Offset: 0)		


If a dispatch or draw is using 2 HLSL spaces, the patched SPIR-V will require the following descriptor set layout. Notice the shifted offsets for the SRVs, Samplers, and UAVs, and how the dynamic constant buffers have been hoisted out to their own descriptor sets.

References

- [Wihlidal 2018] Graham Wihlidal, Colin Barré-Brisebois. "Modern Graphics Abstractions & Real-Time Ray Tracing".
[available online](#)
- [Wihlidal 2018] Graham Wihlidal. "Halcyon + Vulkan".
[available online](#)
- [Stachowiak 2018] Tomasz Stachowiak. "Towards Effortless Photorealism Through Real-Time Raytracing".
[available online](#)
- [Andersson 2018] Johan Andersson, Colin Barré-Brisebois. "DirectX: Evolving Microsoft's Graphics Platform".
[available online](#)
- [Harmer 2018] Jack Harmer, Linus Gisslén, Henrik Holst, Joakim Bergdahl, Tom Olsson, Kristoffer Sjöö and Magnus Nordin. "Imitation Learning with Concurrent Actions in 3D Games".
[available online](#)
- [Opara 2018] Anastasia Opara. "Creativity of Rules and Patterns".
[available online](#)
- [O'Donnell 2017] Yuriy O'Donnell. "Frame Graph: Extensible Rendering Architecture in Frostbite".
[available online](#)
- [Mui 2012] Chunka Mui, Forbes. "How Kodak Failed".
[available online](#)



On one last note, I would like to point out that we're hiring for multiple positions at SEED. If you're interested, please give us a shout!




SEED // Halcyon: Rapid Innovation using Modern Graphics

Thank You

- Reboot Develop Blue & Khronos
- Halcyon Team @ SEED
 - Andrew Lauritzen
 - Colin Barré-Brisebois
 - Jasper Bekkers
 - Henrik Halén
 - Carlos Gonzalez-Ochoa
 - Graham Wihlidal

Thank you to the Reboot Develop Blue event organizers and Khronos for inviting me to speak, and also to the awesome team at SEED working on Halcyon.




Graham Wihlidal

graham@ea.com

[@gwhlidal](#)

Questions?

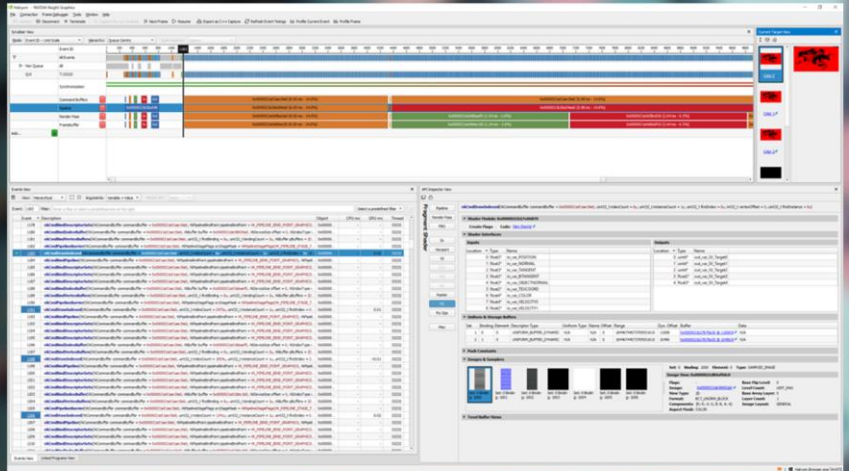
And with that I'd like to open it up for any questions you may have.



I will briefly mention some useful tools used for the Vulkan implementation

Tools

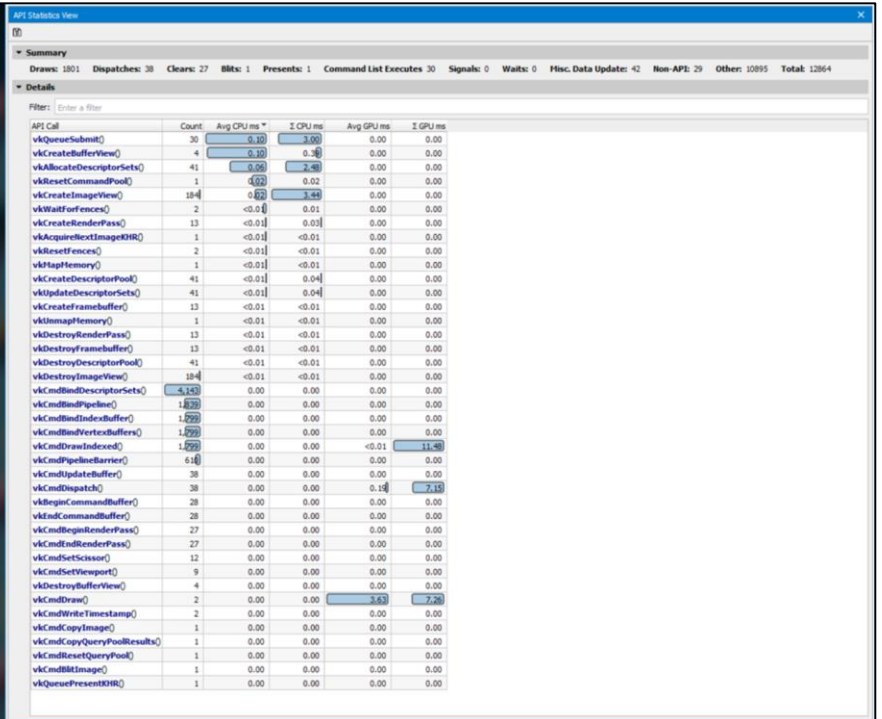
- RenderDoc
- NV Nsight
- AMD RGP



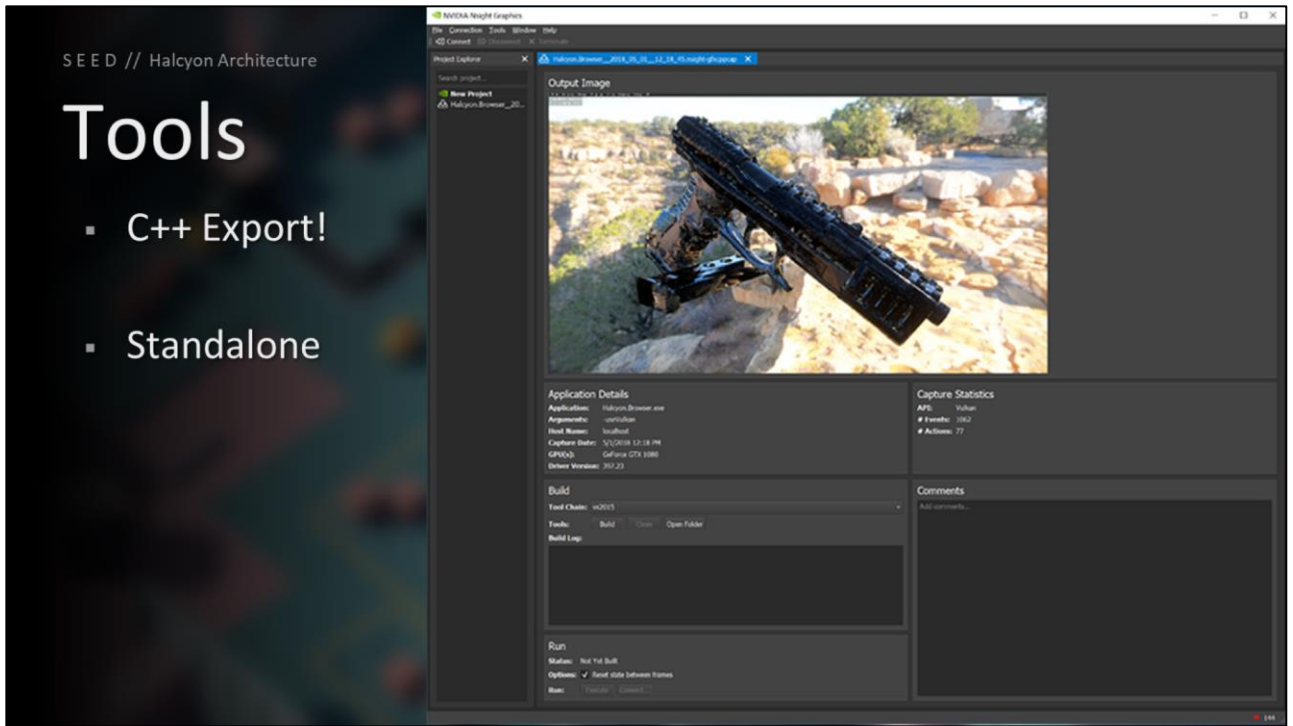
For debugging and profiling, the usual suspects were quite helpful, and used extensively.

For debugging and profiling, the usual suspects were quite helpful, and used extensively.

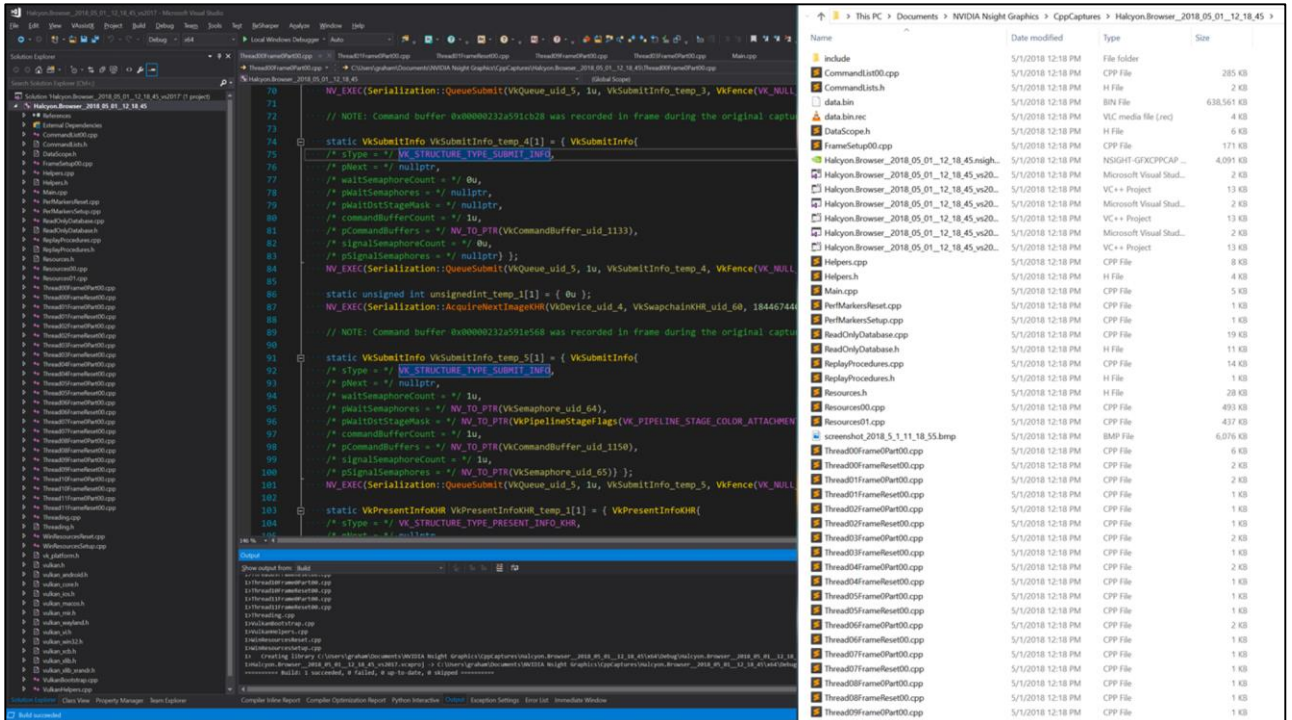
Tools



The API statistics view in Nvidia Nsight is a great way to look at the count and cost of each API call.



Another awesome feature with Nsight is the ability to export a capture as a standalone C++ application. Nsight will write out binary blobs of your resources in the capture, and write out source code issues your API calls. You can build this app and debug problems in an isolated solution.




Another awesome feature with Nsight is the ability to export a capture as a standalone C++ application. Nsight will write out binary blobs of your resources in the capture, and write out source code issues your API calls. You can build this app and debug problems in an isolated solution.



Here is our bug being reproduced in the standalone C++ export

Dear ImGui + ImGuizmo

- Live tweaking
- Very useful!



With our goal of having everything be live-reloadable and live-tweakable, we used DearImGui and ImGuizmo extensively for a number of useful overlays

