



# Making SPIR-V Modules

John Kessenich

January 2017



# Talk Overview

## 1. Making a SPIR-V Module

What's in a module

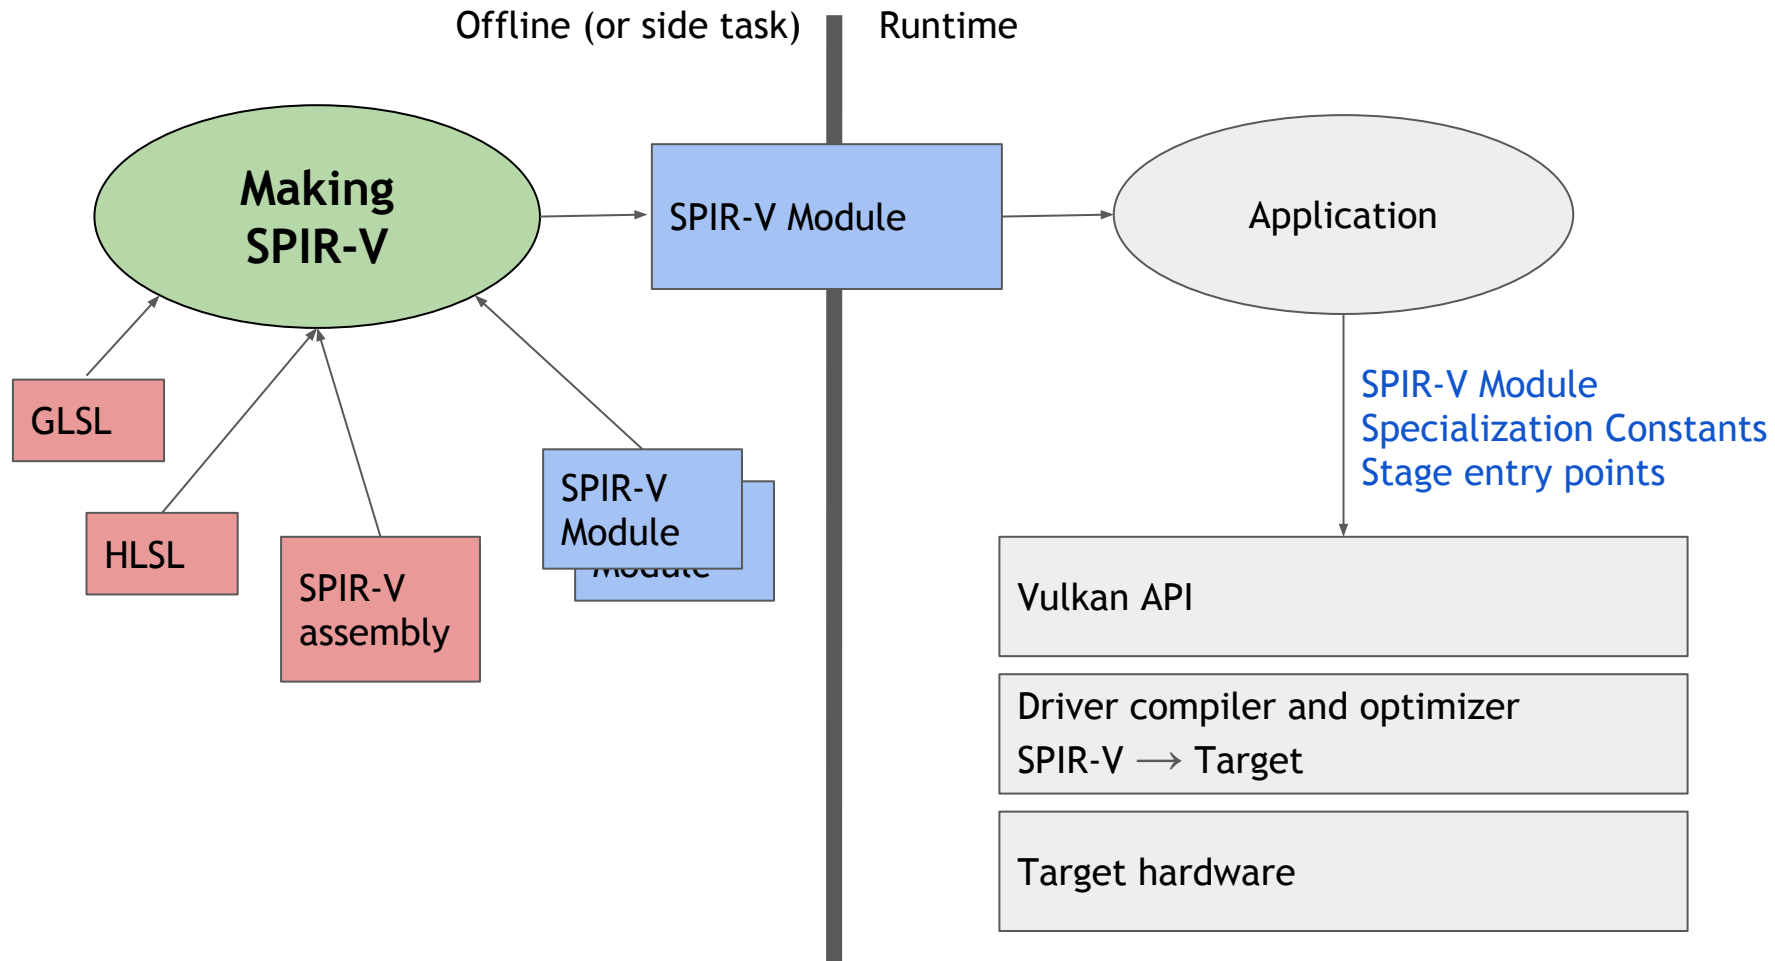
Tools to make modules

## 2. Managing Size of a Large Collection of Modules

SPIR-V features: Specialization Constants

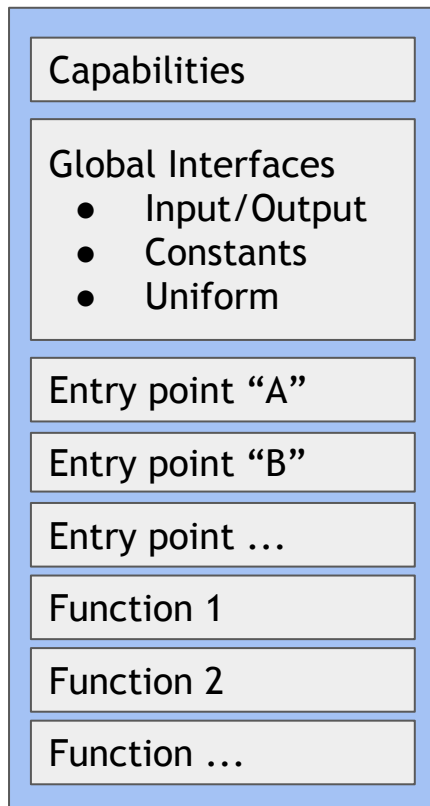
Normalization and compression

# 1. Making a SPIR-V Module



# A SPIR-V Module

Application specifies  
{stage, entry point}  
pairs to subset  
module



Target must support all  
declared capabilities, shared  
by all entry points

Specialization constants

Multiple entry points share  
interfaces and functions

one .spv file == one SPIR-V module

# Projects on GitHub

## Each contains multiple tools

Will discuss today:

- **Glslang:** <https://github.com/KhronosGroup/glslang>
- **SPIR-V Tools:** <https://github.com/KhronosGroup/SPIRV-Tools>
- **Shaderc:** <https://github.com/google/shaderc>

Also see

- **SPIR-V Cross:** <https://github.com/KhronosGroup/SPIRV-Cross>
  - SPIR-V reflection and translation to higher-level languages
- **SMOL-V:** <https://github.com/aras-p/smol-v>
  - Compression

# Glslang

Contains multiple tools:

1. **Khronos reference GLSL validator**
  - Base of original command-line options
2. **GLSL/ESSL → SPIR-V compiler**
3. **HLSL → SPIR-V compiler**
  - Newly developed by LunarG, Google, Valve and other ISVs
4. **Reflection of HLL**
5. **Remapper, a compression normalizer and size optimizer**
  - Independent tool
  - moving to tools

# Glslang: GLSL/ESSL → SPIR-V

- `glslangValidator -V -o module.spv shader.frag`
- Follows `GL_KHR_vulkan_glsl`:
  - No loose uniforms; need to use blocks
    - `uniform blockName { <uniform members> };`
  - All uniform/buffer blocks, samplers, etc. needing bindings and sets
    - `layout(binding = 0, set = 1) <resource declaration>`
  - All in/out variables need locations
    - `layout(location = 5) <variable declaration>`
  - Based on `#version`
  - Precision qualifiers work with desktop shaders
- **Multiple files → single stage → single module**



# Glslang: HLSL → SPIR-V

- **Active project**

- All SteamVR HLSL shaders are working without modification
- Additional features in progress for other shaders
- Many additional ISVs feeding input into the project

- **layout(binding = 5, set = 2) works in Glslang HLSL**

- **Many command-line options to manage mapping of I/O to Vulkan**

```
-D          input is HLSL
-e          specify SPIR-V entry-point name
--source-entypoint name          specify HLSL entry point name
-S <stage> uses explicit stage specified
--shift-sampler-binding [stage] num          set base binding number for samplers
--shift-texture-binding [stage] num          set base binding number for textures
--shift-image-binding [stage] num           set base binding number for images
--shift-UBO-binding [stage] num             set base binding number for UBOs
--auto-map-bindings                  automatically bind uniform variables
--flatten-uniform-arrays              flatten uniform texture & sampler arrays
--no-storage-format                  use Unknown image format
```

# Glslang as a Library

- **Glslang is mostly libraries**
  - Can be linked into by other tools
  - glslangValidator is just an example provided wrapper, easy to add others
  - See `glslang/StandAlone/StandAlone.cpp`

```
glslang::InitializeProcess();
```

```
glslang::TShader shader;  
shader.setStrings(...file content...);  
shader.parse(...);
```

```
glslang::TProgram program;  
program.addShader(shader);  
program.link();
```

```
glslang::FinalizeProcess();
```

# SPIR-V Tools

## spirv-dis: Disassembler

- Binary SPIR-V file → human readable assembly form
- Based on a reusable SPIR-V binary-parser API (use it if you make tool!)

## spirv-as: Assembler

- Human-readable assembly form → binary form

⇒ round-trip editing works


## spirv-val: Validator

- Individual module
- Within-module validation (there is different SPIR-V validation done in Vulkan API validation layers)

## spirv-opt: Optimizer, work in progress

# Shaderc

- **Wraps glslang and SPIR-V Tools**
  - Command-line options based on other command-line compilers
  - glslc ... (see next slide)
- **Uses a copy of glslang**
  - Can be missing latest features
- **Ahead of glslang in #include functionality**



# Shaderc (continued)

# Preprocess

```
glslc -E shader.glsl
```

# Compile

```
glslc -c shader.vert
```

# Disassemble

```
glslc -S shader.vert
```

# Optimize

```
glslc -c -Os shader.vert
```

# Specify shader stage

```
glslc -c -fshader-stage=vertex shader.glsl
```

# Specify language version and profile

```
glslc -c -std=310es shader.vert
```

# Specify target environment

```
glslc -c --target-env=vulkan shader.vert
```

# Specify output file name

```
glslc -c shader.vert -o shader.spv
```

# Specify output format

# E.g., output SPIR-V binary code as

# a C-style initializer list

```
glslc -c -mfmt=c shader.vert
```

# Define

```
glslc -E -DVALUE=42 shader.glsl
```

# Include

```
glslc -E -I../include shader.glsl
```

# Generate dependencies for builds

```
glslc -MD -c shader.glsl
```

# Warnings and errors

```
glslc -c -Werror shader.vert
```

# Shaderc as a Library

```
1 #include <shaderc/shaderc.hpp>
2
3 shaderc::Compiler compiler;
4 shaderc::CompileOptions options;
5
6 const std::string source =
7     "#version 310 es\nvoid main() { int v = VALUE; }\n";
8
9 options.AddMacroDefinition("VALUE", "42");
10 options.SetOptimizationLevel(shaderc_optimization_level_size);
11
12 auto pp = compiler.PreprocessGls1(source.c_str(), source.size(),
13     source, shaderc_gls1_vertex_shader, "shader.gls1", options);
14 std::cout << "Preprocessed shader:\n"
15     << std::string(pp.cbegin(), pp.cend()) << std::endl;
16 // Preprocessed shader:
17 // #version 310 es
18 // void main(){ int v = 42; }
19
20 auto compiling = compiler.CompileGls1ToSpvAssembly(
21     source, shaderc_gls1_vertex_shader, "shader.gls1", options);
22 std::string assembly(compiling.cbegin(), compiling.cend());
```

```
23 std::cout << "Compiled SPIR-V assembly:\n" << assembly << std::endl;
24 // Compiled SPIR-V assembly:
25 // ...
26 //             OpCapability Shader
27 //     %1 = OpExtInstImport "GLSL.std.450"
28 //             OpMemoryModel Logical GLSL450
29 //             OpEntryPoint Vertex %4 "main"
30 // %void = OpTypeVoid
31 // ...
32
33 auto assembling = compiler.AssembleToSpv(assembly.c_str(),
34     assembly.size());
35 std::cout << "Final binary # words: " << spirv.size() << std::endl;
36 // Final binary # words: 54
37
38 const std::string bad_source =
39     "#version 310 es\n void main() { the_ultimate_shader }";
40 auto error = compiler.CompileGls1ToSpv(
41     source, shaderc_gls1_vertex_shader, "shader.gls1", options);
42 std::cerr << error.GetErrorMessage() << std::endl;
43 // shader.gls1:2: error: 'the_ultimate_shader' : undeclared identifier
44 // shader.gls1:2: error: '' : syntax error
```


## 2. Managing Size of a Large Collection of Modules

# Size

- **SPIR-V is verbose, designed**
  - For ease of processing by SPIR-V tools and drivers
  - To be explicit, not inference based, e.g., type redundancy
  - a portable non-lossy standard
- **Individual shaders are small enough**
  - Glslang output includes **OpName**, redundant load/store
  - `spirv-remap --dce all --opt all --strip all -o out_dir -i module.spv`
  - `spirv-opt ...`
  - `glslc -Os ...`



# Issue: Making a large number of variations of similar shaders



Offline

# Managing Large Collections

Two completely different approaches:

## 1. Recommended: Less SPIR-V, using SPIR-V features

- SPIR-V was designed to represent variations efficiently
- Requires using the features provided

## 2. Legacy: Compression

- Works best on collections of related modules
- Involves normalization followed by compression

# Less SPIR-V, Using SPIR-V Features

- **Specialization constants**
  - Fewer SPIR-V modules
  - Defer some options/sizes to runtime
- **Lots of entry points in a single SPIR-V module**
  - Share functions
  - Share uniform interface
  - Same capabilities
  - Tools still immature

**Will focus on Specialization Constants here...**

# Specialization Constants

1. Declare specialization constants in GLSL or HLSL:

```
layout(constant_id = 13) const int size = 9;
```

- Provide default values (9 above)
  - Generates SPIR-V with specialization constant ids (13 above)
2. Distribute SPIR-V with specialization constants
  3. Specialize at runtime
    - Set constants when creating pipeline
    - Those not set use their default
  4. Driver compiler will optimize knowing the specialized constant value

# Vulkan 1.0.31

## 9.7. Specialization Constants

`VkPipelineShaderStageCreateInfo.pSpecializationInfo`

```
typedef struct VkSpecializationInfo {
    uint32_t                mapEntryCount;
    const VkSpecializationMapEntry* pMapEntries;
    size_t                  dataSize;
    const void*             pData;
} VkSpecializationInfo;
```

## Change size 9 → 8

```
struct SpecializationData {
    int32_t data0;
    ...
};

const VkSpecializationMapEntry entries[] =
{
    {
        13, // constant_id
        offsetof(SpecializationData, data0), // offset
        sizeof(SpecializationData::data0) // size
    },
    ...
};

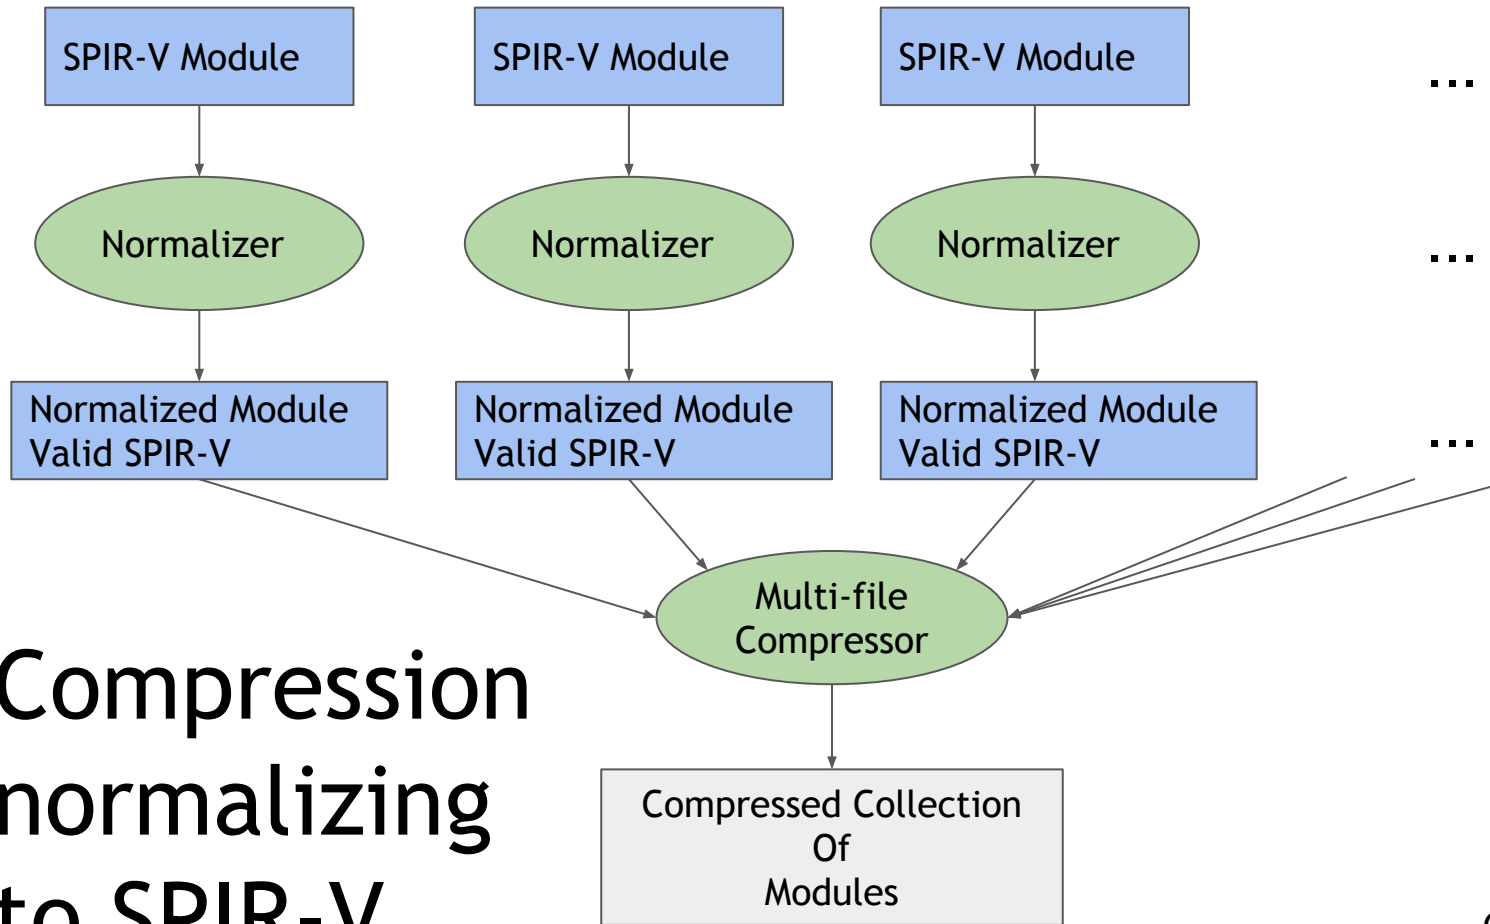
SpecializationData data;
data.data0 = 8; // new value for the constant
...

const VkSpecializationInfo info =
{
    <number of constants being set>, // mapEntryCount
    entries, // pMapEntries
    sizeof(data), // dataSize
    &data, // pData
};
```

# Compression

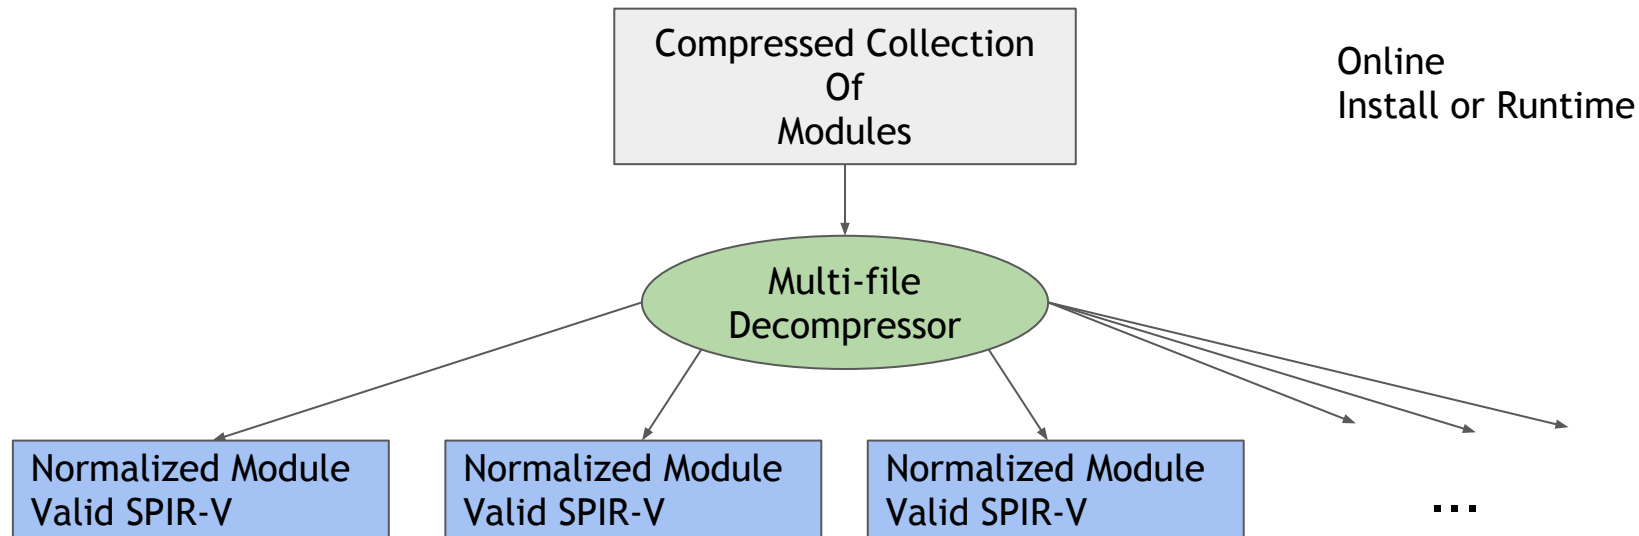
- **Better across multiple modules**
- **Better if normalized first**
  - Can normalize to SPIR-V or something else
  - If something else, need to denormalize before decompression

# Compression normalizing to SPIR-V



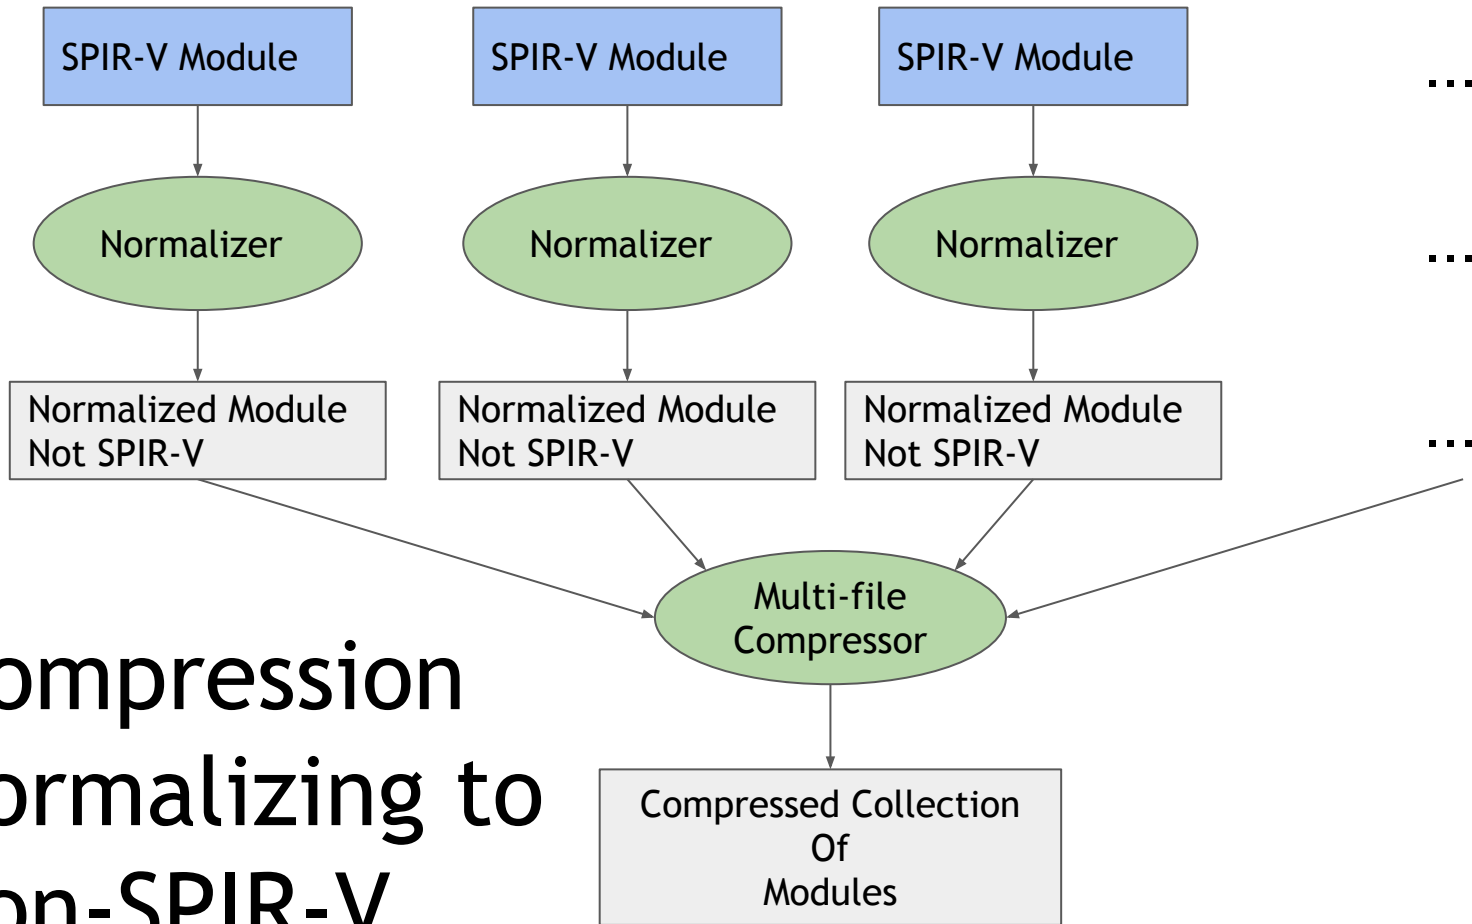
Offline





# Decompressing SPIR-V-normalized SPIR-V

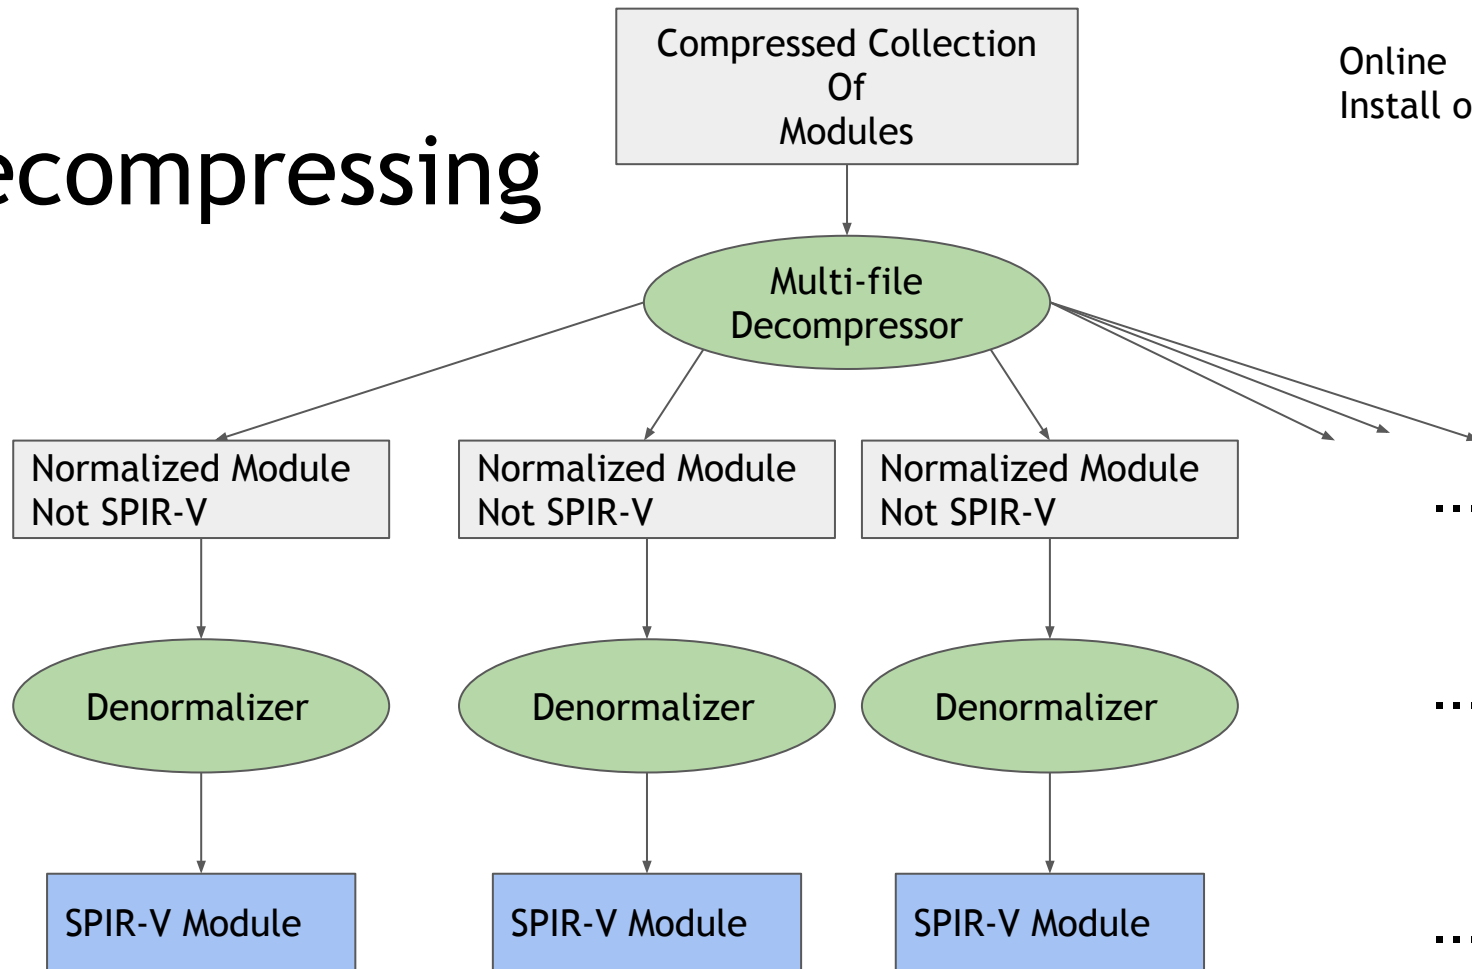
# Compression normalizing to non-SPIR-V



Offline

# Decompressing

Online  
Install or Runtime



# Normalize and Compress with Glslang's Remapper

## 1. Normalize step (offline):

```
spirv-remap --map all -o out_dir -i mod1.spv mod2.spv ...
```

## 2. Compress

```
tar -cf - out_dir | lzma -z > compressed.lzma
```

## 3. Distribute

## 4. Decompress

```
lzma -d < compressed.lzma | tar -xvf -
```

## 5. No denormalization needed

Also see SMOL-V: <https://github.com/aras-p/smol-v>

- Smaller than SPIR-V normalizer
- Needs denormalizer

# Future work

- **More optimizations**
- **Simplify compression, put into SPIRV Tools**
- **Merge multiple modules into a single module**
  - Same capabilities
  - Share same utility functions

End