

# *OPENCL ECOSYSTEM UPDATES*

Developer UX at the forefront

# TABLE OF CONTENTS

## GFX & INTERCONNECT

A tale of HW and SW

11/9/2021

GPU-Day 2021 – Budapest, Hungary

3

## COMPUTE APIS

Status quo

11/9/2021

GPU-Day 2021 – Budapest, Hungary

8

## WHAT HAS CHANGED?

Why do standards matter?

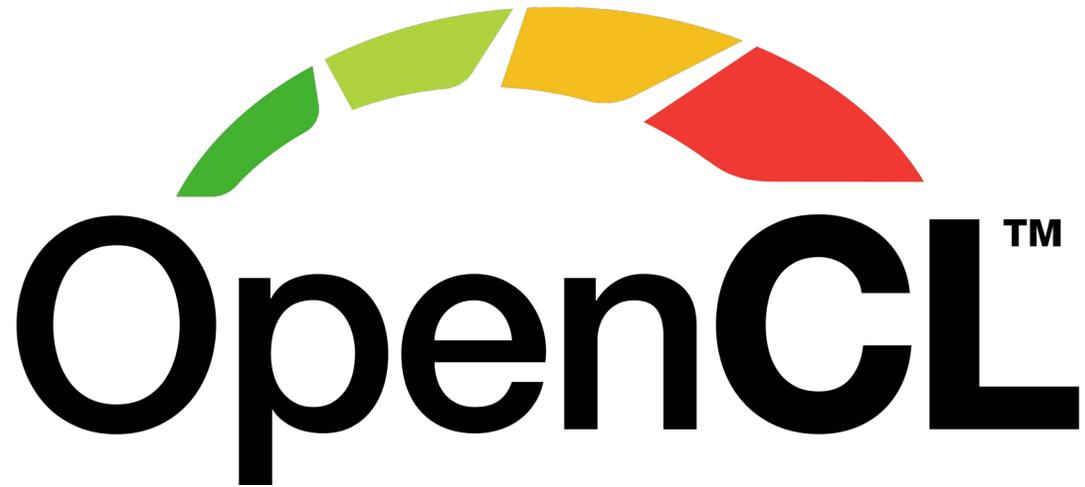
11/9/2021

GPU-Day 2021 – Budapest, Hungary

14

# 1ST PARTY OPENCL-SDK

Where it all begins



- Re-imagination of the API
- Builds on top of the last widely adopted version, 1.2
- Adopts the Vulkan-like „slim core API with everything else extensions” approach
- Lowers the bar for adoption
  - Newcomers
  - Solves issues for existing adopters

**1ST PARTY  
OPENCL-SDK**

# OPENCL 3.0 ECOSYSTEM REBOOT

- New ICD Layers capability allows augmenting/modifying runtime behavior
  - Add missing features, validate correctness, trace, profile, etc
- OpenCL SDK serves as a „one-stop shop” for devs
  - Initial batch of native samples & utilities awarded to Stream
    - Coming to a [GitHub](#) near you
- Defines the semantics of layered implementations to fill in the gaps of platform support
  - Most notably [OpenCLOn12](#), enabling OpenCL across the entire Windows ecosystem, regardless of vendor support.

OpenCL™ and OpenGL® Compatibility Pack

Microsoft Corporation • Utilities & tools

Free

Get

See System Requirements

This compatibility pack allows more of your favorite OpenCL™ and OpenGL® apps to run on a Windows 10 PC that doesn't have OpenCL and OpenGL hardware drivers installed by default. If a DirectX 12 driver is installed, supported apps will run with hardware acceleration for better performance. This package supports apps that use OpenCL version 1.2 and earlier and OpenGL version 3.3 and earlier.

More

EVERYONE

<https://www.microsoft.com/en-us/p/opencl-and-opengl-compatibility-pack/9nqpsl29bfff>

# OPENCL 3.0 ECOSYSTEM REBOOT

- New ICD Layers capability allows augmenting/modifying runtime behavior
  - Add missing features, validate correctness, trace, profile, etc
- OpenCL SDK serves as a „one-stop shop” for devs
  - Initial batch of native samples & utilities awarded to Stream
    - Coming to a [GitHub](#) near you
- Defines the semantics of layered implementations to fill in the gaps of platform support
  - Most notably [OpenCLOn12](#), enabling OpenCL across the entire Windows ecosystem, regardless of vendor support.

OpenCL™ and OpenGL® Compatibility Pack

Microsoft Corporation • Utilities & tools

Free

Get

See System Requirements

This compatibility pack allows more of your favorite OpenCL™ and OpenGL® apps to run on a Windows 10 PC that doesn't have OpenCL and OpenGL hardware drivers installed by default. If a DirectX 12 driver is installed, supported apps will run with hardware acceleration for better performance. This package supports apps that use OpenCL version 1.2 and earlier and OpenGL version 3.3 and earlier.

More

EVERYONE

<https://www.microsoft.com/en-us/p/opencl-and-opengl-compatibility-pack/9nqpsl29bfff>

# WHAT'S INSIDE THE BOX?

- OpenCL-Headers
- OpenCL-ICD-Loader
- OpenCL-CLHPP
- Utility libraries
- Sample codes
- Documentation
  - OpenCL-Guide

# WHAT'S INSIDE THE BOX?

- **OpenCL-Headers**
- OpenCL-ICD-Loader
- OpenCL-CLHPP
- Utility libraries
- Sample codes
- Documentation
  - OpenCL-Guide
- The definitive C bindings to the OpenCL API
- Relatively new helper headers for
  - Creating your own ICD
  - Creating your own layer

# WHAT'S INSIDE THE BOX?

- OpenCL-Headers
- **OpenCL-ICD-Loader**
- OpenCL-CLHPP
- Utility libraries
- Sample codes
- Documentation
  - OpenCL-Guide

- The canonical Installable Client Driver loader
  - libOpenCL.so/OpenCL.dll
- Responsible for loading vendor runtimes

```
clGetPlatformIDs(numPlatforms, platforms,  
NULL);
```

- Exposed to
  - Let users tap into latest features
  - Help implementers debug

# WHAT'S INSIDE THE BOX?

- OpenCL-Headers
- OpenCL-ICD-Loader
- **OpenCL-CLHPP**
- Utility libraries
- Sample codes
- Documentation
  - OpenCL-Guide
- Canonical C++ bindings to the OpenCL API
  - Reduces verbosity
  - Adds safety guards
- C++11 is the minimum

# WHAT'S INSIDE THE BOX?

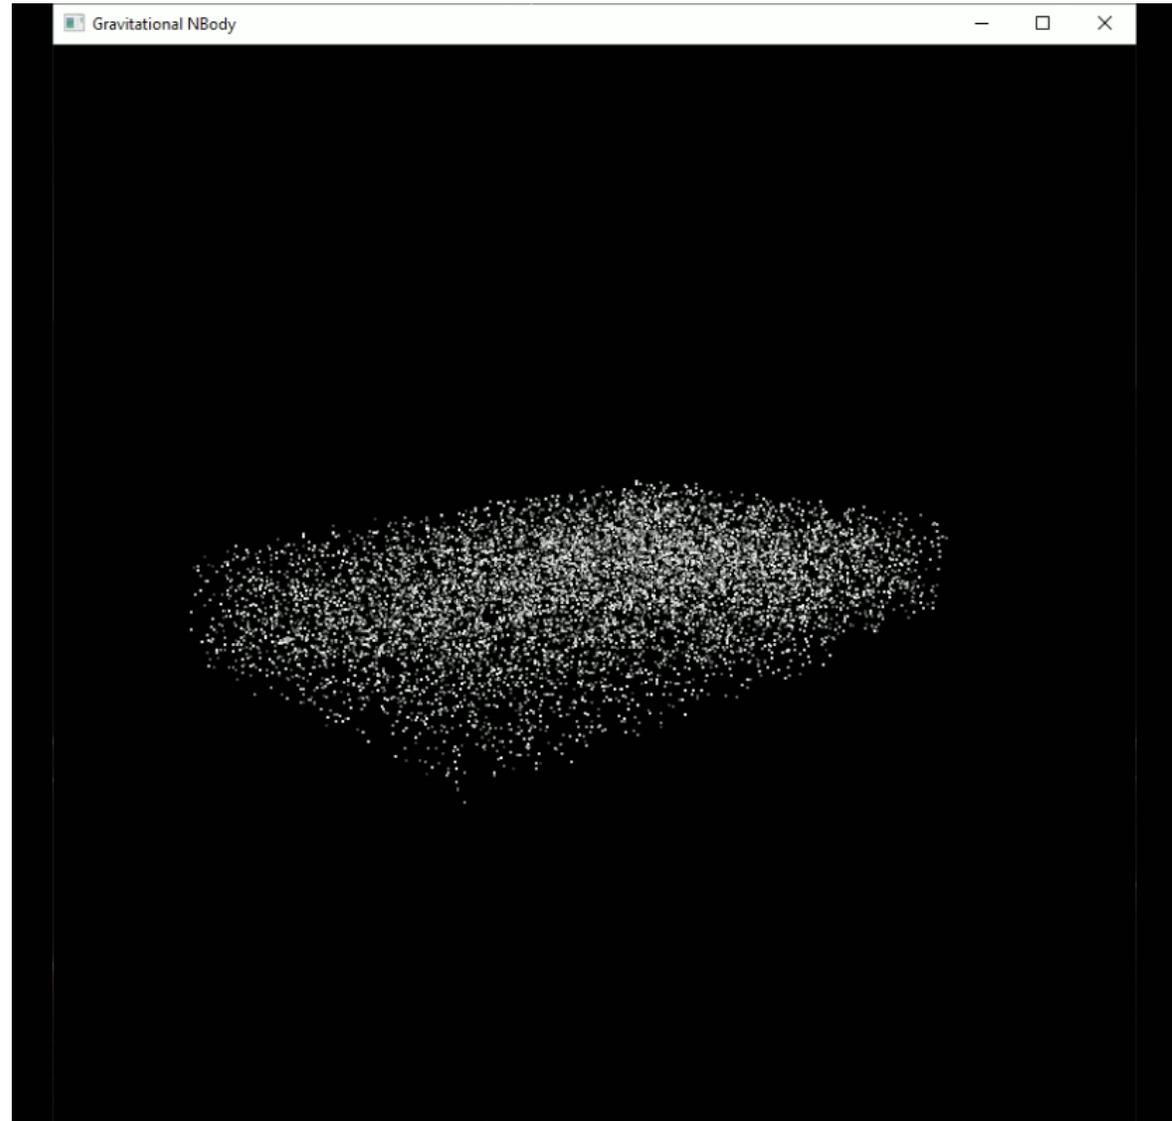
## 1ST PARTY OPENCL-SDK

- OpenCL-Headers
- OpenCL-ICD-Loader
- OpenCL-CLHPP
- **Utility libraries**
- Sample codes
- Documentation
  - OpenCL-Guide

- Utility libraries to help developers with common tasks, for eg.
  - Give me a context with device N on platform M
  - Obtain profiling data in std::chrono-friendly format
- Contributions are welcome
  - Convert error code -48 to „CL\_INVALID\_KERNEL”?

# WHAT'S INSIDE THE BOX?

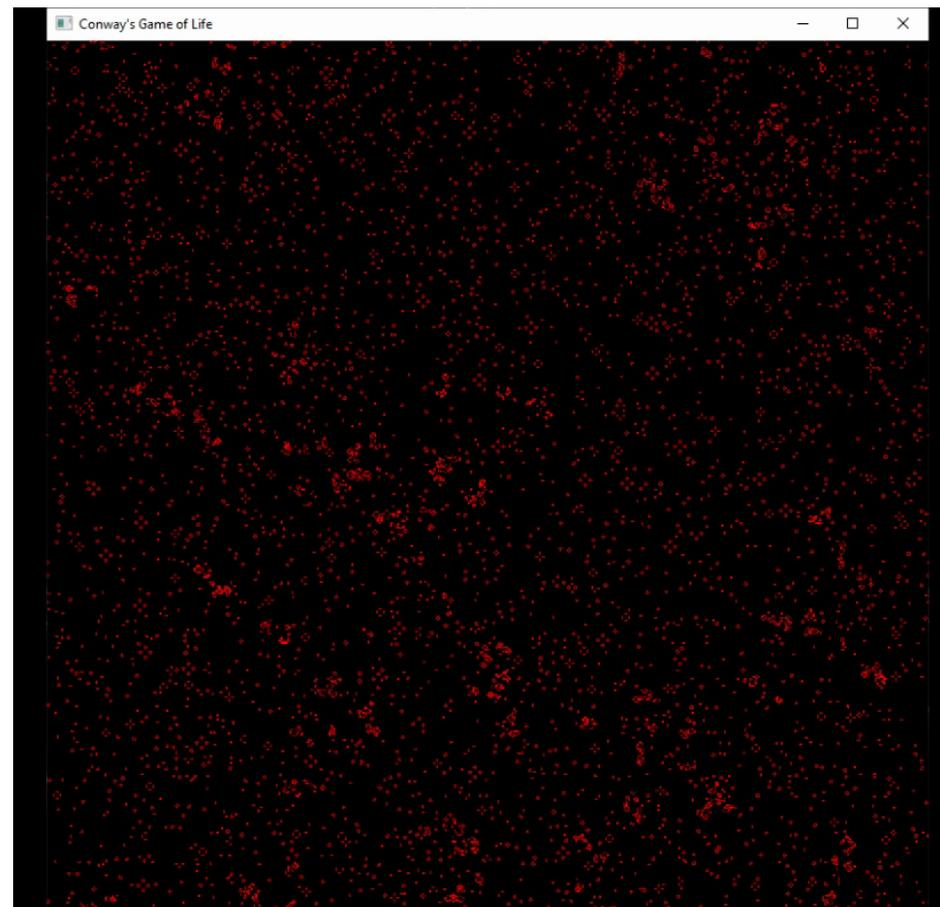
- OpenCL-Headers
- OpenCL-ICD-Loader
- OpenCL-CLHPP
- Utility libraries
- **Sample codes**
- Documentation
  - OpenCL-Guide



**1ST PARTY  
OPENCL-SDK**

# WHAT'S INSIDE THE BOX?

- OpenCL-Headers
- OpenCL-ICD-Loader
- OpenCL-CLHPP
- Utility libraries
- **Sample codes**
- Documentation
  - OpenCL-Guide



**1ST PARTY**  
**OPENCL-SDK**

# WHAT'S INSIDE THE BOX?

- OpenCL-Headers
  - OpenCL-ICD-Loader
  - OpenCL-CLHPP
  - Utility libraries
  - Sample codes
  - **Documentation**
    - OpenCL-Guide
- All documentation is written in Markdown
  - Cross-linking with the OpenCL-Guide
    - Getting started guides
      - Linux
      - Windows
    - CMake support
      - More on this later

# HIGHER STANDARDS

**1ST PARTY  
OPENCL-SDK**

- The entire ecosystem is thoroughly tested
  - Warning free using highest warn levels
    - 3 platforms (Windows, Linux, MacOS)
    - 3 compilers (MSVC, GCC, Clang)
    - Multiple language C/C++ standards
    - With/without compiler extensions
  - Public facing CI scripts



# HIGHER STANDARDS

- The entire ecosystem is thoroughly tested
  - Warning free using highest warn levels
    - 3 platforms (Windows, Linux, MacOS)
    - 3 compilers (MSVC, GCC, Clang)
    - Multiple language C/C++ standards
    - With/without compiler extensions
  - Public facing CI scripts
- Continuous Deployment in place
  - Quarterly source & binary releases



v2022.05.18 Latest

Synchronize with OpenCL v3.0.11 specification release.

This release of the OpenCL SDK includes Windows binaries and complete source code (including submodules).

**What's Changed**

- Fix utility library install by [@MathiasMagnus](#) in #46

Full Changelog: [v2022.04.01...v2022.05.18](#)

# LAYERS ARE KEY

- Ogres are like onions.
- They smell and turn white if you leave them on the sun?
- No Donkey! They got layers!

# OPENCL 3.0 ECOSYSTEM REBOOT

- New ICD Layers capability allows augmenting/modifying runtime behavior
  - Add missing features, validate correctness, trace, profile, etc
- OpenCL SDK serves as a „one-stop shop” for devs
  - Initial batch of native samples & utilities awarded to Stream
    - Coming to a [GitHub](#) near you
- Defines the semantics of layered implementations to fill in the gaps of platform support
  - Most notably [OpenCLOn12](#), enabling OpenCL across the entire Windows ecosystem, regardless of vendor support.



OpenCL™ and OpenGL® Compatibility Pack

Microsoft Corporation • Utilities & tools

This compatibility pack allows more of your favorite OpenCL™ and OpenGL® apps to run on a Windows 10 PC that doesn't have OpenCL and OpenGL hardware drivers installed by default. If a DirectX 12 driver is installed, supported apps will run with hardware acceleration for better performance. This package supports apps that use OpenCL version 1.2 and earlier and OpenGL version 3.3 and earlier.

More

Free

Get

△ See System Requirements

EVERYONE

ESRB

<https://www.microsoft.com/en-us/p/opencl-and-opengl-compatibility-pack/9nqpsl29bfff>

# OPENCL 3.0 ECOSYSTEM REBOOT

- New ICD Layers capability allows augmenting/modifying runtime behavior
  - Add missing features, validate correctness, trace, profile, etc
- OpenCL SDK serves as a „one-stop shop” for devs
  - Initial batch of native samples & utilities awarded to Stream
    - Coming to a [GitHub](#) near you
- Defines the semantics of layered implementations to fill in the gaps of platform support
  - Most notably [OpenCLOn12](#), enabling OpenCL across the entire Windows ecosystem, regardless of vendor support.



OpenCL™ and OpenGL® Compatibility Pack

Microsoft Corporation • Utilities & tools

Free

Get

See System Requirements

This compatibility pack allows more of your favorite OpenCL™ and OpenGL® apps to run on a Windows 10 PC that doesn't have OpenCL and OpenGL hardware drivers installed by default. If a DirectX 12 driver is installed, supported apps will run with hardware acceleration for better performance. This package supports apps that use OpenCL version 1.2 and earlier and OpenGL version 3.3 and earlier.

More

EVERYONE

<https://www.microsoft.com/en-us/p/opencl-and-opengl-compatibility-pack/9nqpsl29bfff>

# UNDERSTANDING LAYERS

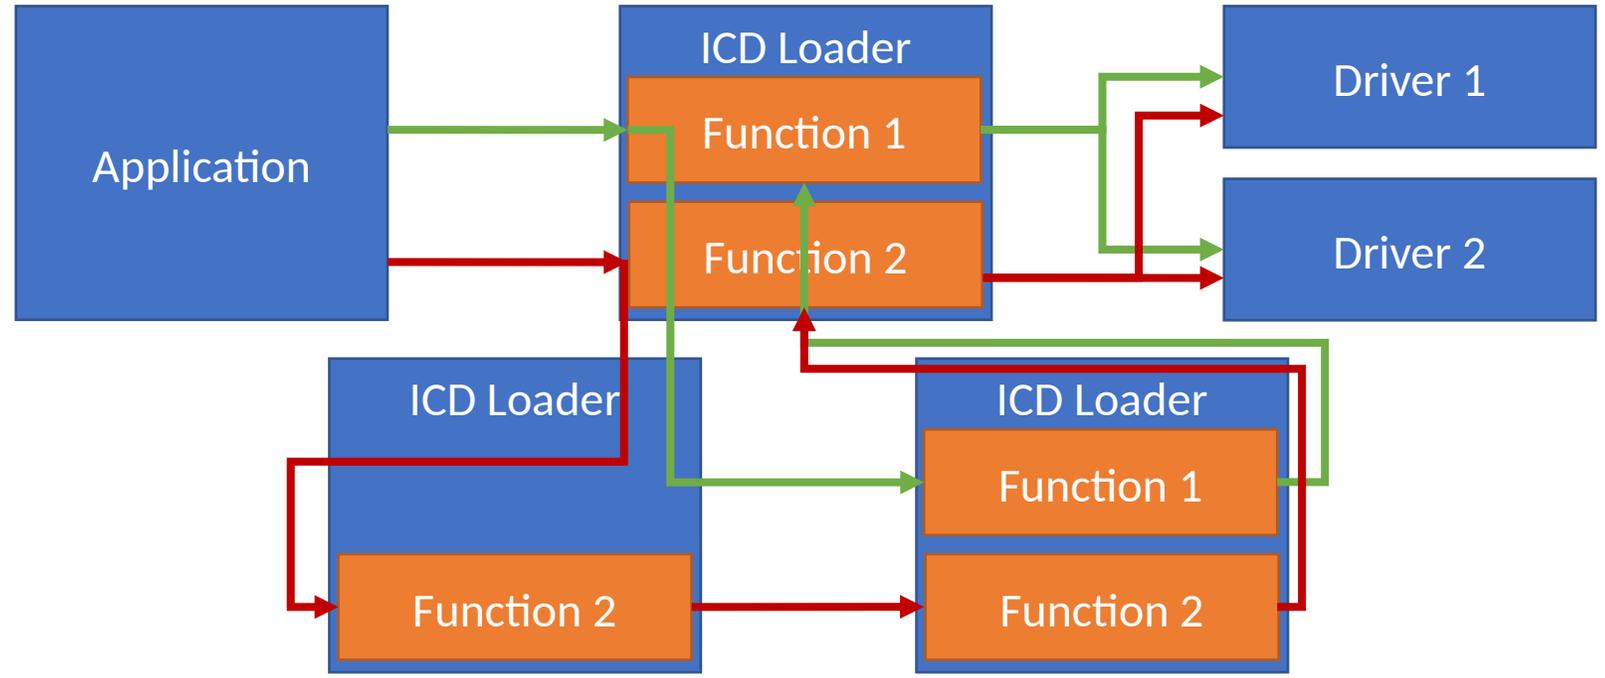
- Without layers support enabled, the ICD Loader will
  - dispatch the call to the correct vendor driver



Layers for OpenCL - Brice Videau @ IWOCL 2021  
(<https://raw.githubusercontent.com/Kerilk/OpenCL-Layers-Tutorial/main/presentation/LayersForOpenCL.pdf>)

# UNDERSTANDING LAYERS

- With layers support enabled, the ICD Loader will
  - first redirect calls to different active layers
  - then dispatch the call to the correct vendor driver



# OBJECT LIFETIME TRACKING

- Leaking objects

```
clReleaseEvent(evt[0]);  
clReleaseEvent(evt[1]);  
clReleaseMemObject(bufferSrc);  
clReleaseMemObject(bufferDst);  
clReleaseMemObject(bufferRes);  
clReleaseKernel(kernel);  
clReleaseProgram(program);  
clReleaseCommandQueue(queue);  
clReleaseContext(context);  
clReleaseDevice(device);
```

**LAYERS ARE  
KEY**

# OBJECT LIFETIME TRACKING

- Leaking objects

```
cl_ulong total_time = 0;
for (int i = 0 ; i < steps ; ++i)
{
    cl_event event;
    clEnqueueNDRangeKernel(
        queue, kernel, 1, NULL,
        &global_work_size, &local_work_size,
        0, NULL, &event);
    clGetEventProfilingInfo(
        event, CL_PROFILING_COMMAND_START,
        sizeof(cl_ulong), &start, NULL);
    clGetEventProfilingInfo(
        event, CL_PROFILING_COMMAND_END,
        sizeof(cl_ulong), &end, NULL);
    total_time += end - start;
}
```

# OBJECT LIFETIME TRACKING

- Leaking objects

```
cl_ulong total_time = 0;
for (int i = 0; i < steps; ++i)
{
    cl_event event;
    clEnqueueNDRangeKernel(
        queue, kernel, 1, NULL,
        &global_work_size, &local_work_size,
        0, NULL, &event);
    clGetEventProfilingInfo(
        event, CL_PROFILING_COMMAND_START,
        sizeof(cl_ulong), &start, NULL);
    clGetEventProfilingInfo(
        event, CL_PROFILING_COMMAND_END,
        sizeof(cl_ulong), &end, NULL);
    total_time += end - start;
    clReleaseEvent(event) ←
}
}
```

# OBJECT LIFETIME TRACKING

- Leaking objects
- Double free
  - Shortcircuit an option

```
cl_context ctx_a = clCreateContext(...);
clReleaseContext(ctx_a);
// Do a whole bunch of other stuff
/*
 *
 *
 *
 *
 *
 *
 *
 *
 *
 *
 */
clReleaseContext(ctx_a);
```

# OBJECT LIFETIME TRACKING

**LAYERS ARE  
KEY**

- Leaking objects
- Double free
  - Shortcircuit an option
- Use after free

```
cl_context ctx_a = clCreateContext(...);
clReleaseContext(ctx_a);
// Do a whole bunch of other stuff
/*
 *
 *
 *
 *
 *
 *
 *
 *
 *
 *
 */
clGetContextInfo(ctx_a,
CL_CONTEXT_DEVICES,
0, NULL, &devices_size);
```

# OBJECT LIFETIME TRACKING

- Leaking objects
- Double free
  - Shortcircuit an option
- Use after free
  - Use of dangling handle

```
cl_context ctx_a = clCreateContext(...);  
clReleaseContext(ctx_a);  
  
cl_context ctx_b = clCreateContext(...);  
clReleaseContext(ctx_a);
```

# OBJECT LIFETIME TRACKING

- Leaking objects
- Double free
  - Shortcircuit an option
- Use after free
  - Use of dangling handle
    - Needs a dictionary between ICD and user code
    - Translate whenever the user uses a handle

```
cl_context ctx_a = clCreateContext(...);  
clReleaseContext(ctx_a);  
  
cl_context ctx_b = clCreateContext(...);  
clReleaseContext(ctx_a);
```

# INPUT ARGUMENT VALIDATION

- API entry points have various pre/post-conditions users ought not violate
  - If conditions are violated, an error is raised
  - Worst case scenario, the app crashes
- There are (core + extension) functions
  - Writing checks for each is *a lot* of work
  - That work has to be done for every language projection
    - C, C++, C# (OpenCL.Net), Python (PyOpenCL), etc
  - Can we do better?

# INPUT ARGUMENT VALIDATION

```
cl_int clEnqueueWriteBuffer(cl_command_queue command_queue, cl_mem
buffer,
    cl_bool blocking_write, size_t offset, size_t size, const void* ptr,
    cl_uint num_events_in_wait_list, const cl_event* event_wait_list,
    cl_event* event);
```

- command\_queue is a valid host command-queue in which the read / write command will be queued. command\_queue and buffer must be created with the same OpenCL context.
- buffer refers to a valid buffer object.
- blocking\_read and blocking\_write indicate if the read and write operations are blocking or non-blocking (see below).
- offset is the offset in bytes in the buffer object to read from or write to.
- size is the size in bytes of data being read or written.
- ptr is the pointer to buffer in host memory where data is to be read into or to be written from.
- event\_wait\_list and num\_events\_in\_wait\_list specify events that need to complete before this particular command can be executed. If event\_wait\_list is NULL, then this particular command does not wait on any event to complete. If event\_wait\_list is NULL, num\_events\_in\_wait\_list must be 0. If event\_wait\_list is not NULL, the list of events pointed to by event\_wait\_list must be valid and num\_events\_in\_wait\_list must be greater than 0. The events specified in event\_wait\_list act as synchronization points. The context associated with events in event\_wait\_list and command\_queue must be the same. The memory associated with event\_wait\_list can be reused or freed after the function returns.
- event returns an event object that identifies this read / write command and can be used to query or queue a wait for this command to complete. If event is NULL or the enqueue is unsuccessful, no event will be created and therefore it will not be possible to query the status of this command or to wait for this command to complete. If event\_wait\_list and event are not NULL, event must not refer to an element of the event\_wait\_list array.

# INPUT ARGUMENT VALIDATION

```
cl_int clEnqueueWriteBuffer(cl_command_queue command_queue, cl_mem
buffer,
    cl_bool blocking_write, size_t offset, size_t size, const void* ptr,
    cl_uint num_events_in_wait_list, const cl_event* event_wait_list,
    cl_event* event);
```

- CL\_INVALID\_COMMAND\_QUEUE if command\_queue is not a valid host command-queue.
- CL\_INVALID\_CONTEXT if the context associated with command\_queue and buffer are not the same or if the context associated with command\_queue and events in event\_wait\_list are not the same.
- CL\_INVALID\_MEM\_OBJECT if buffer is not a valid buffer object.
- CL\_INVALID\_VALUE if the region being read or written specified by (offset, size) is out of bounds or if ptr is a NULL value.
- CL\_INVALID\_EVENT\_WAIT\_LIST if event\_wait\_list is NULL and num\_events\_in\_wait\_list > 0, or event\_wait\_list is not NULL and num\_events\_in\_wait\_list is 0, or if event objects in event\_wait\_list are not valid events.
- CL\_MISALIGNED\_SUB\_BUFFER\_OFFSET if buffer is a sub-buffer object and offset specified when the sub-buffer object is created is not aligned to CL\_DEVICE\_MEM\_BASE\_ADDR\_ALIGN value for device associated with queue. This error code is missing before version 1.1.
- CL\_EXEC\_STATUS\_ERROR\_FOR\_EVENTS\_IN\_WAIT\_LIST if the read and write operations are blocking and the execution status of any of the events in event\_wait\_list is a negative integer value. This error code is missing before version 1.1.
- CL\_MEM\_OBJECT\_ALLOCATION\_FAILURE if there is a failure to allocate memory for data store associated with buffer.
- CL\_INVALID\_OPERATION if clEnqueueReadBuffer is called on buffer which has been created with CL\_MEM\_HOST\_WRITE\_ONLY or CL\_MEM\_HOST\_NO\_ACCESS.
- CL\_INVALID\_OPERATION if clEnqueueWriteBuffer is called on buffer which has been created with CL\_MEM\_HOST\_READ\_ONLY or CL\_MEM\_HOST\_NO\_ACCESS.
- CL\_OUT\_OF\_RESOURCES if there is a failure to allocate resources required by the OpenCL implementation on the device.
- CL\_OUT\_OF\_HOST\_MEMORY if there is a failure to allocate resources required by the OpenCL implementation on the host.

# GENERATE CHECKS FROM XML

## Before input arg validation

```
<command  
suffix="CL_API_SUFFIX__VERSION_1_2  
">  
  <proto>  
    <type>cl_int</type>  
    <name>clReleaseDevice</name>  
  </proto>  
  <param>  
    <type>cl_device_id</type>  
    <name>device</name>  
  </param>  
</command>
```

**LAYERS ARE  
KEY**

# GENERATE CHECKS FROM XML

## Before input arg validation

```
<command  
suffix="CL_API_SUFFIX__VERSION_1_2"  
>  
  <proto>  
    <type>cl_int</type>  
    <name>clReleaseDevice</name>  
  </proto>  
  <param>  
    <type>cl_device_id</type>  
    <name>device</name>  
  </param>  
</command>
```

## After input arg validation

```
<command  
suffix="CL_API_SUFFIX__VERSION_1_2"  
>  
  <proto>  
    <type>cl_int</type>  
    <name>clReleaseDevice</name>  
  </proto>  
  <param>  
    <type>cl_device_id</type>  
    <name>device</name>  
  </param>  
  <if>  
    <object_is_invalid  
      name="device"/>  
  </if>  
  <then>  
    <log>device is not a valid  
      device</log>  
    <name>clReleaseDevice</name>  
    <value>CL_INVALID_DEVICE</value>  
  </then>  
</command>
```

LAYERS ARE  
KEY

# IN ACTUAL NEWS

What happens in OpenCL land?

# OPENCL ON GPUINFO

- Popular database of GPU capabilities
- Supports
  - OpenGL, Vulkan, OpenGL ES and now OpenCL!
  - Windows, Linux, Android
- Invaluable if you wish to
  - check the capabilities of a device you don't own
  - gauge portability when depending on a feature/extension

The screenshot shows the OpenCL Hardware Database website. The page title is "Listing devices". There are tabs for "All platforms", "Windows", "Linux", and "Android". Below the tabs is a table with the following columns: "Device", "Max. API version", "Latest Driver version", "Last submission", "Count", and "compare". The table contains several rows of GPU data.

Device	Max. API version	Latest Driver version	Last submission	Count	compare
AMD Radeon Pro WX 5100 Graphics (pola...	1.1	22.1.1	2022-06-18 13:19:22	1	<input type="checkbox"/>
AMD Radeon Pro WX 5100 Graphics (pola...	1.1	22.1.1	2022-06-18 02:12:35	5	<input type="checkbox"/>
QUALCOMM YotaDevices YOTA 3+	2.0	OpenCL 2.0 QUALCO...	2022-06-17 20:23:54	1	<input type="checkbox"/>
NVIDIA GeForce RTX 3090	3.0	516.40	2022-06-15 23:38:18	10	<input type="checkbox"/>
NVIDIA GeForce GTX 1060 6GB	3.0	516.40	2022-06-15 21:02:40	7	<input type="checkbox"/>
Intel(R) HD Graphics 530	3.0	31.0.101.1999	2022-06-14 23:48:09	14	<input type="checkbox"/>
Intel(R) Iris(R) Xe Graphics	3.0	30.0.101.3109	2022-06-14 17:44:56	7	<input type="checkbox"/>
Intel(R) UHD Graphics	3.0	30.0.101.1994	2022-06-14 16:40:23	3	<input type="checkbox"/>
NVIDIA T1200 Laptop GPU	3.0	472.91	2022-06-14 16:40:08	1	<input type="checkbox"/>
Intel(R) UHD Graphics 630	3.0	31.0.101.1999	2022-06-14 14:55:00	8	<input type="checkbox"/>
Quadro RTX 4000	3.0	516.25	2022-06-14 14:52:44	1	<input type="checkbox"/>
Blackview BV5800 PRO	1.2	1.0@4802505	2022-06-13 21:10:20	1	<input type="checkbox"/>

<https://opencl.gpuinfo.org/>

# SPIRV2CLC

- Kévin Petit has done some awesome work by releasing a SPIR-V to OpenCL C transpiler
  - Calling it disassembling would be a mild overstatement
  - „If you squint really, you can see the SPIR-V”

**IN ACTUAL  
NEWS**

# SPIRV2CLC

- Kévin Petit has done some awesome work by releasing a SPIR-V to OpenCL C transpiler
  - Calling it disassembling would be a mild overstatement
  - „If you squint really, you can see the SPIR-V”

```
kernel void saxpy(  
    float a,  
    global float* x,  
    global float* y)  
{  
    int gid = get_global_id(0);  
    y[gid] = a * x[gid] + y[gid];  
}
```

```
kernel void saxpy(float v20, float global* v21, float global* v22){  
v23;;  
    float __attribute__((aligned(4))) v5_storage; float * v5 =  
&v5_storage;  
    float global* __attribute__((aligned(8))) v6_storage; float global*  
* v6 = &v6_storage;  
    float global* __attribute__((aligned(8))) v7_storage; float global*  
* v7 = &v7_storage;  
    uint __attribute__((aligned(4))) v8_storage; uint * v8 =  
&v8_storage;  
    *v5 = v20;  
    *v6 = v21;  
    *v7 = v22;  
    ulong v25 = get_global_id(0);  
    uint v26 = convert_uint(v25);  
    *v8 = v26;  
    float v27 = *v5;  
    float global* v28 = *v6;  
    uint v29 = *v8;  
    ulong v30 = convert_long(as_int(v29));  
    float global* v31 = &v28[v30];  
    float v32 = *v31;  
    float global* v33 = *v7;  
    uint v34 = *v8;  
    ulong v35 = convert_long(as_int(v34));  
    float global* v36 = &v33[v35];  
    float v37 = *v36;  
    float v38 = mad(v27, v32, v37);  
    float global* v39 = *v7;  
    uint v40 = *v8;  
    ulong v41 = convert_long(as_int(v40));  
    float global* v42 = &v39[v41];  
    *v42 = v38;  
    return;  
}
```

- It is the PTX of OpenCL
- Was born out of necessity of vendors wanting to obfuscate proprietary kernel code
  - The API at some point needs device code as an unencrypted stream of ASCII char array (easy to hijack)
  - Also solves some vendor compiler non-conformance issues
- SPIR turned out to be possibly the biggest mistake OpenCL has ever made
  - A mistake to be repeated by Microsoft with DXIL
- SPIR-V saves the day („V” stands for Vulkan, not 5)
  - Vulkan needed an IR badly
  - New IR is adopted by both OpenCL and OpenGL

# STANDARD PORTABLE IR



**IN ACTUAL  
NEWS**

- It is the PTX of OpenCL

Aspect	SPIR 1.2	SPIR 2.0	SPIR-V 1.X
LLVM Interaction	Uses LLVM 3.2	Uses LLVM 3.4	100% Khronos defined Round-trip lossless conversion
Compute Constructs	Metadata/Intrinsics	Metadata/Intrinsics	Native
Graphics Construct	No	No	Native
Supported Language Feature Supported	OpenCL C 1.2	OpenCL C 1.2 OpenCL C 2.0	OpenCL C 1.2 / 2.X C++ for OpenCL GLSL
OpenCL Ingestion	OpenCL 1.2 Extension	OpenCL 2.0 Extension	OpenCL 2.1/2.2 Core OpenCL 3.0 Extension
Graphics API Ingestion	-	-	Vulkan 1.X OpenGL 4.6 Core

Source: <https://www.khronos.org/spir/>

# STANDARD PORTABLE IR



- It is the PTX of OpenCL

Aspect	SPIR 1.2	SPIR 2.0	SPIR-V 1.X
LLVM Interaction	Uses <b>LLVM</b> 3.2	Uses <b>LLVM</b> 3.4	100% Khronos defined Round-trip lossless conversion
Compute Constructs	Metadata/Intrinsics	Metadata/Intrinsics	Native
Graphics Construct	No	No	Native
Supported Language Feature Supported	OpenCL C 1.2	OpenCL C 1.2 OpenCL C 2.0	OpenCL C 1.2 / 2.X C++ for OpenCL GLSL
OpenCL Ingestion	OpenCL 1.2 Extension	OpenCL 2.0 Extension	OpenCL 2.1/2.2 Core OpenCL 3.0 Extension
Graphics API Ingestion	-	-	Vulkan 1.X OpenGL 4.6 Core

Source: <https://www.khronos.org/spir/>

**IN ACTUAL  
NEWS**

## SPIR-V FOR GRAPHICS

- 100% Khronos defined
- Bi-directionally translatable to/from LLVM
- Can represent any shading language
- Rich, open-source tooling

## SPIR-V FOR COMPUTE

- Everything SPIR-V for graphics does and more
- Unstructured control-flow
- Pointer arithmetic
  - Optimizer looks very different
  - Both extensions can be transformed away
- Lots of optimizer work
- Done in Mesa
  - OpenCLOn12, clvk

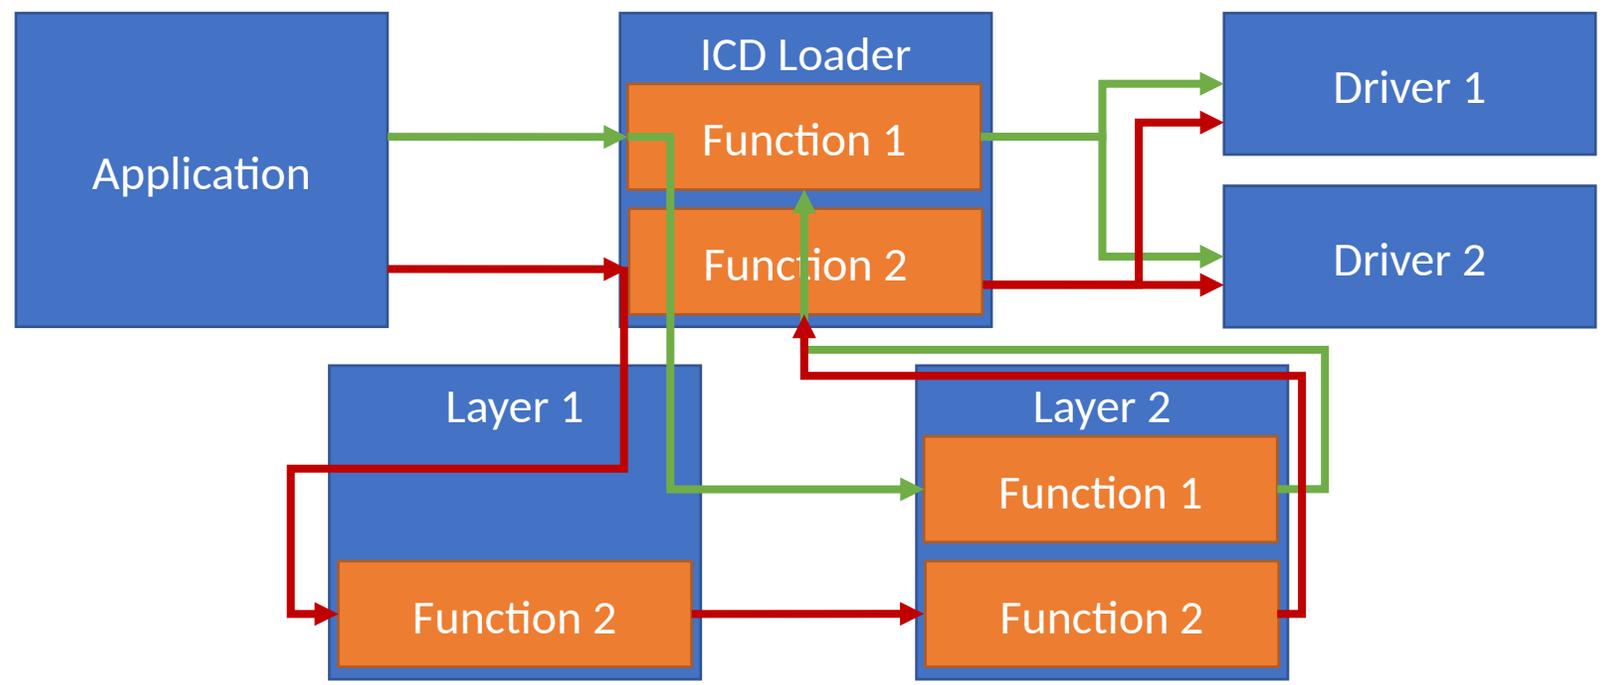
**IN ACTUAL  
NEWS**

# THE STATE OF SPIR-V

- SPIR-V is a crucial, yet notoriously non-portable feature
  - It's one of the reasons why Vulkan is such a huge success
  - Tooling around SPIR-V is thriving
  - Greatly improves portability of shaders and tooling
- On desktop, AMD and Nvidia still don't ship SPIR-V compilers for OpenCL
  - Even OpenCLOn12 supports it
  - cl\_khr\_il\_program is the extension to look for
- Can we teach runtimes that don't support SPIR-V ingestion to somehow... compile them?
  - I thought you'd never ask

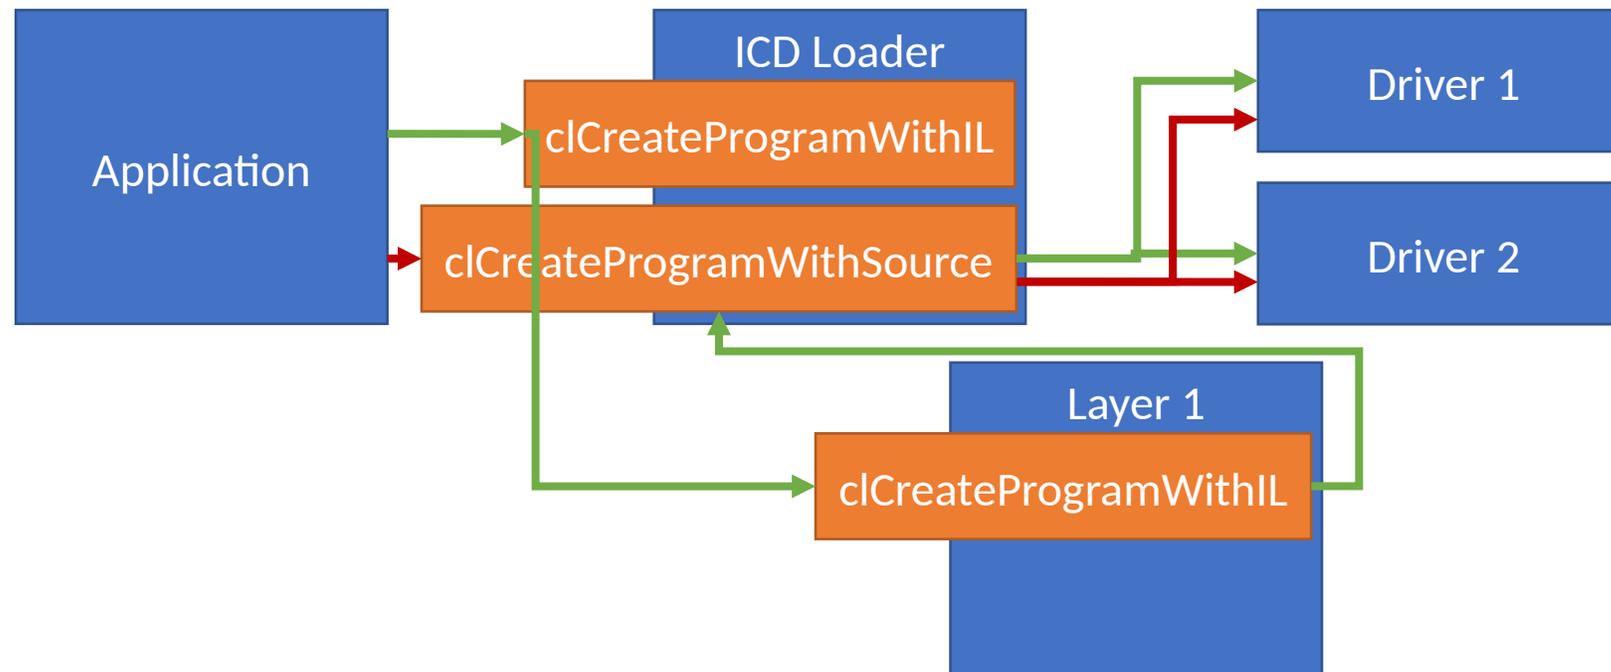
# SPIR-V FOR THE MASSES

- Are you thinking what I'm thinking? Layers!
- Nothing prevents a layer from calling into a different API function than the one being called.



# SPIR-V FOR THE MASSES

- Are you thinking what I'm thinking? Layers!
- Nothing prevents a layer from calling into a different API function than the one being called.



# C++ FOR OPENCL

- Language built on top of OpenCL C 3.0 unified and C++17 enabling most of regular C++ features in OpenCL kernel code.
  - List of restricted C++ features are the same as always
    - Virtual functions, dynamic\_cast, refs to func, ptr to member func, RTTI, exceptions, thread\_local, non-placement new/delete, STL
- Upstream LLVM 14 with experimental support
  - Need to compile kernels offline
  - Load IL as binary and feed to clCreateProgramWithIL

```
clang -cl-std=CLC++2021 kernel.clcpp
```

```
std::ifstream binary{ location, std::ios::binary };
cl::Program prog{ ctx, std::vector{ std::istreambuf_iterator<char>{ binary },
                                     std::istreambuf_iterator<char>{} } };
prog.build({ device });
```

# C++ FOR OPENCL

- Language built on top of OpenCL C 3.0 unified and C++17 enabling most of regular C++ features in OpenCL kernel code.
  - List of restricted C++ features are the same as always
    - Virtual functions, dynamic\_cast, refs to func, ptr to member func, RTTI, exceptions, thread\_local, non-placement new/delete, STL
- Upstream LLVM 14 with experimental support
  - Need to compile kernels offline `clang -cl-std=CLC++2021 kernel.clcpp`
  - Load IL as binary and feed to `clCreateProgramWithIL`

```
std::ifstream binary{ location, std::ios::binary };
cl::Program prog{ ctx, std::vector{ std::istreambuf_iterator<char>{ binary
},
                                std::istreambuf_iterator<char>{} } } };
prog.build({ device });
```

# **WE'RE HIRING**

<https://streamhpc.com/jobs/>