

X.Org Developers Conference 2020

Report of Contributions

Contribution ID: 1

Type: **not specified**

State of text input on Wayland

Wednesday, 16 September 2020 20:15 (5 minutes)

Between the last impromptu talk at GUADEC 2018, text input on Wayland has become more organized and more widely adopted. As before, the three-pronged approach of `text_input`, `input_method`, and virtual keyboard still causes confusion, but increased interest in implementing it helps find problems and come closer to something that really works for many use cases.

The talk will mention how a broken assumption causes a broken protocol, and why we're not done with Wayland input methods yet.

It's recommended to people who want to know more about the current state of input methods on Wayland. Recommended background: aforementioned GUADEC talk, `wayland-protocols` repository, my blog: https://dcz_self.gitlab.io/

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: DCZ, Dorota

Session Classification: Demos / Lightning talks I

Track Classification: Lightning Talk

Contribution ID: 2

Type: **not specified**

IGT GPU Tools 2020 Update

Wednesday, 16 September 2020 20:00 (5 minutes)

Short update on IGT - what has changed in the last year, where are we right now and what we have planned for the near future.

IGT GPU Tools is a collection of tools and tests aiding development of DRM drivers. It's widely used by Intel in its public CI system. IGT has targeted test for v3d, panfrost, amdgpu, etc.

I am one of the maintainers.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: HILER, Arkadiusz (Intel)

Session Classification: Demos / Lightning talks I

Track Classification: Lightning Talk

Contribution ID: 3

Type: **not specified**

Overview of the open source Vulkan driver for Raspberry Pi 4

Wednesday, 16 September 2020 13:20 (45 minutes)

Igalia has been developing a new open source Mesa driver for the Raspberry Pi 4 since December 2019. This talk will discuss the development story and current status of the driver, provide a high level overview of the major design elements, discuss some of the challenges we found in bringing specific aspects of Vulkan 1.0 to the V3D GPU platform and finally, talk about future plans and how to contribute to the on-going development effort.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: TORAL, Iago (Igalia)

Session Classification: Main Track

Track Classification: Talk (full slot) (Closed)

Contribution ID: 4

Type: **not specified**

Why is Peer to Peer DMA so hard on Linux?

Wednesday, 16 September 2020 19:30 (20 minutes)

Whether it is HPC or gaming, peer to peer DMA is an important part of improving IO throughput and performance on servers and workstations and yet, it has only recently become barely functional on Linux. This talk delves into the history of peer to peer DMA on Linux, why it is so challenging, what the current landscape looks like, and ways we can improve in the future.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

Primary author: DEUCHER, Alex (AMD)

Session Classification: Main Track

Track Classification: Talk (half slot) (Closed)

Contribution ID: 5

Type: **not specified**

Secure Buffer Object Support with Trusted Memory Zone

Wednesday, 16 September 2020 17:05 (20 minutes)

Memory encryption is an important part of content protection schemes like Widewine L1. This talk will delve into the details of TMZ (Trusted Memory Zone) on AMD GPUs touching on the software implementation details and hardware requirements and limitations that needed to be addressed to support this in the kernel, mesa, and userspace applications.

References:

- 1.https://www.phoronix.com/scan.php?page=news_item&px=AMD-Trusted-Memory-Zone
- 2.<https://lists.freedesktop.org/archives/amd-gfx/2019-September/039928.html>
- 3.https://linuxreviews.org/Trusted_Memory_Zone_Support_Coming_To_AMD_APUs_in_Linux_Kernel_5.6
- 4.https://gitlab.freedesktop.org/mesa/mesa/-/merge_requests/4401

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: Mr HUANG, Ray (AMD)

Session Classification: Main Track

Track Classification: Talk (half slot) (Closed)

Contribution ID: 6

Type: **not specified**

WSL Graphics Architecture

Wednesday, 16 September 2020 14:15 (45 minutes)

Microsoft announced at //build2020 that support for GPU hardware acceleration, through virtual GPU, was coming to the Windows Subsystem for Linux (WSL). This support enables Linux applications running in a WSL VM to leverage and share the host GPU through a variety of well-known graphics and compute APIs.

This talk will give an overview of the architecture, all the way from the Windows kernel, to the Linux kernel, to Linux userspace: How the various pieces fit together to enable GPU acceleration in various scenarios, from ML and AI compute tools and framework to accelerating rendering of GUI applications. We will go through some of the design choices we made and how we're striving toward making WSL a great environment to experience Linux applications.

This will also be a good opportunity to provide feedback on this design directly to the engineers at Microsoft, and help ensure that the right thing is being built and maintained.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary authors: NATALIE, Jesse (Microsoft); PRONOVOST, Steve (Microsoft)

Session Classification: Main Track

Track Classification: Talk (full slot) (Closed)

Contribution ID: 7

Type: **not specified**

Mesa for D3D12 Mapping Layers

Wednesday, 16 September 2020 16:05 (20 minutes)

Mesa already is host to at least one API mapping layer: Zink. Building on the success of that layer, Microsoft has partnered with Collabora to build another mapping layer as a Gallium driver in Mesa: OpenGLOn12. At the same time, Microsoft has built a small OpenCLOn12 runtime, and is re-using and improving Clover's compiler stack, combined with the NIR to DXIL translator built for OpenGL, to provide a story for OpenCL support as well. This talk will discuss architecture, status, and future plans.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: NATALIE, Jesse (Microsoft)

Session Classification: Main Track

Track Classification: Talk (half slot) (Closed)

Contribution ID: 8

Type: **not specified**

X11 and Wayland applications in WSL

Wednesday, 16 September 2020 15:10 (45 minutes)

Microsoft announced at //build2020 that support for running Linux GUI applications, X11 and Wayland, was coming to the Windows Subsystem for Linux (WSL). This will give developers choosing to use Windows as their desktop of choice, the ability to run their preferred Linux applications in a unified, integrated and seamless desktop experience.

In this talk we take a deep dive into the architecture that enables this support. We will go over details of the Weston based Wayland compositor we are building and how we are teaching it about application level remoting across the VM (WSL) to host (Windows) boundary. How we integrate remote applications into a unified desktop experience and give them that local application feel. How GUI applications will be able to leverage our WSL virtual GPU projection to accelerate their rendering through native Linux rendering API. We will explore how the architecture of WSL is evolving to host this compositor, how it will be delivered to user and how it will enable GUI application across WSL distros.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: PRONOVOST, Steve (Microsoft)

Session Classification: Main Track

Track Classification: Talk (full slot) (Closed)

Contribution ID: 9

Type: **not specified**

Improving Khronos CTS tests with Mesa code coverage

Friday, 18 September 2020 13:20 (20 minutes)

The Khronos Conformance Test Suite is an open-source testing suite developed by the Khronos Group to certify that a given driver is conformant to the respective graphics API specification (OpenGL, OpenGL ES, Vulkan). As this testing suite is publicly available on Github, many Mesa driver developers use it, together with piglit and other tools, to make sure the driver follows the specification, there are no regressions when adding a new change, or to test new features under development.

However, the Khronos CTS tests are not perfect. Sometimes they miss checking some SPIR-V opcodes, or all the different data type options for the arguments to a given opcode, or they don't call all the API functions... to name a few things.

In this talk, we will introduce the work done by Igalia to easily detect low-hanging fruit missing test coverage in this testing suite. Thanks to this work, we have added more test coverage to many Vulkan CTS tests that will ultimately benefit all of Mesa's open-source Vulkan drivers. We will explain how we did it and the lessons learned from that work.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: IGLESIAS GONSÁLVIZ, Samuel (Igalia)

Session Classification: Main Track

Track Classification: Talk (half slot) (Closed)

Contribution ID: 10

Type: **not specified**

How the Vulkan VK_EXT_extended_dynamic_state extension came to be

Friday, 18 September 2020 13:50 (20 minutes)

VK_EXT_extended_dynamic_state is an interesting Vulkan extension that was released recently. This talk will explain the extension purpose and the role it can play in making Vulkan pipelines more flexible and simplifying many Vulkan applications. It will also focus on the events that sparked the effort to create the extension inside the Khronos Group, making it an interesting case study, covering the process from the design phase to having support for it landed in RADV. As part of this, the talk will also go over the preferred way to contribute to the Vulkan specification.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

Primary author: GARCIA, Ricardo

Session Classification: Main Track

Track Classification: Talk (half slot) (Closed)

Contribution ID: 11

Type: **not specified**

Status of freedesktop.org gitlab/cloud hosting

Friday, 18 September 2020 14:20 (45 minutes)

In this talk, we will peek behind the curtain of the freedesktop.org infrastructure, its costs and how we attempted to reduce those.

At the end of 2019 we realized that our gitlab hosting on GCE was costing us more than expected. So we analyzed the costs and developed countermeasures to reduce those costs. This talk will explain the various analysis steps, the measures we took and future measures we are contemplating.

This talk doesn't require any technical knowledge of the various technologies in use. However, we will definitely talk about kubernetes, ingress, egress, cloud, storage, CI, and other insanities, but we will always try to explain those terms to have the widest audience possible. The purpose is to disclose how we spend money, why, and what we are doing or will be doing to reduce that bill.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: TISSOIRES, Benjamin (Red Hat)

Session Classification: Main Track

Track Classification: Talk (full slot) (Closed)

Contribution ID: 12

Type: **not specified**

Introducing adriconf, a tool for mesa driver configuration

Friday, 18 September 2020 16:45 (20 minutes)

In this talk we are going over adriconf, where it all started and what is the current state of mesa driver configuration from a end user perspective.

We will also briefly talk about the MESA_query_driver extension and why we need it as well as future Vulkan extension/work on the configuration area.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: HERTEL, Jean

Session Classification: Main Track

Track Classification: Talk (half slot) (Closed)

Contribution ID: 13

Type: **not specified**

Allocation Constraints

Wednesday, 16 September 2020 18:35 (45 minutes)

I will present a proposal for integrating memory constraints into the Linux graphics software stack, including the kernel, userspace graphics drivers, and windowing systems. Constraints, or properties describing the various limitations imposed by devices with direct access to memory, are the second half of the prototype allocator design originally proposed at XDC 2017. I will contrast “capabilities,” which are currently well represented by DRM format modifiers as described at XDC 2019, with constraints, building the case for a separate mechanism for each. Examples, focusing on the constraints imposed by NVIDIA hardware, will be given to further illustrate the proposed design and the motivations behind it.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: JONES, James (NVIDIA)

Session Classification: Main Track

Track Classification: Talk (full slot) (Closed)

Contribution ID: 14

Type: **not specified**

Ray-tracing in Vulkan: A brief overview of the provisional VK_KHR_ray_tracing API

Thursday, 17 September 2020 17:35 (45 minutes)

Earlier this year, Khronos released a provisional VK_KHR_ray_tracing extension for HW-accelerated ray-tracing with the Vulkan API. In this talk, Jason will introduce the basics of ray-tracing and give an overview of the new shader stages, objects, and other concepts used to accelerate ray-tracing via the new Vulkan extension. The talk will be educational and focus on helping others in the X/Mesa community understand the new API concepts and will contain few if any implementation details.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: EKSTRAND, Jason (Intel)

Session Classification: Main Track

Track Classification: Talk (full slot) (Closed)

Contribution ID: 15

Type: **not specified**

GSoC/EVoC Overview

Wednesday, 16 September 2020 17:35 (20 minutes)

A quick talk to give an overview of GSoC/EVoC and XOrg's involvement therein.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: WOERNER, Trevor

Session Classification: Main Track

Track Classification: Talk (half slot) (Closed)

Contribution ID: 16

Type: **not specified**

VKMS improvements using IGT GPU Tools

Wednesday, 16 September 2020 18:05 (20 minutes)

The Virtual Kernel Mode Setting (VKMS) driver allows you to test DRM and run X on systems without a physical display, making it a great candidate for running inside a virtual machine for CI purposes. Its development intends to expand the test coverage of DRM, giving graphic developers greater autonomy to verify the subsystem's expected operation and develop new features. However, to enjoy these benefits, we need to ensure that VKMS performs well on all sets of basic tests provided by IGT. Aiming to bring more consistency to this module, my work in this year's GSoC was to deliver a fully working and bug-free subset of GPU tests. To achieve this, I had to understand, improve, and decide where to act between the two sides: DRM/VKMS and IGT GPU Tools. In this presentation, I will share our progress on VKMS and subsequently on IGT during this summer. As a newcomer, I also want to share my experience of figuring out, developing, and reaching suitable solutions with the community.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

Yes

Primary author: WEN, Melissa (University of São Paulo)

Session Classification: Main Track

Track Classification: Talk (half slot) (Closed)

Contribution ID: 17

Type: **not specified**

A year of ACO: from prototype to default

Thursday, 17 September 2020 13:20 (45 minutes)

ACO is a new compiler backend for AMD GCN/RDNA GPUs, introduced a year ago in summer 2019 as an experimental prototype sponsored by Valve, and has recently become the default compiler backend of RADV (the Mesa Radeon Vulkan driver).

This talk is about our journey of how we evolved the design of ACO as well as the decisions we took along the road towards feature parity with the LLVM backend as we added all the bits and pieces that we needed in order to extend ACO to support all shader stages and extensions on every hardware generation.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: KRISTÓF, Timur (Valve)

Session Classification: Main Track

Track Classification: Talk (full slot) (Closed)

Contribution ID: 18

Type: **not specified**

LiteDIP: bridging the gap between open source hardware, and open source operating systems

Thursday, 17 September 2020 15:40 (20 minutes)

Most GPUs now have open source drivers, and the trend is for all of them to be treated not as a curiosity, but instead being full-featured and providing an excellent user experience. To further push the open source philosophy, we need to look at the next frontier: Open Source Hardware.

While usual hardware development is prohibitively expensive, reconfigurable hardware (FPGA) is accessible to every hobbyists! This type of hardware has historically been very expensive and unable to provide the necessary performance to achieve any sort of satisfactory user experience, but the cost has dropped dramatically in the past 20 years, and the rise of hardware blocks such as PCIe, DDR memory controllers, and ultra-fast transceivers have enabled the creation of open PCIe display controllers capable of reaching 4K and more for a reasonable amount of money.

Writing open source drivers for such hardware is however a little tricky since users will likely want to mix-and-match the different open source blocks to tailor the features to their liking, and even do this at run time!

In this talk, I will introduce the idea behind LiteDIP, my project of creating a library of discoverable IP blocks for FPGAs along with their Linux driver which would enable users to configure and deploy their own System on Chip in ~10 minutes.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: PERES, Martin

Session Classification: Main Track

Track Classification: Talk (half slot) (Closed)

Contribution ID: 19

Type: **not specified**

etnaviv: The wonderful world of performance counters

Thursday, 17 September 2020 14:45 (45 minutes)

Performance counter are somewhat special on Vivante GPUs. It is not possible to read them via cmd stream but only from the CPU/kernel. This needs some extra work in the kernel and the user space.

The final goal is to have per-draw performance counter values for detailed analysis of performance problems and a way to sample performance counters in a cyclic way for perfetto or some kind of gpu-top tool. GPU load values are also quite special and might be of interest. Overall there is quite some work to be done to get it up and running.

I would talk about the problems and the solutions I came up with.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: GMEINER, Christian

Session Classification: Main Track

Track Classification: Talk (full slot) (Closed)

Contribution ID: 20

Type: **not specified**

etnaviv: status update

Thursday, 17 September 2020 14:15 (20 minutes)

A short status update what happens in etnaviv land right now. Would cover a quick recap about the history, current state, some CI news and whats worked on right now.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: GMEINER, Christian

Session Classification: Main Track

Track Classification: Talk (half slot) (Closed)

Contribution ID: 21

Type: **not specified**

Software and hardware images decoding on the RaspberryPi

Friday, 18 September 2020 16:15 (20 minutes)

When it comes to the hardware acceleration on the RaspberryPi (or any other board, really), we often talk about the video encoding/decoding. With the modern ARM CPUs (with NEON support) and libraries (like libjpeg-turbo), usage of the dedicated hardware components for images encoding/decoding becomes less important. However, there is still low-end hardware on the market (like the RaspberryPi Zero) which can greatly benefit from usage of the hardware images decoding.

In this talk we will compare the performance of the software/hardware images decoding on the RaspberryPi devices. We will focus on the RaspberryPi Zero, as in this case the performance gain from using the hardware acceleration is the most significant. We base our experience on the digital-signage usecases, where both low device price and performance matters.

Although the OpenMAX is said to be practically deprecated, there might be no alternative to achieve the same level of performance on the RaspberryPi Zero. We will briefly present how the OpenMAX IL API is used to decode and display JPEG images. Apart from decoding 1080p images, we will also show how it performs when decoding the 4K images or how it can be used to zoom part of the image.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

Primary author: Mr PIJANOWSKI, Maciej (3mdeb Embedded Systems Consulting)

Session Classification: Main Track

Track Classification: Talk (half slot) (Closed)

Contribution ID: 22

Type: **not specified**

Don't bake your Graphics cards!

Thursday, 17 September 2020 19:00 (20 minutes)

Well, gaming enthusiasts might have tried baking their graphics card while tinkering with their old graphics hardware but doing so is not a great idea though some people might still try it!

https://www.reddit.com/r/pcmasterrace/comments/6zfrf1/should_you_bake_your_graphics_card/

On the contrary, this document describes a power saving, thermal efficient feature known as AMD Zero Power Technology. Though most people like its synonym, BACO!

BACO - Bus Alive Chip Off

BACO is an idle state of the dGPU which is employed in idle scenarios for long idle power requirements. BACO is entered when dGPU has been idle for some time and display has gone blank or when there is no compute work load. Driver support is required to save the video memory and other required information as part of BACO entry sequence. More on it later.

There are other related features as well, similar to BACO.

BOCO: Bus Off, Chip Off. For Notebooks that support legacy S3 sleep. (ACPI)

BAMACO: Bus Alive, Memory Alive, Chip Off. For Desktops that support Modern Standby or Linux suspend to idle state or active S0ix states. (Connected Modern Standby)

BOMACO: Bus Off, Memory Alive, Chip Off. For Notebooks that support Modern Standby special case for deep S0ix states (Disconnected Modern Standby)

The purpose of this abstract is to outline the design and implementation details of AMD Zero Core power technology with focus on BACO.

BACO refers to a hardware state that allows the GPU to save as much power as possible on the graphics chip when it is not being used. The main purpose for this state is to keep power consumption in the dGPU as low as possible when it is not being used while keeping the PCIe configuration space alive. Keeping PCIe configuration space alive maintains device presence for the Linux kernel. The System Management Unit (SMU) implements the actual entry-exit sequence algorithm and it needs inputs from amdgpu driver to detect idleness and trigger entry and exit events. When in BACO, SMU turns power off for as many IP blocks, as possible and gates most PLLs such as SPLL, DPLL etc, but keeps for the bus interface intact. BIF maintains PCIe configuration space and OS configuration requests. BIF switches its clock to PCIe ref clk while GPU is in BACO state. In that state, only a part of the SOC logic remains on such as Thermal, SMBUS interface, DFX and most part of NBIO.

amdgpu driver supports Linux runtime power management framework already and the efforts to integrate goodness of BACO with it were going on from quite some time now <https://lists.freedesktop.org/archives/amd-gfx/2019-February/031552.html>. AMDKFD driver which is a part of amdgpu driver and is used for compute and machine learning workloads, didn't implement support for runtime power management but with recent kernels that dependency is resolved. <https://lists.freedesktop.org/archives/amd-gfx/2020-February/045504.html>

Here's is a plausible scenario that explains how BACO fits into standard runtime pm framework and this should also apply to suspend-to-idle flow w.r.t amdgpu driver.

When there is idleness in the system i.e there is no CPU bound work load the CPUs tend to go to deeper sleep states a.k.a C-states and when the devices such as GPU don't have any graphics or compute kernel to act upon they go to device idle states called D-states. D-state can have sub-states like D0, D1, D2, D0i3, D3Hot, D3Cold, and D3 with D0 as active state and D3 as most power efficient state. Devices that support runtime pm can announce their capabilities in their

PCIe config space and they write 1 to their PMCSR bits when in a D state like D3. How a device implements its sleep state is up to the device and one such power optimization is BACO that works with D3Hot / D0i3 states.

SMU IP Block

Platform Security Engine (PSP)
Power Management Engine (PMU)
Integrated Sensor Hub

SMU IP block contains three important sub units and one of those is known as PMU that implements BACO entry / exit sequencing.

BACO Entry Sequence at a high level

Amdgpu driver notify SMU it would like to enter BACO.

If SMU detects GPU is actually not at idle, it doesn't respond to the driver, hence no trigger for BACO entry.

If SMU finds GPU to be idle, it issues an interrupt to driver to start BACO entry process, followed by below steps.

driver take below steps as the real start of BACO entry :

FB (Frame Buffer) content saving

Put THERMAL PHY at low power mode to save BACO mode power

For any scenarios if enabled displaypll, can put displaypll at low power mode as well

Enable doorbell_monitor. More details on it below.

program Dstate change to use bypass mode; hereafter, any Dstate change would be handled by BIF.

driver signals BACO Entry Event to SMU firmware

SMU firmware take below steps and many more :

disables all BACO domain IP's related rSMU interrupt

ramp downs / gates PLLs.

turns off voltage rails to quiescence the device.

Since the BACO state cuts the power on video memory, we have to make sure all contents allocated in the video memory will be saved / restored properly. When changing BACO state, we need to check if audio device is busy for the following scenarios:

When audio device is busy, the GPU shouldn't enter BACO even the video device is idle.

When the GPU is in BACO, it should exit BACO if audio device starts working

BACO Exit Sequence:

For wake up sequence, the doorbell mechanism is used. The GPU doorbell mechanism that was introduced to the Volcanic Island family provides an application or driver to indicate GPU engine that it requires work on the HW. These doorbells can be issued from the software running on the CPU or on the GPU. The hardware supports doorbell mechanism by implementing a watch I/O memory mechanism that is programmed to recognize when a write happens to a special range of address. For BACO, this presents a new problem since these doorbell accesses cannot be detected by amdgpu driver if they originated from the software running on the CPU. Attempting to access the ASIC while it is in the BACO state will result in a hang. In order to prevent this from happening, there are two major design considerations:

The first is driver will now have to wait for ASIC idle status from SMU before entering BACO. This will prevent the driver from attempting to enter BACO when there is outstanding doorbell access. See above for entry sequence details.

The second is SMU will transfer control of monitoring doorbell activities to BIF when in BACO mode. This will allow BIF to detect any doorbell transaction and initiates an interrupt to the driver to exit BACO.

As BIF detects an incoming configuration cycle, it asserts a GPIO to wake-up the off power rails

and the rest of the dGPU. A PCIe link training is not required after normal BACO exit. The doorbell monitor control is already transferred to BIF before ASIC entered BACO. While ASIC is in BACO, amdgpu driver gets notification of any doorbell activities from BIF via an interrupt. This triggers an event to wake ASIC up from BACO. Rest of the steps are nothing but unwinding steps of entry sequence described above.

MISC:

GPU reset using BACO.

https://github.com/RadeonOpenCompute/ROCK-Kernel-Driver/blob/7b4f1de71ea335e965fba590f4d030de52644137/drivers/sys/kernel/debug/dri/N/amdgpu_gpu_recover can be used to manually trigger a gpu reset at the next fence wait and internally it may use BACO if applicable.

There are challenges with virtualization. Power management features such as BACO with Doorbell are not enabled with PCIe-SRIOV because of BIF ring buffer issues which is used for doorbell. Important functions to look for:

```
bool amdgpu_device_supports_boco(struct drm_device dev);
bool amdgpu_device_supports_baco(struct drm_device dev);
bool amdgpu_device_is_peer_accessible(struct amdgpu_device adev,
struct amdgpu_devicepeer_adev);
int amdgpu_device_baco_enter(struct drm_device dev);
int amdgpu_device_baco_exit(struct drm_device dev);
int smu_baco_get_state(struct smu_context smu, enum smu_baco_state state);
int smu_baco_enter(struct smu_context smu);
int smu_baco_exit(struct smu_context smu);
bool smu_baco_is_supported(struct smu_context smu);
bool amdgpu_dpm_is_baco_supported(struct amdgpu_device adev)
```

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary authors: Mr BHARDWAJ, Rajneesh (AMD); Mr DEUCHER, Alexander (AMD)

Session Classification: Main Track

Track Classification: Talk (half slot) (Closed)

Contribution ID: 23

Type: **not specified**

Buffer constraints

Thursday, 17 September 2020 17:35 (1h 25m)

James Jones from NVIDIA has presented a proposal a few years back to redesign the buffer allocation mechanisms on Linux (especially for GPUs, display devices, etc). However this is a pretty ambitious undertaking because it involves rewriting a big part of the graphics stack.

His proposal included several components. The recent work on modifiers allows to solve the “capability” part of the proposal. For instance last year James gave another talk which uses modifiers for Nouveau tiling.

However some other parts still haven’t been implemented. A common issue is that buffer consumers and producers have no way to agree on buffer constraints like alignment, max pitch, contiguous memory and other placement restrictions. Today, buffer producers need to assume a number of constraints when allocating buffers, some of which may be unnecessary for a given usage.

This workshop aims to discuss about potential solutions to this particular issue. I’ll give an overview of the problem we’re trying to solve, and give a few possible start of solutions to kickstart discussions. A goal of the solution would be to integrate well with the existing ecosystem, extending it rather than replacing it completely.

I’d like to gather feedback from various vendors to make sure potential solutions are sensible, and collect more ideas to bring this forward.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: SER, Simon (sway/wlroots)**Session Classification:** Workshop**Track Classification:** Workshop Discussion Topic

Contribution ID: 24

Type: **not specified**

Profiling on AMD GPUs using tracing

Wednesday, 16 September 2020 16:35 (20 minutes)

In this talk I'd like to show how to go beyond per-draw performance counters by using the thread tracing feature on AMD GPUs. This will include instruction-level shader profiling and high frequency streaming performance counters as well as a look at the impact of barriers and other serializing commands.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: NIEUWENHUIZEN, Bas

Session Classification: Main Track

Track Classification: Talk (half slot) (Closed)

Contribution ID: 25

Type: **not specified**

The Libre-SOC Project

Thursday, 17 September 2020 16:10 (45 minutes)

The Libre-SOC Project aims to bring a DRM-free 3D GPU/CPU/VPU processor to fruition, providing the backbone of guaranteed “right to repair” and beyond. Anyone technically familiar with Apple’s new processor knows the true implications: if Apple controls the entire stack right from boot, then with their market share, vendor lock-in on an unprecedented scale becomes the new reality. With Intel losing the plot (Spectre, Meltdown, QA failures) other vendors will likely follow their example.

If we do not wish to see that happen it is our duty and responsibility to provide alternative processor designs that are targetted at mass-volume products: tablets, smartphones, chromebooks and more.

This then defines the technical requirements:

- The processor must be power-efficient
- It must be capable of good 3D graphics
- It must have audio and video acceleration
- There must be good driver support (BSP)
- The entire stack must be Libre
- The processor must be “unbrickable”

With help from NLNet, under their Privacy and Enhanced Trust Programme, we have received seven separate EUR 50,000 Grants targetted at specific areas to make this a reality, covering:

- The core processor design which is to be an augmented POWER9 compliant design, guided by the OpenPOWER Foundation
- Paying for a 180nm test ASIC to be laid out using entirely libre ASIC tools by Sorbonne University (coriolis2)
- Two separately funded 3D Vulkan Drivers: Kazan and MESA
- Audio/Video assembly-level acceleration for inclusion in ffmpeg and gstreamer low-level libraries
- Support for Development of 3D and Vector Processing Standards and submission to the OpenPOWER Foundation for inclusion in PowerISA
- Documentation and openness to suit educational and business needs alike.
- Formal Correctness Proofs for both the low and high level design (including the IEEE754 units)

This latter is critically important for transparency: the processor has to be independently verifiable, and Mathematical Correctness proofs are a good way to achieve that.

This is a massively ambitious and unprecedented project. It is also based on a technically underappreciated historic design: the CDC 6600. With help from Mitch Alsup, the designer of the Motorola 68000, it has been possible to upgrade the 6600 core to multi-issue and precise exceptions with no architectural compromises.

With so much ground to cover, this talk therefore provides an overview and introduction to the project.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: LEIGHTON, Luke

Session Classification: Main Track

Track Classification: Talk (full slot) (Closed)

Contribution ID: 26

Type: **not specified**

Graphics tracing with Perfetto

Friday, 18 September 2020 15:15 (20 minutes)

In this presentation I will talk about *graphics tracing* and a collection of tools useful for profiling and trace analysis.

I will introduce *gfx-pps*, a project Collabora started working on this year, which provides some components that, in conjunction with Google's *Perfetto*, enable you to capture a trace and visualize GPU performance counters, or any kind of timeline, with a nice web-based UI.

This kind of analysis is crucial to identify bottlenecks on the GPU and to get insights on which area of the graphics application to focus your optimization efforts.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: Mr CAGGIANO, Antonio (Collabora)

Session Classification: Main Track

Track Classification: Talk (half slot) (Closed)

Contribution ID: 27

Type: **not specified**

Quick GL and Vulkan tests with shader_runner and Amber

Thursday, 17 September 2020 18:30 (20 minutes)

Normally, writing a CTS or piglit test requires writing a fair amount of C code. But what if you just want to draw a rectangle using a shader? Fortunately, both test suites come with tools to help you do just that with a minimal amount of fuss. Piglit has `shader_runner` and the Vulkan CTS has Amber, which are scripting languages for their respective graphics APIs. This talk will offer a brief introduction to the capabilities and syntax of both tools.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: GOLDMINTS-ORLOV, Arcady (Igalia)

Session Classification: Main Track

Track Classification: Talk (half slot) (Closed)

Contribution ID: 28

Type: **not specified**

About OpenGL and Vulkan interoperability.

Thursday, 17 September 2020 17:05 (20 minutes)

EXT_external_objects and EXT_external_objects_fd are groups of OpenGL extensions that allow OpenGL and Vulkan interoperability. When enabled, Vulkan allocated resources can be accessed and re-used by OpenGL. This talk is about the implementation of the extensions in various drivers, and some common interoperability use cases and examples that have been added to piglit.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: STEA, Eleni Maria (Igalia)

Session Classification: Main Track

Track Classification: Talk (half slot) (Closed)

Contribution ID: 29

Type: **not specified**

State of the X.org

Friday, 18 September 2020 17:15 (20 minutes)

Your secretary's yearly report on the state of the X.org Foundation. Expect updates on the freedesktop.org, internship and student programs, XDC, and more!

Code of Conduct

Yes

GSoC, EVoC or Outreachy

Primary author: PAUL, Lyude (Red Hat Inc.)

Session Classification: Main Track

Track Classification: Talk (half slot) (Closed)

Contribution ID: 30

Type: **not specified**

DRM-backend tests in Weston's GitLab CI

Friday, 18 September 2020 15:45 (20 minutes)

A few years ago the first steps to implement a virtual KMS device were taken, and so VKMS was born. This virtual device is very useful to run DRM-backend tests in headless machines, and so it can be used to extend CI's tests coverage.

In Weston we are already using it to automatically run some very simple DRM-backend tests in its GitLab CI, and that's what we are going to show in this talk. Also, the work in progress of both Weston and VKMS in order to increase the testing capability is going to be discussed.

Leandro Ribeiro is a Brazilian software engineer that works as an intern in Collabora's Graphics Team. Recently he's been contributing to Wayland/Weston, a project that he believes plays a fundamental role for the future of FOSS.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: RIBEIRO, Leandro

Session Classification: Main Track

Track Classification: Talk (half slot) (Closed)

Contribution ID: 31

Type: **not specified**

From witchcraft to production

Wednesday, 16 September 2020 20:25 (5 minutes)

Young students of witchcraft believe understanding magic is an end in itself, that successful reverse-engineering is the pinnacle of a mage's journey. In teenage naïveté, they believe breaking the hex is the hardest challenge they could ever face. Yet they are unprepared for the eldritch abomination waiting on the other side: driver development.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: ROSENZWEIG, Alyssa (Panfrost (Collabora))

Session Classification: Demos / Lightning talks I

Track Classification: Lightning Talk

Contribution ID: 32

Type: **not specified**

Opening session

Wednesday, 16 September 2020 13:00 (10 minutes)

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: SZWICHTENBERG, Radosław (Intel)

Session Classification: Main Track

Track Classification: Talk (half slot) (Closed)

Contribution ID: 33

Type: **not specified**

Opening session

Thursday, 17 September 2020 13:00 (10 minutes)

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: SZWICHTENBERG, Radosław (Intel)

Session Classification: Main Track

Track Classification: Talk (half slot) (Closed)

Contribution ID: 34

Type: **not specified**

Opening session

Friday, 18 September 2020 13:00 (10 minutes)

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: SZWICHTENBERG, Radosław (Intel)

Session Classification: Main Track

Track Classification: Talk (half slot) (Closed)

Contribution ID: 35

Type: **not specified**

Closing session

Friday, 18 September 2020 18:05 (20 minutes)

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: SZWICHTENBERG, Radosław (Intel)

Session Classification: Main Track

Track Classification: Talk (half slot) (Closed)

Contribution ID: 36

Type: **not specified**

X.Org Foundation Board of Directors Meeting

Thursday, 17 September 2020 19:30 (1 hour)

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: IGLESIAS GONSÁLVIZ, Samuel (Igalia)

Session Classification: Main Track

Track Classification: Talk (full slot) (Closed)

Contribution ID: **38**Type: **not specified**

Vulkan Presentation Timing Extension

Wednesday, 16 September 2020 20:20 (5 minutes)

The Khronos Vulkan working group, recognizing substantial research and development efforts in the open source graphics community in the area, has agreed to make development of the work-in-progress Vulkan presentation timing extension public. This talk will give a very, very brief overview of the current spec and point attendees at the github home for development of the specification.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: JONES, James (NVIDIA)

Session Classification: Demos / Lightning talks I

Track Classification: Lightning Talk

Contribution ID: 39

Type: **not specified**

Buffer constraints workshop summary

Friday, 18 September 2020 17:45 (5 minutes)

Summary for the buffer constraints workshop.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary authors: SER, Simon (SourceHut); JONES, James (NVIDIA)

Session Classification: Lightning talks II

Track Classification: Lightning Talk

Contribution ID: 40

Type: **not specified**

Gamescope update and demos

Wednesday, 16 September 2020 20:05 (10 minutes)

Gamescope is an overhaul of steamcompmgr, the GLX compositing window manager for SteamOS. In the move to Wayland and Vulkan, it gained some interesting properties that make it sometimes useful on a normal desktop, which I'll quickly demo.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: GRIFFAIS, Pierre-Loup (Valve Software)

Session Classification: Demos / Lightning talks I

Track Classification: Demo

Contribution ID: 41

Type: **not specified**

Next steps for wroots

Friday, 18 September 2020 17:50 (5 minutes)

wroots is a modular Wayland compositor library. This talk will give a sneak peak at the recent and upcoming architectural changes for the wroots rendering and display pipeline to support features such as hardware planes, new renderers and explicit synchronization.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: SER, Simon (SourceHut)

Session Classification: Lightning talks II

Track Classification: Lightning Talk

Contribution ID: 42

Type: **not specified**

Universal display management with Disman

Friday, 18 September 2020 17:55 (5 minutes)

Quick introduction to Disman, a front end library, D-Bus service and command line utility to manage displays with a multitude of Wayland compositors and in an X11 session.

Also KDisplay is shown as a front end example for Disman.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: GILG, Roman

Session Classification: Lightning talks II

Track Classification: Lightning Talk

Contribution ID: 43

Type: **not specified**

Update on AMD DRM modifiers

Friday, 18 September 2020 18:00 (5 minutes)

I'll give an update on DRM format modifiers for AMD hardware. Why we want it, what the challenges are and what I think will happen in the future.

Code of Conduct

Yes

GSoC, EVoC or Outreachy

No

Primary author: NIEUWENHUIZEN, Bas (Google)

Session Classification: Lightning talks II

Track Classification: Lightning Talk