

# Improvements to the Raspberry Pi GPU driver stack

Ella Stanforth, Juan A. Suárez

XDC 2025 - Vienna

# Who are we?

- We are open-source developers at Igalia working at the Graphics Team.
- We focus on enhancing the Raspberry Pi graphics stack, mainly on the Mesa user-space driver.
- We are representing the work done not only by us, but from other members of the Graphics Team working improving also the Raspberry Pi driver.

# Contents

- New Features (Ella)
- Performance Improvements (Juan)
- Q&A

# New extensions exposed (since 24.2 branchpoint)

- GL\_EXT\_conservative\_depth
- GL\_EXT\_disjoint\_timer\_query
- GL\_EXT\_multi\_draw\_indirect
- GL\_EXT\_sRGB
- **GL\_EXT\_shader\_framebuffer\_fetch**
- GL\_EXT\_shader\_framebuffer\_fetch\_non\_coherent
- GL\_KHR\_blend\_equation\_advanced
- GL\_KHR\_blend\_equation\_advanced\_coherent
- **GL\_{ARB,EXT}\_blend\_func\_extended**
- GL\_{ARB,EXT}\_timer\_query
- GL\_{ARB,KHR}\_robust\_buffer\_access\_behavior

# Framebuffer Fetch

- We implemented this by lowering `load_output` to a V3D specific TLB read intrinsic.
- Gallium provides `GL_KHR_blend_equation_advanced` for US.

# Dual Source Blend

- The hardware does not support dual source blend factors.
- We used the common `nir_lower_blend` pass as a software fallback.



# 16 Bit Normalisation

- 16 bit normalised formats are required by OpenGL 3.2 but not supported by the hardware.
- We added a NIR lowering pass to convert to and from integer formats.
- We use the same software fallback for blending.

# NIR printf

- Very helpful debugging feature.
- We added a new helper to reduce noise.

```
nir_printf_fmt_at_px(b, size, x, y, fmt, ...)
```



# Robustness

- Added so CTS can detect resets.
- We added reset counters to the kernel driver to support this.

# OpenGL 3.2 - Pending issues

- The driver is missing support for non seamless cubemaps:
  - The hardware only supports seamless cubemaps.
  - Zink has a lowering pass but some piglit tests fail.

# Performance improvements

- For last year, we focused on performance improvements on GPU limited scenarios using Full-HD target resolution.
- We have analyzed the performance of V3D using several GLES gfxbench traces, and we have achieved an average of ~**32.29%** FPS improvement in these scenarios since the last XDC, with cases from ~**4.92%** to ~**68.32%** FPS improvement.
- All these performance optimizations are available in stable Mesa 25.2.3

# Avoid load/stores on invalidated framebuffers

- With the information of the invalidated framebuffers we can avoid the stores of the results of tile buffer rendering and the next load if they re-used in following jobs as any read value would be undefined.
- This gets us an **+1.11%** FPS average improvement (between ~0.05% and 3.26%).

*c1: "v3d: avoid load/store of tile buffer on invalidated framebuffer"*

# Take advantage of Early-Z optimization

- Early-Z optimization was disabled when there is a discard instruction in the draw call shader. But we can enable it at draw time if depth updates are disabled and there are no occlusion queries active.
- This gets us an **+13,21%** FPS average improvement (between 4.78% and 17.69%).

*c2: "v3d: Enable Early-Z with discards when depth updates are disabled"*

# Avoid loads/stores with disabled rasterization


- If all draw calls submitted have the rasterizer discard enabled, we can avoid any tile buffer load/stores.
- This is specially helpful in scenarios where transform feedback is used, because the application is only interested in the geometry results.
- This gets us an **+13.32%** FPS average improvement (between 0% and 41.32%)

*c3: "v3d: Don't load/store if rasterizer discard is enabled"*

# Avoid supertile coordination submission without rasterization

- This is another improvement when rasterizer discard is enabled.
- Even avoiding tile buffer load/stores, GPU still executes the FS in all tiles belonging to a supertile.
- This basically emits no supertile to avoid that.
- This gets us an **+1.16%** FPS average improvement (between 0% and 3.77%).

# FPS over time per benchmark





# Q&A

Join us!

<https://www.igalia.com/jobs>



