

# Implementing Vulkan Timeline Semaphores

Jason Ekstrand, XDC 2018

# Option 1: Kernel Magic

In theory, kernel GPU schedulers could hold work until all dependent fences have materialized. This comes with a lot of problems:

- When do you commit resources?
  - Exec call from userspace or submission to the GPU?
- How do you handle over-committed situations?
- How do you handle wait-before-signal without risking deadlock?
- Who's going to fix all the TTM fallout?
- How do you avoid leaking "future" fences through GEM implicit sync logic?

# Option 2: Rewrite the DRM Subsystem

Other platforms have pushed resource management into userspace


- Resource management and dependency tracking are separated.
- Userspace drivers request resources to be committed before they ever submit the work to use them.
- If the kernel can't make progress (due to over-commit or dead-lock), it starts kicking out clients and canceling their work.
- Userspace is expected to work with the kernel to reduce resources, re-commit and re-submit work.

So, yeah, we could rewrite the entire DRM subsystem...

# Option 3: Userspace Queueing (Current Intel and AMD Proposal)

Add a timeline concept to DRM sync objects:

- Fences are assigned to the syncobj with a u64 serial
- Fences are requested from a syncobj with for a u64 serial
- If the driver requests a fence for a time point which is greater than the maximum currently available, -ENOENT is returned.
- Add a new ioctl to query the highest time point completed or available.
- Extend DRM\_SYNCOBJ\_WAIT to allow waiting for a given time point to complete or to become available.
- Userspace drivers will hold work until all dependent time points are available.



Answers?

# Conclusions from Thursday's break-out

- We want to allow speculative submits from userspace
  - Userspace doesn't have to always wait before submitting
  - We should have some better return value than `-EINVAL`
- Need a `WAIT_FOR_SUBMIT` flag
- Want to support one thread for multiple queues
  - Probably don't need `WAIT_FOR_ANY_OF_ALL_OF_GROUP` flag
  - Maybe add a way to return all finished syncobjs with `WAIT_ANY`? Probably not.
- We all agree lol
- Should we use timeline semaphores for explicit sync in X11 and Wayland?