

DRM GPU Scheduler Workshop Summary

- Life time hack in `drm_sched_fini()`: Community agreed that we'll try to remove it. Shouldn't explode on amdgpu anymore.
- `drm_sched_resubmit_jobs()` successor:
 - We will document how amdgpu and Xe solved job resubmissions, for hypothetical new users
 - Other users will likely live on `drm_jobqueue` and / or won't need resubmissions
 - Matthew Brost is going to propose an iterator for obtaining unbroken jobs' `dma_fences` from `drm_sched`
- Philipp kindly asked the community to never-ever access API internals anymore.
- Missing locking for runqueues:

Locks might have been omitted on purpose for performance reasons years ago. Wasn't documented. We don't know whether locks should / could be added. "That it doesn't explode all the time seems to prove that it works ._."
- `cancel_job()` callback: Clarified to the stakeholders that it's no longer necessary that jobs refcount the scheduler to prevent lifetime issues.
- Completely Fair Scheduler: No objections. Green light for RFC to transition to v1 \o/
- Scheduling CGroup RFC: Scheduling far too fragmented for us to provide a unified uAPI.