

A Step Towards Explicit Sync

Bas Nieuwenhuizen, XDC 2022

dma_fence

- Thing that signals the completion of some work
 - Can be waited upon, signalled etc.

Userspace API:

- Syncobj
 - Container of zero or one dma_fence which can be replaced/reset
- Sync fd
 - File descriptor pointing to a dma_fence

How did implicit sync work?

- Each buffer has a `dma_resv`
- `dma_resv`:
 - Contains up to 1 “exclusive” fence
 - Contains up to many “shared” fences

On submit:

- Pass a list of referenced buffers
- Wait on fences from `dma_resvs`.
- Put new fence in `dma_resvs`.

Iteration 1: Ignoring implicit sync in a process

The hardware has different queues that can run in parallel, e.g.

- compute
- graphics
- DMA

So let everything take only a shared reservation, and then let the driver determine things using a fence owner (=VM):

- Fences with the same fence owner are ignored

Problem: VR compositor

- Needed cross-process Vulkan \leftrightarrow Vulkan sharing
- And implicit sync interfered with some of the high priority compute

Iteration 2: excluding buffers from implicit synchronization

- Mark a buffer as "EXPLICIT_SYNC"
- All submissions only wait on memory fences in the reservation
- They still add the usual fences

Alternative Iteration 2: Don't touch the dma_resv

- Some driver just let you not touch the dma_resv
- Requires some aggressive idling if you want to move buffers

Implicit sync remainder in radv+amdgpu

- WSI buffers
 - Including Vulkan compositors importing buffers
- VA map/unmap
 - Have a shared dma_resv for non-inter-process buffers making this hard

Have option to remove now.

Extended dma_resv

More categories for fences:

1. DMA_RESV_USAGE_KERNEL
2. DMA_RESV_USAGE_WRITE
3. DMA_RESV_USAGE_READ
4. DMA_RESV_USAGE_BOOKKEEP

Disable implicit sync everywhere!

- Only wait for KERNEL fences and set BOOKKEEP fences
 - Add a flag to submission (or context/queue/engine)
- Can stop marking buffers EXPLICIT_SYNC.
- Except did we just break most WSI?

Transferring fences to/from dma_resv

New IOCTLs:

- DMA_BUF_IOCTL_EXPORT_SYNC_FILE
- DMA_BUF_IOCTL_IMPORT_SYNC_FILE

Maintain buffer associated dependencies explicitly.

Thank you to

- Christian König
- Jason Ekstrand
- Daniel Vetter