

Bulk moving mechanism on LRU for DRM/TTM

Thursday, 3 October 2019 14:55 (20 minutes)

While investigating a performance issue with the F1 2017 game benchmark, we identified some bottlenecks related to how ttm and amdgpu do buffer validation and LRU handling. This ultimately led to a major redesign of how we handle buffer migration. This talk describes the process that we took to identify and fix the bottleneck and what we learned along the way.

The Talos Principle(Vulkan) Clpeak(OCL) BusSpeedReadback(OCL) /unit: ms

Original 162.1 FPS 42.15 us 0.254 (1K) 0.241 (2K) 0.230(4K) 0.223(8K) 0.204(16K)

Bulk Move 162.4 FPS 44.48 us 0.260 (1K) 0.274 (2K) 0.249(4K) 0.243(8K) 0.228(16K)

Original (move PT bo on LRU) 147.7 FPS 76.86 us 0.319(1k) 0.314 (2K) 0.308(4K) 0.307(8K) 0.310(16K)

Bulk Move (move PT bo on LRU) 163.5 FPS 40.52 us 0.244(1K) 0.252(2K) 0.213(4K) 0.214(8K) 0.225(16K) <--

With the best performance and highest FPS at the same time

Reference:

https://www.phoronix.com/scan.php?page=news_item&px=AMDGPU-LRU-Bulk-Move

<https://lists.freedesktop.org/archives/amd-gfx/2018-August/025014.html>

Code of Conduct

Yes

GSoC, EVoC or Outreachy

Presenter: HUANG, Ray (AMD GPU driver)

Session Classification: Main Track

Track Classification: Talk (half slot) (closed)