GPU Virtualization in Hardware

October 10th, 2018 | using Renesas R-Car H3 as example

Stephen Lawrence

Principal Engineer, Renesas
BIT Lead, Genivi

Hardware GPU Virtualisation

- Various different forms. Example R-Car H3:
- Each OS is presented with its own logical GPU. Typically with its own gfx driver.
- Multiple command input ports give each OS dedicated input
  - OS 2 can not swamp/spam input port of OS 1
- OS ID
  - Allows separation of data and priority setting of data through IP
- Priority-based high speed task switching
  - Draw tasks have TA (polygon processing) front end and 3D back end
  - Time slices between high, e.g. cluster, and low, e.g. map, priority tasks
    - OS 2 can’t stall GPU with huge vertices or shaders in way that effects OS 1. GPU cluster can be restarted if needed.
    - Ensure 60 fps cluster. Even if low priority GPU restarted, cluster is unaffected.
  - Controllable pre-emption deadline
- IPMMU with OS ID
  - Memory protection between OSs
  - OS ID 2 can not access physical memory of OS ID 1

Thank you!

Visit GENIVI at http://www.genivi.org or http://projects.genivi.org

Contact us: help@genivi.org