FireSim on Xilinx U250 and Other Custom Host Platforms

David Christoph Metz & Magnus Själander
Department of Computer Science
Norwegian University of Science and Technology

Why?

- Getting started with FireSim is cheap and easy
- Running long simulations can get expensive
  - Need to be careful not to forget experiments
- Lockin

FireSim Interfaces

MMIO

- Used by `simif_t::read()` and `simif_t::write()`

DMA

- Used by `simif_t::push()` and `simif_t::pull()`

RAM

- Primarily used by FASED for main memory simulation

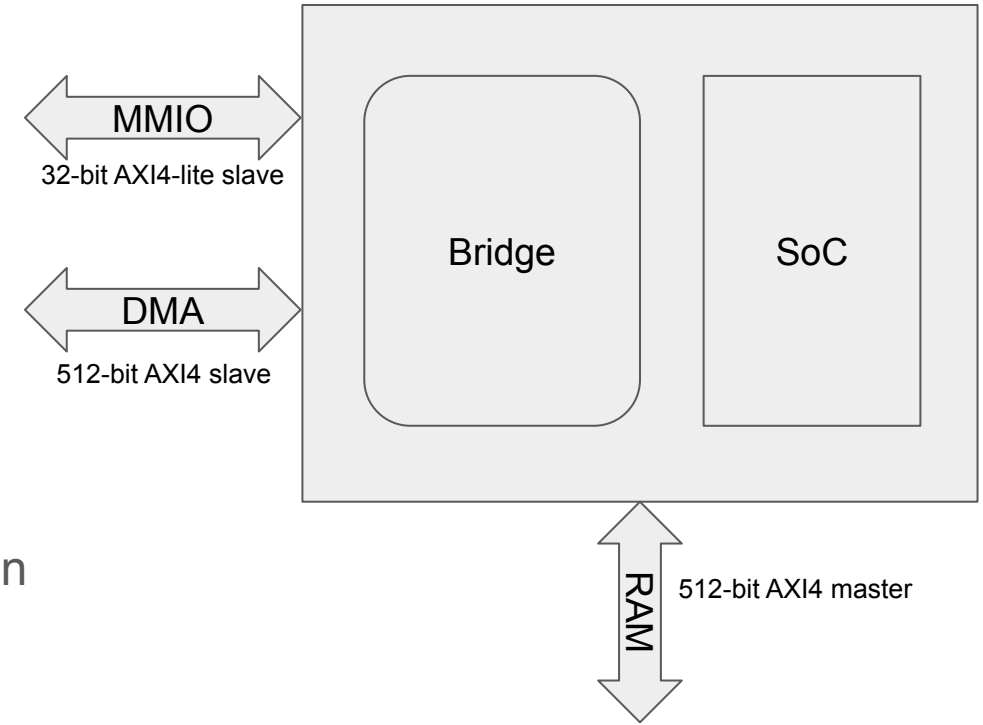

Alveo U250

- Datacenter FPGA
- Similar to AWS F1
  - 4 instead of 3 SLRs
- 1,728K LUTs
- PCIe Gen3x16
- 4x16GB DDR4

Vitis/XRT

- Not available in FireSim when we started this
- Static shell with PCIe, management and partial reconfiguration support
- Reconfiguration of user region via PCIe
- Mainly intended for XRT managed HLS
  - Support for user managed kernels relatively new
- Only supports DMA to memory, not as used by FireSim
  - No TracerV, SimpleNIC, Dromajo, Print or TraceDoctor

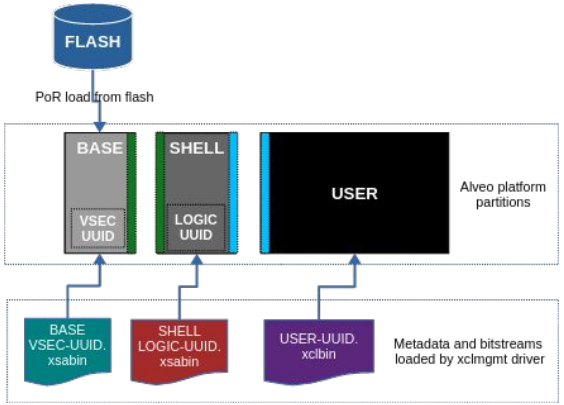

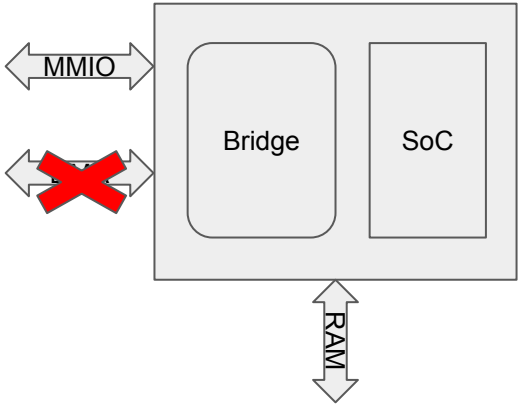

AWS Shell

- More flexible for user logic than XRT
  - Provides DMA
- Marked interfaces used by FireSim
- One RAM channel part of shell
- Several interfaces unused

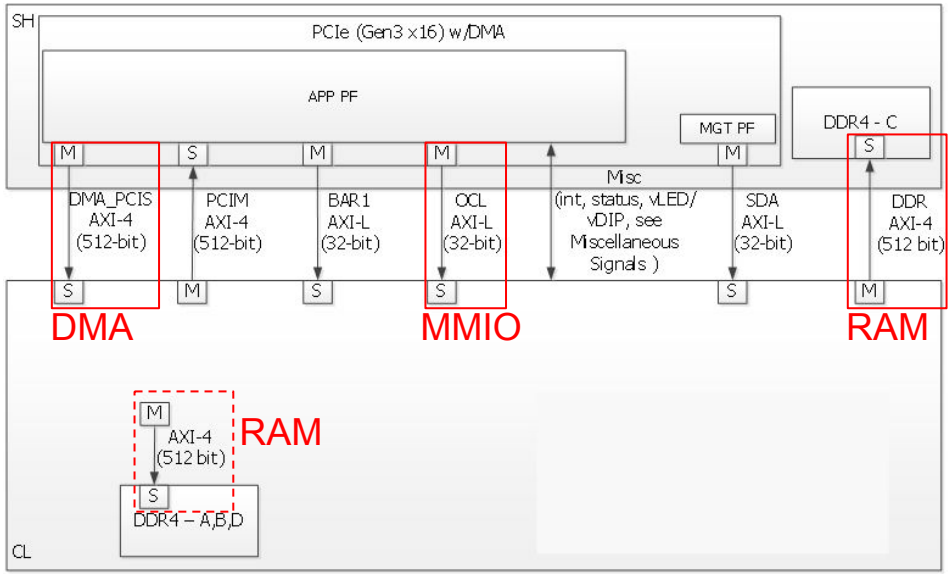

Directly creating top level design

- Use same IPs as the XRT and AWS shells
- Generate bitstream for whole FPGA

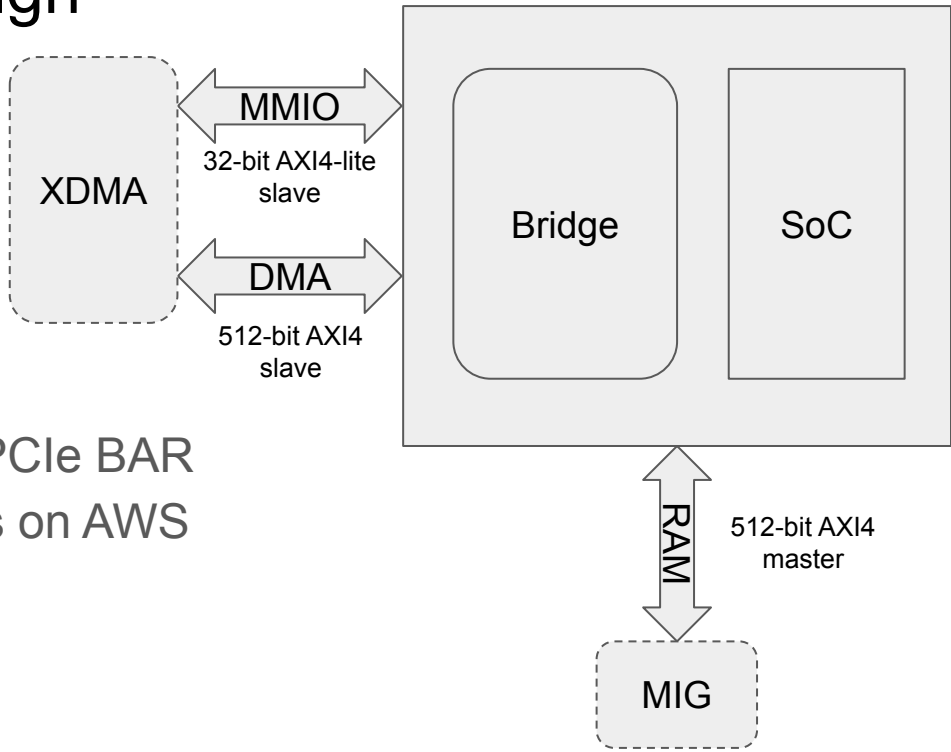

- MMIO is used by directly accessing PCIe BAR
- DMA using XDMA works the same as on AWS

Benefits of FireSim on U250

- >30% larger FPGA than on F1
- Slimmer shell
- High degree of control over implementation possible
  - High frequencies achievable with one memory controller
    - Avoids SLR crossings
  - Possible to for example adjust DMA config further

PCIe Quirks

- PCIe is not hot-swappable
  - Address space is assigned by BIOS at boot time
  - Switching between XRT shell and FireSim requires reboot
- Need to disable PCIe error reporting while flashing bitstream
  - Script for this provided by AMD/Xilinx OpenNIC
  - PCIe link goes down
  - Allows switching between bitstreams with same PCIe configs

USB JTAG on a server FPGA…

- Required to flash complete bitstream
- Only bitstream changed
  - Power-cycling reverts to XRT shell
  - Possible to flash ROM instead

Multi-user/Cluster

- Work largely done by Björn Gottschall
- Simplify flashing bitstreams
  - Single command
- Control allocation of FPGAs to users
  - Acquire and release controlled using file locks
- Don’t want to give sudo to users
  - Required for flashing bitstreams and accessing XDMA/PCIe BAR
  - Required to for allowing access to PCIe BAR and XDMA
  - Enabled by allowing execution of script in sudoers file

Potential Upgrades

- Create real shell
  - Use partial reconfiguration
  - No more PCIe issues
  - Switching to/from XRT still complicated
- Try to match PCIe behavior to XRT shell
  - Would potentially allow switching without rebooting
- Future XRT versions?
  - Would enable full-fledged Firesim on University clusters using XRT
- Upstream?

Future Platforms?

Kria KV260

- Zynq; 256K logic cells; 4GiB RAM; 250$
- Too small for large BOOM - but fine for Rocket SoCs
  - Entry level platform for students
  - Benefit from FASED
- Need cross compilation to ARM for host code
  - Some libraries currently pre-compiled
- Use Vitis?

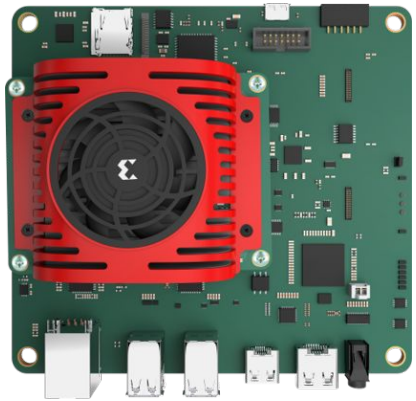

Questions?