

Heterogeneous Multi-Processing for Software-Defined Multi-Tiered Storage Architectures

Xilinx, Fidus Systems and MLE have partnered to address the growing needs in High-Performance Computing and Data Centers to explore “unconventional” data-flow oriented FPGA-based system architecture for acceleration, hyperconvergence, object storage and in-memory compute. The outcome of this is ZU19SN - a high-capacity, hyperconverged, networked storage node with a Zynq UltraScale+ ZU19EG MPSoC.

Contributions

- Xilinx Labs: Streaming Deep-Learning, Key-Value-Store Accelerator.
- Fidus Systems: Sidewinder-100 HPC Accelerator Board / Kit with loopback.
- MLE: Storage & Network Protocol Accelerators, System-Level Integration, Board Support Package, Design Services.

Applications

- Accelerating Memcached servers in OLTP data center applications.
- Object storage for hyper-converged storage nodes.
- Hybrid SSD/HDD Key-Value Drives.

Key Features

- Xilinx ZU19EP with dual NVMe m.2 SSDs and QSFP28 for dual 10/25/50/100 GigE.
- Quad-Core ARM A53 w/ Xilinx PetaLinux.
- PS and PL-attached DDR4LP RAM
- Fully networked with 10/25/50/100 GigE connectivity.
- Integrated System-on-Chip solution for Zynq Ultrascale+, or as PCIe-connected companion FPGA.
- Modular implementation in HDL and C/C++ for Vivado HLS. Supports Xilinx HLx and SDx design flows

Fidus' Sidewinder-100 Xilinx Zynq Ultrascale+ Evaluation System

- Xilinx ZU19EP w/ quad ARM A53 and dual ARM R5.



- QSFP28 for dual 10/25/50/100 GigE
- PCIe Gen3 x16 or Gen4 x8 system i/f
- PCIe Gen3 x8 Host i/f
- 2x PCIe Gen3/4 NVMe m.2 SSDs
- 2x NGFF-8643 i/f for NVMe/SATA/SAS
- 2x 16GB SoDIMM w/ ECC, PS- and PL-attached
- Micro SD-Card, JTAG, UART, I2C, GPIO, ...

Contact Info

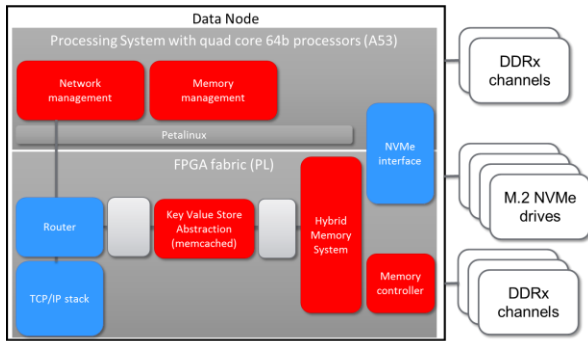
MLE US (San Jose, CA, US):
+1 (408) 475-1490
MLE Europe (Net-Ulm, GER):
+49 (731) 141149-0
www.MLEcorp.com



Heterogeneous Multi-Processing for Software-Defined Multi-Tiered Storage Architectures

Exemplary SW-Defined Storage Architecture

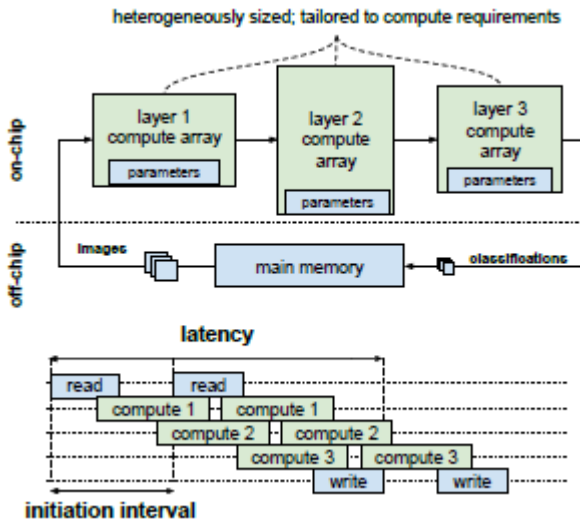
Built on “unconventional” FPGA-based dataflow processing architecture (a.k.a. stream processing) close to network line rates.



Deep-Learning as In-line Processing

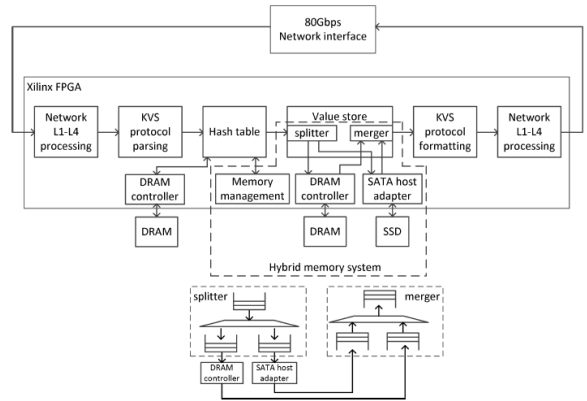
Reduced-Precision (INT8 or Binarized) CNN for on-the-fly data classification:

- Scalable to >5 TOPS
- Very low-power



[Umuroglu et al.: “FINN: A Framework for Fast, Scalable Binarized Neural Network Inference”, FPGA 2017]

Exemplary Memcached Accelerator Architecture

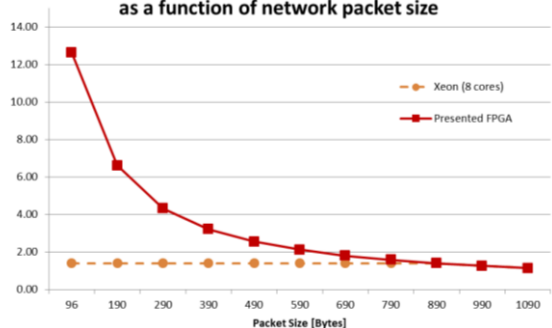


High Performance at Low Power

Many million responses per second (RPS) at over 200k RPS per Watt:

- 13M RPS at 35 Watts board-level, measured at 10 GigE
- 100M RPS, extrapolated for 100 GigE.

Performance for GET operations as a function of network packet size



[Blott et al.: “Scaling out to a Single-Node 80Gbps Memcached Server with 40Terabytes of Memory”, Hot Storage 2015]