THE UNIVERSITY of TULSA

Building an Exotic HPC Ecosystem at The University of Tulsa

John Hale

Peter Hawrylak

Andrew Kongs

Tandy School of Computer Science

Changing Our CS Culture

- Platforms
  - Desktop, workstations, mobile
- Programming
  - Java, python
  - Serial

Changing Our CS Culture

- Platforms
  - Desktop, workstations, mobile += Big iron and HPC
- Programming
  - Java, python += C, etc.
  - Serial += parallel and multi-threaded

Changing Our CS Culture

- Platforms
  - Desktop, workstations, mobile += Big iron and HPC
- Programming
  - Java, python += C and HPC toolchains/APIs.
  - Serial += parallel and multi-threaded

Friends in HPC Places

- OU – Oklahoma Supercomputing Center for Education and Research (OSCER)
  - Wisdom, guidance and models
- OneOklahoma Cyberinfrastructure Initiative (ONEOCII)
  - Community and relationships
- Tandy SuperComputing Center
  - Resources, access and expertise
- Researchers at TU and OU-Tulsa
  - Problem spaces and commiseration

TU HPC Ecosystem Overview

- Facilities
  - Anvi1, Hamm3r, Furn4ce
- Classes
  - High Performance Computing
  - HPC Competition
- Research
  - Security analytics, Bioinformatics, Geosciences
- Philosophy
  - Leverage resources; add value; be different

Bootstrapping

- Re: Friends in HPC Places
  - Sage advice
  - Generous donations of equipment and service
- Instrumentation grants
  - DoD DURIP
  - NSF MRI
- Finding a niche
  - Be different; add value...Heterogeneous computing

Heterogeneous Computing

- Definition: systems using more than one kind of processor, exploiting specialized capabilities in problem solving
- Common heterogeneous compute elements
  - CPU, GPU, MiC, FPGA
- Problem Domains
  - Systems biology, computer vision, image processing

Our First Cluster - Anvil

- Two Alums of our research group went to work at Wikipedia
- Wikimedia/Wikipedia was retiring a large number of servers and giving them to another non profit made the paperwork easy
- We thought we were going to get about five servers, we ended up with forty-five.
- I drove a U-Haul full of servers from Tampa, Florida back to Tulsa, OK
- We decided to build a cluster with the majority of the machines


Tandy School of Computer Science


Tandy School of Computer Science

Anvil – Duct Tape and Bailing Wire

- 32 compute node CPU-only Linux cluster using Dell PowerEdge 1950 III servers
- Dual Xeon L5420 (Quad-Core each), 2.5GHz “Harpertown” CPUs
- 32GB (Upgraded from 8GB) of Fully Buffered DDR2
- Each machine has three Ethernet networks – Compute/Management, Storage, and Lights out management
- Cluster-wide shared storage (3TB), VMware machine for management and services, Dedicated Login and Head node

Software

- Nodes run CentOS 6 for RHEL compatibility (for “science” software)
- Runs SLURM scheduler, single job queue
- Simple shared storage using NFS from NAS
- Machines are configured using puppet and pxe boot using razor
- Paths and environment variables are managed with lmod (lua)

New Machines - HAMM3R and FURN4CE

- 12/16 Nodes, each with Dual Xeon E5-2630 v3 “Haswell” CPUs @ 2.4GHz
- Each node has 64GB of RAM, 240GB SSD, 56Gb/s FDR Infiniband
- HAMM3R -Dual Xeon Phi 31S1P w/ 57 cores @ 1.1GHz, 8GB RAM
- FURN4CE – Single AMD FirePro S9050 GPU, 12GB RAM, 1792 SPs
- Each node will have a Nallatech OpenCL accelerator with either Altera Stratix-V (HAMM3R) FPGA or Arria-10 FPGA (FURN4CE)
- Built using SuperMicro servers assembled by Colfax

HAMM3R Node

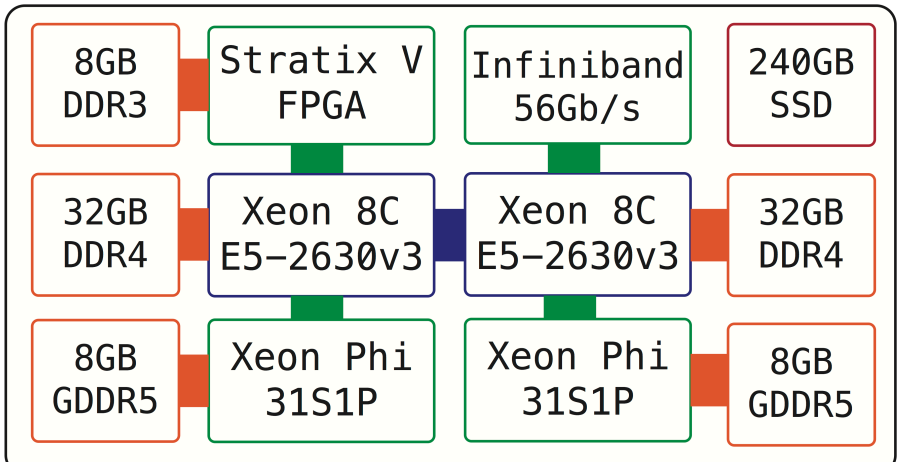

FURN4CE Node

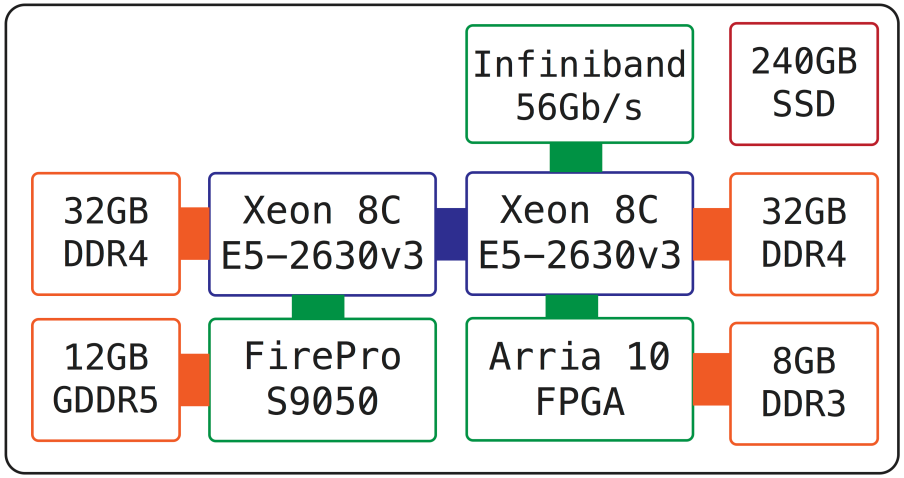

Command and Control Interface

- Called – “Blacksmith”
- Web interface for monitoring and managing the cluster
- Tightly integrated with SLURM, Puppet, and Razor
- Planned to visual and and simplify the using the machine greatly
- VAPORWARE

Computing Environment

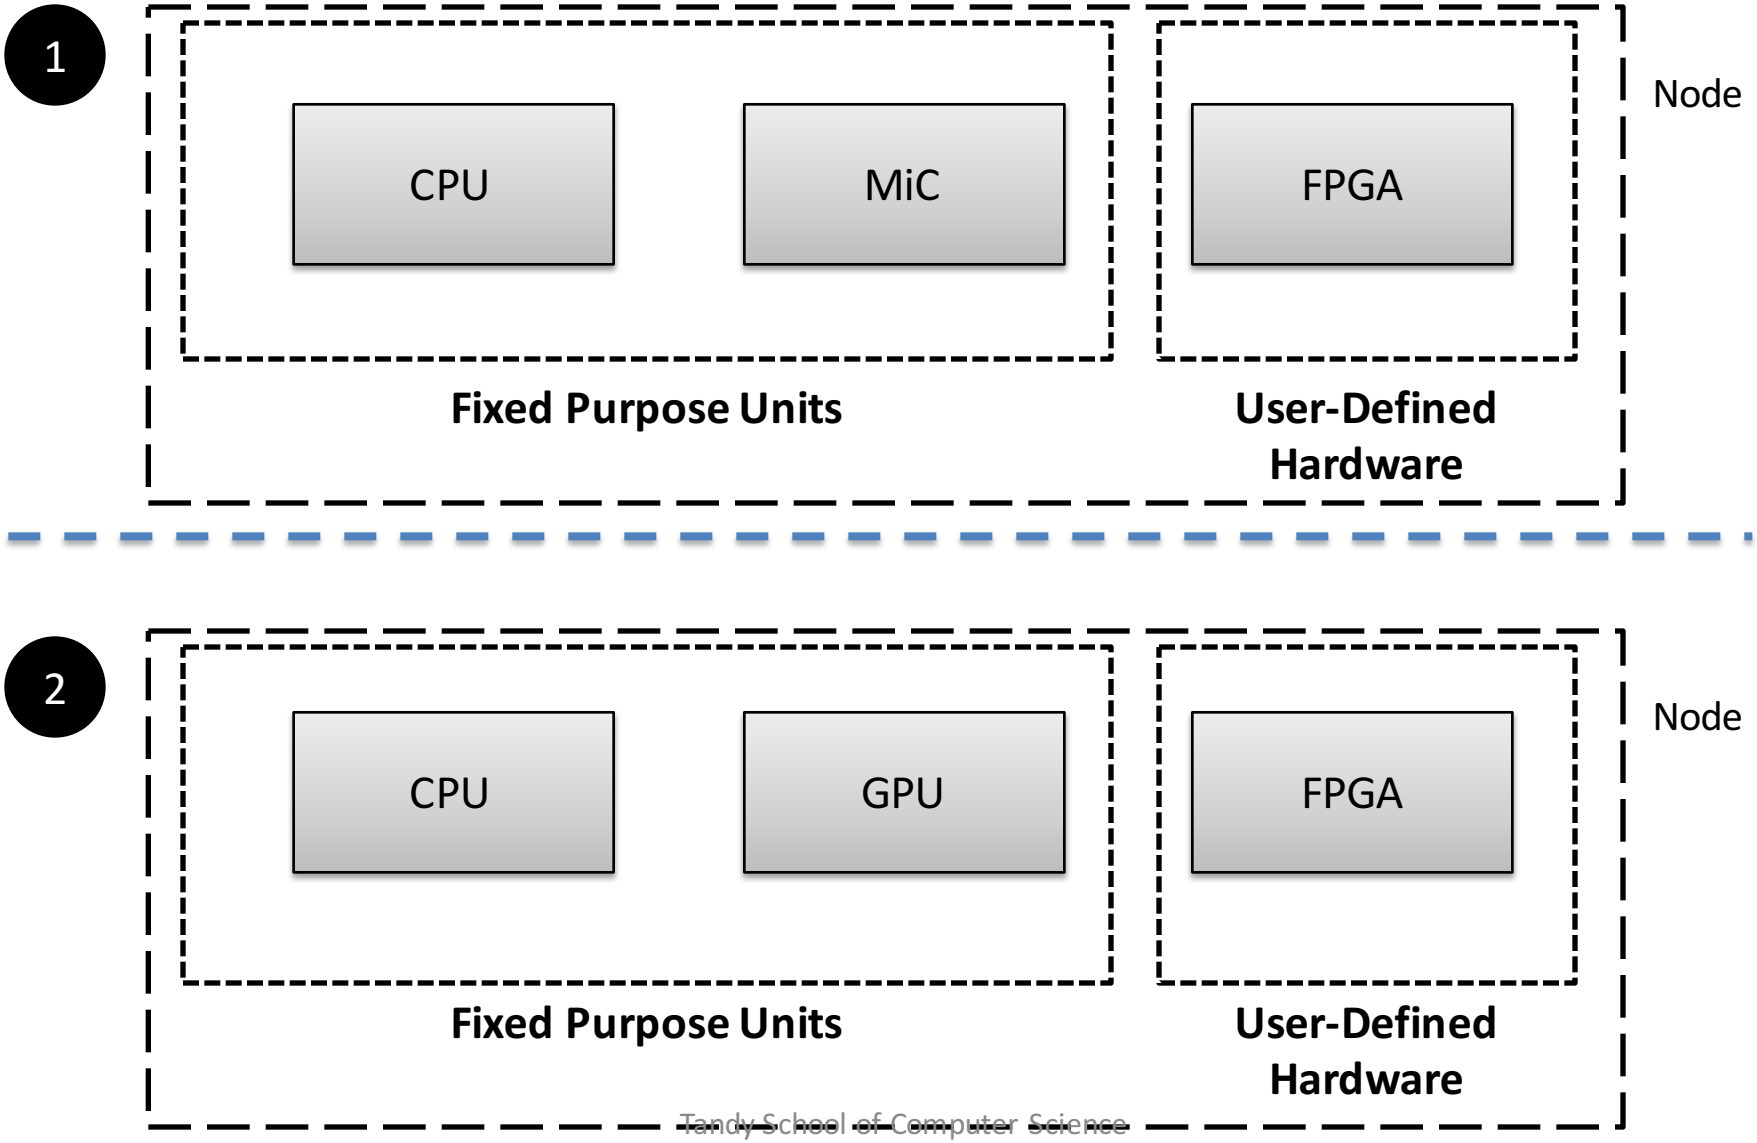

FPGA Benefits

- Customizable
  - "Blank Slate" for hardware
  - Custom pipelines
  - New architectures
- Reconfigurable between jobs
  - Optimally balance resources
- Soft-Core Processor
  - More control over internal components

FPGA Hardware

- Altera Stratix V
  - General purpose logic
- Customizable
- Connectivity
  - InfiniBand
  - Ethernet

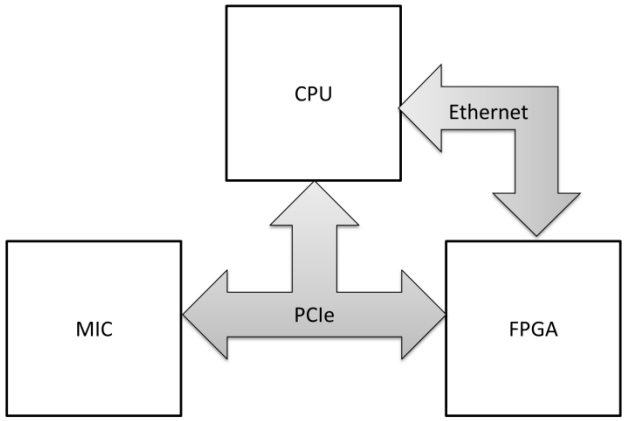

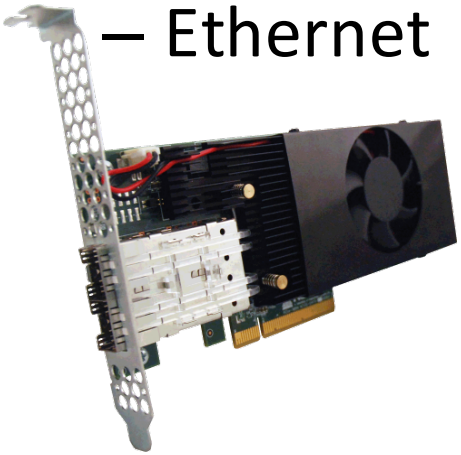

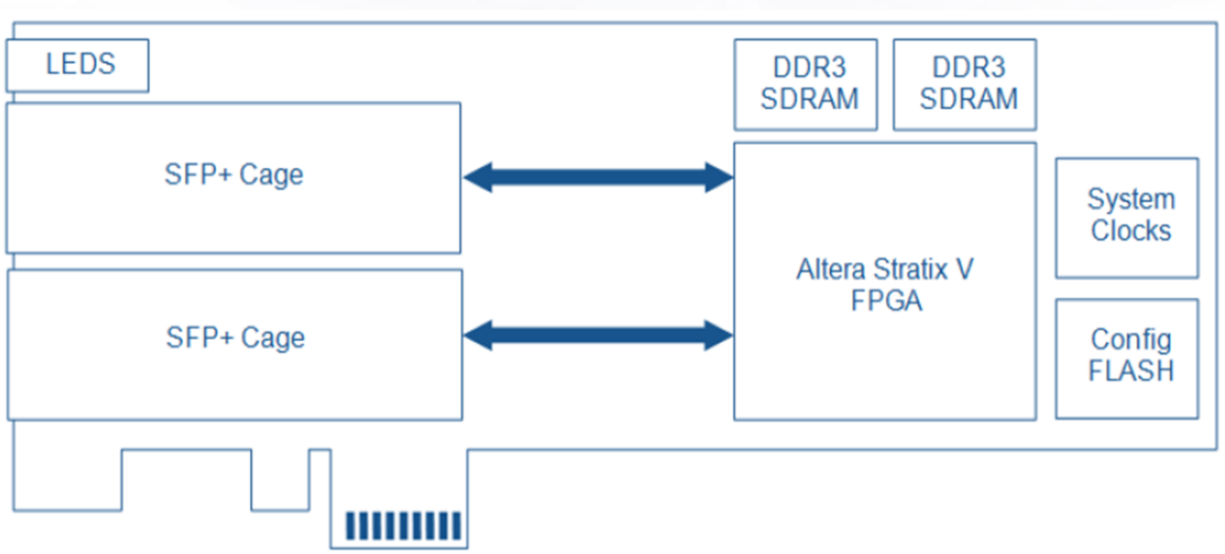

OpenCL Toolchain

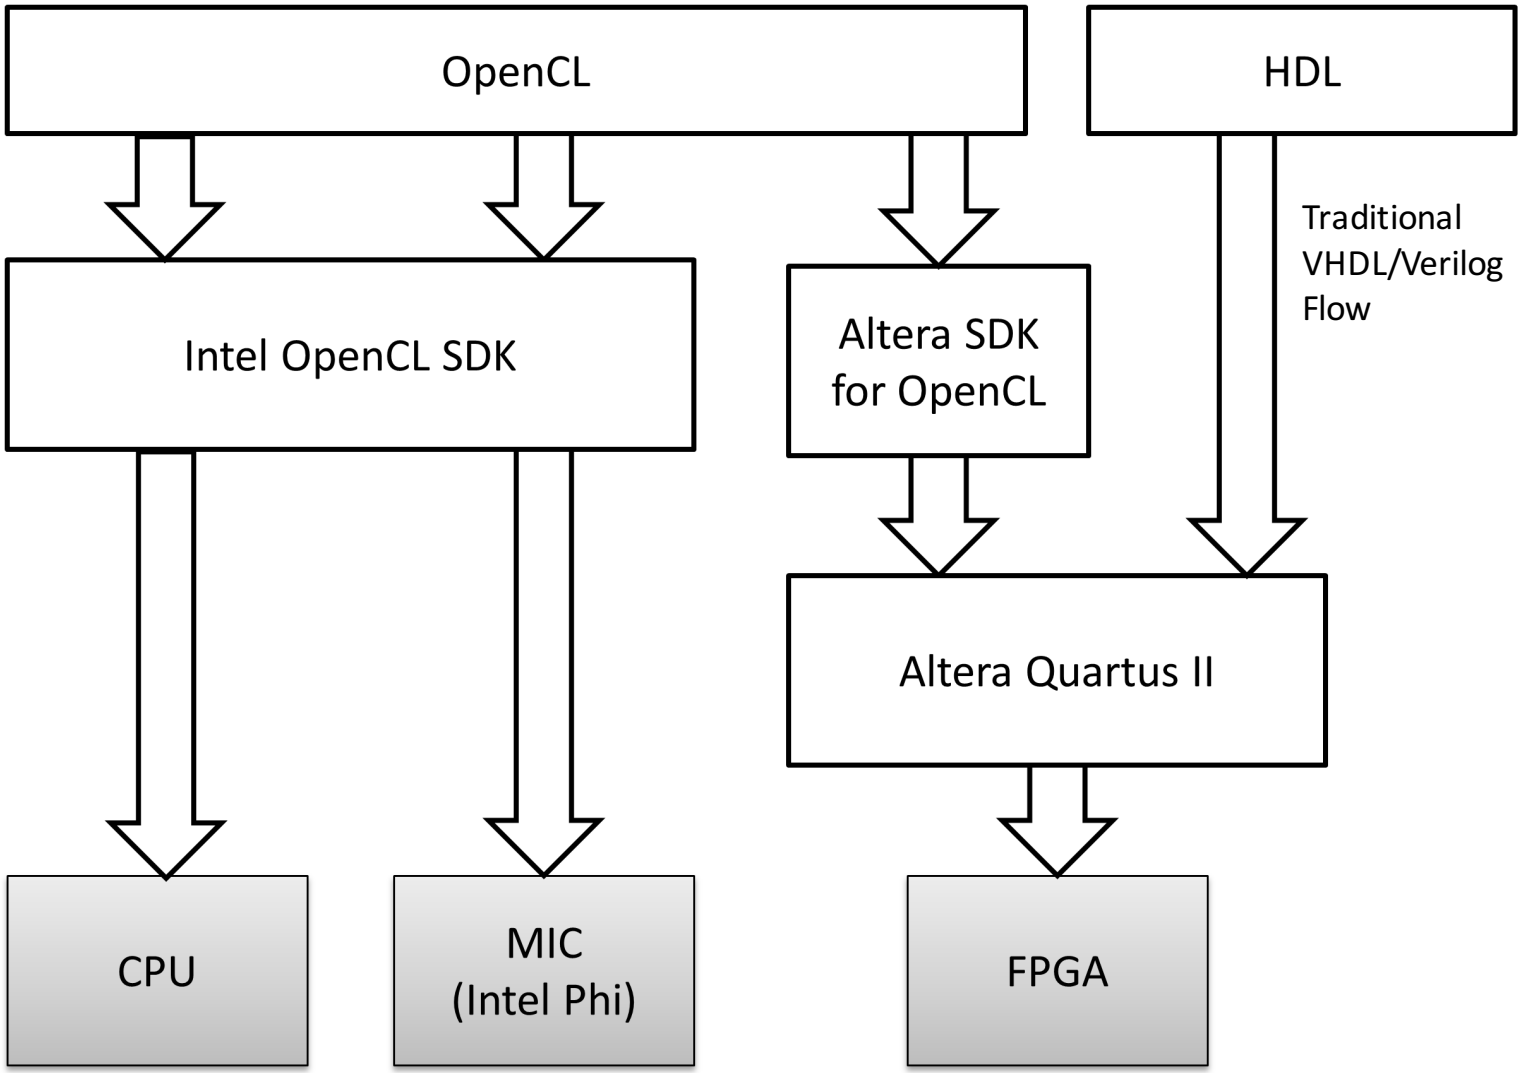

Education and Outreach

- New 1-credit course
  - Basic HPC knowledge
  - Train to participate in HPC competition
  - Freshman to Senior Undergraduates
- HPC Competition
  - ISC 2016 Student Cluster Competition
  - Build a local HPC Competition?

Conclusions

- New Research Capabilities
  - HPC
  - Novel HPC Architecture
  - Software/Hardware Co-Development
  - HPC and Software/Hardware Tool Development
- New Education Capabilities
  - HPC Courses
- New Outreach Capabilities
  - HPC Competition – Internal and External