Composable Systems Opportunities and Challenges

Wen-mei Hwu, Kaoutar El Maghraoui, Brian Pan, Christian Pinto, Carl Pearson, I-Hsin Chung

Traditional Workload

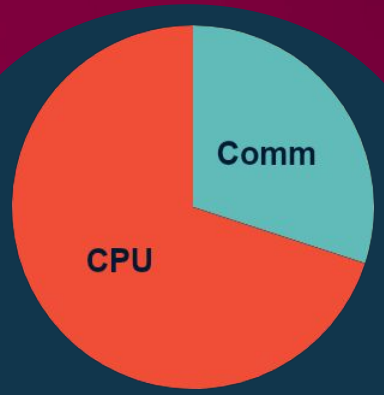

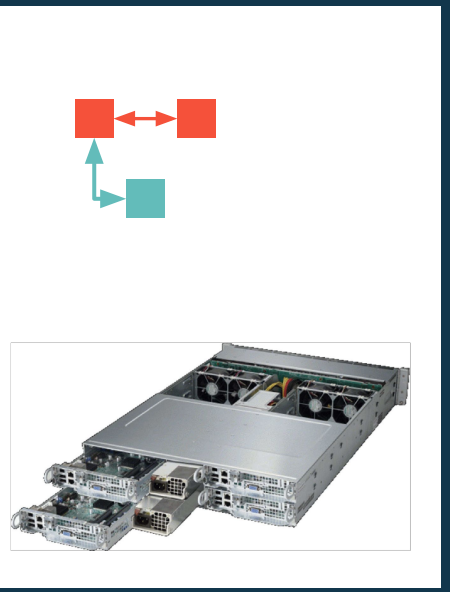

Typical Cloud Server

CPU Offload

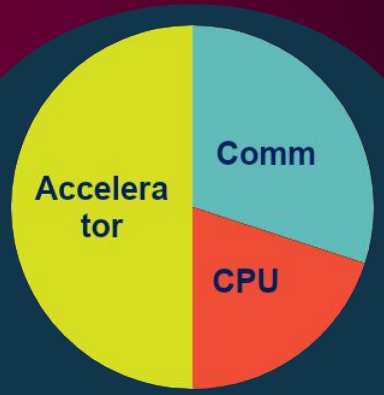

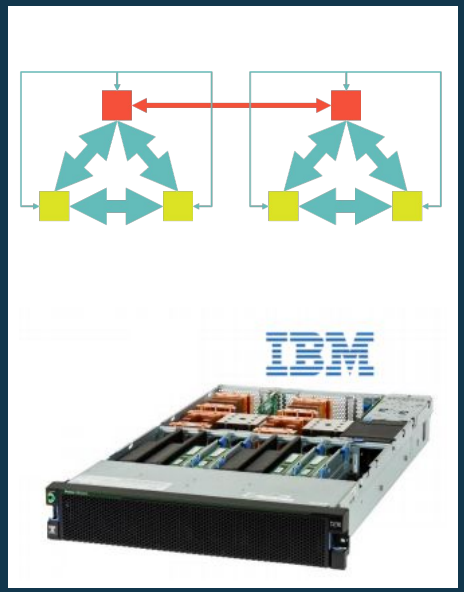

IBM POWER Server

DL Model R & D

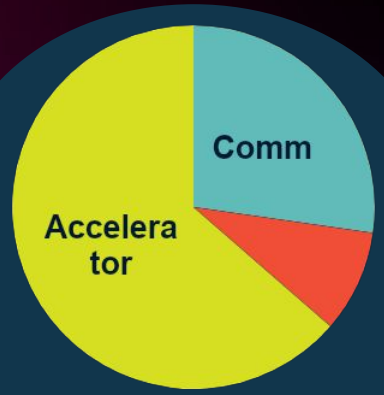

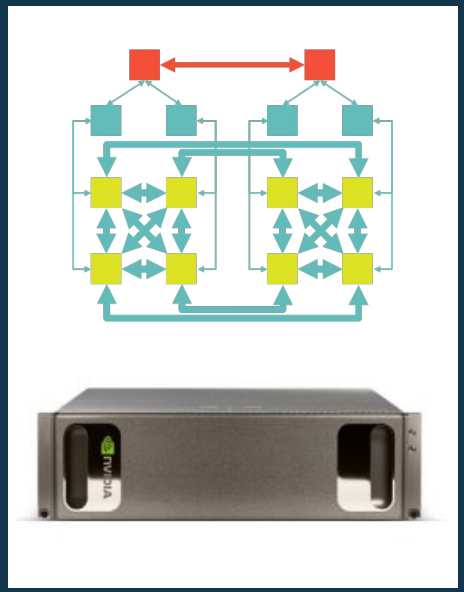

NVIDIA DGX-1

Large Scale Data

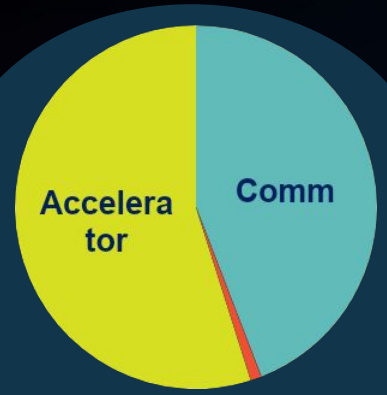

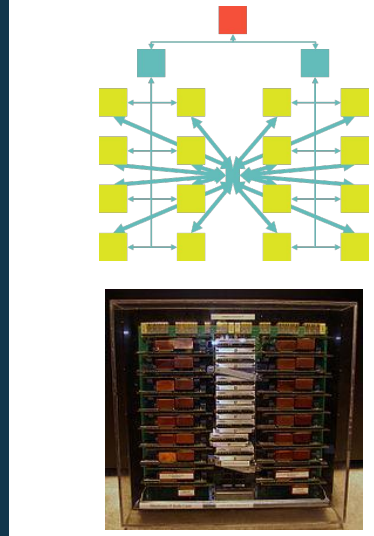

Specialized System e.g. DL Trainer

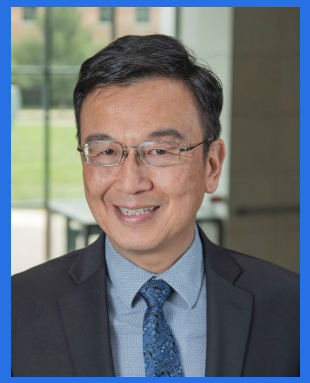					
Wen-mei Hwu	Kaoutar El Maghraoui	Brian Pan	Christian Pinto	Carl Pearson	I-Hsin Chung (Host)
NVIDIA/UIUC	IBM Research	H3 Platform	IBM Research	Sandia National Labs / University of New Mexico	IBM Research

I-Hsin Chung

- Research scientist and Manager at IBM Research
- Adjunct professor, Dept. of Computer Science, Courant Institute at NYU.

- Research interests
  - Performance analysis/modeling/tuning/tooling
  - Cloud infrastructure
  - Composable system
  - High performance computing

Wen-mei Hwu

- Senior Distinguished Research Scientist, NVIDIA
- Professor Emeritus, University of Illinois at Urbana-Champaign
- Disaggregated memory system architecture for GPUs
- Software stack for GPU-Kernel-Initiated Storage and Disaggregated Memory Accesses
- Hi-Performance graph and data-frame libraries for GPUs

Kaoutar El Maghraoui

- Principal Research Scientist at IBM T.J. Watson Research Center
- Adjunct professor, Dept. of Computer Science, Columbia University, NY
- Principal Investigator of IBM Research AI Hardware Center End-user Testbed

- Research interests
  - Next generation Deep Learning Accelerators
  - Composable Systems
  - In-memory Computing
  - High Performance Machine Learning
  - Cloud-native Supercomputing

Brian Pan

- CEO and founder of H3 Platform
- Pioneer of the GPU or NVMe composable solution by using PCIe Gen3, 4, and 5
- CXL composable memory and accelerator
- Implementation in cloud and private data center

Christian Pinto

- Staff Research Scientist at IBM Research Europe – Dublin
- Co-chair Open Fabrics Management Framework Workgroup (Open Fabrics Alliance) - Booth #2840

- Research Interests
  - Next-Generation cloud infrastructure
  - Composable systems modeling and prototyping (https://github.com/OpenCAPI/ThymesisFlow)
  - Control Plane for Composable Systems
  - Operating systems and hybrid memory management

Carl Pearson

- Scalable Algorithms – Sandia National Laboratories,
- Research Assistant Professor – University of New Mexico Computer Science

- Responsibilities & Research Interests
  - Next-generation computer systems evaluation
  - Accelerator kernel performance
  - Distributed linear algebra

- Here in my individual capacity. Not a spokesperson for Sandia or affiliates

Topics

- Overview
- Current state
- Changes needed
- Reliability
- Adoption
- Future outlook
- Q&A

1. What is a composable computing system?

2. What are the advantages and disadvantages?

3. What is the current state of composable computing?

4. What are the changes needed in hardware/software to adopt the technology?

5. How robust are composable systems compared to traditional server clusters?

6. How can existing complex HPC workflows to be combined with composable systems?

7. How can novel/emerging devices be incorporated into a composable system?

8. How to diagnose, analyze, cope with reliability with a large mix of device types?

9. What are the future directions and expectations of composable system technology?

Q&A

- hw/sw change
  - heterogeneous
  - shortcut/shared mem(device)
  - dynamic/intelligent resource allocation/perf diagnosis
  - standardization
  - reliability/hotplug
- user/provider
  - trade-off of perf and flexibility
  - automation
  - transparency

Composable Systems Opportunities and Challenges

Supercomputing 2022 Panel

SC22

Dallas, TX | hpc accelerates.

Cast

- Panelists
  - Wen-mei Hwu (UIUC, NVIDIA)
  - Carl Pearson (Sandia National Laboratory)
  - Brian Pan (H3 Platform)
  - Kaoutar El Maghraoui (IBM Research)
  - Christian Pinto (IBM Research, Ireland)
- Host
  - I-Hsin Chung (IBM Research)

Time and Location

- **10:30am-12pm CST Tuesday, 15 November 2022**
- C147-148-154, Kay Bailey Hutchison Convention Center Dallas

- All panelists need to attend in person
- Please arrive at the location 30 mins prior to the panel event.
- Shall we meet on Monday, November 14?

Planned schedule

- 2 min introduction by host
- 10 min - 2 min self-intro by each panelist
- 60 min - 6 prepared questions. 2 min per panelist
- 18 mins Q&A from the audience

List of questions from the reviews

- Overview
  - What is a composable system? Provide a definition.
  - What are the advantages/disadvantage of composable systems?
- Operation
  - Cost-benefit of any required (significant?) changes to existing datacenters (including human resources)?
  - How to diagnose, analyze, and cope with reliability for a large mixture of device types?
- Reliability
  - How robust are composable systems compared to traditional HPC? (ed: reliable?)
  - What if hardware is not refreshed in a timely manner?
- Usage
  - How can existing complicated HPC workflows be combined with composable systems?
  - Current status of composable system performance, and current / future approaches?
- Outlook
  - What is the current state of composable systems technology?
  - What are the future expectations for composable systems technology?
  - How can novel/emerging devices be incorporated into a composable system?