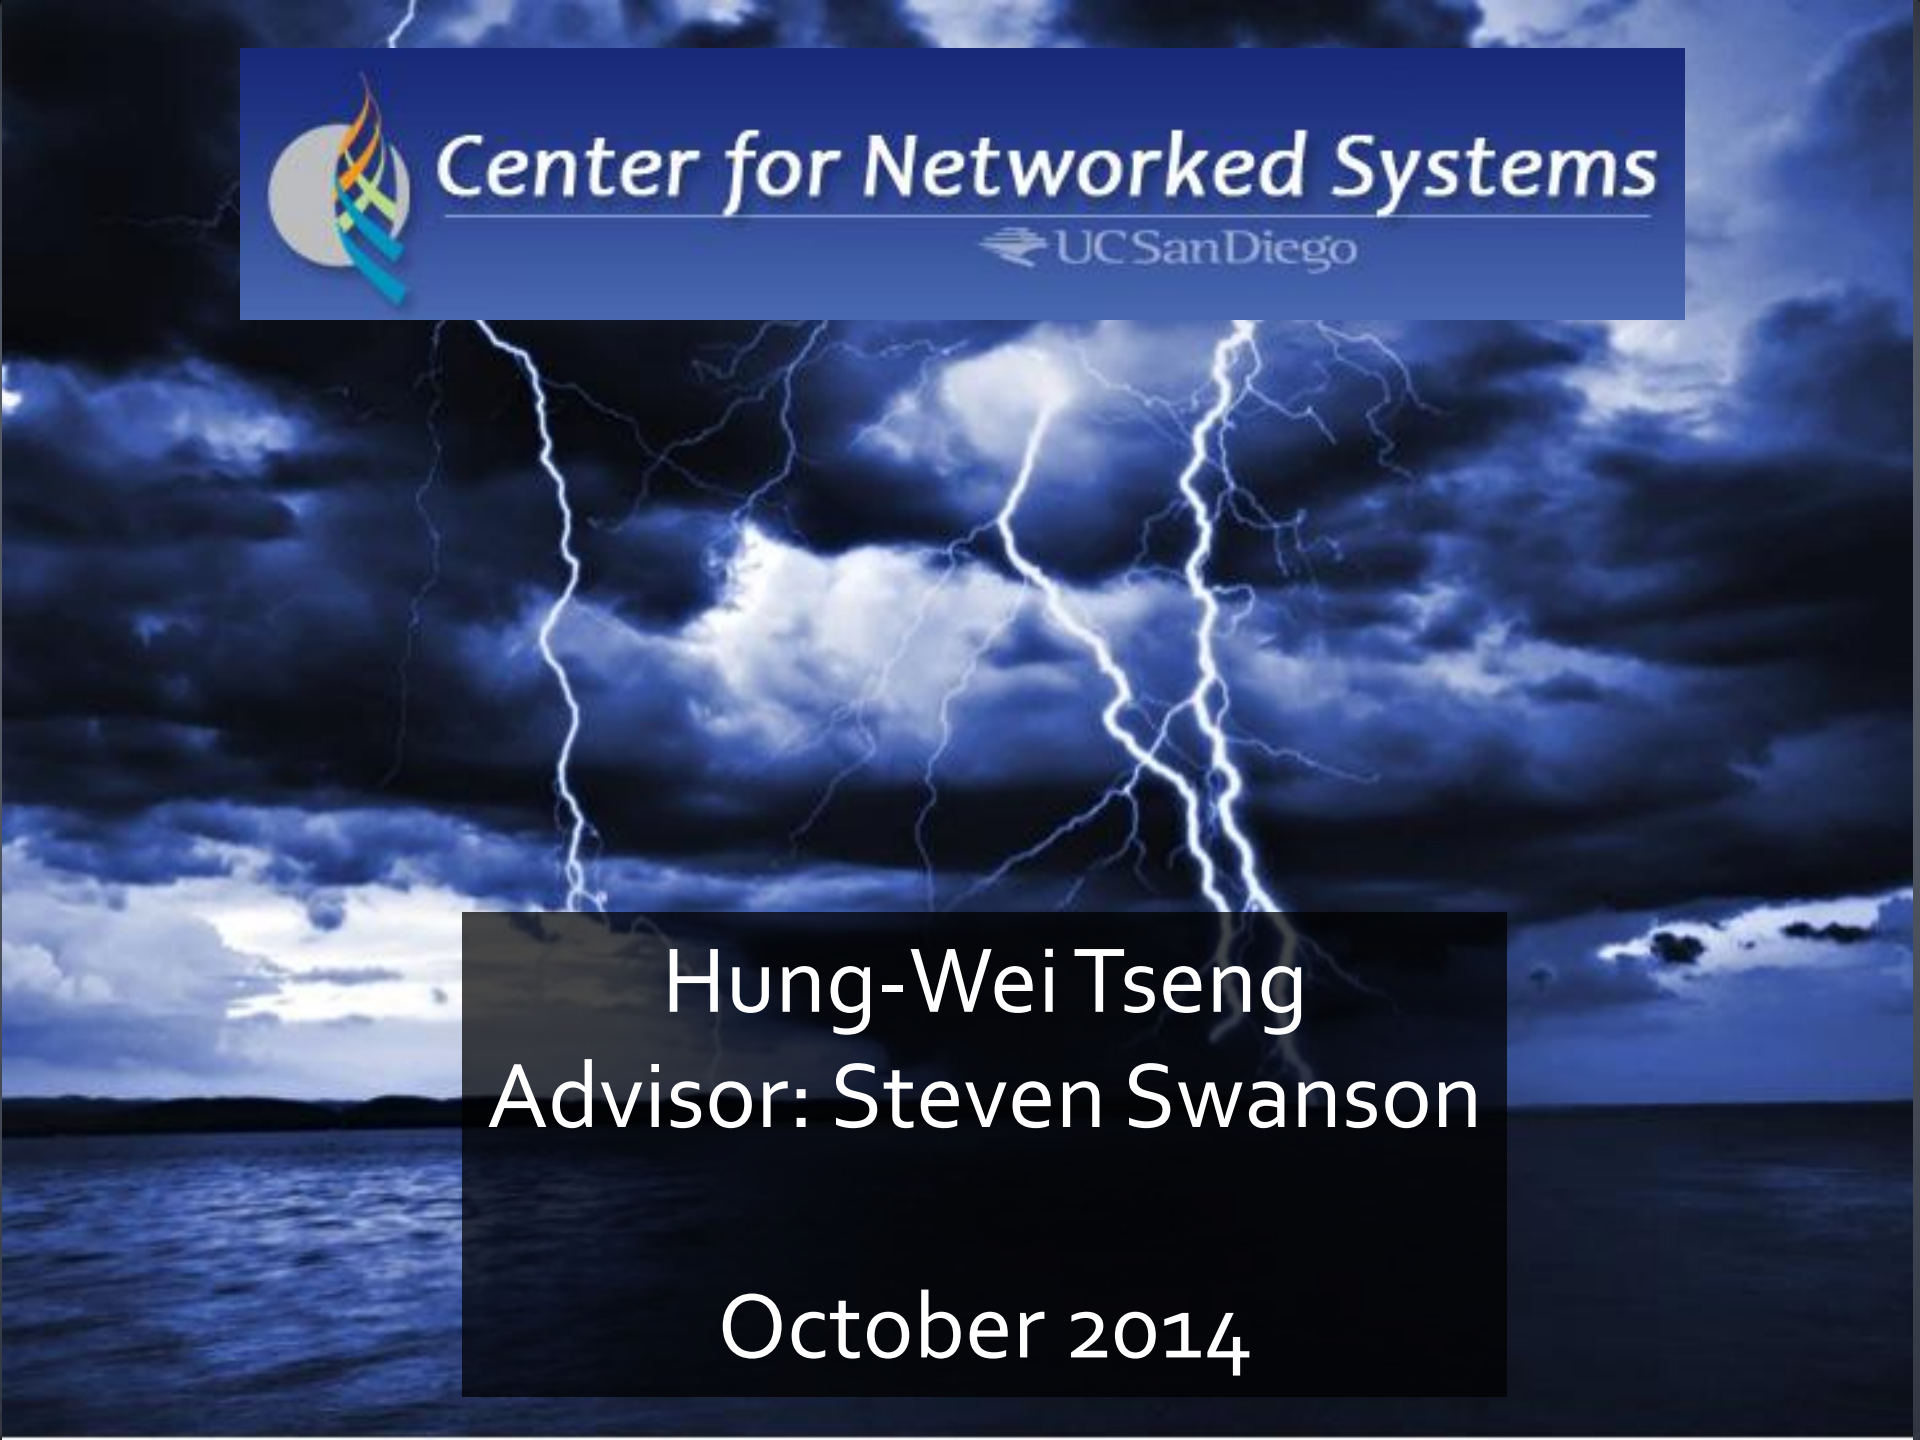

Hung-Wei Tseng
Advisor: Steven Swanson

October 2014

As hardware gets more sophisticated, programmability emerges

Fixed function

Limited programmability:
Programmable Shader

Full-blown programmability

cns
center for networked systems

The rise of fast NVMs increases storage flexibility and performance demands

Read/write

Full-blown programmability

OS offloading

Willow

- Exposing the programmability of modern SSDs
- A generic, easy, safe programing interface for ordinary programmers based on RPC
- Published in OSDI 2014

cns
center for networked systems

Willow: A User-Programmable SSD

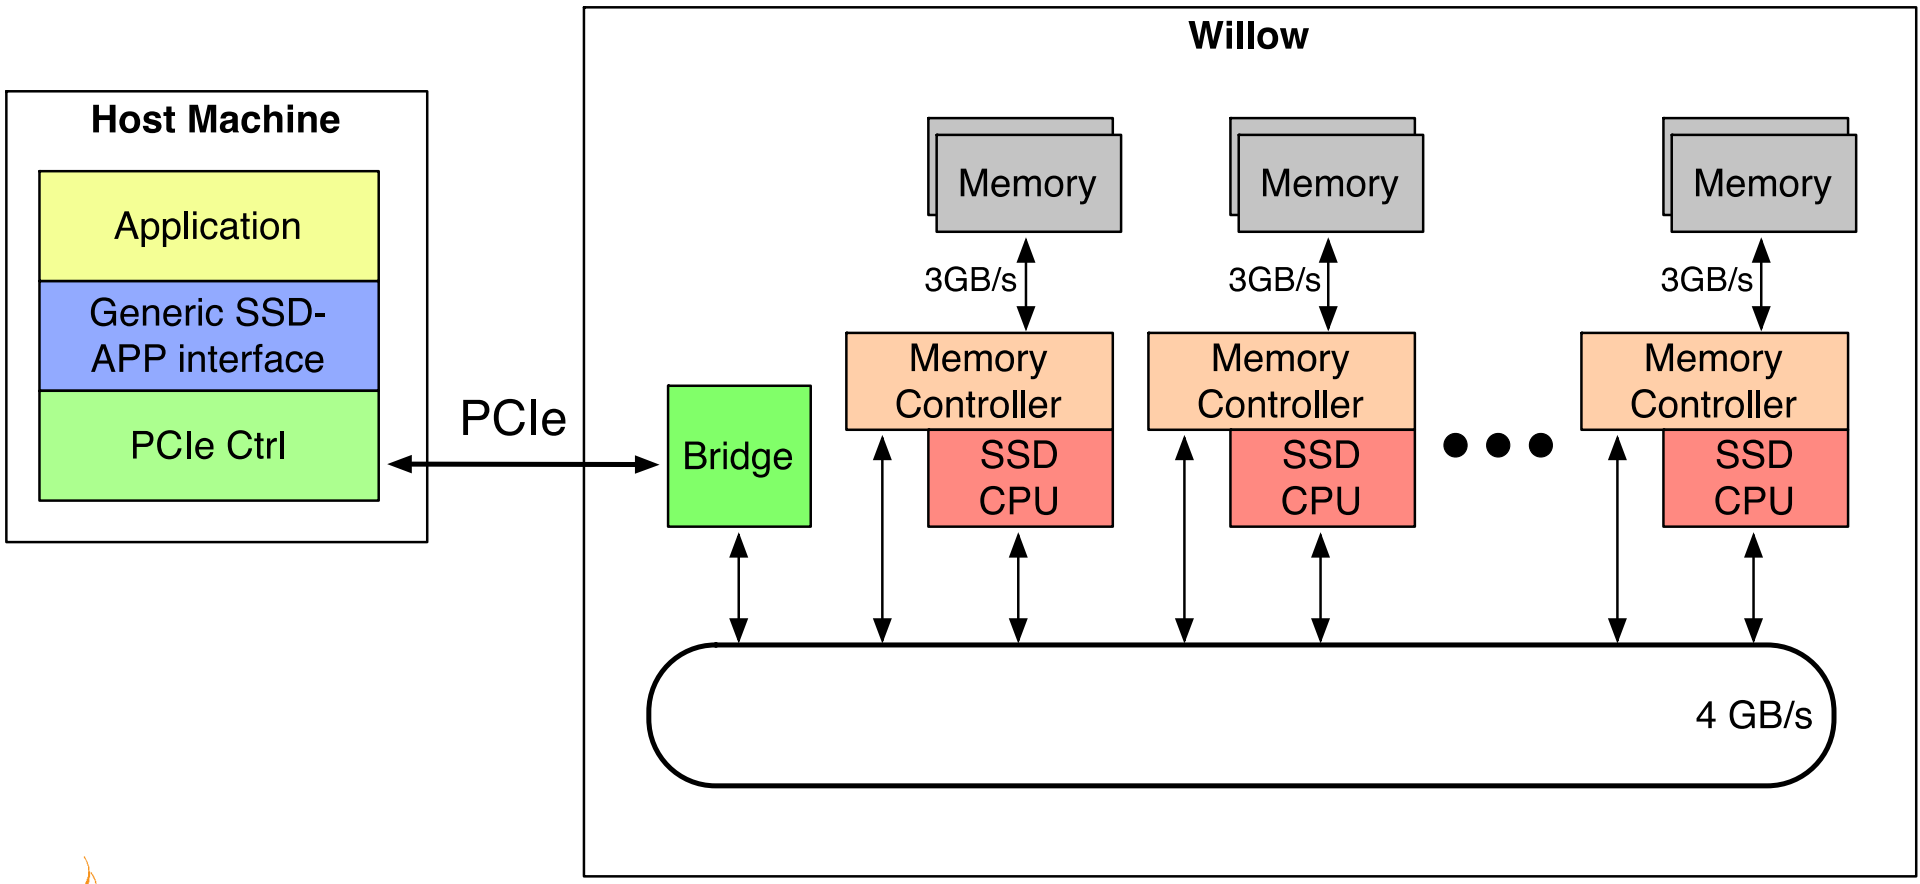

Candidate applications for willow

- Data-dependent accesses
  - e.g. pointer chasing
- Semantic extension
  - e.g. transactions
- Privileged execution
  - e.g. OS offload
- Data-intensive computation
  - Database scans
  - Transcoding
  - Analytics

Hippogriff

cns
center for networked systems

Hippogriff: platform for data-intensive computation

Massive vector processing
Floating point operations

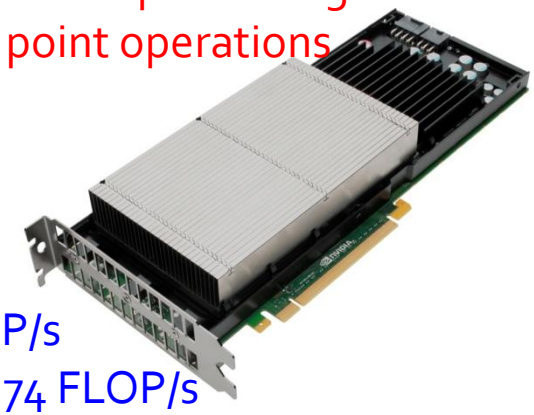

GPU
1.17TFLOP/s
Per byte: 74 FLOP/s

More complex integer operations
OS/application workflow

Processor
96 GIntOps/s
Per byte: 7.68 IntOP/s

Programmable SSD
2+ GB/s read bandwidth
0.5 GIntOps/s
Per byte: 0.25 IntOP/S, software FP

Lightweight data processing
Data-dependent access
OS off-loading

cns
center for networked systems

CPU-centric model

Massive vector processing
Floating point operations

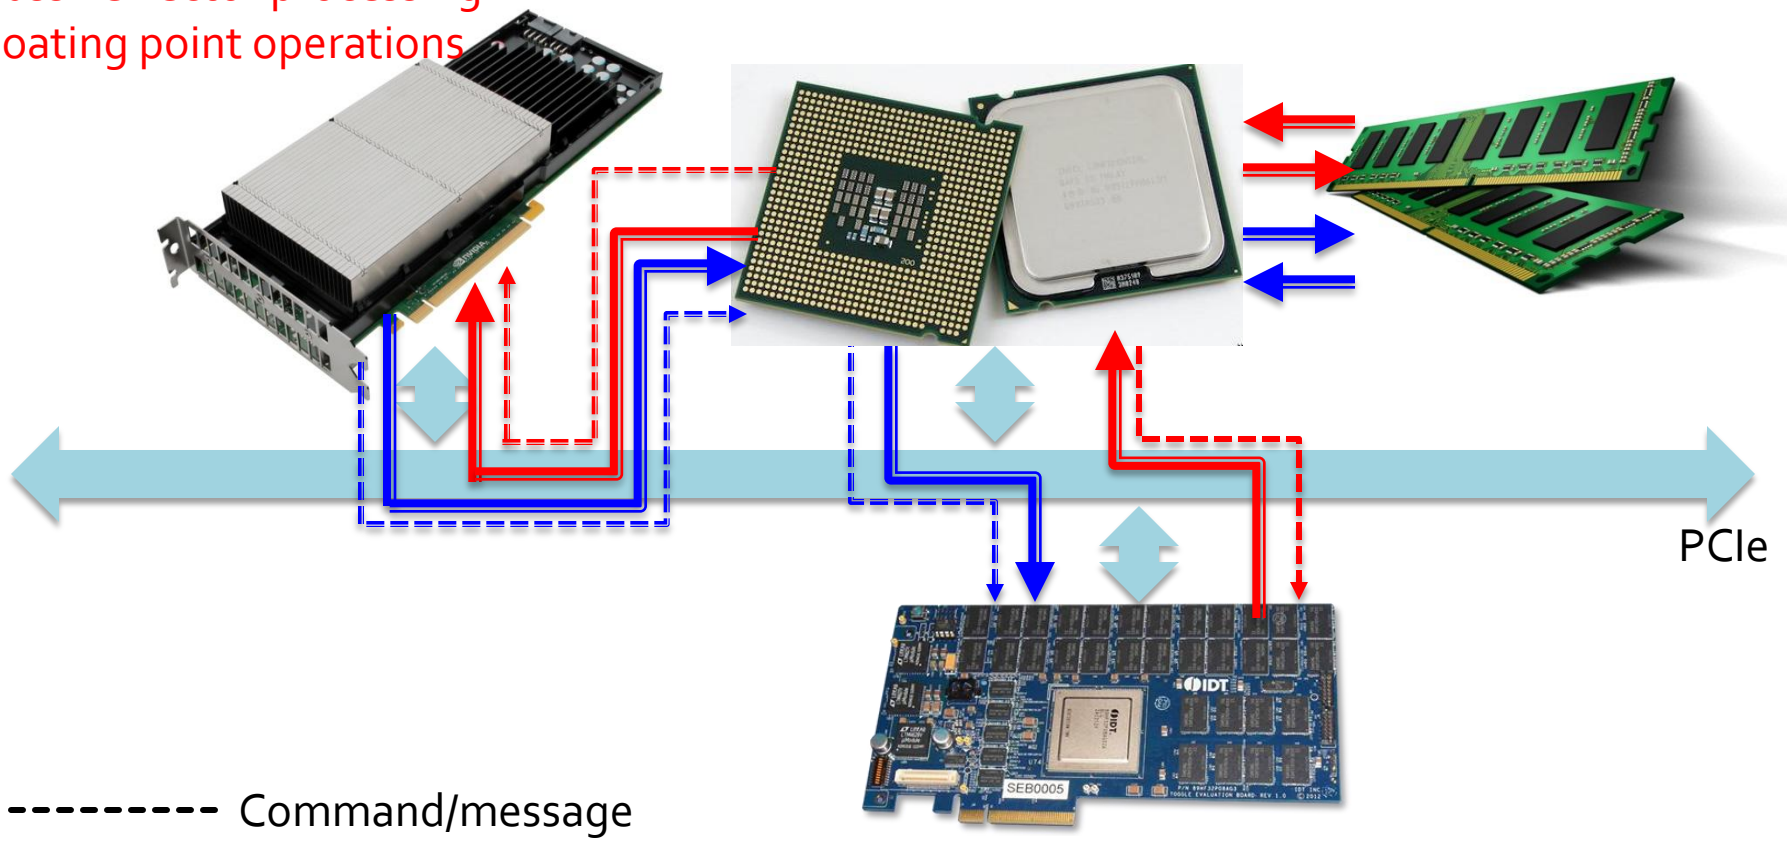

Command/message

Data

Lightweight data processing
Data-dependent access
OS off-loading

cns
center for networked systems

Hippogriff: storage-centric model

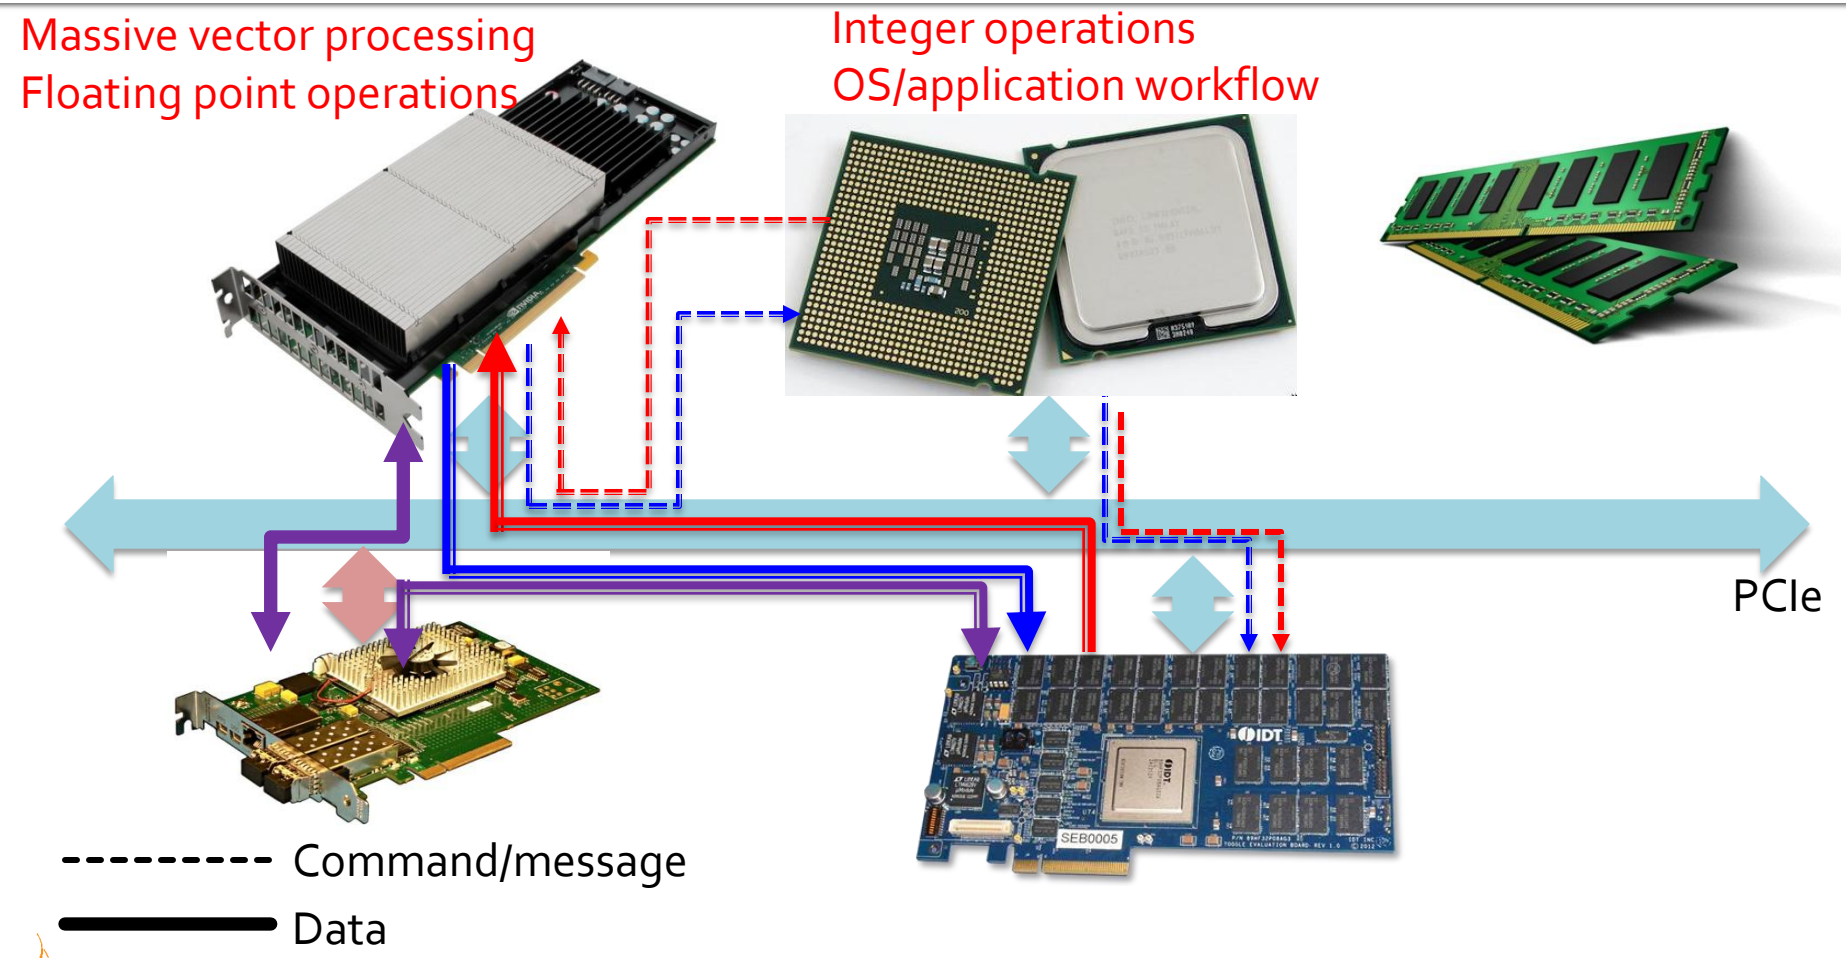

Current progress

- Achieved direct SSD and GPU data transfer
- Willow-like programing interface for SSDs
- MapReduce framework supports GPU
  - MapReduce runtime helps determine the best hardware resource to execute the program
  - Decision based on the amount of computation for each byte of data