Research on High Performance Computing with Shared Persistent Storage

MING-HONG YANG

11/07/2018

Outline

Background and Motivation

Problems

Big Data

- What is big data?
- Social Network
- Internet of Things
- •

By 2020

- 30+ billion connected devices
- 200 billion IoT devices
- 1 billion smart meters

Architecture remains the same over the past 60 years

Limitations

- Compute
- Memory
- Storage
- Network

Performance gap between CPU and DRAM

Storage Class Memory (Non-volatile Memory)

- Non-volatility
- Byte addressability
- Lower power consumption
- •

From processor centric to memory driven ...











Special purpose cores

Massive memory pool





Problems

Operating Systems

Data Stores

Analytics Platforms

Programming Models and Tools

Applications

Algorithms

•••

Problems

Page Cache in Shared Persistent Memory

- Where to cache pages?
- Who should get more space?
- How to do page eviction collaboratively?

Checkpointing

- What to be dumped?
- Where to dump the image?

•

Reference

https://www.mcs.anl.gov/events/workshops/ross/2015/slides/ross2015-keeton.pdf

http://moais.imag.fr/res/slides/20150709-HP-The_Machine.pdf

https://www.slideshare.net/linaroorg/the-hpe-machine-and-genz-bud17503

http://www.pdl.cmu.edu/SDI/2016/slides/keeton-2016-10-19-memory-driven-computing.pdf