Converging Memory and Storage for Big Data Applications

Qiong Cai
Systems Research Lab
HP Labs

MPSoC – July 2015

Agenda

- Big data applications
- Architecture Challenges and Opportunities
- The Machine Project in HP Labs
- Conclusions

Disclaimer: The views in this talk are my opinions and not necessarily those of HP.
Acknowledgement: Paolo Faraboschi

Data explosion outpacing technology

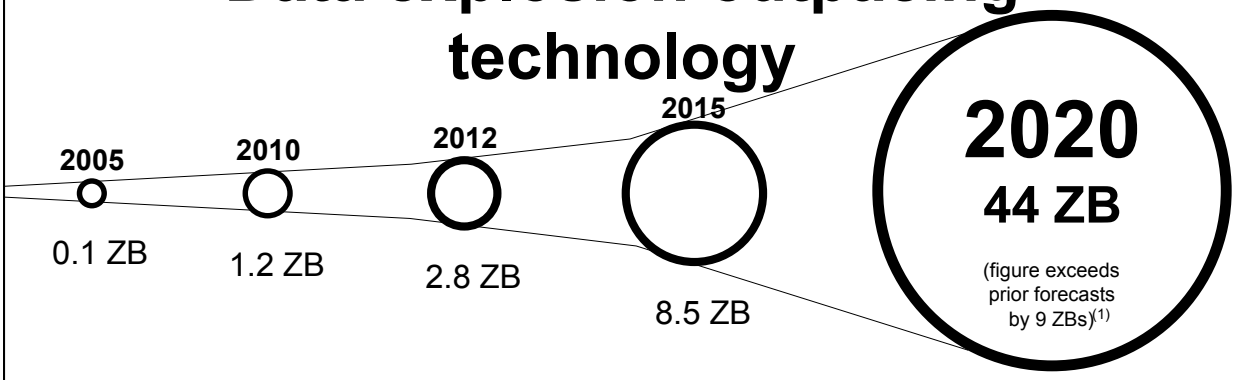

Next-generation competitive advantage delivered through:

- **Business insight at real-life speeds**
- **Personalized content that follows you**
- **Questions that arise automatically from data**

(1) IDC "The Digital Universe of Opportunities: Rich Data and the Increasing Value of the Internet of Things" April 2014

Big Data Applications

- Traditional database queries and processing
- Graph analytic applications
- Machine learning applications

In-memory applications

Big data applications significantly increase the capacity and bandwidth requirements for the memory system.

Architecture Walls

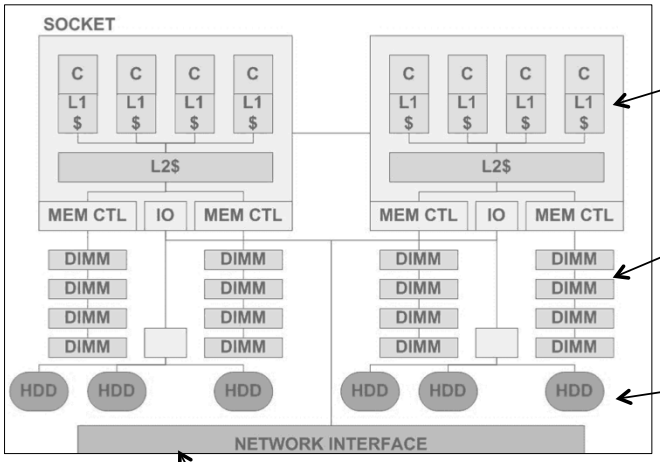

Compute

Single-thread performance wall. Diminishing return of multi-core. Chip-edge bandwidth wall

Memory

DRAM reaching technology scaling wall

Storage

HDD/SSD layer is a significant performance bottleneck. Prevents big data getting closer to compute

Data Movement

Too slow (and cumbersome) for real-time access to shared memory

Memory System Requirements

- High bandwidth
- Large capacity
- Low and predictable latency
- Low power
- High reliability

Current solutions: multi-channel DDR-based main memory systems.

Memory System Challenges

- Bandwidth and capacity increase significantly mainly due to the following reasons
  - More computing nodes in the system
  - Emerging memory-intensive applications such as in-memory Big Data applications
- DRAM scaling may not at 1x nm
  - Current scaling solutions such as relaxed core timing and on-die ECC decrease the performance and increase the complexity

Multi-channel DDR-based main memory system may not be scalable in the long run.

Emerging Memory Technologies

- High bandwidth memory technology
  - Hybrid Memory Cube, High Bandwidth Memory, WIDE-IO
- High performance persistent memory
  - STTRAM, Memristor
- High capacity persistent memory/storage
  - Phase change memory, 3D NAND

We can leverage different properties of memory technologies to architect a flat memory and storage system.

Architecting NVM

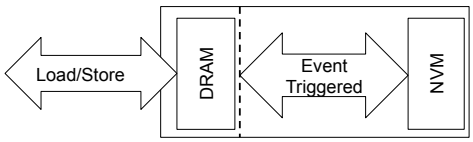

Load/Store, Block **(JEDEC/NVDIMM-N)**
No data copy required during access
Capacity = DRAM
Performance = ~DRAM

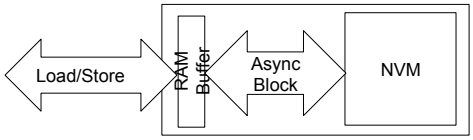

Block, (guarded) Load/Store **(JEDEC/NVDIMM-P)**
Some data copy required during access
Capacity = NVM
Performance = DRAM (hit), NVM (miss)

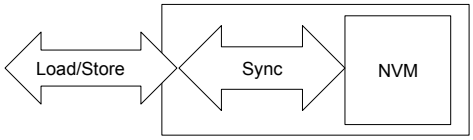

Load/Store, Block **(NVDIMM-x)**
No data copies required during access
Capacity = NVM
Performance = NVM

- DRAM can be combination of traditional DDR memory and emerging high bandwidth memory.
- NVM can be combination of high performance persistent memory and high capacity persistent memory.

The Machine Project

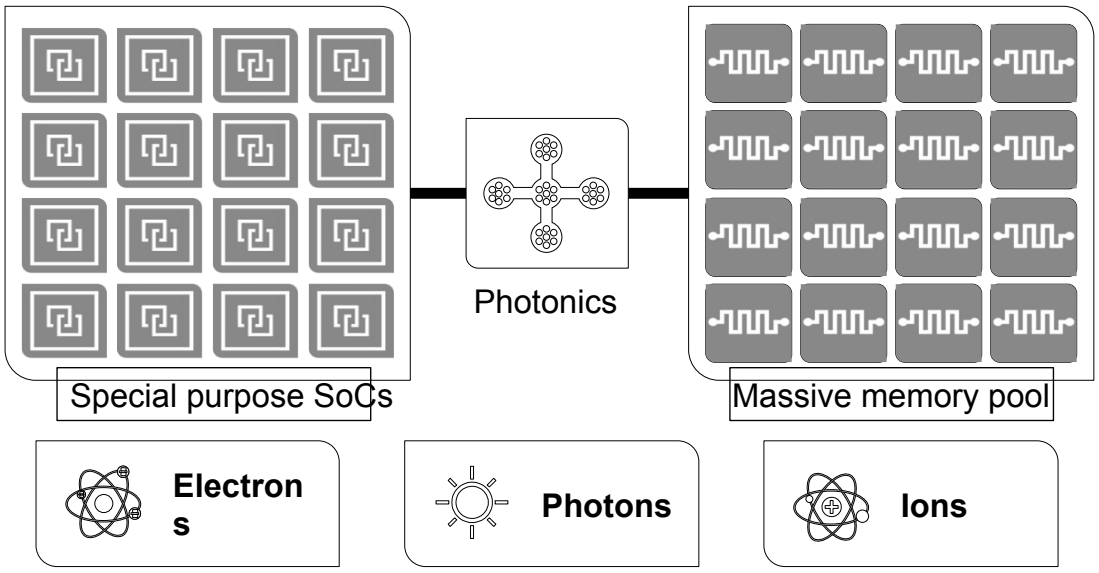

Special purpose cores

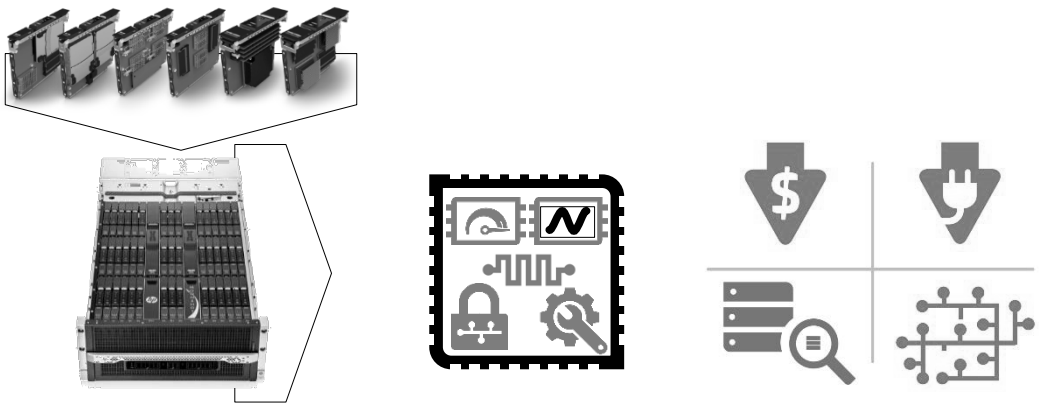

SoCs customized to the workload

Photonics

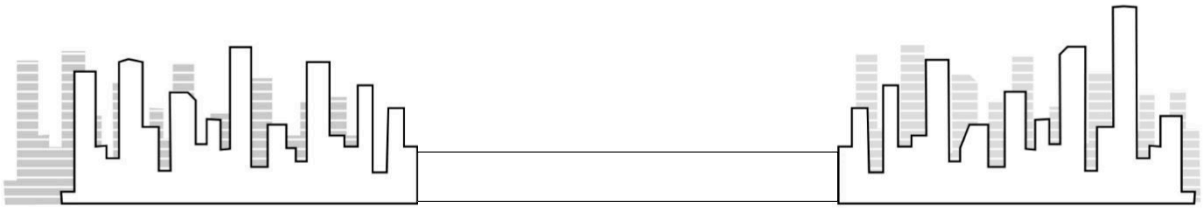

Photonics destroys distance

Massive memory pool

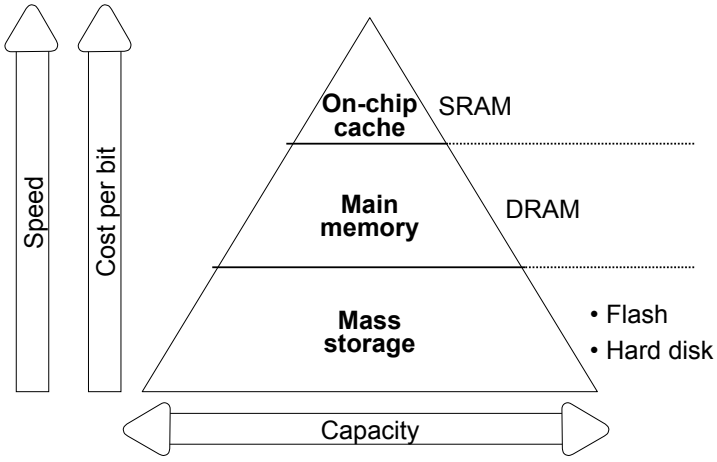

Universal memory obsoletes this hierarchy

NVM at Scale

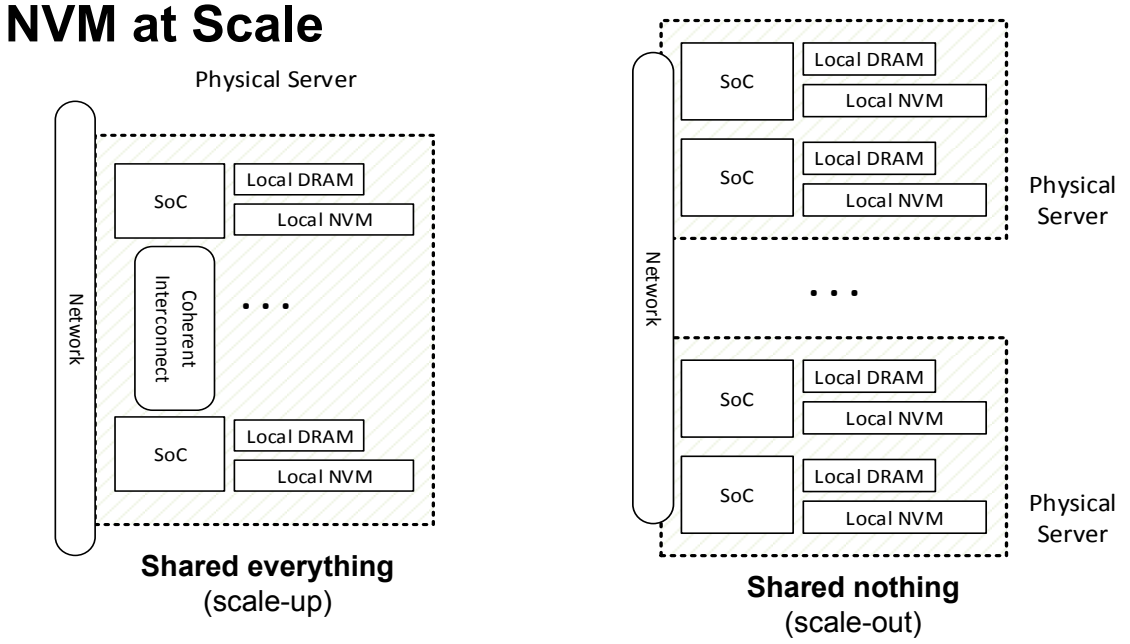

Converging Memory and Storage

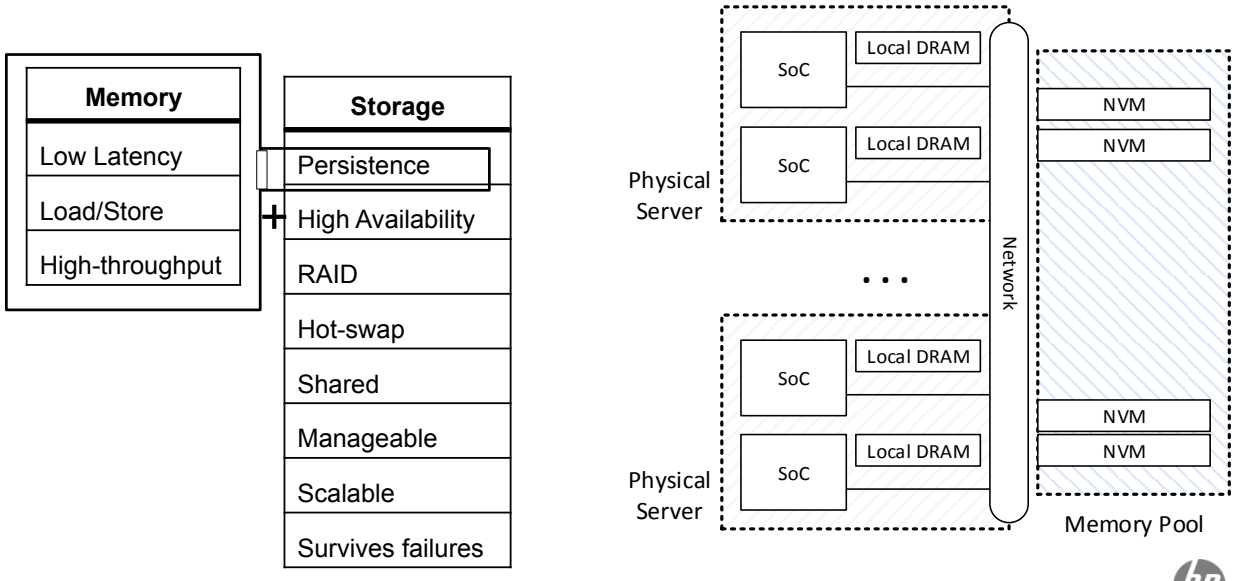

Memory Side Accelerator for NVM– High Level Characteristics

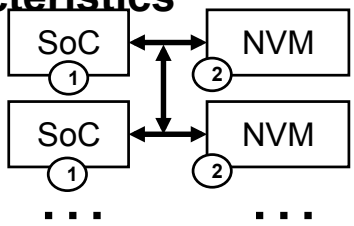

This is a 40-years old idea. Why now?

- MSA can help to decrease the latency. It is critical to reduce NVM latency, since NVM latency is still longer than DRAM latency.
- NVM bandwidth is precious. MSA can utilize NVM bandwidth more effectively.
- MSA provides more energy efficient way to compute memory intensive operations.
- The programming model and tool chain for accelerators are getting mature. MSA can leverage it.

Label	Position	Types	Characteristics	Usages
1	SoC Side	GPU or FPGA accelerator in another socket or on the PCIe	Powerful computing capability (5x-10x more than CPU), memory bandwidth could be bottleneck (GPU has its own local memory)	The whole application
2	NVM Side	Simple cores in the memory controller	Limited computing power (1/10 of SoC), low memory latency and spare memory bandwidth	Dependent and memory intensive code snippet ; background data processing

Conclusions

- We expect to see a converged memory and storage system soon
- It is a whole stack change
  - Hardware: NVM technology, CPU, memory controller, etc
  - System Software: OS, compiler, memory management, runtime framework, etc
  - Applications: algorithms and data structures, use cases, etc