


SSRLabs


Hybrid Memory Cube-Based Unified Memory

August 2015

© 2015 Scalable Systems Research Labs Inc.


Axel Kloth, President & CEO SSRLabs






Overview and Motivation

- Big Data
- HPC as defined and requested by President Obama for the ExaScale Challenge (1 ExaFLOPS at 20 MW)
 - Unstructured data = random distribution of data across all addresses in the address space.
 - Random accesses to random addresses decrease efficiency of CPU caching strategies which rely on spatial and some degree on temporal locality.
 - Worse than lack of locality is the need to swap to disk – even to and from an SSD or PCIe-attached Flash.
 - Plenty of evidence that Big Data and HPC fare better on computers with very large main memories, even if they are slower than DRAM.




Motivation for Big Main Memory

- Why is there a need for a new type of memory?
- The problem size (Big Data, HPC) keeps growing
- Economic considerations rule out SRAM and DRAM
- What we really need is
 - Big
 - Fast
 - Cheap
 - Energy-efficient




Very large capacity Main Memory

- Total Main Memory size must grow to accommodate in-situ processing
- SRAM and DRAM are not dense enough and consume too much power
- SRAM and DRAM are too expensive
- SSDs and PCIe-attached Flash are too slow
- Very big main memory often can avoid swapping
- Really, Big Data means never having to go to Disk

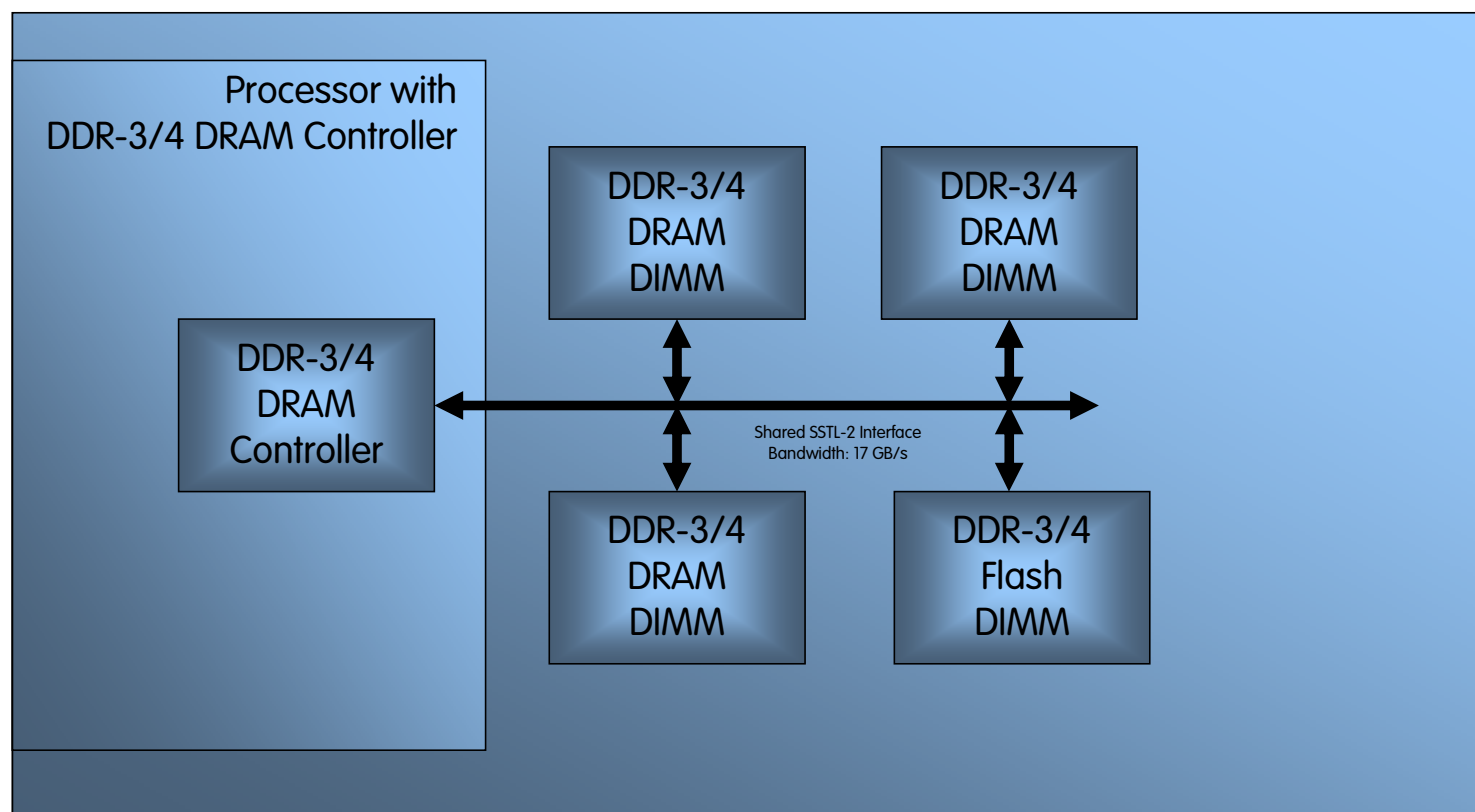



Practical Solutions

- Direct attachment to the CPU is preferred over SAS, SATA or PCIe for latency reasons
- DDR3, DDR4 and HBM rely on outdated buses
- A faster infrastructure is needed, such as Hybrid Memory Cube and its High Speed Serial Links
- The memory controller(s) should reside with the memory, not on the CPU
- 3D XPoint is in its infancy
- It is a material property change in intersecting wires

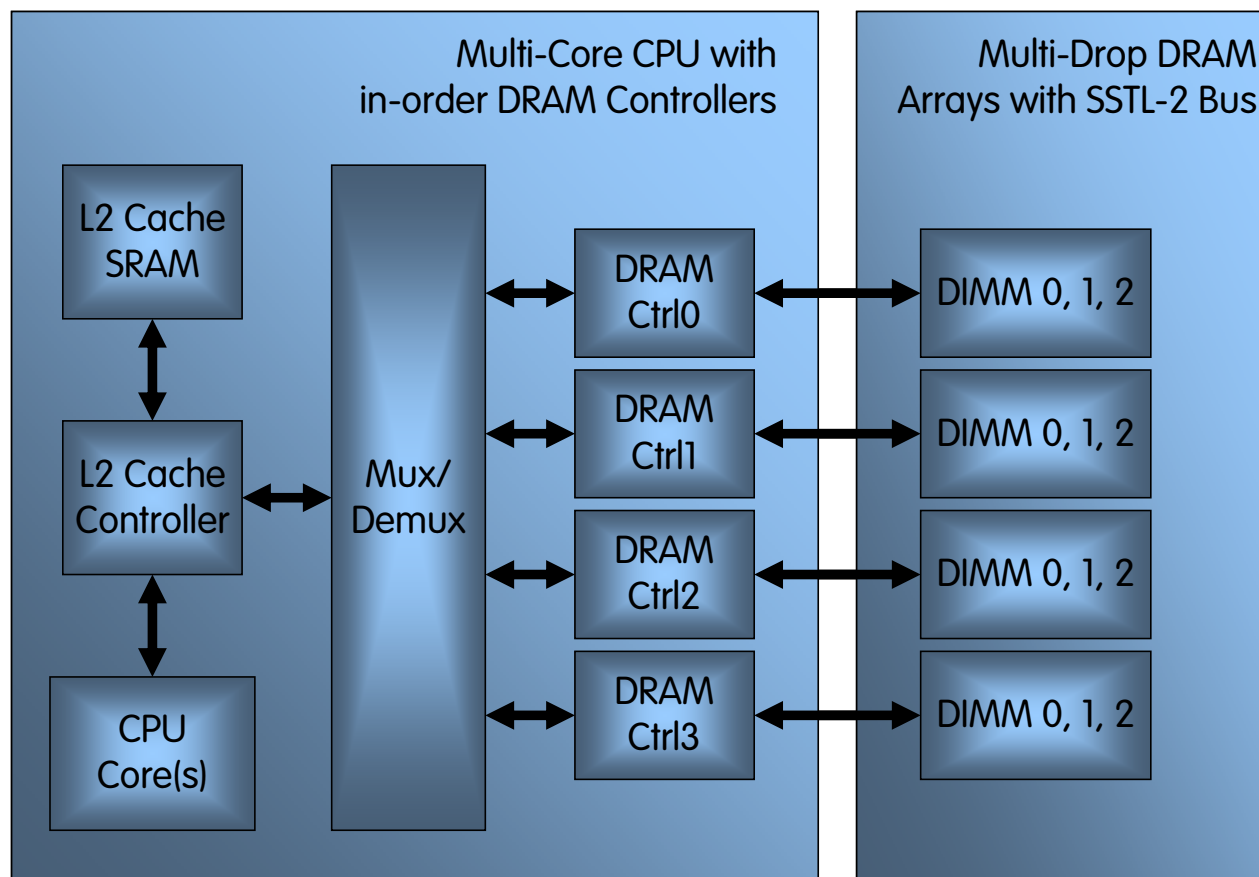



Current CPU & Memory



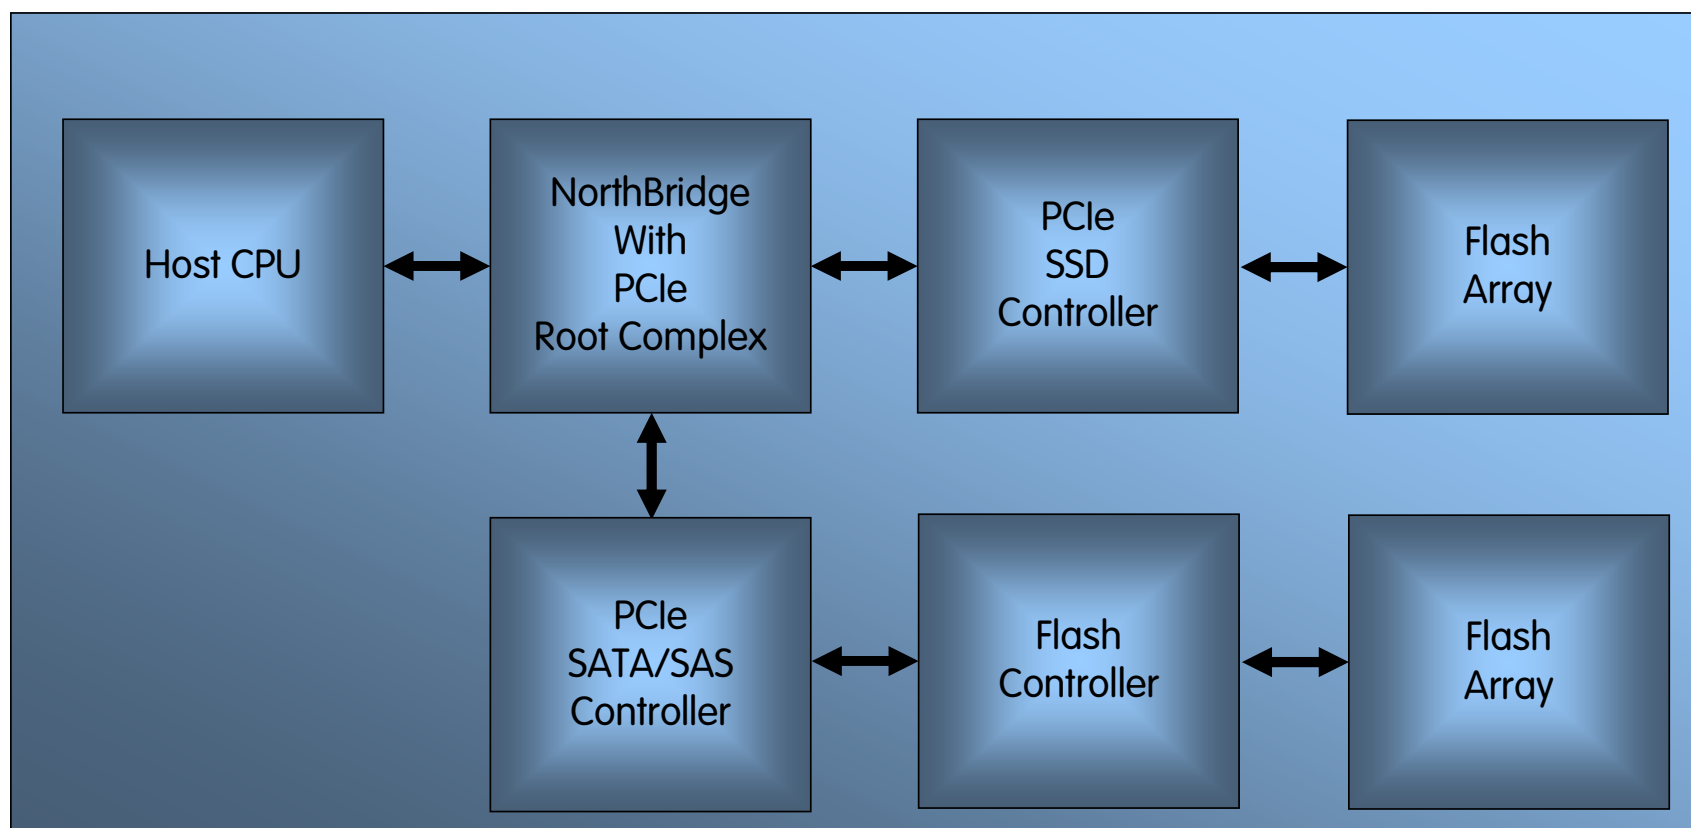


Multi-Core CPU with Memory

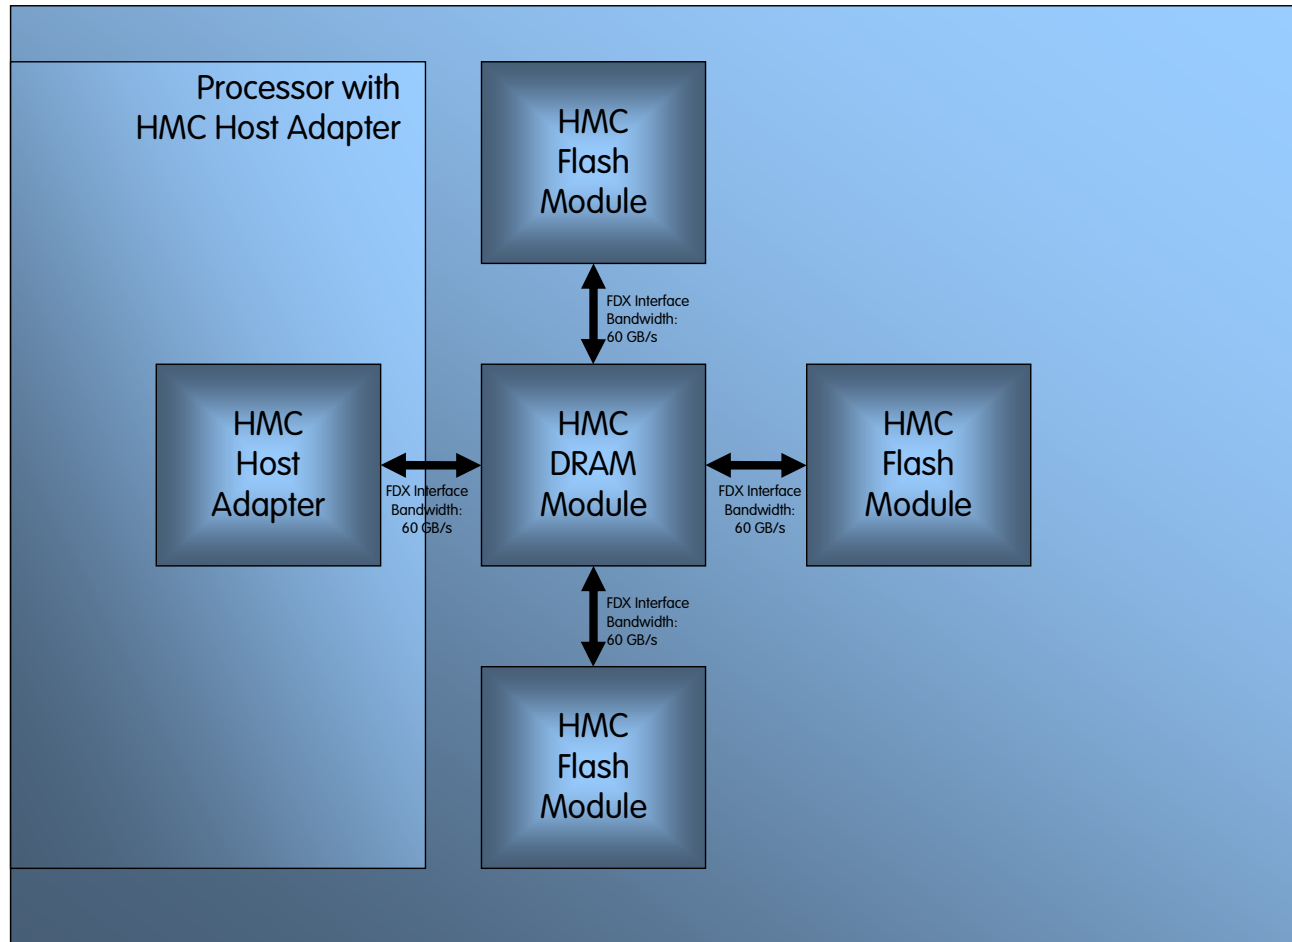






Host CPU to Disk I/O

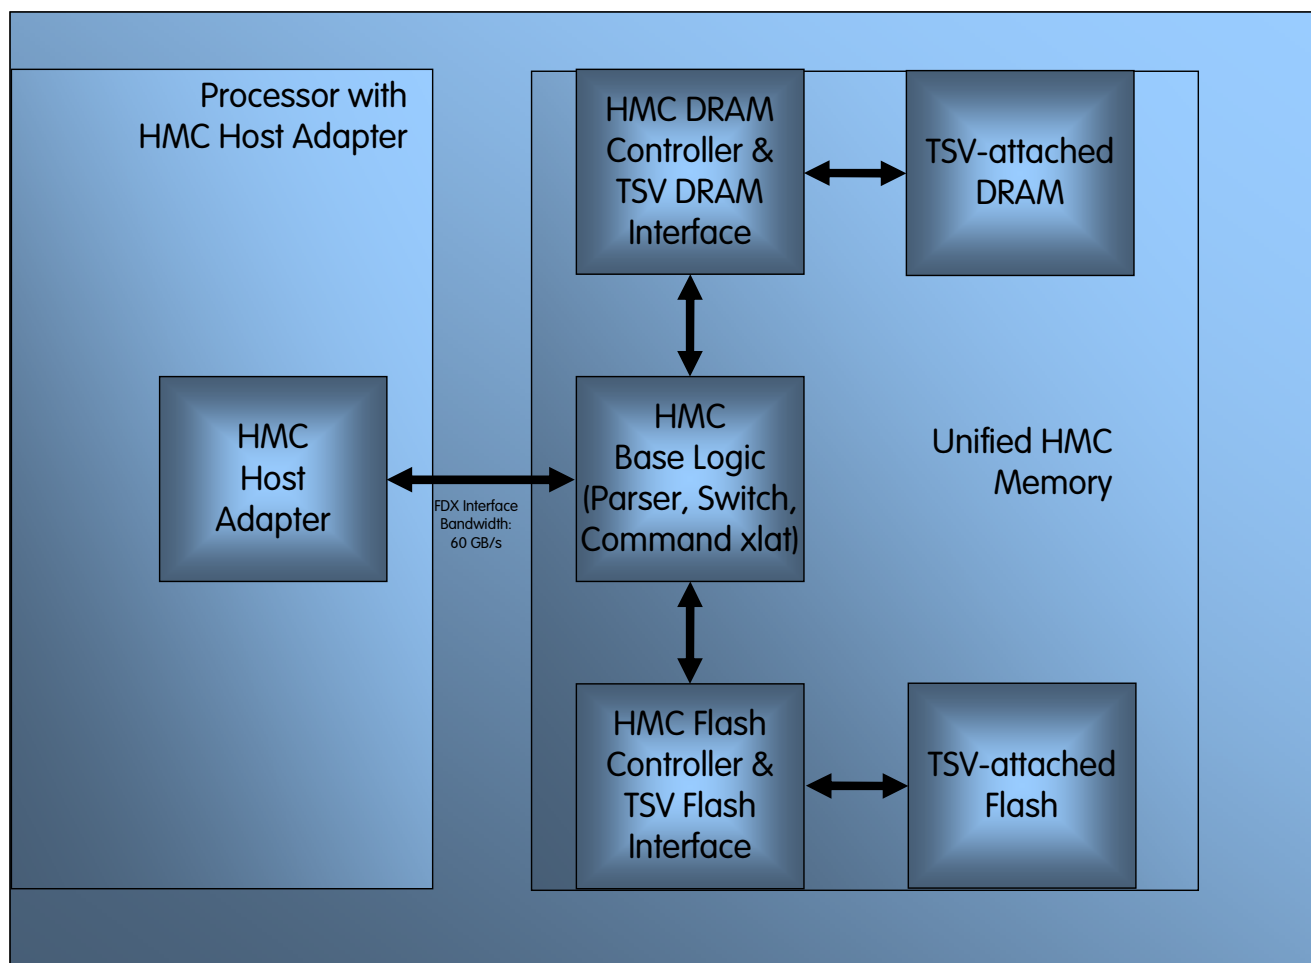



Single-Port HMC-based Memory





SSRLabs Unified HMC Memory






Cost Comparison

- Assumption: 512 GB Memory Array
- Source: DRAMXChange

Type	Per-Unit Cost	Number needed	Total Cost
DDR4 DRAM Chip	\$3.35	1024	\$3,430.40
4 GB Registered DIMM (DDR4)	\$66.99	128	\$8,574.72
32GB DDR4 PC4-17000 Load Reduced ECC 1.2V 4096Meg x 72	\$469.99	16	\$7,519.84



Benefits of a Unified HMC Mem

- 3D and TSV manufacturing is maturing
- All components are readily available
- Internal and port bandwidth exceed all legacy memory architectures
- Better than DDR3/4 DRAM Performance at better price, density and power