SSRLabs

Unified Memory for HPC and Big Data

August 2016

Axel Kloth, President & CEO SSRLabs

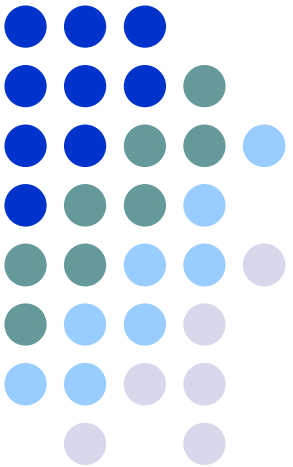

Overview and Motivation

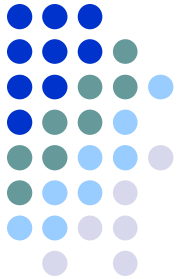

- The requirements for HPC and Big Data are converging for both processors and memory.
- Very large data sets, large growth rates of the existing data sets and less locality (aka "unstructured data") force the deployment of new memories.
- Removing memory bottlenecks and density issues is paramount to solve both Big Data and HPC Challenges.

A new approach is needed

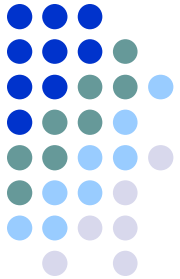

- A novel memory architecture for HPC and Big Data is required.
- It cannot be DDR3, DDR4 or HBM as those keep the Memory Controller inside the CPU (within the UnCore).
- The Memory Controller has to reside within the memory to allow the CPU to be agnostic to the specific type of memory.
- A superset of the Hybrid Memory Cube standard can fit the bill.

SSRLabs' vlcRAM

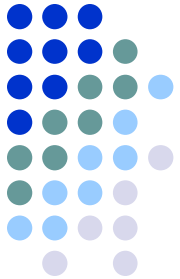

- A novel memory architecture based on a superset of the Hybrid Memory Cube standard.
- Removes the need for any memory (i.e. DRAM) controller in the CPU.
- Swaps out a dedicated set of pins or balls for DDR3, DDR4 or HBM interfaces for a much denser set of trunked High Speed Serial Links with better performance, better power consumption and a completely agnostic versatile Host Port.
- Moves the memory controller(s) into the memory ASIC and thus allows hierarchies of memory to be used to enable the desired density, size, performance and power consumption.

Current Server System

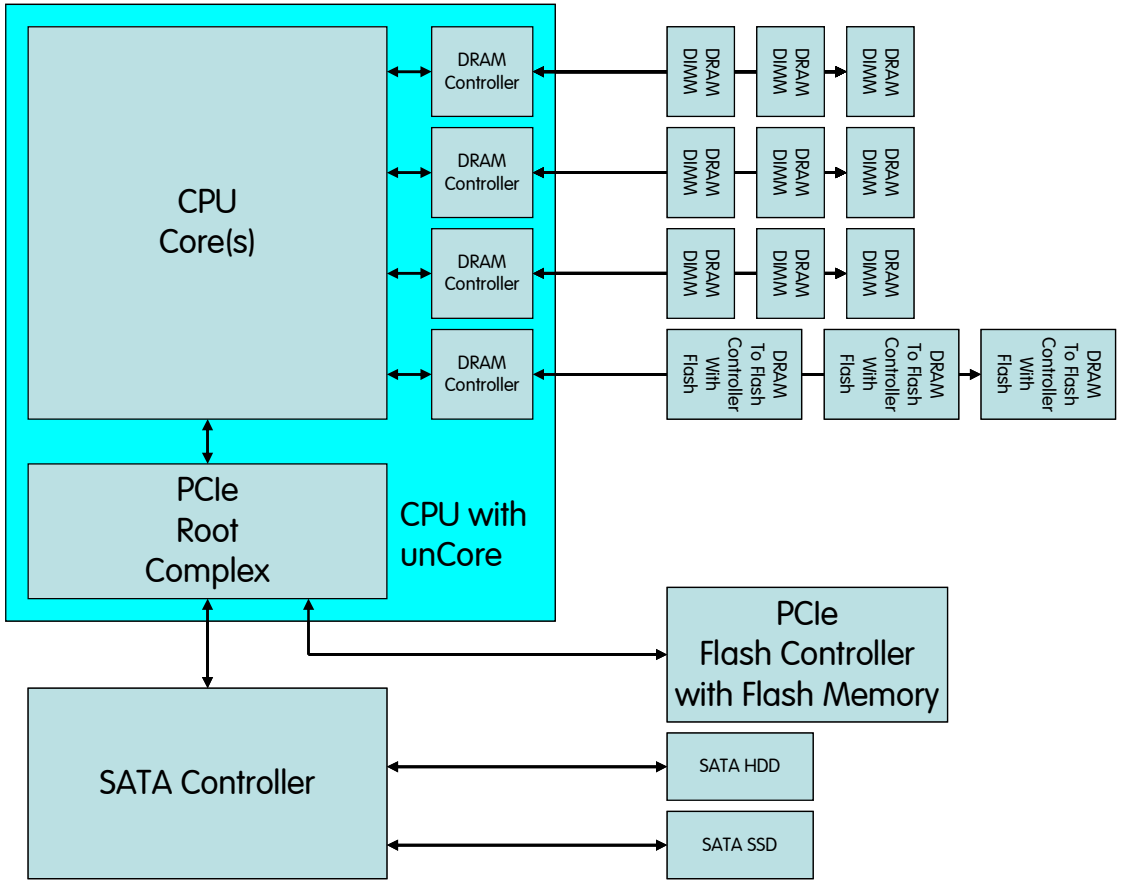

Novel Server System

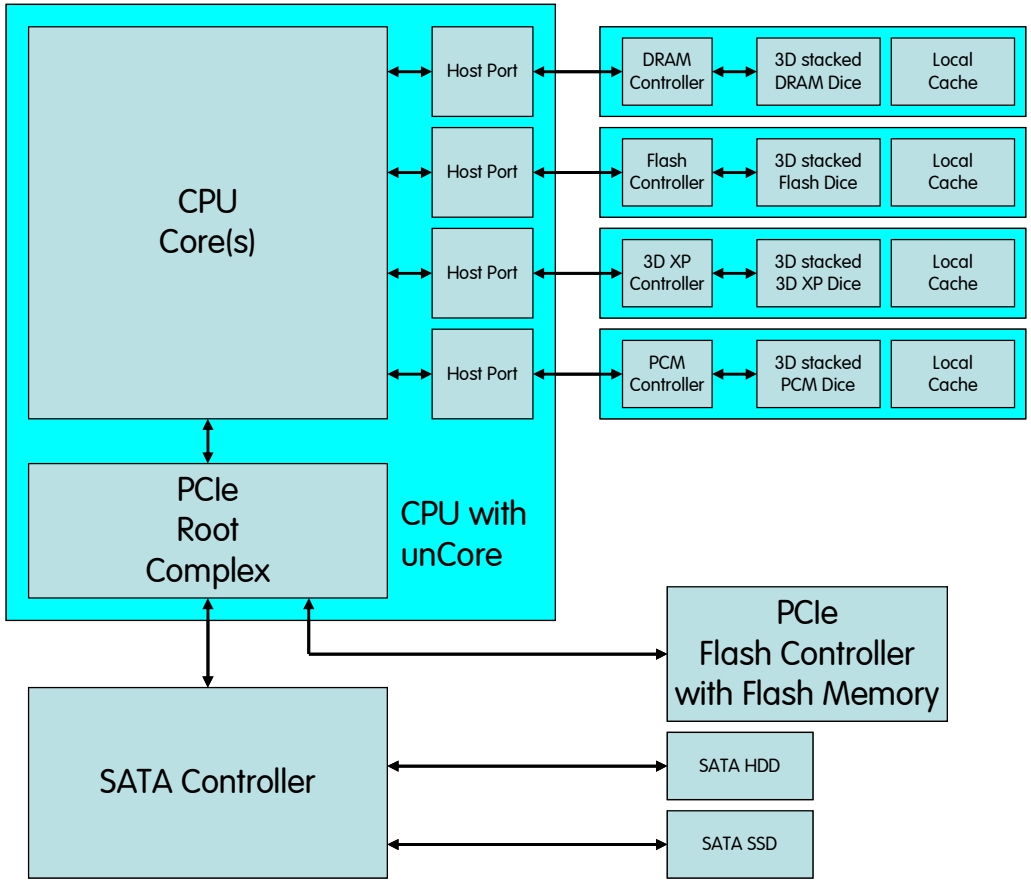

vlcRAM Benefits

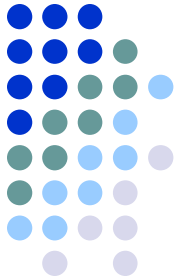

- Memory is completely independent of the CPU.
- CPU is agnostic to the memory and memory type.
- Local caching on the memory improves performance in HPC and Big Data applications as each memory module may add its own local Cache and thus is additive to the CPU Cache.
- Lower locality has less impact on Cache hit and miss rates because there is simply more Cache.
- Much denser memory is possible compared to DRAM.

Future-Proof Design

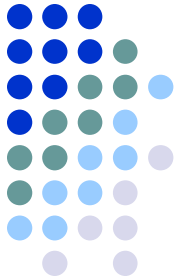

- Even if Flash Memory or 3DXP is used, the memory remains Random Access Memory (RAM) versus File System only Memory on PCIe Flash -> less overhead, lower latency, fewer protocols to traverse or convert.
- Future-proof and naturally able to cope with variable latency, i.e. no patented extension to DDR4.
- In-order (as a TCAM replacement), RMW (atomic cycles for semaphores) and deep OOO fully supported.
- Works well even if 3DXP and Phase Change Memory (or Memristors) become economically feasible and start to replace DRAM and Flash and a combination thereof.

Performance Data

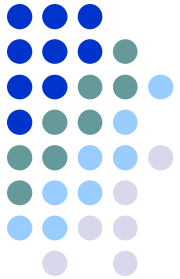

- 60 GB/s FDX versus ~ 17 GB/s HDX per port.
- 208 pins/balls versus ~ 80 pins/balls per port.
- Realistically over 4 times the bandwidth over DDR4 per port and over 12 times the performance over DDR4 per ball/pin on each memory port.
- Same memory power consumption against DDR3/DDR4, but roughly 50% less power per port itself (not including the memory array).
- Reduction of power consumption in the CPU where it really counts.

Applications

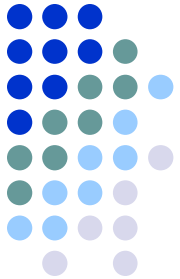

- Big Data
  - A single 1U device can hold two processors with 8 HMC ports each. Each port supports 3 vlcRAM memories of 512 GB each. Total memory in a 1U device therefore is 2 * 24 * 512 GB, or 24 TB.
  - A standard 42U rack therefore holds a maximum of 42 * 24 TB = 1008 TB, or about 1 PB if SSRLabs' vlcRAM is used.
- HPC
  - More focus on math acceleration, so same number of vlcRAMs but more accelerators in a larger device.
  - Devices are 6U and have the same number of the vlcRAMs as in the Big Data device, so 24 TB in 6U and 168 TB in 42U. Contact us for the total PFLOPS in this device.

Summary

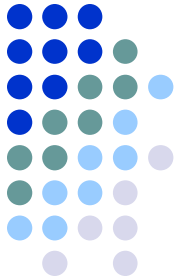

- Big Data and HPC requirements converge.
- Conventional memory solutions cannot fulfill the need.
- A novel memory architecture is required, and it should not involve a memory (i.e. DRAM) Controller in the CPU.
- Hybrid Memory Cube Technology is pretty close, except for density issues.
- An HMC Superset would do, particularly if it addressed the standard HMC shortcomings.
- SSRLabs' vlcRAM does all of that and allows for ~1 PB of RAM in a 42U rack. That memory can be used as RAM or as file system memory.