SAND2011-5219P

Memory for Exascale and …

Micron's new memory component is called

HMC: Hybrid Memory Cube

Redo memory parts and memory system architecture for HPC and ultimately for all systems above cell phones (and maybe those too :-)

Sandia is a Multiprogram Laboratory Operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy Under Contract DE-ACO4-94AL85000.

The Memory Wall and …

- Current DRAM roadmap will not enable achieving exascale systems with anything like the expected needs and goals
- Most all of the semiconductor industry is developing TSV (Through Silicon Via) technology which will greatly improve the energy outlook.
- Most efforts in TSV development for memory are aimed at cell phones and the like—not at HPC, so low power is main goal, not density, bandwidth, or high performance
- Micron has developed a revised memory architecture along with its TSV effort.

Micron's HMC

Uses TSVs to interconnect multiple memory die on a CMOS chip base

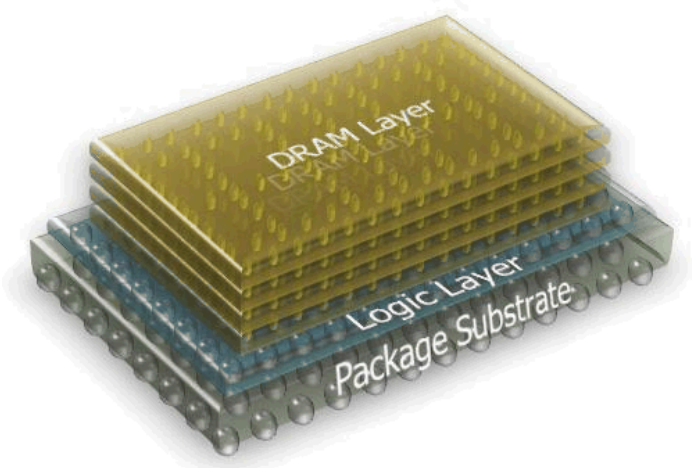

Micron has working demonstration components

HMC High-level View

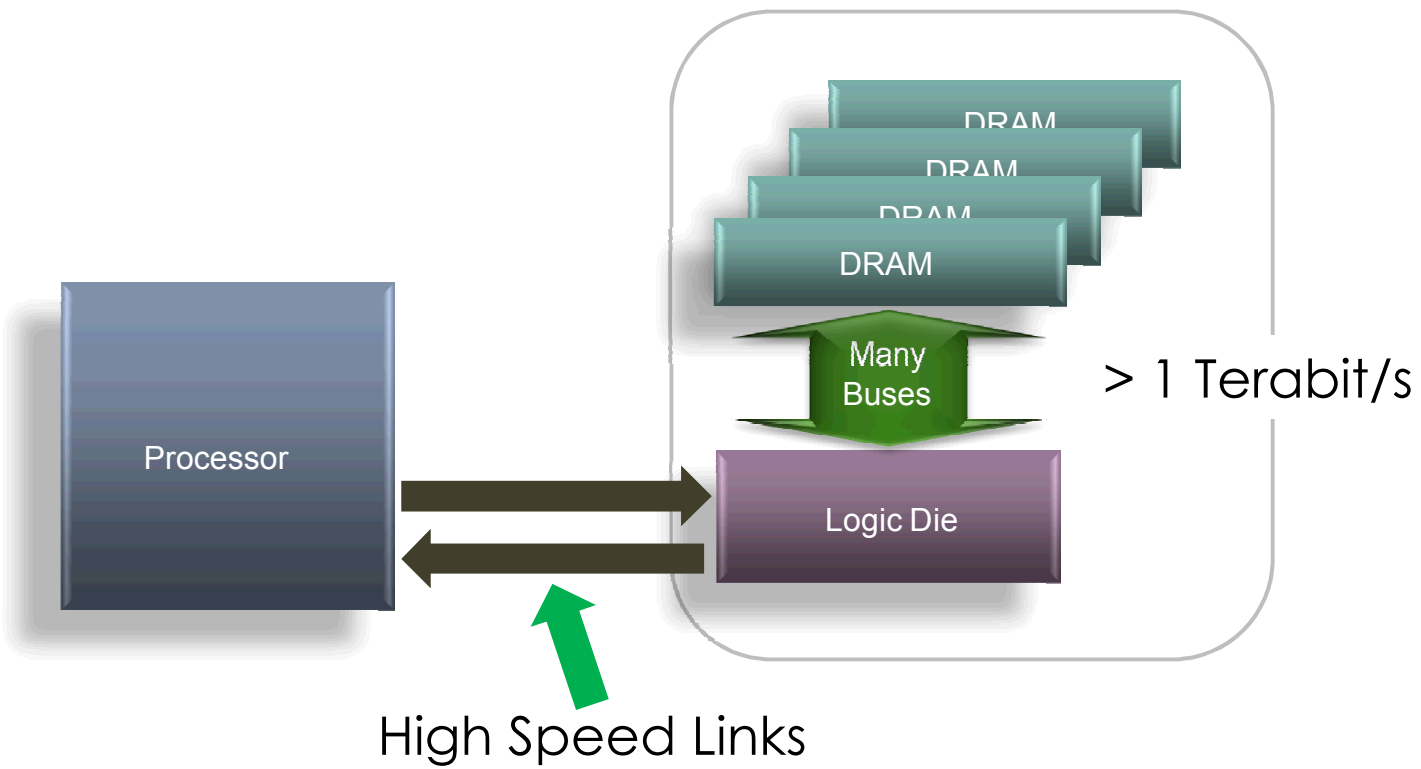

About HMC Technology

- Uses fairly standard DRAM cells
- DRAM controller in the CMOS base
- Internal memory architecture is changed to get higher bandwidth and other benefits
- Interface is totally different, having nothing in common with current DDRn implementations:

  Protocol, Electrical, Packaging

HMC Benefits

- 1 HMC has ~20 times the performance of a DDR3 DIMM (8 to 72 memory parts)
- 1 HMC uses ~10% of the energy per bit compared to current DIMMs and memory channels

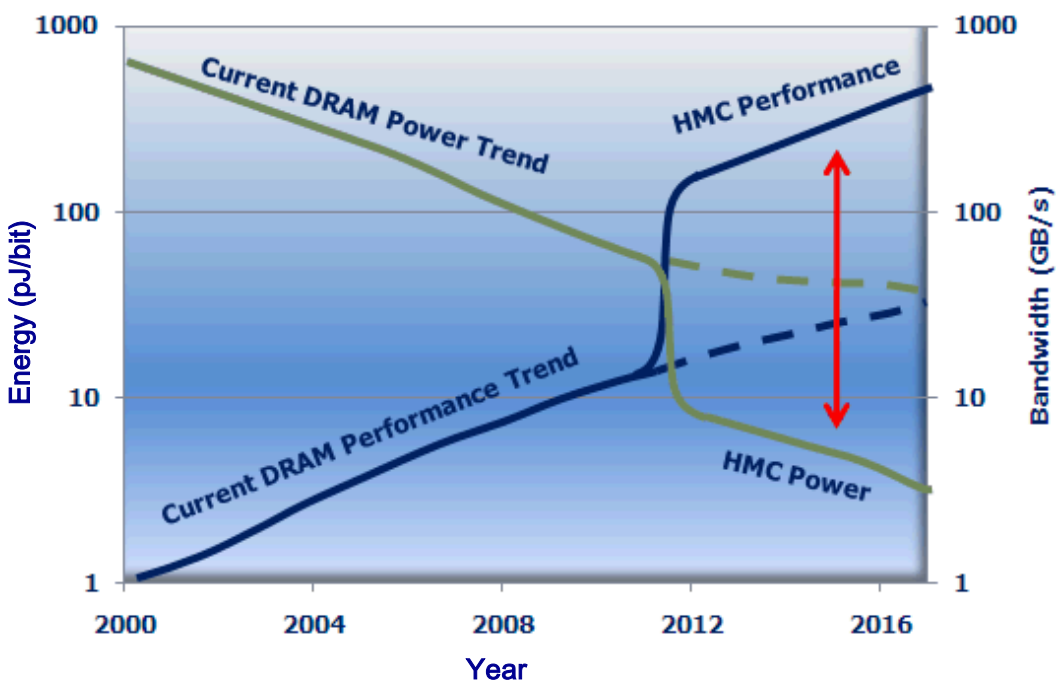

Developing Additional HMC Benefits

- Fairly easy to support new error correction and resiliency capabilities like "CubeKill" and "ModuleKill"
- Supports a wide range of network architectures — to: **"The memory system *IS* the network"** — and can help implement intelligent networks: operations within the network, self repairing,…
- The CMOS base offers great opportunity for new memory functions including including:
  - Gather/Scatter/Move,
  - Coherency,
  - Atomic operations;
  - ALU functions

Ongoing

- Sandia is developing a project, with NSA/ASC and DOE/ASCR support, to engage Micron in the development of HMC versions that become the main memory of high-end systems

Finally…

- Will want to expand the interface, protocol and other features of HMC technology to support non-volatile memory and other functions that can/will become how future systems are designed—for all system levels
- Breaking through the memory wall suggests that system implementations can/should be greatly improved
  - Intelligent Memory systems
  - Upgraded and new Network functions
  - Replace CPU caches
  - Latency tolerant CPU implementations