electronic design
The Authority on Emerging Technologies for Design Solutions

Hybrid Memory Cube Shows New Direction For High Performance Storage

William Wong

Fri, 2013-04-12 12:52

The final spec for the Hybrid Memory Cube (HMC) is done. Designed to improved DRAM performance by a factor of 15, the Hybrid Memory Cube Consortium (HMCC) standard specifies a high-speed, low-overhead memory protocol that targets the next generation of local storage for servers and networking devices. The architecture can employ 3D multilayer chips built using through silicon vias (TSV).

Related Articles

- Magnetic DRAM Arrives
- Non-Volatile DIMMs and NVMe Spice Up The Flash Memory Summit
- Understanding The Ivy Bridge Architecture
- Memory And Processor Advances Redefine Digital Technology

Memory layers typically are stacked on a logic layer that implements the HMC protocol. The first version of the standard supports up to eight memory layers with a maximum address space of 8 Gbytes. It also supports multiple cubes within a system to expand the amount of memory as well as the bandwidth available to the system *(Fig. 1)*. These cubes are designed to replace the ubiquitous dual-inline memory module (DIMM) found in systems now.

The HMC specification defines two physical-layer (PHY) configurations: short reach and ultra-short reach. The short-reach configuration addresses runs up to 8 in. It consumes 6 jJ/bit and runs at 10 Gbits/s or 12.5 Gbit/s. The ultra-short-reach configuration is designed for runs under 3 in. It is energy optimized so it consumes only 1 to 2 pJ/bit. It runs at 10 Gbits/s. Both use a common packet protocol.

The cube includes 16 pairs of links. There are four channels with four links/channel. Each channel has a 40-Gbyte/s bandwidth or 160 Gbytes/s/cube. The packet protocol allows messages to be forwarded through adjacent cubes. This increases latency, though developers are already utilizing non-uniform memory access (NUMA) architectures with multichip, multicore configurations. Multilevel NUMA is more challenging from a software standpoint, but it offers significant advantages. The packet protocol also provides limited atomic operation support.

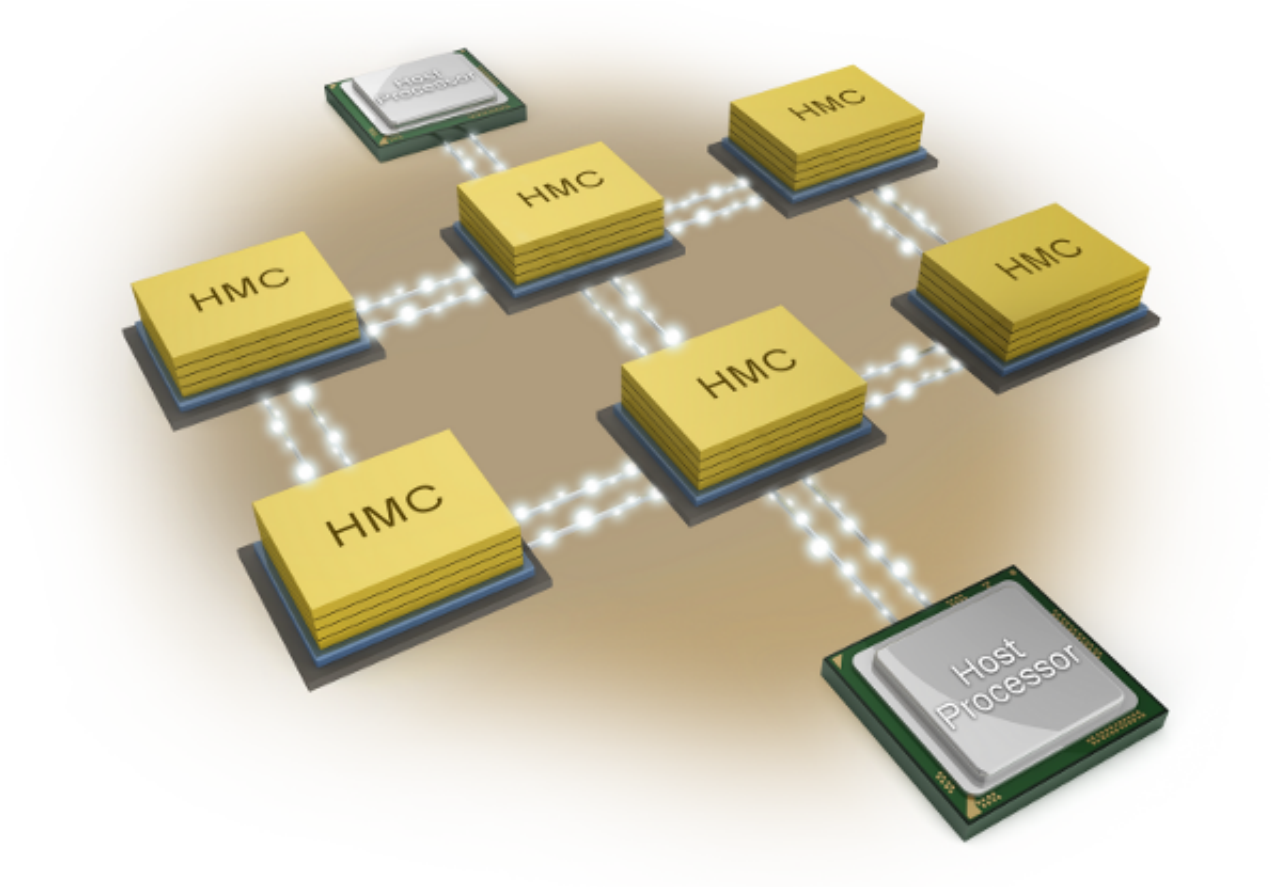

Stacking significantly reduces the footprint for memory. Switching to high-speed serial links reduces the number of connections and, along with linked cubes, greatly simplifies board layout. In theory, it would be possible to create a switch node without storage. Likewise, non-volatile memory could be added to the mix. For now, DRAM is where it's at. HMCC has the support of the major players in storage and computing including Micron, Samsung, and Hynix.

The standard will have a major impact on adoption and interoperability. Initially single HMC designs will dominate but multicube solutions will be in high demand. Of all the emerging memory technologies, HMC is the one that will likely have the most impact in the long run.

Source URL: http://electronicdesign.com/memory/hybrid-memory-cube-shows-new-direction-high-performance-storage