Denali Memory Report

Initial Hybrid Memory Cube short-reach interconnect specification issued to Consortium adopters

Posted on August 14, 2012

The Hybrid Memory Cube Consortium (HMCC), now supported by the three top DRAM vendors (Samsung, SK hynix, and Micron), has just issued an initial draft specification for the Hybrid Memory Cube's "short-reach interconnection across physical layers"—in other words, the short-reach PHY—to adopting members of the Consortium. As the Hybrid Memory Cube has evolved, it's become apparent that two PHY specifications would be needed: a short-reach spec for interconnections on FR4 circuit boards with trace lengths to perhaps 10 inches and a low-power "ultra short-reach" spec for interconnections in multichip modules and in 2.5D assemblies with trace lengths of perhaps 2-3 inches. The energy consumption for the short-reach PHYs is expected to be on the order of 5-10 pJ/bit and the energy for the short-range PHY is expected to be "much less" than 5pj/bit, according to Micron Technology Strategist Mike Black.

The initial target application focus for the Hybrid Memory Cube includes high-performance networking, industrial equipment, and test and measurement. Although these applications share many memory-specific requirements, they do have differing needs that are accommodated through settings in a mode register within the Hybrid Memory Cube. Details are still confidential but an example of an application-specific mode setting might be the serialized packet size used to communicate between the Hybrid Memory Cube and the host system.

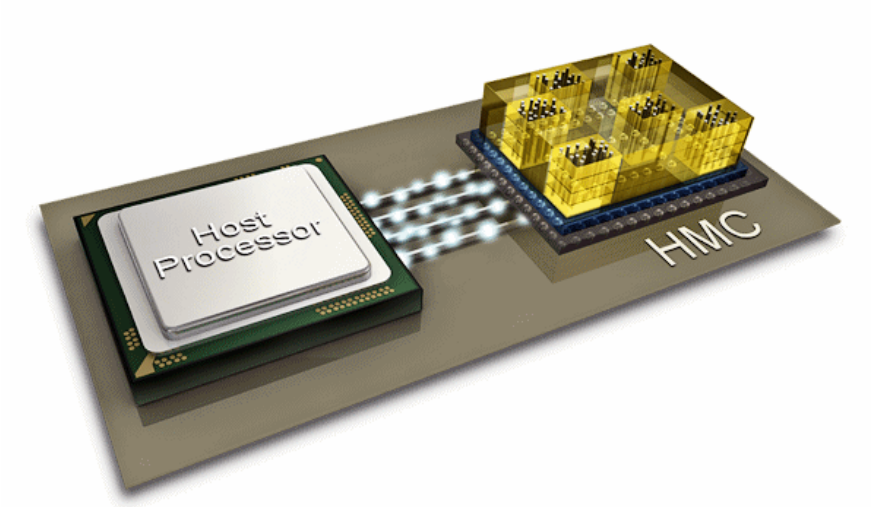

If you're not familiar with the Hybrid Memory Cube, it's a development initiated by Micron to extract much more of the inherent data parallelism in today's DRAMs, which are all based on multiple on-chip DRAM arrays. Conventional DRAM design funnels the I/Os from all of the on-chip arrays through one memory interface—a real bottleneck. Through 3D assembly, the Hybrid Memory Cube gives each of 16 on-chip memory arrays a clear path to a base logic chip where there are memory controllers and SerDes interfaces for high-bandwidth connection to the host system.

Here's an illustration of what a Hybrid Memory Cube might look like:

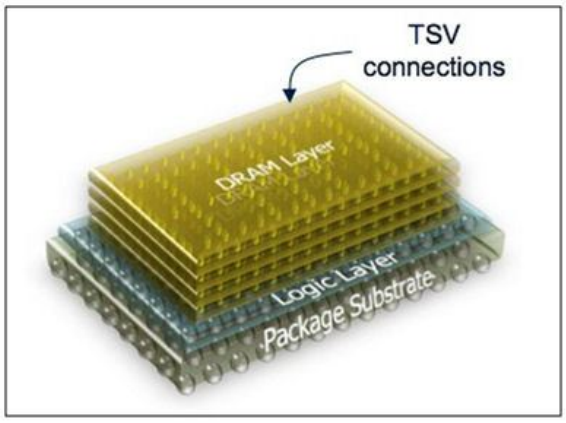

For more information about the Hybrid Memory Cube, see:

- 3D Thursday: Micron's 3D Hybrid Memory Cube delivers more DRAM bandwidth at lower power and in a smaller form factor using TSVs
- Want to know more about the Micron Hybrid Memory Cube (HMC)? How about its terabit/sec data rate?
- 3D Thursday: Hybrid Memory Cube—Does anyone know what's happening with IBM and Micron?
- Is 2012 going to be another breakout year for NAND Flash and Low-Power Design?

About sleibson2
EDA360 Evangelist and Marketing Director at Cadence Design Systems (blog at http://eda360insider.wordpress.com/)
View all posts by sleibson2 →

This entry was posted in 3D, DRAM, HMC, Hybrid Memory Cube, Hynix, Micron, Samsung and tagged Hybrid Memory Cube, IBM, Micron, MicronTechnology, PHY, Samsung, SerDes. Bookmark the permalink.

Denali Memory Report