

New Hybrid Memory Cube spec to boost DRAM bandwidth by 15X

Published: 02 April 13, 14:20 GMT

Backed by 100 tech companies, the three largest memory makers announced the final specifications for three-dimensional DRAM, which is aimed at increasing performance for networking and high performance computing markets.

Micron, Samsung and Hynix are leading the technology development efforts backed by the Hybrid Memory Cube Consortium (HMC). The technology, called a Hybrid Memory Cube, will stack multiple volatile memory dies on top of a DRAM controller.

The DRAM is connected to the controller by way of the relatively new silicon VIA (Vertical Interconnect Access) technology, a method of passing an electrical wire vertically through a silicon wafer.

An illustration of a Hybrid Memory Cube connecting over a bus to a CPU

Mike Black, chief technology strategist for Micron's Hybrid Memory Cube team, said what the developers did was change the basic structure of DRAM.

"We took the logic portion of the DRAM functionality out of it and dropped that into the logic chip that sits at the base of that 3D stack," Black said. "That logic process allows us to take advantage of higher performance transistors ... to not only interact up through the DRAM on top of it, but in a high-performance, efficient manner across a channel to a host processor.

"So that logic layer serves both as the host interface connection as well as the memory controller for the DRAM sitting on top of it," he added.

The first Hybrid Memory Cube specification will deliver 2GB and 4GB of capacity, providing aggregate bi-directional bandwidth of up to 160GBps compared with DDR3's 11GBps of aggregate bandwidth and DDR4, with 18GB to 20GB of aggregate bandwidth, Black said.

Jim Handy, director of research firm Objective Analysis, said the Hybrid Memory Cube technology solves some significant memory issues. Today's DRAM chips are burdened with having to drive circuit board traces or copper electrical connections, and the I/O pins of numerous other chips to force data down the bus at gigahertz speeds, which consumes a lot of energy.

The hybrid memory cube technology reduces the tasks that a DRAM must perform so that it only drives the TSVs, which are connected to much lower loads over shorter distances, he said. A controller at the bottom of the DRAM stack is the only chip burdened with driving the circuit board traces and the processor's I/O pins.

"The interface is 15 times as fast as standard DRAMs ... while reducing power by 70%," Handy said "Basically, the beauty of it is that it gets rid of all the issues that were keeping DDR3 and DDR4 from going as fast as they could."

The Hybrid Memory Cube consortium has defined two physical interfaces from the memory cube back to a host system processor: a short reach and an ultra-short reach. The short reach is similar to most motherboard technologies today where the DRAM is within eight to 10 inches of the CPU. That technology is aimed mainly for use in network applications and has the goal of boosting throughput from as much as 15Gbps per pin up to 28Gbps per pin.

The consortium plans to launch the short reach Hybrid Memory Cube in the second half of this year, and the ultra-short reach next year, Black said.

The ultra short-reach interconnection definition is focused on a low energy, close-proximity memory design support of FPGAs, ASICs and ASSPs, such as high-performance networking, and test and measurement applications. That will have a one to three inch channel back to the CPU, and it has the throughput goal of 15Gbps per pin

"It's optimised at very low energy signaling for multi-chip modules," Black said.

<http://www.computerworlduk.com/news/infrastructure/3438093/new-hybrid-memory-cube-spec-to-boost-dram-bandwidth-by-15x/>