Microsoft joins Micron memory cube effort

Rick Merritt

5/8/2012 10:00 AM EDT

Microsoft became the seventh core member of the Hybrid Memory Cube Consortium led by Micron and Samsung, a sign of the concept's broad technical implications. SAN JOSE – Microsoft became the seventh core member of the **Hybrid Memory Cube Consortium** led by Micron and Samsung, a sign of the broad technical implications for the concept of 3-D memory chip stacks.

Microsoft's participation signals the potential of the Memory Cube to drive changes in the traditional memory hierarchy and systems software for computers and networks. Micron has proposed a Cube of stacked memory die that includes a logic layer to optimize the placement of and access to memory, potentially handling functions carried out by memory controllers typically integrated in server and network processors today.

"The compilers and the whole system stack can be optimized [because] we can put additional functions out in our logic layer," said Mike Black, senior manager of business development for the Hybrid Memory Cube at Micron. "Having the logic layer manage some of your memory lets you localize data movement patterns that traditionally would be in systems memory or elsewhere," he said.

A Microsoft spokesperson was not available by press time. The Cube "represents a major step forward in the direction of increasing memory bandwidth and performance, while decreasing the energy and latency for moving data between the memory arrays and the processor cores," said KD Hallman, general manager of Microsoft's strategic software/silicon architectures group, speaking in a press statement.

"Harvesting this solution for various future systems could lead to better and/or novel experiences," she added.

Hallman oversees a small but significant semiconductor analysis group Microsoft created in recent years as it ramped up plans to support ARM in Windows 8. The group is led by Marc Tremblay, a former microprocessor architect for Sun Microsystems.

The group of seven Memory Cube developers—including Altera, IBM, Open-Silicon and Xilinx--is on track to deliver in June a draft specification for the Cube's interface. A broader but closed group of Cube adopters will review that draft. The consortium expects a final spec to be available to all interested parties by the end of the year.

With the developer's draft just weeks from being completed, Microsoft is formally joining the effort late in its process. "I hesitate to say they changed [the current version of the draft], but they got early look at it and have made valuable input," Black said.

Earlier this year, Black said he hoped to add two high-profile members to the consortium, presumably referring to Intel that showed last year work with Micron on the concept. There was no word on any other new member.

Packaging is another major issue under discussion in the Memory Cube group. "We don't want to design a package for every customer's desire, but a handful that addresses the majority of the market," said Black.

Micron says it will deliver early next year 2 and 4 Gbyte versions of the Cube providing aggregate bi-directional bandwidth of up to 160 Gbytes/second.

Separately, the Jedec standards group is working on a follow on to the 12.8 Gbit/second Wide I/O interface that targets mobile applications processors. The so-called HB-DRAM or HBM effort is said to target a 120-128 Gbyte/second interface and is led by the Jedec JC-42 committee including representatives from Hynix and other companies.