

Tweet 2 Like 4 Share

✉ 🖨 [Permalink](#)

[0 Comments](#)



Tags: [Hybrid Memory Cube](#), [chips](#), [RAM](#), [Hybrid Memory Cube Consortium](#)

Channel: [Servers & Storage](#), [Networking & Mgmt](#)

[See more from this blogger](#)

Hybrid Memory Cube Takes RAM to the Third Dimension

Posted by [Howard Marks](#)
May 14, 2012

As Intel and AMD pack six, eight or more cores into each processor, servers are once again struggling to move data from main memory in and out of those cores fast enough to keep them fully occupied. The [Hybrid Memory Cube Consortium](#), led by Micron, Samsung and Intel, has come up with a memory package that promises higher CPU-to-memory bandwidth by extending integrated circuit technology to the third dimension.

Formerly known as Hyper Memory Cubes, the new [Hybrid Memory Cubes](#) promise memory bandwidth up to 1 Tbps, more than 10 times what today's DDR-3 can deliver, while using about one-eighth the power per gigabyte. The cube is a hybrid package that stacks four or eight DRAM chips on top of a base-level memory controller chip.

MORE INSIGHTS

Webcasts

- [High-Frequency Trading: The Good, The Bad and The Ugly](#)
- [Protecting End Users Against Emerging Threats](#)

[More >>](#)

White Papers

- [Solve the 6 Top Problems in Your Data Center](#)
- [Drive Your Business with Predictive Analytics](#)

[More >>](#)

Reports

- [Fundamentals: 5 Options for Mobile App Development](#)
- [Strategy: Using Google to Find Vulnerabilities](#)

[More >>](#)

- [Panzura Cloud Storage Gateway Refresh Delivers More Control](#)
- [Whiptail Reinvents Performance Flash](#)

- [Commentary](#)
- [On the Web](#)

- [Consolidation Hits All-Solid-State Market](#)
- [Time To Say Goodbye To Static IPs](#)
- [Prediction: OpenFlow Is Dead by 2014; SDN Reborn in Network Management](#)
- [VMware's Lost Source Code: Not A Panic Situation](#)

[More Commentary >>](#)

Stacking the chips in a small package has several advantages. First, the stack keeps the interconnections between the memory chips and logic chip significantly shorter than they could be on a more conventional DIMM. At multigigahertz-frequencies, distances--even those measured in just a few inches--matter, since data signals are limited by the speed of light to traveling about one foot per nanosecond.

Using a stack rather than a single chip required the Hybrid Memory Cube designers to develop what they're calling silicon through-vias, which provide data paths vertically through the extra-thin silicon chips that make up each layer. Without silicon through-vias, signal paths would have to extend to the edges of each chip, where a massive number of interconnection wires would have to be connected. This would slow performance and make the whole thing too complex to manufacture cost effectively.

Most significantly, it allows designers to use completely different integrated circuit manufacturing processes for the memory controller logic chip and the DRAM chips themselves. Combining logic and memory on the same chip means the logic sections of the chip have to be produced using the memory chip process, which significantly limits their performance. A dedicated logic chip provides significantly more horsepower for ECC and other memory management.

We should see Hybrid Memory Cubes appear as preproduction samples some time in 2013, with the technology appearing as an extended processor cache in leading-edge servers, as vendors roll out their next generation of servers in 2014 to 2015. With power players like Micron, Intel, Samsung and most recently Microsoft in the consortium, odds are good that Hybrid Memory Cubes could be the solution to our memory bandwidth woes.

Disclaimer: Micron has provided SSDs for use in DeepStorage Labs.

Related Reading

- [News](#)
- [Most Popular](#)
- [Compuware Tunes APM Tools for Mobile, Cloud and Big Data](#)
- [Pure Storage SSDs Address VDI, Database Challenges, But Is Focus Too Narrow?](#)

[More News >>](#)