

- [AUTOMOTIVE](#)

Oct07

Tags

- [Architecture](#)
- [Hybrid Cube](#)
- [Memory chip](#)
- [Micron's](#)
- [Samsung](#)

Related Posts

- [14 Strange & Interesting Facts Worldwide](#)
- [Climbing Stalin's Red Gates Skyscraper In Moscow](#)
- [Apple has successfully blocked Samsung from selling its latest tablets](#)
- [Guggenheim, Bilbao](#)

Share This



Samsung and Micron's new Hybrid Cube Memory tech is greener and faster

Posted by [priya](#) on Oct 7, 2011 in [AUTOMOTIVE](#) | [0 comments](#)

[Ads by Google](#)

[Crystal Meth](#)

[Geek Support](#)

[Geek Tools](#)

[Samsung](#)

Like

Send

Share

tumblr

Delicious

Digg This Story



Memory chip giants Samsung and Micron have joined forces to create a new type of memory chip designed for high performance computing in a world with much faster broadband networks. The two firms said Thursday that they have formed the Hybrid Memory Cube Consortium to build a chip that can send information from memory chips to the CPU

cores 15 times faster than current memory technology.

To understand how cool this is, you have to understand the problem.

The big power problem

The chip industry is really good at making a CPU that does calculations faster, but it hasn't been able to make memory chips fast and dense enough to feed the cores enough information to keep up with the CPU's capabilities. So chips are left with is a massively large brain that stands idle sometimes while it waits for information to come to it. That idle time burns power and reduces the overall performance of a computer — and it's becoming a bigger deal as both power and performance are being pushed to the edge.

Many companies such as Microsoft, Intel, and even startups such as Tiler are looking for ways to solve the what the industry calls the memory bandwidth problem. Samsung and Micron's contribution with Hybrid Cube Memory is to deliver a chip that can send 15 times more data than a common DDR3 DRAM module used today. The consortium claims in a presentation that its new technology will use 70 percent less energy per bit than existing DDR3 DRAM technologies.

A dramatically different architecture

Some companies try to address the memory bandwidth problem by creating custom fabrics inside the chip to shuttle information around, while others basically create a giant caching system to pull information as rapidly as possible. The Hybrid Memory Cube guys are approaching it with density and a logic layer in a different type of architecture. The logic layer sits on the bottom, and the memory is densely stacked on top in a cube as opposed to a flatter architecture. Stacked memory isn't new either, as startups and large chip firms like IBM have tried that approach — but such a radically new architecture as a densely stacked cube has its pros and cons.

Rethinking the way memory chips are built allows for massive improvements in performance. That, in turn, will help drive faster supercomputers and help computers take full advantage of the coming improved speeds in broadband network bandwidth (pumping a gigabit connection to server doesn't help if the processor inside can't pull information in for processing at gigabit speeds). It will also help support the gear that is needed to deliver that bandwidth.

The chip is just the beginning

But, the chip industry is the first layer in a huge stack of hardware and applications, which means that architecture changes are hard to push without the consent of the folks on top of the stack. A chip's technology may be the most awesome thing ever, but to get it out of the niche market, big time equipment makers and software firms will have to demand the technology. A perfect case for this might be Infiniband, the networking technology that was supposed to take on Ethernet, but instead ended up as an awesomely fast networking technology relegated to the high performance market.

Luckily, the Hybrid Memory Cube guys are solving a huge problem that's been a pain point for the industry for a few years, and will only get worse. It's also starting in a niche market where the technology can be used first, become more widely adopted, and eventually put into production in a manner that might let it gain some economies of scale and drive down the cost of the chips. It's also trying to solicit folks across several industries to join the effort and embrace the technology. Altera [Corporation](#), Open Silicon, and Xilinx are working with the consortium to define a specification to enable applications ranging from large-scale networking to industrial products and high-performance computing to run on the chips. That specification should be out in 2012. Let's see if this licks the memory bandwidth problem, and

whether or not the industry elects to embrace it.



Leave a Reply

Your email address will not be published. Required fields are marked *

Name *

Email *

Website

Comment



You may use these HTML tags and attributes: <abbr title=""> <acronym title=""> <blockquote cite=""> <cite> <code> <del datetime=""> <i> <q cite=""> <strike>