Hybrid Memory Cube specification gets finalised

Published on 3rd April 2013 by Gareth Halfacree

News 9 Comments

The Hybrid Memory Cube Consortium, the industry group responsible for creating three-dimensional high-speed memory based on stacking layers of DRAM connected using through-silicon vias (TSV,) has officially announced the finalisation of the HMC Specification - opening the floodgates for manufacturers to start implementing the technology.

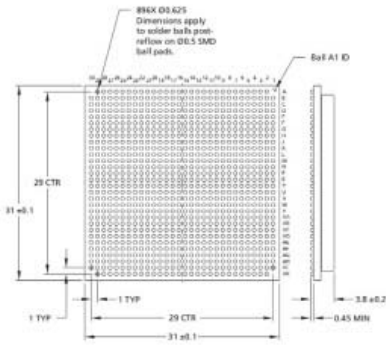

The Hybrid Memory Cube standard has been officially finalised, meaning manufacturers can begin to produce hardware offering ten times the data throughput of today's DDR3 DRAM.

The Hybrid Memory Cube technology is the work of a group of DRAM industry giants including Samsung, Hynix and Micron - the latter of which picked up an award for the technology in January last year, when its prototype modules hit an impressive throughput of around 128GB/s - and recently welcomed Microsoft into its ranks. So far, though, while prototype products have been demonstrating impressive potential, offering data throughput some ten times higher than traditional DDR3 memory, purchasable products have been thin on the ground.

The system works by utilising through-silicon vias (TSV) which allow layers of silicon to be built one on top of another and interconnected using conductive channels - 'vias' - that pass through each layer of silicon to the next. The result is a dense design that packs the capacity of multiple memory modules into the footprint of a single module, while offering extremely short signal paths between layers. In testing, Micron was able to show a boost in peak data throughput of ten times that offered by traditional two-dimensional planar memory construction, while simultaneously dropping the power draw by 70 per cent and the size of the finished module to one-tenth that of current products.

Doing that in the lab is one thing, but actually producing something people can pick up requires the various companies involved - which, in addition to those named above, includes the like of Altera, ARM, HP, IBM, Open-Silicon and Xilinx - to sit down and agree to a finalised specification that manufacturers can use to guarantee interoperability. Back in May last year, the Hybrid Memory Cube Consortium had promised that the finalised specification would be available by the end of 2012 - but it sadly missed its self-imposed deadline when the year ended with no sign of the specification.

Better late than never, the group has now released version 1.0 of the specification (PDF warning) offering manufacturers the chance to start gearing up to produce commercialised implementations of the hardware without fear of taking a wrong turn that will result in an incompatible design.

'*The consensus we have among major memory companies and many others in the industry will contribute significantly to the launch of this promising technology,*' claimed Jim Elliott, vice president of memory planning and marketing at Samsung's semiconductor arm. '*As a result of the work of the HMCC, IT system designers and manufacturers will be able to get new green memory solutions that outperform other memory options offered today.*'

'*This milestone marks the tearing down of the memory wall,*' added Robert Feurle, vice president of DRAM marketing at Micron. '*The industry agreement is going to help drive the fastest possible adoption of HMC technology, resulting in what we believe will be radical improvements to computing systems and, ultimately, consumer applications.*'

'*HMC is a very special offering currently on the radar,*' concluded JH Oh, SK Hynix's vice president of DRAM product planning. '*HMC brings a new level of capability to memory that provides exponential performance and efficiency gains that will redefine the future of memory.*'

A standard does not a product make, however, and while plenty of vice presidential types from HMCC's member companies were only to happy to provide soundbytes as to the breakthrough that has been made and the potential for HMC to change the face of computing forever, none were willing to suggest when the first modules will actually appear in consumer products. The group is already looking to revision 2.0 of the specification, however, promising to double data rates by the time the second specification is formalised in early 2014.

PREVIOUS ARTICLE NEXT ARTICLE

MORE ABOUT...

3d semiconductors dram hmc hmcc hybrid memory cube Hybrid Memory Cube Consortium memory planar through-silicon via tsv

RELATED ARTICLES POPULAR ARTICLES

- GTC 2013: Nvidia unveils Volta, Parker and Faceworks tech
- Intel and ITRI announce high-speed memory array
- Microsoft backs Hybrid Memory Cube tech
- GlobalFoundries announces TSV tech on 20nm
- Invensas to present xFD DRAM breakthrough
- Micron's Hybrid Memory Cubes win tech award

9 Comments

REPLY DISCUSS IN THE FORUMS

alpaca 3rd April 2013, 11:24

A standard does not a product make, young padawan.

Gareth Halfacree 3rd April 2013, 11:25

Quote:

> Originally Posted by **alpaca**
> A standard does not a product make, young padawan.

And where, pray tell, do I claim it does?

Samsung, Micron and Hynix have all committed to producing HMC hardware commercially. *That* makes a product, patronising Jedi. :p

Marquee 3rd April 2013, 11:44

I am glad to see the future is looking good. So for fast memory for the mass, will be great for gamers alike.

forum_user 3rd April 2013, 12:35

I can see the potential in these advances, but say it was to become available this year, how will our real experiences be improved? Will we see those 10x benefits? Would there be something in our machines to cripple that increase in memory speeds? Asking as a gamer, rather than folder :P

Kacela 3rd April 2013, 20:23

Hopefully, one of the manufacturers will re-implement an architecture utilized by Silicon Graphics in the 1990s; where the memory was itself the system buss, thereby having a throughput latency of near-zero. All components connected to each-other through the "memory buss".
I'm very excited to see if this prospect comes about with "3D" memory technology like this!

jb0 4th April 2013, 05:33

Quote:

> Originally Posted by **forum_user**
> I can see the potential in these advances, but say it was to become available this year, how will our real experiences be improved? Will we see those 10x benefits? Would there be something in our machines to cripple that increase in memory speeds? Asking as a gamer, rather than folder :P

It might not be immediately evident, but you WILL see a difference. RAM access is a HUGE bottleneck right now, just because RAM runs so much slower than the processor. It's why processors have large internal caches, and why cache misses are so devastating to performance.

Interestingly, in the dawn of time back in the 70s and 80s, RAM was far faster than processors. To the extent that some processor designs designated a portion of RAM to use as CPU registers because there was no speed penalty to it and you could implement an arbitrarily large number of registers that way.

fluxtatic 4th April 2013, 08:54

Quote:

> Originally Posted by **jb0**
> Interestingly, in the dawn of time back in the 70s and 80s, RAM was far faster than processors. To the extent that some processor designs designated a portion of RAM to use as CPU registers because there was no speed penalty to it and you could implement an arbitrarily large number of registers that way.

Now this I like the sound of. So, what are the odds something like this might become practical once the 2.0 or 3.0 spec comes along? I don't see processors getting hugely faster, and I have to wonder if Intel can keep the pace of ever-increasing IPC up...

At any rate, it's things like this that sort of make me hope I live to be 100, just to see what the tech landscape looks like (assuming I'm in any sort of shape to comprehend it by then.)

ModSquid 5th April 2013, 12:57

I heard that by the time we're a hundred, we'll have moved to socket type LB (Lamby Bridge), whereby the vat-cloned brains of lambs will be wired directly to the motherboards. System memory and processor on one chip (or chop, at least). Cooling will be provided by a small child with a palm frond.

Obviously, at that time, they will be back to working on miniaturisation.

yougotkicked 6th April 2013, 23:10

If I were to venture a guess, I'd say we could see this tech in the supercomputer market in 2-5 years. The performance benefits will be substantial, especially in environments where large amounts of data is being processed with poor locality, requiring the repeated loading of data sets from memory. The impacts in the consumer market may be a bit more muted, as engineers have worked hard on caching strategies over the years to try and take memory performance out of the equation whenever possible.

I think the biggest contribution this has to make for the consumer market will be higher memory density in SSD's. I know this spec is for volatile memory, and therefore not applicable to SSD's, but the refining of 3D circuit printing tech will surely spill over to other sectors in time.

LOG IN

You are not logged in, please login with your forum account below. If you don't already have an account please register to start contributing.

User Name:

Password:

LOG IN

DISCUSS IN THE FORUMS