By | Timothy Prickett Morgan 2nd July 2012 18:38

More revelers amp up hybrid memory cube party

ARM, HP (and more) join Micron, Samsung, Microsoft (and more)

A bunch more IT vendors have picked up shiny new Hybrid Memory Cube sledgehammers and are working with Micron Technology and Samsung Electronics to smash the memory barrier.

Memory performance, density, and lower power consumption are all key to future servers and other exotic high performance computing clusters as well as for graphics cards, networking devices, and other memory-dense devices.

To meet these needs, the Hybrid Memory Cube consortium was formed last fall to create a 3D stack of memory chips that has 15 times the performance of a single DDR3 memory module, uses 70 per cent less energy per unit of capacity, and occupies only a tenth of the space of current memory subsystems. The goal is to get the first 3D chip modules to market by 2013.

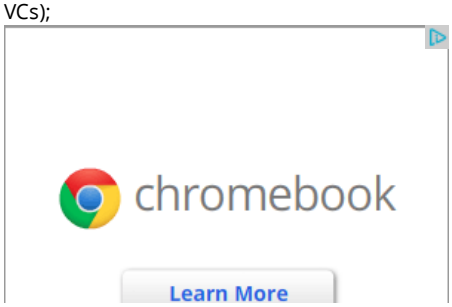

More Reading

Micron 'fesses up: Yes, we plan to eat limping DRAMurai warrior

Resistive Ram cache to make Flash fly, say boffins

New chippery on parade at ISSCC

IBM, Micron tag team on 3D memory breakthrough

Samsung, Micron bake 3D chips for next-gen RAM

Today's processors spin at several billions of cycles per second and have an insatiable need for data. As you add more cores to a processor, even if you don't ramp up the clock speed, you make the memory bandwidth problem worse, seeing as how current memory modules have long traces between memory chips and their pins, which require lots of power and time to send a signal.

The answer that most chip makers have to this problem is to go vertical – to stack up chips in three dimensions. In the case of HMC memory, you stack up chips in cubic arrays and connect them in parallel to a logic layer and a package substrate – like the pin-out on a CPU – to build a structure that has shorter paths, can run faster, and uses less energy. Prototype HMC blocks that have been manufactured to date can deliver around 128GB/sec of bandwidth into and out of the memory block using DDR3 memory chips, compared to 12.8GB/sec for 1.33GHz DDR3 memory sticks made in a conventional 2D layout.

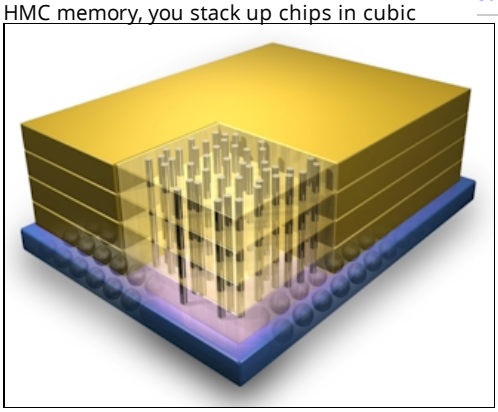

Micron and Samsung formed the HMC consortium along with field programmable gate array makers Altera and Xilinx as well as Open Silicon in October 2011, and in December the two primary vendors behind the effort said they had tapped IBM Microelectronics for manufacturing breakthroughs relating to the Through Silicon Vias, or TSVs, to link DRAM chips together and to controllers that implement a crossbar interconnect that play traffic cop across the cube of memory circuits.

Micron showed off HMC memory modules back at last years's Intel Developer Forum – even ahead of when the consortium was announced – and seems confident that it can get modules into the field in the second half of 2013.

This May, Microsoft joined the HMC consortium. The company has a direct interest in

memory technology because of its Xbox game consoles – and now its Surface fondleslabs – and an indirect interest because of its dominant position in the PC and server rackets and its status as a provider of raw computing from its Azure cloud.

At the time, the consortium said that 75 companies had contacted it with interest in joining the HMC effort and participating in specification development.

Now, ARM Holdings, which hopes to get its foot in the server door soon with its low-power RISC processors, has joined the HMC effort and so have server and switch maker HP and memory maker SK Hynix. The number of HMC tire kickers is now up to 90, according to a statement put out by Micron. ®