Denali Memory Report

Want another opinion about the Hybrid Memory Cube? Michael Feldman of HPCwire.com weighs in

Posted on July 11, 2012

Michael Feldman over at HPCwire.com has just published his own analysis of the Hybrid Memory Cube (HMC), which I've covered extensively in the **EDA360 Insider** and the **Denali Memory Report** (see below). Feldman reiterates many of the same points I've made in the past, but his perspective is one that comes from high-performance computing so I think it's interesting to see what he's saying.

Briefly, the HMC as currently envisioned is a 3D stack of DRAM die atop a logic chip. Each DRAM chip in the stack consists of 16 separate DRAM arrays and each array has its own RAS/CAS interface port. Stacking these DRAM die interconnects the DRAM arrays in the Z direction using through-silicon vias (TSVs) to create 16 parallel DRAM stacks.

Here's a diagram of the physical stack:

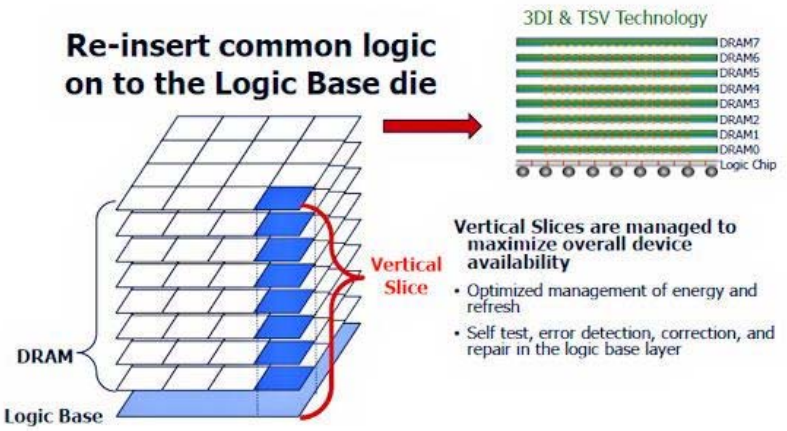

This stack then sits atop a "logic chip" that contains memory controllers for each DRAM array stack plus SERDES interfaces to connect the HMC to a host processing complex. Here's a block diagram of logical design of the HMC:

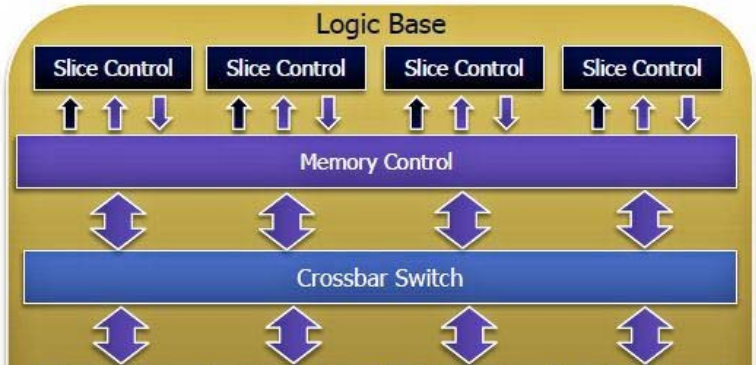

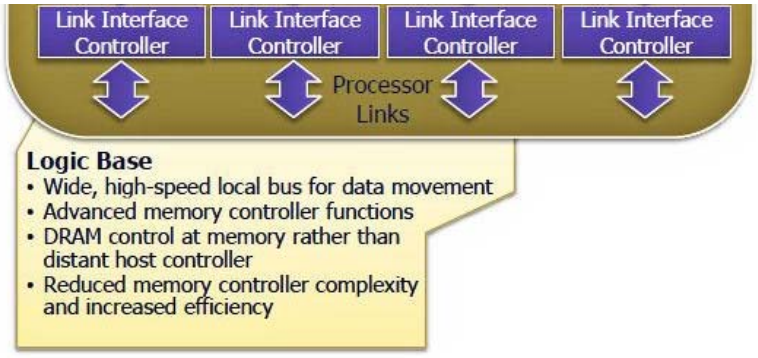

The reason Micron felt the need to take this design approach is because large DRAMs (and other semiconductor memories) are bandwidth-constrained by the limited number of interface pins. The consequence of pin limitations in memory design is the need to perennially push transfer rates into the GHz range. One way to solve this problem, taken by the JEDEC Wide I/O committee, is to adopt a very wide interface—512 data bits in the case of Wide I/O.

The HMC takes a different approach. Although it too employs wide interfaces to the memory die, the HMC itself presents a high-speed serial interface to the host processors. Feldman writes:

"The HMC design gets around those limitations by going vertical and using the TSV technology to parallelize communication to the stack of memory chips, which enables much higher transfer rates. Bandwidth between the logic and the DRAM chips are projected to top a terabit per second (128 GB/second), which is much more in line with exascale needs."

Then, Feldman makes an additional and very insightful comment:

"Another important aspect of the design is that the interface abstracts the notion of reads and writes. That means a microprocessor's memory controller doesn't need to know about the underlying technology that stores the bits. So one could build an HMC device that was made up of DRAM or NAND flash, or even some combination of these technologies."

I'll just note here that the HMC concept isn't limited to existing semiconductor memory technologies. MRAMs and memristors, should they become commercially viable and competitive, are just as suitable in an HMC stack. That's certainly one of the advantages of the HMC design.

For additional insight into the HMC see:

3D Thursday: Micron's 3D Hybrid Memory Cube delivers more DRAM bandwidth at lower power and in a smaller form factor using TSVs

3D Thursday: Hybrid Memory Cube—wide I/O only more so—gets an industry consortium

Want to know more about the Micron Hybrid Memory Cube (HMC)? How about its terabit/sec data rate?

Is 2012 going to be another breakout year for NAND Flash and Low-Power Design?

ARM, HP, and SK hynix join Hybrid Memory Cube Consortium (HMCC). First spec due by end of year

Share this:

Like this: Be the first to like this.

About sleibson2

EDA360 Evangelist and Marketing Director at Cadence Design Systems (blog at http://eda360insider.wordpress.com/)

View all posts by sleibson2 →

This entry was posted in DDR, DRAM, HMC, Hybrid Memory Cube, Micron and tagged DRAM, Dynamic random-access memory, Flash memory, HMC, Michael Feldman, Micron. Bookmark the permalink.

Denali Memory Report

Theme: Twenty Ten *Blog at WordPress.com.*

Follow