MoSys Launches New DRAM Architecture

Memory Chips Will Compete with Rambus for Graphics Designs

by Michael Slater

Startup memory-chip company MoSys (previously known as Most) has unveiled its high-performance DRAM design, initially aimed at graphics applications. MoSys claims that its multibank architecture will deliver higher bandwidth and lower latency than competing alternatives, such as Rambus DRAMs (RDRAMs) or synchronous DRAMs (SDRAMs), and will cost less.

MoSys (San Jose, Calif.) is a 16-person firm headed by Chairman and CEO Fu-Chieh Hsu, formerly a VP at IDT; VP of Design Wing-yu Leung, formerly with Rambus and IDT; and VP of Marketing Gary Banta, formerly with National. IDT is a corporate partner and is providing office space and design facilities, as well as funding, but will not manufacture the chips. MoSys has three foundry and second-source partners, but their identities have not yet been revealed. Total funding to date is $7.5 million.

Unlike Rambus, which is strictly a technology licensing company, MoSys will operate as a fabless DRAM company, selling chips made by its foundries as well as licensing its designs to other chip makers to sell as alternate sources. MoSys believes that this will enable the company to establish its new architecture more rapidly.

A chip announcement is planned for the fall, with volume production late this year. Graphics controller makers pledging support for the architecture so far include Tseng Labs, Trident, and S3. Future products will expand the application range beyond graphics to include 3D games, network hubs and routers, and set-top boxes.

The heart of the MoSys architecture is its multibank design. Each chip is divided into independent banks, each of which is a complete 32-Kbyte memory. All the banks are bused together, and the only common interface logic is a bus repeater. By building an array of small, 256-Kbit DRAMs on a single chip using 16-Mbit technology, MoSys achieves not only high bandwidth but also low latency—the missing ingredient in most other DRAMs. The reduced latency results from the small bank size and proprietary circuit techniques. Column-access latency is 15 ns for reads and 6 ns for writes.

MoSys has not revealed the details of the interface but says that it is "SDRAM-like," is 16 bits wide, and uses a total of 26 pins. Various clock speeds will be offered, initially up to 166 MHz (6-ns period). Data is transferred on each clock edge, resulting in a peak bandwidth of 660 Mbytes/s vs. 500 Mbytes/s for Rambus.

The multibank design dramatically increases the likelihood of a page hit, so the fast column access will be used more frequently. In a single-bank design, every time the access stream changes (as the system switches tasks), it is likely to move to a different page. When the access pattern returns to an earlier stream, the page address has already been changed, so no hit occurs. In the multibank design, accesses can jump around from bank to bank, and a page hit can still occur within a bank on a subsequent access.

MoSys claims that this design, combined with its lower latency, allows a MoSys DRAM (MDRAM) to deliver 74% of its peak bandwidth, or 490 Mbytes/s, in typical applications. A Rambus chip, according to MoSys, delivers only 41% of its peak bandwidth, or 206 Mbytes/s, in the same applications (assuming 32-byte block transfers). Rambus contends that graphics applications tend to use longer transfers, for which RDRAMs deliver a larger fraction of their peak bandwidth.

The multibank design also increases yield, because banks can be individually disabled. This is especially valuable early in the learning curve. Initially, MoSys plans to offer three device capacities—512 Kbytes, 768 Kbytes, and 1024 Kbytes—but the three versions will be made from only two different silicon designs. The middle size can be built either with fully functional 24-bank chips or with partially functional 32-bank devices.

All common screen resolutions up to $800 \times 600 \times 16$ can be supported with one chip, or up to $1600 \times 1200 \times 8$ with two chips. MoSys claims that its chips will carry a price premium of only 10% per bit over standard ×16 DRAMs, and that the ability to match common graphics memory sizes with fewer chips often will result in memory system costs that are lower than DRAM. Rambus claims that RDRAM prices are 15–20% higher than commodity ×4 DRAMs (which are cheaper than ×16 parts), so the net result (according to Rambus) is that RDRAMs will be less expensive than MDRAMs. Rambus gains some cost advantage from its lower pin count.

MoSys and Rambus appear to be headed for a major battle over the graphics memory market, with VRAM as the clear loser. MoSys claims that it has the support of many graphics chip makers and that Rambus will be relegated to a small niche; Rambus makes the opposite claim. As Rambus has learned, getting design wins for a new DRAM architecture is a frustratingly slow process. Rambus has a considerable head start and the benefit of established partners and demonstrated technology; until MoSys announces its partners and begins to ship chips, the strength of its challenge won't be clear. ♦

