

# MoSys Reveals MDRAM Architecture

## Radical Alternative Is Not Just Another VRAM

by Steven Przybylski

MoSys Inc., a four-year-old start-up in San Jose (Calif.), has recently disclosed technical details of a new DRAM architecture that it first announced in the summer of 1994. The MoSys Multibank DRAM (MDRAM) is explicitly aimed at graphics and video applications. To meet the aggressive bandwidth, latency, and cost needs of these applications within the context of a small total memory size, MoSys has reengineered the DRAM interface at all levels: logical, temporal, electrical, and physical.

### A Radically New Multibank Architecture

MoSys's biggest departure from all previous architectures is the most fundamental. Traditionally, a DRAM or VRAM is logically organized as a single, monolithic bank. Some new DRAMs have two or four banks: conceptually, these DRAMs are equivalent to a small number of semi-independent DRAMs on the same die. MDRAMs, however, are constructed of a large number of small (32-Kbyte) banks. The largest MDRAM, at 10 Mbits, is divided into 40 banks.

There are several advantages to many small banks:

- Because the banks are physically smaller, row and column accesses are faster than those of large banks.
- The bits available via column accesses can be spread across more distinct regions of the screen or address space, so concurrent operations in separate portions of the screen can be completed almost entirely with column accesses.
- MDRAMs can, in theory, be any size that is an integral multiple of 32 Kbytes instead of being restricted to the traditional binary sizes.

This last advantage is perhaps the most significant for a target market characterized by nonbinary display resolutions, flexible off-screen memory requirements, and extreme cost sensitivity. Starting in early 1996, MoSys will offer 0.5-Mbyte, 0.75-Mbyte, and 1.0-Mbyte MDRAMs with 1.25-Mbyte devices to follow by midyear. For example, a  $768 \times 1024$  true color (24-bit) graphics system requires 2.3 Mbytes for the frame buffer plus some extra memory for off-screen storage. If 256K $\times$ 16 DRAMs and a 64-bit bus are used, the only workable memory size that accommodates this frame buffer is 4 Mbytes, constructed of two ranks of four chips each. However, with a MDRAM memory system of 2.5 Mbytes constructed of two or three chips, the total memory cost can be significantly reduced, despite a per-bit price premium.

### Fast Synchronous Bus

Externally, MDRAMs are connected to the graphics controller via a 16-bit bidirectional bus (ADQ0..15) that carries address and data, and to a collection of 11 clock and command signals (see Figure 1). The bus operates at frequencies of up to 167 MHz, though parts are binned for operation between 100 MHz and 167 MHz. With data being transmitted on both clock edges, the peak data transfer rate is 667 Mbytes/s at 167 MHz.

A transfer is initiated by presenting a column read or write command to the MDRAM on the four-bit command bus (V0..3) and an address on the ADQ bus at the rising edge of clock. All data transfers are burst transfers. However, unlike in other synchronous memories, the burst length is neither encoded into the command nor programmed into a register. In the MDRAM protocol, bursts are initiated and sustained by read or write commands on the command bus and terminated by explicit stop commands on that same bus; bursts can be any length between 4 and 128 bytes. A two-bit byte mask accompanies write data to provide byte-write capability and byte masking. The address sent to the MDRAM explicitly addresses the bank and row or column that is being accessed.

The external buses are received and buffered at one edge of the MDRAM and propagated down its center aisle, as Figure 2 shows. All banks tap into this central aisle and continuously monitor the bus, looking for commands with a matching bank address. Many commands,

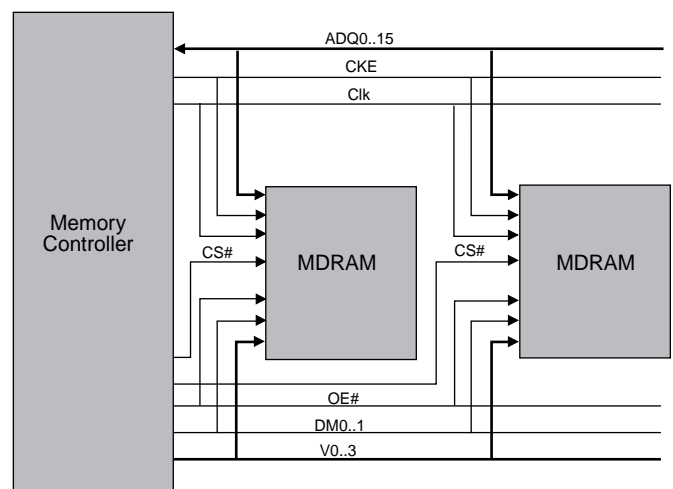


Figure 1. Up to four MDRAMs share a high speed, 16-bit synchronous address/data bus and nine clock and control signals.




Figure 2. The MDRAM block diagram shows many banks connected to an interface unit via a shared central bus.

such as Activate (i.e., row access) and Precharge, take a single cycle on the internal and external bus, freeing it for other commands and data transfers during the time the precharge or row access is being performed. Because the banks are independent and capable of simultaneous operation, minimizing the use of the bus by such commands inherently increases the opportunities for parallelism within the MDRAM.

DRAM designers have traditionally stayed away from multibank designs because of the area overhead associated with replicating the column decoders and all the associated timing and control logic. However, the die-size penalty due to MDRAM's plethora of banks is very small. One way that this area overhead has been kept low is by limiting the number of column address bits to only five. Each 32-Kbyte bank is logically 32 bits wide, with 256 rows and 32 columns. The bank overhead is also minimized by reducing the functionality of the banks. Each bank is capable of row accesses (Activate command), read and write column bursts, precharge operations, and little else. Refresh, for example, is entirely the responsibility of the graphics controller and must be performed via Activate and Precharge commands. The absence of commands to automate refresh is not a prob-

lem, since it takes only two cycles on the bus to initiate a refresh of one row in one precharged bank. Even at 100 MHz, only 1.3% of the total bus bandwidth is needed to entirely refresh 40 banks every 16 ms.

### High Bandwidth at Moderate Frequency

The goal of the MDRAM temporal, electrical, and physical external interface is to provide high sustained bandwidth on all lengths of bursts without unduly restricting or burdening the controller design or board layout. This is achieved through a combination of strategies. For example, high sustained bandwidth derives in part from low latency and minimal intervals in which the bus is not being used.

The MDRAM column-access latency ranges from 24 ns for the slowest grade to 12 ns for the fastest, and is programmed as an integral number of half-clocks. Row accesses are similarly fast: the mandated delay between an Activate command and a Read or Write command to the same bank ranges from 30 ns for the slowest speed grade to only 18 ns for the fastest. In other words, the slowest MoSys part is slightly faster than a 60-ns DRAM, while the fastest MDRAM is roughly equivalent to a 30-ns DRAM.

Most significant, since Activate and Precharge commands take a single cycle on the bus, the degradation in the sustained bandwidth due to the interference between these commands and read and write transfers is small: multiple independent banks provide opportunities for parallelism, and the protocol doesn't get in the way of exploiting that parallelism.

### Unusual Packaging and Electrical Interface

All MDRAMs are packaged in either a 128-pin PQFP or a 68-pin PLCC package. In both cases, the bus and control signals are clustered around one edge, minimizing skew. The surface-mounted PQFP package provides better electrical characteristics, while the PLCC facilitates in-the-field upgradability.

MDRAMs use a unique and nonstandard electrical interface.  $V_{ILmax}$  and  $V_{IHmin}$  are defined to be 0.5 V above and below one-half the I/O  $V_{CC}$ . All signals are parallel terminated to both  $V_{CC}$  and ground. The voltage swing generated depends on the amount of drive current and the termination resistors. For example, with standard 220-ohm resistors, generic 8-mA CMOS output drivers generate about 3 V of signal swing with a 5-V I/O supply voltage. The maximum operating frequency depends on the characteristics of the output drivers and the size of the termination resistors, while the maximum length of the bus depends on the operating frequency. At 167 MHz, a two-inch channel will support up to four MDRAMs. Lower operating frequencies allow more MDRAMs and/or greater voltage swing. Though the core supply voltage is 5 V for all current MDRAMs, there are

versions that support either 5-V or 3.3-V I/O supply voltages, and pure 3.3-V parts will be available in 1996.

Another important side effect of having many small banks is that a new level of redundancy is introduced into the DRAM. DRAMs have always had spare rows and columns to be fused into place at or after wafer sort to replace bad ones. In MDRAMs, the default bank address of the individual banks is also fuse-programmed, and whole banks can be fused in to replace bad banks. Therefore, a 1-Mbyte die with fewer than 32 good banks can be appropriately fused and sold as a 0.75-Mbyte or even 0.5-Mbyte MDRAM. Even though the yield on these relatively small DRAMs should be very good to start with, this ability should help to partly offset the tremendous economy-of-scale cost advantage currently held by the standard DRAMs.

Though the capability is not initially being used, each MDRAM package can have one or two die in it. The packages have two chip selects, however, the bank addresses on the die can be remapped dynamically to provide a contiguous block of memory without the use of chip selects. Only if there are more than 256 banks on the bus (8 Mbytes) do the chip selects have to be used to unambiguously address all the memory.

### Silicon Partners, Initial Design Wins

MoSys is a fabless DRAM company. It has contracted foundry services from IDT, Oki, TSMC, and Siemens. Though IDT is not currently a DRAM vendor, it has developed a competitive 4-Mbit DRAM technology. Parts from these foundries will be tested, packaged, and sold by MoSys. However, MoSys is also an intellectual-property company. Siemens is the first company to purchase a license to independently sell MDRAMs and will begin doing so by mid-1996. MoSys hopes that, by both licensing the technology and participating directly in the market, it can facilitate multiple sources of supply and ensure that the price reaches an attractive level so that GUI-accelerator and systems companies will adopt the technology.

None of these DRAM vendors is currently in the top ten by market share. Even with four sources of supply, the question remains whether, in this time of an industry-wide capacity shortage, these companies will be able to supply all of MoSys's needs if the architecture becomes widely adopted.

At last month's Comdex, Tseng Labs showed its new ET6000 GUI accelerator, which uses MDRAMs for its frame buffer and video store. The ET6000/MDRAM combination is expected to be available in GUI-accelerator boards from several sources in the first quarter of 1996. Other announced MoSys design wins are Lockheed Martin's R3D100 high-performance PC 3D graphics chip set and I-Cube's PCI switch. There is also the usual complement of unannounced design wins.

MDRAMs have had their problems coming to market. The technology was originally previewed in the summer of 1994. At that time, volume production was expected in the fourth quarter of that year. Subsequent press releases promised second-quarter and fourth-quarter 1995 volume shipments. As is normal in such circumstances, the announced list of companies designing with MDRAMs has ebbed and flowed somewhat over the past 18 months.

### How MDRAMs Stack Up

The big question is whether the world needs yet another graphics-oriented RAM. MDRAMs have a number of important characteristics that differentiate them from both the old devices and other new ones:

- Highest peak bandwidth out of a single device (667 Mbytes/s)
- Highest sustained bandwidth on a typical mixture of transfer lengths
- Low-latency row and column accesses
- Only solution with many banks and non-power-of-two memory sizes

With the standard solutions of wide EDO DRAMs and VRAMs becoming increasingly uncompetitive in either performance or cost, most of the future midrange and high-end PC GUI-accelerator market will be shared among several of the new entrants: synchronous graphics RAMs (SGRAMs), Window RAMs (WRAMs), Rambus DRAMs (RDRAMs), and now MDRAMs.

Compared with an SGRAM, an MDRAM provides lower latencies, higher sustained bandwidth, and smaller granularity with fewer control pins. SGRAMs also don't have as many banks as is desirable in a multimedia environment with up to half a dozen concurrent demands on the frame buffer. In SGRAM's favor are its forthcoming wide availability, its write-per-bit and block-write functions and large number of suppliers (see Table 1).

The comparison with Samsung's single-bank Window RAM (WRAM) is similar on most fronts. The MDRAM again provides better latency and bandwidth with fewer pins. The WRAM has an added serial port to send pixel data to a RAMDAC, which increases its effective bandwidth above that of SGRAMs but not to the level of MDRAMs. The WRAM also has high-bandwidth intra-chip aligned-data-movement capabilities that are not present in MDRAMs: to move a pixel from one location in the MDRAM to another, it must be read into the GUI accelerator, then written to its destination.

Both MDRAMs and Rambus RDRAMs are revolutionary in that a single bus is used to transmit address and data. In the Rambus case though, the bus also carries control information, while in the MoSys design, sideband signals are used instead. The MDRAM bus is wider (16 bits vs. 8 or 9) and slower (up to 167 MHz/667 Mbytes/s instead of 266 MHz/533 Mbytes/s). Despite

their many similarities, these two devices reflect fairly different philosophies:

- The Rambus approach moves some traditional memory-controller functions into the DRAM. RDRAMs are responsible for performing row accesses when appropriate and for deciding which address bits are bank, row, and column address bits. With MDRAMs, the memory controller is completely responsible for managing all aspects of the memory banks.
- The Rambus architecture was originally conceived as a main-memory architecture and has evolved to serve the graphics domain as well. From the outset, the MoSys DRAM architecture is specifically aimed at only the graphics and embedded domains.
- The Rambus protocol is more complicated than that of MoSys and includes a full complement of graphics-oriented writes and random-access reads.
- Efficiency of the RDRAM bus suffers if transfer lengths are short, while the MDRAM's lower latency and simpler protocol make for higher bus utilization with small and medium transfer lengths.
- Most significant, at the heart of all RDRAMs is a one- or two-bank DRAM core that is relatively unchanged from that of a standard DRAM. In contrast, MDRAM designers started with the premise of a new and radically different DRAM core.

### Tailored to Graphics

MDRAMs are highly tailored to graphics and video applications. Consequently, they are less than ideal for other applications. For example, the maximum amount of memory that can be addressed directly on an MDRAM bus is 8 Mbytes. To go beyond this requires using the chip selects or multiple buses. Though there is a low-power mode in which the on-chip PLL is turned off, there are no self-refresh capabilities: the memory controller must remain awake enough to keep the MDRAM refreshed if the frame buffer contents are to be preserved.

## Price & Availability

MoSys will be the primary source of MDRAMs. Beginning in early 1996 it will offer 0.5-Mbyte, 0.75-Mbyte and 1-Mbyte devices. A 1.25-Mbyte device will follow by mid-year. In addition to being a foundry, Siemens has licensed the technology and will be an independent vendor, also beginning in mid-1996. Initial pricing for 10,000-piece volumes will be \$21.25 for the 0.5-Mbyte device, up to \$50 for the 1.25-Mbyte MDRAM. As is standard for the industry, substantial high-volume discounts apply.

For further information, contact MoSys (San Jose, Calif.) at 408.456.2370; fax 408.321.0780, or via e-mail at [cognac@mosys.com](mailto:cognac@mosys.com).

Because MDRAMs are dramatically different from conventional DRAMs and VRAMs, a substantial change in the memory-controller design is needed to take full advantage of the memory's capabilities. For example, because there are so many banks, the standard strategy of keeping track of which row is open in each bank may not be the best choice. Instead, controllers that keep track of which banks are being used by each of several requesting functional units may be more effective.

The need for a major redesign of the memory controller can represent a significant hurdle, given the very short product life cycles in the graphics-accelerator marketplace. But once this design effort is complete, the result is a low-latency, high-bandwidth, pin-efficient memory technology that is technically very well suited to its target marketplace. ♦

*Steven Przybylski is an independent consultant based in San Jose (Calif.) specializing in DRAM architecture and memory-system design. He is the author of the research report New DRAM Technologies, published by MicroDesign Resources.*

	MDRAM		x32 SGRAM		WRAM		RDRAM	
	Fast (-166)	Slow (-100)	Fast (-12)	Slow (-15)	Fast (-60)	Slow (-80)	Fast (533MBps)	Slow (500MBps)
Column access latency	12 ns	24 ns	36 ns	45 ns	25 ns	35 ns	37.5 ns	40 ns
Row access latency	30 ns	54 ns	70 ns	90 ns	60 ns	80 ns	120 ns*	128 ns*
Peak bandwidth	664 MBps	400 MBps	333 MBps	264 MBps	200/142 MBps	132/90 MBps	533 MBps	500 MBps
Sustainable bandwidth	350– 550 MBps	210– 333 MBps	200– 333 MBps	176– 264 MBps	66– 200 MBps†	50– 135 MBps†	177– 405 MBps	166– 380 MBps
Controller pins	Under 40		About 90		About 115		31	
Effective granularity	256 Kbyte		1 Mbyte		1 Mbyte		1 Mbyte	
No. of banks per 1 Mbyte	32		2		1		1	
Graphics features	None		Write-per-bit; Block write		Aligned Move; Fill; Masked Writes; Serial-Port		Random Read; Bit & Byte Masked Write	
Number of vendors/ sources of supply	2/4		>8		1		6	

Table 1. MDRAMs have more banks and shorter access times than the other new DRAMs for graphics applications. MoSys has deliberately traded graphics features for raw sustainable bandwidth and good pin efficiency. \*Includes precharge time. †Random port only.