DRAMs For New Memory Systems (Part 3)

New Designs to Play Strong Role in Next-Generation Systems

By Steven Przybylski, Consultant, San Jose, CA

*The first part of this series (see **070205.PDF**) described the new evolutionary CDRAM, EDRAM, and SDRAM designs, while the second part (see **070304.PDF**) focused on the more revolutionary Rambus and RamLink alternatives. In this third and final part, we compare these five devices to standard DRAM offerings in a system context.*

Several new DRAM architectures have been proposed over the past year to meet the increasing bandwidth requirements of modern processors. Meaningful comparison of these alternative designs requires consideration of their operation in a complete memory system, not just as standalone parts. Table 1 compares typical 8 Mbyte memory systems for a medium- to high-end PC, each built using one of the different DRAM parts. The comparison assumes that memory system cost is important but not paramount. For example, to reduce cost and design complexity, no interleaving between multiple banks was considered.

This comparison is based on currently available parts in the 4-Mbit generation, except for the SDRAM and RamLink memory systems. Although only 16M SDRAMs are currently sampling, the comparison assumes a hypothetical 4M part running at 66 MHz with a low-voltage (LVTTL) interface. (The JEDEC specification does not include a 4M definition, but some DRAM vendors are contemplating 4M parts.) For the RamLink design, the definition is still preliminary and no DRAMs have been announced. For the basis of comparison, we postulate a generic 4M RamLink part with a 60-ns DRAM core and a 2-ns per-slave ring delay.

One source of uncertainty in this comparison is the estimate of the maximum frequency at which the various parts can operate reliably in a noisy system environment. Clock distribution, PCB layout, supply decoupling, and signal quality all become more critical as clock rates rise. TTL, and to a lesser extent LVTTL, parts are particularly susceptible to significant delays in address/control distribution and data bus settling. Though clever engineers could undoubtedly design faster memory systems, we assume that fairly standard design practices are used, limiting operation to roughly 60–80 MHz (depending on the design parameters) regardless of the availability of higher-speed DRAMs. It is widely acknowledged that higher operating frequencies rely on changing the electrical interface to use terminated signals with small voltage swings.

Cache/Buffer Hits

Memory system performance depends on many factors, including the primary cache miss rate, the secondary cache miss rate and speed, and the average read and write access times of the main memory. DRAM cache organization ranks highly among these factors. For all caches, size is by far the most significant organizational variable since it strongly impacts the cache miss rate; for most general-purpose workloads, doubling a cache's size decreases its miss rate by about 30%. (For a DRAM cache, the size is the product of the cache per DRAM and the number of chips in the memory system.) On average, increasing the cache size from 4 Kbytes to 32 Kbytes results in one-third as many row accesses.

The large size of the Rambus and CDRAM caches can therefore compensate for the large differences between their cache-hit and cache-miss access times. In contrast, the EDRAM and SDRAM memories rely on fast DRAM cache miss accesses—rather than on sustaining fewer misses—to achieve high performance.

After cache size, the second most significant organizational characteristic is the cache's block size. For the cache sizes involved here, the block size that minimizes the cache miss ratio is on the order of 128 to 1K bytes, depending largely on workload and primary cache characteristics. Smaller caches perform better with smaller blocks because when a cache has too few blocks, unrelated consecutive references are more likely to compete for use of the same cache block.

The DRAM caches are different from most caches in that the cache miss penalty is constant regardless of the block size, since the very wide data path from the DRAM array to its cache can fill a block in a single array access. Consequently, unlike most caches, DRAM-resident caches yield the best performance with the block size that minimizes the cache miss rate*.

The CDRAM is unique among the alternatives in that its block size is quite small (64 bytes versus up to 4 Kbytes in other memory systems). It is the only part in which the block size is dramatically less than the number of sense amps active during a row access. In contrast, the RDRAM's 1-Kbyte block size corresponds to the number of sense amps in one bank of one part; it is the only alternative in which every bit retrieved out of the DRAM array is saved in the row cache.

With the exception of the CDRAM, all the caches are direct-mapped. The CDRAM's cache can be configured as either direct-mapped or set-associative, since the

*See S. Przybylski, *Cache and Memory Hierarchy Design: A Performance-Directed Approach*, Morgan Kaufmann, 1990, Section 4.3.

	Conventional		Evolutionary			Revolutionary	
	Generic DRAM	Wide DRAM	SDRAM	Mitsubishi CDRAM	Ramtron EDRAM	Rambus RDRAM	RamLink
DRAM Organization	1M × 4	256K × 16	512K × 8	1M × 4	1M × 4	4.5 Mbit	4 Mbit
Memory Organization	2 banks of 8 chips	8 banks of 2 chips	4 banks of 4 chips	2 banks of 8 chips	2 banks of 8 chips	1 Rambus with 16 chips	1 ringlet with 16 chips
Expansion Mechanism	×32 SIMMs	×32 SIMMs	×32 Buffered SIMMs	×32 Buffered SIMMs	×32 Buffered SIMMs	RDRAM & RModule sockets	Unspecified
Assumptions	Single bus	4 banks per bus	LVTTL interface 2 banks per bus	Separate bus for each bank	Separate bus for each bank	No transceivers	60 ns DRAM access
Peak Transfer Rate	32 bits @ 33 MHz = 1.1 Gbps[1]	32 bits @ 33 MHz = 1.1 Gbps[1]	32 bits @ 66 MHz = 2.1 Gbps	32 bits @ 66 MHz = 2.1 Gbps[2]	32 bits @ 50 MHz = 1.6 Gbps[3]	8 bits @ 500 MHz = 4.0 Gbps[4]	8 bits @ 500 MHz = 4.0 Gbps
Cache/Buffer Size	8 Kbytes	8 Kbytes	16 Kbytes	32 Kbytes	4 Kbytes	32 Kbytes	Vendor specific
Cache/Buffer Organization	2 blocks of 4 Kbytes	8 blocks of 1 Kbyte	8 blocks of 2 Kbytes	Various	2 blocks of 2 Kbytes	32 blocks of 1 Kbyte	Unknown
Cache Hit Access Time (16-byte Access)[5]	4 × 30 ns = 120 ns	4 × 30 ns = 120 ns	(1 + 5) × 15 ns = 90 ns	(1 + 4) × 15 ns = 75 ns	(1 + 4) × 20 ns = 100 ns	80 ns	16 × 2 ns + 15 ns + 16 × 2 ns ≈ 79 ns
Cache Miss Access Time (16-byte Access)[5]	80 ns + 4 × 30 ns = 200 ns	80 ns + 4 × 30 ns = 200 ns	60 ns + (1 + 4) × 15 ns = 135 ns	70 ns + (1 + 4) × 15 ns = 145 ns	35 ns + (1 + 4) × 20 ns = 135 ns	184 ns[6]	16 × 2 ns + 60 ns + 16 × 2 ns ≈ 124 ns
Cache Hit Access Time (32-byte Access)[5]	8 × 30 ns = 240 ns	8 × 30 ns = 240 ns	(1 + 9) × 15 ns = 150 ns	(1 + 8) × 15 ns = 135 ns	(1 + 8) × 20 ns = 180 ns	112 ns	16 × 2 ns + 15 ns + 32 × 2 ns ≈ 111 ns
Cache Miss Access Time (32-byte Access)[5]	80 ns + 8 × 30 ns = 320 ns	80 ns + 8 × 30 ns = 320 ns	60 ns + (1 + 8) × 15 ns = 195 ns	70 ns + (1 + 8) × 15 ns = 205 ns	35 ns + (1 + 8) × 20 ns = 215 ns	216 ns[6]	16 × 2 ns + 60 ns + 32 × 2 ns ≈ 156 ns
Granularity of Expansion	4 Mbytes	1 Mbyte	2 Mbytes	4 Mbytes	4 Mbytes	0.5 Mbytes	0.5 Mbytes

1. Speed derated due to high fanout bidirectional bus.
2. Speed derated due to TTL interface overhead.
3. Speed derated due to asynchronous TTL interface.
4. Derated from 9 to 8 bits per cycle due to no parity.
5. Estimated
6. Includes precharge time.

Table 1. Comparison of hypothetical 8-Mbyte memory systems using various 4-Mbit DRAM designs. The memory systems are all 32 bits wide with no interleaving and no parity. All parts are currently announced or available except for the SDRAM, which is available only at 16M, and the RamLink part, which is a hypothetical design based on preliminary specifications.

tags are kept in the memory controller. For most programs, two-way set-associative caches have roughly 20% to 30% fewer misses than direct-mapped caches of the same size. Consequently, a CDRAM memory system configured with a set-associative cache of 32K would have the same performance as one with a 64K direct-mapped cache. Whether or not the implementation of set associativity within the memory controller imposes additional delay that erodes the benefit of the lower miss rate requires careful consideration on a case-by-case basis.

Of course, the access time for a cache/buffer hit is as important as the cache miss ratio in determining the overall performance. The access times shown in Table 1 are based on the fastest available speed grade and the expected number of cycles needed to fetch 16 or 32 bytes of data. An additional cycle of latency is added to the access time for the high-speed evolutionary designs to account for buffering of the address and data buses.

Table 1 shows that the revolutionary designs hold a small advantage in buffer-hit access time. For a cache hit, access latency is dwarfed by transfer time, so the higher bandwidth of the revolutionary designs translates into a smaller overall access time, especially for larger transfer sizes. All of the conventional or revolutionary designs, however, could match or beat the Rambus access time by changing to a wider memory organization. The cost of such a change would be increased granularity and more pins on the memory controller.

Cache/Buffer Misses

In the absence of a primary cache on the processor, these 8K–32K caches with large block sizes would have miss rates between 0.5% and 5% for most applications. As primary caches get larger, however, the ratio of DRAM cache misses to DRAM accesses increases dramatically, because the primary cache absorbs most of the locality in the reference stream. As a result, the DRAM's cache-miss access time is more important than would be indicated by the DRAM cache and block size alone.

Table 1 shows a range in the cache-miss access times of 124 ns to 200 ns for a 16-byte access, and 156 ns to 320 ns for a 32-byte access. As with cache hits, band-

	Conventional		Evolutionary			Revolutionary	
	Generic DRAM	Wide DRAM	SDRAM	CDRAM	EDRAM	Rambus	RamLink
DRAM Organization	4M × 4	1M × 16	2M × 8	4M × 4	4M × 4	16/18 Mbit	16 Mbit
Memory Size	16 Mbytes	8 Mbytes	8 Mbytes	16 Mbytes	16 Mbytes	8 Mbytes	8 Mbytes
Memory Organization	1 bank of 8 chips	2 banks of 2 chips	1 bank of 4 chips	1 bank of 8 chips	1 bank of 8 chips	1 Rambus with 4 chips	1 ringlet with 4 chips
Cache/Buffer Size	4 Kbytes	2 Kbytes	4 Kbytes	16 Kbytes	16 Kbytes	16 Kbytes	Vendor specific
Cache/Buffer Organization	1 block of 4 Kbytes	2 blocks of 1 Kbyte	2 blocks of 2 Kbytes	Various	4 block of 4 Kbytes	8 blocks of 2 Kbytes	Unknown
Granularity of Expansion	16 Mbytes	4 Mbytes	8 Mbytes	16 Mbytes	16 Mbytes	2 Mbytes	2 Mbytes

Table 2. Comparison of hypothetical 8-Mbyte (or 16-Mbyte, if that is the minimum size) memory systems using various 16-Mbit DRAM designs. These memory systems are 32 bits wide with no interleaving, no parity, and assume a medium-cost design.

width plays a significant role in determining the overall access time. Consequently, the impact of the EDRAM's very fast DRAM-array access time (35 ns) is somewhat muted. Of course, this advantage would largely be restored if either the transfer size were reduced, the width of the memory system were increased, or if Ramtron wrapped a high-speed synchronous interface around its fast core.

Despite its higher bandwidth, the RDRAM's access time is comparable to that of the evolutionary alternatives because its protocol requires two Rambus accesses, one to expose the row miss and another to retrieve the data. Also, current implementations only begin precharging the DRAM array when a row miss occurs. So although the other RDRAMs can be accessed during the precharge and row-access times, all of the precharge time is reflected in the cache-miss access time. The other alternatives hide some or all of the precharge time during idle time or underneath other transfers. In general, the CDRAM and EDRAM make it easier to hide this precharge time, facilitating accesses to the cache even after precharge of the array has begun.

DRAM Writes

The new designs' different ways of handling writes are especially important when the primary cache is a write-through cache. In this circumstance, up to 50% of the accesses reaching main memory are single-word writes. Again, the CDRAM and EDRAM designs are particularly noteworthy. Both allow writes into the DRAM array to optionally bypass the DRAM cache structures and complete in parallel with subsequent cache read operations. This is especially beneficial in graphics applications, since frame buffer writes can bypass the DRAM cache. Also, the EDRAM's fast core further accelerates the retiring of these writes.

For a write-back primary cache, the frequency of writes to main memory is greatly reduced, but those writes tend to be to different addresses than are being read. Consequently, these write-back references have poor hit rates in the DRAM cache. Worse yet, since the write-back data and the associated read reference share the same index address bits in the primary cache, the two can regularly collide and cause thrashing in a direct-mapped DRAM cache.

At least three of the alternatives (CDRAM, EDRAM and RDRAM) are being marketed as technologies that help eliminate discrete secondary caches. Though all provide at least comparable performance to a 256-KByte secondary cache backed by generic DRAM memory, the EDRAM is particularly successful in outperforming secondary caches in PCs. The key to its success in this area is the good match between the 486's write-though primary cache and the EDRAM's efficient write interface and high write bandwidth into the DRAM array.

The Impact of 16-Mbit Devices

In recent months, 16M SDRAMs have begun sampling, and details of 16M CDRAM and RDRAM parts have been released. Table 2 contrasts these new, higher-density parts and the memory systems that they can construct. As in Table 1, a medium-cost, low-complexity 8-Mbyte memory system is the basis of comparison.

Most significantly, all the ×4 organizations have a memory granularity, and thus a minimum size, greater than 8 Mbytes. The finer granularity of the revolutionary alternatives will be important both in small systems and in high-performance systems that require 64-bit data buses and/or interleaved memory arrays to attain the desired throughput.

Second, with the exception of the EDRAM, the amount of cache per megabyte of memory declines as the level of integration increases. Despite the increase in memory sizes over time, the total buffer or cache size will likely stay constant or even decrease slightly. The continued increase in primary cache sizes will reduce the effectiveness of these small DRAM caches. As performance requirements grow, this phenomenon will benefit those devices that have a larger amount of cache to begin with: CDRAMs, EDRAMs, and RDRAMs.

Designers of the evolutionary and revolutionary DRAMs are caught between the need to decrease the area overhead relative to generic DRAMs, and the desire to increase the amount of cache per memory bit. At the 16M level, the vendors have, without exception, emphasized decreasing the area overhead as much as possible.

High-Performance Interfaces

The common element underlying the new alternatives is an increase in the DRAM interface speed and effectiveness. The three basic techniques used to achieve these ends are:

- Larger or more useful buffers and caches
- Wider and/or synchronous interfaces
- Higher-frequency electrical interfaces

The revolutionary designs use all three approaches, while the evolutionary alternatives apply these techniques incrementally or less aggressively. As performance needs increase, these techniques will be used more frequently in the evolutionary and even conventional camps.

The Rambus and RamLink designs illustrate two different approaches to high-frequency design. In the RamLink solution, all signals are terminated, differential point-to-point links not limited by physical constraints. A consequence, however, is that *n* buses are needed to connect *n* slaves to one master. Each slave needs access to two differential byte-wide buses, as opposed to Rambus' one single-ended bus.

The Rambus approach of a physically-constrained bus links the maximum frequency of operation to the physical length and properties of the bus. In this environment, high frequency is attainable only through careful design and appreciation of the factors limiting reliable operation. As levels of integration increase and more circuitry can be placed close together, these fundamental limits will become less onerous. The flip side is that further increases in the frequency and bandwidth of a byte-wide Rambus-style channel necessarily involve further restrictions on the physical and electrical properties of the bus.

Interestingly, the same laws of nature that limit the length of a Rambus channel also limit the physical size of evolutionary memory systems. At 66 MHz, the maximum trace length in an LVTTL SDRAM system is comparable to the maximum Rambus channel length.

One problem with the combined address/data bus common to Rambus and RamLink is that addresses, control, and NACKs all take bandwidth. Consequently, sustainable bandwidth is highly dependent on factors such as transfer size, row-cache hit rate, number of DRAMs, and the controller's intelligence. For example, in graphics applications with good cache-hit rates and large blocks, well over 300 Mbytes per second can be achieved on a Rambus channel. On small random references, however, sustainable bandwidth can easily drop to 150 Mbytes/s. Admittedly, only the EDRAM would fare any better in this environment.

Prefetching of instructions and data from most memory systems made of generic DRAMs has little impact on performance. The long latency and mediocre bandwidth of generic DRAMs significantly inhibit the performance gains offered by these techniques. As primary cache sizes increase, however, the higher bandwidth of the evolutionary and revolutionary designs will increase the effectiveness of prefetching. The short latency of the EDRAM, the pipelined interfaces of the CDRAM and SDRAM, and the multiple independent banks of the RDRAM all greatly reduce the interference between prefetch references and demand misses.

Market Affinities

Clearly, no one alternative is superior to the others in all situations. Each approach has strengths and weaknesses that make it more or less appropriate for various applications.

Portable systems, for example, are sensitive to power consumption, physical size, and (to a lesser extent) cost. Typically, their main memories are smaller than in comparable desktop systems, and performance is not a

Price and Availability

NEC and Samsung are the only companies that have announced SDRAMs. Both are sampling 16M parts at up to 100 MHz, with production available in 3Q93. Neither company has announced production pricing, although Samsung expects its parts to carry a 20% premium over standard DRAM. Also, Oki will second-source the Samsung part.

Contact NEC at 401 Ellis Street, MS MV4570, Mt. View, CA 94039; 415/965-6002. Contact Samsung Semiconductor at 800/446-2760 or 408/954-7229.

Mitsubishi is shipping 4M CDRAMs at a price of $13 in quantities of 1000. Contact Mitsubishi Electronics America at 1050 E. Arques Avenue, Sunnyvale, CA 94086; 408/730-5900 x2226.

Ramtron's 4M EDRAM is currently available in ×1 and ×4 configurations at a price of $15 in quantities of 10,000. Contact Ramtron at 1850 Ramtron Drive, Colorado Springs, CO 80921; 719/481-7000.

RDRAMs with a 4.5-Mbit capacity are available from Toshiba, Fujitsu, and NEC. Pricing varies among vendors, but is about 25% above generic DRAM. 16/18 Mbit RDRAMs are expected to sample in 1Q94 and will eventually carry a 15% price premium over standard parts. Contact Toshiba America at 9775 Toledo Way, Irvine, CA 92718; 714/455-2000. Contact Fujitsu Microelectronics at 3545 N. First Street, Bldg. 1, San Jose, CA 95134; 408/922-9345. See above for NEC.

major concern. Consequently, wide generic DRAMs will continue to be popular in these systems for the next few years. Ultimately, a switch to Rambus or something comparable will be motivated by growing performance demands and the availability of Rambus interfaces on relevant processors. One drawback of the 4M RDRAM is its high per-bit power consumption; this problem should be largely remedied in the 16M generation, improving its applicability to portable solutions.

The low-end desktop market is almost entirely cost-driven. As a result, it will probably continue to use the conventional organizations for some time. In this domain, off-chip secondary caches are a marketing feature, not a performance driver. As a result, the argument proposed by Ramtron, Mitsubishi, and Rambus that their new architectures facilitate elimination of the secondary cache may meet with resistance. As with portable systems, PC manufacturers will not adopt a new memory technology until chip set designers incorporate the necessary interfaces.

For high-end PCs, some sophisticated buyers look beyond processor type and frequency when deciding on a purchase. The marketability of a "high-performance" memory technology may combine with a new design's cost/performance advantage to coax a daring vendor or two away from conventional approaches to one of the evolutionary or even revolutionary choices. The synchronous alternatives show significant performance gains over wide generic DRAMs only at frequencies above 50 MHz; the PC market, however, is monopolized by x86 processors with interfaces below that speed. In this arena, the move to new memory designs will be motivated only when new processor designs (such as Pentium) exceed this frequency.

The workstation market requires high performance with good cost/performance. Processors in these systems have wide interfaces operating at high frequencies. Especially in multiprocessor configurations, memory bandwidth requirements are increasingly formidable. Although typical main memories are several times larger than those in PCs, workstations require a higher bandwidth per DRAM. In addition, total system cost is higher than in the PC domain.

As a result, workstation manufacturers are willing to pay for large SRAM-based secondary caches. Their users, however, remain fairly sensitive to the cost of extra memory. Consequently, the workstation market will likely adopt the evolutionary and revolutionary techniques, especially SDRAMs and CDRAMs, if the cost premium can be reduced to 10% or so. Although workstations represent a comparatively small market for DRAMs, they may create a foothold from which these devices can gain wider acceptance.

A significant though often neglected market is the graphics arena. Despite a significant price premium, VRAMs currently account for about 10% of the DRAM market. RDRAMs could become a higher-performance, lower-cost replacement for VRAMs. If Rambus becomes the graphics/video memory of choice, the blurring distinction between video and main memory in multimedia systems may vault RDRAMs into greater prominence.

Conclusions

Increasing chip densities and main memory bandwidth requirements are making it more difficult to build small, high-performance memory systems from generic narrow DRAMs. Each of the alternatives to the generic DRAM changes one or more of the physical, electrical or logical interfaces in order to at least maintain current bandwidths per bit.

Wide generic DRAMs maintain the current interface while widening the data path. This safe route prolongs the life of the interface by a generation or two, but eventually a new interface will be needed. One option is the SDRAM; as a JEDEC standard backed by a variety of vendors, SDRAM is likely to gain acceptance in the workstation arena and may spread to other markets.

The CDRAM offers a highly effective combination of on-chip resources providing high bandwidth, flexible control of SRAM transfers, and a configurable cache. One disadvantage is that cache tags must be stored and managed in the memory controller; another is its status as a proprietary organization under Mitsubishi's control.

Ramtron's EDRAM provides good memory system performance at low frequencies without a secondary cache, and with a high degree of compatibility with existing memory controllers. The main disadvantages are its proprietary nature and speed limitations imposed by the asynchronous TTL interface.

Rambus' completely new physical, electrical and logical interface yields high-performance memory systems from as few as one RDRAM. Although the maximum sustained bandwidth can be significantly less than its 500-Mbytes/s peak bandwidth, the RDRAM's large cache and fast cache-hit accesses yield good performance, especially when large transfers are appropriate. Rambus is well suited to applications such as graphics frame buffers and devices requiring small memory systems. Its apparent disadvantage as a proprietary technology is mitigated by a strategy of broad licensing to multiple vendors.

RamLink is an interesting technology, but is several years away from being commercially available.

None of the new DRAM architectures will take over the world. Most likely, several will find niche markets to sustain them over several generations of DRAMs. Price will largely determine market acceptance, and volumes will drive price. A few significant design wins will mean more in determining ultimate market dominance than any of the technical issues. ♦

