A new era of fast dynamic RAMs

Chip makers are turning to innovative circuit designs to bring inexpensive dynamic memories up to speed

Way back when, the microprocessor's need for data outstripped the rate at which dynamic RAM could access the information; so to catch up, system designers turned to cache subsystems based on static RAM. But static RAM, though faster, is also the costlier form of chip memory. Now an alternative has arrived. Until recently, only a few suppliers focused on the high-speed dynamic RAM market; now many are developing high-speed I/O concepts, in a move away from the high-density, low-speed memory-chip-as-commodity market. For system designers, that means fast, inexpensive new dynamic RAMs (DRAMs) are here, while static RAM (SRAM) is being displaced or relegated to the DRAM chip.

For example, three generations (256K through 4M bits) of address-multiplexed DRAMs with row address access times as short as 40 ns have emerged from NMB Semiconductor Co. While the Tateyama, Japan, company employed standard CMOS technology, Tokyo's Hitachi Ltd. used the more complex and expensive biCMOS process to extract a 25-ns access time from a 1M-bit-by-1 nonmultiplexed (broadside-addressed) DRAM.

Seldom do these speed-optimized devices require larger chips to achieve their higher performance. Though Hitachi's new fast DRAMs are fabricated with specially optimized processes, others resort to imaginative circuit designs to obtain high speed.

Several new chips turn the wide data buses within DRAMs to advantage; on most of them, 1024 to 4096 memory locations may be accessed in parallel. Examples are cache DRAMs, enhanced DRAMs, synchronous DRAMs, and Rambus DRAMs, all described later in this report. While all these ICs capitalize on wide on-chip buses, they vary greatly in their die sizes, terminal characteristics, and performance benefits [Table 1, p. 44].

Fred Jones United Memories Inc.

FIRST APPLICATION. Video was the high-data-rate application that spurred the development of application-specific DRAMs that utilize wide on-chip buses. Video RAMs take wide data fields within the addressed row and transfer them, in parallel, into parallel-to-serial registers. When loaded, the registers independently clock serial data out of the serial port to refresh the video screen, while the processor or graphics processor is accessing the random access port to update the image.

Often, split registers load half the serial register for the on-chip video port, all the while clocking serial data out of the other half-register. A split register thus loads data into the serial registers without interrupting the data stream from the video RAM's serial port.

Other features on video RAMs, such as block and flash writes, simplify and speed up many graphics operations. In general, the dual-port architecture of video RAMs is their prime attraction, but they rarely demonstrate blazing speed on either of the I/O ports. Rather, the high video pixel rate is produced by accessing data broadside from the serial ports of several chips and by boosting the data rate with high-speed parallel-to-serial shift registers or other bit acceleration techniques.

> Once considered too slow for high-performance applications, dynamic RAM can now replace fast but expensive static RAM

WIDE BUSES. Like video RAMs, cache DRAMs and enhanced DRAMs fill on-chip caches from their internal buses. Their SRAM caches are integrated into a standard CMOS high-density DRAM chip, yielding a much higher data transfer rate between dynamic memory and static RAM cache than if they were on separate chips. The on-chip cache is also more effective and less expensive. These single-chip cached dynamic RAMs produce full output voltage swings, either CMOS- or TTL-compatible, but may also use limited-swing I/Os in the fiture.

The enhanced DRAM from Ramtron Corp., Colorado Springs, Colo., uses 8K bits, distributed on chip, of static RAM cache (2048 bits are accessible during a given cache read) ["Enhanced dynamic RAM," p. 49]. The cache is direct-mapped and row (address) associative. Since 2048 bits are transferred in parallel (from DRAM array to cache) when the cache is loaded, the fill bandwidth per device is 7.3 gigabytes per second.

Access time for a cache hit is 15 ns (giving a total cycle time of 15 ns), and for a cache miss, 35 ns (65-ns cycle time). Read or write burst cycle time is 15 ns. All memory cycles are asynchronous. Refresh cycles and precharge operations can be hidden during cache read operations.

A write-through cache is used to maintain coherency between cache and main memory: data is written to the main memory array through the cache for a cache hit, but written to the array directly for a miss. With the enhanced DRAM, this practice incurs none of the delay normally associated with write-through architectures, since both the memory write and cache write operations together require only a 15-ns write cycle.

SPLIT SECOND. Cache can be read immediately after a write hit (write to cache), but a write to read latency (30 ns maximum) occurs when a read miss or write miss (data not in cache) occurs after any write operation. This seldom affects system performance because of traditional bus protocol limitations. With write-through cache, writes (except for posted writes) occur in real time, without the additional latencies associated with copy-back cache.

The cache DRAM from Mitsubishi Electric Corp., Tokyo, like Ramtron's enhanced DRAM, utilizes on-chip cache, but in a different role ["Dynamic RAM as secondary cache," p. 49]. The Japanese DRAM contains two 8K-bit sections of localized on-chip static RAM (256 rows of 16 four-bit words); each may be configured either as direct-mapped, two-way set-associative, or as four-way set-associative. In a cache transfer, 64 bits (16 by 4) are moved in parallel, giving the cache DRAM a cache-fill bandwidth of 114 megabytes per second.

For cache hits, access time is 10 ns (10-ns cycle), while for cache misses and direct array access, it is 70 ns (280-ns cycle time for cache operations, 140 ns for direct array

0018-9235/92/$3.00©1992 IEEE

High-speed dynamic RAMs—a comparison

Type of dynamic RAM	Enhanced	Cache	Synchronous	Rambus
I/O width, bits	X1, X4	X4	X4, X8, X9	X8, X9
Data rate, single-hit, MHz	67	50–100	50–100	500
First-access latency Cache/bank hit, ns Cache/bank miss, ns	 15–20 35–45	 10–20 70–80	 30–40 60–80	 36 112
Cache-fill bandwidth, Mbytes/s	7314	114	8533[a]	9143[a]
Cache/bank size, bits	2048	8192	4096[a]	8192[a]
Area penalty, percent[b]	5	7	5–10	10–20
Output level	CMOS/TTL	CMOS/TTL	CMOS/TTL, GTL/CTT	600-mV swing, terminated
Access method	Asynchronous, DRAM-like	Synchronous, proprietary	Synchronous, pulsed/RAS	Synchronous, proprietary
Access during refresh	Yes	Yes	Undecided	No
Pin count/package	28/SOJ	44/TSOP	44/TSOP	32/VSMP
Density, bits	4M	4M	16M	4M

CTT = center-tap termination; GTL = Gunning transceiver logic; RAS = row access strobe; SOJ = small-outline J-lead package; TSOP = thin small-outline package; VSMP = nonstandard vertically mounted package.
a Synchronous and Rambus DRAMs store data in sense amplifier latches, not in separate synchronous RAM cache.
b Area penalty is relative to the manufacturer's standard die size, so that the figures are not directly comparable.

accesses). All memory cycles are synchronous. Refresh cycles and precharge operations can be hidden during cache read operations.

The cache DRAM uses fast copy-back to maintain cache/main memory coherency. If a cache miss occurs, the old data is moved into an on-chip temporary buffer until the new block is moved, in 70 ns, from the DRAM into the cache. The old block, if it needs updating, is moved to main memory, updated, and returned to cache while the processor is accessing the cache. Except for multiple back-to-back misses, this method hides the latency of write-back cycles and takes one-third the time of traditional write-back methods.

NEW CLASS. Meanwhile, a new class of memory chips, synchronous DRAMs, is being developed by several suppliers ["Synchronous dynamic RAM," pp. 44–48]. The synchronous DRAM bursts data sequentially from its wide internal bus at synchronous clock rates of 50 to 100 MHz. Many believe it will become the next standard DRAM.

Synchronous DRAMs provide high sustained data rates, but impose long delays for single random reads or writes. All memory functions of the synchronous DRAM are timed from the rising edge of a master high-speed clock (up to 100 MHz). Pipelining ensures fast, sustained sequential accessing and queued writing. Synchronous DRAMs are optimized for sequential data transfers (sequential binary Intel burst sequences) rather than random access. Precharge operations can be hidden by using interleaved accesses to separate memory banks within the synchronous DRAM. Depending on the performance range of the synchronous DRAM, I/O levels will be either full swing (CMOS/TTL) or restricted swing (Gunning transceiver logic, called GTL, or center-tapped termination, called CTT).

Rambus DRAMs, based on a bus concept developed and licensed by Rambus Inc., Mountain View, Calif., are being developed by several large manufacturers ["A fast path to one memory," pp. 50–51]. They share some characteristics with cache, enhanced, and synchronous DRAMs. Rambus DRAMs are intended for use in a very controlled and specifically defined Rambus environment, in which both transitions of a 250-MHz clock are used, for an effective clock rate of 500 MHz (clocked at a transition every 2 ns). The serial burst data rate is expected to be at the effective clock rate—that is, 500 megabytes per second.

This data rate is phenomenal for DRAM. But it applies to data packets once they start being transferred—it does not include the delay associated with the request protocol, particularly when a miss and re-request occur; these operations take a disproportionately long time.

Moreover, the transfer time is 2 ns within the Rambus environment only; it applies to the transfer of data between master and slave portions of the local memory bus, and not between the Rambus environment and the subsystem (a central processing unit or graphics processor, say) that uses the high-speed data.

Even so, system and equipment designers have a far wider choice of memory devices than they did just a year ago.

Synchronous dynamic RAM

Betty Prince, Roger Norwood,
Joe Hartigan, Wilbur C. Vogley
Texas Instruments Inc.

Data rates of 500 megabytes per second are expected for a dynamic RAM with a new architecture, new interfaces, and a new name: the synchronous dynamic RAM. In other words, the 100–200-MHz microprocessor of the mid-1990s could function nonstop if this kind of dynamic RAM (DRAM) were used for main memory. In addition, in simple cases, main memory might once again handle special functions like graphics, which currently need specialty DRAMs to achieve high enough data rates.

Yet the synchronous DRAM is an evolutionary approach to high-speed memory, in that it ties together, in a single device, many of the diverse techniques developed over the years to augment memory transfers. The plan is for it to become a multi-sourced, commodity chip.

SPEED MISMATCH. Standard DRAMs and the new microprocessors have increasingly diverged in speed. Basic DRAM architecture has not changed since the early 1970s, whereas design and architectural innovations have improved microprocessors to the point where speeds of 500 MHz or more look possible by the end of the decade. DRAMs, on the other hand, currently take about 60 ns to access a row address, or 30 ns from address to data available in page mode, which is about a 25-MHz cycle time. The problem is compounded by the increasing density of dynamic RAMs, which further adds to the time needed to access the total memory through a memory bus of a given width.

CHIP SOLUTIONS. Still, as DRAM chips have become larger, chip-level solutions have begun to appear. The cache, for example, can be included on the DRAM. Fast-access modes such as page mode, static column mode, or nibble mode have helped to nearly double basic DRAM speed. Page and static column modes work by keeping the active row data latched in the sense amplifiers, and merely selecting new random column addresses. Nibble mode uses wider internal architectures (by 4) and fast registers to make a 4-bit parallel-to-serial conversion. The result is fast access to 4 consecutive bits after a random address.

Although these modes can enlarge DRAM bandwidth to twice its normal, non-mode width, they still fall far short of today's processor needs. Wider external data buses (8 or 16 bits wide, for example) can also allow more information to be retrieved from the DRAM on a single access. The penalties here are larger chips and packages and greater DRAM output noise, all of which tend to slow down the read access.

NONPROPRIETARY. In the synchronous DRAM, many manufacturers, including Texas Instruments, Samsung, Hitachi, Toshiba, Mitsubishi, and NEC are addressing these speed concerns in an evolutionary, nonproprietary approach. Initially, they are developing synchronous DRAMs of 16M-bit capacity. These are currently being considered for open standardization in the Electronic Industries Association/Joint Electron Devices Engineering Council (EIA/Jedec) JC42.3 DRAM Standards Committee. (While synchronous static RAMs have been around for some years, the idea of using a dynamic RAM synchronously has only recently gained acceptance.)

Historically, DRAMs have been controlled asynchronously: the processor presents addresses and control levels to the memory, indicating that a set of data at a particular location in memory should be either read from or written into the DRAM (for a write, the processor also presents the data it wants written). After a delay—the access time—the RAM either writes the new information from the processor into its memory or else provides the information on its outputs for the processor to read.

During the access-time delay, the DRAM performs various internal functions, such as activating the high capacitance of the word lines and bit lines, sensing the data, and routing the data out through the output buffers. The delay creates a wait state for the processor; it simply waits for the RAM's response, and the entire system slows down as a result.

Under synchronous control, on the other hand, the DRAM latches information in and out under the control of the system clock. The processor drops off the instructions for the DRAM in a set of latches, which stores the addresses, data, and control signals on the DRAM inputs until the memory can process the request. The DRAM responds after a set number of clock cycles, which can be programmed by the user in a special configuration cycle.

NO WAITING. Since the processor knows how many clock cycles it takes for the DRAM to respond, it can safely go off and do other tasks while the RAM is processing its requests. An example is a DRAM that has a 60-ns read delay after initial addressing and is operated with a 10-ns (100-MHz) clock. If the RAM is asynchronous, the processor waits the full 60-ns access time for the information. But if the DRAM is synchronous, the processor can strobe the addresses into a set of input latches and do other tasks while the memory does the read operation. Then, when the processor clocks the outputs of the RAM six cycles (60 ns) later, the data it wants is there.

In the synchronous DRAM timing diagram, the row address is strobed in on the rising edge of the system clock, activating a row or word line of memory bits. A column address is then clocked in after three clock cycles (30 ns), sufficient time to activate the word line. The byte of data then appears on the outputs after three more cycles (another 30 ns) to decode the new address and get the data from the sense amplifiers through the output buffers. Six clock cycles after the row address has been clocked in, the processor can expect the requested information from the output buffers of the synchronous DRAM.

The architecture of the synchronous DRAM can further shorten its average access time by pipelining addresses. The input latch stores the next address the processor will want while the RAM is operating on the previous address. Normally the addresses to be accessed are known several cycles in advance by the processor; therefore, after the address for the first access has been sent and the RAM has started working on it, the processor can send the following address to the input latch, so that it is available as soon as the first address has moved on to the next stage of processing. The processor need not wait a full access cycle before starting its next access to the DRAM.

Various techniques can also be used inside the synchronous DRAM to hide the components of the internal timing delay. The address setup time and word and bit line precharge time can be eliminated after the first access by using a burst mode, in which a series of data bits can be clocked out rapidly after the first bit has been accessed.

The burst mode in the synchronous DRAM is similar to the old nibble mode, in which 4 bits of sequential data are provided in rapid succession without inputting new address information to the DRAM—but now as much as a full page of data can be provided [Fig. 1].

Burst mode is only useful if all the bits to be accessed are in sequence and in physically the same row of cells as the initial access. Likely applications include high-speed memory functions such as video support requiring 100-MHz data rates.

Burst mode can be combined with a "wrap" feature. The wrap gives access, for example, to strings of bits stored both before and after the initial bit location in rapid succession after the initial access. This feature is useful for cache filling, since the most likely bits to be wanted next are those physically close to the current bit. The user can program both the type of wrap and the number of bytes available on each wrap.

If data from different rows is needed, it is still possible to hide some of the row precharge time if the two rows lie in different banks. With the synchronous DRAM, the multiple-bank interleaving previously done at the system level may be moved onto the chip. Then one bank can be precharged while another is being accessed. An example is the "nibbled-page" architecture of Tokyo's Toshiba Corp., in which the data from 8-bit sections of different columns is interleaved on-chip to give byte-level random access at a 100-Mb/s rate.

Another advantage of multiple-bank synchronous DRAMs is that the active rows (potentially one in each bank) may serve as a cache. In more detail, once a given row in a given bank is accessed, it is held active and may be accessed again simply by supplying a new column address. This method has been used in page-mode devices, but had only limited success because only one row in an asynchronous DRAM could be held active.

ADDING ASSETS. All these features—burst and wrap modes and interleaved banks—can be combined on a single synchronous DRAM that runs at up to 400–500 MHz. And still more features may be added. For example, a data mask control can be used as an enable/disable pin to ignore inputs or turn the outputs off for a single clock cycle. This could be useful, especially in write cycles, if the user wants to access a string of bits, but not all in the string.

Another feature, clock enable/disable, turns off the system clock inside the RAM, thereby suspending the device in its current state, or puts it into low-power standby mode, saving energy in battery-powered equipment.

Synchronous DRAMs need a refresh cycle, since they are still composed of dynamic cells that lose their charge. But another option, self-refresh, appears on many of the newer byte-wide DRAMs. This new, simpler refresh mode is completely controlled on the chip and retains data in a low-

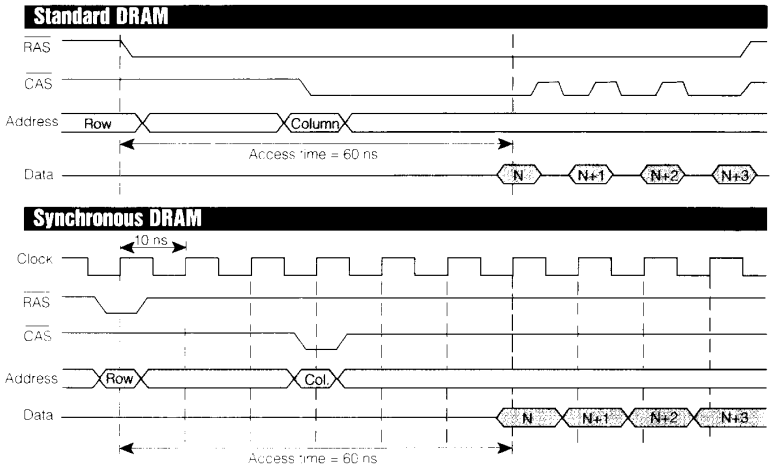

[1] In nibble mode on an ordinary DRAM, 4 bits of data from locations N, N+*1,* N+*2, and* N+*3 appear after location* N *has been addressed. Burst mode in a synchronous DRAM, however, can produce as much as a full page of data (typically 512 bits per output) after only* N *has been addressed.*

power mode. The self-refresh option is expected to be available on synchronous DRAMs from many vendors.

The final speed attained by this kind of DRAM in a system depends not only on its internal architecture but also on the interface signal levels it uses. A 5- or 3.3-V interface can function smoothly in, say, a desktop computer running at 50 MHz. If, however, these same interfaces, with their relatively large 2-V output swings (0.4 V low to 2.4 V high), run in a 125-MHz system, unterminated lines longer than a few centimeters could cause delays. Therefore, higher-speed synchronous DRAMs will have low-swing interfaces such as Gunning transceiver logic (GTL) or center-tap-terminated (CTT) ["Fast interfaces for DRAMs," pp. 54–57] to compensate for transmission-line effects.

SHORTER LEADS. The high speed of synchronous DRAMs influences their packaging. Most appropriate for them are new miniature packages such as the thin small-outline package (TSOP) or various vertical surface-mount packages. These reduce the effective length of the package wiring and leads, and devices may be tightly spaced on a circuit board.

Advances in fast DRAM architecture, high-speed interfaces, and miniature packaging combine in the synchronous DRAM into a widely sourced device type that could become the next-generation commodity DRAM. Whether the promise is fulfilled depends on several factors: producers must standardize their products and they must produce them in the high volume needed to bring costs down. No less important, many mainstream DRAM manufacturers must offer synchronous DRAMs so that users may have alternative sources to rely on.

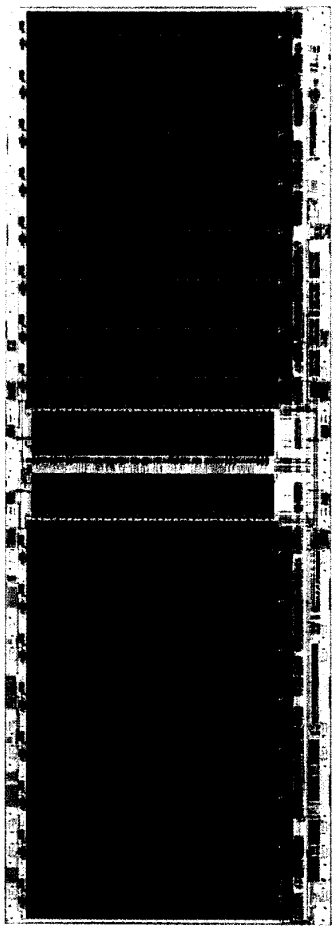

[2] A cache DRAM contains static RAM in two shorter banks at center of die.

Dynamic RAM as secondary cache

Charles A. Hart
Mitsubishi Electronics America Inc.

Virtually all future microprocessors will have a first-level cache memory on chip, where it boosts efficiency by allowing access hits 80–90 percent of the time. But as processor clock frequencies rise, so does the deleterious effect of cache misses on system performance. This is where cache DRAM [Fig. 2] can play a key role: it provides a second-level cache that complements the on-chip cache.

For example, a 200-MHz microprocessor whose on-chip cache has a 90 percent hit ratio will operate at only 40 percent of its potential if each miss cycle takes 80 ns or 16 clocks. But if the processor has access to a secondary off-chip cache with a 20-ns cycle time and a hit rate of 90 percent, the processor can attain 70 percent of its potential.

The two most important parameters for a secondary cache are access time and hit rate. The secondary cache hit rate must be optimized and its speed must be as great as possible.

OPTIMAL HITS. A cache DRAM optimizes its hit rate by including a small static RAM and linking its static and dynamic sections by a wide data path [Fig. 3]. Depending on the benchmark, a cache DRAM improves processor utilization by 42 percent and boosts hit rate to 96.9 percent. At the same time, adding a 32K-byte static RAM cache to an 8M-byte dynamic RAM increases the die size by only 7 percent and raises the already low chip cost by only 10–20 percent.

A single bank of 4M-byte dynamic RAM on a 64-bit-wide bus can transfer data to and from a small on-chip static RAM at a rate of 1.6G bytes/s. As the number of static RAM banks is increased, data transfer rates grow even higher—eight banks transfer 12.8G bytes/s—and hit rates increase as well.

Cache DRAM does not duplicate the function of the primary cache on a processor. The primary cache usually has some associativity and a modest-size block (usually 4). A cache DRAM secondary cache usually will be direct-mapped (although associativity is allowed) and the block size will range from 16 to 64 or even larger, depending on the number of banks. Thus the primary cache works well for tight instruction loops, while the secondary cache works well as an instruction and data prefetcher. The cache DRAM secondary cache will handle context switching and data transfers very efficiently.

At least as important as the hit rate of a secondary cache is its access time. The cache DRAM architecture allows a small memory of this kind to function like a large one. Since it is easier to make a small cache fast, a 10–15-ns access time is easily achieved by the 4M-byte memory made by standard, low-cost dynamic RAM processes.

Future generations of the cache DRAM will most probably employ some type of new I/O standard (such as Gunning transceiver logic, or GTL) to further shorten the access time from processor to second-level cache. The synchronous interface that cache DRAMs now employ will probably be retained in future generations.

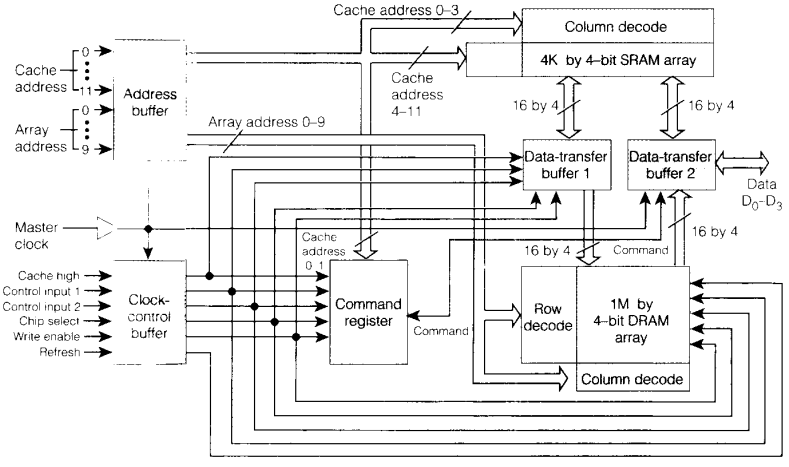

[3] The cache memory in a cache DRAM is a small, fast static RAM. The main memory is a standard dynamic RAM. The two memory sections communicate over a wide data path through data transfer buffers. Miss data is transferred into one buffer, while the expected data is transferred simultaneously from the dynamic RAM to the static RAM through the other buffer.

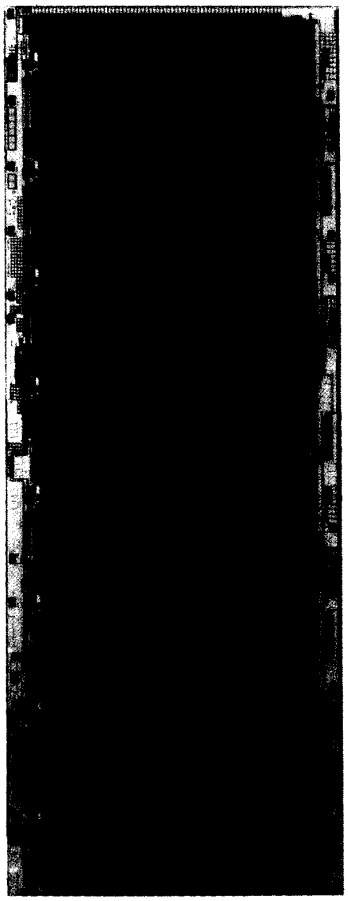

[4] An enhanced DRAM spreads static RAM cells throughout the die; they surround the light horizontal strips.

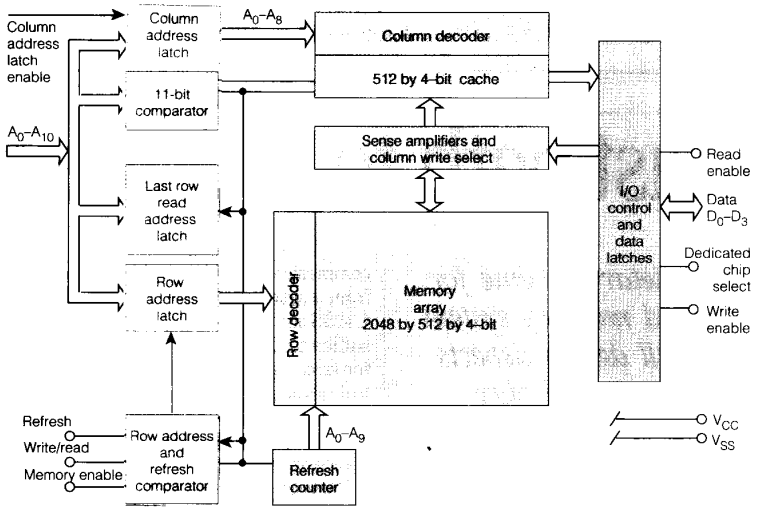

[5] In the enhanced dynamic RAM, static RAM cache directly maps to the most recent row accessed from the dynamic RAM array. This feature simplifies cache control and saves on the chip area needed for the static RAM. Data is available asynchronously and randomly from the static RAM cache at a rate of 67 MHz.

Enhanced dynamic RAM

David Bondurant
Ramtron International Corp.

When computer designers are asked what features they want in the next generation of dynamic RAMs, they will generally specify three: zero wait states on all bus cycles (not just burst reads); standard dynamic RAM (DRAM) architecture and packaging for an easy system upgrade; and no significant cost increase over standard DRAM. Enhanced DRAMs are aimed at satisfying these wants.

Having no wait states is especially important, since all waits impair microprocessor efficiency. This fact is made painfully obvious by new clock-doubled architectures where every wait state on the bus translates into two within the processor. Enlarging burst-read bandwidth without doing anything for other bus cycles solves only part of the problem. With its novel device and circuit designs, the enhanced DRAM improves performance on all bus cycles [Fig. 4].

The CMOS fabrication process for enhanced DRAMs introduces a field-shield structure that actively isolates transistors in the silicon wafer and reduces their junction capacitance; the transistors switch faster as a result. A 4M-bit enhanced DRAM has a row access time of 35 ns and a cycle time of 65 ns—twice as fast as most conventional DRAMs.

An enhanced DRAM chip also includes a 2K-bit static RAM (SRAM) cache row register that is accessed in 15 ns. This cache provides zero-wait-state read-hit and burst-read cycles at clock rates of up to 33 MHz without interleaving and greater than 50 MHz with two-way interleaving.

DATA TRANSFER. A 256-byte-wide bus joins the 4M-bit DRAM array to the 2K-bit on-chip SRAM cache [Fig. 5]. All 2048 bits of data in the cache are updated in a single 35-ns row access, for an effective cache-fill bandwidth of 7.3 gigabytes per second. This high data transfer rate is made possible by the close coupling of the DRAM and SRAM. It reduces wait states due to read, burst-read, and multiple burst-read misses to a far lower level than is attained by employing conventional combinations of external SRAM and interleaved DRAM.

Writing data to the enhanced DRAM is speeded by an on-chip write-posting register. With it, writes may be posted to the chip in 15 ns, and then the cache may be read while the 65-ns write cycle is completed in the background. Write bursts proceed within a page at 15 ns per word. This is three times faster than standard fast-page-mode DRAMs. An on-chip hit/miss comparator automatically maintains cache coherency with on-chip DRAM for write hits, without incurring a speed penalty.

Since all reads occur from the cache row register, precharging is also performed in the background. The next bus cycle can start immediately without a precharge delay. Similarly, DRAM refresh can also be performed in the background of cache read operations.

An enhanced DRAM has the same address and data path as a standard DRAM with page mode and static column operations from cache. An enhanced DRAM needs only five more pins than a standard DRAM—one extra address and four new control pins. It is housed in a small-outline J-lead (SOJ) plastic package, like a standard DRAM.

Eight enhanced DRAMs can be combined in a single in-line memory module (SIMM) that directly replaces a standard DRAM SIMM; a computer may then be upgraded to faster memory simply by reconfiguring the memory controller. The 72-pin SIMM's capacity of 36M bits is joined to 16K bits of SRAM by a 16K-bit-wide bus, resulting in an extemely high cache-fill bandwidth of 58.6 gigabytes/s.

The SRAM cache and wide bus add only a little area to the enhanced DRAM chip—about 5 percent. The enhanced DRAM chip therefore costs only slightly more than a standard DRAM of equal size.

ABOUT THE AUTHORS. Fred Jones [pp. 43–48] is concept engineering manager at United Memories Inc., Colorado Springs, Colo. He previously was associate director of memory semiconductor service at Dataquest Inc., San Jose, Calif.

The authors of "Synchronous dynamic RAMs" are all with Texas Instruments Inc., Houston, Texas. Betty Prince is manager, new products, MOS Memory Worldwide Marketing. Roger Norwood is manager, DRAM design. Joe Hartigan and Wilbur C. Vogley are DRAM systems specialists.

Charles A. Hart [p. 48] is memory marketing manager at Mitsubishi Electronics America Inc., Sunnyvale, Calif.

David Bondurant [p. 49] is director of applications engineering at Ramtron International Corp., Colorado Springs, Colo. ◆