

SYSTEM EVOLUTIONS, TECHNOLOGY CAPABILITIES, AND ECONOMIC REALITIES COMBINE TO MAKE DRAM THE INTRIGUING AND UNCONVENTIONAL BLACK SHEEP OF AN OTHERWISE RELATIVELY STAID AND PREDICTABLE MEMORY FAMILY. THE EVER-INCREASING DIVERSITY OF ALTERNATIVES, ALL JOCKEYING FOR A LEADERSHIP POSITION, COMPLICATES BOTH YOUR AND THE VENDORS' LIVES, EVEN AS IT INCREASES YOUR OPTIONS.

THE SLAMMIN,' JAMMIN' DRAM SCRAMBLE

“DRAM IS MOVING BEYOND its conservative, predictable, multi-sourced roots to an era of greater controversy, increased innovation, and more

product options.” So I stated more than two years ago ([Reference 1](#)). The trend hasn't stopped: I probably even *underestimated* the amount of debate that would ensue as various DRAM candidates wrestled each other for your attention.

One 1997 memory option, SDRAM, has fallen by the wayside, although pieces of the technology may live on under other acronyms. Standard synchronous DRAM (SDRAM) is now running beyond 133 MHz, a speed that in 1997 barely existed on any of the manufacturers' road maps. Several new architectures, including virtual-channel memory (VCM), fast-cycle RAM (FCRAM), and high-speed DRAM (HSDRAM), have entered the picture, oth-

ers, such as double-data-rate SDRAM have transitioned from vendor-proprietary to industry-standard status, and at least one, Direct Rambus DRAM (DRDRAM) is slower to come out of the gate than its backers originally anticipated.

Economic factors are fundamentally behind all this confusion. In 1997, the DRAM industry was in only the early stages of what has become a prolonged industry recession. Microsoft's repositioning and introduction of Windows NT 5.0 (now Windows 2000) is partly to blame, because its absence significantly decreases the average amount of memory that computer manufacturers ship with each system. Remember that the construction and qualification of a new semiconductor-fabrication facility is a multiyear, multibillion-dollar process.

The mid-1990s saw the beginning stages of a tremendous capacity buildup, partly as an over-reaction to the supply-constrained—and, therefore, lucrative—business climate and partly in anticipation of the greater memory demands that the supposedly soon-to-appear

Photo courtesy Hitachi

operating system upgrade would place on PCs. With supply exceeding demand by a wide margin a few years later and Windows 2000 nowhere in sight, prices crashed (**Figure 1**).

Responses to the oversupply crisis have been inconsistent. Some companies have narrowed their focus to a few mainstream architectures, densities, and bus widths and concentrated on reducing costs by shrinking die on existing fabrication equipment. Many of these same vendors have simultaneously cut back on more significant next-generation investments. Others have turned their backs on commodity memory

in the hope that differentiated products will prove more profitable.

A few manufacturers have bucked the trend and accelerated their R&D programs in the hope that, when the long-awaited industry upturn finally happens, they'll be better prepared to exploit it than their more fiscally conservative competitors. And foundries, stung by overcapacity problems of their own, have entered the DRAM business in an attempt to fill their originally ASIC-intended facilities.

Previously bitter enemies have forged joint technology- and product-development alliances, and some vendors actively selling DRAM in 1997 are now nearly or completely out of that business, having auctioned off their market share to the highest bidders.

This wave of consolidation will undoubtedly continue.

HOW LATE CAN YOU WAIT?

DRAM designers have plenty of ideas for improving memory speed, but commodity suppliers operate under a restrictive set of constraints. Drive the internal signal lines with larger transistors or reduce the transistor switching threshold, and you've just increased power consumption and refresh requirements. Subdivide an array to reduce signal-line length and impedance, toss an on-chip cache and controller alongside the DRAM array, or widen the internal data bus to fetch more information at once, and you've just blown up the die size. Turbocharge the bus-toggle rate, and the chip becomes exponentially more complex to test. And, use nonstandard, expensive process steps, such as copper interconnect, or esoteric capacitor dielectric materials, and see what happens.

Moore's Law trends don't help DRAM speed to the degree that they accelerate logic circuits, and density growth—along with corresponding array-size-driven internal-interconnect-length increases—counterbalance these trends. Nevertheless, incremental

At a glance.....	70
A murky crystal ball.....	72
Can't get enough?.....	74
For more information.....	79



improvements have occurred. Late '97's asynchronous extended-data-out DRAMs are a distant memory (pun intended), PC-66 SDRAMs are quickly disappearing from all but low-end notebook PCs and high-end servers (which derive much their performance from ultrawide system buses), and PC133 memories are rapidly ramping into production. Specialized SDRAMs for lightly loaded point-to-point connections, such as graphics frame buffers, operate at speeds beyond 200 MHz.

The term "PC133" encompasses a number of specifications and can include both column-address-strobe (CAS) latency 3 (CAS-3) and CAS latency 2 (CAS-2) variants. Even the common term "CAS-2" is imprecise because it can refer to the number of required clocks for any of three device specifications, frequently documented in an "X, Y, Z" fashion, where X=CAS to valid data, Y=row-address-strobe (RAS) to CAS, and Z=CAS high time, or "precharge."

Nowadays, the vendors all yield nearly 100% to PC100 DRAM with two-clock specifications for all three parameters. However, many of today's PC133 parts are CAS-3, and some of the claimed CAS-2 parts, until their next lithography reduction, have at least one three-clock CAS-related delay, often the critical CAS-to-valid data specification. Off-the-shelf DRAM controllers and cores for ASICs and programmable-logic devices should be able to handle the multiple DRAM variants, but slower memories will at least minimally impact system performance. If you design your own DRAM controller, make sure it's equally flexible, or ask your supplier enough questions to be sure of what you're getting.

Per-chip density growth has continued its torrid pace, surpassing the predictions of a number of industry observers in the process. In late '97, manufacturers were only beginning to move from 16- to 64-Mbit devices. Nowadays, many companies no longer make 16-Mbit DRAMs, regardless of what their Web sites say, and they're rapidly shifting away from the 64-Mbit density as well. Near- or sub-0.2- μ m lithographies are to thank (if you can use the density) or blame (if you need cost-effective smaller parts) for this trend. It's nearly impossible to figure out from press releases just where each company is in a given density's development: Does an announce-

AT A GLANCE

- ▷ The diversity of DRAM architecture, density, and interface choices continues to increase with no end in sight.
- ▷ DRAM manufacturers focus on reducing latency even as CPUs do their utmost to make slow random access a trivial issue.
- ▷ Rambus is attempting to justify its memory's bandwidth in PCs, whereas double-data-rate synchronous DRAM looms as a long-term threat.
- ▷ BGAs are the future DRAM packages of choice, but incompatible alternatives may have some designers yearning for days when they could use standard TSOPs.
- ▷ Benchmarks, like standards, are becoming plentiful. Supplement their data with your own analysis to pick the right DRAM for you.

ment for a 256-Mbit part mean that it is in volume production or limited sampling, or does it just mean that the vendor has a photoworthy but nonfunctional piece of silicon?

Skepticism aside, announcements, such as Infineon's mid-August 1999 claim that it had produced 1 million 256-Mbit parts, are notable, and many companies are also in production with half-step 128-Mbit memories. Traditionally, DRAMs made four-times density jumps (1 to 4 to 16 to 64 Mbits) from one generation to another, partially because this trend retained row- and column-address symmetry and a roughly square die. Follow-on, 128-Mbit devices have proved

easier for some companies to develop on existing fab equipment and fit into industry-standard packages, however, and, for similar reasons, 512-Mbit parts appear on some vendors' future road maps. These half-step memories not only extend the lifetime of expensive manufacturing lines, but also keep the vendors from overshooting your system-memory requirements.

The stacked-versus-trench-capacitor debate remains heated, though one of the participants may be switching sides (**Reference 2**). Previously, Toshiba had partnered with trench-capacitor advocates IBM and Infineon on joint technology development. Toshiba recently announced, however, that it would in the future be working with Fujitsu. Will Toshiba move its future products to a Fujitsulike stacked-capacitor technology? Perhaps to deflect any doubt about the long-term viability of the trench approach, IBM and Infineon announced at December's IEEE International Electronic Devices Meeting (IEDM) in Washington that they'd successfully fabricated test DRAM arrays using a novel, miniaturized cell with a rotated transistor element (**Figure 2**). High-volume production of 1-Gbit and larger memories based on this approach should begin in four years or so.

The biggest debate in DRAM, though, remains the dispute over whether random-access latency or burst bandwidth is the more important performance parameter. Much of the discussion centers on PC main memory, although the answer depends on the operating system and other software's code- and data-access profiles. Expand the list to other applications, and conclusions become even

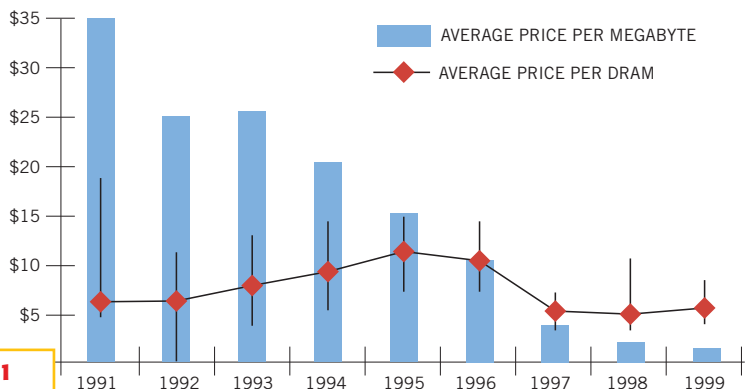


Figure 1

The last few years have been difficult times for DRAM suppliers and happy days for memory users, although cost per device has held fairly constant (courtesy Micron Technology).

more elusive. But the following general observations might help guide your analysis and selection:

- The shorter the average burst-access length, the less you can amortize an extended initial latency over much shorter subsequent burst accesses. For example, accesses consisting primarily of cache-line fills prioritize latency more than those dealing with much longer data streams.
- The more effective the CPU's caching scheme before the DRAM array, the more random the DRAM-code accesses to fill cache lines. In such a scenario, the system-memory controller is less likely to reaccess already-open DRAM pages.
- When a CPU, particularly one without pipelining or prefetch support, has to read information from main memory, the CPU stalls, wasting precious clock cycles, until the first

access completes. Any shorter time taken to fill the remainder of the cache line is a don't-care if the CPU never uses this information.

- The fewer system masters accessing main memory, the less likely that they will consume a significant amount of a memory's peak bandwidth. Some benchmarking studies suggest that cache-line fills in a PC, for example, use less than 5% of the total available PC100 SDRAM bandwidth.

Compare this reality with the claimed 95% usage capability of the high-speed Direct Rambus Channel, and the need for multiple access sources and types to exploit the DRDRAM (or, for that matter, double-data-rate (DDR) SDRAM) potential becomes clear.

CUTTING TO THE CORE

MoSys, one of the first advocates of low-latency DRAM, has redirected its MoSys DRAM (MDRAM) architecture,

which never achieved more than niche graphics-memory status, in two interesting directions. First, the company is using the MDRAM core in a line of memories that it claims are functional- and pinout-compatible but lower cost, synchronous-SRAM replacements. MoSys is also aggressively pursuing an embedded-memory-licensing program. Nintendo plans to use an embedded-MDRAM array in its next-generation *Dolphin* game console, representing MoSys' most notable design win to date.

Enhanced Memory continues to evolve its plan for enhanced-DRAM (EDRAM)-based devices, again in a number of directions. Enhanced synchronous DRAMs (ESDRAMs), available in densities as large as 16 Mbits, use an SDRAM-compatible synchronous interface and retain EDRAM's combination of fast-cycle DRAM core and higher speed, on-chip, four-way-direct-map cache and controller logic (**Figure 3**). ESDRAM's command set is a superset of the Joint

A MURKY CRYSTAL BALL

What's going to happen with DRAM in the near- and long-term future? PC133 will be highly successful in PC main memory, though for how long is unclear. At the fall 1999 Developer Forum, even Intel announced that it was belatedly developing a PC133-targeted chip set for first-quarter-2000 availability. The company doesn't say what this chip set will look like, but it will probably be a variant of the i810, targeting low-end, low-cost, low-performance systems.

Skeptics suspected that Intel would attempt to derail Via's chip-set momentum by dictating changes to already-developed PC133 specifications, but this scenario hasn't—at least not yet—occurred. And those of you not designing PC main-memory subsystems, who correspondingly have shorter, more lightly loaded bus configurations, should be able to use the even higher speed single-data-rate synchronous-DRAM (SDRAM) variants that will naturally appear, thanks

to shrinking lithographies.

Intel and Rambus are in a quagmire. They mostly missed the lucrative 1999 holiday and end-of-budget-year buying season. Internet-based and mainstream business applications don't give a significant performance edge to the higher cost Direct Rambus DRAM (DRDRAM)-based systems over their cheaper PC100 and PC133 counterparts. A 133-MHz Pentium III processor's maximum local-bus bandwidth is only 1 Gbyte/sec, after all, the same as with AGP 4X. The intrachip bus peak bandwidth is 266 Mbytes/sec between the i820's north- and south-bridge chips. (The north bridge interfaces to the CPU, AGP, and DRDRAM, and the south bridge handles Ultra ATA/66, 32-bit PCI, Universal Serial Bus, and audio interfaces.) Multiprocessor systems for the i840 chip set, whose north-bridge chip also supports 64-bit PCI traffic through a companion interface chip, haven't yet hit average users' desktops.

Media-content developers performing, for example, audio, still-image and video-image editing might notice a difference with a DRDRAM-powered system, according to system reviews. However, these developers represent a fairly small segment of the PC-buying population. Don't overlook Intel's ability to help its anointed memory technology along via chip-set availability and pricing. Note, though, that the i820 chip set, via its optional 82805AA memory-translation-hub companion chip, can also talk to PC100 SDRAM. Also, don't underestimate the 1.6- and 3.2-Gbyte/sec marketing hype in its ability to influence consumers who already buy faster and more expensive CPUs than they really need (**Reference A**). DRDRAM- and i820-based systems with prices starting at less than \$2000 are cheaper than many folks predicted.

AMD's Athlon microprocessor's local bus operates at 200 MHz, and Pentium local buses will also soon extend beyond 133 MHz. At

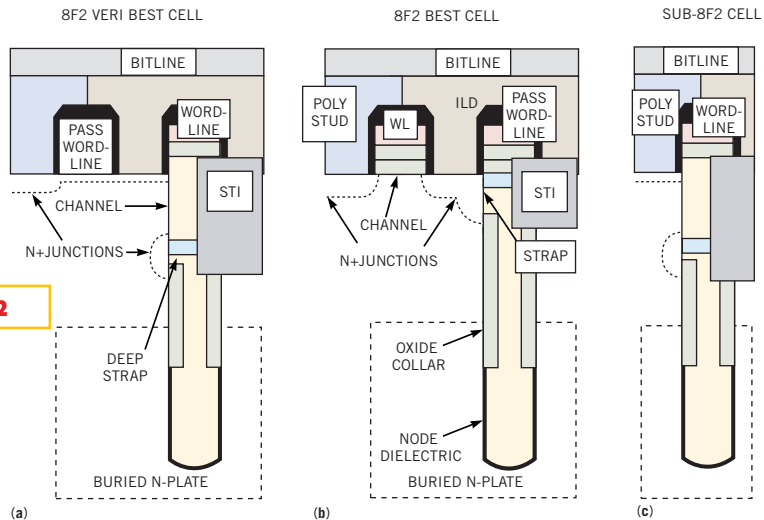
that point, PC133 runs out of steam, and some variant of double-data-rate (DDR) DRAM will take over. Which one—DDR SDRAM or DRDRAM—is to some extent up to Intel, given its dominant position both as a CPU vendor and a chip-set provider in the desktop- and notebook-computer segments. Intel's dominance is less secure, though, at the server and workstation and "free"-PC ends of the PC spectrum, which currently see the highest memory-bit growth rates. Intel's first server chip set for the Itanium (formerly, Merced) CPU will use DDR SDRAM. Will the company's DDR SDRAM-based chip sets extend into even lower cost systems, thereby butting heads with its DRDRAM-based workstation and mainstream PC chip sets?

DRDRAM probably has at least a year's lead on first-generation DDR SDRAM. SDRAM's developers only recently finalized its specs, and the only DDR SDRAM-chip-set developer, Micron, admits that SDRAM isn't yet ready

Electronic Device Engineering Council (JEDEC)-approved SDRAM standard, although ESDRAM costs more than SDRAM, limiting ESDRAM's use in PCs. However, ESDRAM also offers a lower cost alternative to synchronous SRAM. The chips' CAS-1 support in the integrated memory controllers of Motorola's MPC860 embedded processors, Analog Devices' ADSP-21065, and other DSPs also helps you easily use ESDRAMs in non-PC applications. SRAM manufacturers are taking notice: In November, Enhanced Memory signed a joint technology- and product-development agreement with Cypress Semiconductor (Reference 3).

Enhanced Memory also offers HS-DRAMs, which retain ESDRAM's fast core but strip out the cache and controller. The HSDRAMs are available in densities as great as 256 Mbits and provide CAS-2 performance at 133 MHz. Enhanced Memory's marketing strategy for these devices is somewhat unique:

Figure 2



IBM and Infineon hope to put any lingering trench-capacitor-scaling concerns to rest with their innovative DRAM cell (a), which modifies a conventional cell transistor (b) and may lead to even more compact structures (c) (courtesy IBM Microelectronics and Infineon Technologies).

The company sells HSDRAMs in DIMM-mounted form directly from its Web site, targeting chip sets that support

133-MHz memory. The HSDRAMs also appeal to game-playing and other per-

(continued on pg 76)

for high-volume production. Creating a few stable-running systems at room temperature with silicon from one supplier is far simpler than extending the process to a few dozen suppliers and a few million PCs. Characterization and compatibility work remains, and this work will probably result in at least a few tweaks to the specs, the memory design, and the chip-set design.

The newly formed Rambus-DRAM (RDRAM) Implementers Forum had better use this time wisely and quickly get memory costs closer to DDR SDRAM levels. If Intel, Rambus, and their memory partners had hit their original mid-1999 production schedules, DDR SDRAM would have had a greatly diminished competitive position. Two easy cost-reduction opportunities are for the companies to use a cheaper package, as Micron and Toshiba have done, and to reduce the number of on-chip banks, as long as this approach doesn't adversely affect performance.

Assuming that DDR SDRAM's developers fix any bugs it may

have, it will deliver, at 133 MHz, higher peak bandwidth than one Rambus channel's worth of 400-MHz DRDRAM. For roughly the same pin count as that of a DDR SDRAM interface, however, a chip set can support *two* Rambus channels. This approach, which Intel uses on the i840, would return the theoretical bandwidth lead to DRDRAM. The pin-count issue has another interesting ramification for low-end, single-chip PCs. It's a good bet that Via is basing its plans on DDR SDRAM. However, which memory interface will Intel's upcoming, highly integrated Timna μ P support? And with fine-pitch BGAs now mainstream packaging approaches, do a few dozen extra pins matter that much?

What about Itanium and other future high-end PC microprocessors? Longer cache-line-fill sequences, should they appear, will make bandwidth a higher priority than latency. Deeper pipelining, better use of prefetching within compilers, speculative execution, dynamic instruction re-ordering, and explicit parallelism

will further isolate the CPU from long DRAM latencies. These features match DRDRAM's packet-based protocol, allowing simultaneous system requests to multiple DRDRAM devices. But would a virtual-channel-memory (VCM) core-based DDR SDRAM deliver comparable benefits?

Mobile computing will probably be the last PC segment to move to DDR memory. After all, high-end notebook PCs are only now moving from PC-66 to PC100 SDRAM. Both DDR SDRAM and DRDRAM have power-consumption issues to resolve, and neither technology has a clear advantage. Intel's spring 1999 Developers Forum unveiled the small-outline Rambus inline-memory module (SO-RIMM), and presentations at the fall 1999 Forum hinted that a DRDRAM-based notebook PC chip set would appear in the second half of this year. However, Intel subsequently embraced PC133 SDRAM, which may alter the plans for this chip set's introduction. DRDRAM's narrow channel reduces its memory granulari-

ty, a benefit in mobile computing. However, because it wakes from sleep mode slower than DDR SDRAM, DRDRAM could limit notebook-PC performance. If the BIOS and operating-system drivers choose instead to keep the DRDRAM subsystem awake, battery life could suffer.

DRDRAM's best fit may be in areas other than PC main memory, such as in Sony's upcoming Playstation 2. DRDRAM's narrow 16-bit channel delivers top speeds close to those of a 64-bit DDR SDRAM interface but requires one-fourth the minimum density of an array comprising same-sized DDR SDRAM components. Enhanced synchronous DRAM (ESDRAM) and fast-cycle RAM (FCRAM) appear headed for success, because, even with their higher than commodity DRAM costs, vendors can still price them lower than SRAM and make a profit.

REFERENCE

A. Dipert, Brian, "The high-end PC looks for a home," *EDN*, Nov 24, 1999, pg 145.

CAN'T GET ENOUGH?

If the combination of this and other *EDN* articles hasn't quenched your thirst for DRAM knowledge, here are some more good sources of information.

BENCHMARKS: BAH HUMBUG!

In early November 1999, analyst Bert McComas published a series of benchmarks comparing a Micron Samurai-based double-data-rate (DDR) synchronous-DRAM (SDRAM) PC with a preproduction i820- and 400-MHz Direct Rambus DRAM (DRDRAM)-based counterpart. Not surprisingly, DDR SDRAM's high peak bandwidth translated to faster numbers on the stream-intensive memory-access and memory-to-AGP-transfer benchmarks he used. How those numbers translate to real-life results, to what extent they reflect chip-set differences *other* than those related to the main memory subsystem, and their relevance to non-PC memory applications is anyone's guess. Interestingly, a 32-component DRDRAM configuration performed worse than its 16-chip counterpart, probably because of power-management-caused latencies, but again the performance degradation's relevance to real-life memory usage is unclear and might reflect immature chip-set drivers. Anyway, head to www.inqst.com, and have a look.

Not surprisingly, Rambus shot back with its own set of benchmarks at November 1999's Comdex. Comparing a 600/133-MHz local-bus Pentium III system using PC133 virtual-channel memory (VCM) with a 400-MHz DRDRAM-based counterpart shows a slight performance boost. Independent reviews of DRDRAM-equipped PCs support Rambus' conclusions. The applications that Rambus chose, however—AutoCad 2000, PhotoShop 5.5, and texture-intensive *Quake 3 Arena Test*—are data-rich, make heavy use of both the CPU and the graphics subsystem, and are the same types of applications Intel highlights with its Pentium III processor. You probably wouldn't see the difference running *Solitaire*, an Office ap-

plication, or a Web browser. Rambus likes to point out the even better results with a 733/133-MHz local-bus Pentium III machine, results that look better yet with a shift from an i820-based PC to one running with the i840 chip set. However, the company neglects to give PC133 VCM comparison numbers. You'll find it all at www.rambus.com.

Intel has its own set of numbers at <http://developer.intel.com/update/archive/issue23/stories/top2.htm>, along with a lot of SDRAM and DRDRAM data at developer.intel.com/design/chipsets/memory. In fairness to the Rambus camp, today's benchmarks, which their developers created with today's system capabilities in mind, don't fully take advantage of the DRDRAM memory-subsystem capabilities, specifically for handling multiple concurrent transactions. This situation has also occurred in the graphics world, in which advanced chip architectures have to wait for the application-programming interfaces, benchmarks, and real-life applications to catch up. Programmable logic has the same problem when a special feature of a CPLD or FPGA goes to waste because a compiler doesn't bother to use it.

To find out more about the Stream benchmark, point your browser at www.cs.virginia.edu/stream. Several benchmarking-discussion threads are also on the Internet newsgroup comp.arch. Search on www.deja.com for topics "Stream on i820/Coppermine" and "Fuel on the Rambus versus DDR fire." Finally, keep an eye out for PC reviews in your favorite computer magazine, especially those comparing Micron's system based on PC133 VCM and Via's Apollo Pro133A chip set with DRDRAM- and i820-based alternatives and comparing conventional workstations with i840-based ones.

BITS AND BYTES

Incompatibilities among vendors and a general lack of systems understanding marked the early days of SDRAM. These problems led to insufficiently detailed and impossible-to-design-with specifications. Intel stepped in, defined the PC66 and PC100 specification sets, and worked with third-party testing labs to ensure a

high degree of system and memory compatibility.

When Intel focused its efforts on DRDRAM, the SDRAM industry was left with a void, just as they were ramping PC133 and finalizing their initial DDR SDRAM work. Advanced Memory International (AMI2, www.ami2.com) has filled that void by concentrating on areas that the Joint Electronic Device Engineering Council standardization committees don't tackle, such as module and system-board Gerber layouts and infrastructure coordination. In addition to vendor-independent AMI2's Web site, you can find standards information at www.jedec.org. Via Technology has validated a number of PC133 memories for use with its chip set; www.viatech.com/news/pc133develop.htm contains the list and other PC133 info.

Several DRAM manufacturers have taken the lead in advocating next-generation SDRAM architectures. IBM's PC133 and DDR SDRAM Web sites, respectively, at www.chips.ibm.com/products/memory/memory_enablement/register.html and www.chips.ibm.com/products/memory/ddr_memory_enablement/register.html are worth checking out, as is Micron's PC133 site at www.micron.com/mti/msp/html/pc133index.html, which includes links to free HDL models.

Curious to find out how much memory the DRAM manufacturers think you need in your PC? The Council on Computing Power offers a RAM configurator at www.rammatters.com/configurator, and you'll find Kingston's version at www.kingston.com/ecom/assessor/index.htm. Speaking of Kingston, if you'd like to read a good basic tutorial on DRAM, take a look at their Ultimate Memory Guide at www.kingston.com/tools/umg.

Want to see what other pundits think about the future of DRAM? Head over to SiliconTech's Web site, especially www.silicontech.com/pinnacle.html, and watch presentations delivered via RealVideo streams. Finally, if you'd like to wade through the legalese and read all the gory details on the resolution of Enhanced Memory's lawsuit against NEC, you'll find them at http://dockets.usitc.gov/eol/publ_under_document_type_337,number_421,ID_199911120036.

formance-oriented computer users who have overclocked their 100-MHz systems. Enhanced has also submitted its ESDRAM-Lite proposal to the JEDEC committee developing the DDR-II specification, and, like MoSys, Enhanced is also beginning to pursue the embedded-licensing business with its 24-nsec, random-cycle enhanced-SRAM (eSRAM) core.

NEC's VCM, which it introduced in late 1997, focuses not on improving the inherent DRAM-array speed, but, like EDRAM, on transforming the on-chip sense amp array into a more robust cache (Reference 4). Unlike EDRAM, however, the system controller, not the memory itself, manages cache-maintenance functions, such as activate, prefetch, restore, and precharge. The system controller can, if appropriate, assign multiple associative-mapped cache arrays, or "virtual channels," to portions of the same memory bank (Figure 4). Having more virtual channels than a standard SDRAM's single per-bank row cache reduces page thrashing in multimaster environments, such as when one or multiple CPUs, a network adapter, several mass-storage subsystems, the graphics accelerator, and other peripherals all contend for memory control.

The caching similarities between Enhanced's and NEC's architectures led to a recently resolved patent-infringement lawsuit that Enhanced initiated. NEC's 64-Mbit VCM includes 16 1-kbit virtual-channel pages. The company plans to begin to make available for sampling a follow-on single-data-rate (SDR), 128-Mbit device sometime this quarter, with a

DDR version closely following it. Of all the low-latency variants, VCM requires the most extensive DRAM-controller modifications. A simple VCM port of an SDRAM controller might actually *reduce* performance because of VCM's additional clock cycle of nonchannel access latency and the need for an explicit write-back cycle. NEC has therefore put significant efforts into both core-logic support and alternative sourcing.

The company obtained approval for VCM from JEDEC as an open standard and claims it will seek no licensing fees from other interested manufacturers. NEC signed a second-sourcing agreement with Infineon in late 1998. (No Infineon-developed parts have yet appeared.) In early 1999, NEC entered a joint-marketing and -sales agreement with Hyundai. Note, however, the careful wording of the partnership; Hyundai is unsure about whether it will design and manufacture its own VCMs or simply resell NEC-developed parts. NEC has also secured chip-set support or future planned support from Acer Laboratories, Silicon Integrated Systems, and Via Technologies. And it's a safe bet that the recently announced Hitachi/NEC alliance will further develop the VCM concept.

Whereas NEC is pursuing main-memory applications for VCM, Fujitsu is content for now to focus its FCRAM efforts on graphics, networking and other non-PC-main-memory applications. FCRAM forgoes VCM's onboard cache and instead mimics the other half of ESDRAM's bag of tricks: a high-speed DRAM core. Like ESDRAM, FCRAM can obtain its performance gains with only

minimal enhancements to a memory controller intended for SDRAM. Fujitsu achieves fast random access, like the MoSys approach, by segmenting each bank into miniarrays (Figure 5). A three-stage pipelined architecture also enables the next random access to begin while the part is outputting previous data.

The initial FCRAM test chip included an SRAMlike demultiplexed bus (Reference 5). Fujitsu's first commercial device, an eight-bank, DDR, 64-Mbit memory with a 32-bit data bus, uses a more conventional RAS/CAS multiplexed approach. It comes in a DDR SDRAM-compatible "P" version and an optimized "N" version with better read/write-turn-around performance and modified stub-series-terminated-logic (SSTL) I/O buffers that require no off-chip resistors. Fujitsu estimates that the 64-Mbit FCRAM's die is 5% larger than that of a comparably sized conventional DDR SDRAM. The FCRAM specifies a 30-nsec random-access cycle, compared with a standard SDRAM's 60-nsec cycle, and Fujitsu targets clock rates as high as 200 MHz, translating to 400-Mbps/pin burst-transfer rates. Fujitsu believes that, by powering up only a portion of the array, FCRAMs will have lower active power consumption than standard SDRAMs. Toshiba plans to act as an alternative source for the Fujitsu device, and both companies will in the future design and market additional FCRAMs.

WHAT'S THE FREQUENCY, KENNETH?

With CPU local-bus clock rates now topping 133 MHz, main-memory-bus burst bandwidth must keep pace to avoid sapping the performance of the all-important cache-line-fill operation. Multi-CPU configurations do their part to keep the main-memory pipeline busy. And plenty of other high-speed buses, either here or on the way, will also contend for memory attention.

Double a standard DRAM's number of on-chip banks, add a digital delay-locked loop (DLL), and make the internal data bus twice as wide as the external interface would indicate, and you've got DDR SDRAM with data transferred on both edges of the input clock. You can find DDR techniques in devices other than SDRAMs: SLDRAM, MDRAM, and DRDRAM use a conceptually similar approach. Synchronous SRAM also employs this performance-boosting tech-

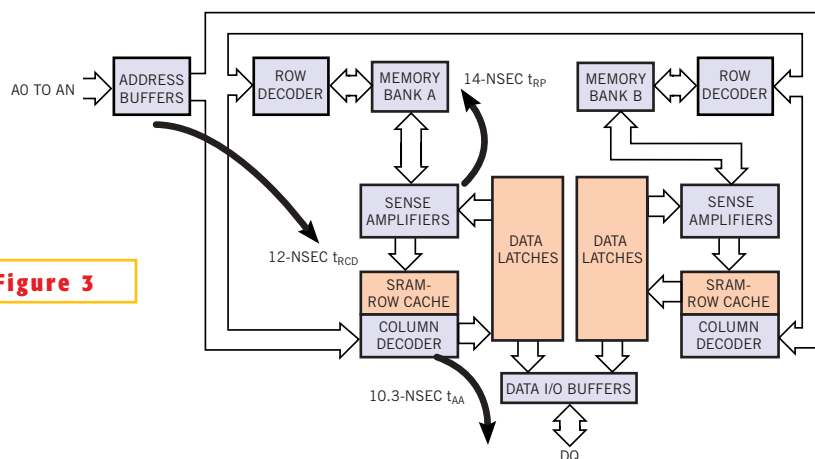


Figure 3

ESDRAMs, combining a fast-access array with on-chip cache and control logic, target SRAM (courtesy Enhanced Memory Systems).

nique. Therein lies one key to DDR SDRAM's appeal: The approach is well-understood and to at least some degree proven.

Vendor-proprietary DDR SDRAM variants began appearing in 1997. The DRAM vendors, when they haven't been frantically working on their DRDRAM designs, have spent the last 30 months or so ironing out industry-standard DDR SDRAM specifications, including resolving the unidirectional-versus-bidirectional data-strobe debate (see sidebar "Can't get enough?"). Initial main-memory devices, operating at 133 MHz, deliver a burst bandwidth of 2.1 Gbytes/sec across a 64-bit data bus—33% faster than a single-channel, 400-MHz DRDRAM interface and 75% faster than a 300-MHz DRDRAM channel—but requiring significantly more pins to do so. Even higher speed DDR SDRAMs, such as Hyundai's 64-Mbit, 166-MHz parts and Infineon's 32-Mbit, 200-MHz devices, target lightly loaded memory arrays.

Graphics companies such as Nvidia, one of the few remaining mainstream vendors still using fully featured synchronous-graphics RAMs (SGRAMs)—with its GeForce and Quadro—are gobbling up as many chips as the DRAM manufacturers can cost-effectively supply. A few memory companies, such as IBM, Mitsubishi, and Samsung, create high-density SDRAMs and DDR SDRAMs from the same die, configuring the part during testing. An 8-bit-interface SDRAM and a 4-bit-interface DDR SDRAM, for example, both require an 8-

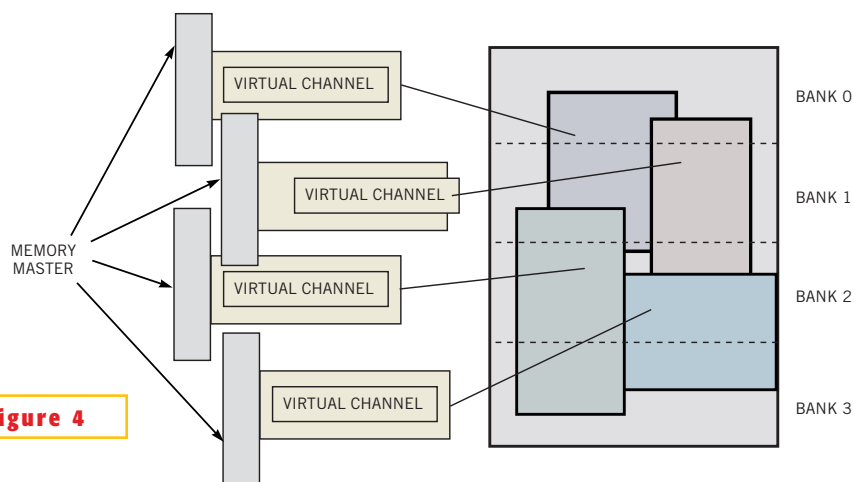


Figure 4

VCM's virtual channels can span multiple DRAM banks, and the system can allocate multiple channels to the same DRAM row (courtesy NEC Electronics).

bit internal data bus, and these companies are willing to place a cost burden on each—at least currently—higher volume SDRAM. This burden comes from DDR SDRAM's additional bank overhead and on-chip DLL. In exchange, the vendors can leverage one design effort across multiple products.

DDR SDRAM advocates tout the memory's high degree of backward-compatibility with the SDR chips that vendors have been shipping for more than five years and with the infrastructure to support them. DDR-SDRAM-DIMM pinouts have moved from 168 to 184 contacts, but the DIMMs themselves have the same dimensions as in the past, and the contact-to-contact pitch also remains the same. JEDEC added signals for differential clocks and more grounds by

replacing one of the keying notches and placing extra contacts on both sides of the SDR-SDRAM-module pinout. You may see DDR-SDRAM modules specified with a designation of "PC2100" instead of the previous "PC133." This name represents an attempt to go beyond the clock rate and DRDRAM's component spec number and advertise the module's peak 64-bit bandwidth capability.

You cannot use the low-voltage TTL interface at these high data-transition speeds on a heavily loaded bus, so DDR SDRAM moves to the low-swing SSTL approach, which high-speed SRAMs also use. For speeds beyond 133 MHz, vendors envision a move from TSOP to a more performance-friendly BGA package, an interim step that many of them are calling DDR-1.5 and that may also include a differentially driven data strobe. And, for DDR-II, additional evolutionary features will probably make their appearance. These features may include 1.8V I/O buffer and core voltages; an electrical interface derived from SLDRAM (but not a packetized bus); and a modification of the "posted-CAS" write protocol, which overcomes wasted clock cycles during write-to-read bus transitions.

Micron not only did its own SLDRAM design, but also its own PC-core-logic chip sets, both for internal validation and for use by its systems division. Micron's latest north-bridge effort, code-named Samurai, bundles a number of leading-edge features, including support for both AGP 4X and PC2100 DDR SDRAM. InQuest and Micron's benchmarks indicate impressive




Figure 5

FCRAM's array subdivision approach (a) both improves random-access latency and reduces chip power consumption compared over standard DDR SDRAM (b) (courtesy Fujitsu Microelectronics).

performance, although the data-stream-focused numbers they chose to publish, as the benchmark's Stream name implies, heavily exploit DDR SDRAM's high peak throughput. Micron has licensed the Samurai design to Via Technologies, a company that is also leveraging DDR SDRAM

and other design work that AMD did for its Athlon CPU's companion chip sets.

And what about RDRAM? Congratulations to Intel, Rambus, and their DRAM partners are in order, because systems based on the single-channel i820 and dual-channel i840 chip sets began shipping late in the fourth quarter of 1999 (Reference 6). This introduction marks a three-year time span from the first public disclosure of the companies' partnership and an approximately six-month delay from the originally planned production schedules. In spite of the development dollars Intel tossed at the DRAM manufacturers, only a few of them are in production with qualified parts. Due in part to the production delays and in part to a desire to make RDRAM die size and cost as close to SDRAM as possible, vendors offer the parts in 128- and 144-Mbit densities, not the original 64- and 72-Mbit versions.

A few memory companies claim that they're completing their Rambus designs only to fulfill the terms of their licensing agreements and have no plans to take the parts to production. Acer is the only other PC-core-logic-chip-set vendor to have taken a Rambus license, and neither it nor CPU-vendor licensees such as AMD and National Semiconductor plan to support the technology. Other PC-core-logic manufacturers are, at least for now, content to support PC133 and DDR SDRAM. What's going on? (See sidebar "A murky crystal ball.")

First is the two-versus-three-slot issue. Some of the original i820-based motherboard designs included slots for three Rambus inline-memory modules (RIMMs). Intel and Rambus couldn't resolve the rarely seen, but existent, data corruption that occurred when the system was executing certain combinations

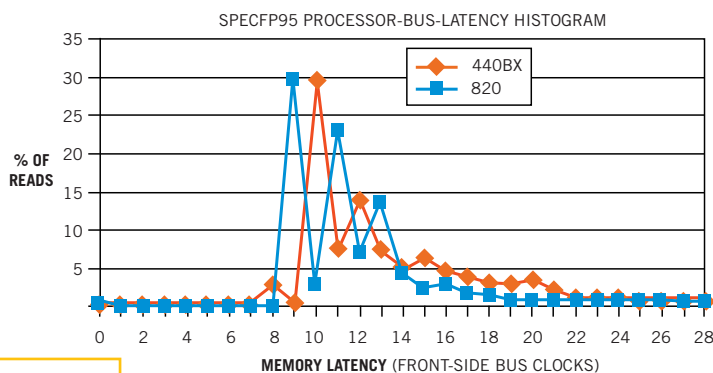


Figure 6

PC100 best-case latency beats 400-MHz DRDRAM, but DRDRAM is the overall better option, according to this benchmark (courtesy Intel).

of operations. Depending on whom you talked to, the companies blamed this problem on signal attenuation due to impedance mismatches or to lost signal-to-clock edge synchronization. The companies' testing revealed that, with two-slot configurations, they could eliminate the problem by reducing the variability across the range of multiple vendors' devices, modules, density combinations, and motherboard layouts.

Not wanting to give the new technology even a tinge of a black eye, Intel delayed the launch by six weeks and dictated consistent two-slot designs to its customers. Given that each Rambus channel has a slot-independent 32-device limitation, the two-slot restriction means nothing more than a perceived decrease in system-upgrade flexibility. Single-channel configurations based on 128-Mbit DRDRAMs can support as much as 512 Mbytes of system memory, and use of the 82804AA memory-repeater hub bumps single-channel densities as much as 1 Gbyte. One positive angle to the two-slot move, pragmatically speaking, is that the approach requires one fewer continuity RIMM to place an additional cost burden on Rambus-based systems.

Intel and Rambus have appended two more sets of specifications to the 400-MHz PC800, which the October 1997 architecture announcement trumpeted, and PC800 now comes in both 40- and 45-nsec t_{RAC} variants. The 300-MHz PC600 delivers 600 Mbps/pin or 1.2 Gbytes/sec/channel maximum bandwidth, whereas the 355.5-MHz PC700, which i840-based motherboards don't use, delivers 711-Mbps/pin and 1.4-Gbyte/sec/channel speeds. In compari-

son, the maximum bandwidth of a 64-bit PC133 bus is slightly less than 1.1 Gbytes/sec, identical to the i820's slowest memory option, the PC600 DRDRAM running at 266 MHz. But bandwidth isn't everything. Latency is also important, and here the Rambus advantage is less clear. Latency is particularly a problem when a design uses the wait-state-adding memory-repeater hub. Ram-

bus' assertions that DRDRAM latency is less than that of SDRAM or DDR SDRAM apply only to the fastest specified DRDRAM component version. These claims also ignore signal-flight time to and from the memory and the additional system overhead a system needs to convert the address and data buses between serial (memory) and parallel (system) formats.

At the fall 1999 Developer Forum, Intel indicated that, in a best-case scenario, a PC100 CAS-2 SDRAM's initial access latency using the i440BX chip set would be one clock faster than RDRAM coupled with the i820 chip set. This scenario has both chip sets hooked to a 400-MHz Pentium II CPU with a 100-MHz local bus (Figure 6). Extrapolations of this data to 133-MHz-local-bus CPUs and corresponding PC133 CAS-2 SDRAM, PC-2100 CAS-1.5, or CAS-2 DDR SDRAM would probably show an even wider best-case gap. The data, however, also indicated that, according to Intel's benchmarks, RDRAM performed better than SDRAM across the range of main-memory latencies a system would typically see. Latencies came from access contention, refresh collisions, and similar conditions, and RDRAM's performance advantage became particularly apparent with streaming data-intensive applications.

In retrospect, why did Intel go down the Rambus road? For some clues, look at the Pentium III processor, also still being defined when Intel and Rambus began seriously working together in 1996 (Reference 7). Features such as the streaming single-instruction-multiple-data extensions; enhanced buffering, prefetching, and multiple-transaction queuing; the ability to bypass the cache

for noncode transfers; and the multi-processor "hooks" imply that Intel was predicting that average DRAM-access characteristics would differ greatly from the simple code fetches of the past. Some guesses at the type of rich multimedia environment Intel assumed would be mainstream and pervasive in 1999 include voice recognition, 3-D graphical

user interfaces, high-resolution still- and video-image capture and editing, and software-based DVD and high-definition-TV decoding.

AGP was also under development in 1996. The i740 and i752 graphics chips and the i810 core-logic chip set with integrated graphics functions all rely on main memory for 3-D-texture storage.

Low-end i810 variants use main memory for *all* frame-buffer functions. AGP execute mode both increases the amount of AGP and main-memory traffic and shifts the average access pattern toward long data streams instead of short cache-line fills. No other graphics vendor has to date adopted this exclusive texture-storage reliance on AGP execute mode, however.

FOR MORE INFORMATION...

For more information on products such as those discussed in this article, circle the appropriate numbers on the Information Retrieval Service card or use EDN's InfoAccess service. When you contact any of the following manufacturers directly, please let them know you read about their products in EDN.

CHIP MANUFACTURERS

Advanced Memory International Inc

1-408-436-1300
www.ami2.com
Circle No. 324

Alliance Semiconductor Corp

1-408-855-4900
www.alsc.com
Circle No. 325

Austin Semiconductor Inc

1-512-339-1188
www.austinsemiconductor.com
Circle No. 326

Council on Computing Power

www.computingpower.org
www.rammatters.com
Circle No. 327

Enhanced Memory Systems Inc

A Division of Ramtron International Corp
1-719-481-7000
www.edram.com
Circle No. 328

Fujitsu Microelectronics Inc

1-408-922-9000
www.fujitsumicro.com
Circle No. 329

G-Link Technology

1-408-492-9068
www.glinktech.com
Circle No. 330

Hitachi Semiconductor America Ltd

1-408-895-3045
www.hitachi.com/
semiconductor
Circle No. 331

Hyundai Electronics America

1-408-232-8000
www.he.com
Circle No. 332

IBM Microelectronics Corp

1-914-765-1900
www.chips.ibm.com
Circle No. 333

Infineon Technologies

1-408-777-4500
www.infineon.com
Circle No. 334

Integrated Silicon Solution Inc

1-800-379-4774
www.issusa.com
Circle No. 335

Kingmax Micro Technology Inc

1-949-454-9888
www.kingmax.com
Circle No. 336

Micron Technology

1-208-368-3900
www.micron.com
Circle No. 337

Mitsubishi Electronics America Inc

1-408-730-5900
www.mitsubishichips.com
Circle No. 338

Mosel Vitelic Inc

1-408-433-6000
www.moselvitelic.com
Circle No. 339

MoSys Inc

1-408-731-1800
www.mosys.com
Circle No. 340

NEC Electronics Inc

1-408-588-6000
www.necel.com
Circle No. 341

Oki Semiconductor

1-408-720-1900
www.okisemi.com
Circle No. 342

Rambus Inc

1-650-944-8000
www.rambus.com
Circle No. 343

Samsung Semiconductor

1-408-544-4000
www.samsungsemi.com
Circle No. 344

Silicon Magic

1-408-331-8000
www.simagic.com
Circle No. 345

Toshiba America Electronics Components Inc

1-949-455-2000
www.toshiba.com/taec
Circle No. 346

MODULE MANUFACTURERS

Apacer Memory America Inc

1-408-434-3311
www.apacer.com
Circle No. 347

Kingston Technology Co

1-714-435-2600
www.kingston.com
Circle No. 348

Memory Card Technology

1-760-603-0150
www.memory-card-
technology.com
Circle No. 349

SiliconTech Inc

1-949-476-1130
www.sti-oem.com
Circle No. 350

Smart Modular Technologies Inc

A Division of Solectron Corp
1-510-623-1231
www.smartm.com
Circle No. 351

Viking Components Inc

1-949.643.7255
www.vikingcomponents.com
Circle No. 352

Wintec Industries

1-510-498-4165
www.wintecind.com
Circle No. 353

OTHER COMPANIES MENTIONED IN THIS ARTICLE

Acer Laboratories Inc
www.acerlabs.com

AMD Inc
www.amd.com

Apple Computer Corp
www.apple.com

Cypress Semiconductor Corp
www.cypress.com

InQuest Market Research
www.inqst.com

Intel Corp
www.intel.com

National Semiconductor Corp
www.national.com

Nintendo of America Inc
www.nintendo.com

Nvidia Corp
www.nvidia.com

Semico Research Corp
www.semico.com

Silicon Integrated Systems Corp
www.sis.com.tw

Sony Corp of America
www.sony.com

Tessera Inc
www.tessera.com

Transmeta
www.transmeta.com

Via Technologies Inc
www.viatech.com

SUPER CIRCLE NUMBER

For more information on the products available from all of the vendors listed in this box, circle one number on the reader service card.
Circle No. 354

Considering that DRAM was a supply-constrained, expensive commodity in early 1996, Intel probably underestimated the severe price-per-bit decline that would soon follow, a tumble that would only accentuate the difference in cost between SDRAM and DRDRAM.

Chip cost is a complex combination of variables: die size, package, and module; test; yield and inventory. The last few years have clearly shown little to no correlation between DRAM price and cost. You also need to consider related system costs when comparing various memory options, such as terminators, clock drivers, connectors, DRAM controllers, pc-board layers, and trace pitch. The 128- and 256-Mbit DRDRAMs will include 32 internal banks versus DDR DRAM's four, and DRDRAM's even wider internal data bus also increases the incremental die size—currently averaging approximately 10% more than that of DDR SDRAM at equivalent densities, bus widths, and technologies.

Vendors report that, for the fastest 400-MHz DRDRAM chips now leaving fabs, only around half of them and, in some cases, only a quarter of them run at 400 MHz. The next process step, which will take the devices to 0.2- μ m or lower lithographies, should significantly ease this limitation on yield. Costs will also naturally decrease as vendors ramp into high-volume production. And treat the DRAM manufacturers' claims with a healthy dose of skepticism. Companies are prone to overstating costs and understating yields in an attempt to slow or completely quash the Rambus volume ramp-up or in the desire to keep prices high as long as possible and make as much return on their DRDRAM investments as they can.

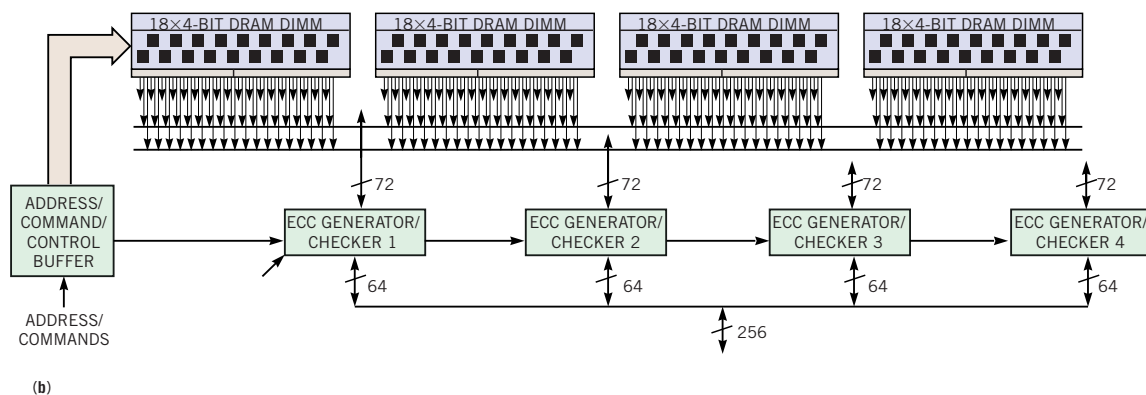

Rambus has addressed one of its key initial criticisms: the ability for a high-end system to continue operating if an entire DRAM device fails (Figure 7). The company claims that its interleaved-data scheme incurs no die-size penalty for 256-Mbit and larger devices. But the limited per-Rambus-channel-component count requires multiple-channel server chip sets. Having gigabytes of DRAM per system means that memory pricing is as critical here as it is in a low-end "free" PC. Per-server DRAM density is increasing more quickly than in any other segment of computing, thanks to exploding popularity of the Internet. And server- and workstation-focused chip-set vendors, such as Reliance Computer Corp, are in no hurry to pay Rambus royalties, either, although Compaq is supposedly working on an Alpha- and DRDRAM-based server design. Apple, which designs its own chip sets, is also likely in the future to use DDR SDRAM.

THE REST OF THE PACKAGE

The latency-versus-bandwidth concern is secondary in PCs to the most important criterion: that the system contain enough DRAM. If the computer doesn't have enough main memory to hold all simultaneously running applications and

their data, the operating system swaps them out of the virtual-memory partition on the hard drive. This reliance on slow mass storage is enough to bring even the highest powered machine to its knees. The DRAM vendor-comprised Council on Computing Power is attempting to educate PC end users, who normally focus only on CPU clock frequencies. In attempting to squeeze more chips onto a DIMM, DRAM and module manufacturers are going down two packaging paths. One option stacks multiple memory die, such as Hitachi's multichip, double-density packaging approach, or packaged chips, such as Hitachi's tape-carrier package and stacked TSOPs, on each other (Reference 8). This approach takes advantage of the established packaging infrastructure, but it cannot shake the reputation for being a costly, low-yielding approach.

Companies such as Kingmax are touting a board-space-reducing alternative: BGA packaging (Figure 8). Because BGA packages have lower impedance than TSOPs, this approach can increase the percentage of modules that meet PC133 and DDR2100 specifications and may also offer thermal-dissipation benefits. However, although BGA-packaged-module functional yields may be higher than



Interleaved data mode gives DRDRAM-based systems (a) chip-killing capability equivalent to those with SDRAM (b) (courtesy Rambus Inc).

those of stacked TSOPs, the BGA packages and modules themselves are inherently more expensive. Regardless of the approach, high-density DDR SDRAM modules not only register the address and control buses, which incurs an additional one-clock latency, to minimize system loading, but also buffer the data lines with fast-responding FET switches.

DRDRAMs already use BGA packaging, mostly in the form of Tessera's μ BGA, and DDR SDRAMs are also headed toward using BGA. Micron's and Toshiba's DRDRAMs, with their respective FBGA and chip-sized packages, are, the vendors claim, lower cost alternatives to μ BGA with comparable electrical and thermal characteristics. The μ BGA packages create headaches for third-party module manufacturers because of the packages' die-specific bond-pitch and layout characteristics, which vary among DRAM manufacturers and even among devices from the same manufacturer. Encapsulated-die BGA packages also include a small internal translation board that enables a die-independent pinout and pitch. Packaging is the least of the module vendors' RDRAM concerns, though, as they struggle with the necessary equipment-upgrade investments to test the fast parts. Expect a shakeup in the memory-module industry if DRDRAM becomes as successful as Rambus hopes.

Designers of embedded systems understandably view DRAM as equal parts curse and blessing. Its cost per bit is more attractive than that of SRAM. As DRAM gets faster, SRAM's traditional performance lead also becomes less important, especially if refresh-generated power consumption is a secondary concern or the memory is accessed so frequently that it rarely sees low-power standby modes. However, with main memory in PCs, workstations, and servers consuming a dominant percentage of all DRAM, the manufacturers are focusing their product- and technological-development efforts on these customers to obtain the most bang for their buck.

Historically, you align your memory plans with the trends of your PC-engineering brethren to avoid the hassles of limited suppliers, high prices, and device obsolescence. However, your per-chip density




Figure 8

Using stacked die (a), stacked chips (b) (courtesy Hitachi Semiconductor America), and small-footprint chips (c) (courtesy Kingmax Micro Technology Inc) let you squeeze more memory on a module.

needs might lag behind those of PCs, and your decreased system-memory granularity also might require a wider memory bus or an interface with higher operating frequency. Manufacturers that claim that their per-chip prices remain flat, regardless of per-chip density growth provide little consolation to those who would like to follow the decreasing cost-per-bit trends to reduce system cost. They also don't help if package or pinout changes at higher densities force you to redesign your board layout or if your memory controller can't handle the bigger chips.

Application-specific memory-access patterns, which may radically diverge from those of a PC, define your relative priority of initial access latency versus subsequent burst performance. Power consumption may be more critical to you than to someone designing a wall outlet-powered workstation, and your board-space requirements may also be more stringent. The vendors are all at least beginning to seriously talk about servicing opportunities outside the PC, although in most cases they have few deeds to back up their words.

Chips with granularity-friendly 32-bit data buses and fast-core memories, such as ESDRAM and FCRAM are notable exceptions, though. You might be surprised by how well a memory for, say, point-to-point interconnect-graphics-memory subsystems also serves your needs. Your range of DRAM options will likely continue to improve, both as ASIC-embedded memory becomes a mainstream technology and as discrete-memory vendors react to

predictions that computing as a percentage of total yearly DRAM consumption will shrink. □

REFERENCES

1. Dipert, Brian, "Advanced DRAM puts you in the fast lane," *EDN*, Oct 9, 1997, pg 52.
2. Dipert, Brian, "Embedded memory: the all-purpose core," *EDN*, March 13, 1998, pg 34.
3. Dipert, Brian, "SRAMs strive to specialize," *EDN*, Nov 5, 1998, pg 62.
4. Dipert, Brian, "DRAM gets to the core of the problem," *EDN*, April 1, 1998, pg 14.
5. Dipert, Brian, "DRAM redesign: not just plastic surgery," *EDN*, July 2, 1998, pg 20.
6. Dipert, Brian, "Rambus reality, finally," *EDN*, Sept 30, 1999, pg 19.
7. Dipert, Brian, "Fight fashion with foresight," *EDN*, Sept 16, 1999, pg 52.
8. Dipert, Brian, "Wider is better," *EDN*, July 22, 1999, pg 16.
9. Collins, Hansel A, and Ronald E Nikel, "DDR-SDRAM, high-speed, source-synchronous interfaces create design challenges," *EDN*, Sept 2, 1999, pg. 63.
10. Resso, Michael J, "The right technique yields critical Rambus signal-integrity measurements," *EDN*, Aug 5, 1999, pg. 61.

ACKNOWLEDGEMENTS

Special thanks to Bill Gervasi from Transmeta, Bert McComas from InQuest Market Research, and Desi Rhoden from Advanced Memory International for sharing their analyses and perspectives with me.

You can reach
 Technical Editor
 Brian Dipert at
 1-916-454-5242,
 fax 1-916-454-5101,
 bdipert@
 pacbell.net

