


# Embedded memory: the all-purpose core

**BRIAN DIPERT, TECHNICAL EDITOR**



The era of the system on a chip is here. So say the ASIC vendors, and along with them a herd of semiconductor analysts and journalists. Time for a reality check. It'll be a few years yet before you can download intellectual-property cores and data books from multiple vendors' Web sites, mix and match the functions you need, add your own proprietary circuits, and tie the whole thing together in one ASIC as "easily" as you do today with multiple chips on a system board.

The single- or few-chip system design is, however, a natural extension of the integration that began when early semiconductor pioneers hooked a few transistors together on a silicon substrate. Back then they called it "small-scale integration"; now, the name has become "ultralarge-scale integration." There's little debate anymore about *whether* system-scale integration will occur, yet industry pundits still argue over *when* it will become widespread. Industry consortiums, such as the Virtual Socket Interface Alliance, have emerged to speed this process.

A significant portion of today's system-on-a-chip hype comes from the semiconductor vendors themselves. These companies are quickly ramping production capability on 0.25- $\mu\text{m}$ -drawn (sub-0.2- $\mu\text{m}$ -effective) and even finer processes. Simultaneously, they're worrying about where they'll find customers with multimillion-gate designs that can fill the silicon in a way that is not bondpad-limited. Given that the average logic core consumes only 10,000 or so gates, you can easily understand their concern. You can also see why embedded memory is such an

attractive answer to their—and your—problems.

Memory arrays can contain millions or tens of millions of transistors. Memory is also the most universal “core,” giving vendors maximum return on their R&D investment. Almost all systems contain some type of both volatile and nonvolatile storage. And, if ASIC vendors also sell discrete memory chips, embedded-memory capability gives them a convenient conversion message when they tell you that they’re obsoleting older low-density memory chips.

### Lots of positives, a few caveats

Forget what the vendors want you to do. If the embedded-memory concept doesn’t hold water, then marketing hype alone won’t ensure its success. Fortunately, embedded memory *does* make sense today for a few—and in the future even more—applications. If you heed Moore’s Law (which says that the number of transistors that you can cost-effectively put on a piece of silicon doubles every 18 months), you know that chips built on advanced lithographies are most cost-effective when measured by cost-per-bit at higher transistor counts. This trend toward higher and higher densities is no problem for PC main DRAM with software that seems able to consume all of whatever density chips are available. However, plenty of embedded applications also use DRAM, and their density, voltage, interface, and other needs lag—sometimes significantly—behind those of the PC. The same scenario holds for most other memory technologies. Digital cellular phones today use 1 to 2 Mbytes of flash memory, but plenty of embedded systems are content with 128 kbytes or less. Embedded memory gives you an escape plan should your list of discrete-component vendors begin to dwindle.


Board space and reliability are two other key embedded-memory advantages. Fewer chips on a board result in a smaller system footprint to implement a given function. Mobile-computing and communications devices, disk-drive-controller ICs, PCMCIA-card chip sets, and automotive-electronics modules are just a few of the

applications reduced board space motivates. An integrated memory-plus-logic, single-die IC avoids the esoteric, costly packaging used in multichip modules. Lower chip count might also mean higher system-manufacturing yields and a lower probability of premature system failure—either within a device or in the interdevice connections.

Embedded memory also neatly solves the granularity problem. For example, graphics frame buffers, whose sizes depend on required resolution, color depth, and degree of 3-D support, rarely fall into tidy megabyte partitions (Reference 1). RAM densities also normally quadruple from generation to generation, accelerating the gap

between what you might need and what is economically available. On-chip RAM densities are more flexible, and you can sometimes even manipulate the aspect ratio, orientation, and location of the RAM array (or multiple subarrays) for the simplest, fastest interconnection with embedded logic or minimal ASIC die size.

System software provides another illustration: Its ideal embedded size, allowing for upgrades through a system’s lifetime, might be 1.3 Mbytes. Good luck finding a chip in exactly that density. Unless you use one of a few specialty ICs, such as Catalyst’s (Sunnyvale, CA) 1.5-Mbit 28F015 and 28F150 flash memories, you have to either strip features to make the code fit



**Trident Microsystems’ Cyber9520 3-D mobile graphics controller combines approximately 280,000 logic gates, analog functions, and 2 Mbytes of SDRAM on one die (courtesy Samsung).**

## EMBEDDED MEMORY

or pay for memory bits you don't need. Or you could choose a third path: embedding the memory on the ASIC along with the system processor or DSP core in whatever density you require.

What about performance? Even though intrachip routing delays are becoming dominant in defining ASIC speeds, these delays are still significantly shorter than the alternative of driving signals off-chip to a separate memory bank. Widening the off-chip i n t e r - connect bus to increase bandwidth negatively impacts the ASIC pin count. You could instead try running that wide bus between memory and logic inside the chip, where the bus operates even faster. Instead of using expensive external SRAM, you might be able to get by with slower embedded DRAM. If you previously used embedded SRAM because it was the only available technology, the advent of cheaper one-transistor embedded DRAM may encourage a shift in strategy, although you might also need to include on-chip refresh circuitry.

Embedded memory gives you flexibility in how you implement DRAM refresh. For the lowest average power consumption, you could choose a minimal-refresh array that trades off density by using larger cell capacitors. Ultra-high-density arrays minimize cell-capacitor size with a corresponding increase in refresh frequency. If the data stored to the DRAM needs to be valid for only a short time, such as in video-buffering applications, you might even choose to completely eliminate refresh. Additional customization options include refresh modes, burst length and order, and access latencies. You include only what you need and strip the logic you don't require.

Many graphics controllers with embedded DRAM use 128-bit and wider buses, and ASIC vendors claim that their 0.25- $\mu$ m embedded-DRAM processes will allow internal buses as wide as 1024 bits. Try doing that in a multichip design! You might even choose to separate the internal input and output-data buses for improved overall throughput. Alternatively, you can take advantage of the fast intra-ASIC delays to craft a narrower inter-

### @ a glance

- Embedded memory is the first step on the path to a system on a chip, the ultimate destination of semiconductor integration.
- Low power, high performance, enhanced flexibility, minimal board space, and increased reliability highlight embedded memory's advantages.
- Embedded-memory ASICs may not cost less than a multichip alternative at the chip level, but you might save on system costs.
- Logic and memory don't naturally coexist on the same silicon, requiring vendors to make trade-offs.
- A number of competing technologies characterize both volatile and non-volatile embedded memory.

connect bus than you might need with an off-chip alternative, a decision that simplifies the design, shrinking die size and lowering cost. With embedded DRAM, you need no longer use a pinout- and bandwidth-restrictive multiplexed address bus, although if the vendor's embedded-DRAM approach uses already-available array designs, you may not have a choice.

All those off-chip interconnections require hefty I/O buffers to overcome package and board-trace impedances. The resultant increased power draw means limited battery life, excessive heat dissipation, and reduced reliability. NeoMagic (Santa Clara, CA) estimates that a graphics controller with embedded DRAM consumes 500 to 750 mW—roughly 25% the power of its multichip alternative, which consumes approximately 2.5W. Although you typically still need to drive the embedded-memory ASIC with 2.5V or lower voltages, because of the advanced lithographies these chips use, you might be able to omit a heat sink and a few fans your system ordinarily needs to keep from self-destructing. Eliminating fast-switching external signals also reduces radiated EMI.

One factor that you should not count on, however, is that a memory-plus-logic ASIC design is cheaper than its multichip alternative. IC cost stems

from a complex combination of factors, including die size, volume, and test time. The price you pay relates only peripherally to silicon cost. With embedded-memory ASICs, you integrate technologies with different sets of potential silicon defects, so logic and memory-yield impacts multiply. Because both the number of die and the die yield per wafer exponentially decrease with increased die size, two smaller chips, memory and logic, might be cheaper than their larger, integrated alternative, even though you pay twice for packaging costs.

A more integrated and, therefore, customer-specific chip cannot benefit from the same volume-manufacturing efficiencies as an off-the-shelf alternative. This fact may not be a problem for vendors that already derive much of their business from standard-cell ASICs. These vendors are simultaneously manufacturing numerous unique, quick-turnaround, relatively low-volume, and short-lived wafer sets in a factory for many different customers. However, a logistics nightmare could occur for a company entering the embedded-memory market from a commodity-memory background. The commodity-memory market has much longer design cycles for products that are in production for several years. The larger the percentage of die size that the memory array defines, the more influence the memory market has on the chip's prices. Depending on which way the market is headed, this situation could be good or bad for you.

Keep in mind that oversupply and the resulting rapidly falling memory prices over the last year or so have left semiconductor economic conditions in an "unnatural" state. Rumor has it that some discrete-memory manufacturers are selling silicon below cost just to keep their fab lines full and avoid losing market share. Poor economic conditions are driving the silicon vendors' interest in embedded memory, as they strive to find more profitable uses for their engineers and equipment. System engineers, on the other hand, might find it more economically attractive to open their memory-sourcing options to more than one company by staying with a discrete-memory

design.

If discrete-memory supply and demand return to a more balanced state and if demand exceeds supply, as was the case in 1995, the resulting impact on embedded-memory-ASIC prices and supply sources may not be to your liking. Reducing chip count cuts parts inventory, system-manufacturing complexity, board space, and power consumption, so an embedded-memory ASIC could reduce your system cost as well as your chip cost.

### Divergent needs and integration

Design some logic, add a memory array, hook it all up, and you're done, right? Well, it's not quite that easy (see box "Libraries and compilers complete the design"). However, the substantial amount of logic, such as registers, counters, comparators, and state machines, on today's high-volume synchronous DRAMs (SDRAMs) shows that merging logic and memory is possible, although you must consider some trade-offs. Logic and memory are fundamentally different technologies with sometimes-contradictory demands, beginning with their interconnect requirements. Look at the die plot for a controller or processor, and you see a number of discrete logic blocks scat-

tered across the die in an irregular fashion (Figure 1). A memory array, on the other hand, is symmetrical and repetitive and requires far less transistor interconnection than does logic. This difference explains why logic designs of a few million gates garner large amounts of industry attention, whereas DRAMs with 64 million array transistors plus other circuitry are commonplace.

The more interconnect required, the more metal layers are necessary to keep die sizes reasonable. Leading-edge logic processes provide a minimum of four metal layers, with as many as six in some cases, but normally include only one or two polysilicon layers. Standard memory processes, on the other hand, require four or more polysilicon layers but only one or two metal layers. You can use polysilicon in the floating gate of EPROM, EEPROM, and flash memory; the resistors of a four-transistor SRAM cell; and most vendors' stacked-capacitor DRAM cells. Each metal or polysilicon layer you add to a base semiconductor substrate, plus its corresponding isolation dielectric, adds cost, complexity, and fabrication time. Metal interconnect requires a wider pitch in logic to minimize impedance and signal-routing delays. That same pitch would make a DRAM array

unnecessarily large, so memory vendors typically use narrower, slower metal lines.

Transistor-gate oxide thicknesses also differ for logic and memory cells. Standard logic transistors have thin oxides with low turn-on threshold voltages to minimize switching time and maximize performance. Many memory technologies prefer slower, thicker oxide transistors. High turn-on threshold voltages minimize off-transistor leakage current, a primary determinant of the required DRAM refresh frequency. Because memory arrays of given linear dimensions contain significantly more transistors than their logic equivalents, low leakage current is also critical to minimizing device power consumption. Thick oxides improve data-retention characteristics for EPROM, EEPROM, and flash memory. Equally important, they help the cells withstand the effects of high voltages and corresponding electrical fields during repeated programming and erasing cycles.

Most memory technologies retain information within the active array transistors. Examples include the SRAM cell's multiple-transistor differential configuration and the presence or absence of stored electron charge on the

## LIBRARIES AND COMPILERS COMPLETE THE DESIGN

Once you decide to include embedded memory in your next ASIC project, choose the right technologies and vendors, and settle on the densities you need, what's the next step? Face it, you'd rather spend your time working on custom logic, not crafting a memory array. Consider, then, obtaining an externally developed memory core. For simulation purposes, your design requires functional, timing, and electrical (Spice or equivalent) models. And, when you're ready to go to layout, the dimensions and pinouts of the memory black box are also useful.

If you are in a partnership with an ASIC vendor, you typically obtain these libraries and compilers from that vendor. In most cases, you obtain a library suite for the exact memory or memories your design requires. However, if you plan to eventually design multiple ASICs using the same supplier and process generation, you might want to consider spending the extra money on a memory compiler, which lets you input variables, such as density, bus width, and aspect ratio, and outputs an optimized model set. If your design parameters change, you can also recompile the memory model instead of begging or buying additional libraries. Brian Barrera, marketing director for Mentor Graphics' Inventra Division, offers a rule of thumb that compilers make three or more designs per manufacturer and process.

Model compilers typically support standard memory cores, such as ROM and six-transistor SRAM. DRAM, EEPROM, and flash memory fur-

ther limit your options. With DRAM, this situation occurs because the cell designs are vendor-proprietary, and, with EEPROM and flash, this situation occurs because they require more complex, location-inflexible peripheral circuitry. Virage Logic, one exception, is in a partnership with Tower Semiconductor to supply embedded-EPROM compilers.

Obtaining silicon from a foundry rather than an ASIC supplier typically limits the available library and compiler options. However, several third-party model manufacturers can fill the gap. Ironically, these same companies may be the developers of the models you obtain from ASIC vendors.

Such options include the Memory Modeler from Denali Software; Memory Builder and Memory Model Builder from the Inventra Division of Mentor Graphics; MemPro from the Logic Modeling Group of Synopsys; Custom-Touch Memory Compilers from Virage Logic; and products from Aspec Technology, Cascade Design Automation, and Virtual-Silicon Technology. Other memory-library vendors include Azalea Microelectronics, Duet Technologies, Nara Technologies, Nurlogic Design, Phoenix Technologies' Virtual Chips Division, and Puyallup Integrated Circuit Co. (More seem to appear every week.) Artisan Components, another memory-library provider, supplies single- and dual-port SRAM cores targeting 0.25- $\mu\text{m}$  ASIC processes with 300-MHz speed and low-power options. The company also recently signed a licensing agreement with VLSI Technology.

## EMBEDDED MEMORY

EPROM, EEPROM, or flash-memory transistor floating gate. DRAM, on the other hand, stores data in a passive capacitor, relying on signal boosting for accurate sensing. For this reason, DRAM also requires thick-oxide transistors and is comparatively sensitive to coupled and injected noise, which can cause problems when you combine DRAM with fast-switching and noisy logic transistors and interconnect. The boosted sensing scheme also makes it more difficult to scale DRAM to lower operating voltages. In contrast, logic prefers low voltage because it enables smaller, faster, less power-hungry transistors.

EPROM, EEPROM, and flash memory all require numerous high voltages for programming and erasing. The required isolation circuitry to switch, place, and remove these voltages within a subset of the memory array and to isolate them from the rest of the chip, including the logic subsystems, can significantly complicate fabrication and chip design. If external high-voltage generation is not an option, on-chip charge pumps consume precious silicon real estate. Internal-scan ATPG, 1149.1, and JTAG, built-in self-test, quiescent current, fault grading, and other test methods, although necessary in the highest density embedded-memory arrays, also use incremental logic resources. Redundancy is another controversial addition. Although this feature is necessary to boost yields in the early stages of manufacturing, ASIC vendors are more likely to remove it and reduce die size than they would be with discrete-memory chips. Smaller die may equal lower cost, unless re-

moving redundancy adversely affects the silicon yield—a tricky balancing act, especially with DRAM's current slim profit margins.

When evaluating various vendors' logic-performance claims on an embedded-memory process, make sure you compare worst-case, not typical, specifications. Operating voltage, trace impedance, and fan-out are all factors that can significantly alter propagation delays. Some companies offer several logic-cell configurations; faster cells may be less silicon-efficient than slower alternatives from the same or another vendor's library. Also, ensure that the gates themselves are similar. Some companies specify performance using inverters, and others use inherently slower, but more functionally robust, two-input NAND gates.

### Volatile-memory options


Historically, ASIC vendors have emulated DRAM cells using multitransistor-logic cell structures with transistor gates used to form capacitors (Figure 2a). This approach lets you use a standard logic process without degrading transistor switching performance. Some ASIC vendors offer only multitransistor-cell DRAM. However, although multitransistor-cell DRAM is smaller than a four-transistor, two-resistor, or six-transistor SRAM cell, it consumes many times the silicon area of a one-transistor, one-capacitor DRAM equivalent. For designs requiring more than a few megabits, you need an alternative.

Embedded DRAM has reignited the long-smoldering controversy between

advocates of stacked-capacitor and buried-trench-capacitor DRAMs. In a stacked-transistor construction, the capacitor, comprising multiple polysilicon layers, lies above the transistor (Figure 2b). At advanced lithographies, complex "folding" design techniques ensure that the capacitor continues to get smaller while retaining adequate charge-storage characteristics to limit refresh requirements. Trench-capacitor DRAMs take advantage of the third wafer dimension, depth, and project the charge-storage device down into the silicon substrate, using either diffusion or ion implanting (Figure 2c).

Trench-capacitor defenders, such as IBM, Siemens, and Toshiba, point out that the approach delivers higher capacitance per given linear silicon area, resulting in smaller cells that are crucial in an integrated logic-plus-memory ASIC. Because the capacitor is at or below the level of the transistor, additional masking steps can focus exclusively on metal interconnect for on-chip logic, and the technology requires no complicated polysilicon planarization to avoid metal deformation. Polysilicon deposition is also a high-temperature process that potentially alters the beta characteristics of the already-constructed transistors underneath the stacked capacitor.

Stacked-capacitor-DRAM promoters, such as Mitsubishi and Samsung, don't deny the trench vendors' claims of smaller cells and fewer required polysilicon layers. However, constructing the trench capacitors requires high precision. They believe that this need for precision makes the capacitors unman-



**Multitransistor (a), one-transistor stacked-capacitor (b), and one-transistor buried-trench-capacitor (c) cell structures offer alternative methods of embedding DRAM.**

## EMBEDDED MEMORY

ufacturable in processes beyond approximately 0.18  $\mu\text{m}$ . They also point out that most DRAM vendors use the stacked-capacitor approach.

Mitsubishi was one of the first vendors to aggressively promote embedded-RAM capability, whereas other manufacturers remained content with their high-margin, discrete-DRAM component business. The company reports that, by the end of last year, it had shipped 5 million of its eDRAM embedded-DRAM ASICs. For design processes larger than 0.25  $\mu\text{m}$ , Mitsubishi could use normal thin-oxide logic transistors

to construct the DRAM array, relying on lower operating voltages to counteract the higher leakage current. At 0.25  $\mu\text{m}$  and smaller, however, Mitsubishi's HyperDRAM process contains a dual-oxide thickness—vs standard logic's single oxide—and adds three processing steps. Standard logic and HyperDRAM share the same metal pitch and therefore the same logic-layout libraries but have different logic timing. The 0.25- $\mu\text{m}$ , dual-oxide, two-input NAND gate delay is 85 psec with a fan-out of 2, still substantially shorter than the 150-psec delay in the previous-generation, 0.35-

$\mu\text{m}$ , single-oxide process.

In addition to providing embedded-DRAM to its ASIC customers, Mitsubishi leverages its eDRAM processes among other internal groups. Some of the products the company builds with these processes include the 3D-RAM graphics chip, combining 1.25 Mbytes of DRAM and a high-performance ALU, and the M32R/D 32-bit RISC CPU with 2 Mbytes of DRAM. The M32R/D was an early pioneer of what may prove to be a common  $\mu\text{P}$  design technique (see box "Intelligent RAM"). Mitsubishi's technology-exchange agreement with

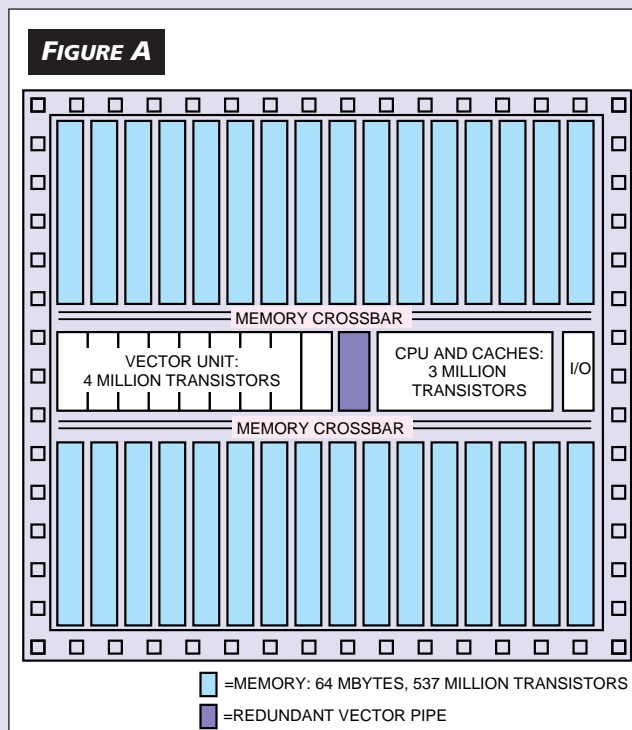
## INTELLIGENT RAM

Anyone who has taken a computer-architecture class in the last decade or so or who has worked with Sun's (Mountain View, CA) SPARC CPU should be familiar with the name David Patterson. Patterson, a professor at the University of California—Berkeley since the late 1970s, is co-author of the seminal work on  $\mu\text{P}$  design, **Computer Architecture: A Quantitative Approach**. He was an early RISC-processor advocate, and his work on the RISC 1 and 2 ultimately led to the SPARC  $\mu\text{P}$ . Patterson was also a pioneer in RAID, network clusters, and multi-processor systems.

Now that the electronics industry has widely adopted his RISC, RAID, networking, and multi-CPU techniques, Patterson has turned his attention—and graduate students' efforts—to what he believes will be the next major evolution in  $\mu\text{P}$  design, intelligent RAM (IRAM). He provides several reasons for this belief:

- On-chip SRAM consumes significantly more die area than does embedded DRAM for the same storage density.
- Cache memory does an imperfect job of eliminating external access latency, and increasing cache size produces diminishing returns.
- Some programs and data types have poor locality of reference and subsequent low cache performance even with large cache sizes.
- Interchip interfaces between  $\mu\text{Ps}$  and discrete memory not only underperform the CPU core but also burn large amounts of power.
- Sub-0.2- $\mu\text{m}$  lithographies will enable semiconductor vendors to cost-effectively squeeze a billion or so transistors onto a chip.

Why, Patterson wonders, continue to cache a subset of DRAM contents



**The Berkeley V-IRAM-2 gives a glimpse into the potential architecture of future  $\mu\text{Ps}$ .**

insufficient for full-featured computers. In this case, the CPU could page off-chip secondary DRAM contents into and out of the embedded-DRAM array, much like today's CPUs swap information between the SRAM cache and external DRAM. High-speed connections between multiple parallel processors, each with embedded DRAM, might be another method of accomplishing the same objective. Patterson hopes to fabricate a 256-Mbit IRAM within the next two years, assuming that he can find a semiconductor vendor interested enough to provide the necessary financing and resources to make his dream a reality. (For more details, see <http://iram.cs.berkeley.edu>.)

in on-processor SRAM? Why not simply put the DRAM on the CPU?

The Berkeley V-IRAM processor targets a 0.13- $\mu\text{m}$  process lithography and 400 sq-nm die (Figure A). This theoretical case study combines a vector processor running at 1 GHz and a scalar unit on one device. On-chip memory includes 1 Gbit of high-speed DRAM divided into 32 sections, each with 16 2-Mbit banks and crossbar switching. The complex pipelined synchronous-DRAM architecture, with estimated 20-nsec latency and 4-nsec cycle times, attempts to minimize random-access delays, much like the approach MoSys uses with its MDRAM. Each vector-processor element has its own corresponding DRAM subarray.

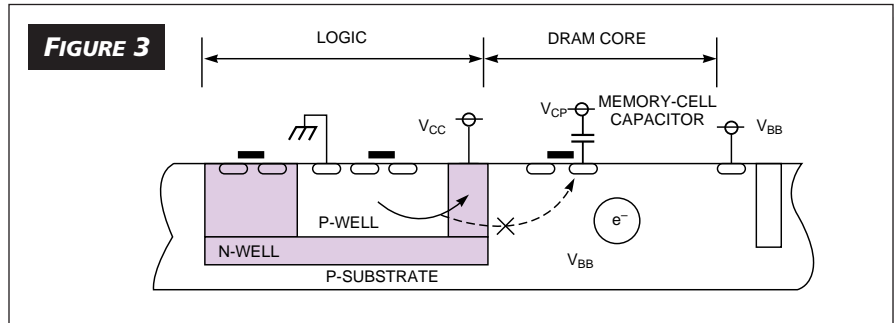
With IRAM, Patterson primarily focuses on mobile devices, which need energy- and space-efficient  $\mu\text{Ps}$  to run real-time, multimedia applications. Patterson admits that by the time 0.13- $\mu\text{m}$  processes are in production, 1 Gbit of DRAM may be

## EMBEDDED MEMORY

Motorola, currently covering 0.35- $\mu\text{m}$  embedded DRAM and versions 2 and 3 of Motorola's ColdFire CPU core, should lead to some interesting integrated chips as well. Mitsubishi predicts that, for internal customers, products employing the 0.25- $\mu\text{m}$  HyperDRAM process with four layers of metal and supporting as much as 64 Mbits of DRAM will see first silicon by this month. Mitsubishi's triple-well approach isolates the DRAM substrate from bias and injected noise originating in the logic and standard SRAM circuits and also contains any required DRAM-substrate bias (Figure 3). The company also hopes to begin manufacturing embedded-DRAM wafers on its upcoming 0.18- $\mu\text{m}$  process by year's end.

Samsung, another stacked-capacitor-DRAM supporter, is shipping merged DRAM and logic on MDL90, a 3.3V, 0.35- $\mu\text{m}$ , three-layer- or four-layer-metal process. The company based this process on the one Samsung uses to manufacture its 500-MHz Alpha 21164 CPU. MDL90 has 150-psec gate delays with a fan-out of 2 and three layers of polysilicon. It supports as much as 24 Mbits of extended-data-out DRAM or SDRAM. A four-well approach ensures isolation between the logic and DRAM subsections. Like some other ASIC suppliers, Samsung offers an assortment of logic cores, such as the Oak (Sunnyvale, CA) DSP and ARM7, the Zilog (Campbell, CA) Z-80 CPU, and the Philips (Sunnyvale, CA) 80C52 CPU, to accompany Samsung's embedded-DRAM capability. Key customers include graphics-controller vendors Chips and Technologies (San Jose, CA) and Trident Microsystems (Mountain View, CA).

Toshiba has historically been the most visible embedded-trench-capacitor-DRAM advocate. Production shipments for the company's first 0.5- $\mu\text{m}$  embedded-DRAM ASICs began in 1995. Early customers included Silicon Graphics (Mountain View, CA), which used the process in its Indigo Impact chip set. This chip set has 70,000 logic gates and 8 Mbits of DRAM with a 128-bit data bus and has a 1.6-Gbyte/sec bandwidth. The company estimates that its 2.5V TC240D 0.25- $\mu\text{m}$  process, which debuted last April, delivers 80-



**Processes supporting more than two wells enable electrical isolation of high-speed logic from the noise-sensitive DRAM array (courtesy Mitsubishi).**

psec high-speed gate delays with a fan-out of 5 and typical power consumption of 0.17  $\mu\text{W}$ /gate/MHz. Memory density, as with other manufacturers, depends on internal-bus width, including optional parity support, array aspect ratio, the number of memory macros, and logic-gate count (Figure 4 and Table 1).

Products using the 3.3V, 0.35- $\mu\text{m}$  TC210D process can have embedded-DRAM densities as high as 32 Mbits and high-speed gate delays of 176 psec with a fan-out of 5. TC240D and TC210D offer three performance and power-logic-gate options and come in embedded-array and standard-cell variants. Toshiba also licenses its embedded-DRAM technology to foundry Chartered Semiconductor Manufacturing. This license also potentially lets Chartered manufacture wafers for Toshiba. NeoMagic, which formed an exclusive partnership with Mitsubishi for NeoMagic's early graphics controllers, recently added Toshiba as an alternative source for NeoMagic's new ICs that integrate as much as 2 Mbytes of frame-buffer DRAM. NeoMagic develops its chips using silicon from 100%-DRAM processes. The company created its own logic libraries for both the Mitsubishi and Toshiba designs, relying on both vendors for the embedded-memory arrays.

Siemens, another member of the trench-capacitor-DRAM consortium, takes the jointly developed technology in a somewhat different direction. First, the company is not a broad-based ASIC supplier, although it supports a limited number of foundry opportunities. Second, Siemens' embedded-DRAM-ASIC

process is virtually identical to the one the company uses with its discrete DRAM components. Bob Pierce, Siemens' director of emerging memories, admits that using the same process on both products results in slower logic. He estimates 100-psec gate delays with a fan-out of 2 on the company's 0.24- $\mu\text{m}$  process with two polysilicon layers and four-layer metal. However, because the approach uses thicker oxide transistors, the embedded-memory ASICs can operate at a system-friendly 3.3V, not 2.5V or lower, as with merged-logic alternatives. The company plans to focus on applications requiring high-density memory, minimal logic, and moderate performance. Siemens, like Mitsubishi, also supplies stand-alone devices with embedded DRAM, such as the SDA-9388 picture-in-picture decoder, the Scarabaeus scan converter with 5.2 Mbits of DRAM, and the Voice Engine speech-recognition IC containing 1 Mbit of DRAM and the Oak DSP core.

IBM offers three-transistor DRAM from its cell-based-ASIC group and one-transistor trench-capacitor DRAM from its foundry division. As with Siemens, the trench-capacitor architecture is virtually identical to the DRAM-component process, and the 0.25- $\mu\text{m}$  basic memory building block is a 1-Mbit array with a 256-bit data-bus interface. IBM Fellow Howard Kalter believes that, because of high DRAM volumes, using a DRAM process yields low silicon costs. Also, using the process for embedded DRAM gives the company extensive knowledge about the process and its characteristics. The technology delivers 90-psec gate delays with a fan-out of 3, and Kalter is confident that

## EMBEDDED MEMORY

IBM can easily scale its trench-capacitor approach to processes lower than 0.15  $\mu\text{m}$ . Kalter disputes other vendors' assertions that isolating logic noise from the DRAM substrate *always* requires triple-well process technology. The company also offers metallized ROM, EPROM, and EEPROM cells.

Many other ASIC vendors are working closely and quietly with a few customers, slowly ramping up their embedded-DRAM yields, or are still performing R&D. Silicon Magic (Santa Clara, CA) has leveraged Oki's memory expertise on the 1.25-Mbyte DRAM F/X256 and MSM-7680 (Oki part number) graphics controllers. NEC and Nintendo (Kyoto, Japan) are reportedly collaborating on a next-generation game controller containing a MIPS core and embedded DRAM. NEC's current chip set uses external Rambus (Mountain View, CA) DRAM to ensure sufficient bandwidth between the CPU and memory. NEC also presented a derivative of its Virtual Channel DRAM at this year's International Solid State Circuits Conference. Hitachi's H673M 0.35- $\mu\text{m}$  process offers flexible, 256-kbit embedded-DRAM-module granularity, alongside the company's H8 and SH CPU cores. Accelerix (Ottawa), another graphics-chip-set manufacturer, working with Mosaid Technology (Ontario), will be one of Taiwan Semiconductor Manufacturing Co's (TSMC's) first customers for TSMC's 0.35- $\mu\text{m}$  embedded-DRAM process. Silicon Motion (San Jose, CA) plans to use United Microelectronics Corp's (UMC's) process for Silicon Motion's LynxE, containing 2 Mbytes of resident DRAM and a 64-bit external SDRAM bus for frame-buffer expansion.

Taking NeoMagic's design techniques a step further, S3 (Santa Clara, CA) crafts both its own proprietary logic and RAM libraries. Although S3's approach requires more upfront design work, the company claims that the resultant proprietary memory compiler

can quickly generate an optimized S3SRAM array for any foundry, giving S3 widespread sourcing flexibility.

Notebook-computer graphics controllers have dominated much of the early activity in embedded DRAM because their manufacturers ship them in high volume, which is attractive to ASIC vendors, and because these controllers support lower resolutions and pixel color depth than do desktops. As a result, they require only 1- to 2-Mbyte frame-buffer densities. Notebooks also value low power and minimal board space, two qualities in which embedded DRAM shines over multichip alternatives. Desktop PCs might also later adopt embedded memory, because with the advent of Intel's (Santa Clara, CA) Accelerated Graphics Port, graphics-resident frame-buffer densities may begin to flatten.

Looking beyond graphics, Sony, which offers no discrete DRAMs, has formed a partnership with Oak Technology on Oak's OTI-9220 CD-ROM controller, containing 1 Mbit of Sony-developed embedded DRAM. Meanwhile, Yamaha (Shizuoka, Japan) intends to replace embedded SRAM with DRAM in future generations of audio and video processors, reducing die size and cost. Integrated Device Technology (IDT, Santa Clara, CA) has also begun limited use of Fusion, which

employs MoSys' (Sunnyvale, CA) MDRAM (MoSys DRAM) core as an SRAM alternative in networking chip sets, in IDT's 77V400 1.24-Gbps switching memory. Other potential applications for both DRAM and other embedded memories include cellular phones and pagers, fax machines, digital answering machines, PC core-logic chip sets, set-top boxes, digital cameras, video and arcade games, and ink-jet and laser printers.

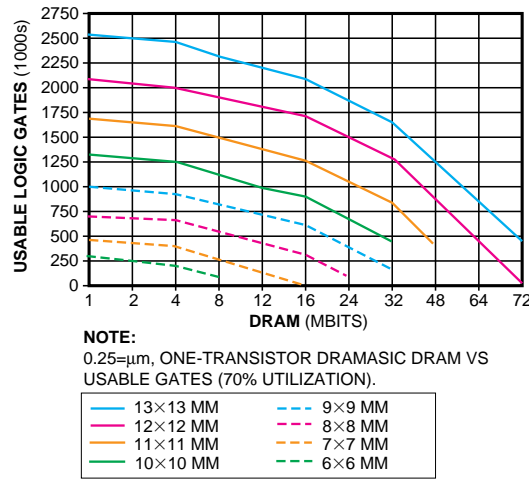
For DRAM companies without ASIC divisions, the path to embedded DRAM may involve a partnership or a subcontractor relationship with a foundry. LSI Logic and Micron Technology (Boise,

ID) take the partnership route, basing this alliance on LSI's new 0.25- $\mu\text{m}$ , six-layer-metal G11 logic process and Micron's 0.25- $\mu\text{m}$  DRAM process. Under the arrangement, wafers travel from LSI Logic to Micron for embedded-DRAM fabrication and then back to LSI Logic for metal deposition and final processing. Both companies can use the merged technology in custom and standard products. For LSI Logic, the agreement provides access to leading-edge DRAM capability without the need for internal R&D. Micron's interest hinges on the ability to diversify its product line beyond commodity DRAM into more stable, higher margin logic ICs.

### Plenty of SRAM choices

You can relatively easily implement six-transistor SRAM cells on a standard logic process, as the common practice of integrating Level 1 cache inside a  $\mu\text{P}$  demonstrates. Implementation alternatives include compiled and customized cells. Compiled cells offer fast turnaround times, whereas customized cells deliver the most efficient array architectures. Using four-transistor, two-resistor SRAM or multiport-SRAM configurations limits your ASIC-vendor and -foundry options. Four-transistor, two-resistor cells can be smaller than six-transistor alternatives, an

**FIGURE 4**



**Embeddable-memory density depends on lithography, logic complexity, and desired die size (courtesy Toshiba).**

## EMBEDDED MEMORY

attractive attribute in high-density arrays. Four-transistor, two-resistor SRAMs can also sometimes operate faster, but, because of passive-resistor-generated current, they burn more power than six-transistor alternatives. However, four-transistor, two-resistor configurations require additional polysilicon plug layers to implement the resistors. Because it's more difficult to scale down resistors, such as DRAM capacitors, than transistors during lithography shrinks, the four-transistor, two-resistor size advantage may disappear in time.

The primary limitation of SRAMs having two or more ports is not process compatibility but library-model availability. Although vendors claim to support multiport SRAM in their marketing collateral, make sure that models exist in the densities, configurations, and array aspect ratios that your design requires or ensure that the vendor can quickly create a model within your budget constraints. This limitation occurs to an even greater extent with content-addressable-memory cells popular in networking designs, which are more logic-intensive than standard SRAMs.

Chip Express, a gate-array-ASIC manufacturer, provides dedicated, on-chip, diffused SRAM arrays on its CS2000 and CX2001 families. Single-port block sizes for the two families are 32 and 8 kbits, respectively, and dual-port configurations combine adjacent single-port modules within a block. Block configuration involves using the vendor-supplied memory compiler, which can implement both ROM and RAM functions. Modules have separate read and write ports, variable bus widths, and optional synchronous operation. Worst-case module access times on the 0.6- $\mu\text{m}$  process are 5 nsec at 5V and 6 nsec at 3.3V.

Ferroelectric RAM (FRAM) is finally moving into limited production after having long lingered in the R&D labs and is also experiencing activity in the embedded-memory arena (**Reference 2**). Fujitsu and Motorola are integrating FRAM alongside their respective processor cores. Fujitsu reported its efforts at the recent International Electron Devices Meeting in Washington. Rohm also promotes FRAM as part of its

embedded-memory portfolio. Much of the interest in FRAM stems from its combination of the traditional benefits of volatile DRAM (bit alterability, a single-transistor cell, and high read and write performance) and nonvolatile ROM (low power). FRAM is more complex to manufacture than DRAM because of the specialized capacitor-dielectric material FRAM uses. With sufficient industry focus and if DRAM moves to similar dielectrics at advanced lithographies, this difference could disappear.

### Nonvolatile memory?

ASIC vendors have for years supported ROM, as they have six-transistor SRAM. Manufacturers configure diffused ROM along with the rest of the device transistors as one of the first steps in wafer fabrication. Metallized ROM configuration occurs during later metal-deposition steps or even during postfabrication testing. Diffused ROM uses silicon more efficiently but results in longer leadtimes from order to device delivery and less flexible vendor response to customer code changes than does metallized ROM.

The availability of user-programmable microcontrollers from a variety of vendors shows that embedding EPROM is possible on a logic process. However, compare prices of an EPROM or even one-time-programmable (OTP) controller with its ROM equivalent, and you quickly determine that this technique is not cost-efficient. Some vendors' slow nonvolatile memories require redundant, embedded, shadow RAM for reasonable code-execution performance, resulting in an expensive multimemory die. In reality, high

EPROM microcontroller prices probably result from both higher EPROM cost and the manufacturers' desire to maximize profit margin on a historically low-volume device.

Tower Semiconductor licenses the Alternate Metal Virtual Ground (AMG) technology from Waferscale Integration's (WSI, Fremont, CA) with plans to improve embedded EPROM's cost-effectiveness (**Reference 3**). Fairchild Semiconductor (Santa Clara, CA) and SGS-Thomson Microelectronics (Lincoln, MA) also license AMG, which has one metal ground per two bit lines vs the traditional approach of one ground per bit line and which eliminates field isolation in the EPROM core. Halving the number of required metal lines, thereby reducing metal-pitch limitations, and eliminating diffusion-pitch restrictions means that the polysilicon pitch is the primary remaining determinant of transistor size (**Figure 5**).

WSI estimates that on a 0.6- $\mu\text{m}$  process, an AMG EPROM and flash memory would have cell sizes of 1.56 and 2.08 sq  $\mu\text{m}$ , respectively, compared with 2.9 sq  $\mu\text{m}$  for traditional EPROM and 1.44 sq  $\mu\text{m}$  for mask ROM. Tower Semiconductor is producing embedded-EPROM ASICs on a process the company also uses to manufacture Fairchild Semiconductor's advanced EPROMs. Tower Strategic Marketing Manager Eran Liron points out that cost is critical because most of his customers prefer to avoid expensive windowed ceramic packages. Using embedded EPROM as an OTP memory in cheaper plastic packages lets the company's technology compete more directly with masked ROM.

Tower Semiconductor also has high

**TABLE 1—EMBEDDED-DRAM OPTIONS  
IN TOSHIBA'S 0.25- $\mu\text{m}$  TC240D DRAMASIC PROCESS**

Memory capacity	Memory organization (bits)	No. of banks
1 to 32 Mbits/macro (1-Mbit increment/bank)	$\times 64/72$ , $\times 128/144$	One, two, or four
2 to 64 Mbits/macro (2-Mbit increment/bank)	$\times 128/144$ , $\times 256/288$	One, two, or four
4 to 128 Mbits/macro (4-Mbit increment/bank)	$\times 256/288$ , $\times 512/576$ , $\times 1024/1152^1$	One, two, or four

<sup>1</sup>  $\times 512/576$  using two macros,  $\times 1024/1152$  using four macros.

## EMBEDDED MEMORY

hopes for its upcoming embedded-flash-memory process, which will see first silicon in the second quarter of this

year. Tower based its flash-memory approach not on AMG, but on the results of a licensing agreement the

company signed with technology-development partner Saifun Semiconductors (Natanya, Israel). The Saifun/

## FOR MORE INFORMATION...

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

**Embedded-memory-ASIC vendors and foundries****Atmel Corp**

San Jose, CA  
1-408-441-0311  
fax 1-408-436-4300  
www.atmel.com

**Circle No. 360****Chartered Semiconductor Manufacturing Ltd**

Singapore  
1 011 65 362 2838  
fax 1 011 65 362 2903  
www.csminc.com

**Circle No. 361****Chip Express Corp**

Santa Clara, CA  
1-408-235-7353  
fax 1-408-988-0513  
www.chipexpress.com

**Circle No. 362****Fujitsu Microelectronics Inc**

San Jose, CA  
1-408-922-9000  
fax 1-408-432-9044  
www.fujitsumicro.com

**Circle No. 363****GEC Plessey**

San Jose, CA  
1-408-451-4724  
fax 1-408-451-4710  
www.gpsemi.com

**Circle No. 364****Hitachi America Ltd**

Brisbane, CA  
1-415-244-7848  
fax 1-415-583-4207  
www.halsp.hitachi.com

**Circle No. 365****Hyundai Electronics America**

San Jose, CA  
1-408-232-8000  
fax 1-408-232-8125  
www.he.com

**Circle No. 366****IBM Microelectronics Corp**

Armonk, NY  
1-914-765-1900  
fax 1-914-892-5334  
www.chips.ibm.com

**Circle No. 367****Kawasaki LSI USA Inc**

San Jose, CA  
1-408-570-0555  
fax 1-408-570-0567  
www.klsi.com

**Circle No. 368****LG Semicon**

San Jose, CA  
1-408-432-5000  
fax 1-408-432-6067  
www.lg.co.kr

**Circle No. 369****LSI Logic Corp**

Milpitas, CA  
1-408-433-8000  
fax 1-408-433-8989  
www.lsillogic.com

**Circle No. 370****Lucent Technologies**

Allentown, PA  
1-610-712-4331  
fax 1-610-712-4209  
www.lucent.com

**Circle No. 371****Mitsubishi Electronics America Inc**

Sunnyvale, CA  
1-408-730-5900  
fax 1-408-732-9382  
www.mitsubishichips.com

**Circle No. 372****Motorola Corp**

Phoenix, AZ  
1-602-732-2852  
fax 1-602-732-5020  
www.design-net.com

**Circle No. 373****NEC Electronics Inc**

Santa Clara, CA  
1-408-986-1020  
fax 1-408-588-6374  
www.nec.com

**Circle No. 374****Oki Semiconductor**

Sunnyvale, CA  
1-408-720-1900  
fax 1-408-720-1918  
www.okisemi.com

**Circle No. 375****Rohm Electronics**

Antioch, TN  
1-615-641-2020  
fax 1-615-641-2022  
www.rohmelectronics.com

**Circle No. 376****Samsung Semiconductor**

San Jose, CA  
1-408-954-7000  
fax 1-408-954-7870  
www.samsung.com

**Circle No. 377****Siemens Components Inc**

Cupertino, CA  
1-408-777-4500  
fax 1-408-777-4988  
www.siemens.com

**Circle No. 378****S-MOS Systems**

San Jose, CA  
1-408-922-0200  
www.smos.com

**Circle No. 379****Sony Semiconductor Co of America**

San Jose, CA  
1-408-955-6569  
www.sel.sony.com

**Circle No. 380****Symbios**

Fort Collins, CO  
1-970-223-5100  
www.symbios.com

**Circle No. 381****Taiwan Semiconductor Manufacturing Co Ltd**

Hsinchu, Taiwan  
1 011 886 3 5780221  
fax 1 011 886 3 5781546  
www.tsmc.com.tw

**Circle No. 382****Texas Instruments Inc**

Austin, TX  
1-800-477-8924, ext 4500  
www.ti.com

**Circle No. 383****Toshiba America Electronics Components Inc**

Irvine, CA  
1-714-455-2000  
fax 1-714-859-3963  
www.toshiba.com

**Circle No. 384****Tower Semiconductor USA Inc**

San Jose, CA  
1-408-551-6500  
fax 1-408-551-6509  
www.towersemi.com

**Circle No. 385****United Microelectronics Corp**

Hsinchu, Taiwan  
1 011 035 782258  
fax 1 011 035 774767  
www.umc.com.tw

**Circle No. 386****VLSI Technology Inc**

San Jose, CA  
1-408-434-3100  
fax 1-408-434-7584  
www.vlsi.com

**Circle No. 387****Winbond Electronics Corp**

Hsinchu, Taiwan  
1 001 886 3 5770066  
fax 1 001 886 3 5792668  
www.winbond.co.tw

**Circle No. 388****Embedded-memory model libraries and compilers****Artisan Components Inc**

Sunnyvale, CA  
1-408-734-5600  
fax 1-408-734-5050  
www.artisan.com

**Circle No. 389****Aspec Technology**

Sunnyvale, CA  
1-408-774-2199  
fax 1-408-522-9450  
www.aspec.com

**Circle No. 390****Avant! Corp**

Fremont, CA  
1-510-413-8000  
fax 1-510-413-8080  
www.avanticorp.com

**Circle No. 391****Azalea Microelectronics**

Santa Clara, CA  
1-408-982-9141  
www.azaleamc.com

**Circle No. 392****Cascade Design Automation**

Bellevue, WA  
1-425-643-0200  
fax 1-425-649-7600  
www.cdac.com

**Circle No. 393****Denali Software Inc**

Palo Alto, CA  
1-650-325-7241  
fax 1-650-325-5724  
www.denalisoft.com

**Circle No. 394****Duet Technologies**

San Jose, CA  
1-408-432-9200  
www.duettech.com

**Circle No. 395****Mentor Graphics Corp**

Wilsonville, OR  
1-503-685-7000  
fax 1-503-685-1202  
www.mentorg.com

**Circle No. 396****Nara Technologies**

San Jose, CA  
1-408-954-1700

**Circle No. 397****Nurlogic Design**

San Diego, CA  
1-408-455-7570  
www.nurlogic.com

**Circle No. 398****Phoenix Technologies Virtual Chips Division**

San Jose, CA  
1-408-570-1000  
www.phoenix.com

**Circle No. 399****Puyallup Integrated Circuit Co**

Federal Way, WA  
1-253-927-6910  
www.picco.com

**Circle No. 400****Synopsys Inc**

Mountain View, CA  
1-650-962-5000  
fax 1-650-694-4249  
www.synopsys.com

**Circle No. 401****Virage Logic Corp**

Milpitas, CA  
1-408-263-7700  
fax 1-408-263-9523  
www.virlog.com

**Circle No. 402****Virtual-Silicon Technology**

Sunnyvale, CA  
1-408-747-1850  
fax 1-408-747-1950  
www.virtual-silicon.com

**Circle No. 403****VOTE . . .**

Please also use the Information Retrieval Service card to rate this article (circle one):

**High Interest 586**  
**Medium Interest 587**  
**Low Interest 588**

**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. 404**


## EMBEDDED MEMORY

Tower process has a cell smaller than 0.7  $\mu\text{m}$  on the initial 5V, 0.5- $\mu\text{m}$  process and adds only four mask steps to a base CMOS manufacturing flow. Liron predicts that access times will be less than 40 nsec, ideal for DSP applications. Although programming and erasing performance will lag behind that of other flash-memory technologies, 512-byte blocking will make Tower's approach applicable to some data-storage designs.

Mitsubishi has ported its 32-bit M32R CPU core not only to embedded DRAM with the M32R/D, but also to flash memory with the M32R/E family. Now available for sampling, the M32160F4UFP and M32160F3UFP include 512 and 384 kbytes, respectively, of embedded DINOR, as well as 16 kbytes of SRAM, multiple timers, 10-bit A/D converters, DMA, serial-I/O channels, and an integrated real-time debugger. Mitsubishi builds the devices on a 0.5- $\mu\text{m}$  process, in which the company specifies 25-MHz performance across the  $-5$  to  $+110^\circ\text{C}$  and 3.5 to 5V ranges.

Silicon Storage Technology (SST, Sunnyvale, CA) is another company with great expectations for embedded flash memory. Although currently mired in legal battles with Intel, SST CEO Bing Yeh has set an aggressive goal for the company to be the preferred embedded-flash-memory provider for the semiconductor industry. SST has licensing agreements with Analog Devices (Norwood, MA) and Information Storage Devices (San Jose, CA), which also has a license from flash-memory R&D vendor Bright Microelectronics (Sunnyvale, CA). SST's licensees also include Samsung, Sanyo (Osaka, Japan), Seiko Epson (Nagano, Japan), and Winbond Electronics. TSMC, yet another SST licensee planning to implement embedded flash memory beginning at a 0.5- $\mu\text{m}$  process, targets less than 15% incremental wafer cost over a base CMOS-logic process. SST asserts that its split-gate approach lets ASIC suppliers insert high-voltage and memory modules without modifying logic processes—before logic fabrication and with only four to five incremental masking steps.

Other vendors offering embedded flash memory built on an EEPROM



**The Alternate Metal Virtual Ground architecture enables dense embedded-EPROM arrays (courtesy WSI).**

foundation include Atmel, Hyundai, Lucent Technologies, Motorola, and Texas Instruments. EEPROM variants, although they have more complex cell structures than NOR alternatives, offer efficient programming and erasing that minimize the required size of on-chip charge pumps, scale down well to low-voltage operation, and, unlike NAND flash memory, are appropriate for both code and data. Atmel's embedded-flash-memory process is identical to that the company uses in its flash-memory-based 8051 and AVR microcontrollers. According to Atmel Marketing Manager Dennis Kish, Atmel can cost-effectively integrate as much as 1.2 Mbits of embedded flash memory on the current 0.35- $\mu\text{m}$  process, with an 8-Mbit target on the 0.25- $\mu\text{m}$  process, which will see first silicon by the end of this year. Lucent Technologies has high hopes for a new 0.25- $\mu\text{m}$  flash-memory approach that came out of Bell Labs research. This approach combines small cells, low-voltage operation, and low-current programming and erasing. Integrated Silicon Solution (Santa Clara, CA), which recently acquired flash-memory vendor Nexcom Technology (Sunnyvale, CA), plans both to embed the technology within its microcontrollers and other stand-alone products and to license the core to ASIC suppliers. □

## References

1. Dipert, Brian, "Advanced DRAM puts you in the fast lane," *EDN*, Oct 9, 1997, pg 52.
2. Dipert, Brian, "FRAM: ready to

ditch niche?," *EDN*, April 10, 1997, pg 93.

3. "EPROM virtual ground array," US Patent #5,204,835, granted April 20, 1993, IBM patent server <http://patent.womplex.ibm.com/cgi-bin/viewpat.cmd/5204835>.

4. Dipert, Brian, "Data storage in a flash," *EDN*, July 3, 1997, pg 65.

5. "Scalable processors in the billion-transistor era: iRAM," *IEEE Computer Society*, September 1997, pg 75.

6. Dipert, Brian, "EEPROM: survival of the fittest," *EDN*, Jan 15, 1998, pg 77.

7. Lalchandani, Amrit, and Frank Krupecki, "DRAMASIC: the marriage of memory and logic," *EDN Products Edition*, May 14, 1997, pg 23.

8. Lipman, Jim, "Not just your basic ASIC libraries," *EDN*, June 5, 1997, pg 52.

9. Przybylski, Steven, "Embedded DRAMs: what, why and when?" pg 319, DesignCon '98 Conference, High-Performance System Track.

## Acknowledgments

*Special thanks go to Brian Barrera from Mentor Graphics, John Garner and Julianne Whitelaw representing Mitsubishi, Eran Liron from Tower Semiconductor, and Betty Prince from Memory Strategies International (Sugarland, TX) for their assistance during the research phase of this article. Steven Przybylski from the Verdande Group (San Jose, CA) and David Patterson from the University of California—Berkeley also provided valuable feedback.*



You can reach Technical Editor Brian Dipert at 1-916-454-5242, fax 1-916-454-5101, [edndipert@worldnet.att.net](mailto:edndipert@worldnet.att.net), URL <http://members.aol.com/bdipert>.

## VOTE

Please use the Information Retrieval Service card to rate this article (circle one):

High  
Interest  
586

Medium  
Interest  
587

Low  
Interest  
588