

Embedded DRAM: Technology platform for the Blue Gene/L chip

S. S. Iyer
J. E. Barth, Jr.
P. C. Parries
J. P. Norum
J. P. Rice
L. R. Logan
D. Hoyniak

*The Blue Gene®/L chip is a technological tour de force that embodies the **system-on-a-chip** concept in its entirety. This paper outlines the salient features of this 130-nm complementary metal oxide semiconductor (CMOS) technology, including the IBM unique embedded dynamic random access memory (DRAM) technology. Crucial to the execution of Blue Gene/L is the simultaneous instantiation of multiple PowerPC® cores, high-performance static random access memory (SRAM), DRAM, and several other logic design blocks on a single-platform technology. The IBM embedded DRAM platform allows this seamless integration without compromising performance, reliability, or yield. We discuss the process architecture, the key parameters of the logic components used in the processor cores and other logic design blocks, the SRAM features used in the L2 cache, and the embedded DRAM that forms the L3 cache. We also discuss the evolution of embedded DRAM technology into a higher-performance space in the 90-nm and 65-nm nodes and the potential for dynamic memory to improve overall memory subsystem performance.*

Introduction

System-on-a-chip (SoC) technology plays a key role in the implementation of the Blue Gene*/L supercomputer. Central to the machine is the Blue Gene/L chip shown in **Figure 1**. It consists of twin IBM PowerPC* 440 processors (PU0, PU1) and their local static random access memory (SRAM) caches, two floating-point units (FPU0, FPU1), five network controllers, a memory control system, and 4 MB of on-chip dynamic memory (L3). The integration of the dynamic memory on-chip requires the development of embedded dynamic random access memory (DRAM) technology [1]. With embedded DRAM, the advantages of dense, reasonably fast but very-high-bandwidth DRAM can be integrated with the logic used to build the rest of the processor chip. Since memory continues to be the single largest component of die area in most high-performance processor chips,

developing denser and higher-performing memory capable of being integrated with high-performance logic is perhaps the crux of both the SoC and the memory subsystem optimization problems.

Increasingly, as exemplified by Blue Gene/L, more of this memory is becoming dynamic. (Both dynamic and static memory are volatile; the term *dynamic* refers to the fact that a dynamic memory cell can lose the memory state, even with the power on, and must be refreshed periodically.) There are four important reasons for integrating dynamic memory on the processor chip:

- *Cell size*: Dynamic cells tend to be significantly smaller (5× to 7×) than static cells in a given technology. Traditionally, static memory cells are composed of six transistors, including a cross-coupled pair that toggles between the two states of the

©Copyright 2005 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

0018-8646/05/\$5.00 © 2005 IBM

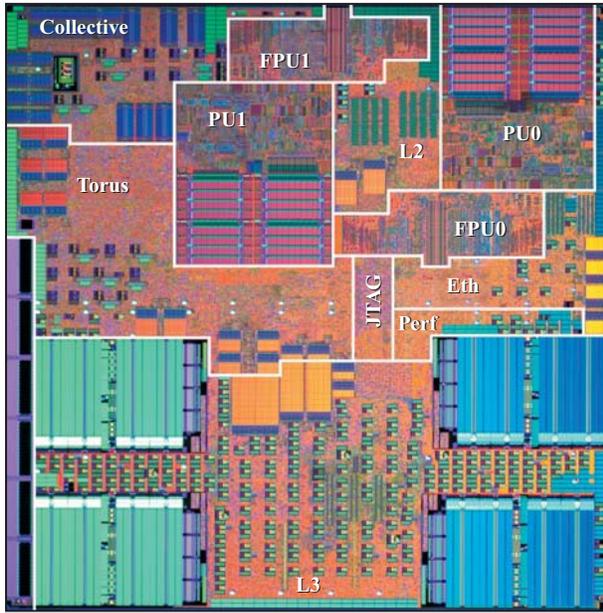


Figure 1

The Blue Gene/L compute chip, showing the twin PowerPC 440 processors (PU0 and PU1) and their associated floating-point units (FPU0 and FPU1), network interface controllers (collective, torus, JTAG, Eth), and the memory hierarchy (L1 caches internal to the processor units, L2, and L3). The four large arrays in the L3 region are the 4 MB of embedded DRAM. The small rectangles distributed throughout the L3 region and elsewhere are I/O cells and decoupling capacitors. ©2005 IEEE. Reprinted with permission from A. A. Bright et al., "Creating the Blue Gene/L Supercomputer from Low Power System-on-a-Chip ASICs," *Digest of Technical Papers*, 2005 IEEE International Solid-State Circuits Conference.

memory cell [Figure 2(a)]. Dynamic cells consist of a single transistor and a storage capacitor [Figure 2(b)]. The presence or absence of charge on the storage node determines the state of the memory cell. Figure 3 shows the scaling trends for both static and dynamic memories when they are embedded in bulk logic technology. Memory cell size, however, does not translate directly to die size savings. The charge stored on the node of the dynamic cell is susceptible to leakage and must be refreshed periodically. Additionally, dynamic cells lose their state during the read operation and have to be rewritten after reading. Furthermore, dynamic memory requires the access transistor gate voltages to be boosted to ensure maximum stored charge. Thus, there is additional circuit overhead that causes the area advantage of dynamic memory over static memory to be about 3 to 4, depending on the size of the memory being considered. Larger memory blocks are more area-efficient.

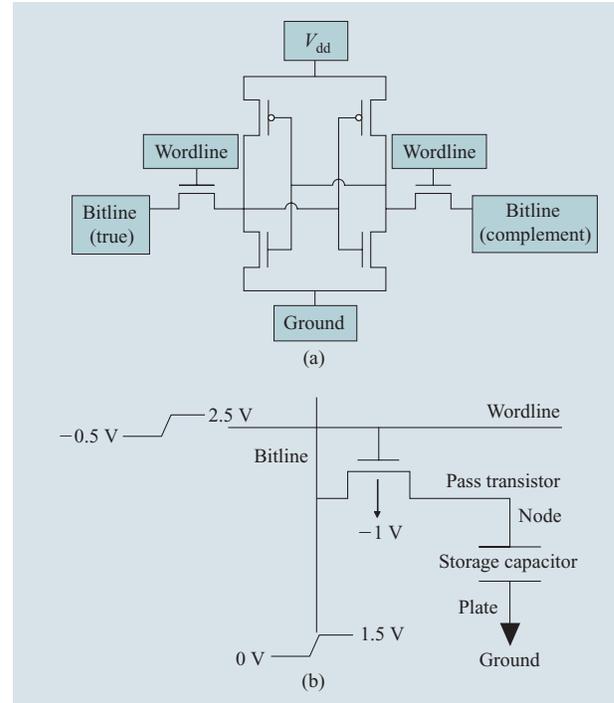


Figure 2

(a) Six-transistor static memory cell. (b) Dynamic memory cell.

- Standby power:** As semiconductor scaling proceeds beyond the 130-nm generation, the device off currents show an alarming increase; as shown in Figure 4 [2], the standby power density begins to approach the active power density of the chip. In the case of memory, a static memory cell with six transistors tends to have ~1000 times more leakage current on a per-cell basis than a dynamic cell. At the memory macro level, one has to consider the standby current of the peripheral devices and the various internally generated power-supply circuits, and the refresh current needed for the dynamic memory. When all of these are included, embedded DRAM tends to have a 6× to 8× advantage over embedded SRAM. Active power tends to be comparable and is dictated by the performance and memory bandwidth used, since the output data lines tend to be heavy loads.
- Soft-error rate (SER):** SER refers to the transient single-cell upsets caused by the penetration of high-energy particles, such as cosmic rays, into the silicon. The carriers generated by these particles can cause individual bits to be upset. The susceptibility of the cell to these events is exponentially dependent on the stored charge. Since embedded DRAMs store approximately 20× more charge, they tend to have

SER-induced failure rates that are $1000\times$ to $5000\times$ lower than those seen on SRAM cells. While SER can be addressed through error detection and correction codes, such codes result in added complexity and latency and only partially mitigate the problem.

- Performance:* It is generally assumed that embedded DRAMs provide lower performance than embedded SRAMs. This must be qualified, as there are density and performance tradeoffs. There are multiple metrics for performance. The most important is latency, which refers to the time interval between presenting the memory macro with a random address and receiving the data. Another important metric, cycle time, is the time interval before the next address can be presented. A significant component of latency arises from time of flight across the macro. For large memory sizes (several megabytes), this component of latency can dominate. In fact, for a 9-MB embedded SRAM in 90-nm technology, substituting embedded DRAM reduces the memory size by about 70% and the associated wiring component of latency to the farthest block by more than 30%. The case for embedded DRAMs for large embedded memories is clear. What is not well known, however, is that as technology scales, advances in embedded DRAM architecture and design techniques allow for better intrinsic performance, even for 1-Mb-type memory blocks. This results directly from the fact that the signal levels at the bitline for both types of memory are becoming comparable, so the smaller cell size (hence, time of flight) becomes a defining parameter for all but the smallest memory blocks. Cycle time tends to be longer for embedded DRAMs because the read is destructive and the time to write a full charge on the storage node can be long, but the use of multibank organization, as described later, offers a significantly reduced bank cycle time. Another important point to note is that aggressively scaled SRAMs exhibit new instability modes, but discussion of this point is beyond the scope of this paper [3].

From a processor design perspective, optimizing memory subsystem performance remains the biggest challenge. Processor frequency, until now an accepted performance metric, is no longer adequate; furthermore, it does not show the dramatic increases witnessed in the last several generations. Two contributors to this shift are the saturation of device-level performance brought about by the difficulties in scaling through technology alone and the dramatic increase in chip power dissipation to almost unsustainable levels. A finer granularity in the on-chip memory hierarchy, including the use of a third level of

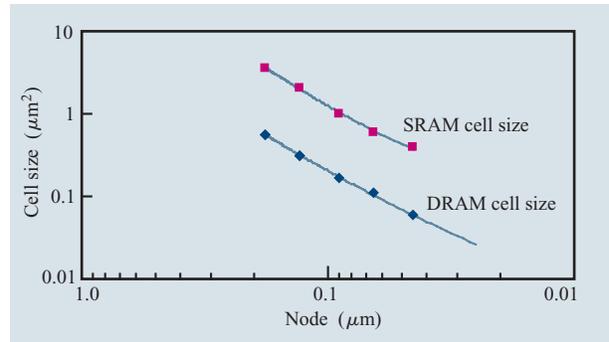


Figure 3

Cell size trends for static and dynamic memory cells when embedded in logic technology. Node generation is as shown. Cell size scaling is expected to saturate for a variety of reasons outlined in the text.

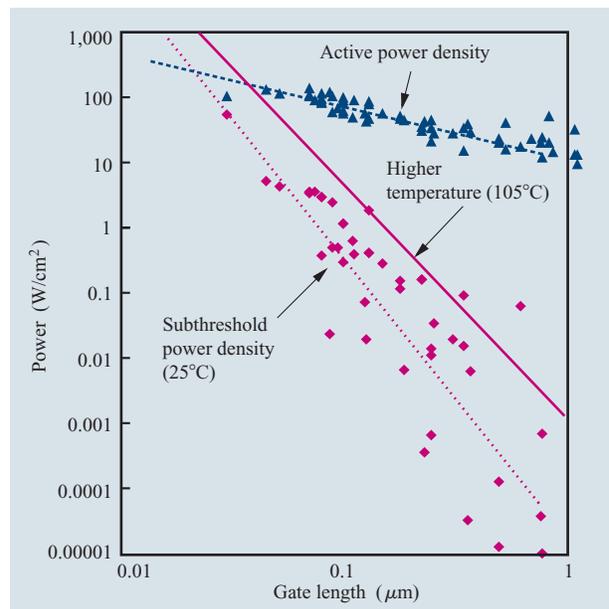


Figure 4

Active and standby power as a function of gate length. Not shown is the gate leakage component, which becomes significant beyond the 90-nm node. Both embedded SRAMs and embedded DRAMs show high standby currents in the peripheral circuits. However, DRAM subarrays will have a significant (more than $6\times$ to $8\times$) reduction in standby current compared with SRAMs. Active powers for both are dominated by data rates and are comparable. Adapted from [2] with permission.

on-chip memory, is increasingly common. To be useful, however, this third level must be large and dense, with low enough latency and high bandwidth. The bandwidth criteria should be understood in terms of what can be achieved on-chip using on-chip wiring pitches, but cannot

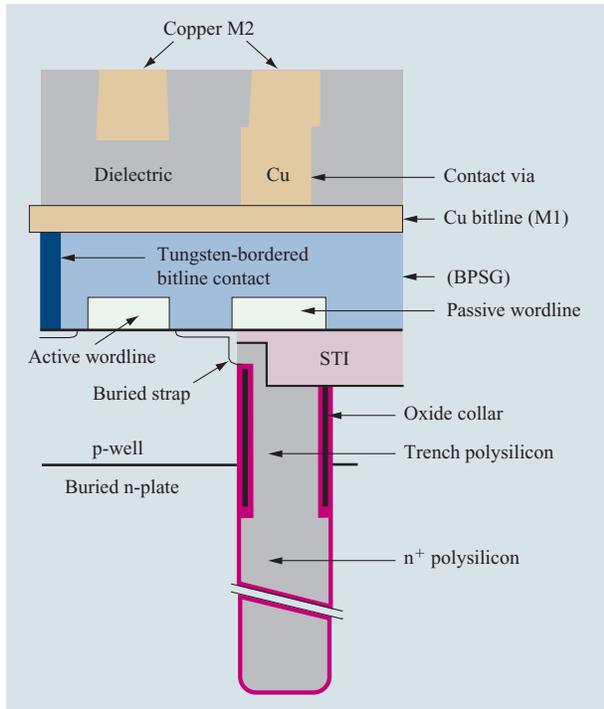


Figure 5

Schematic cross section of an embedded DRAM cell.

be met by off-chip memory, which is limited by input/output (I/O) wiring density and off-chip drivability and power considerations. Embedded DRAM meets these criteria for the on-chip L3 and promises to meet these criteria for increasingly demanding L2 applications in future generations.

Migrating from the commodity DRAM base to the logic base

Historically, embedded DRAM developed from a DRAM base. Initial work at IBM and elsewhere focused on leveraging DRAM technology primarily to fabricate custom DRAMs for use in off-chip cache and graphic DRAM applications [4, 5]. The main rationale for this approach was the belief that the complexity of DRAM process architecture should not be modified, and any logic functions should be contained within the available menu of devices offered by the DRAM technology. The yield and line learning of this part would have to depend on the presence of a commodity part with the same subarray running in volume in the semiconductor line. At the 0.5- μm to 0.25- μm nodes, this meant relatively poor performance for the logic devices and, consequently, limited logic function. In 1998, IBM introduced a DRAM-based embedded DRAM [6] at the 0.25- μm node called 7LD, with additional performance-oriented

changes in the devices. These included a shorter channel logic device for higher performance and the introduction of logic, such as multilevel back-end technology, with up to six aluminum wiring levels and a variety of logic-like packaging options. In addition, a limited custom library was also offered in this technology.

It is worth summarizing some important conclusions from this exercise:

- The embedded DRAM macro performance was modest. In fact, it was comparable to commodity memory performance. This resulted from the fact that the subarray was designed to meet the standardized performance metrics required of the commodity part.
- The relatively poor performance of the peripheral logic devices resulted in a somewhat marginal library. This limited the widespread use of the technology.
- Density suffered on two counts: DRAM processes were optimized to print array features at very tight pitches, but random patterns were significantly looser than logic technology at the time; second, the low drive of the devices required them to be sized large enough to drive the on-chip load. As a result, the logic density lagged by more than a generation.
- The assumption of DRAM line learning did not apply as well as initially envisioned. The on-chip voltage variations, junction temperatures, and retention characteristics were different enough from standard DRAM that additional process sensitivity learning was required. Furthermore, a strategic redirection of focus to logic rather than commodity DRAM meant that there was no commodity DRAM learning effect on the embedded DRAM parts in any case.

On the basis of these findings, subsequent embedded DRAM technologies were developed to be integrable with high-performance logic [7]. Furthermore, the embedded DRAM design intellectual property was integrated into the IBM Blue Logic* libraries and followed the same integration and test methodologies for ease of use. Both of these factors have contributed immensely to the success of embedded DRAM at IBM and have facilitated realization of the Blue Gene/L SoC, including the integration of the PowerPC 440 cores. Nevertheless, it should be emphasized that the commodity DRAM pedigree of IBM embedded DRAM has been a crucial element of its success. By incorporating the learning of several generations of commodity DRAMs and retaining crucial elements of the process architecture, our embedded DRAM macros have avoided the pitfalls of the well-known complex DRAM pattern

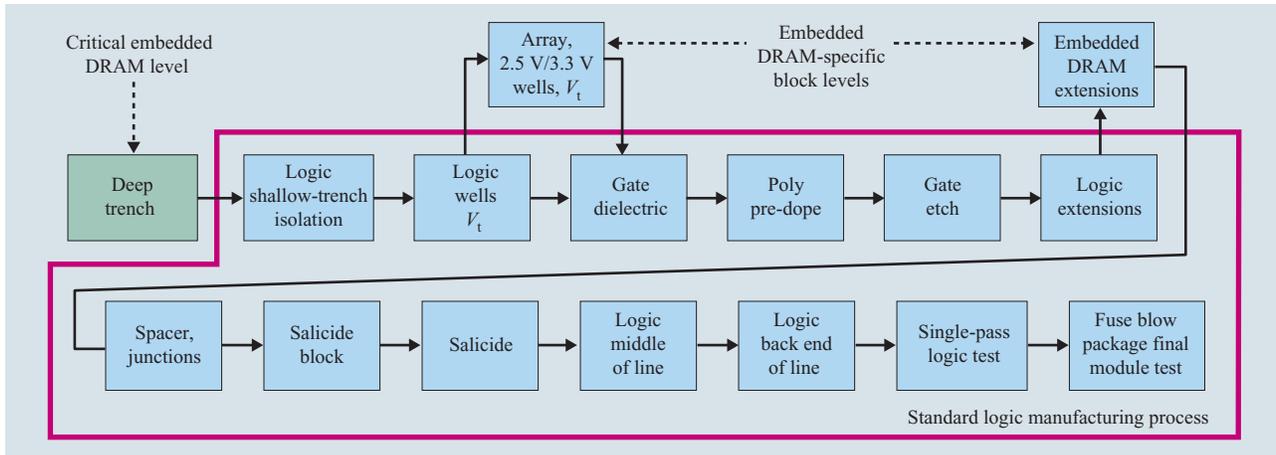


Figure 6

Process flow used to manufacture the Blue Gene chip.

sensitivities and retained a robust character over a wide range of operating conditions.

Embedded DRAM technology

Blue Gene/L currently uses 130-nm technology [8]. This is second-generation IBM embedded DRAM technology, the first being at 180 nm. This has also been extended to the 90-nm and 65-nm nodes. At 130 nm, the embedded DRAM platform allows for the integration of high-performance logic devices, embedded SRAMs with two different cell sizes, content addressable memories (CAMs), and a variety of passive devices, such as precision resistors, inductors, and capacitors, if needed.

The embedded DRAM technology, shown in **Figure 5**, is based on IBM deep-trench (DT) technology [9]. The details of the trench fabrication have been described elsewhere [10]. The key feature to be emphasized here is the fact that the trench process is the very first processing that the wafer undergoes. After DT processing, the wafer presents a planar surface and, for all practical purposes, the DT processing can be considered transparent to the rest of the process. This should be contrasted with other processes, such as the stacked capacitor or metal-insulator-metal capacitor (MIM cap) approach, followed by others [11, 12] in which the capacitor is created in the middle-of-the-line process (between the device and wiring levels). In the latter case, novel materials must be used, and complex three-dimensional capacitor structures must be built. These require additional planarizing steps and, with advanced back-end wiring, even a minimal degradation of planarity can result in degraded yields. Very often, the presence of these tall capacitor structures in the middle-of-the-line process requires unique first-level-metal processing that may also require changes in

the electrical timing of logic circuits for designs with and without embedded DRAM. It must be pointed out that the embedded DRAM cell layout follows all of the layout rules required by the technology, and no special waivers are granted. This methodology allows the embedded DRAM processing to have minimal impact on yield. This should be contrasted with the approach followed in the dense SRAM cells, in which the SRAM array is shrunk beyond the layout rules.

Another approach to embedded DRAM is the so-called “logic-only” process, in which the gate of the device is used as the storage node. Typically, the capacitance obtainable by this approach is of the order of 2–3 fF and is usually woefully inadequate for robust DRAM operation as required in Blue Gene/L types of applications.

Figure 6 shows the process flow used in the manufacture of the Blue Gene/L chip. Note that after the DT processing, the rest of the process allows for the seamless integration of the logic and SRAM devices. The pass transistor used in the array to access the deep trench is threshold-tailored for low leakage. A triple-well process is used for the embedded DRAM subarray. Typical leakage currents for the array are in the fA regime. In 130-nm technology, cobalt disilicide is formed on the gate, source, and drain areas of only the logic field-effect transistors (FETs) to reduce their resistance. The array areas are blocked from forming silicide. This is done to ensure long retention. The total leakage associated with the cell is a few fA. The median retention time is of the order of a few seconds, and cell leakage is dominated by junction leakage in this technology. This being the case, retention is sensitive to temperature, and the median retention time shows an increase of approximately 2× for

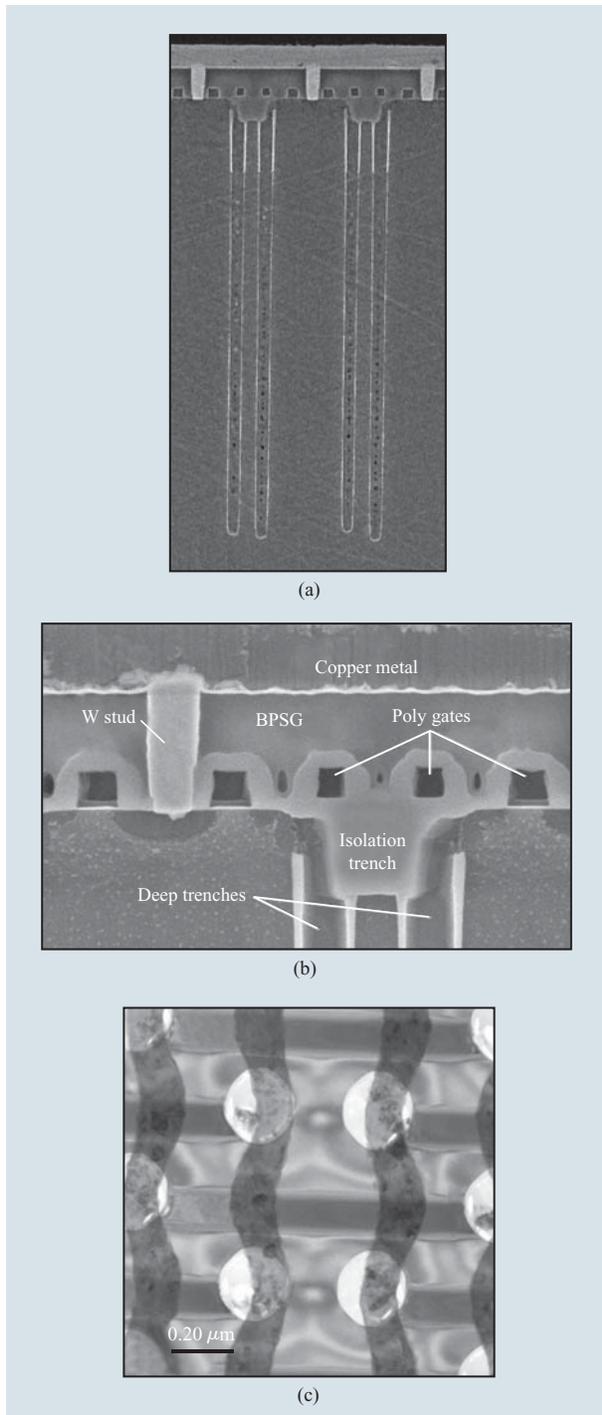


Figure 7

(a) and (b) SEM cross sections of cell used in the Blue Gene chip. (c) Top view of the $0.31\text{-}\mu\text{m}^2$ cell.

every 10°C drop in temperature. The refresh criteria for the macro are designed for a retention time of 3.2 ms at

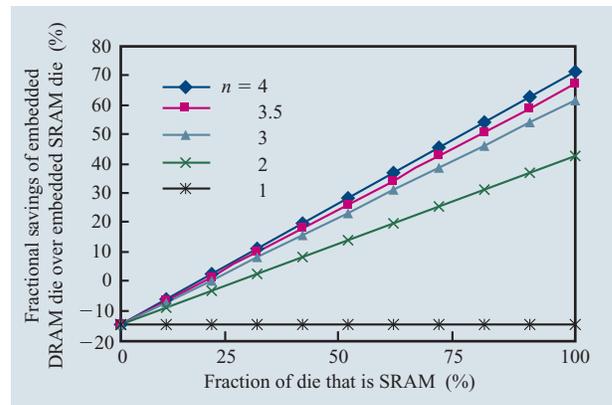


Figure 8

Fractional cost savings obtained by replacing embedded SRAM with embedded DRAM, as a function of the fraction of the die that is replaceable embedded SRAM, for different values of embedded DRAM efficiency n , which represents the density improvement of embedded DRAM over embedded SRAM. The process complexity adder is assumed to be 15%, which is representative of current technology. Basically, if 25% of the die is replaceable SRAM, the embedded DRAM solution is cost-effective.

application conditions ($T_j = 105^\circ\text{C}$; worst-case voltage and worst-case process). The biasing conditions for the cell are shown in Figure 2(b). All required voltages are generated on-chip using the supplied V_{dd} of 1.5 V. The redundancy scheme using laser fuses is discussed in a later section. **Figures 7(a) and 7(b)** show scanning electron microscope (SEM) cross sections of the cell, and **Figure 7(c)** a top view. The deep trench, the buried strap that connects the node to the pass transistor, the pass transistor, and the bitline contact are clearly seen.

Cost and process complexity

A major misconception in the industry is that embedded DRAM is more complex and, therefore, more expensive. While the additional processing associated with the fabrication of the capacitor and DRAM access device does mean that an embedded DRAM wafer is more expensive to process, this does not mean that an embedded-DRAM-based die is more expensive than an embedded-SRAM-based die. The discussion must be focused on cost-effectiveness. Consider the following argument:

1. Start with an SRAM-based die.
2. L = logic area; S = SRAM memory area.
3. SRAM-based die size = $L + S$ = cost.
4. Replace SRAM with DRAM $D = \alpha S$, where α is the ratio of the areas of the DRAM and SRAM on a per-Mb basis. (In **Figure 8**, $n = 1/\alpha$.)

5. $L =$ logic area; $D =$ DRAM area.
6. DRAM-based die size $= L + D = L + \alpha S$.
7. DRAM-based cost $= \beta(L + D) = \beta(L + \alpha S)$, where β is the fractional increase in processing cost per wafer for the embedded DRAM wafer compared with one with no embedded DRAM.
8. Cost savings $= (L + S) - \beta(L + \alpha S)$
 $= L(1 - \beta) + S(1 - \alpha\beta)$.
9. Fractional savings $= [L(1 - \beta) + S(1 - \alpha\beta)] / (L + S)$.

This has been plotted in Figure 8 for different values of $n = 1/\alpha$. With respect to Blue Gene/L, for 130-nm technology, a die with approximately 25% of contiguous SRAM is an excellent candidate for replacement with embedded DRAM. If the Blue Gene/L chip was designed with embedded SRAM, approximately two thirds of the die would be composed of replaceable SRAM, so that the embedded DRAM solution is, in fact, 40% more cost-effective—a significant cost savings. Interestingly, this cost savings has been achieved with no significant increase in latency. This is because the smaller footprint of the DRAM cache compared with the hypothetical SRAM cache results in a shorter time-of-flight delay from the farthest bits in the macro.

Unfortunately, misguided ideas on the perceived cost adder of embedded DRAM without a clear understanding of its value proposition has led to the development of so-called 1T cells of dubious value and limited extendability. (Such cells are offered by several foundries without DRAM experience.) These cells lack the robustness, including noise immunity, of the IBM trench technology and do not offer the same die area savings. Additionally, as logic technologies build more complexity into the base process—such as additional device types, additional metal levels, precision resistors and capacitors, and silicon-on-insulator substrates—the cost and complexity adder for embedded DRAM continues to fall and is expected to be well below 10% at the 65-nm high-performance node.

Embedded DRAM macro design considerations

Embedding DRAM into an IBM Cu-11¹ application-specific integrated circuit (ASIC) design extends the on-chip memory capacity to more than 40 MB, allowing memory that has historically been off-chip to be integrated on-chip. In fact, chips with as much as 344 Mb of DRAM are currently manufactured in this technology.

¹IBM bulk complementary metal oxide semiconductor (CMOS) technologies are associated with an IBM standard cell library. Thus, the 0.18- μ m node is called CMOS 7SF, and the associated library is called SA27E; the 0.13- μ m node is called CMOS 8SF and the library is called Cu-11; the 90-nm node is called CMOS 9SF and the library is called Cu-08. More details are available at <http://www.chips.ibm.com>. The standard cell libraries employ only a subset of the technology options available in general at a given node.

With the memory on-chip, applications can leverage the high bandwidth naturally offered by a wide I/O DRAM and achieve data rates previously limited by pin count and off-chip data rates. Applications for this memory include standalone L3 cache chips for IBM pSeries* and xSeries* eServers*, network processors, and digital signal processors. The integration of embedded DRAM into ASIC designs has intensified the focus on how best to architect, design, and test a high-performance, high-density macro as complex as dynamic RAM in an ASIC logic environment.

The ASIC environment itself presents many difficult challenges that have historically affected DRAMs—specifically, wide voltage and temperature operating ranges and uncertainties in surrounding noise conditions. These challenges dictate a robust architecture that is noise-tolerant and can operate at high voltage for performance and low voltage for reduced power.

With the advent of embedded DRAM offerings in a logic-based ASIC technology, the performance of embedded DRAM macros has improved significantly over that in DRAM-based technologies. Subsequently, users are increasingly replacing SRAM implementations with embedded DRAM, placing additional pressure on macro performance and random cycle time. This pressure extends into test, where the use of traditional direct memory access (DMA) is costly in silicon area and wiring complexity, and it introduces uncertainty in performance-critical tests.

A more attractive solution to this test problem is the use of a built-in self-test (BIST) system that is adapted to provide all of the necessary elements required for high fault coverage on DRAM, including the calculation of a two-dimensional redundancy solution, pattern programming flexibility, at-speed testing, and test mode application for margin testing [13, 14]. The Cu-11 embedded DRAM macro has been developed around the idea of user simplicity, while including a high degree of flexibility, function, and performance [15]. For application flexibility, the embedded DRAM can grow in 1-Mb increments to provide macro sizes from a 1-Mb minimum to a 16-Mb maximum; it offers a 256-I/O width and a 292-I/O width for applications requiring parity. The wide I/O was chosen to provide maximum bandwidth; for applications not requiring the full width, bit-write control was included to facilitate masking.

Multiple embedded DRAM macros can be instantiated on an ASIC die, enabling customers to make the performance and die area tradeoff specific to their application. **Figure 9** shows a high-level floorplan of the embedded DRAM. The embedded DRAM is constructed from building blocks: a 1-Mb array core, a power system for generating boosted voltage levels used by the array core, a control system for buffering and generating

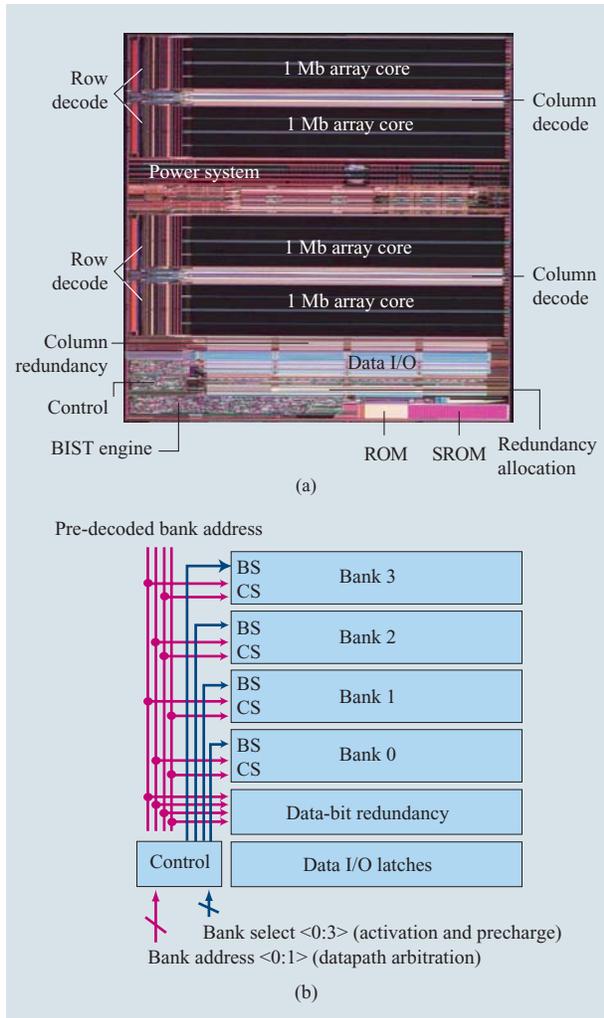


Figure 9

(a) Floorplan of the 0.5-MB macro. It shows the macro utilities at the bottom—including the control logic, BIST engine, redundancy analysis logic, etc.—that can be shared over the entire macro (up to 16 Mb). The 4-Mb subarray is built with 1-Mb array cores and is mirrored to implement a 1-MB macro. Four such macros are integrated in the Blue Gene chip (see Figure 1). (b) Floorplan multibank architecture.

the array core timing signals, column redundancy for replacing defective data bits, data I/O for receiving and transmitting off-macro data, and BIST for testing the embedded DRAM macro. BIST is composed of a microprocessor-based engine, instruction memory [read-only memory (ROM) and scannable read-only memory (SROM)], a data comparator, and a redundancy allocation unit. The 1-Mb array and its support circuitry are replicated to construct the desired macro size. Each embedded DRAM macro contains a single control system, a common power system, and a BIST.

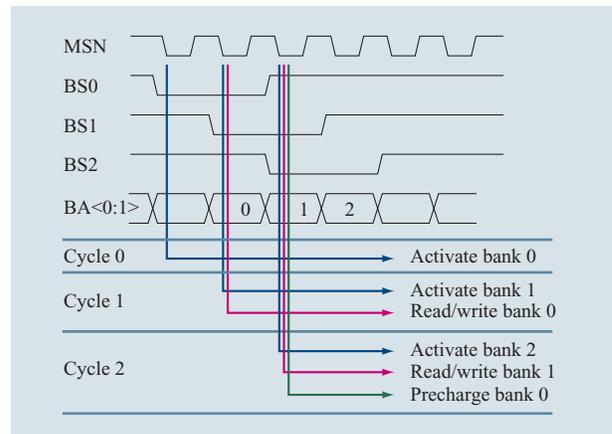


Figure 10

Timing diagrams showing multibank operation over three cycles.

This architecture lends itself well to providing two modes of macro operation: single-bank and multibank interleave modes. The single-bank operation provides a simple SRAM replacement function, while the multibanking mode extends the macro performance by allowing concurrent operations to independent banks. Bank operation was intended to resemble an embedded SRAM, supporting simple broadside addressing with read and write control. To improve bandwidth, the user can optionally use page mode, which was carried over from conventional DRAM. This was the only mode supported in SA27E; in Cu-11, however, it was decided to support multibank operation.

For the multibank-mode configuration, each 1-Mb block of the macro acts as an independent bank that shares a common address and data bus with all other 1-Mb blocks within the macro. The number of banks within a macro is determined by the macro size. Figure 9(a) shows the floorplan, and Figure 9(b) shows a 4-Mb macro with four banks. A bank select (BS) pin is associated with each bank (1-Mb block) and controls activation and precharge of that bank. The bank address (BA) is decoded by control logic and arbitrates which bank has control of the datapath.

In multibank configuration, the macro does not employ broadside addressing. Rather, the embedded DRAM macro operates similarly to a synchronous DRAM (SDRAM) whose addressing is performed in a manner similar to a row-address-strobe/column-address-strobe (RAS/CAS). The macro select input (MSN) is treated like a master input clock, latching the state of all other input pins with each falling MSN edge. Figure 10 shows three cycles. Cycle 0 activates bank 0, cycle 1 activates bank 1 and reads or writes bank 0, and finally cycle 2 activates bank 2, reads or write bank 1, and precharges bank 0. The

MSN input can be cycled at a maximum rate of 250 MHz (4 ns, assuming a nominal 50/50 clock duty cycle). This protocol supports simultaneous activate, read and write, and precharge to three different banks. Maximizing the number of banks in a macro improves the probability of avoiding an open (or busy) bank and maintaining the pseudo-random peak bandwidth of 8 GB/s.

Test considerations

Since the bit cells used in the embedded DRAM macro are derived from their commodity DRAM predecessors, they will likely have the same type of sensitivities, which are well known from the development of commodity DRAM and require identification at test. Many of the interactions in the DRAM cell matrix are complex and only activated with certain combinations of defects and test patterns. To deliver the complex test patterns, commodity DRAMs use specialized test equipment with algorithmic pattern capability for generating the test sequences, and they employ large and fast data-capture memory with redundancy allocation hardware to identify and repair faults.

Considering how to test a DRAM embedded in logic creates a dilemma: logic tester or memory tester [16–18]? The logic test platform that has developed for past generations of ASICs without embedded DRAM can be characterized as a low-cost reduced-pin-count tester with no algorithmic pattern generation or redundancy allocation hardware; it is, therefore, unable to test DRAM without assistance. The logic test patterns implemented are automatically generated with software based on the customer’s netlist. The test strategy comes down to either a two-tester solution (memory tester and logic tester) or comprehensive BIST. The two-tester approach suffers from the following issues:

- There are multiple test gates with the associated increase in wafer handling.
- Cumbersome requirements are placed on the customer to multiplex the macro I/O to package pins.
- Part-number-specific test pattern development is required and is typically difficult to automate.

In contrast to the two-tester approach, BIST also allows rapid isolation of faults to either the logic or embedded DRAM in SoCs, allowing for fault resolution.

In the high-part-number ASIC environment, it is essential to implement a single-tester platform using BIST for memory test and automated test generation for logic test. The design goal of the BIST is to provide a test engine operable in the logic test environment of low-cost, low-pin-count testers that stimulates the control, datapaths, and array of the embedded DRAM and provides fault coverage equivalent to that traditionally supplied to

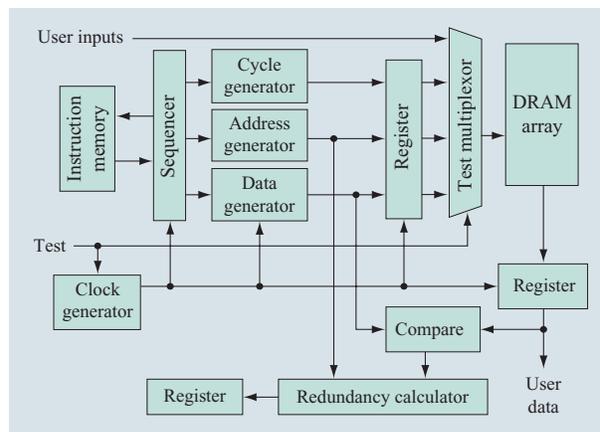


Figure 11

Built-in self test (BIST) engine components.

discrete DRAM by high-cost memory testers. The flexibility of the BIST system, made possible through the use of an SRAM, enables the test development engineers to alter the instruction memory to create new or modified test patterns or to change the sequence or number of patterns applied at each manufacturing test gate. The BIST components are shown in **Figure 11**. Ultimately, the BIST locates all faults in each 1-Mb array segment, calculates the two-dimensional redundancy solution required to repair these faults, and reports this solution via standard scan string methods [19]. The redundancy solution is permanently stored in a remote fuse memory (nonvolatile) programmed with a laser after testing.

A key enhancement to the BIST schemes previously used for embedded SRAM macros is the inclusion of a redundancy calculator, also referred to as *redundancy allocation logic (RAL)*, for two-dimensional redundancy. The function of the redundancy calculator is to compare data read from the array with the data expected from the BIST engine and optimally allocate redundancy for array fails. The BIST processor calculates row and data-bit redundancy for wide-I/O embedded DRAM macros. The system is described in [20]. Each 1-Mb array contains its own redundant elements, which may not be shared with other 1-Mb arrays. For this reason, BIST calculates and stores only 1 Mb worth of a redundancy solution at a time. Calculation of the full 16-Mb repair solution would require 16 times the number of fail counters and address registers, increasing the silicon area required for the BIST to an unacceptable level.

An additional enhancement to the redundancy calculator is the capability to reload the SRAM instruction memory discussed earlier. By segmenting the scan string and reloading only the SRAM, neither the

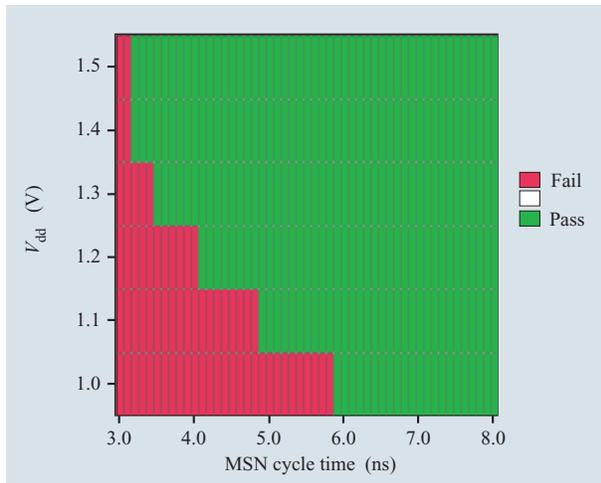


Figure 12

“Schmoo” plot of random cycle time as a function of chip V_{dd} . A 4-ns cycle time (250-MHz clock frequency) can be supported with manufacturing margin over the entire process window and operating conditions of temperature, including retention at 115°C.

state of the BIST engine nor the current state of the RAL calculated redundancy solution is upset. Thus, additional patterns can continue to be applied one right after the other with the redundancy solution being calculated on the cumulative failure set from all patterns.

To maximize system performance, 250-MHz timing rules are provided to describe the random access memory (RAM) core performance. Testing at speed (high-frequency or ac testing) becomes essential to ensure the high-performance timings, especially in the presence of the speed-sensitive fails commonly found in DRAMs. As the complexity of chip-level integration increases, so do the challenges for ac test. Effective ac testing requires that the stimuli be provided with sufficient speed and accuracy while meeting the constraints for cost-competitive manufacturing.

The ac test development is greatly simplified if BIST units are placed in such a way that the delays between the BIST and the RAM core boundary are minimized and predictable. An example of such a design is one in which each instance of a RAM core comes with its own BIST. While this approach may use more silicon area than others, the savings in design, test development, and test cost can be realized through the reuse of the same integrated RAM and BIST core across a wide range of applications. In some environments, the results that can be derived from a finite design and test development resource are enhanced by focusing on the development of a reusable core rather than the development of a specific

customer part number. The RAM and BIST integrated core design point enables effective ac test by reducing the delay between the RAM and BIST, which minimizes the timing uncertainties introduced by process variability and typical tester hardware. Efficient ac test is realized by leveraging the self-test concept. Performing data compare and storing fail locations in the BIST reduces the demands on tester resources to input only. Eliminating the need for expensive high-speed data-capture hardware in the test head greatly simplifies requirements from both bandwidth and calibration points of view. Containing critical timings within a BIST represents another large step toward reducing tester demands when clock multiplication capability is added.

In the ASIC environment, in which the BIST has all the capability and flexibility to test the DRAM with minimal tester interaction, there still exists the need to get results from the BIST back to the tester. Upon completion of a functional pattern, there are several types of data that must be acquired from the BIST: pass/fail data, repair data needed for fusing, and bit-fail mapping to identify defective bits. Fail mapping is critical for locating defective bits for failure analysis and follow-on yield learning. In order to generate statistically significant data, large volumes of data are required, and these must be generated in an automated fashion. Providing easy access to a synchronized set of address and compare states makes it possible to create bitmaps for any BIST pattern with a minimum amount of offline processing. In this scheme, BIST design-for-test efforts have had a significant positive impact on both test overhead costs and diagnostic capabilities.

Hardware measurements on 16-Mb macros verify 3.3-ns (300-MHz) operation and 6.6-ns random access and 3.3-ns bank access at 1.5 V for high-performance applications, while also providing reasonable performance at 1.0 V. These results are shown in **Figure 12**, a “schmoo” pass-fail plot for random cycle compared with chip V_{dd} . These performance values easily support the application specification of 250 MHz at 1.5 V with high yield. Bit-fail mapping using BIST has been implemented in the manufacturing flow, generating data for failure analysis and yield learning. The general-purpose nature of the macro has enabled use across multiple product applications, including the POWER5* L3 cache, and cumulatively provides yield learn approaching that of a high-volume single-standard product.

Device considerations

To ensure the best possible performance in bulk silicon technology at competitive cost, the Blue Gene/L chip was designed and fabricated in the IBM standard 0.13- μm bulk CMOS offering [21]. The basic device characteristics

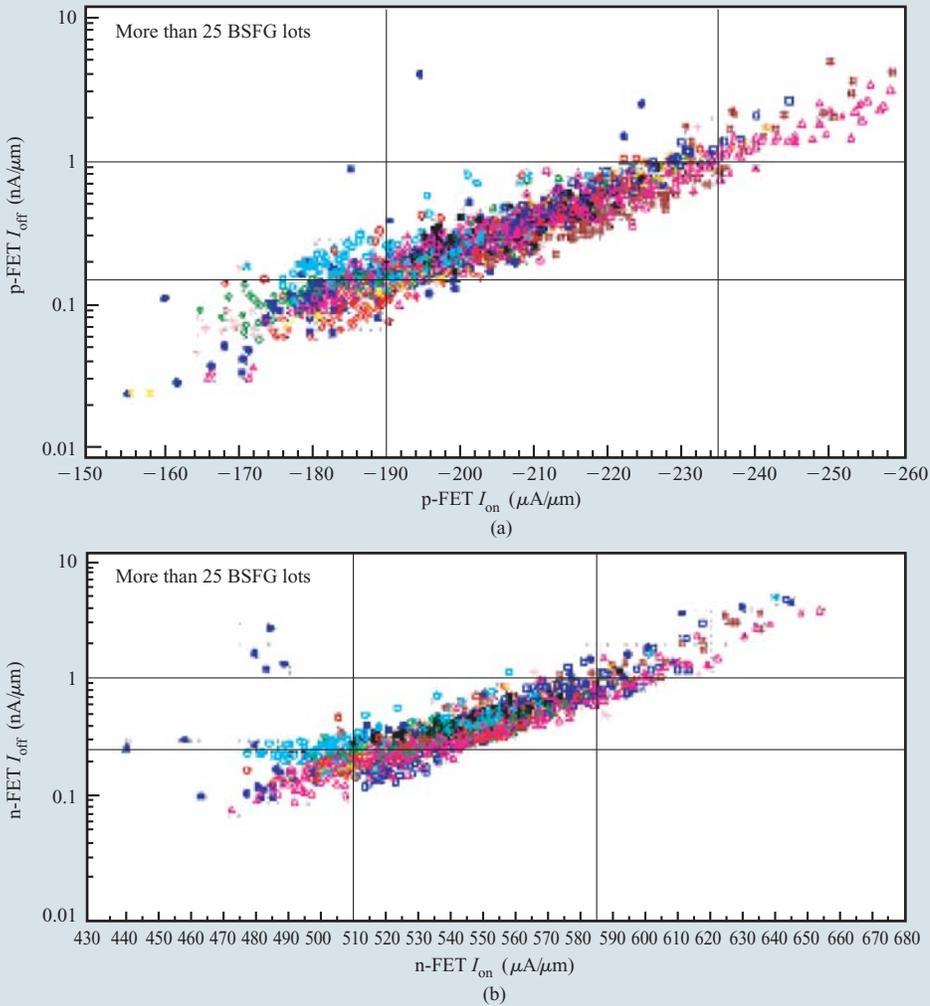


Figure 13

(a) Plot of p-FET I_{off} for different I_{on} taken from a variety of different products in Cu-11 technology; (b) plot of n-FET I_{off} for different I_{on} taken from a variety of different products. Control of I_{off} is crucial to ensure the power budget at a given ring oscillator performance target.

are as follows: the minimum polysilicon gate length on-wafer is 92 nm, with an oxynitride gate dielectric of 22 Å physical thickness, and n-MOS and p-MOS effective thickness in inversion of 31 Å and 32 Å, respectively. A variety of technology features are available that span applications over a range of leakage and performance requirements, and the device design is compatible with 1.2-V or 1.5-V (maximum 1.6-V) power-supply voltage. Blue Gene/L operates nominally at 1.5 V. At the 1.2-V benchmark, the room-temperature current drives for n-FET and p-FET are 590 $\mu\text{A}/\mu\text{m}$ and 235 $\mu\text{A}/\mu\text{m}$ when measured at 1 nA/ μm off current, as shown in **Figures 13(a)** and **13(b)**.

The technology specifications for off current are approximately 0.3 nA/ μm for both devices, with maximum off currents of 2 nA/ μm and 1 nA/ μm for n-FET and p-FET at a 6σ short gate length of 70 nm. Controlling the worst-case off currents requires excellent line width control, both across chip and across wafer. Finally, field-effect transistor (FET) parasitic capacitances are optimized as low as possible, taking into account the tradeoffs involving FET current drive, off current, and junction leakages. The culmination of the above is an 18.5-ps delay-per-stage performance on a fan-out-1 inverter ring oscillator as measured and modeled including all parasitics. **Figure 14** shows cross-sectional images of the fully processed devices.

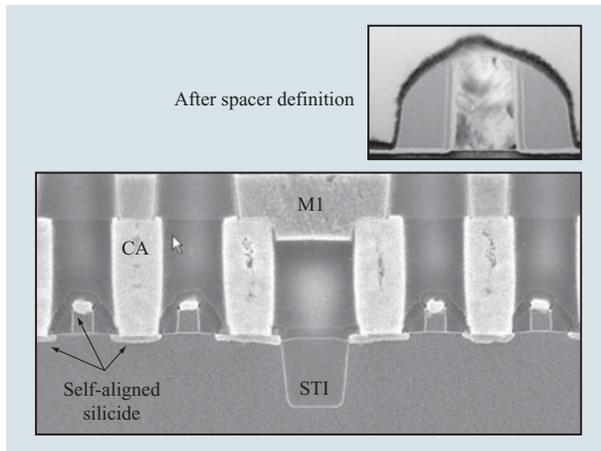


Figure 14

SEM cross sections of the logic devices showing the shallow-trench isolation (STI), completed device with CoSi_2 self-aligned silicide, contact metallurgy (CA), and metal 1 (M1). The inset shows the spacer definition. The spacer architecture determines the placement of the source and drain extensions, and—together with the polysilicon gate length, post-spacer thermal cycles, and stress generated from the isolation regions—determines the device characteristics.

In addition to the base FET devices described above, there is a dual-gate-oxide feature available for high-speed I/O devices. This feature is employed for the array FET and wordline system in the embedded DRAM macro. This is accomplished through a single-mask process rendering both the 22-Å and 52-Å oxides within the chip. The thick-oxide FET nominal design gate length is $0.24 \mu\text{m}$, with an on-wafer dimension of approximately $0.21 \mu\text{m}$. Independent threshold-voltage centering is also available on these devices, which are used with a 2.5-V (max 2.7-V) power supply. This same 52-Å oxide is employed in both the pass transistor in the array and the wordline driver circuits. The only additional device used in the embedded DRAM macro is the array device, which has its own optimized well and threshold tailoring. The on current for the array transistor is $40 \mu\text{A}$ per cell at a wordline boost of 2.5 V. A final note on the technology device feature set are optional high- and low- V_t devices derived from the base FET through an additional mask for implant tailoring. Such FETs often find applications

in arrays and high-speed logic critical paths. Typical offsets with respect to base FETs are approximately -90 mV and $+70 \text{ mV}$ for the low- and high- V_t devices, respectively.

For passive devices, besides the deep-trench decoupling capacitors, reference and electrostatic discharge (ESD) diodes, the technology has two types of optional-precision resistors. One is based on an unsilicided n-type diffusion, the other on an unsilicided p-type polysilicon line; they are specified at $73 \text{ } \Omega/\square$ and $340 \text{ } \Omega/\square$, respectively. Such resistors find a large range of critical applications, from high-speed analog applications to front-end ESD protection. The embedded DRAM uses this feature to block silicide from the array.

Sophisticated device models that account for topographical details and mechanical stress were developed to accurately model circuits over a wide variety of process conditions. Such models, including layout extraction capabilities, become important at the $0.13\text{-}\mu\text{m}$ node and beyond, and are key to optimizing design and closing timing on a complex SoC, such as Blue Gene/L.

As a whole, the technology delivers a versatile CMOS offering to a large variety of product applications. These range from components in high-end server chipsets to a realm of products under a common ASIC library, and finally to various custom foundry applications. Recent aggressive defect-reduction activities provide yield learning approaching that of a single high-volume standard product.

Embedded SRAM considerations

The Blue Gene/L chip was designed using 26 instances of single-port SRAM arrays. These SRAM arrays use a dense SRAM cell developed for 130-nm ASIC applications. This section discusses the design features of this SRAM cell, the methodology used to customize the design to allow the cell size to be reduced by 16% from the standard SRAM cell, and the electrical characteristics of the devices used in the cell. In Blue Gene/L, SRAM cells at a given technology node are somewhat relaxed initially. As the technology matures, however, the cell is shrunk by tightening the layout rules in the SRAM array through cell waivers. In the 130-nm node, a 22% shrinkage in the SRAM cell area was achieved using selective layout rule

Table 1 Comparison of 130-nm SRAM cell dimensions.

SRAM	Area (μm^2)	Wordline pitch (μm)	Bitline pitch (μm)	Bitline	Wordline
Standard SRAM	2.47	1.90	1.30	M2	PC with M3 stitch
Dense SRAM	2.04	1.70	1.20	M2	PC with M3 stitch

waivers. This shrinkage comes at the cost of a modest performance loss, as explained below.

Cell features

The dense SRAM cell used is a single-port six-transistor SRAM cell. The cell consists of two n-FET pass-gate devices, two n-FET pull-down devices, and two p-FET load (or pull-up) devices. The dense SRAM cell size is $2.04 \mu\text{m}^2$. This is a 16% shrinkage of the standard SRAM cell ($2.47 \mu\text{m}^2$) offered in the base technology. **Table 1** shows a comparison of several key dimensions for both the $2.47\text{-}\mu\text{m}^2$ and $2.04\text{-}\mu\text{m}^2$ cells. The bitline connections to the cell are made through metal 2. The wordlines are continuous polycrystalline conductor (PC) lines but, for performance reasons, they are strapped at metal 3 (M3) to reduce the wordline resistance.

Array features

The dense SRAM arrays used in the Blue Gene/L chip range in size from 256×72 (words \times bits) to 2048×78 . They are all designed with wordline redundancy, 4:1 decode, and one subarray. The access time and cycle time depend on the array size, process corner, operating temperature, and operating voltage. For the nominal process, 65°C and 1.5 V , the access and cycle times for the largest and smallest arrays are shown in **Table 2**. Within Blue Gene/L, these single-port dense SRAM arrays are used primarily in the L3 directory and network interface circuits.

Device design considerations

For SRAM cells, there are three parameters that are often considered when comparing cells: cell performance, cell size, and cell leakage. The dense SRAM cell was designed to achieve the maximum performance possible for the smallest cell size. **Table 3** shows a comparison of the device sizes used in the standard and dense SRAM cells. From **Table 3**, it can be seen that to achieve the smaller cell size, the width of all of the transistors had to be reduced. However, because the cell size was reduced, the transient delay due to the resistance and capacitance of the circuit across the cell was also lowered, so some of the performance loss due to the smaller device widths can therefore be regained. Because of reliability considerations, no attempt was made to redesign the devices to achieve either a shorter channel length or a narrower channel width than what was already supported in the technology. No attempt was made to reduce the leakage characteristics of the dense SRAM cell below what was already achievable through the base technology.

To achieve the optimized cell size, in addition to the previously mentioned device width reductions, aggressive layout rule shrinks [22] were required. These layout rule shrinks were applied to the following levels: active

Table 2 Access and cycle times for the largest and smallest arrays.

Array size (word \times bit)	Access time (ns)	Cycle time (ns)
256×72	1.06	1.36
$2,048 \times 78$	1.68	1.96

Table 3 SRAM device dimensions.

Device	Ground-rule minimum		Standard SRAM		Dense SRAM	
	Length (nm)	Width (nm)	Length (nm)	Width (nm)	Length (nm)	Width (nm)
Pass gate	120	160	130	170	130	160
Pull-down	120	160	120	270	110 [†]	220
Pull-up	120	160	130	160	120	150 [†]

[†]These design dimensions have been adjusted to account for optical proximity correction and on-wafer printing. The final on-wafer dimensions will be as if they were designed at ground-rule minimum.

diffusion area, gate area, contacts, metal 1, and via 1. No layout rule shrinks were required for metal 2 and above. The final dense SRAM design (up to metal 1) is shown in **Figure 15(a)**. The custom optical proximity correction (OPC) shapes for the active area, gate conductor, and metal 1 levels are shown in **Figure 15(b)** and the on-wafer SEM images in **Figures 15(c)** and **15(d)**.

Future considerations

The multiprocessor core hierarchical memory architecture embodied in the specialized Blue Gene/L processor is in line with the current trend for more general-purpose processors. Processor chips will continue to add multiple independent cores, and the local caches used by these cores will require ever-larger amounts of high-bandwidth embedded memory. The use of SRAM memory to realize these chips is being limited by available chip size, power constraints, and time-of-flight performance considerations. Furthermore, SRAM cell scaling is causing diminishing cell functionality margins because of reduced operating voltage and variations in device matching caused by statistical fluctuations in the number of dopant atoms contained in the device channel [23]. If embedded DRAM memory can continue to improve performance and maintain its large density advantage over SRAM, it can play a critical role in future processor systems.

The basis of embedded DRAM is a one-transistor, one-capacitor cell. It is important to be able to scale both the transistor and the capacitor in area while maintaining or

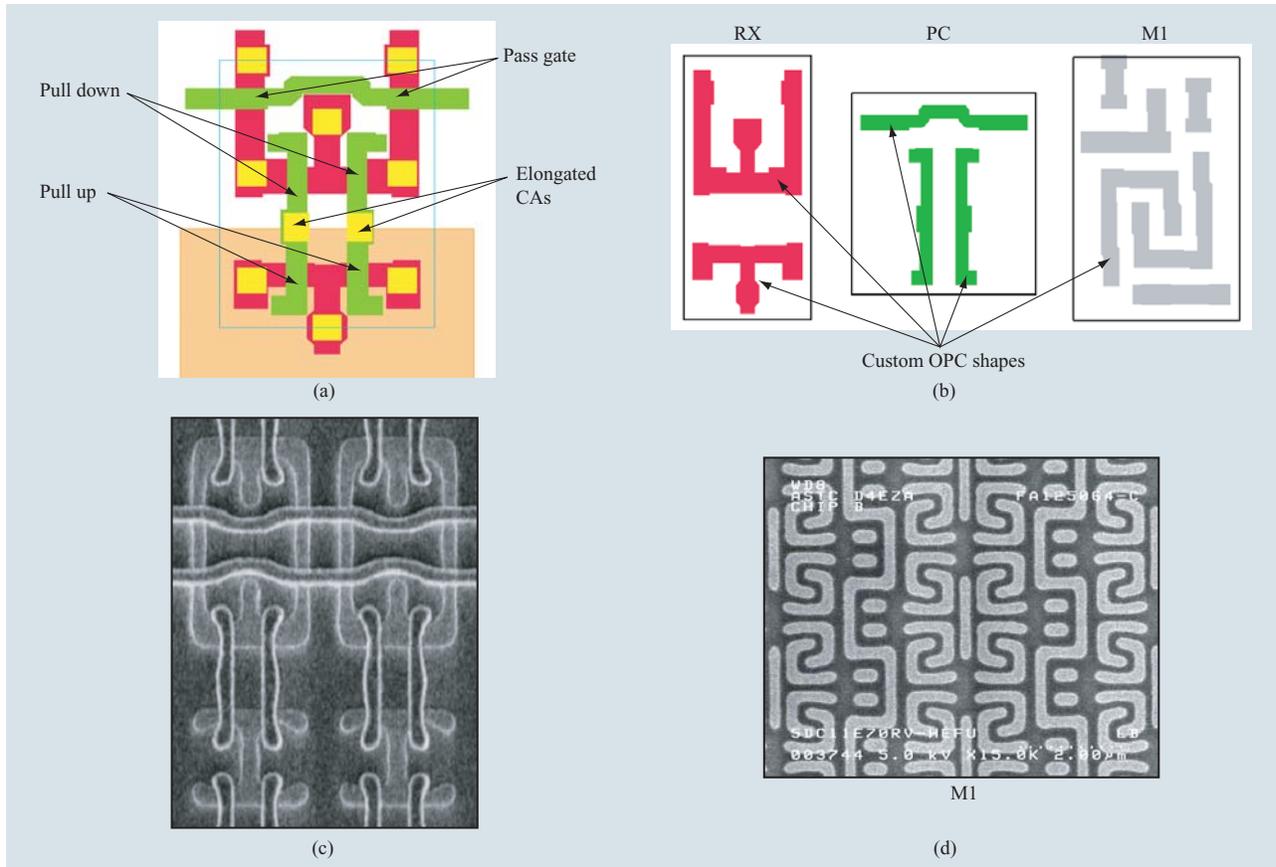


Figure 15

(a) Layout of the dense SRAM cell showing the special features. (b) The active area (RX), gate level (PC), and metal 1 (M1) designs for the dense SRAM cell. Custom OPC shapes have been added to the design shapes to ensure printability. (c) and (d) SEM micrographs respectively showing (c) RX and PC and (d) M1 for the dense SRAM cell.

improving the inherent performance of the cell. Let us examine some of the challenges we face if we are to continue scaling embedded DRAM cells and circuits.

The capacitor

The heart of a DRAM cell is its capacitor. Whether the capacitor is formed as a deep trench or a stacked type, most of the challenge in scaling a DRAM to the next technology node involves achieving the same high capacitance, but in a smaller space. This leads to aggressive scaling of the aspect ratio of the fabricated capacitor structures and to reducing the equivalent thickness of the capacitor insulator through physical thinning or increasing the dielectric constant of the material.

The basic operation of reading a DRAM cell involves turning on the array device and allowing the charge stored on the capacitor on one side of the device to pass onto a metal wire connected to the other side of the device. This metal line, or bitline, is shared with as many

as 500 other nonselected cells. The voltage generated on the bitline is detected or sensed by circuitry at the end of the bitline. A key metric in DRAM design is the amount of voltage or signal generated when a data cell is read by dumping its stored charge onto its bitline. Since the charge initially stored in the cell capacitor equilibrates with the bitline capacitance, the voltage ultimately created is given by

$$signal = V_{cell} \times \frac{C_{cell}}{C_{bitline} + C_{cell}},$$

where V_{cell} is the voltage initially stored in the cell capacitor, C_{cell} is the capacitance of the cell capacitor, and $C_{bitline}$ is the capacitance of the bitline. The quotient of the capacitance terms in the preceding equation is known as the *transfer ratio* and is typically of the order of 0.2.

As technology shrinks, the capacitance of the bitline wire typically remains relatively constant, since the wire thickness is not scaled and the capacitance reduction due

to the shrinkage in the length of the line is offset by the capacitance increase due to the reduced spacing between the lines. As long as the number of cells per bitline remains constant, the capacitance of the cell capacitor must remain constant to maintain the transfer ratio and the signal available for sensing cell data correctly. Unfortunately, as technology scales, the devices that detect the bitline signal suffer from the same matching degradation as was mentioned for the SRAM cell above, requiring high levels of signal to detect the correct data state of the cell.

Since simple dimensional scaling causes the capacitance of a cell to decrease, maintaining a constant cell capacitance from generation to generation is possible only with continuous improvements in fabrication technology or increasing permittivity of the dielectric materials. Since new invention is not always assured, it is desirable to explore options that allow capacitance to be reduced. The obvious option is to reduce the bitline capacitance by reducing the number of bits per bitline. This reduces the capacitance on the bitline and allows a reduction of the cell capacitance while maintaining a high signal transfer ratio. This solution, unfortunately, degrades the density of the memory system by causing the fixed area overhead of the sensing circuitry at the end of the bitline to be amortized over fewer memory cells. The key to enabling fewer cells per bitline is to have a very efficient design of the array sensing and control circuitry. With proper design area, efficient designs using 64 or fewer cells per bitline are feasible.

Reducing cell and bitline capacitance has additional performance benefits. A lower-capacitance bitline is faster to switch, improving array speed. Similarly, a lower-capacitance cell charges more quickly, enabling faster write time when switching the cell to its opposite data polarity state. This is shown in **Figure 16**. It should be noted that reducing cell capacitance too far will limit the data retention time and increase the susceptibility of the cell to SER caused by energetic background radiation, as discussed above.

The transistor

Because CMOS logic device technology has been aggressively scaled to create the fastest possible transistor switching, we have now arrived at a point at which the channel length of the conventional DRAM array transistor is significantly longer than the minimum used in the logic devices, and is limiting the continued area scaling of the array cell.

DRAM array transistors must maintain off currents in the range of 10^{-14} A, which allows data to be retained for the several-millisecond interval required between cell refreshes. (In comparison, logic device leakage ranges from 10^{-12} A to 10^{-7} A per device.) A gate voltage swing

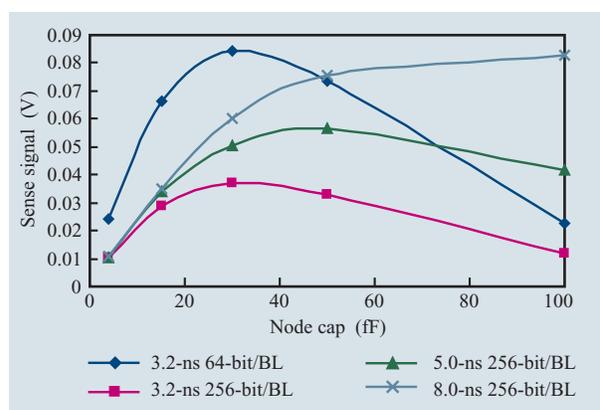


Figure 16

Depending on the performance target, the choice of node capacitance and bitline (BL) length may optimize differently.

of approximately 3 V is necessary to achieve both the low off current and high-gate overdrive that enables writing a high voltage to the storage node. This high operating voltage requires the use of a gate oxide of approximately 5 nm to 6 nm compared with the 1 nm to 1.5 nm seen in the state-of-the-art logic devices of today.

In addition, the low-leakage requirement necessitates careful grading of the doping profiles of the junctions used in the source and drain of the device in order to limit the electric fields that increase leakages in the junctions. As a result, such DRAM cells have reached their minimum channel limit. Further scaling of the cell area causes the width and available drive current of these devices to be significantly reduced. Such reduced drive current becomes a limitation in the read and write performance for the cell.

If it were possible to relax the leakage constraints for the embedded DRAM by significantly decreasing the interval between cell refreshes, the techniques of gate oxide reduction and device halo implantation used for logic device scaling could be employed to drive continued reduction of the embedded DRAM array device area. During a cell refresh cycle, the DRAM memory is unavailable to the system, and just increasing the rate of refresh for the DRAM would reduce this availability unacceptably. However, innovative circuit designs have enabled reduction of the retention interval to less than 100 μ s while maintaining 99% memory availability to the system [24]. With data retention in the tens of μ s range, the off-current requirements can be relaxed to several pA per cell. Such off currents are achieved today by the low-power logic FET designs using a gate oxide of 2 nm to 2.5 nm and a gate voltage swing of 1.5 V (**Figure 17**).

The use of aggressive oxide thickness scaling can enable shrinking of the array device length and width while still

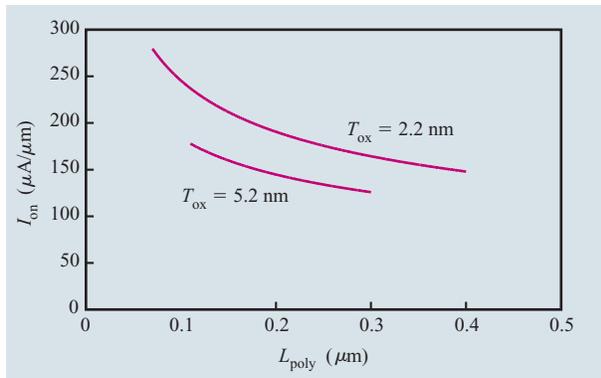


Figure 17

Comparison of drain current for 52-Å DRAM-type device and low-power logic 22-Å device. Measurements taken at $V_{ds} = 1.2$ V for both cases. $V_{gs} = 2.5$ V for the 52-Å device and 1.5 V for the 22-Å device.

maintaining drive current. The higher doping levels of the junctions in the scaled transistors also improve the resistance of the connection to the trench.

The switching speeds of the circuits driving the DRAM gate also improve as their oxide thickness and voltage swing decrease. In a DRAM memory system, the DRAM gate voltage is typically generated internally with a system of charge pumps. Reducing the level of the generated voltage significantly improves the efficiency and the area required for these pump circuits.

The leakage current through the gate of the device will, however, ultimately limit the data retention for the array cell. For the 64- μ s retention, this limit will be approximately 1.8 nm to 2.0 nm for conventional oxide gate dielectrics. These same device considerations can be applied to adapt the silicon-on-insulator devices used in the highest-performance logic technologies for use as an array device [25].

Conclusion

We have described the embedded DRAM platform that provides the basis for the Blue Gene/L SoC. The focus of this technology platform is to deliver high-density, high-bandwidth dynamic memory with true high-performance logic and SRAM technology. This second-generation embedded DRAM technology provides the basis for advanced cache chips for IBM eServer applications as well. It has also found widespread use in network applications and digital signal processors [26]. This technology does not differ from our standard logic technology, which is a subset of the embedded DRAM platform, in either performance or reliability. Embedded DRAM technology has been shown to be an excellent and cost-effective solution.

Going forward, using the capacitor and device scaling approaches outlined here with continued innovation in circuit techniques, it is anticipated that embedded DRAM memory systems can achieve cycle times of <2.5 ns and data rates of >1.5 GHz, while still maintaining a 3 \times to 4 \times density advantage over an SRAM solution. This level of performance and density will be ideal for providing embedded data storage for the next generation of multicore processor designs.

Acknowledgments

The authors acknowledge the members of the embedded DRAM technology development team, the DRAM and SRAM design and characterization teams, and the IBM Advanced Semiconductor Technology Center line in East Fishkill, New York, for building the prototypes and early user hardware.

*Trademark or registered trademark of International Business Machines Corporation.

References and notes

1. S. S. Iyer and H. L. Kalter, "Embedded DRAM Technology: Opportunities and Challenges," *IEEE Spectrum* **36**, No. 4, 56–64 (April 1999).
2. E. J. Nowak, "Maintaining the Benefits of CMOS Scaling When Scaling Bogs Down," *IBM J. Res. & Dev.* **46**, No. 2/3, 169–180 (March/May 2002).
3. K. Muellera, S. Gupta, S. Pae, M. Agostinelli, and P. Aminzadeh, "6-T Cell Circuit Dependent GOX SBD Model for Accurate Prediction of Observed vccmin Test Voltage Dependency," *Proceedings of the IEEE 42nd Annual International Reliability Physics Symposium*, 2004, pp. 426–429.
4. NeoMagic Corporation, 3250 Jay St., Santa Clara, CA 95054. The NeoMagic memory and logic graphics processors were introduced in 1993; see <http://www.neomagic.com/about/history.asp>.
5. G. Giacalone, R. Busch, F. Creed, A. Davidovich, S. Divakaruni, C. Drake, C. Ematrudo, J. Fifield, M. Hodges, W. Howell, P. Jenkins, M. Kozyrczak, C. Miller, T. Obremski, C. Reed, G. Rohrbaugh, M. Vincent, T. von Reyn, and J. Zimmerman, "A 1 MB, 100 MHz Integrated L2 Cache Memory with 128b Interface and ECC Protection," *Proceedings of the IEEE International Solid-State Circuits Conference (ISSCC)*, 1996, pp. 370–371.
6. A guide to IBM embedded DRAM offerings can be found at http://www-306.ibm.com/chips/techlib/techlib.nsf/products/Embedded_DRAM.
7. S. Crowder, R. Hannon, H. Ho, D. Sinitzky, S. Wu, K. Winstel, B. Khan, S. R. Stiffler, and S. S. Iyer, "Integration of Trench DRAM into a High-Performance 0.18 μ m Logic Technology with Copper BEOL," *Proceedings of the IEEE International Electron Devices Meeting (IEDM)*, 1998, pp. 1017–1020.
8. V. Klee, J. Norum, R. Weaver, S. S. Iyer, C. R. Kothandaraman, J. Chiou, M. Chen, N. Kusaba, S. Lasserre, C. Liang, J. Liu, A. Lu, P. R. Parries, B. J. Park, J. Rice, N. Robson, D. Shum, B. Khan, Y. Liu, A. Sierkowski, C. Waskiewicz, P. Wensley, T. Wu, and J. Yan, "A 0.13 μ m Logic-Based Embedded DRAM Technology with Electrical Fuses, Cu Interconnect in SiLK, sub-7 ns Random Access Time and Its Extension to the 0.10 μ m Gene/Lration," *Proceedings of the IEEE International Electron Devices Meeting (IEDM)*, 2001, *Technical Digest*, pp. 18.5.1–18.5.4.

9. E. Adler, J. K. DeBrosse, S. F. Geissler, S. J. Holmes, M. D. Jaffe, J. B. Johnson, C. W. Koburger III, J. B. Lasky, B. Lloyd, G. L. Miles, J. S. Nakos, W. P. Noble, Jr., S. H. Voldman, M. Armacost, and R. Ferguson, "The Evolution of IBM CMOS DRAM Technology," *IBM J. Res. & Dev.* **39**, No. 1/2, 167–188 (January/March 1995).
10. T. V. Rajeevakumar, T. Lii, Z. A. Weinberg, G. B. Bronner, P. McFarland, P. Coane, K. Kwietniak, A. Megdanis, K. J. Stein, and S. Cohen, "Trench Storage Capacitors for High Density DRAMs," *Proceedings of the IEEE International Electron Devices Meeting (IEDM)*, 1991, *Technical Digest*, pp. 835–838.
11. S.-I. Kimura, Y. Kawamoto, N. Hasegawa, A. Hiraiwa, Y. Nakagome, M. Aoki, T. Kisu, H. Sunami, and K. Itoh, "An Optically Delineated 4.2- μm^2 Self-Aligned Isolated-Plate Stacked-Capacitor DRAM Cell," *IEEE Trans. Electron Devices* **35**, No. 10, 1591–1595 (October 1988).
12. M. Takeuchi, K. Inoue, M. Sakao, T. Sakoh, T. Kitamura, S. Arai, T. Iizuka, T. Yamamoto, H. Shirai, Y. Aoki, M. Hamada, R. Kubota, and S. Kishi, "A 0.15 μm Logic Based Embedded DRAM Technology Featuring 0.425 μm^2 Stacked Cell Using MIM (metal-insulator-metal) Capacitor," *Proceedings of the Symposium on VLSI Technology*, 2001, *Digest of Technical Papers*, pp. 29–30.
13. T. Obremski, "Advanced Non-Concurrent BIST Architecture for Deep Sub-Micron Embedded DRAM Macros," Ph.D. dissertation, University of Vermont, Burlington, May 2001.
14. N. Watanabe, F. Morishita, Y. Taito, A. Yamazaki, T. Tanizaki, K. Dosaka, Y. Morooka, F. Igaue, K. Furue, Y. Nagura, T. Komoike, T. Morihara, A. Hachisuka, K. Arimoto, and H. Ozaki, "An Embedded DRAM Hybrid Macro with Auto Signal Management and Enhanced-On-Chip Tester," *Proceedings of the IEEE International Solid-State Circuits Conference (ISSCC)*, 2001, *Digest of Technical Papers*, pp. 388–389.
15. J. Barth, D. Anand, J. Dreibelbis, and E. Nelson, "A 300 MHz Multi-Banked DRAM Macro Featuring GND Sense, Bit-Line Twisting and Direct Reference Cell Write," *Proceedings of the IEEE International Solid-State Circuits Conference (ISSCC)*, 2002, *Digest of Technical Papers*, vol. 1, pp. 156–157.
16. R. McConnell, U. Moller, and D. Richter, "How We Test Siemens Embedded DRAM Cores," *Proceedings of the International Test Conference (ITC)*, 1998, pp. 1120–1126.
17. R. C. Aitken, "On-Chip Versus Off-Chip Test: An Artificial Dichotomy," *Proceedings of the International Test Conference (ITC)*, 1998, p. 1146.
18. J. Dreibelbis, J. Barth, H. Kalter, and R. Kho, "Built-In Self Test for Embedded DRAM," *Proceedings of the 6th IEEE North Atlantic Test Workshop*, 1997, pp. 19–27.
19. E. B. Eichelberger and T. W. Williams, "A Logic Design Structure for LSI Testability," *J. Design Automation & Fault-Tolerant Computing* **2**, 165–178 (May 1978).
20. J. Dreibelbis, J. Barth, H. Kalter, and R. Kho, "Processor-Based Built-in Self-Test for Embedded DRAM," *IEEE J. Solid-State Circuits* **33**, No. 11, 1731–1740 (November 1998).
21. L. K. Han, S. Biesemans, J. Heidenreich, K. Houlihan, C. Lin, V. McGahay, T. Schiml, A. Schmidt, U. P. Schroeder, M. Stetter, C. Wann, D. Warner, R. Mahnkopf, and B. Chen, "A Modular 0.13 μm Bulk CMOS Technology for High Performance and Low Power Application," *Proceedings of the Symposium on VLSI Technology*, 2000, *Digest of Technical Papers*, pp. 12–13.
22. O. Bula, R. Mih, E. Jasinski, D. Hoyniak, A. Lu, J. Harrington, and A. McGuire, "Pushing SRAM Densities Beyond 0.13- μm Technology in the Year 2000," *Proceedings of the 20th Annual BACUS Symposium on Photomask Technology and Management*, 2001, pp. 601–608.
23. T. B. Hook, M. Breitwisch, J. Brown, P. Cottrell, D. Hoyniak, Chung Lam, and R. Mann, "Noise Margin and Leakage in Ultra-Low Leakage SRAM Cell Design," *IEEE Trans. Electron Devices* **49**, No. 8, 1499–1501 (August 2002).
24. T. Kirihata, P. Parries, D. Hanson, H. Kim, J. Golz, G. Fredeman, R. Rajeevakumar, J. Griesmer, N. Robson, A. Cestero, M. Wordeman, and S. Iyer, "An 800MHz Embedded DRAM with a Concurrent Refresh Mode," *Proceedings of the IEEE International Solid-State Circuits Conference (ISSCC)*, 2004, *Digest of Technical Papers*, vol. 1, pp. 206–523.
25. M. Kumar, M. D. Steigerwalt, B. L. Walsh, T. L. Doney, D. Wildrick, K. A. Bard, D. M. Dobuzinsky, P. A. McFarland, C. E. Schiller, B. Messenger, S. E. Rathmill, A. R. Gasasira, P. C. Parries, S. S. Iyer, S. E. Chaloux, and H. L. Ho, "A Simple and High-Performance 130 nm SOI EDRAM Technology Using Floating-Body Pass-Gate Transistor in Trench-Capacitor Cell for System-on-a-Chip (SoC) Applications," *Proceedings of the IEEE International Electron Devices Meeting (IEDM)*, 2003, *Technical Digest*, pp. 17.4.1–17.4.4.
26. Y. Adelman, D. Agur, T. Bennun, O. Chalak, Z. Greenfield, R. Holzer, M. Jalfon, A. Kadry, E. Kraus, F. Lange, H. Meirov, A. Olofsson, O. Raikhman, D. Treves, S. Zur, and R. Talmudi, "A 600MHz DSP with 24Mb Embedded DRAM with an Enhanced Instruction Set for Wireless Communications," *Proceedings of the IEEE International Solid-State Circuits Conference (ISSCC)*, 2004, *Digest of Technical Papers*, vol. 1, pp. 418–419.

Received July 20, 2004; accepted for publication September 10, 2004; Internet publication April 12, 2005

Subramanian S. Iyer *IBM Systems and Technology Group, 2070 Route 52, Hopewell Junction, New York 12533 (sSIyer@us.ibm.com)*. Dr. Iyer received his B.Tech. degree in electrical engineering from the Indian Institute of Technology, Bombay, in 1977, and his M.S. and Ph.D. degrees in electrical engineering from the University of California at Los Angeles in 1978 and 1981, respectively. He joined the IBM Thomas J. Watson Research Center in 1981 and was manager of the Exploratory Structures and Devices Group until 1994, when he founded SiBond LLC to develop and manufacture silicon-on-insulator materials. Since 1997 he has been with the IBM Microelectronics Division, Semiconductor Research and Development Center, where he is currently a Distinguished Engineer and manages 90-nm bulk logic development including the embedded DRAM and autonomic chip projects. He has coauthored several book chapters and more than 150 articles in technical journals. Dr. Iyer is a Fellow of the IEEE.

John E. Barth, Jr. *IBM Systems and Technology Group, Essex Junction Development Laboratory, 3000 River Road, Essex Junction, Vermont 05452 (jbarth@us.ibm.com)*. Mr. Barth is a Senior Technical Staff Member in the High Performance Embedded DRAM and SRAM Development Department. He received a B.S.E.E. degree from Northeastern University in 1987 and an M.S.E.E. degree from National Technological University in 1992. After receiving his B.S. degree, he joined the IBM Development Laboratory and was involved with the design of a 16-Mb DRAM product featuring embedded ECC and SRAM cache. He subsequently worked on the array design for the 16/18-Mb DRAM product. In 1994, he began work in his current field of wide-I/O, high-performance DRAM macros for embedded applications. Mr. Barth has coauthored numerous IEEE papers on embedded DRAM design and built-in self test of DRAM structures for ASIC systems.

Paul C. Parries *IBM Systems and Technology Group, 2070 Route 52, Hopewell Junction, New York 12533 (parries@us.ibm.com)*. Dr. Parries is a Senior Technical Staff Member working in embedded DRAM technology and product development. He has had nearly 20 years of experience working on commodity DRAM and embedded DRAM technologies and has expertise in process development, DRAM process integration, circuit design, and characterization. Dr. Parries received his B.S. and Ph.D. degrees in electrical engineering from the University of Minnesota prior to joining IBM in 1985.

James P. Norum *IBM Systems and Technology Group, 2070 Route 52, Hopewell Junction, New York 12533 (norum@us.ibm.com)*. Mr. Norum received his B.S. degree in electrical engineering from Yale University in 1987 and his M.E. degree in electrical engineering from Cornell University in 1989. He joined IBM in 1987 and worked on process integration issues related to IBM advanced bipolar transistor technology. He next worked in manufacturing as a reactive ion etch (RIE) engineer. He managed Plasma Equipment Engineering Services and subsequently Manufacturing Engineering RIE until 2000. Since then, Mr. Norum has been involved in embedded DRAM integration; he currently manages the bulk CMOS process integration effort.

James P. Rice *IBM Systems and Technology Group, 2070 Route 52, Hopewell Junction, New York 12533 (ricejp@us.ibm.com)*. Mr. Rice received A.A.S. and B.S. degrees in computer science from Park College and an M.S. degree in electrical engineering from Columbia University. He joined

IBM in 1997 as a functional test engineer, focusing on DRAM characterization, after working for Micron Technology for several years. He acted as lead characterization engineer for a number of IBM 64-MB and 256-MB DRAM projects, and later became the manager of the Embedded DRAM Device Design and Characterization Department, focusing on development and qualification of IBM embedded DRAM technology. Mr. Rice is currently the project manager for 300-mm-in-line characterization, focusing on yield ramp of IBM leading-edge technologies.

Lyndon R. Logan *IBM Systems and Technology Group, 2070 Route 52, Hopewell Junction, New York 12533 (llogan@us.ibm.com)*. Dr. Logan joined IBM in 1990 after graduating with a Ph.D. degree in physics from the University of Western Ontario. He spent his first two years conducting research in electron-transport theory and the thermodynamics of solid interfaces. He subsequently spent several years on TCAD-software development and applications. Since 1997 he has worked in the areas of CMOS technology development and manufacturing, focusing on semiconductor device design in 0.18- μm , 0.13- μm , and 0.09- μm technologies. Dr. Logan is currently working on 65-nm silicon-on-insulator (SOI) device and technology design.

Dennis Hoyniak *IBM Systems and Technology Group, 2070 Route 52, Hopewell Junction, New York 12533 (hoyniak@us.ibm.com)*. Mr. Hoyniak received a B.S. degree in electrical engineering and an M.S. degree in nuclear engineering from Pennsylvania State University in 1975 and 1977, respectively, and an M.S. degree in materials science from the University of Vermont in 1997. Since joining IBM in 1980, Mr. Hoyniak has been engaged in the development of silicon-based devices. His current work interest is in the area of SRAM cell development.