

IBM POWER6 SRAM arrays

D. W. Plass
Y. H. Chan

The IBM POWER6™ microprocessor presented new challenges to array design because of its high-frequency requirement and its use of 65-nm silicon-on-insulator (SOI) technology. Advancements in performance (2X to 3X improvement over the 90-nm generation) and design margins (cell stability, writability, and redundancy coverage) were major focus areas. Key elements of the POWER6 processor chip arrays include paradigm shifts such as thin memory cell layout, large signal read (without a sense amplifier), segmented bitline structure, unclamped column-half-select scheme, multidimensional programmable timing control, and separate elevated static random access memory (SRAM) power supply. There are two main array categories on the POWER6 microprocessor chip: core and nest. Processor core arrays use a single-port, 0.75- μm^2 , six-transistor (6T) cell and operate at full frequency, whereas the surrounding nest arrays use a smaller 0.65- μm^2 cell that operates at half or one-quarter of the core frequency in order to achieve better density and power efficiency. The core arrays include the 96-KB instruction cache (I-cache) and the 64-KB data cache (D-cache), with associate lookup-path SRAM macros. The I-cache is a four-way set-associative, single-port design, whereas the D-cache is an eight-way design with dual read ports to handle multithreading capability. The lookup-path arrays contain content-addressable memory (CAM) and RAM macros with integrated dynamic hit logic circuitry. In the nest portion, an 8-MB level 2 (L2) D-cache and a level 3 (L3) directory (1.2 MB) make up the largest arrays. The latter macro designs use longer bitlines and orthogonal word-decode layouts to achieve high array-area efficiency.

Introduction

IBM entered the gigahertz (GHz) frequency regime with the POWER4* processor and extended this trend beyond 2 GHz with the POWER5* processor, which was introduced in 2003. Going well beyond a simple evolution with technology scaling, the design goal for the POWER6* processor was to enable array operation at 2X to 3X the frequency of previous designs while simultaneously increasing the array capacities.

This led to the conclusion that the conventional array structure of earlier chips, with long bitlines and small-signal dynamic sense amplifiers, should be replaced with short, low-capacitance local bitlines (LBLs) driving logic circuits that in turn drive longer global bitlines (GBLs). This arrangement is sometimes referred to as a *domino sense scheme*, and the key benefit lies in the ability to write the cells and restore the LBLs in a shorter cycle time.

As well as raising the bar for performance and density, the POWER6 processor array designs must also address the issues with technology scaling, which are more apparent in circuits that use the smallest devices, such as the static random access memory (SRAM) cells themselves. Device variability, including the effect of random dopant fluctuation, is an increasingly difficult reality that must be neutralized by more-stringent layout and circuit restrictions, as well as increased array redundancy to circumvent any marginal cells.

Array technology overview

The arrays are designed using the IBM 65-nm complementary metal-oxide semiconductor (CMOS) technology, which uses partially depleted silicon-on-insulator (SOI) devices that come in various threshold voltages in order to optimize performance, power, and

©Copyright 2007 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.

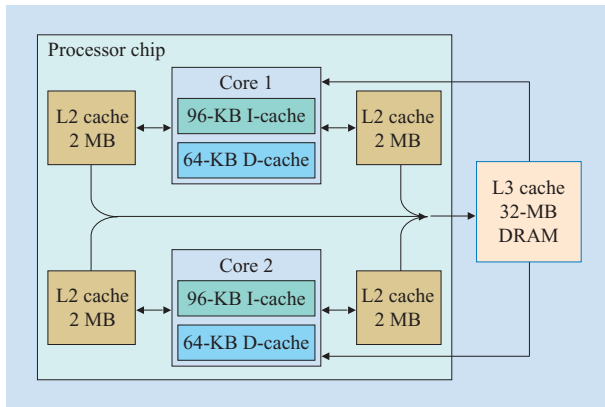


Figure 1

Cache block diagram.

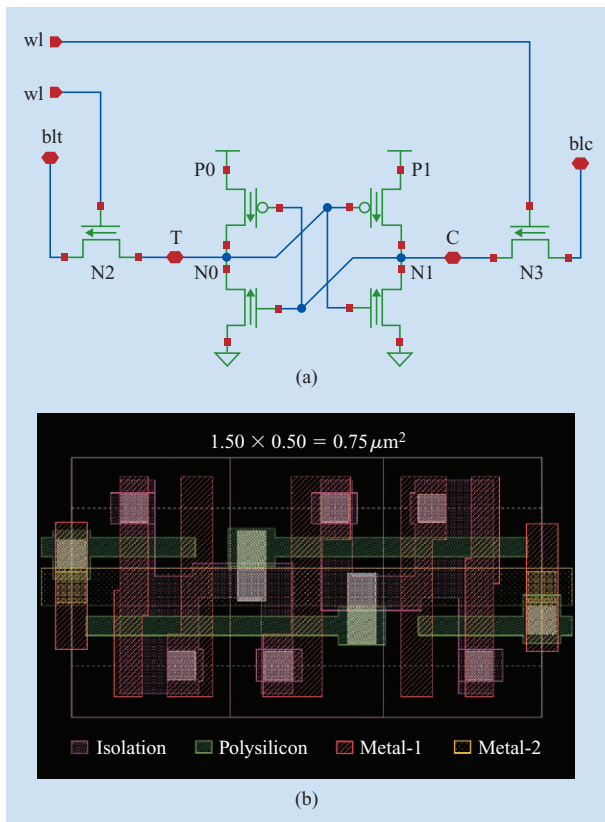


Figure 2

(a) Six-transistor (6T) cell schematic and (b) SRAM cell layout. (wl: wordline; blt: bitline true; blc: bitline complement.)

noise margins. In order for the SRAM operation to be tuned independently from the logic, separate threshold voltage (V_t) tailoring is used in conjunction with an additional dedicated power supply (V_{cs}) for the memory

cells and critical array signals. The POWER6 processor features two simultaneous multithreading (SMT) processor cores and a surrounding nest structure, with a different SRAM approach for each category. To obtain higher density and lower power consumption, the nest arrays use a smaller cell and more-compact support circuits, and they extensively use devices with high V_t . The larger array macros use up to six levels of metal internally to improve density and reduce the *resistance · capacitance* (RC) time constant delay of long wires.

Cache hierarchy

The POWER6 processor cache hierarchy consists of three levels (Figure 1). The first level is composed of a 96-KB level 1 (L1) instruction cache (I-cache) that resides in the instruction fetch unit (IFU) and a 64-KB L1 data cache (D-cache) that is located in the load/store unit (LSU), both being part of the core. The second level consists of a private 4-MB level 2 (L2) cache for each core, split into left and right halves that straddle the central core regions. The third level consists of a 32-MB L3 off-chip cache that is shared by the two processors. The controller for the L3 cache is integrated into the POWER6 chip, but the data resides on another chip employing embedded DRAM (dynamic random access memory) for high density.

Each level of cache has its own performance/density optimization level. The very-high-frequency processor cores in the POWER6 chip represent a significant challenge compared with those of prior designs; the subarrays (SAs) must be capable of cycling in less than 200 ps, which requires very short bitlines and higher power circuitry. On the other hand, the large private L2 caches (8 MB total) are a significant upgrade from the somewhat smaller (<2 MB) shared L2 cache employed by the POWER4 and POWER5 processors. The POWER6 design is optimized for higher density and lower power and performance. Because the L2 cache unit is designed to operate at half the frequency of the processor core, the widths of the load and store interfaces were doubled, driving 64 bytes of data per L2 cache cycle into the processing core and receiving 16 bytes of store data.

SRAM cell technology

The POWER6 processor SRAM cell designs take advantage of technology leverage provided by the IBM CMOS 65-nm SOI process. To provide the best possible control over device channel length, the polysilicon (polySi) gates for both logic and arrays all run in the same (horizontal) direction on a fixed pitch. The cell layout that conforms to this arrangement is referred to as a *thin cell* [1, 2] (Figure 2). The thin-cell configuration is narrow in the vertical direction, so the bitline wiring length is minimized.

SOI technology provides device isolation with no underlying well structure, thus enabling a butted junction with a salicide bridge that connects the source–drain region of the n-FETs to that of the p-FETs without resorting to the first-layer metal, as shown in the layout view of the cell (Figure 2). Although the device-isolation shapes are more complex, the wiring of the internal cell node is not blocking the path of a metal-1 bitline, thus enabling a lower bitline wiring capacitance than that which is possible with a stacked via and metal-2 approach. The wiring of the internal cell node is completed by an elongated contact that shorts the polySi gate to the adjacent source–drain in the p-FET region, providing a dense layout without an additional level of interconnect.

With a high-frequency core that demands a high-performance cell, and with a large nest (8-MB L2) requiring high density and reduced power, the optimal solution is to employ two cell sizes. A $0.75\text{-}\mu\text{m}^2$ cell ($1.5\ \mu\text{m} \times 0.5\ \mu\text{m}$) provides the foundation for the core arrays, whereas a $0.65\text{-}\mu\text{m}^2$ cell ($1.3\ \mu\text{m} \times 0.5\ \mu\text{m}$) with reduced device widths is the basis for the nest array designs. Since the larger cell contains wider devices and is never used with more than 16 cells per bitline, it is inherently more resistant to read disturbs¹; hence, the pass-gate² V_t is lowered in order to provide additional performance without decreasing stability.

Circuit techniques

Dual supply voltage

To better optimize the SRAM performance, stability, and power, a separate power supply voltage (V_{cs}) powers the memory cells and localizes control signals such as the wordline and bitline precharge. This dedicated SRAM supply is targeted to be about 150 mV higher than the main processor voltage (V_{dd}), although a larger delta is desirable for low- V_{dd} operation (reduced power). Because of technology scaling, the SRAM cells have increased V_t variability; the number of dopant atoms that influence the V_t is reduced as the device area diminishes. Along with a degraded stability margin, these V_t variations cause leakage problems at the low extreme and poor performance at the high extreme. Centering the SRAM six-device cell V_t higher and keeping the V_{cs} sufficiently high with a dedicated supply mitigates these problems while still allowing the logic circuits to scale further with respect to voltage for ac power reduction.

Domino sense scheme

SRAM designs in previous generations of IBM servers use sense amplifiers for memory array reading [3]. Key

attributes of the sense amplifier design include fast read performance and relatively low peripheral area overhead in the column circuitry. As technology scales beyond 90 nm, new challenges arise. First, sense amplifier read access speed does not scale well with technology. Second, the slow slew rate³ on long bitlines in a sense amplifier array aggravates cell read stability. Third, variability in performance and functionality becomes worse in sub-90-nm technology because of sense device offset, particularly in SOI technology in which the floating-body history effect drives the need for body contacts and the effectiveness of the body contact is lower. To overcome these challenges, the SRAMs in the POWER6 processor have switched from a sense amplifier to the large-signal domino sense scheme [4, 5]. A segmented hierarchical bitline structure replaces traditional long bitlines. The hierarchical bit column structure consists of short LBLs with domino read inverters that are connected to form longer GBLs. In the POWER6 processor core, to support the very high frequency, LBLs of the SRAM macros are only 16 or even 8 cells deep. The local and global column circuitry is designed and optimized for fastest read access in order to meet the goal of 2X to 3X improvement in performance over previous systems. Outside the core, functional units operate at fractions of the core frequency. For better array area and power efficiency, LBLs of the non-core SRAM are longer and contain either 32 or 64 cells. These slower SRAMs typically run at 2X to 4X the core cycle time.

In addition to having faster read performance and better technology scalability, the POWER6 processor SRAM domino read scheme also improves memory cell stability. With the segmented bitline structure, LBLs are shorter and have faster slew rates than before. The cell read current decays at a faster rate, which minimizes the duration of the *cell disturb* condition. This allows the memory cell devices to be tuned and optimized for faster access without causing a *read disturb*.

Elimination of clamped bitlines

To further improve cell stability, the design approach for the POWER6 processor SRAM was to eliminate the clamped half-selected bitlines. Half-selected bitlines are bit columns that are not selected by the bit decoder in an active SA. In a conventional design, bitline restore is usually controlled by the bit decoder signals. When bit columns are not selected, the bitline restore devices are turned on, clamping the bitlines to the high-power supply (e.g., at V_{dd}). When a wordline is selected, the bit half-selected cells along this wordline conduct at a higher read current than the fully selected cells, since the unselected bitlines are clamped high and are not allowed to

¹A *disturb* is an interrupt or stability failure.

²The wider *pass-gate* devices control reading from and writing to the CMOS SRAM cell.

³The *slew rate* represents the maximum rate of change of signal at any point in a circuit.

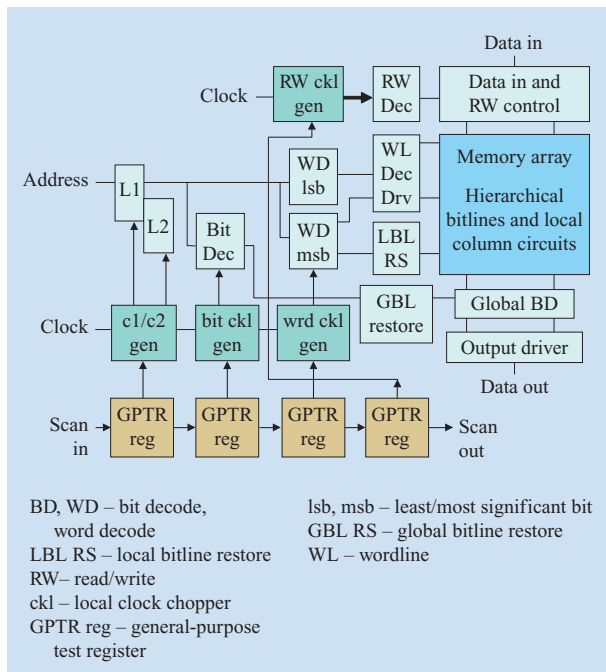


Figure 3

Programmable SRAM timing control diagram.

discharge. The read disturb problem in these half-selected cells is worse. In the POWER6 processor segmented bitline scheme, LBL restore is controlled by a word-path decode signal rather than the bit decode. When a wordline is selected in a local SA group, its local bit restore signal is shut off. The LBLs are not clamped high, but they are allowed to fall, hence reducing the amount of read current pumping into the cells. As a result, a read disturb is less likely and cell stability is improved.

Multiprogrammable array timing

In high-end microprocessor applications, arrays usually dominate critical paths of the machine, particularly in the core units. Like the L1 cache and its associate lookup-path arrays, SRAM macros typically have very aggressive timing in order to meet more demanding cycle-time goals. Internal timing control of these SRAM arrays, therefore, becomes crucial to avoid an impact on final chip yield and system frequency. For these reasons, previous generations of the IBM server SRAM design had programmable internal timing control to ensure hardware functionality. In the POWER6 processor, the introduction of a second SRAM-only power supply (V_{cs}) provides a new level of flexibility for array performance and functionality tuning. To couple this added power supply dimension, and to further support hardware optimization for yield, V_{min} (minimum V_{dd}/V_{cs}), V_{diff} (maximum V_{dd}/V_{cs}

differential), and f_{max} (maximum frequency), the SRAM programmable timing design was improved. In addition to having fully adjustable local c1/c2 clocks for boundary L1/L2 latches within the array macros, all critical array internal timing for bit, word, and data paths have multiprogrammable local clock generators. Each of these programmable clock generators is controlled by a dedicated mode block containing a series of general-purpose test register (GPTR) latches that are scan initialized. By setting values of the GPTR latches, all critical internal timings of the array (such as clock chopper pulse width, word/bit/data path alignments) can be adjusted for wafer, module, or system tests to aid in hardware characterization and debugging. Yield enhancement, system maximum frequency (f_{max}) tuning, and V_{min}/V_{diff} optimization are attained. **Figure 3** shows a simple block diagram of the POWER6 processor SRAM programmable timing control scheme.

Core arrays

L1 caches

The 96-KB L1 I-cache is a single-port SRAM, capable of either one 32-byte read or one 32-byte write for each core cycle. It is composed of six instances of the 16-KB array macro and is four-way set associative (512 entries \times 4 ways \times 68 bits wide). The 64-KB L1 dual-read D-cache can support either two 1-byte to 8-byte reads or one 1-byte to 32-byte write per processing core cycle. It is composed of four copies of the macro and is eight-way set associative (256 \times 8-ways \times 72 writes, 512 \times 8-way \times 36 reads). It is a store-through cache, so store instruction data is also sent to the L2 cache.

D-cache

The 16-KB D-cache macro features a two-stage pipeline that supports two independent reads or one write per cycle. Each read port provides 4 bytes and is 512 entries deep, whereas the double-bandwidth write operation provides individual control of the 8 bytes to the two macro halves (upper and lower) in parallel. Write data is passed through to the outputs, and a byte-steering signal selects which output port the upper- and lower-order bytes are routed to on a write cycle. A wordline disable feature provides an effective bypass mode with no additional pipeline latches, since no cells are disturbed.

The macro floorplan is shown in **Figure 4**. The memory cells are grouped into four vertical quadrants. Each vertical quadrant contains 72 active columns of cells and 2 spare columns for use in skip-over replacement, which uses additional FET switches to avoid a defective region. The 72 active columns are organized as nine groups of eight interleaved sets (9 bits = 1 byte). The macro is also divided horizontally, with the right and left halves each

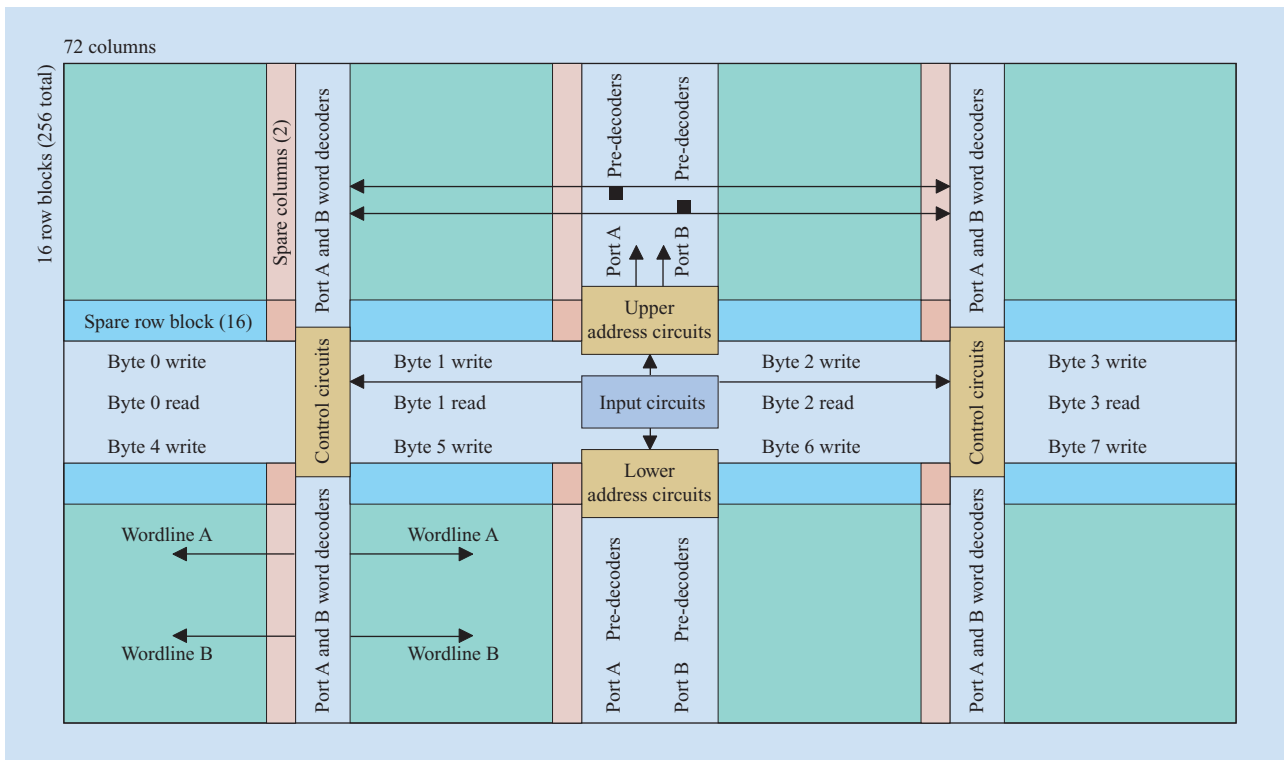


Figure 4

Data-cache macro floorplan.

having its own central row decoder region. Each side contains 16 usable SAs and 1 spare SA for use in skip-over replacement. The combined column- and SA-based redundancy allows six repair actions per macro.

Read and write addresses are received in the center of the macro and combined into pre-decoded address signals. There are two identical sets of decoding circuits (one for each port). The first-level pre-decoder multiplexes one of the read addresses with the write address. This first-stage pre-decoder uses a very narrow capture pulse and a wider enable pulse to create dynamic return-to-zero pulses that are propagated throughout the decoder tree. The narrow capture pulse eliminates the need for a separate inline latch on the address inputs within the macro.

The second-level pre-decoder generates an SA-select (SAS) signal and a row-in-SA (RIS) signal. There are 16 RIS signals, 16 SAS signals, and 1 spare SAS signal. In addition to the SAS signals, a dummy SAS signal is generated for each port. This dummy SAS signal is active whenever an SA is active within the given macro half (upper or lower) for a given port. This signal is sent to the remote timing circuit regions, where it is used to gate and control timings associated with the GLBs. The second-

level SAS pre-decoder incorporates a skip-over redundancy multiplexer (mux), enabling replacement of a defective SA with the spare SA. The word decoder combines an SAS and an RIS signal to select a wordline, which activates one pass transistor on one side of each cell in a row of cells. During a write, the write address feeds the decode circuitry for both ports, activating both pass transistors for a given row of cells. This scheme permits the use of a traditional word decoder circuit, which can be shared with single-port SRAMs.

The memory cell is the same six-transistor (6T) $0.75\text{-}\mu\text{m}^2$ cell used elsewhere in the core, but with each pass transistor connected to a separate wordline. Single-ended reads are triggered by activating a wordline and bitline precharge connected to a given side of the cell. Writes are possible when both wordlines for a given row are active along with the appropriate write controls.

Figure 5 shows the details of the datapath. The read path uses a hierarchical bitline approach. The pass transistors of 16 memory cells are connected to form an LBL pair, with one LBL serving each port. Unlike designs that use a static NAND circuit to combine the LBLs from a pair of SAs, before driving a GBL, this design reduces the LBL load by replacing the n-FET

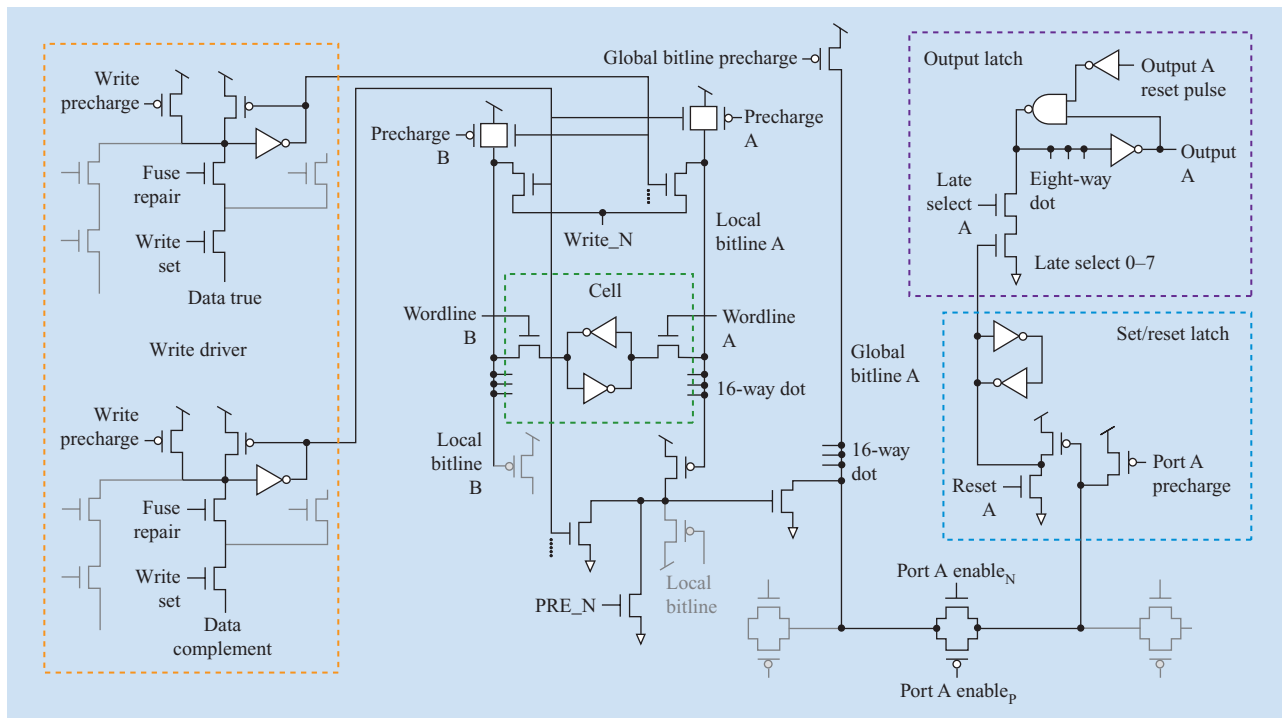


Figure 5

Data-cache datapath.

stack with a single device controlled by a signal (PRE_N) that mimics a bitline transition, based on the logical OR of the two adjacent SAS signals. In addition to reducing the load on the LBL, this approach slows the read access time of extremely fast cells while improving the delay for marginal cells because the p-FET turn-on and the n-FET turn-off are skewed. With programmable control of the slew rate and the down level of the PRE_N signal, the response of the single n-FET pull-down can be varied to perform read margin testing of the memory cells. This LBL read structure for each SA pair drives an n-FET pull-down on a GBL, which in turn is routed through a mux to a set/reset latch per column.

The final stage in the read path is a dynamic eight-way mux/latch. This latch is reset to a “1” at the beginning of the cycle and then sampled using one of eight chopped clocks based on the select signals received from the set-predict macro. Since sets are interleaved within a bit position, a compact 8:1 mux can be located near the dynamic-to-static conversion latches. This saves power and wire resources by eliminating the need to drive all 288 latch outputs to a remote set mux.

I-cache

The I-cache macro is similar to the D-cache macro, with the dual-read and byte-write capability removed. The

split wordline cell used on the D-cache to provide dual single-ended reads is replaced by a standard single wordline. The address path is also simplified with just one set of inputs instead of three, and the eight-way associativity is converted to four way.

Lookup-path arrays

Access to the POWER6 processor L1 cache is by way of a three-cycle path: A0, A1, and A2. For the D-cache in the LSU (as with the I-cache in the IFU), cycle A0 is the effective address (EA) generation, the output of which is fed to the D-cache, the data directory (D-directory), and the set-predict arrays. Cycle A1 is parallel SRAM access in the cache, directory, and set-predict arrays. Cycle A2 is the late select of the cache output set and the path to the formatters, along with directory hit or miss detection.

Figure 6 shows a functional block diagram of the LSU D-cache and lookup array structures. Two of the top timing-critical paths in this cache lookup include the *set-predict to late-select data out* and the *D-directory hit/miss generation*. These two critical paths dictate the LSU cycle-time capability.

Set-predict array

The set-predict/late-select function is implemented in the set-predict array (SETP) macro, which contains a

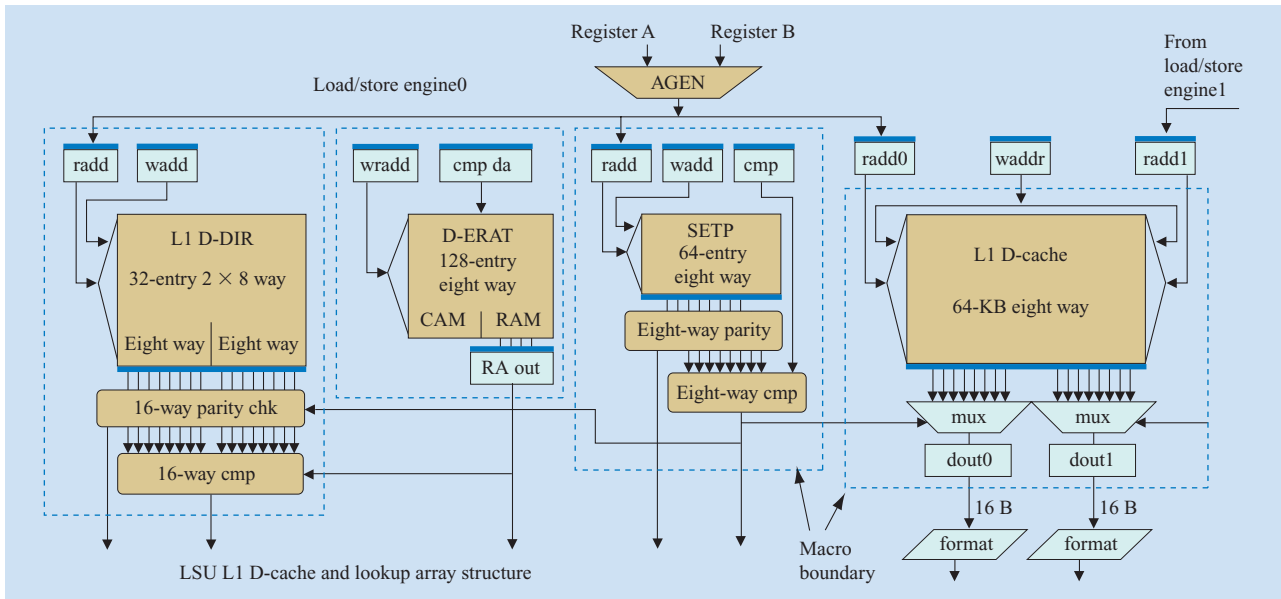


Figure 6

Block diagram of the load/store unit arrays. (radd: read address; wadd: write address; wradd: write read address; cmp da: compare data (containing effective address); RA out: real address out.)

64-entry \times 8-way \times 16-bit single-port array coupled to 8-way \times 16-bit dynamic comparators for hit logic generation. The SRAM is physically implemented as two half-arrays of 32 wordlines \times 4 ways \times 32 bits for fast read performance. The 32 wordlines are further divided into four segments of eight cells per LBL to yield the fastest possible read access time. It provides dual-phase, dynamic output directly to the 16-bit dynamic comparators. The comparators are pseudo-dynamic circuits with robust yet fast switching performance. The comparator/hit drivers send one out of eight dynamic late-select signals, feeding into the D-cache data-out muxes for set selection [4].

D-ERAT

The other array-limited critical path in the LSU is the directory hit/miss detection logic. It consists of two highly complex array macros: the data effective-to-real address translation (D-ERAT) and the D-directory (D-DIR) arrays. The D-ERAT performs the first-level translation for the data address. It is an 8-way \times 128-entry associative design implemented with content-addressable memory (CAM) and RAM array topology. It supports three page sizes (4 KB, 64 KB, and 64 MB) concurrently. The CAM array is 60 bits wide, and it contains the EA, page-size bits, and the thread valid bits. The RAM array is 59 bits wide and contains the real address (RA) and miscellaneous bits. The D-ERAT is a high-performance

single-cycle design accessed in the A1 cycle. The first half of the cycle is the dynamic CAM search and hierarchical match line generation. Twelve to fourteen CAM cells are connected to form the local match lines, four of which are then connected to form the global match/hit output. The second half of the cycle is the RAM access driven by the CAM hit lines. The RAM output containing an RA is cycle-staged and sent in the A2 cycle to the D-DIR macro for the RA hit/miss compare function. The D-ERAT array is the only SRAM macro on the POWER6 processor that does not use the standard 6T thin cells; the larger SRAM devices are shown schematically in Figure 7. The CAM cell is custom designed with 6T for non-CAM read/write operation. Internal latch nodes of the 6T design are coupled to a bit XOR with n-FET drain match line pull-down. The CAM memory cell size is $2.0 \mu\text{m} \times 1.75 \mu\text{m}$. The D-ERAT RAM cell is an 8T design with standard 6T topology in order to support normal read/write plus two stacked n-FET open-drain pull-downs for the CAM hit line read. The RAM cell layout is also custom, at $1.8 \mu\text{m} \times 1.75 \mu\text{m}$ and with a pitch matching the CAM cell height.

D-directory

The D-DIR macro is a large array containing both RAM and logic functions. It is a 16-way, 32-entry by 48-bit-wide RAM, with a 16-way by 38-bit-wide dynamic comparator for hit logic, plus 16-way static parity

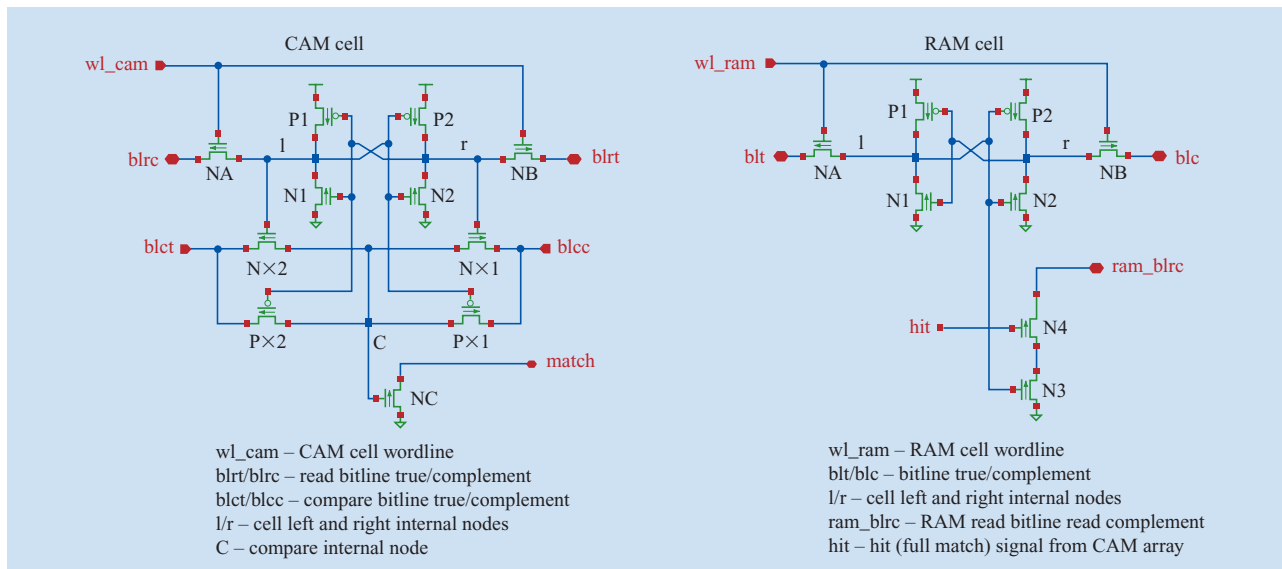


Figure 7

Content-addressable memory (CAM) and RAM (random access memory) cell circuits in the data effective-to-real address translation (D-ERAT) array.

checkers. The RAM contains 34 bits of RA and miscellaneous data. It is accessed in the A1 cycle. It provides both static and dual-phase dynamic outputs that are cycle-staged to drive, respectively, the parity checkers and comparators in the A2 cycle. To provide fast performance, the D-DIR 16 SAs are each physically implemented as 16 words \times 48 \times 2 bits wide. The bit column is one level deep with 16 cells along the bitlines. This is the only SRAM on the POWER6 processor chip that does not use the hierarchical bitline structure. The dynamic comparators and parity checkers are placed directly adjacent to the SA output and are connected with short local metal-4 wires to minimize critical path delay.

Nest arrays

L2 cache

The L2 cache building block for the POWER6 processor is a 32-KB SRAM macro (2 \times 8 ways \times 18 bits). The array runs at one-quarter of the core speed, or half the nest frequency, and is optimized for density. It features 64 cells per LBL, bidirectional GBLs, and fully independent two-dimensional redundancy. It is an eight-way set-associative design that accesses all sets simultaneously for low latency, with late-select multiplexing occurring in the last quarter of the SRAM cycle (fourth core cycle). The late-select signals come from the L2 directory compare circuits. The L2 macro floorplan is shown in **Figure 8** and is similar to that of the slower L3 design [6].

The macro area is 0.26 mm² with an array efficiency (usable cell area/macro area) of 74%. There are no vertical circuit columns, resulting in a narrow design that allows for more flexible placement of repeaters and chip I/O circuits. The row decoders are combined with the column circuitry, and the wordlines are connected with orthogonal wires, thereby eliminating the empty space normally associated with a vertical spine of row circuits. The memory cells are grouped into 16 vertical sections (or SAs), each containing 288 active columns of cells and 4 spare columns. The skip-over-column redundancy uses series pass-gate muxes to avoid the defective elements; this implementation allows two random repair actions. There is a smaller seventeenth SA that contains eight spare wordlines that are used for four random row repair actions, bringing the two-dimensional total to six. Each SA has a local area that includes the bit-select, local word decode, bitline precharge, local write, and local bit decode circuits. All other circuits are located in the short central I/O block, forming an efficient layout.

Within an SA, there is a single synchronization signal (SS), and the various local array support circuits are timed from a single common bus, in contrast to having individual control wires for each of the SA circuit functions. This provides better synchronization of the row and column functions, which in turn improves performance. In addition, an area savings is realized by eliminating extra logic devices utilized for individual sinking of the various decoding, precharging, and write

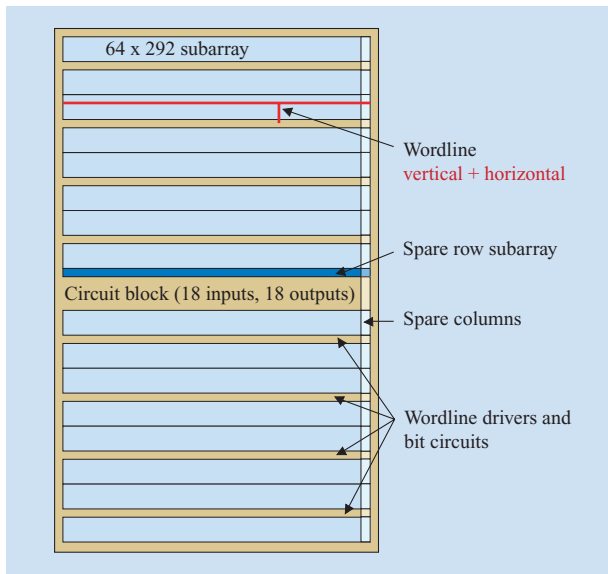


Figure 8

Level 2 cache macro floorplan.

control activation nodes. There are 16 SS signals, one for each of the 16 SAs. The local word decode is a two-input decoder with one input coming from one of the 64 global word decoders and the second input coming from one of the 16 SS decoders, both of which are generated in the central I/O block and wired on a combination of metal-4 and metal-6 levels.

L3 directory

The L3 directory macro (4×2 ways $\times 43$) is a widened derivative of the L2 cache design, employing the same 64-cell-per-bitline SA design and redundancy scheme. Unlike with the cache, there is no late-select mux; only an additional 2:1 column decode in the central I/O block provides 86 direct outputs to the directory compare circuitry.

L2 directory

The L2 directory macro (2×48) is a full-speed nest array, operating at one-half of the core frequency. To achieve this level of performance, the LBLs are shortened to 33 cells, and the separate SA with fully re-locatable spares is replaced by a 32-out-of-33 skip-over-row redundancy scheme.

Summary

The 65-nm POWER6 microprocessor arrays represent a significant paradigm shift relative to the earlier POWER5 processor designs. Key elements that allow the POWER6 processor arrays to achieve 2X to 3X frequency gains

include features such as thin memory cell layout, large signal read (without sense amplification), segmented bitline structure, unclamped column-half-select scheme, multidimensional programmable timing control, and separate elevated SRAM power supply. Under favorable conditions, the core arrays have been fully functional up to 7 GHz. By separately optimizing array cells, circuits, and layouts for the core and nest arrays, very-high performance and high-density arrays are realized on the same chip.

Acknowledgments

The authors would like to acknowledge the dedication of the many engineers who helped to design and test these arrays, including the following: Paul Bunce, John Davis, James Dawson, Bill Huott, Rajiv Joshi, Tom Knips, Mike Lee, Kavita Nair, Pradip Patel, Tony Pelella, Ken Reyer, Dan Rodko, Rolf Sautter, Uma Srinivasan, Art Tuminaro, Shei-ei Wang, and Tobias Werner.

*Trademark, service mark, or registered trademark of International Business Machines Corporation in the United States, other countries, or both.

References

1. K. J. Kim, J. M. Youn, S. B. Kim, J. H. Kim, S. H. Hwang, K. T. Kim, and Y. S. Shin, "A Novel $6.4 \mu\text{m}^2$ Full-CMOS SRAM Cell with Aspect Ratio of 0.63 in a High-Performance $0.25 \mu\text{m}$ -Generation CMOS Technology," *1998 Symposium on VLSI Technology, Digest of Technical Papers*, June 9–11, 1998, pp. 68–69.
2. S. Nakai, M. Kojima, N. Misawa, M. Miyajima, S. Asai, S. Inagaki, Y. Iba, et al., "A 65 nm CMOS Technology with a High-Performance and Low-Leakage Transistor, a $0.55 \mu\text{m}^2$ 6T-SRAM Cell and Robust Hybrid-ULK/Cu Interconnects for Mobile Multimedia Applications," *IEEE International Electron Devices Meeting, Technical Digest*, December 8–10, 2003, pp. 285–288.
3. B. W. Curran, Y. H. Chan, P. T. Wu, P. J. Camprese, G. A. Northrop, R. F. Hatch, L. B. Lacey, J. P. Eckhardt, D. T. Hui, and H. H. Smith, "IBM eServer z900 High-Frequency Microprocessor Technology, Circuits, and Design Methodology," *IBM J. Res. & Dev.* **46**, No. 4/5, pp. 631–644 (2002).
4. A. R. Pelella, A. D. Tuminaro, R. T. Freese, and Y. H. Chan, "A 8Kb Domino Read SRAM with Hit Logic and Parity Checkers," *Proceedings of the 31st European Solid-State Circuits Conference*, September 12–16, 2005, pp. 359–362.
5. J. Davis, D. Plass, P. Bunce, Y. Chan, A. Pelella, R. Joshi, A. Chen, et al., "A 5.6GHz 64kB Dual-Read Data Cache for the POWER6™ Processor," *Solid-States Circuits, 2006 IEEE International Conference, Digest of Technical Papers*, San Francisco, CA, February 5–9, 2006, pp. 622–623.
6. J. Wu, D. Weiss, C. Morganti, and M. Dreesen, "The Asynchronous 24MB On-Chip Level-3 Cache for a Dual-Core Itanium®-Family Processor," *International Solid-States Conference, Digest of Technical Papers*, February 6–10, 2005, pp. 488–489.

Received January 19, 2007; accepted for publication July 14, 2007; Internet publication October 12, 2007

Donald W. Plass *IBM Systems and Technology Group, 2455 South Road, Poughkeepsie, New York 12601 (plassd@us.ibm.com).* Mr. Plass is a Distinguished Engineer in the IBM Systems and Technology Group. He has been responsible for SRAM technology and designs for several generations of IBM server designs, including iSeries*, pSeries*, and zSeries* microprocessors, with a focus on the larger arrays. He joined IBM in 1978 at the Poughkeepsie facility, and in addition to CMOS SRAM, his research and development interests have included DRAM, gallium arsenide (GaAs), and BiCMOS. He received Outstanding Technical Achievement, Outstanding Innovation, and Corporate Awards for his contributions to POWER4, POWER5, and S/390* microprocessors. Mr. Plass has achieved his ninth IBM Invention Achievement Award plateau.

Yuen H. Chan *IBM Systems and Technology Group, 2455 South Road, Poughkeepsie, New York 12601 (chany@us.ibm.com).* Mr. Chan is a Senior Technical Staff Member at the IBM Poughkeepsie Development Laboratory, working on custom very-large-scale integration (VLSI) circuit and SRAM designs. He joined IBM at the East Fishkill facility in 1977 and has worked on high-performance bipolar, BiCMOS, and CMOS logic and array development. He received a B.S. degree in electrical engineering from Union College in 1977 and an M.S.E.E. degree from Syracuse University in 1984. He received an IBM Corporate Award for S/390 CMOS high-frequency techniques and several IBM Outstanding Technical Achievement and Outstanding Innovation Awards for high-performance array development. He was recognized with the title Master Inventor in the IBM Poughkeepsie site. He has reached the fourteenth IBM Invention Achievement plateau, and he holds more than 30 U.S. circuit patents. Mr. Chan is currently a Technical Team Leader of the IBM Server Group array design teams, supporting both the zSeries and the pSeries microprocessor applications.