A New Memory System Design for Commercial and Technical Computing Products

This new design is targeted for use in a wide range of HP commercial servers and technical workstations. It offers improved customer application performance through improvements in capacity, bandwidth, latency, performance scalability, reliability, and availability. Two keys to the improved performance are system-level parallelism and memory interleaving.

by Thomas R. Hotchkiss, Norman D. Marschke, and Richard M. McClosky

Initially used in HP 9000 K-class midrange commercial servers and J-class high-end technical workstations, the J/K-class memory system is a new design targeted for use in a wide range of HP's commercial and technical computing products, and is expected to migrate to lower-cost systems over time. At the inception of the memory design project, there were two major objectives or themes that needed to be addressed. First, we focused on providing maximum value to our customers, and second, we needed to maximize HP's return on the development investment.

The primary customer value proposition of the J/K-class memory system is to maximize application performance over a range of important cost points. After intensive studies of our existing computing platforms, we determined that memory capacity, memory bandwidth, memory latency, and system-level parallelism were key parameters for improving customer application performance. A major leap in memory bandwidth was achieved through system-level parallelism and memory interleaving, which were designed into the Runway bus and the memory subsystem. A system block diagram of an HP 9000 K-class server is shown in Fig. 1 in *Article 1*. The Runway bus (*Article 2*) is the "information superhighway" that connects the CPUs, memory, and I/O systems. System-level parallelism and memory interleaving means that multiple independent memory accesses can be issued and processed simultaneously. This means that a CPU's access to memory is not delayed while an I/O device is using memory. In a Runway-based system with the J/K-class memory system, multiple CPUs and I/O devices can all be accessing memory in parallel. In contrast, many of HP's earlier computing platforms can process only one memory transaction at a time.

Another important customer value proposition is investment protection through performance scalability. Performance scalability is offered in two dimensions: symmetric multiprocessing and processor technology upgrades to the forthcoming PA 8000 CPU. The J/K-class memory system provides the memory capacity and bandwidth needed for effective performance scaling in four-way multiprocessing systems. Initially, Runway-based systems will be offered with the PA-7200 CPU (see *Article 3*), and will be upgradable to PA 8000 CPU technology with a simple CPU module exchange. The J/K-class memory system will meet the demanding performance requirements of the PA 8000 CPU.

Performance is only one part of overall system value. Another major component of system value is cost. For example, the use of commodity DRAM technology was imperative because competitive memory pricing is an absolute requirement in the cost-sensitive workstation marketplace. The J/K-class memory system provides lasting performance with commodity memory pricing and industry-leading price/performance. Low cost was achieved by using mature IC processes, commodity DRAM technology, and low-cost chip packaging. A closely coupled system design approach was combined with a structured-custom chip design methodology that allowed the design teams to focus custom design efforts in the areas that provided the highest performance gains without driving up system cost. For example, the system PC boards, DRAM memory modules, and custom chip I/O circuits were designed and optimized together as a single highly tuned system to achieve aggressive memory timing with mature, low-cost IC process and chip packaging technologies.

A further customer value important in HP computing products is reliability and availability. The J/K-class memory system delivers high reliability and availability with HP proprietary error detection and correction algorithms. Single-bit error detection and correction and double-bit error detection are implemented, of course; these are fairly standard features in modern, high-performance computer systems. The J/K-class memory system provides additional availability by detecting single DRAM part failures for ×4 and ×8 DRAM topologies, and by detecting addressing errors. The DRAM part failure detection is particularly important because single-part failures are more common than double-bit errors. Extensive circuit simulation and margin and reliability testing ensure a high-quality electrical design that minimizes the occurrence of errors.

Fig. 1. *Entry-level memory system.*

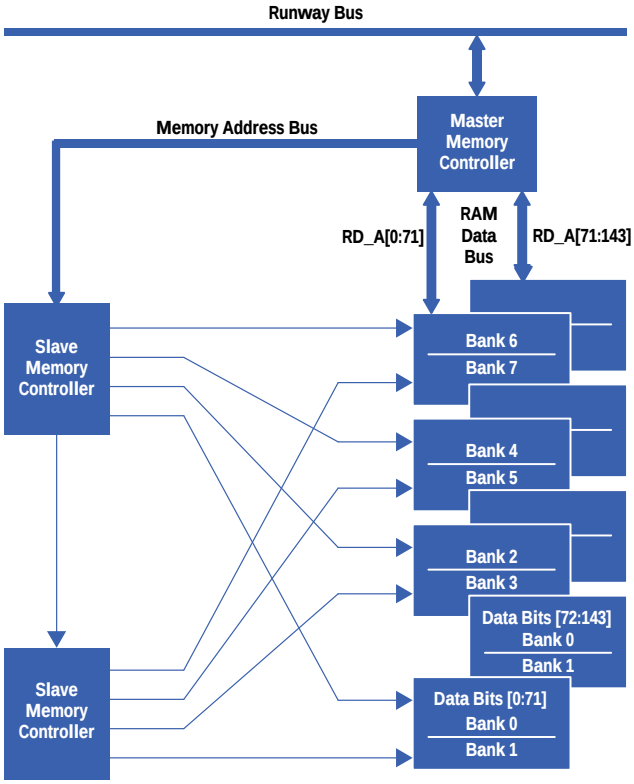

Finally, greater system reliability is achieved through complete memory testing at system boot time. Given the large maximum memory capacity, memory test time was a major concern. When full memory testing takes a long time, customers may be inclined to disable complete memory testing to speed up the boot process. By using custom firmware test routines that capitalize on system-level parallelism and the high bandwidth capabilities of the memory system, a full 2G bytes of memory* can be tested in less than five minutes.

Return on Investment

Large-scale design projects like the J/K-class memory system typically have long development cycles and require large R&D investments. To maximize the business return on large projects, designs need to provide lasting value and cover a wide range of products. Return on investment can be increased by improving productivity and reducing time to market. Leveraging and outsourcing should be used as appropriate to keep HP engineers focused on those portions of the design that provide maximum business value. The J/K-class memory system project achieved all of these important objectives.

A modular architecture was designed so that different memory subsystem implementations can be constructed using a set of building blocks with simple, DRAM-technology-independent interfaces. This flexible architecture allows the memory system to be used in a wide range of products, each with different price and performance points. Given the long development cycles associated with large VLSI design projects, changing market conditions often require VLSI chips to be used in products that weren't specified during the design cycle. The flexible, modular architecture of the J/K-class memory system increases the design's ability to succeed in meeting unforeseen market requirements. Simple interfaces between modules allow components of the design to be leveraged easily into other memory design projects. Thus, return on investment is maximized through flexibility and leverage potential.

Reducing complexity is one of the most powerful techniques for improving time to market, especially with geographically diverse design teams and concurrent engineering. A single critical bug discovered late in the development cycle can easily add a month or more to the schedule. In the J/K-class memory project, design complexity was significantly reduced by focusing on the business value of proposed features. The basic philosophy employed was to include features that provide 80% to 90% of the customer benefits for 20% of the effort; this is the same philosophy that drove the development of RISC computer architectures.

* HP Journal memory size convention:

1 kbyte = 1,000 bytes	1K bytes = 1,024 bytes
1 Mbyte = 1,000,000 bytes	1M bytes = 1,048,576 bytes
1 Gbyte = 1,000,000,000 bytes	1G bytes = 1,073,741,824 bytes

The dedication to reduced complexity coupled with a strong commitment to design verification produced excellent results. After the initial design release, only a single functional bug was fixed in three unique chips.

Several methods were employed to increase productivity without sacrificing performance. First, the memory system architecture uses a "double-wide, half-speed" approach. Most of the memory system runs at half the frequency of the high-speed Runway bus, but the data paths are twice as wide so that full Runway bandwidth is provided.

This approach, coupled with a structured-custom chip design methodology, made it possible to use highly automated design processes and a mature IC process. Custom design techniques were limited to targeted portions of the design that provided the greatest benefits. Existing low-cost packaging technologies were used and significant portions of the design were outsourced to third-party partners. Using all these techniques, high performance, low cost, and high productivity were achieved in the J/K-class memory system design project.

Wide Range of Implementations

The memory system design for the J/K-class family of computers covers a wide range of memory sizes from 32M bytes in the entry-level workstation (Fig. 1) to 2G bytes in the fully configured server (Fig. 2). Expandability is achieved with plug-in dual-inline memory modules that each carry 36 4M-bit or 16M-bit DRAMs. (Note: Even though the memory modules are dual-inline and can be called DIMMs, we usually refer to them as SIMMs, or single-inline memory modules, because this terminology is so commonly used and familiar.) Because each DRAM data bus in the memory system is 16 bytes wide, the 8-byte wide SIMMs are always installed in pairs. Using the 4M-bit DRAMs, each pair of SIMMs provides 32M bytes of memory; and with 16M-bit DRAMs, a pair of SIMMs provides 128M bytes of memory.

Fig. 2. *High-performance HP 9000 Model K400 memory system.*

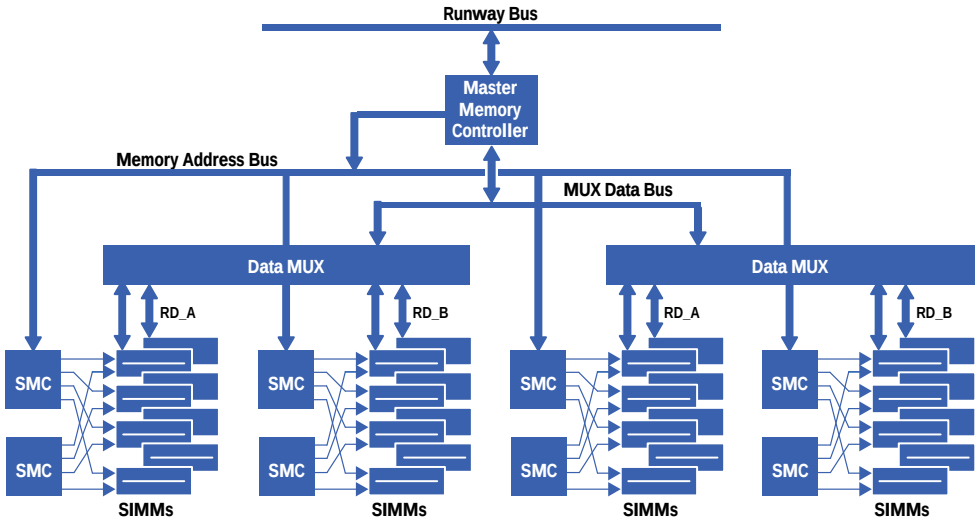

In an entry-level workstation, the memory can start at 32M bytes (one pair of SIMMs with 4M-bit DRAMs) and be expanded up to 512M bytes (using 4 pairs of SIMMs with 16M-bit DRAMs) as shown in Fig. 1. The HP 9000 J-class workstation can be expanded to 1G bytes of memory using eight pairs of SIMMs with 16M-bit DRAMs. The HP 9000 Model K400 server can be expanded up to 2G bytes using 16 pairs of SIMMs with 16M-bit DRAMs installed in two memory carriers.

Design Features

The J/K-class memory system design consists of a set of components or building blocks that can be used to construct a variety of high-performance memory systems. A primary goal of the memory system is to enable system designers to build high-bandwidth, low-latency, low-cost memory systems. Major features of the design are:

- High performance
- 36-bit real address (32G bytes)
- Support for 4M-bit, 16M-bit, and 64M-bit DRAM technology
- Proven electrical design up to 2G bytes of memory with 16M-bit DRAMs (up to 8G bytes with 64M-bit DRAMs when available)
- Logical design supports up to 8G bytes with 16M-bit DRAMs and up to 32G bytes with 64M-bit DRAMs
- Minimum memory increment of 32M bytes with 1M×4-bit (4M-bit) DRAMs (2 banks of memory on 2 SIMMs)
- 32-byte cache lines
- Memory interleaving: 4-way per slave memory controller, electrically proven up to 32-way with logical maximum of 128-way interleaving

- Single-bit error correction, double-bit error detection, and single-DRAM device failure detected for ×4 and ×8 parts
- Address error detection
- Error detection, containment, and reporting when corrupt data is received
- Memory test and initialization less than 5 minutes for 2G bytes of memory
- Soft-error memory scrubbing and page deallocation implemented with software
- 16-byte and 32-byte write access to memory
- IEEE 1149.1 boundary scan in all VLSI parts.

Memory System Description

A block diagram for a high-performance HP 9000 Model K400 memory system is shown in Fig. 2. The memory system has four major components: the master memory controller (MMC), multiple slave memory controllers (SMC), a data accumulator/multiplexer (DM), and plug-in memory modules (SIMMs). The memory system design allows many possible configurations, and the high-performance Model K400 system is an example of a large configuration.

The basic unit of memory is called a bank. Each bank of memory is 16 data bytes wide and can be addressed independently of all other banks. A 32-byte processor cache line is read or written in two segments using fast-page-mode accesses to a bank. Two 16-byte segments are transferred to or from a bank to make up one 32-byte cache line.

Each slave memory controller (SMC) supports up to four independent memory banks. Memory performance is highly dependent on the number of banks, so the SIMMs are designed so that each SIMM contains eight bytes of two banks. Since a bank is 16 bytes wide, the minimum memory increment is two SIMMs, which yields two complete banks. An additional 16 bits of error correction code (ECC) is included for each 16 bytes (128 bits) of data. Thus, a memory data bus carrying 16 bytes of data requires 144 bits total (128 data bits + 16 ECC bits).

The 16-byte-wide memory data bus, which connects the master memory controller (MMC) to the data multiplexer (DM) chip set, operates at 60 MHz for a peak bandwidth of 960 Mbytes/s. Memory banks on the SIMM sets are connected to the DM chip set via 16-byte RAM data buses (RD_A and RD_B), which operate at 30 MHz, yielding a peak bandwidth of 480 Mbytes/s. However, these data buses are independent, so if memory access operations map to alternate buses, the peak bandwidth available from RD_A and RD_B equals that of the memory data bus. The actual bandwidth will depend on the memory access pattern, which depends on the system workload.

The set of signals connecting the MMC to the SMC chips and DM chips is collectively known as the memory system interconnect (MSI) bus. It is shown in Fig. 2 as the memory address bus and the MUX data bus.

A 32-byte single cache line transfer requires two cycles of data on the RAM data and MSI buses. Since the RAM data bus operates at one-half the frequency of the MSI bus, the data multiplexer chips are responsible for accumulating and distributing data between the MUX data bus and the slower RAM data buses. To reduce the cost of the data MUX chips, the design of the chip set is bit-sliced into four identical chips. Each DM chip handles 36 bits of data and is packaged in a low-cost plastic quad flat pack.

Two sets of DM chips are shown in Fig. 2, and four SMC chips are associated with each DM set. Logically, the MSI protocol supports up to 32 total SMC chips, up to 32 SMC chips per DM set, and any number of DM sets. Presently, memory systems having up to eight SMC chips and two DM sets have been implemented.

VLSI Chips

Master Memory Controller. Each memory system contains a single master memory controller. The MMC chip is the core of the memory system. It communicates with the processors and the I/O subsystem over the Runway bus, and generates basic memory accesses which are sent to the slave memory controllers via the MSI bus.

Slave Memory Controller. A memory system contains one or more SMC chips, which are responsible for controlling the DRAM banks based on the accesses generated by the MMC. The partitioning of functionality between the MMC and SMC has been carefully designed to allow future support of new types of DRAMs without modification to the MMC. The motivation for this is that the MMC is a large complex chip and would require a large engineering effort to redesign or modify. The SMC and DM chips are much simpler and can be redesigned with less effort on a faster schedule. The memory system design is partitioned so that the MMC does not contain any logic or functionality that is specific to a particular type of DRAM. The following logic and functionality is DRAM-specific and is therefore included in the SMCs:

- DRAM timing
- Refresh control
- Interleaving
- Memory and SIMM configuration registers
- DM control.

Each SMC controls up to four banks of memory.

Operation of the DRAMs is controlled by multiple slave memory controllers which receive memory access commands from the system bus through the master memory controller. Commands and addresses are received by all SMCs. A particular SMC responds only if it controls the requested address and subsequently drives the appropriate DRAMs with the usual row address strobe (RAS) and column address strobe (CAS).

The slave memory controller chips have configuration registers to support the following functions:

- Interleaving
- Bank-to-bank switching rates
- Programmable refresh period
- SIMM sizes
- Programmable DRAM timing
- SMC hardware version number (read only)
- SMC status registers for latching MSI parity errors.

Memory refresh is performed by all of the SMCs in a staggered order so that refresh operations are nonsimultaneous. Staggered refresh is used to limit the step load demands on the power supply to minimize the supply noise that would be caused by simultaneously refreshing all DRAMs (up to 1152 in a Model K400 system). This lowers overall system cost by reducing the design requirement for the power supply.

Data Multiplexer Chip Set. The DM chips are responsible for accumulating, multiplexing, and demultiplexing between the 16-byte memory data bus and the two independent 16-byte RAM data buses. They are used only in high-performance memory systems with more than eight banks of memory.

Dual-Inline Memory Modules

The dual-inline memory modules (called SIMMs) used in this design are 72 bits wide (64 data bits + 8 ECC bits) organized into two half-banks as shown in Fig. 3. With 72 bits of shared data lines and two independent sets of address and control lines, they hold twice as much memory and supply twice as many data bits as the common single-inline memory modules used in personal computers. Two SIMMs are used to form two fully independent 16-byte banks of memory. Each SIMM holds 36 4-bit DRAMs—18 DRAMs for each half-bank. Using 36 1M×4-bit DRAMs, each SIMM provides 16M bytes of memory. With 4M×4-bit DRAMs, each SIMM holds 64M bytes. To connect all of the data, address, and control signals, a 144-pin dual-inline socket is used (not compatible with the 72-pin SIMMs used in PCs).

Fig. 3. *Diagram of the dual-inline memory module (called a SIMM because its function is similar to the familiar single-inline memory modules used in PCs).*

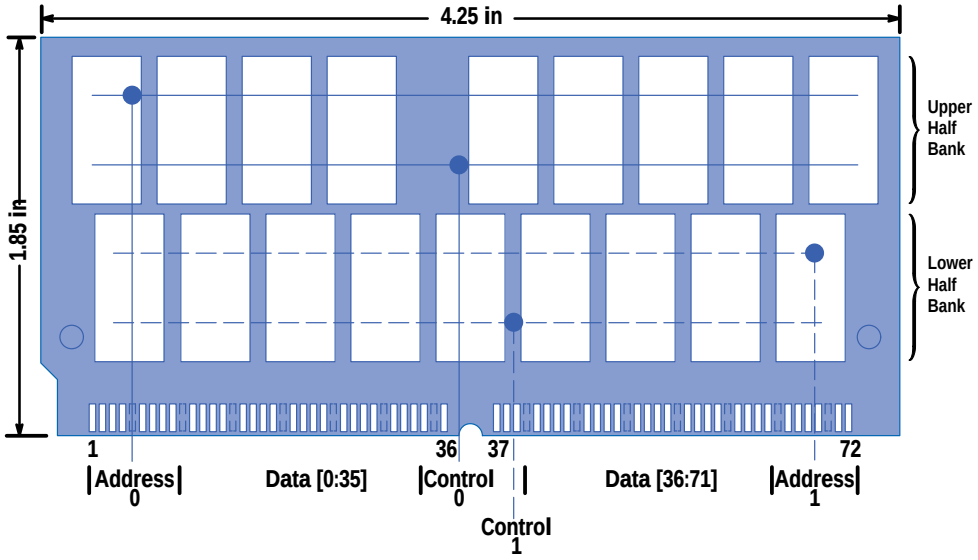

The motivation for designing our own memory module rather than using an industry-standard SIMM was memory capacity and performance. We would have needed twice as many PC SIMMs for an equivalent amount of memory. This would create a physical space problem because twice as many connectors would be needed, and would create performance problems because of the increased printed circuit board trace lengths and unterminated transmission line stubs.

However, our custom SIMM is expected to become “industry available.” There are no custom VLSI chips included on our SIMMs, so third-party suppliers can clone the design and offer memory to HP customers. This is very important because

having multiple suppliers selling memory to our customers ensures a market-driven price for memory rather than proprietary pricing. HP customers now demand competitive memory pricing, so this was a "must" requirement for the design.

Banks and Interleaving

Each SMC has four independent bank controllers that perform a double read or write operation to match the 16-byte width of the RAM data path to the 32-byte size of the CPU cache line. Thus each memory access operation to a particular bank is a RAS-CAS-CAS* sequence for reading two successive memory locations, using the fast-page-mode capability of the DRAMs. A similar sequence to another bank can overlap in time so that the -CAS-CAS portion of the second bank can follow immediately after the RAS-CAS-CAS of the initial bank. This is interleaving.

N-way interleaving is implemented, where N is a power of 2. The total number of banks in an interleave group is not necessarily a power of 2. When the number of banks is a power of 2, then the bank select for a given physical address is determined by the N low-order bits of the physical address. All banks within an interleave group must be the same size. Memory banks from different-size SIMMs can be installed in the same memory subsystem, but they must be included in different interleave groups.

When the number of banks installed is not a power of 2, the interleaving algorithm is specially designed to provide a uniform, nearest-power-of-2 interleaving across the entire address range. For example, if you install six banks of the same size, you will get 4-way interleaving across all six banks rather than 4-way interleaving across 4/6ths of the memory and 2-way interleaving across 2/6ths of the memory. This special feature prevents erratic behavior when nonpower-of-2 numbers of banks are installed.

Soft Errors and Memory Scrubbing

DRAM devices are known to lose data periodically at single-bit locations because of alpha particle hits. The rate of occurrence of these soft errors is expected to be one every 1 million hours of operation. A fully configured 2G-byte memory system uses 1152 (32×36) DRAM devices. Thus, a soft single-bit error is expected to occur once every 868 hours or ten times per year in normal operation. Single-bit errors are easily corrected by the MMC when the data is read from the memory using the ECC bits. Single-bit errors are corrected on the fly with no performance degradation. At memory locations that are seldom accessed, the occurrence of an uncorrectable double-bit error is a real threat to proper system operation. To mitigate this potential problem,

memory-read operations are periodically performed on all memory locations to find and correct any single-bit errors before they become double-bit errors. This memory scrubbing software operation occurs in the background with virtually no impact on system performance.

Sometimes when a particular DRAM device has a propensity for soft errors or develops a hard (uncorrectable) error, then that area of memory is deemed unusable. The memory is segmented into 64K-byte pages. When any location within a particular page is deemed unusable, then that entire page is deallocated from the inventory of available memory and not used. Should the number of deallocated pages become excessive, the respective SIMM modules are deemed faulty and must be replaced at a convenient service time.

Memory Carrier Board

The memory carrier board (Fig. 4) is designed to function with HP 9000 K-class and J-class computer systems. The larger K-class system can use up to two memory carrier boards, while the smaller J-class system contains only one board, which is built into the system board. Each memory carrier board controls from two to sixteen SIMMs. This allows each memory carrier board to contain from 32M bytes to 1G bytes of memory.

There are four data multiplexer chips on the memory carrier board. These multiplex the two 144-bit RAM data buses (RD_A and RD_B) to the single 144-bit MSI data path to the MMC chip. They also provide data path timing. Four SMC components on the memory carrier board provide the MSI interface control, DRAM control and timing, data MUX control, refresh timing and control, and bank mapping.

The memory carrier board is designed with maximal interleaved memory access in mind. Each SMC controls four SIMM pairs (actually only four banks because there are two banks on each SIMM pair) and one data bus set (two of four 72-bit parallel RAM data buses). Each data MUX controls 36 bits of each RAM data bus and 36 bits of the MSI bus. Each SIMM has two address buses (one for each bank) and one 72-bit RAM data Bus. For example, SMC 0 controls SIMM pairs 0a/b, 5a/b, 6a/b, and 3a/b, using the RD_A0(0:71) bus for the SIMMs in the "a" connectors and the RD_A1(0:71) bus for the SIMMs in the "b" connectors.

* RAS = Row address strobe.
CAS = Column address strobe.

Fig. 4. *Memory carrier board architecture. Each SIMM has two address buses and one 72-bit data bus.*

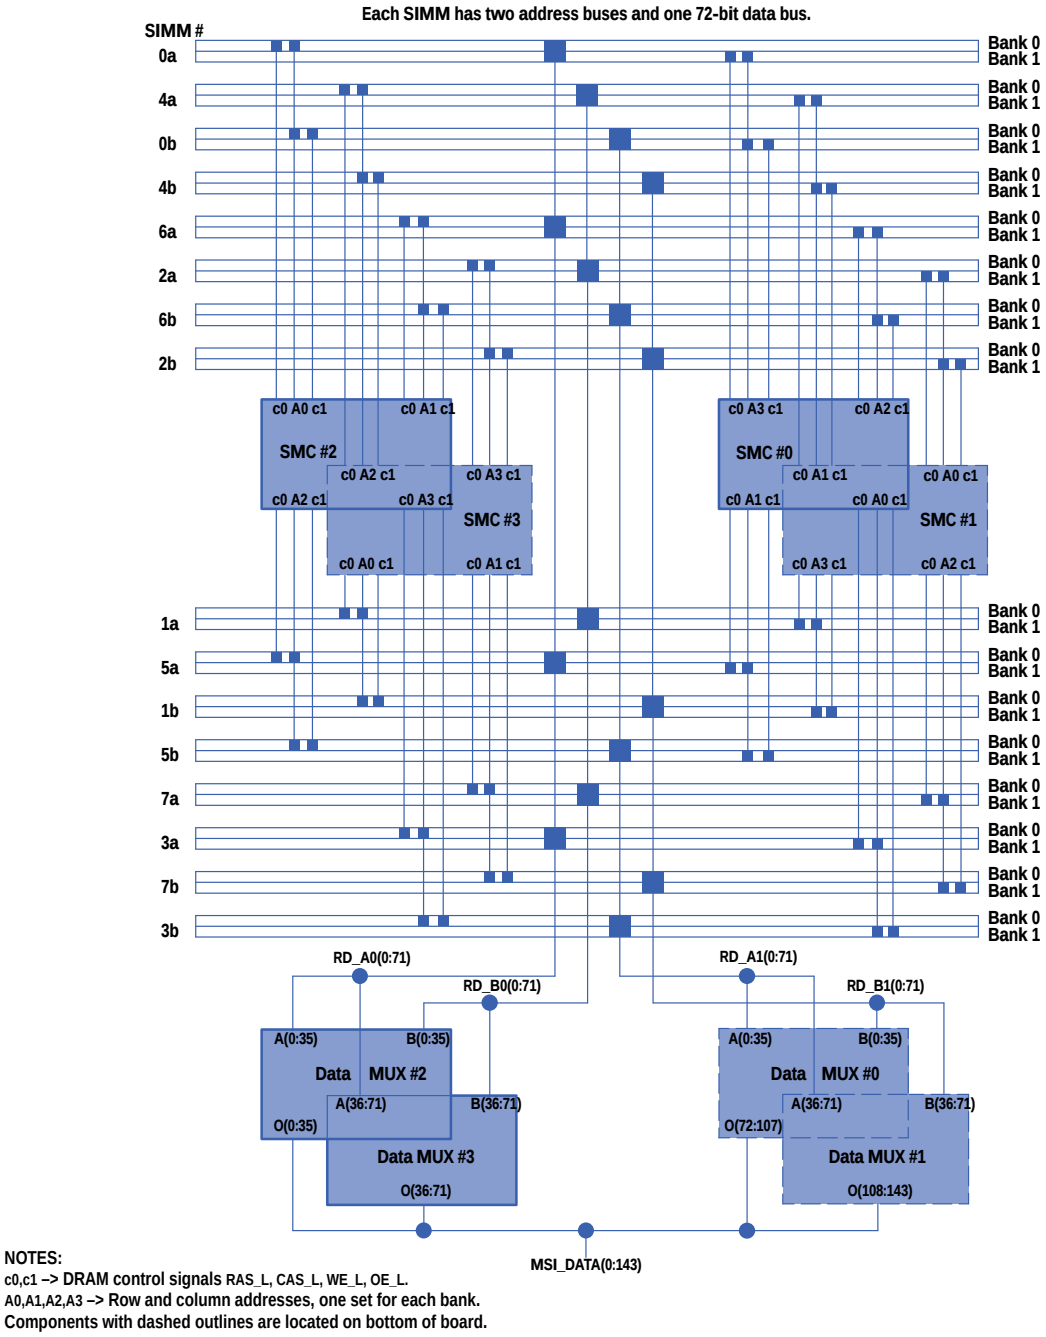

Memory Read or Write Operations

The memory carrier board operates on basically one kind of memory transaction: a 32-byte read or write to a single bank address. This board will also handle a 16-byte read or write operation; the timing and addressing are handled just like a 32-byte operation except that the CAS signal for the unused 16-byte half is not asserted.

To perform a memory read or write operation the following sequence of events occurs. First, an address cycle on the Runway bus requests a memory transaction from the memory subsystem. This part of the operation is taken care of by the MMC. The MMC chip then places this address onto the MSI bus along with the transaction type code (read or write). All the SMCs and the MMC latch this address and transaction code and place this information into their queues. Each SMC and the MMC chip must keep an identical set of transaction requests in their queues; this is how data is synchronized for each memory operation (the SMCs operate in parallel and independently on separate SIMMs).

Once the address is loaded into their queues, the SMCs check to see if the address matches one of the banks that they control. The SMC that has a match will then start an access to the matching bank if that bank is not already in use. The other SMCs could also start other memory accesses in parallel provided there are no conflicts with banks currently in use.

The memory access starts with driving the row address to the SIMMs followed by the assertion of RAS. The SMC then drives the column address to the same SIMMs. This is followed by the assertion of CAS and output enable or write enable provided that the data buses are free to be used. At this time the SMC sends the MREAD signal or the MWRITE signal to the data MUX chips to tell them which direction the data will be traveling. The TACK (transaction acknowledged) signal is toggled by the SMC to tell the MMC chip to supply data if this is a write operation or receive data if this is a read operation. When TACK changes state all of the other SMCs and the MMC step up their queues because this access will soon be completed.

Once the MMC chip supplies the write data to the data MUXs or receives the data from the data MUXs on a read operation, it completes the transaction on the Runway bus. The memory system can have up to eight memory transactions in progress at one time and some of them can be overlapped or paralleled at the same time.

The timing for a single system memory access (idle state) in a J/K-class system breaks down as follows (measuring from the beginning of the address cycle on the Runway bus to the beginning of the first data cycle and measuring in Runway cycles):

Cycles	Operations during these Cycles
1	Address cycle on Runway bus
2	Address received at MMC to address driven on MSI bus (actually 1.5 or 2.5 cycles, each with 50% probability)
2	Address on MSI bus
2	SMC address in to RAS driven to DRAMs
6	RAS driven to output enable driven
4	Output enable driven to first data valid at data MUX input
4	First data valid at data MUX input to data driven by data MUX on MSI bus. With EDO (extended data out) DRAMS this time is reduced to 2 cycles.
2	Data on MSI bus
2.5	Data valid at MMC input to data driven on Runway bus (includes ECC, synchronization, etc.)
25.5	Total cycles delay from address on Runway bus to data on Runway bus for a read operation.

Register accesses to SMC chips are very similar to memory accesses except that the register data values are transferred on the MSI address bus instead of the MSI data bus.

Board Design Challenges

Early in the board design it became clear that because of the number of SIMMs and the physical space allocated to the memory carrier board, the design would not work without some clever board layout and VLSI pinout changes. After several different board configurations (physical shape, SIMM placements, through-hole or surface mount SIMM connectors, SMCs and data MUX placements) were evaluated, the final configuration of 16 SIMMs, four data multiplexers, and four SMCs on each of two boards was chosen.

Given the very tight component spacing required with this configuration, the pinouts of the data MUX and SMC chips had to be chosen carefully. The pinout of the data MUX chip was chosen so that the RAM data buses from the SIMMs and the MSI data bus to the connector were “river routable” (no trace crossing required). The pinout of the SMC chip was chosen with the layout of the SIMMs in mind. It also had to be possible to mount both chips on the backside of the board and still meet the routing requirements. Being able to choose chip pinouts to suit board layout requirements is one of the many advantages of in-house custom chip designs. Without this ability it is doubtful that this memory carrier board configuration would be possible. This is another example of closely coupled system design.

One of the goals for this design was to have a board that could be customer shippable on first release (no functional or electrical bugs). To meet this goal a lot of effort was placed on simulating the operating environment of the memory subsystem. By doing these simulations, both SPICE and functional (Verilog, etc.), electrical and functional problems were found and solved before board and chip designs were released to be built.

For example, major electrical cross talk problems were avoided through the use of SPICE simulations. In one case, four 72-bit buses ran the length of the board (about 10.5 inches) in parallel. Each trace was 0.005 inch wide and the traces were spaced 0.005 inch apart on the same layer (standard PCB design rules) with only 0.0048 inch of vertical separation between layers. Five-wire mutually coupled memory carrier board and SIMM board models for SPICE were created using HP VLSI design tools. When this model set was simulated, the electrical cross talk was shown to be greater than 60% and would have required a major board redesign to fix when found after board release. The solution was to use a 0.007-inch minimum spacing between certain wires and to use a nonstandard board layer stack construction that places a ground plane between each pair of signal layers.

The memory carrier board uses several unusual technologies. For example, the board construction (see Fig. 5) is designed to reduce interlayer cross talk between the RD_A and RD_B data buses. As a result of this board layering, the characteristic impedance of nominal traces is about 38 ohms for both the inner and outer signal layers. The nominal trace width for both

inner and outer signal layers is 0.005 inch with 0.005-inch spacing on the outer layers and 0.007-inch spacing on the inner layers to reduce coupling between long parallel data signals.

The early SPICE and functional simulation effort paid off. No electrical or functional bugs were found on the memory carrier board, allowing the R&D revision 1 board design to be released as manufacturing release revision A.

Another new technology used on the memory carrier board is the BERG Micropax connector system. The Micropax connector system was selected because of its controlled impedance for signal interconnect and the large number of connections per inch. However, for these very same reasons the connector system requires extremely tight tolerances when machining the edge of the board containing the connectors.

A new manufacturing process used by the memory carrier board, the HP 2RSMT board assembly process, was developed to allow the surface mounting of extra-fine-pitch VLSI components on both sides of the board along with the through-hole SIMM connectors.

Fig. 5. *Printed circuit board construction.*

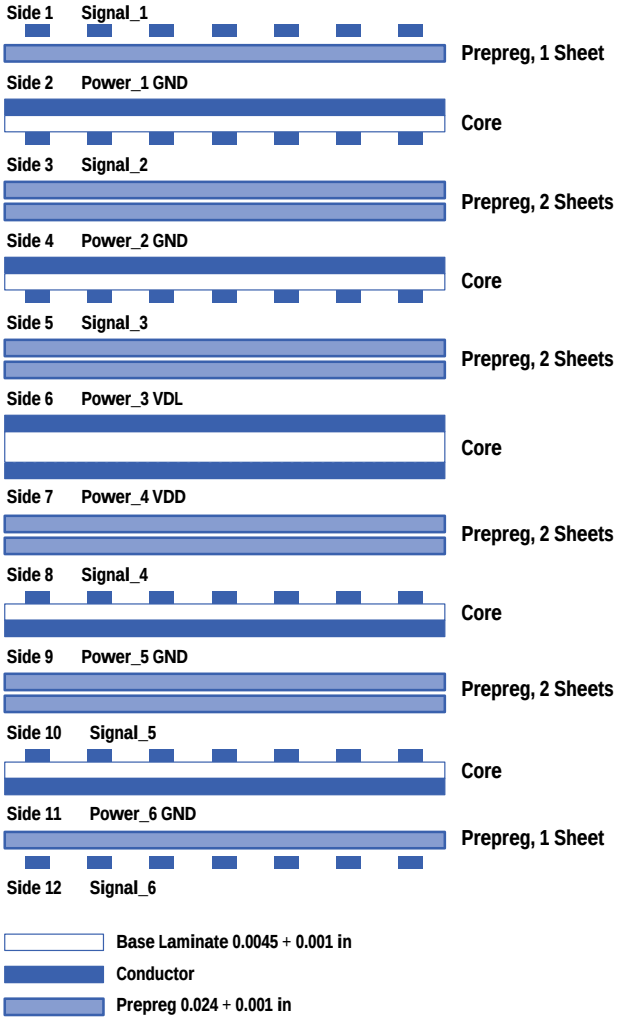

Acknowledgments

Development and manufacture of the J/K-class memory system was carried out at multiple sites within HP. The numerous contributors are thanked for their dedication and hard work in the timely completion of this highly successful memory system. At Chelmsford, Massachusetts: Ruya Atac, Dave Chou, Shawn Clayton, Bill Clouser, Dan Danecki, Tom Dickson, Tom Franco, Craig Frink, Joy Han, Larry Kaylor, Milton Makris, Dale Montrone, Jack Murphy, Mike Ricchetti, Bob Ryan, Dan Schumacher, Jim Williams, John Wood, Steve Zedeck. At Cupertino, California: Tak Watanabe, Ed Jacobs, Elio Toschi, Sudhir Chandratreya, John Youden, Margaret Au. In Puerto Rico: Rodolfo Mazo, Roberto Falcon, Enid Davila, Edgardo Aviles, Joel Vientos. At Guadalajara, Mexico: Adolfo Gascon, Marco Gonzales, Ruben Kleiman, Manuel Rincon. At Corvallis, Oregon: Pete Fricke, Kathleen Kollitz, Trung Vo-Vu. At Roseville, California: Dan Cromwell, Matt Harline, Hani Hassoun, Andy Levitt.

▶ Go to *Article 6*
▶ Go to Table of Contents
▶ Go to HP Journal Home Page