RAMBUS

64-MEGABIT RAMBUS DRAM TECHNOLOGY DIRECTIONS

64M Rambus DRAM Technology Directions

In September of 1995, an unprecedented announcement hit the PC industry. Six leading DRAM suppliers — Hitachi, LG Semicon, NEC, Oki, Samsung and Toshiba — came together to announce their commitment to a new generation of compatible, high-speed memory devices based on the high bandwidth Rambus™ interface. The announcement that each would provide volume shipments of 64Mbit Rambus DRAMs starting in 1997 was directed to PC manufacturers. The announcement affirmed DRAM suppliers' understanding of the performance issues faced by PC designers, convincing them to provide the high speed 533MHz Rambus DRAM (RDRAM®), in sufficient volume, to populate main memory systems of re-architected low cost, multimedia-rich PCs.

Over the past several years, PC OEMs have faced DRAM shortages, limiting their ability to meet the dramatic growth underway in the PC market. The move to high frequency, sophisticated P6 processors and to data-intensive multimedia applications have raised the bandwidth requirements of the memory subsystem. Conventional DRAMs, even if available in wider organizations, would be underpowered to take on these increased bandwidth requirements. PC designers have a choice: to continue implementing conventional, complex memory subsystems, or to simplify their designs and save costs with the streamlined Rambus architecture.

Rambus DRAMs are already well accepted in computer graphics and consumer video system products. Providing a high speed interface to the DRAM core, Rambus DRAMs are able to transfer data at up to ten times faster than conventional DRAMs. Based on a byte-wide narrow bus, the Rambus DRAMs offer several advantages for decreasing system cost, such as a low pin-count interface, finer granularity, low cost memory expansion, compact form factor and lower power consumption. With Rambus technology, PC designers have several options for achieving higher performance, richer multimedia functionality and lower system costs.

This document reviews the PC main memory technology trends that drove the six DRAM suppliers to simultaneously announce their commitment to providing 64M Rambus DRAMs. The first section provides an update on Rambus Inc. and partner programs; the second section provides an overview of the 64M Rambus DRAM; and the last section provides more details on the characteristics of PC memory workloads and the performance of the 64M Rambus DRAM under those workload conditions.

List of Topics

- Rambus Inc. Update
- 64M Rambus DRAM Overview
- Performance of the 64M Rambus DRAM in PC Main Memory

Rambus Inc. Update

Rambus Inc. has developed a revolutionary DRAM architecture and a very high-speed chip-to-chip data transfer technology — the first interface standard able to be directly implemented on CMOS DRAMs, memory controllers, graphics/video ICs and other high performance VLSI components. The Rambus solution achieves a tenfold increase in component throughput, while utilizing fewer ICs and assuring a modular/scalable solution. A single Rambus DRAM (RDRAM) delivers 533 megabytes of data per second over the Rambus Channel — a narrow high-speed bus.

The adoption of Rambus technology dramatically simplifies computer architectures, since hardware traditionally used to increase the speed of the processor to memory interface become unnecessary. Rambus products use standard CMOS processes, low cost IC packaging and conventional PC board technologies in order to take advantage of high-volume, low-cost manufacturing processes. The results of these factors will enable faster, smaller and lower cost systems for personal, portable and multimedia applications.

DRAM Market: High Growth Driven by PC Industry

The DRAM market continues its high rate of growth, largely due to the growth in the PC industry. While Dataquest predicts that 58 million PCs will ship worldwide in 1995, their worldwide forecast for 1997 is now 79 million PCs, and they expect worldwide PC shipments to reach 107 million units during the year 1999. Over 75 percent of the DRAMs shipped are used in personal computer and consumer video products. As a result of the growth in the PC market, DRAM shipments will move upward from over $35B in 1995 to more than $60B in 1999.

DRAM component density has traditionally quadrupled every three years. The 64Mbit generation of DRAMs will begin shipping in volume in 1997. Dataquest's projections show that for memory systems of at least 8Mbytes, the 64Mbit DRAMs will offer best economics during the 1998-2000 timeframe.

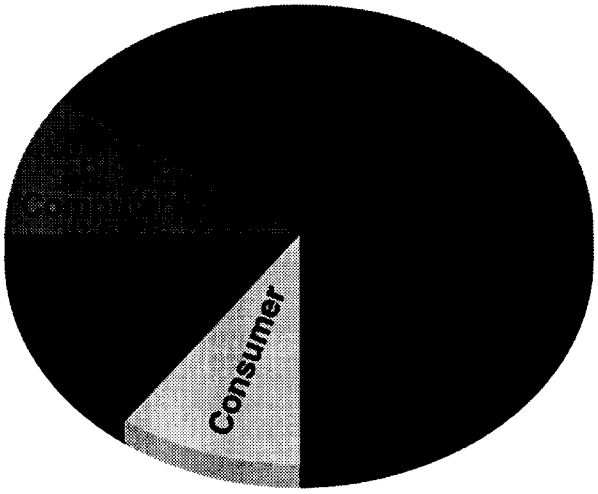

- 1997 DRAM Market Revenue: $44B (Source: Dataquest 5/95)
- 75% of DRAMs shipped to PCs and Consumer products.

Figure 1: 1997 DRAM Market by Application

Approaches for Providing High Memory Bandwidth

As processor performance requirements continue to increase, PC designers are faced with two approaches for providing greater bandwidth from their memory subsystem. The conventional approach is to gang lots of DRAMs in parallel on wide data buses. Conventional memory systems of high performance PCs and graphics/video subsystems use multiple DRAMs in 64-bit wide data paths. The Rambus approach to higher performance is to greatly increase the bandwidth of the interface to the DRAM core. The Rambus interface transfers data at 533 megabytes per second, which provides ten times the throughput over the time-multiplexed, RAS/CAS interface on conventional DRAMs. With the Rambus approach, even an individual RDRAM provides sufficient bandwidth for multimedia subsystems or PC main memory.

Another consideration in memory system design is the number of pins and data wires needed. The low pin-count interface of Rambus DRAMs uses only 31 pins on the memory controller, a savings of over 80 pins compared to the 64-bit wide memory array. The narrow byte-wide Rambus Channel requires only a few square inches of board space, about a tenth of what might be needed with conventional

memories. The Rambus approach uses fewer components, smaller memory controllers, smaller form factor and less power.

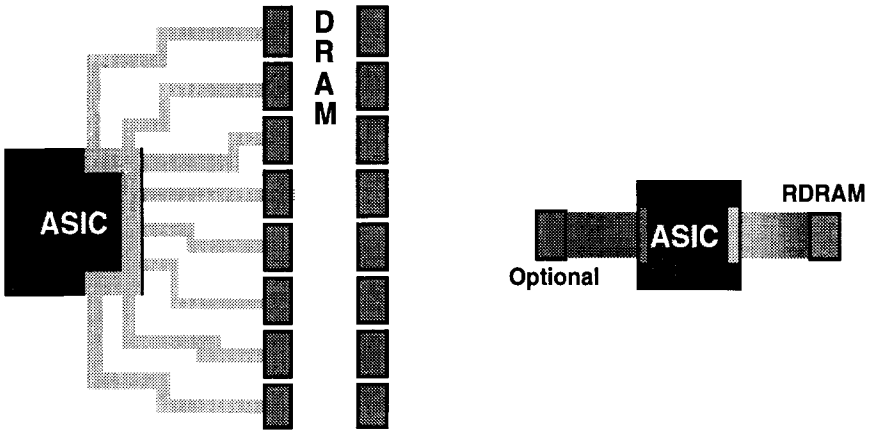

- 533 MB/sec
- 8 bytes @ 66 MHz
- 110 pins to memory
- Big system solution

- 533 MB/sec
- 1 byte @ 533MHz
- 31 pins to memory
- Personal and portable

Figure 2: High Bandwidth Memory Solutions

Rambus System Components Availability

Rambus memory technology is an open standard that Rambus licenses to integrated circuit companies. The IC companies manufacture, market and sell Rambus-based components to system companies. This allows Rambus to provide a superior solution to system designers and users from multiple, high-quality manufacturers. In addition, this approach enables Rambus to establish business relationships with, and deliver technology to, a wide range of partners.

Availability of Rambus System Components	1995	1996
8M, 16M RDRAMs	NEC, Toshiba, Oki, LG Semicon	Samsung
ASICs	NEC, Toshiba, Hitachi, LSI, TSMC	TBA
Multimedia	Cirrus, TBA	TBA
Core Logic		TBA
Sockets	Augat, Molex: Vertical Augat - Horizontal	Molex - Horizontal
Clock Generators	Macronix, ICS, Chrontel, AMCC	

Rambus technology has been broadly licensed to leading DRAM, ASIC and PC peripheral chip set suppliers. NEC Corporation, Toshiba Corporation, Samsung Corporation, Hitachi Ltd., LG Semicon Co., Ltd. and Oki Electric Industry Co., Ltd. are providing 16M generation Rambus DRAMs. Rambus engineers continue to work closely with a wide range of semiconductor, socket, test and CAD companies to better orchestrate the delivery of a complete memory solution to system designers.

In 1996 five suppliers for the 16Mbit Rambus DRAMs will be shipping in volume. Currently, NEC and Toshiba are providing 0.35µ ASICs with Rambus ASIC Cells (RACs) for customer programs. Several companies are developing system core logic chips based on Rambus technology for Pentium main memory and unified memory architectures. In addition to Rambus IC components, all of the system completion components, such as sockets and clock generator chips, are available from multiple sources.

Applications for Rambus Technology

Rambus-based integrated circuit programs are underway for a wide range of system applications such as PC graphics, 3-D workstations, PC main memory, 3-D home video games, TV settop cable boxes and ATM (Asynchronous Transfer Mode) communications. Silicon Graphics, Inc., Cirrus Logic, Inc. and Nintendo have announced systems using Rambus technology. These system products demonstrate the application breadth of Rambus technology, ranging from $30,000 3D graphics workstations to $2000 PCs and $250 consumer games — all of which benefit from the high bandwidth, low pin-count Rambus interface.

- The Silicon Graphics Indigo^{2}™ IMPACT™ 3D graphics workstations use a minimum of six RDRAMs to achieve peak bandwidth of three gigabytes per second. This configuration allows the IMPACT workstations to deliver performance of big computers costing hundreds of thousands of dollars to the $30-40,000 price point. The IMPACT workstations deliver unprecedented textured fill rate for a desktop system. Silicon Graphics expects the IMPACT workstations to generate over a billion dollars in revenue during 1996.
- Cirrus Logic is the leading supplier of graphics controller products to the PC industry. They are developing a family of VisualMedia™ accelerators based on Rambus technology for the graphics framebuffer. Their CL-GD5462 product gives PC users the unique ability to simultaneously display multiple video windows at 30 frames per second along with high-resolution (up to 1600 x 1200) graphics. The low pin-count interface allows Cirrus Logic family of controllers to all use the same 208-pin plastic package. Cirrus Logic forecasts that Rambus-based controllers will penetrate at least 40 percent of the PC graphics market by the end of 1996.

- Nintendo chose Rambus technology for use in their next generation Ultra 64 3D home video game system to make possible unparalleled three-dimensional imagery. With the help of the Rambus technology, video game players will be able to play in the kind of three-dimensional worlds that previously have been available only in high-end graphics workstations. The use of Rambus technology was key to meeting the Nintendo Ultra 64's target price point of under $250.

Memory Standards with Critical Mass

Over the past few years, the requirements for faster and more cost-effective DRAMs have spawned a variety of proposals for new DRAM alternatives. Most of the new alternatives are based on incremental performance improvements or address niche applications. The system manufacturers have set the criteria for determining which DRAM alternatives will be considered mainstream, and thus will enjoy high volume at competitively low prices. The criteria for becoming a mainstream memory technology are:

- Commitments by at least 4 of the top-ten DRAM suppliers
- 100 percent compatibility across the suppliers' DRAM components
- Proven, full-speed systems in production
- System products representing multiple high volume applications, as opposed to addressing a niche application
- Major committed systems companies who are leaders in their market

The three memory technologies that meet these criteria are Page mode/EDO DRAMs, Synchronous DRAMs and Rambus DRAMs. Of these technologies, Rambus DRAMs provide the highest bandwidth. The Rambus approach with its highest bandwidth per pin and low pin-count interface provides the additional benefits of lower pin count memory controllers, simplified memory design and overall lower system costs.

64 Megabit RDRAM Overview

Taking advantage of improvements in silicon processing, processor performance has doubled every year, while memory density has quadrupled every three years. Users need for PC memory storage double every two years.The impact of these trends is that memory subsystems consisting of fewer DRAM components must deliver far greater bandwidth.

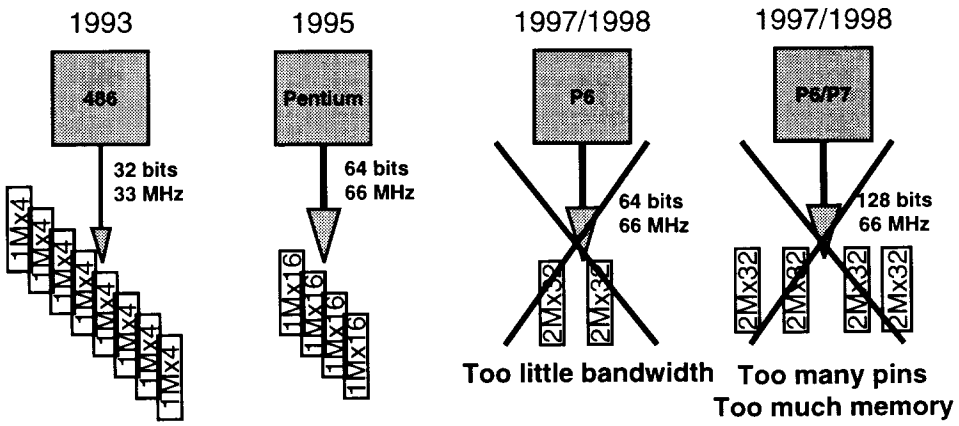

Figure 3: PC Main Memory Trends

The trend toward higher performance from fewer DRAM components dramatically drives up the bandwidth requirements of the individual DRAM component. In response, conventional DRAMs will have to move to wider x32 organizations and to a synchronous DRAM interface. These evolutionary approaches have led to higher component costs, more complex system designs and higher system costs. Despite the incremental improvements, these memory systems are not able to keep up with the PC's exploding bandwidth requirements.

The high speed Rambus interface yields cost-effective, compact memory subsystems that have sufficient bandwidth for PC main memory subsystems. A single Rambus 64Mbit RDRAM provides adequate bandwidth for low-end 8-megabyte subsystems, while two RDRAMs provide over a gigabyte per second of bandwidth for higher performance systems. A Rambus-based memory outperforms EDO DRAM memory systems in multimedia applications. The Rambus-based PCs can be lower cost, as the SRAM cache can be eliminated while still providing high performance. The finer-grain granularity of the Rambus memory allows PC system OEMs greater configuration flexibility.

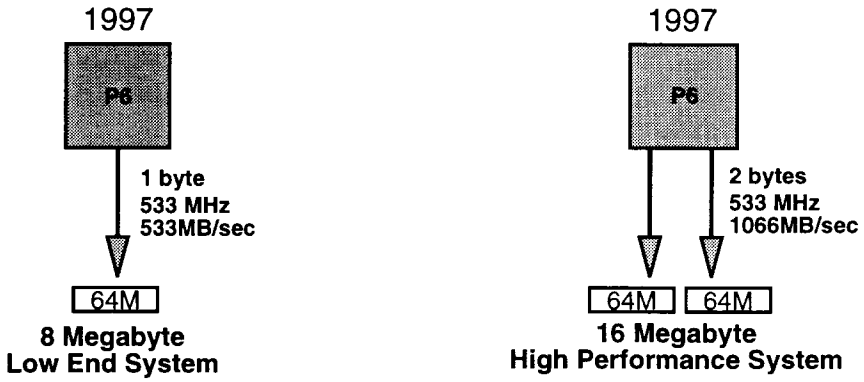

Figure 4: Memory Solutions Using the 64M Rambus DRAM

64M Rambus DRAM Features

The feature set of the 64M generation of Rambus DRAMs is optimized toward PC main memory applications. Rambus engineers developed extensive modeling tools to characterize performance of various memory system configurations running typical PC applications. As a result of this optimization, the 64M RDRAM achieves highest effective bandwidth, low latency and improved power management for PC main memory subsystems.

- **100 percent upward-compatible with the 16Mbit Rambus DRAMs.** Systems can be designed to accept either or both the 16M and 64M RDRAMs. The 64M generation leverages the investment in system component infrastructure, including Rambus Interface Cells (RACs), sockets, clocks and board designs.
- **533 MHz operation.** The Rambus subsystem is easily synchronized with processor buses operating from 50 to 267 MHz. The 533MB/sec bandwidth of a single RDRAM is sufficient bandwidth for entry level PC systems.
- **8-megabyte granularity.** Memory can be expanded by adding a single RDRAM: a 16M RDRAM can be used for two-megabyte granularity, and a 64M RDRAM yields 8-megabyte granularity. The relatively small granularity allows lower cost memory configurations and greater configuration flexibility to the PC OEM.
- **High sustained bandwidth, particularly for main memory application workloads.** The 64M RDRAM uses a four-bank architecture internally, each bank is supported by a row of sense amps, which is referred to as an *open page*. Support for concurrent RAS/CAS operations allows pipelining of data transfers in multiple transaction environments. The result is 425MB/sec sustained bandwidth for 32-byte transfers to random memory locations.

- **Latency improvements at the component level and system level.** At the component level the 64M RDRAM has latency comparable to the EDO DRAM. The 64M RDRAM's latency timings are 26ns for a hit, 60ns row access time in the case of a miss operation and less than 2ns/byte burst time. The component latency can be hidden in multi-transaction processing environments, such as those found in unified or P6 systems. In these environments (which can take advantage of the concurrency between bank operations), the effective latency becomes almost zero.
- **Power management features to address portable applications.** Due to its compact design and high bandwidth, portable system designers find Rambus memory attractive. The 64M RDRAM's power management capabilities allow lower effective power consumption. Only accessed RDRAMs consume power.
- **Cost parity to x32 EDO.** The 64M RDRAM die size is expected to be comparable to the x32 EDO die size.

Sourcing for the 64M Rambus DRAMs

Six suppliers —Hitachi Ltd., LG Semicon Co., Ltd., NEC Corporation, Oki Electric Industry Co., Ltd., Samsung Electronics and Toshiba Corporation — simultaneously announced their support for the Rambus DRAM. These six include the top four suppliers to the DRAM market.

Rambus Partner	1994 Market Share (Source Dataquest 5/95)	Ranking
Samsung	12.7%	1
Hitachi	9.7%	2
NEC	9.2%	3
Toshiba	8.9%	4
LG Semicon	4.8%	8
Oki	2.6%	12

System Design Features

Rambus Inc. provides a totally engineered solution that reduces design time and results in higher performance, lower cost, more compact system products. Using Rambus technology, system architects have several options to re-engineer the PC in order to integrate more features, reduce form factors or address new markets. Following are example system design opportunities for Rambus-based systems:

- A single RDRAM has sufficient performance to meet needs of a 150MHz Pentium or P6-class processor, allowing very compact low-end consumer PC.

- Two RDRAMs can deliver over a gigabyte per second of bandwidth, which is sufficient for a unified memory system based on a Pentium or P6-class processor and display resolution of up to 1280x1024x24 bits per pixel.
- One Rambus memory controller can support one to four Rambus Channels for 533 MB/sec to over 2 GB/sec in bandwidth. A single board design based on this controller footprint is able to address a wide performance range and variety of system models.
- Finer memory granularity for systems using 64Mb DRAM technology or 64-bit wide buses, thus providing system manufacturers with greater configuration flexibility. Systems can support 8MB, 16MB, 24MB etc. of memory storage. The 16M RDRAM allows two megabyte memory expansion.
- Fully engineered solution supports 16M or 64M generation RDRAMs interchangeably across density levels and suppliers.
- Lower system cost accomplished through:
  - Single-chip P5 or P6 memory controllers, due to its low pin-count interface.
  - Low cost memory expansion using SIMM-like memory modules
  - Fewer components (no glue; SRAM cache can be offered as an option)
  - Fewer printed circuit board layers
  - Reduced board space (room to add multimedia features to the motherboard)
  - Reduced EMI

Benefits to PC System Manufacturers

Rambus system design features allow improved price/performance, richer feature sets and more configuration flexibility to system manufacturers. Since Rambus memory subsystems are easily expanded to multiple Rambus Channels, a PC OEM is able to leverage one printed circuit board design across a wide range of system products and price/performance points.

- Rambus technology allows the lowest cost memory solutions for P5/P6/P7 class processors and most flexibility in configuration and board layout.
- Due to Rambus' compact design and low pin-count interface, the OEM is able to integrate more multimedia features (audio, video, compression, 3-D graphics, etc.) onto the PC motherboard.
- Security of supply and competitive pricing for Rambus DRAMs, which are widely sourced by suppliers representing leaders in DRAM manufacturing. RDRAMs are compatible and interchangeable across vendors.
- Much simpler board layout and system packaging as compared to x32 DRAM solutions, resulting in lower system costs, shortened design time and more reliable systems.

Benefits to PC End-Users

Rambus technology's high performance, which is particularly suited toward bandwidth-intensive multimedia and graphics applications, provides a more robust processing environment and richer user interface for personal system users.

- **Consumer and home PCs:** Rambus technology enables games and multimedia applications to run faster on low-cost consumer PCs, while offering full compatibility to run all office applications.
- **Corporate Desktops:** Highest-performance PCs for corporate applications. A Rambus-based memory subsystem in conjunction with an SRAM cache provide highest performance for traditional compute-intensive corporate applications. Rambus technology allows the corporate PC to best take advantage of emerging video-based applications, such as video teleconferencing, video playback from CDs or video streaming over the internet.
- **Consumer Video Games**: Rambus provides the best performance for 3D graphics, complex imagery and sophisticated visualization effects.
- **Settop Boxes:** Rambus provides the best performance and integration of features (audio, video, graphics, MPEG decompression) for the $200 settop box.

Rambus DRAM Product Roadmap

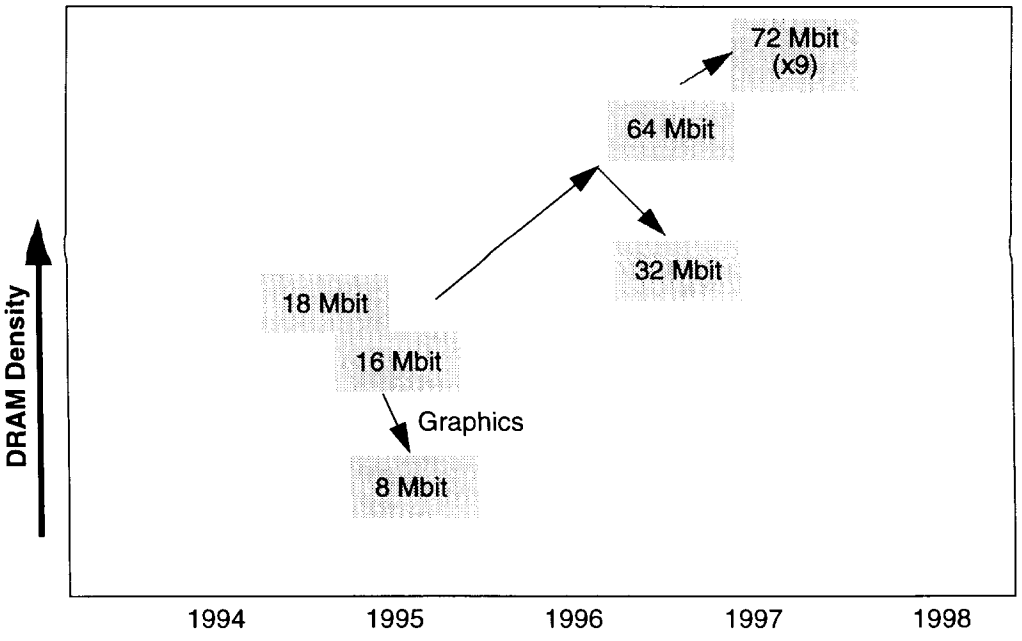

Figure 5: RDRAM Product Roadmap

The Rambus interface technology address multiple generations of DRAM cores. As DRAMs move to higher density levels, the need for high bandwidth per DRAM becomes more compelling — and the proliferation of Rambus DRAMs into PCs more pervasive.

Performance of the 64M RDRAM in PC Main Memory Systems

PC manufacturers will choose the lowest-cost memory that can meet their bandwidth requirements. Two factors are driving up the bandwidth requirements of PCs: CPU processing speeds and user applications. Microprocessor units (MPUs) have been moving to higher internal frequencies, wider internal data paths and sophisticated instruction processing techniques resulting in dramatically increased performance.

P6- and P7-class processors include very sophisticated instruction execution units that can process multiple memory transactions in parallel. Using many of the techniques demonstrated in advanced RISC processors, the P6 processor can issue multiple requests to the memory subsystem. At the same time these processors use 64-bit wide internal data paths that are clocked at up to 200-300MHz. Thus the bandwidth to the memory system is substantial. As an example, a 64-bit bus operating at 100MHz can demand up to 800 megabytes per second in bandwidth from its memory system. The P6 is projected to sustain bandwidth of over 250 megabytes per second.

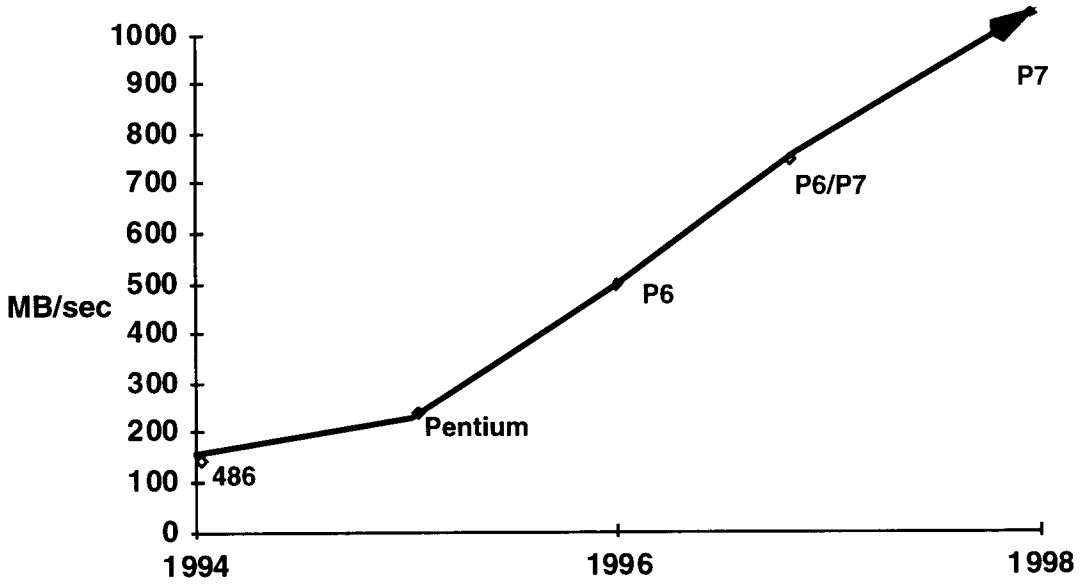

Figure 6: Processor Bandwidth Trends

The second trend is that user applications have moved beyond integer computationally based applications to data streaming multimedia based applications. Applications are taking advantage of high resolution full color displays, 3-D graphics and video based programs. To keep up with the bandwidth requirements, system

designers have moved to large synchronous SRAM caches and wide DRAM memory arrays. These raise the entry cost of the PC, and increase the cost for memory expansion.

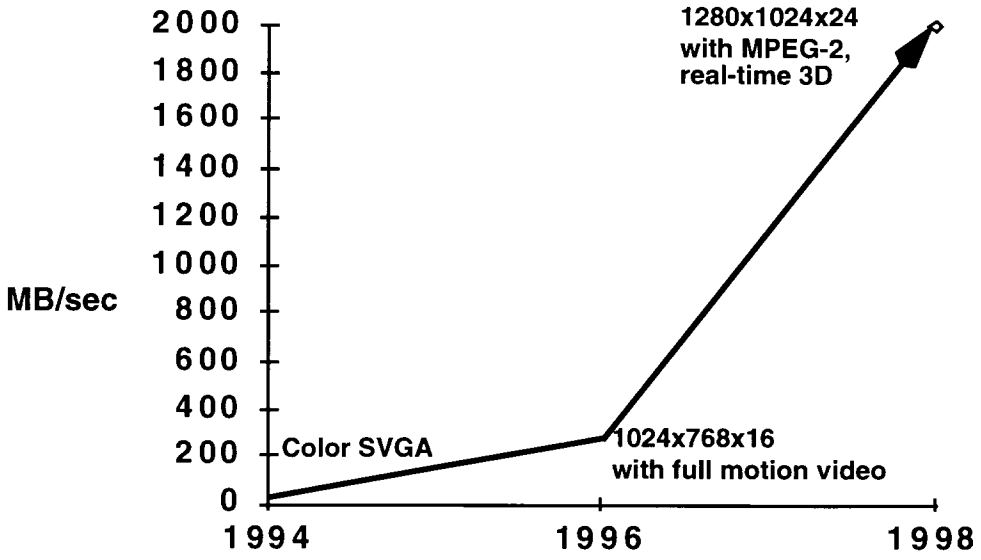

Figure 7: Graphics / Multimedia Subsystem Bandwidth Requirements

Many PC manufacturers are looking at unified memory architectures (UMAs) for providing low-cost PCs. The UMA saves memory costs as it combines the CPU main memory with the graphics frame buffer memory. A UMA system based on a Pentium processor is also characterized as a multiple-requestor environment as the memory system must respond to both processor as well as graphics display requests. In the 1997 timeframe as users move to P6 processors and 1280x1024 x16

bits per pixel displays, the memory subsystem of a UMA must be able to sustain up to 800 megabytes per second transfer rates.

Bandwidth Requirements (MB/sec)	UMA Systems 1995	UMA Systems 1997
Screen Refresh	57 (1024x768x8, 72Hz)	180 (1280x1024x16, 72Hz)
2D Drawing BW	40	80
Video	25	75
3D Graphics	Not Supported	300
CPU Processing	100	200
Total	200 MB/sec	Up to 800 MB/sec

Multiple-Memory-Transaction Processing Systems

The processor-to-memory interface is changing to support multiple-transaction memory systems. High clock rate Pentium-class CPUs running write-intensive multimedia applications, sophisticated RISC processors, next generation P6- and P7-class processors, video games systems and set-top boxes all support high CPU processing rates that can take advantage of a multiple-transaction memory system.

The memory subsystem designer's goal is to provide sustained high bandwidth for the memory transfers demanded by the processor. A memory transaction on Pentium and P6 processor consists of a request for a 32-byte data transfer from the memory subsystem. These requests can have individual locality, but are random relative to each other. The challenge is to deliver lowest latency, which is the time required for the first access, and maximum bandwidth. These are competing design goals.

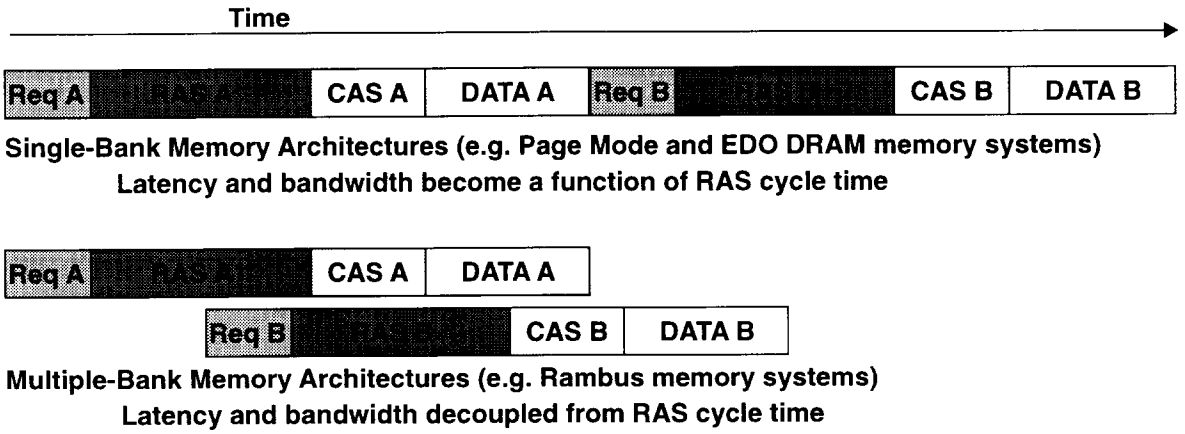

Figure 8: Single Versus Multiple-Bank Memory Architectures

Page mode and EDO DRAMs are based on single-bank designs. In single-bank memory systems, latency depends entirely on the RAS cycle time. The bandwidth depends on the RAS time, precharge time and time to transfer the 32-byte block. The Rambus DRAM has a multiple-bank architecture, in which access to a second data block can take place in parallel with the first data transfer. The latency associated with the second data transfer is thereby hidden.

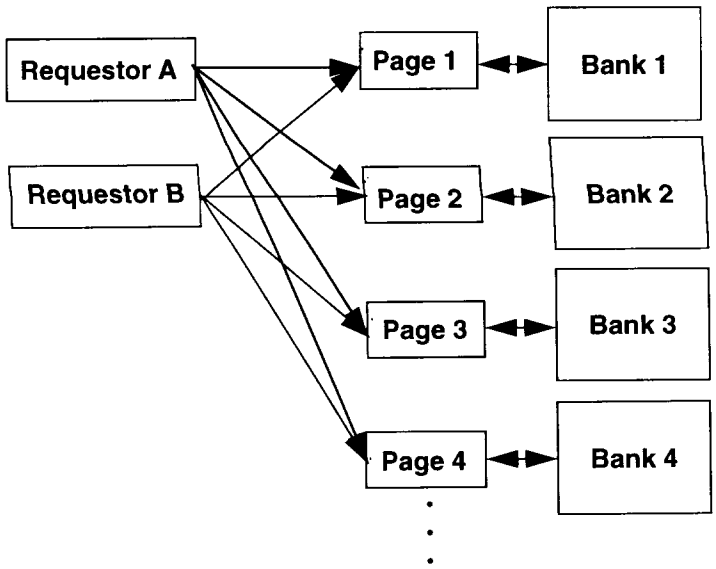

Figure 9: Example of a Multiple Requestor Memory Subsystem

A memory subsystem with multiple banks/pages allows overlapping of transactions to hide the component RAS/CAS latency, resulting in higher effective band-

width. While adding more banks to a DRAM reduces the chances of two requests contending for the same bank, it tends to increase design complexity, die size, and, hence, cost. Rambus' analysis of application workloads for 8- and 16-megabyte memory systems established that a four-bank architecture balances the cost versus performance trade-off.

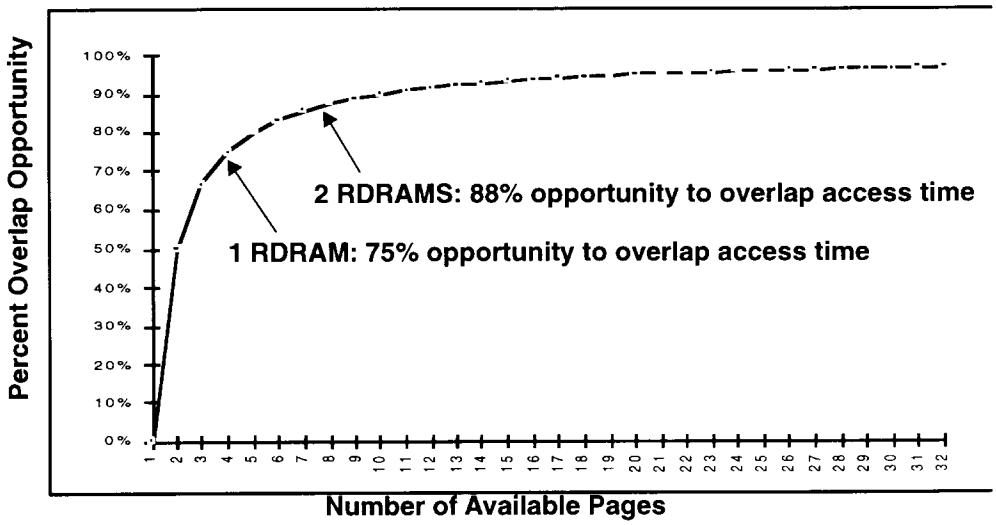

Figure 10: Optimizing the Number of Banks for Minimal Interference and Low DRAM Cost

In processor environments supporting two requestors, even a few pages can greatly decrease the probability of interference. With four banks and open pages, only 25 percent of the random requests to memory will experience bank interference. A 16-megabyte memory system based on two RDRAMs for a total of eight open pages allows an 88 percent opportunity for overlapping accesses. Doubling the number of banks per RDRAM provides an incremental 6-7 percent opportunity for overlap. Thus the four-bank architecture was optimal for providing high sustained bandwidth with minimal silicon overhead.

64M RDRAM Performance

Based on PC bandwidth requirements and applications workload performance modeling, the 64M RDRAM architecture is optimized to provide high effective bandwidth for small, 32-byte block transfers. The 64Mbit Rambus DRAM supports four banks with four open pages. Internally the RDRAM supports simultaneous RAS/CAS to interleave transactions. The four bank architecture increases opportunities for concurrency. At the same time the RDRAM's RAS cycle time is as low as a

-50 EDO DRAM. The 533 MB/sec bandwidth allows easy synchronous connection to 66/100/133MHz CPU buses.

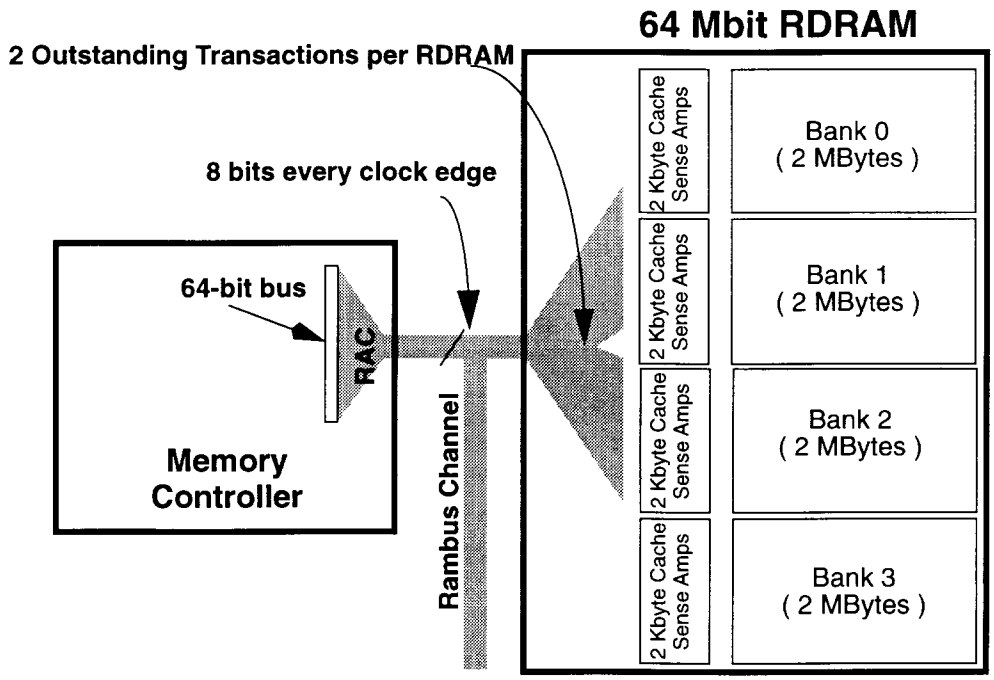

Figure 11: Example System Using the 64M RDRAM

The RAS/CAS concurrency allows access to a second page to take place during access to the current page. This pipelining of requests hides the access latency. The sustained data traffic across the Rambus Channel almost reaches the 533MB/sec peak rate. In typical CPU workloads, the 64M RDRAM is able to achieve 425MB/second bandwidth even on worst case main memory workloads, such as 32-byte transfers to random locations (100 percent misses).

The 64M RDRAM's four bank architecture also helps to minimize bank interference. In DRAMs with two-bank architectures, 50 percent of the accesses may be to the same bank, and will need to incur the miss penalty on the second access. The Rambus DRAM has a couple of features to offset this situation.

- Support for four open pages per RDRAM.
- Use of multiple RDRAMs provides greater number of banks and open pages.
- Four banks and flexible addressing allow tiling of memory to minimize bank interference.

With two RDRAMs providing eight open pages, only 12 percent of the accesses will experience bank interference in two-requestor environments. The controller designer can exploit these features to minimize bank interference.

Effective Bandwidth Comparison

The comparison of the effective bandwidth across DRAM alternatives are based on workloads of 32-byte random transfers. The random accesses results in 100 percent misses to the open page. Thus each transfer requires the DRAM complete a full RAS cycle.

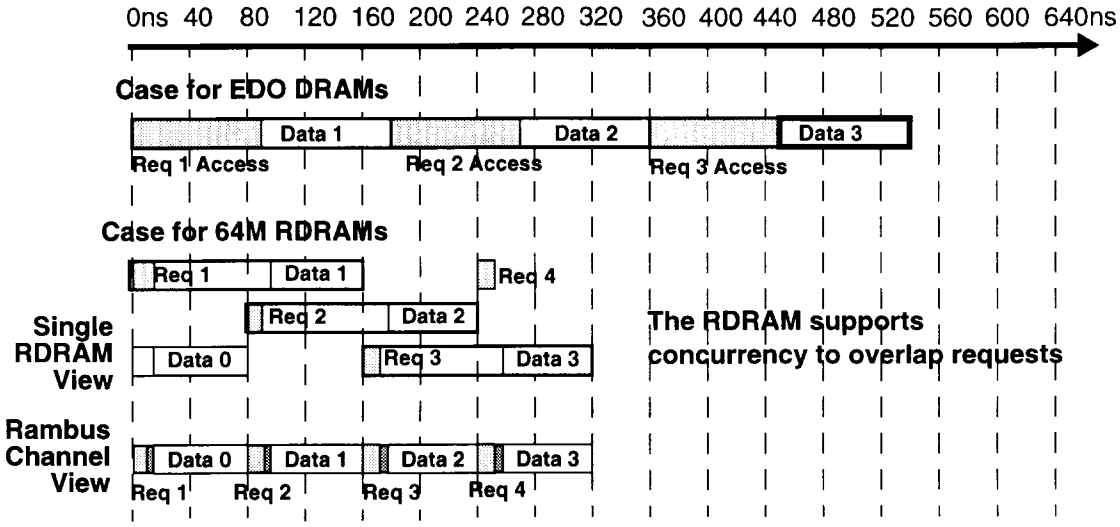

Figure 12: EDO Versus Rambus DRAM: Timing to complete three 32-byte transfers

The above figure shows the time required to complete three 32-byte transfers to random locations in memory. It is assumed that the accesses are to closed pages ("misses"), so that the access requires a full RAS cycle. The shaded "access" timeslots for the EDO case account for precharge time from the previous access, as well as RAS time for the next block transfer. Since the EDO DRAM is a single-bank design, the transactions are processed serially. In the Rambus case the multiple transfers are proceed in parallel, and so almost every timeslot on the Rambus Channel is filled with data.

The Rambus 64MB RDRAM is ideally suited for P6/P7 based systems. The P6 is the first split transaction x86 CPU, Rambus is the highest bandwidth split transaction

memory device. The P6 operating at 66MHz has a peak bandwidth equal to the Rambus peak bandwidth of 533MB/sec. A Rambus 64M RDRAM's four-bank architecture supports the P6's memory bus with its four outstanding transactions. The Rambus 64M RDRAM satisfies P6 requests in their minimum seven bus cycles for all page hits. It is the only memory technology able to satisfy P6 bandwidth from a single DRAM. Rambus technology enables 8- and 16Mbyte base P6 systems.

Performance: 8Mbyte System (1 DRAM)

Performance under worst-case workload conditions:
32-byte random requests, 100% page misses

Memory Type	Number of Controller Pins	Effective Bandwidth (w/o bank interference)	Effective Bandwidth (w/ bank interference)
64M EDO DRAM - 50ns, 32-bit bus width	55-60	107 MB/sec	107 MB/sec
64M SDRAM @ 66 MHz, 32-bit bus width	55-60	264 MB/sec	179 MB/sec
64M Rambus DRAM	31	425 MB/sec	380 MB/sec **

** The Rambus DRAM architecture provides more opportunity for minimizing bank interference

For low-end PC configurations, the Rambus based subsystem provides over three times the performance of the EDO subsystem and twice the performance of the SDRAM subsystem.

Performance: 16Mbyte System (2 DRAMs)

Performance under worst-case workload conditions:
32-byte random requests, 100% page misses

Memory Type	Number of Controller Pins	Effective Bandwidth (w/o bank interference)	Effective Bandwidth (w/ bank interference)
64M EDO DRAM - 50ns, 64-bit bus width	110-120	178 MB/sec	178 MB/sec
64M SDRAM @ 66 MHz, 64-bit bus width	110-120	528 MB/sec	385 MB/sec
64M Rambus DRAM	62	850 MB/sec	760 MB/sec **

** The Rambus DRAM architecture provides more opportunity for minimizing bank interference.

For mid-range PC configurations, the Rambus based subsystem using two Rambus Channels provides over four times the performance of the EDO subsystem and twice the performance of the SDRAM subsystem. The savings in up to 64 pins on the controller can be used to lower controller costs or for additional interface functions, such as multimedia or graphics support.

Memory Expansion Options for Rambus Subsystems

Rambus memory subsystems are easily expanded using horizontal or vertical RSockets and low cost SIMM-like RModules. The horizontal RSockets/RModules are for low profile add-in cards or laptop systems. The vertical RSockets/RModules are for systems (such as desktop main memory systems) that need to support many RDRAMs and upgrade sockets. Due to the narrow Rambus interface, the RSockets use only 32 pins, which allows for very low cost memory expansion as compared to

the 168-pin DIMM sockets that have been proposed for SDRAMs on 64-bit wide data buses.

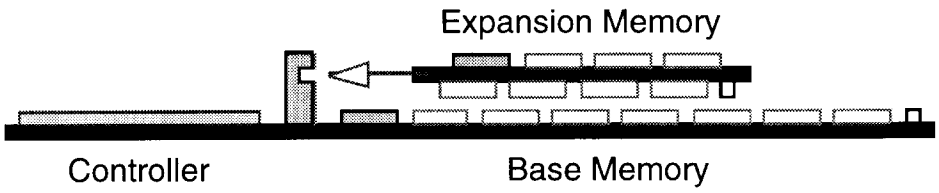

Figure 13: Horizontal Rambus RSocket / RModule Expansion

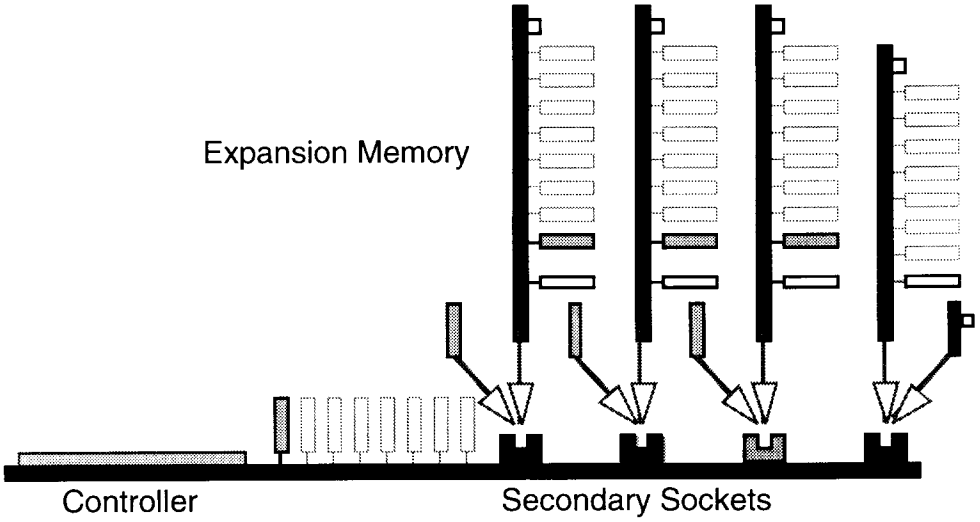

Figure 14: Vertical Rambus RSocket / RModule Expansion

Memory System Cost Considerations

Several factors influence the cost of the memory system. The first is the cost of the individual DRAM components, which are determined by die size, package costs and test time, as well as by competition. The Rambus 64M RDRAM is expected to be no larger than a x32 64M DRAM. Lower pin count and fewer bond wires yield a less expensive package. The 64M RDRAM will be available in a plastic surface mount package with 32 pins, in contrast to the more than 70 pins required for a x32 64M EDO DRAM.

Memory system costs include costs for the DRAM components, glue logic, clock source, board area, power, impact of EMI and memory expansion.

- Rambus memory systems are glueless and use low cost clock generators.
- Requiring no more than four square inches, the Rambus memory array can save up to 20 square inches of board space when compared to conventional memory systems.
- Rambus memory can be expanded using low-cost SIMM-like connectors and modules that need only 32 pins.
- When compared to a conventional high bandwidth memory system based on a 64-bit data path and clocked at 66MHz, the Rambus system results in up to a ten-fold reduction in electromagnetic emissions.
- Since Rambus memory systems typically use fewer DRAM components, it uses less power.

These factors — lower cost, lower power, smaller form factor and low-cost memory upgrade — make Rambus DRAMs extremely attractive for lower-end consumer PCs and laptop computers.

9-bit and 8-bit Organizations Supported

The 16M and 64M Rambus DRAMs are available in x8 and x9 organizations. While most PCs have eliminated error detection on the DRAM array in order to save costs and would use the x8 organization, higher end systems still require some degree of error detection/correction. Main system memory arrays can use the ninth bit to support single-bit error detection across eight-bit bytes, or to support single-bit correction, double bit detection across 64-bit data word. Systems designers can use the ninth bit for larger grained error detection schemes. Graphics and Video applications use the ninth bit to support a number of graphics techniques. These include keying for multiple color spaces; keying for masking of multiple source; Z-buffer storage with pixels or very large pixels; compressed 24-bit color; or for non-standard pixel formats.

The 64M Rambus DRAM: Optimal for PC Main Memory Systems

The use of the 64M Rambus DRAM in PC main memory systems provides the highest performance of any DRAM alternative. Users benefit from snappy applications performance, superior video and vibrant, color graphics. The eight-megabyte granularity, low cost sockets and SIMM-like memory modules allow cost-effective memory upgrades. Entry-level PCs provide good performance at the lowest price. Using multiple DRAM Channels allows performance upgrades to the highest performance memory subsystem. Users of Rambus based PCs will benefit from the cost-savings, increased performance and improved multimedia feature set for improved productivity, enriched user interface and expanded applications.