

DRAMs For New Memory Systems (Part 2)

Rambus, Ramlink Offer Revolutionary New Interfaces

By Steven Przybylski, Consultant, San Jose, CA

This is Part 2 of a three-part series on the next generation of DRAM designs. Part 1 discussed EDRAM, CDRAM, and SDRAM (see [070205.PDF](#)). This article covers Rambus and RamLink. It concludes with a comparison of the chip-level features of all of these new designs. Part 3 will provide an extended comparison of the new parts and their impact on the design of complete memory systems (see [070405.PDF](#)).

Over the past year, several new DRAM architectures have been proposed to overcome the growing mismatch between memory bandwidth and processor requirements. The two most revolutionary of these alternatives are Rambus and RamLink. Both of these designs dramatically improve bandwidth by significantly changing the physical, electrical, and logical interfaces. Both use high-speed, byte-wide paths to transfer address and data between the DRAMs and the memory controller. Consequently, moderately high-performance memory systems can be built from a single DRAM.

Rambus DRAM (RDRAM)

The RDRAM technology was developed by Rambus (Mt. View, CA) and has been licensed by major DRAM vendors such as Fujitsu, Toshiba, and NEC. The heart of the new interface is the Rambus channel—a bidirectional, current-mode bus with a peak bandwidth of 500 MBytes per second. Although the maximum bus length is approximately 10 cm, up to 32 RDRAMs can be connected by mounting them on edge. By using transceivers to extend the bus, up to 10 modules of 32 RDRAMs each can be connected into a single memory system. One of the significant features of the Rambus channel is that the peak memory bandwidth is constant regardless of

whether the memory system consists of one RDRAM or 320. This makes the Rambus solution particularly attractive for video memories and for small main memories using 16Mbit and larger DRAMs.

To achieve this high bandwidth without resorting to point-to-point connections, the Rambus design relies on:

- Terminated, controlled-impedance signal traces
- Small (600 mV) voltage swings
- 0.6-ns rise and fall times
- 0.3-ns setup and hold times
- Dual 250-MHz clocks

Rambus transfers data on both edges of the 250-MHz clock to achieve its high bandwidth. Figure 1 shows the connections for the two clock signals, ClockToMaster and ClockFromMaster, that synchronize the flow of information to and from the bus masters, which must be together at one end of the bus. Clock-skew problems are greatly reduced since the data and its clock always travel in the same direction down the bus. A serial communication loop formed by S_{in} and S_{out} is used for initialization.

Internally, RDRAMs have two banks, each with a single-line row cache that stores the row last retrieved from its DRAM bank, shown in Figure 2. Reads and writes that hit the row cache are acknowledged immediately. If an access misses the row cache, the RDRAM sends a “negative acknowledge” and then brings the desired row into its cache. During this DRAM-array access, the bus can be used to access other RDRAMs. The memory controller can optionally be programmed with the row access time so that it does not repeat its request until the access completes, thereby minimizing the bus traffic needed to retrieve a block of data.

Rambus transactions begin with a request packet consisting of a header containing the opcode, address,

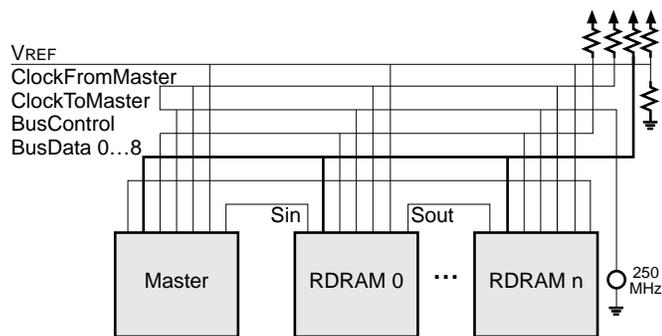


Figure 1. The Rambus channel allows up to 32 RDRAMs to be connected with a data transfer rate of 500 Mbytes per second.

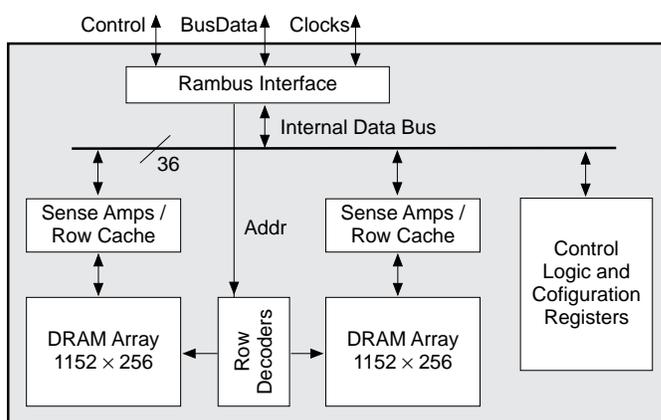


Figure 2. Block diagram of a 4.5 Mbit RDRAM.

First SDRAMs Announced

Samsung and NEC became the first companies to announce synchronous DRAMs. NEC is now sampling 16 Mbit parts at 33, 66, and 100 MHz. These parts are available in $\times 4$, $\times 8$, $\times 9$, $\times 16$, and $\times 18$ organizations. Samples of the 33-MHz parts are priced at \$180 to \$220 depending on the width. NEC expects to be in full production by 3Q93. Although the JEDEC standard for SDRAM is not yet complete, NEC is committed to delivering compliant parts once the specification is final.

Samsung is also sampling its SDRAM parts, but they are not JEDEC compliant, using a single-bank design with a level-sensitive RAS. The Samsung SDRAMs are offered at speeds up to 100 MHz in a $2M \times 8$ organization only. The company expects the chip to carry a 20% price premium over traditional DRAMs.

and packet size. Although the amount of data transferred is always an integral number of quadbytes (four 9-bit bytes), byte masks for the first and last quadbyte are included, providing support for byte addressability and byte-sized transfers. Protocol signalling occurs on a dedicated trace with the same timing characteristics as the data lines.

Unquestionably, system design with 250-MHz clocks is tricky. The Rambus company provides a "cook-book" solution to dramatically reduce this burden. They have specified the physical and electrical properties of all traces as well as the placement and value of the bypass capacitors and terminating resistors. Although the recommended PC board is a bit more expensive than some, it requires no unusual technology. By following these guidelines, a system designer is assured of reliable operation regardless of the vendor or size of RDRAMs used.

Although the RDRAM technology is proprietary, Rambus has licensed its technology to a number of companies. Toshiba, Fujitsu, and NEC will manufacture and sell RDRAMs as well as Rambus ASICs and other products. Additional partners include Augat and Molex for socket design, Biomation for logic analyzer support, and Toshiba/Vertex for a Rambus master ASIC megacell. Hitachi has announced a licensing agreement with Rambus but has not disclosed any specific product plans.

The RamLink Interface

RamLink is a revolutionary DRAM interface under development by the IEEE Computer Society P1596.4 working group. The group is applying the techniques adopted in the Scalable Coherent Interface (SCI) standard to develop a ring-topology interconnect for memory and I/O devices that provides high bandwidth and low latency.

Figure 3 shows a RamLink memory system that consists of one or more *ringlets*, each with a single con-

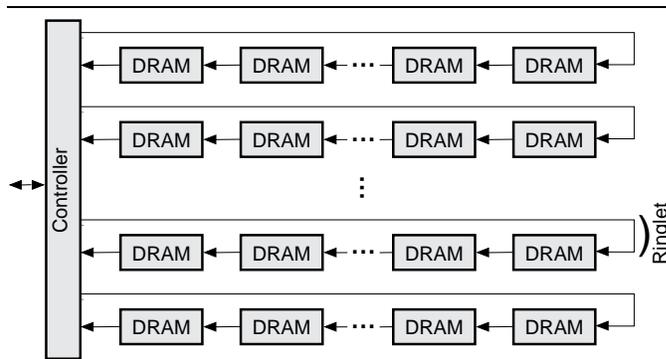


Figure 3. RamLink memory system with multiple ringlets.

troller (master) and up to 60 slaves. It uses 8 or 9 data lines to achieve a peak bandwidth of 500 Mbytes per second, equivalent to Rambus. In addition to the data lines, each link includes a flag signal and a 250-MHz clock. The design uses small voltage swings and differential signals to limit any physical or frequency constraints on future implementations.

RamLink packets come in four flavors: request packets, retry packets, response packets, and idle packets. Request packets initiate memory transactions. They are sent by the controller and contain a command header, address (6 bits of slave identifier plus 32 or 48 bits of per-slave address), checksum, and, in the case of write commands, the data to be written. The command header consists of type, size, and control information and contains either a specific response time or the maximum time allowed for the slave to respond. The control information includes a "sequential" bit that indicates whether subsequent requests will be to sequential addresses. Up to four transactions per device can be active simultaneously; thus, all packets have a two-bit transaction ID to unambiguously match request and response packets.

Retry and response packets are sent by slaves to indicate success or failure of a request. For a successful read, the response packet also includes the read data. For an unsuccessful request, the slave can indicate how much additional time it needs to complete the transaction. Idle packets are used by the controller to fill up otherwise unused cycles; non-DRAM slaves may use them to transmit interrupt requests.

RamLink supports both small and large transfers. Small transfers are 4 bytes long, while large transfers are 8, 16, 32 or 64 bytes long. Transfer addresses are always aligned to the transfer size. Small writes specify byte enables to facilitate single-byte transfers. The entire packet, including the address and control fields, is protected by a single check byte.

The RamLink protocol is designed so that a slave can forward a packet without examining it first, minimizing the latency through each slave. Since a ringlet can have up to 60 slaves, each nanosecond of per-slave delay can translate into a significant latency for a complete transaction.

Market Acceptance Inhibitors

The diversity of future DRAM options significantly complicates planning and design tasks. Many new technologies face a chicken-and-egg problem: a new part cannot be cost-competitive without significant volume, but companies are reluctant to adopt it without guarantees of competitive costs. Especially in the cost-sensitive PC arena, there is strong impetus for companies to wait to see which alternative will become broadly favored, thus avoiding being stuck with a product that is not cost-competitive.

Often a single large player or market segment can break this impasse by committing to a new alternative, causing an avalanche of design wins for that solution. Another way around this problem is if multiple vendors simultaneously enter an emerging market with plug-compatible products, keeping prices low through close competition. Unfortunately, superior technical solutions often end up as niche products for lack of the volumes needed to drive the price down.

Cost is such an important issue that DRAM vendors must dispel any hint of added cost in their products. Although all of the new interfaces add a modest amount of area to the generic DRAM, this area penalty is too small to inhibit market acceptance. Another concern with RDRAMs is the Rambus license fee for each part. Rambus will not reveal the exact amount, but it is estimated at a few percent of the part cost. By far the most significant factor in determining the comparative price of the new DRAMs will be sales volume, making future prices very difficult to predict.

Intangible perceptions can also influence market acceptance. For example, a revolutionary system inherently has a greater barrier to overcome due to a perception of greater risk and general unease with the unknown. Regardless of a vendor's size or resources, failure to develop broad corporate relationships can make some customers reluctant to become dependent on a single supplier. This is especially true today, as corporate alliances and consortia are often seen as imperative to corporate survival.

RamLink is still in a very preliminary stage. Although several DRAM vendors have indicated interest, none has committed to build RamLink devices. This interface is not expected to be available until 64-Mbit parts emerge. As such, it is an interesting technology for system designers to be aware of, but it is not of immediate or even medium-term significance.

Chip-Level Comparison

Each of the six new alternatives has a unique set of characteristics that makes it more or less appropriate for any particular application. Table 1 summarizes the basic specifications of the alternative DRAMs. It describes the

organization and speed of the currently-available devices and announced future organizations. For the conventional solutions and the SDRAMs, access times are representative of what is currently available for 4M parts.

The primary difference between the evolutionary and revolutionary alternatives is whether they retain the split address/data interface and the multiplexed address bus. For example, the evolutionary SDRAM retains the existing interface but adds a synchronizing clock signal to smooth the flow of data. Mitsubishi's CDRAM uses a similar synchronous interface and also includes a small (16 Kbit) SRAM cache. Ramtron's EDRAM modifies the interface only slightly to improve the efficiency of writes. The primary advantage of Ramtron's part is its fast DRAM core.

On the other hand, by completely breaking with tradition, the revolutionary solutions provide dramatically more per-chip or per-bit bandwidth. This bandwidth apparently comes at the expense of greater access latency, at least for 4M parts. This is somewhat deceptive because the Rambus access latency, for example, includes the row precharge time, which is not counted in the access latency of other DRAMs. For the evolutionary designs, row precharge will at least occasionally interfere with accesses, increasing the actual average latency. This increase is difficult to quantify, since the number of conflicts depends on the DRAM's cache characteristics, memory-controller design, and access patterns.

Another significant differentiator is the size and configuration of the cache on the DRAM. The Rambus and CDRAM devices contain significantly more cache (or row buffer) than the others, which can greatly improve the cache hit rate. Depending on the workload and system configuration, increased cache size can have a significant influence on the overall performance of the memory system. Only the CDRAM partitions the cache into relatively small chunks. It is also the only organization that facilitates a set-associative cache, improving performance when the access pattern has poor locality and there is no large external cache in the system. The flip side of this organization is that fewer bits are transferred from the DRAM array to the cache on every row access. (The issues of cache block and transfer sizes will be discussed further in Part 3 of this series.)

As the number of transistors used by the DRAM core continues to increase, the relative size of the interface logic becomes small. Even the 16-Kbit SRAM included in the CDRAM interface does not take up a significant portion of the die. In fact, for 16-Mbit parts, the new DRAM alternatives all suffer less than a 15% area penalty compared to generic DRAMs. This area variation is less than the die size differences among the various vendors for the generic DRAMs. The only significant cost difference between the various alternatives is package cost, which will be higher for solutions that use

	Conventional		Evolutionary			Revolutionary	
	Generic DRAM ¹	Wide DRAM ¹	SDRAM ²	CDRAM	EDRAM	RDRAM	RamLink
Proponent		Various	JEDEC Committee	Mitsubishi	Ramtron	Rambus	IEEE CS Working Group
Vendors/Manufacturers	Many	Various	NEC, Samsung, others	Mitsubishi + second source	NMB	Toshiba, Fujitsu and NEC	None Committed Yet
Interface	Asynchronous RAS/CAS	Asynchronous RAS/CAS	Synchronous RAS/CAS	Synchronous RAS/CAS + SRAM access	Enhanced asynchronous RAS/CAS	Custom high-speed sync. 9-bit bus	Custom synchronous token ring
Current Sizes and Organizations	4 or 16 Mbit	4 or 16 Mbit	16M samples	4 Mbit	4 Mbit	4.5 Mbit	None
	1M × 4, 4M × 1, 4M × 4, 16M × 1	×8, ×9, ×16, and ×18	×4, ×8, ×9, ×16, ×18	1M × 4	1M × 4, 4M × 1	512K × 9	
Future Organizations		×32 and ×36		256K × 16, 4M × 4	512K × 8, 4M × 4, 2M × 8	2M × 8, 2M × 9	First availability likely at 64 Mbits
Best Current Row Access Latency	60 ns	60 ns	60 ns	70 ns	35 ns	136 ns ⁵	Unknown
Cache/Buffer Organization (4Mbit part ³)	1 Row Buffer of 4 Kbits	1 Row Buffer of 4 Kbits	1 or 2 Row Buffers of 4 Kbits (8 Kbits ⁴)	256 Blocks of 64 bits (16 Kbits)	1 Row Cache with fast writes (2 Kbits)	2 Row Buffers of 9 Kbits (18 Kbits)	Vendor specific
Cache/Buffer Access Latency	30 ns	30 ns	30 ns	10 ns	15 ns	36 ns	Unknown
Peak Chip Bandwidth	133 Mbps (×4 part)	600 Mbps (×18 parts)	800 Mbps @ 100 MHz	400 Mbps	267 Mbps (×4 part)	4.5 Gbps	4.0 Gbps or 4.5 Gbps
Package	26/28 pin SOJ/TSOP	40/44 pin SOJ/TSOP (×16 & ×18 parts)	44 pin TSOP (2M × 8)	44 pin TSOP	28 pin SOJ	32 pin VSMP	Unspecified
Electrical Interfaces	TTL	TTL	LVTTL and GTL/CTT @ 100 MHz	TTL (4M), GTL or LVTTL (16M)	TTL	Single-ended, small swing, current mode	Differential low voltage swing

1. Access times are representative of the class of parts as a whole.

2. Access times assume the same DRAM core as represented in the Generic DRAM column.

3. Except the SDRAM which is a 16 Mbit part.

4. For 2 bank organization.

5. RDRAM latency includes precharge time.

Table 1. Summary of specifications for conventional and alternative DRAM approaches.

wider data paths. This added cost applies mainly to the wider generic DRAM, as all of the new interfaces keep the data path to a reasonable width, at least in their initial implementations.

The cache/buffer access times vary significantly across the designs. The minimum access latencies at the chip level can be deceiving, however, since each organization imposes additional overhead and transfer times. Even as a representation of the peak bandwidth, these numbers cannot be used as a basis of direct comparison because of the difficulties in designing reliable, high-speed systems with TTL or even LVTTL electrical interfaces. Until systems move to a terminated, controlled-impedance interface with low signal swings (such as GTL), it will be difficult to operate even the synchronous parts at much above 66 MHz. This difficulty is exacerbated by the fact that none of the conventional or evolutionary designs has reduced the pin input capacitance below the traditional 5 pF to 10 pF. In contrast, the RDRAM input capacitance is just 2 pF.

Most PCs today are designed with SIMM-based memory systems. Since more than 50% of all DRAMs currently go into PCs, there has been a lot of emphasis recently on flat SOJ and TSOP packages. As conventional and evolutionary designs grow beyond 40 pins, they are moving from 1.27-mm (50 mil) pin pitches to 0.8-mm (31 mil) pitches to retain a roughly comparable package size. In contrast, RDRAMs use a vertical, surface-mount package similar to the ZIP (Zigzag In-line Package). This vertical package maximizes the number of DRAMs that can be placed on the short Rambus.

Each of the new DRAM alternatives changes one or more of the physical, electrical, or logical interfaces of the generic, narrow DRAM with its asynchronous, multiplexed address bus. These changes impact the applicability of the organizations to different types of memory systems. In our next issue, we'll conclude with a look at the characteristics of the memory systems that result from use of these new devices, and look ahead to the 16M generation and beyond. ♦