Rambus Unveils Revolutionary Memory Interface

Rambus Channel Provides 500 Mbyte/s Memory Interface

By Michael Slater

Rambus Inc. has unveiled its radical new processor-to-memory interface and DRAM architecture, which promise to create the most significant change in processor/memory system architecture since the introduction of the DRAM two decades ago. The Rambus technology includes a new interface, the Rambus Channel, and a revised internal architecture for DRAMs. It eliminates the familiar RAS/CAS control signals and multiplexed address, replacing them with a packet-oriented bus.

Rambus is venture-funded and currently has 31 employees. The company was founded in March 1990 by Mike Farmwald, a former Associate Professor of Electrical and Computer Engineering at the University of Illinois, and by Mark Horowitz, an Associate Professor of Electrical Engineering at Stanford University. Farmwald now serves as Chief Scientist at Rambus; Horowitz is a VP and member of the technical staff but has returned to full-time teaching at Stanford. The VP of Engineering is Allen Roberts, formerly at MIPS and Elxsi. All three were key participants in the design of the R6000-based ECL system at MIPS, and the Rambus Channel design has its heritage in the system bus used in the MIPS RC6280 system. The president is Geoff Tate, formerly a senior VP at AMD where he was responsible for all microprocessors and peripherals.

The Rambus Channel is a 500-Mbyte/s interface, operating with a 250-MHz clock and transferring a byte of data on each clock edge—one byte every 2 ns. The channel provides a 9-bit data path and includes 28 signals, including power and ground. There are no address lines; addressing is performed by packets transmitted via the data lines. Despite its very high clock rate, the interface uses standard CMOS circuits and is intended to be a high-volume, low-cost technology.

The Rambus approach is somewhat counter-intuitive, in that it achieves its high bandwidth through fast data rates rather than wide buses. One benefit is that a minimum system can use a single Rambus memory device, regardless of the system's word width. The narrow bus also minimizes the number of pins and amount of board area required. While designers may be wary of the high clock rates, the intent is that all the high-speed circuitry will be provided in Rambus devices, and the system designer will only need to follow simple layout guidelines to ensure reliable operation.

Rambus Inc. is a technology provider, not a manufacturer. It has licensed its technology to Toshiba, NEC, and Fujitsu—three of the top five manufacturers of DRAMs and ASICs. The agreements are non-exclusive, and more licensees are expected to join in. In addition to DRAM makers, the interface technology will be licensed to ASIC and microprocessor suppliers.

The role of Rambus as an independent technology provider may enable a degree of change that would be difficult for any semiconductor company to achieve. Any single DRAM maker would naturally be hesitant to introduce a radically new interface, especially one that requires special circuit design support from ASIC and processor vendors.

Because of the difficulty of introducing a new standard, DRAM interfaces have not evolved along with microprocessor technology. Nearly all DRAMs today still use the same interface designed in the early '70s, with the address multiplexed in two parts. While DRAMs have increased modestly in performance, microprocessor speeds have increased much more quickly. The nature of microprocessor memory access patterns has also changed dramatically due to the introduction of on-chip caches as a standard feature of high-end microprocessors. The Rambus technology revamps the DRAM interface to better exploit the bandwidth available inside DRAMs and to better match the needs of contemporary microprocessors and the capabilities of sub-micron CMOS technology.

Rambus Technology

The Rambus technology includes several aspects that combine to make the high-bandwidth interface practical:

- The Rambus Channel, a 500-Mbyte/s, 9-bit-wide, multiplexed interface designed to connect one or more masters (typically processors) to one or more slaves (typically Rambus DRAMs).
- CMOS circuit designs for the Rambus interface. This interface can be integrated into the pad ring of an interface ASIC or microprocessor.
- Logic designs for the Rambus master and slave interfaces, to be incorporated into processors (or interface ASICs) and DRAMs.
- A new package design for Rambus DRAMs (RDRAMs).
- A socket design for RDRAM modules, designed in collaboration with Augat.

Figure 1 shows a block diagram of a Rambus system. The Rambus Channel has nine data lines, which most commonly will be used for 8 data bits plus parity.

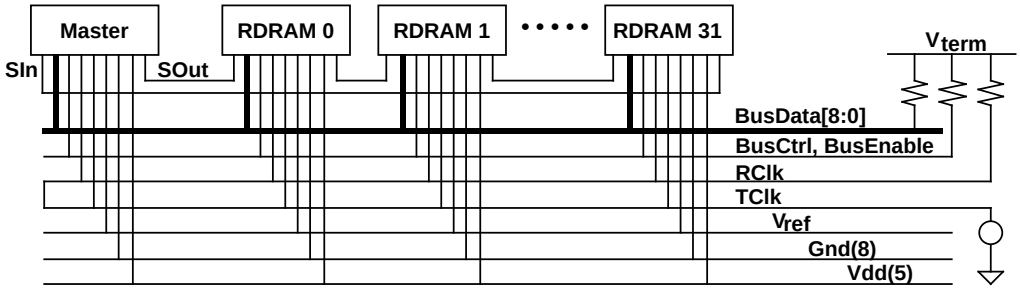

Figure 1. Rambus system block diagram.

The ninth bit is used in request packets and is therefore required by the interface, but it need not be implemented in the memory array if parity is not required.

The high clock rate is made possible by a combination of special CMOS interface circuits, careful circuit board layout, short trace lengths, RAM packages with low parasitic capacitance and inductance, and low voltage swings. Rambus provides an example circuit-board layout that is designed to work with standard FR4 epoxy-glass printed-circuit boards. The only additional circuit board cost beyond a basic DRAM system is that a four-layer board must be used to provide the controlled-impedance, transmission-line characteristics and low-noise power and ground required. The channel is limited to a maximum length of about 4 inches.

The 250-MHz clock is routed as a U, driven from the DRAM end of the channel to form the transmit (slave-to-master) clock and looped back to the receive (master-to-slave) clock line at the master end. With this arrangement, the clock signal travels along with the data, minimizing skew. A phase-locked loop on each Rambus device limits clock skew within the chip.

Only Rambus devices connect to the Rambus channel, so no compatibility with existing logic-level standards is needed; using standard logic levels, practical data rates are a factor of ten slower. Rambus signals are pulled up with termination resistors to a V_{term} level of about 2.5 V. A reference voltage, V_{ref}, is about 300 mV below V_{term}; it is distributed with the bus signals to simplify the receiver design and provide some common-mode noise immunity. The drivers sink current to pull signal lines about 600 mV below V_{term} for a logic 1, providing a signal swing of ±300 mV around V_{ref}.

RDRAMs will be the first Rambus devices to become available, but EPROMs and flash memories are also possible. For early system designs, the Rambus master will typically be a memory interface ASIC that translates a microprocessor's native bus to the Rambus protocol. The initial three Rambus licensees are also leading ASIC vendors, and they will provide the required interface cell as part of their ASIC libraries. In 1993, microprocessors incorporating a direct Rambus interface are expected.

Rambus DRAMs

The first RDRAMs will be 4.5-Mbit devices. (The size is 4.5 Mbits instead of the usual 4 Mbits because of the ninth data bit.) The memory is divided into two banks, each with 256 rows and 1024 columns, each of which is 9 bits deep. Within the memory banks, the chip design is the same as a conventional DRAM.

RDRAMs incorporate the functions of a DRAM controller into the RDRAM, eliminating the need for logic between the Rambus master and the RDRAM. This simplifies the system design and also reduces the latency introduced by the DRAM controller.

In addition to the Rambus Channel interface, the major difference between a standard DRAM and an RDRAM is the way the sense amplifiers are used. In any DRAM, when a row is accessed, the data from all columns is latched into the sense amplifiers. Page-mode DRAMs take advantage of this to provide faster access to locations within the same page; RDRAMs use the sense amp latches in a similar way, but the division of RDRAMs into banks and the fact that each RDRAM is addressed separately allow RDRAMs to make more effective use of the sense amp latches as a cache.

In a conventional page-mode DRAM system, all the DRAMs provide bits for a single memory word, so the addressed page acts as one very long cache line distributed across the DRAMs. In a Rambus system, each RDRAM provides all the bits for a given address, so each sense-amp array acts as a separate cache line. The 4.5-Mbit RDRAM includes a total of 2 Kbytes of cache, organized as two lines of 1K bytes each. In a system with 8 RDRAMs (providing 32 Mbytes plus parity), there is a total of 16 Kbytes of cache, organized as 16 lines of 1K each.

Each RDRAM has mapping registers that determine its address assignment. Other control registers allow the master to set the RDRAM's refresh mode and put the device in a power-down mode. Read-only registers within the RDRAM provide manufacturer ID, device type, and device size information, in addition to the device's required minimum timing parameters.

The second generation of RDRAMs will be 18-Mbit devices, expected in 1993, using comparable technology to the 16-Mbit DRAMs that are just now going into volume production. The higher-density RDRAMs will be divided into more banks, increasing the number of on-chip cache lines, and the access latency will be cut approximately in half. Since there are no explicit address pins on an RDRAM, no pinout changes are required as

the density is increased.

The 4.5-Mbit RDRAMs operate from a 5-V power supply. The 18-Mbit devices will be available in both 3.3-V and 5-V versions. The Rambus signal levels are the same, however, regardless of the power-supply voltage, so RDRAMs using either power supply voltage can be easily mixed in a system.

Because of the high clock rate, RDRAMs have higher power consumption than conventional DRAMs. Power consumption for a 5-V, 4.5-Mbit RDRAM is expected to be about 300 mW in standby, increasing to 600 mW when it is active and 1100 mW when data is being transferred. Until 18-Mbit, 3.3-V RDRAMs become available, Rambus technology will not be suitable for low-power systems.

Rambus Protocol

A Rambus system initializes the slave devices at power-up using the S_{in} and S_{out} signals. Each slave keeps its S_{out} signal negated until it has been initialized. The processor asserts its S_{out} signal, which connects to the first slave's S_{in} pin, enabling that slave's configuration registers to be accessed. When the first slave has been initialized, it asserts its S_{out} output, which enables the next slave. This process continues until all the slaves have been initialized. As part of this process, the master sets the address range to which each slave will respond. After initialization, the only other time S_{in} and S_{out} are used is to provide a refresh timing signal if power-down mode is used.

Figure 2 shows a Rambus read transaction, assuming a hit in the RDRAM's cache. A 16-byte read takes 80 ns: 12 ns to transmit the request, 36 ns minimum latency before a response can be sent, and 32 ns to transfer 16 bytes of data.

The six-byte request packet encodes a 36-bit address, a 4-bit operation code, and an 8-bit transfer length count (in bytes). Byte addressing and block sizes of up to 256 bytes are supported.

All RDRAMs must monitor all request packets, and each device compares the address to its own address range, as set in its configuration registers. The addressed device responds with a positive acknowledge (Ack) if the request hits in its cache, or a negative acknowledge (Nack) if it misses. The Ack/Nack signals are transmitted on the BusCtrl line during two bit times.

In case of a miss (and the resulting Nack), as shown in Figure 3, the master must retry the access after waiting for a minimum period to allow the RDRAM to write its cache line into the DRAM array and fetch the requested line. For the initial 4.5-Mbit RDRAMs, the delay before the access can be retried is 160 ns, so the total time from the first request to the completion of a 16-byte transfer in the case of a read miss is 240 ns. (Faster versions may also be available.) A Nack is also sent by the RDRAM if it is busy for some other reason, such as an internal refresh cycle.

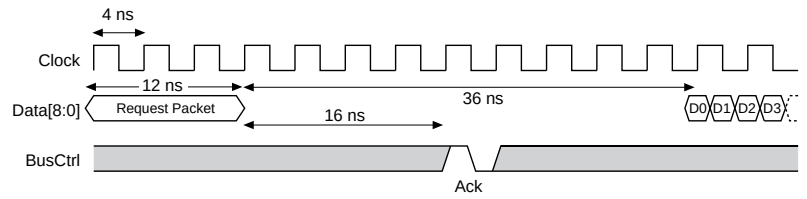

Figure 2. Timing diagram for a 16-byte read with a hit in the RDRAM's sense-amp cache.

For comparison, a fast-page-mode conventional DRAM system can process a page miss in about 100 ns and a page hit in about 40 ns. If the memory system is only 32 bits wide, a 16-byte access would require four transfers, taking a minimum of 160 ns in the case of a page hit and 220 ns for a page miss. Thus, a conventional, 32-bit-wide page-mode system can be slightly faster than the Rambus system when a miss occurs, but considerably slower when there is a hit.

The key advantage of the Rambus approach is that the same bandwidth is achieved for an 8-bit-wide system, requiring only a single RDRAM and providing a granularity of 0.5 Mbyte with 4.5-Mbit chips. To outperform a Rambus system on a page hit, a conventional DRAM system would have to be 128 bits wide, resulting in a granularity and minimum memory size of 16 Mbytes using $1M \times 4$ DRAMs.

The other advantage of the Rambus design is that the hit rate is much higher than for a page-mode DRAM system, so the longer latency in the case of a miss has less of an effect. According to Rambus, a 4-Mbyte page-mode system using conventional $1M \times 4$ DRAMs, which has a 2-Kbyte page size, has a page hit rate of only 30–40%, while a comparable RDRAM system has a hit rate of about 95%. The RDRAM system has a major bandwidth advantage after the initial latency. The latency after a miss is significantly greater for an RDRAM than for a conventional DRAM, however, because data in the sense-amp cache must be written to the memory array before the requested data can be read. The Rambus advantage is greatest if large block sizes can be used.

For write cycles, a delay of only 4 ns is required after

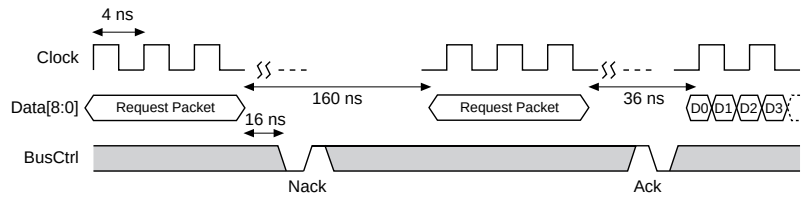

Figure 3. Timing diagram for a 16-byte read with a miss in the RDRAM's sense-amp cache.

Price & Availability

Fujitsu and Toshiba plan to ship 4.5-Mbit RDRAMs later this year, with 18-Mbit devices to follow sometime in '93. NEC will not produce the 4.5-Mbit devices, but expects to ship 16- and 18-Mbit (×8 and ×9) devices in the first half of '93. Rambus expects pricing to be about 25% per bit higher than standard DRAMs for 4.5-Mbit chips, and 15% higher for 18-Mbit chips. Toshiba and Fujitsu both plan to provide interface ASICs this year. In 1993, Fujitsu plans to provide Rambus support for its SPARC processors, and NEC expects to provide ASIC support and an R4000-to-Rambus interface chip.

For more information on Rambus technology, contact Dave Mooring at Rambus Inc., 2465 Latham Street, Mountain View, CA 94040; 415/903-3820; fax 415/965-1528; e-mail mooring@rambus.com.

For information on Rambus DRAMs and ASICs, contact the three semiconductor licensees:

Charlie Shafton, Fujitsu Microelectronics, 3545 N. First St., Bldg. 1, San Jose, CA 95134; 408/922-9625; fax 408/432-9044.

Hank Bardsley, NEC Electronics, P.O. Box 7241, Mountain View, CA 94039; 415/965-6358; fax 415/965-6374.

Avo Kanadjian, Toshiba America, 9775 Toledo Way, Irvine, CA 92718; 714/455-2257; fax 714/859-3963.

the request packet before sending the data; the master does not have to wait for an acknowledge. If the RDRAM cannot process the write because of a cache miss or a refresh conflict, it ignores the data and sends a Nack, and the master must retry the transfer later.

The Rambus supports up to 4 masters, although physical and electrical considerations limit initial systems to two masters; it is not intended for symmetric multiprocessor systems. Examples of masters other than the main processor include graphics controllers and network interfaces. Arbitration typically will be provided separately from the Rambus signals.

Note that while the Rambus uses a split-transaction protocol, in which the bus is not active between a request and a response, the bus is effectively locked while a request is pending; multiple requests cannot be pending simultaneously. This simplifies the design of the slave devices, since they do not need to include arbitration logic; the bus is guaranteed to be free when they are ready to respond.

Figure 4 shows the RDRAM package, which provides 32 contacts on 1 mm centers. RDRAMs can be mounted on 100-mil centers, allowing a 4-Mbyte memory system (using 4.5-Mbit chips) to be implemented in less than one square inch. The RDRAM package looks much like a conventional single-in-line package, except that it is a surface-mount device. The package is a key enabler of the high-speed interface; it is designed to minimize lead capacitance and inductance by putting all the signal pads on one side of the die, keeping all the lead lengths uniformly short.

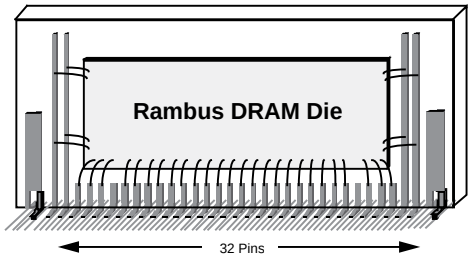

Figure 4. Rambus DRAM package.

Rambus has also defined a standard connector for Rambus modules. This connector has been specially designed to meet the signal integrity requirements for the high-speed Rambus. A Rambus module can be a master or a slave. Examples of Rambus master modules include display controllers and network interfaces. More typically, Rambus modules will be slaves, providing up to 32 RDRAMs each. A Rambus transceiver isolates the module from the primary Rambus; up to 10 modules can be connected to a single Rambus. With 32 RDRAMs per module, this allows a maximum system size of 320 RDRAMs, which provides 160 Mbytes using 4.5-Mbit devices or 640 Mbytes with 18-Mbit devices.

System Architecture Implications

If the Rambus technology lives up to its expectations, it will have a dramatic effect on computer system architecture. Because of their high bandwidth, RDRAMs may eliminate the need for a second-level processor cache in uniprocessor systems. Rambus claims that a 50-MHz 486 system using RDRAMs but no secondary cache performs nearly identically to a system with a 128-Kbyte, write-back secondary cache and a high-performance, conventional DRAM system. Assuming that the RDRAMs carry only a modest price premium over conventional DRAMs, the Rambus solution should be considerably less expensive.

If the Rambus interface is integrated into the processor, additional cost savings can be realized. For the MIPS R4000, for example, the secondary cache support adds 207 signal pins—plus additional power and ground pins—to the processor. A Rambus interface requires only 28 pins, and it could replace the main system bus connection as well as the secondary cache interface. In this case, an external interface chip would translate the Rambus to an I/O bus for devices that don't connect directly to the Rambus.

Continued on page 21

ads will presumably become more specific, but so far they have been of questionable value.

The 486DX2 can be used to upgrade some 486DX systems by removing the CPU and replacing it with the DX2, but there are several potential problems:

- The chip's power consumption is substantially higher, so the cooling in some systems may not be adequate. The DX-25 draws 550 mA typical and 700 mA maximum, while the DX2-50 draws 775 mA typical and 950 mA maximum.
- While the interface timing *specifications* are identical, the *actual* timing is slightly different. This can cause problems in some marginal system designs.
- Some BIOS programs include speed-dependent timing loops.

Intel says that its testing revealed that about one system in four encountered problems. Making a list of systems that can be safely upgraded isn't as easy as it might seem, since it sometimes depends on which revision of the system board and BIOS is present. Some system vendors may certify their systems as upgradeable. Computer dealers may offer unauthorized upgrades, and sophisticated end-users may be willing to try the upgrade themselves, but the potential for problems is significant.

From the system maker's perspective, upgrade chips cut them out of the upgrade sale. Thus, some may prefer to sell upgrade CPU cards, offering larger caches or other features in addition to the faster processor. Some system makers more focused on their profits than on benefits to their customers might *ensure* that their BIOS contains speed-sensitive code as a way of making users come to them for an authorized upgrade, so they could charge more than the street price of a 486DX2 chip for the new processor and a new BIOS ROM.

Conclusions

Intel's 486DX2-50 will give system makers an easy way to upgrade their 25-MHz system designs, providing a cost-effective alternative to true 50-MHz systems. When the 486DX2-66 becomes available, it will provide a new pinnacle of 486 performance without stressing the ability of PC makers to build systems around it. The end-user versions of these processors—the OverDrive chips—will expand Intel's market by giving them a chance to sell a second processor to 486 system buyers.

This technology will also be important for Intel because it will help keep them one step ahead of 486 cloners. With the 386, Intel let its design age without upgrading it, making it easy for other semiconductor vendors to provide faster versions. With the 486, Intel will not give them this opportunity, and the clock-doubler chips are one more weapon in this battle. ♦

Rambus

Continued from page 18

RDRAMs also threaten to replace VRAMs in frame buffers. With the higher bandwidth of the RDRAMs, it is practical to use part of main memory for the frame buffer. This would have some impact on the CPU performance, since transfers from the frame buffer to the video output logic would contend with CPU requests. For maximum performance, the frame buffer controller could have its own private Rambus with a dedicated frame buffer. The separate video port provided by VRAMs would not be needed, since the Rambus interface provides more than enough bandwidth for processor access and screen refresh, even in a high-resolution, full-color system. Unlike VRAMs, RDRAMs are not specific to video applications, so they have the potential for higher volume and the resulting lower costs. RDRAM chips will also be smaller than VRAMs.

Conclusions

Rambus has high expectations for its technology—the company predicts that RDRAMs will account for over 50% of DRAMs within five years. While it is conceivable that this degree of success will be realized, there are several obstacles.

The price of RDRAMs is a key issue. Rambus claims that RDRAMs are about 20–25% larger than conventional DRAMs for 4-Mbit devices, with the area premium dropping to about 10–15% for 16-Mbit chips and 5% for 64-Mbit chips. Rambus DRAMs will initially have a much lower volume than conventional DRAMs, however, and there will not be as much competition, so the price premium could be higher.

Some designers are skeptical of whether Rambus will meet its 250-MHz clock rate goal, but Rambus says that test chips have been fabricated and they have had no difficulty reaching this rate. Until RDRAMs and systems using them are in production, however, some question will remain regarding the practicality of this speed. Another question is whether the high hit rates will be achieved in systems with only a few RDRAMs; as the hit rate drops, performance drops drastically.

There may also be some resistance among system vendors to using a technology that requires royalties to be paid to Rambus. This royalty will be buried in the cost of RDRAMs, interface ASICs, and eventually processors, but it is nonetheless present. System makers will also have a limited choice of vendors for interface ASICs, since they must go to a vendor that has licensed the Rambus interface cell.

If these barriers are overcome, however, the Rambus technology has the potential to reshape memory system design, graphics system architecture, and microprocessor bus interfaces. ♦