From the companion CD-ROM to the IEEE CS Press book,
"The Anatomy of a Microprocessor: A Systems Perspective,"
by Shriver & Smith

DIRECT RAMBUS TECHNOLOGY: THE NEW MAIN MEMORY STANDARD

Richard Crisp

Rambus Inc.

Providing three times the memory bandwidth of the 66-MHz SDRAM subsystem, Direct RDRAM modules fit seamlessly into the existing mechanical space and airflow environment of the industry-standard PC chassis.

Conventional DRAM architectures have reached their practical upper limit in operating frequency and bus width. Mass-market CPUs operating at over 200 MHz and media processors executing more than 2 GOPs (gigaoperations per second)[1,2] are now in production. Their external memory bandwidth of approximately 500 Mbytes/s cannot meet increasing application demands. In addition, no longer does just the CPU consume the majority of main memory bandwidth. A modern multimedia PC's graphics accelerator, media processor, and system I/O all consume significant memory bandwidth.

Efforts to extend conventional DRAMs have included scaling the SDRAM's memory clock from 66-MHz to 100-MHz operation. However this adaptation created numerous system design issues while offering only 33% additional peak bandwidth. This article explores the memory bandwidth scaling problem and then describes our solution, the Direct RDRAM device, which successfully meets multimedia requirements and fits seamlessly into the PC chassis.

Bandwidth scaling problem

A user's perception of interactivity and performance of a multimedia computer is largely determined by processing throughput. Fast processors are necessary, but memory bandwidth also plays a key, though often-overlooked, role.

Because current memory subsystems can only transfer data for one requester at a time, the length of time required to finish a transfer in progress adds to the latency of any pending requests. For a given bus width and clock frequency, the amount of time the bus is occupied depends on the transfer size and the memory bus bandwidth. Therefore, memory bandwidth directly affects memory system latency.

In multimedia computing, bandwidth-dependent latency is a dominant factor in the memory subsystem's performance. The focus of mass market computing on bandwidth-intensive multimedia-oriented applications further increases the memory bandwidth requirements. So, in next-generation multimedia PCs, memory bandwidth will largely determine a user's perception of interactivity and performance.

Traditional approaches to increasing memory bandwidth include speeding up the memory clock, increasing the bus width, or both. For conventional DRAMs, these approaches are reaching their practical limits.

Clock rate scaling. This approach is the most technically challenging. The legacy matrix interconnection topology of SDRAM-based systems simply does not lend itself to economic scaling beyond 100 MHz. Even the transition from 66-MHz to 100-MHz system operation is expected to be challenging due to stringent system timing requirements that dictate precise component and PCB modeling.[3]

There are several classes of system nets in a conventional memory system. For example, an SDRAM-based system may have an address net, a clock net, a data net, a DQM net, and a control net (CS, WE, RAS, CAS) (Figure 1). Each of the nets has a different loading and settling time from the other nets. A key issue limiting memory bus frequency in these systems is the fact that the loading on these nets increases nonuniformly from net to net as memory modules are added to the system (Figure 2a).

Motherboards are designed to operate reliably at both minimum and maximum system memory capacities. The system timing depends on the signal loading, which, in turn, depends on the number and storage capacity of modules inserted. Since the delay of the various nets scales nonuniformly, the system's timing margin degrades.

0272-1732/97/$10.00 © 1997 IEEE

INFINEON 1019

Another frequency-limiting factor in SDRAM-based systems results from the fact that the SDRAM modules (DIMMs) are connected in parallel to the primary bus transmission lines routed on the motherboard. Because each DIMM signal either has a heavy capacitive load or long module routing trace (or both), each DIMM signal represents a significant stub load on the motherboard. These stubs cause troublesome signal reflections if left unterminated.

Systems that operate significantly faster than 66 MHz need faster DRAMs to deliver balanced performance. Often the DRAM modules must be buffered, either on the module or the motherboard. Though buffered modules reduce the dependence of the motherboard timing on the module loading, and reduce the effect of the stubs, they have a disadvantage. Buffered modules require additional components, PCB area, routing, and system power. They also add one or two clock cycles to every memory access, depending on the extent of the buffering.

Data transfer. A second approach to increasing memory bandwidth is transferring memory data on both clock edges without changing the properties of any other nets. Since the address net has the highest loading dependent delay, leaving that network unchanged simplifies the design task. Yet one of the critical problems is meeting the required setup and hold specifications for the data bus at each device. Changing to a rising- and falling-edge clocked data bus necessarily requires improved clock access time specifications.

Current SDRAMs require nearly a whole clock cycle to establish valid data on the output pins. For example, a SDRAM with a 10-ns cycle time has a worst-case output delay from the rising clock edge of 9 ns. To trigger the output buffer to drive data on both clock edges while maintaining a 10-ns minimum clock period, such SDRAMs must feature a reduced output delay of at least a factor of two.

Due to the difficulty in meeting bus-timing constraints, the maximum system clock frequency must be reduced from that of a single-edge clocked system to avoid violating critical timing specifications.[4] Any clock rate reduction would therefore come at the expense of memory control bandwidth. In many cases, memory control bandwidth limits performance because each memory word may come from a nearby but different address such as in the nonsequential accesses characteristic of texture map rendering.

Increased bus width. A third scaling approach involves increasing the bus width to 128 bits. Although a simple idea electrically, it comes at the expense of doubling pins, memory, word width, I/O power, and memory granularity. Furthermore, doubling the bus width creates a host of mechanical and PCB layout problems.

Since a wide, high-speed bus can generate large transient currents in the driver elements, a significant number of ground and power pins are needed on the controller to support a large number of bus I/O pins. Because a 64-bit SDRAM or page-mode/EDO bus interface typically uses between 110 and 130 pins (counting power and ground pins), a 128-bit-wide bus will have significantly more than 200 pins.

Core logic chipsets used in PCs have ports for the CPU, the graphics, one or more system I/O buses, and a 64-bit memory interface. These core logic chips require as many as 472 pins today,[5] so doubling the memory interface pins is not an attractive option. The extra pins take more silicon area, increase package cost, and increase on-chip supply noise. Wide buses also consume more power. For example, a 128-bit LVTTL bus operating at 100 MHz and driving a 3.3-V swing into an 80-pF load/pin consumes over 5.5 watts versus 2.75 watts for a 64-bit bus.

Increasing the bus width also increases the memory granularity. For a 128-bit bus using 64-Mbit devices (4M×16), the minimum memory capacity is 64 Mbytes. If ×8 devices are used, the granularity jumps to 128 Mbytes. When the 256-Mbit generation reaches cost parity with the 64-Mbit generation, the granularity issue will worsen by a factor of four. The granularity issue is particularly important in applications that require only a small amount of high-bandwidth memory such

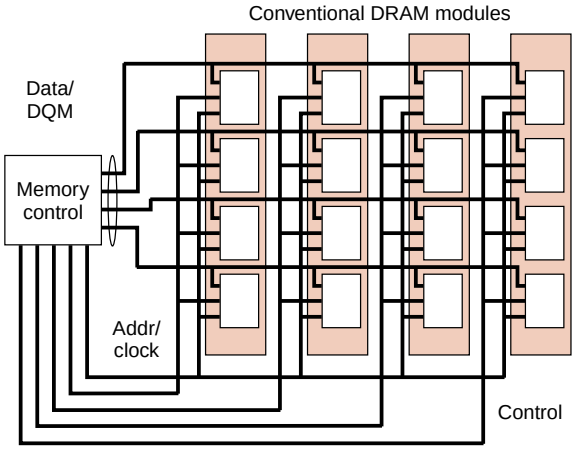

Figure 1. Nets in a conventional DRAM system.

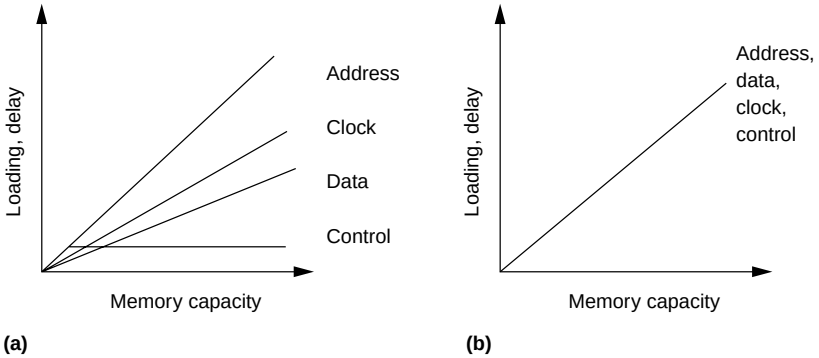

Figure 2. Signal loading versus memory capacity for SDRAM (a) and Direct RDRAM (b) systems.

Definitions

ASIC application-specific integrated circuit
CPU central processing unit, usually a microprocessor these days
CS,WE,RAS,CAS Chip select, write enable, row address strobe, column address strobe
CTT Center tapped termination, signaling technology implemented with totem pole MOSFET drivers operating in the voltage mode employing resistive parallel termination to half the power supply voltage
DDR Double data rate, a variant of SDRAMs in which the data pins transfer data on both clock edges instead of only on the rising edge without changing the sample rate of the other pins
Differential Signals represented in both true and complement form; that is, one is active high and the other is active low
DIMM Dual in-line memory module, a commonly used DRAM module type
DLL Delay-locked loop, subset of PLL (phase-locked loop) in which a variable delay element adds delay to a periodic signal such that its transition edges align with an external signal's transition edges. Used instead of using a variable frequency oscillator to accomplish such alignment or locking.
DQM Data mask signal used by SDRAMs to provide byte masking during write operations. There is one DQM signal for every 8 bits of data width.
DVD Digital versatile disk, formerly digital video disk
ECC Error-correcting code, a coding scheme using redundant logic states such that certain classes of errors that may arise can be automatically detected and corrected
EDO Extended data out, DRAM type that operates asynchronously while the output buffer remains in active state when CAS is pulsed; provides higher operating frequency than page-mode DRAMs with same process generation
GOP Gigaoperations per second, 1,000,000,000 operations per second
Granularity Minimum increment of memory that can be added to a system. If a system has a 64-bit bus, and 8M×8 bit wide devices are used, the granularity is 64 Mbytes.
Green PC A PC that has certain power-saving standby operating modes used to save power during periods of inactivity
I/O Input and output
JEDEC Joint Electron Devices Engineering Council, a committee comprised of semiconductor manufacturers that establish standards for certain electron devices
Jitter A measure of stability of a periodic time-varying signal
Latency A measure of the time it takes to complete a memory request. In a multiple master system, requests may be queued if more than one is allowed to exist at a time.
LVTTL Low-voltage transistor-transistor logic (TTL), commonly used chip-to-chip signaling technology, typically implemented with MOSFET totem pole output

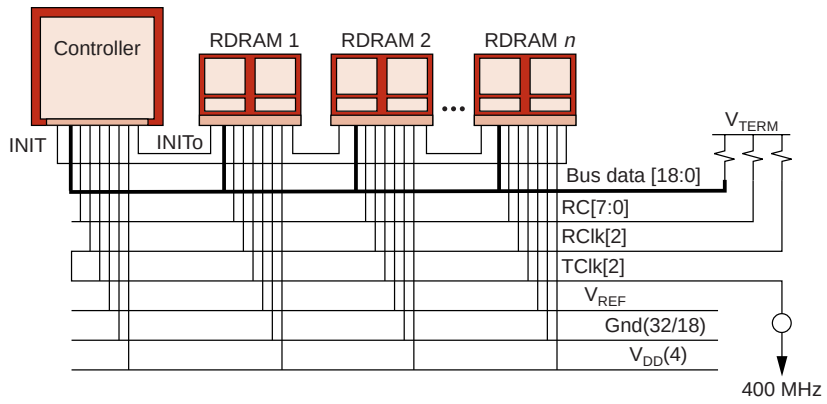

Figure 3. Direct RDRAM system.

as 3D graphics and DVD playback.

Direct Rambus technology

Our solution, the Direct Rambus DRAM (RDRAM), takes another approach that provides 1.6-Gbytes/s bandwidth from a single DRAM. It nears 95% efficiency when subjected to typical multimedia PC main memory workloads. Using a 16-bit data field and a separate 8-bit address and control field, a Direct RDRAM independently controls and schedules all row and column resources as well as I/O data. Direct RDRAMs, while using conventional PCB and connector technology, bring high speed and low-power operating modes to serve the needs of both line-operated and portable products.

Our technology uses a narrow bus topology operating at a high clock rate to solve the memory bandwidth problem. A Direct Rambus channel includes a controller and one or more Direct RDRAMs connected together via a common bus. The controller is located at one end, and the RDRAMs are distributed along the bus, which is parallel terminated at the far end (Figure 3). The signaling technology used for the high-speed channel signals is called Rambus signaling logic. Each RSL signal wire has equal loading and fan-out and is routed parallel to each other on the top trace of a PCB with a ground

structures. External signal voltage swings are typically equal to the power supply voltage, that is, 3.3 V.

Mbytes/s Megabytes per second, a unit of information bandwidth equal to 1,048,576 bytes per second

MOSFET Metal-oxide-semiconductor field-effect transistor, type of insulated gate transistors used in the majority of semiconductor microprocessors and memory devices, as opposed to BJT (bipolar junction transistor), JFET (junction field-effect transistor), or MESFET (metal-epitaxy-semiconductor field-effect transistor)

Open-drain driver Output driver stage consisting of at least one transistor connected between the power supply rail or ground and an output pin in which the transistor can only drive the pin to the same voltage rail to which the other node of the transistor is connected. Typically used in a bused wired-OR configuration using multiple parallel-connected drivers and an external resistor connected to a supply voltage. This way, any of the drivers may drive the bus low at any time without creating a fault condition that may arise if any two drivers attempt to drive the bus to opposite states at the same time.

Page mode DRAM type that operates asynchronously and permits output buffer transitions into a high impedance state when CAS is negated

Parallel termination. Method used to terminate a transmission line that uses a resistor with a value equal to the line's characteristic impedance. The resistor is connected between at least one end of the line and an fixed voltage supply such that direct current (DC) will flow in at least one logic state.

PCB Printed circuit board

PLL Phase-locked loop, closed-loop feedback system in which a variable frequency oscillator is adjusted until its output frequency is locked to a separate reference frequency

RSL Rambus signaling logic, signaling technology used by Rambus DRAMs implemented with open-drain MOSFET pull-down drivers operated in the current mode. Uses resistive parallel termination to a termination supply voltage greater than or equal to the high signal level.

Series termination. Method used to terminate a transmission line that uses a resistor in series with the line driver such that during the switching transient, a voltage is developed across the line. Used as the sole termination technique and has no DC component.

Single-ended Signals represented in one form only such as in true or in complement form; the opposite of differential

SSTL Stub series terminated logic, a variation of CTT that uses damping resistors in series with all outputs to minimize the influence of the impedance of the stubs connecting the bus drivers to the bus and to increase the output impedance of the drivers)

Totem pole driver Output driver stage consisting of at least two transistors in series connected between the power supply voltage rail and ground. The intermediate node connects to an output pin such that one transistor drives the output pin high and the other drives it low.

plane located on the layer underneath.

The channel uses 18 data pins cycling at 800 Mbps per pin to provide a low-cost, high-bandwidth bus. The bus requires a total of 76 controller pins including all signals and power supply pins. The Direct Rambus controller interface cell (RAC) (Figure 4) provides 1,600 Mbytes/s of bandwidth and uses slightly more than half as many pins as the 64-bit SDRAM or DDR interface found on PC core logic chipsets today.

Because a single Direct RDRAM spans the entire width of the channel, the loading on each bus pin increases uniformly as memory is added (Figure 2b). This ensures a constant timing relationship between signal pins independent of the total loading. This key property of the Rambus channel permits much higher frequency operation than the matrix topology used by SDRAMs.

Because a full-bandwidth channel can be constructed using a single Direct RDRAM, the minimum memory granularity is a single chip. This characteristic is growing in importance as DRAMs progress to the 256-Mbit generation and beyond.

Direct RDRAM architecture

The Direct RDRAM has a pipelined microarchitecture (Figure 5, next page). Like its predecessors, it has a wide internal bus connected via a high-speed interface to a narrow external bus.

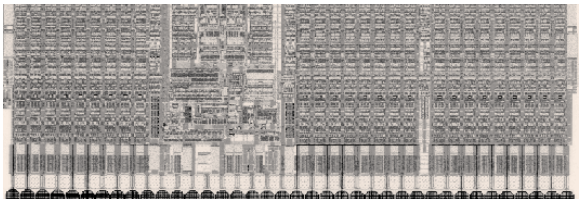

Figure 4. Direct Rambus controller interface cell (4 × 1.6-mm, 0.35-micron technology).

The narrow on-chip bus is serialized and deserialized to provide a 144-/128-bit data path into the core, which provides 16 bytes every 10 ns internally. The Rambus interface transforms the 10-ns internal bus into an external 1.25-ns bus that is 2 bytes wide to yield the 1,600-Mbyte/s bandwidth.

Because the core timing requirements are no more difficult than those used for a 100-MHz SDRAM, Direct RDRAMs leverage existing cores and process technology. As a result, Direct RDRAMs are compatible with common semiconductor manufacturing technology to assure low cost.

The Direct Rambus channel includes an 18-bit-wide bidirectional data field and an 8-bit-wide field carrying commands and row and column addresses. Like its predecessors,

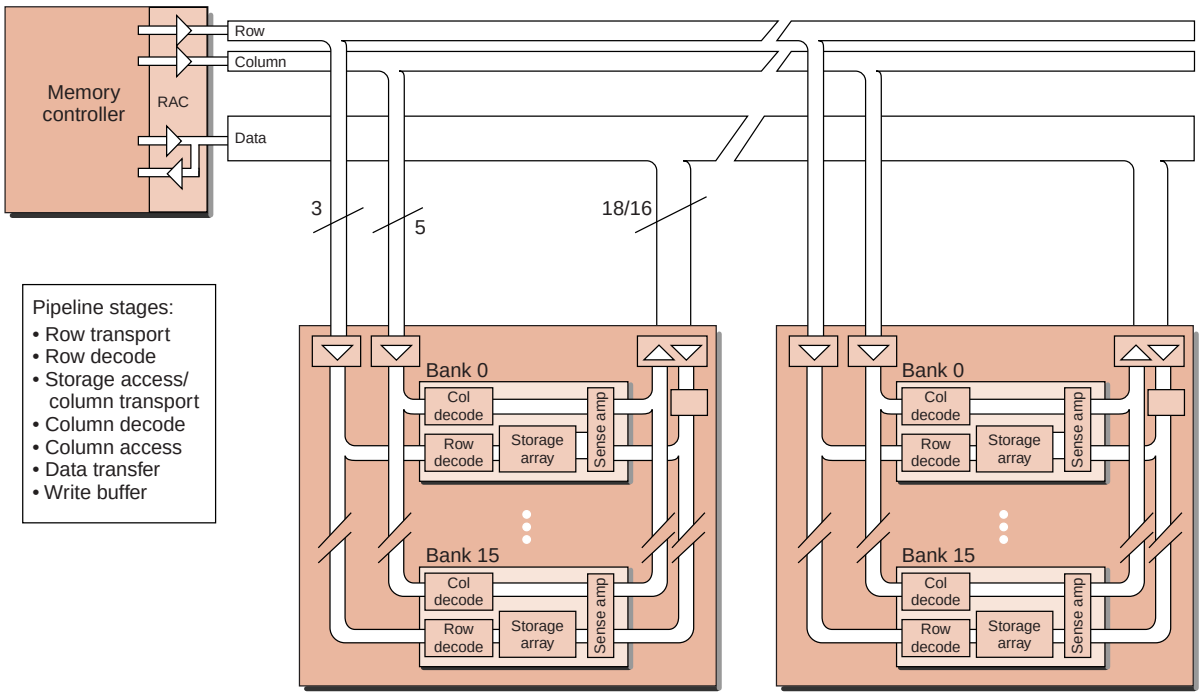

Figure 5. Direct RDRAM microarchitecture.

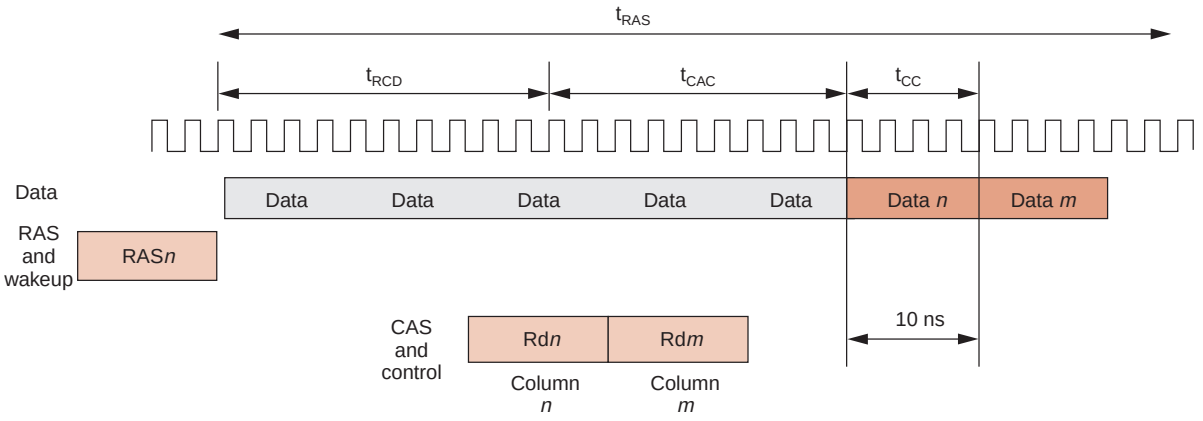

Figure 6. Direct RDRAM basic read cycle.

random column addresses can be supplied to the Direct RDRAMs while data is being transferred.[6] However the Direct RDRAM protocol introduces direct control of all row and column resources concurrently with data transfer operations (Figure 6)—hence the name "Direct."

Because the Direct Rambus protocol supports fully concurrent RAS and CAS operation in a pipelined microarchitecture that includes write buffering, each device can service up to four outstanding requests (Figure 7). Because of the streamlined microarchitecture, Direct RDRAMs avoid the empty time slots, or "bubbles," that frequently occur in single clocked SDRAM systems. Bubbles result from inadequate control bandwidth necessary to support page manipulation and scheduling while transferring data to and from random locations.[3] Doubled data rate schemes only aggravate the bubble problem.

The Direct RDRAM's high control bandwidth permits optimized data scheduling to provide approximately 95% efficiency over a wide range of workloads.[7]

Direct RDRAMs support explicit control of precharge and row-sensing operations as well as data scheduling during concurrent column operations. A Direct RDRAM can therefore perform row precharging and sensing operations concurrently with column operations to provide on-chip

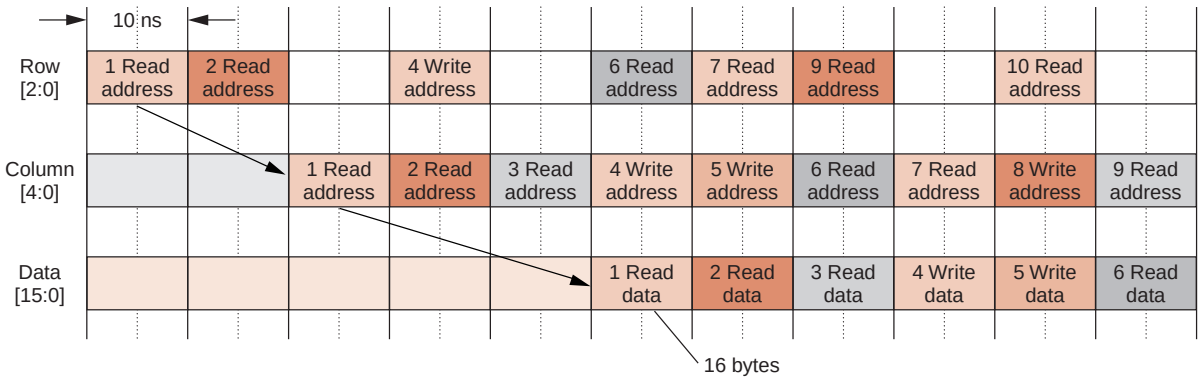

Figure 7. Direct RDRAM interleaved memory transactions at full-memory bandwidth (16 bytes/10 ns).

interleaving. Users can schedule the data resulting from the row operation to appear immediately after the column operation completes. This highly interleaved condition greatly improves the efficiency of the channel.

This interleaving can only happen when the requests target different banks in either the same Direct RDRAM or a different RDRAM on the channel. The more banks in a system, the better the chances are that any two requests are mapped to different banks. The more interleaving that is possible, the more the memory system performance improves.

The Direct RDRAM's memory array is divided into banks. To permit core design optimization, the specifications left the number of banks and the page size to the individual DRAM manufacturer. However all 64-Mbit Direct RDRAMs in development have 16 banks with a page size of 1 Kbyte.

In typical system configurations RDRAMs provide more system memory banks than SDRAMs or other conventional DRAMs on a per-megabyte basis. In bandwidth-intensive applications, several conventional DRAMs are frequently ganged together in parallel to provide the necessary aggregate bandwidth. Despite the fact that each DRAM may have four banks, this parallel combination does not increase the number of system banks; instead it just adds to their size.

Because a Direct RDRAM spans the entire channel, the CPU accesses each RDRAM independently. So each RDRAM directly adds to the number of memory banks accessible to the memory controller.

With a minimum of eight banks per Direct RDRAM, a 32-Mbyte system has at least 32 system banks. An SDRAM system of the same capacity constructed from 4-bank, JEDEC-standard 64-Mbit SDRAMs has only four system banks (Figure 8).

If in any DRAM the bank to be precharged is the same as the one being accessed, a bank conflict condition occurs, and bank precharging must be deferred until the current access completes. This precharge deferral results in diminished system bandwidth. Since an RDRAM system has more banks per megabyte than an SDRAM or DDR system, RDRAM systems boast lower bank conflict rates.

Statistically, the more banks that are included in a system, the less the probability of a bank conflict. When a Direct RDRAM is transferring data, any of the banks not being accessed can be precharged concurrently with column operations to provide a hidden precharge.

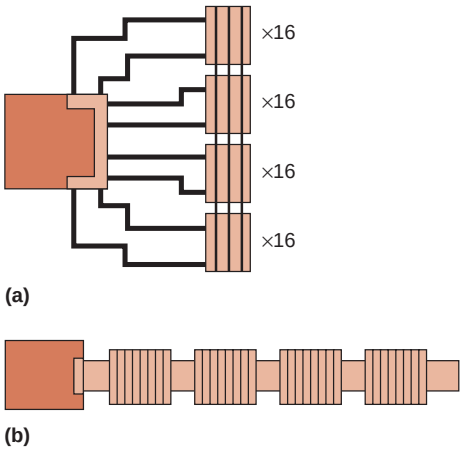

Figure 8. Bank counts: a 32-Mbyte, 64M SDRAM system with four large banks (a) versus a 32-Mbyte, 64M Direct RDRAM system with 32 small banks (b).

Error-correction support

Direct RDRAMs come in both 16-bit- and 18-bit-wide versions. The 18-bit-wide device can support 16-bit ECC over a 128-bit word without increasing the number of memory devices. Besides ECC or parity applications, a 9-bit byte effectively supports graphics and video applications by providing more bandwidth (12.5% increase over an 8-bit byte). It also supports multiple occluded-windows, antialiasing, *z*-buffer extensions, and alpha blending.

Backward compatibility

The Direct RDRAM incorporates the same physical layer as its predecessors (Rambus Base and Concurrent DRAMs). The primary differences are the channel width, which is 18 bits instead of 9, and the address and control information, which is no longer multiplexed onto the data field. Despite

these differences, designers can develop a Direct Rambus system logic controller that connects directly to a single Direct Rambus channel or to a parallel pair of Rambus concurrent channels. By offering an upwardly compatible design philosophy, this provides system designers with flexibility.

Rambus channel clocking

Like its predecessors, the Direct Rambus Channel has two different clock signals, which are separated at the far end of the bus and connected together at the controller. One of the clock signals is driven from the far end, and the other one is terminated at that end. This horseshoe arrangement means the clock passes through the array twice: once traveling toward the controller and once traveling away from the controller (Figure 3).

Each Direct RDRAM connects to both clocks and synchronizes its transfers to the clock traveling in the same direction as the information packet. Since each pin has the same loading, each signal wire has the same propagation velocity. This causes the phase relationship between the source clock and the packet to remain in lockstep as both propagate along the bus. Moreover if additional memories are installed, this key property is maintained.

Designers can maintain precise control over clock-to-data delay and bus sample points by compensating the internal clock skew and duty cycle with delay locked loops. DLLs allow all bus transfers to operate so that they are synchronized to both edges of a 400-MHz clock. This provides an 800-MHz data rate per pin.

Each Direct RDRAM and controller contains two DLLs for locking the internal clocks to the external clocks. One samples the input receivers, and the other triggers the output drivers. DLLs permit the internal clocks to avoid delays relative to the external reference clock (see the clock skew box). DLLs also ensure that the internal clocks have a 50% duty cycle. These properties greatly enhance the RDRAM system's operating frequency.

The transmit DLL introduces a 90-degree phase shift between the external clock and the data output signal. As a result, the channel data is precisely centered about the transmit clock. Since the receiving chip accepts the bus transmit clock as the receiving clock, this centering simplifies the receiving device's sampling task. The inherently low clock jitter provided by the DLLs[8] permits the system to operate reliably with a packet validity window of less than 300 ps for outputs and a sample of less than 200 ps for inputs.

RSL bus drivers and switching transients

Most bus drivers have either open-drain or totem pole output structures. Because a totem pole driver has at least two transistors connected to the output pin, it typically has higher pin capacitance than the single transistor of an open-drain driver. That increases bus loading, which increases system I/O power. RSL drivers are open-drain structures, while LVTTL, CTT, and SSTL are all totem pole structures.

Since totem pole drivers actively drive their output to either logic state, they dissipate power every time they switch. Open-drain structures actively drive the external signal in one direction only (usually low) and rely on external means to pull the bus high. So open-drain drivers only dissipate power in one of the two logic states. All else being equal, an open-drain driver has a 2:1 power advantage when compared to a totem pole driver.

Because an active RSL driver can only sink a current, the resulting bus voltage swing is a function of the driver current and termination resistance. Active output current regulation is incorporated into each RSL driver to assure that each

Compensating clock skews

Whenever a digital integrated circuit receives an external clock signal through its input pins, a finite delay occurs passing through the input receiver circuit. The delay for a conventional clock receiver will vary depending on the particular chip's process characteristics, operating temperature, and its power supply voltage. Noise on the chip's power supplies can also introduce delay uncertainty.

For synchronous inputs such as address, control, or data, a clocked input receiver is generally used. Every clocked receiver has a finite setup and hold time requirement referenced to the clock that triggers it (that is, the internal clock). If the internal clock has a range of delay (see Figure A), the required minimum validity window as referenced to the external clock will have to be stretched to assure the required minimum setup and hold time specifications are met. Likewise for output delays, the timing uncertainty of the internal clock triggering the output buffer will cause the data validity window to shrink and increase the time the output is in an indeterminate state (Figure B).

We can compensate for the delay through the clock input receiver (Figure C) in two ways. We can use open-loop schemes such as passing the buffered clock through controlled delay elements (that is, timing verniers commonly used in automatic test equipment) or use closed-loop circuits such as PLLs or DLLs. In each case, the basic idea is to have the circuit exploit the periodic nature of the clock by adding enough delay to the clock path so that the transition point of the compensated internal clock is aligned with the time of the external reference clock.

Timing verniers are generally simpler to implement and often require less power than PLL or DLL circuits, but suf-

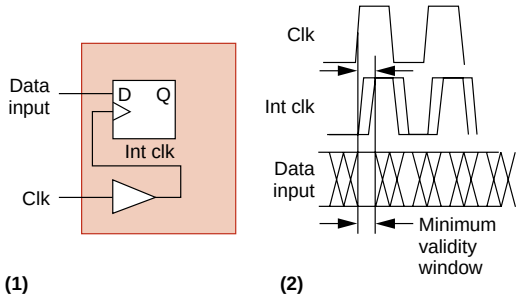

Figure A. Input sampling (1) with uncompensated clocks (2).

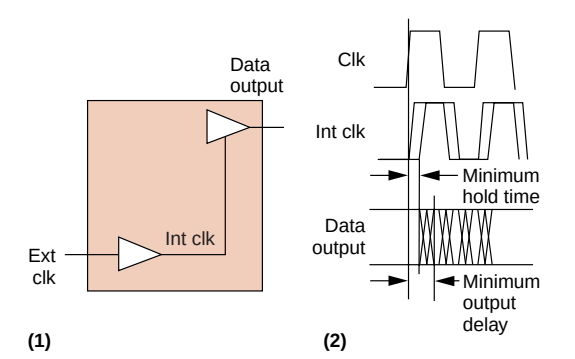

Figure B. Output drive windows (1) with uncompensated clocks (2).

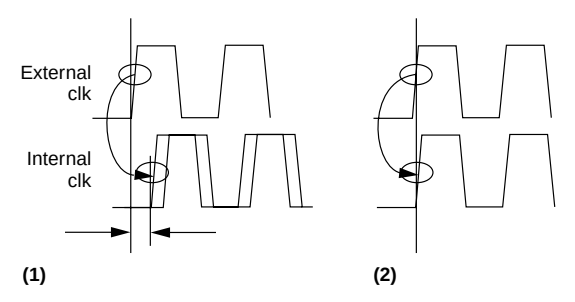

Figure C. Uncompensated (1) versus compensated (2) clocks.

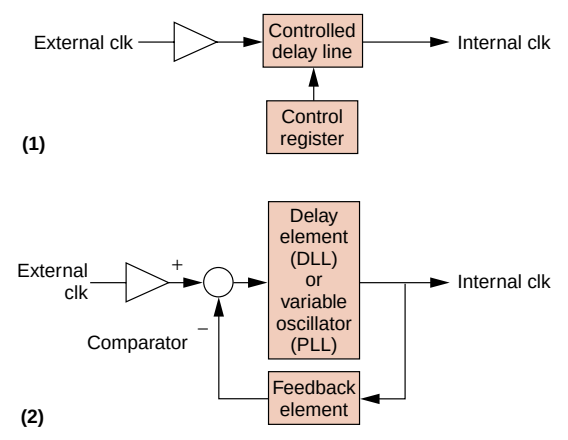

Figure D. Open-loop (1) versus closed-loop (2) clock compensation.

fer from sensitivity to temperature, power supply variations, and on-chip noise. They are classified as open-loop systems from a circuit topology perspective because there is no hardware feedback path used to compare the circuit's output with its input. A timing vernier's delay is usually set by a numerical value stored in a control register. The delay is therefore set manually, and it must be explicitly changed to account for drift of the circuits.

In applications such as automatic test equipment, the timing verniers are generally located within the test head. The test head environment is very stable from a temperature and power supply perspective. The result is that once calibrated, a timing vernier in an ATE environment is quite stable. However it may take as long as a half hour or more for a tester to stabilize when it is first turned on.

Because the power dissipated by a DRAM is related to its activity, application software or Green PC functions can affect the operating temperature of a DRAM used in a computer system. Therefore, depending on dynamic behavior such as memory access patterns changing, a changing workload may require adjustment of the vernier. Each recalibration may take several hundred or several thousand nanoseconds and may result in poor performance of the vernier circuit in DRAM or memory-controller applications under real-world workloads such as 3D texture mapping or cycling into and out of power-saving standby modes.

In contrast, PLLs and DLLs are closed-loop circuits. PLLs and DLLs both feature a hardware feedback loop within the circuit that is used to continuously sample its output for comparison against its input. Therefore these circuits automatically adjust their timing parameters on a cycle-by-cycle basis (Figure D). The result is that PLLs and DLLs exhibit far less sensitivity to thermal and operating voltage variations than do open-loop circuits such as timing verniers. This makes PLLs and DLLs more suitable for high-speed memory and controller applications.

DLLs and PLLs are currently used on today's highest bandwidth production microprocessors, DRAMs, and SRAMs to ensure system reliability at the highest operating frequencies.

device drives the bus to the same signal level.

Because the Direct Rambus channel operates in the current mode on a transmission line, the rate the driver is switched off determines the rise time of the bus. Applying the principle of superposition reveals that switching off the current mode driver is equivalent to sending an opposite polarity current wave back down the bus.

This reverse current wave cancels the DC already flowing in the line as the wave front propagates along the bus. Due to the finite impedance of the bus, the reverse current wave induces a voltage wave that adds to the voltage at each point during this transient. The more rapidly the current is switched off, the faster the rise time of the induced voltage step. Steady state is reached when the voltage wave reaches the termination resistor and the line is fully charged to the termination supply.

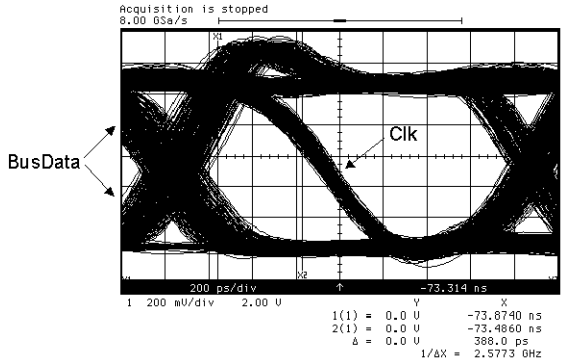

Figure 9. RDRAM eye diagram.

In a transmission line environment such as used by RSL, the rise time of the signal pin is set by the rate its MOSFET pin driver turns off instead of being set by the R-C time constant of a lumped circuit. Therefore, RSL-based systems not only have symmetric rise and fall times on a Rambus channel, but they also offer a power/bandwidth advantage over totem pole drivers such as LVTTL or SSTL. Figure 9[9] is an eye diagram showing clock and data transfers measured on an operating Rambus channel. Note that the bus data waveforms have nearly identical rise and fall times.

When an RSL bus line is driven from the controller end, the signal develops a full voltage swing on its incident wave and maintains its magnitude as the wave travels down the channel to be absorbed by the termination resistor. When the channel is driven by an RDRAM in the middle section of the channel, the effective bus impedance is the parallel combination of the two bus segments connecting the RDRAM to the controller and the RDRAM to the terminated end. Since this parallel combination of transmission line segments results in a 50% reduction of load impedance, the RSL driver can only drive a half-amplitude incident wave. This wave splits into two components with one propagating toward the controller and the other toward the terminator end. The wave that travels toward the terminator end is absorbed upon arrival, but the one traveling toward the controller reflects.

Since the controller end is unterminated, the signal's voltage reflection coefficient is approximately +1, and the current reflection coefficient is approximately −1. Therefore when the incident wave reflects off the controller end, the magnitude of the voltage at the controller pin doubles, causing a full level to form at the pin in a single transition. The reflected wave, also a half step, then travels back down the channel in the opposite direction toward the controller end. The behavior is the same for rising or falling transitions.

When the reflected wave reaches the initiating driver, the high dynamic output impedance of the current source driver combined with the very short signal stub associated with the packaging and/or module represents a negligible impedance discontinuity on the channel. As a result, there are no further significant reflections, and the signal passes by the driver undisturbed. It travels to the terminator end where it is finally absorbed.

These switching characteristics permit the Direct Rambus channel to be terminated at only one end. The single-ended termination gives RSL a significant DC power advantage over double-ended terminated buses such as CTT or SSTL.

Finally, applying the principle of superposition reveals that the Direct Rambus channel is inherently immune to the potential inter-symbol interference resulting from the reflected waves that interact with subsequent incident ones.

RSL I/O levels

All high-speed Direct Rambus signals are single-ended except the clocks, which are differential. The high level for the bus is 1.8 V, and the low level is 1.0 V. A common bused reference voltage of 1.4 V feeds into all input receivers, which are differential amplifiers. Because it has fewer simultaneously switching signals, each Direct RDRAM or controller experiences less ground noise than other high-bandwidth offerings.

Since the channel can only be actively driven low, the RSL drivers do not induce large V_{cc} noise transients such as seen with totem pole drivers. Additionally, the input receivers are more immune to noise than conventional input receivers due to both the relatively high bias point of the external bus and the high common mode rejection inherent in well-designed differential amplifiers.

Power modes

Direct RDRAMs include power management modes to address the needs of both the environmentally protective Green PC and the portable computing markets. A low-latency transition from the low-power standby state to the active condition assures high system performance when using power-saving modes. The result is the Direct RDRAM is well suited for both portable and line-powered green applications.

Memory expansion

Systems based on Direct RDRAMs can be readily upgraded via the use of memory modules. Bearing a strong appearance to traditional DIMMs, Direct Rambus DRAM modules (RIMMs) fit into sockets similar to standard DIMMs. Therefore, they fit within the standard mechanical and thermal envelope of all modern industry-standard PC chassis configurations.

Despite their similar appearance to DIMMs, RIMMs are fundamentally different. Instead of being connected in parallel, RIMMs are connected in series when installed in a system.

Because each Direct RDRAM uses only 30 high-speed signals and is connected in a bused topology, both ends of a Direct Rambus channel segment can be routed to a single edge of the RIMM while maintaining matched signal electrical lengths. Despite this unusual physical configuration, a RIMM fits completely within the footprint of the 168-pin DIMM in common use today (Figure 10) using standard connector technology. Yet it operates at 800-MHz data rates (Figure 11).

Because RIMMs are connected in series, the primary current-carrying conductors are routed on the RIMMs, not on the motherboards. This reduces the signal stub length to essentially the same length as the Direct RDRAM package's signal lead. In effect, an RDRAM mounted on a RIMM behaves as if it were soldered directly to the motherboard instead of being

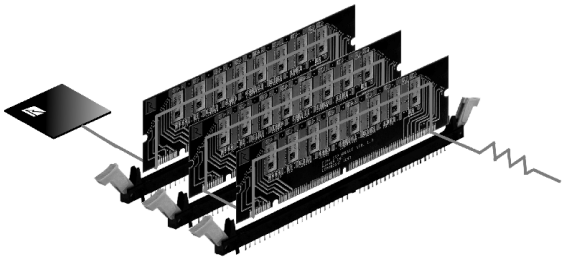

Figure 10. RIMM subsystem used on a motherboard.

connected via long traces through a DIMM socket.

This key difference virtually eliminates signal integrity concerns associated with DIMMs because there are no long stubs on a RIMM. (As mentioned earlier, the long stubs of DIMMs cause signal reflections if left unterminated.)

The low inductance offered by modern low-profile socket technology makes the series connection of the RIMMs practical, even at their unprecedented data rates. RIMM sockets use all the manufacturing infrastructure of their DIMM cousins to assure low cost and plentiful supply.

PCB physical design

The physical design for any Rambus PCB is quite simple. The channel requires only two PCB conductor layers. All signal traces are located on the top layer. A ground plane located directly underneath forms what RF engineers call a microstrip. The mass-produced Nintendo 64 game console uses base RDRAMs on a two-layer motherboard PCB operating at 500 MHz.

The Rambus channel requires a uniform value of loaded impedance at any point along the channel. Because the pin capacitance of the RDRAM dominates the PCB impedance and propagation velocity, Rambus systems tolerate the customary high-volume PCB manufacturing tolerances of +/– 15%. PCB design therefore consists of following simple common-sense rules to determine PCB thickness and the trace widths, spacing, and thickness.

Direct RDRAMs are packaged in a low-cost chip-scale package called a micro ball grid array (µBGA). The inherently small size of the µBGA minimizes the materials required for its manufacture while providing outstanding electrical and thermal properties.

Sourcing and manufacturing infrastructure

Intel has selected the Direct Rambus technology to become its next PC main-memory standard and plans mass-market shipments in 1999.[3,10]

All of the world's top 13 DRAM suppliers are in active development of Direct RDRAMs. There also are four major ASIC suppliers producing Rambus ASIC technology. Joining these suppliers are a number of leading semiconductor peripheral component manufacturers developing Direct Rambus controllers.

Besides semiconductor suppliers, a host of memory module manufacturers, connector, clock chip, and test equipment manufacturers have rallied behind the Direct Rambus technology. As a group they will be supplying key manufacturing infrastructure components to assure rapid deployment of the technology and ease the integration of high-performance PCs and other products.

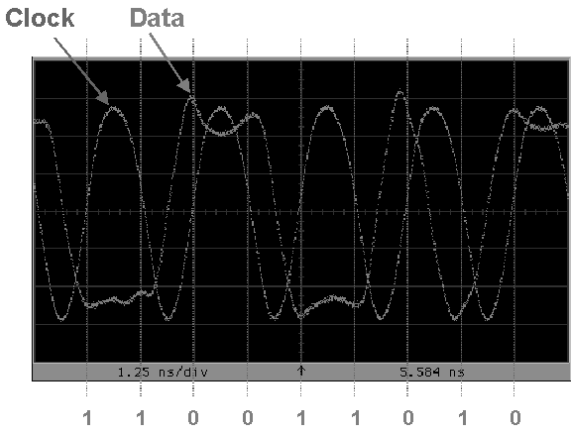

Figure 11. A 32-RDRAM RIMM with waveforms measured at a 800-MHz data rate.

Included in the list of module makers are Fujitsu Ltd., Hitachi Ltd., Hyundai Electronics Industries Co., Kingston Technology, Co., LG Semicon Co. Ltd., Mitsubishi Electric Corp., NEC Electronics Inc., Samsung Electronics Co. Ltd., Smart Modular Technologies Inc., and Texas Instruments Inc. Connector makers include Berg Electronics and several unannounced suppliers. Clock-chip makers include Cypress Semiconductor Corp., Integrated Circuit Systems Inc., IC Works, International Microcircuits Inc., and NEC Electronics Inc. Finally, semiconductor automatic test equipment manufacturers include Advantest Corp., Hewlett-Packard Co. and Teradyne Inc.

For background information on the Direct Rambus, visit our Web site at www.rambus.com. Among other things, it contains Allen Roberts' presentation at the October 1997 Microprocessor Forum in which he discussed the Direct RDRAM architecture.

DIRECT RAMBUS TECHNOLOGY offers significant performance and system integration advantages over older-generation EDO/SDRAM technology. In addition to performance gains, Direct Rambus technology offers longevity. It applies to a minimum of three generations (64 Mbit, 256 Mbit, and 1 Gbit). Furthermore, we plan to increase flexibility in the use of the technology with half-generation devices (32 Mbit, 128 Mbit, and 512 Mbit). Each device is fully interchangeable using a common pinout. Besides density scaling, future plans include enhancing the operating frequency.

For background information on the Direct Rambus, visit our Web site at www.rambus.com. Among other things, it contains Allan Roberts' presentation at the October 1997 Microprocessor Forum in which he discussed the Direct RDRAM architecture.

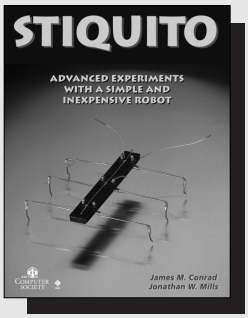

Stiquito™

Advanced Experiments with a Simple and Inexpensive Robot

by James M. Conrad and Jonathan Mills

The Stiquito robot is an small, inexpensive, six-legged robot that is intended for use as a research and educational tool. This book, describes how to assemble and build Stiquito, provides information on the design and control of legged robots, illustrates its research uses, and includes the robot kit. The experiments in the text lead the reader on a tour of the current state of robotics research. The hobbyist with some digital electronics background will also find this book challenging.

Contents: Stiquito Introduction and History
• Walking Robots • Control of Walking Robots
• Using Stiquito for Research • The Future of Stiquito

328 pages. 7" x 10" Softcover. Nov. 1997. ISBN 0-8186-7408-3.
Catalog # BP07408 — $28.00 Members / $35.00 List
Price includes the robot kit

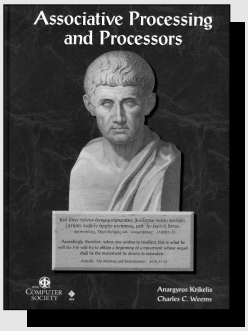

Associative Processing and Processors

Anargyros Krikelis and Charles C. Weems

Covers recent research on associative processing and processors and details the unique features they offer for cost-effective system solutions. The book explores the distinct advantages that associative processing systems have over other parallel processors. The text illustrates associative processing techniques for both traditional architectures and architectures that support multiassociative processing. It details several general purpose associative processing architectures, as well as a dataflow architecture that uses associative processing elements.

Contents: Associative Processing • Associative Processing Architectures and Systems • Associative Processing Software • Associative Processing Applications • Associative Processing and Neural Computation

360 pages. 8fi" x 11" Softcover. July 1997. ISBN 0-8186-7661-2.
Catalog # BP07661 — $40.00 Members / $50.00 List

Call toll-free:
+1.800.CS.BOOKS
Online Bookstore:
http://computer.org

IEEE COMPUTER SOCIETY

References

1. S. Purcell, "A VLIW and SIMD Vector Processor for PC Multimedia," Microprocessor Forum, Oct. 1995; reprints available from Chromatic Research Inc., Sunnyvale, Calif.
2. J. Forbes, "DVD Technology, New MPUs, AGP, Media Processors, and New Audio Strategies Spell Change," *Electronic Buyer's News,* July 21, 1997; http://www.techweb.com/se/directlink.cgi?EBN19970721S0063.
3. W. Hunting, "PC SDRAM Platform Considerations," Intel Developer's Forum, Sept. 30, 1997; CD-ROM available from Intel Corp., Santa Clara, Calif.
4. R. Crisp et al., "Physical Design Comparison of 128 bit SGRAM-Based and Rambus-Based 3D Controllers," to be presented at the Dec. 1997 Multimedia Technology and Applications Conf.; contact author for copies.
5. M. Lapedus, "VIA Chipset Supports DDR SDRAM," *Electronic Buyer's News,* Sept. 22, 1997; http://www.techweb.com/se/directlink.cgi?EBN19970922S0153.
6. "16/18Mbit (2Mx8/9) and 64/72Mbit (8Mx8/9) Concurrent RDRAM Data Sheet," DL0029-05, Rambus Inc., Mountain View, Calif., Jan. 1997.
7. "64M/72M Direct RDRAM Data Sheet," DL 0035-00.c0.5.28, Rambus Inc., May 1997.
8. T. Lee et al., "A 2.5V DLL for an 18Mb, 500MB/s DRAM," *Int'l Solid-State Circuits Conf. Digest of Tech. Papers,* IEEE, Piscataway, N.J., Feb. 1994, pp. 300-301.
9. R. Crisp et al., "Development of Single Chip Multi-GB/s DRAMs," *Int'l Solid-State Circuits Conf. Digest of Tech. Papers,* slide supplement, Feb. 1997, p. 177.
10. Intel World Wide Web site, http://developer.intel.com/design/pcisets/rambus.htm.

Richard Crisp has been responsible for DRAM development, applications engineering, and business development at Rambus, Inc. His current interests involve new product development exploiting Rambus technology. Crisp is a graduate of Texas A&M University and a member of the IEEE, Eta Kappa Nu, and Phi Eta Sigma. Currently serving his second year as Memory Program Subcommittee chair for the International Solid-State Circuits Conference, he has been a program committee member since 1990.

Direct questions or comments concerning this article to Richard Crisp, Rambus Inc., 2465 Latham Street, Mountain View, CA 94040; crisp@rambus.com.

Reader Interest Survey

Indicate your interest in this article by circling the appropriate number of the Reader Interest Card.

Low 156 Medium 157 High 158