

Slide #01

**Direct Rambus™
Architecture and
Measurements**



Slide #02

Agenda

- *Direct Rambus Overview*
- **Direct Rambus Channel Electrical Architecture**
- **Direct RDRAM Device Physical Architecture**
- **Direct Rambus Logical Architecture**
 - **Read and Write Transaction Protocol**
 - **Transaction Pipeline**
 - **Packet Format**
 - **Unique Concepts**
 - **Maintenance Operations**
 - **Configuration**
- **Measurements made with HP Instruments**

This seminar will give you a good understanding of the Rambus system architecture. The overview will help you become familiar with the top-level system architecture. Keeping in mind that the Rambus channel is a transmission line, the electrical architecture portion of the presentation describes the electrical characteristics of this high-speed and high-bandwidth interface. The physical architecture discussion describes the architecture of the RDRAM chip while explaining novel features not supported in traditional DRAM's. The logical architecture portion of the presentation describes the read/write operation and the transaction packet formats. Unique concepts in the RDRAM architecture are explained. A brief understanding of the configuration process is presented.

Finally, we will show you measurements made on the Rambus channel using HP test instruments. We show you how to check for signal integrity and transmission line characteristics. Using a HP logic analyzer, we will demonstrate the strategy for following a transaction originating on the CPU or PCI bus until the transaction reaches the Rambus channel. The inverse assembler capability of the logic analyzer is also demonstrated.

Slide #03

Direct Rambus Overview

- **1.6 GB/s peak (up to 1.5 GB/s effective) data transfer rate per channel**
- **>95% channel utilization with transaction pipelining**
- **32 RDRAM devices per channel**
- **256 MB memory per channel using 8MB RDRAM devices**
- **DRAM-like architecture with multiple banks**

RDRAM memory is the highest bandwidth system memory available today. Direct RDRAM channel data bus width is 2 Bytes with an optional parity/ECC bit per Byte of data. The channel clock frequency is 400 MHz, but data is clocked on both edges of clock. Assuming 100% utilization of the data bus, the peak bandwidth is 1.6 GB/s.

Because fewer pins are required on the interface, the design is comparatively cost effective and consumes less board real estate. Transactions on the channel are pipelined, resulting in greater than 95 % bus utilization.

RDRAM devices offer 8 MB and higher storage capacity. Each channel can support 32 RDRAM chips resulting in memory capacity of 256 MB and higher per channel.


RDRAM devices contain a large (16 in our example) number of banks. As a result, assuming random access to the address space, the effective data transfer rate is in the order of 1.5 GB/s. The overhead for maintenance type transactions (which includes refresh, current calibrate etc) is small.

In summary, the Rambus memory technology provides a high bandwidth, high memory capacity, and cost-effective system memory solution.

Slide #04

Direct Rambus Overview

- Topology of the Rambus Channel



This diagram is a topology view of the Rambus channel showing the memory controller, the connectors, RIMM modules and RDRAM devices. Notice that there are 3 RIMM modules, the maximum allowed. Each RIMM can have up to 16 RDRAM devices, but a maximum of 32 devices are supported per channel.

A connector cannot be left empty. A continuity module is used to maintain the continuity of the channel which starts at the memory controller (RMC) and ends with termination at the opposite end.

Slide #05

Direct Rambus Overview

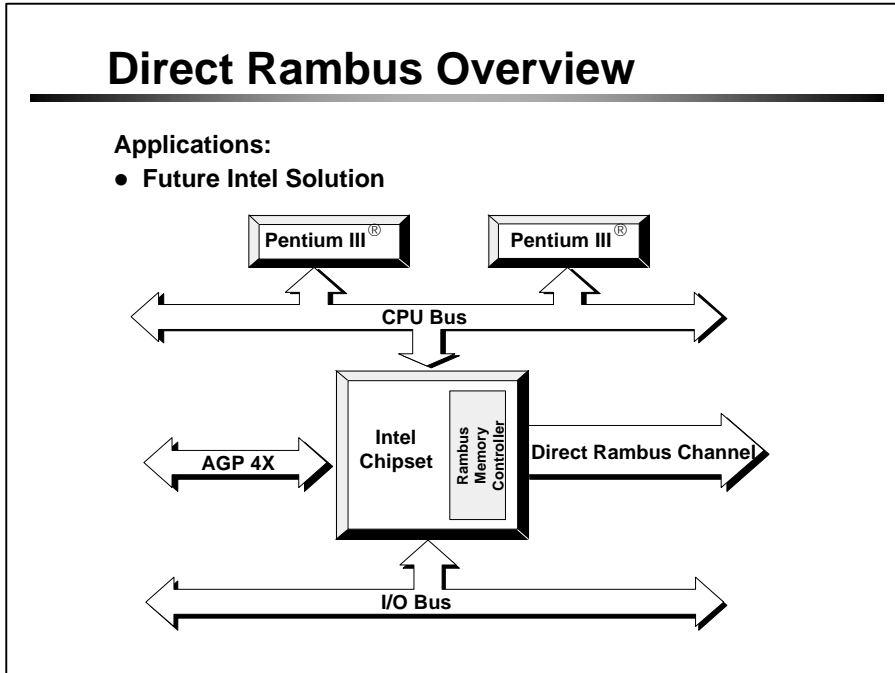
- Complete system design consists of:
 - Direct RDRAM™ Device
 - RIMM™ modules
 - Connectors
 - Clock Generator (DRCG)
 - Rambus ASIC Cell (Direct RAC)
 - Rambus Memory Controller (Direct RMC)

Rambus memory system design consists of the components listed in the slide above.

- An example of a Direct RDRAM device is the 64-Mb (8-MB) memory chip. It contains 16 banks of memory cells. Each bank contains 512 pages of memory. Each page contains 64 dualocts of storage. A dualoct is defined as 16 bytes, which is the smallest addressable entity in the RDRAM device.
- RIMM modules hold up to 16 RDRAM devices. A channel can be designed with up to 3 RIMM modules. However, each channel can hold a maximum of 32 RDRAM devices. A RIMM module without RDRAM devices is referred to as a continuity module and is used to fill an empty connector.
- The connector has a DIMM connector-like form factor. Notice the two groups of pins separated by some space. One group corresponds to the pins associated with the channel coming into the RIMM module. The other set of pins corresponds to the channel leaving the module.
- The Direct Rambus Clock Generator (DRCG) produces the 400 MHz differential clock. Data and clock travel together.
- The Rambus ASIC Cell (RAC) is the physical analog interface to the channel. It encodes/decodes packets of information transmitted over the channel.

- The Rambus Memory Controller (RMC) sits between the host system and the RAC. Both RMC and RAC are typically included in one silicon. The RMC generates the transactions necessary to communicate with the RDRAM devices. Both RAC and RMC reference designs are available from Rambus, Inc.

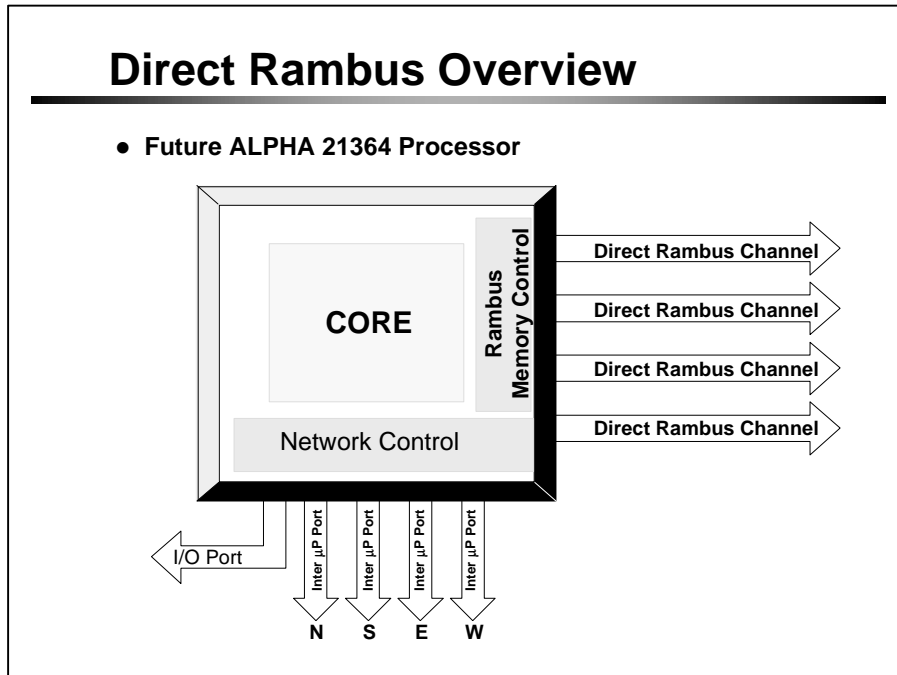
Slide #06



The above picture is a block diagram of a future Intel system architecture. A future Intel chipset will support the Pentium® III processor running at 133 MHz on the host bus. The processor requires a high bandwidth path to main memory. Therefore, Rambus memory architecture is the solution.

Because specific details about the chipset are not public-domain information, they cannot be discussed at the time of this presentation.

Slide #07



The above block diagram depicts the future 21364 Alpha processor due out in the year 2000. It will run at internal core frequencies greater than 1 GHz. The processor is system memory bandwidth hungry and needs access to large amounts of memory. Rambus system memory architecture is the solution. The processor accesses four channels, resulting in a peak bandwidth of 6.4 GB/s and more than 1 GB memory storage capacity.


The processor can be connected in an array fashion to four other processors for a maximum of 16 processors in the system. The memory associated with each processor can be stored in cache and shared by each processor. The Inter Processor Ports are used for inter-processor communication. The I/O port is the interface to peripheral buses through a chipset.

Slide #08

Agenda

- **Direct Rambus Overview**
- *Direct Rambus Channel Electrical Architecture*
- **Direct RDRAM Device Physical Architecture**
- **Direct Rambus Logical Architecture**
 - **Read and Write Operations**
 - **Transaction Pipeline**
 - **Packet Format**
 - **Unique Concepts**
 - **Maintenance Operations**
 - **Configuration**
- **Measurements made with HP Instruments**

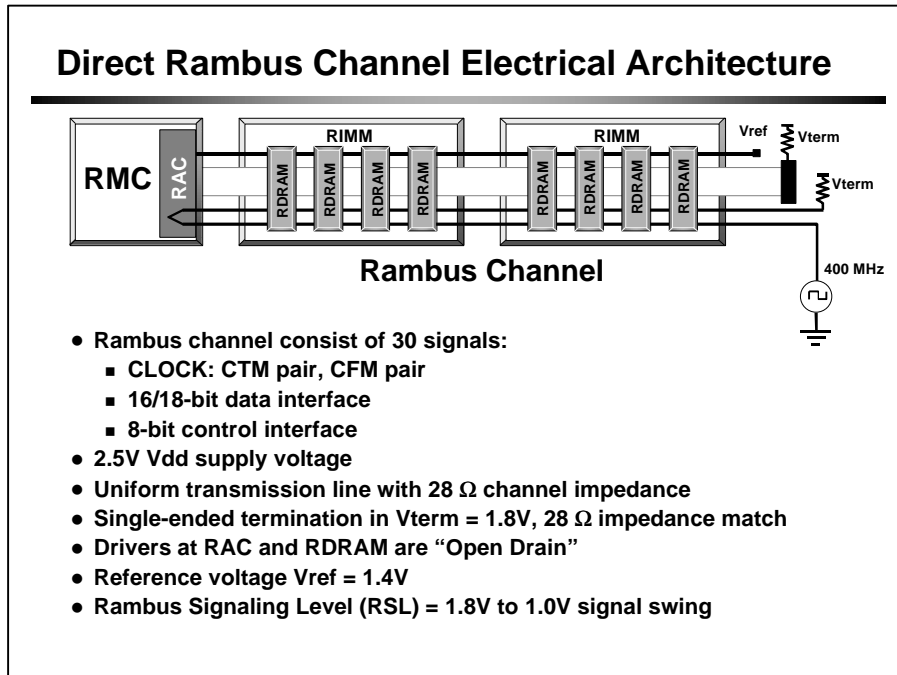
Slide #09



The block diagram above is a representation of a Rambus channel, showing all the components discussed previously, RMC, RAC, RIMM modules, RDRAM devices, DRCG.

Note the termination resistors that are impedance matched to the impedance of the transmission line and the reference voltage source.

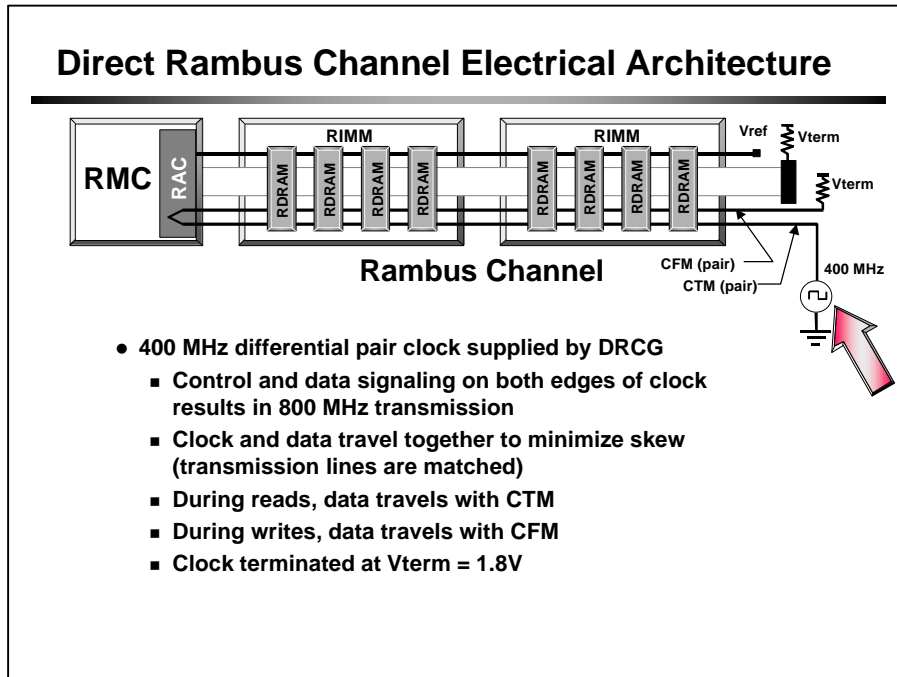
Slide #10



The Rambus channel consists of 30 signals: 18 (or 16 if parity/ECC is not supported) data lines, 8 control pins, and 4 clock signals. The channel impedance is a uniform 28 ohms, and stubs are not allowed. To guarantee that there are no reflections from any part of the channel with the exception of the RAC, strict design rules must be followed. The ground plane and signals are carefully routed to eliminate signal cross talk and ground bounce noise. Rambus, Inc. provides design rules to help you avoid these problems when laying out your board.

All signals are properly terminated at 1.8 V with 28-ohm resistors. Signal drivers are electrically “Open Drain” in current source mode which means that signals are asserted electrically low. All signals swing uniformly between 1.8 V and 1.0 V. The reference voltage (1.4 V) shown in the diagram provides a reference to comparators present on all inputs. When a signal is at 1.8 V it is considered to be at Logical 0 and when a signal is at 1.0 V it is considered to be at Logical 1. These signaling levels are referred to as “Rambus Signaling Level” (RSL)


Slide #11



The DRCG at the far end of the channel supplies the 400 MHz differential clock pair. Control and data signals are clocked on both edges of the clock resulting in an 800 MHz effective transfer. Clock and data travel together. Therefore, the skew between clock and data is kept to a minimum (< 100 ps). If the clock and data do not travel together, it is not possible to meet setup and hold times at this high of a frequency.

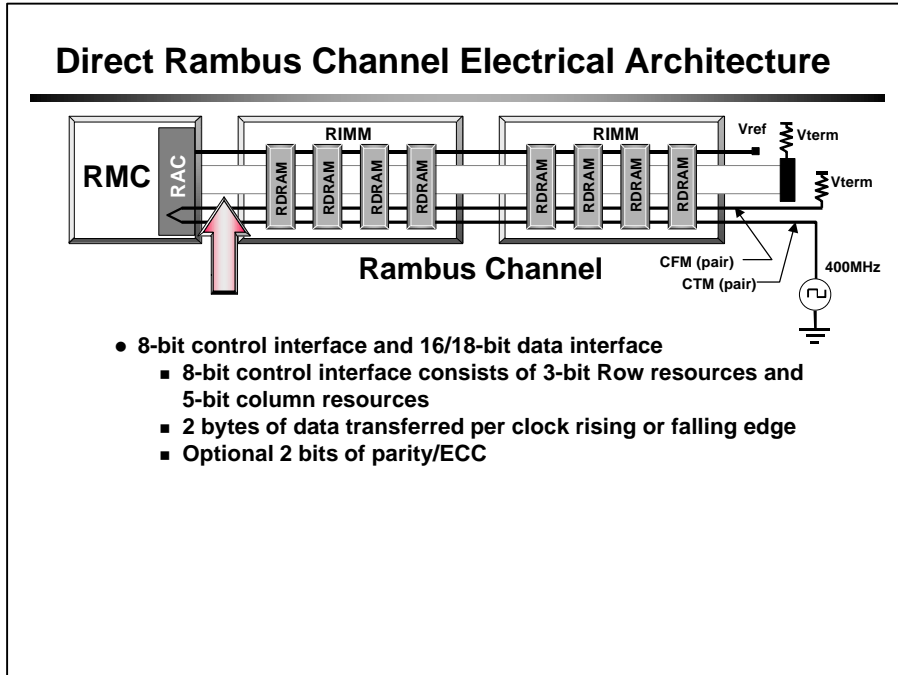
The clock is a differential clock pair. This means that when one clock signal is at 1.8 V, the other clock in the differential pair is at 1.0 V. The clock pair that originates at the DRCG and goes towards the RAC is labeled as clock-to-master (CTM/CTMN pair). Read Data from the RDRAM devices is driven on this clock and is clocked into the RAC using this clock. This clock is re-driven by the RAC and is terminated at the far end of the channel. This portion of the clock signal is referred to as clock-from-master (CFM/CFMN pair). Write Data and the Control signals are driven by the RAC on this clock and are clocked into the RDRAM devices using this clock.

Slide #12



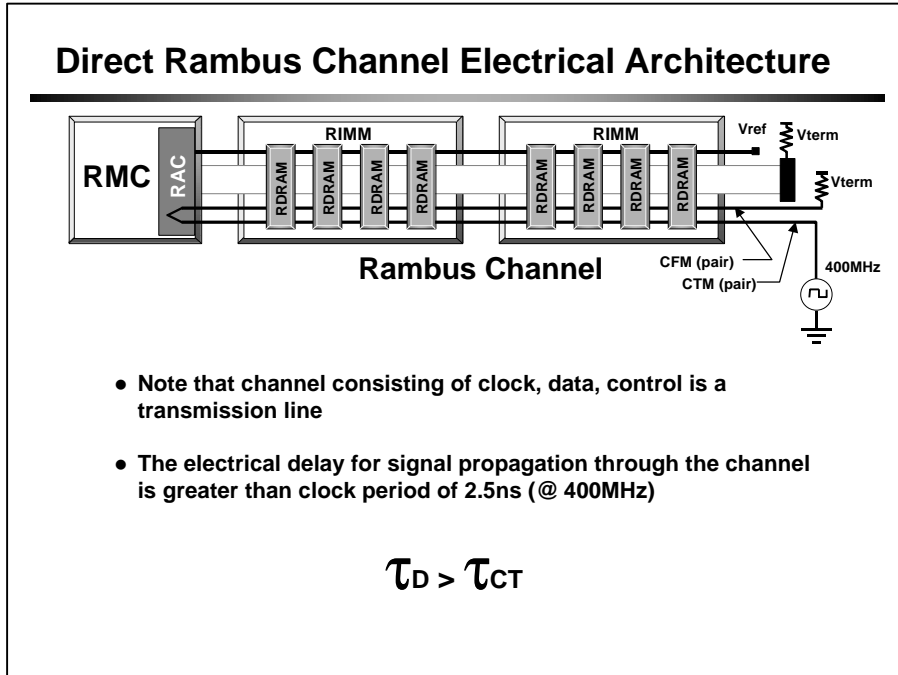
RSL signal swing is uniform around $1.4 V_{ref}$. V_{ref} is an input to all devices on the channel. When an RSL signal is asserted (1.0 V) the driver sinks ~ 30 mA of current. This current is uniform across RSL signaling levels as shown in the graph above. The drivers are current calibrated periodically as explained later in the presentation to ensure that the signal swing remains uniform over time.

Slide #13



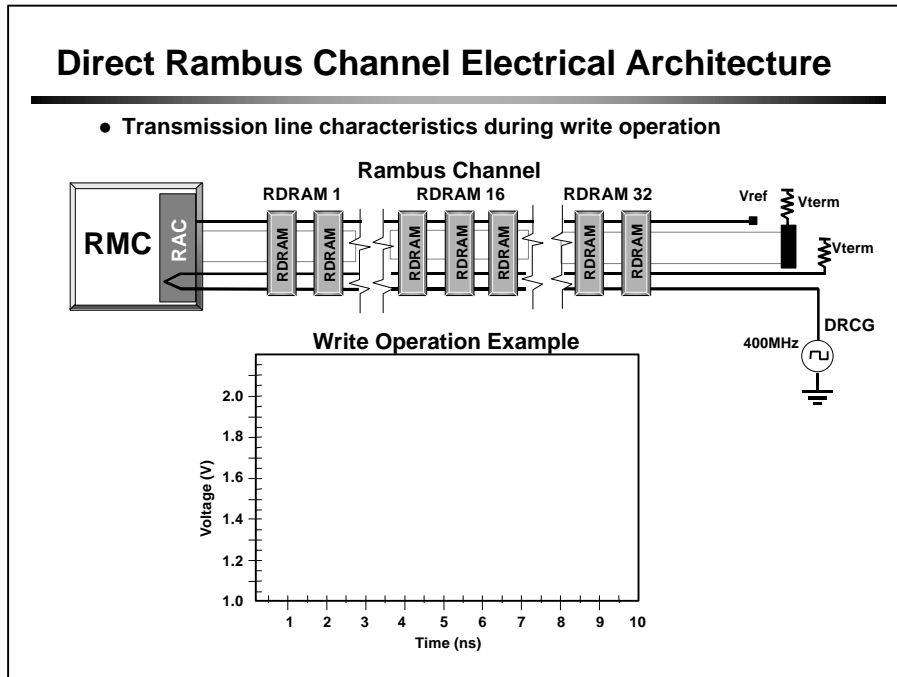
There are a total of 8 Control signals consisting of 3 Row resources and 5 Column resources. The Data bus consists of 2 bytes of data with optional parity/ECC bit per byte.

Slide #14



Note that the channel is considered a transmission line because the time that it takes a signal to propagate through the channel is greater than the clock period of 2.5 ns. The propagation time of a signal through the channel is 5-10 ns.

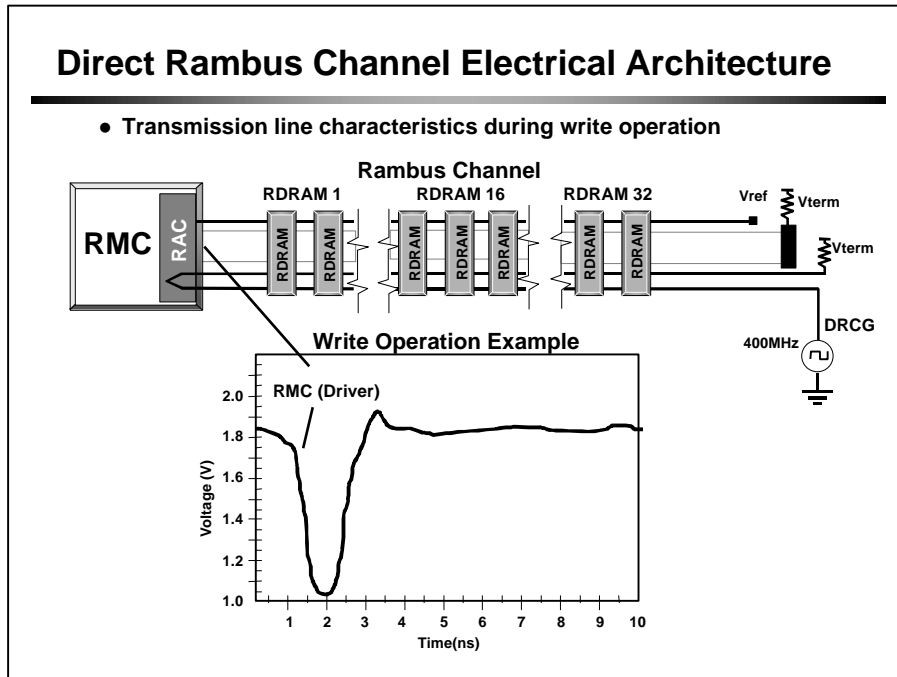
Slide #15



Transmission line during write operations

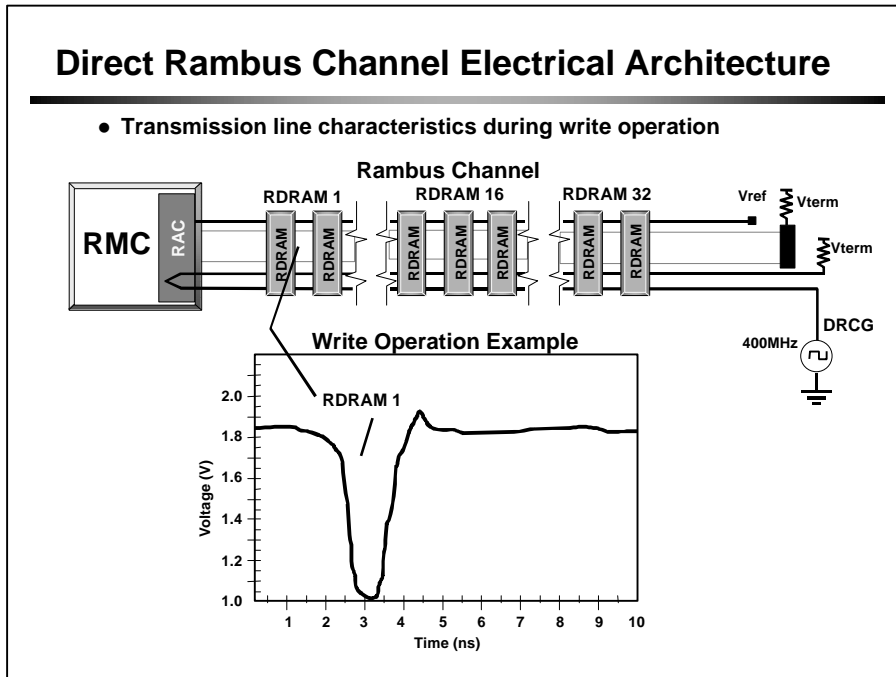
In these next set of slides, the transmission line characteristics observed on the Control signals and the Data signals during write operations is described. The RAC is driving the channel.

Slide 16



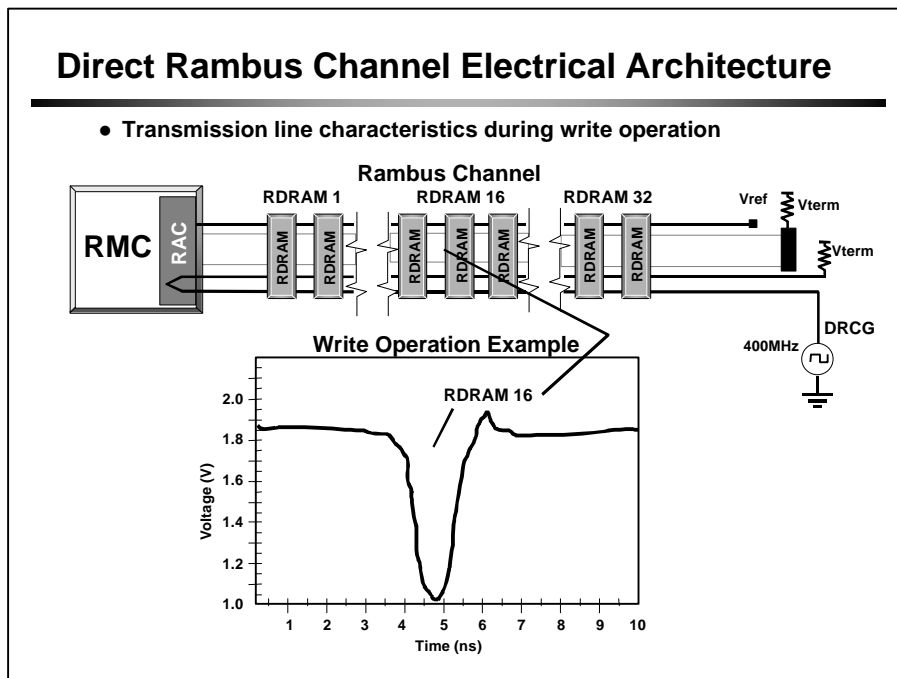
At time $t \approx 1.5$ ns, the RAC drives the signal low as shown. This signal will propagate down the transmission line

Slide #17



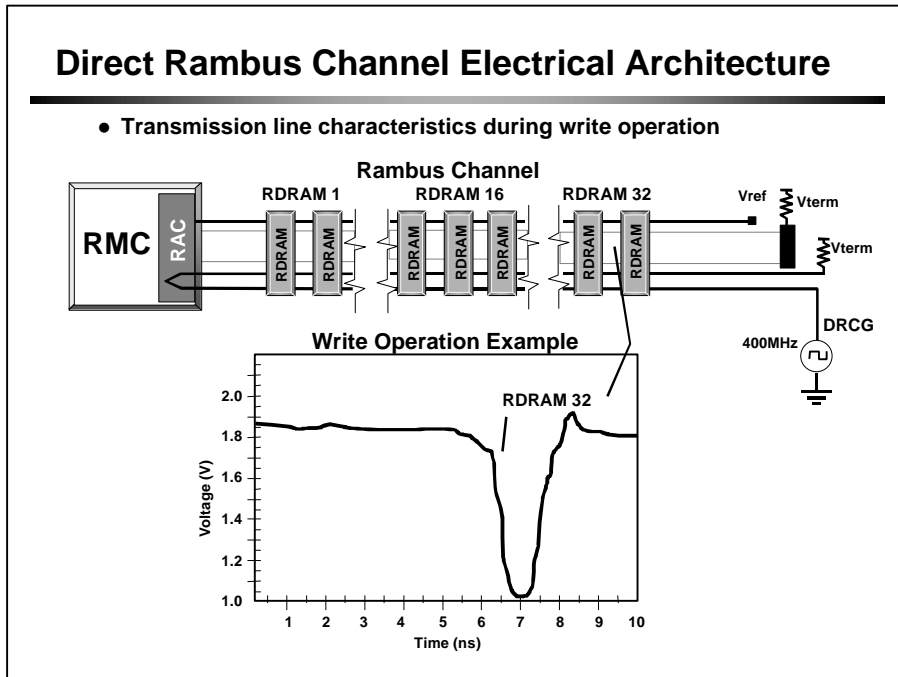
At time $t \approx 3.0$ ns, the signal reaches the first RDRAM device. The transmission line delay is about 1.5 ns at this point.

Slide #18



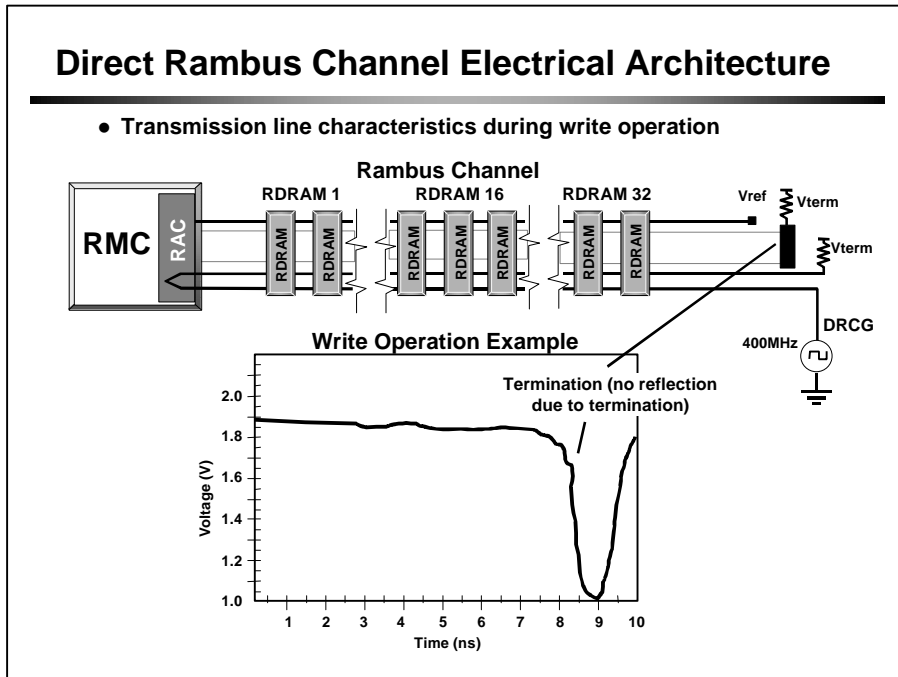
At time $t \approx 5.0$ ns, the signal reaches RDRAM device 16, meaning that the transmission line delay thus far is about 3.5 ns.

Slide #19



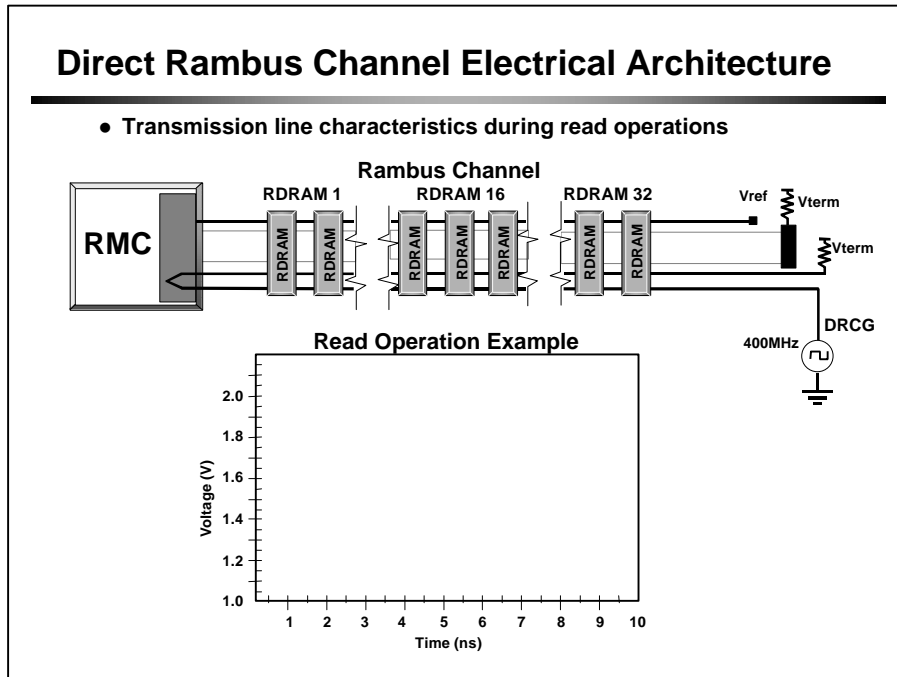
At time $t \approx 7.0$ ns, the signal reaches RDRAM device 32,. Now the transmission line delay thus far is about 5.5 ns.

Slide #20



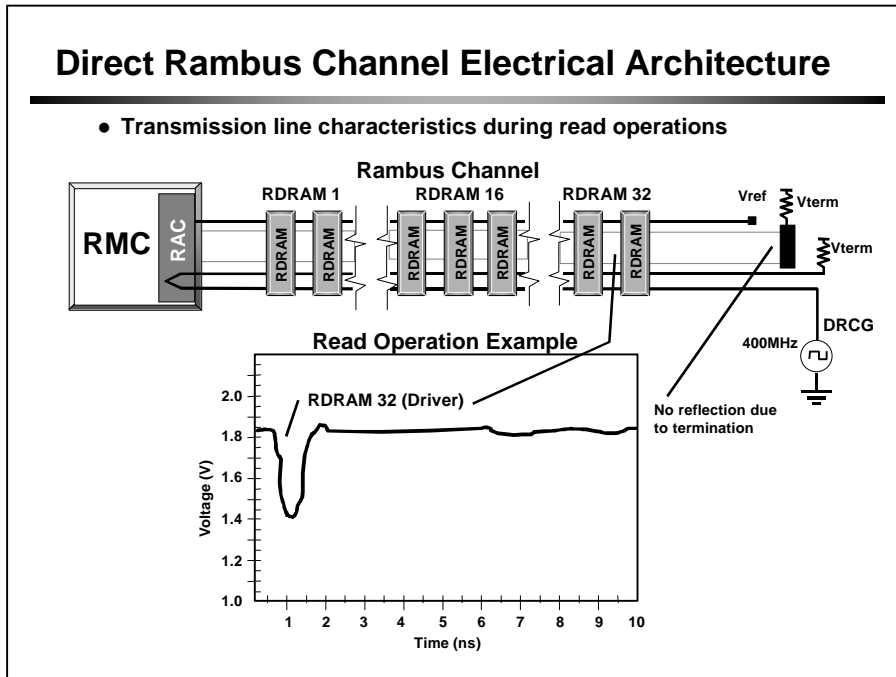
Finally, at time $t \approx 9.0$ ns, the signal reaches the end of the transmission line and the termination point. There is no reflection at this point. The transmission line delay thus far is about 7.5 ns.

Slide #21




This next set of slides demonstrates the transmission line characteristics on the Data bus during read transactions where the Data bus is driven by the RDRAM device and sampled by the RAC.

Slide #22



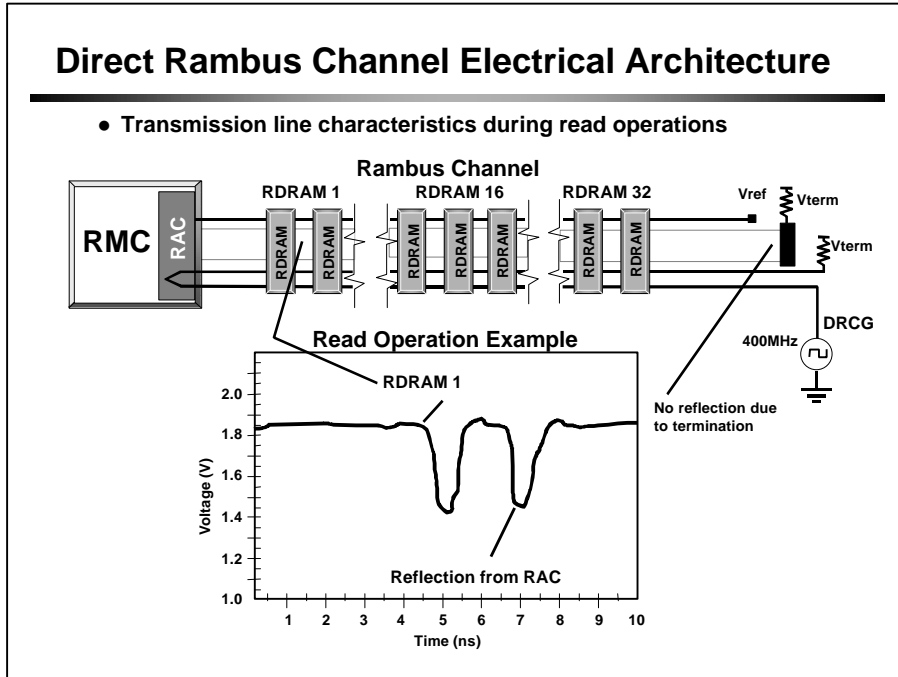
At time $t \approx 1.0$ ns, RDRAM device number 32 drives a logical 1. However, note that the signal is only driven from 1.8 V to 1.4 V. The signal is not driven to 1.0 V. The wave now propagates through the transmission line towards the RAC.

Slide #23




At time $t \approx 3$ ns, the signal reaches RDRAM device 16, meaning that the transmission line delay thus far is about 2.0 ns. Also, at time $t \approx 9$ ns notice a pulse which is the reflection of the incident wave from the RAC.

Slide #24



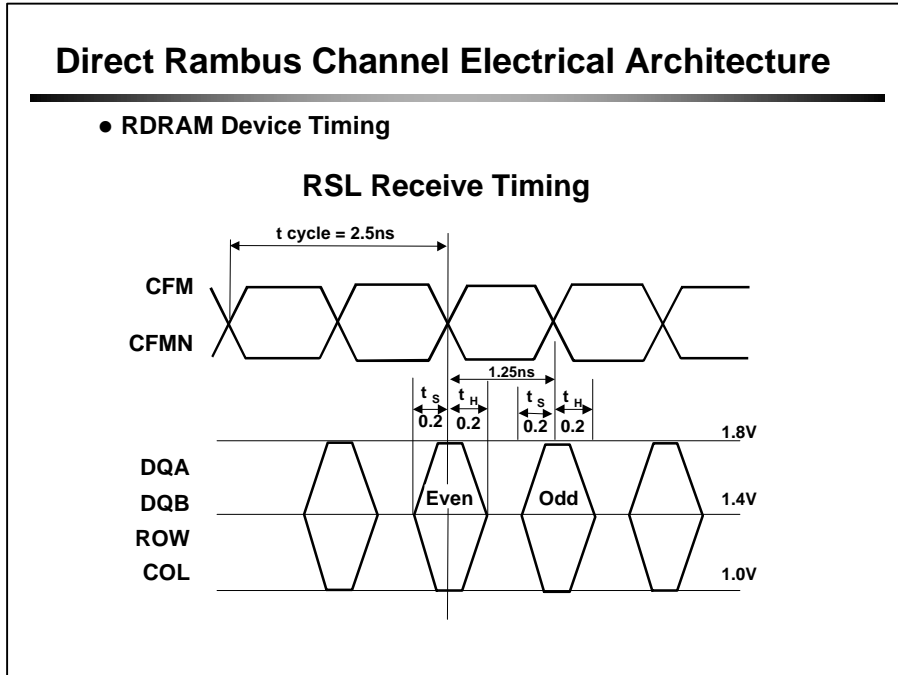
At time $t \approx 5$ ns, the signal reaches RDRAM device 1, meaning that the transmission line delay thus far is about 4.0 ns. Also, at time $t \approx 7$ ns notice a pulse which is the reflection of the incident wave from the RAC.

Slide #25



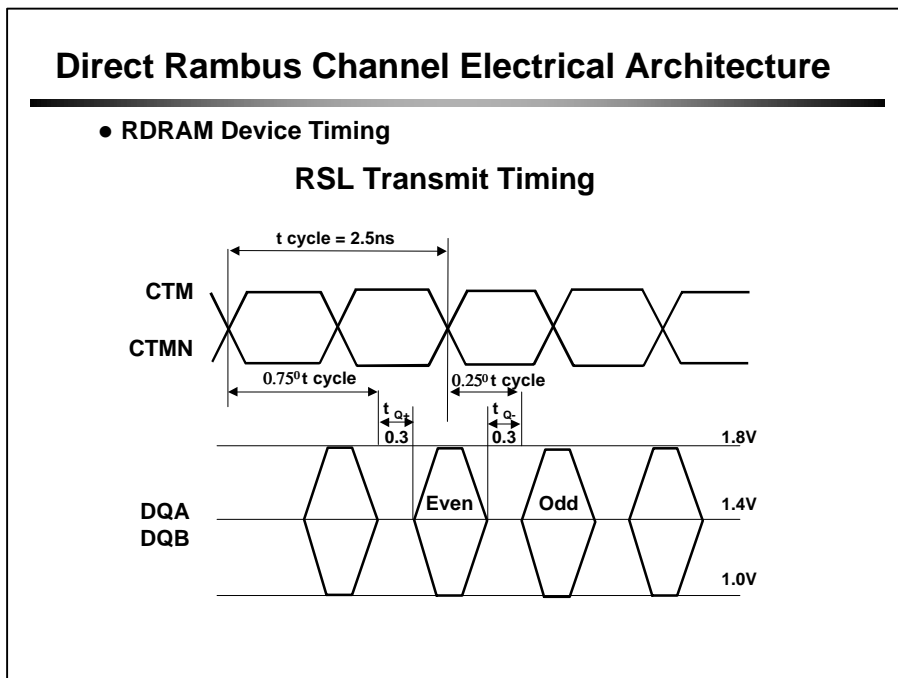
Finally, at time $t \sim 6.5$ ns, the wave reaches the RAC which has a high input impedance (no termination). Hence the incident wave is reflected and we notice the signal doubling. This voltage transition is detected at the RAC. Note that there are no reflections at the far end of the channel because of termination.

Slide #26



This timing parameter applies to the control signals and data bus during a write transaction. At the RDRAM device, signal setup and hold timing parameters are specified with respect to the falling edge of CFM. The setup and hold parameters are centered around the 0% and 50% points within a cycle time.

Slide #27



This timing parameter applies to read transactions from the RDRAM device. These timing parameters are specified for the data bus only. Transmit timing is specified with respect to 75% and 25% points within a cycle time from the falling edge of CTM.

Slide #28

Agenda

- **Direct Rambus Overview**
- **Direct Rambus Channel Electrical Architecture**
- *Direct RDRAM Device Physical Architecture*
- **Direct Rambus Logical Architecture**
 - **Read and Write Operations**
 - **Transaction Pipeline**
 - **Packet Format**
 - **Unique Concepts**
 - **Maintenance Operations**
 - **Configuration**
- **Measurements made with HP Instruments**

Slide #29


Direct RDRAM Physical Architecture

RDRAM Device Signals Consist of:

- **30 “RSL” Signals**
 - 3 ROW pins, 5 COL pins
 - 8/9 DQA Data pins, 8/9 DQB Data pins
 - CFM, CTM differential pair clocks
- **4 CMOS signals used for RDRAM initialization**
 - SIO0, SIO1, SCK, CMD
- **Static high-impedance, low current input**
 - Vref
- **DC power pins**
 - Vdd, Vterm and GND pins

The above slide lists the names of all signals on an RDRAM chip. In addition to the RSL signals, there are 4 CMOS signals, which are used during access to control registers, within each RDRAM device. V_{ref} input is the 1.4 V signal used as the compare input-to-input signal comparators. All other pins are power (2.5 V supply voltage) and ground inputs.

Slide #30




Reference design, VHDL/Verilog models and design guidelines for the RMC and the RAC are available from Rambus Inc. The RMC translates memory read and write cycles generated by a bus master into read and write channel commands. The memory controller initiates the maintenance operations, such as, refresh and current calibration and temperature calibration. The RMC manages power thermal envelope.

Key decisions a designer has to make when designing the RMC include the RDRAM device page management policy, based on factors such as locality of traffic during access to memory and the latency an application can tolerate during access. A close page policy might be recommended when memory accesses are not localized. Controller design is simpler in this case. An open page policy reduces latency when memory accesses have locality. The controller design requires page-tracking logic such as a Row Cache. Another decision might be pipeline depth support. The greater the pipeline depth, the higher the effective bandwidth needed, at the expense of increased design complexity. A typical RMC design may synthesize into 13K gates.

The RAC is the analog interface to the channel. It encodes the packets to be sent over the channel and decodes read data packets received from the channel. Designers must obtain accurate models of the RAC. Rambus Inc. provides initial models.

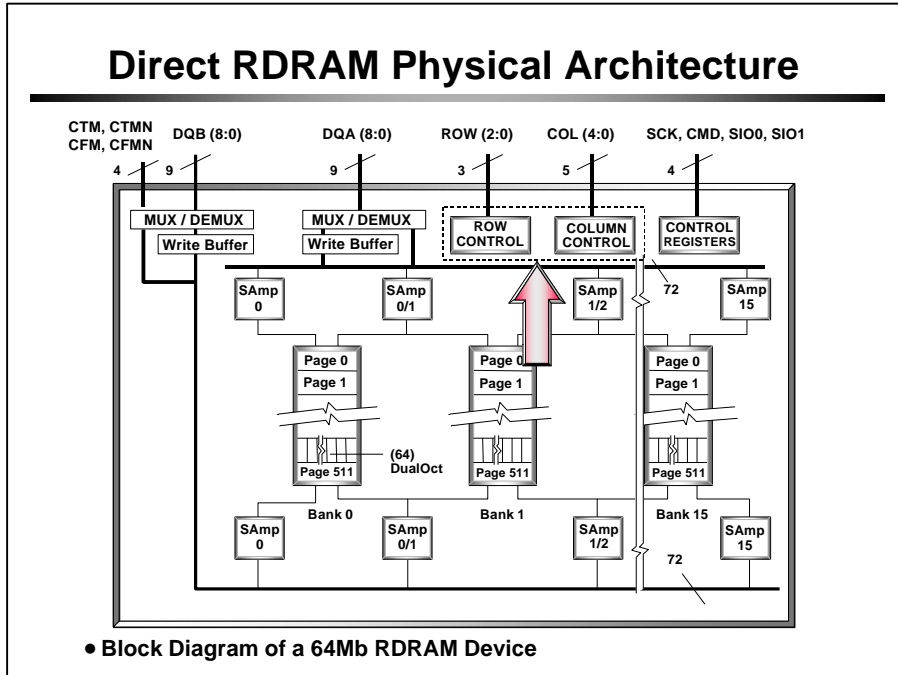
Slide #31



The above figure is a block diagram of a 64 Mb RDRAM device. The arrow points to all the signals present on the RDRAM chip. Their functions have already been described.

Major sub-units within the RDRAM device include the row control logic, column control logic, multiplexers and de-multiplexers, write buffers, sense amps, banks, pages or rows, data storage elements, control registers which contain configuration information and, finally, all the data paths interconnecting these units. The next set of slides describe each one of these blocks.

Slide #32




The Row Control and Column Control blocks decode packet information received through the ROW and COL pins.

The ROWA and ROWR (row-operation) packets are received via the ROW pins. These packets are de-multiplexed into 24-bits. The Row Control block primarily manages data transfers between the banks and the sense amps of the RDRAM device.

The COLC, COLM and COLX (column-operation) packets types are received via the COL pins. The COLC packet de-multiplexes into 23-bits. The COLM and COLX packets de-multiplex into 17-bits. The primary role of the Column Control block is to manage data transfers between the DQA/DQB pins and the sense amps of the RDRAM device.

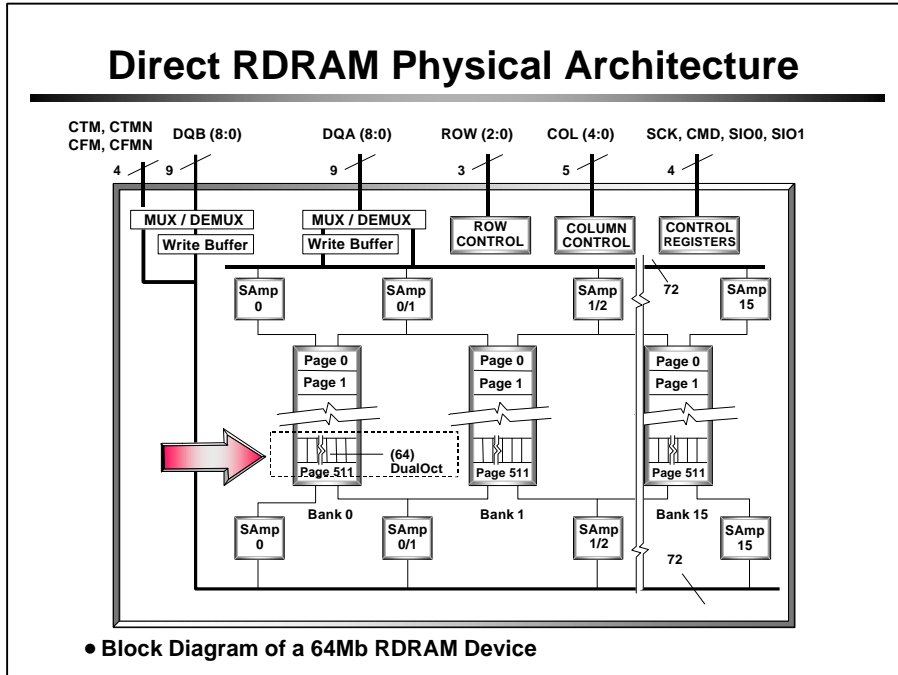
Slide #33



The write buffers shown accept write data through the DQA/DQB bits. Write buffers reduce the delay needed for the internal data paths to turn around, speeding up back-to-back write-read transactions and write-write transaction completion. A dualoct (16 bytes) of data is loaded into the write buffer during a write operation. The write buffer also holds bank address information, column address information and optional byte mask information. This information is automatically retired into the appropriate sense amp during a subsequent column-operation. Retires take place during subsequent column read/write operation to another device or a write operation to the same device. As a result, the write buffer is capable of holding a maximum of 2 dualocts of data. Note that the write buffer does not retire during reads to the same device.

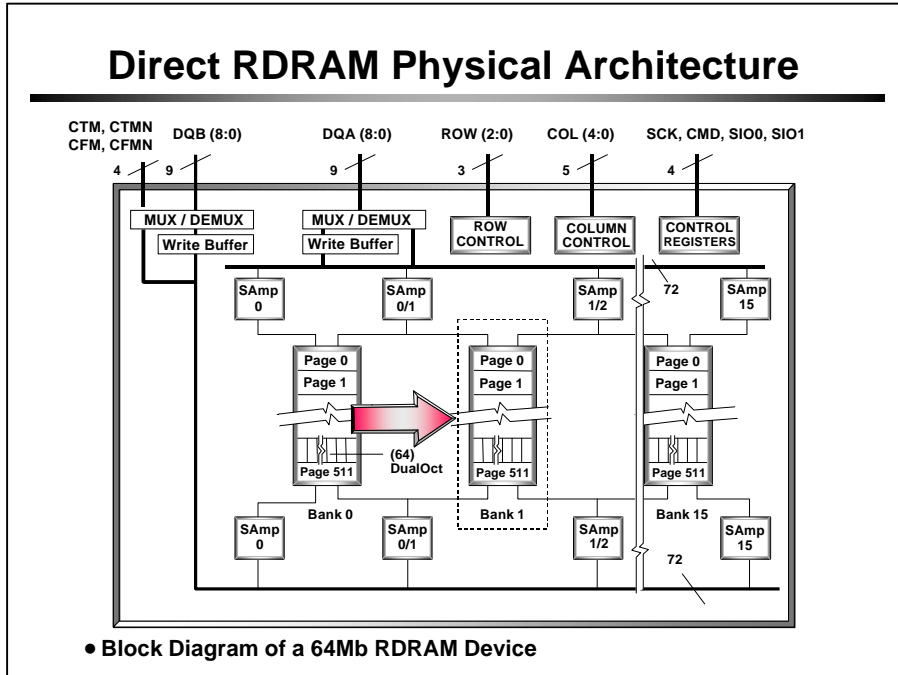
The 18-bit DQA/DQB pins carry read (Q) and write data (D) across the channel. These data packets are multiplexed/de-multiplexed (by the MUX/DEMUX unit) from/to two 72-bit internal data paths. Thus, the internal operating frequency is one-eighth the external frequency of 400 MHz.

Slide #34



The smallest addressable unit within the RDRAM device is a dualoct. A dualoct is defined as 16 bytes of data. In a 64-Mb device a row shown in the diagram above consists of 64 dualocts (1 KB) of data. This row of information is loaded into sense amps before it can be accessed during read or write operations.

Slide #35

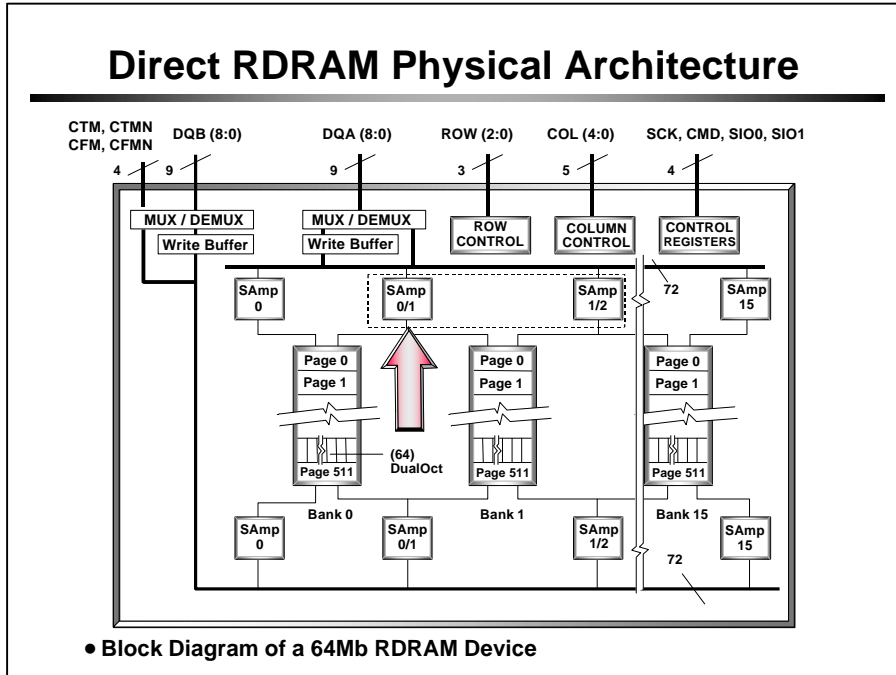


A 64 Mb RDRAM device consists of 16 banks. Each bank contains 512 rows. Thus, each bank holds 0.5 MB of data. A dualoct of data is addressed by specifying the following sub-address:

- 1) Device address. This is a 5-bit field specifying the address of 1 of 32 devices.
- 2) Banks address. This is a 4-bit field specifying 1 of 16 banks
- 3) Row address. This is a 9-bit field specifying 1 of 512 rows.
- 4) Column address. This is a 6-bit field specifying 1 of 64 dualocts.

The address is a 24-bit address with which one can address a 16 M dualoct address space per channel.

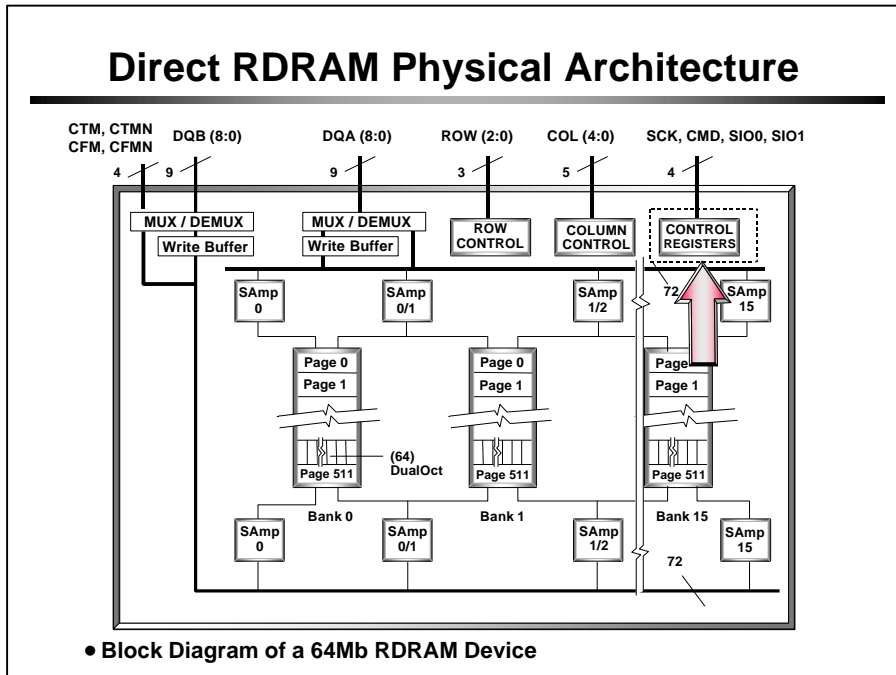
Slide #36



The RDRAM device contains a total of 17 sets of sense amps. Each set of sense amps holds 512 bytes (256 bytes on DQA internal data path and 256 bytes on DQB internal data path) of fast storage. Note that this is only half the amount of data that needs to be loaded from a row of 1 KB of data during a bank activate operation. The reason for this is that each sense amp is shared between two adjacent banks (with the exception of sense amp 0 and 15). In other words, a row of a bank gets loaded into two sets of adjacent sense amps. This sense amp sharing optimizes the RDRAM device physical size, reducing overall die size and cost. A pair of adjacent sense amps (for example sense amp 0/1 and 1/2) holds 1 KB of data from a row of an open bank.


A side effect to sharing sense amps is that two adjacent banks cannot be open at the same time. This means that a row within one bank must be precharged before a row within an adjacent bank can be opened.

Slide #37



The control register block holds the RDRAM device configuration information which is accessed through CMOS pins SCK, CMD, SIO0 and SIO1. These registers hold RDRAM device information, which is read by the RMC for configuration purposes. These registers are written to in order to select the operating mode of the RDRAM device. The DEVID register in this block holds the device address of the RDRAM device on the channel. The DEVID register value holds the address of the last refreshed row. The SCK is a 1 MHz slow running clock.

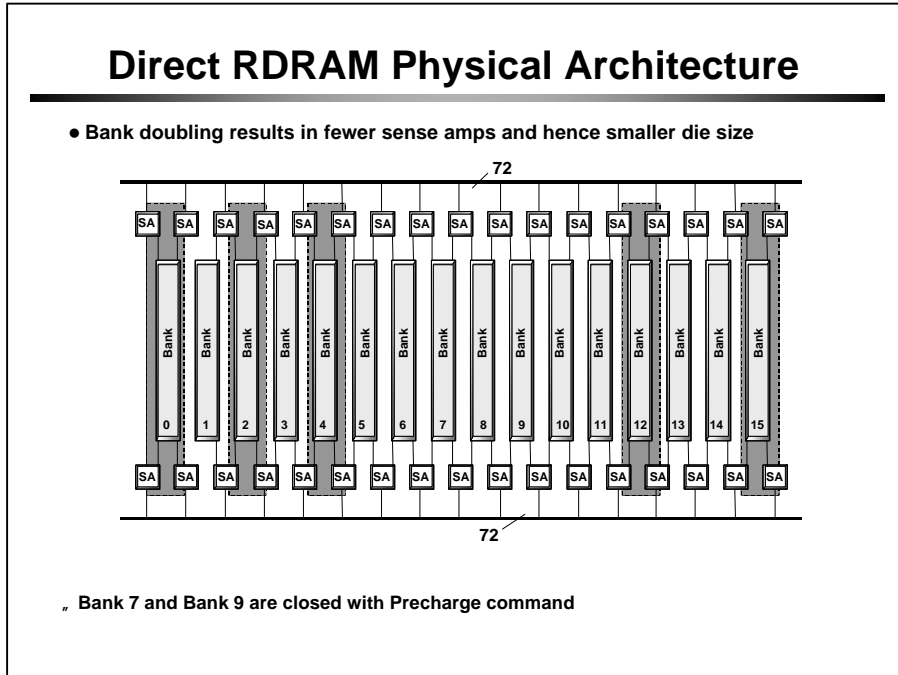
Slide #38



This example demonstrates the concept of bank doubling. The shaded banks are open or activated and a row from each of the shaded banks is contained in adjacent sense amps.


To activate Bank 8, Bank 7 and 9 must be precharged first.

Slide #39



Bank 7 and 9 are precharged with a precharge command, which can be sent either through a row-operation or column-operation command packet.

Slide #40



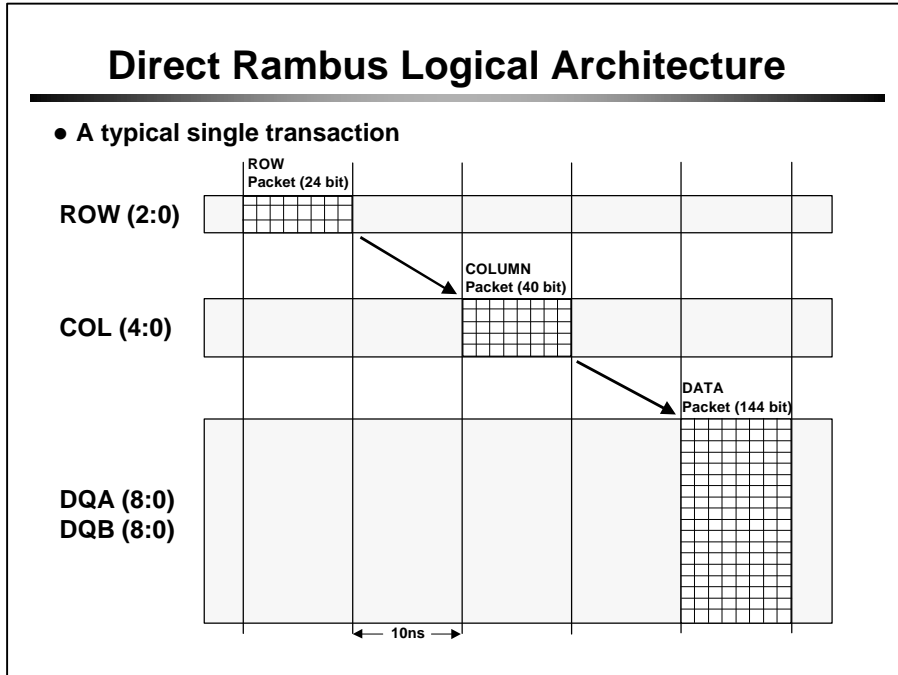
Bank 8 is activated with an Activate row-operation command packet (ROWA) which causes 1 of 512 rows of the selected bank to be loaded into the associated sense amps 7/8 and 8/9.

Slide #41

Agenda

- **Direct Rambus Overview**
- **Direct Rambus Channel Electrical Architecture**
- **Direct RDRAM Device Physical Architecture**
- ***Direct Rambus Logical Architecture***
 - **Read and Write Operations**
 - **Transaction Pipeline**
 - **Packet Format**
 - **Unique Concepts**
 - **Maintenance Operations**
 - **Configuration**
- **Measurements made with HP Instruments**

Slide #42




A typical read or write operation initiated by the memory controller consists of the transmission of a ROW packet followed by the transmission of a COL packet. If the operation is a write operation, then the memory controller also sends a DATA packet. If the operation is a read operation, the DATA packet is transferred by the RDRAM device that is being read.

Note that each of the packets are sent during a 10 ns window and consist of 8 columns of data. The first bit transmitted always starts on the falling edge of the appropriate 400 MHz clock. This bit time is called the “even” bit time. A packet is transferred within 4 clocks.

The ROW packet is always a 24-bit packet, the COL packet a 40-bit packet and the DATA packet a 144-bit (or 128-bits of a dualoct if no parity/ECC is supported) packet.

Keep in mind that this is a typical operation, only. More specific operations are described later in the presentation.

Slide #43



The transaction protocol supports two types of ROW packets transferred by the memory controller:

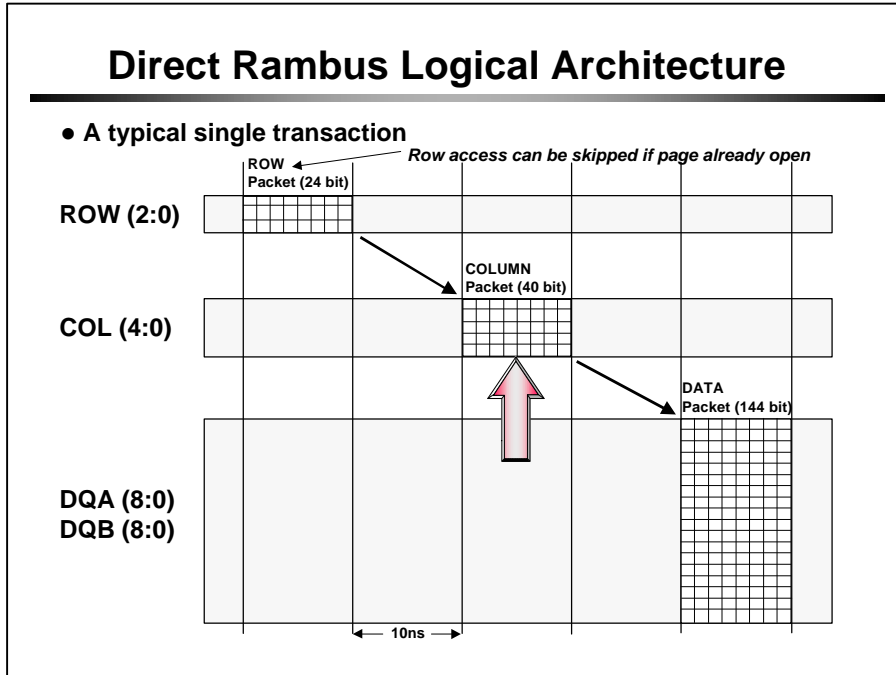
- 1) The ROWA (row-activate). This packet is distinguished by the fact that a bit within this packet, called the Activate bit (AV) in packet, is equal to 1.
- 2) The ROWR (row-operation). This packet is distinguished by the fact that the AV bit in the packet is equal to 0.

The specific contents of these two types of packets are discussed later.

The ROWA packet specifies the activation of a row of a bank of an RDRAM device into a sense amp pair that holds 1 KB (64 dualocts) of data.

The ROWR packet specifies the precharge of a bank of an RDRAM device whose row is loaded in a sense amp pair. This ROWR packet may also contain control information such as Refresh, Power Management and Temperature Calibration.

Slide #44



A COL packet transferred by the memory controller is actually split into two fields:

- first field is COLC (column-command);
- second field is either a COLM (column-mask) or COLX (column-extended) field.

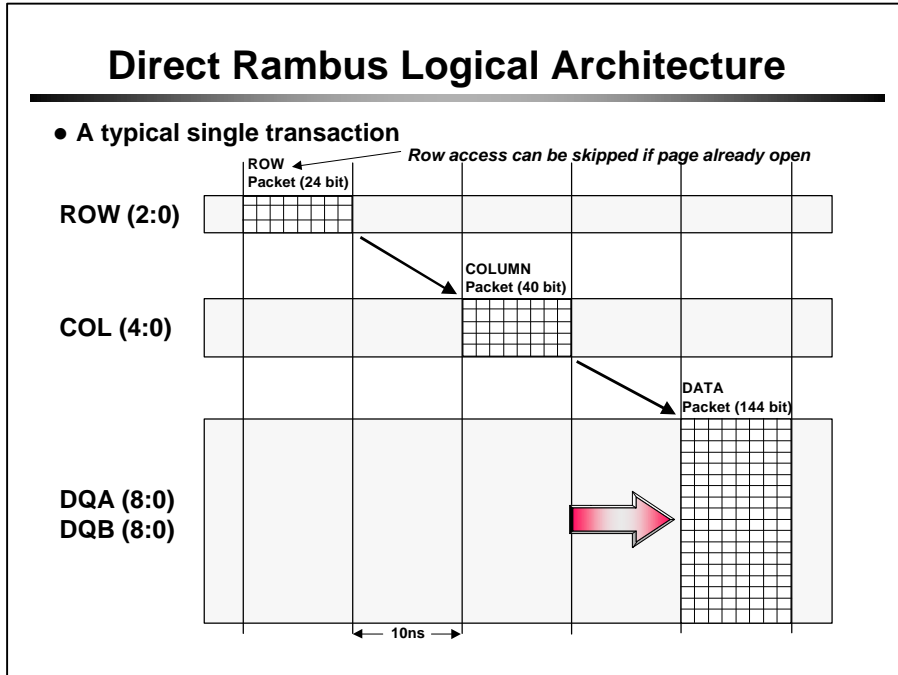
COLC field is distinguished by the Start bit (S) in the packet being equal to 1. The COLC portion of the COL packet specifies:

- Read or Write operation of a dualoct of data from/to an active row of data within a sense amp pair of an RDRAM device;
- Precharge or power management operation associated with an RDRAM device.

COLM field portion of a COL packet is distinguished by the Mask bit (M) in the packet being equal to 1. This optional portion of the COL packet specifies Byte Masks for write operations only. There is no COLM packet supported for read operations during which the assumption is that all 16 bytes of data within a dualoct are read. If no COLM packet is specified during a write operation, the RDRAM device assumes that all 16 bytes of the write DATA packet contains valid data.

COLX field is distinguished by the Mask bit (M) in the packet being equal to 0. This portion of the COL packet specifies No Operation, Power Management, Current Calibration or Precharge operations.

Slide #45




A DATA packets consists of up to 16 bytes of data with 8- or 9-bits per byte. Two bytes of data are transferred per clock rising or falling edge.

During write operations the Byte Mask field in COLM packet indicates the associated bytes that contain valid write data. During read operations, all bytes of data are valid.

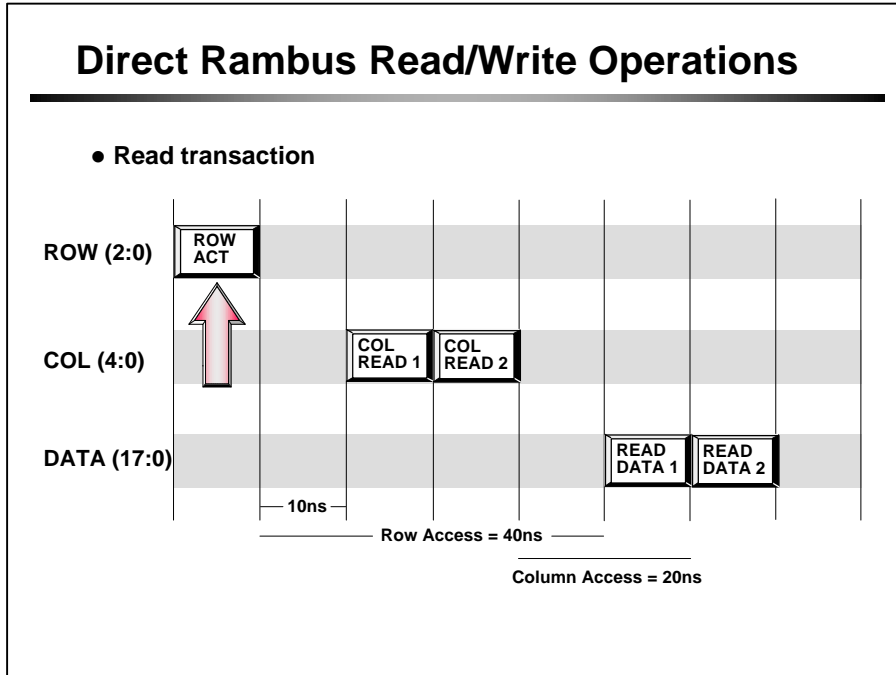
This packet is driven by the RMC to RDRAM devices during writes operations and is driven by the RDRAM device back to RMC during read operations.

Slide #46



This slide demonstrates two pipelined back-to-back read operations. The ROWA packet, which is optional if the accessed row is already open, is transmitted first. Then two back-to-back COL packets are transferred indicating the read of two dualocts of data from open rows. After a time delay, the read data arrives at the memory controller. Two back-to-back data packets are transferred with no gaps on the data bus, thus showing efficient use of the data bus during pipelined operations. Notice that there is a 10 ns or 4-clock gap between a ROW packet and a COL packet. This gap allows sufficient time for a row to be activated. Given that the operations are pipelined, no gaps are noticed on the COL bus.

Slide #47

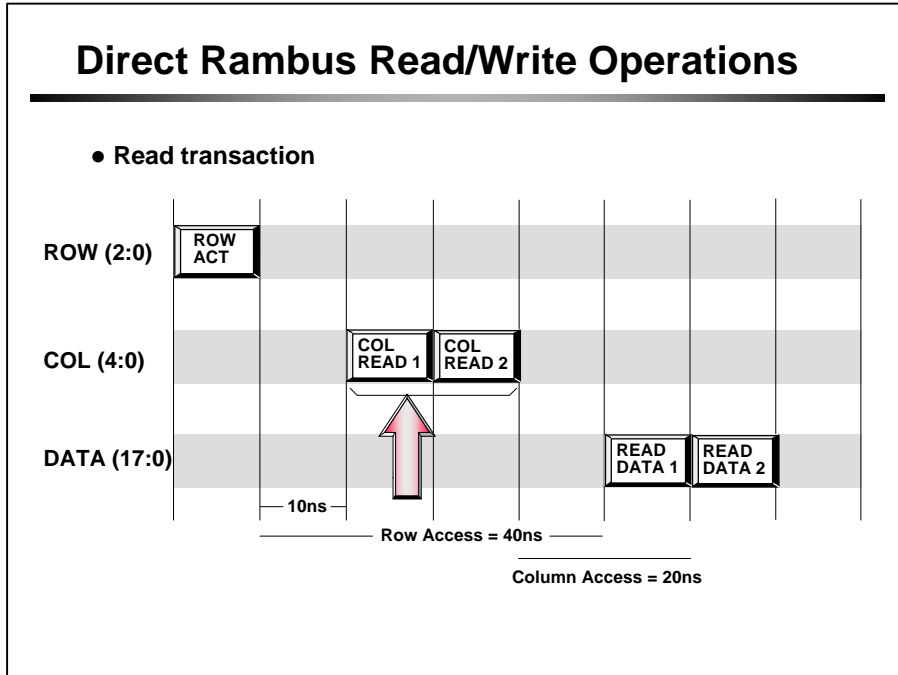


The ROWA activate packet contents include:

- AV bit = 1 indicating ROWA Activate packet;
- 6-bit Device Address (DR) pointing to either 1 of 32 devices, or indicating a broadcast operation to all RDRAM devices on the channel;
- 4-bit Bank Address (BR) pointing to 1 of 16 banks within an RDRAM device; and a
- 9-bit Row Address (R) pointing to 1 of 512 rows that is to be activated.

An 11-bit Opcode field (ROP), which is not included in a ROWA packet but is included in a ROWR packet, specifies the type of operation, such as, precharge, refresh, temperature calibration, no operation and power management functions

Slide #48




COL packet content includes:

- S bit = 1 for framing COLC packet;
- 4-bit Opcode field (COP), which, in this example, specifies Read type Packet;
- 5-bit Device Address (DC);
- 4-bit Bank Address (BC); and a
- 6-bit Column Address (C), which specifies dualoct address of data to be read within an active row contained within a sense amp pair.

Note that row address is not specified in this packet because the row was activated by a previous ROWA packet.

The second COLC packet has no previously associated ROWA packet because the read operation is from an already activated row.

Slide #49

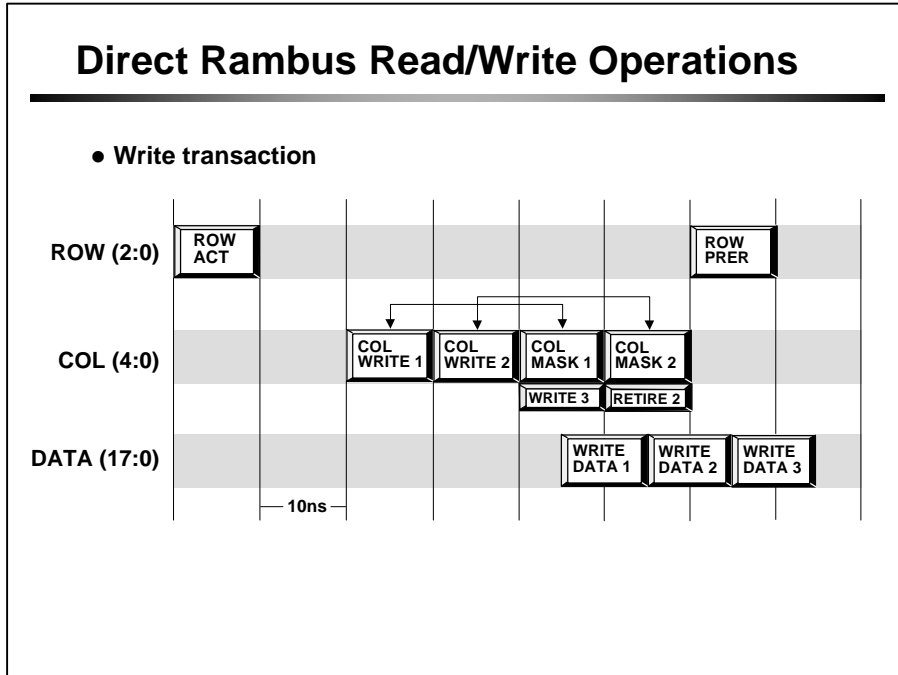


Read Data is available approximately 8 clocks (20ns) after the associated COLC packet. The exact time from COLC before data is returned to RAC depends on:

- RDRAM device column access time; and
- Channel length where “multiple clock domains” exist.

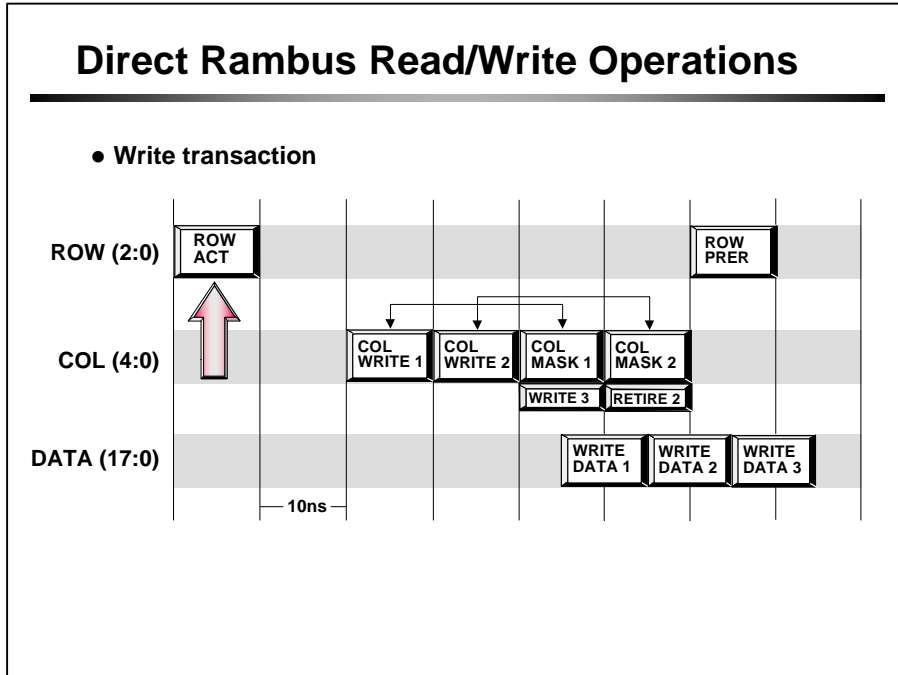
RDRAM devices are initialized with varying TRDLY delay values. TRDLY value, which is contained in a control register, represents column operation to read data latencies. The TRDLY value is initialized so that data arrives at RAC from the RDRAM device with the same delay from column operation no matter what distance the RDRAM device is from RAC. The closest RDRAM device to the RAC is initialized with the largest latency and the farthest RDRAM device is initialized with the smallest (zero) latency.

Slide #50



This slide demonstrates three pipelined back-to-back write operations. The ROWA packet, which is optional if the accessed row is already open, is transmitted first. Then, two back-to-back COLC packets are transmitted. A third COLC packet with a COLM packet concatenated is transferred. After a time delay from the COL packets, the memory controller transfers the write data in 3 back-to-back write DATA packets. Notice that there is 10 ns or 4-clock gap between a ROW packet and a COL packet. This gap allows sufficient time for a row to be activated. But, given that the operations are pipelined, no gaps are noticed on the COL bus.

Slide #51

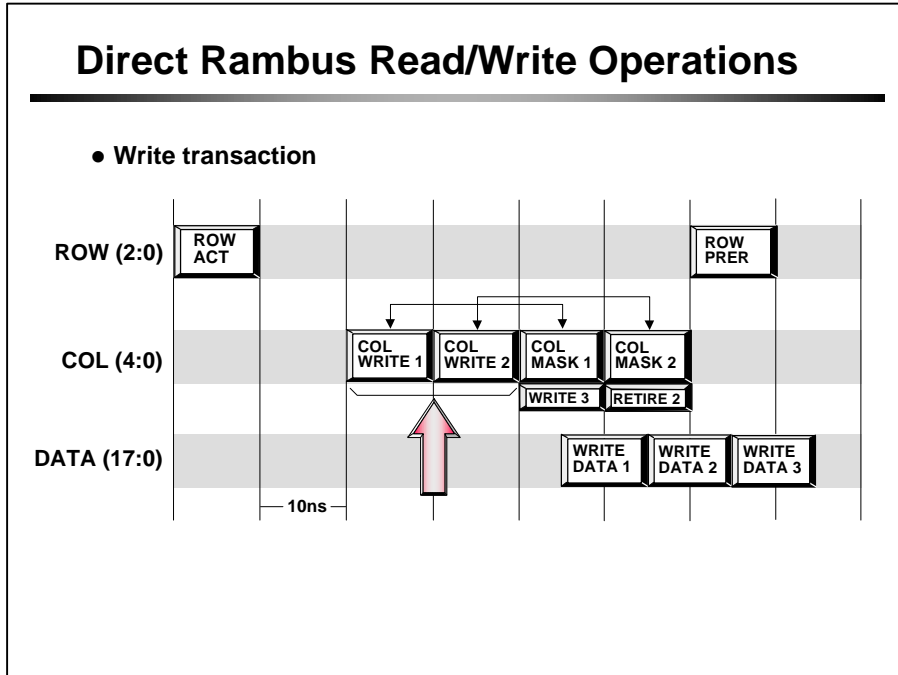


The ROWA activate packet contents include:

- AV bit = 1 indicating ROWA Activate packet;
- 6-bit Device Address (DR) pointing to either 1 of 32 devices, or indicating a broadcast operation to all RDRAM devices on the channel;
- 4-bit Bank Address (BR) pointing to 1 of 16 banks within an RDRAM device; and
- 9-bit Row Address (R) pointing to 1 of 512 rows that is to be activated.

An 11-bit Opcode field (ROP), which is not included in a ROWA packet but is included in a ROWR packet, specifies the type of operation such as precharge, refresh, temperature calibration, no operation and power management functions.

Slide #52



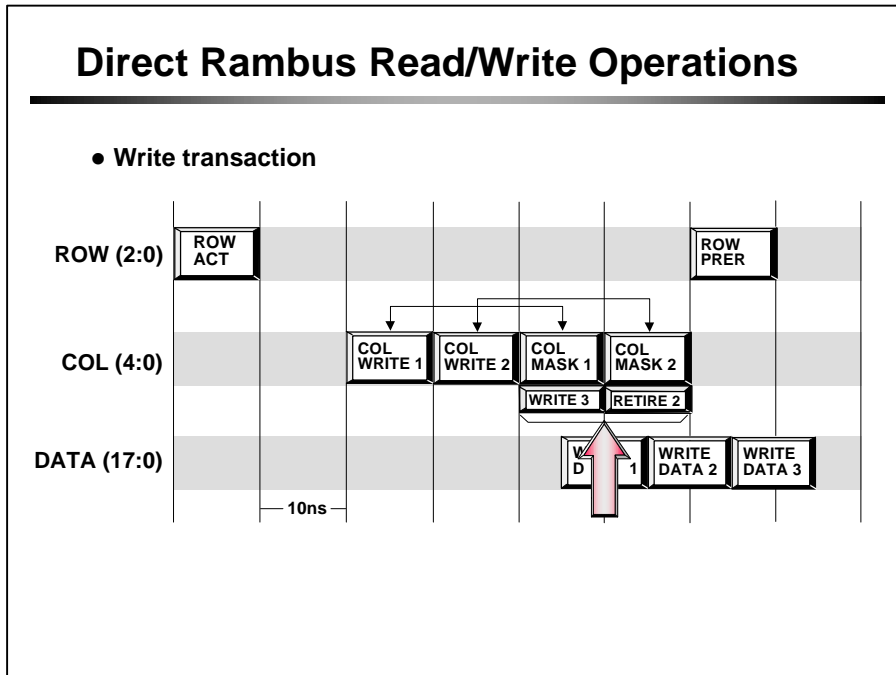
COLC packet contents include:

- S bit = 1 for framing COLC packet;
- M bit = 0 indicating no concatenated COLM packet to this COLC packet;
- 4-bit Opcode field (COP), which, in this example, specifies Write type packet;
- 5-bit Device Address (DC);
- 4-bit Bank Address (BC); and
- 6-bit Column Address (C) specifies dualoct address of the data to be written within an active row contained within a sense amp pair.

Note that row address is not specified in a COLC packet because the row is already activated by a previous ROWA packet.

The second COLC packet has no previously associated ROWA packet because the write operation is to an already activated row.

Slide #53

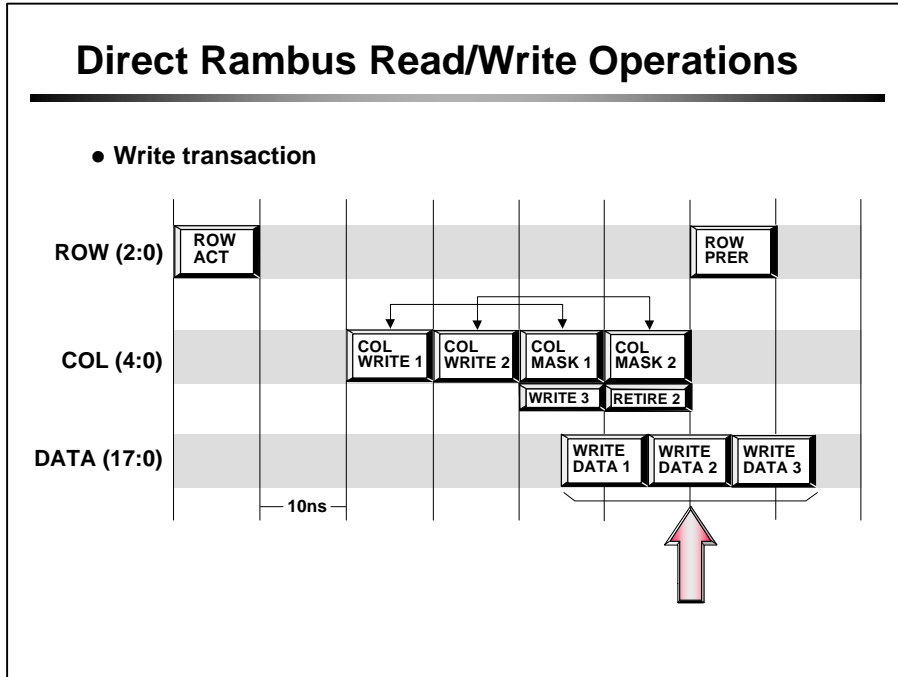


The first COLC packet, indicated by the arrow in the slide above, is the third write packet in this diagram. This COLC packet will cause the write buffer to retire of the first data packet written. Concatenated to this packet is a COLM packet. The COLM portion of this COLC packet is associated with the previous COLC packet, 8 clocks earlier. The COLM packet specifies Byte Masks for the dualoct data to be written.

The second COLC packet in the diagram is a retire packet, which will cause the second dualoct of data in the write buffer to be retired also. This COLC packet also has a concatenated COLM packet, which is associated with the previous COLC packet 8 clocks earlier.

If, at these time frames, there was no concatenated COLM packet ($M = 0$) the assumption made by the RDRAM device is that all Byte Masks associated with the previous two COLC packets are enabled and all data bytes are unconditionally written.

Slide #54

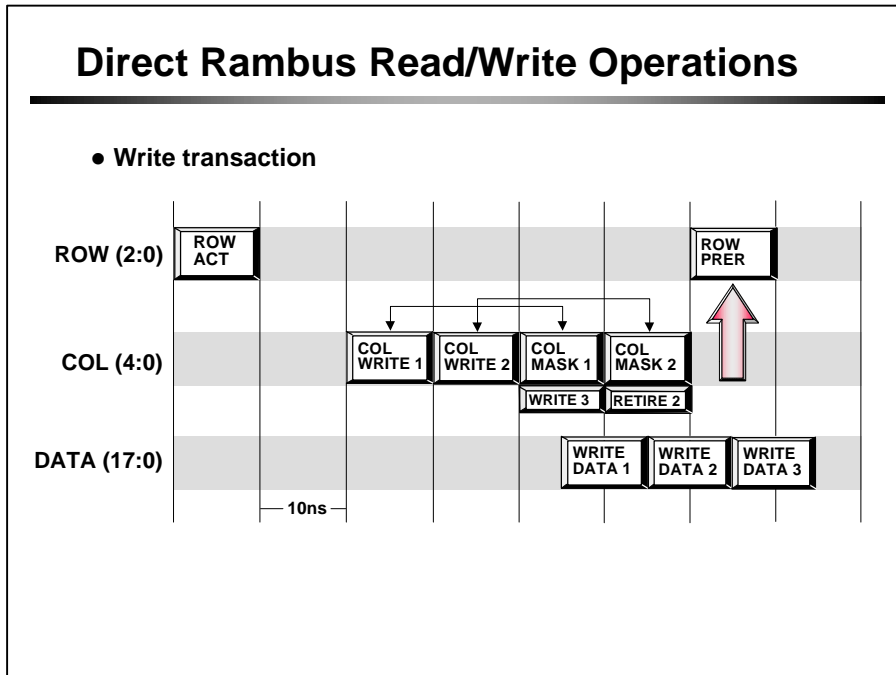


Data is written 6 clocks after the associated COLC packet. But the data is actually written to a write buffer in the RDRAM device first. The write buffer is then retired to sense amps at a later time. The exact time of write buffer retirement is based on a complex set of rules

Note that after a write operation has completed to the write buffer, a read from the same bank must not be initiated unless:

- Write buffer has been retired or
- Memory controller guarantees that the read is not to the same dualoct address as the previous write. Optionally, a comparator is built within the memory controller to determine if the read address is the same as previous write address. If so, the read operation is completed internal to the memory controller.

Slide #55



This row precharge packet (ROWR) is used to precharge the open row, which was activated a little earlier by the ROWR packet shown. The ROWR packet contents include:

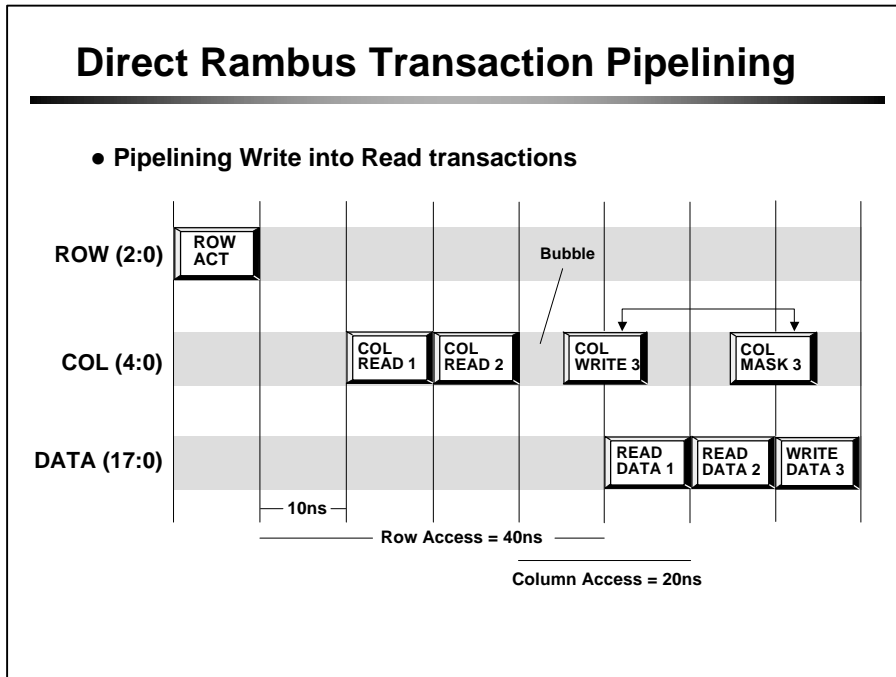
- AV bit = 0 indicating ROWR packet;
- 6-bit Device Address (DR) pointing to either 1 of 32 devices, or indicating a broadcast operation to all RDRAM devices on the channel;
- 4-bit Bank Address (BR) pointing to 1 of 16 banks within an RDRAM device; and an
- 11-bit Opcode field (ROP) specifies the type of ROWR operation, which, in this example, is a precharge operation. There is no Row Address (R) field for this packet because the RDRAM device remembers to precharge a row contained in the sense amps to the same row it was activated from in a bank.

A row is precharged because:

- Another row within the same bank is to be opened or
- A row within an adjacent bank is to be opened or
- A row is open for maximum allowed row active time (64µs).

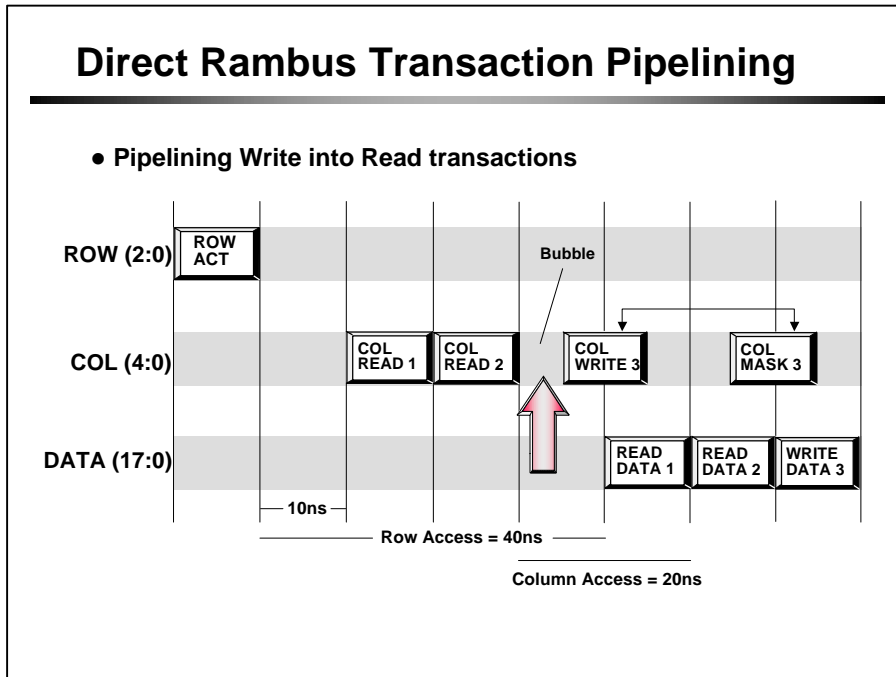
Minimum latency must be guaranteed between activation of a row and precharge of same row.

Slide #56



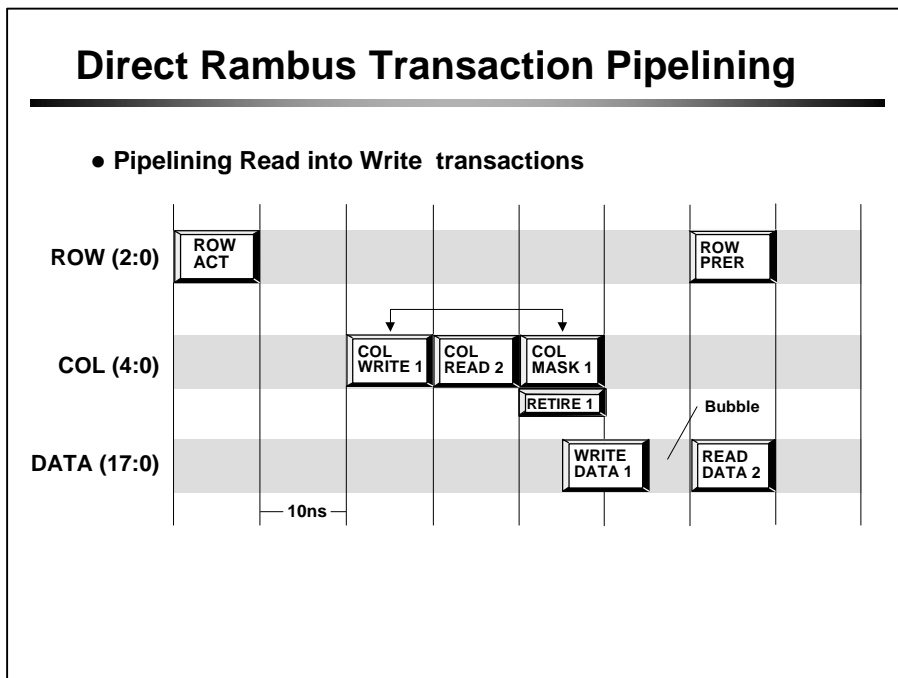
The above diagram shows two back-to-back read transactions followed by a write transaction pipelined into the read transaction. Notice that there are no bubbles on the data bus.

Slide #57




Notice the bubble between read and write COLC packet which is inserted by the memory controller. This bubble is needed to delay the write data until the read data from the previous read has arrived at the memory controller. The data bus is efficiently utilized.

Slide #58



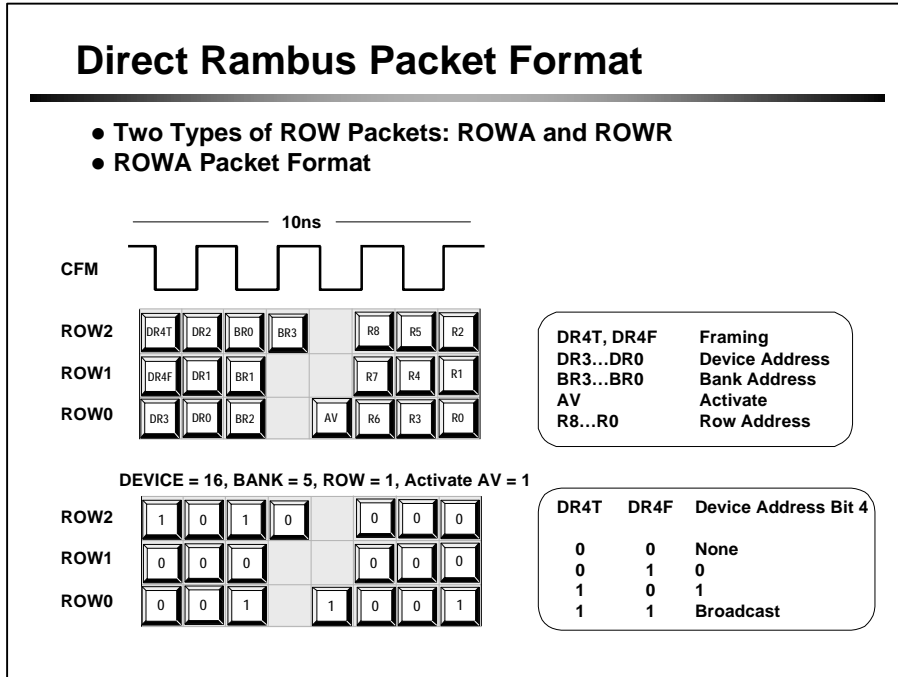
The above diagram shows a write transaction followed by a back-to-back read transaction. Notice that there is no bubble inserted by the memory controller on the COL bus.

Slide #59



Notice the bubble between write and read data. This bubble exists because of inherent write-read data turn-around time. The data bus is efficiently utilized with the exception of gaps on the data bus that are a result of write operations followed by read operations. Typically, the data bus utilization is greater than 95 %.

Slide #60



Two types of ROW packets are defined, the ROWA and ROWR packet. The above diagram shows a ROWA (or row activate) packet and the definition of each bit. Like all packets transmitted, the first column of bits of a packet is transferred on the falling edge of CFM clock. Clock and the packet move through the channel together.

The AV bit = 1 for a ROWA packet. The DR4T and DR4F bits indicate either broadcast operation or transmission to an RDRAM device whose most significant device address bit is either 0 or 1. All other bits are self-explanatory.

Slide #61

Direct Rambus Packet Format

- Two Types of ROW Packets: ROWA and ROWR
- ROWR Packet Format

10ns

CFM

ROW2	DR4T	DR2	BR0	BR3	ROP10	ROP8	ROP5	ROP2
ROW1	DR4F	DR1	BR1		ROP9	ROP7	ROP4	ROP1
ROW0	DR3	DR0	BR2		AV	ROP6	ROP3	ROP0

DR4T, DR4F Framing
 DR3...DR0 Device Address
 BR3...BR0 Bank Address
 AV Activate
 ROP10...ROP0 Row Op Code

DEVICE = BROADCAST, BANK = 5, COMMAND = PRECHARGE

ROW2	1	0	1	0	1	0	0	0
ROW1	1	0	0		1	0	0	0
ROW0	0	0	1		0	0	0	1

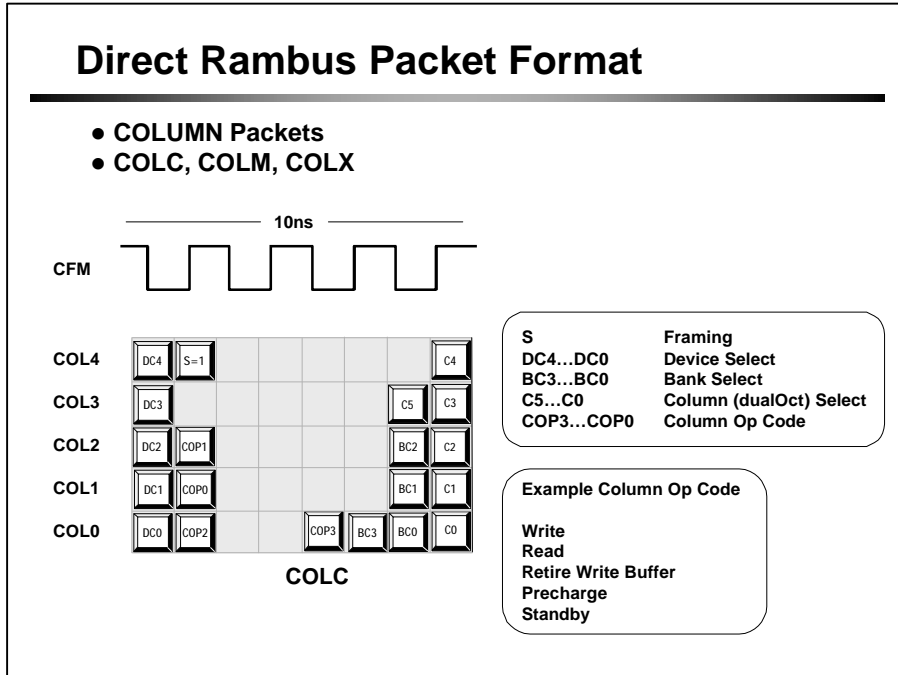
Example Row Op Code

Precharge
 Temperature Calibrate
 Refresh
 Power Down
 Nap
 Relax

For a ROWR packet, the AV bit = 0. This bit is located at the same position as in a ROWA packet. Also notice that instead of the R bits in the ROWA packet, they are replaced by 11-bit ROP bits. Definition of these bits is shown in the table above. All other bits are self-explanatory.

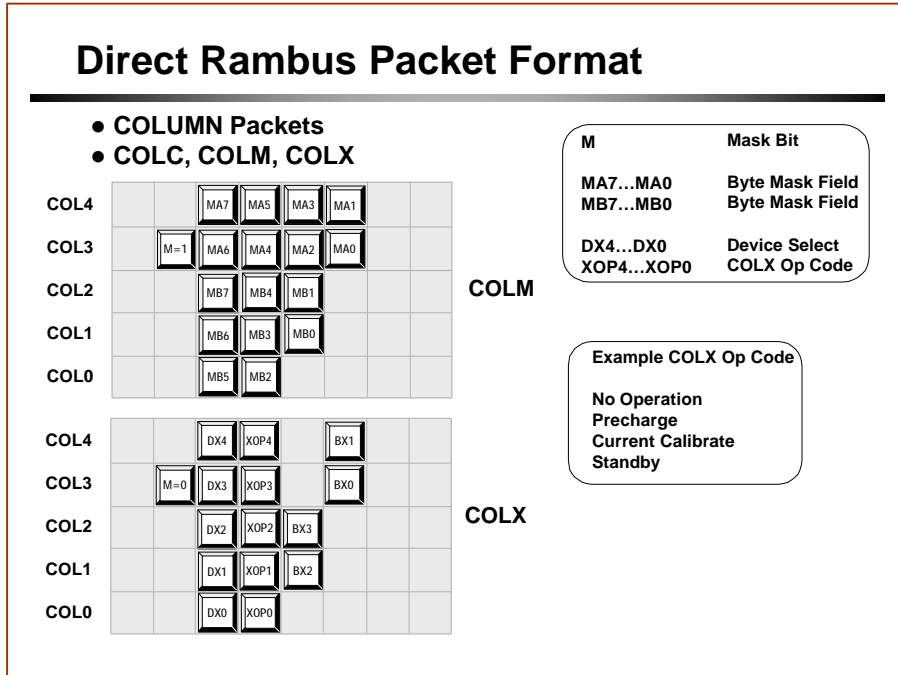
A ROWR packet is also used in general to precharge a row within the sense amps.

Slide #62



Now, let's look at COL packets. The above diagram shows the COLC portion of the COL packet. A valid COLC packet contains the S bit = 1. The 4-bit COP field indicates the kind of column operation being initiated, such as, read or write etc. The 6-bit Column Address (C) field points to the exact dualoct (1 of 64) within the row contained in the sense amps that are being accessed.

Slide #63




Whenever a Column packet is transmitted it contains a COLC portion as well as either a COLM or COLX portion concatenated. When the M bit = 1, the COLM packet is concatenated, when M bit = 0, the COLX packet is concatenated.

In the COLM packet, there are 16 mask bits corresponding to the 16 bytes in a dualoct. During read operations, all 16 bytes are read unconditionally.

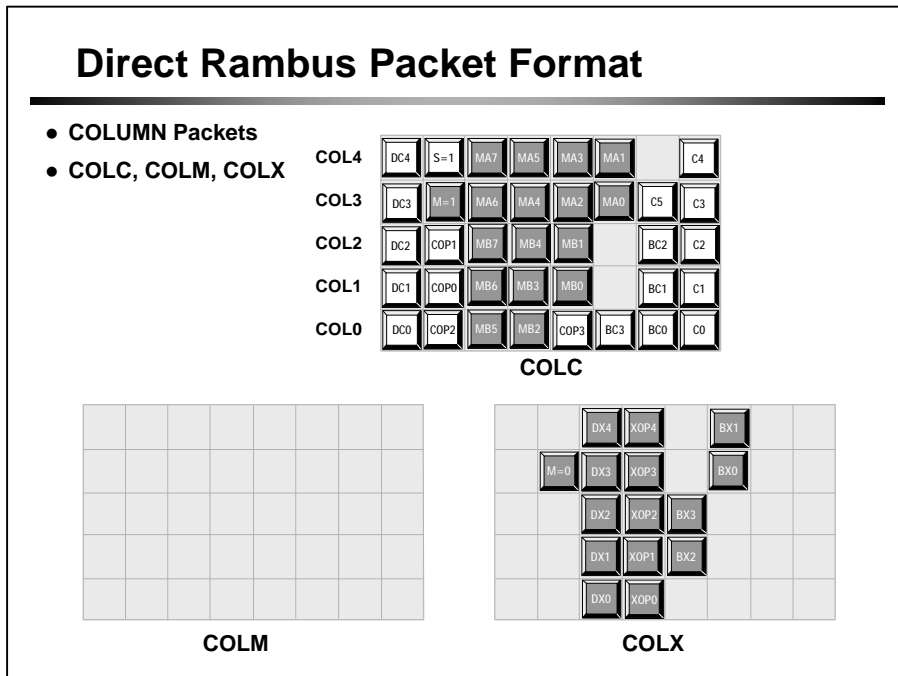
The COLX packet contains a 5-bit Device Address (DX) field, which points to 1 of 32 devices. The 5-bit XOP field indicates the type of operation for which this packet is used. The table above shows the type of operations represented by the XOP packet.

Slide #64



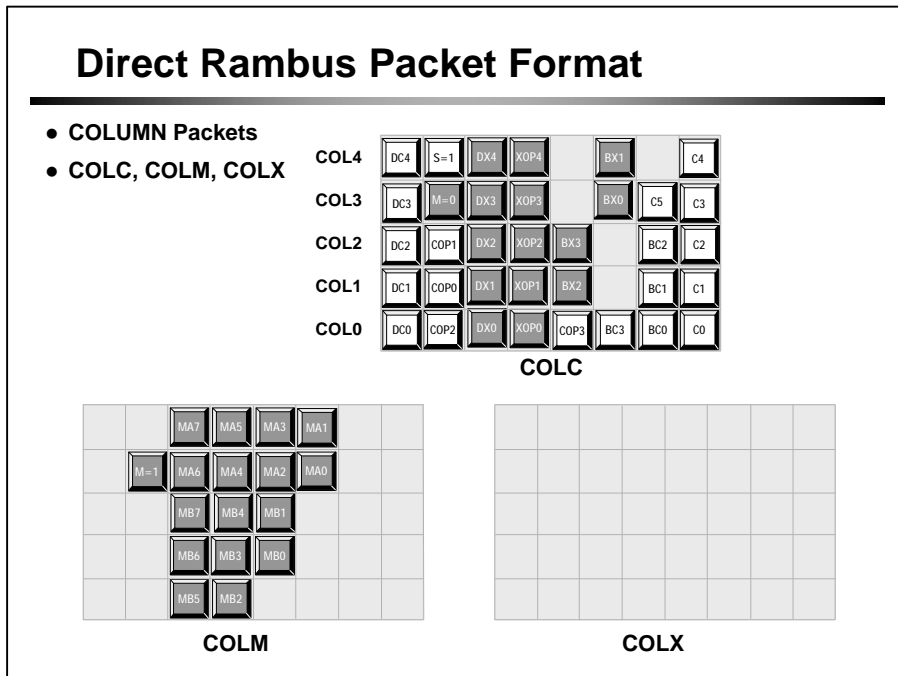
These next set of slides shows you how a COLC packet is concatenated with a COLM or COLX packet

Slide #65



COLC packet concatenated with COLM packet.

Slide #66



COLC packet concatenated with COLX packet.

Slide #67

Direct Rambus Unique Concepts

- **Some Unique Direct RDRAM Device Concepts**
 - **Bank Doubling**
 - **Write Buffer**
 - **Read Domains and Levelization**
 - **Power Management Modes**
 - **Active (Attention)**
 - **Standby**
 - **Nap**
 - **Power Down**
 - **Maintenance Operations**
 - **Self Refresh**
 - **Active Refresh**
 - **Current Control**
 - **Temperature Calibration**

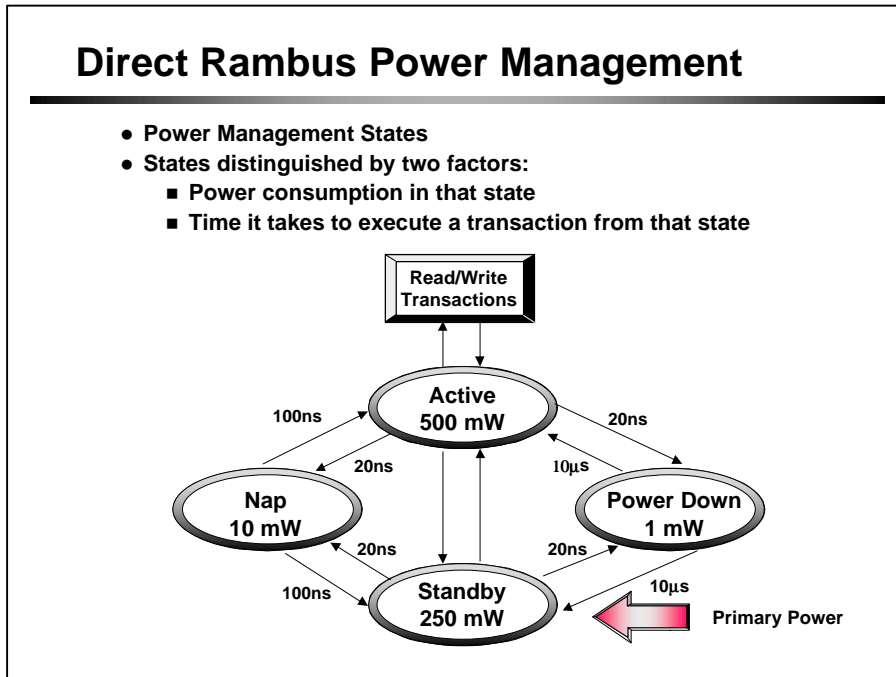
The concept of “bank doubling” allows for more banks, lower die area, and lower power consumption. We discussed bank-doubling architecture earlier.

The write buffers within the RDRAM device improve back-to-back write-write and write-read performance by allowing pipelining for these transactions. Without write buffers, it would not be possible to efficiently pipeline write-write and write-read transactions. Write operations first complete to the write buffer. The contents of the write buffer are retired to the sense amps on a subsequent operation provided the operation is a real operation to another RDRAM device or any non-read operation to any RDRAM device. In other words, the contents of the write buffer of an RDRAM device will not retire if a subsequent operation is a read operation to the same RDRAM device as the previous write operation.

During read operations, the data arrival delay from COL packet transmission is twice the flight time from memory controller to the RDRAM device. As a result, the memory controller sees a longer delay for read data from a farther RDRAM device than from a nearer RDRAM device. To “levelize” this delay across all RDRAM devices, a TRDLY delay value is programmed into the Control Registers of each RDRAM device. In essence, the memory controller sees equal delays between COL packet transmission and read data no matter how far the RDRAM device is from the memory controller.

Power management and maintenance operations are discussed in the following slides.

Slide #68



The primary power state of an RDRAM device is the standby state. When in the standby state, the RDRAM device is consuming about 250 mW of power. Whenever an RDRAM device is responding to a transaction, it transitions to the active state. At any point in time, only one RDRAM device is in the active state, unless a broadcast transaction is targeting all RDRAM devices.

The nap and power down modes are lower-power consumption modes in which select logic within the RDRAM device is turned off. The power down mode consumes less power (1 mW) than the nap mode (10 mW). But the power down state exit latency is larger than the nap state exit latency.

The memory controller manages these power states to optimize overall system power consumption.

Slide #69

Direct Rambus Maintenance Operations

- **Maintenance Operations**
 - **Active Refresh**
 - Memory controller directed refresh
 - SIO0gle bank of each RDRAM Device refreshed per refresh cycle
 - Multi bank refresh in future
 - Broadcast
 - ~2% of bandwidth
 - **Self Refresh**
 - No memory controller support
 - In Power Down or Nap modes
 - **Current Control**
 - Once every 100ms
 - Ensure constant current over PVT
 - **Temperature Calibrate**
 - Once every 100ms
 - Ensure I/O slew rate within 200-400ps spec limit over PVT
 - Broadcast

The four maintenance operations associated with RDRAM devices are outlined in the slide above.

The memory controller initiates a refresh by sending a REFA command using a ROWR packet followed by a REFP command also using a ROWR packet. Each row has to be refreshed once every t_{REF} period, which is specified in the RDRAM device data book. This period is about 32 ms. The memory controller has to initiate a refresh every period equal to $t_{REF} / (\# \text{ of banks} * \# \text{ rows per bank})$. The refresh command is sent as a broadcast command, thus all RDRAM devices are refreshed at the same time. The memory controller keeps track of which banks to refresh by incrementing the bank address each time a refresh is generated. The RDRAM device keeps track of which row in the bank to refresh by incrementing the row address in the RDRAM device after a row within the last bank of the RDRAM device has been refreshed. When in the nap or power down power states, RDRAM devices refresh themselves with no memory controller overhead. Every period equal to $t_{CTRL} / (\# \text{ of RDRAM devices})$, the memory controller sends current calibration (CAL/SAM) COLX packets to a specified RDRAM device. The RDRAM device adjusts its I_{OL} current. t_{CTRL} is about 100 ms.

Every t_{TEMP} period the memory controller broadcasts a TCEN/TCAL command in a ROWR packet. All RDRAM devices adjust the slew rate of their output drivers to compensate for temperature drift. t_{TEMP} is about 100 ms.

Slide #70

Direct Rambus Configuration

- RDRAM Device has two CMOS input pins: SCK (1 MHz) and CMD
- RDRAM Device has two CMOS I/O pins: SIO0 and SIO1
- Pins provide memory controller access to RDRAM device control registers
- Register Read or Write transactions consist of 4 packets
- Initialization establishes:
 - SDEVID device ID of 32 possible RDRAM devices
 - DEVID for each RDRAM device
 - TRDLY value which sets latency value for memory read data

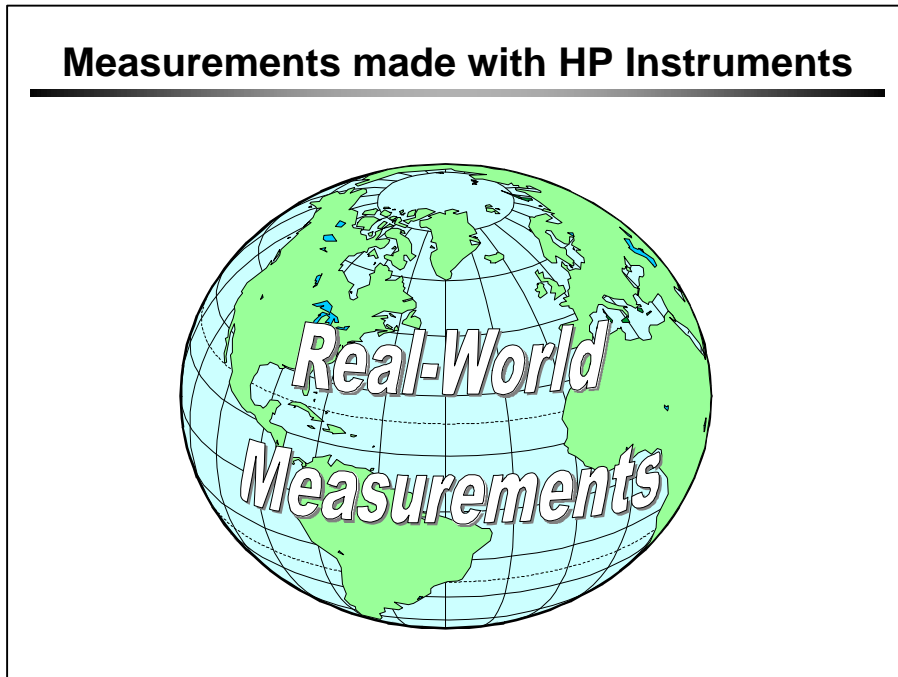
The memory controller initiates configuration read requests to read the read-only control registers of an RDRAM device. The memory controller processes this information and then writes to the read/write registers to place the RDRAM devices in the correct operating state. The important control registers that are initialized include the DEVID and TRDLY register.

Slide #71

Agenda

- **Direct Rambus Overview**
- **Direct Rambus Channel Electrical Architecture**
- **Direct RDRAM Device Physical Architecture**
- **Direct Rambus Logical Architecture**
 - **Read and Write Transaction Protocol**
 - **Transaction Pipeline**
 - **Packet Format**
 - **Unique Concepts**
 - **Maintenance Operations**
 - **Configuration**
- ***Measurements made with HP Instruments***

Slide #72



Measurements made with Hewlett-Packard Instruments

Slide #73

HP 16700A Logic Analysis System




The HP 16700A series logic analysis systems share an intuitive, easy-to-use, multi-window interface and common capabilities. In multiple-bus target systems, cross-bus measurements are key to designing and validating system performance. HP's powerful analysis tools have been designed to accommodate targets, such as, computer motherboards.

Slide #74

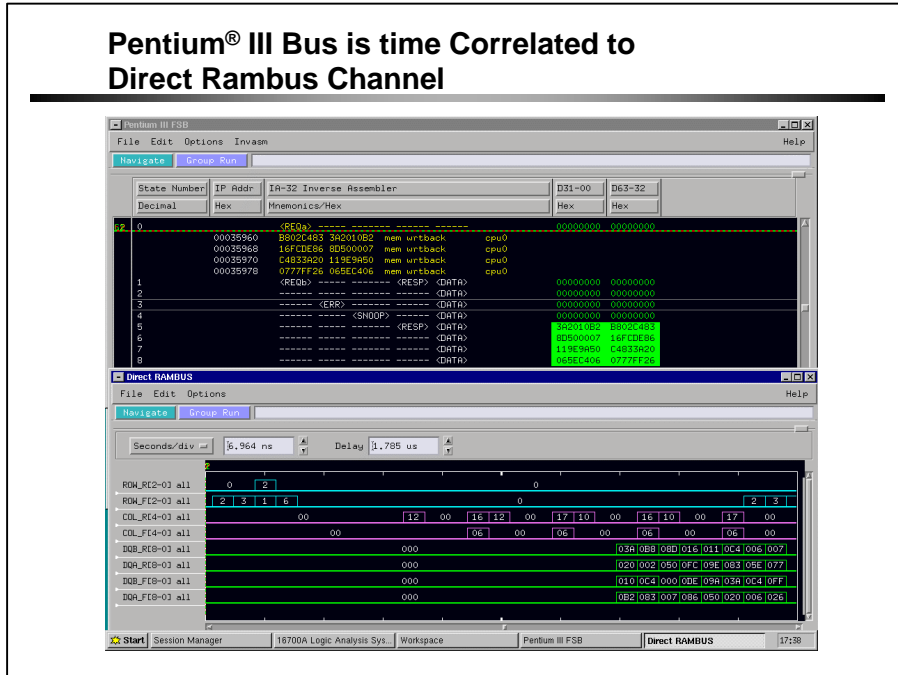
HP E2487C Analysis Probe

HP E2487C Analysis Probe for Intel Pentium® III,
Pentium II and Xeon processor-based systems



The HP E2487C analysis probe and HP E2492B/C probe adapters are used with the HP 16700A logic analysis system to observe Pentium® III, Pentium II and Xeon™ processor-based systems. The HP E2487C analysis probe provides a transaction tracker and inverse assembler that can decode Intel processor transactions into bus operation mnemonics. Filter options allow you to selectively list transactions by agent and transaction type.


Slide #75



The screen shot above was taken from an Intel Pentium III processor system and displays time-correlated data from the CPU 0 bus to the Direct Rambus channel.

The top part of the screen shows several transactions of a memory write back from CPU 0. The data that is being written back is shown in the display. The bottom part of the screen shows Direct Rambus signals in a waveform format. If you look closely you can see the same data that was on the CPU 0 processor bus.

Slide #76



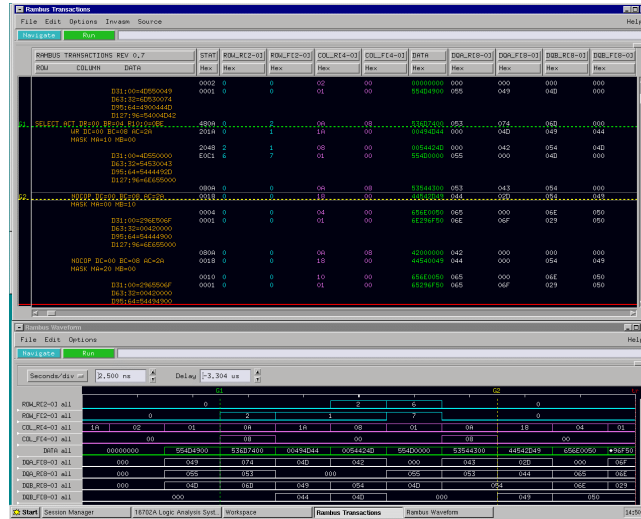
The Direct Rambus analysis probe from FuturePlus Systems (A HP Premier Channel Partner) samples and demultiplexes the Direct Rambus channel. (www.futureplus.com)

For Direct Rambus analysis the following components are needed:

- One FS2222 (HP FSI-60033) Direct Rambus analysis probe;
- One HP 16700A or HP 16702A logic analysis system;
- One HP 16517A 4GHz timing/1GHz state logic analyzer master module; and
- Three HP 16518A 4GHz timing/ 1GHz state logic analyzer expansion modules.

Slide #77


Rambus Analysis Probe provides decoding Direct Rambus instructions with Inverse Assembly



The screen shot above also was taken from an Intel Pentium III processor system. The Direct Rambus analysis probe can show bus transactions displayed in Direct Rambus mnemonics. The top part of the screen has several packs of data that have been decoded by the analysis probe. The Direct Rambus probe captures the data at full channel speed (400 MHz clock speed or 800 MHz data acquisition speed). The bottom part of the screen displays the state data in a waveform format.

Slide #78

Computer Verification Tool



The HP E2920 is the most comprehensive PCI test solution in the industry, encompassing the entire development cycle from bring-up and debug, to I/O system performance characterization, and chip and system validation.


The hardware consists of a powerful PCI analyzer/exerciser on a single card, which supports all flavors of PCI, such as, 32 and 64 bit, 0 to 66 MHz.

The System Validation Package (SVP) is the latest software for the HP E2920 family. This software includes a test suite for verification of all system data paths including the Direct Rambus channel.

The SVP uses the PCI bus as a test port within the system.

Slide #79

HP's System Validation Package helps reveal problems as early as possible!



The screenshot above is from the HP E2920 family System Validation Pack (SVP), showing 2 test scenarios (top), 5 test card resources (middle) and test status report window (bottom).

The HP E292X test cards are distributed within the I/O architecture PCI slots. When the software (which resides on the target system) is started, each HP card is automatically configured as a test resource.


Then, the user can associate a test from the test suite with a test card from the list of available test resources. Next, each test card automatically stresses the data path and checks for errors.

For example, test scenario 1 shown above creates traffic from the CPU to system memory and simultaneously from the HP test card to the same system memory page. Addresses are interleaved so the paths should not interfere.

Stress is created by varying the traffic profile from the I/O side. For example: burst size, command type, cache alignment, and wait states are all pseudo-randomly permuted to create worst case scenarios while checking for data integrity. Timing and protocol are also checked on the PCI side.

Slide #80

HP Infiniium 8GSa/s Single-shot Oscilloscope

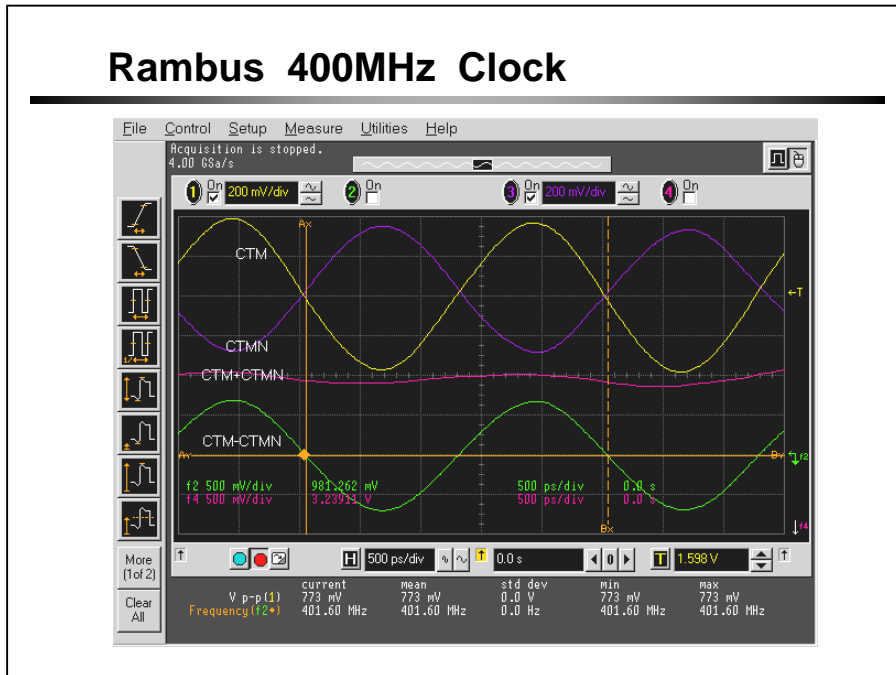


Debugging high-speed bus designs, such as Direct Rambus, often requires a logic analyzer and a high-bandwidth oscilloscope for characterizing the signal integrity of the bus. However, a limitation of high-bandwidth (>2 GHz) oscilloscopes, is that a repetitive signal is needed to take advantage of their measurement capabilities.

Often, a single-shot oscilloscope is needed to capture elusive problems that may occur only once. The HP 54845A Infiniium oscilloscope is a four-channel, 1.5 GHz oscilloscope that offers a maximum sample rate of 8 GSa/sec. This scope is capable of capturing single-shot events up to the full analog bandwidth of the scope (1.5 GHz).

You can use the extensive triggering capability of a logic analyzer to send an external trigger to the HP Infiniium oscilloscope to tell the scope when to start sampling the signal. Logic analyzers, which can be used to locate faults, can trigger on a sequence of events across many channels. Then, the logic analyzer can trigger the HP Infiniium oscilloscope so that it can capture the high-speed anomaly with its 8 GSa/sec sample rate.

Slide #81




The HP Infiniium scope also has extensive internal triggering to help track problems in high-speeds designs. Some examples of its triggering capabilities are:

- Glitch
- Pattern
- State
- delay by time or events
- runt pulse
- setup/hold time
- pulse width
- transition (rise & fall time)

You also can easily document and share measurement results with others by connecting the HP Infiniium scope's built-in LAN interface to your network. You can annotate waveforms with the scope's keyboard, and then save the screen in electronic format so that it can be incorporated into documentation packages later. In addition, you can print the screen to any compatible printer on your network. Lastly, you can transfer setups to other HP Infiniium oscilloscopes on your network.

Slide #82

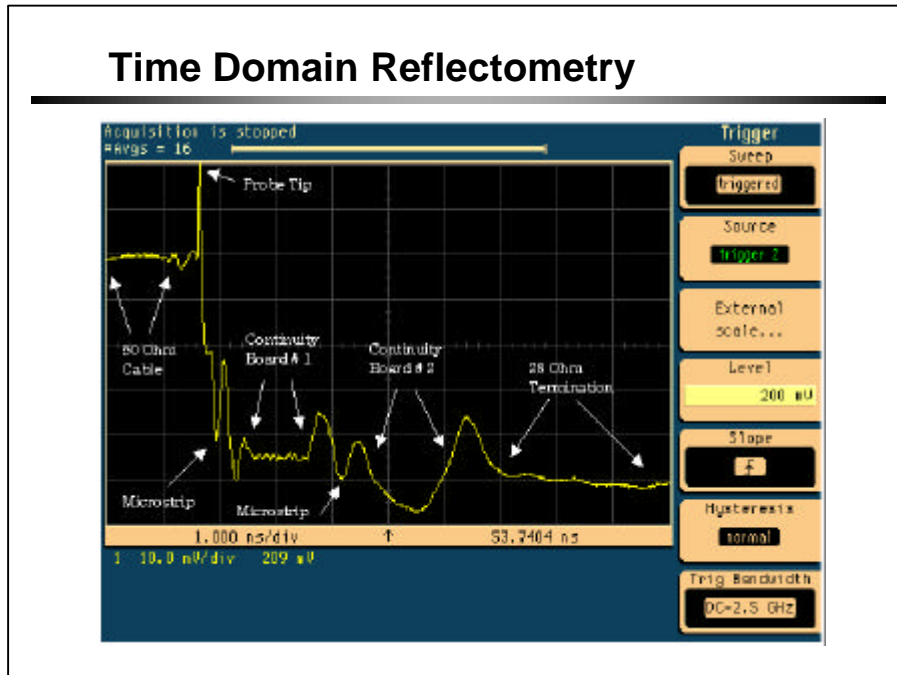
Verify RAMBUS Signal Integrity with HP 54750A Digitizing Oscilloscope



Over the last 10 years, computer processor architects and designers have done a remarkable job by increasing CPU performance 200 fold. Main memory subsystems haven't kept pace, only improving by a factor of 20 in the same timeframe. This gap has been bridged in part by improved and more sophisticated caches. However, larger data structures, associated with images, are demanding more performance from the memory system.

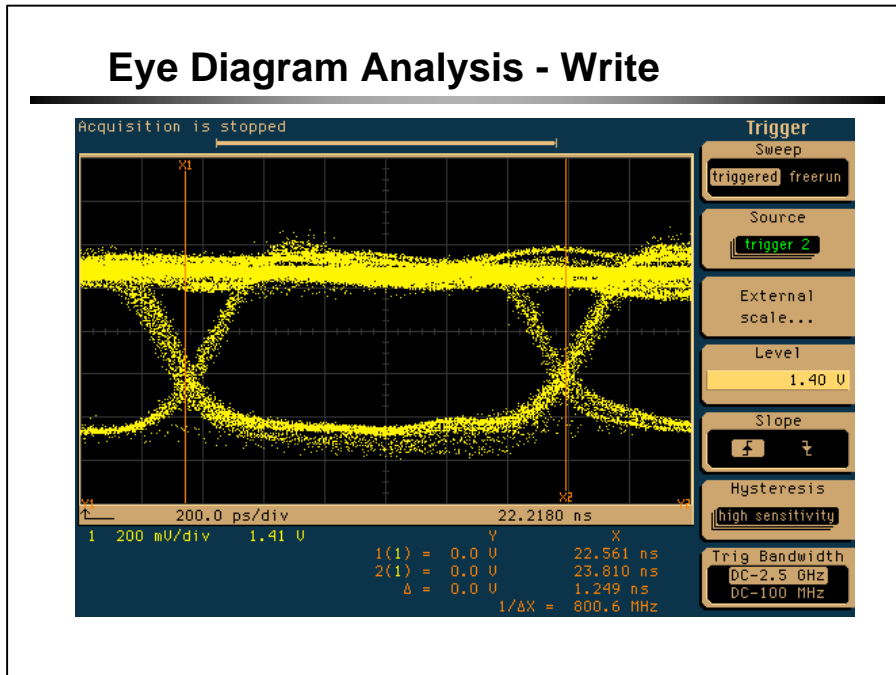
Direct Rambus is a high-speed digital bus which requires circuit board traces that are impedance controlled. The traditional logic analysis measurement technique now needs to be complemented with a signal integrity measurement tool. For digital design engineers involved in signal integrity design, time domain reflectometry (TDR) is the tool of choice. New insight can be gathered by using a digitizing oscilloscope and a plug-in TDR module to measure eye diagrams. Viewing jitter in this manner is helpful when trying to minimize clock and data skew over each line of the bus.

Slide #83



The TDR waveform shown above was measured on the Direct Rambus channel from a computer platform system. It shows the physical layout as seen by the 45 ps rise time pulse generator in the HP TDR module. As you can see, we start with the 50 ohm coax cable connected to the probe. The waveform indicates an inductive impedance discontinuity at the probe. Then, the waveform drops down to a short segment of motherboard into the Rambus connector on the motherboard. Plugged into this connector is a Rambus Continuity Module. The continuity module is similar to a DIMM (dual-in-line-memory module) and exhibits the same characteristic impedance as the Direct Rambus itself (28 ohms).

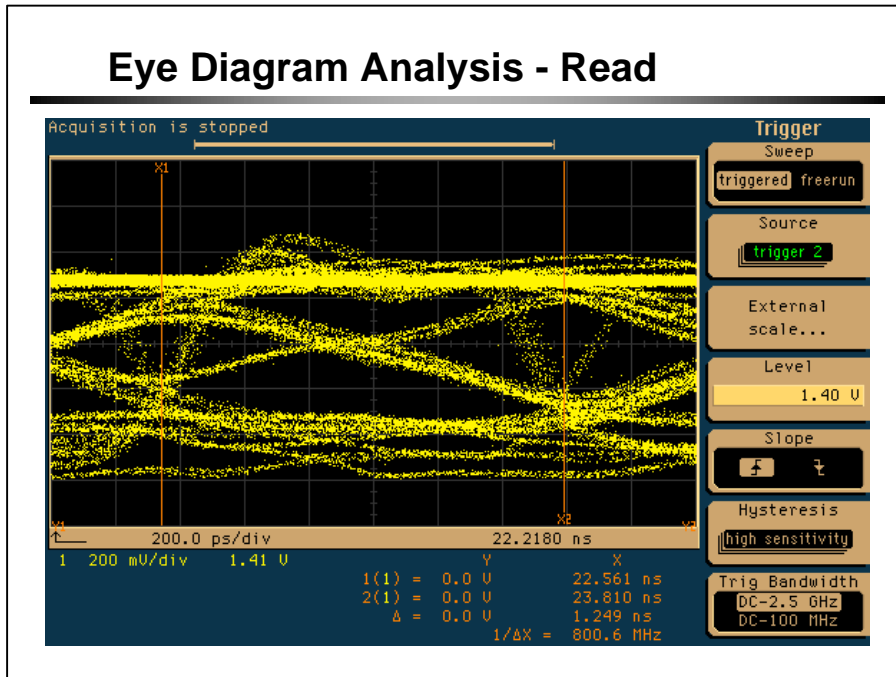
Slide #84



Signal jitter can be measured on a Direct Rambus line by using eye diagram analysis. The above screenshot shows approximately two unit intervals (two bit cells). Common eye diagram measurements include eye height, eye width, crossing %, RMS jitter, and duty cycle distortion.

With a high-persistence display, heavy data traffic is indicated by a brighter line across the top of the eye. Because we know that Rambus is an active low logic, this tells us that we have long periods of time when only zeros are written to the RIMM from the processor. Less common data transitions seen in this unit interval are the 010 (high-low-high) and 101 (low-high-low) patterns. This eye diagram is characteristic of the Direct Rambus write operation.

Slide #85



The Direct Rambus read operation is complex and, therefore, more complicated to interpret using eye diagrams. In fact, the only meaningful information that can be obtained from this Rambus line is that it is performing both reads and writes.

The Direct Rambus read operation yields an eye diagram which is a half-height eye compared to the write eye diagram. Also, complicating the eye diagram is the fact that read data reflexes back into the channel after it reaches the controller. The eye diagram above indicates data traffic which has both reads and writes occurring on the same line within the bus.

Slide #86

To learn more

- **Books by MindShare:**
 - ◆ **Pentium® II/ Pentium Pro System Architecture**
 - ◆ **AGP System Architecture**
- **On-site training offered by MindShare**
 - ◆ **1-800-633-1440**
- **Visit MindShare's web site at www.mindshare.com**
- **email: mindshar@interserv.com**