



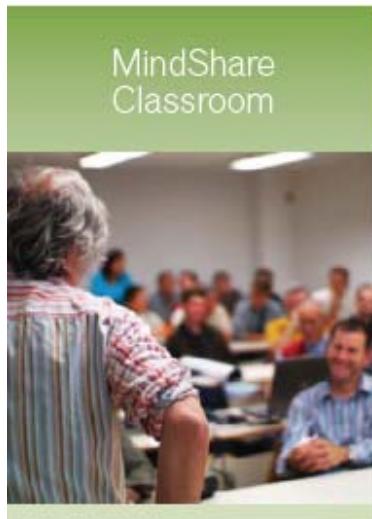
Modern DRAM (DDR2/DDR3) Architecture



This presentation is copyrighted material and is intended solely for the use of students who have attended this MindShare course. Any other distribution is not allowed.

Please do not copy or distribute our material without permission.

training@mindshare.com
1-800-633-1440



MindShare Classroom



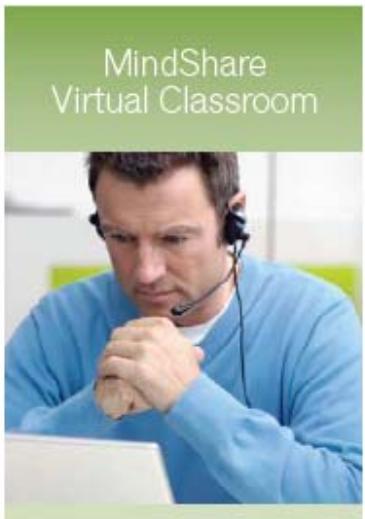
In-House Training



Public Training

Classroom Training

Invite MindShare to train you in-house, or sign-up to attend one of our many public classes held throughout the year and around the world. No more boring classes, the 'MindShare Experience' is sure to keep you engaged.



MindShare Virtual Classroom



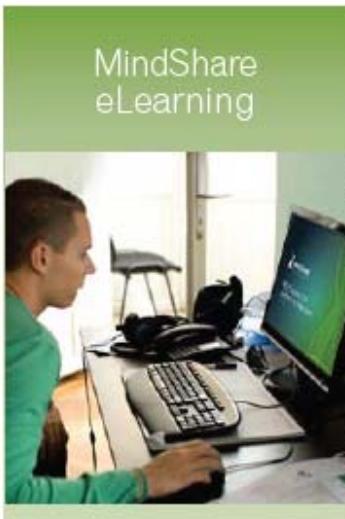
Virtual In-House Training



Virtual Public Training

Virtual Classroom Training

The majority of our courses live over the web in an interactive environment with WebEx and a phone bridge. We deliver training cost-effectively across multiple sites and time zones. Imagine being trained in your cubicle or home office and avoiding the hassle of travel. Contact us to attend one of our public virtual classes.



MindShare eLearning



Intro eLearning Modules



Comprehensive eLearning Modules

eLearning Module Training

MindShare is also an eLearning company. Our growing list of interactive eLearning modules include:

- **Intro to Virtualization Technology**
- **Intro to IO Virtualization**
- **Intro to PCI Express 2.0 Updates**
- **PCI Express 2.0**
- **USB 2.0**
- **AMD Opteron Processor Architecture**
- **Virtualization Technology**
- ...and more



MindShare Press



Books



eBooks

MindShare Press

Purchase our books and eBooks or publish your own content through us. MindShare has authored over 25 books and the list is growing. Let us help make your book project a successful one.



Intel Architecture

- Intel Core 2 Processor (Penryn)
- Intel Core Processor (Nehalem)
- Intel QuickPath Interconnect (QPI)
- Intel 64 and IA-32 Architecture
- Intel PC and Chipset Architecture
- Intel Core 2 Processor and Chipset Combo

AMD Architecture

- AMD Opteron Processor (Barcelona)
- AMD64 Architecture

x86 Programming

- Assembly Language Programming
- System Programming for the x86 Architecture

IO Buses

- PCI Express 2.0
- USB 3.0
- Embedded USB 2.0 & Workshop
- PCI / PCI-X

Memory Technology

- Modern DRAM Architecture (DDR2/DDR3)

Virtualization Technology

- PC Virtualization
- IO Virtualization (IOV)

Storage Technology

- SAS Architecture
- Serial ATA Architecture



➤ In-House classroom 

➤ Virtual classroom 

➤ eLearning Courses 

➤ Check our website for Public course offerings

➤ PCI Express 2.0

➤ USB 3.0

➤ Embedded USB 2.0 & Workshop

➤ Protected / Long Mode Programming

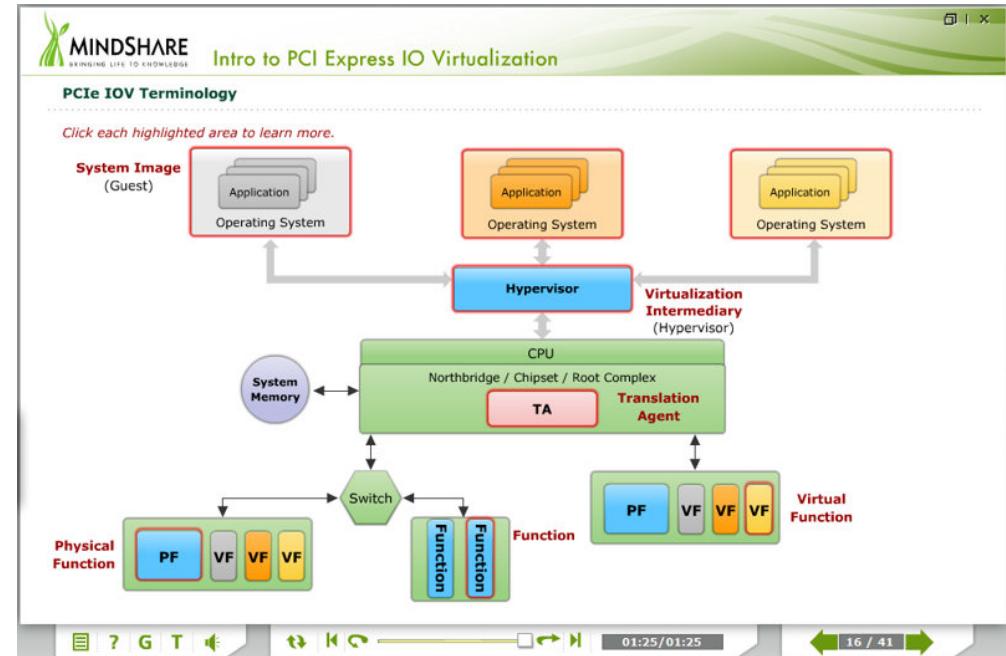
➤ Modern DRAM Technology

➤ High-Speed Design

➤ We can customize any of our classes to meet your budget and content needs



- Some of our eLearning courses include:
 - Comprehensive PCI Express
 - Intro to PCI Express Changes
 - Intro to PCI Express IO Virtualization
 - Intro to Virtualization Technology
 - and more



- Visit www.mindshare.com for a full list of our eLearning courses.



MINDSHARE
BRINGING LIFE TO KNOWLEDGE

CART CONTACT US MY ACCOUNT

Search MindShare

LEARN a subject with the experts SHOP for a learning experience CONNECT with the standard in training ENGAGE with content creation with MindShare

Shop > SATA Storage Technology

Books + eBooks

EMAIL TO A FRIEND HOW eBooks WORK eBooks FAQ

All Available Titles

- HyperTransport 3.1 Interconnect Technology
- SATA Storage Technology**
- SAS Storage Architecture
- PCI Express System Architecture
- The Unabridged Pentium 4
- PCI System Architecture (4th Edition)
- PCI-X System Architecture
- InfiniBand Network Architecture
- Universal Serial Bus System Architecture
- HyperTransport System Architecture
- FireWire System Architecture (2nd Edition)

Most Popular Titles

- PCI Express System Architecture
- HyperTransport 3.1 Interconnect Technology
- Universal Serial Bus System Architecture

Most Popular eBooks

- The Unabridged Pentium 4
- PCI Express System Architecture
- SATA Storage Technology

SATA Storage Technology

Author(s) Don Anderson
Publisher MindShare Press
Pages 496
Retail Price \$59.99

ADD BOOK TO CART Book Price \$55.99

ADD eBook TO CART eBook Price \$32.00

ADD TO CART Book + eBook Price \$71.99

PEEK INSIDE

"This SATA Storage Technology book from MindShare is the ultimate tutorial on the SATA specification and its implementation. Any SATA based product developer will find the clear and concise explanations of the various features found in this protocol to be invaluable. Particular emphasis is placed on the sections of the spec that enable SATA to be such an important interface for enterprise storage systems."

- Mark Adams, Product Marketing Manager of the LeCroy Corporation

MindShare's **SATA Storage Technology** book provides ATA background information necessary to understand the Serial ATA implementation. This book provides detailed information regarding the serial protocol, commands, and electrical interface. A tutorial approach is taken to facilitate thorough understanding, using numerous illustrations, tables and examples. The text is also designed to provide easy to access valuable reference information.

The SATA (Serial ATA) interface is intended to expand and eventually replace the parallel ATA interface that has been the standard connection for hard drives and CD-ROM/DVD drives in Personal Computers for nearly twenty years. SATA, a high-speed serial version of ATA, is designed to maintain compatibility with software written for the standard ATA implementation, thus easing the transition to SATA. New advances in SATA provide new capabilities such as

PCI EXPRESS SYSTEM ARCHITECTURE

UNIVERSAL SERIAL BUS SYSTEM ARCHITECTURE

SECOND EDITION

SAS Storage Architecture

HyperTransport 3.1 Interconnect Technology



Register on MindShare's Website

8

- Receive our quarterly newsletter
- Download course presentations and eBooks
- Shop for eLearning modules and Books/eBooks
- Register for Public courses

The screenshot displays the homepage of the MindShare website. At the top, there are navigation links for 'LEARN', 'SHOP', 'CONNECT', and 'ENGAGE'. The main header reads 'MINDSHARE' and 'THE MINDSHARE QUARTERLY NEWSLETTER'. Below this, the 'sept 08 technologybrief' is shown. The page is divided into several sections:

- ALUMNI MODULES:** Intro to PCI Express® 2.0 Changes, Public Courses, Memory, SATA, USB, PCI, and more.
- IN THIS ISSUE:** PCI Express 2.0 Changes, Public Courses, Memory, SATA, USB, PCI, and more.
- UPCOMING PUBLIC ENROLLMENT COURSES:** Q3, 2008, including PCI Express 2.0 Architecture, Embedded USB 2.0 Architecture, Embedded USB 2.0 Workshop, PCI Express 2.0 Architecture, and PCI Express 2.0 Compliance.
- ALUMNI MODULES:** Intro to PCI Express® 2.0 Changes, Purchase Module, Learn More.
- BOOKS + eBOOKS:** HyperTransport 3.1 Interconnect Technology, Authors: Brian Holden, Dan Anderson, Publisher: Mindshare Press, Price: \$95.99, Purchase Book/eBook, Learn More.
- MINDSHARE COURSE LIST:** World-Class Technical Training, including PCI Express 2.0, Intel® Core™ Processor Architecture, AMD Opteron Processor Architecture, Intel® 64 & IA-32 Software Architecture, Intel® PC and Chipset Architecture, PCI Virtualization, USB 2.0, and Wireless USB. Courses Coming Soon: Intel® QuickPath Interconnect, Intel® Core™ (Nehalem) Processor, and USB 3.0.
- TECHNICAL WHITE PAPERS:** PCI Express® Compliance Using a Predictable Metrics Based Verification Closure Methodology, by Mike Berlin and Johnathan ClearSpeed Technology Pte. This article from ClearSpeed Technology and Cedex describes the methodology, tools used, and implementation guidelines employed in designing a PCI Express based SOC. The paper also highlights technical and business benefits of using this approach.
- NEWSROOM:** Press Release - 8/15/2008, Intel® and Mindshare to develop courses and train customers on the new Intel® QuickPath Interconnect (QPI) Technology and Intel® Core™ (Nehalem) Processor. Press Release - 8/16/2008, Mindshare's partner, Agilent Technologies, introduces Active State Power Management support in their PCI Express testing tools.
- MINDSHARE BROCHURE:** Now Available MindShare's Service Brochure, Download Now.



System Architecture

DRAM Feature Summary

Intro to DRAM

 Historical Background

 Why DRAM?

DRAM Cell Architecture

DRAM Chip Architecture

DRAM Modules

Commands and Waveforms

DDR Initialization

 SMBus Overview

Electrical Specifications

DDR1 and DDR2 Routing

On-Die Termination

Errors

Min Huang(min.huang@ lecroy.com)

Do Not Distribute

Additional DDR3 Topics

 Timing and Electrical Differences

 Fly-by Routing Read Example

 Read Calibration

 Fly-by Routing Write Example

 Write Leveling

 DDR3 On-Die Termination

 ZQ Calibration

 Reset

 On-DIMM Address Mirroring

DRAM Controller Basics

 Addresses

Alternative DRAM Solutions



Upon completion, the student should:

- Understand the DRAM bank state diagram
- Understand the JEDEC initialization procedure
- Understand the timing waveforms
- Know why the DRAM controller has to do so much of the work
- Know what motivates the move to DDR3 and the associated problems
- Know what elements of the system determine the address of a DRAM cell
- Know these terms and their concepts:

Prefetch width

Refresh vs. auto refresh vs. self refresh vs. auto self refresh

Activate

Precharge vs. auto precharge vs. precharge all

Additive latency

Bank vs. rank

SSTL

ODT, dynamic ODT

ZQ calibration

Fly-by routing, read calibration, and write leveling

SPD and Mode Registers

On-the-fly burst chop (the Austin Powers mode)

Unbuffered vs. registered vs. fully-buffered

1T vs. 2T timing

Remember the
quick reference
handout



Where's the Money?

DRAM Maker	2007 Revenue \$ billions
Samsung	8.7
Hynix	6.7
Qimonda	4.0
Elpida	3.8
Micron	3.2
Nanya	1.6
Powerchip	1.4
ProMos	1.1
Etron	0.4
Elite	0.2
Others	0.4
Total	31.5

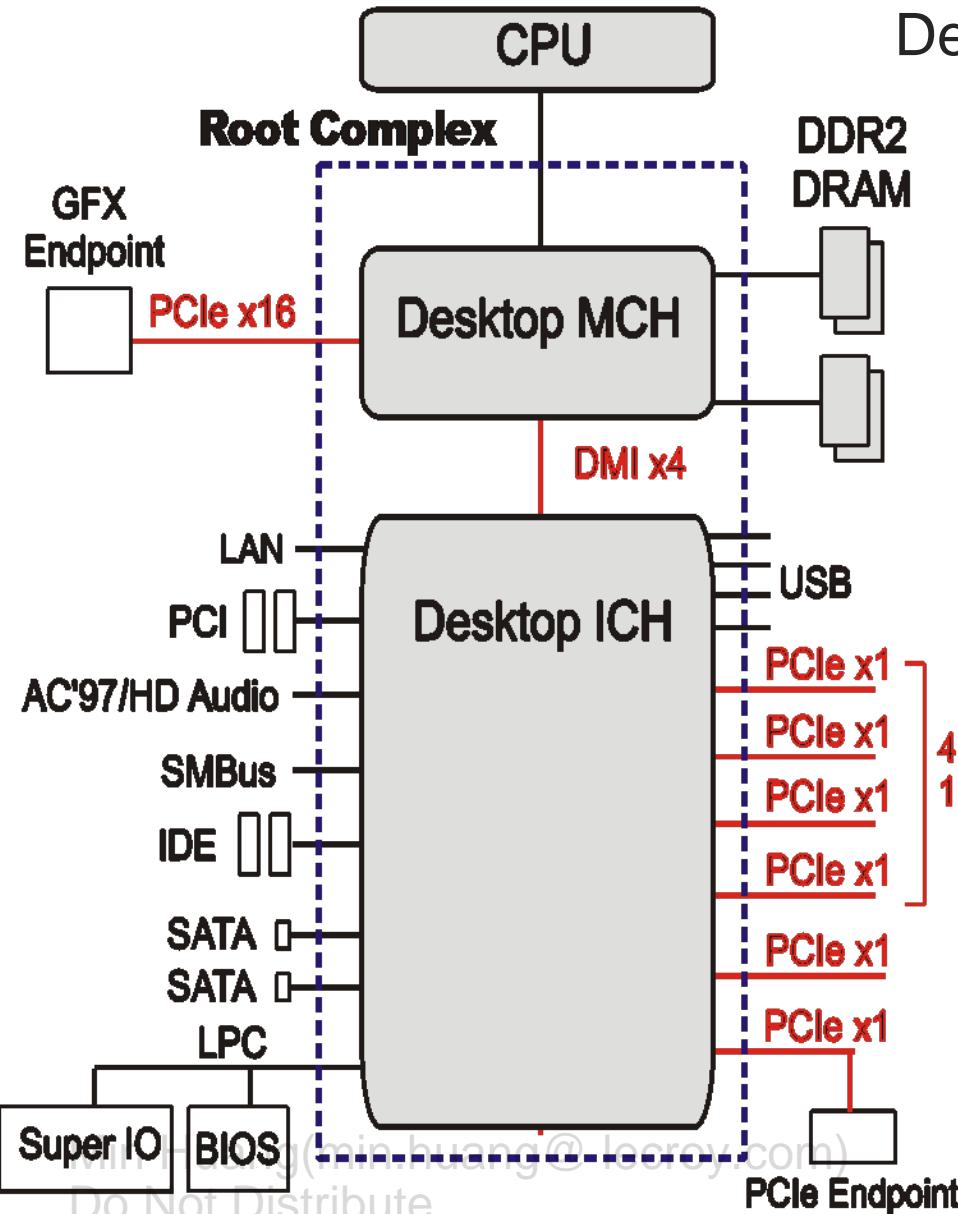
Logic Maker	2007 Revenue \$ billions
Intel	39.2
AMD	6.3

So who has the greatest responsibility
for new features such as fly-by?
Follow the money.



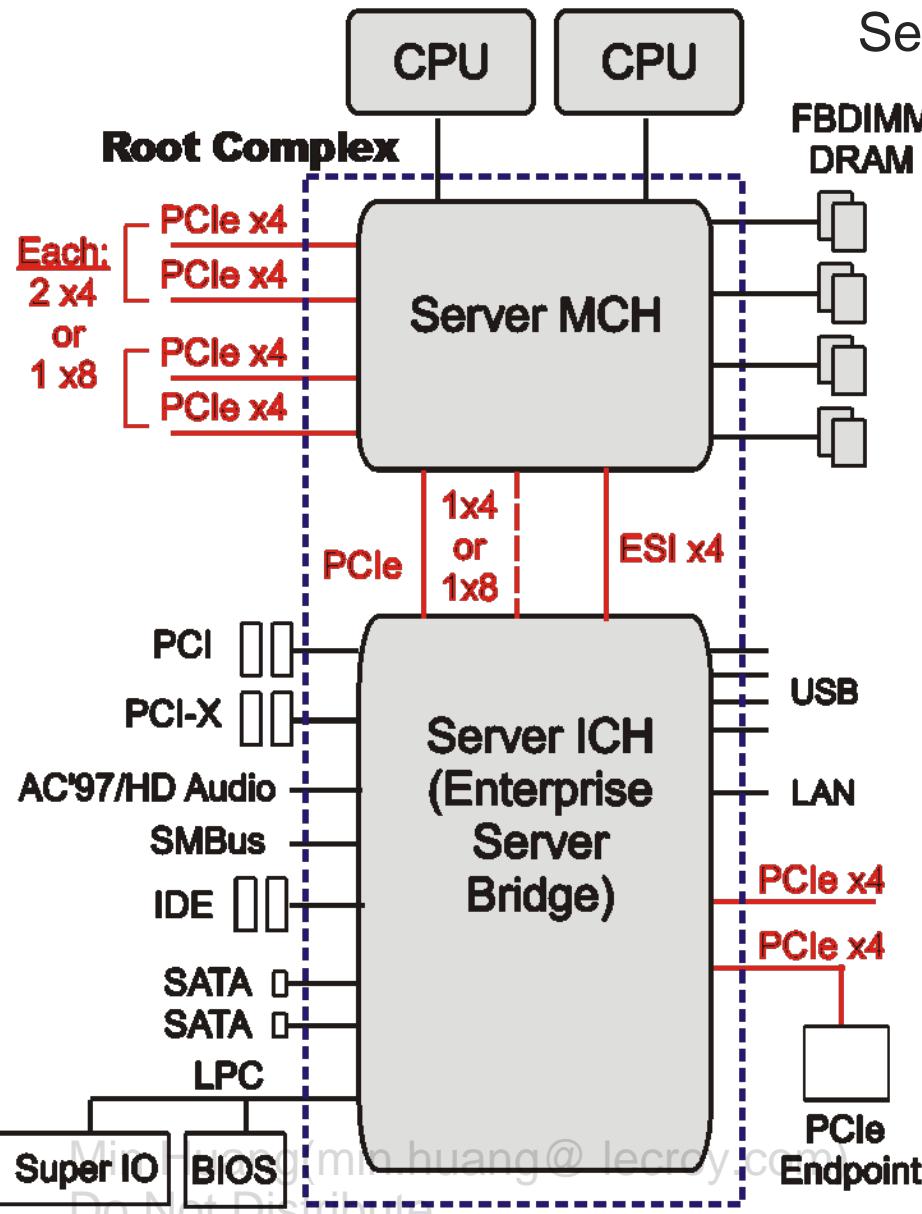
System Architecture

Min Huang(min.huang@ lecroy.com)
Do Not Distribute



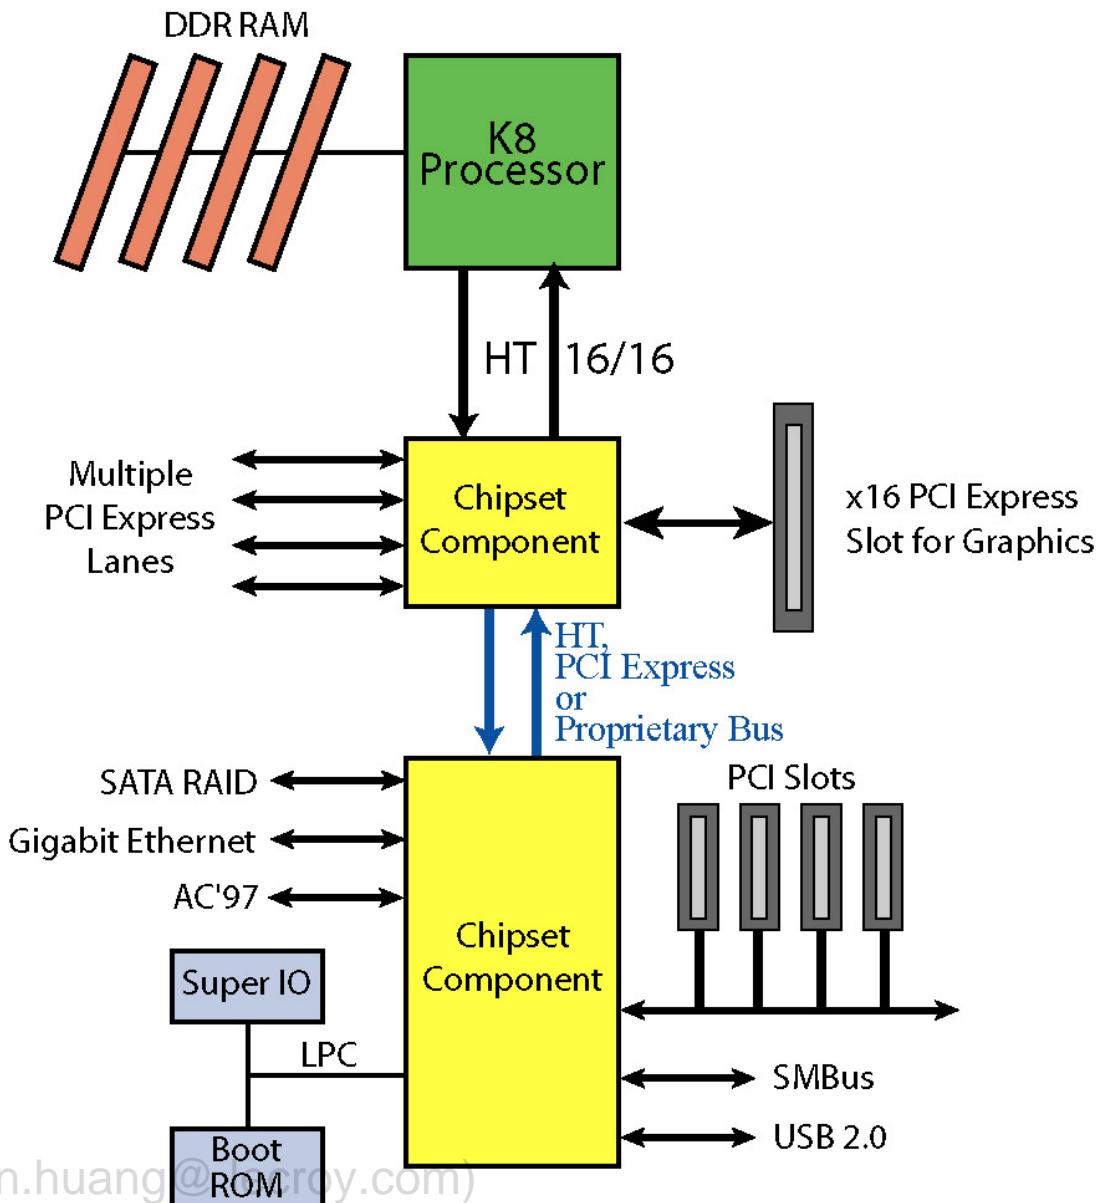
Desktop Chipset PCIe Support

- MCH hosts a 16-bit PCIe Link interface for high-performance graphics.
- DMI (Direct Media Interface) is an Intel-specific version of x4 PCIe.
- ICH (IO Controller Hub) supports multiple x1 PCIe Links, four of which may operate as x1 Links or ganged together as a single x4 Link.



Server Chipset PCIe Support

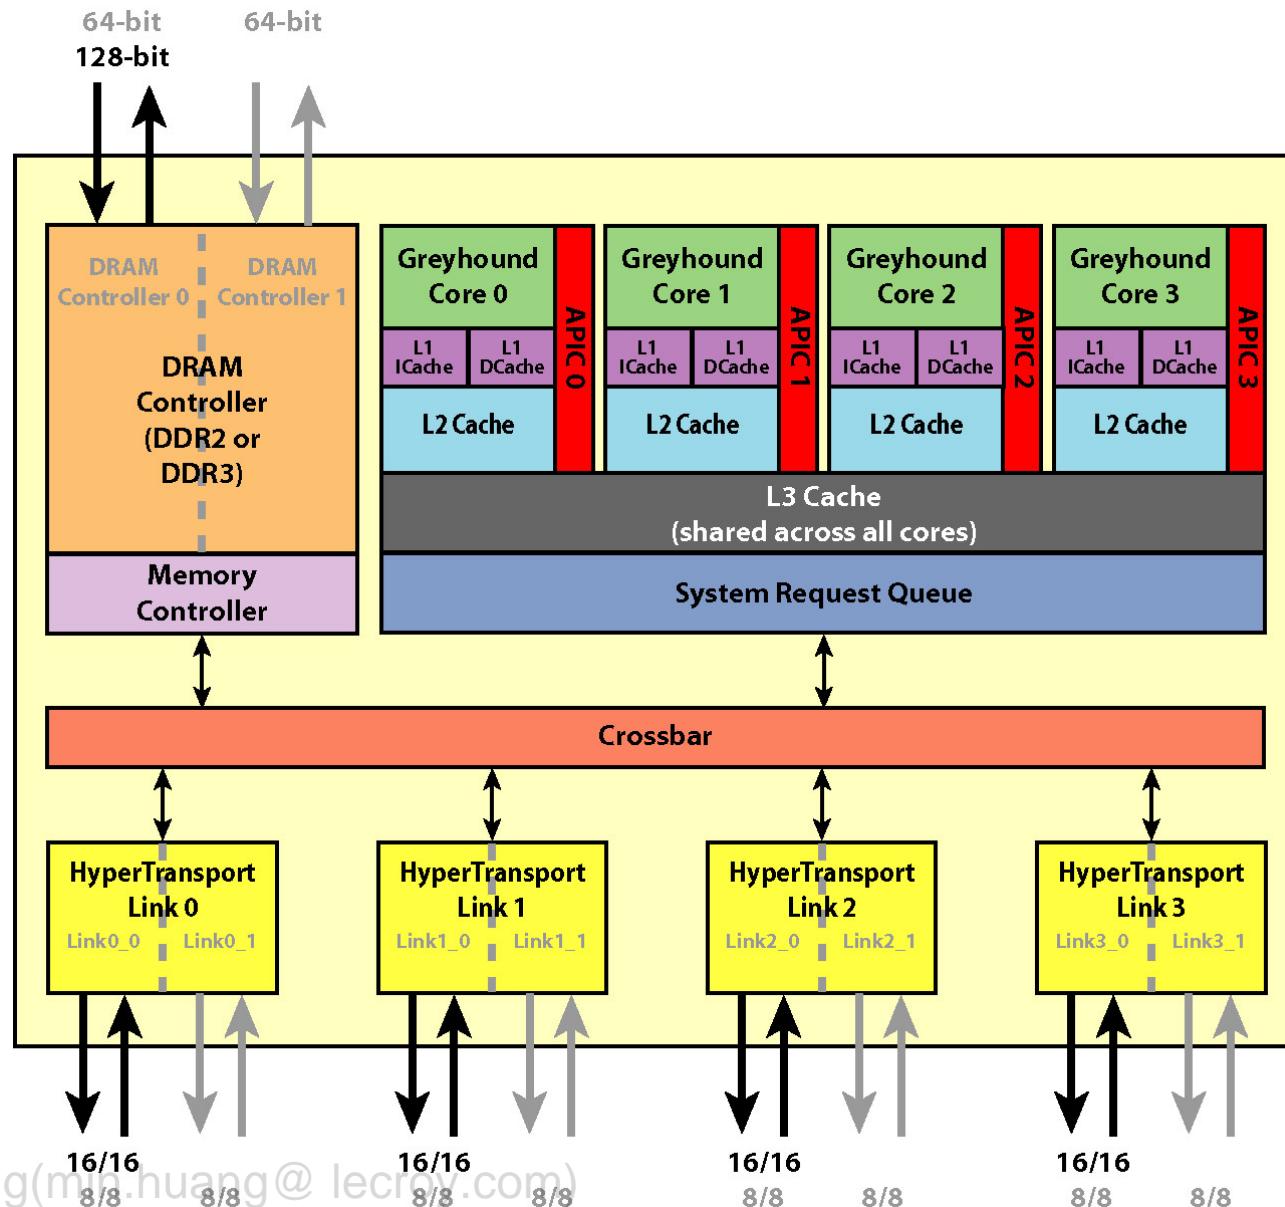
- MCH hosts 16 PCIe Lanes, organized as 2x8, 4x4, etc.
- ESI (Enterprise Server Interface) is a variant of x4 PCIe used for general chipset traffic.
- An additional x4 or x8 Link can connect the MCH and ICH for improved DMA performance.
- The Enterprise Server Bridge ICH supports two additional PCIe x4 Links.
- Note presence of both PCI and PCI-X slots.





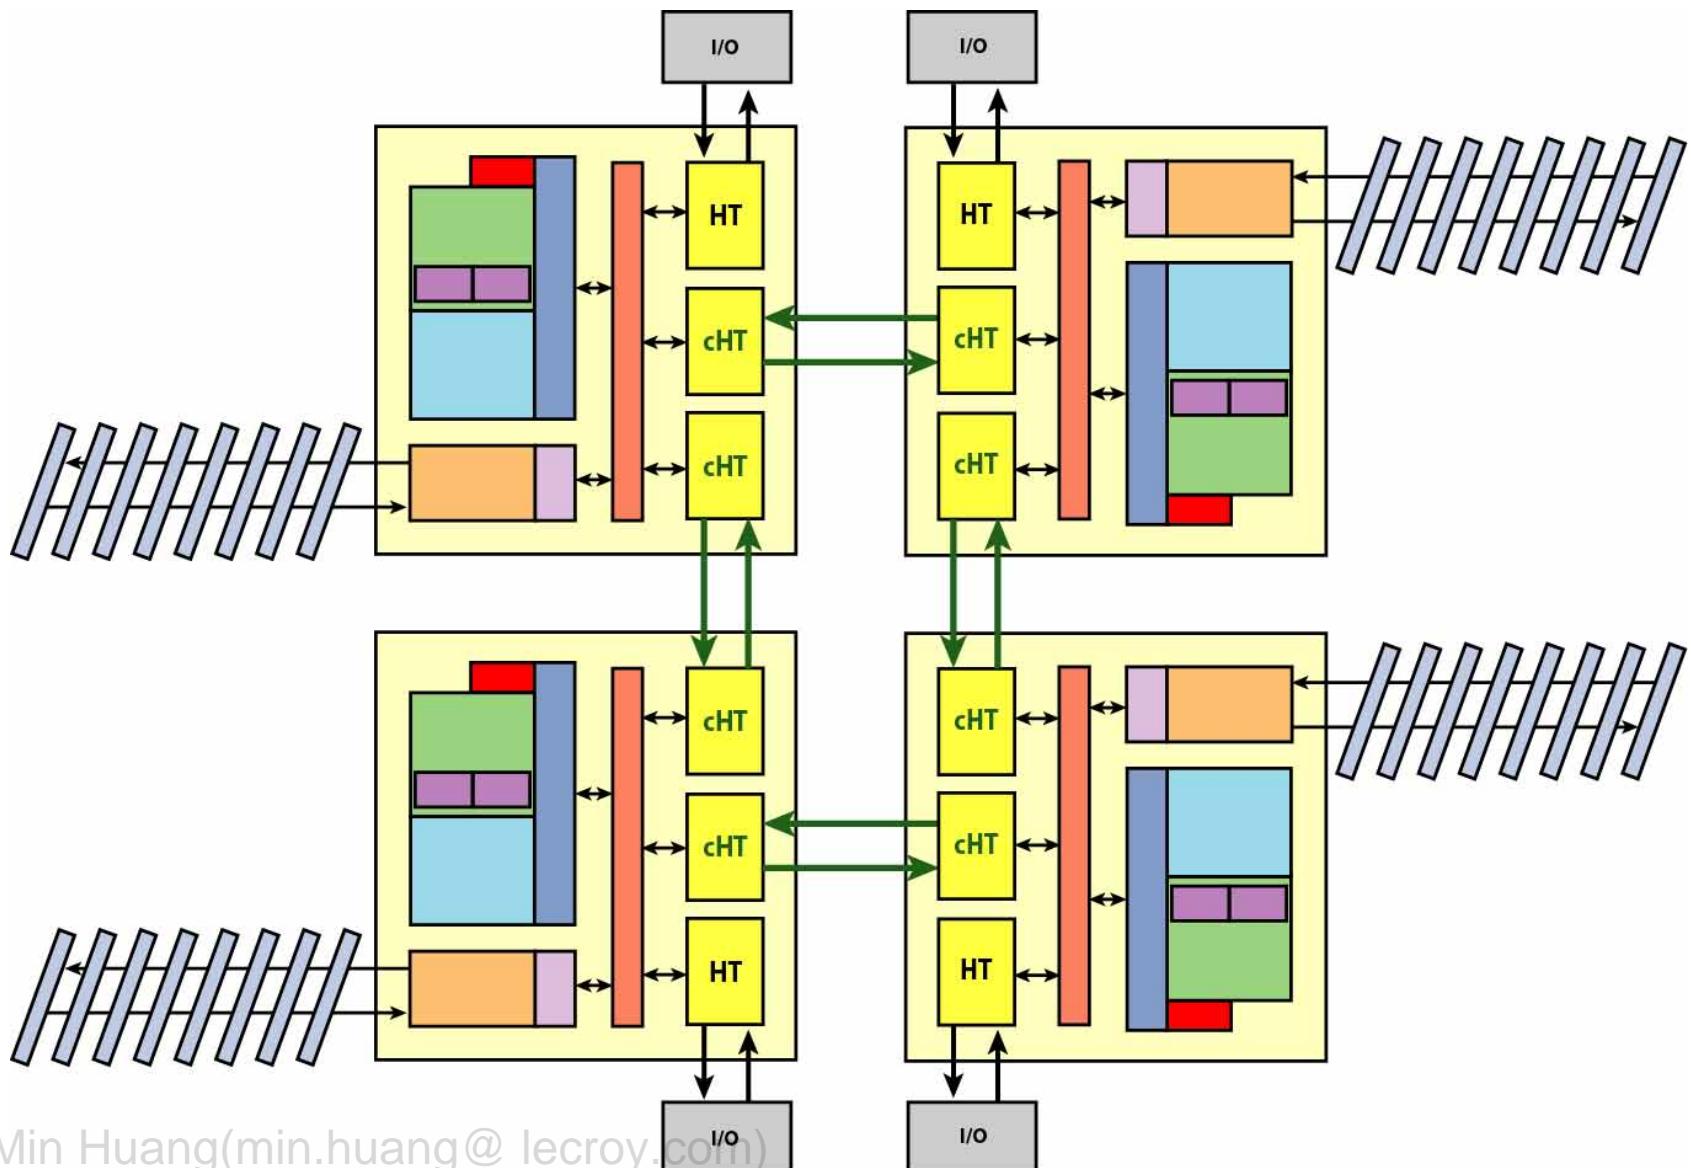
AMD Quad Core Greyhound Processor

16



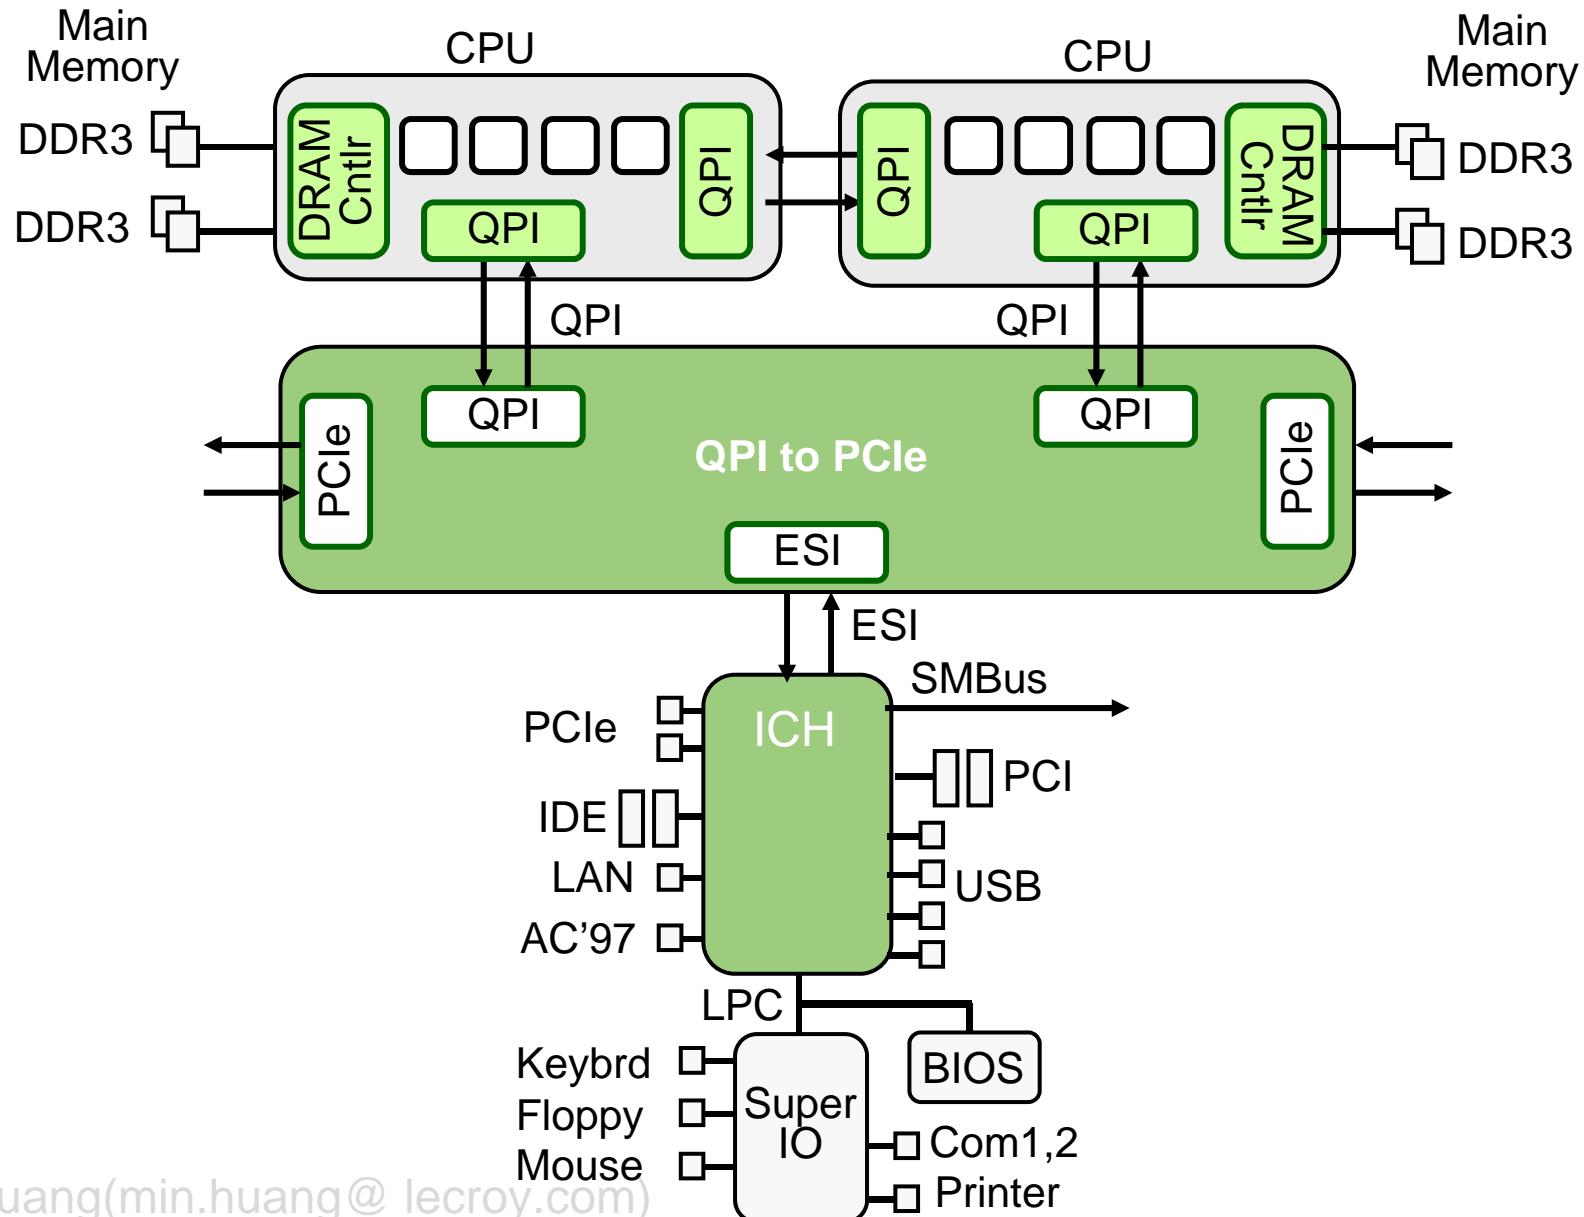


Example 4-Way System





Possible Intel Nehalem Topology





DRAM Feature Summary

Min Huang(min.huang@ lecroy.com)
Do Not Distribute



Device Architecture	Explanation	Clock Frequency (MHz)	Data Bus Rate (MT/s)	Bandwidth per 64-bit channel (MB/s)
FPM DRAM	<u>Fast Page Mode DRAM</u>	NA	-	-
EDO DRAM	<u>Extended Data Out DRAM</u>	NA	-	-
SDR SDRAM PC66 – PC133	<u>Single Data Rate Synchronous DRAM</u>	66 – 133	66 – 133	533 – 1066
DDR SDRAM DDR200 – DDR400 PC-1600 – PC-3200	<u>Double Data Rate Synchronous DRAM</u>	100 – 200	200 – 400	1600 – 3200
DDR2 SDRAM DDR2-400 – DDR2-800 PC-3200 – PC-6400	<u>Double Data Rate 2 Synchronous DRAM</u>	200 – 400	400 – 800	3200 – 6400
DDR3 SDRAM DDR3-800 – DDR3-1600 PC-6400 – PC-12800	<u>Double Data Rate 3 Synchronous DRAM</u>	400 – 800	800 – 1600	6400 – 12800



	SDRAM	DDR1	DDR2	DDR3
Performance	66-133MT/s	200-400MT/s	400-800MT/s	800-1600MT/s
V_{DDQ}	3.3 Volts	2.5 Volts	1.8 Volts	1.5 Volts
V_{TT}	NA	$\frac{1}{2} V_{DDQ}$	$\frac{1}{2} V_{DDQ}$	$\frac{1}{2} V_{DDQ}$
IO Interface Logic	LVTTL	SSTL_2	SSTL_18	SSTL_15
Organization	x4, x8, x16	x4, x8, x16	x4, x8, x16	x4, x8, x16
Density	16Mb-512Mb	64Mb-2Gb	256Mb-4Gb	512Mb-8Gb
Number of Banks	4	4	4 (256Mb-512Mb) 8 (1Gb-4Gb)	8
Package	TSOP	TSOP2/BGA	BGA	BGA (mirrored option)
Prefetch	1	2	4	8
Burst Length	1, 2, 4, 8, Page	2, 4, 8	4, 8	8 (chop 4)
Clock	Single Ended	Differential	Differential	Differential
Strobes	NA	Single-Ended (SE)	SE or Differential	Differential
DQ Driver Strength	Wide Envelope	Narrow Envelope	18 Ohm OCD	ZQ cal
Termination	NA	Mother Board	Mo Bo, Dyn ODT	DIMM, Dyn ODT
Read Latency	CL=1,2,3	CL=1.5, 2, 2.5, 3	CL=2, 3, 4, 5, 6	CL=5, 6, 7, 8, 9, 10, 11
Additive Latency	NA	NA	AL=0, 1, 2, 3, 4	AL=0, CL-2, CL-1
Burst Interrupts	Yes	Yes	R-R, W-W 4n only	Burst Chop

Since the mid 1990s, JEDEC (was the Joint Electron Device Engineering Council, now the Solid State Technology Association) has controlled the DRAM standards. Specs can be found at:

<http://www.jedec.org>

- DDR1 spec is JESD79E
- DDR2 spec is JESD79-2C
- DDR3 spec is JESD79-3C
- Package specs are MO-207 (for example)
- DIMM and SPD specs are JESD21C



Intro to DRAM

Min Huang(min.huang@ lecroy.com)
Do Not Distribute



Historical Background

Min Huang(min.huang@ lecroy.com)
Do Not Distribute

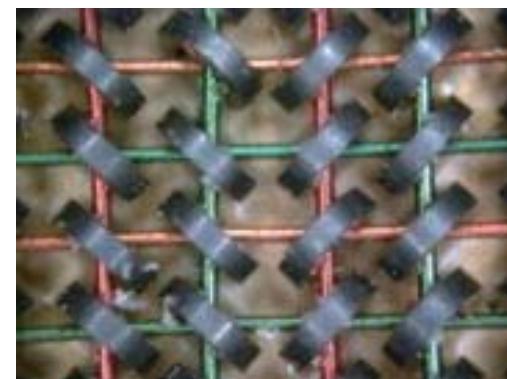
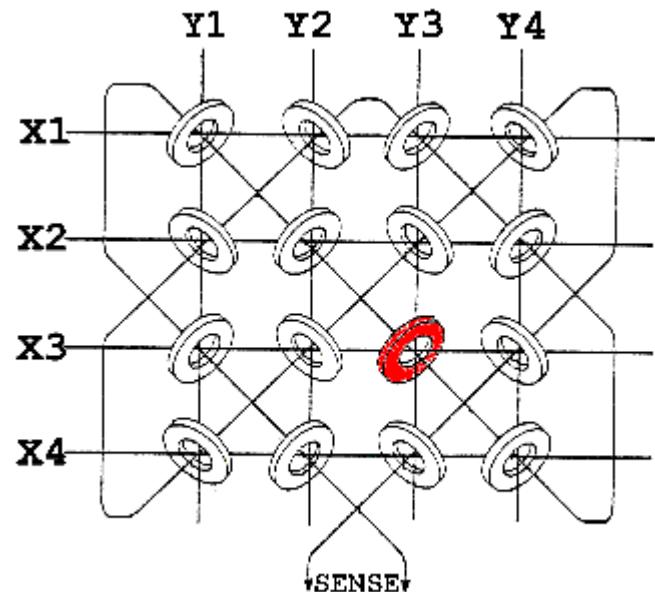


	1950	1960	1970	1980
	1949 Pulse Transfer core invented	Ferrite core memory moves into high-volume production	Intel announces the 1103 1K device	Page Mode introduced
Tubes primary source of early memory	Coincident current core invented by Forester for flight simulator		First silicon memory devices invented by IBM	Core memory used in the first space shuttle for its robustness

- Vacuum tubes were the first electronic devices used for memory. Tubes were plagued with problems like excessive heat, size, and voltage as well as being unreliable and temperamental. This set the stage for a magnetic “core memory”.



- Magnetic core memories dominated the market for more than two decades.
- The basic concept is changing the magnetic properties of a ferrite ring using current.
- It takes 3 wires through each ring to read, write, and erase a bit.





1980	1990	2000	2010
Fast Page Mode introduced at 28.5 MHz	EDO memory becomes available	DDR introduced at 200MT/s	DDR3 spec released
Intel exits DRAM business	SDRAM introduced JEDEC holds memory standards from here on out	DDR2 introduced at 400MT/s	

- With the invention of the PC in 1981, silicon-based memory demand skyrocketed. Other countries like Japan begin to manufacture low cost DRAMs. By 1985, Japan has nearly dominated the market. US lawmakers team up with silicon manufacturers to equalize the playing field. New laws against price gouging are instituted.



Why DRAM?

Min Huang(min.huang@ lecroy.com)
Do Not Distribute



- What is RAM?
- Random Access Memory.
- Why is it called Random access memory?
- Previously there were other kinds of memory that were sequentially accessed. These other kinds of memory were usually in the form of magnetic tape or drum. To my knowledge the term SAM was never coined.



- What are the 2 types of RAM?
- Static RAM and Dynamic RAM.
- Which is better?
- Well it depends on the application.

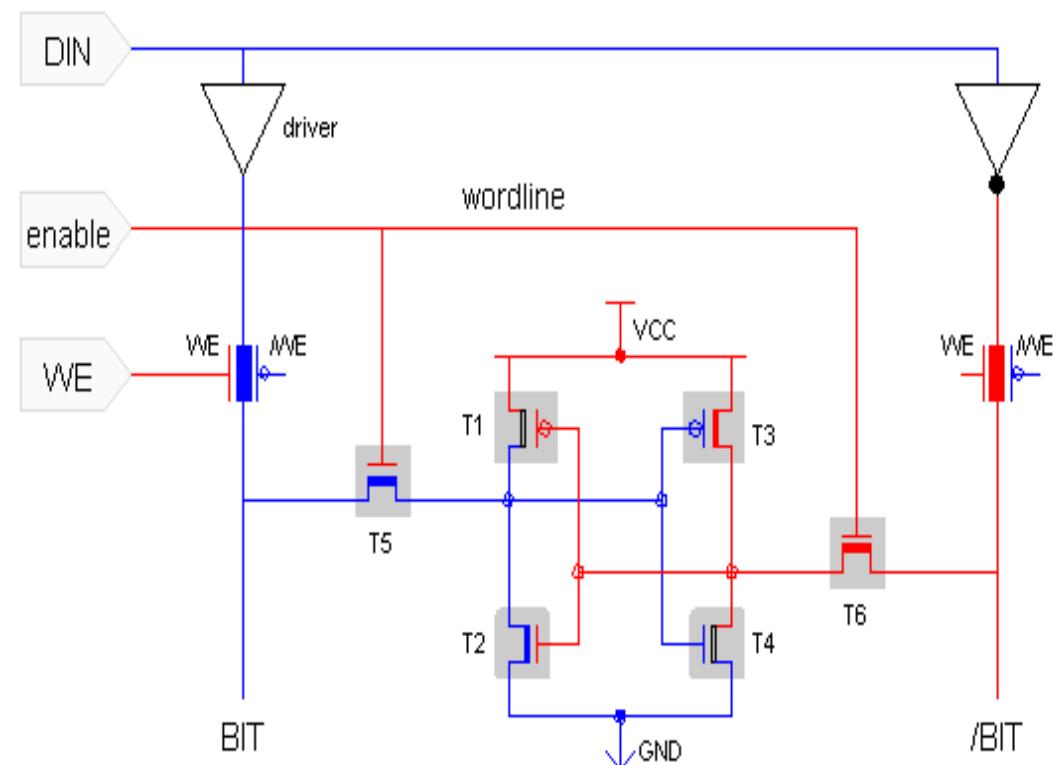
Type of RAM	Cost per cell	Size of each cell	Power Dissipated	Speed
Dynamic RAM	Low	Very small	High	Slow
Static RAM	High	Large	Low	High



➤ SRAM Cell

- A SRAM cell is composed of many transistors.
- The cell consists of two cross-coupled CMOS inverters that store one bit of information, and two N-type transistors that connect the cell to the bitlines.
- To read the information, the word line is activated while the external bit line drivers are switched off. Therefore, the inverters inside the SRAM cell drive the bitlines, whose value can be read-out by external logic.
- To write new data into the cell, the big (external) tristate drivers are activated to drive the bitlines. Next, the word line transistors are enabled. Because the external drivers are much bigger than the small transistors used in the 6T SRAM cell they easily override the previous state of the cross-coupled inverters.

CMOS: 6T-SRAM Zelle



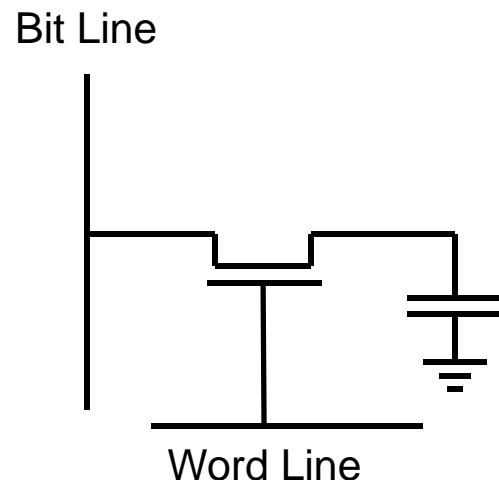


DRAM Cell Architecture

Min Huang(min.huang@ lecroy.com)
Do Not Distribute



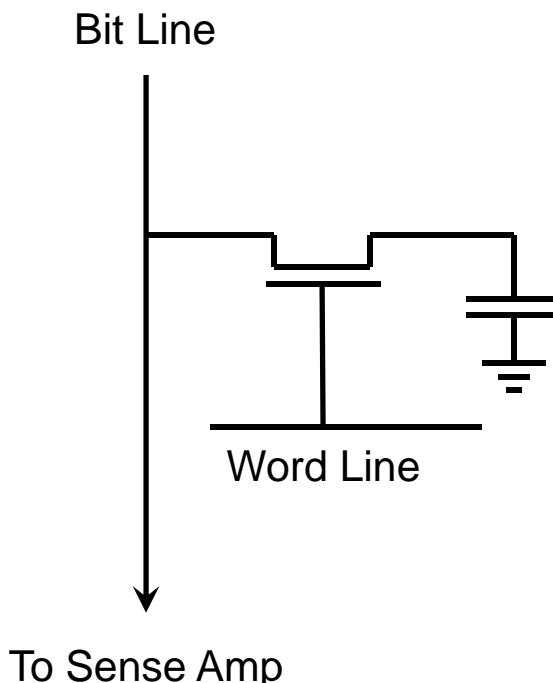
- DRAM Cell (conceptual)
 - A DRAM cell is composed of a capacitor and a transistor.
 - The data is stored in the capacitor.
 - Capacitors lose charge over time due to leakage (dissipation), so it is necessary to recharge them periodically. This is called Refresh and is controlled by the chipset.
 - The act of reading the capacitor is destructive because the charge in the cap is so small (<30fF). The data in the cap is always read out to a Sense Amp and then written back at a later time (Precharge).





➤ DRAM Cell Read

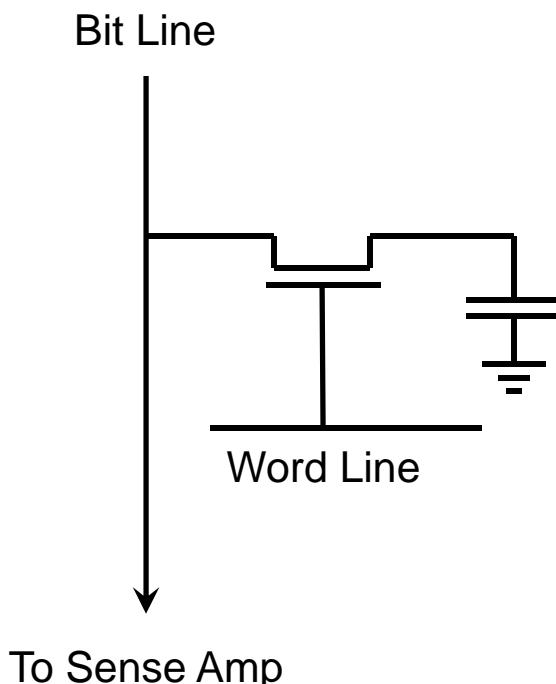
1. Bit line is precharged to Sense Amp threshold voltage ($V_{DD} / 2$).
2. Word line turns on transistor to allow charge to flow from capacitor to bit line. (Read from Cell)
3. Sense Amp latches electrical high or low.





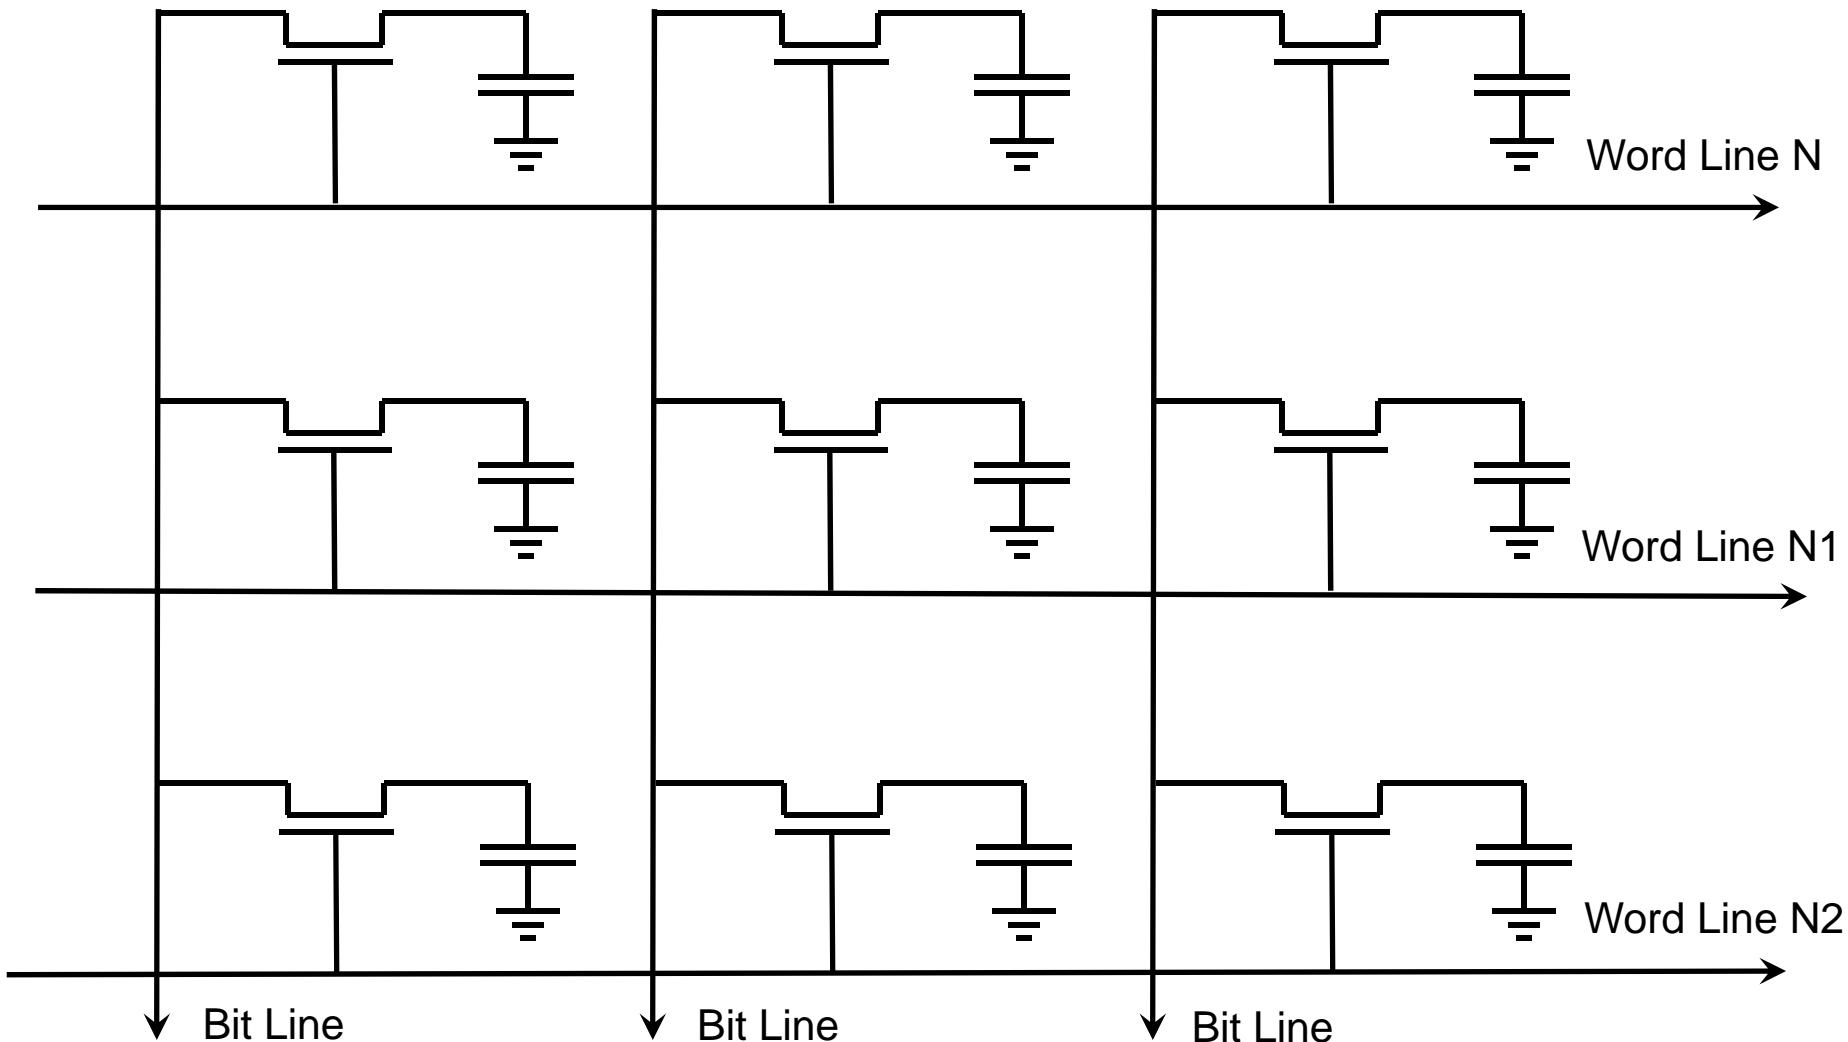
➤ DRAM Cell Write or Precharge

1. Charge or discharge the bit line to core voltage or GND as needed to store 0 or 1. DRAM manufacturer decides what state (1 or 0) is represented by a charged cell.
2. Word line turns on transistor to allow charge to flow from bit line to capacitor.
3. During writes the word line may need to be forced above V_{DD} to take care of voltage drop.



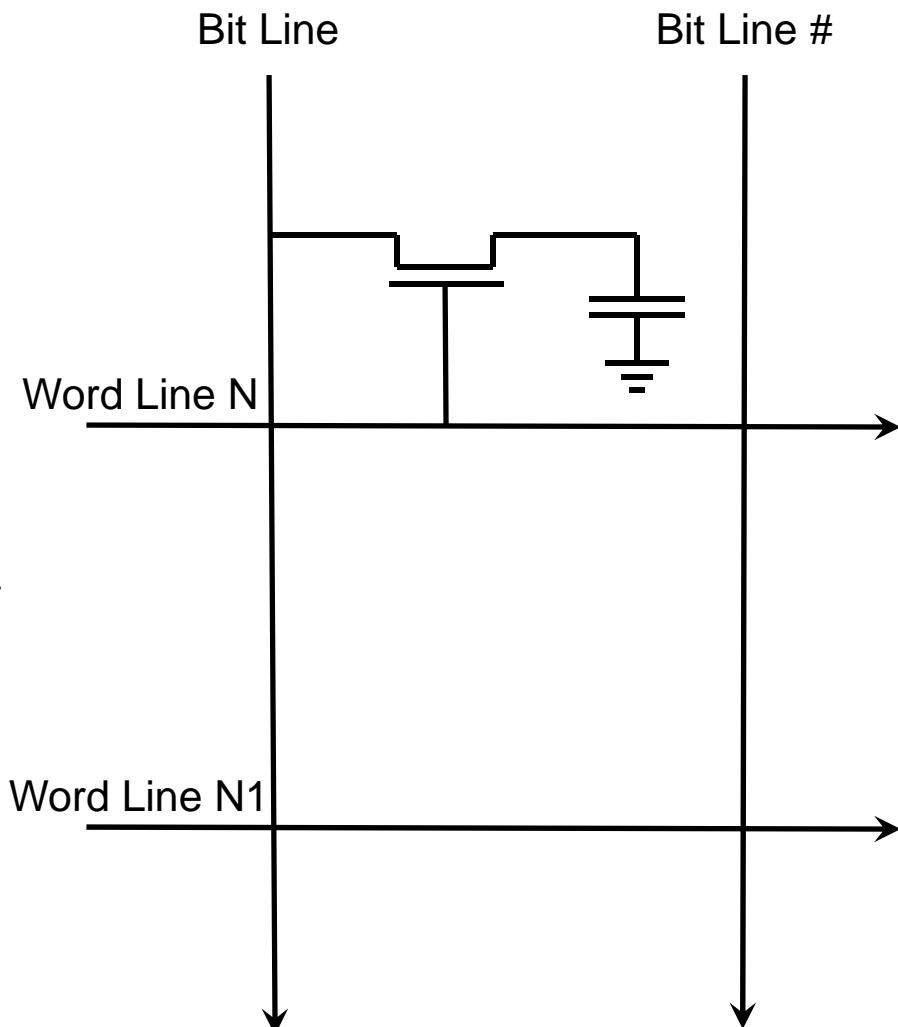


Open Cell Array Architecture



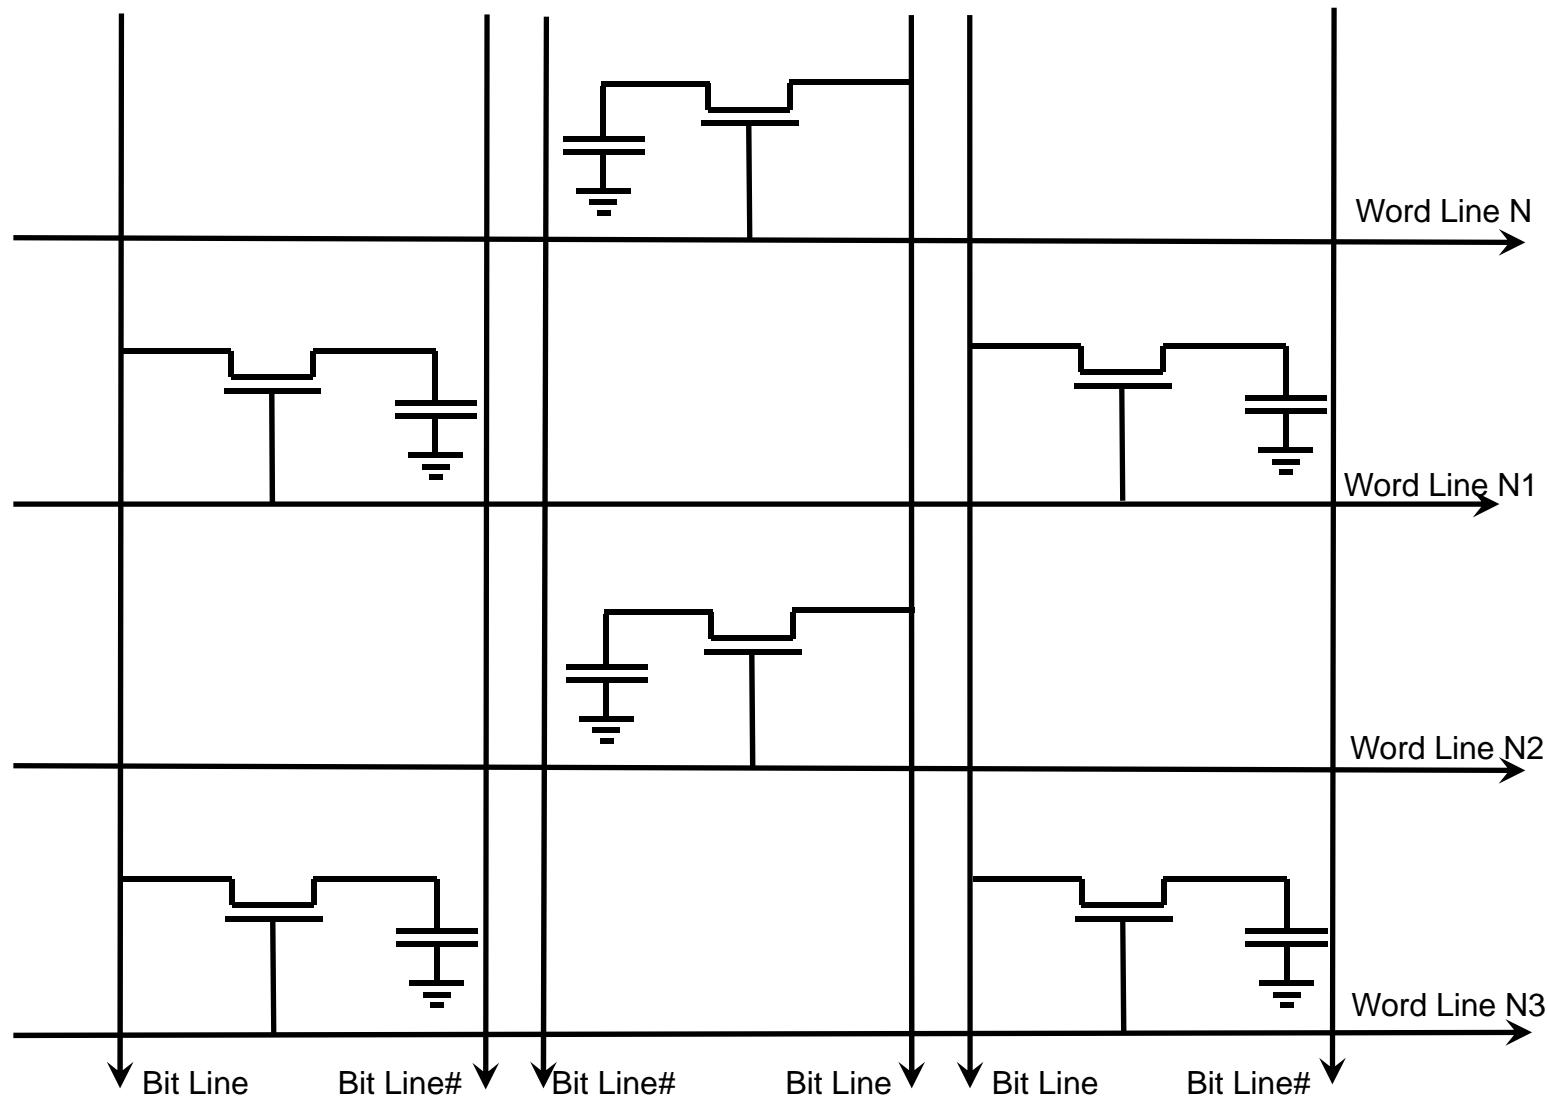


- Folded DRAM Cell (conceptual)
 - Adjacent bit lines are organized as differential inputs to the sense amplifier.
 - During a read operation one bit line acts as the reference input to the sense amp.
 - Folded bit lines have better noise immunity due to equal coupling of noise on the neighboring reference bit line.



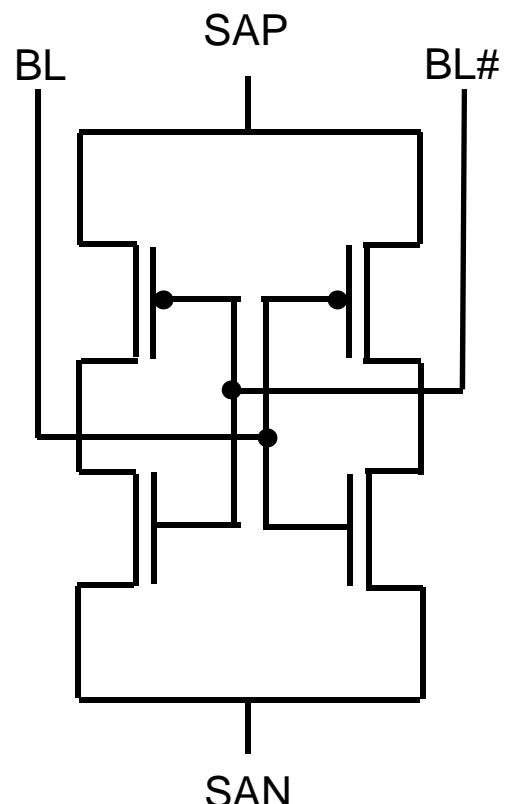


Folded Cell Array Architecture





- Sense Amplifier (conceptual)
 - The sense amplifier is a differential latch.
 - Example of reading a “0”
 - Bit lines are charged to $V_{DD}/2$, SAN (Sense Amp Negative) = $V_{DD}/2$, SAP (Sense Amp Positive) = 0. The latch is off.
 - When word line is asserted, one of the bit lines changes by delta V while the other bit line remains unchanged.
 - SAN is discharged to 0
 - As SAN falls below the higher of the bit lines, the corresponding NMOS transistor turns on and starts pulling the lower bit line to 0.
 - After a small delay, SAP is pulled high. The cross coupled PMOS transistor pulls the higher bit line voltage up to V_{DD}
 - Isolation and equalization circuits not shown





DRAM Chip Architecture

Min Huang(min.huang@ lecroy.com)
Do Not Distribute



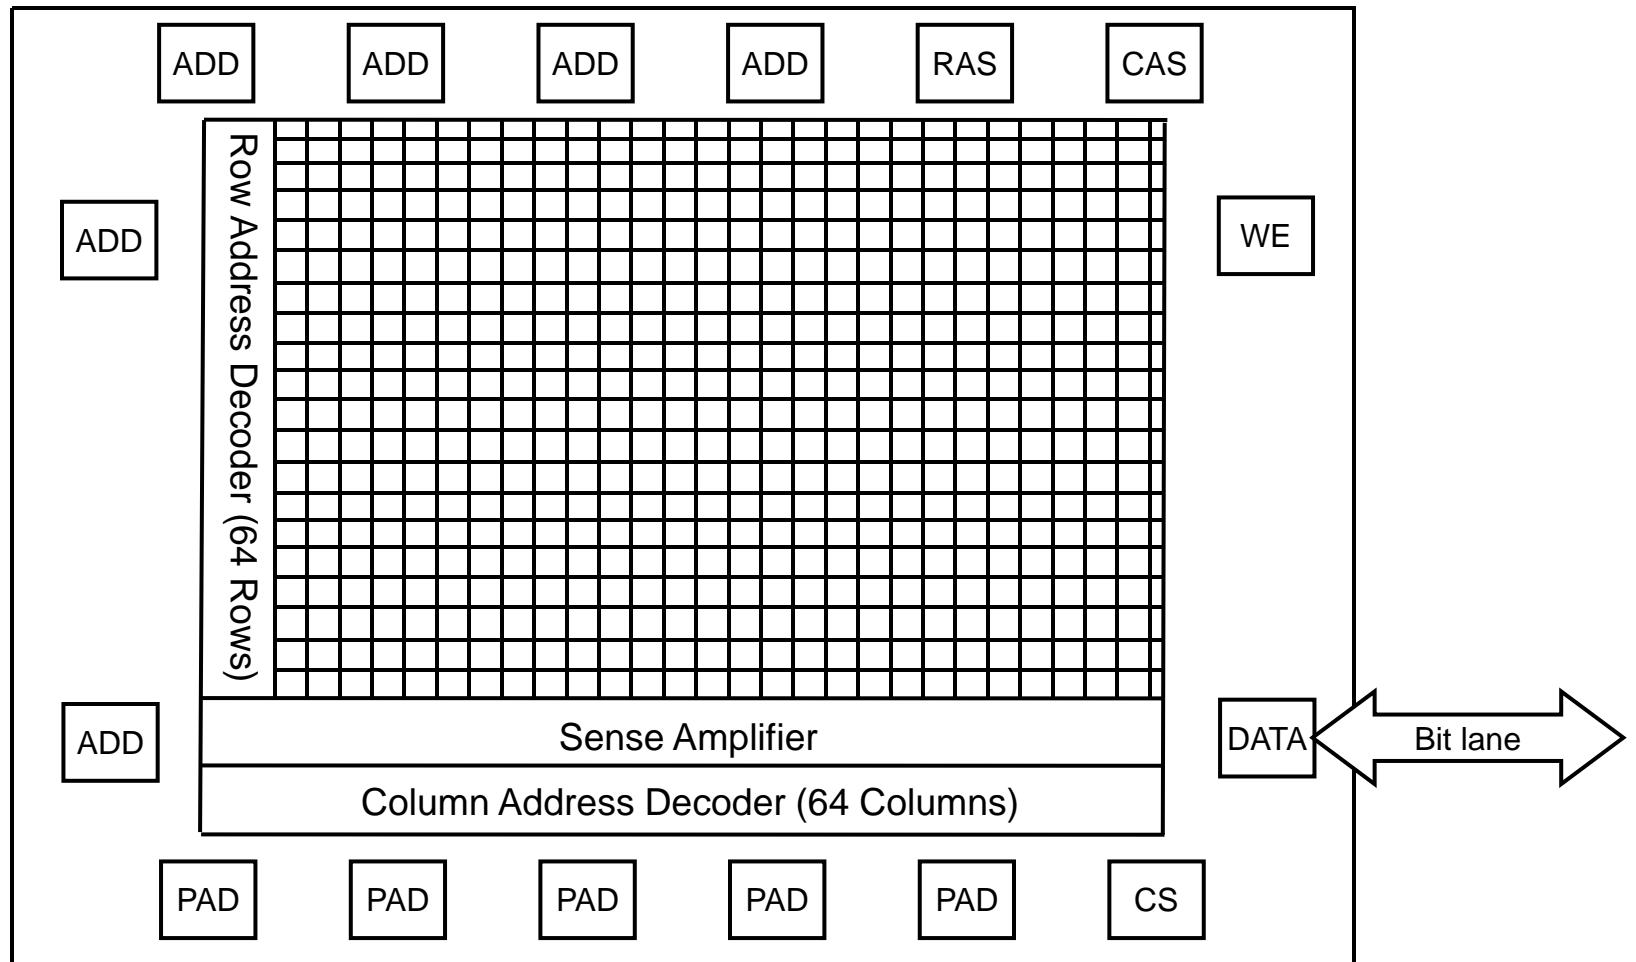
- The first LSI (Large Scale Integration) DRAMs were manufactured by Intel in 1970 starting with the 1103 chip.
- The 1103 had a 1,024-bit array with one data line. The following examples show the evolution of the DRAM.



- The core of the chip is the Array block.
- The total size of the chips is determined by multiplying the number of Rows by the number of Columns by the number of Banks.
- There may be more than one bit at each individual column address.
- This total size is referred to as the Chip Technology and is always stated in bits not bytes.

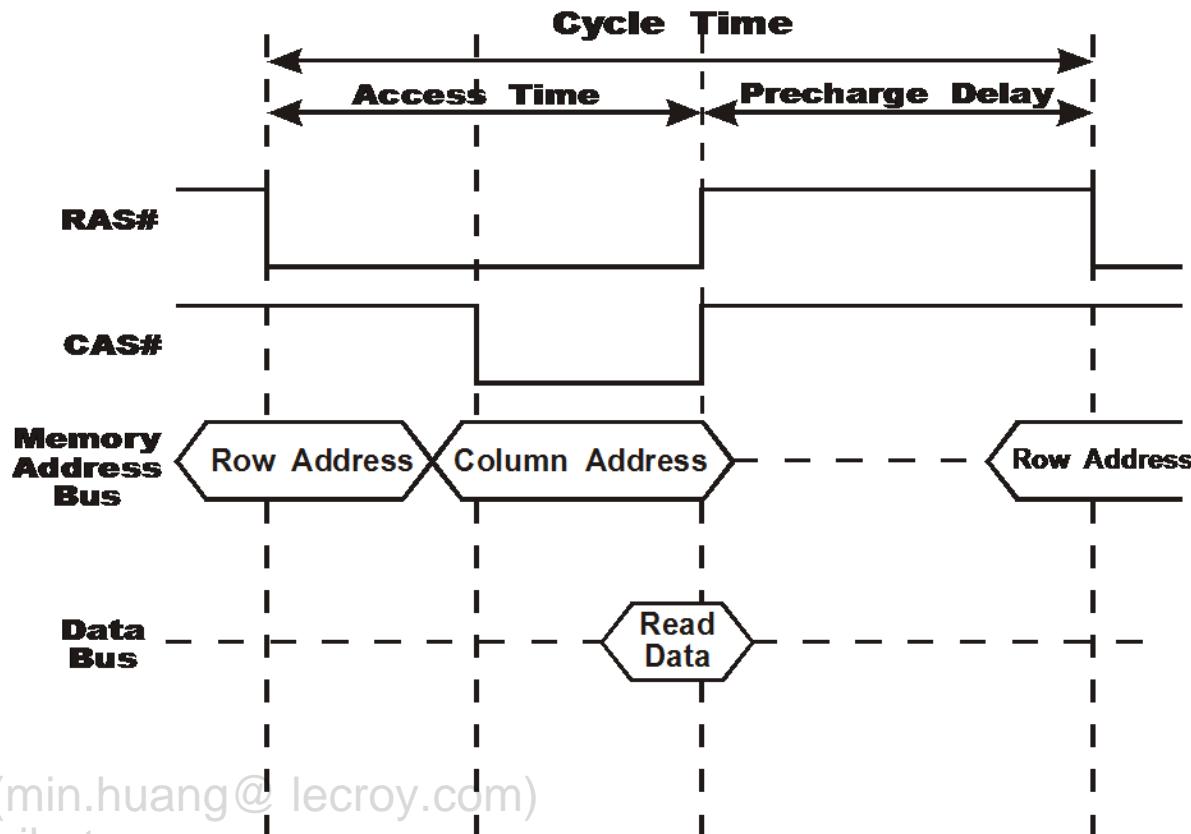


- 4096 bit Array X1 example (2^{12} Bits)



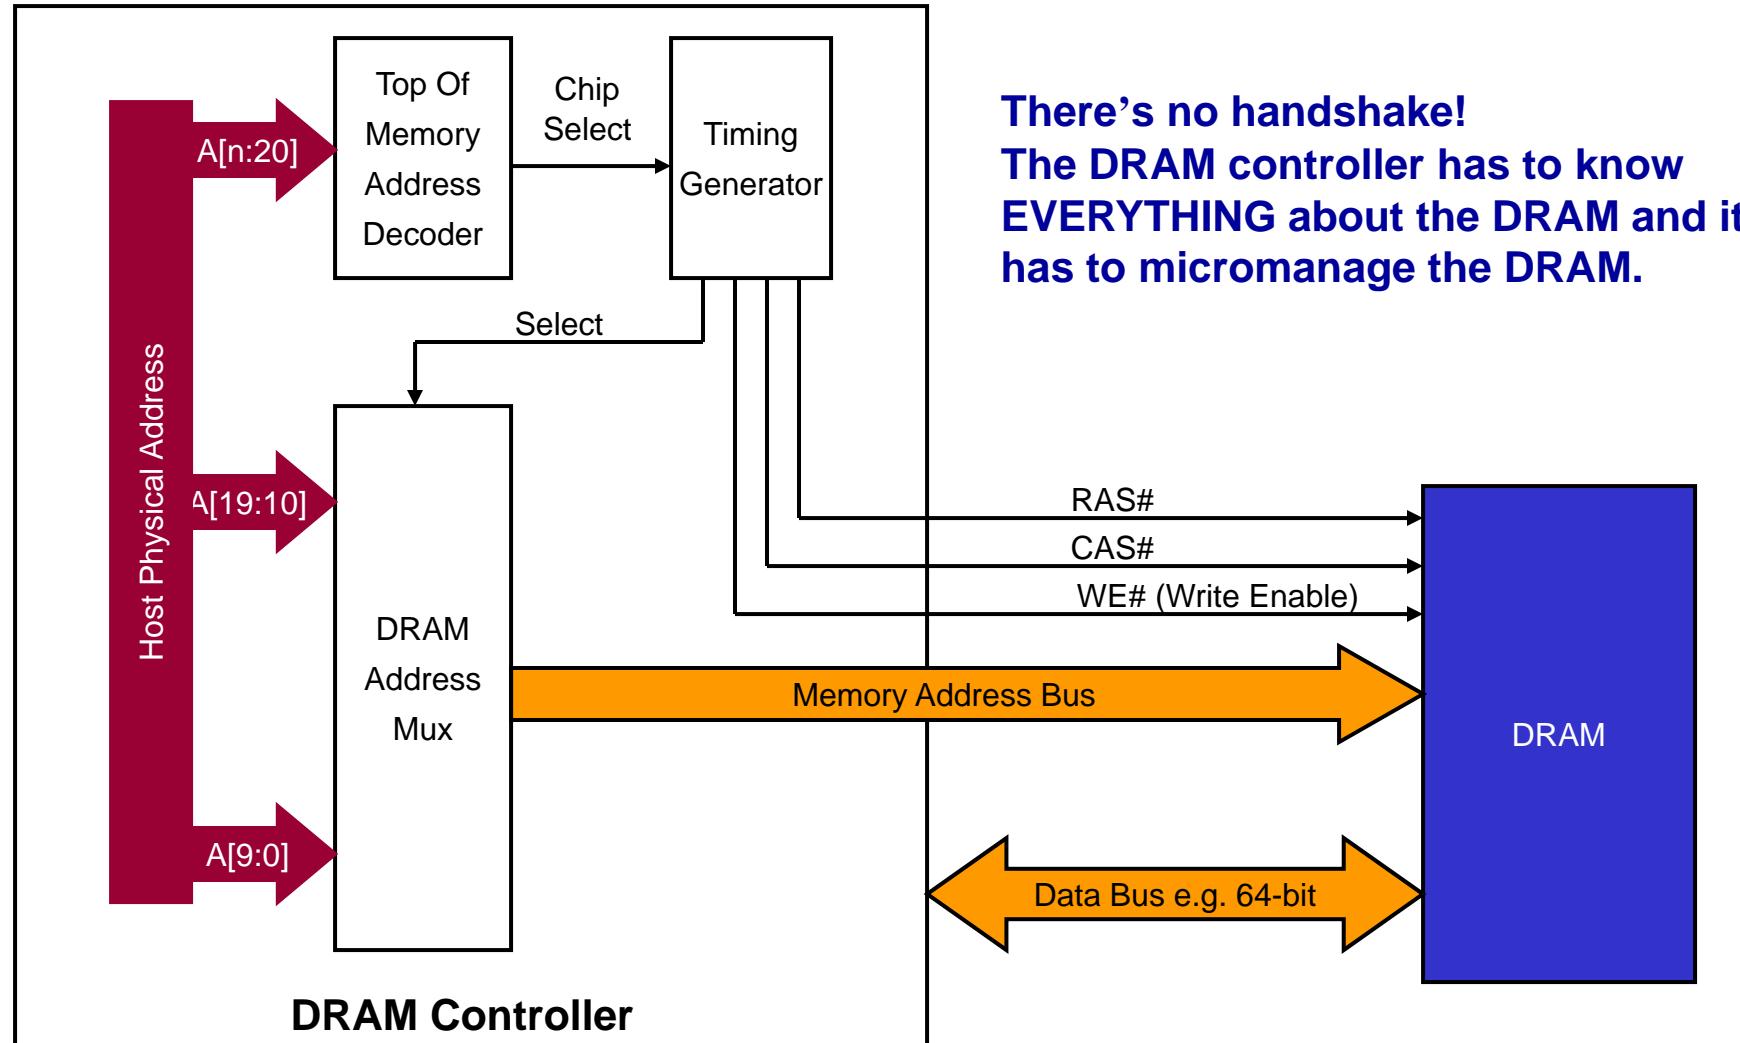


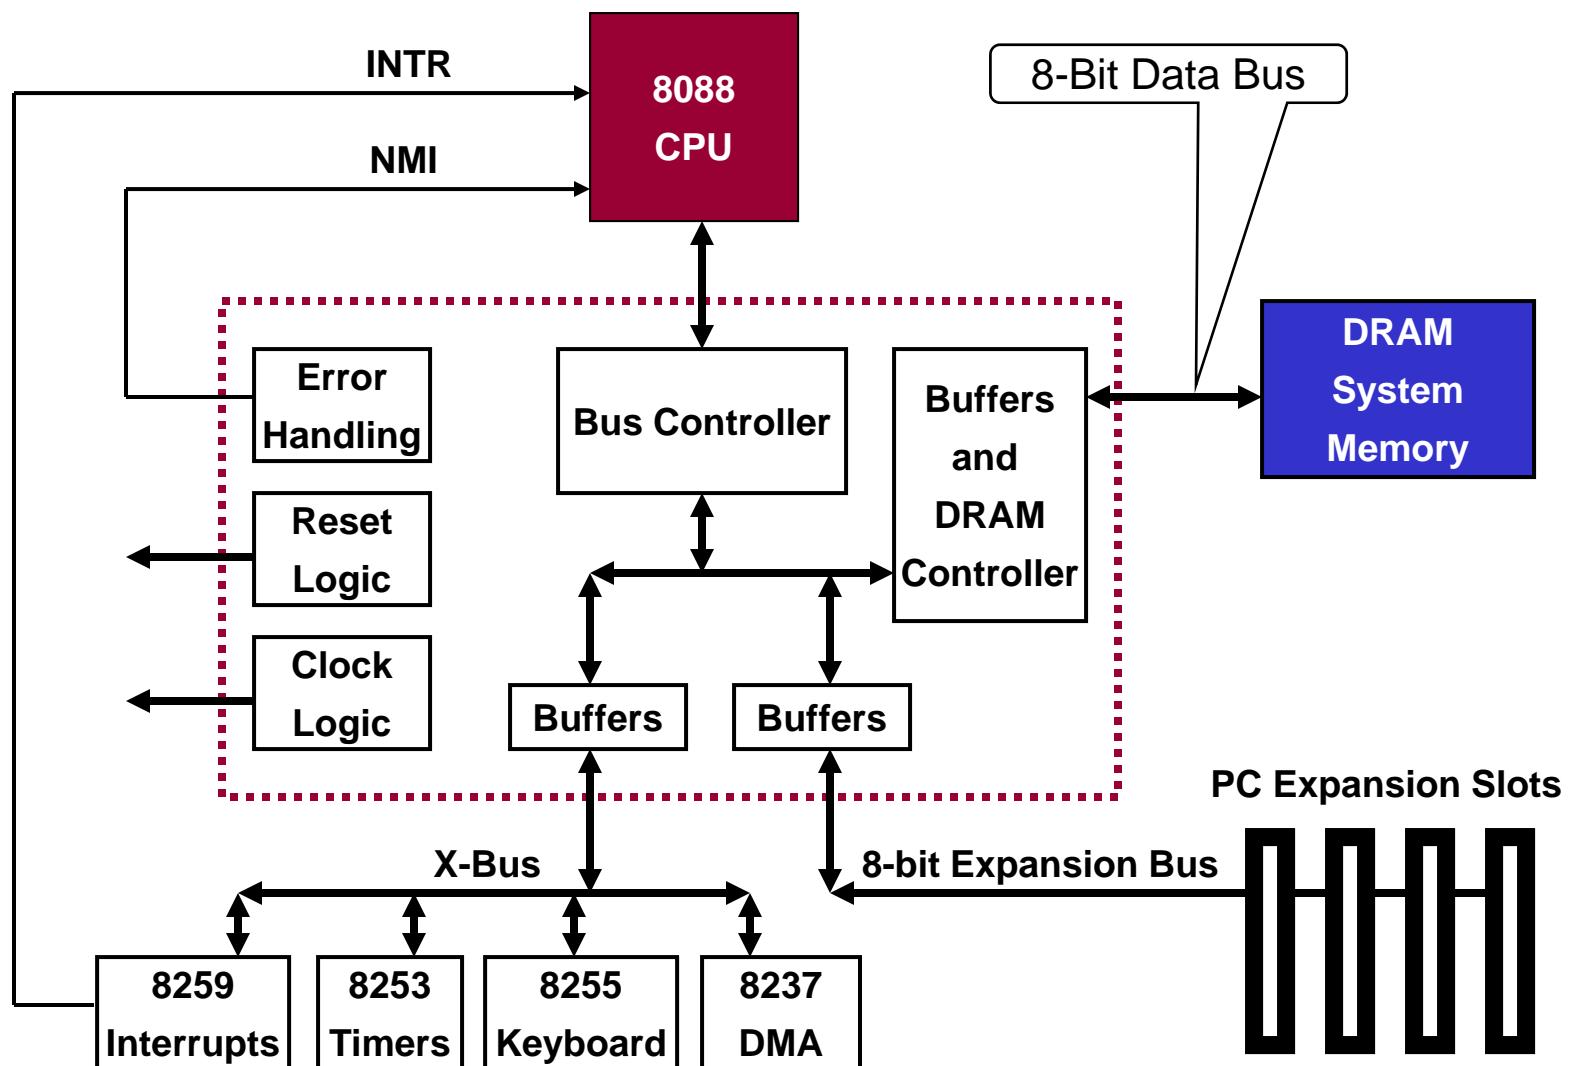
Each time a DRAM location is read, the capacitor associated with that cell is discharged. Before the device is read again, DRAM logic must recharge these locations, which takes some time—called the precharge delay.





Simple Async DRAM Controller Design



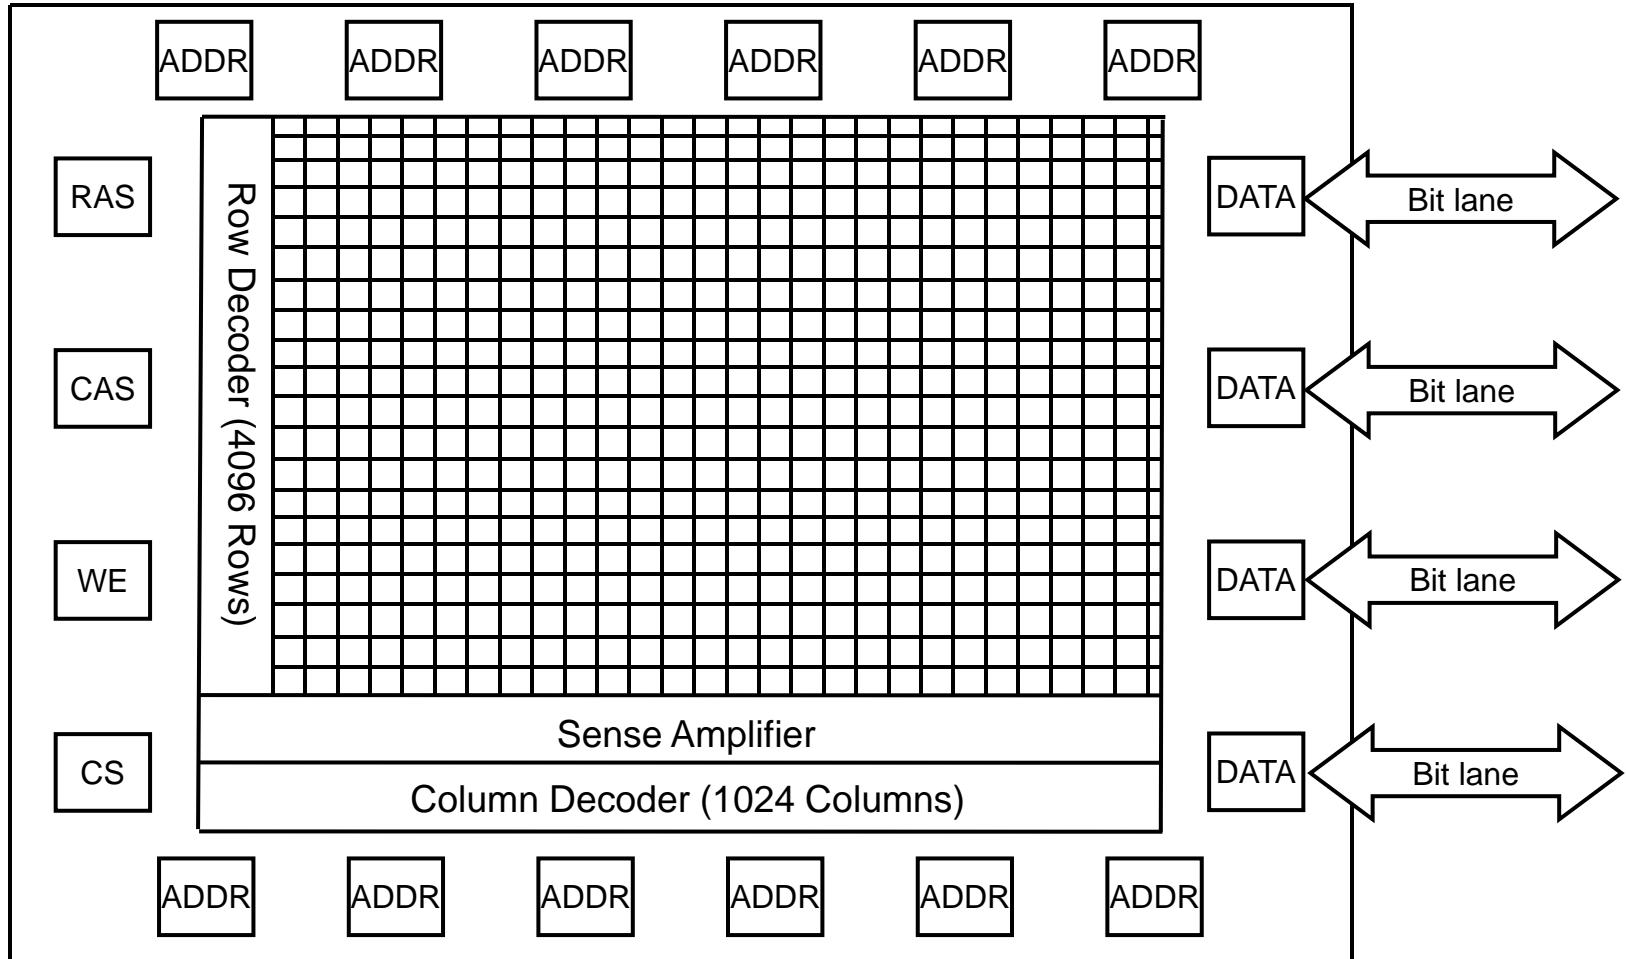




- The 8088 memory bus was 8 bits wide.
- By using 8 chips with 1 data line (X1) and a single chip select to all 8 chips, the chips would perform in unison to drive the 8-bit bus.
- As technology evolved, the DRAM controllers demanded wider buses to increase speed.
- In response, the DRAM manufacturers made the arrays bigger and added more data lines to the chips.

16M bit Array example; 4M X 4

Sometimes shown as density X data width X #banks (4M X 4 X 1)

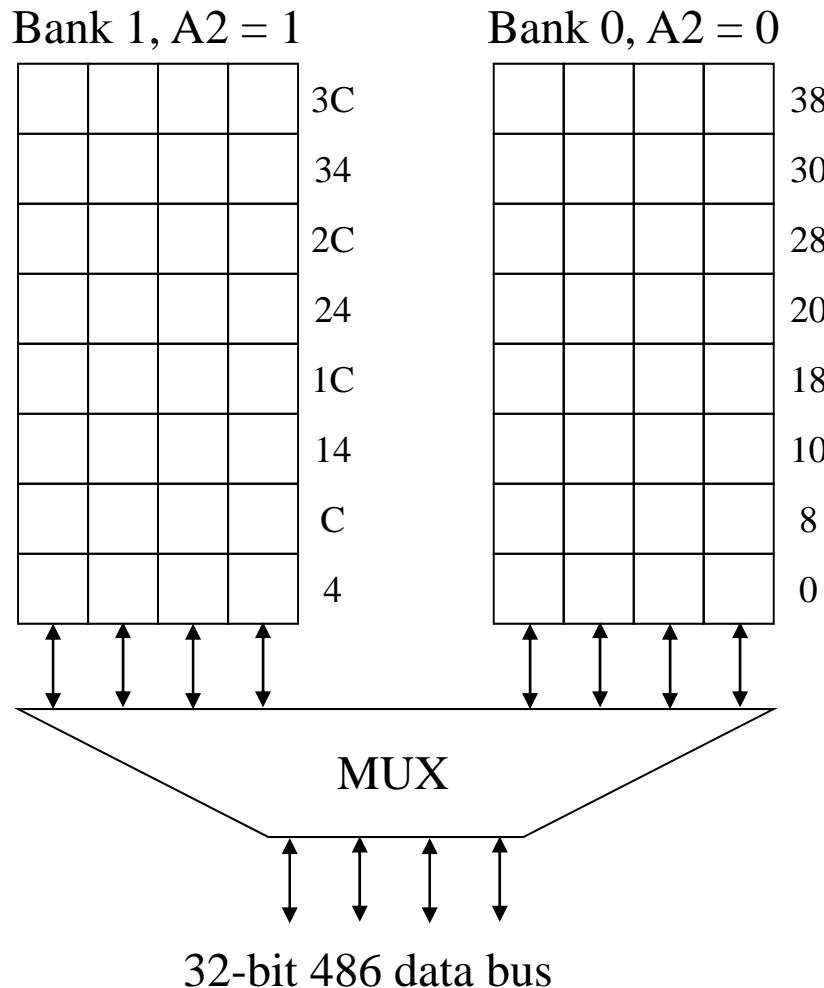




- One of the first changes in architecture to increase the speed of modern chips was to split the array into multiple banks.
- This allows more than one bank to be active at the same time.
- This can effectively “hide” precharge/activate and other timings related to single bank implementations.
- Refresh timing is not affected since all the banks are refreshed at the same time.



Banks and Toggle-Mode Addressing



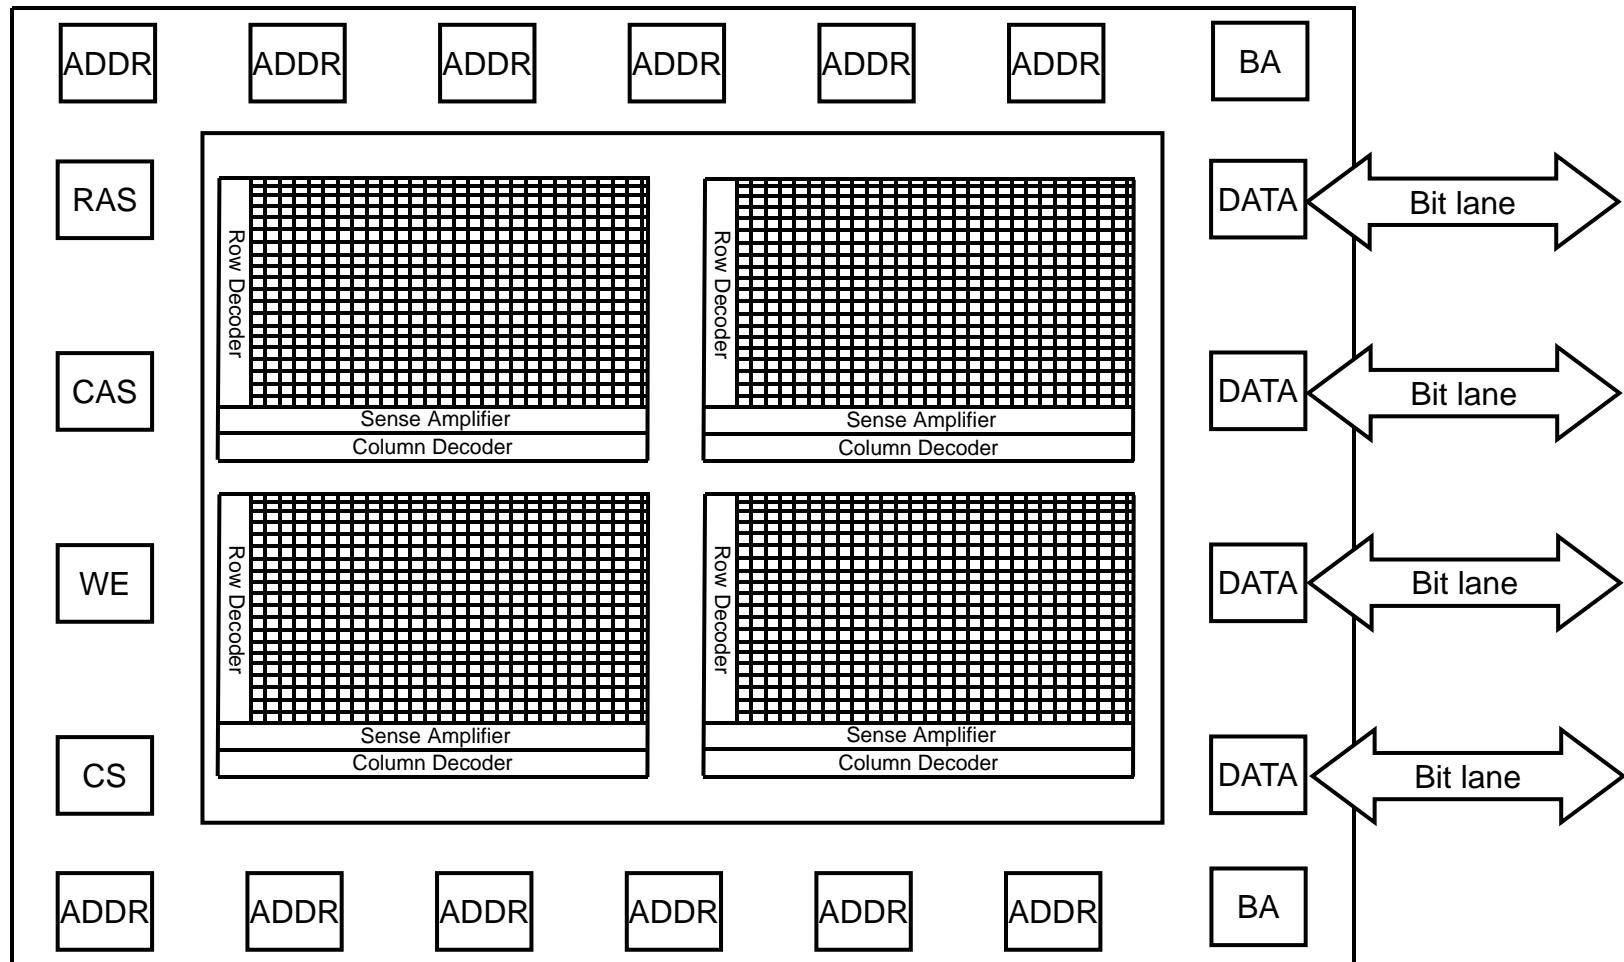
Intel 80486 has a 32-bit data bus and 16-byte cache lines.

Intel's toggle-mode burst line fill address sequence optimized DRAM performance regardless of whether DRAM was single-bank 32-bit memory (single SIMM channel), dual-bank 32-bit memory, or 64-bit memory. The sequence is shown in the Commands and Waveforms section of this slide show.



Multiple Banks

16M bit Array example; $1M \times 4 \times 4$
Shown as density X data width X #banks



- Chips can be organized into data bus widths of 4, 8, 16, and 32.
- This was shown in the previous example as number of bits X (by) number of data lines X (by) number of banks. The number of bits is referred to as the density.
- Here is an exercise of how the same technology can yield different chip organizations.

For this example,
use a 512 Mb
technology chip.

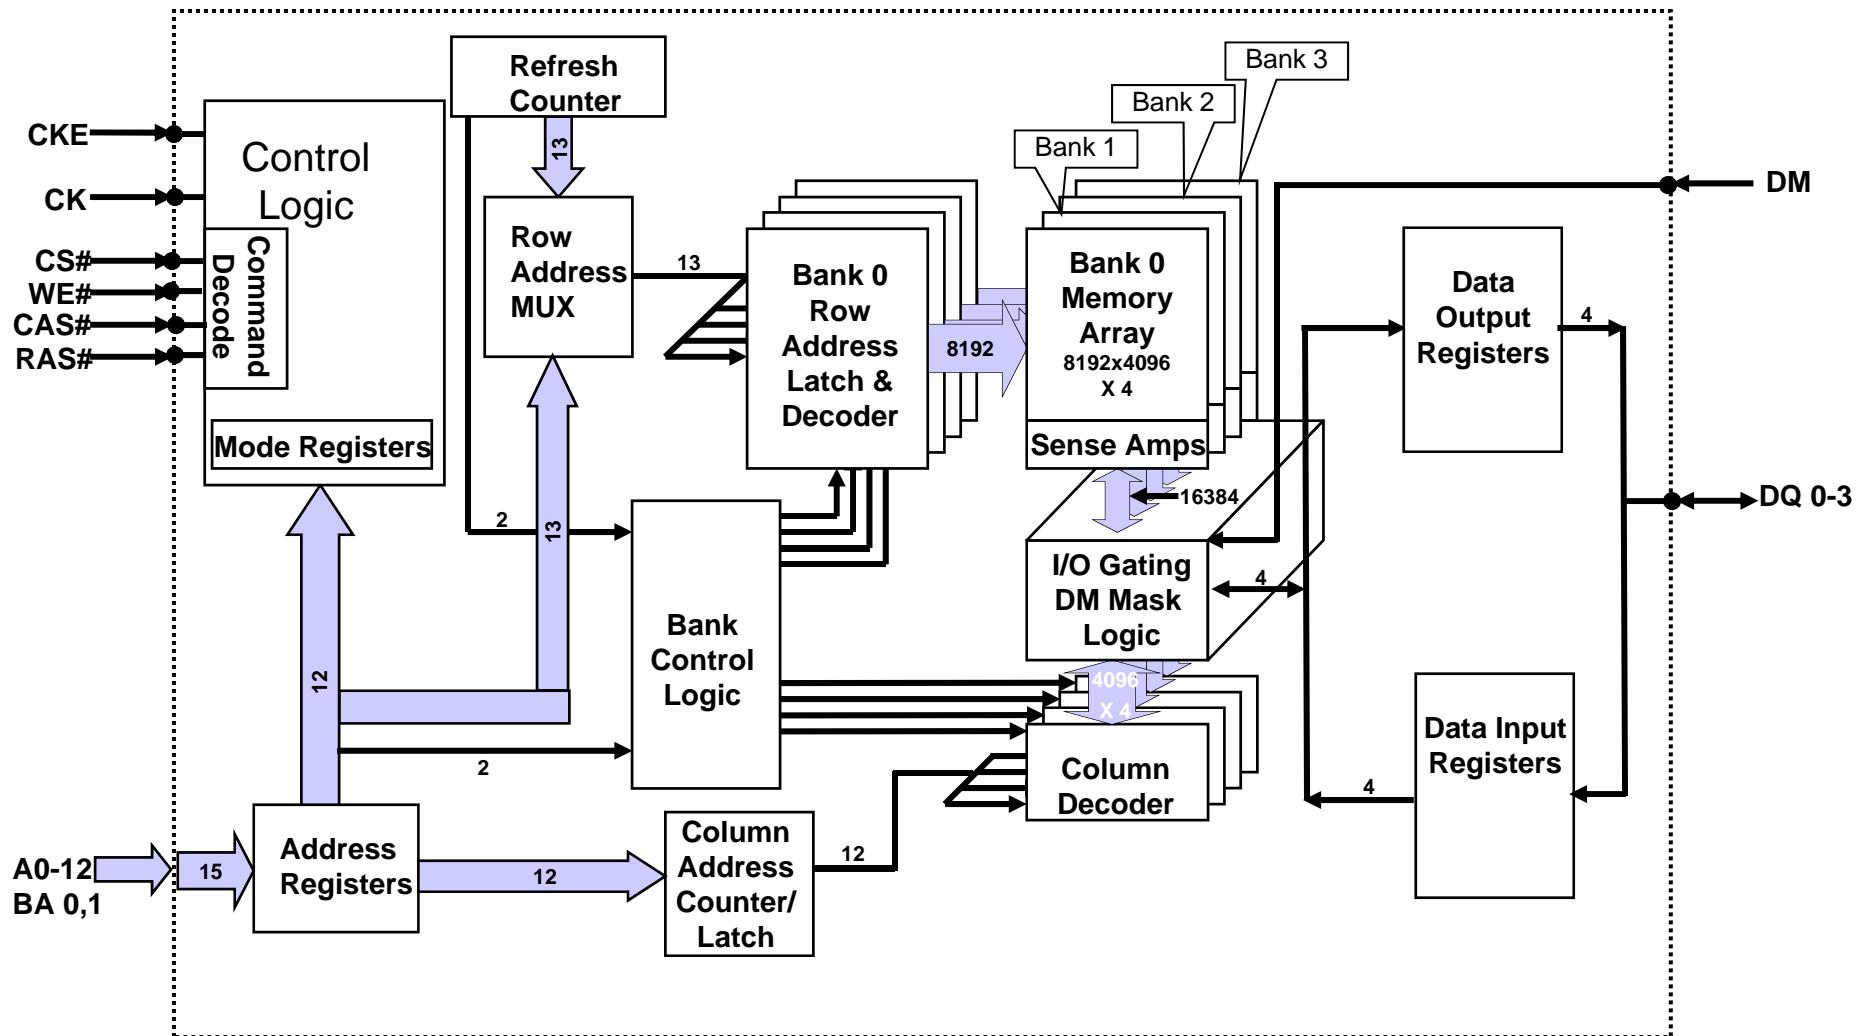
Density	Data Lines	Banks
32M	4	4
16M	8	4
8M	16	4
4M	32	4



Features of an SDRAM Chip

- PC 100, PC 133
- Fully Synchronous
- Internally pipelined column address can be changed every clock. (1T timing)
- Multiple internal banks for hiding row access/precharge timing. Single bank precharge or all bank precharge.
- Programmable burst length 1, 2, 4, 8, or full page (Row)
- Self refresh and low power states
- LVTTL 3.3 volt operation

Architecture of a 512 Mb SDRAM chip organized as 32M X 4 X 4





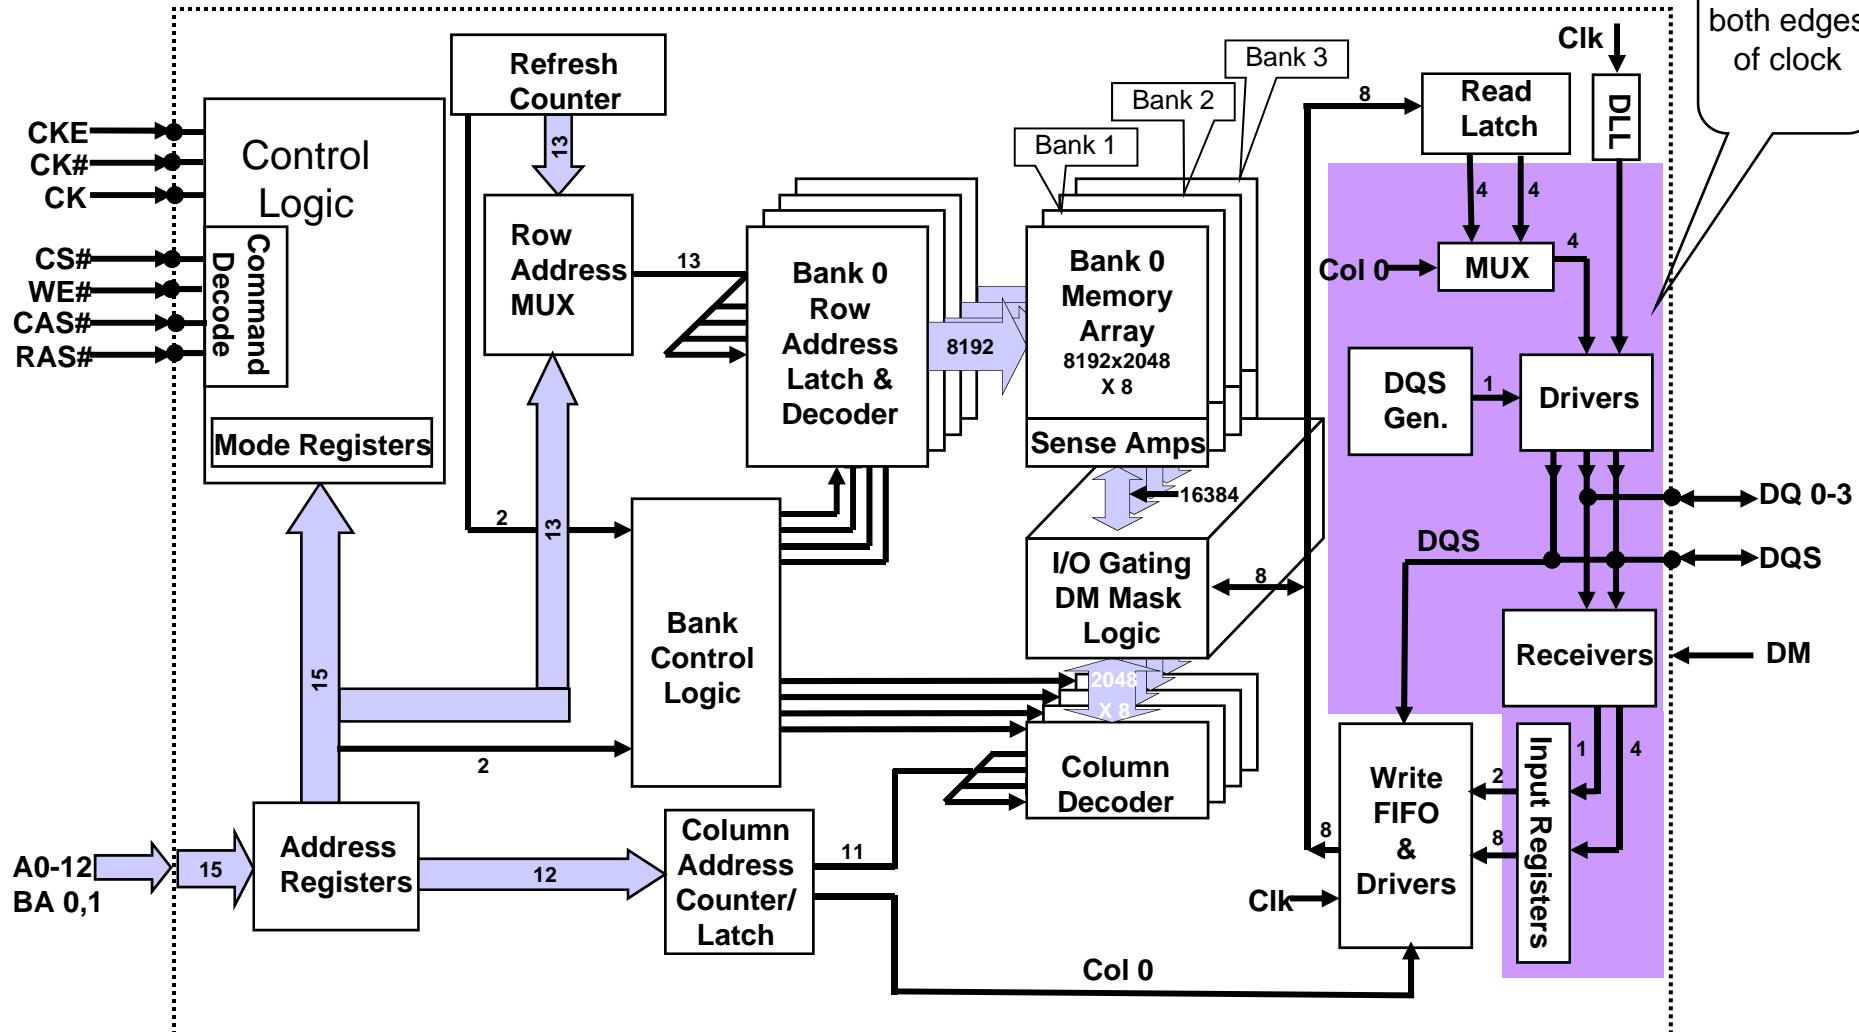
Features of a DDR1 Chip

- DDR 200, 266, 333, and 400
- Bidirectional Data Strobes for source synchronous capture (x16 has 2)
- DDR architecture two data accesses per clock
- 2n prefetch architecture
- Multiple internal banks for hiding row access/precharge timing.
- Differential clock inputs (CK and CK#)
- Commands active on rising edge of clock.
- DQS edge aligned for data reads and center aligned for writes.
- Addition of DLL (Delay Lock Loop) to align DQS and Data
- Programmable burst length of 2, 4, and 8
- Auto refresh and Self refresh modes
 - “Auto refresh” is called “Refresh” for DDR3.
- 2.5 volt operation for I/O (SSTL_2 compatible) 2.4 volts for 400MT/s DRAMs

DDR1 SDRAM Architecture

56

Architecture of a 512 Mb DDR1 chip organized as 32M X 4 X 4



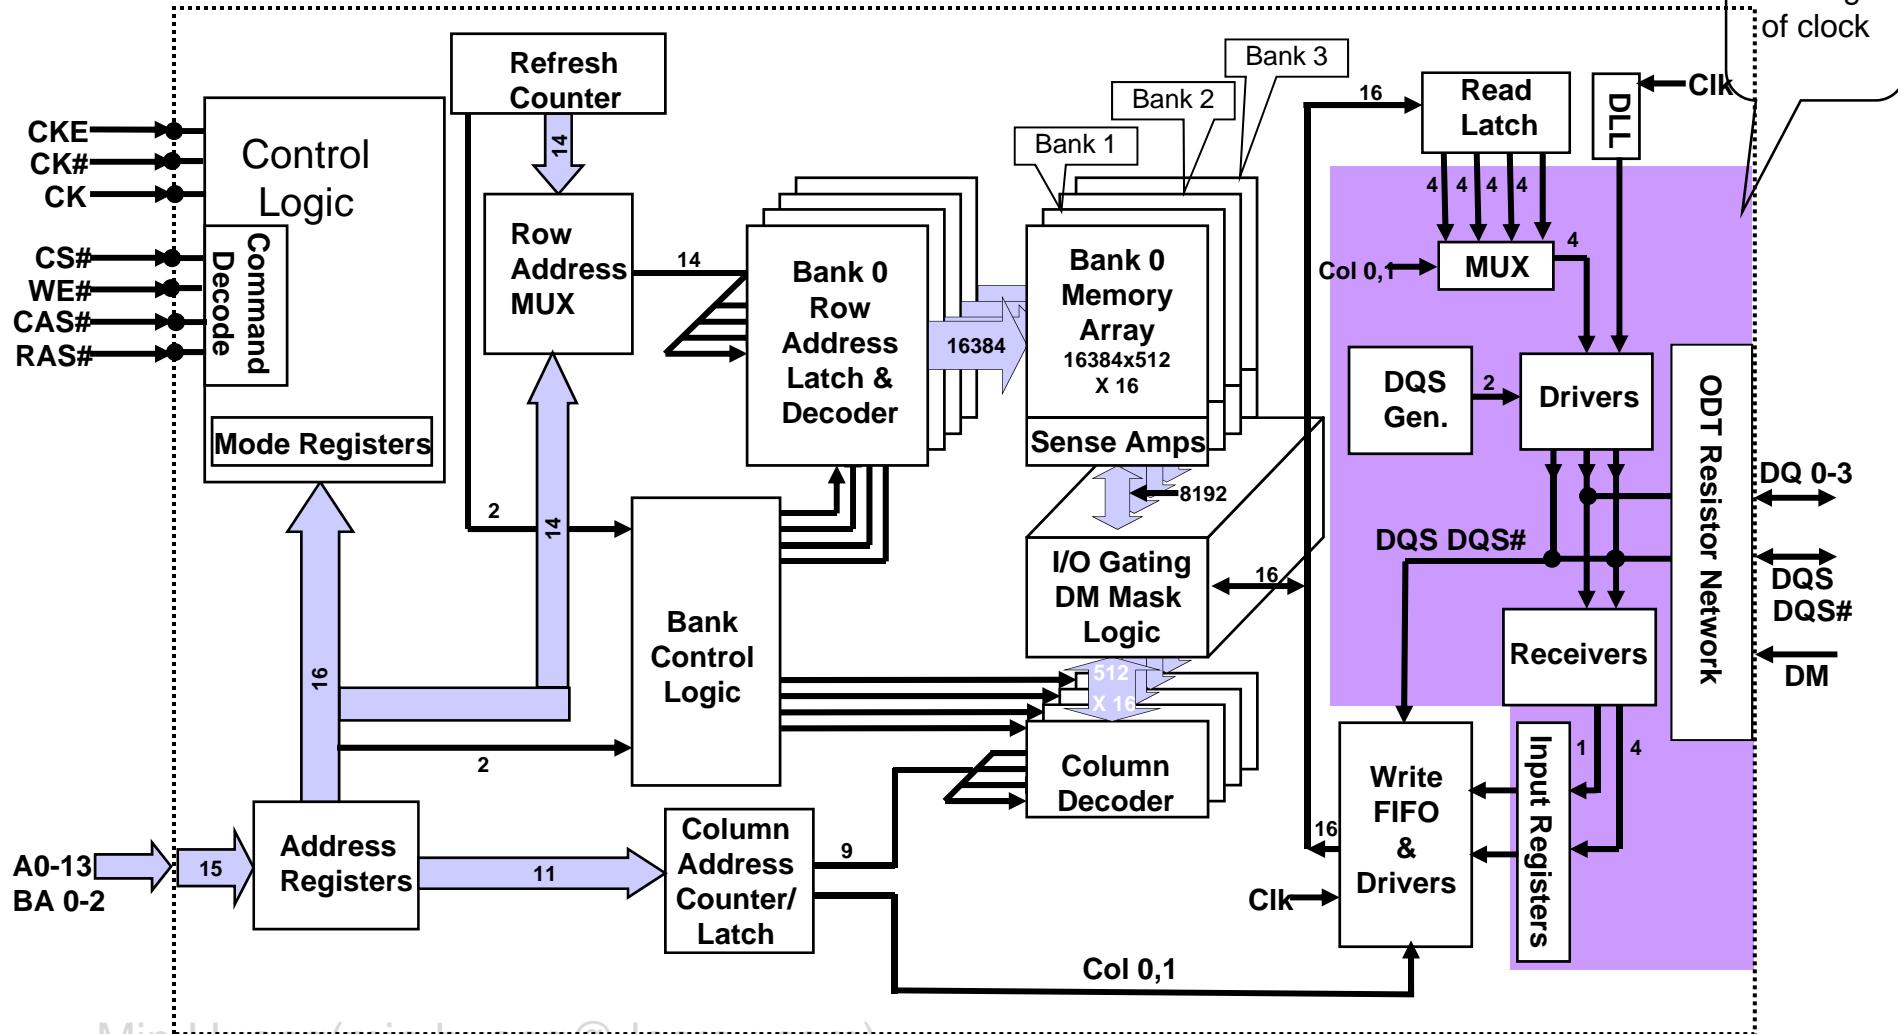


Features of a DDR2 Chip

- DDR 400, 533, 667 and 800
- Bidirectional Differential Data Strobes (DQS, DQS#) for source synchronous capture (x16 has 2 pairs)
- Duplicate output strobes (RDQS) for x8 devices
- 4n prefetch architecture
- 4 Banks for concurrent operation 1Gb devices have 8 Banks
- Differential clock inputs (CK and CK#)
- Programmable CAS latency
- Posted CAS additive latency
- On Die Termination (ODT)
- Programmable burst length of 4 or 8
- Auto refresh and Self refresh modes
“Auto refresh” is called “Refresh” for DDR3.
- 1.8 volt operation for I/O (SSTL_18 compatible)

DDR2 SDRAM Architecture

Architecture of a 512 Mb DDR2 chip organized as 32M X 4 X 4





Features of a DDR3 Chip

- 8n prefetch architecture
- Fly-by routing which requires read delay calibration and write leveling.
- On-die termination (ODT) improvements
- ZQ calibration (OCD) redefined
- Speeds up to 1600MT/s (800 MHz)
- Burst Chopping
- MRS definition changed
- V_{DDQ} goes to 1.5 volts and V_{REF} and V_{TT} are $\frac{1}{2} V_{DDQ}$
- Addition of asynchronous reset pin
- Package changes to accommodate pin mirroring and support balls



Core Speed vs. I/O Speed

Memory Technology	Input Clock	Prefetch Width	Bus Speed	DRAM Core Speed
SDRAM-100	100MHz	1	100MT/s	100MHz
DDR-200	100MHz	2	200MT/s	100MHz
DDR-400	200MHz	2	400MT/s	200MHz
DDR2-400	200MHz	4	400MT/s	100MHz
DDR2-800	400MHz	4	800MT/s	200MHz
DDR3-800	400MHz	8	800MT/s	100MHz
DDR3-1600	800MHz	8	1600MT/s	200MHz



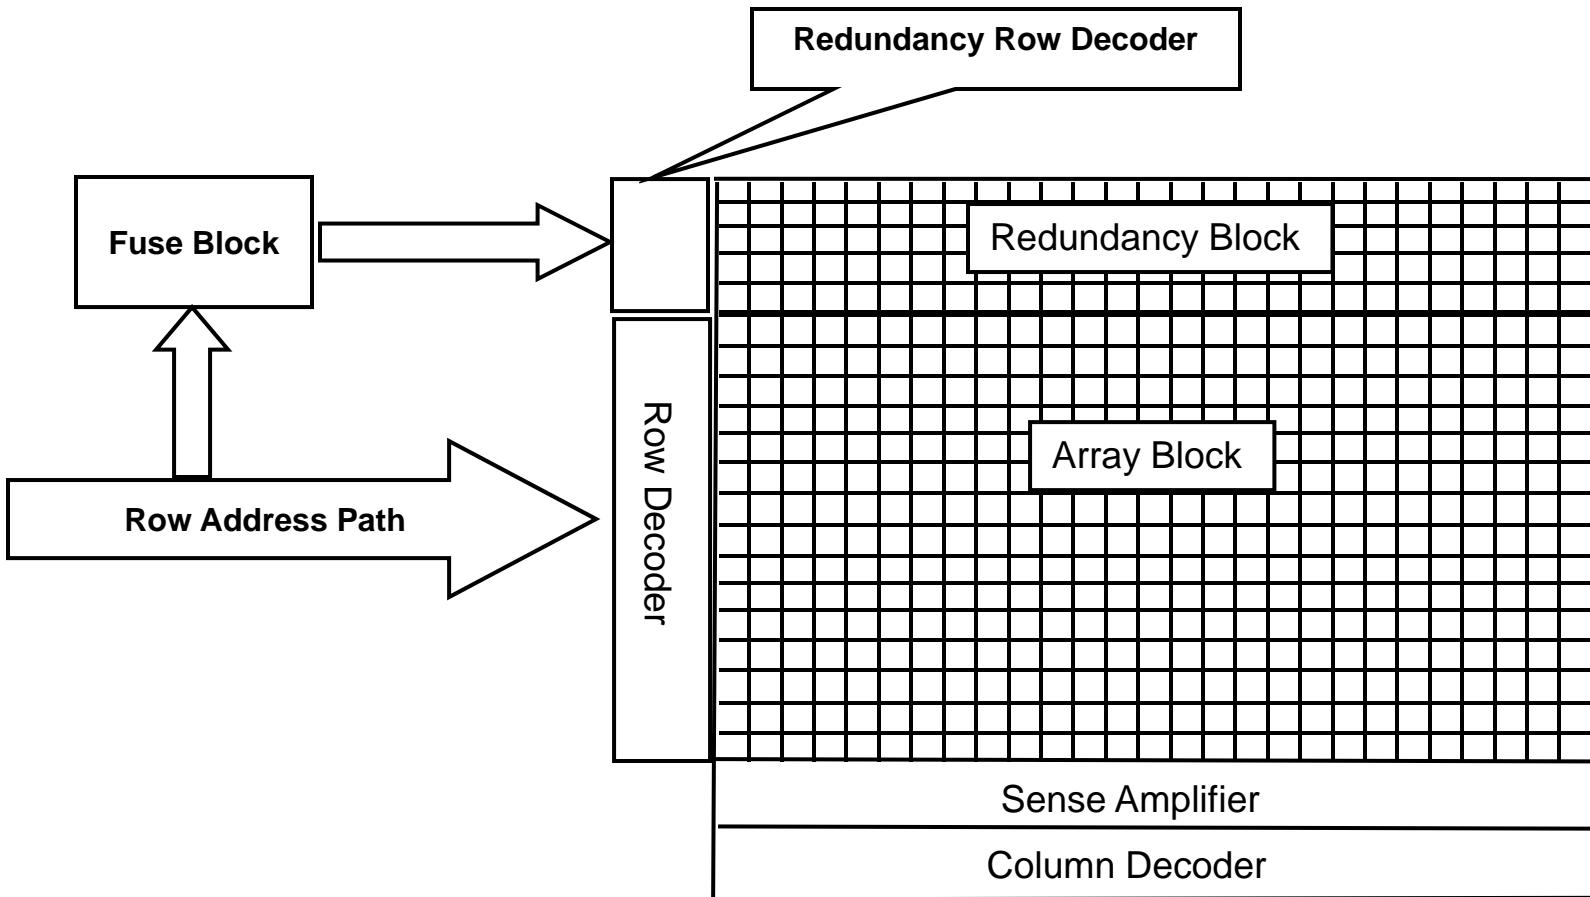
- Many fabrication processes have evolved over the last 35 years for DRAMs.
- One of the biggest challenges are creating the capacitors. Some of these basic concepts are:
 - Planar
 - Deep Trench



- As with all silicon products flaws occur during the manufacturing process.
- To keep costs low DRAM manufacturers build redundant Rows into the array to increase yields.
- When the Sort/Test process is done “bad” Rows are identified and a new row is assigned in the redundant region of the array.



Redundancy



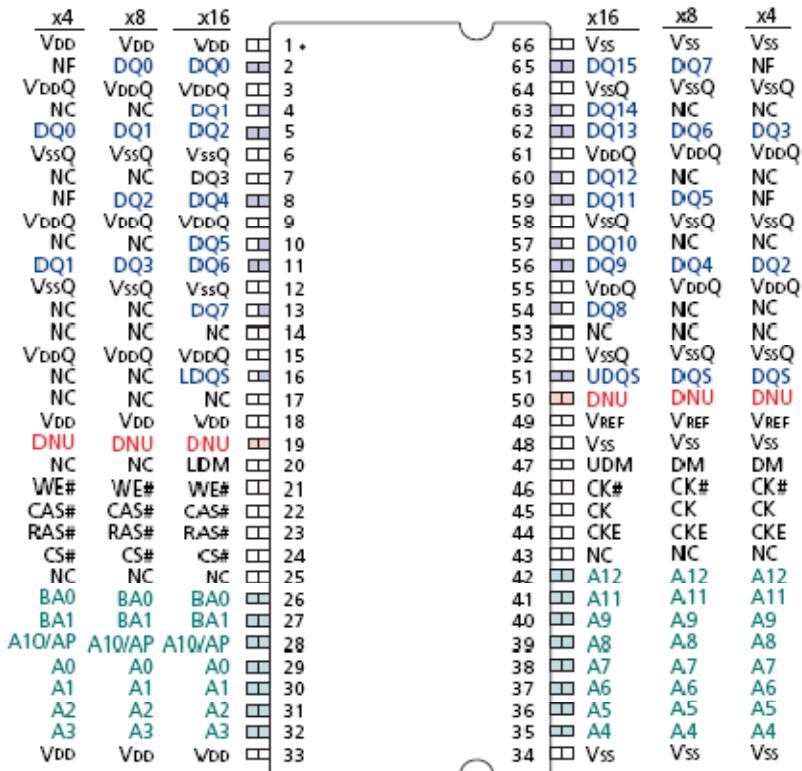


- After silicon manufacturing the die are ready for packaging.
- Package type is going to affect electrical performance of the chip. TSOPs (Thin Small Outline Package) have been used up until the higher speed DDR required a change to BGA (Ball Grid Array)



TSOPs Still in Production

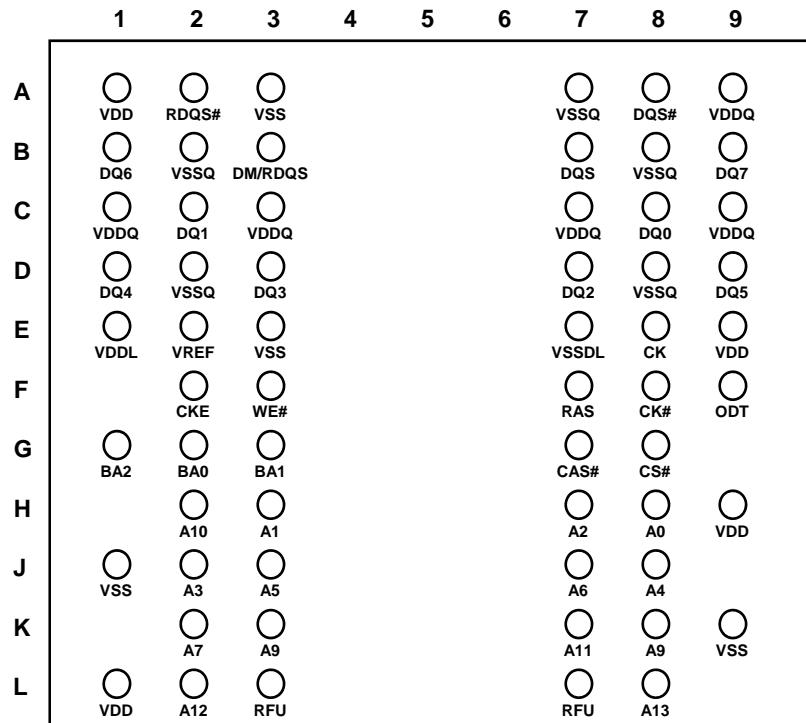
- This 66-pin TSOP is used in DDR1 for x4, x8, and x16.
- With DDR1, it was not necessary to go to a BGA package for electrical performance issues.
- Many manufacturers did use BGA packages for DDR1 in order to build high density modules.



Top View



- This is a 60-pin package used for x4 and x8 DDR2.
- JEDEC wrote in the DDR2 spec that TSOP packages were not to be used.
- This package is JEDEC standard and is referred to as MO-207.
- Ball-out depends on die organization.



Looking through the package



BGA Packages in Production

- This is the 84-pin version of the MO-207 package used for x16 DDR2 devices.
- A DDR3 BGA package is shown in the On-DIMM Mirroring section of this presentation.

	1	2	3	4	5	6	7	8	9
A	○ VDD	○ NC	○ VSS				○ VSSQ	○ UDSQ#	○ VDDQ
B	○ DQ14	○ VSSQ	○ UDM				○ UDQS	○ VSSQ	○ DQ15
C	○ VDDQ	○ DQ9	○ VDDQ				○ VDDQ	○ DQ8	○ VDDQ
D	○ DQ12	○ VSSQ	○ DQ11				○ DQ10	○ VSSQ	○ DQ13
E	○ VDD	○ NC	○ VSS				○ VSSQ	○ LDQS#	○ VDDQ
F	○ DQ6	○ VSSQ	○ LDM				○ LDQS	○ VSSQ	○ DQ7
G	○ VDDQ	○ DQ1	○ VDDQ				○ VDDQ	○ DQ0	○ VDDQ
H	○ DQ4	○ VSSQ	○ DQ3				○ DQ2	○ VSSQ	○ DQ5
J	○ VDDL	○ VREF	○ VSS				○ VSSDL	○ CK	○ VDD
K		○ CKE	○ WE#				○ RAS	○ CK#	○ ODT
L	○ BA2	○ BA0	○ BA1				○ CAS#	○ CS#	
M		○ A10	○ A1				○ A2	○ A0	○ VDD
N	○ VSS	○ A3	○ A5				○ A6	○ A4	
P		○ A7	○ A9				○ A11	○ A9	○ VSS
R	○ VDD	○ A12	○ RFU				○ RFU	○ A13	

Looking Through from Top



DRAM Modules

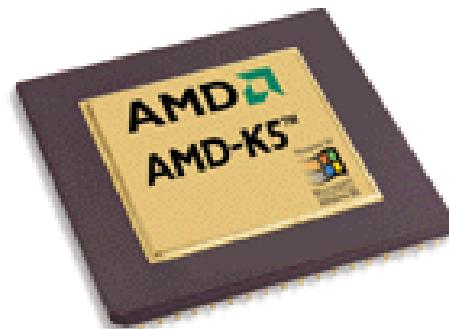
Min Huang(min.huang@ lecroy.com)
Do Not Distribute



- In the 486 and earlier processors, as the memory data bus width grew, the need to integrate monolithic DRAMs onto a module grew with it.
- This started with the SIMM or Single Inline Memory Module.
- SIMMs only have gold finger connections on one side of the module as well as all of the necessary Command/Address and PWR/GND.
- The module matched up with the 386 and 486 32-bit processor buses.

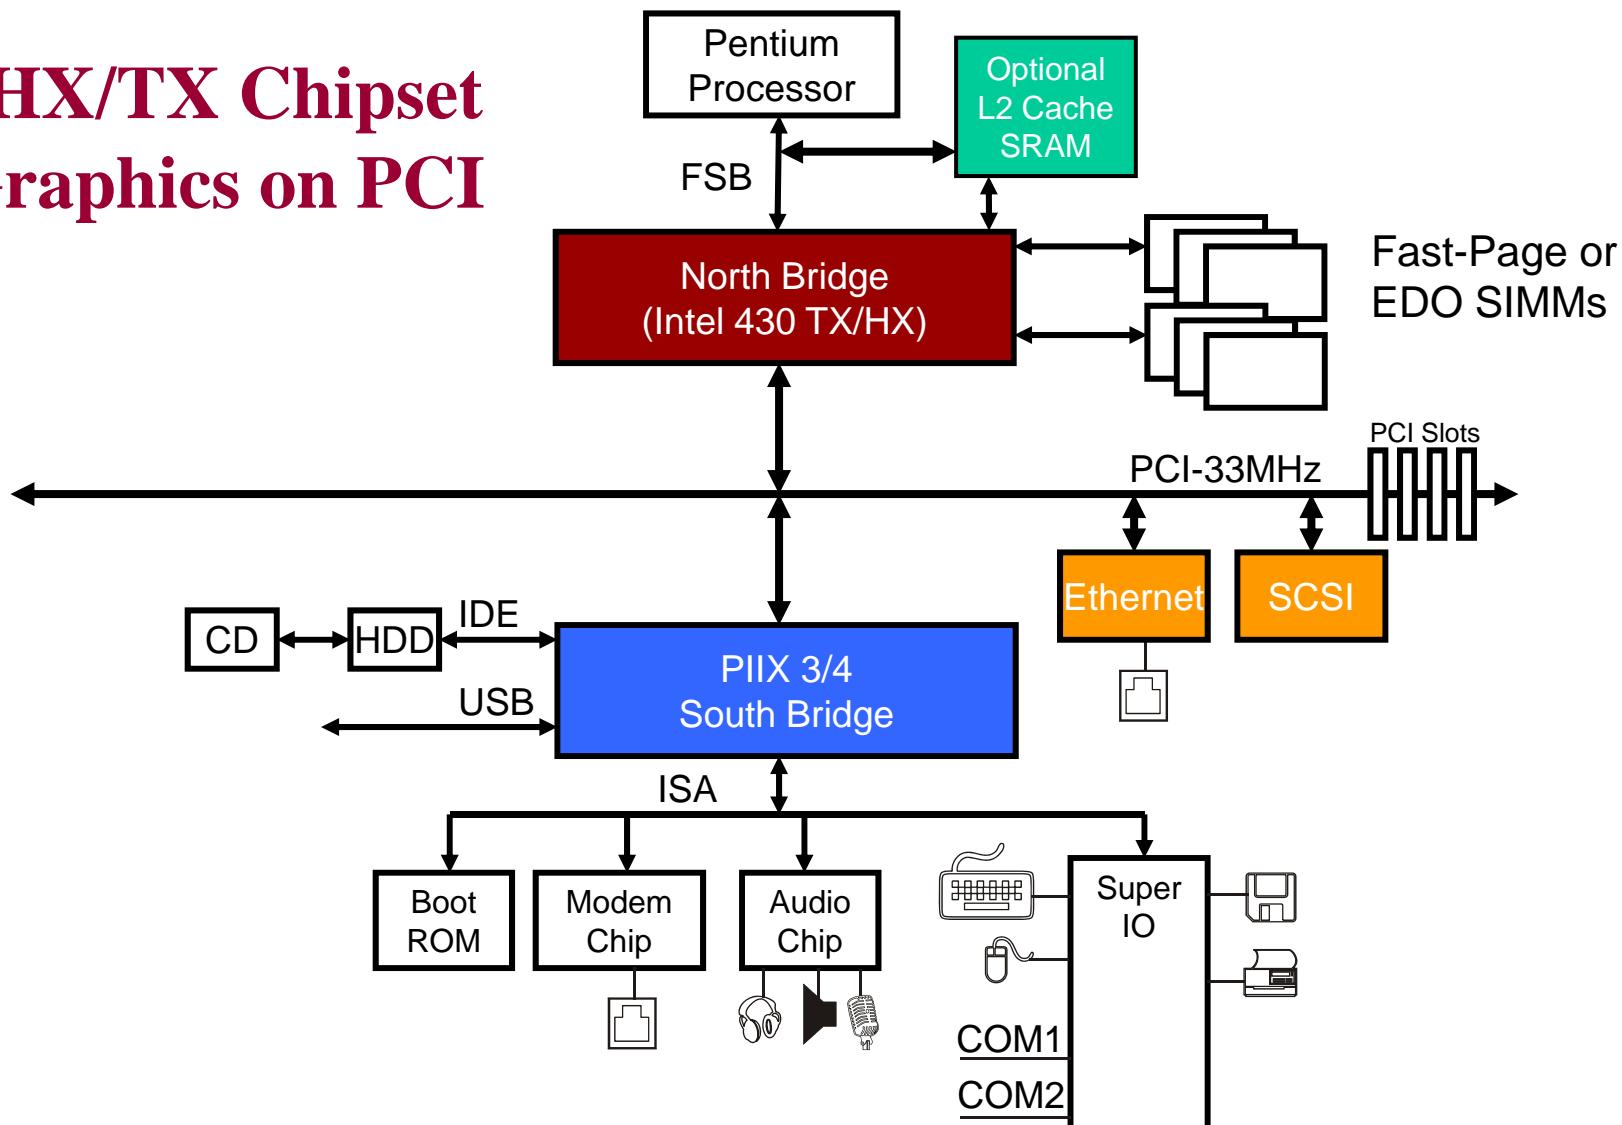


- To speed the Pentium and K5 to market, the same 32-bit SIMMs were used.
- Due to the Pentium and K5's 64-bit memory data bus two SIMMs were used.





430 HX/TX Chipset w/ Graphics on PCI





- Gold fingers on both sides accommodating a full 64-bit-wide data bus and 8 bits of ECC as well as all of the necessary Command, Address, power and ground pins.
- 168-pin DIMMs, used for SDRAM
- 184-pin DIMMs, used for DDR SDRAM
- 240-pin DIMMs, used for DDR2 SDRAM
- 240-pin DIMMs, used for DDR3 SDRAM
 - Key is different than DDR2 DIMMs
- Beware that pins are numbered differently (such as location of pin 2) for SO-DIMMs than for UDIMMs.

- The most common DIMMs are Unbuffered.
- These DIMMs have 3 basic components:
 - The PCB substrate
 - The DRAMs
 - The SPD (Serial Presence Detect)
- These are mostly for the Desktop and Mobile markets.

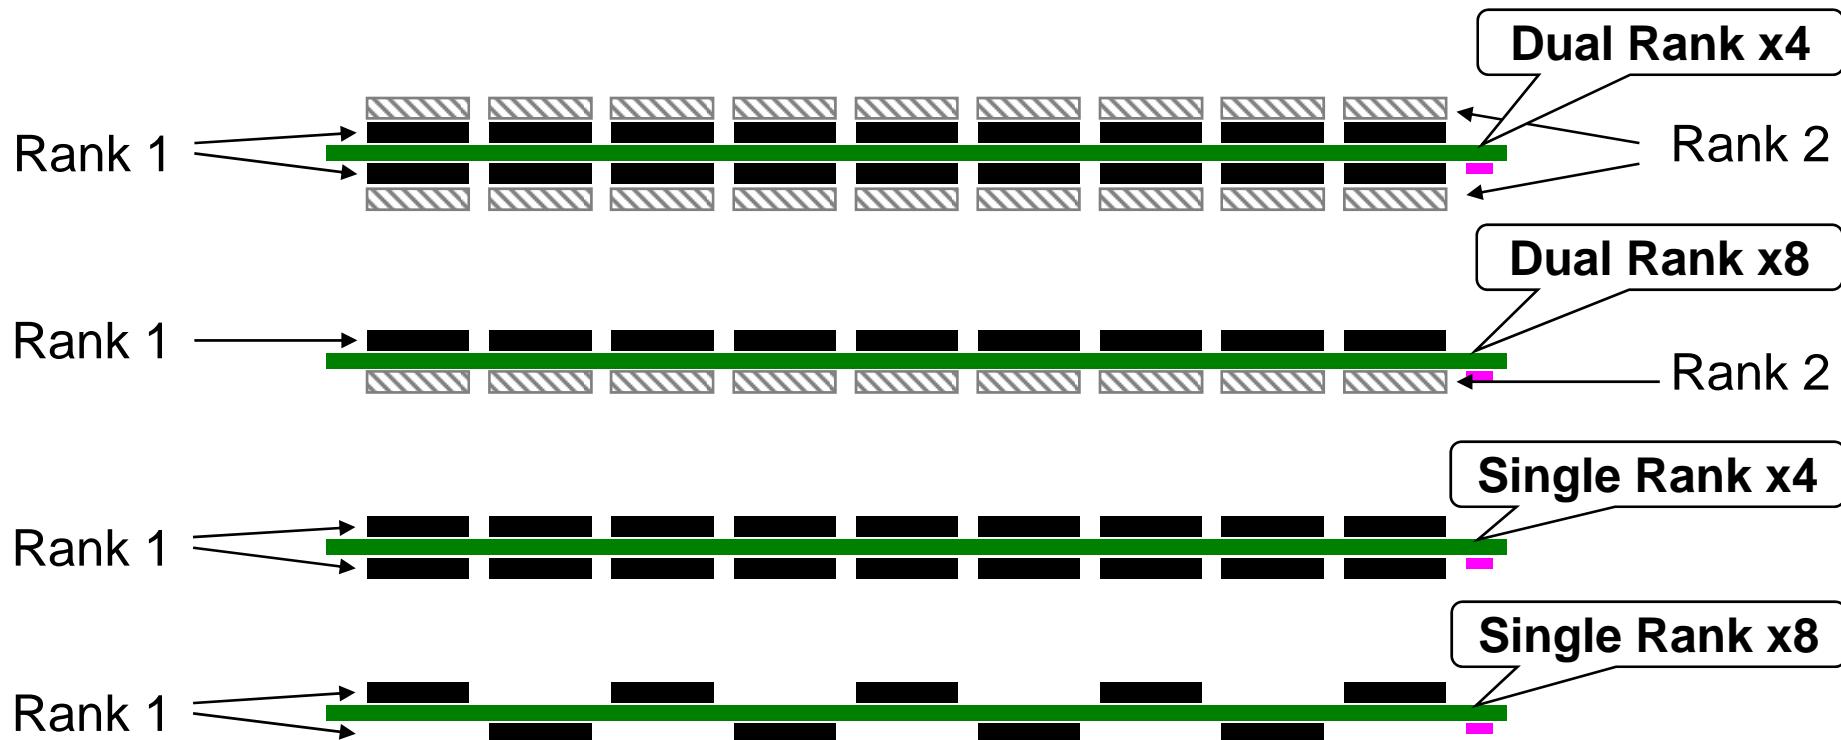
- The other very common type of DIMM used in server and workstation PCs is the Registered DIMM. Unlike an Unbuffered DIMM where all of the signals are routed directly to the DRAMs, the Registered DIMM uses registers and a PLL for the command, address, and clocks.
 - The PCB substrate
 - The DRAMs
 - The SPD (Serial Presence Detect)
 - Registers
 - PLL
- The registers and PLL reduce electrical loading on the most heavily loaded signals.



- The DRAMs on DIMMs are ordered into Ranks.
- A Rank consists of enough DRAMs to drive the full width of the system data bus. In most systems this is 64 bits wide or 72 bits with ECC.
- A quadword is 64 bits (eight bytes) using the Intel Architecture definition that one word is 16 bits.
- This concept used to be referred to as single sided and double sided.

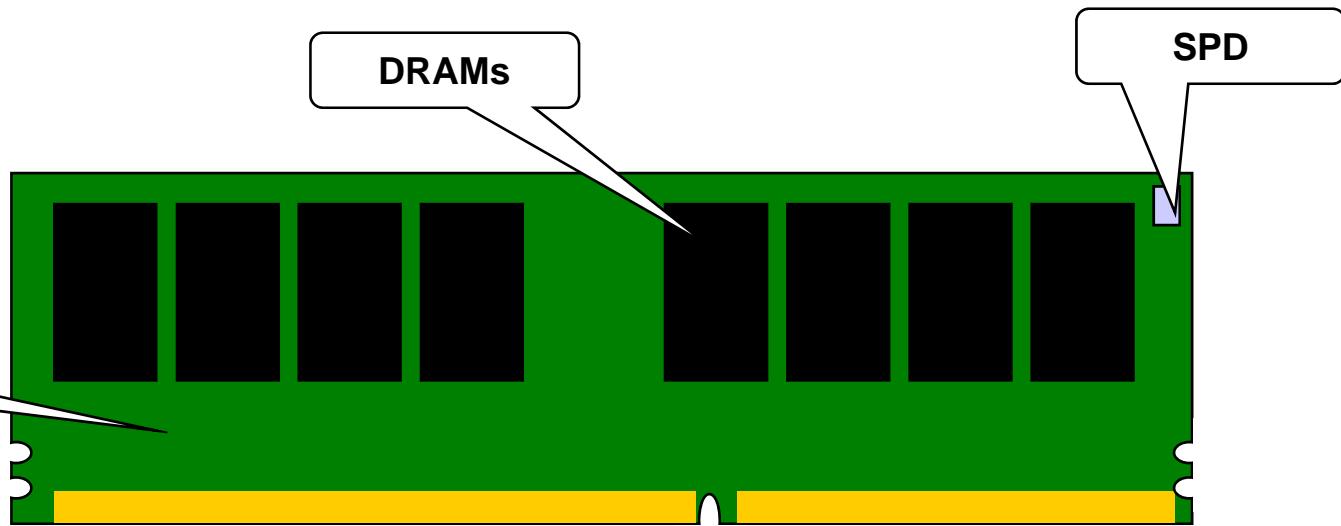


- Examples of Rank Configurations:
 - Shown with ECC

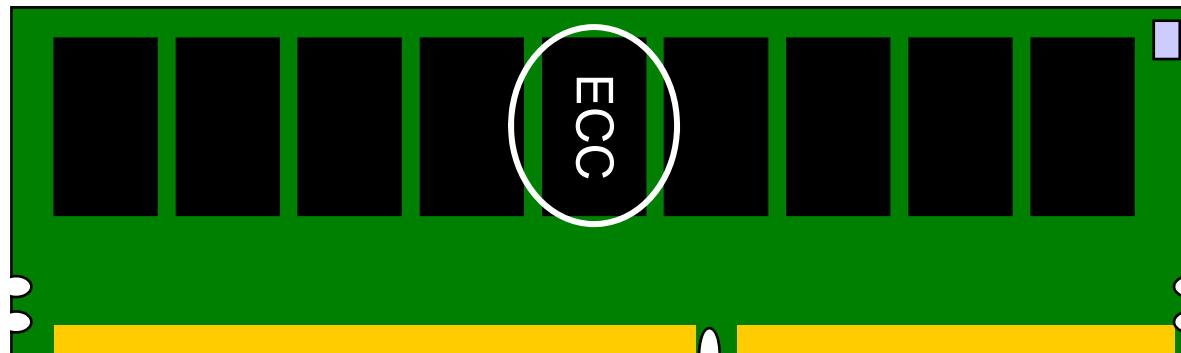




Module Types Unbuffered

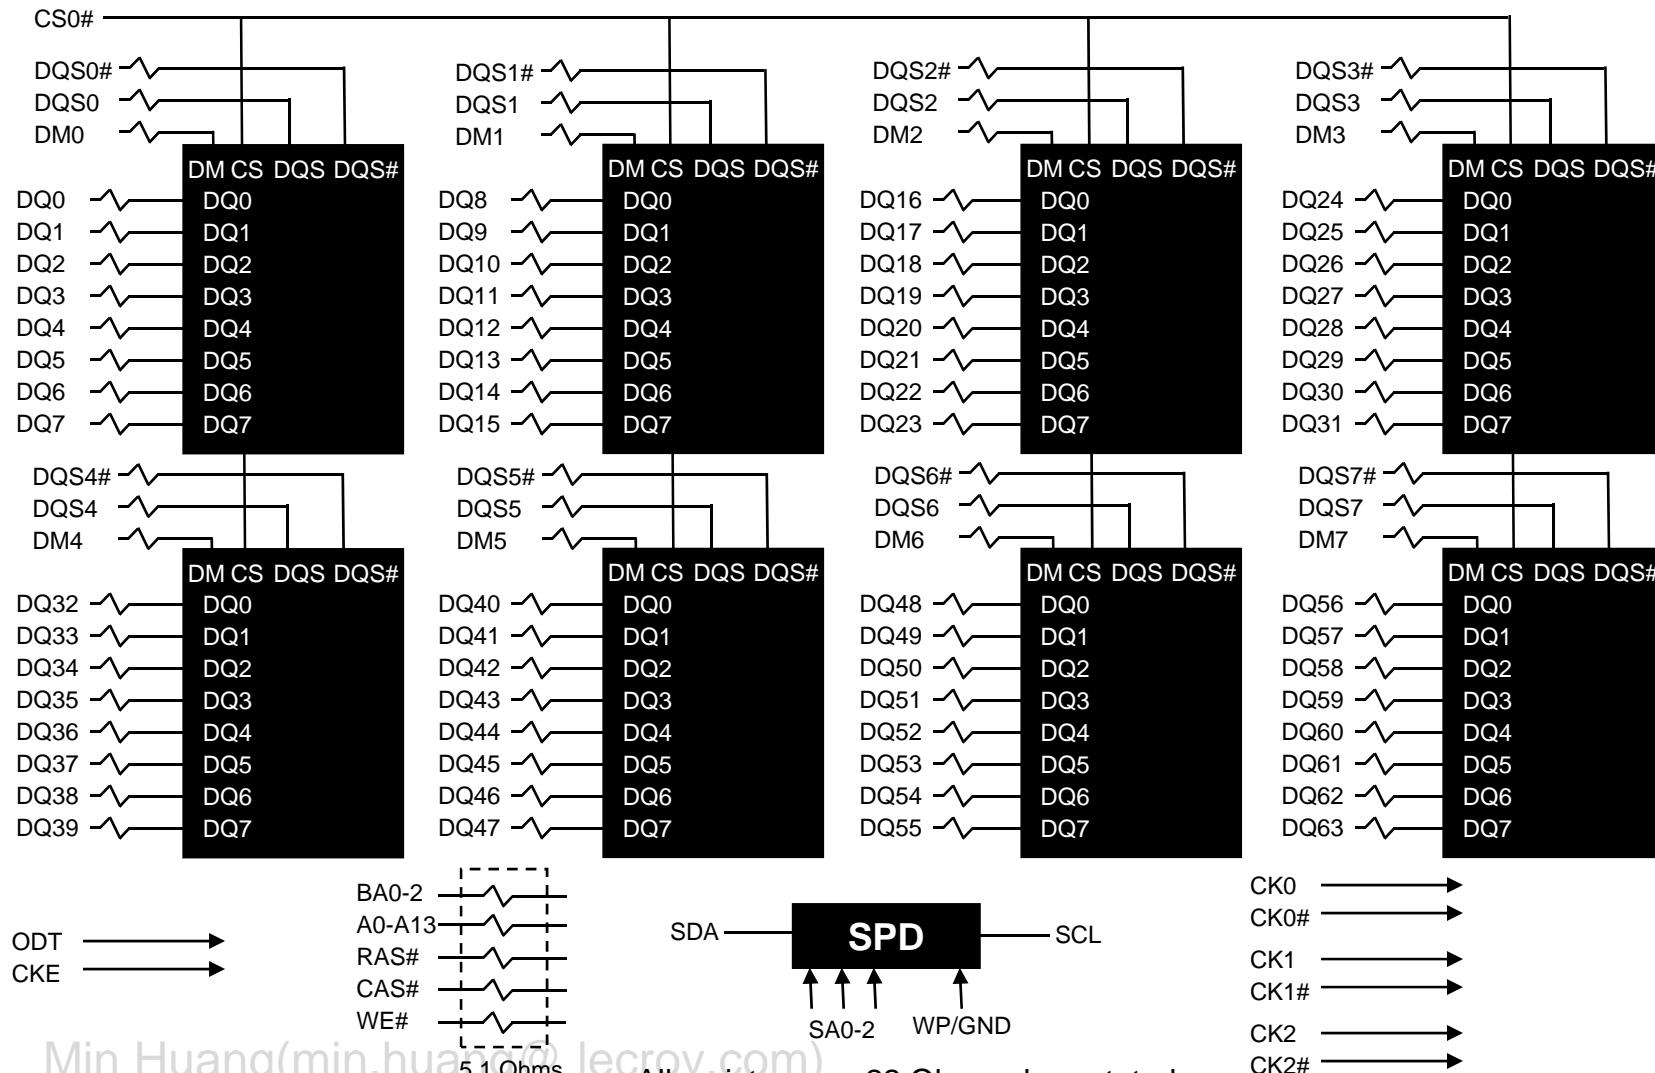


Shown with optional ECC chip. Typically unbuffered DIMMs do not have ECC support.





Unbuffered Logical Example





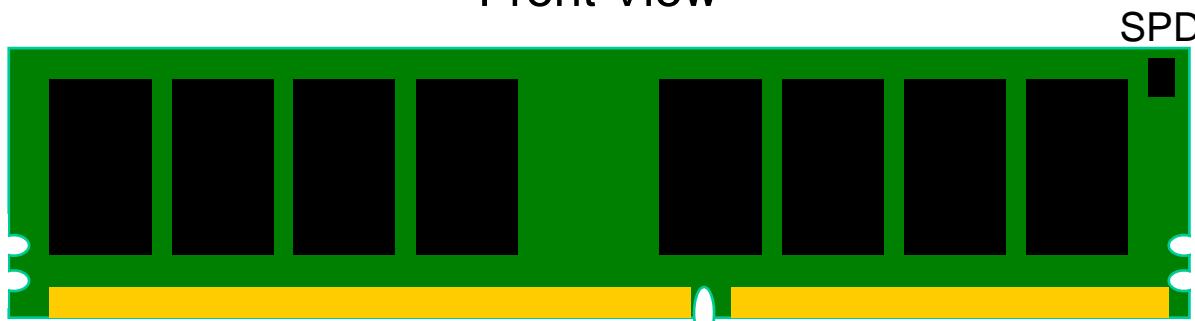
- This example assumes 512Mb DRAMs on the DIMMs.
- Total size = number of chips X size of chip / 8

	Configuration	# of chips	Ranks	Total Size
Unbuffered Mobile Desktop	x16	4	Single	256MB
	x16	8	Dual	512MB
	x8	8	Single	512MB
Registered Server Workstation	x8	16	Dual	1GB
	x4	16	Single	1GB
	x4	32	Dual	2GB



x8 Single Rank, 8 components

Front View



Side View



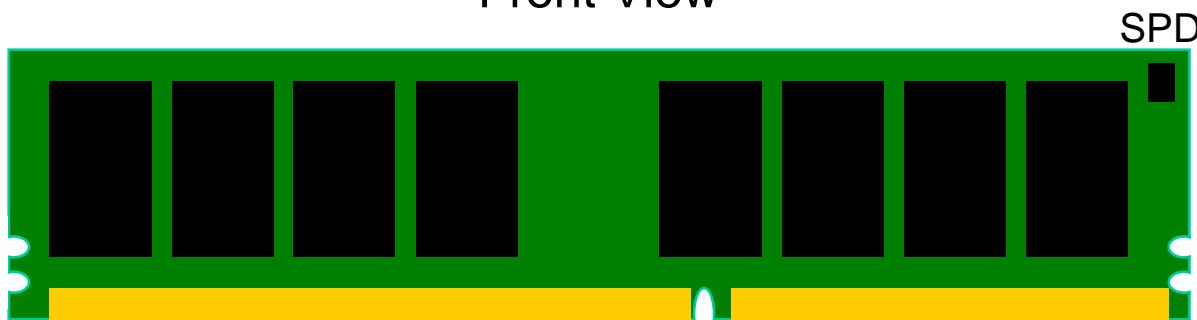
Top View





x8 Dual Rank, 16 components

Front View



Side View



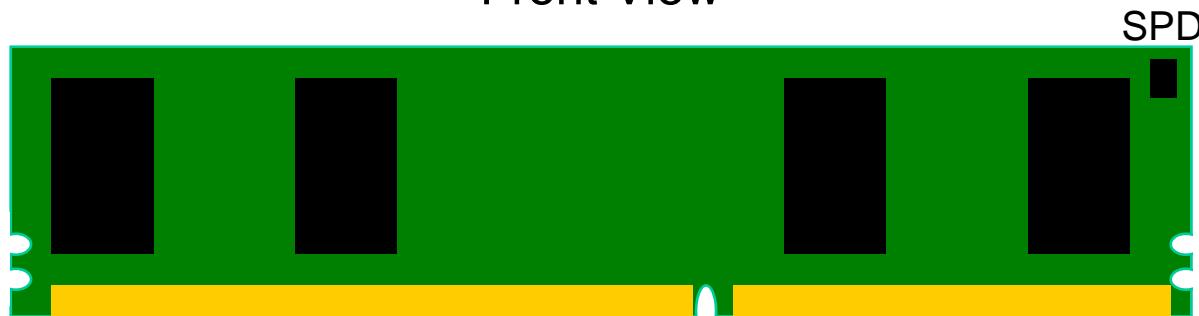
Top View





x16 Single Rank, 4 components

Front View



Side View

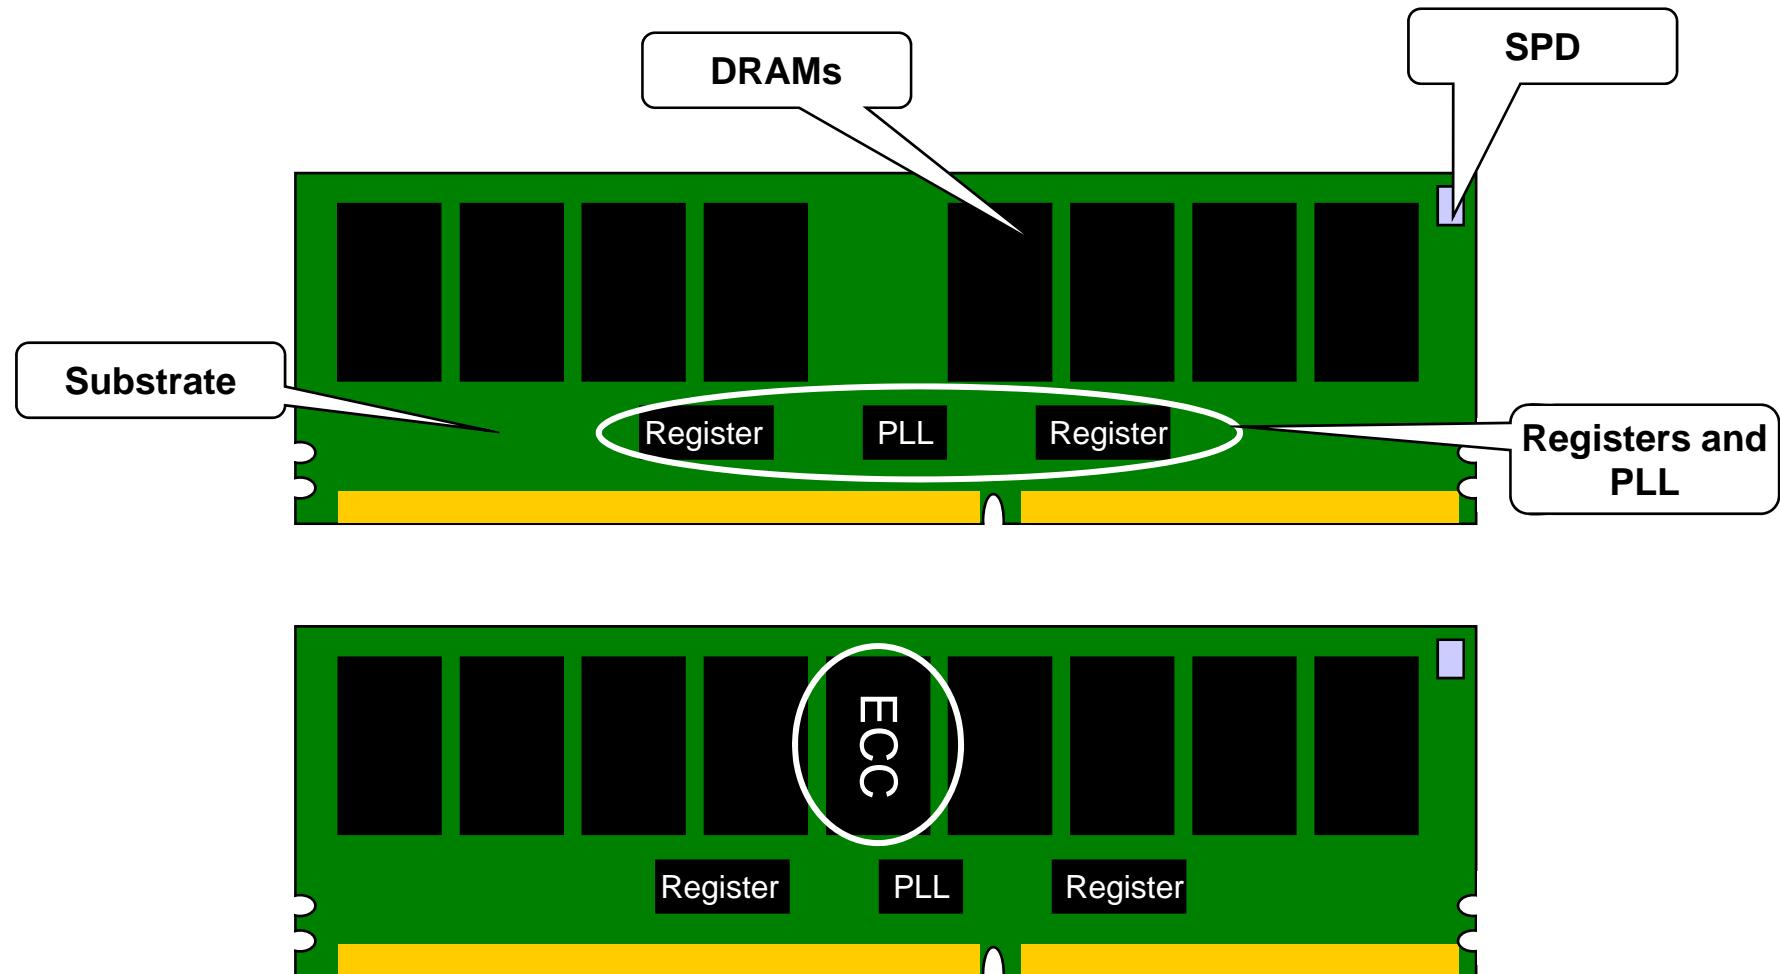


Top View



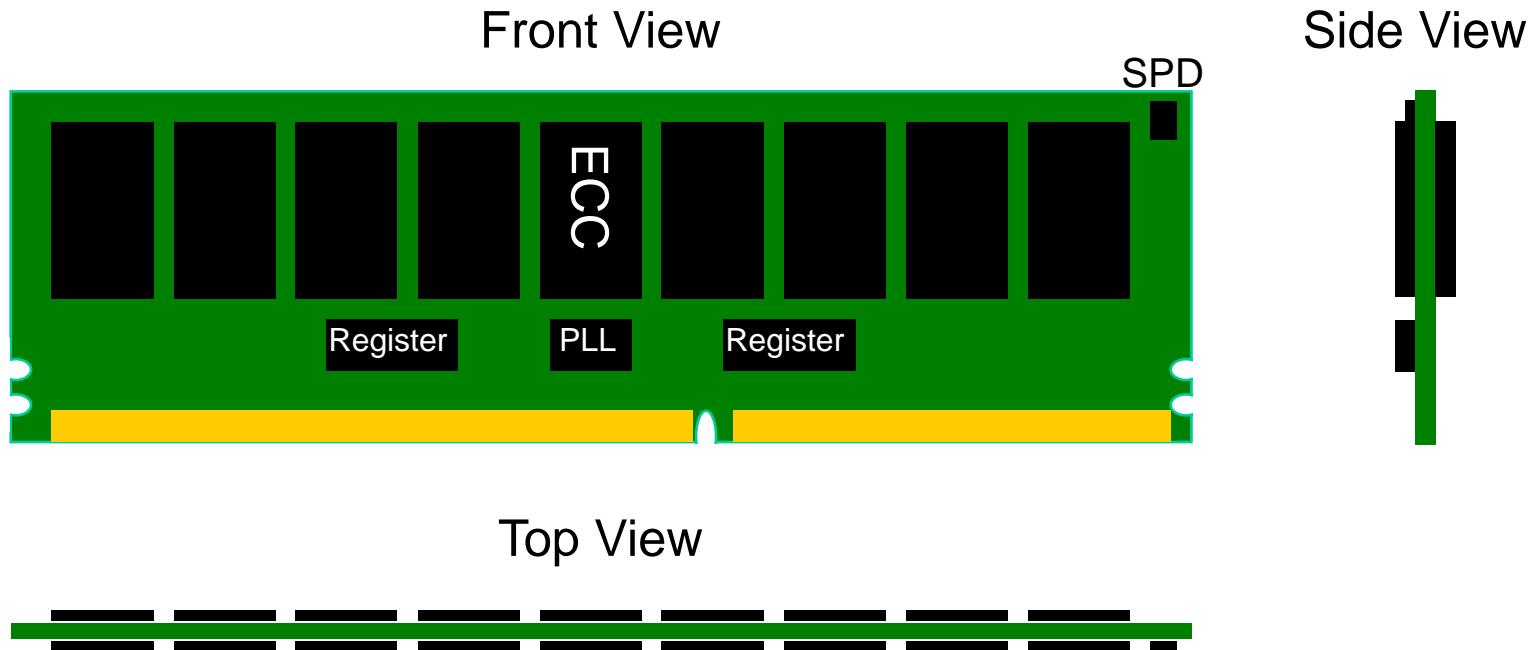


Module Types Registered





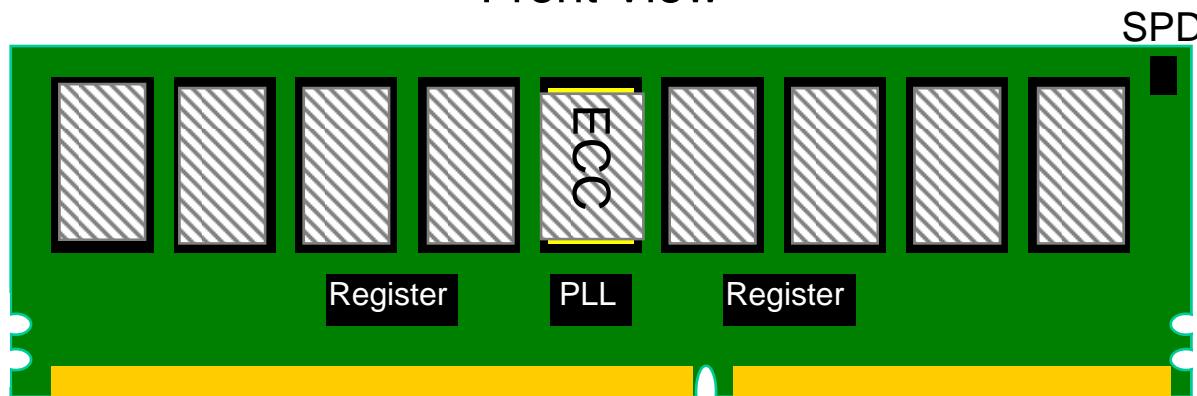
x4 Single Rank, 18 components (with ECC)



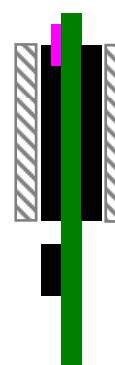


x4 Dual Rank, 36 components (with ECC)

Front View



Side View

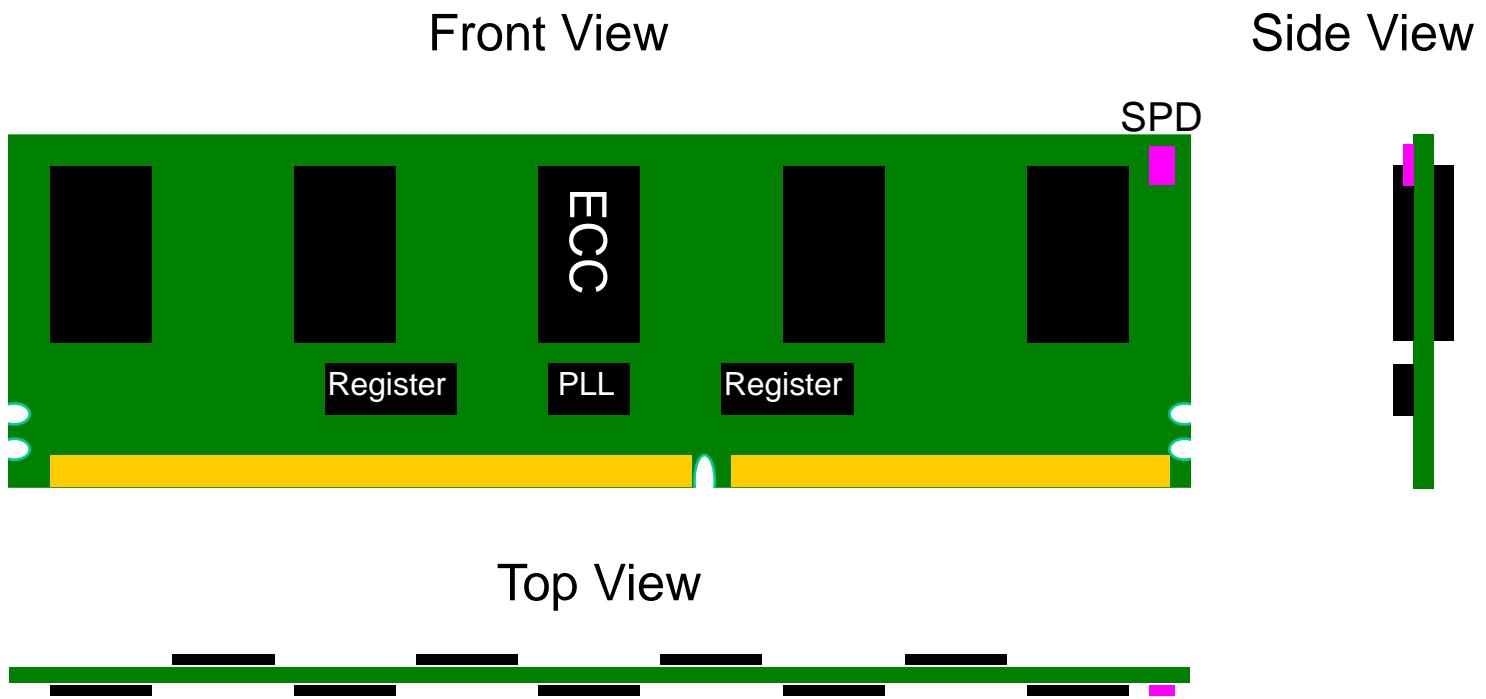


Top View





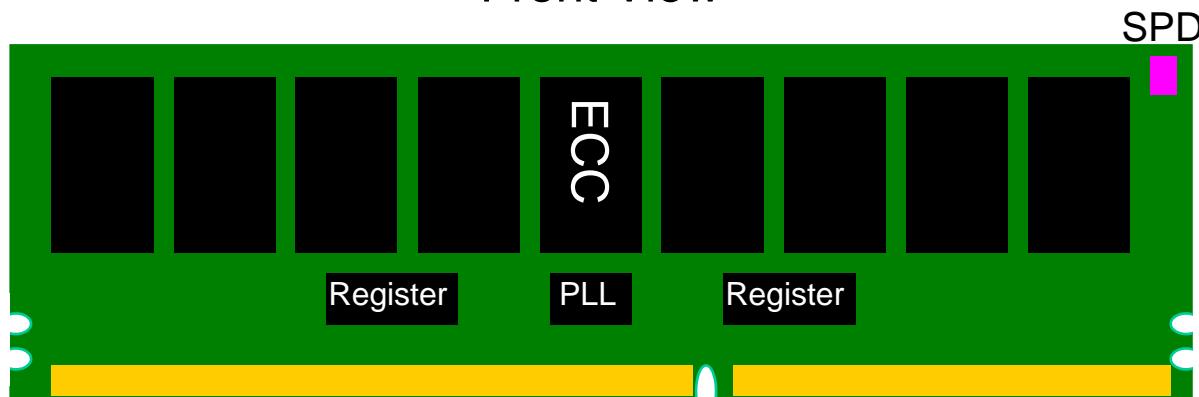
- x8 Single Rank, 8 components (with ECC)





- x8 Dual Rank, 18 components (with ECC)

Front View



Side View



Top View





Module Types

UDIMM: Unbuffered Desktop standard



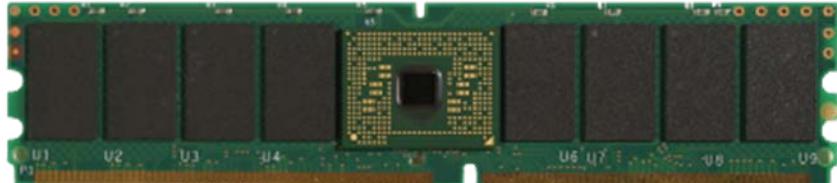
RDIMM: Registered Server standard



SODIMM: Notebook standard



FBDIMM: Fully Buffered Server



VLP RDIMM: Very Low Profile Computing and Networking



MiniDIMM:
Computing and Networking



VLP MiniDIMM:
Computing and Networking





Clock (CK, CK#)

- Input Differential For DDR2 single ended for DDR1
- Address & control signals sampled on the crossing of the positive edge of CK & the negative edge of CK#
- Output (read) data is referenced to the crossing of CK and CK# (both directions)
- One pair to each DIMM in Registered DIMM and 3 pairs for Unbuffered. For SODIMMs one of the clocks can be shut off

Clock Enable (CKE)

- Input High
- HIGH activates, LOW deactivates internal CK signals, input buffers and output drivers.
- Must remain high throughout read/write accesses. Primarily used for self refresh and sleep states like active power down.
- Each Rank has its own CKE



Chip Select (CS0#-CS3#)

- Input Low
- All commands are ignored when CS# is inactive (considered part of each command)
- Each rank has its own chip select.
- Which CS# is asserted for a given physical address is based on the programming of the memory controller (DRB's – DRAM Rank Boundary Registers)
- How many ranks can be supported on a single DIMM.

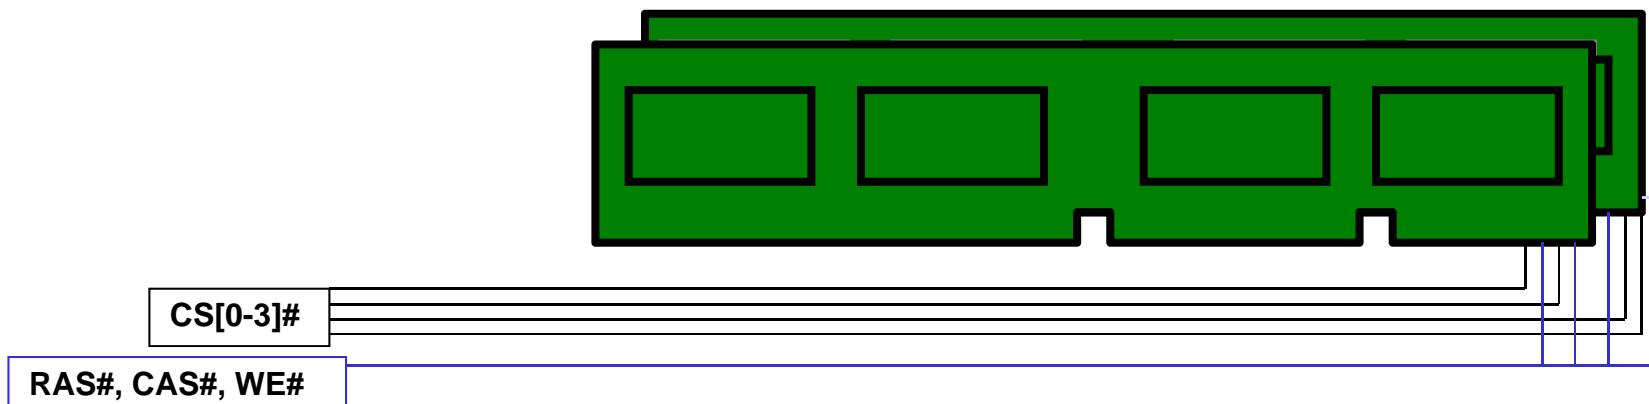
ODT

- Input High
- On Die Termination enables internal termination resistors to the following signals: DQ, DQS, DQS#, CB and Data Mask.
- Every Rank has its own ODT signal.
- ODT will be discussed more later



Row Address Strobe (RAS#), Column Address Strobe (CAS#), Write Enable (WE#)

- Input Low
- Define (along with CS#) the command being entered.
- Command truth table defined in JEDEC spec
- RAS#, CAS#, and WE# are broadcast to each chip.





Bank Address (BA 0-2)

- Input High
- Define to which bank an Active, Read, Write or Precharge command is being applied.
- Determines if the mode register or extended mode register is being accessed during an MRS or EMRS cycle.

Address (A 0-15)

- Input High
- Defines the Row Address for Active commands
- Defines the Column Address and auto precharge bit for read/write commands.
- A10 is Row Address during an Active (Activate) command.
- A10 is sampled during a Precharge command.
 - A10 Low = Precharge one bank
 - A10 High = Precharge all banks
- A10 is Auto Precharge (AP) during a Read or Write command.
 - A10 Low = No Precharge
 - A10 High = Auto Precharge when command is complete.

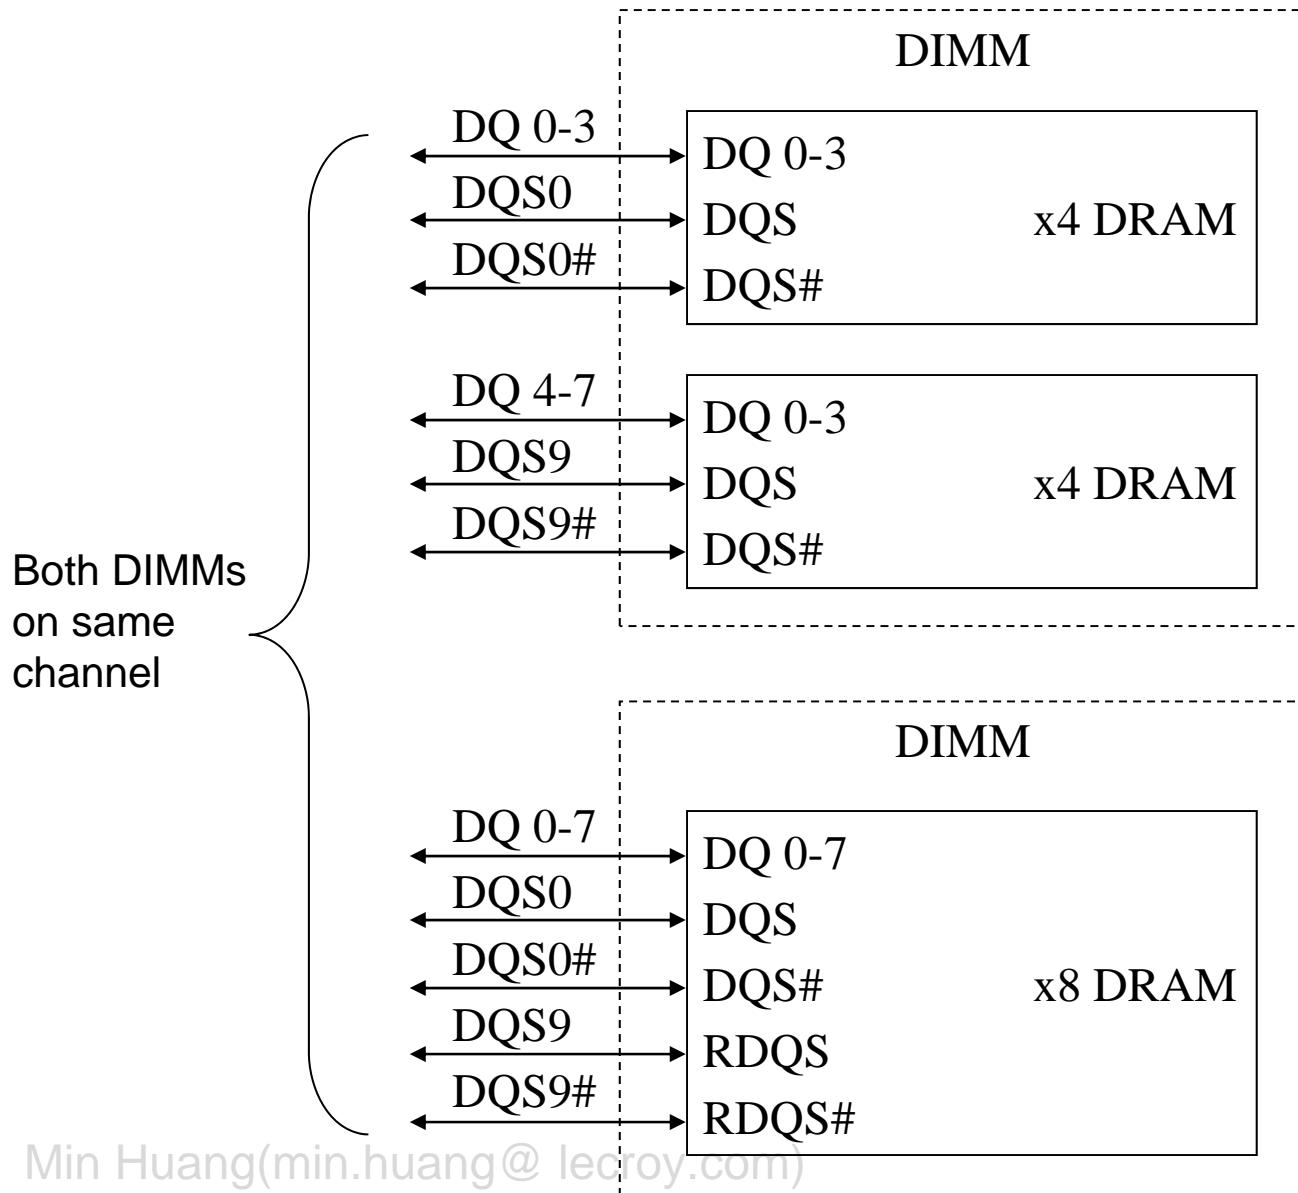


Data Input/Output (DQ 0-63)

- Bi-directional data bus.
- DQ is a bused signal so every rank is seen as a load on the bus.

Data Strobes (DQS, DQS#, UDQS, UDQS#, LDQS, LDQS#, RDQS, RDQS#) 0-17, 0-17#

- Bi-directional Differential in some cases (EMRS Programming)
- Used to latch data signals at the receiver.
- Edge-aligned with data (DQ) during reads (DRAM driving)
- Center-aligned with data (DQ) on writes (Chipset driving)
- For x4 devices, 16 or 18 pairs are used, one pair per chip (DQS, DQS#)
- For x8 devices, 8 or 9 pairs are used, one pair per chip (DQS, DQS#), except that in RDQS mode, 16 or 18 pairs are used (DQS, DQS#, RDQS, RDQS#)
- For x16 devices, 8 or 9 pairs are used, 2 pairs per chip (LDQS, LDQS#, UDQS, UDQS#). For x16, the ECC chip is x16, not x8.
- DQS is bused, so every rank is seen as a load on the bus.





Data Mask (DM 0-7, UDM, LDM)

- Input High
- Input mask signal for write data. When DM is high, write DQS is ignored. When DM is low, DQS is valid.
- Same pins as DQS 9-17 on DIMMs.
- Pins are DM for DIMMs using x8 and x16 devices
- Pins are DQS for DIMMs using x4 devices. DM of x4 devices are tied low on DIMMs.
- For x16 devices, these signals are named UDM and LDM. On DIMMs with x16 devices, there is an unconnected set of DM pins, which are the pins that would have been the DQS of the x4 devices.
- DM is bused, so every rank is seen as a load on the bus.



Check Bits (CB 0-7)

- Bi-directional data bus.
- Check bits used for ECC.
- CB 0-7 are strobed with DQS 8 on x8 devices, and with DQS 8 and 17 on x4 devices.

Par_In

- Parity for the address and control bus. Not used on most Intel controllers. Implemented in the chipset and register.

Err_Out#, QERR#

- Indicates parity error found on the address or control bus.
- QERR# is register pin. Err_Out# is DIMM pin.

Serial Presence Detect (SPD) Clock Input (SCL)

- Serial clock used to synchronize the SPD logic.

SPD Data (SDA)

- Bidirectional pin to transmit address and data in and out of the SPD.

SPD Address (SA 0-2)

- Must be hardwired to different addresses for each module.



RESET# - Active when low

- CKE must be pulled low first.
- When RESET# is deasserted, CKE may be deasserted 500uS later. During this time internal initialization is started.
- RESET# forces DRAM into a defined state.
- RESET# can be asserted asynchronously during any operation.
- Data in DRAM may be lost and DRAM must be re-initialized, which includes (but not limited to) load mode registers and DLL reset.

A12/BC# Burst Chop - Active when low

- BC# during Read and Write commands.
- A12 during Active (Activate) command.

TDQS/TDQS# - Termination Data strobe

- Only applicable to x8 DRAMs (replaces RDQS/RDQS#).

V_{REFDQ} , V_{REFCA} – Separate V_{REF} for DQ and CMD/ADDR

- This decoupling reduces noise on the reference planes.

ZQ – Reference pin for ZQ calibration

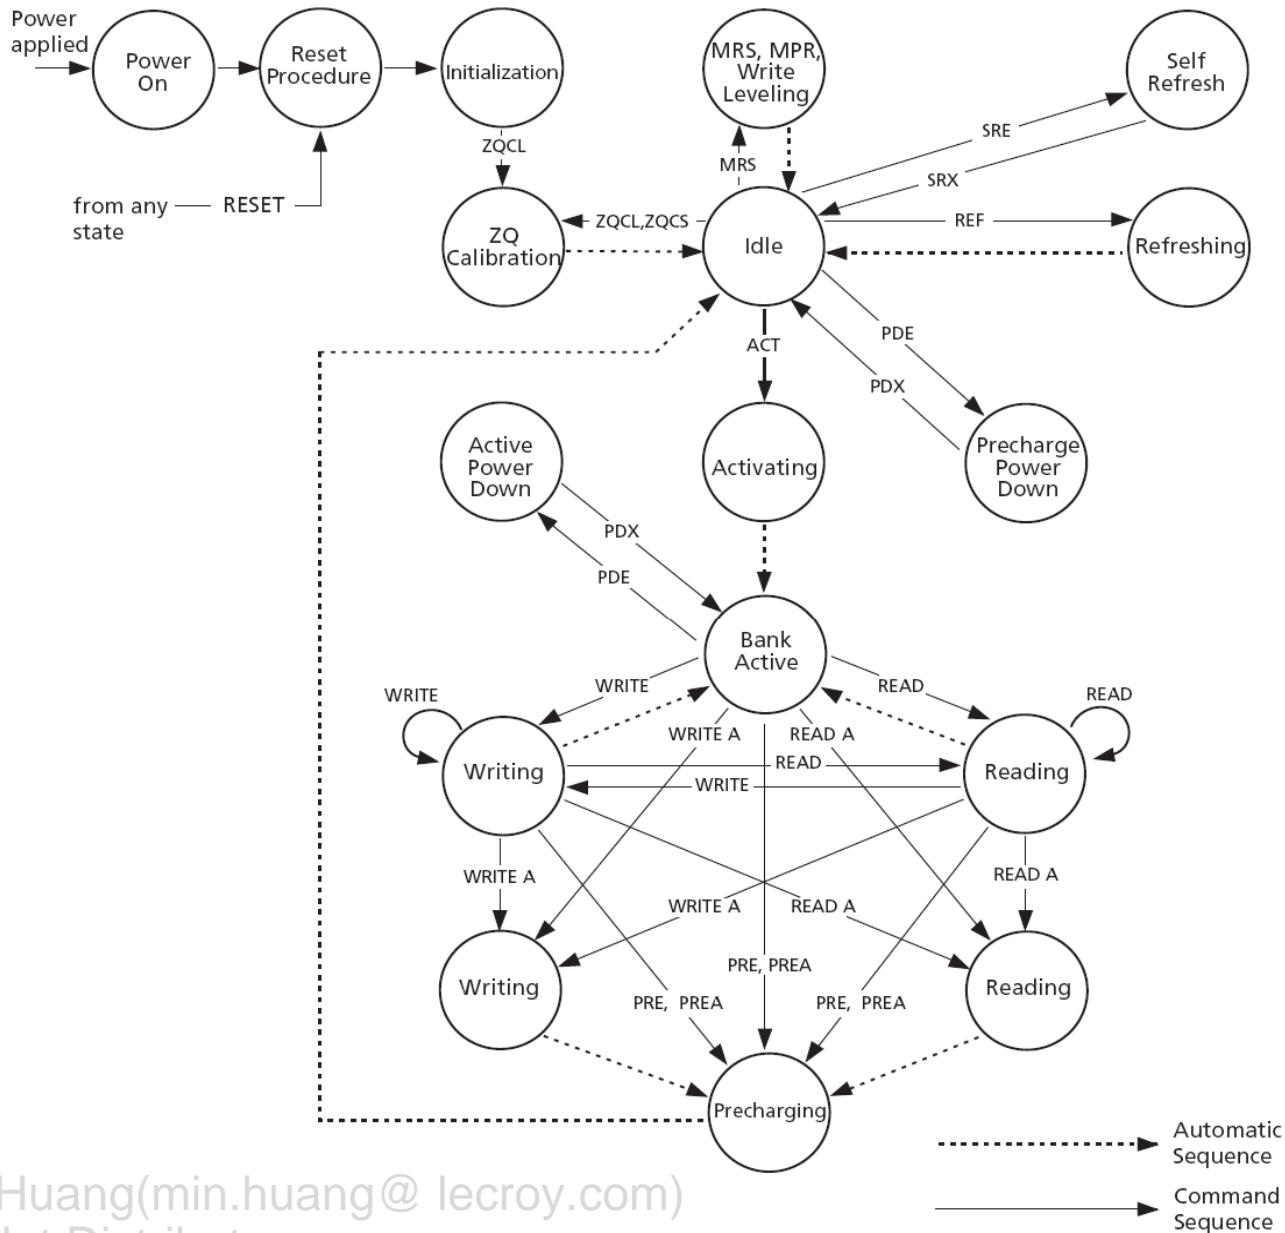
CK,CK# - Unbuffered DDR3 DIMMs only have 2 pairs instead of 3.



Commands and Waveforms

Min Huang(min.huang@ lecroy.com)
Do Not Distribute

DDR3 Bank States





DDR2 Commands

Command	CKE		CS#	RAS#	CAS#	WE#	BA2 BA1 BA0	A13 A12 A11	A10 AP	A9-0
	Previous Cycle	Current Cycle								
Mode Register Set	H	H	L	L	L	L	MRn	Op Code		
Refresh	H	H	L	L	L	H	X	X	X	X
Self Refresh Entry	H	L	L	L	L	H	X	X	X	X
Self Refresh Exit	L	H	L	H	H	H	X	X	X	X
Single Bank Precharge	H	H	L	L	H	L	BA	X	L	X
All Banks Precharge	H	H	L	L	H	L	X	X	H	X
Bank Activate	H	H	L	L	H	H	BA	Row Address		
Write	H	H	L	H	L	L	BA	Col Addr	L	Col Addr
Write with Auto Precharge	H	H	L	H	L	L	BA	Col Addr	H	Col Addr
Read	H	H	L	H	L	H	BA	Col Addr	L	Col Addr
Read with Auto Precharge	H	H	L	H	L	H	BA	Col Addr	H	Col Addr
No Operation	H	X	L	H	H	H	X	X	X	X
Device Deselect	H	X	H	X	X	X	X	X	X	X
Power Down Entry	H	L	L	H	H	H	X	X	X	X
Power Down Exit	H	L	H	H	H	H	X	X	X	X



DDR3 Command Additions

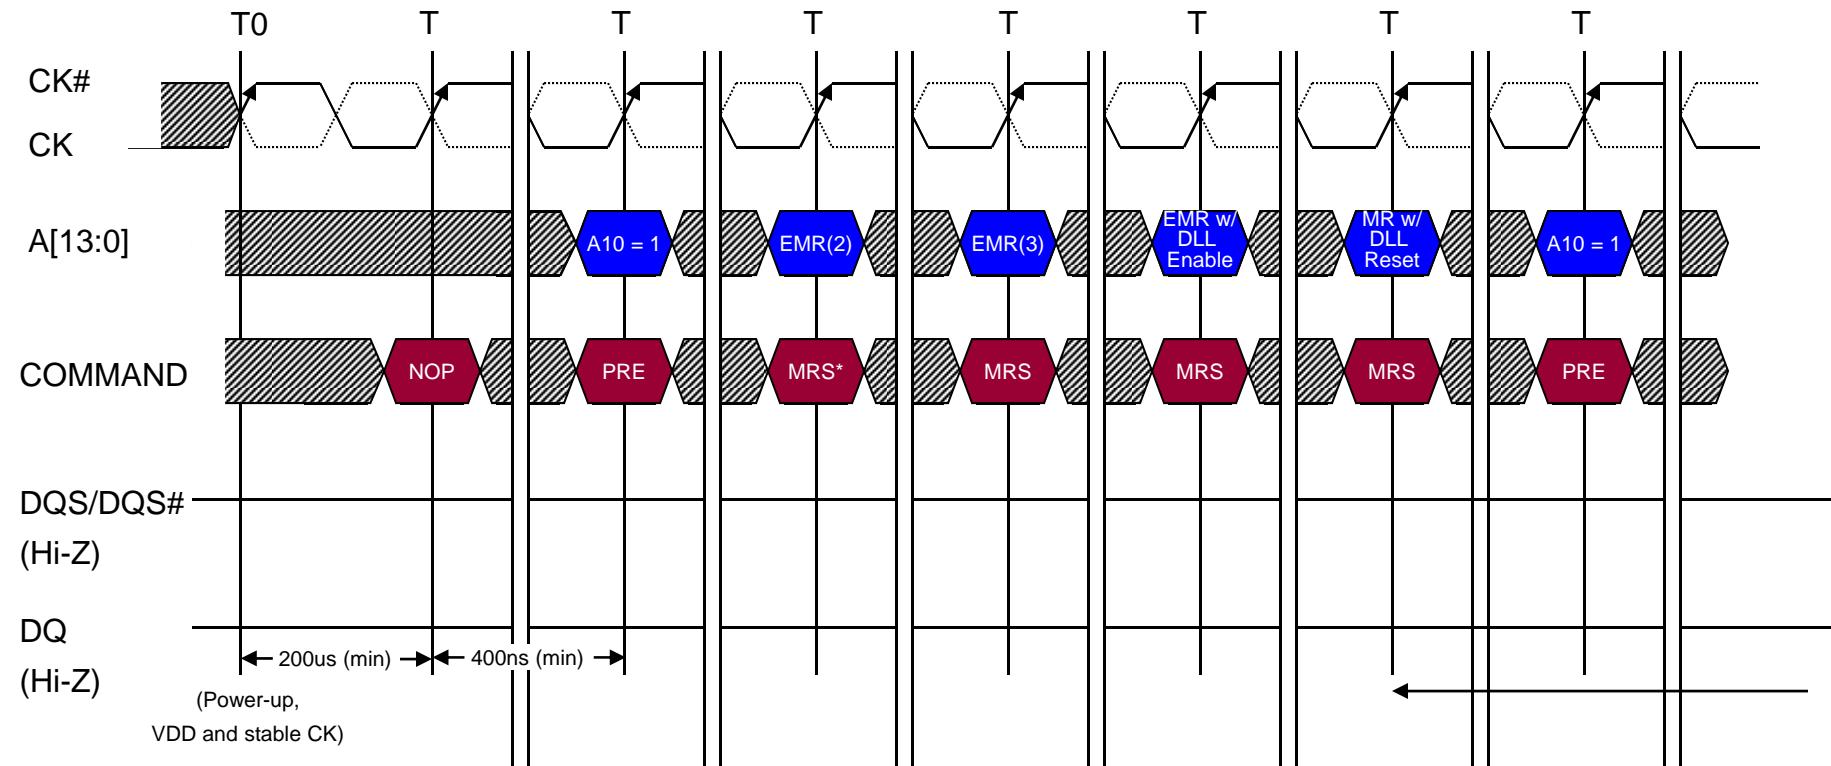
Command	CKE		CS#	RAS#	CAS#	WE#	BA0- BA3	A13 A14 A15	A12 BC#	A10 AP	A0-9, A11
	Previous Cycle	Current Cycle									
Write BL= 4 (Burst Chop)	H	H	L	H	L	L	BA	CA	L	L	CA
Write BL= 8	H	H	L	H	L	L	BA	CA	H	L	CA
Write with Auto Precharge BL=4	H	H	L	H	L	L	BA	CA	L	H	CA
Write with Auto Precharge BL=8	H	H	L	H	L	L	BA	CA	H	H	CA
Read BL= 4 (Burst Chop)	H	H	L	H	L	H	BA	CA	L	L	CA
Read BL= 8	H	H	L	H	L	H	BA	CA	H	L	CA
Read with Auto Precharge BL=4	H	H	L	H	L	H	BA	CA	L	H	CA
Read with Auto Precharge BL=8	H	H	L	H	L	H	BA	CA	H	H	CA
ZQ Calibration Long	H	H	L	H	H	L	X	X	X	H	X
ZQ Calibration Short	H	H	L	H	H	L	X	X	X	L	X



- Mode Register Set – Sometimes referred to as Load Mode Register.
- Mode Register Set issued to load the MRS and EMRS values via the Bank Address and Address lines.
- The MRS command can only be used when all banks are idle. No command can be issued after MRS command until t_{MRD} is met.

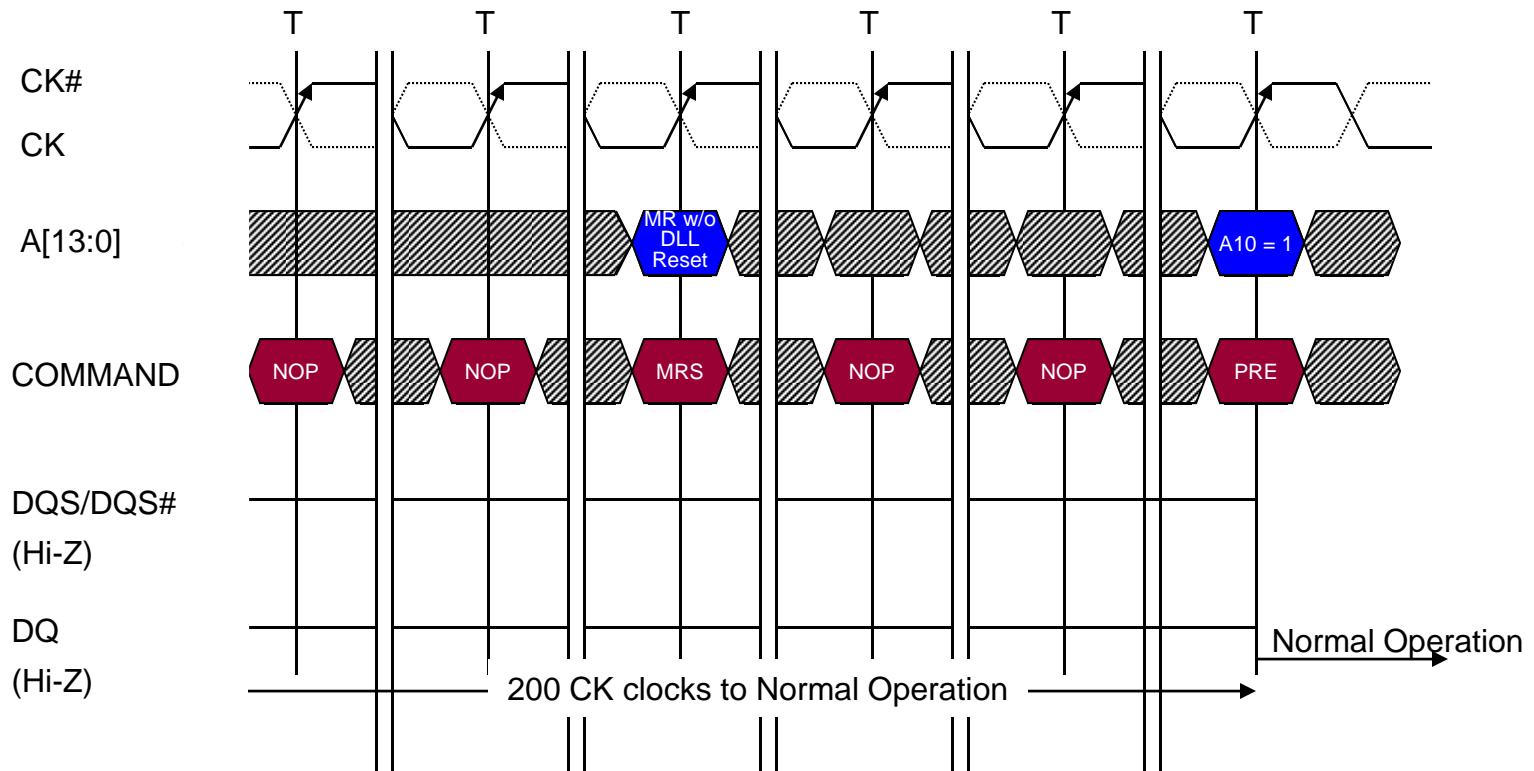


Mode Register Set Waveforms



*MRS = Mode Register Set Command used to initialize SDRAM registers

4 SDRAM registers to initialize consist of: MR (Mode Register), EMR (Extended-MR), EMR(2) and EMR(3)



OCD = Off-Chip Driver Impedance Calibration.

- The Activate command is used to open (activate) a row in a particular bank for a subsequent access.
- The values of BA0 through BA2 select the bank to be activated.
- The row remains active for accesses until a Precharge (or Read or Write with Autoprecharge) is issued to that bank.
- Only one row per bank can be open at a time.

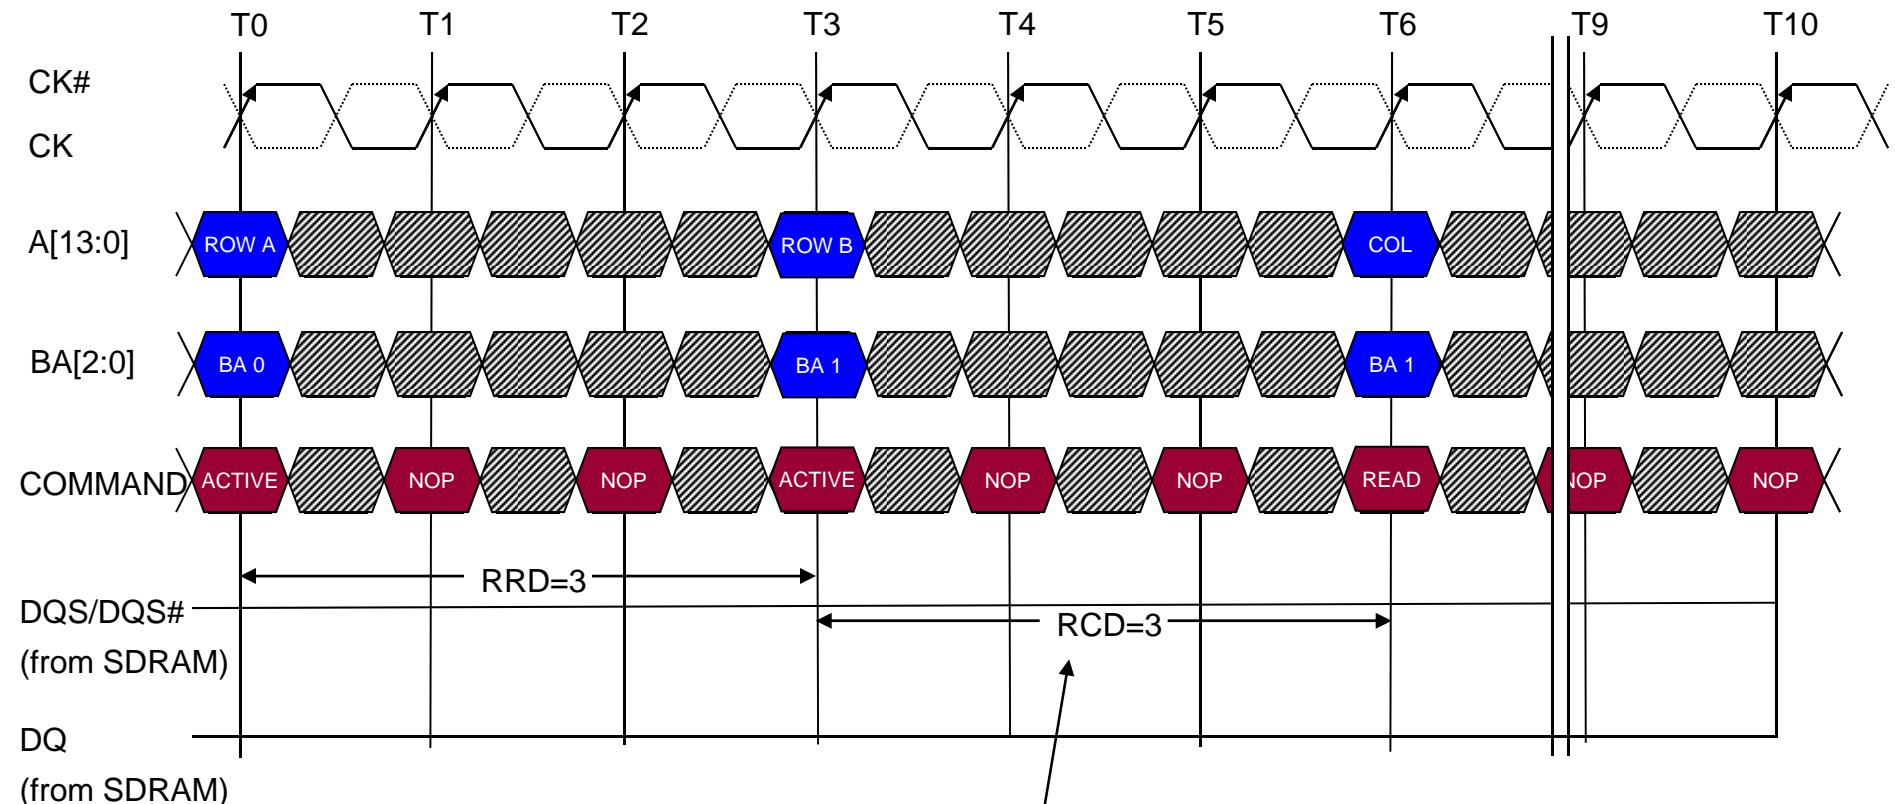


- There are typically 3 timings that are referred to for DIMM speed besides the frequency
 - CL - Column Address Strobe Latency is the amount of time in base clocks from when CAS is asserted until data should be valid.
 - DDR2 requires the clock interval to be in whole clocks.
 - CAS latency is a function of the DRAMs internal speed. The faster the DRAM the lower the CAS Latency.
 - RCD – RAS-to-CAS Delay is the time in base clocks required from an Activate to a Read or Write.
 - RP – Time in base clocks required to precharge or write back a row.
 - Example: DIMM might say on it 4 – 4 – 4
CL – RCD – RP



Activate Waveform

107



t_{RRD} is the Minimum time interval from one Bank Activate to another Bank (RAS to RAS delay)

t_{RCD} is the Minimum time from activate to a read or write command (RAS to CAS delay)

DDR2 400 3 - 3 - 3



- The Read command is used to issue a burst read access to an active row. The value of BA0 through BA2 will determine the bank to be accessed.
- A0-Ax provide the column address.
- The value on A10 during the Read command determines whether or not to use Auto Precharge.
- If Auto Precharge is selected, the row currently active will be precharged at the end of the burst.



- All DDR SDRAM (DDR1, DDR2, DDR3) are burst oriented.
- Accesses start at a selected location and continue for a programmed number of locations in a programmed sequence.
- Two parameters define burst operation:
 - Burst Type: Interleaved or Sequential
 - Burst Length: 2, 4, or 8
- DDR1 supports burst of 2, 4, or 8.
- DDR2 supports burst of 4 or 8 only.
- DDR3 supports burst of 8 with option to chop to 4.
- When issuing back-to-back bursts of 4, CAS# can be asserted every other clock. In back-to-back bursts of 8, CAS# can be asserted every 4th clock.



- Programmed during initialization
- Burst type is programmed to be sequential or interleaved by bit A3 in the Mode Register (MR0).
- Interleaved is used on the FSB of all Intel chipsets and is also referred to as toggle mode.
- Sequential mode is used in all other PC-based designs, i.e., AMD.
- Burst length is programmed by A0-A2 in MR0 and determines the maximum number of column locations that can be accessed for a Read or Write command.
- All accesses for a burst take place within a block. The burst wraps within the block if a boundary is reached.
- The least significant address bit(s) is (are) used to select the starting location within the block.
- BL/2 is the minimum number of clocks to wait until the next CAS command



DDR2 Burst Order

Burst Length	Starting Column Address (A2,A1,A0)	Order of Access within a Burst	
		Burst Type = Sequential (Linear wrap)	Burst Type = Interleaved (Toggle Intel)
4	000	0,1,2,3	0,1,2,3
	001	1,2,3,0	1,0,3,2
	010	2,3,0,1	2,3,0,1
	011	3,0,1,2	3,2,1,0
8	000	0,1,2,3,4,5,6,7	0,1,2,3,4,5,6,7
	001	1,2,3,0,5,6,7,4	1,0,3,2,5,4,7,6
	010	2,3,0,1,6,7,4,5	2,3,0,1,6,7,4,5
	011	3,0,1,2,7,4,5,6	3,2,1,0,7,6,5,4
	100	4,5,6,7,0,1,2,3	4,5,6,7,0,1,2,3
	101	5,6,7,4,1,2,3,0	5,4,7,6,1,0,3,2
	110	6,7,4,5,2,3,0,1	6,7,4,5,2,3,0,1
	111	7,4,5,6,3,0,1,2	7,6,5,4,3,2,1,0



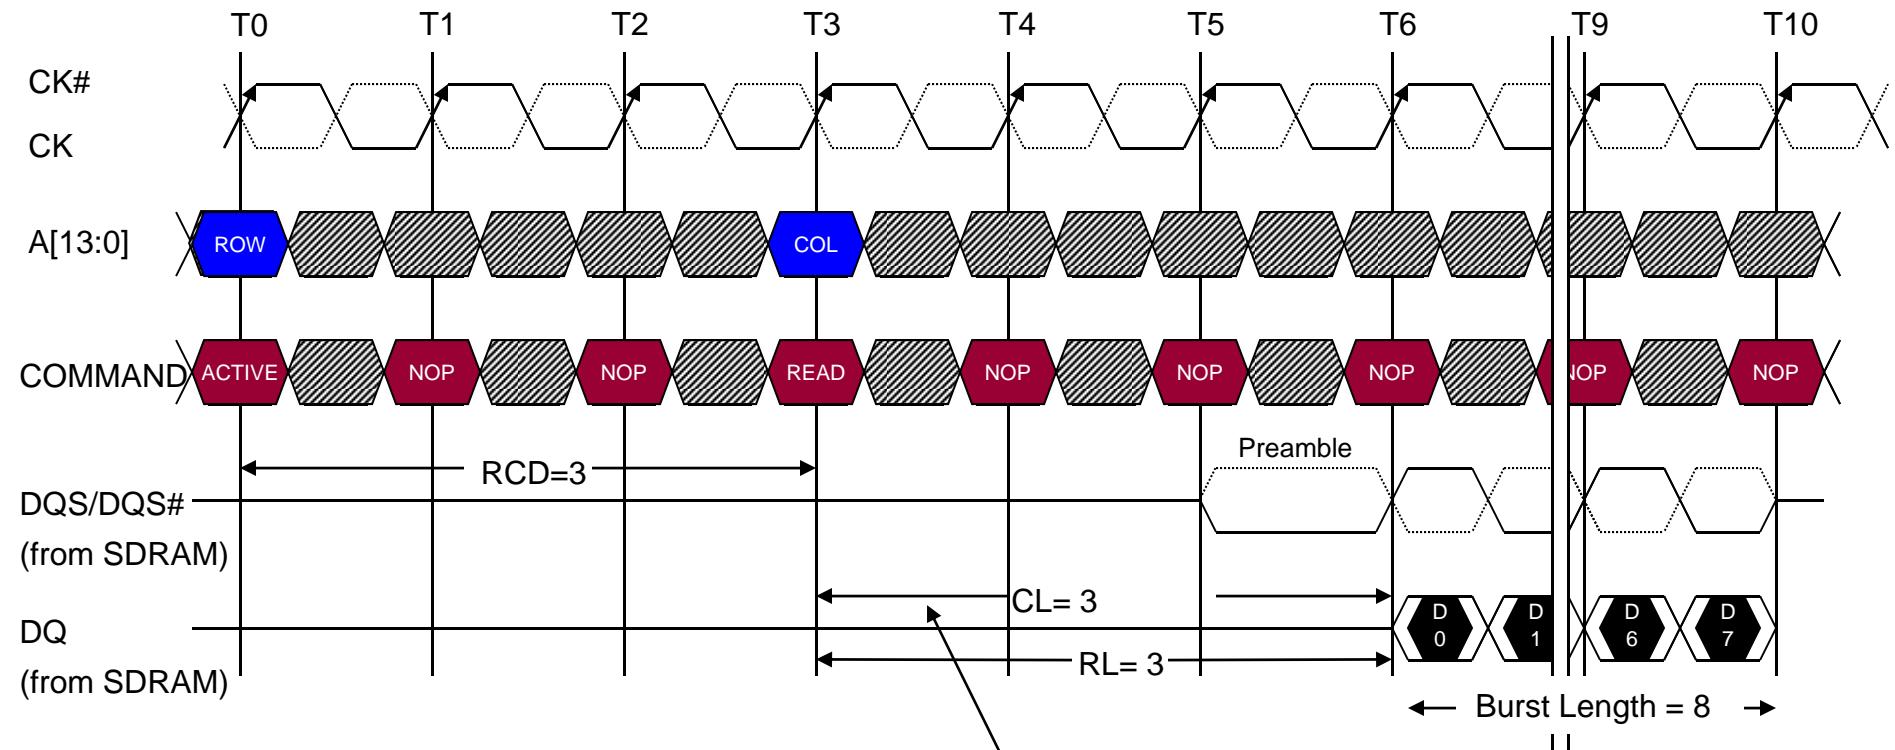
DDR3 Burst Order

Burst Length	Read/ Write	Starting Column Address (A2,A1,A0)	Order of Access within a Burst	
			Burst Type = Sequential (Linear wrap)	Burst Type = Interleaved (Toggle)
4 Chop	Read	000	0,1,2,3,Z,Z,Z,Z	0,1,2,3,Z,Z,Z,Z
		001	1,2,3,0,Z,Z,Z,Z	1,0,3,2,Z,Z,Z,Z
		010	2,3,0,1,Z,Z,Z,Z	2,3,0,1,Z,Z,Z,Z
		011	3,0,1,2,Z,Z,Z,Z	3,2,1,0,Z,Z,Z,Z
		100	4,5,6,7,Z,Z,Z,Z	4,5,6,7,Z,Z,Z,Z
		101	5,6,7,4,Z,Z,Z,Z	5,4,7,6,Z,Z,Z,Z
		110	6,7,4,5,Z,Z,Z,Z	6,7,4,5,Z,Z,Z,Z
		111	7,4,5,6,Z,Z,Z,Z	7,6,5,4,Z,Z,Z,Z
	Write	0VV	0,1,2,3,X,X,X,X	0,1,2,3,X,X,X,X
		1VV	4,5,6,7,X,X,X,X	4,5,6,7,X,X,X,X
8	Read	000	0,1,2,3,4,5,6,7	0,1,2,3,4,5,6,7
		001	1,2,3,0,5,6,7,4	1,0,3,2,5,4,7,6
		010	2,3,0,1,6,7,4,5	2,3,0,1,6,7,4,5
		011	3,0,1,2,7,4,5,6	3,2,1,0,7,6,5,4
		100	4,5,6,7,0,1,2,3	4,5,6,7,0,1,2,3
		101	5,6,7,4,1,2,3,0	5,4,7,6,1,0,3,2
		110	6,7,4,5,2,3,0,1	6,7,4,5,2,3,0,1
		111	7,4,5,6,3,0,1,2	7,6,5,4,3,2,1,0
	Write	VVV	0,1,2,3,4,5,6,7	0,1,2,3,4,5,6,7



Read Burst without Additive Latency

113



CL = CAS# Latency

AL = Additive Latency

RCD = RAS#-to-CAS# delay

RL = Read Latency

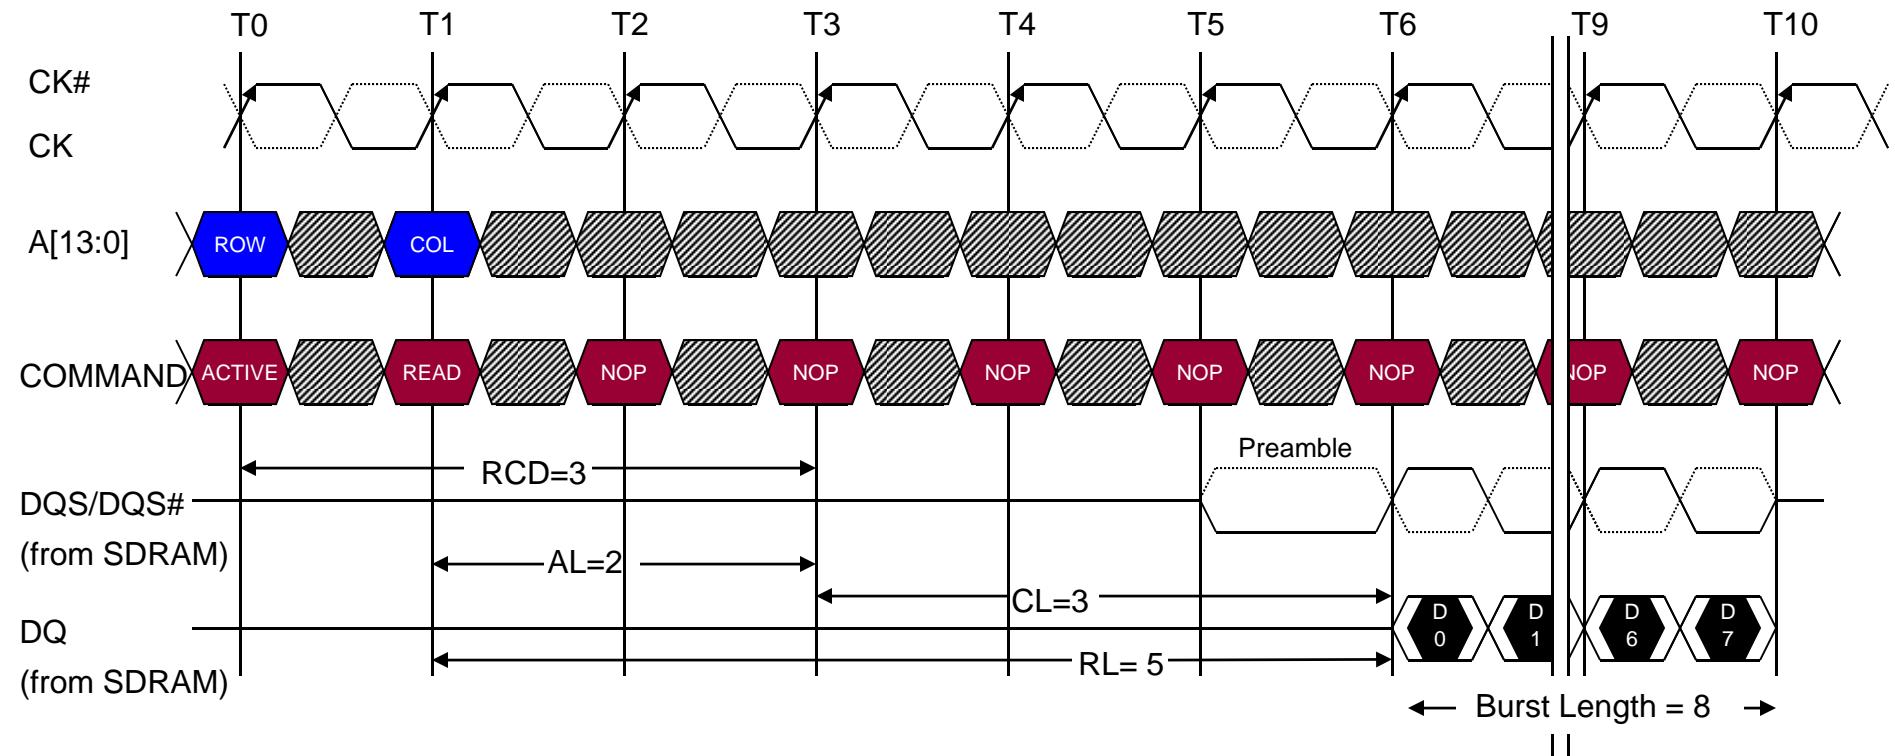
Burst Length = 8

DDR2 400 3 - 3 - 3



Read Burst with Additive Latency

114



CL = CAS# Latency

AL = Additive Latency

RCD = RAS#-to-CAS# delay

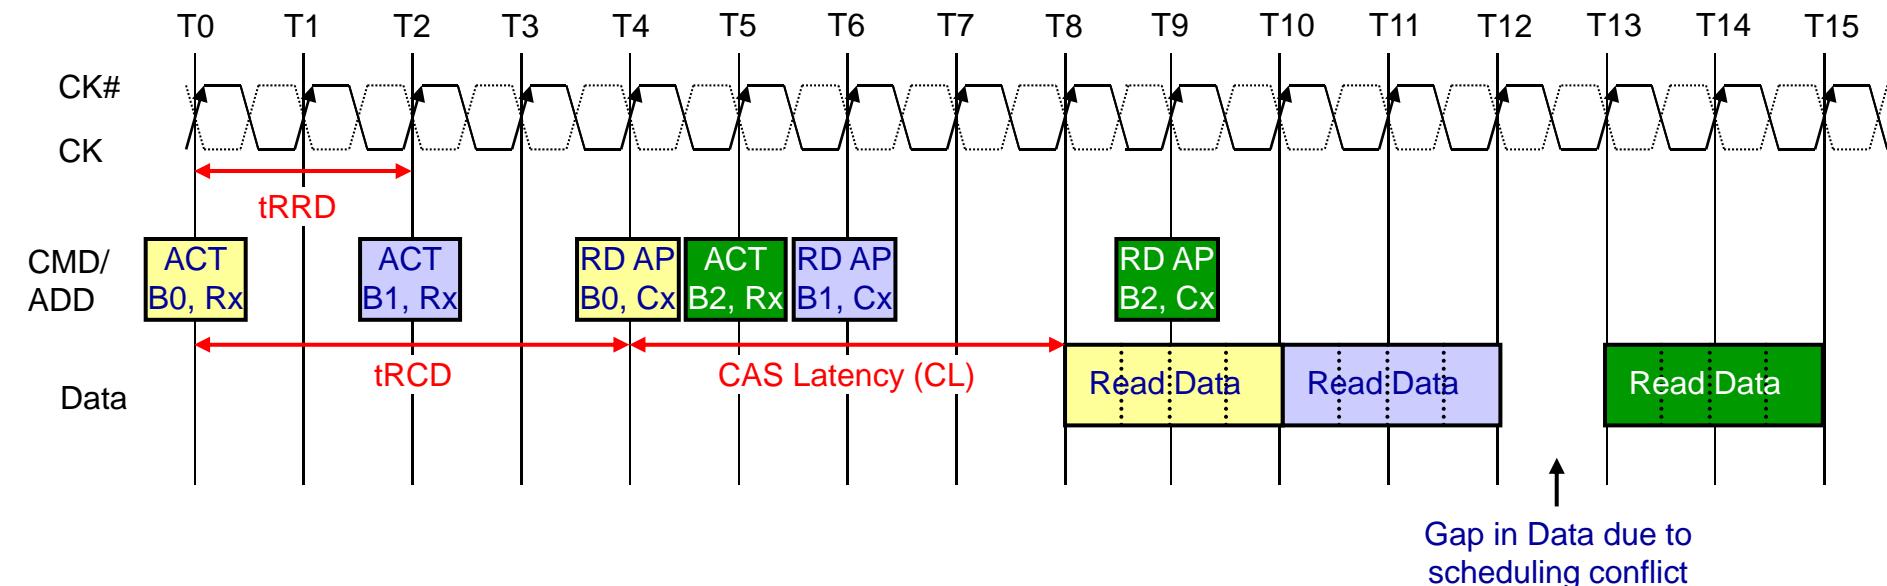
RL = Read Latency

Burst Length = 8



Why Have Additive Latency?

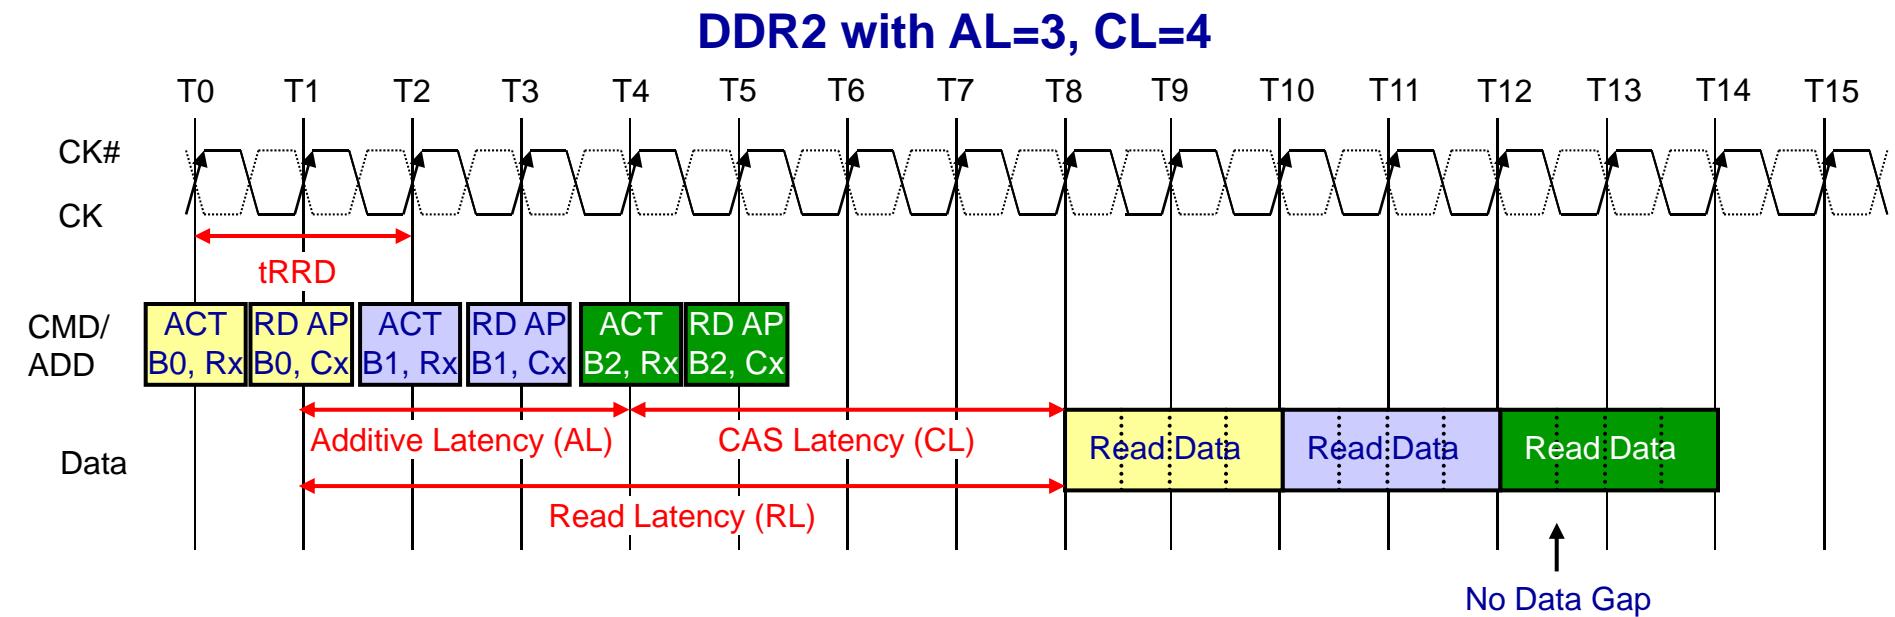
DDR2 with AL=0, CL=4



- The Read Command for Bank0 was sent during cycle T4.
- The memory controller would have liked to send the Activate Bank2 command on cycle T4 instead of T5.
- Due to this delay in activating Bank2, it results in a gap in the data stream being returned by the SDRAM because of this scheduling conflict.



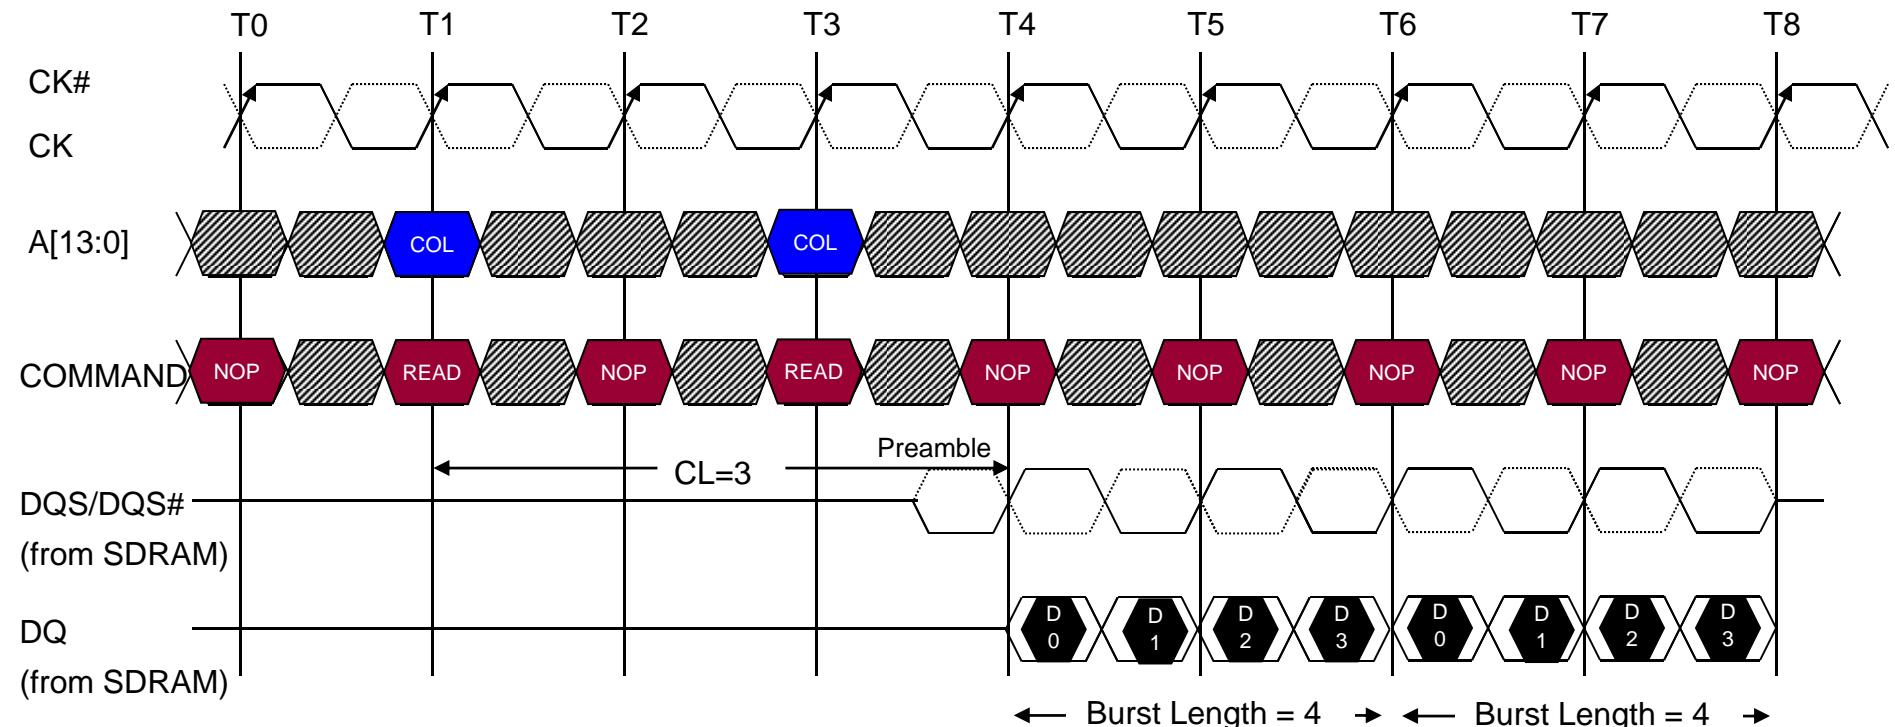
Why Have Additive Latency?



- DDR2 SDRAM can queue commands and schedule them at the appropriate time based on the programmed value of Additive Latency (AL).
- The memory controller no longer has a scheduling conflict because the commands can be sent to SDRAM back-to-back and the SDRAM will schedule them at the appropriate time.
- The data gap seen in the previous example no longer exists.



Read Burst Consecutive



CL = CAS# Latency

AL = Additive Latency

RCD = RAS#-to-CAS# delay

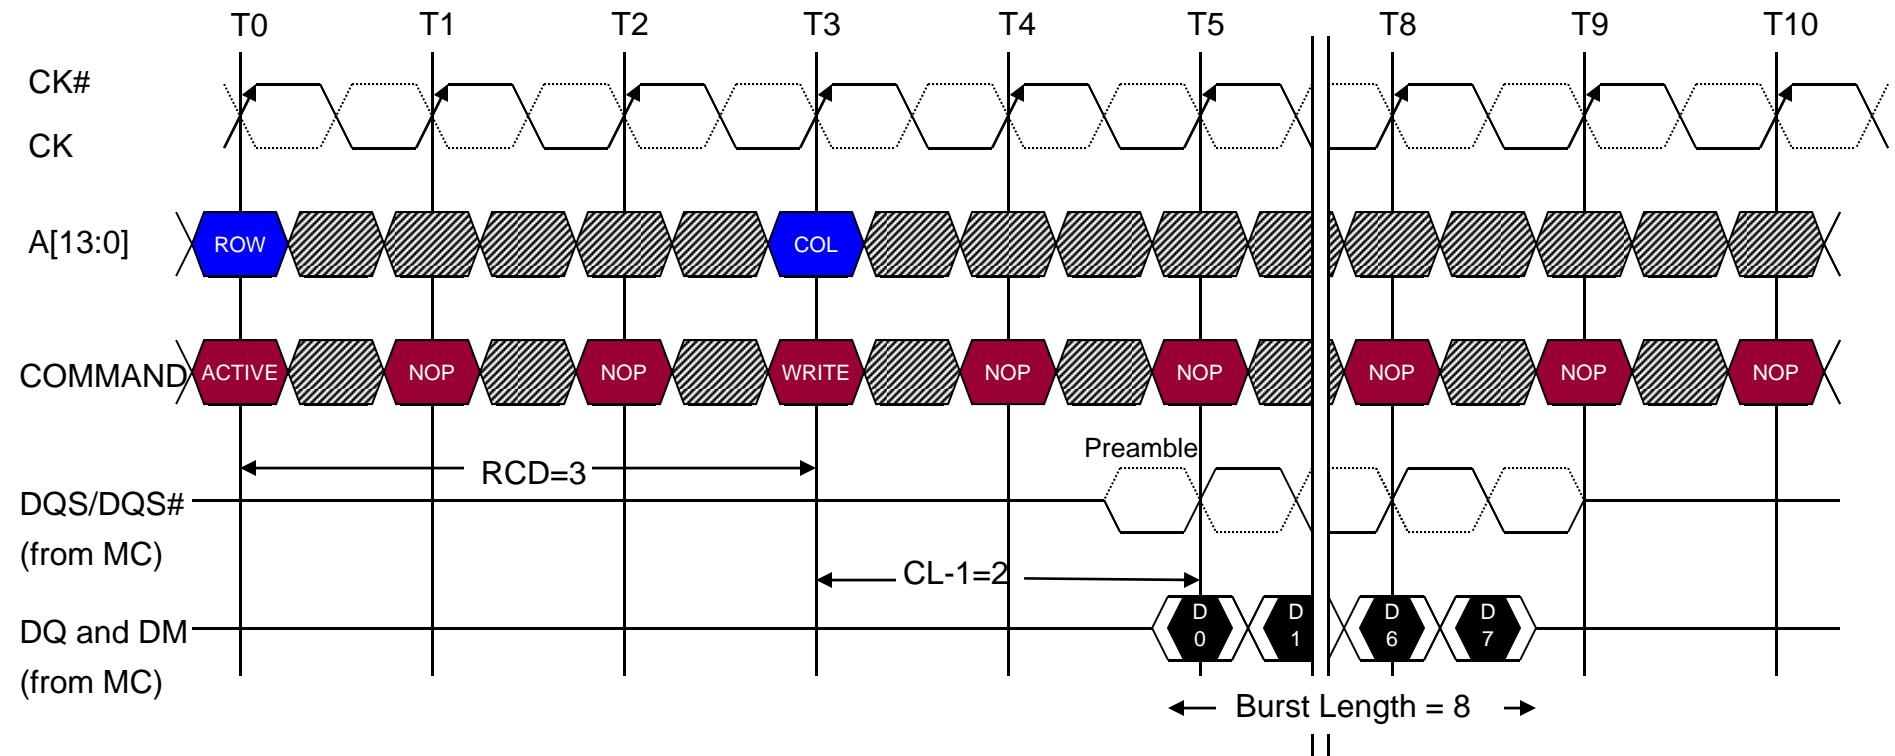
Burst Length = 4



- The Write command is used to issue a burst write to an active row.
- The value of BA0 through BA2 will determine the bank to be accessed. A0-Ax provide the column address.
- The value on A10 determines whether or not Auto Precharge is to be used. If Auto Precharge is selected the row currently accessed will be precharged at the end of the burst.
- Write data appearing on the DQ pins are written to the memory array if DM is registered low. If DM is registered high the data is ignored.



Write Waveform



CL = CAS# Latency

AL = Additive Latency

RCD = RAS#-to-CAS# delay

WL = Write Latency

Burst Length = 8



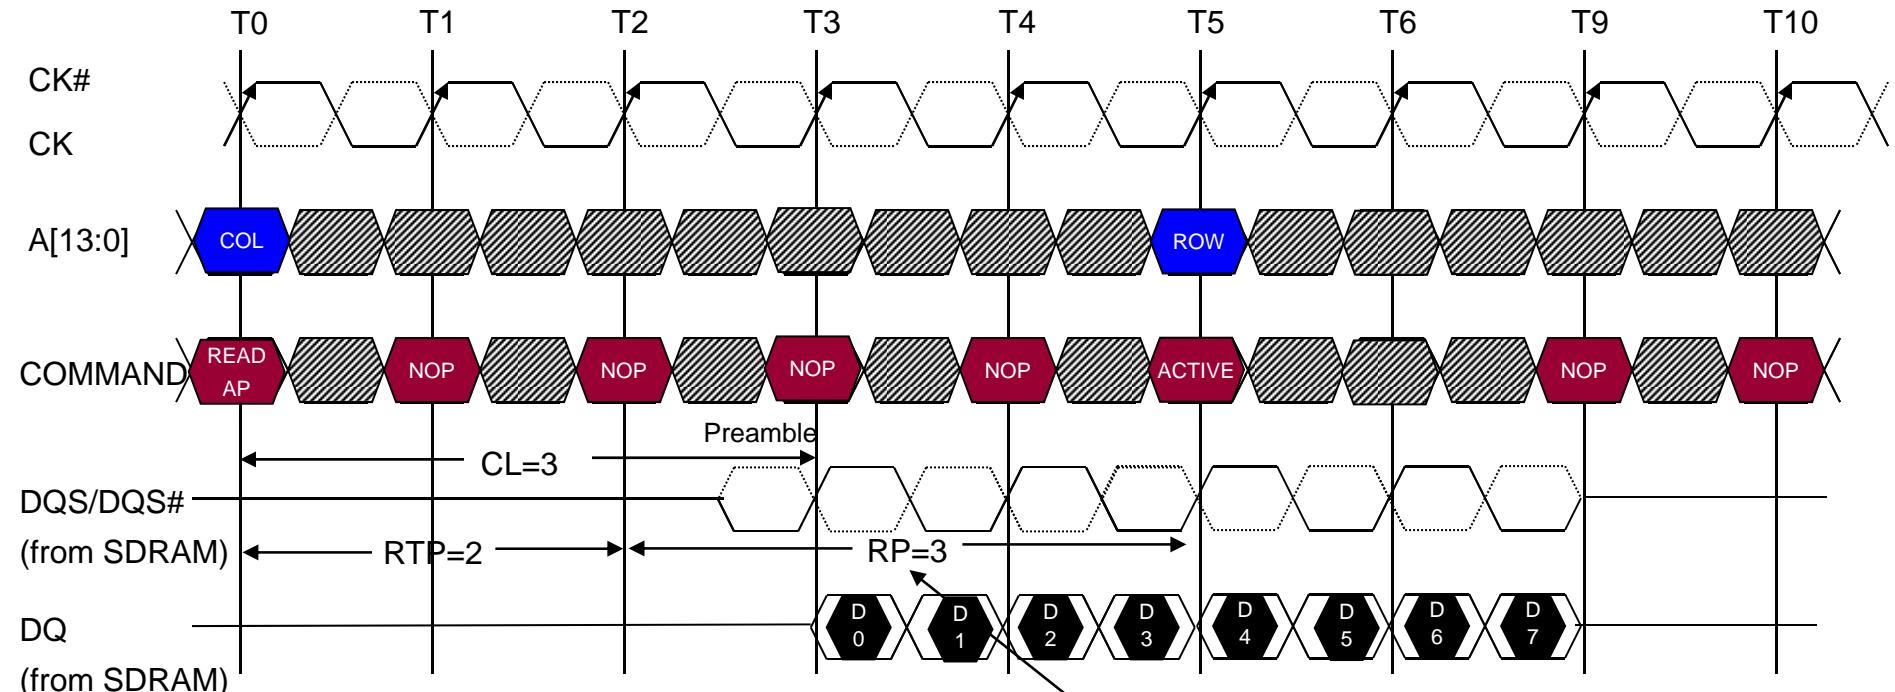
- The Precharge command deactivates the open (active) row in a particular bank, or in all banks.
- The bank(s) will be ready for an Activate command after t_{RP} is met. A10 determines whether all or one bank is to be precharged.
- In the case where only one bank is to be precharged, BA0-BA2 will select the bank. Once a bank is precharged it is in the idle state.
- If no bank is active a Precharge is seen as a NOP.



- The same effect as with a separate Precharge command, but doesn't require a separate command.
- Indicated by asserting A10 during a Read or Write command.
- Auto Precharge ensures that the Precharge is going to happen at the earliest possible time after the Read or Write command completes.



Read with Auto Precharge



t_{RP} is the time required to internally precharge an active Row until the next command

Min Huang (min.huang@leecy.com)
Do Not Distribute



Refresh

Controller told async DRAM chips which row to refresh.

Auto Refresh

SDRAM feature. CAS-before-RAS refresh.

Row counters are inside the SDRAM chips.

Controller only tells chips when to refresh, not which row.

Simply called “refresh” now.

Self Refresh

System powered off; DRAM fully powered.

Refresh interval timer inside SDRAM chips.

Interval is fixed before starting self refresh.

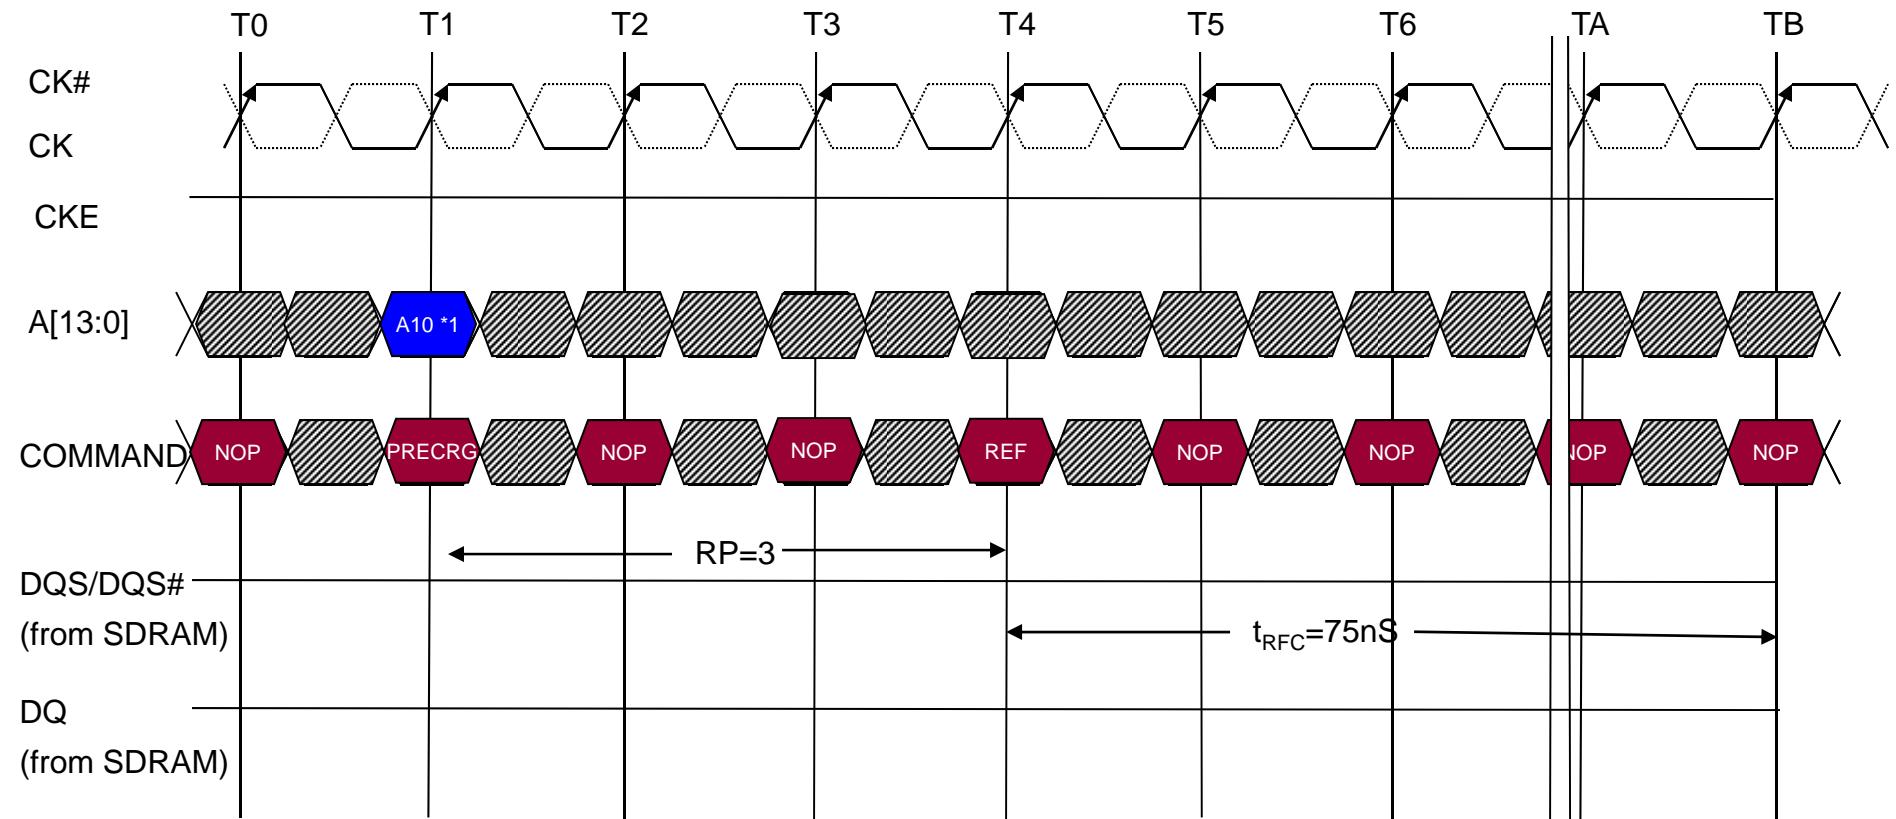
Auto Self Refresh

Self refresh where DDR3 chips automatically choose their refresh interval based on their temperature.

- Refresh is used in normal operation and is the same as CAS-before-RAS (CBR) refresh in older DRAM.
- This command is non-persistent and must be issued every time a refresh is required. No address is required.
- The address is internally generated.
- Refresh can be posted up to 8 times. This means the max time can be as much as $9 \times t_{REFI}$. That is the Refresh Interval (t_{REFI}) \times 8 plus the current Refresh command (giving 9).
- No more than 16 Refresh commands may be issued within $2 \times t_{REFI}$.
- The minimum time between Refresh commands is Refresh Cycle Time (t_{RFC}).

- The refresh period is 64ms (32ms for high temp) which equals one refresh every 7.8us (t_{REFI}) for a device with 8192 rows.
- A single refresh command may refresh one or more Rows. For example here are the refresh requirements for different densities.
- t_{RFC} Refresh Cycle Time is how long it takes for the DRAM to finish the refresh command.

Density	DDR2 t_{RFC}	Estimated number of Rows Refreshed, depending on organization
256Mb	75ns	1 (8192 Rows)
512Mb	75-105ns	1-2 (8192,16384 Rows)
1Gb	105-127.5ns	1-2 (8192,16384 Rows)
2Gb	127.5-197.5ns	2-4 (16384, 32768 Rows)



*1 A10 must be high for more than 1 bank to get Precharged. Precharge all must be done before entering the Refresh mode.

t_{RFC} is the refresh cycle time

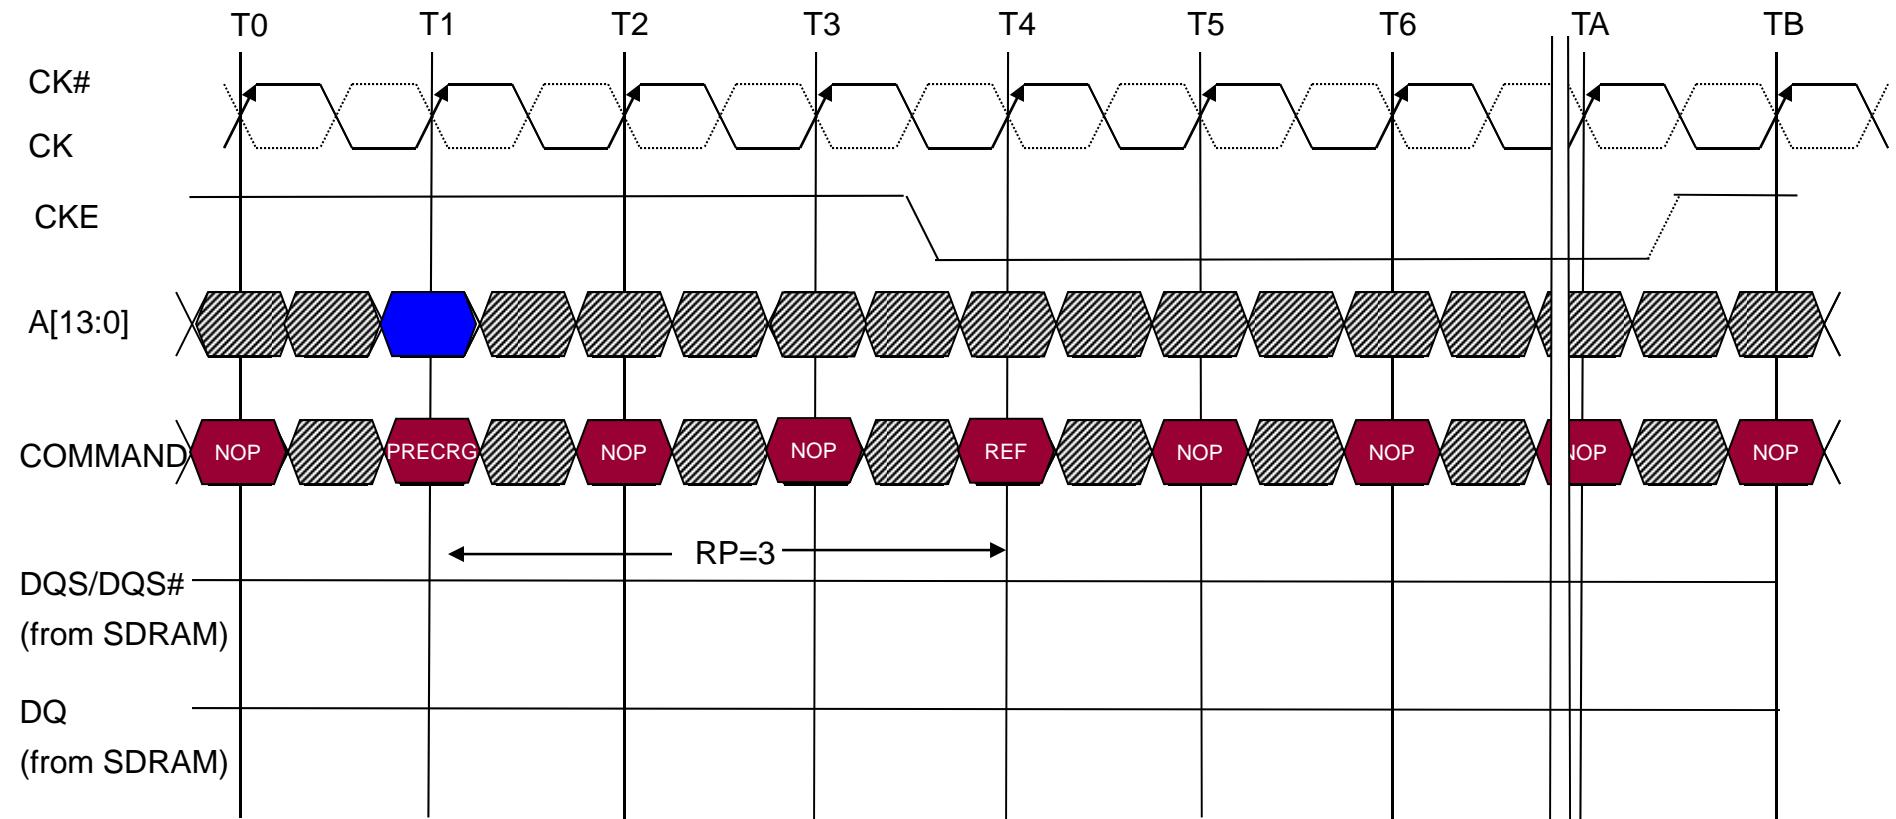


- Self Refresh is the same command as Refresh except CKE is low. This command is used to keep data integrity while the rest of the system is powered down. The DLL is automatically disabled going into self refresh and automatically enabled coming out of a Self Refresh state. CKE must stay low during Self Refresh.
- Here is the procedure for coming out of Self Refresh:
 - CK and CK# must be stable before CKE goes high.
 - CKE goes high
 - NOP command issued for tXSNR (Exit Self Refresh to a Non Read Command) because time is required for internal refreshes to complete
 - Refresh command is recommended
 - No DLL reset is required



Self Refresh Waveform

128



All banks must be precharged before entering the Self Refresh mode.

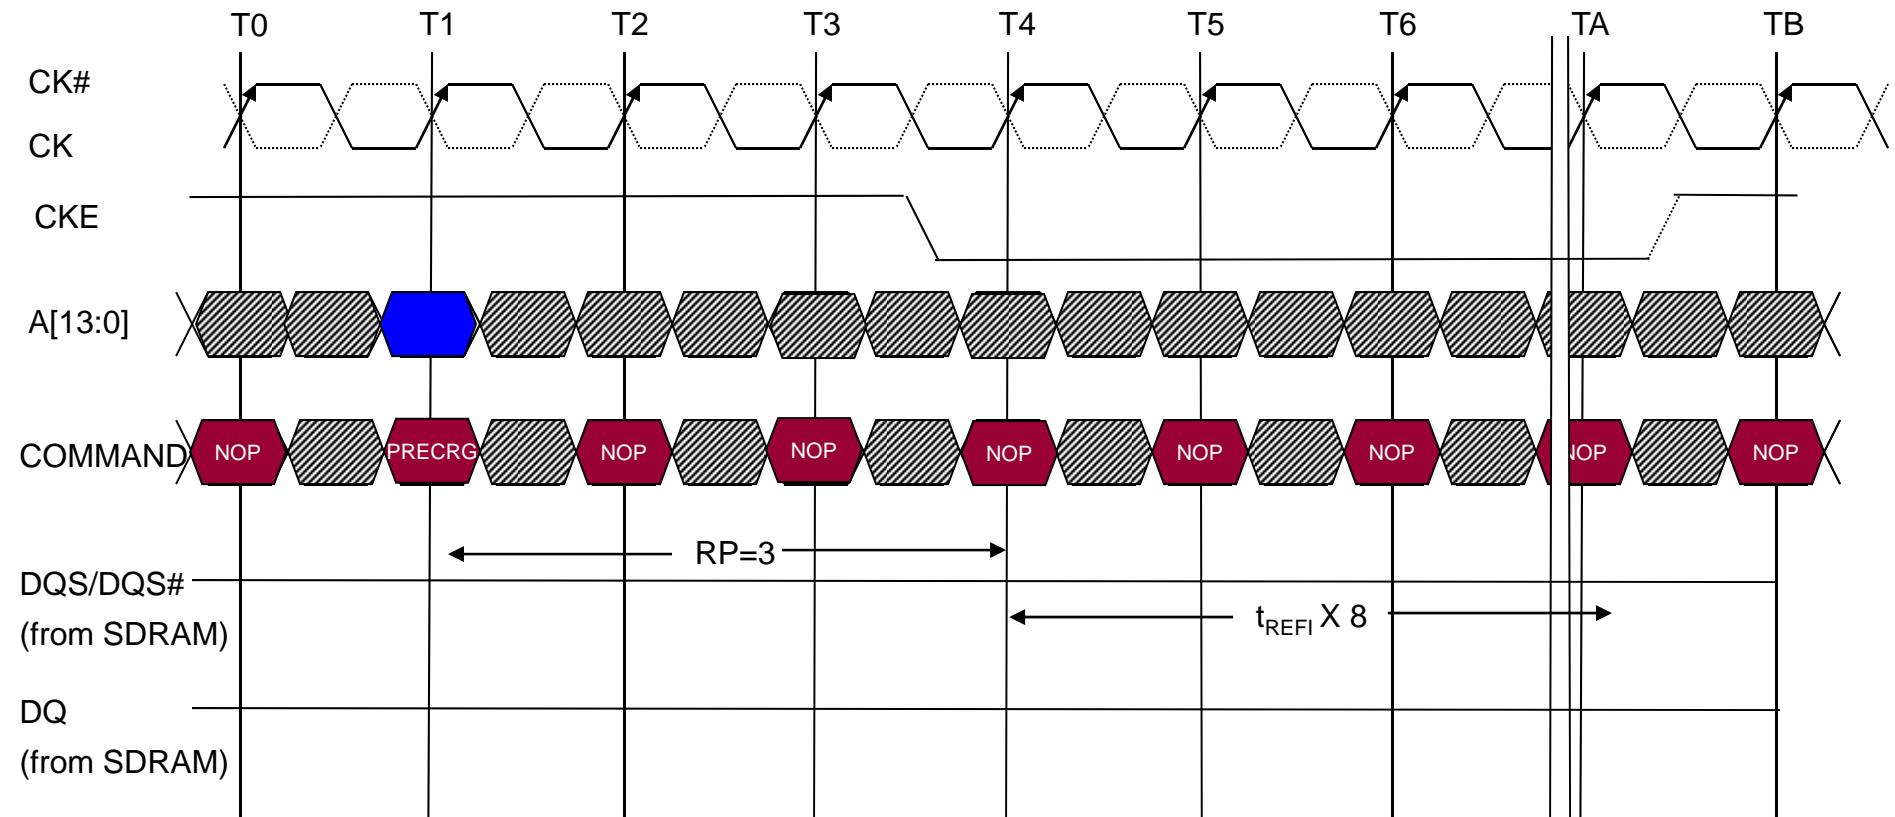
Refresh with CKE low will cause the DRAM to go into a Self Refresh state.

t_{RP} is the time required to internally precharge an active Row until the next command

- No Operation - NOP command is used to perform a NOP to the selected DRAM (DRAM whose CS# = low). Operations in progress are not affected.
- Deselect - Deselect (CS# = high) prevents new commands from being executed. DRAM is effectively disabled.
 - Prevents unwanted commands from being registered in an idle state.



- Power Down – Also called CKE power down. CKE is brought low synchronously with CS and the other commands high. This low power state can only be held for $8 \times t_{REFI}$ of the device because no refreshes can happen in this state. The $8 \times t_{REFI}$ is due to the fact that refresh can be posted up to 8 times.
- Clock must be active in this state.
- Fast exit refers to leaving the DLL on and Slow Exit turns the DLL off. If Slow exit is used the DLL must relock.



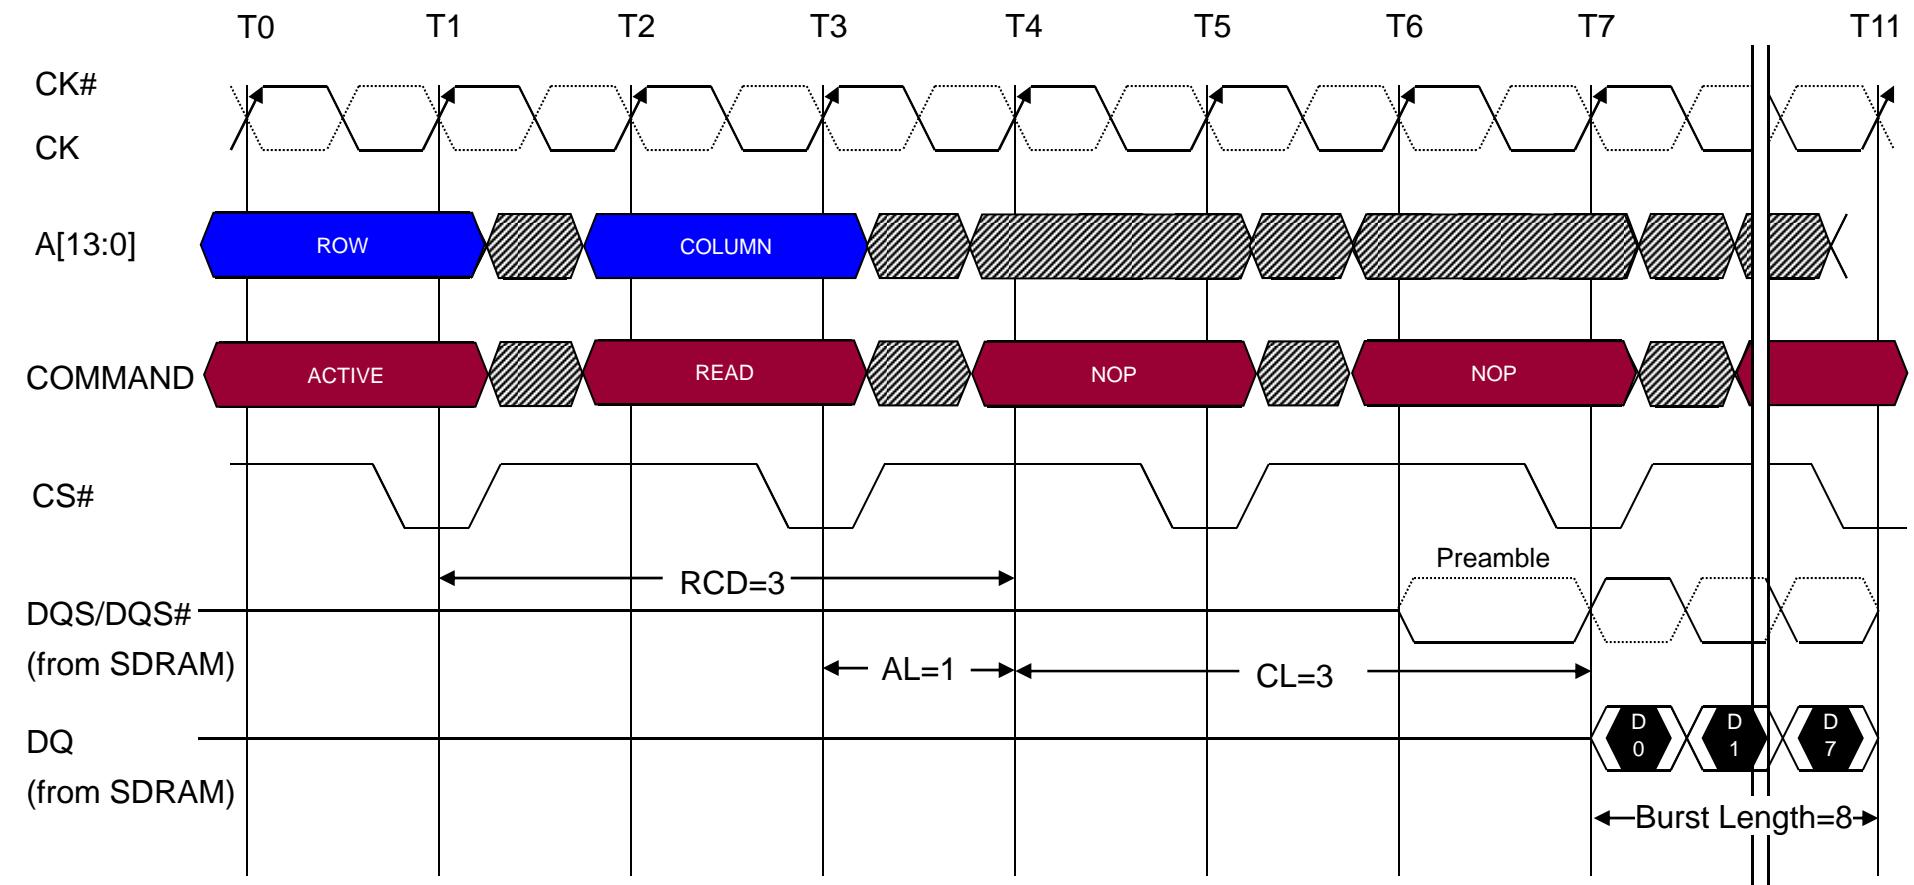
Precharge Power Down shown. If any banks are left open, this becomes Active Power Down.

t_{REFI} is Refresh interval time times the 8 posted refreshes until exit of Power Down is required.



- 1T and 2T timings
 - Depending on how many DIMMs the system needs to support, the DRAM controller may use one of 2 different address and command timing schemes.
 - 1T: The address and command signals are held active by the controller for 1 clock. This allows for faster turnaround times.
 - 2T: The address and command signals are held active by the controller for 2 clocks. For systems with more loads, this allows for longer setup and hold times. Control signals (CS#, CKE, ODT) must still obey 1T timing.

2T Timing Example





DDR Initialization

Min Huang(min.huang@ lecroy.com)
Do Not Distribute



- DDR SDRAMs must be powered up and initialized in a very specific sequence to ensure stable working parts.
- Every time power is lost the procedure must be repeated.
- The initialization code is typically held in firmware and is often referred to as JEDEC initialization.



➤ Step 1

- Apply power. Keep CKE below $0.2 \times \text{VDDQ}$ and ODT LOW. All other inputs may be undefined.
- Refer to the JESD79-2C standard for details about the allowable voltage ramp times.



➤ Step 2

- Start clock and maintain stable condition with CKE held low.

➤ Step 3

- For the minimum of 200 us after stable power and stable clock (CK, CK#), then apply NOP or Deselect & take CKE HIGH.



➤ Step 4

- Wait minimum of 400 ns then issue precharge all command. NOP or Deselect applied during 400 ns period.



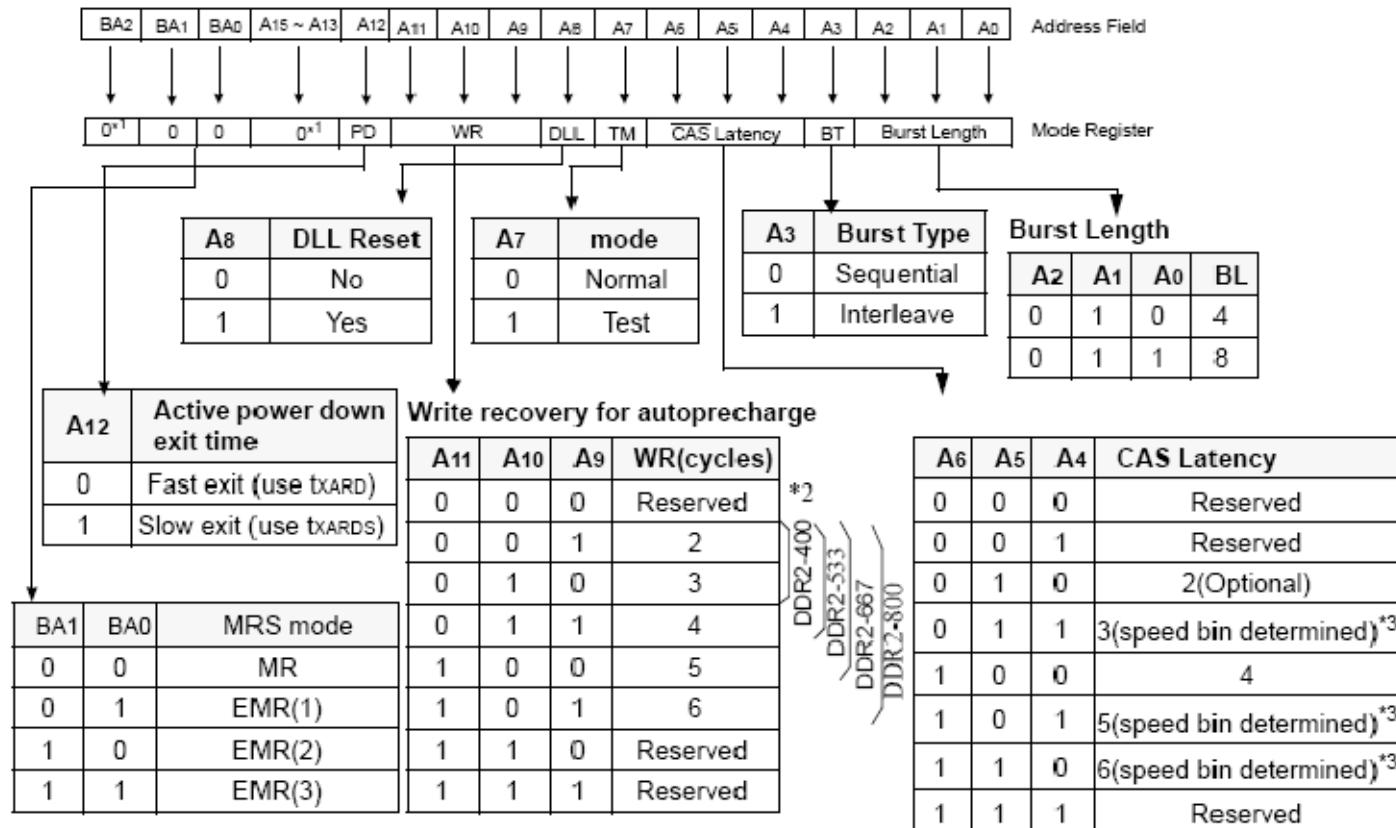
- Step 5
 - Issue MRS command to EMR2.
- Step 6
 - Issue MRS command to EMR3.
- Let's see what's in the mode registers and also see how we learn what values to program into the mode registers.



- Mode Register (MR) & Extended Mode Register (EMR)
- The mode registers are used to define the specific mode of operation of the SDRAM
 - The default value for these registers is not defined.
 - Mode registers will retain all information until rewritten or power is removed from the device except bit A8 (DLL reset) which is self clearing
 - Reprogramming the mode registers during operation will not alter the information held in the DRAM
 - The Bank Address bits are used to determine which mode register is being written to.
 - The Address bus is used to write to each of the registers.
 - Why is the Address bus used to program Mode Registers??



Step by Step DDR2 Initialization MR0



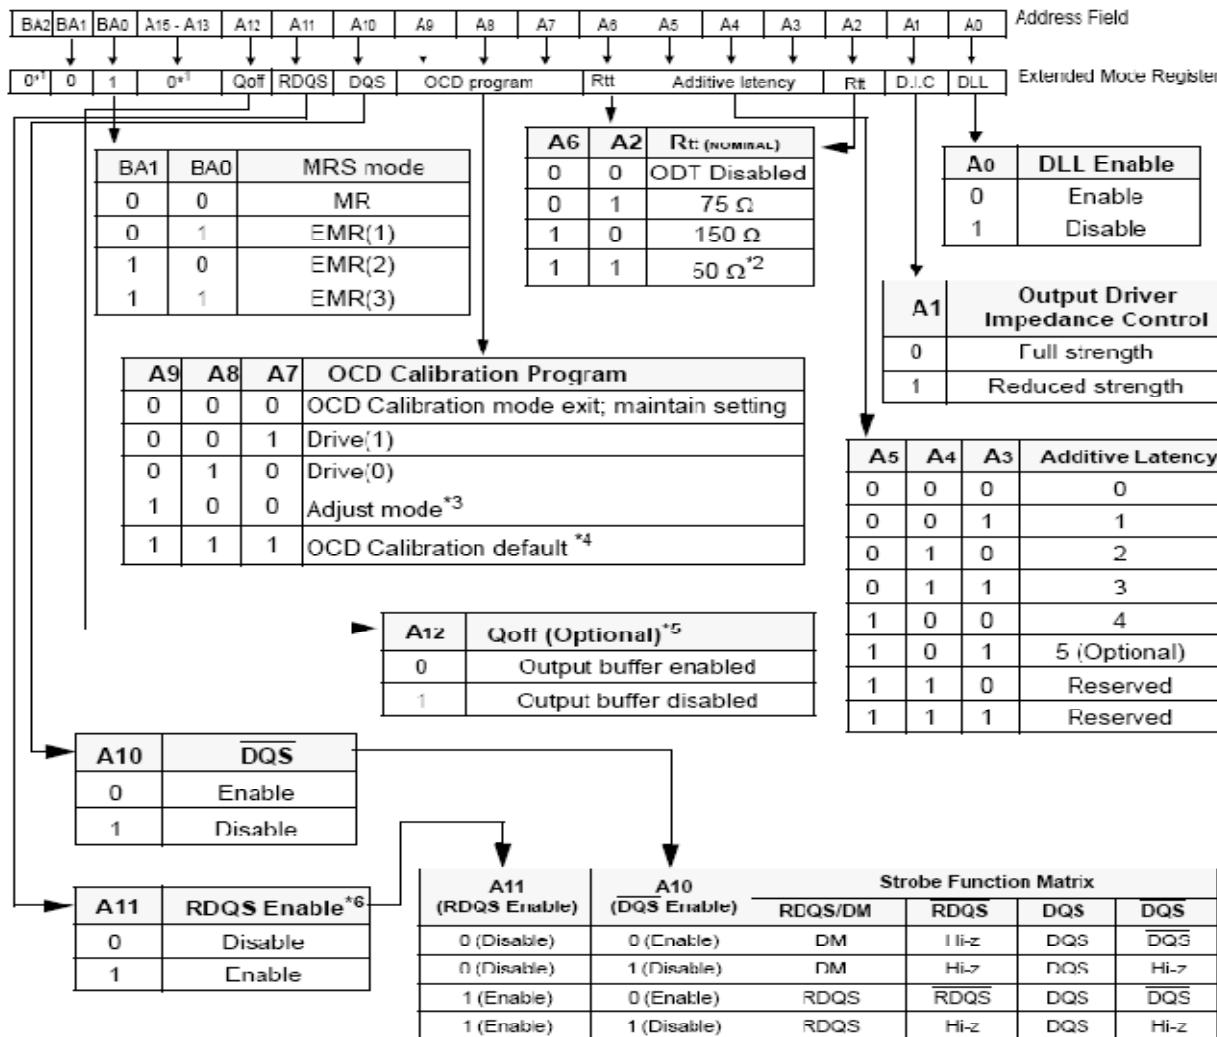
NOTE 1 BA2 and A13-A15 are reserved for future use and must be set to 0 when programming the MR.

NOTE 2 For DDR2-400/533, WR (write recovery for autorecharge) min is determined by tCK max and WR max is determined by tCK min. WR in clock cycles is calculated by dividing tWR (in ns) by tCK (in ns) and rounding up to the next integer ($WR[\text{cycles}] = RU\{ tWR[\text{ns}] / tCK[\text{ns}] \}$, where RU stands for round up). For DDR2-667/800, WR min is determined by tCK(avg) max and WR max is determined by tCK(avg) min. WR[cycles] = RU { tWR[ns] / tCK(avg)[ns] }, where RU stands for round up. The mode register must be programmed to this value. This is also used with tRP to determine tDAL.

NOTE 3 Speed bin determined. Not required on all speed bins.



Step by Step DDR2 Initialization EMR1



NOTE 1 BA2 and A13-A15 are reserved for future use and must be set to 0 when programming the EMR(1).

NOTE 2 Optional for DDR2-400/533/667, mandatory for DDR2-800.

NOTE 3 When Adjust mode is issued, AL from previously set value must be applied.

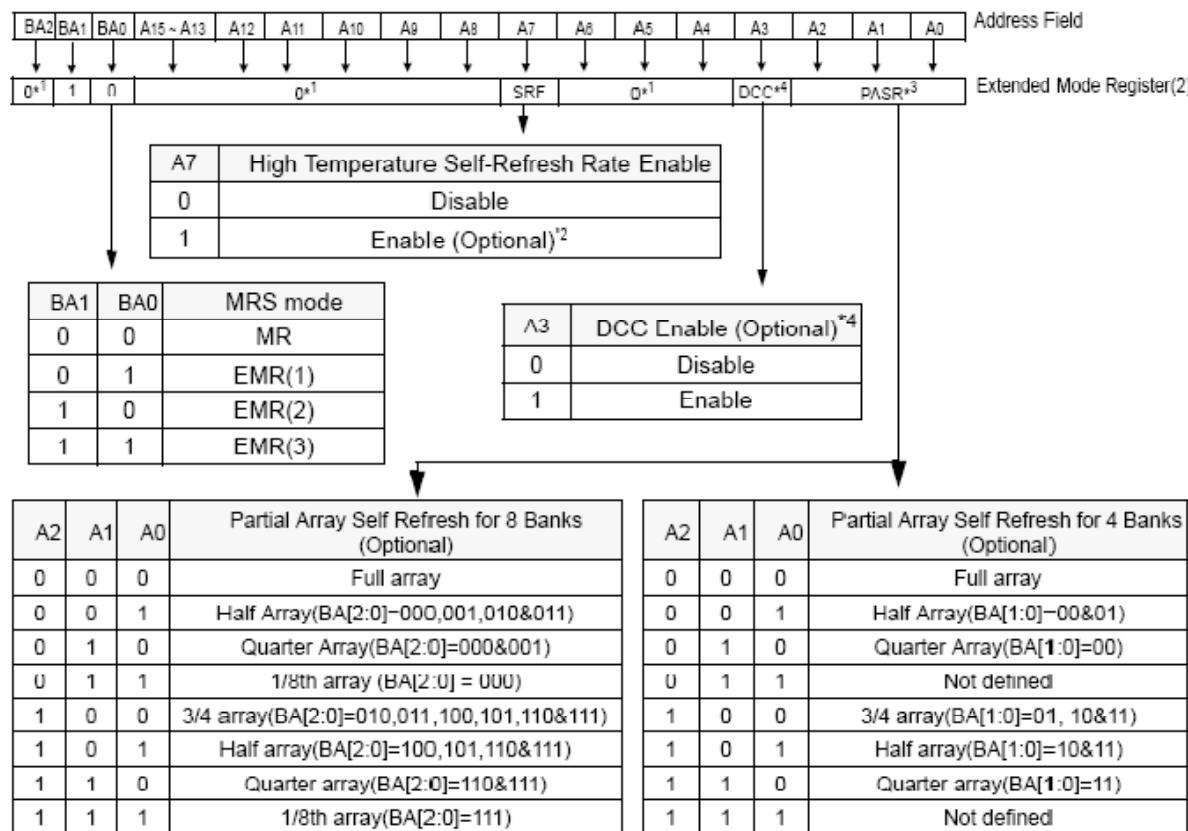
NOTE 4 After setting to default, OCD calibration mode needs to be exited by setting A9-A7 to 000. Refer to section 2.4.3 for detailed information.

NOTE 5 Output disabled - DQs, DQ9s, DQ3s, RDQS, RDQS. This feature is used in conjunction with DIMM IDD measurements when IDDQ is not desired to be included.

NOTE 6 If RDQS is enabled, the DM function is disabled. RDQS is active for reads and don't care for writes.



Step by Step DDR2 Initialization EMR2



NOTE 1 BA2 and A4-A6, A8-A15 are reserved for future use and must be set to 0 when programming the EMR(2).

NOTE 2 As industry adoption of high temperature parts proceeds, users need to determine if a DRAM supports High Temperature Self-Refresh Rate Enable mode before attempting to use it in that mode. JEDEC standard DDR2 SDRAM Module user can look at DDR2 SDRAM Module SPD field Byte 49 bit [0]. If the high temperature self-refresh mode is supported then controller can set the EMR(2)[A7] bit to enable the self-refresh rate in case of higher than 85 °C temperature self-refresh operation. For the loose part user, please refer to DRAM Manufacturer's part number and data sheet to check the high temperature self-refresh rate availability.

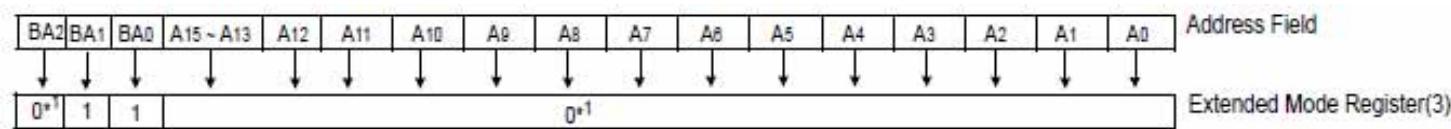
NOTE 3 Optional in DDR2 SDRAM. If PASR (Partial Array Self Refresh) is enabled, data located in areas of the array beyond the specified address range will be lost if self refresh is entered. Data integrity will be maintained if tRFF conditions are met and no Self Refresh command is issued. If the PASR feature is not supported, EMR(2)[A0-A2] must be set to 000 when programming EMR(2).

NOTE 4 Optional in DDR2 SDRAM. JEDEC standard DDR2 SDRAM may or may not have DCC (Duty Cycle Corrector) implemented, and in some of the DRAMs implementing DCC, user may be given the controllability of DCC thru EMR(2)[A3] bit. JEDEC standard DDR2 SDRAM users can look at manufacturer's data sheet to check if the DRAM part supports DCC controllability. If optional DCC Controllability is supported, user may enable or disable the DCC by programming EMR(2)[A3] accordingly. If the controllability feature is not supported, EMR(2)[A3] must be set to 0 when programming EMR(2).



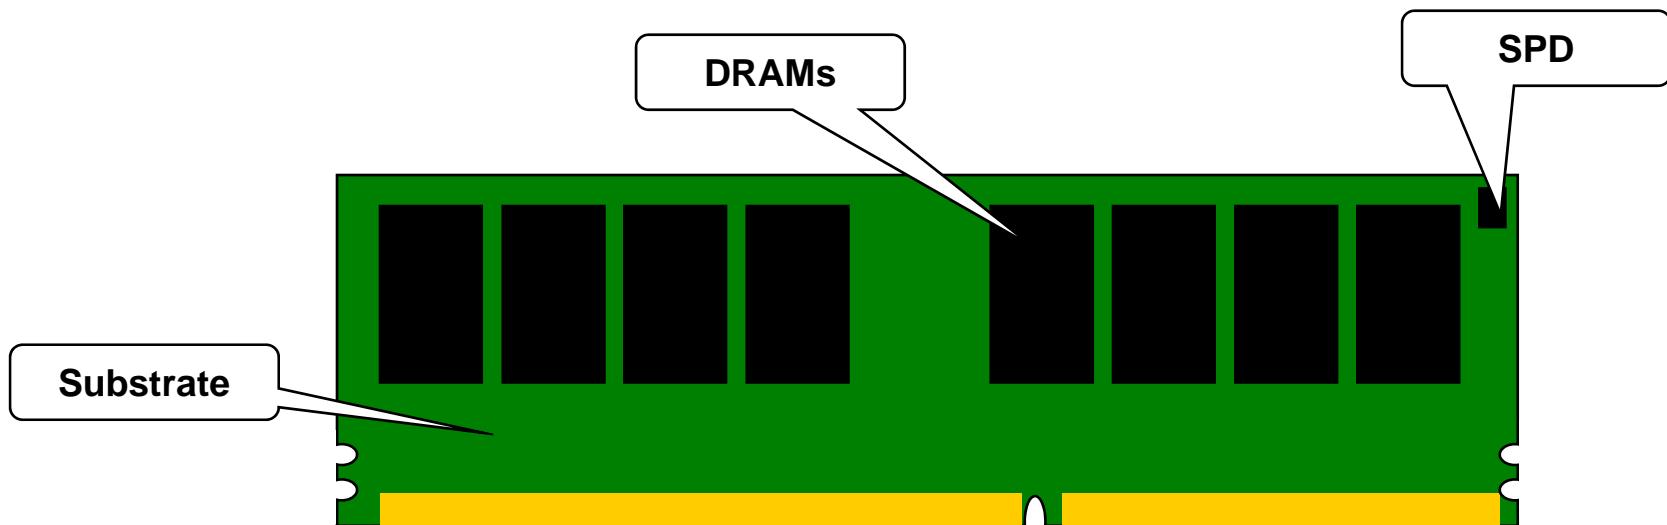
Step by Step DDR2 Initialization EMR3

144





The information required from each DIMM is held in the SPD.





- Prior to JEDEC initialization, the BIOS/Firmware has read all data from SPD and worst-case timings have been calculated. Example:
 - If two DIMMs are installed where one has a CAS latency of 2 and one with a CAS latency of 2.5, then CAS latency 2.5 will be used for both DIMMs.
- A DLL reset will be performed during this step. Anytime DLL reset occurs 200uS must be provided for it to relock before issuing a read command.



This is an example of the data stored in the SPD.

- The first 128 bytes are defined by the memory manufacturer.
- The last 128 bytes can be used by the customer.
- Each SPD has a unique SMBus address hardwired on the motherboard typically on the lower 3 address bits of the SMBus.
- The SMBus can attach to the MCH or the ICH.

BYTE	TITLE	VALUE
0	NUMBER OF BYTES USED BY MICRON	80
1	TOTAL NUMBER OF SPD MEMORY BYTES	08
2	MEMORY TYPE	07
3	NUMBER OF ROW ADDRESSES	0D
4	NUMBER OF COLUMN ADDRESSES	0C
5	NUMBER OF MODULE RANKS	02
6	MODULE DATA WIDTH	48
7	MODULE DATA WIDTH (CONTINUED)	00
8	MODULE VOLTAGE INTERFACE LEVELS	04
9	DDR SDRAM CYCLE TIME (CAS LATENCY = 2.5)	75
10	DDR SDRAM ACCESS FROM CLOCK (CAS LATENCY = 2.5)	75
11	MODULE ERROR CORRECTION CONFIGURATION TYPE	02
12	MODULE REFRESH RATE AND TYPE	82
13	SDRAM DEVICE WIDTH	04
14	ERROR CHECKING WIDTH	04
15	MIN CLOCK DELAY FOR BACK-TO-BACK RANDOM COLUMN ADDR	01
16	BURST LENGTHS SUPPORTED	0E
17	NUMBER OF BANKS INTERNAL TO DISCRETE SDRAM DEVICES	04
18	CAS LATENCIES SUPPORTED	0C
19	CS LATENCY	01
20	WE LATENCY	02
21	SDRAM MODULE ATTRIBUTES	26
22	SDRAM DEVICE ATTRIBUTE: GENERAL	C0
23	DDR SDRAM CYCLE TIME (TCK) AT CL = 2	A0
24	DDR SDRAM ACCESS FROM CLOCK (TAC) AT CL = 2	75
25	DDS SDRAM CYCLE TIME (TCK) AT CL = 1	00
26	DDR SDRAM ACCESS TIME FROM CLOCK (TAC) AT CL = 1	00
27	SDRAM: MINIMUM ROW PRECHARGE TIME (TRP)	50
28	MINIMUM ROW ACTIVE TO ROW ACTIVE	3C
29	MINIMUM RAS TO CAS DELAY	50



BYTE	TITLE	VALUE
30	MINIMUM RAS PULSE WIDTH	2D
31	MODULE RANK DENSITY	01
32	COMMAND/ADDRESS SETUP	A0
33	COMMAND/ADDRESS HOLD	A0
34	DATA SIGNAL INPUT SETUP	50
35	DATA SIGNAL INPUT HOLD	50
36-40	RESERVED 5	0000000000
41	DEVICE MINIMUM ACTIVE/AUTO-REFRESH TIME (TRC)	41
42	DEVICE MINIMUM AUTO-REFRESH TO ACTIVE/AUTO-REFRESH	4B
43	DEVICE MAXIMUM DEVICE CYCLE TIME (TCK MAX)	34
44	DEVICE DQS-DQ SKEW FOR DQS AND ASSOCIATED DQ SIGNAL	32
45	DEVICE READ DATA HOLD SKEW FACTOR (TQHS)	75
46	RESERVED (BYTE 46)	00
47	DIMM HEIGHT	01
48-61	RESERVED BYTES 48-61	0000...00
62	SPD REVISION	10
63	CHECKSUM FOR BYTES 0 THRU 62	EB
64	MANUFACTURER'S JEDEC ID CODE	2C
65-71	MANUFACTURER'S JEDEC ID CODE (CONTINUED)	FFFFFFFFFFFF
72	MANUFACTURING LOCATION	00
73-90	MODULE PART NUMBER	36VDDF25672G265C2
91	PCB IDENTIFICATION CODE	02
92	PCB IDENTIFICATION CODE (CONTINUED)	00
93	YEAR OF MANUFACTURE	00
94	YEAR OF MANUFACTURE	00
95-98	MODULE SERIAL NUMBER	00000000
99-127	MANUFACTURER SPECIFIC DATA (RSVD) 99-127	0000...00
128-191	UNUSED	FFFF...FF
192-255	UNUSED2	FFFF...FF



- Step 7
 - Issue MRS command to EMR1 to enable DLL.
- Step 8
 - Issue MRS command to MR0 to reset DLL.



- Step 9
 - Issue a Precharge All command.
- Step 10
 - Issue 2 or more Refresh commands.
- Step 11
 - Issue a MRS command to MR0 with LOW to A8 to program the desired device operation without resetting the DLL.



➤ Step 12

➤ At least 200 clocks after resetting the DLL, execute OCD Calibration (Off Chip Driver impedance adjustment). If OCD calibration is not used, issue a MRS command to EMR1 to set OCD Calibration Default followed by issuing a MRS command to EMR1 to exit OCD Calibration Mode while also setting other operating parameters of EMR1.

➤ Done!

➤ The standard says the DRAM is now ready for normal operation, but the controller still needs to train the timing.



➤ Step 1

- Apply power. Assert RESET# for at least 200 us with stable power. RESET# is recommended to be less than $0.2 \times VDD$. All other inputs may be undefined. Pull CKE Low at least 10 ns anytime before deasserting RESET#. The VDD ramp time between 300 mv to VDDmin must be no more than 200 ms. During the ramp, $VDD > VDDQ$ and $(VDD - VDDQ) < 0.3$ volts.
- The voltage levels on all pins other than VDD, VDDQ, VSS, VSSQ must be less than or equal to VDDQ and VDD on one side and must be larger than or equal to VSSQ and VSS on the other side.
- VDD and VDDQ are driven from a single power converter output, AND
- VTT is limited to 0.95 V max once power ramp is finished, AND
- Vref tracks VDDQ/2.

OR

- Apply VDD, without any slope reversal, with or before VDDQ.
- Apply VDDQ, without any slope reversal, with or before VTT & Vref.



➤ Step 2

- After RESET# is de-asserted, wait 500 us until CKE becomes active. During this time, the DRAM will start internal state initialization independently of external clocks.



➤ Step 3

- Start and stabilize CK and CK# for at least 10 ns or 5 tCK (whichever is larger) before CKE goes active. Since CKE is a synchronous signal, the corresponding set up time to clock (tIS) must be met. Also, a NOP or Deselect command must be registered (with tIS set up time to clock) before CKE goes active. Once the CKE is registered “High” after RESET#, CKE needs to be continuously registered “High” until the initialization sequence is finished, including expiration of tDLLK and tZQinit.



➤ Step 4

- The DDR3 SDRAM keeps its on-die termination (ODT) in high-impedance state as long as RESET# is asserted. Further, the SDRAM keeps its ODT in high impedance state after RESET# deassertion until CKE is registered “High”. The ODT input signal may be in undefined state until tIS before CKE is registered “High”. When CKE is registered “High”, the ODT input signal may be statically held at either “Low” or “High”. If RTT_NOM is to be enabled in MR1, the ODT input signal must be statically held “Low”. In all cases, the ODT input signal remains static until the power up initialization sequence is finished, including the expiration of tDLLK and tZQinit.



➤ Step 5

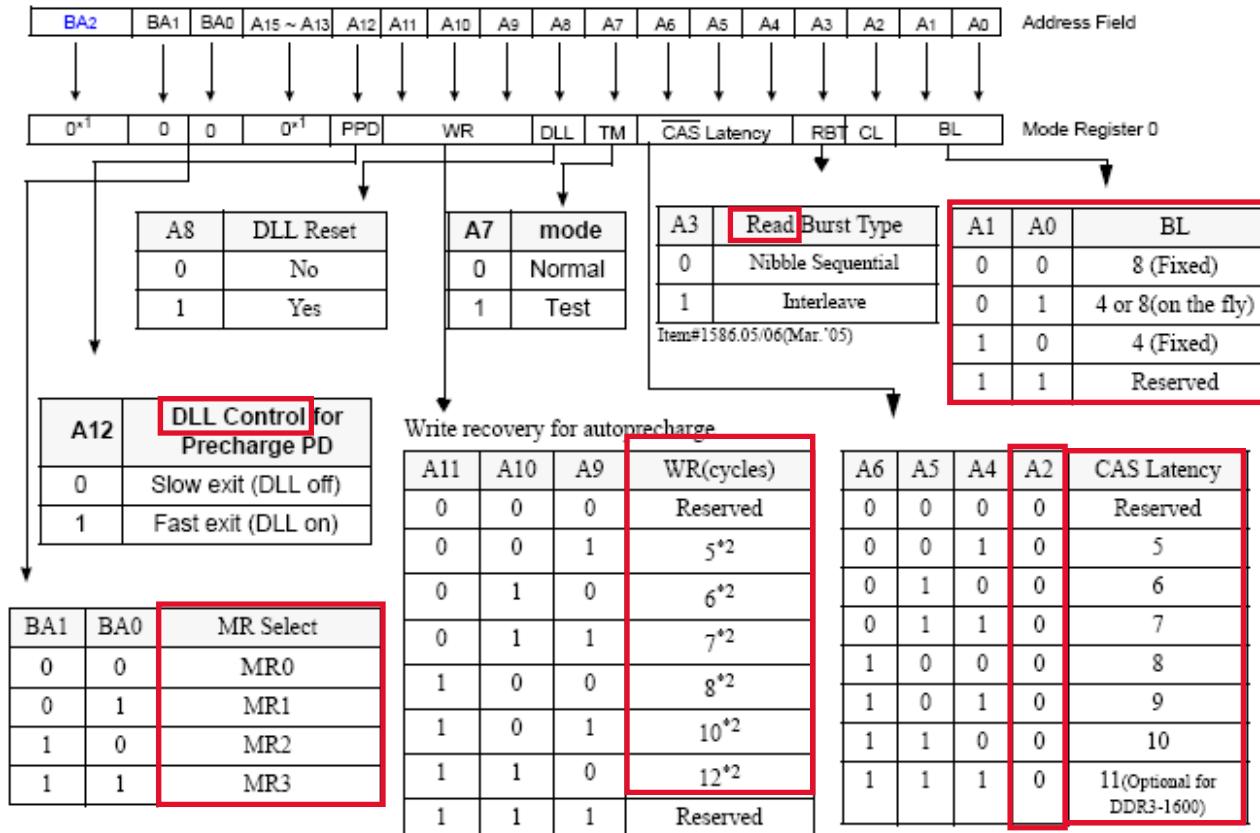
- After CKE is registered “High”, wait minimum of Reset CKE Exit time, tXPR, before issuing the first MRS command to load mode register. (tXPR=max (tXS; 5 x tCK))



- DDR3 Mode Registers are numbered consistently to the addressing:
 - MR0: Mode Register 0
 - MR1: Mode Register 1
 - MR2: Mode Register 2
 - MR3: Mode Register 3



Mode Register 0

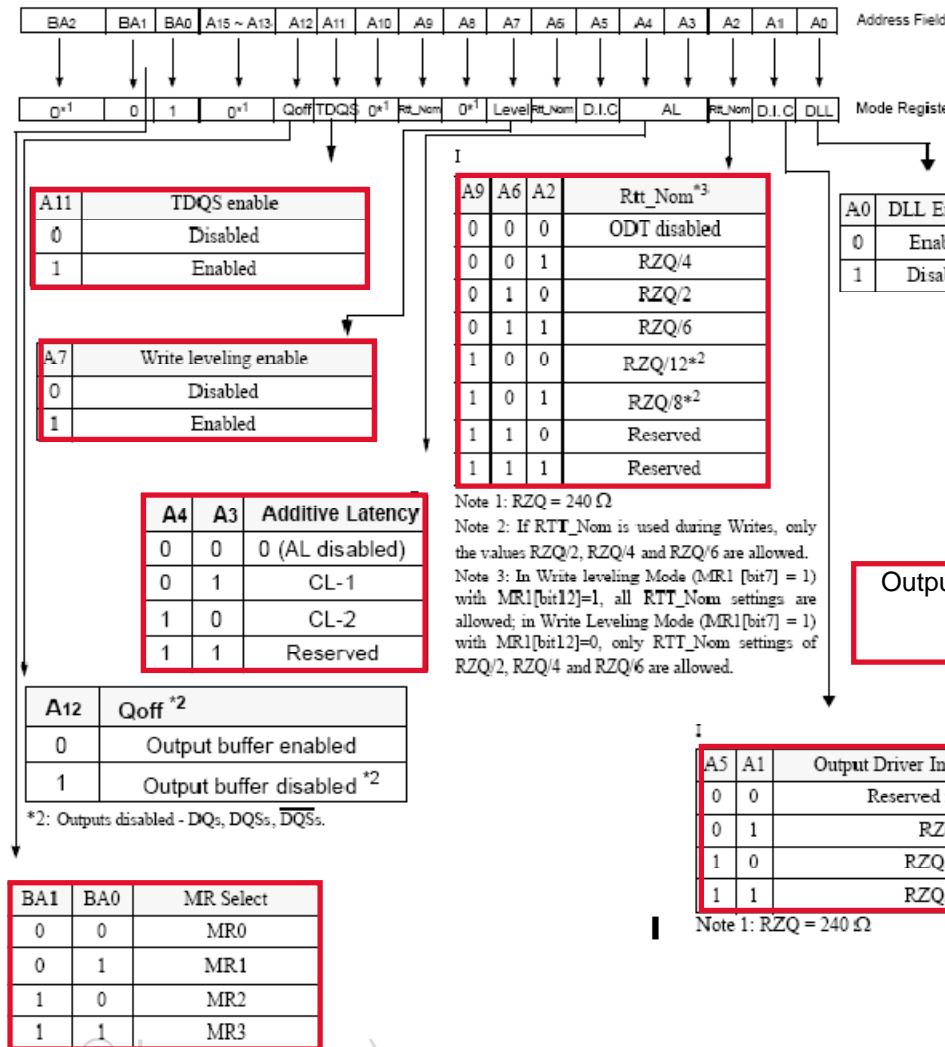


New/changed features compared to DDR2 are marked in red.



DDR3 Mode Registers

Mode Register 1

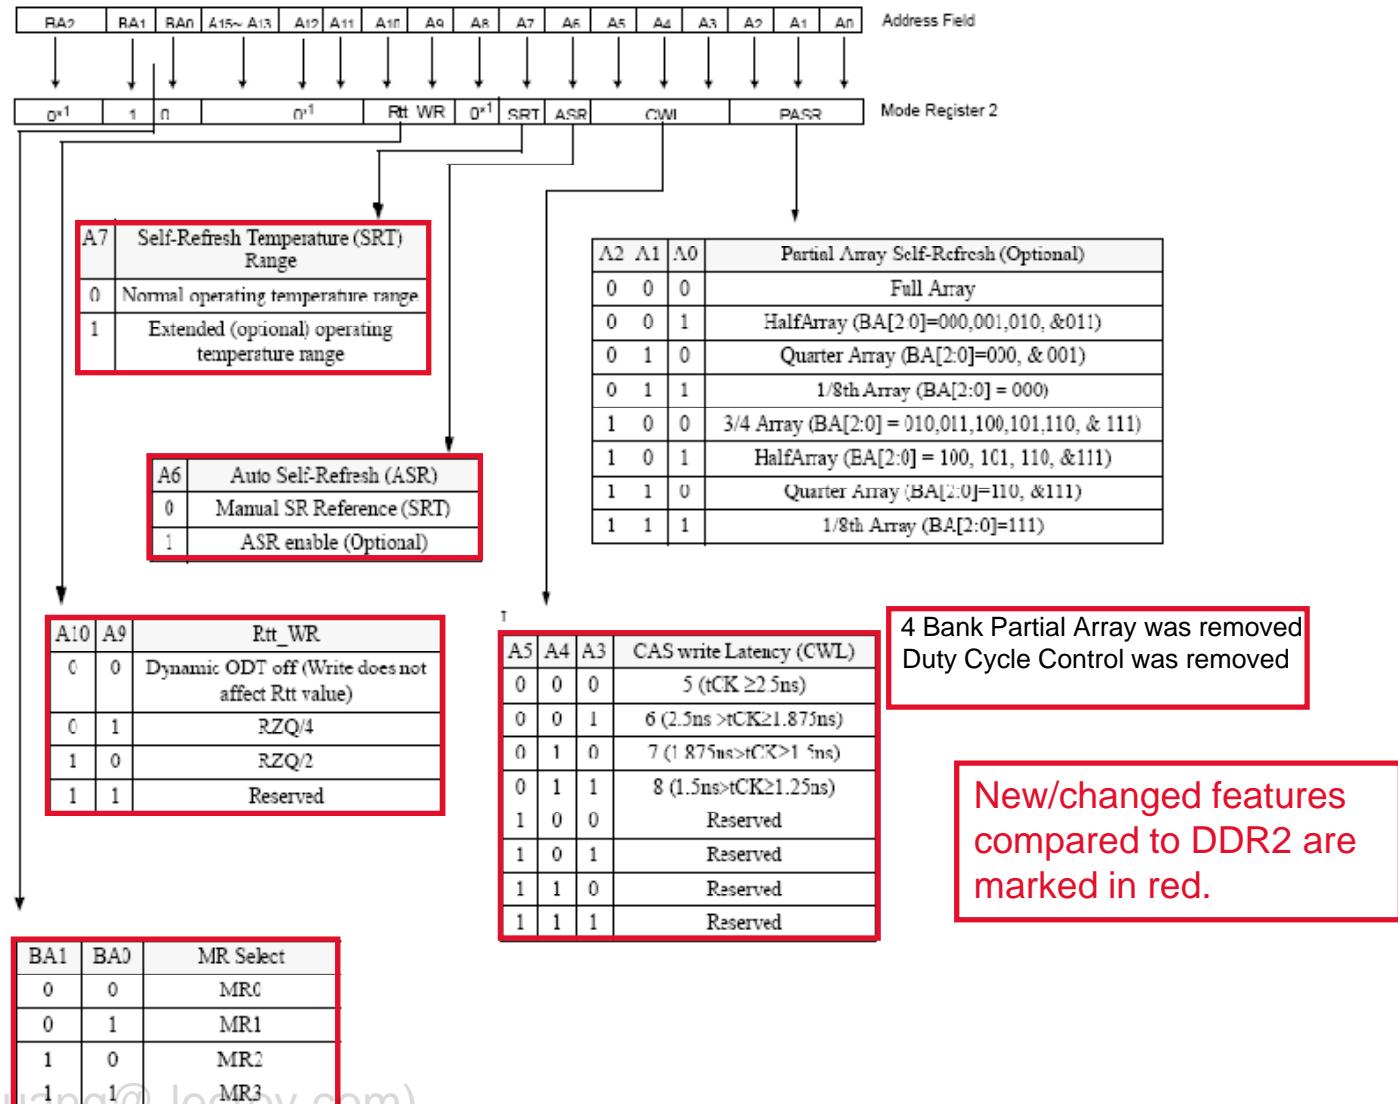


New/changed features compared to DDR2 are marked in red.



DDR3 Mode Registers

Mode Register 2

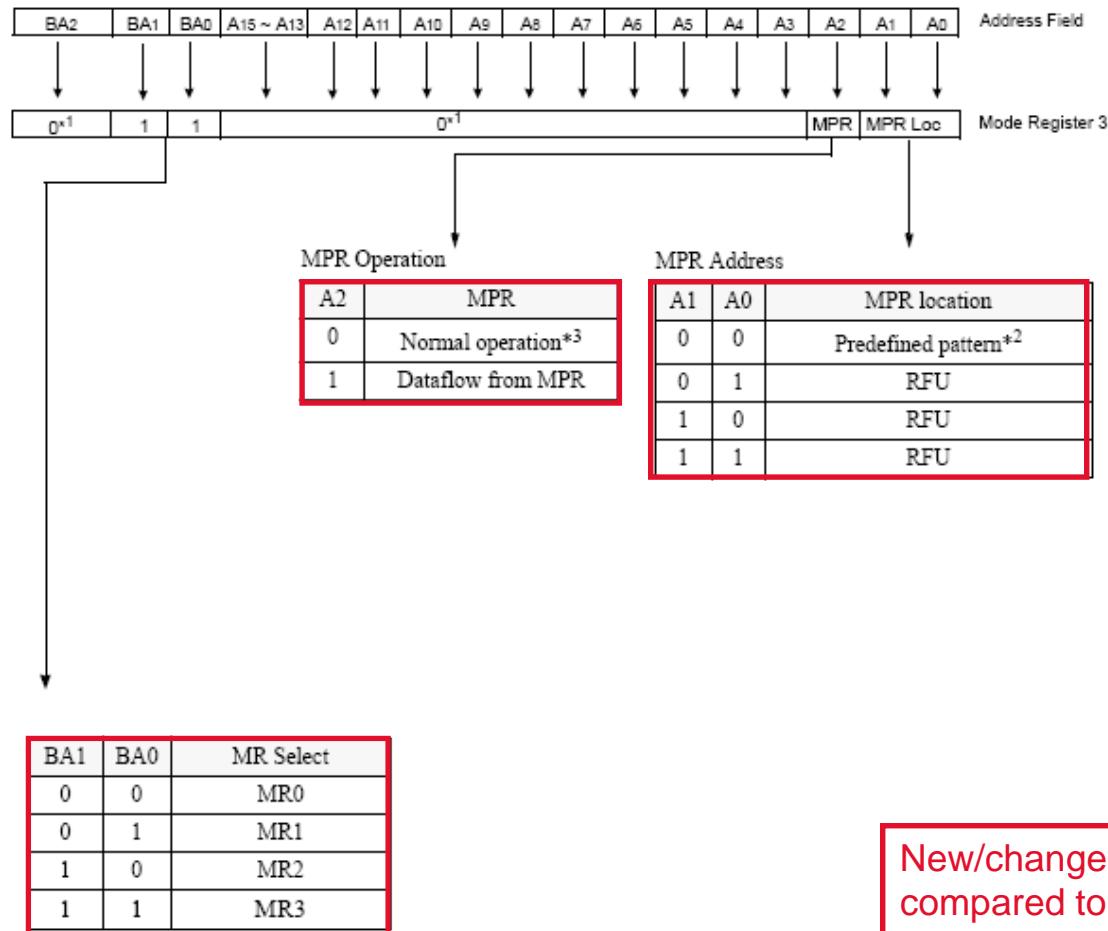




- Self refresh temperature range – Allows the controller access to the check to see if the DRAM is capable of extended temperature ranges with a reduced refresh cycle time.
- Auto Self Refresh – a DDR3 option allowing the DRAMs themselves to initiate refresh.



Mode Register 3





- Step 6
 - Issue MRS Command to load MR2 with all application settings.
- Step 7
 - Issue MRS Command to load MR3 with all application settings.
- Step 8
 - Issue MRS Command to load MR1 with all application settings and DLL enabled.
- Step 9
 - Issue MRS Command to load MR0 with all application settings and “DLL reset”.



- Step 10
 - Issue ZQCL command to start ZQ calibration.
- Step 11
 - Wait until both tDLLK and tZQinit complete.
- Done!
 - The standard says the DRAM is now ready for normal operation, but the controller still needs to train the timing. This is covered in the Read Calibration and Write Leveling sections later.

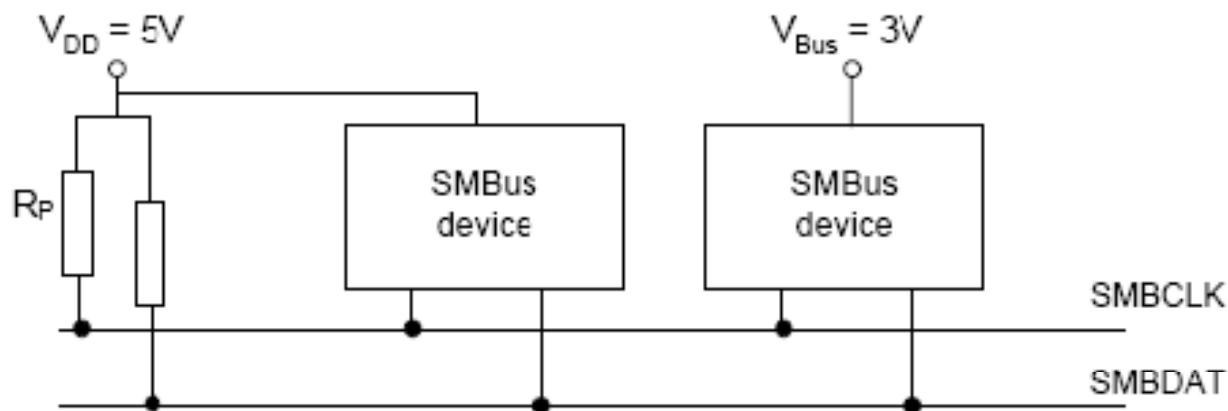


SMBus Overview

Min Huang(min.huang@ lecroy.com)
Do Not Distribute



SMBus is a 2-wire Multi-Master and Multi-Slave device interface based on the OSI model used in high speed network protocols.





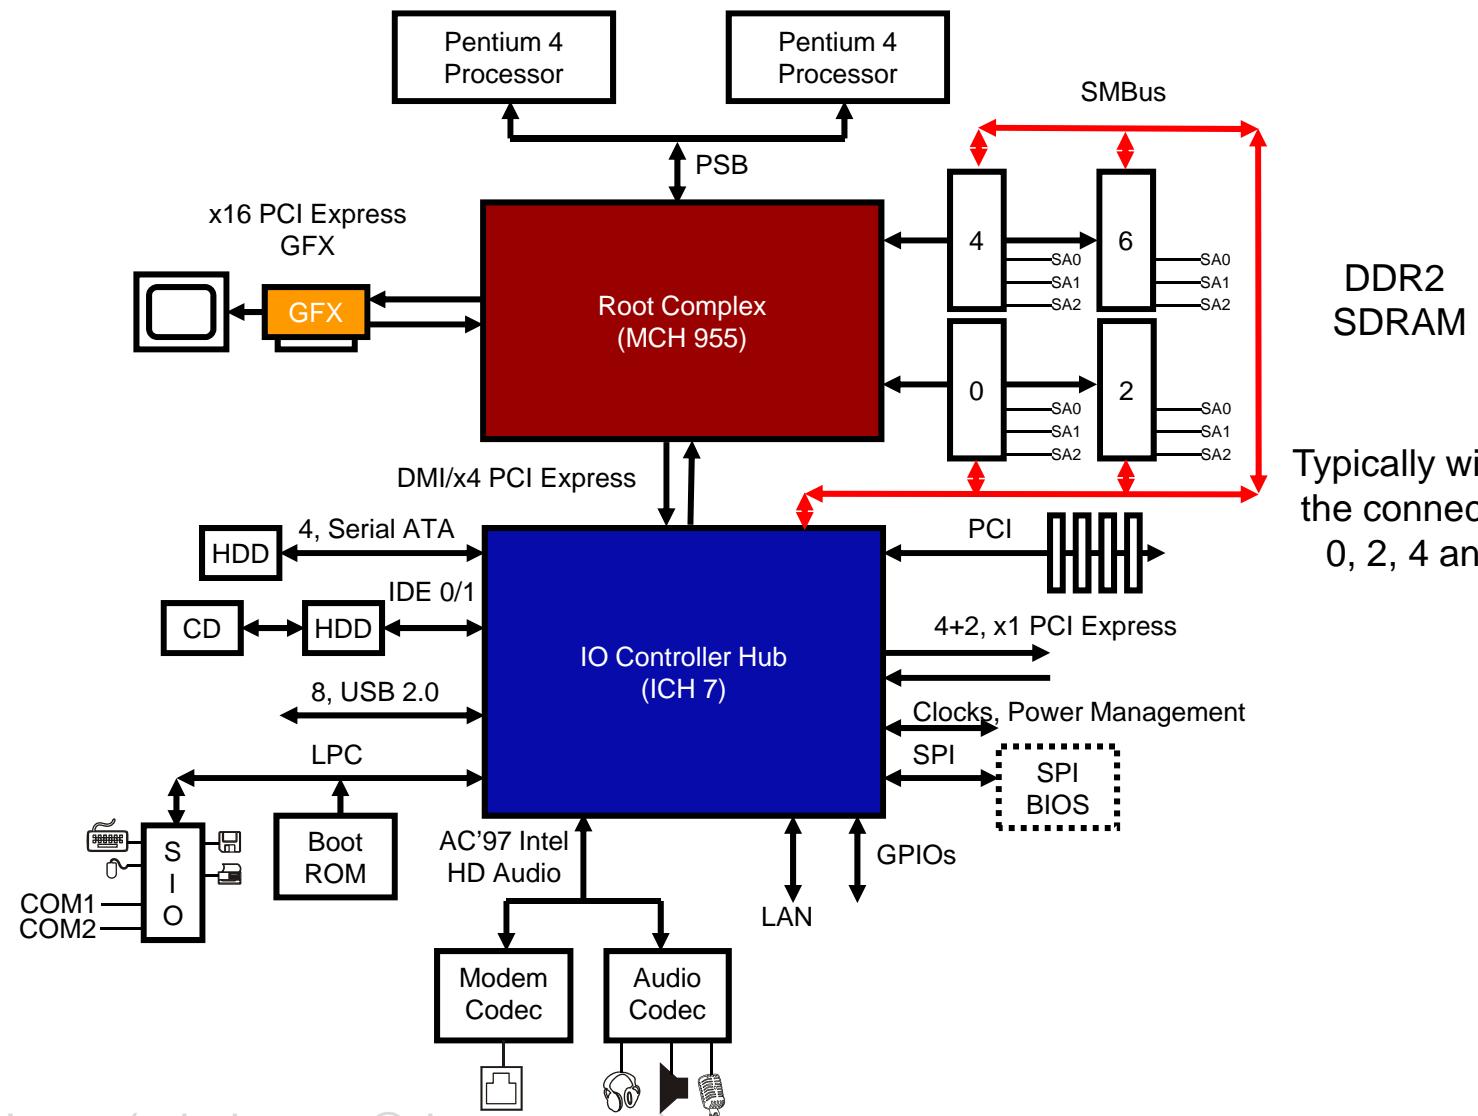
- Bus masters initiate transactions to slave devices and provide the clock.
- 7-bit addresses (DRAM usually uses the lower 3 bits for its hardwired DIMM SPD addresses in a PC. The upper 4 bits are usually Ah, 1010b)
- Min speed is 10KHz and Max is 100KHz
- Devices may be powered by 3 or 5 volts.

- The three layers pertinent to reading the SPD:
 - Network Layer
 - Data Link Layer
 - Physical Layer
- To read the SPD through an operating system an application will be needed to access the Bus Master that is attached to the SPDs.
- This Bus Master may reside in the MCH or ICH and there may be multiple Masters, such as for multiple memory channels.



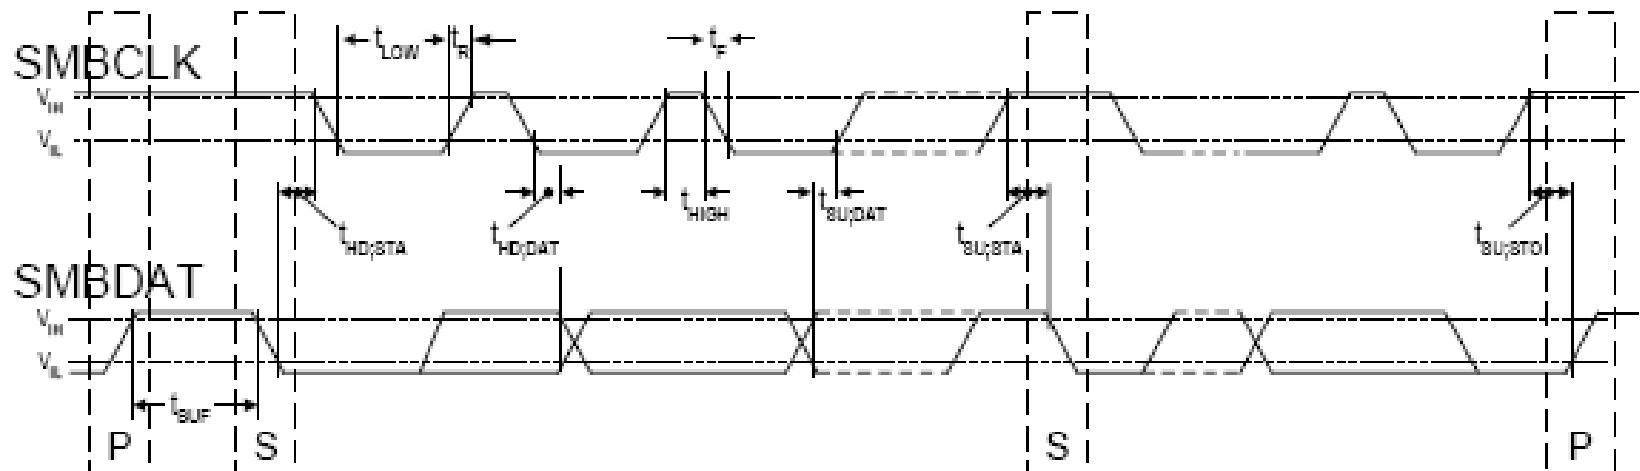
System SMBus Layout

169



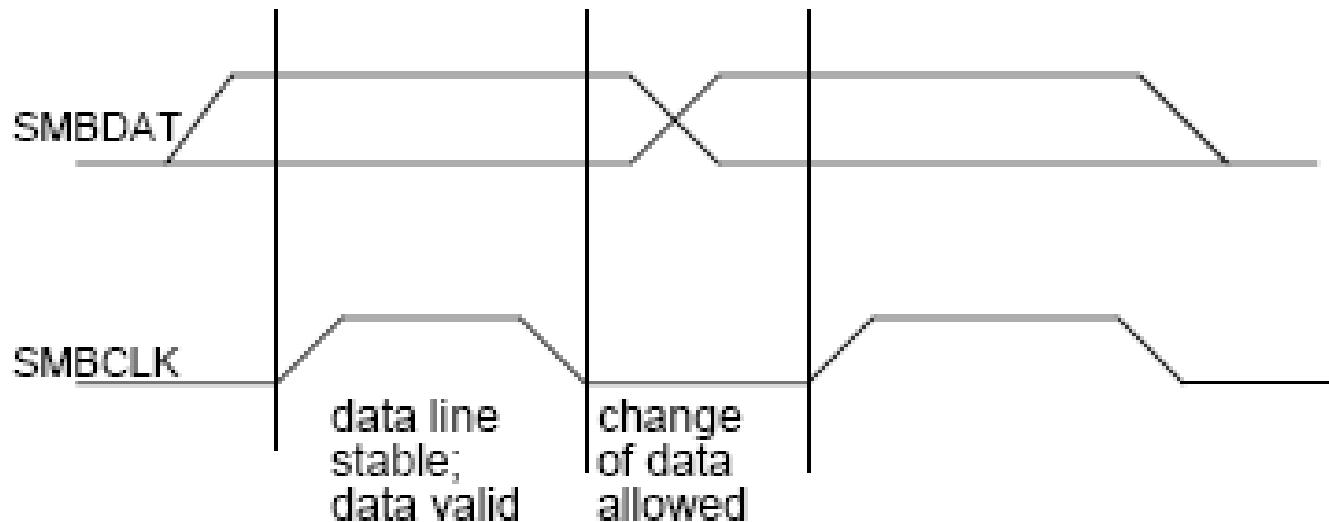


- The Physical layer consists of the Drive and Receive buffer for the clock and data.
- When the bus is idle SMBCLK is at 0KHz and pulled high.



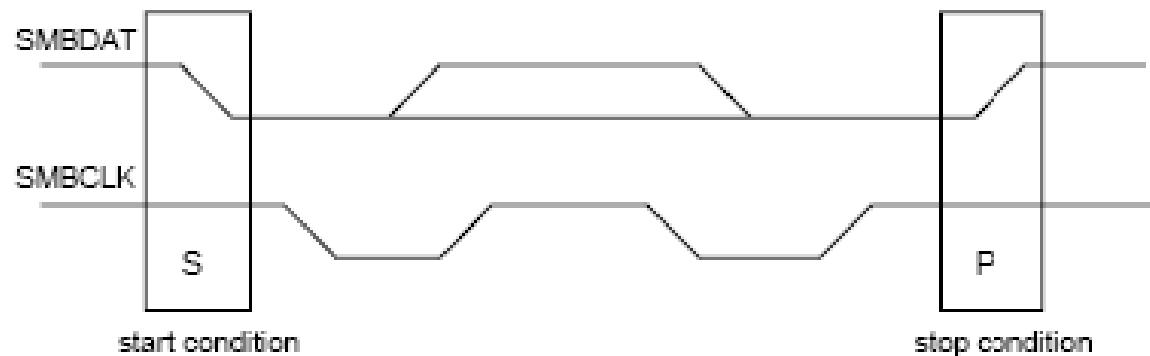


- The SMBus uses level-sensitive logical high and low signaling.
- SMBDAT must be stable when SMBCLK is high.
- Transitions only occur when SMBCLK is low.
- **Data is latched only when SMBCLK is high.**



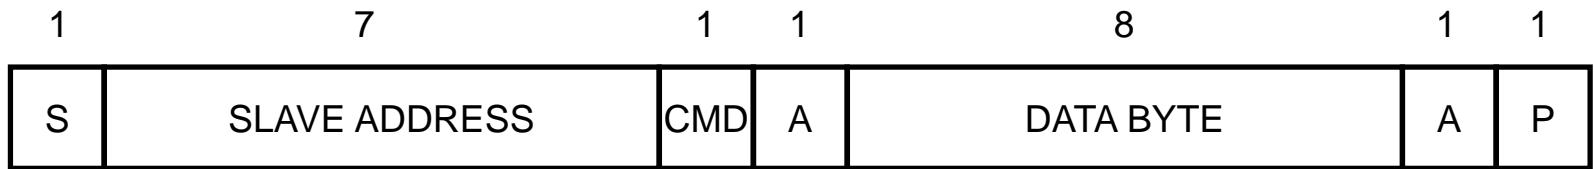


- Definition of a start and stop condition:
 - Start (S) is indicated by a high to low transition on SMBDAT while SMCLK is high. Once the Start transaction has occurred the bus is in the Busy state.
 - Stop (P) is indicated by a low to high transition on SMBDAT while clock is high. Start and Stop are only initiated by a Bus Master. After the Stop transaction the bus is the Idle state.





- A generic packet is made up of Start and Stop bits, Slave Address, Data Byte, Command bit and Ack/Nack bits.





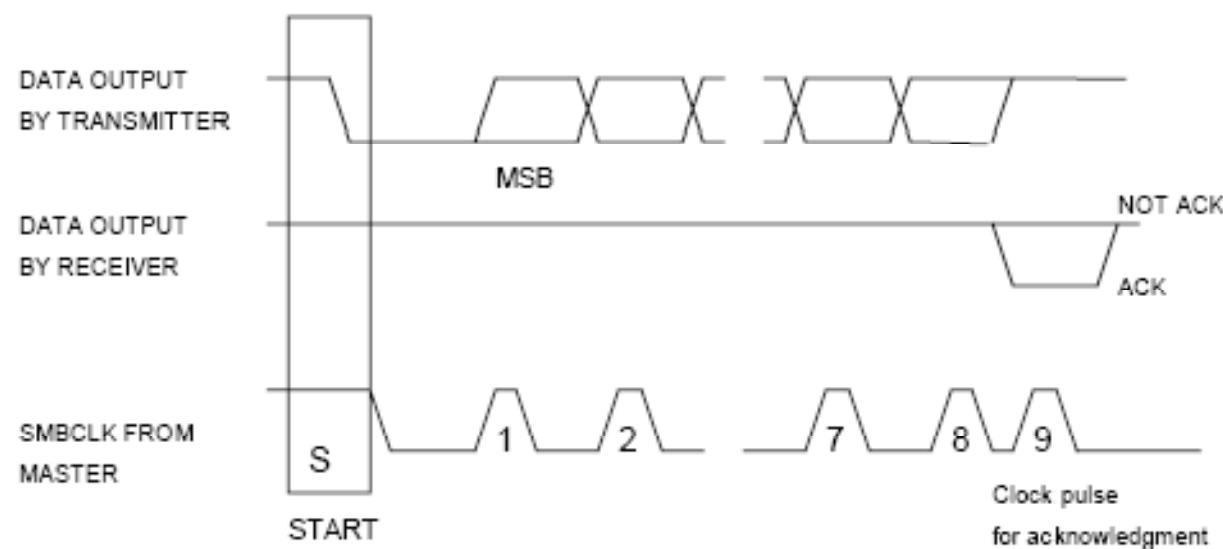
➤ Ack/Nack overview:

- An Ack or Nack must be generated for every byte that is transferred.
- The transmitter must release the SMBDAT during the Ack/Nack clock period.
- An Ack will be signaled by asserting the SMBDAT line low. If SMBDAT is not actively driven low the external pull-ups will keep it high signaling a Nack.



➤ Ack/Nack overview:

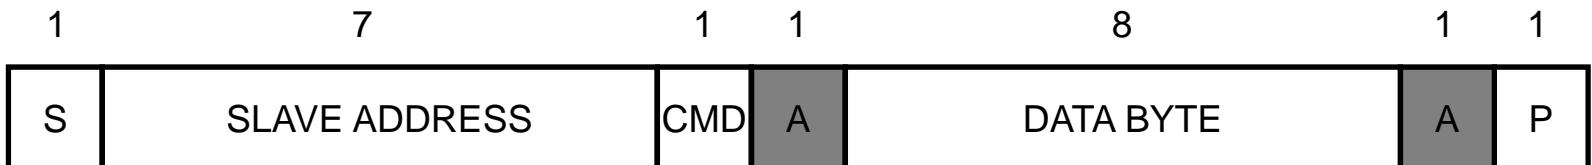
- An Ack merely indicates that a byte was received within the correct timing window. This does not mean there was any error checking done like Parity, ECC, or CRC.





➤ Usage Model Overview

- There are 3 types of devices
- 1. Master devices initiate transfers as well as drive SMBCLK.
- 2. Slave devices act as receivers and respond to commands.
- 3. Host devices are a Master and Slave combination that act as the interface to the systems CPU.



S Start Condition

Sr Repeat Start Condition

Rd Read (bit value of 1)

Wr Write (bit value of 0)

A Acknowledge (this bit position may be 0 for Ack and 1 for Nack)

P Stop Condition

Slave to Master

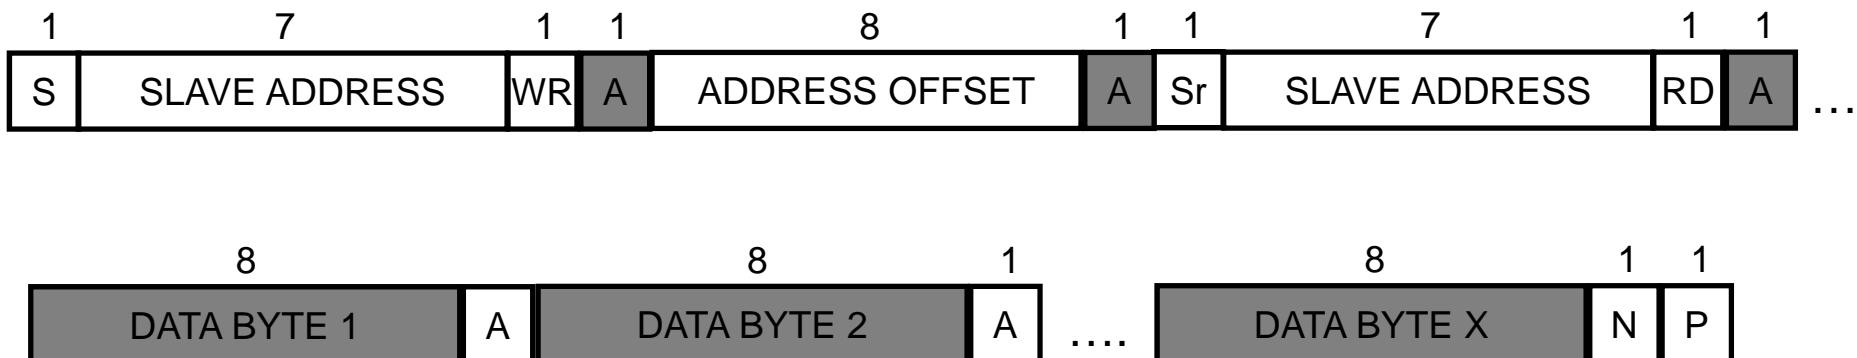
 Master to Slave





Example of Block Read

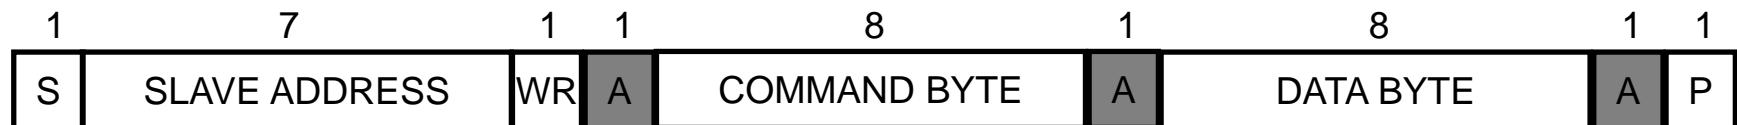
- Block Read starts write to the slave.
- The data payload is the address offset into the SPD.
- Now that the address pointer has been loaded the Start repeat is used to continue
- Nack can be signaled before Stop to indicate the end of the transfer.
- No shade indicates from Master. Shaded indicates from the slave device.





Example of Write Byte

- The Master asserts the Slave address followed by the write bit. The device acknowledges and the Master sends the command code. The slave acknowledges again and the Master transfers the Byte/Word Low byte first to the device.
- No shade indicates from Master. Shaded indicates from the slave device.



- Quick Command
- Send Byte
- Receive Byte
- Read Byte/Word
- Write Byte/word
- Block Read
- Block Write
- Process call
- Host Notify
- All of these can support the PEC packet (Packet Error Checking)



Electrical Specifications

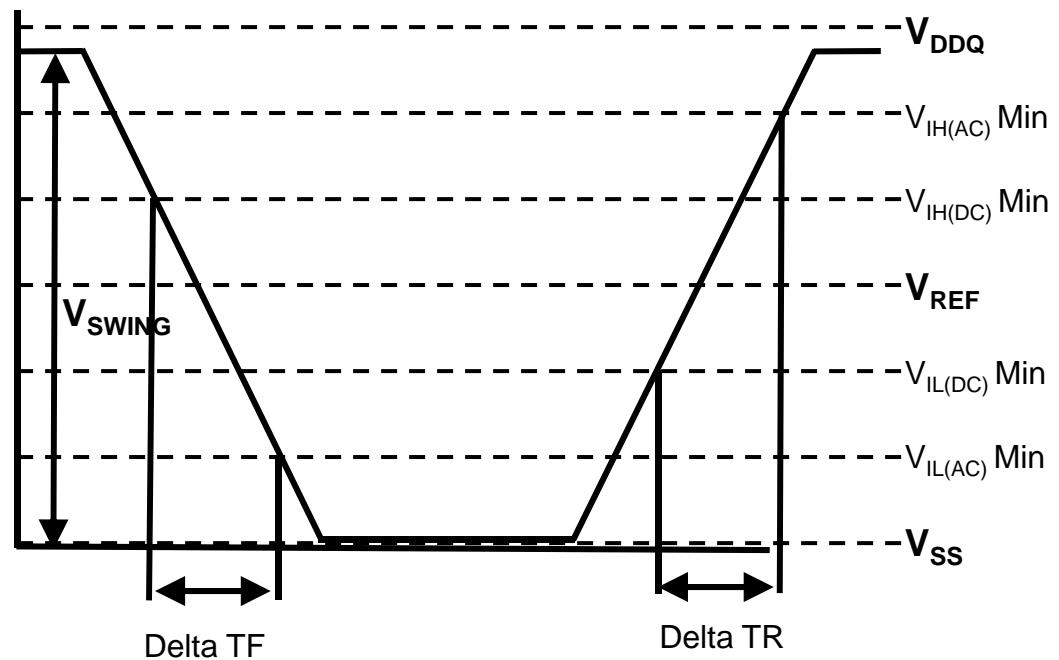
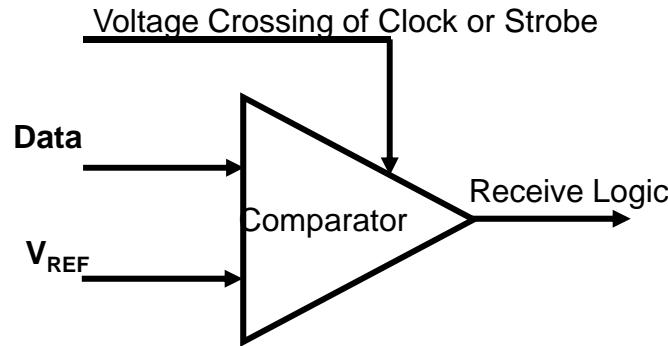
Min Huang(min.huang@ lecroy.com)
Do Not Distribute



- DDR is based on reference comparator input logic. This logic is defined by the Stub Series Terminated Logic (SSTL) spec for each of the DDR voltage operational modes.
- For DDR2, JESD8-15A is the spec.

- Voltage requirements overview
 - DDR1 is SSTL_2. (Stub Series Terminated Logic)
 - Main power is V_{DDQ} 2.5 Volts
 - Ground is V_{SS}
 - V_{REF} is $V_{DDQ}/2$
 - V_{TT} and V_{REF} are 1.25 Volts tracking with V_{DDQ}

- DDR2 is SSTL_18. (Stub Series Terminated Logic)
- Main power is V_{DDQ} 1.8 Volts
- Ground is V_{SS}
- V_{REF} is $V_{DDQ}/2$
- V_{TT} and V_{REF} are 0.9 Volts tracking with V_{DDQ}



- JEDEC specifies all of the states that vendors must use to define the min and max current drawn by their parts.
- Many of the vendors provide tools to help specify the Total Design Power or TDP.
- The following is a list of all the parameters that must be included.

Symbol	Conditions
IDD0	Operating one bank active-precharge current; CKE is high, CS is high between valid commands; Address bus inputs and switching; Data bus inputs are switching
IDD1	Operating one bank active-read-precharge current; BL (Burst Length)=4; CKE is high, CS is high between valid commands; Address bus inputs and switching; Data bus inputs are switching
IDD2P	Precharge power down current; All banks idle; CKE is low; Address and control bus inputs are stable; Data inputs are floating
IDD2Q	Precharge quiet standby current; All banks idle; CKE is high, CS is high; Address and control bus inputs are stable; Data inputs are floating
IDD2N	Precharge standby current; All banks idle; CKE is high, CS is high; Address and control bus inputs are switching; Data inputs are switching

Symbol	Conditions
IDD3P	Active power down current; (typically broken into fast or slow power down) All banks open (active); CKE is low; Address and control bus inputs are stable; Data inputs are floating
IDD3N	Active standby current; All banks open (active); CKE is high, CS is high between valid commands; Address and control bus inputs are switching; Data inputs are switching
IDD4W	Operating burst write current; BL (Burst Length) =4 All banks open (active), Continuous burst writes; CKE is high, CS is high between valid commands; Address and control bus inputs are switching; Data inputs are switching
IDD4R	Operating burst read current; BL (Burst Length) =4 All banks open (active), Continuous burst reads; CKE is high, CS is high between valid commands; Address and control bus inputs are switching; Data inputs are switching

Symbol	Conditions
IDD5B	Burst refresh current; Refresh command at every t_{RFC} interval; All banks open (active), Continuous burst reads; CKE is high, CS is high between valid commands; Address and control bus inputs are switching; Data inputs are switching
IDD6	Self refresh current; CK and CK# are at 0Volts; CKE is less than or equal to 0.2 volts; Address and control bus inputs are floating; Data inputs are floating
IDD7	Operating bank interleave read current; All banks interleaving reads; BL (Burst Length)=4 CKE is high, CS is high between valid commands; Address and control bus inputs are stable during deselects; Data inputs are switching



DDR1 and DDR2 Routing

Min Huang(min.huang@ lecroy.com)
Do Not Distribute

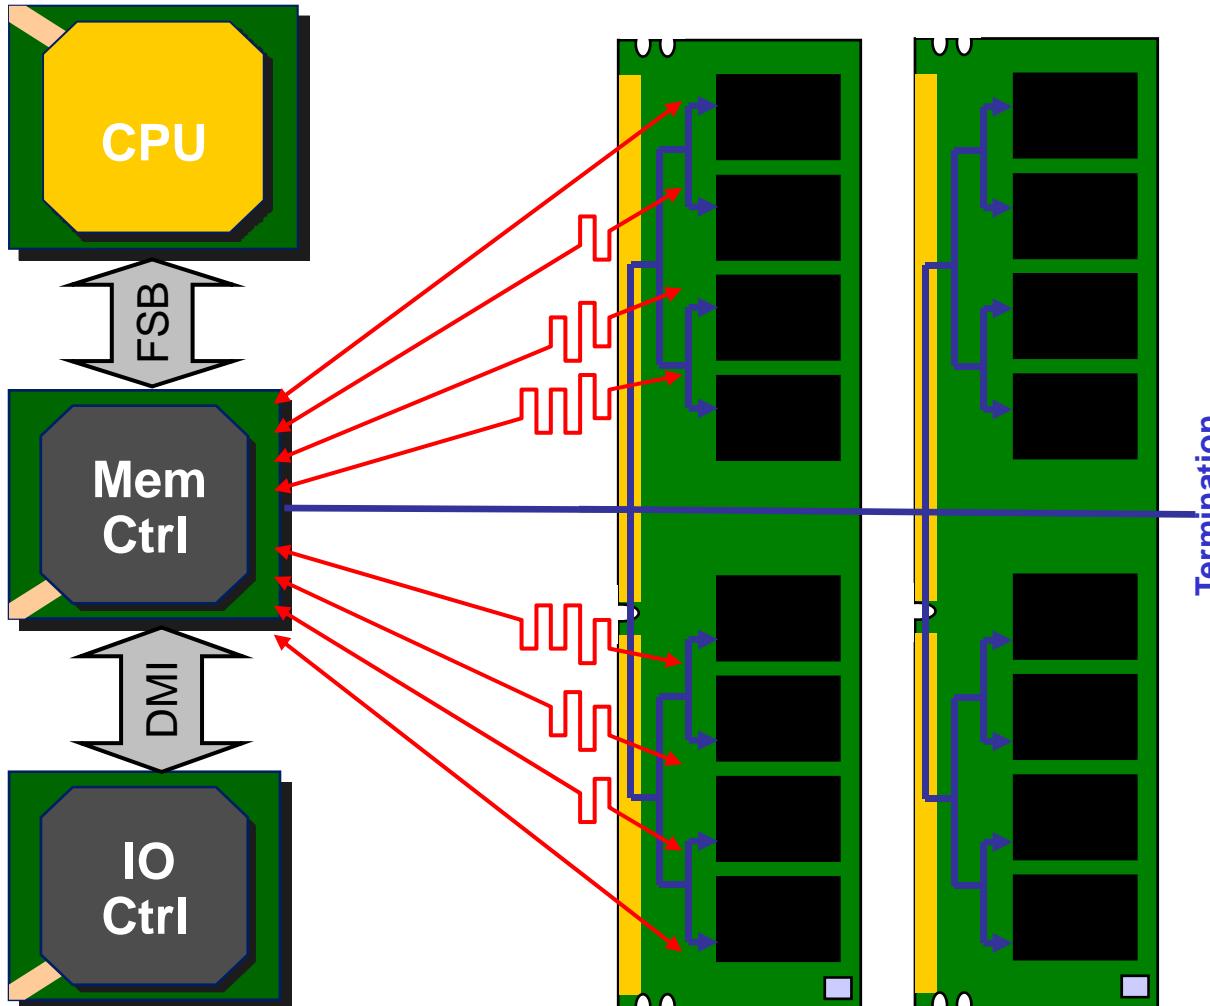


- There are many things to consider when designing a DDR2 system.
- Source synchronous clocking schemes and how to route them.
- ODT as well as Command and Address termination
- Input delay to center DQS to data

- Always refer to the chipset design guide for lead-in length and DIMM-to-DIMM spacing.
- All Command and Address signals must be end terminated as well as length-matched within some tolerance. They are all in the same time domain and are latched with clock.



- Data and strobe will be On-Die Terminated
- If the system supports x4 DRAMs, then each group of 4 data lines and their associated strobe must be length matched. These signals are source synchronous.
- If system does not support x4 DRAMs, then every byte lane needs to be length matched to its associated strobe.

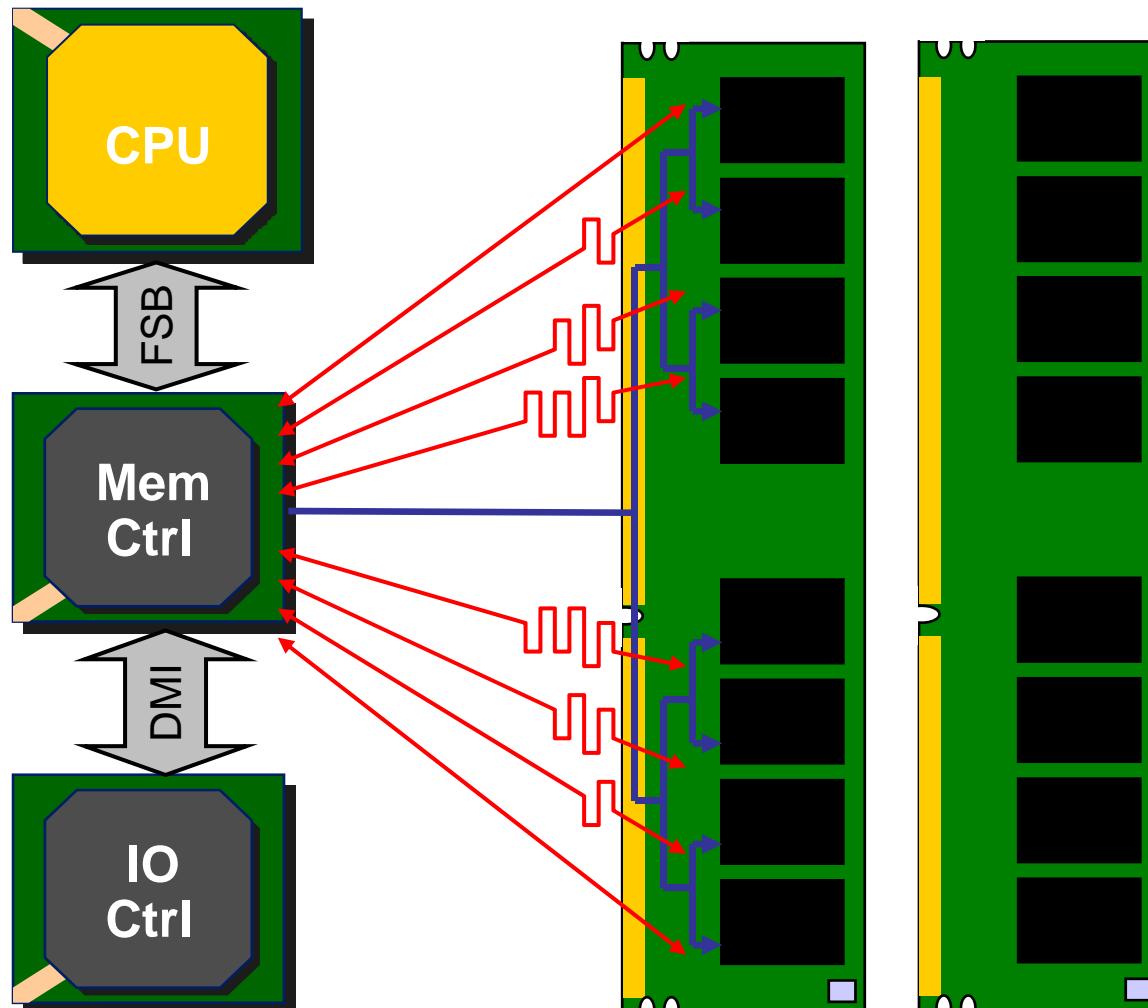


- CMD and ADDR are terminated on Motherboard.
- Data Bus is terminated on die.
- Byte lanes and their associated DQS must be length matched. For x4 chips, nibbles must be length matched with their DQS.
- Command and address are also length matched to ensure that all commands will get to each DRAM at the same time.

Key: Red is Data Bus

Blue is Address and Command Bus

Min Huang/min.huang@leecroy.com)
Do Not Distribute

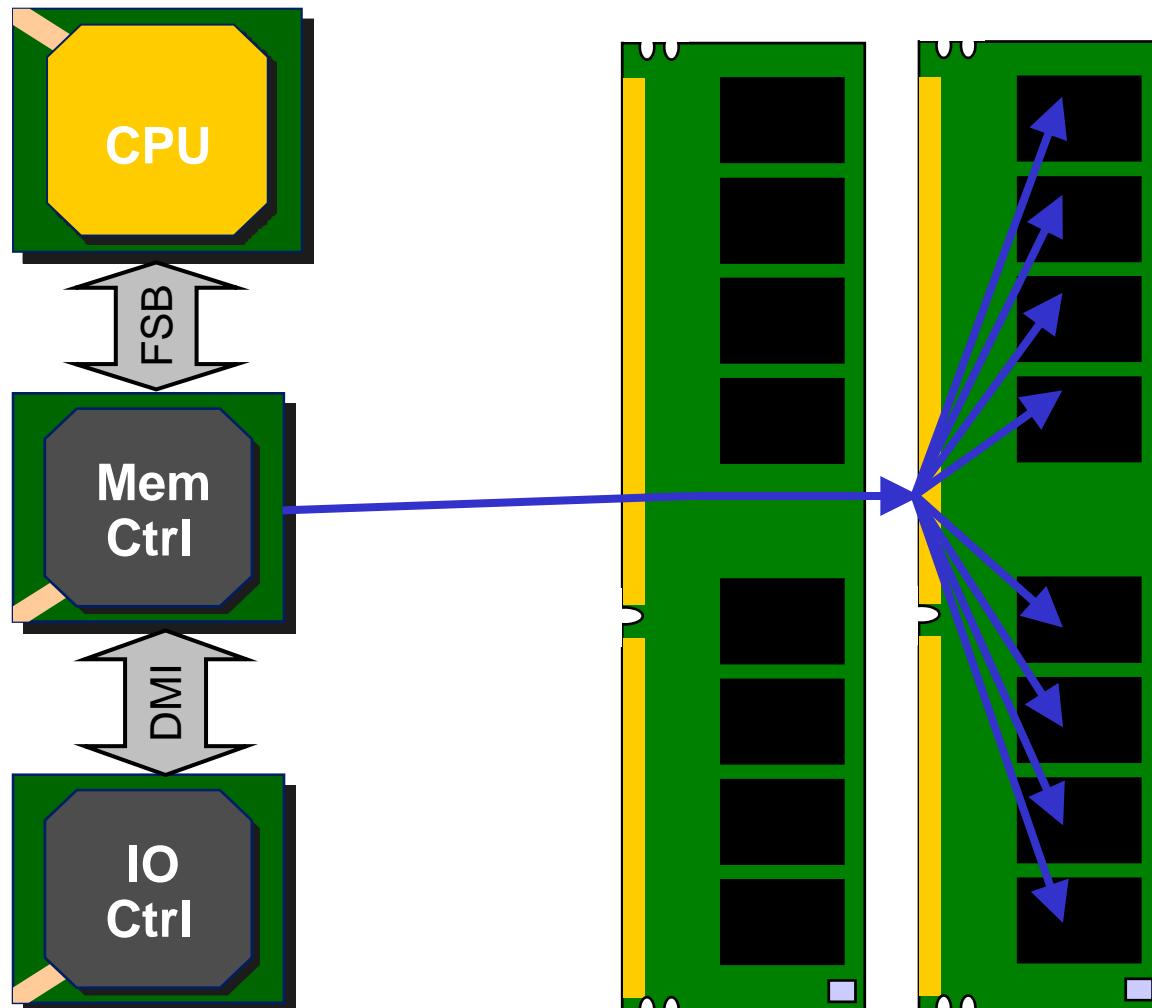


Key: Red is Data Bus

Blue is Address and Command Bus

Min Huang (min.huang@lectroy.com)
Do Not Distribute

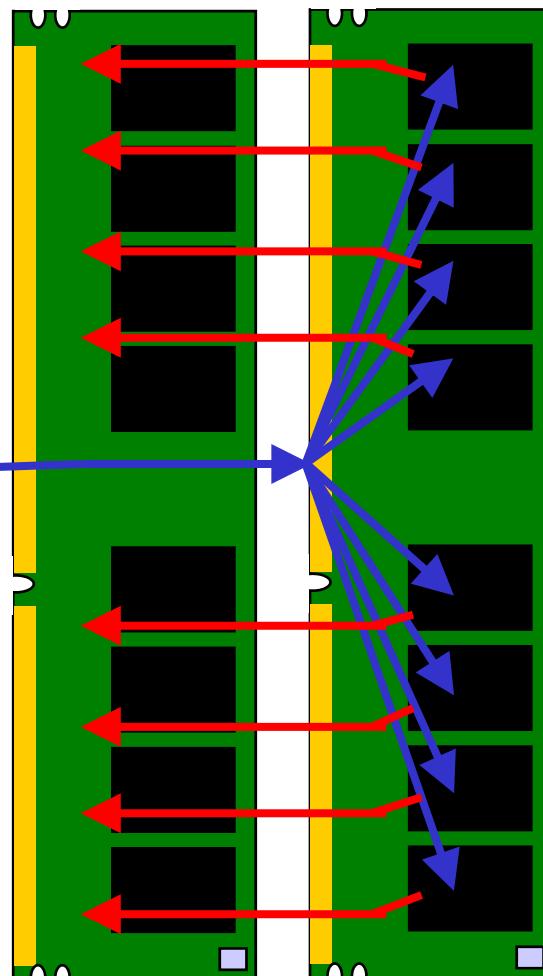
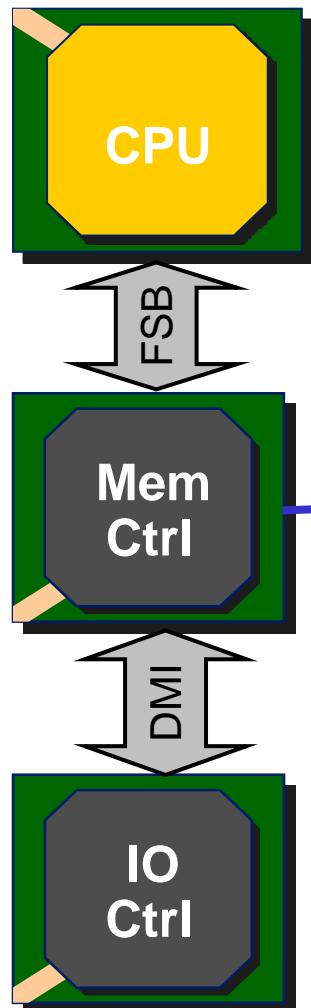
- **Traditional DDR1 and DDR2 routing.**
- Data Bus is flight-time matched for all data groups.
- Command and address are also length matched to ensure that all commands will get to each DRAM at exactly the same time.





DDR2 Read Overview

196

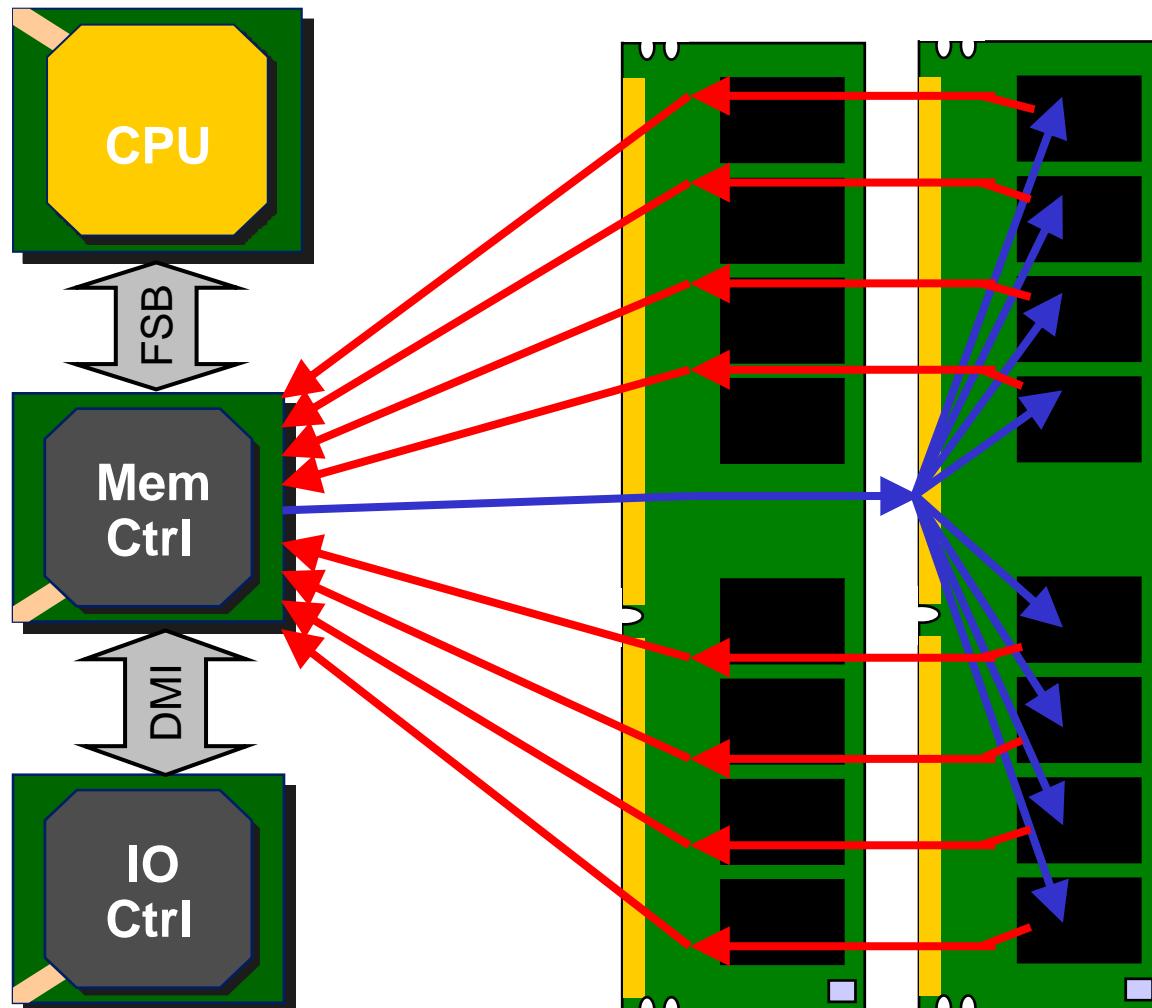


- Read Cycle in DDR2
- DRAM controller sends read command to second DIMM
- Command propagates to all DRAM with the same timing due to the length matched routing.
- After RL, the DRAMs will drive the requested data back to the controller.

Key: Red is Data Bus

Blue is Address and Command Bus

Min Huang (min.huang@leecoy.com)
Do Not Distribute



- Read Cycle in DDR2
- DRAM controller sends read command to second DIMM
- Command propagates to all DRAM with the same timing due to the length matched routing.
- After RL, the DRAMs will drive the requested data back to the controller.
- Read data arrives at the controller at the same time.

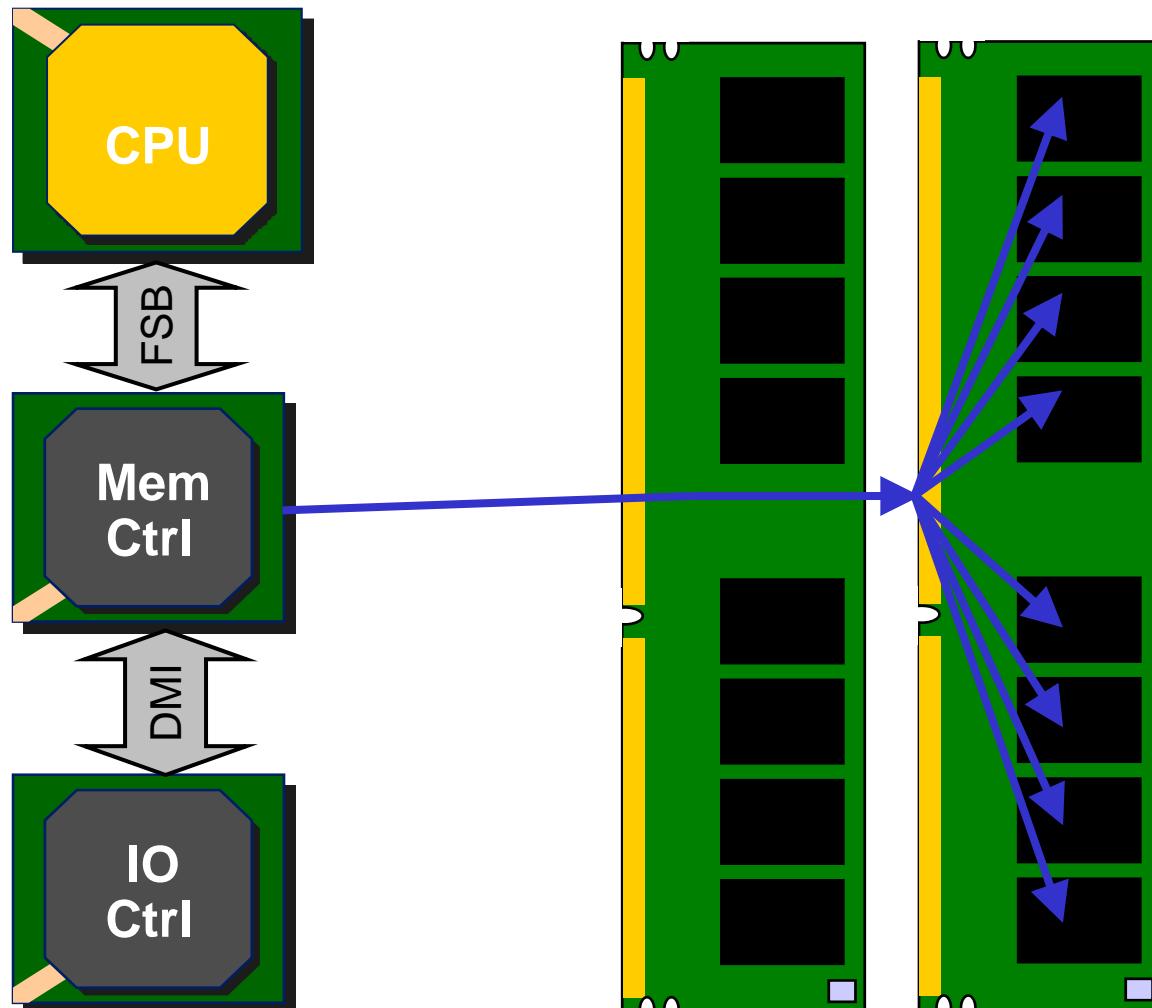
Key: Red is Data Bus

Blue is Address and Command Bus

Min Huang(min.huang@leecoy.com)
Do Not Distribute



- How does the memory controller know
 - when to launch the write data to the clock
 - and ensure t_{DSH} , t_{DSS} and t_{DQSS} are satisfied when the write data arrives at the DRAM?
- One way is to rely on the flight-time matching between the clock routing and the data bus routing on the mother board and on the DIMM.
 - In other words, if the write data is launched correctly at the controller, it *hopefully* will arrive correctly at the DRAM
- This works for DDR2 but not for DDR3

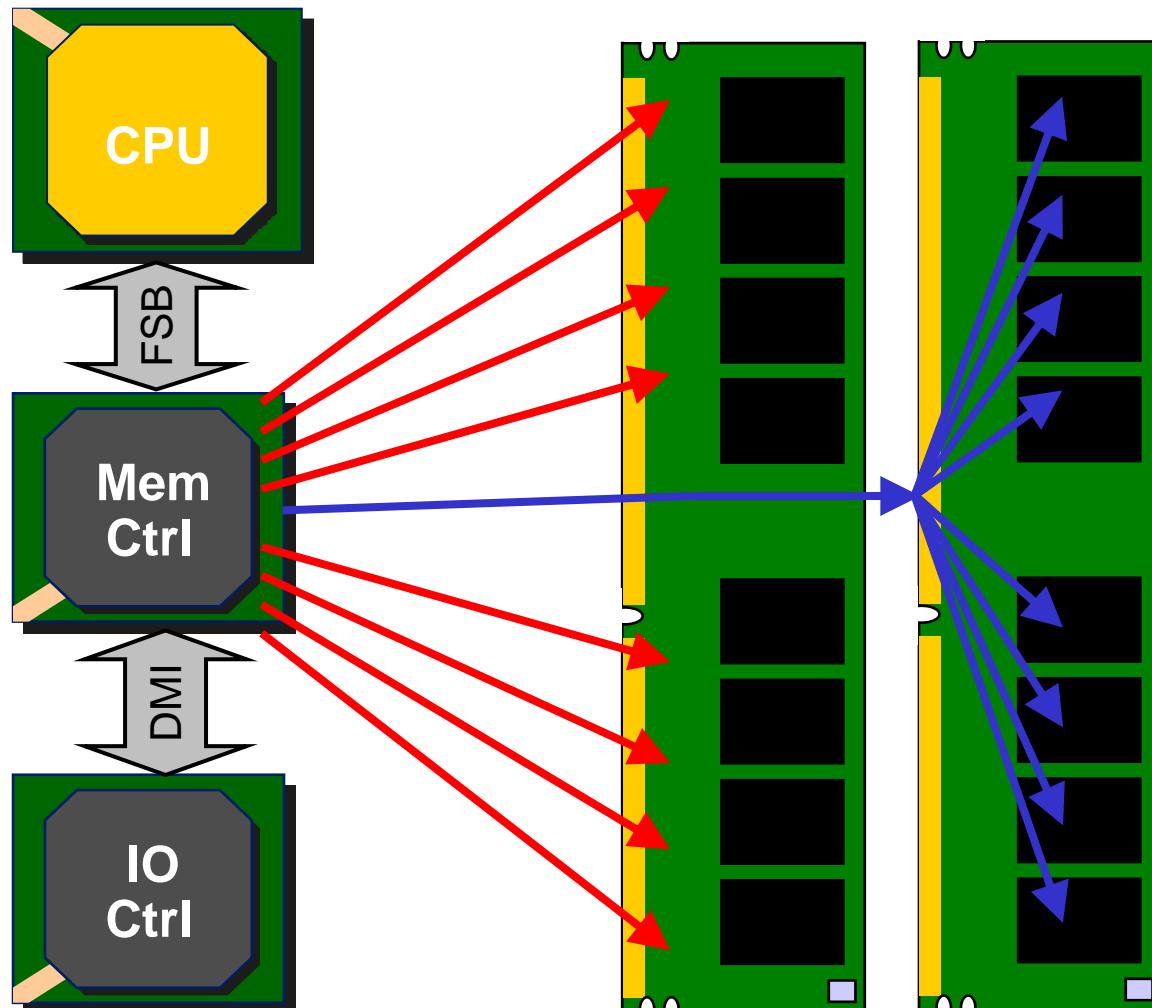


- **Write Cycle in DDR2**
- DRAM controller sends Write command to second DIMM
- Command propagates to all DRAMs with the same timing due to the length-matched routing.



DDR2 Write Overview

200

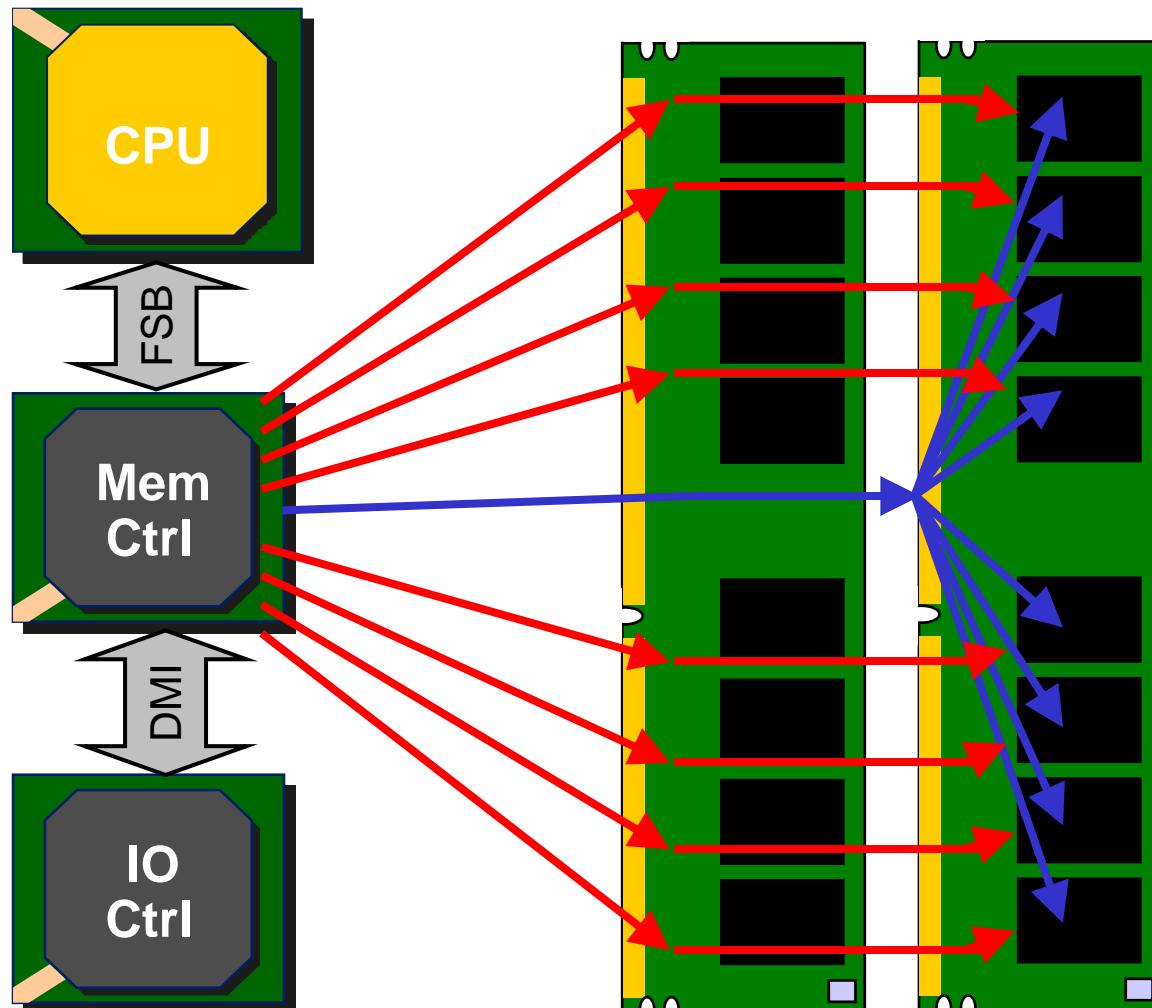


Key: Red is Data Bus

Blue is Address and Command Bus

Min Huang (min.huang@leecoy.com)
Do Not Distribute

- **Write Cycle in DDR2**
- DRAM controller sends Write command to second DIMM
- Command propagates to all DRAMs with the same timing due to the length-matched routing.
- Write data follows the command by one clock synchronously.



Key: Red is Data Bus

Blue is Address and Command Bus

Min Huang (min.huang@leecoy.com)
Do Not Distribute

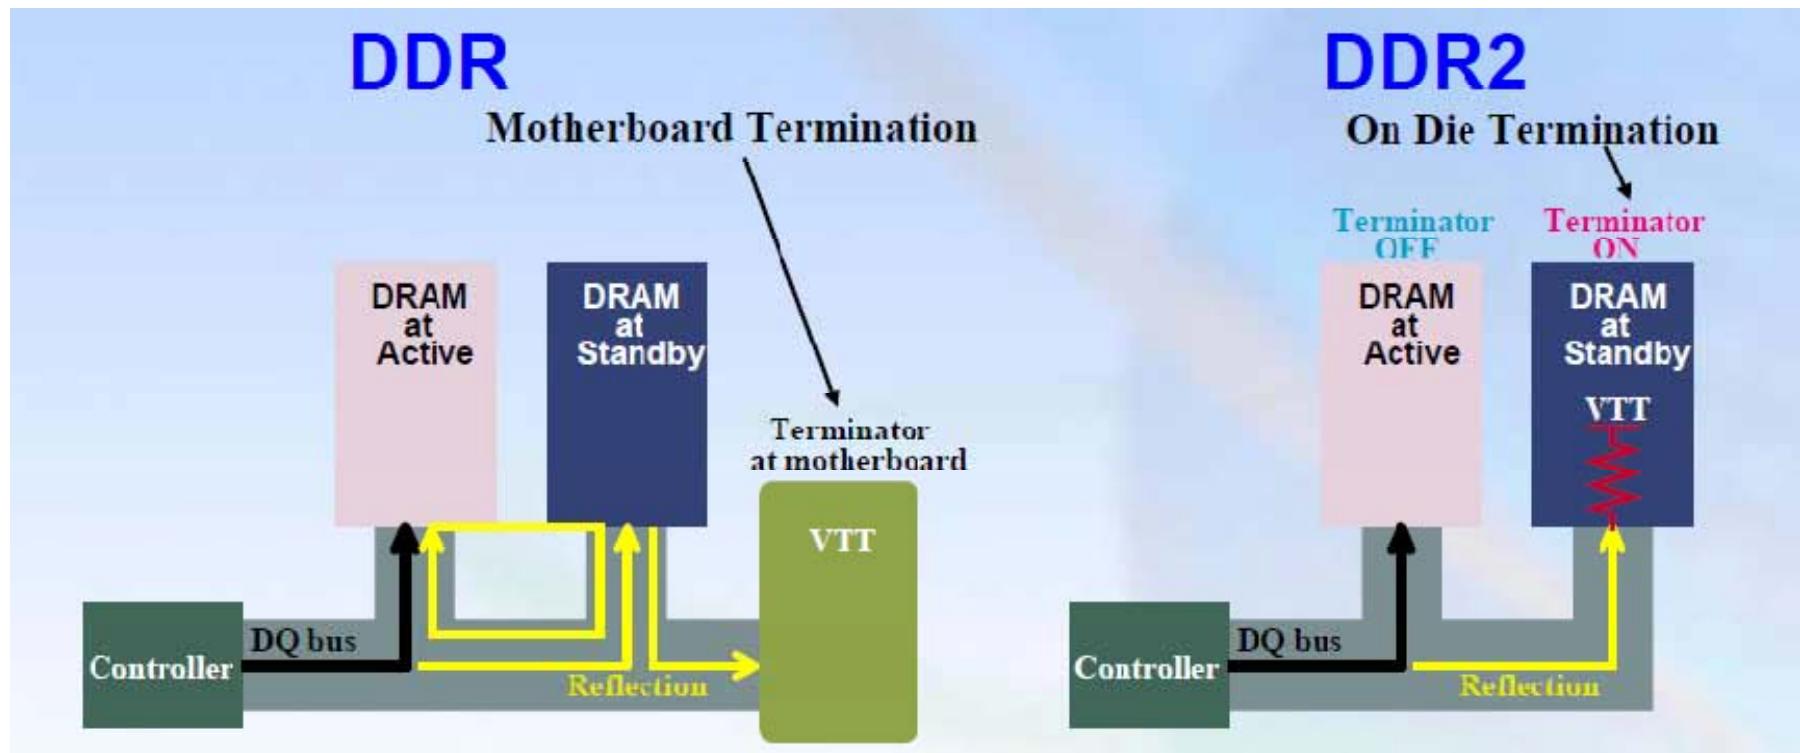
- **Write Cycle in DDR2**
- DRAM controller sends Write command to second DIMM
- Command propagates to all DRAMs with the same timing due to the length-matched routing.
- Write data follows the command by one clock synchronously.
- Write data reaches all of the DRAMs at the same time.



On-Die Termination

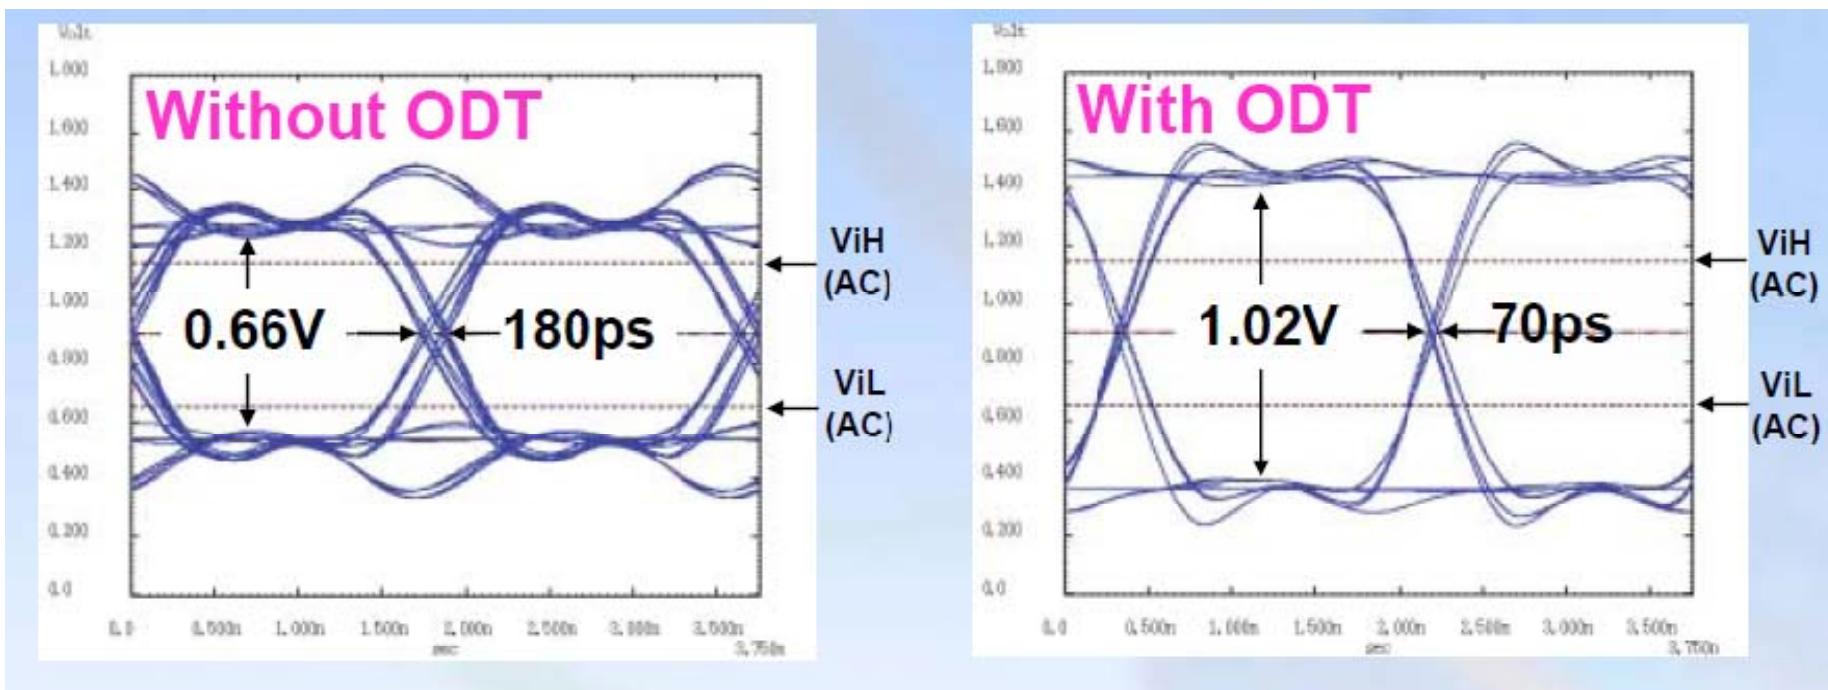
Min Huang(min.huang@ lecroy.com)
Do Not Distribute

- ODT Example: ODT reduces reflections from the stubs to DIMMs not being addressed during write cycles. ODT is also used by most DRAM controllers during read cycles. Settings are not intuitive when more than two DIMMs are used. There are 3 settings 50, 75 and 150 Ohms
- One ODT signal per Rank.





More noise margin, approx. 400mv gain





- Off Chip Driver Calibration was implemented on the DIMMs when DDR2 was introduced. The drive strength granularity was set too large on most designs so it was not recommended that the chipset run the calibration sequence. Resurrected in DDR3.
- This feature is not used in DDR2.



Errors

Min Huang(min.huang@ lecroy.com)
Do Not Distribute



Two basic types of error:

- Hard Errors – These exist in the cell array and are thought to be caused by a bad cell, damage to the DRAM etc. These errors account for less than 10% of the errors in a DDR subsystem.
- Soft Errors – Occur when electrical anomalies are present on the transmission line. These could be caused by poor design and are expected in higher speed systems.



- Error Checking and Correction is used in systems that require better reliability.
- Many servers use ECC DIMMs to guarantee data integrity.
- This is done by adding an extra 8 bits for every 64 bits. These 8 bits are referred to as the check bits.
- This feature is referred to as SEC (Single Error Correction) and DED (Double Error Detection)



- ECC is totally system dependent. What a system does when a double or multi-bit error is detected is up to the system designers.
- Since most errors are soft errors, a retry may fix the problem.
- Many controllers will just allow the error to propagate and cause a “blue screen of death”.

- Many servers use two 64-bit channels with 16 bits of ECC.
- This enables Chip Kill (IBM) or S4EC and D4ED (Intel)
- The feature is able to fix 4 consecutive nibble-aligned bits. An example of this would be most servers that use x4 DRAMs.
- If one of the DRAMs were to go bad, the controller could correct the errors on the fly and alert the administrator that errors were happening.



- This is a mandatory step for ECC-based systems, especially those with automatic scrubbing that reads memory.
- All of memory is written with correct ECC before being read.



- Scrubbing can be done while the system is running to correct single-bit errors (which are correctable).
- This reduces the chances of multi-bit errors (which might not be correctable).
- The entire DRAM subsystem is periodically scrubbed while making sure not to interfere with system performance.



- Parity is for Registered DIMMs only.
- This feature is added to the register on DDR2 DIMMs at speeds of 533MT/s and higher.
- When a command is sent to the DIMM, the register performs a parity check on the Command and Address bus.
- If a parity error is detected the QERR# out pin is asserted low.



Additional DDR3 Topics

Min Huang(min.huang@ lecroy.com)
Do Not Distribute



Timing and Electrical Differences

Min Huang(min.huang@ lecroy.com)
Do Not Distribute



- t_{CCD} or CAS-to-CAS delay is no longer variable. In DDR2 it was $\frac{1}{2}$ the burst length. DDR3 it is set to 4. (BL = 8)
- Transmitter impedance went up to approx. 34 ohms. DDR2 was approx. 18 ohms. This equates to less power consumption.

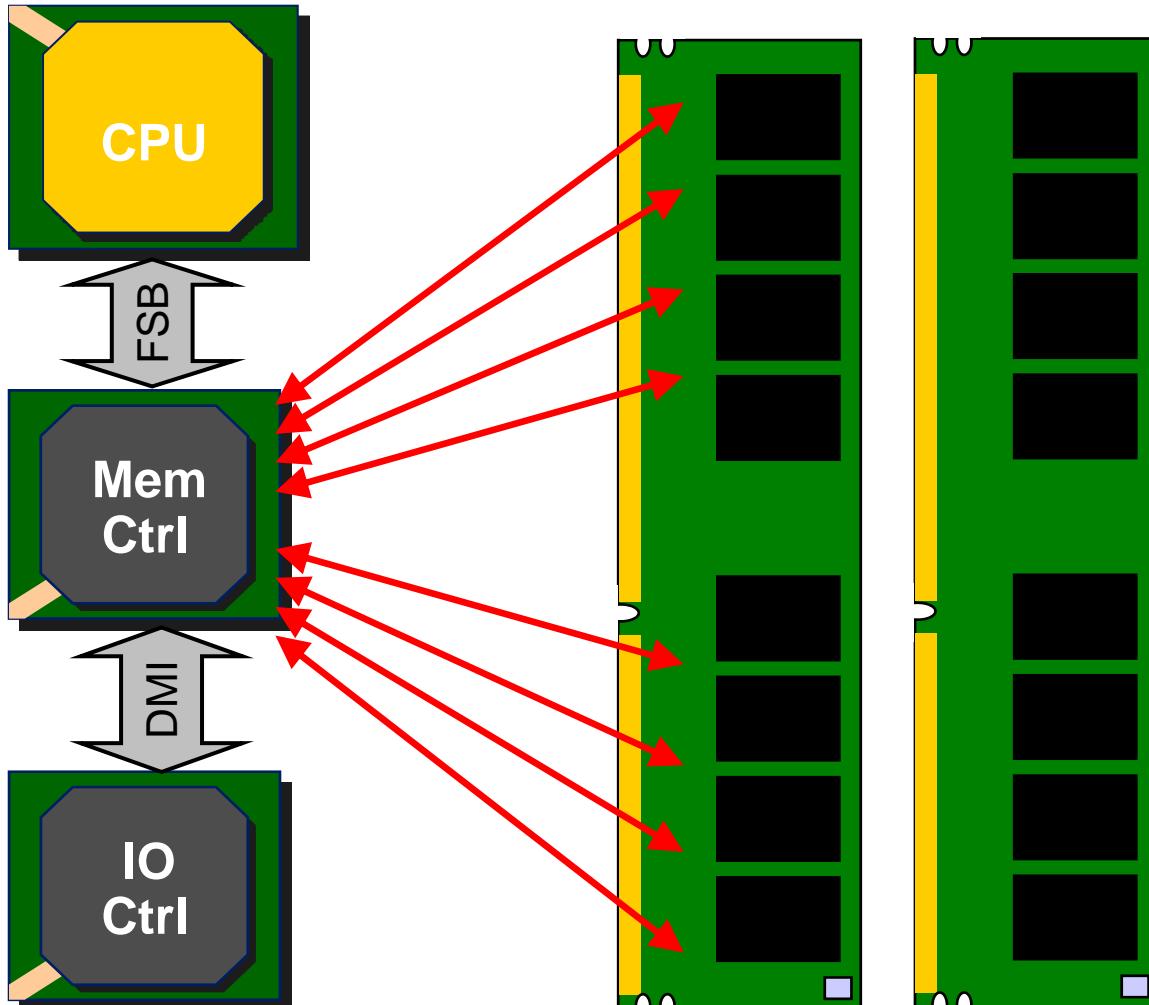


Fly-by Routing Read Example

Min Huang(min.huang@ lecroy.com)
Do Not Distribute



- DDR3 looks completely different?
- How different?
- The following slides show an example of the read data arrival time at the memory controller in a DDR3 UDIMM system.



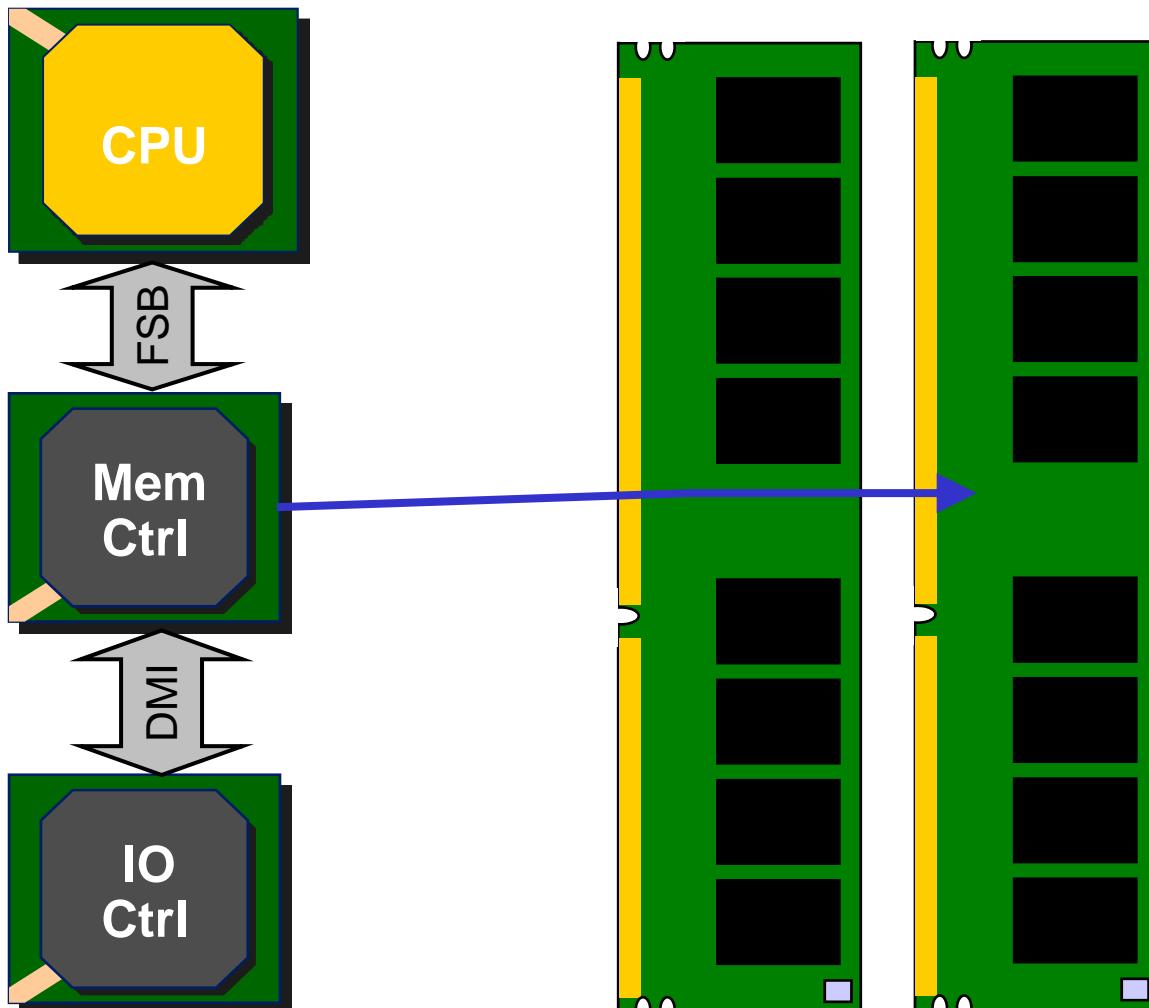
Key: Red is Data Bus
Blue is Address and Command Bus
Min Huang (min.huang@jccrccy.com)
Do Not Distribute

- **New fly-by routing**
- Data bus is still length matched within each data group to its associated strobe. There does not seem to be any implied length matching across the data groups



DDR3 Read Overview

220

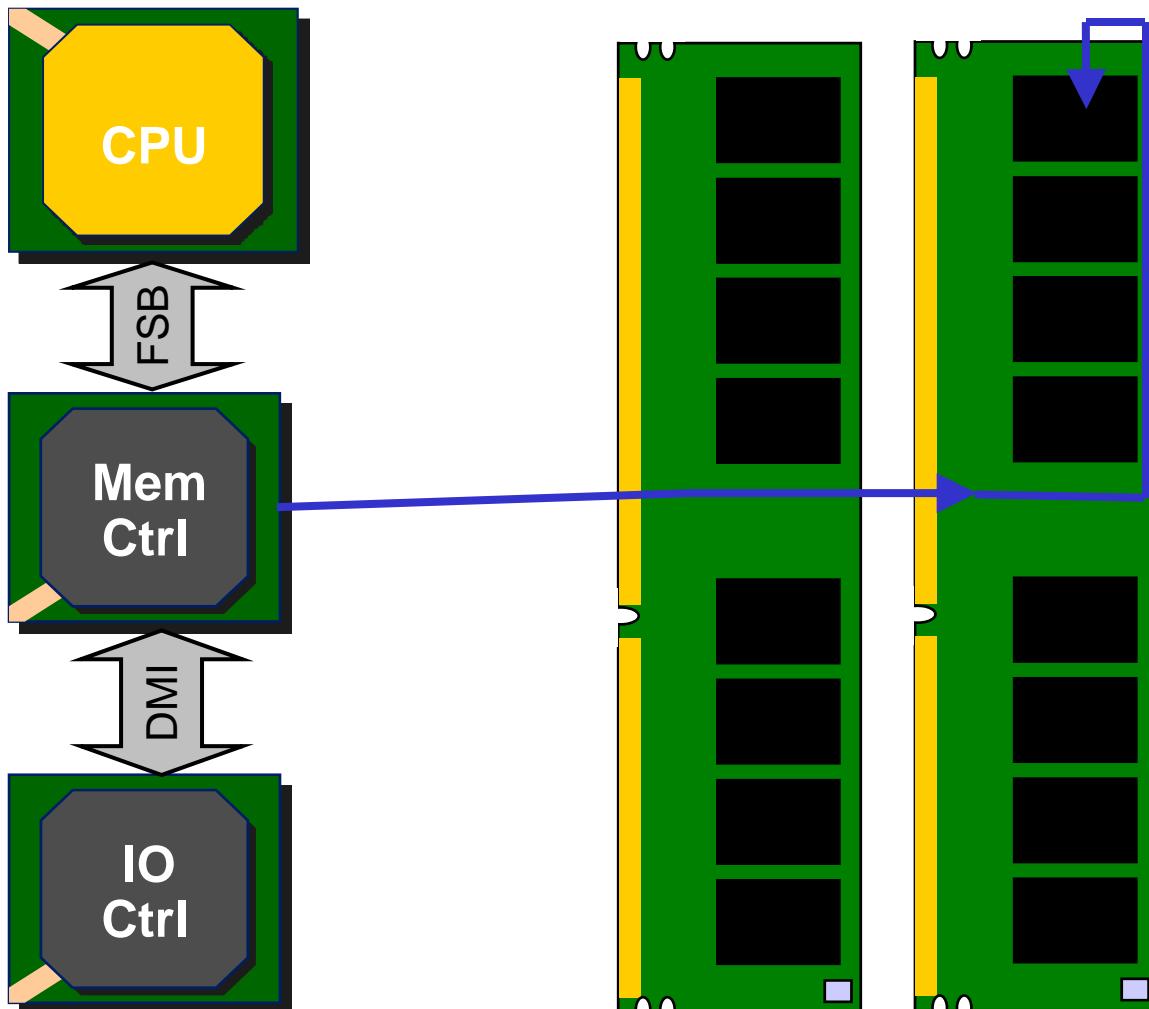


- **Read Cycle**
- **Read command is sent to the DIMM over the address and command bus.**

Key: Red is Data Bus

Blue is Address and Command Bus

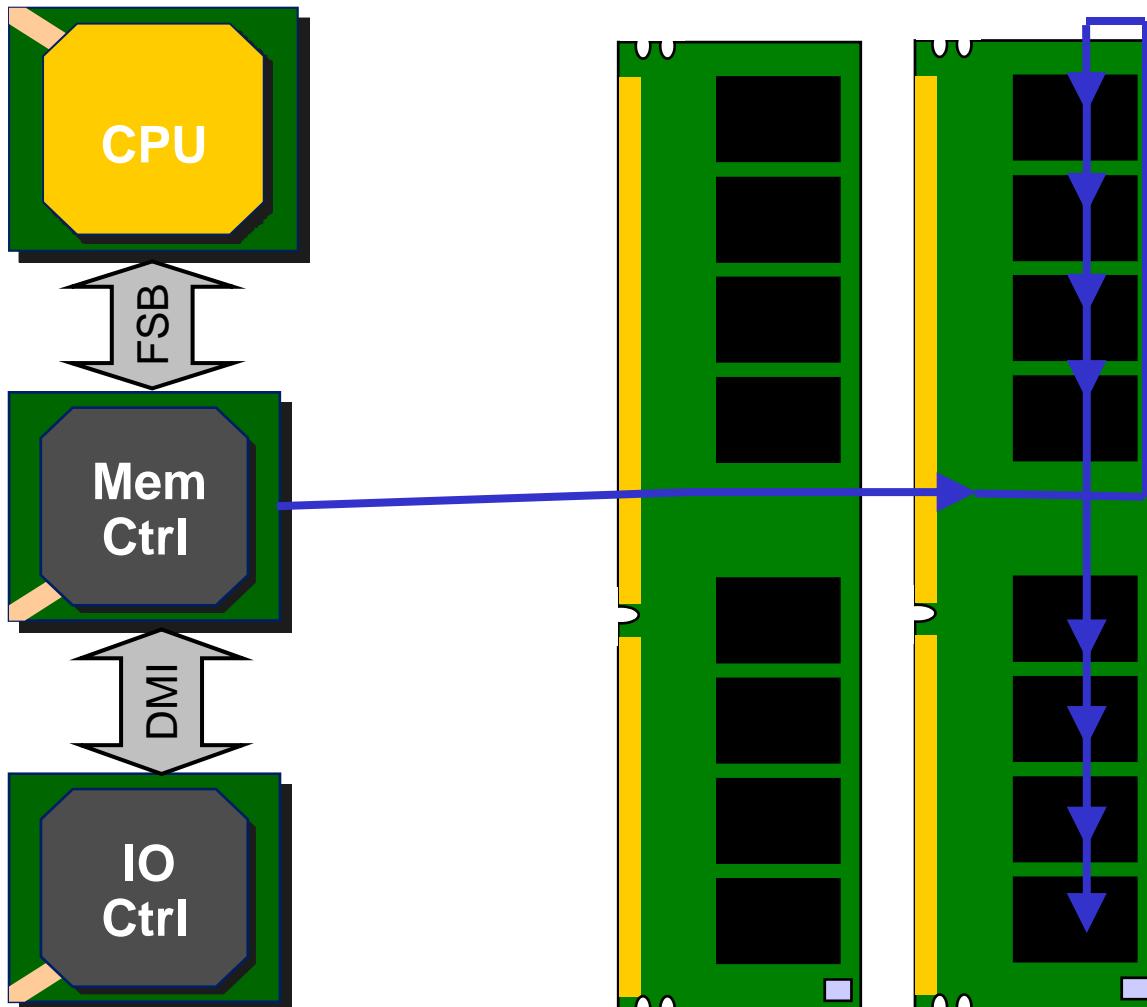
Min Huang/min.huang@leecroy.com)
Do Not Distribute



- Read Cycle
- Read command is sent to the DIMM over the address and command bus.
- The command gets to the first DRAM on the new fly-by routing chain.

Key: Red is Data Bus

Key: Red is Data Bus
Min Huang(min.huang@jrcroy.com)
Blue is Address and Command Bus
Do Not Distribute



- Read Cycle
- Read command is sent to the DIMM over the address and command bus.
- The command gets to the first DRAM on the new fly-by routing chain.
- The command propagates to the remaining DRAMs

Key: Red is Data Bus

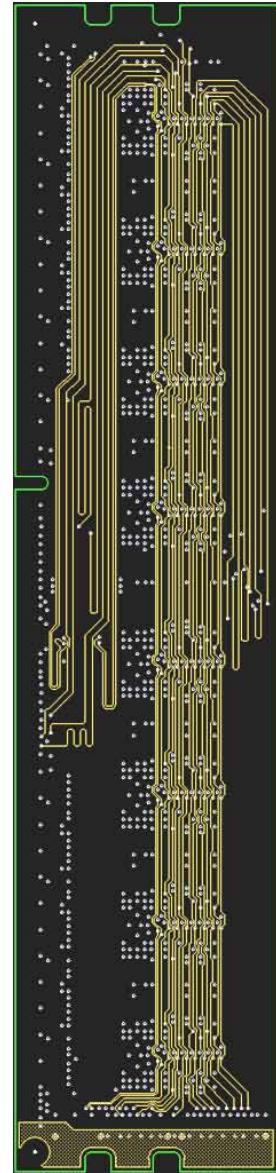
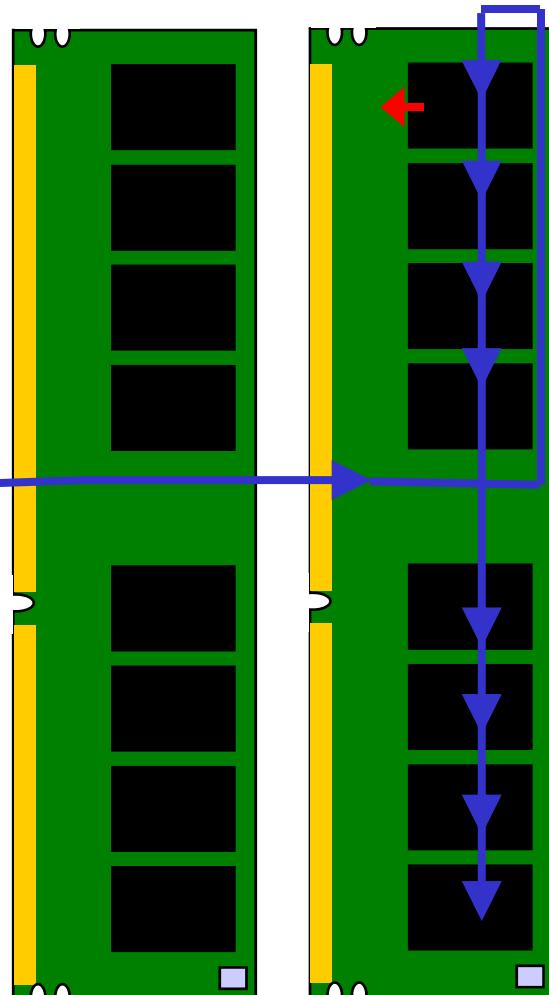
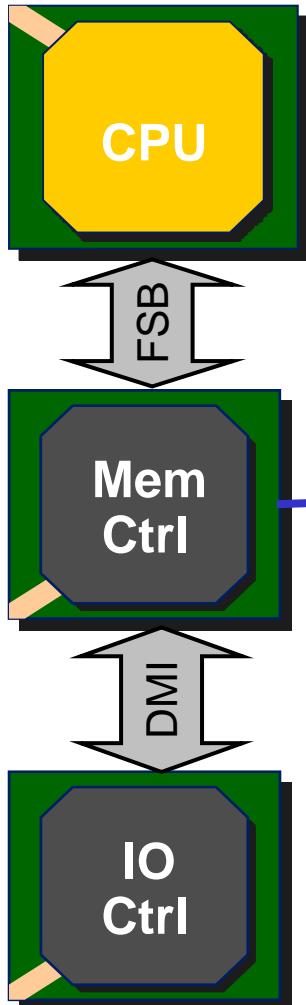
Blue is Address and Command Bus

Min Huang (min.huang@leecroy.com)
Do Not Distribute



DDR3 Read Overview

223

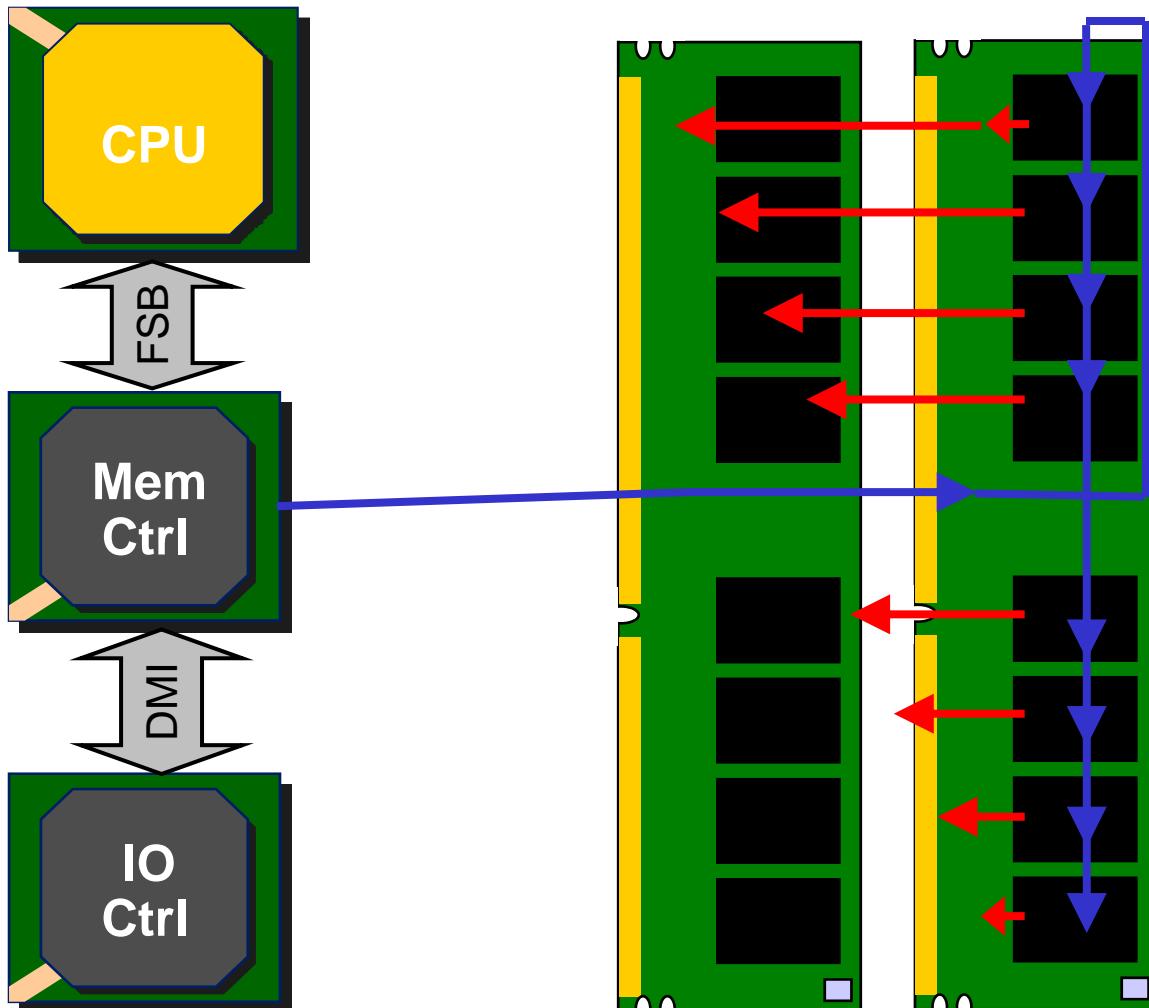


Key: Red is Data Bus

Blue is Address and Command Bus

Min Huang (min.huang@leecroy.com)
Do Not Distribute

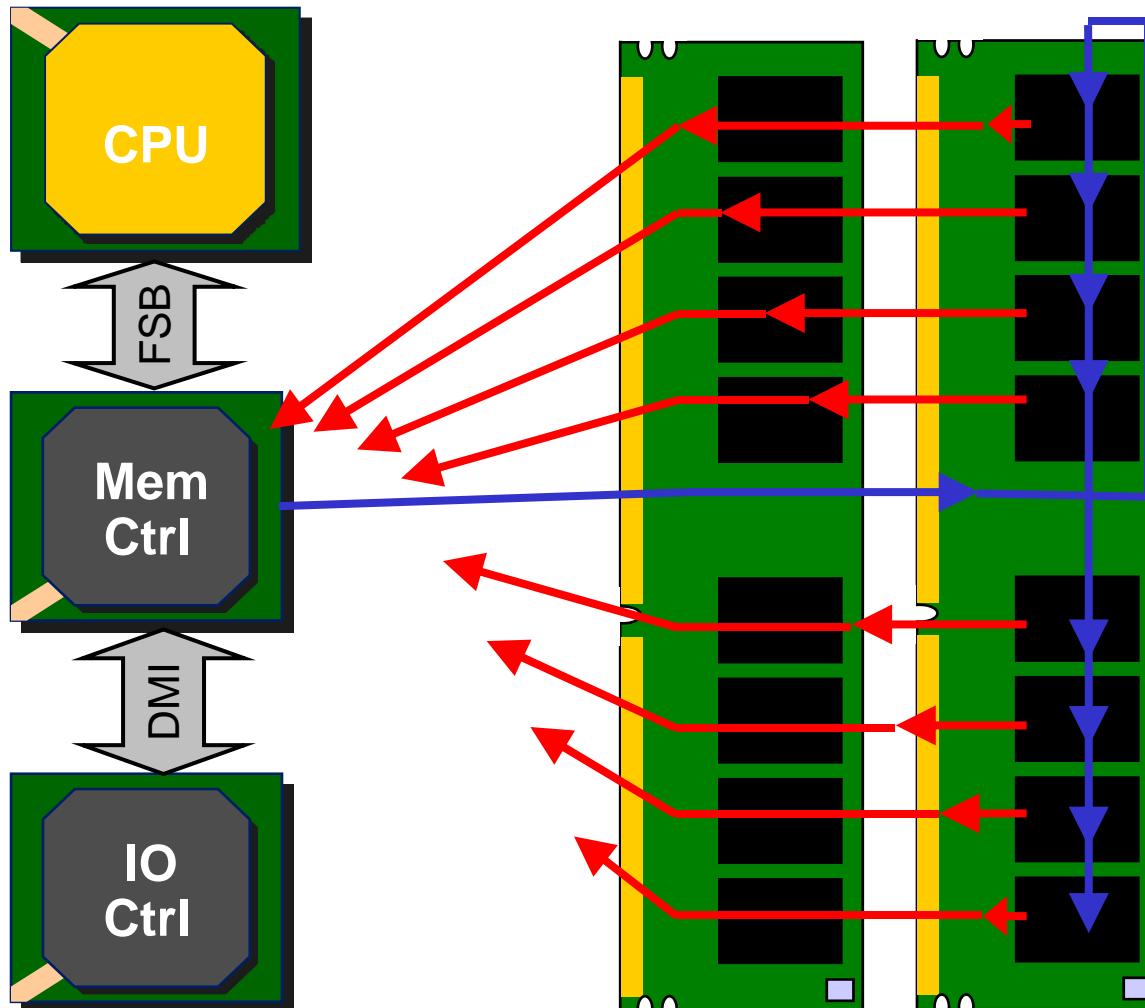
Termination on DIMM



- Read Cycle cont.
- After RL, data begins streaming back from the first DRAM.
- Data propagates from the remaining DRAMs in order that the command was received.

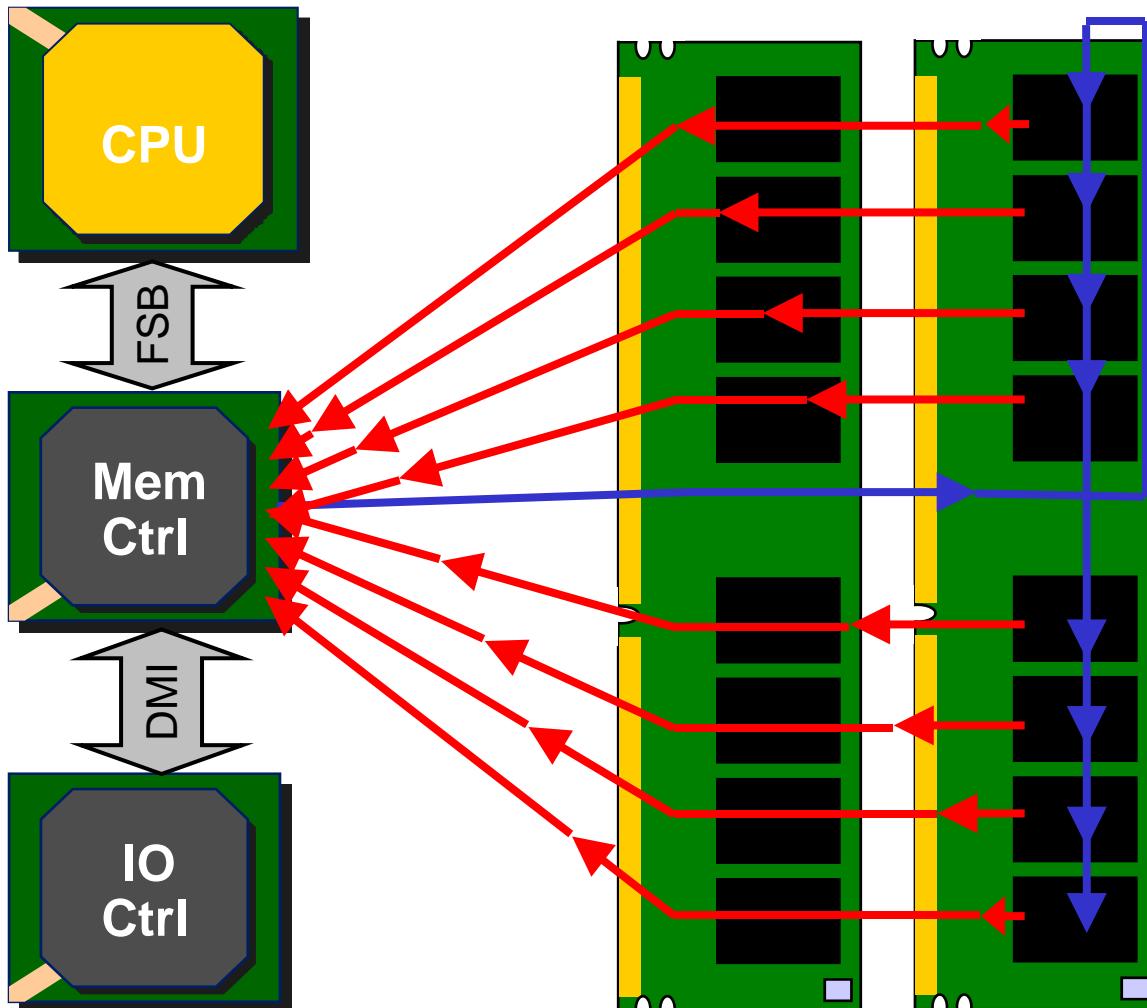
Key: Red is Data Bus

Min Huang/min.huang@leecroy.com)
Do Not Distribute



Key: Red is Data Bus
Blue is Address and Command Bus
Min Huang (min.huang@leecroy.com)
Do Not Distribute

- **Read Cycle cont.**
- Data propagates from the remaining DRAMs in order that the command was received.
- Read data from the first byte lane arrives.
- The remaining data lanes arrive one after the other.



Key: Red is Data Bus
Blue is Address and Command Bus
Min Huang (min.huang@jccrccy.com)
Do Not Distribute

- **Read Cycle cont.**
- Data propagates from the remaining DRAMs in order that the command was received.
- Read data from the first byte lane arrives.
- The remaining data lanes arrive one after the other.
- What is the problem with this?



Read Calibration

Min Huang(min.huang@ lecroy.com)
Do Not Distribute

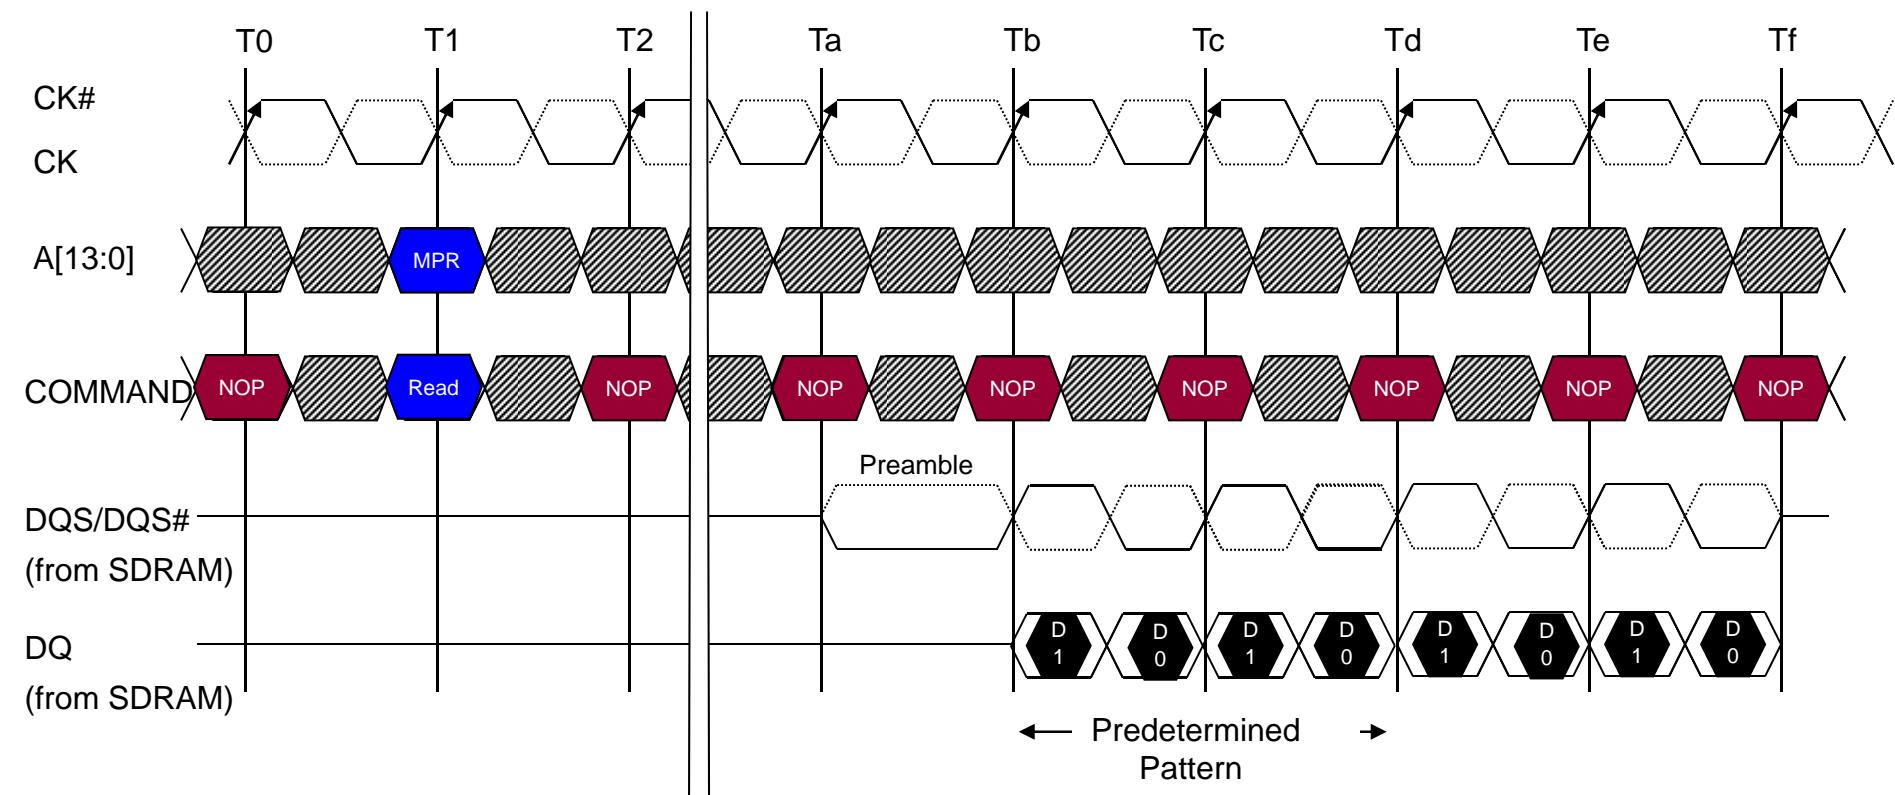


- To overcome this skewing of data the memory controller must calibrate the read data.
- Most of the Read calibration will be done by the memory controller.
- There are predetermined patterns built into the DRAM's Multi-Purpose Register (MPR) to facilitate the calibration. The read de-skewing must be done on all byte lanes of every rank every time the system reboots.



Read Calibration

During read calibration the controller initiates a predetermined pattern via MR3. The pattern can continue as long as the controller needs to calibrate.



MPR is the new multi-purpose register

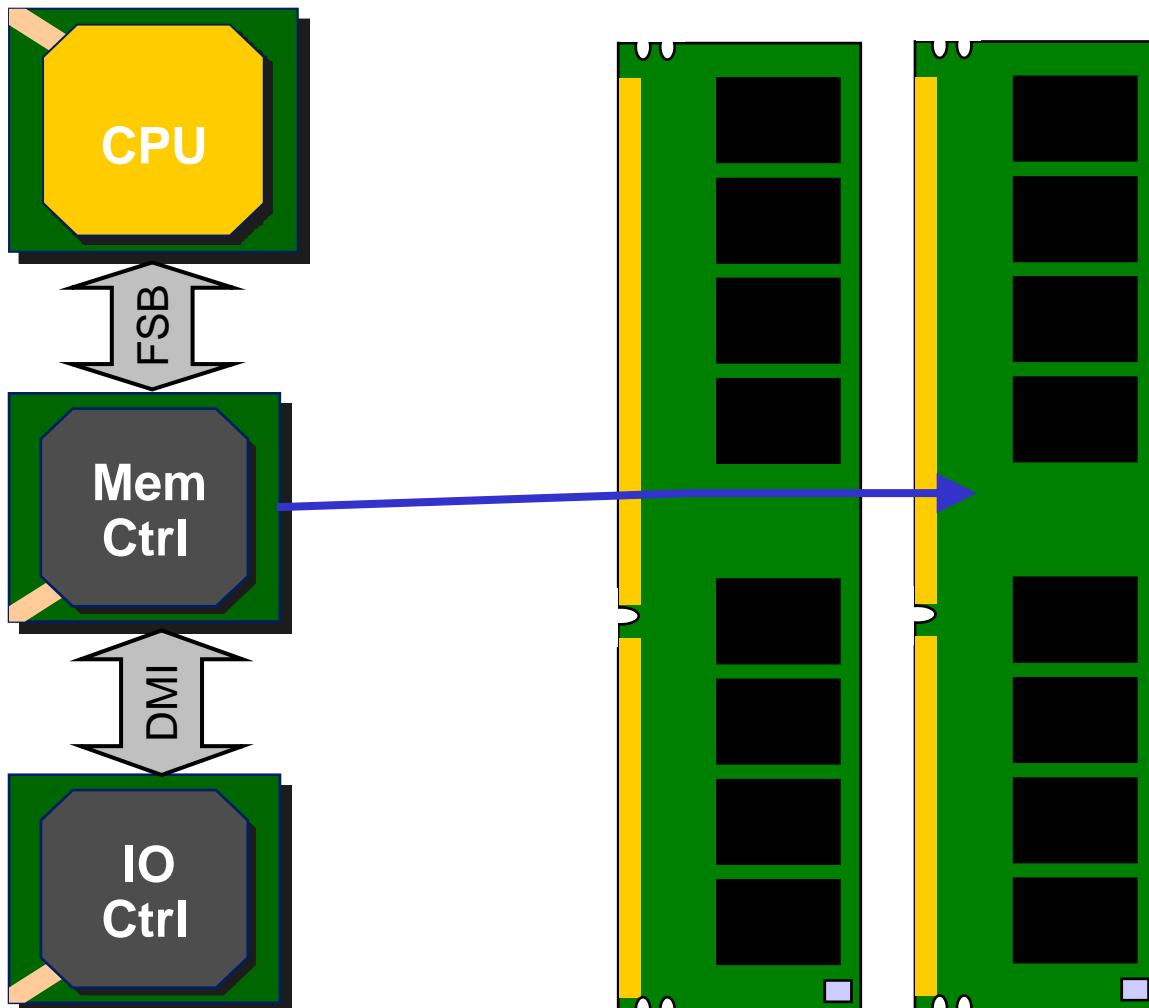
Min Huang(min.huang@ lecroy.com)
Do Not Distribute



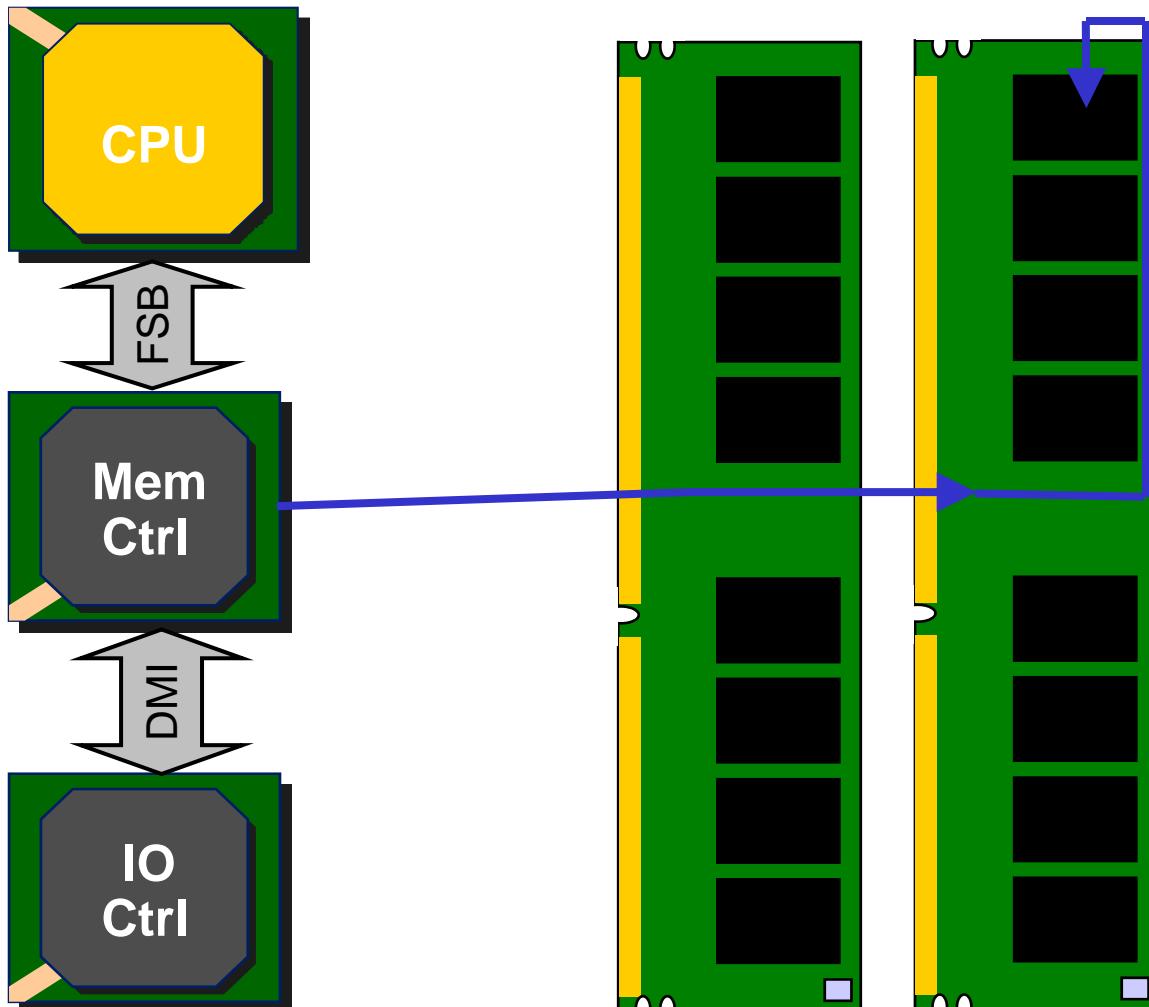
Fly-by Routing Write Example

Min Huang(min.huang@ lecroy.com)
Do Not Distribute

- As in DDR1 and DDR2, DDR3 SDRAMs require a certain phase relation between the strobe signals and the clock signals during writes.
 - This results in a certain write data launch time requirement at the memory controller.
- In DDR1 and DDR2, the write data launch time is equal for all byte lanes of a DIMM.
 - This is achieved via flight time length matching on the mother board and on the DIMM.
 - Depending on the channel routing, it may even be equal among 2 DIMMs.
- In DDR3, the write data launch time will be different
 - across the byte lanes for a single DIMM
 - from one DIMM to another DIMM



- **Write Cycle**
- Write command is sent to the DIMM over the address and command bus.

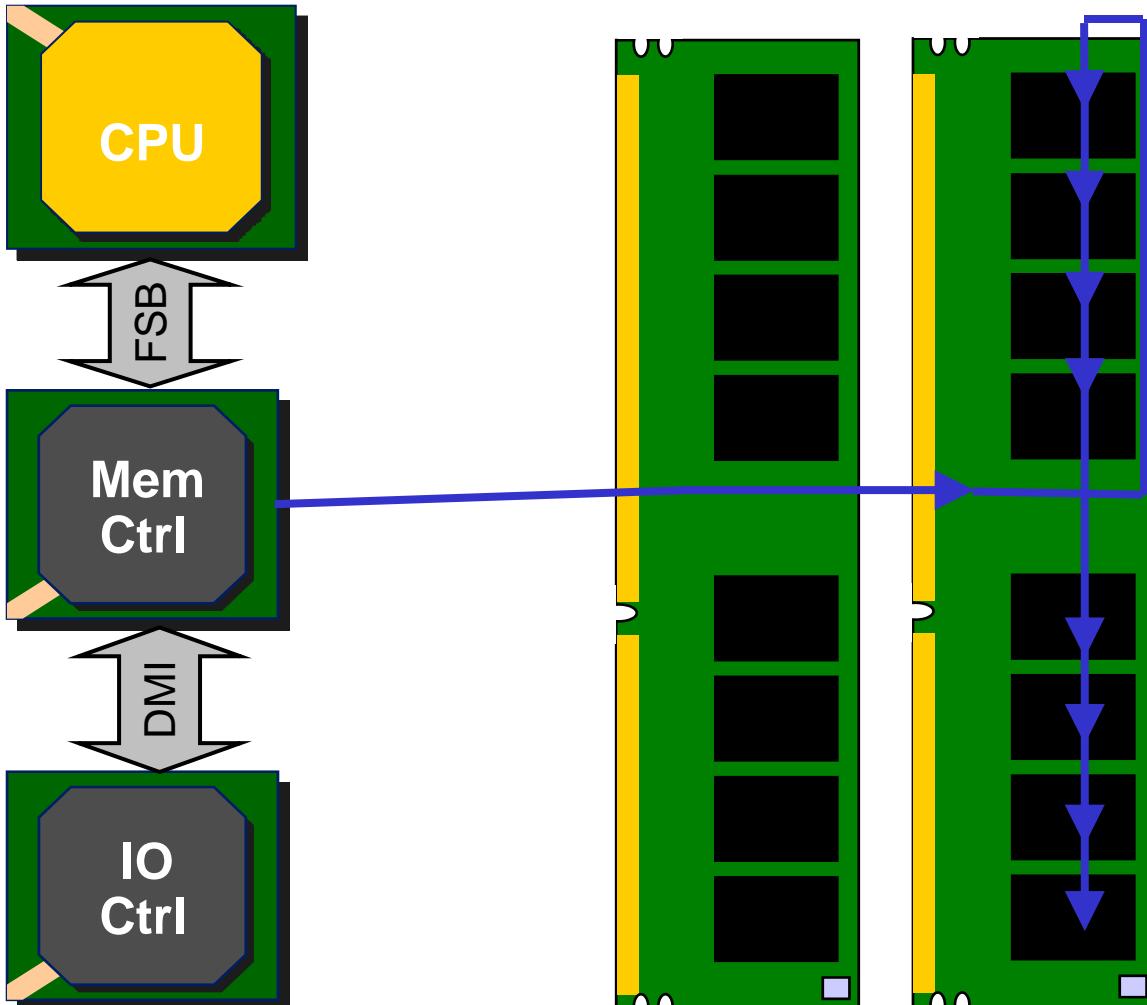


- **Write Cycle**
- Write command is sent to the DIMM over the address and command bus.
- Command reaches the first DRAM due the fly-by routing.

Key: Red is Data Bus

Blue is Address and Command Bus

Min Huang (min.huang@leecroy.com)
Do Not Distribute

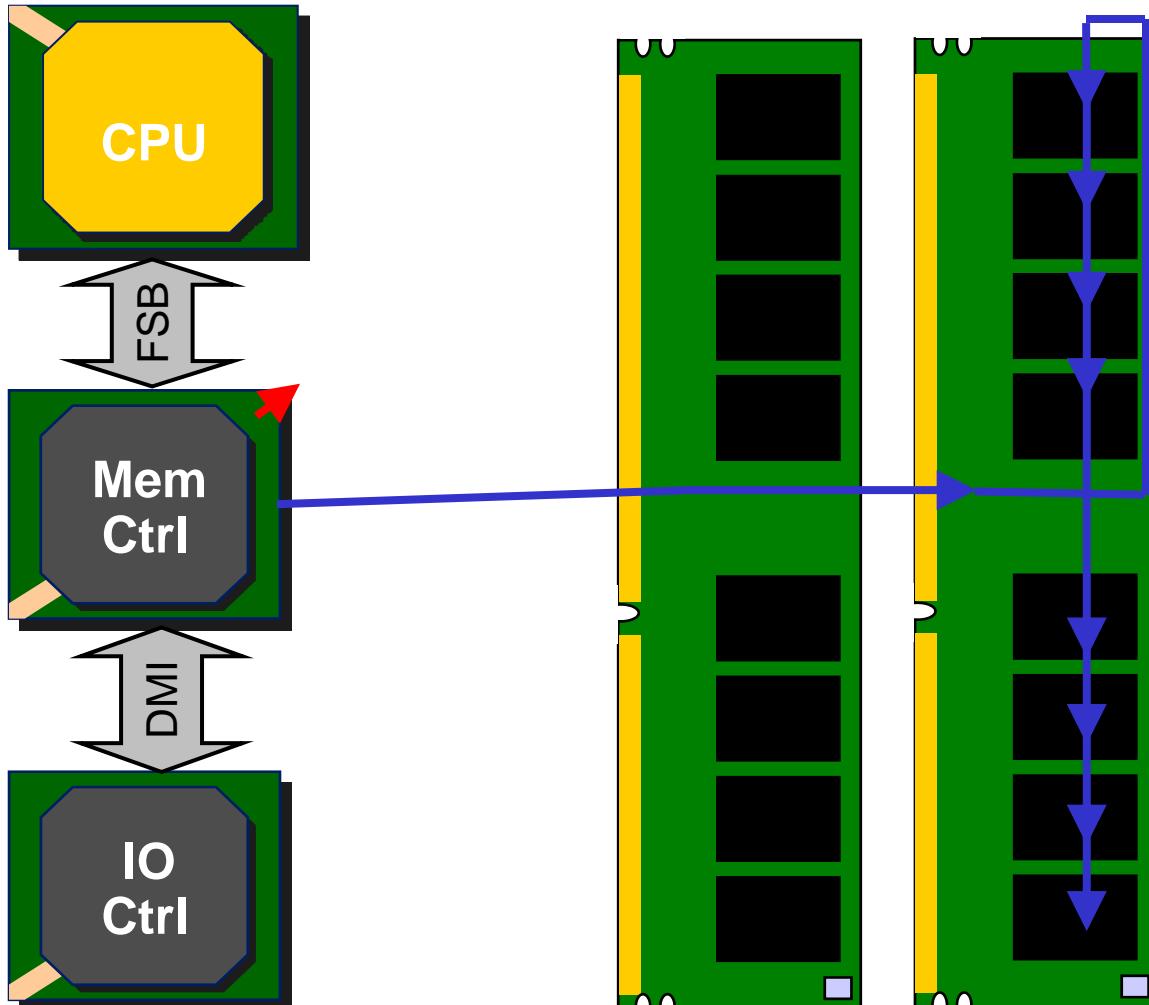


- **Write Cycle**
- Write command is sent to the DIMM over the address and command bus.
- Command reaches the first DRAM due the fly-by routing.
- Command propagates to the remaining DRAMs.

Key: Red is Data Bus

Blue is Address and Command Bus

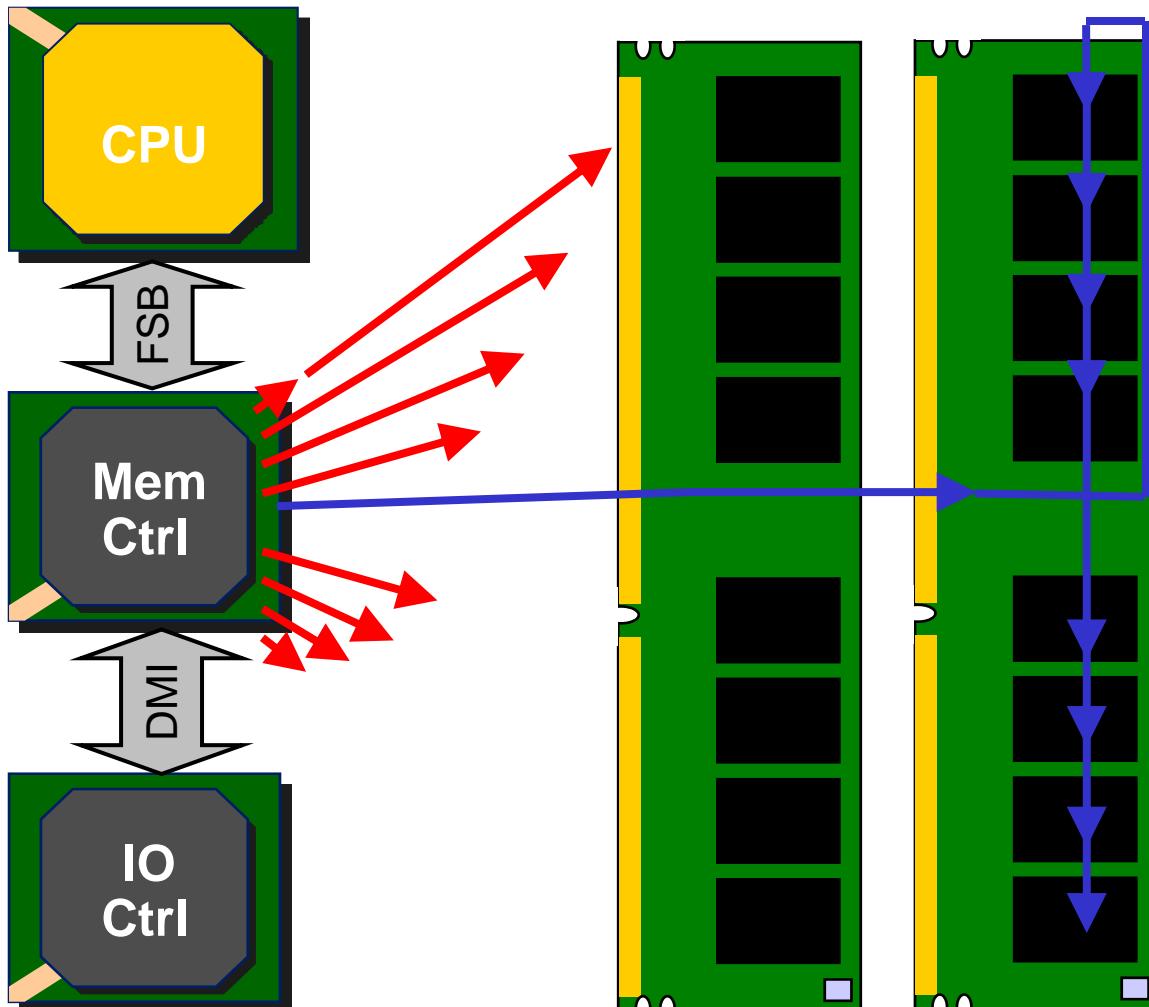
Min Huang (min.huang@leecroy.com)
Do Not Distribute



- **Write Cycle**
- Write command is sent to the DIMM over the address and command bus.
- Command reaches the first DRAM due to the fly-by routing.
- Command propagates to the remaining DRAMs.
- After WL, the controller starts driving the first data lane.

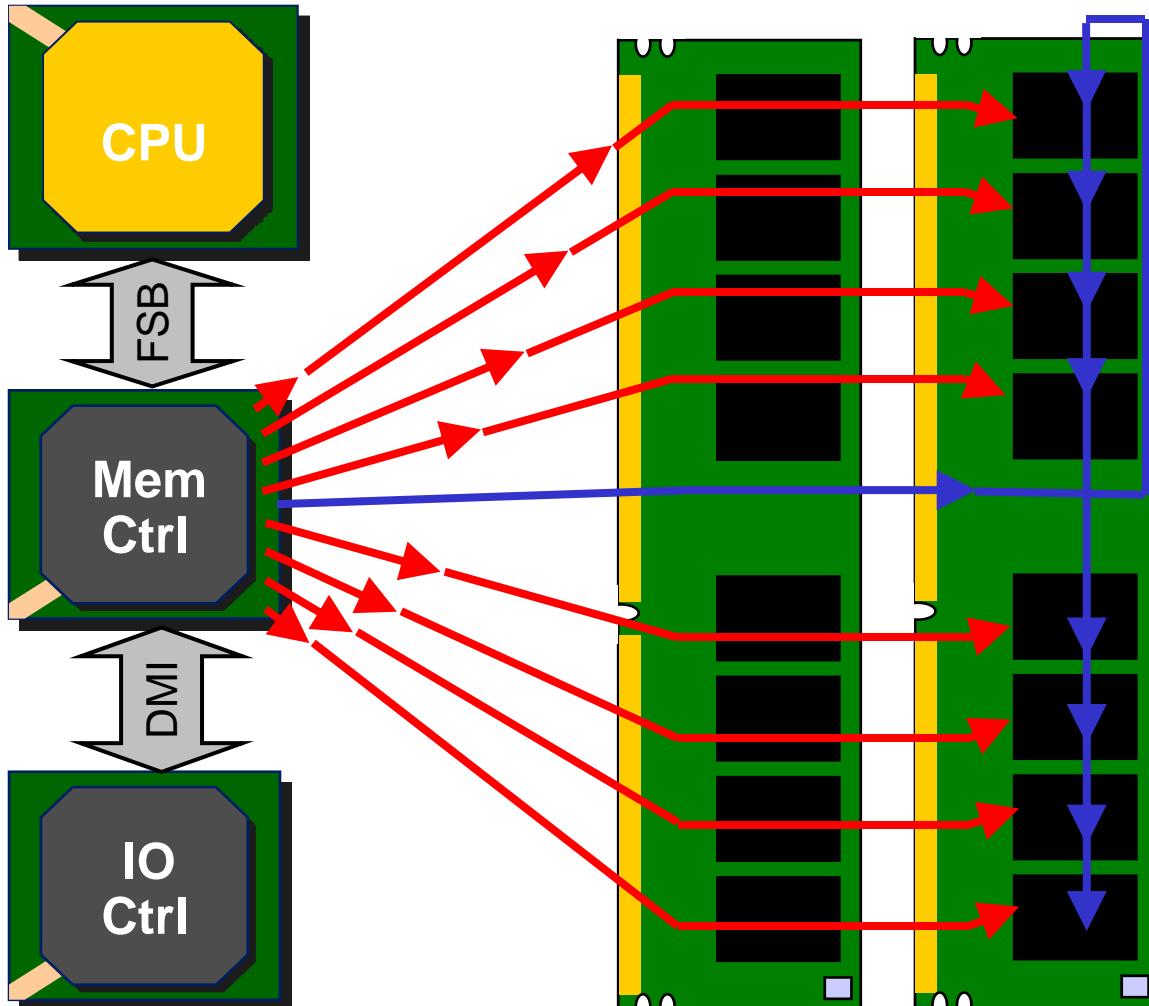
Key: Red is Data Bus

Min Huang (min.huang@leecroy.com)
Do Not Distribute



Key: Red is Data Bus
Blue is Address and Command Bus
Min Huang (min.huang@leecroy.com)
Do Not Distribute

- **Write Cycle**
- Write command is sent to the DIMM over the address and command bus.
- Command reaches the first DRAM due to the fly-by routing.
- Command propagates to the remaining DRAMs.
- After WL, the controller starts driving the first data lane.
- The other data lanes follow in order.



Key: Red is Data Bus
Blue is Address and Command Bus
Min Huang (min.huang@leecroy.com)
Do Not Distribute

- **Write Cycle**
- Write command is sent to the DIMM over the address and command bus.
- Command reaches the first DRAM due to the fly-by routing.
- Command propagates to the remaining DRAMs.
- After WL, the controller starts driving the first data lane.
- The other data lanes follow in order.
- Data does not arrive at the same time to all DRAMs.

- The write data launch time at the controller varies
 - from byte lane to byte lane
 - from DIMM to DIMM
 - from system to system.
- Therefore, it must be calibrated at the memory controller.

- How can this calibration be done in a DDR3 System?
- How does the DDR3 SDRAM support this calibration?

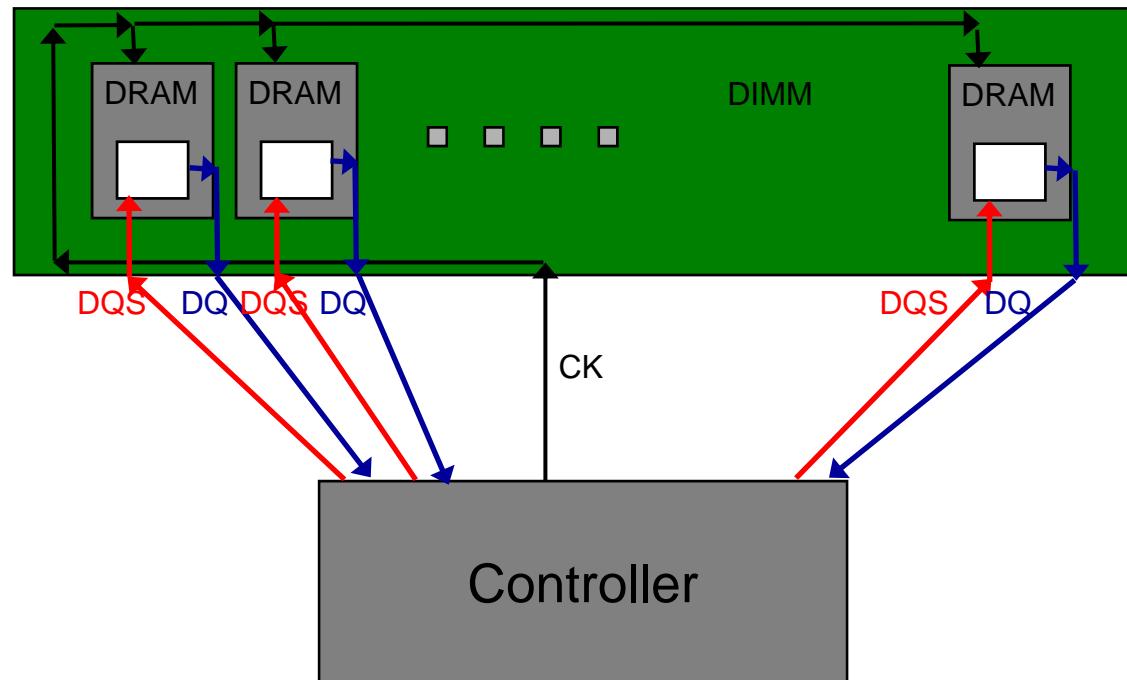


Write Leveling

Min Huang(min.huang@ lecroy.com)
Do Not Distribute



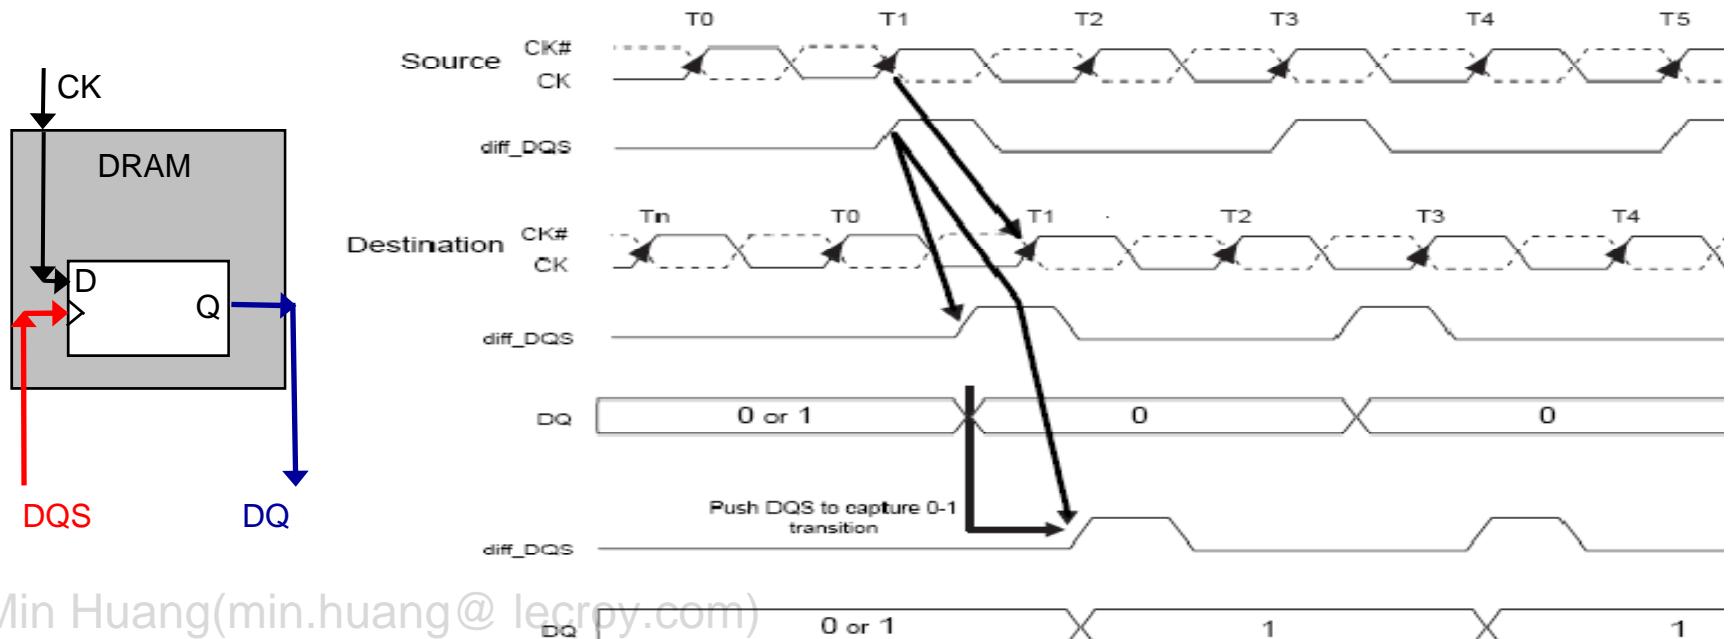
If each data lane is being launched at different times how can the controller know when to launch them?



By calibrating each byte lane on every DIMM utilizing a phase detector available inside the DDR3 SDRAM during "Write Leveling Mode"



- The phase detector on the DDR3 SDRAM can be understood as a flip-flop:
 - DRAM samples the clock status with rising edge of the strobe
 - and provides the sample result on the DQ pins after an asynchronous delay of t_{WLO} .
 - No strobe information is driven out with the DQ information since the strobes are used as trigger input for the flip-flop.
- Like every flip-flop, the phase detector flip-flop on the DDR3 SDRAMs has a requirement for minimum setup and hold times.
- Besides the pure flip-flop timing requirements, the imperfect delay matching between clock pads to flip-flop and strobe pads to flip-flop imposes additional timing constraints.
- As a result, the clock signals must be stable.





- In DDR2, the write data launch time is equal for all byte lanes of a DIMM, sometimes even among two DIMMs within a channel
 - This is achieved through flight-time length matching on the mother board and on the DIMM.
- In DDR3, this is completely different due to the fly-by command/address/control/clock bus topology on the DIMM:
 - The write data launch time is different across the byte/nibble lanes of a DIMM.
 - The write data launch time is different from one DIMM to another DIMM.
- Therefore, a DDR3 memory controller must be able to calibrate the launch time for every byte/nibble lane for every DIMM
 - The recommended timing resolution at the memory controller is $1/16 t_{CK}$
 - The launch times may be spread over the boundaries of a clock cycle at the memory controller.
 - A DDR3 memory controller
 - must store different launch time settings per byte/nibble lane and per slot
 - must dynamically switch between slot settings!



DDR3 On-Die Termination

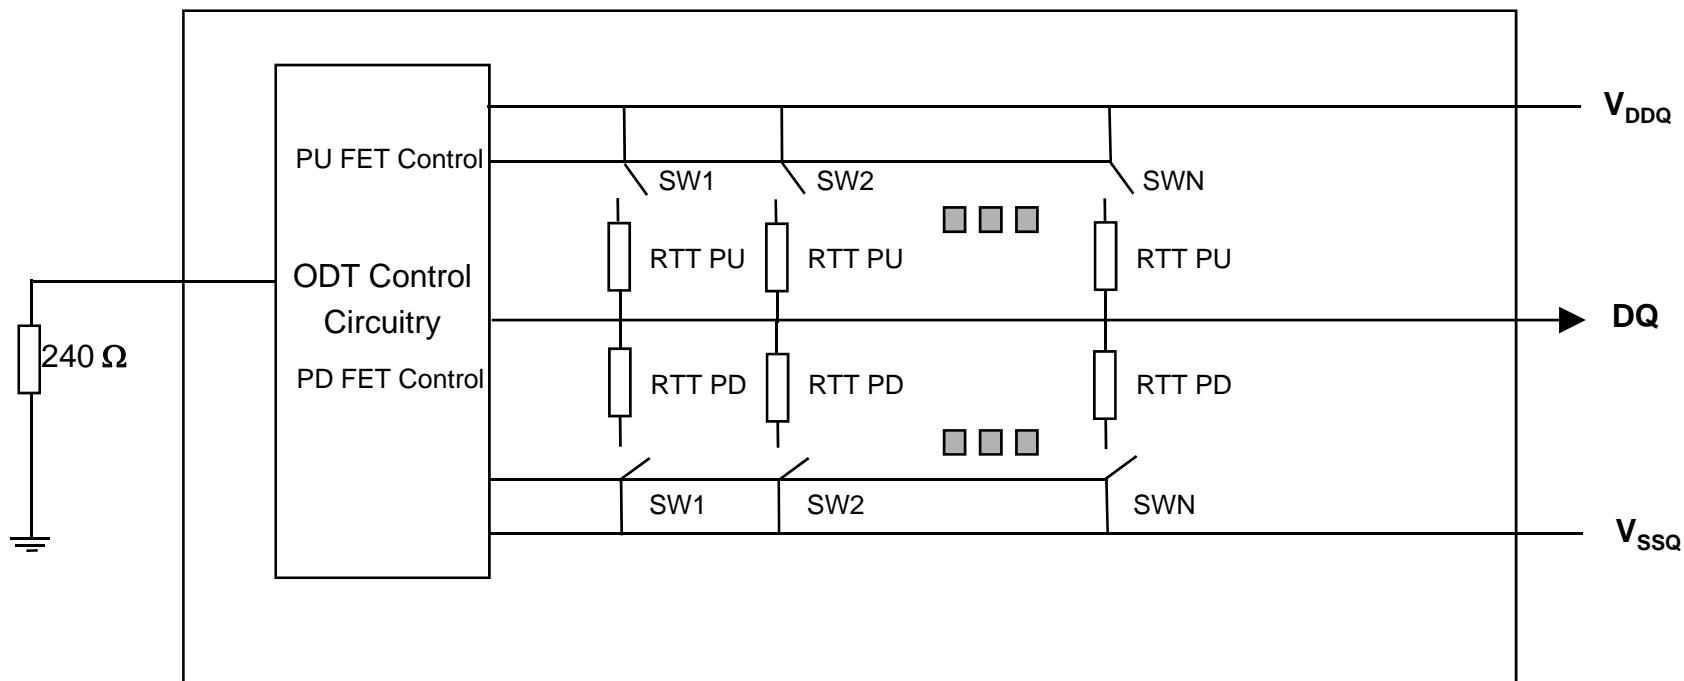
Min Huang(min.huang@ lecroy.com)
Do Not Distribute



- ODT is designed to improve signal integrity of the memory channel by allowing the memory controller to independently turn on/off termination resistance at the DQ bus interface for any or all SDRAM devices via the ODT pin.
- ODT is implemented on the DDR3 SDRAM as selectable, center-tapped termination resistance on the following DQ bus pins:
 - x16 DRAMs: DQU, DQL, DQSU, DQSU#, DQSL, DQSL#, DMU, DML
 - x4 and x8 DRAMs: DQ, DQS, DQS#, DM
 - and TDQS, TDQS# for x8 DRAMs, when enabled via A11=1 in MR1



- Conceptual Circuit Diagram
- ODT can be compensated using an external resistor
- It is not clearly stated if actual silicon resistors are used or it is the R_{on} of the Pull Up and Pull Down transistors.

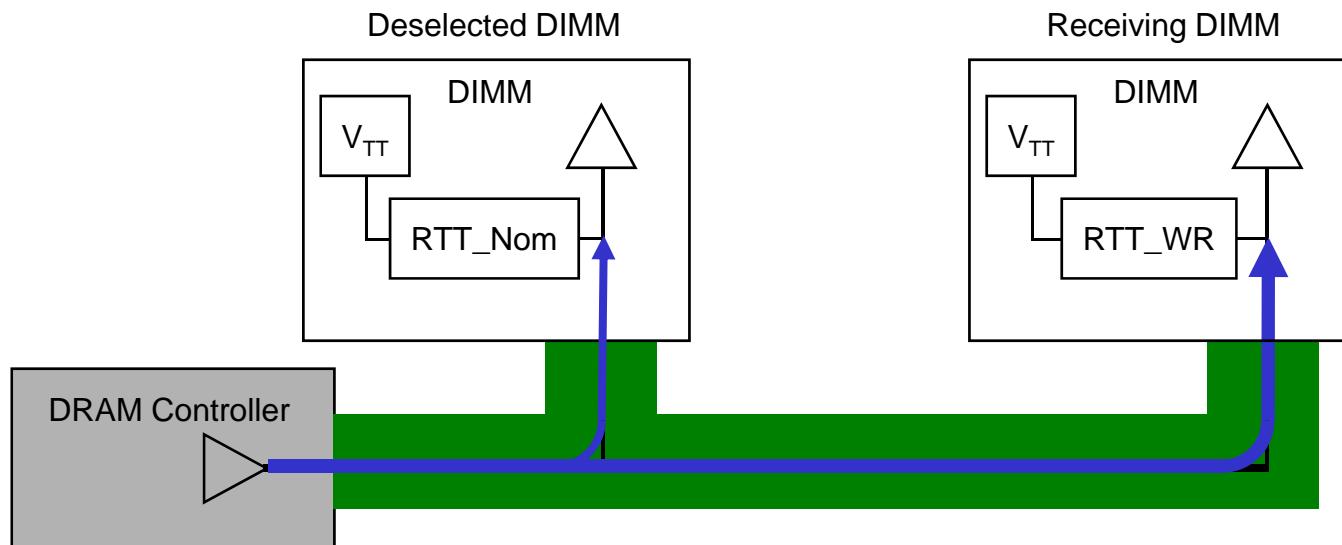




- The following ODT modes are available:
 - Synchronous ODT
 - Selected whenever the DLL is turned on and locked:
 - Active mode
 - Idle mode with CKE high
 - Active power down mode (regardless of MR0 bit A12)
 - Precharge power down mode if DLL is enabled during precharge power down by MR0 bit A12.
 - Asynchronous ODT
 - Selected when DRAM runs in DLL on mode, but DLL is temporarily disabled.
 - Precharge power-down with slow exit (selected by MR0 bit A12).



- Write to furthest slot from Controller
- RTT_Nom 20 Ω
- RTT_WR 120 Ω





- In dual slot DDR3 systems, the slot which is not accessed during a read or write transaction needs to terminate the bus with low impedance (e.g. 20 or 30 Ohms).
- When writing to the slot which is configured to terminate with low impedance, a higher impedance is required to achieve open data eye.
- Thus, not only two termination options are needed (RTT on/off), but three options must be available without MRS interaction:
 - RTT turned off
 - RTT = Low impedance during reads/writes from/to the deselected slot
 - RTT = High impedance during writes to the active slot



Mode Register 1 (MR1)				
A9	A6	A2	RTT_Nom (nominal)	RTT_Nom if RZQ=240 Ω
0	0	0	disable (also dyn. ODT)	off
0	0	1	RZQ / 4	60
0	1	0	RZQ / 2	120
0	1	1	RZQ / 6	40
1	0	0	RZQ / 12	20
1	0	1	RZQ / 8	30
1	1	0	RFU	RFU
1	1	1	RFU	RFU

- The values 1-3 can only be used for the deselected Ranks during writes.
- Values 4 and 5 can be used on the deselected Ranks during reads (although not specifically stated by JEDEC)



Mode Register 2 (MR2)			
A10	A9	RTT_WR (during WR)	RTT_WR if RZQ=240 Ω
0	0	Dynamic ODT off: Write does not affect RTT value	
0	1	RZQ / 4	60
1	0	RZQ / 2	120
1	1	RFU	RFU

- The values 1 and 0 are used for the selected Rank during writes only.



ZQ Calibration (OCD)

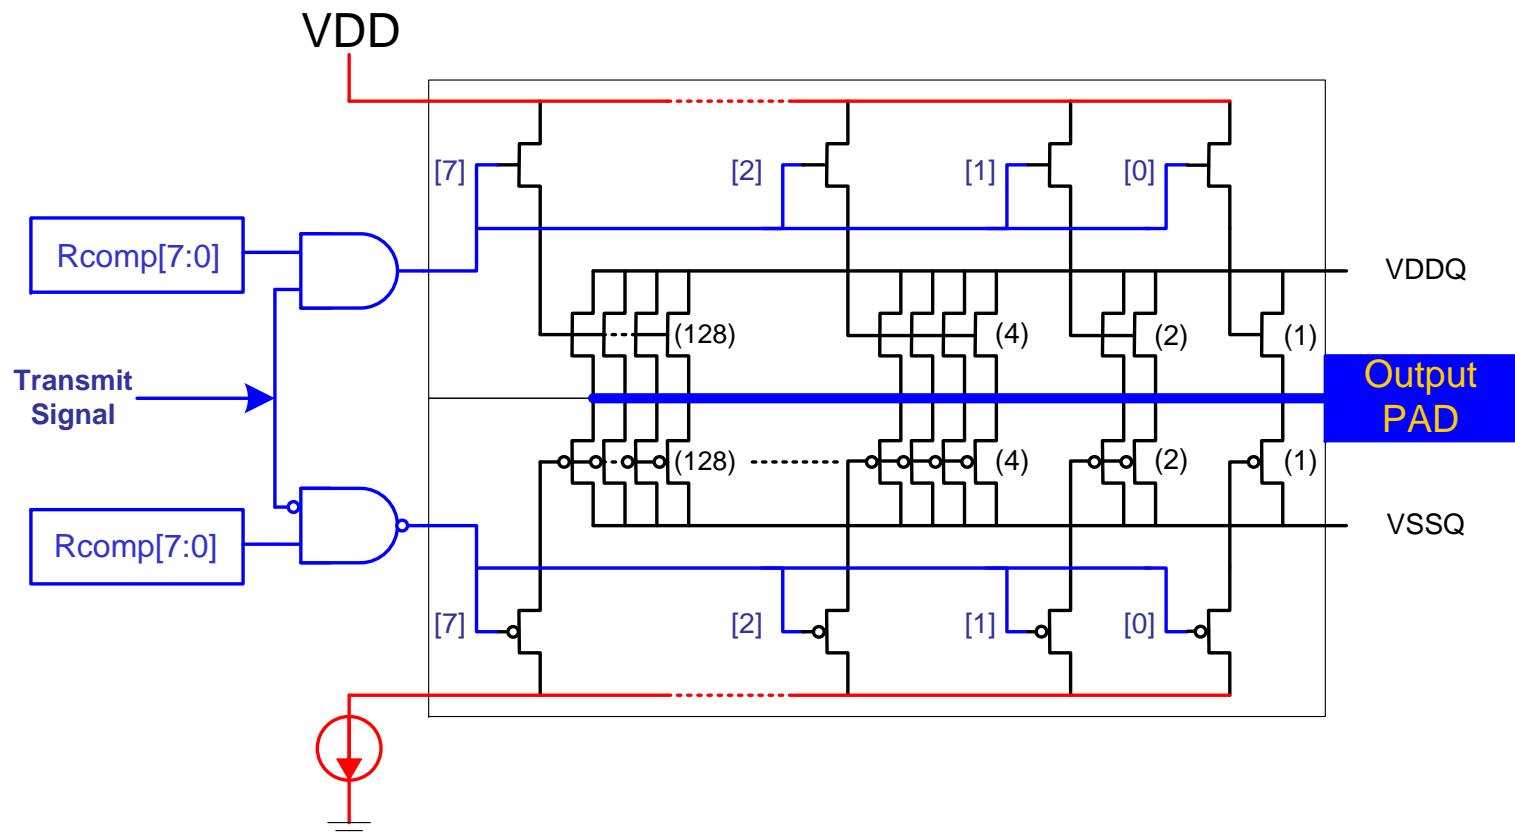
Min Huang(min.huang@ lecroy.com)
Do Not Distribute



- The Off-Chip Driver (OCD) calibration protocol for DDR2 was not used because the granularity of the drive strength was too large.
- Both OCD and ODT are calibrated and adjusted against an external reference resistor connected to the ZQ pin of DDR3 chips.
- Calibration is triggered by a calibration command from the memory controller.



ZQ calibration is done after reset and is recalibrated periodically to compensate for voltage and temperature fluctuations. DRAM vendors recommend that calibration be done every 128mS. Below is an example of an 8-bit Rcomp binary-legged circuit.





Reset

Min Huang(min.huang@ lecroy.com)
Do Not Distribute



- DDR3 DRAM implements a RESET# pin which is an asynchronous external control signal of the DDR3 SDRAM.
- This RESET# signal forces DRAM from operational or non-operational conditions into a defined state.
- RESET# can be applied asynchronously to an on going DRAM operation and can be asserted at any time.
- After RESET# is asserted, data in DRAM may be lost (though RESET# is **not** designed to destroy the data) and DRAM needs to be re-initialized, which includes (but not limited to) load mode registers and DLL reset.
 - Note for Application: In order to maintain data inside the DDR3 SDRAMs, the RESET# signal must not be applied upon EXIT from S3 (suspend to RAM).
 - PCI RST# should not be used.
- It is mandatory for system to hold RESET# asserted at beginning of power ramp until power supplies reach stability.
- Minimum pulse width of RESET# at DRAM input:
 - At power ramp up: min. 200 μ s
 - After power ramp up, during stable supply conditions: min. 100 ns
 - Since there may be many DRAMs connected to the RESET# signal, system needs to generate a pulse with a larger width to ensure minimum pulse width is achieved at DRAM RESET# input.

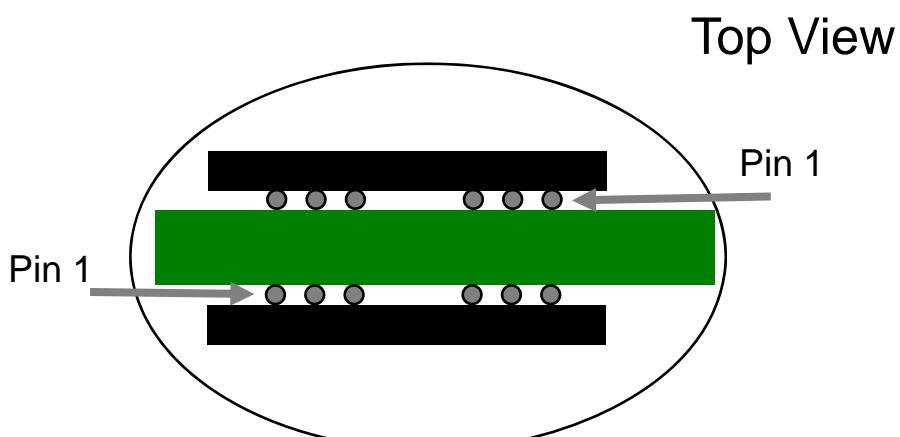
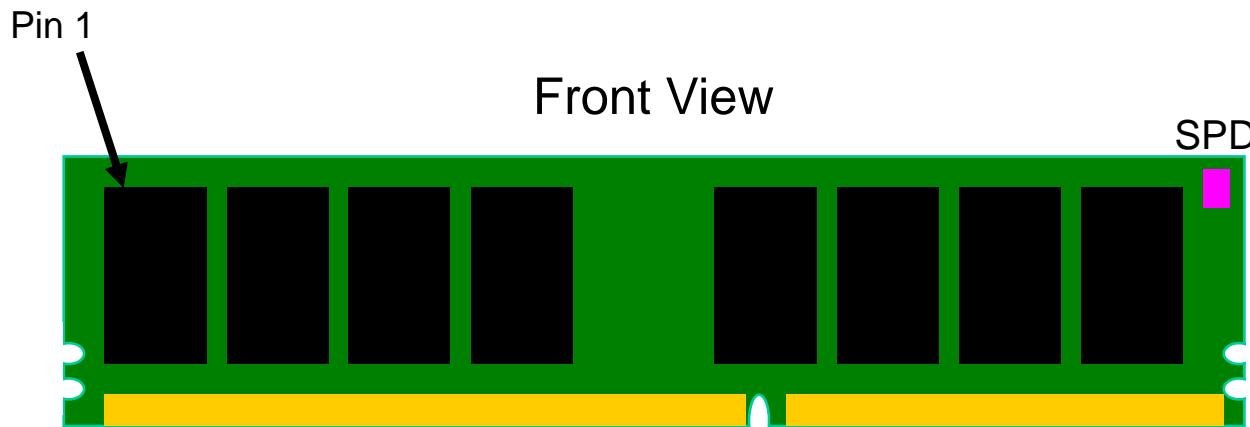


On-DIMM Address Mirroring

Min Huang(min.huang@ lecroy.com)
Do Not Distribute



➤ What is the motivation?

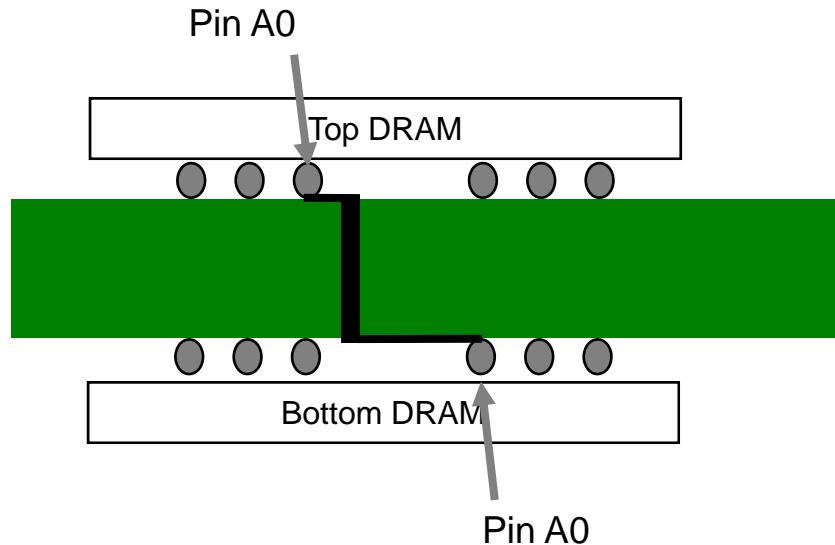




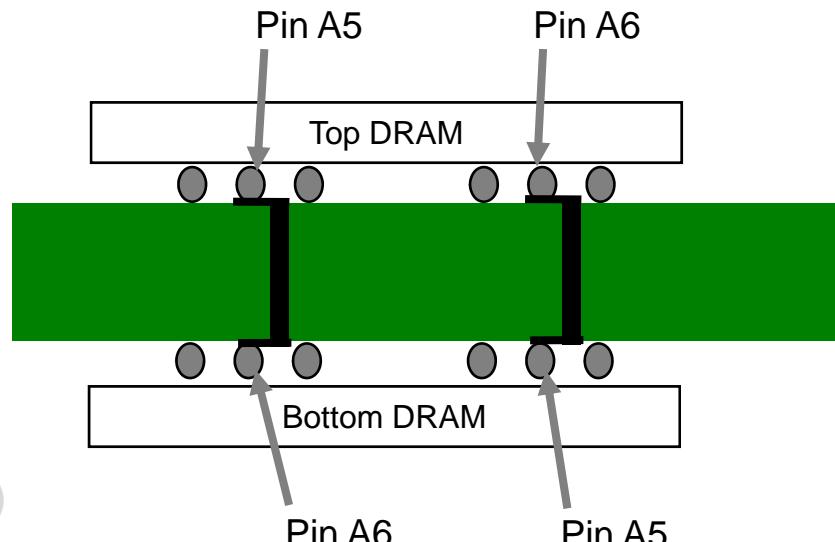
On-DIMM Address Mirroring

258

Problem: Traditional DIMM has longer stub and complex routing due to un-mirrored top and bottom ball-out



Solution: Mirrored DIMM has simpler routing that decreases stub length





- Only these pins may be mirrored:
 - BA0 and BA1
 - A3 and A4
 - A5 and A6
 - A7 and A8
- Controller must be told whether a DRAM is mirrored so that it can access the Mode Registers correctly. SPD contains this information. DDR3 SPD spec was released June, 2008.

	1	2	3	4	5	6	7	8	9
A	VSS	VDD	NC				NC	VSS	VDD
B	VSS	VSSQ	DQ0				DM	VSSQ	VDDQ
C	VDDQ	DQ2	DQS				DQ1	DQ3	VSSQ
D	VSSQ	NC	DQS#				VDD	VSS	VSSQ
E	VREFDQ	VDDQ	NC				NC	NC	VDDQ
F	NC	VSS	RAS#				CK	VSS	NC
G	ODT	VDD	CAS#				CK#	VDD	CKE
H	NC	CS#	WE#				A10/AP	ZQ	NC
J	VSS	BA0	BA2				A15	VREFCA	VSS
K	VDD	A3	A0				A12/BC#	BA1	VDD
L	VSS	A5	A2				A1	A4	VSS
M	VDD	A7	A9				A11	A6	VDD
N	VSS	RESET#	A13				A14	A8	VSS

**x4 DDR3 SDRAM,
looking through the package**



DRAM Controller Basics

Min Huang(min.huang@ lecroy.com)
Do Not Distribute



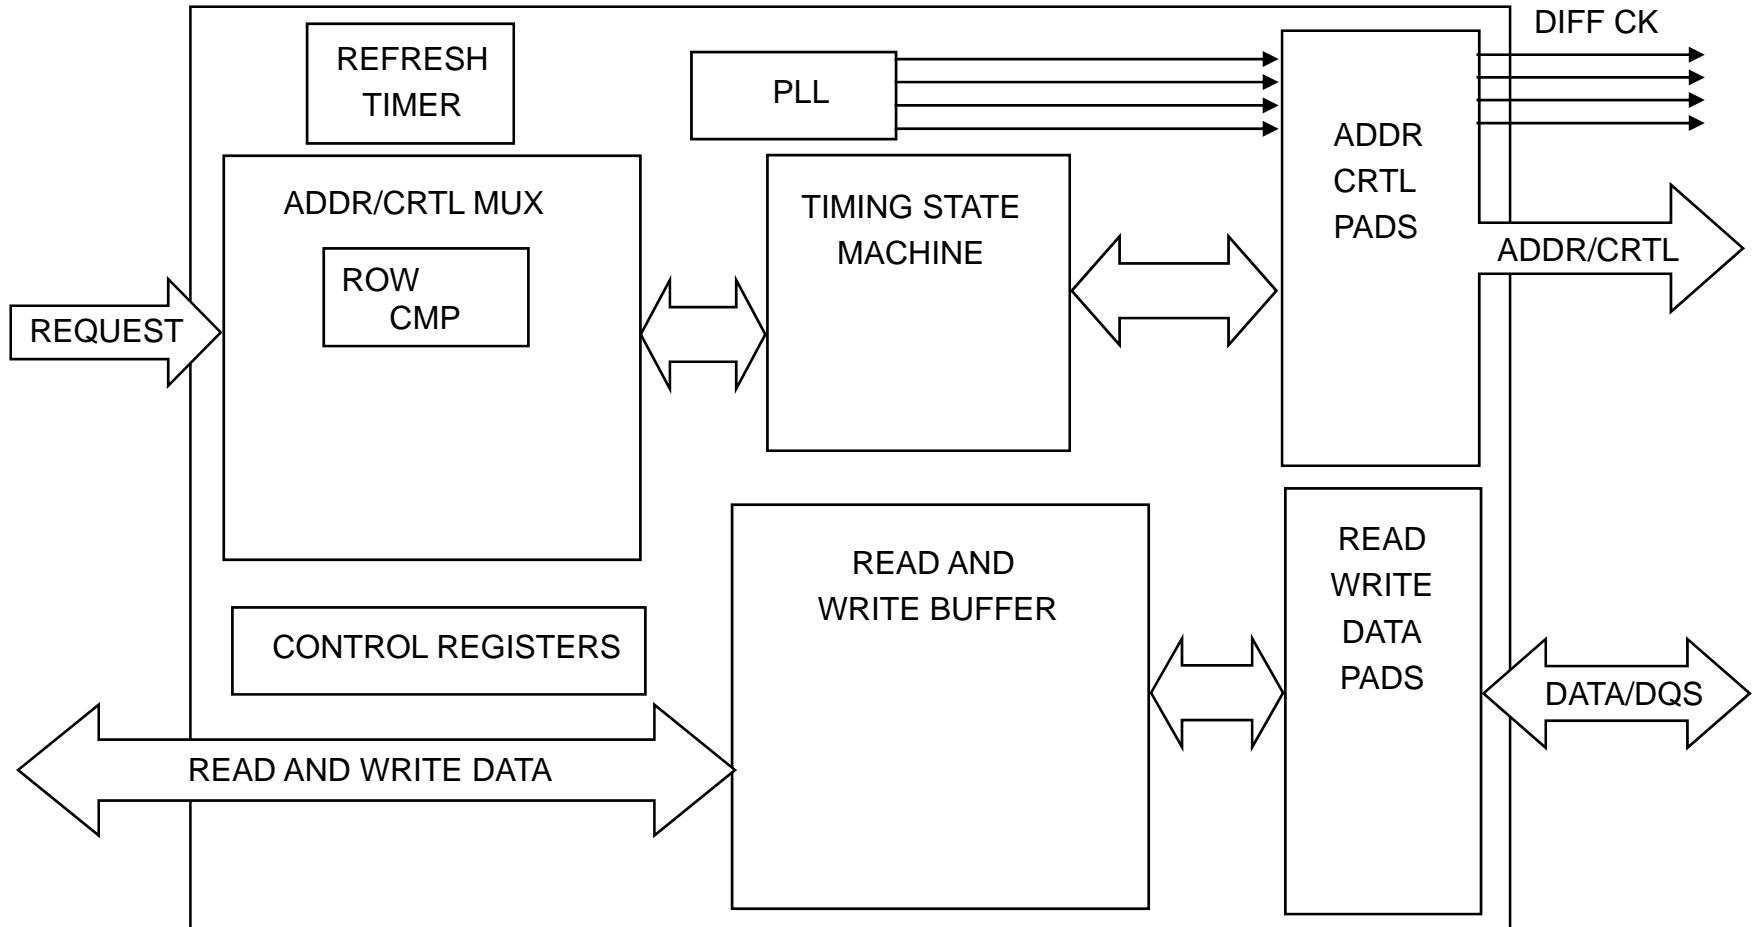
- Controller design is completely implementation-specific. JEDEC specifies AC timings and DC levels as well as implied DQS receive delay.
- There are some common building blocks that a controller must have and some blocks that are optional. The focus in the chapter will be on PC-based DIMM controllers.



- Address and Control Mux
- Refresh Timer
- PLL
- Timing Generator
- Control Register
- Read and Write Buffer
- IO Buffer Pads



DDR Controller Block Diagram





- This unit uses the physical address to create the Bank address, Row address and Column address.
- It also drives the correct CS# and generates the command sequence.
- This unit also does Row compare to see what Rows need to be open (active).

- The PLL receives an input clock from an external source and delivers multiple outputs to the functional units of the controller.
- The PLL makes multiple copies to send to the DIMM.
 - 3 copies for each Unbuffered DIMM
 - One copy for each Registered DIMM



- The Timing Generator connects directly to the IO Pads.
- This unit holds all of the timing parameters collected during initialization.
- Typically this unit would work with a calibration unit to optimize timing and turning the IO Pads on and off.



- These are PCI configuration registers as well as MMIO registers that most of the units access to gain information about the memory that is attached.
- BIOS programs these registers during initialization.



- The read and write queues hold data going into and out of memory.
- Typically this is a combining queue that can merge the changed bits with the existing bits and write them back to memory when bandwidth is available.
- Each buffer entry is a cache line in size.



- The IO Pads are the analog circuits used to drive the buses.
- Rcomp Scomp are found in these circuits.



- Additional units that are system dependent:
 - Arbitration/Gearing for multiple request inputs and different clock domains
 - ECC generation and checking
 - Address and Command Parity checking usually accompanied by a retry buffer
 - Pad Calibration Unit



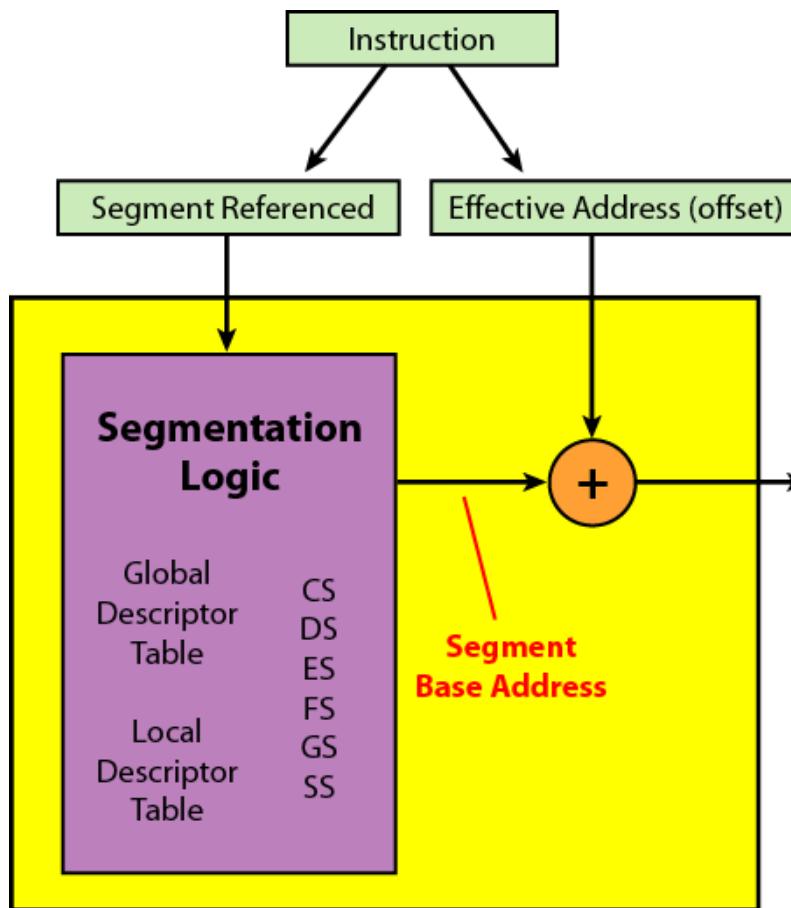
Addresses

Min Huang(min.huang@ lecroy.com)
Do Not Distribute



What is the Address?

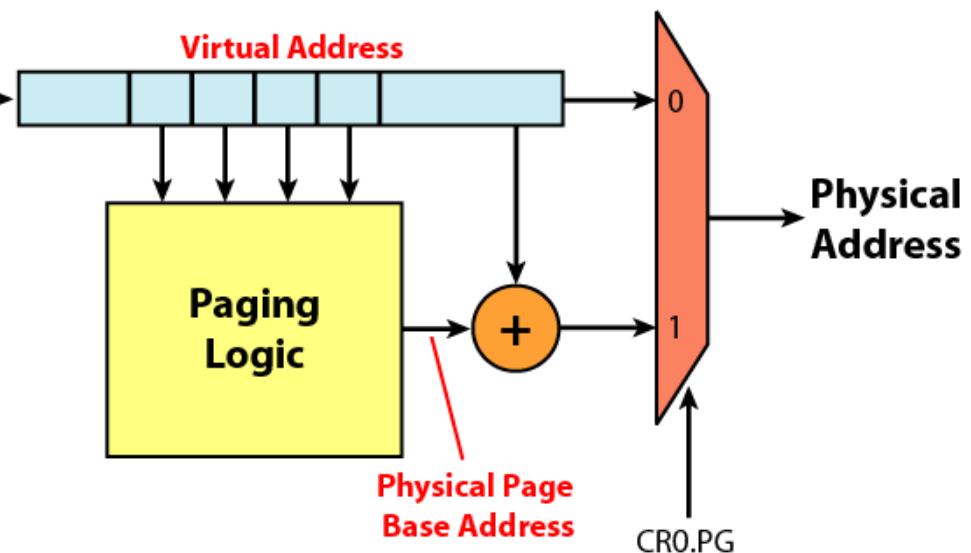
272

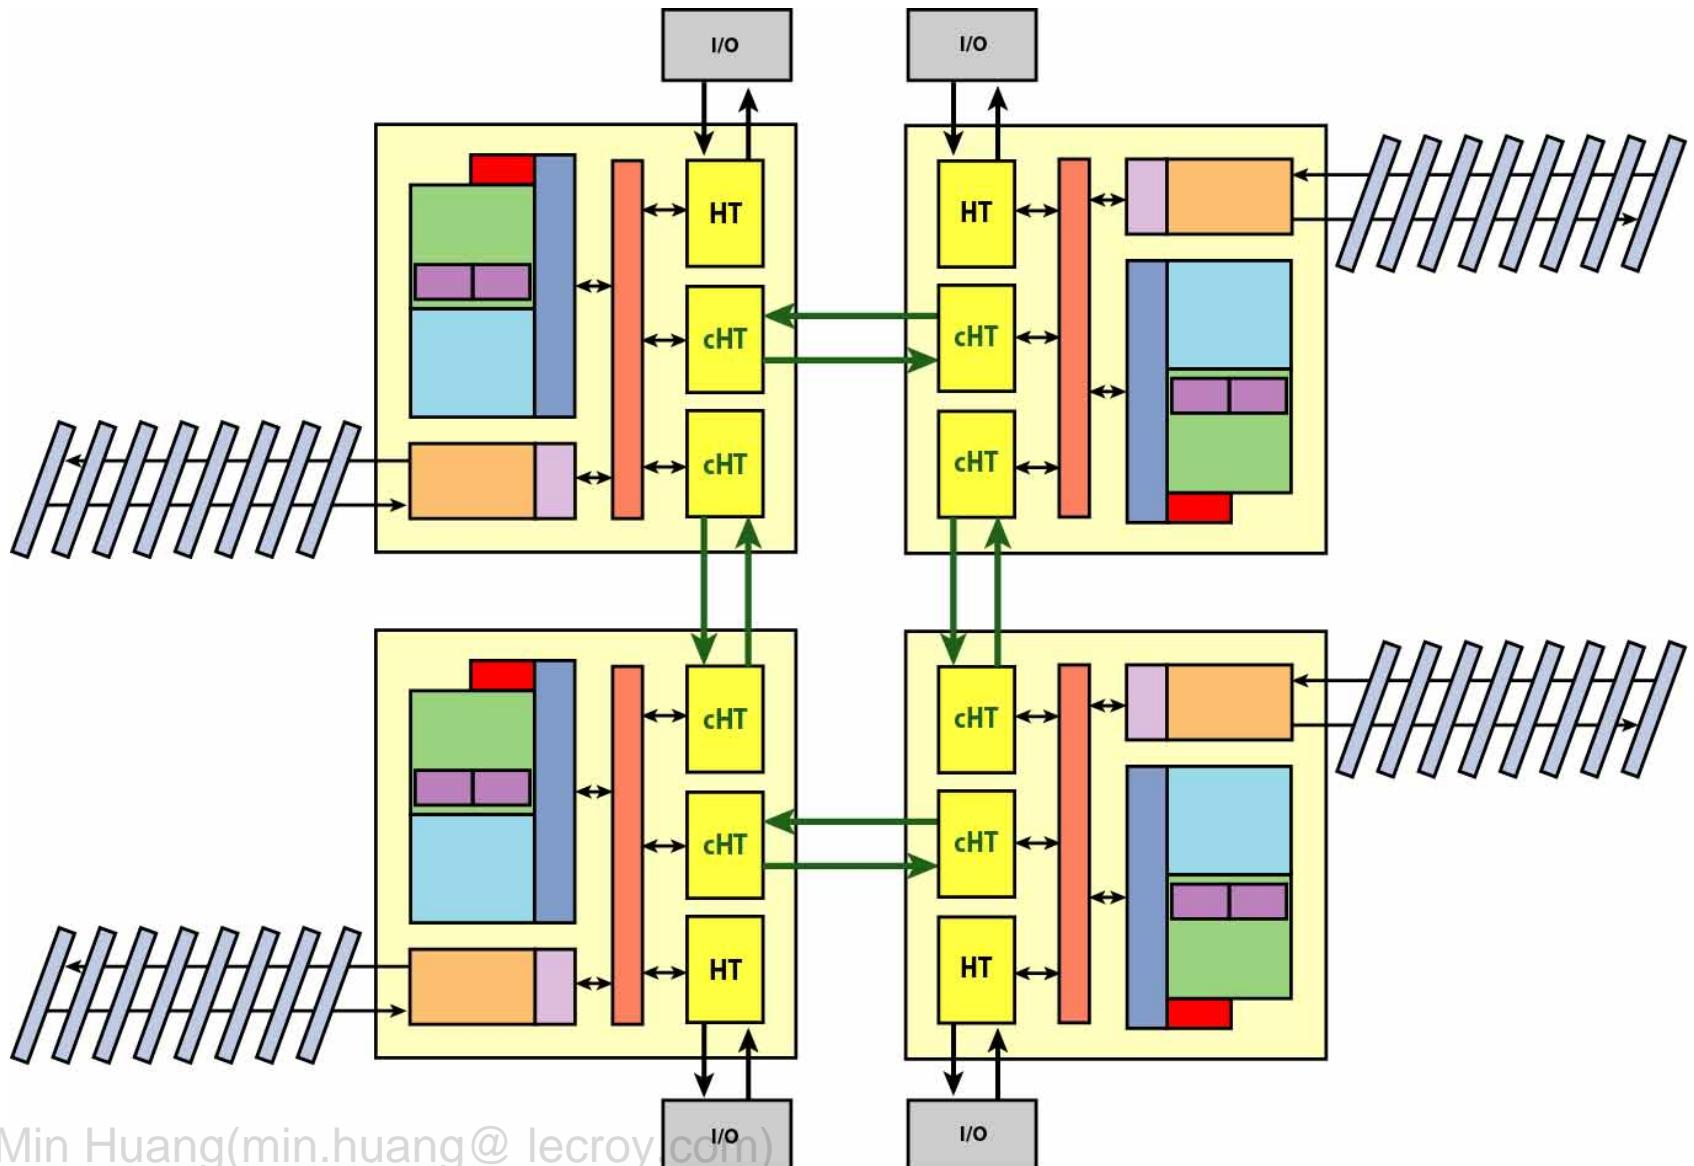


Beware that the address that the programmer uses (the offset) is not the physical address on the FSB nor on cHT, and the physical address is not the address on the memory bus.

Beware the three “interleaved addresses”:

- Intel’s toggle-mode addressing
- Dual channels ganged
- NUMA interleaving





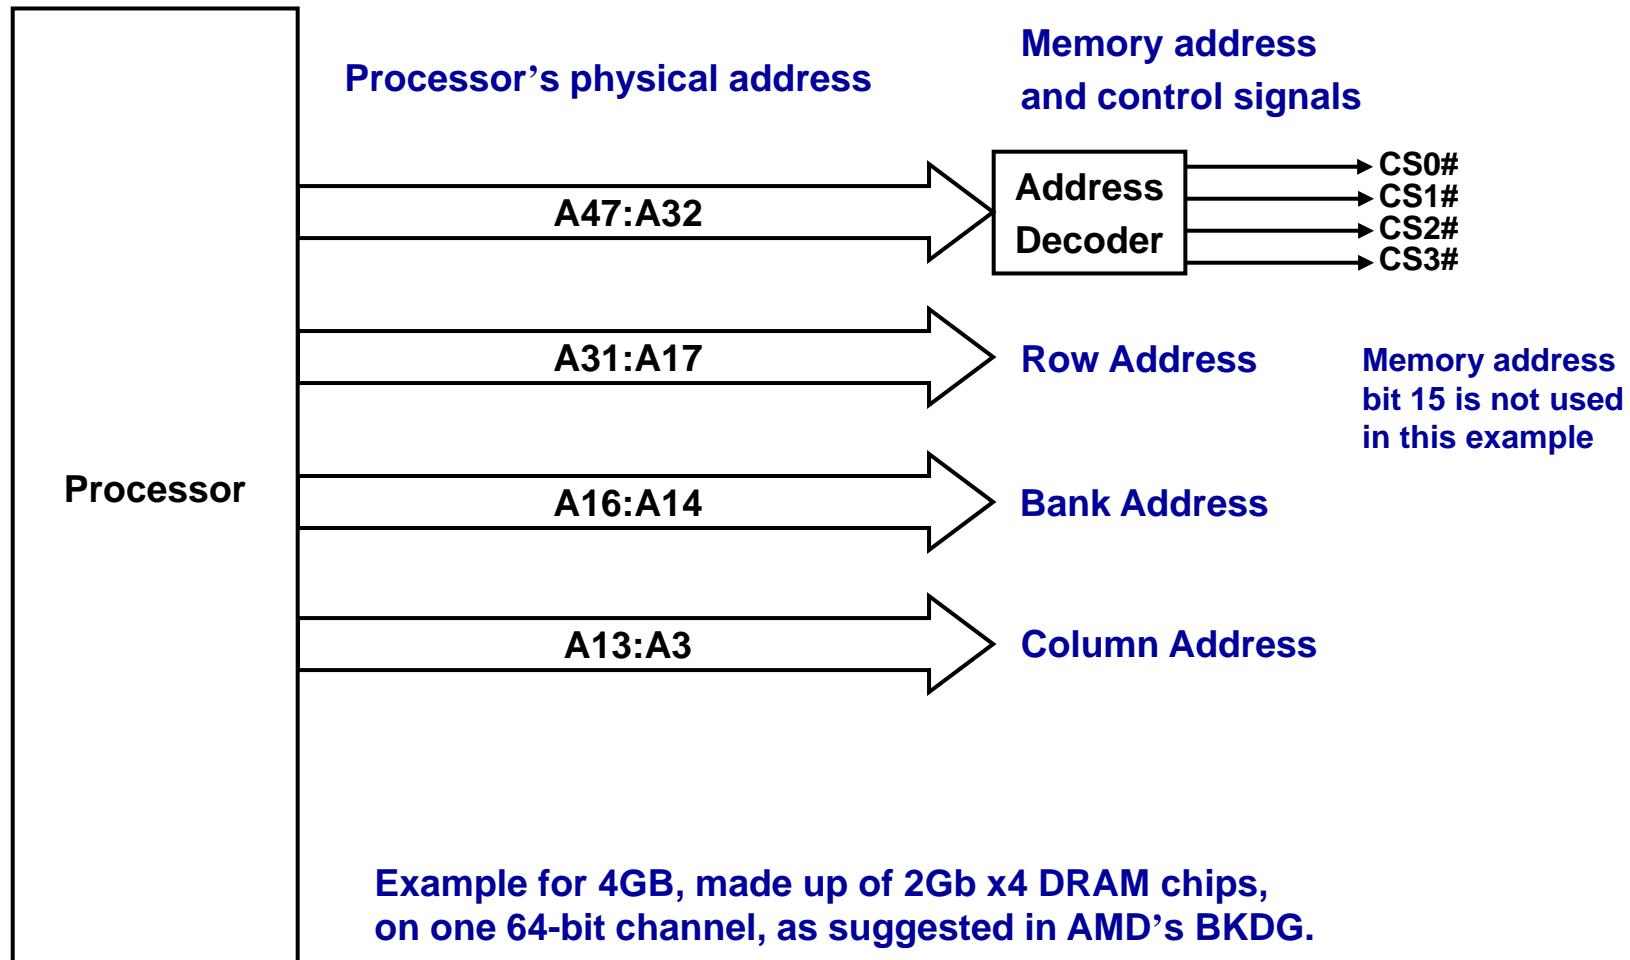


Dual channel can be done several different ways.

- Interleave and Lock Step (Ganged)
 - If like pairs of DIMMs are populated across each channel.
 - The MCH will interleave cache lines (Interleave) or partial cache lines (Lock Step) across each channel.
- Asymmetric
 - If the DIMM are not exact pairs then the MCH will align the addresses sequentially through memory from the furthest DIMM to the nearest in one channel then the other.



Example of Address Translation



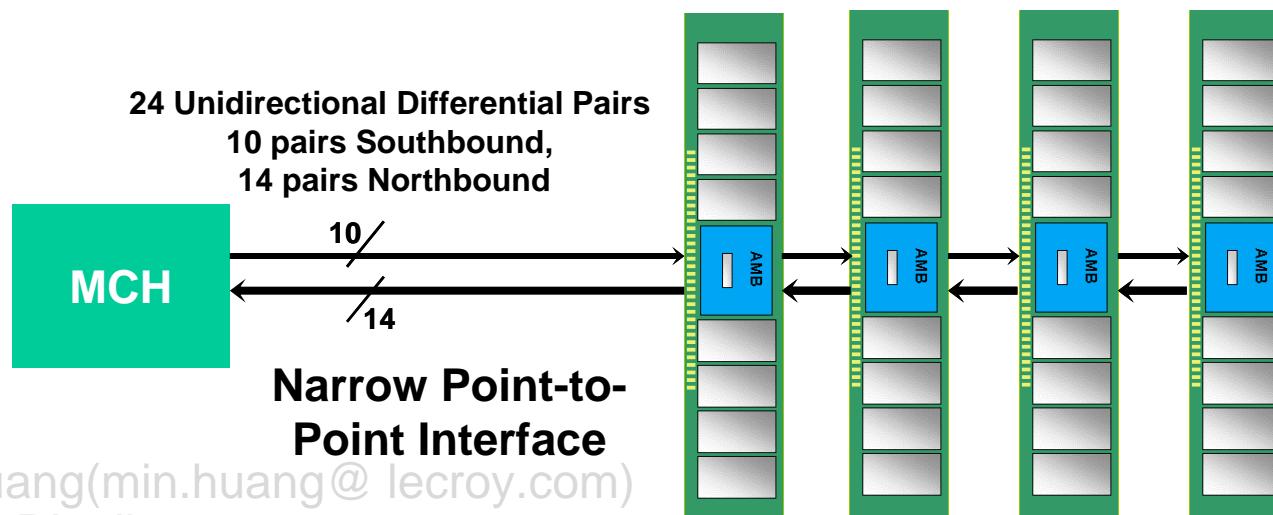
See handout showing 8GB DDR2 with Intel 955 MCH.



Alternative DRAM Solutions

Min Huang(min.huang@ lecroy.com)
Do Not Distribute

- Allows more DIMMs (more memory) per channel while maintaining signal integrity
 - Existing RDIMM solutions only support two DIMMs per channel. FBDIMMs allow more.
 - Address, Control and Data buffered
 - Data buffers isolate the DRAM voltage and data stubs
 - All interface signalling is differential
- Performance
 - Faster processors require higher memory throughput
 - Simultaneous Read/Writes to two Fully-Buffered DIMMs
 - Up to 36 devices behind each buffer
 - 256, 512, 1 & 2 Gbit DRAM support
 - Input clock is $\frac{1}{2}$ Dram base clock. Bit rate per lane is 6X the DRAM speed ie. 533MT/s = 3.2G bit times.
- Cost Sensitive Market
 - No DRAM changes, uses commodity DRAM chips
 - DDR2 DIMM form factor/connector with industry available reference design and Gerbers





- Used in high-end graphics cards.
- GDDR1 and GDDR2 (backed by NVIDIA)
- GDDR3 (backed by ATI)
- GDDR4 first samples out by Samsung
- Don't be confused: GDDR3 is not JEDEC DDR3.
- GDDR uses faster clocks and lower voltages using specially binned production DRAM parts as well as changing internal timing parameters.
- Interfaces are proprietary from GDDR to the GPU.
- One architecture routes 2 sets of address and command lines, one to each rank, with a shared data bus.



- Reduced Latency DRAM is a Micron product aimed at cache-like memory subsystems for communications and graphics.
- RL DRAM is a spin off of DDR2.
- The key difference is that the time from the request to the data transfer is much shorter, and RL DRAM is cost-effective than other high speed solutions



- XDR and XDR2 are follow-on to Rambus.
- Used in high-end, compact systems such as game stations and some graphics cards.
- Bidirectional (16 lane) differential Rambus signaling levels (200mV)
- Programmable ODT
- Adaptive impedance compensation
- Up to 4 bank interleaved transactions at once.
- Early read-after-write capability for higher bandwidth.
- Zero overhead for refresh



Thank you!

Please send questions to eLearning@mindshare.com