MINDSHARE

BRINGING LIFE TO KNOWLEDGE

DRAM Technology Architecture

Presentation Rev g1 Agilent

DRAM Topics

- 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
- 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

Objectives

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.

MINDSHARE

BRINGING LIFE TO KNOWLEDGE

System Architecture

Intel Basic Desktop

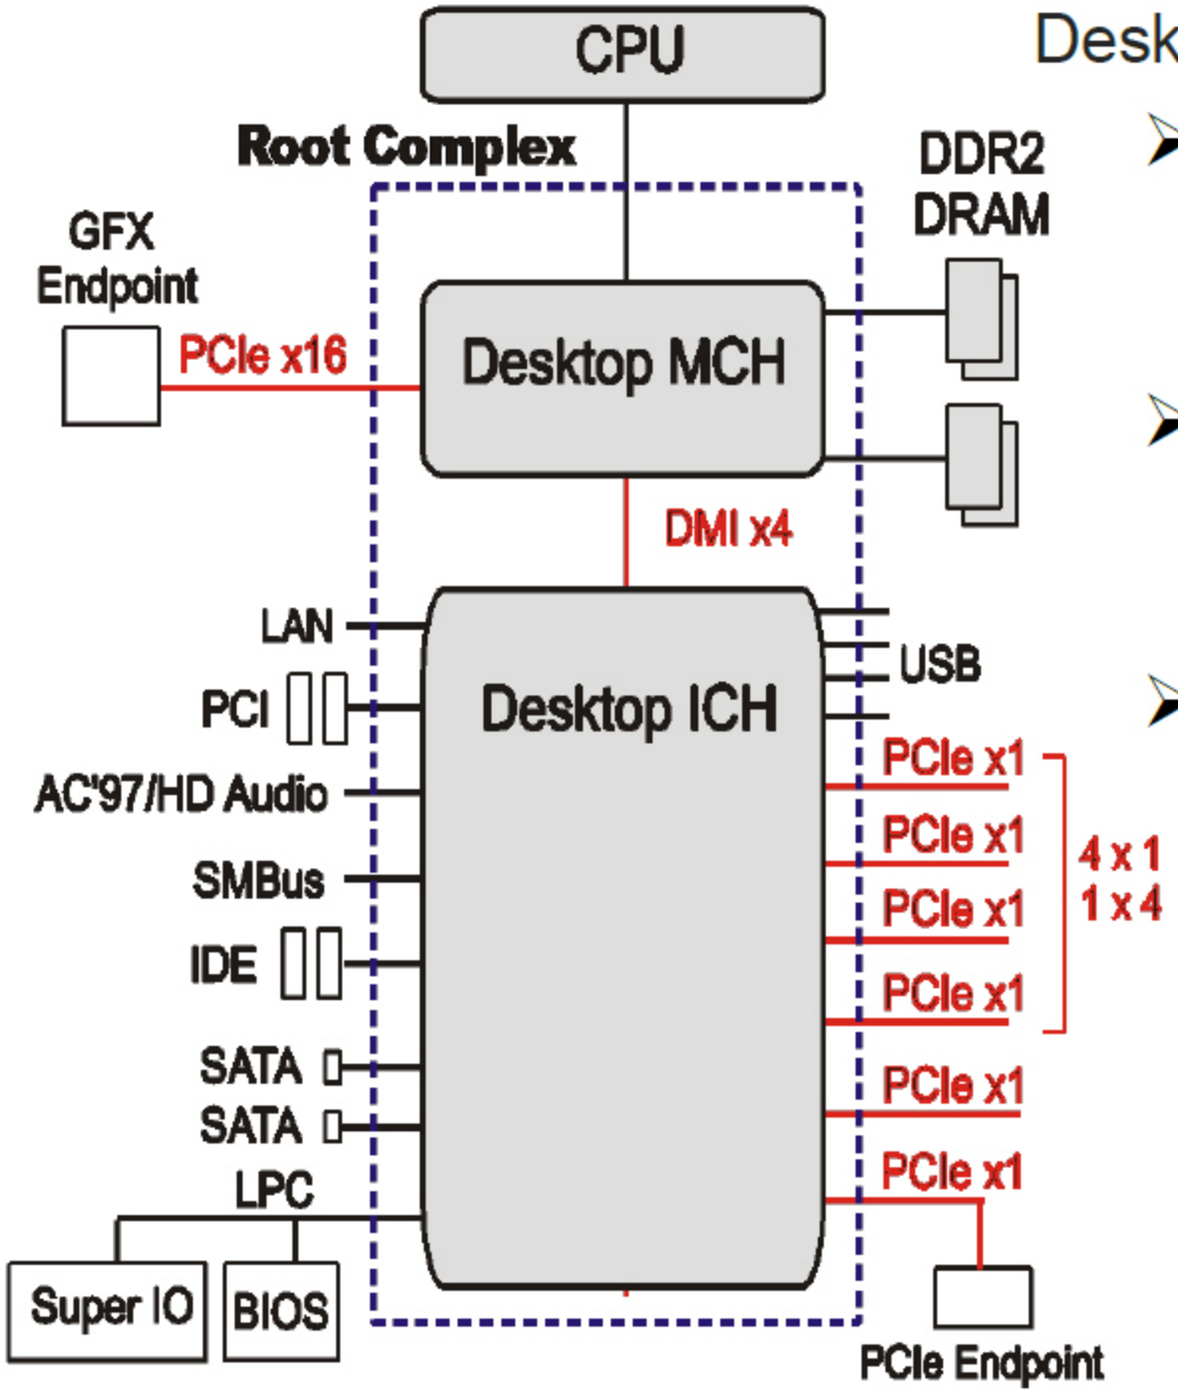

Desktop Chipset PCIe Support

- 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.

Intel Server System

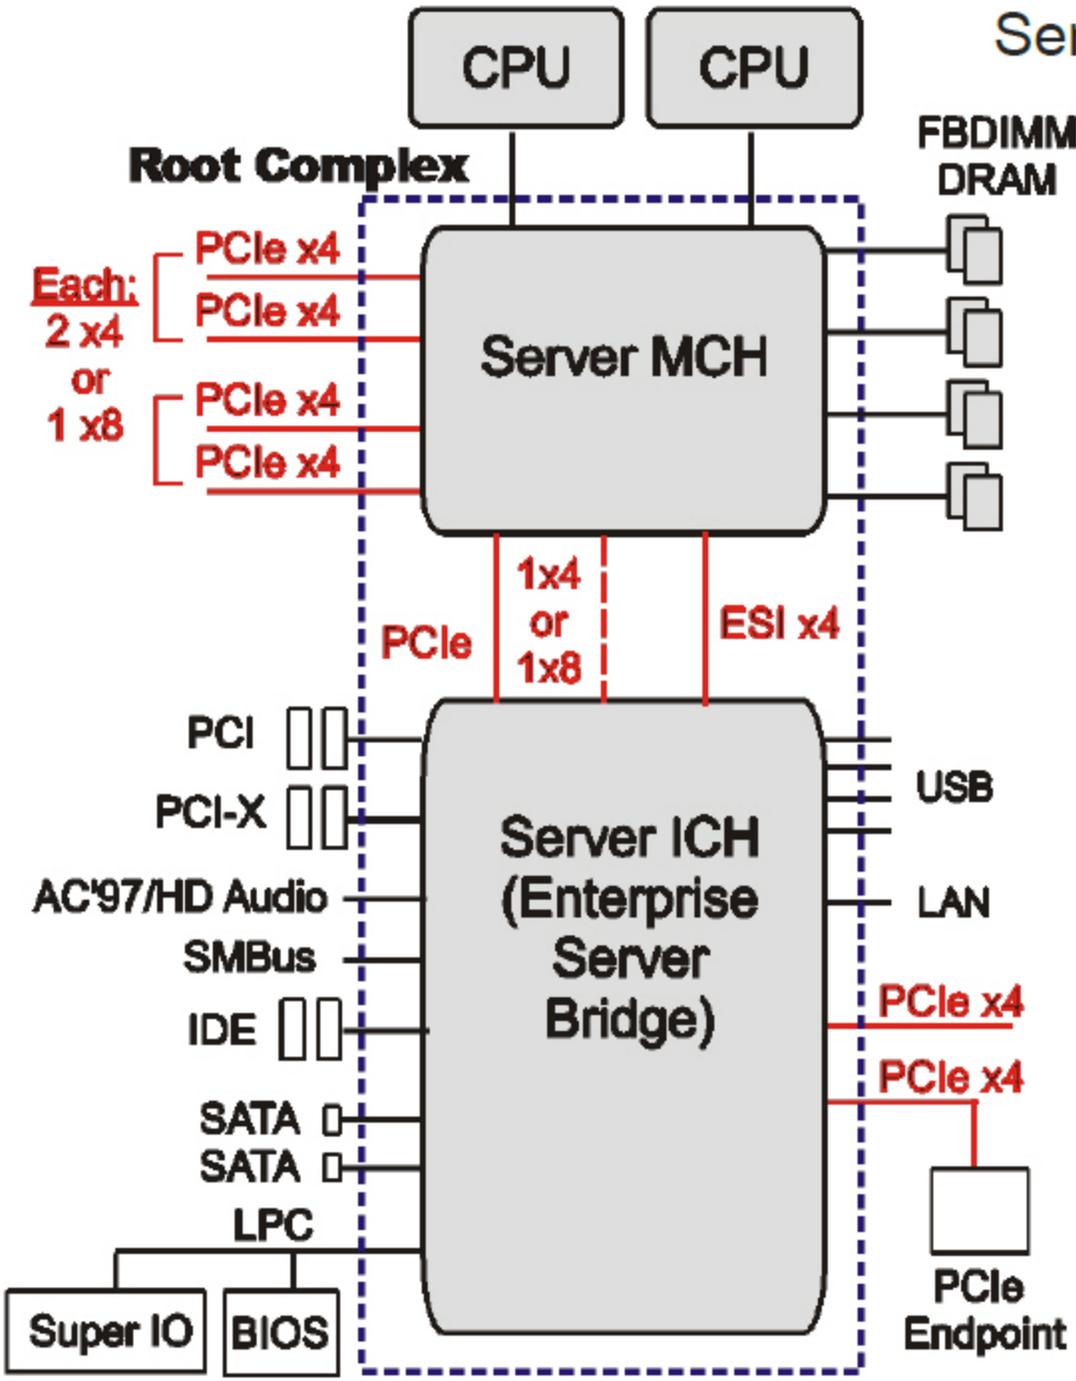

Server Chipset PCIe Support

- 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 System Overview

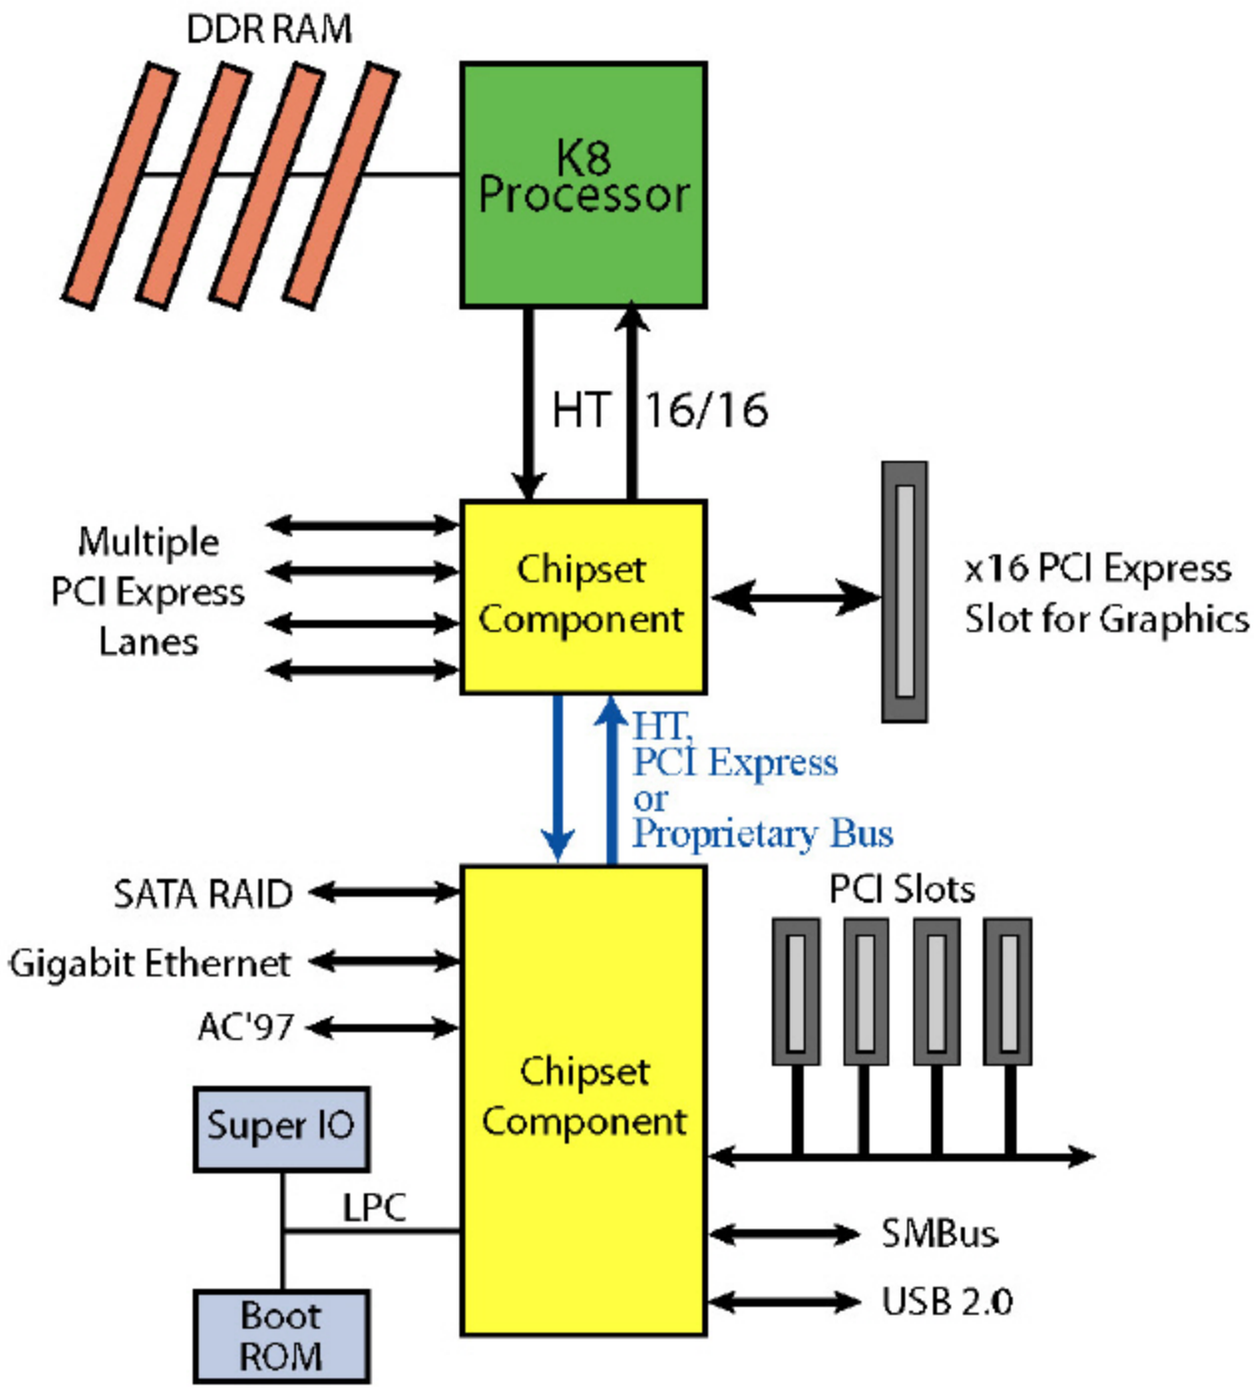

Quad Core Greyhound Processor

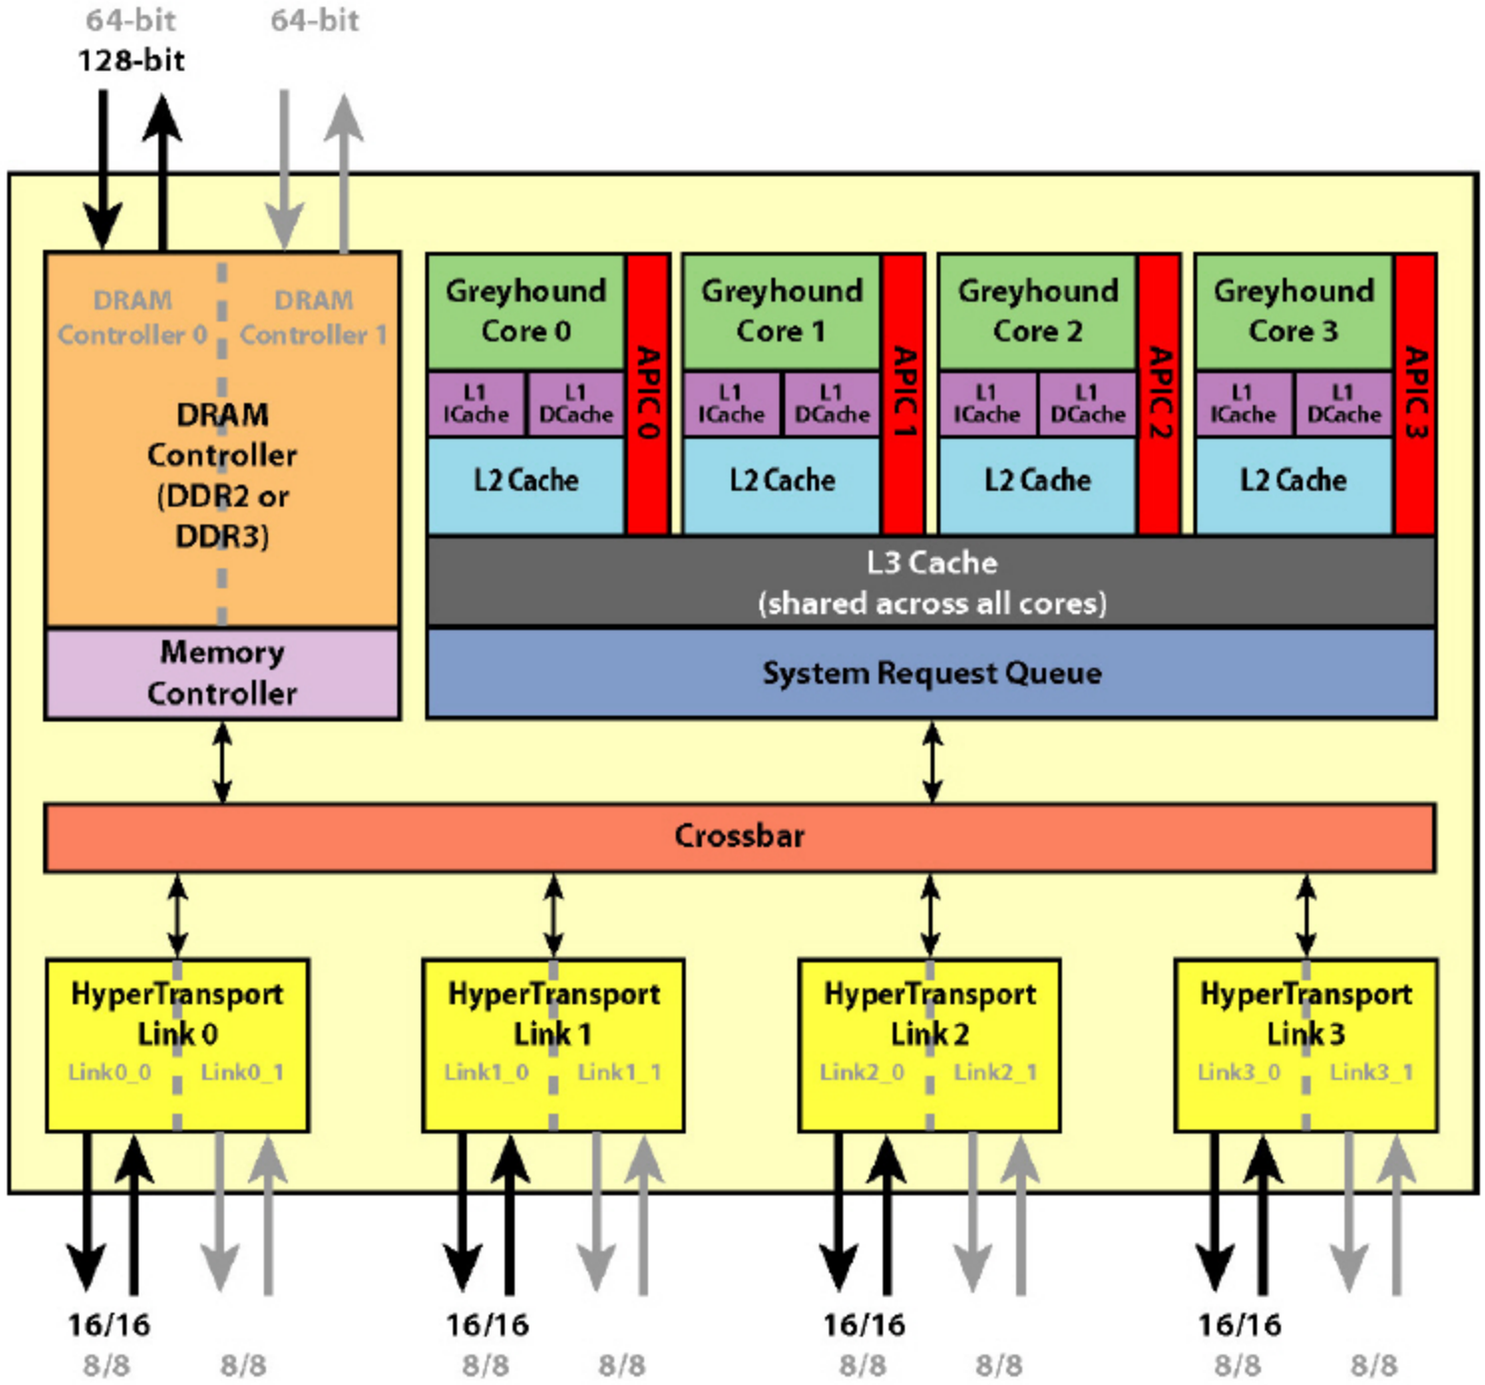

Example 4-Way System

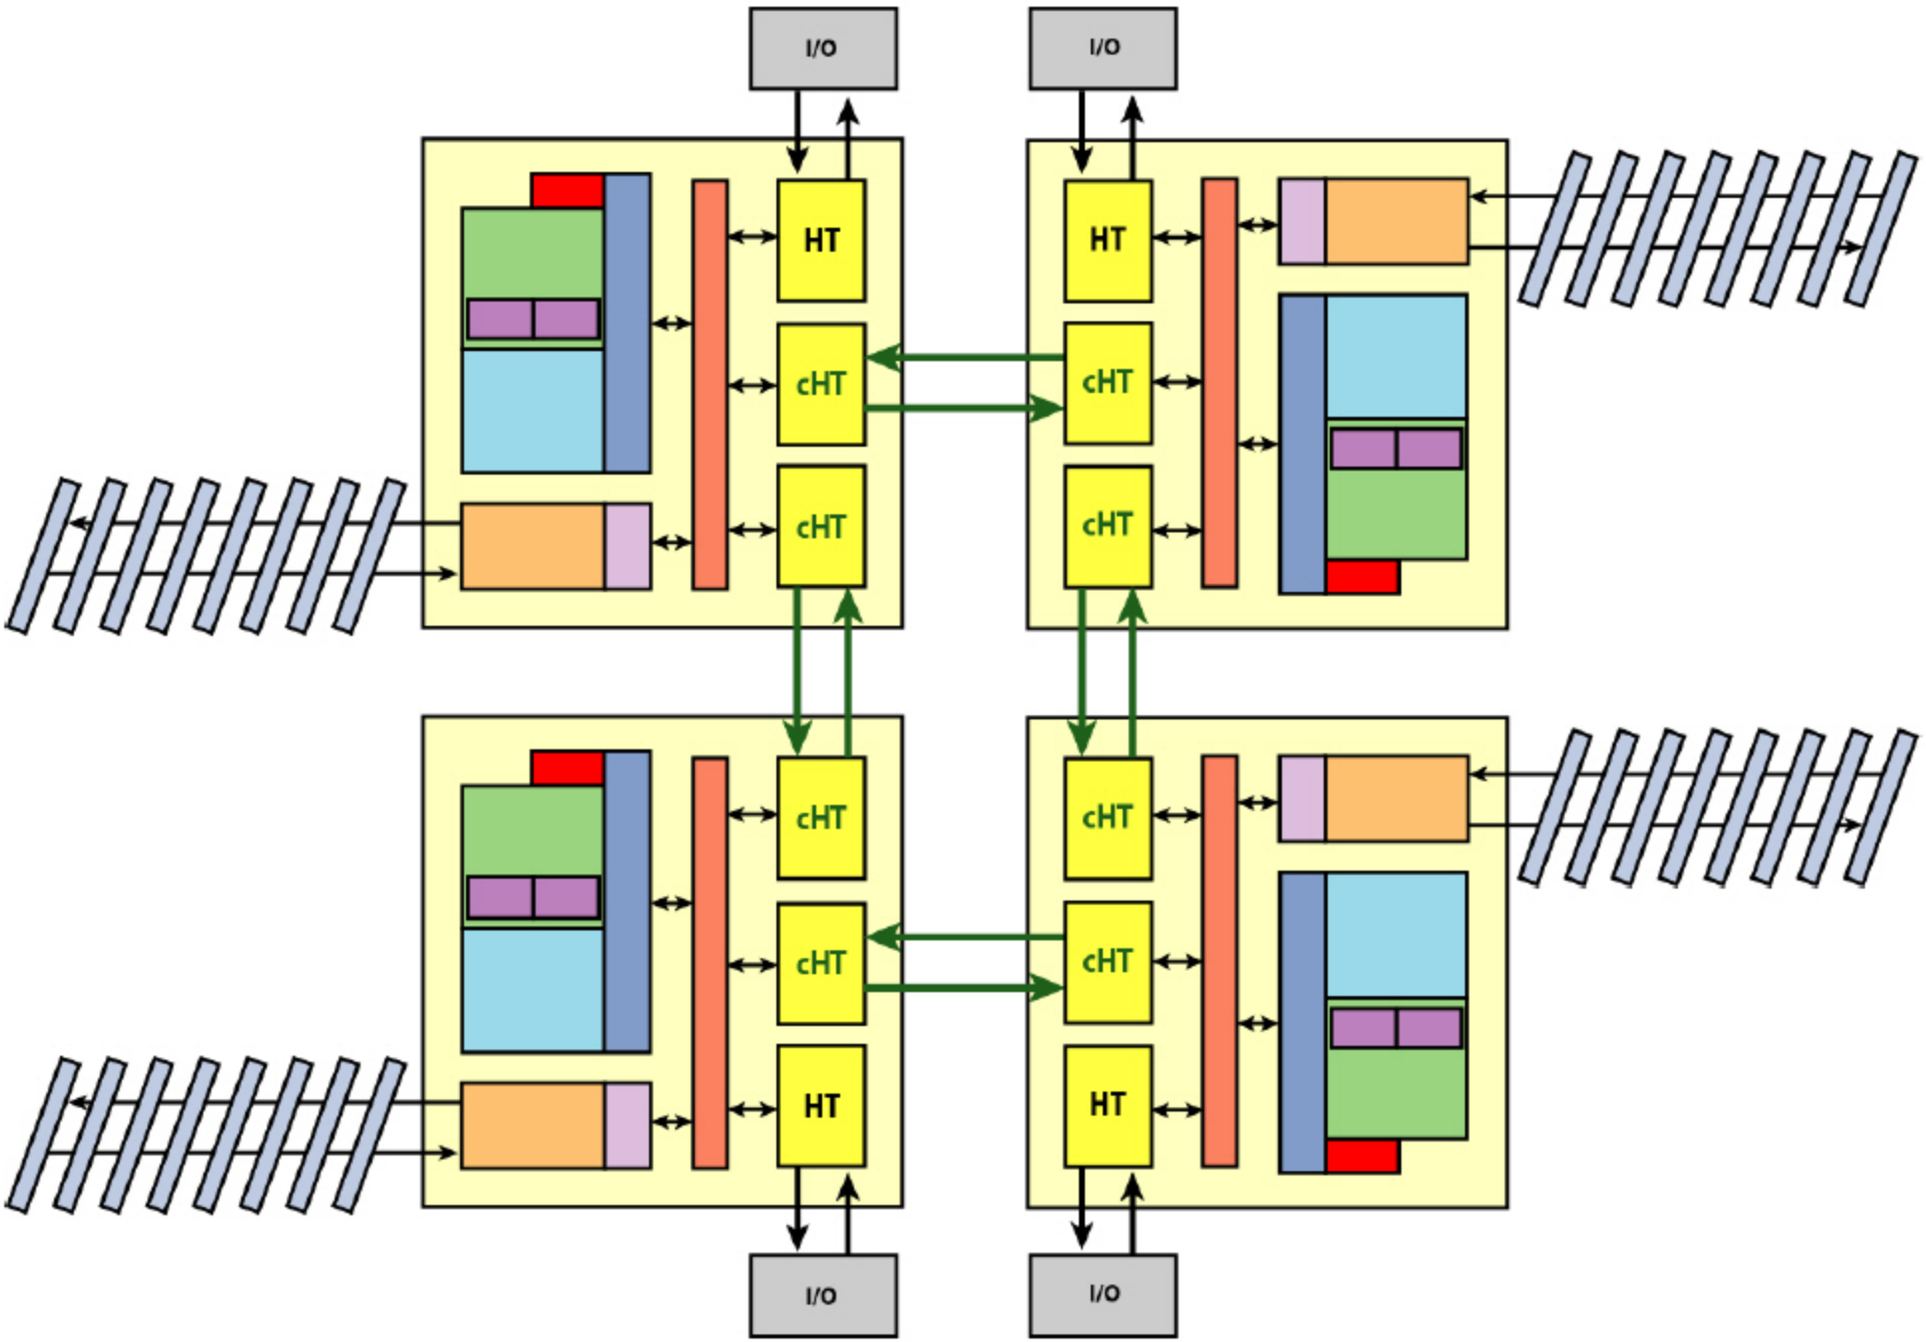

MINDSHARE

BRINGING LIFE TO KNOWLEDGE

DRAM Feature Summary

DRAM Overview

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

DRAM Overview

	SDRAM	DDR1	DDR2	DDR3
Performance	100-133MT/s	200-400MT/s	400-1000MT/s	800-1600MT/s
V_{DDQ}	3.3 Volts	2.5 Volts	1.8 Volts	1.5 Volts
V_{TT}	NA	½ V_{DDQ}	½ V_{DDQ}	½ 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 (TBD)
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, etc.
Additive latency	NA	NA	AL=0, 1, 2, 3, 4	TBD
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 JESD79-3B
- Package specs are MO-207 (for example)
- DIMM and SPD specs are JESD21C

MINDSHARE

BRINGING LIFE TO KNOWLEDGE

Intro to DRAM

MINDSHARE

BRINGING LIFE TO KNOWLEDGE

Historical Background

History of DRAM

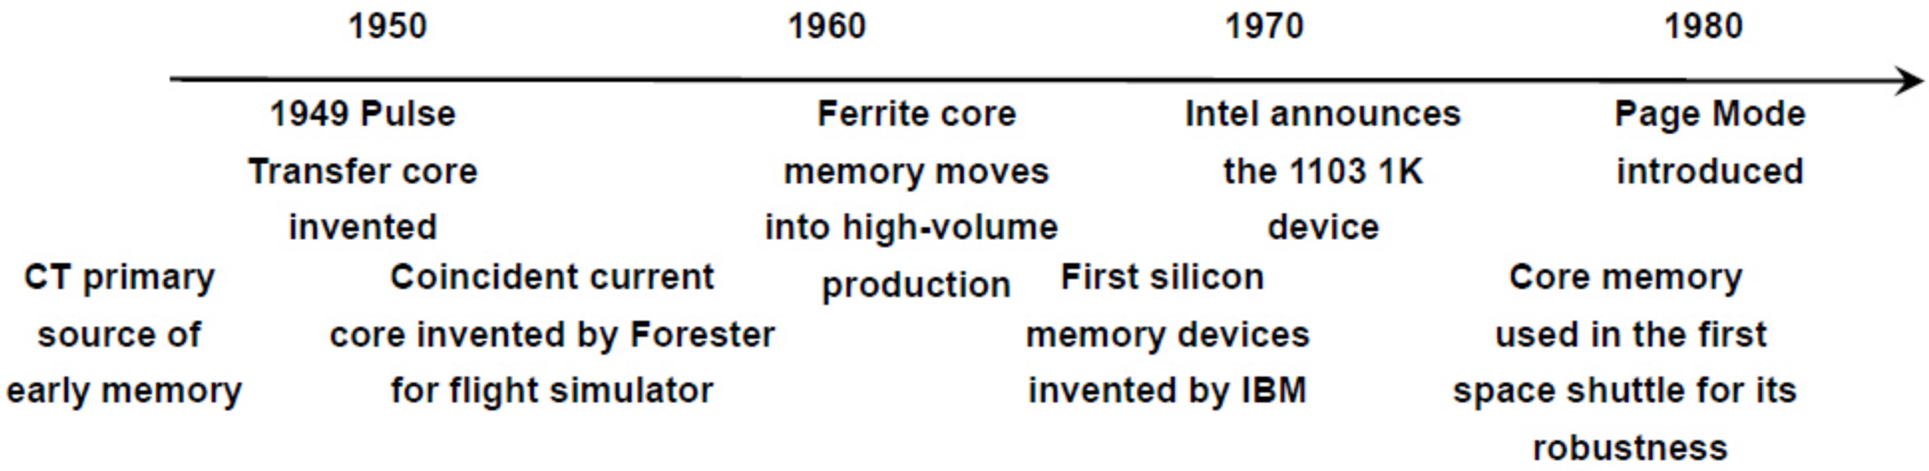

- Cathode Tubes were the first electronic devices used for memory. CT's were plagued with problems like excessive heat, size, and voltage as well as being unreliable and temperamental. This set the stage for an electro-mechanical memory called "core memory".

History of DRAM

- 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.

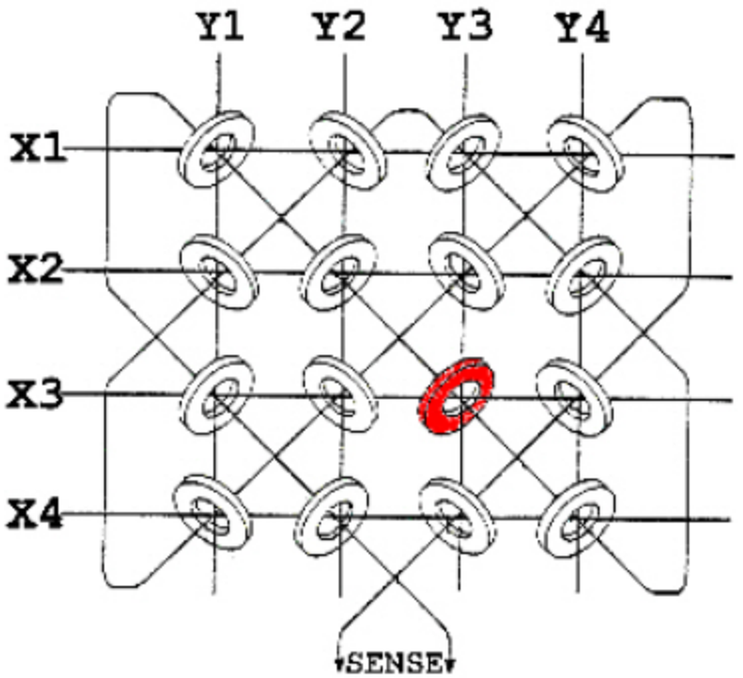

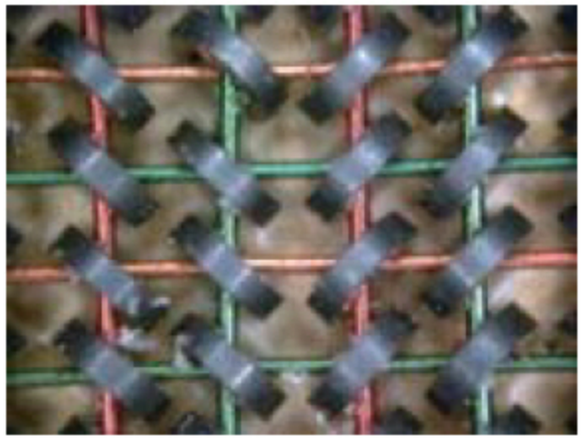

History of DRAM

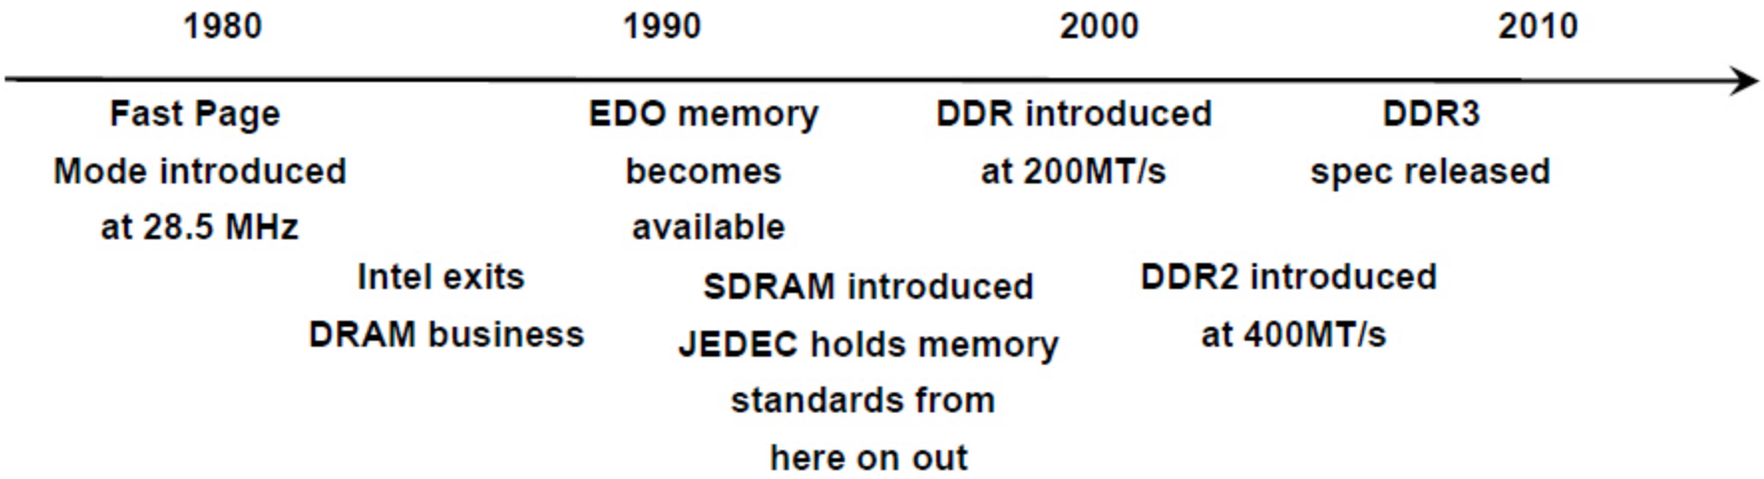

- With the invention of the PC in 1981, silicon-based memory demand skyrocketed. Other countries like Japan begin to manufacture low cost DRAMs. By 1985, Japan has nearly dominated the market. US lawmakers team up with silicon manufacturers to equalize the playing field. New laws against price gouging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

MINDSHARE

BRINGING LIFE TO KNOWLEDGE

Why DRAM?

RAM

- 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.

Types of RAM

- 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

MINDSHARE

BRINGING LIFE TO KNOWLEDGE

DRAM Cell Architecture

Cell Architecture

- 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).

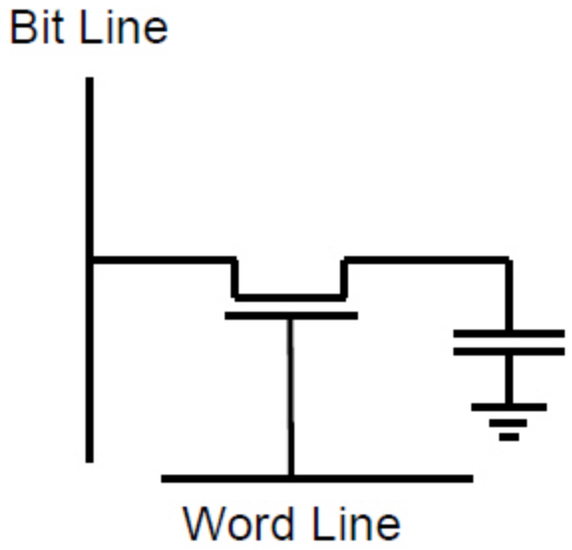

- DRAM Cell Read
  1. Bit line is precharged to Sense Amp threshold voltage (V_{DD} / 2).
  2. Word line turns on transistor to allow charge to flow from capacitor to bit line. (Read from Cell)
  3. Sense Amp latches electrical high or low.

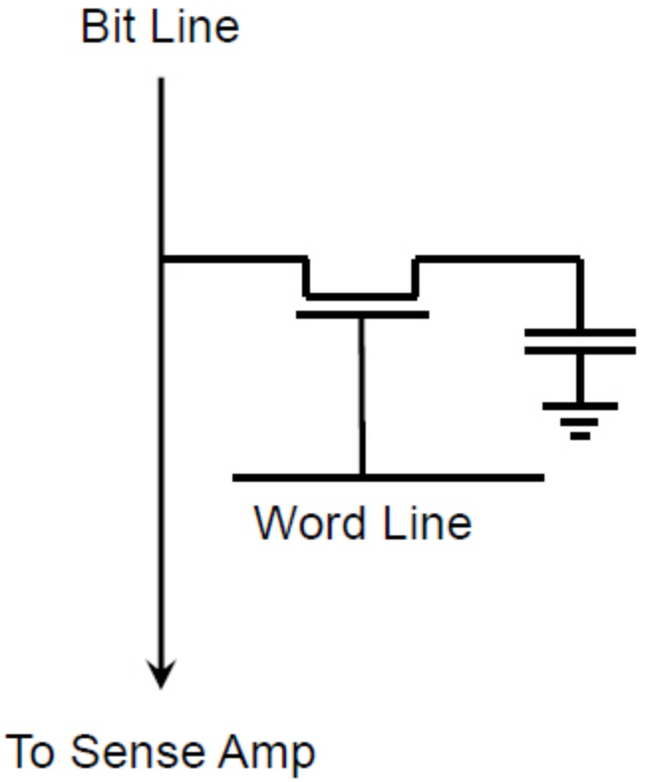

DRAM Cell Write or Precharge

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.

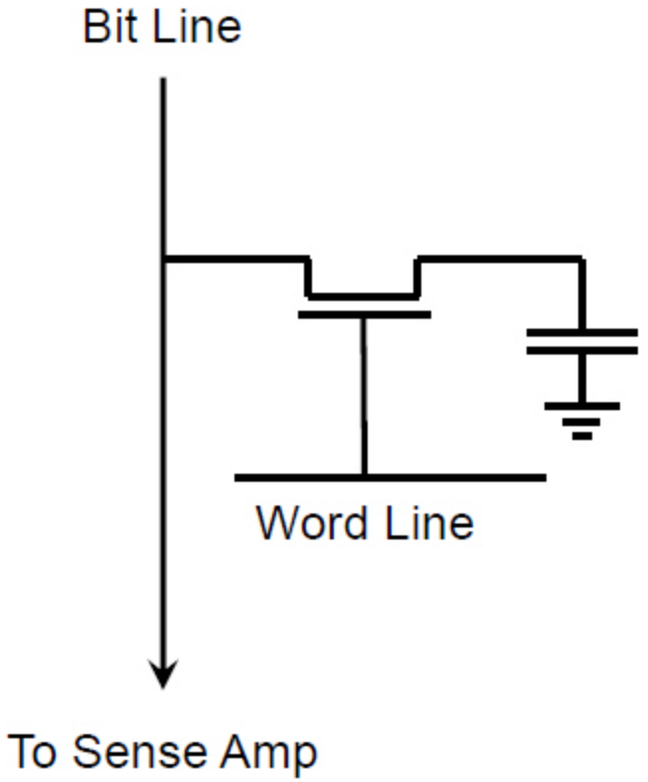

Folded Cell Array Architecture

- 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.

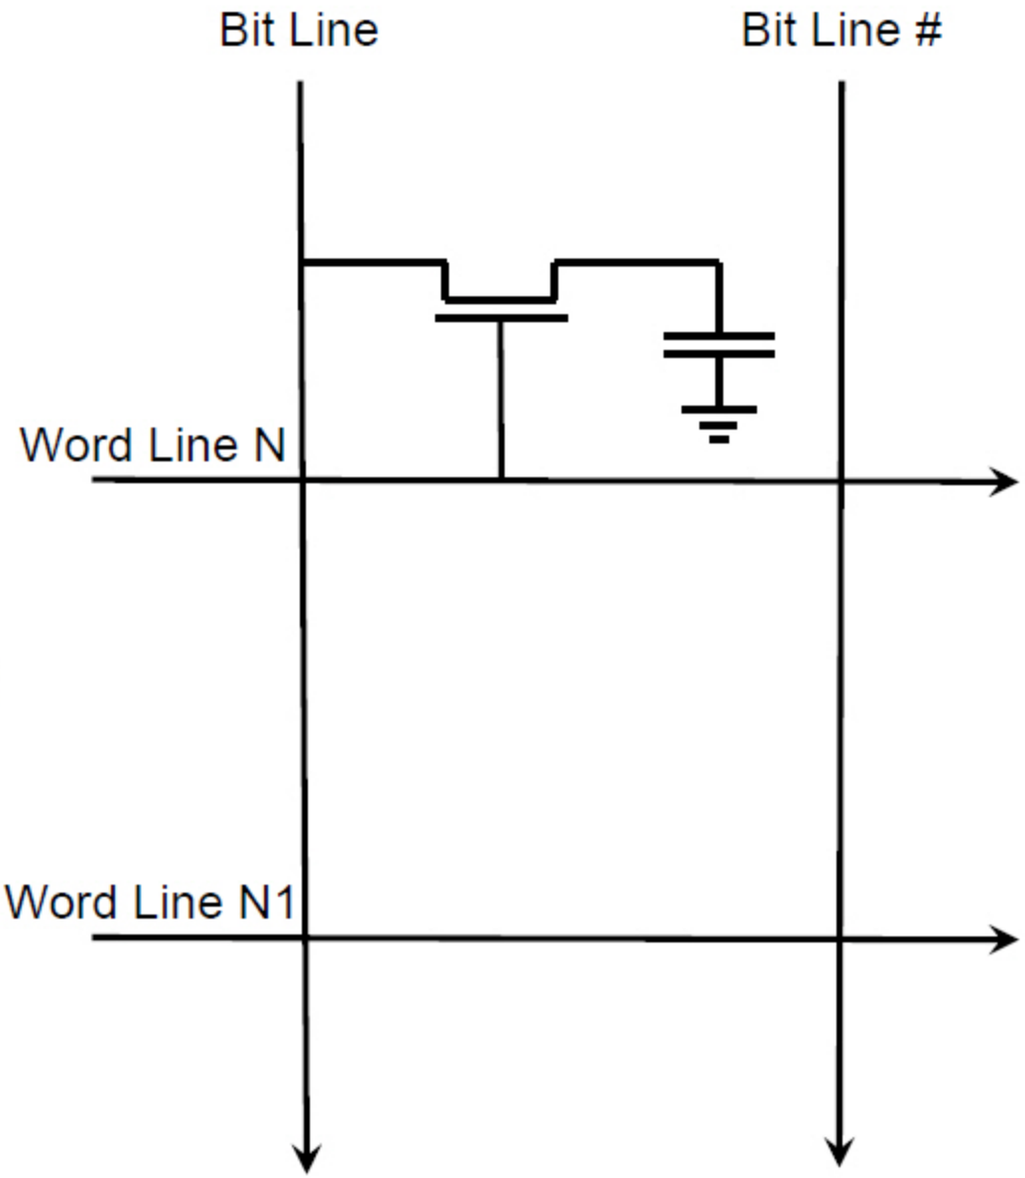

Folded Cell Array Architecture

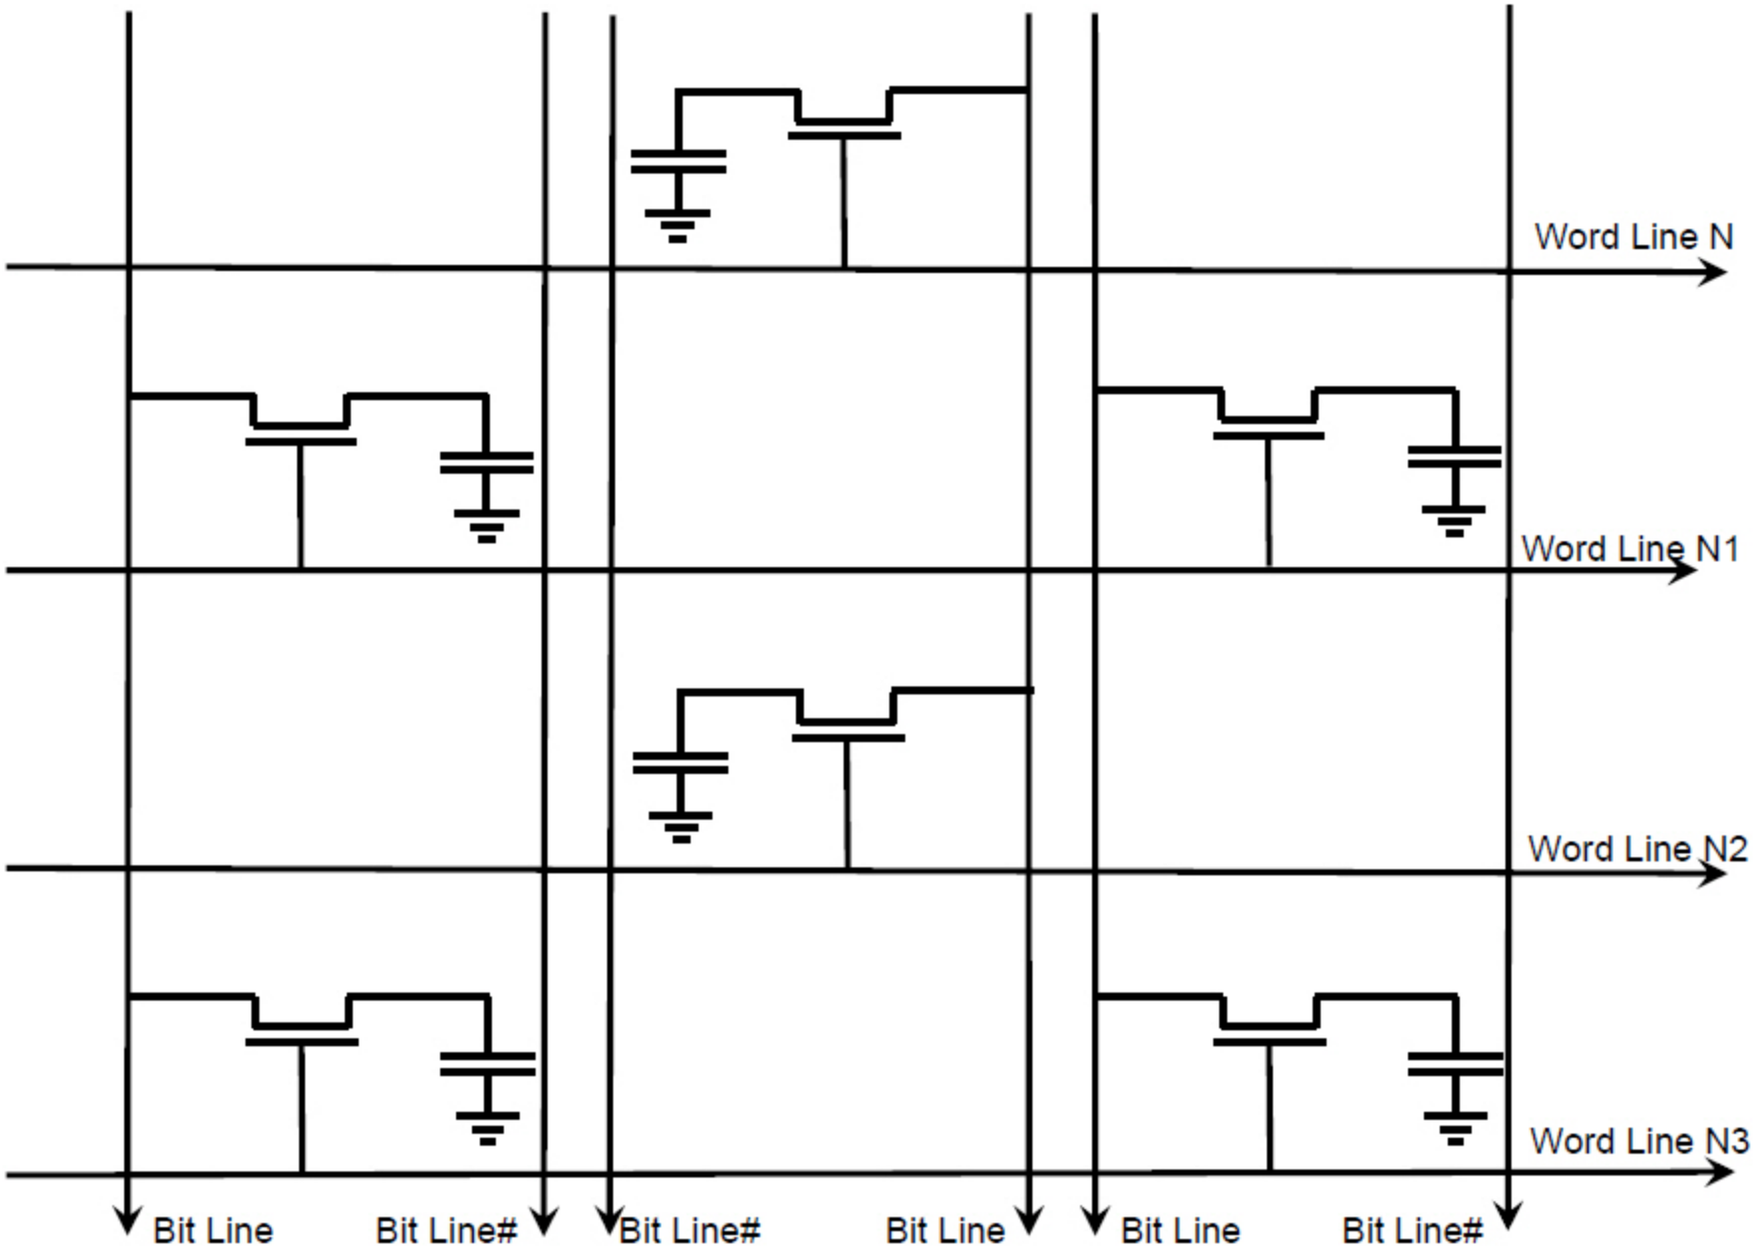

To Sense Amp

MINDSHARE

BRINGING LIFE TO KNOWLEDGE

DRAM Chip Architecture

The First DRAMs

- The first LSI (Large Scale Integration) DRAMs were manufactured by Intel in 1970 staring with the 1101 chip.
- The 1101 had a 256-bit array with one data line. The following examples show the evolution of the DRAM.

The Array

- 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.

Split Address Scheme

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

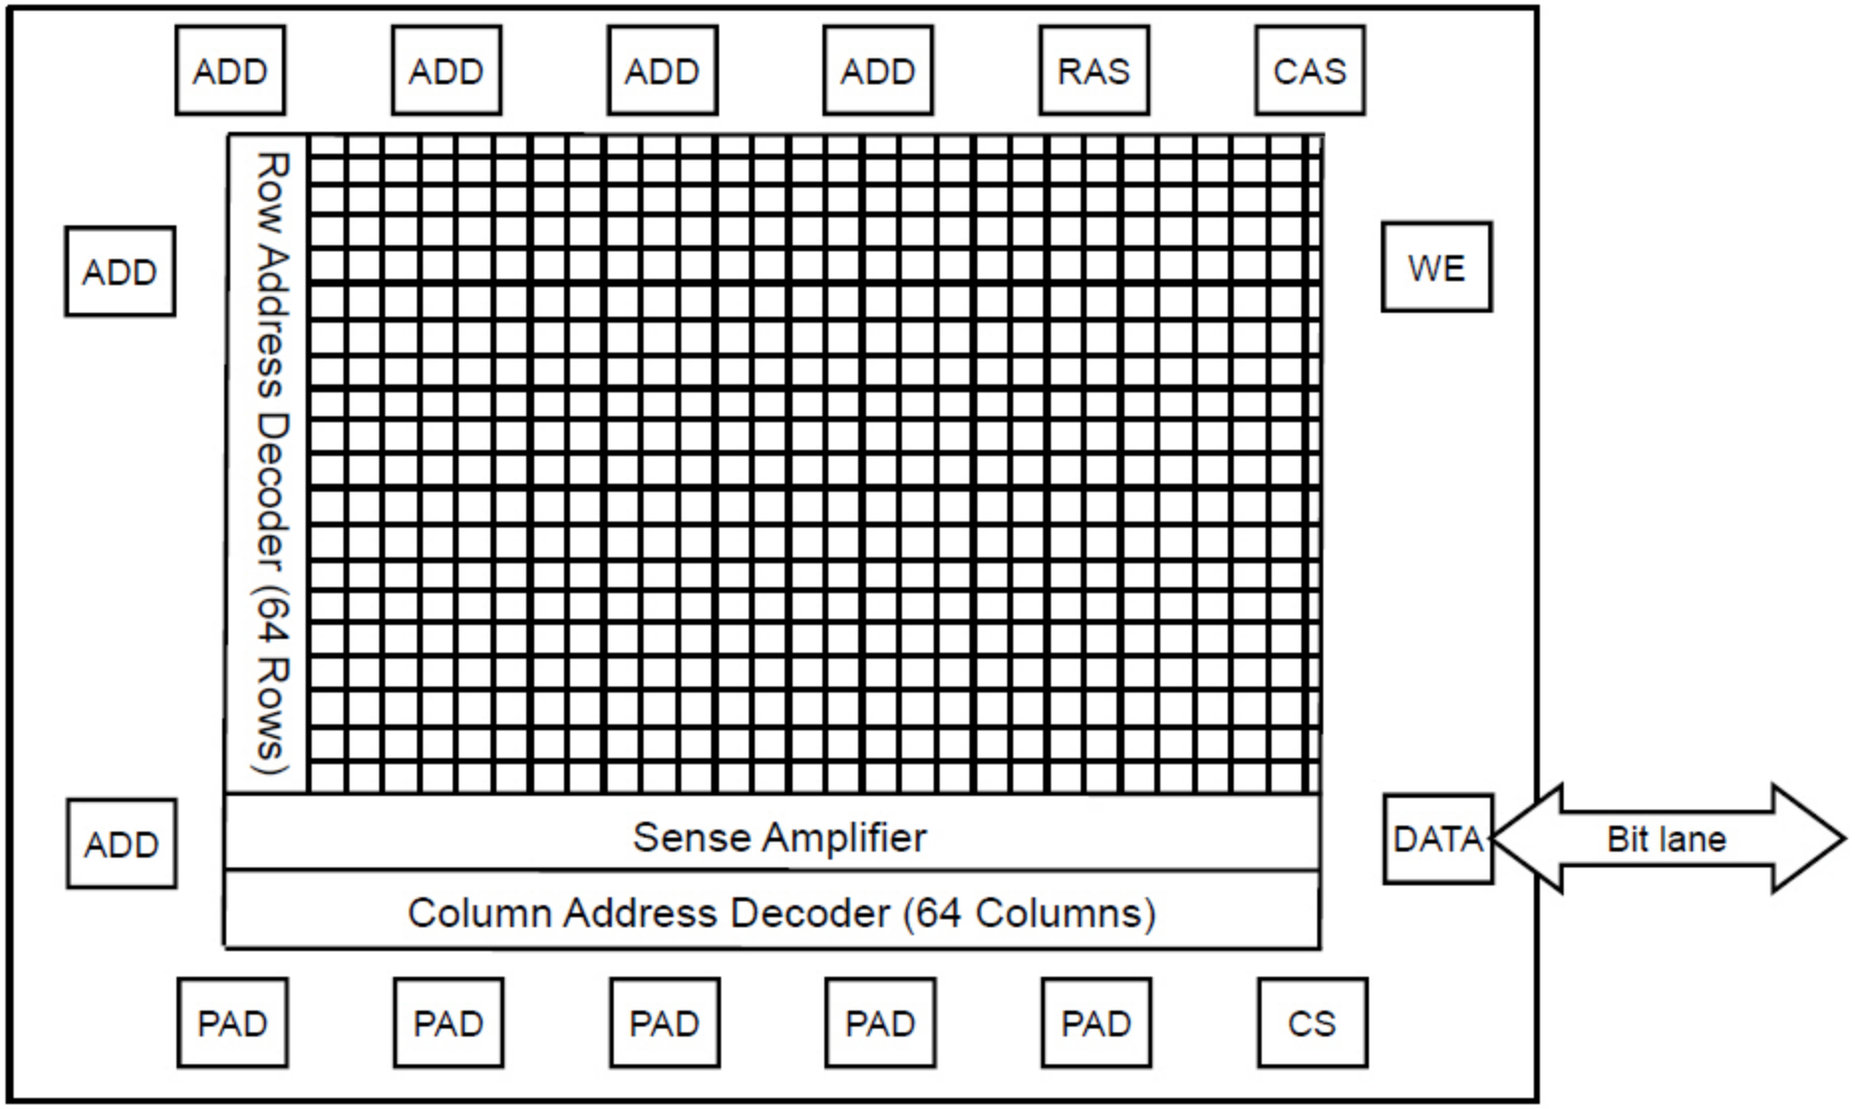

Simple Async DRAM Controller Design

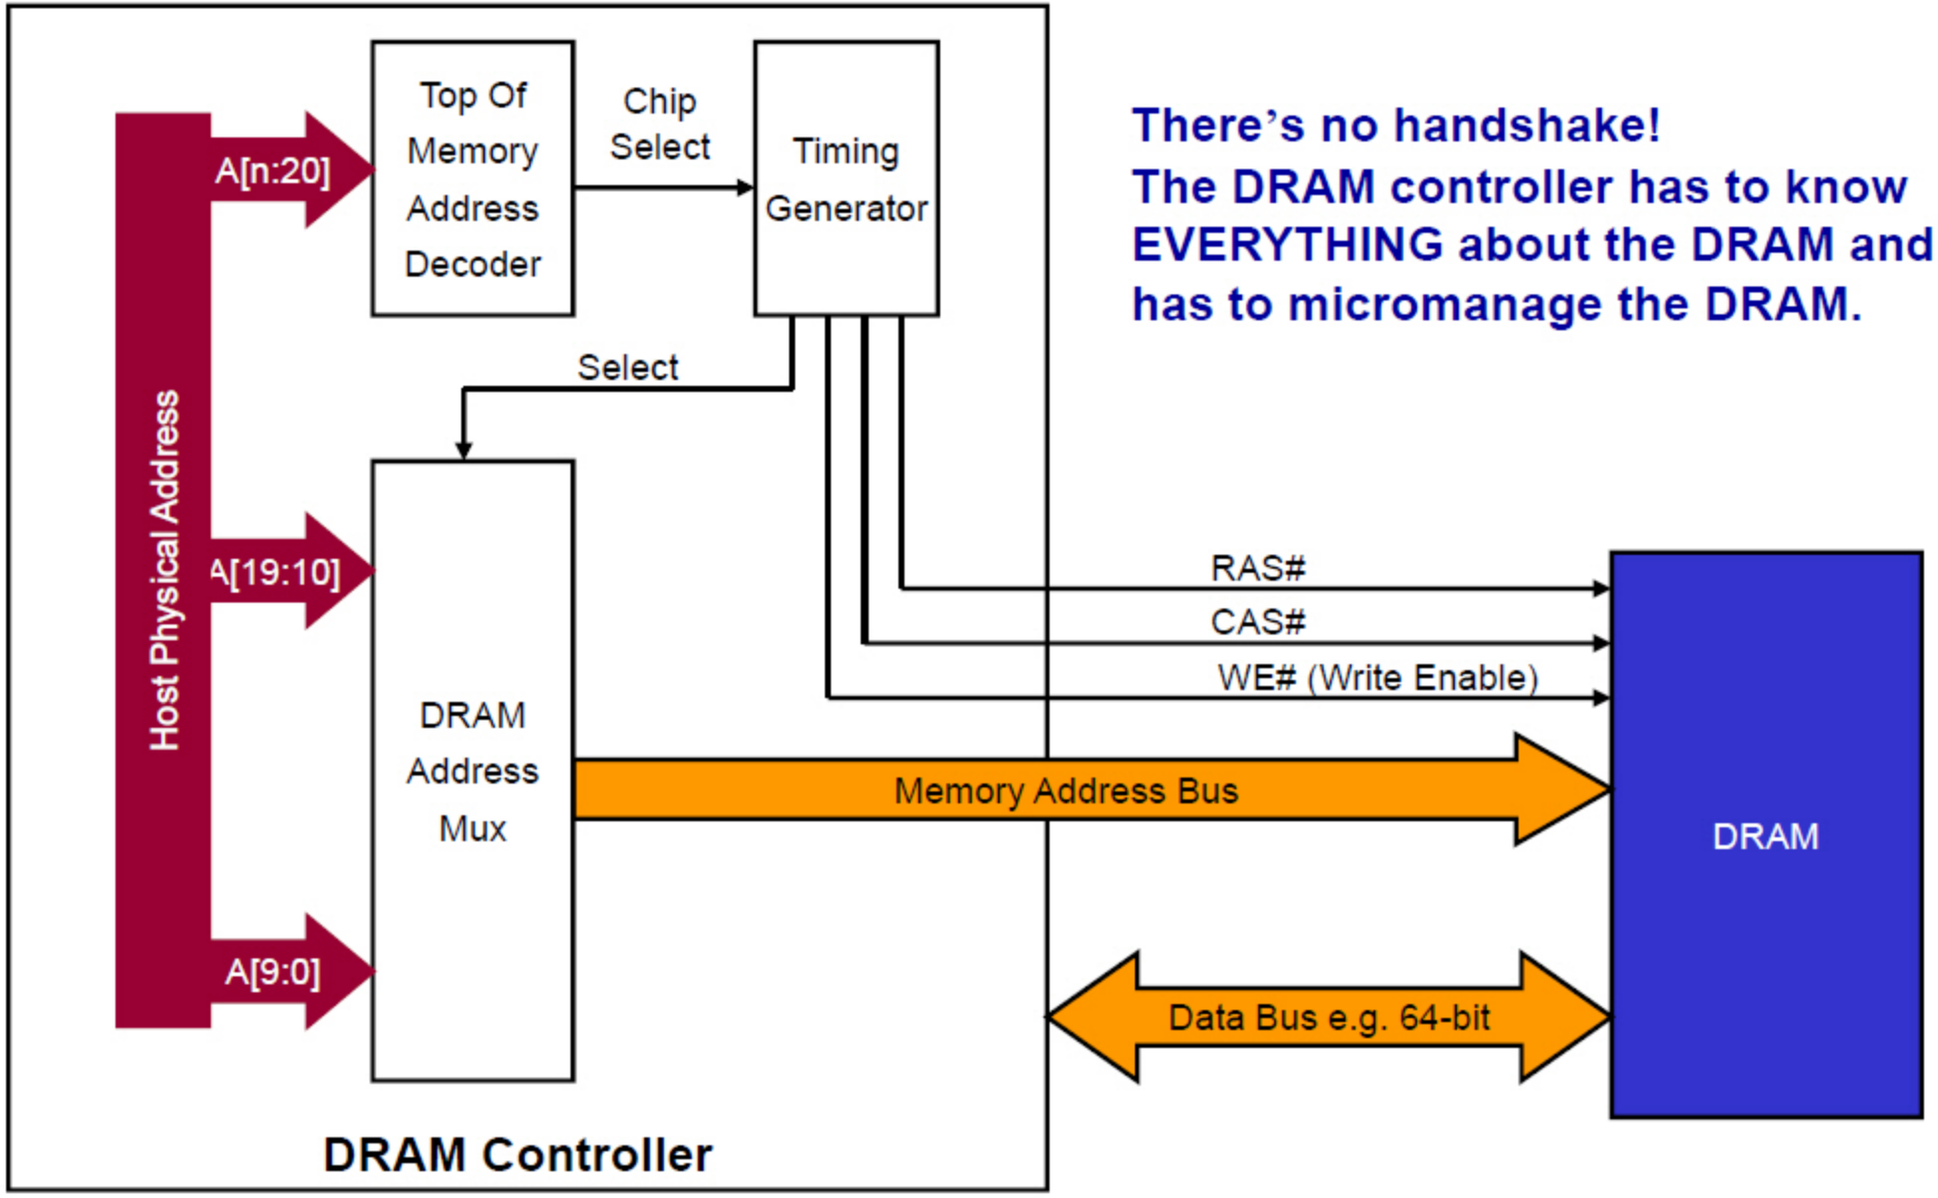

There's no handshake!
The DRAM controller has to know EVERYTHING about the DRAM and it has to micromanage the DRAM.

Problems: Muxed Address, Precharge, Refresh

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**.

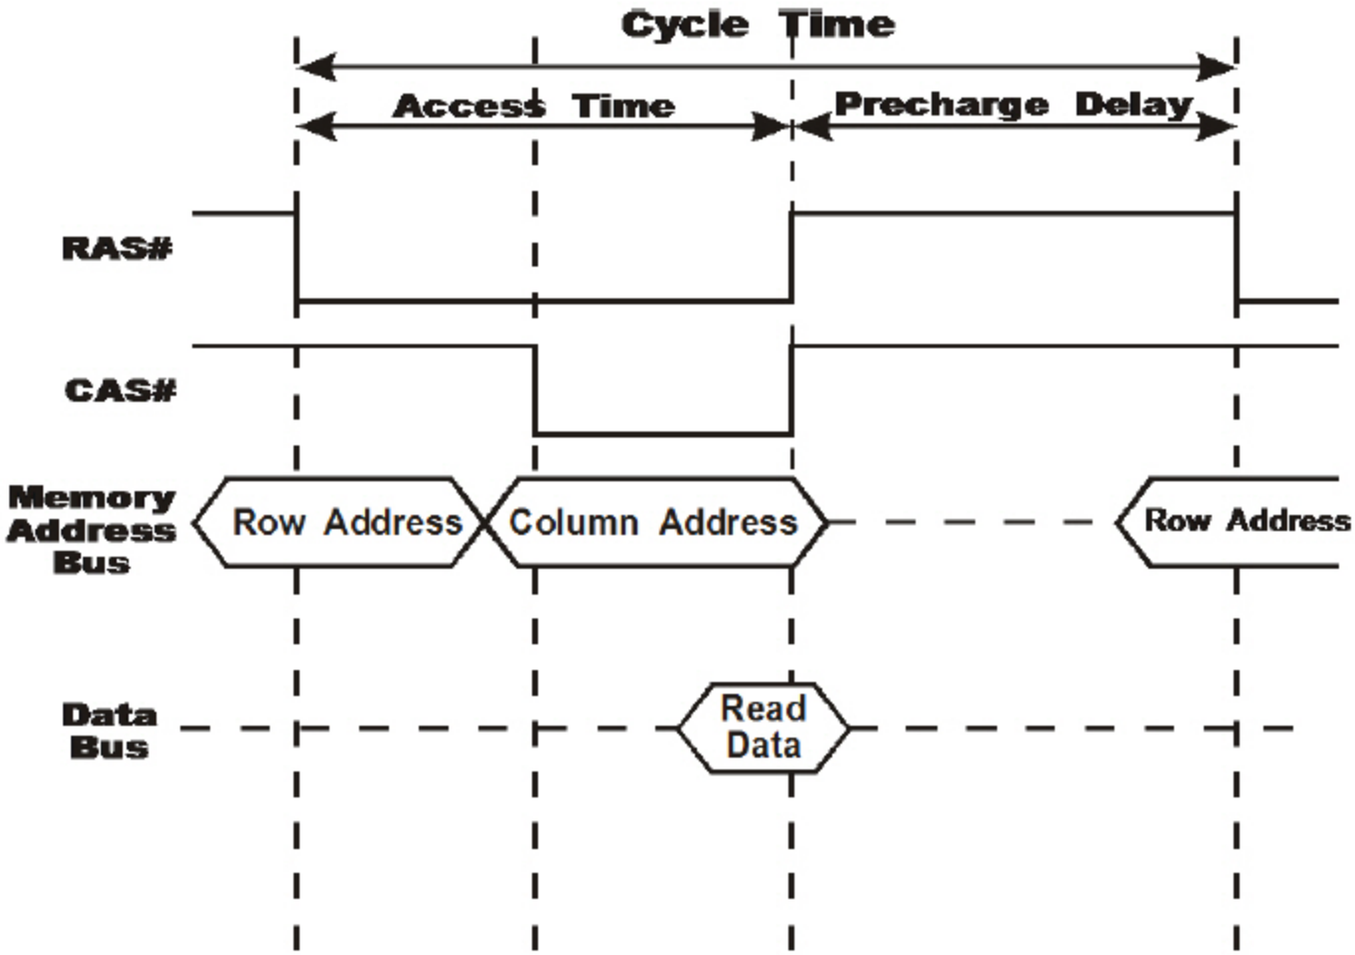

8088 System Memory Bus

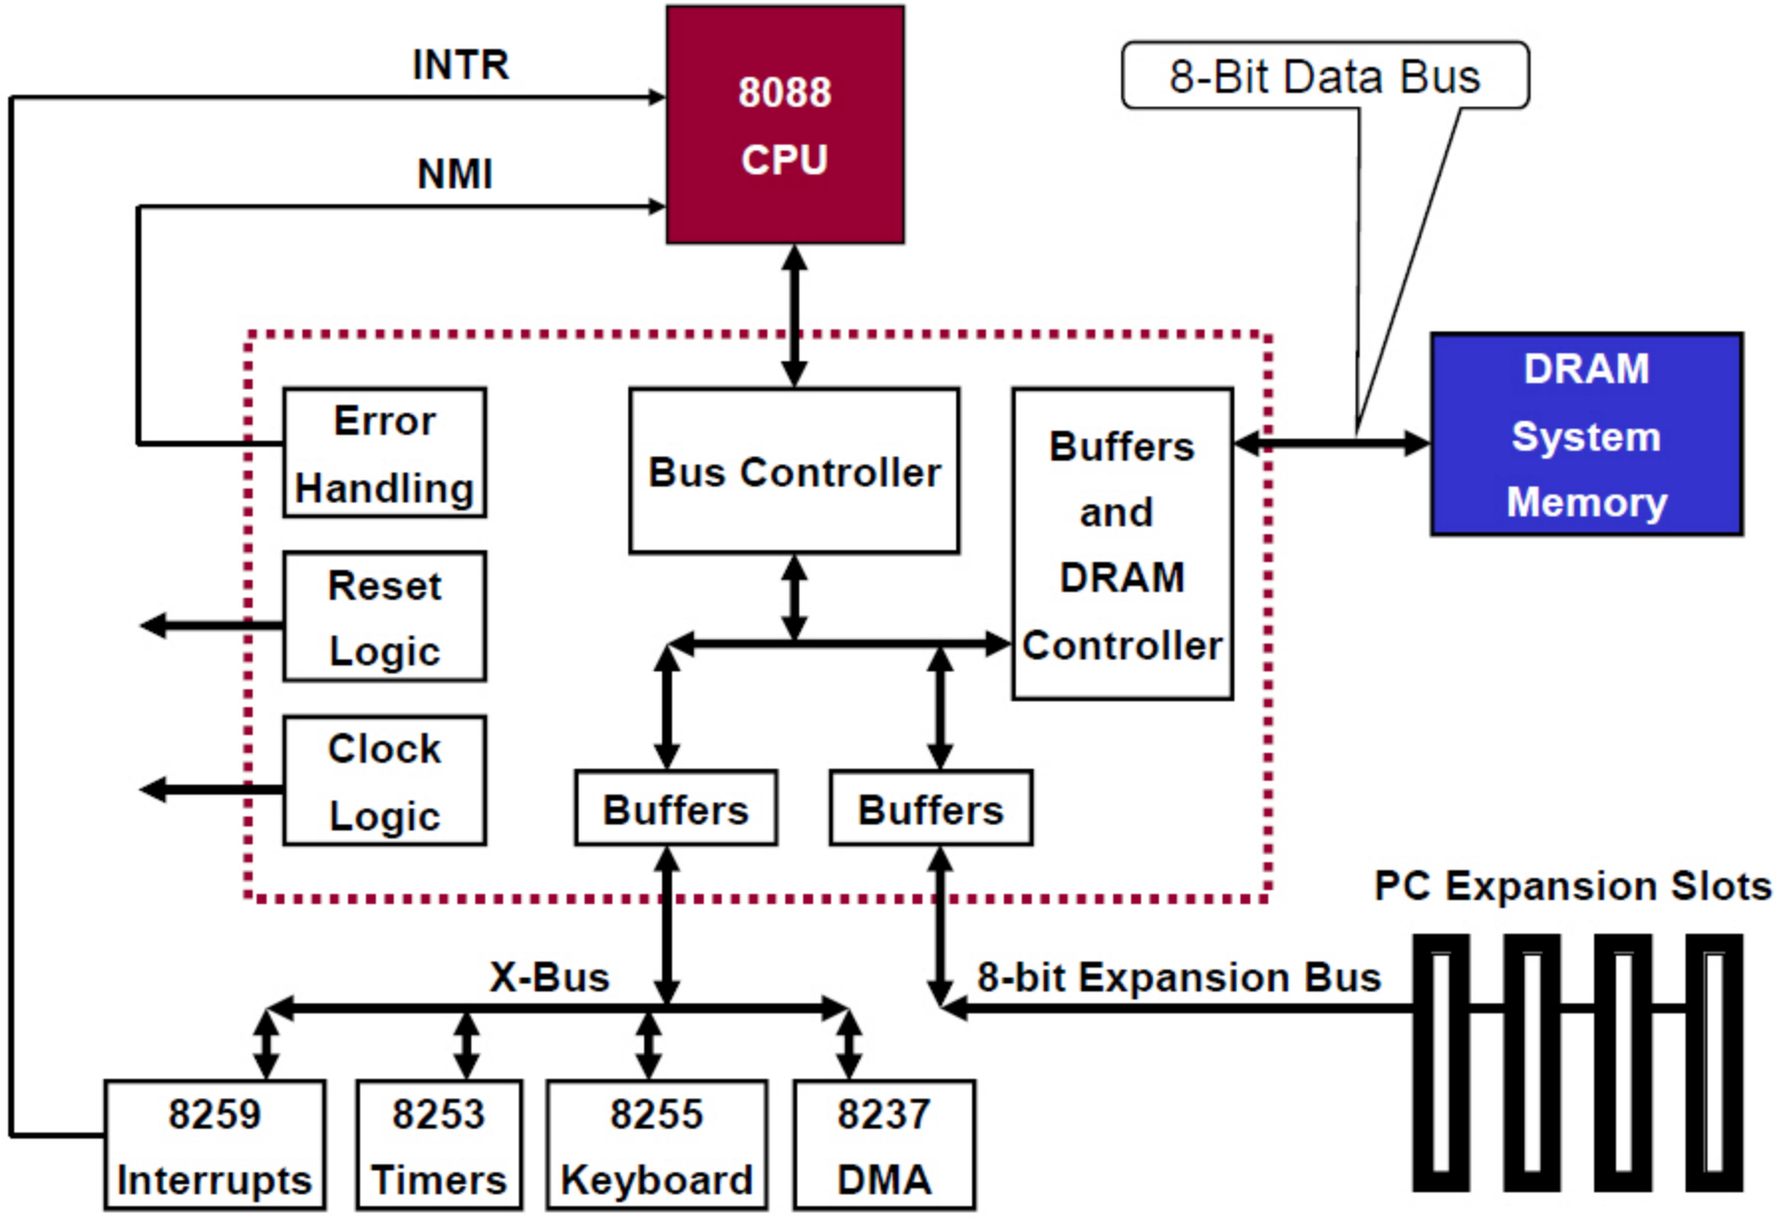

Increasing the Bus Size

- 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.

Multiple Data Lines

16M bit Array example; 4M X 4

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

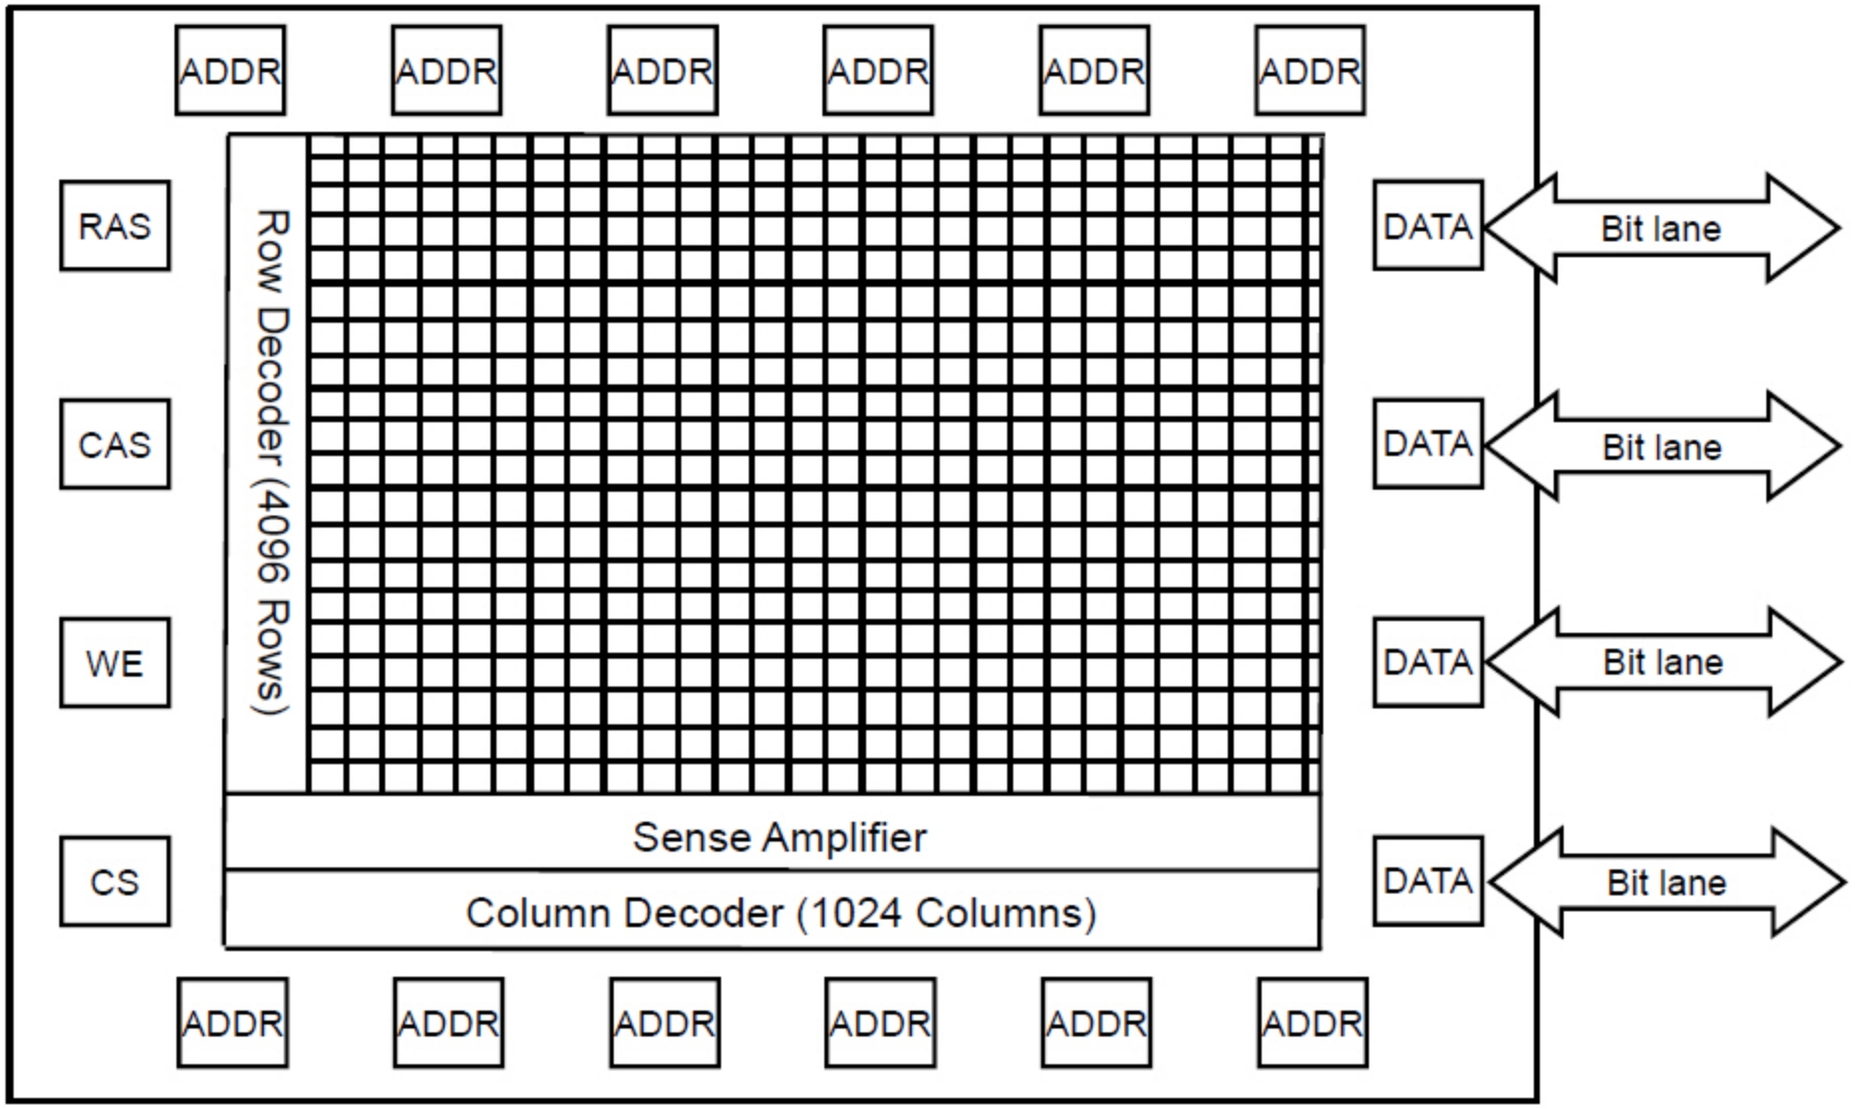

How to Increase Speed of the Chip

- 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.

Multiple Banks

16M bit Array example; 1M X 4 X 4

Shown as density X data width X #banks

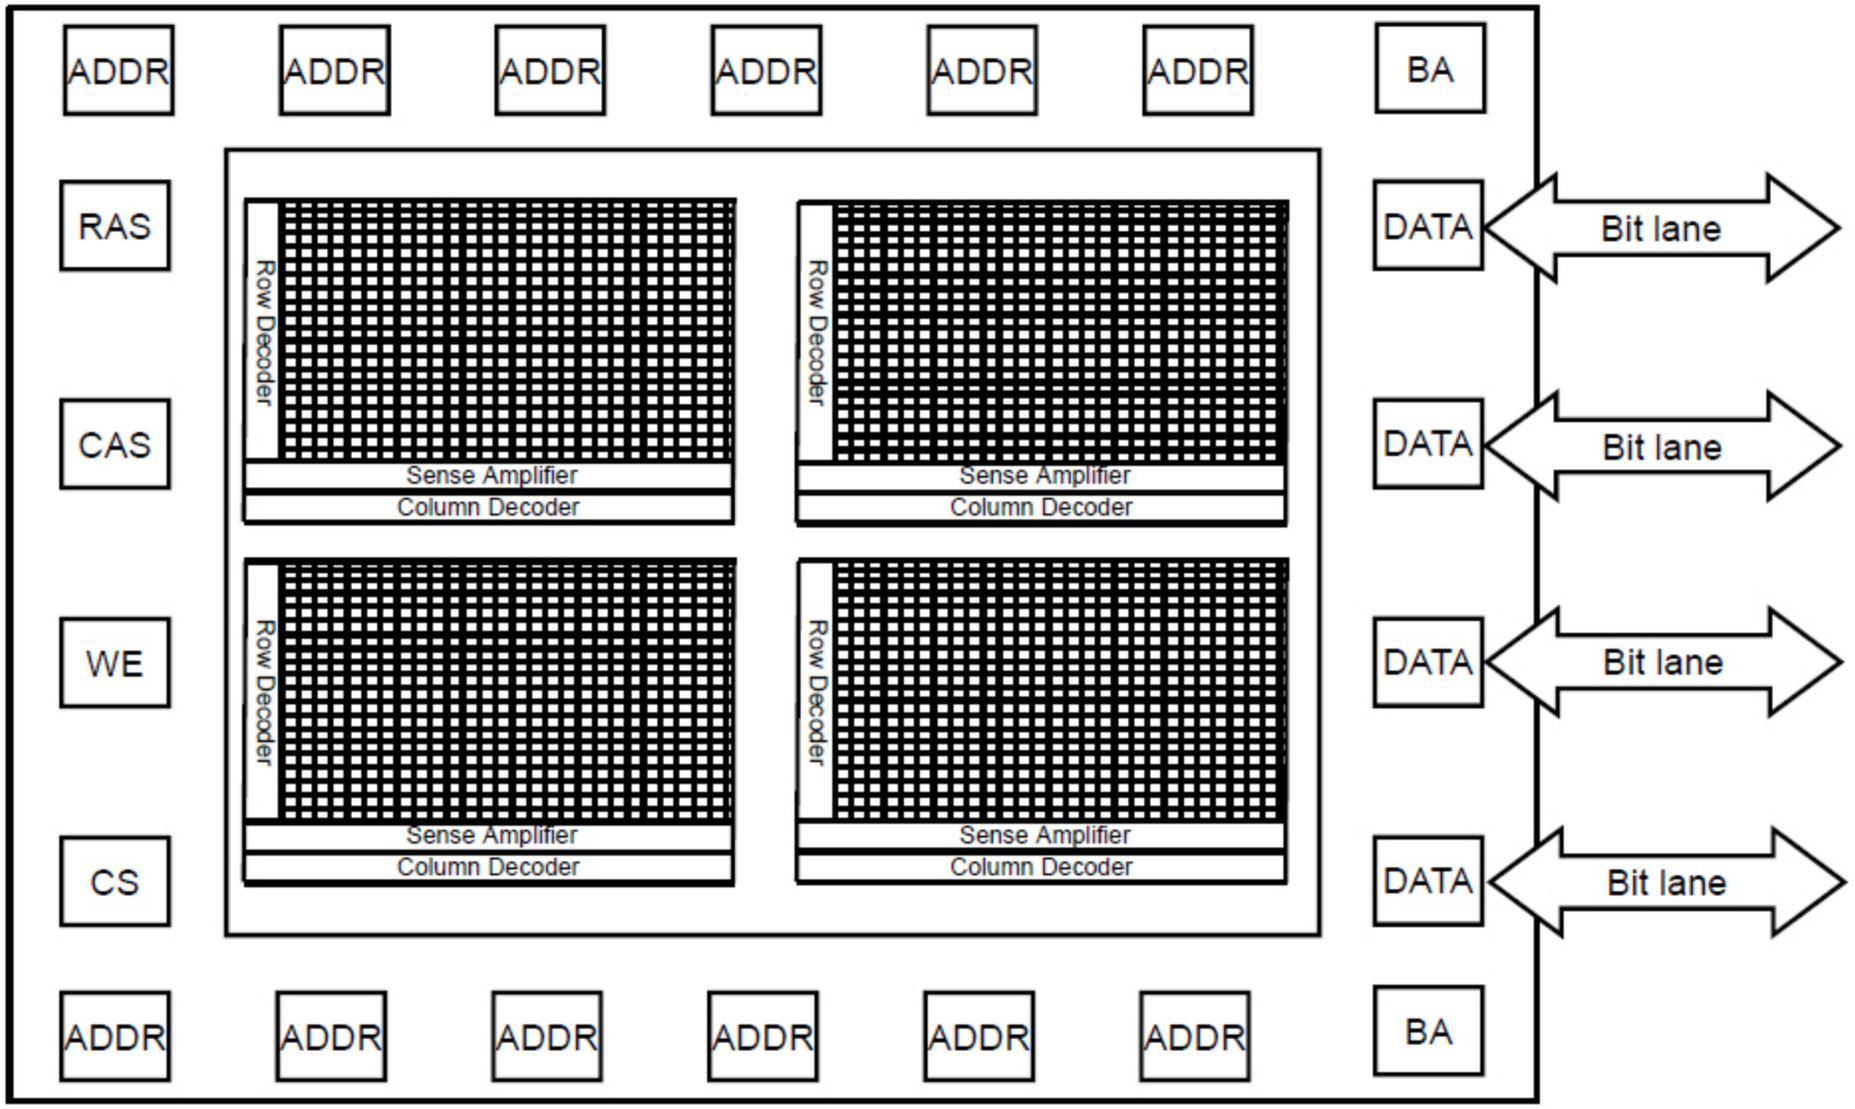

Multiple Data Lines

- 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

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

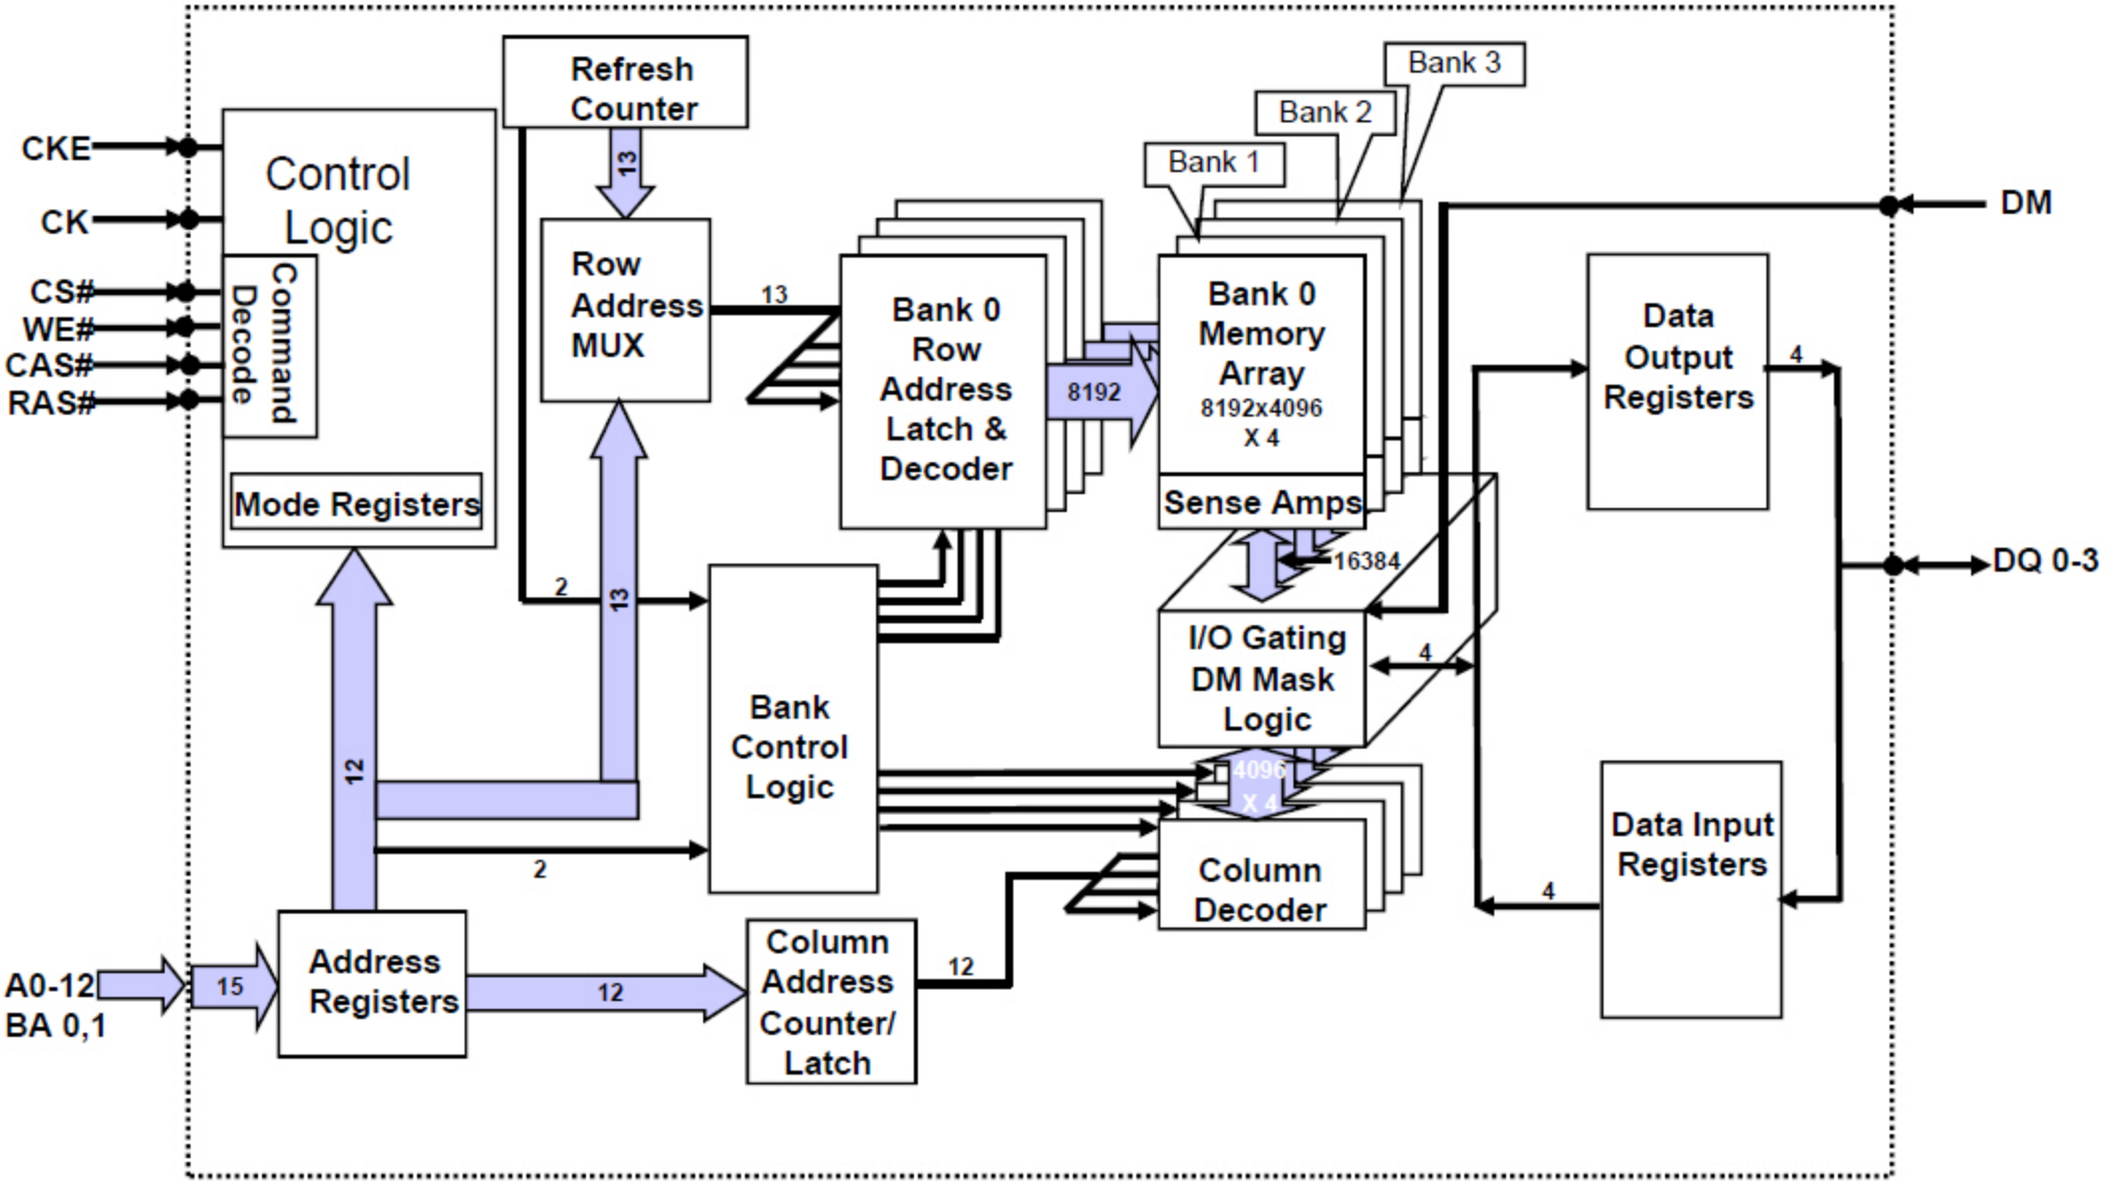

Features of a DDR1 Chip

- 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

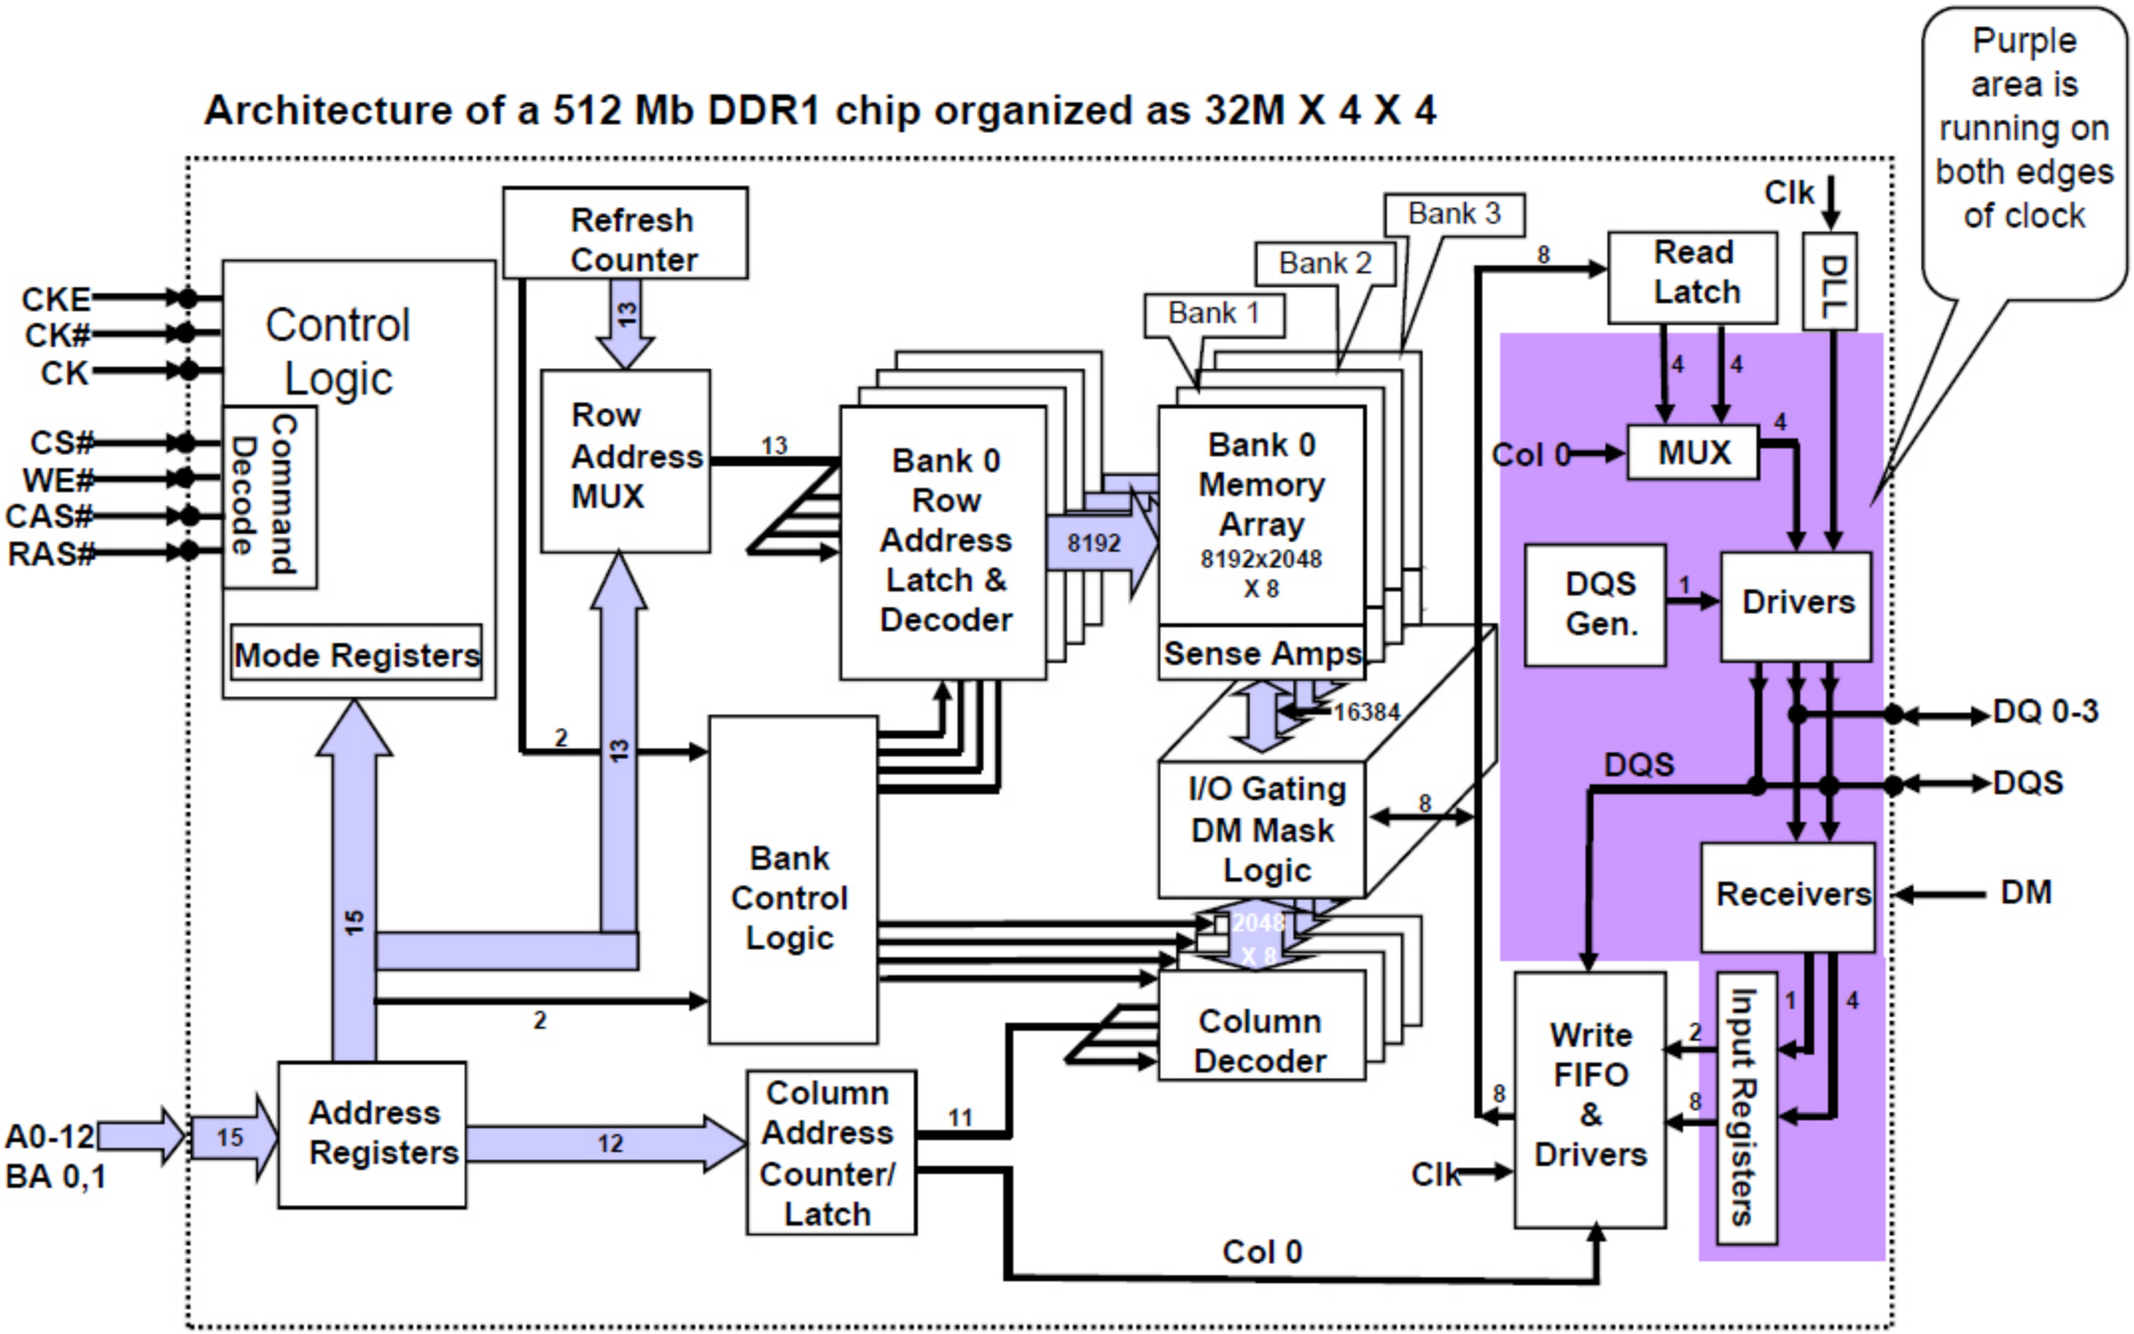

Features of a DDR2 Chip

- 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

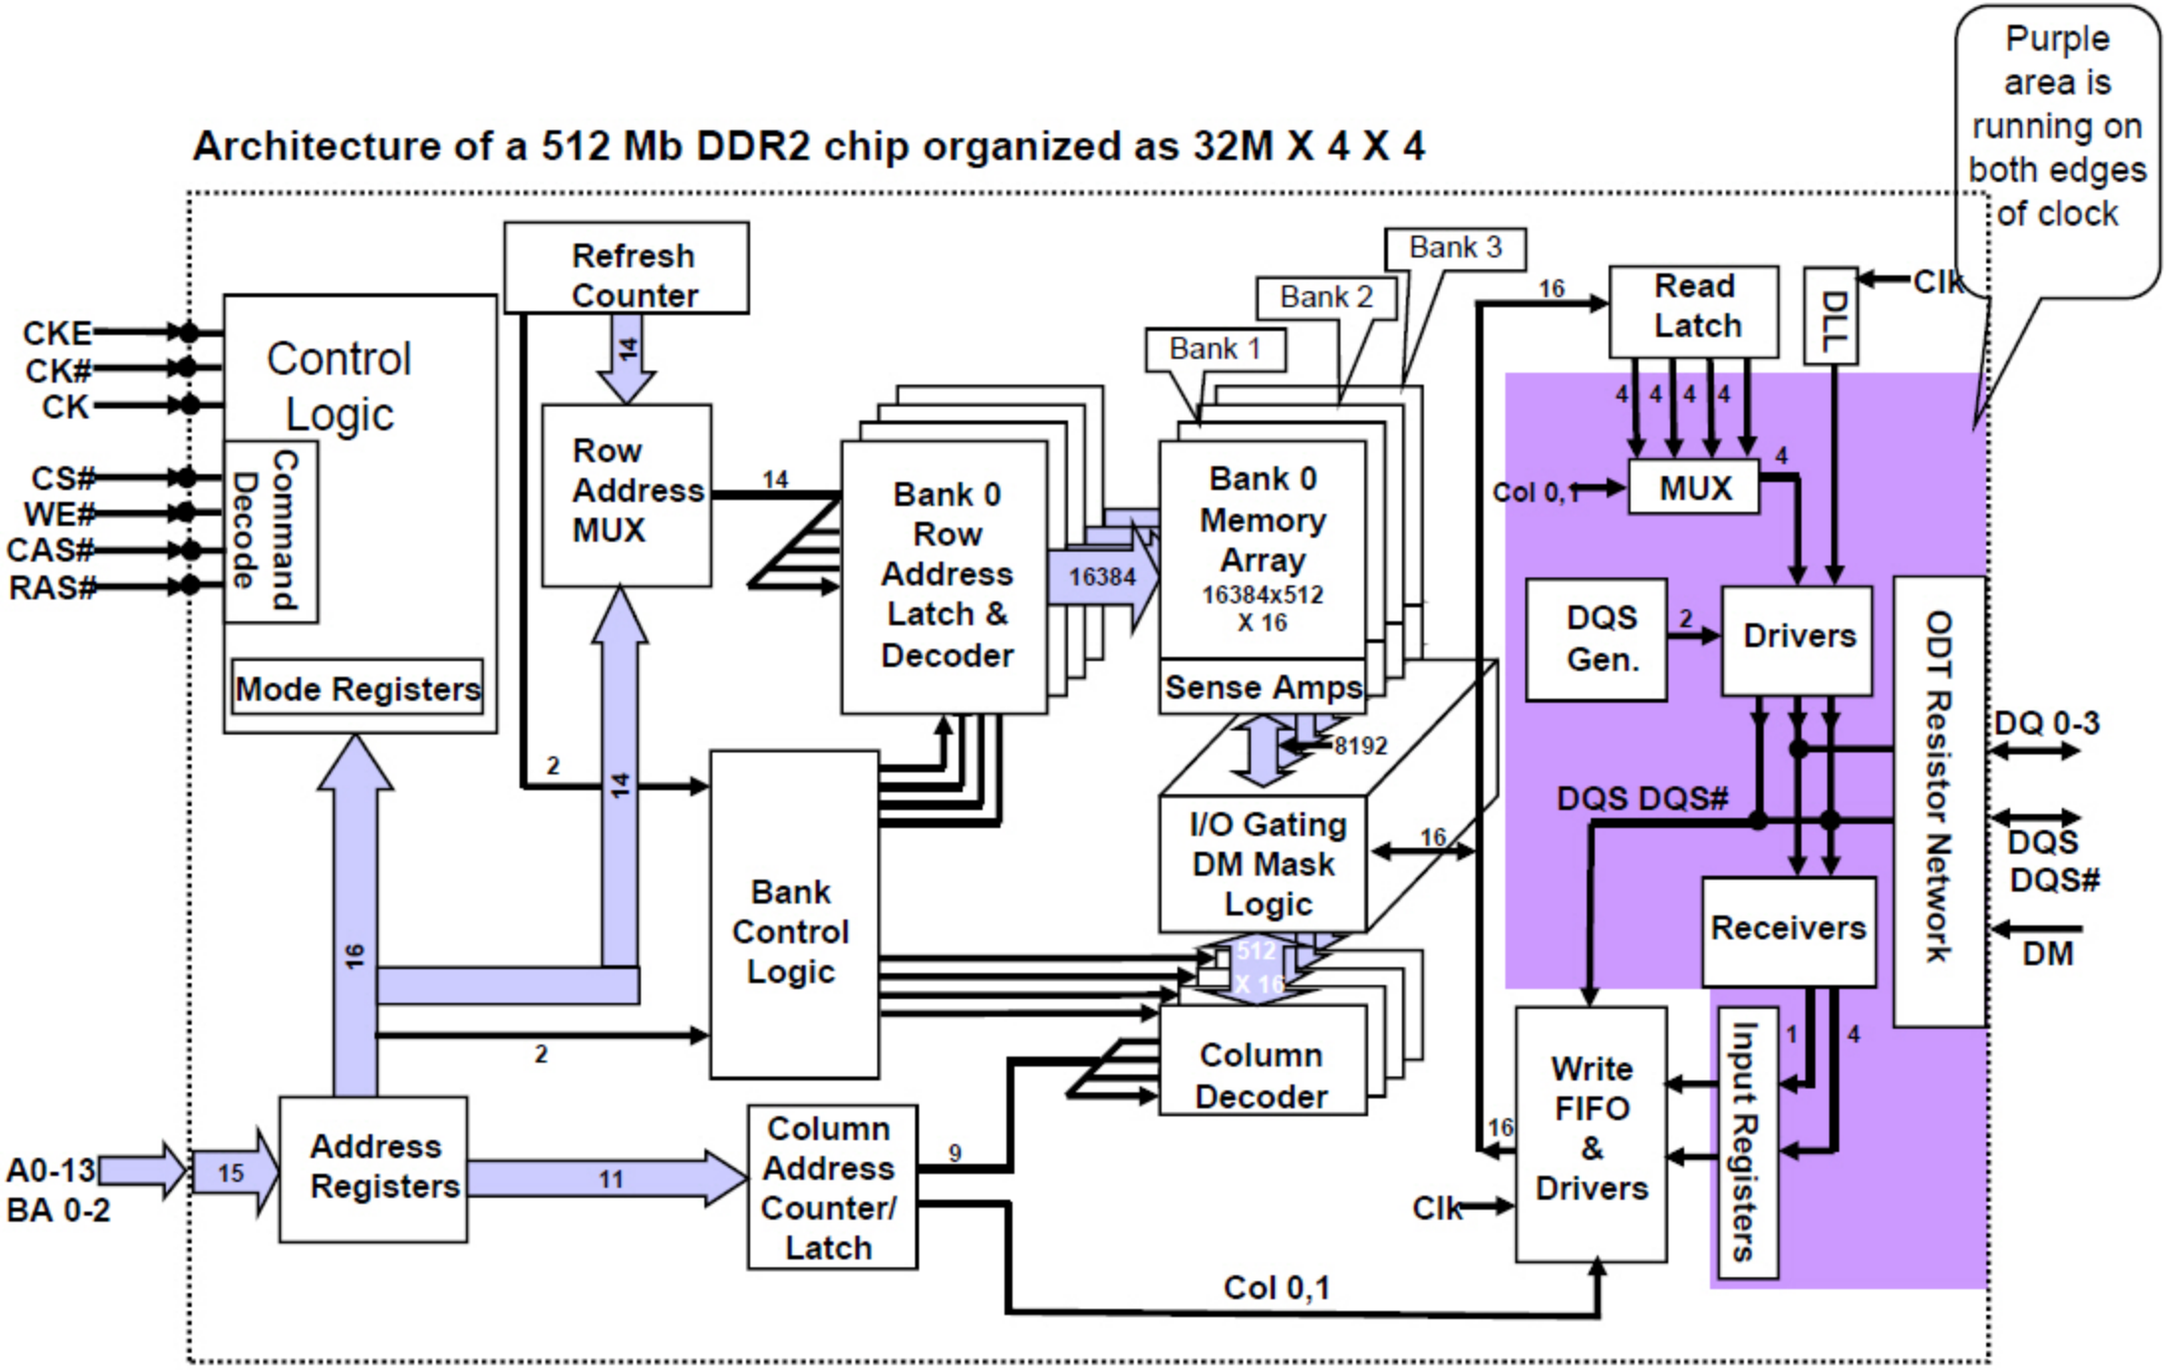

Features of a DDR3 Chip

- 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 ½ 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

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.

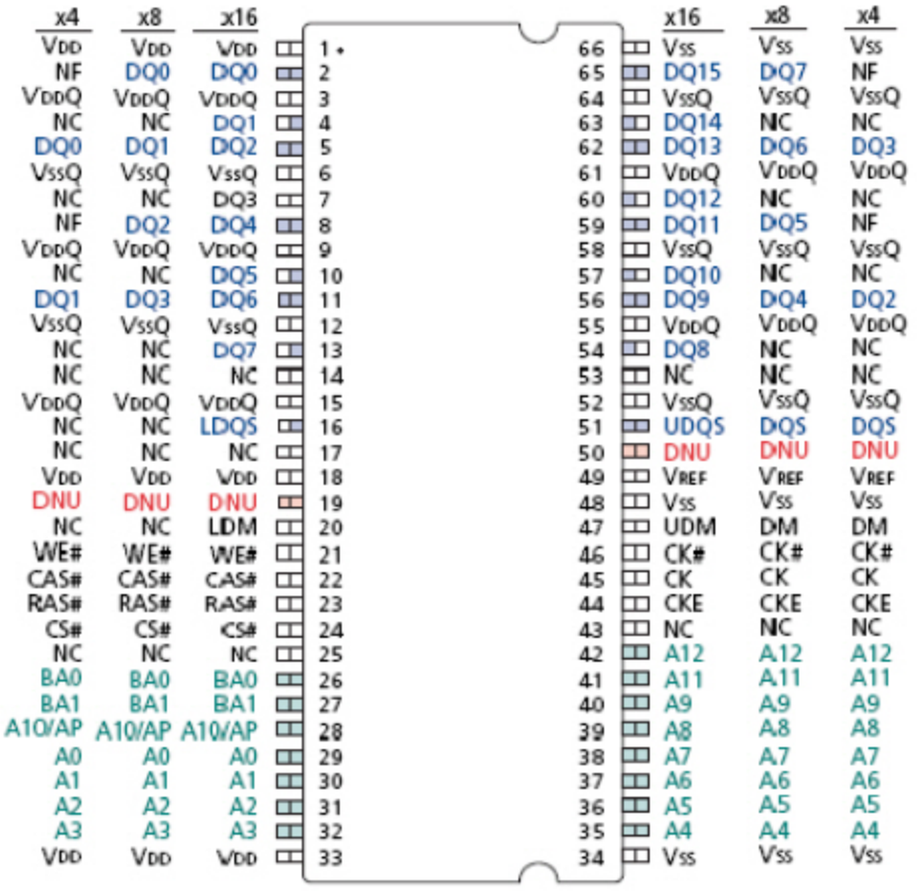

Top View

BGA Package

- 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.

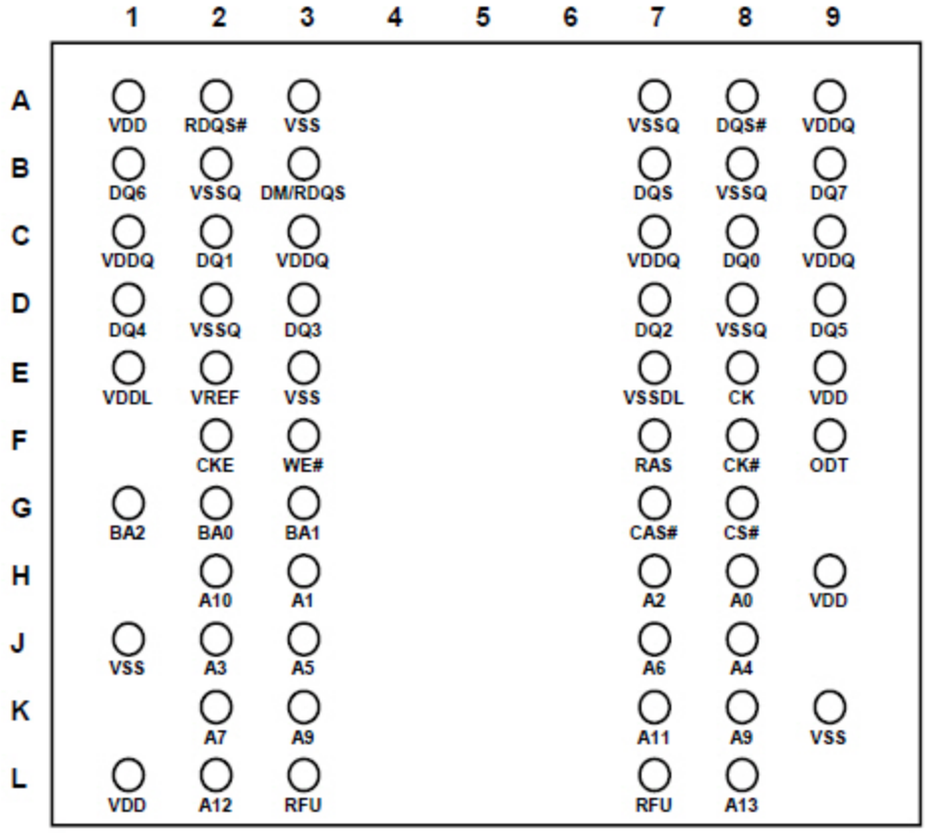

Looking through the package

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.

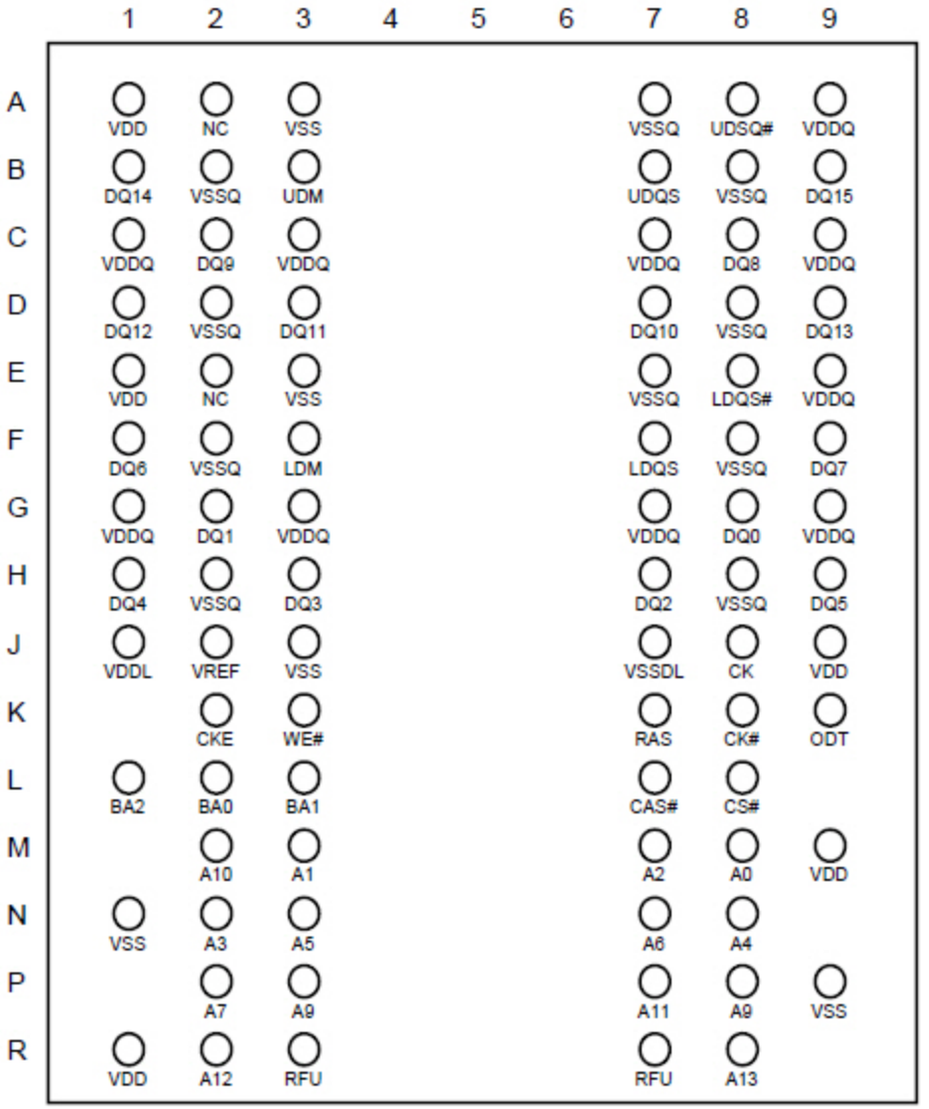

Looking Through from Top

MINDSHARE

BRINGING LIFE TO KNOWLEDGE

DRAM Modules

Dual Inline Memory Module (DIMM)

- 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.

Unbuffered Modules

- 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.

Registered Modules

- 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.

Ranks

- 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

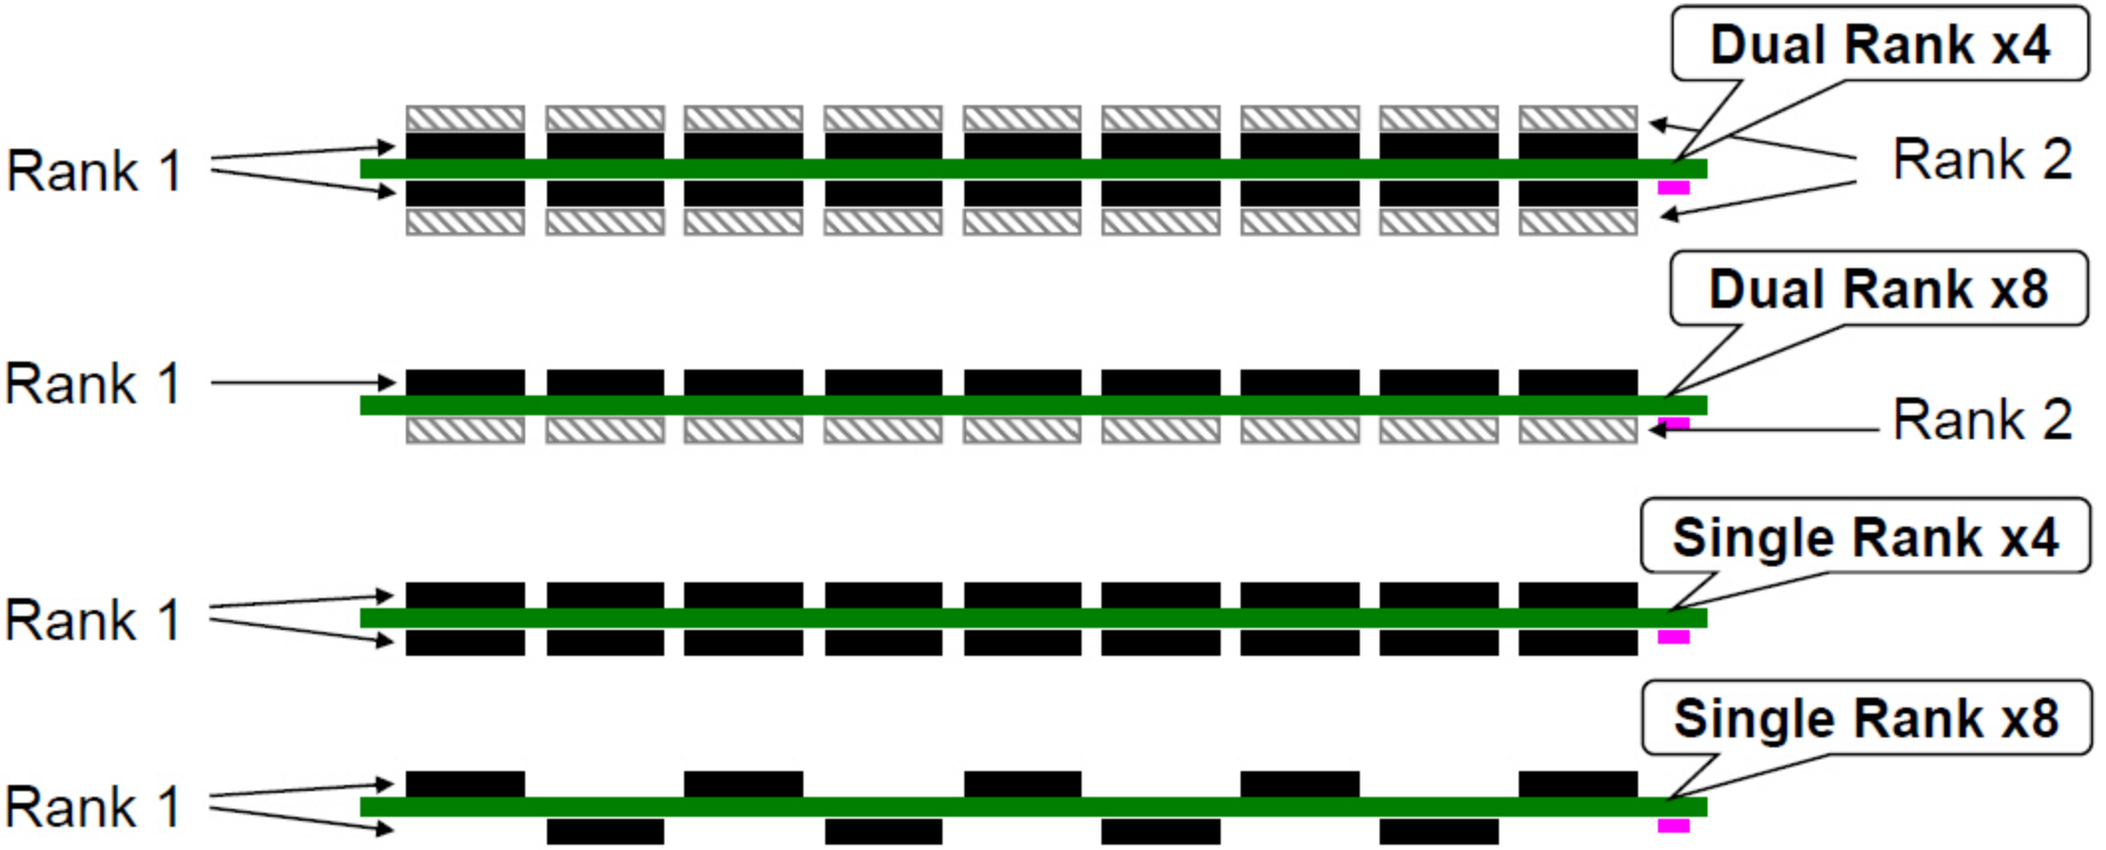

Module Types Unbuffered

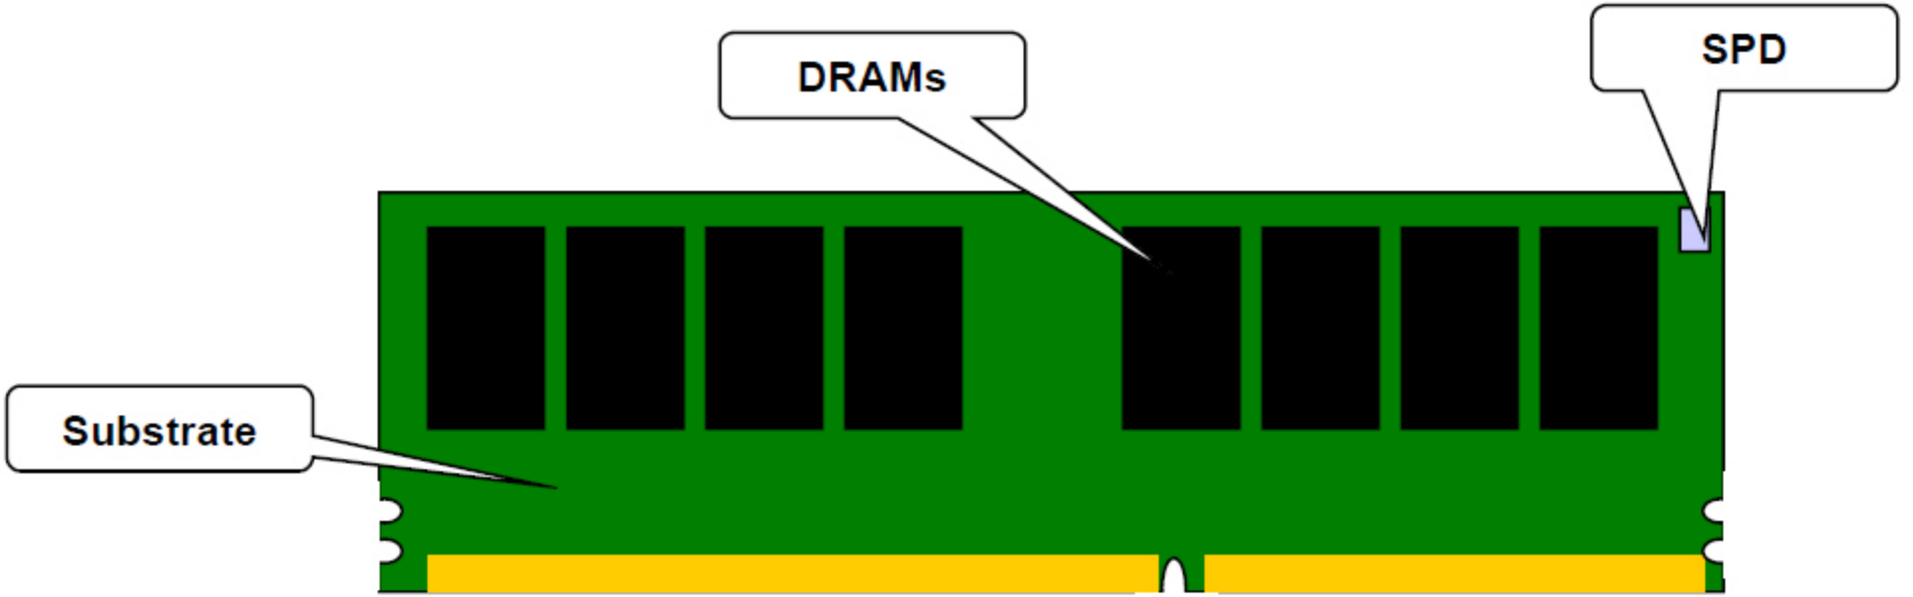

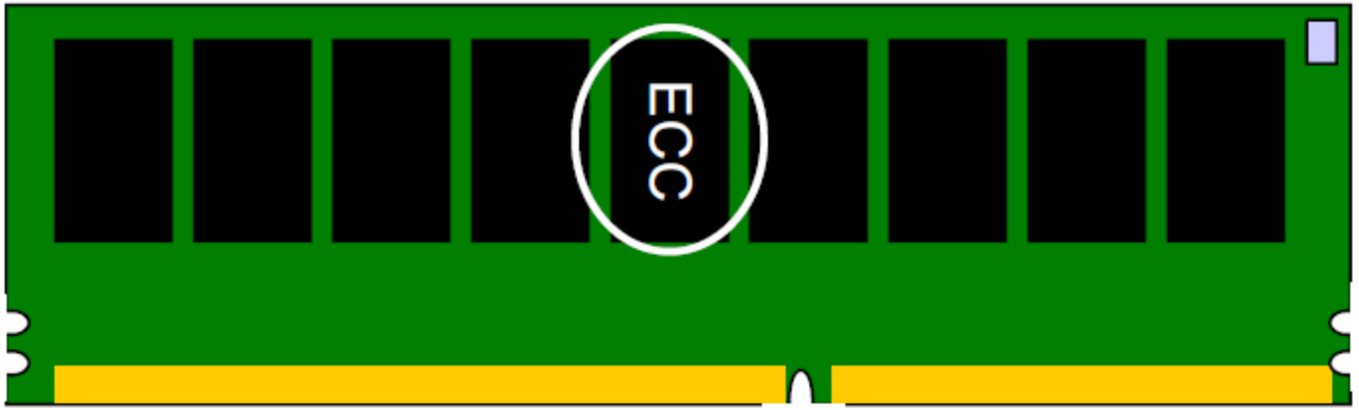

Shown with optional ECC chip. Typically unbuffered DIMMs do not have ECC support.

Unbuffered Logical Example

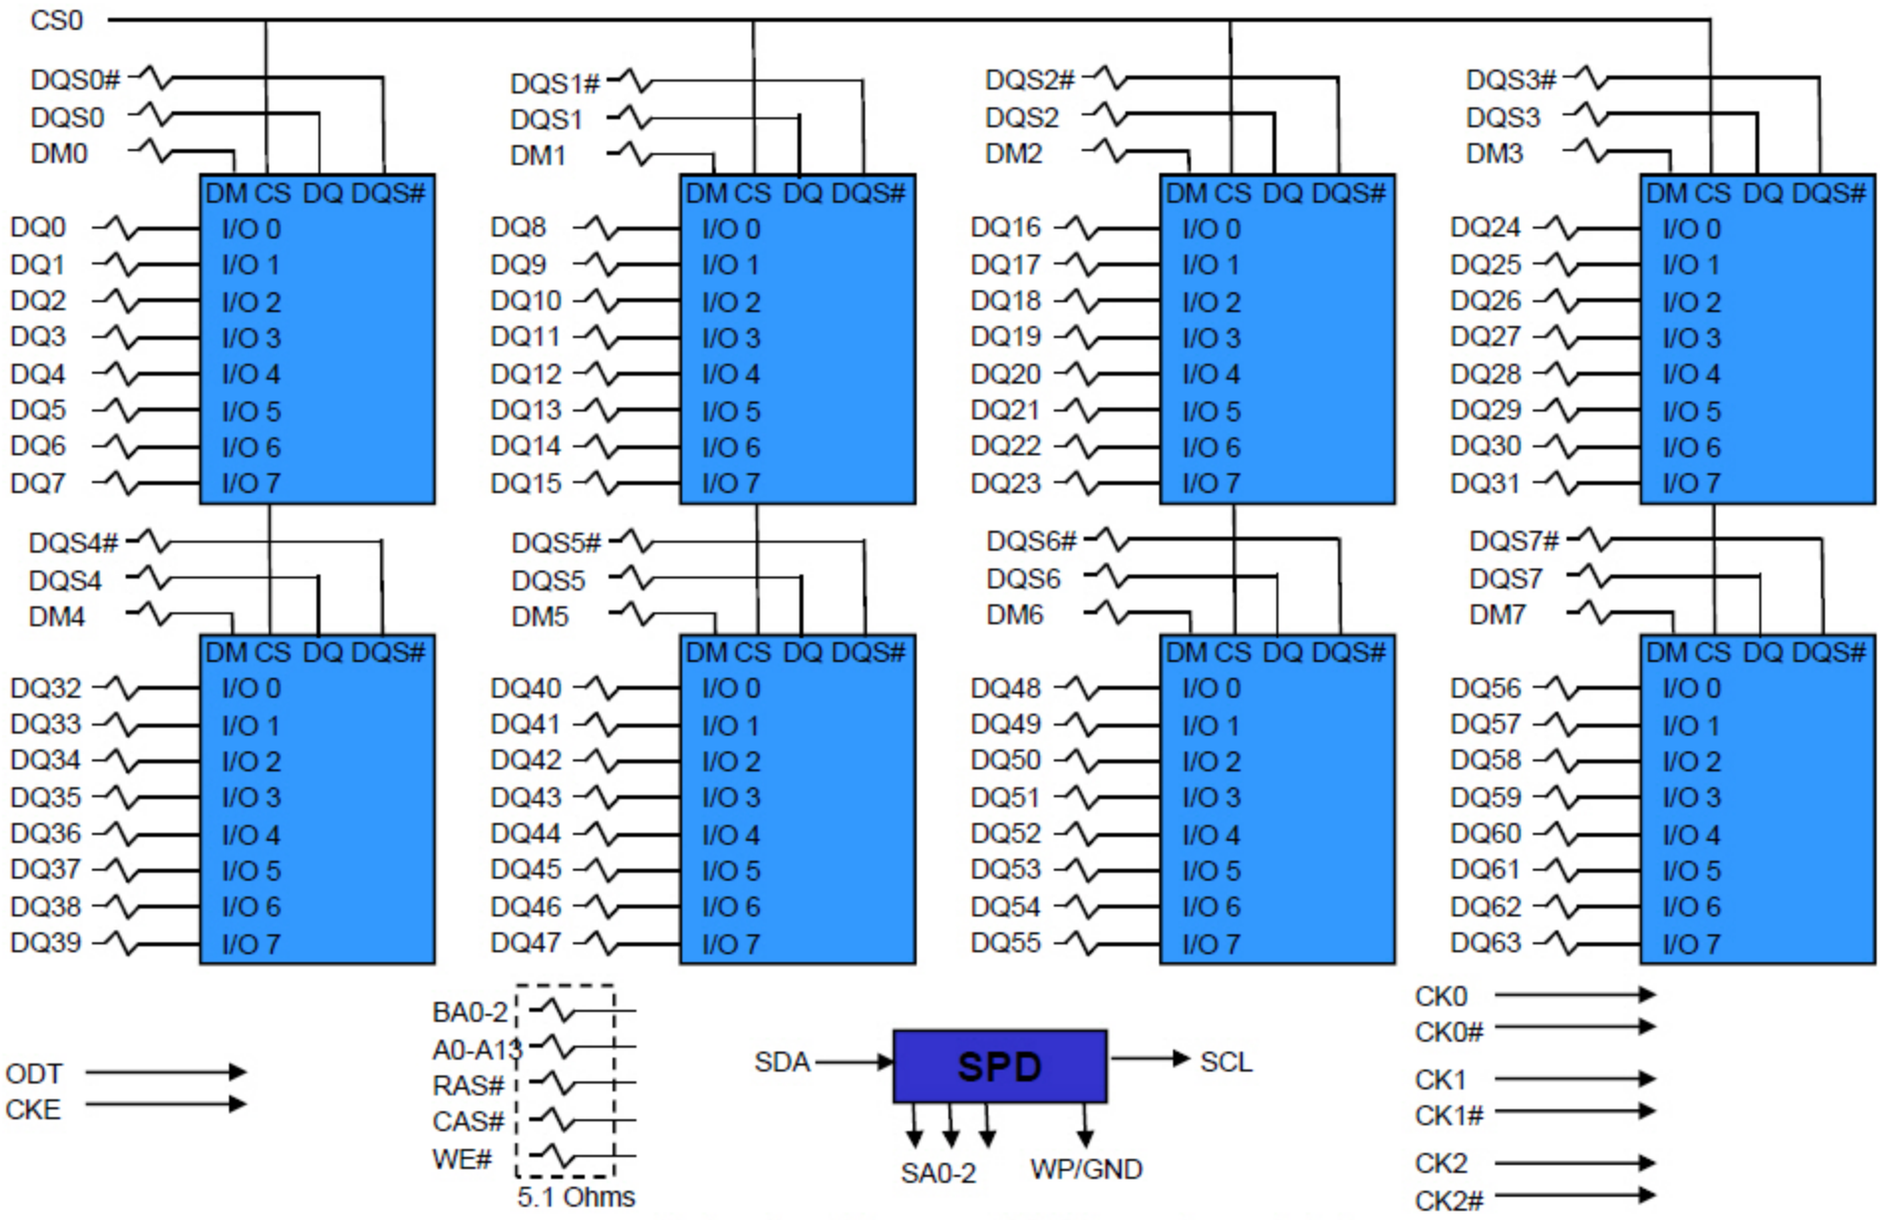

- Note all resistors are 22 Ohm unless stated

Rank Example

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

Unbuffered
Mobile Desktop

Registered
Server
Workstation

Configuration	# of chips	Ranks	Total Size
x16	4	Single	256MB
x16	8	Dual	512MB
x8	8	Single	512MB
x8	16	Dual	1GB
x4	16	Single	1GB
x4	32	Dual	2GB

x8 Single Rank, 8 components

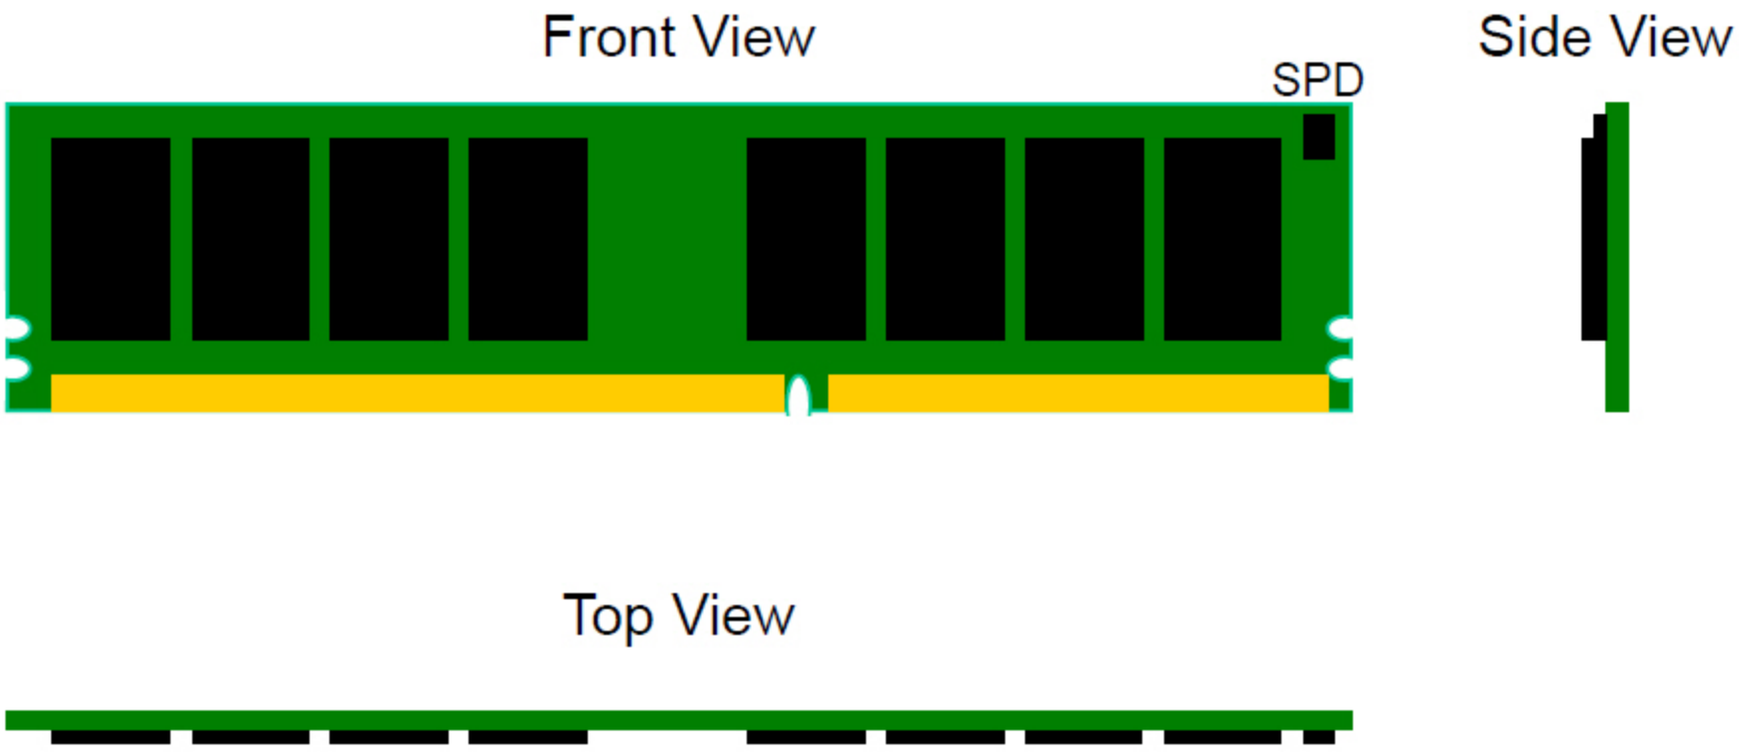

x8 Dual Rank, 16 components

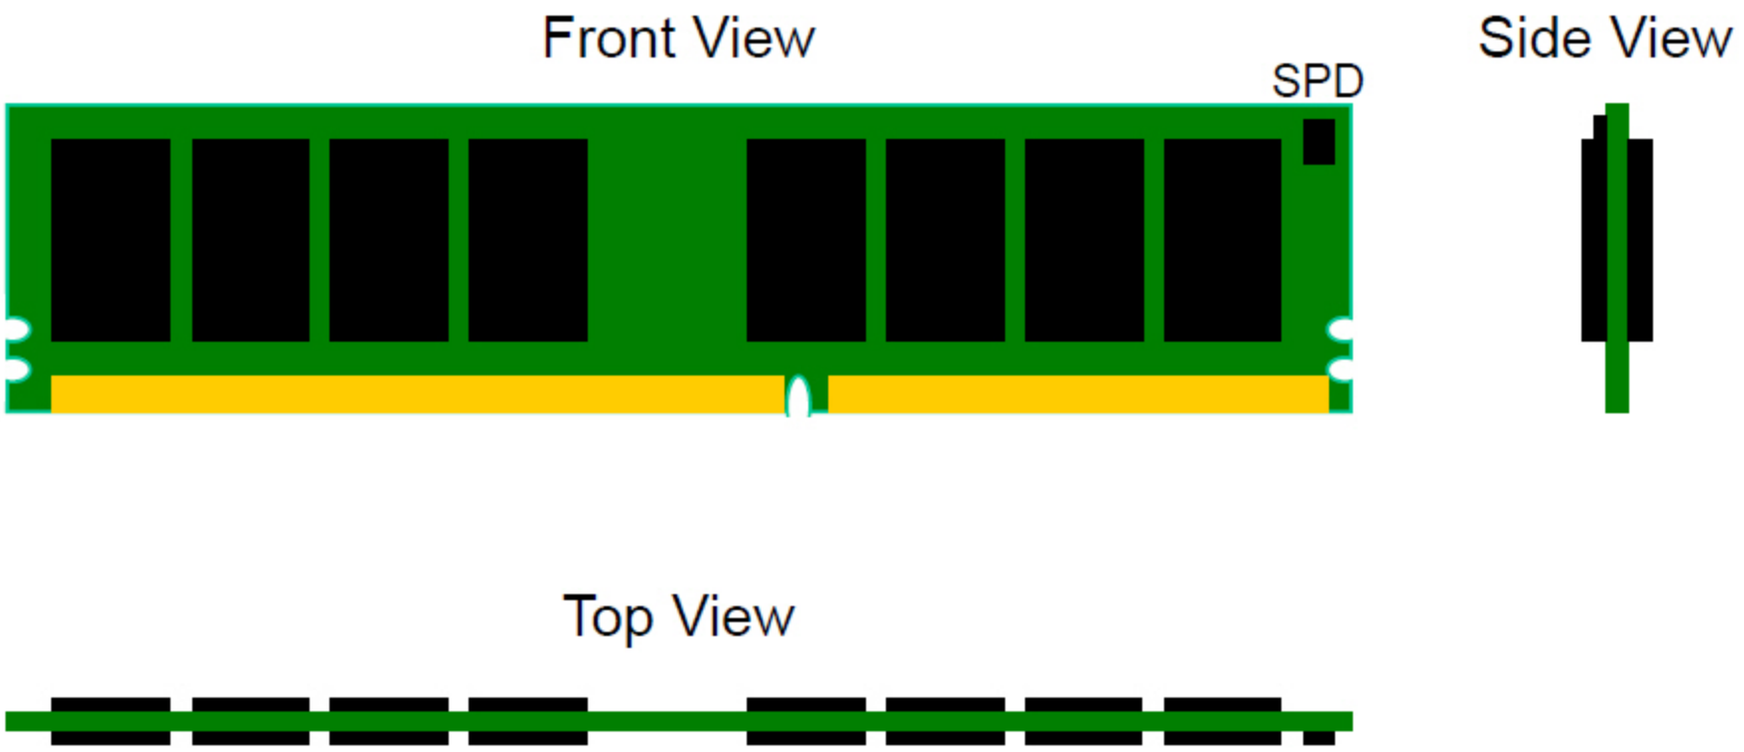

x16 Single Rank, 4 components

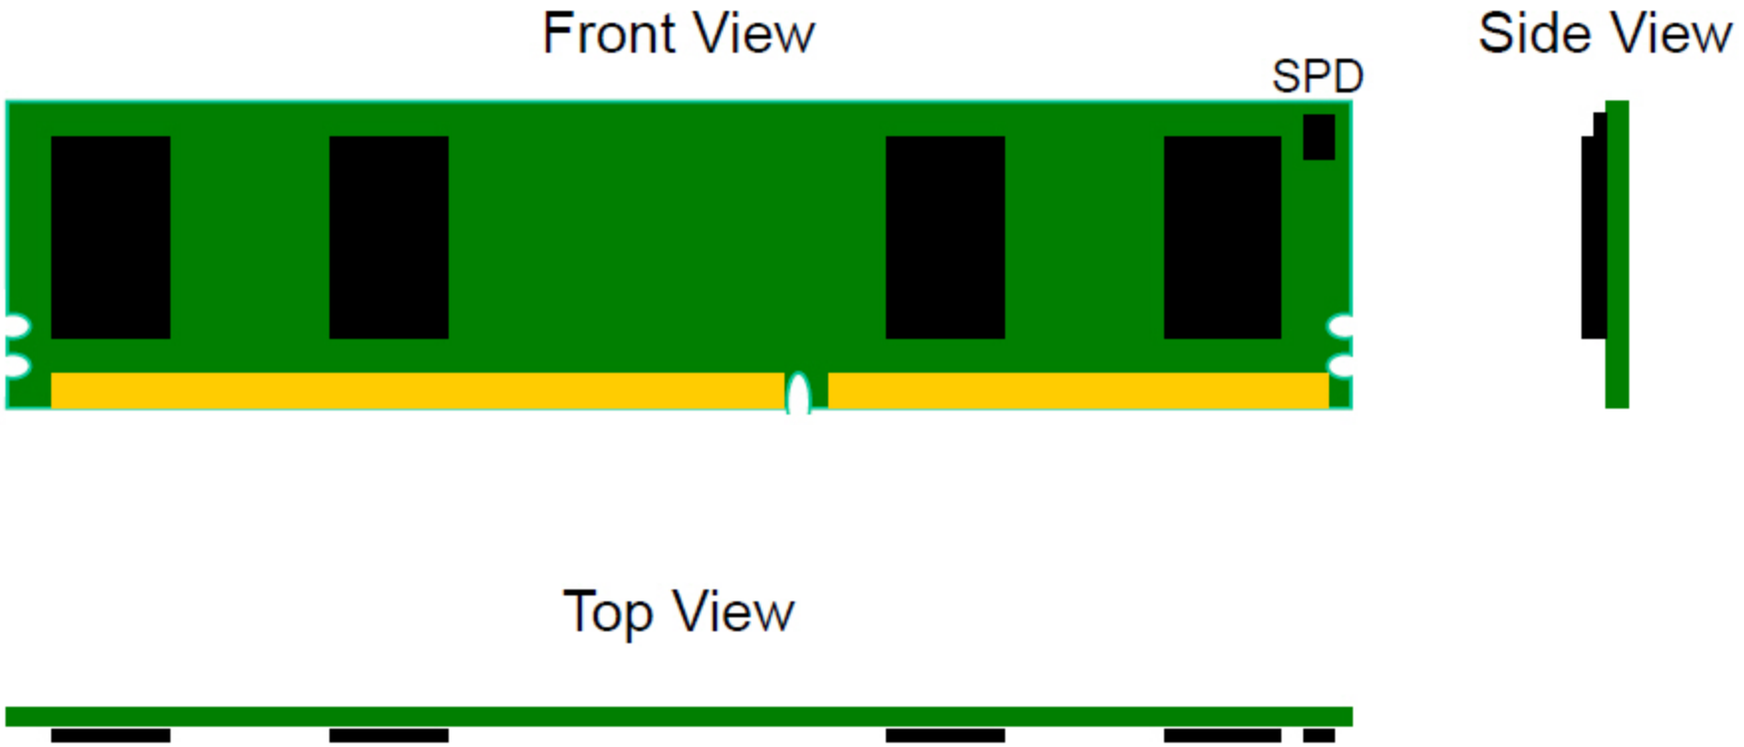

Module Types Registered

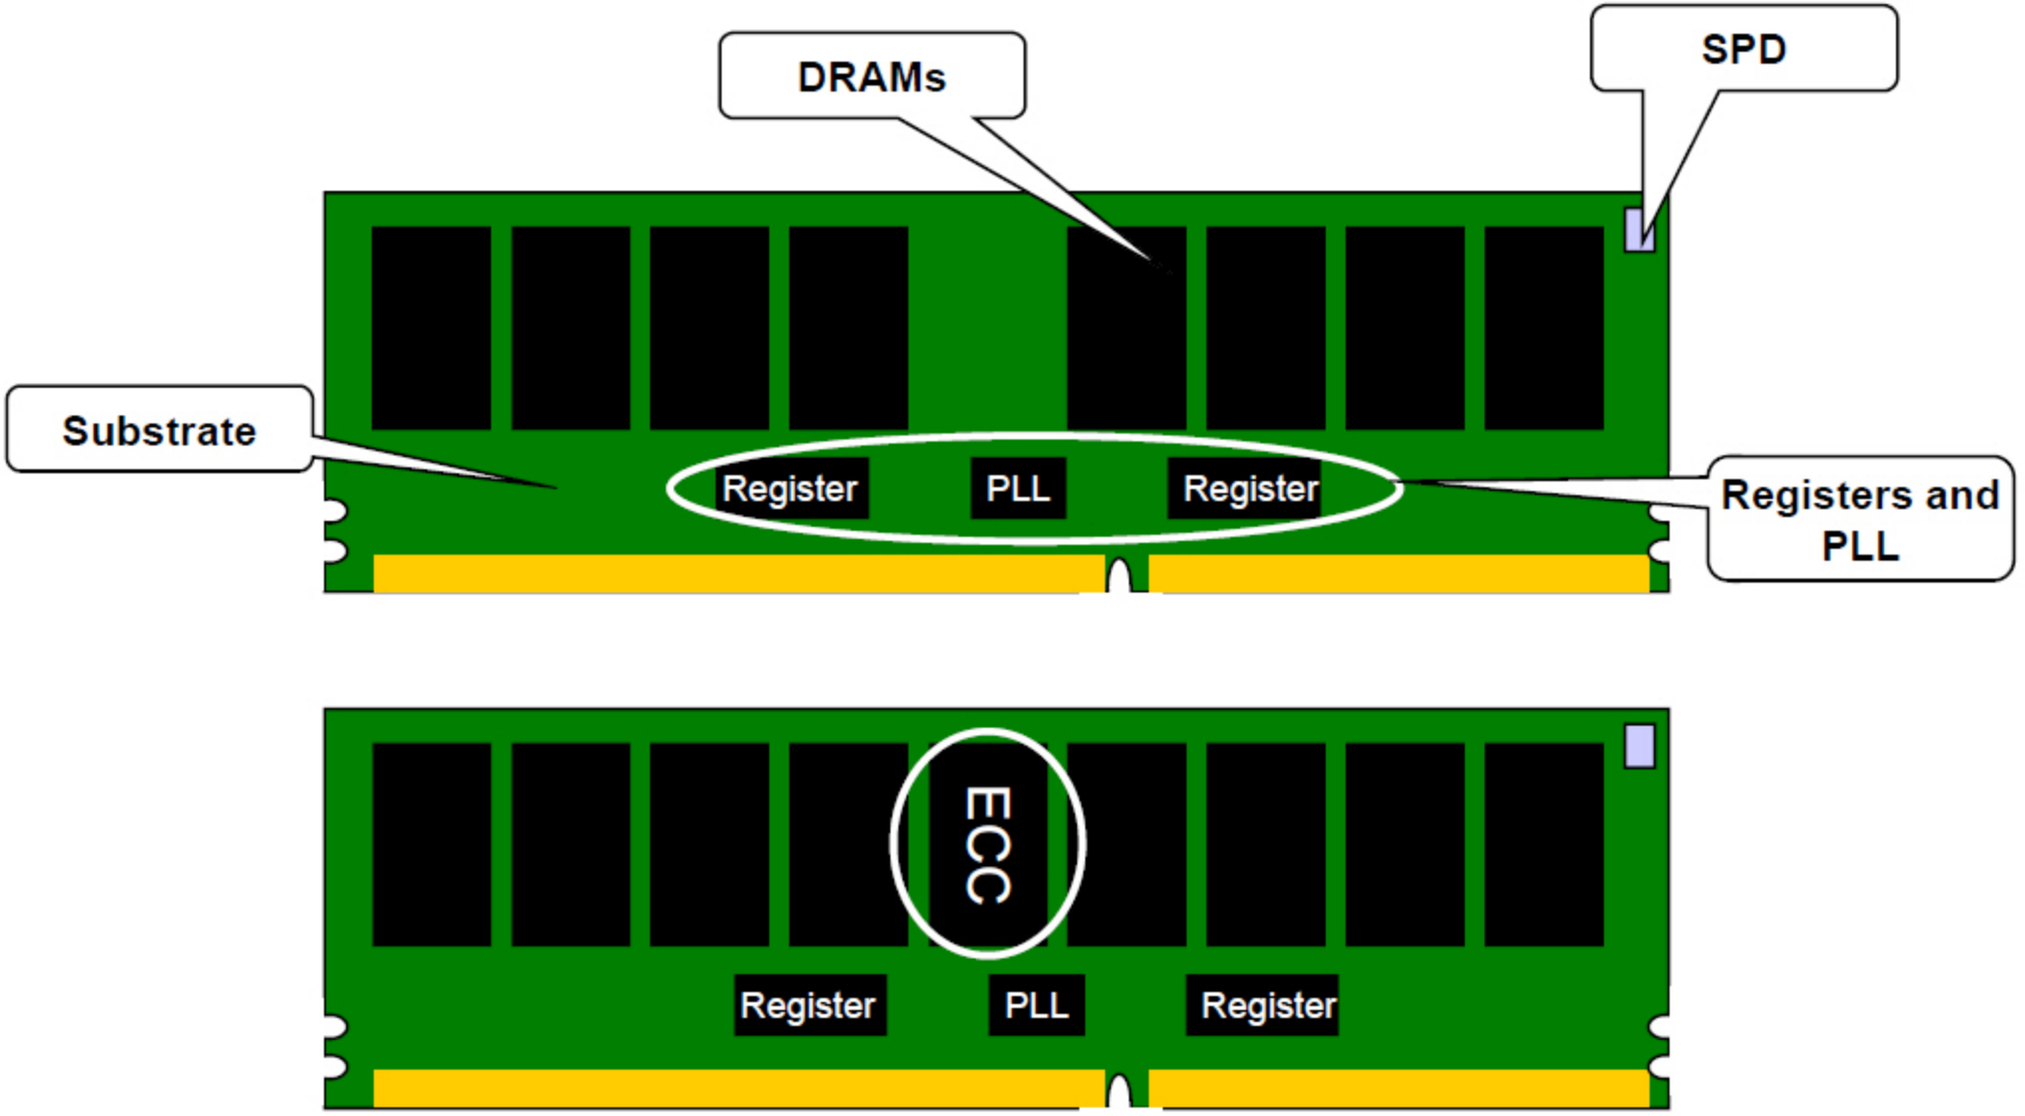

x4 Single Rank, 18 components (with ECC)

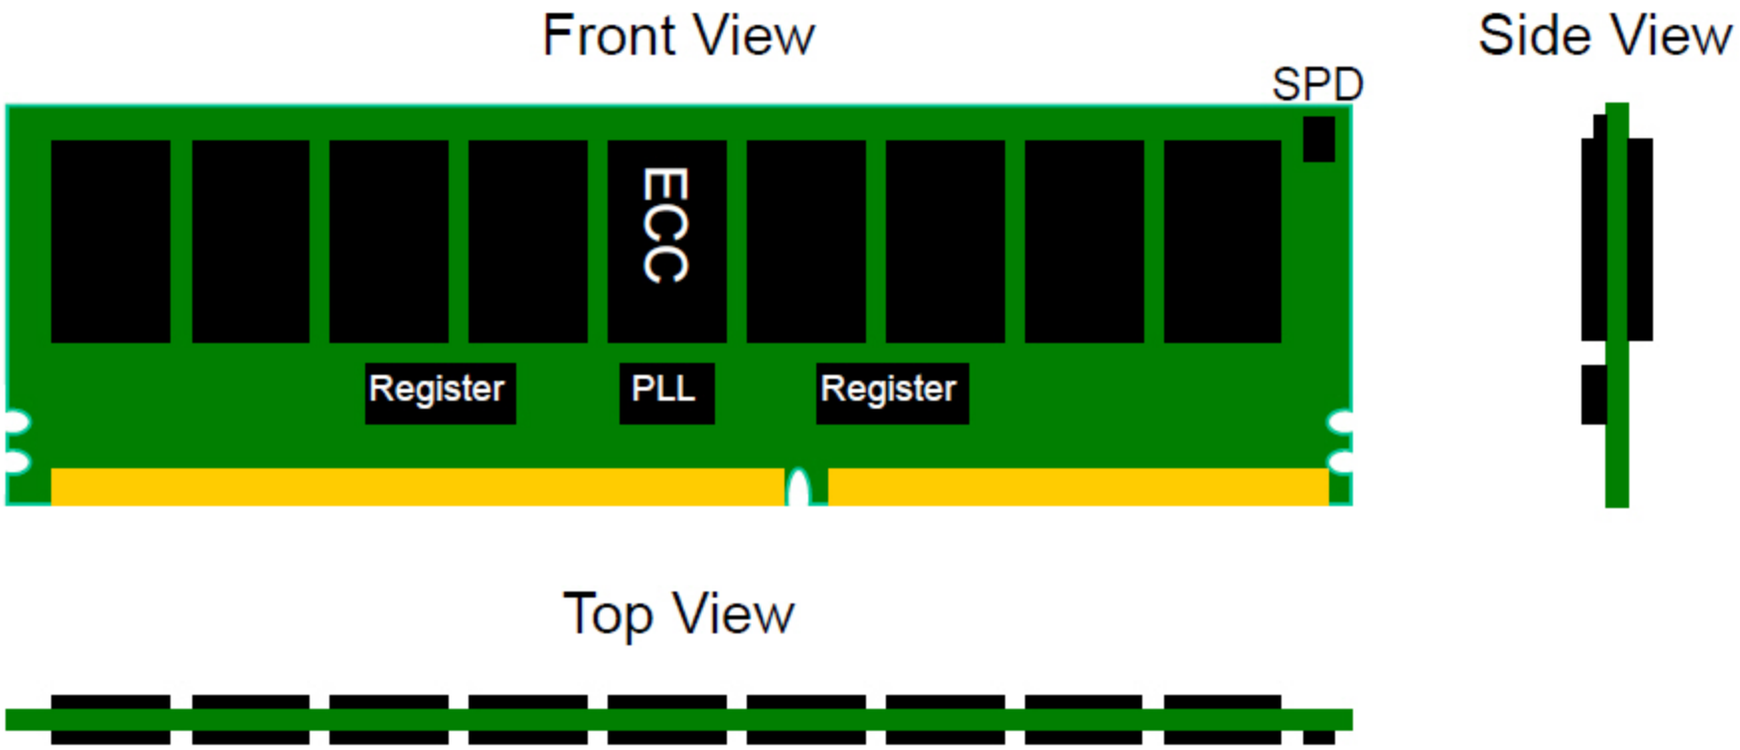

x4 Dual Rank, 36 components (with ECC)

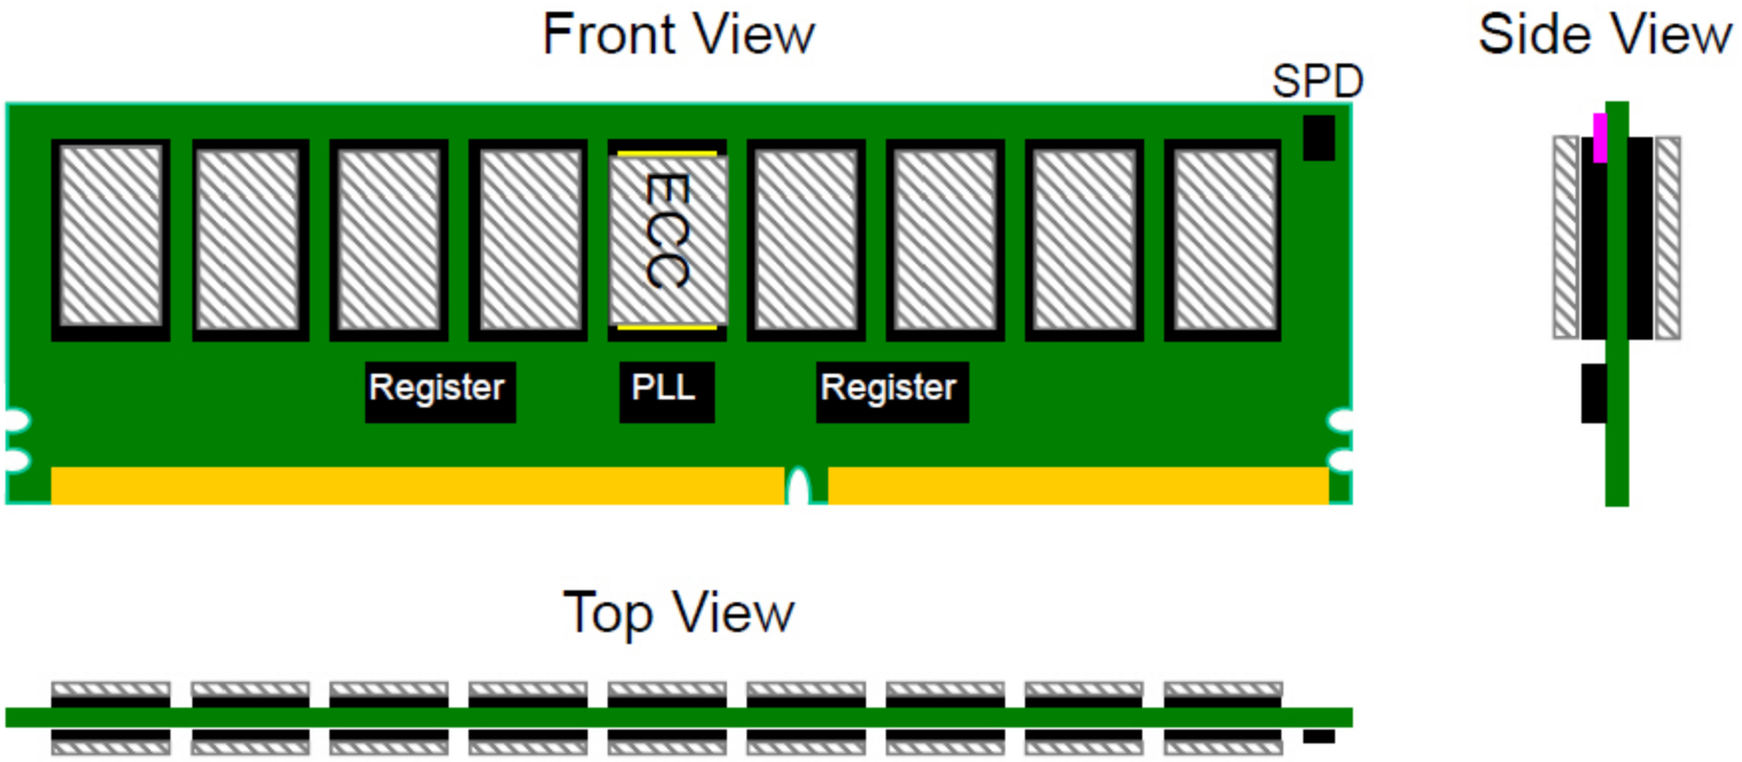

DIMM Configuration Registered

- x8 Single Rank, 8 components (with ECC)

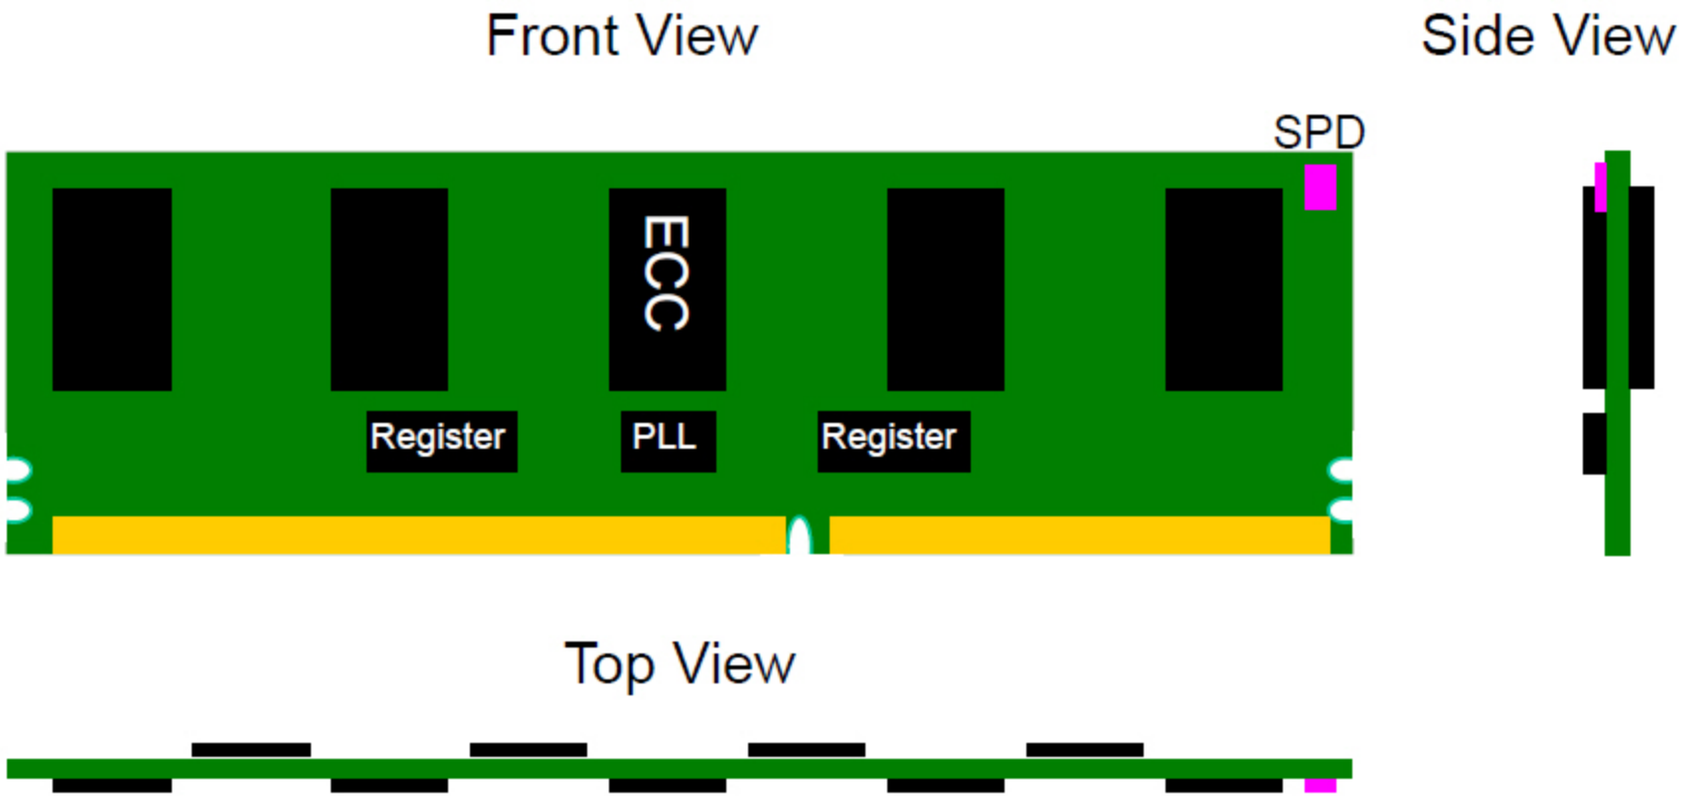

DIMM Configuration Registered

- x8 Dual Rank, 18 components (with ECC)

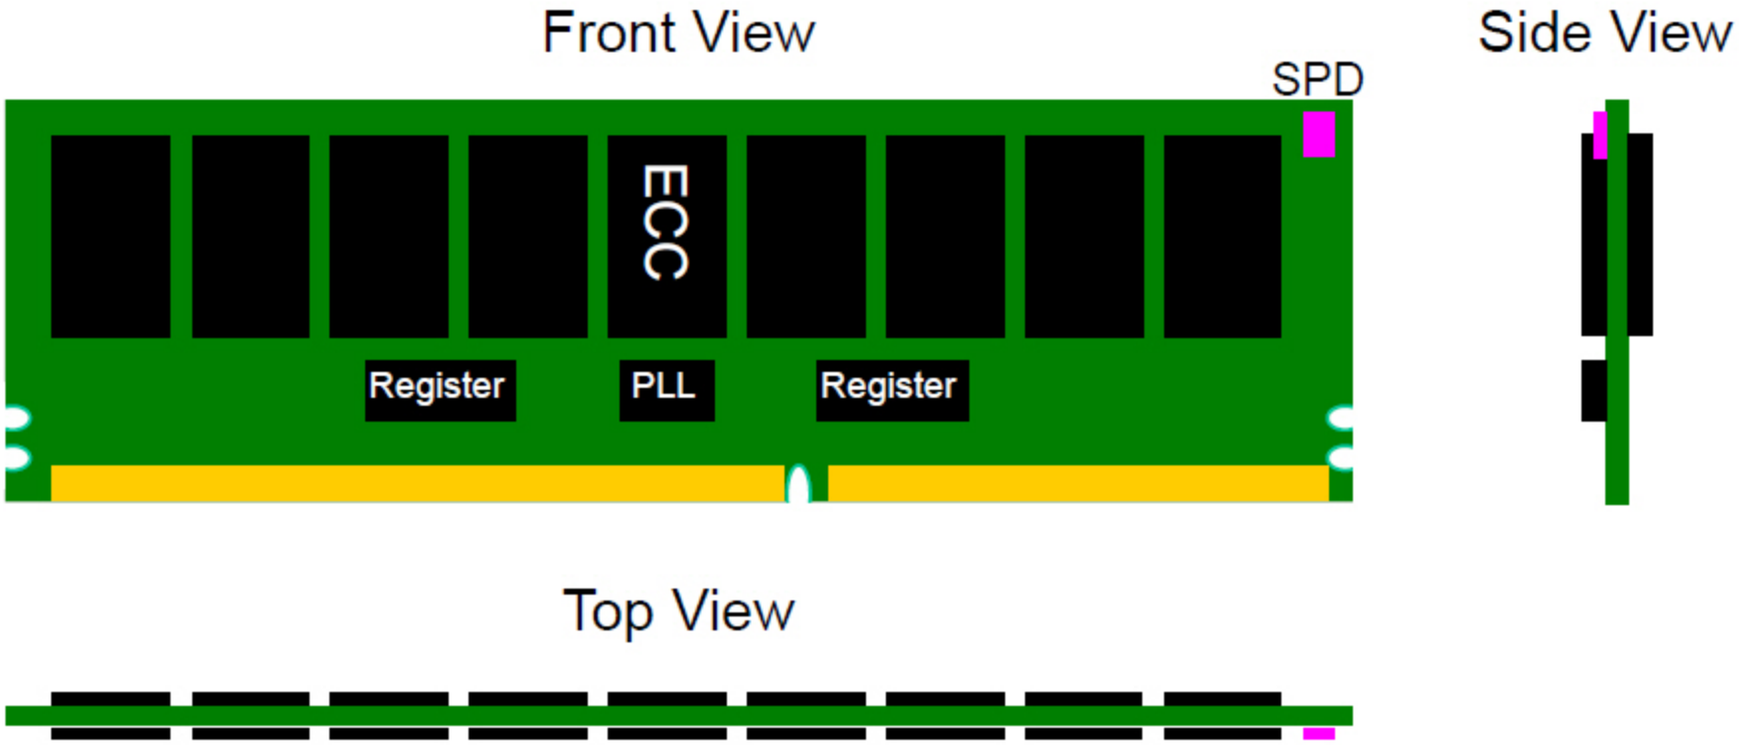

Module Types

UDIMM: Unbuffered Desktop standard

RDIMM: Registered Server standard

SODIMM: Notebook standard

FBDIMM: Fully Buffered Server

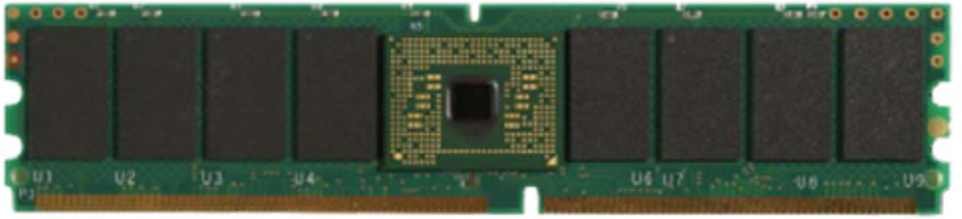

VLP RDIMM: Very Low Profile Computing and Networking

MiniDIMM: Computing and Networking

VLP MiniDIMM: Computing and Networking

DDR2 Module Pin Description

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

DDR2 Module Pin Description

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.

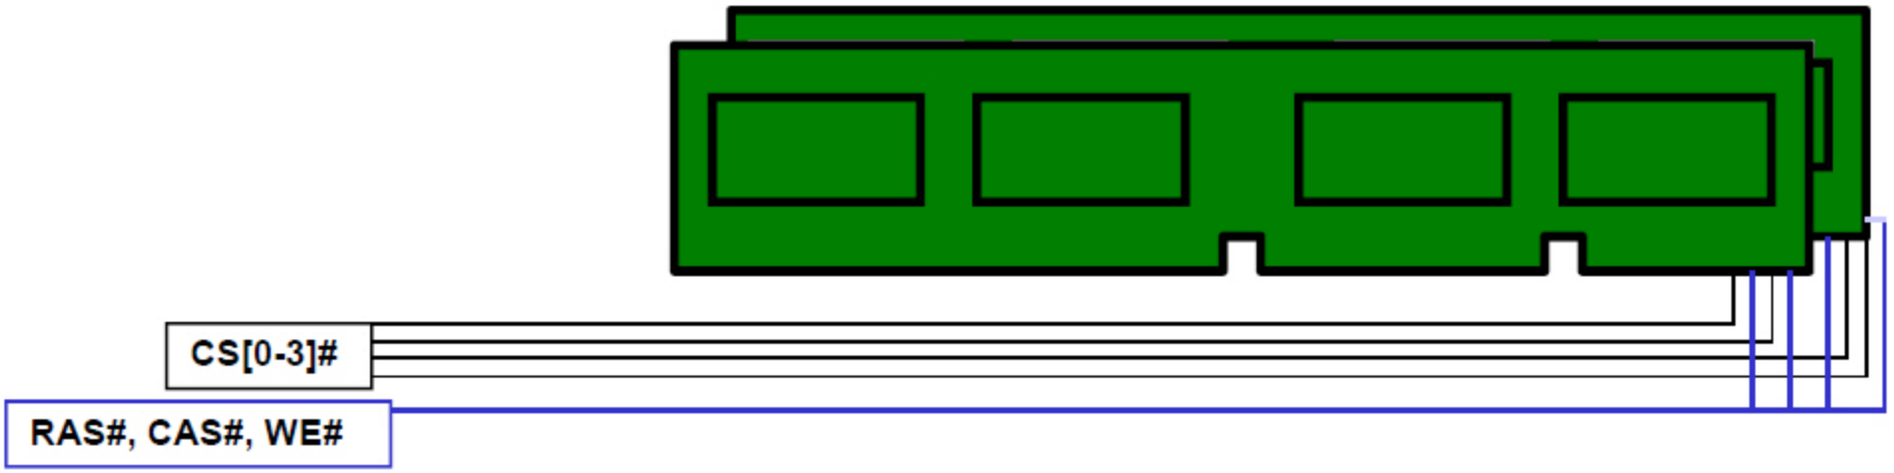

DDR2 Module Pin Description

Data Mask (DM 0-7)

- Input high
- Muxed with upper DQS (DQS9-17)
- Input mask signal for write data only used in x8 and x16 modes
- For x16 they are referred to UDM and LMD there is a second unconnected set used in x16 devices
- DM is tied low in x4 mode
- When DM is high write DQS is ignored when DM is low DQS is valid.
- DM is a bused signal so every rank is seen as a load on the bus.

DDR2 Module Pin Description

Bank Address (BA0-2)

- 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 (A0-15)

- Input High
- Define the Row Address for Active commands
- Define the Column Address and auto precharge bit for read/write accesses.
- A10 is normal Address info during a RAS
- A10 is sampled during a Precharge command.
  - A10 Low = Precharge one bank
  - A10 High = Precharge all banks
- A10 is sampled during a Read or Write command.
  - A10 Low = No Precharge
  - A10 High = Auto Precharge when command is complete.
- Define the A10 operation in Mode Register.

DDR2 Module Pin Description

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 signal on read data (DRAM driving)
- Center aligned with data signal on write data (Chipset driving)
- Differential example
  - x4 mode all 16 are used one per chip (DQS, DQS#)
  - x8 mode 8 are used one per chip (DQS, DQS#) Except in RDQS mode then all 16 are used. (DQS, DQS#) (RDQS, RDQS#)
  - x16 mode 8 (per rank) are used 2 per chip. (LDQS, LDQS#) (UDQS, UDQS#)
- DQS is a bused signal so every rank is seen as a load on the bus.

DDR2 Module Pin Description

CB 0-7

- Bi-directional data bus.
- Check bits used for ECC.
- The upper DQS is used for CB 0-7 either 9 or 9 and 17 depending on if it is x8 or x4 respectively.

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.

SCL

- Serial Presence Detect (SPD) Clock Input
- Serial clock used to synchronize the SPD logic.

SDA

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

SA 0-2

- SPD Address pins. Must be hardwired to different addresses for each module.

DDR3 Pin Additions/Changes

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 is lost and DRAM must be re-initialized, which includes (but not limited to) load mode registers and DLL reset.

A12/BC# - Active when low

- Active during Read and Write cycle commands.
- Used as A12 during CAS 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

Unbuffered DDR3 DIMMs only have 2 clocks (CK, CK# pairs) instead of 3.

MINDSHARE

BRINGING LIFE TO KNOWLEDGE

Commands and Waveforms

DDR3 Bank States

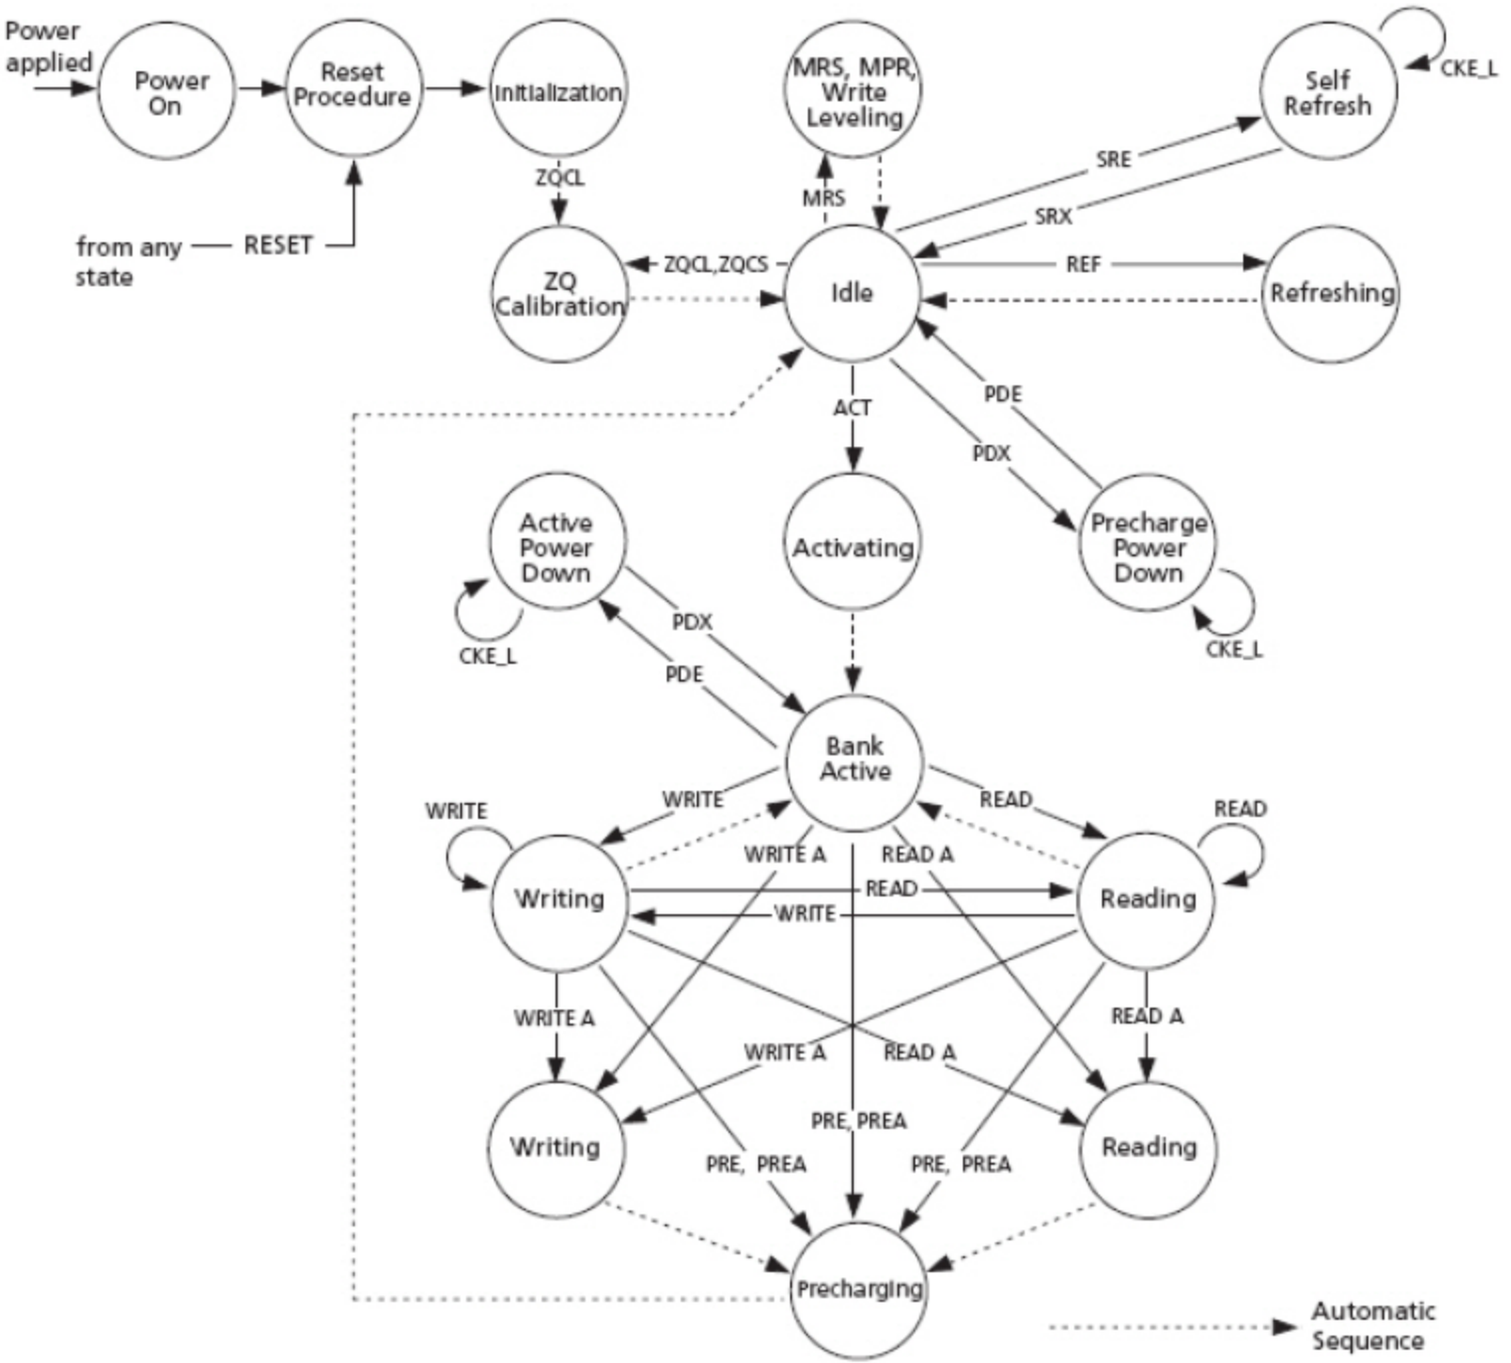

DDR2 Commands

Function	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	BA	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	H	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	L	H	L	H	H	H	X	X	X	X

DDR3 Command Additions

Function	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	0	L	CA
Write BL= 8	H	H	L	H	L	L	BA	CA	1	L	CA
Write with Auto Precharge BL=4	H	H	L	H	L	L	BA	CA	0	H	CA
Write with Auto Precharge BL=8	H	H	L	H	L	L	BA	CA	1	H	CA
Read BL= 4 (Burst Chop)	H	H	L	H	L	H	BA	CA	0	L	CA
Read BL= 8	H	H	L	H	L	H	BA	CA	1	L	CA
Read with Auto Precharge BL=4	H	H	L	H	L	H	BA	CA	0	H	CA
Read with Auto Precharge BL=8	H	H	L	H	L	H	BA	CA	1	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

Initialization Commands

- 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

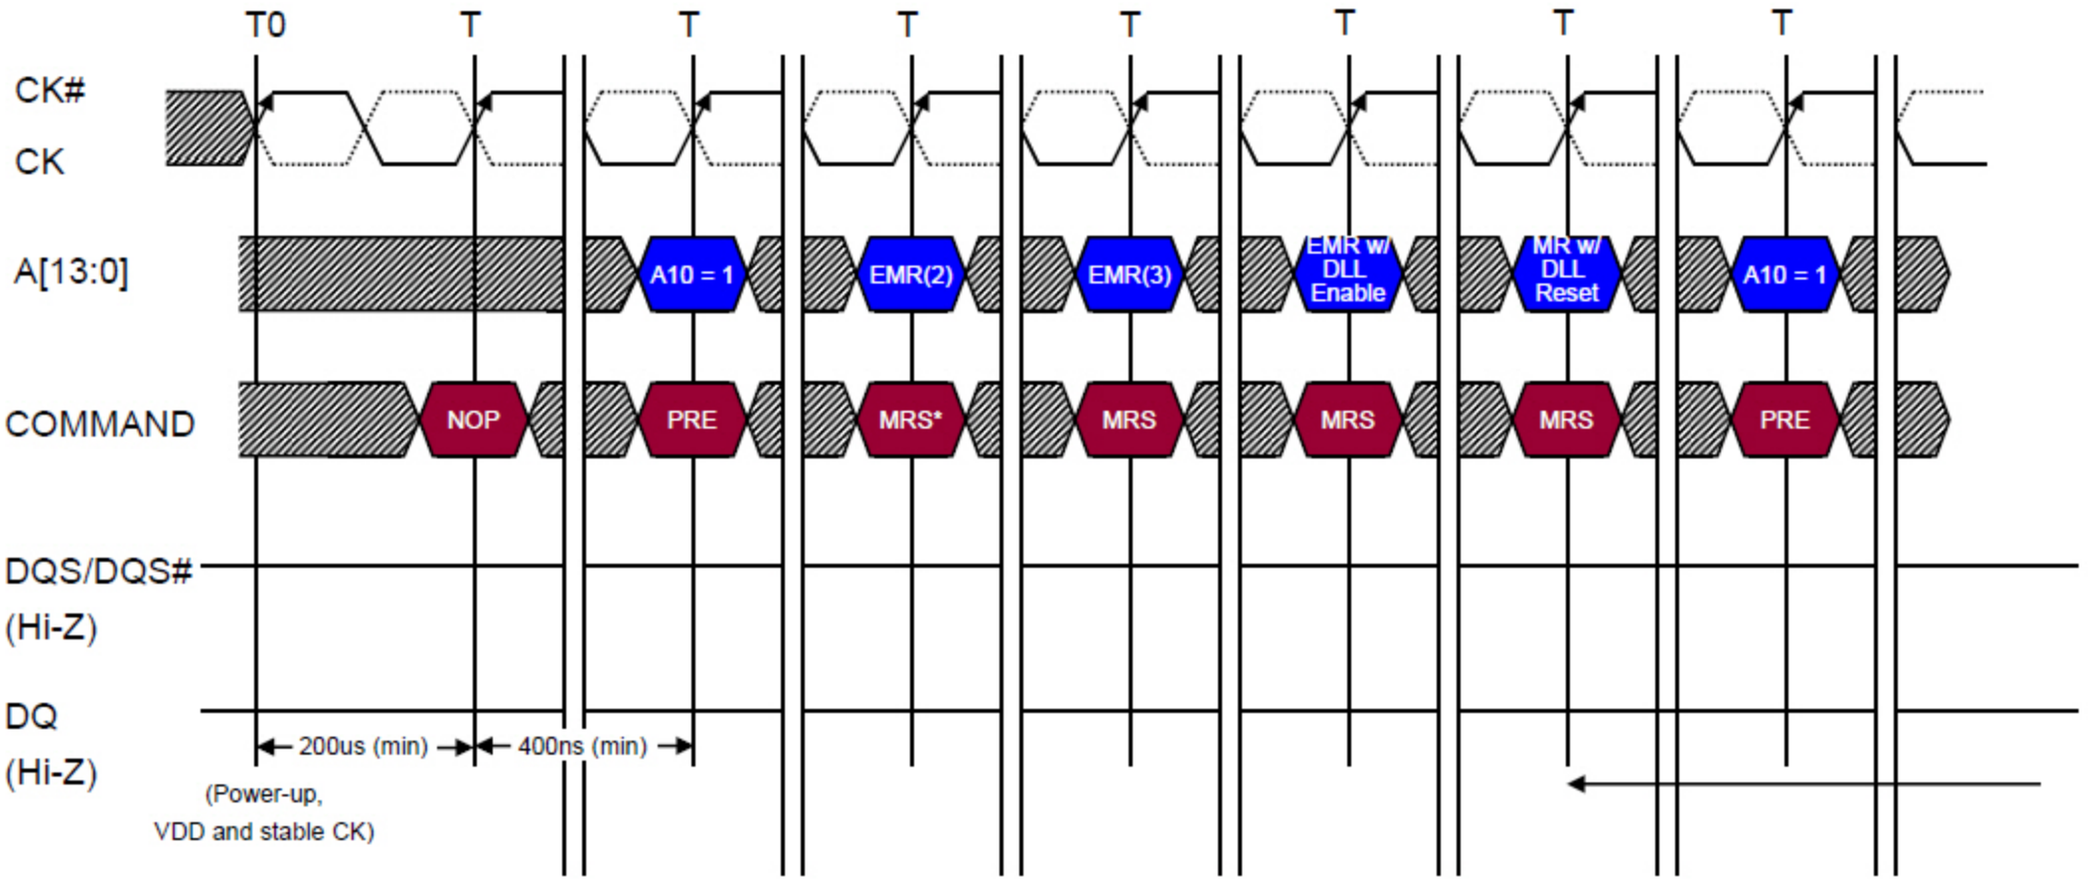

*MRS = Mode Register Set Command used to initialize SDRAM registers

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

MRS Waveforms cont.

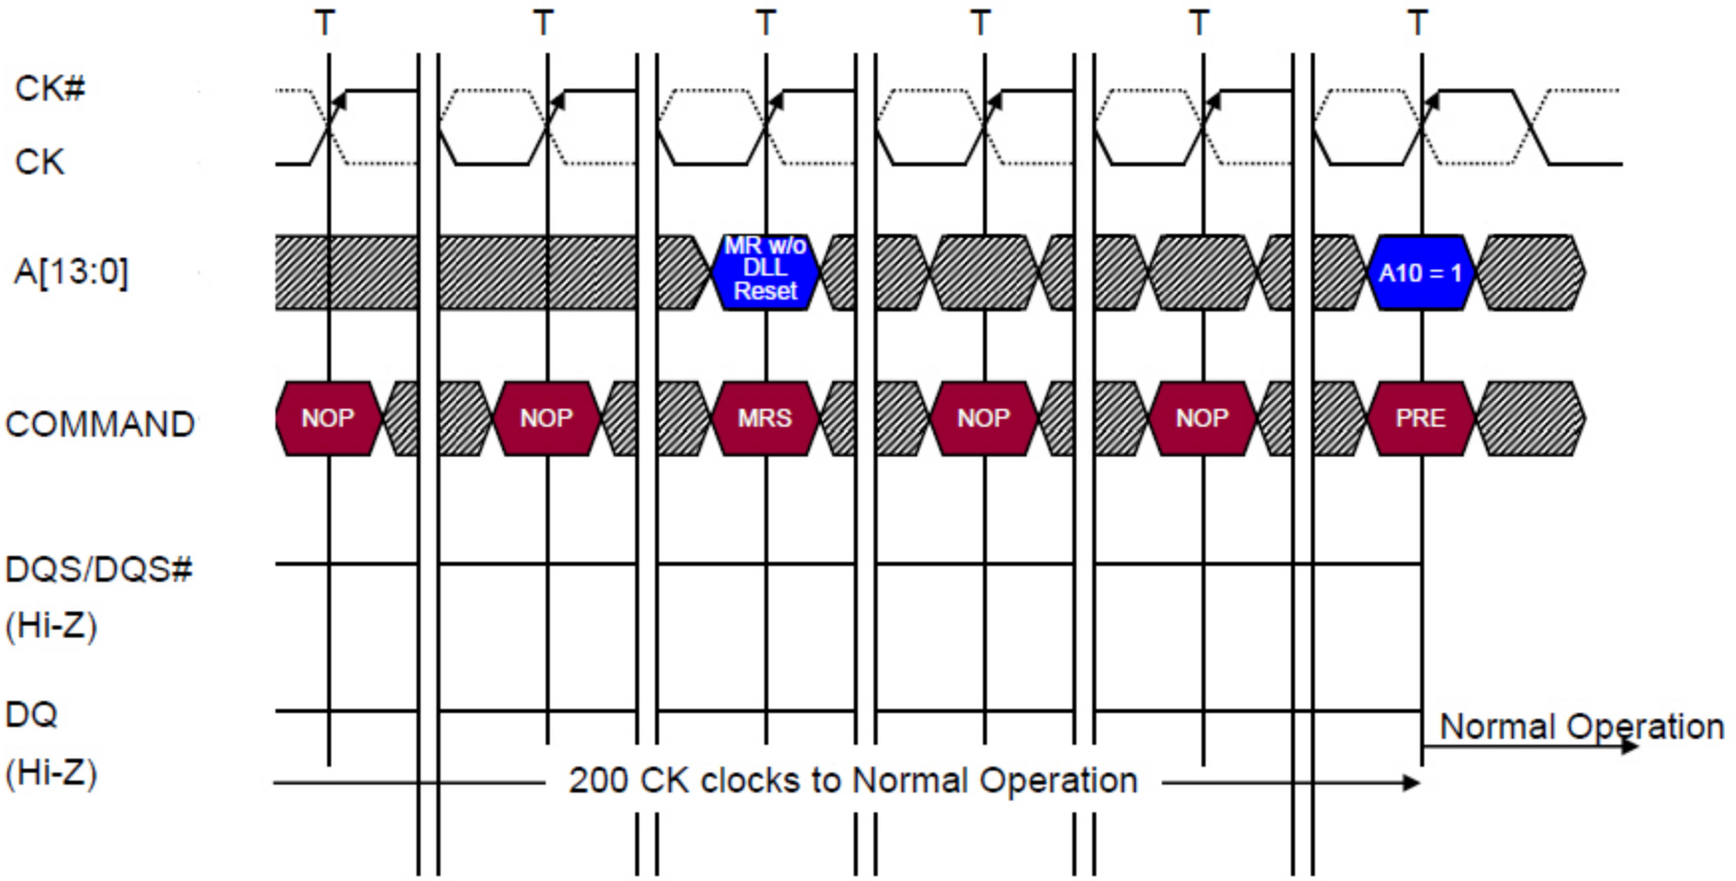

OCD = Off-Chip Driver Impedance Calibration.

Activate Command

- 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.

Typical Timings Shown

- 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

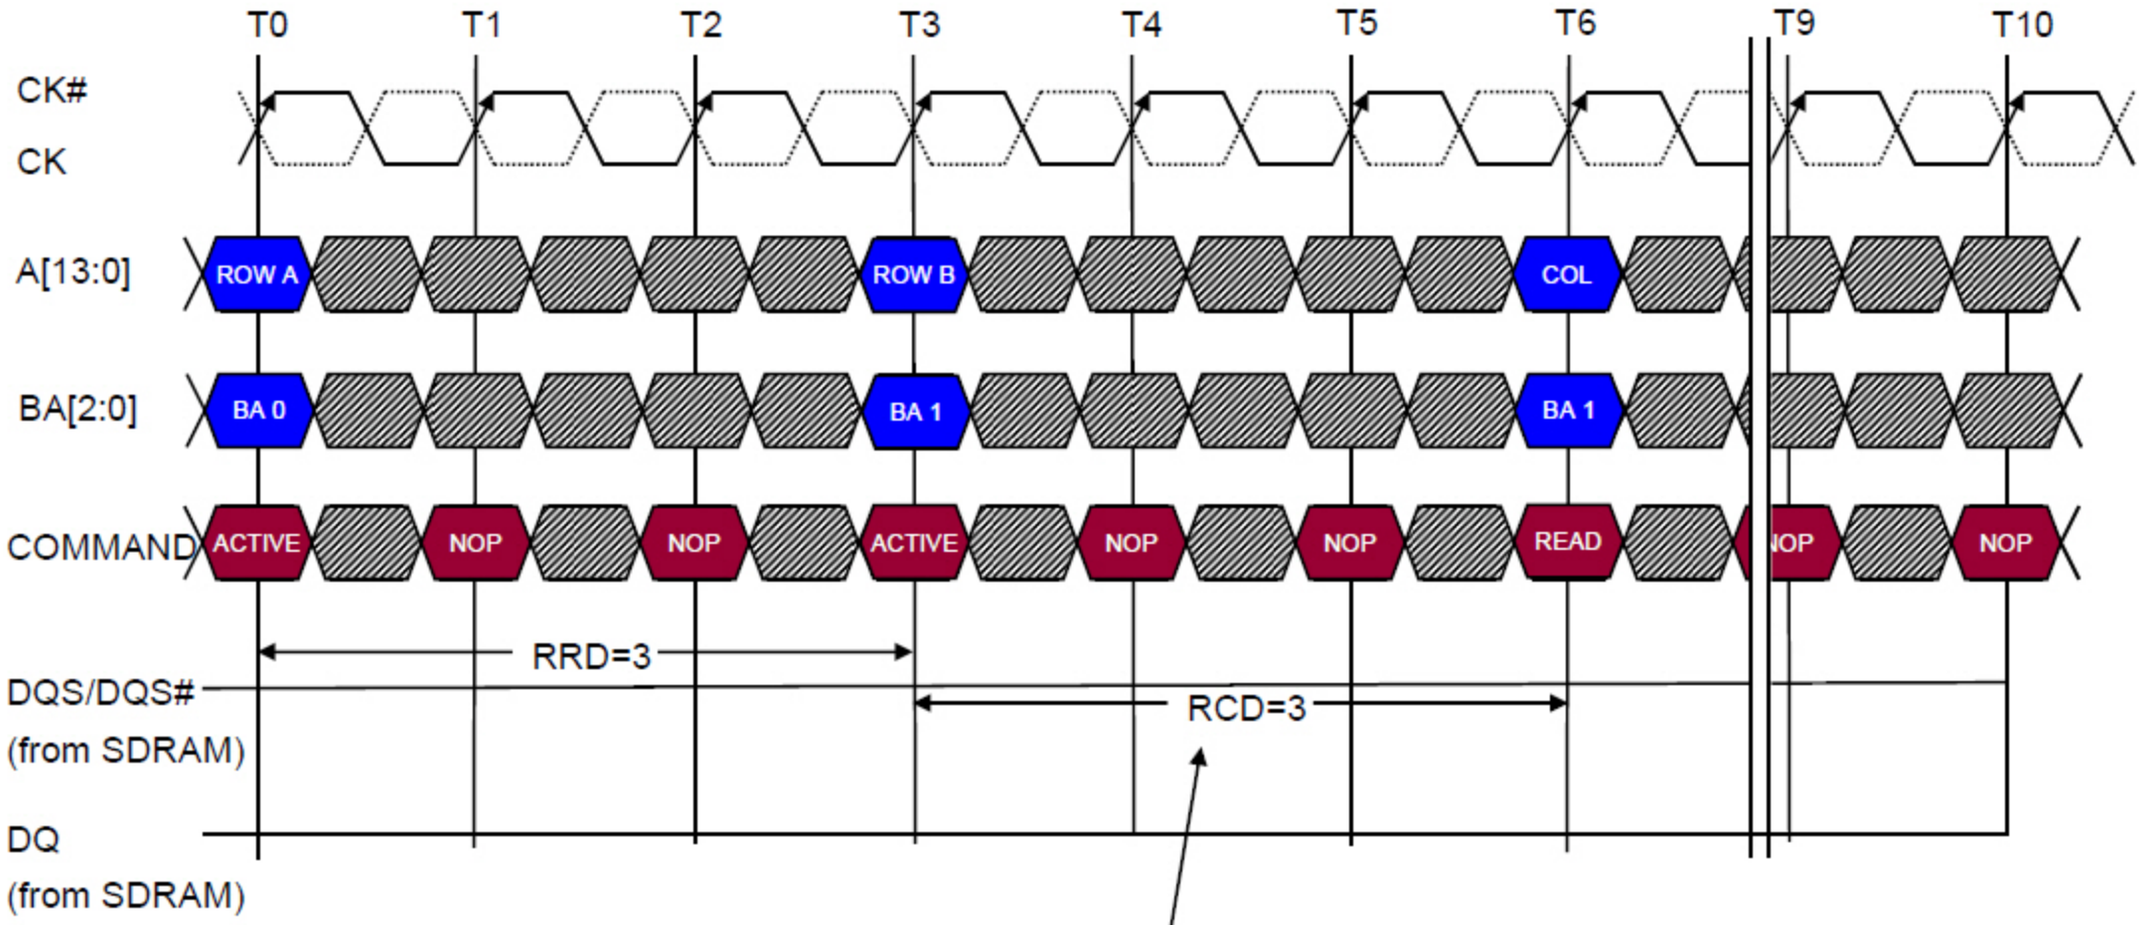

t_{RRD} is the Minimum time interval from one Bank Activate to another Bank (RAS to RAS delay)

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

DDR2 400 3 - 3 - 3

Read Command

- 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.

DDR1, DDR2, DDR3 Burst Orientation

- 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.

Burst Type and Length

- 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

V = Valid (stable) 0 or 1

Read Burst without Additive Latency

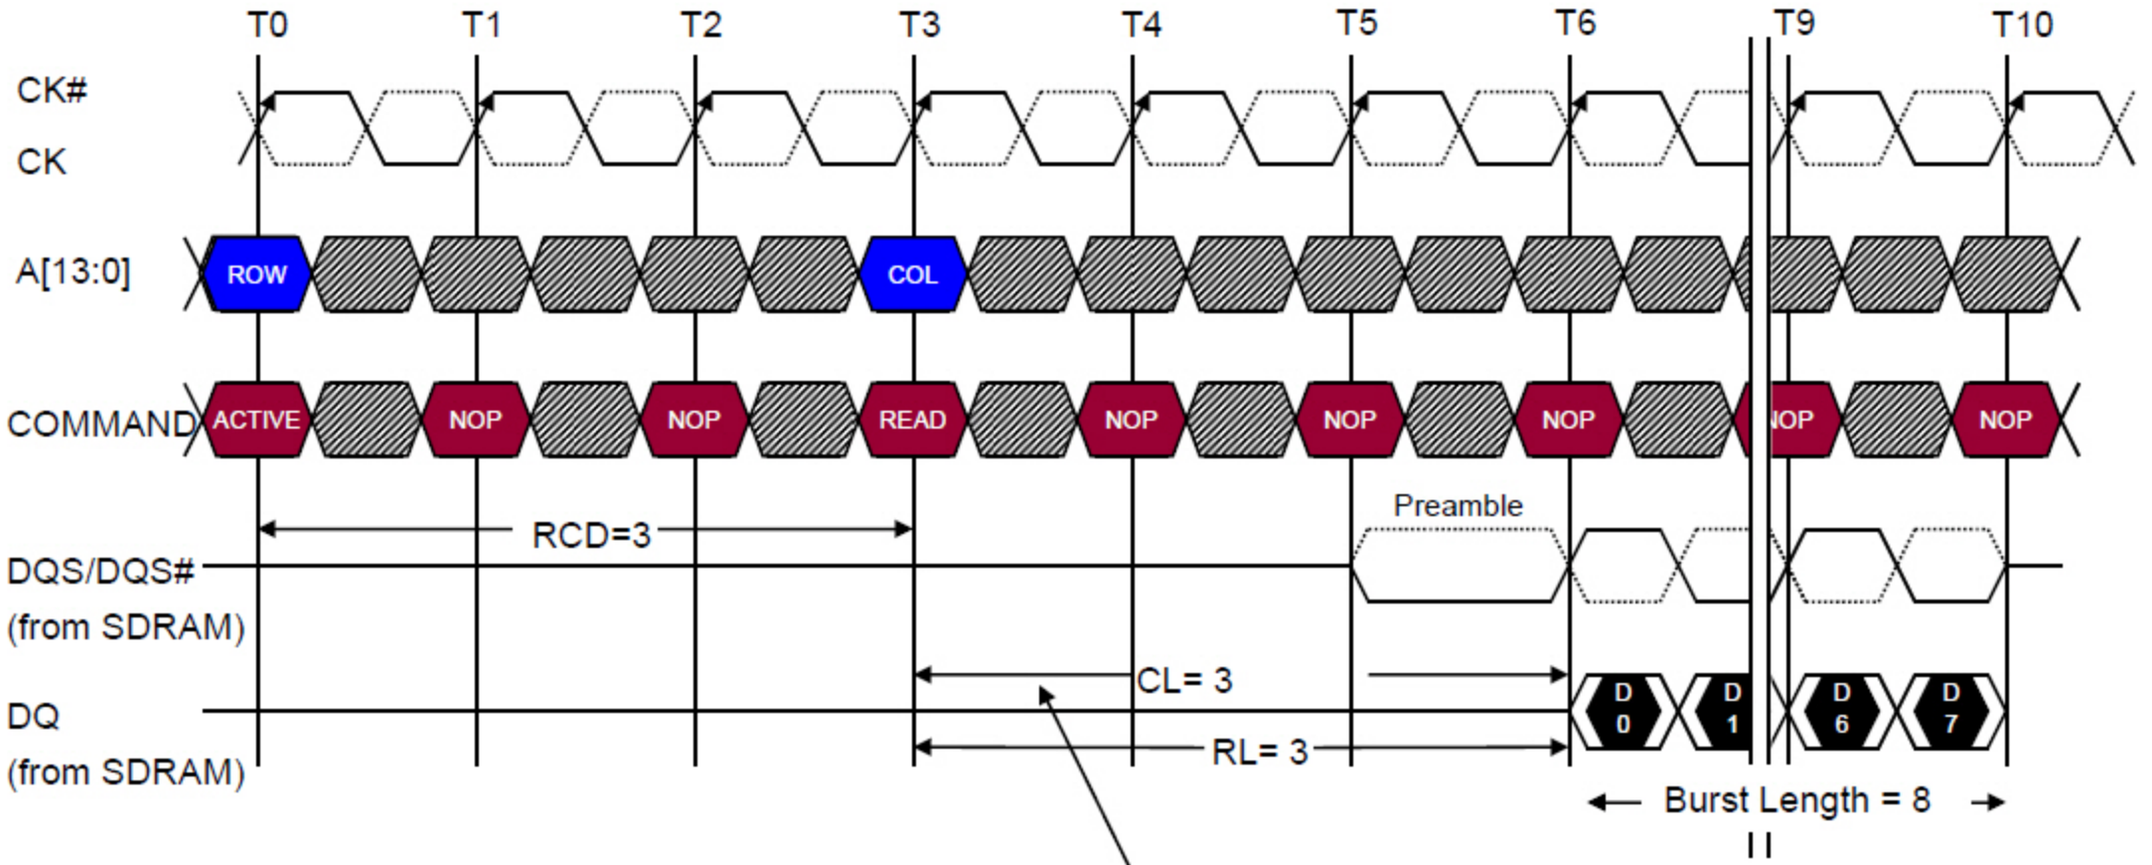

CL = CAS# Latency

AL = Additive Latency

RCD = RAS#-to-CAS# delay

RL = Read Latency

Burst Length = 8

DDR2 400 3 - 3 - 3

Read Burst with Additive Latency

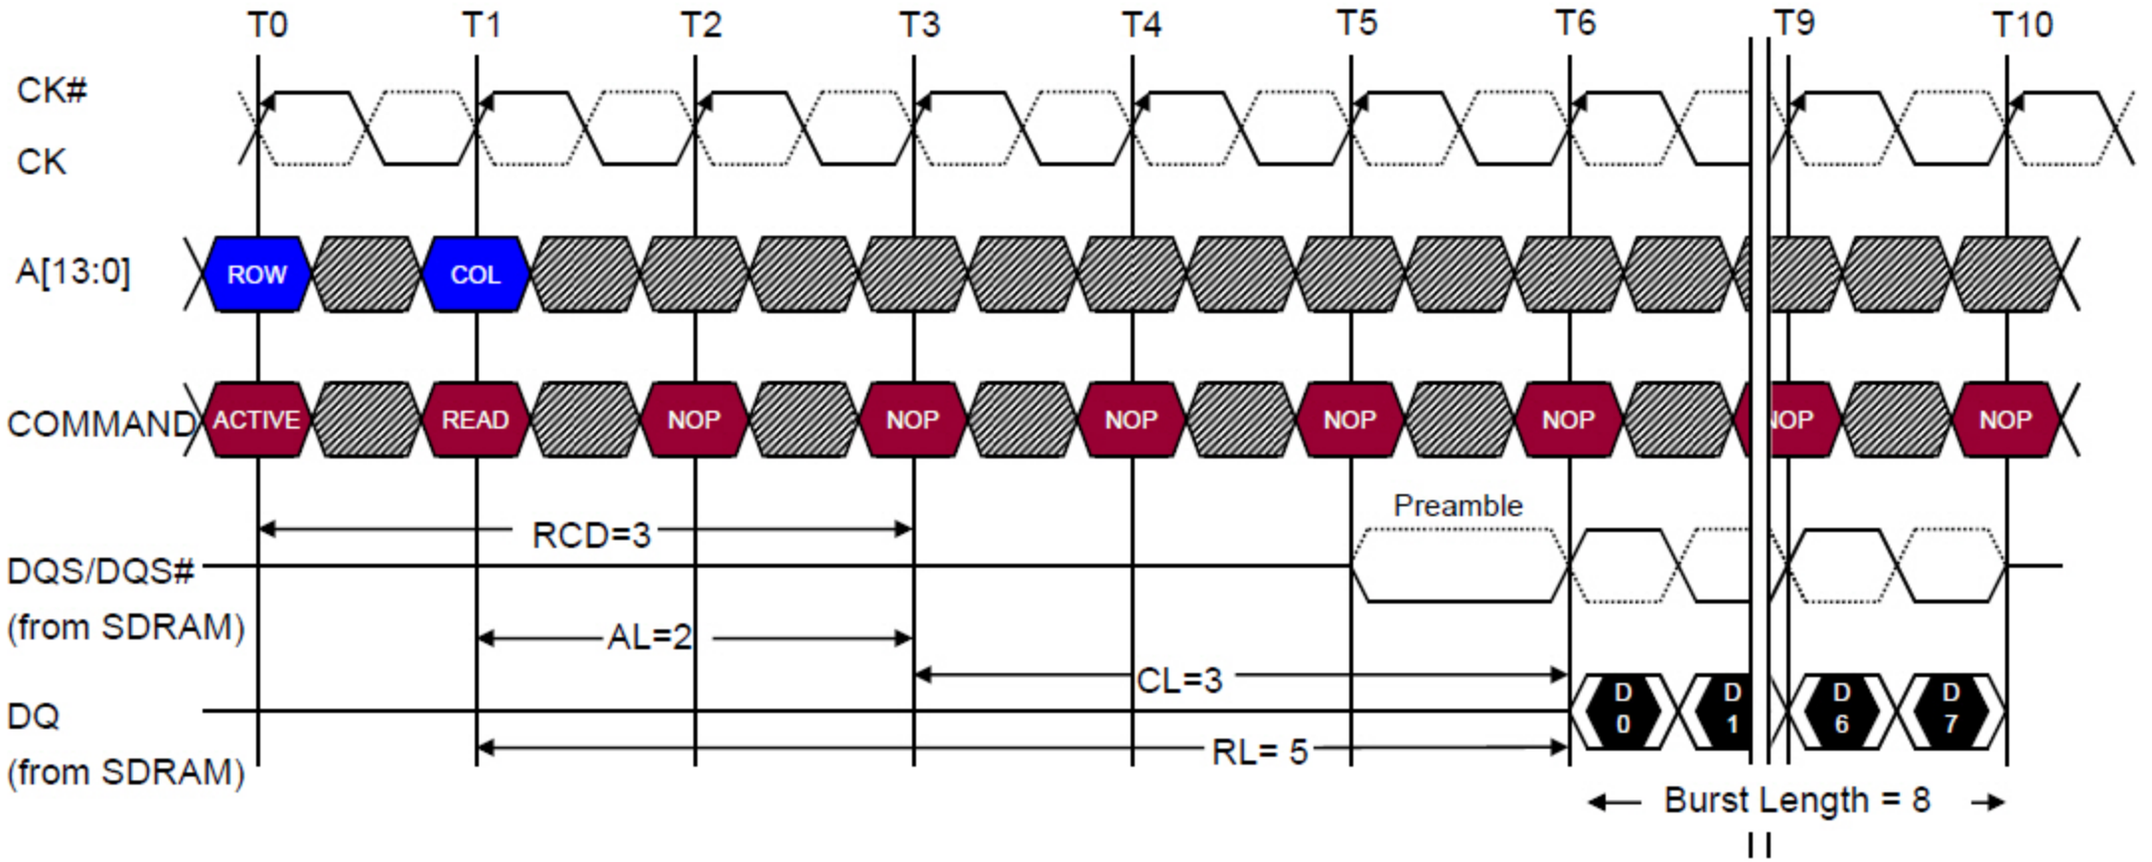

CL = CAS# Latency

AL = Additive Latency

RCD = RAS#-to-CAS# delay

RL = Read Latency

Burst Length = 8

Why Have Additive Latency?

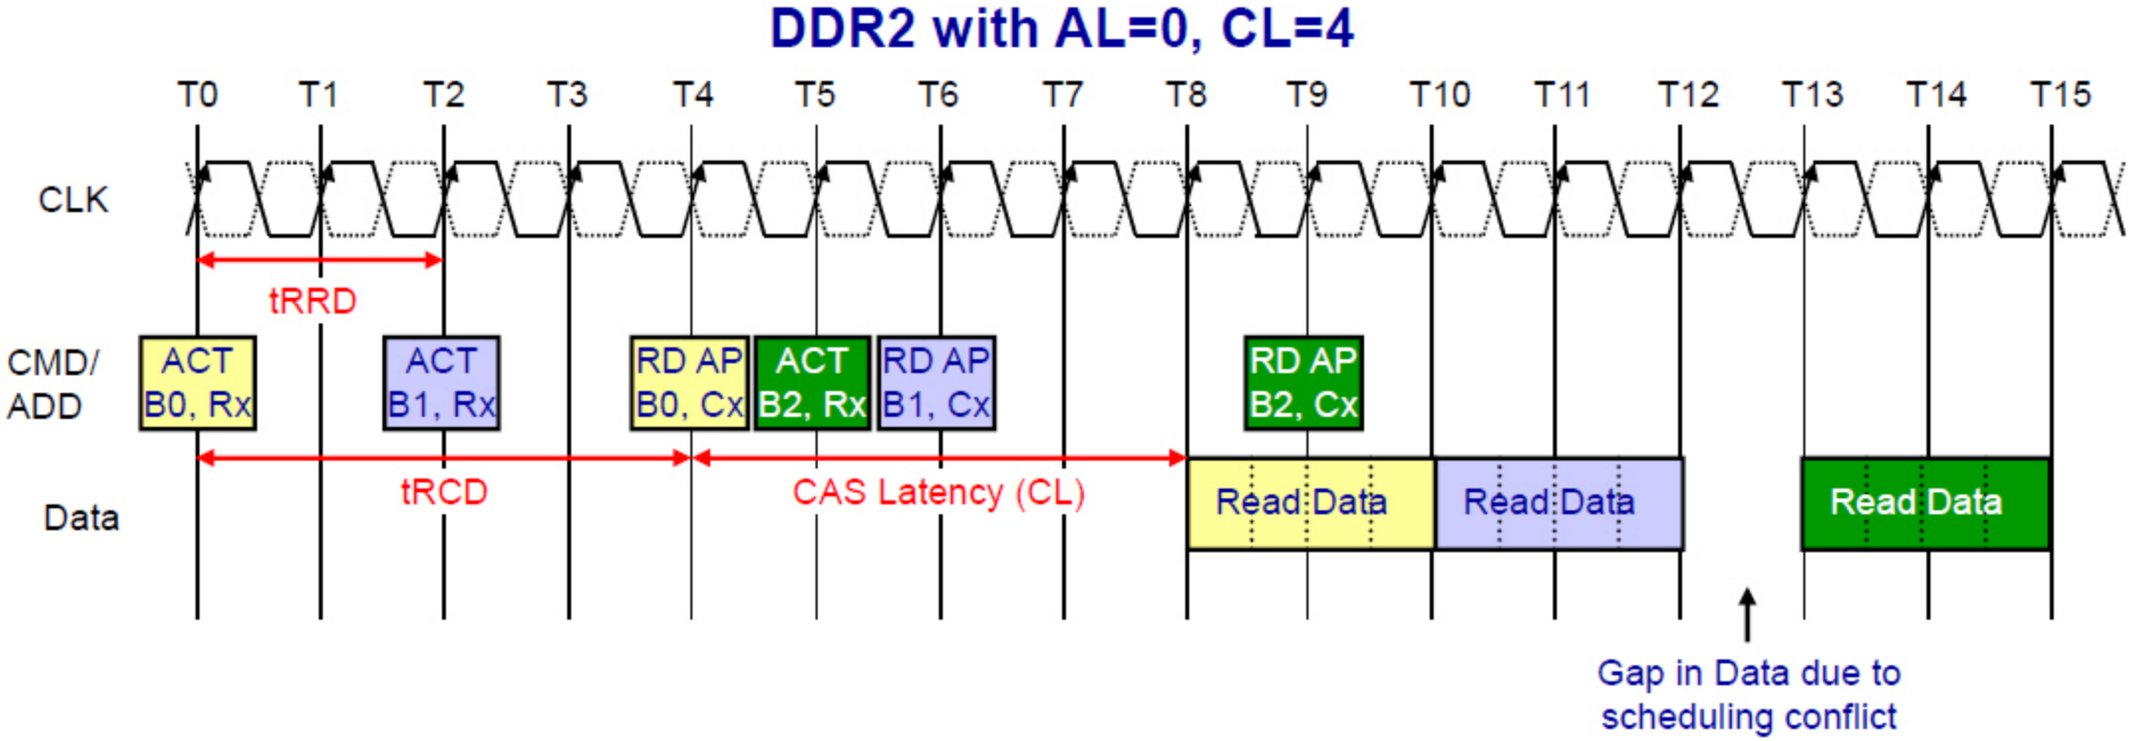

- The Read Command for Bank0 was sent during cycle T4.
- The memory controller would have liked to send the Activate Bank2 command on cycle T4 instead of T5.
- Due to this delay in activating Bank2, it results in a gap in the data stream being returned by the SDRAM because of this scheduling conflict.

Why Have Additive Latency?

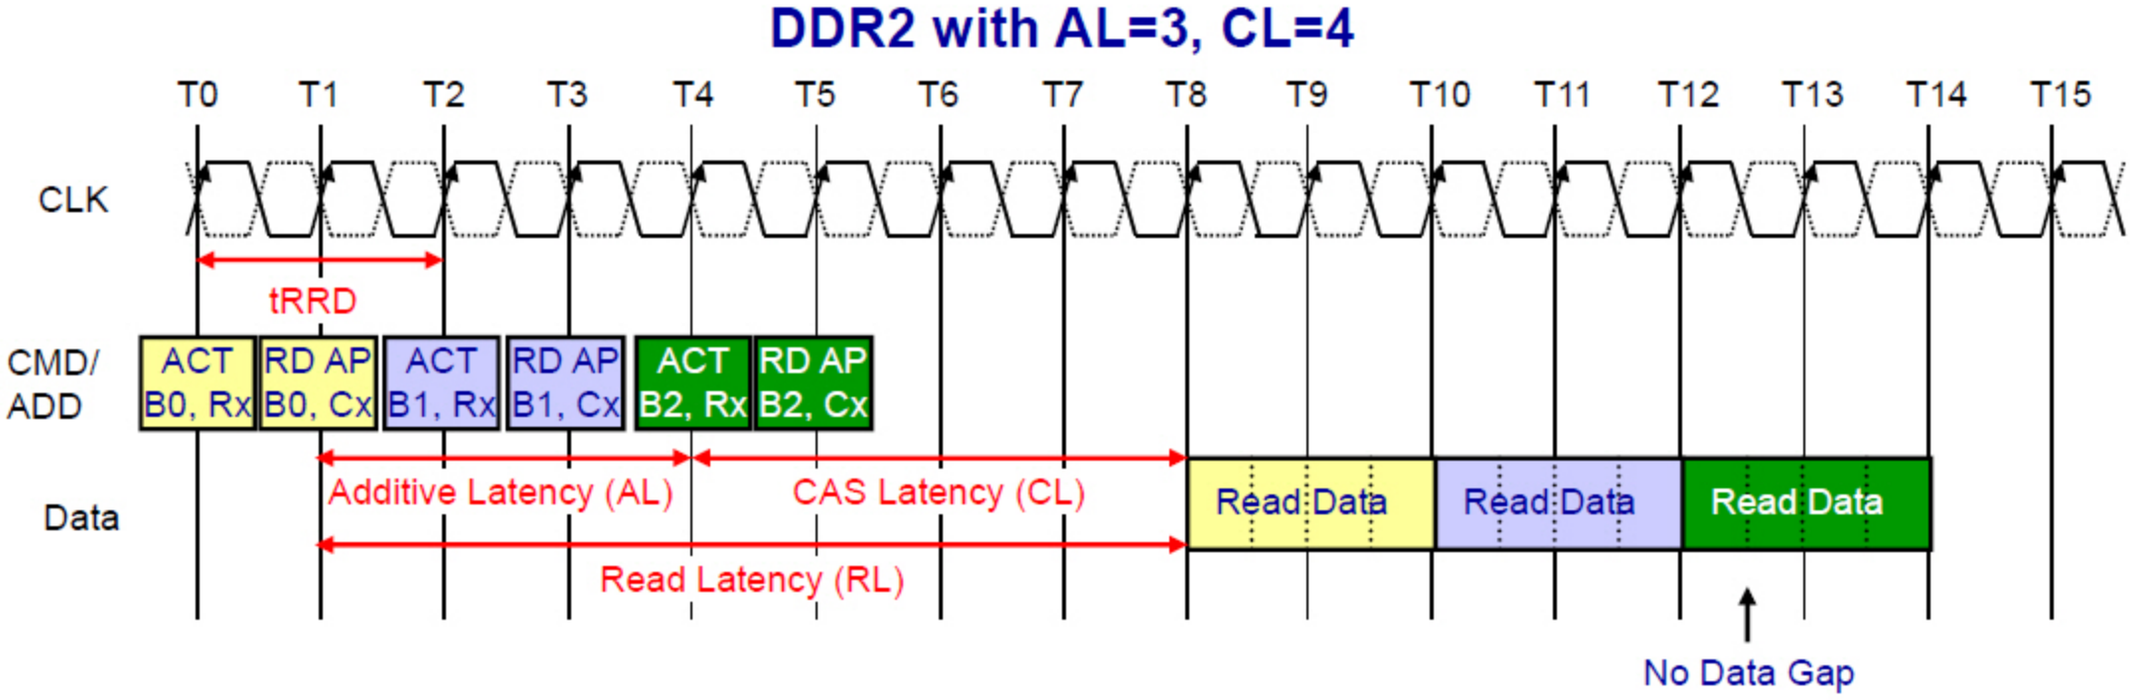

- DDR2 SDRAM can queue commands and schedule them at the appropriate time based on the programmed value of Additive Latency (AL).
- The memory controller no longer has a scheduling conflict because the commands can be sent to SDRAM back-to-back and the SDRAM will schedule them at the appropriate time.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

Read Burst Consecutive

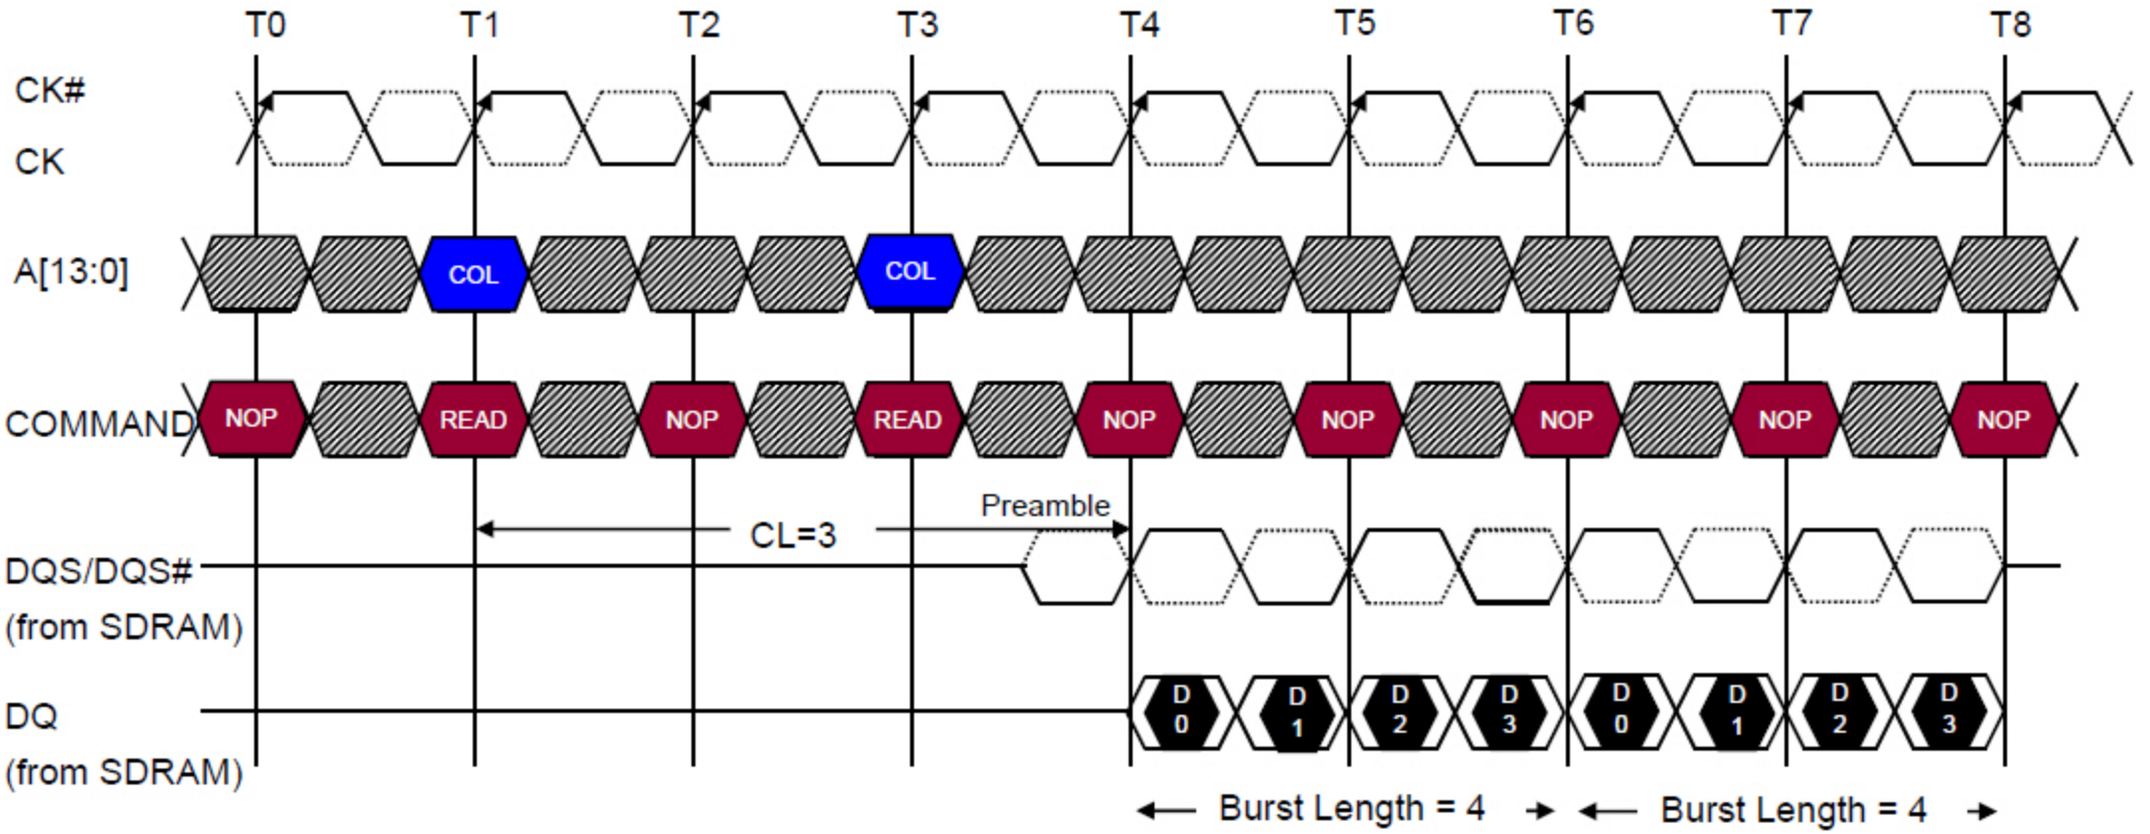

CL = CAS# Latency

AL = Additive Latency

RCD = RAS#-to-CAS# delay

Burst Length = 4

Write Command

- 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

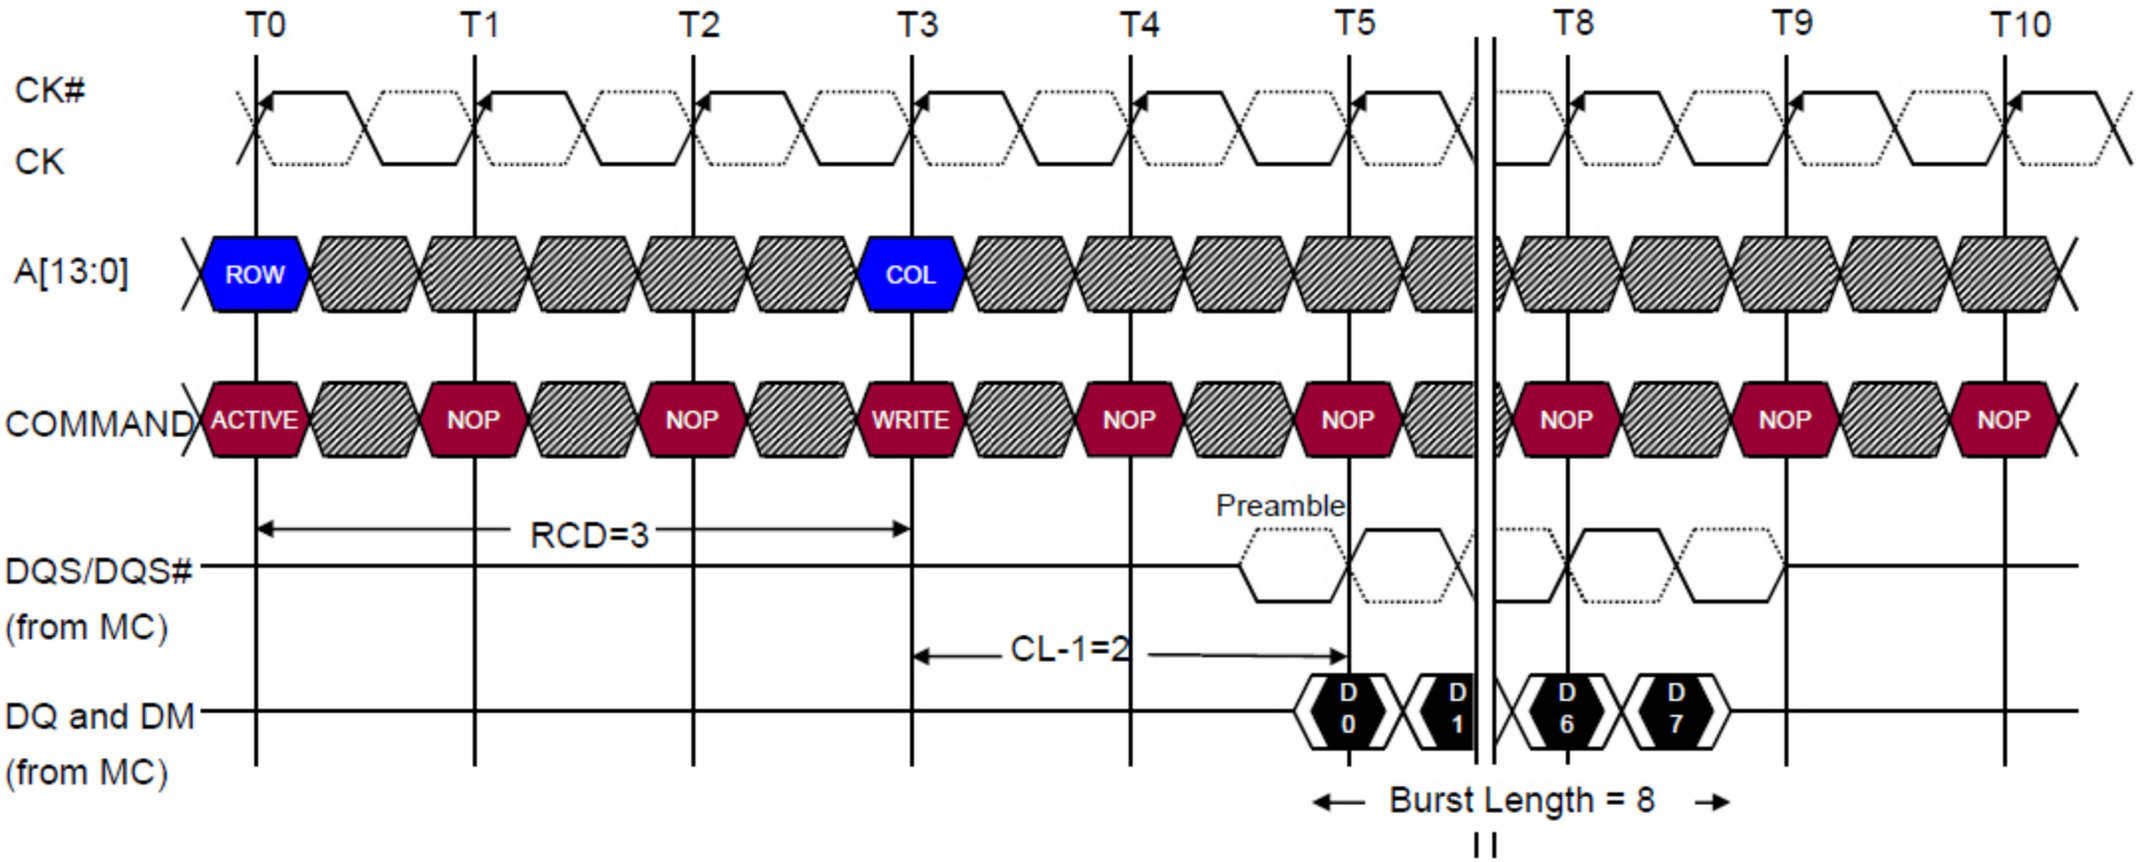

CL = CAS# Latency

AL = Additive Latency

RCD = RAS#-to-CAS# delay

RL = Read Latency

Burst Length = 8

Precharge Command

- 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.

Auto Precharge

- 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

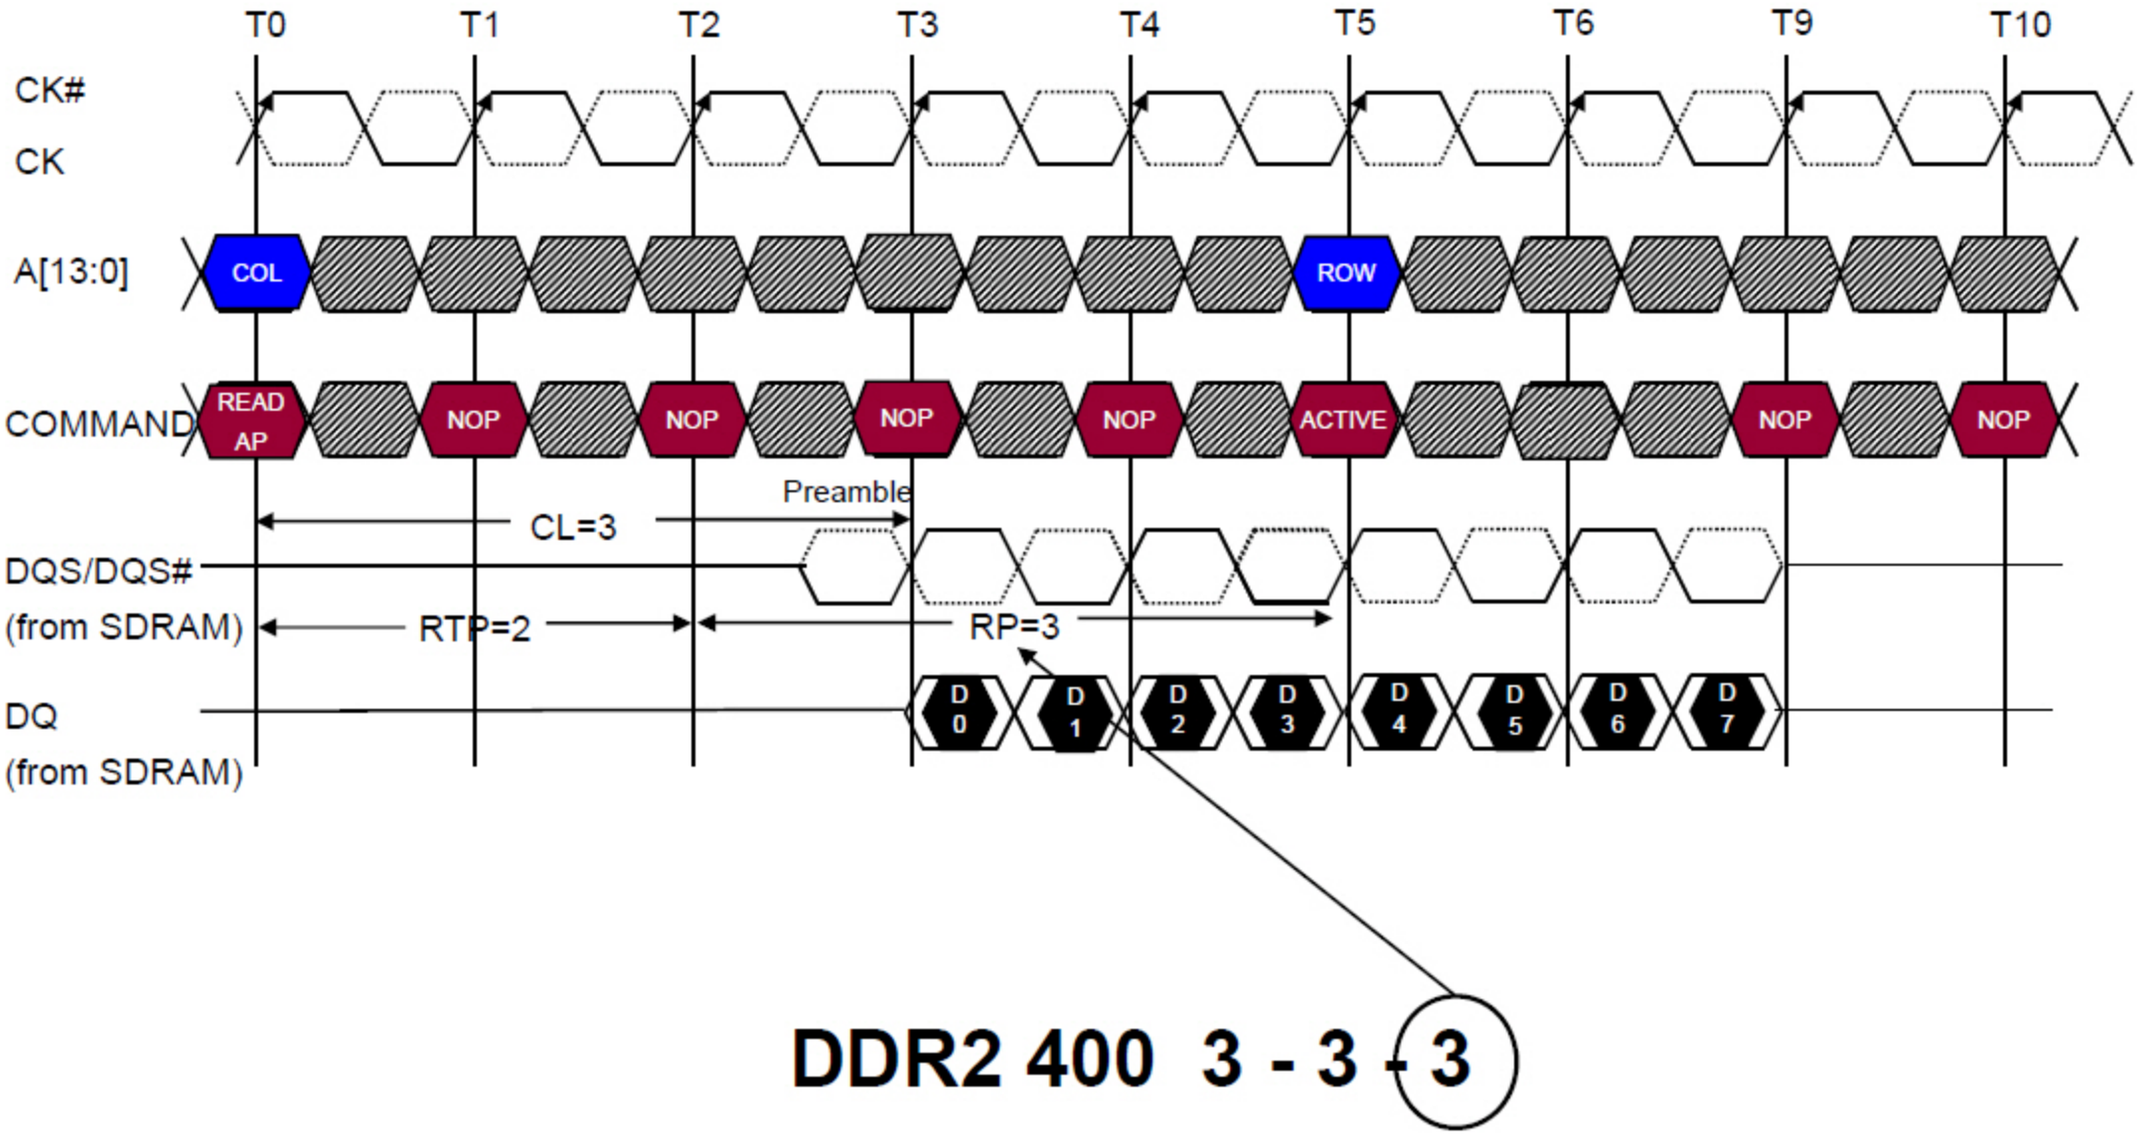

DDR2 400 3 - 3 - 3

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

t_{RTP} is the time from Read to Precharge

Refresh Commands

- 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 t_{REFC} (Refresh Cycle Time) X 8 Auto Refresh to next Refresh is t_{REFC} min.

Refresh Commands

- 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)

Longest DDR3 t_{RFC} is 350ns

Refresh Waveforms

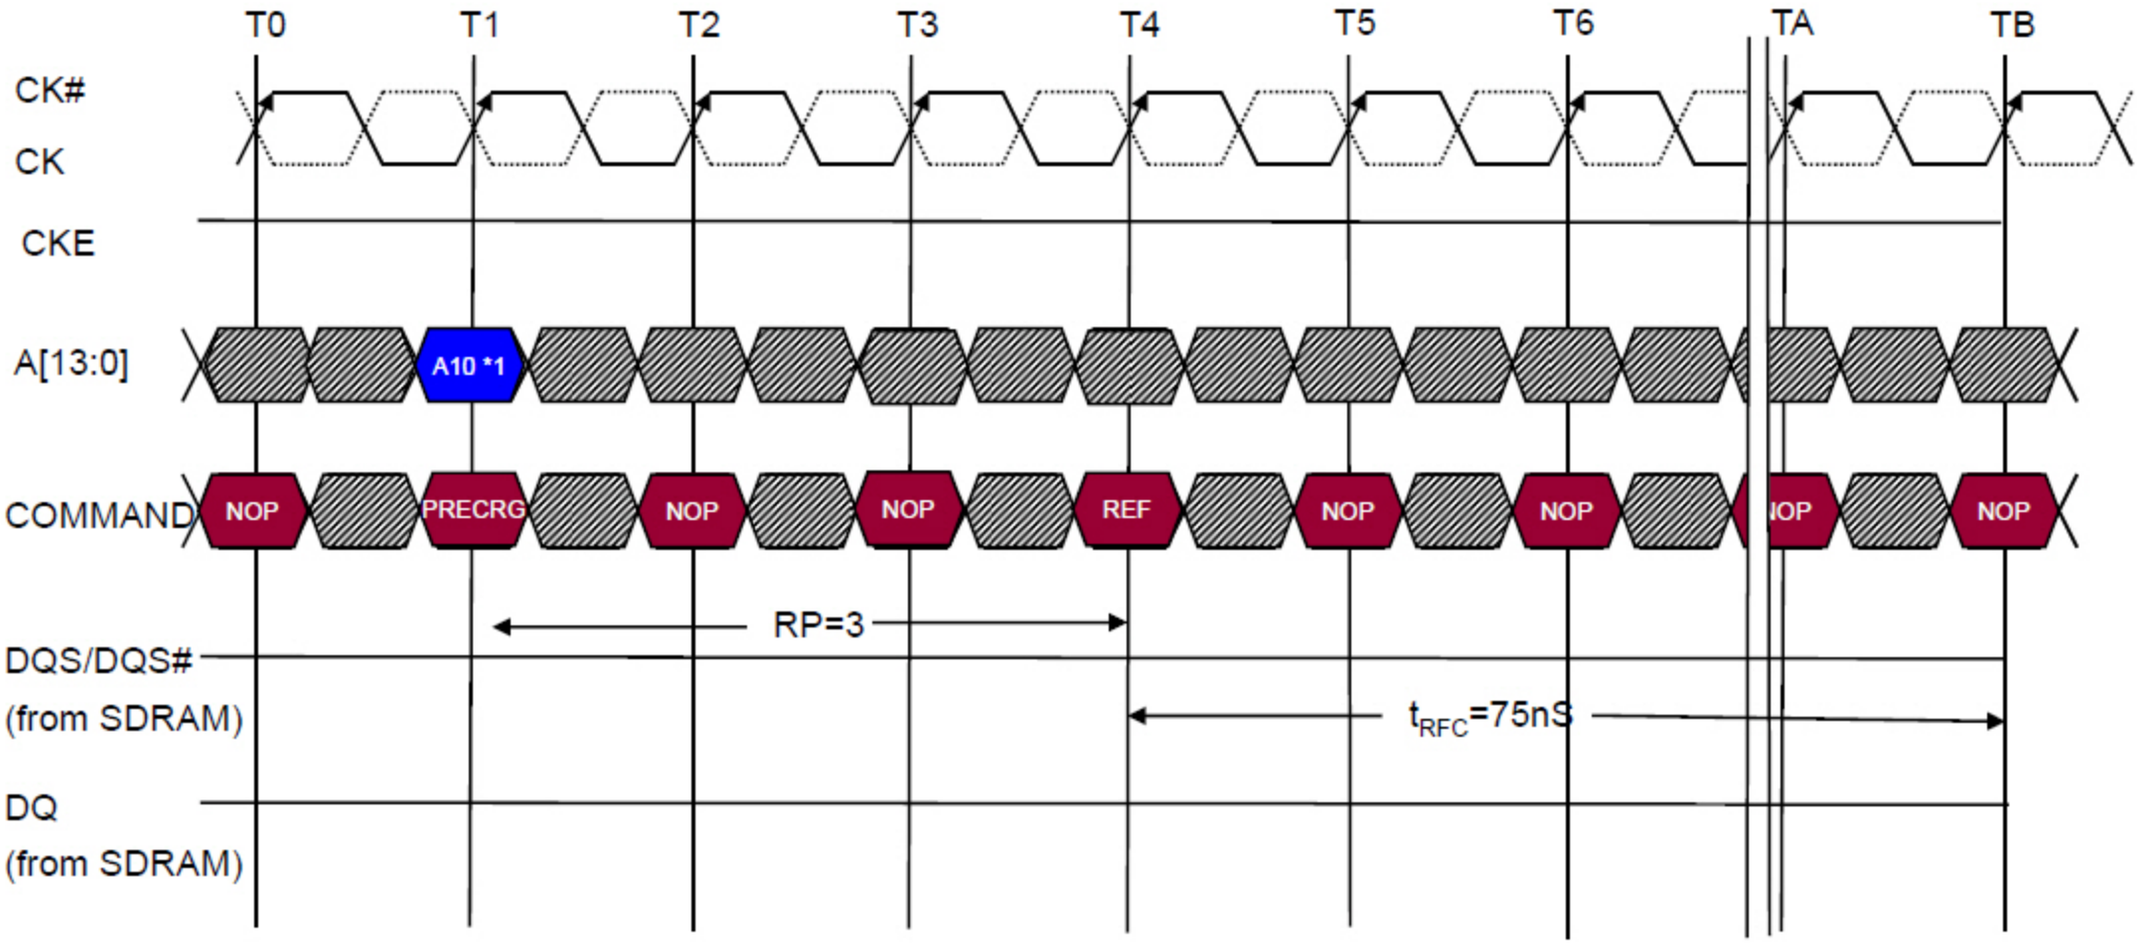

*1 A10 must be high for more than 1 bank to get Precharged. Precharge all must be done before entering the Refresh mode.

t_{RFC} is the refresh cycle time

Self Refresh

- 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

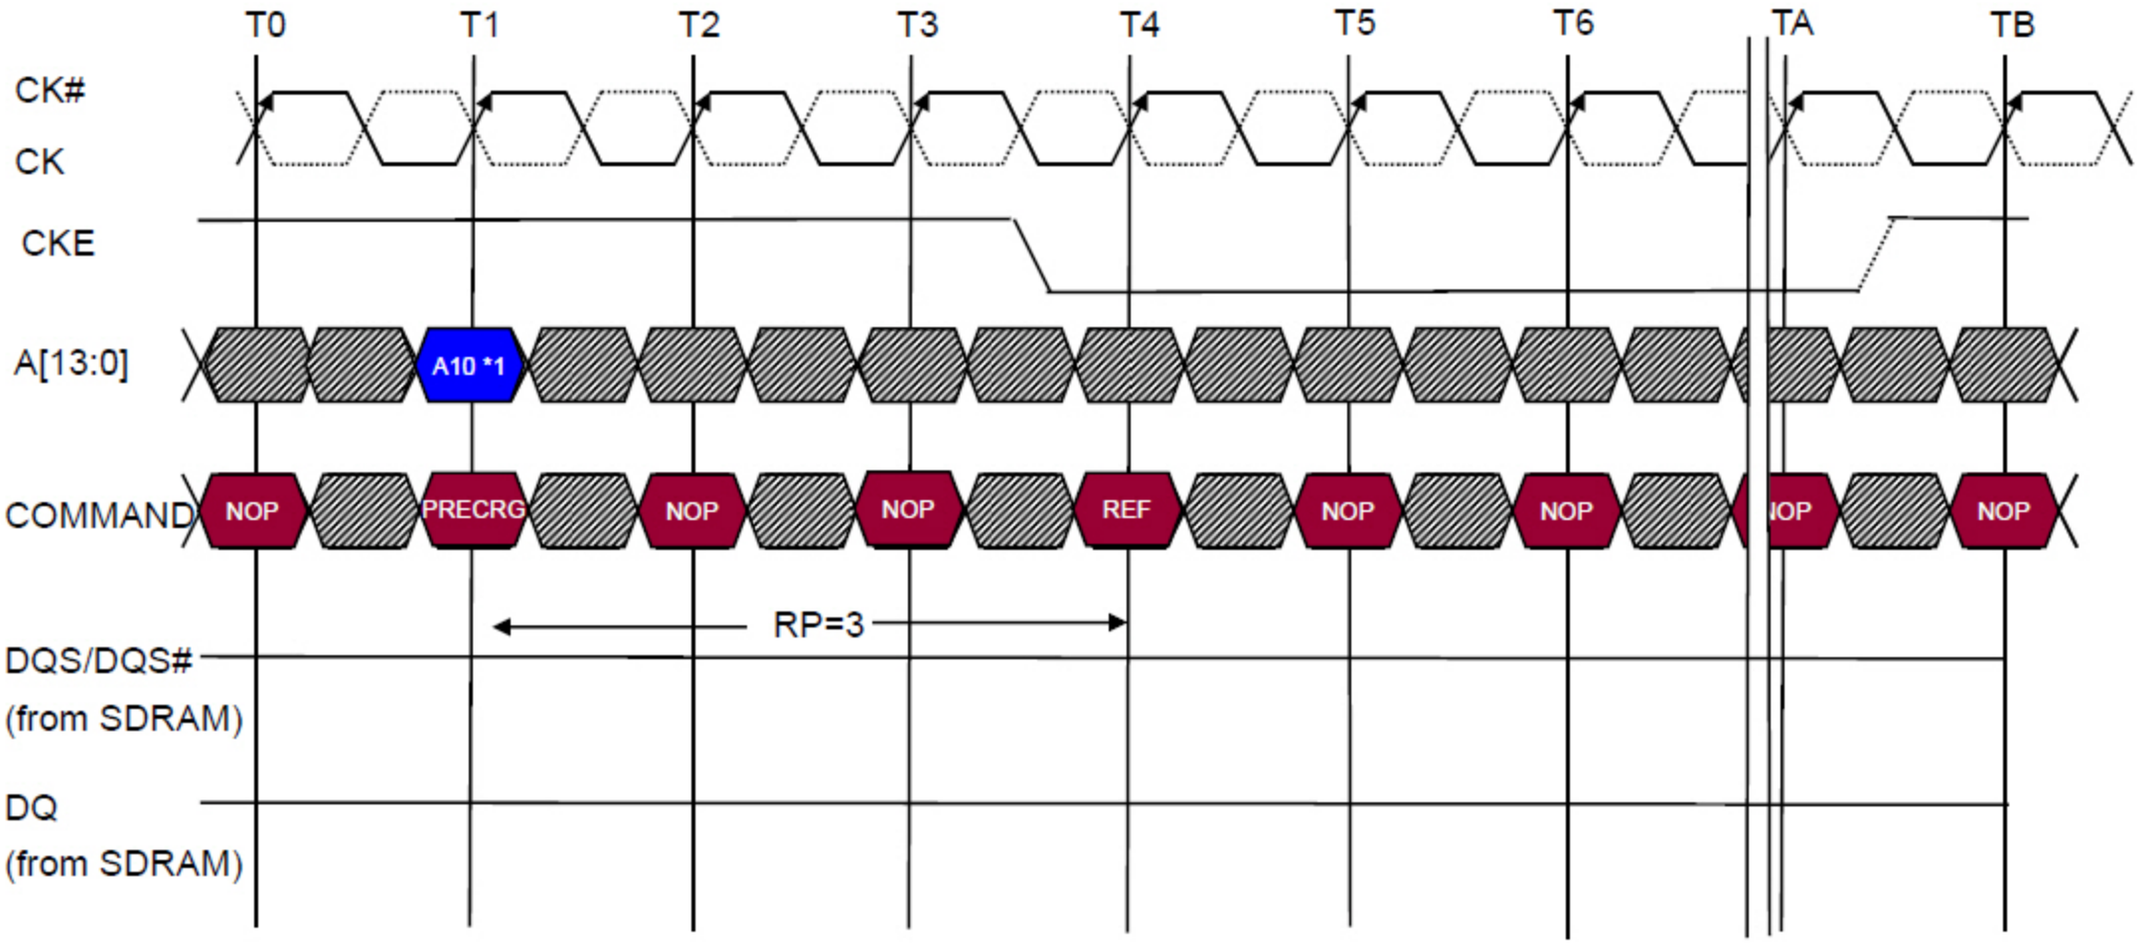

*1 A10 must be high for more than 1 bank to get Precharged. Precharge all must be done before entering the Self Refresh mode.

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

Refresh

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.

NOP and Deselect

- 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

- 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 X t_{REFI} of the device because no refreshes can happen in this state. The 8 X 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.

Power Down

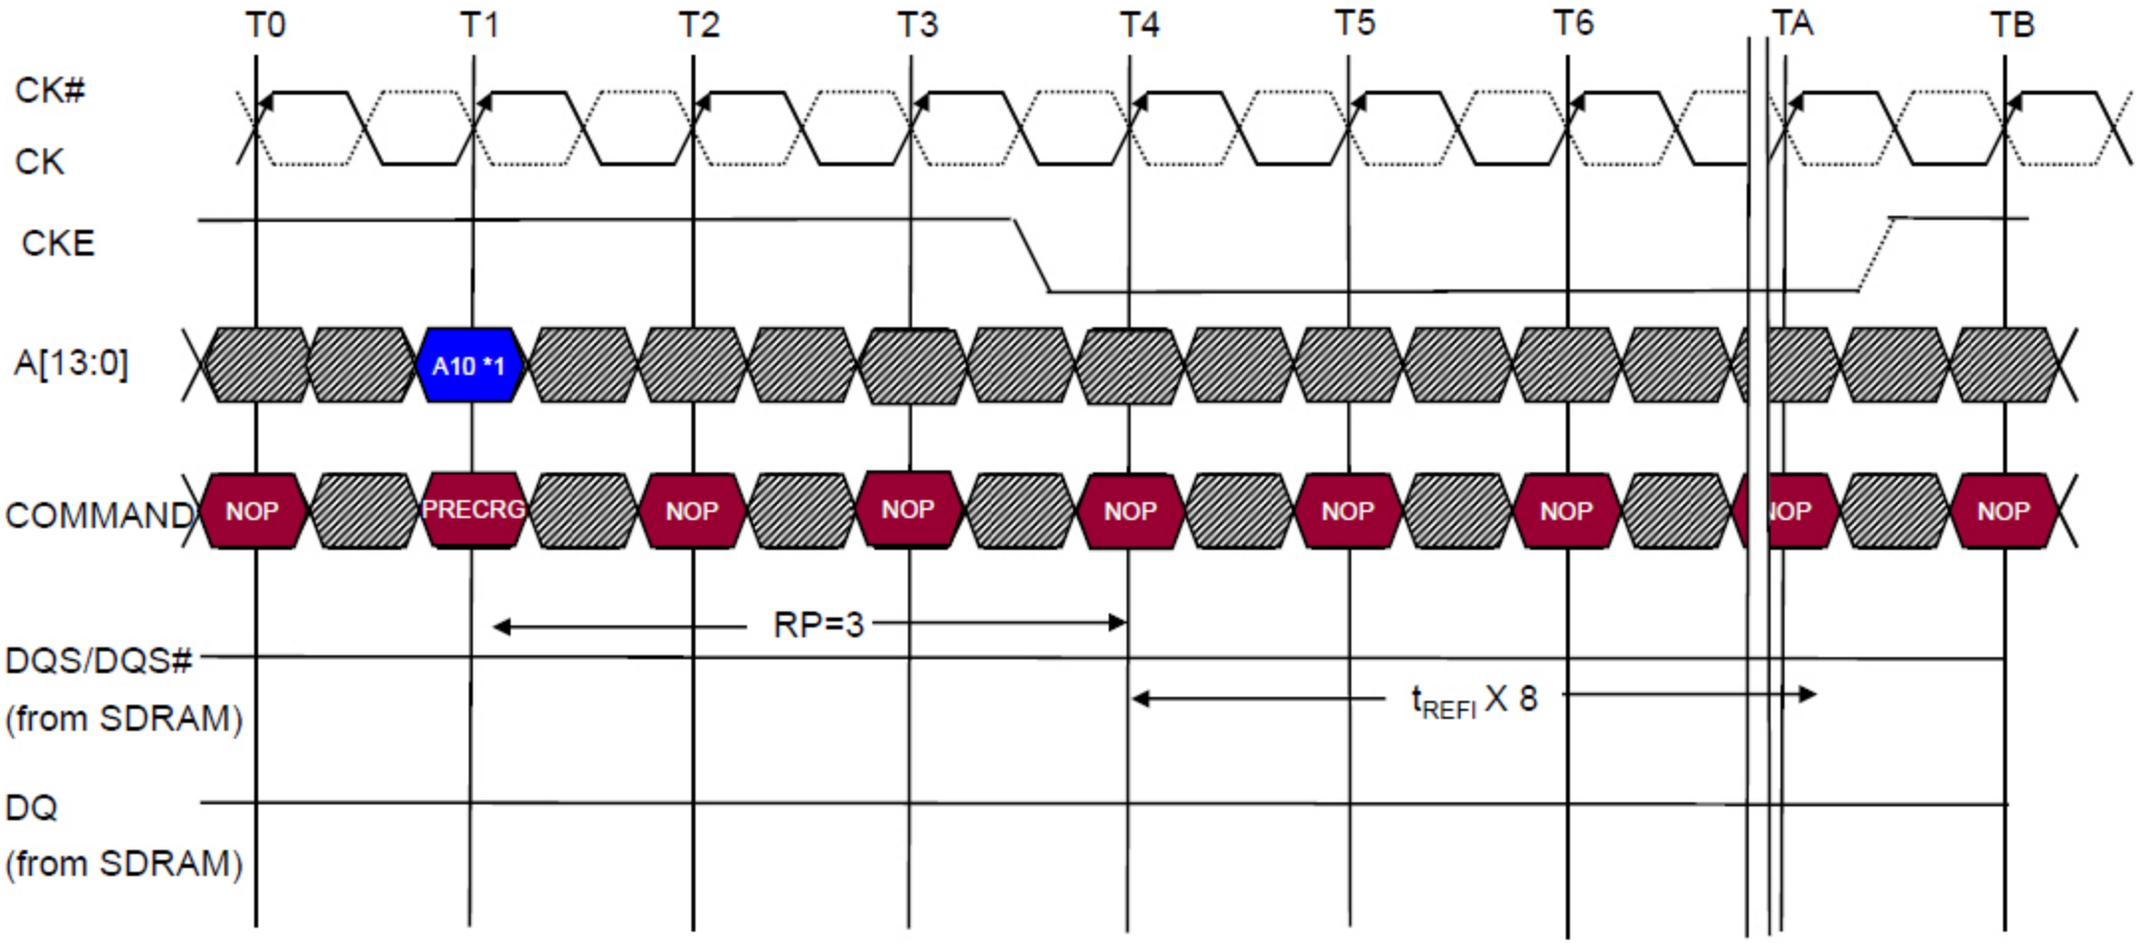

*1 A10 must be high for more than 1 bank to get Precharged.

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

Address and Command Timing

- 1T and 2T timings
  - Depending on how many DIMM modules 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 turn around 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

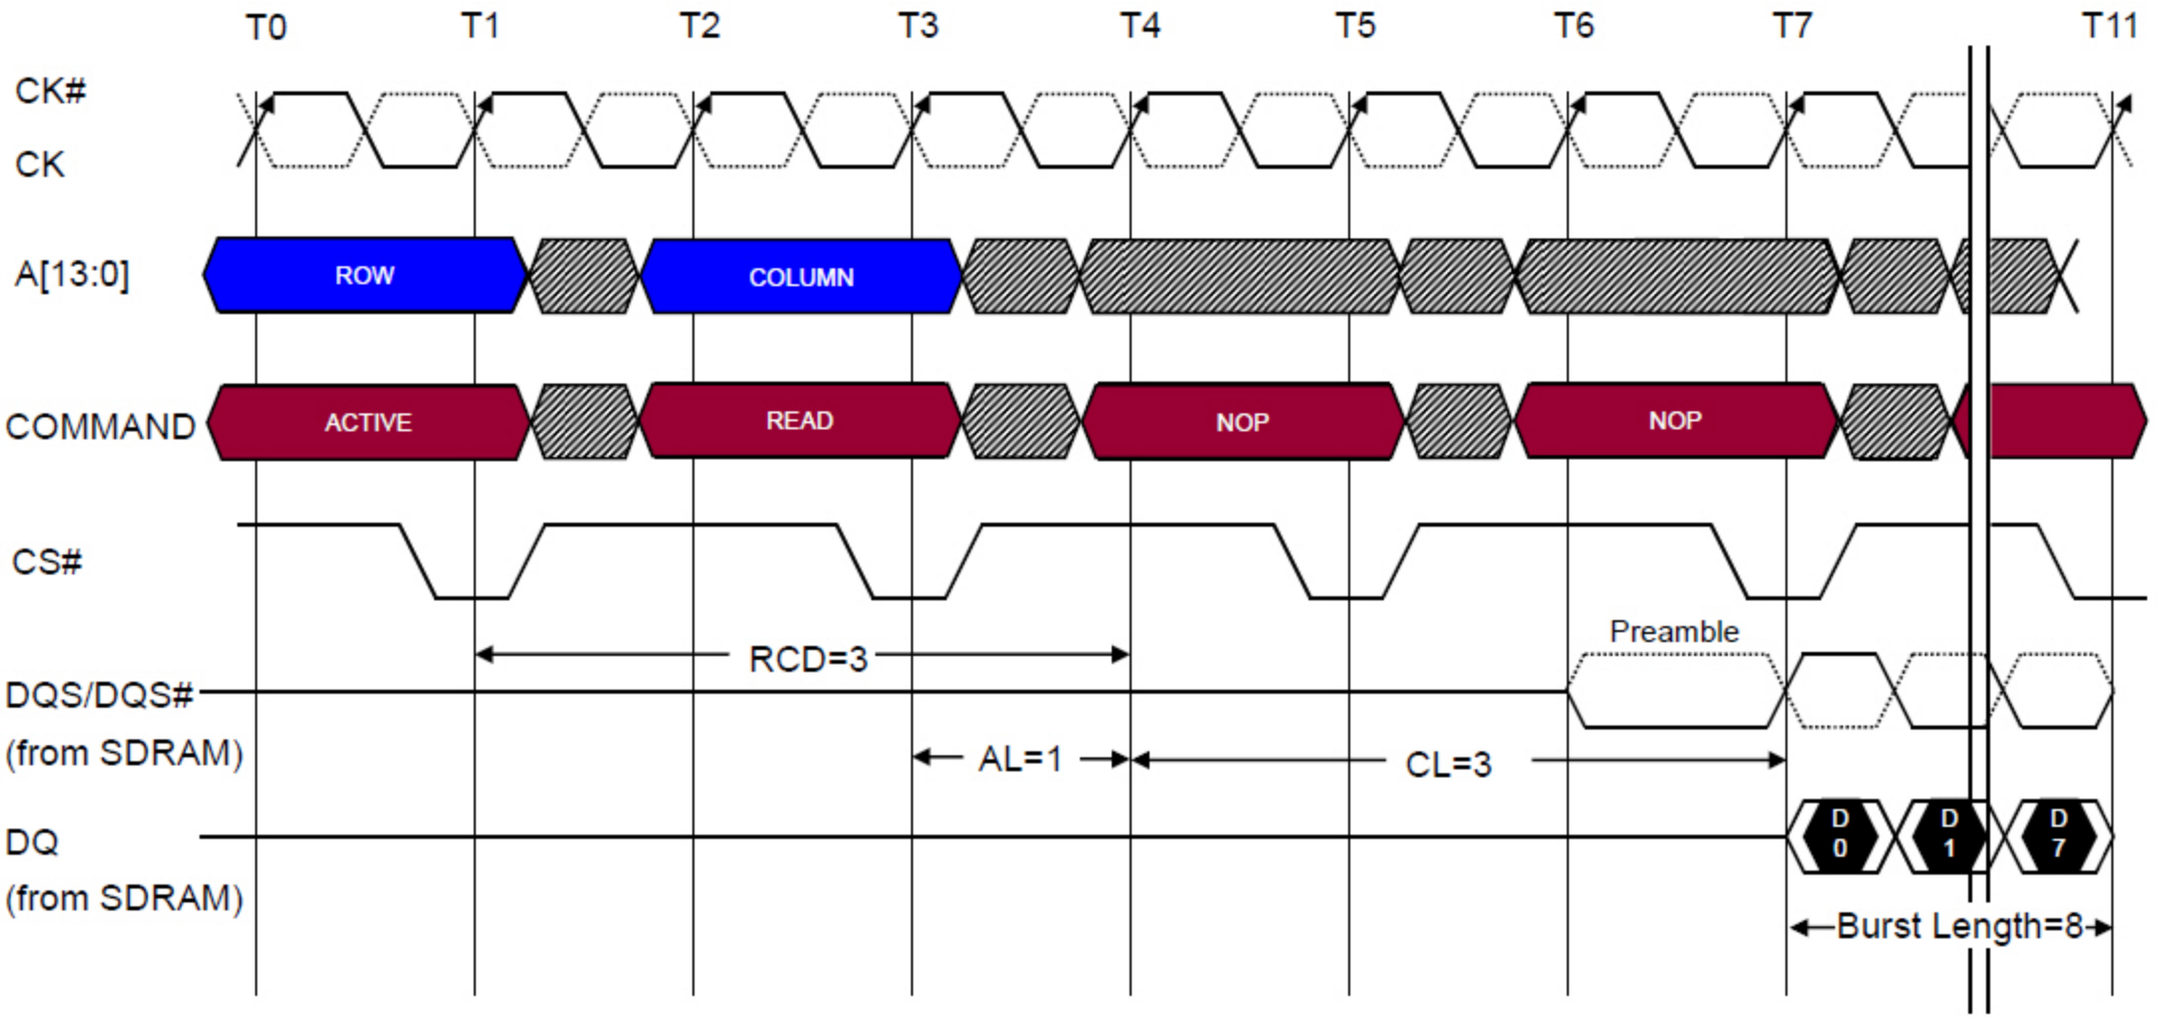

MINDSHARE

BRINGING LIFE TO KNOWLEDGE

DDR Initialization

DDR Initialization Routine

- 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.
- For comparison, DDR1 initialization is shown followed by DDR3 initialization.

Step by Step DDR1 Initialization

- Step 1
  - Provide power. This includes the device core (V_{DDQ}) and IO power (V_{TT}) according to the SSTL spec. Both supplies must be brought up simultaneous to avoid latch-up.
- Step 2
  - Apply the reference voltage (V_{REF}). During this time it is important that the IO voltage does not exceed the core. Typically this is not a problem due to the termination on the IO pins.

Step by Step DDR1 Initialization cont.

- Step 3
  - CKE must be driven and held in a logic level low. This will keep the DRAM from receiving unwanted commands during initialization.
- Step 4
  - Once reliable power levels have been reached and CKE has been driven low it is OK to apply a stable clock.
- Step 5
  - After clock is valid for greater than 200uS, CKE will go high and the DRAMs are ready to receive valid commands.

Step by Step DDR1 Initialization cont.

- Step 6
  - To initialize the internal logic after bringing CKE high, you must provide a NOP or Deselect command.
- Step 7
  - Issue a Precharge All command.
- Step 8
  - Provide NOP or Deselect command until t_{RP} is met.

Step by Step DDR1 Initialization cont.

- Step 9
  - Using the Mode Register Set command program the Delay Locked Loop (DLL) and IO drive strength. This is the first 2 bits of EMR1. All other bits must be set to 0.
- Step 10
  - Provide NOP or Deselect command for at least t_{MRD}.
- Step 11
  - Using the Mode Register Set command program all of the Mode Registers to the desired operating conditions

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

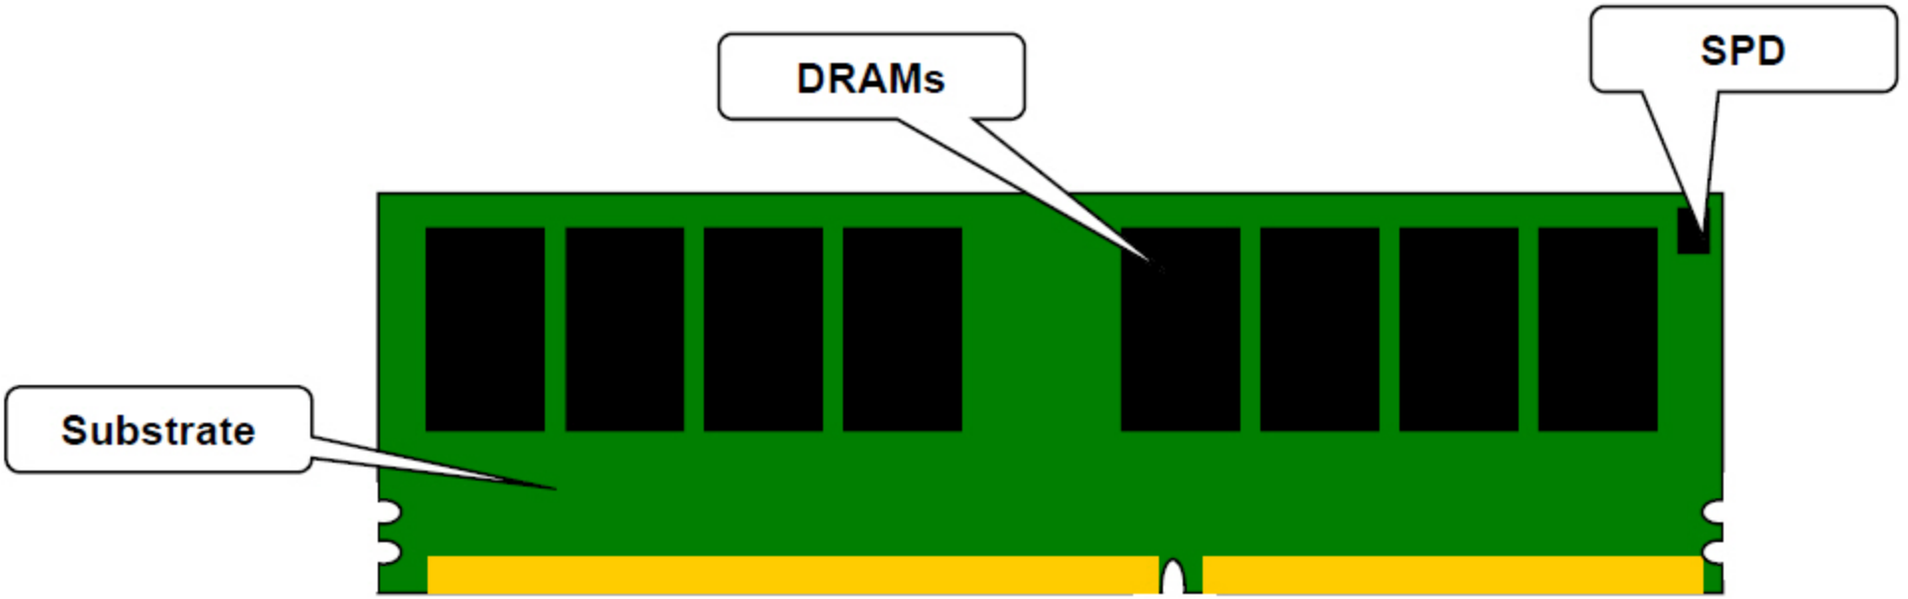

Step by Step DDR1 Initialization cont.

- During Step 11, the BIOS/Firmware has read all data from SPD and worst-case timings are 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.

SPD Data

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

SPD Data cont.

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)	FFFFFFFFFFFFFF
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

DDR Initialization cont.

- 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 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 DDR1 Initialization MR0

A8	DLL Reset
0	No
1	Yes

A7	mode
0	Normal
1	Test

A3	Burst Type
0	Sequential
1	Interleave

Burst Length

A2	A1	A0	BL
0	1	0	4
0	1	1	8

A12	Active power down exit time
0	Fast exit (use tXARD)
1	Slow exit (use tXARDS)

Write recovery for autoprecharge

A11	A10	A9	WR(cycles)
0	0	0	Reserved
0	0	1	2
0	1	0	3
0	1	1	4
1	0	0	5
1	0	1	6
1	1	0	Reserved
1	1	1	Reserved

*2
DDR2-400
DDR2-533
DDR2-667
DDR2-800

A6	A5	A4	CAS Latency
0	0	0	Reserved
0	0	1	Reserved
0	1	0	2(Optional)
0	1	1	3(speed bin determined)*3
1	0	0	4
1	0	1	5(speed bin determined)*3
1	1	0	6(speed bin determined)*3
1	1	1	Reserved

BA1	BA0	MRS mode
0	0	MR
0	1	EMR(1)
1	0	EMR(2)
1	1	EMR(3)

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 autoprecharge) 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[cycles] = RU{ tWR[ns] / tCK[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 DDR1 Initialization EMR1

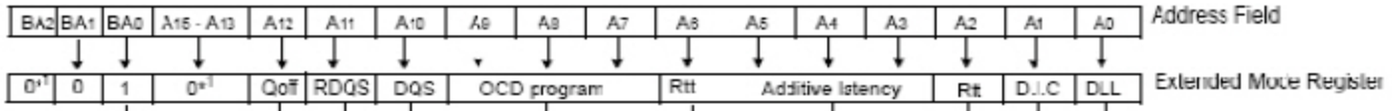

BA1	BA0	MRS mode
0	0	MR
0	1	EMR(1)
1	0	EMR(2)
1	1	EMR(3)

A6	A2	Rtt (NOMINAL)
0	0	ODT Disabled
0	1	75 Ω
1	0	150 Ω
1	1	50 Ω*2

A0	DLL Enable
0	Enable
1	Disable

A1	Output Driver Impedance Control
0	Full strength
1	Reduced strength

A9	A8	A7	OCD Calibration Program
0	0	0	OCD Calibration mode exit; maintain setting
0	0	1	Drive(1)
0	1	0	Drive(0)
1	0	0	Adjust mode*3
1	1	1	OCD Calibration default *4

A5	A4	A3	Additive Latency
0	0	0	0
0	0	1	1
0	1	0	2
0	1	1	3
1	0	0	4
1	0	1	5 (Optional)
1	1	0	Reserved
1	1	1	Reserved

A12	Qoff (Optional)*5
0	Output buffer enabled
1	Output buffer disabled

A10	$\overline{DQS}$
0	Enable
1	Disable

A11	RDQS Enable*6
0	Disable
1	Enable

A11 (RDQS Enable)	A10 ($\overline{DQS}$ Enable)	Strobe Function Matrix			
		RDQS/DM	$\overline{RDQS}$	DQS	$\overline{DQS}$
0 (Disable)	0 (Enable)	DM	Hi-z	DQS	$\overline{DQS}$
0 (Disable)	1 (Disable)	DM	Hi-z	DQS	Hi-z
1 (Enable)	0 (Enable)	RDQS	$\overline{RDQS}$	DQS	$\overline{DQS}$
1 (Enable)	1 (Disable)	RDQS	Hi-z	DQS	Hi-z

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, DQSs, $\overline{DQS}$s, RDQS, $\overline{RDQS}$. 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 DDR1 Initialization EMR2

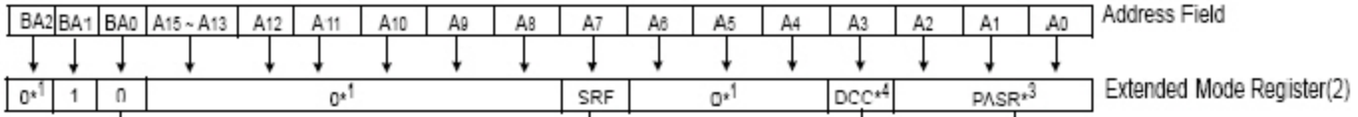

A7	High Temperature Self-Refresh Rate Enable
0	Disable
1	Enable (Optional)*2

BA1	BA0	MRS mode
0	0	MR
0	1	EMR(1)
1	0	EMR(2)
1	1	EMR(3)

A3	DCC Enable (Optional)*4
0	Disable
1	Enable

A2	A1	A0	Partial Array Self Refresh for 8 Banks (Optional)
0	0	0	Full array
0	0	1	Half Array(BA[2:0]=000,001,010&011)
0	1	0	Quarter Array(BA[2:0]=000&001)
0	1	1	1/8th array (BA[2:0] = 000)
1	0	0	3/4 array(BA[2:0]=010,011,100,101,110&111)
1	0	1	Half array(BA[2:0]=100,101,110&111)
1	1	0	Quarter array(BA[2:0]=110&111)
1	1	1	1/8th array(BA[2:0]=111)

A2	A1	A0	Partial Array Self Refresh for 4 Banks (Optional)
0	0	0	Full array
0	0	1	Half Array(BA[1:0]=00&01)
0	1	0	Quarter Array(BA[1:0]=00)
0	1	1	Not defined
1	0	0	3/4 array(BA[1:0]=01, 10&11)
1	0	1	Half array(BA[1:0]=10&11)
1	1	0	Quarter array(BA[1:0]=11)
1	1	1	Not defined

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 tREF 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 DDR1 Initialization EMR3

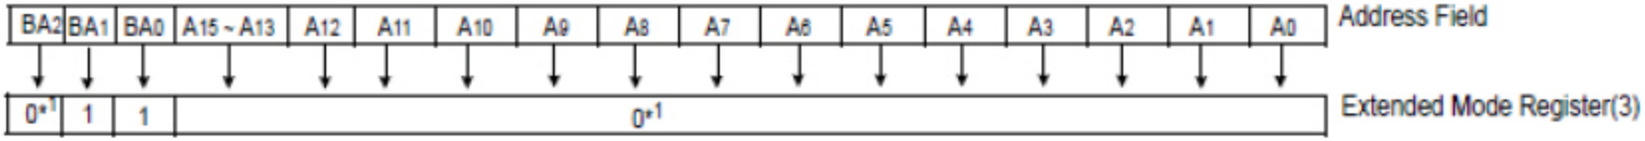

DDR3 Mode Registers

- 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

A8	DLL Reset
0	No
1	Yes

A7	mode
0	Normal
1	Test

A3	Read Burst Type
0	Nibble Sequential
1	Interleave

Item#1586.05/06(Mar.'05)

A1	A0	BL
0	0	8 (Fixed)
0	1	4 or 8(on the fly)
1	0	4 (Fixed)
1	1	Reserved

A12	DLL Control for Precharge PD
0	Slow exit (DLL off)
1	Fast exit (DLL on)

Write recovery for autoprecharge

A11	A10	A9	WR(cycles)
0	0	0	Reserved
0	0	1	5*2
0	1	0	6*2
0	1	1	7*2
1	0	0	8*2
1	0	1	10*2
1	1	0	12*2
1	1	1	Reserved

A6	A5	A4	A2	CAS Latency
0	0	0	0	Reserved
0	0	1	0	5
0	1	0	0	6
0	1	1	0	7
1	0	0	0	8
1	0	1	0	9
1	1	0	0	10
1	1	1	0	11(Optional for DDR3-1600)

BA1	BA0	MR Select
0	0	MR0
0	1	MR1
1	0	MR2
1	1	MR3

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

Mode Register 1

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

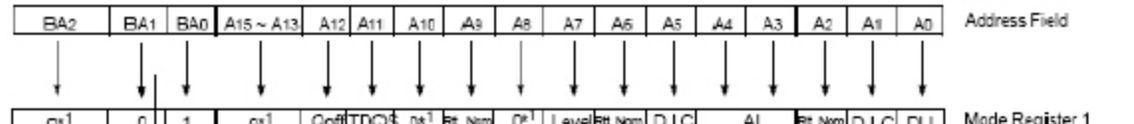

A11	TDQS enable
0	Disabled
1	Enabled

A7	Write leveling enable
0	Disabled
1	Enabled

A4	A3	Additive Latency
0	0	0 (AL disabled)
0	1	CL-1
1	0	CL-2
1	1	Reserved

A12	Qoff *2
0	Output buffer enabled
1	Output buffer disabled *2

*2: Outputs disabled - DQs, DQSs, $\overline{DQS}$s.

BA1	BA0	MR Select
0	0	MR0
0	1	MR1
1	0	MR2
1	1	MR3

A9	A6	A2	Rtt_Nom *3
0	0	0	ODT disabled
0	0	1	RZQ/4
0	1	0	RZQ/2
0	1	1	RZQ/6
1	0	0	RZQ/12*2
1	0	1	RZQ/8*2
1	1	0	Reserved
1	1	1	Reserved

Note 1: RZQ = 240 Ω
Note 2: If RTT_Nom is used during Writes, only the values RZQ/2, RZQ/4 and RZQ/6 are allowed.
Note 3: In Write leveling Mode (MR1 [bit7] = 1) with MR1[bit12]=1, all RTT_Nom settings are allowed; in Write Leveling Mode (MR1[bit7] = 1) with MR1[bit12]=0, only RTT_Nom settings of RZQ/2, RZQ/4 and RZQ/6 are allowed.

A0	DLL Enable
0	Enable
1	Disable

Output Driver strength removed

A5	A1	Output Driver Impedance Control
0	0	Reserved for RZQ/6
0	1	RZQ/7
1	0	RZQ/TBD
1	1	RZQ/TBD

Note 1: RZQ = 240 Ω

Mode Register 2

A7	Self-Refresh Temperature (SRT) Range
0	Normal operating temperature range
1	Extended (optional) operating temperature range

A6	Auto Self-Refresh (ASR)
0	Manual SR Reference (SRT)
1	ASR enable (Optional)

A2	A1	A0	Partial Array Self-Refresh (Optional)
0	0	0	Full Array
0	0	1	HalfArray (BA[2:0]=000,001,010, &011)
0	1	0	Quarter Array (BA[2:0]=000, & 001)
0	1	1	1/8th Array (BA[2:0] = 000)
1	0	0	3/4 Array (BA[2:0] = 010,011,100,101,110, & 111)
1	0	1	HalfArray (BA[2:0] = 100, 101, 110, &111)
1	1	0	Quarter Array (BA[2:0]=110, &111)
1	1	1	1/8th Array (BA[2:0]=111)

A10	A9	Rtt_WR
0	0	Dynamic ODT off (Write does not affect Rtt value)
0	1	RZQ/4
1	0	RZQ/2
1	1	Reserved

A5	A4	A3	CAS write Latency (CWL)
0	0	0	5 (tCK ≥2.5ns)
0	0	1	6 (2.5ns >tCK≥1.875ns)
0	1	0	7 (1.875ns>tCK>1.5ns)
0	1	1	8 (1.5ns>tCK≥1.25ns)
1	0	0	Reserved
1	0	1	Reserved
1	1	0	Reserved
1	1	1	Reserved

BA1	BA0	MR Select
0	0	MR0
0	1	MR1
1	0	MR2
1	1	MR3

4 Bank Partial Array was removed
Duty Cycle Control was removed

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

Mode Register 2 Changes

- 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

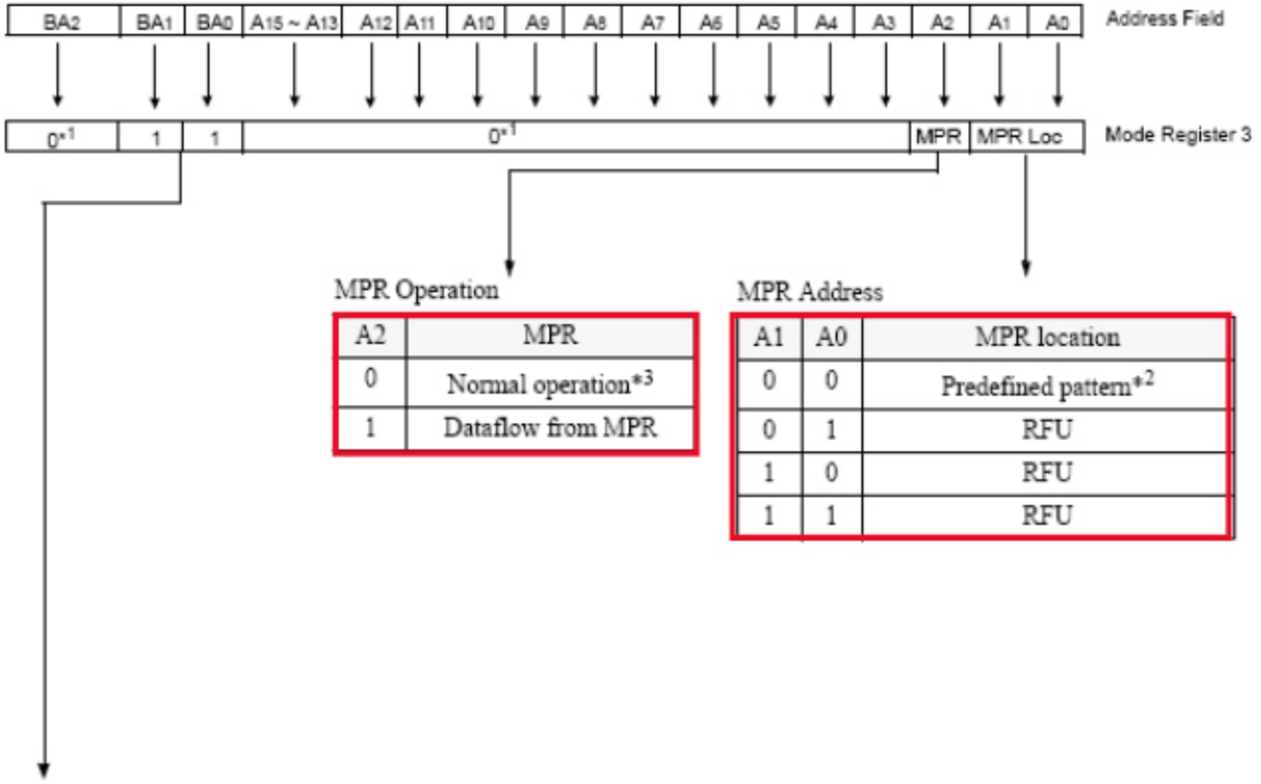

MPR Operation

A2	MPR
0	Normal operation*3
1	Dataflow from MPR

MPR Address

A1	A0	MPR location
0	0	Predefined pattern*2
0	1	RFU
1	0	RFU
1	1	RFU

BA1	BA0	MR Select
0	0	MR0
0	1	MR1
1	0	MR2
1	1	MR3

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

Step by Step DDR1 Initialization cont.

- Step 12
  - Provide NOP or Deselect command for at least t_{MRD}.
- Step 13
  - Issue a Precharge All command with A10 set high.
- Step 14
  - Provide NOP or Deselect command until t_{RP} is met.

Step by Step DDR1 Initialization cont.

- Step 15
  - Issue a Refresh command. JEDEC requires that 2 Refresh commands be given after step 10.
- Step 16
  - Provide a NOP or Deselect command for at least t_{RFC}.
- Step 17
  - Issue the second Refresh.
- Step 18
  - Provide a NOP or Deselect command for at least t_{RFC}.

Step by Step DDR1 Initialization cont.

- Step 19
  - Although not required by all DRAM devices JEDEC requires a DLL reset at this time.
- Step 20
  - Provide NOP or Deselect command for at least t_{MRD}.
- Done!
  - The spec 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 x 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 by Step DDR3 Initialization cont.

- 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)

Step by Step DDR3 Initialization cont.

- 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 by Step DDR3 Initialization cont.

- Step 10
  - Issue ZQCL command to start ZQ calibration.
- Step 11
  - Wait until both tDLLK and tZQinit complete.

- Done!
  - The spec 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.

MINDSHARE

BRINGING LIFE TO KNOWLEDGE

SMBus Overview

Hidden during presentation

MINDSHARE

BRINGING LIFE TO KNOWLEDGE

DDR1 and DDR2 Routing

Design Challenges for DDR2

- 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

Basic Routing Rules

- 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.

Basic Routing of Data

- 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.

DDR1 and DDR2 Routing

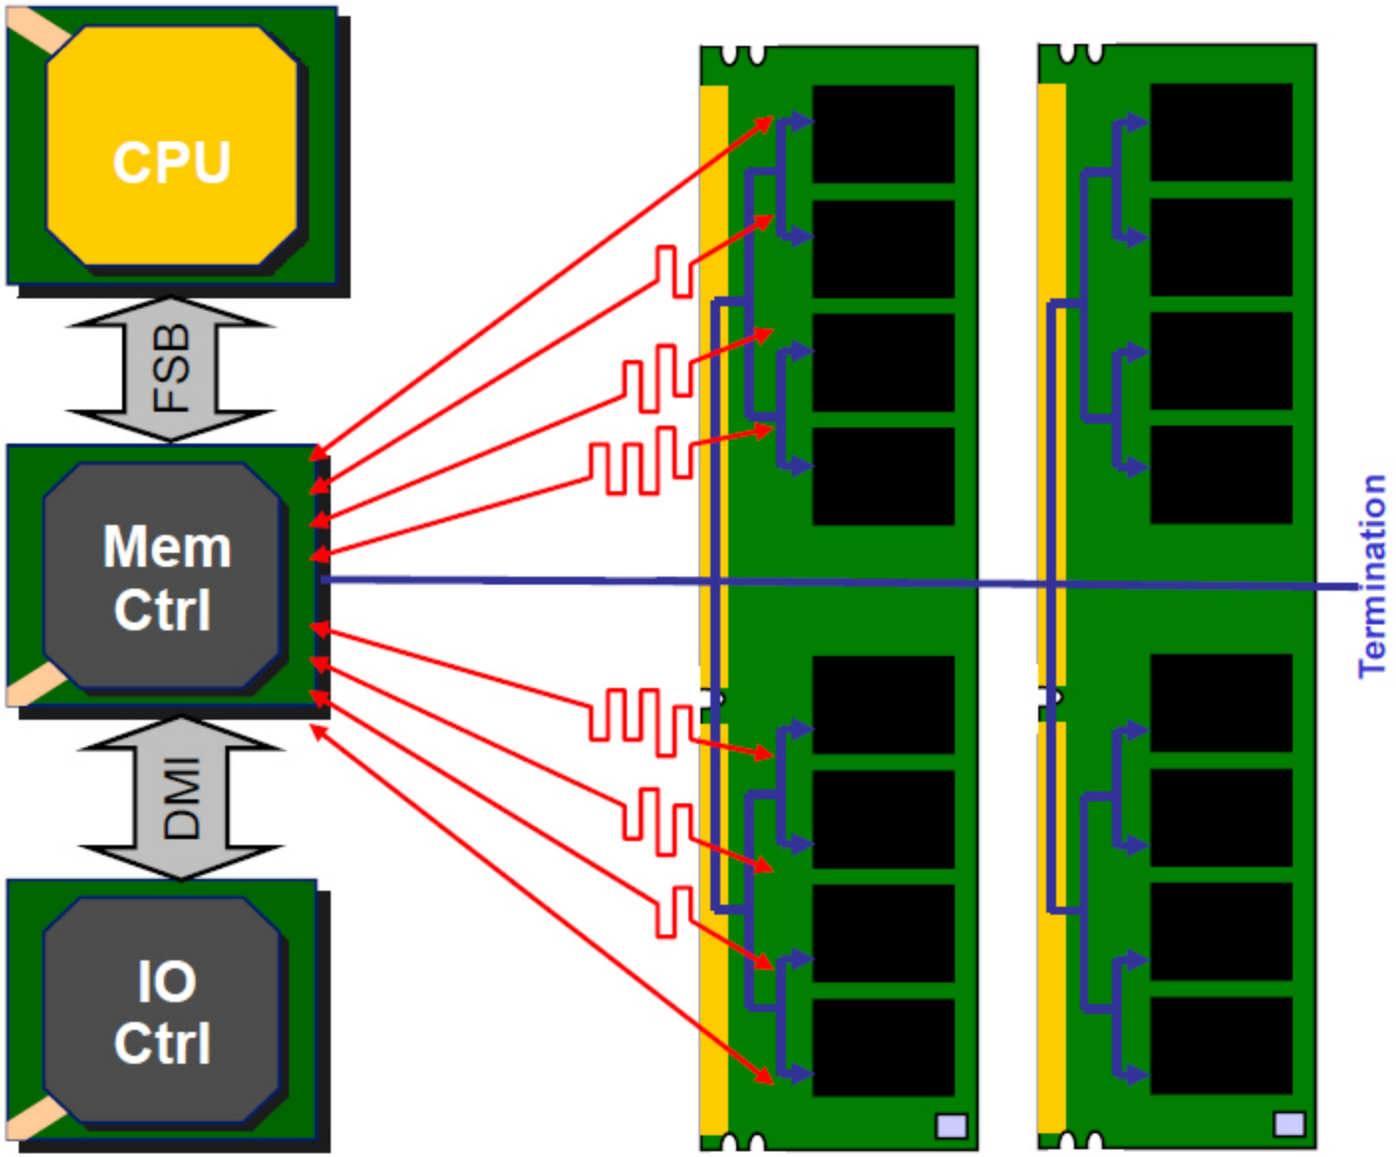

- CMD and ADDR are terminated on Motherboard.
- Data Bus is terminated on die.
- Byte lanes and their associated DQS must be length matched. For x4 chips, nibbles must be length matched with their DQS.
- Command and address are also length matched to ensure that all commands will get to each DRAM at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

Key: Red is Data Bus
Blue is Address and Command Bus

DDR2 Read Overview

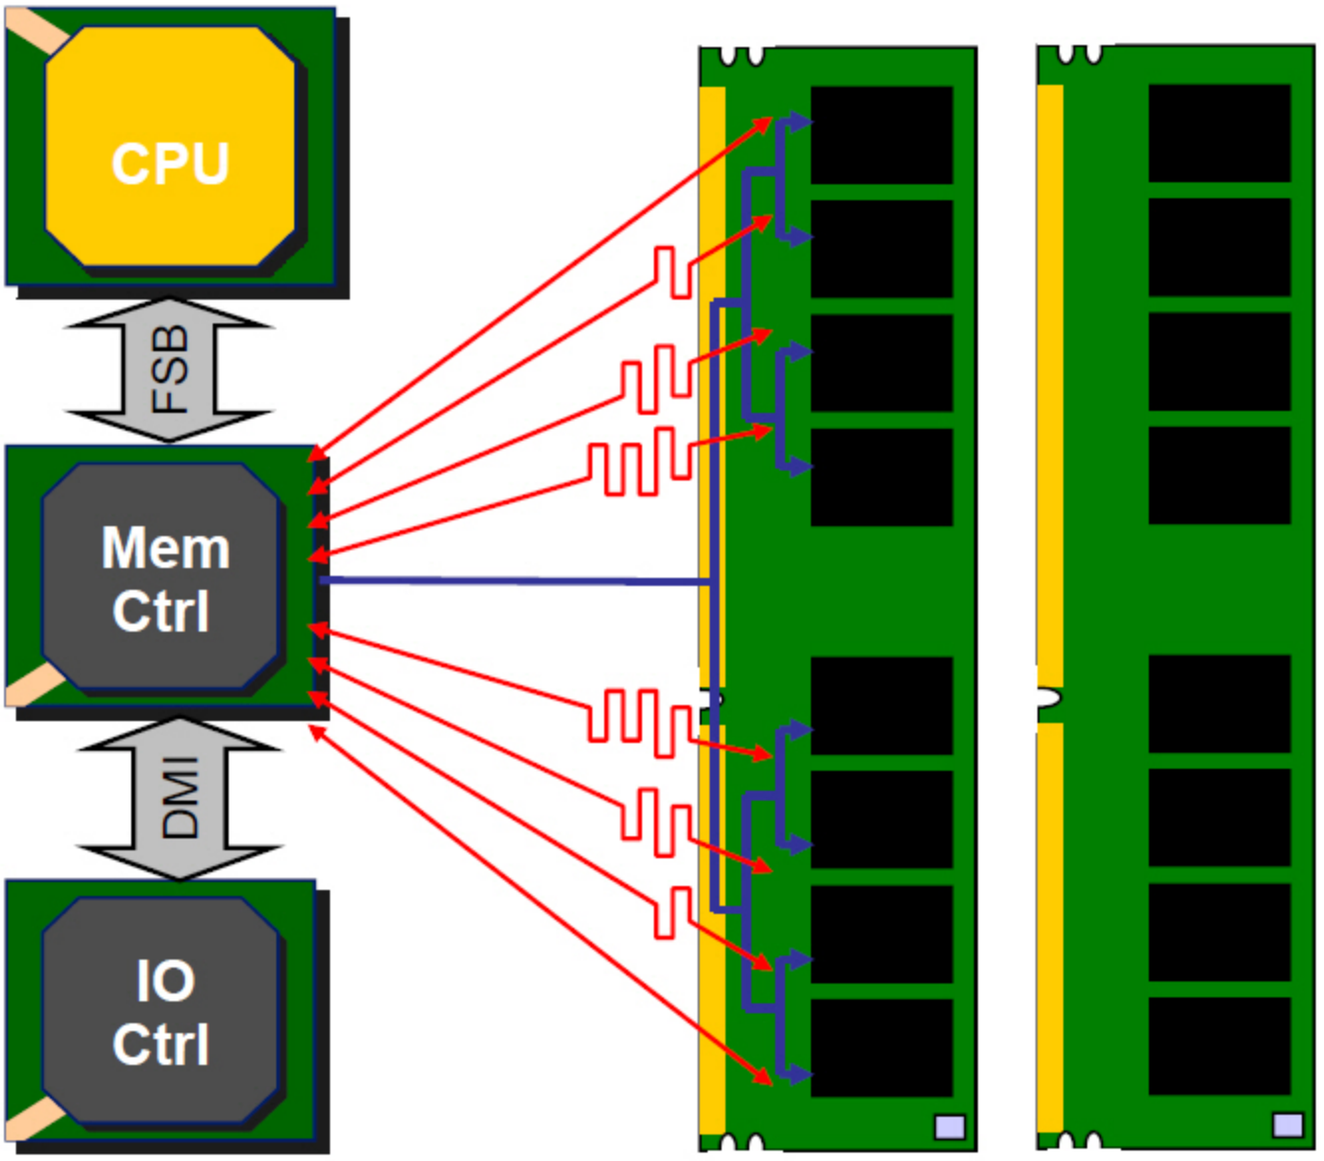

Key: Red is Data Bus
Blue is Address and Command Bus

- **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

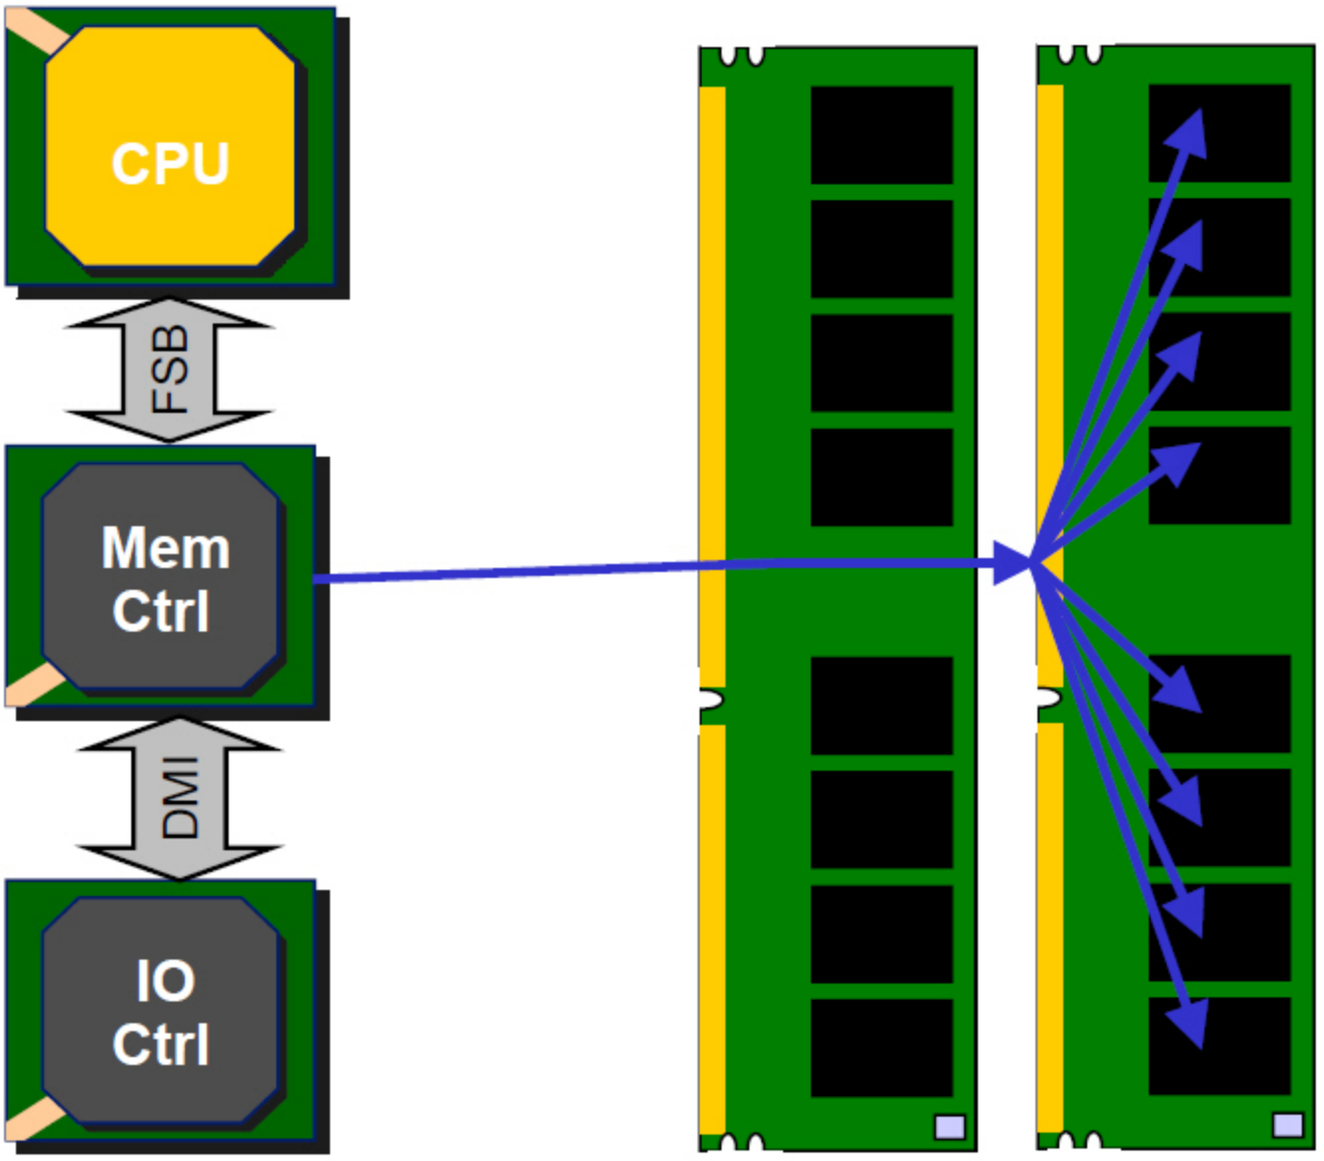

Key: Red is Data Bus
Blue is Address and Command Bus

- 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.

DDR2 Read Overview

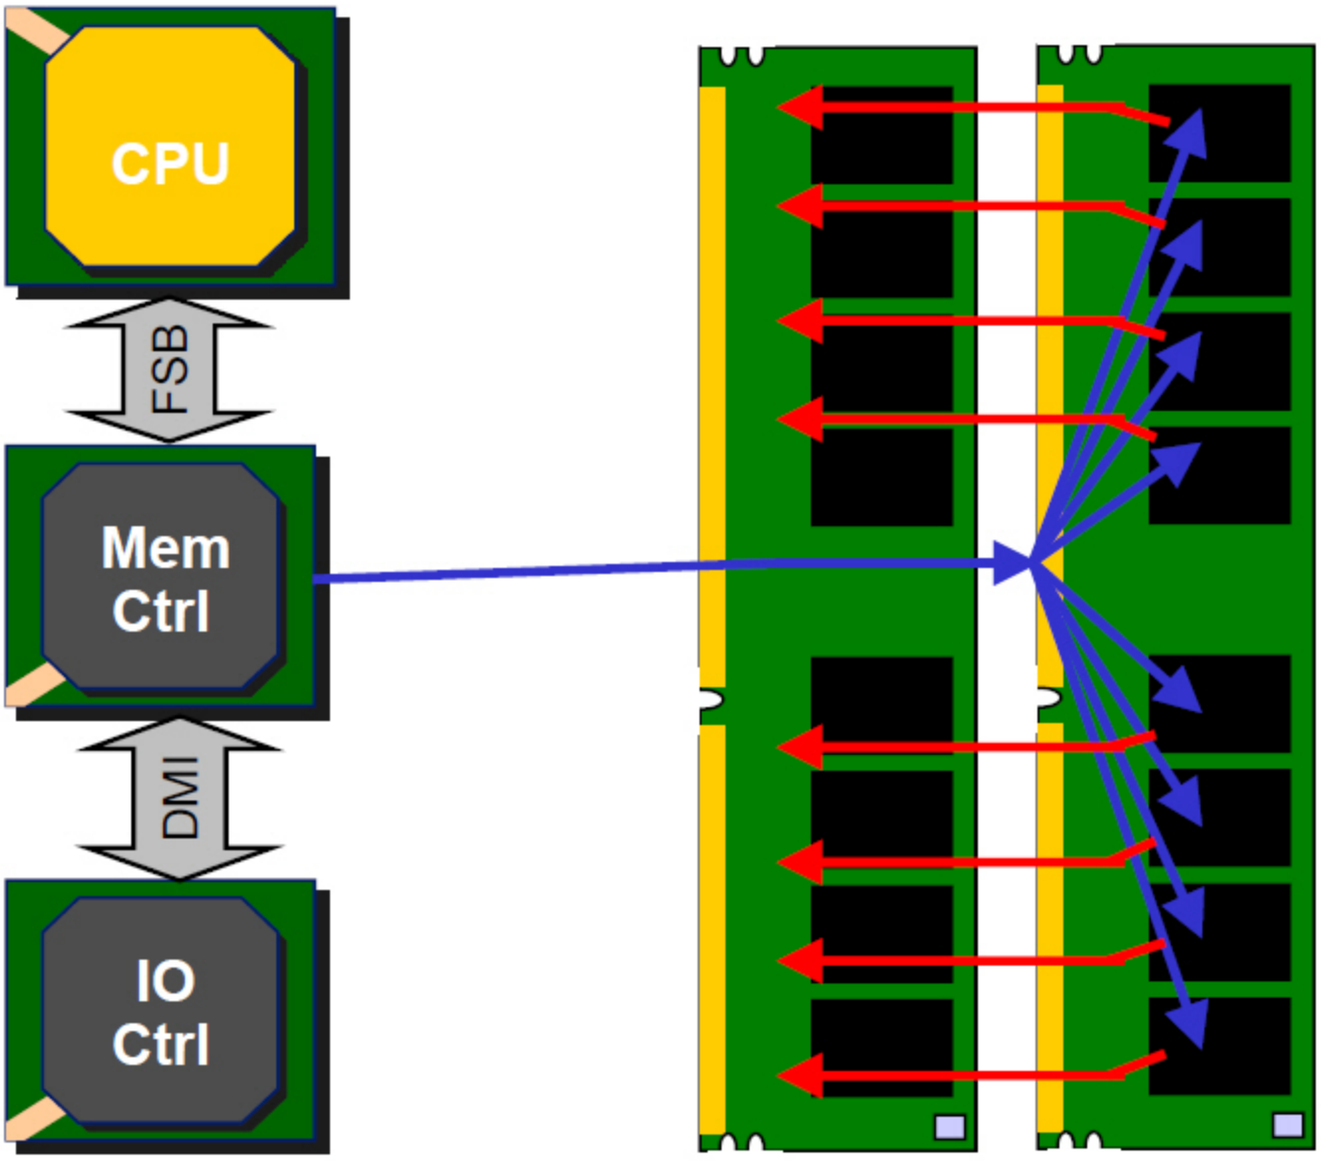

Key: Red is Data Bus
Blue is Address and Command Bus

- 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.

DDR2 Read Overview

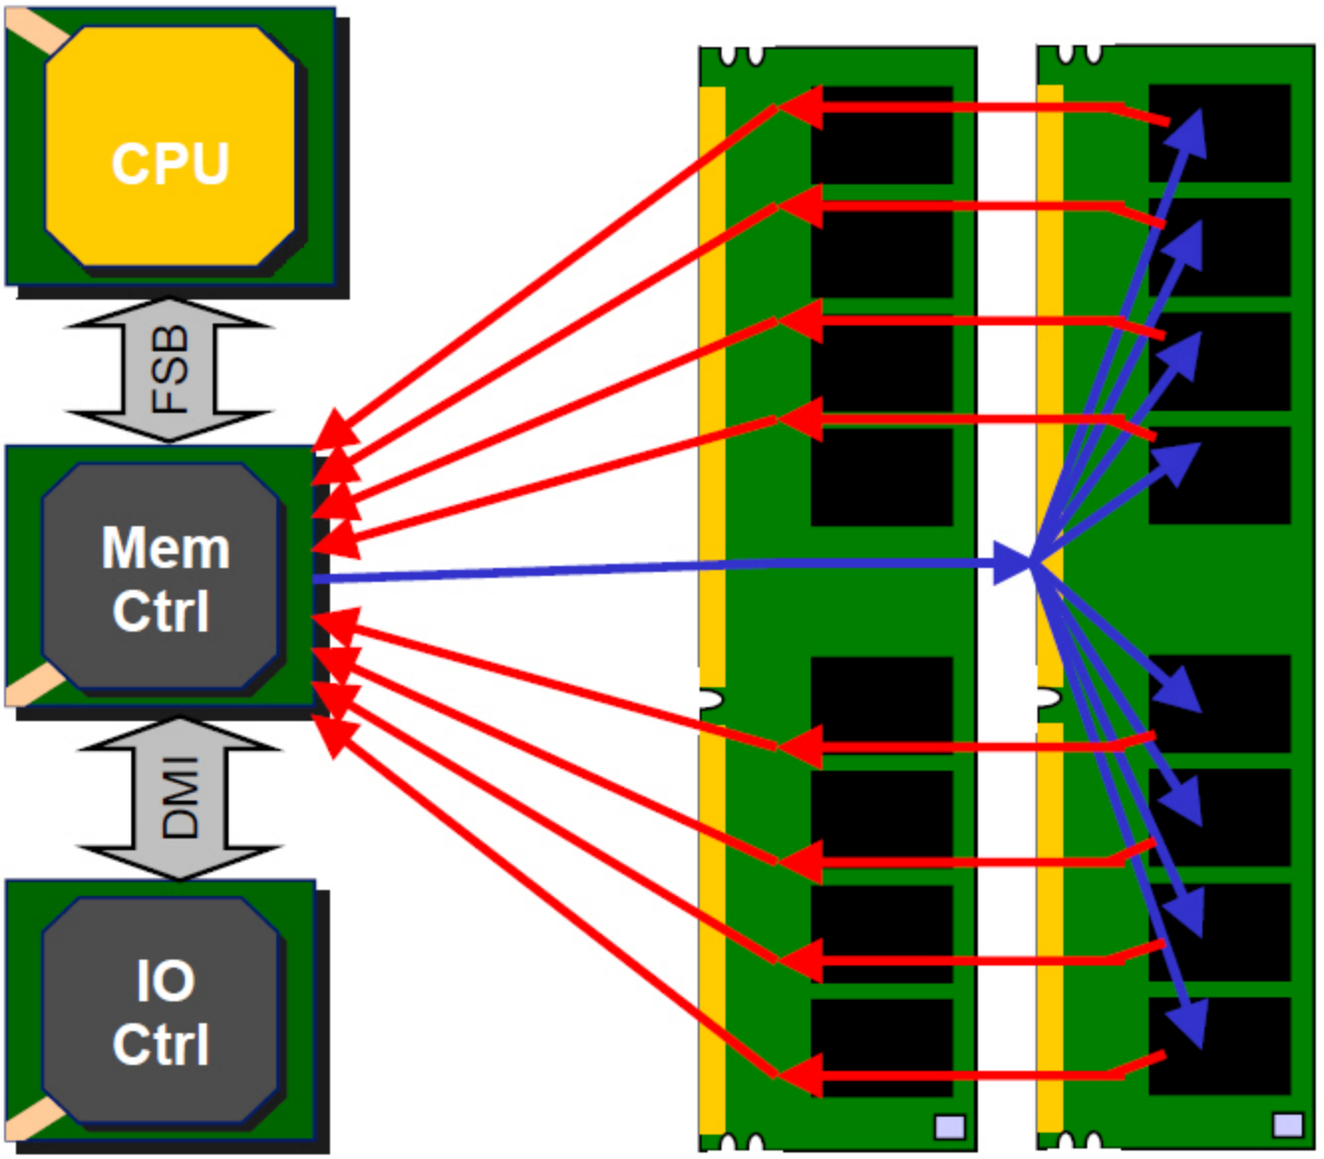

Key: Red is Data Bus
Blue is Address and Command Bus

- 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.

DDR2 Write Overview

- 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

DDR2 Write Overview

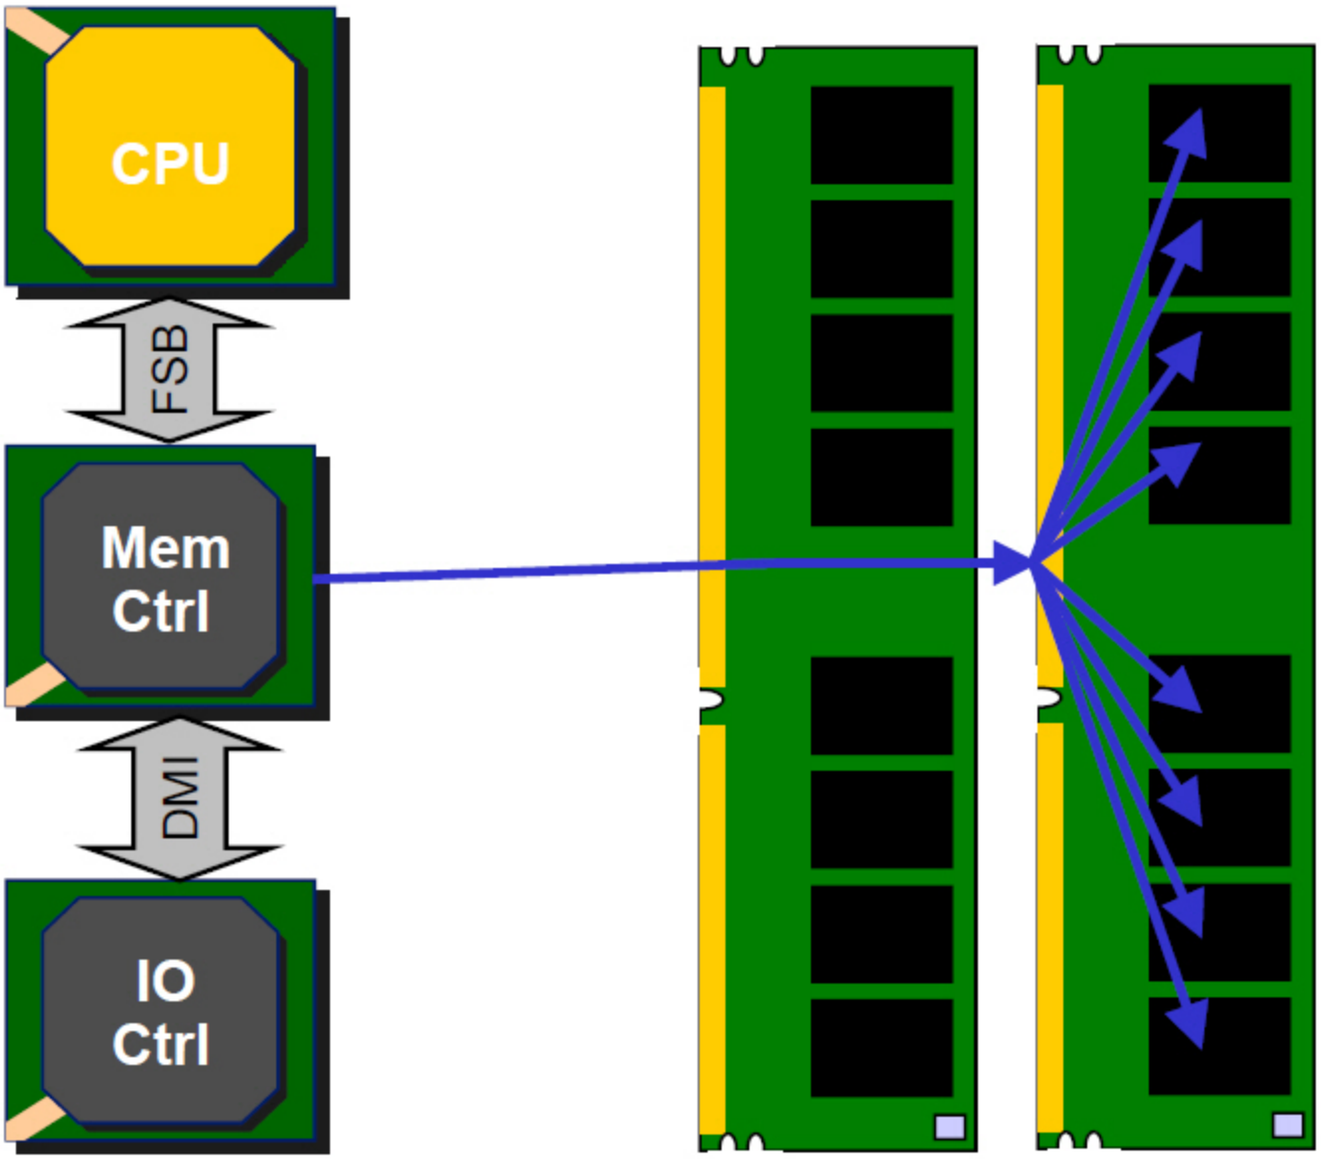

Key: Red is Data Bus
Blue is Address and Command Bus

- **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

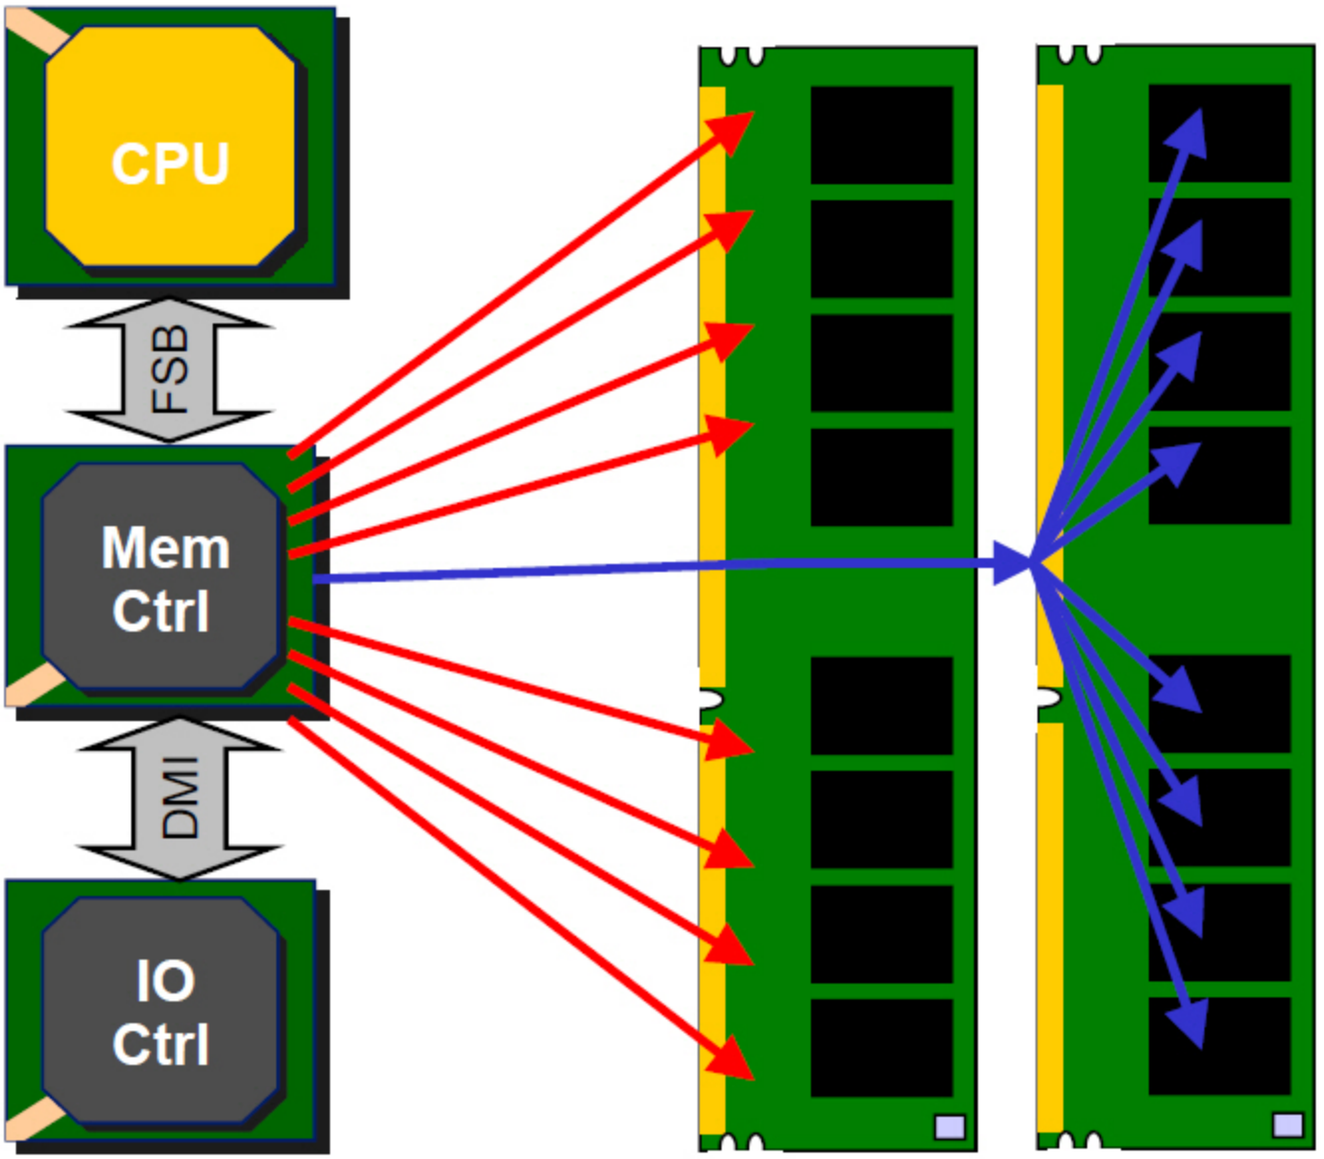

Key: Red is Data Bus
Blue is Address and Command Bus

- **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.

DDR2 Write Overview

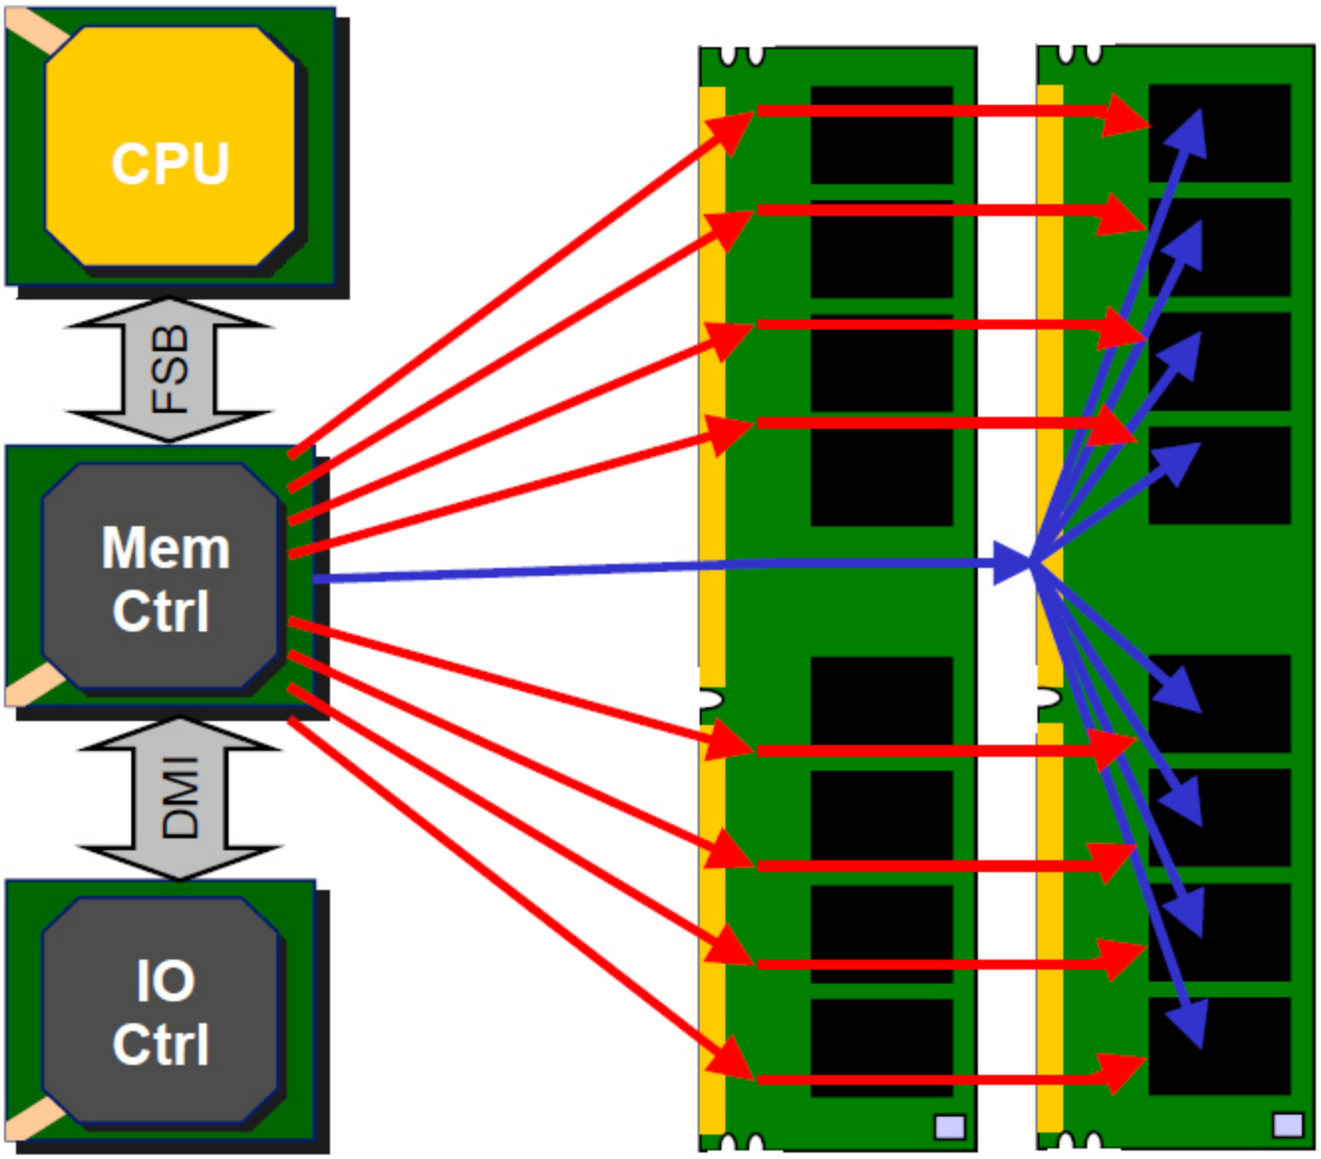

Key: Red is Data Bus
Blue is Address and Command Bus

- **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.

MINDSHARE

BRINGING LIFE TO KNOWLEDGE

On-Die Termination

On-Die Termination

- 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.

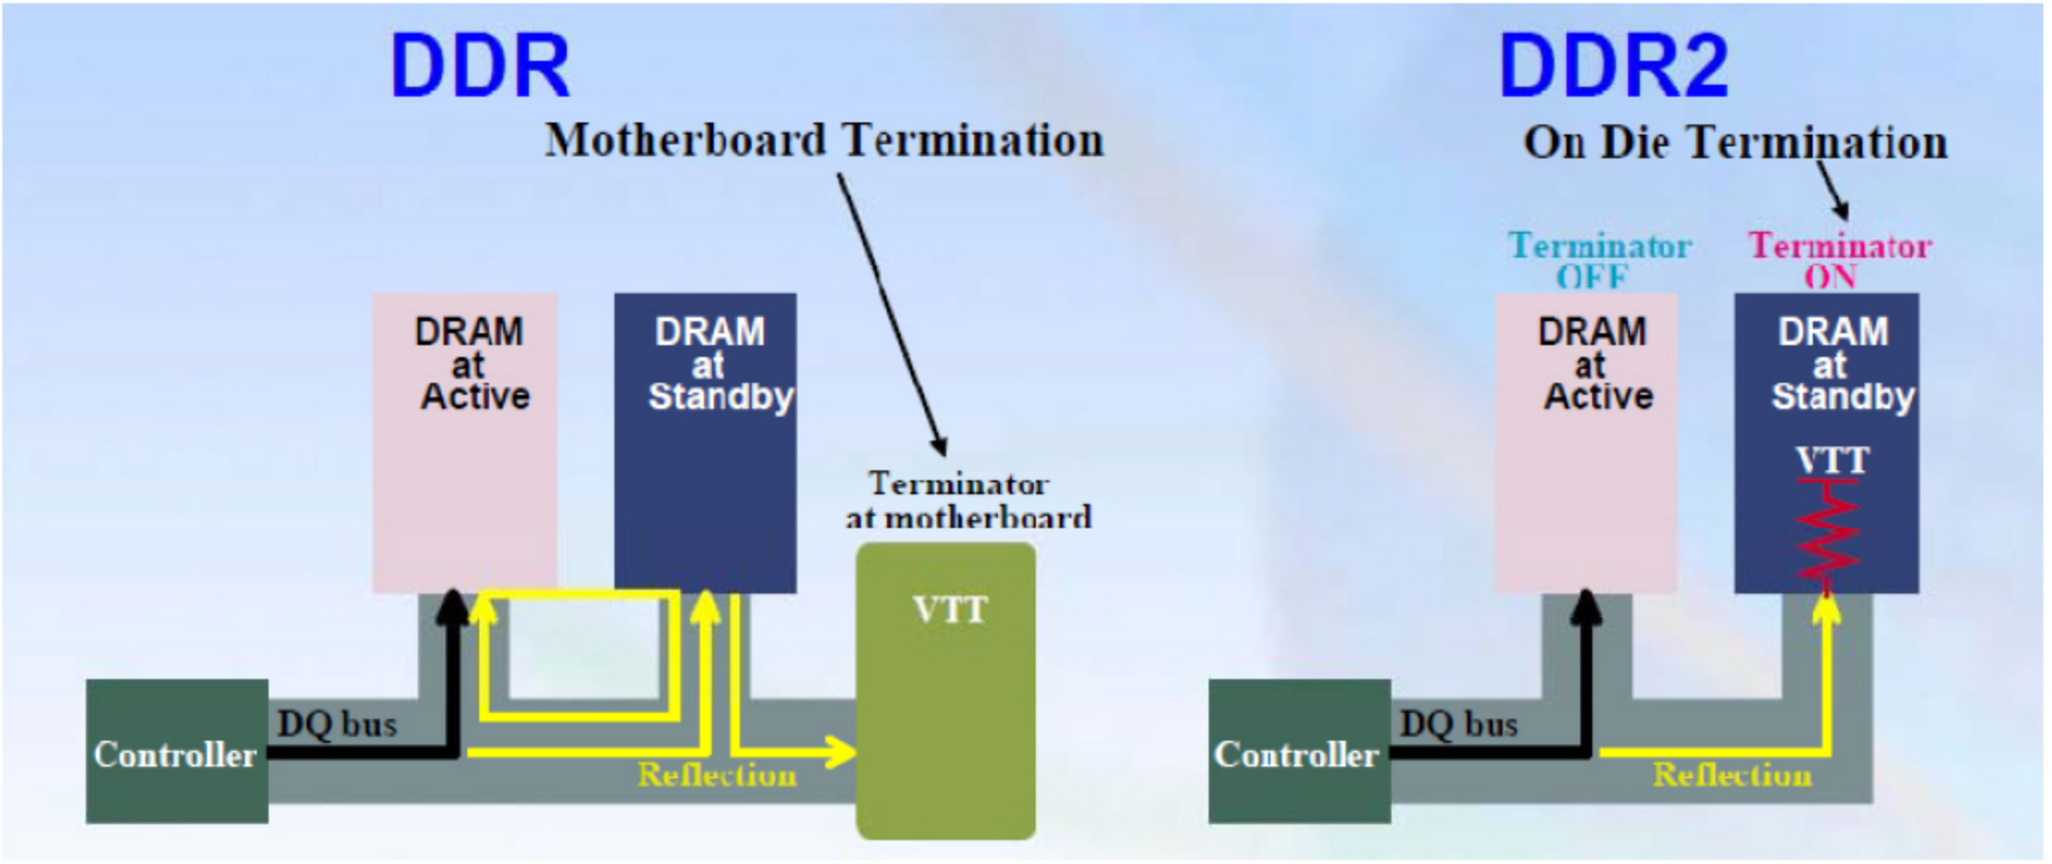

Signal Integrity with ODT

More noise margin, approx. 400mv gain

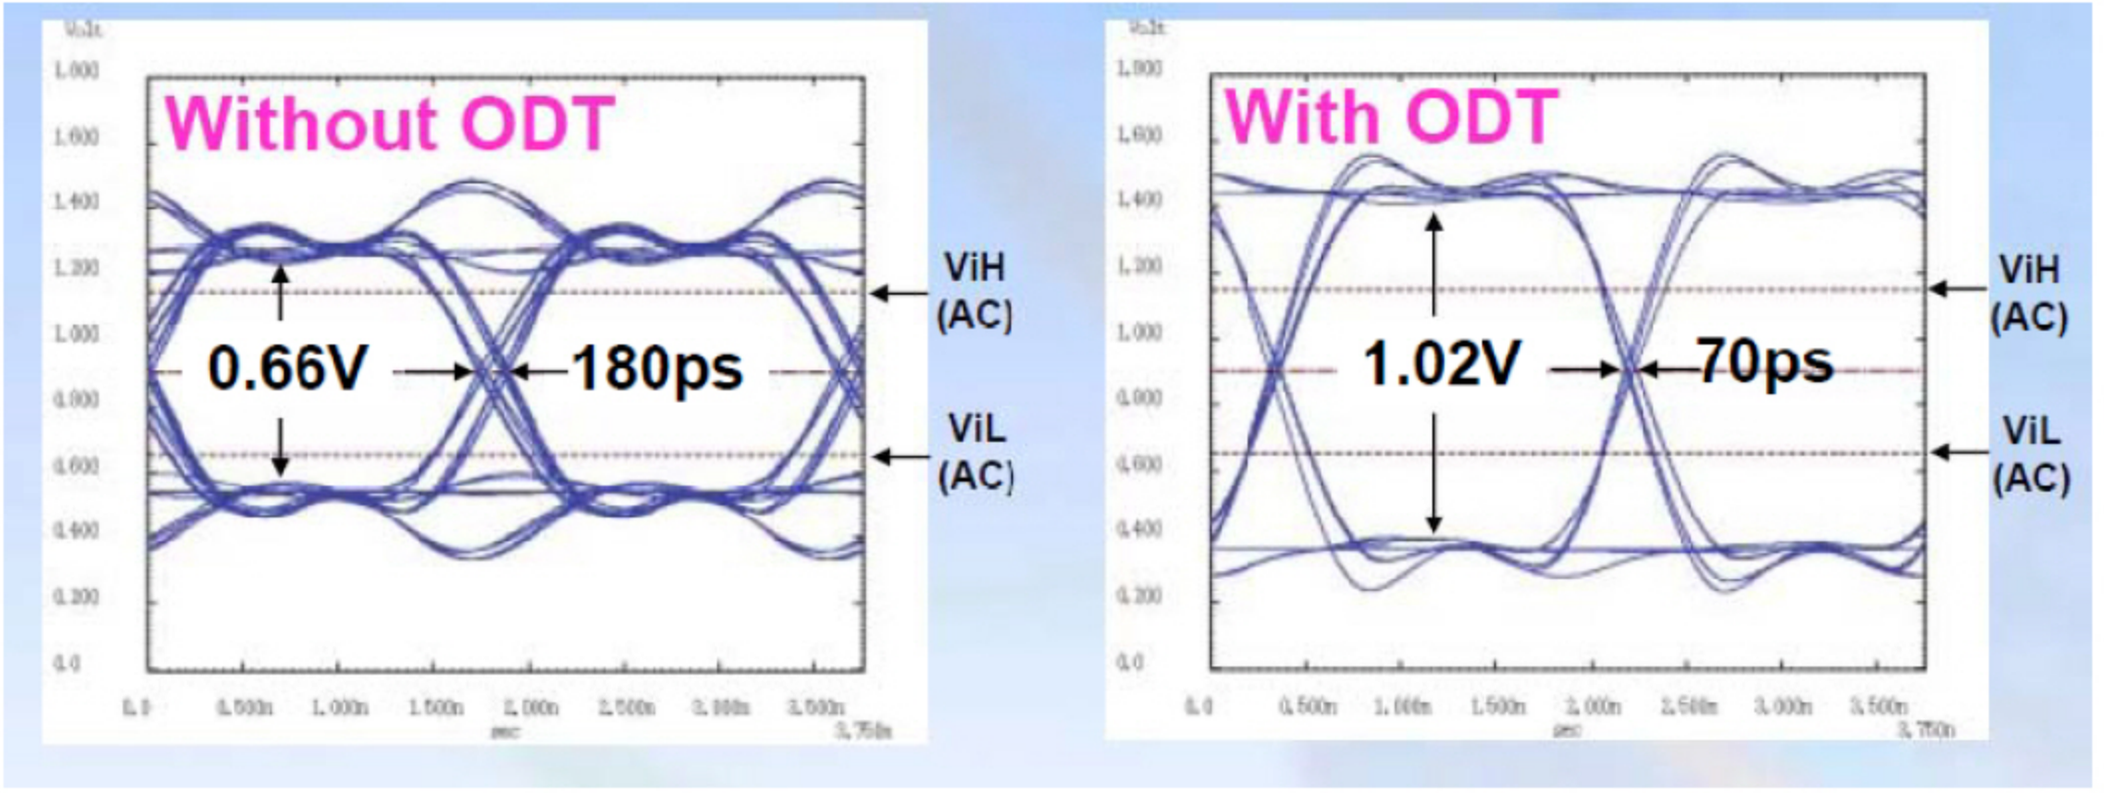

Off Chip Driver Calibration

- 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.

MINDSHARE

BRINGING LIFE TO KNOWLEDGE

Errors

Error Checking and Correction (ECC)

- 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)

Dual-Channel ECC

- 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.

ECC and Scrubbing

- 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.

DIMM Parity

- 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.

MINDSHARE

BRINGING LIFE TO KNOWLEDGE

Additional DDR3 Topics

MINDSHARE

BRINGING LIFE TO KNOWLEDGE

Timing and Electrical Differences

- t_{CCD} or CAS-to-CAS delay is no longer variable. In DDR2 it was ½ the burst length. DDR3 it is set to 4. (BL = 8)
- Receiver impedance went up to approx. 34 ohms. DDR2 was approx. 18 ohms. This equates to less power consumption.

MINDSHARE

BRINGING LIFE TO KNOWLEDGE

Fly-by Routing Read Example

DDR3 Read Overview

- 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.

DDR3 Read Overview

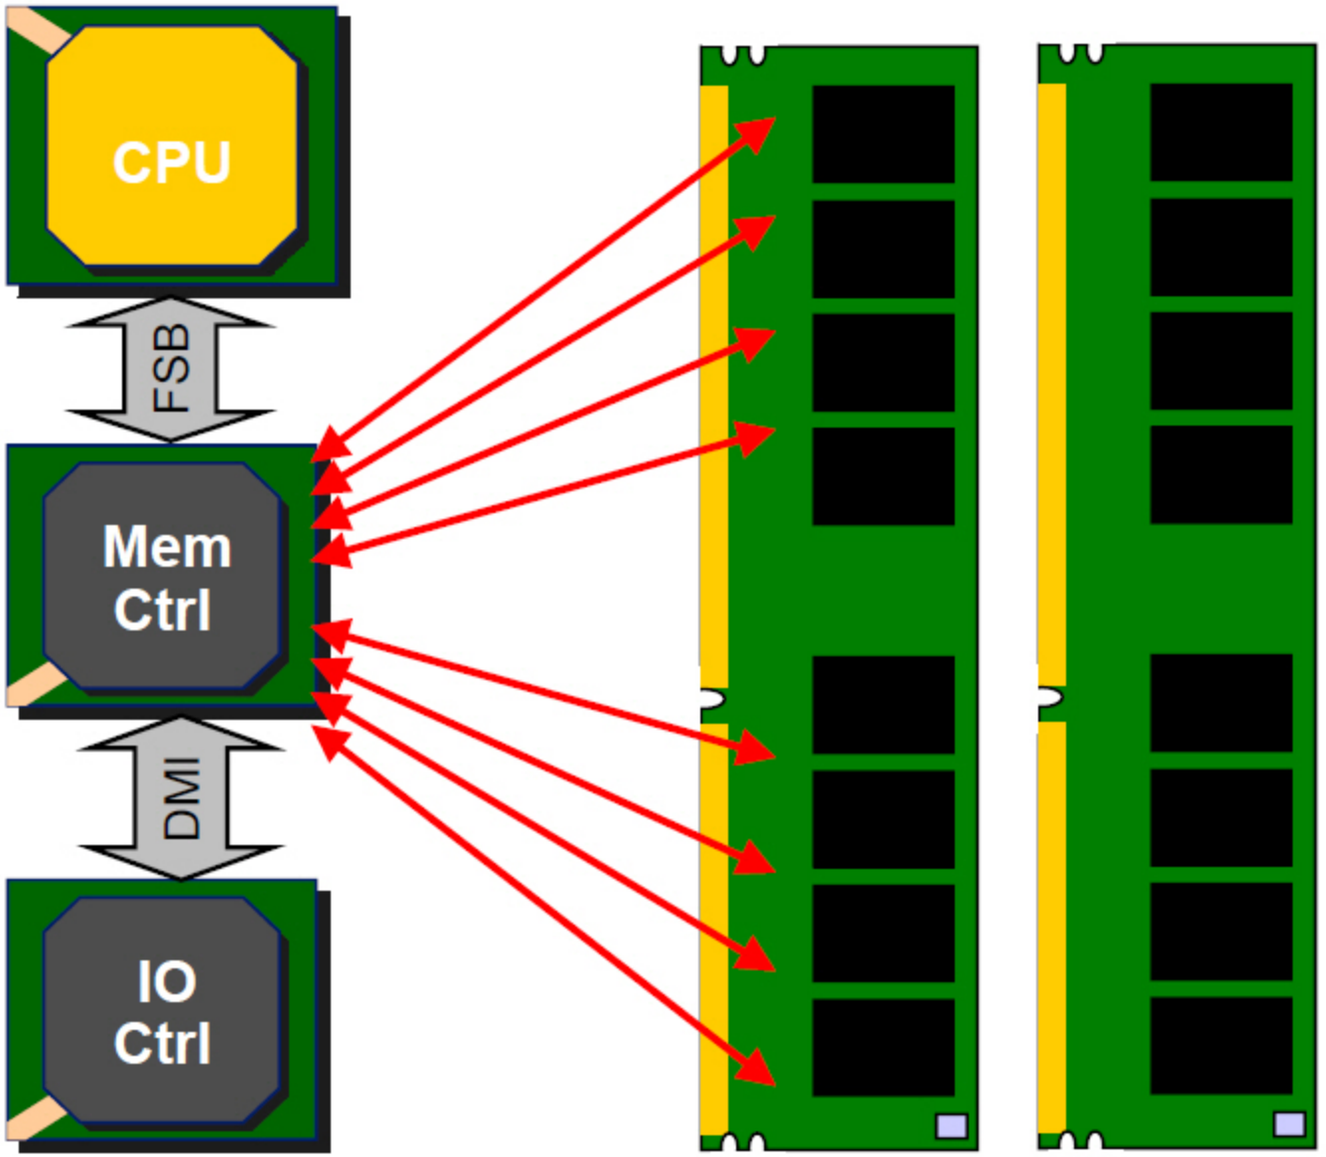

- **New fly-by routing**
- Data bus is still length matched within each data group to its associated strobe. There does not seem to be any implied length matching across the data groups

Key: Red is Data Bus
Blue is Address and Command Bus

DDR3 Read Overview

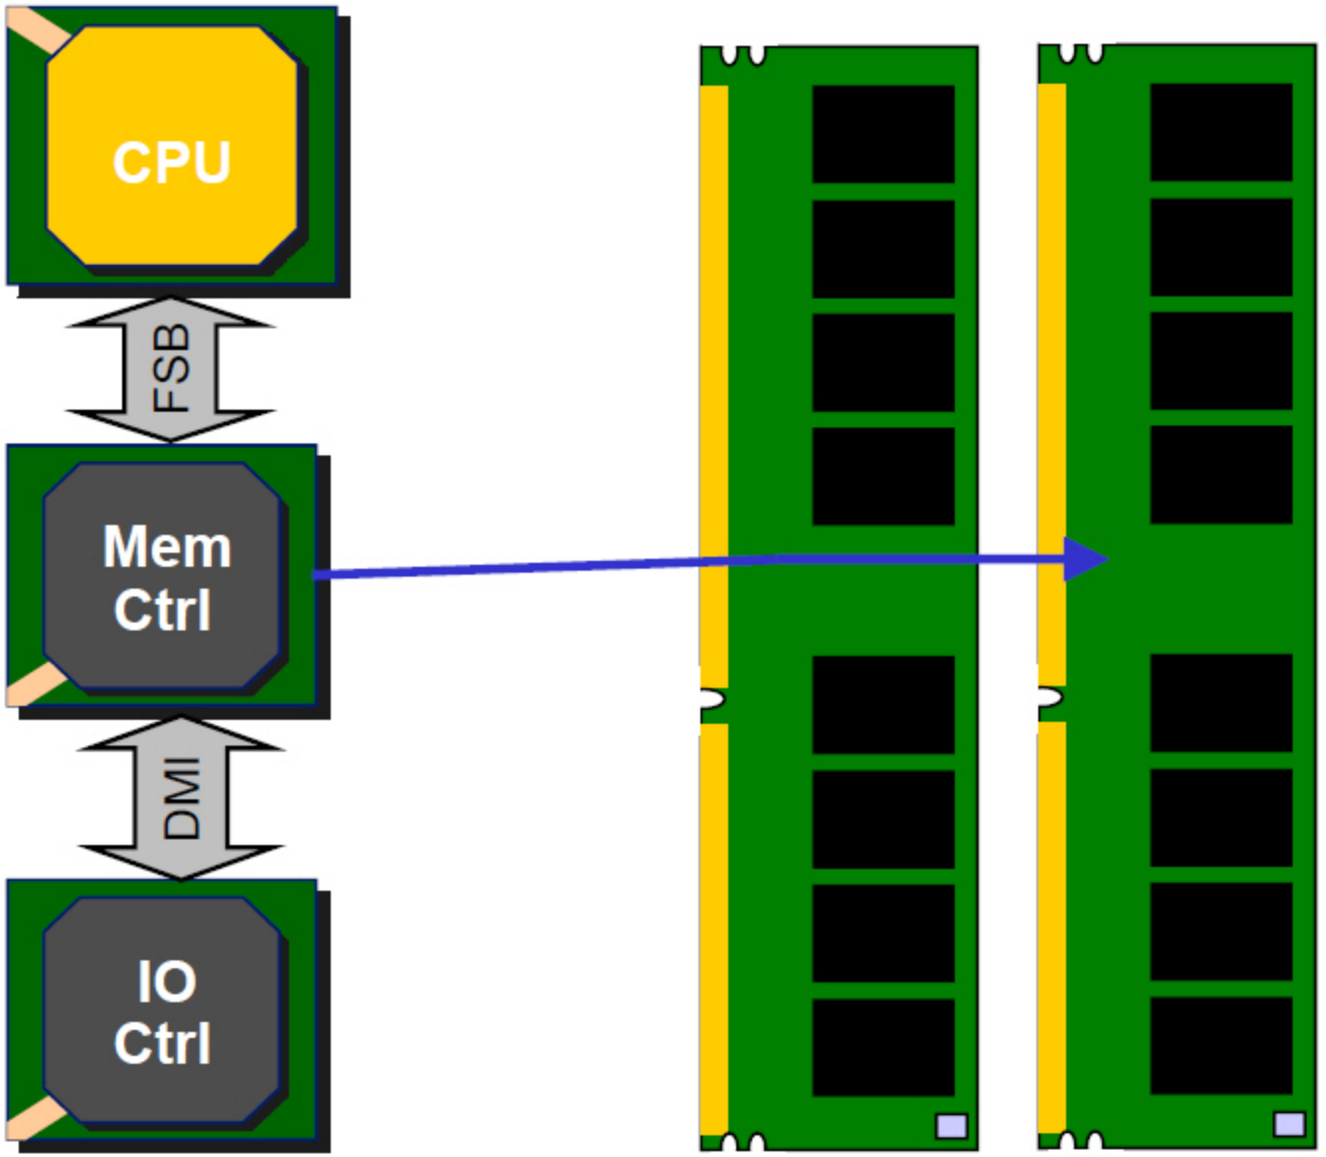

- Read Cycle
- Read command is sent to the DIMM over the address and command bus.

Key: Red is Data Bus
Blue is Address and Command Bus

DDR3 Read Overview

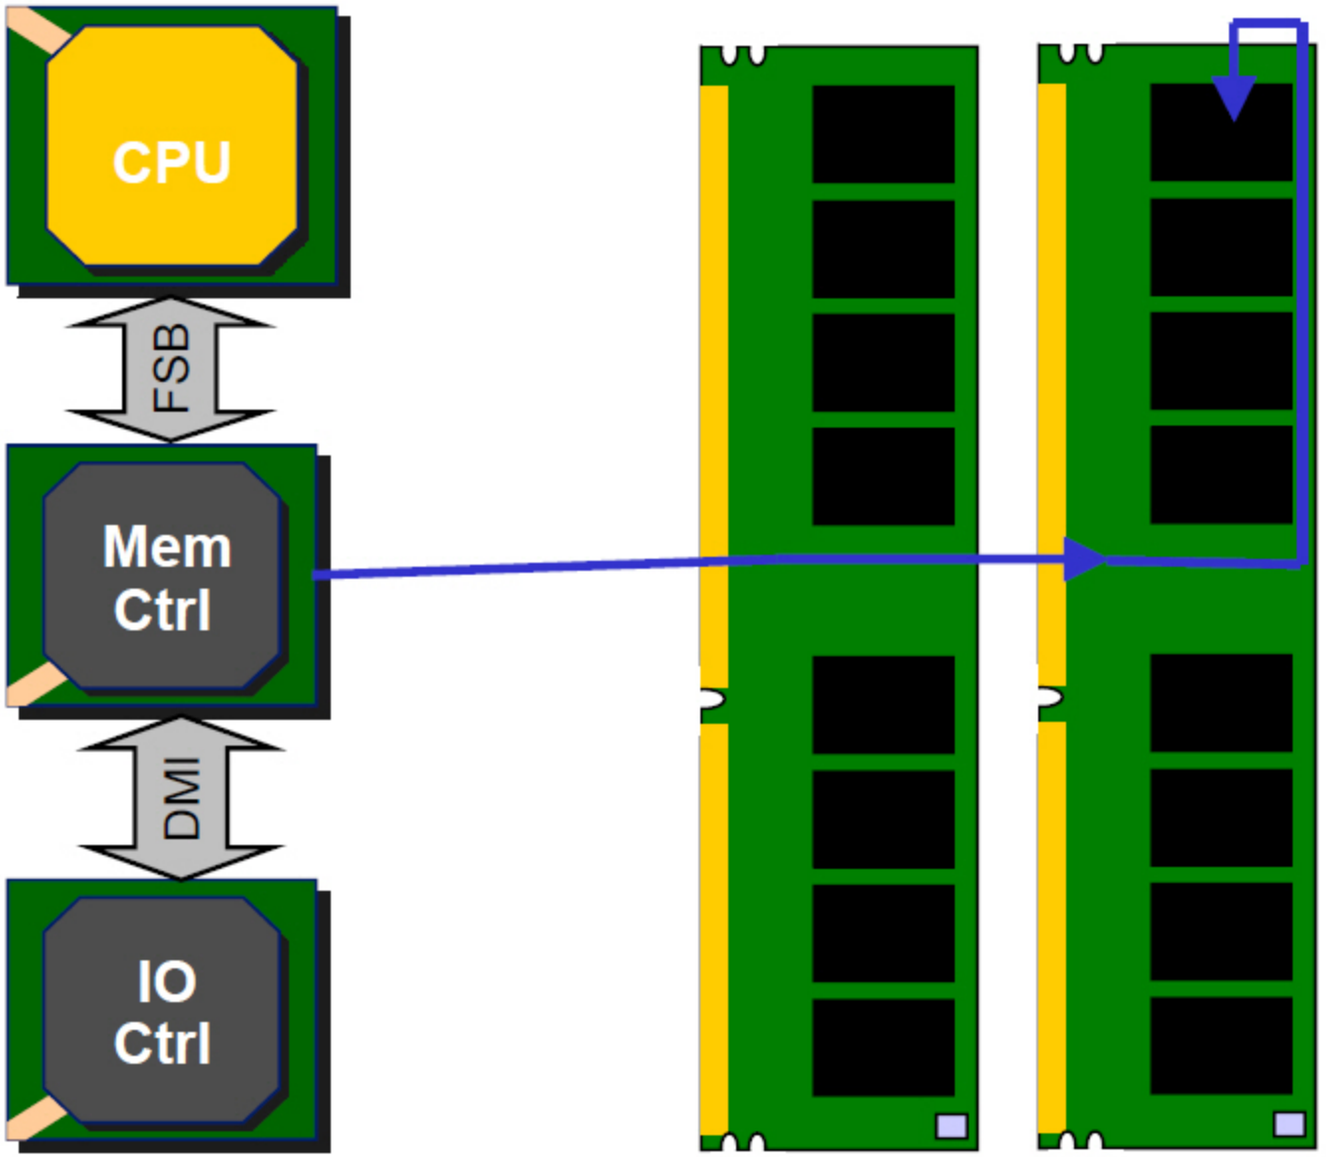

- Read Cycle
- Read command is sent to the DIMM over the address and command bus.
- The command gets to the first DRAM on the new fly-by routing chain.

Key: Red is Data Bus
Blue is Address and Command Bus

DDR3 Read Overview

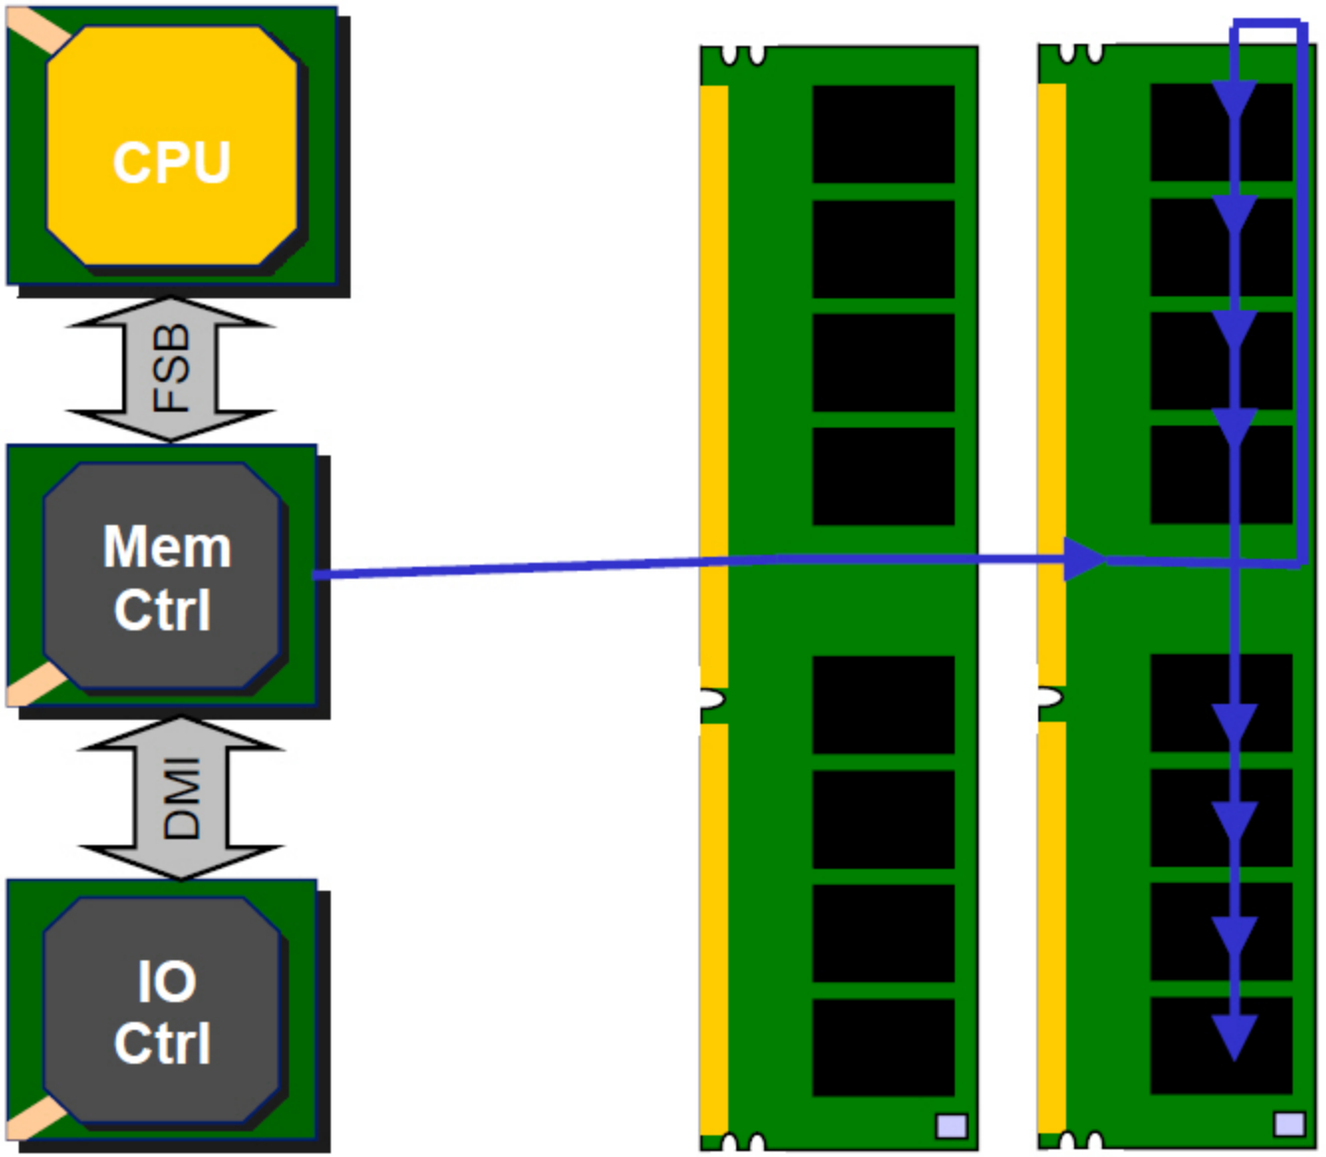

- Read Cycle
- Read command is sent to the DIMM over the address and command bus.
- The command gets to the first DRAM on the new fly-by routing chain.
- The command propagates to the remaining DRAMs

Key: Red is Data Bus
Blue is Address and Command Bus

DDR3 Read Overview

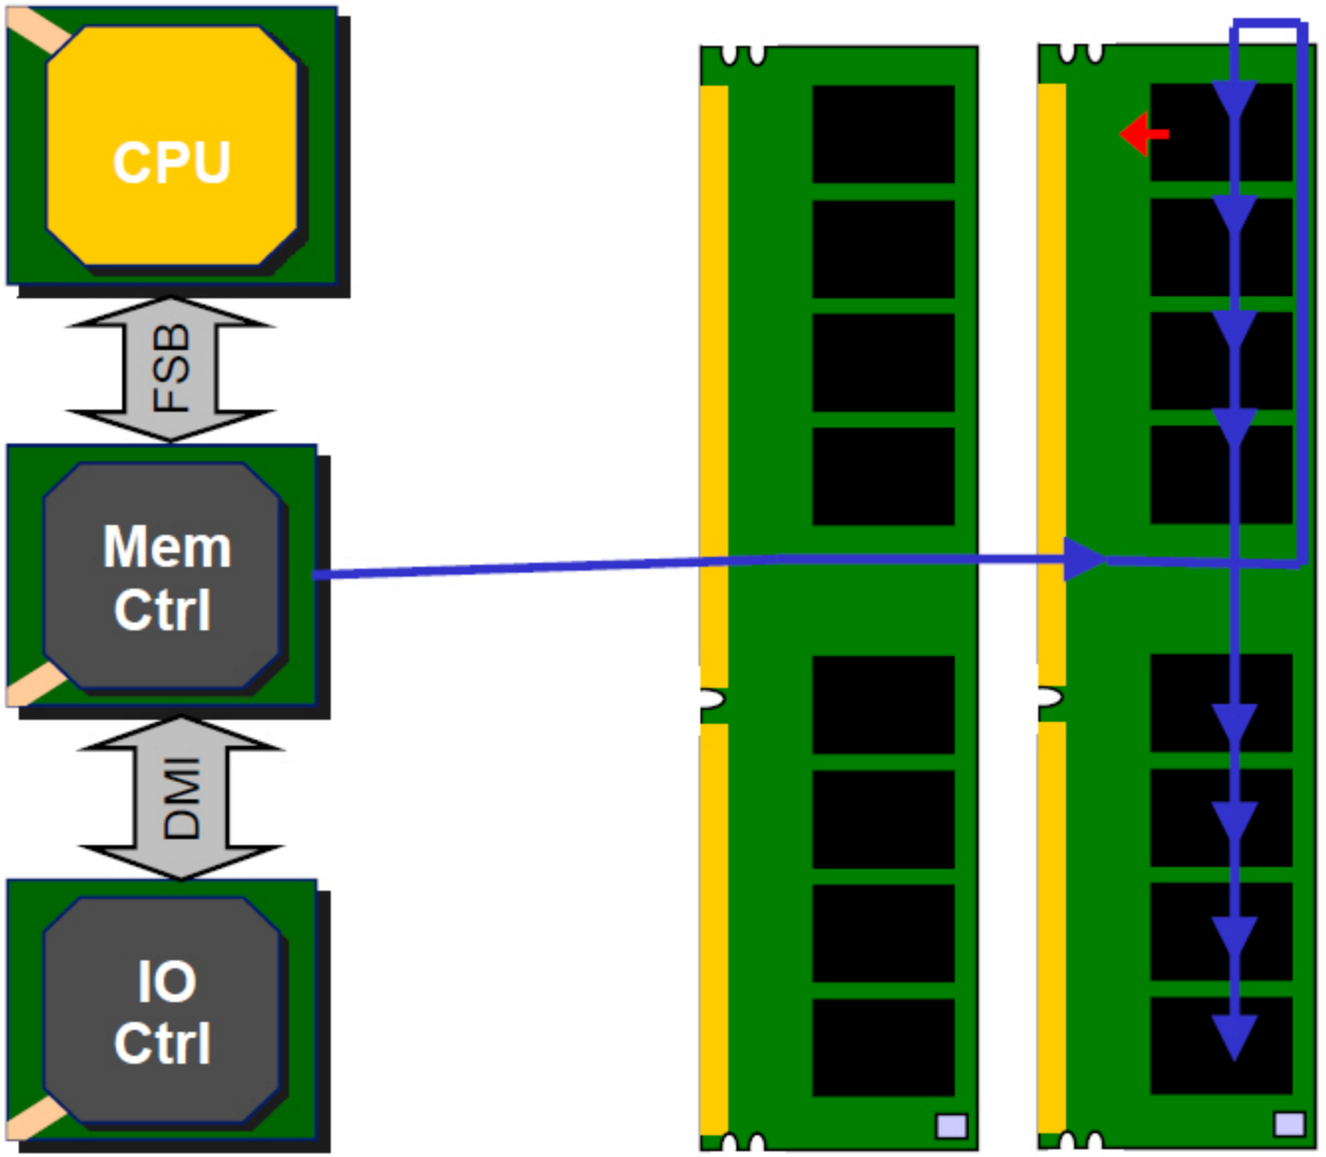

- Read Cycle
- Read command is sent to the DIMM over the address and command bus.
- The command gets to the first DRAM on the new fly-by routing chain.
- The command propagates to the remaining DRAMs
- After RL, data begins streaming back from the first DRAM.

Key: Red is Data Bus
Blue is Address and Command Bus

DDR3 Read Overview

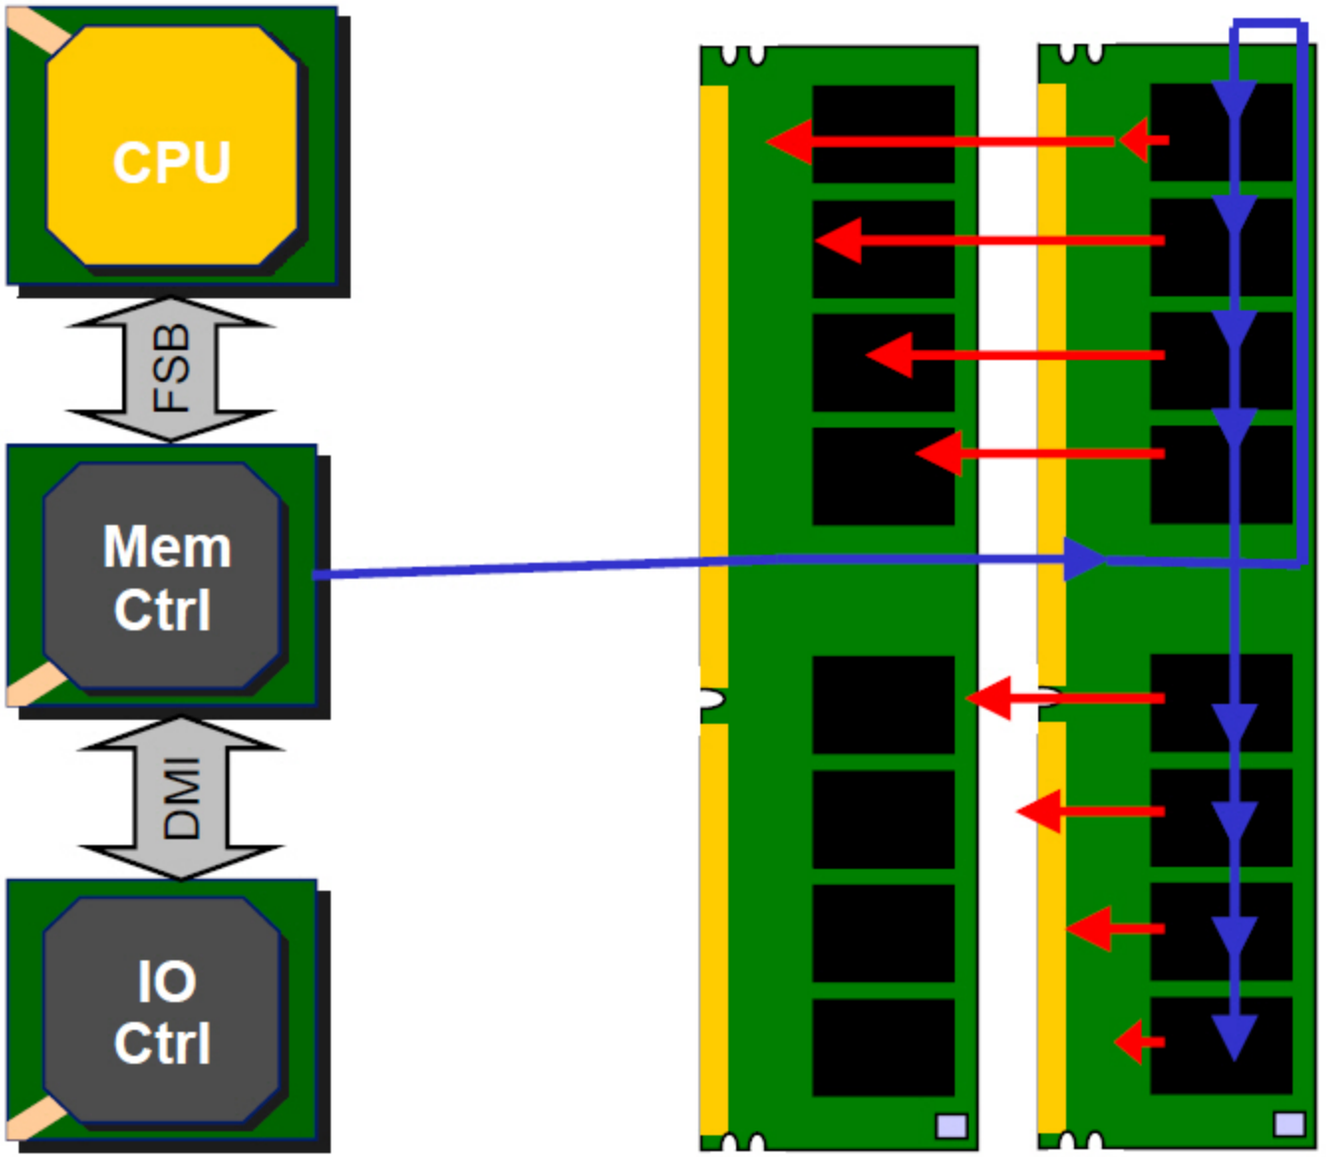

- **Read Cycle cont.**
- Data propagates from the remaining DRAMs in order that the command was received.

Key: Red is Data Bus
Blue is Address and Command Bus

DDR3 Read Overview

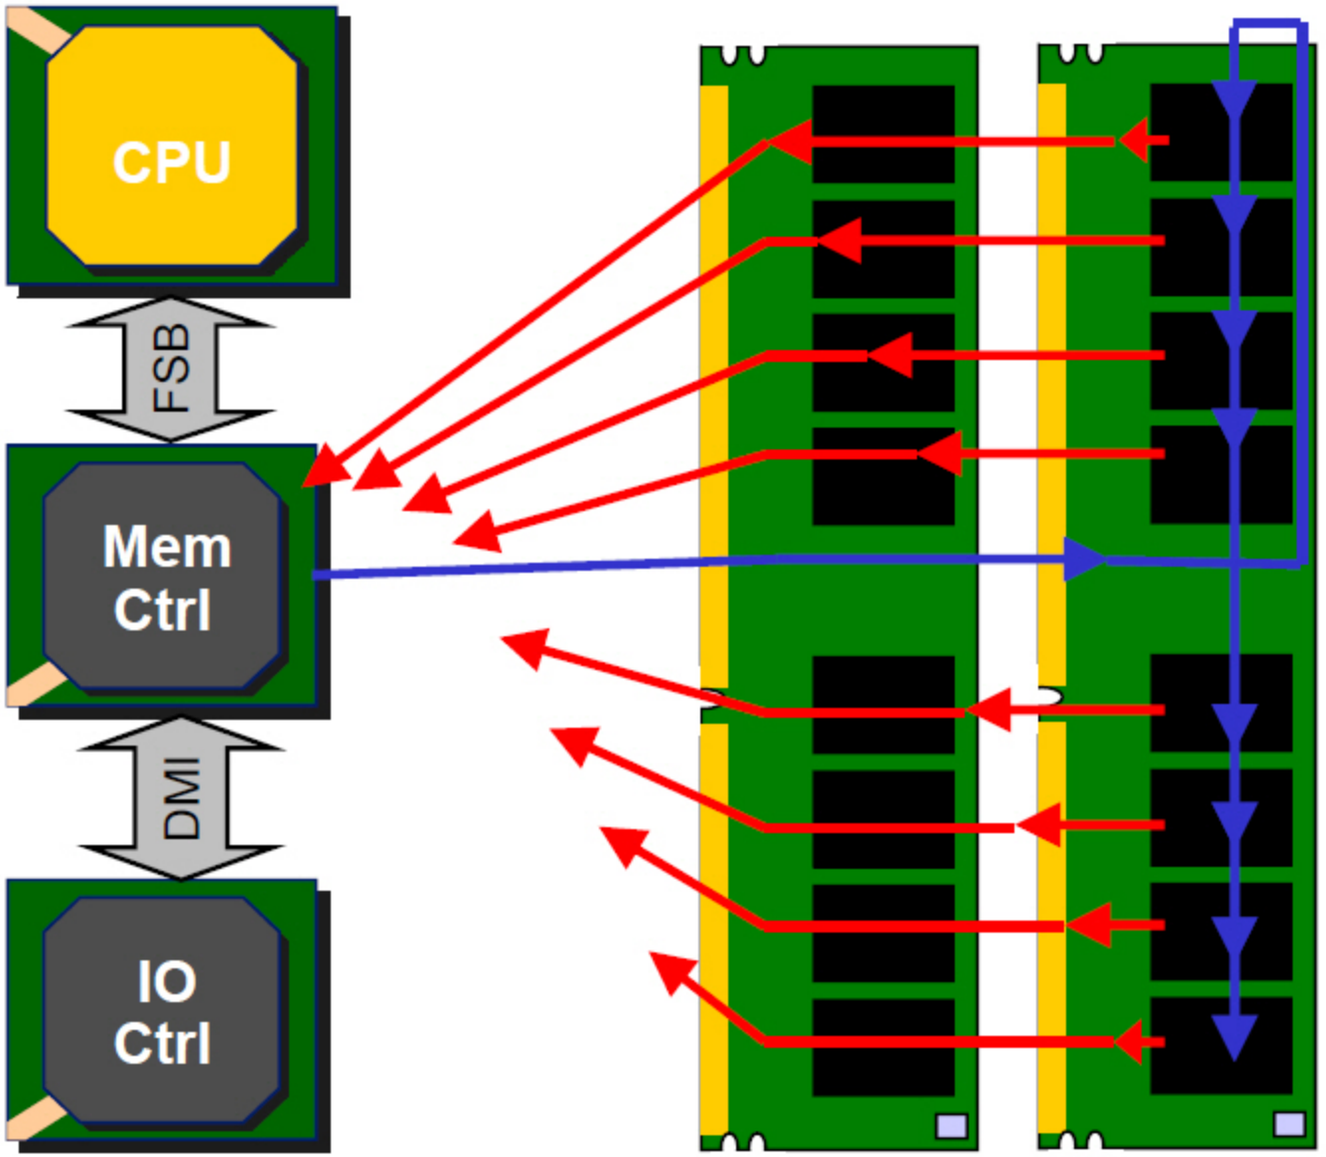

- **Read Cycle cont.**
- Data propagates from the remaining DRAMs in order that the command was received.
- Read data from the first byte lane arrives.
- The remaining data lanes arrive one after the other.

Key: Red is Data Bus
Blue is Address and Command Bus

DDR3 Read Overview

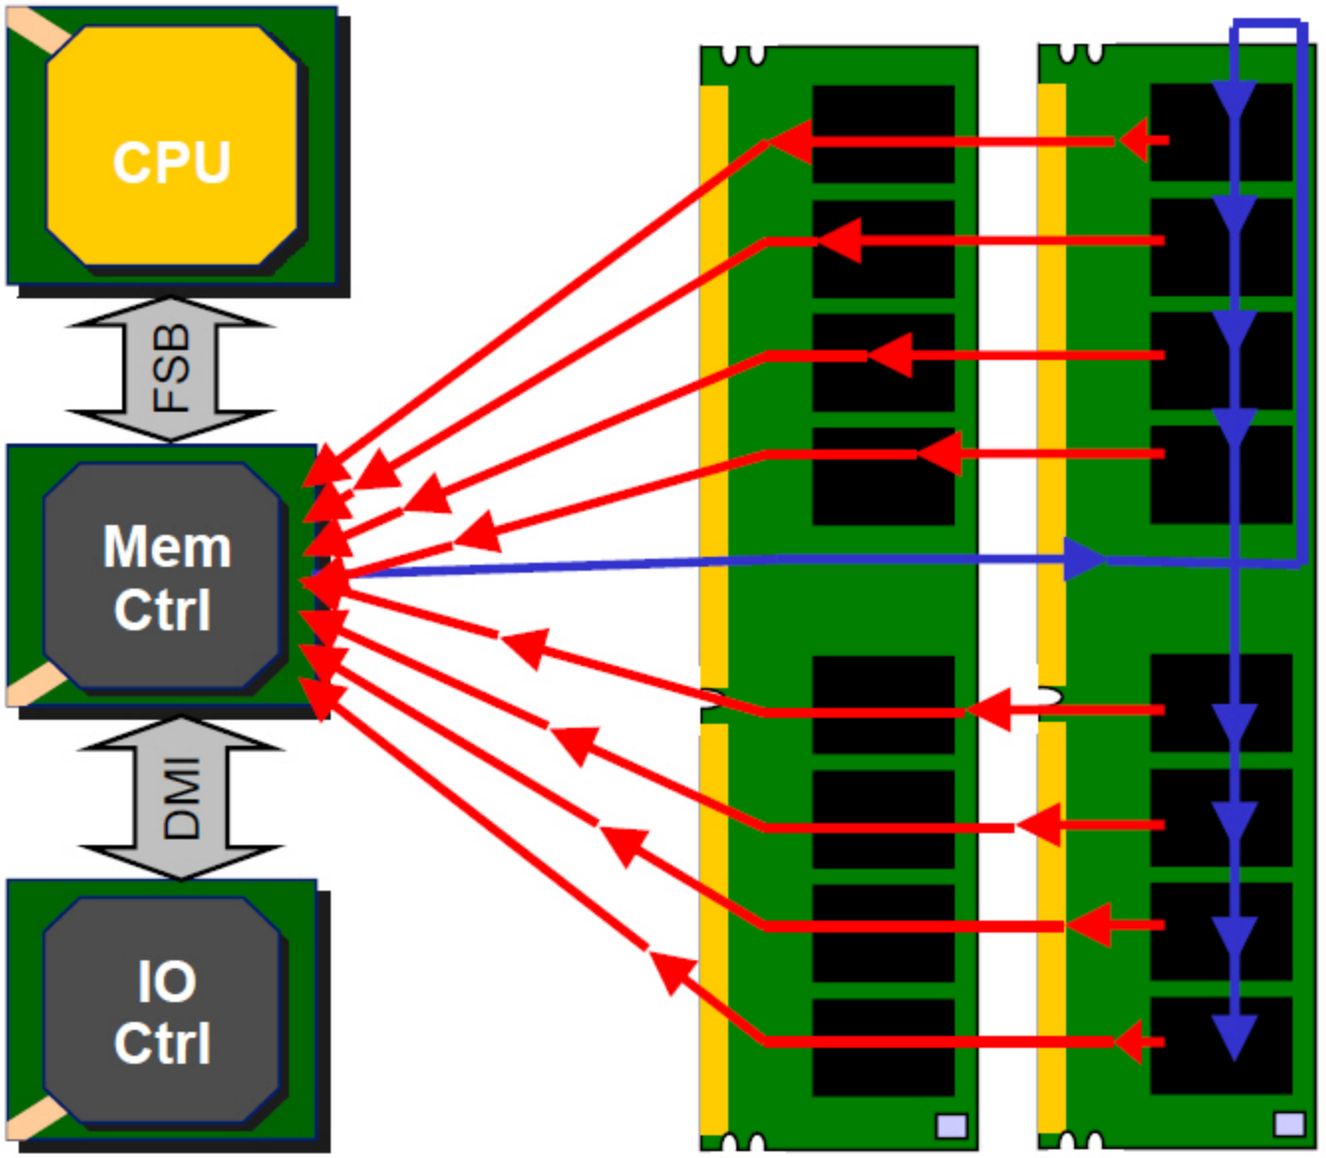

Key: Red is Data Bus
Blue is Address and Command Bus

- **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?

MINDSHARE

BRINGING LIFE TO KNOWLEDGE

Read Calibration

Read Calibration

- 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.

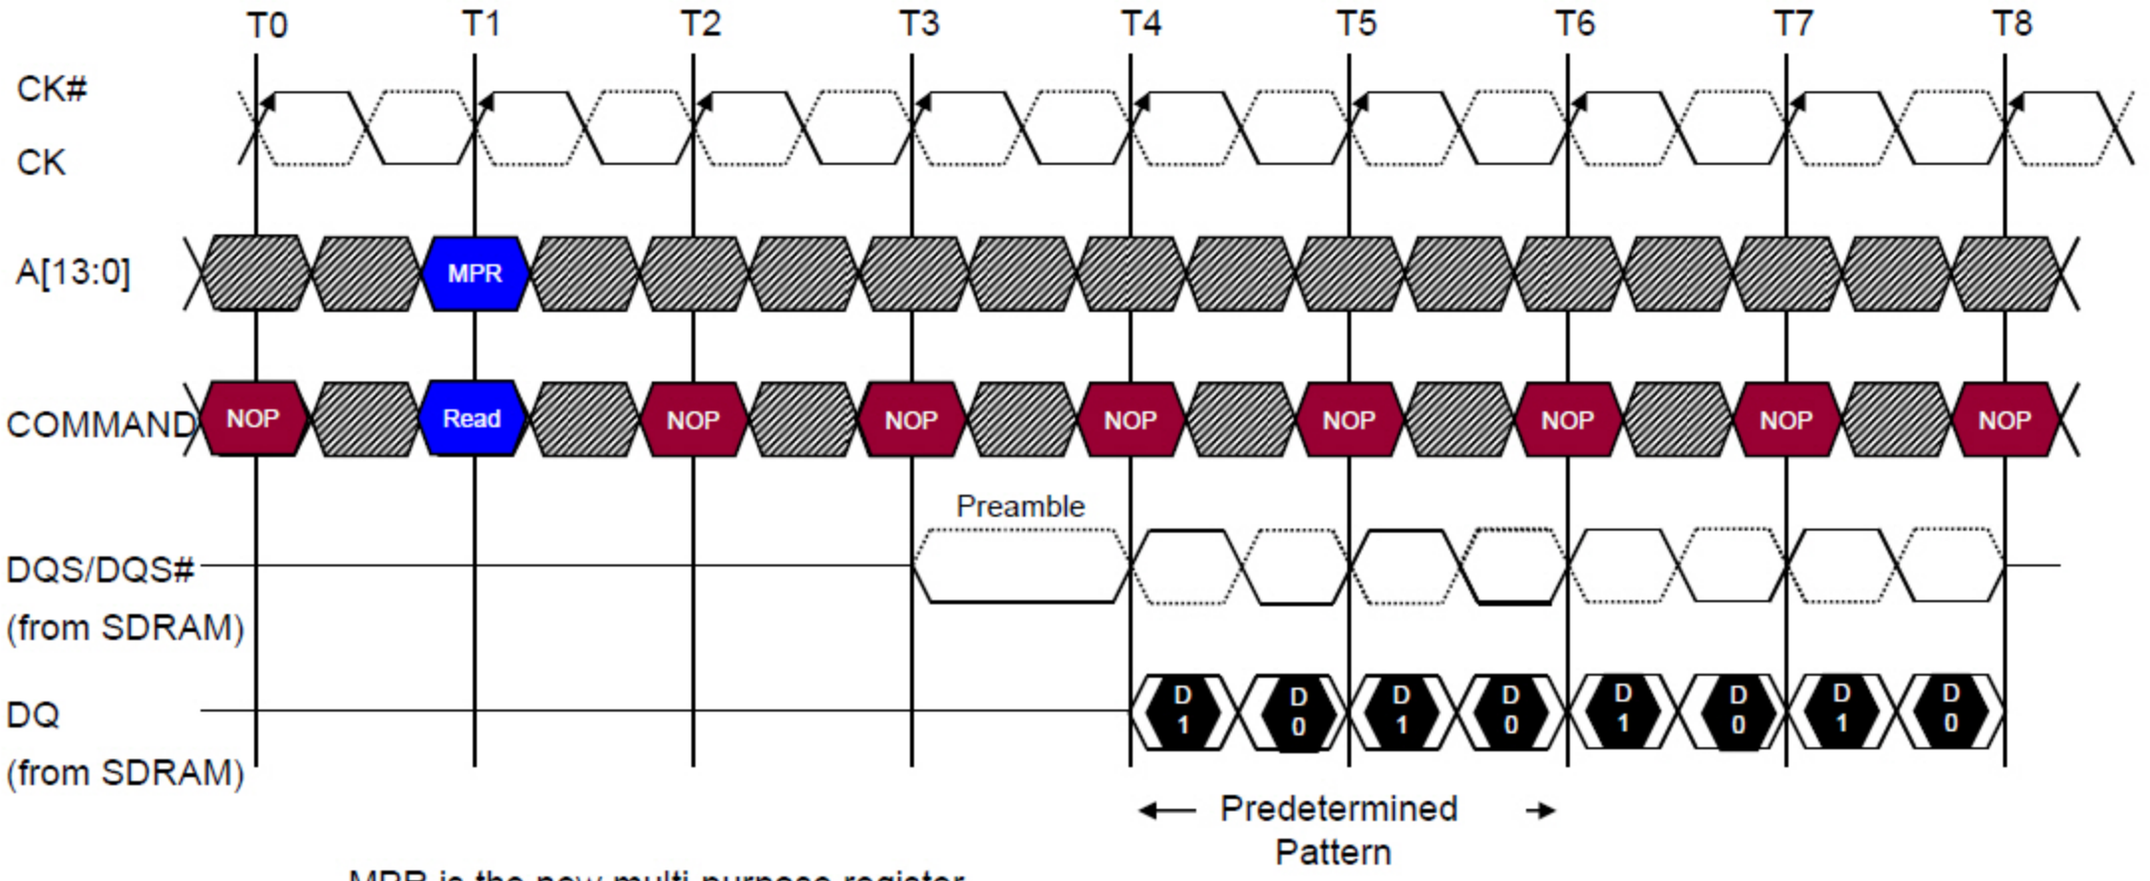

MPR is the new multi-purpose register

MINDSHARE

BRINGING LIFE TO KNOWLEDGE

Fly-by Routing Write Example

DDR3 Write Overview

- 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

DDR3 Write Overview

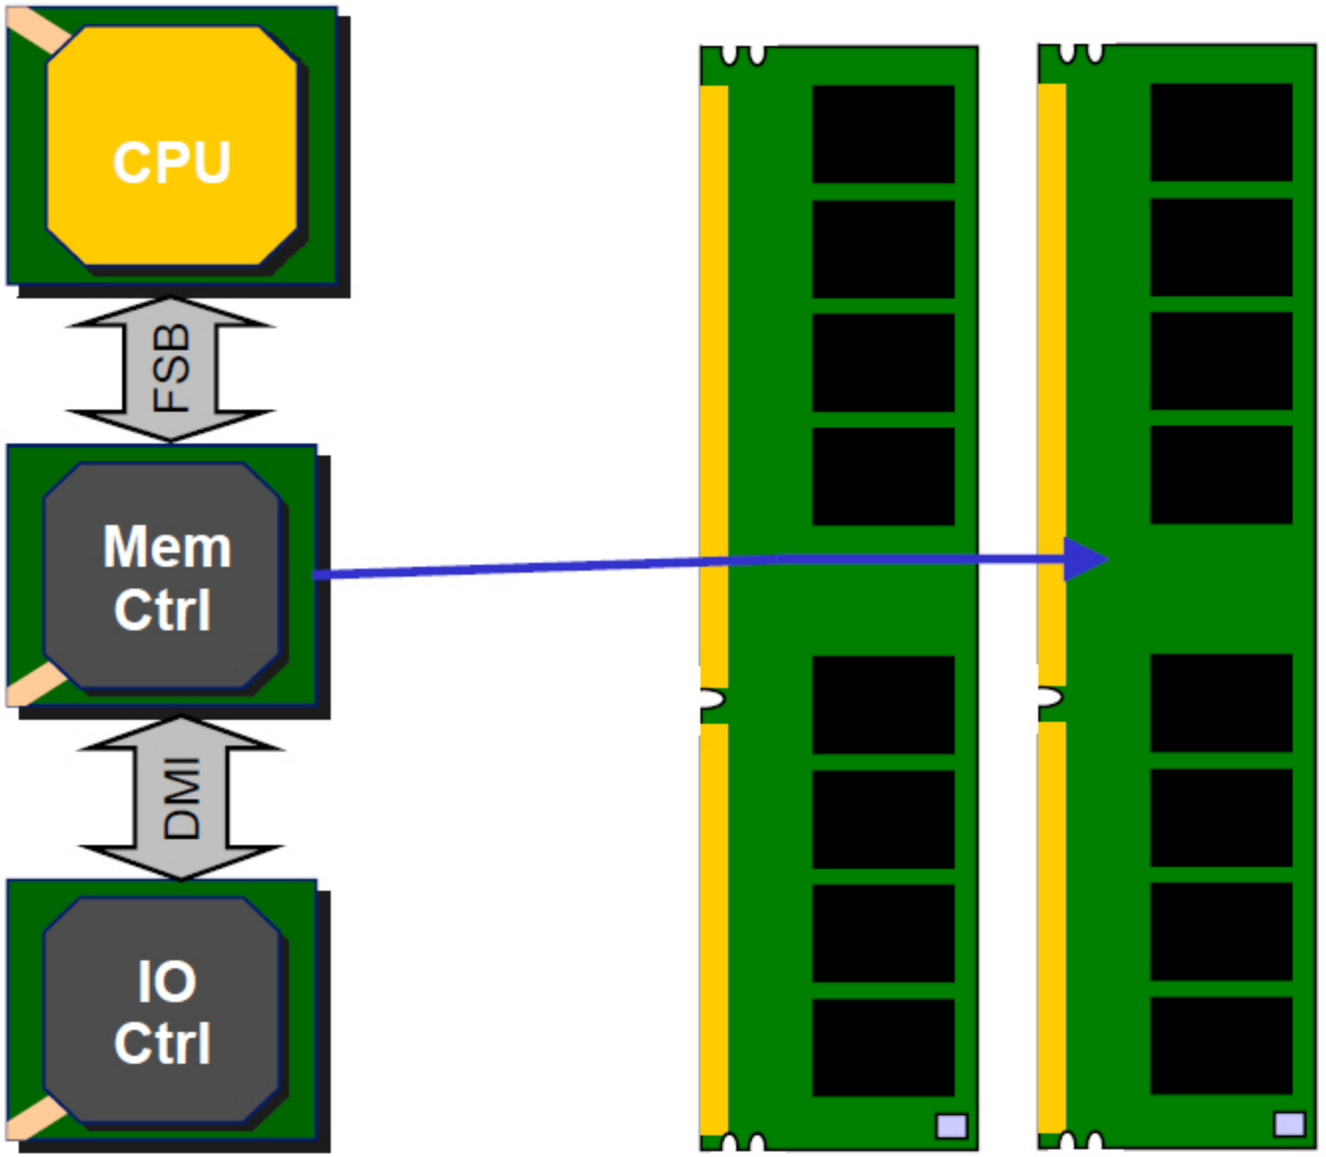

- **Write Cycle**
- Write command is sent to the DIMM over the address and command bus.

Key: Red is Data Bus
Blue is Address and Command Bus

DDR3 Write Overview

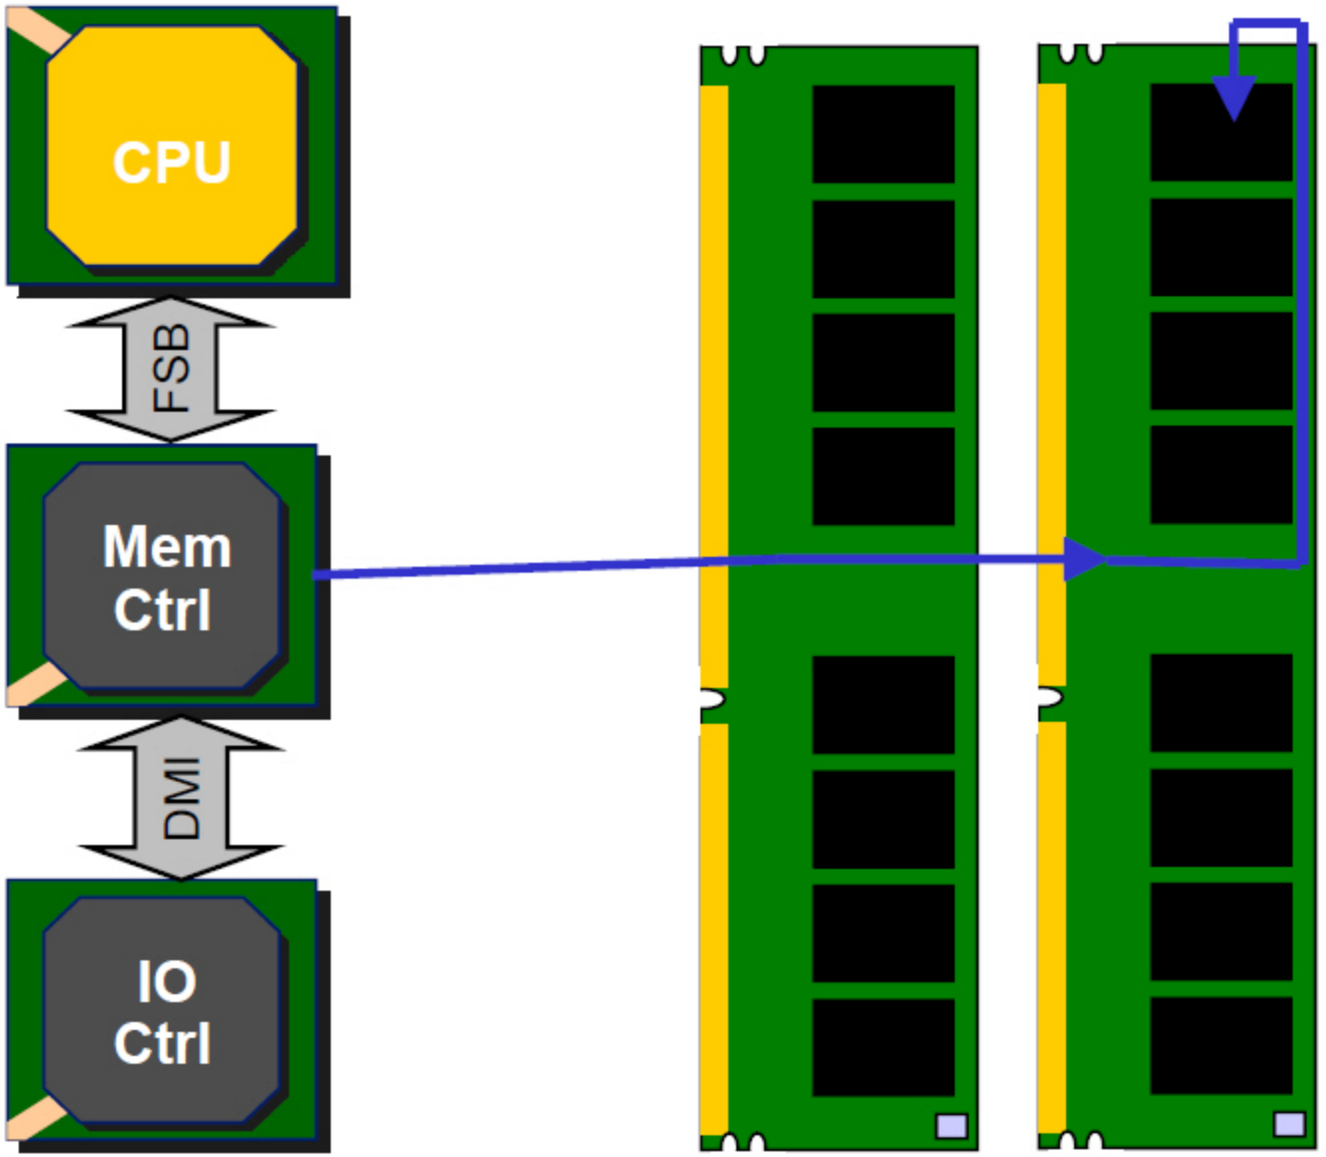

- **Write Cycle**
- Write command is sent to the DIMM over the address and command bus.
- Command reaches the first DRAM due the fly-by routing.

Key: Red is Data Bus
Blue is Address and Command Bus

DDR3 Write Overview

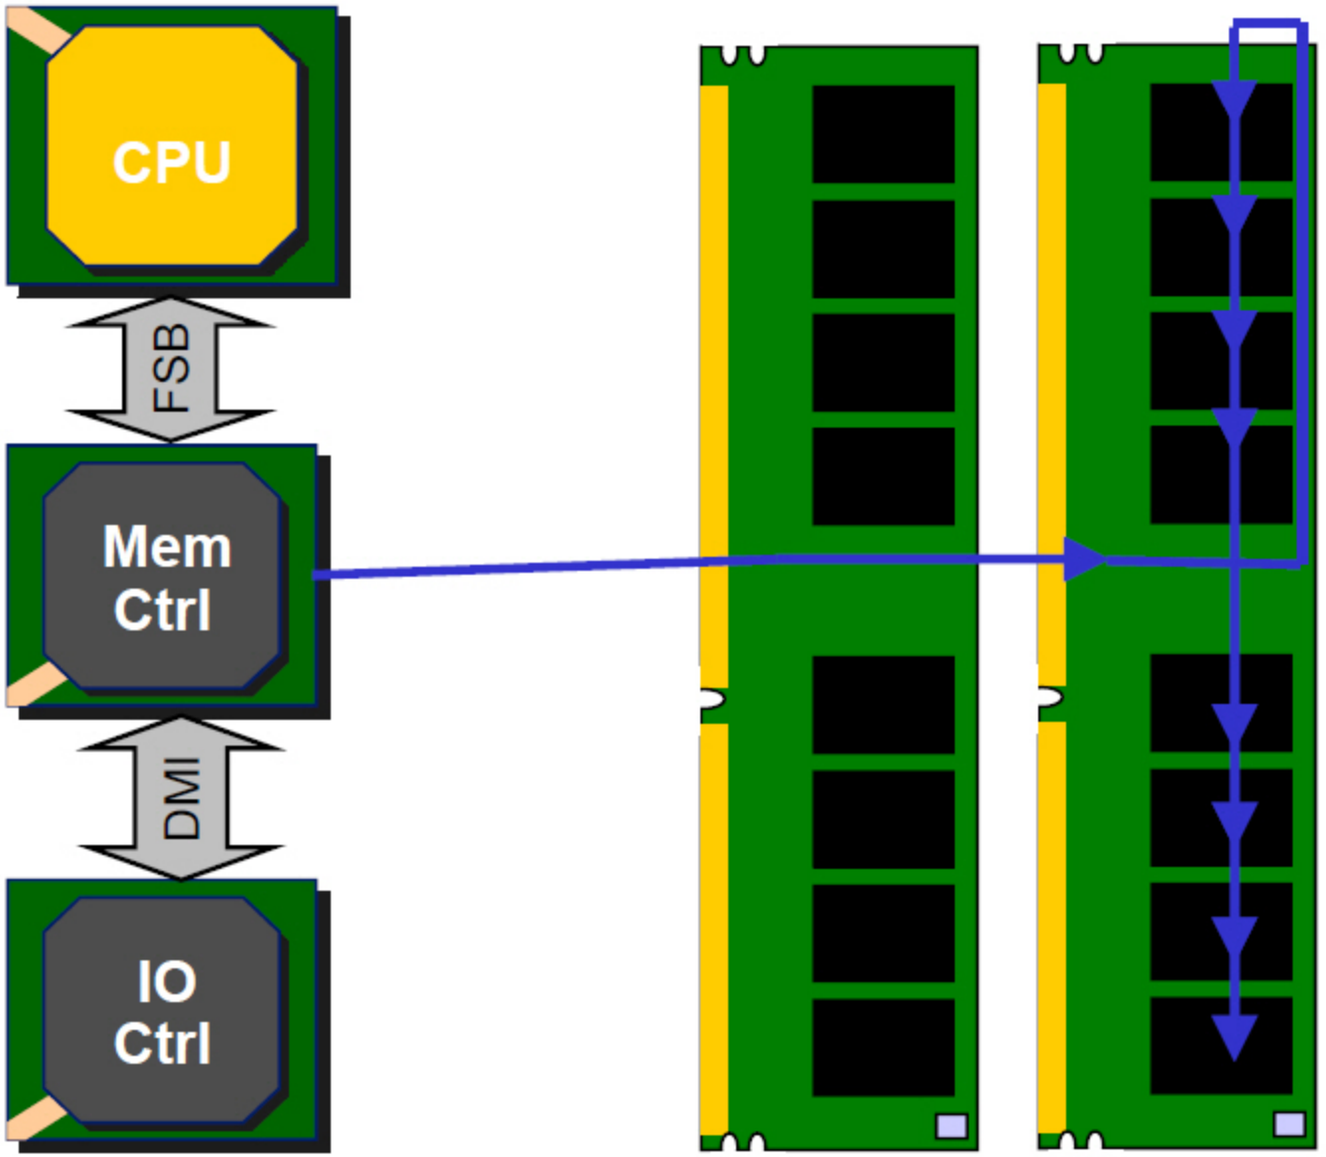

- **Write Cycle**
- Write command is sent to the DIMM over the address and command bus.
- Command reaches the first DRAM due the fly-by routing.
- Command propagates to the remaining DRAMs.

Key: Red is Data Bus
Blue is Address and Command Bus

DDR3 Write Overview

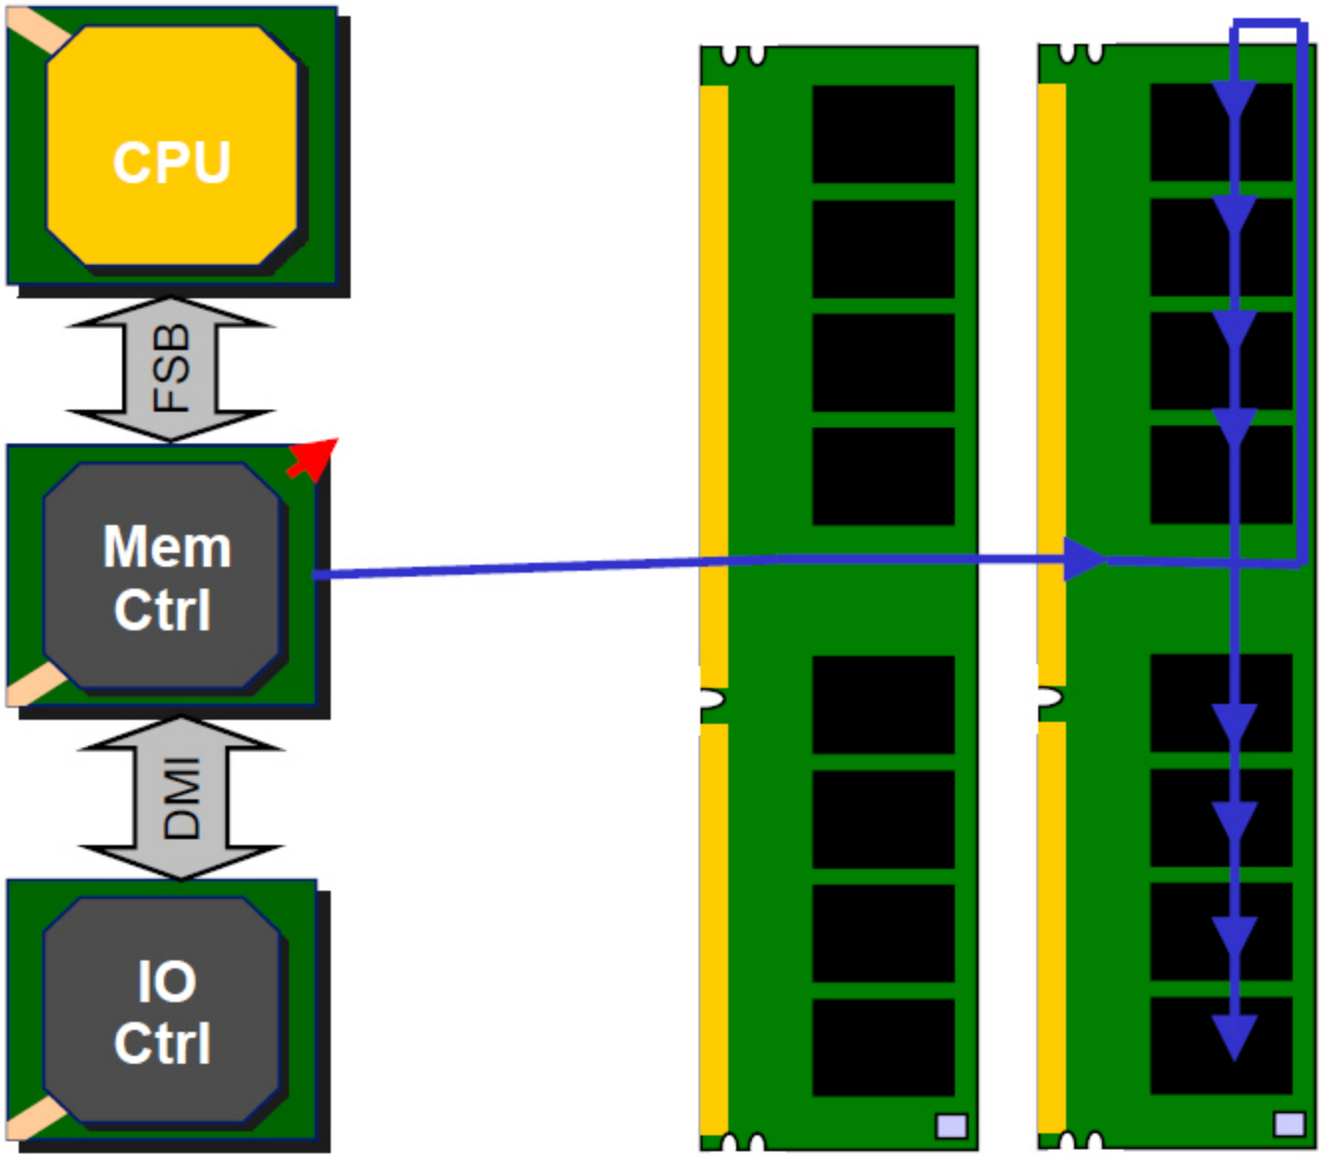

- **Write Cycle**
- Write command is sent to the DIMM over the address and command bus.
- Command reaches the first DRAM due the fly-by routing.
- Command propagates to the remaining DRAMs.
- After WL, the controller starts driving the first data lane.

Key: Red is Data Bus
Blue is Address and Command Bus

DDR3 Write Overview

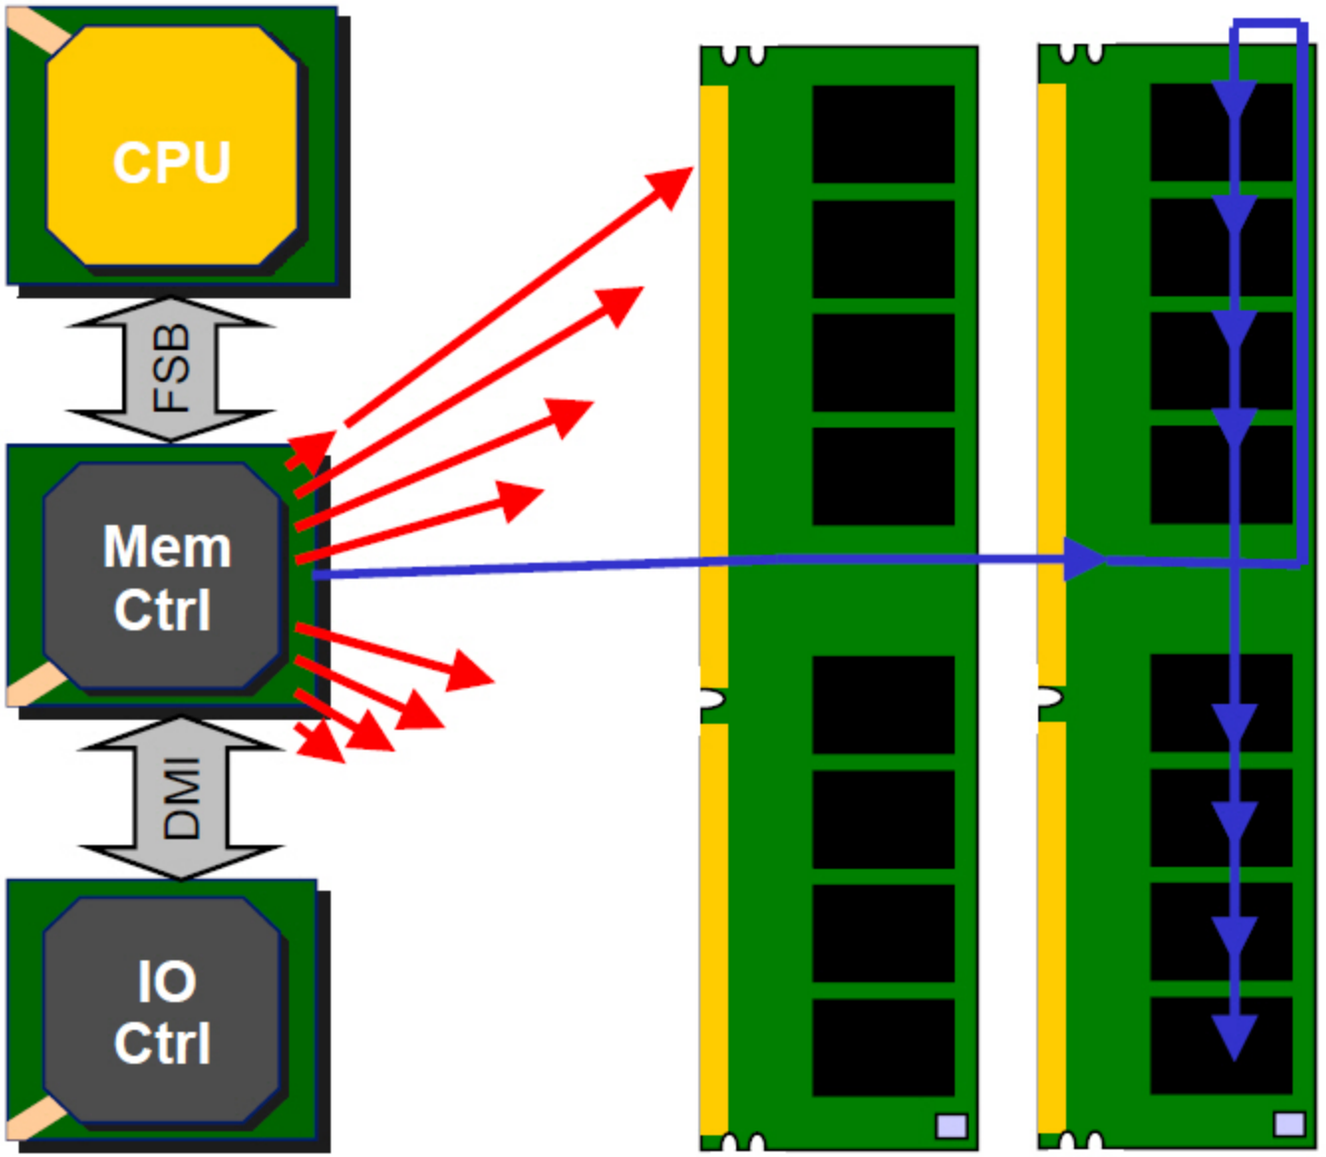

- **Write Cycle**
- Write command is sent to the DIMM over the address and command bus.
- Command reaches the first DRAM due the fly-by routing.
- Command propagates to the remaining DRAMs.
- After WL, the controller starts driving the first data lane.
- The other data lanes follow in order.

Key: Red is Data Bus
Blue is Address and Command Bus

DDR3 Write Overview

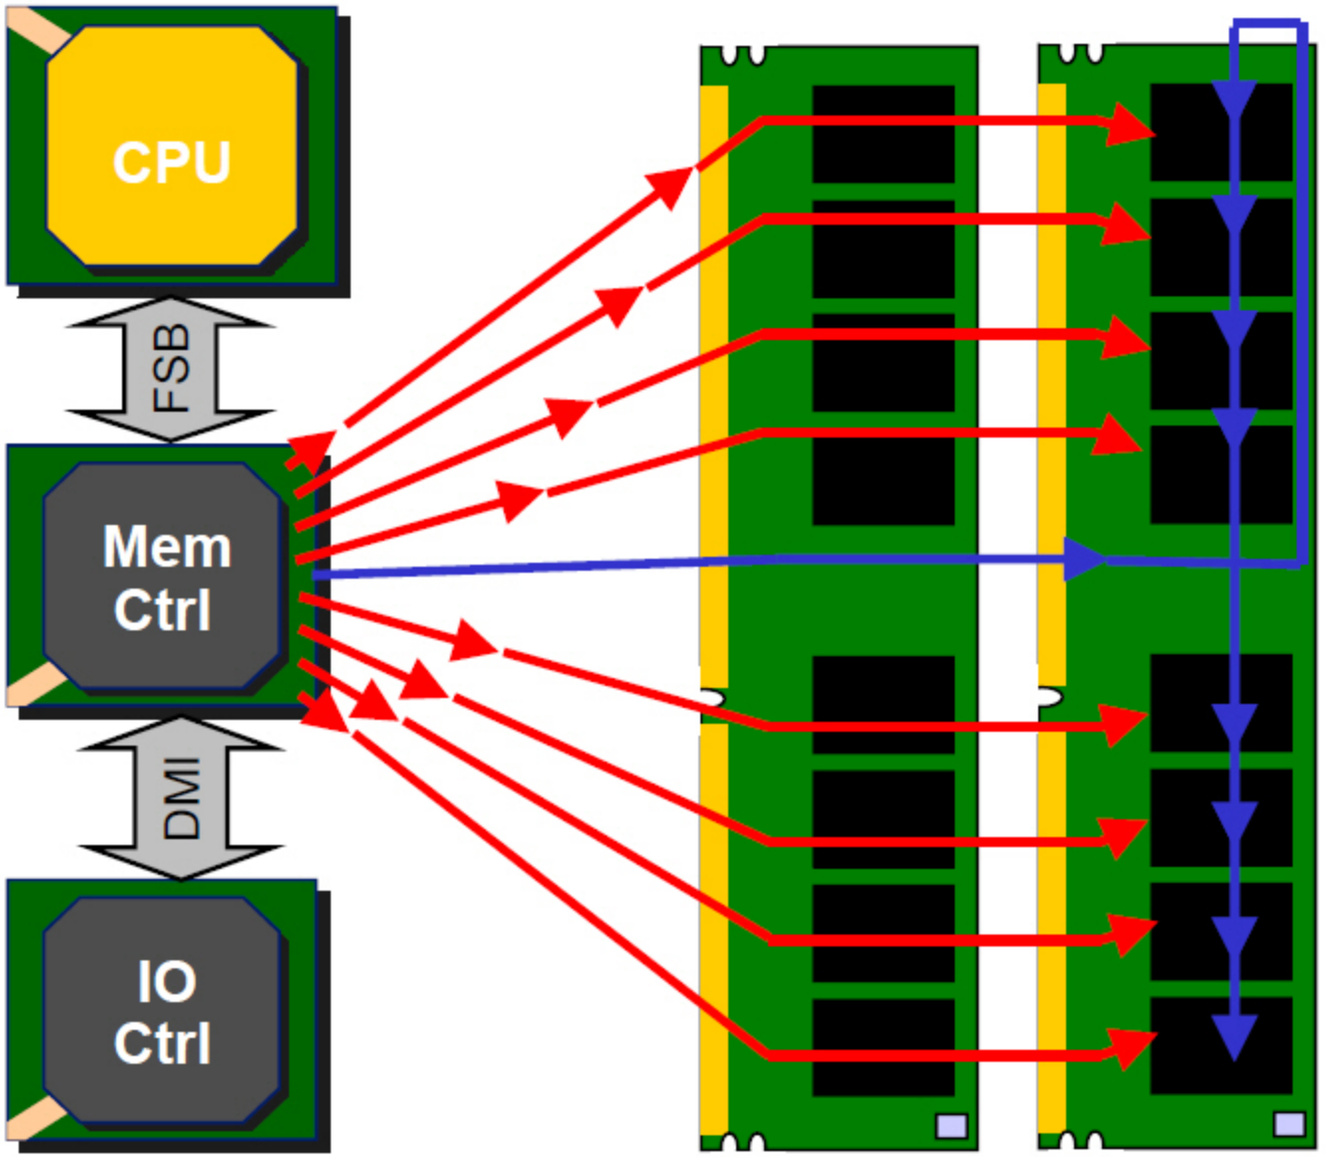

Key: Red is Data Bus
Blue is Address and Command Bus

- **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.
- 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.

DDR3 Write Leveling

- 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?

MINDSHARE

BRINGING LIFE TO KNOWLEDGE

Write Leveling

DDR3 Write Leveling

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"

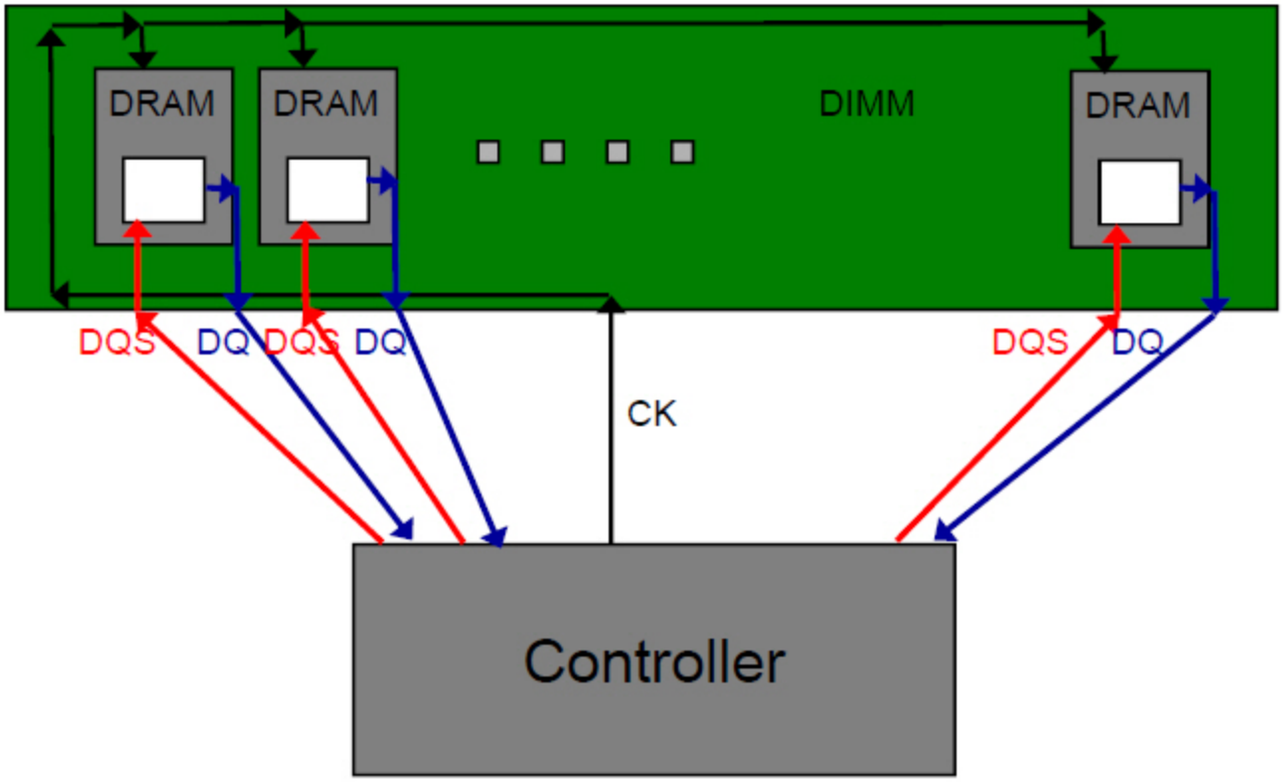

DDR3 Write Leveling

- 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.

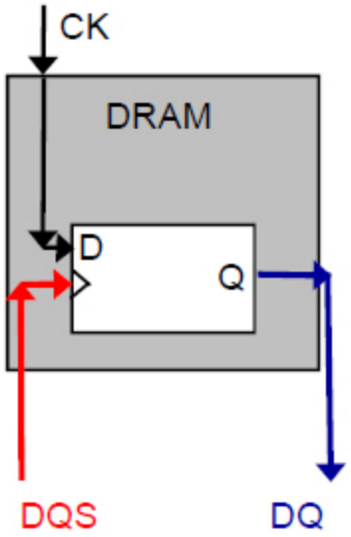

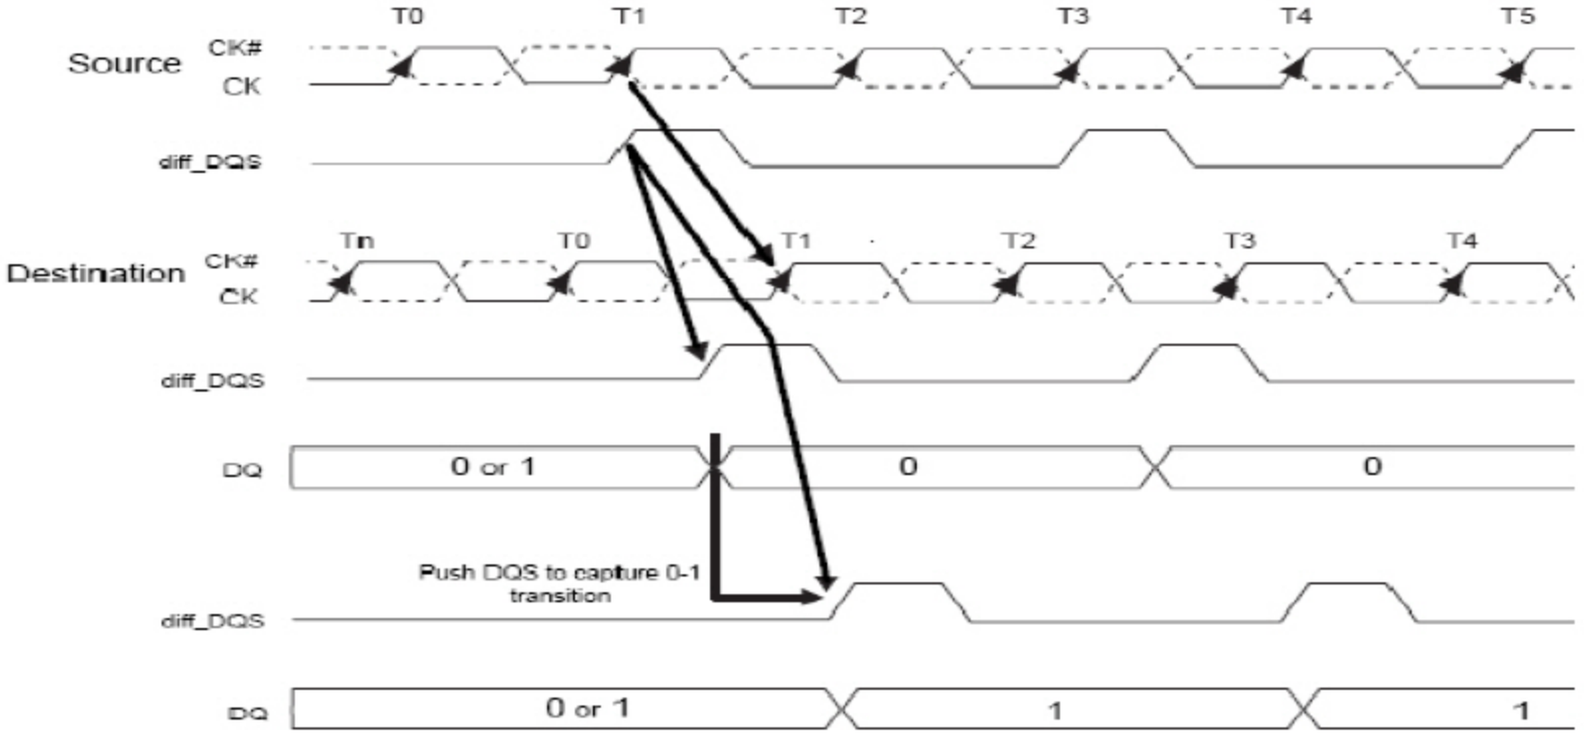

DDR3 Write Leveling

- 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

DDR3 ODT

- 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

ODT Circuitry

- 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 Ron of the Pull Up and Pull Down transistors.

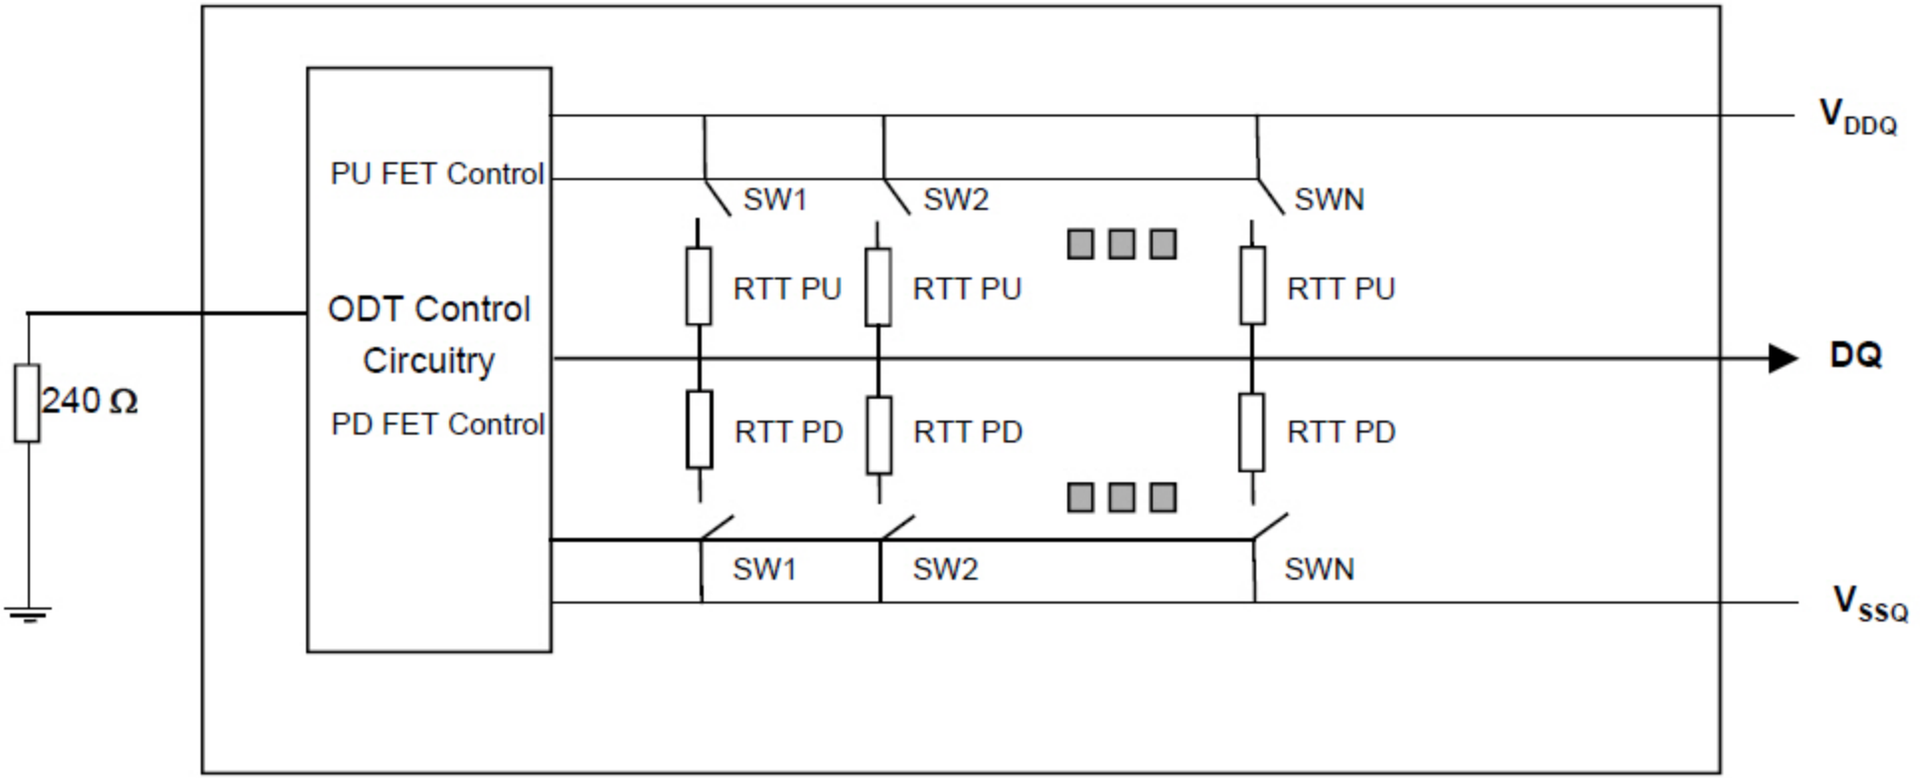

ODT Changes

- 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).

ODT Write Example Synchronous Mode

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

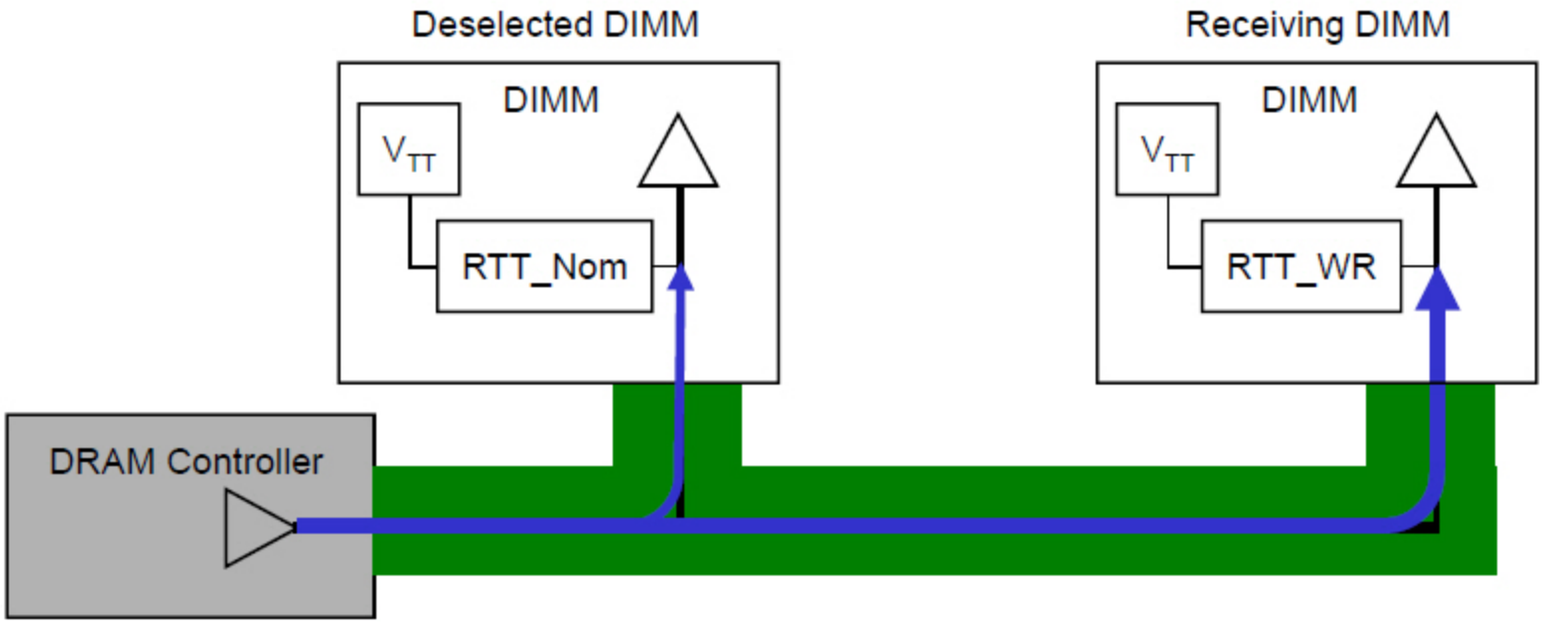

ODT Changes

- 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

ODT Register Settings

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)

ODT Register Settings

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.

MINDSHARE

BRINGING LIFE TO KNOWLEDGE

ZQ Calibration (OCD)

ZQ Calibration

- 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

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.

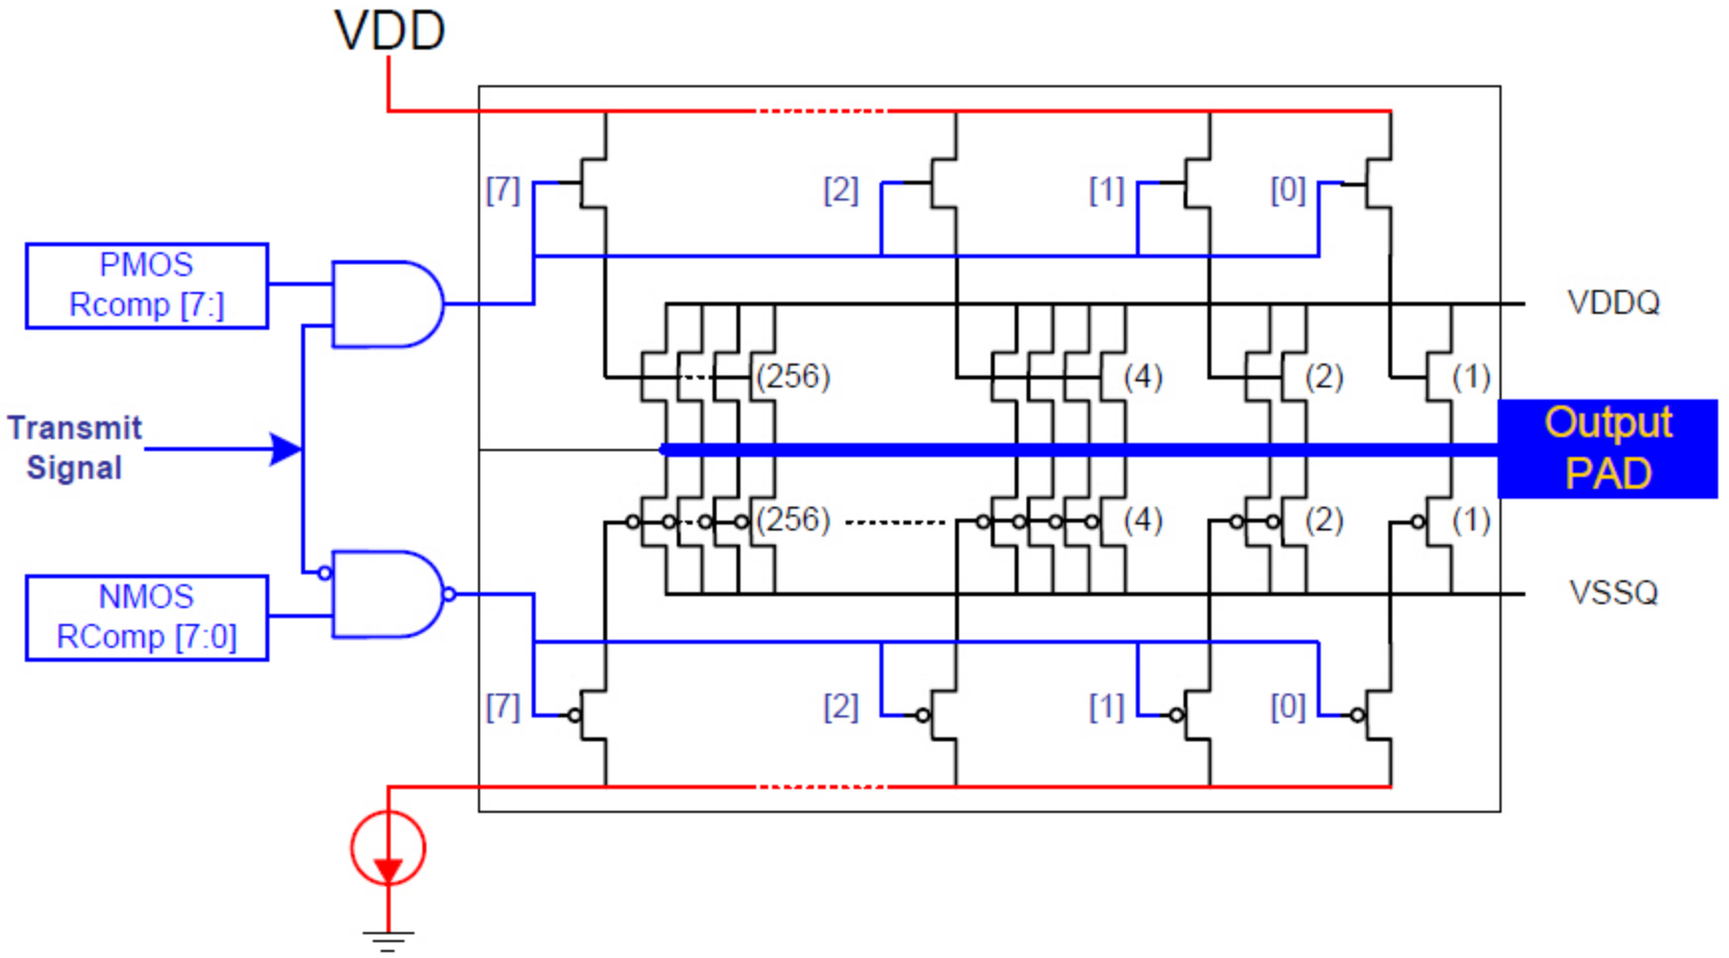

MINDSHARE

BRINGING LIFE TO KNOWLEDGE

Reset

Asynchronous Reset

- 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 µs
  - 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.

MINDSHARE

BRINGING LIFE TO KNOWLEDGE

On-DIMM Address Mirroring

On-DIMM Address Mirroring

- What is the motivation?

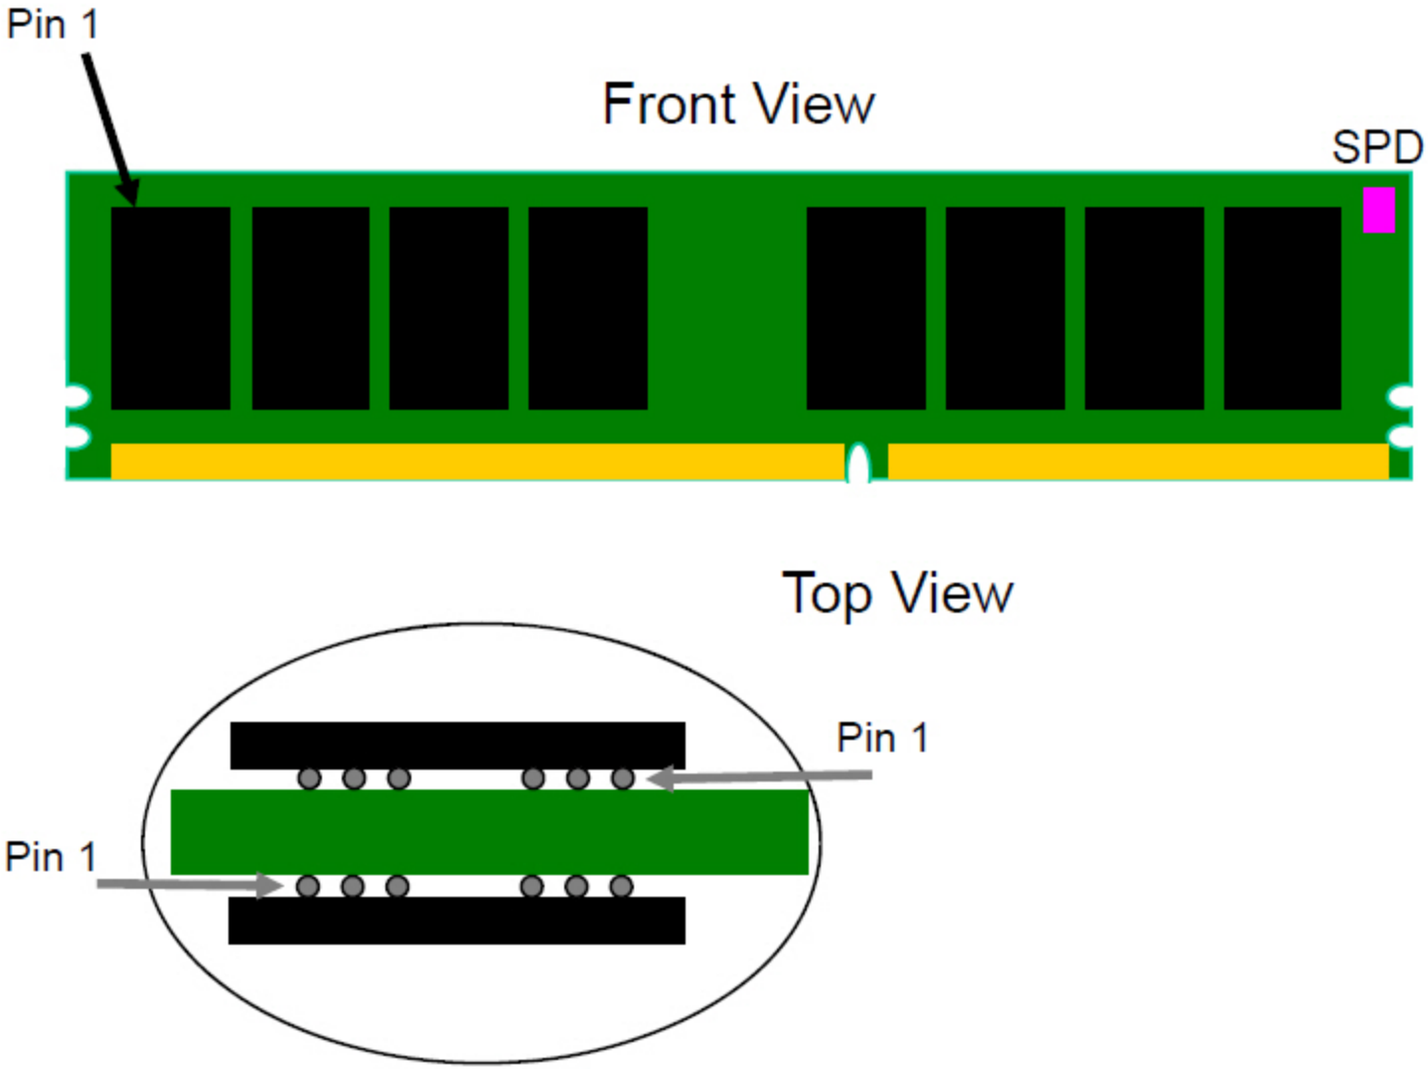

On-DIMM Address Mirroring

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

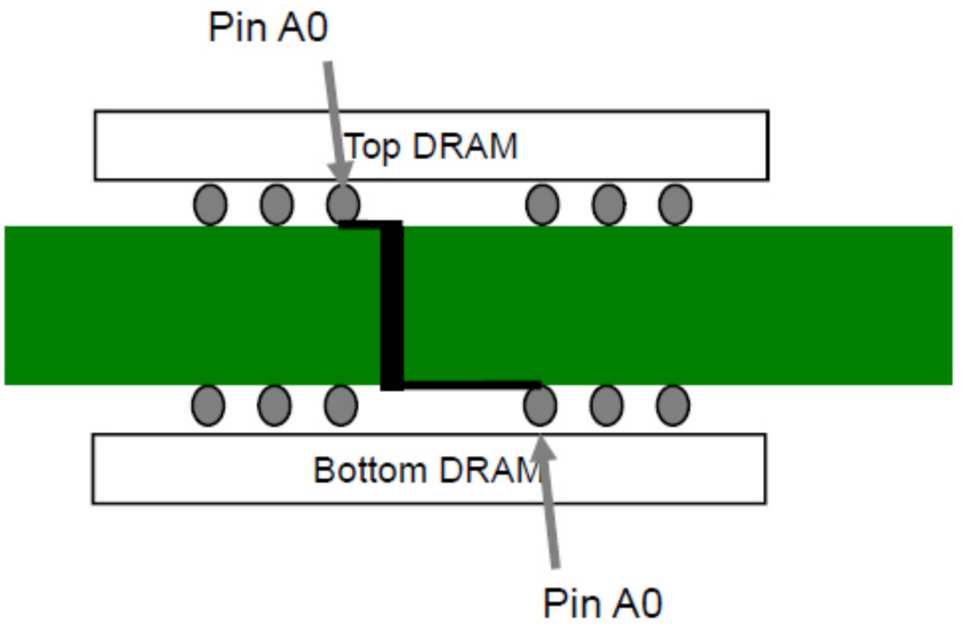

Solution: Mirrored DIMM has simpler routing that decreases stub length

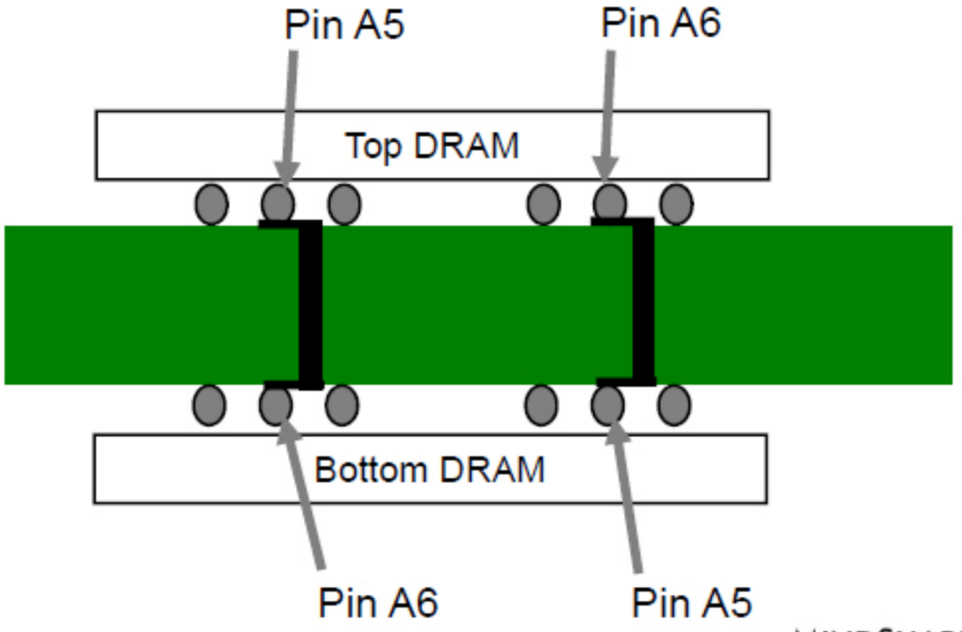

On-DIMM Address Mirroring

- 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.

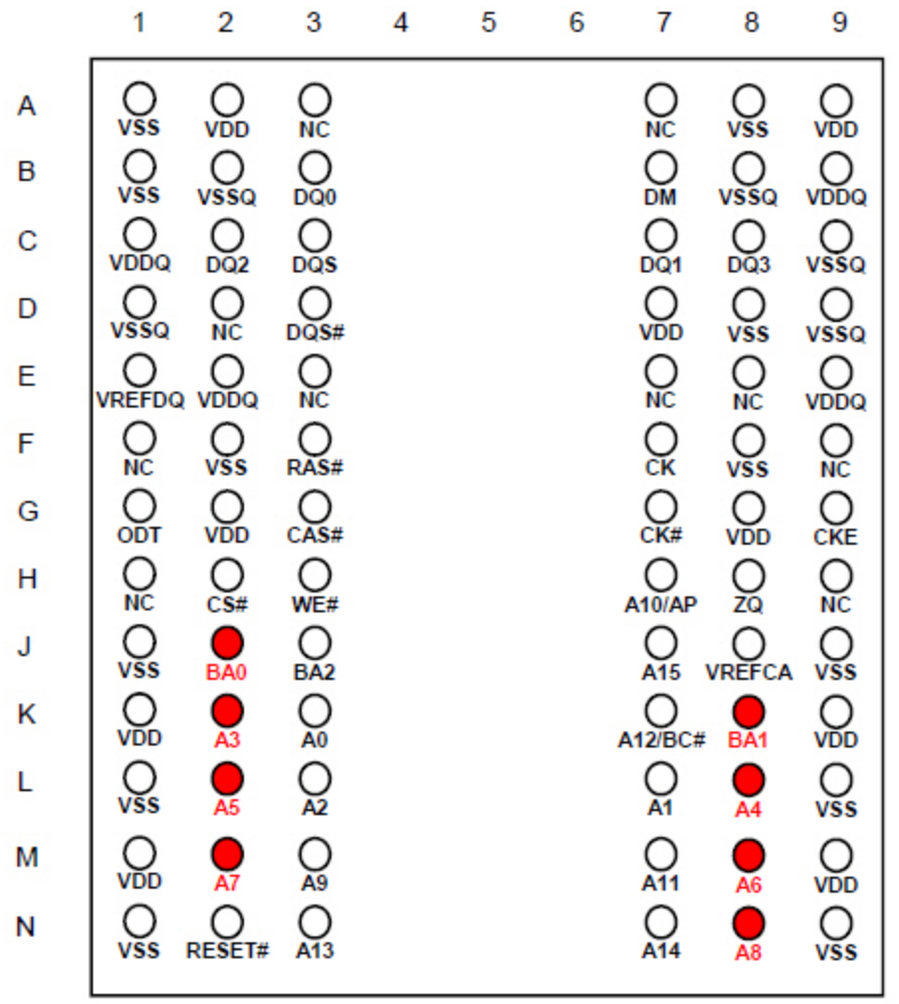

x4 DDR3 SDRAM,
looking through the package

MINDSHARE

BRINGING LIFE TO KNOWLEDGE

DRAM Controller Basics

Controller Basics

- 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.

Required Blocks

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

DDR Controller Block Diagram

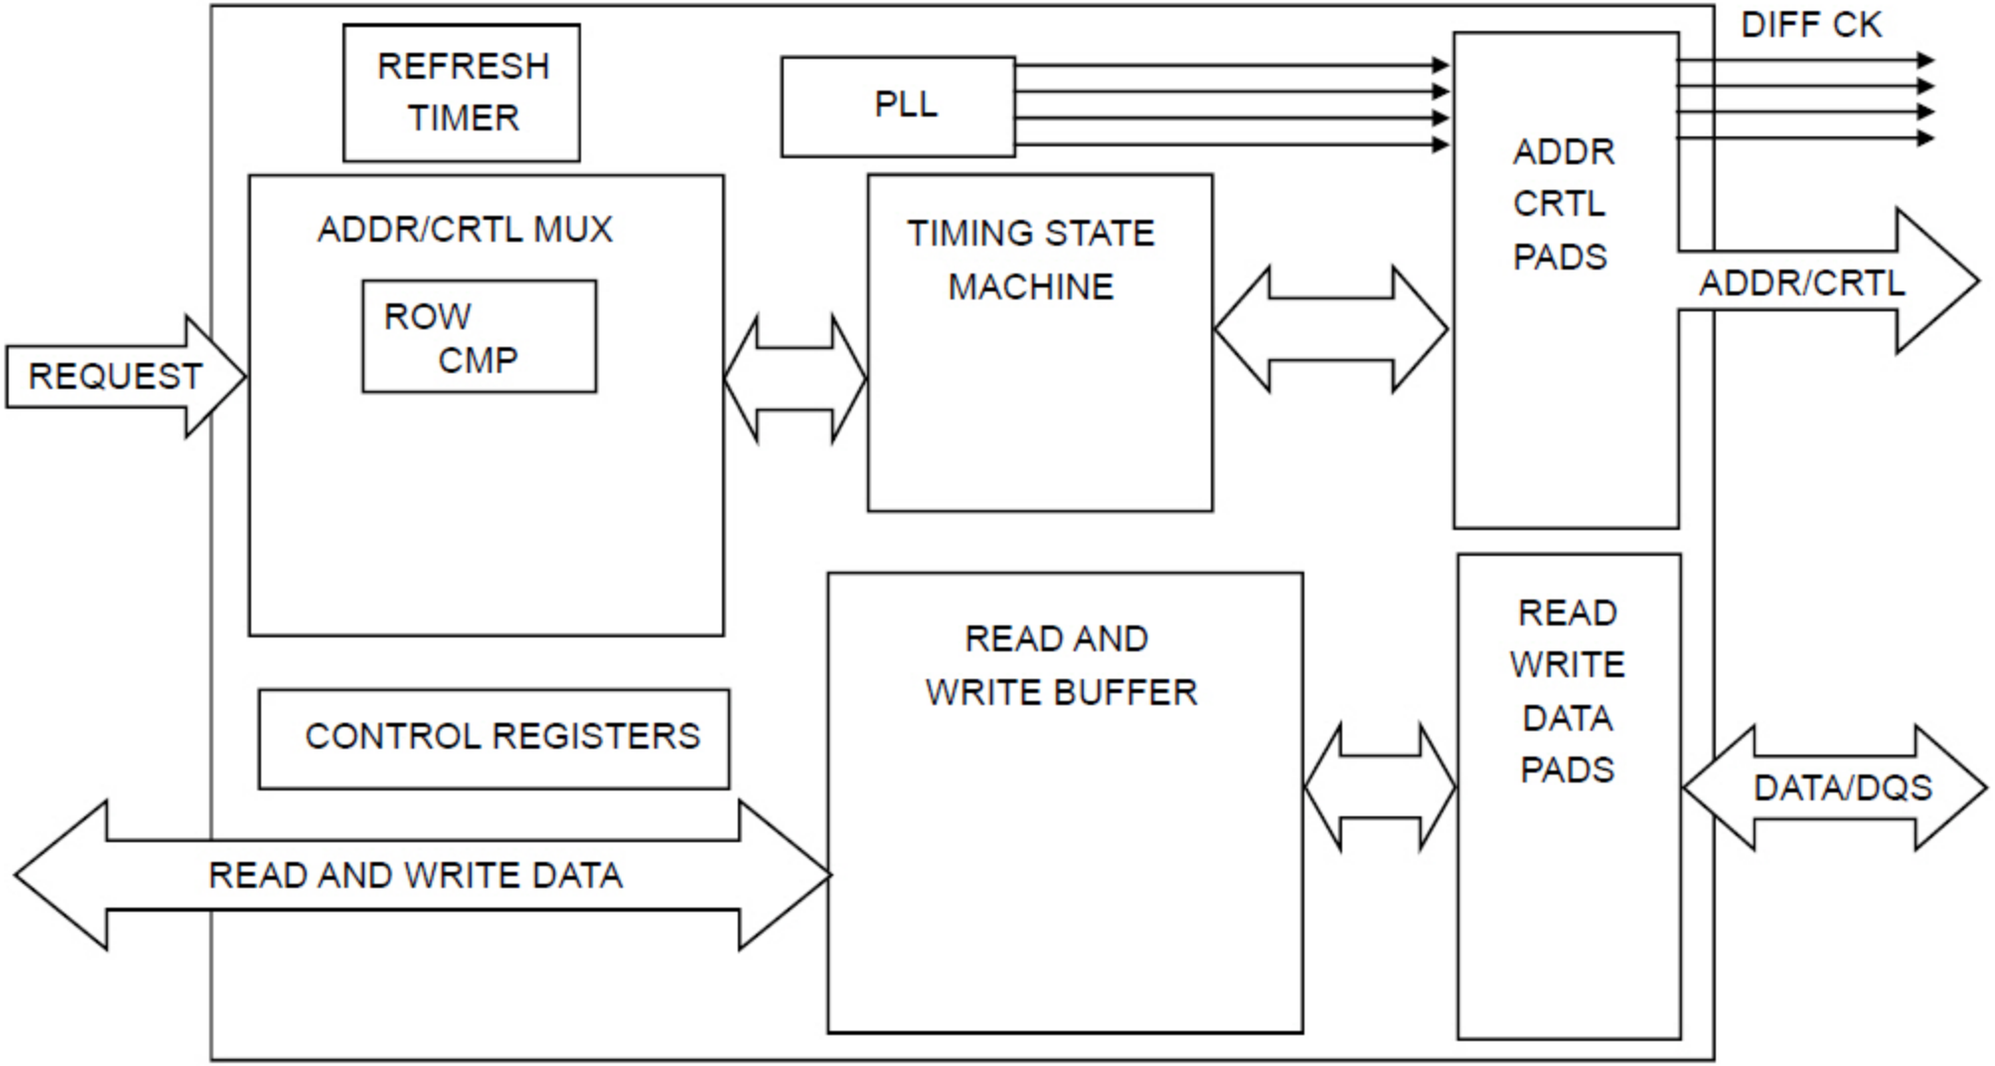

Address and Control Mux

- 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).

PLL

- 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

Timing/State Generator

- 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.

Control Register

- 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.

Read and Write Buffer

- 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.

IO Pads

- 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

MINDSHARE

BRINGING LIFE TO KNOWLEDGE

Addresses

What is the Address?

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

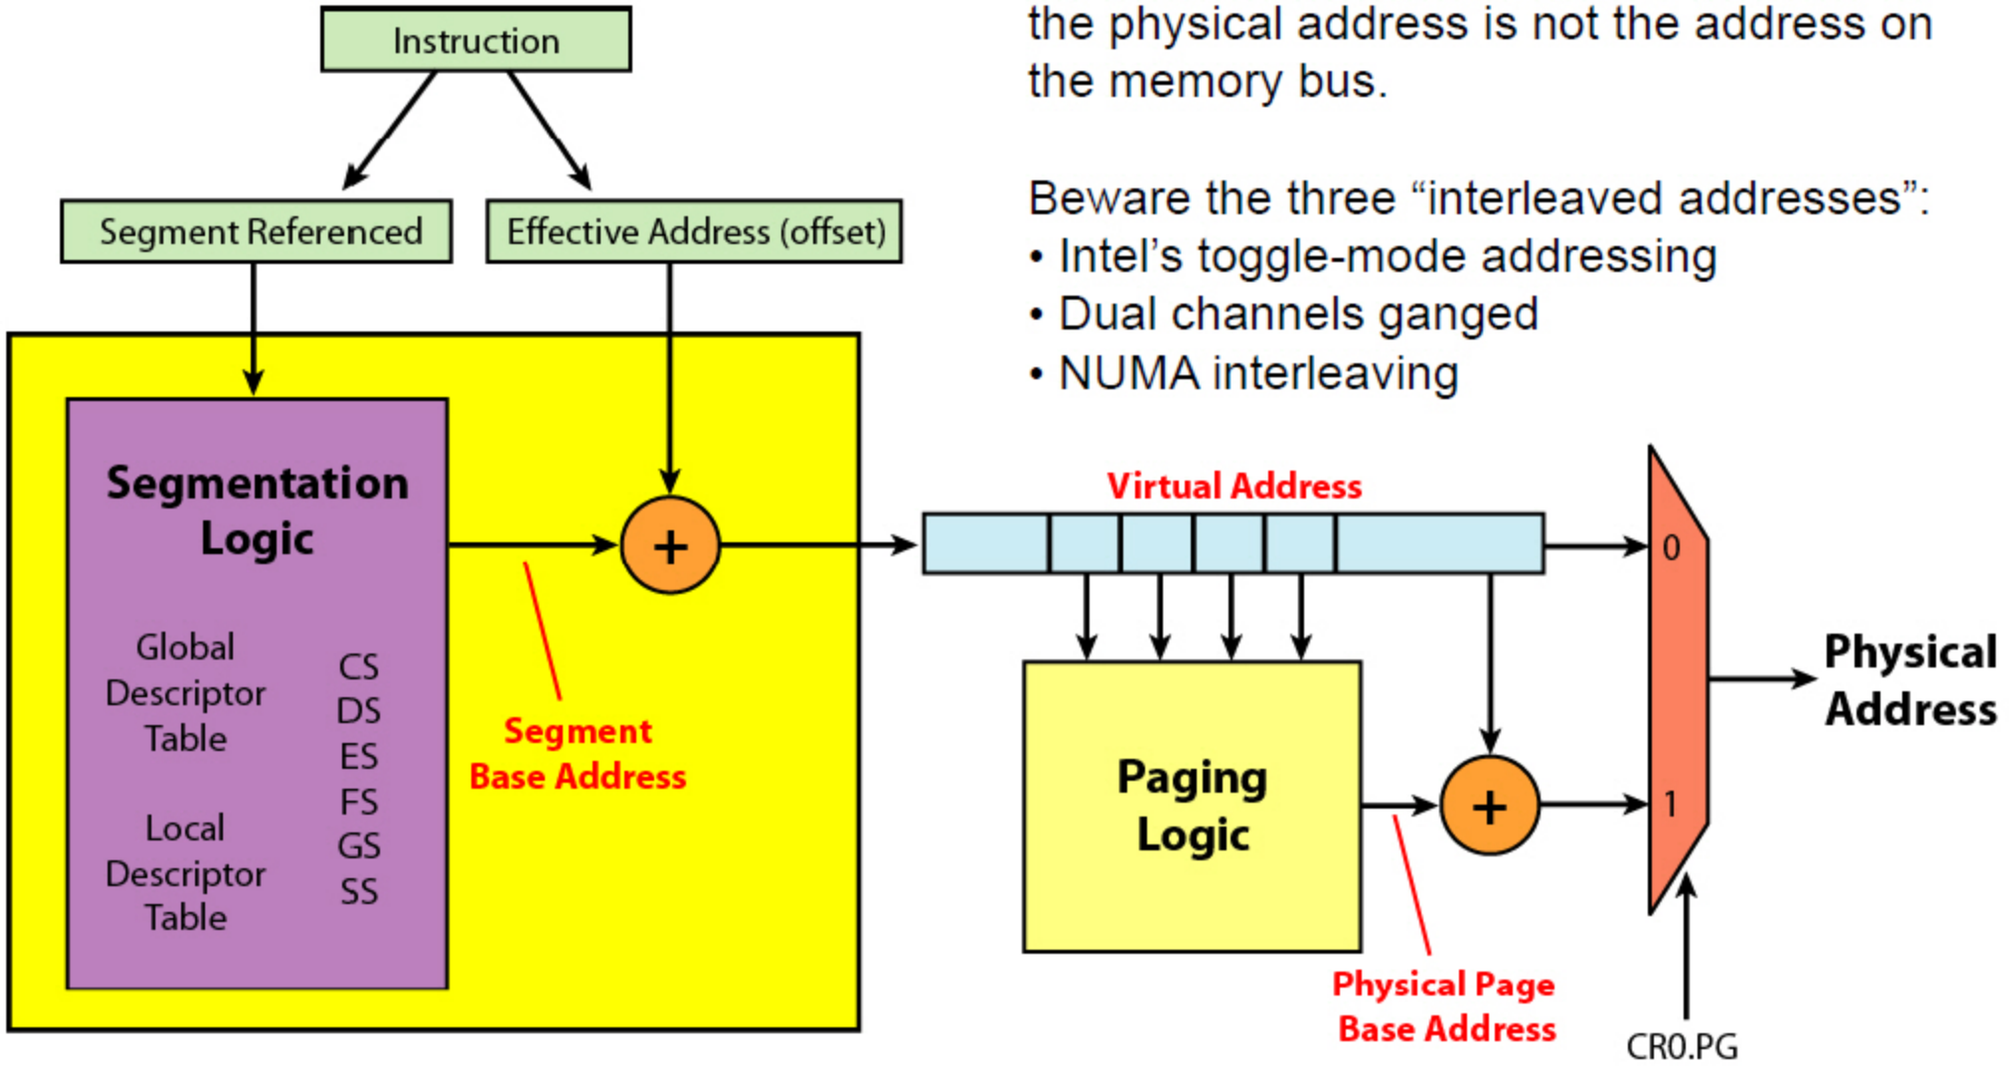

NUMA Addresses

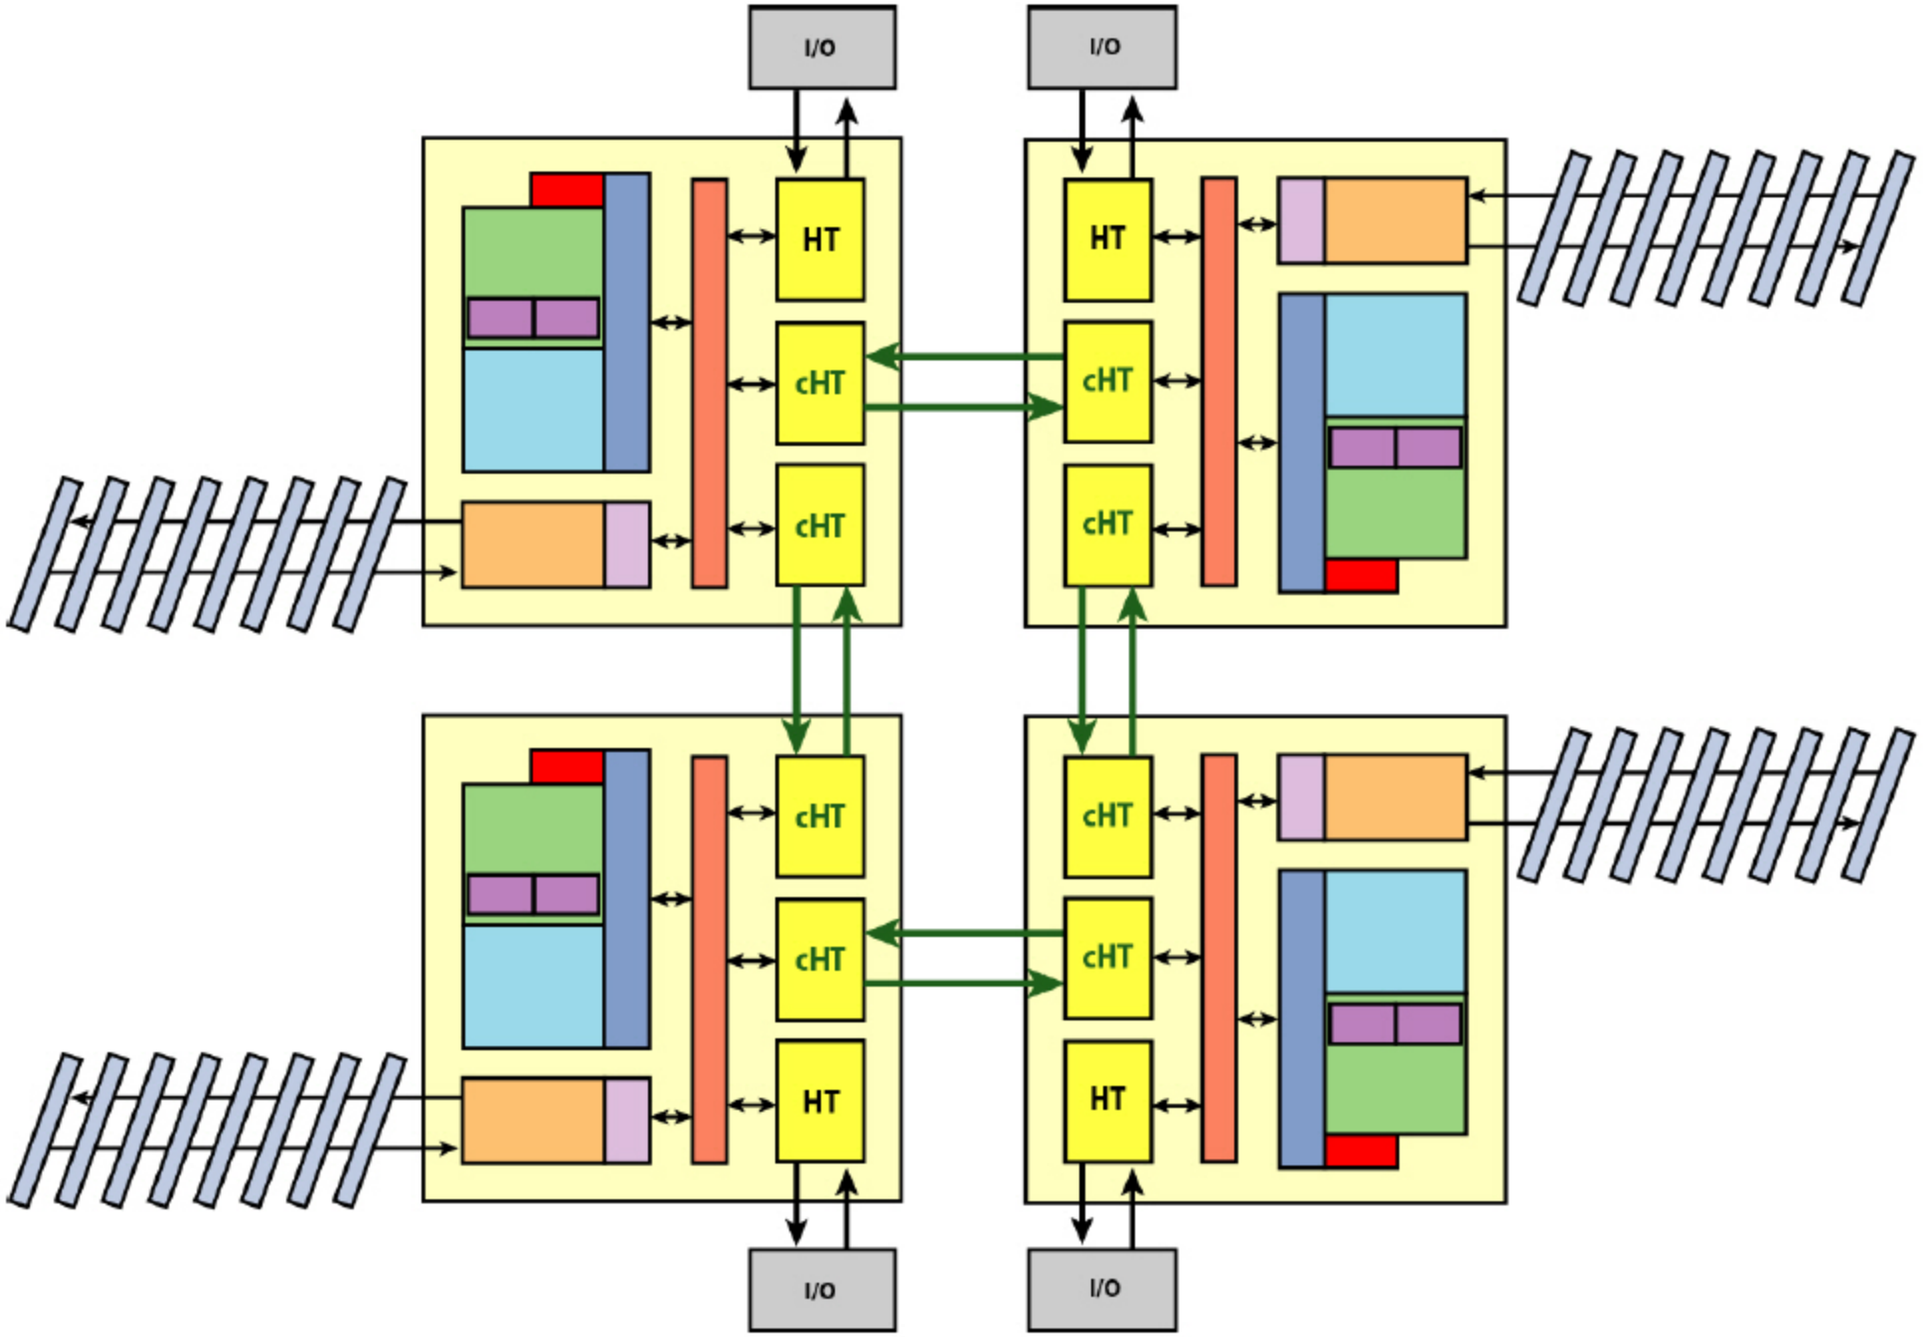

Dual-channel Address Schemes

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

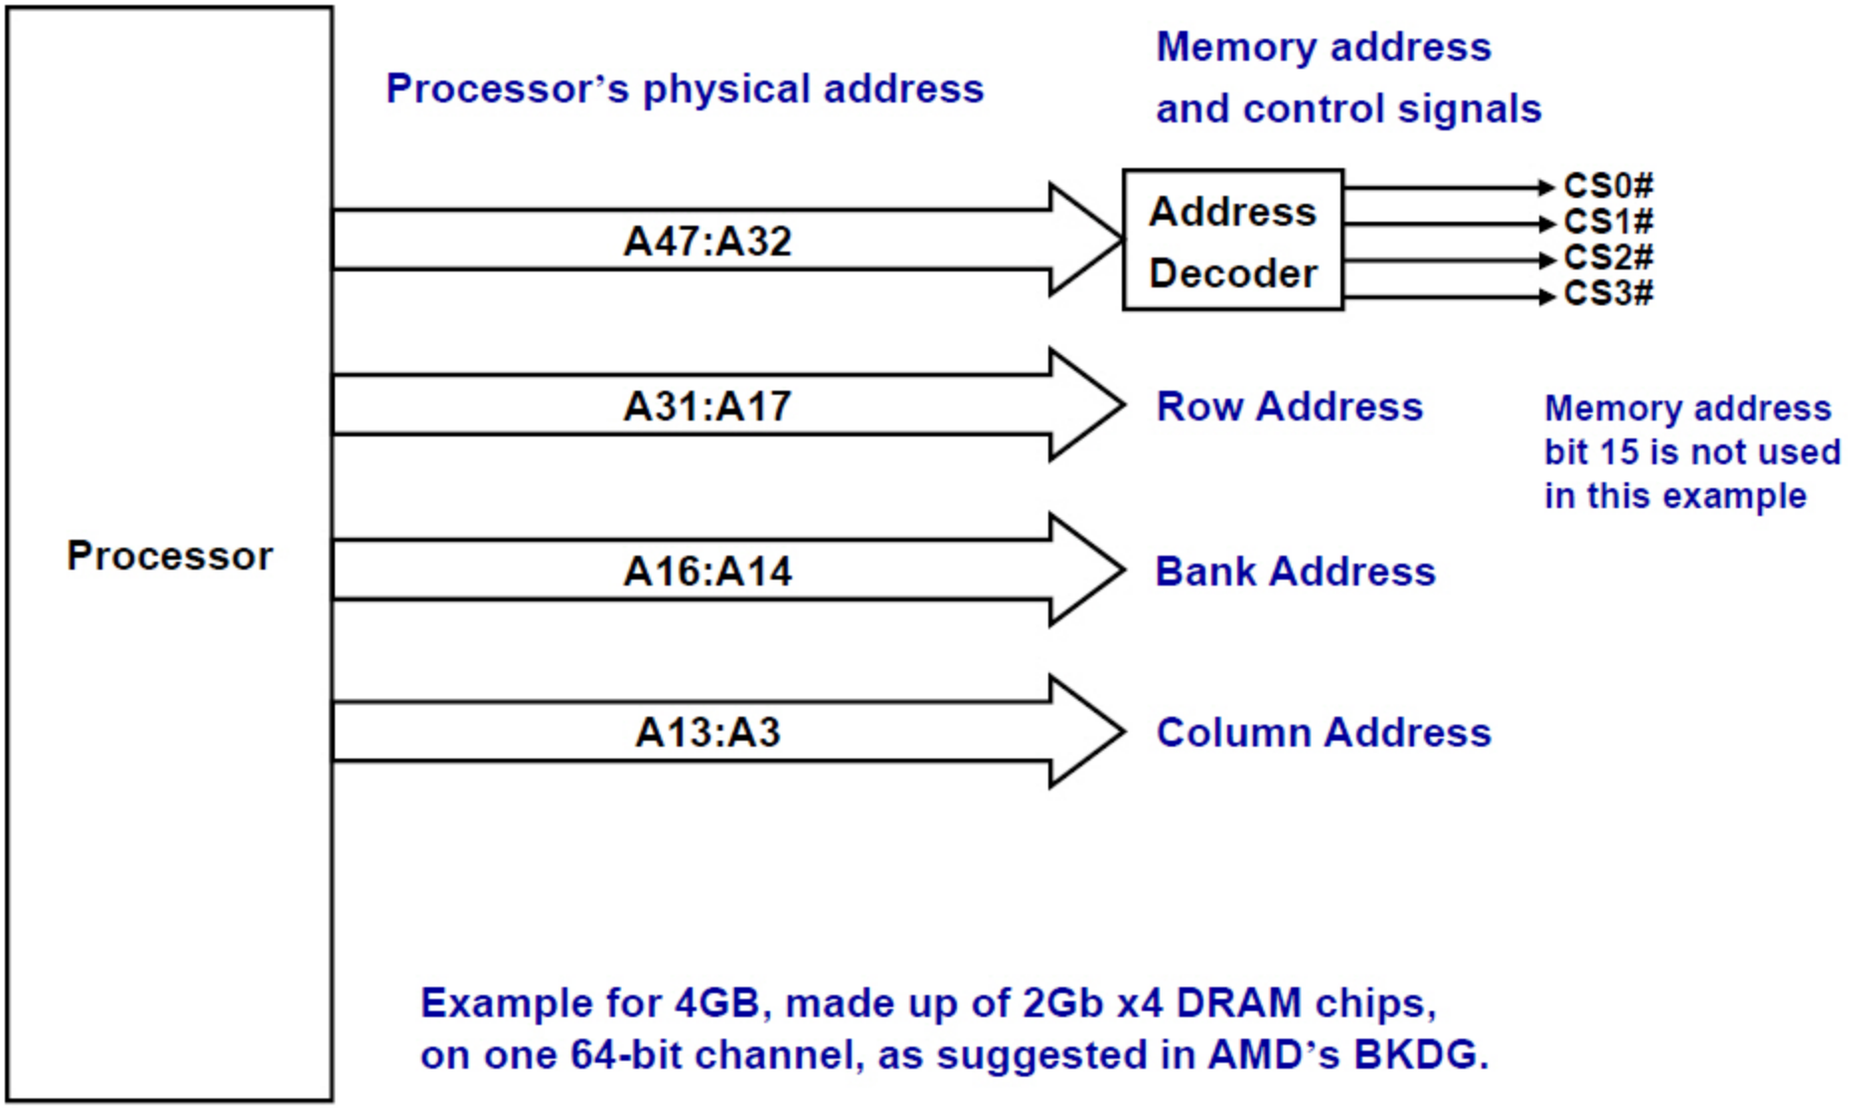

Example for 4GB, made up of 2Gb x4 DRAM chips, on one 64-bit channel, as suggested in AMD's BKDG.

See handout showing 8GB DDR2 with Intel 955 MCH.

Thank you!

MINDSHARE

explore ***the possibilities***

Courses

- Public Course Registration
- Full Course List
- Intel Architecture Courses
- AMD Architecture Courses
- Memory Technology Courses
- Virtualization Technology Courses
- Firmware Courses
- IO Bus Courses
- Storage Courses
- Linux Courses
- Windows Courses
- High-Speed Courses
- Electronic/Digital Courses
- SMT Courses
- Management Courses

WHERE MINDSHARE HAS BEEN LATELY

WHERE CAN WE GO FOR YOU?

MINDSHARE & LOGIC TECHNOLOGY

PC Virtualization Overview Webinar

MindShare and Logic Technology are teaming up to provide a webinar on the topic of PC Virtualization. This **FREE** 1-hour webinar will give an overview of what virtualization is, how it works, and the advantages of virtualization.

This online event will be given on Monday, November 5th at 8:30am Central time.

Click here for more information.