Bruce Jacob

University of
Maryland
ECE Dept.

ENEE 359a
Digital VLSI Design

CMOS Memories and Systems: Part I, DRAM Systems

Prof. Bruce Jacob

blj@ece.umd.edu

Credit where credit is due:
Slides contain original artwork (© Jacob 1999–2004, Wang 2003/4).

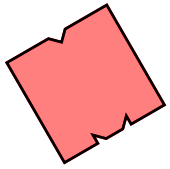

UNIVERSITY OF MARYLAND

Bruce Jacob

University of
Maryland
ECE Dept.

Overview

DRAM:

- DRAM systems
- DRAM circuits

SRAM:

- SRAM systems
- SRAM circuits
- Register files

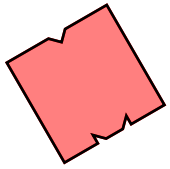

DRAM

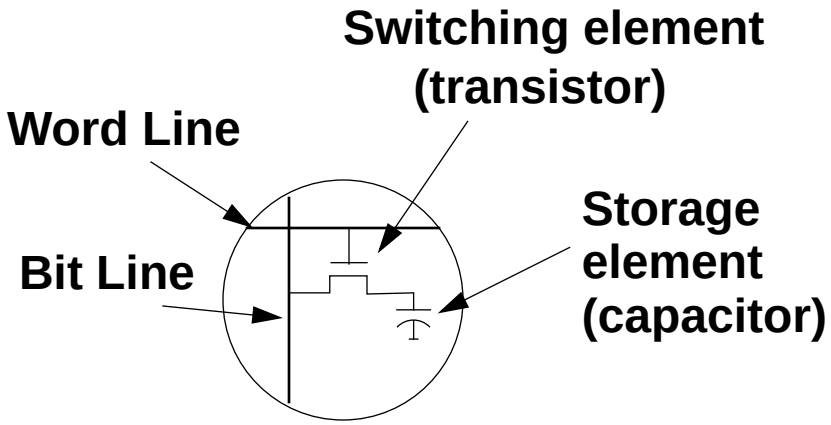

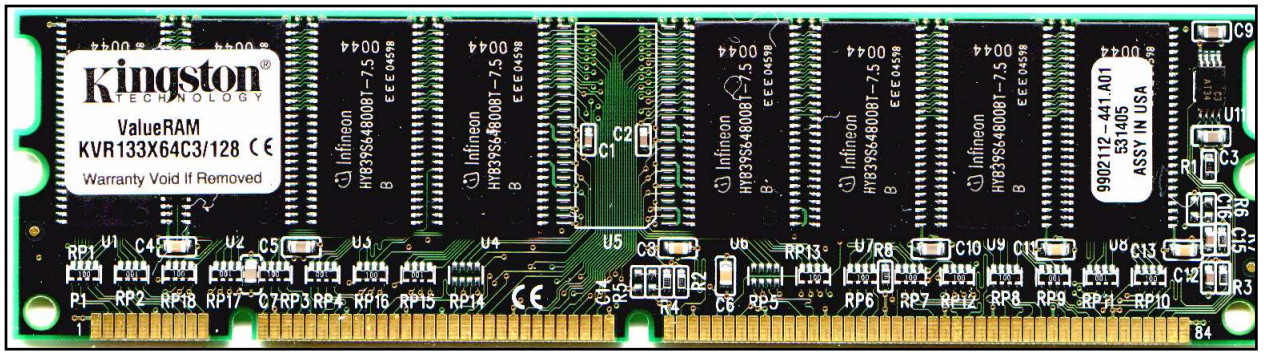

Dual In-line Memory Module (DIMM)

(printed circuit board w/ DRAM chips on it)

The Memory System

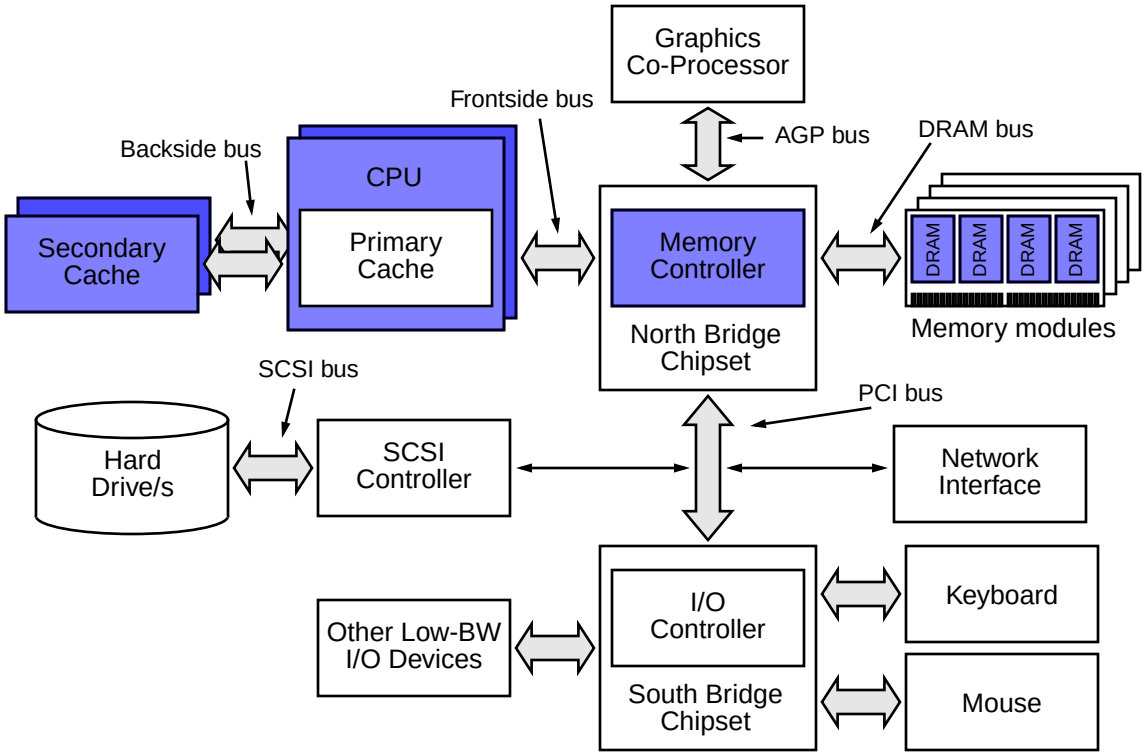

… and DRAM's place within it.

(typical PC-style desktop system)

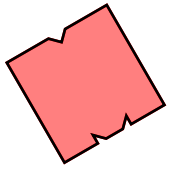

Bruce Jacob

University of
Maryland
ECE Dept.

DRAM-System Closeup

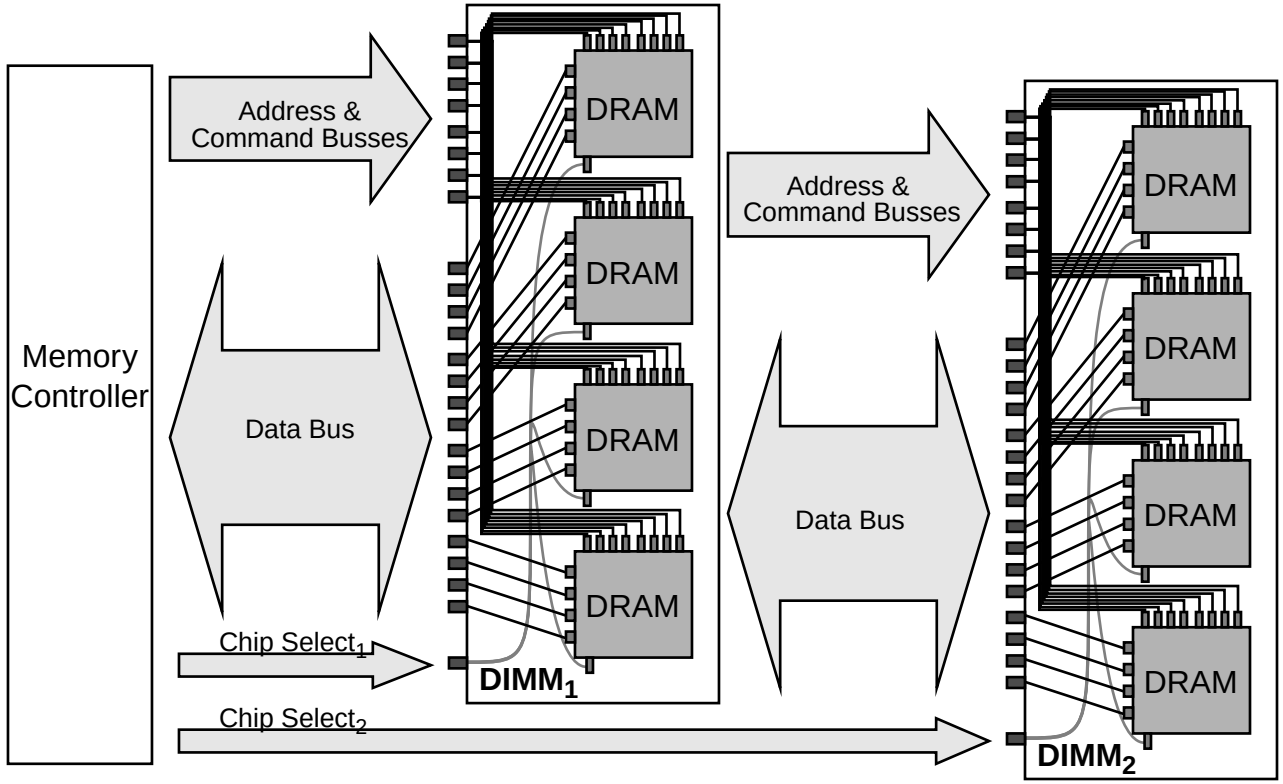

Traditional "JEDEC-Style" DRAM system

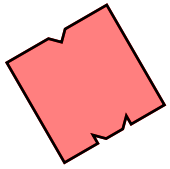

Memory Request Overview

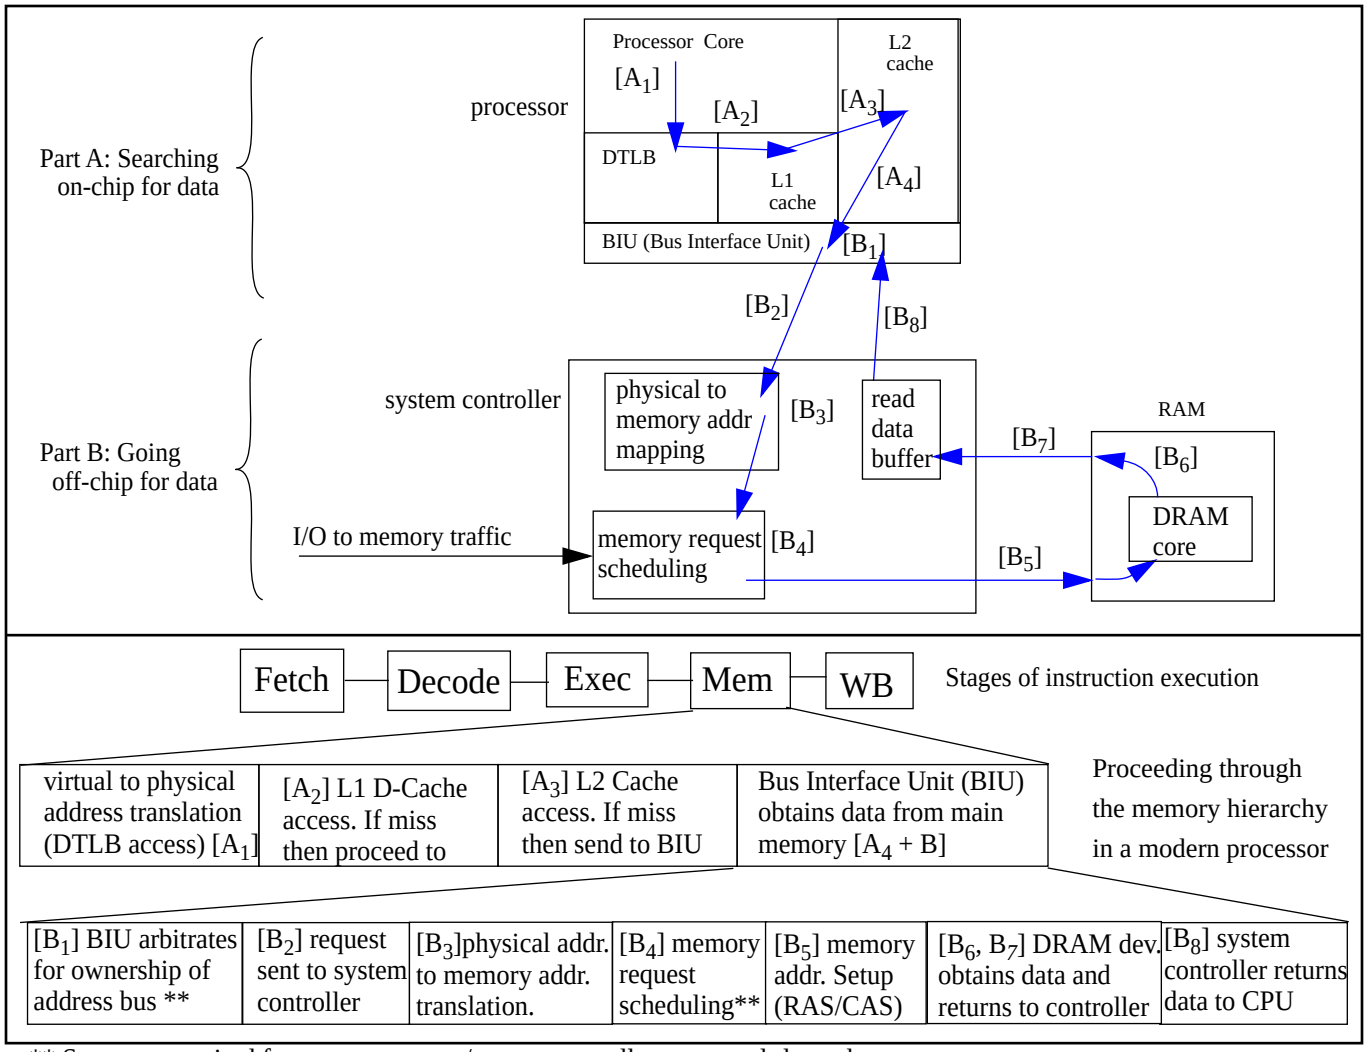

** Steps not required for some processor/system controllers. protocol-dependent.

Progression of a Memory Read Transaction Request Through Memory System

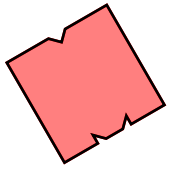

Access-Protocol Basics

DRAM ORGANIZATION

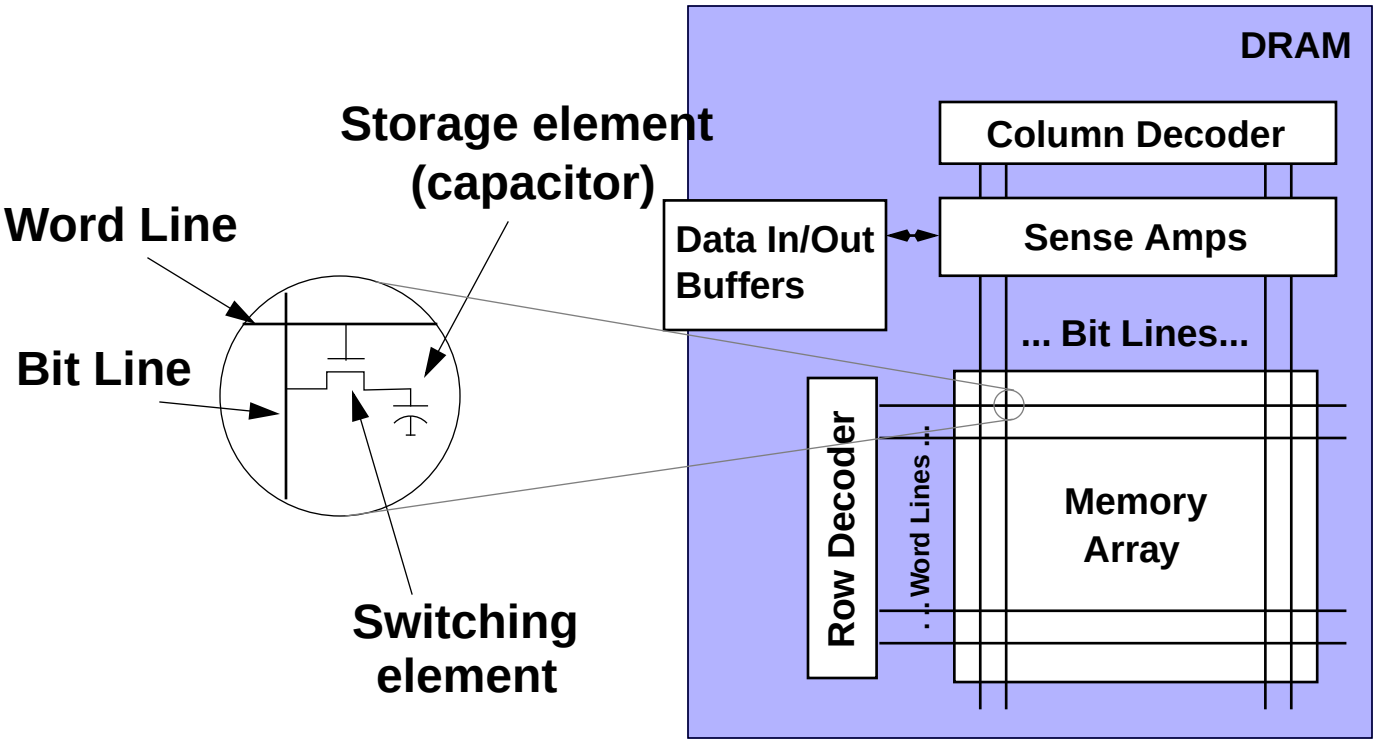

Access-Protocol Basics

BUS TRANSMISSION

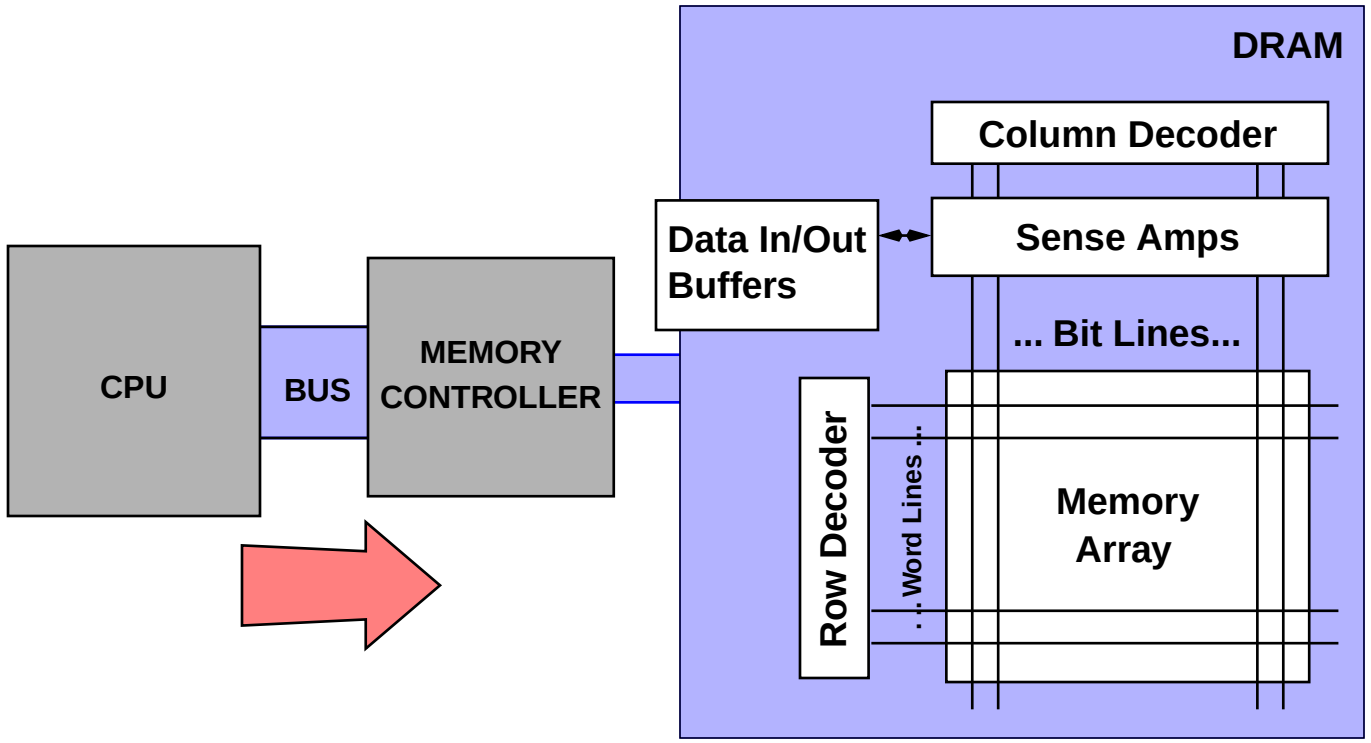

Access-Protocol Basics

[PRECHARGE and] ROW ACCESS

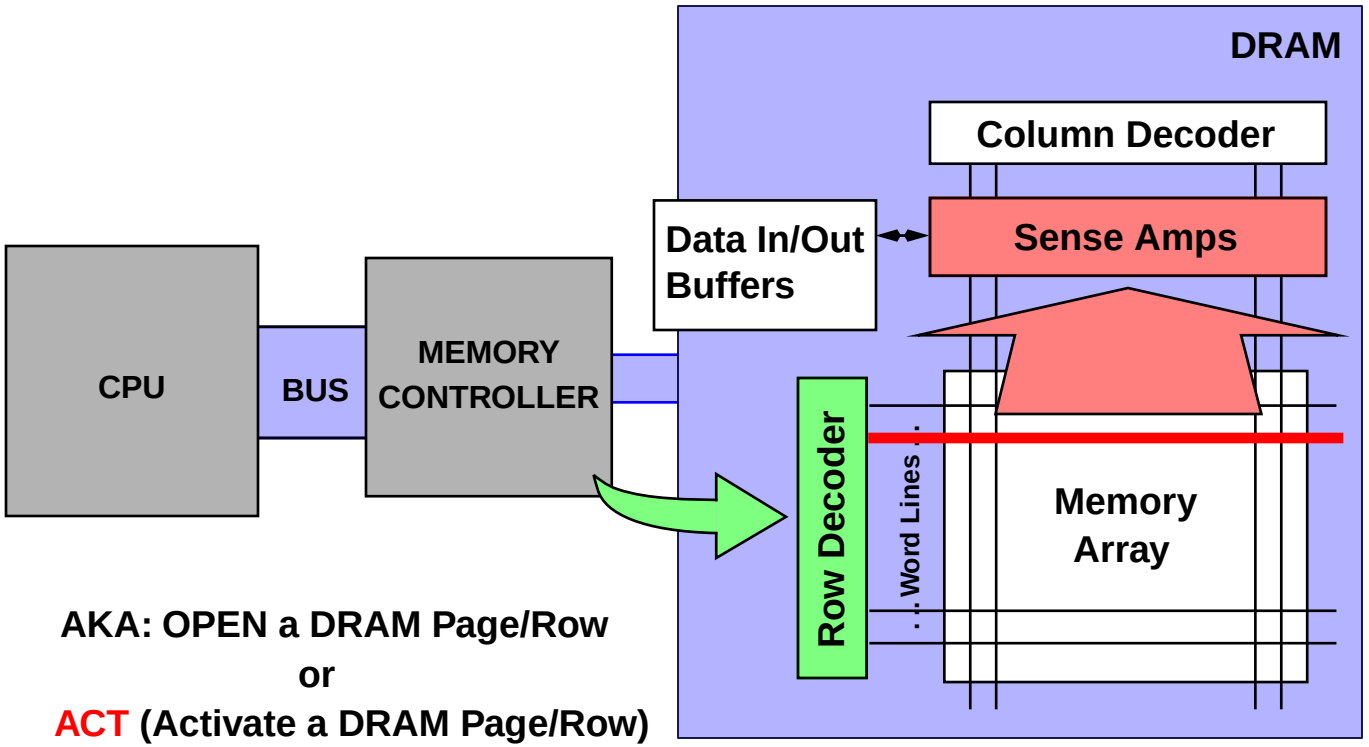

AKA: OPEN a DRAM Page/Row
or
ACT (Activate a DRAM Page/Row)
or
RAS (Row Address Strobe)

Access-Protocol Basics

COLUMN ACCESS

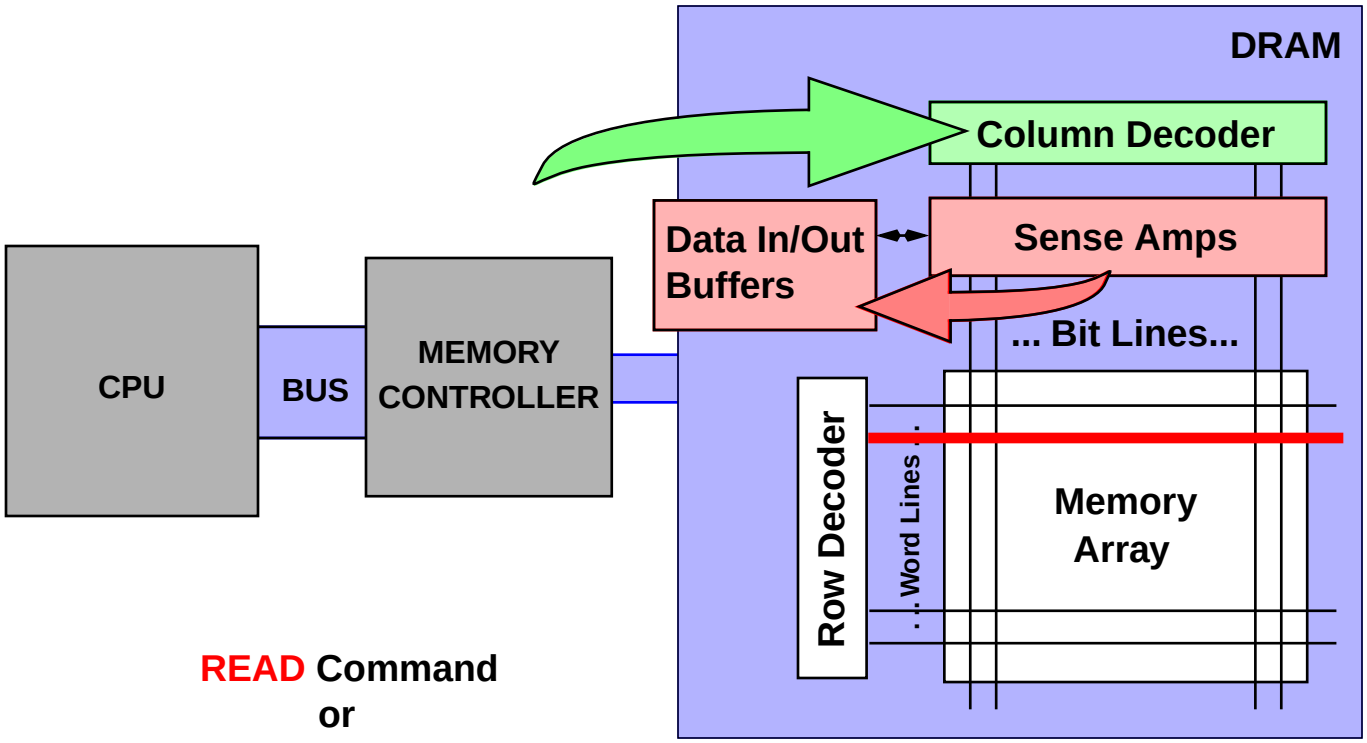

READ Command
or
CAS: Column Address Strobe

Access-Protocol Basics

DATA TRANSFER

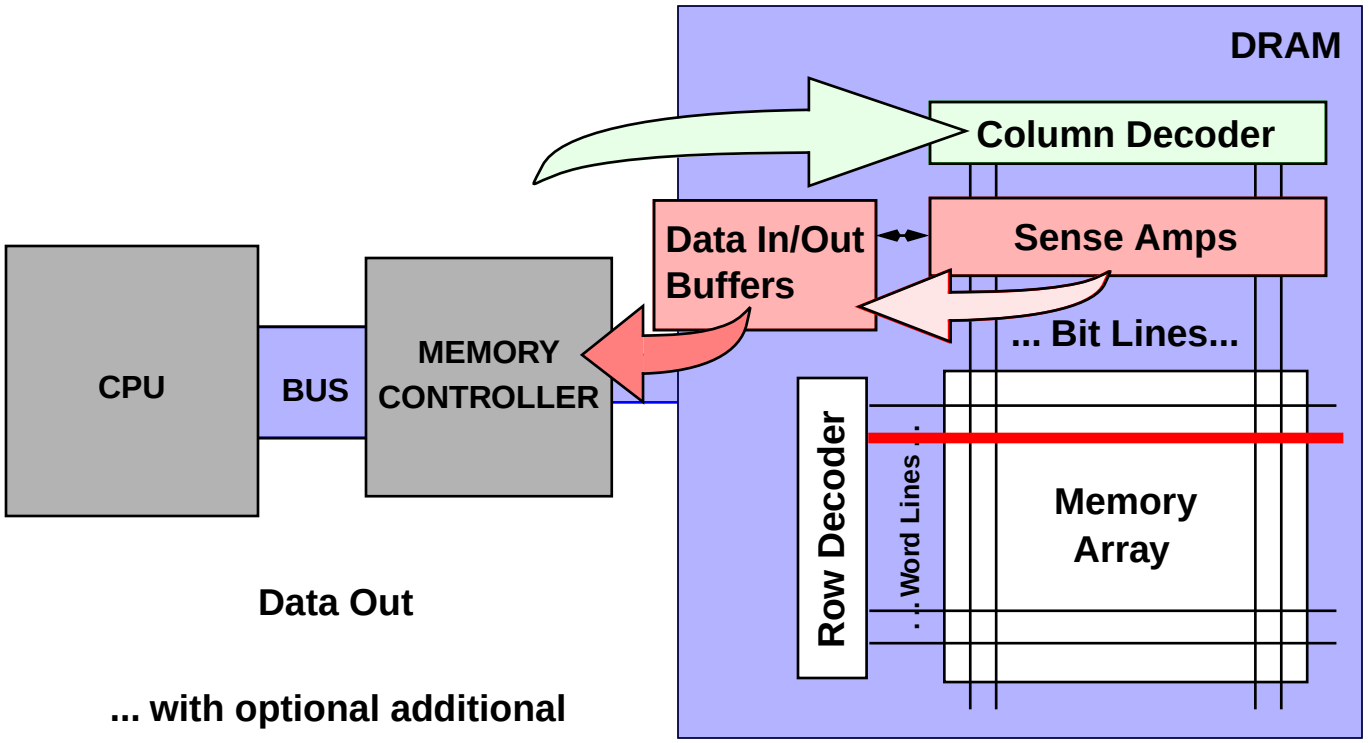

... with optional additional
CAS: Column Address Strobe

note: page mode enables overlap with CAS

Access-Protocol Basics

BUS TRANSMISSION

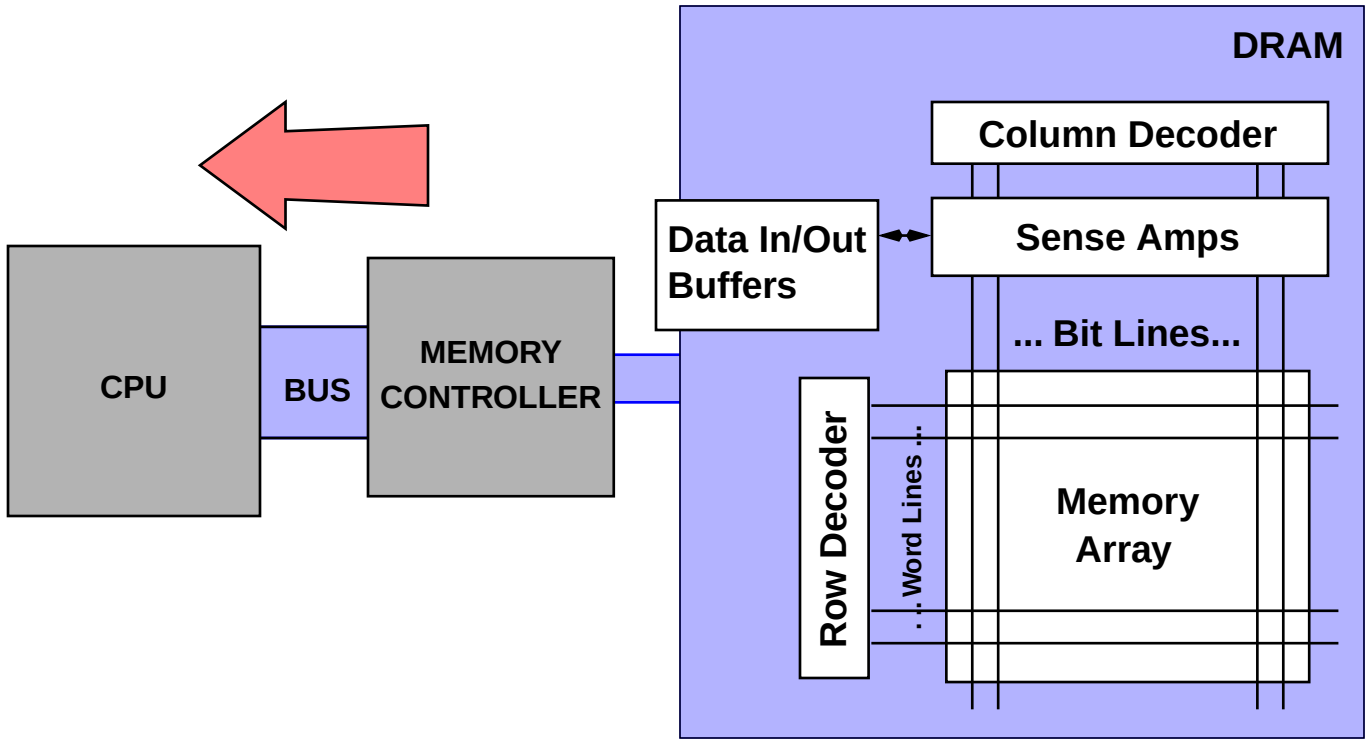

Bruce Jacob

University of
Maryland
ECE Dept.

Access-Protocol Basics

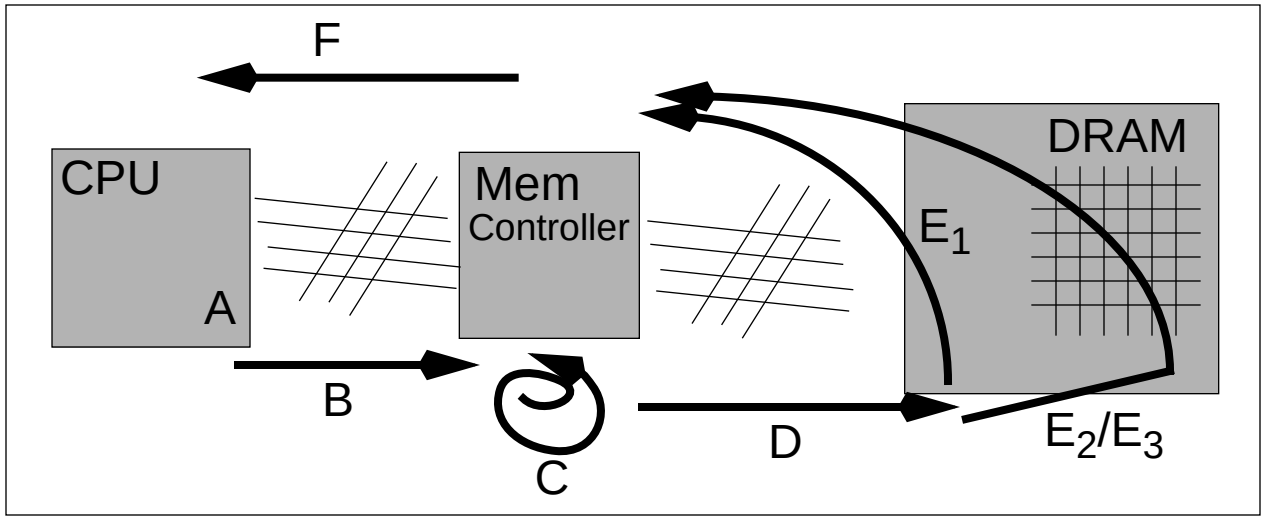

A: Transaction request may be delayed in Queue
B: Transaction request sent to Memory Controller
C: Transaction converted to Command Sequences (may be queued)
D: Command/s Sent to DRAM
E_1: Requires only a **CAS** or
E_2: Requires **RAS + CAS** or
E_3: Requires **PRE + RAS + CAS**
F: Transaction sent back to CPU

"DRAM Latency" = A + B + C + D + E + F

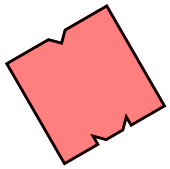

Bruce Jacob

University of
Maryland
ECE Dept.

Access-Protocol Basics

Read Timing for Conventional DRAM

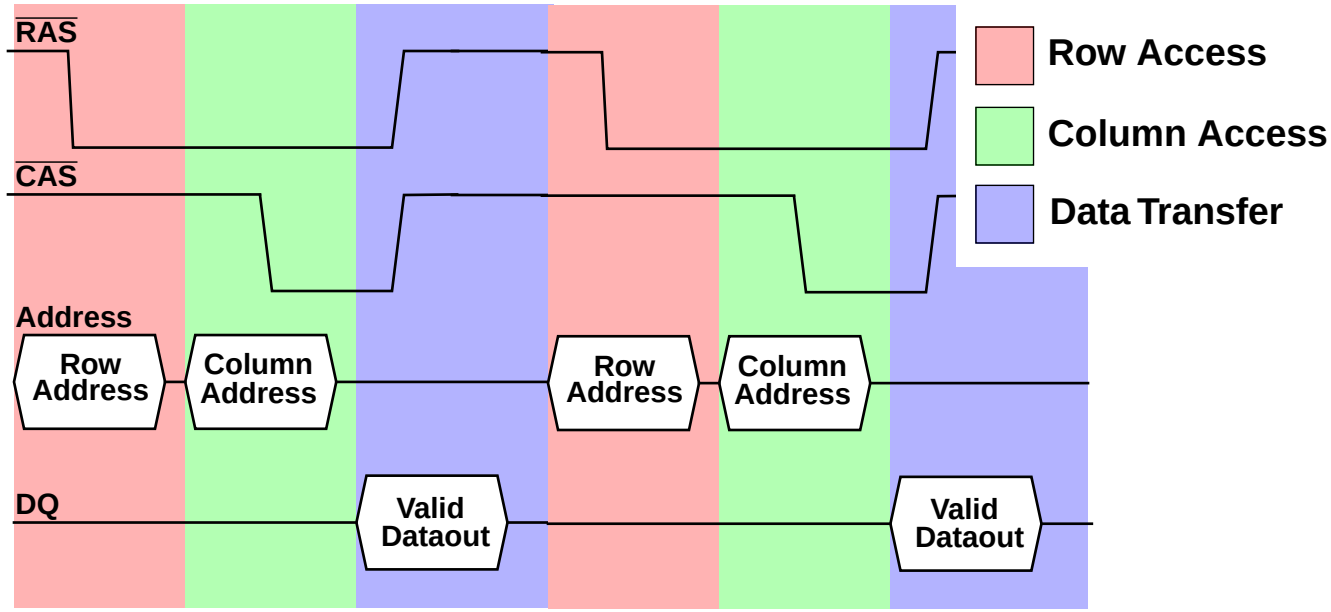

Access-Protocol Basics

Read Timing for Synchronous DRAM

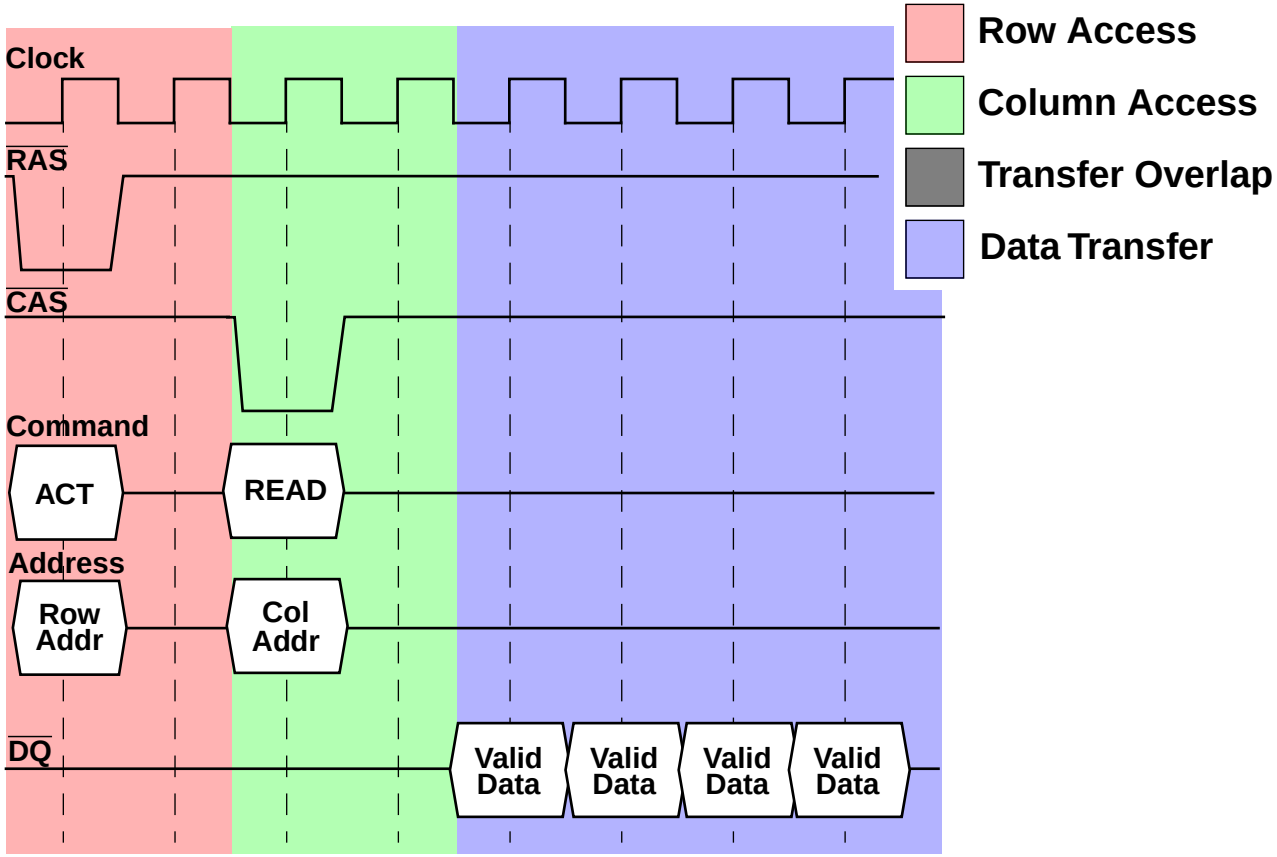

($\overline{RAS}$ + $\overline{CAS}$ + $\overline{OE}$ ... == Command Bus)

DRAM Circuit Basics

“Row” Defined

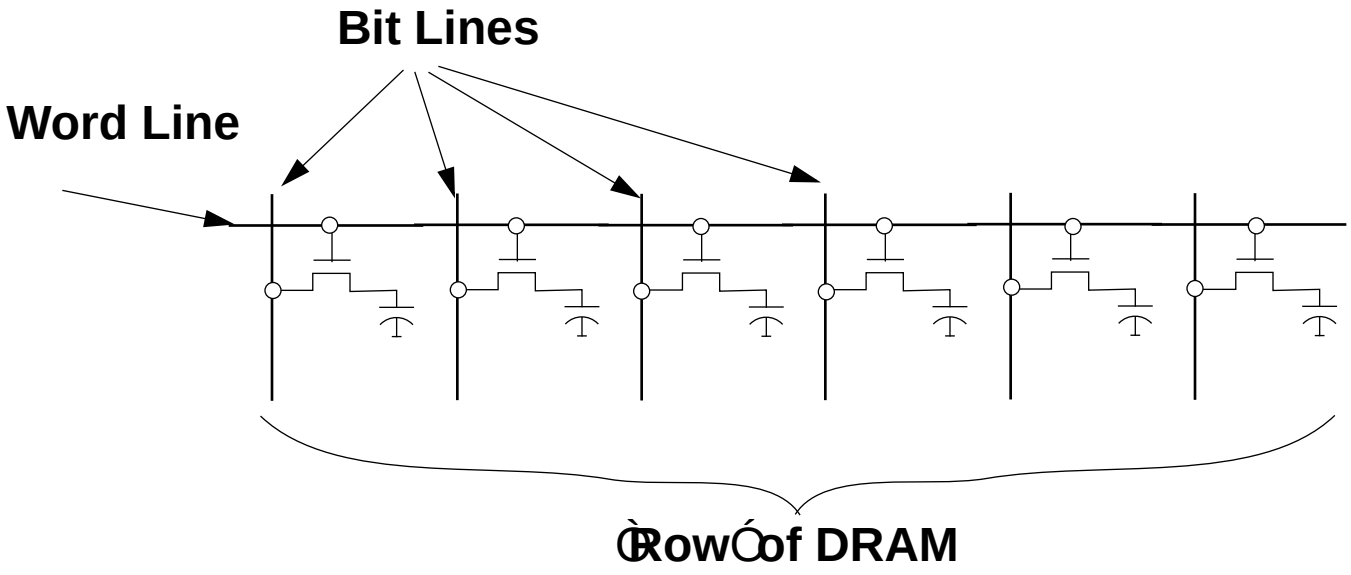

Row Size: 8 Kb @ 256 Mb SDRAM node
4 Kb @ 256 Mb RDRAM node

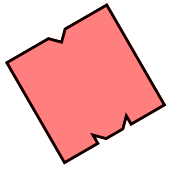

Bruce Jacob

University of
Maryland
ECE Dept.

DRAM Circuit Basics

Sense Amplifier I: *6 rows shown*

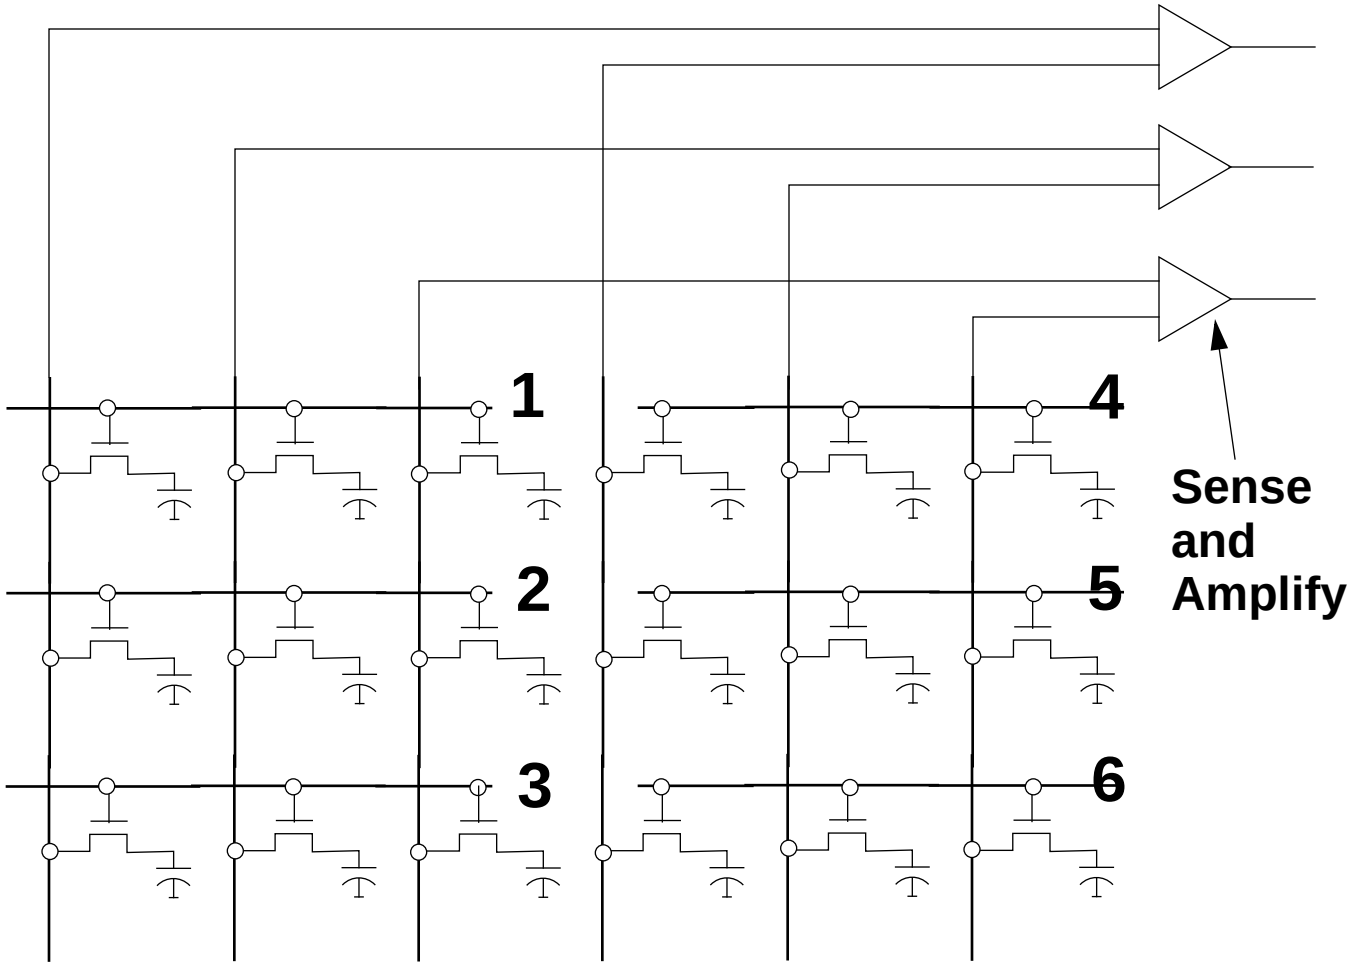

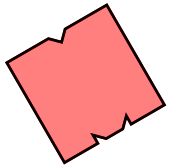

Bruce Jacob

University of
Maryland
ECE Dept.

DRAM Circuit Basics

Sense Amplifier I: *6 rows shown*

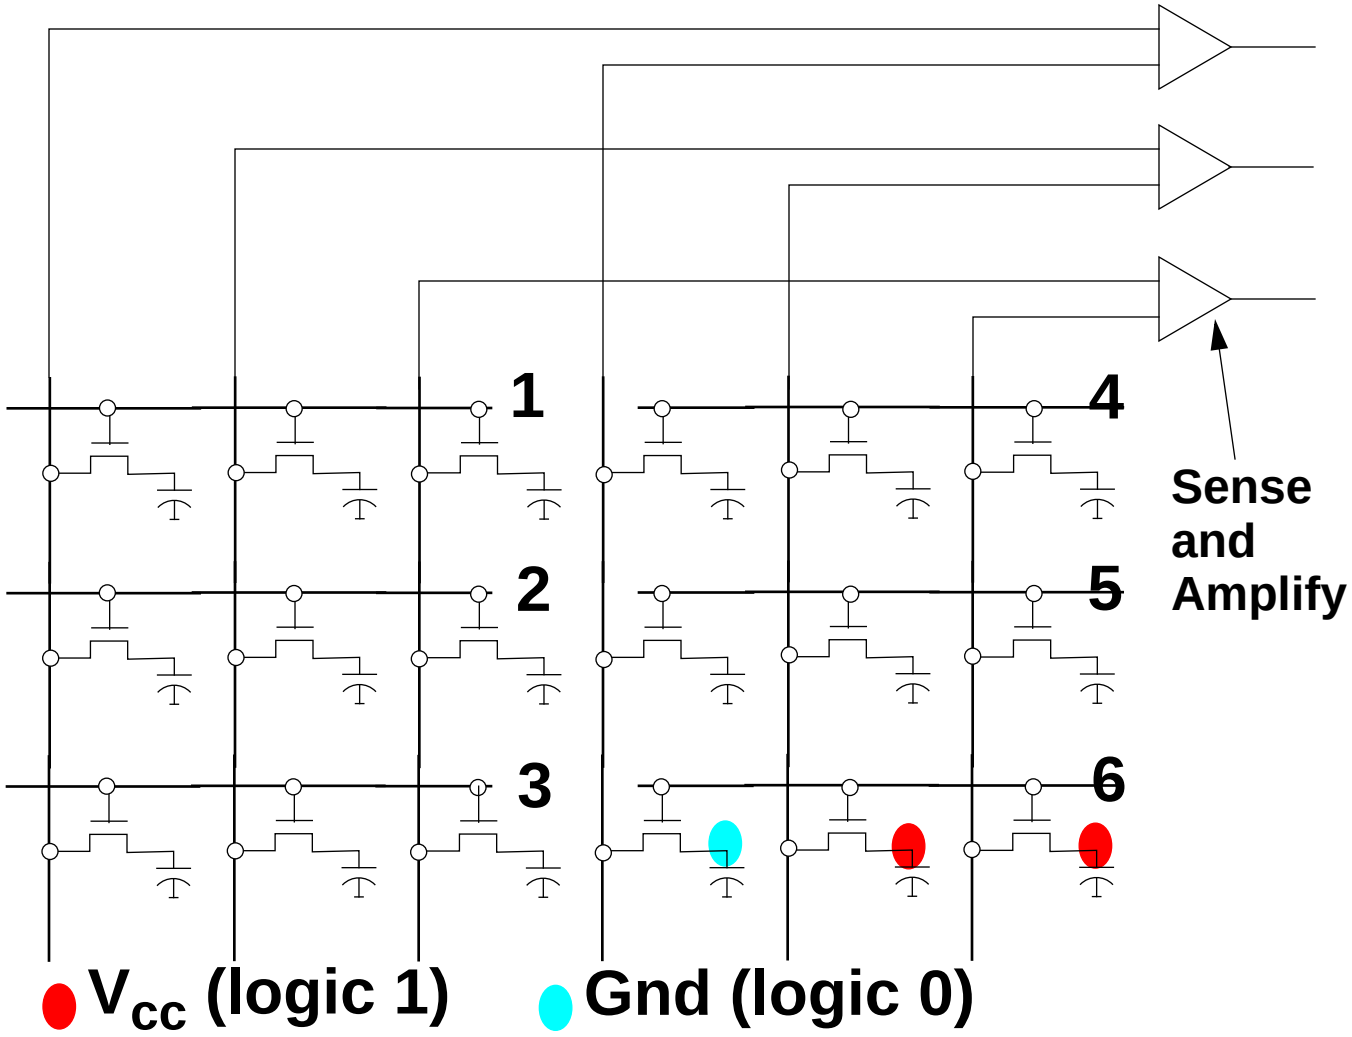

DRAM Circuit Basics

Sense Amplifier II : Precharged

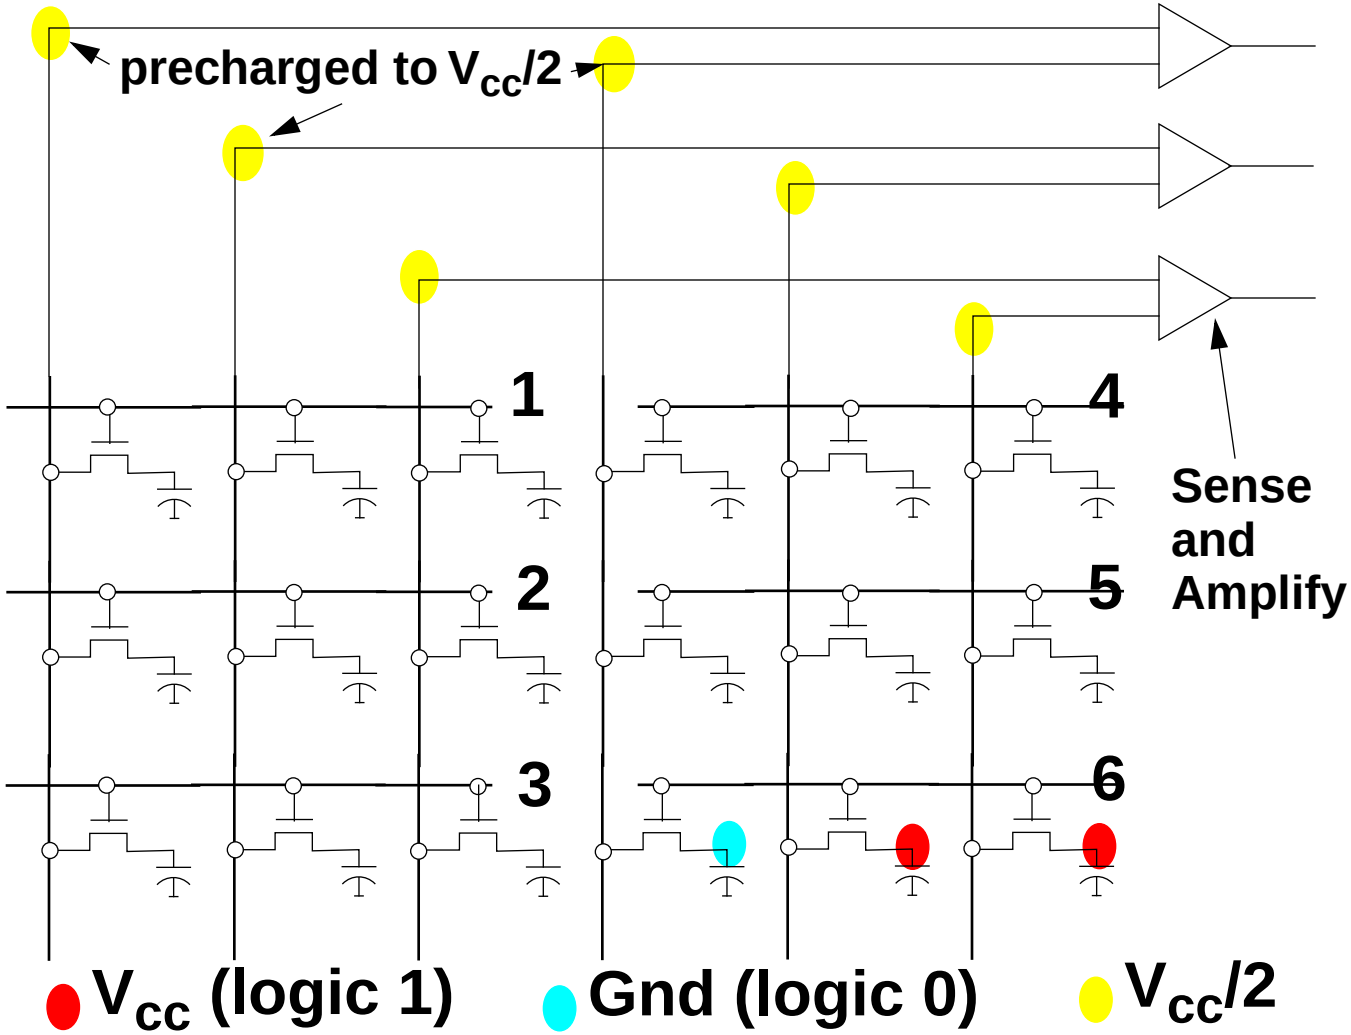

Bruce Jacob

University of
Maryland
ECE Dept.

DRAM Circuit Basics

Sense Amplifier III : Destructive Read

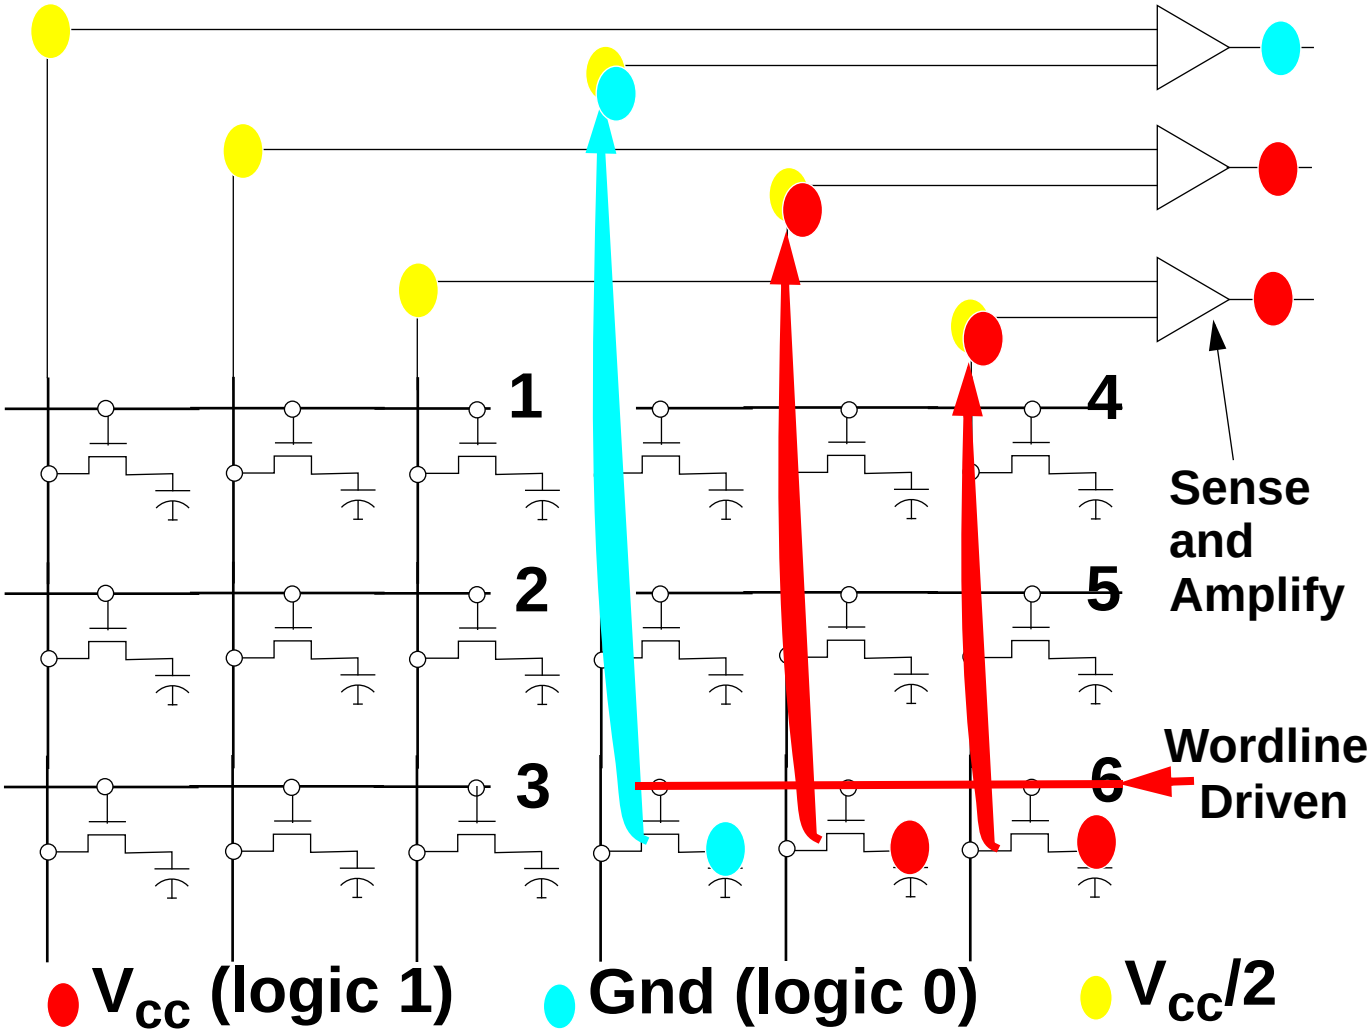

DRAM Circuit Basics

“Column” Defined

Column: Smallest addressable quantity of DRAM on chip

SDRAM*: column size == chip data bus width (4, 8,16, 32)
RDRAM: column size != chip data bus width (128 bit fixed)

SDRAM*: get ÒnÓ columns per access. n = (1, 2, 4, 8)
RDRAM: get 1 column per access.

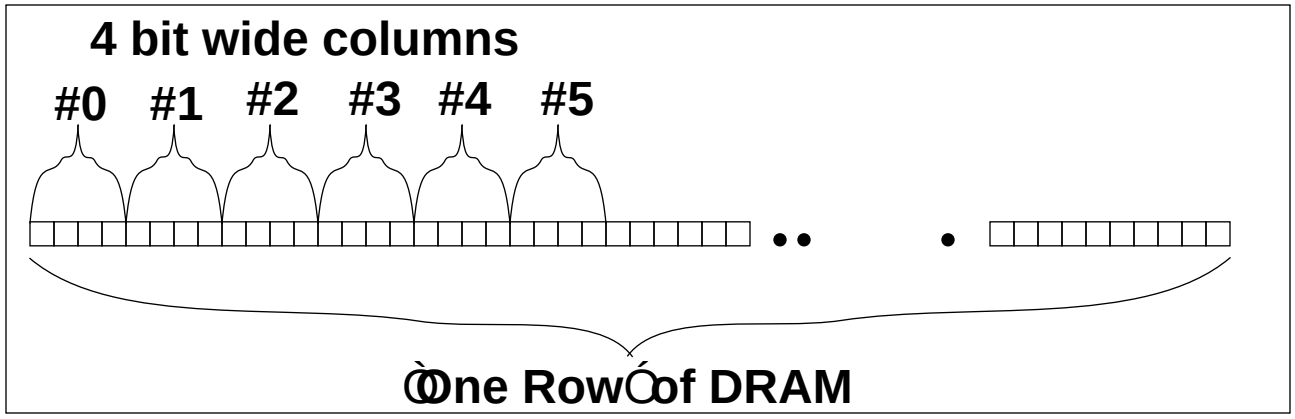

*** SDRAM means SDRAM and variants. i.e. DDR SDRAM**

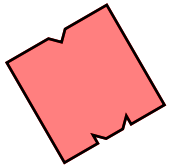

DRAM Architecture Basics

PHYSICAL ORGANIZATION

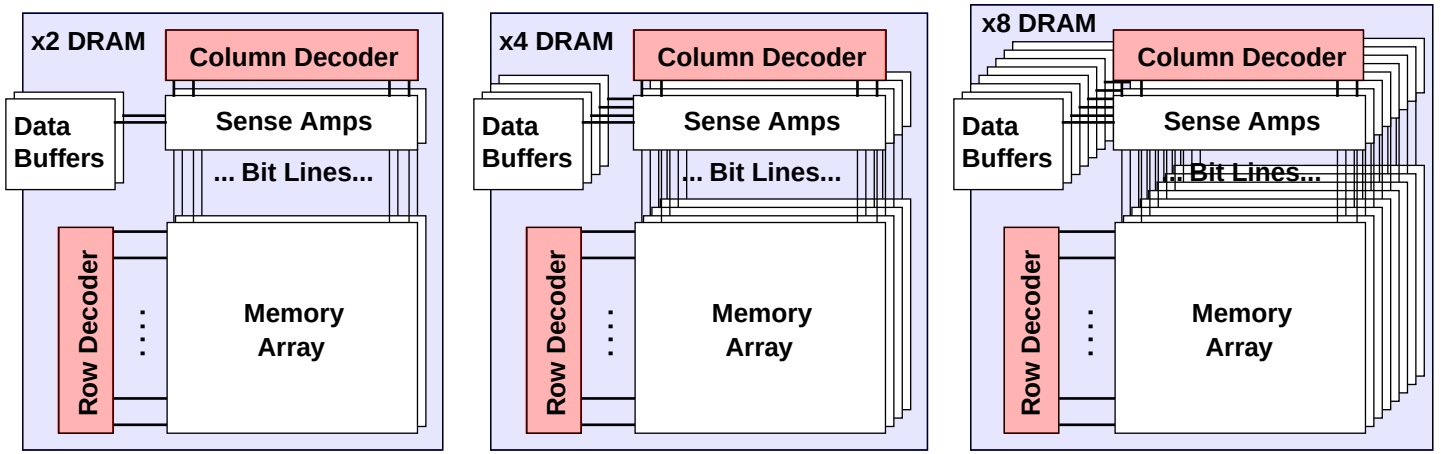

x2 DRAM **x4 DRAM** **x8 DRAM**

This is per bank …

Typical DRAMs have 2+ banks

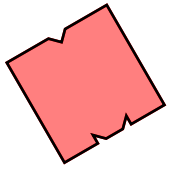

DRAM "Speed" Part I

How fast can I move data from DRAM cell to sense amp?

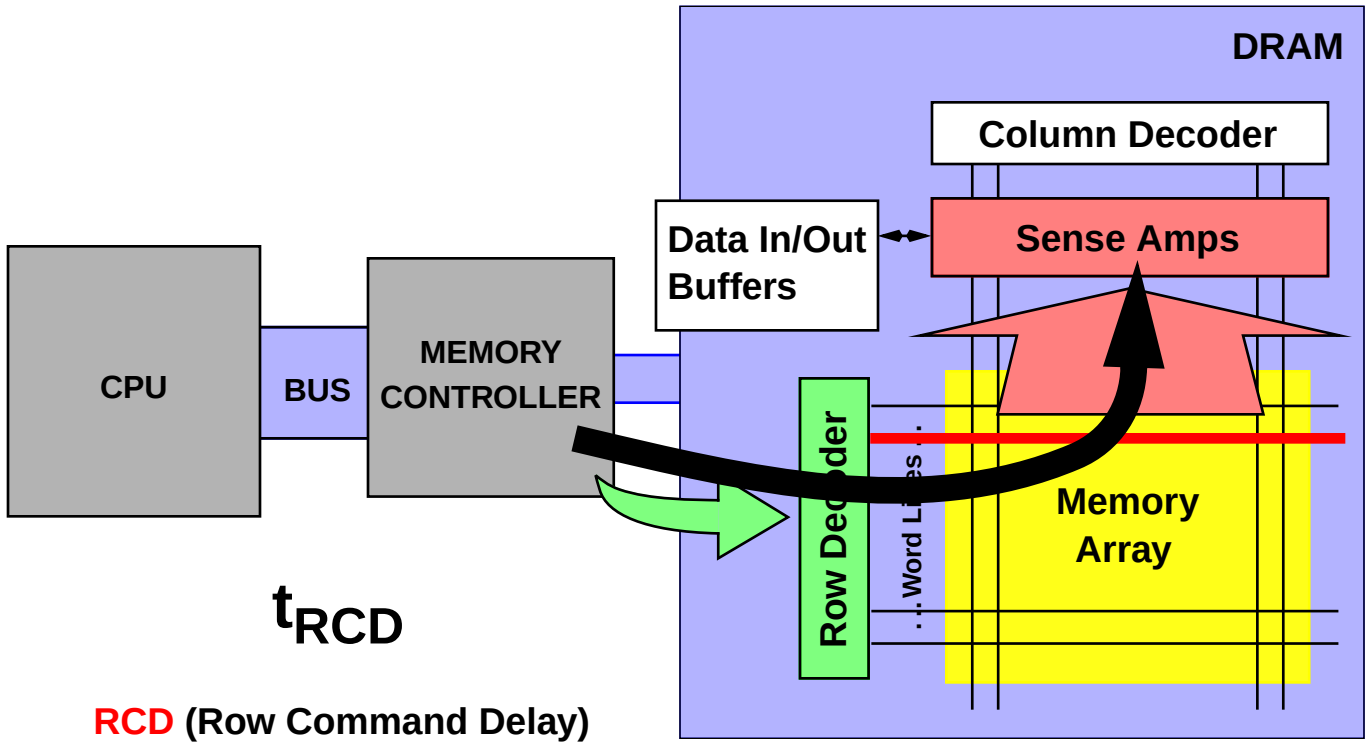

RCD (Row Command Delay)

Bruce Jacob

University of
Maryland
ECE Dept.

DRAM "Speed" Part II

How fast can I get data out of sense amps back into memory controller?

t_{CAS} aka
t_{CASL} aka
t_{CL}

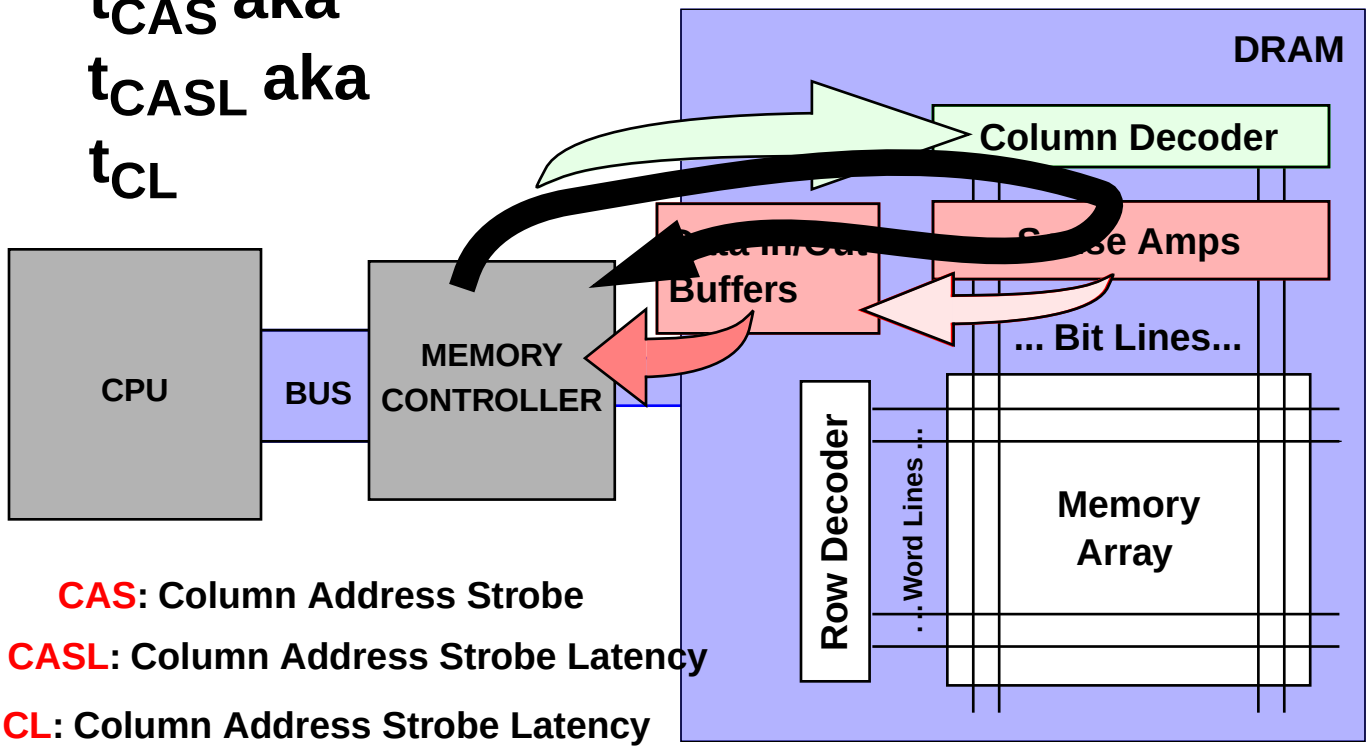

CAS: Column Address Strobe

CASL: Column Address Strobe Latency

CL: Column Address Strobe Latency

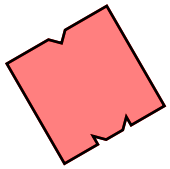

DRAM "Speed" Part III

How fast can I move data from DRAM cell into memory controller?

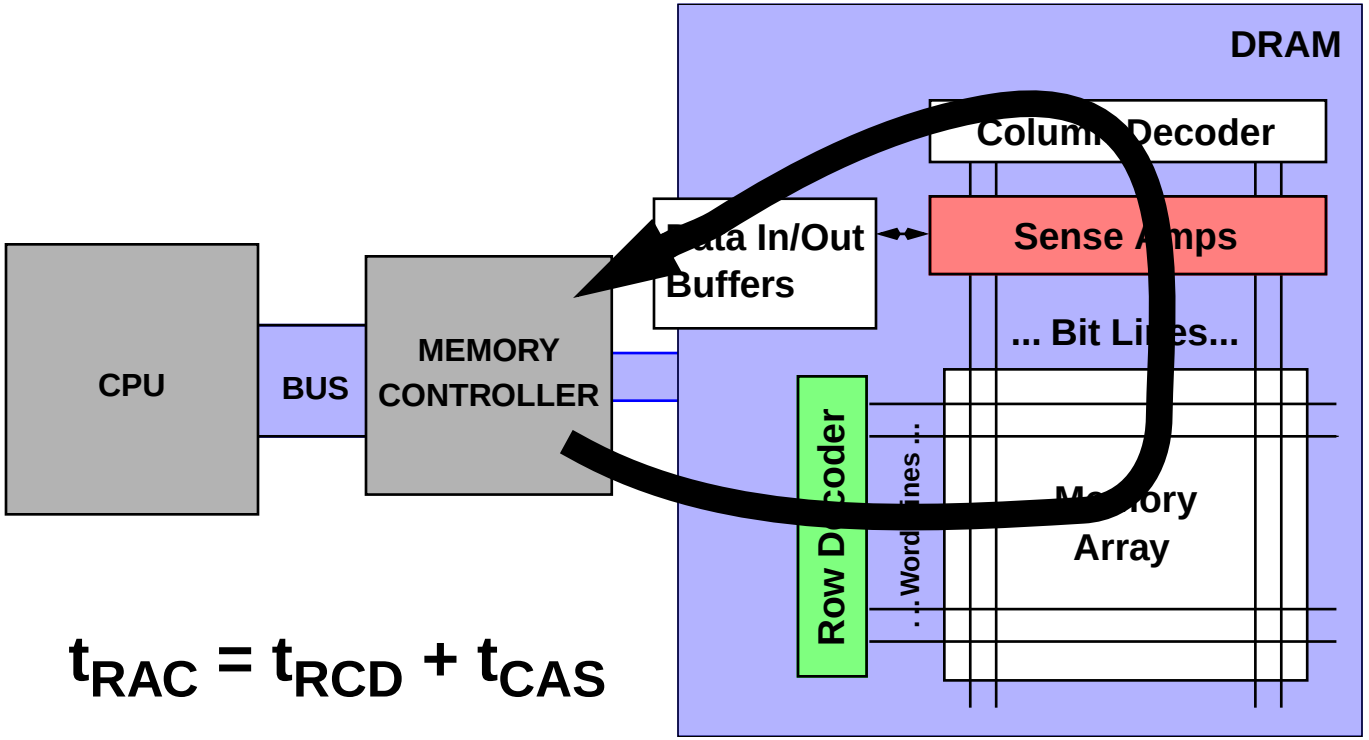

$$t_{RAC} = t_{RCD} + t_{CAS}$$

RAC (Random Access Delay)

DRAM "Speed" Part IV

How fast can I precharge DRAM array so I can engage another RAS?

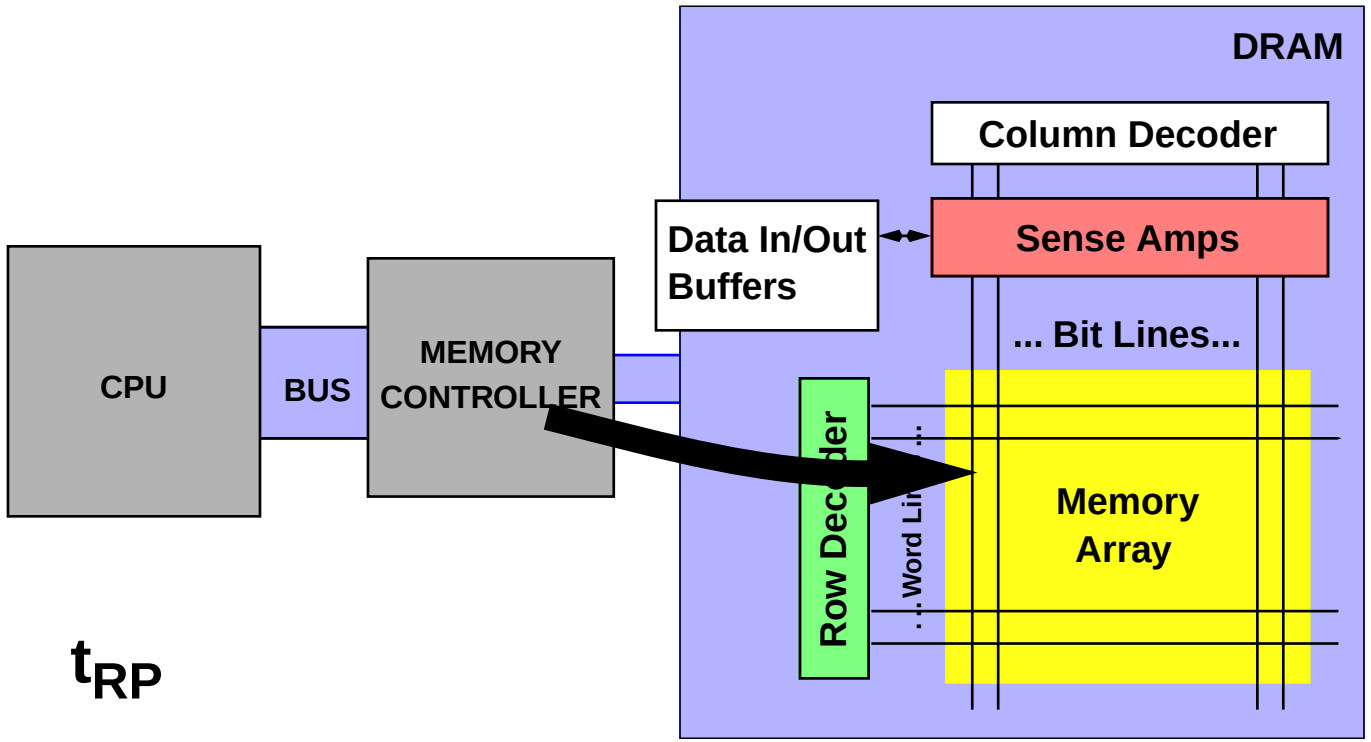

t_{RP}

RP (Row Precharge Delay)

DRAM "Speed" Part V

How fast can I read data from two different rows?

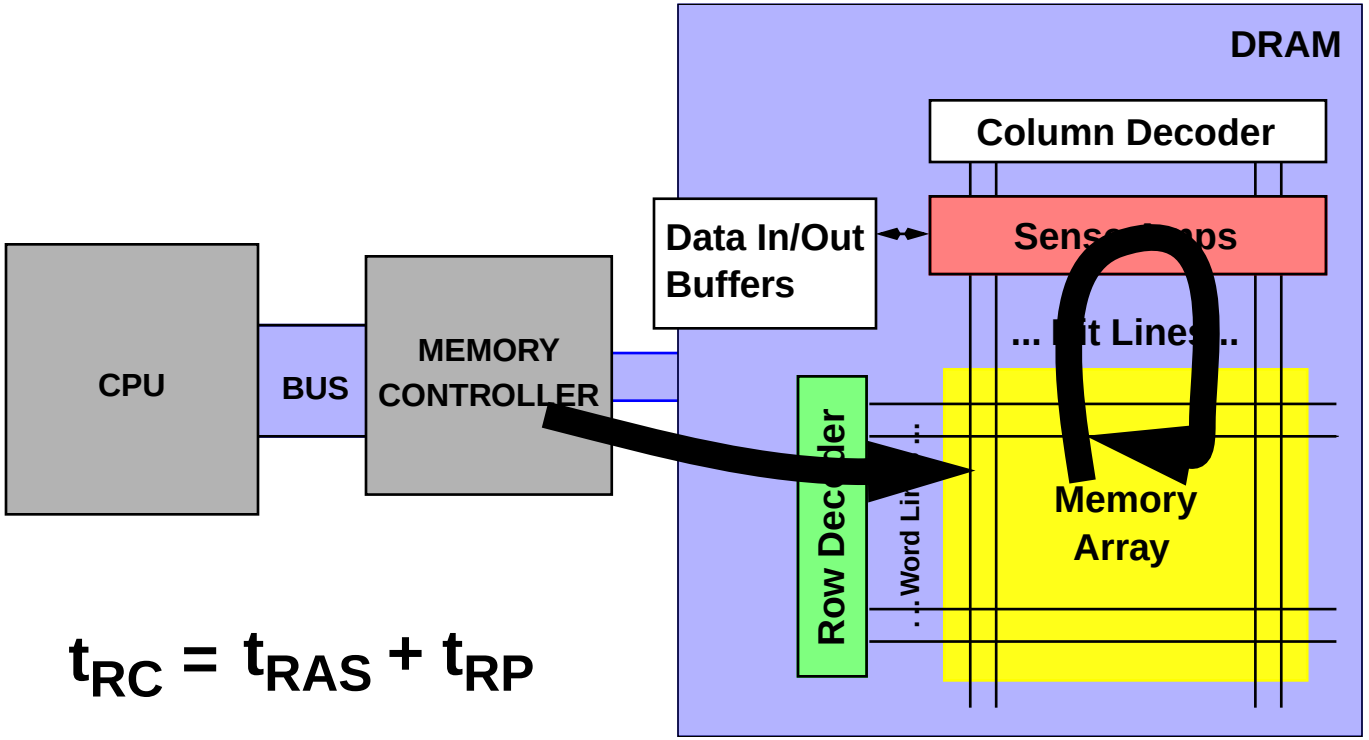

$$t_{RC} = t_{RAS} + t_{RP}$$

RC (Row Cycle Time)

DRAM "Speed" Summary I

What do I care about?

t_{RCD}

t_{CAS} ← Seen in ads. Easy to explain Easy to sell

t_{RP}

$t_{RC} = t_{RAS} + t_{RP}$ ← Embedded systems designers DRAM manufactuers

$t_{RAC} = t_{RCD} + t_{CAS}$ ← Computer Architect: Latency bound code i.e. linked list traversal

RAS: Row Address Strobe

CAS: Column Address Strobe

RCD: Row Command Delay

RAC :Random Access Delay

RP :Row Precharge Delay

RC :Row Cycle Time

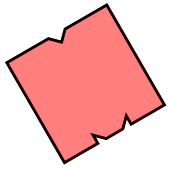

DRAM "Speed" Summary II

DRAM Type	Frequency	Data Bus Width (per chip)	Peak Data Bandwidth (per Chip)	Random Access Time (t_{RAC})	Row Cycle Time (t_{RC})
PC133 SDRAM	133	16	200 MB/s	45 ns	60 ns
DDR 266	133 * 2	16	532 MB/s	45 ns	60 ns
PC800 RDRAM	400 * 2	16	1.6 GB/s	60 ns	70 ns
FCRAM	200 * 2	16	0.8 GB/s	25 ns	25 ns
RLDRAM	300 * 2	32	2.4 GB/s	25 ns	25 ns

data: Dec. 2002

DRAM is "slow"
But doesn't have to be
t_{RC} < 10ns achievable

Higher die cost → Not adopted in standard

Not commodity → Expensive

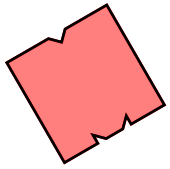

Signal Propagation

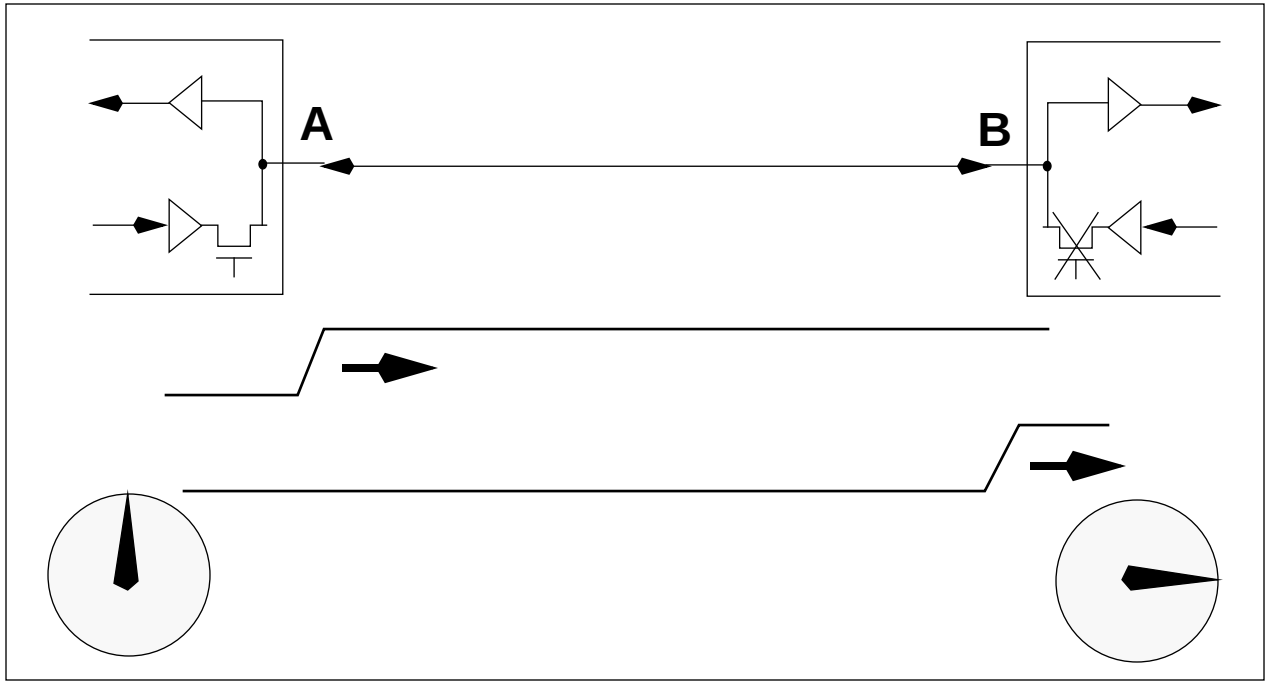

Ideal Transmission Line

~ 0.66c = 20 cm/ns

PC Board + Module Connectors +
Varying Electrical Loads

= Rather non-Ideal Transmission Line

DRAM Interface: Protocol

The Digital Fantasy

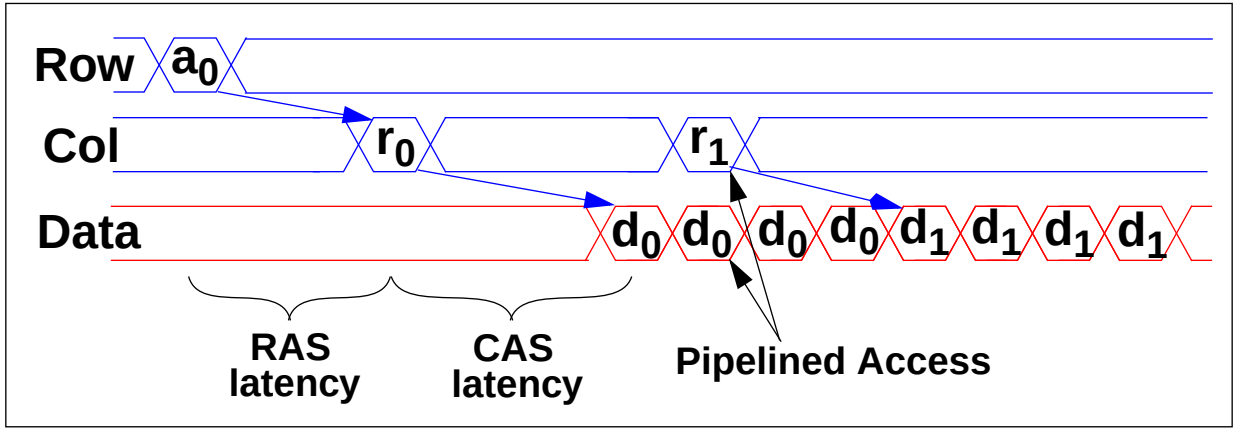

Pretend that the world looks like this

But...

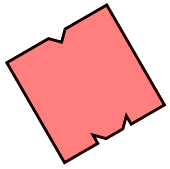

DRAM Interface: Signals

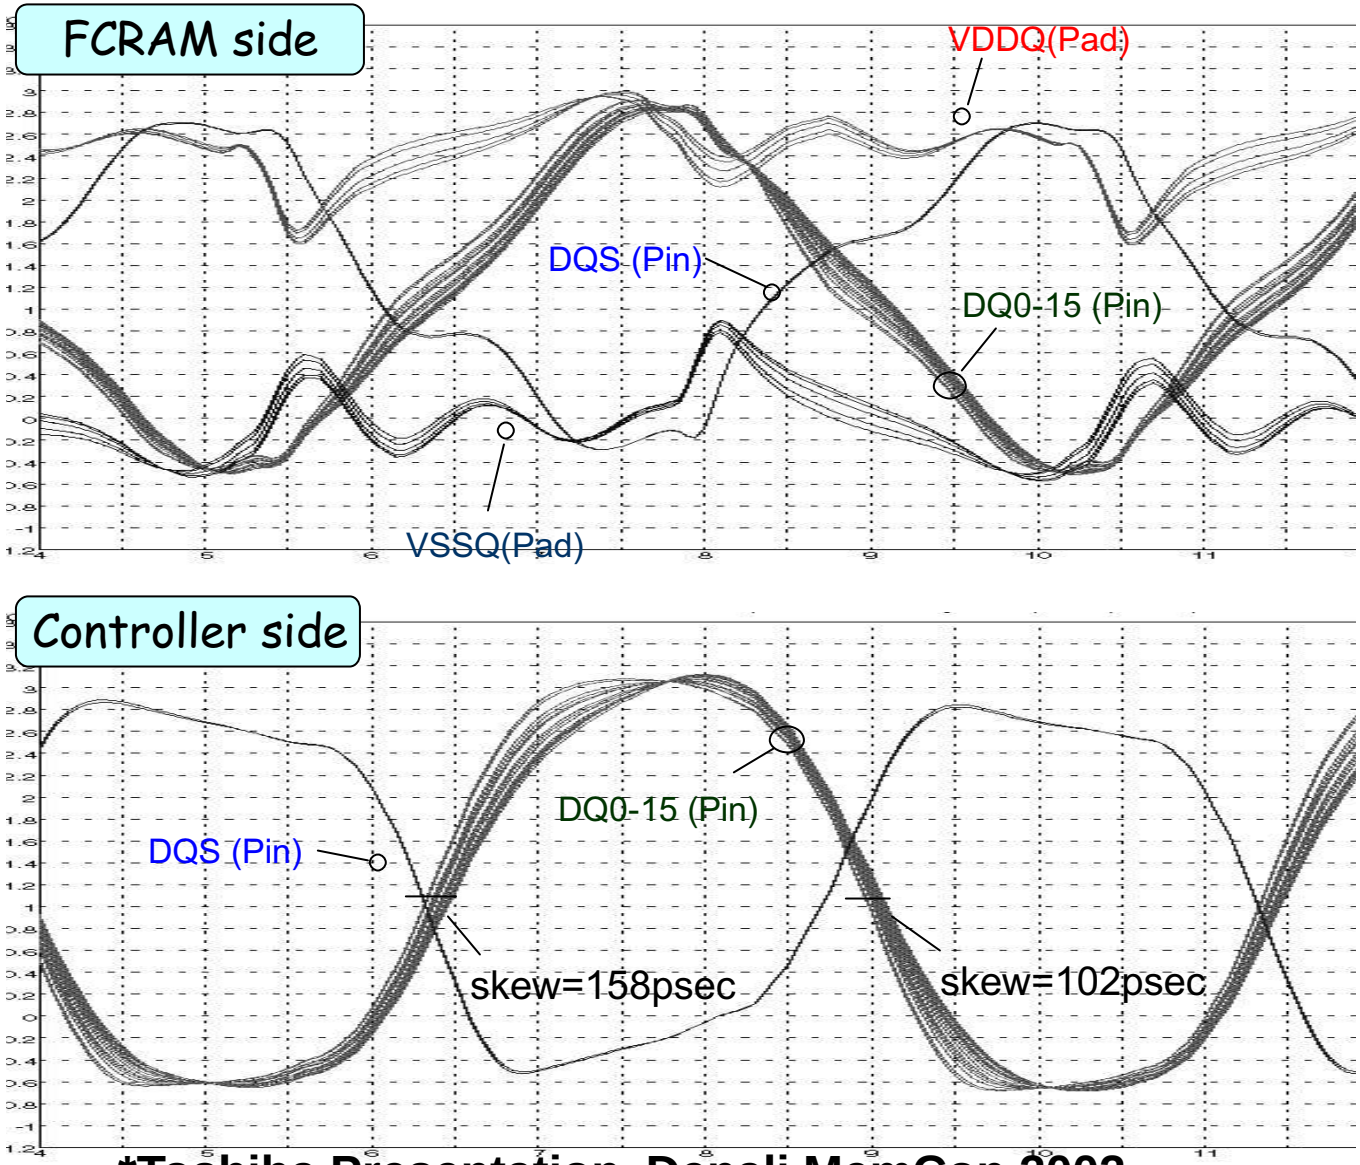

***Toshiba Presentation, Denali MemCon 2002**

Interface: Clocking Issues

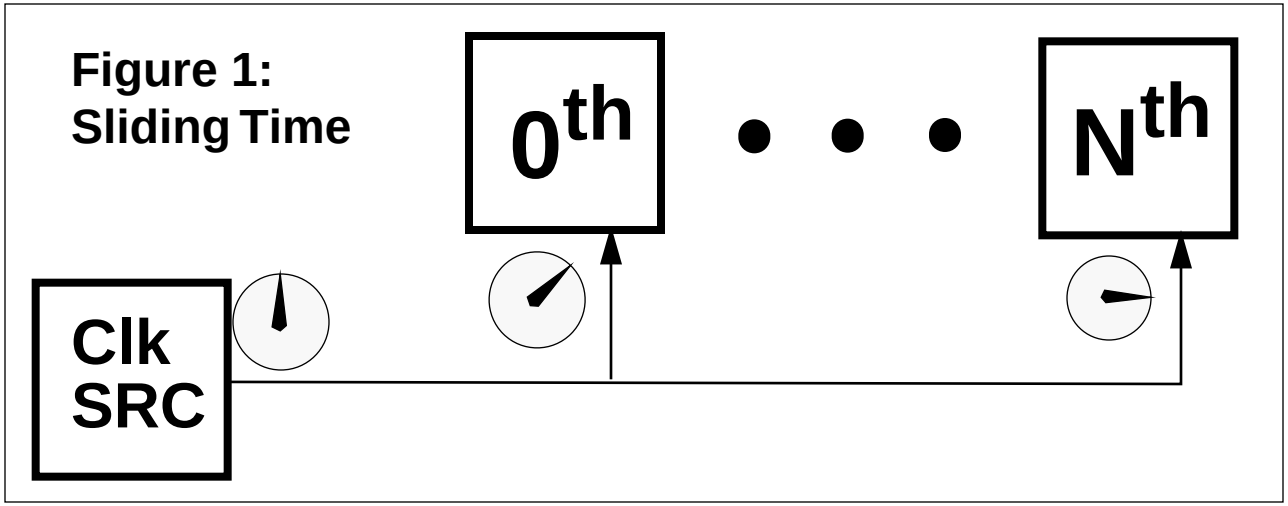

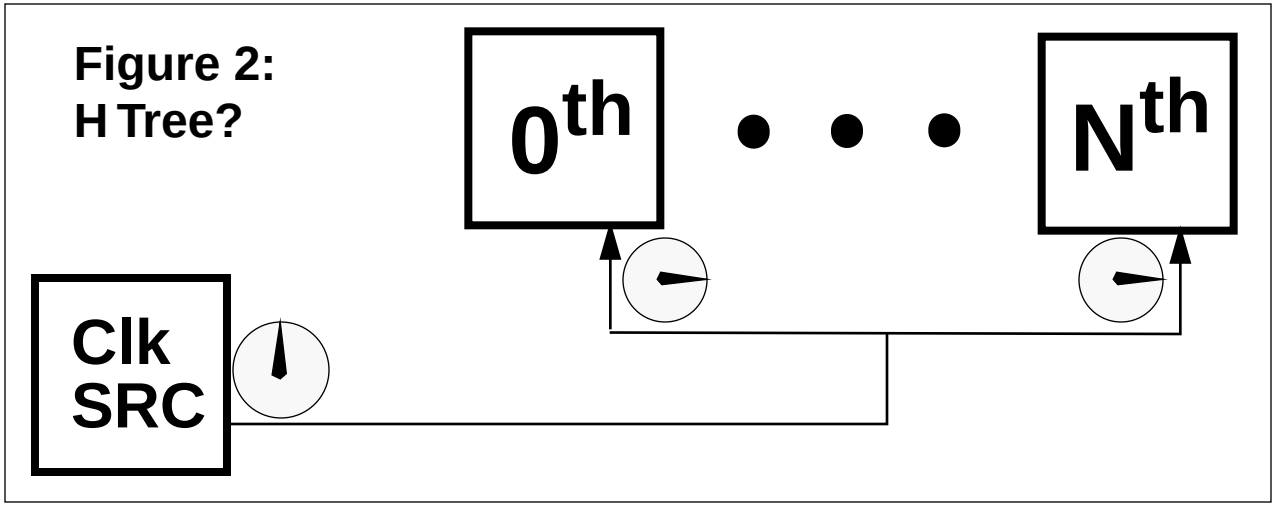

What Kind of Clocking System?

Path Length Differential

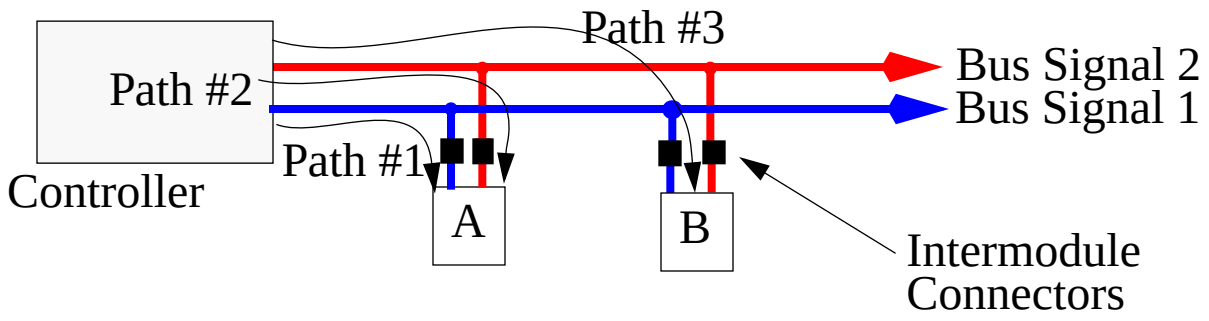

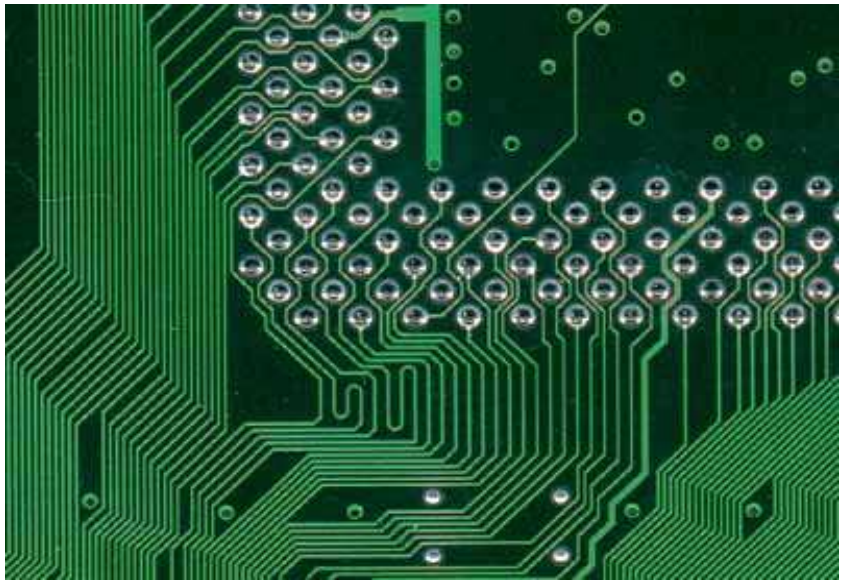

High Frequency AND Wide Parallel Busses are Difficult to Implement

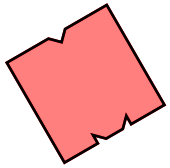

Timing Variations

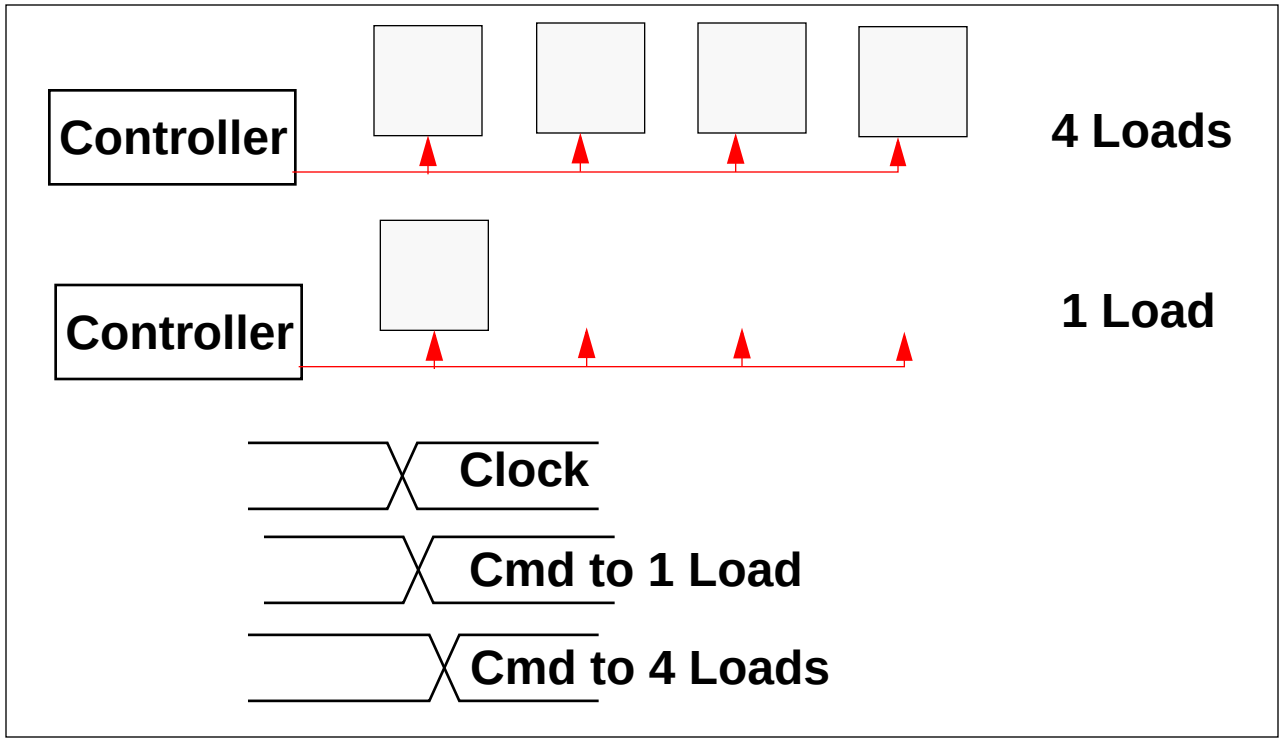

How many DIMMs in System?

How many devices on each DIMM?

Who built the memory module?

Infinite variations on timing!

Topology

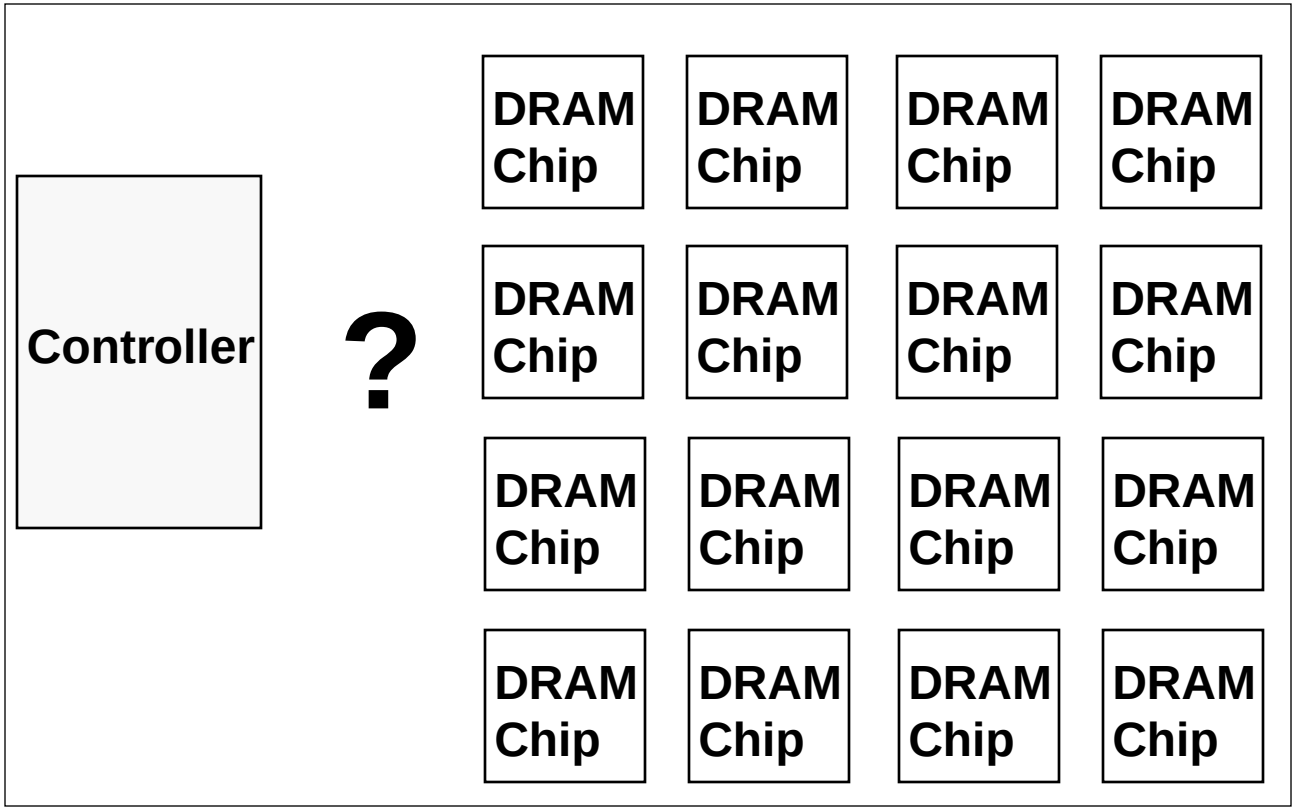

DRAM System Topology Determines

Electrical Loading Conditions

and Signal Propagation Lengths

SDRAM Topology Example

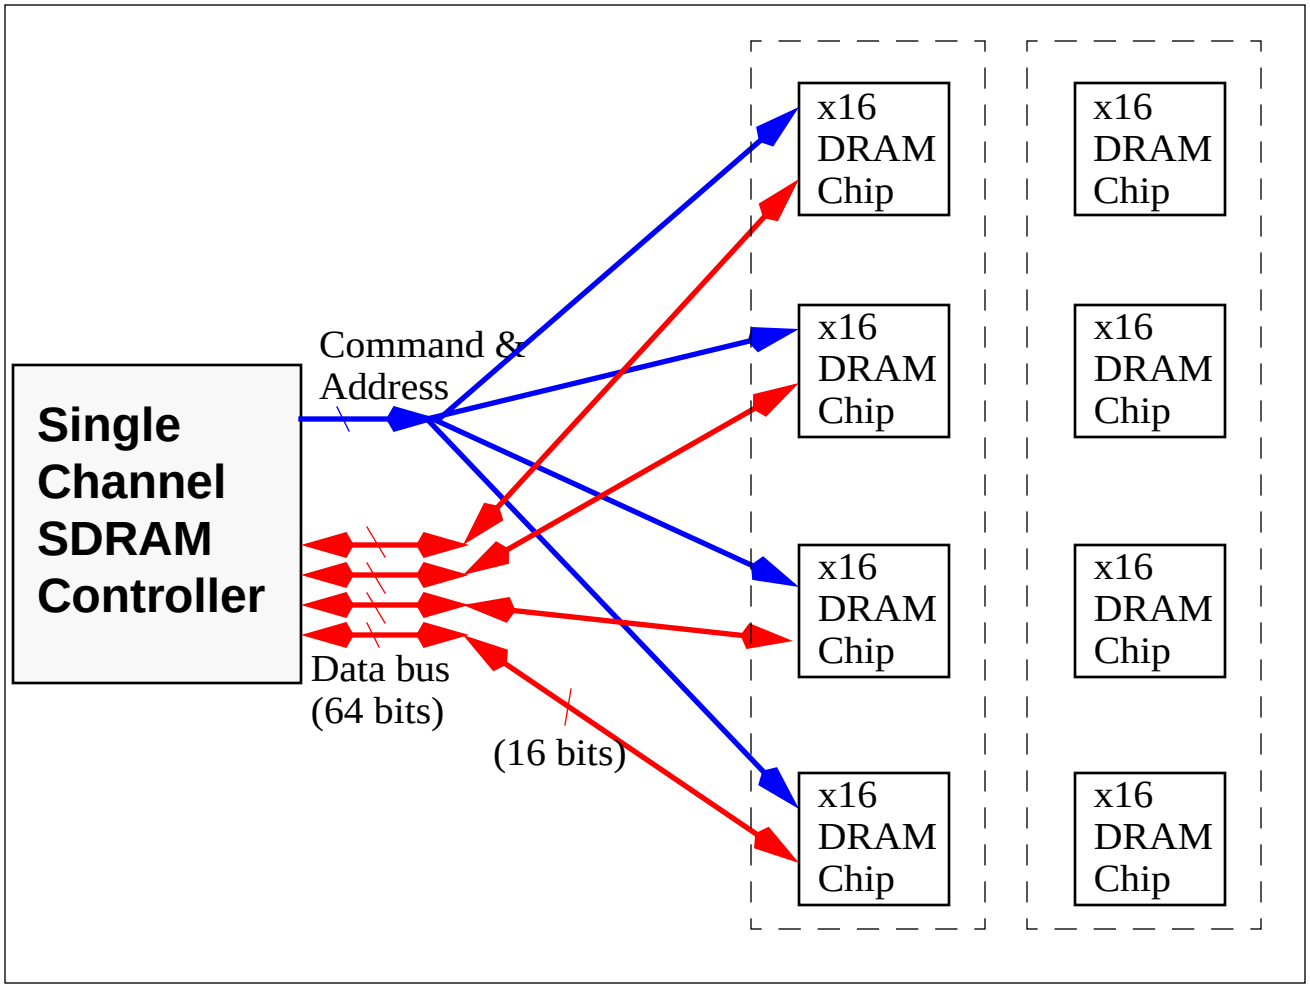

Loading Imbalance

SDRAM Topology Example II

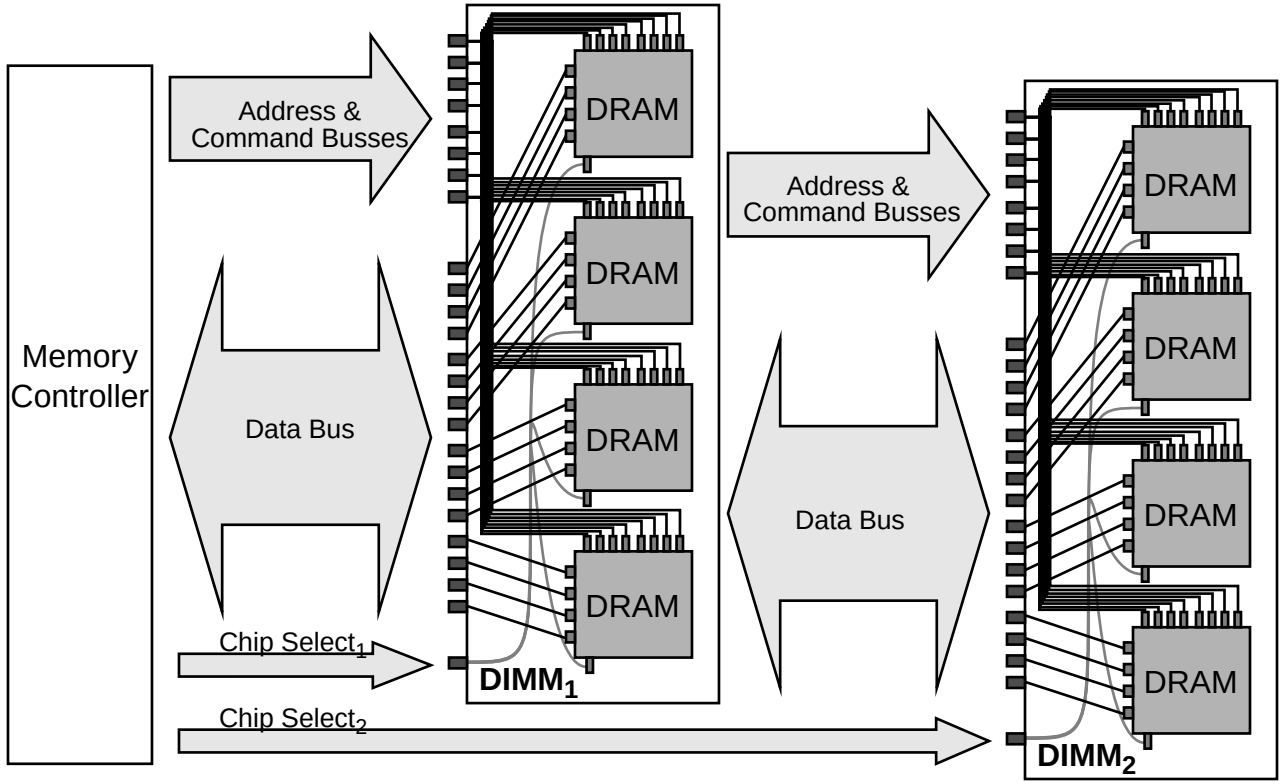

(Same topology, different drawing, a little more detail)

RDRAM Topology Example

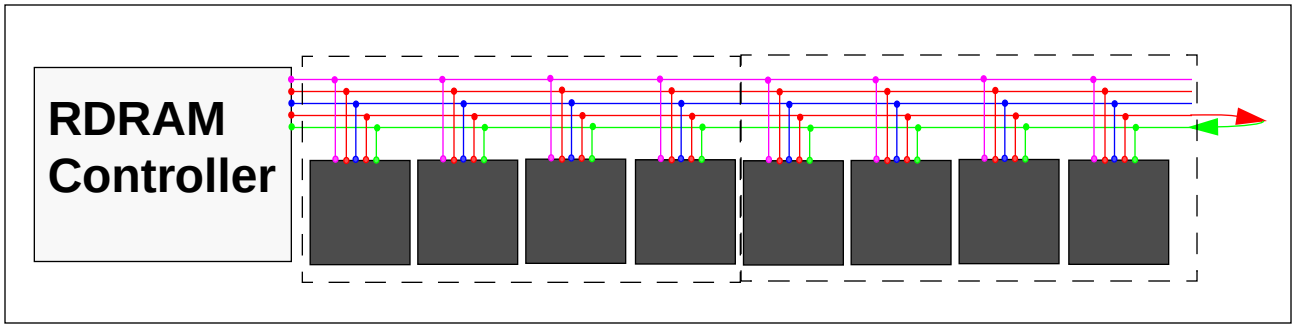

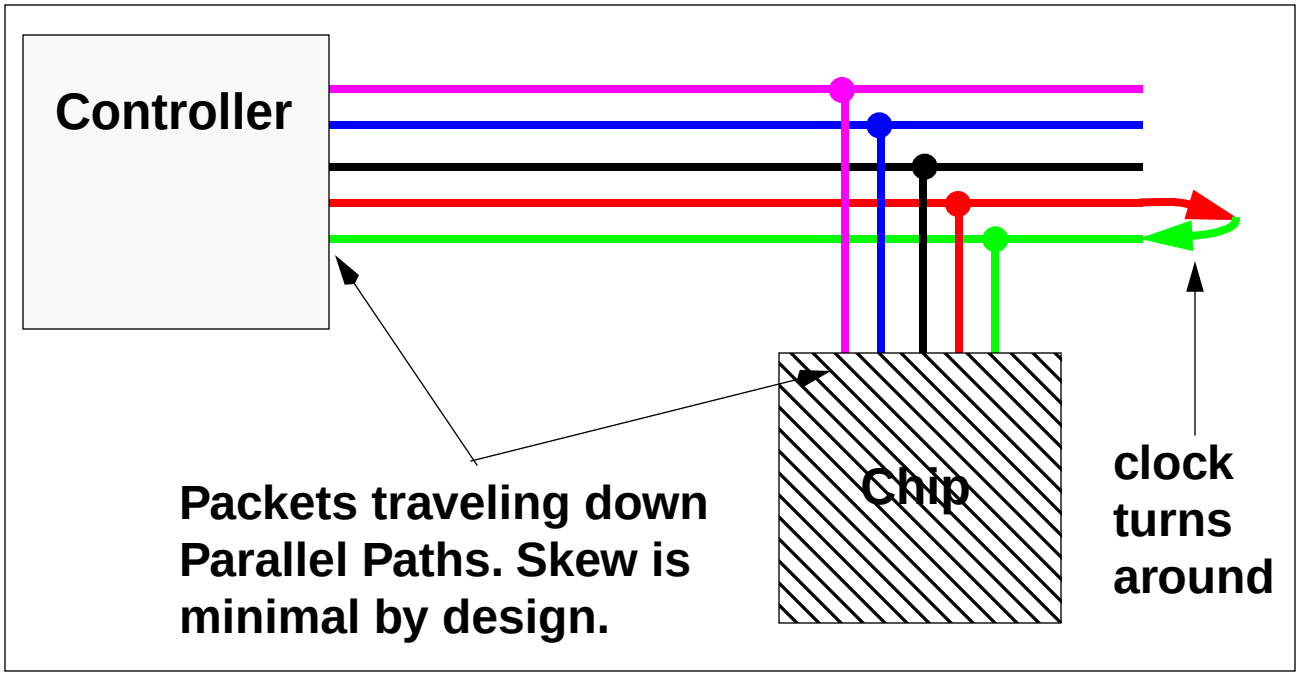

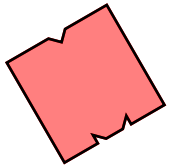

I/O - Differential Pair

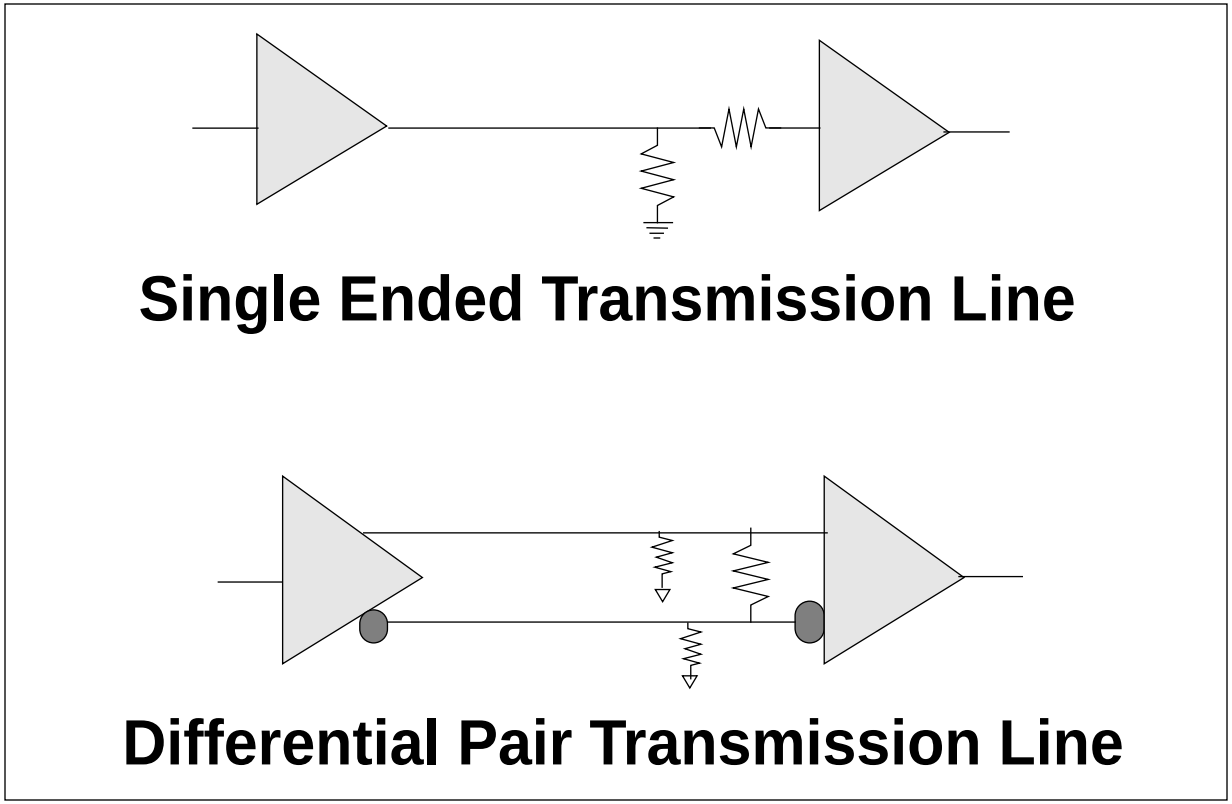

Increase Rate of bits/s/pin ?

Cost Per Pin?

Pin Count?

I/O - Multi Level Logic

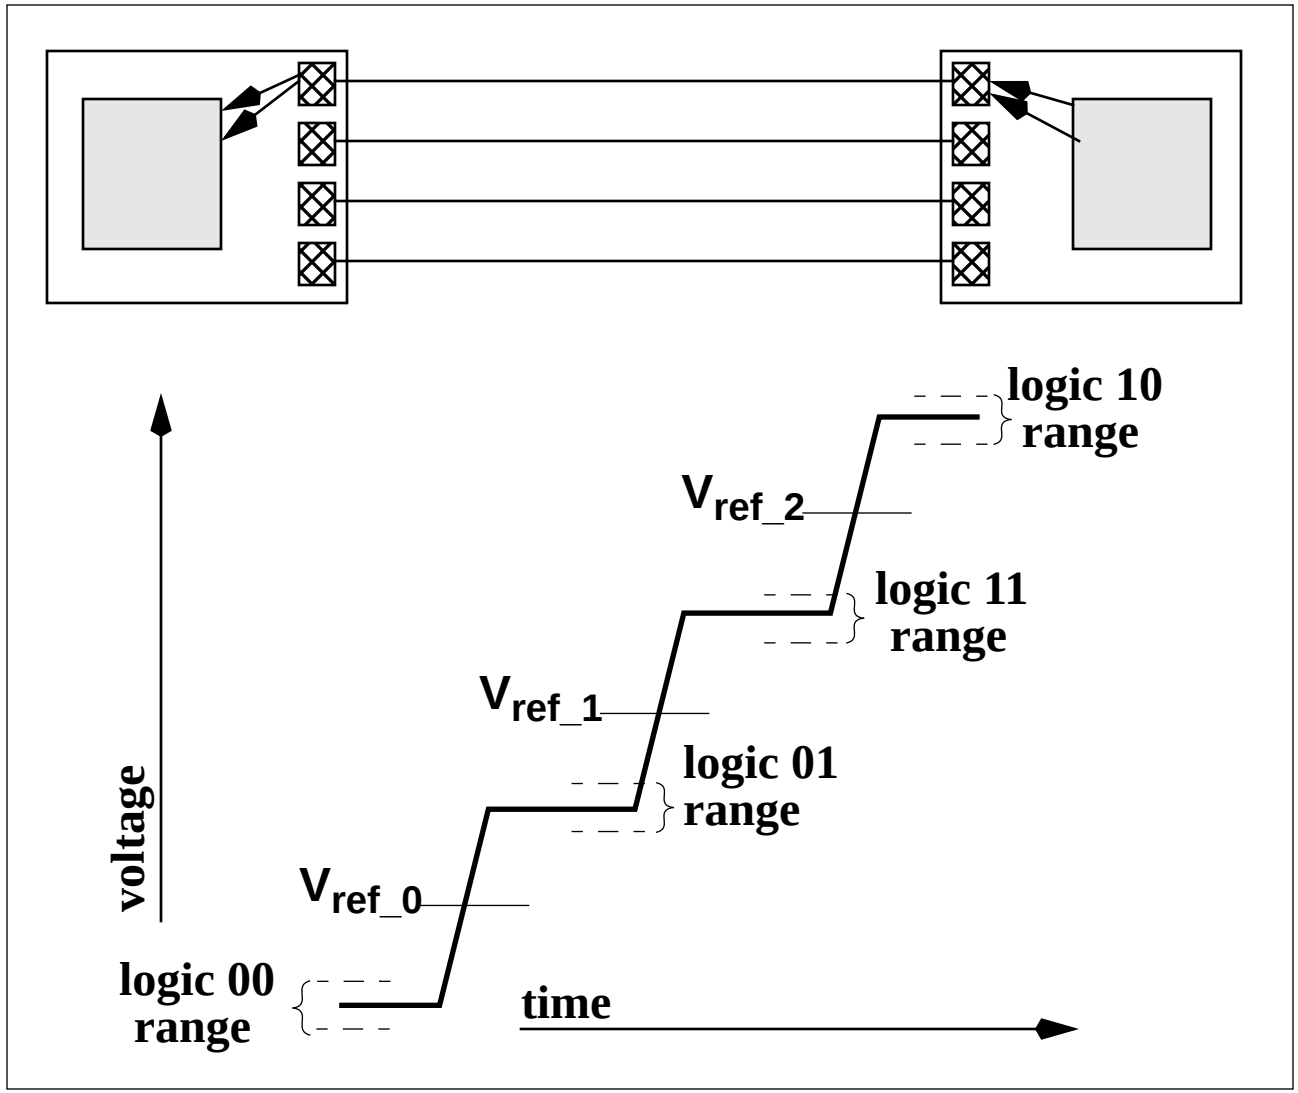

Increase Rate of bits/s/pin

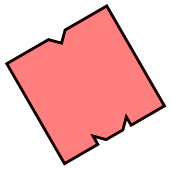

Packaging

DIP
"good old days"

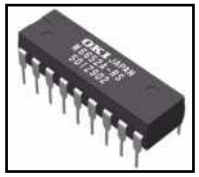

SOJ
Small Outline J-lead

TSOP
Thin Small Outline Package

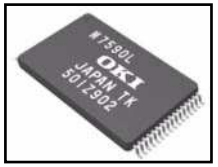

LQFP
Low Profile Quad Flat Package

FBGA
Fine Ball Grid Array

Features	Target Specification	
Package	FBGA	LQFP
Speed	800MBps	550Mbps
Vdd/Vddq	2.5V/2.5V (1.8V)	
Interface	SSTL_2	
Row Cycle Time t_{RC}	35ns	

Memory Roadmap for Hynix NetDDR II

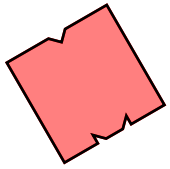

Access Protocol

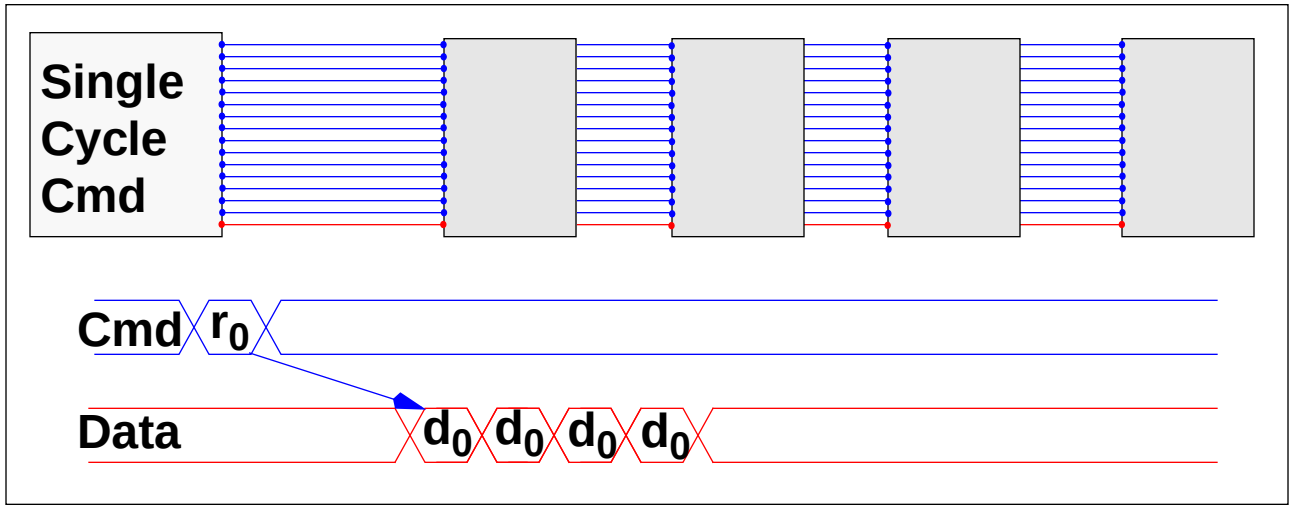

Single Cycle Command

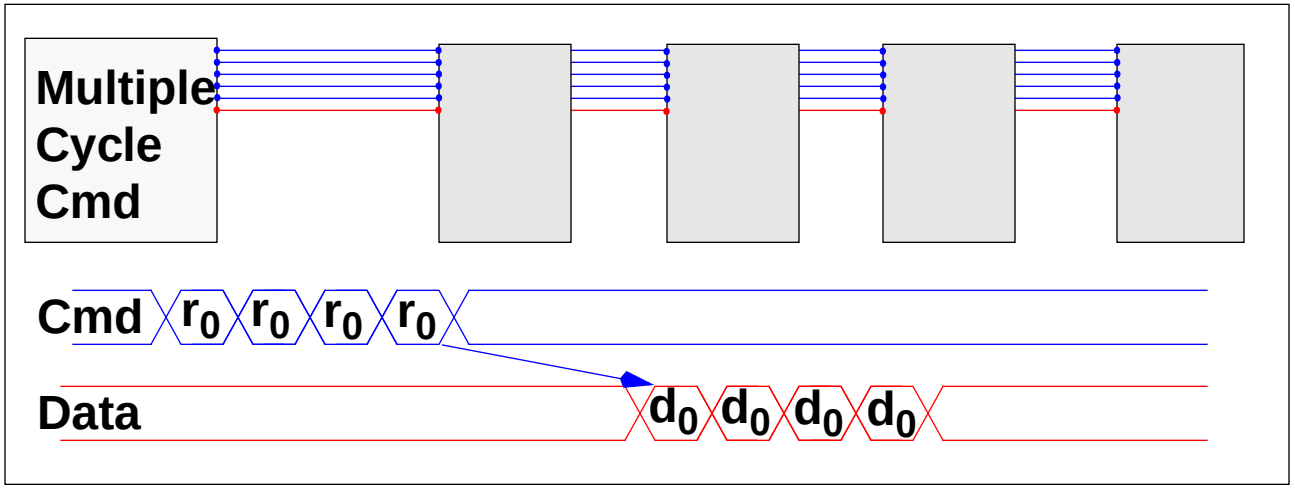

Multiple Cycle Command

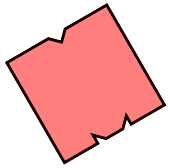

Access Protocol (r/r)

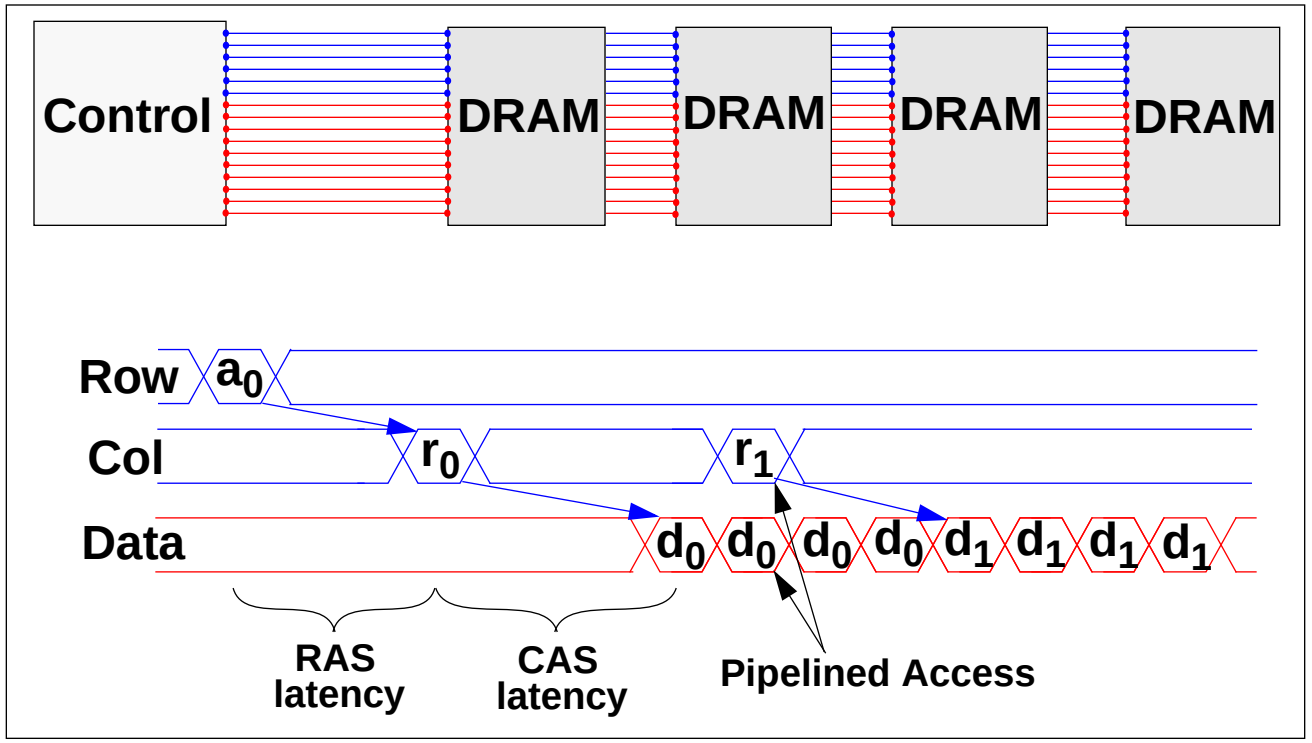

Consecutive Cache Line Read Requests to Same DRAM Row

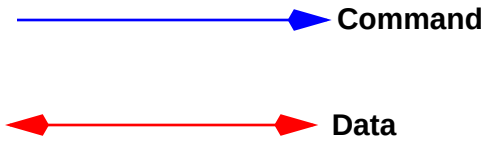

a = Active (open page)
r = Read (Column Read)
d = Data (Data chunk)

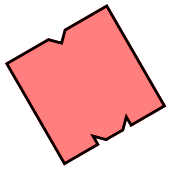

Access Protocol (r/w)

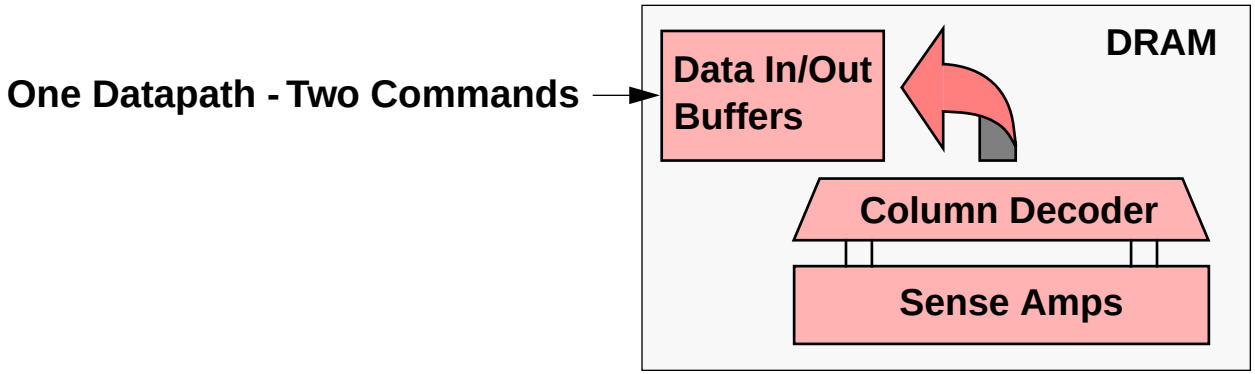

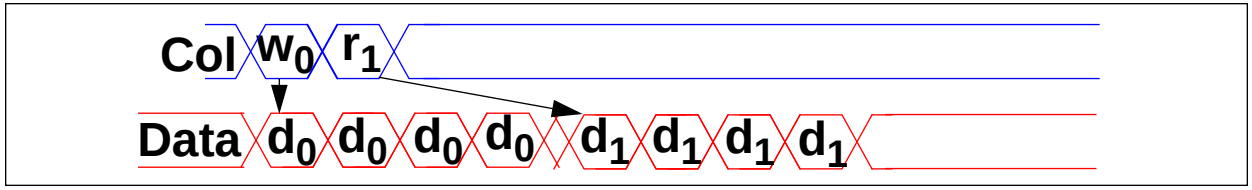

Case 1: Read Following a Write Command to Different DRAM Devices

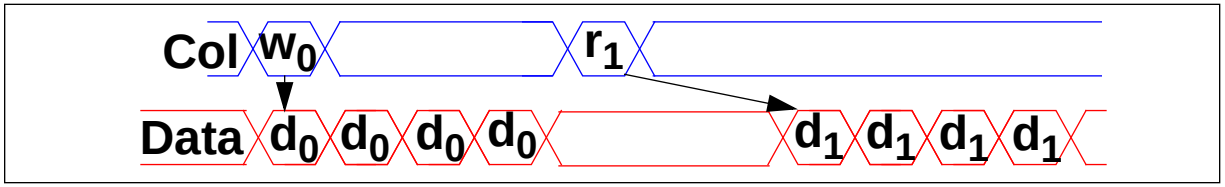

Case 2: Read Following a Write Command to Same DRAM Device

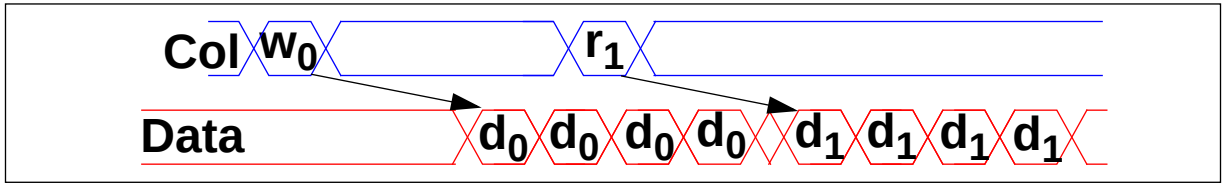

Soln: Delay Data of Write Command to match Read Latency

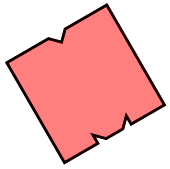

Address Mapping

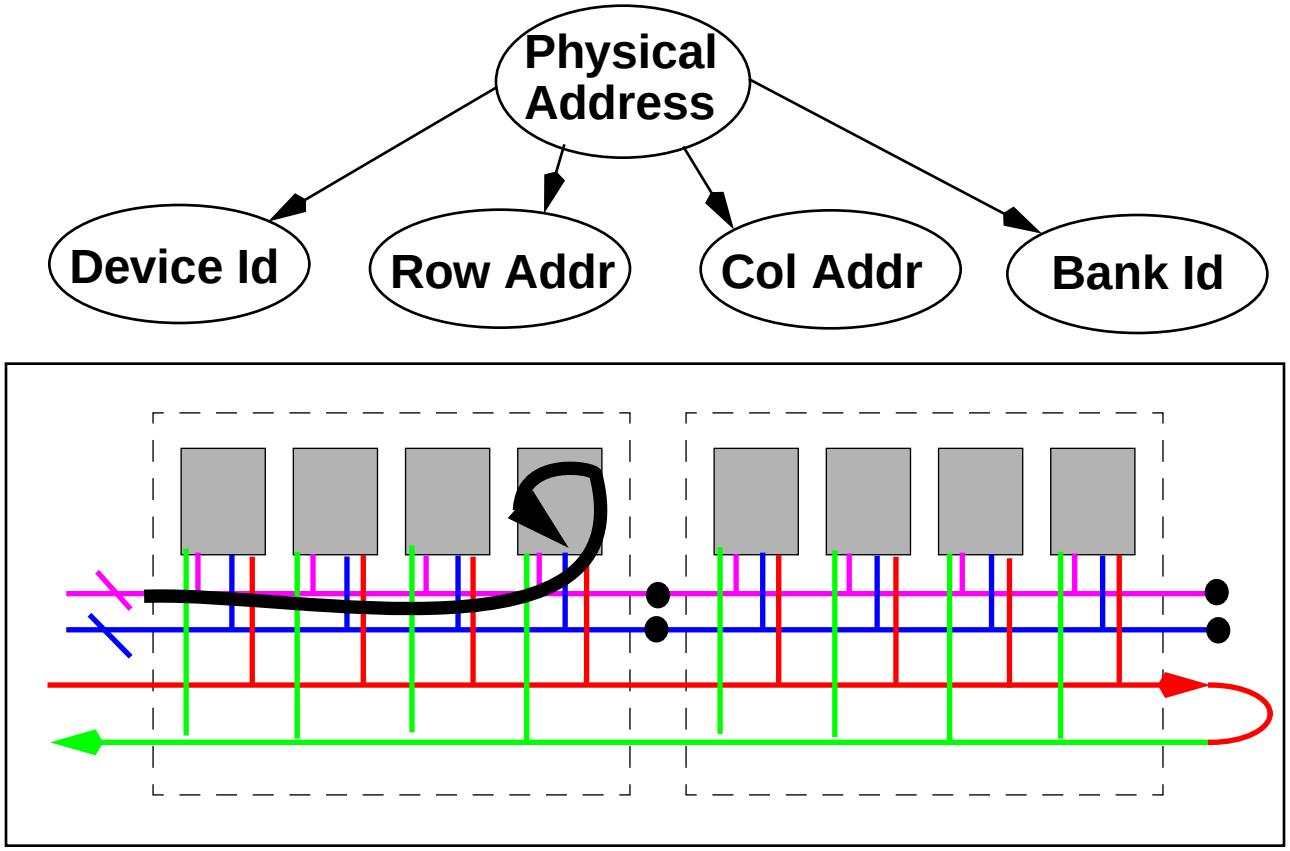

Access Distribution for Temp Control

Avoid Bank Conflicts

Access Reordering for performance

Example: Bank Conflicts

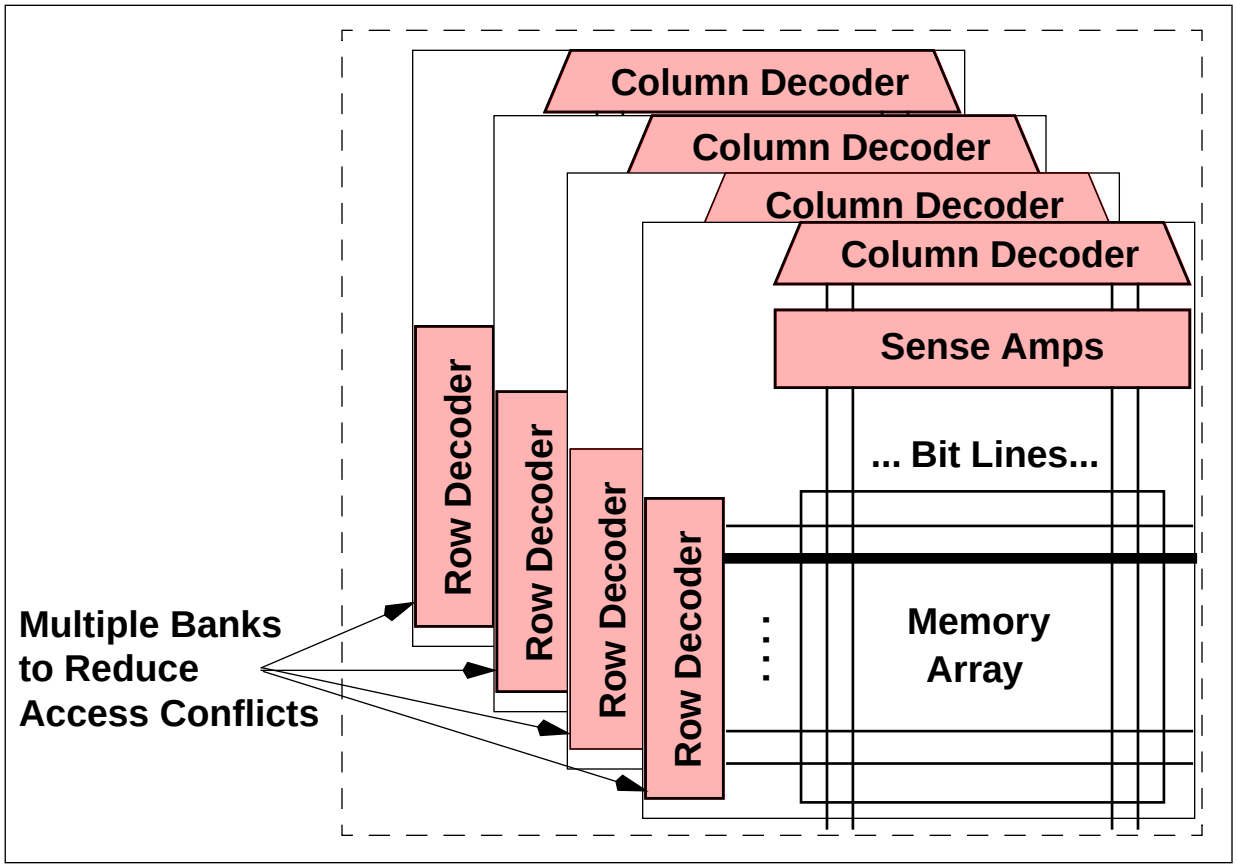

Read 05AE5700 → Device id 3, Row id 266, Bank id 0
Read 023BB880 → Device id 3, Row id 1BA, Bank id 0
Read 05AE5780 → Device id 3, Row id 266, Bank id 0
Read 00CBA2C0 → Device id 3, Row id 052, Bank id 1

More Banks per Chip == Performance == Logic Overhead

Example: Access Reordering

1 Read 05AE5700 → Device id 3, Row id 266, Bank id 0
2 Read 023BB880 → Device id 3, Row id 1BA, Bank id 0
3 Read 05AE5780 → Device id 3, Row id 266, Bank id 0
4 Read 00CBA2C0 → Device id 1, Row id 052, Bank id 1

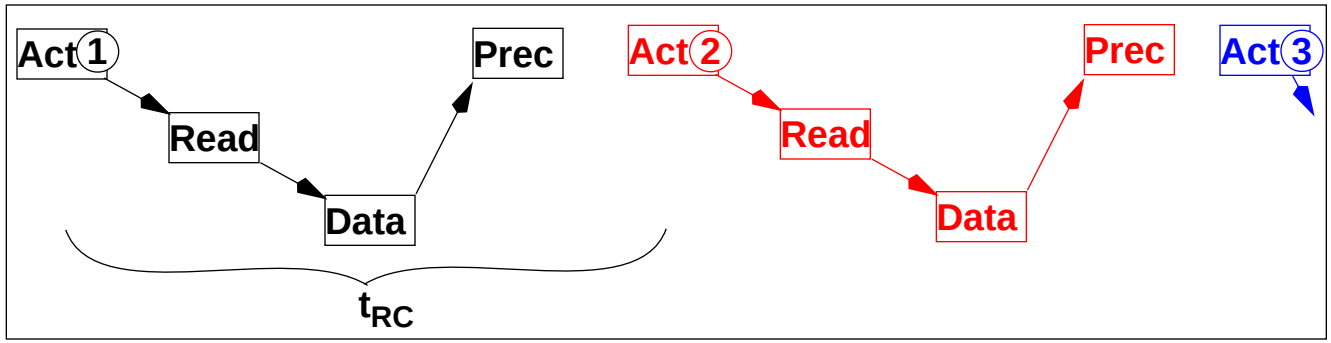

Strict Ordering

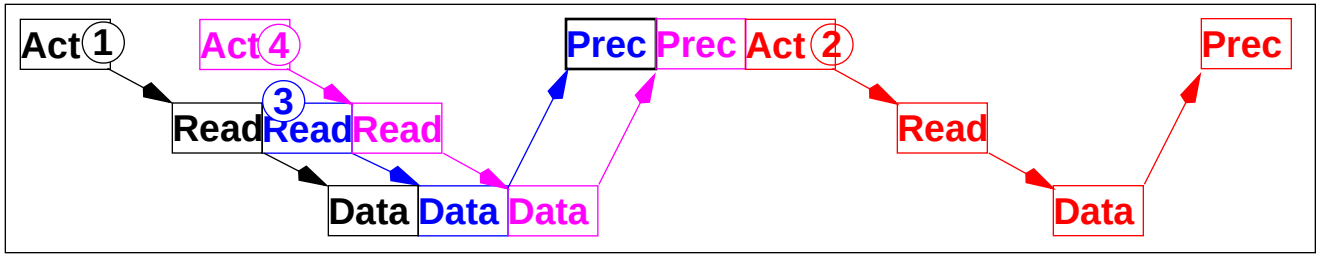

Memory Access Re-ordered

Act = Activate Page (Data moved from DRAM cells to row buffer)
Read = Read Data (Data moved from row buffer to memory controller)
Prec = Precharge (close page/evict data in row buffer/sense amp)

Technology Roadmap (ITRS)

	2004	2007	2010	2013	2016
Semi Generation (nm)	**90**	**65**	**45**	**32**	**22**
CPU MHz	**3990**	**6740**	**12000**	**19000**	**29000**
MLogicTransistors/cm^2	**77.2**	**154.3**	**309**	**617**	**1235**
High Perf chip pin count	**2263**	**3012**	**4009**	**5335**	**7100**
High Performance chip cost (cents/pin)	**1.88**	**1.61**	**1.68**	**1.44**	**1.22**
Memory pin cost (cents/pin)	**0.34 - 1.39**	**0.27 - 0.84**	**0.22 - 0.34**	**0.19 - 0.39**	**0.19 - 0.33**
Memory pin count	**48-160**	**48-160**	**62-208**	**81-270**	**105-351**

Trend: **Free Transistors & Costly Interconnects**

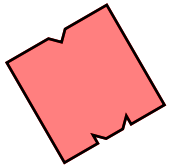

Choices for Future

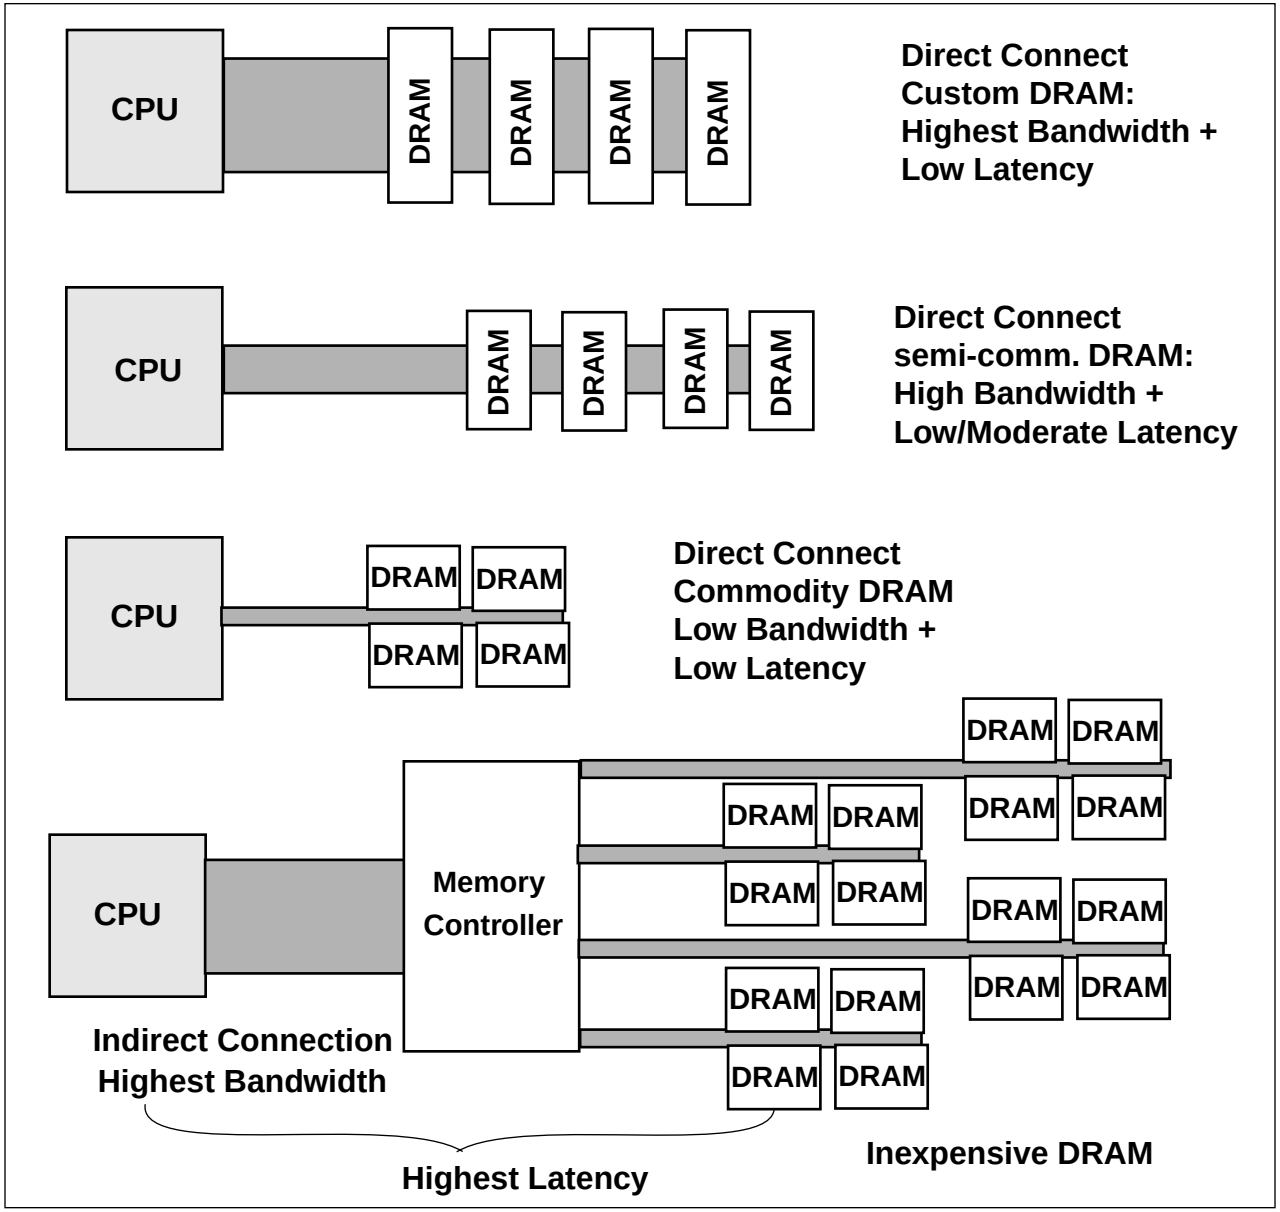

Bruce Jacob

University of
Maryland
ECE Dept.

DRAM Evolutionary Tree

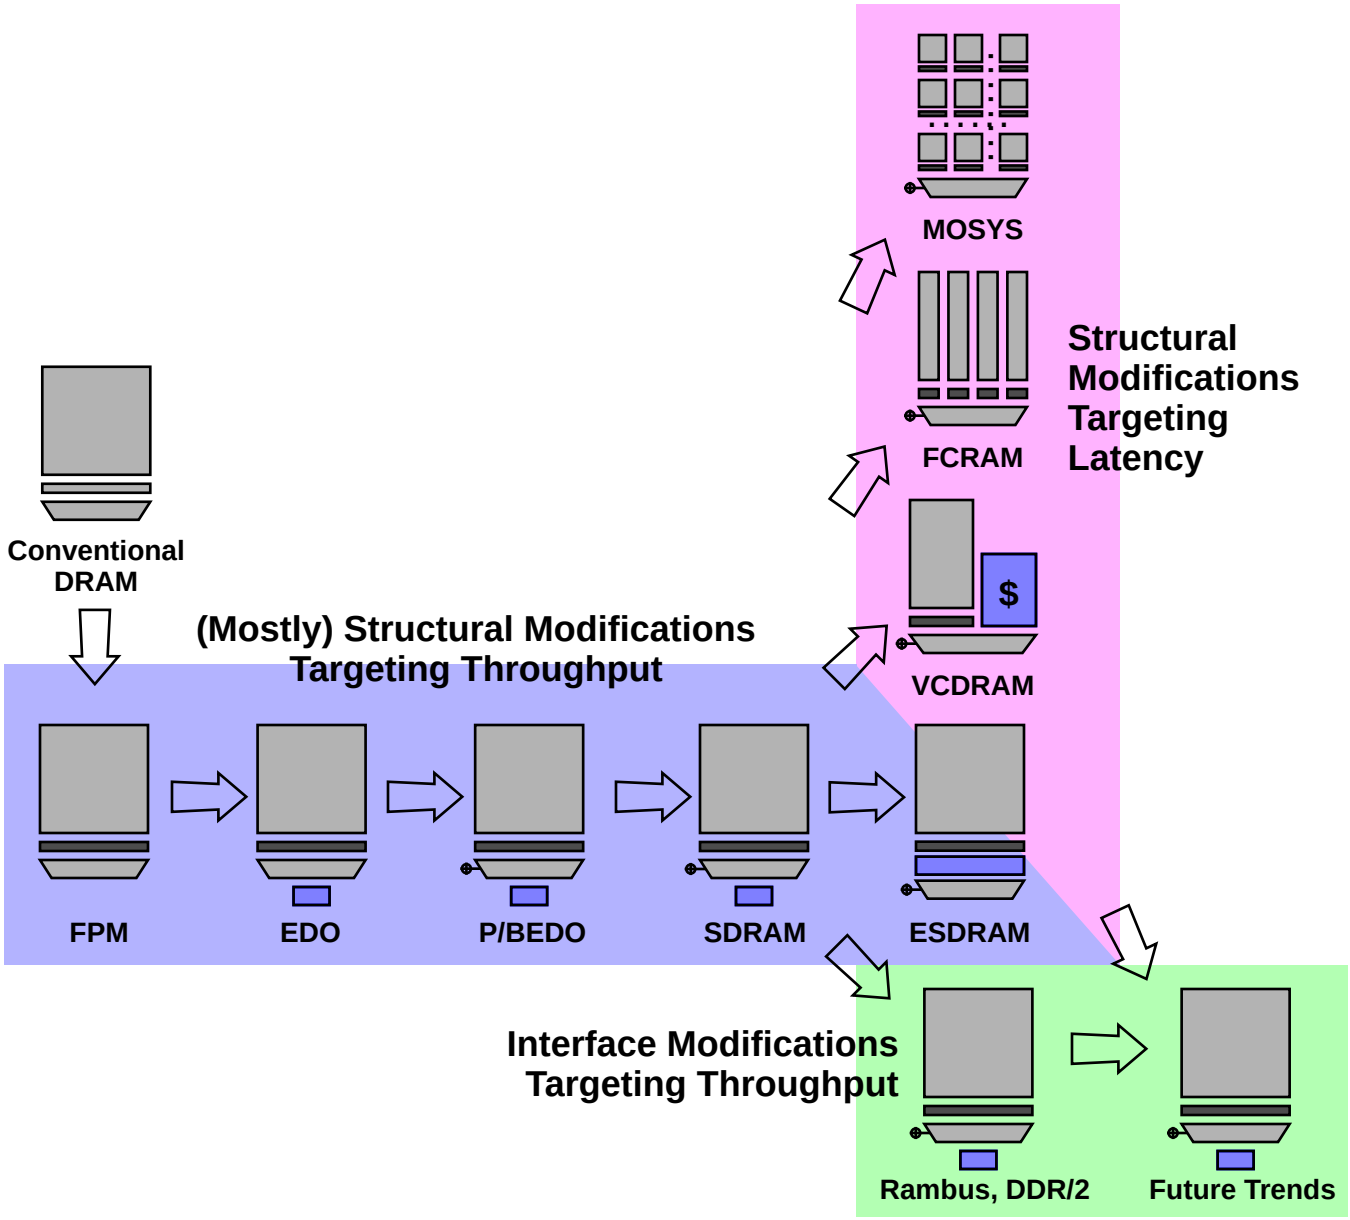

Bruce Jacob

University of
Maryland
ECE Dept.

DRAM Evolution

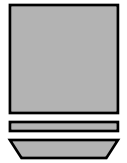

Read Timing for Conventional DRAM

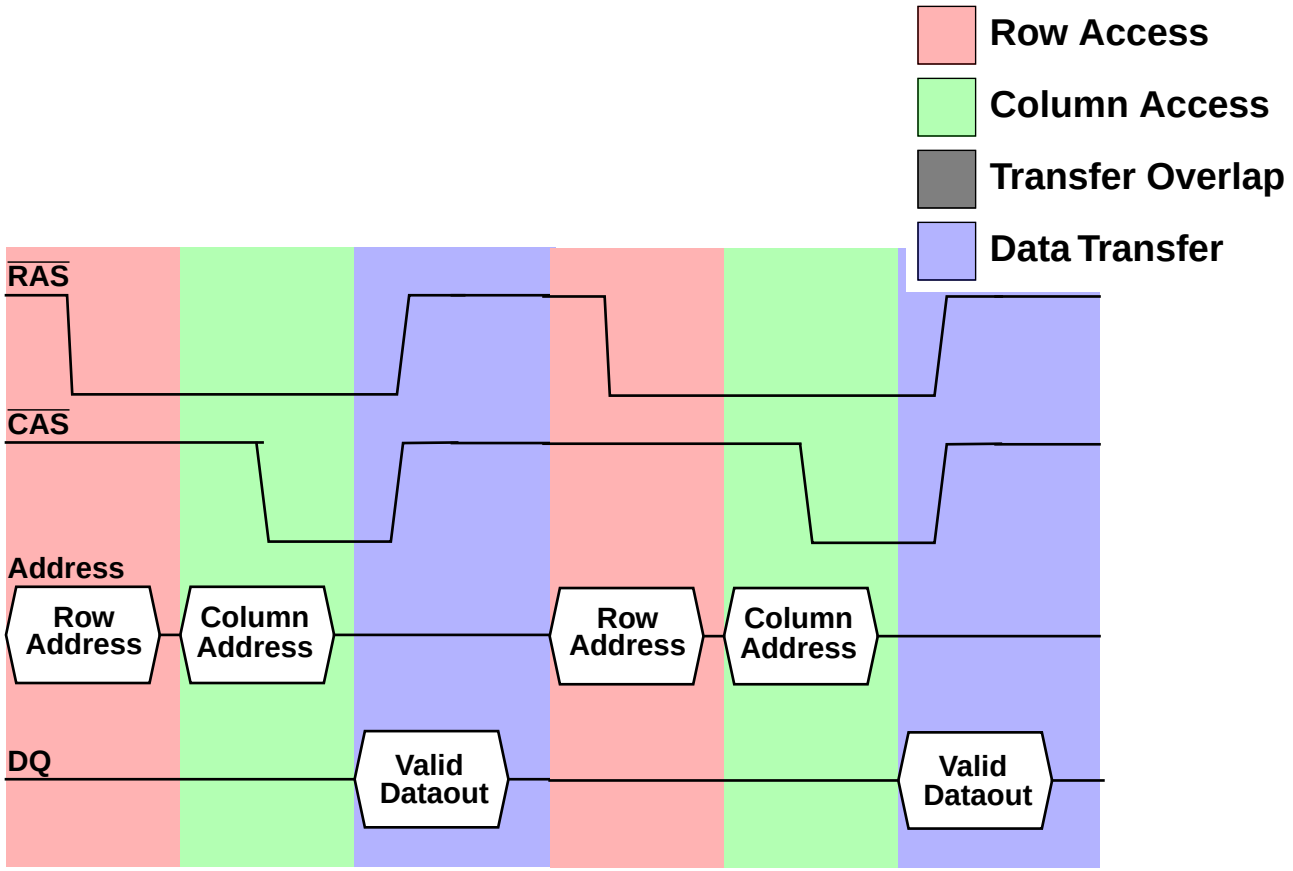

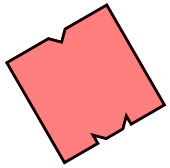

Bruce Jacob

University of
Maryland
ECE Dept.

DRAM Evolution

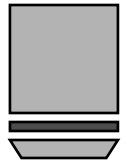

Read Timing for Fast Page Mode

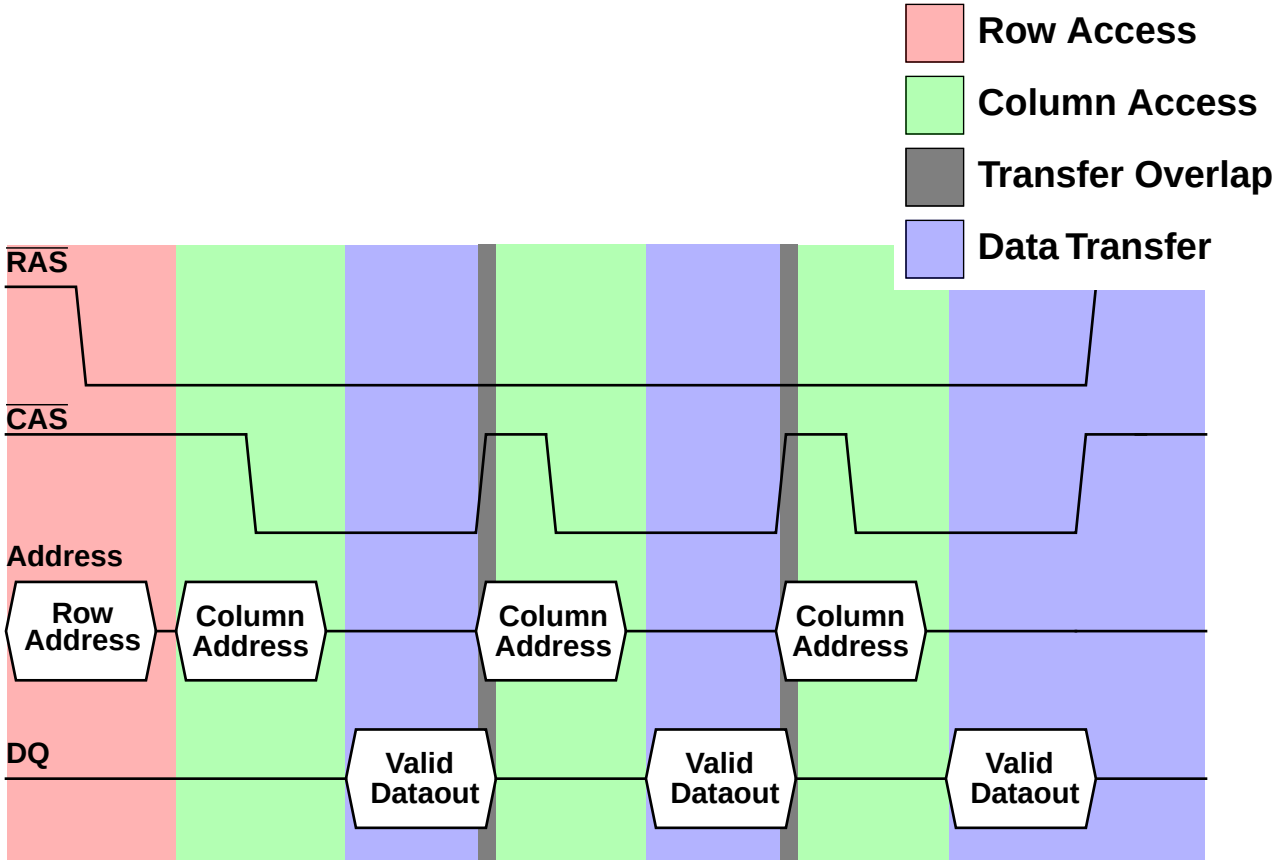

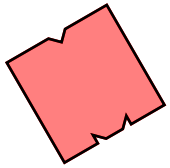

DRAM Evolution

Read Timing for Extended Data Out

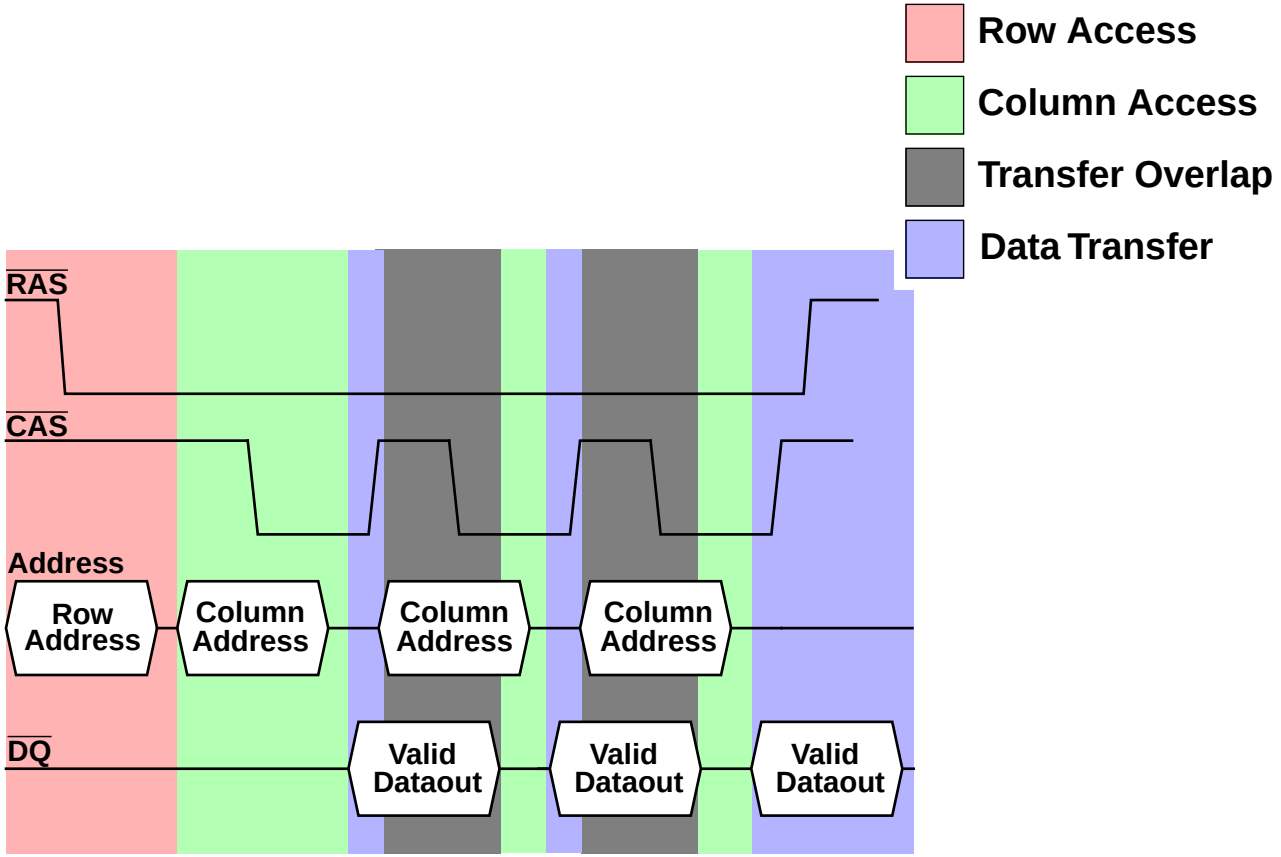

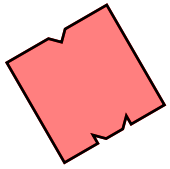

Bruce Jacob

University of
Maryland
ECE Dept.

DRAM Evolution

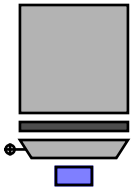

Read Timing for Burst EDO

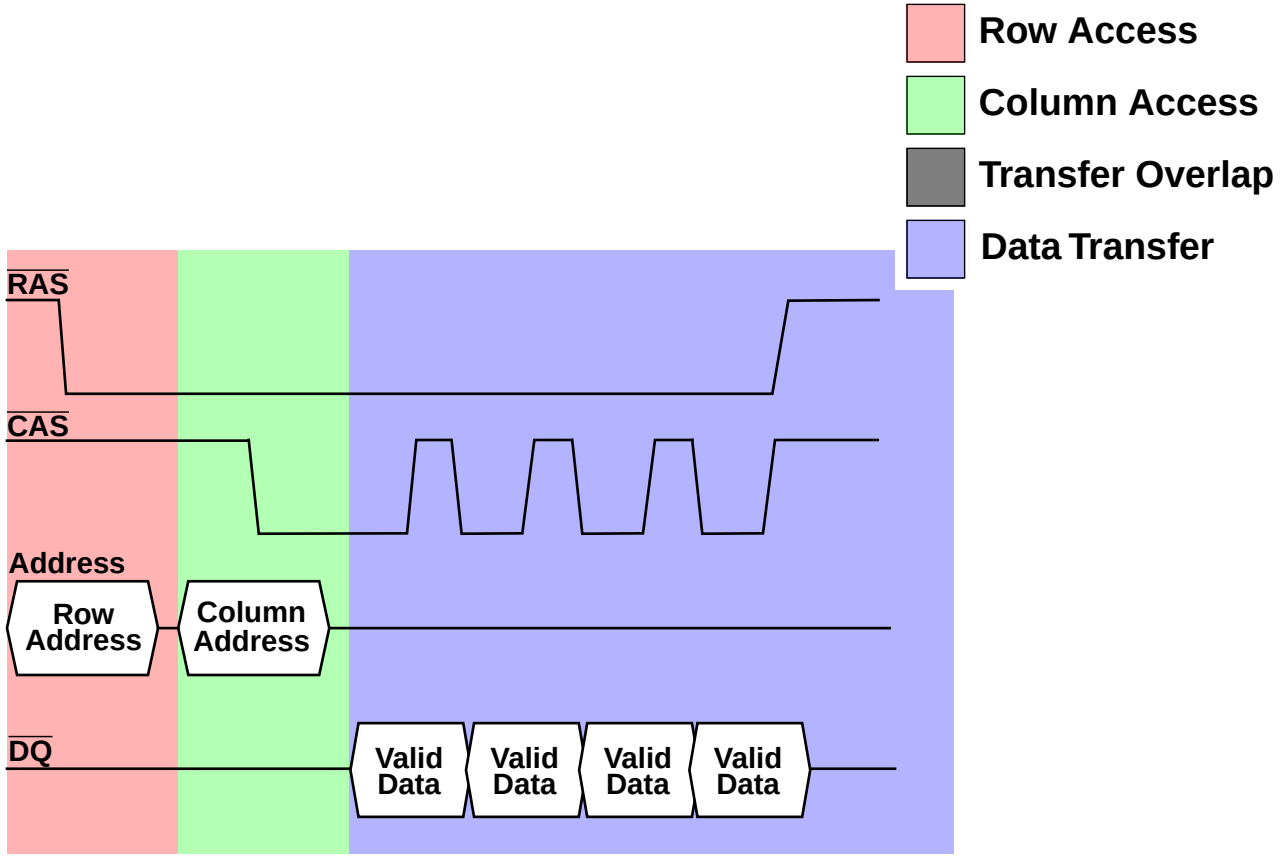

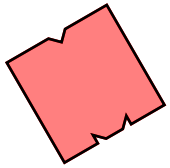

Bruce Jacob

University of
Maryland
ECE Dept.

DRAM Evolution

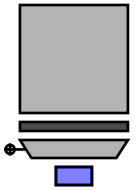

Read Timing for Pipeline Burst EDO

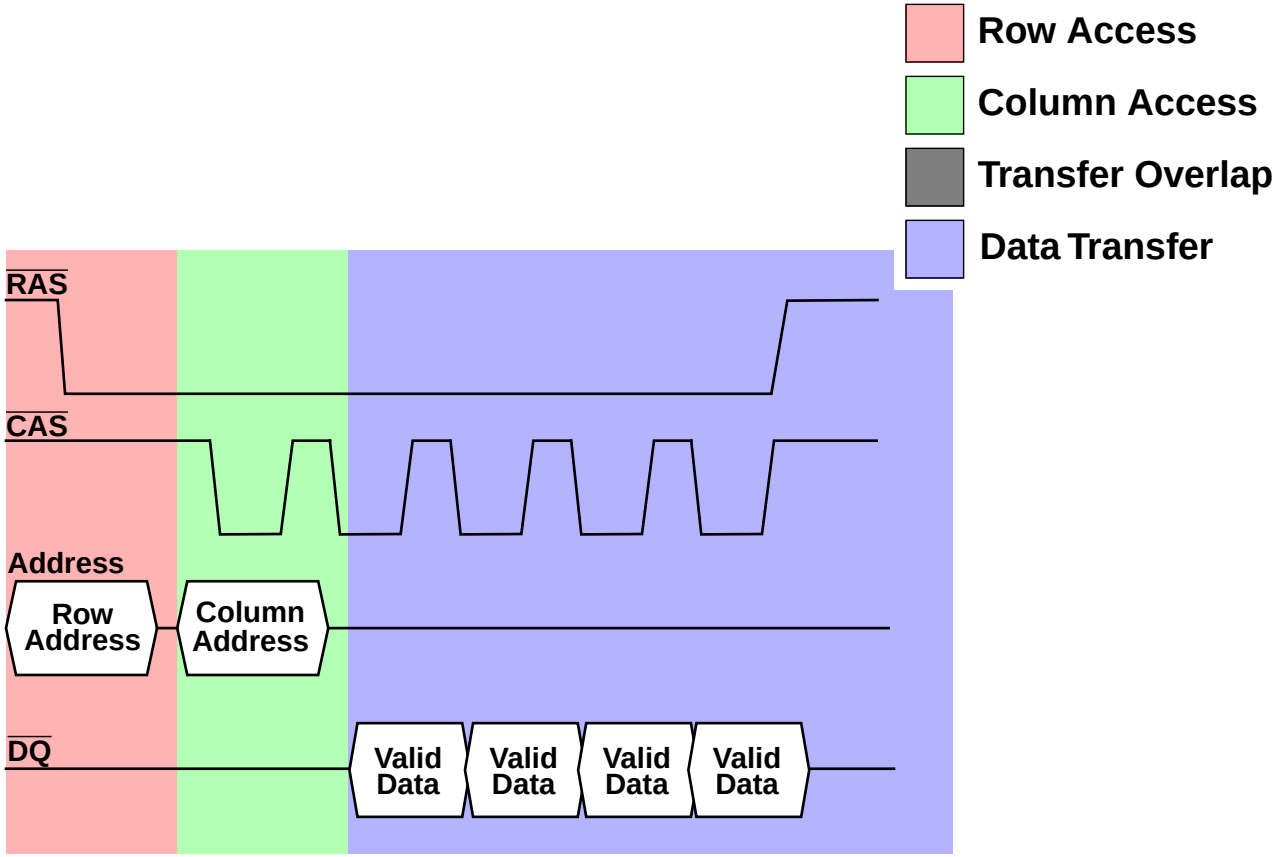

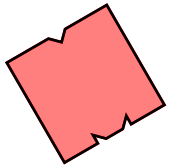

DRAM Evolution

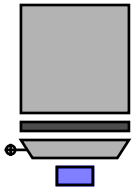

Read Timing for Synchronous DRAM

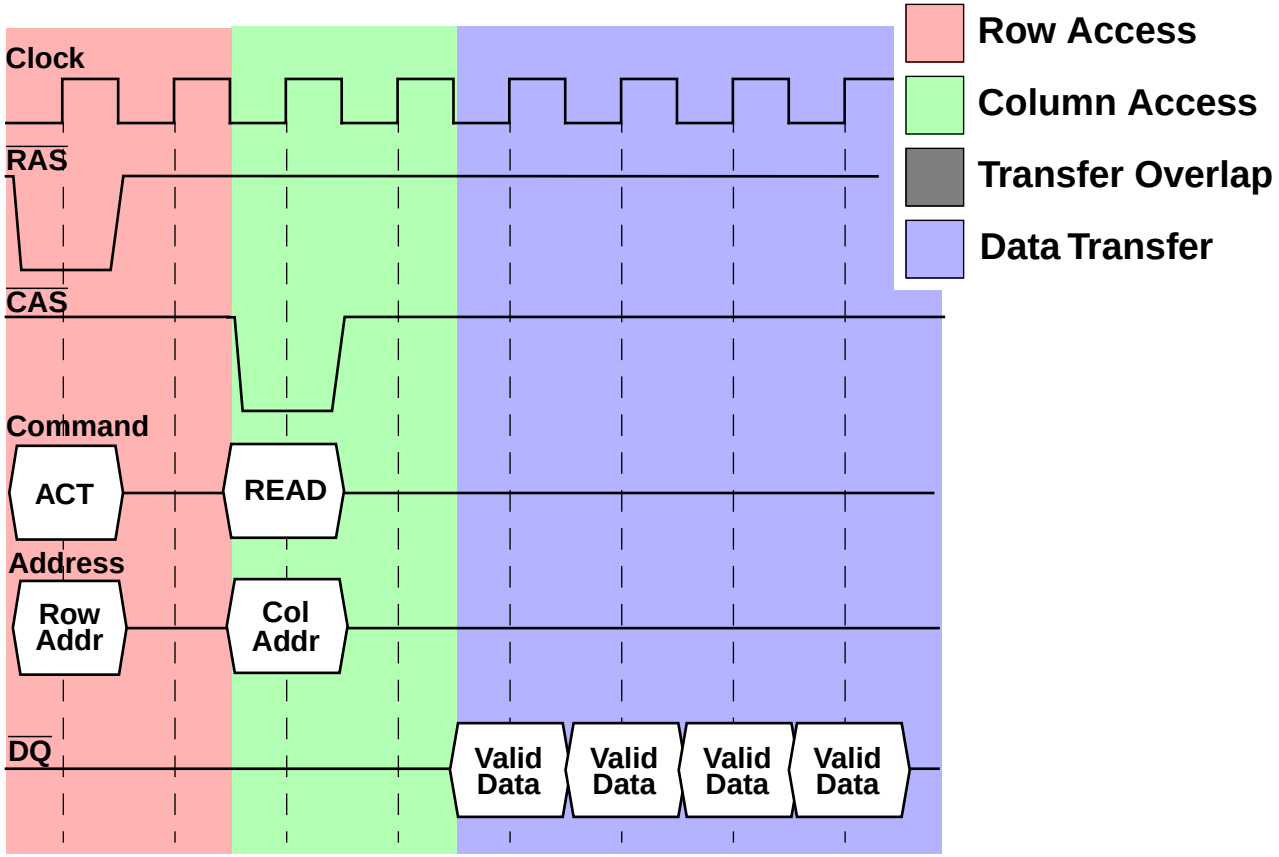

($\overline{RAS}$ + $\overline{CAS}$ + $\overline{OE}$ ... == Command Bus)

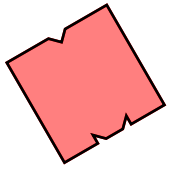

DRAM Evolution

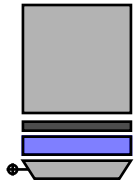

Inter-Row Read Timing for ESDRAM

ÒRegularÓ CAS-2 SDRAM, R/R to same bank

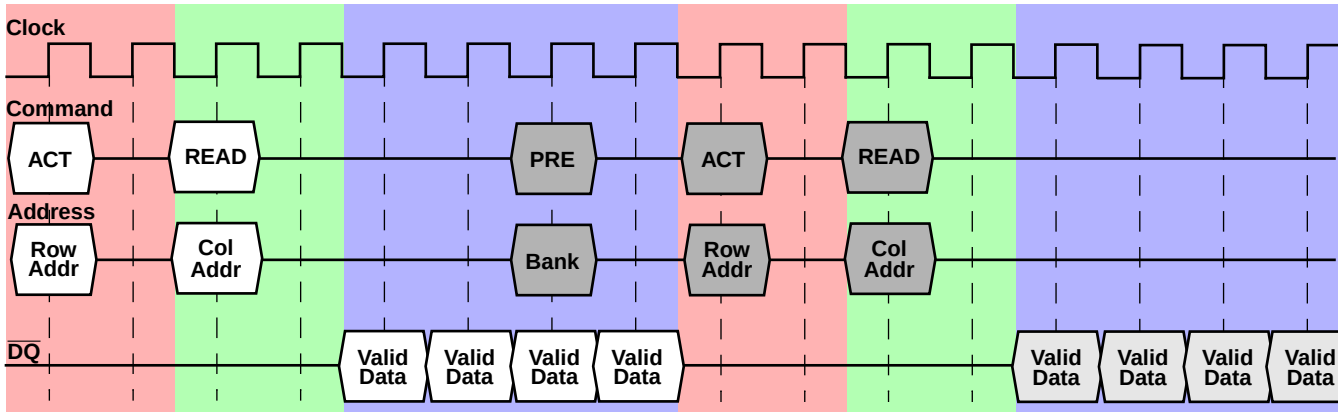

ESDRAM, R/R to same bank

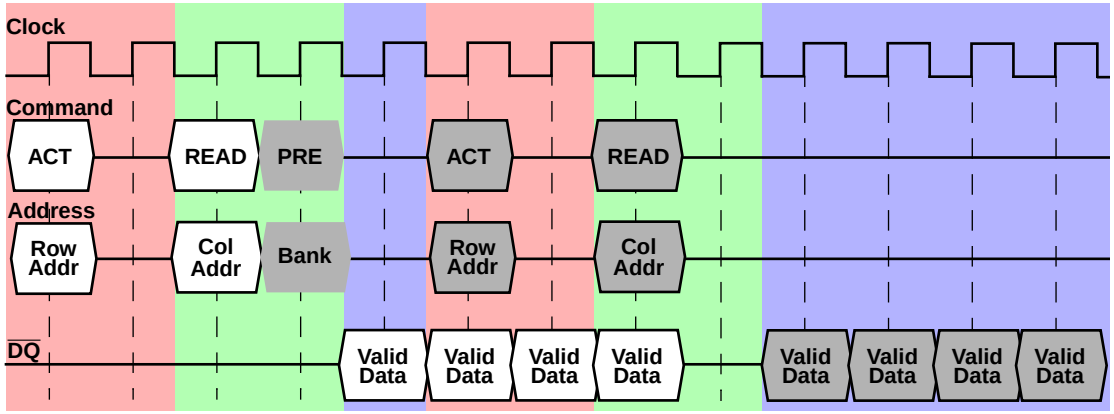

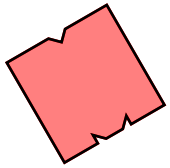

DRAM Evolution

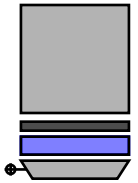

Write-Around in ESDRAM

ÒRegularÓ CAS-2 SDRAM, R/W/R to same bank, rows 0/1/0

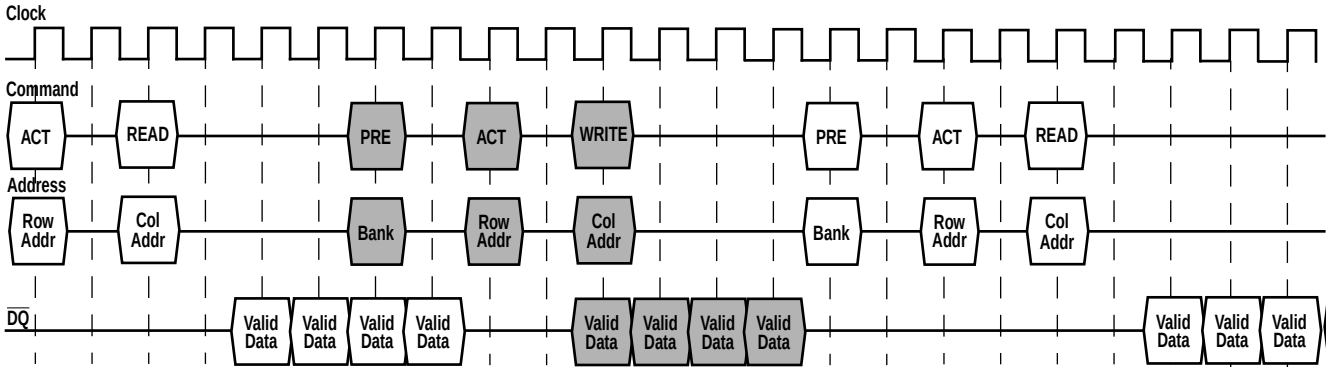

ESDRAM, R/W/R to same bank, rows 0/1/0

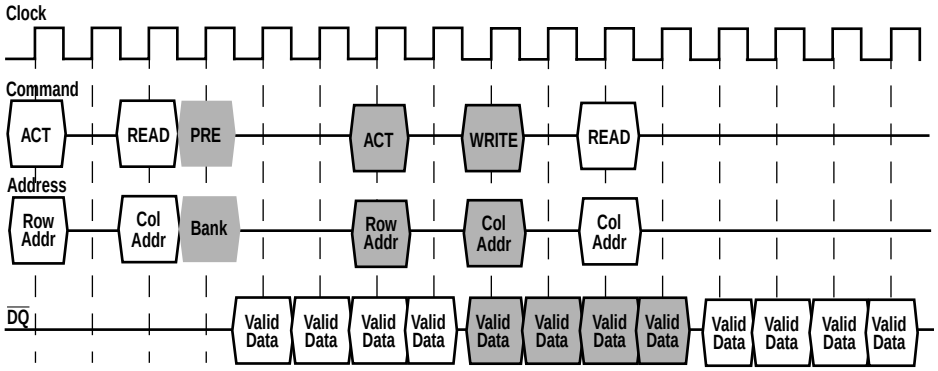

(can second READ be this aggressive?)

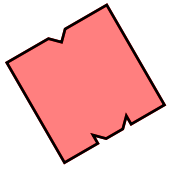

DRAM Evolution

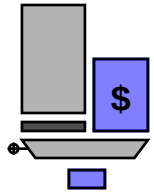

Internal Structure of Virtual Channel

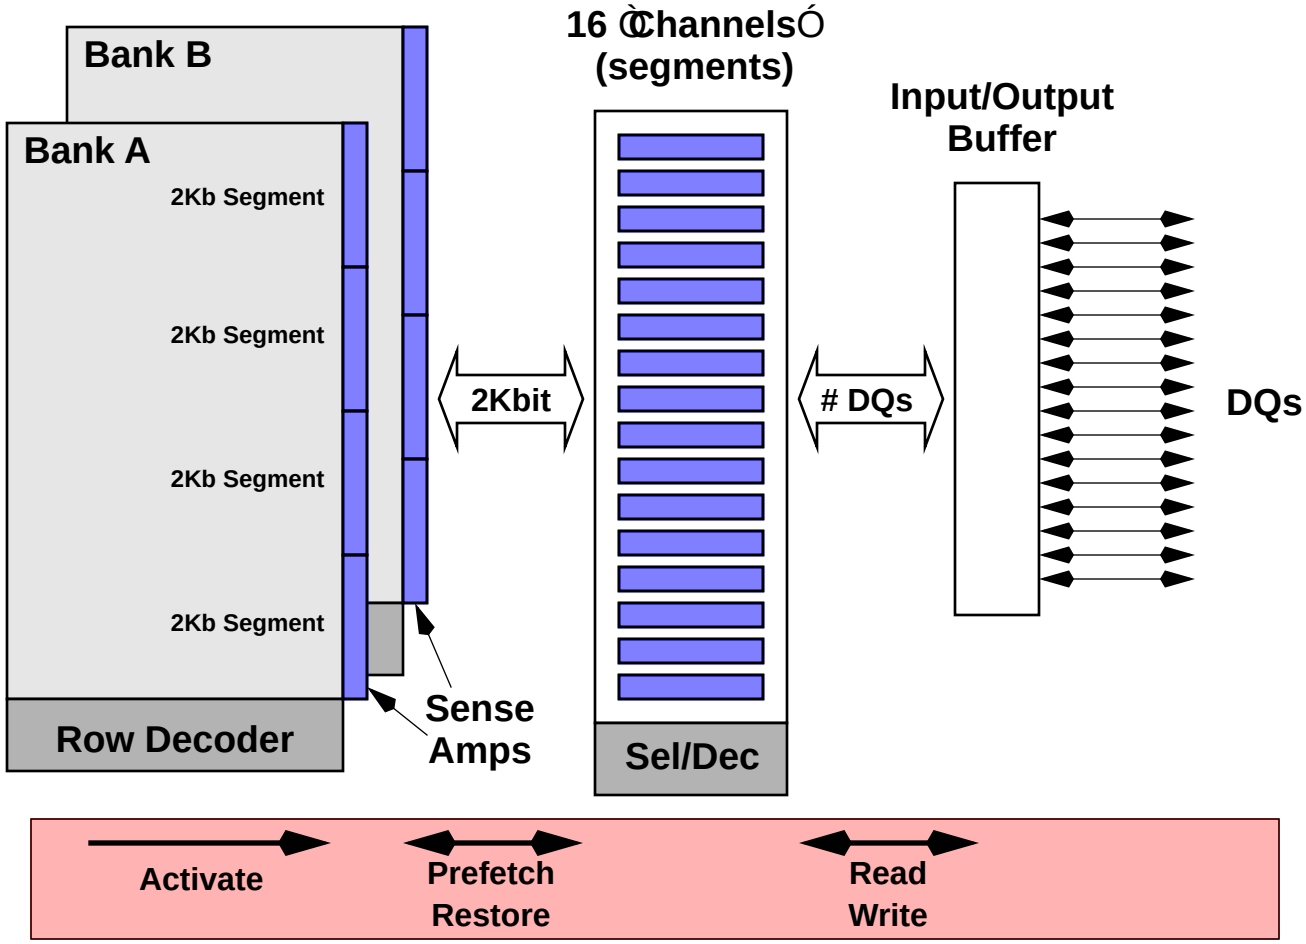

Segment cache is software-managed, reduces energy

DRAM Evolution

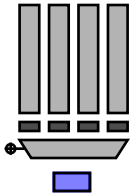

Internal Structure of Fast Cycle RAM

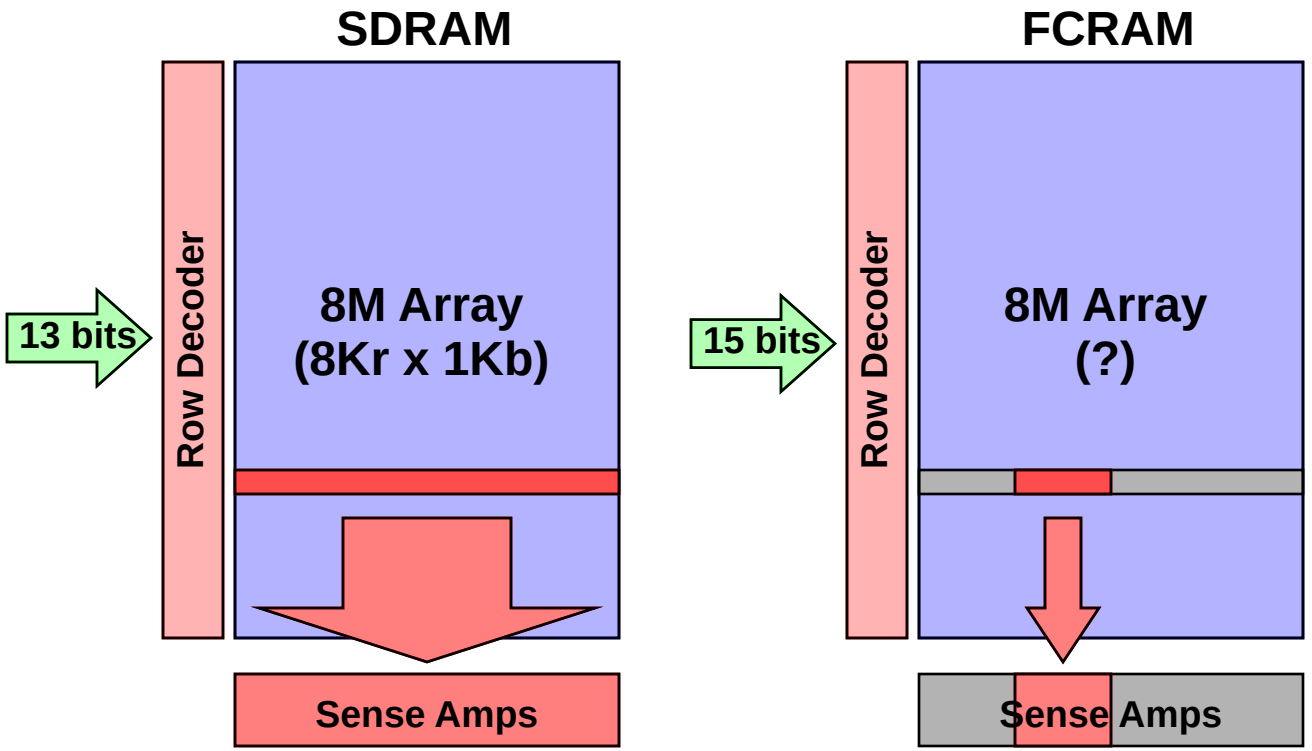

t_{RCD} = 15ns
(two clocks)

t_{RCD} = 5ns
(one clock)

Reduces access time and energy/access

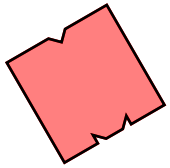

DRAM Evolution

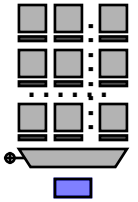

Internal Structure of MoSys 1T-SRAM