Bruce Jacob
David Wang

University
of Crete

CS-590.26, Spring 2014

High Speed Memory Systems: Architecture and Performance Analysis

Introduction

Credit where credit is due:
Slides contain original artwork (© Jacob, Wang 2005)

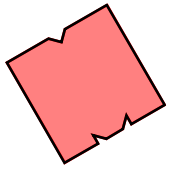

Stuff that will* be covered in this class:

- DRAM Device Architecture
- Memory System Organization
- System Controller
- DRAM Access Protocol
- Performance Analysis
- Reliability (Error detection/correction)
- Signalling (Data transport/reception)
- New System Architecture (HMC, etc.)
- Non Volatile RAM (Flash)
- A lot of other stuff...

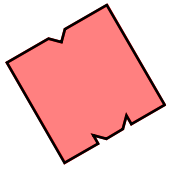

Stuff you need to know in order to understand the stuff that will be covered in this class:

- **Computer Architecture**
- **Computer Organization**
- **Basic Circuit-Design Concepts**
- **Caches (buffering)**
- **Pipelines (buffering & forwarding)**
- **Scheduling (queueing, out-of-order, etc.)**
- **Concurrency & Parallelism**
- **CPU Design — in particular, Control**
- **HW/SW Performance Analysis**
- **Hopefully not a lot of other stuff...**

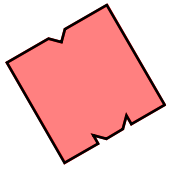

Stuff that will* be covered in this class:

- **DRAM Device Architecture**
- **Memory System Organization**
- **System Controller**
- **DRAM Access Protocol**
- **Performance Analysis**
- **Reliability (Error detection/correction)**
- **Signalling (Data transport/reception)**
- **New System Architecture (HMC, etc.)**
- **Non Volatile RAM (Flash)**
- **A lot of other stuff...**

*** to be read as "can"** (list covers full semester)

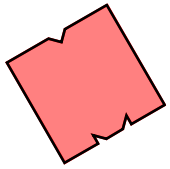

DRAM Device Architecture I

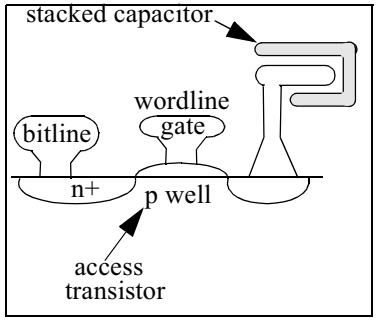

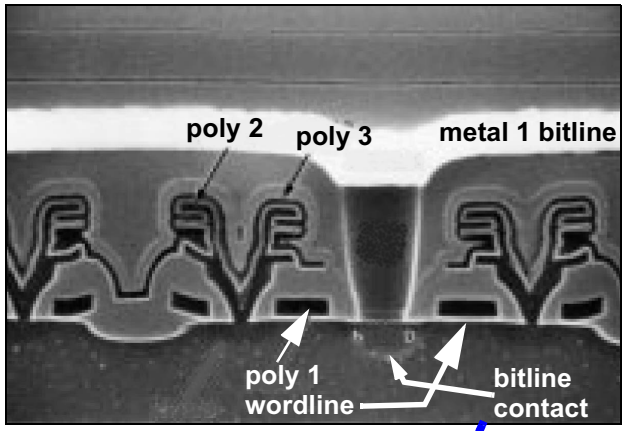

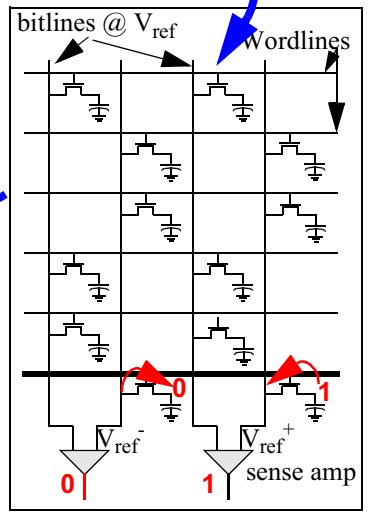

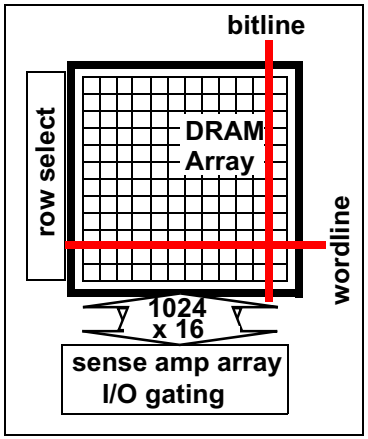

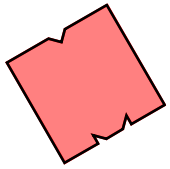

UNIVERSITY OF MARYLAND

DRAM Device Architecture II

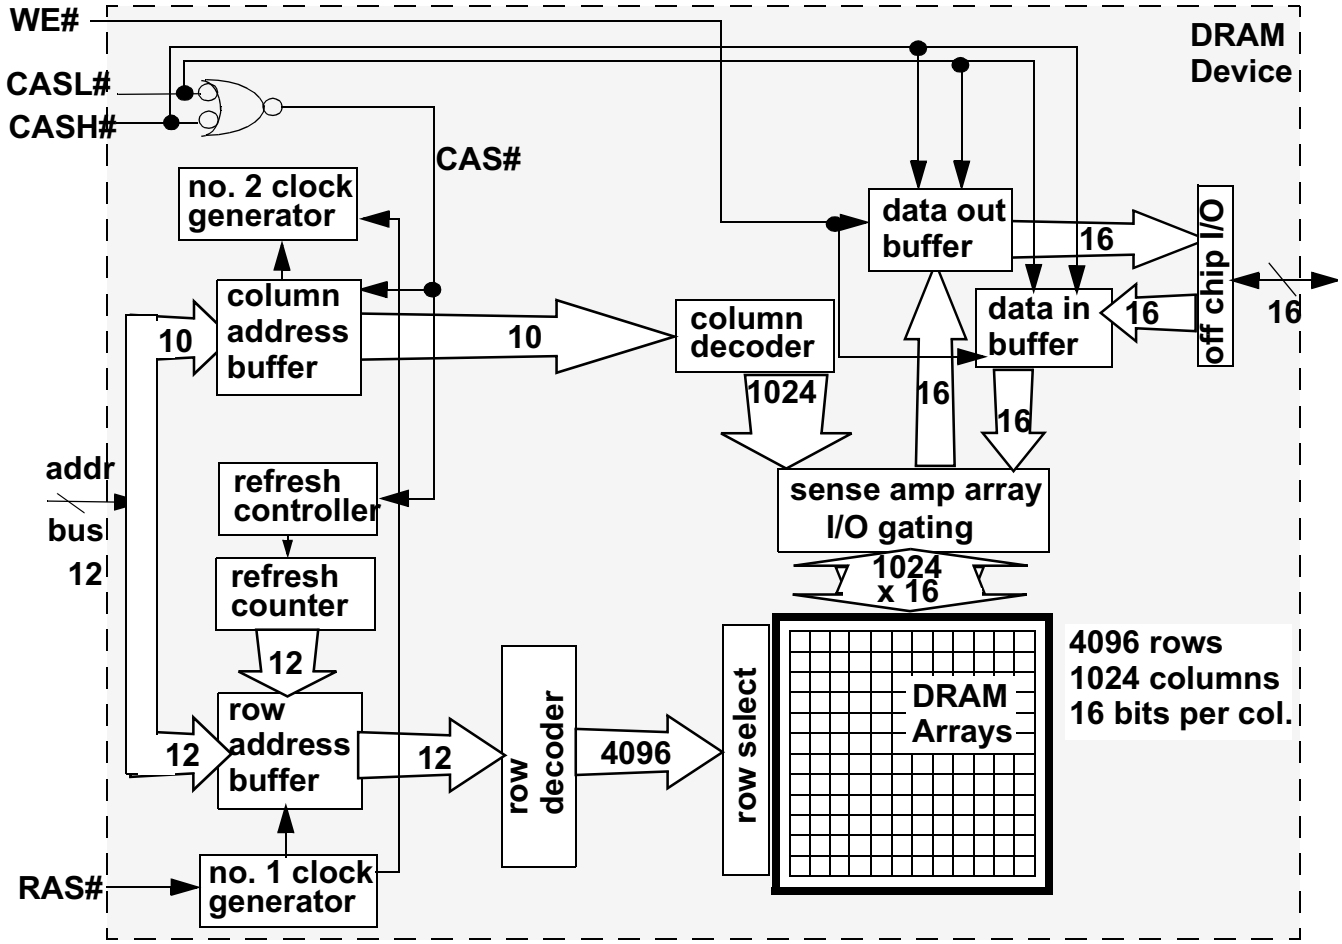

Today, DRAMs have multiple banks internally

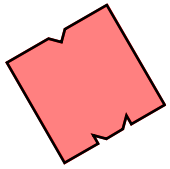

Memory System Organization

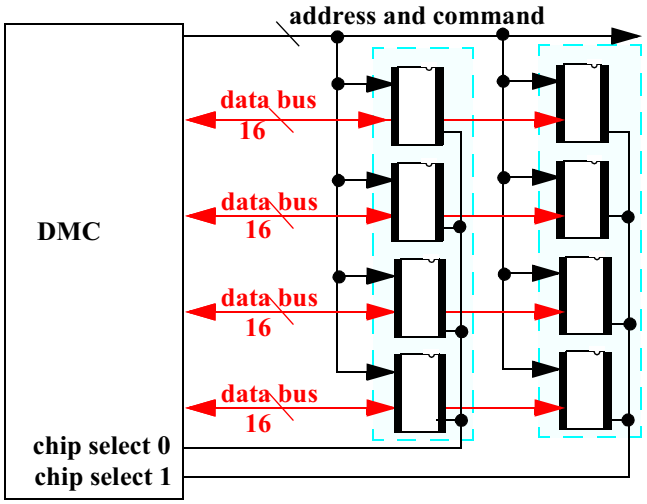

Channel ?
Rank ?
Bank ?
Row ?
Column ?

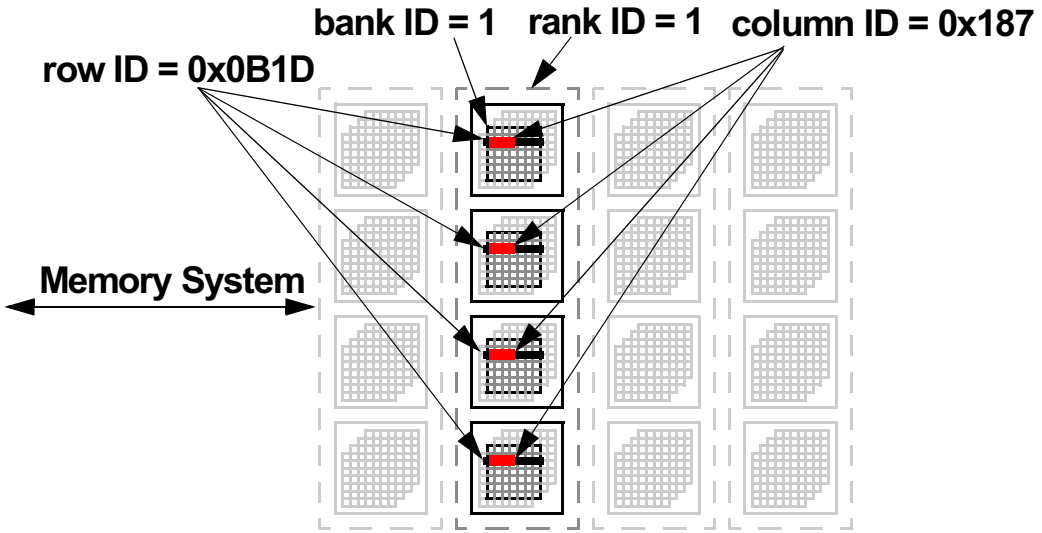

System Controller

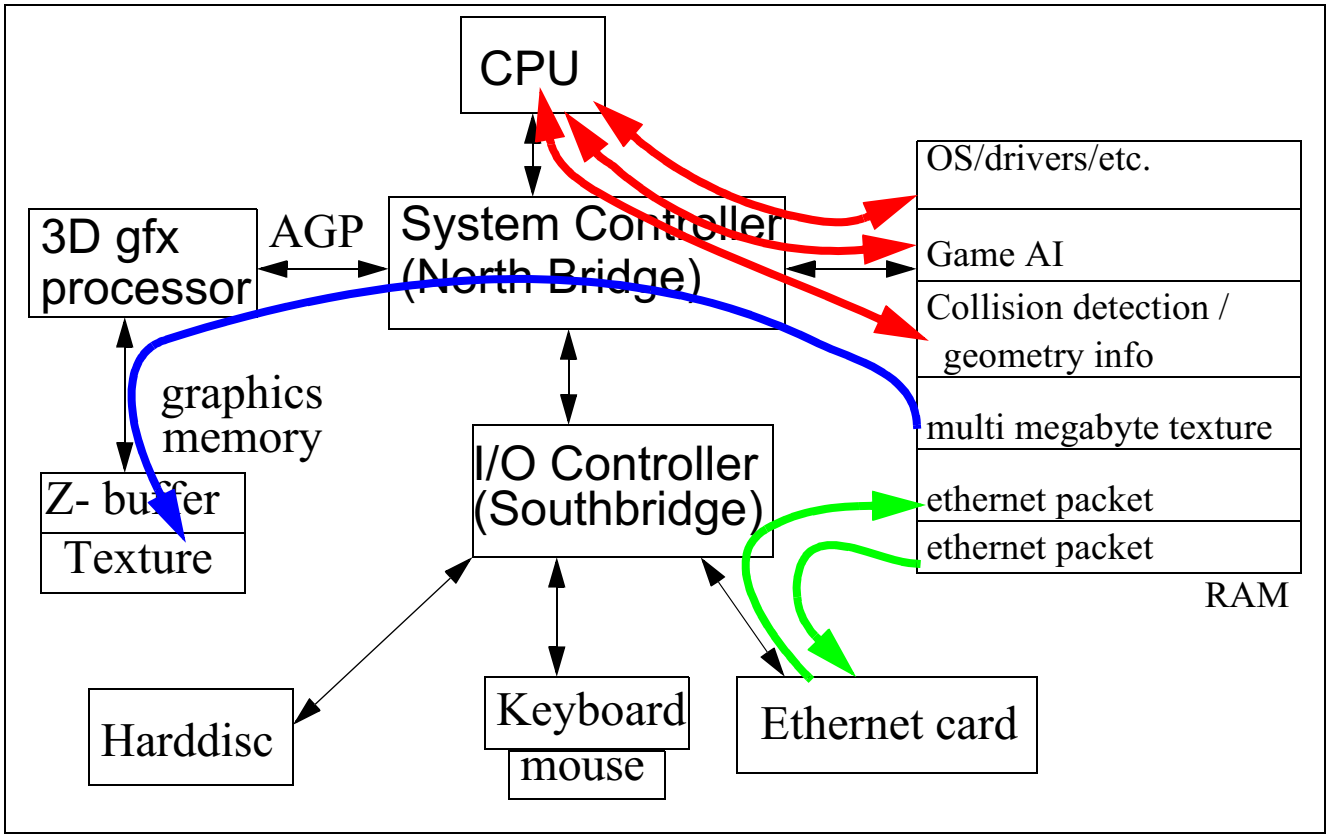

Heavy demand placed on memory system

Heavier still in SMP/SMT/CMP system

System Controller == System traffic cop

DRAM Access Protocol

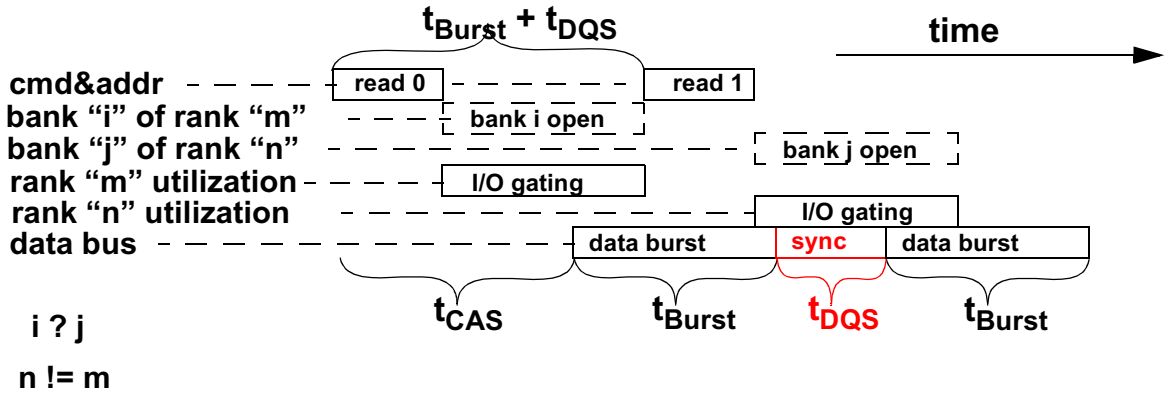

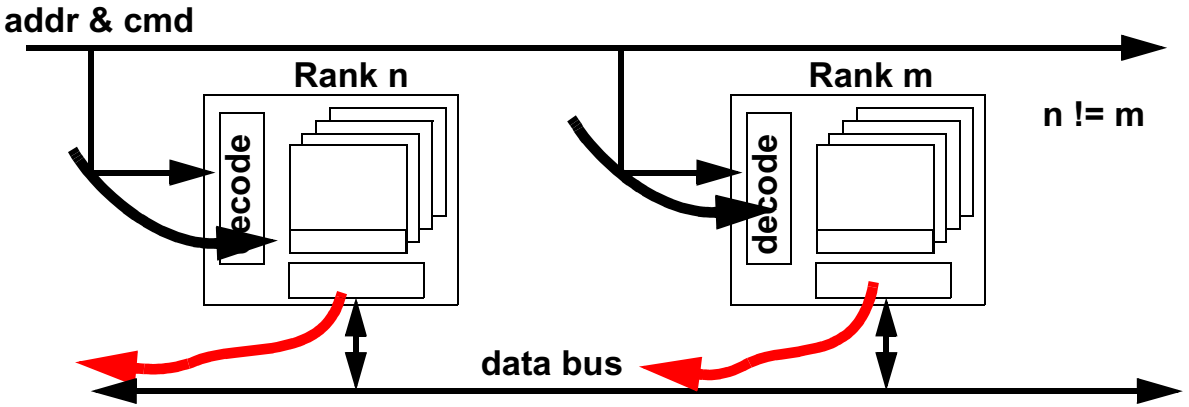

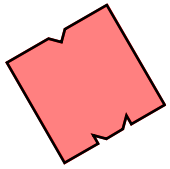

Memory Request Overview

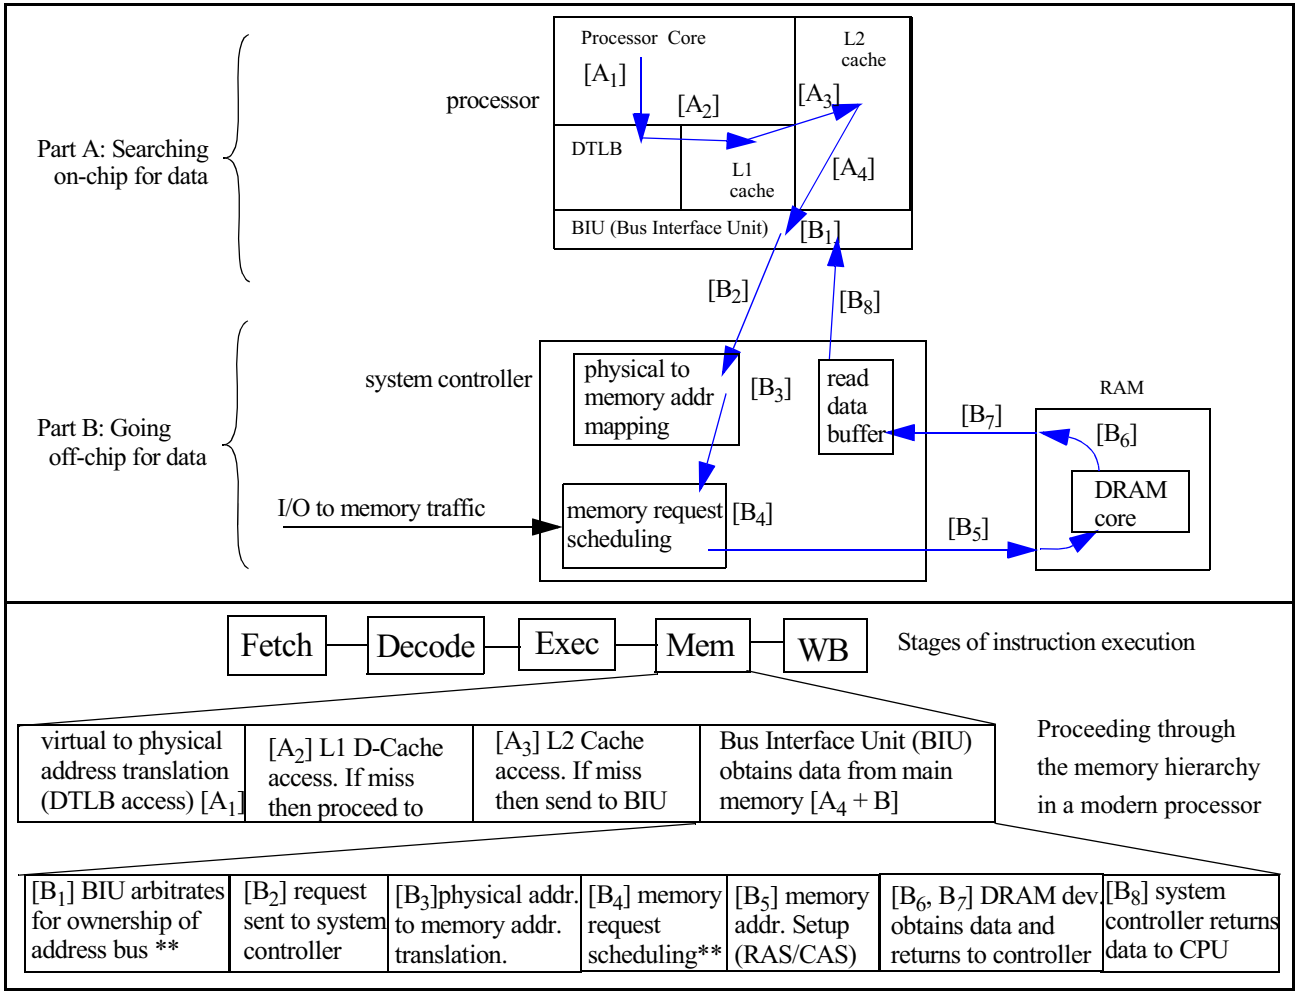

** Steps not required for some processor/system controllers. protocol dependant.

Progression of a Memory Read Transaction Request Through Memory System

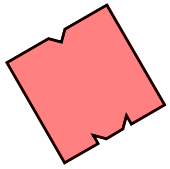

"Memory Latency"

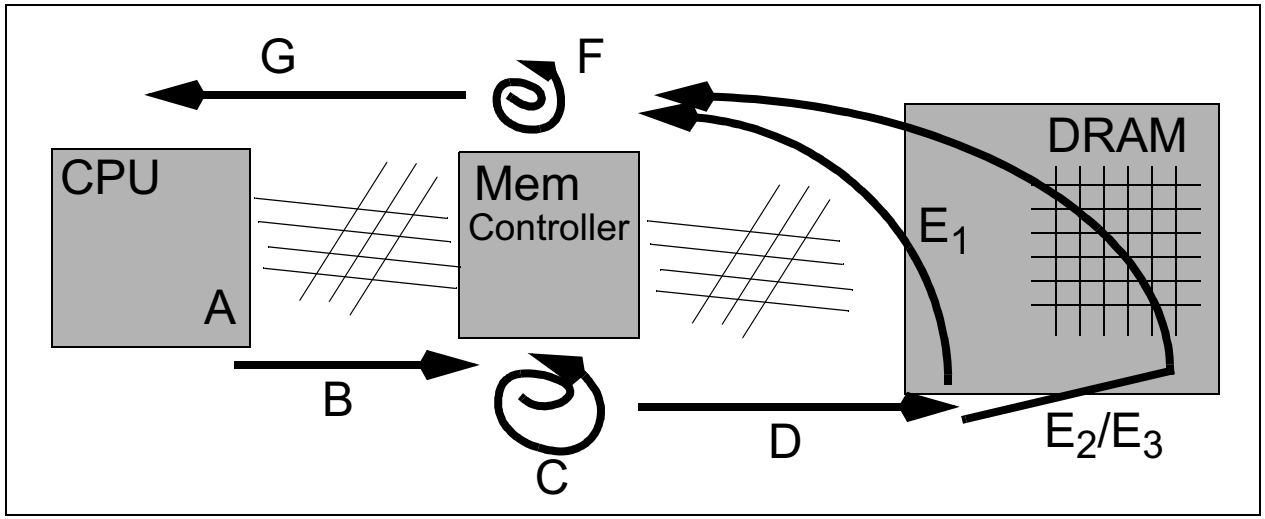

A: Transaction request may be delayed in Queue
B: Transaction request sent to Memory Controller
C: Transaction converted to Command Sequences (may be queued)
D: Command/s Sent to DRAM
E_1: Requires only a **CAS** or
E_2: Requires **RAS + CAS** or
E_3: Requires **PRE + RAS + CAS**
F: Data is staged at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

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

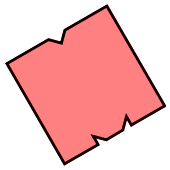

Performance Analysis

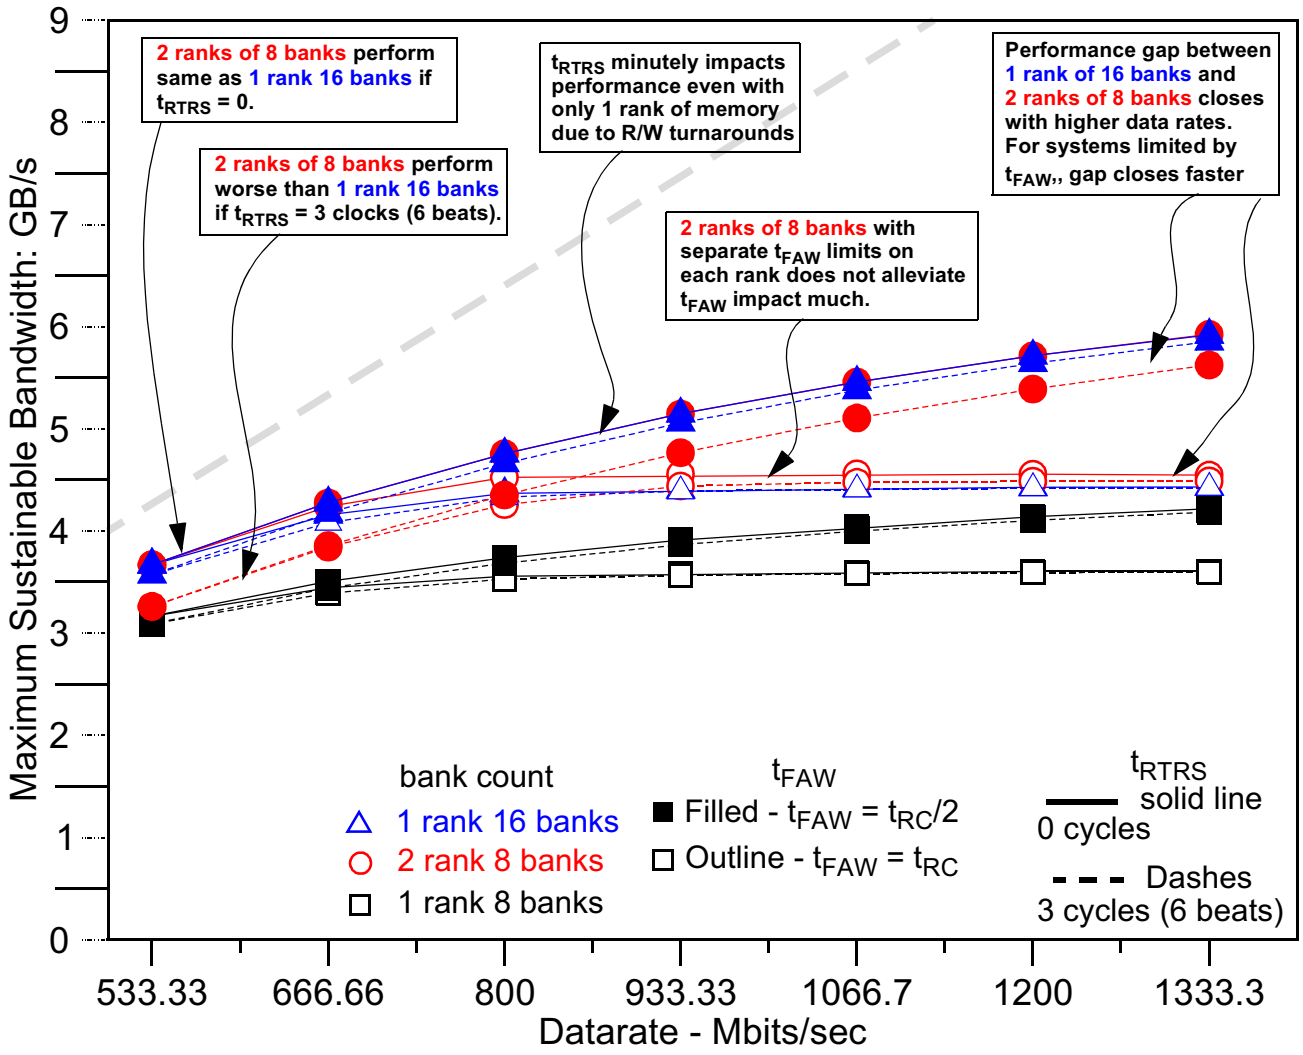

t_{RC} = 60ns, burst of eight, 8B wide channel

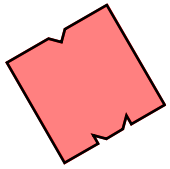

Reliability

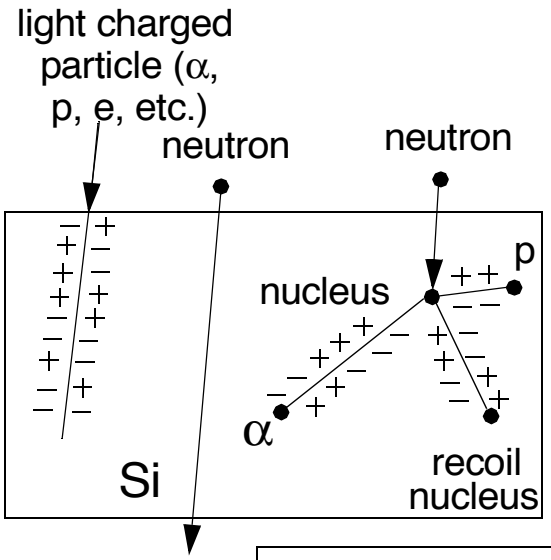

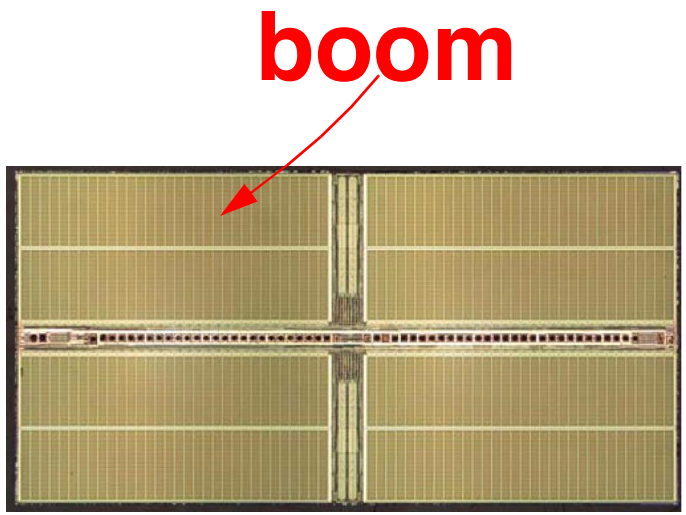

R = { 0 1 1 0 1 0 0 1 1 1 1 0 }

R = { 0 1 1 0 1 0 0 1 1 1 0 0 }

One bit error. Can we detect and correct?

Recompute check bits

$R_{0001} = R_{0011} \oplus R_{0101} \oplus R_{0111} \oplus R_{1001} \oplus R_{1011} = 1 \oplus 1 \oplus 0 \oplus 1 \oplus 0 = 1$

$R_{0010} = R_{0011} \oplus R_{0110} \oplus R_{0111} \oplus R_{1010} \oplus R_{1011} = 1 \oplus 0 \oplus 0 \oplus 1 \oplus 0 = 0$

$R_{0100} = R_{0101} \oplus R_{0110} \oplus R_{0111} \oplus R_{1100} = 1 \oplus 0 \oplus 1 \oplus 0 = 0$

$R_{1000} = R_{1001} \oplus R_{1010} \oplus R_{1011} \oplus R_{1100} = 1 \oplus 1 \oplus 0 \oplus 0 = 0$

XOR old check bits against new check bits

	R_{1000}	R_{0100}	R_{0010}	R_{0001}	
	1	0	1	0	Old
⊕	0	0	0	1	New
	1	0	1	1	ECC Syndrome

Syndrome != 0000

Bit position 11 is rotten

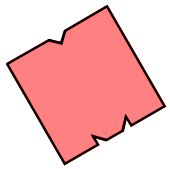

Signalling

The Digital Fantasy

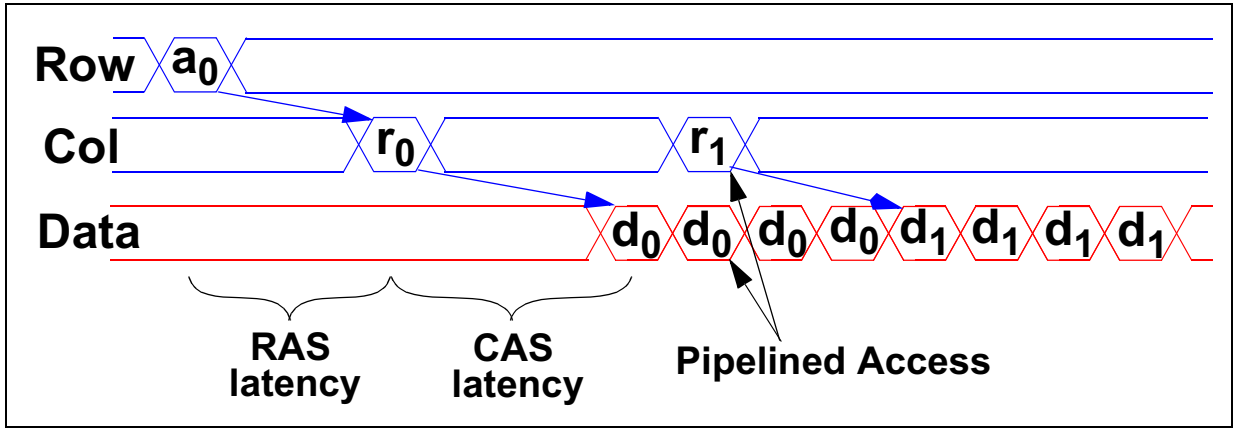

Pretend that the world looks like this

But...

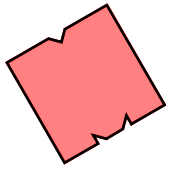

DRAM Interface: Signals

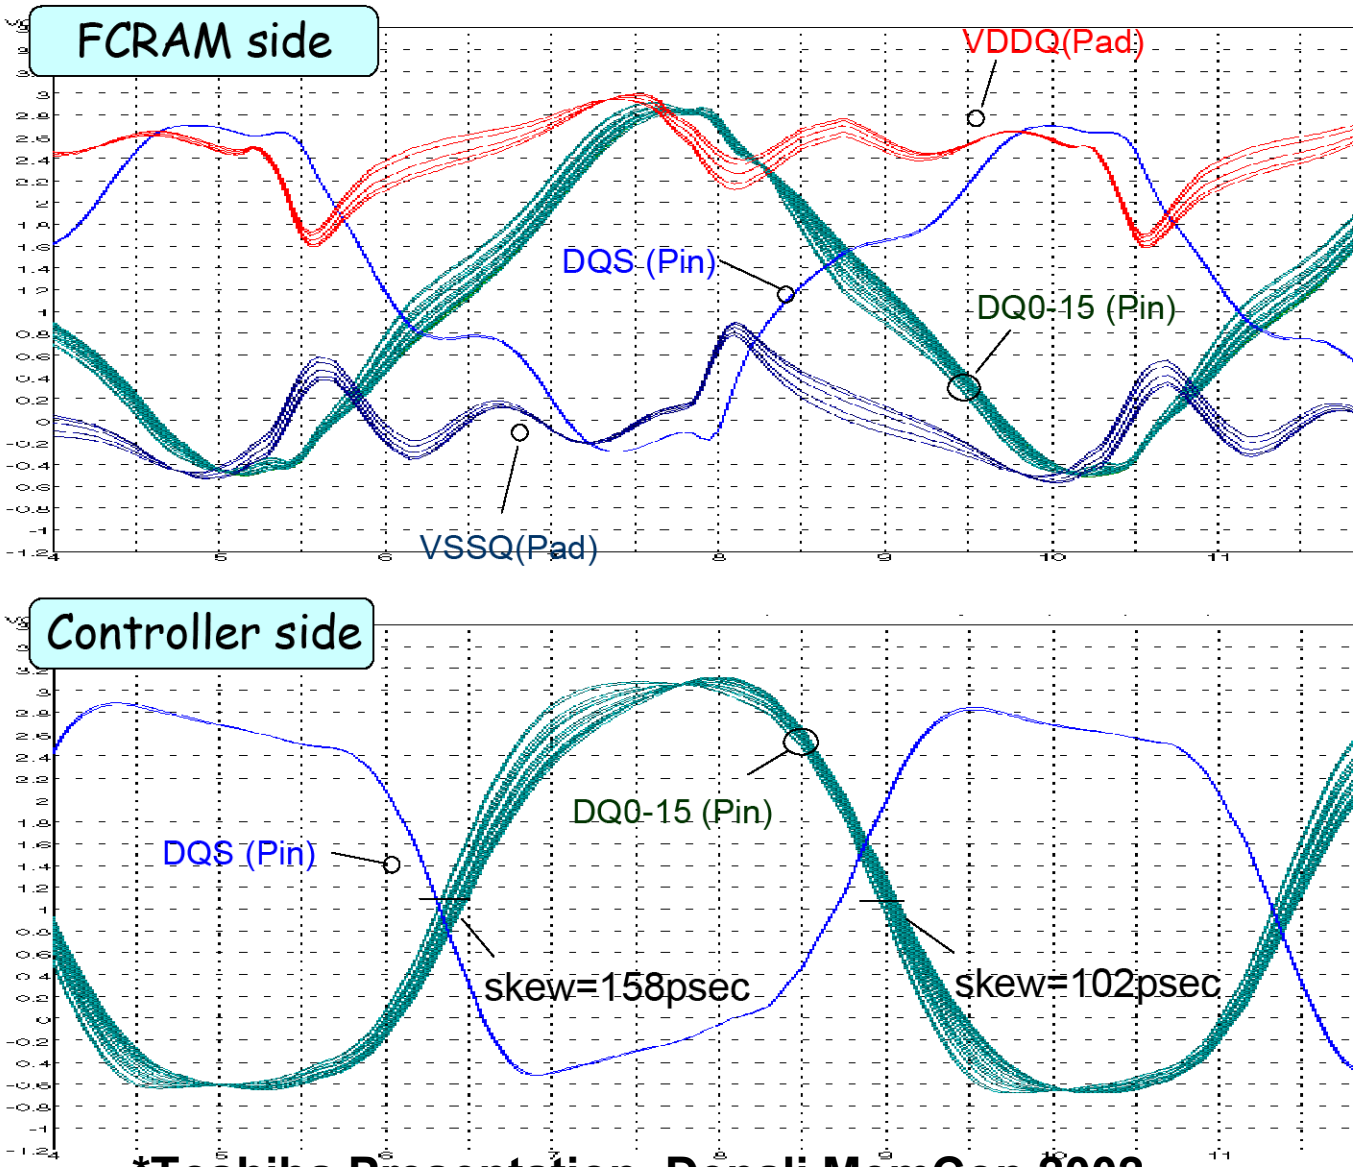

*Toshiba Presentation, Denali MemCon 2002

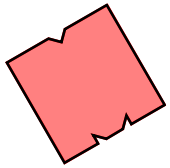

Interface:Signal Propagation

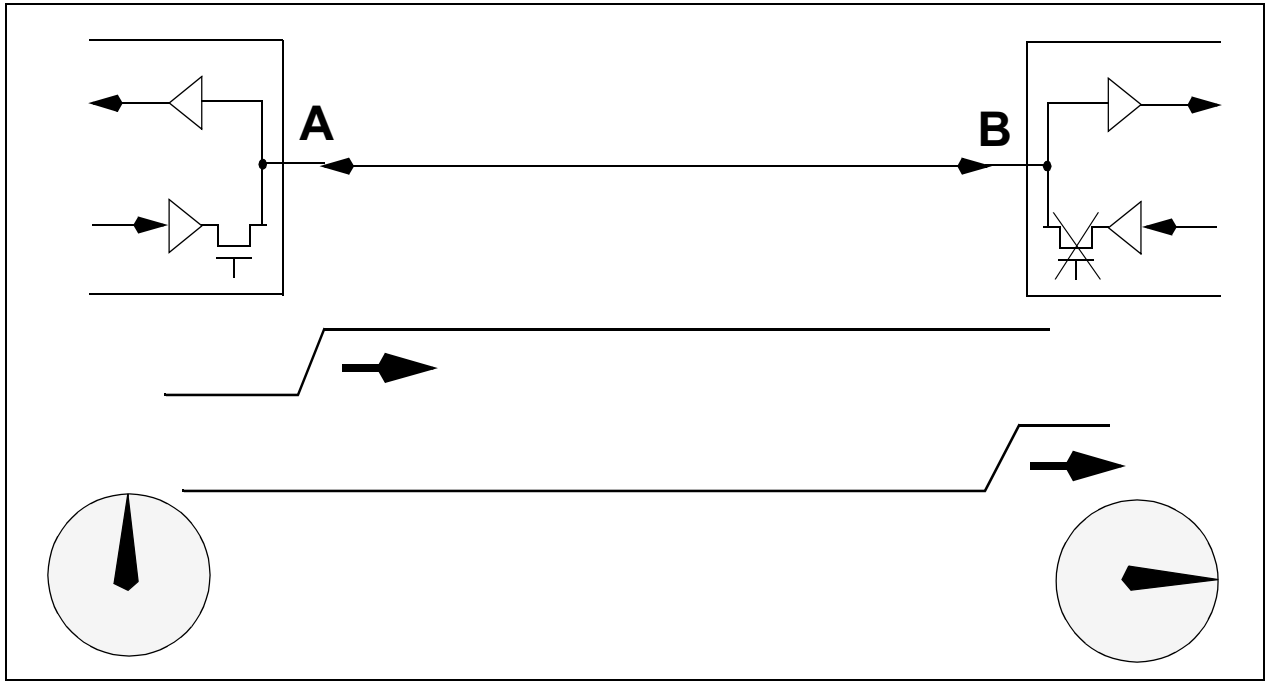

Ideal Transmission Line

~ 0.66c = 20 cm/ns

PC Board + Module Connectors +
Varying Electrical Loads

= Rather non-Ideal Transmission Line

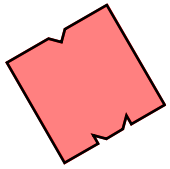

Interface: Clocking Issues

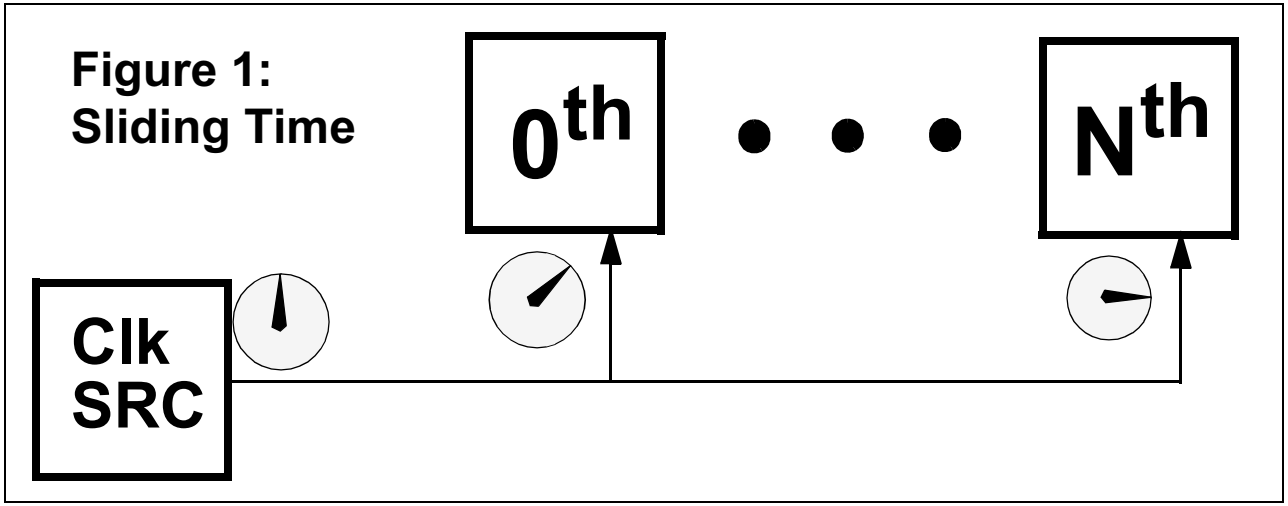

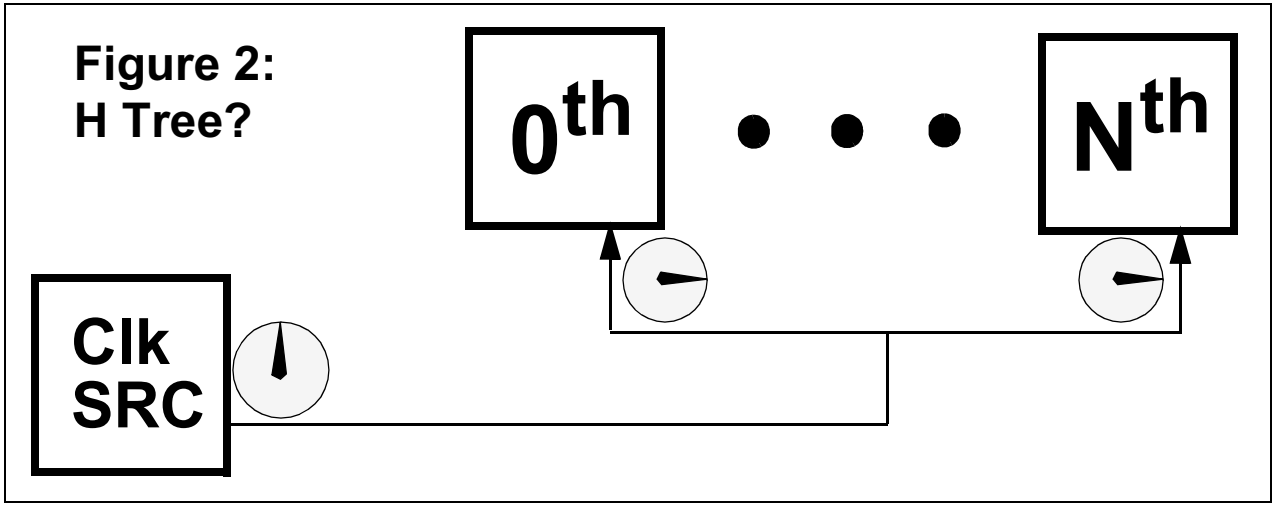

What Kind of Clocking System?

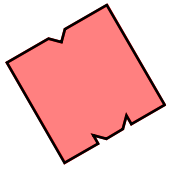

Path Length Differential

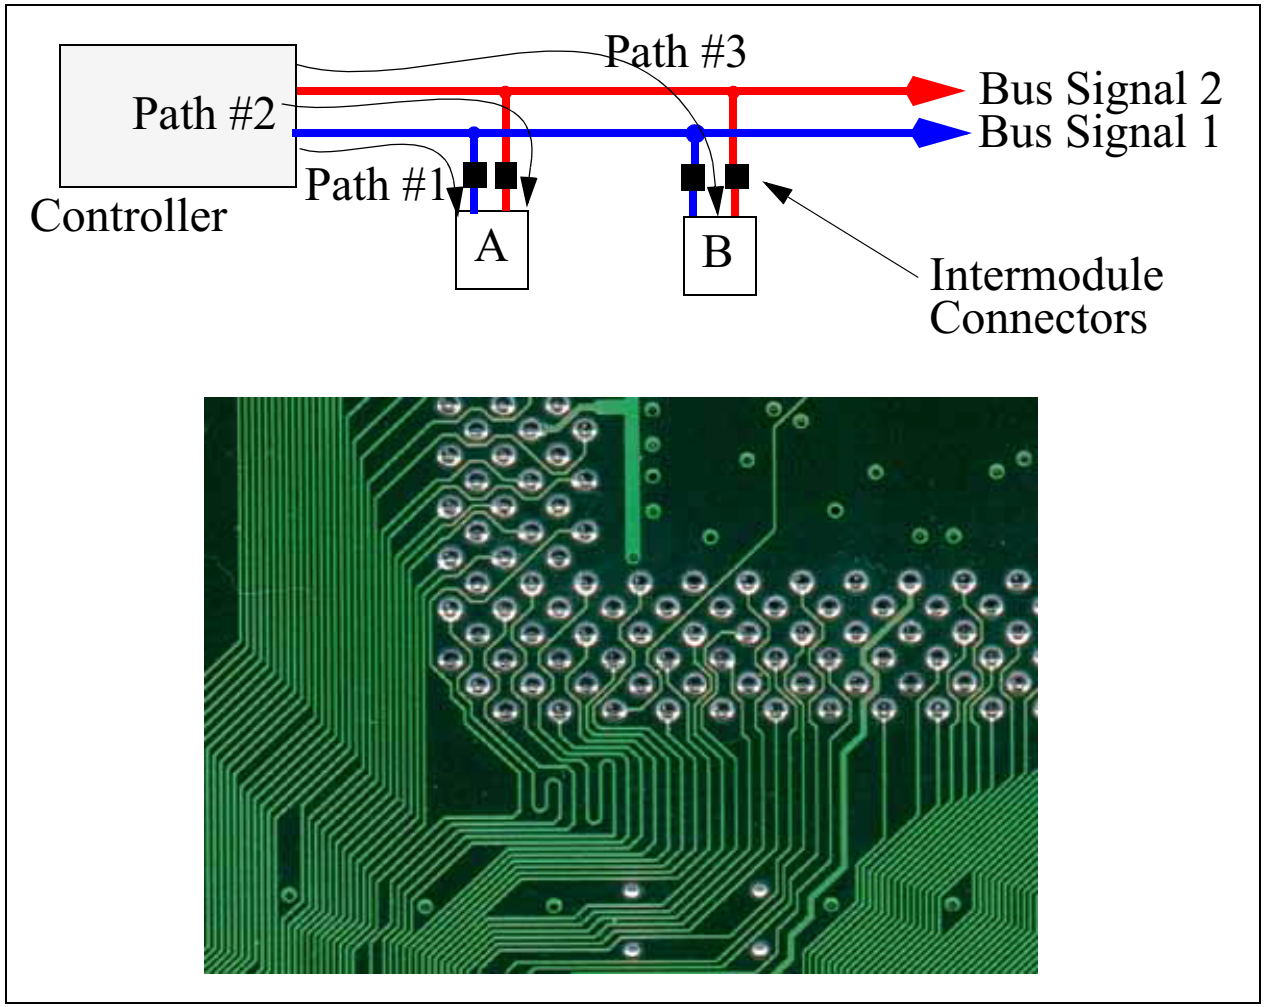

High Frequency AND Wide Parallel Busses are Difficult to Implement

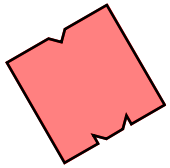

Future Trends I

	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**

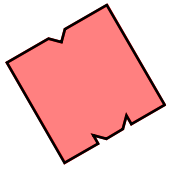

Future Trends II

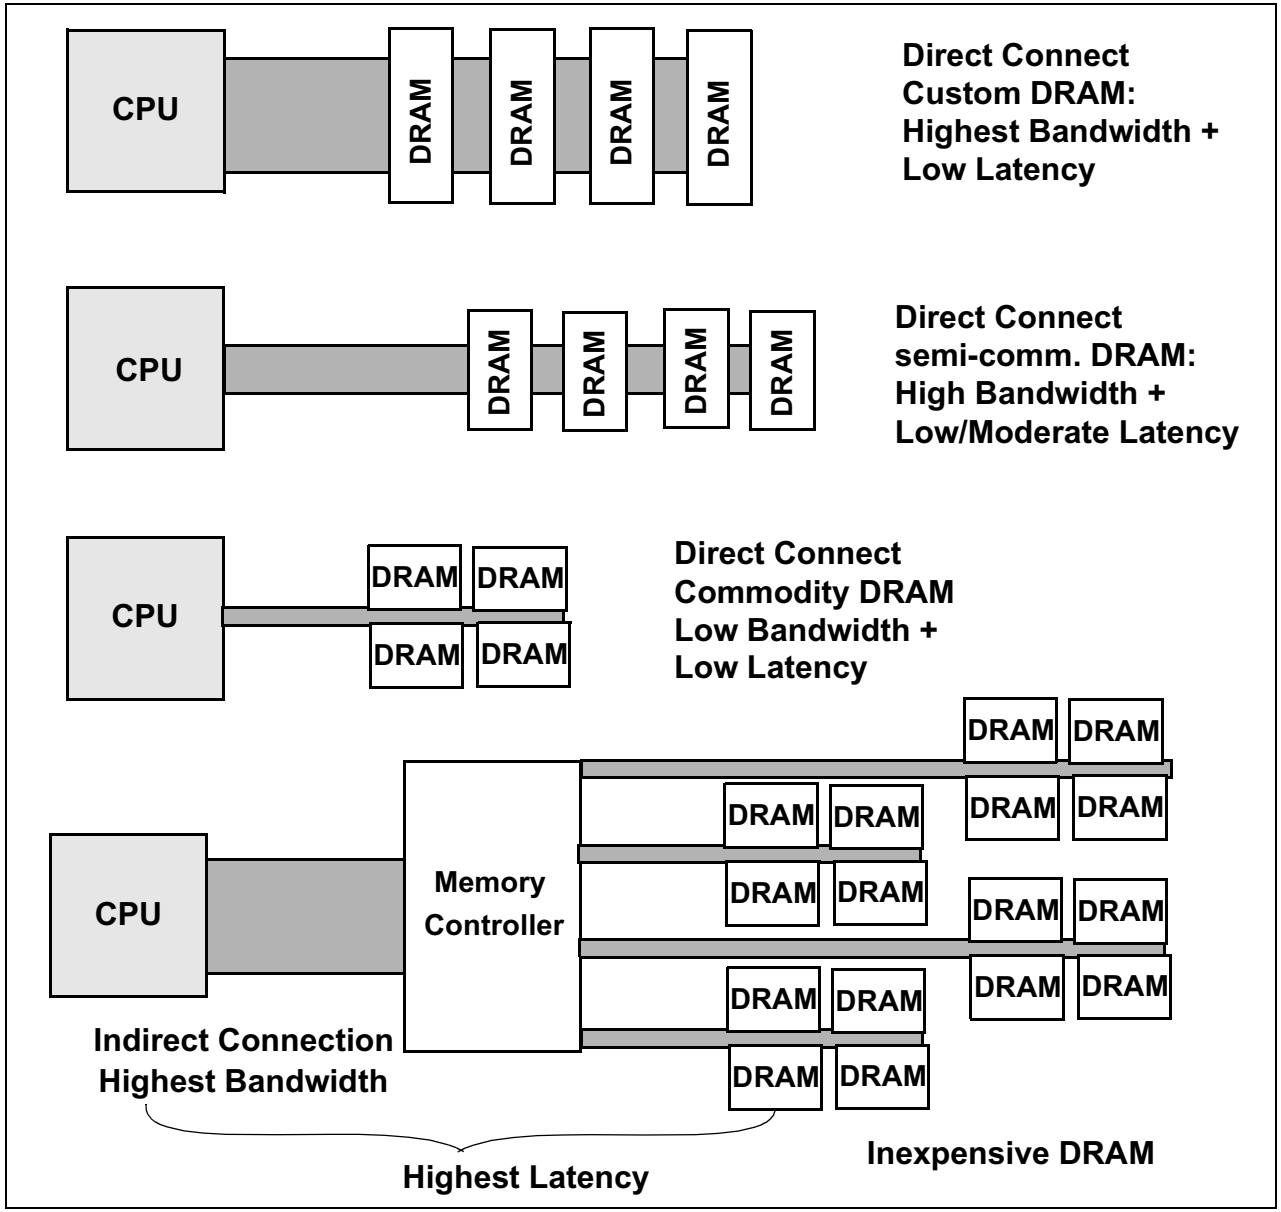

Research Areas: Topology

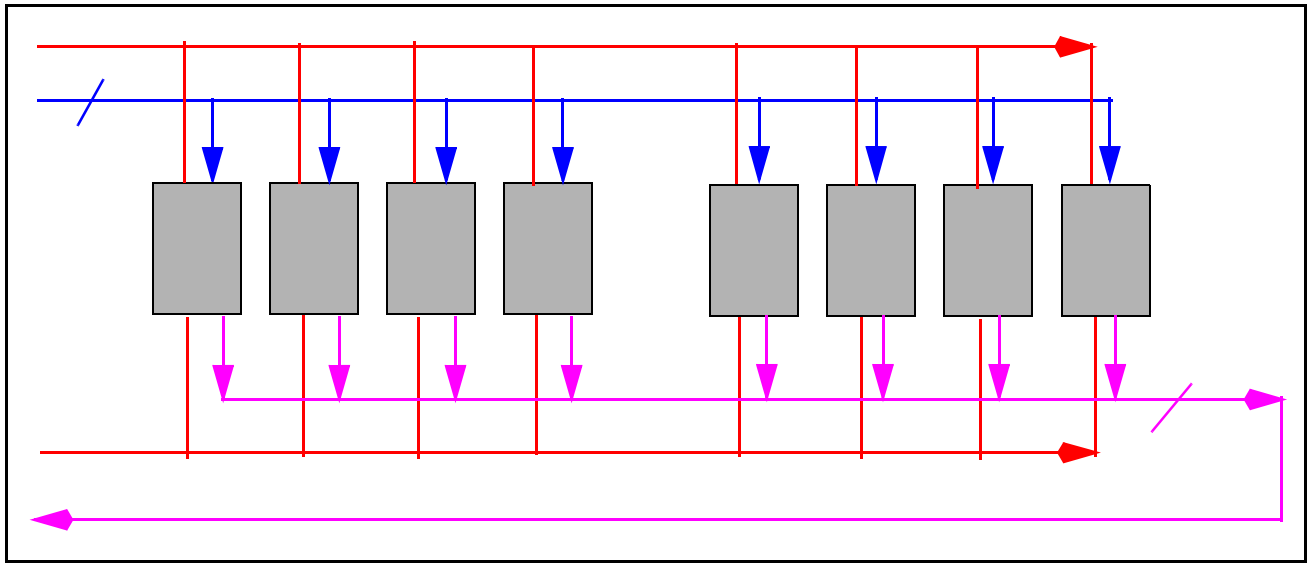

Unidirectional Topology:

- **Write Packets sent on Command Bus**
- **Pins used for Command/Address/Data**
- **Further Increase of Logic on DRAM chips**

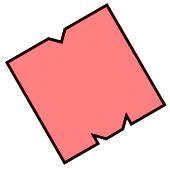

Memory Commands?

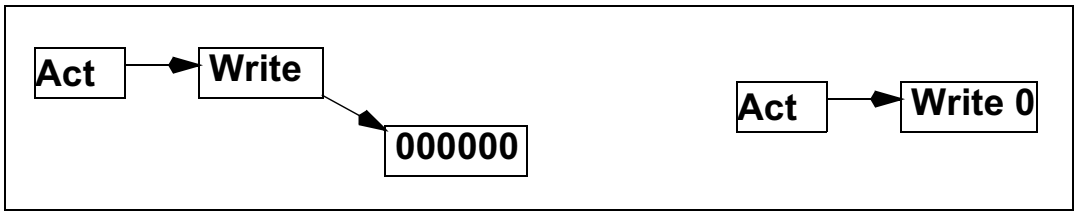

Instead of A[ ] = 0; Do "write 0"

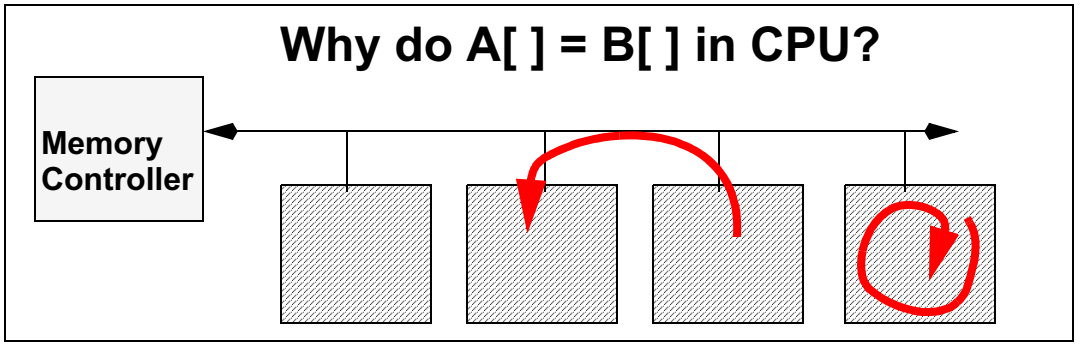

Move Data inside of DRAM or between DRAMs.

Why do STREAMadd in CPU?

A[ ] = B[ ] + C[ ]

Active Pages *(Chong et. al. ISCA '98)

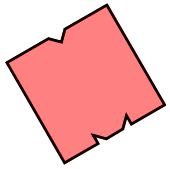

- **Grading**
- **Projects**
- **Textbook**