CS 250
VLSI System Design

Lecture 11 – DRAM

2009-10-1

John Wawrzynek and Krste Asanovic´

with John Lazzaro

TA: Yunsup Lee

www-inst.eecs.berkeley.edu/~cs250/

Today's Lecture: DRAM

- **Bottom-up: DRAM core cells**
- **Top-down: SDRAM commands**
- **DRAM controller design ideas**

Dynamic Memory Cells

Recall: Capacitors in action

Because the dielectric is an insulator, and does not conduct.

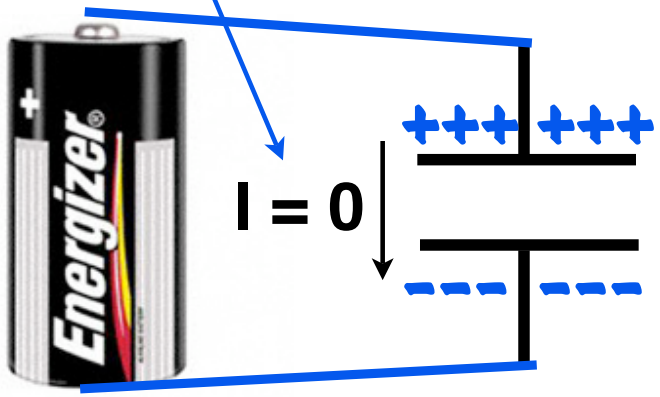

After circuit "settles" ...

Q = C V = C * 1.5 Volts (D cell)

Q: Charge stored on capacitor

C: The capacitance of the device: function of device shape and type of dielectric.

After battery is removed:

Still, Q = C * 1.5 Volts

Capacitor "remembers" charge

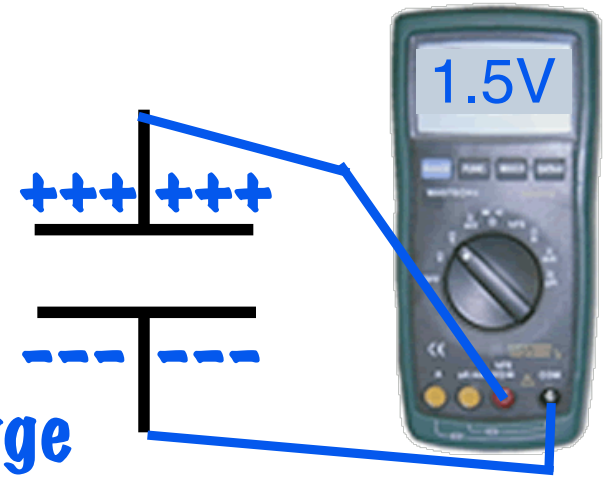

DRAM cell: 1 transistor, 1 capacitor

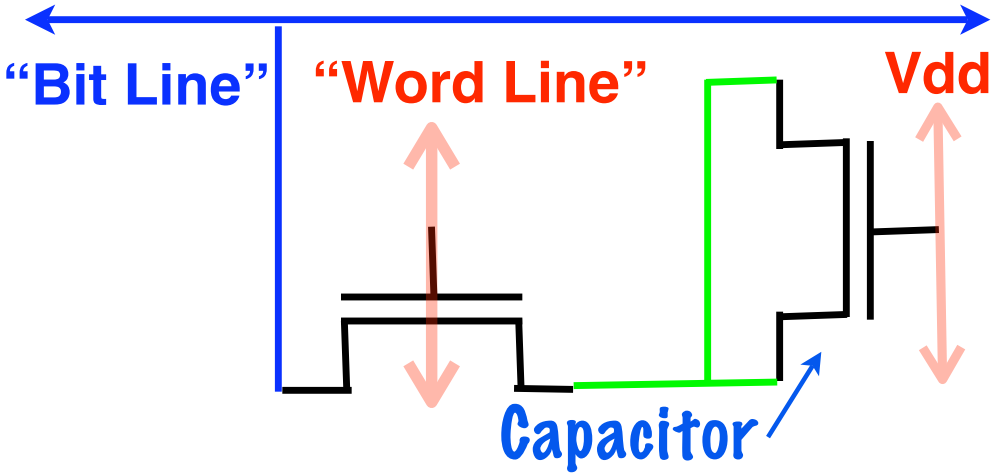

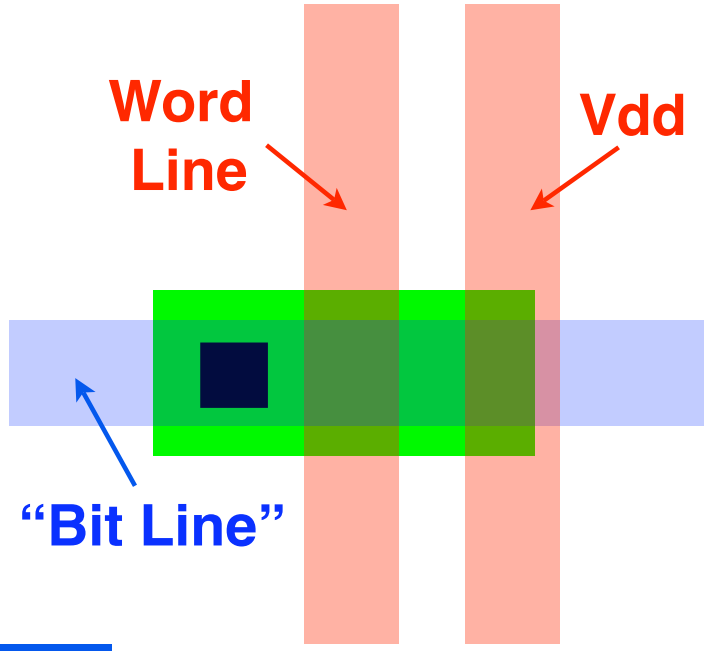

"Bit Line"

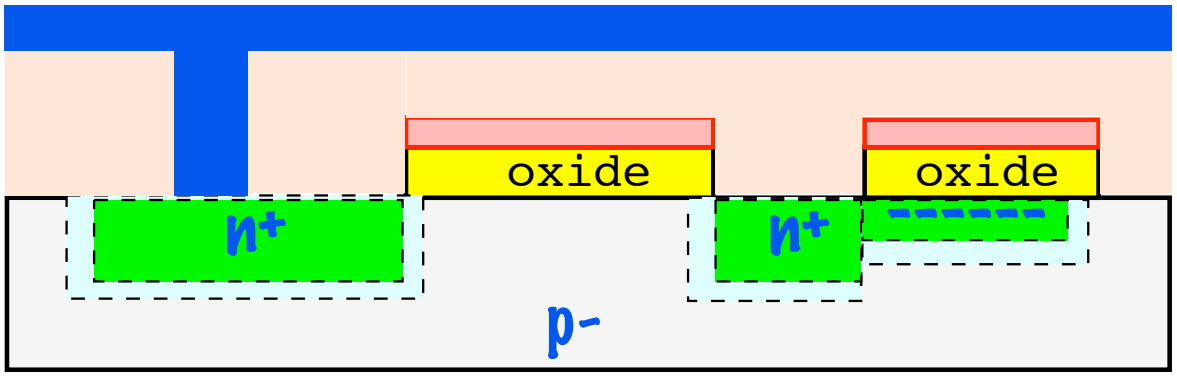

Word Line and Vdd run on "z-axis"

Why Vcap values start out at ground.

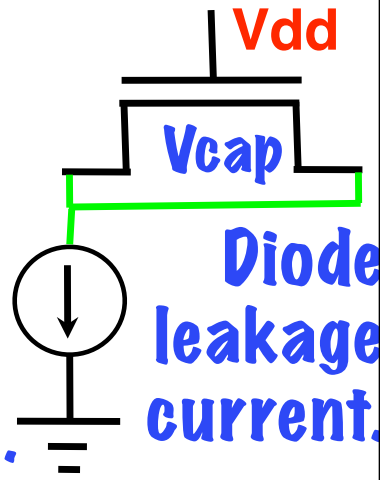

A 4 x 4 DRAM array (16 bits) ....

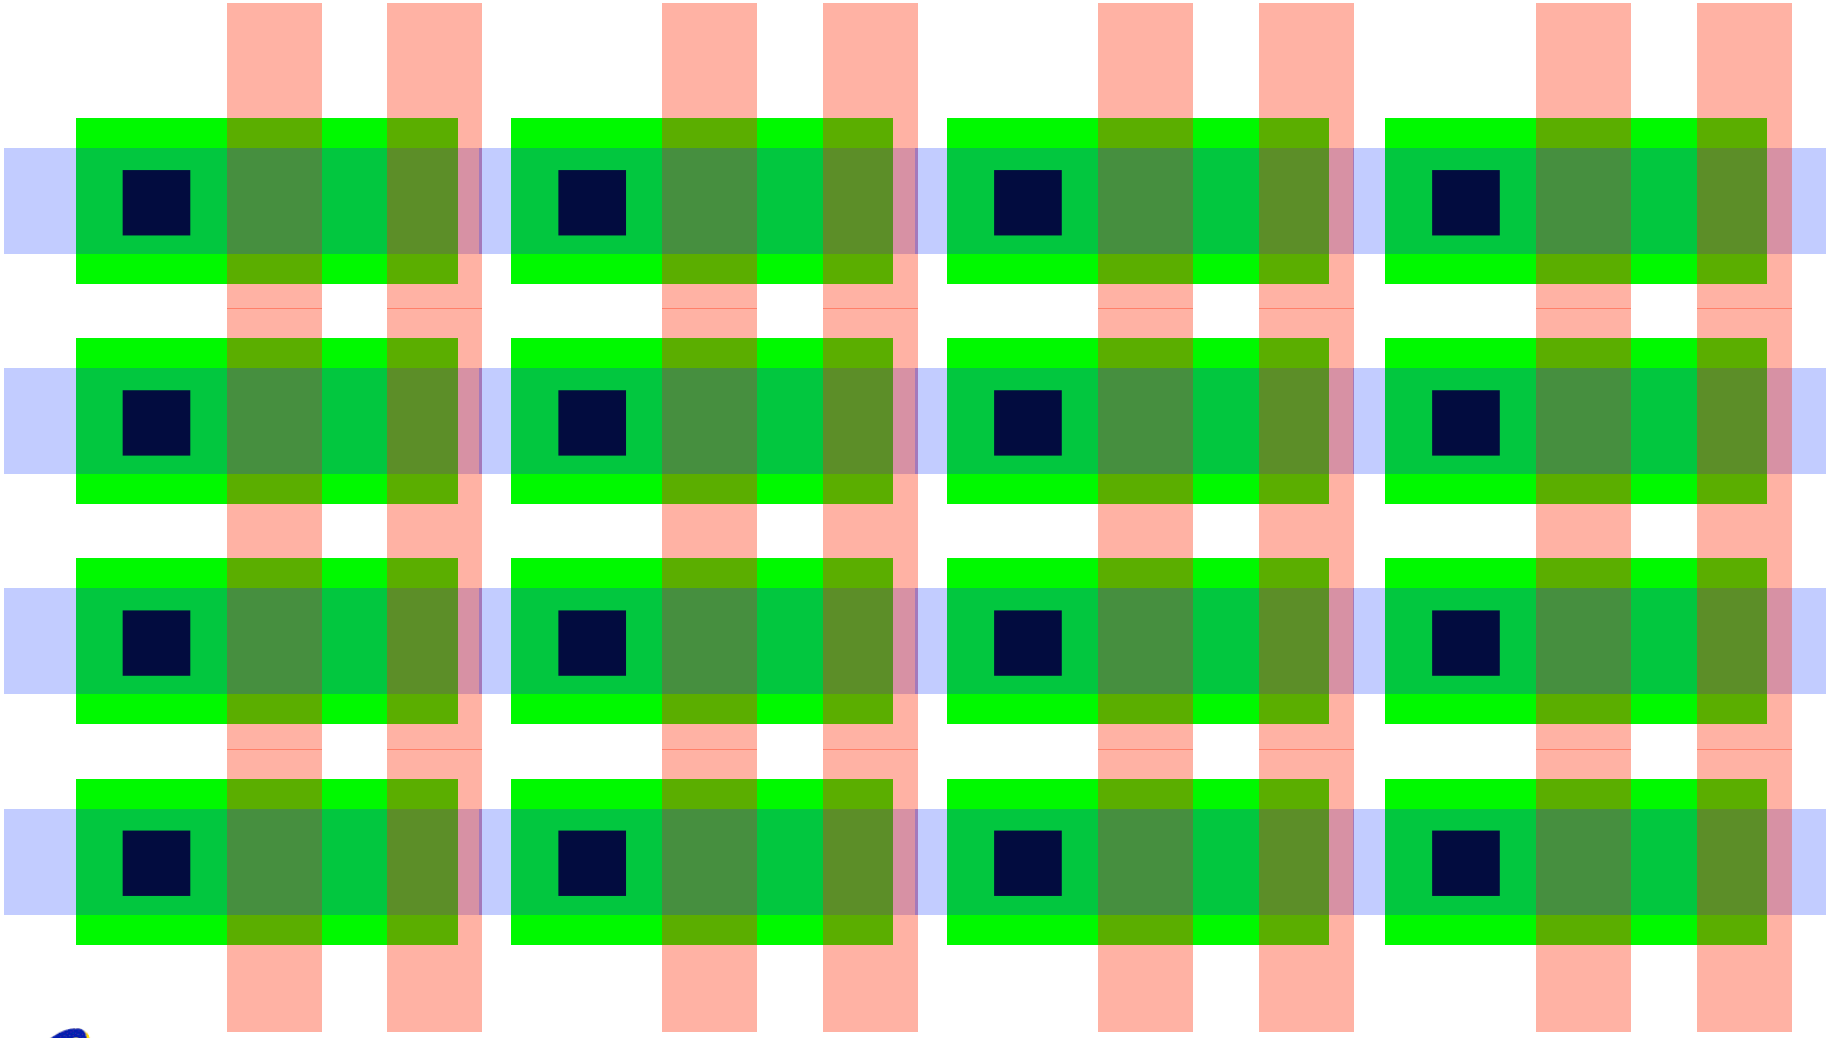

Invented after SRAM, by Robert Dennard

United States Patent Office

3,387,286
Patented June 4, 1968

3,387,286
FIELD-EFFECT TRANSISTOR MEMORY
Robert H. Dennard, Croton-on-Hudson, N.Y., assignor to International Business Machines Corporation, Armonk, N.Y., a corporation of New York
Filed July 14, 1967, Ser. No. 653,415
21 Claims. (Cl. 340—173)

tinent in disclosing various concepts and structures which have been developed in the application of field-effect transistors to different types of memory applications, the primary thrust up to this time in conventional read-write random access memories has been to connect a plurality of field-effect transistors in each cell in a latch configuration. Memories of this type require a large number of active devices in each cell and therefore each cell re-

FIG. 1

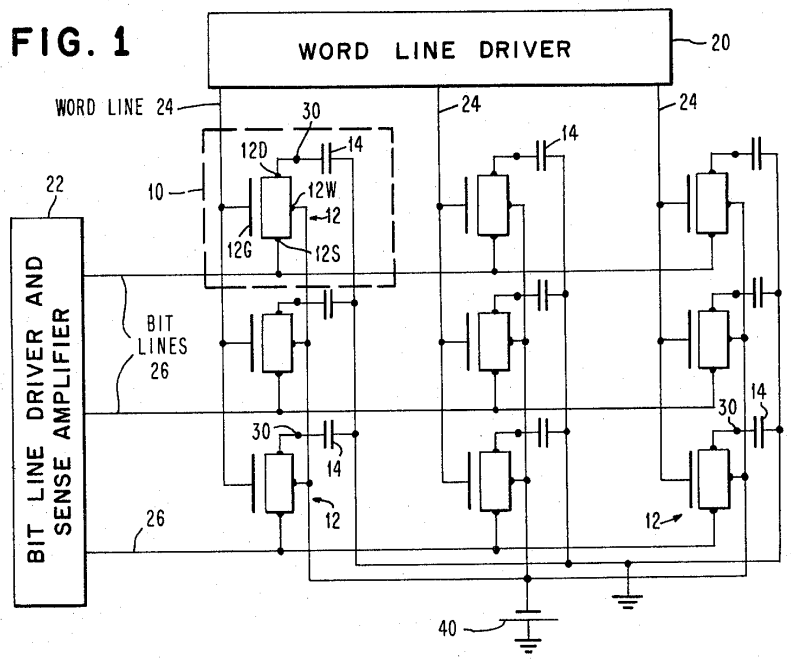

FIG. 2

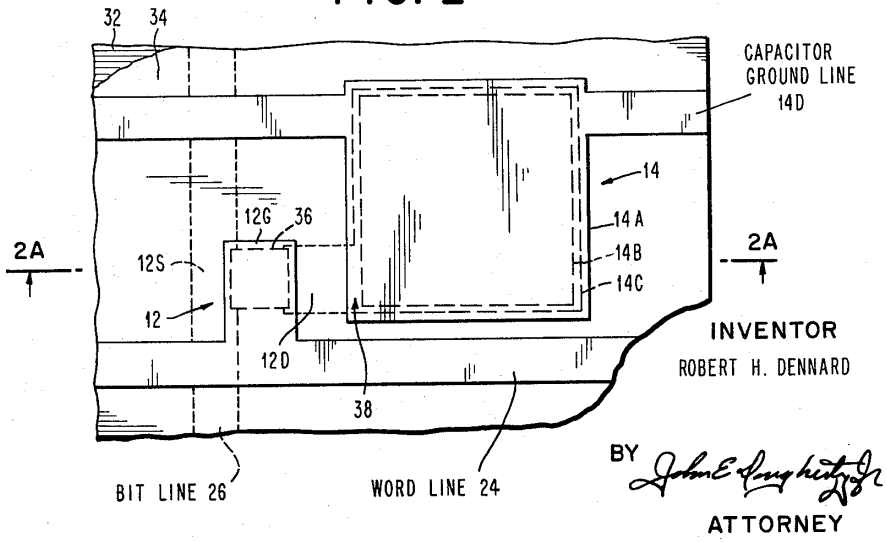

DRAM Circuit Challenge #1: Writing

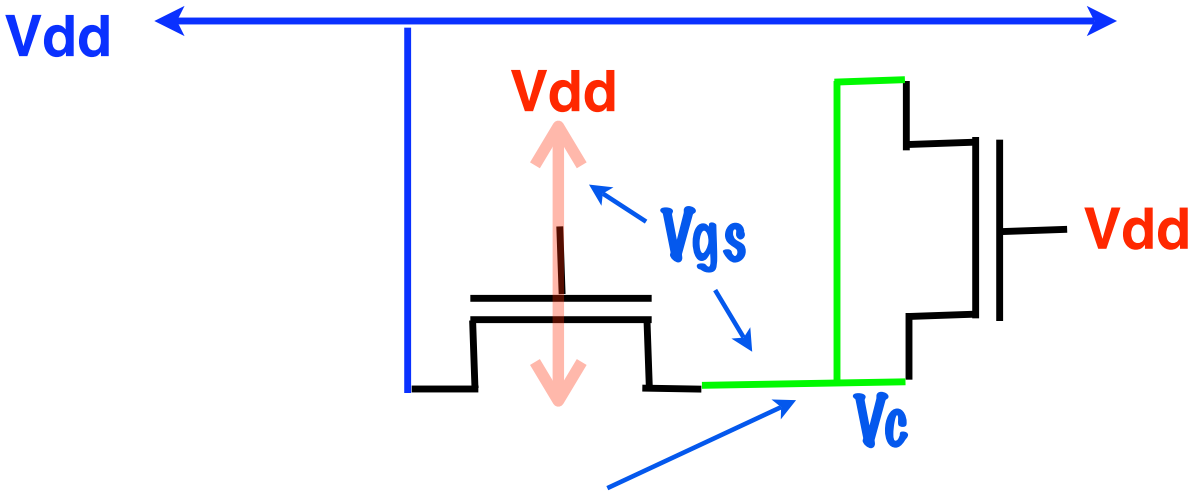

Vdd - Vth. Bad, we store less charge. Why do we not get Vdd?

Ids = [(μϵW)/(2LD)] [Vgs -Vth]^2 , but "turns off" when Vgs <= Vth!

Vgs = Vdd - Vc. When Vdd - Vc == Vth, charging effectively stops!

DRAM Challenge #2: Destructive Reads

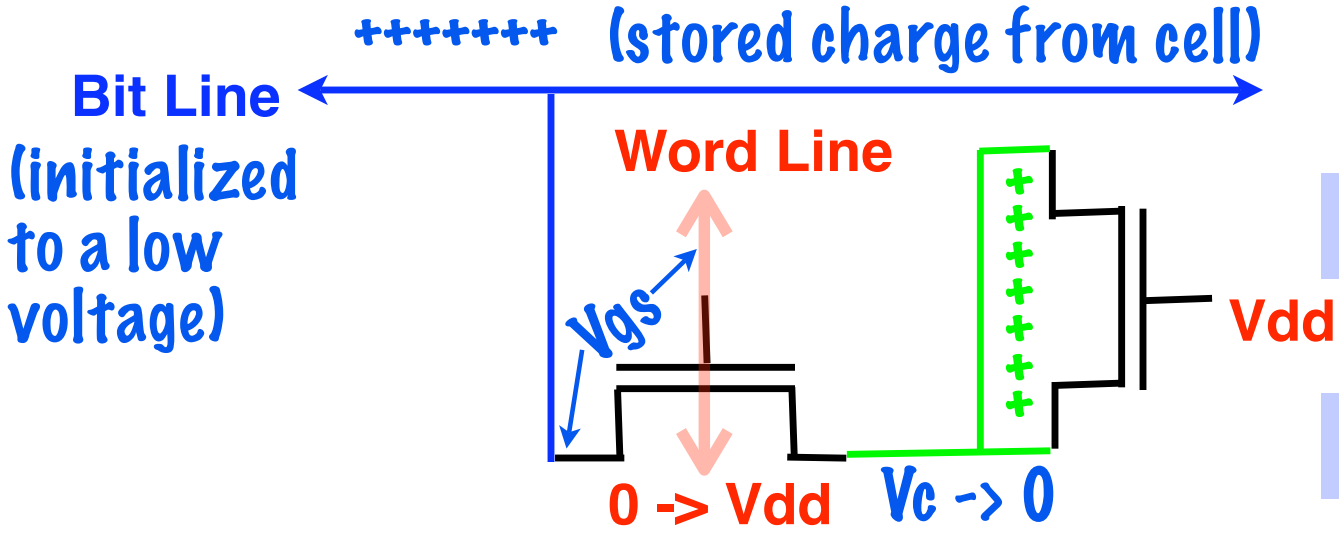

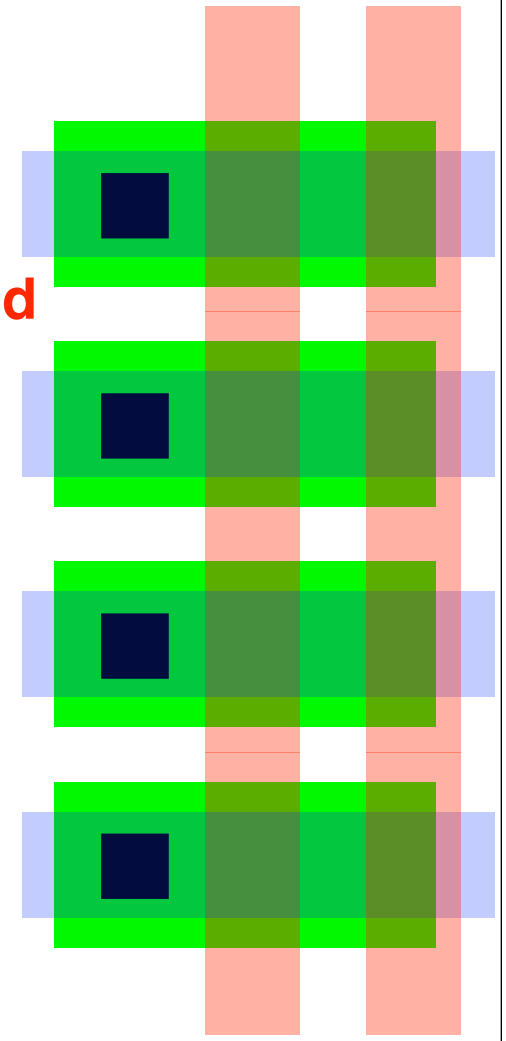

Raising the word line removes the charge from every cell it connects to!

Must write back after each read.

DRAM Circuit Challenge #3a: Sensing

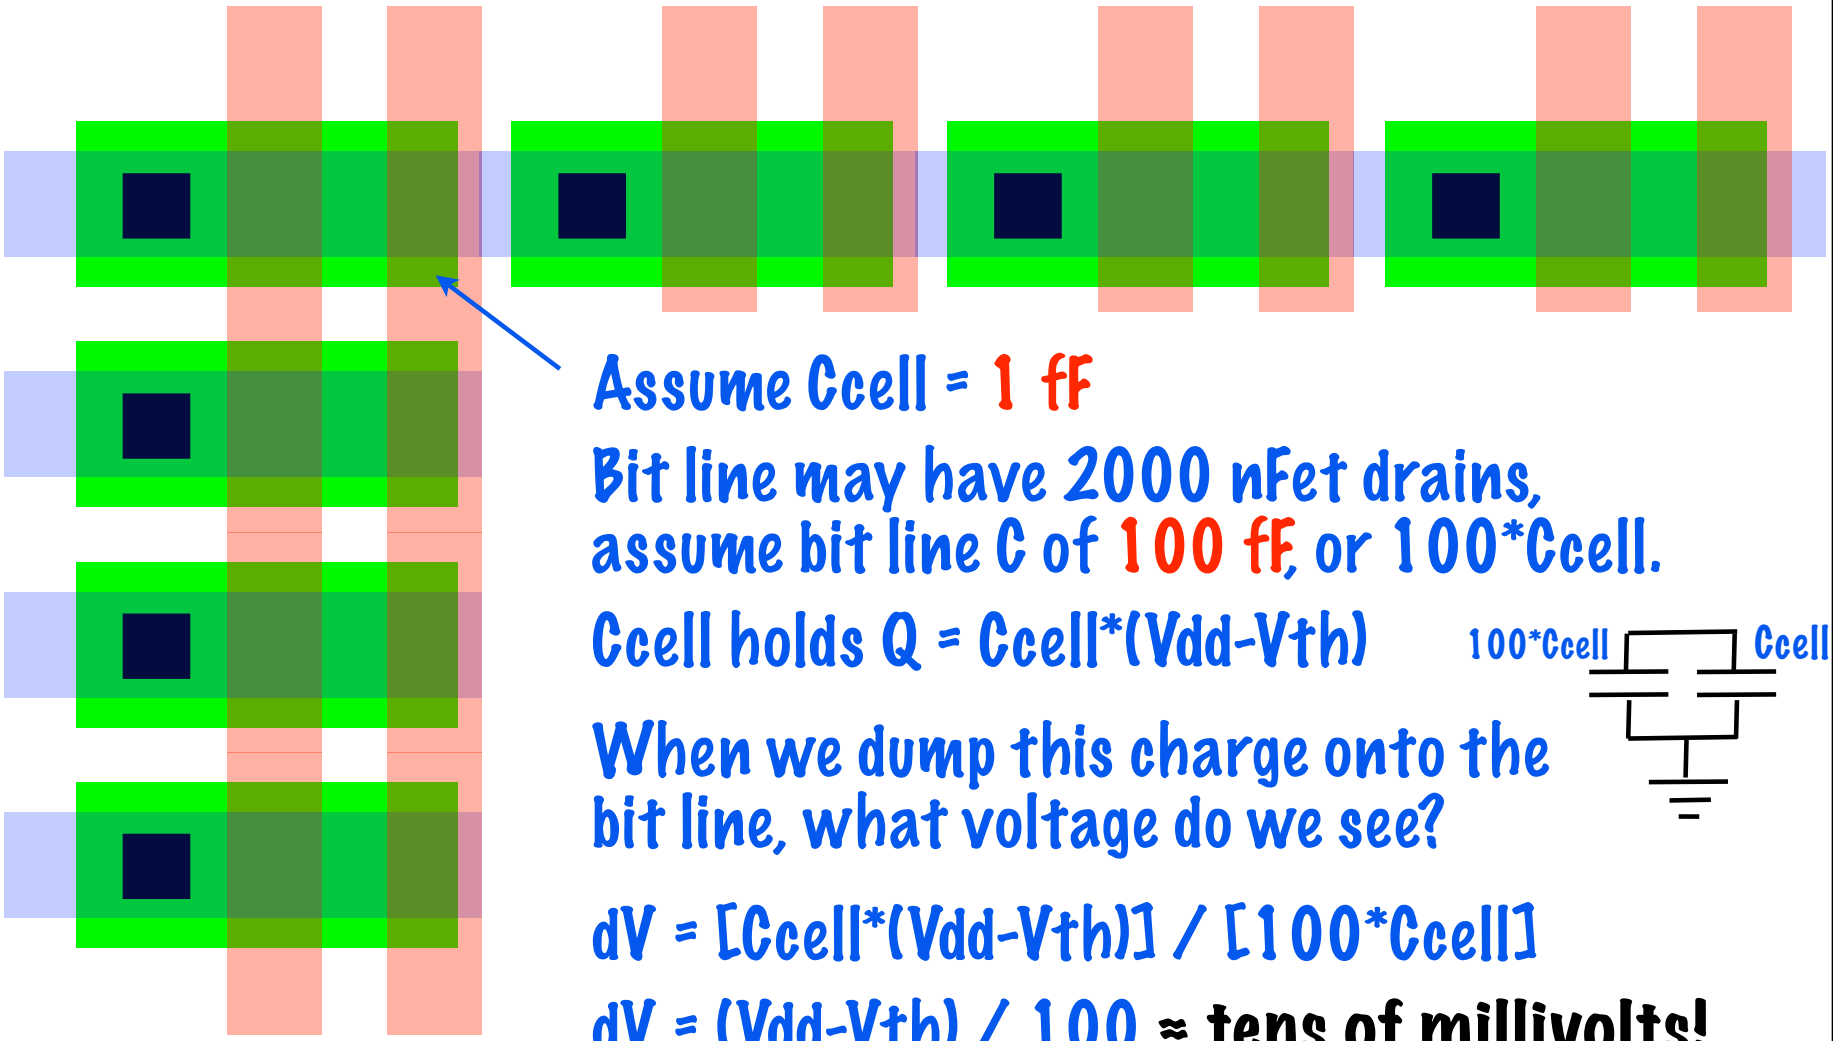

Assume Ccell = 1 fF

Bit line may have 2000 nFet drains, assume bit line C of 100 fF, or 100*Ccell.

Ccell holds Q = Ccell*(Vdd-Vth)

When we dump this charge onto the bit line, what voltage do we see?

dV = [Ccell*(Vdd-Vth)] / [100*Ccell]

dV = (Vdd-Vth) / 100 ≈ **tens of millivolts!**

In practice, scale array to get a 60mV signal.

DRAM Circuit Challenge #3b: Sensing

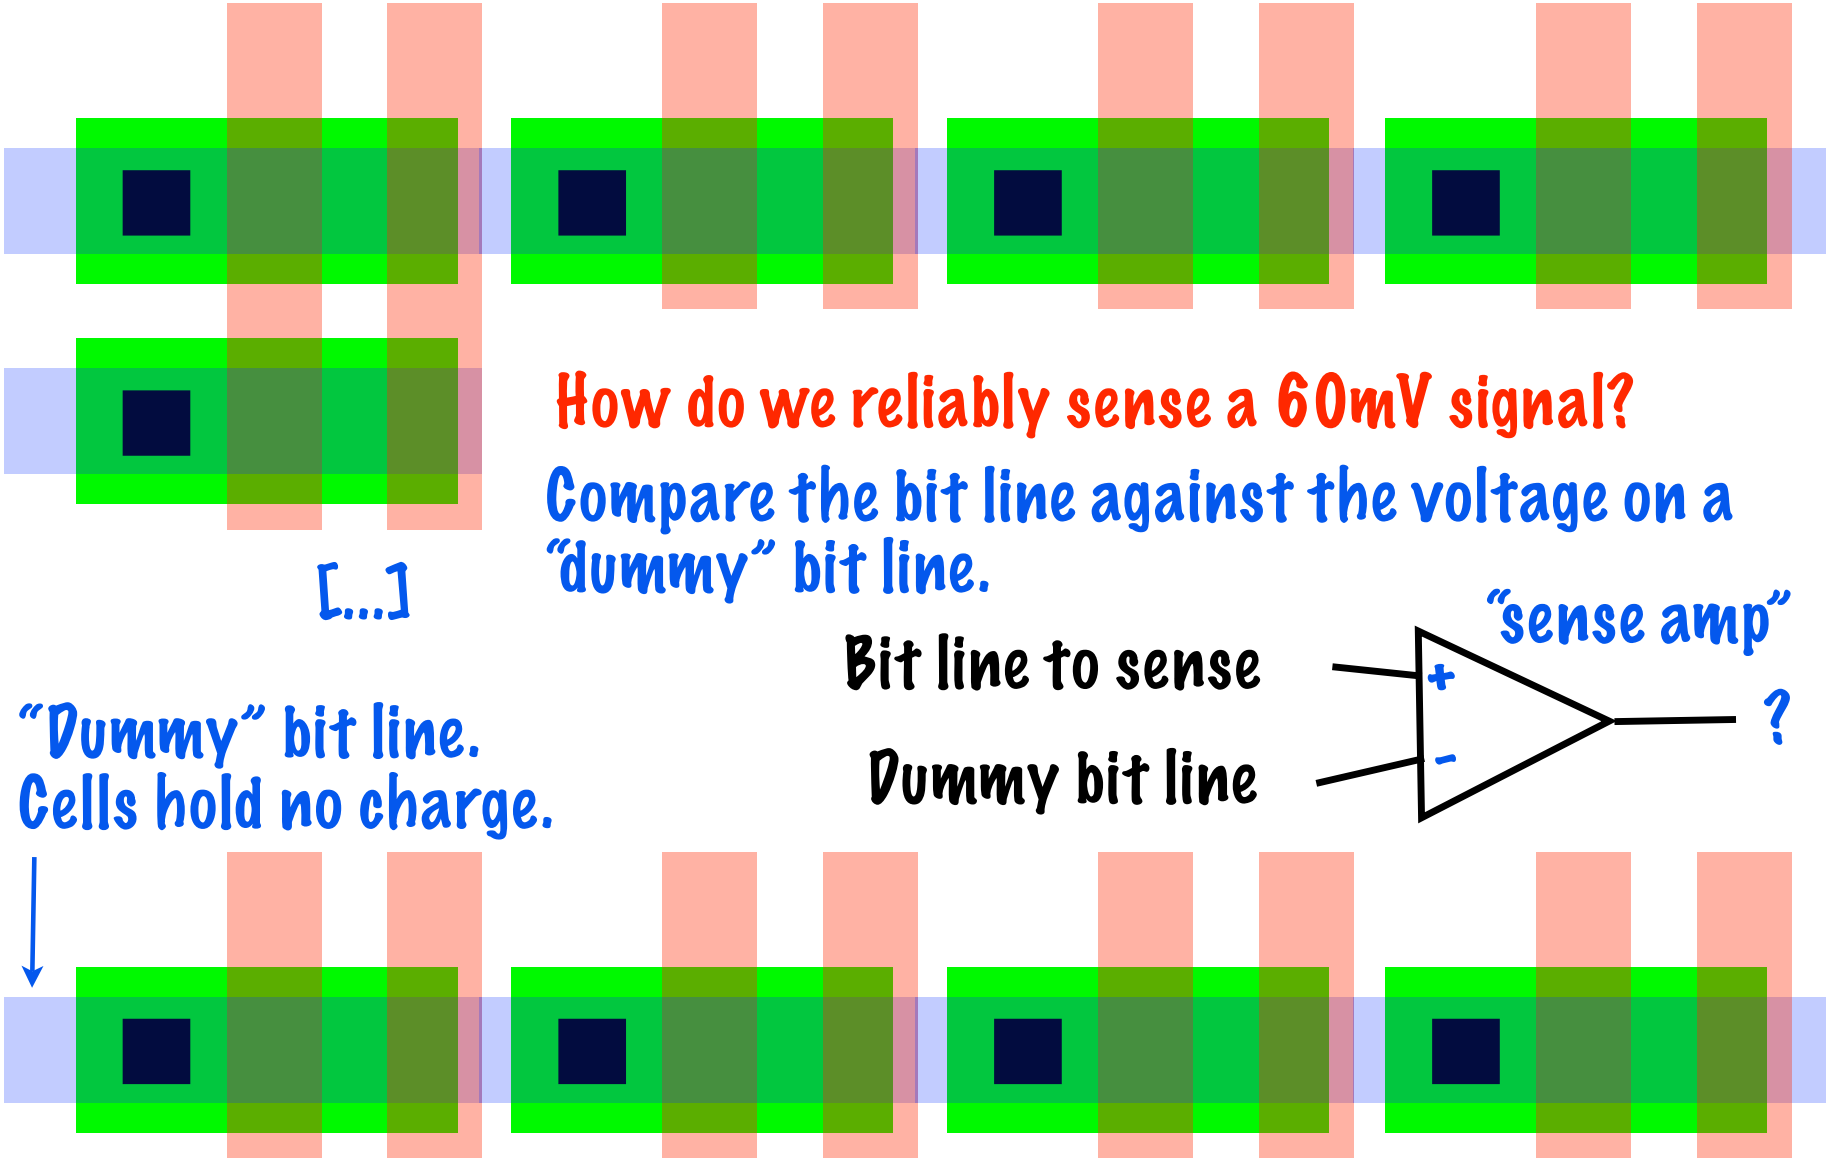

How do we reliably sense a 60mV signal?

Compare the bit line against the voltage on a "dummy" bit line.

[...]

"Dummy" bit line.
Cells hold no charge.

Bit line to sense

Dummy bit line

"sense amp"

DRAM Challenge #4: Leakage ...

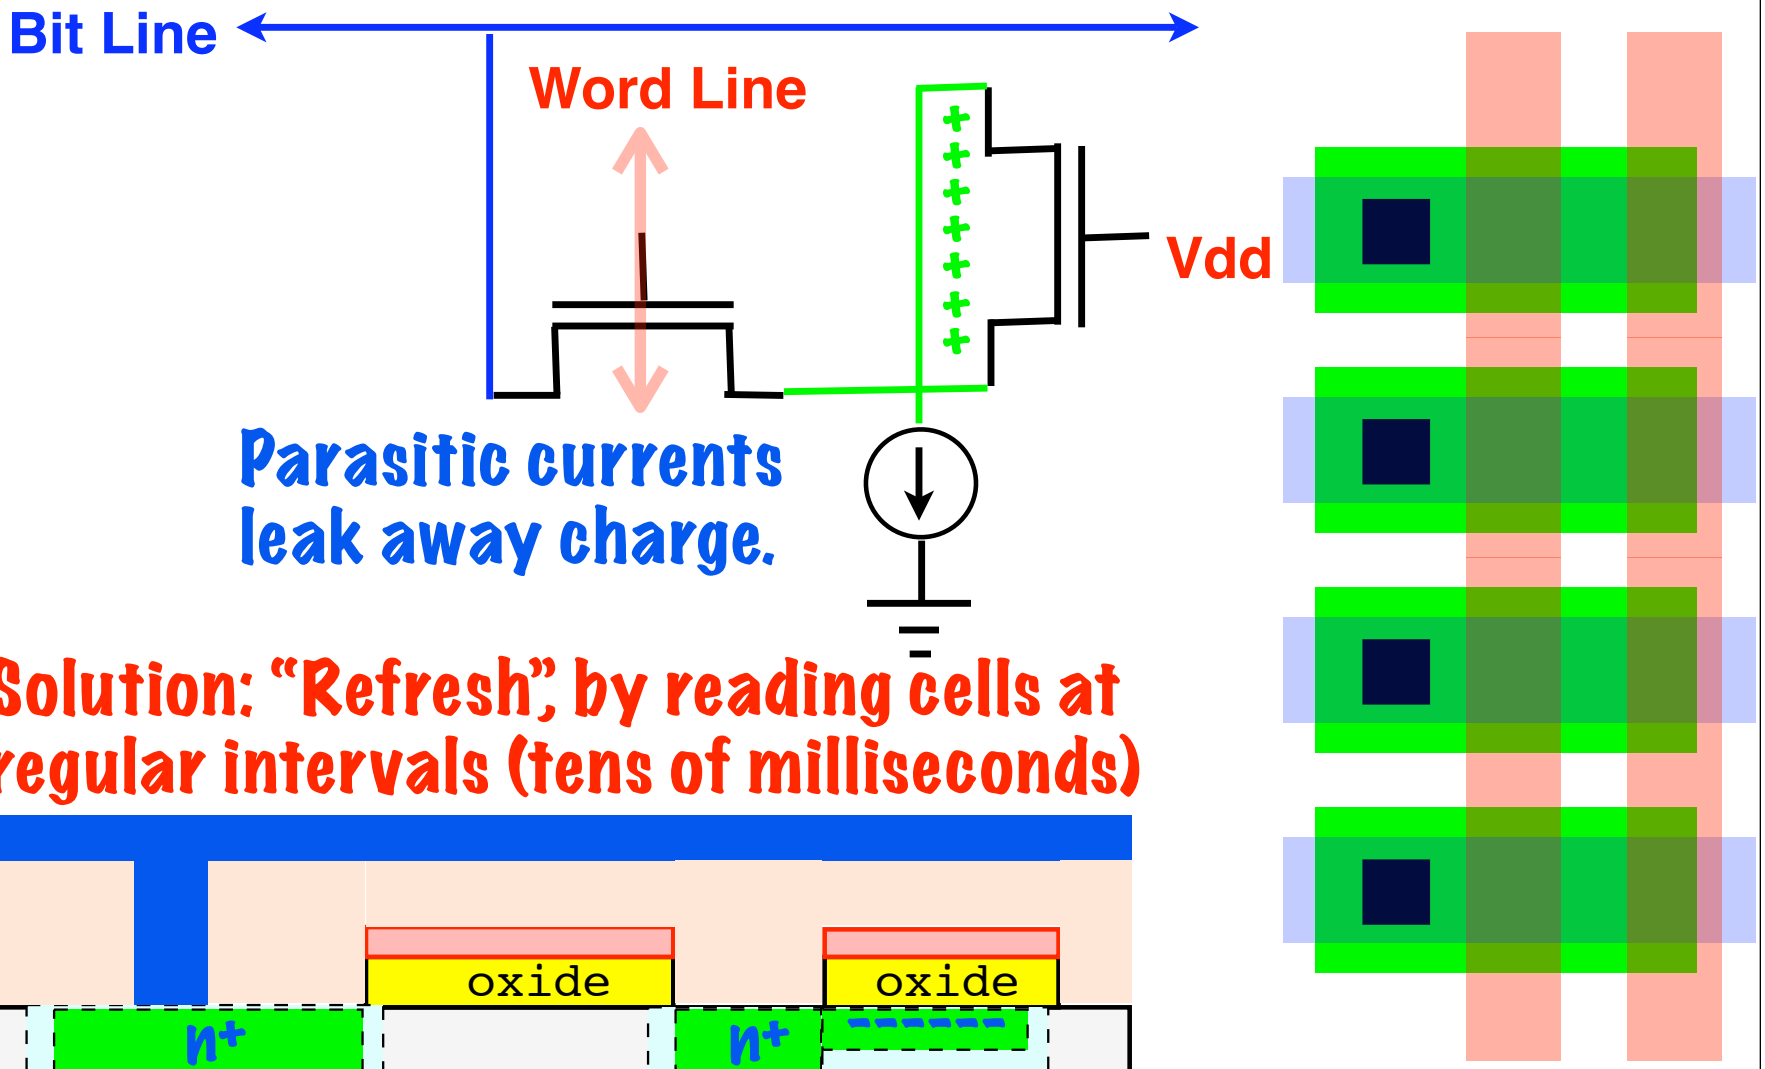

Solution: "Refresh", by reading cells at regular intervals (tens of milliseconds)

Diode leakage ...

DRAM Challenge #5: Cosmic Rays ...

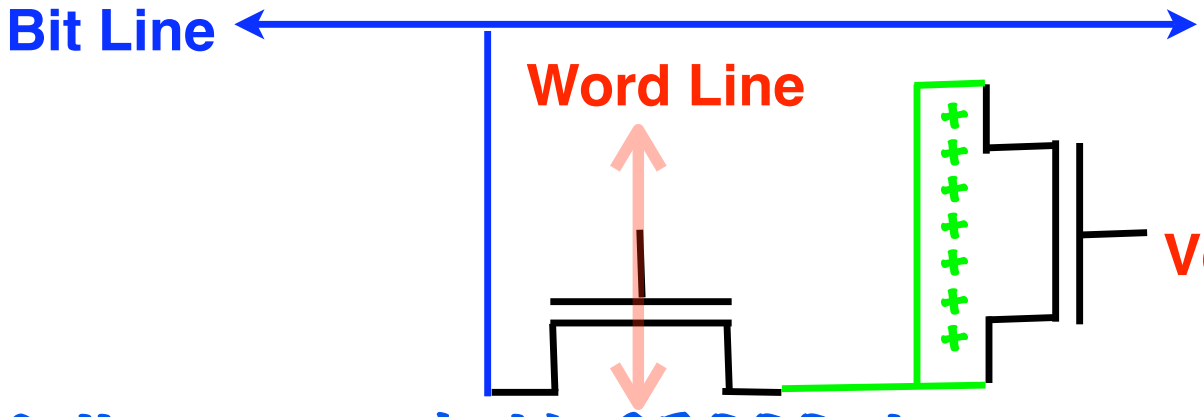

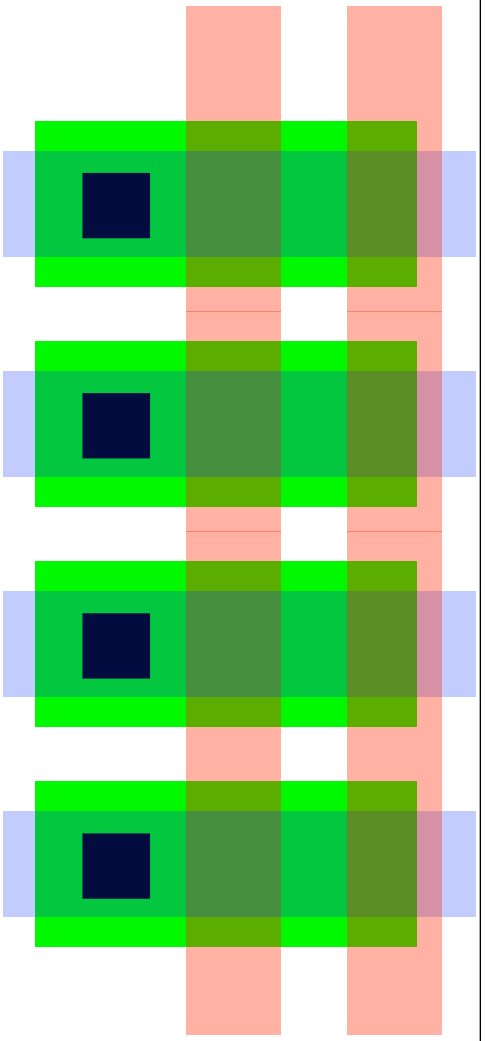

Cell capacitor holds 25,000 electrons (or less). Cosmic rays that constantly bombard us can release the charge!

Solution: Store extra bits to detect and correct random bit flips (ECC).

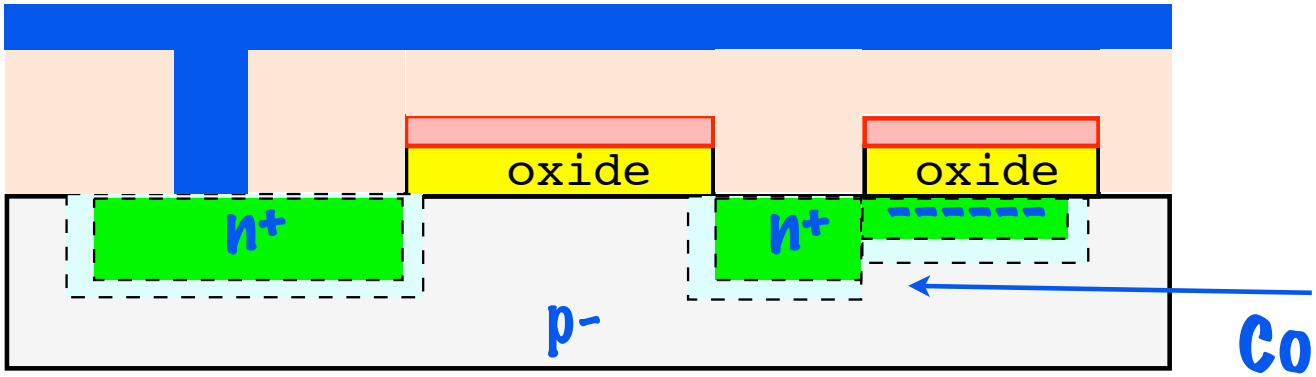

Cosmic ray hit.

DRAM Challenge 6: Yield

[...]

If one bit is bad, do we throw chip away?

Solution: add extra bit lines (i.e. 80 when you only need 64). During testing, find the bad bit lines, and use high current to burn away "fuses" put on chip to remove them.

Extra bit lines.
Used for "sparing".

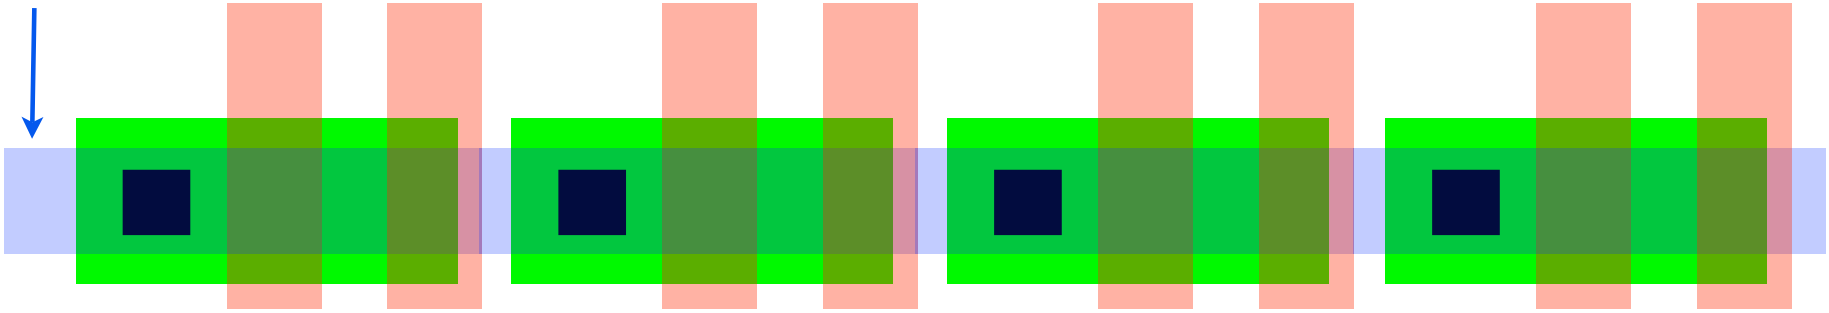

Recall: Process Scaling

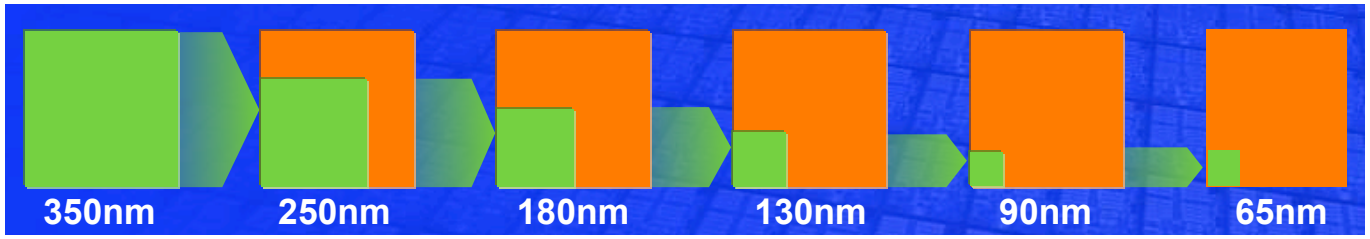

Recall process scaling ("Moore's Law")

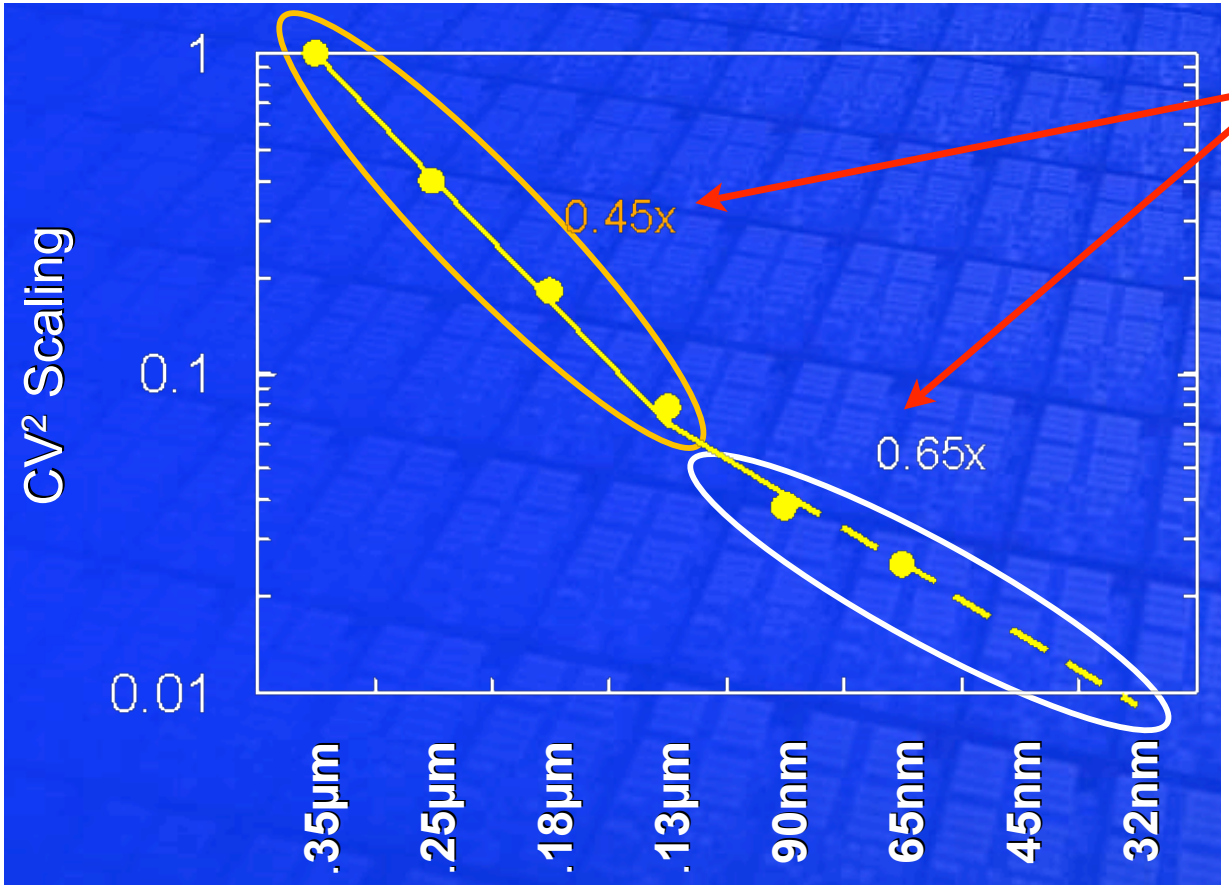

Due to reducing V and C (length and width of Cs decrease, but plate distance gets smaller).

Recent slope more shallow because V is being scaled less aggressively.

From: "Facing the Hot Chips Challenge Again", Bill Holt, Intel, presented at Hot Chips 17, 2005.

DRAM Challenge 7: Scaling

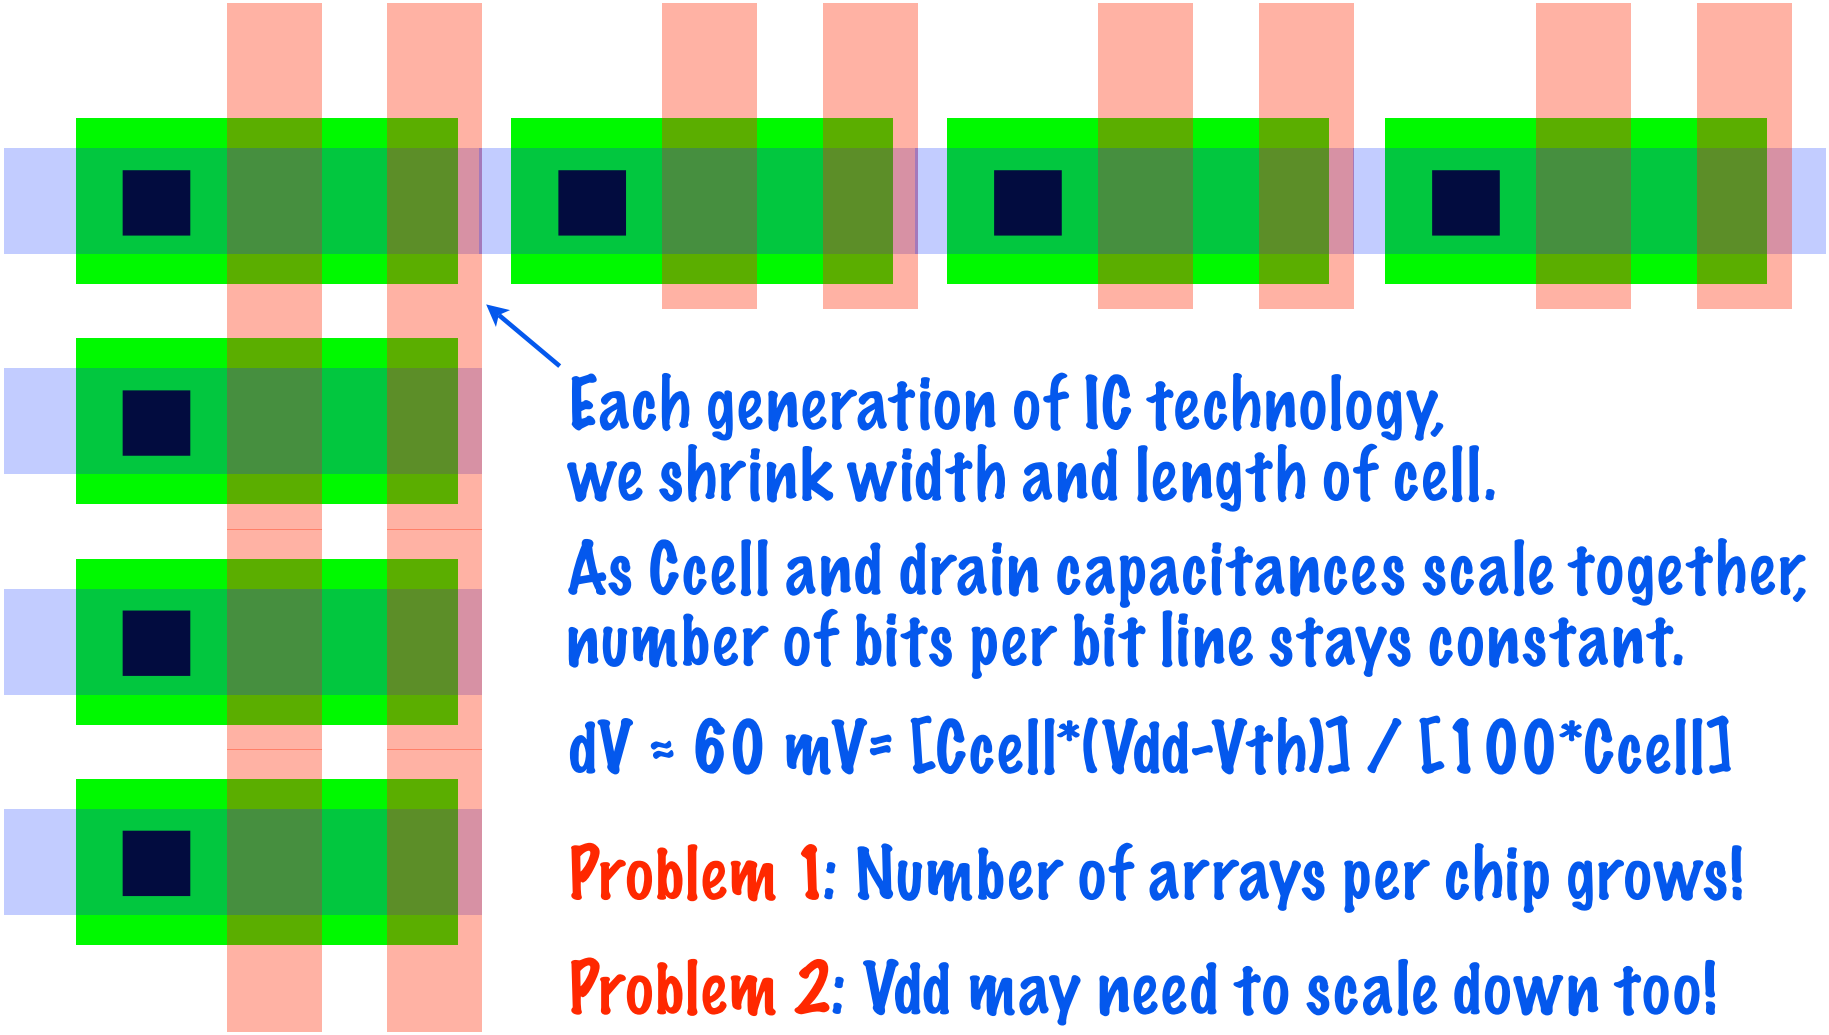

Each generation of IC technology, we shrink width and length of cell.

As Ccell and drain capacitances scale together, number of bits per bit line stays constant.

dV ≈ 60 mV= [Ccell*(Vdd-Vth)] / [100*Ccell]

Problem 1: Number of arrays per chip grows!

Problem 2: Vdd may need to scale down too!

Solution: Constant Innovation of Cell Capacitors!

Poly-diffusion Ccell is ancient history

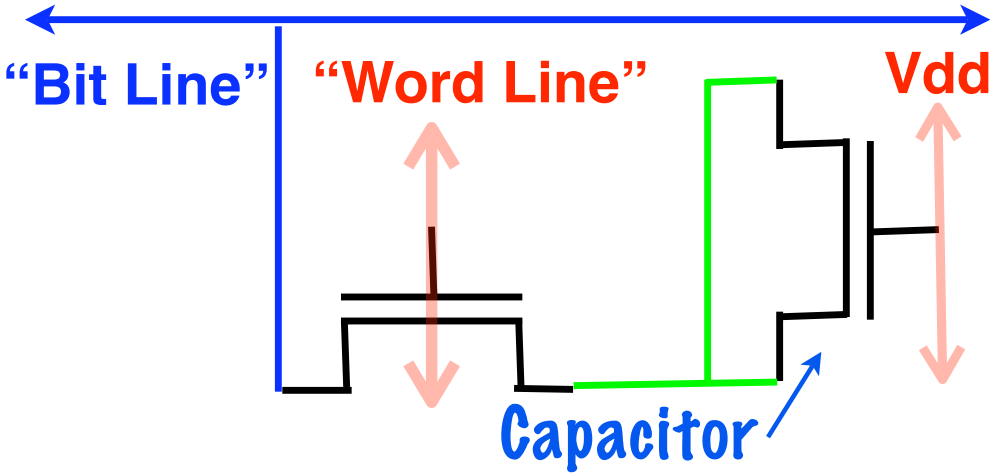

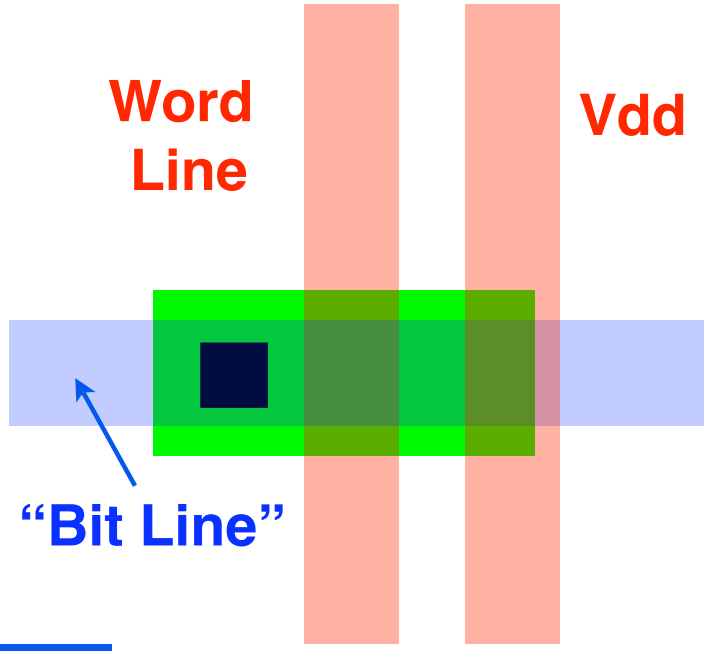

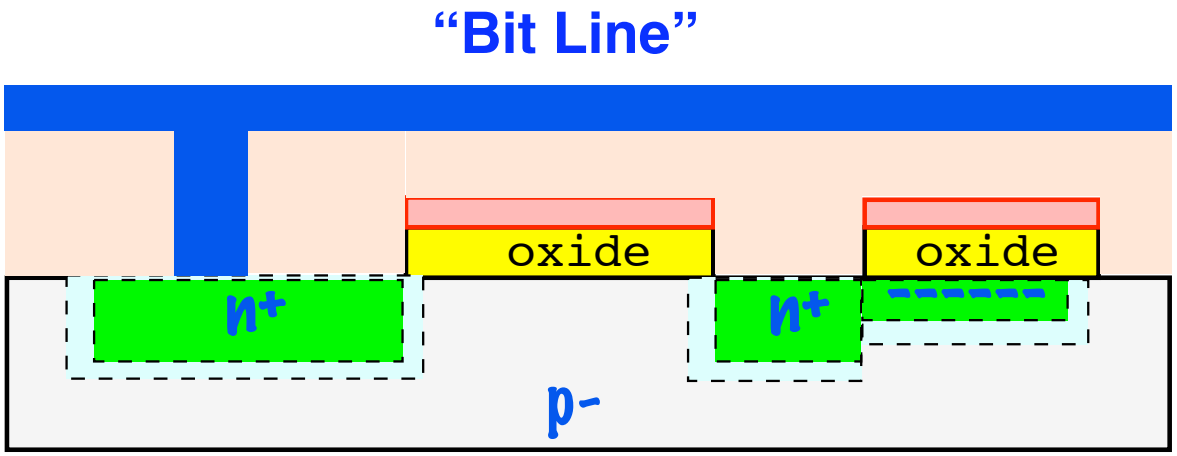

Word Line and Vdd run on "z-axis"

Early replacement: “Trench” capacitors

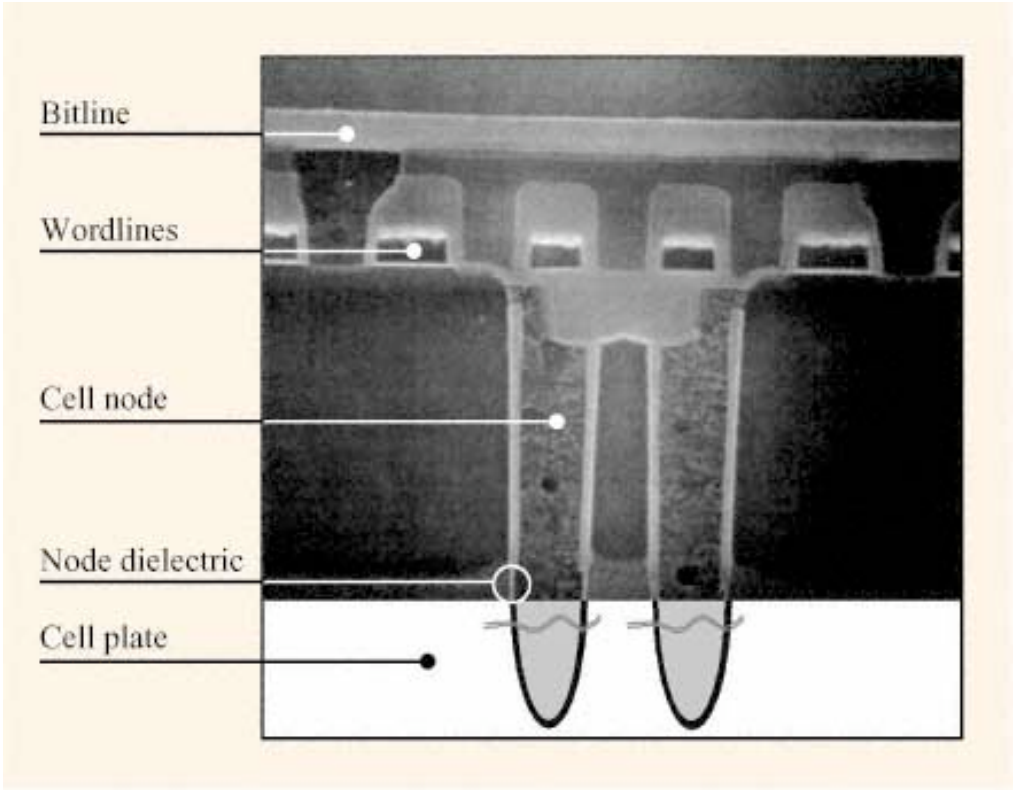

Figure 4

SEM photomicrograph of 0.25-μm trench DRAM cell suitable for scaling to 0.15μm and below. Reprinted with permission from [17]; © 1995 IEEE.

Final generation of trench capacitors

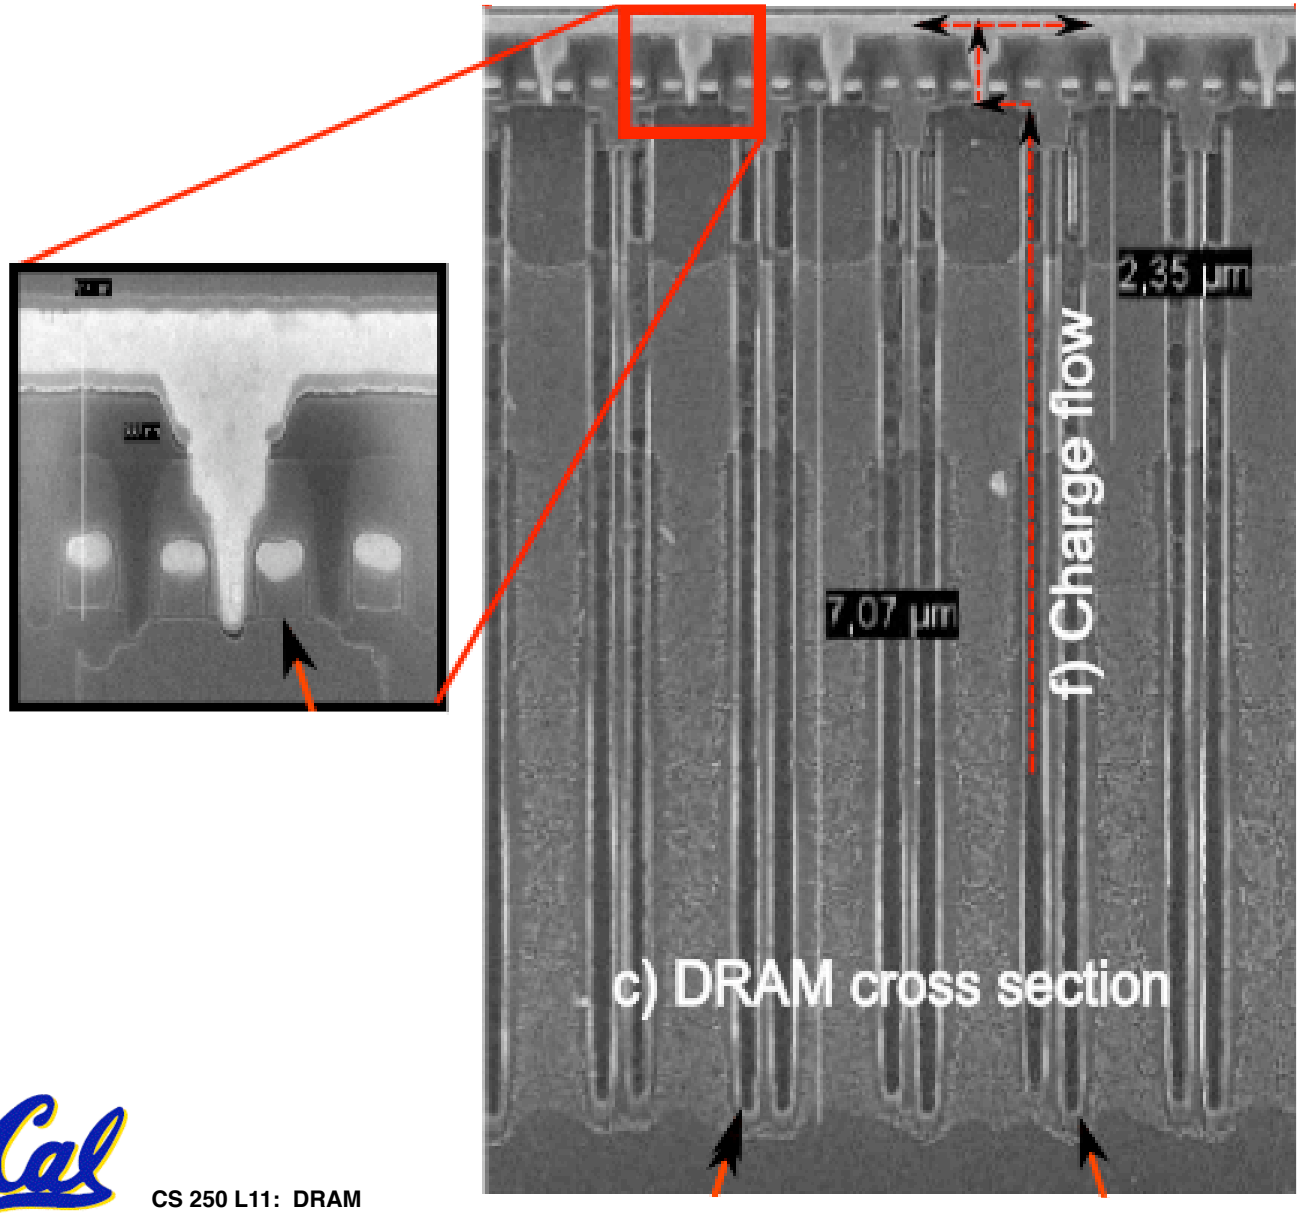

The companies that kept scaling trench capacitors are in the process of going out of business …

Modern cells: “stacked” capacitors

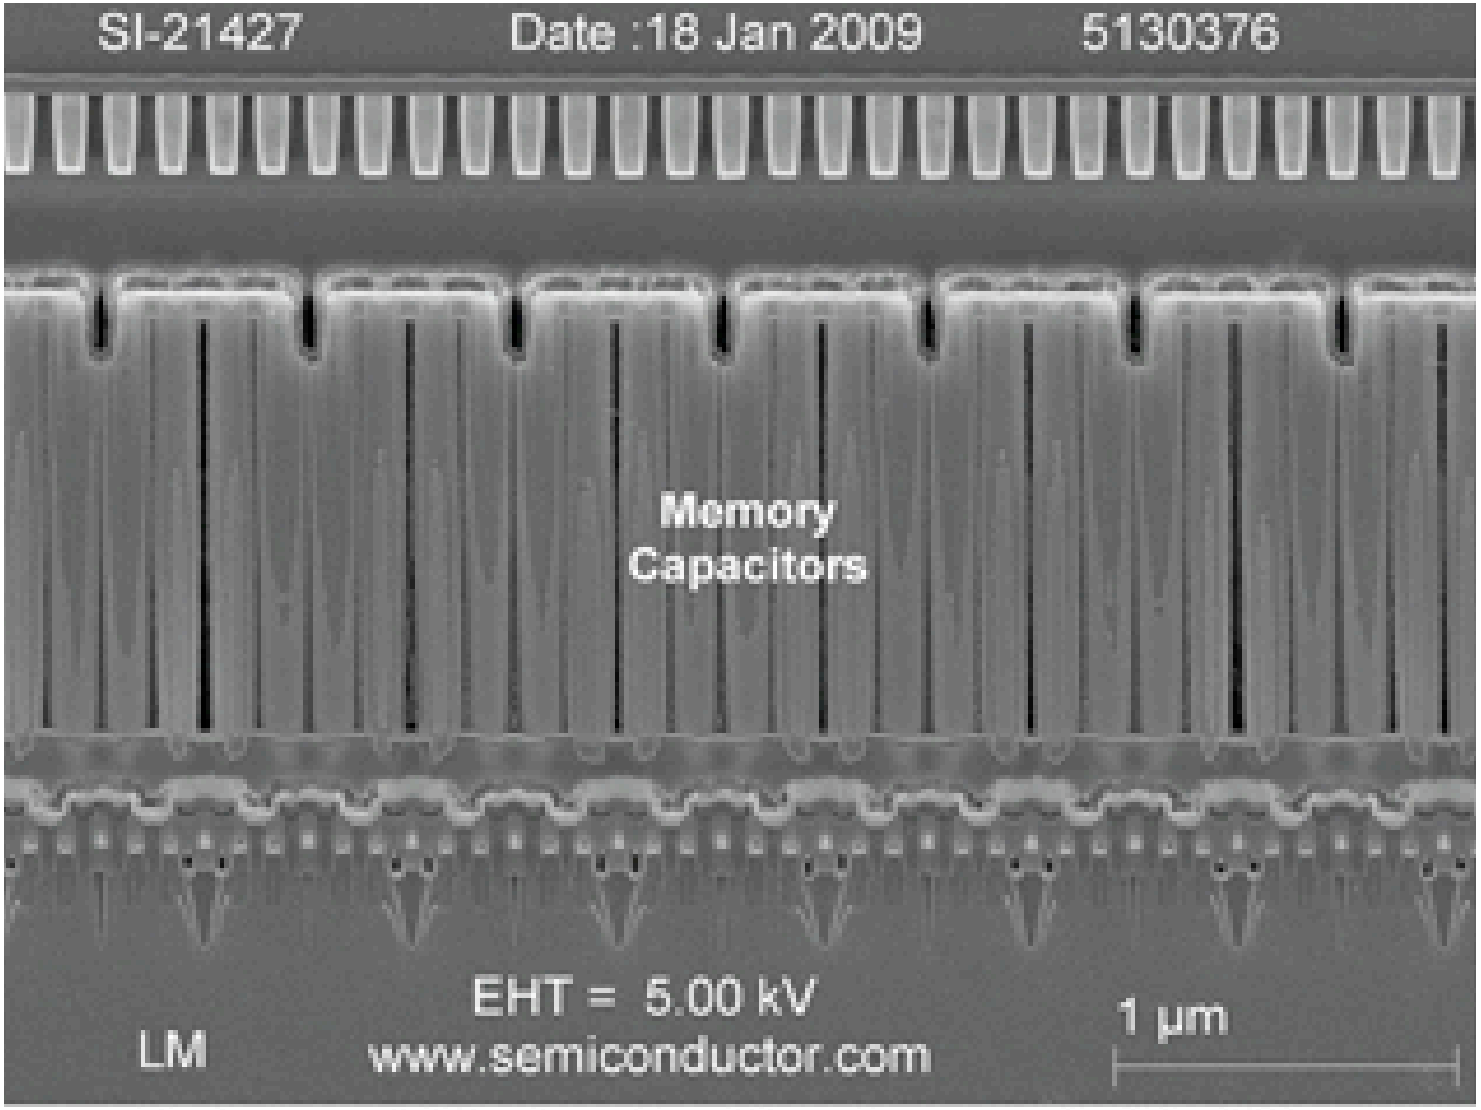

Micron 1-Gbit DDR2 50-nm SDRAM

Micron 50nm 1-Gbit DDR2 die photo

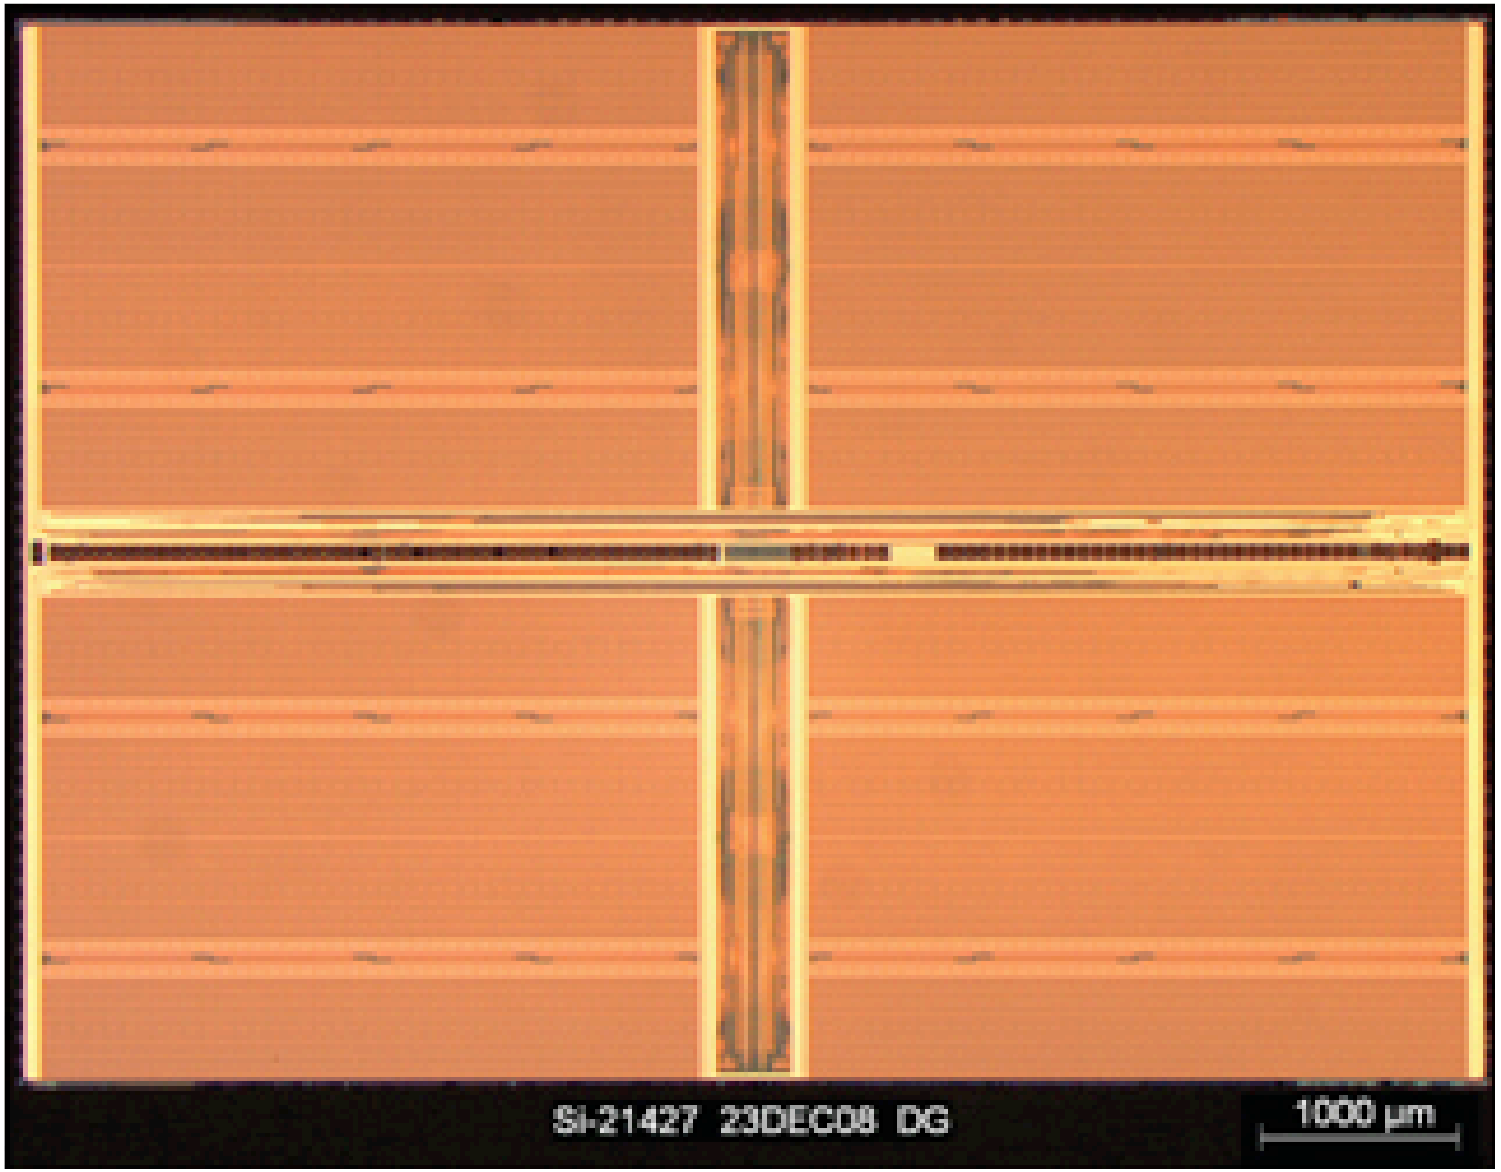

Memory Arrays

Older SDRAM part: 133 Mhz, 128 Mb

SYNCHRONOUS DRAM

MT48LC32M4A2 – 8 Meg x 4 x 4 banks
MT48LC16M8A2 – 4 Meg x 8 x 4 banks
MT48LC8M16A2 – 2 Meg x 16 x 4 banks

For the latest data sheet, please refer to the Micron Web site: www.micron.com/dramds

A "bank" of 32 Mb (128Mb chip -> 4 banks)

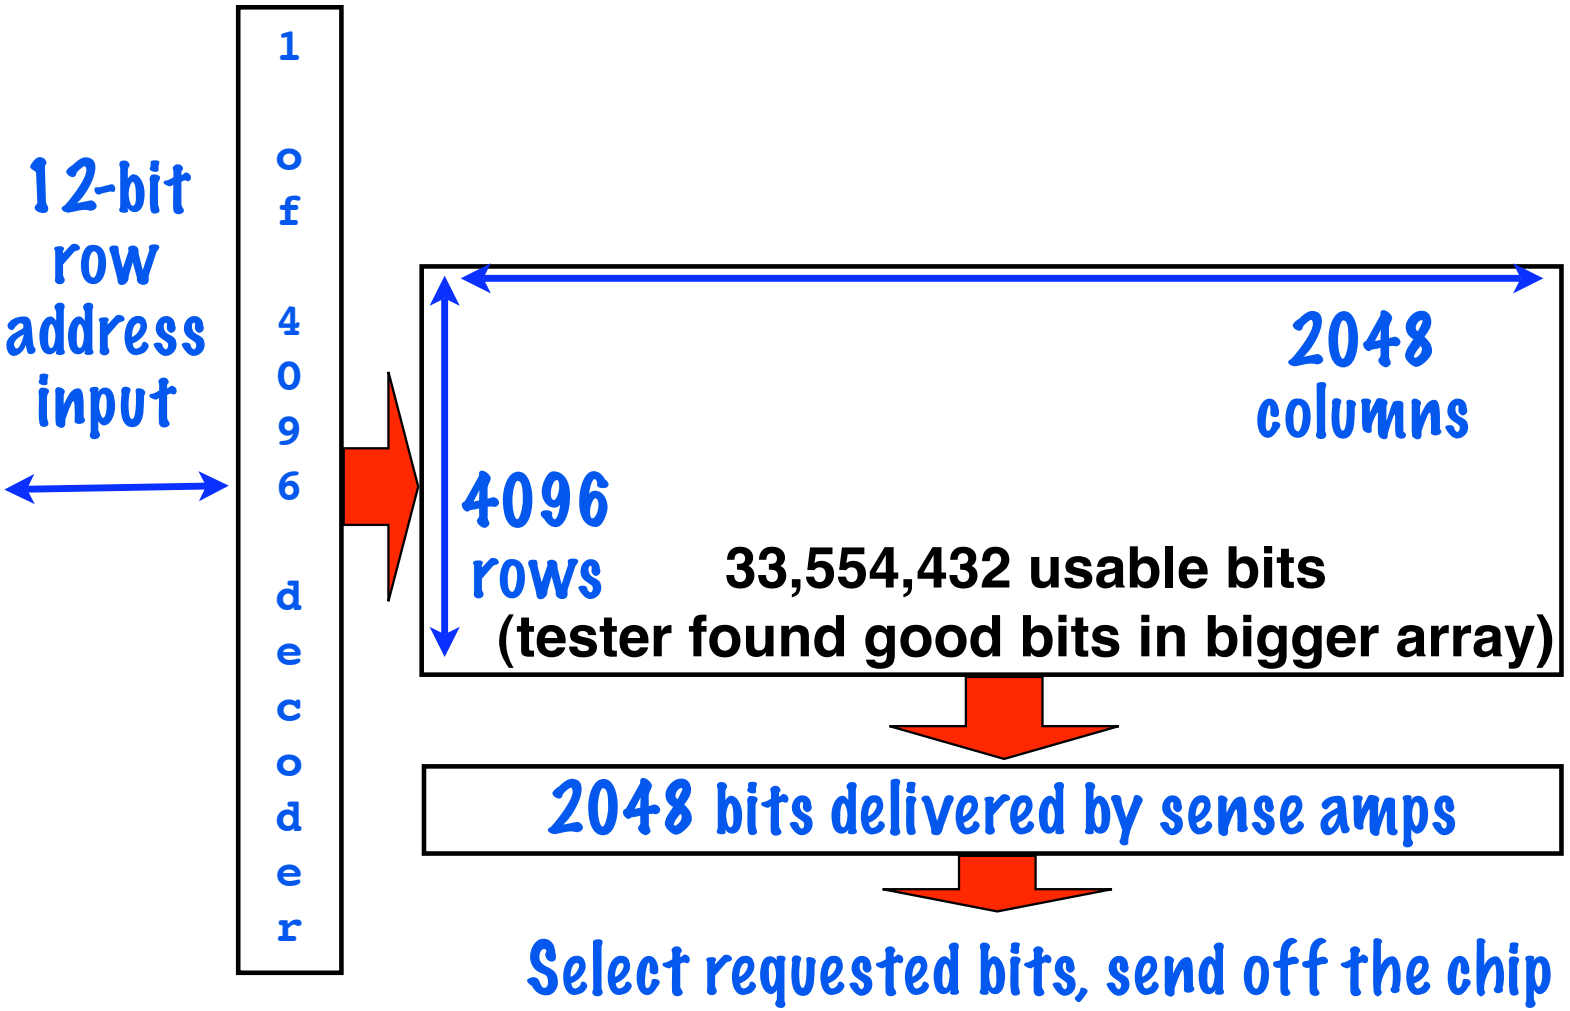

Recall DRAM Challenge #3b: Sensing

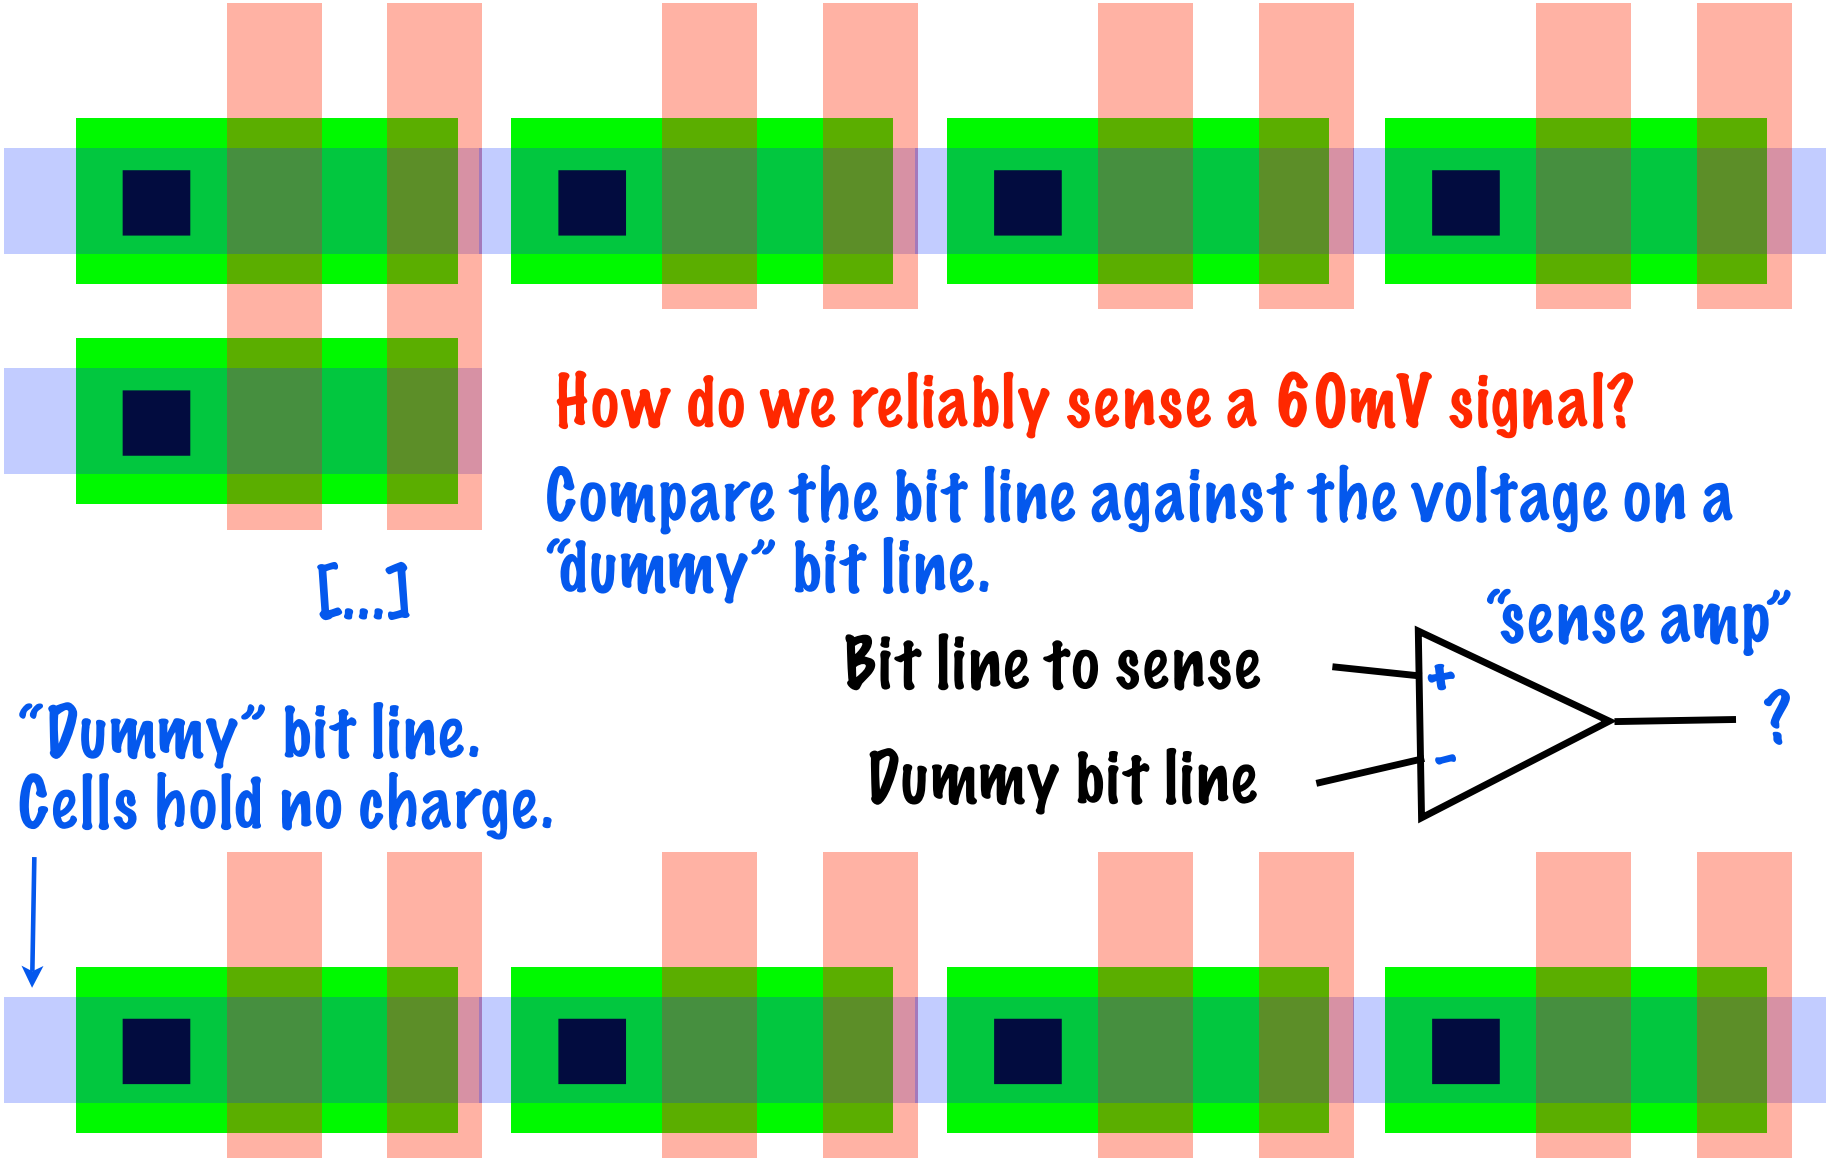

How do we reliably sense a 60mV signal?

Compare the bit line against the voltage on a "dummy" bit line.

[...]

"Dummy" bit line.
Cells hold no charge.

Bit line to sense

Dummy bit line

"sense amp"

"Sensing" is row read into sense amps

Slow! This 7.5ns period DRAM (133 MHz) can do row reads at only 75 ns (13 MHz).

DRAM has high latency to first bit out. A fact of life.

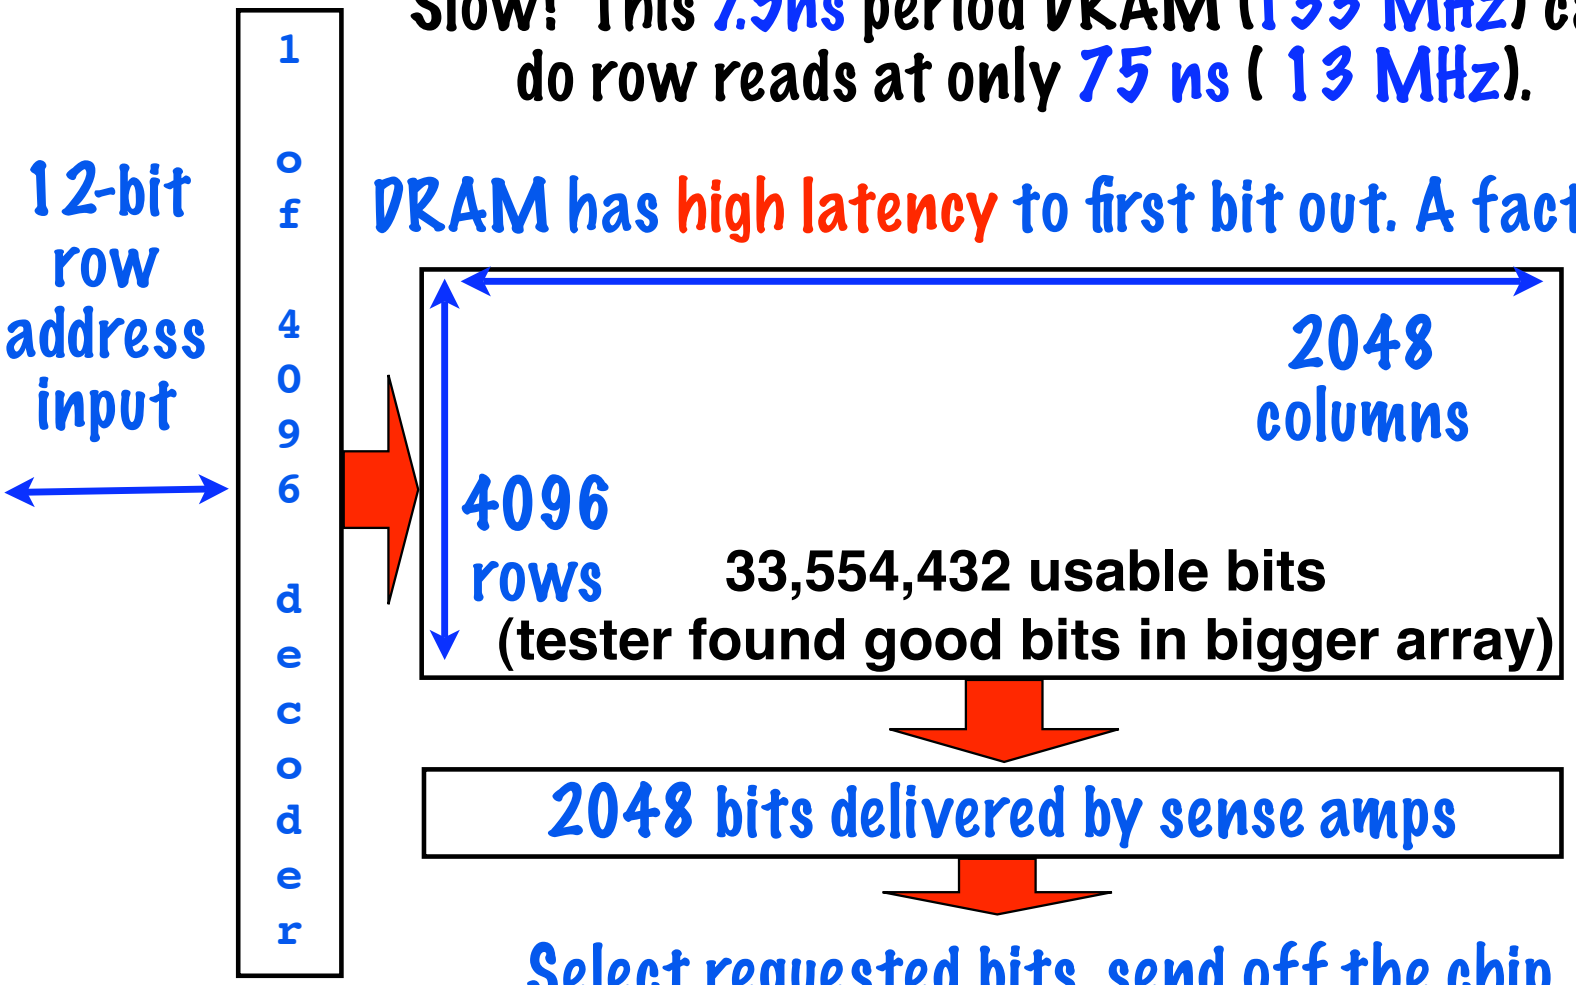

An ill-timed refresh may add to latency

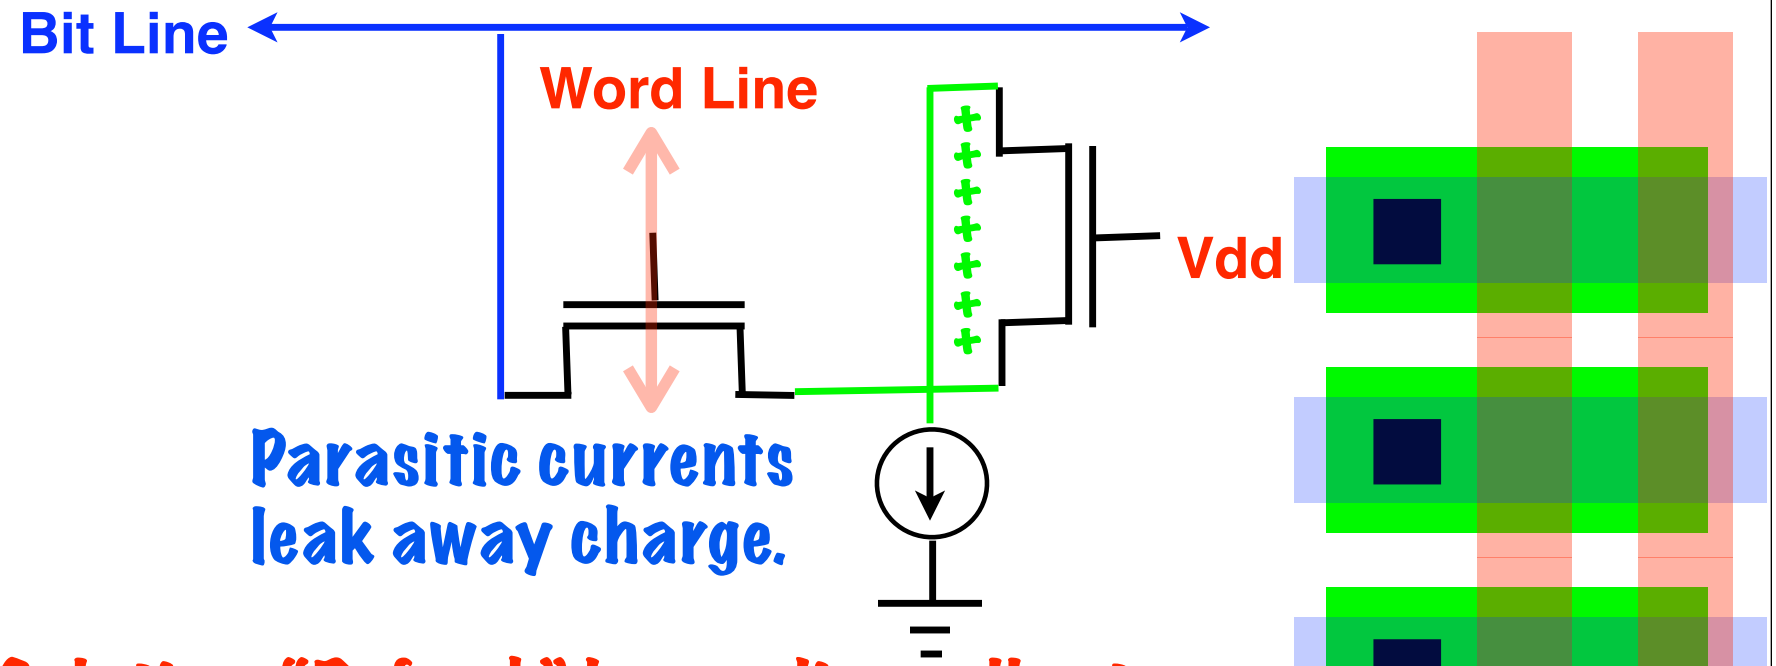

Solution: "Refresh", by reading cells at regular intervals (tens of milliseconds)

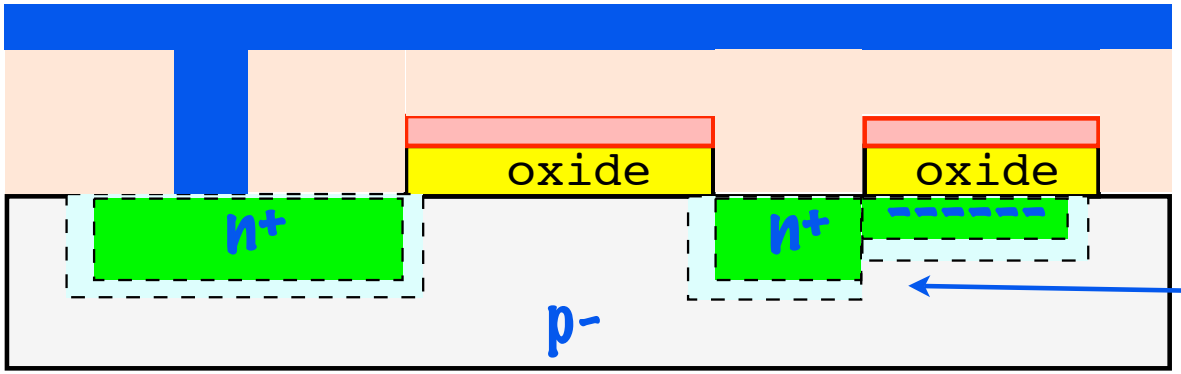

Diode leakage ...

Latency is not the same as bandwidth!

Thus, push to faster DRAM interfaces

What if we want all of the 2048 bits?

In row access time (75 ns) we can do 10 transfers at 133 MHz.

8-bit chip bus -> 10 x 8 = 80 bits << 2048

Now the row access time looks fast!

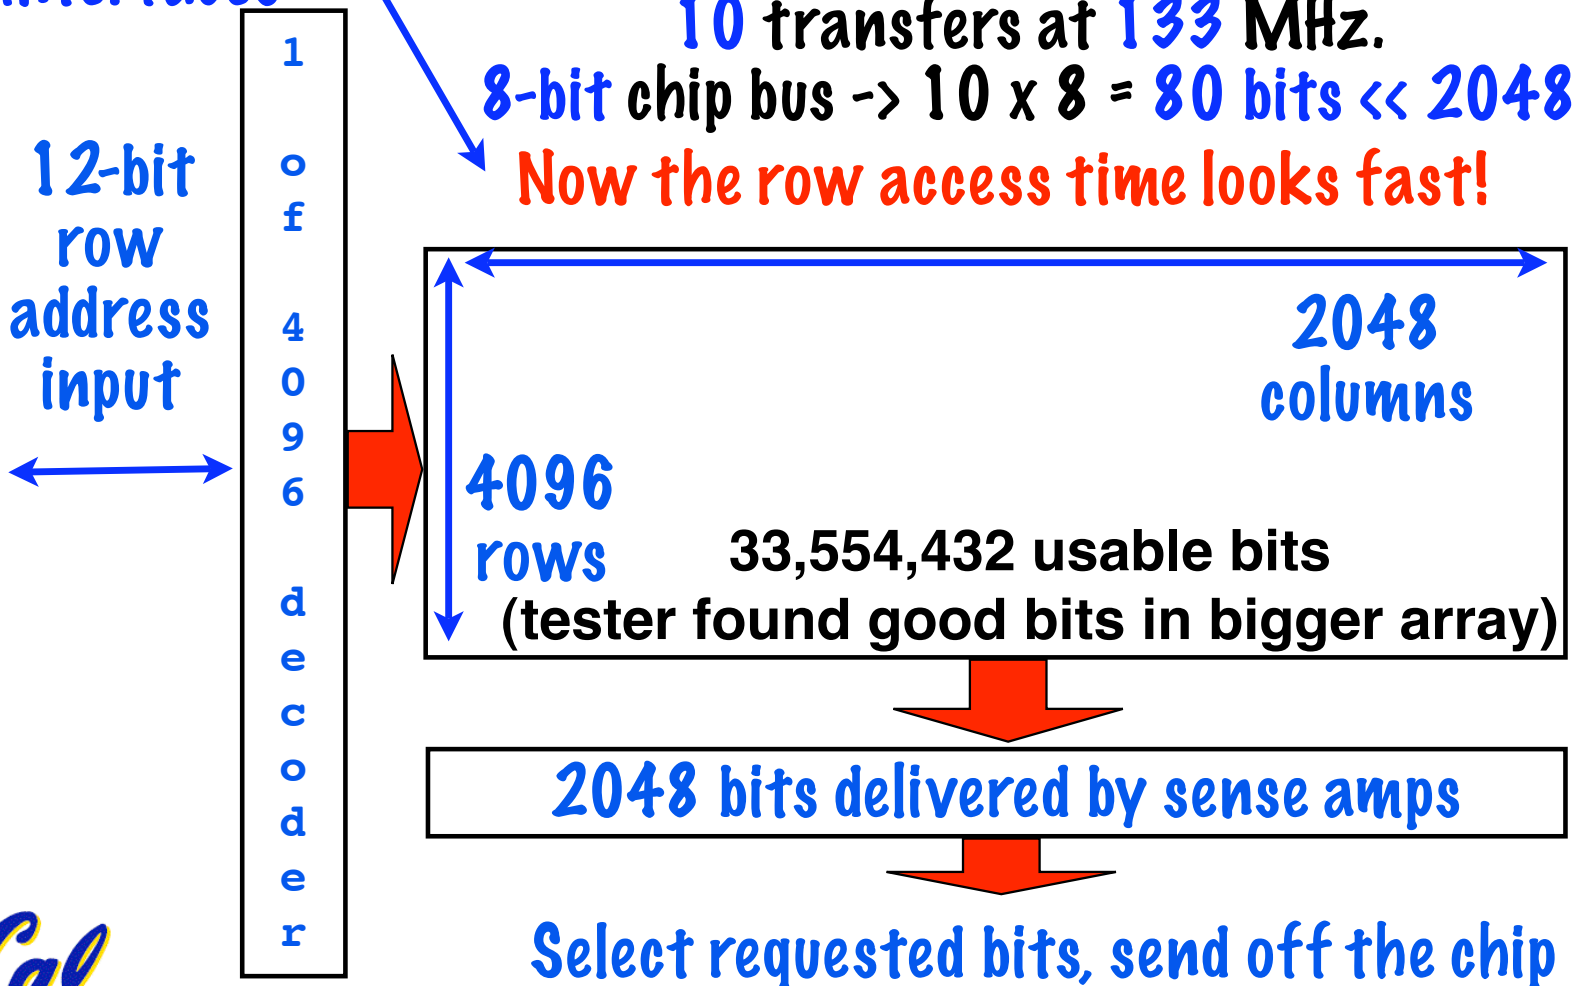

Sadly, it's rarely this good ...

What if we want all of the 2048 bits?

The "we" for a CPU would be the program running on the CPU.

Recall Amdalh's law: If 20% of the memory accesses need a new row access ... not good.

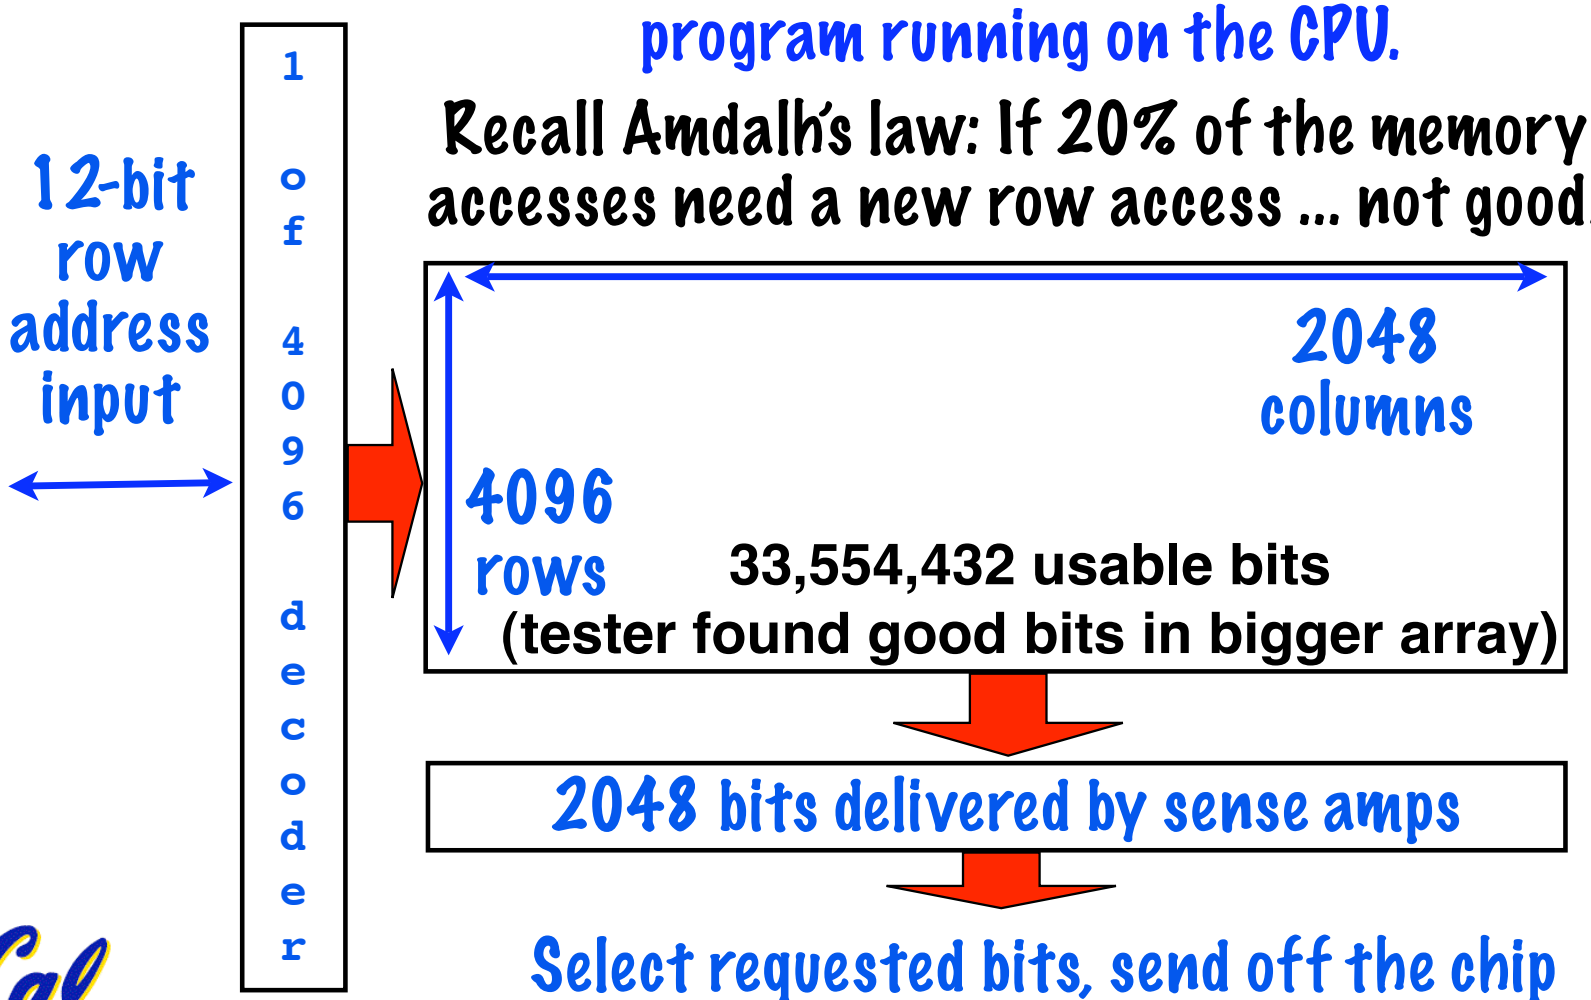

DRAM latency/bandwidth chip features

- **Columns:** Design the right interface for CPUs to request the subset of a column of data it wishes:

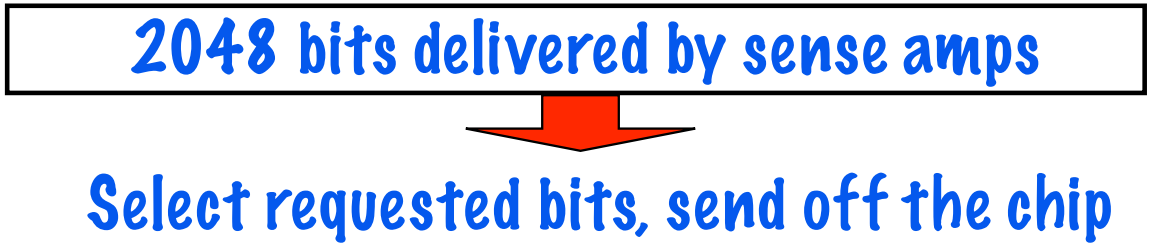

- **Interleaving:** Design the right interface to the 4 memory banks on the chip, so several row requests run in parallel.

Off-chip interface for the Micron part ...

A clocked bus protocol (133 MHz)

Note! This example is best-case! To access a new row, a slow ACTIVE command must run before the READ.

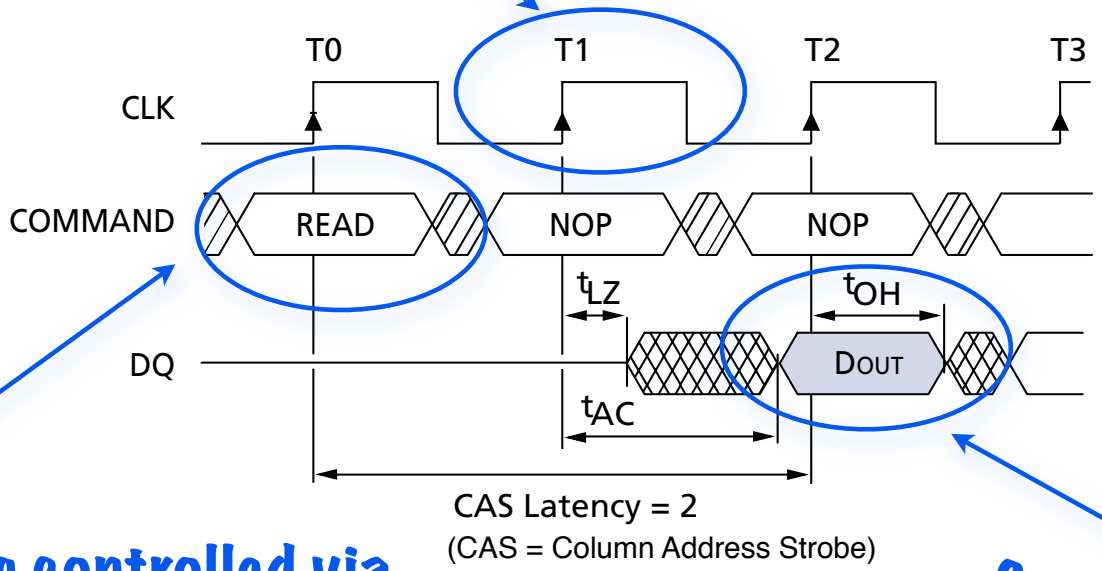

DRAM is controlled via commands (READ, WRITE, REFRESH, ...)

Synchronous data output.

From Micron 128 Mb SDRAM data sheet (on "resources" web page)

Opening a row before reading ...

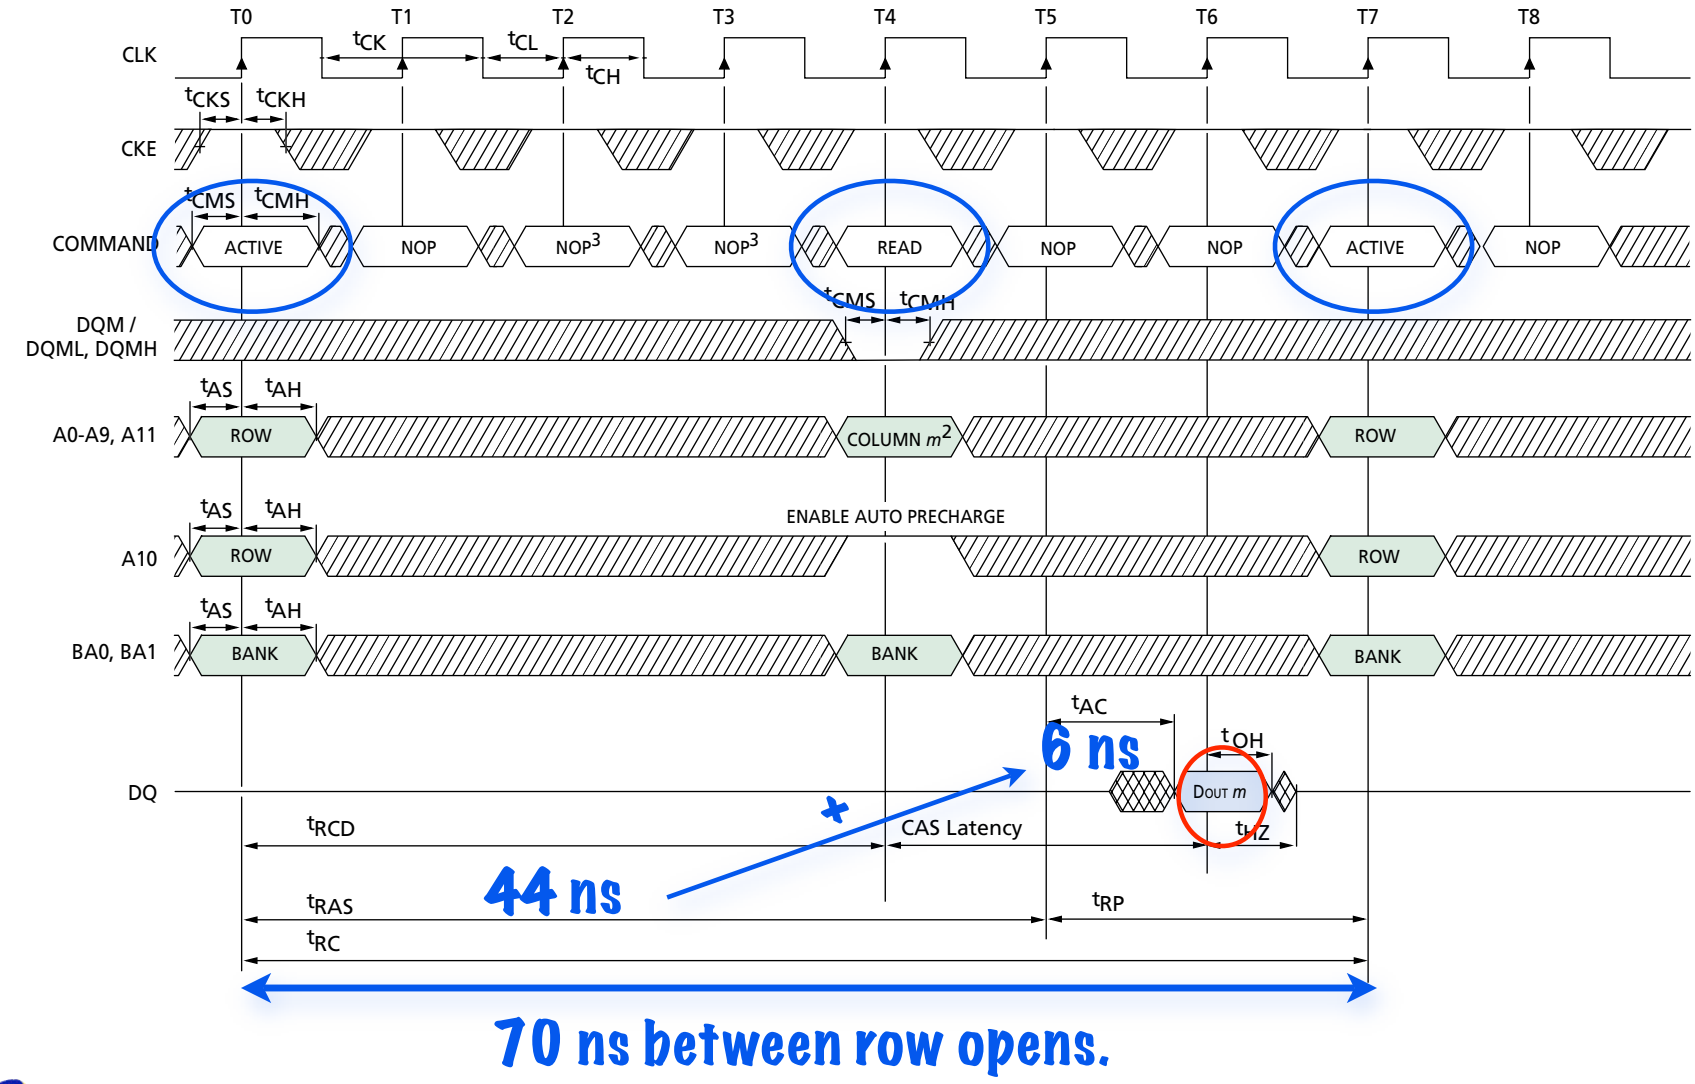

Note: Example uses Auto-Precharge, 100Mhz part

However, we can read columns quickly

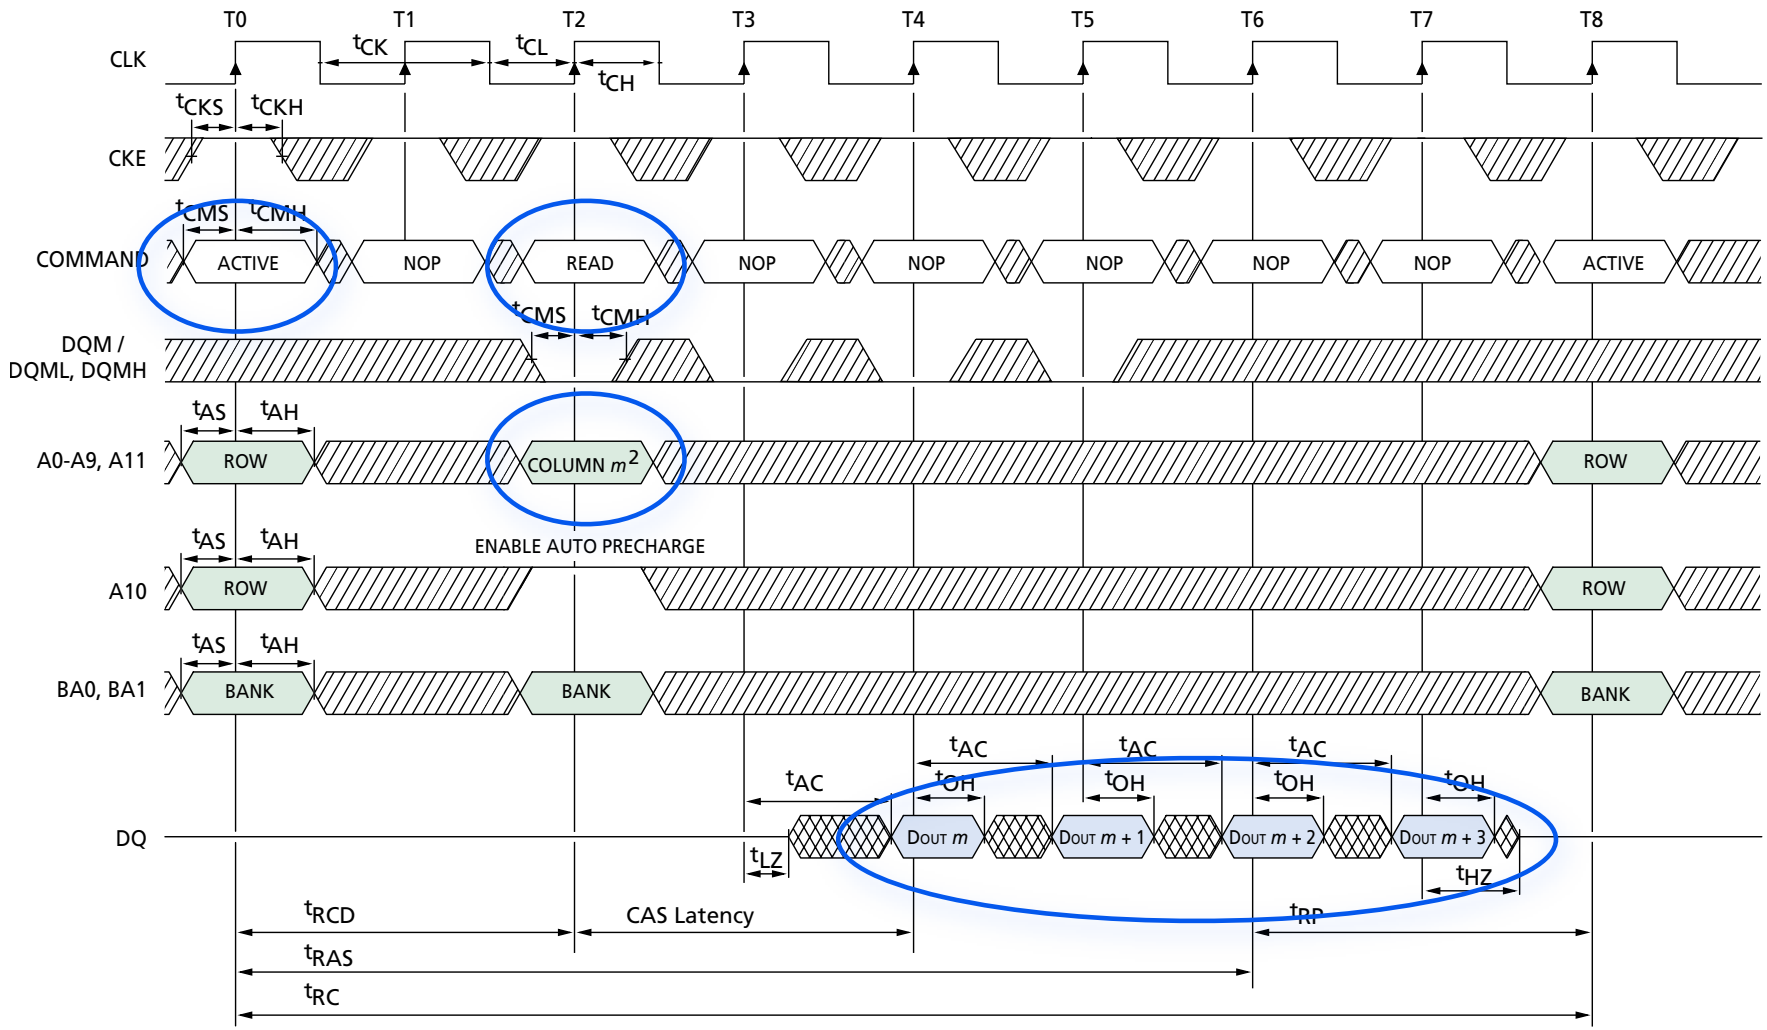

Why? Reading "delivered bits" is fast.

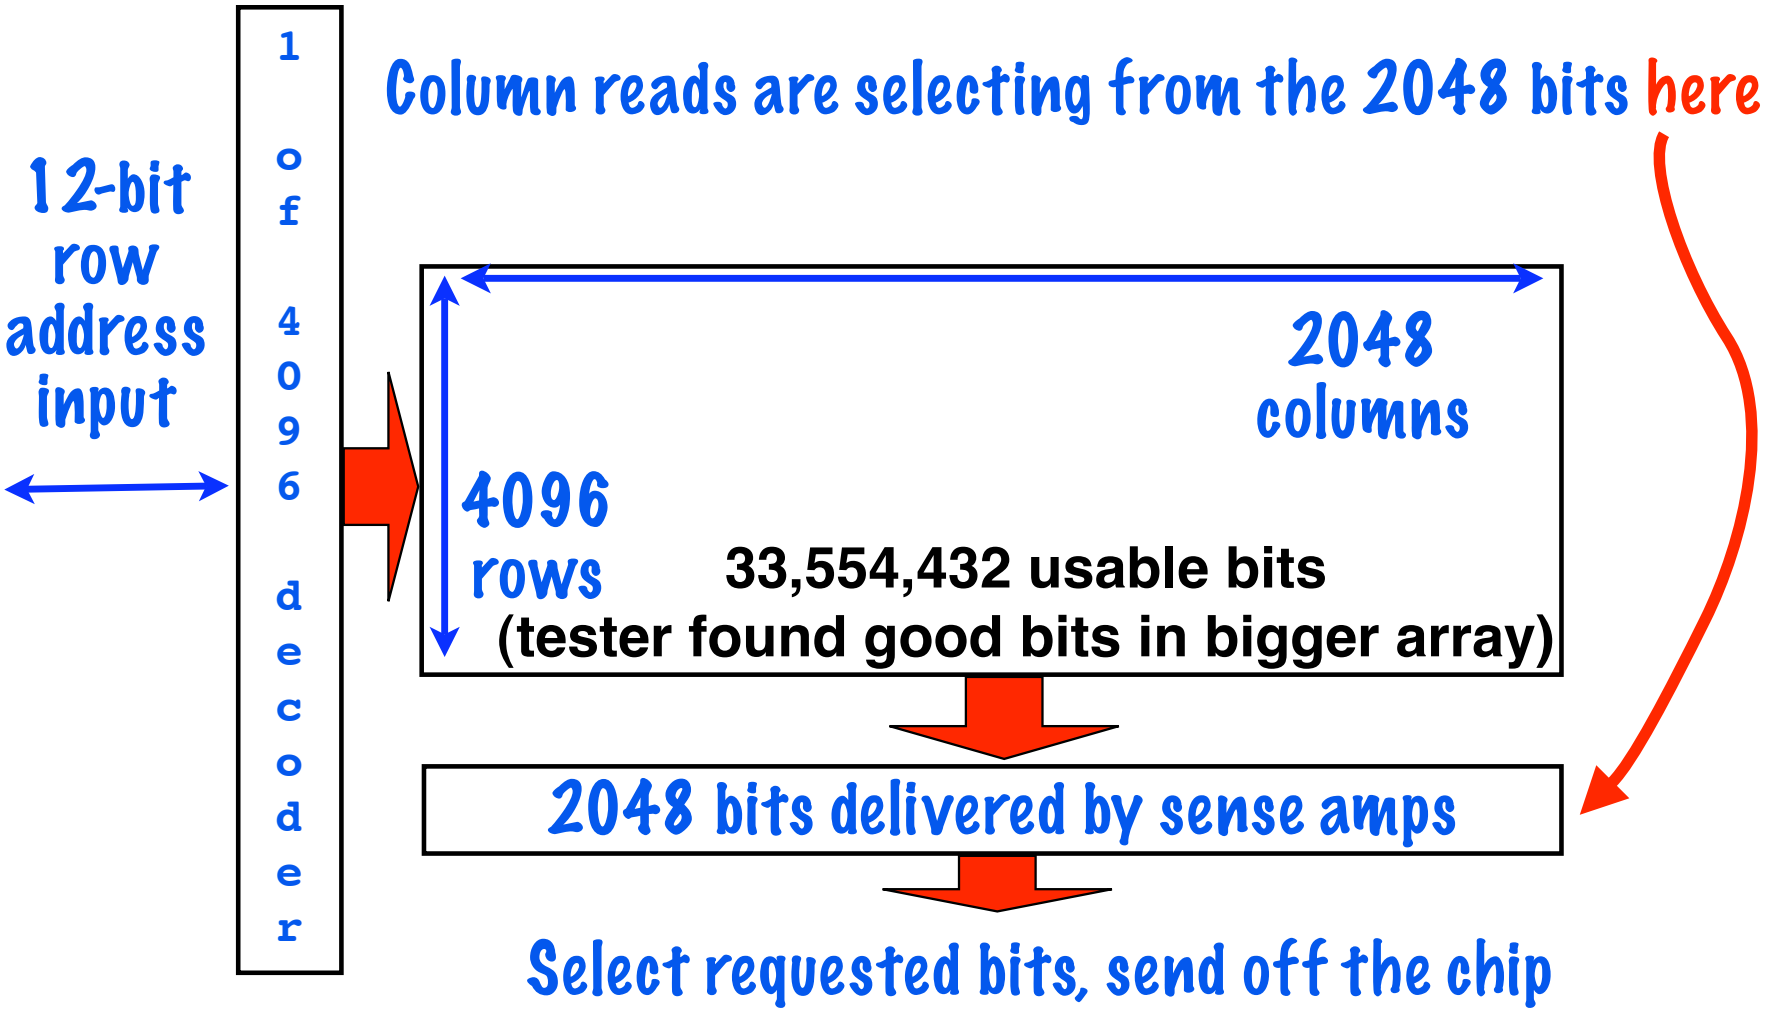

Precharge can also be done manually

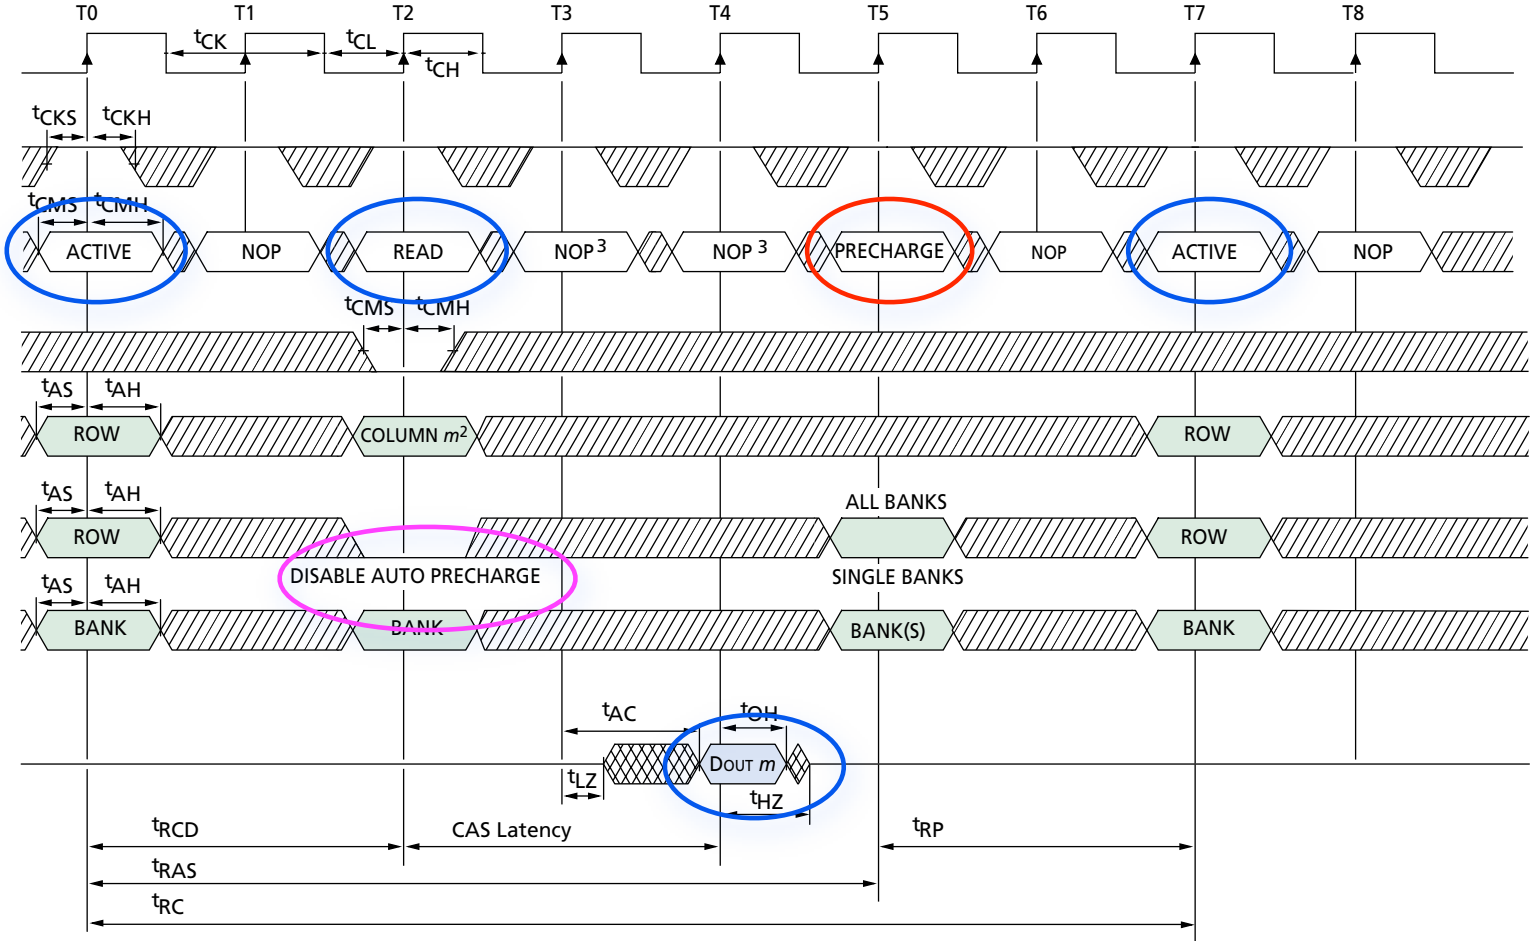

Between first ACTIVE and PRECHARGE, we can do as many READ and WRITE commands on a row as we wish ...

What happens during a PRECHARGE?

Recall: A DRAM read cycle ...

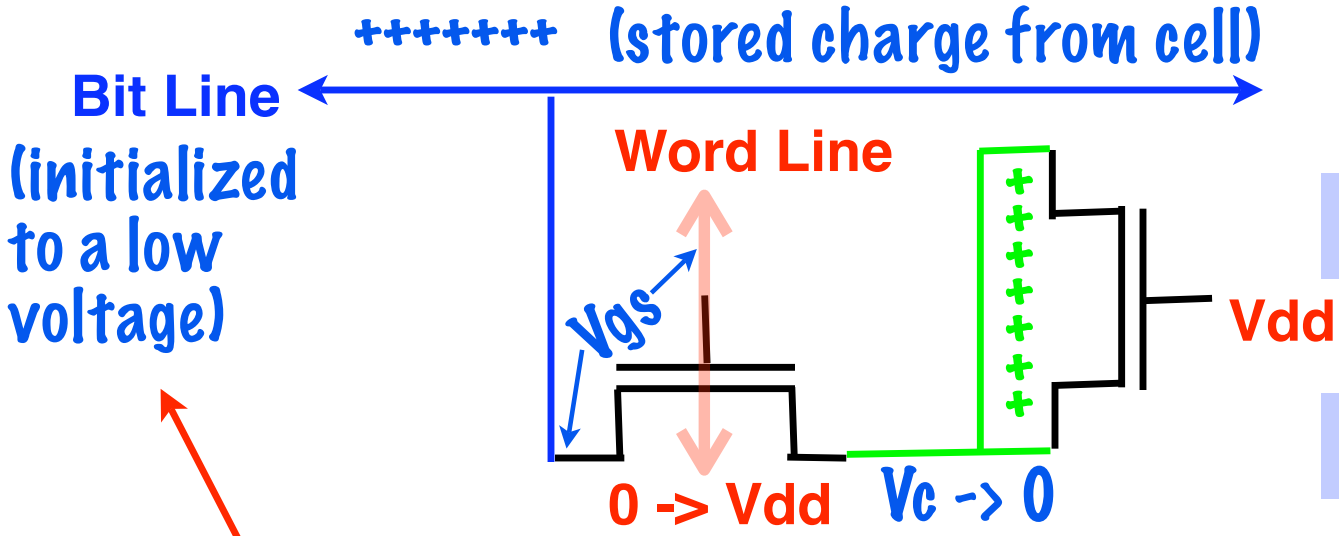

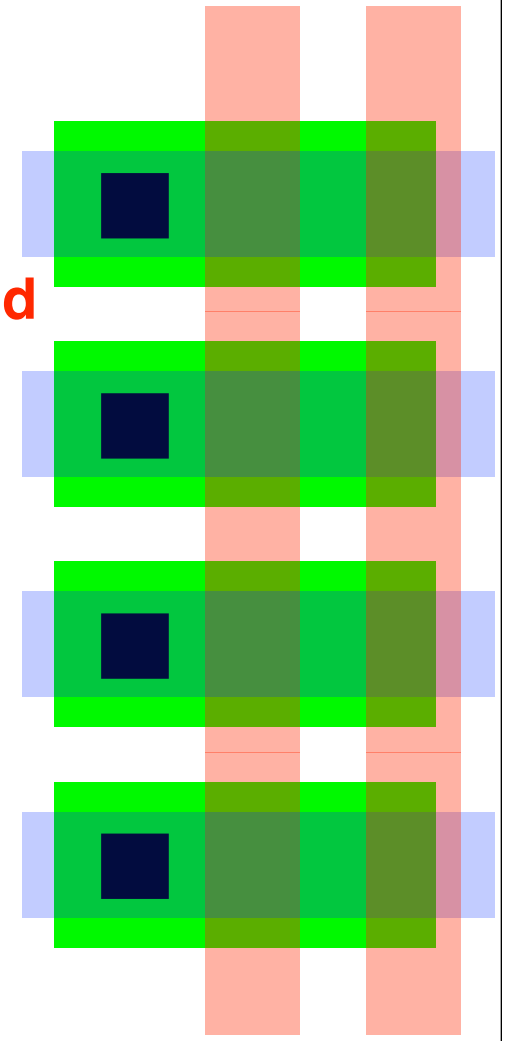

This "initializing" takes time to do ... and is done during Precharge.

Interleave: Access all 4 banks in parallel

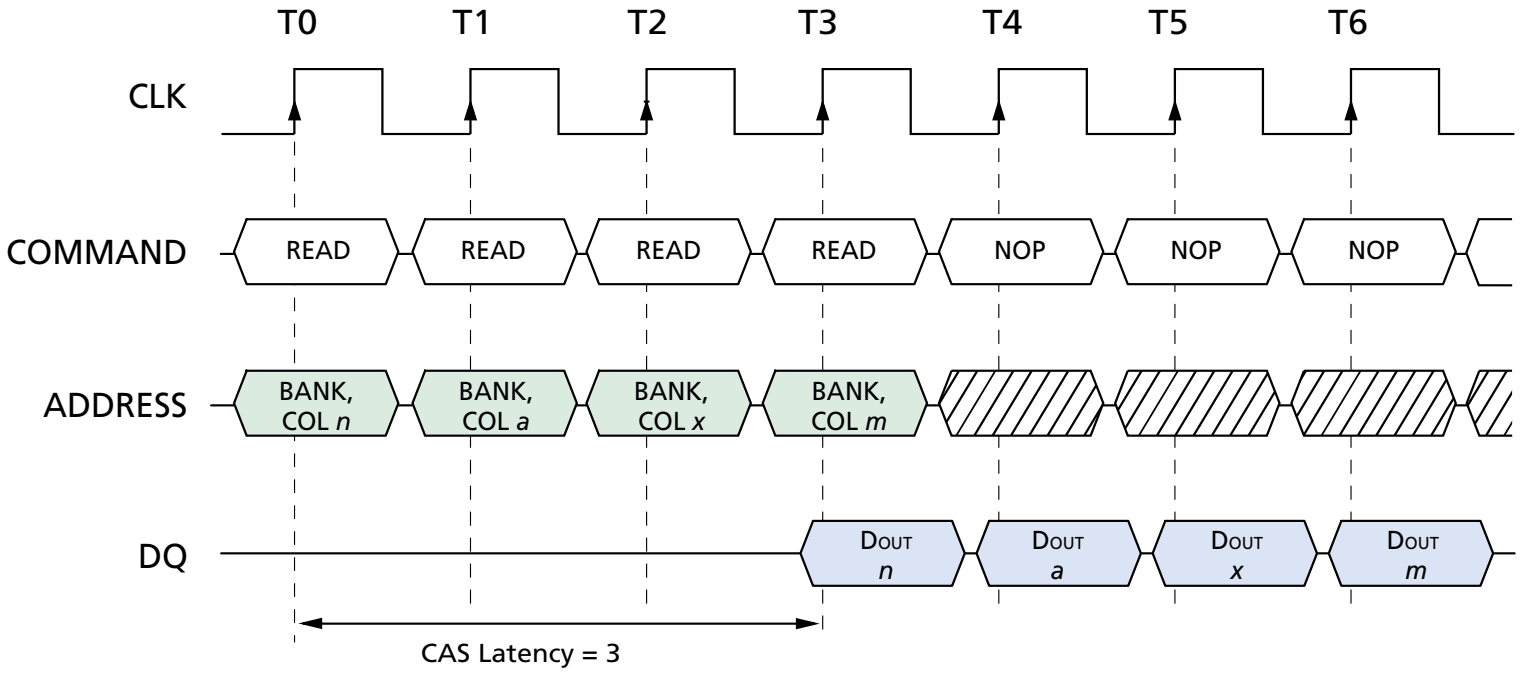

NOTE: Each READ command may be to any bank. DQM is LOW.

DON'T CARE

Figure 8
Random READ Accesses

Can also WRITE, PRECHARGE, ACTIVE banks concurrently.

From DRAM chip to DIMM module ...

Each RAM chip responsible for 8 lines of the 64 bit data bus (U5 holds the check bits).

Commands sent to all 9 chips, qualified by per-chip select lines.

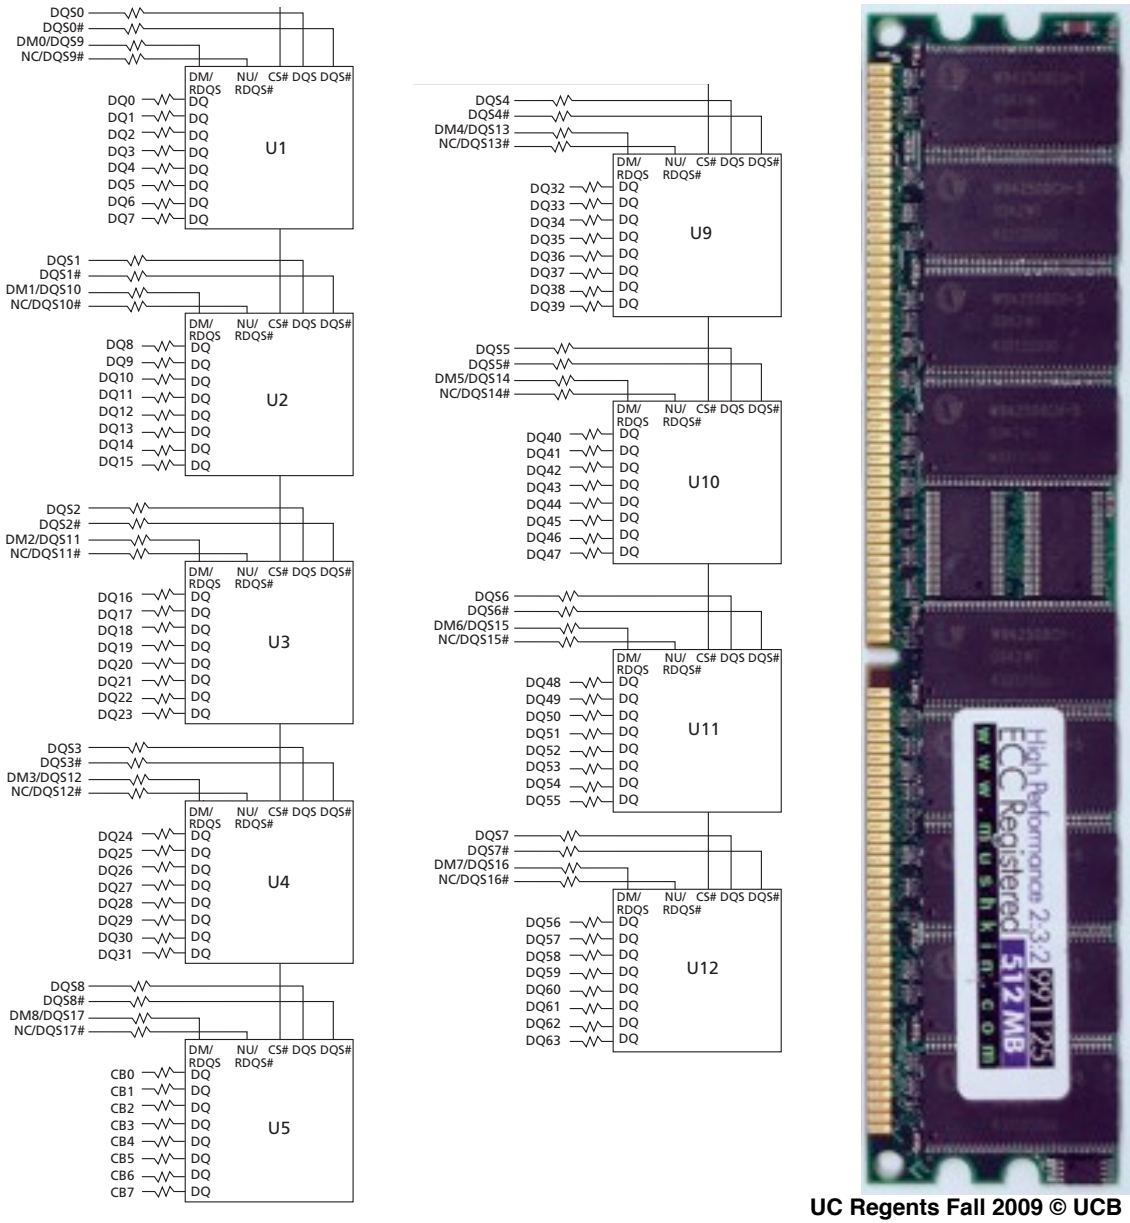

DRAM Controller Ideas

Reorder accesses for concurrency

(A) Without access scheduling (56 DRAM Cycles)

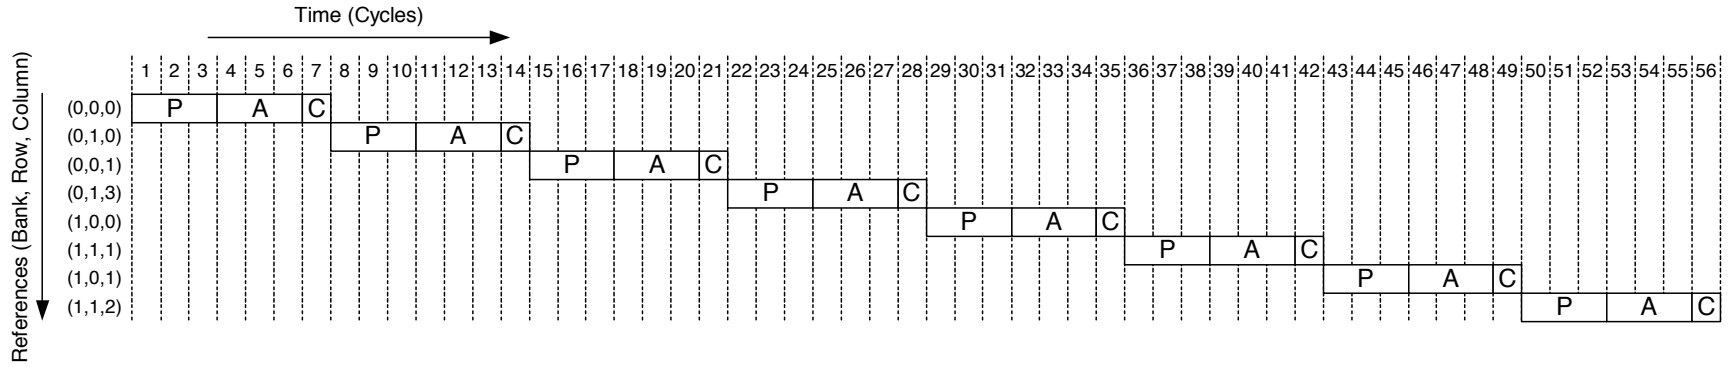

(B) With access scheduling (19 DRAM Cycles)

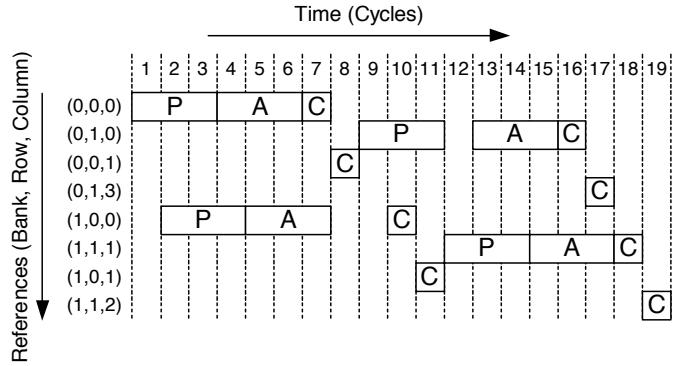

DRAM Operations:

P: bank precharge (3 cycle occupancy)
A: row activation (3 cycle occupancy)
C: column access (1 cycle occupancy)

From: **Memory Access Scheduling**

Scott Rixner[1], William J. Dally, Ujval J. Kapasi, Peter Mattson, and John D. Owens

Model SDRAM command semantics ...

(A) Simplified bank state diagram

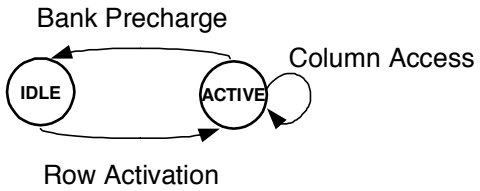

(B) Operation resource utilization

	Cycle	1	2	3	4
Precharge:	Bank	■	■	■	
	Address	■			
	Data				
Activate:	Bank	■	■	■	
	Address	■			
	Data				
Read:	Bank				
	Address	■			
	Data				■
Write:	Bank				
	Address	■			
	Data	■			

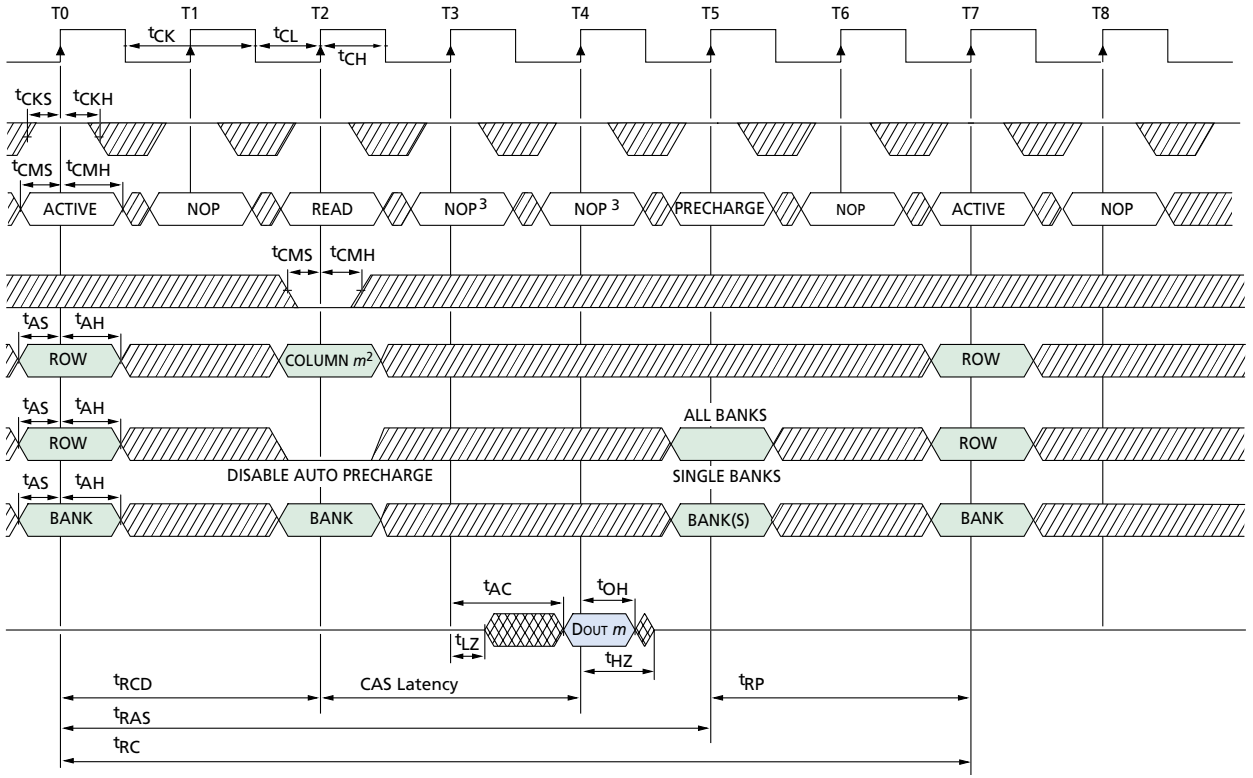

From: **Memory Access Scheduling**

Scott Rixner[1], William J. Dally, Ujval J. Kapasi, Peter Mattson, and John D. Owens

Set scheduling algorithms in gates ...

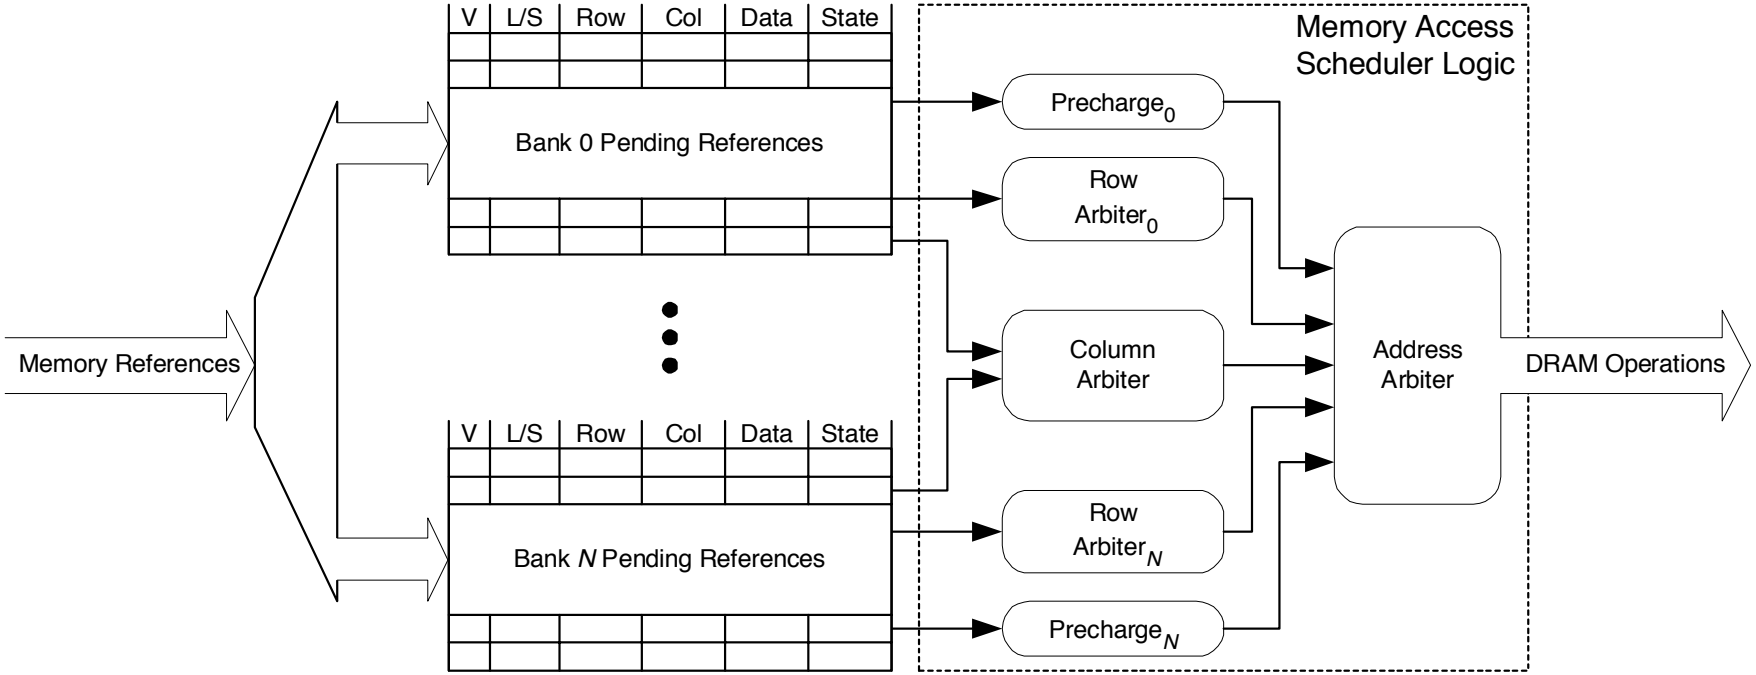

Figure 4. Memory access scheduler architecture.

From: **Memory Access Scheduling**

Scott Rixner[1], William J. Dally, Ujval J. Kapasi, Peter Mattson, and John D. Owens

Scheduling algorithms ...

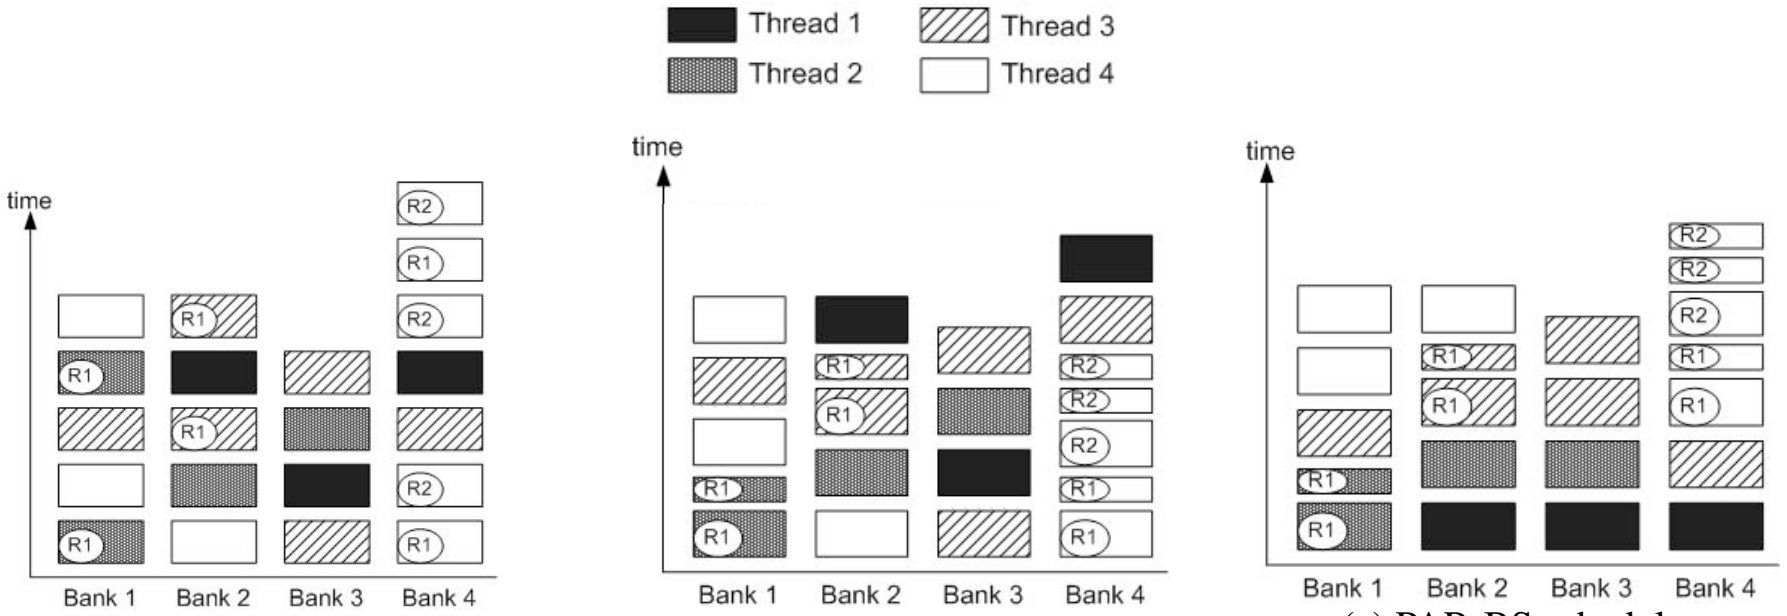

a) Arrival order (and FCFS schedule) (b) FR-FCFS schedule (c) PAR-BS schedule

FCFS schedule batch-completion (stall) times					FR-FCFS schedule batch-completion (stall) times					PAR-BS schedule batch-completion (stall) times				
Thread 1	Thread 2	Thread 3	Thread 4	AVG	Thread 1	Thread 2	Thread 3	Thread 4	AVG	Thread 1	Thread 2	Thread 3	Thread 4	AVG
4	4	5	7	5	5.5	3	4.5	4.5	4.375	1	2	4	5.5	3.125

From:

Parallelism-Aware Batch Scheduling:
Enhancing both Performance and Fairness of Shared DRAM Systems

Onur Mutlu Thomas Moscibroda
Microsoft Research
{onur,moscitho}@microsoft.com

For more concurrency: rethink DIMMs?

Traditional DIMMs minimize pins by sharing wires.

As SDRAM chip interfaces get faster, should we optimize concurrency instead?

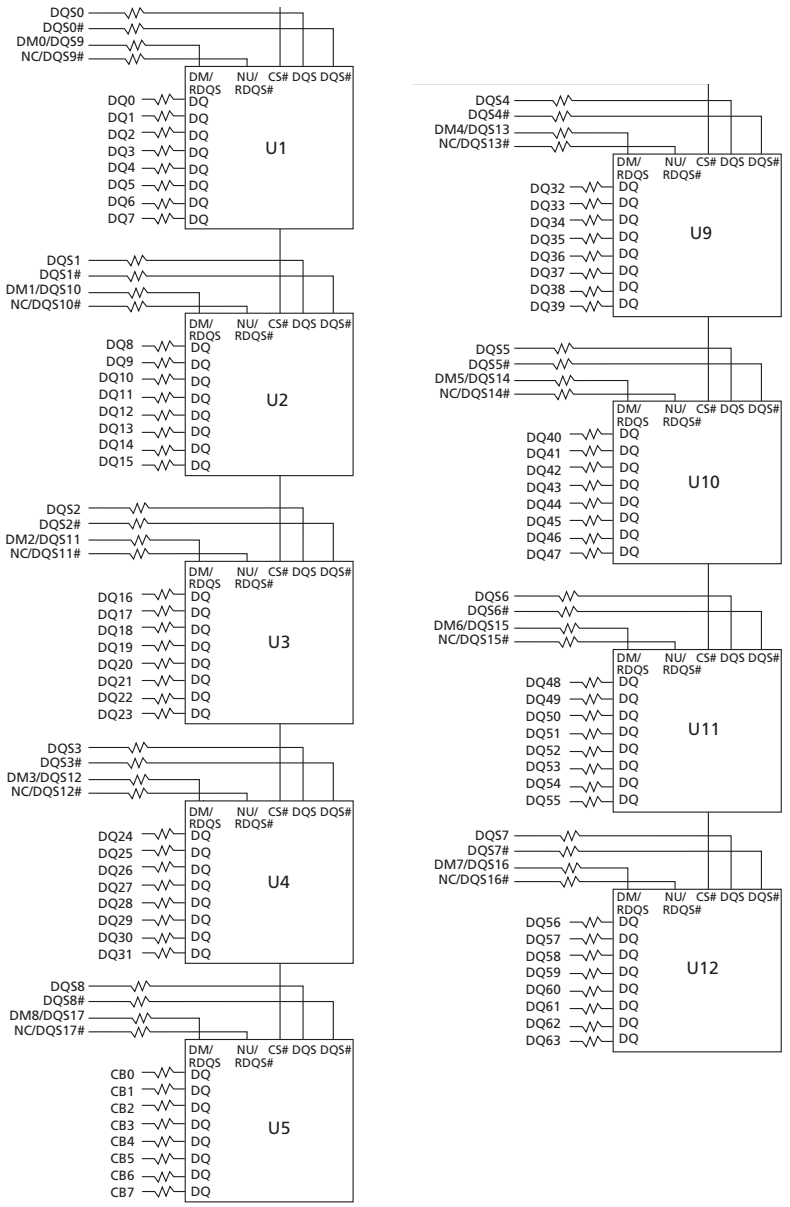

Web server data ...

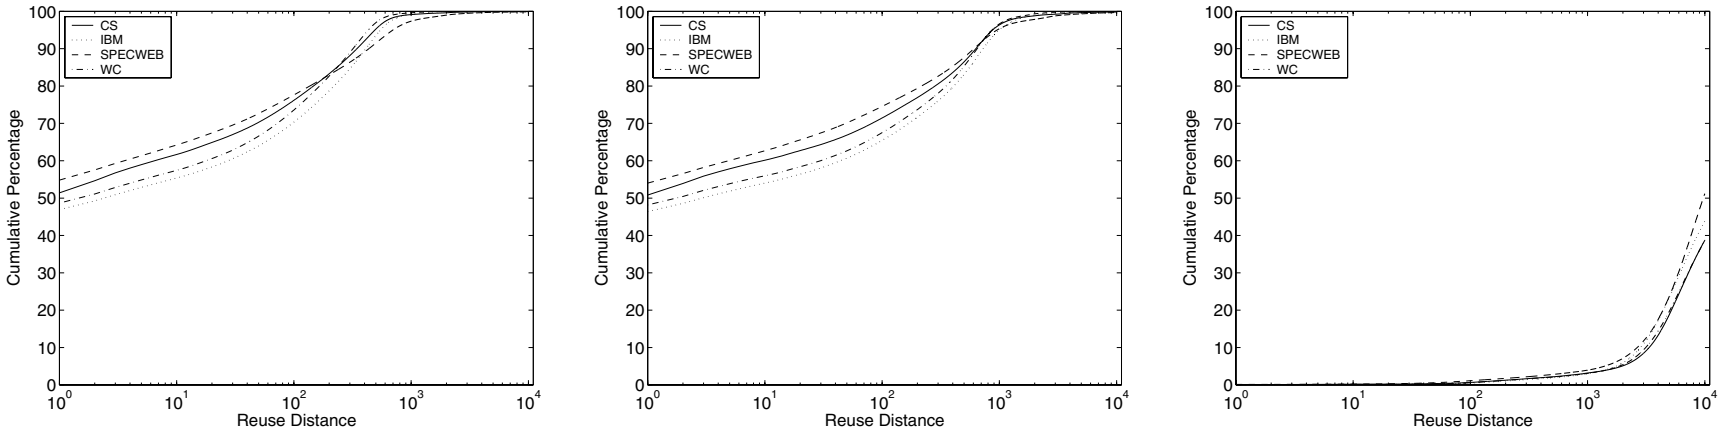

(A) 32 KB Row Reuse (B) 8 KB Segment Reuse (C) 256 B Block Reuse

Figure 1. Cumulative distribution of SDRAM access reuse distances.

From: **Memory Controller Optimizations for Web Servers**

Scott Rixner
Rice University

... yields DRAM chip architecture ideas

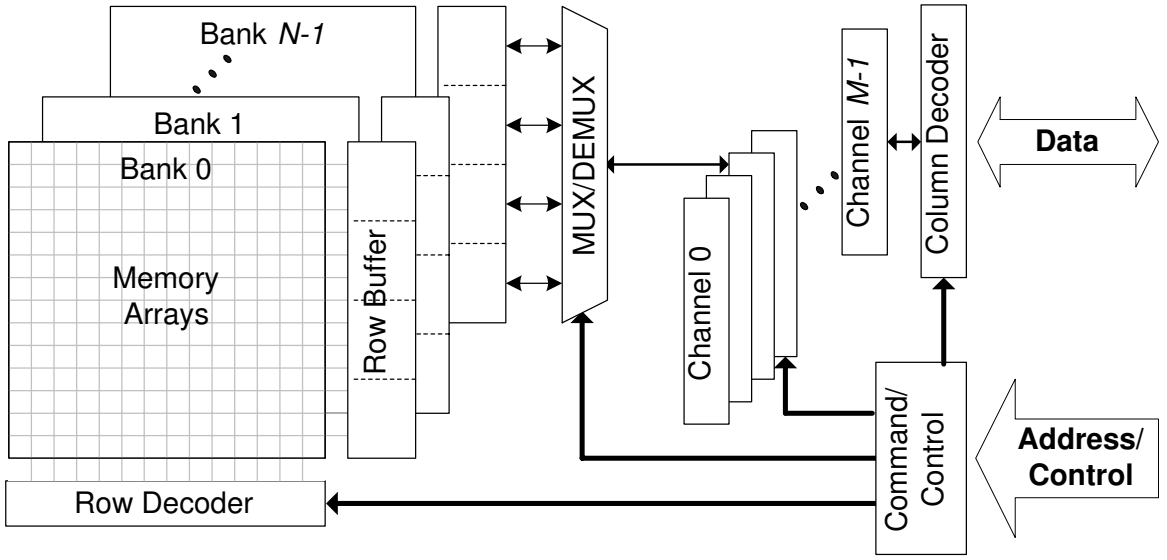

Figure 2. Virtual Channel SDRAM organization (N Banks and M Channels).

From: **Memory Controller Optimizations for Web Servers**

Scott Rixner
Rice University

Next Week: Practical design techniques

- Tuesday: Micro-architecture
- Thursday: Signalling, etc.