CS 150
Digital Design

Lecture 12 – DRAM

2012-2-23

Professor John Wawrzynek

today's lecture by John Lazzaro

TAs: Shaoyi Cheng, Daiwei Li, James Parker

www-inst.eecs.berkeley.edu/~cs150/

Today's Lecture: DRAM

- **DRAM, Xilinx, and You**
- **DRAM: Bottom-up**
- **DRAM: Top-down**

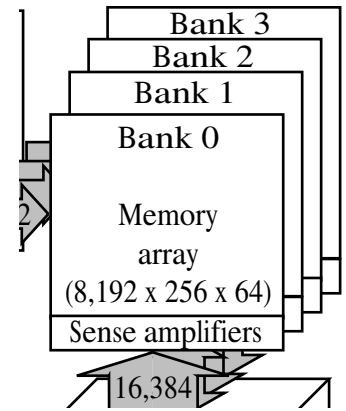

DDR2 SO-DIMM on ML505 Board

DDR2: Double-Data Rate, 2nd generation

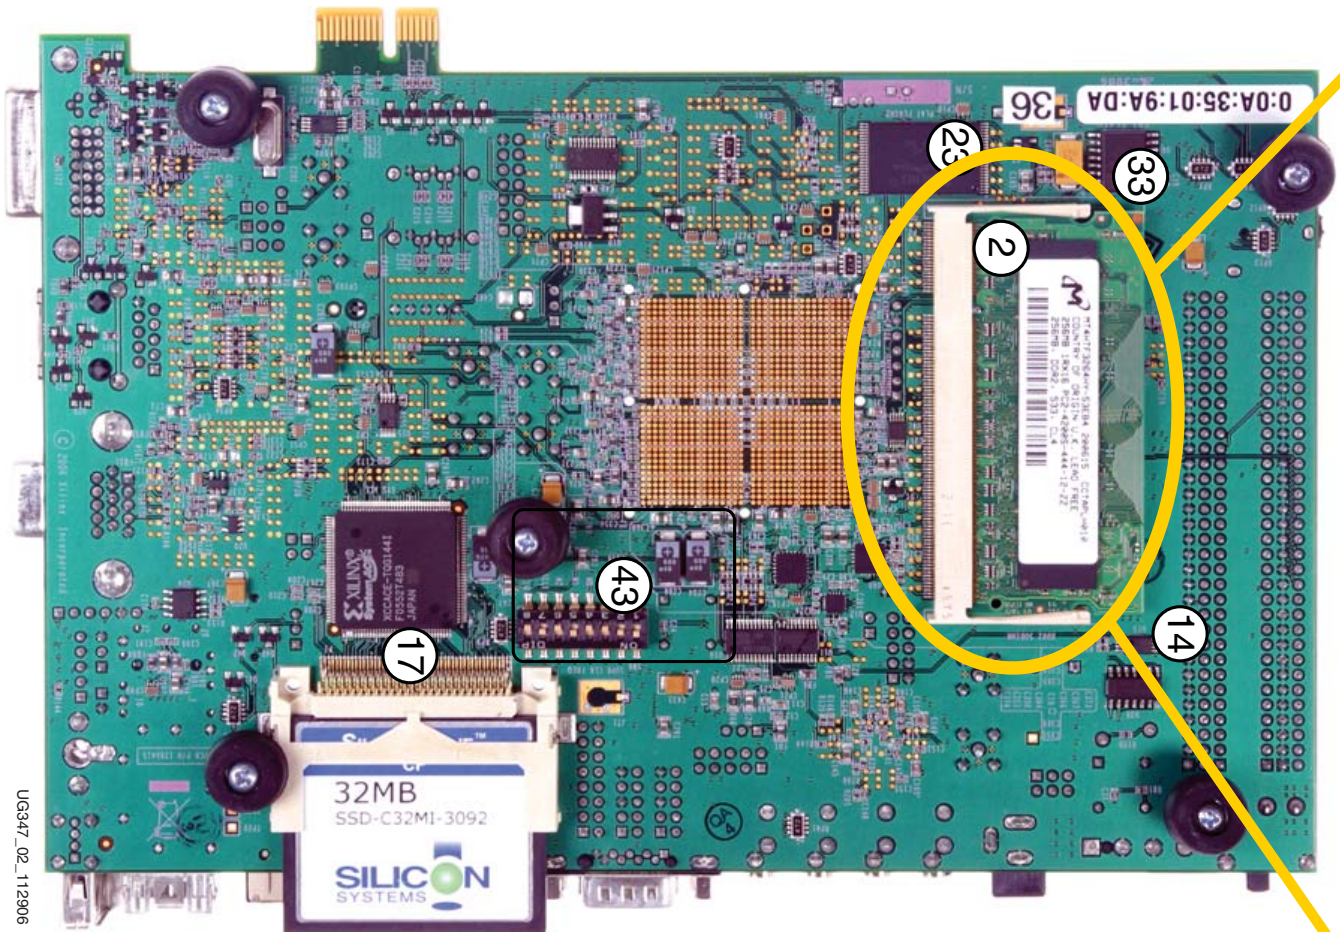

UG347_02_112906

SO-DIMM: Small-Outline, Dual Inline Memory Module

DDR2 SO-DIMM Module

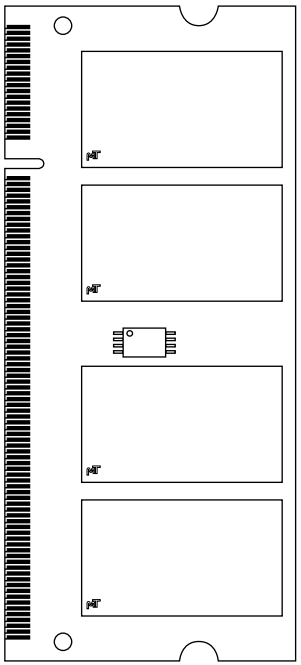

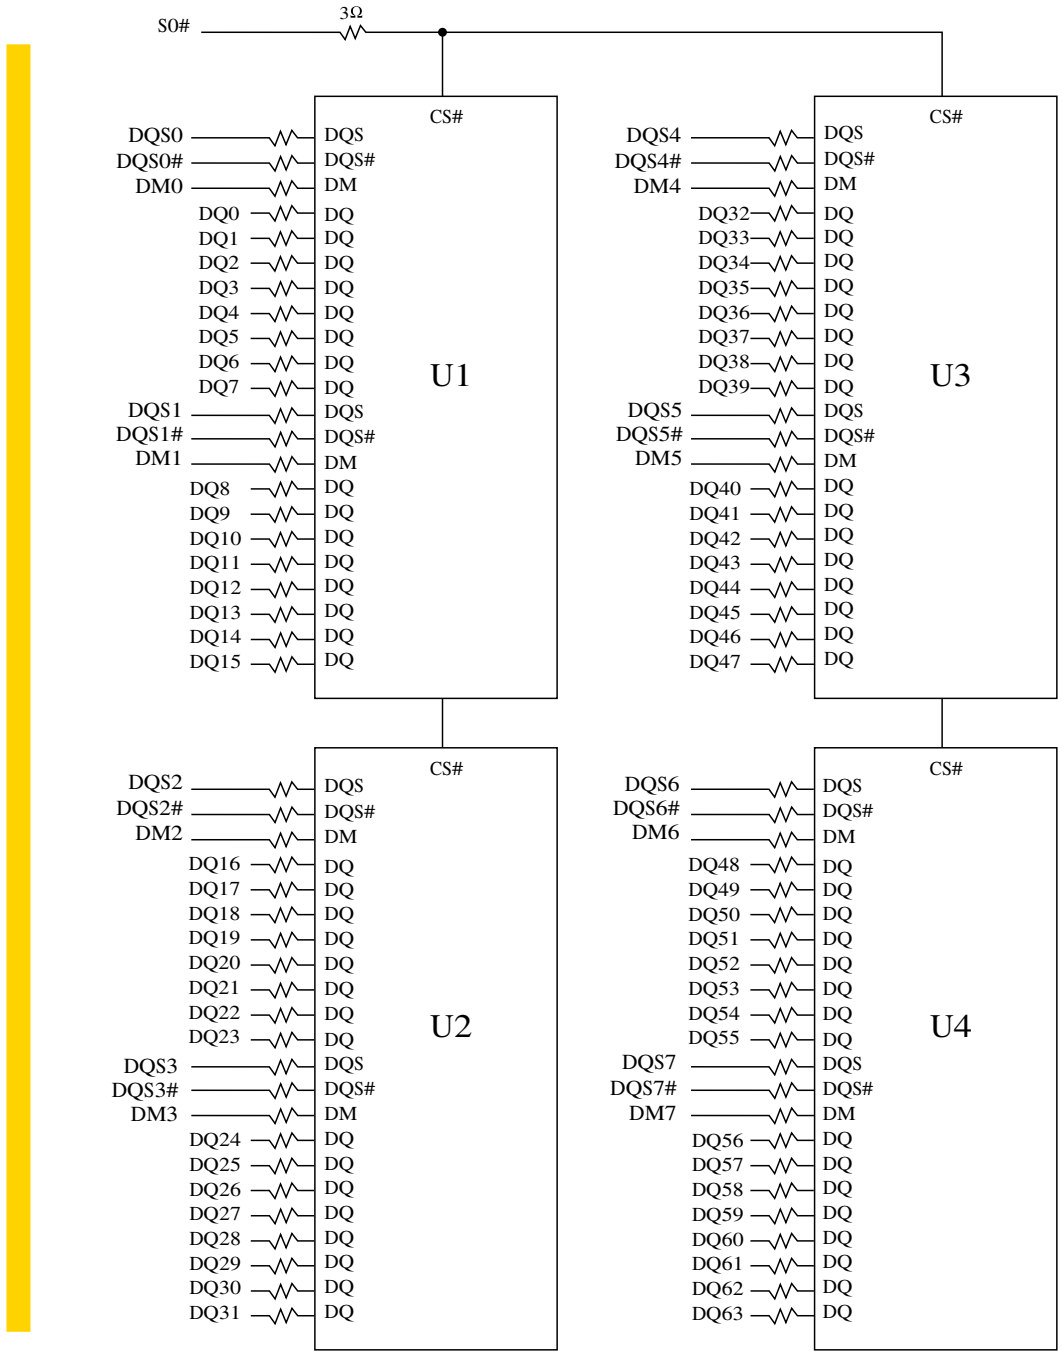

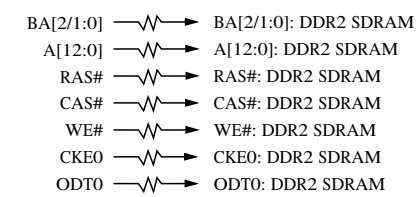

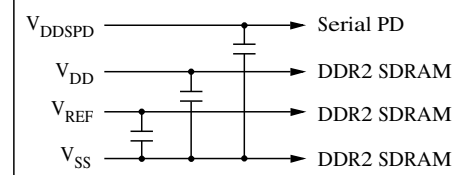

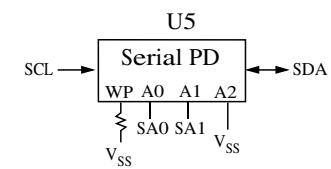

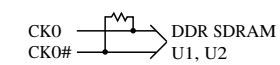

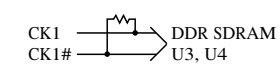

DRAM chips are wired in parallel and run in lockstep.

Project controller: Xilinx-supplied IP

Your project's Verilog code sees a FIFO R/W interface.

Xilinx IP translates FIFO requests to DRAM commands.

DDR2 SO-DIMM

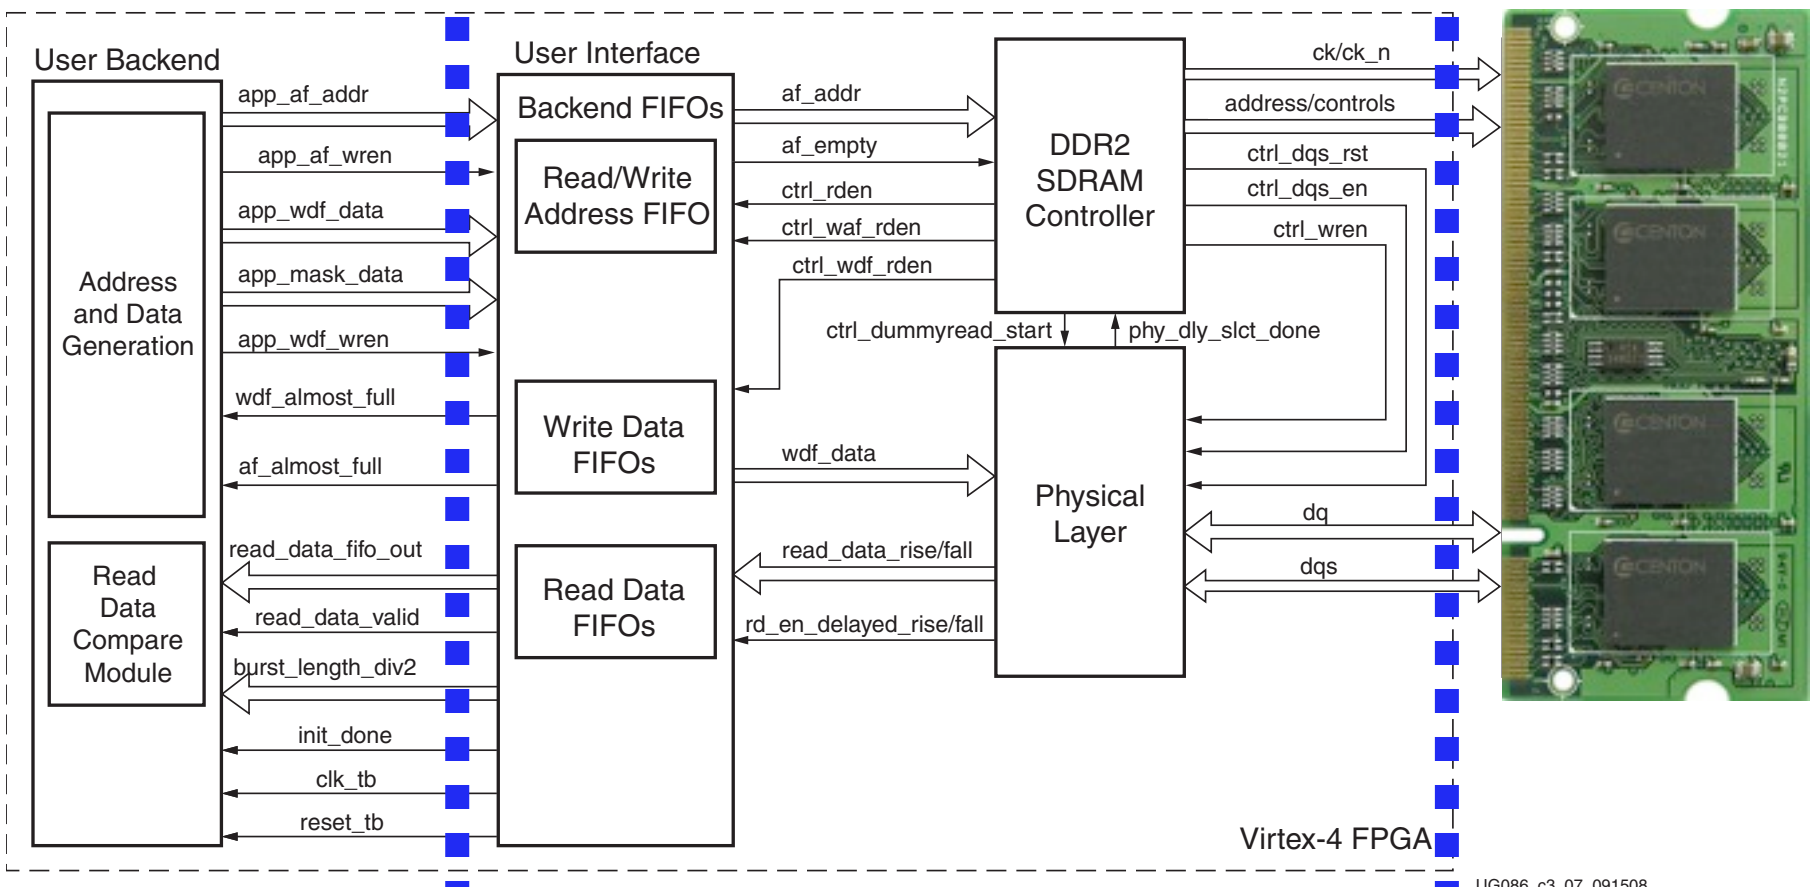

Today's Lecture: DRAM

- DRAM, Xilinx, and You
- **DRAM: Bottom-up**
- DRAM: Top-down

To understand the DRAM controller, you need to understand how a DRAM chip works. Otherwise, it just seems like magic.

Capacitance

Recall: Building a capacitor

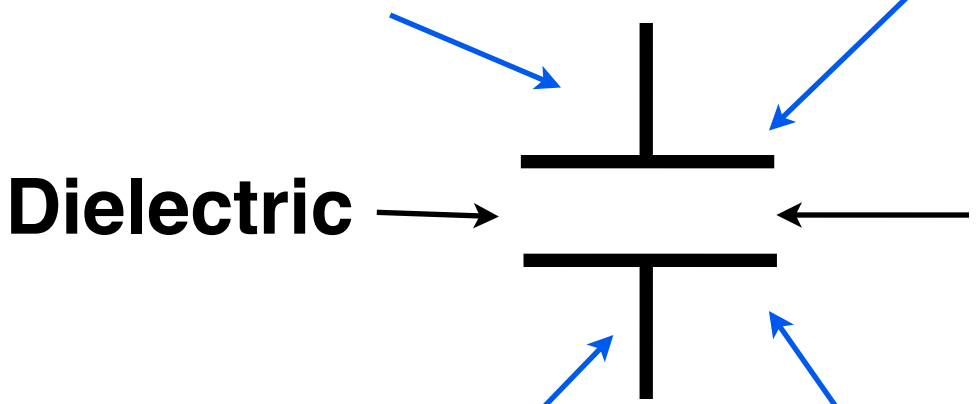

Top Plate

Conducts electricity well. (metal, doped polysilicon)

Dielectric

An insulator. Does not conducts electricity at all. (air, glass (silicon dioxide))

Bottom Plate

Conducts electricity well (metal, doped polysilicon)

Recall: Capacitors in action

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

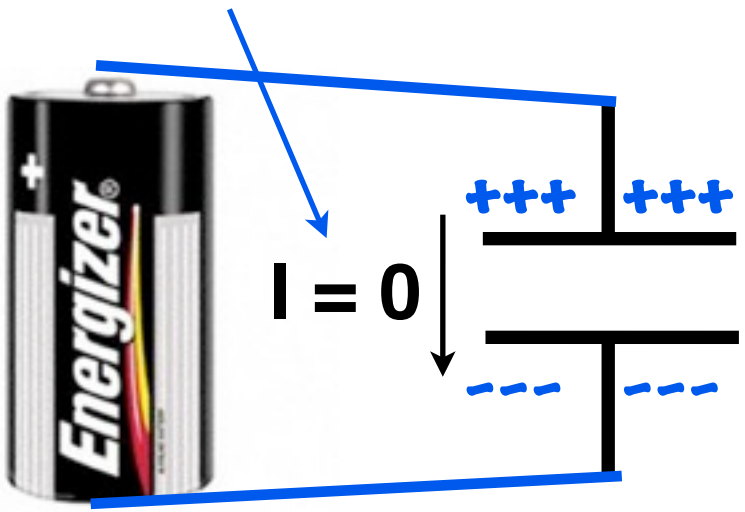

After circuit "settles" ...

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

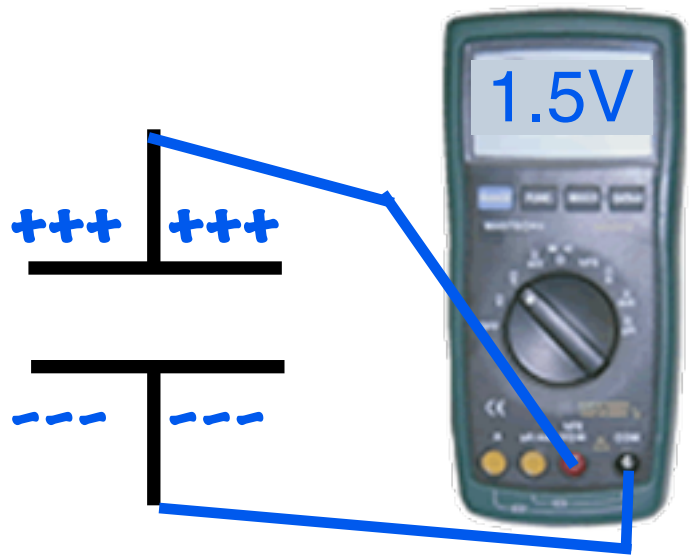

Storing computational state as charge

State is coded as the amount of energy stored by a device.

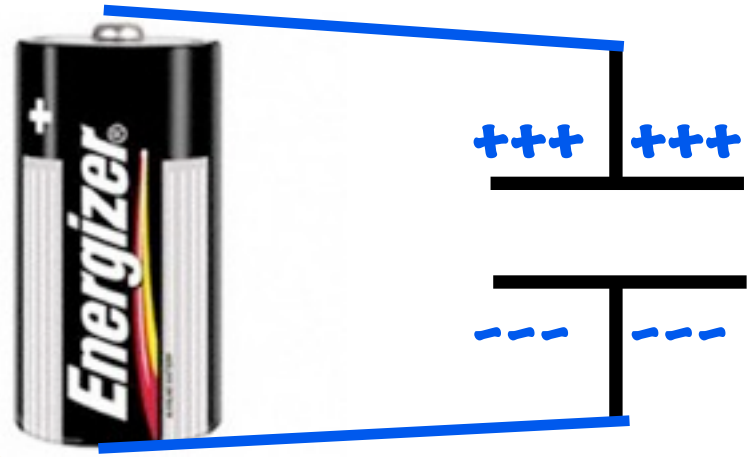

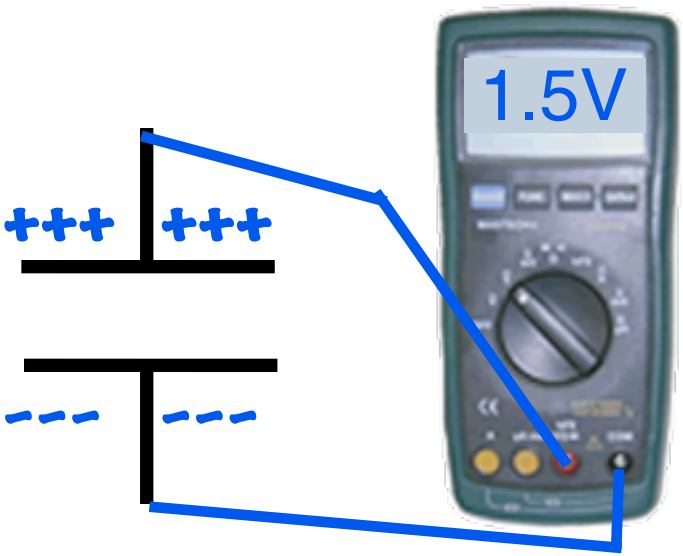

State is read by sensing the amount of energy

Problems: noise changes Q (up or down), parasitics leak or source Q. Fortunately, Q cannot change instantaneously, but that only gets us in the ballpark.

MOS Transistors

Two diodes and a capacitor in an interesting arrangement. So, we begin with a diode review …

Diodes in action ...

Resistor

Light emitting diode (LED)

Light on?

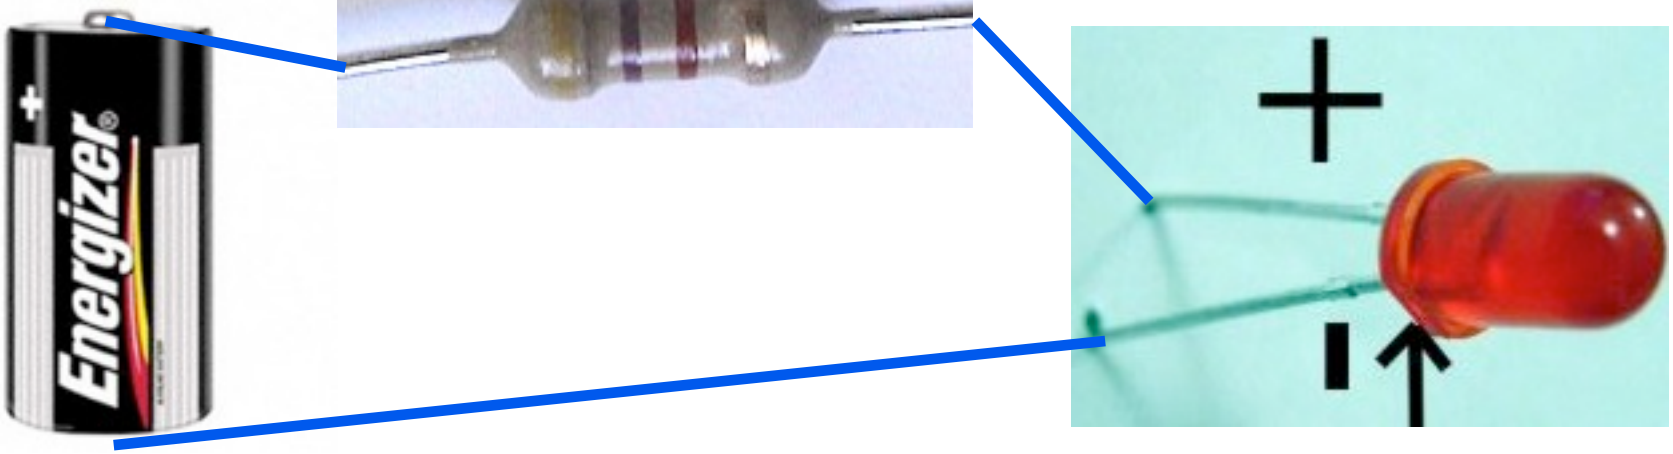

Yes!

Light on?

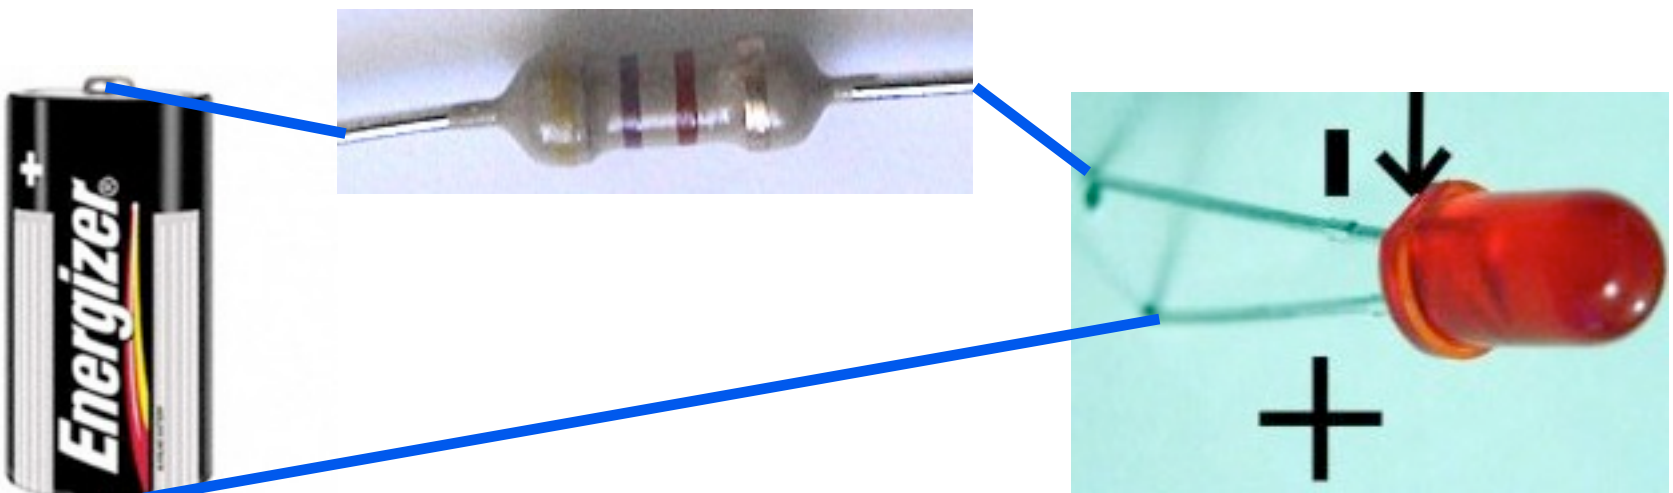

No!

Diodes: Current vs Voltage

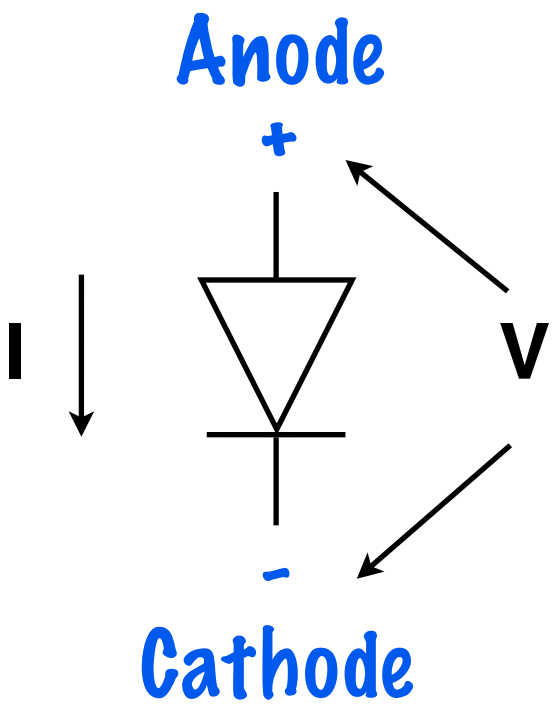

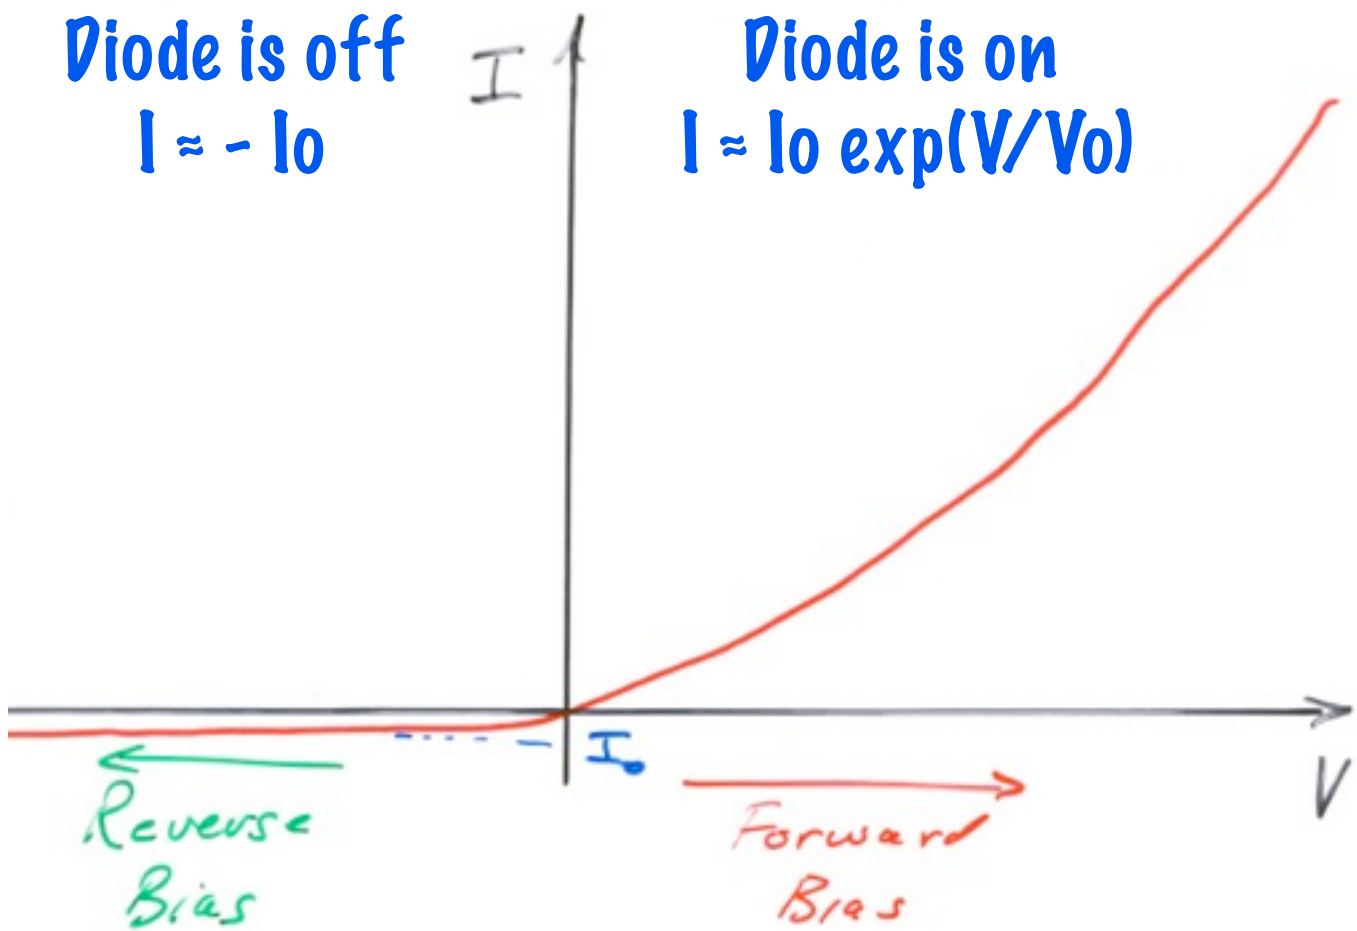

$$I = Io\,[\exp(V/Vo) - 1]$$

Io range: 1fA to 1nA

Vo range: 25mV to 60 mV

How to make a silicon diode ...

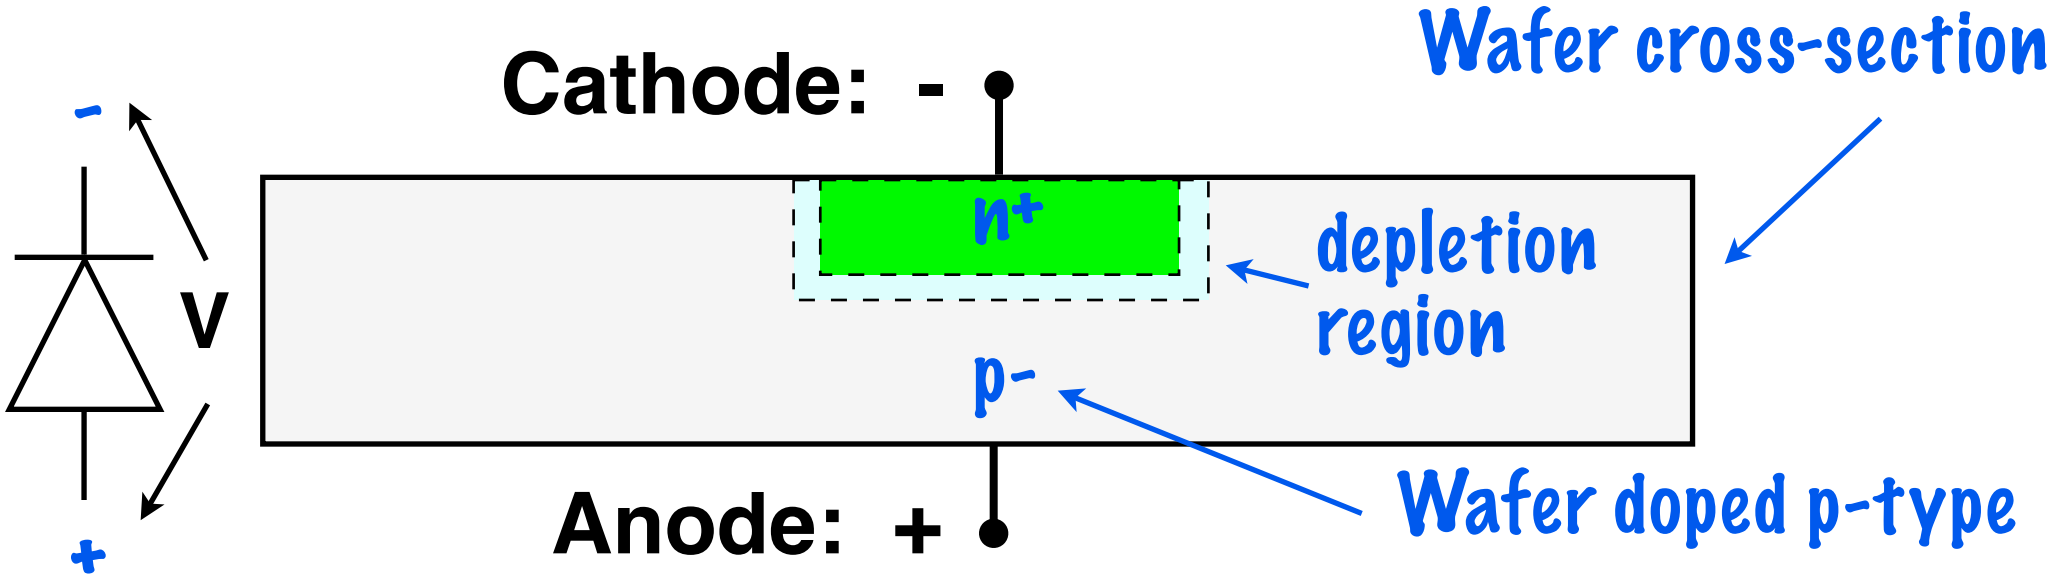

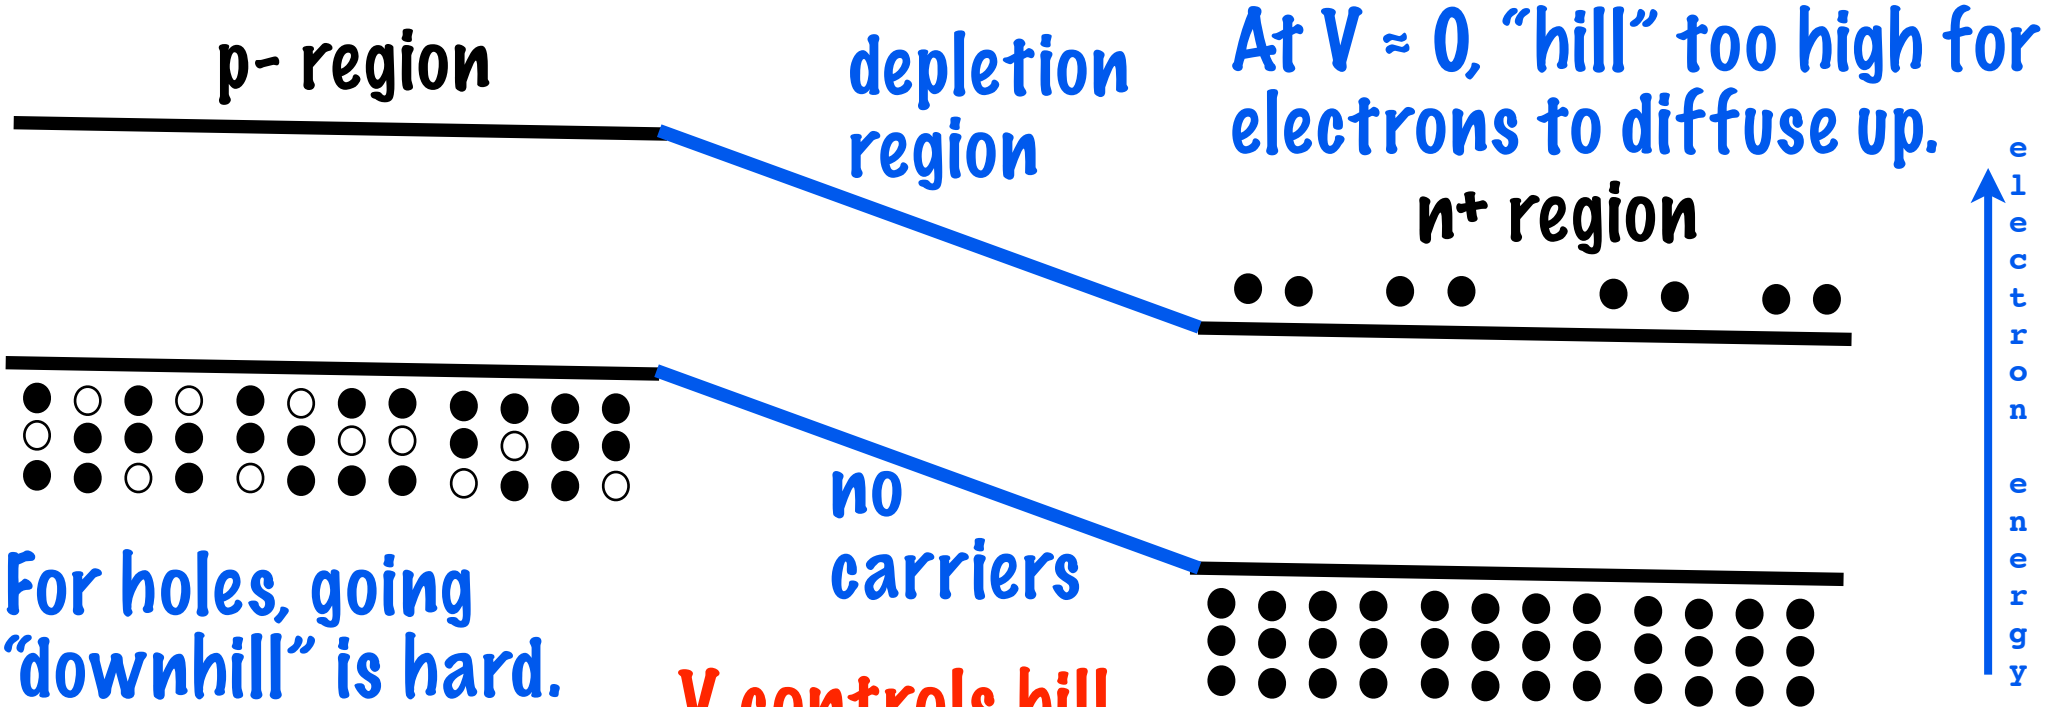

Note: IC Diodes are biased "off"!

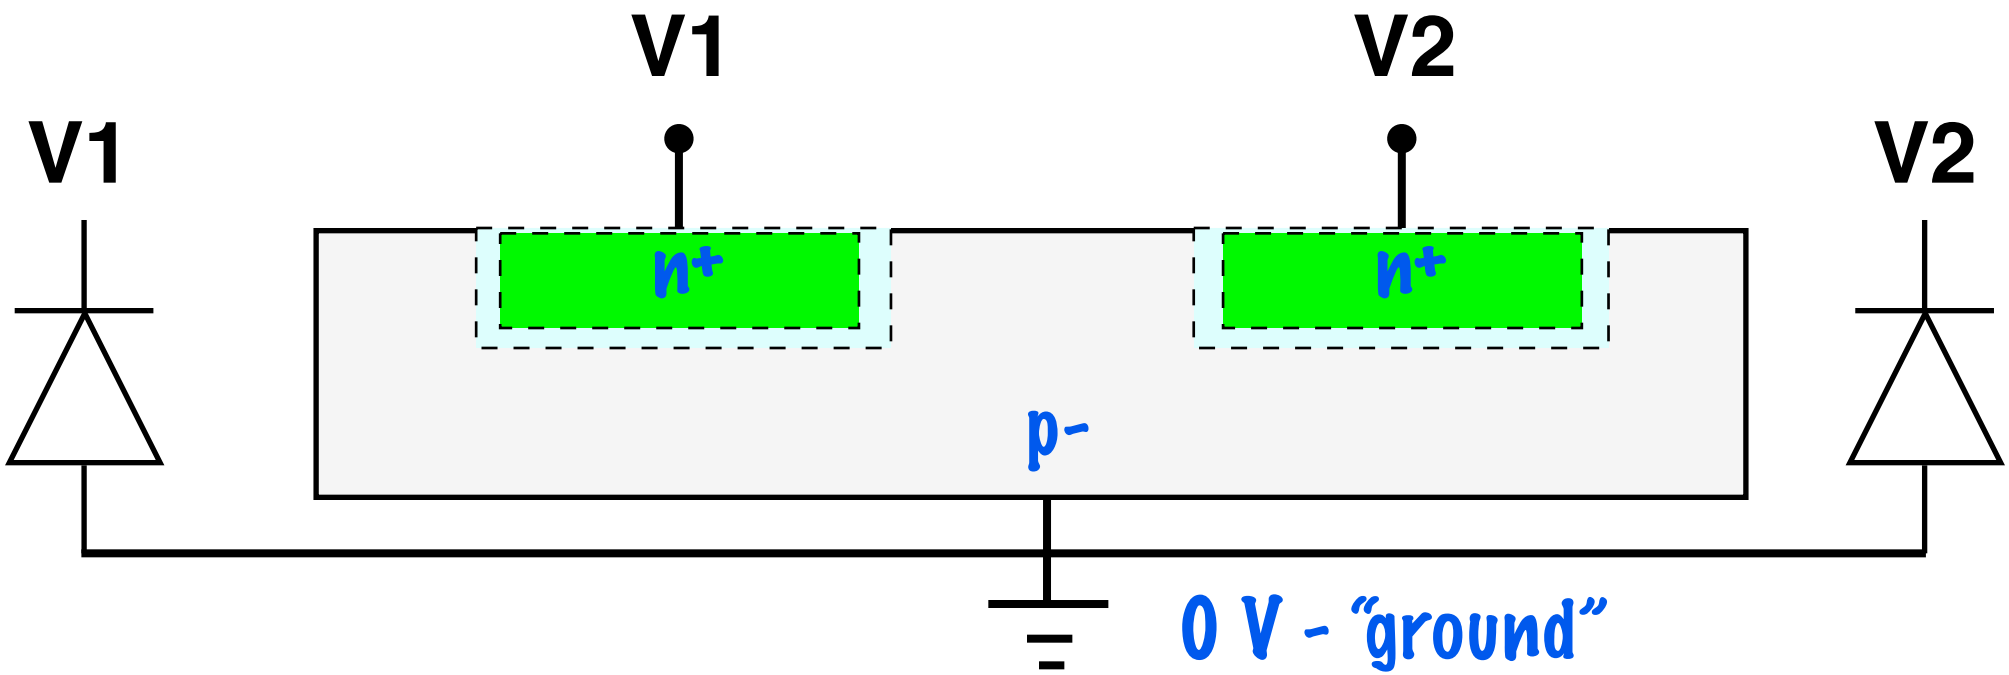

V1, V2 > 0V. Diodes "off", only current is Io "leakage".

$$I = Io\,[\exp(V/Vo) - 1]$$

Anodes of all diodes on wafer connected to ground.

MOS Transistors

Two diodes and a capacitor in an interesting arrangement ...

What we want: the perfect switch.

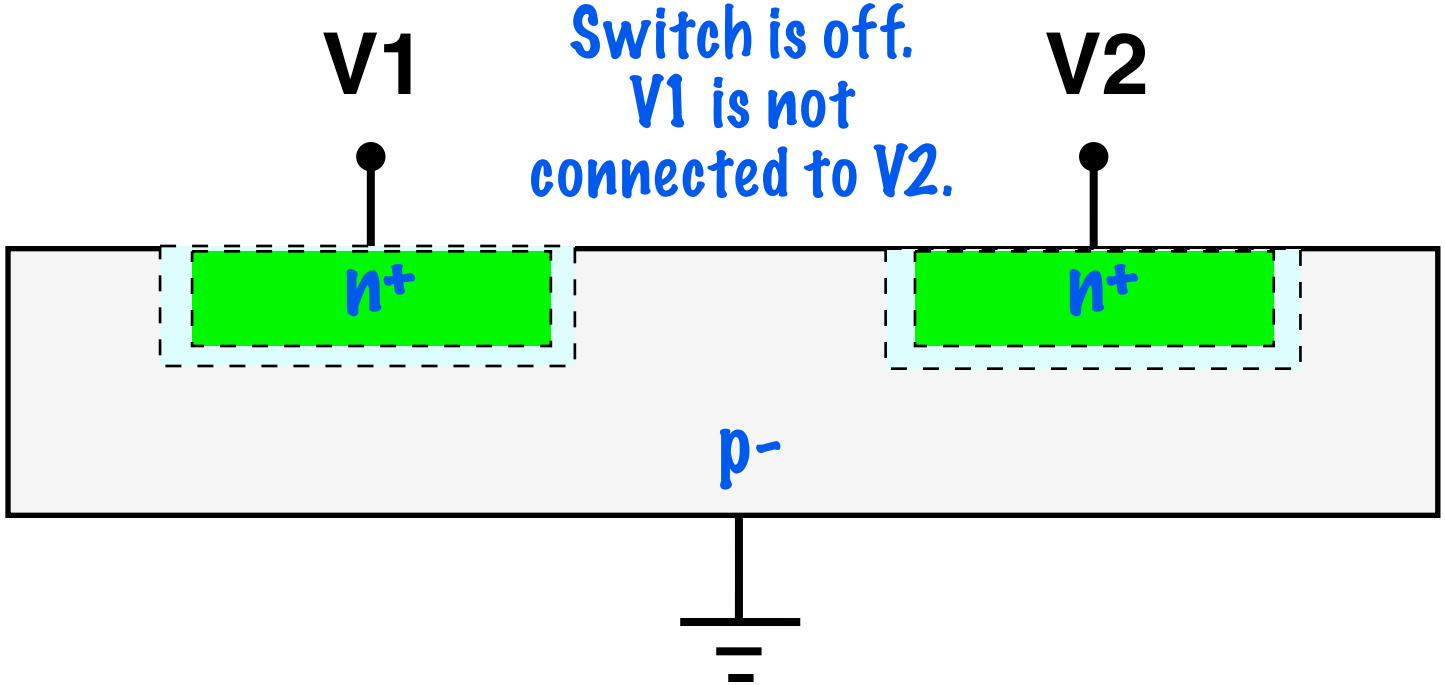

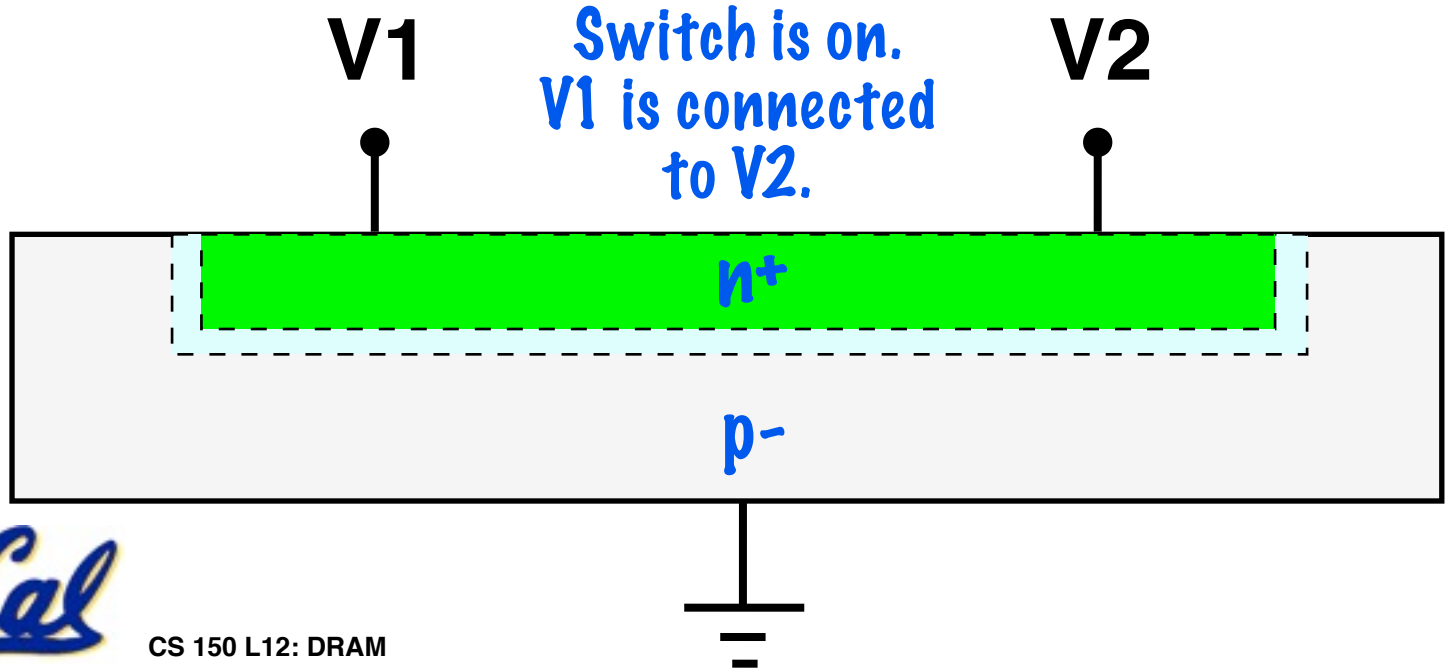

We want to turn a p-type region into an n-type region under voltage control.

We need electrons to fill valence holes and add conduction band electrons

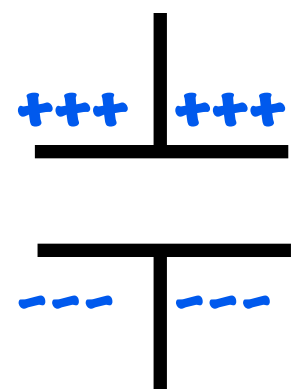

An n-channel MOS transistor (nFET)

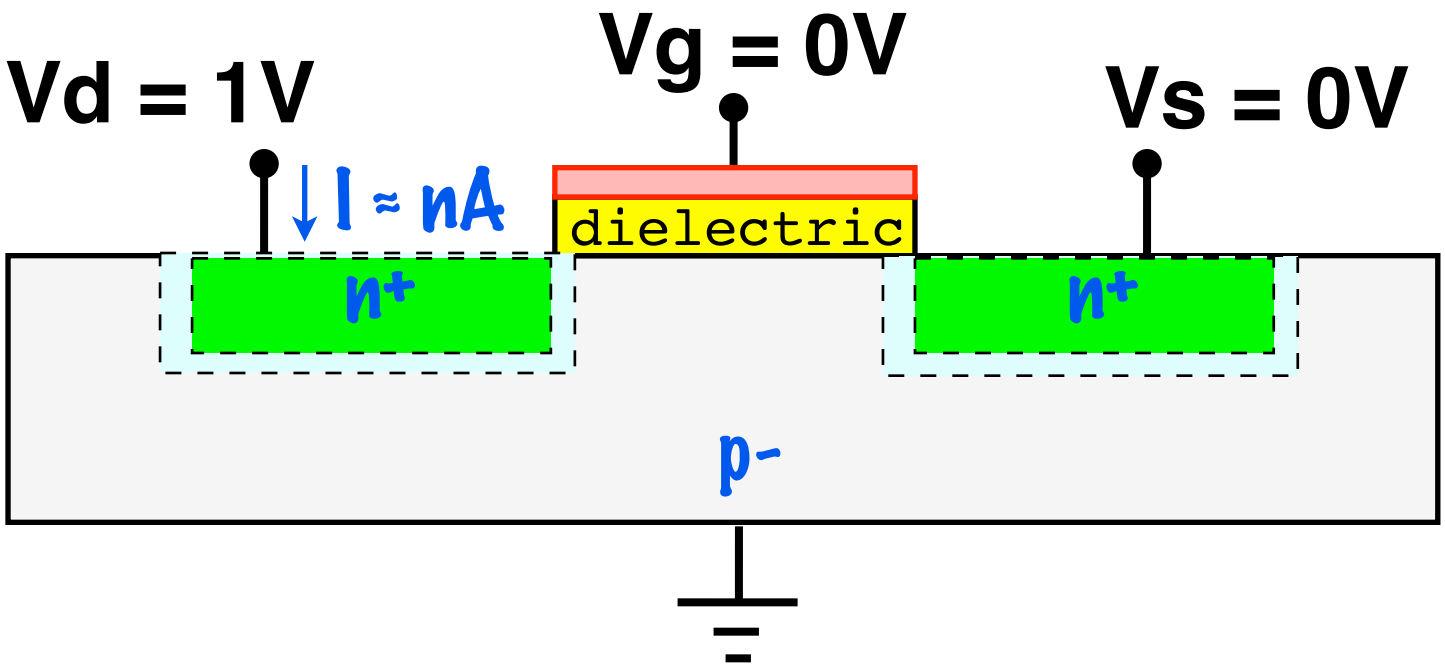

Polysilicon gate, dielectric, and substrate form a capacitor.

nFet is off
(I is "leakage")

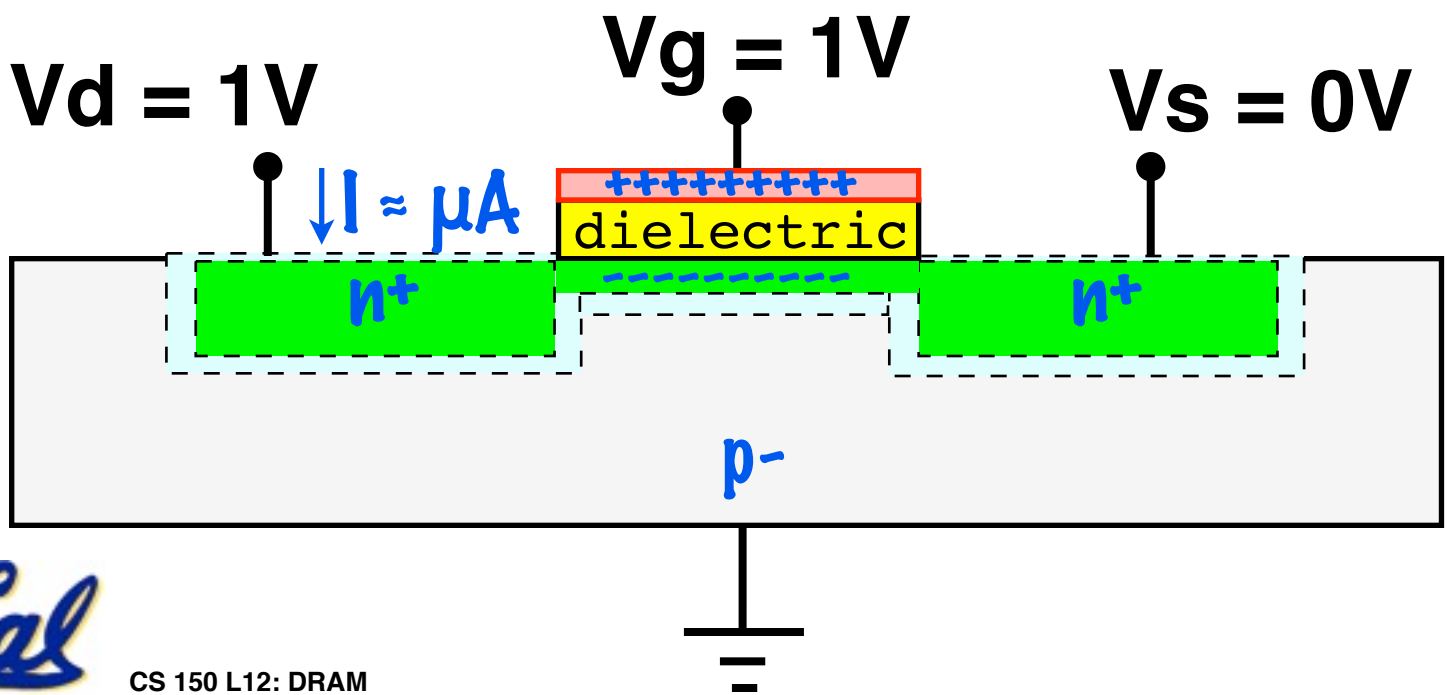

Vg = 1V, small region near the surface turns from p-type to n-type.

nFet is on

Mask set for an n-Fet (circa 1986)

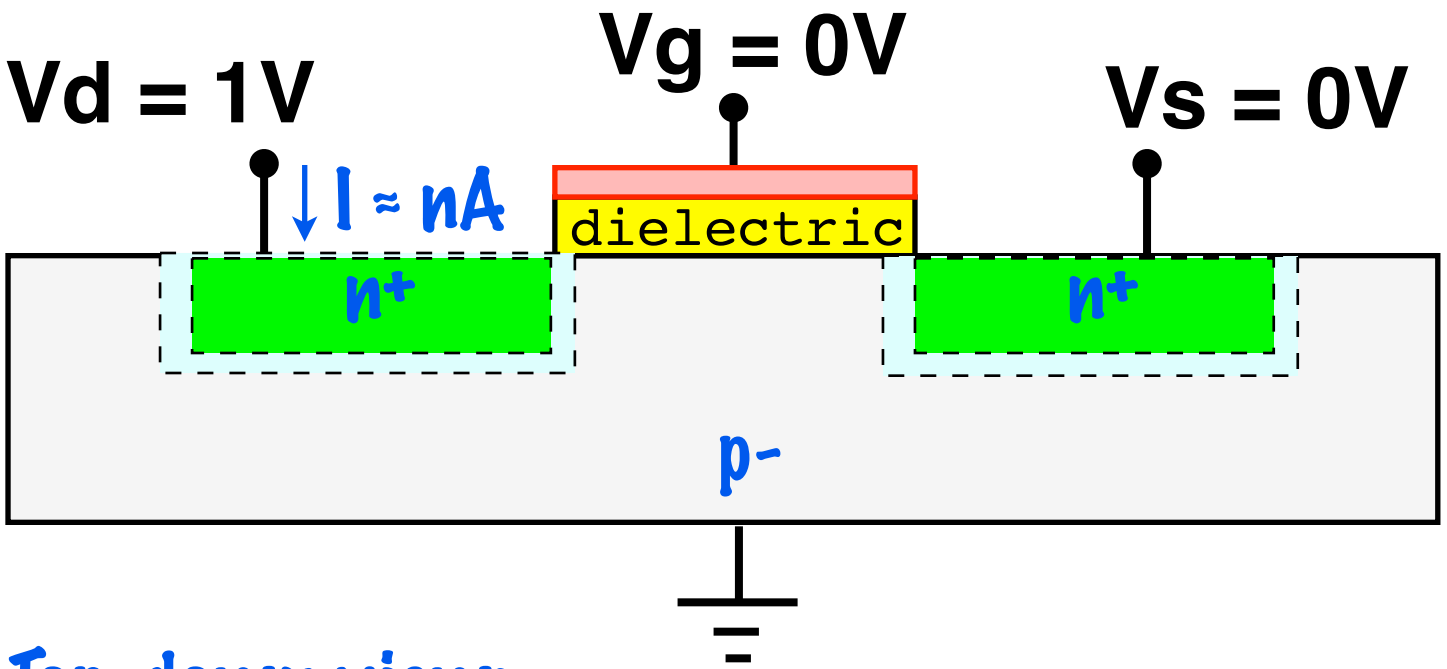

Top-down view:

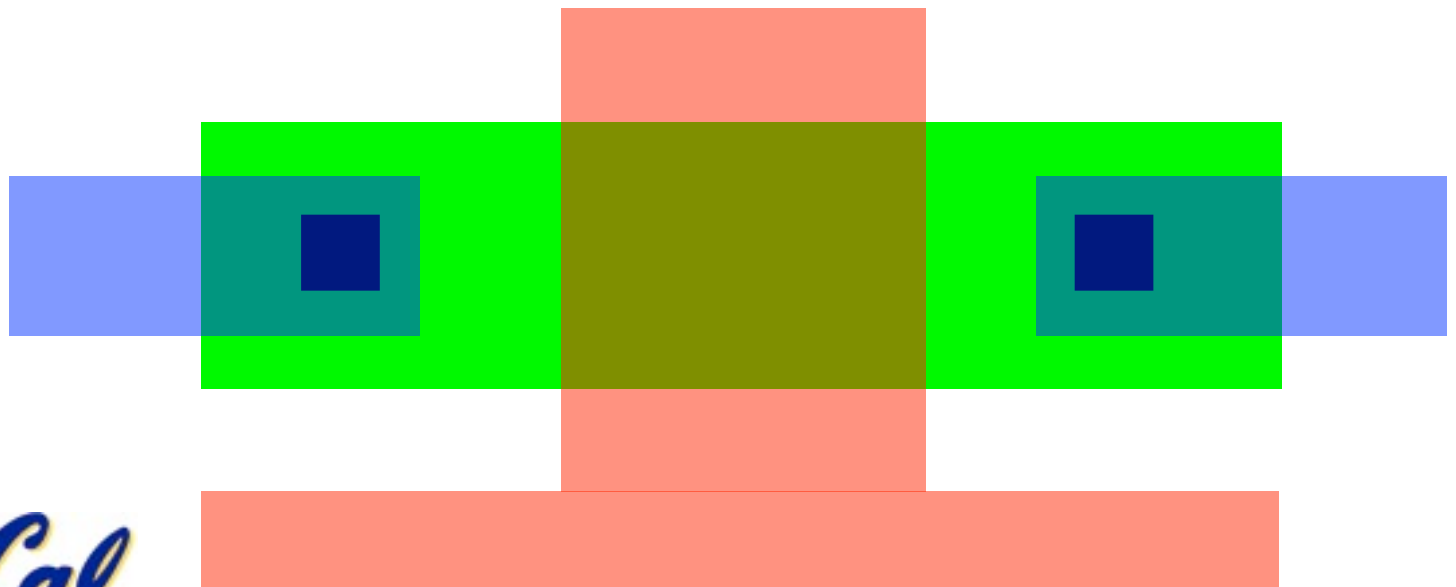

Masks

#1: n+ diffusion

#2: poly (gate)

#3: diff contact

#4: metal

Layers to do p-Fet not shown. Modern processes have 6 to 10 metal layers (or more) (in 1986: 2).

Dynamic Memory Cells

Recall: Capacitors in action

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

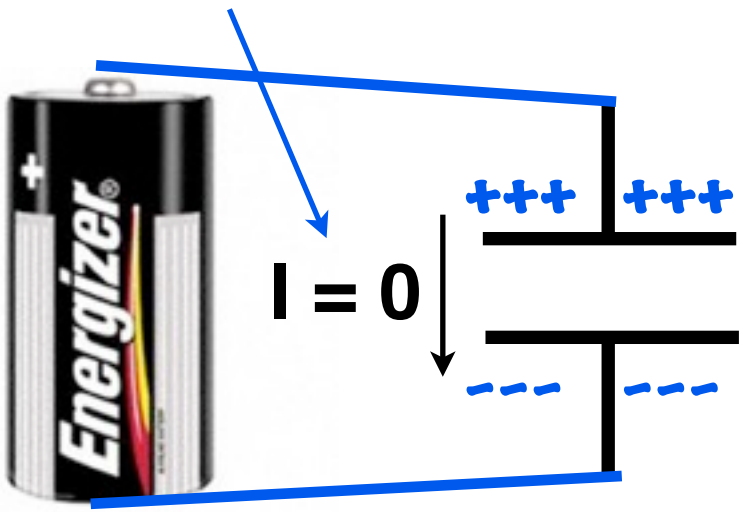

After circuit "settles" ...

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

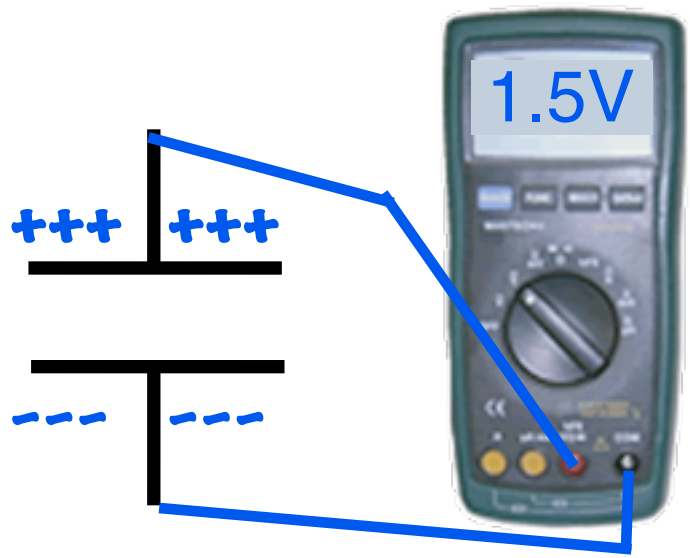

DRAM cell: 1 transistor, 1 capacitor

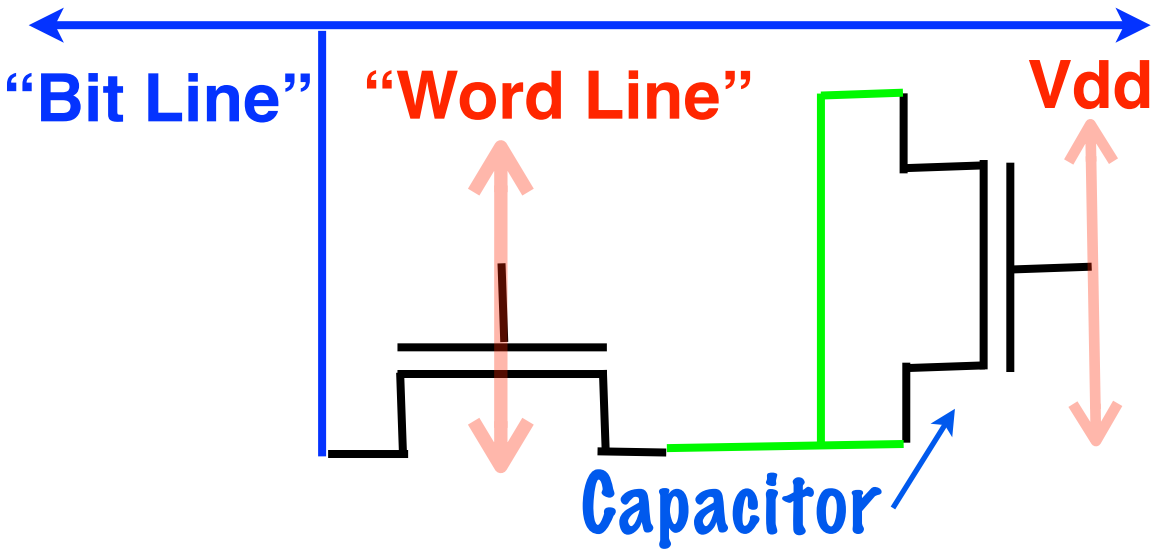

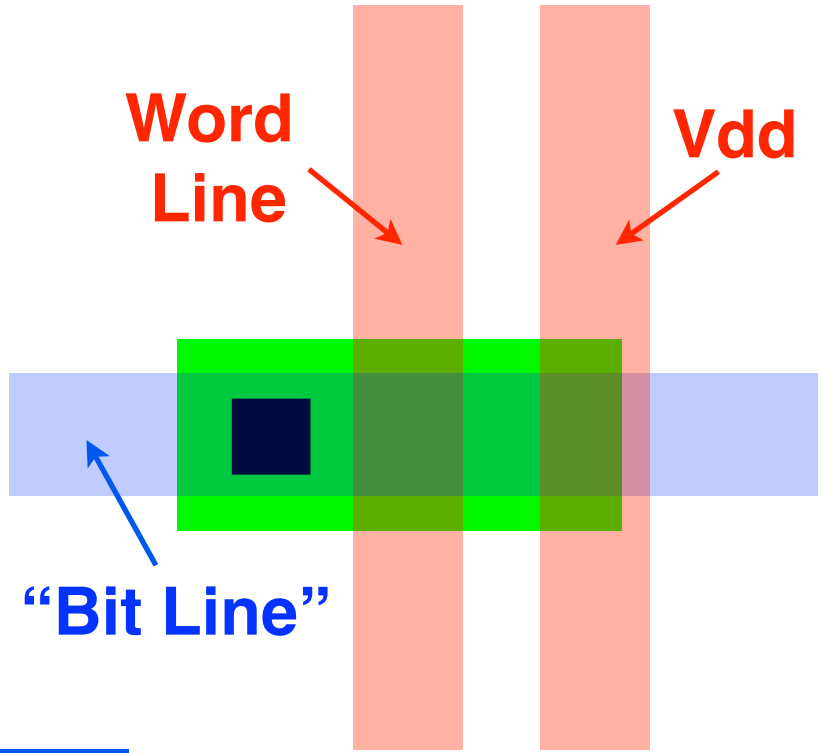

"Bit Line"

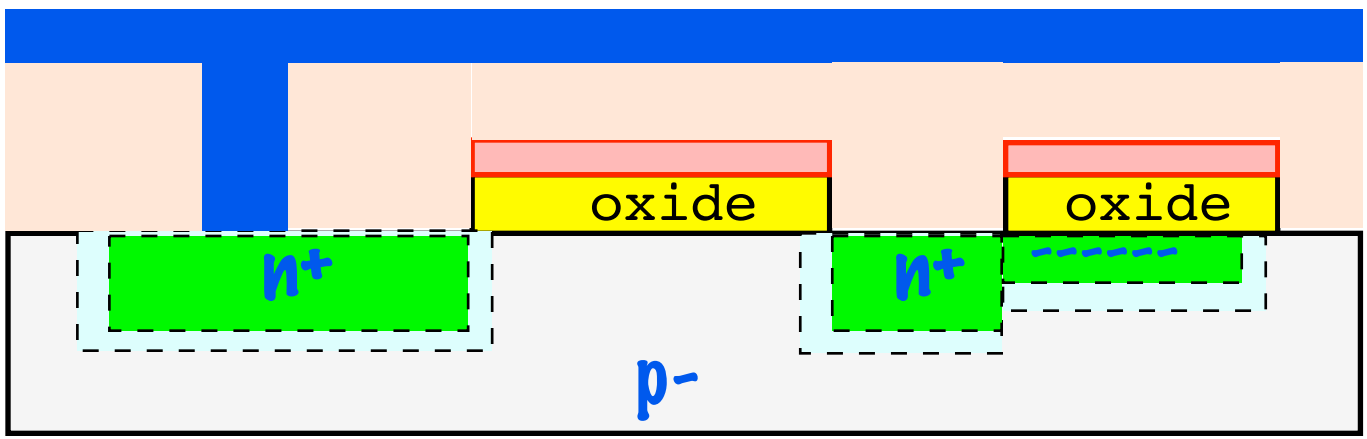

Word Line and Vdd run on "z-axis"

Why Vcap values start out at ground.

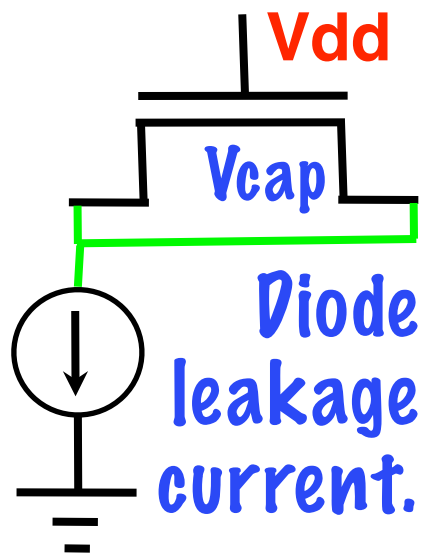

Diode leakage current.

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

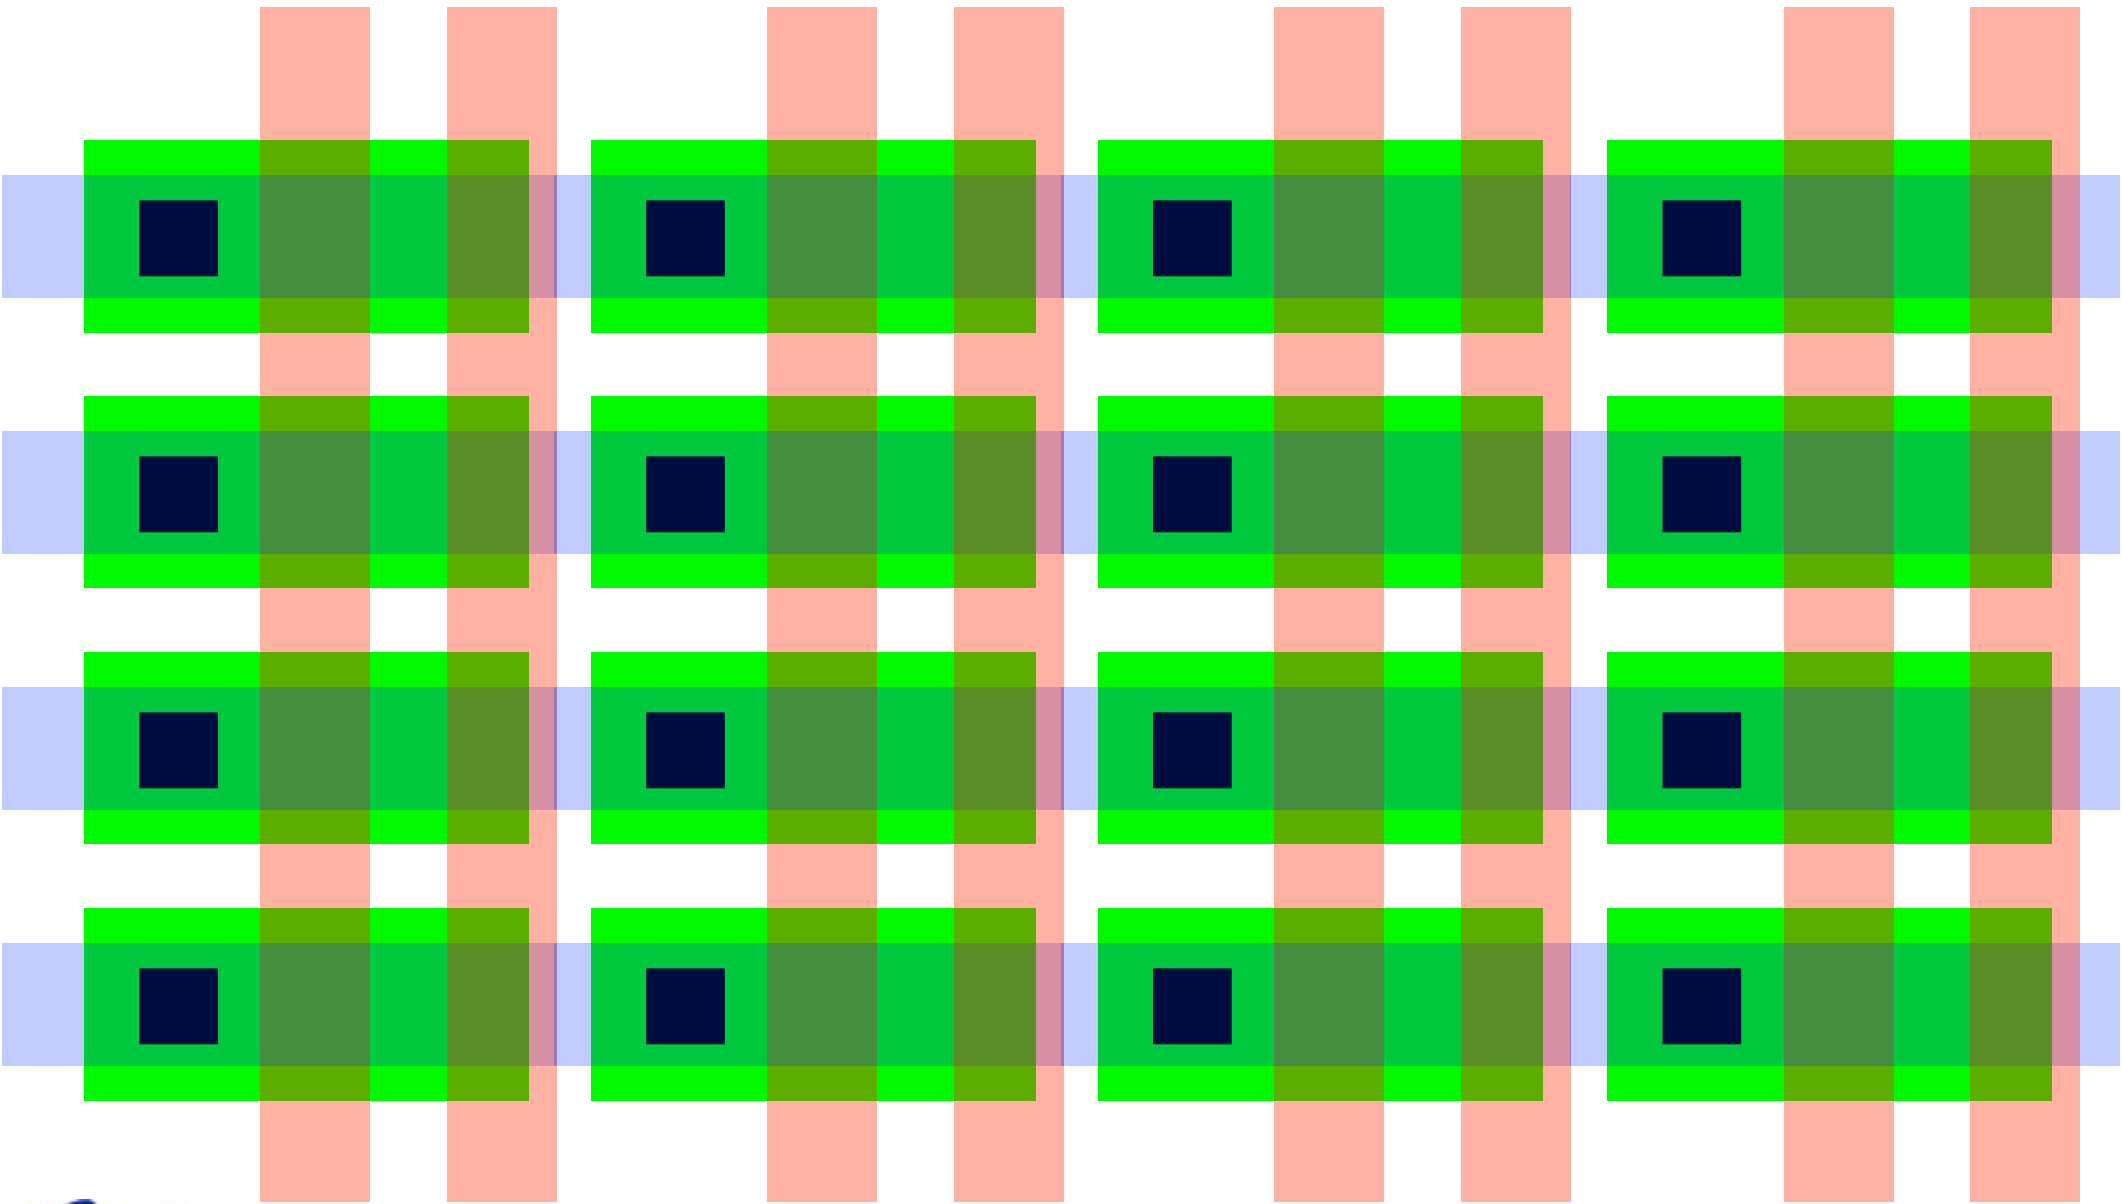

Invented after SRAM, by Robert Dennard

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

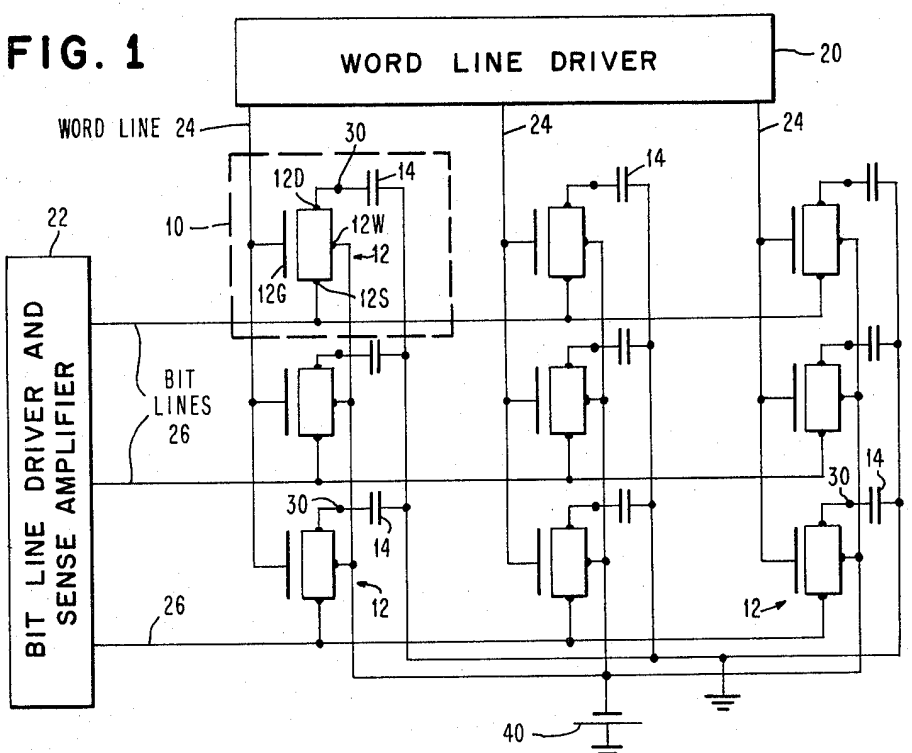

FIG. 2

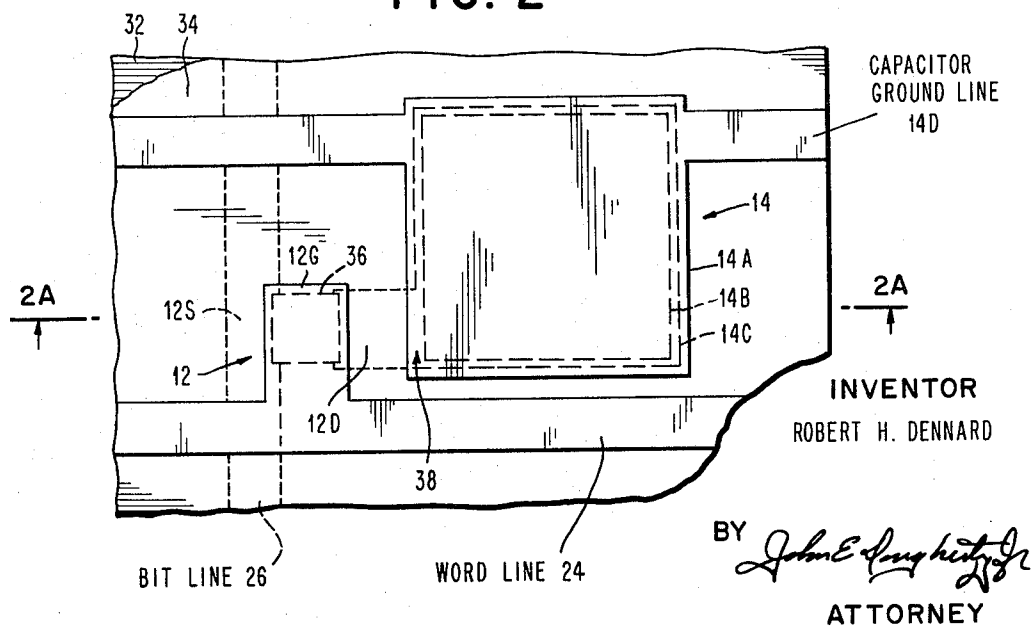

DRAM Circuit Challenge #1: Writing

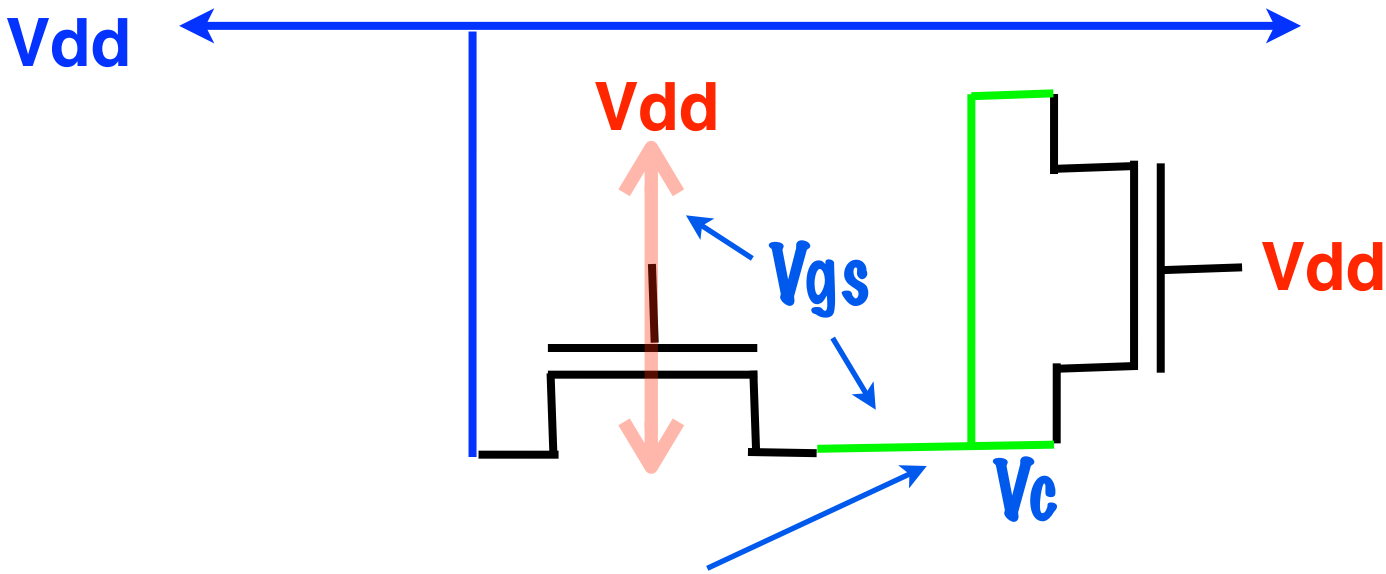

Vdd - Vth. Bad, we store less charge. Why do we not get Vdd?

Ids = k [Vgs -Vth]^2 ,
but "turns off" when Vgs <= Vth!

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

DRAM Challenge #2: Destructive Reads

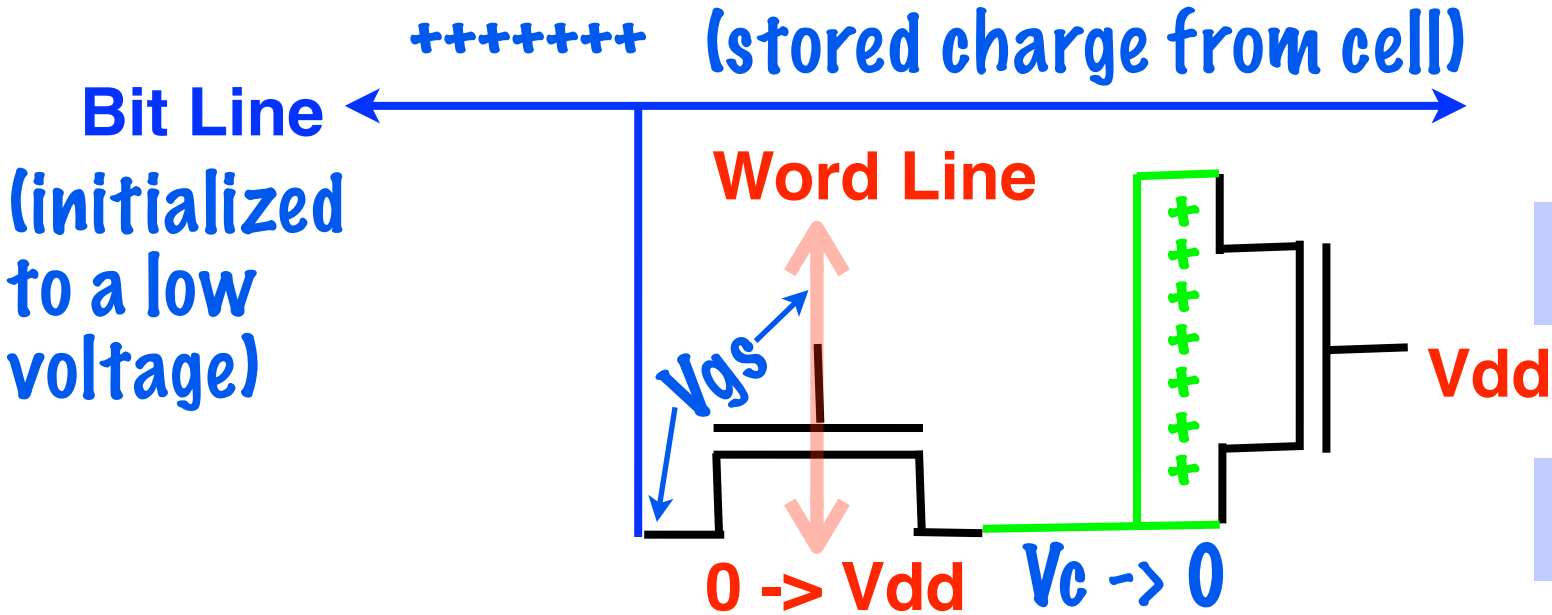

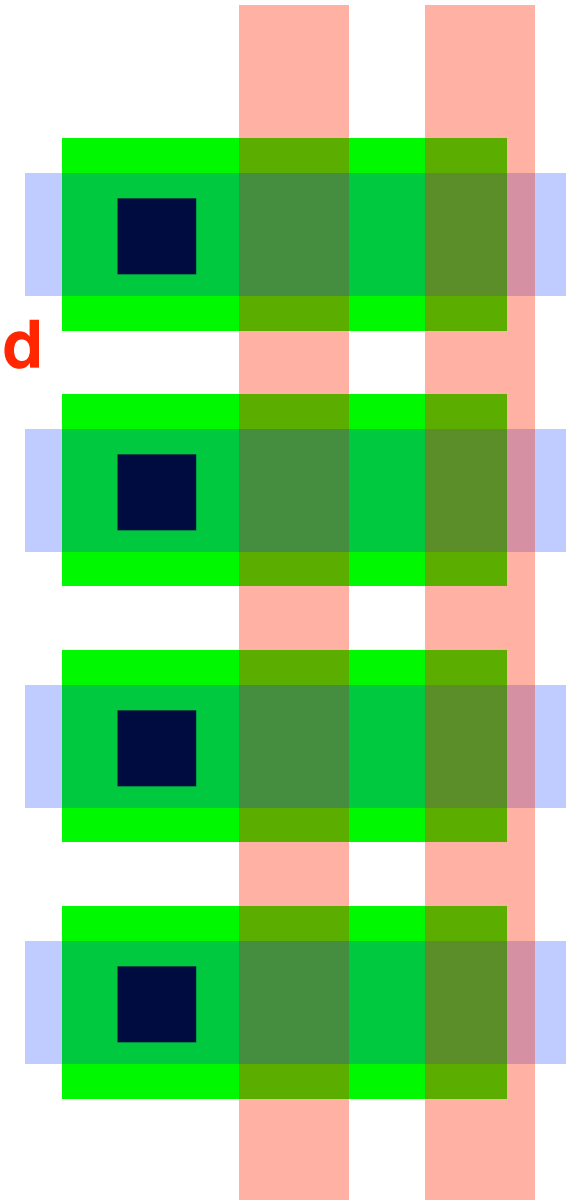

Raising the word line removes the charge from every cell it connects to!

DRAMs write back after each read.

DRAM Circuit Challenge #3a: Sensing

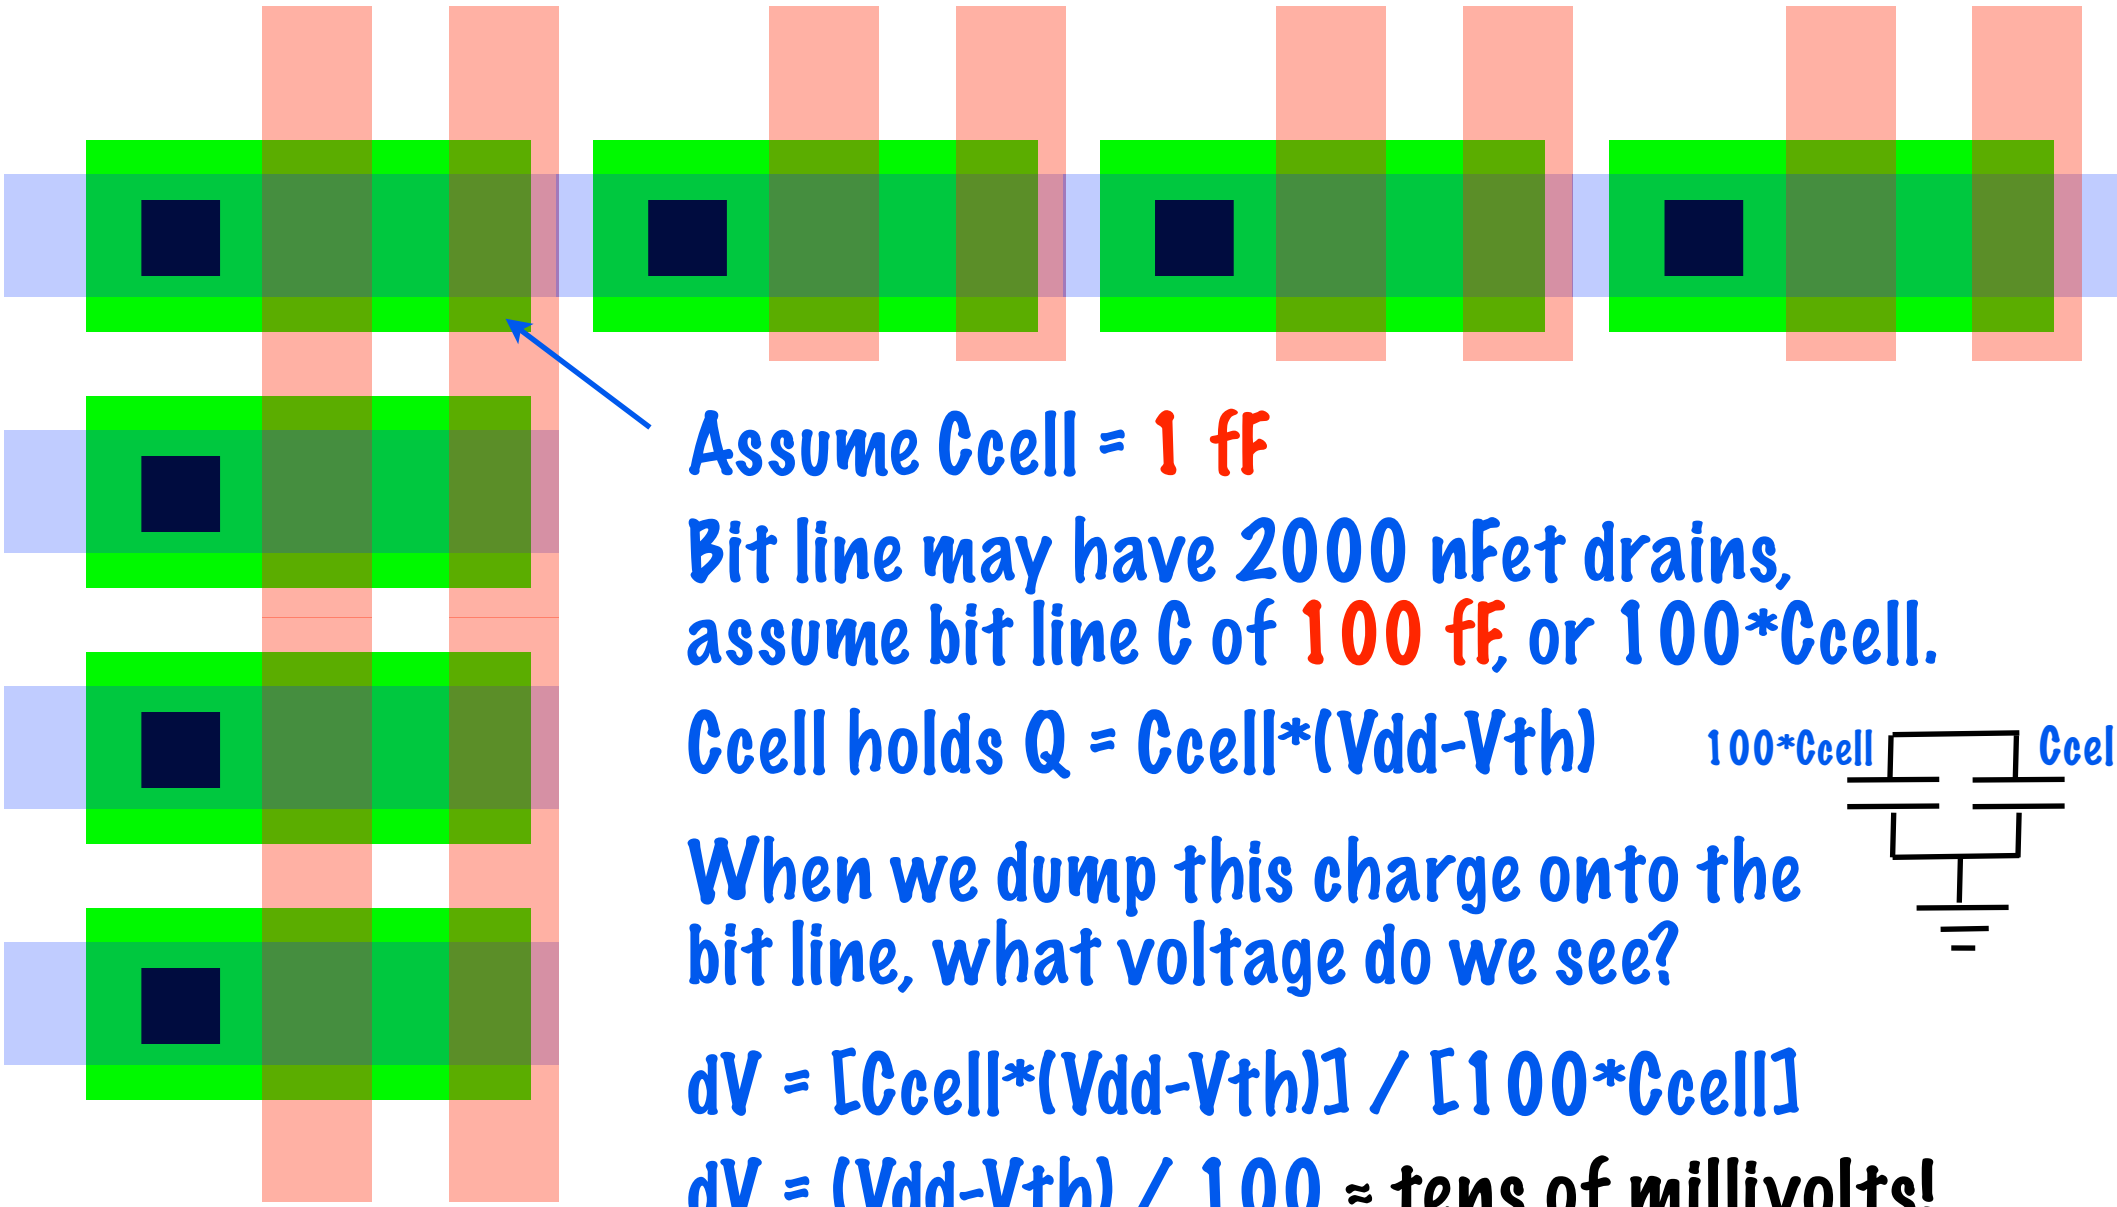

Assume Ccell = 1 fF

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

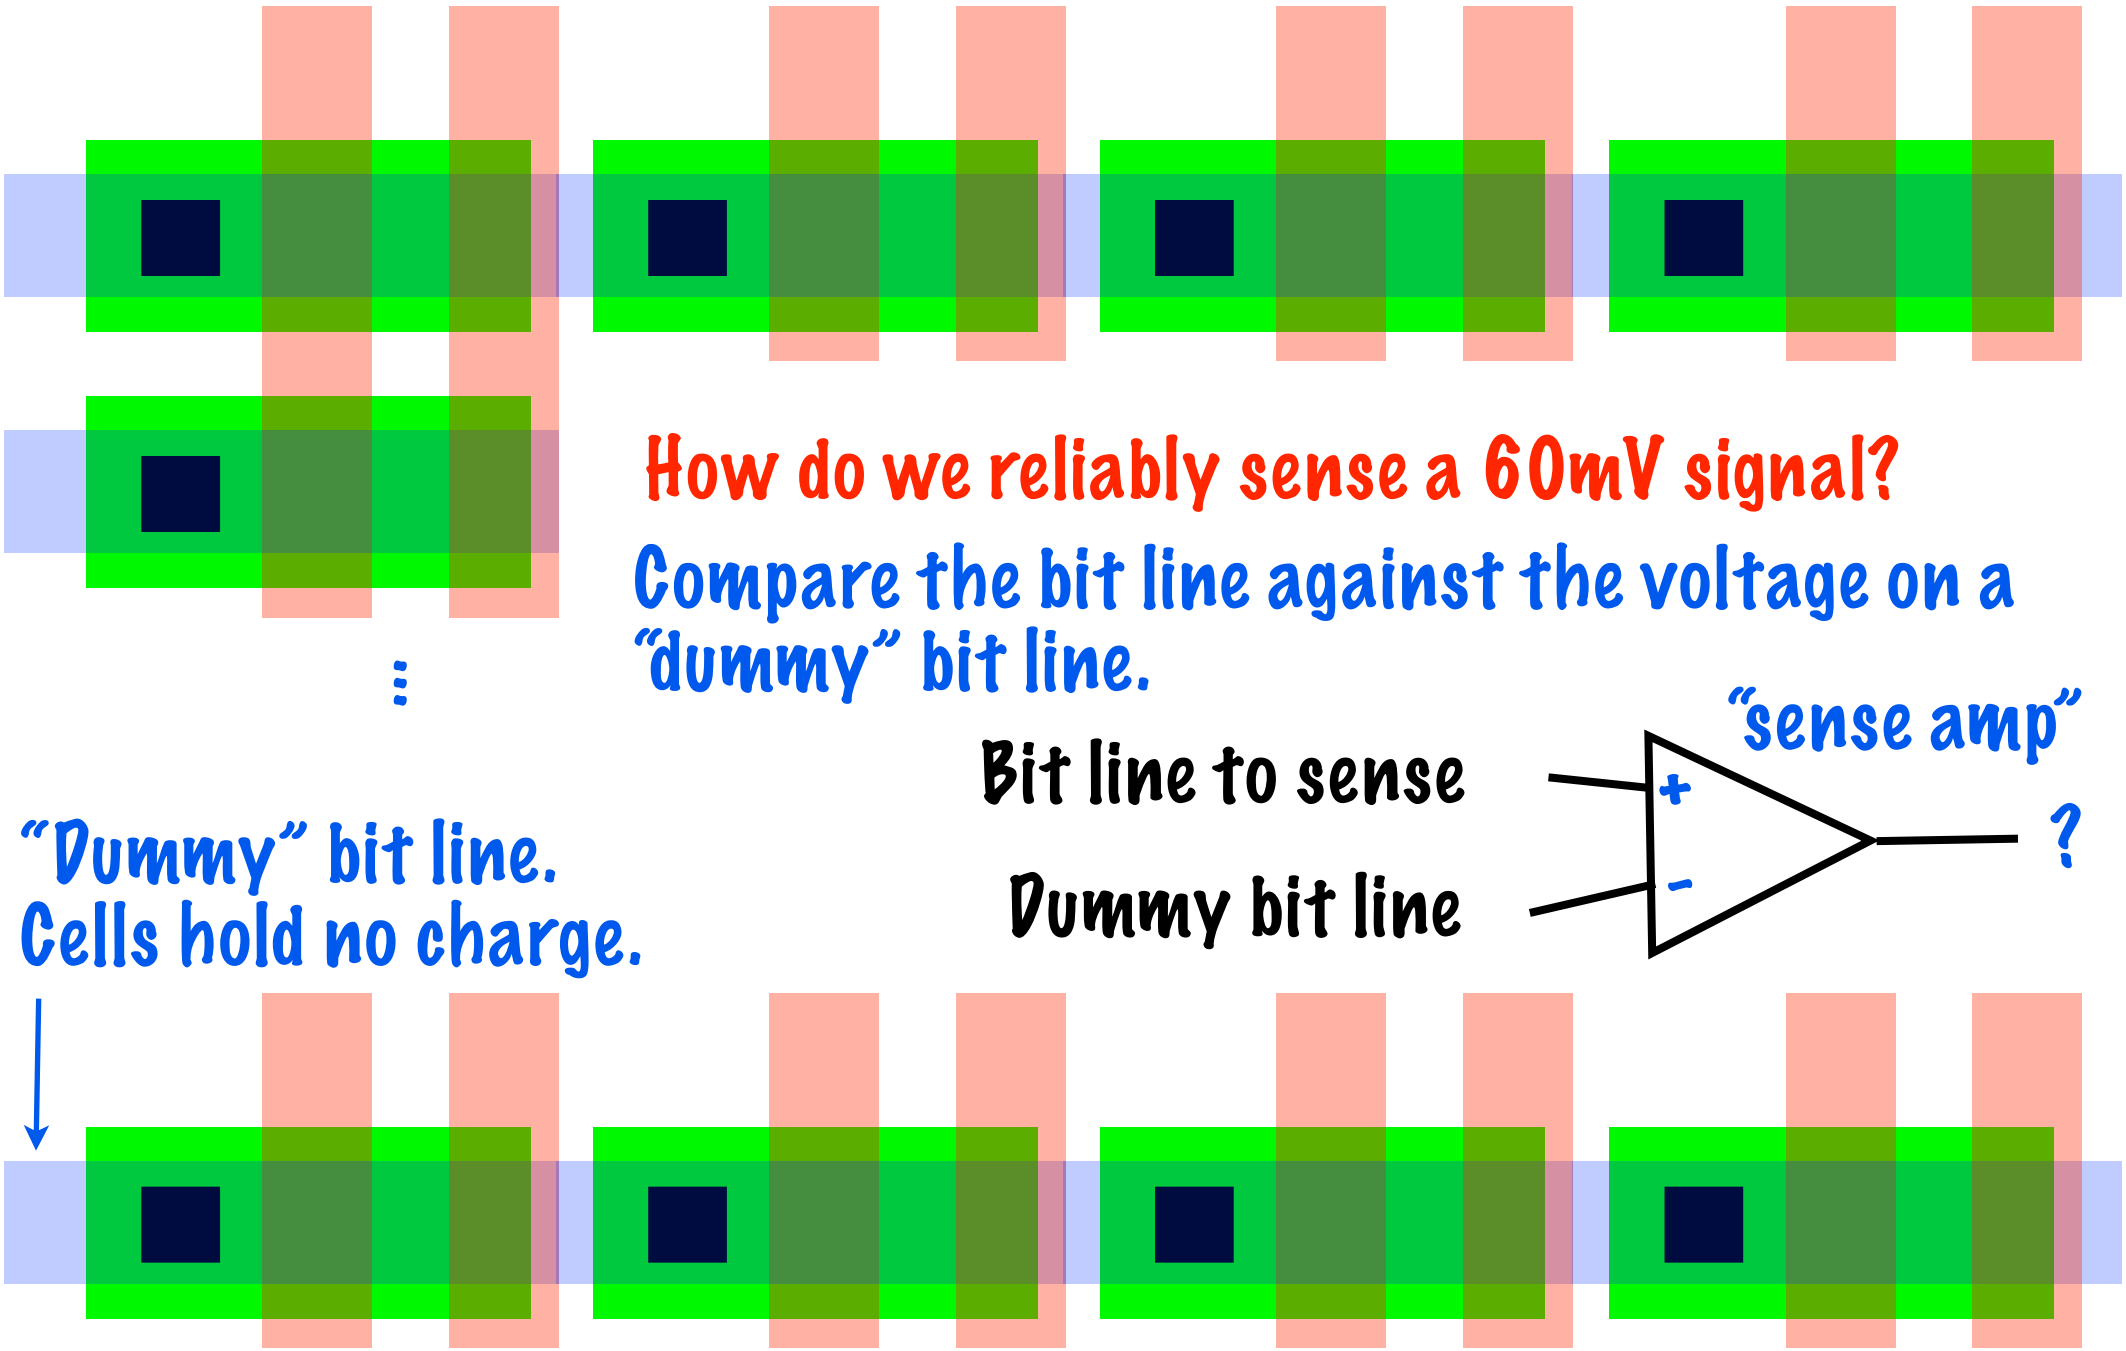

How do we reliably sense a 60mV signal?

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

"sense amp"

Bit line to sense

Dummy bit line

+

-

?

"Dummy" bit line.
Cells hold no charge.

DRAM Challenge #4: Leakage ...

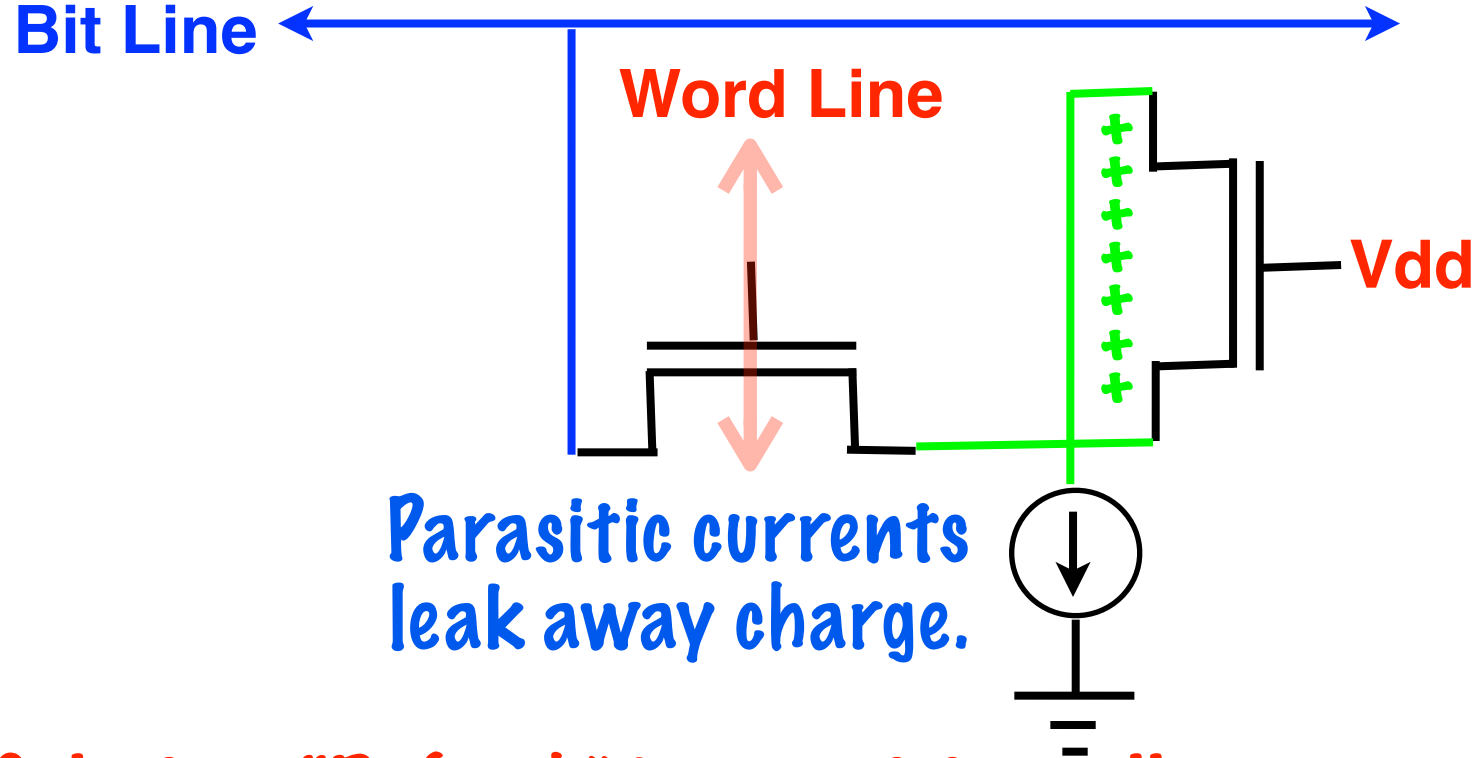

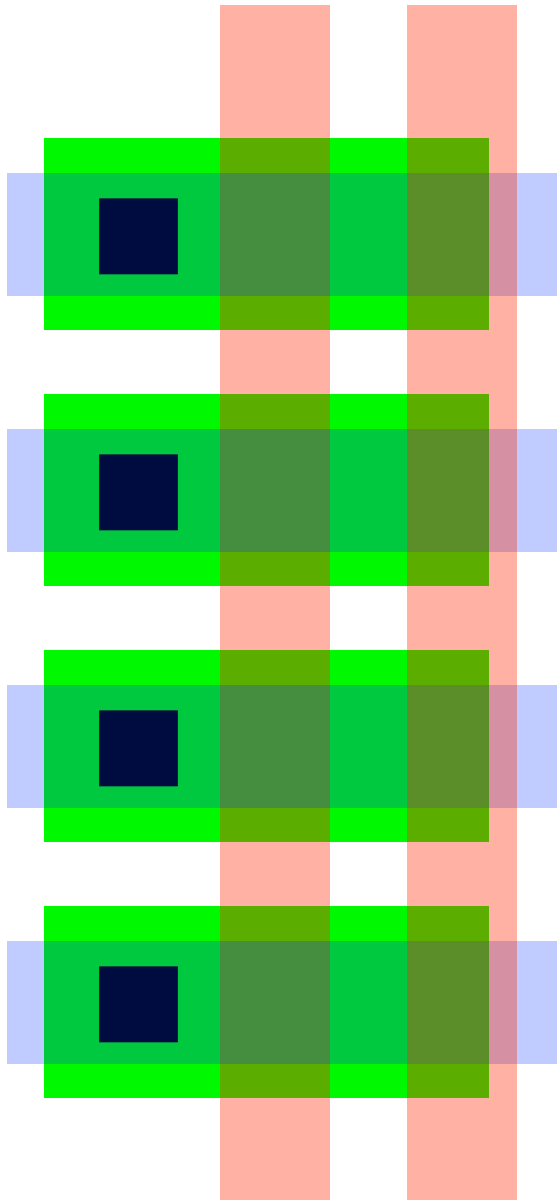

Solution: "Refresh", by rewriting cells at regular intervals (tens of milliseconds)

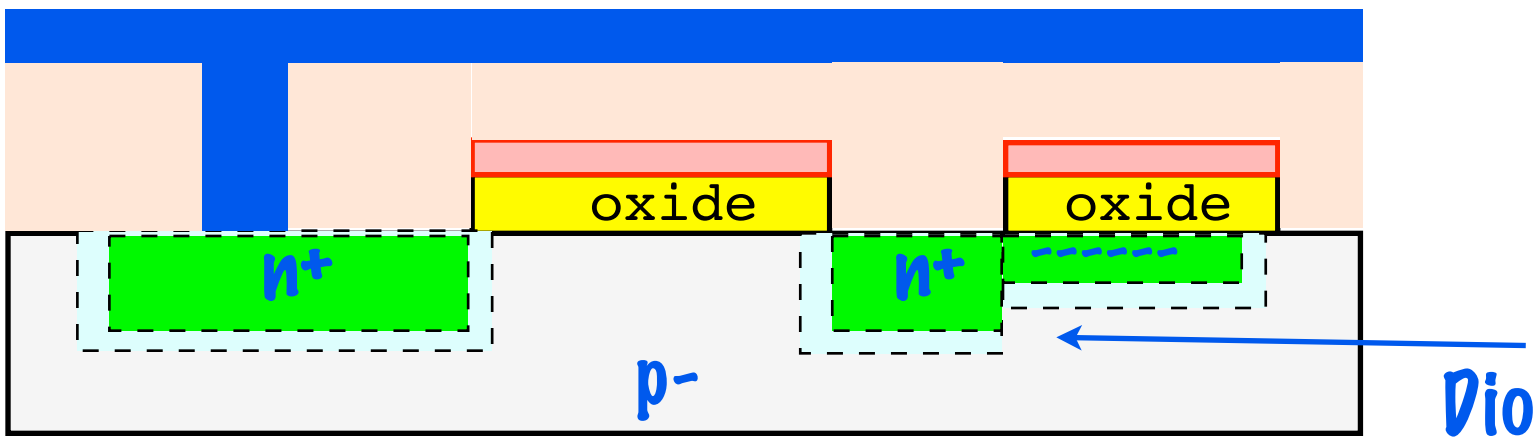

Diode leakage ...

DRAM Challenge #5: Cosmic Rays ...

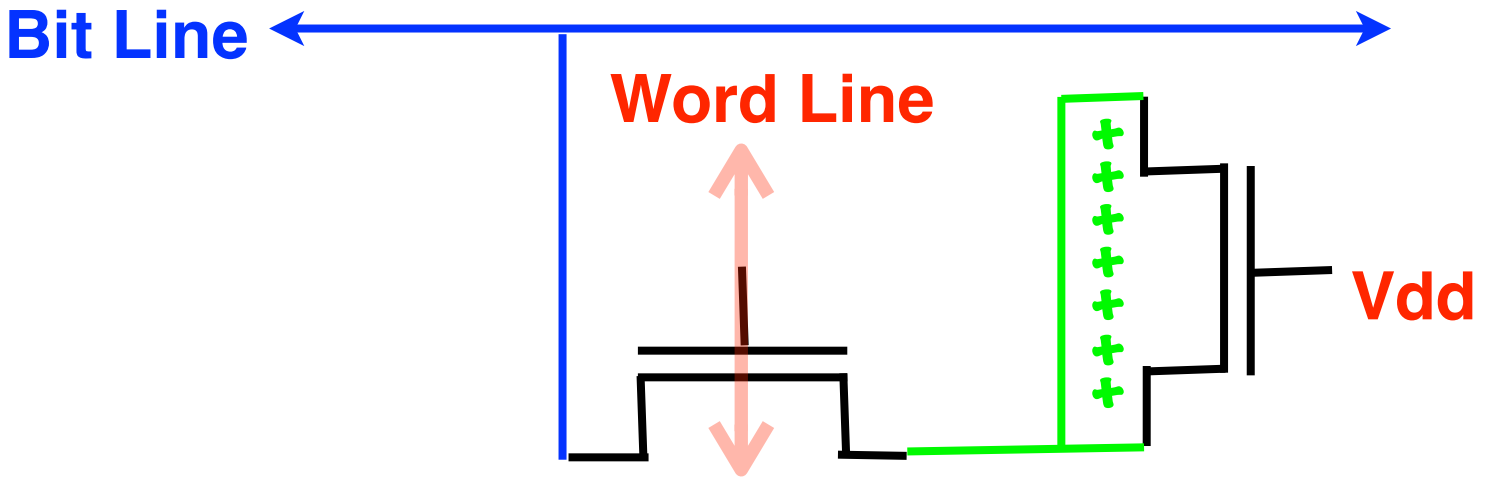

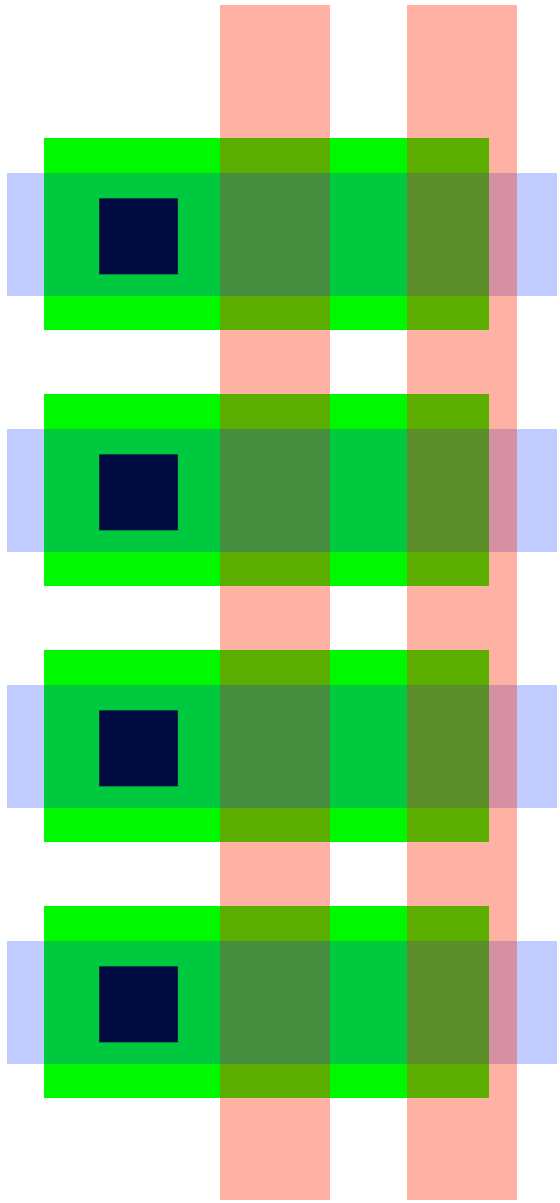

Cell capacitor holds 25,000 electrons (or less). Cosmic rays that constantly bombard us can release the charge!

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

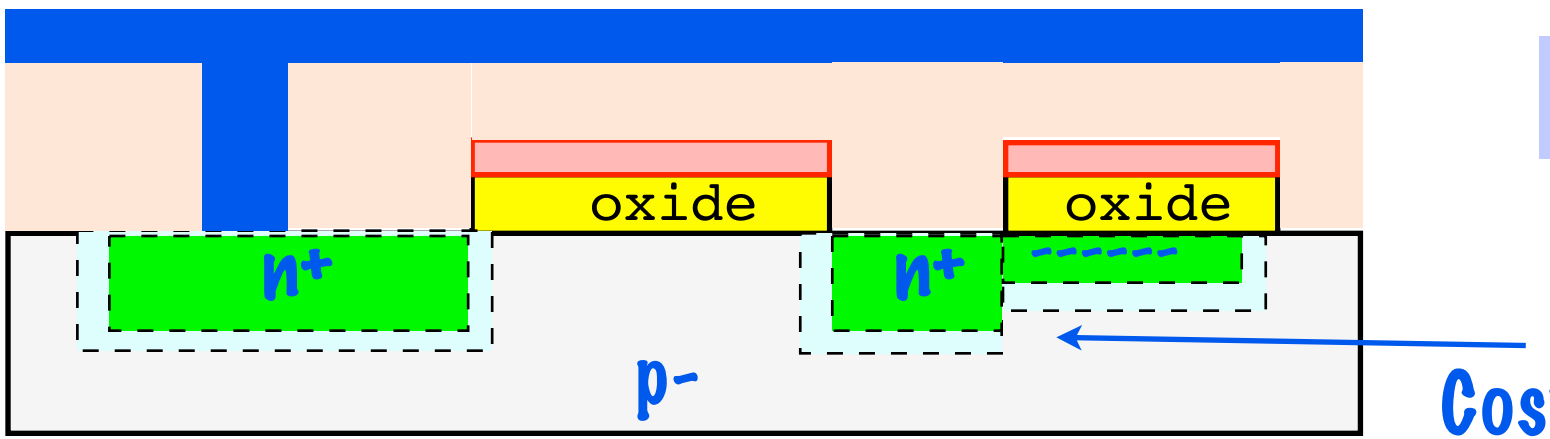

DRAM Challenge 6: Yield

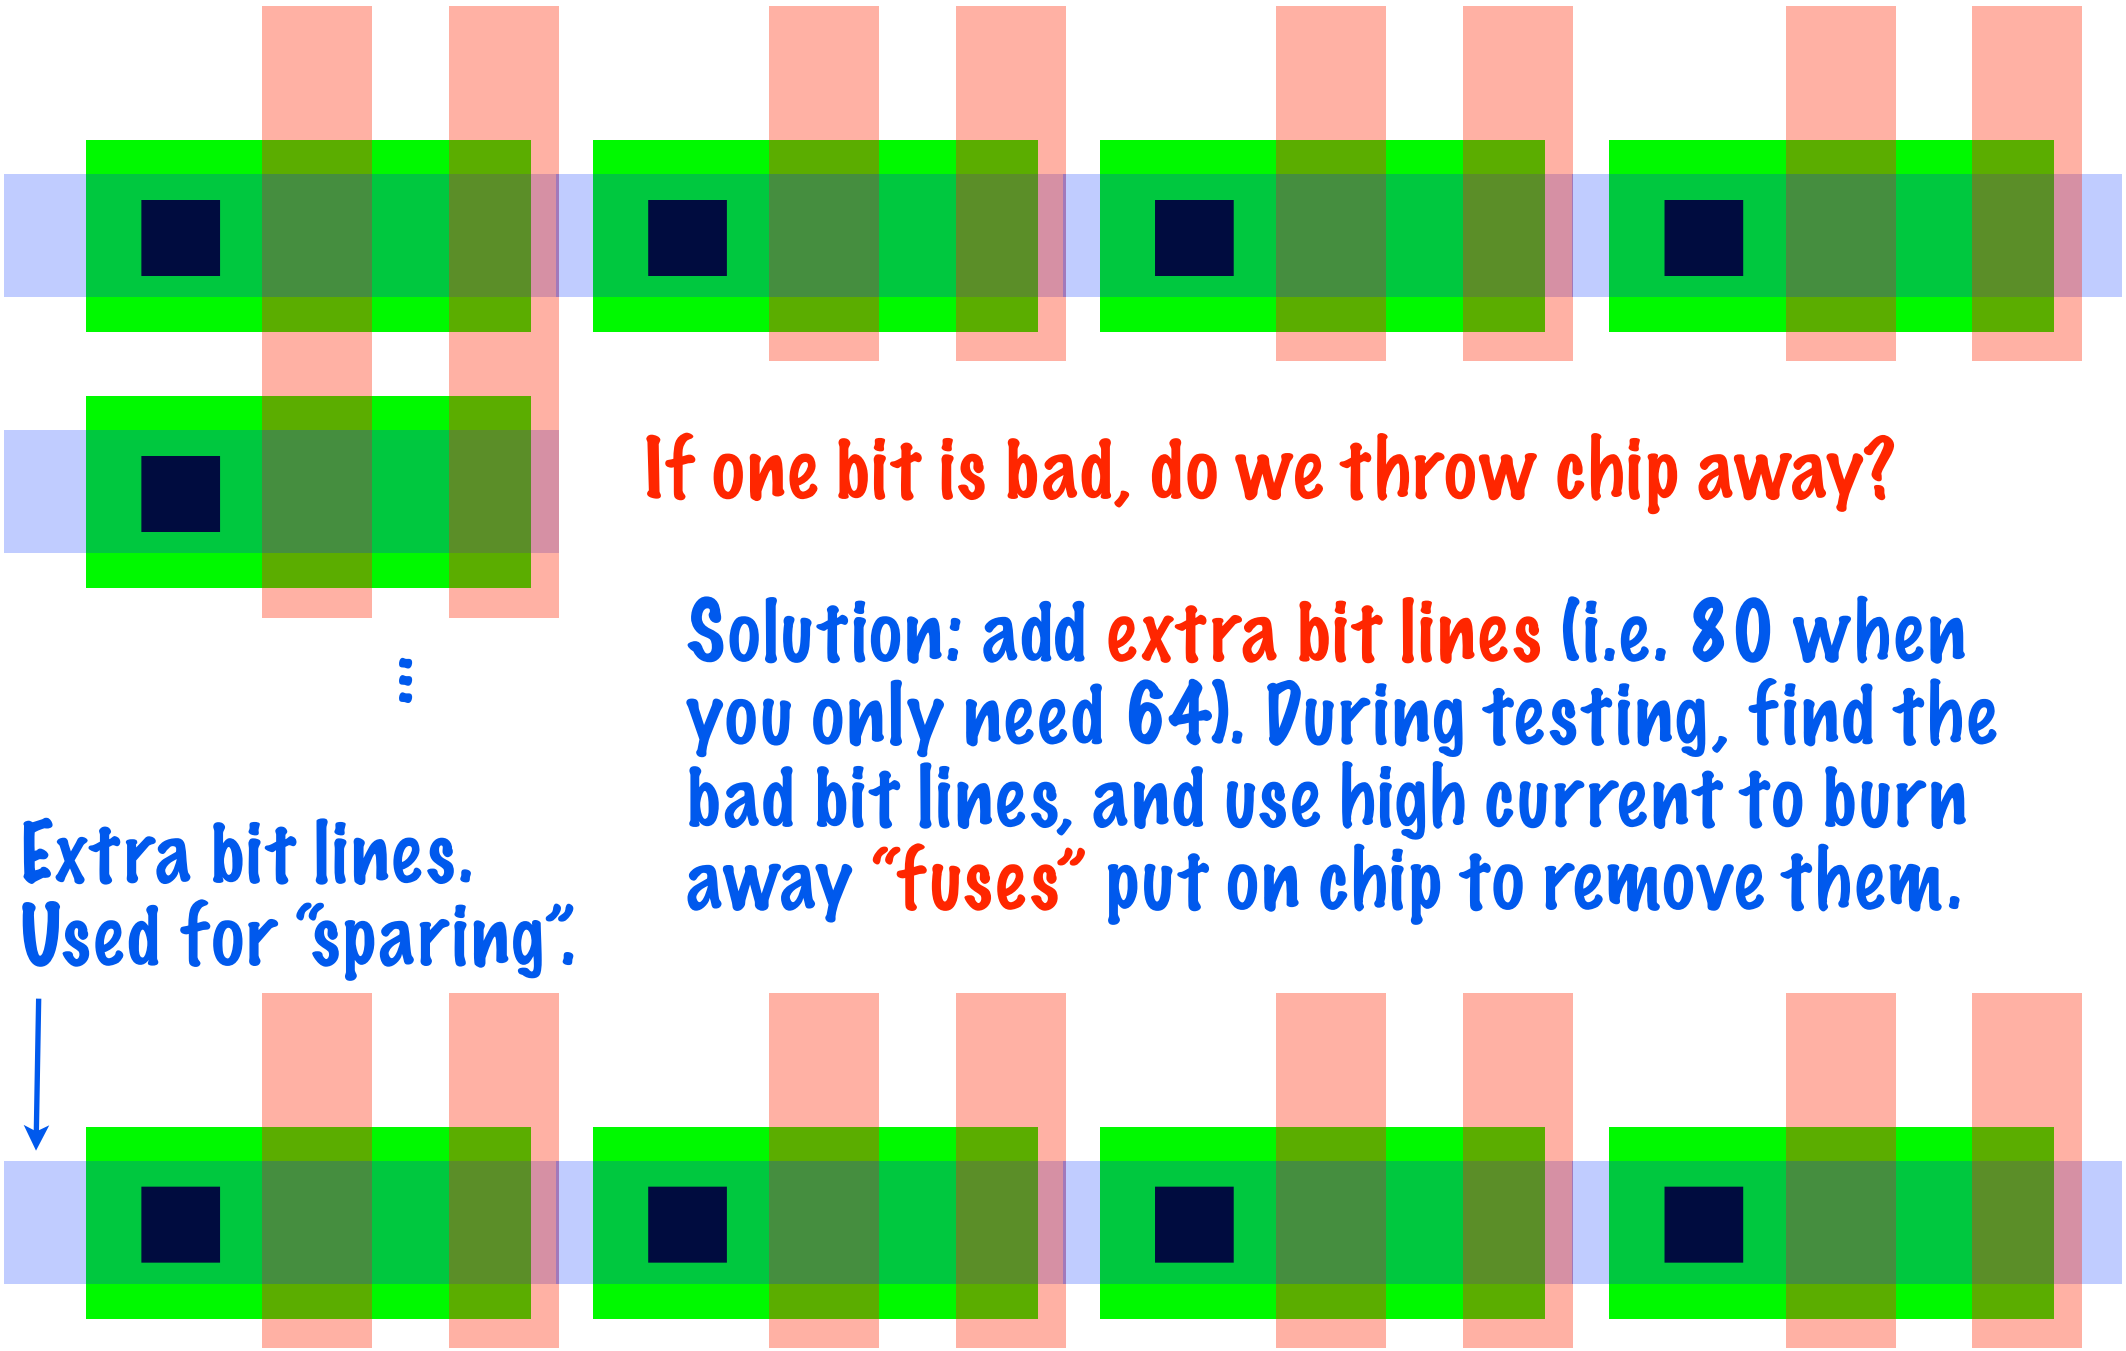

If one bit is bad, do we throw chip away?

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".

Moore's Law for CPUs and DRAMs

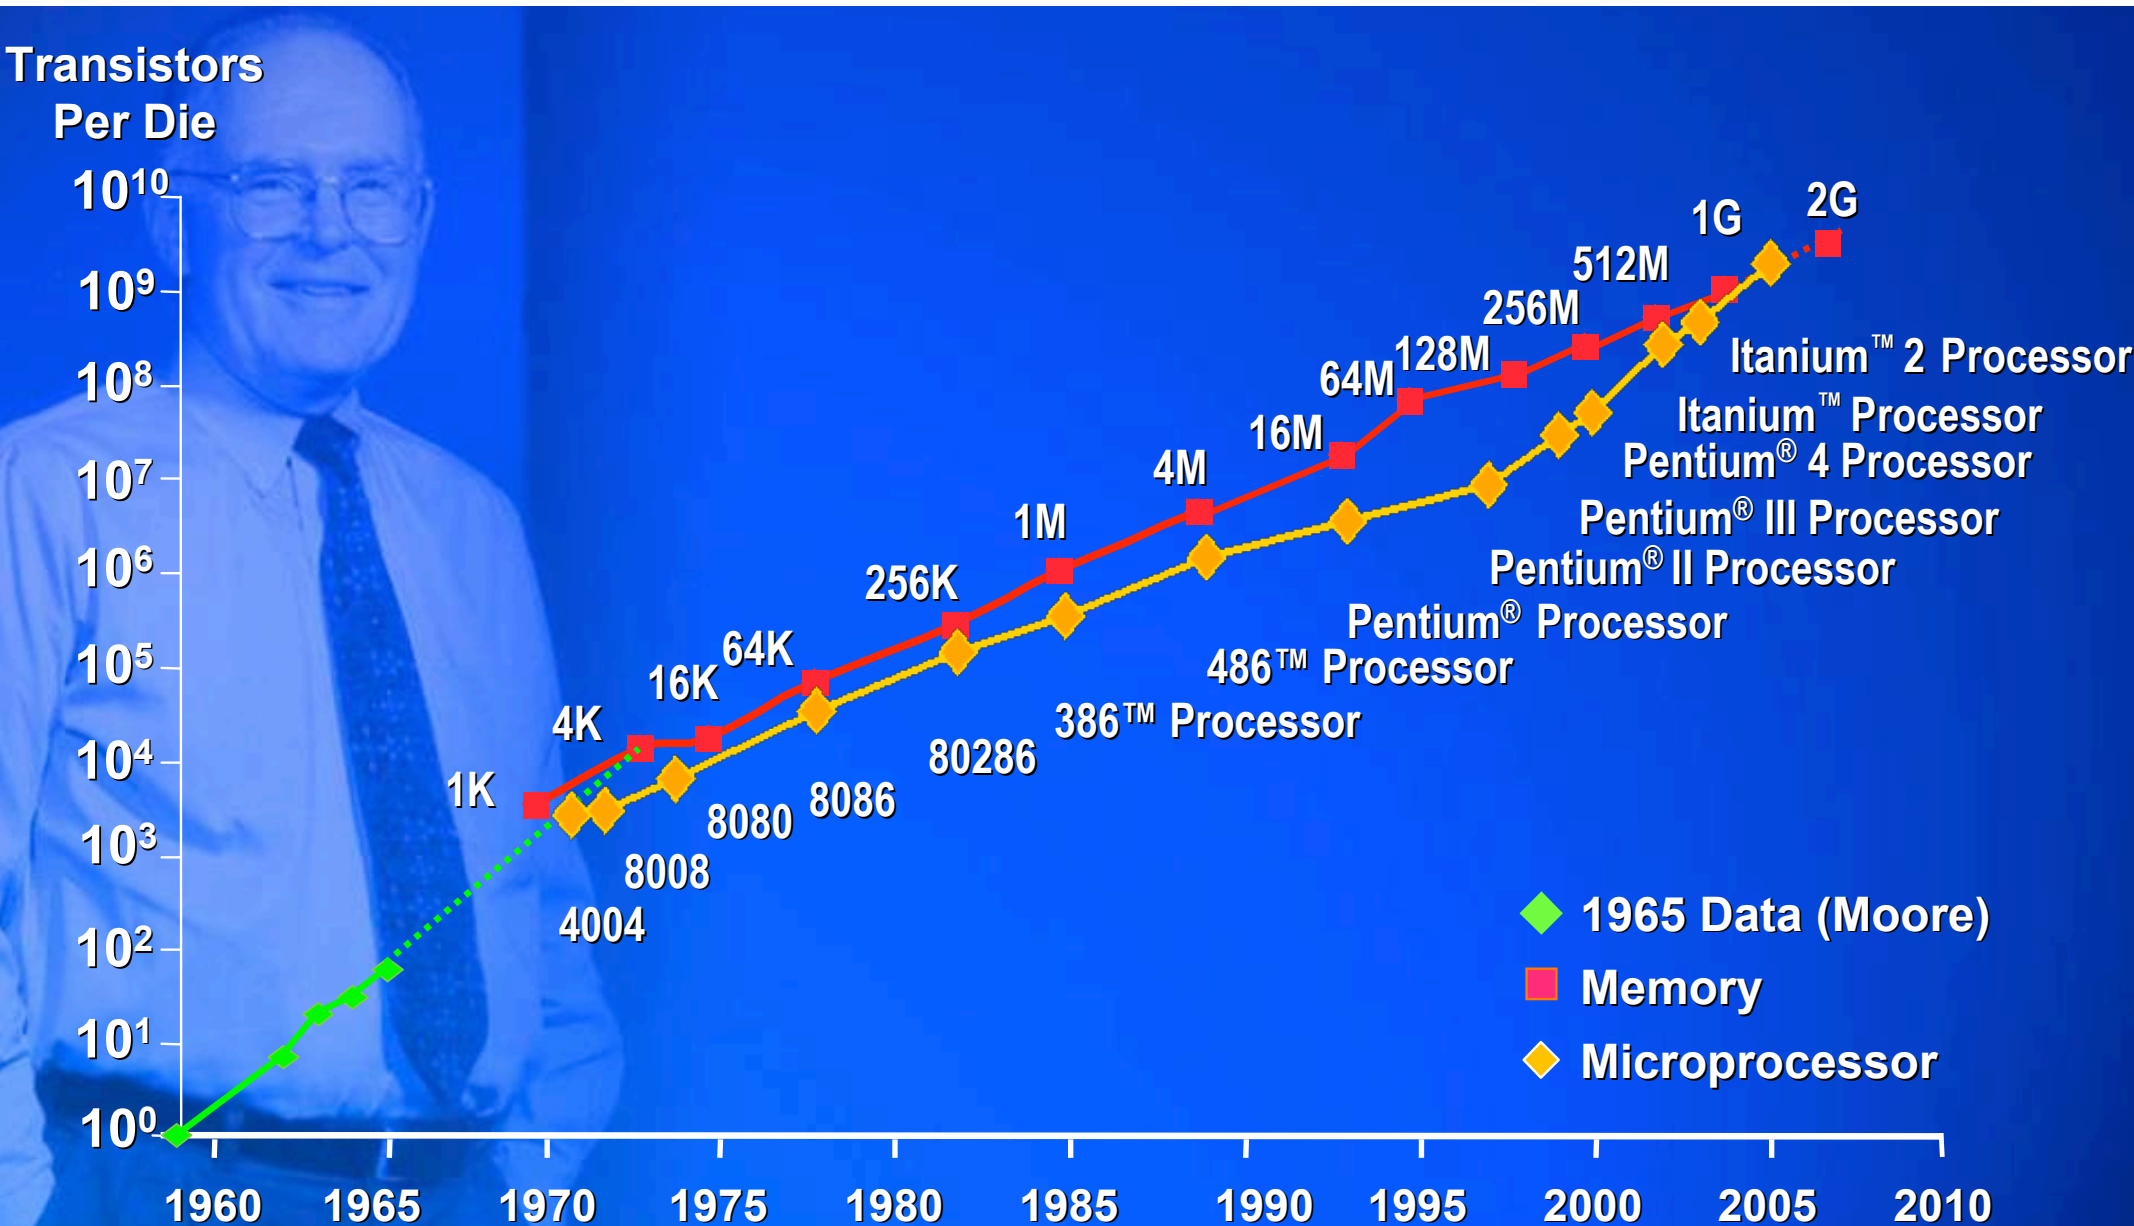

From: "Facing the Hot Chips Challenge Again", Bill Holt, Intel, presented at Hot Chips 17, 2005.

Main driver: device scaling ...

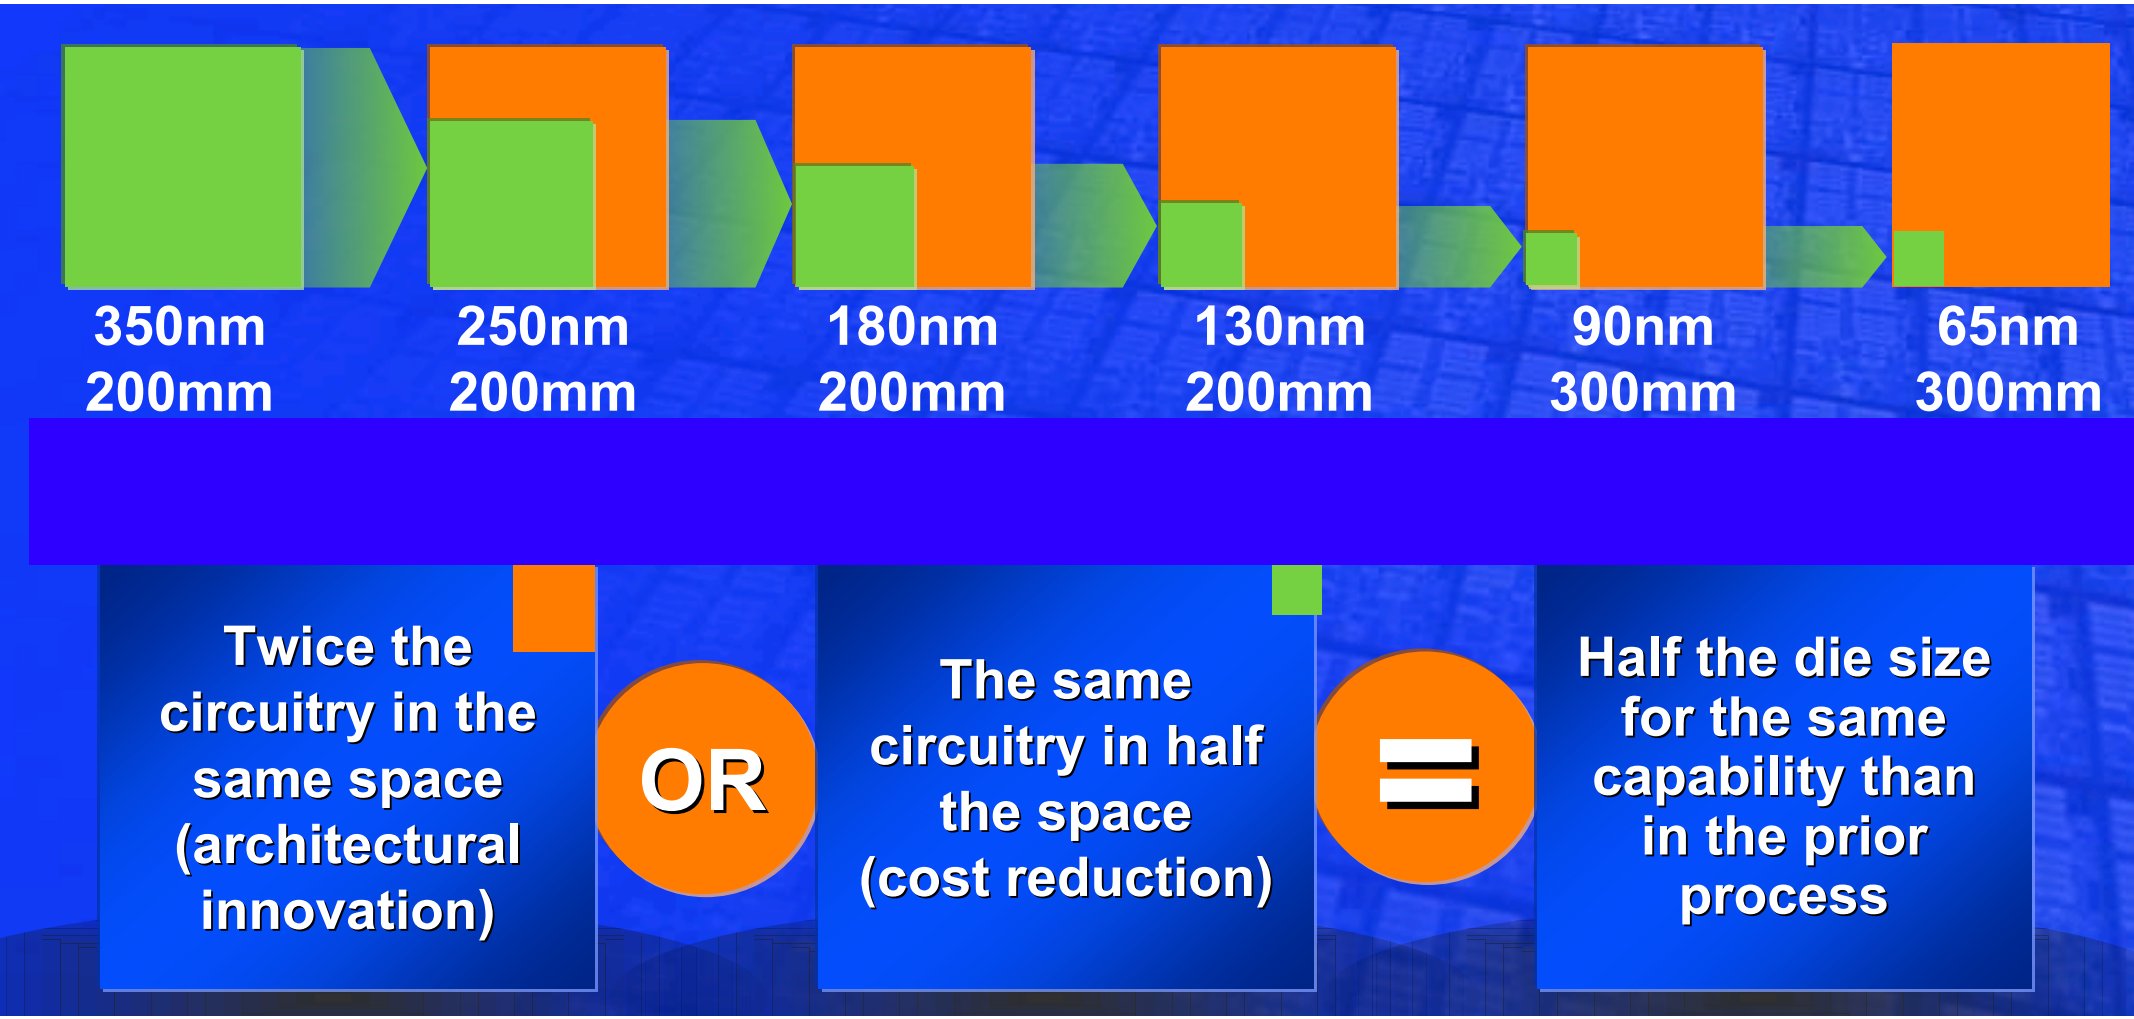

From: “Facing the Hot Chips Challenge Again”, Bill Holt, Intel, presented at Hot Chips 17, 2005.

Process Scaling: Why chips don't fry

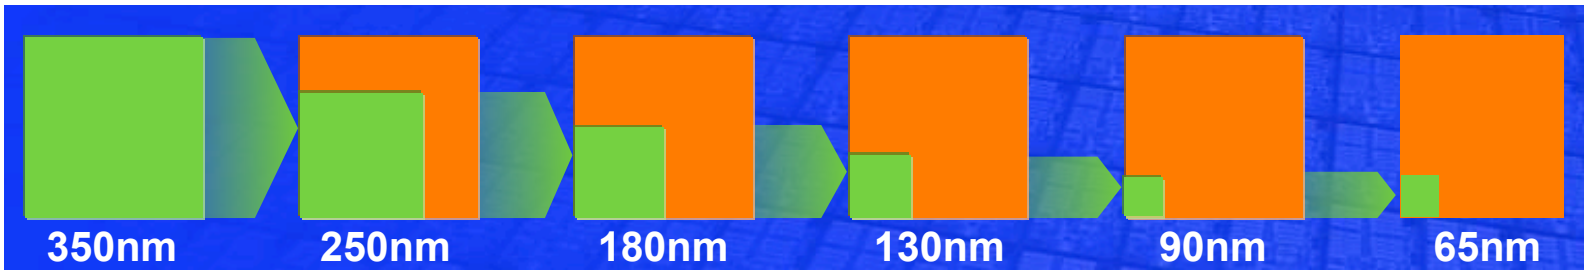

IC process scaling ("Moore's Law")

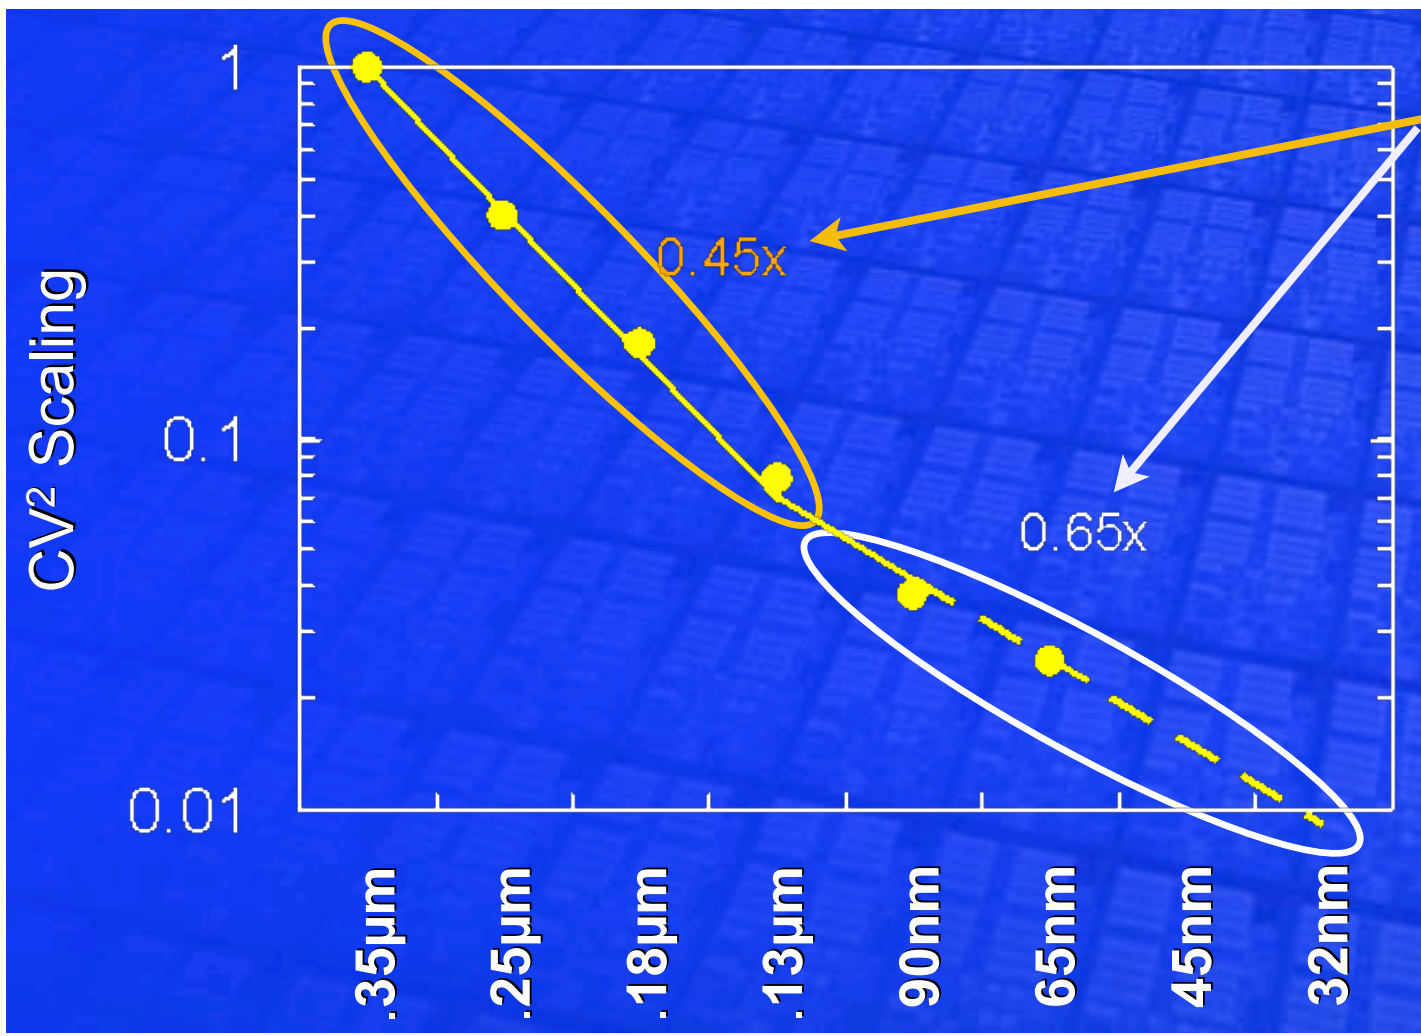

Due to reducing V and C (length and width of Cs decrease, but plate distance gets smaller).

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

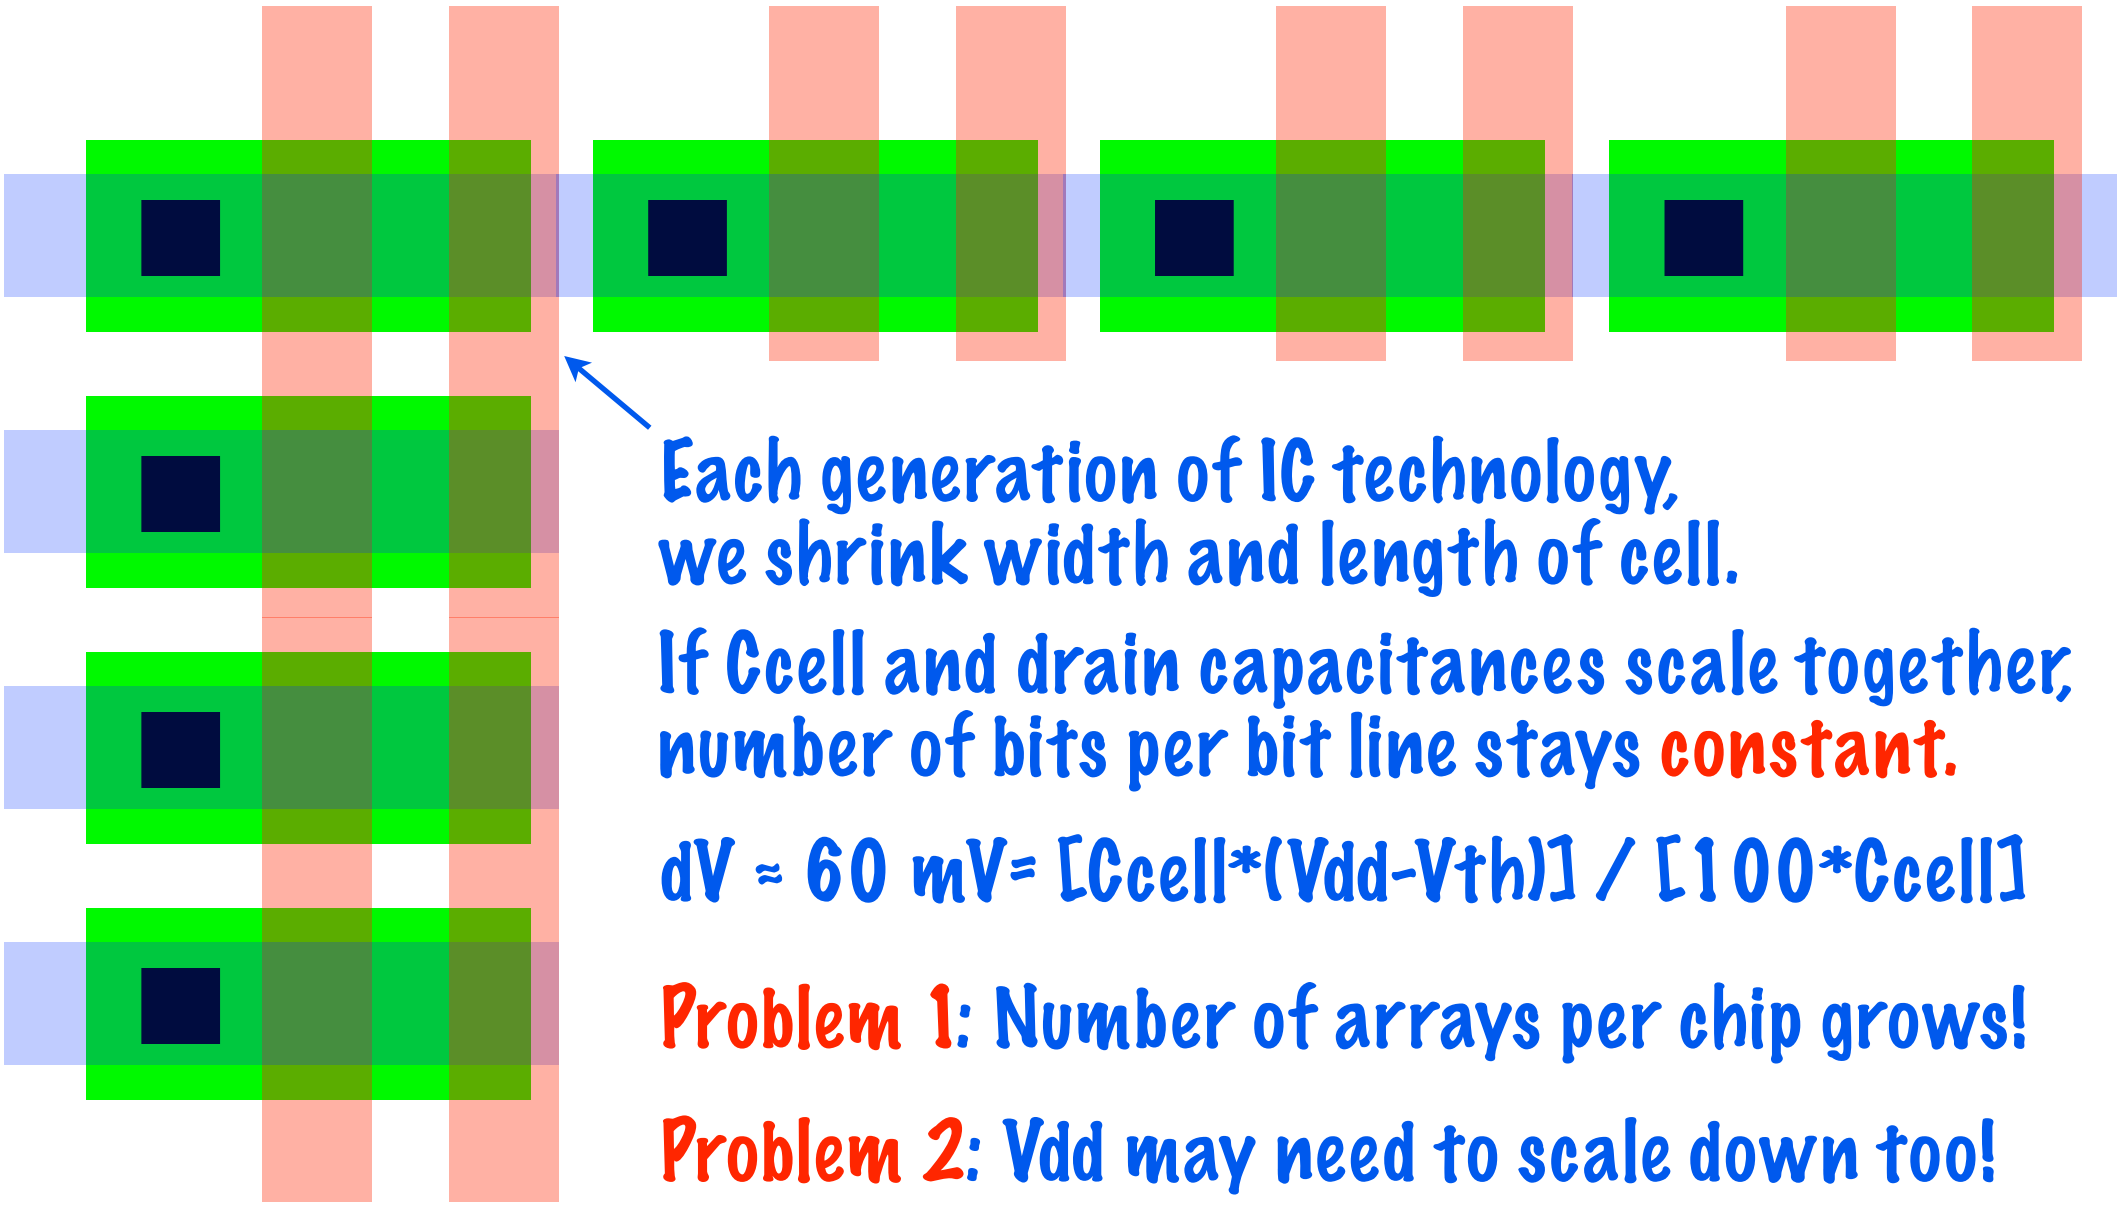

Each generation of IC technology, we shrink width and length of cell.

If 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

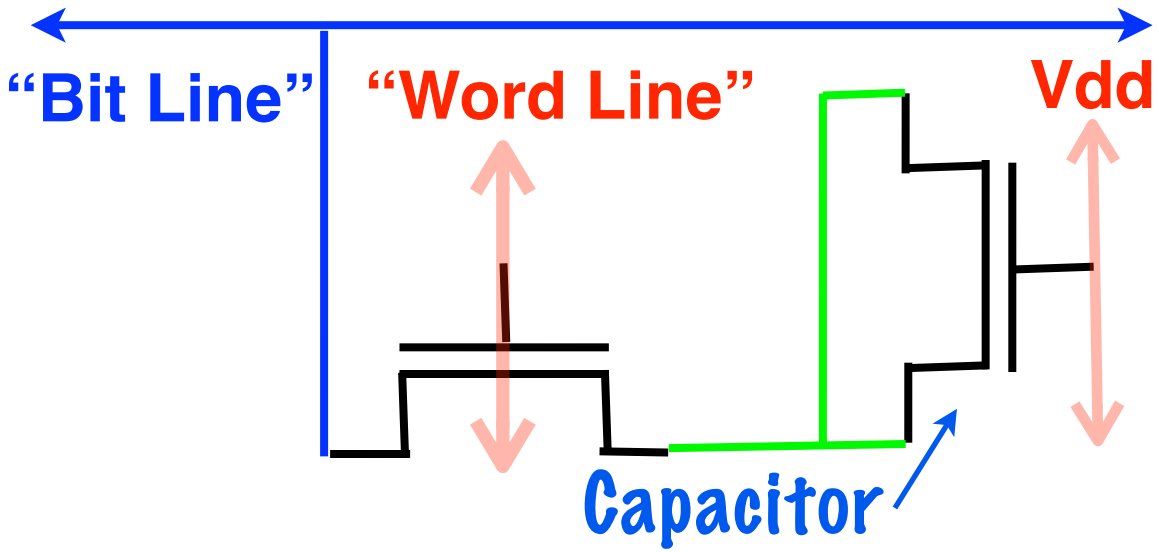

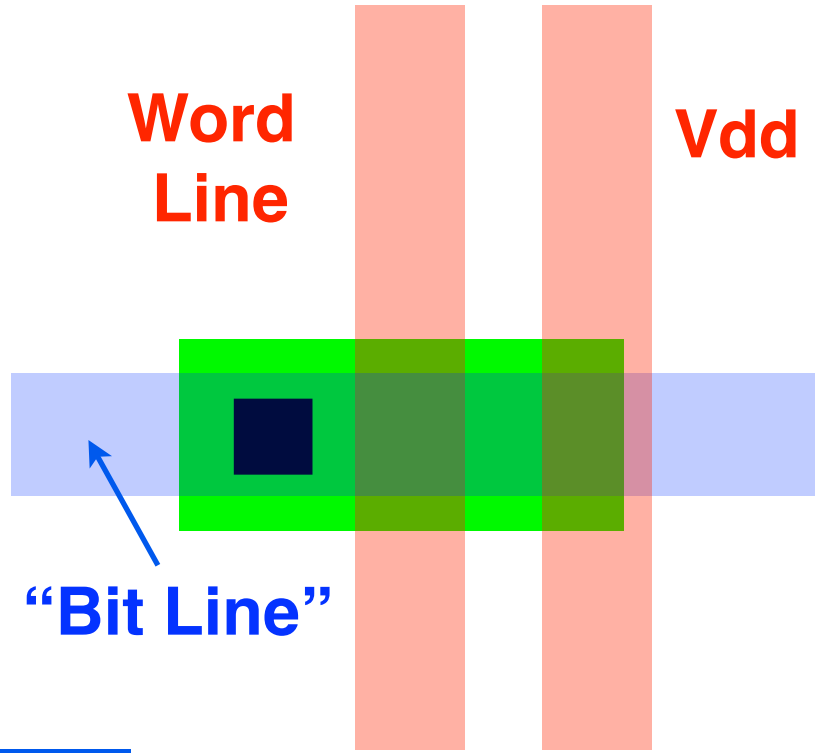

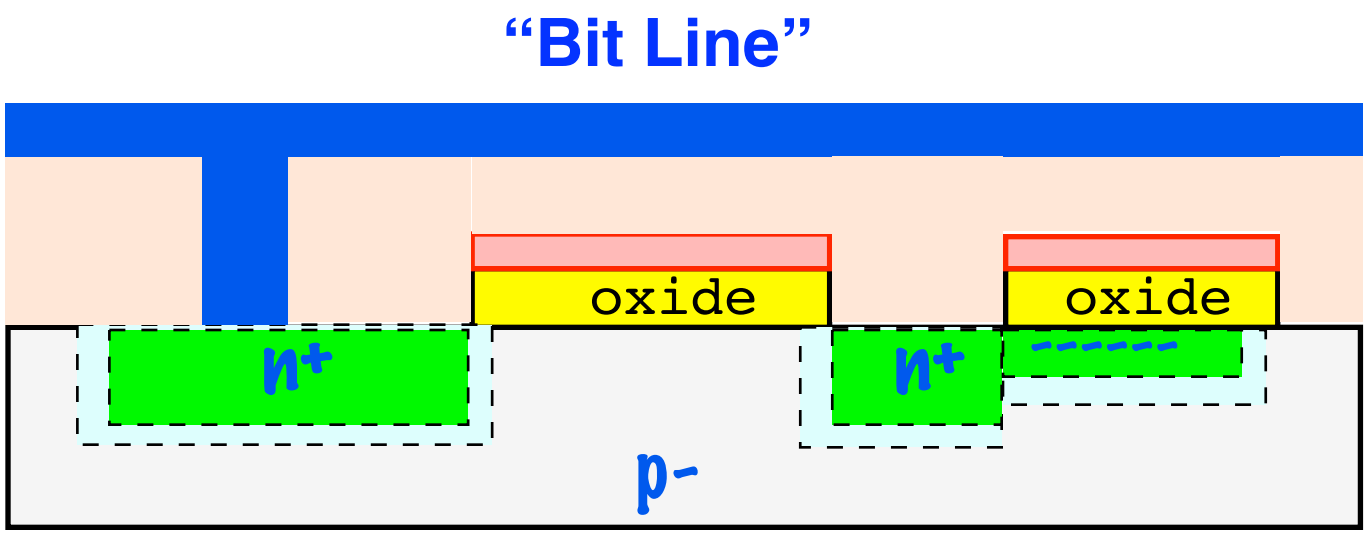

Word Line and Vdd run on "z-axis"

Early replacement: “Trench” capacitors

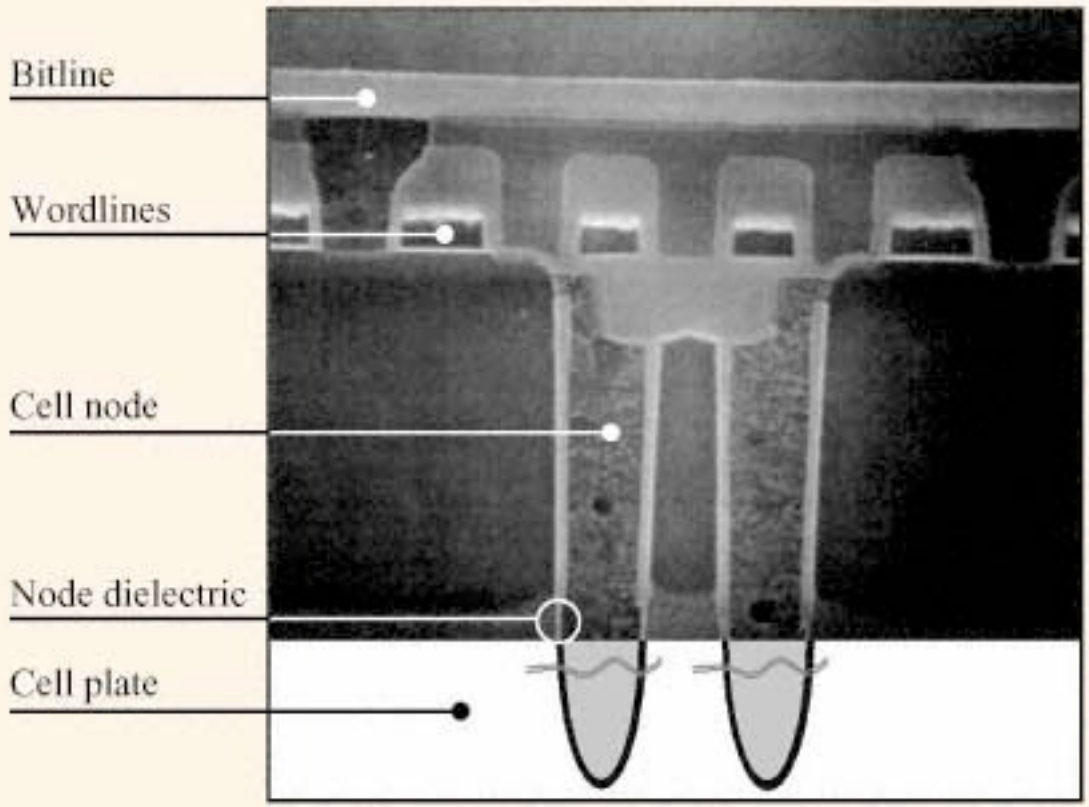

Figure 4

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

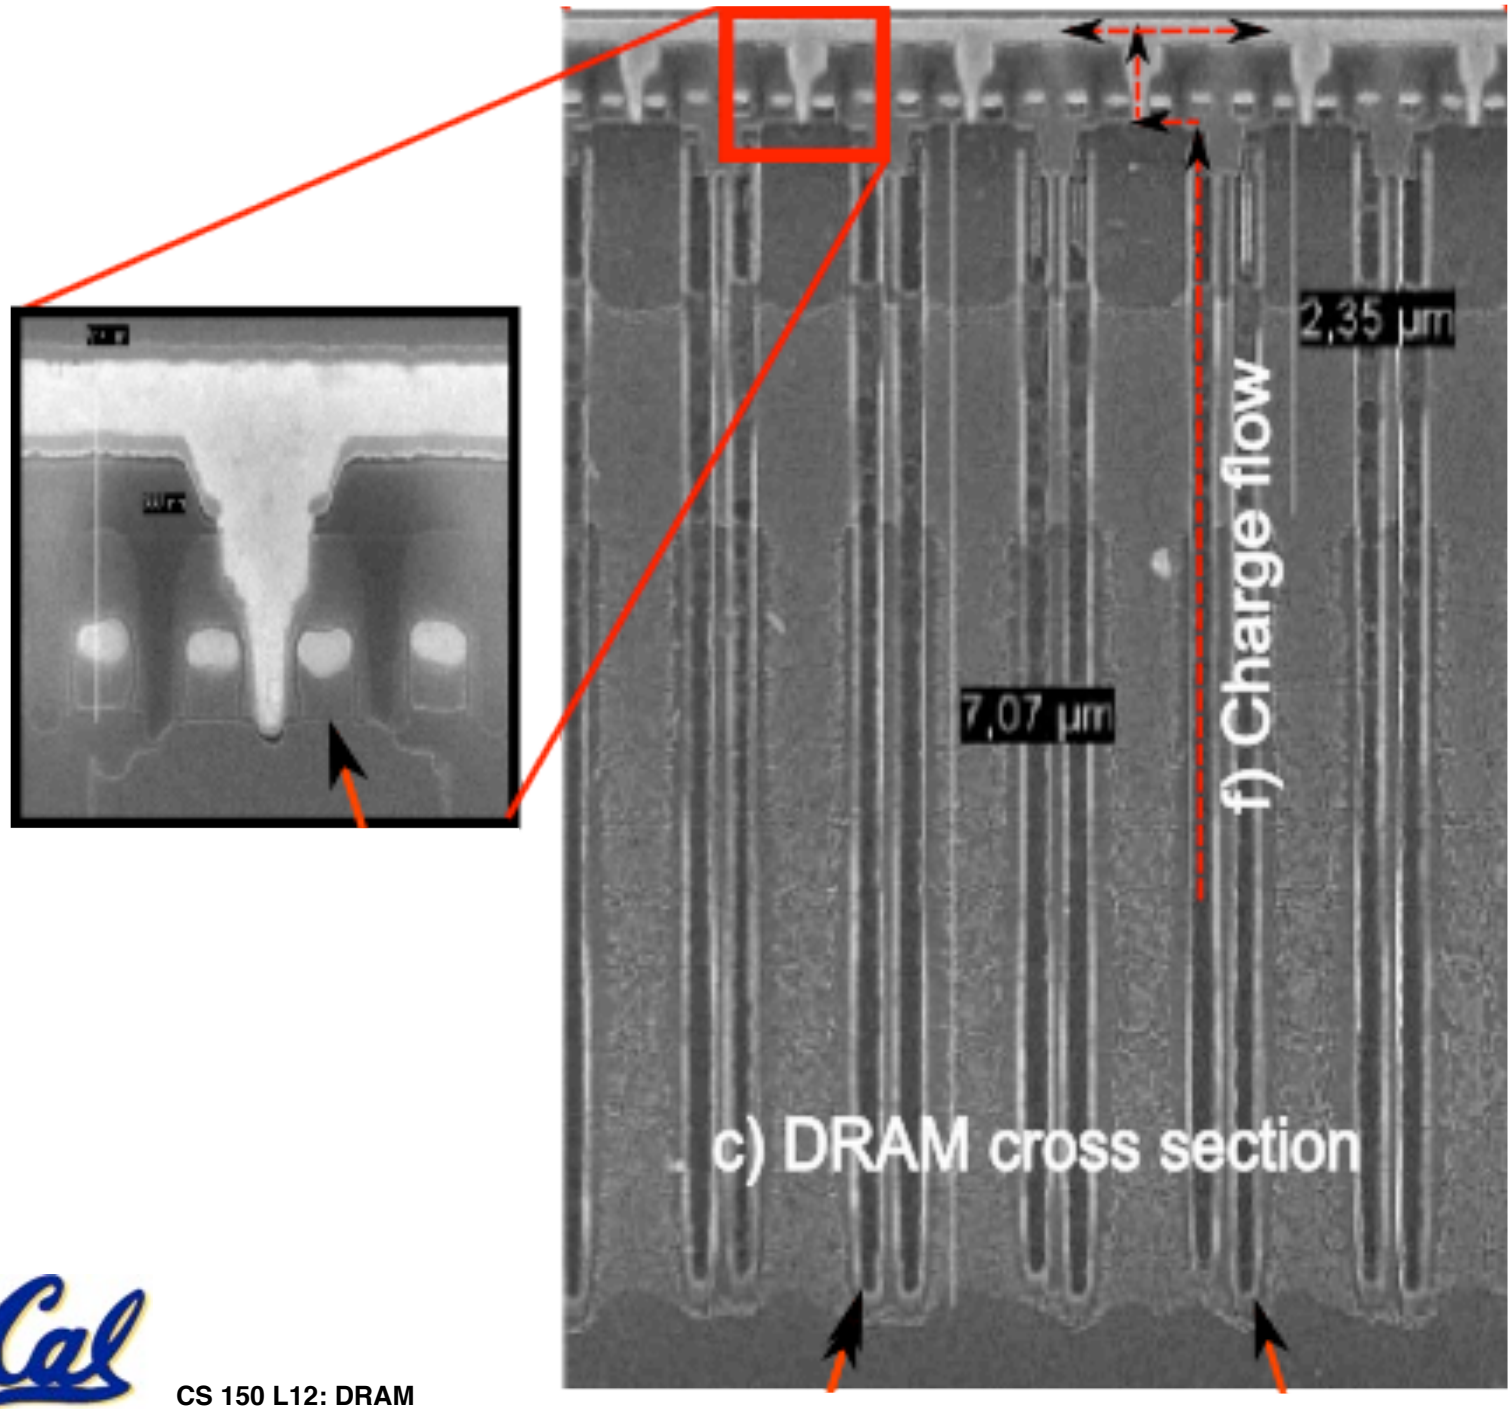

The companies that kept scaling trench capacitors for commodity DRAM chips went out of business.

Modern cells: “stacked” capacitors

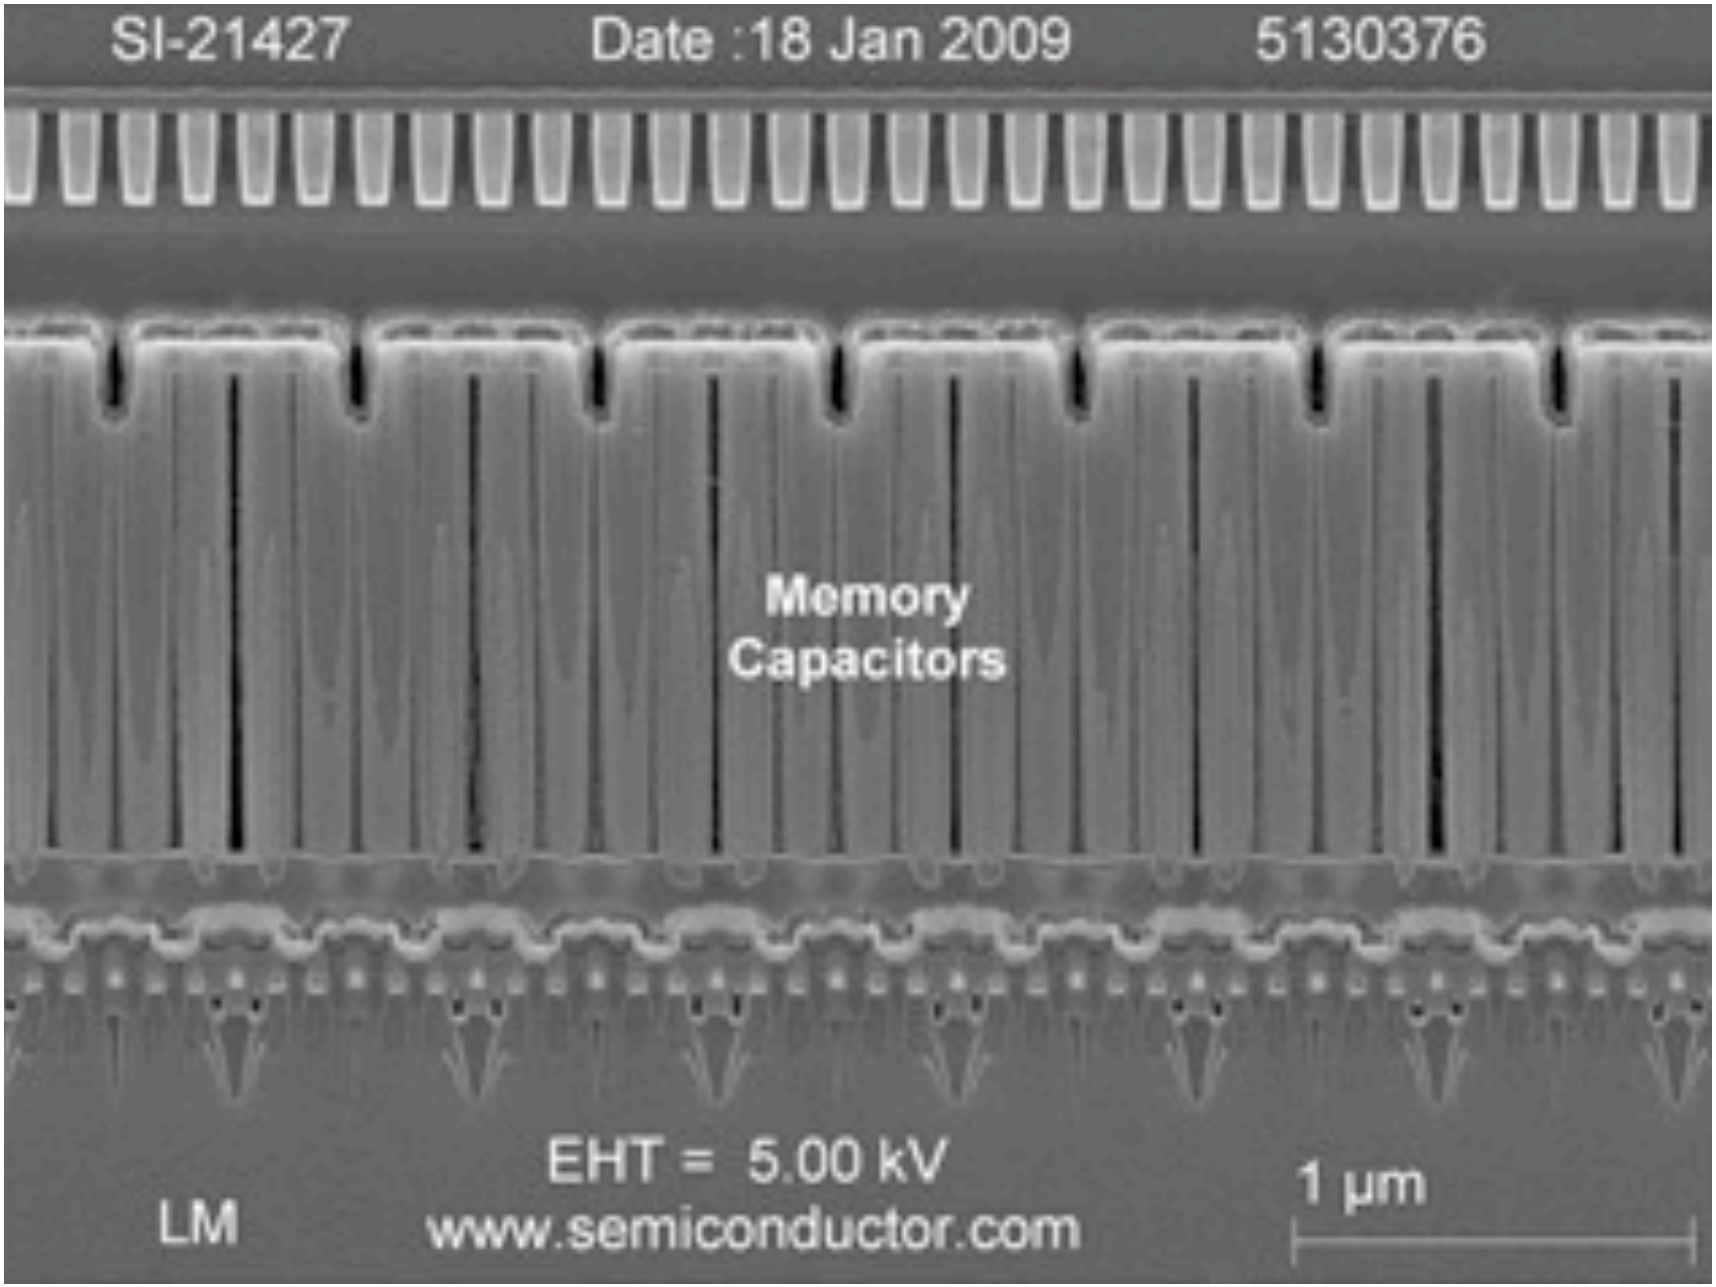

Micron 1-Gbit DDR2 50-nm SDRAM

In the labs: Vertical cell transistors ...

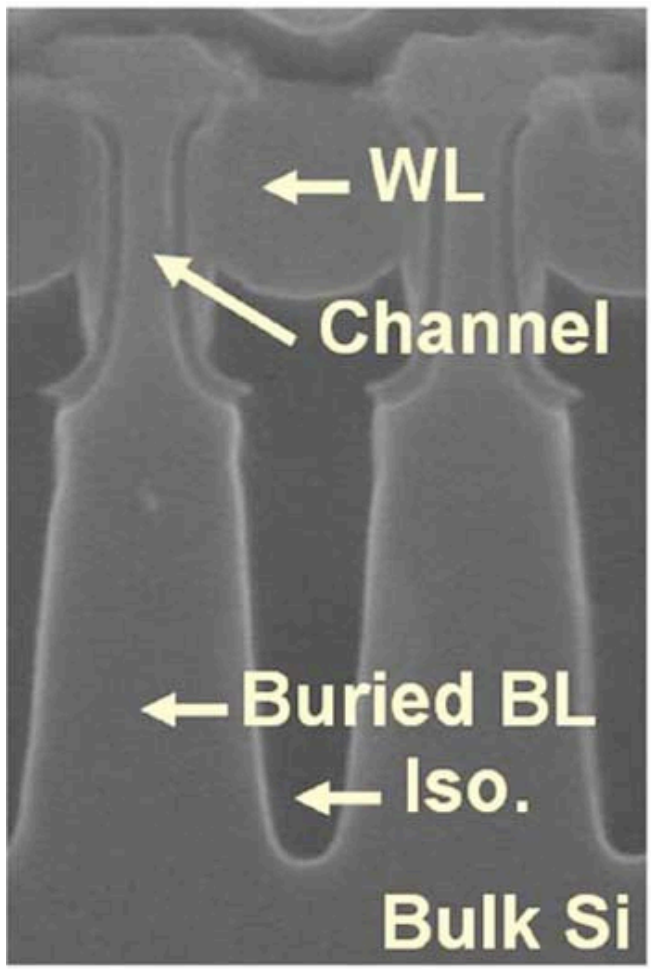

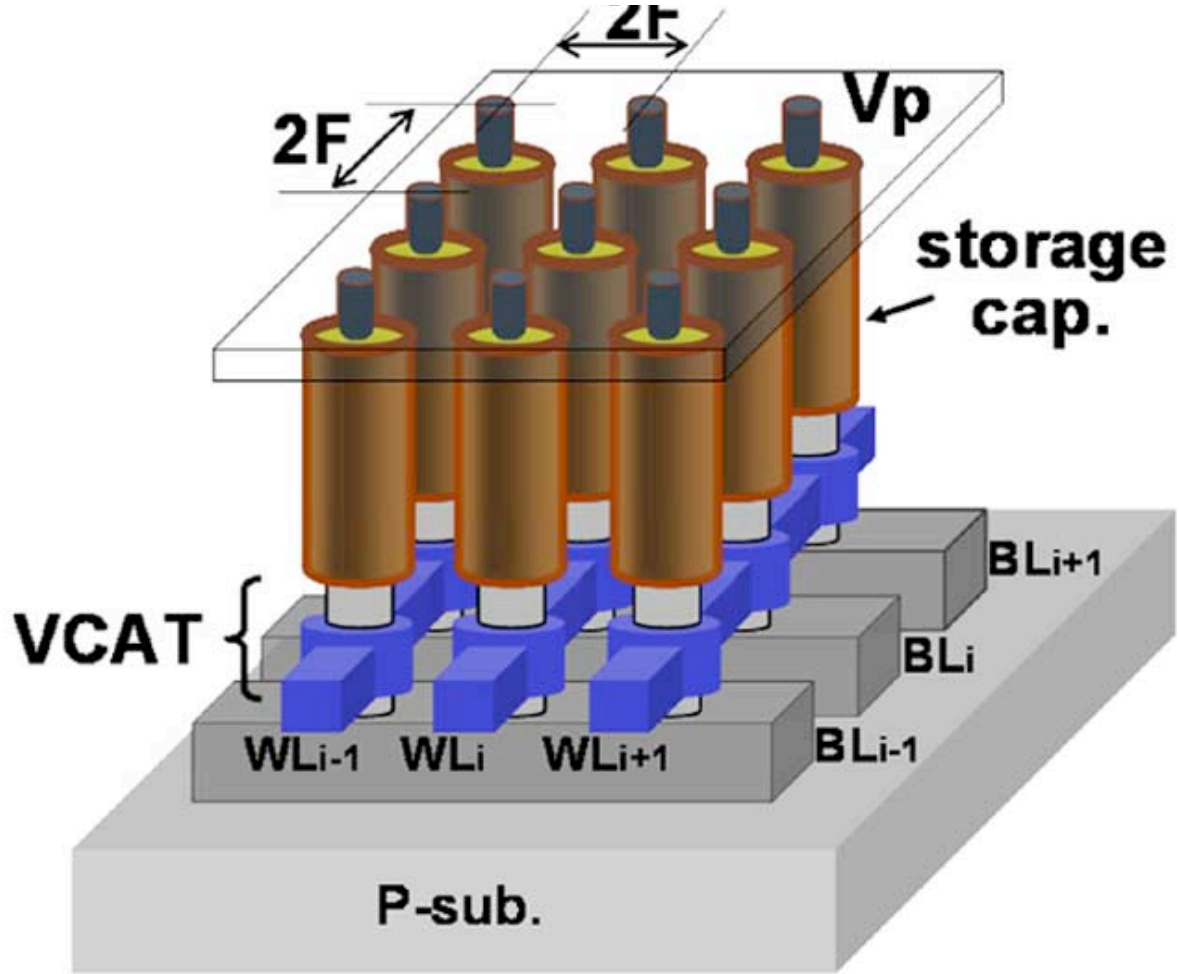

880 IEEE JOURNAL OF SOLID-STATE CIRCUITS, VOL. 45, NO. 4, APRIL 2010

A 31 ns Random Cycle VCAT-Based $4F^2$ DRAM With Manufacturability and Enhanced Cell Efficiency

Ki-Whan Song, Jin-Young Kim, Jae-Man Yoon, Sua Kim, Huijung Kim, Hyun-Woo Chung, Hyungi Kim, Kanguk Kim, Hwan-Wook Park, Hyun Chul Kang, Nam-Kyun Tak, Dukha Park, Woo-Seop Kim, *Member, IEEE*, Yeong-Taek Lee, Yong Chul Oh, Gyo-Young Jin, Jeihwan Yoo, Donggun Park, *Senior Member, IEEE*, Kyungseok Oh, Changhyun Kim, *Senior Member, IEEE*, and Young-Hyun Jun

Micron 50nm 1-Gbit DDR2 die photo

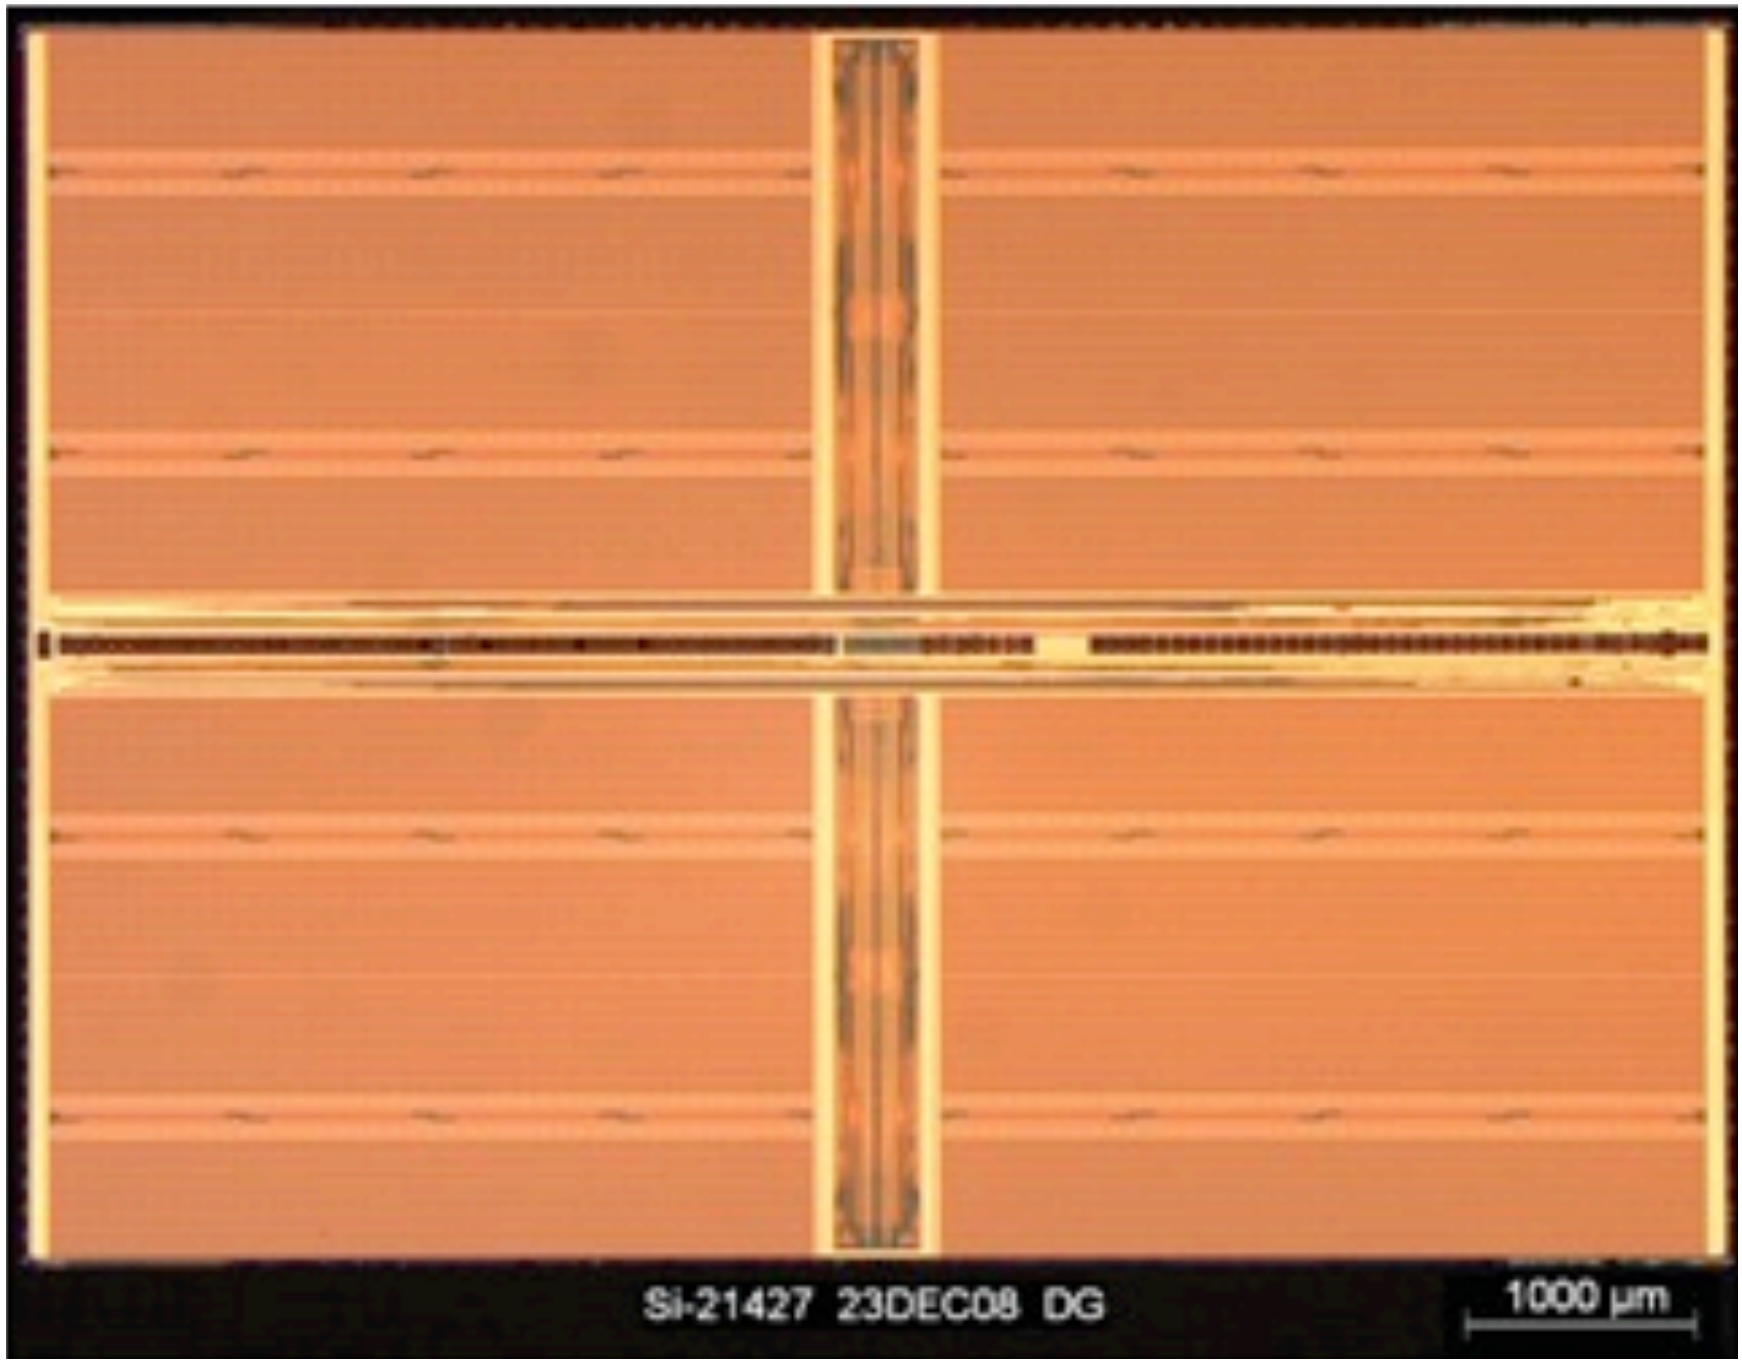

Today's Lecture: DRAM

- DRAM, Xilinx, and You
- DRAM: Bottom-up
- **DRAM: Top-down**

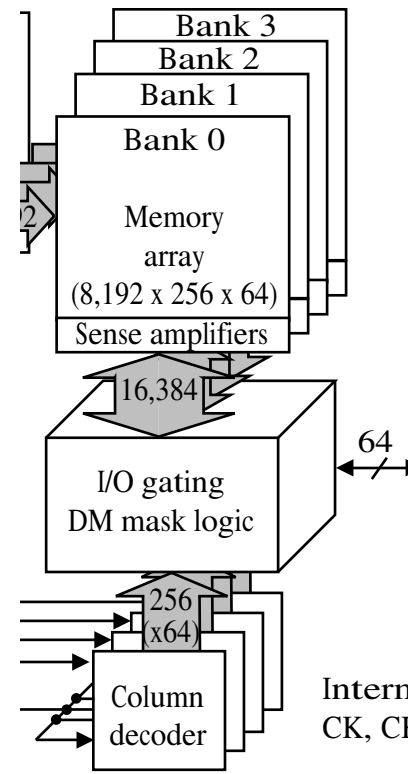

Memory Arrays

512Mb: x4, x8, x16 DDR2 SDRAM
Features

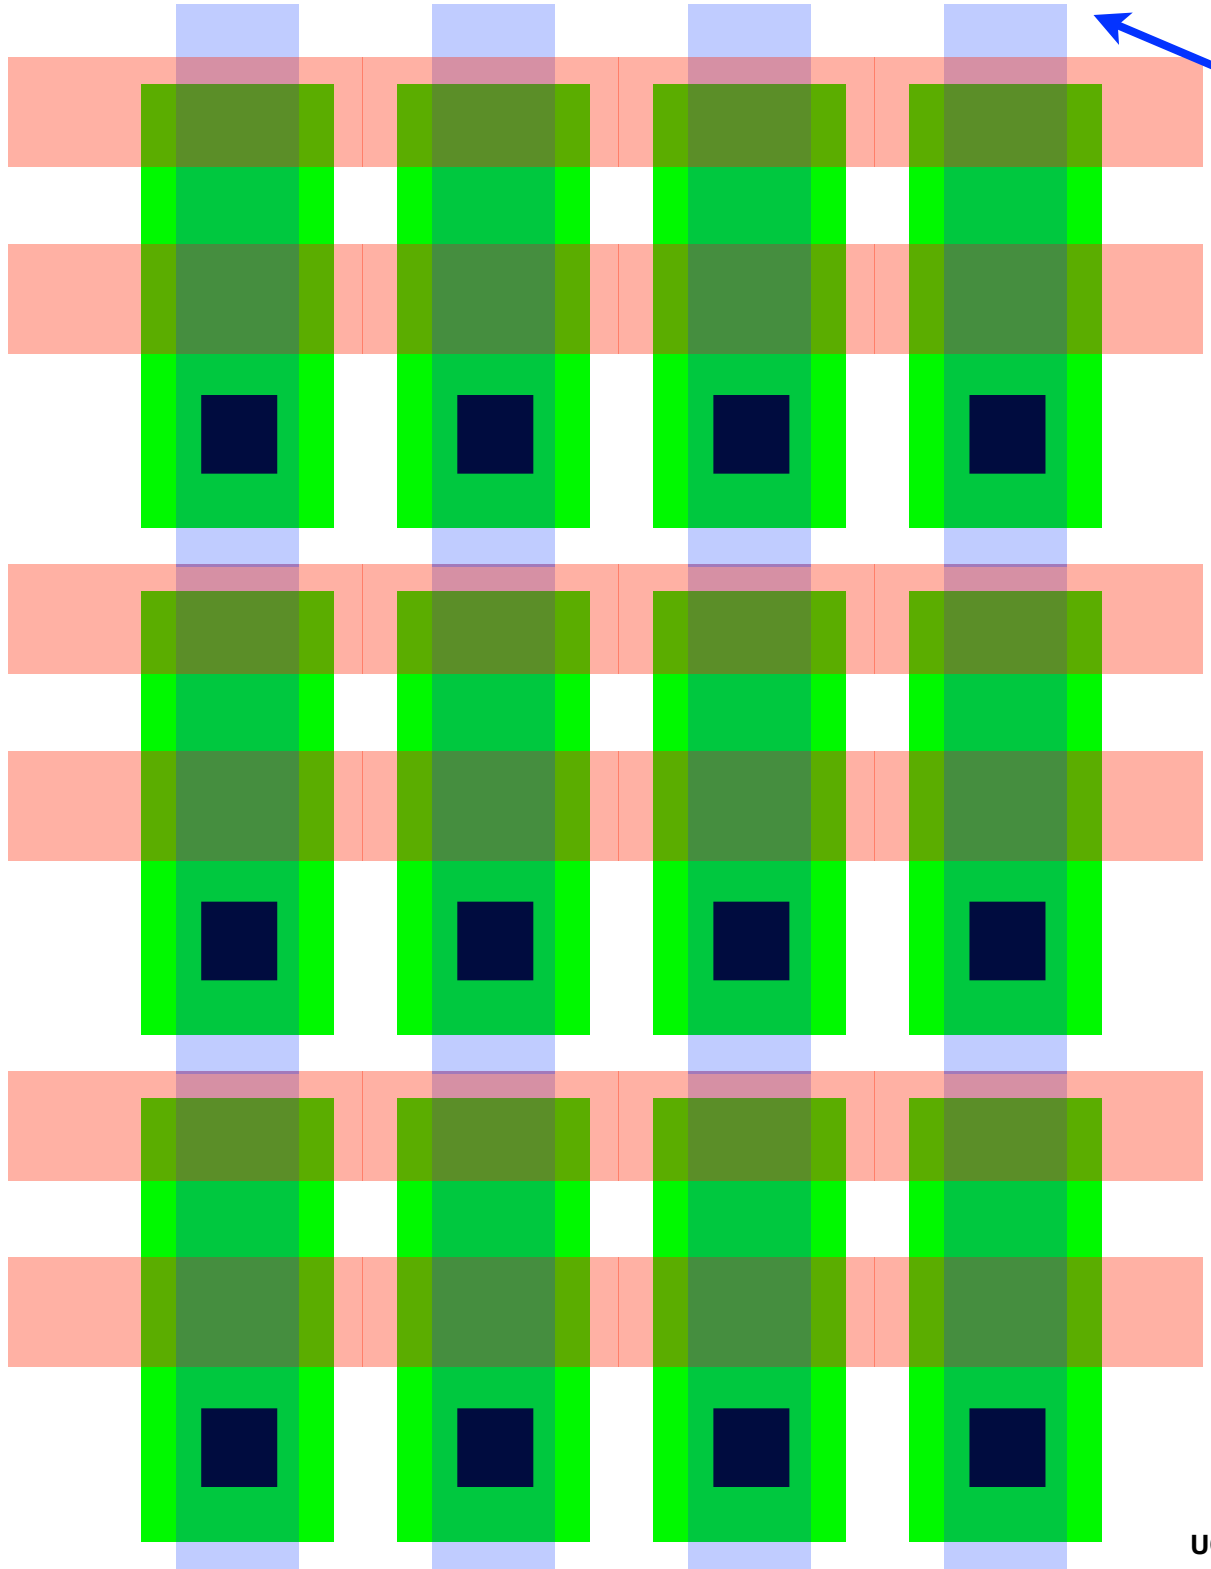

People buy DRAM for the bits. "Edge" circuits are overhead.

So, we amortize the edge circuits over big arrays.

A "bank" of 128 Mb (512Mb chip -> 4 banks)

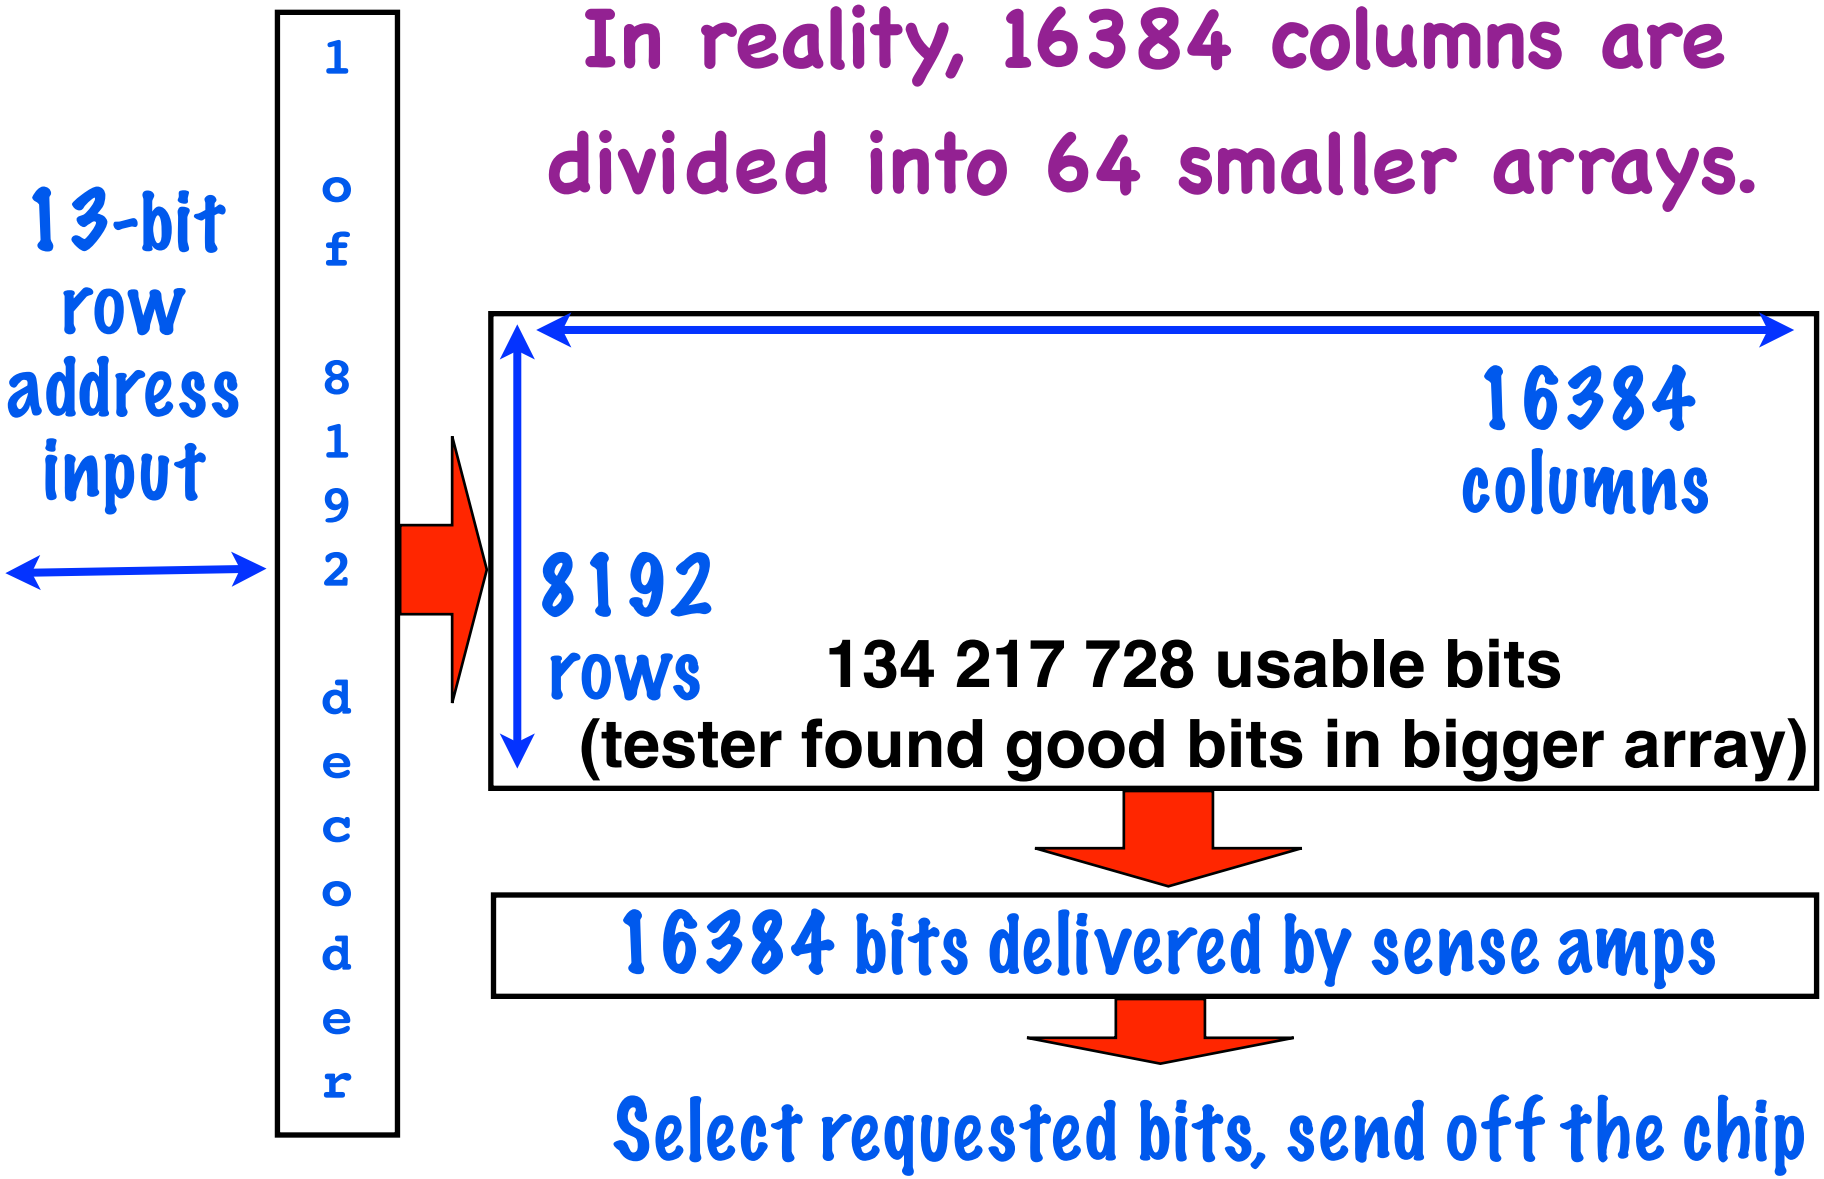

Recall DRAM Challenge #3b: Sensing

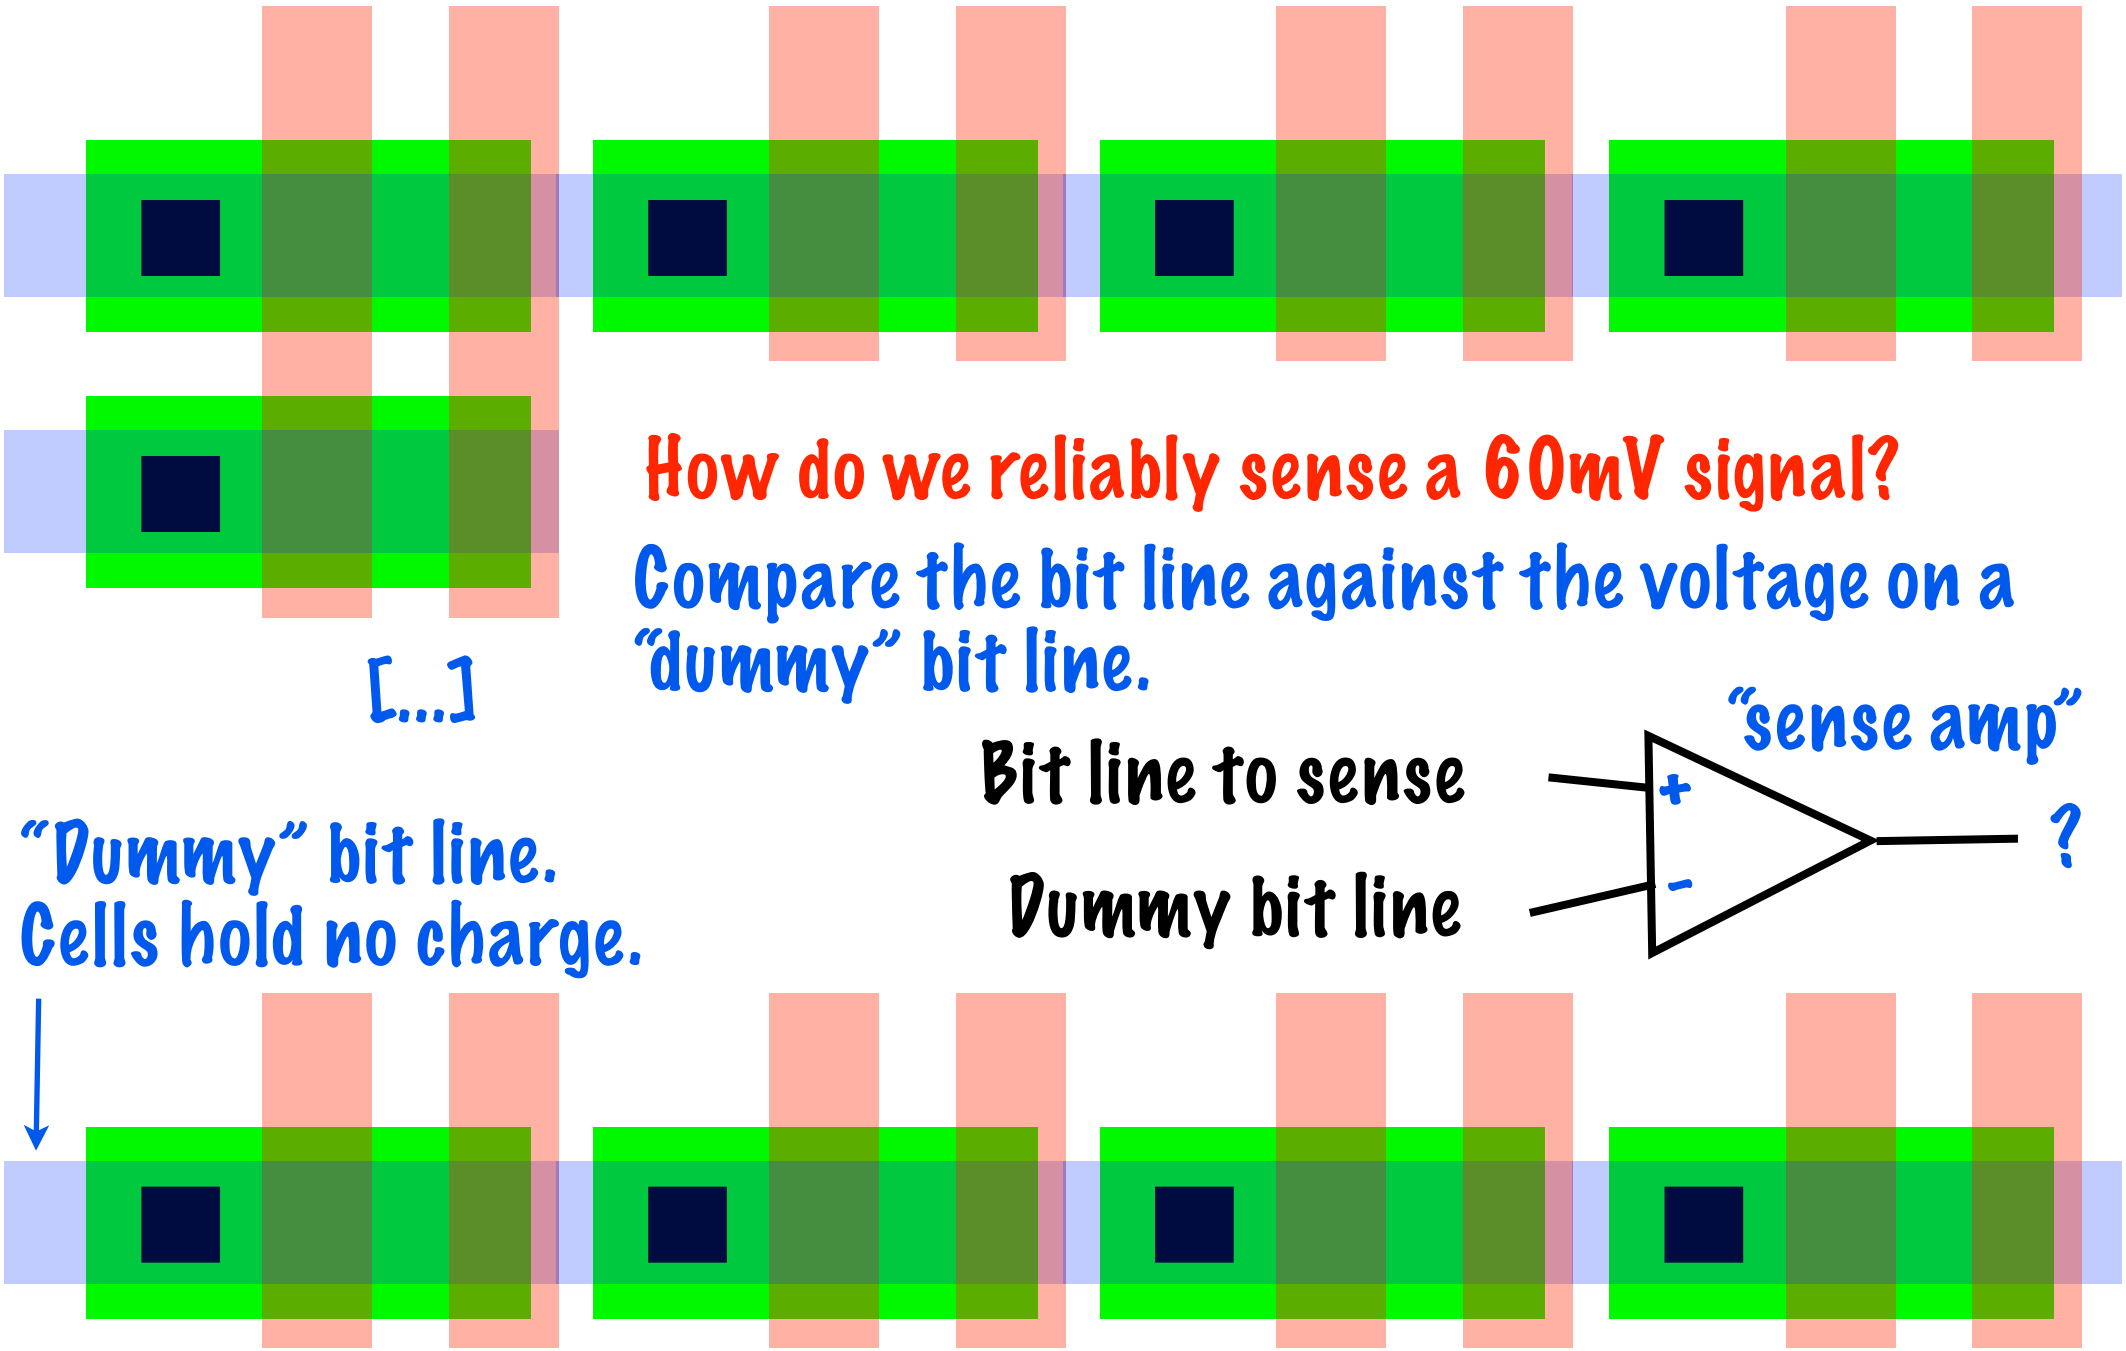

How do we reliably sense a 60mV signal?

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 2.5ns period DRAM (400 MT/s) can do row reads at only 55 ns (18 MHz).

13-bit row address input

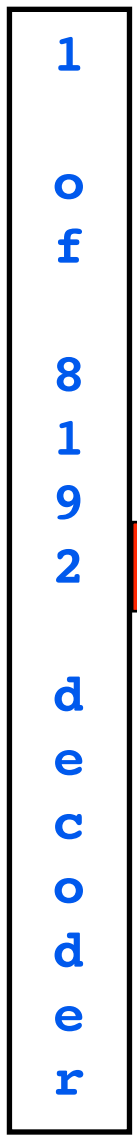

DRAM has high latency to first bit out. A fact of life.

16384 columns

8192 rows

134 217 728 usable bits
(tester found good bits in bigger array)

16384 bits delivered by sense amps

Select requested bits, send off the chip

An ill-timed refresh may add to latency

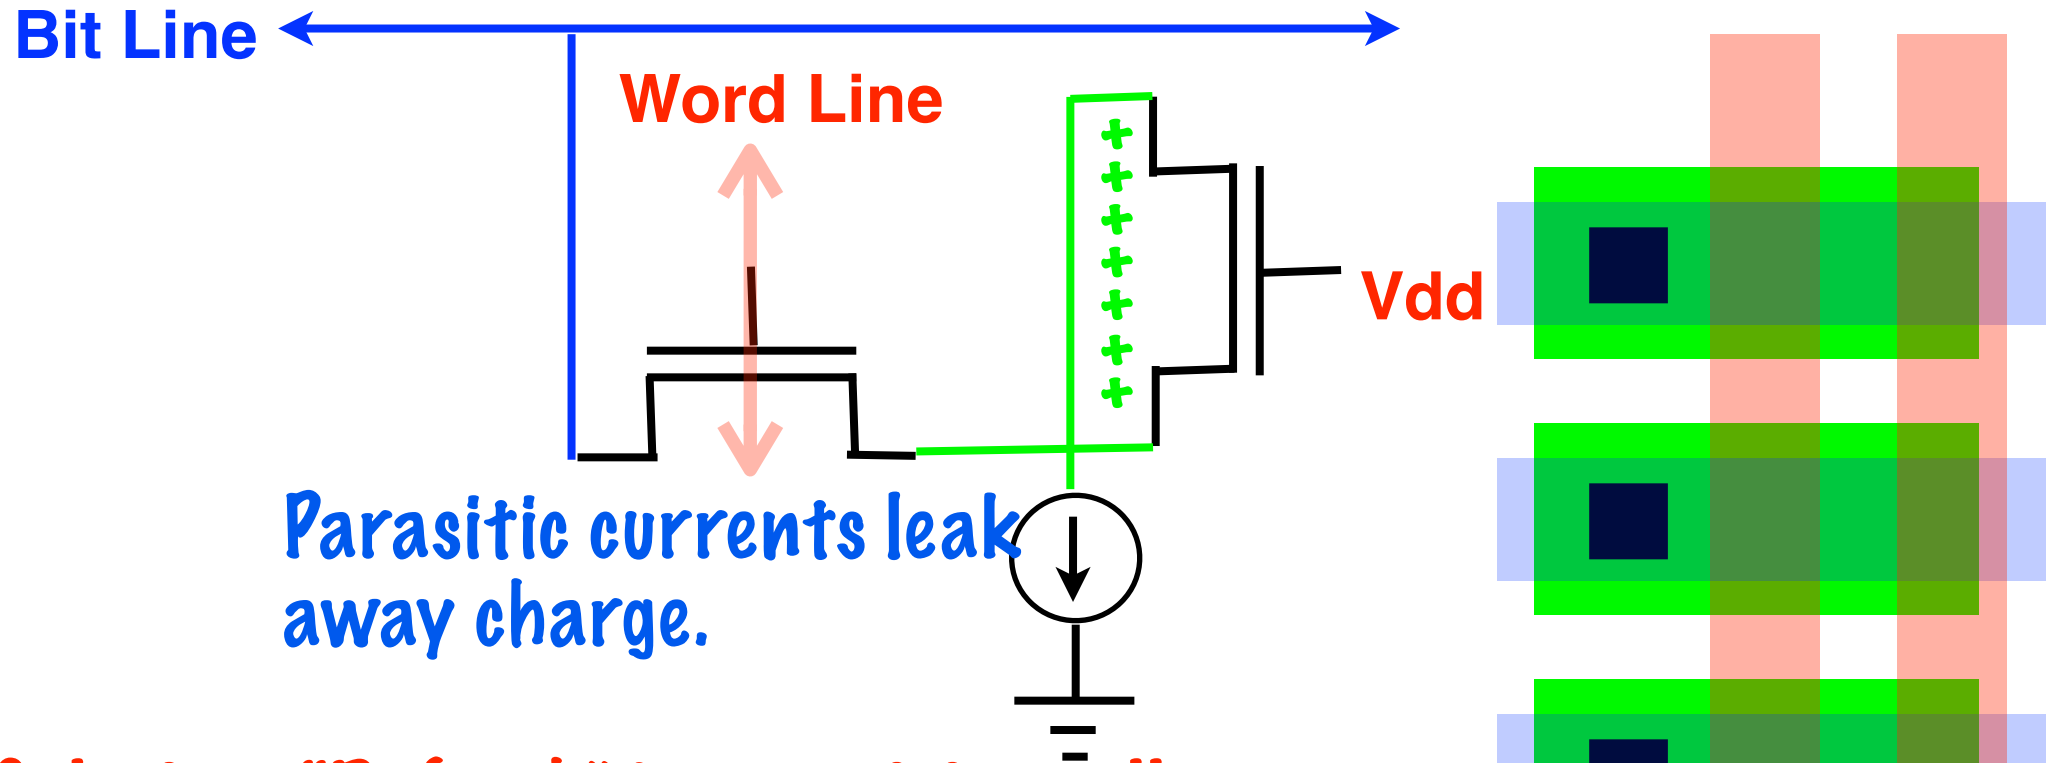

Solution: "Refresh", by rewriting cells at regular intervals (tens of milliseconds)

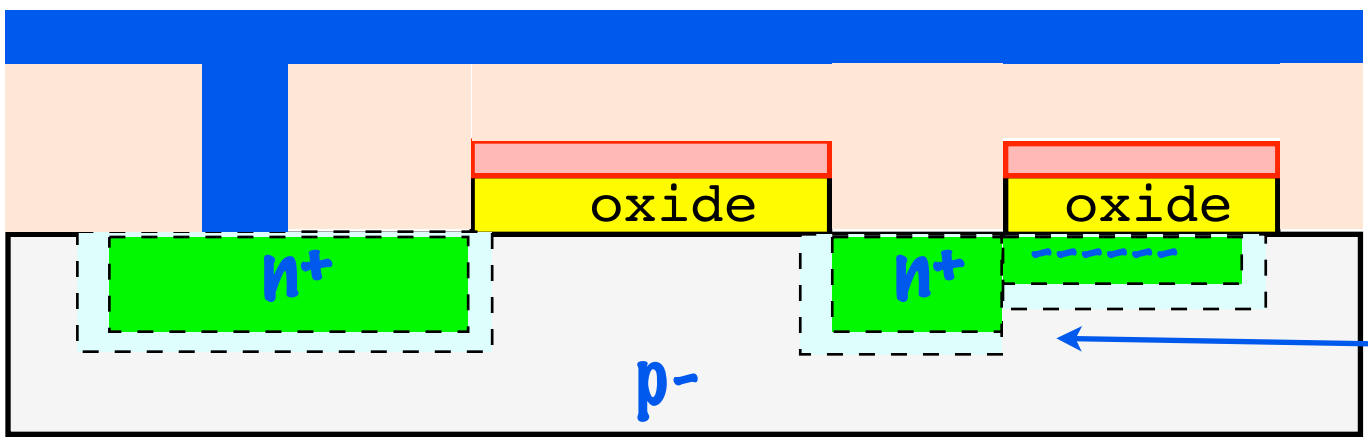

Diode leakage ...

Latency is not the same as bandwidth!

Thus, push to faster DRAM interfaces

What if we want all of the 16384 bits?

In row access time (55 ns) we can do 22 transfers at 400 MT/s.

16-bit chip bus -> 22 x 16 = 352 bits << 16384

Now the row access time looks fast!

13-bit row address input

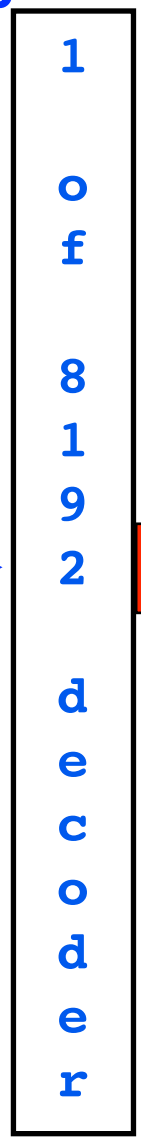

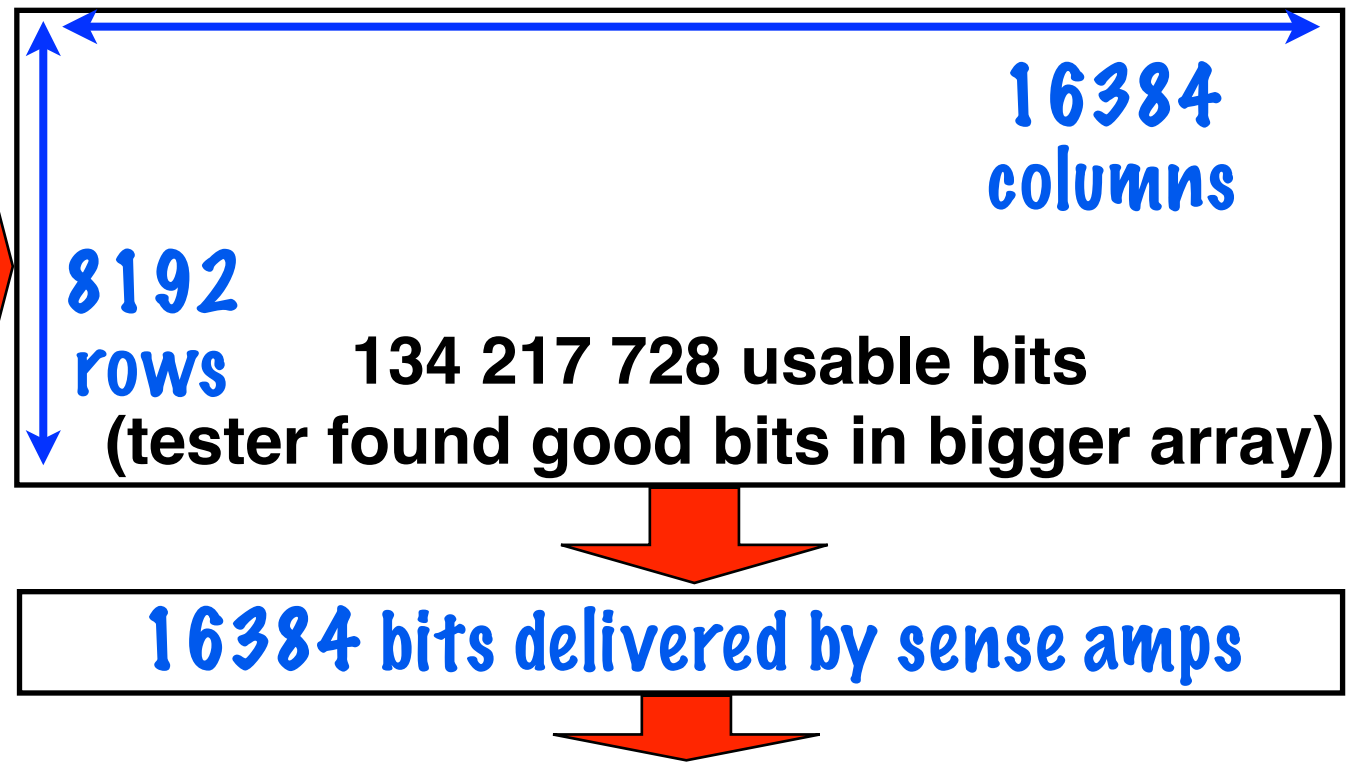

Select requested bits, send off the chip

Sadly, it's rarely this good ...

What if we want all of the 16384 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.

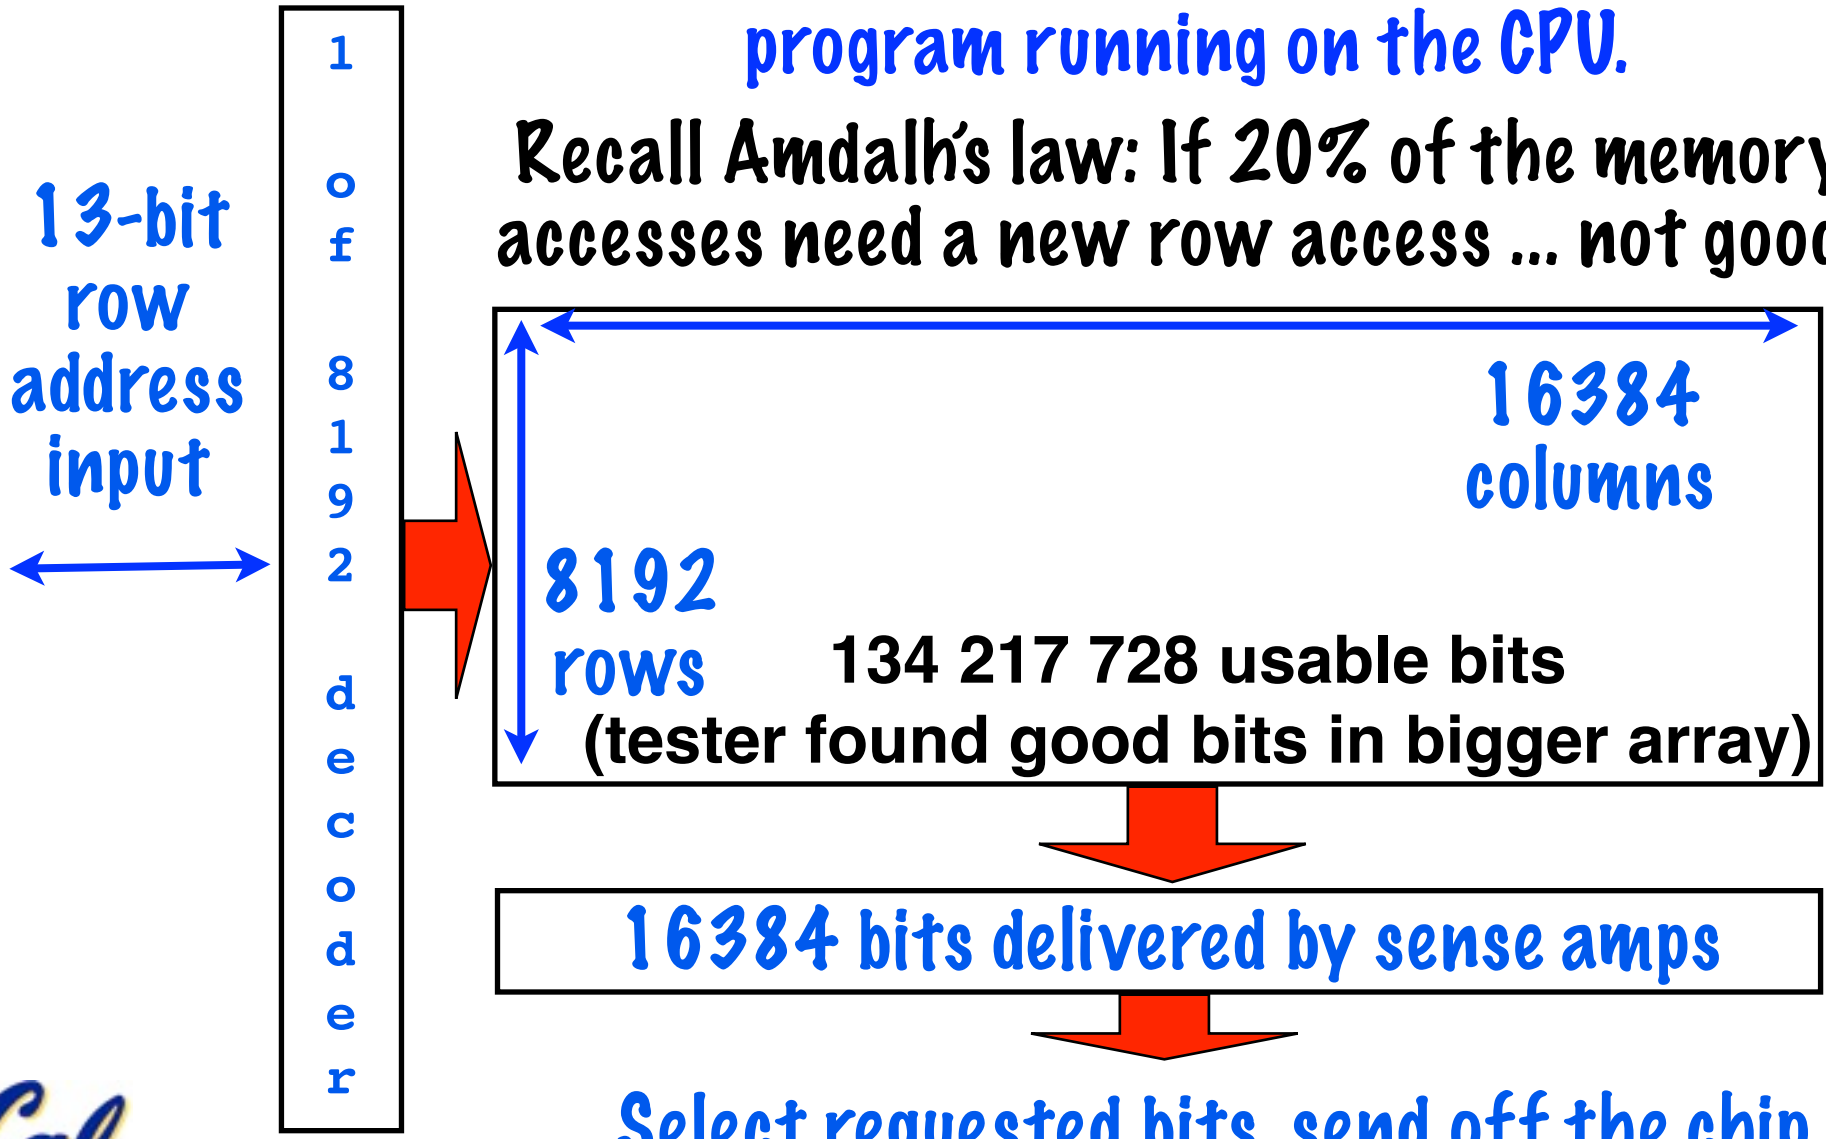

DRAM latency/bandwidth chip features

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

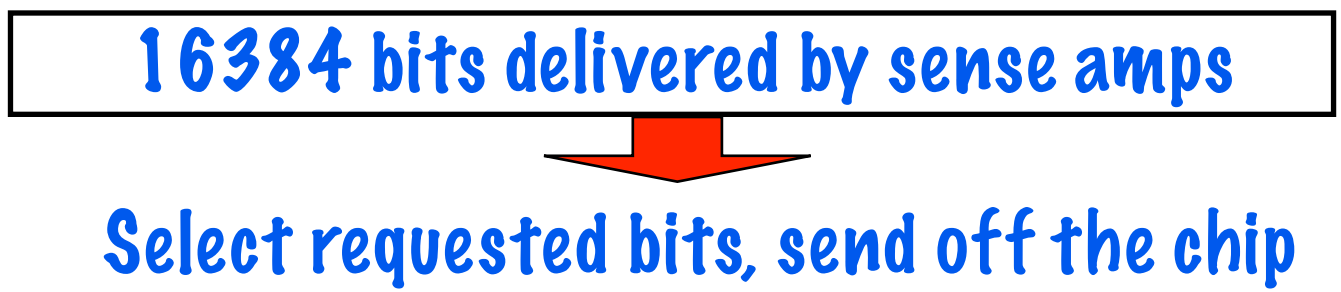

- **Interleaving:** Design the right interface to the 4 memory banks on the chip, so several row requests run in parallel.

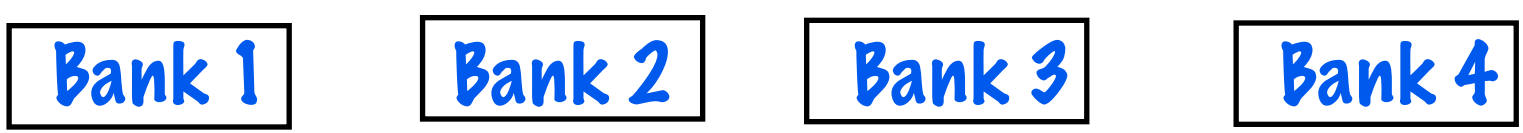

Off-chip interface for the Micron part ...

A clocked bus: 200 MHz clock, data transfers on both edges (DDR).

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

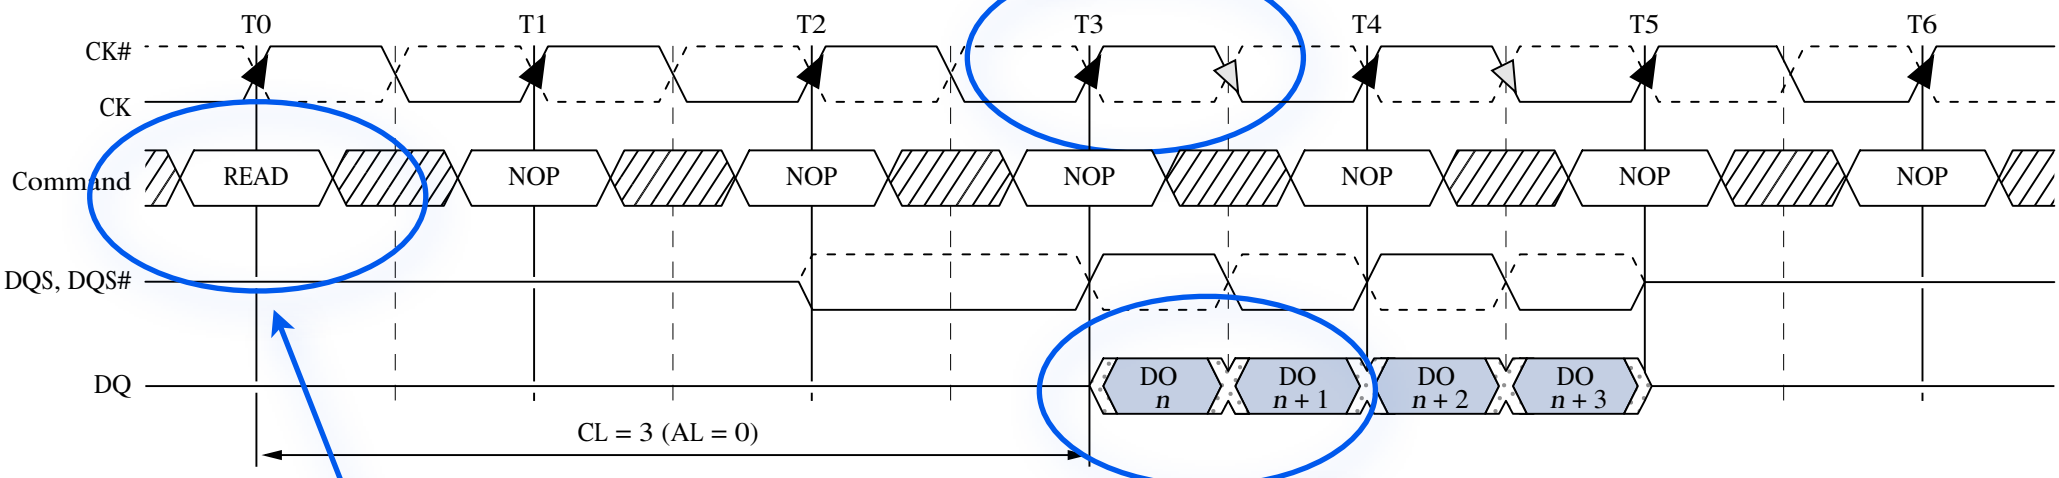

DRAM is controlled via commands (READ, WRITE, REFRESH, ...)

Synchronous data output.

Opening a row before reading ...

Auto-Precharge READ

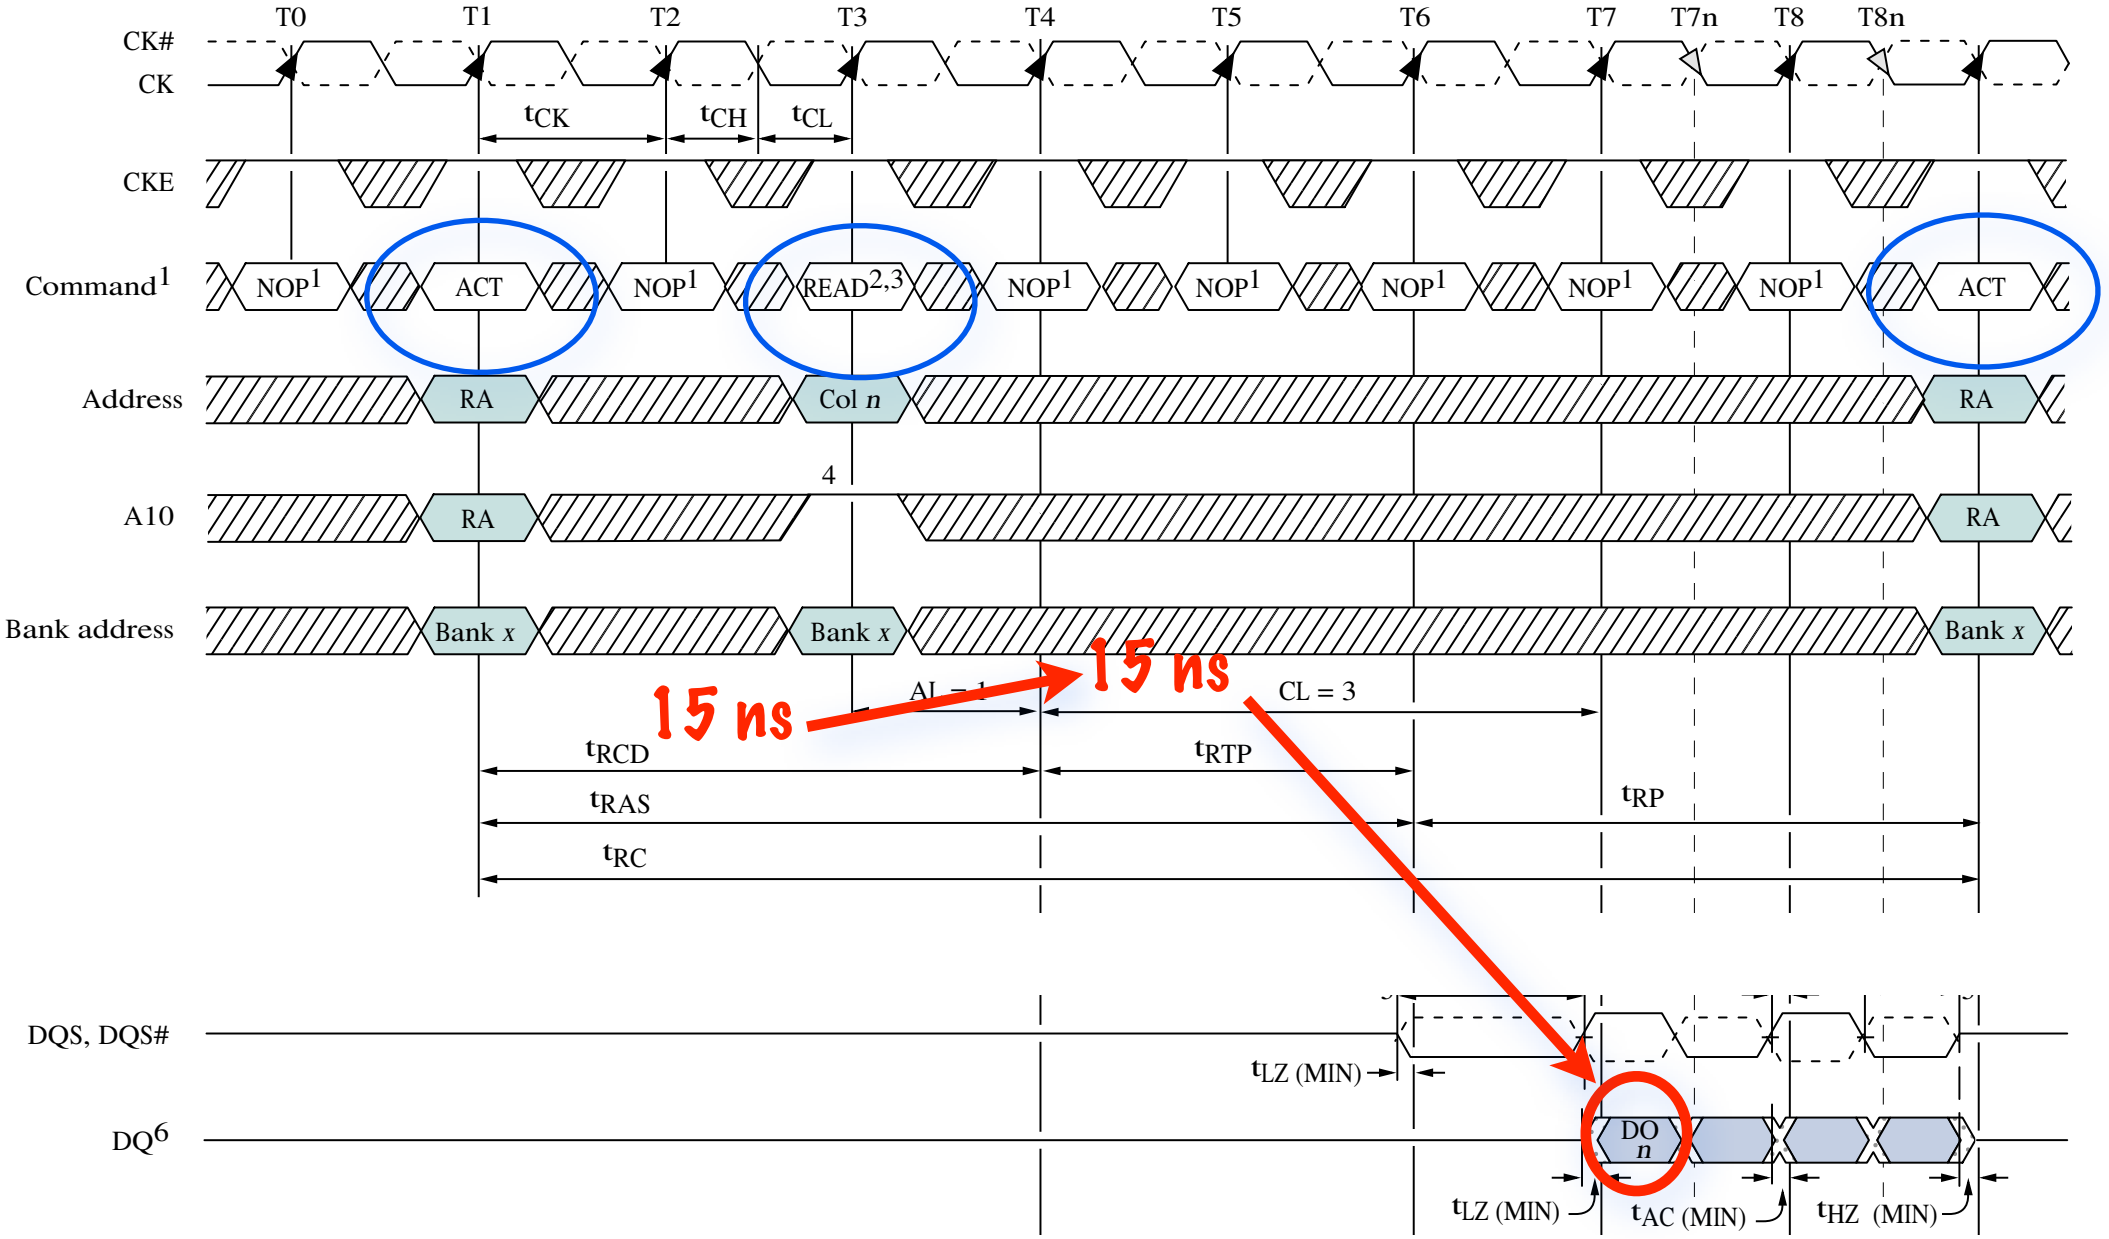

55 ns between row opens.

However, we can read columns quickly

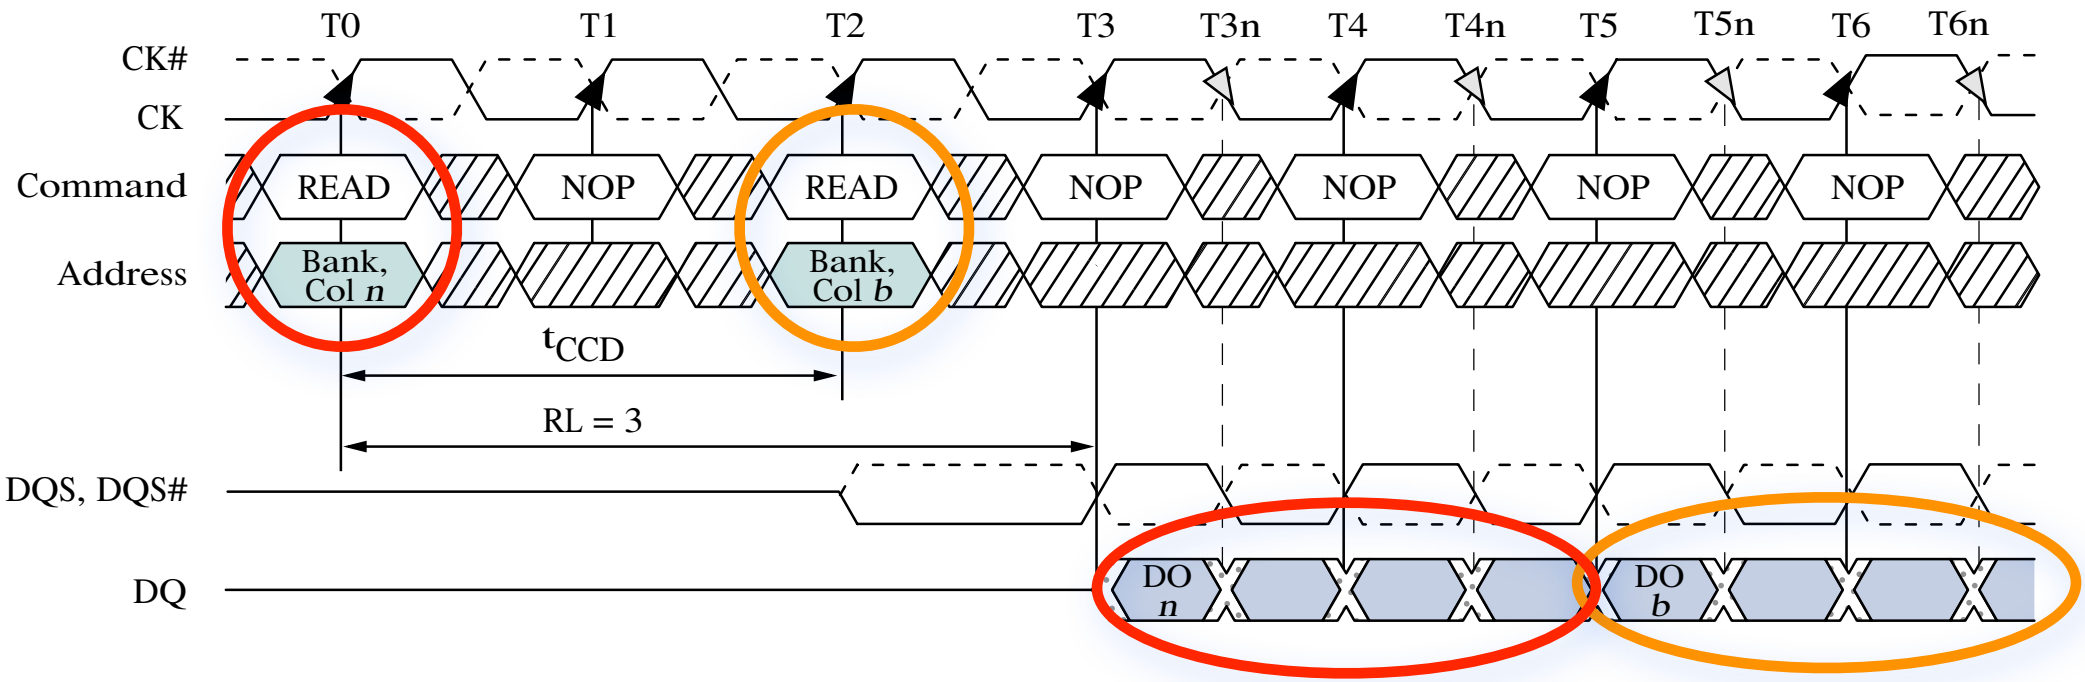

Note: This is a "normal read" (not Auto-Precharge).
Both READs are to the same bank, but different columns.

Why can we read columns quickly?

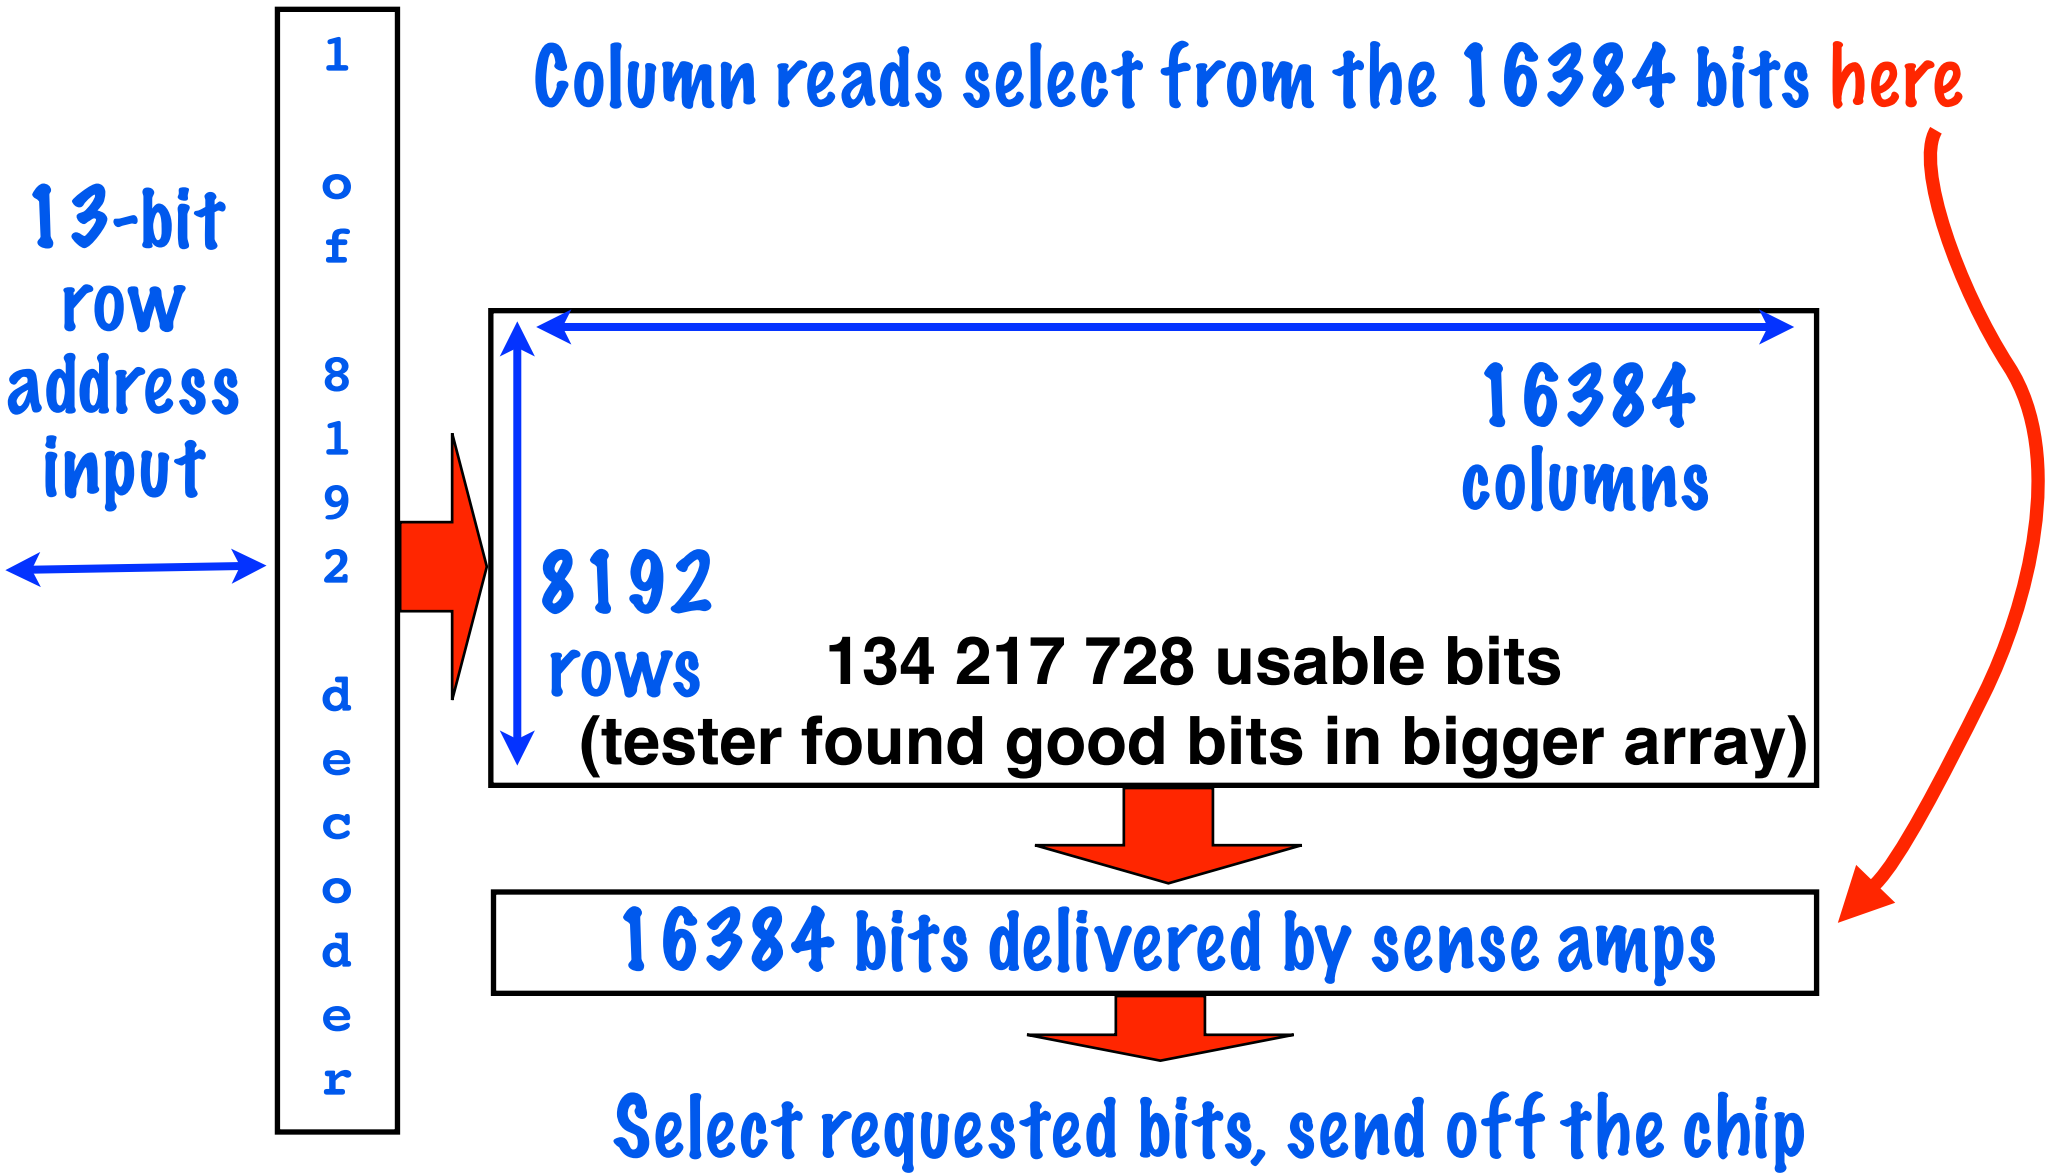

Interleave: Access all 4 banks in parallel

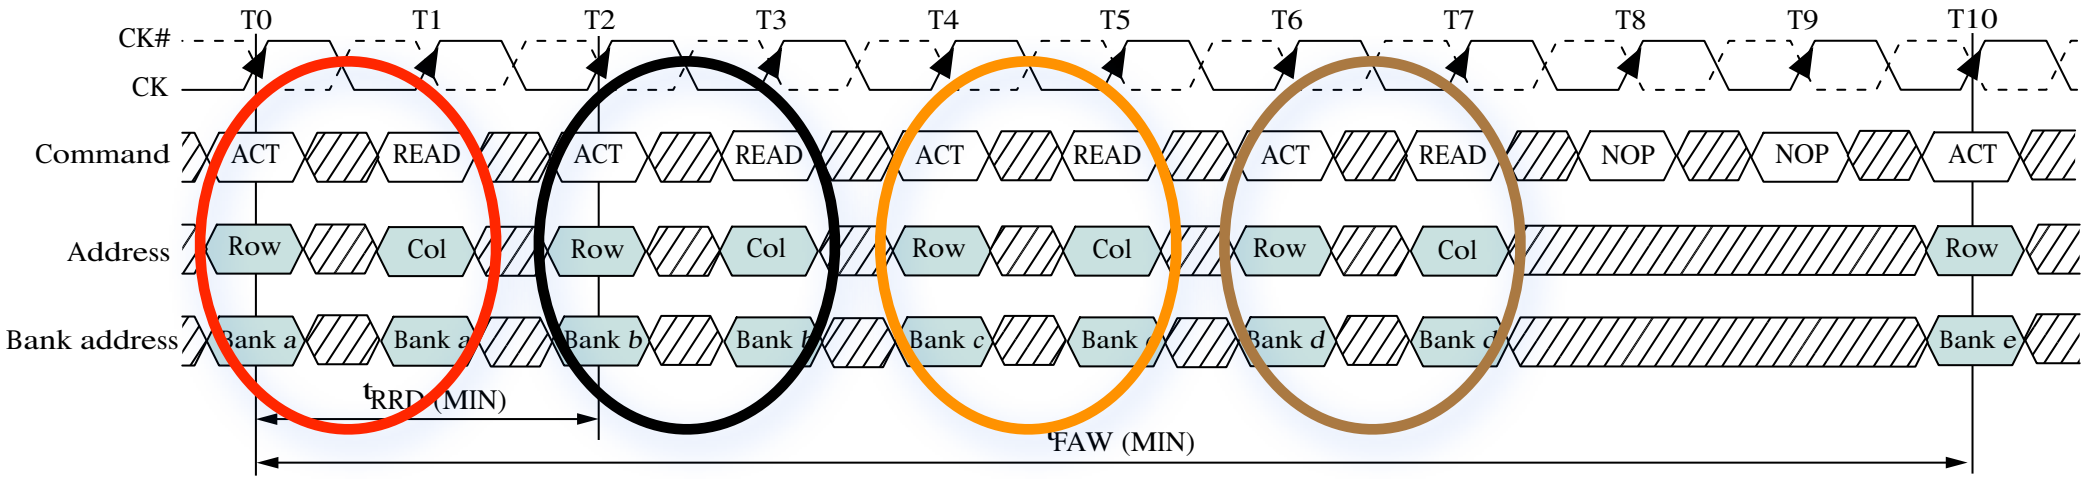

Interleaving: Design the right interface to the 4 memory banks on the chip, so several row requests run in parallel.

Bank a | Bank b | Bank c | Bank d

Can also do other commands on banks concurrently.

Only part of a bigger story ...

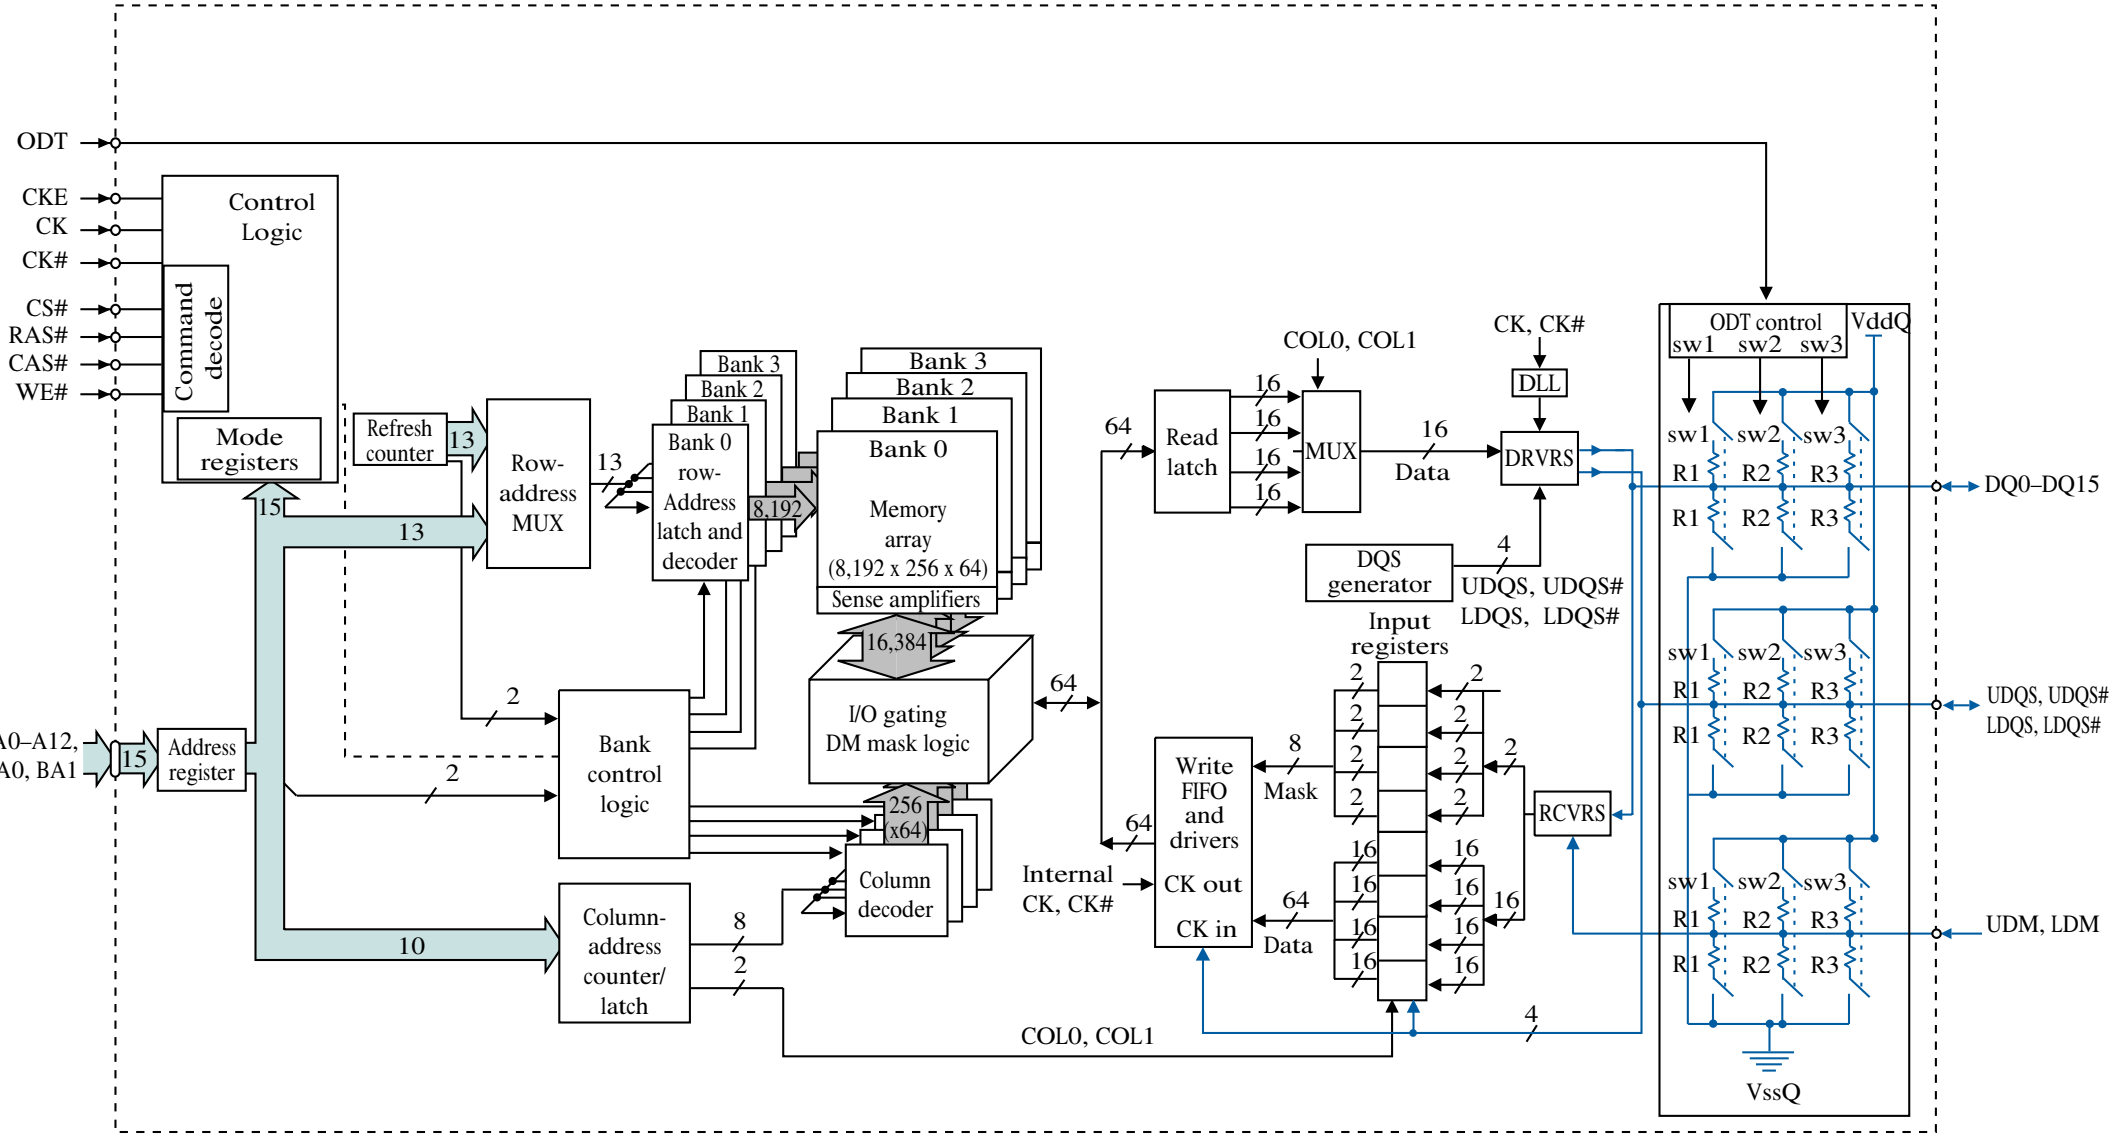

Only part of a bigger story ...

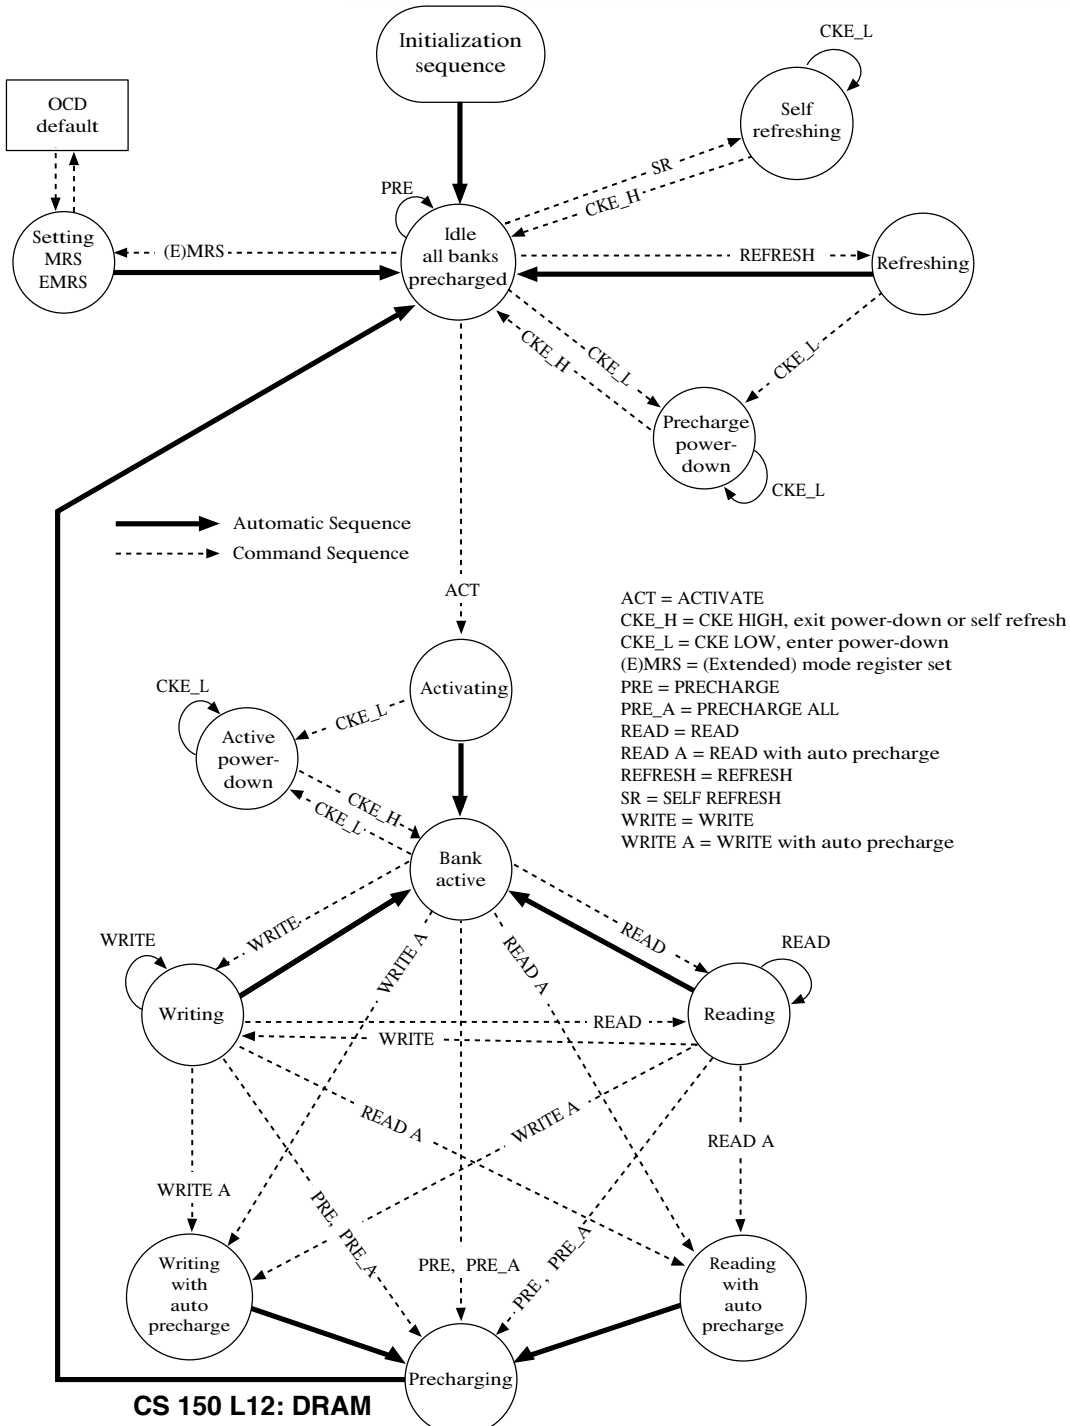

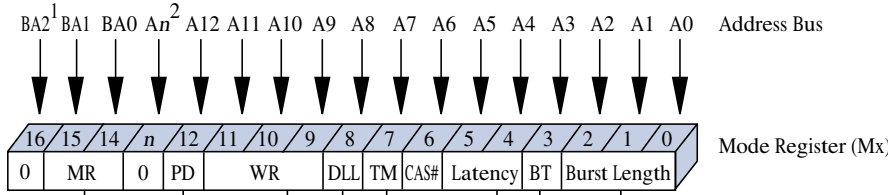

M12	PD Mode
0	Fast exit (normal)
1	Slow exit (low power)

M7	Mode
0	Normal
1	Test

M8	DLL Reset
0	No
1	Yes

M2	M1	M0	Burst Length
0	0	0	Reserved
0	0	1	Reserved
0	1	0	4
0	1	1	8
1	0	0	Reserved
1	0	1	Reserved
1	1	0	Reserved
1	1	1	Reserved

M11	M10	M9	Write Recovery
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	7
1	1	1	8

M3	Burst Type
0	Sequential
1	Interleaved

M15	M14	Mode Register Definition
0	0	Mode register (MR)
0	1	Extended mode register (EMR)
1	0	Extended mode register (EMR2)
1	1	Extended mode register (EMR3)

M6	M5	M4	CAS Latency (CL)
0	0	0	Reserved
0	0	1	Reserved
0	1	0	Reserved
0	1	1	3
1	0	0	4
1	0	1	5
1	1	0	6
1	1	1	7

DRAM controllers: reorder requests

(A) Without access scheduling (56 DRAM Cycles)

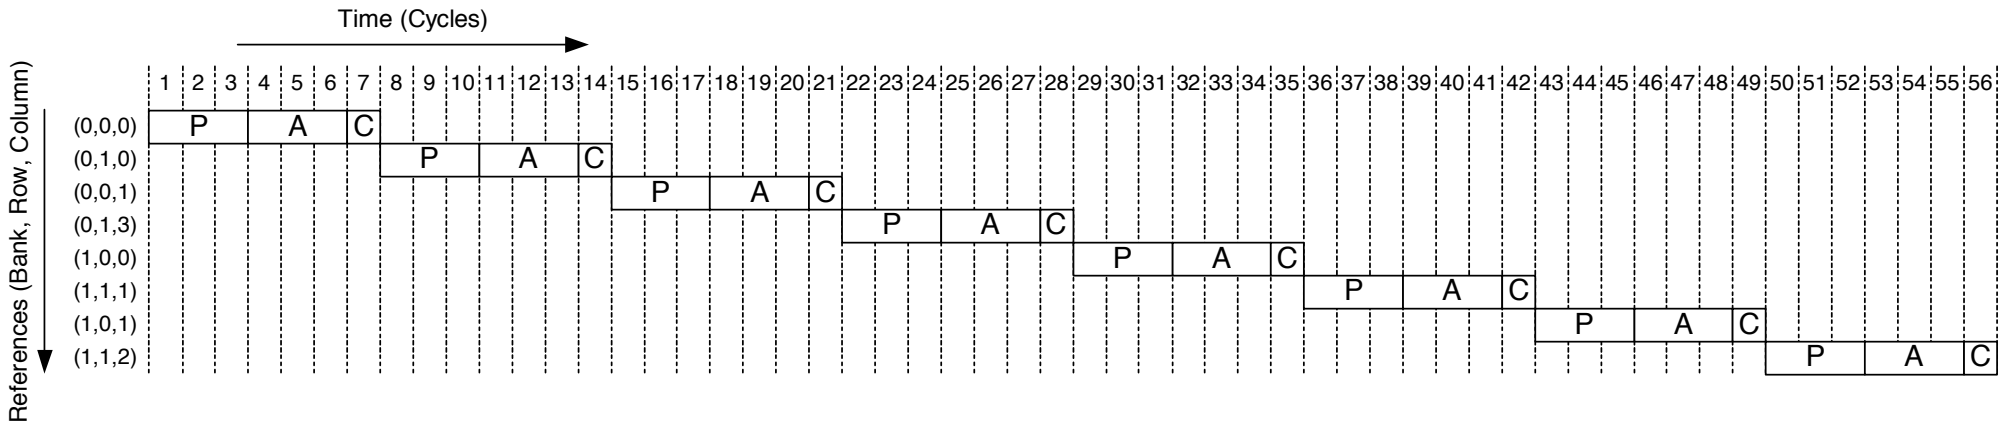

(B) With access scheduling (19 DRAM Cycles)

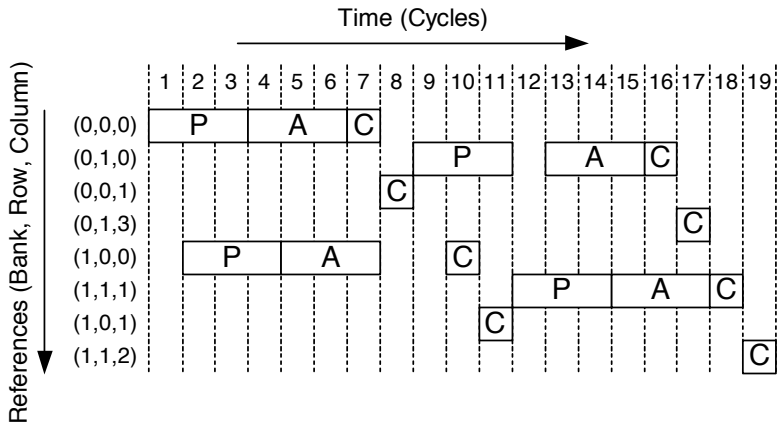

DRAM Operations:

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

Present and Future ...

MacBook Air … too thin to use DIMMs

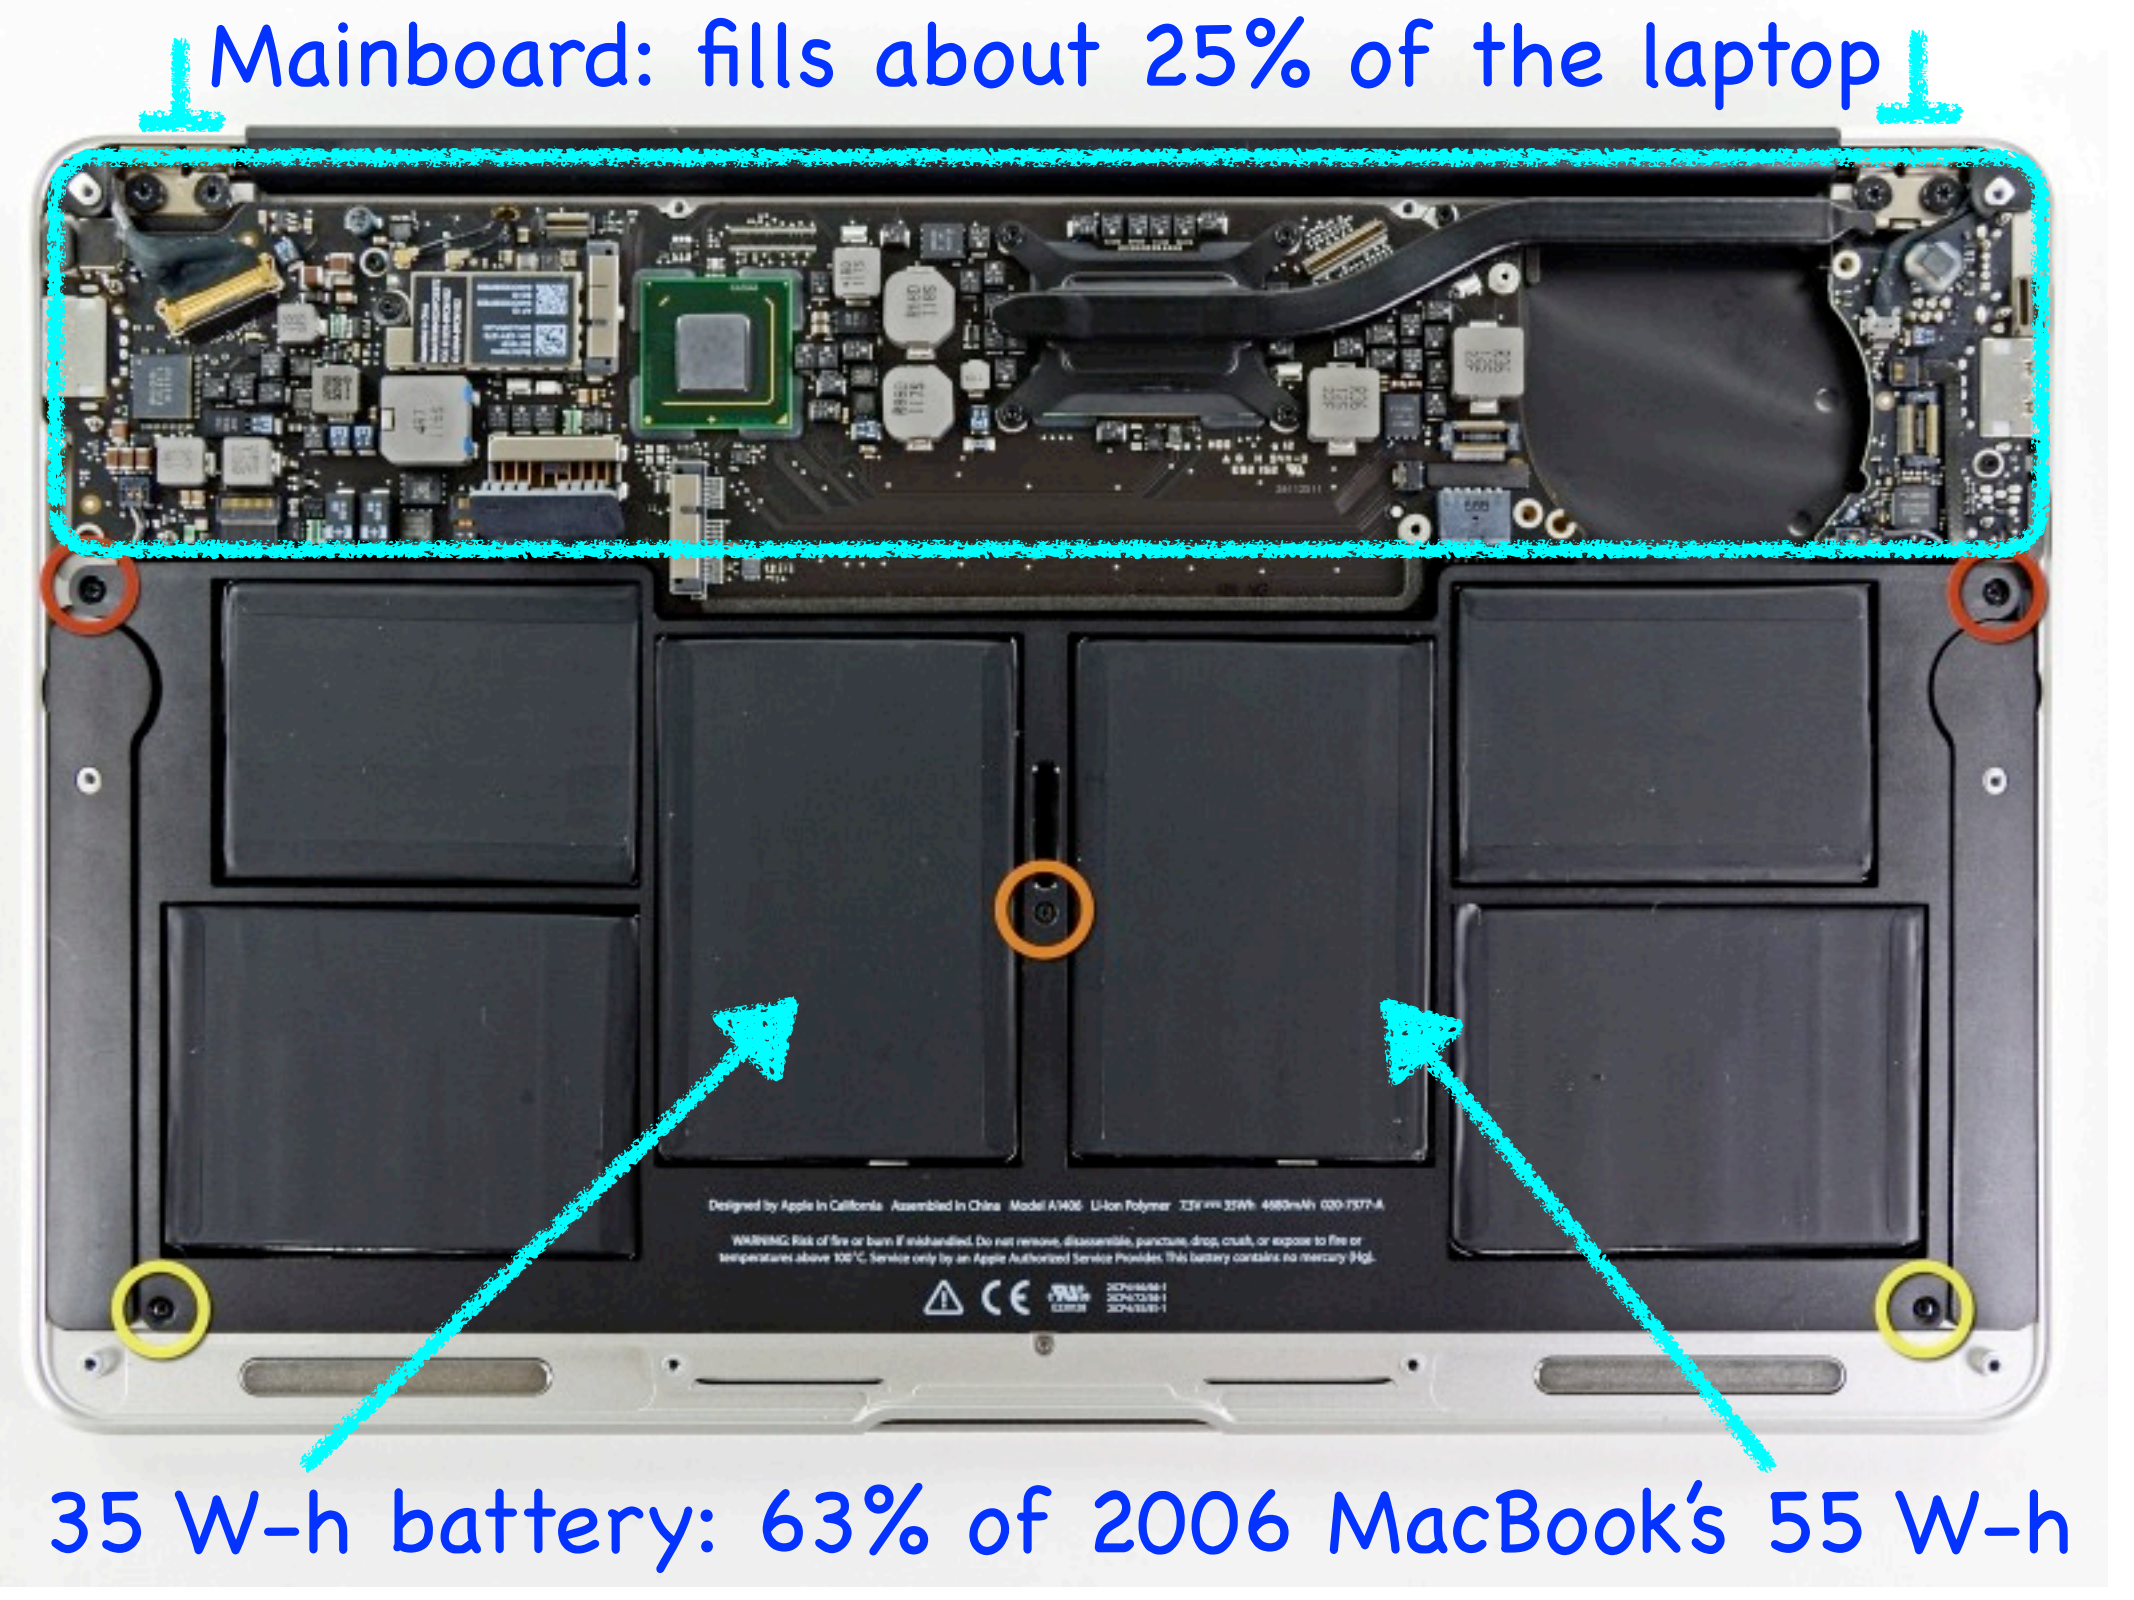
Mainboard: fills about 25% of the laptop
35 W-h battery: 63% of 2006 MacBook's 55 W-h

Main board

Core i5: CPU + DRAM controller

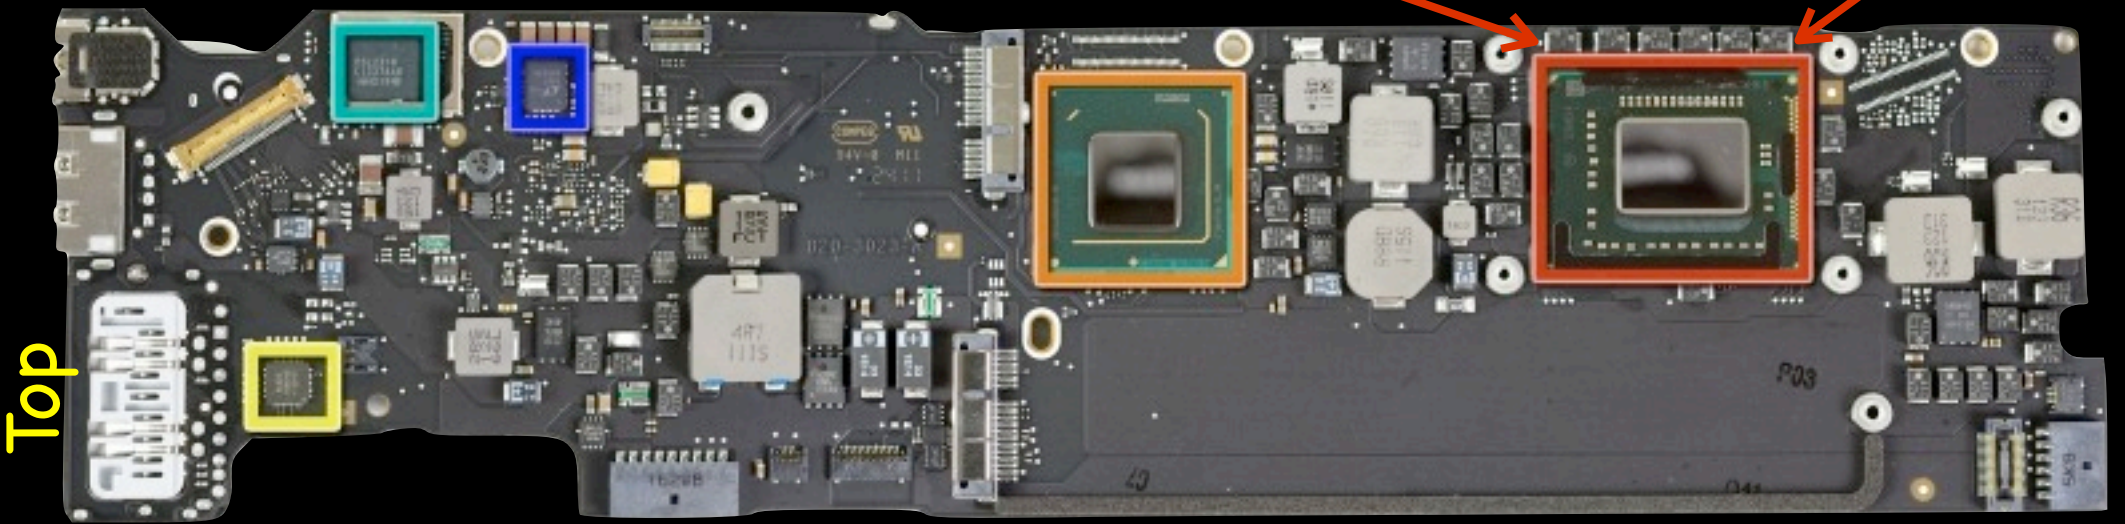

4GB DRAM soldered to the main board

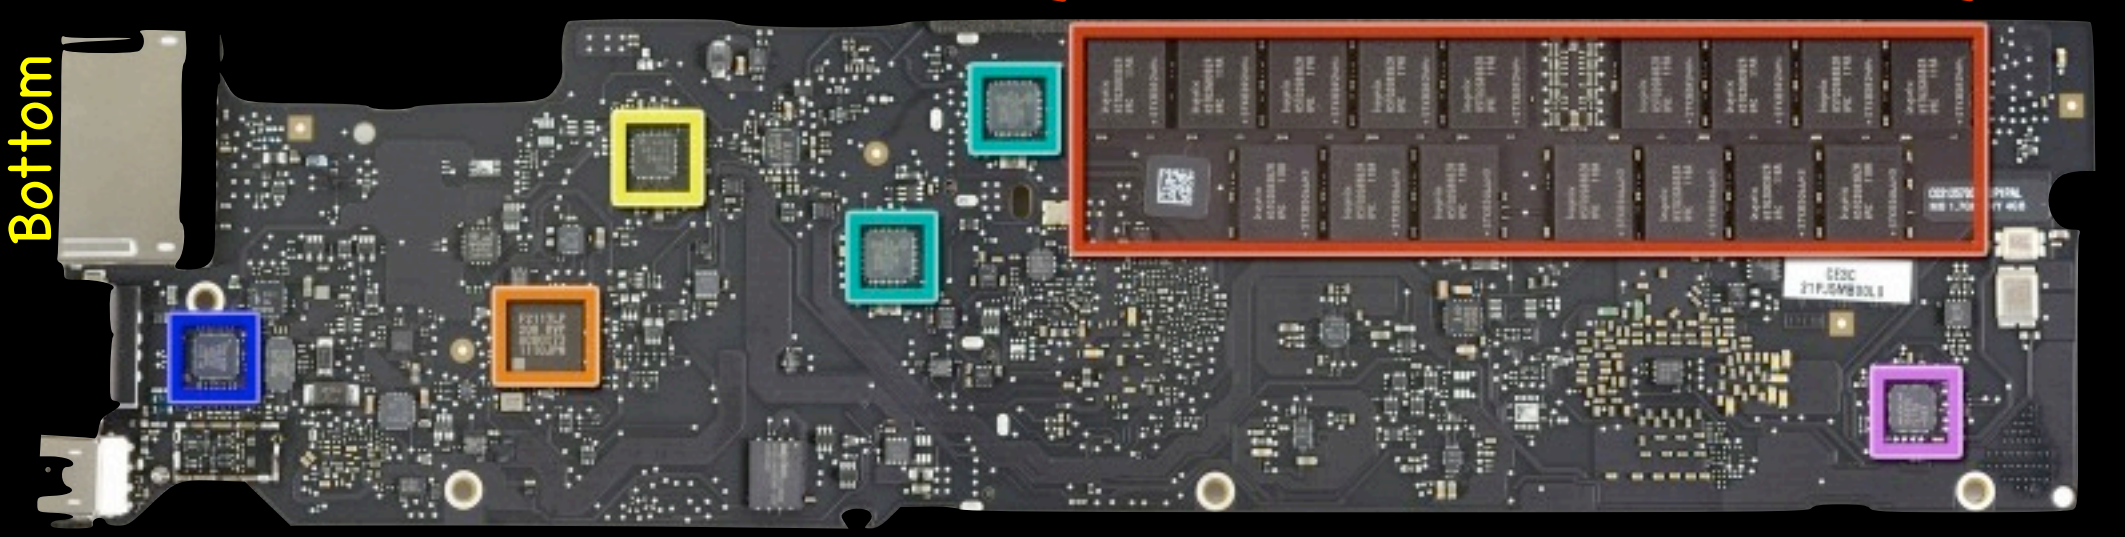

3-D memory stack

DRAM die size	10.7 mm × 13.3 mm
DRAM die thickness	50 μm
TSV count in DRAM	1,560
DRAM capacity	512 Mbit/die × 2 strata
CMOS logic die size	17.5 mm × 17.5 mm
CMOS logic die thickness	200 μm
CMOS logic bump count	3,497
CMOS logic process	0.18 μm CMOS
DRAM-logic FTI via pitch	50 μm
Package size	33 mm × 33 mm
BGA terminal	520 pin / 1mm pitch

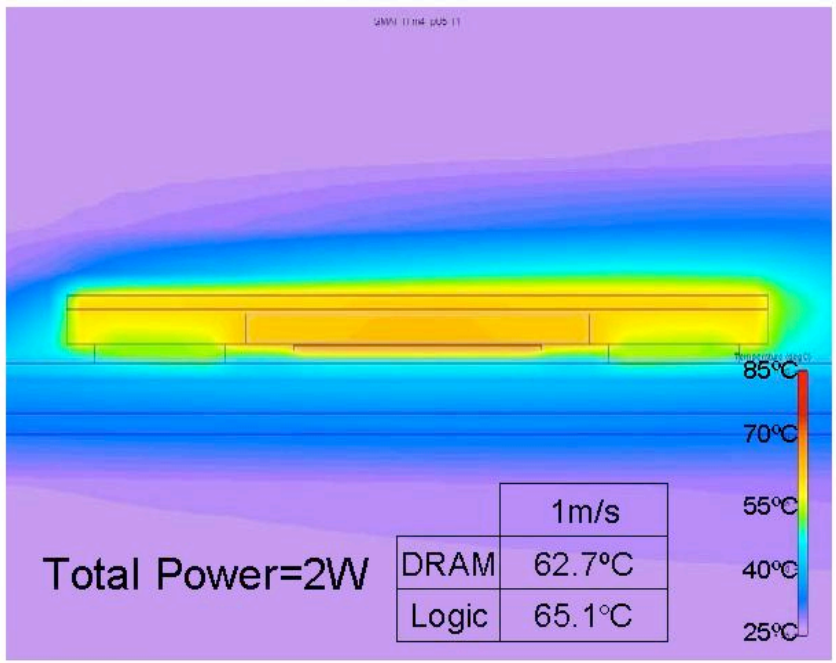

1 Gbit stacked DRAM with TSV (512 Mbit × 2 strata)

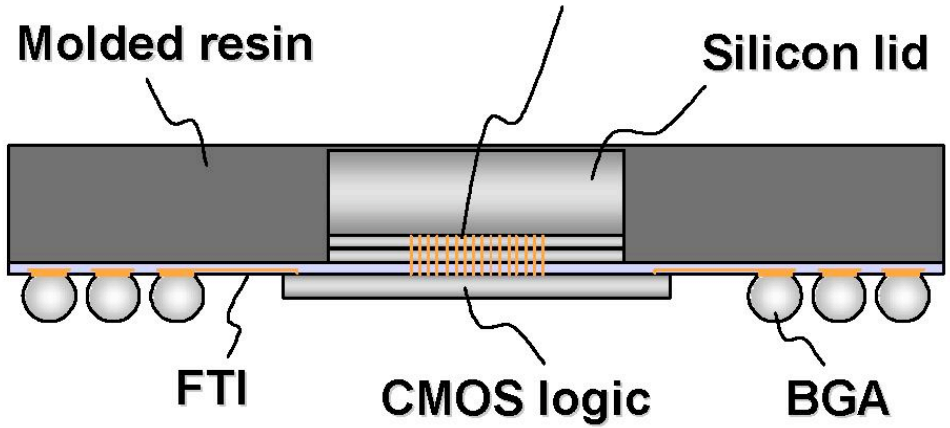

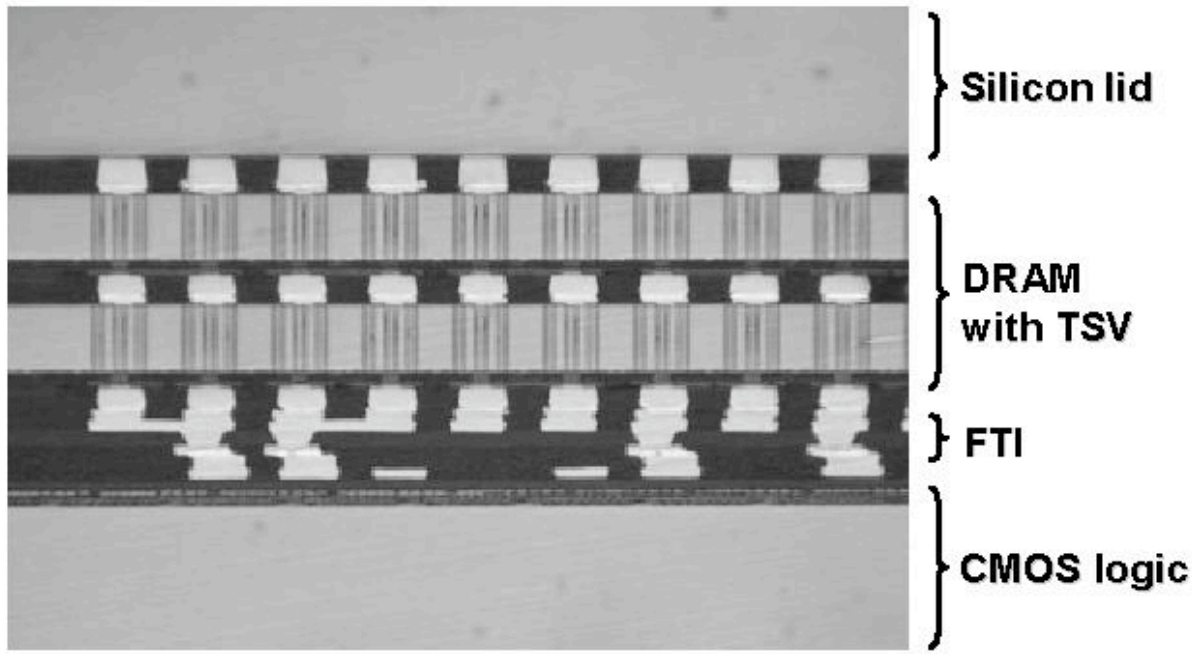

A 3D Stacked Memory Integrated on a Logic Device Using SMAFTI Technology

Yoichiro Kurita[1], Satoshi Matsui[1], Nobuaki Takahashi[1], Koji Soejima[1], Masahiro Komuro[1], Makoto Itou[1], Chika Kakegawa[1], Masaya Kawano[1], Yoshimi Egawa[2], Yoshihiro Saeki[2], Hidekazu Kikuchi[2], Osamu Kato[2], Azusa Yanagisawa[2], Toshiro Mitsuhashi[2], Masakazu Ishino[3], Kayoko Shibata[3], Shiro Uchiyama[3], Junji Yamada[3], and Hiroaki Ikeda[3]
[1]NEC Electronics, [2]Oki Electric Industry, and [3]Elpida Memory
1120 Shimokuzawa, Sagamihara, Kanagawa 229-1198, Japan
y.kurita@necel.com

Next week's lectures: Timing ...

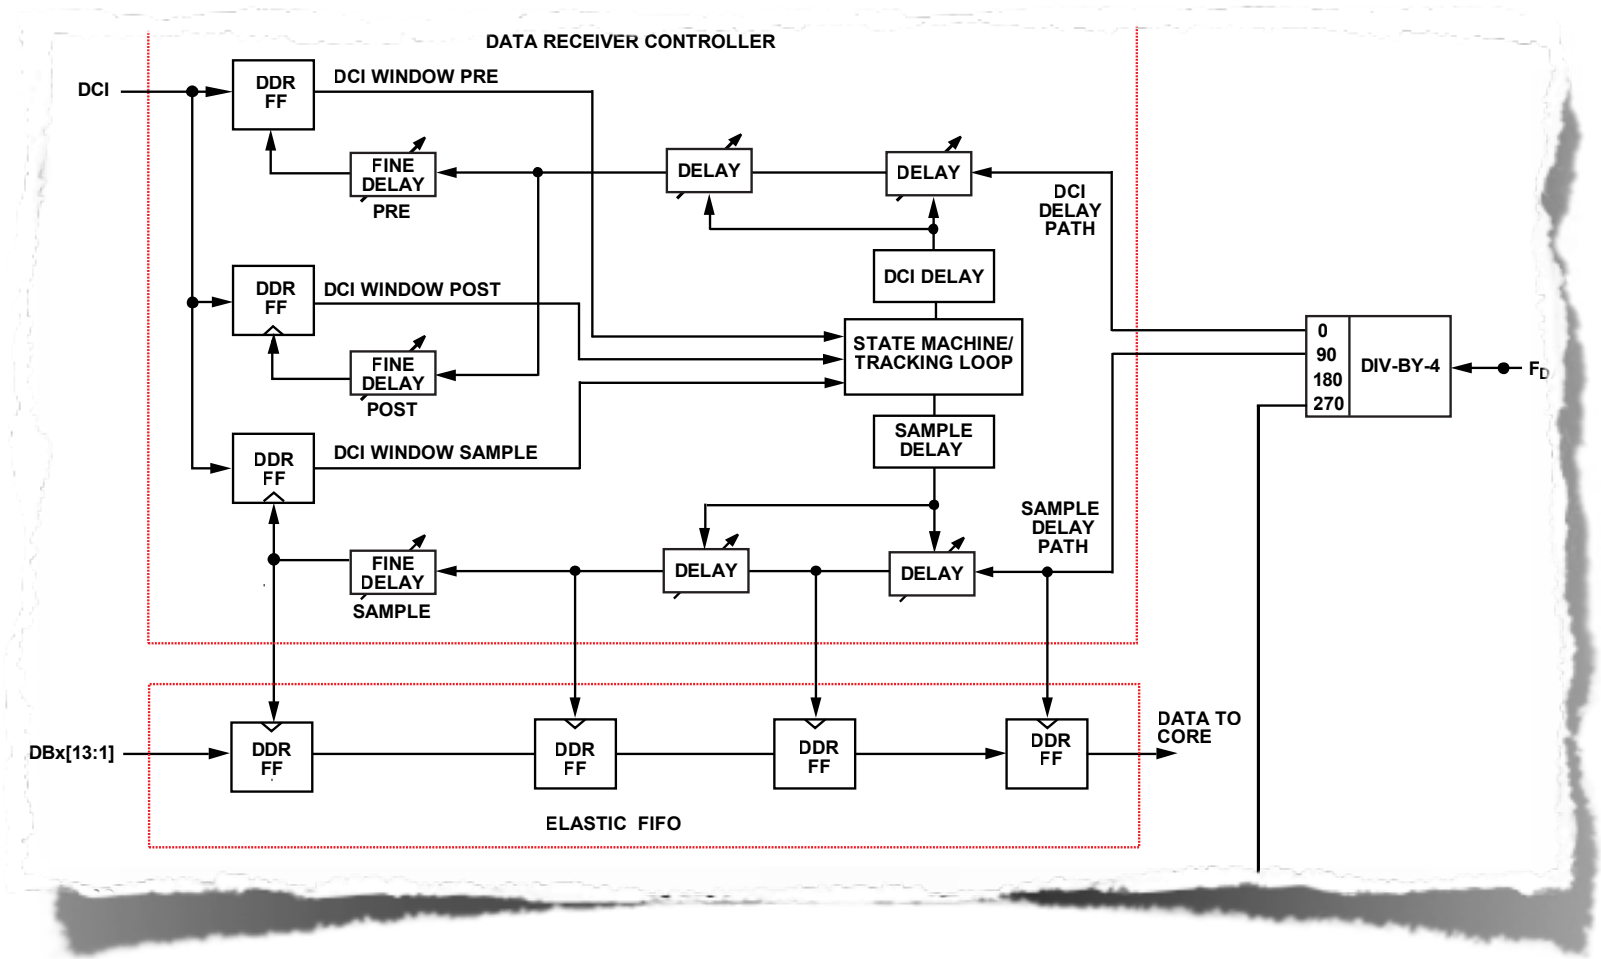

7	Tue 2/28	Lec #13: Timing (1) (Slides) Reading: 3.5-3.6
	Thu 3/1	Lec #14: Timing (2) (Slides)