

ECE 5730
Memory Systems
Spring 2009

More on DRAMs

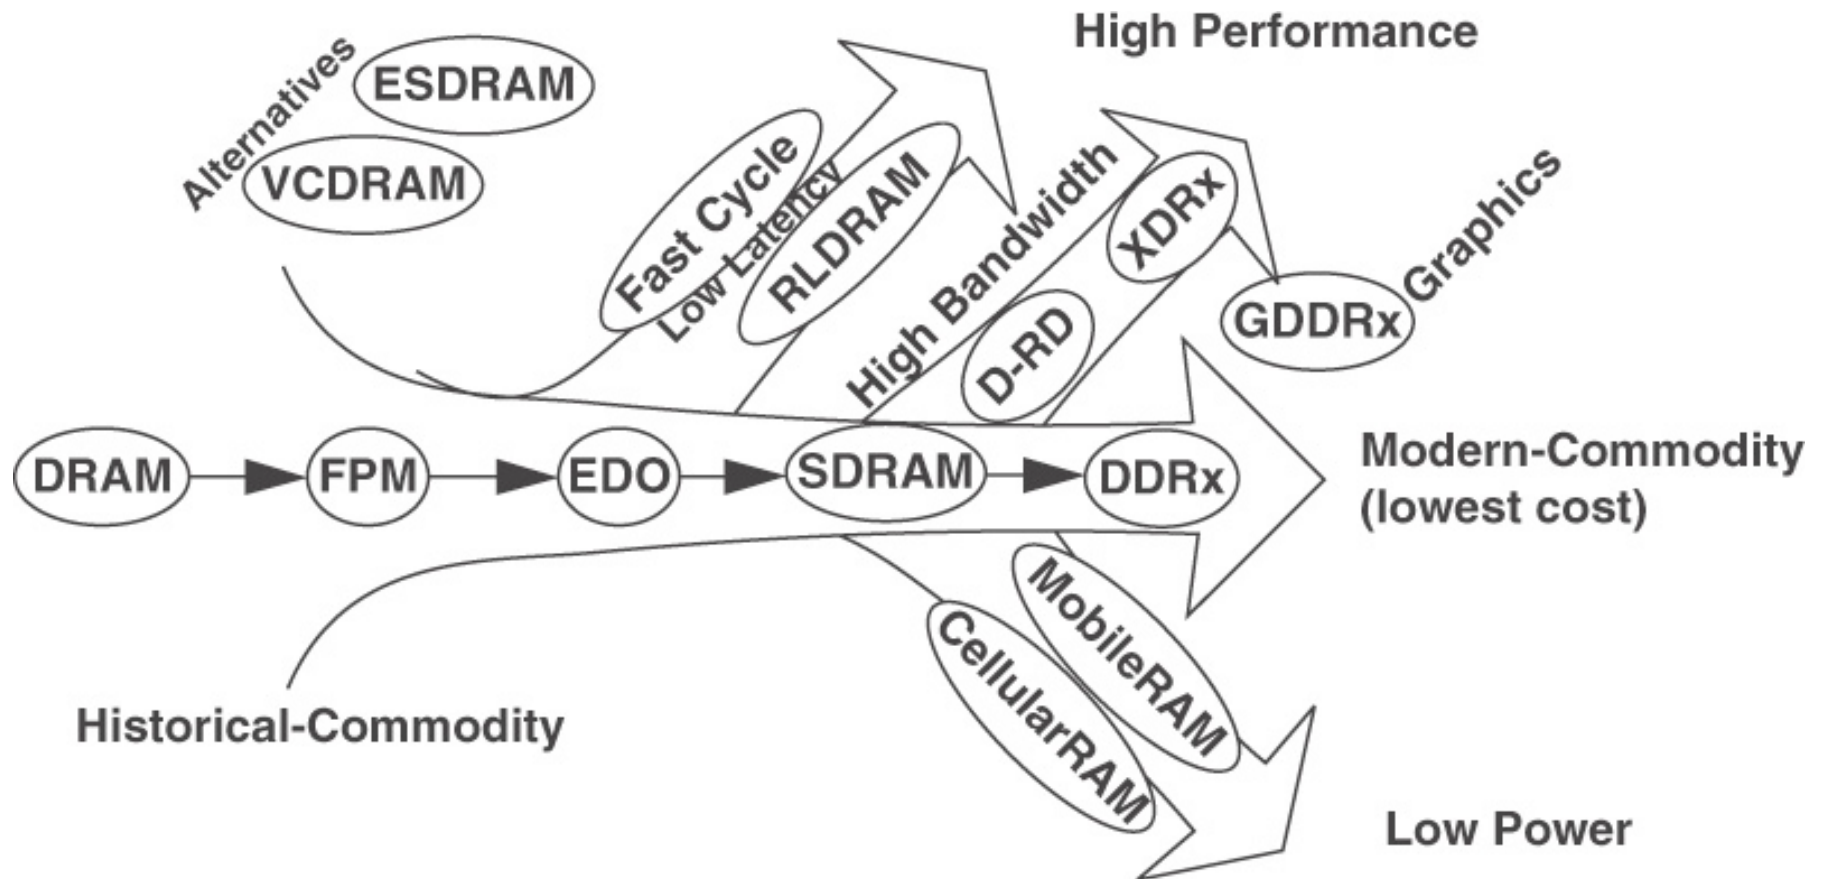


Cornell University

Announcements

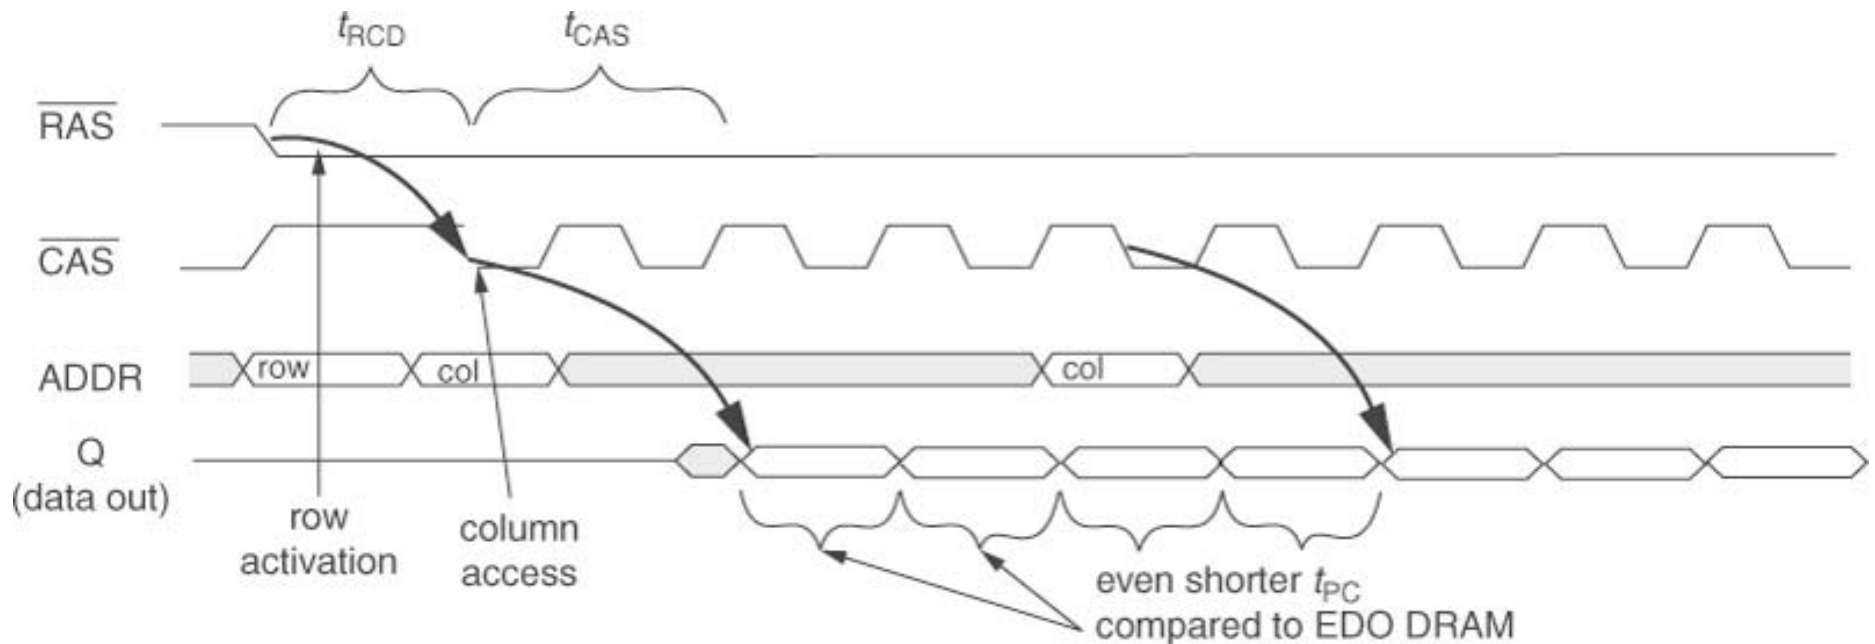
- **Exam I**
 - **You may attend either date**
 - **Wednesday, March 11, 6:30-9:30pm**
 - **Friday, March 13, 12-3pm**
 - **Send me an email TODAY saying which you will attend**
 - **Those taking the exam on Wednesday are prohibited from discussing the exam until after the Friday exam**
- **No office hours today**

The Evolution of the Modern DRAM



Burst EDO DRAM

- DRAM automatically increments column address internally to burst multiple units of data



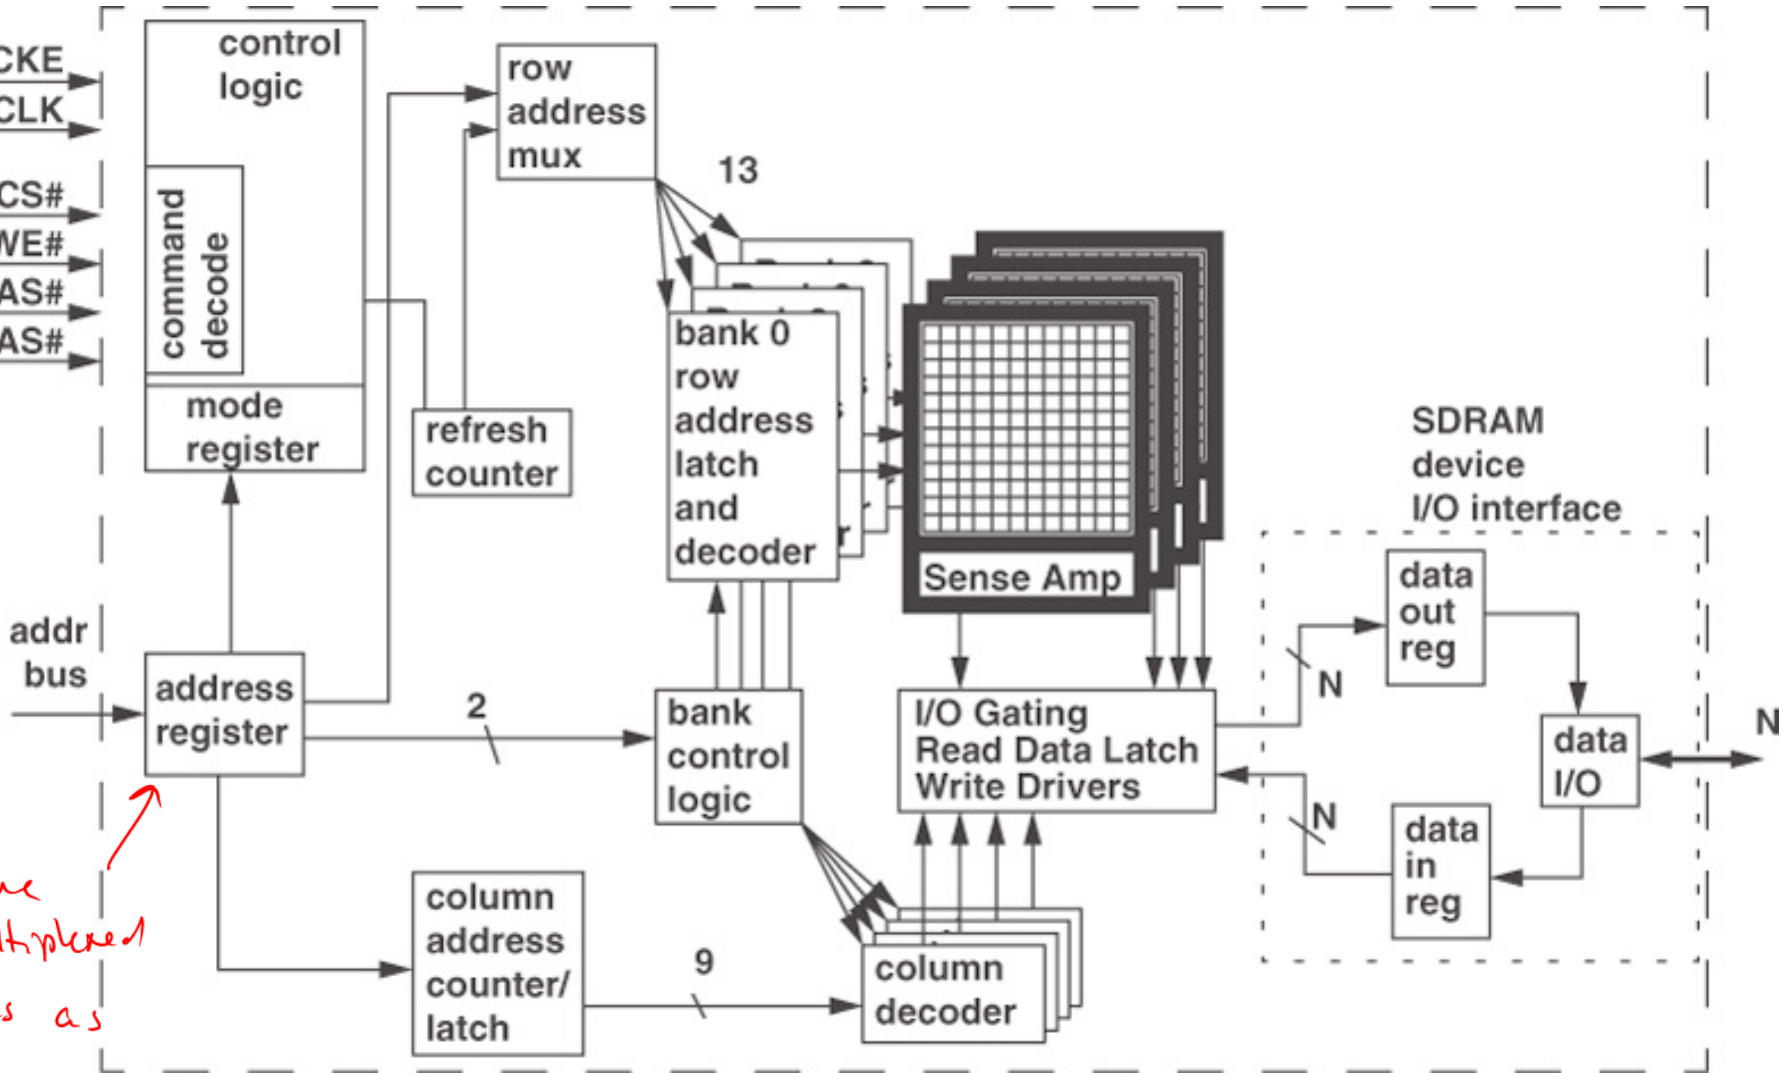
Synchronous DRAM (SDRAM)

- **Asynchronous DRAMs had limited bandwidth and concurrency**
 - Limited overlap of address and data phases of consecutive accesses
 - Single bank of arrays
- **SDRAMs greatly increase concurrency and bandwidth by**
 - Registering the inputs and outputs
 - Incorporating multiple independent banks
 - Increasing the amount of on-chip control intelligence

SDRAM Device Architecture

form command, clock


Same multiplexed bus as for



SDRAM Commands

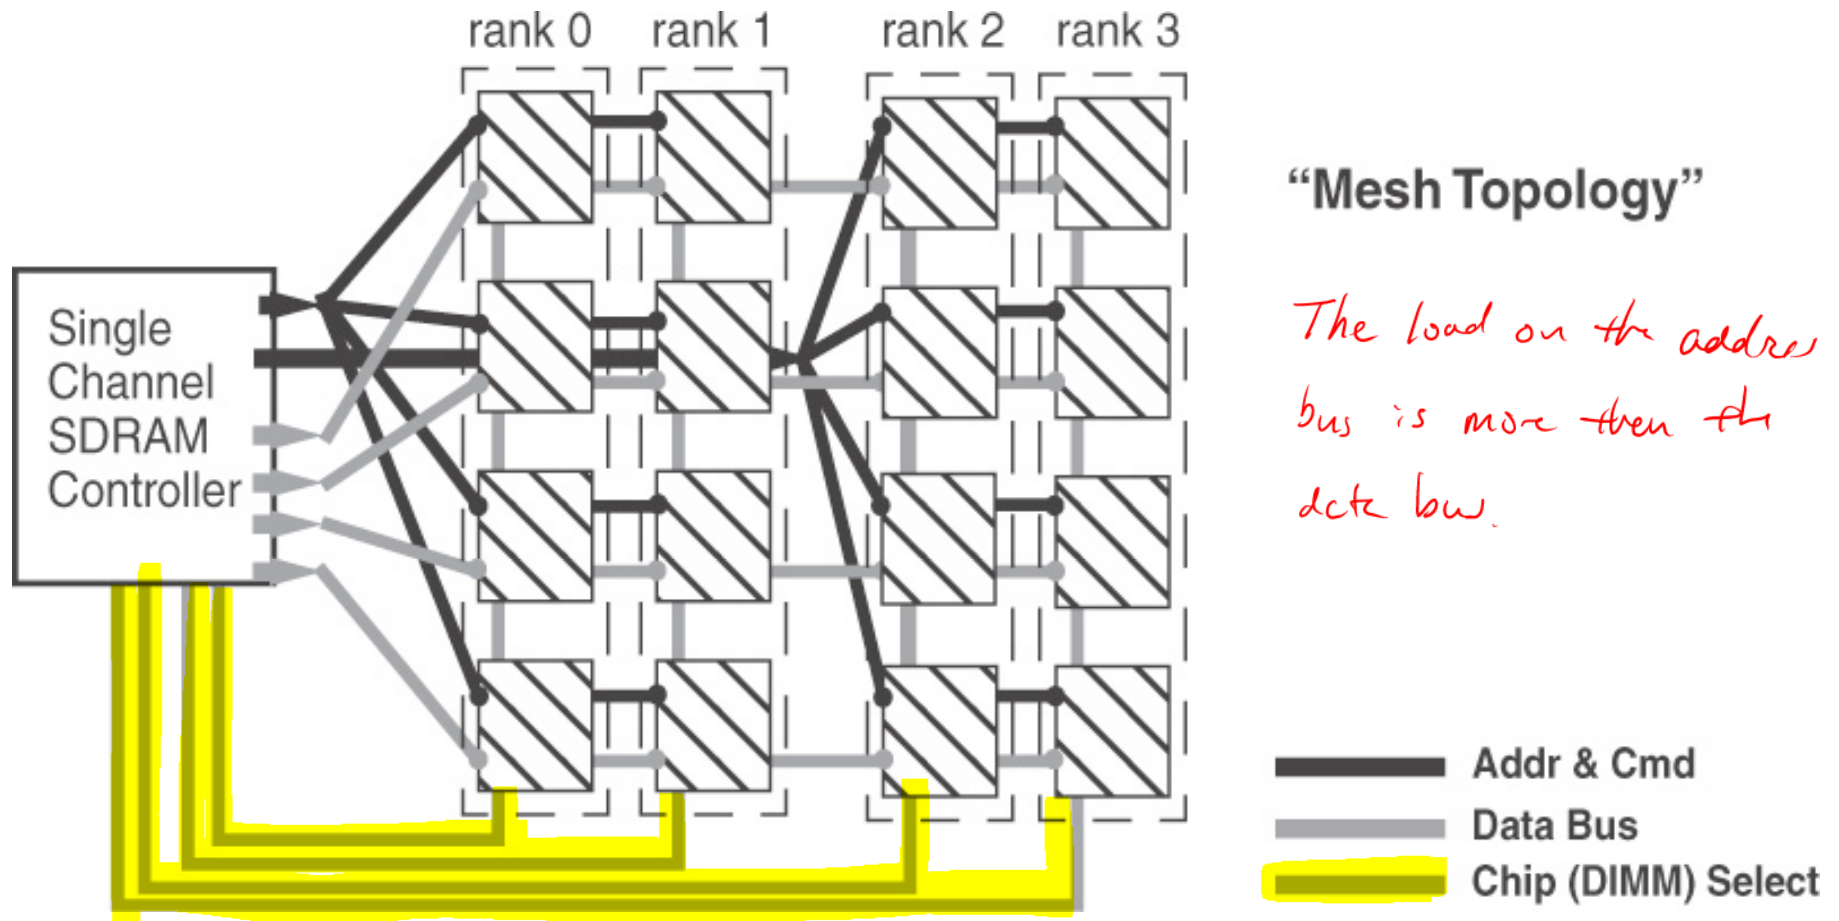
Name (Function)	CS#	RAS#	CAS#	WE#
COMMAND INHIBIT (NOP)	H	X	X	X
NO OPERATION (NOP)	L	H	H	H
ACTIVE (Select bank and activate row)	L	L	H	H
READ (Select bank and column, and start READ burst)	L	H	L	H
WRITE (Select bank and column, and start WRITE burst)	L	H	L	L
BURST TERMINATE	L	H	H	L
PRECHARGE (Deactivate row in bank or banks)	L	L	H	L
AUTO REFRESH or SELF REFRESH (Enter self refresh mode)	L	L	L	H
LOAD MODE REGISTER	L	L	L	L

SDRAM Access Protocol



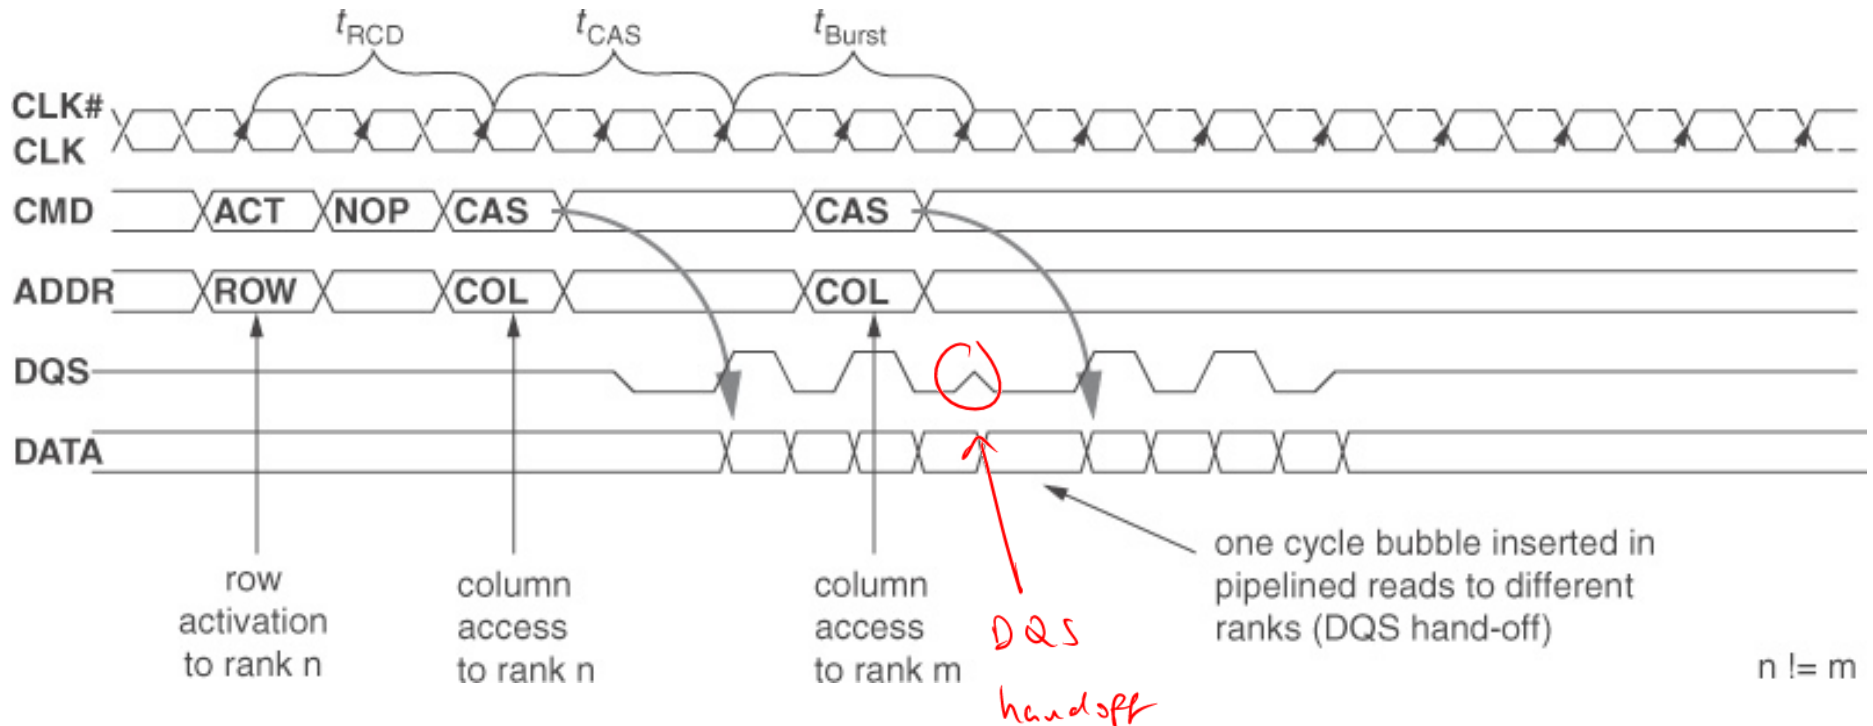
Double Data Rate (DDR) SDRAM

- Address and Command buses more heavily loaded than data bus



Double Data Rate (DDR) SDRAM

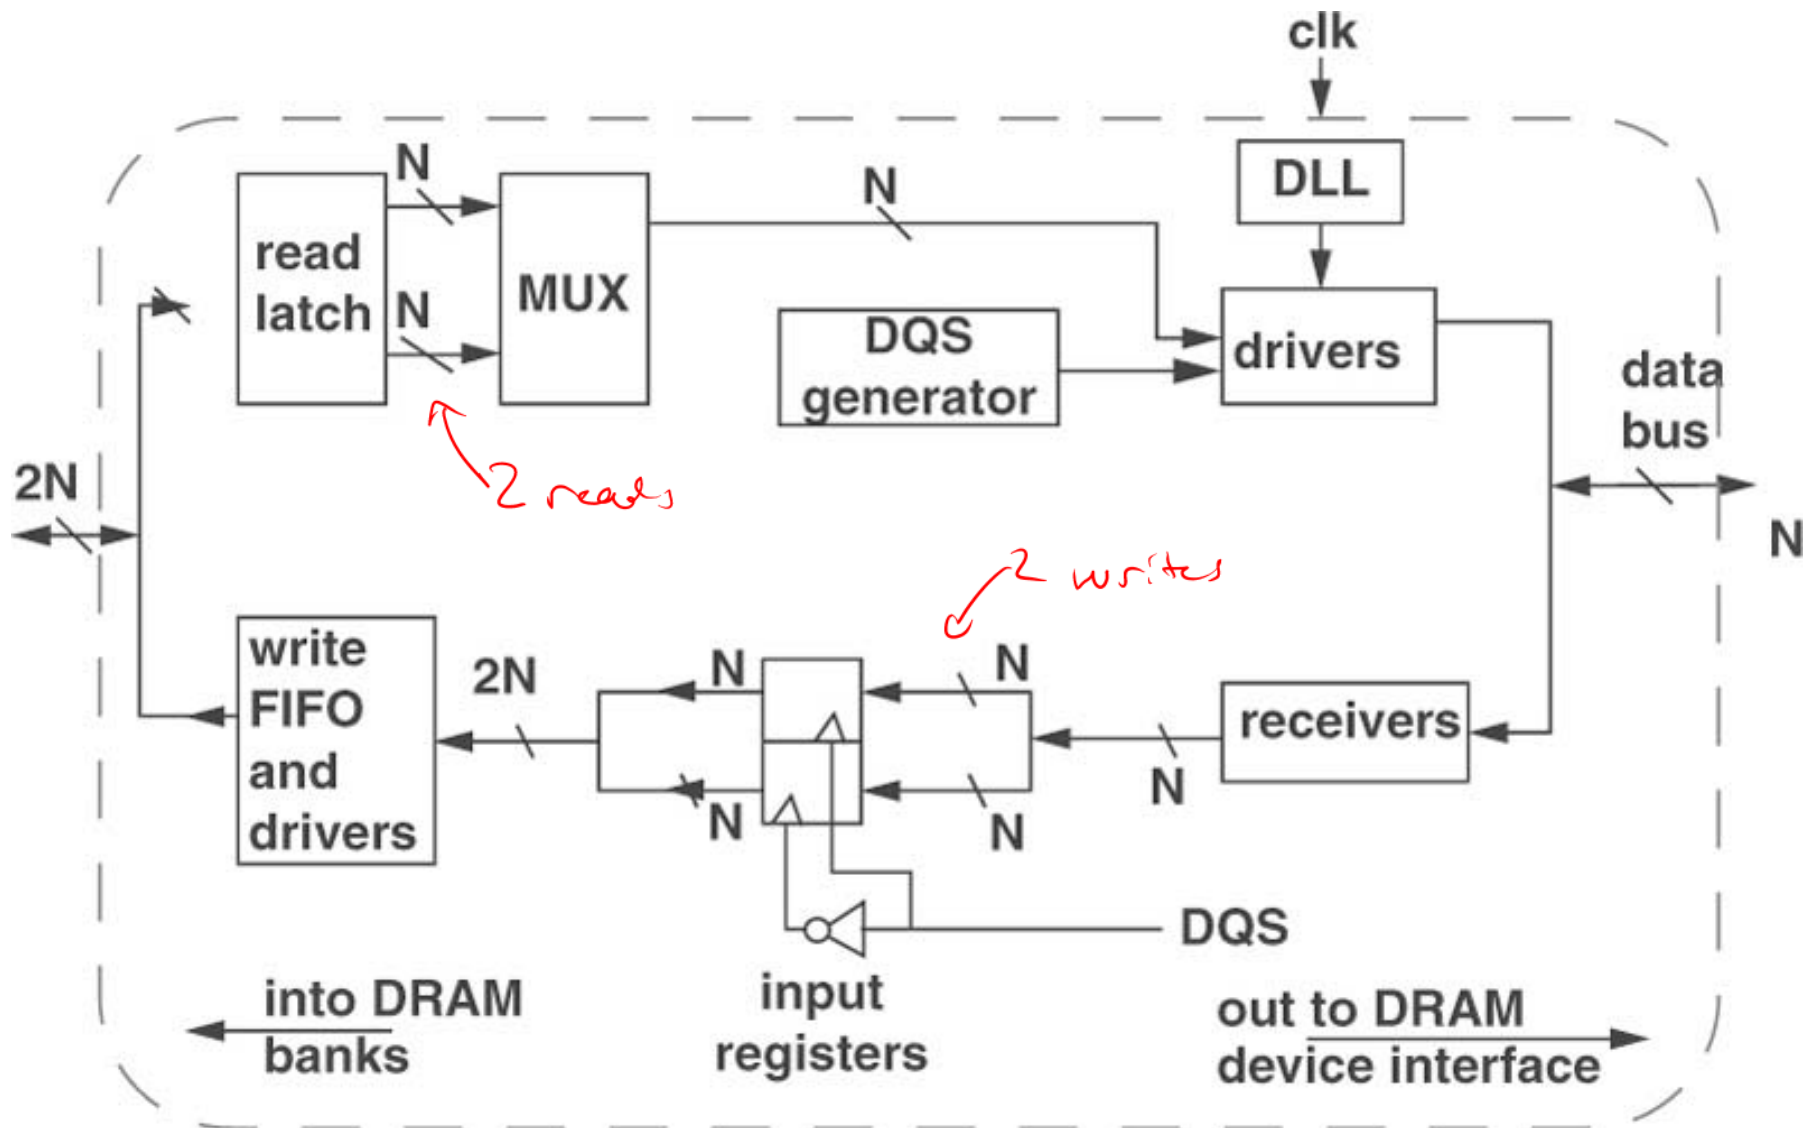
- Run data bus on both edges of the clock



DQS: (shared) data strobe signal controlled by the source of the data

→ sent from the source of data

DDR SDRAM I/O Architecture

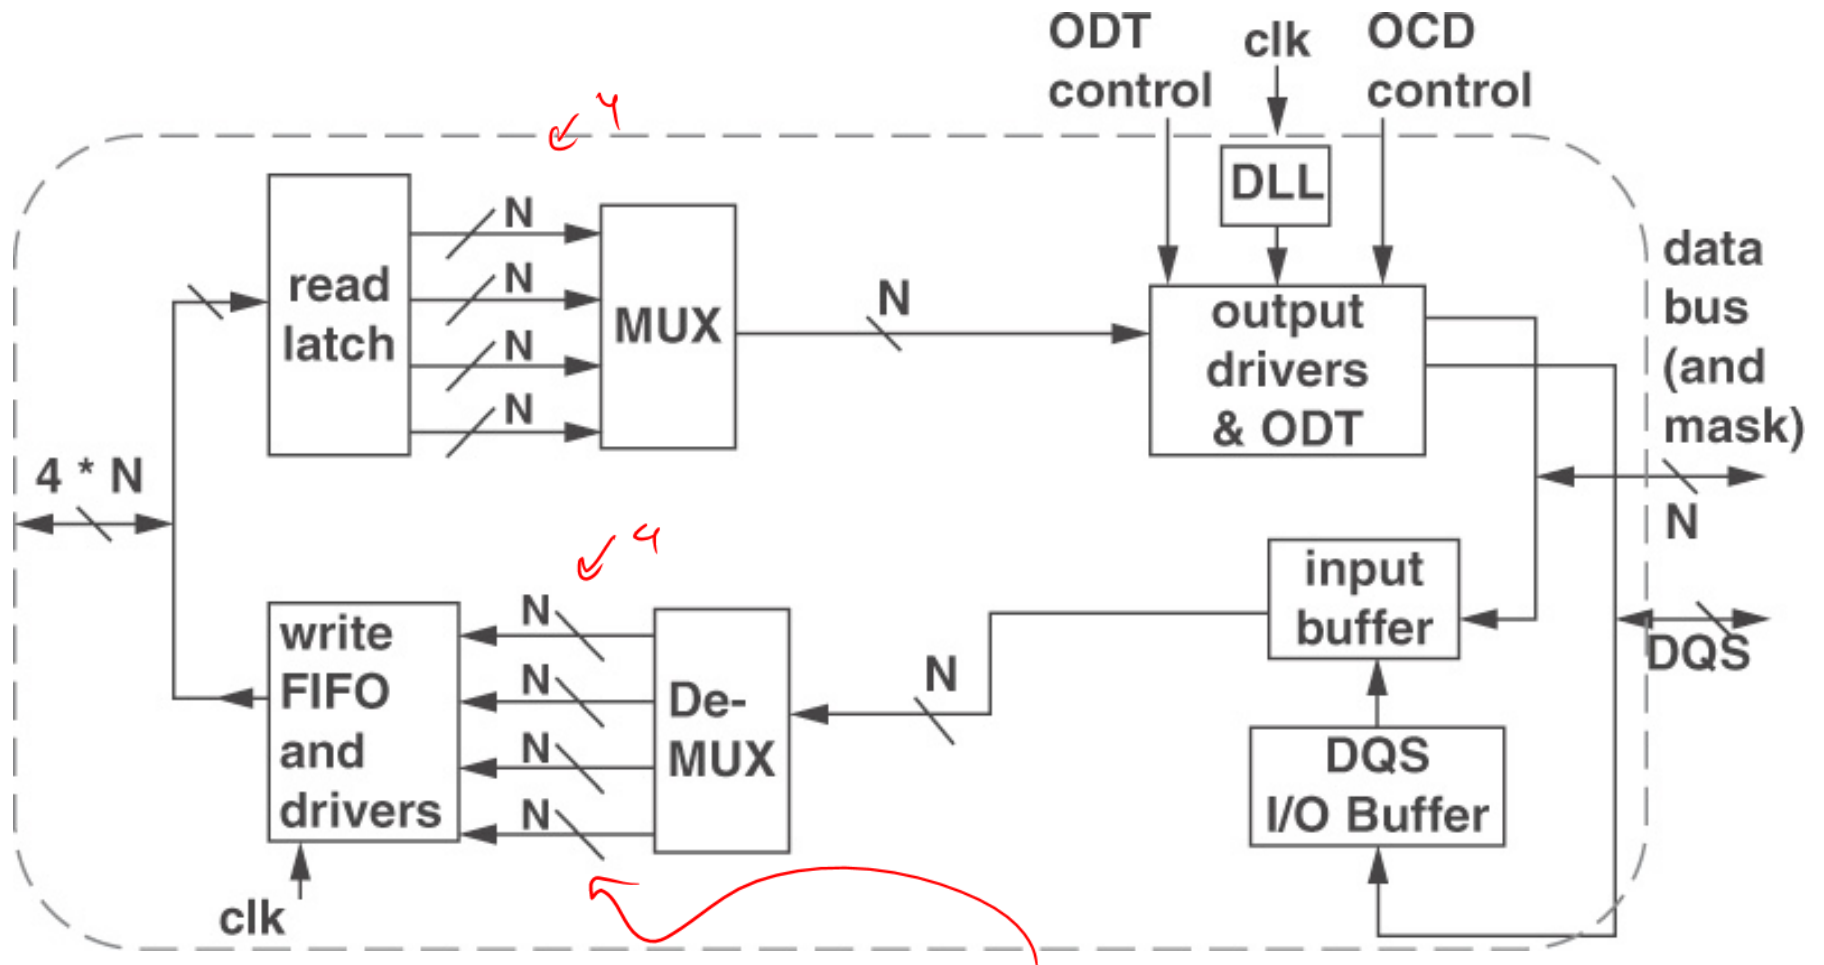


2-bit prefetch architecture

DDR2 and DDR3

- **Increased prefetch length**
 - 4 in DDR2 and 8 in DDR3
- **Some interface signaling changes**
 - E.g., differential DQS
- **Some additional commands**

DDR2 I/O Architecture



basically the same, just 4 instead of 2

SDRAM and DDRx Comparison

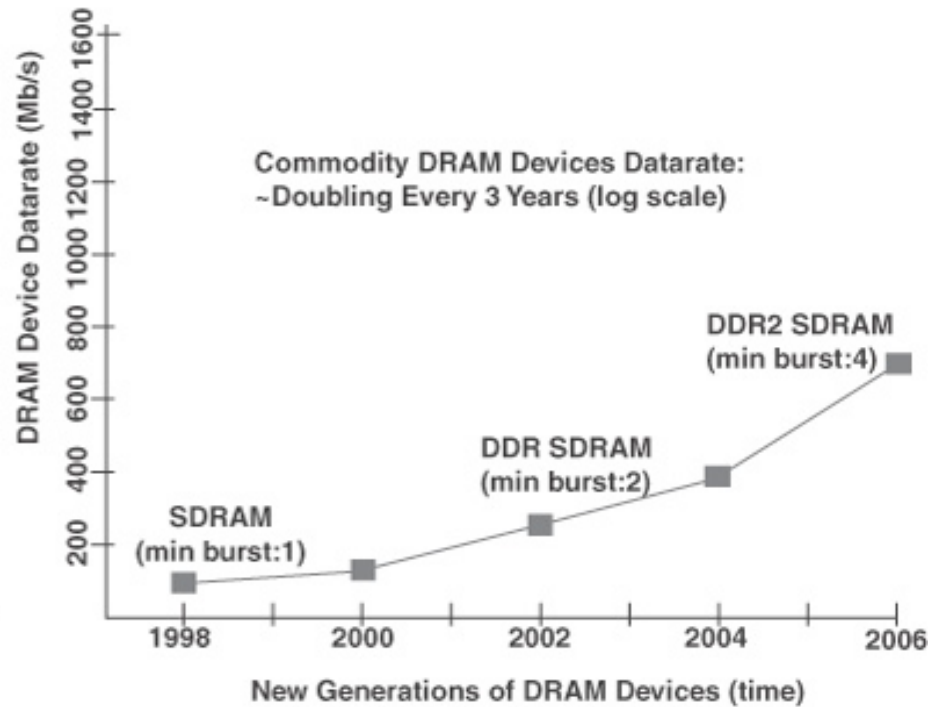
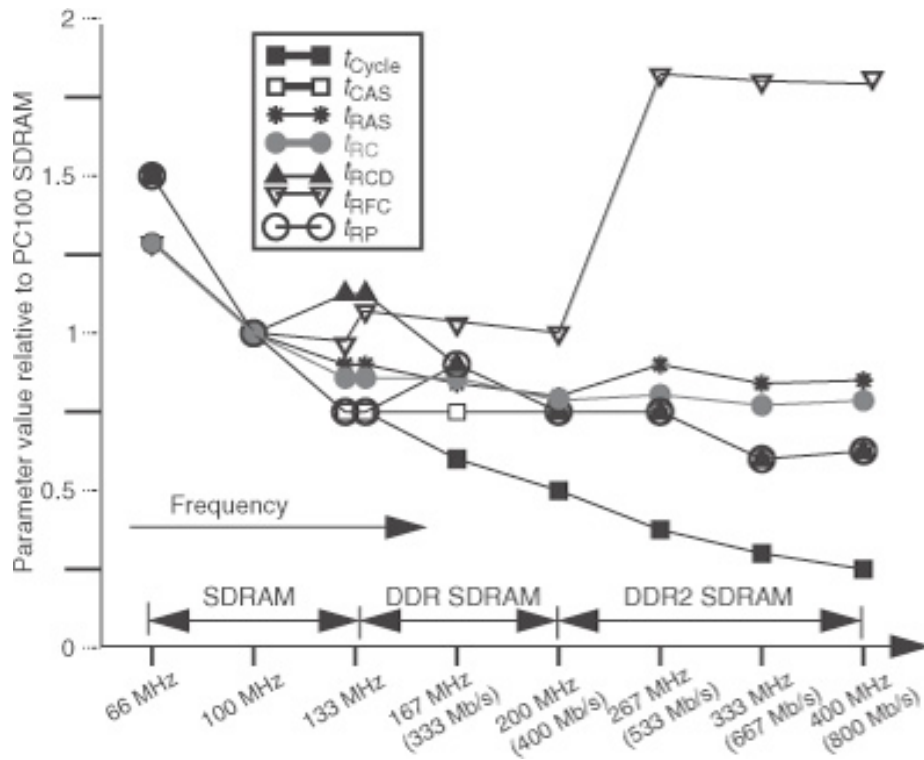
Variables	SDRAM	DDR1	DDR2	DDR3
Clock	100/133/166 MHz	100/133/166/200 MHz	200/266/333/400 MHz	400/533/667/800 MHz
Transfer Data Rate	100/133/166 Mbps	200/266/333/400 Mbps	400/533/667/800 Mbps	800/1066/1333/1600 Mbps
I/O width	x16/x32	x4/x8/x16/x32	x4/x8/x16	x4/x8/x16/x32
Prefetch bit width	1 bit	2 bits	4 bits	8 bits
Clock Input	Single Clock	Differential Clock	Differential Clock	Differential Clock
Burst Length	1, 2, 4, 8, full page	2, 4, 8	4, 8	8, 4 (Burst chop)
Data Strobe	Unsupported	Single data strobe	Differential data strobe	Differential data strobe
Supply Voltage	3.3V/2.5V	2.5V	1.8V	1.5V
Interface	LVTTL	SSTL_2	SSTL_1.8	SSTL_1.5
$\overline{\text{CAS}}$ latency (CL)	2, 3 clock	2, 2.5, 3 clock	3, 4, 5, clock	5, 6, 7, 8, 9, 10 clock

← how many bits do you get?

← need better clock

variable

DRAM Scaling Trends




Random row access cycle time has decreased 7% per year
Transfer data rate has doubled every three years

DRAM Internals

- **Cell structure**
- **Bitline array structure**
- **Sense amplifier**
- **Decoders and redundancy**

DRAM Cell Structures

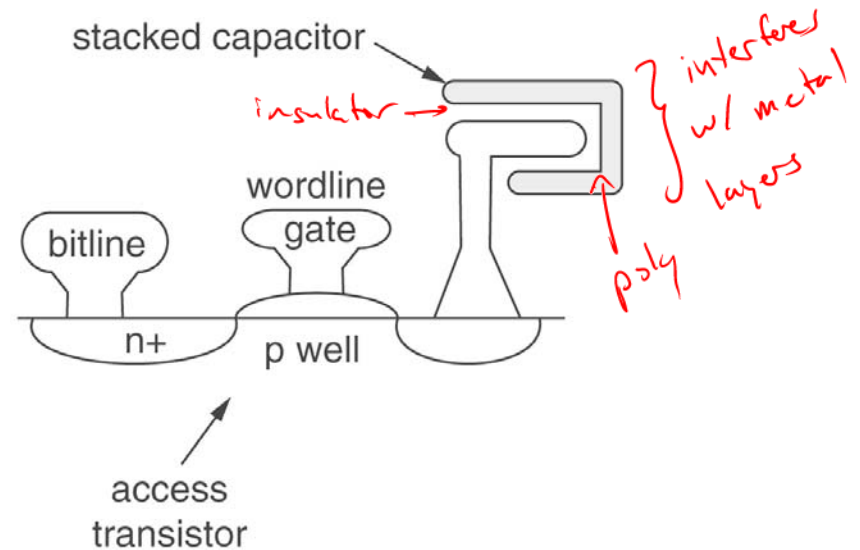



trench capacitor

eases metallization
(embedded DRAM)

→ will play nice with metal layers as it's stuck on the substrate

poly
insulator
50 times deeper than it is wide. Not to scale



stacked capacitor

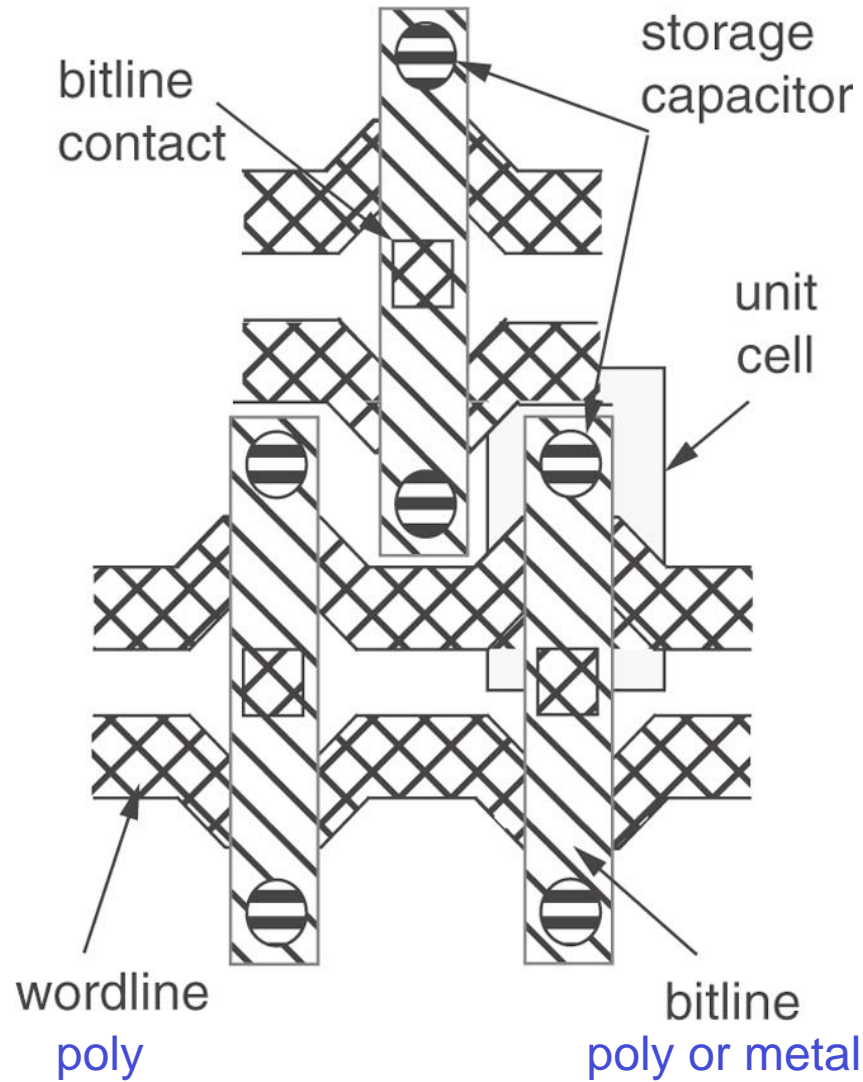
more reliable
(commodity DRAM)

→ way better, less leakage

interferes w/ metal layers

DRAM Cell Layout

shared between
2 caps



30fF in 90nm

cell area = $8F^2$
F = feature size
cell scales with F

DRAM Refresh

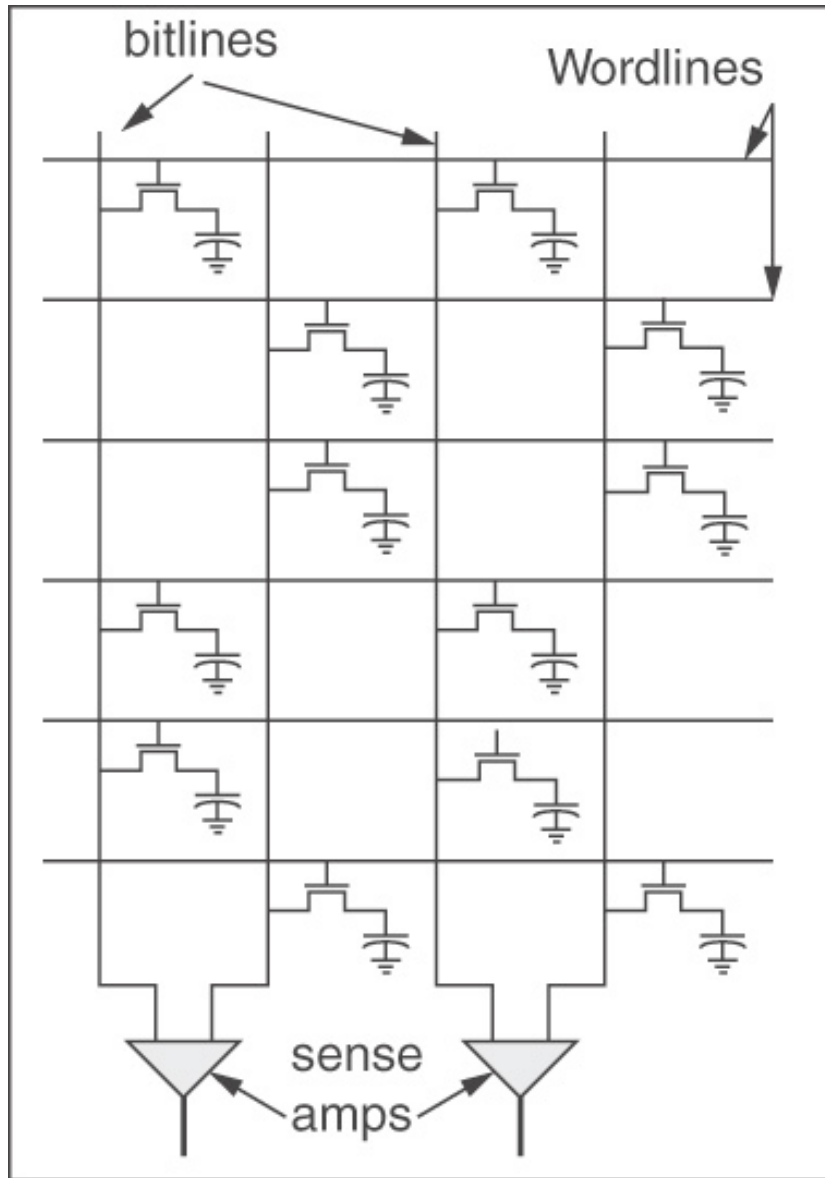
- Cell capacitive charge leaks away over time
- Charge must be periodically restored through a *refresh* operation (row activation)
- Refresh rate has remained pretty constant across generations
 - 8192 refresh commands must be made in 64ms
→ otherwise potential for data loss

bitlines will be precharged to some reference voltage.

when a wordline goes high, either we drain charge to/from the cap, and the

Bitline Array Structure


sense amp is differential, so it picks that up



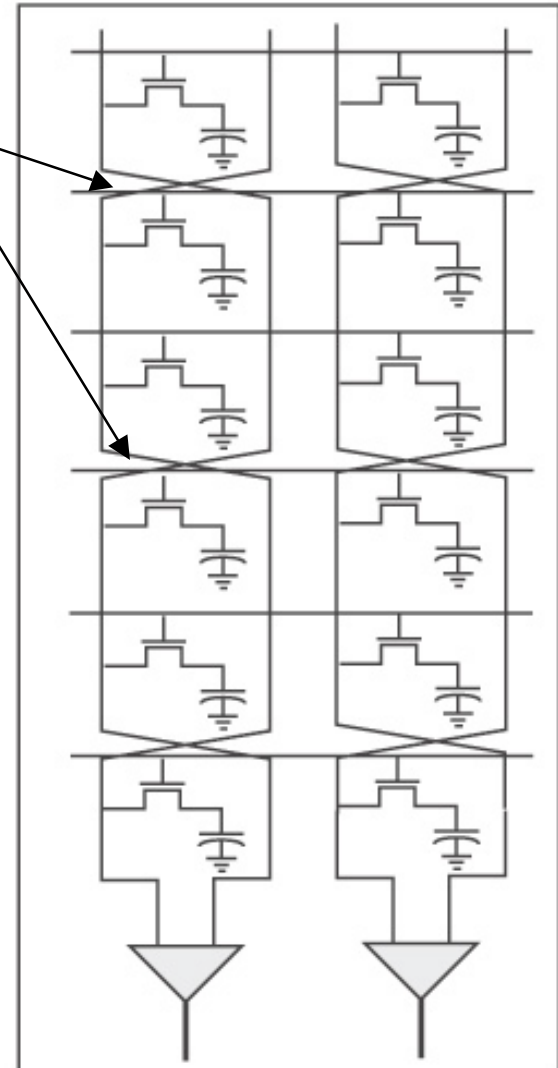
periodic "twists"
increase noise immunity

(addressing cross talk)

common-mode noise is ok



differential mode is BAD



folded bitline

DRAM Sense Amplifier

- Capacitance of storage capacitor is about 1/10 that of the bitline



- Accessing a cell causes tiny bitline voltage changes

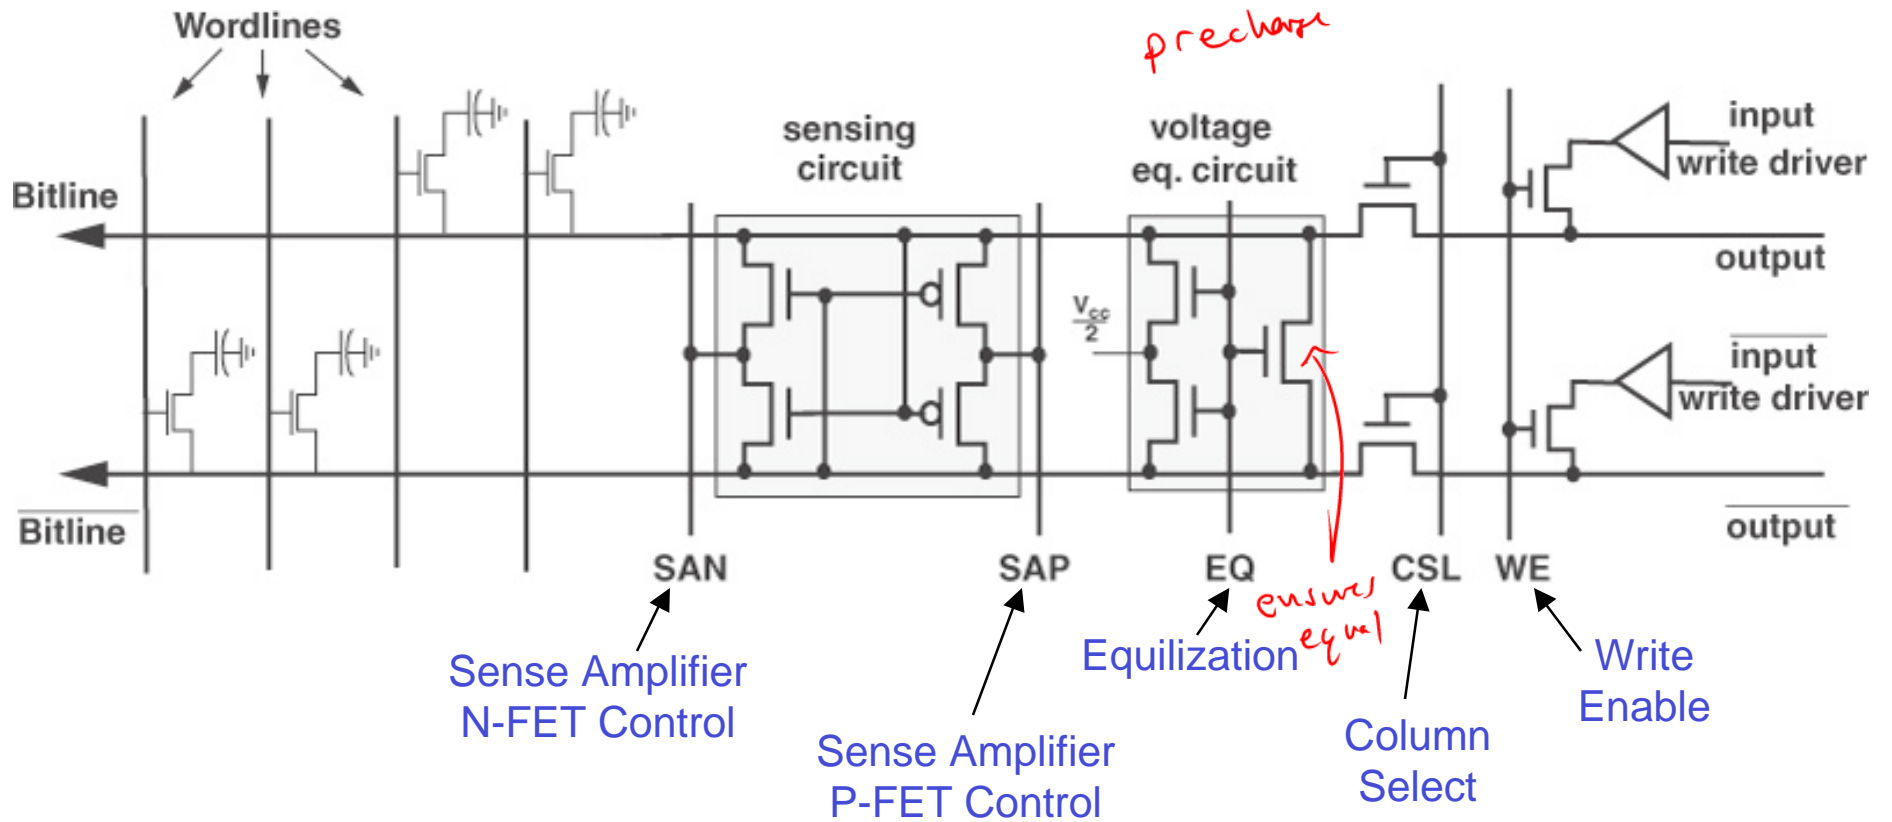
- Differential sense amplifier compares bitline voltage against reference bitline *(huge gain, probably positive feedback)*

- Result is amplified to digital 0 or 1

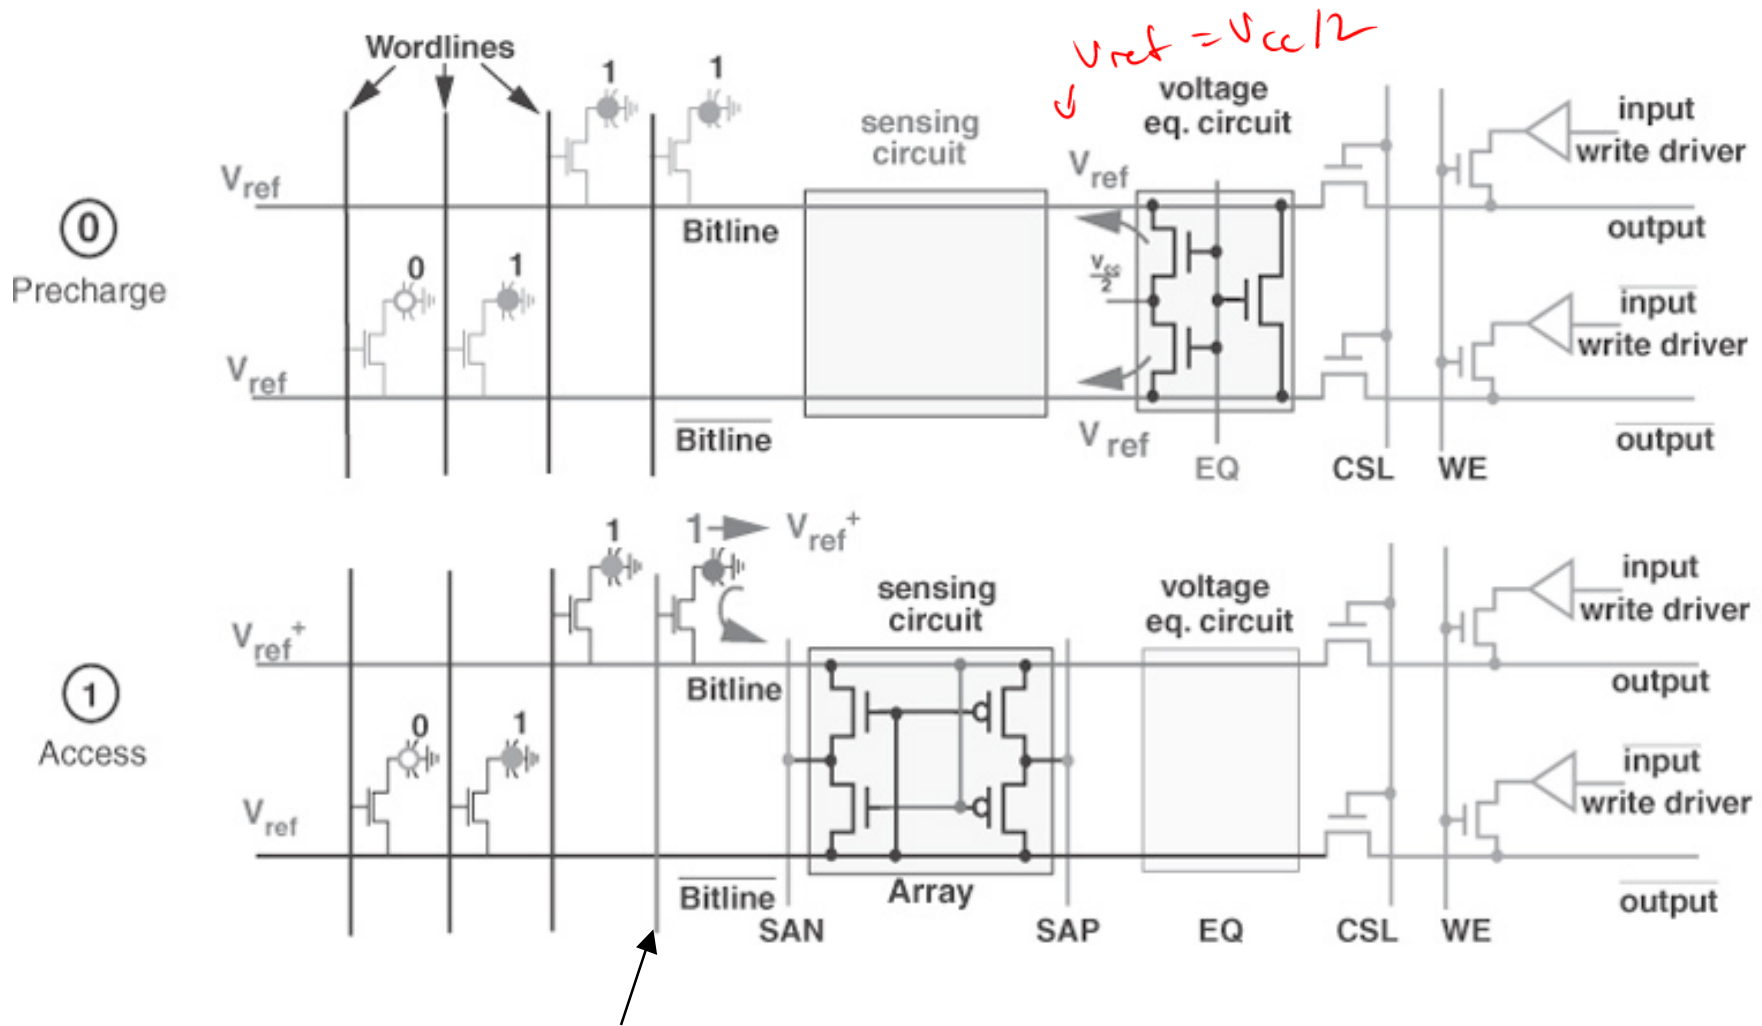
- Charge is restored to the capacitor (if a 1)

⇒ must restore the cap

Basic Sense Amplifier

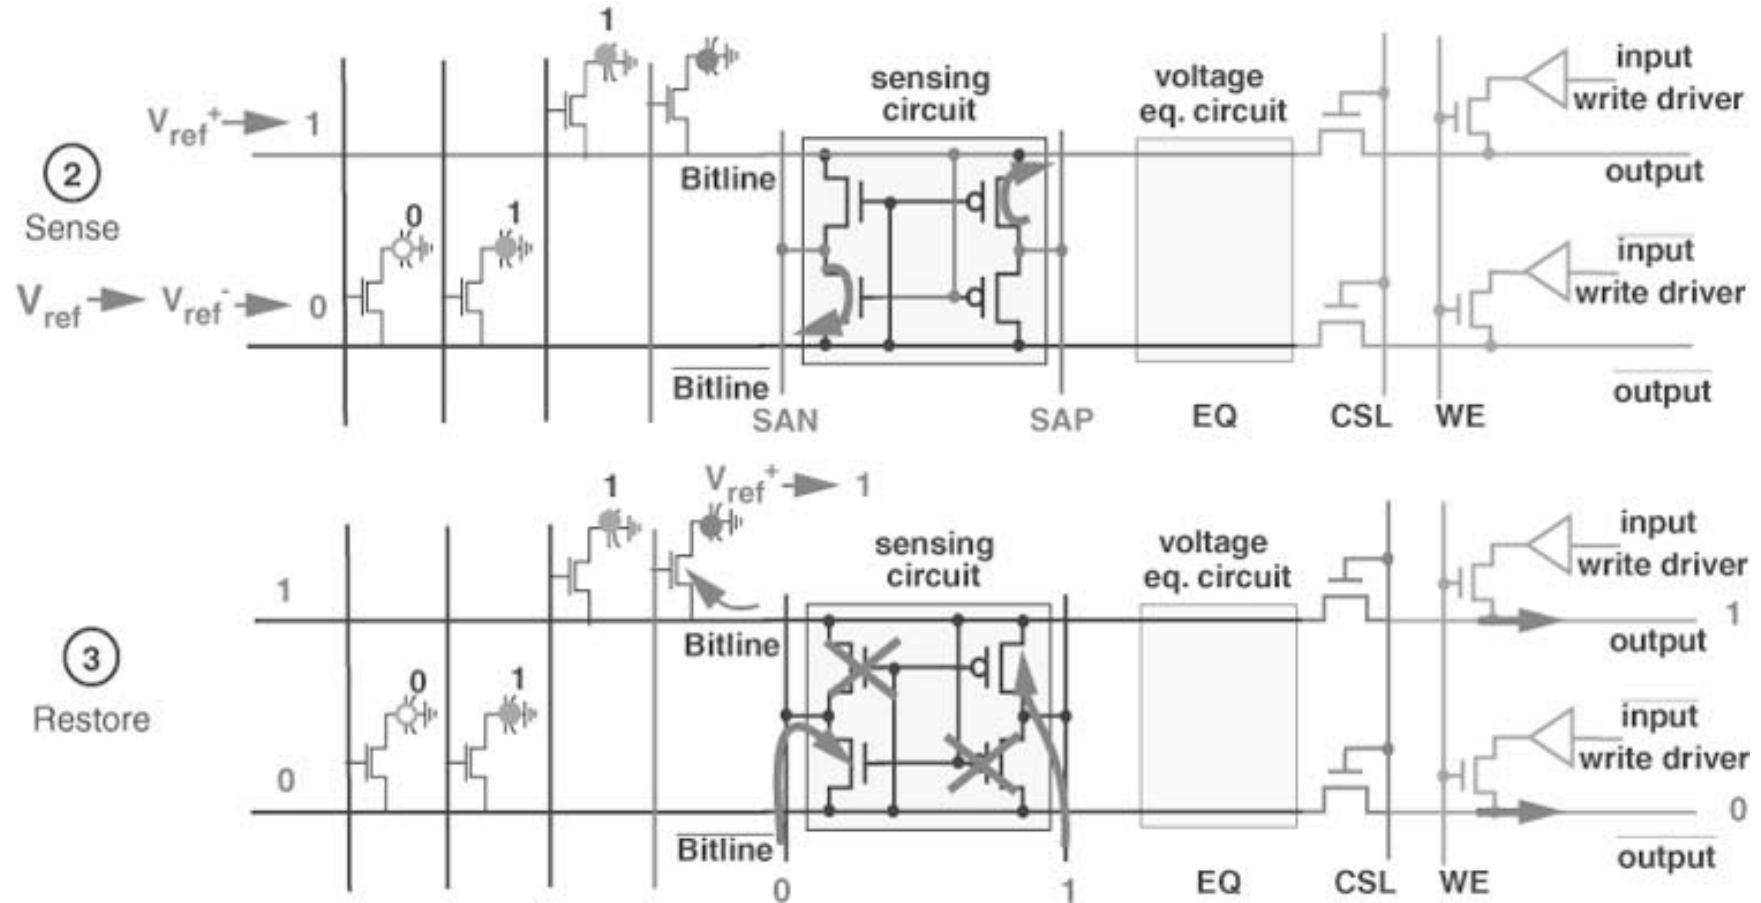


Sense Amplifier Operation

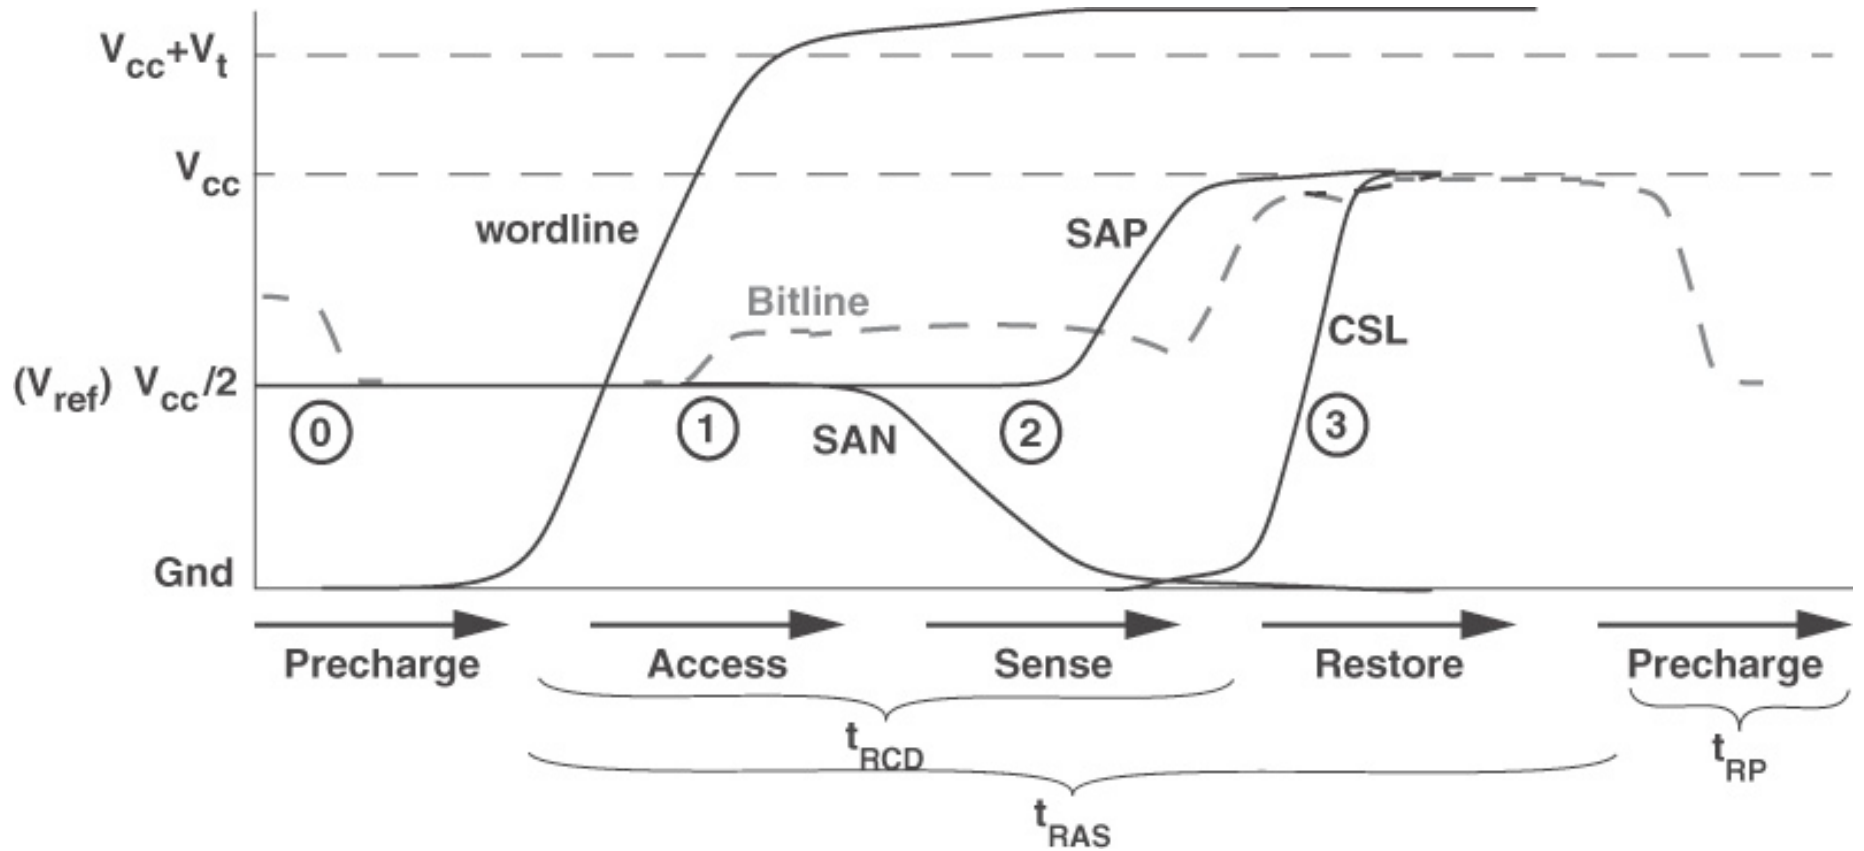


overdriven to $V_{cc} + V_t$
(permits cap to be restored to V_{cc})

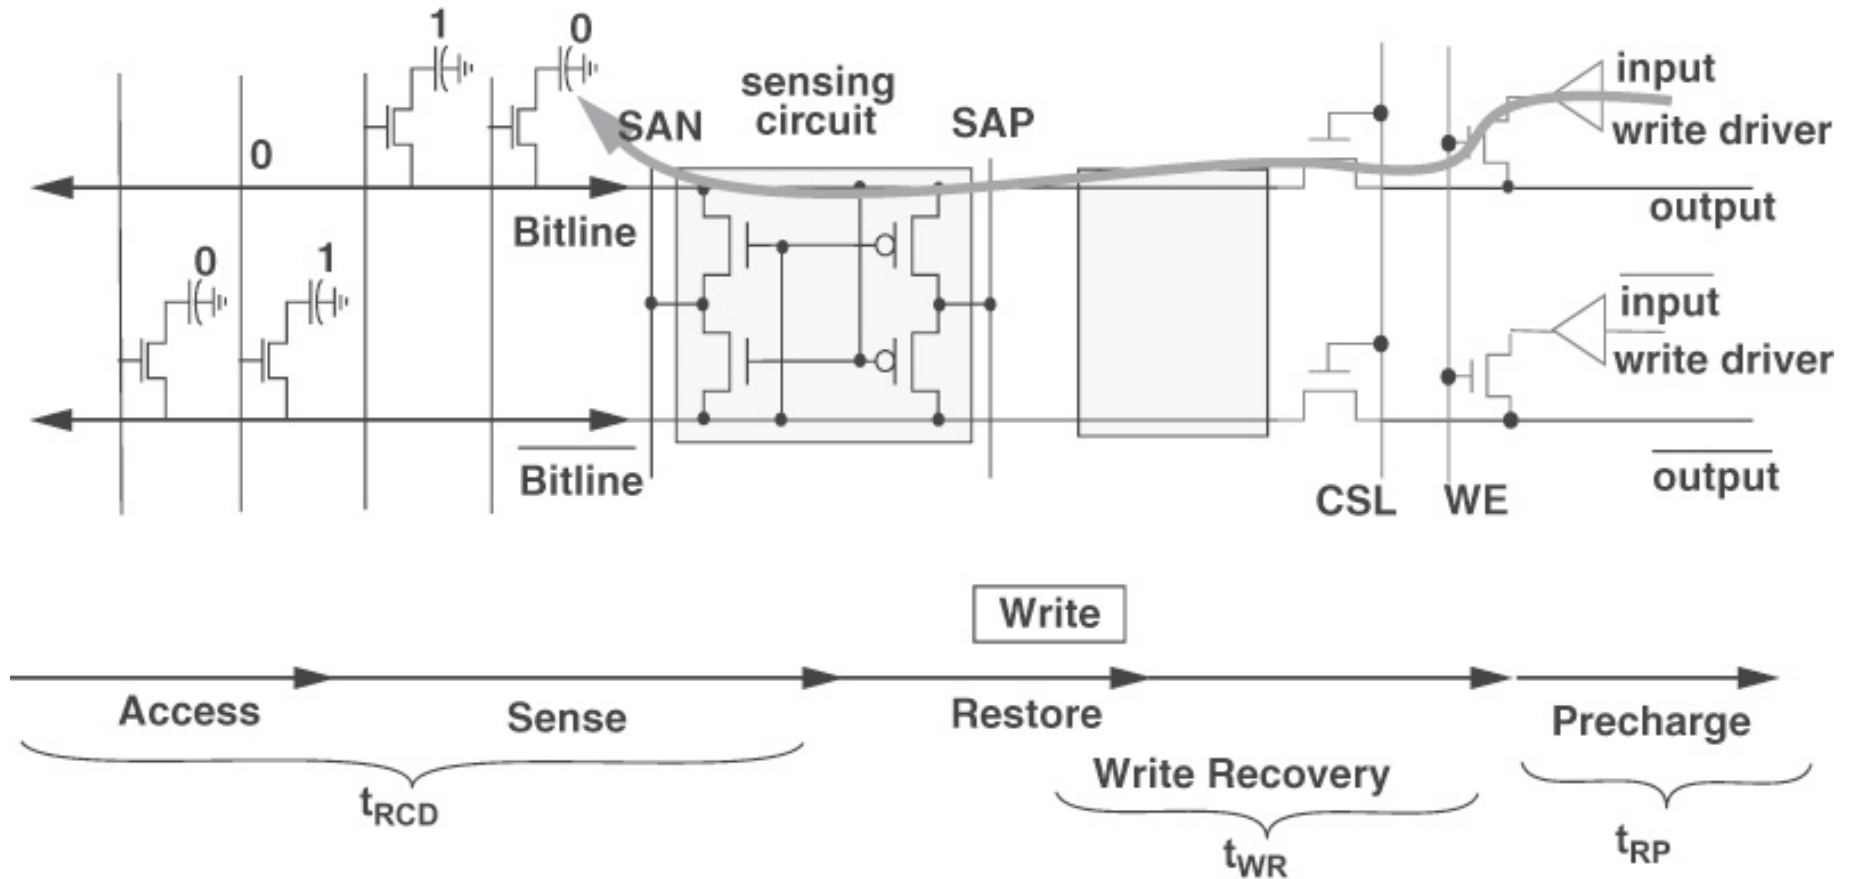
Sense Amplifier Operation




(Simplified) Sense Amp Waveforms



DRAM Write Operation



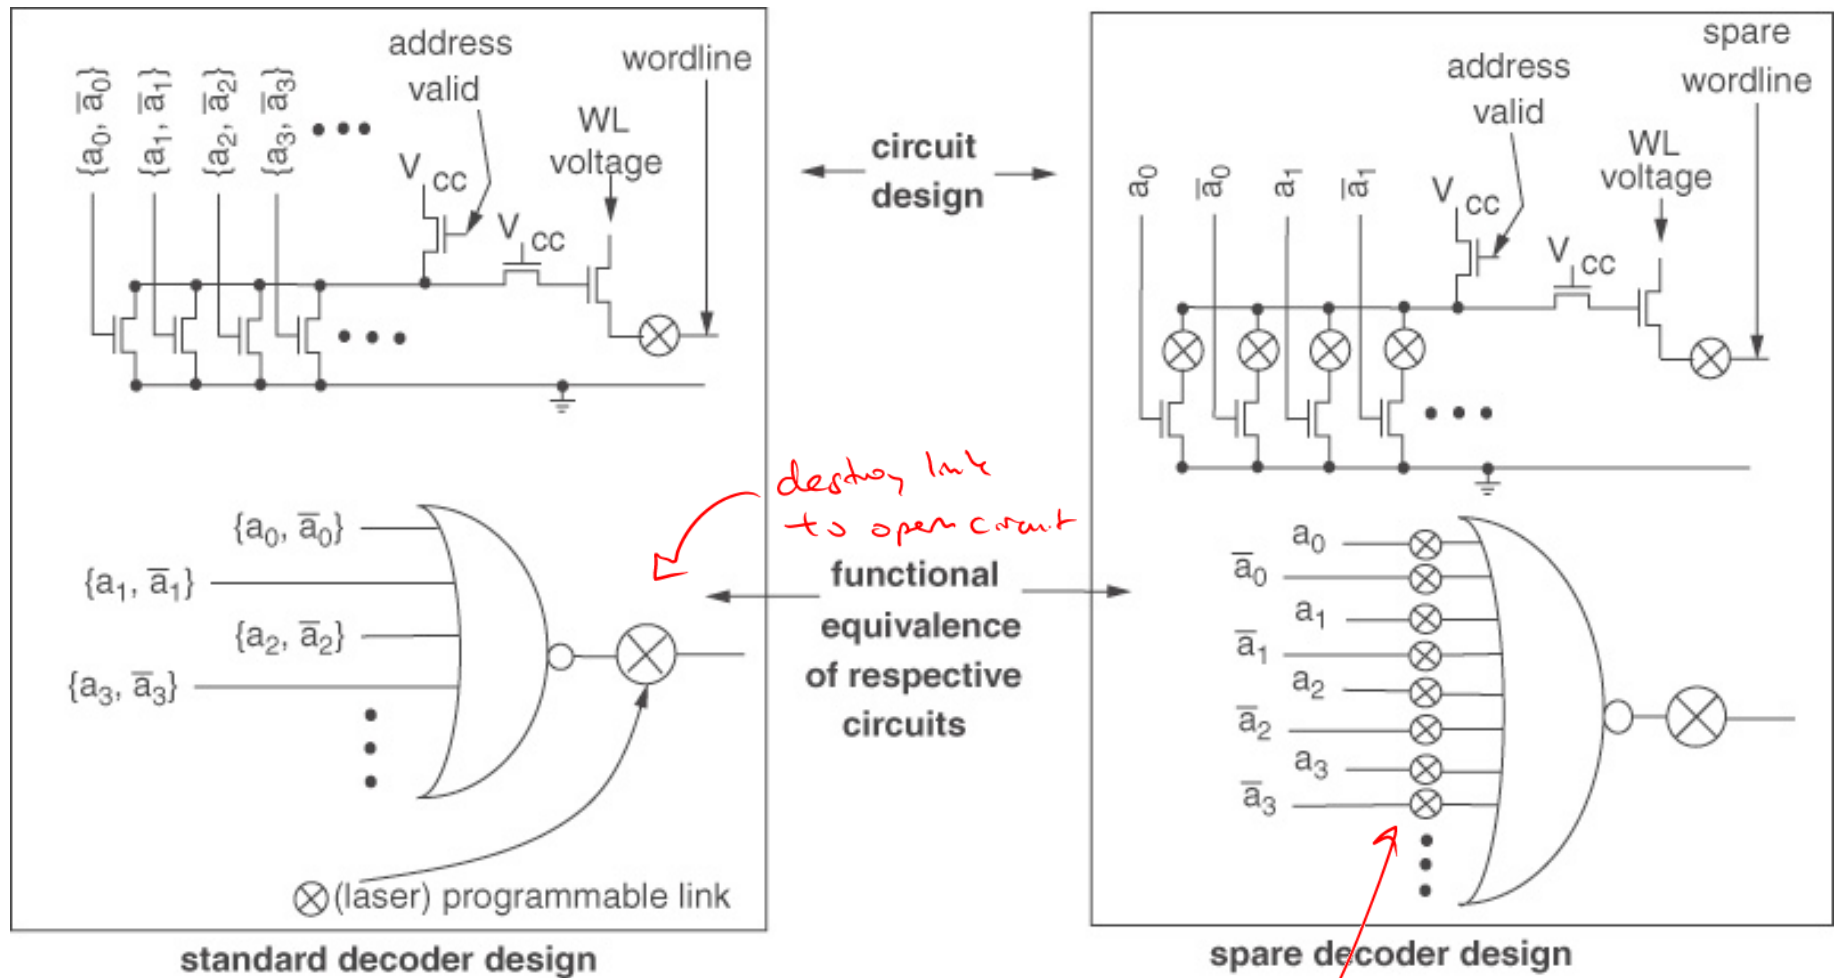
DRAM Row and Column Redundancy



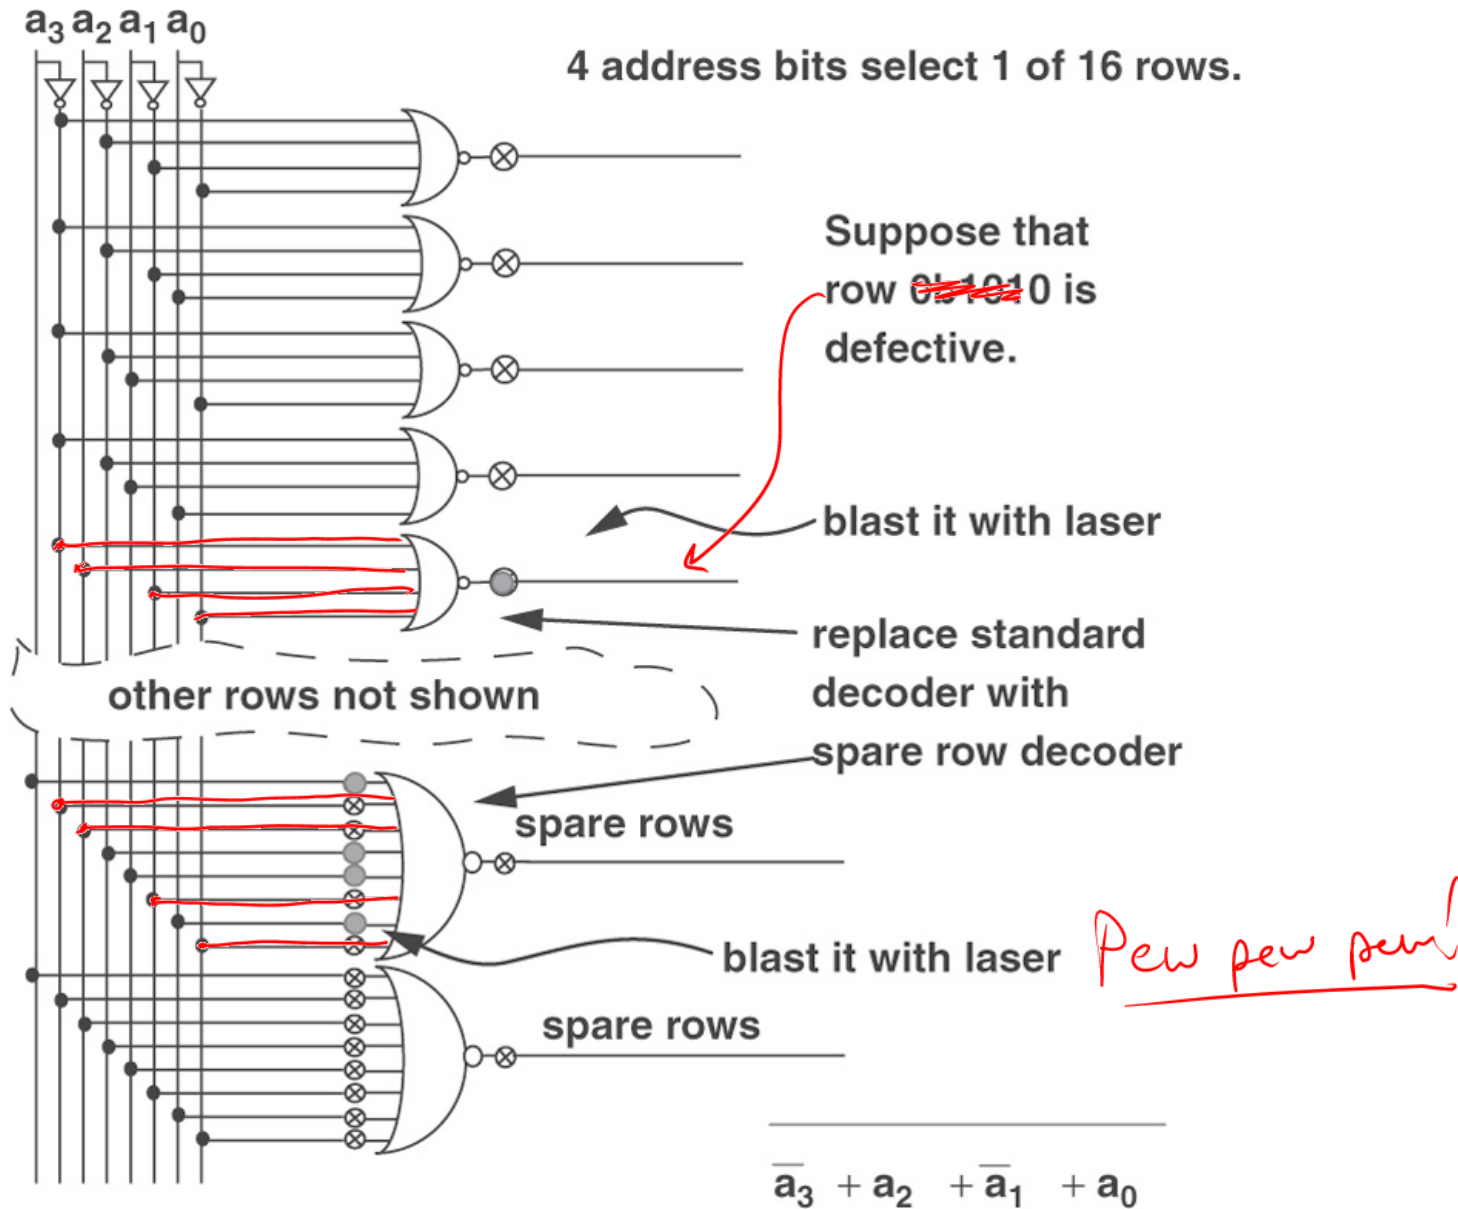
↑
↓
decoders can fix this

has extra rows/columns to address defects

Standard and Spare Decoders



Standard and Spare Decoders



Next Time

DRAM Access Protocol