CS698Y: Modern Memory Systems
Lecture-16 (DRAM Timing Constraints)

Biswabandan Panda

biswap@cse.iitk.ac.in

https://www.cse.iitk.ac.in/users/biswap/CS698Y.html

Accessing a Row

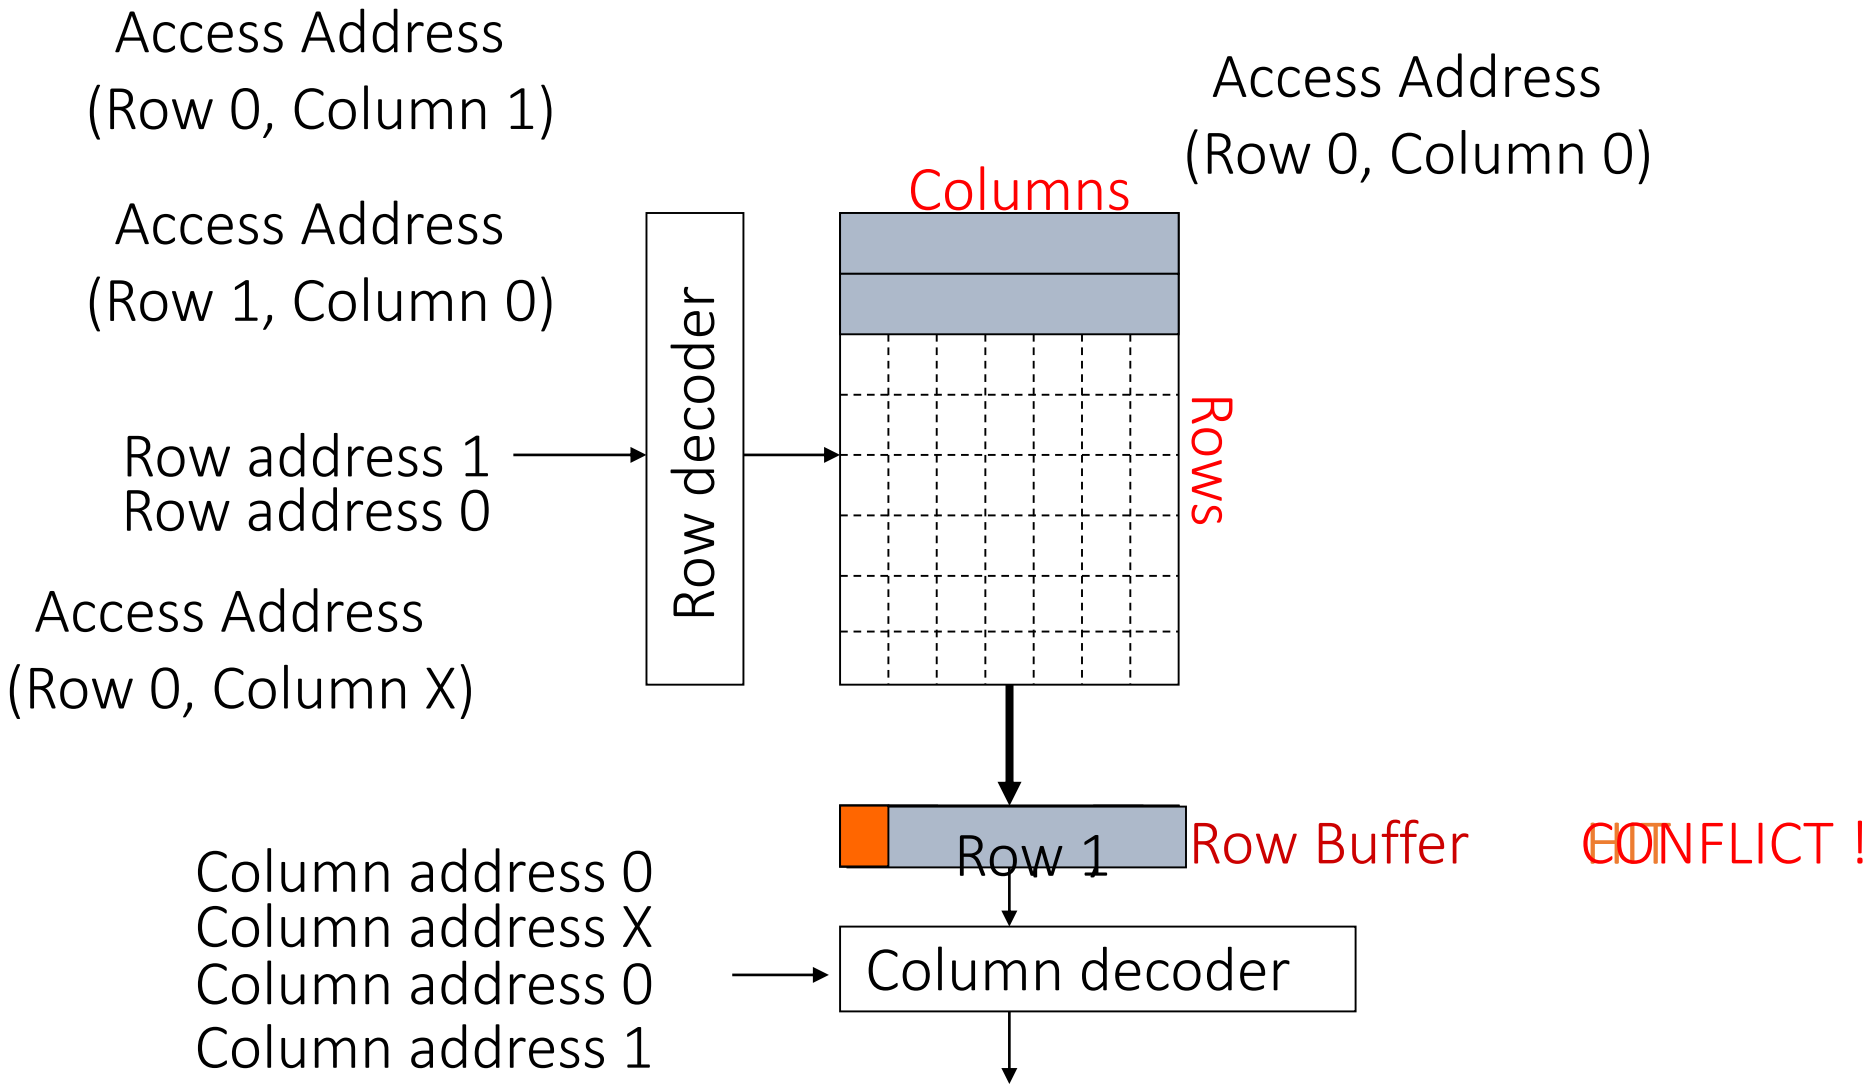

Row Buffer Operations

Access to a closed row (Sequence of Commands):

ACT: Activate the row and place it into row buffer
READ/WRITE: Reads or writes a column
PRECHARGE: Close the row

Access to an open row (Sequence of Commands):

READ/WRITE: Reads or writes a column
PRECHARGE: Close the row

DRAM Timing Constraints

tRAS: Latency for Row Address Strobe

tRP: Latency for PRECHARGE

tRC: Row cycle time. Time gap between two row accesses (tRAS + tRP)

tCAS (tCL): Latency for Column Address Strobe (data from column to bus)

tRCD: Row to Column Delay, Latency of an ACT (gap between row access and data in sense amplifiers)

tCCD: Column to Column Delay

tRRD: Row to Row Delay (between two row ACTs)

tBURST: Burst length (how many bytes read/written)

tRAS

tRAS: Latency for Row Address Strobe

The time interval between row access command and data restoration in a DRAM array. A DRAM bank can not be precharged until at least tRAS time after the previous bank activation

Minimum time a row must open

DRAM Refresh

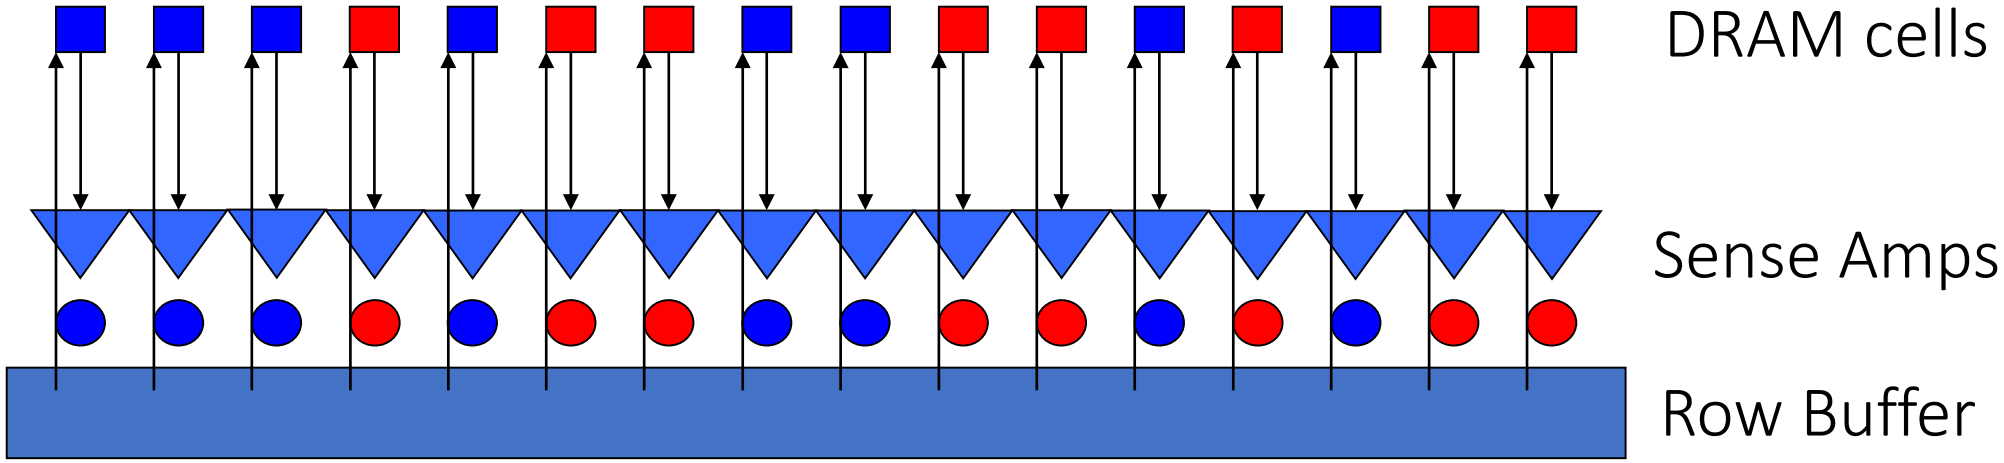

DRAM cells lose contents after a while,
Refresh command refreshes all rows (different avatars like all rows in one bank, all banks)

How to implement refresh?

What is the latency? How frequent? Look at tRFC and tREFI

OPEN/CLOSED Row Policies

OPEN PAGE

After an access:
Keep the page in the row-buffer
Consecutive accesses to the same page : Row-buffer Hit
On an access to different page: Close the row and open the new one

CLOSED PAGE

After an access:
Close the page
Consecutive accesses to different page : Low latency
On an access to different page: No need to close the row

OPEN/CLOSED Row Policies

OPEN :
Exploits spatial and temporal locality
Access patterns ?
Latency is limited tCAS

CLOSED :
Exploits random access pattern
Access patterns ?

Let’s Revisit the Latency

PAGE EMPTY

ACT + CAS

PAGE HIT

CAS

PAGE MISS

PRE + ACT + CAS

Given a DRAM request,
what is the best/worst-case delay ?

Solve It [Courtesy: CS6810, Utah]

CL: 20ns
RP + ACT + CL : 60ns
ACT + CL : 40ns

For the following access stream, estimate the finish times

Req	Time of arrival	Open	Closed
X	0 ns		
Y	10 ns		
X+1	100 ns		
X+2	200 ns		
Y+1	250 ns		
X+3	300 ns		

Note that X, X+1, X+2, X+3 map to the same row and Y, Y+1 map to a different row in the same bank. Ignore bus and queuing latencies. The bank is pre-charged at the start.

Cheat Sheet (Timing constraints)

A=row access
R=col_rd
W=col_wr
P=precharge
F=Refresh
s=same
d=different
a=any

Prev	Next	Rank	Bank	Min. Timing	Notes
A	A	s	s	tRC	
A	A	s	d	tRRD	plus tFAW for 5th RAS same rank
P	A	s	d	tRP	
F	A	s	s	tRFC	
A	R	s	s	tRCD-tAL	tAL=0 unless posted CAS
R	R	s	a	Max(tBURST, tCCD)	tBURST of previous CAS, same rank
R	R	d	a	tBURST+tRTRS	tBURST prev. CAS diff. rank
W	R	s	a	tCWD+tBURST+tWTR	tBURST prev CASW same rank
W	R	d	a	tCWD+tBURST+tRTRS-tCAS	tBURST prev CASW diff rank
A	W	s	s	tRCD-tAL	
R	W	a	a	tCAS+tBURST+tRTRS-tCWD	tBURST prev. CAS any rank
W	W	s	a	Max(tBURST, tCCD)	tBURST prev CASW same rank
W	W	d	a	tBURST+tOST	tBURST prev CASW diff rank
A	P	s	s	tRAS	
R	P	s	s	tAL+tBURST+ tRTP-tCCD	tBURST of previous CAS, same rank
W	P	s	s	tAL+tCWD+tBURST+tWR	tBURST prev CASW same rank
F	F	s	a	tRFC	
P	F	s	a	tRFC	

JEDEC Standard

Standard name	DRAM cell array clock (MHz)	Cycle time (ns)	I/O bus clock (MHz)	Data rate (MT/s)	Module name	Peak transfer rate (MB/s)	Timings (CL-tRCD-tRP)	CAS Latency (ns)
DDR3-800D DDR3-800E	100	10	400	800	PC3-6400	6400	5-5-5 6-6-6	12.5 15
DDR3-1066E DDR3-1066F DDR3-1066G	133.33	7.5	533.33	1066.67	PC3-8500	8533.33	6-6-6 7-7-7 8-8-8	11.25 13.125 15
DDR3-1333F* DDR3-1333G DDR3-1333H DDR3-1333J*	166.67	6	666.67	1333.33	PC3-10600	10666.67	7-7-7 8-8-8 9-9-9 10-10-10	10.5 12 13.5 15
DDR3-1600G* DDR3-1600H DDR3-1600J DDR3-1600K	200	5	800	1600	PC3-12800	12800	8-8-8 9-9-9 10-10-10 11-11-11	10 11.25 12.5 13.75
DDR3-1866J* DDR3-1866K DDR3-1866L DDR3-1866M*	233.33	4.286	933.33	1866.67	PC3-14900	14933.33	10-10-10 11-11-11 12-12-12 13-13-13	10.56 11.786 12.857 13.929

Some More Events

(Refer Chapter 11 of Memory Systems by Bruce Jacob)
Try to understand TABLE 11.1
Find latencies for the following

Consecutive reads/writes to same rank

Consecutive reads to different ranks

Consecutive reads to different rows of same bank

Consecutive reads to different banks

Conflicts of Interest