Lecture-3
(DRAM organization: 10K feet view)
CS422-Spring 2019

Biswa@CSE-IITK

Why DRAM?

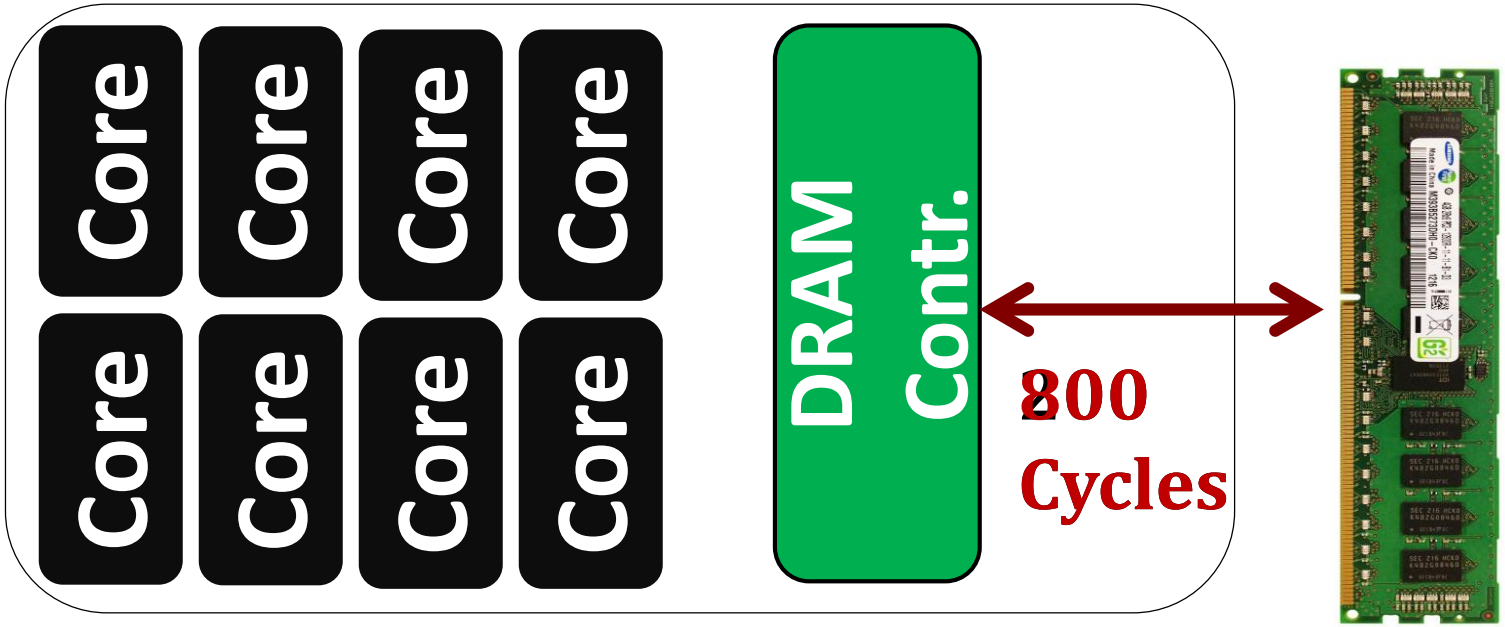

Core count doubling every 2 years

Memory wall = Latency wall + Bandwidth wall

DRAM Organization

Channel

DIMM

Rank

Chip

Bank

Row

Column

Rank 1 with 8 chips

Rank 0 with 8 chips

Ranks, Banks, Rows, and Columns

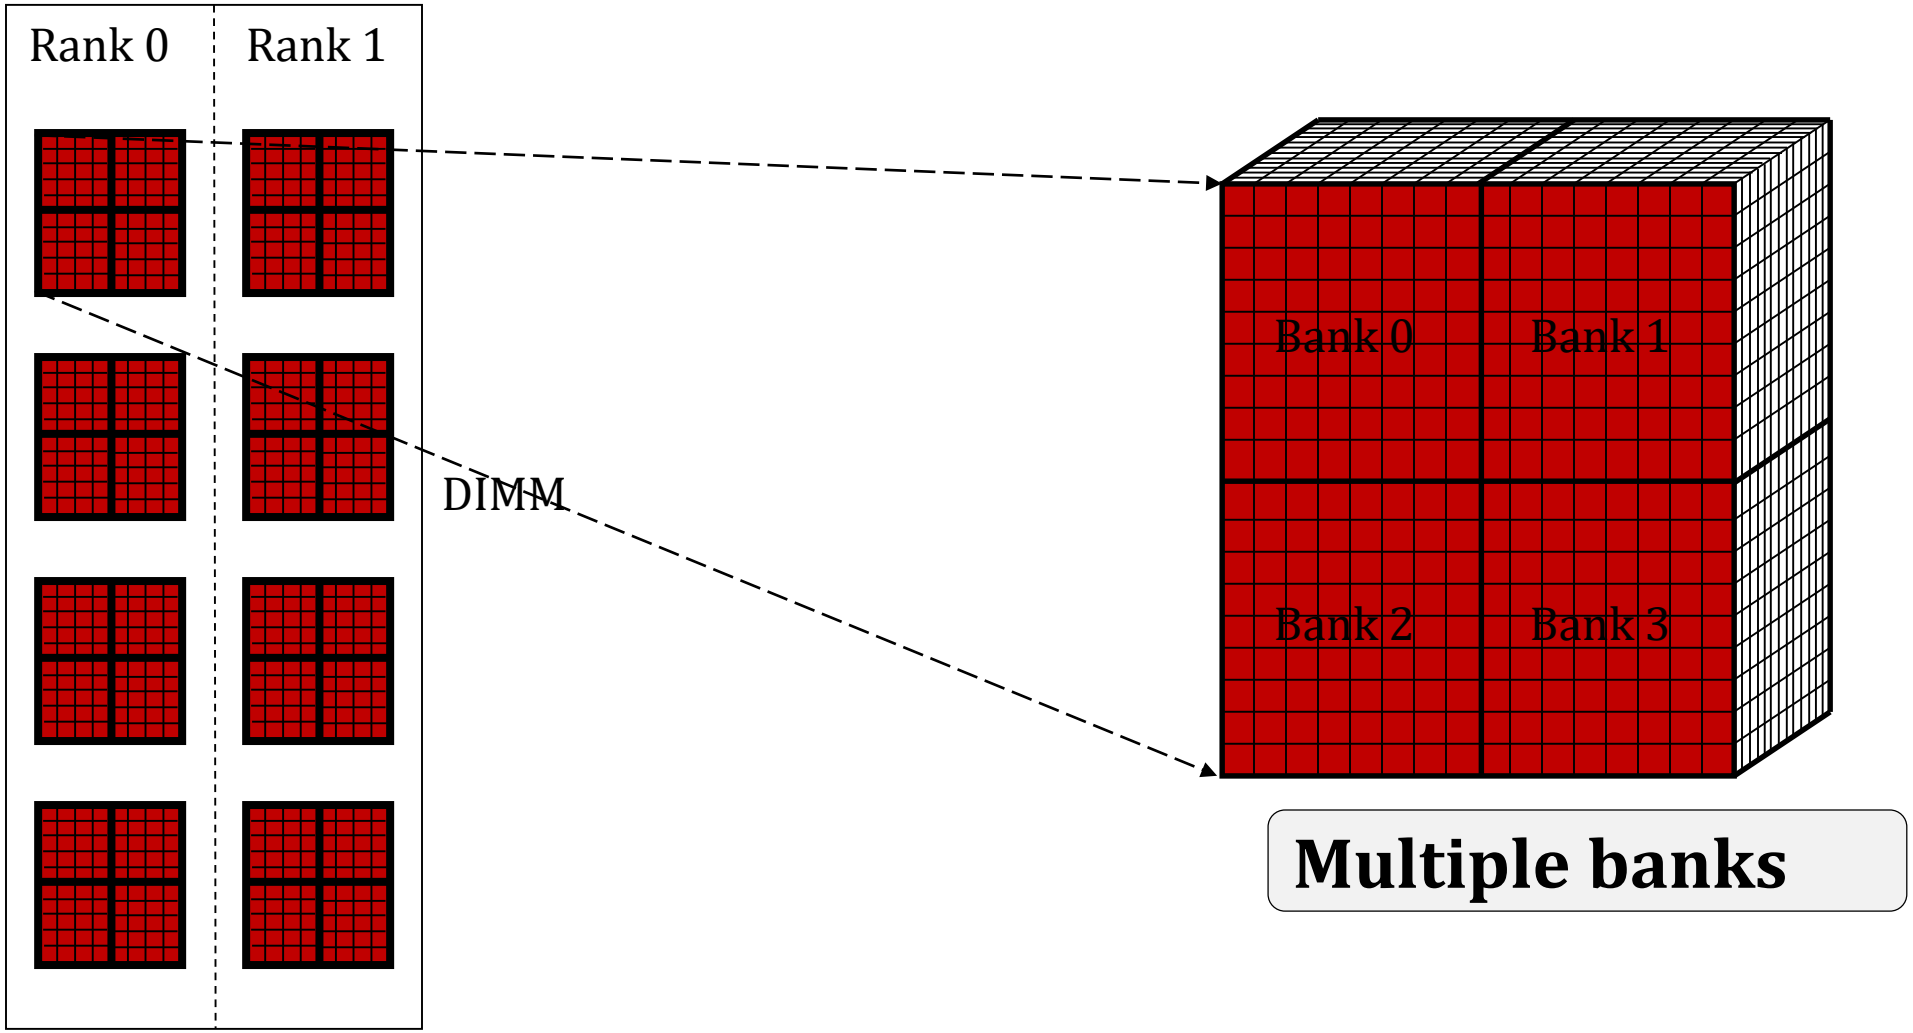

Ranks, Banks, Rows, and Columns

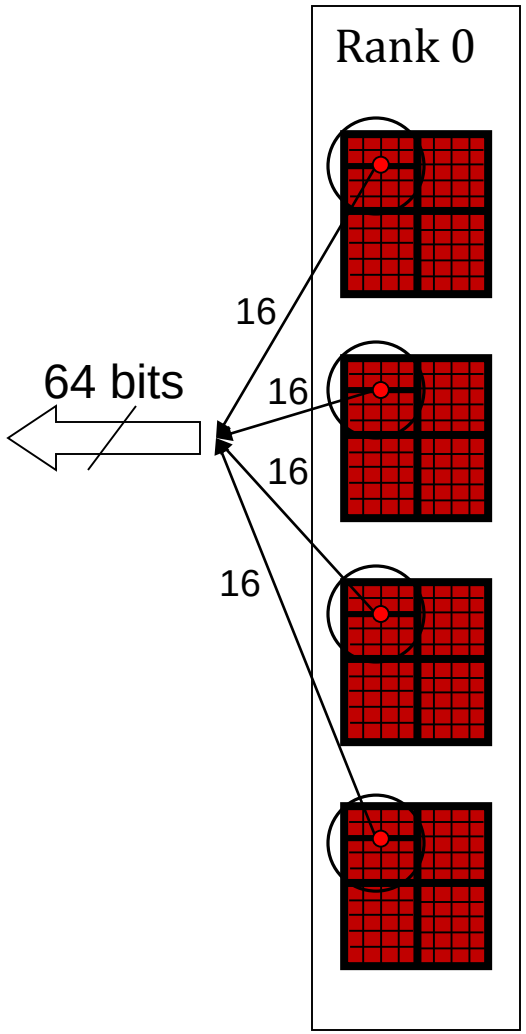

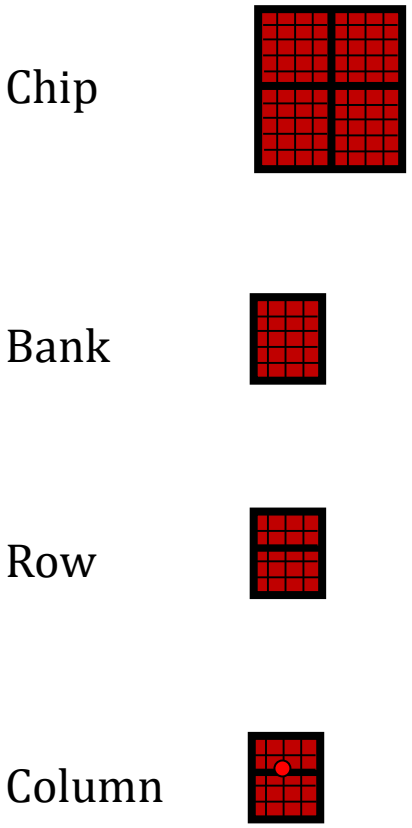

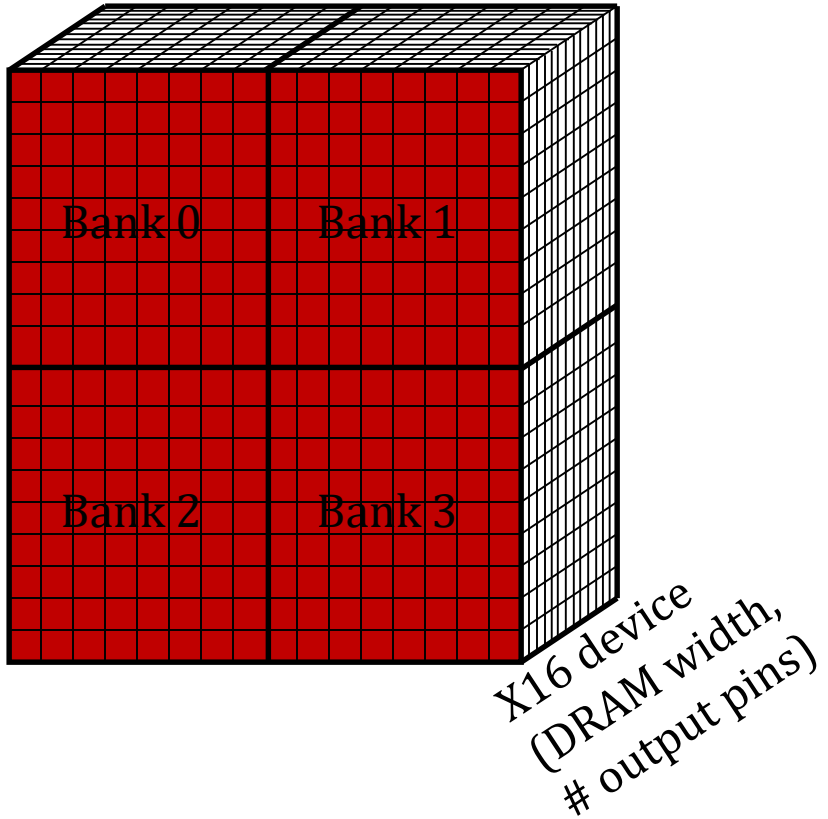

16-bit interface: 16 bits from each chip in one go

Let’s Dig Deep

Each rank has 64-bit wide data bus

If a rank is of width x8 then # DRAM chips ??

What about x4, # DRAM chips ??

If a rank is of width x8 then # DRAM chips ?? 8

What about x4, # DRAM chips ?? 16

Where is Your 1 bit?

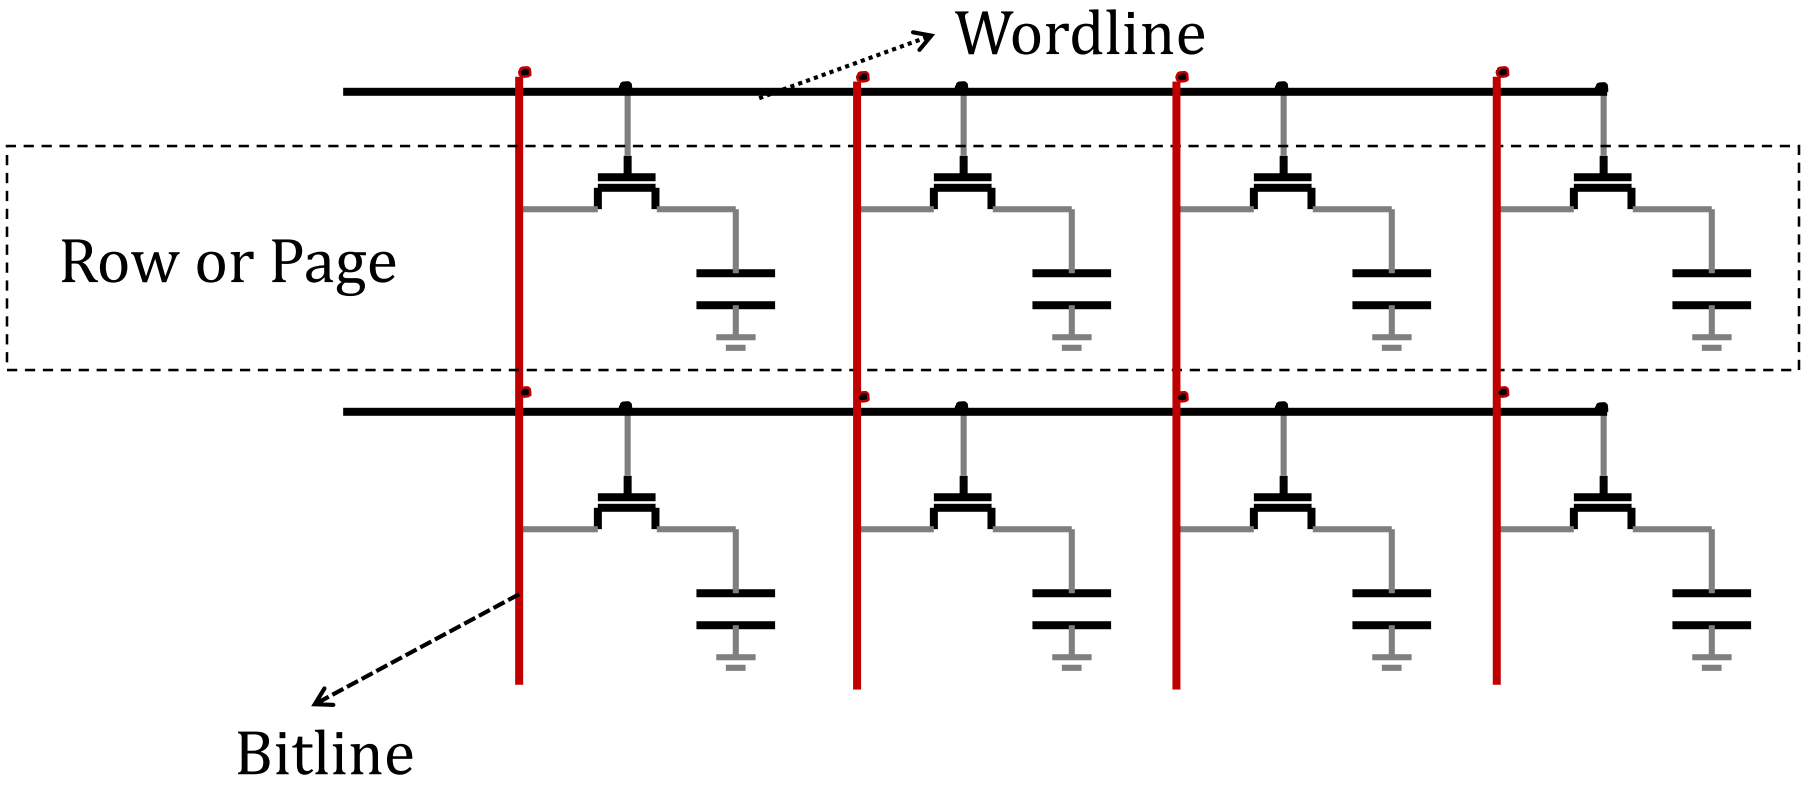

Row (page) and Row buffer (Sense Amplifier)

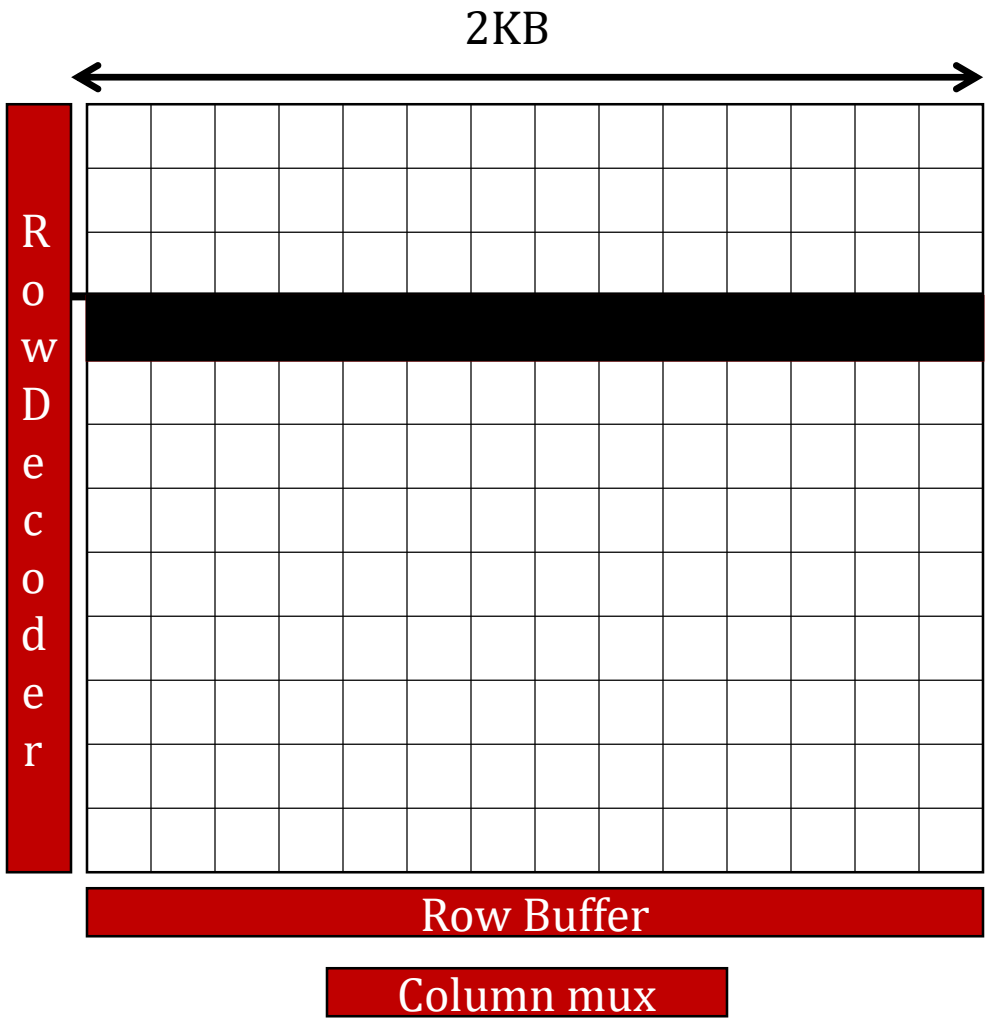

Logically

Each bank has a row buffer

Stores the last used row

Bank: Collection of DRAM Arrays

- DRAM Width
  - x4 device
  - x8 device
- Other possible widths
  - x16
  - x32
  - x48
  - x72

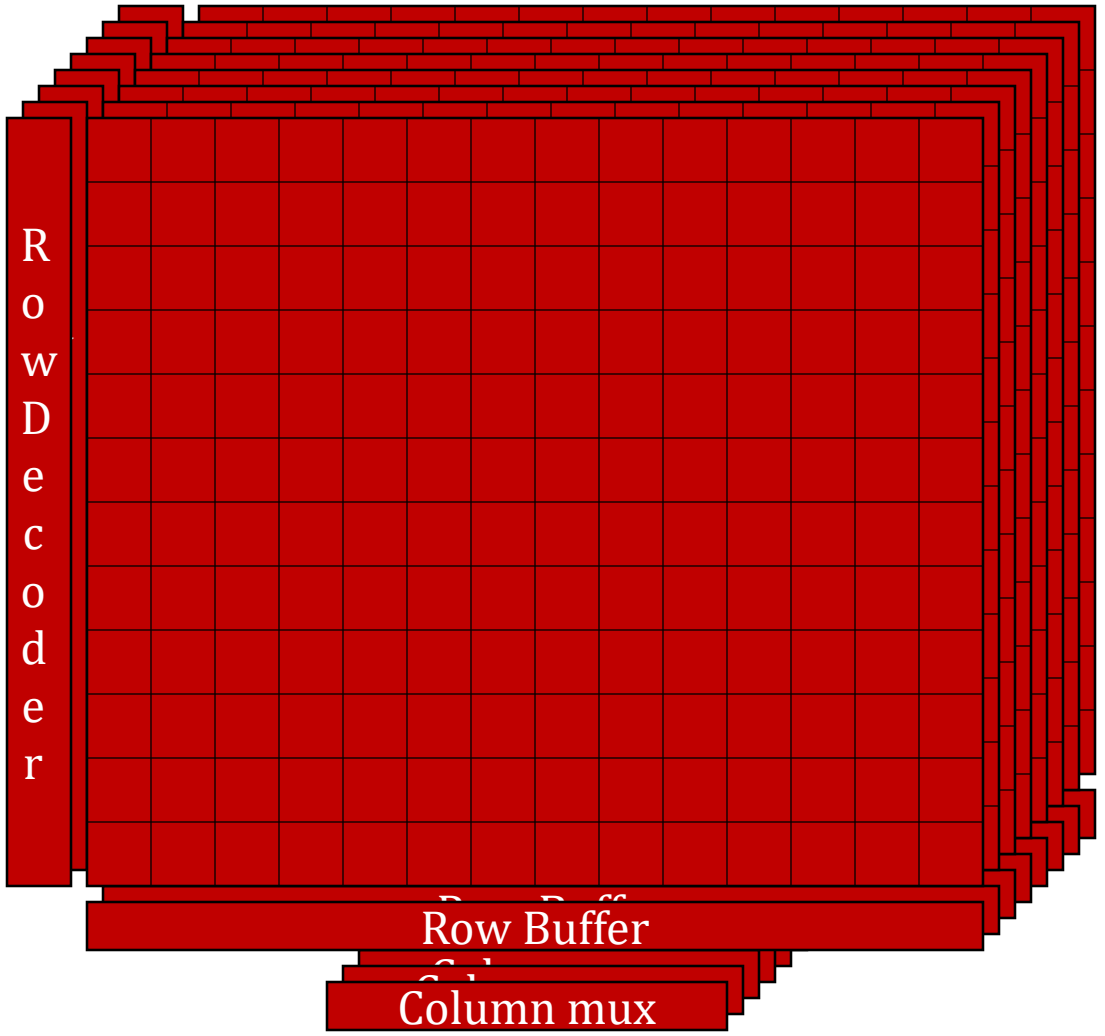

An Example – 4GB DIMM

2Gb * 8 DRAM Chips (one side of the rank)

Total 16 chips + 2 chips for ECC (for both the ranks)

64 bit + 8 bit ECC interface (72 bit wide DIMM)

Transferring a 64B cache line will take 8 transfers of 8B each

8B will come from 8 chips (8 bits from one chip)

1 bit from each DRAM array assuming 8 DRAM arrays per bank

Another View

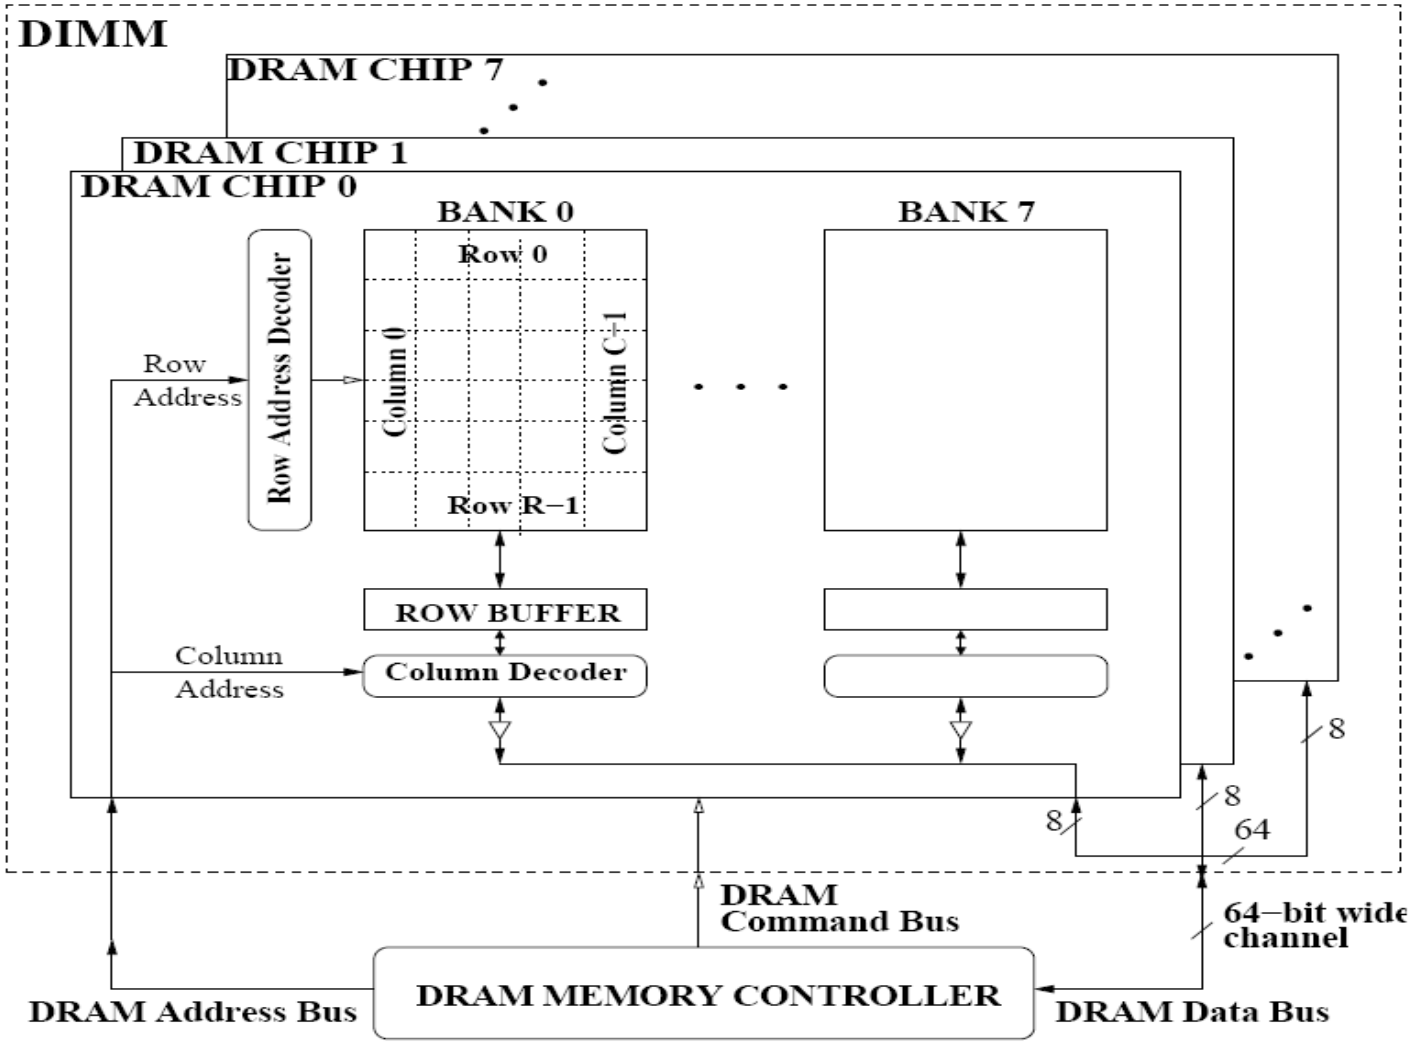

DRAM Channels

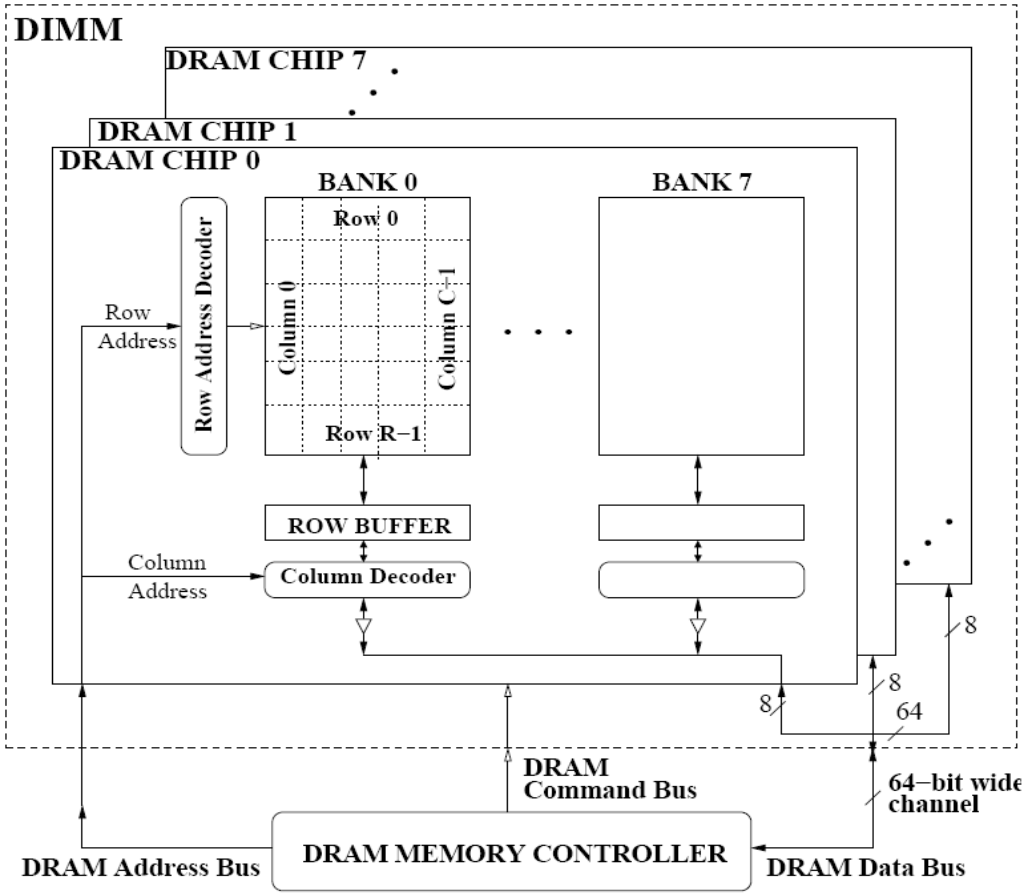

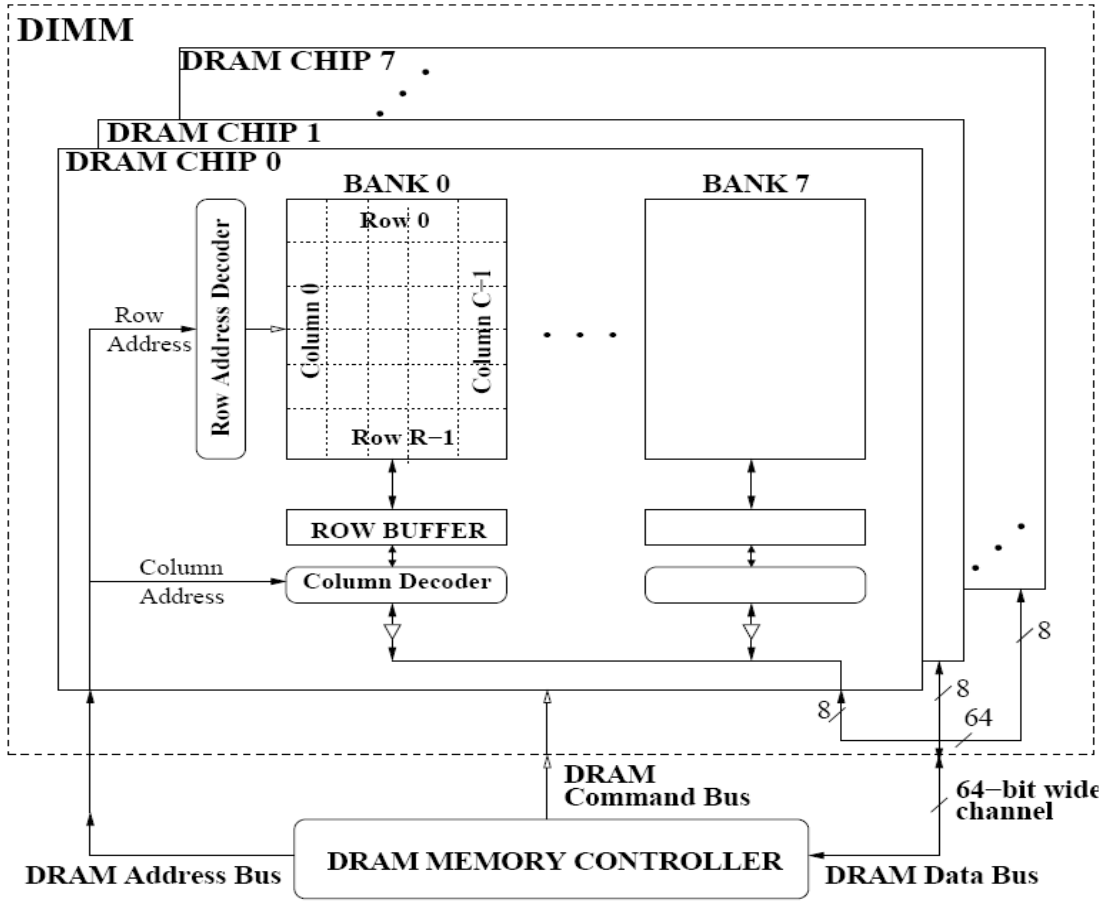

2 channels: 1 channel per controller

DRAM to LLC

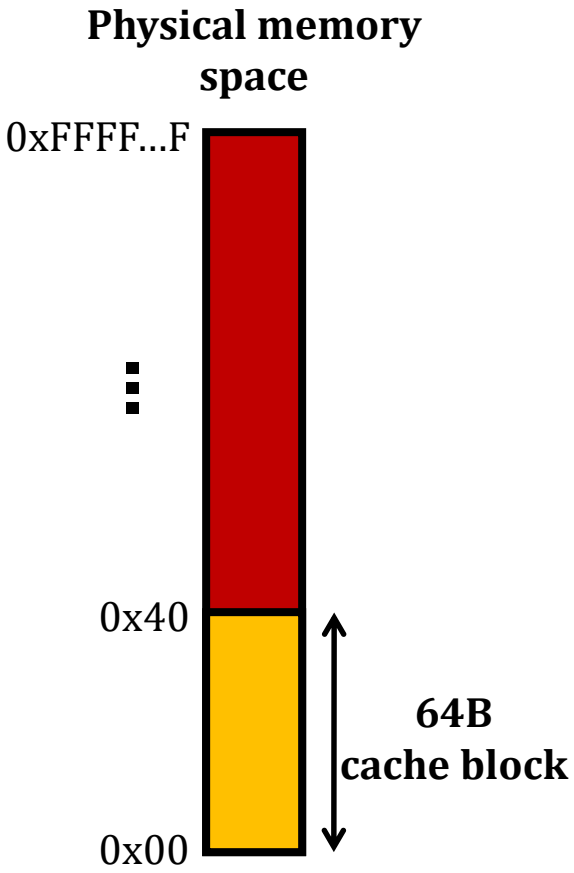
Physical memory space
0xFFFF…F
0x40
0x00
64B cache block

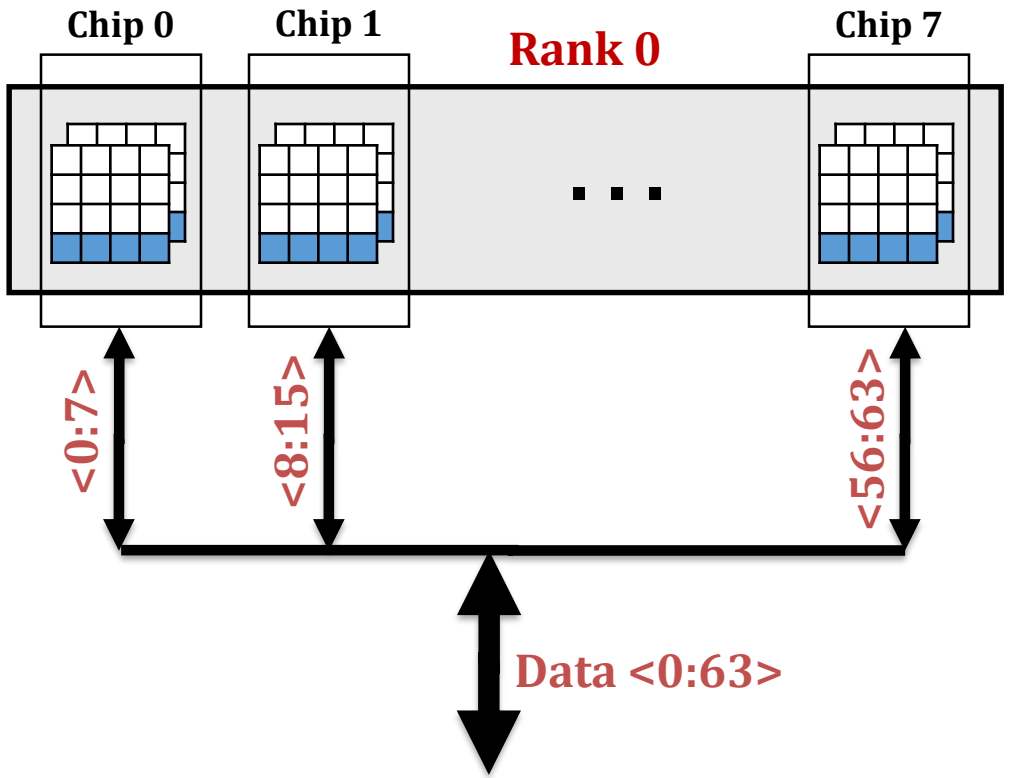
Chip 0
Chip 1
Rank 0
Chip 7
<0:7>
<8:15>
<56:63>
Data <0:63>

DRAM to LLC [CMU 18-447]

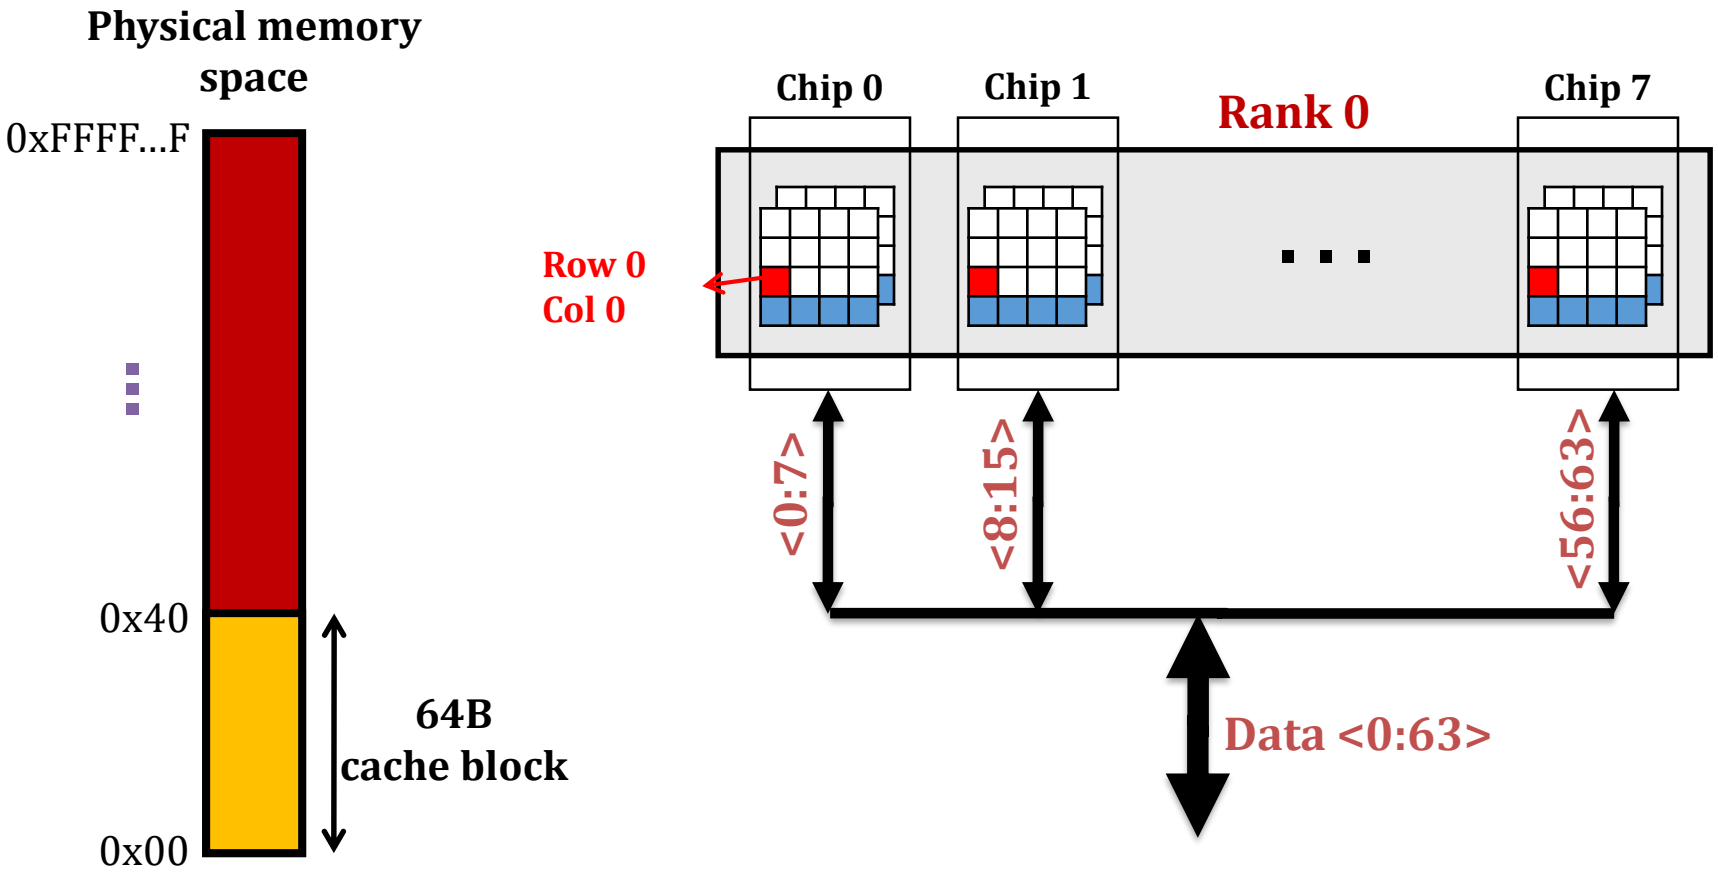

DRAM to LLC [CMU 18-447]

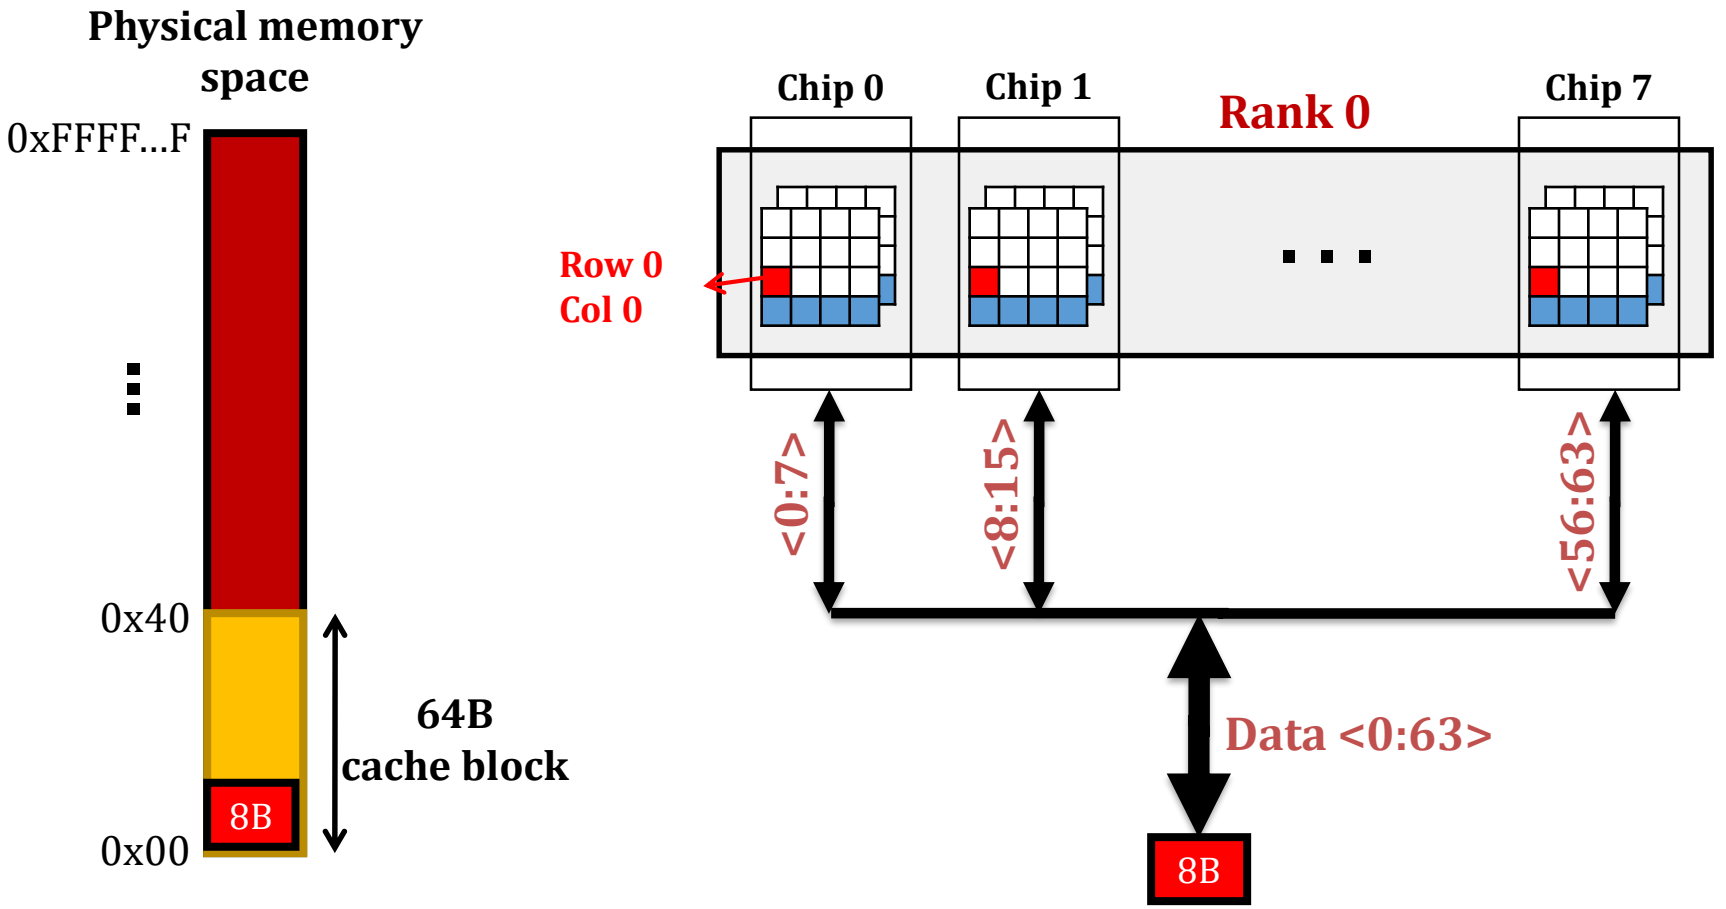

DRAM to LLC

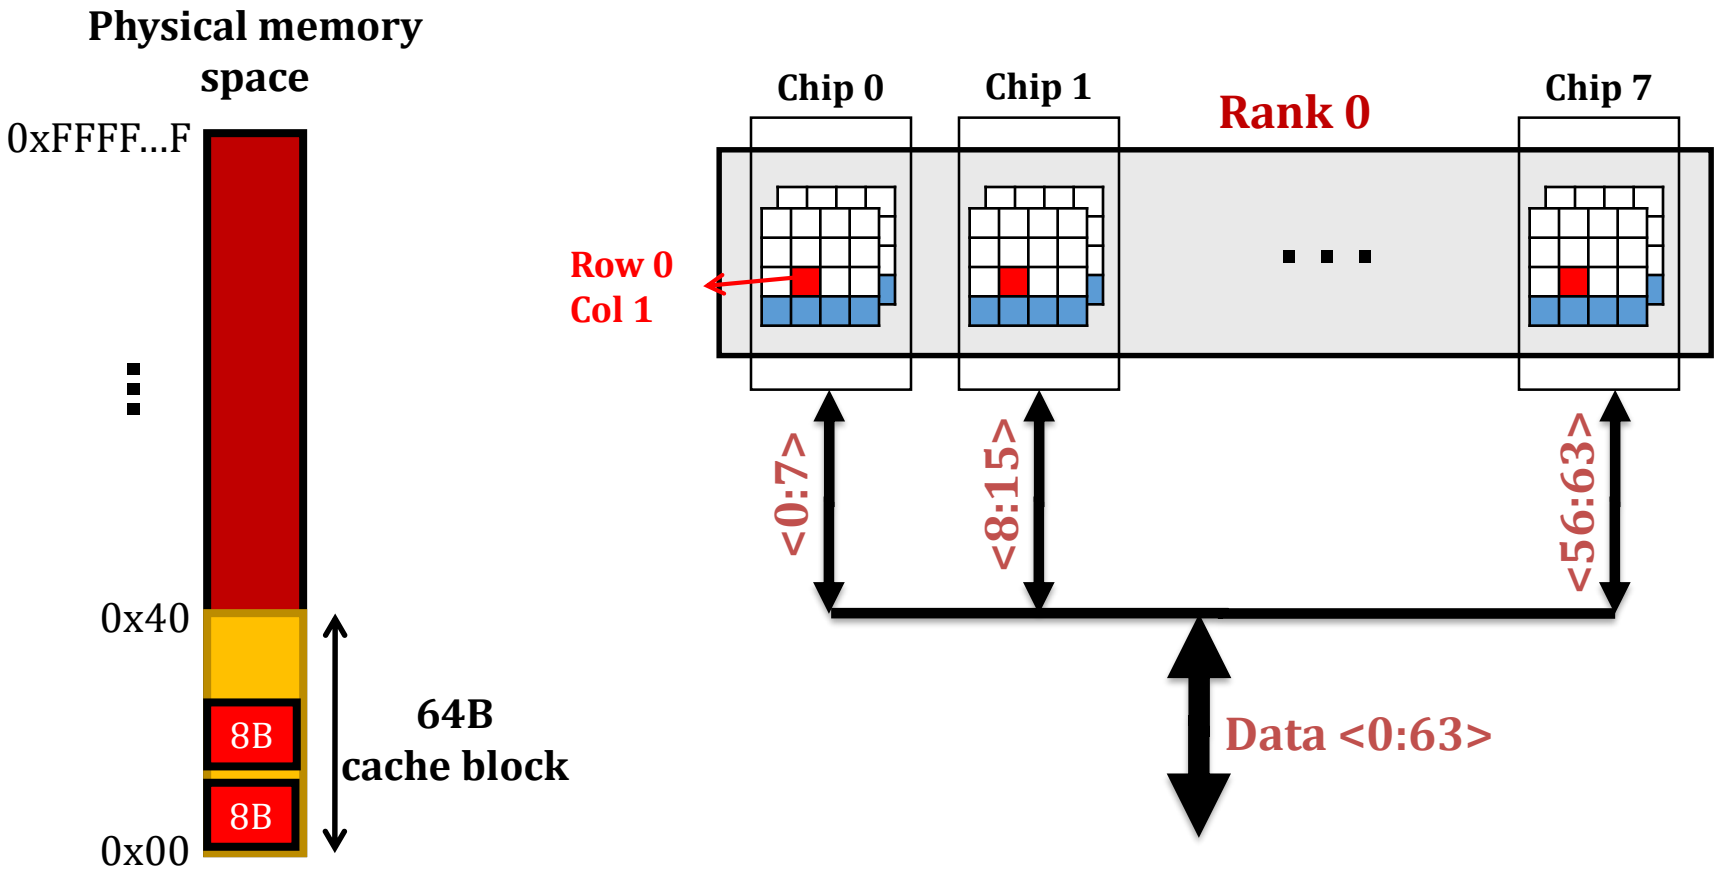

8 cycles (DRAM IO): 1 cycle transfers 8 bytes from a column

DRAM Address Mapping (1 Channel)

2GB DRAM, 8 Banks, 16K rows, 2K Columns per bank

Cache Interleaving: Consecutive cache blocks in consecutive banks

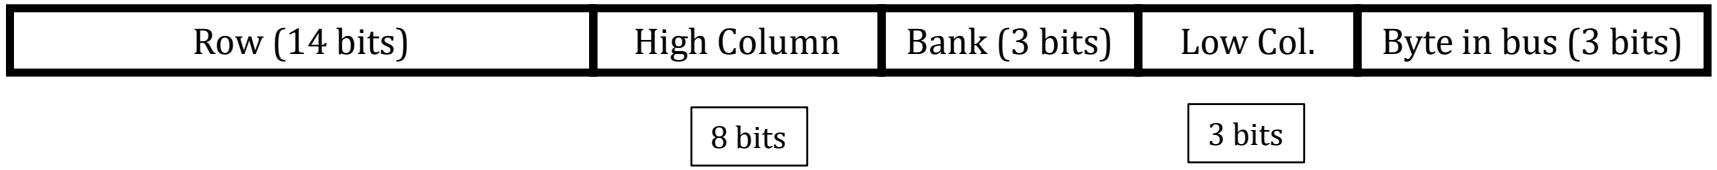

Row Interleaving: Consecutive rows in consecutive banks

What about Multiple Channels?