CASPER

CS773: Computer Architecture for Performance and Security

Lecture 5: It's the Memory Stupid

https://www.cse.iitb.ac.in/~biswa/

DIMM

Channel

DIMM

Rank

Chip

Bank

Row

Column

Rank 1 with 8 chips

Rank 0 with 8 chips

Row Buffer

Access Address:
(Row 0, Column 0)
(Row 0, Column 1)
(Row 0, Column 85)
(Row 1, Column 0)

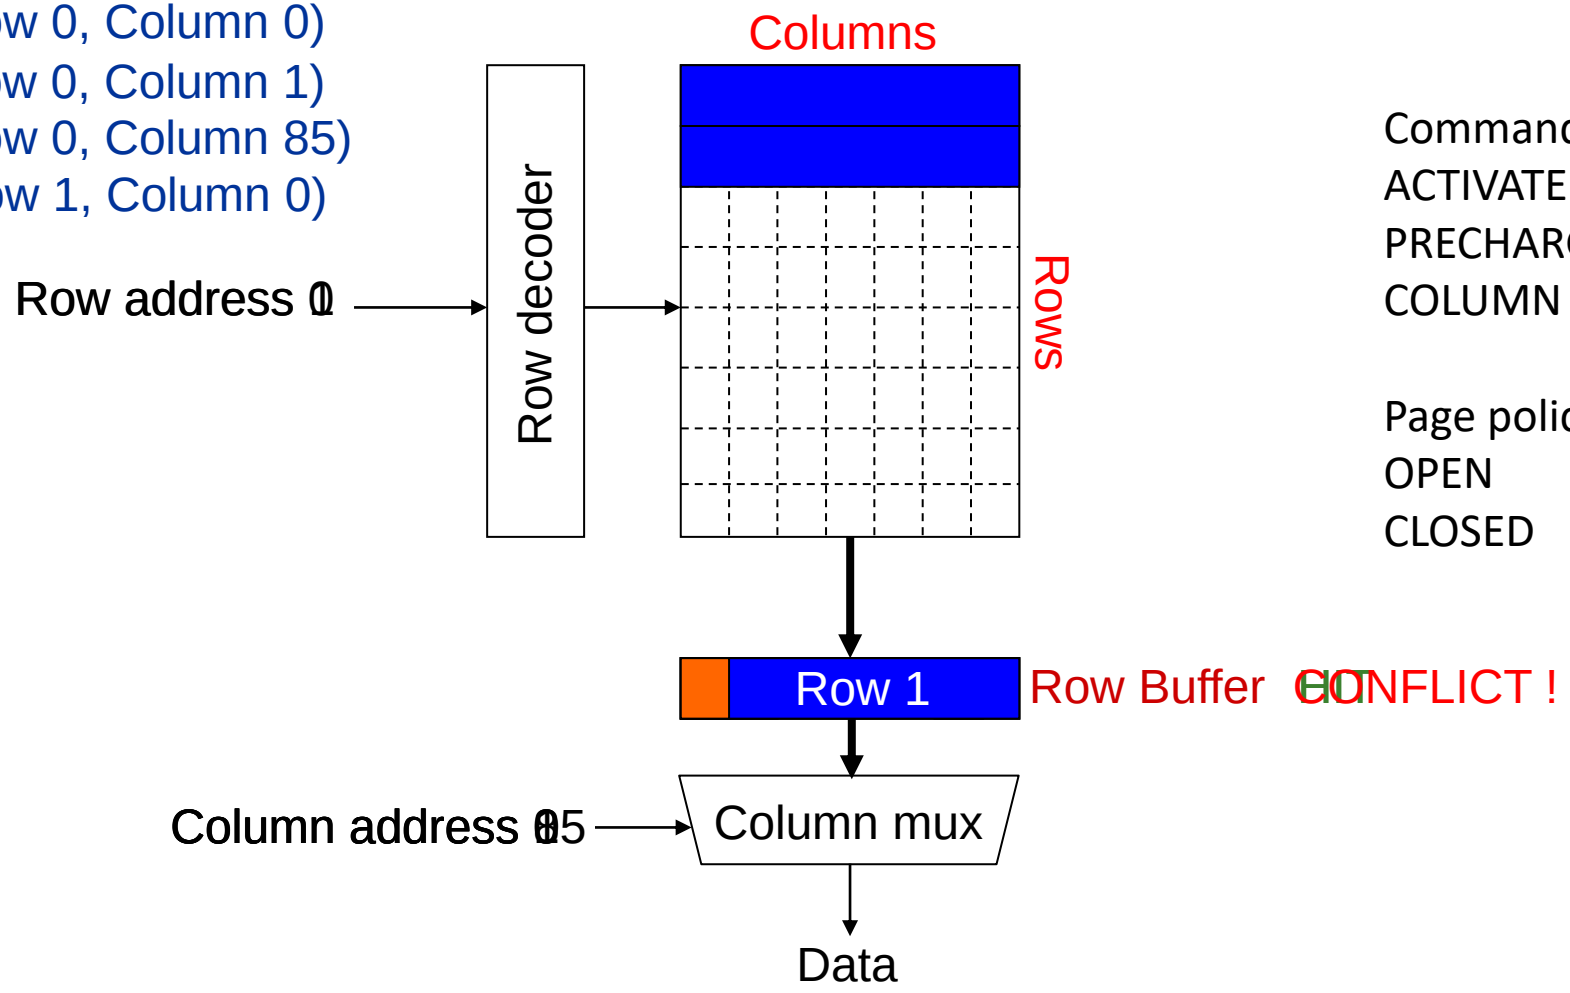

Commands of interest:
ACTIVATE
PRECHARGE
COLUMN READ

Page policy:
OPEN
CLOSED

DRAM

- Capacity
- Latency
- Bandwidth
- Power
- Reliability
- Security
- ..

DRAM Addressing

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

Cache Interleaving: Consecutive cache blocks in consecutive banks

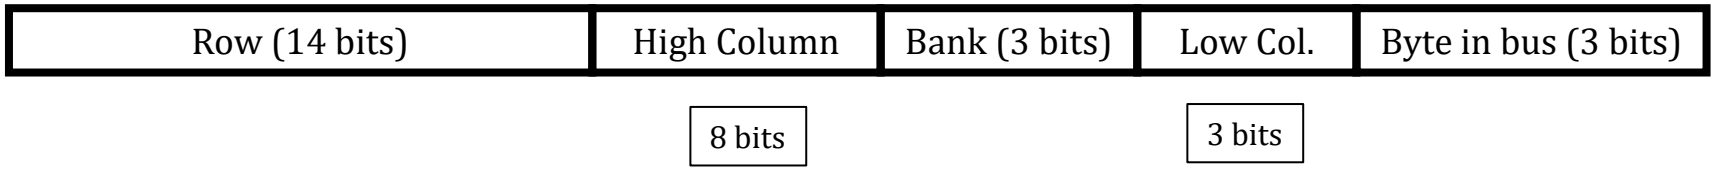

Row Interleaving: Consecutive rows in consecutive banks

DRAM Controller

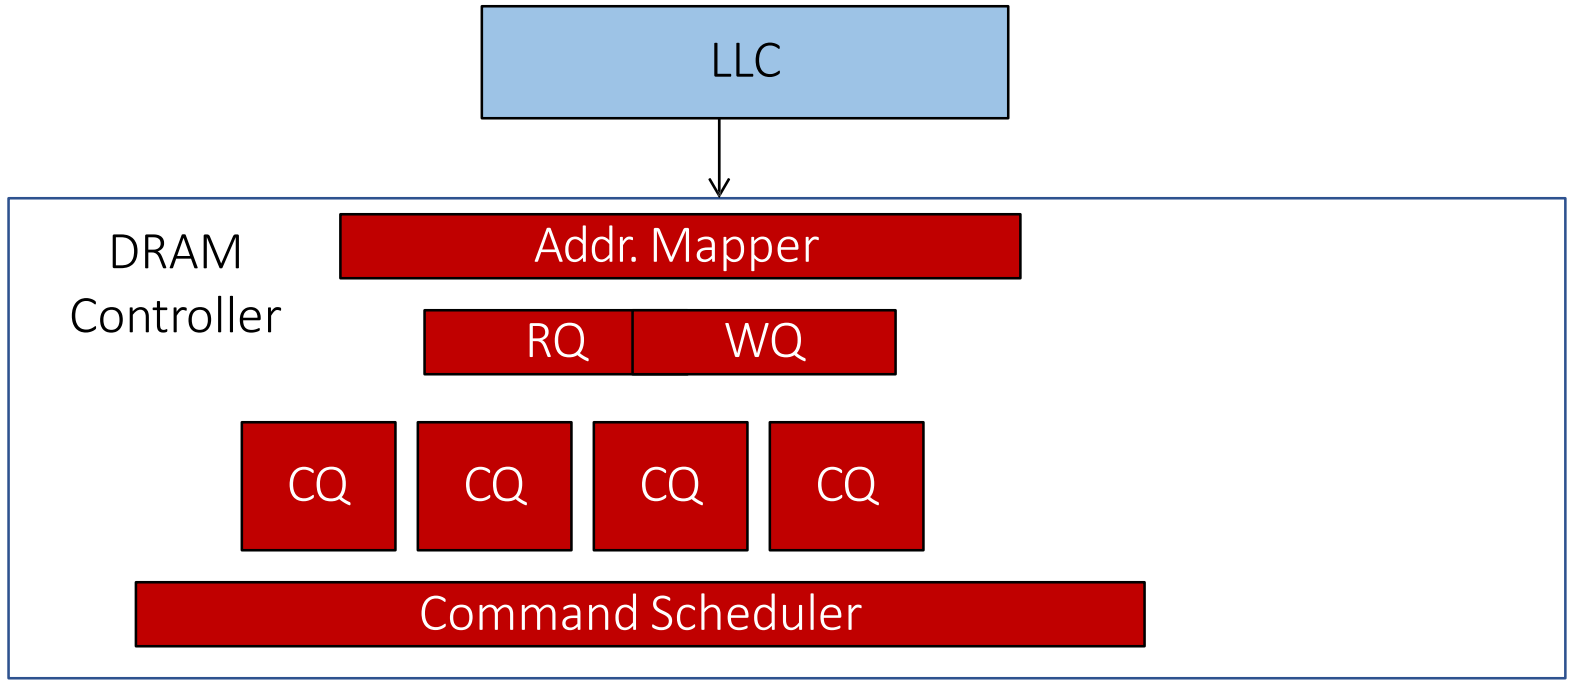

Command and Address bus

Double Data Rate (DDR)
Data bus (burst length= eight beats each of 64-bits)

Frequency
800MHz
1066MHz
1333MHz

Frequency
1600MHz
2133MHz
2666MHz

DRAM Bandwidth

Names	Memory clock	I/O bus clock	Transfer rate	Theoretical bandwidth
DDR-200, PC-1600	100 MHz	100 MHz	200 MT/s	1.6 GB/s
DDR-400, PC-3200	200 MHz	200 MHz	400 MT/s	3.2 GB/s
DDR2-800, PC2-6400	200 MHz	400 MHz	800 MT/s	6.4 GB/s
DDR3-1600, PC3-12800	200 MHz	800 MHz	1600 MT/s	12.8 GB/s
DDR4-2400, PC4-19200	300 MHz	1200 MHz	2400 MT/s	19.2 GB/s
DDR4-3200, PC4-25600	400 MHz	1600 MHz	3200 MT/s	25.6 GB/s
DDR5-4800, PC5-38400	300 MHz	2400 MHz	4800 MT/s	38.4 GB/s
DDR5-6400, PC5-51200	400 MHz	3200 MHz	6400 MT/s	51.2 GB/s

High Bandwidth Memory (HBM)

- Base logic layer (blue), CPU and Memory glued by TSV
- Multiple DRAM chip layers
- Significantly higher bandwidth (200 to 800 GB/s)
- What is the problem then? Capacity ☹
- HBM3: 819.2 GB/sec Capacity: 24GB ☹

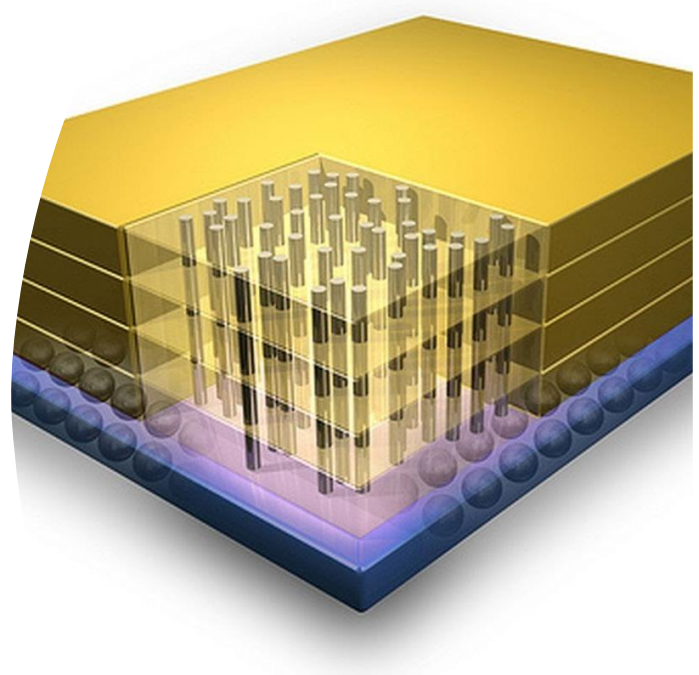

DRAM/DRAM-Controller for Performance

Heavily mined research area: DRAM scheduling, DRAM address mappers, fairness, QOS, prefetch awareness, ……

The Performance Perspective

- All of Google's Data Center Workloads (ISCA 2015):

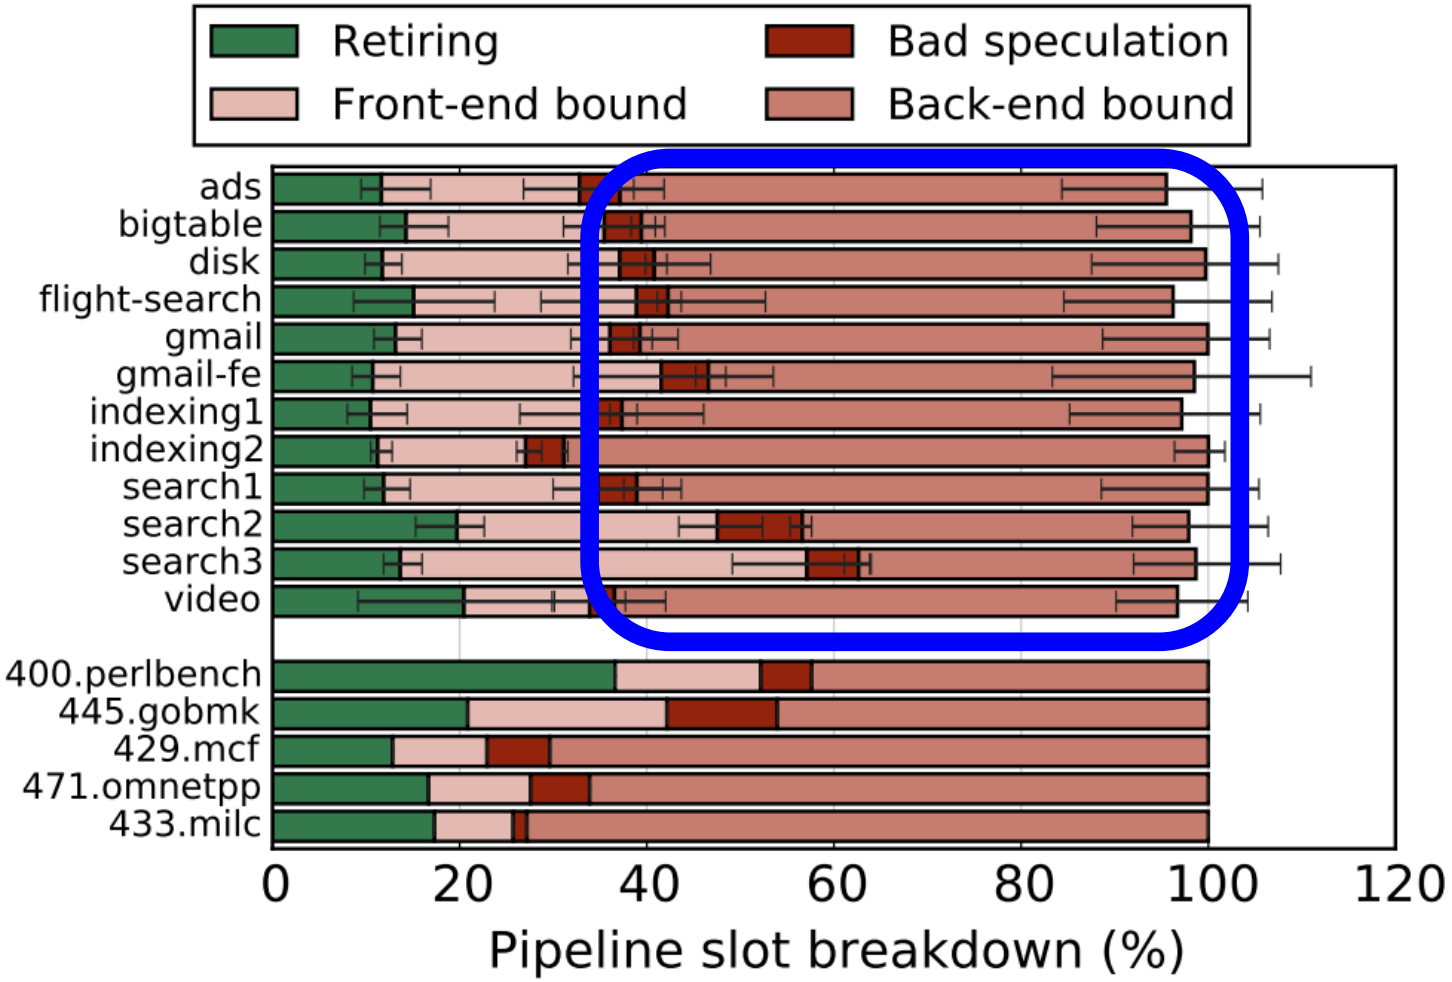

Data Movement vs. Computation Energy

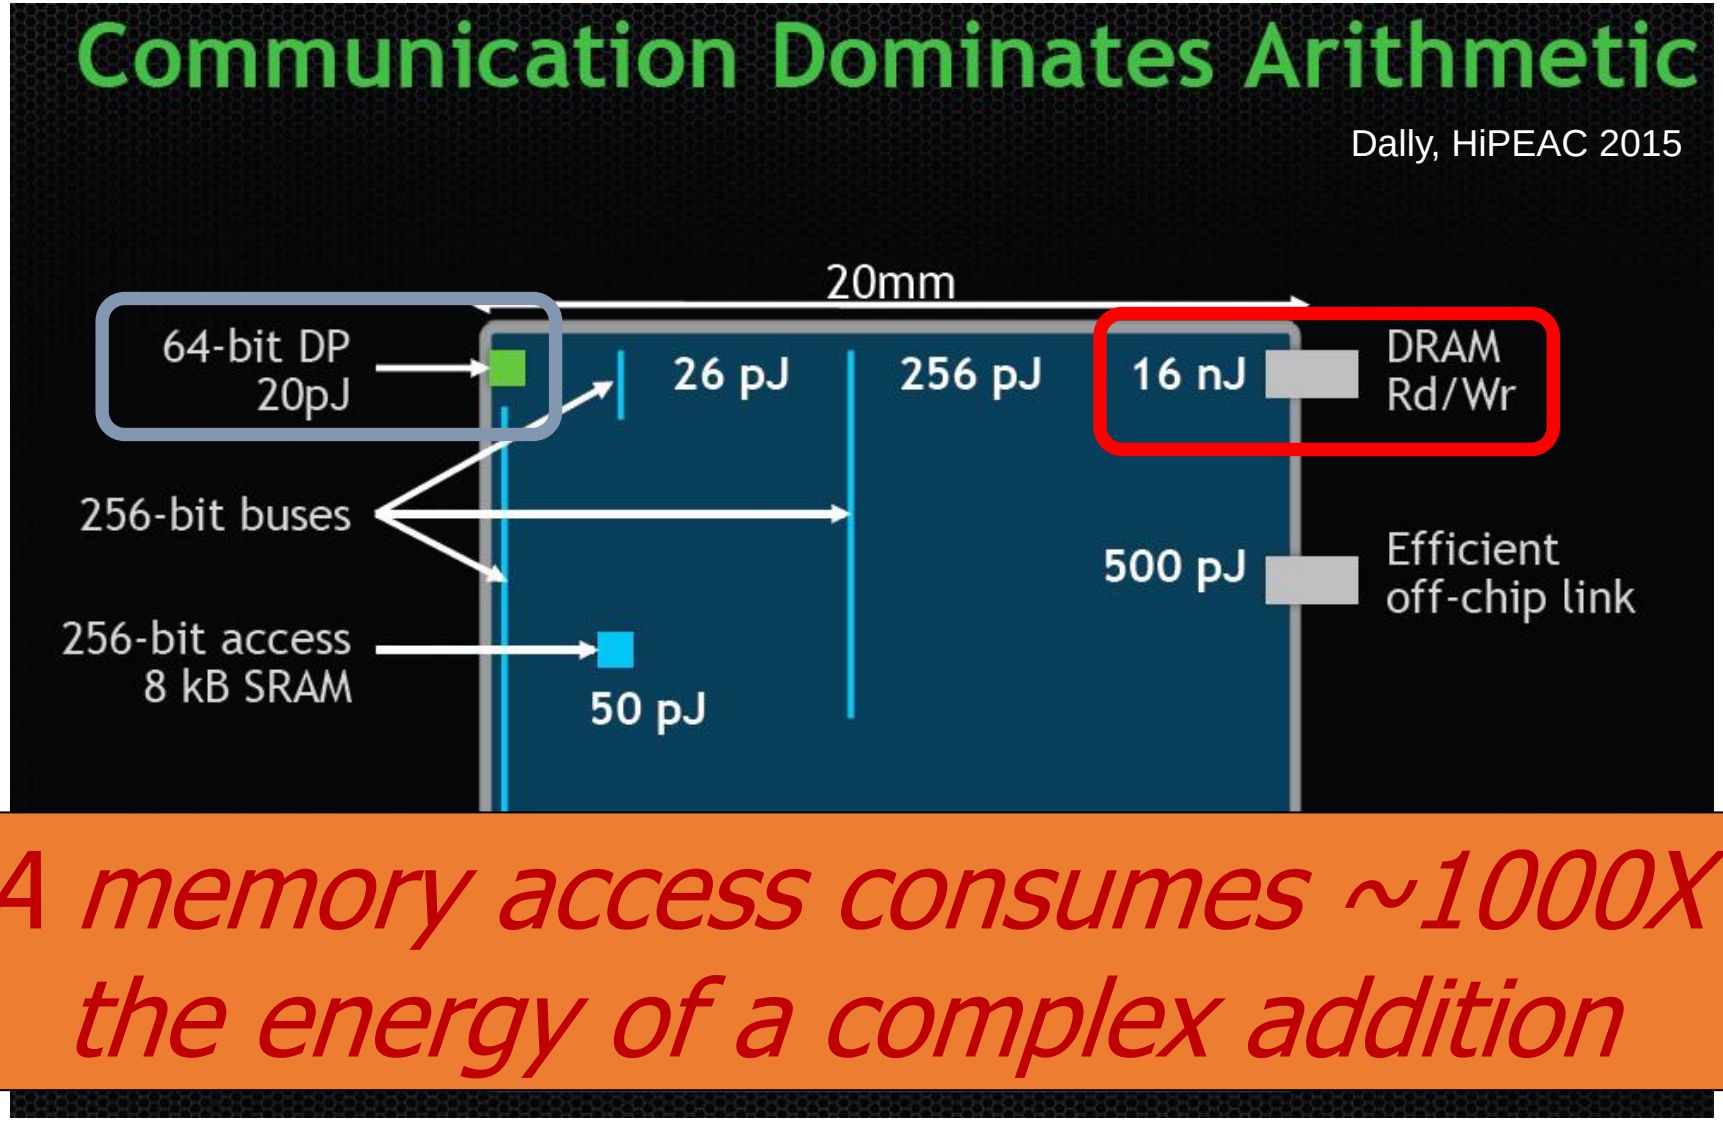

A memory access consumes ~1000X the energy of a complex addition

Data Movement vs. Computation Energy

- **Data movement** is a major system energy bottleneck
  - Comprises 41% of mobile system energy during web browsing [2]
  - Costs ~115 times as much energy as an ADD operation [1, 2]

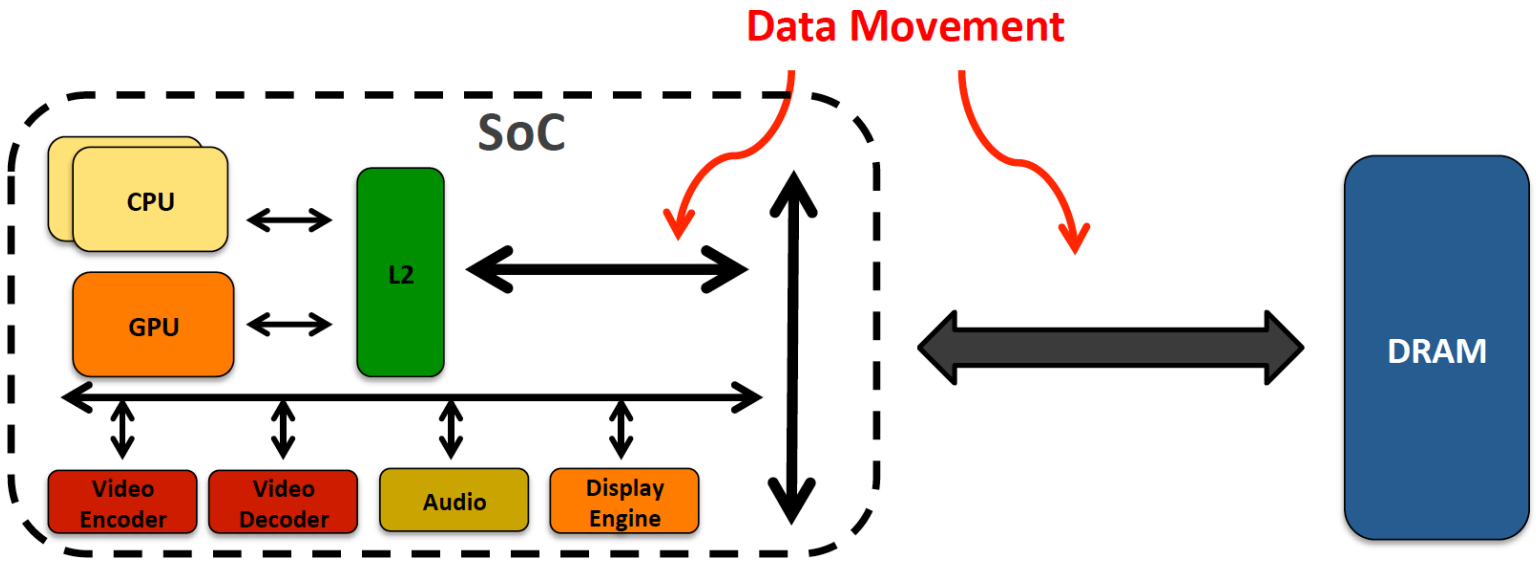

[1]: Reducing data Movement Energy via Online Data Clustering and Encoding (MICRO'16)

[2]: Quantifying the energy cost of data movement for emerging smart phone workloads on mobile platforms (IISWC'14)

Processing Near Memory

Move compute near memory

What compute to move? (simple functions)

Mapping computation (heterogenous computing) ?

Virtual memory ?

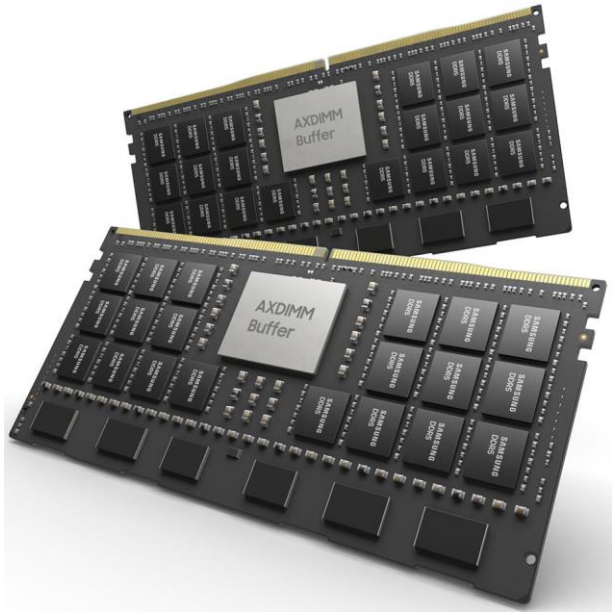

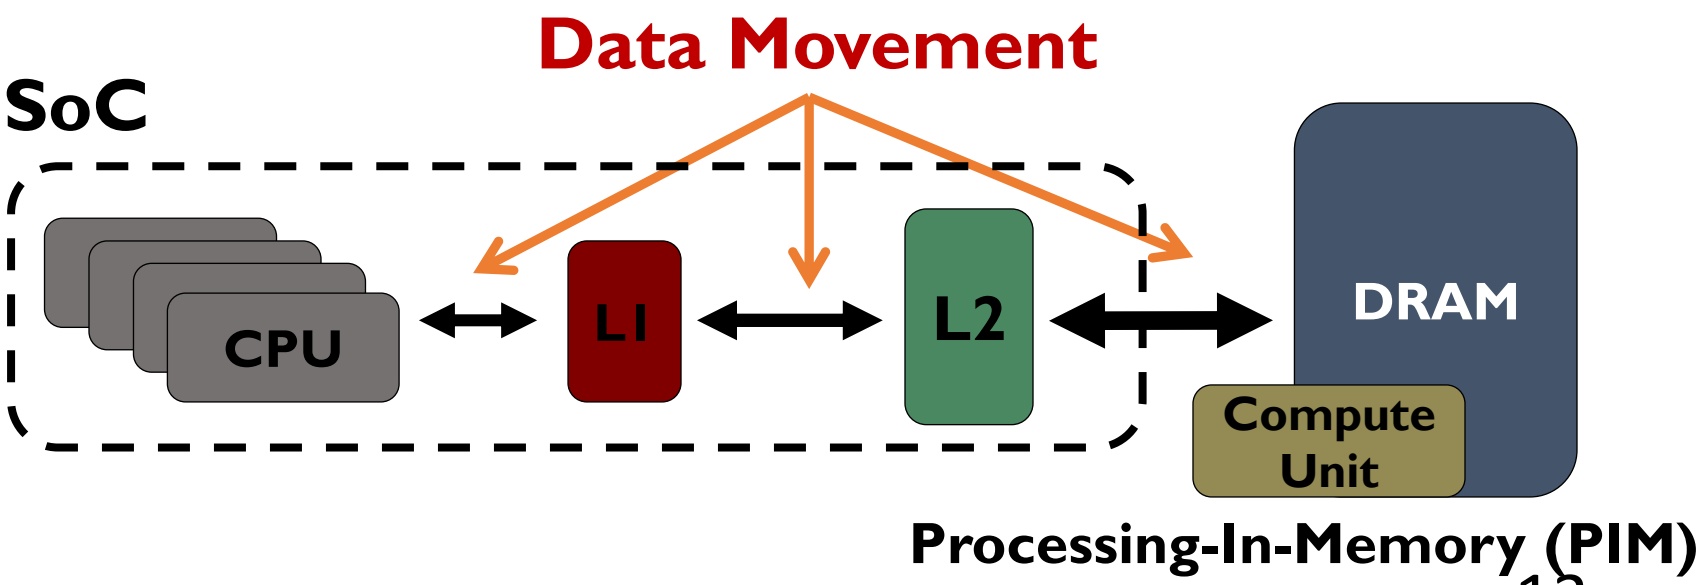

Non-volatile Persistent Memory

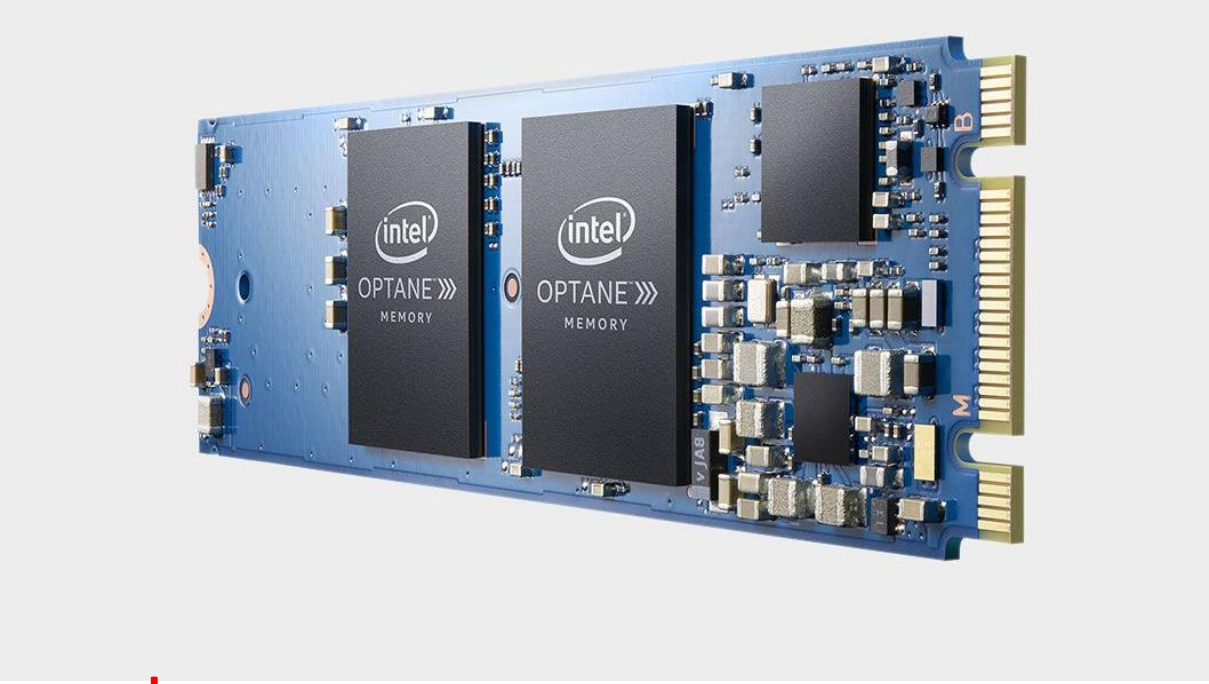

Two modes

Application mode: Non-volatile, app/OS can decide what/where..

Memory mode: Volatile, DRAM acts like a cache ☺ More capacity

Non-volatile LLC too ☺ Write-latency is higher than volatile LLC ☹

PAUSE

From Performance to Security

DRAM Row Shared ☹

Sandy Bridge /w 1 DIMM

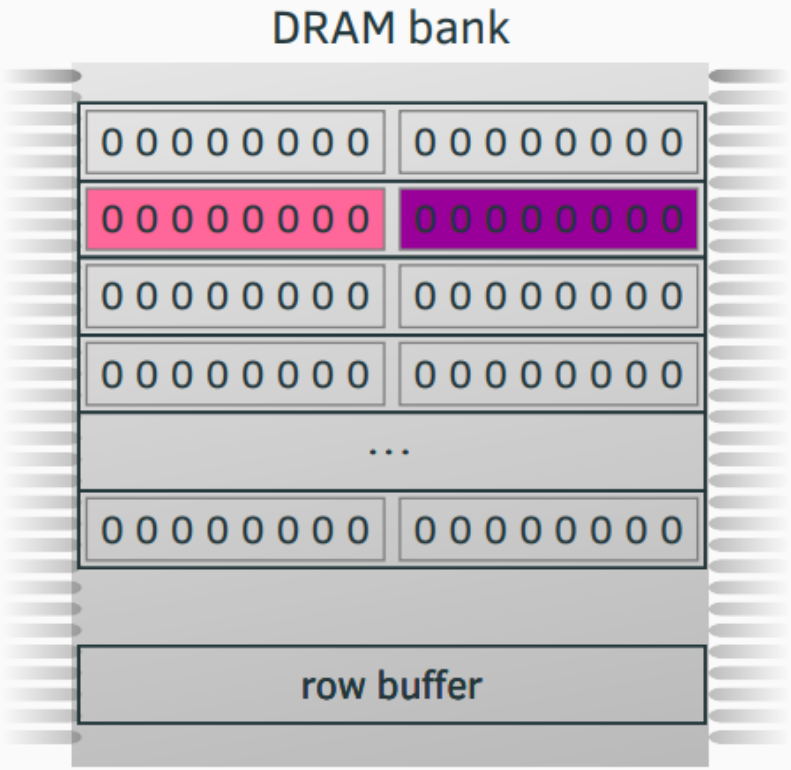

2 pages per row

Side Channel

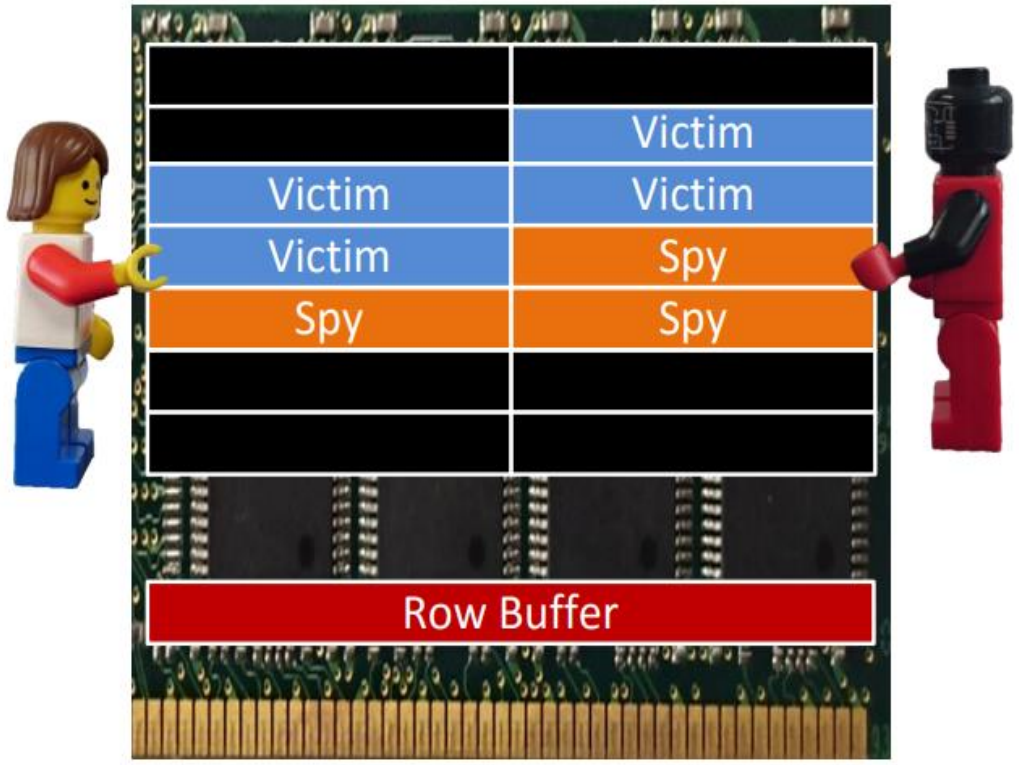

Side-channel attacks

Row-hits: Fast access
Conflicts: Slow access

Covert Channel

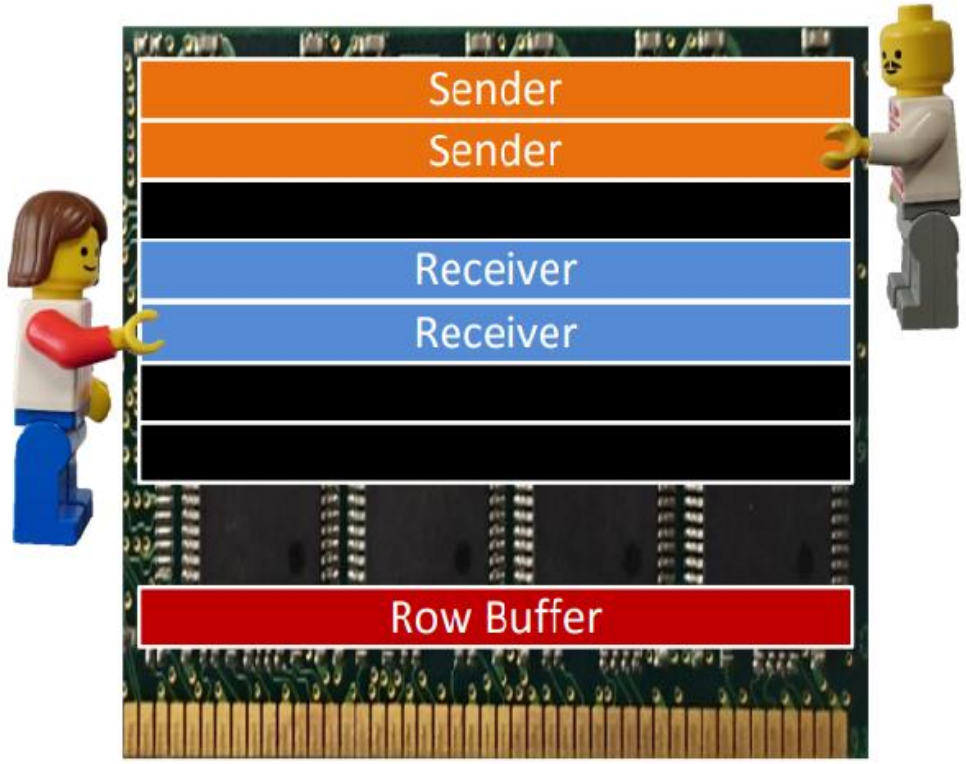

Covert-channel attacks

Something more: Integrity attacks

Rowhammer

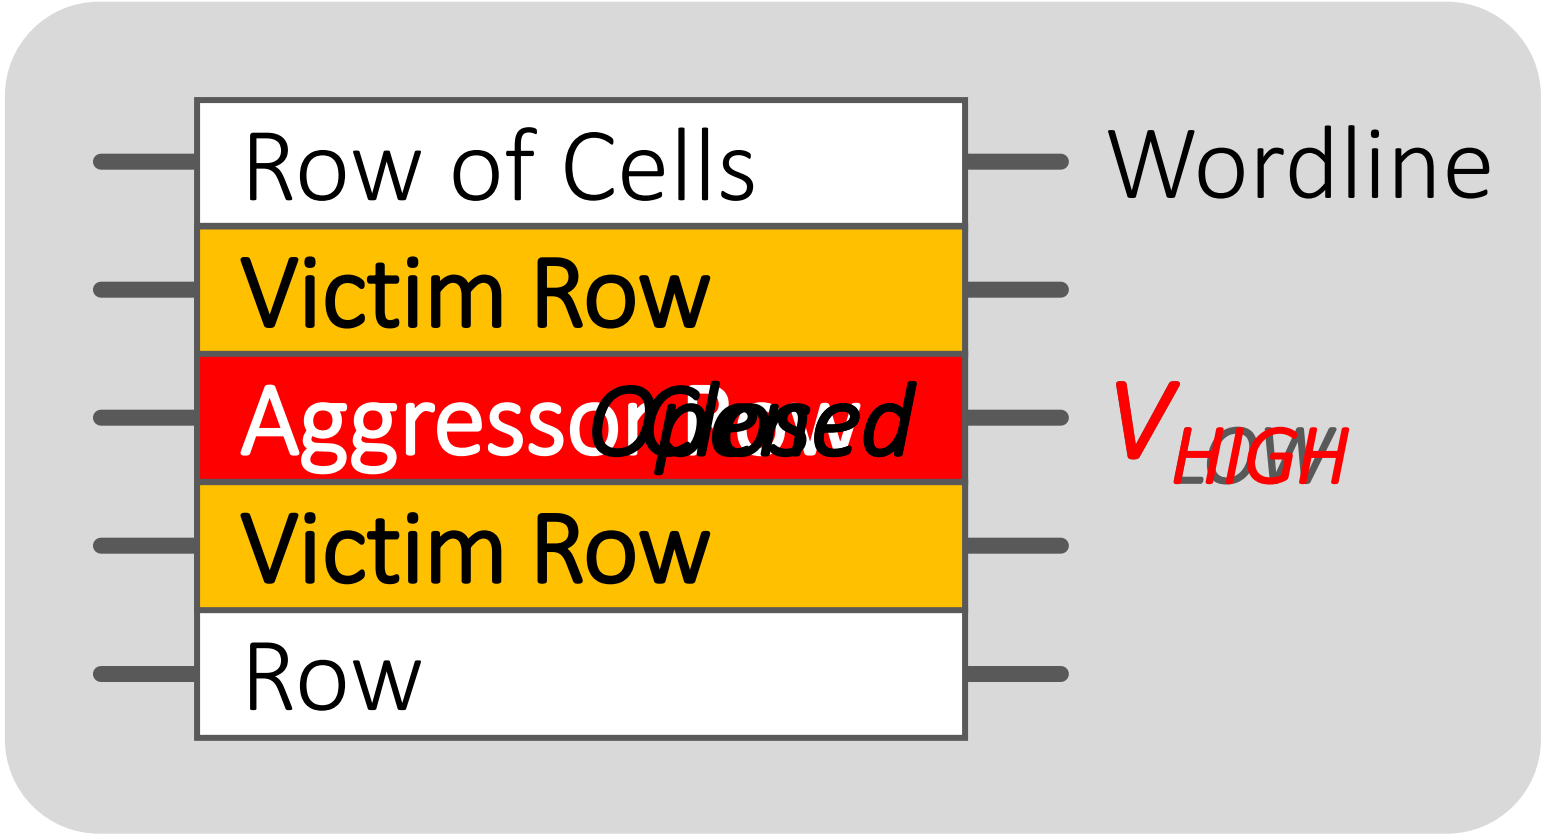

*Repeatedly opening and closing a row induces **disturbance errors** in adjacent rows*

Rowhammer

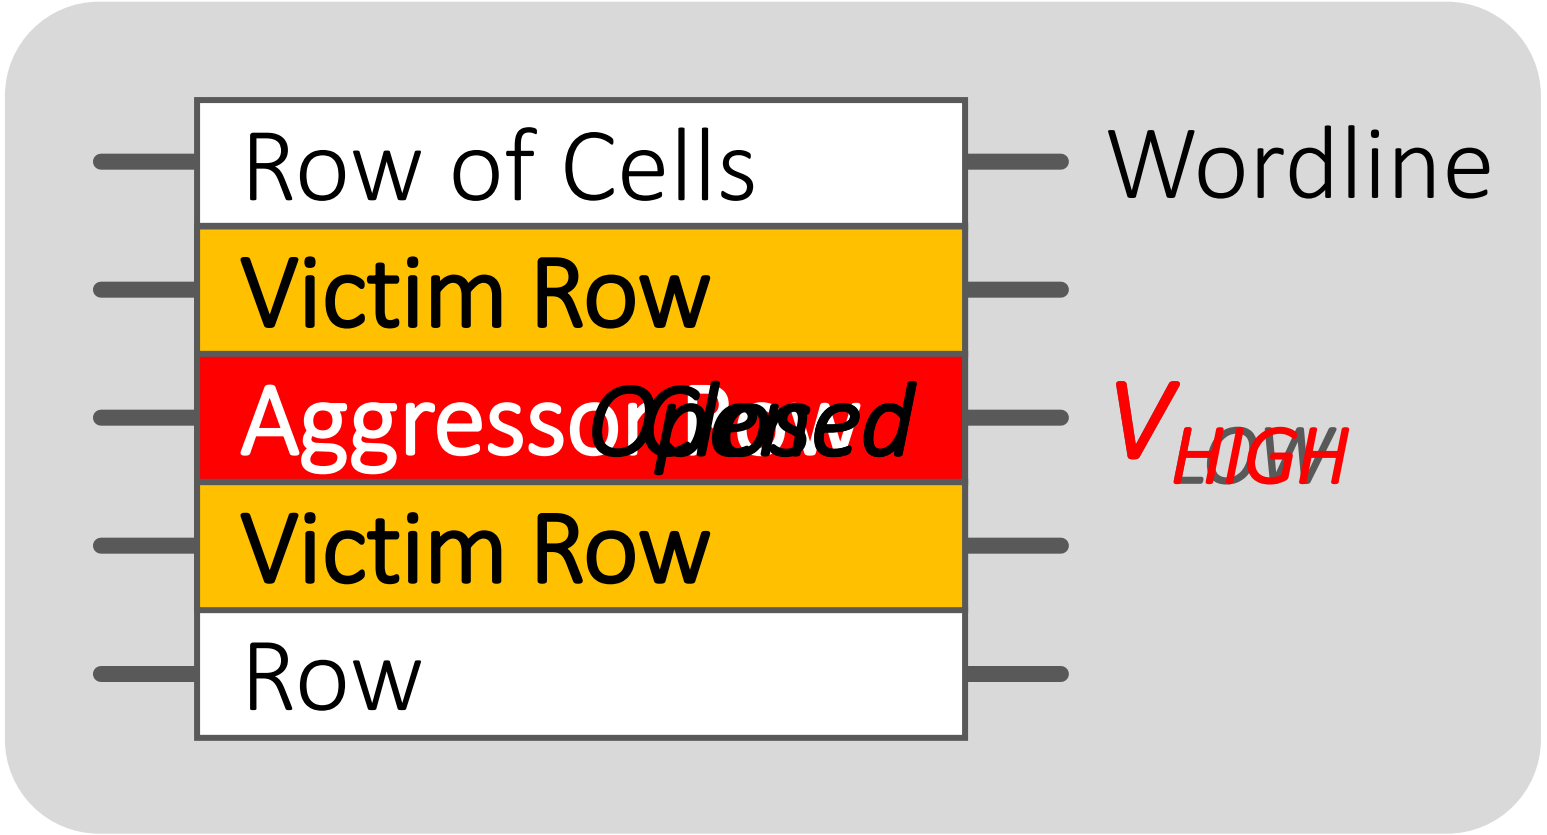

"It's like breaking into an apartment by repeatedly slamming a neighbor's door until the vibrations open the door you were after" – Motherboard Vice

The Code Please

1. Avoid ***cache hits***
   - Flush **X** from cache

2. Avoid ***row hits*** to **X**
   - Read **Y** in another row

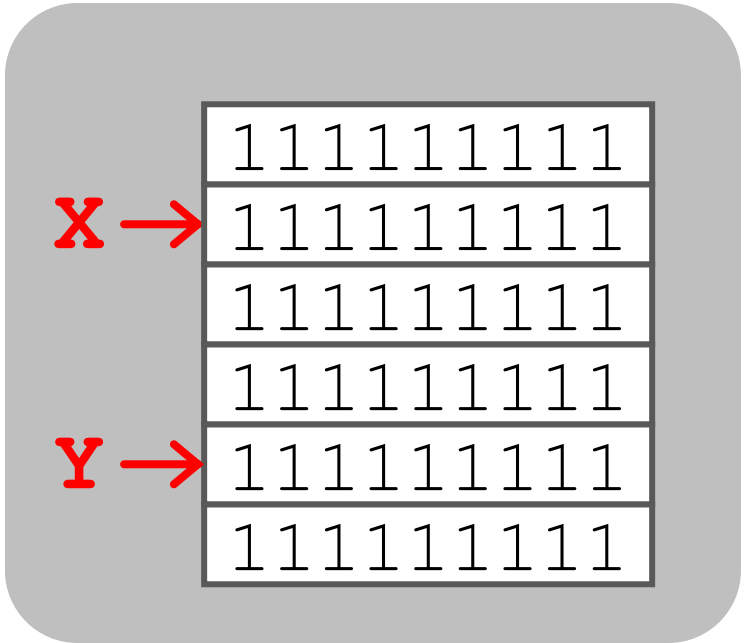

The code please

```
loop:
   mov (X), %eax
   mov (Y), %ebx
   clflush (X)
   clflush (Y)
   mfence
   jmp loop
```

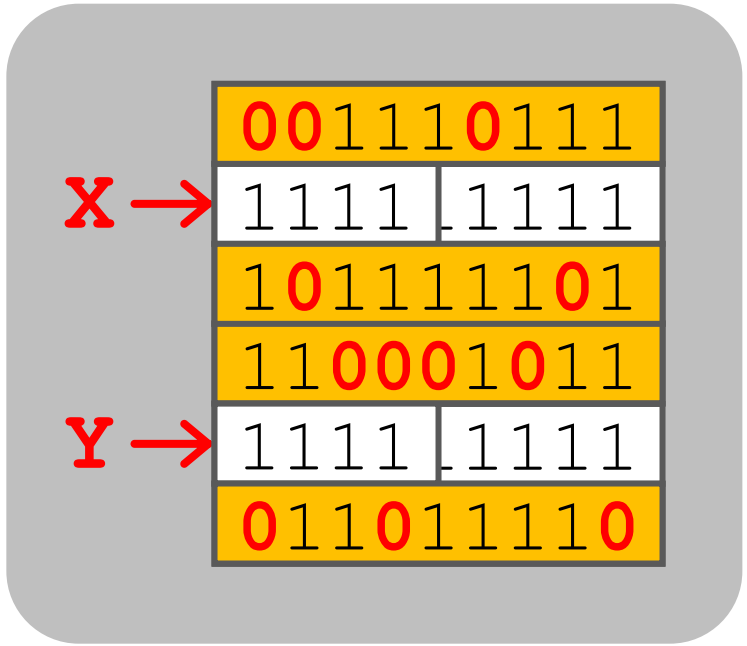

Why? Electromagnetic Coupling

Toggling the wordline voltage briefly increases the voltage of adjacent wordlines

Slightly opens adjacent rows → Charge leakage

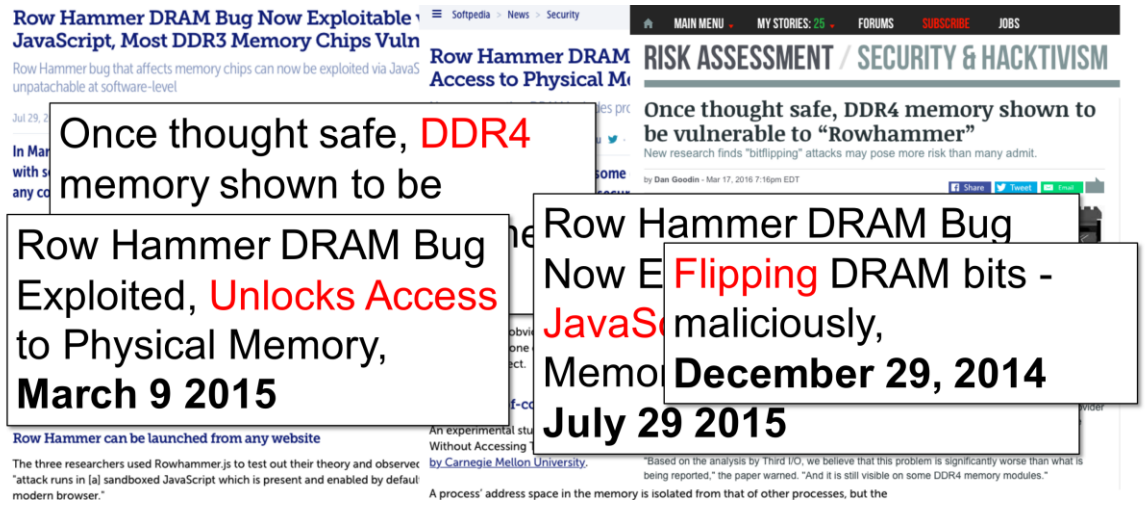

Two Days Before

Spechammer ☹

Cold boot attacks

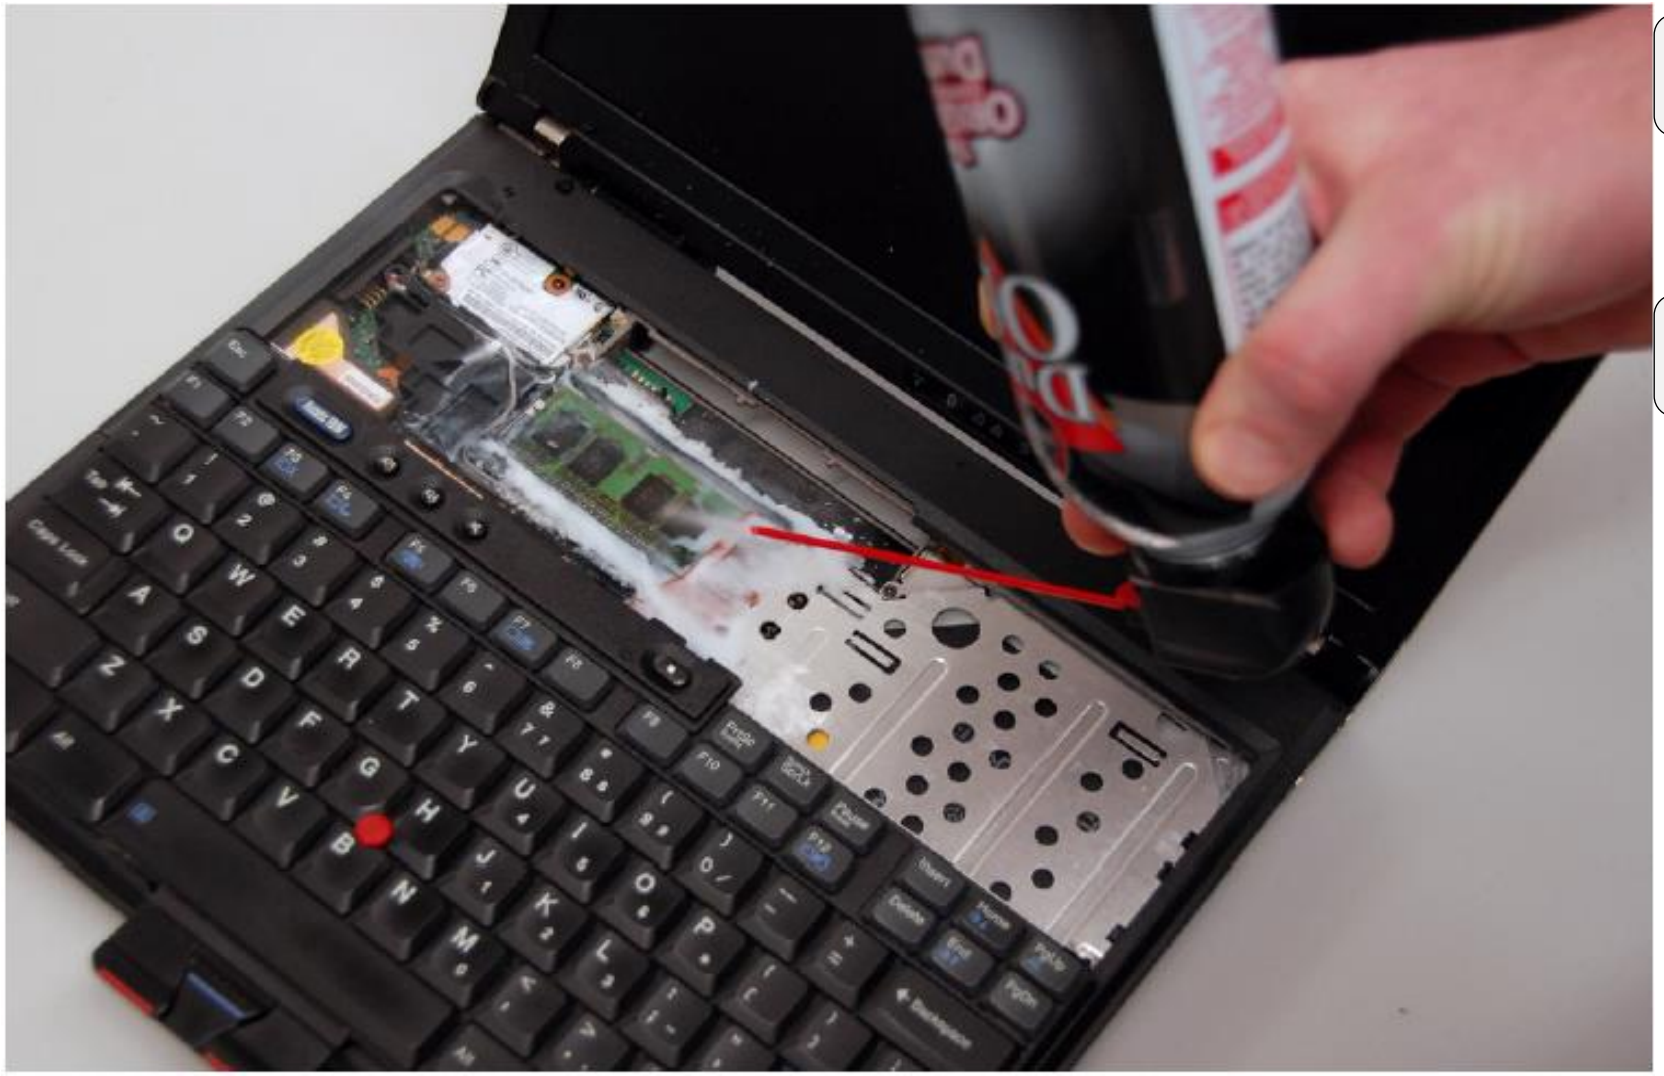

Before powering off

Freeze it to -50 ° C

Cool it

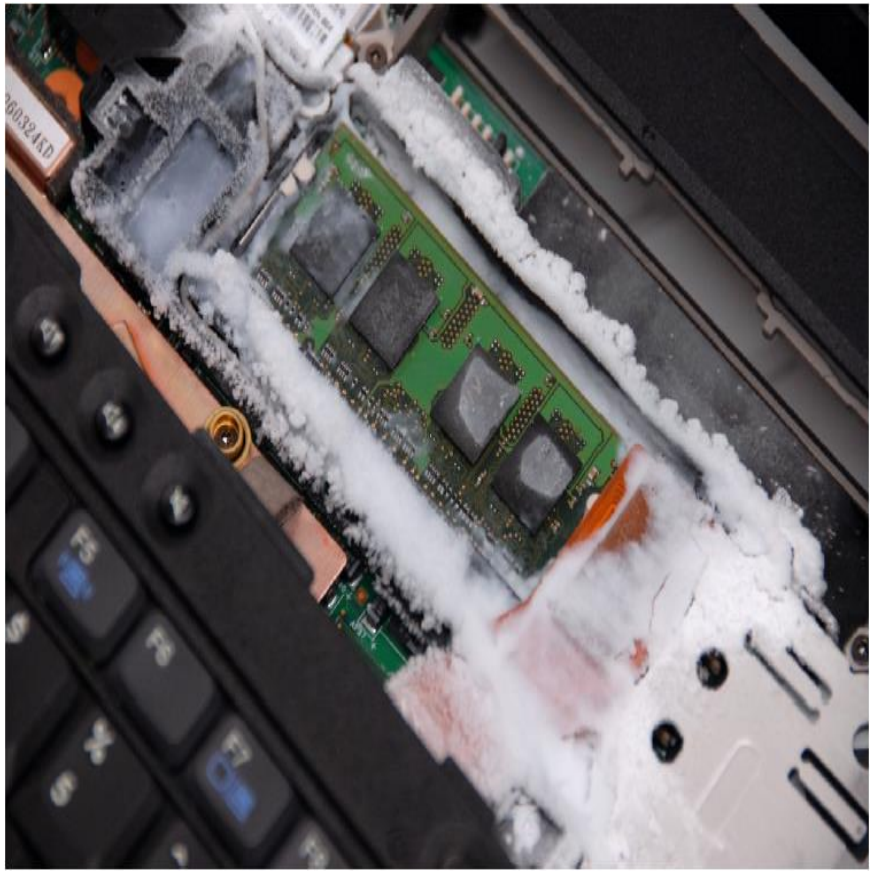

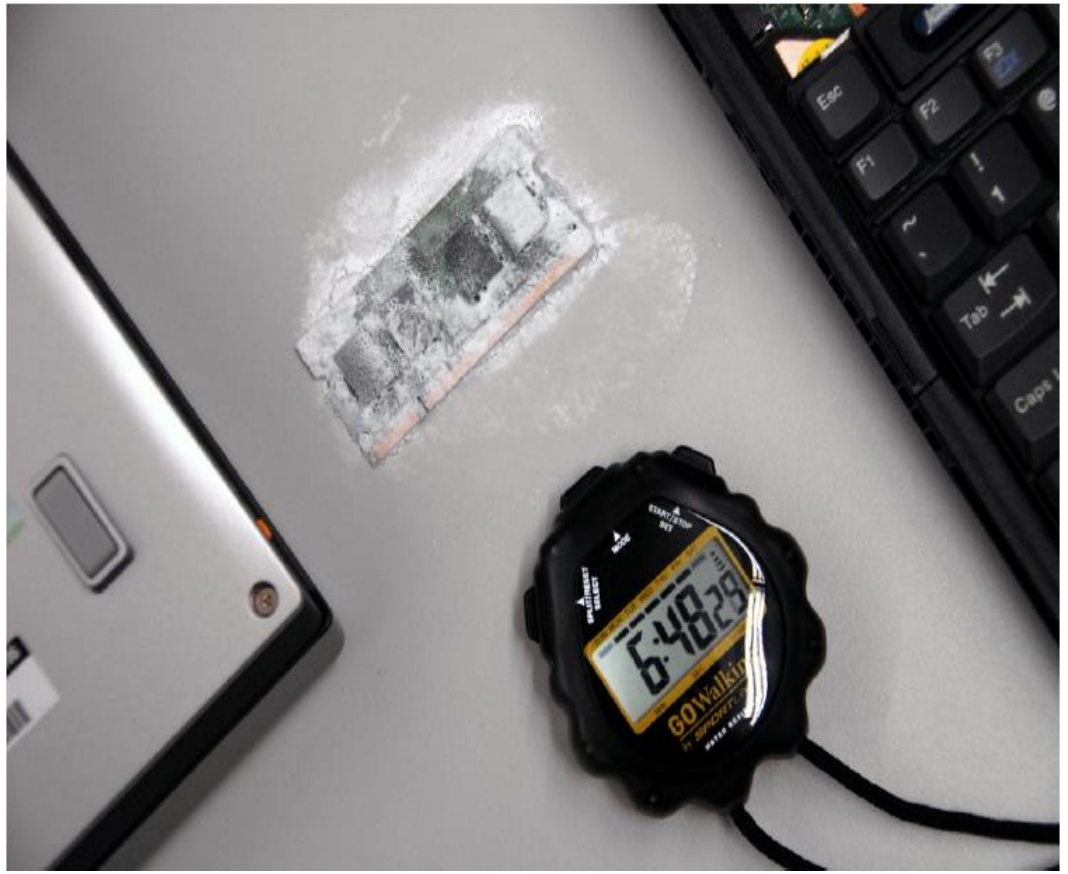

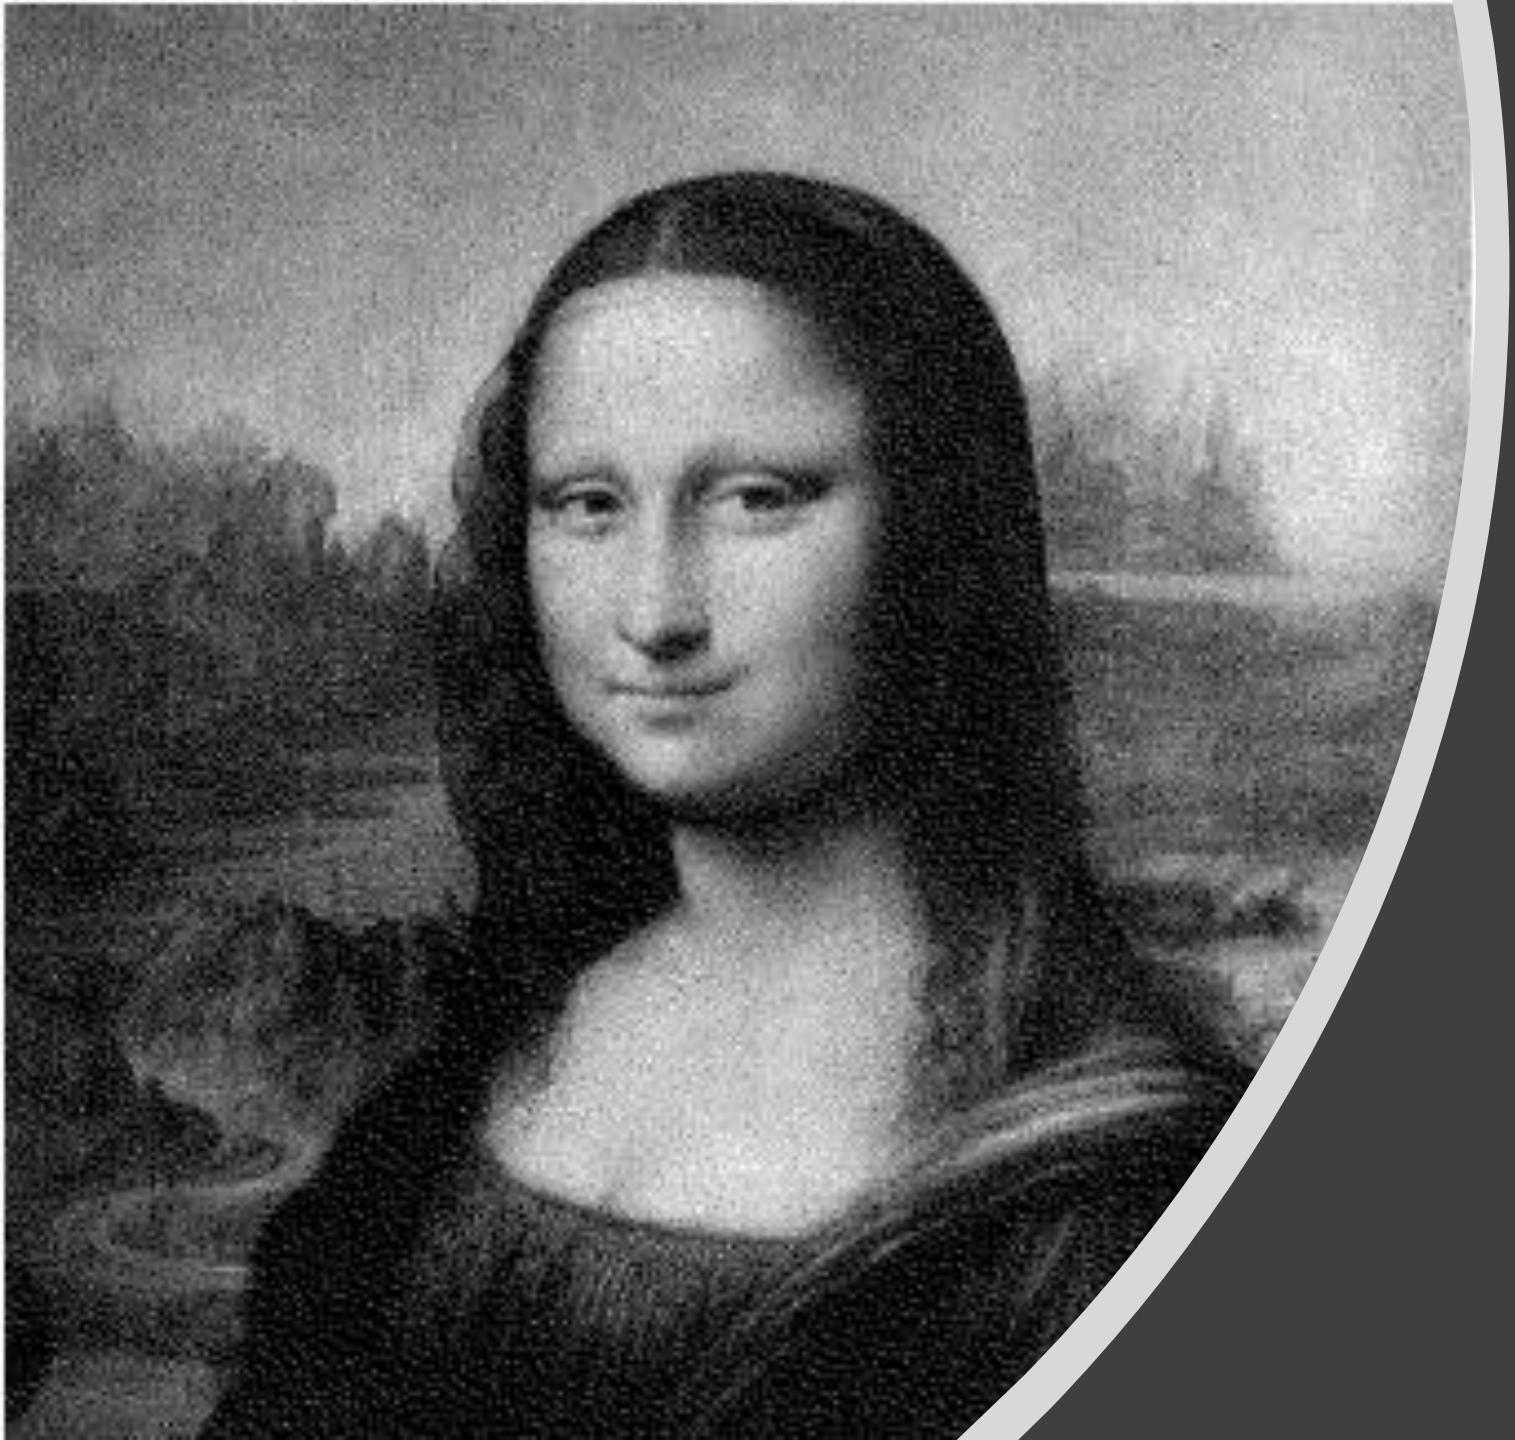

After 5 Seconds

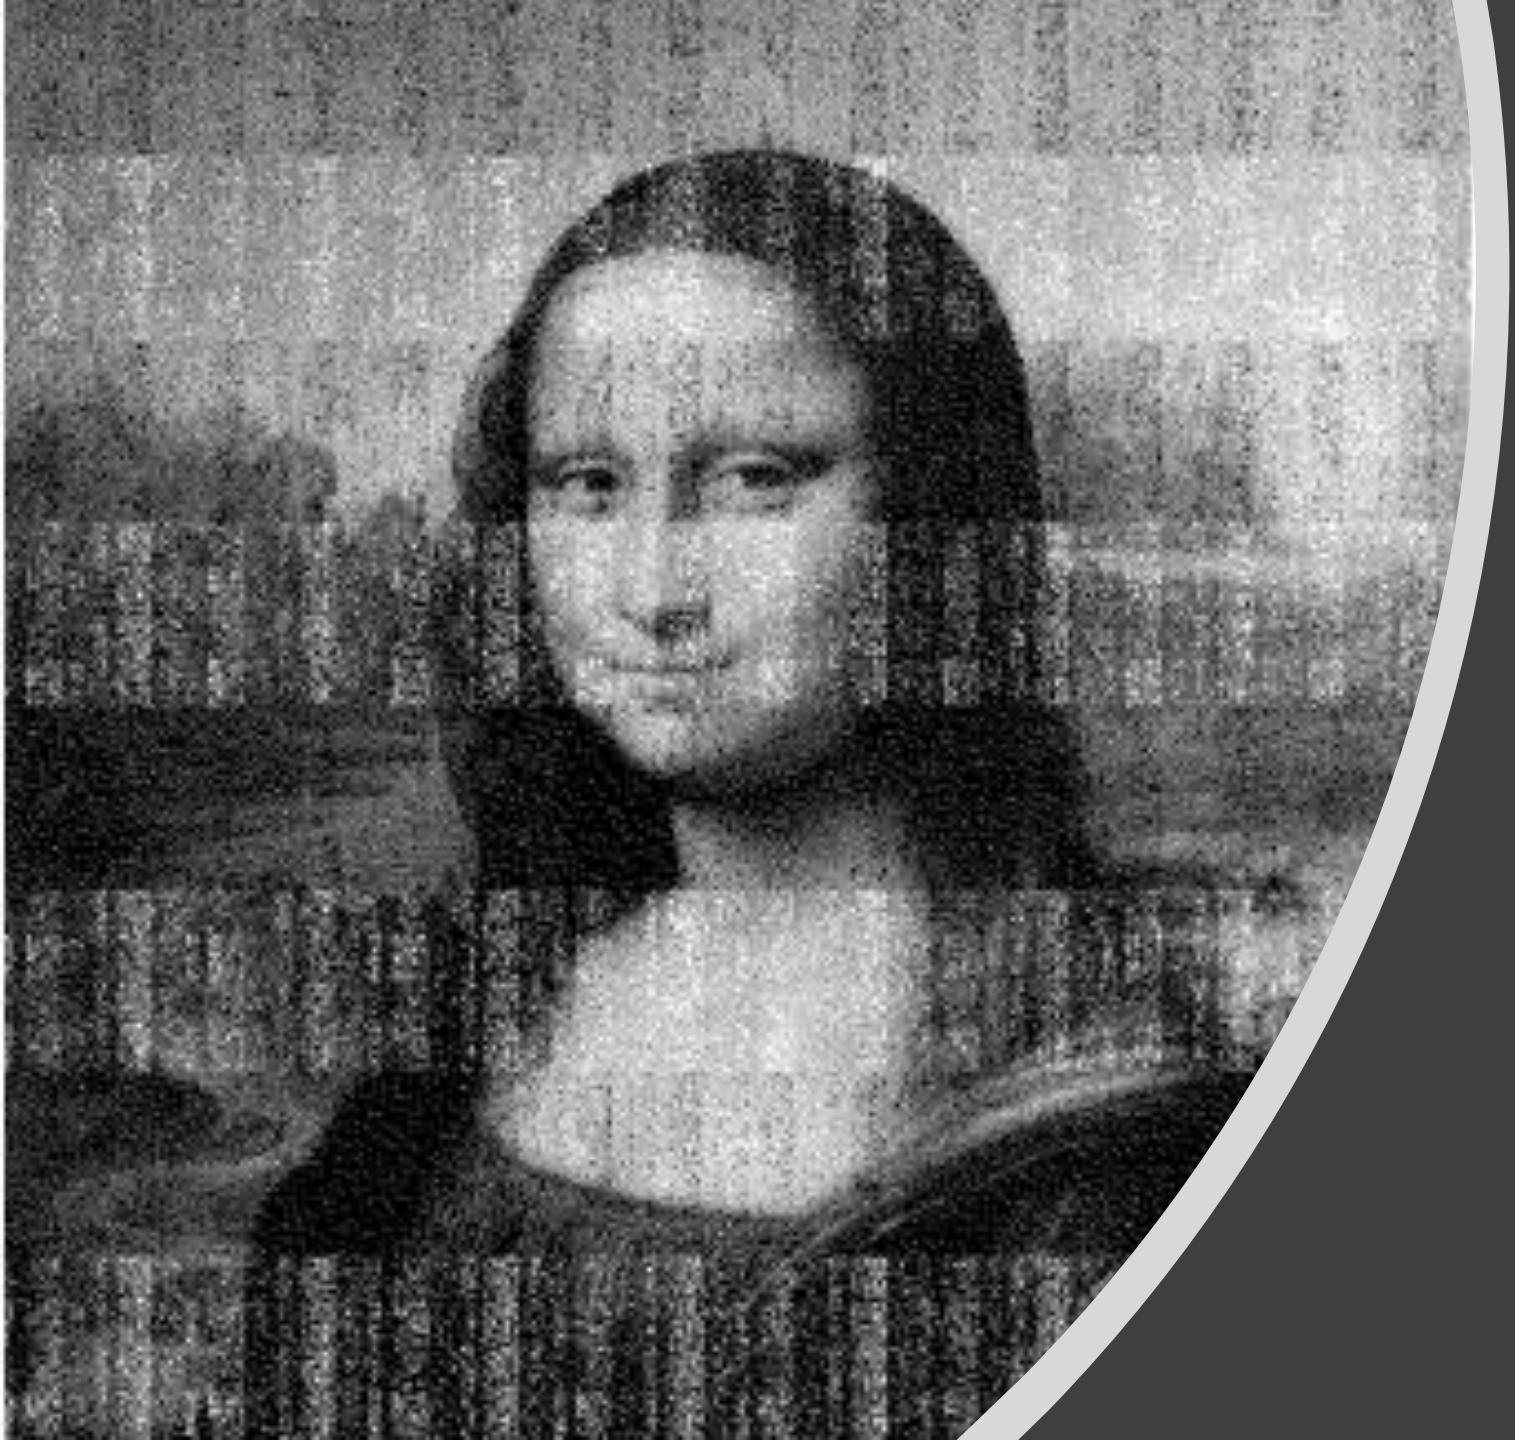

30 Seconds

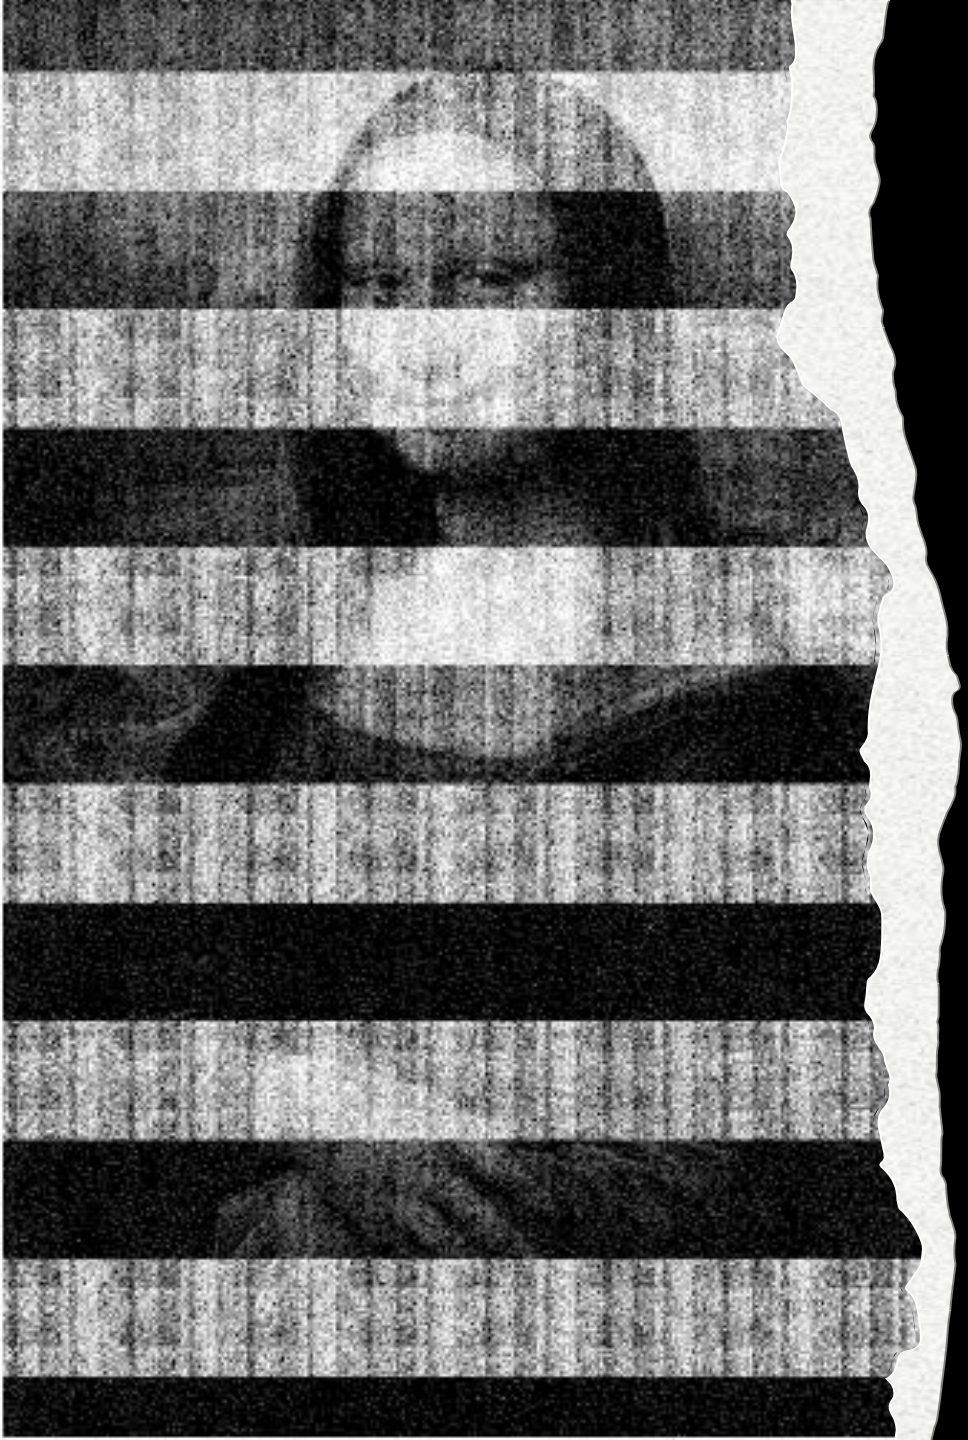

60 Seconds

300 Seconds Think about non-volatile memory

Thanks