Lecture-17
(DRAM Timing attacks & Row Hammer)
CS665-Fall 2018
Secure Memory Systems

Biswa@CSE-IITK

Now at the DRAM

Control flow/pattern of operations in RSA depends upon the secret key.

These operations cause memory accesses

Memory intensity of victim can leak number of ‘1’s in the key.

By measuring it’s own memory access latency, attacker can infer memory intensity of the victim.

Now at the DRAM: An Example

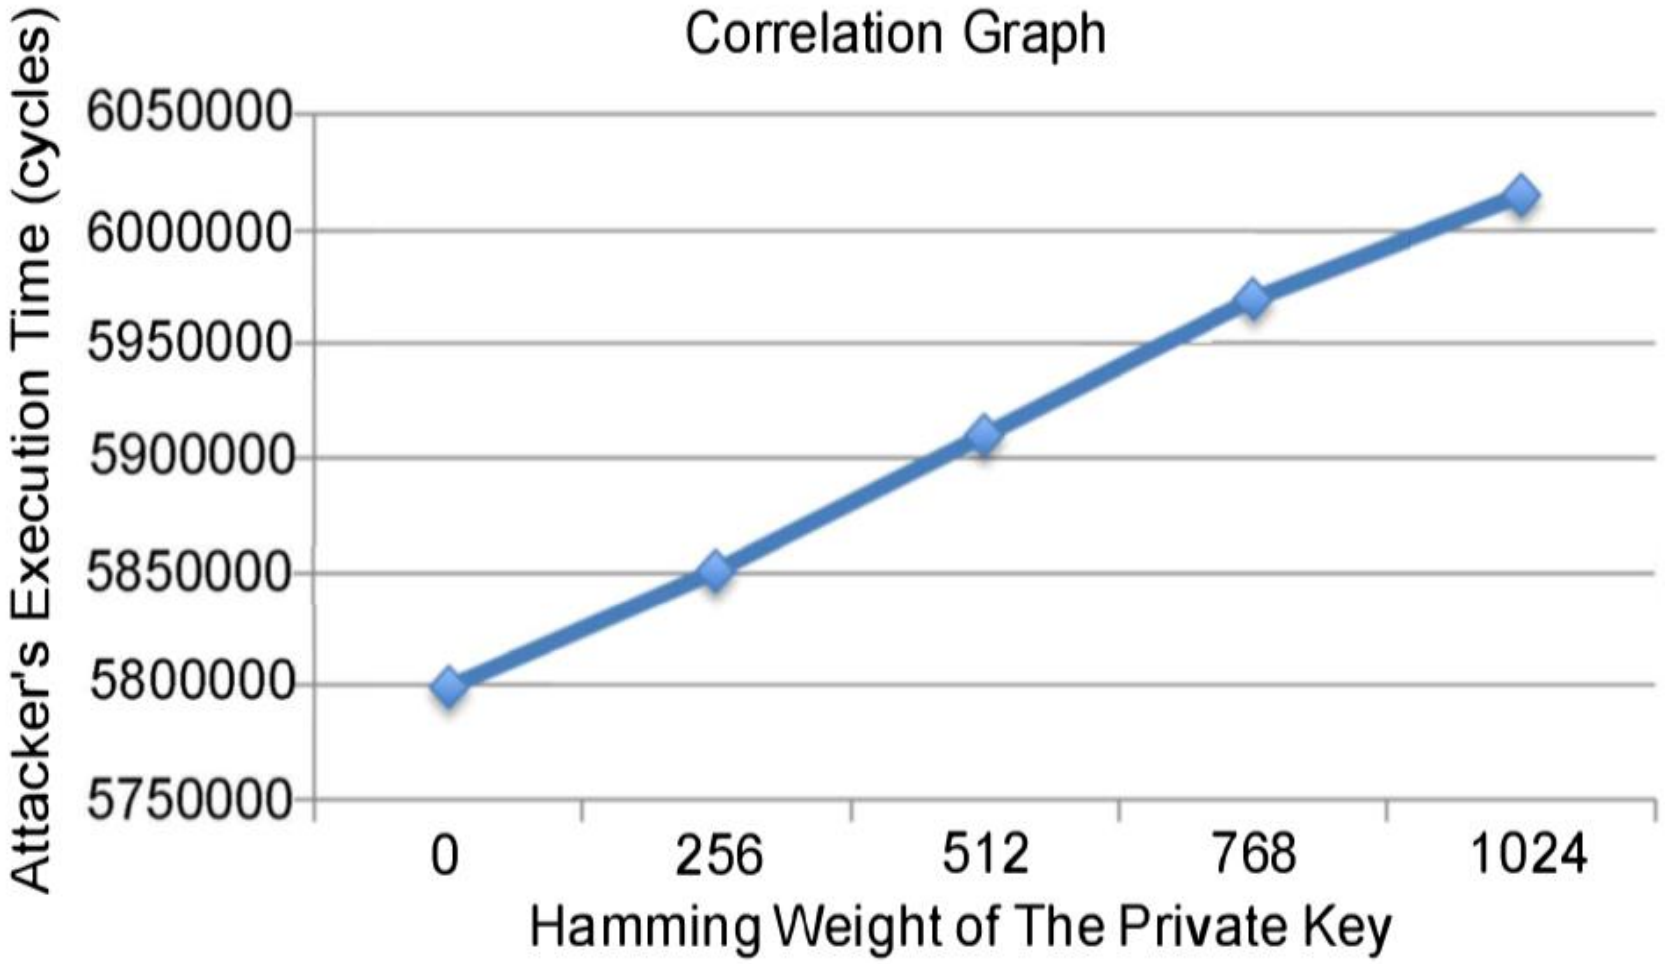

Bank Triple Alteration [MICRO '15]

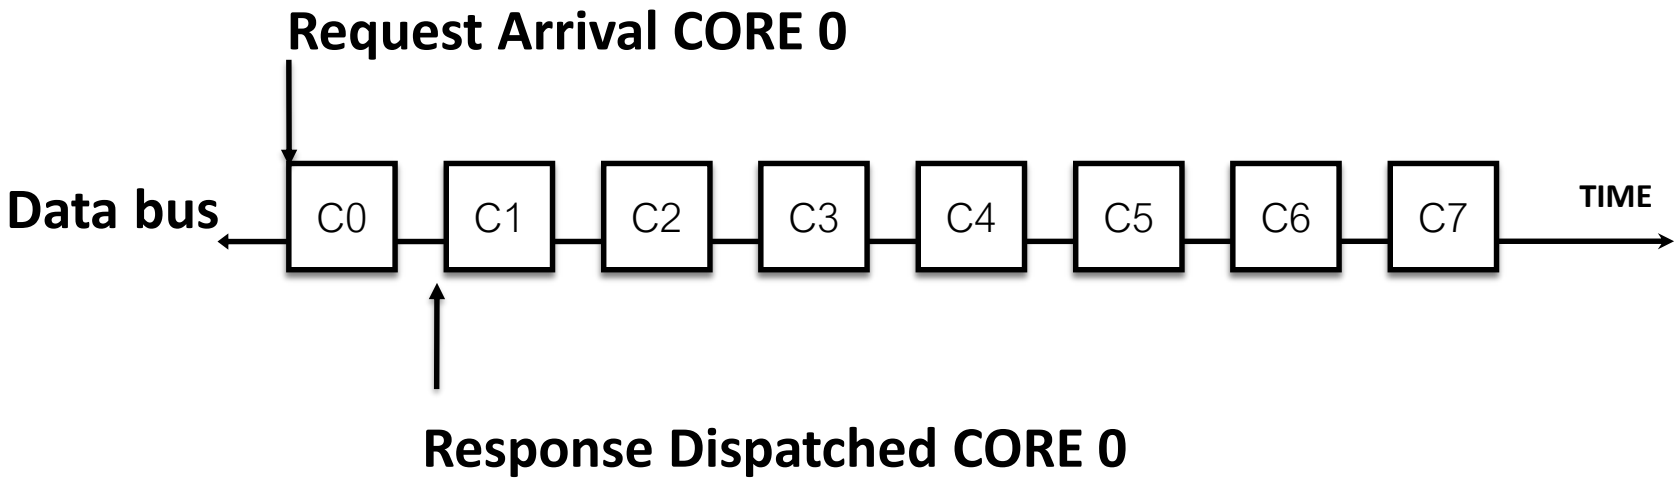

STEP 1. Schedule cores alternatively (Round Robin fashion)

Bank Triple Alteration [MICRO '15]

Number of Banks = 8

STEP 2. Divide banks into 3 sets

Bank Triple Alteration [MICRO '15]

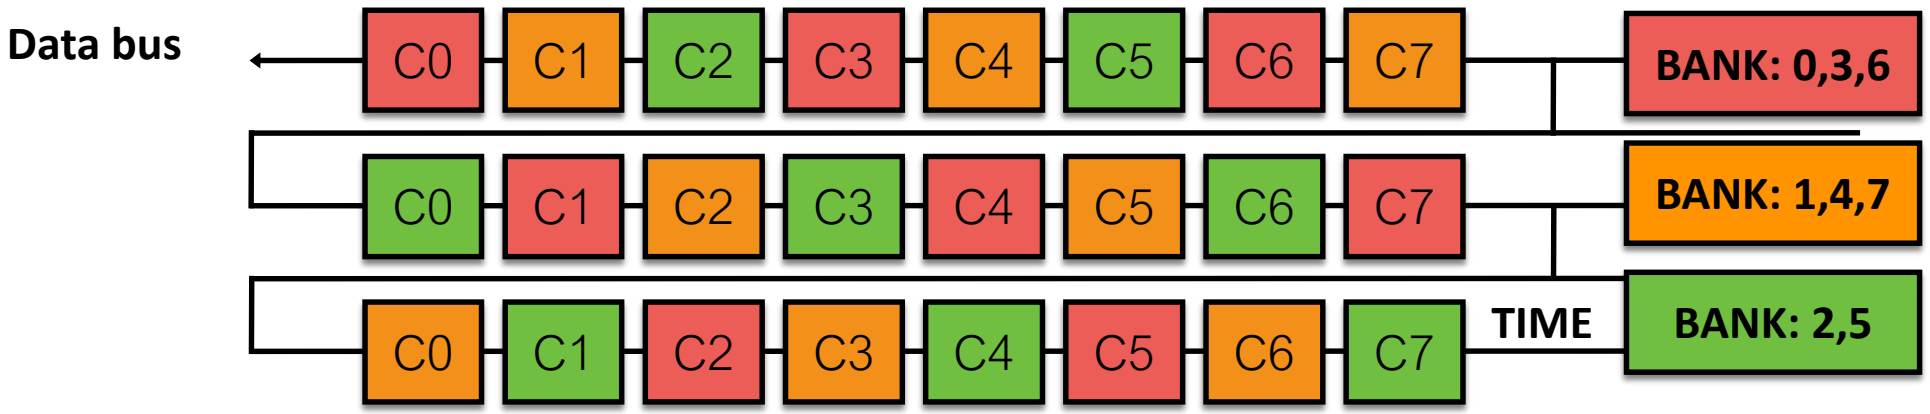

STEP 3. Alternatively schedule request from each bank set.

Bank Triple Alteration [MICRO '15]

Strict scheduling constraints —> Lots of empty slots

Request arriving at a "bad time" needs to wait for the whole interval.

Bank interleaving 18 cycles. Performance loss ?

Now think ☺

Row-hammer [ISCA '14]

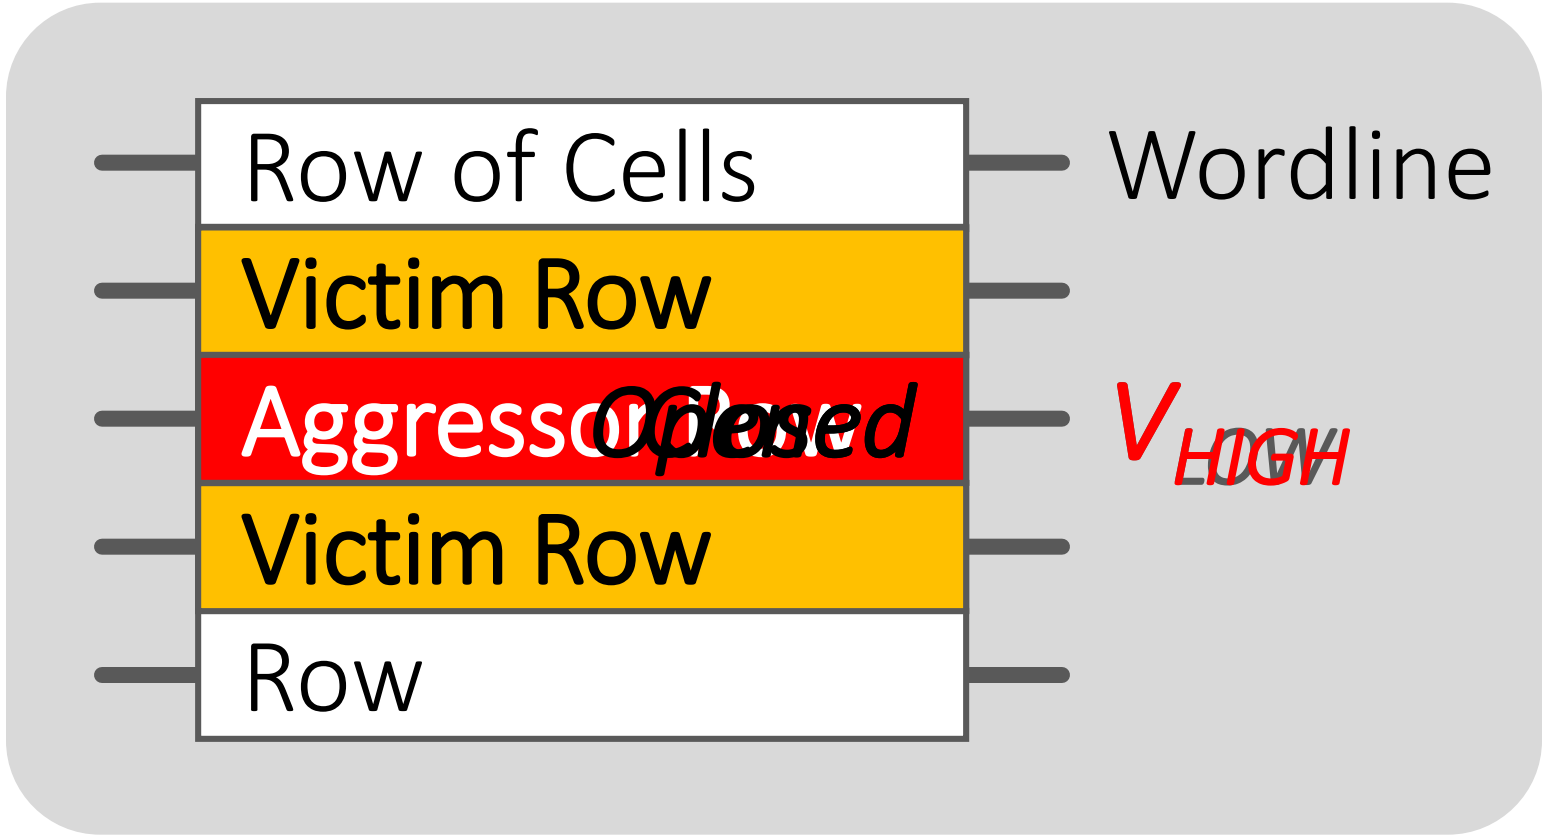

*Repeatedly opening and closing a row induces **disturbance errors** in adjacent rows*

Row-hammer [ISCA '14]

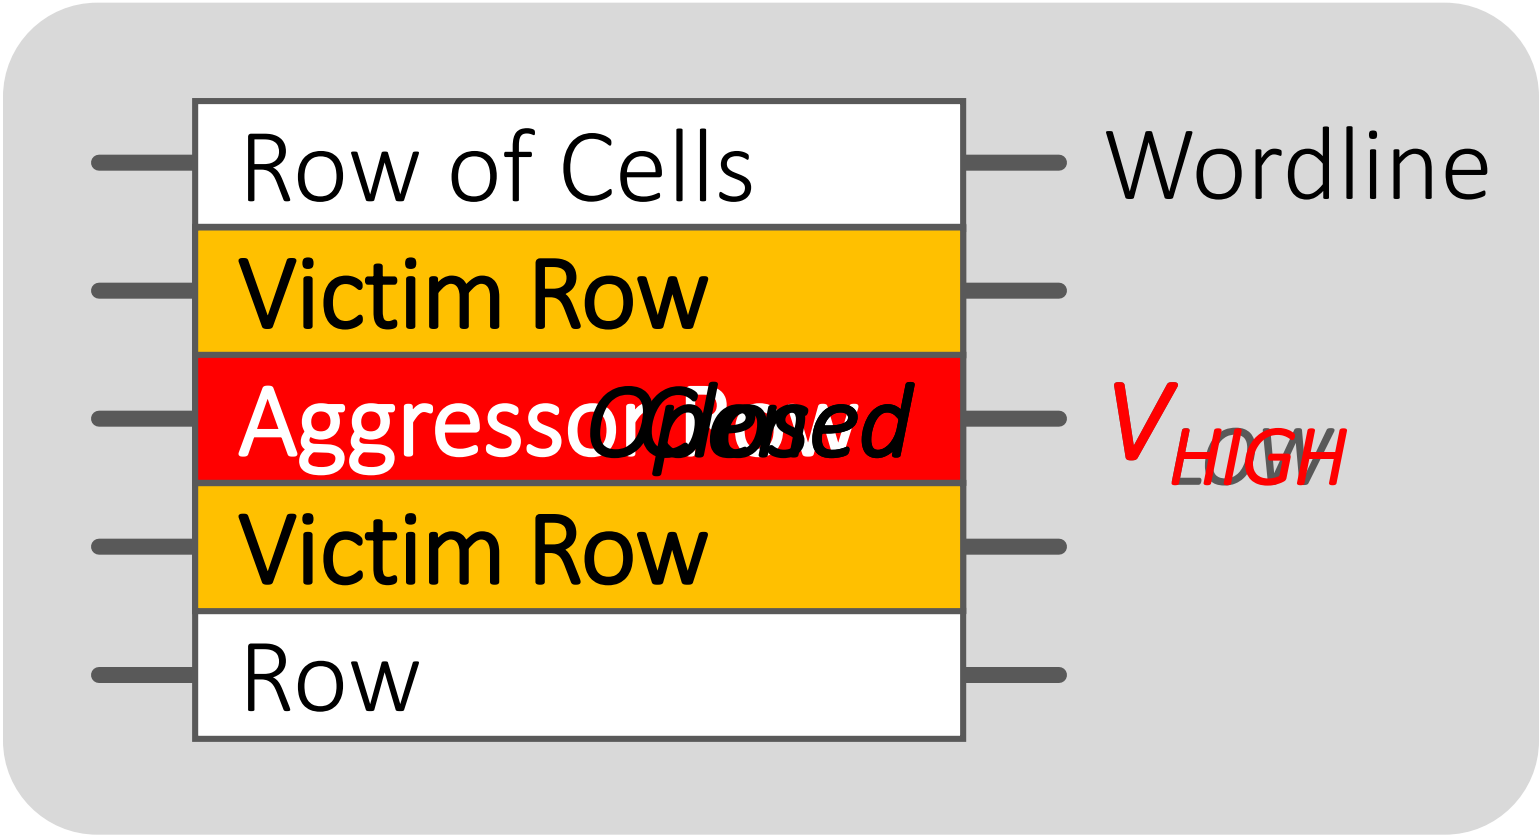

"It's like breaking into an apartment by repeatedly slamming a neighbor's door until the vibrations open the door you were after" – Motherboard Vice

Row-hammer [ISCA '14] ()

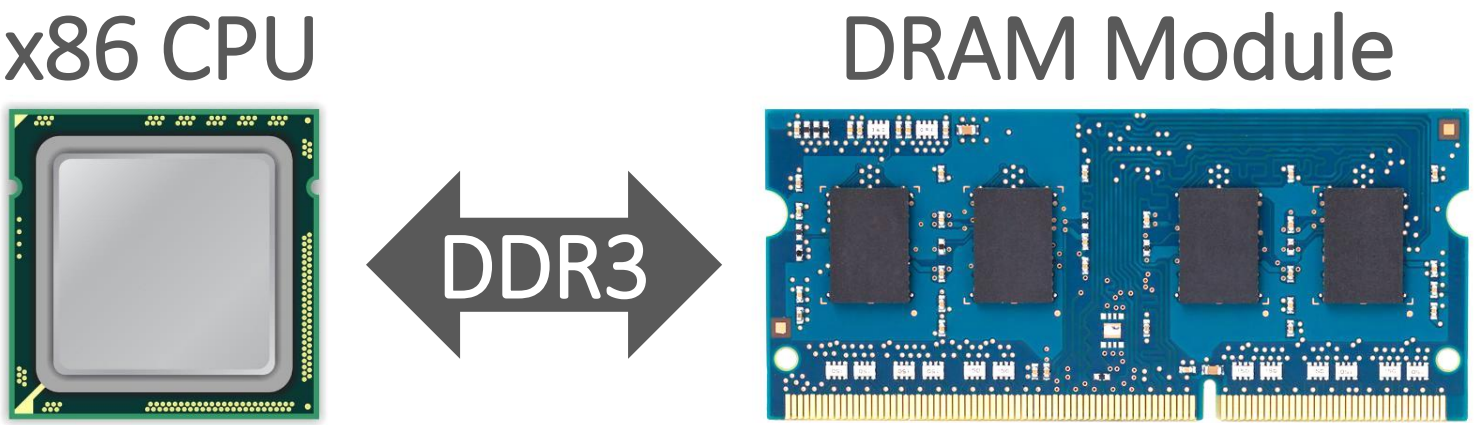

1. Avoid ***cache hits***
   - Flush **X** from cache

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

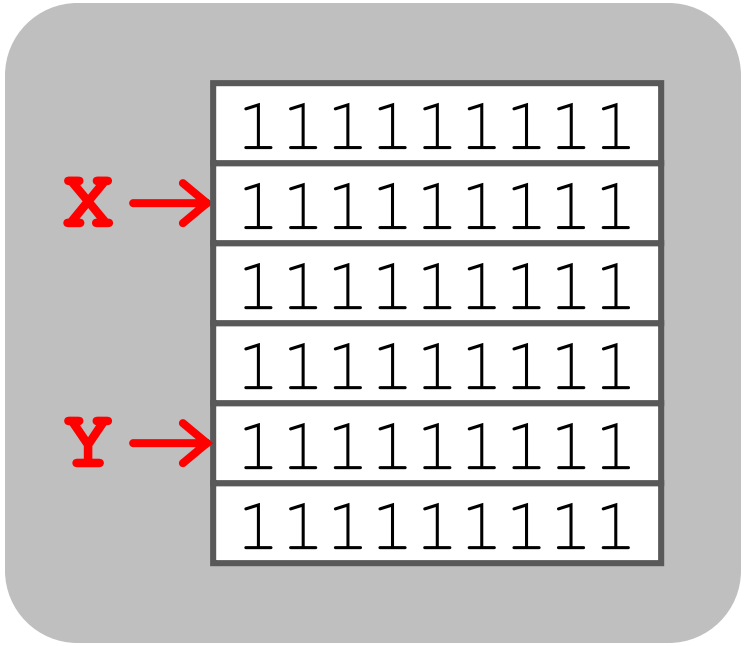

Row-hammer [ISCA '14] (Can you do without Clflush?)

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

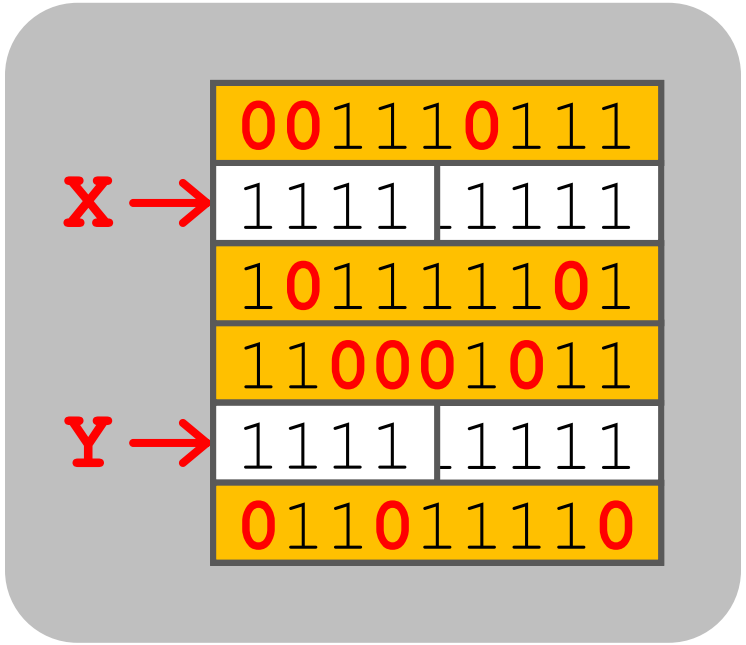

CPU Architecture	Errors	Access-Rate
Intel Haswell (2013)	22.9K	12.3M/sec
Intel Ivy Bridge (2012)	20.7K	11.7M/sec
Intel Sandy Bridge (2011)	16.1K	11.6M/sec
AMD Piledriver (2012)	59	6.1M/sec

But Why?

- *Cause 1: Electromagnetic coupling*
  - Toggling the wordline voltage briefly increases the voltage of adjacent wordlines
  - Slightly opens adjacent rows → Charge leakage
- *Cause 2: Conductive bridges*
- *Cause 3: Hot-carrier injection*

Confirmed by at least one manufacturer

Oh!!

A company	B company	C company
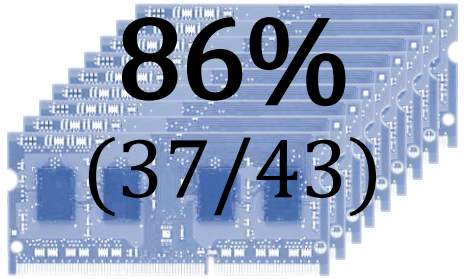 86% (37/43)	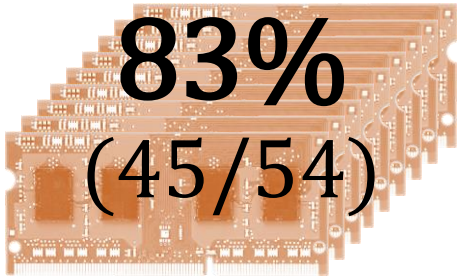 83% (45/54)	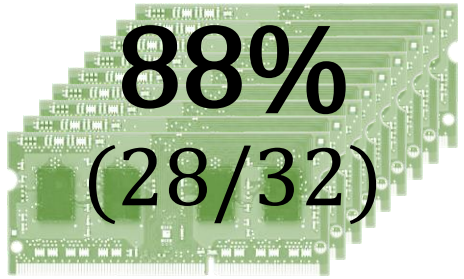 88% (28/32)
Up to 1.0×10^{7} errors	Up to 2.7×10^{6} errors	Up to 3.3×10^{5} errors

Oh!!

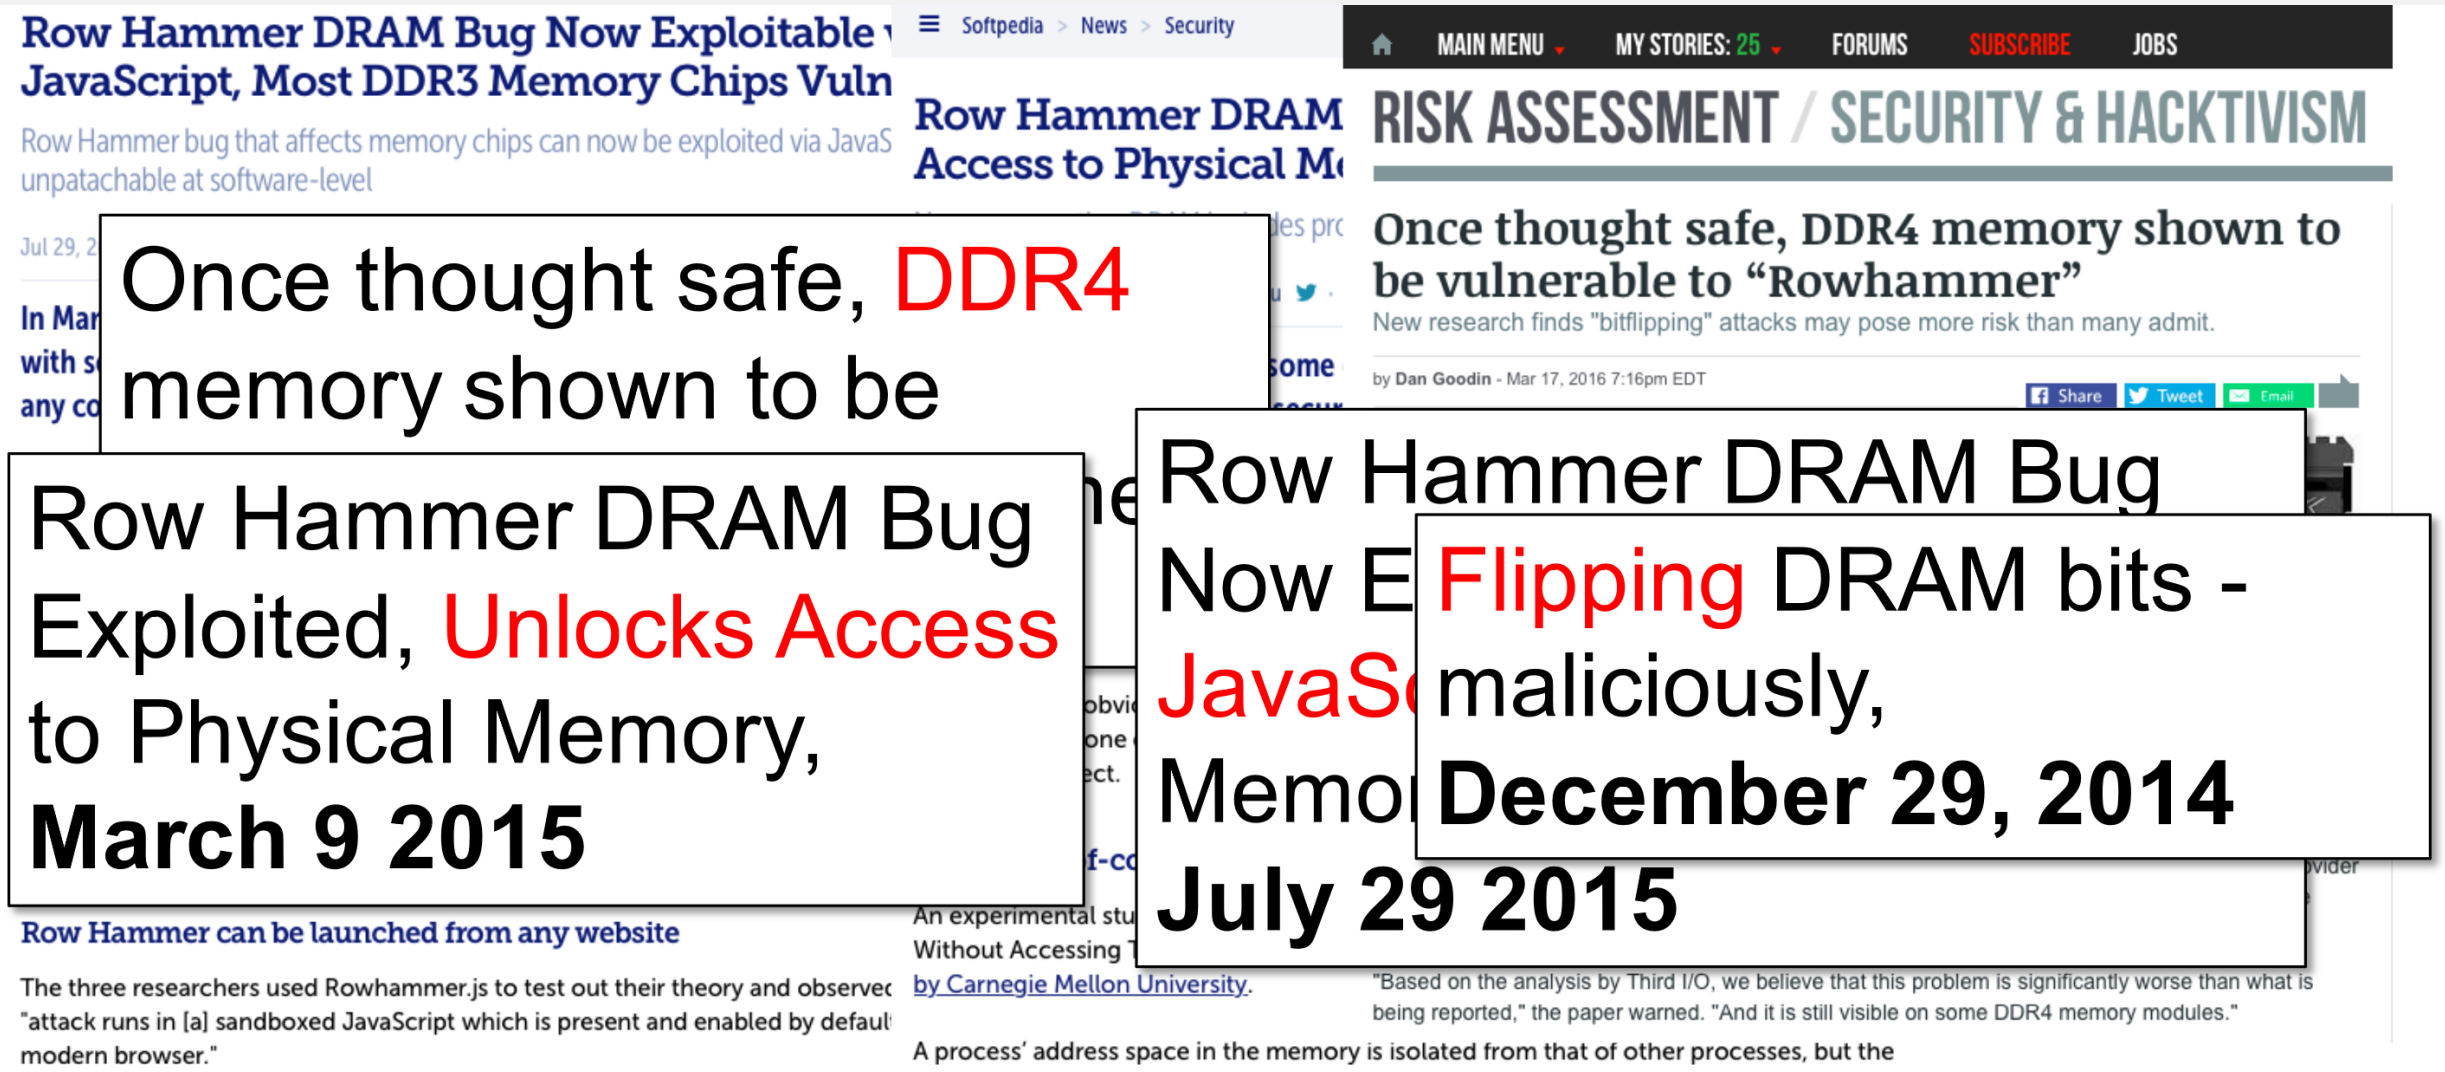
Row Hammer DRAM Bug Now Exploitable
JavaScript, Most DDR3 Memory Chips Vuln
Row Hammer bug that affects memory chips can now be exploited via JavaS
unpatachable at software-level
Softpedia > News > Security
Row Hammer DRAM
Access to Physical M
MAIN MENU
MY STORIES: 25
FORUMS
SUBSCRIBE
JOBS
RISK ASSESSMENT / SECURITY & HACKTIVISM
Once thought safe, DDR4 memory shown to be vulnerable to "Rowhammer"
New research finds "bitflipping" attacks may pose more risk than many admit.
by Dan Goodin - Mar 17, 2016 7:16pm EDT
Share
Tweet
Email
Once thought safe, DDR4 memory shown to be
Row Hammer DRAM Bug Exploited, Unlocks Access to Physical Memory, March 9 2015
Row Hammer DRAM Bug Now E
JavaS
Memo
July 29 2015
Flipping DRAM bits - maliciously, December 29, 2014
Row Hammer can be launched from any website
The three researchers used Rowhammer.js to test out their theory and observed
"attack runs in [a] sandboxed JavaScript which is present and enabled by defaul
modern browser."
An experimental stu
Without Accessing
by Carnegie Mellon University.
"Based on the analysis by Third I/O, we believe that this problem is significantly worse than what is being reported," the paper warned. "And it is still visible on some DDR4 memory modules."
A process' address space in the memory is isolated from that of other processes, but the

Naïve Solutions

1. *Throttle accesses to same row*
   - Limit access-interval: $\geq 500\text{ns}$
   - Limit number of accesses: $\leq 128\text{K}$ $(=64\text{ms}/500\text{ns})$

2. *Refresh more frequently*
   - Shorten refresh-interval by $\sim 7\text{x}$

Both naive solutions introduce significant overhead in performance and power