


# CPU Timing Attacks

①

Side channels : reveal info  
in an "unintended" way  
request



(timing), RF signals, power, audio  
size of messages

Also called metadata

Cache affected by req/reply  
at time  $t$ , affects req/reply  
at time  $t + T$ .

# Examples

(2)

TEMPEST: Screen showing info (CRT screen). Electron beam produced EM emissions.

Linux dm-crypt: encrypt contents of disk on the fly. CPU cache timing attack: Derbat kernel key leak

Smart cards: power analysis. CPU runs code depending on the key. "0" and "1" use different ops & power.

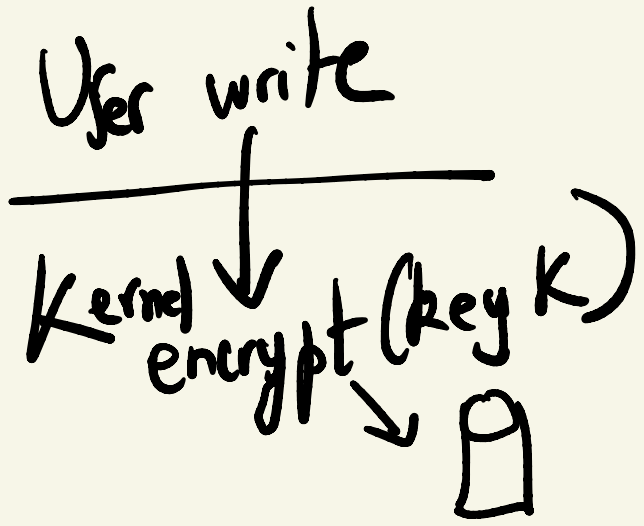
## Hertz Bleed (2022)

Browser page cache: use timing in Javascript to decide if user has visited a page or not.

# Spectre

- complicated, will build up to it.
- Cache timing (dm-crypt)
  - speculative execution (melt down)
  - branch prediction (Spectre)
- prog. control flow ←  
faulting instr. ←


# Linux dm-crypt



# AES Encryption

Key schedule

key K →



data  
 ↓  
 encrypt →  $b \text{ XOR } \text{kschedule}_i = b'$   
 $b' \text{ XOR } \text{kschedule}_j = b''$   
 ...

①

Assume adversary knows  $b$

②

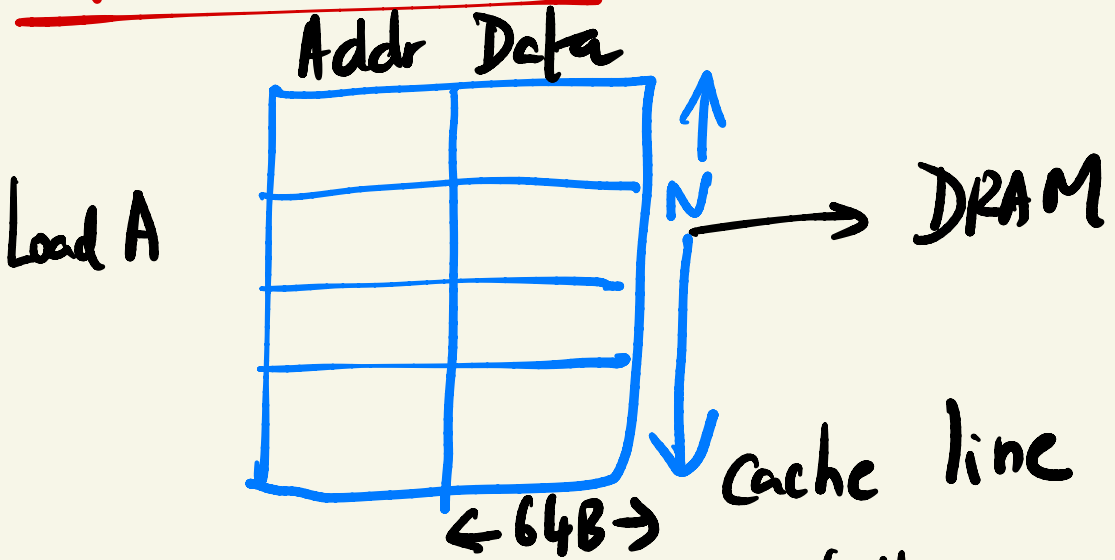
Adversary sees memory accesses  
 i.e.,  $i, j$

$i, j$  depend on  $K$  and  $b, b''$

figure this out!

# CPU Caches

5



A : any place in cache, fully associative : expensive!

"hash" A :  $A \% N$  is going to be in cache. where it

Cache timing affects exploit collisions  
only 64B can be stored in a row (in general a multiple of 64B)

# AES Timing Attack.

Addr

① Fill the cache

② Issue a write to kernel for data  $d$

Kernel

runs encryption on  $d$  and brings some data into cache from key schedule

③ Access the same memory locations as in ①.

noisy game

repeat

Record how long they take. If fast, wasn't evicted by kernel. If slow, was evicted. Figure out address of data that kernel accessed.

info about  $i, j$

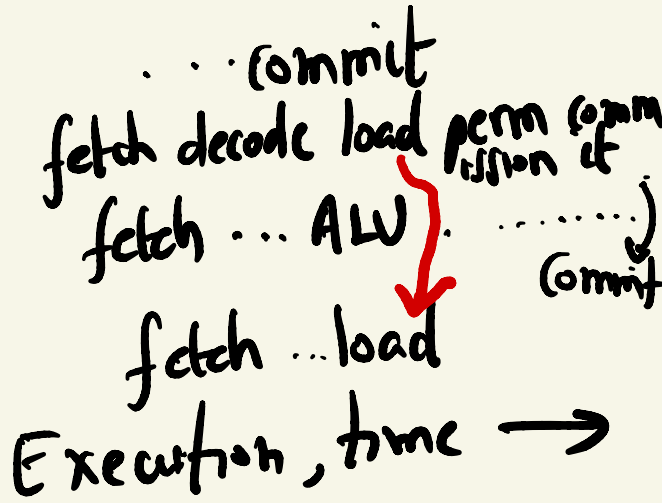
# Speculative Execution

```

add r1, r1, r2
mov [a] -> r1
add r2, r3, r4
mov [r1] -> r3

```

can take 100's of cycle



commits always in order

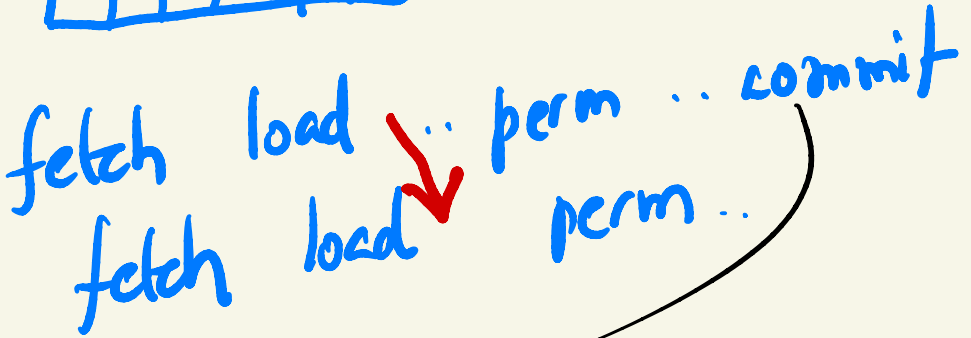


- ① mis-speculation → undo r4
- ② CPU forwards values between instructions → in pic. above

# Meltdown

speculative protection checks

User

```
mov [a], r1  
mov array[r1], r2
```



will commit, but next instr is kernel exception handler, not mov array[r1] modifies cache contents!

side effects not undone



# Back to Cache Timing Attack

⑨

→ Fill cache

```
[ mov [a], r1  
  mov array[r1], r2
```

→ Check what's in cache.

Can scan kernel memory. Kernel maps all of physical memory to its space → know this. Access well known memory locations, e.g., screen.

→ don't see `array[r1]` but know part of `r1` value, which is value at address `a`

Secret-dependent memory accesses are a no-no.

# Mitigations for M.D.

⑩

Hardware: don't do perm check after load, don't forward values

Software: Linux has 2 page tables, one for user processes and some kernel pages, other for kernel.  
sys call  $\rightarrow$  change the page table

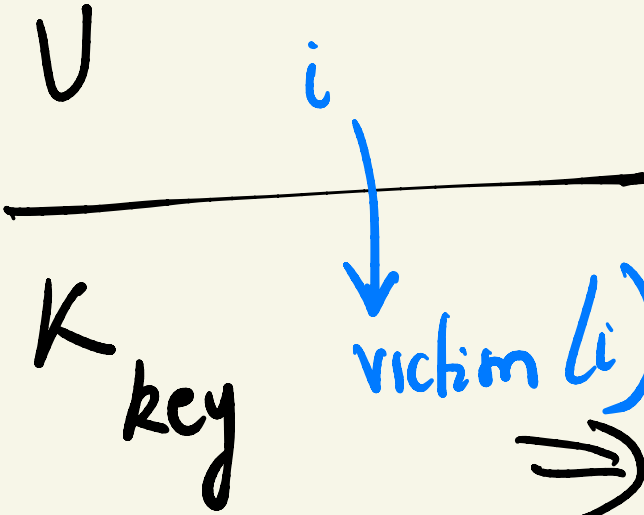
## Branch Prediction

cmp r1, r2      fetch...cmp.  
beg foo      ??  
add r1, r2, r3      ↑  
foo: sub ...      f(PC of beg)

need to abort on misprediction

# Spectre Attack

11



victim is a gadget

meltdown

$\left\{ \begin{array}{l} \text{mov } [i] \rightarrow r1 \\ \text{mov } \text{array}[r1] \rightarrow r2 \end{array} \right\}$

similar

if  $(i < SZ)$  {  
 $v = \text{array}[i];$   
 $w = \text{array2}[v];$

- a) attacker wants to go inside the if when  $i \geq SZ$ !
- b) Make sure  $SZ$  is not cached.
- c) Massage branch predictor to predict  $i < SZ$
- d)  $\text{array2}[v]$  loaded into cache

getting info on v

# Mitigations for Spectre

(12)

## Software

C1: reduce accuracy of branch in JavaScript by adding jitter.

C2: Use ternary instructions like ifence on both outcomes of a branch.

C3: WebKit replaces array bound checks with index masking ( $> \text{page val}$ )

Linux: array  $[i \% N]$  software change