# Advanced Computer Architectures

Spectre, Meltdown and others.. (2018)

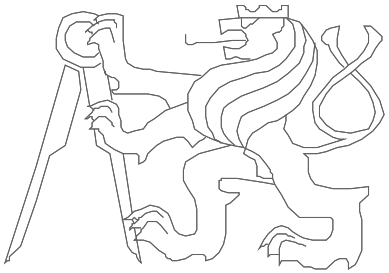

Czech Technical University in Prague, Faculty of Electrical Engineering
Slides authors: Karel Kočí

# What you already know

- There are multiple tasks sharing processor time
- Memory mapping and CPU modes are used to separate memory access between tasks
- Processors use pipeline, cache, speculative execution and other techniques to improve performance
- When speculation is following invalid branch then such execution has to be abandoned and all changes reverted

# What if all changes are not reverted?

# Spectre

- Disclosed: January 2018
- CVE-2017-5753 (1. 2. 2017)
- At least 4 variants and multiple sub-variants
- Fully affects Intel and some ARM processors
- AMD and some ARM processors are affected by some of the variants
- Mitigation in kernel → performance lost
- Requires CPU redesign

# Spectre V1

Attacker

```
for(i=0; i<1000; i++)
        syscall(0);

clfush();
syscall(1000);

for(i=0; i<CACHE_SIZE; i++)
        measure_access_time(i);
```

Kernel

```
int syscall(int idx) {
        if (idx < 16)
                t = a[b[idx]];
}
```

# Spectre V2

Attacker

```
for(i=0; i<1000; i++)
        struct->func();

clflush();
syscall();

for(i=0; i<CACHE_SIZE; i++)
        measure_access_time(i);
```

Kernel

```
int syscall() {
        object->method()
}
```

# Spectre V1 and V2 mitigation

- V1: Mathematical limiting array indexing

  a[b[idx & 0xF]]

- V2: Branch prediction flush and Retpoline

```
RM_RETPOLINE:
    call MAKE_TARGET
CAPTURE_SPEC:
    pause
    jmp CAPTURE_SPEC
MAKE_TARGET:
    mov %R11, (%RSP)
    RET
```

# Meltdown

- Disclosed: January 2018
- CVE-2017-5754 (1. 2. 2017)
- Affects Intel and ARM Cortex-A75
- Exploits check for supervisor bit and speculation
- Mitigation by changing CR3 (Page Table) on context switch
- Mitigation results to big performance lost

```
uint8_t exploit_meltdown(long idx) {
        clflush();
        if (mispredicted)
                a[kernel_addr[idx]] = 1;
        return measure_cache(a);
}
```

# LazyFP

- Lazy FPU switching:
  - CPU saves state of FPU coprocessor when tasks are switched
  - Instead of doing that always it disables FPU support and on exception it enables/restores FPU
- On speculation CPU does not check FPU state
- With double-indirection attacker can snoop FPU register content of previous process
- CVE-2018-3665 (21. 6. 2018)
- Solution: always save, always restore
- Thanks to FXSAVE and FXRSTOR it does not cost that much

# POP SS

- CVE-2018-8897 (5. 8. 2018)
- Kernel has to use interrupt stack for all inter.

---

POP SS
Interrupt? → crash! → magically disable interupts
POP SP

---

TF=1
POP SS
SYSCALL

---

Kernel: ?? exception!

# TLBleed

- Disclosed: June 2018
- Experimentally demonstrated on Intel

- Similar to: Percival attack (2005)
- Allows follow program execution on second CPU thread

- Exploits TLB (transaction lookaside buffer)

- Won’t be fixed and is probably not exploitable in real live
- Cryptography code should be secured because of other exploits

# L1TF (L1 Terminal Fault / Foreshadow)


- Disclosed: 14. 8. 2018
- CVE-2018-3615 (28. 12. 2017)
- Affects Intel CPUs
- Similar to Meltdown but instead of exploiting supervisor bit it exploits present bit
- Easy enough to fix for processes (all non-valid cache lines point to non-existent addresses)
- Hard to mitigate in case of virtualization (flush L1 cache before entering virtual machine)
- No mitigation in case of HyperThreading (disable HyperThreading or VM-VM scheduling)

# MDS (cpu.fail: Zombieload, RIDL, Fallout)

Intel 8th generation affected

- Fallout (CVE-2018–12126):
  Attack on store and forward buffer
  (data to be written to L1)
- Zombieload (CVE-2018–12130):
  Attack on Line Fill buffer
  (data moved from L2 to L1)
- RIDL (CVE-2018–12127):
  Attack on Load buffer
  (data coming from L1)

# Intel Management Unit vulnerabilities

- Because prediction is not only Intel’s problem
- IMU has access to all memory and
- IMU:
  - Full access to hardware (including memory, net..)
  - Not documented, features not disclosed
  - Running MINIX (BSD Unix-like OS)
- Multiple exploits in 2017
- 2018:
  - CVE-2018-3628
  - CVE-2018-3629
  - CVE-2018-3632
- No patch for Intel 3rd generation and older..

# 2019...

- MLPDS/MDSUM (CVE-2019-11091)
- Zombieload 2 (CVE-2019–11135)
- iTLB multihit (CVE-2018–12207)
- Spectre SWAPGS (CVE-2019-1125)
- ...

Intel* in 2018 be like:


*And some other CPU manufacturers

Source: http://gunshowcomic.com/648

# Sources

- https://www.youtube.com/watch?v=rwbs-PN0Vpw
- https://www.youtube.com/watch?v=mIKSXv0Cgjg
- https://www.youtube.com/watch?v=smkzWwtjzkE
- https://meltdownattack.com/
- https://foreshadowattack.eu/
- https://arxiv.org/abs/1806.07480
- https://www.vusec.net/projects/tlbleed/
- https://cpu.fail/