

Software-based Microarchitectural Attacks

Daniel Gruss

April 19, 2018

Graz University of Technology

- Daniel Gruss
- Post-Doc @ Graz University of Technology
- Twitter: @lavados
- Email: `daniel.gruss@iaik.tugraz.at`

- Both vulnerabilities existed for many years

- Both vulnerabilities existed for **many years**
- No one discovered it before

- Both vulnerabilities existed for **many years**
- No one discovered it before
- Suddenly, **4** independent teams discover it within **6** months

- Both vulnerabilities existed for **many years**
- No one discovered it before
- Suddenly, **4** independent teams discover it within **6** months
- Let's create an evidence board

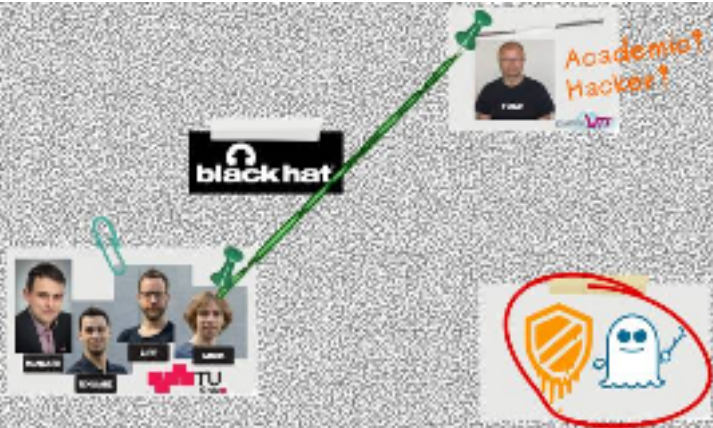







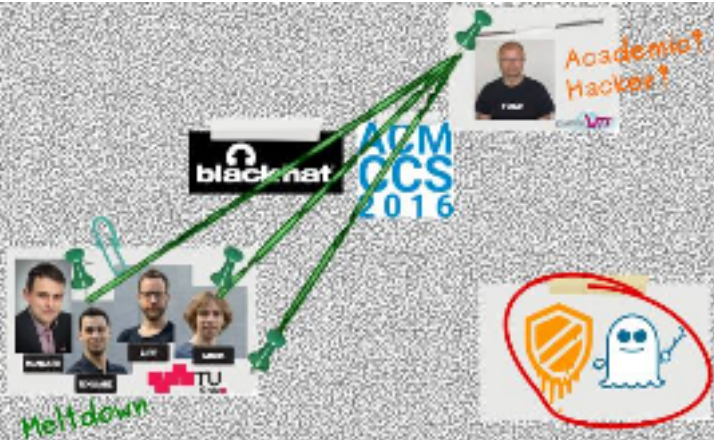






























Why two names, two papers, etc?

- Two different problems





Why two names, two papers, etc?

- Two different problems
- They only have a very loose connection



Why two names, two papers, etc?

- Two different problems
- They only have a very loose connection
- Two different teams had already quite matured drafts ready when learning of each other



Why two names, two papers, etc?

- Two different problems
- They only have a very loose connection
- Two different teams had already quite matured drafts ready when learning of each other
- Initially we tried to merge, but all co-authors quickly agreed that it would mix things that don't belong together

→ More on that after we understand the attacks



You realize it is something big when...



You realize it is something big when...

- it is in the **news**, all over the world









SECURITY FLAW REVEALED

Intel (Prev)

45.26

-1.59

[-3.39%]

Intel (After Hours)

44.85

-0.41

[-0.91%]

CAPITAL
CONNECTION

SHROUT: ISSUE NOT UNIQUE TO
INTEL, BUT IT'S AFFECTED THE MOST

CNBC



You realize it is something big when...

- it is in the **news**, all over the world
- you get a **Wikipedia** article in multiple languages



WIKIPEDIA
The Free Encyclopedia

Main page

Contents

Featured content

Current events

Random article

Donate to Wikipedia

Wikipedia store

Interaction

Help

About Wikipedia

Community portal

Recent changes

Contact page

Tools

Wikipedia links

Wikipedia changes

Wikipedia files

Not logged in - Talk - Contributions - Create account - Log in

Article - Talk

Read - Edit - View history

Search Wikipedia



Meltdown (security vulnerability)

From Wikipedia, the free encyclopedia

Meltdown is a hardware vulnerability affecting Intel x86 microprocessors and some ARM-based microprocessors.^{[1][2][3]} It allows a rogue process to read all memory, even when it is not authorized to do so.

Meltdown affects a wide range of systems. At the time of disclosure, this included all devices running any but the most recent and patched versions of x86,^[4] Linux^{[5][6]}, macOS,^[4] or Windows. Accordingly, many servers and cloud services were impacted,^[7] as well as a potential majority of smart devices and embedded devices using ARM-based processors (mobile devices, smart TVs and others), including a wide range of networking equipment. A purely software workaround to Meltdown has been assessed as slowing computers between 5 and 30 percent in certain specialized workloads,^[8] although companies responsible for software correction of the exploit are reporting reduced impact from general benchmark testing.^[9]

Meltdown was issued a Common Vulnerabilities and Exposures ID of CVE-2017-

0754 et al., also known as *Rogue Data Cache Load*,^[10] in January 2018. It was disclosed in conjunction with another exploit, Spectre, with which it shares some, but not all, characteristics. The Meltdown and Spectre vulnerabilities are considered "catastrophic"



The logo used by the team that discovered the vulnerability



Spectre (security vulnerability)

From Wikipedia, the free encyclopedia

Spectre is a vulnerability that affects modern microprocessors that perform branch prediction.^{[1][2][3]} On most processors, the speculative execution resulting from a branch misprediction may leave observable side effects that may reveal private data to attackers. For example, if the pattern of memory accesses performed by such speculative execution depends on private data, the resulting state of the data cache constitutes a side channel through which an attacker may be able to extract information about the private data using a timing attack.^{[4][5][6]}

Two CoreMark Vulnerabilities and Exposures (CvEs) related to Spectre, CVE-2017-5753^[7] (bounds check bypass) and CVE-2017-5715^[8] (branch target injection), have been issued.^[9] JIT engines used for JavaScript were found vulnerable. A website can read data stored in the browser for another website, or the browser's memory itself.^[10]

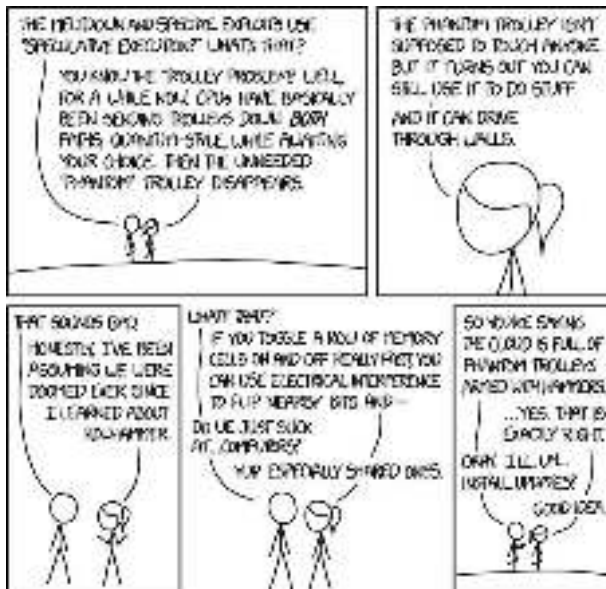
Several procedures to help protect home computers and related devices from the Spectre (and Meltdown) security vulnerabilities have been published.^{[11][12][13]} Spectre patches have been reported to significantly slow down performance, especially on older computers; on the newer 8th generation Core platforms, benchmark performance drops of 2–14 percent have been measured.^[14] Meltdown patches may also produce performance loss.^{[15][16][17]} On January 18, 2018, unwanted reboots, even for newer Intel chips, due to






You realize it is something big when...

- it is in the **news**, all over the world
- you get a **Wikipedia** article in multiple languages
- there are **comics**, including xkcd






Cartoonist: [illegible]



You realize it is something big when...

- it is in the **news**, all over the world
- you get a **Wikipedia** article in multiple languages
- there are **comics**, including xkcd
- you get a lot of **Twitter** follower after Snowden mentioned you



Edward Snowden

@Snowden



You may have heard about Ughie's horrific file down bug. But have you watched it in action? When your computer asks you to apply updates this month, don't click "not now" (via [specifications.com](#) & [@misc0110](#))


23:27 - 4 Jun 2018

👁 162 🗨 6643 ❤ 6,542







- Kernel is isolated from user space




- Kernel is isolated from user space
- This **isolation** is a combination of hardware and software



- Kernel is isolated from user space
- This **isolation** is a combination of hardware and software
- User applications cannot access anything from the kernel



- Kernel is isolated from user space
- This **isolation** is a combination of hardware and software
- User applications cannot access anything from the kernel
- There is only a well-defined interface → **syscalls**











1337 4242

FOOD CACHE

Revolutionary concept!

Store your food at home,
never go to the grocery store
during cooking.

Can store **ALL** kinds of food.



ONLY TODAY INSTEAD OF ~~\$1,300~~

\$1,299

ORDER VIA PHONE: +555 12345





```
printf("%d", i);  
printf("%d", i);
```




```
printf("%d", i);  
printf("%d", i);
```

Cache miss




```
printf("%d", i);  
printf("%d", i);
```

Cache miss




Request





```
printf("%d", i);  
printf("%d", i);
```

Cache miss




Request

Response




```
printf("%d", i);  
printf("%d", i);
```

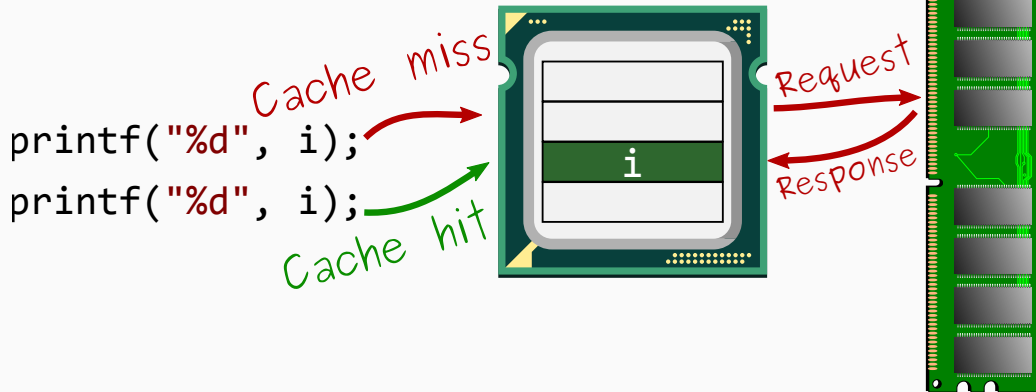
Cache miss




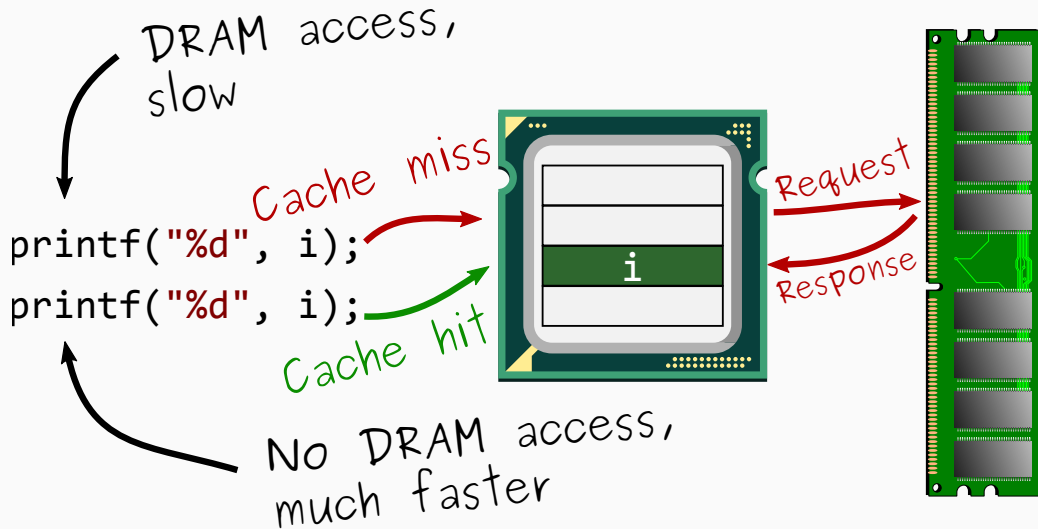
Request

Response










Shared Memory


ATTACKER

flush
access



VICTIM


access



Shared Memory

ATTACKER

flush
access




VICTIM

access

Shared Memory

ATTACKER

flush
access




VICTIM

access

Shared Memory

ATTACKER

flush
access




VICTIM

access

Shared Memory

ATTACKER

flush
access



VICTIM


access



Shared Memory

ATTACKER

flush
access



VICTIM


access

Shared Memory

ATTACKER

flush


access





VICTIM

access

fast if victim accessed data,
slow otherwise











6. Cook everything until
vegetables are soft

6. Add spices to taste
and stir for 10 minutes

7. Serve with cooked
and peeled potatoes





Wait for an hour



Wait for an hour

LATENCY



1. Wash and cut
vegetables

2. Pick the basil leaves
and set aside

3. Heat 2 tablespoons of
oil in a pan

4. Fry vegetables until
golden and softened



Dependency

1. Wash and cut
vegetables

2. Pick the basil leaves
and set aside

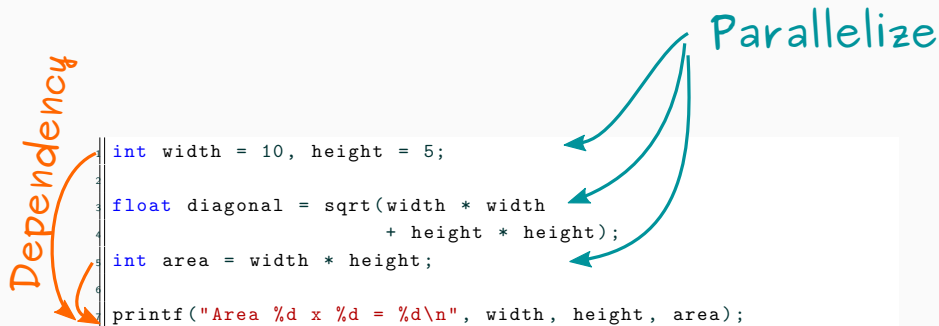
3. Heat 2 tablespoons of
oil in a pan

4. Fry vegetables until
golden and softened

Parallelize



```
1 int width = 10, height = 5;
2
3 float diagonal = sqrt(width * width
4                       + height * height);
5 int area = width * height;
6
7 printf("Area %d x %d = %d\n", width, height, area);
```



```
1 | char data = *(char*)0xffffffff81a000e0;  
2 | printf("%c\n", data);
```



```
1 char data = *(char*)0xffffffff81a000e0;  
2 printf("%c\n", data);
```



```
1 segfault at ffffffff81a000e0 ip 0000000000400535  
2          sp 00007ffce4a80610 error 5 in reader
```

```
1 char data = *(char*)0xffffffff81a000e0;  
2 printf("%c\n", data);
```



```
1 segfault at ffffffff81a000e0 ip 0000000000400535  
2          sp 00007ffce4a80610 error 5 in reader
```

- Kernel addresses are not accessible

```
1 char data = *(char*)0xffffffff81a000e0;  
2 printf("%c\n", data);
```



```
1 segfault at ffffffff81a000e0 ip 0000000000400535  
2          sp 00007ffce4a80610 error 5 in reader
```

- Kernel addresses are not accessible
- Are privilege checks also done when executing instructions out of order?



- Adapted code

```
1 | *(volatile char*)0;  
2 | array[84 * 4096] = 0; // unreachable
```




- Adapted code

```
1 | *(volatile char*)0;  
2 | array[84 * 4096] = 0; // unreachable
```

- Static code analyzer is not happy

```
1 | warning: Dereference of null pointer  
2 |     *(volatile char*)0;
```



- Flush+Reload over all pages of the array



- “Unreachable” code line was actually executed



- Flush+Reload over all pages of the array



- “Unreachable” code line was actually executed
- Exception was only thrown afterwards



- Combine the two things

```
1 char data = *(char*)0xffffffff81a000e0;  
2 array[data * 4096] = 0;
```



- Combine the two things

```
1 char data = *(char*)0xffffffff81a000e0;  
2 array[data * 4096] = 0;
```

= sending end of a cache covert channel

- Then check whether any part of array is cached



- Combine the two things

```
1 char data = *(char*)0xffffffff81a000e0;  
2 array[data * 4096] = 0;
```

= sending end of a cache covert channel

- Then check whether any part of array is cached
= receiving end of a cache covert channel



- Flush+Reload over all pages of the array




- Index of cache hit reveals data



- Flush+Reload over all pages of the array



- Index of cache hit reveals data
- Permission check is in some cases not fast enough



A man and a woman are shown in a close-up, looking off-camera with serious expressions. The scene is dimly lit with a blueish-grey tone. A small, bright red light source is visible near the man's ear. The background is dark and out of focus, showing some architectural lines.


CAN YOU
ENHANCE THAT

```
meltdown@meltdown ~/ppm2 % taskset 1 ./imgdump 0x375a00000 14919 > output.flif
```

```
Reading from 0xffff880375a00000
```

```
[
```

I



AND IN OTHER NEWS...




WE'RE ALL DOOMED, SANDRA.
HOW ABOUT THE WEATHER?



Not so fast...

- Kernel addresses in user space are a problem

- Kernel addresses in user space are a problem
- Why don't we take the kernel addresses...





- ...and remove them if not needed?



- ...and remove them if not needed?
- User accessible check in hardware is not reliable



- Let's just unmap the kernel in user space



- Let's just unmap the kernel in user space
- Kernel addresses are then no longer present



- Let's just unmap the kernel in user space
- Kernel addresses are then no longer present
- Memory which is not mapped cannot be accessed at all





Kernel **A**ddress **I**solation to have **S**ide channels **E**fficiently **R**emoved


KAISER /ˈkʌɪzə/

1. [german] Emperor, ruler of an empire
2. largest penguin, emperor penguin



Kernel **A**ddress **I**solation to have **S**ide channels **E**fficiently **R**emoved

Userspace




Applications

Kernelspace

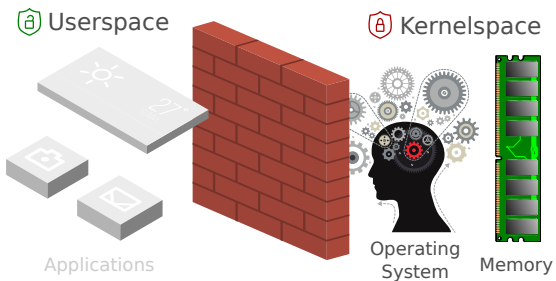


Operating
System




Memory


Kernel View




User View




context switch




- We published **KAISER** in July 2017




- We published **KAISER** in July 2017
- Intel and others improved and merged it into Linux as **KPTI** (Kernel Page Table Isolation)



- We published **KAISER** in July 2017
- Intel and others improved and merged it into Linux as **KPTI** (Kernel Page Table Isolation)
- Microsoft implemented similar concept in Windows 10



- We published **KAISER** in July 2017
- Intel and others improved and merged it into Linux as **KPTI** (Kernel Page Table Isolation)
- Microsoft implemented similar concept in Windows 10
- Apple implemented it in macOS 10.13.2 and called it “**Double Map**”



- We published **KAISER** in July 2017
- Intel and others improved and merged it into Linux as **KPTI** (Kernel Page Table Isolation)
- Microsoft implemented similar concept in Windows 10
- Apple implemented it in macOS 10.13.2 and called it “**Double Map**”
- All share the same idea: switching address spaces on context switch

A cartoon illustration of a man with spiky orange hair and glasses, looking upwards and to the right with a thoughtful expression. The background is a solid blue color.

WAIT A MOMENT...

DUPLICATING EVERYTHING? THAT
SOUNDS REALLY SLOW



- Depends on how often you need to switch between kernel and user space



- Depends on how often you need to switch between kernel and user space
- Can be slow, 40% or more on old hardware




- Depends on how often you need to switch between kernel and user space
- Can be slow, 40% or more on old hardware
- But modern CPUs have additional features



- Depends on how often you need to switch between kernel and user space
- Can be slow, 40% or more on old hardware
- But modern CPUs have additional features
- \Rightarrow Performance overhead on average below 2%




MELTDOWN



SPECTRE



MELTDOWN



SPECTRE







Prosciutto




Funghi



Diavolo



Diavolo



Diavolo



Diavolo

»A table for 6 please«






Speculative Cooking



»A table for 6 please«











- Mistrains branch prediction



- Mistrains branch prediction
- CPU speculatively executes code which should not be executed



- Mistrains branch prediction
- CPU speculatively executes code which should not be executed
- Can also mistrain indirect calls




- Mistrains branch prediction
 - CPU speculatively executes code which should not be executed
 - Can also mistrain indirect calls
- Spectre “convinces” program to execute code



```
index = 0;
```

```
char* data = "textKEY";
```



```
index = 0;
```

```
char* data = "textKEY";
```

```
if (index < 4)
```

then

```
LUT[data[index] * 4096]
```



Prediction

else

```
0
```

```
index = 0;
```

```
char* data = "textKEY";
```

```
if (index < 4)
```

then

```
LUT[data[index] * 4096]
```



```
index = 0;
```

```
char* data = "textKEY";
```

```
if (index < 4)
```

Execute

then

```
LUT[data[index] * 4096]
```



Prediction

else

0

```
index = 1;
```

```
char* data = "textKEY";
```

```
if (index < 4)
```

then



Prediction

else

```
LUT[data[index] * 4096]
```

```
0
```

```
index = 1;
```



```
char* data = "textKEY";
```

```
if (index < 4)
```

then




```
LUT[data[index] * 4096]
```



Prediction

else



```
0
```

```
index = 1;
```

```
char* data = "textKEY";
```

```
if (index < 4)
```

Speculate

then

```
LUT[data[index] * 4096]
```



Prediction

else

0


```
index = 1;
```



```
char* data = "textKEY";
```

```
if (index < 4)
```

then




```
LUT[data[index] * 4096]
```



Prediction

else



```
0
```

```
index = 2;
```

```
char* data = "textKEY";
```

```
if (index < 4)
```

then

else



Prediction

```
LUT[data[index] * 4096]
```

```
0
```

```
index = 2;
```

```
char* data = "textKEY";
```

```
if (index < 4)
```

then

```
LUT[data[index] * 4096]
```



Prediction

else

```
0
```

```
index = 2;
```

```
char* data = "textKEY";
```

```
if (index < 4)
```

Speculate

then

```
LUT[data[index] * 4096]
```



Prediction

else

0

```
index = 2;
```

```
char* data = "textKEY";
```

```
if (index < 4)
```

then

```
LUT[data[index] * 4096]
```




Prediction

else

```
0
```

```
index = 3;
```

```
char* data = "textKEY";
```



```
index = 3;
```

```
char* data = "textKEY";
```

```
if (index < 4)
```

then

```
LUT[data[index] * 4096]
```



else

```
0
```

```
index = 3;
```

```
char* data = "textKEY";
```

```
if (index < 4)
```

Speculate

then

```
LUT[data[index] * 4096]
```



else

0


```
index = 3;
```

```
char* data = "textKEY";
```

```
if (index < 4)
```

then

```
LUT[data[index] * 4096]
```




Prediction

else

```
0
```

```
index = 4;
```

```
char* data = "textKEY";
```



```
index = 4;
```

```
char* data = "textKEY";
```

```
if (index < 4)
```

then

```
LUT[data[index] * 4096]
```




Prediction

else

```
0
```

```
index = 4;
```

```
char* data = "textKEY";
```



```
index = 4;
```

```
char* data = "textKEY";
```

```
if (index < 4)
```

then

```
LUT[data[index] * 4096]
```



else

Execute

0

```
index = 5;
```

```
char* data = "textKEY";
```

```
if (index < 4)
```

then

else



Prediction

```
LUT[data[index] * 4096]
```

```
0
```

```
index = 5;
```

```
char* data = "textKEY";
```

```
if (index < 4)
```

then

```
LUT[data[index] * 4096]
```



Prediction

else

```
0
```

```
index = 5;
```

```
char* data = "textKEY";
```

```
if (index < 4)
```

Speculate

then

```
LUT[data[index] * 4096]
```



else

0


```
index = 5;
```

```
char* data = "textKEY";
```

```
if (index < 4)
```

then

```
LUT[data[index] * 4096]
```



else

Execute

0

```
index = 6;
```

```
char* data = "textKEY";
```

```
if (index < 4)
```

then

else



Prediction

```
LUT[data[index] * 4096]
```

```
0
```

```
index = 6;
```

```
char* data = "textKEY";
```

```
if (index < 4)
```

then

```
LUT[data[index] * 4096]
```




Prediction

else

```
0
```

```
index = 6;
```


```
char* data = "textKEY";
```




```
index = 6;
```

```
char* data = "textKEY";
```


```
if (index < 4)
```




```
LUT[data[index] * 4096]
```



```
Animal* a = bird;
```




```
Animal* a = bird;
```



```
Animal* a = bird;
```


a->move()



LUT[data[index] * 4096]


0


```
Animal* a = bird;
```



```
Animal* a = bird;
```


a->move()



LUT[data[index] * 4096]


0

```
Animal* a = bird;
```



```
Animal* a = bird;
```

a->move()




LUT[data[index] * 4096]

0

```
Animal* a = fish;
```


a->move()



LUT[data[index] * 4096]


0

```
Animal* a = fish;
```



```
Animal* a = fish;
```


a->move()



LUT[data[index] * 4096]


0

```
Animal* a = fish;
```




```
Animal* a = fish;
```

a->move()



LUT[data[index] * 4096]

0



- Trivial approach: disable speculative execution



- Trivial approach: disable speculative execution
- No wrong speculation if there is no speculation



- Trivial approach: disable speculative execution
- No wrong speculation if there is no speculation
- Problem: massive performance hit!



- Trivial approach: disable speculative execution
- No wrong speculation if there is no speculation
- Problem: massive performance hit!
- Also: How to disable it?



- Trivial approach: disable speculative execution
- No wrong speculation if there is no speculation
- Problem: massive performance hit!
- Also: How to disable it?
- Speculative execution is deeply integrated into CPU





- Workaround: insert instructions stopping speculation




- Workaround: insert instructions stopping speculation
→ insert after every bounds check




- Workaround: insert instructions stopping speculation
- insert after every bounds check
- x86: LFENCE, ARM: CSDB




- Workaround: insert instructions stopping speculation
→ insert after every bounds check
- x86: LFENCE, ARM: CSDB
- Available on all Intel CPUs, retrofitted to existing ARMv7 and ARMv8







- Speculation barrier requires compiler supported



- Speculation barrier requires compiler supported
- Already implemented in GCC, LLVM, and MSVC



- Speculation barrier requires compiler supported
- Already implemented in GCC, LLVM, and MSVC
- Can be automated (MSVC) → not really reliable



- Speculation barrier requires compiler supported
- Already implemented in GCC, LLVM, and MSVC
- Can be automated (MSVC) → not really reliable
- Explicit use by programmer: `__builtin_load_no_speculate`


```
// Unchecked  
  
int array[4];  
  
int get_value(unsigned int n) {  
    int tmp;  
  
    if (n < 4) {  
        tmp = array[n];  
    } else {  
        tmp = 0; // FAIL!  
    }  
  
    return tmp;  
}
```

```
// Unprotected
```

```
int array[N];
```

```
int get_value(unsigned int n) {  
    int tmp;
```

```
    if (n < N) {  
        tmp = array[n];  
    } else {  
        tmp = FAIL;
```

```
    }  
    return tmp;
```

```
}
```

```
// Protected
```

```
int array[N];
```

```
int get_value(unsigned int n) {
```

```
    int *lower = array;  
    int *ptr = array + n;  
    int *upper = array + N;
```

```
    return  
        __builtin_load_no_speculate  
        (ptr, lower, upper, FAIL);
```

```
}
```





- Speculation barrier works if affected code constructs are known




- Speculation barrier works if affected code constructs are known
- Programmer has to fully understand vulnerability



- Speculation barrier works if affected code constructs are known
- Programmer has to fully understand vulnerability
- Automatic detection is not reliable




- Speculation barrier works if affected code constructs are known
- Programmer has to fully understand vulnerability
- Automatic detection is not reliable
- Non-negligible performance overhead of barriers




Intel released microcode updates

- Indirect Branch Restricted Speculation (IBRS):




Intel released microcode updates

- Indirect Branch Restricted Speculation (IBRS):
 - Do not speculate based on anything before entering IBRS mode




Intel released microcode updates

- Indirect Branch Restricted Speculation (IBRS):
 - Do not speculate based on anything before entering IBRS mode
 - lesser privileged code cannot influence predictions




Intel released microcode updates

- Indirect Branch Restricted Speculation (IBRS):
 - Do not speculate based on anything before entering IBRS mode
→ lesser privileged code cannot influence predictions
- Indirect Branch Predictor Barrier (IBPB):




Intel released microcode updates

- Indirect Branch Restricted Speculation (IBRS):
 - Do not speculate based on anything before entering IBRS mode
→ lesser privileged code cannot influence predictions
- Indirect Branch Predictor Barrier (IBPB):
 - Flush branch-target buffer



Intel released microcode updates


- Indirect Branch Restricted Speculation (IBRS):
 - Do not speculate based on anything before entering IBRS mode
→ lesser privileged code cannot influence predictions
- Indirect Branch Predictor Barrier (IBPB):
 - Flush branch-target buffer
- Single Thread Indirect Branch Predictors (STIBP):




Intel released microcode updates

- Indirect Branch Restricted Speculation (IBRS):
 - Do not speculate based on anything before entering IBRS mode
→ lesser privileged code cannot influence predictions
- Indirect Branch Predictor Barrier (IBPB):
 - Flush branch-target buffer
- Single Thread Indirect Branch Predictors (STIBP):
 - Isolates branch prediction state between two hyperthreads

Retpoline (compiler extension)




Retpoline (compiler extension)



```
1      push <call_target>
2      call 1f
3
4      2:                                ; speculation will continue here
5      lfence                            ; speculation barrier
6      jmp 2b                            ; endless loop
7
8      1:                                ; the actual call to <call_target>
9      lea 8(%rsp), %rsp ; restore stack pointer
10     ret
```

→ always predict to enter an endless loop

Retpoline (compiler extension)




```
1      push <call_target>
2      call 1f
3
4      2:                                ; speculation will continue here
5      lfence                            ; speculation barrier
6      jmp 2b                            ; endless loop
7
8      1:                                ; the actual call to <call_target>
9      lea 8(%rsp), %rsp ; restore stack pointer
10     ret
```

→ always predict to enter an endless loop

- instead of the correct (or wrong) target function

Retpoline (compiler extension)




```
1      push <call_target>
2      call 1f
3
4      2:                                ; speculation will continue here
5      lfence                            ; speculation barrier
6      jmp 2b                            ; endless loop
7
8      1:                                ; the actual call to <call_target>
9      lea 8(%rsp), %rsp ; restore stack pointer
10     ret
```

→ always predict to enter an endless loop

- instead of the correct (or wrong) target function → performance?

Retpoline (compiler extension)




```
1      push <call_target>
2      call 1f
3
4      2:                                ; speculation will continue here
5      lfence                            ; speculation barrier
6      jmp 2b                            ; endless loop
7
8      1:                                ; the actual call to <call_target>
9      lea 8(%rsp), %rsp ; restore stack pointer
10     ret
```

→ always predict to enter an endless loop

- instead of the correct (or wrong) target function → performance?
- On Broadwell or newer:

Retpoline (compiler extension)




```
1      push <call_target>
2      call 1f
3
4      2:                                ; speculation will continue here
5      lfence                            ; speculation barrier
6      jmp 2b                            ; endless loop
7
8      1:                                ;
9      lea 8(%rsp), %rsp ; restore stack pointer
10     ret                               ; the actual call to <call_target>
```

→ always predict to enter an endless loop

- instead of the correct (or wrong) target function → performance?
- On Broadwell or newer:
 - **ret** may fall-back to the BTB for prediction


Retpoline (compiler extension)




```
1      push <call_target>
2      call 1f
3
4      2:                                ; speculation will continue here
5      lfence                            ; speculation barrier
6      jmp 2b                            ; endless loop
7
8      1:                                ;
9      lea 8(%rsp), %rsp ; restore stack pointer
10     ret                               ; the actual call to <call_target>
```

→ always predict to enter an endless loop


- instead of the correct (or wrong) target function → performance?
 - On Broadwell or newer:
 - **ret** may fall-back to the BTB for prediction
- microcode patches to prevent that




- ARM provides hardened Linux kernel




- ARM provides hardened Linux kernel
- Clears branch-predictor state on context switch



- ARM provides hardened Linux kernel
- Clears branch-predictor state on context switch
- Either via instruction (BPIALL)...



- ARM provides hardened Linux kernel
- Clears branch-predictor state on context switch
- Either via instruction (BPIALL)...
- ...or workaround (disable/enable MMU)



- ARM provides hardened Linux kernel
- Clears branch-predictor state on context switch
- Either via instruction (BPIALL)...
- ...or workaround (disable/enable MMU)
- Non-negligible performance overhead ($\approx 200\text{-}300\text{ ns}$)



- Prevent access to high-resolution timer



- Prevent access to high-resolution timer
- Own timer using timing thread



- Prevent access to high-resolution timer
- Own timer using timing thread
- Flush instruction only privileged



- Prevent access to high-resolution timer
- Own timer using timing thread
- Flush instruction only privileged
- Cache eviction through memory accesses



- Prevent access to high-resolution timer
- Own timer using timing thread
- Flush instruction only privileged
- Cache eviction through memory accesses
- Just move secrets into secure world



- Prevent access to high-resolution timer
- Own timer using timing thread
- Flush instruction only privileged
- Cache eviction through memory accesses
- Just move secrets into secure world
- Spectre works on secure enclaves

Meltdown

Spectre

Meltdown

- Out-of-Order Execution

Spectre

- Speculative Execution (subset of Out-of-Order Execution)

Meltdown

- Out-of-Order Execution
- has **nothing to do with branch prediction**

Spectre

- **Speculative** Execution (subset of Out-of-Order Execution)
- fundamentally builds on branch (mis)prediction

Meltdown

- Out-of-Order Execution
- has **nothing to do with branch prediction**
- turning off speculative execution **entirely** has no effect on Meltdown

Spectre

- **Speculative Execution** (subset of Out-of-Order Execution)
- fundamentally builds on branch (mis)prediction
- turning off speculative execution **entirely** would work

Meltdown

- Out-of-Order Execution
 - has **nothing to do with branch prediction**
 - turning off speculative execution **entirely** has no effect on Meltdown
- **melts down** the isolation provided by the `user_accessible-bit`

Spectre

- **Speculative Execution** (subset of Out-of-Order Execution)
- fundamentally builds on branch (mis)prediction
- turning off speculative execution **entirely** would work
- has nothing to do with the `user_accessible-bit`

Meltdown

- Out-of-Order Execution
 - has **nothing to do with branch prediction**
 - turning off speculative execution **entirely** has no effect on Meltdown
- **melts down** the isolation provided by the `user_accessible-bit`
- in theory: OoO not required, pipelining can be sufficient

Spectre

- **Speculative Execution** (subset of Out-of-Order Execution)
- fundamentally builds on branch (mis)prediction
- turning off speculative execution **entirely** would work
- has nothing to do with the `user_accessible-bit`
- KAISER has no effect on Spectre at all

Meltdown

- Out-of-Order Execution
 - has **nothing to do with branch prediction**
 - turning off speculative execution **entirely** has no effect on Meltdown
- **melts down** the isolation provided by the `user_accessible-bit`
- in theory: OoO not required, pipelining can be sufficient
 - mitigated by KAISER

Spectre

- **Speculative Execution** (subset of Out-of-Order Execution)
- fundamentally builds on branch (mis)prediction
- turning off speculative execution **entirely** would work
- has nothing to do with the `user_accessible-bit`
- KAISER has no effect on Spectre at all

Meltdown

Spectre

Meltdown

- performs illegal memory accesses → we need to take care of processor exceptions

Spectre

- performs only legal memory accesses

Meltdown

- performs illegal memory accesses → we need to take care of processor exceptions
 - exception handling

Spectre

- performs only legal memory accesses
 - has nothing to do with exception handling

Meltdown

- performs illegal memory accesses → we need to take care of processor exceptions
 - exception handling
 - exception suppression with TSX

Spectre

- performs only legal memory accesses
 - has nothing to do with exception handling or suppression

Meltdown

- performs illegal memory accesses → we need to take care of processor exceptions
 - exception handling
 - exception suppression with TSX
 - exception suppression with branch misprediction

Spectre

- performs only legal memory accesses
 - has nothing to do with exception handling or suppression


Meltdown

- performs illegal memory accesses → we need to take care of processor exceptions
 - exception handling
 - exception suppression with TSX
 - exception suppression with branch misprediction


Spectre


- performs only legal memory accesses
 - has nothing to do with exception handling or suppression

→ two papers, two names, etc.




... why were they named variant 1, 2 and 3 by Google?






... why were they named variant 1, 2 and 3 by Google?

- “How can you use speculative execution maliciously?”




... why were they named variant 1, 2 and 3 by Google?

- “How can you use speculative execution maliciously?”
- Intel had much interest in not fancy-naming them ;)



... why were they named variant 1, 2 and 3 by Google?


- “How can you use speculative execution maliciously?”
- Intel had much interest in not fancy-naming them ;)



... why were they named variant 1, 2 and 3 by Google?

- “How can you use speculative execution maliciously?”
- Intel had much interest in not fancy-naming them ;)

... why were they presented on the same date and on the same website?




... why were they named variant 1, 2 and 3 by Google?

- “How can you use speculative execution maliciously?”
- Intel had much interest in not fancy-naming them ;)

... why were they presented on the same date and on the same website?

- We did not choose the date



... why were they named variant 1, 2 and 3 by Google?

- “How can you use speculative execution maliciously?”
- Intel had much interest in not fancy-naming them ;)

... why were they presented on the same date and on the same website?

- We did not choose the date
- We did not want to have one of them overshadow the other immediately



We have ignored microarchitectural attacks for many many years:



We have ignored microarchitectural attacks for many many years:

- attacks on crypto



We have ignored microarchitectural attacks for many many years:

- attacks on crypto → “software should be fixed”



We have ignored microarchitectural attacks for many many years:

- attacks on crypto → “software should be fixed”
- attacks on ASLR



We have ignored microarchitectural attacks for many many years:

- attacks on crypto → “software should be fixed”
- attacks on ASLR → “ASLR is broken anyway”



We have ignored microarchitectural attacks for many many years:

- attacks on crypto → “software should be fixed”
- attacks on ASLR → “ASLR is broken anyway”
- attacks on SGX and TrustZone



We have ignored microarchitectural attacks for many many years:

- attacks on crypto → “software should be fixed”
- attacks on ASLR → “ASLR is broken anyway”
- attacks on SGX and TrustZone → “not part of the threat model”



We have ignored microarchitectural attacks for many many years:

- attacks on crypto → “software should be fixed”
 - attacks on ASLR → “ASLR is broken anyway”
 - attacks on SGX and TrustZone → “not part of the threat model”
- for years we solely optimized for performance



After learning about a side channel you realize:




After learning about a side channel you realize:

- the side channels were documented in the Intel manual



After learning about a side channel you realize:

- the side channels were documented in the Intel manual
- only now we understand the implications



Motor Vehicle Deaths in U.S. by Year



A unique chance to

- rethink processor design



A unique chance to

- rethink processor design
- grow up, like other fields (car industry, construction industry)



A unique chance to

- rethink processor design
- grow up, like other fields (car industry, construction industry)
- dedicate more time into identifying problems and not solely in mitigating known problems

Software-based Microarchitectural Attacks

Daniel Gruss

April 19, 2018

Graz University of Technology