

LVI



Hijacking Transient Execution through Microarchitectural Load Value Injection

Jo Van Bulck¹ **Daniel Moghimi**² **Michael Schwarz**³ **Moritz Lipp**³ **Marina Minkin**⁴
Daniel Genkin⁴ **Yuval Yarom**⁵ **Berk Sunar**² **Daniel Gruss**³ **Frank Piessens**¹

41st IEEE Symposium on Security and Privacy (digital edition) – May 18, 2020

¹imec-DistriNet, KU Leuven ²Worcester Polytechnic Institute ³Graz University of Technology

⁴University of Michigan ⁵University of Adelaide and Data61

Hijacking transient execution?



Spectre



Meltdown

Hijacking transient execution?



Spectre

v1, v2, v4, v5,
Spectre-BTB,
Spectre-RSB,
ret2spec,
SGXPectre,
SmotherSpectre,
NetSpectre?



Meltdown

v3, v3.1, v3a,
RDCL?



ZombieLoad, MDS?



Foreshadow

Foreshadow-NG,
L1TF?




RIDL, Fallout?



A new perspective. . .




A new perspective. . .




BACK TO THE FUTURE



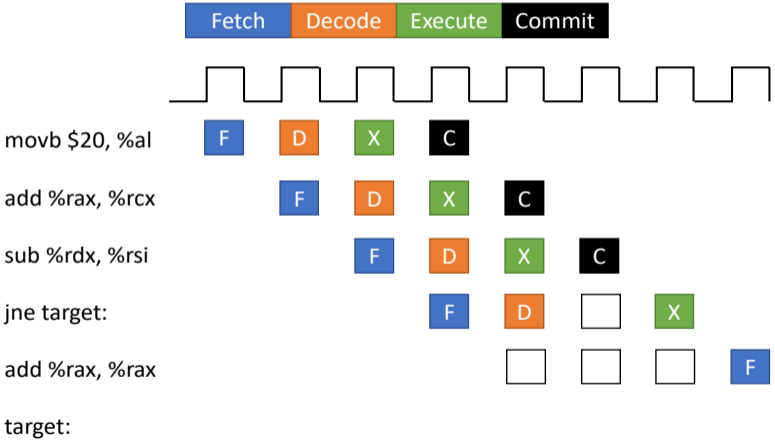
Motivation: Branch prediction to avoid pipeline stalls



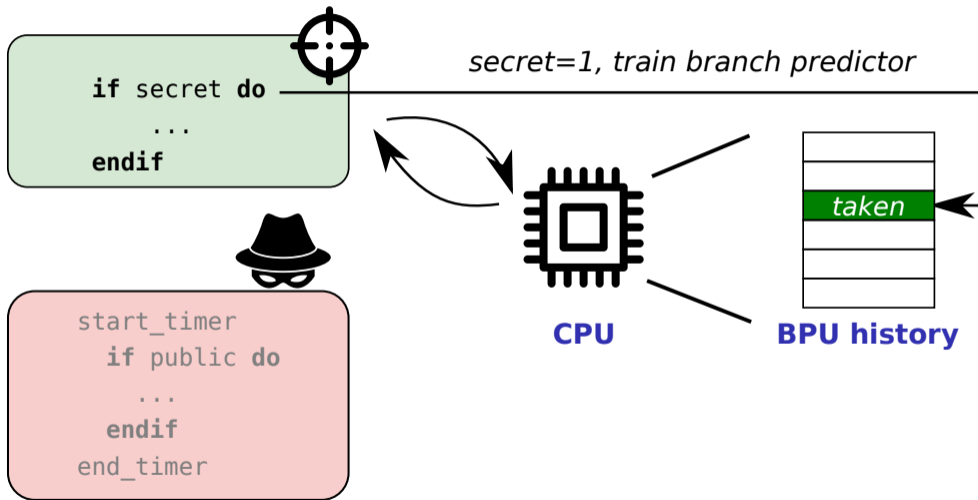
Motivation: Branch prediction to avoid pipeline stalls



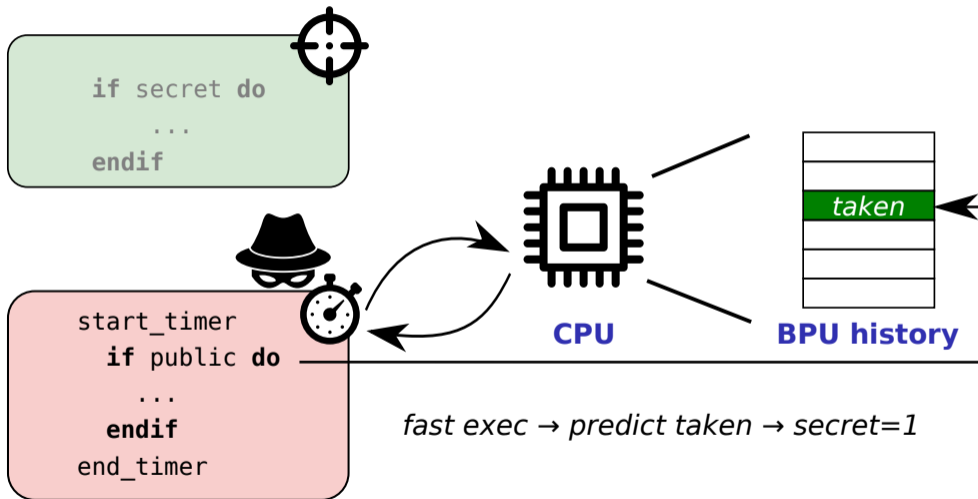
Motivation: Branch prediction to avoid pipeline stalls




Branch prediction as a side channel



Branch prediction as a side channel







?

?

?



?

?

THE WHITE HOUSE
6:14 PM

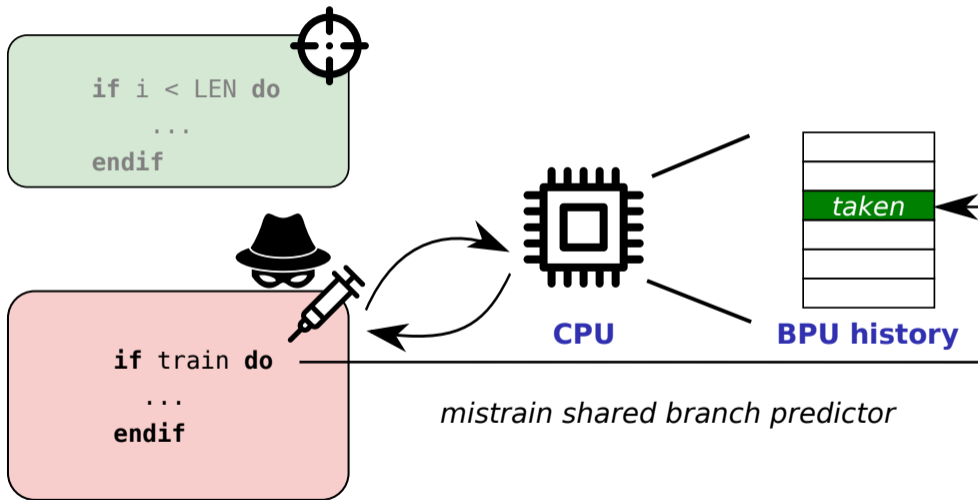


BREAKING NEWS

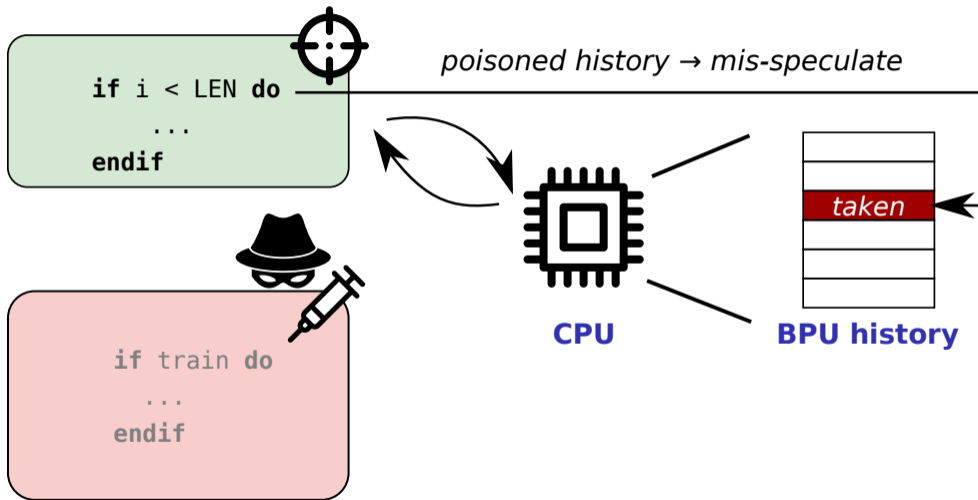
PRES. TRUMP UPDATES PUBLIC ON FEDERAL RESPONSE TO VIRUS

 **MSNBC**


Branch prediction as a transient fault injection primitive



Branch prediction as a transient fault injection primitive




Spectre v1: Speculative buffer over-read



```
if (idx < LEN)
{
  s = buffer[idx];
  t = lookup[s];
  ...
}
```

- Programmer *intention*: no out-of-bounds accesses


Spectre v1: Speculative buffer over-read



```
if (idx < LEN)
{
    s = buffer[idx];
    t = lookup[s];
    ...
}
```

- Programmer *intention*: no out-of-bounds accesses
- **Mistrain gadget** to **speculatively** “ahead of time” execute with $idx \geq LEN$ in the transient world

Spectre v1: Speculative buffer over-read





```
if (idx < LEN)
{
  s = buffer[idx];
  t = lookup[s];
  ...
}
```

- Programmer *intention*: no out-of-bounds accesses
- **Mistrain gadget** to **speculatively** “ahead of time” execute with $idx \geq LEN$ in the transient world
- **Side channels** may leave traces after roll-back!

Spectre take-away


- CPU **transiently** executes wrong code paths
- **Confused-deputy gadgets** encode secrets via side channels






?

?



?

Meltdown: Transiently encoding unauthorized memory



Unauthorized access


Listing 1: x86 assembly

```
1 meltdown:  
2  // %rdi: oracle  
3  // %rsi: secret_ptr  
4  
5  movb (%rsi), %al  
6  shl $0xc, %rax  
7  movq (%rdi, %rax), %rdi  
8  retq
```

Listing 2: C code.

```
1 void meltdown(  
2     uint8_t *oracle,  
3     uint8_t *secret_ptr)  
4 {  
5     uint8_t v = *secret_ptr;  
6     v = v * 0x1000;  
7     uint64_t o = oracle[v];  
8 }
```


Meltdown: Transiently encoding unauthorized memory



Unauthorized access



Transient out-of-order window

Listing 1: x86 assembly.

```
1 meltdown:  
2 // %rdi: oracle  
3 // %rsi: secret_ptr  
4  
5 movb (%rsi), %al  
6 shl $0xc, %rax  
7 movq (%rdi, %rax), %rdi  
8 retq
```

Listing 2: C code.


```
1 void meltdown(  
2     uint8_t *oracle ,  
3     uint8_t *secret_ptr)  
4 {  
5     uint8_t v = *secret_ptr;  
6     v = v * 0x1000;  
7     uint64_t o = oracle[v];  
8 }
```

oracle array



secret idx

Meltdown: Transiently encoding unauthorized memory



Unauthorized access



Transient out-of-order window



Exception

(discard architectural state)


Listing 1: x86 assembly.

```
1 meltdown:  
2 // %rdi: oracle  
3 // %rsi: secret_ptr  
4  
5 movb (%rsi), %al  
6 shl $0xc, %rax  
7 movq (%rdi, %rax), %rdi  
8 retq
```

Listing 2: C code.

```
1 void meltdown(  
2     uint8_t *oracle,  
3     uint8_t *secret_ptr)  
4 {  
5     uint8_t v = *secret_ptr;  
6     v = v * 0x1000;  
7     uint64_t o = oracle[v];  
8 }
```

Meltdown: Transiently encoding unauthorized memory



Unauthorized access



Transient out-of-order window



Exception handler

Listing 1: x86 assembly.

```
1 meltdown:  
2  // %rdi: oracle  
3  // %rsi: secret_ptr  
4  
5  movb (%rsi), %al  
6  shl $0xc, %rax  
7  movq (%rdi, %rax), %rdi  
8  retq
```

Listing 2: C code.


```
1 void meltdown(  
2     uint8_t *oracle,  
3     uint8_t *secret_ptr)  
4 {  
5     uint8_t v = *secret_ptr;  
6     v = v * 0x1000;  
7     uint64_t o = oracle[v];  
8 }
```

oracle array



cache hit

Meltdown variants: Address dependencies



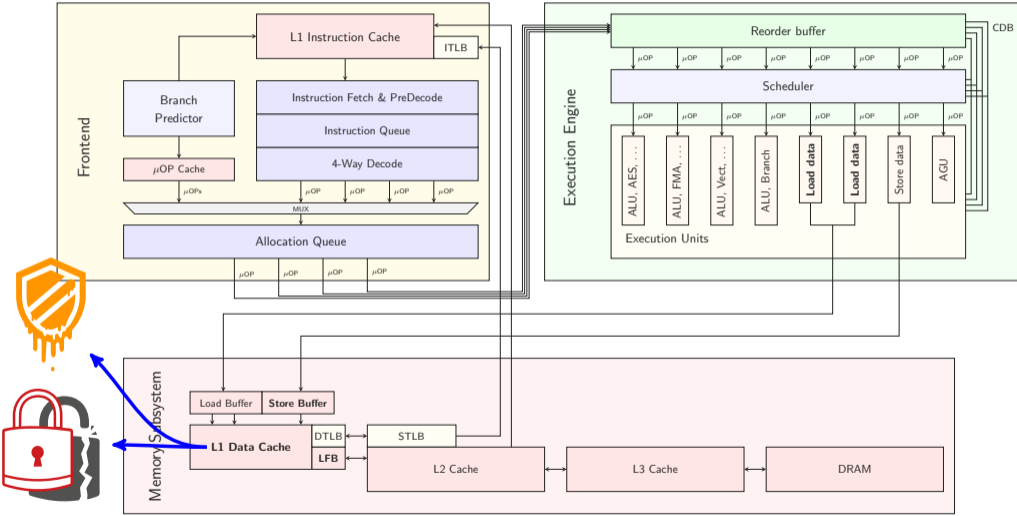
Meltdown variants: Address dependencies

		Page Number		Page Offset
Meltdown	51	Physical	12	11 0
	47	Virtual	12	
Foreshadow	51	Physical	12	11 0
	47	Virtual	12	

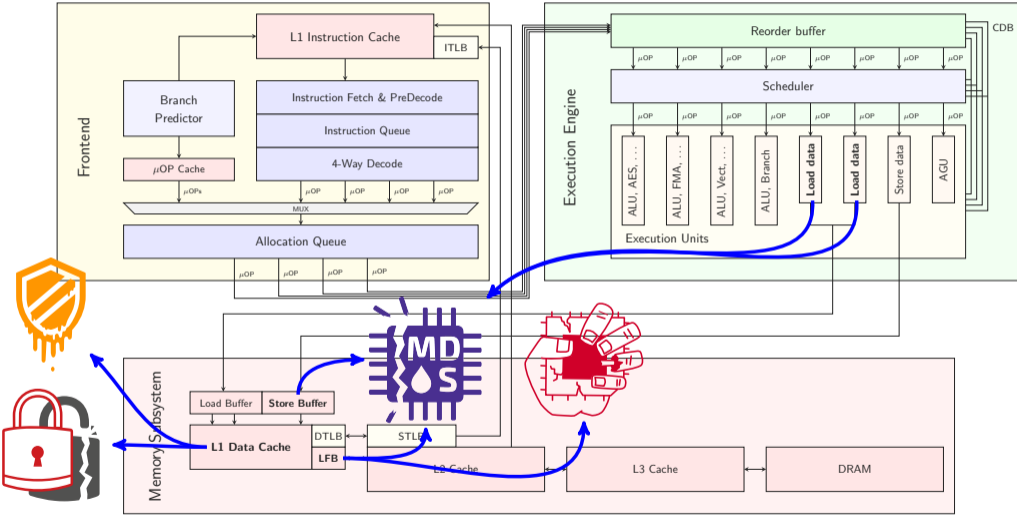
Meltdown variants: Address dependencies

		Page Number		Page Offset			
Meltdown	51	Physical	12	11 0			
	47	Virtual	12				
Foreshadow	51	Physical	12	11 0			
	47	Virtual	12				
Fallout	51	Physical	12	11 0			
	47	Virtual	12				
ZombieLoad/ RIDL	51	Physical	12	11	6	5 0	
	47	Virtual	12				

Meltdown variants: Microarchitectural buffers





Meltdown variants: Microarchitectural buffers





Meltdown take-away



Faulting (or assisted) loads transiently forward **unrelated data** from various microarchitectural buffers






?

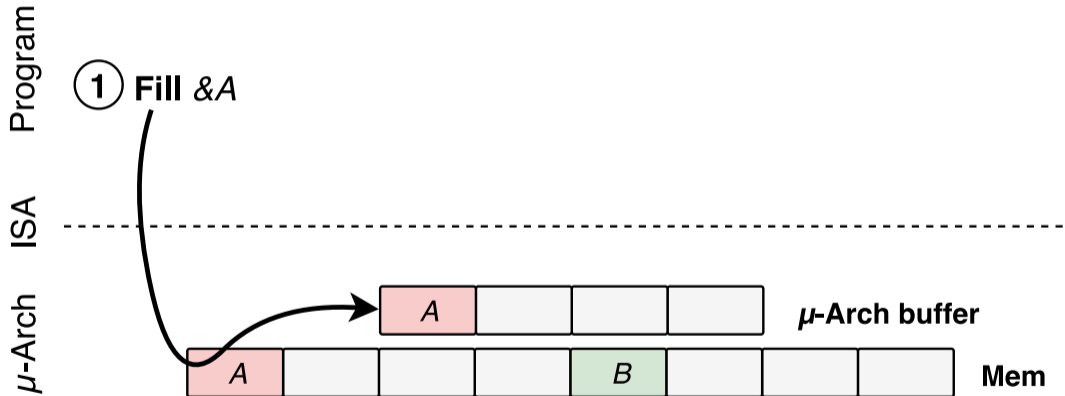




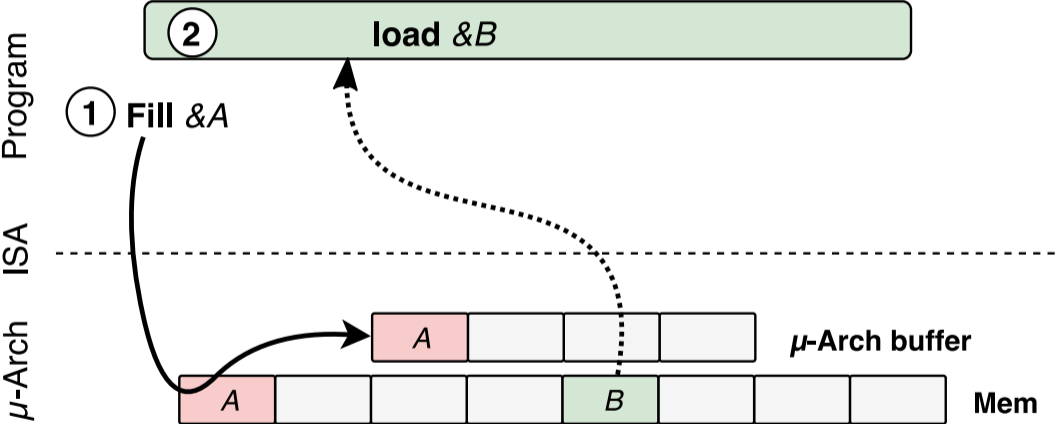
LVI: The basic idea



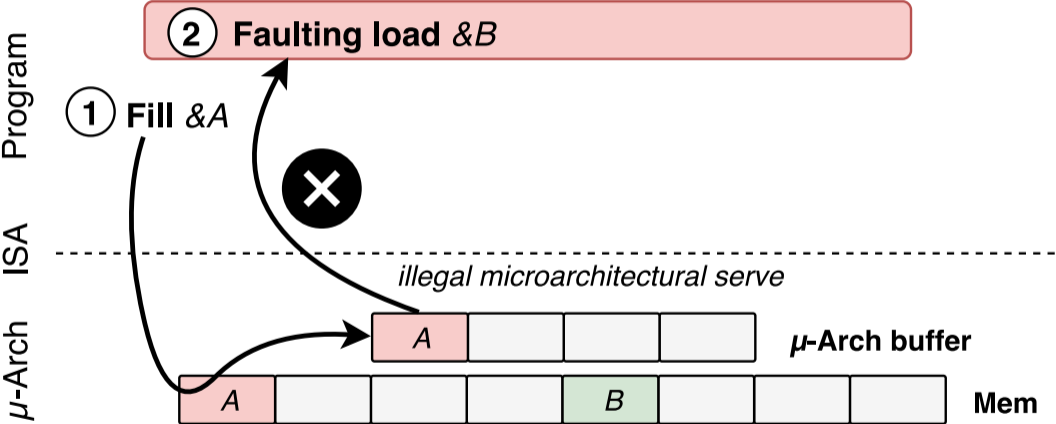
LVI: The basic idea



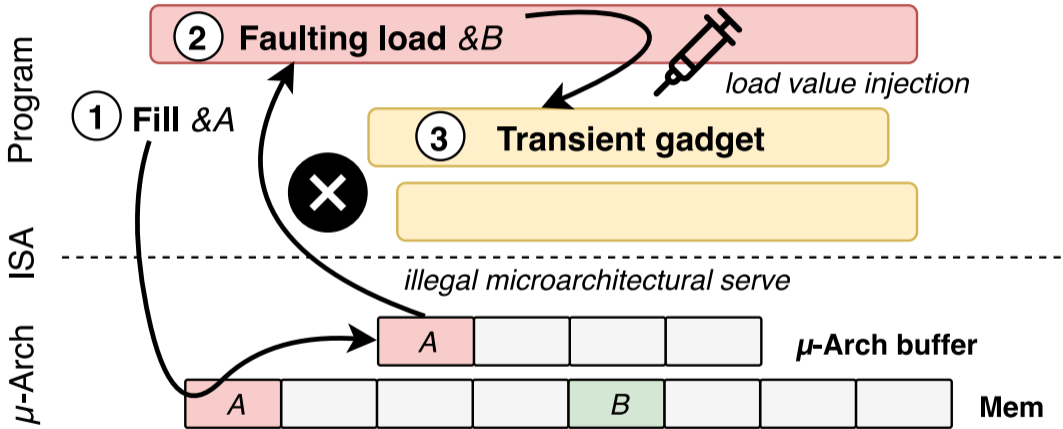
LVI: The basic idea



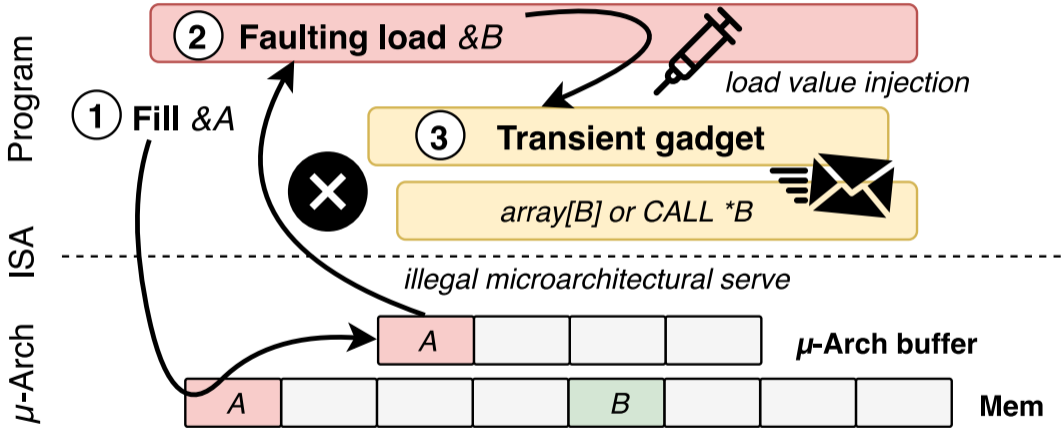
LVI: The basic idea



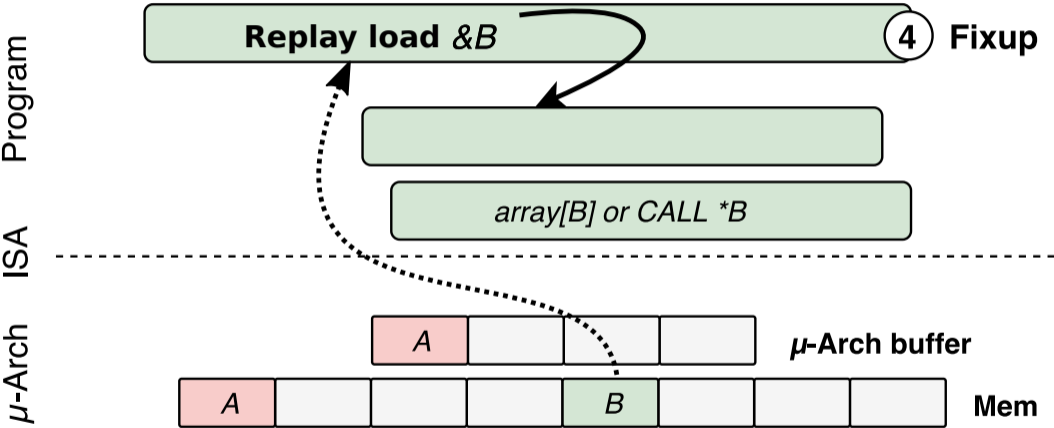
LVI: The basic idea



LVI: The basic idea




LVI: The basic idea



BUT WAIT...


A PAGE FAULT IN THE VICTIM??

Enclaves to the rescue!




Intel SGX promise: hardware-level **isolation and attestation**

Intel SGX: A look under the hood




- **SGX machinery** protects against direct address remapping attacks

Intel SGX: A look under the hood



- **SGX machinery** protects against direct address remapping attacks
- ... but untrusted address translation may **fault** during enclaved execution (!)

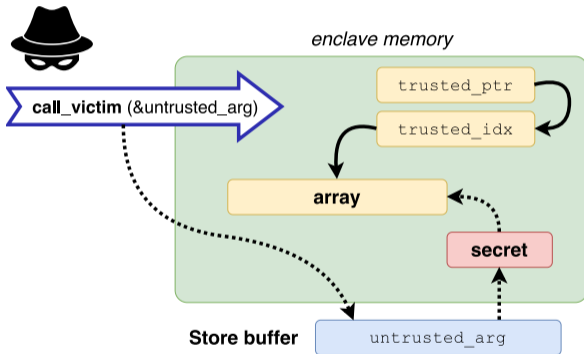
Intel SGX: A look under the hood




We can arbitrarily provoke page faults for trusted enclave loads!

LVI toy example: Hijacking transient data flow

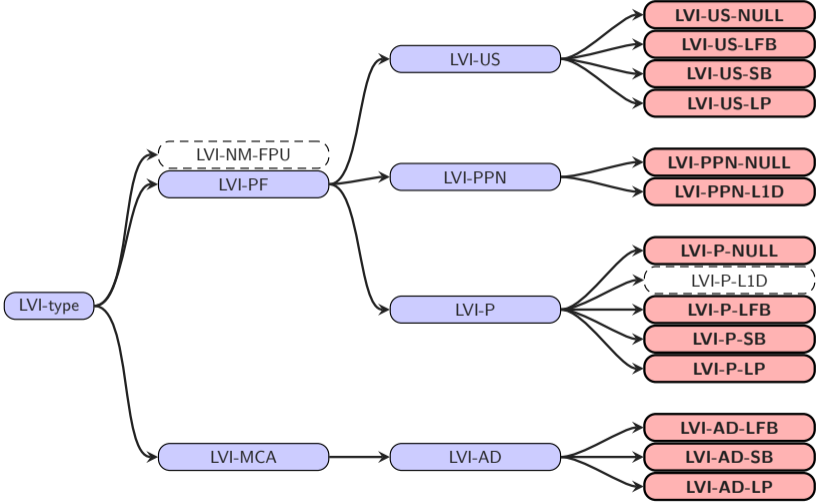
```
1 void call_victim(size_t
2   untrusted_arg)
3 {
4   *arg_copy = untrusted_arg;
5   array[**trusted_ptr * 4096];
6 }
```



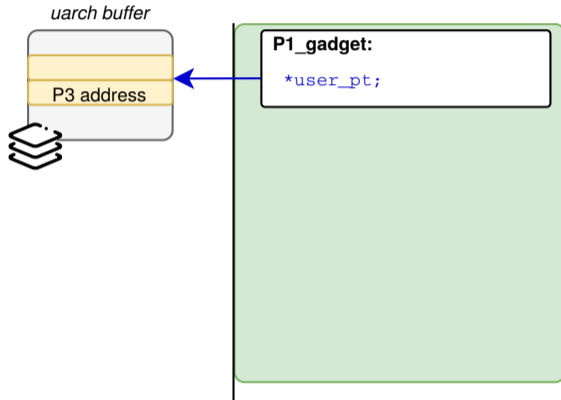
LVI toy example: Recovering arbitrary secrets



Taxonomy of LVI variants: Many buffers, many faults...

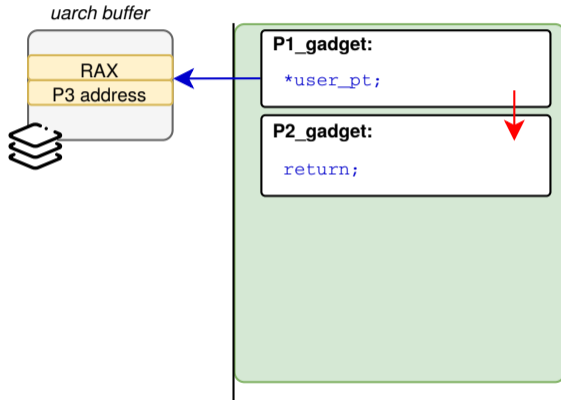


LVI-based transient control-flow hijacking




1. Victim fills μ -arch buffer with attacker-controlled data

LVI-based transient control-flow hijacking



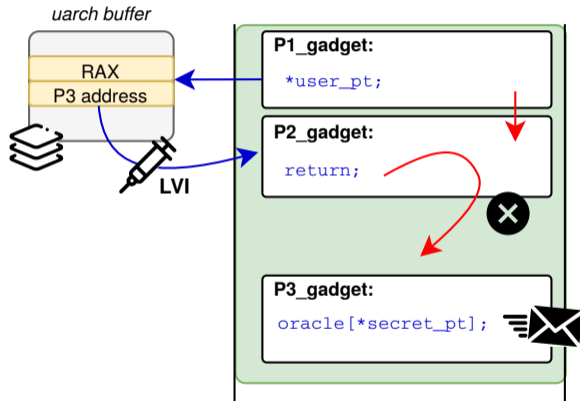
1. Victim fills μ -arch buffer with attacker-controlled data
2. Victim executes indirect branch (JMP/CALL/RET)

LVI-based transient control-flow hijacking



1. Victim **fills μ -arch buffer** with attacker-controlled data
2. Victim executes **indirect branch** (JMP/CALL/RET)
3. **Faulting load** \rightarrow inject incorrect attacker values(!)

LVI-based transient control-flow hijacking




1. Victim fills μ -arch buffer with attacker-controlled data
2. Victim executes indirect branch (JMP/CALL/RET)
3. Faulting load \rightarrow inject incorrect attacker values(!)
4. Redirect transient control flow

```
E/asm.S main.c
28 .global ecall_lvi_sb_rop
29 # %rdi store_pt
30 # %rsi oracle_pt
31 ecall_lvi_sb_rop:
32 mov %rsp, rsp_backup(%rip)
33 lea page_b(%rip), %rsp
34 add $0FFSET, %rsp
35
36 /* transient delay */
37 clflush dummy(%rip)
38 mov dummy(%rip), %rax
39
40 /* STORE TO USER ADRS */
41 movq $'R', (%rdi)
42 lea ret_gadget(%rip), %rax
43 movq %rax, 8(%rdi)
44
45 /* HIJACK TRUSTED LOAD FROM ENCLAVE STACK */
46 /* should go to do_real_ret; will transiently go to ret_gadget if we fault on the stack loads */
47 pop %rax
48 #if LFENCE
49 notq (%rsp)
50 notq (%rsp)
51 lfence
52 ret
53 #else
54 ret
55 #endif
56
57 1: jmp 1b
58 mfence
59
60 do_real_ret:
61 mov rsp_backup(%rip), %rsp
62 ret
63
```


BUT ARE RECENT INTEL CPUS NOT

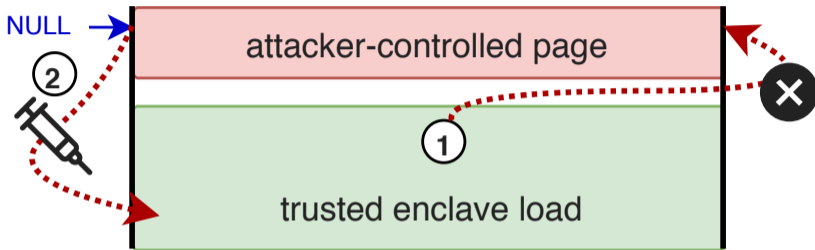
MELTDOWN-RESISTANT?

LVI-NULL: Why 0x00 is not a safe value




- Recent Intel CPUs forward **0x00 dummy values** for faulting loads

LVI-NULL: Why `0x00` is not a safe value



- Recent Intel CPUs forward `0x00` **dummy values** for faulting loads
- ... but NULL is a **valid virtual memory address**, under attacker control

LVI-NULL: Why 0x00 is not a safe value



- Recent Intel CPUs forward **0x00 dummy values** for faulting loads
- ... but `NULL` is a **valid virtual memory address**, under attacker control
- ... hijack **pointer values** (e.g., function pointer-to-pointer)

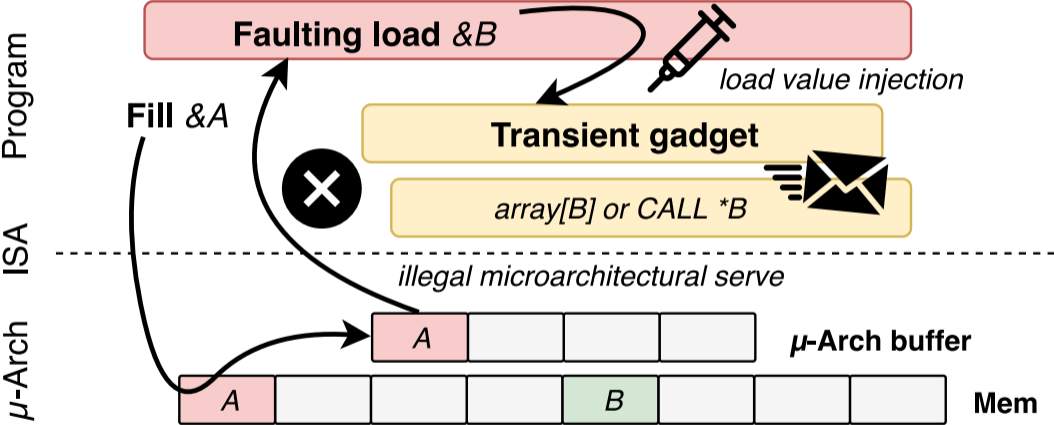


Mitigating LVI?

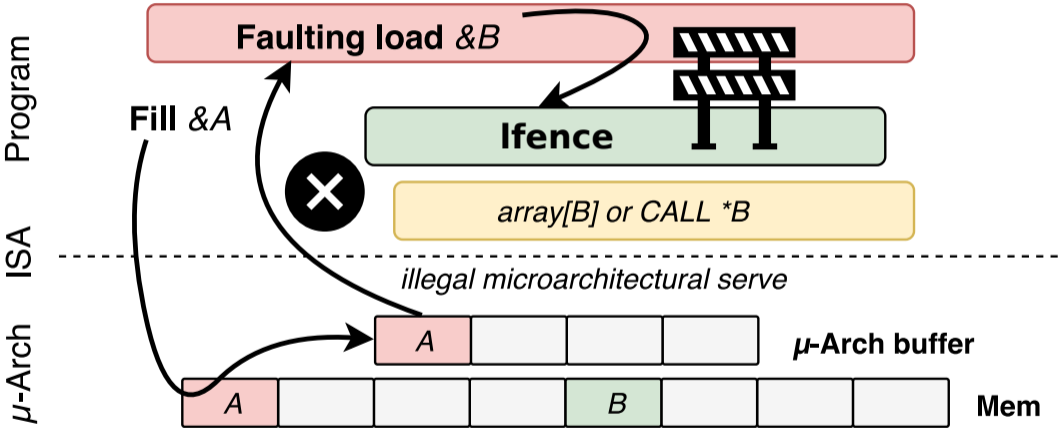


Gadget-based exploits → flushing μ -arch buffers does not suffice!

Mitigation idea: Fencing vulnerable load instructions



Mitigation idea: Fencing vulnerable load instructions




Serializing indirect branches



LVI ↔ Spectre: no control-flow prediction; **every load can be hijacked**

Instruction	Possible emulation	Clobber-free
ret	pop %reg; lfence; jmp *%reg	X
ret	not (%rsp); not (%rsp); lfence; ret	✓
jmp (mem)	mov (mem),%reg; lfence; jmp *%reg	X
call (mem)	mov (mem),%reg; lfence; call *%reg	X



23 fences

October 2019—“surgical precision”



23 fences

October 2019—“surgical precision”




49,315 fences

March 2020—“big hammer”







Performance overheads: OpenSSL (our prototype mitigation)



Performance overheads: OpenSSL (Intel's mitigation)



Performance overheads: SPEC (Intel's mitigation)



GNU Assembler Adds New Options For Mitigating Load Value Injection Attack

Written by [Michael Larabel](#) in [GNU](#) on 11 March 2020 at 02:55 PM EDT. [14 Comments](#)

The Brutal Performance Impact From Mitigating The LVI Vulnerability

Written by [Michael Larabel](#) in [Software](#) on 12 March 2020. **Page 1 of 6.** [76 Comments](#)

LLVM Lands Performance-Hitting Mitigation For Intel LVI Vulnerability

Written by [Michael Larabel](#) in [Software](#) on 3 April 2020. **Page 1 of 3.** [20 Comments](#)

Looking At The LVI Mitigation Impact On Intel Cascade Lake Refresh

Written by [Michael Larabel](#) in [Software](#) on 5 April 2020. **Page 1 of 5.** [10 Comments](#)

- ⇒ LVI **gadgets** reversely exploit Meltdown-type effects
- ⇒ **Short-term:** extensive **lfence compiler mitigations** for Intel SGX enclaves
- ⇒ **Long-term:** improved **silicon patches** in new CPUs



LVI



Hijacking Transient Execution through Microarchitectural Load Value Injection

Jo Van Bulck¹ **Daniel Moghimi**² **Michael Schwarz**³ **Moritz Lipp**³ **Marina Minkin**⁴
Daniel Genkin⁴ **Yuval Yarom**⁵ **Berk Sunar**² **Daniel Gruss**³ **Frank Piessens**¹

41st IEEE Symposium on Security and Privacy (digital edition) – May 18, 2020

¹imec-DistriNet, KU Leuven ²Worcester Polytechnic Institute ³Graz University of Technology

⁴University of Michigan ⁵University of Adelaide and Data61