

CustomProcessingUnit: Reverse Engineering and Customization of Intel Microcode

OffensiveCon 2023

Pietro Borrello

Sapienza University of Rome

Roland Czerny

Graz University of Technology

Catherine Easdon


Dynatrace Research & Graz University of Technology

Michael Schwarz


CISPA Helmholtz Center for Information Security

Martin Schwarzl


Graz University of Technology




- PhD from Sapienza University of Rome (defended two days ago)
- Interested in (very) low-level research
- 3x BlackHat speaker
- 2x Pwnie Award recipient




1. Deep dive on CPU μ code



1. Deep dive on CPU μ code
2. μ code Software Framework



1. Deep dive on CPU μ code
2. μ code **Software Framework**
3. Case Studies: x86 PAC, μ software bp, constant-time div

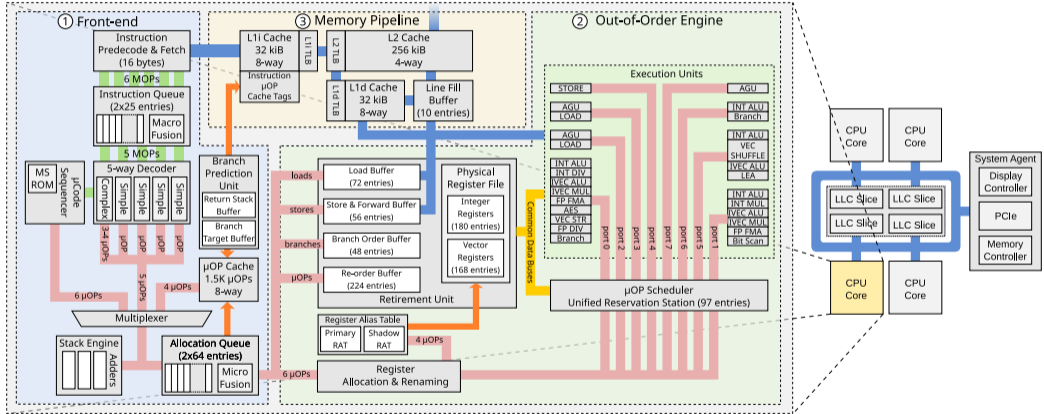


1. Deep dive on CPU μ code
2. μ code **Software Framework**
3. Case Studies: x86 PAC, μ software bp, constant-time div
4. Reverse Engineering of the secret **μ code update algorithm**



- This is based on **our understanding** of CPU Microarchitecture.
- In theory, it may be **all wrong**.
- In practice, a lot **seems right**.

How do CPUs work?





- **Red Unlock** of Atom Goldmont (GLM) CPUs



- **Red Unlock** of Atom Goldmont (GLM) CPUs
- **Extraction** and **reverse engineering** of GLM μ code format




- **Red Unlock** of Atom Goldmont (GLM) CPUs
- **Extraction** and **reverse engineering** of GLM μ code format
- Discovery of undocumented control instructions to access **internal** buffers


**What can you do with access
to microarchitectural buffers?**




...
cpuid
...
...
...
...

XLAT











OP1	OP2	OP3	SEQW
09282eb80236	0008890f8009	092830f80236	0903e480



```
U1a54: 09282eb80236          CMPUJZ_DIRECT_NOTTAKEN(tmp6, 0x2, U0e2e)
U1a55: 0008890f8009          tmp8:= ZEROEXT_DSZ32(0x2389)
U1a56: 092830f80236          SYNC-> CMPUJZ_DIRECT_NOTTAKEN(tmp6, 0x3, U0e30)
U1a57: 000000000000          NOP
SEQW:      0903e480          SEQW GOTO U03e4
```




- **Architectural**: rax, rbx, rcx, ..., r8-r15
- **Temporary**: tmp0-tmp15 + flag0-flag15
- **Remappable Temporary**: tmpv0-tmpv3
- **Instruction Dependent**: r64dst, r64src, r64base, r64idx
- **μ code IP**: uip0, uip1




- Physical Memory (virtually or physically addressable)
- URAM
- Staging Buffer
- Constants ROM



```
1 |
2 | void rc4_decrypt(ulong tmp0_i,ulong tmp1_j,byte *ucode_patch_tmp5,int len_tmp6,byte *S_tmp7,
3 |                 long callback_tmp8)
4 |
5 | {
6 |     byte bVar1;
7 |     byte bVar2;
8 |
9 |     do {
10 |         tmp0_i = (ulong)(byte)((char)tmp0_i + 1);
11 |         bVar1 = S_tmp7[tmp0_i];
12 |         tmp1_j = (ulong)(byte)(bVar1 + (char)tmp1_j);
13 |             /* swap S[i] and S[j] */
14 |         bVar2 = S_tmp7[tmp1_j];
15 |         S_tmp7[tmp0_i] = bVar2;
16 |         S_tmp7[tmp1_j] = bVar1;
17 |         *ucode_patch_tmp5 = S_tmp7[(byte)(bVar2 + bVar1)] ^ *ucode_patch_tmp5;
18 |         ucode_patch_tmp5 = ucode_patch_tmp5 + 1;
19 |         len_tmp6 += -1;
20 |     } while (len_tmp6 != 0);
21 |     (*(code *) (callback_tmp8 * 0x10))();
22 |     return;
23 | }
24 |
```




CPU controls its internal units through the CRBUS




CPU controls its internal units through the CRBUS

- MSRs → CRBUS addr




CPU controls its internal units through the CRBUS

- MSRs → CRBUS addr
- **Control** and **Status** registers



CPU controls its internal units through the CRBUS

- MSRs → CRBUS addr
- **Control** and **Status** registers
- **SMM** configuration



CPU controls its internal units through the CRBUS

- MSRs → CRBUS addr
- **Control** and **Status** registers
- **SMM** configuration
- **Post Silicon Validation** features (LDAT)



- The μ code Sequencer manages the access to μ code ROM and RAM



- The μ code Sequencer manages the access to μ code ROM and RAM
- The LDAT has access to the μ code Sequencer



- The μ code Sequencer manages the access to μ code ROM and RAM
- The LDAT has access to the μ code Sequencer
- We can access the LDAT through the CRBUS



- The μ code Sequencer manages the access to μ code ROM and RAM
- The LDAT has access to the μ code Sequencer
- We can access the LDAT through the CRBUS
- If we can access the CRBUS we can control μ code!



Mark Ermolov, Maxim Goryachy & Dmitry Sklyarov discovered the existence of two secret instructions that can access (RW):



- System agent
- URAM
- Staging buffer
- I/O ports
- Power supply unit



Mark Ermolov, Maxim Goryachy & Dmitry Sklyarov discovered the existence of two secret instructions that can access (RW):



- System agent
- URAM
- Staging buffer
- I/O ports
- Power supply unit
- **CRBUS**



```
def CRBUS_WRITE(ADDR, VAL):  
    udbgwr(  
        rax: ADDR,  
        rbx|rdx: VAL,  
        rcx: 0,  
    )
```



Reverse engineer patterns of how the CPU itself accesses the CRBUS

```
//Decompile of: U2782 - part of ucode update routine  
write_8 (crbus_06a0, (ucode_address - 0x7c00));  
MSLOOPCTR = (*(ushort *)((long)ucode_update_ptr + 3) - 1);  
syncmark();  
if ((in_ucode_ustate & 8) != 0) {  
    syncfull();  
    write_8 (crbus_06a1, 0x30400);  
    ucode_ptr = (ulong *)((long)ucode_update_ptr + 5);  
    do {  
        ucode_qword = *ucode_ptr;  
        ucode_ptr = ucode_ptr + 1;  
        write_8 (crbus_06a4, ucode_qword);  
        write_8 (crbus_06a5, ucode_qword >> 0x20);  
        syncwait();  
        MSLOOPCTR --;  
    } while (-1 < MSLOOPCTR);
```



```
def ucode_sequencer_write(SELECTOR, ADDR, VAL):  
    CRBUS[0x6a1] = 0x30000 | (SELECTOR << 8)  
    CRBUS[0x6a0] = ADDR  
    CRBUS[0x6a4] = VAL & 0xffffffff  
    CRBUS[0x6a5] = VAL >> 32  
    CRBUS[0x6a1] = 0
```

with SELECTOR:


- 2 -> SEQW PATCH RAM
- 3 -> MATCH & PATCH
- 4 -> UCODE PATCH RAM




Redirects execution from μ code ROM to μ code RAM to execute patches.


```
patch_off = (patch_addr - 0x7c00) / 2;
```

```
entry:
```






Leveraging `udbgrd/wr` we can patch μ code via software




Leveraging `udbgrd/wr` we can patch μ code via software

- Completely `observe` CPU behavior



Leveraging `udbgrd/wr` we can patch μ code via software

- Completely **observe** CPU behavior
- Completely **control** CPU behavior



Leveraging `udbgrd/wr` we can patch μ code via software

- Completely **observe** CPU behavior
- Completely **control** CPU behavior
- All within a **BIOS** or **kernel** module



Patch μcode



Patch μcode



Hook μcode




Patch μcode




Hook μcode



Trace μcode




We can change the CPU's behavior.



We can change the CPU's behavior.

- **Change** microcoded instructions



We can change the CPU's behavior.

- **Change** microcoded instructions
- **Add** functionalities to the CPU



```
.patch 0x0428 # RDRAND ENTRY POINT
.org 0x7c00
rax:= ZEROEXT_DSZ64(0x6f57206f6c6c6548) # 'Hello Wo'
rbx:= ZEROEXT_DSZ64(0x21646c72) # 'rld!\x00'
UEND
```




```
.patch 0x0428 # RDRAND ENTRY POINT
.org 0x7c00
rax:= ZEROEXT_DSZ64(0x6f57206f6c6c6548) # 'Hello Wo'
rbx:= ZEROEXT_DSZ64(0x21646c72) # 'rld!\x00'
UEND
```

1. Assemble µcode
2. Write µcode at 0x7c00
3. Setup Match & Patch: 0x0428 → 0x7c00
4. rdrand → "Hello World!"

fs0:\EFI> █



rdrand returns random data, what if we make it return SMM memory?

```
.patch 0x0428 # RDRAND ENTRY POINT
.org 0x7c00
tmp1:= MOVEFROMCREG_DSZ64(CR_SMRR_MASK)
tmp2:= ZEROEXT_DSZ64(0x0)
MOVETOCREG_DSZ64(tmp2, CR_SMRR_MASK) # DISABLE SMM MEMORY RANGE

rax:= LDPPHYS_DSZ64(0x7b000000) # SMROM ADDR

MOVETOCREG_DSZ64(tmp1, CR_SMRR_MASK)
UEND
```



rdrand returns random data, what if we make it return SMM memory?

```
.patch 0x0428 # RDRAND ENTRY POINT
.org 0x7c00
tmp1:= MOVEFROMCREG_DSZ64(CR_SMRR_MASK)
tmp2:= ZEROEXT_DSZ64(0x0)
MOVETOCREG_DSZ64(tmp2, CR_SMRR_MASK) # DISABLE SMM MEMORY RANGE

rax:= LDPPHYS_DSZ64(0x7b000000) # SMROM ADDR

MOVETOCREG_DSZ64(tmp1, CR_SMRR_MASK)
UEND
```



rdrand returns random data, what if we make it return SMM memory?

```
.patch 0x0428 # RDRAND ENTRY POINT
.org 0x7c00
tmp1:= MOVEFROMCREG_DSZ64(CR_SMRR_MASK)
tmp2:= ZEROEXT_DSZ64(0x0)
MOVETOCREG_DSZ64(tmp2, CR_SMRR_MASK) # DISABLE SMM MEMORY RANGE

rax:= LDPPHYS_DSZ64(0x7b000000) # SMROM ADDR

MOVETOCREG_DSZ64(tmp1, CR_SMRR_MASK)
UEND
```



rdrand returns random data, what if we make it return SMM memory?

```
.patch 0x0428 # RDRAND ENTRY POINT
.org 0x7c00
tmp1:= MOVEFROMCREG_DSZ64(CR_SMRR_MASK)
tmp2:= ZEROEXT_DSZ64(0x0)
MOVETOCREG_DSZ64(tmp2, CR_SMRR_MASK) # DISABLE SMM MEMORY RANGE

rax:= LDPPHYS_DSZ64(0x7b000000) # SMROM ADDR

MOVETOCREG_DSZ64(tmp1, CR_SMRR_MASK)
UEND
```

fs0:\EFI> █

x86 PAC



ARM mitigation:

protect pointers from corruption using a cryptographic signature



ARM mitigation:

protect pointers from corruption using a cryptographic signature

- `pacia ptr, ctx` → **sign** pointer `ptr` using salt `ctx`
`0x1337` → `0xaced000000001337`



ARM mitigation:

protect pointers from corruption using a cryptographic signature


- pacia ptr, ctx → **sign** pointer ptr using salt ctx
0x1337 → 0xaced000000001337
- autia ptr, ctx → **authenticate** pointer ptr using salt ctx
0xaced000000001337 → 0x1337
0xaced0000000013ff → 0xdead0000000013ff



ARM x86 mitigation:


protect pointers from corruption using a cryptographic signature

- `pacia verw (ptr, ctx, 1) → sign pointer ptr using salt ctx`
`0x1337 → 0xaced000000001337`
- `autia verr (ptr, ctx, 1) → authenticate pointer ptr using salt ctx`
`0xaced000000001337 → 0x1337`
`0xaced0000000013ff → 0xdead0000000013ff`




Bruteforce PAC using speculative execution

```
if (condition):  
    auth_ptr <- aut(ptr) // speculatively executed  
    load(auth_ptr)       // if correct -> address loaded  
                        // if wrong -> address not loaded
```



Bruteforce **x86** PAC using speculative execution


```
if (condition):  
    auth_ptr <- aut(ptr) // speculatively executed  
    load(auth_ptr)      // never loaded
```



Bruteforce **x86** PAC using speculative execution

```
if (condition):  
    auth_ptr <- aut(ptr) // speculatively executed  
    load(auth_ptr)      // never loaded
```

our 54 μ ops fill the speculative window!



Attack a **weaker** version of **x86** PAC (1/2 round SipHash)

```
if (condition):  
    auth_ptr <- aut(ptr) // speculatively executed  
    load(auth_ptr)      // if correct -> address loaded  
                        // if wrong -> address not loaded
```

fs0:\EFI> █



- Speculation Barrier



- ~~Speculation Barrier~~ → high overhead + partial protection

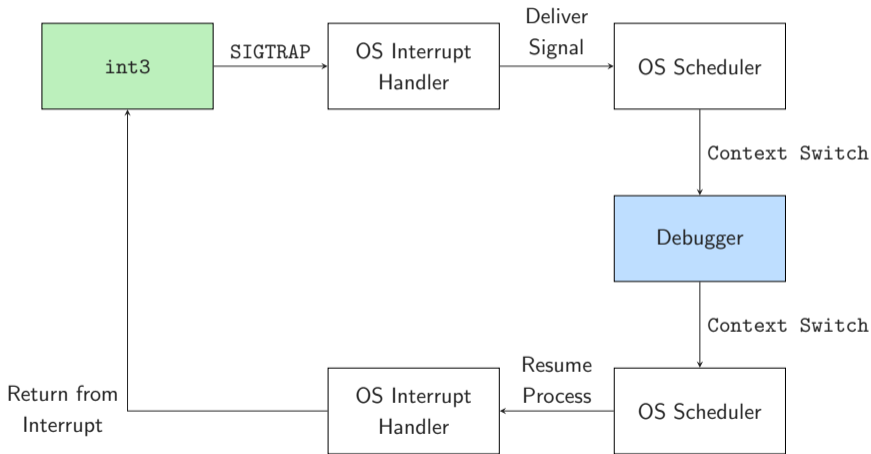


- ~~Speculation Barrier~~ → high overhead + partial protection
- Make aut **fault** when invalid PAC



- ~~Speculation Barrier~~ → high overhead + partial protection
- Make aut **fault** when invalid PAC
- **Always** remove signature from pointer
→ pointer access always valid on speculative paths

μsoftware breakpoints





Patch `int3` to implement breakpointing logic directly in μcode

- Avoid **interrupt** overhead ($\sim 50\times$ faster)
- Avoid **context switch** overhead ($\sim 1000\times$ faster)
- \rightarrow 10 cycles for a breakpoint



Implement fuzzing **coverage collection** in μcode

```
.patch 0xa78 # int3 entry point
let [cov_map] := tmp1
let [rip] := tmp0

# load address of coverage map from staging buffer
[cov_map] := LDSTGBUF_DSZ64_ASZ16_SC1(0xba00)
# get instruction pointer low bits
[rip] := ZEROEXT_DSZ64(IMM_MACRO_ALIAS_RIP) !m0
[rip] := AND_DSZ64(0xffff, [rip])
# set coverage for basic block
STADPPHYS_DSZ8_ASZ64_SC1([cov_map], [rip], 0x01)
```

Constant-Time division



The `div` instruction latency depends on the input data



The `div` instruction latency depends on the input data

- Side channel attacks can **infer** input data



The `div` instruction latency depends on the input data

- Side channel attacks can **infer** input data
- $2/1 \rightarrow 22$ cycles



The `div` instruction latency depends on the input data

- Side channel attacks can **infer** input data
- $2/1 \rightarrow 22$ cycles
- $0x113371337/1 \rightarrow 40$ cycles



Software `div` implementation mitigates the issue



Software `div` implementation mitigates the issue

- **High overhead** $\rightarrow \sim 700$ cycles



Software `div` implementation mitigates the issue

- **High overhead** → ~ 700 cycles
- Need compilation pass or binary patching



```
.org 0x7c00
.patch 0x6c8 # div entry point
.entry 0

let [dividend] := rax;      let [temp1] := tmp3
let [divisor] := rcx;      let [temp2] := tmp4
let [size] := 0x3f;       let [temp3] := tmp5
let [quotient] := tmp0;    let [temp4] := tmp7
let [temp] := tmp1;       let [temp5] := tmp8
let [i] := tmp2;         let [comp] := tmp6
[temp] := ZEROEXT_DS264(0x0); [i] := ZEROEXT_DS264([size])
[quotient] := ZEROEXT_DS264(0x0)

<loop>
# if (i < 0) goto end;
UJMPCC_DIRECT_NOTTAKEN_CONDB([i], <end>)

# temp = (temp << luLL) | ((dividend >> i) & 1);
[temp1]:= SHL_DS264([temp], 0x1)
[temp2]:= SHR_DS264([dividend], [i])
[temp2]:= AND_DS264([temp2], 0x1)
[temp] := OR_DS264([temp1], [temp2])

# comp = (temp >= divisor);
[comp] := SUB_DS264([divisor], [temp])

# temp -= comp? divisor : 0;
[temp3]:= SELECTCC_DS264_CONDB([comp], [divisor])
[temp] := SUB_DS264([temp3], [temp])

# quotient |= comp ? luLL << i : 0;
[temp4]:= SHL_DS264(0x1, [i])
[temp5]:= SELECTCC_DS264_CONDB([comp], [temp4])
[quotient] := OR_DS264([quotient], [temp5])

# i--; goto loop
[i] := SUB_DS264(0x1, [i]) SEQW GOTO <loop>

<end>
# return quotient, ignore the remainder for simplicity
rax := ZEROEXT_DS264([quotient])
rdx := ZEROEXT_DS264(0x0)
```

We can patch the div instruction to be constant time



```
.org 0x7c00
.patch 0x6c8 # div entry point
.entry 0

let [dividend] := rax;      let [temp1] := tmp3
let [divisor] := rcx;      let [temp2] := tmp4
let [size] := 0x3f;       let [temp3] := tmp5
let [quotient] := tmp0;    let [temp4] := tmp7
let [temp] := tmp1;       let [temp5] := tmp8
let [i] := tmp2;         let [comp] := tmp6
[temp] := ZEROEXT_DS264(0x0); [i] := ZEROEXT_DS264([size])
[quotient] := ZEROEXT_DS264(0x0)

<loop>
# if (i < 0) goto end;
UJMPCC_DIRECT_NOTTAKEN_CONDB([i], <end>)

# temp = (temp << luLL) | ((dividend >> i) & 1);
[temp1]:= SHL_DS264([temp], 0x1)
[temp2]:= SHR_DS264([dividend], [i])
[temp2]:= AND_DS264([temp2], 0x1)
[temp] := OR_DS264([temp1], [temp2])

# comp = (temp >= divisor);
[comp] := SUB_DS264([divisor], [temp])

# temp -= comp? divisor : 0;
[temp3]:= SELECTCC_DS264_CONDB([comp], [divisor])
[temp] := SUB_DS264([temp3], [temp])

# quotient |= comp ? luLL << i : 0;
[temp4]:= SHL_DS264(0x1, [i])
[temp5]:= SELECTCC_DS264_CONDB([comp], [temp4])
[quotient] := OR_DS264([quotient], [temp5])

# i--; goto loop
[i] := SUB_DS264(0x1, [i]) SEQW GOTO <loop>

<end>
# return quotient, ignore the remainder for simplicity
rax := ZEROEXT_DS264([quotient])
rdx := ZEROEXT_DS264(0x0)
```

We can patch the div instruction to be constant time

- **Reduced overhead** → ~ 400 cycles



```
.org 0x7c00
.patch 0x6c8 # div entry point
.entry 0

let [dividend] := rax;      let [temp1] := tmp3
let [divisor]  := rcx;      let [temp2] := tmp4
let [size]     := 0x3f;     let [temp3] := tmp5
let [quotient] := tmp0;     let [temp4] := tmp7
let [temp]     := tmp1;     let [temp5] := tmp8
let [i]        := tmp2;     let [comp]  := tmp6
[temp] := ZEROEXT_DS264(0x0); [i] := ZEROEXT_DS264([size])
[quotient] := ZEROEXT_DS264(0x0)

<loop>
# if (i < 0) goto end;
UJMPCC_DIRECT_NOTTAKEN_CONDB([i], <end>)

# temp = (temp << luLL) | ((dividend >> i) & 1);
[temp1]:= SHL_DS264([temp], 0x1)
[temp2]:= SHR_DS264([dividend], [i])
[temp2]:= AND_DS264([temp2], 0x1)
[temp] := OR_DS264([temp1], [temp2])

# comp = (temp >= divisor);
[comp] := SUB_DS264([divisor], [temp])

# temp -= comp? divisor : 0;
[temp3]:= SELECTCC_DS264_CONDB([comp], [divisor])
[temp] := SUB_DS264([temp3], [temp])

# quotient |= comp ? luLL << i : 0;
[temp4]:= SHL_DS264(0x1, [i])
[temp5]:= SELECTCC_DS264_CONDB([comp], [temp4])
[quotient] := OR_DS264([quotient], [temp5])

# i--; goto loop
[i] := SUB_DS264(0x1, [i]) SEQW GOTO <loop>

<end>
# return quotient, ignore the remainder for simplicity
rax := ZEROEXT_DS264([quotient])
rdx := ZEROEXT_DS264(0x0)
```

We can patch the div instruction to be constant time

- **Reduced overhead** → ~ 400 cycles
- **Transparent** to the running program



Install μcode hooks to observe events.

- Setup Match & Patch to execute custom μcode at certain events
- Resume execution



We can make the CPU to react to certain μ code events, e.g., `verw` executed

```
.patch 0xXXXX # INSTRUCTION ENTRY POINT
.org 0x7da0

tmp0:= ZEROEXT_DSZ64(<counter_address>)
tmp1:= LDPPHYSTICKLE_DSZ64_ASZ64_SC1(tmp0)
tmp1:= ADD_DSZ64(tmp1, 0x1) # INCREMENT COUNTER
STADPPHYSTICKLE_DSZ64_ASZ64_SC1(tmp0, tmp1)

UJMP(0xXXXX + 1) # JUMP TO NEXT UOP
```



We can make the CPU to react to certain μ code events, e.g., `verw` executed


```
.patch 0xXXXX # INSTRUCTION ENTRY POINT
.org 0x7da0

tmp0:= ZEROEXT_DSZ64(<counter_address>)
tmp1:= LDPPHYSTICKLE_DSZ64_ASZ64_SC1(tmp0)
tmp1:= ADD_DSZ64(tmp1, 0x1) # INCREMENT COUNTER
STADPPHYSTICKLE_DSZ64_ASZ64_SC1(tmp0, tmp1)

UJMP(0xXXXX + 1) # JUMP TO NEXT UOP
```





Trace μcode execution leveraging hooks.






Trace μcode execution leveraging hooks.






Trace μcode execution leveraging hooks.






Trace μcode execution leveraging hooks.






Trace μcode execution leveraging hooks.





Trace μcode execution leveraging hooks.




hook:

1. dump timestamp
2. disable hook
3. continue



Trace μcode execution leveraging hooks.




hook:

1. dump timestamp
2. disable hook
3. continue



Trace μcode execution leveraging hooks.




hook:

1. dump timestamp
2. disable hook
3. continue



Trace μcode execution leveraging hooks.




hook:

1. dump timestamp
2. disable hook
3. continue



Trace μcode execution leveraging hooks.




hook:

1. dump timestamp
2. disable hook
3. continue



Trace μcode execution leveraging hooks.

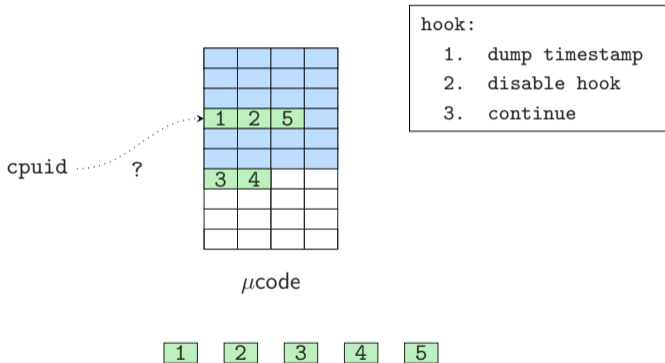


hook:

1. dump timestamp
2. disable hook
3. continue



Trace μcode execution leveraging hooks.





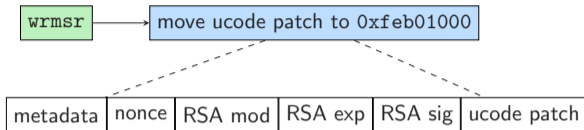
μ code update algorithm has always been kept **secret** by Intel
Let's trace the execution of a μ code update!


- Trigger a μ code update
- **Trace** if a microinstruction is executed
- **Repeat** for all the possible μ code instructions
- Restore order

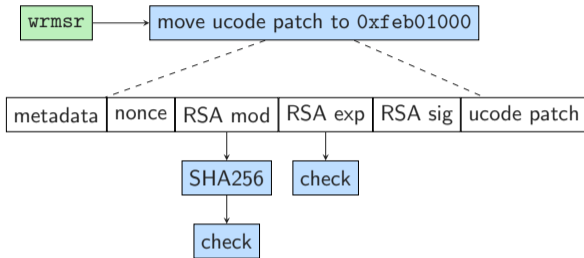


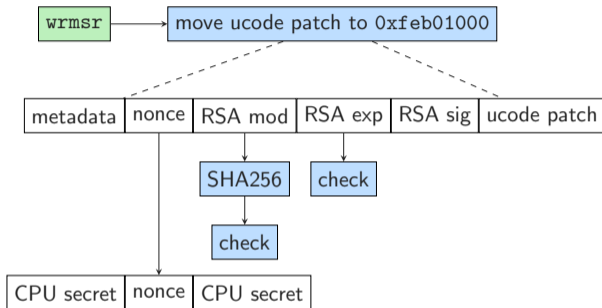
wrmsr

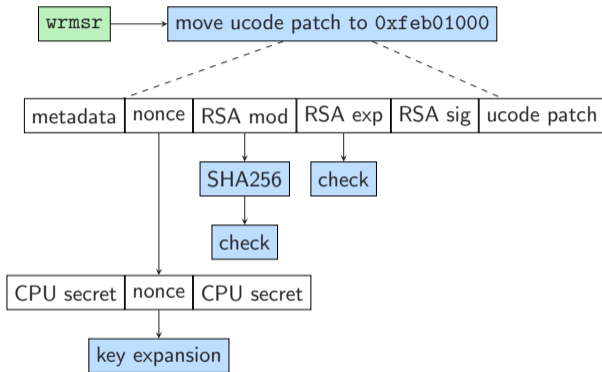


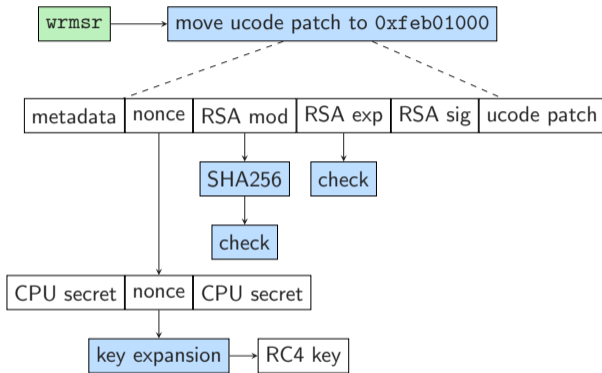


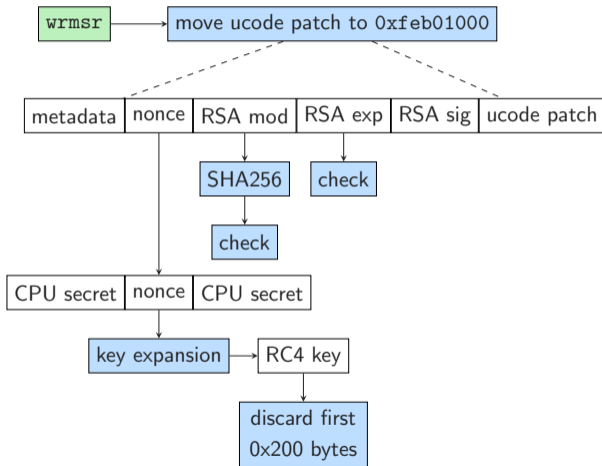







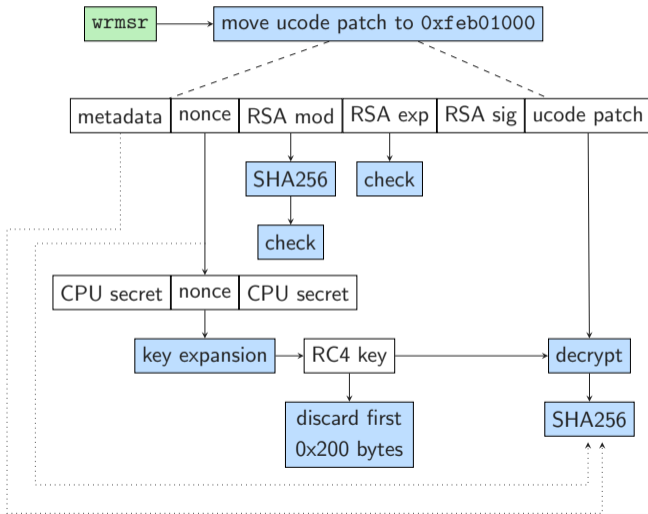




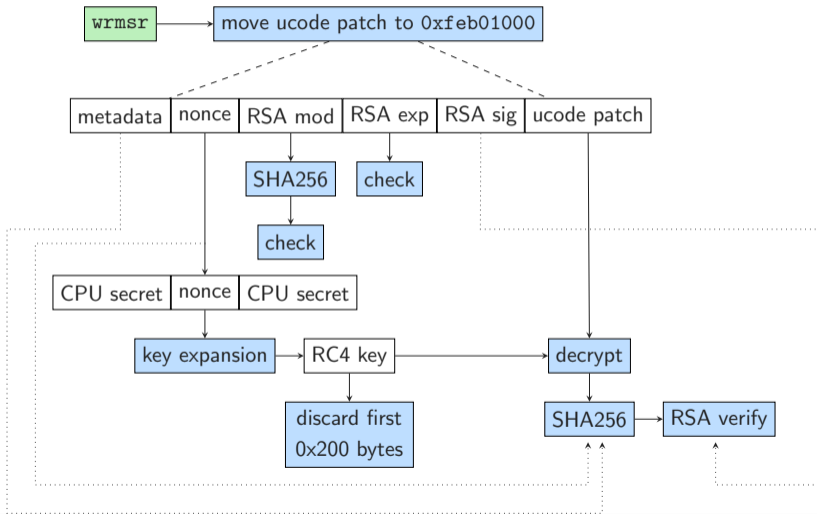




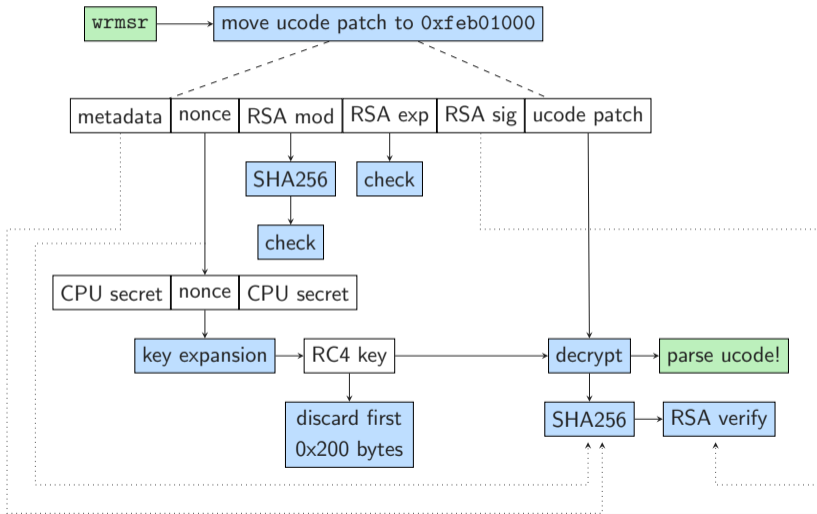
GLM μ code update algorithm



GLM μ code update algorithm



GLM μ code update algorithm





Used to decrypt updates

- 16B **per-CPU-model** key, 32B **per-update** nonce → expanded to 256B
- Discard first 512 bytes to reduce bias
- Attacks on RC4 use $> 1_000_000$ ciphertexts → only 453 public updates :(



Used to check update signature

- 2048 bits, PKCS#1 v1.5 padding
- RSA modulus and exponent **hardcoded**
- Signature checks:
 - security revision
 - cpuid values
 - nonce
 - decrypted update



The temporary physical address where μ code is decrypted.
→ Used as secure memory



The temporary physical address where μ code is decrypted.

→ Used as secure memory

```
> sudo cat /proc/iomem | grep feb00000
```

```
:(
```



The temporary physical address where μ code is decrypted.

→ Used as secure memory

```
> sudo cat /proc/iomem | grep feb00000
```

```
:(
```

```
> read_physical_address 0xfeb01000
```

```
00000000: ffff ffff ffff ffff ffff ffff ffff ffff
```

```
00000010: ffff ffff ffff ffff ffff ffff ffff ffff
```

```
00000020: ffff ffff ffff ffff ffff ffff ffff ffff
```

```
00000030: ffff ffff ffff ffff ffff ffff ffff ffff
```



- **Dynamically** enabled by the CPU



- Dynamically enabled by the CPU
- Access time: about 20 cycles



- Dynamically enabled by the CPU
- Access time: about 20 cycles
- Content not shared between cores



- **Dynamically** enabled by the CPU
- Access time: about **20 cycles**
- Content **not shared** between cores
- Can fit 64-256Kb of valid data



- **Dynamically** enabled by the CPU
- Access time: about **20 cycles**
- Content **not shared** between cores
- Can fit 64-256Kb of valid data
- **Replacement policy** on the content?!



- **Dynamically** enabled by the CPU
- Access time: about **20 cycles**
- Content **not shared** between cores
- Can fit 64-256Kb of valid data
- **Replacement policy** on the content?!
- It's a special CPU view on the **L2 cache!**



Can we corrupt the μcode after decryption but before being applied

- Each core has a **private** Secure Memory area



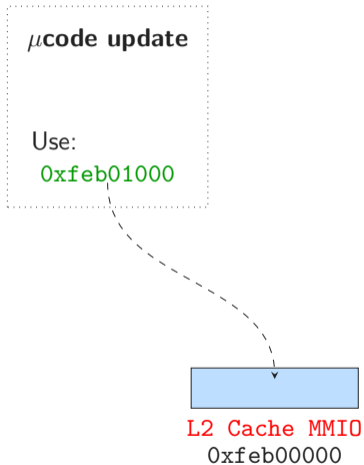
Can we corrupt the μcode after decryption but before being applied


- Each core has a **private** Secure Memory area
- `wrmsr` → Hyperthreading **disabled** during update

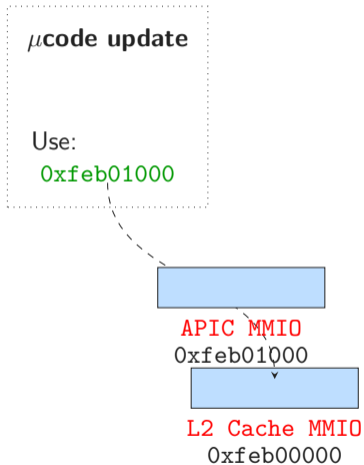


Can we corrupt the μcode after decryption but before being applied

- Each core has a **private** Secure Memory area
- `wrmsr` → Hyperthreading **disabled** during update
- Secure Memory only enabled during updates → read `0xff`









Can we **transiently** leak μcode updates?

- Fallout/MDS → No **hyperthreading**



Can we **transiently** leak μcode updates?

- Fallout/MDS → No **hyperthreading**
- L1TF → Only leaks of **internal** buffers (L1)



Can we **transiently** leak μcode updates?

- Fallout/MDS → No **hyperthreading**
- L1TF → Only leaks of **internal** buffers (L1)
- Inverse \AE PIC Leak → Internal buffers **flushed**



Can we induce **faults** and skip checks?

- We found no real fault-injection protection



Can we induce **faults** and skip checks?

- We found no real fault-injection protection
- **Signature check** operations seems the most profitable



Can we induce **faults** and skip checks?


- We found no real fault-injection protection
- **Signature check** operations seems the most profitable
- The update is not persistent → repeat attack at each boot



```
00000000: 0102 007c 3900 0a00 3f88 4bed c000 080c ...|9...?.K.....
00000010: 0b01 4780 0000 0a00 3f88 4fad 0003 0a00 ..G.....?.0.....
00000020: 2f20 4b2d 8002 080c 0322 4740 a903 0a00 / K-....."G@....
00000030: 2f20 4f6d 1902 0002 0353 6380 c000 3002 / Om.....Sc...0.
00000040: b8a6 6be8 0000 0002 0320 63c0 0003 f003 ..k..... c.....
00000050: f8a6 6b28 c000 0800 03c0 0bed 0000 0b10 ..k(.....
00000060: 7f00 0800 8001 3110 0300 a140 c000 310c .....1....@..1.
00000070: 0300 0700 0000 4012 0b30 6210 0003 4b1c .....@..0b...K.
00000080: 7f00 0440 c000 3112 0310 2400 0000 310c ...@..1...$...1.
00000090: 0300 01c0 0003 0800 03c0 0fad 0002 00d2 .....
```



A μ code update is bytecode: the CPU interprets commands from the μ code update





- Create a **parser** for μcode updates



- Create a **parser** for μcode updates
- Automatically collect existing μcode (s) for GLM



- Create a **parser** for μcode updates
- Automatically collect existing μcode (s) for GLM
- **Decrypt** all GLM updates





- Create a **parser** for μcode updates
- Automatically collect existing μcode (s) for GLM
- **Decrypt** all GLM updates

github.com/pietroborrello/CustomProcessingUnit/ucode_collection




- Deepen understanding of modern CPUs with `μcode` access






- Deepen understanding of modern CPUs with μ code access
- Develop a static and dynamic analysis framework for μ code:



- Deepen understanding of modern CPUs with μ code access
- Develop a static and dynamic analysis framework for μ code:
 - μ code decompiler
 - μ code assembler
 - μ code patcher
 - μ code tracer



- Deepen understanding of modern CPUs with **μcode** access
- Develop a static and dynamic analysis framework for **μcode**:
 - **μcode** decompiler
 - **μcode** assembler
 - **μcode** patcher
 - **μcode** tracer
- Let's **control** our CPUs!

`github.com/pietroborrello/CustomProcessingUnit`