

Finding and Exploiting CPU Features using MSR Templating

IEEE Symposium on Security and Privacy 2022

Andreas Kogler

Graz University of Technology

Moritz Lipp

Amazon Web Services

Daniel Weber

CISPA Helmholtz Center for Information Security

Daniel Gruss

Graz University of Technology

Martin Haubenwallner

Graz University of Technology

Michael Schwarz

CISPA Helmholtz Center for Information Security



- **Motivation**
- **Framework**
 - Detection
 - Classification
 - Extensions
- **Case Studies**



- **Model Specific Registers (MSRs)**

- 2^{32} 64-bit Registers
- Documented
- Undocumented



- **Model Specific Registers (MSRs)**
 - 2^{32} 64-bit Registers
 - Documented
 - Undocumented
- **Influences** on instructions



- **Model Specific Registers (MSRs)**
 - 2^{32} 64-bit Registers
 - Documented
 - Undocumented
- **Influences** on instructions
- **Security** patches



- **Model Specific Registers (MSRs)**
 - 2^{32} 64-bit Registers
 - Documented
 - Undocumented
- **Influences** on instructions
- **Security** patches
- **Hidden** features (e.g., Domas [1])



MSR Detection

MSR Classification



- Scan **all** MSR addresses
 - rdmsr → GP-Fault?
 - wrmsr → GP-Fault?



- **Scan all** MSR addresses

- `rdmsr` → GP-Fault?

- `wrmsr` → GP-Fault?

✓ Complete MSR list



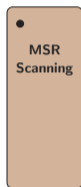
- **Scan all** MSR addresses

- `rdmsr` → GP-Fault?
- `wrmsr` → GP-Fault?

- ✓ Complete MSR list

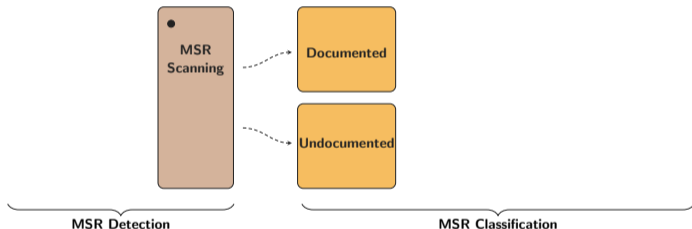
- ✓ *R, W, RW* or *not-present*

The Framework: Documented vs Undocumented

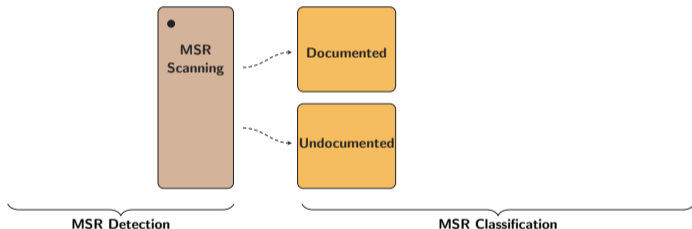


MSR Detection

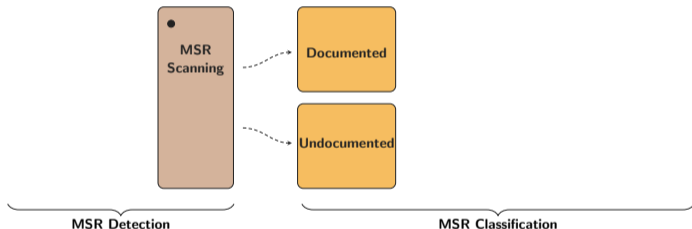
MSR Classification



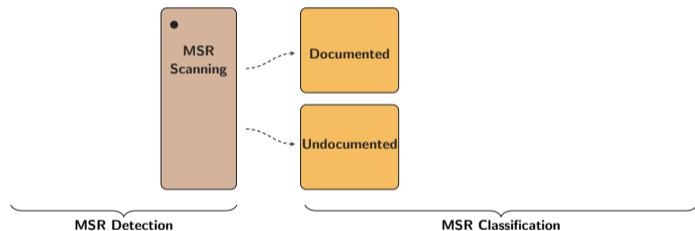
- Parse **official** PDFs
 - AMD's Reference
 - Intel's SDM

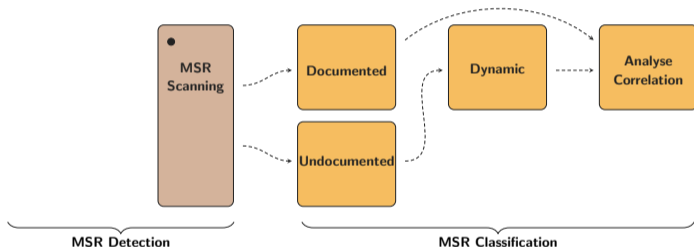


- Parse **official** PDFs
 - AMD's Reference
 - Intel's SDM
- **Extract** table structures
 - Python script

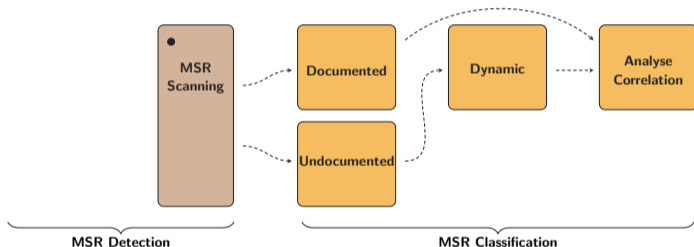


- Parse **official** PDFs
 - AMD's Reference
 - Intel's SDM
- **Extract** table structures
 - Python script
- ✓ Documented MSRs
- ✓ Undocumented MSRs

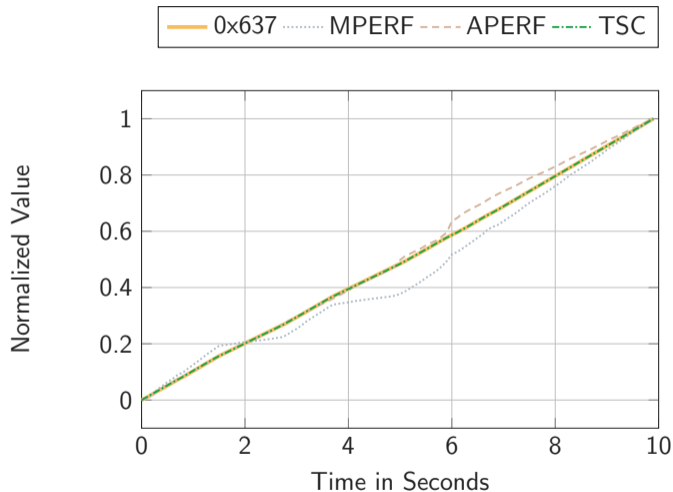




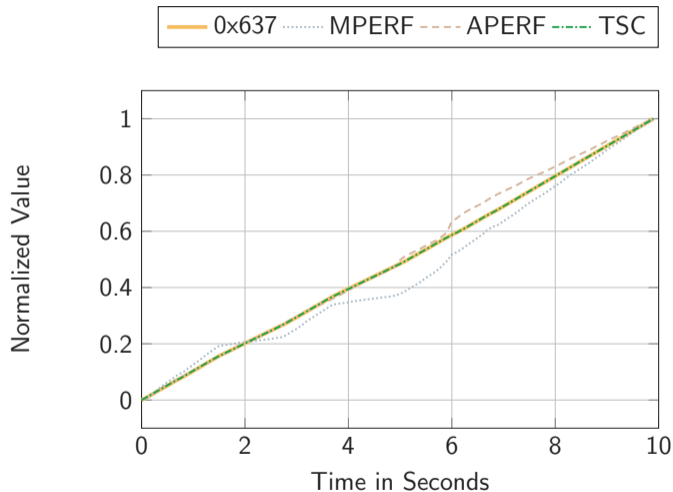
- **Dynamic MSR:**
 - Changing **signals**



- **Dynamic MSR:**
 - Changing **signals**
- **Correlation analysis**
 - Similarity
 - Source

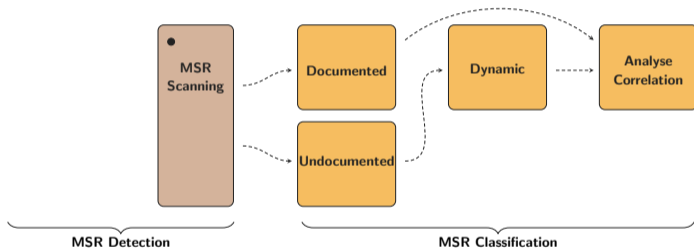


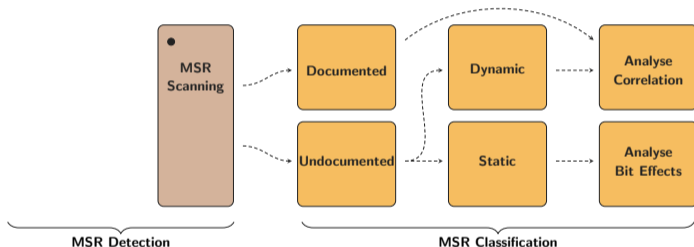
- **Dynamic MSR:**
 - Changing **signals**
- **Correlation** analysis
 - Similarity
 - Source
- **Example:** MSR 0x637



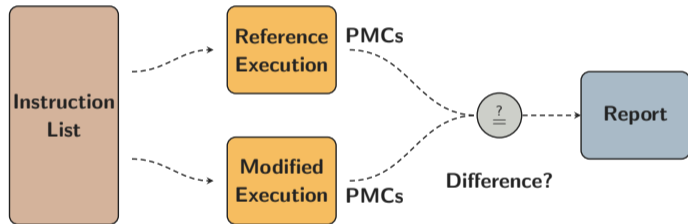
- **Dynamic MSR:**
 - Changing **signals**
- **Correlation** analysis
 - Similarity
 - Source
- **Example:** MSR 0x637
 - ✓ Similar MSRs
 - ✓ Source hints

The Framework: Static Analysis

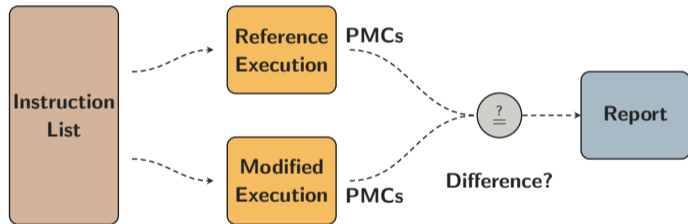




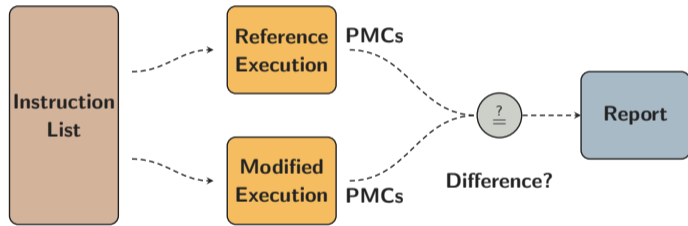
- **Static MSR:**
 - Configuration bits



- **Static MSR:**
 - Configuration **bits**
- **Execute** instruction twice
 - Reference
 - Modified

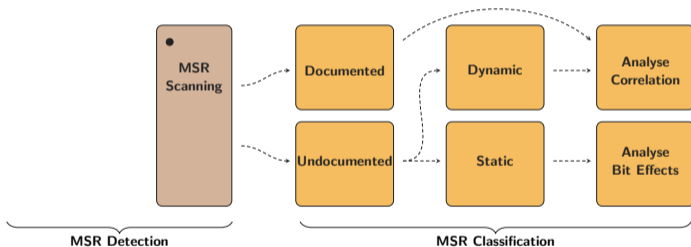


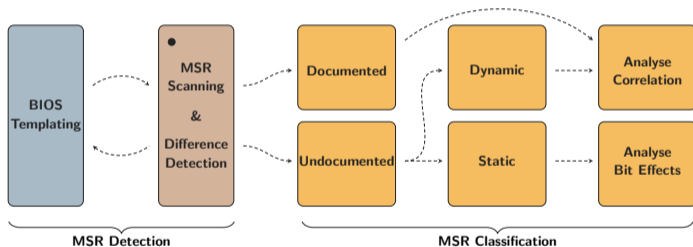
- **Static MSR:**
 - Configuration **bits**
- **Execute** instruction twice
 - Reference
 - Modified
- **Analyze** PMC differences



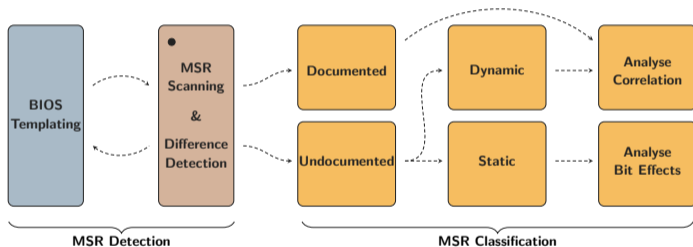
- **Static MSR:**
 - Configuration **bits**
- **Execute** instruction twice
 - Reference
 - Modified
- **Analyze** PMC differences
 - ✓ **Influenced** instructions

The Framework: BIOS Templating

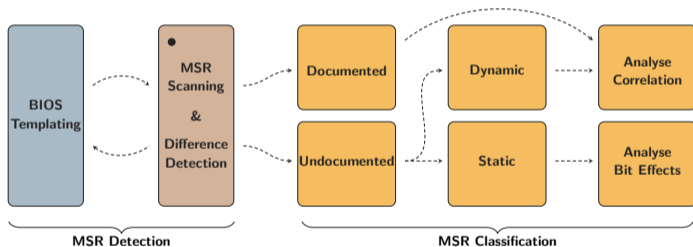




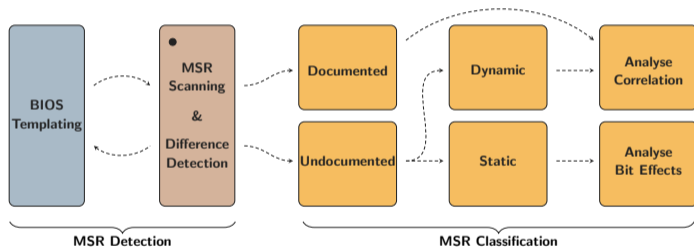
- **Extend search space**



- **Extend** search space
- **Change** BIOS feature

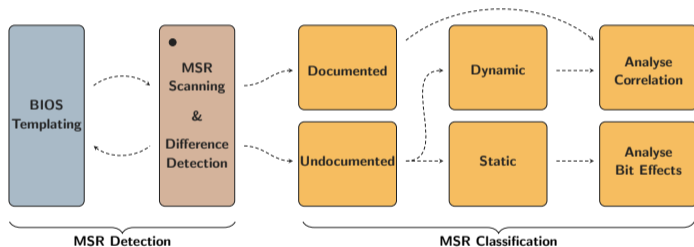


- **Extend** search space
- **Change** BIOS feature
- **Trace** differences

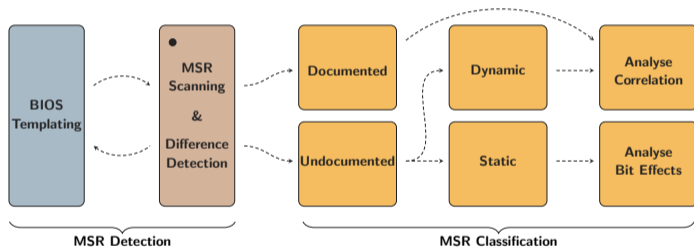


- **Extend** search space
- **Change** BIOS feature
- **Trace** differences
- ✓ **Changed** MSRs

The Framework: Summary

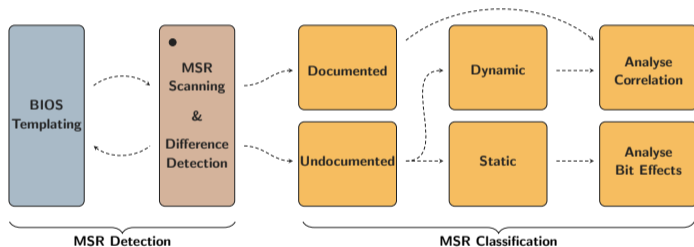


The Framework: Summary

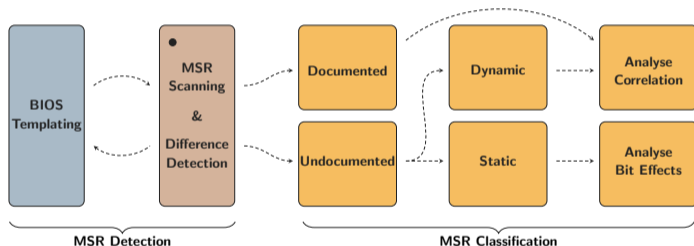


✓ List (*R*, *W*, *RW*, or *NP*)

The Framework: Summary

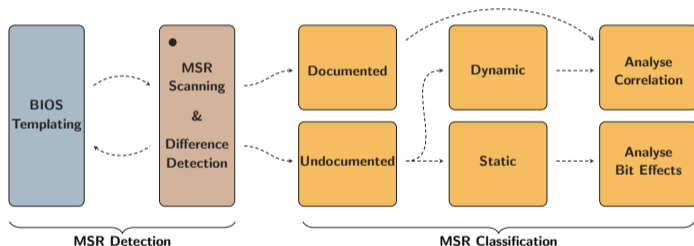


- ✓ List (*R*, *W*, *RW*, or *NP*)
- ✓ Dynamic: **similar** MSRs

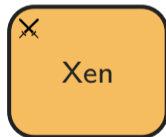


- ✓ List (*R*, *W*, *RW*, or *NP*)
- ✓ Dynamic: **similar** MSRs
- ✓ Static: **influenced** instruction

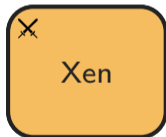
The Framework: Summary



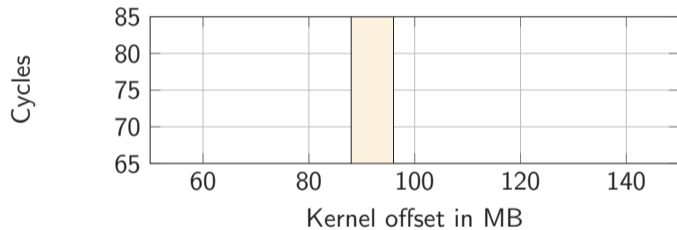
- ✓ List (*R*, *W*, *RW*, or *NP*)
- ✓ Dynamic: **similar** MSRs
- ✓ Static: **influenced** instruction
- ✓ BIOS: **changed** MSRs



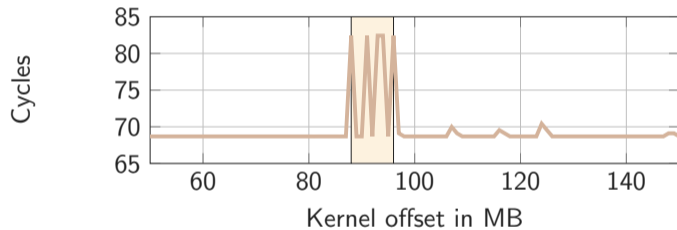
- **Attack** case studies



- **Attack** case studies
- **Defense** case studies



- **Prefetch-based** attacks [2]



- **Prefetch-based attacks** [2]



Instruction	MSR	PMC Effect
PREFETCHNTA	Bit 2	-1 LdDispatch
PREFETCHT0	Bit 3	-1 LdDispatch
PREFETCHT1	Bit 4	-1 LdDispatch
PREFETCHT2	Bit 5	-1 LdDispatch
PREFETCHW	Bit 6	-1 LdDispatch
PREFETCH	Bit 7	-1 LdDispatch

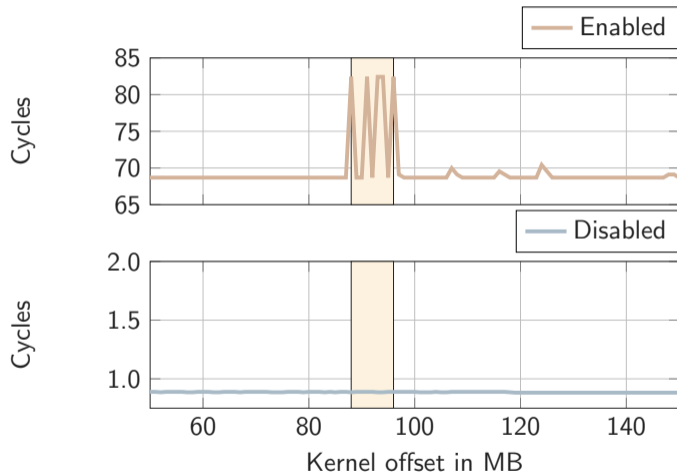
- **Prefetch-based** attacks [2]
- **Search** configuration bits



Instruction	MSR	PMC Effect
PREFETCHNTA	Bit 2	-1 LdDispatch
PREFETCHT0	Bit 3	-1 LdDispatch
PREFETCHT1	Bit 4	-1 LdDispatch
PREFETCHT2	Bit 5	-1 LdDispatch
PREFETCHW	Bit 6	-1 LdDispatch
PREFETCH	Bit 7	-1 LdDispatch

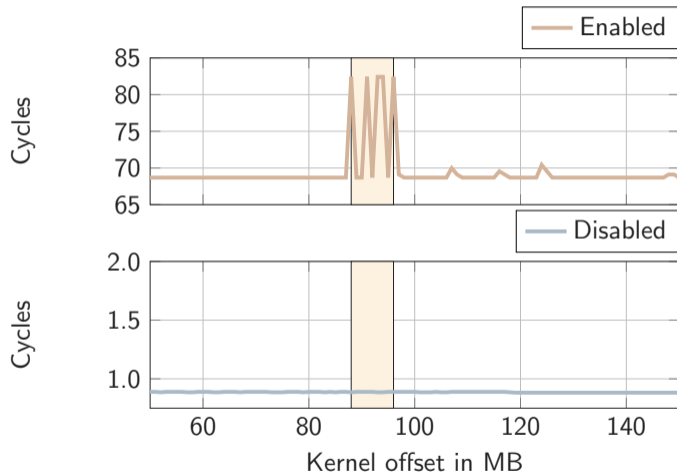
- **Prefetch-based** attacks [2]
- **Search** configuration bits
- **Disable** prefetch*

Case Study: Prefetch



- **Prefetch-based** attacks [2]
- **Search** configuration bits
- **Disable** prefetch*
- ✓ **No prefetch-based** attacks

Case Study: Prefetch



- **Prefetch-based** attacks [2]
- **Search** configuration bits
- **Disable** prefetch*
- ✓ **No** prefetch-based attacks
- ✓ 1% Binaries → 0.04% SPEC



- **Lock bit**



- Lock bit
- Disable at runtime

```
/* ... */
if( mbedtls_aesni_has_support( MBEDTLS_AESNI_AES ) )
    return( mbedtls_aesni_setkey_enc( ctx->rk, key, keybits ) );
/* ... */
switch( ctx->nr ) {
case 10:
    for( i = 0; i < 10; i++, RK += 4 ) {
        RK[4] = RK[0] ^ RCON[i] ^
            ( FSb[ ( RK[3] >> 8 ) & 0xFF ]          ) ^
            ( FSb[ ( RK[3] >> 16 ) & 0xFF ] << 8 ) ^
            ( FSb[ ( RK[3] >> 24 ) & 0xFF ] << 16 ) ^
            ( FSb[ ( RK[3]          ) & 0xFF ] << 24 );

        RK[5] = RK[1] ^ RK[4];
        RK[6] = RK[2] ^ RK[5];
        RK[7] = RK[3] ^ RK[6];
    }
    break;
    /* additional cases for different key lengths */
}
/* ... */
```

- Lock bit
- Disable at runtime
- MbedTLS in SGX

```
/* ... */
if( mbedtls_aesni_has_support( MBEDTLS_AESNI_AES ) )
    return( mbedtls_aesni_setkey_enc( ctx->rk, key, keybits ) );
/* ... */
switch( ctx->nr ) {
    case 10:
        for( i = 0; i < 10; i++, RK += 4 ) {
            RK[4] = RK[0] ^ RCON[i] ^
                ( FSb[ ( RK[3] >> 8 ) & 0xFF ]          ) ^
                ( FSb[ ( RK[3] >> 16 ) & 0xFF ] << 8 ) ^
                ( FSb[ ( RK[3] >> 24 ) & 0xFF ] << 16 ) ^
                ( FSb[ ( RK[3]          ) & 0xFF ] << 24 );

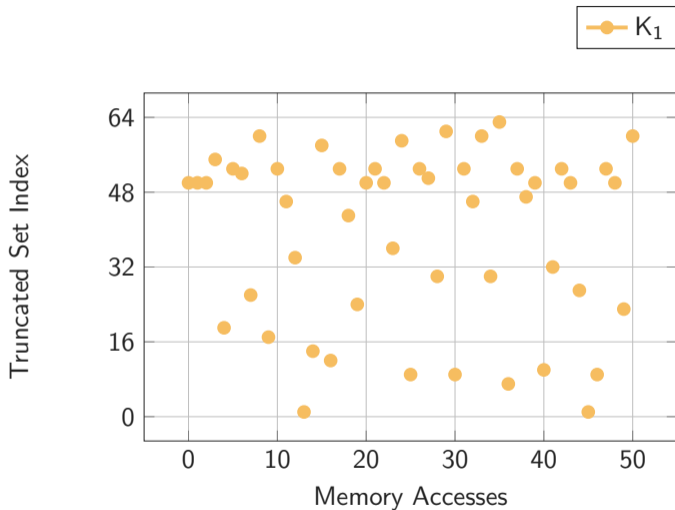
            RK[5] = RK[1] ^ RK[4];
            RK[6] = RK[2] ^ RK[5];
            RK[7] = RK[3] ^ RK[6];
        }
        break;
    /* additional cases for different key lengths */
}
/* ... */
```

- Lock bit
- Disable at runtime
- MbedTLS in SGX
 - Check AES-NI

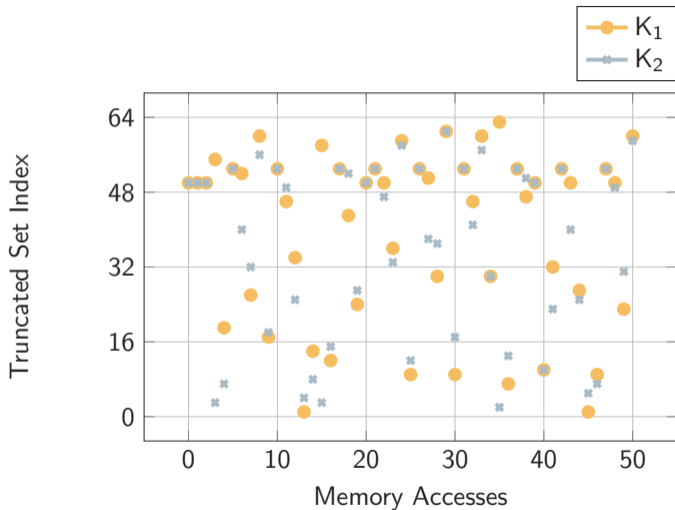
```
/* ... */
if( mbedtls_aesni_has_support( MBEDTLS_AESNI_AES ) )
    return( mbedtls_aesni_setkey_enc( ctx->rk, key, keybits ) );
/* ... */
switch( ctx->nr ) {
    case 10:
        for( i = 0; i < 10; i++, RK += 4 ) {
            RK[4] = RK[0] ^ RCON[i] ^
                ( FSb[ ( RK[3] >> 8 ) & 0xFF ]          ) ^
                ( FSb[ ( RK[3] >> 16 ) & 0xFF ] << 8 ) ^
                ( FSb[ ( RK[3] >> 24 ) & 0xFF ] << 16 ) ^
                ( FSb[ ( RK[3]          ) & 0xFF ] << 24 );

            RK[5] = RK[1] ^ RK[4];
            RK[6] = RK[2] ^ RK[5];
            RK[7] = RK[3] ^ RK[6];
        }
        break;
    /* additional cases for different key lengths */
}
/* ... */
```

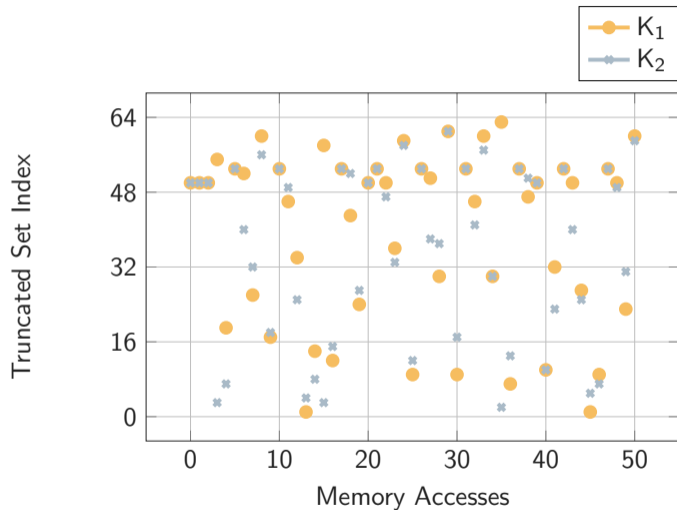
- Lock bit
- Disable at runtime
- MbedTLS in SGX
 - Check AES-NI
 - Fallback T-Tables



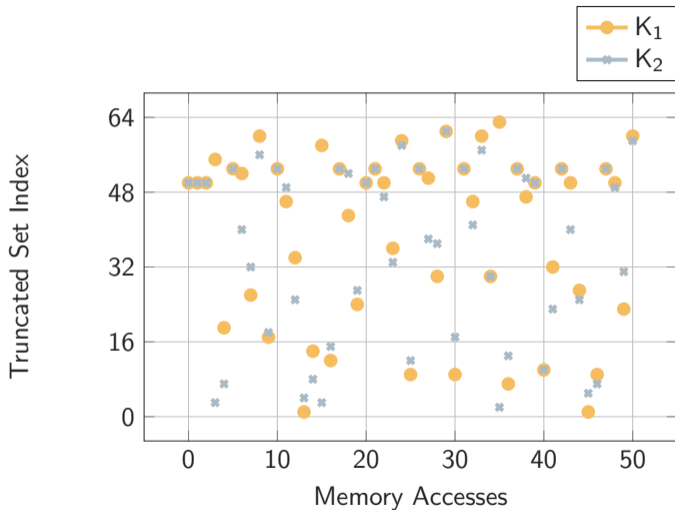
- Lock bit
- Disable at runtime
- MbedTLS in SGX
 - Check AES-NI
 - Fallback T-Tables
 - LLC P+P



- Lock bit
- Disable at runtime
- MbedTLS in SGX
 - Check AES-NI
 - Fallback T-Tables
 - LLC P+P



- Lock bit
- Disable at runtime
- MbedTLS in SGX
 - Check AES-NI
 - Fallback T-Tables
 - LLC P+P
 - Z3 Solver



- Lock bit
 - Disable at runtime
 - MbedTLS in SGX
 - Check AES-NI
 - Fallback T-Tables
 - LLC P+P
 - Z3 Solver
- ✓ Full key



- **CrossTalk** attack [3]



- **CrossTalk** attack [3]
- **Unprivileged** leakage
 - cpuid → 88.9%
 - rdseed → 0.4%



- **CrossTalk** attack [3]
- **Unprivileged** leakage
 - cpuid → 88.9%
 - rdseed → 0.4%
- **Search** configuration bits

- **CrossTalk** attack [3]
- **Unprivileged** leakage
 - cpuid → 88.9%
 - rdseed → 0.4%
- **Search** configuration bits
- **CPUID** trap

- **CrossTalk** attack [3]
- **Unprivileged** leakage
 - `cpuid` → ~~88.9%~~
 - `rdseed` → 0.4%
- **Search** configuration bits
- **CPUID** trap
- ✓ **Reduced** by **211.4** times

Hardware:

Xen HV:

Guest:

`rdmsr`

- **Hypervisor** handles MSR

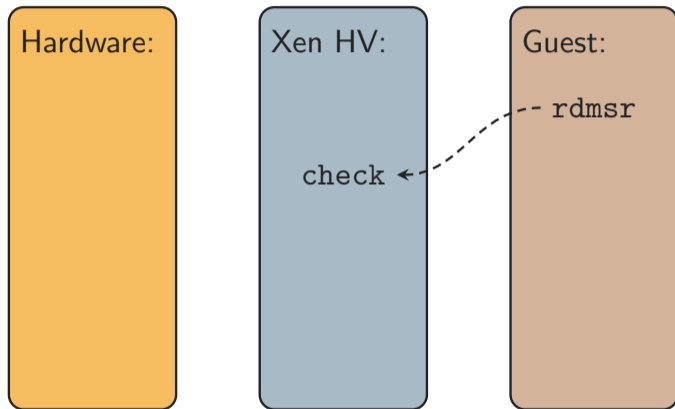
Hardware:

Xen HV:

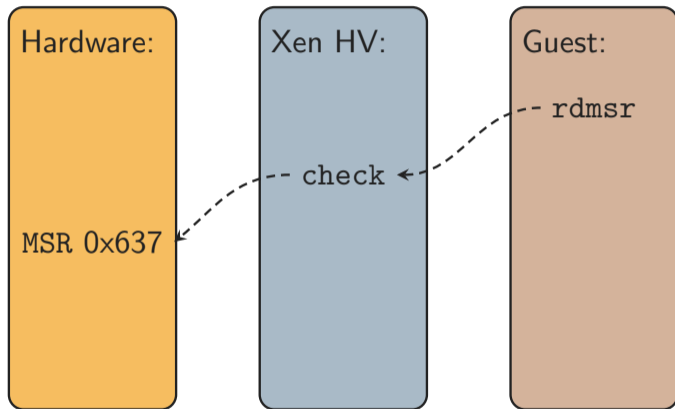
Guest:

`rdmsr`

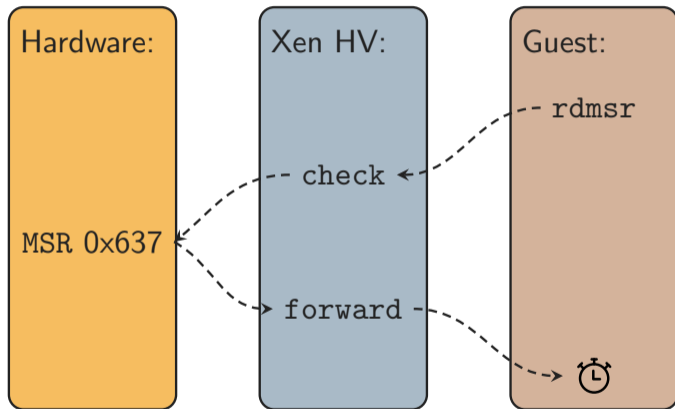
- **Hypervisor** handles MSR
- **XEN deny** list



- **Hypervisor** handles MSR
- **XEN deny** list

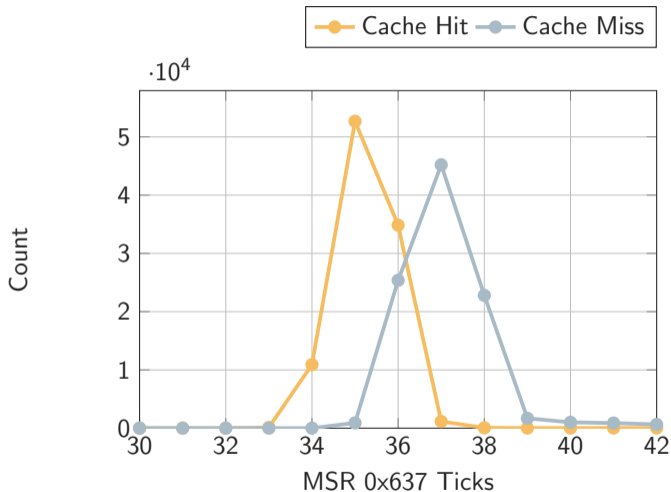


- **Hypervisor** handles MSRs
- **XEN deny** list
- **Unrestricted** read access



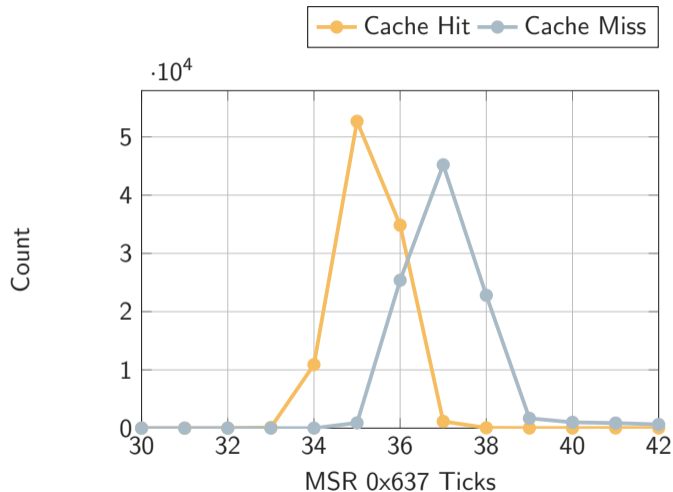
- **Hypervisor** handles MSRs
- **XEN deny** list
- **Unrestricted** read access
- **Timer** MSR

Case Study: Xen Foreshadow



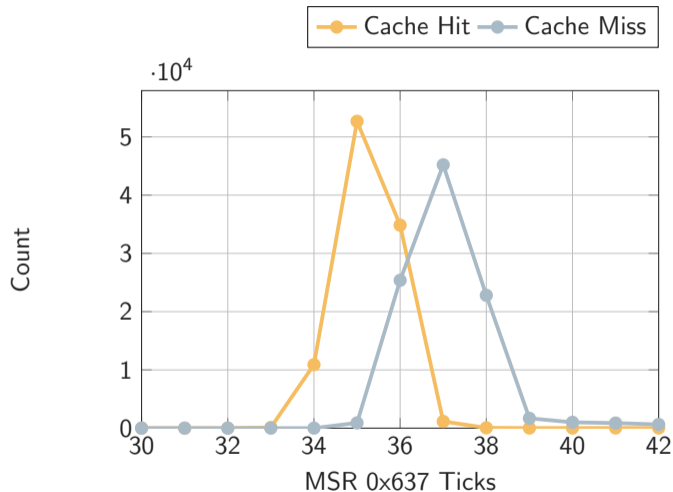
- Hypervisor handles MSR
- XEN deny list
- Unrestricted read access
- Timer MSR
 - Cache hit vs miss

Case Study: Xen Foreshadow

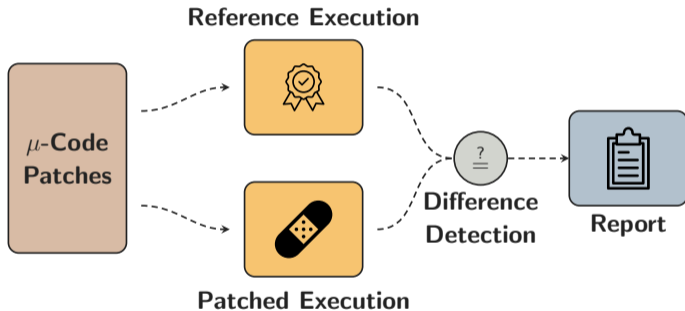


- Hypervisor handles MSRs
- XEN deny list
- Unrestricted read access
- Timer MSR
 - Cache hit vs miss
 - Foreshadow attack [4]

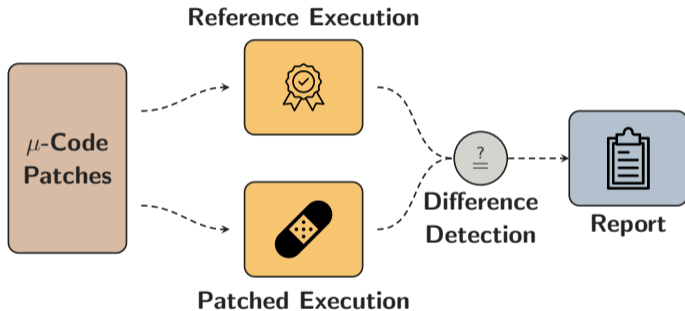
Case Study: Xen Foreshadow



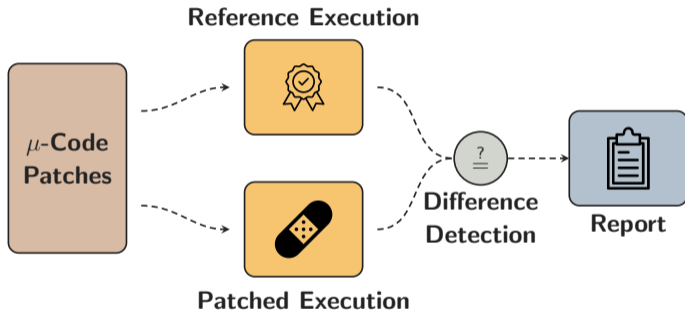
- Hypervisor handles MSRs
 - XEN deny list
 - Unrestricted read access
 - Timer MSR
 - Cache hit vs miss
 - Foreshadow attack [4]
- ✓ Leak **214** Byte/s



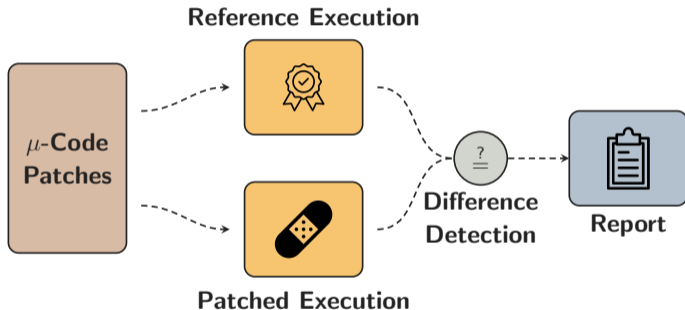
- Analyze μ -Code Patches



- Analyze μ -Code Patches
- Detect **new**



- Analyze μ -Code Patches
- Detect **new**
- Detect **affected** instructions




- Analyze μ -Code Patches
- Detect **new**
- Detect **affected** instructions
- ✓ **Before** public disclosure




- Framework  <https://github.com/IAIK/msrevelio>




- **Framework**  <https://github.com/IAIK/msrevelio>
- **Case Studies**




- **Framework**  <https://github.com/IAIK/msrevelio>
- **Case Studies**
- **MSRs enable defenses**



- **Framework**  <https://github.com/IAIK/msrevelio>
- **Case Studies**
- MSR enable defenses
- MSR open new attack vectors



- **Framework**  <https://github.com/IAIK/msrevelio>
- **Case Studies**
- MSR enable defenses
- MSR open new attack vectors
- For more details ...







- **Framework**  <https://github.com/IAIK/msrevelio>
- **Case Studies**
- MSR enables defenses
- MSR opens new attack vectors
- For more details, see the paper

Read the Paper

Finding and Exploiting CPU Features using MSR Templating

IEEE Symposium on Security and Privacy 2022

 **Andreas Kogler**  **@0xhilbert**  **andreas.kogler@iaik.tugraz.at**

-  Christopher Domas. Hardware Backdoors in x86 CPUs. In: Black Hat US (2018).
-  Daniel Gruss, Clémentine Maurice, Anders Fogh, Moritz Lipp, and Stefan Mangard. Prefetch Side-Channel Attacks: Bypassing SMAP and Kernel ASLR. In: CCS. 2016.
-  Hany Ragab, Alyssa Milburn, Kaveh Razavi, Herbert Bos, and Cristiano Giuffrida. CrossTalk: Speculative Data Leaks Across Cores Are Real. In: S&P. 2021.
-  Ofir Weisse, Jo Van Bulck, Marina Minkin, Daniel Genkin, Baris Kasikci, Frank Piessens, Mark Silberstein, Raoul Strackx, Thomas F Wenisch, and Yuval Yarom. Foreshadow-NG: Breaking the virtual memory abstraction with transient out-of-order execution. In: (2018).

Overall Results

CPU	AMD	Intel			
	Threadripper 1920X	i7-6700k	i7-8700k	i9-9900k	Xeon Silver 4208
μ -Arch	Zen	Skylake	Coffee Lake	Coffee Lake	Cascade Lake
μ -Code	0x8001137	0x9e	0xb4	0xde	0x5003102
# Found ¹	5244 (5223, 17, 4)	477 (363, 108, 5)	517 (388, 122, 7)	537 (413, 117, 7)	1109 (957, 142, 10)
# Undoc ¹	4876 (4873, 2, 1)	105 (68, 35, 2)	126 (89, 35, 2)	136 (99, 35, 2)	647 (591, 52, 4)
# Static ²	4873 (4871, 2)	99 (68, 31)	121 (89, 32)	132 (99, 33)	601 (553, 48)
# Dynamic ²	2 (2, 0)	4 (0, 4)	3 (0, 3)	2 (0, 2)	42 (38, 4)
# Similar	0	2	3	2	42

¹ \sum (RW, RO, WO) ² \sum (RW, RO)