

# Beyond Controlled-Channel Attacks: Information Leakage from Elementary CPU Behavior

*Jo Van Bulck*   Raoul Strackx   Frank Piessens

imec-DistriNet, KU Leuven, June 20, 2017


Distributed Trust Workshop on Data Protection and Privacy

# Road Map

- 1 Introduction
- 2 Basic Attack
- 3 Sancus PMA
- 4 Intel SGX
- 5 Conclusions



[thehackernews.com/2015/10/windows-patch-update.html](http://thehackernews.com/2015/10/windows-patch-update.html)



[thehackernews.com/2016/03/android-root-hack.html](http://thehackernews.com/2016/03/android-root-hack.html)

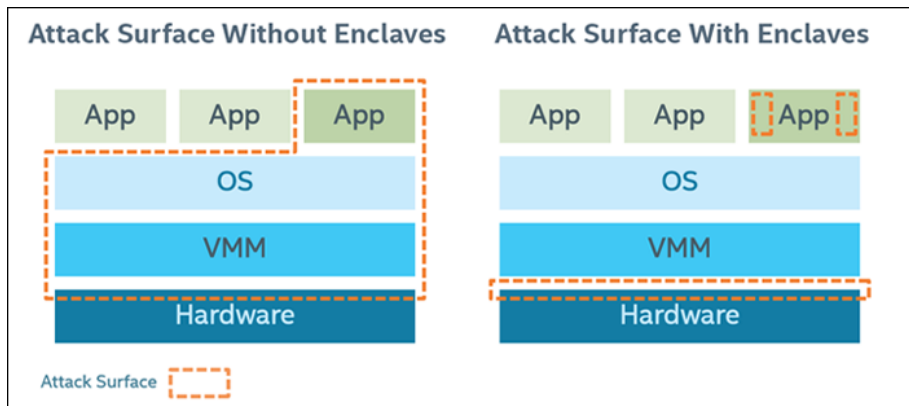


[thehackernews.com/2017/02/linux-kernel-local-root.html](http://thehackernews.com/2017/02/linux-kernel-local-root.html)




[thehackernews.com/2015/04/rootpipe-mac-os-x-vulnerability.html](http://thehackernews.com/2015/04/rootpipe-mac-os-x-vulnerability.html)

# Motivation: Application Attack Surface



<https://software.intel.com/en-us/articles/intel-software-guard-extensions-tutorial-part-1-foundation>

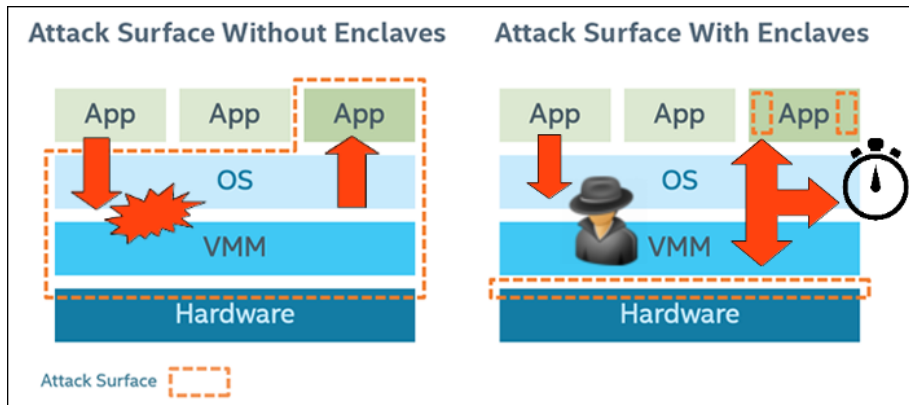
# Motivation: Application Attack Surface



<https://software.intel.com/en-us/articles/intel-software-guard-extensions-tutorial-part-1-foundation>

Layered architecture ↔ **hardware-only TCB**

# Motivation: Application Attack Surface



<https://software.intel.com/en-us/articles/intel-software-guard-extensions-tutorial-part-1-foundation>


Untrusted OS → new class of powerful **side-channels**

# Side-Channel Attack Principle



Source: <https://commons.wikimedia.org/wiki/File:WinonaSavingsBankVault.JPG>

# Side-Channel Attack Principle




Source: <https://flic.kr/p/69sHDa>



# Road Map


- 1 Introduction
- 2 Basic Attack**
- 3 Sancus PMA
- 4 Intel SGX
- 5 Conclusions

# Fetch-Decode-Execute CPU Operation



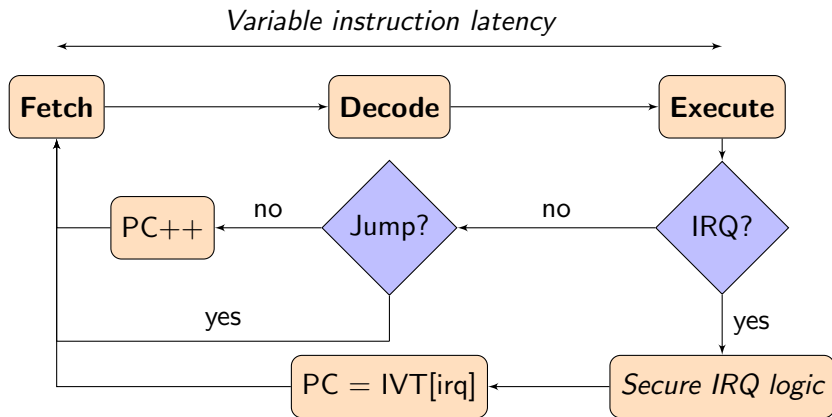
# Fetch-Decode-Execute CPU Operation

**Note:** IRQ only served *after current instruction* has completed




# Wait a Cycle ...

⇒ **IRQ latency leaks instruction execution time (!)**



# Interrupt Latency as a Side-Channel




# Road Map

- 1 Introduction
- 2 Basic Attack
- 3 Sancus PMA**
- 4 Intel SGX
- 5 Conclusions

# Sancus Protected Module Architecture

## Memory

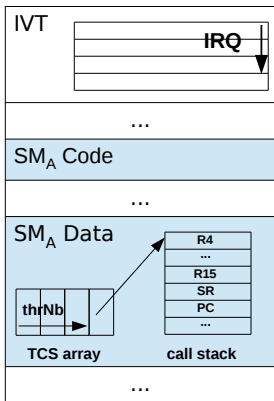


Low-cost embedded processor:

- **IoT** device: no pipeline/cache/MMU
- Extended **openMSP430** instruction set
- SM **isolation/authentication** primitives

# Sancus Protected Module Architecture

## Memory



## Low-cost embedded processor:

- **IoT** device: no pipeline/cache/MMU
- Extended **openMSP430** instruction set
- **SM isolation/authentication** primitives

## Secure interrupts:

- HW-level **interrupt engine** saves and clears CPU registers
- **Multithreading SW** extensions

Noorman et al.: "Sancus 2.0: A low-cost security architecture for IoT devices", TOPS 2017 [NVBM<sup>+</sup>17].

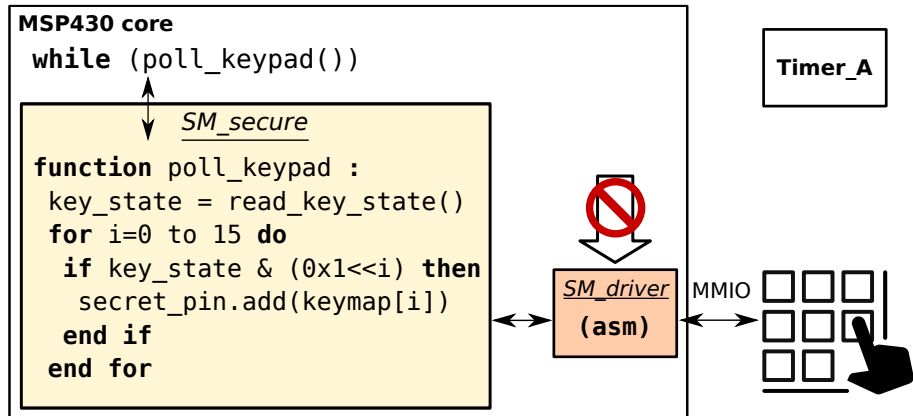
Koerberl et al.: "Trustlite: A security architecture for tiny embedded devices", EuroSys 2014 [KSSV14].

De Clercq et al.: "Secure interrupts on low-end microcontrollers", ASAP 2014 [DCPSV14].

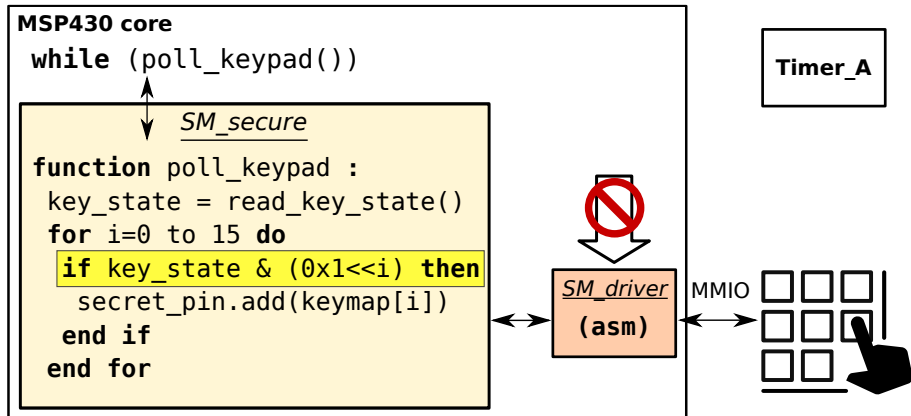
Van Bulck et al.: "Towards availability and real-time guarantees for protected module architectures", MASS 2016 [VBNMP16].



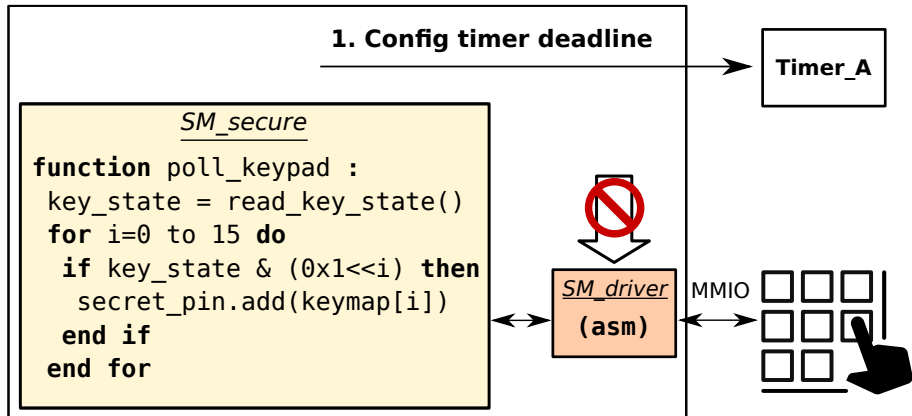
# Secure Keypad Application Scenario



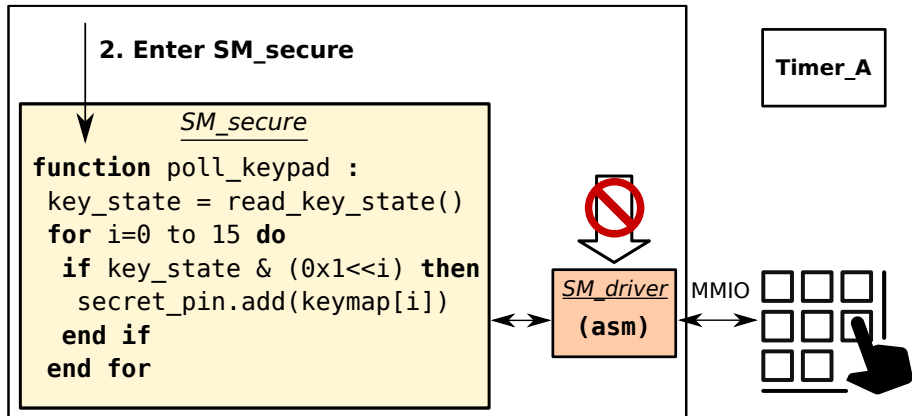
# Secure Keypad Application Scenario



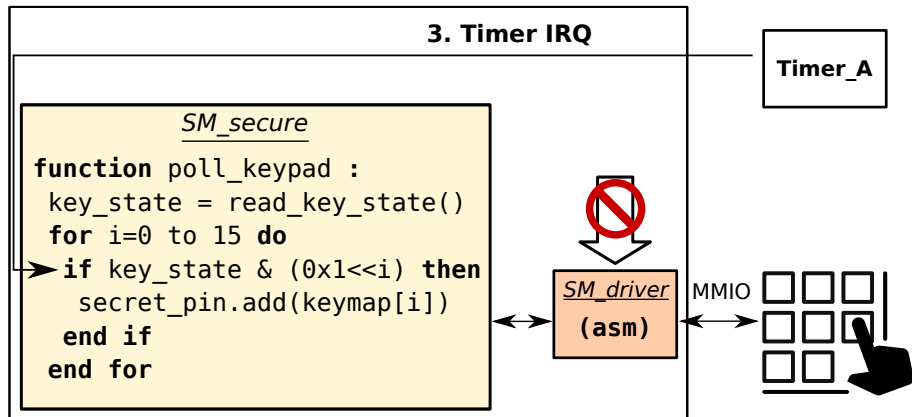
# Secure Keypad Application Scenario



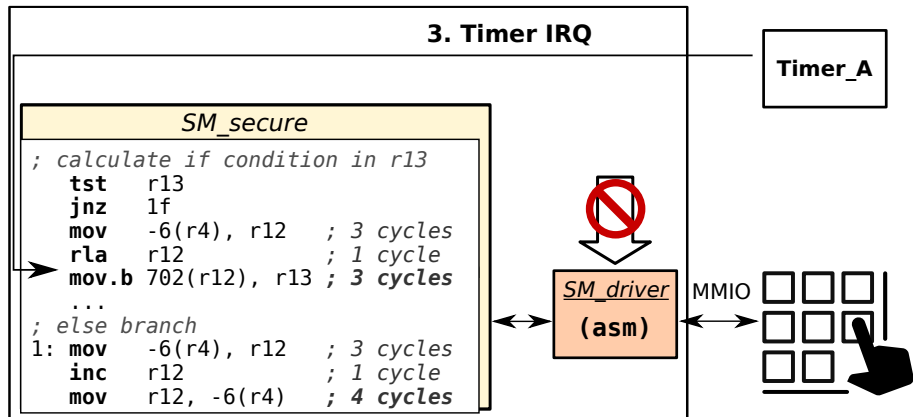
# Secure Keypad Application Scenario



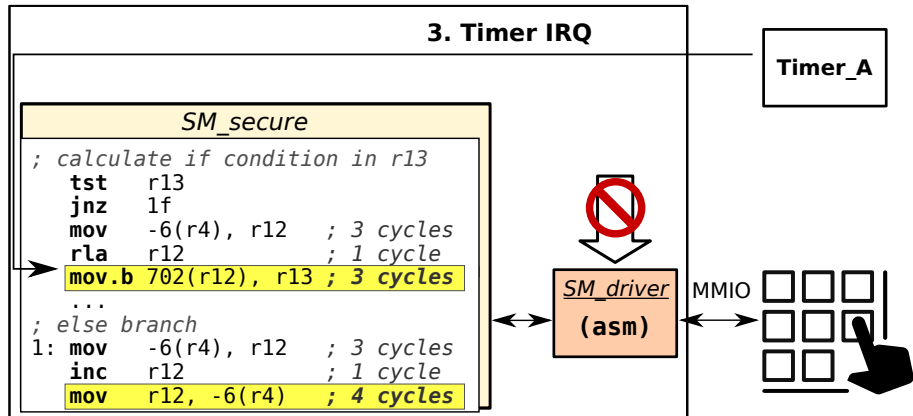
# Secure Keypad Application Scenario



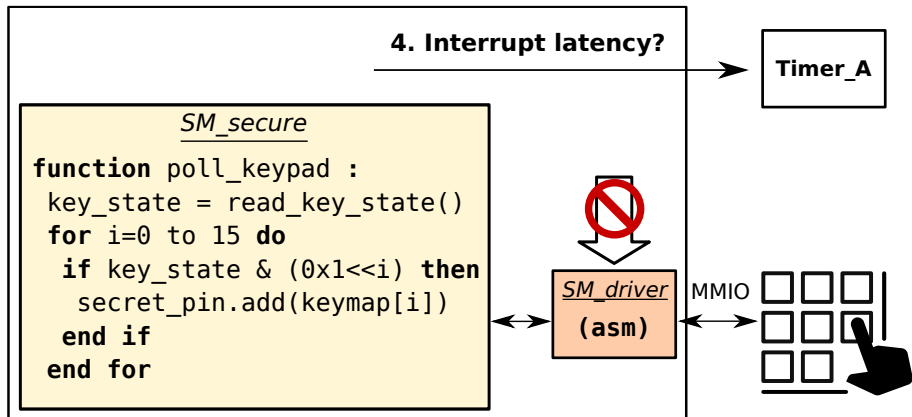
# Secure Keypad Application Scenario



# Secure Keypad Application Scenario



# Secure Keypad Application Scenario

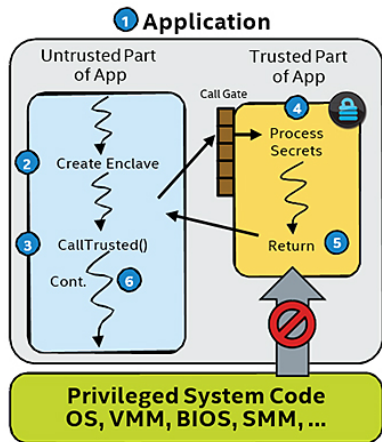




# Road Map

- 1 Introduction
- 2 Basic Attack
- 3 Sancus PMA
- 4 Intel SGX**
- 5 Conclusions

# Intel SGX Helicopter View



- Protected enclave in application's **virtual address space**
- x86** CPU:  $\exists$  pipeline, cache, out-of-order execution, ...
- Secure **interrupt** hardware mechanism: AEX/ERESUME

<https://software.intel.com/en-us/sgx/details>

# Interrupting and Resuming Enclaves

**Goal:** single-step through SGX enclave: interrupt each instruction sequentially and record corresponding *IRQ latency trace*

# Interrupting and Resuming Enclaves

**enclave**

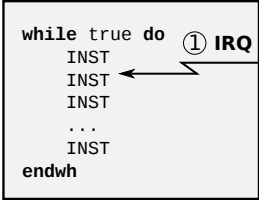
```
while true do  
  INST  
  INST  
  INST  
  ...  
  INST  
endwh
```



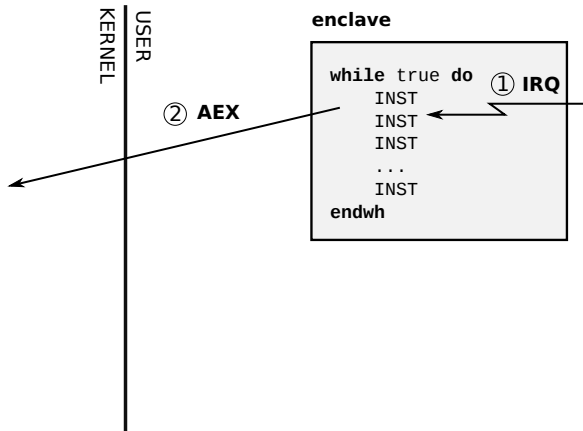
# Interrupting and Resuming Enclaves

**enclave**

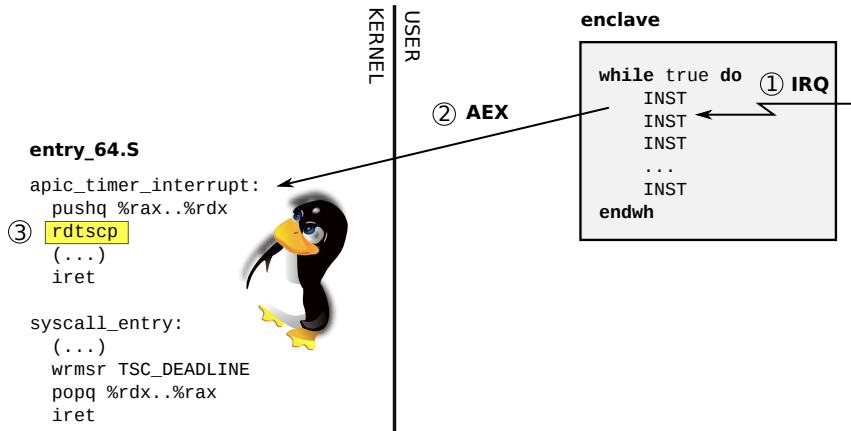
```
while true do ① IRQ  
  INST  
  INST  
  INST  
  ...  
  INST  
endwh
```

A diagram showing an enclave loop. The code is enclosed in a box. The code consists of a 'while true do' loop containing several 'INST' instructions, followed by an 'endwh' instruction. An interrupt signal, labeled '① IRQ', enters from the right side of the box. A zigzag arrow points from the IRQ signal to the first 'INST' instruction inside the loop, indicating an interruption.

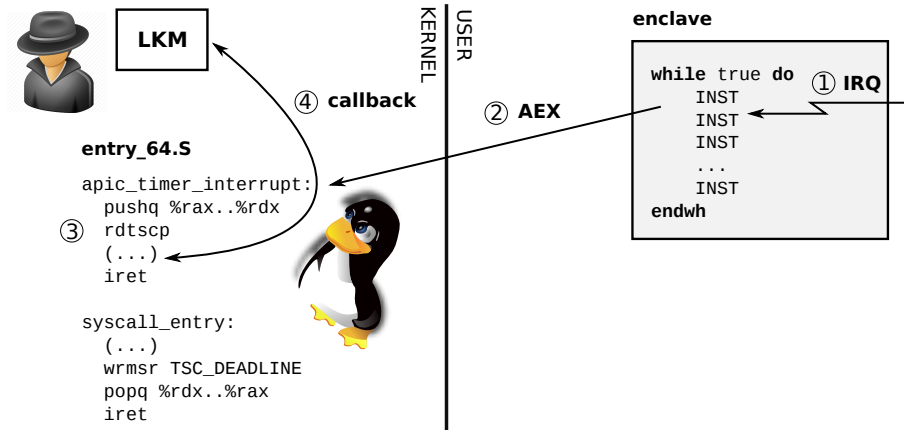
# Interrupting and Resuming Enclaves



# Interrupting and Resuming Enclaves

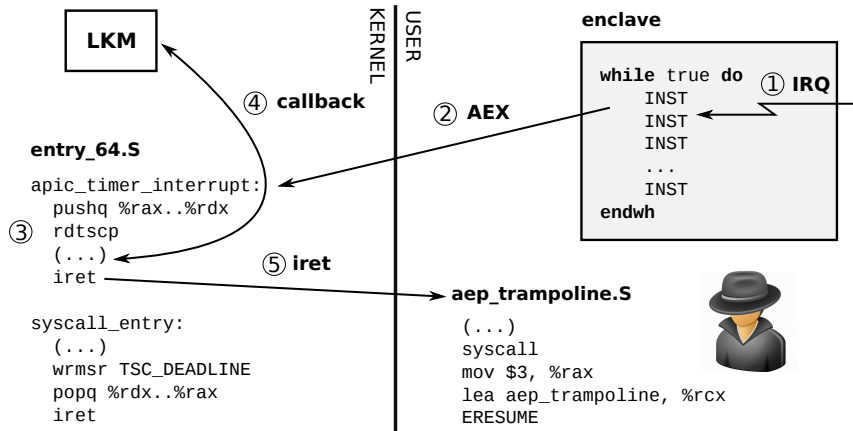


# Interrupting and Resuming Enclaves







# Interrupting and Resuming Enclaves




# Interrupting and Resuming Enclaves




# Interrupting and Resuming Enclaves



# Interrupting and Resuming Enclaves

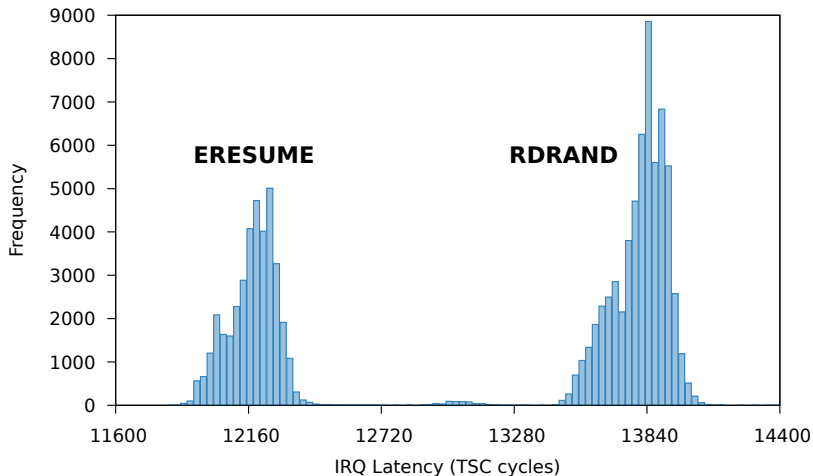


# Interrupting and Resuming Enclaves




# Configuring the Timer Interrupt

⇒ *RDRAND* execution time  $\gg$  cycles to complete *ERESUME*




# Microbenchmarks: x86 Latency Distributions


**Note:** IRQ latency leaks *interrupted instruction type*




Register increment




Hardware-level random number



Floating point operation




No-operation




Spin loop hint (delaying nop)

# Microbenchmarks: Data Caching Behavior


**Note:** IRQ latency leaks *micro-architectural cache state*




Enclave store



Enclave store non-temporal



Enclave load




Unprotected flush + load




# Microbenchmarks: Address Translation Latency


**Note:** IRQ latency leaks *memory page accesses*




No-operation



Nop + flush code PTE



Enclave load




Load + flush code/data PTE

# Macrobenchmark: Modular Exponentiation


```
function SQUARE_AND_MULTIPLY(c,d,e,n)  
  r  $\leftarrow$  rand()  
  c  $\leftarrow$  c * re mod n  
  m  $\leftarrow$  1  
  for most to least significant bit b in d do  
    m  $\leftarrow$  m2 mod n  
    if b then  
      m  $\leftarrow$  m * c mod n  
    end if  
  end for  
  return m * r-1 mod n  
end function
```

# Extracted IRQ Latency Trace




- “X-ray” extracted from a single **dummy RSA decryption**

# Extracted IRQ Latency Trace




- “X-ray” extracted from a single **dummy RSA decryption**
- **Distinct instructions** for stack canary + blinding: RDRAND

# Extracted IRQ Latency Trace




- “X-ray” extracted from a single **dummy RSA decryption**
- **Distinct instructions** for stack canary + blinding: RDRAND
- Sharply defined **algorithm phases**

# Extracted IRQ Latency Trace



- “X-ray” extracted from a single **dummy RSA decryption**
- **Distinct instructions** for stack canary + blinding: RDRAND
- Sharply defined **algorithm phases**
- Full 16-bit **key recovery**

# Extracted IRQ Latency Trace



Flush page table entry for *global variable accessed every loop iteration*

# Road Map

- 1 Introduction
- 2 Basic Attack
- 3 Sancus PMA
- 4 Intel SGX
- 5 Conclusions**



# Conclusion

⇒ (First) remote side-channel for **embedded + high-end** trusted computing hardware

# Conclusion

⇒ (First) remote side-channel for **embedded + high-end** trusted computing hardware

IRQ latency trace reveals **micro-architectural** behavior:

- Lots of *noise/non-determinism* on modern CPUs
- Abuse subtle timing differences with *machine learning*?

# Conclusion

⇒ (First) remote side-channel for **embedded + high-end** trusted computing hardware

IRQ latency trace reveals **micro-architectural** behavior:

- Lots of *noise/non-determinism* on modern CPUs
- Abuse subtle timing differences with *machine learning*?

**Defense techniques:**

- Eliminate *secret-dependent control flow* ↔ practice
- Sancus secure *hardware patch* to mask IRQ latency

Thank you! Questions?

# References I



A. Baumann, M. Peinado, and G. Hunt.

Shielding applications from an untrusted cloud with Haven.

In *11th USENIX Symposium on Operating Systems Design and Implementation*, pp. 267–283. USENIX Association, 2014.



V. Costan, I. Lebedev, and S. Devadas.

Sanctum: Minimal hardware extensions for strong software isolation.

In *25th USENIX Security Symposium*, pp. 857–874. USENIX Association, 2016.



B. Coppens, I. Verbauwhede, K. De Bosschere, and B. De Sutter.

Practical mitigations for timing-based side-channel attacks on modern x86 processors.

In *2009 IEEE Symposium on Security and Privacy*, pp. 45–60. IEEE, 2009.



R. De Clercq, F. Piessens, D. Schellekens, and I. Verbauwhede.

Secure interrupts on low-end microcontrollers.

In *Application-specific Systems, Architectures and Processors (ASAP), 2014 IEEE 25th International Conference on*, pp. 147–152. IEEE, 2014.



P. Koeberl, S. Schulz, A.-R. Sadeghi, and V. Varadharajan.

TrustLite: A security architecture for tiny embedded devices.

In *Proceedings of the Ninth European Conference on Computer Systems*, pp. 10:1–10:14. ACM, 2014.



F. McKeen, I. Alexandrovich, A. Berenson, C. V. Rozas, H. Shafi, V. Shanbhogue, and U. R. Savagaonkar.

Innovative instructions and software model for isolated execution.

In *Proceedings of the 2nd International Workshop on Hardware and Architectural Support for Security and Privacy*, pp. 10:1–10:1. ACM, 2013.

# References II

 P. Maene, J. Götzfried, R. de Clercq, T. Müller, F. Freiling, and I. Verbauwhede.

Hardware-based trusted computing architectures for isolation and attestation.

*IEEE Transactions on Computers*, (99), 2017.



J. Noorman, J. Van Bulck, J. T. Mühlberg, F. Piessens, P. Maene, B. Preneel, I. Verbauwhede, J. Götzfried, T. Müller, and F. Freiling.

Sancus 2.0: A low-cost security architecture for IoT devices.

*ACM Transactions on Privacy and Security (TOPS)*, 2017.

Accepted for publication.



S. Shinde, Z. L. Chua, V. Narayanan, and P. Saxena.

Preventing page faults from telling your secrets.

In *Proceedings of the 11th ACM on Asia Conference on Computer and Communications Security, ASIA CCS '16*, pp. 317–328. ACM, 2016.



M. Schwarz, S. Weiser, D. Gruss, C. Maurice, and S. Mangard.

Malware guard extension: Using SGX to conceal cache attacks.

In *Detection of Intrusions and Malware, and Vulnerability Assessment (DIMVA)*, DIMVA, 2017.



J. Van Bulck, J. Noorman, J. T. Mühlberg, and F. Piessens.

Towards availability and real-time guarantees for protected module architectures.

In *MODULARITY Companion Proceedings '16*, pp. 146–151, New York, 2016. ACM.



Y. Xu, W. Cui, and M. Peinado.

Controlled-channel attacks: Deterministic side channels for untrusted operating systems.

In *2015 IEEE Symposium on Security and Privacy*, pp. 640–656. IEEE, 2015.

# Sancus: Configuring Timer IRQ

⇒ Near-exact copy “spy” SM to leak intermediary timings

9760		9762			
JNE		MOV x(r4), r12			IRQ 8
	JUMP	SRC AD	SRC RD	EXEC	IRQ 0
781	782	0	1	2	3
200290	200291	200292	200293	200294	200295

```
[main] spy SM execution time report:
```

```
-----  
738 | first key comparison  
233 | reti if to subsequent comparison  
208 | reti else to subsequent comparison
```

```
[main] enter secure PIN...
```

```
[isr] key '3' was pressed!
```

```
[isr] key '5' was pressed!
```

# Intel SGX: Reducing Noise

**BIOS** Maximize **execution time predictability**: disable Hyper-Threading, C-States, TurboBoost, SpeedStep

**Kernel** Reserve **dedicated CPU core** for victim enclave:  
`isolcpus + /proc/smp_affinity + CONFIG_NO_HZ_FULL`

**Hardware** “sub-cycle accurate” APIC **TSC-deadline** timer

