

Side Channel Risks in Hardware Trusted Execution Environments (TEEs)



中国科学院 信息工程研究所
INSTITUTE OF INFORMATION ENGINEERING, CAS

Wenhao Wang (王文浩)
May 2019



About Me

- ❑ Research Associate Professor (副研究员) of Institute of Information Engineering, CAS
- ❑ Previously worked as a visiting researcher in Indiana University (with Prof. XiaoFeng Wang)
- ❑ Research Interests
 - System Security
 - Computer Architectural Security
 - Isolation with Hardware Features
 - Privacy Preserving Computing Technologies
 - Cryptography (esp. Symmetric Cryptanalysis)
- ❑ Email: wangwenhao@iie.ac.cn

Scenarios when Hardware TEEs are needed

- ❑ Users' private data are delegated to untrusted (public) cloud servers
- ❑ Multi-sources (federated) deep learning training
- ❑ Machine-Learning As A Service
- ❑ Data sharing of genomic data or big data

Scenarios when Hardware TEEs are needed


❑ Crypto Techniques

- FHE、 MPC、 Searchable Encryption、 ZK *etc.*
- **Extremally High Communication and Computation Overhead**

❑ Hardware Techniques

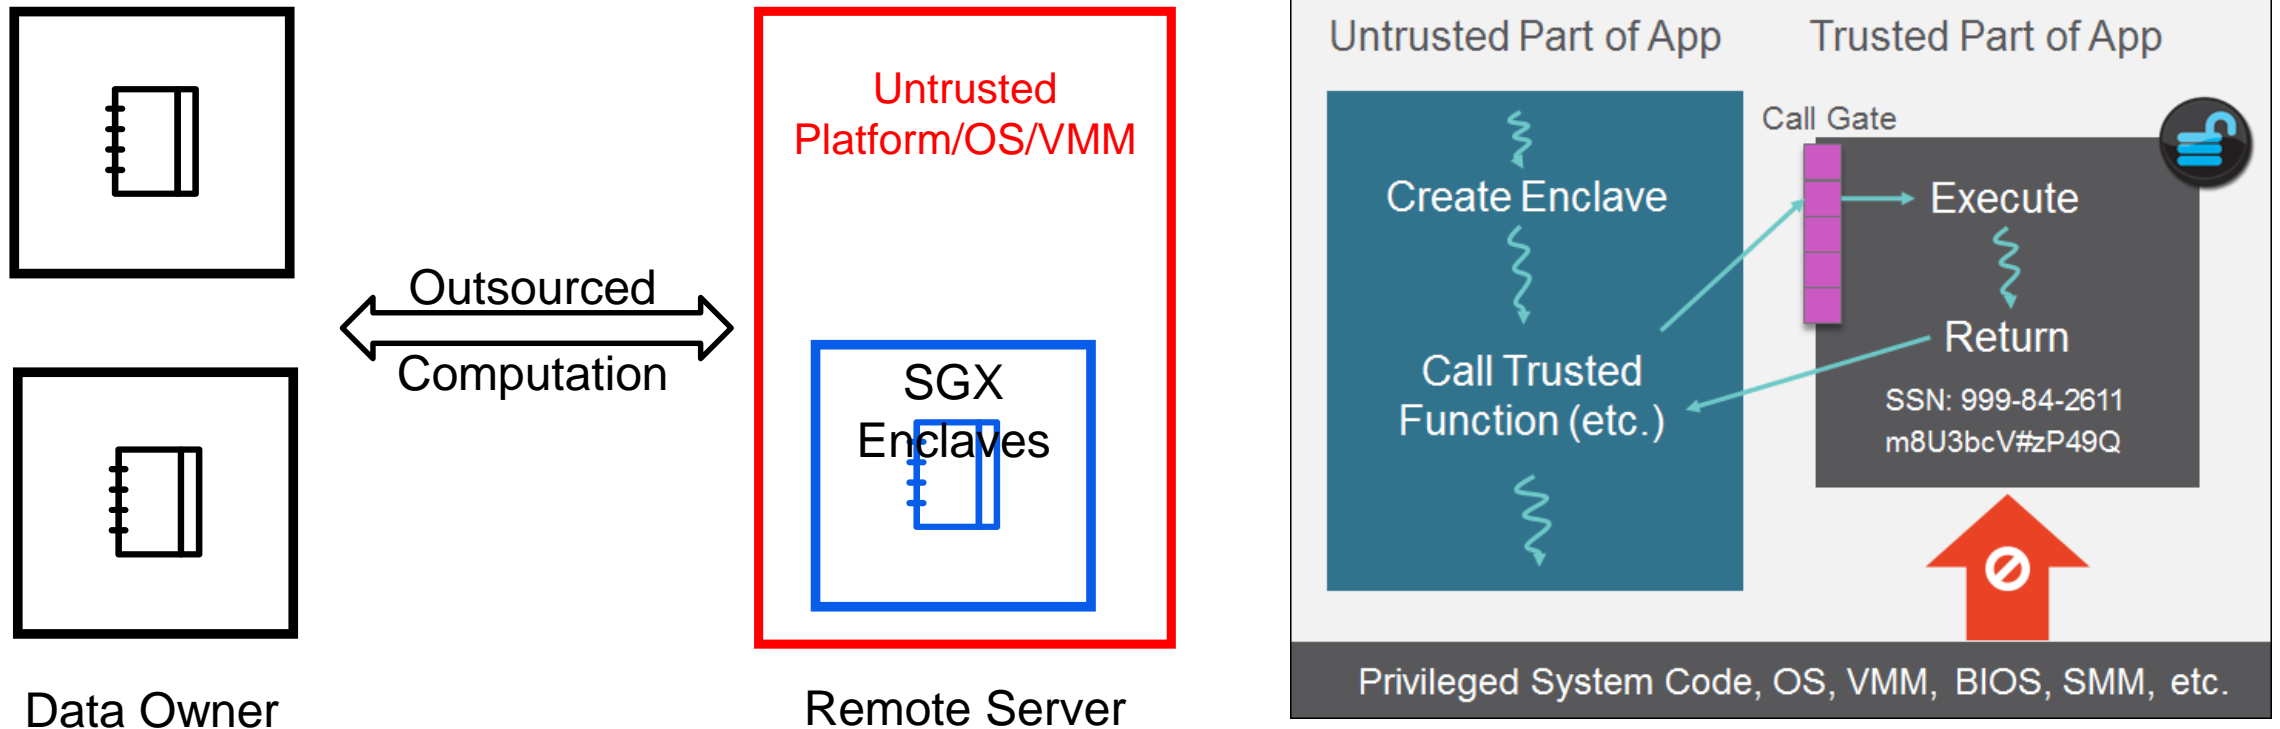
- Intel TXT, ARM Trustzone, Intel SGX, AMD SEV *etc.*

Hardware TEEs – A Review



Excluded: Intel CAT/CET/SMEP/SMAP/VT-x/PT

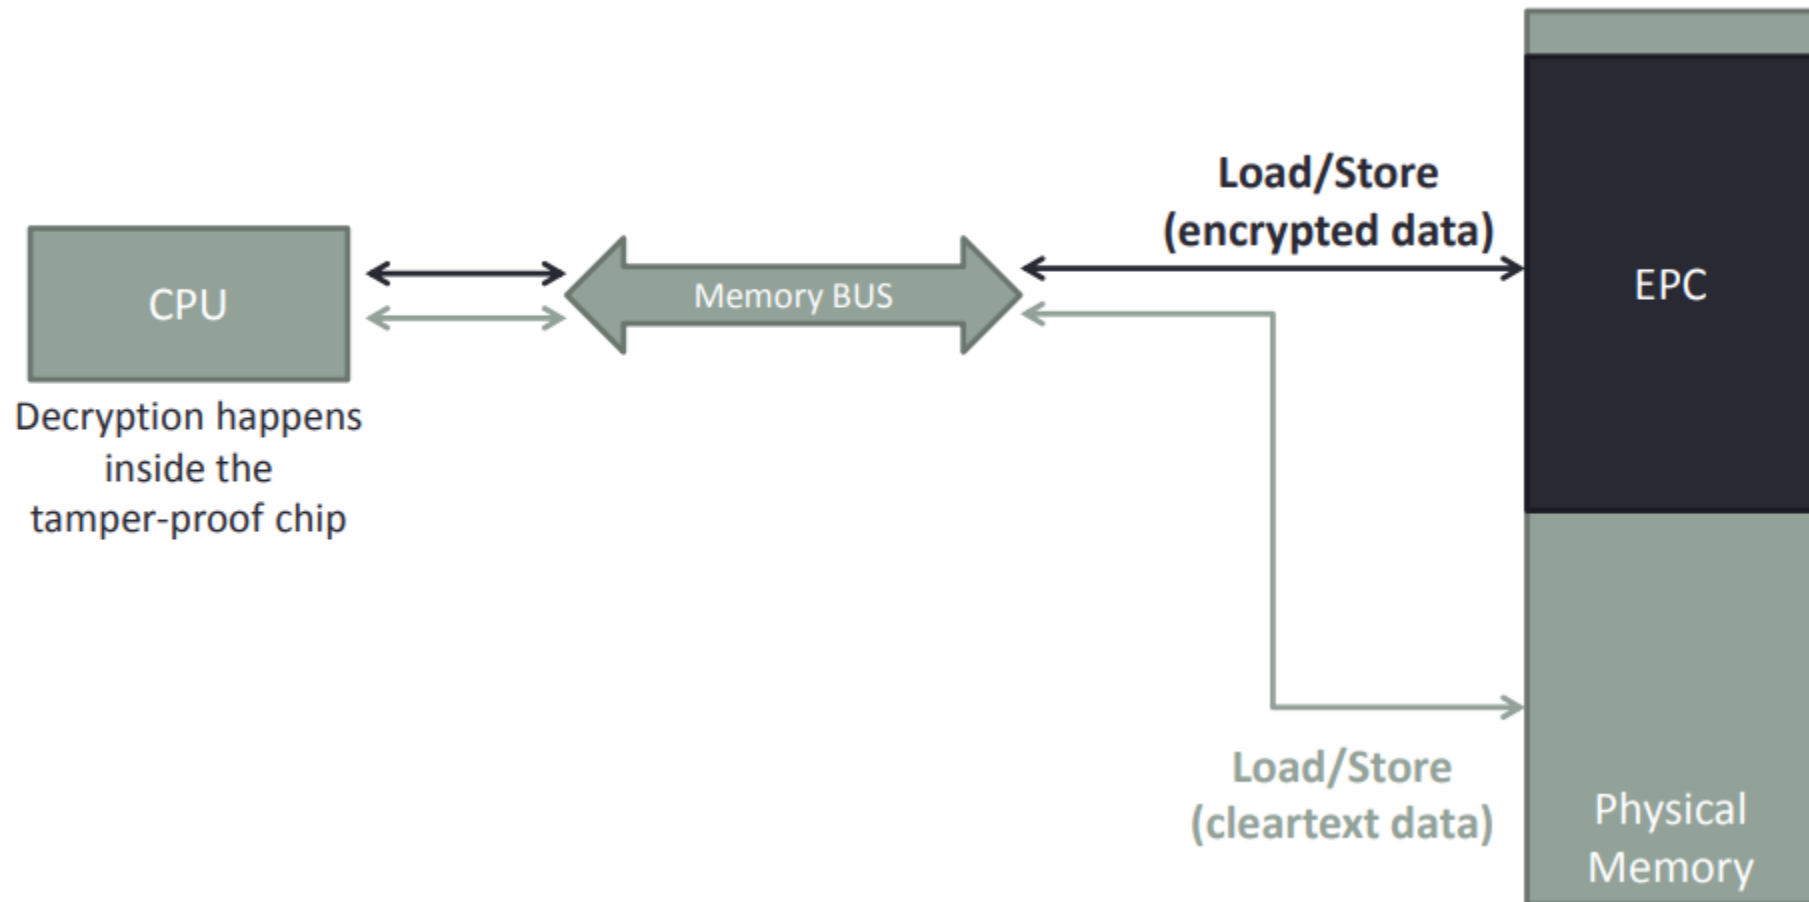
Intel SGX



- Memory Encryption
- Access Control
- Remote Attestation

Intel SGX

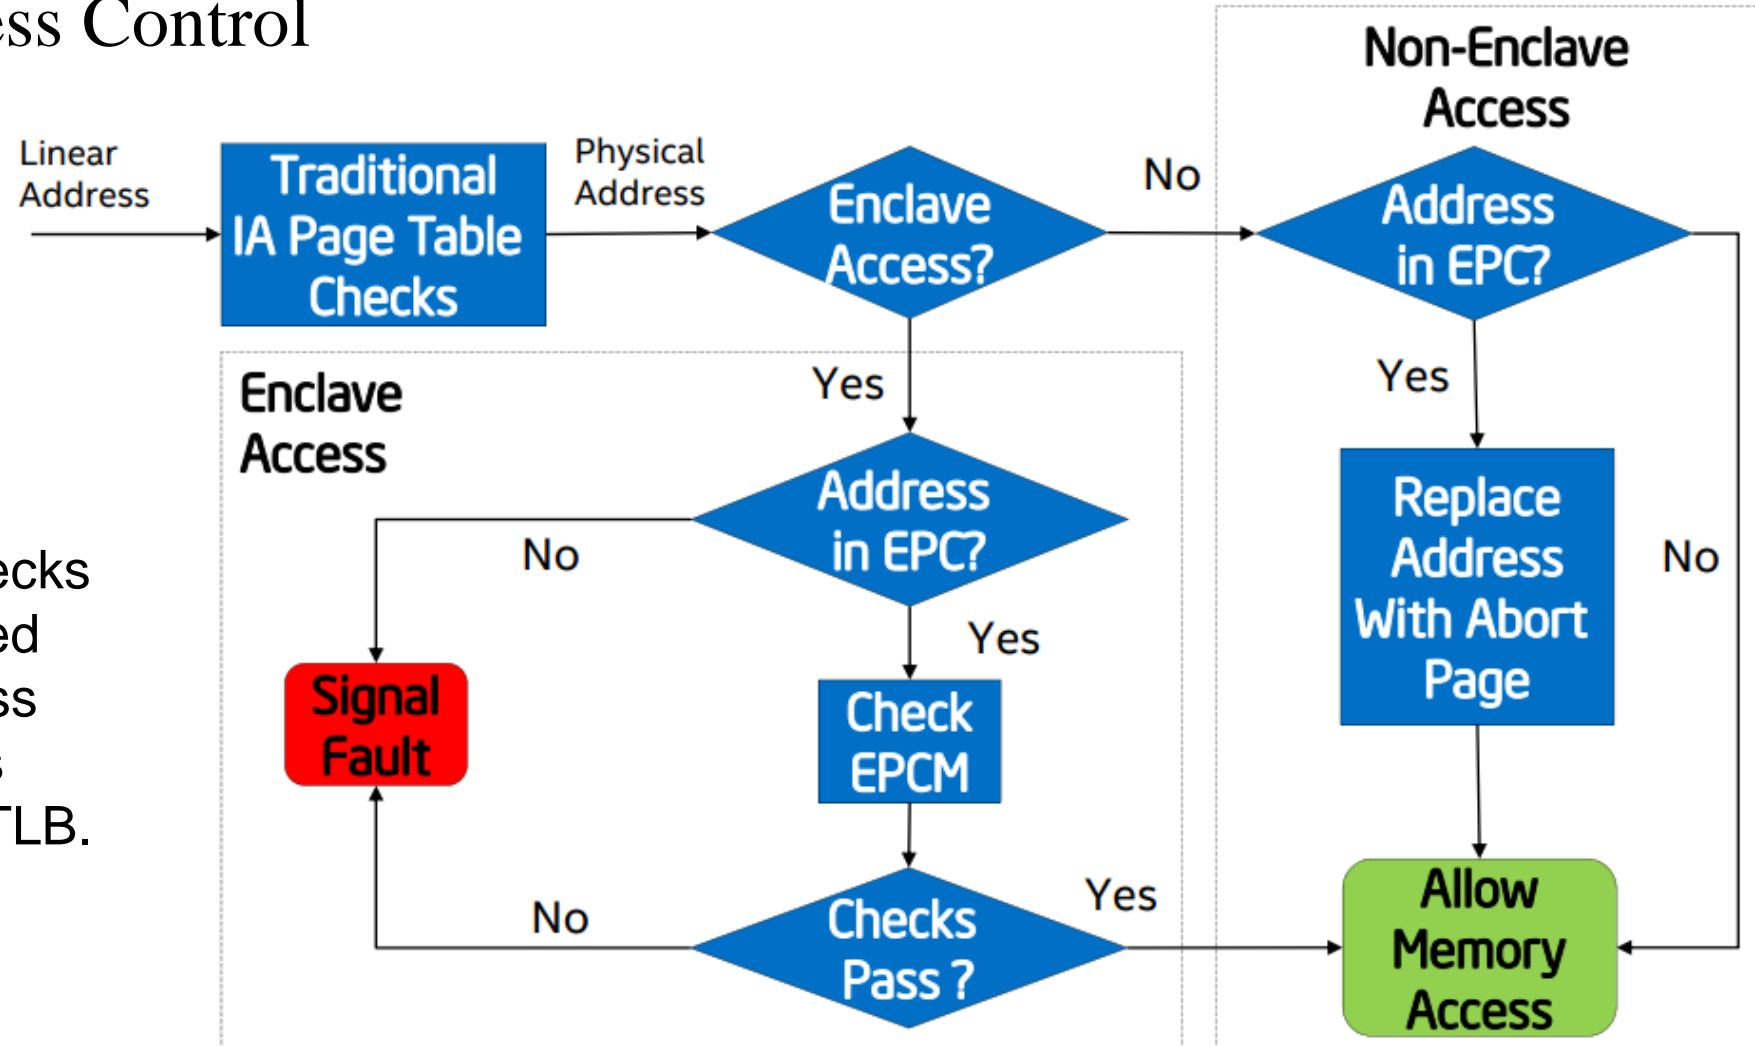
- ❑ Enclave memory is stored within the Enclave Page Cache



Intel SGX

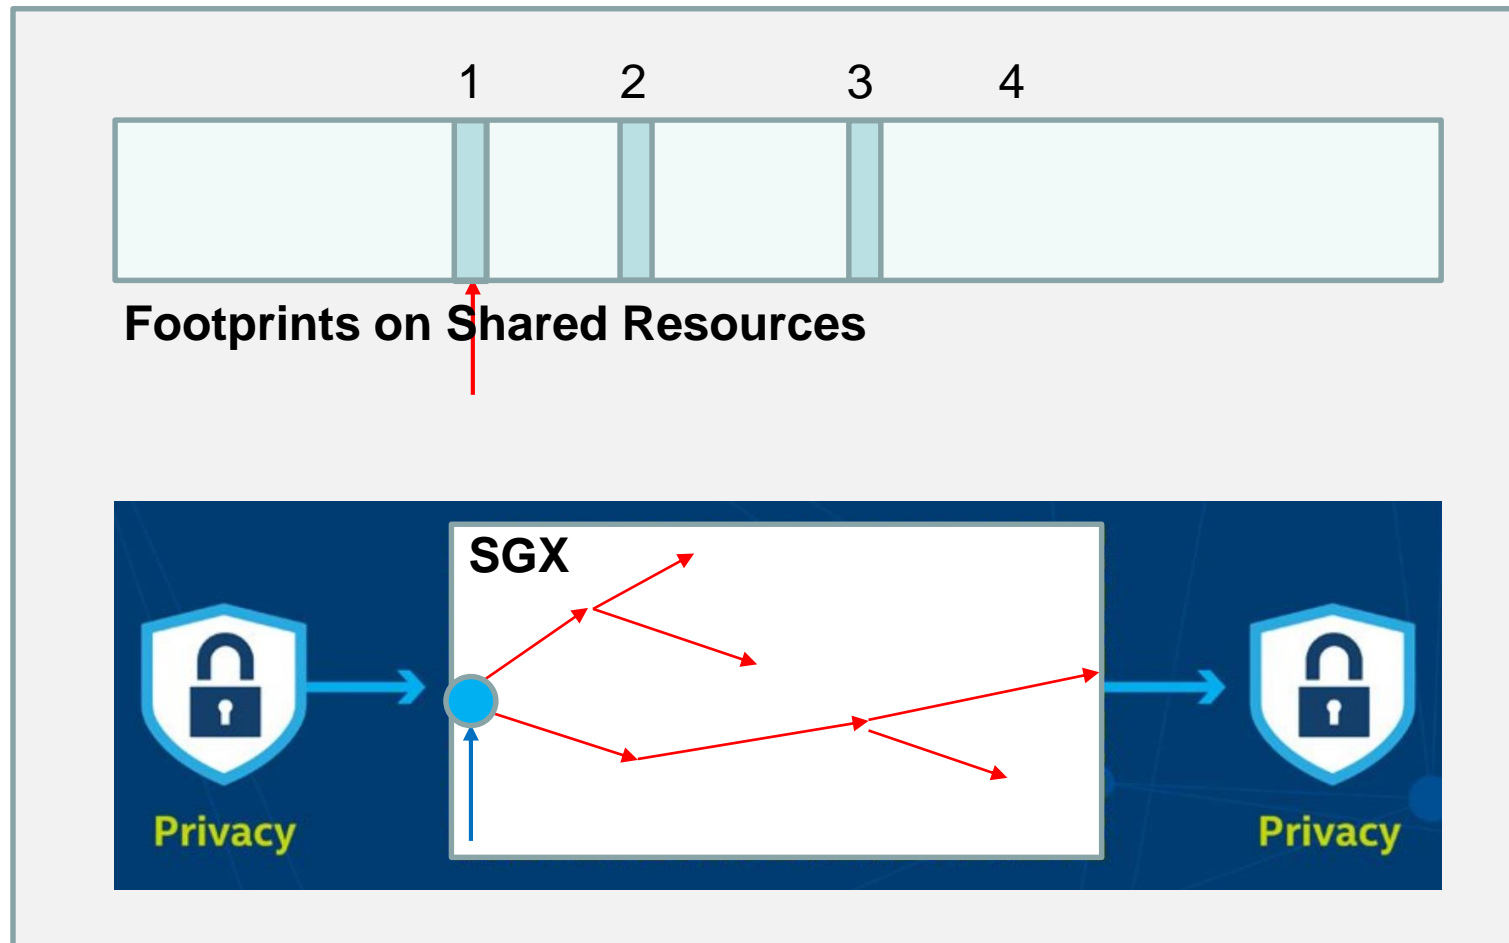
Access Control

Security Checks are performed when address translation is loaded into TLB.




What is a side channel?

- ❑ Side channels from resources shared crossing multi-domains




Side channels – An example (Cache Timing Attacks)


- ❑ The cache holds copies of aligned blocks of B bytes in main memory (blocks).
- ❑ When a memory access instruction is processed, memory cell is searched in the cache first.
- ❑ If a cache miss occurs, a full memory block is copied into the appropriate set (S possible sets) into one of the W cache lines.



Side channels – An example (Cache Timing Attacks)

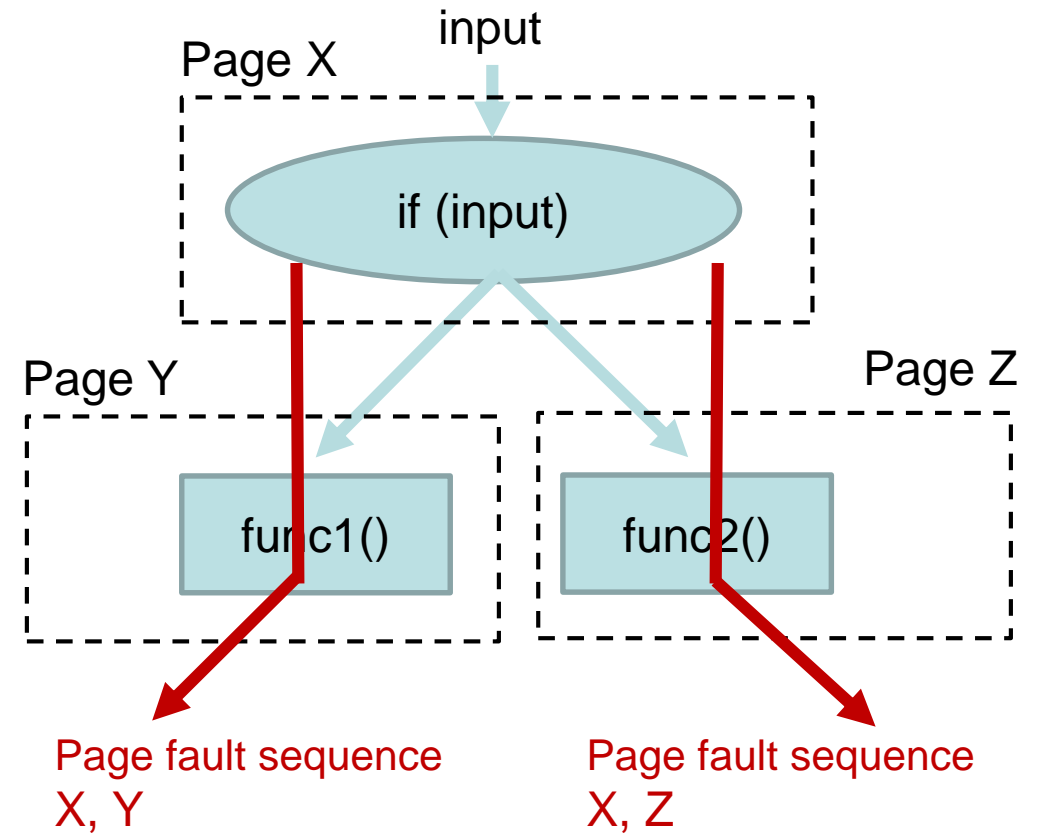
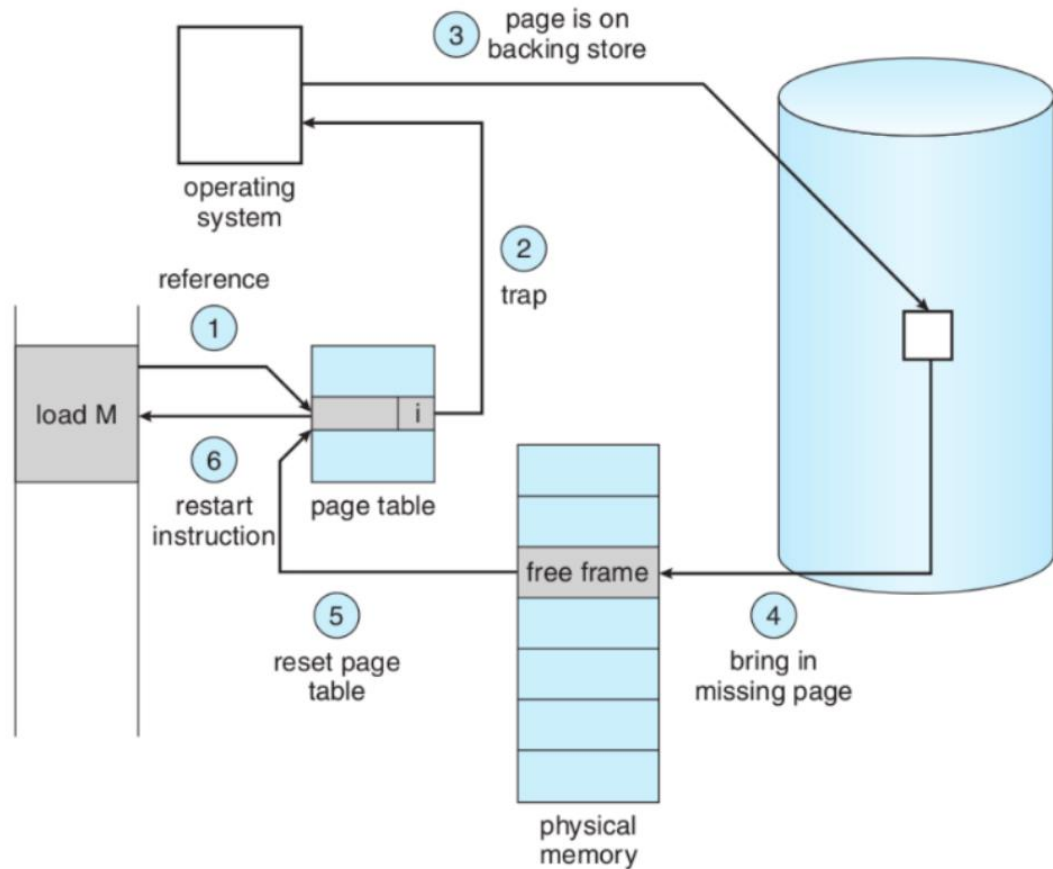


Side channels – An example (Cache Timing Attacks)



1. Completely evict victim data from cache
2. Trigger a victim data access
3. Access attacker memory again and see which cache sets are slow

Side channels – An example (Controlled-channel Attacks)



Side channels – Others?

❑ Memory Hierarchy

- Data Caching creates fast and slow execution paths, leading to timing differences depending on whether data is in the cache or not

❑ Function Unit Contention

- Sharing of hardware leads to contention, whether a program can use some hardware leaks information about other programs

❑ Stateful Functional Units

- Program's behavior can affect state of the function units (e.g. branching target), and other programs can observe the output (which depends on the state)

❑ Variable Instruction Execution Timing

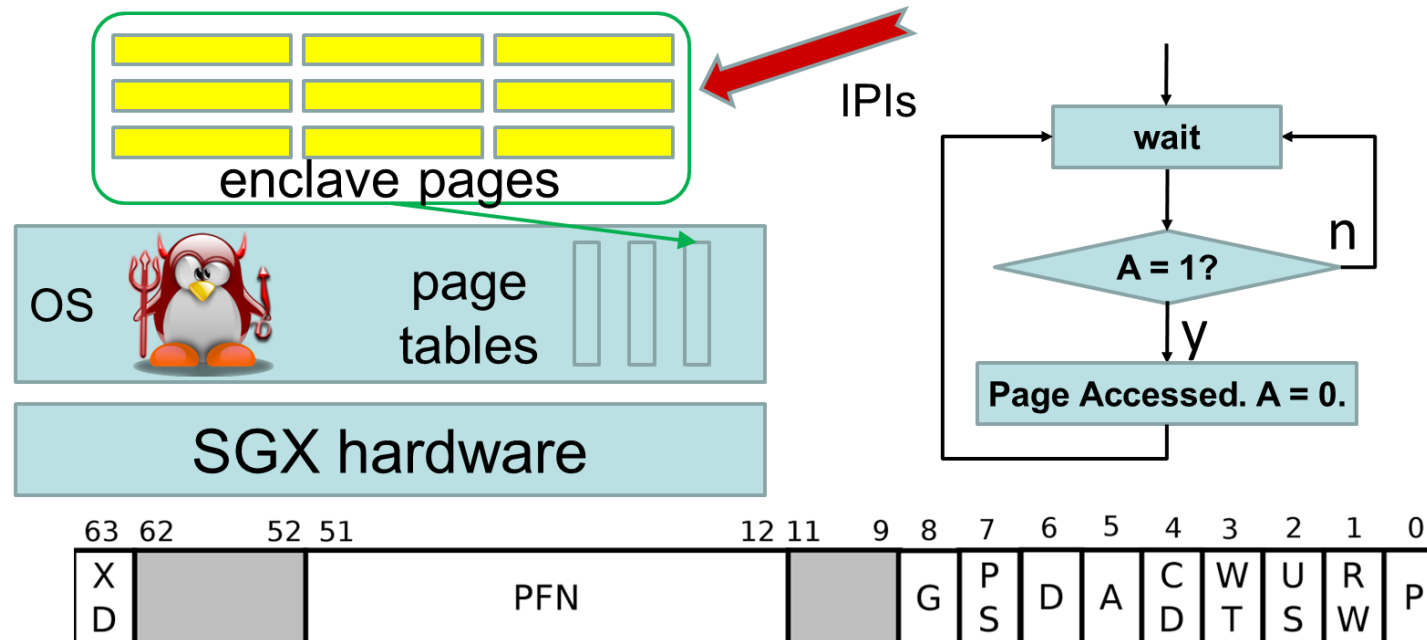
- Execution of different instructions or same instruction with different operands takes different amount of time

❑ Physical Emanations

- Execution of programs affects physical characteristics of the chip, such as thermal changes (e.g. avx512), which can be observed


Can we reduce the interrupts for page based attacks?

❑ 1. Passive observation over the Access bit of a PTE



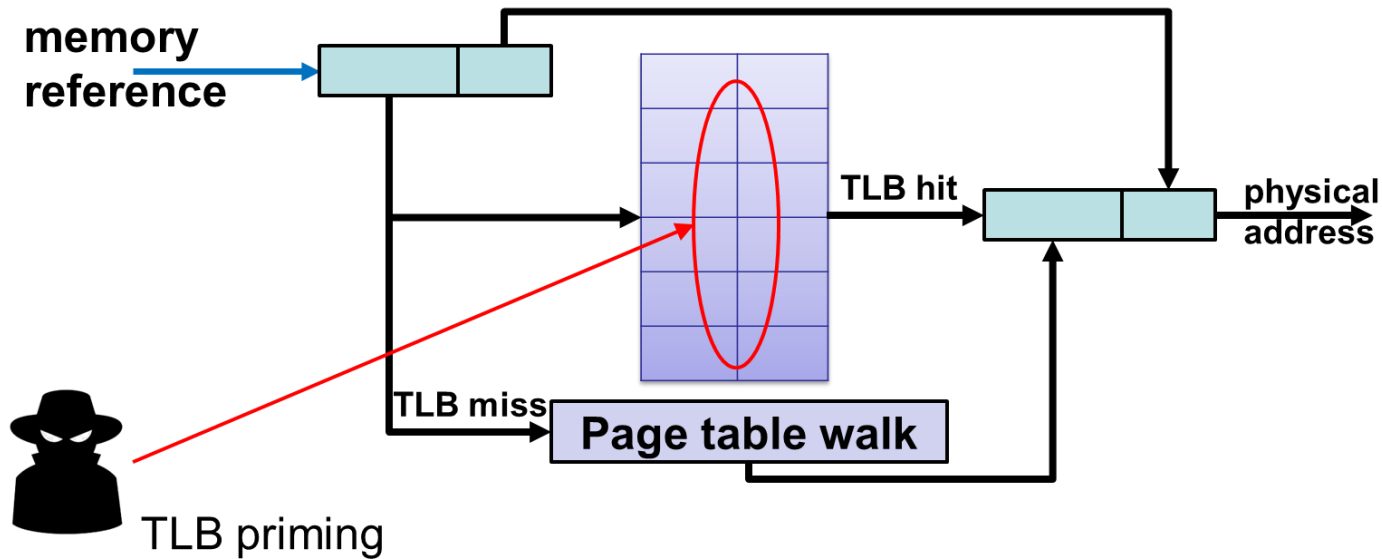
Can we reduce the interrupts for page based attacks?

□ 2. Measuring the time between accesses to pages



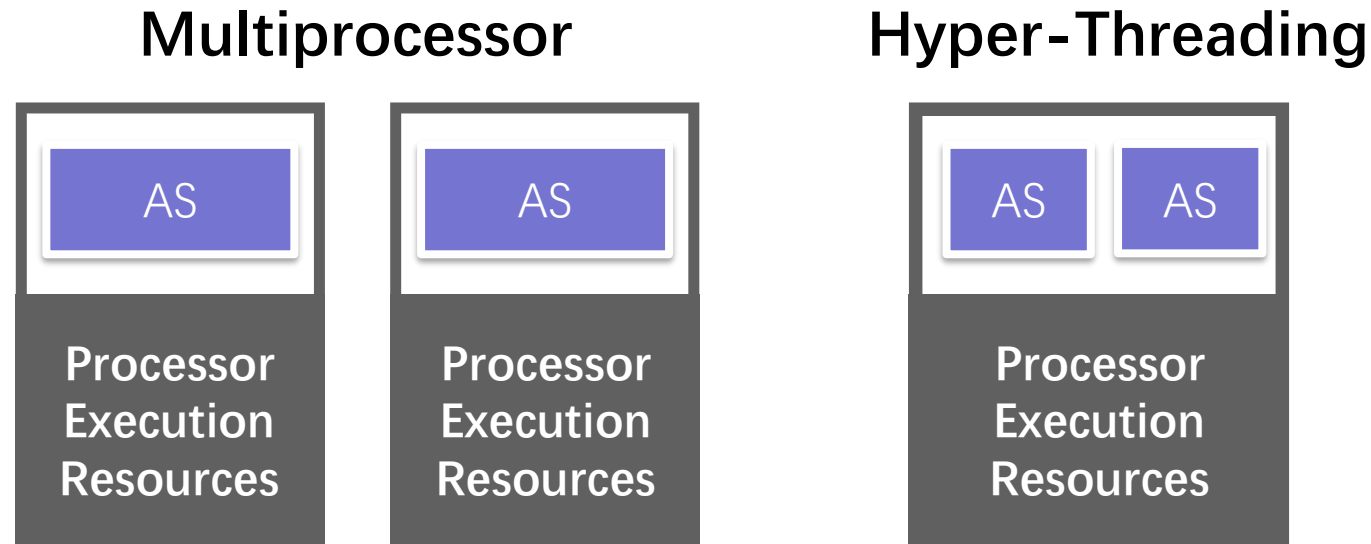
Can we reduce the interrupts for page based attacks?

- 3. Clearing TLB entries from the other Hyper-thread to force a page table walk



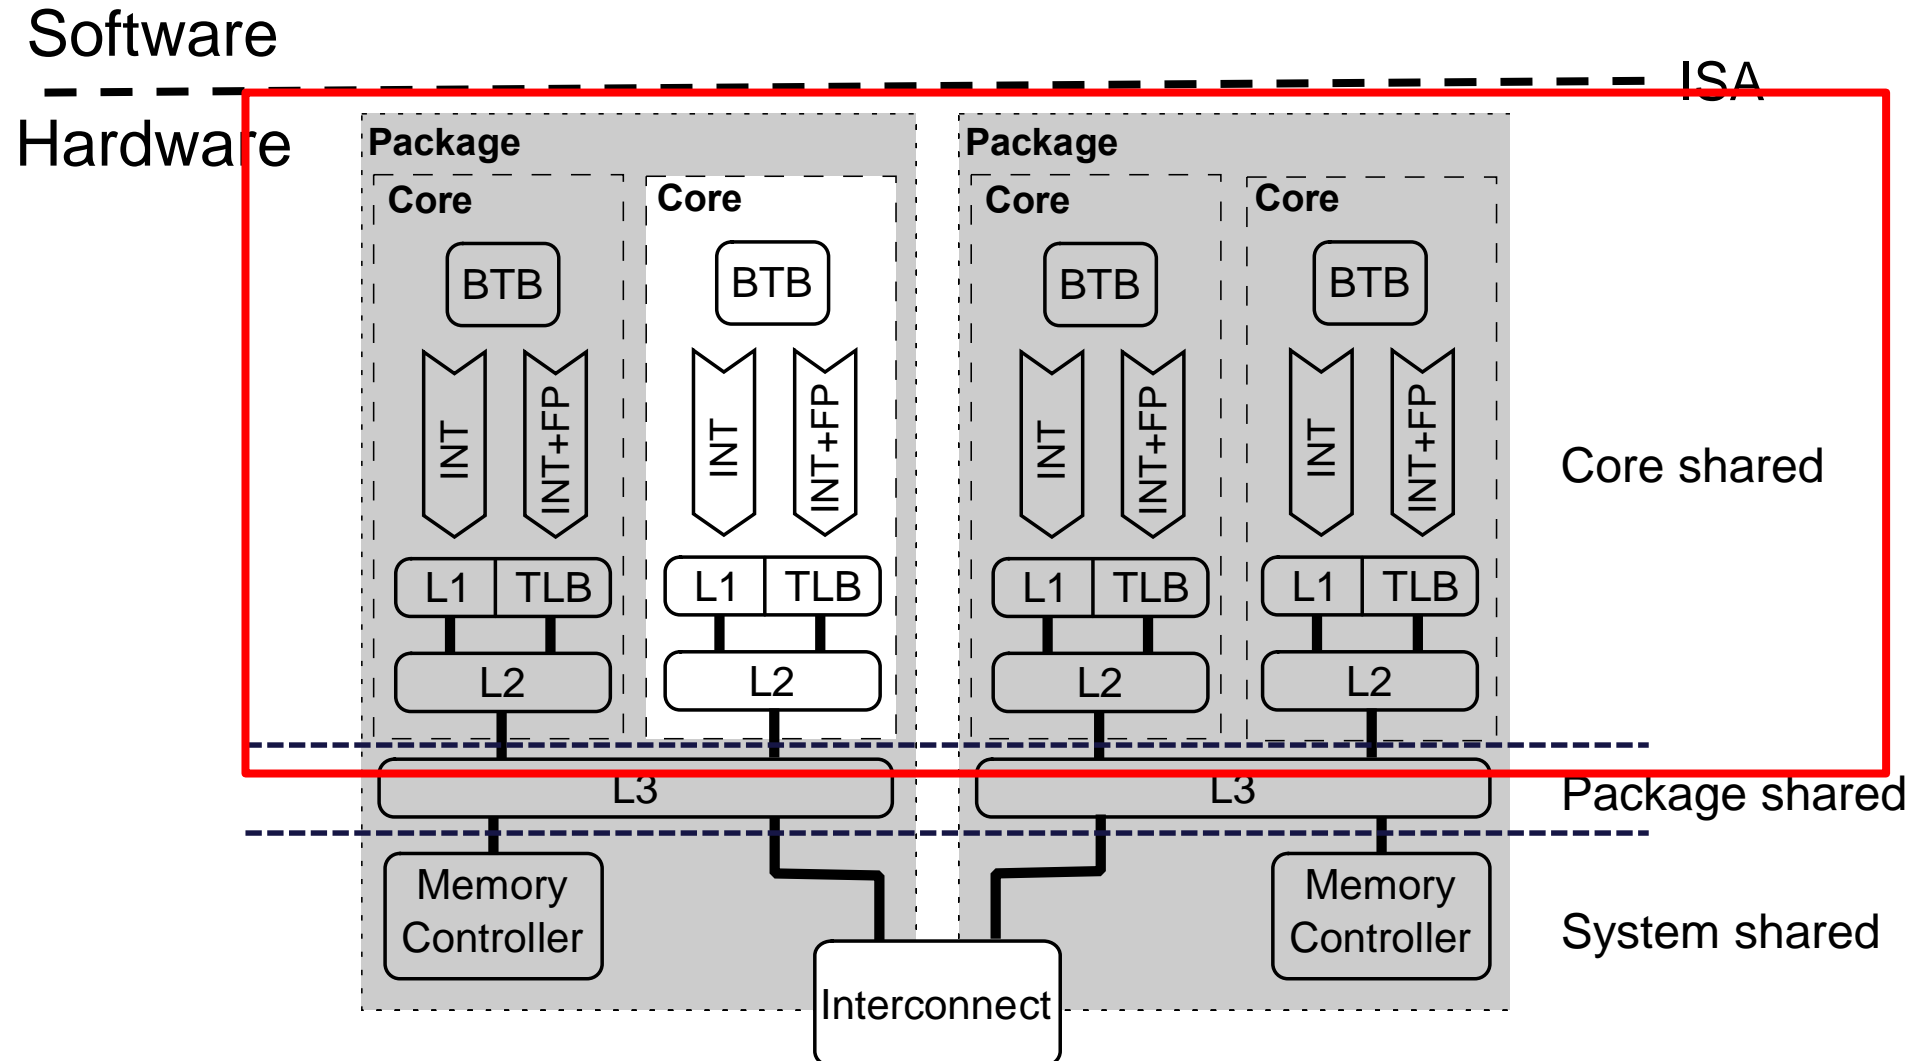
Hyper-threading (SMT)

- ❑ Hyper-Threading enables new side channel attack surfaces



AS: architectural state (eax, ebx, control registers, etc.)

Problems with Hyper-Threading

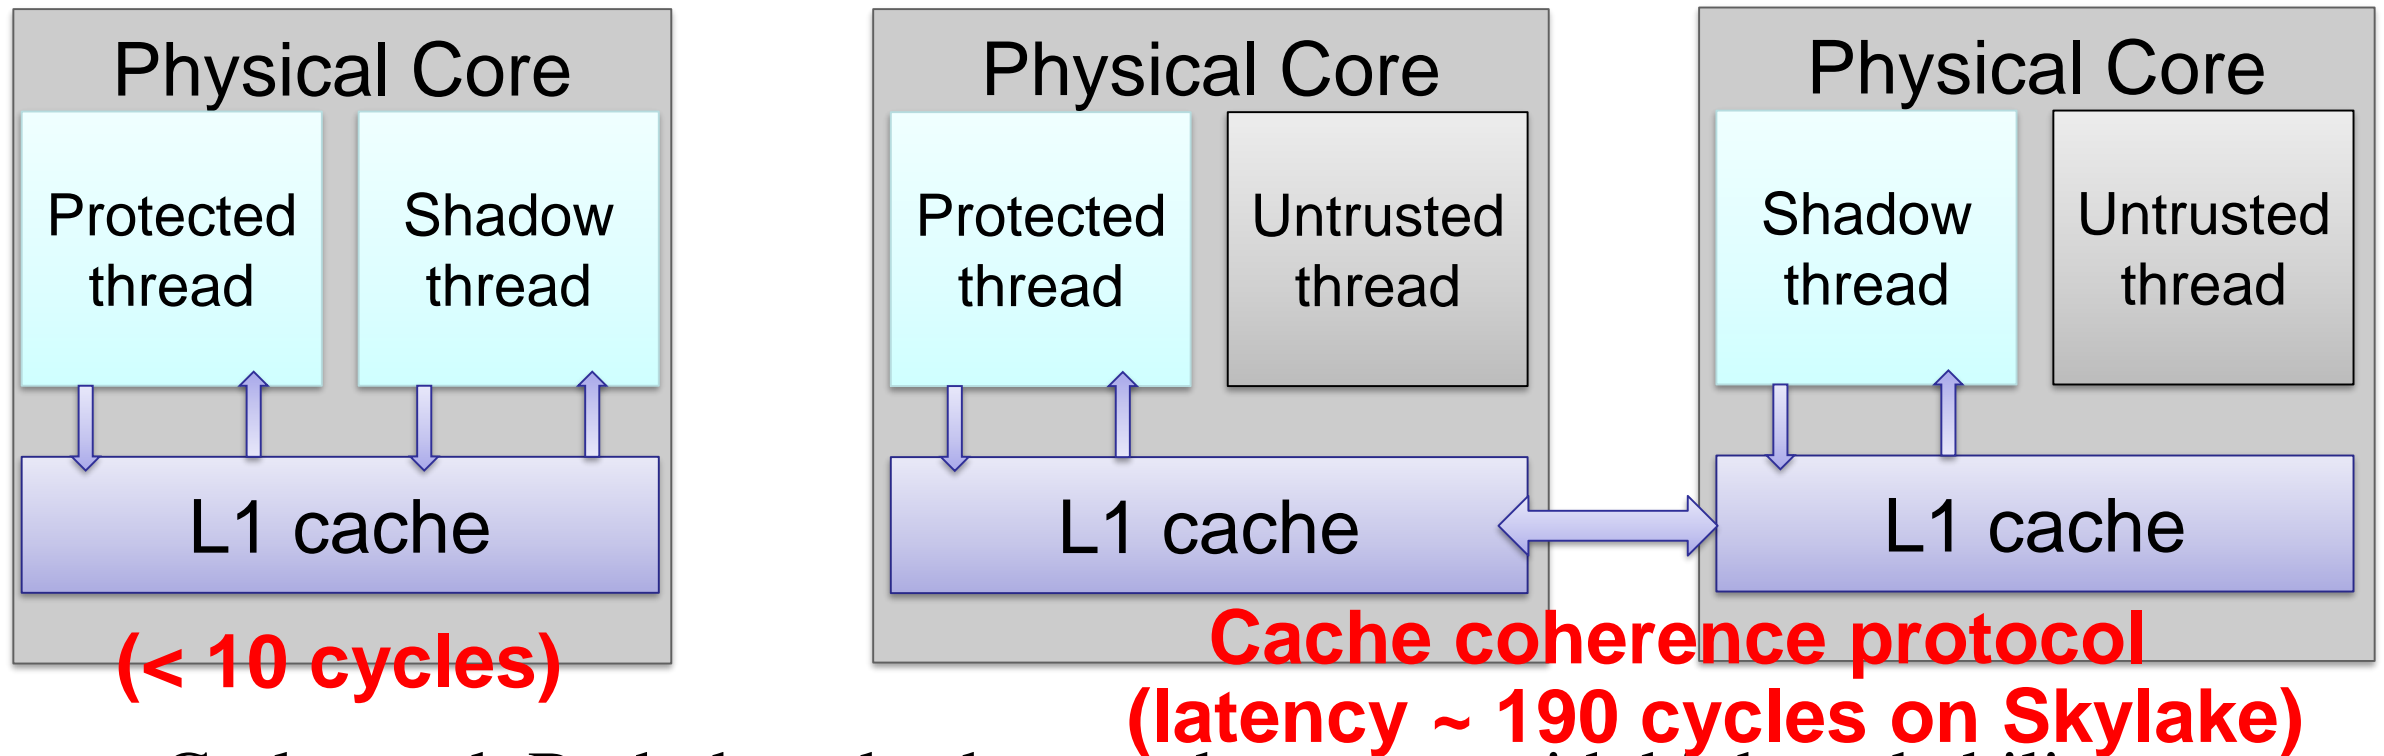


Naïve Solutions do not work

- ❑ Simply disabling Hyper-Threading
 - No effective way to verify
 - `cpuid`, `rdtscp` and `rdpid` are not supported in enclave mode
 - Remote attestation
 - Does not contain information about Hyper-Threading (before our work)
- ❑ Create a shadow thread from the enclave program to occupy the other hyper-thread
- ❑ How to reliably verify the physical-core co-location?

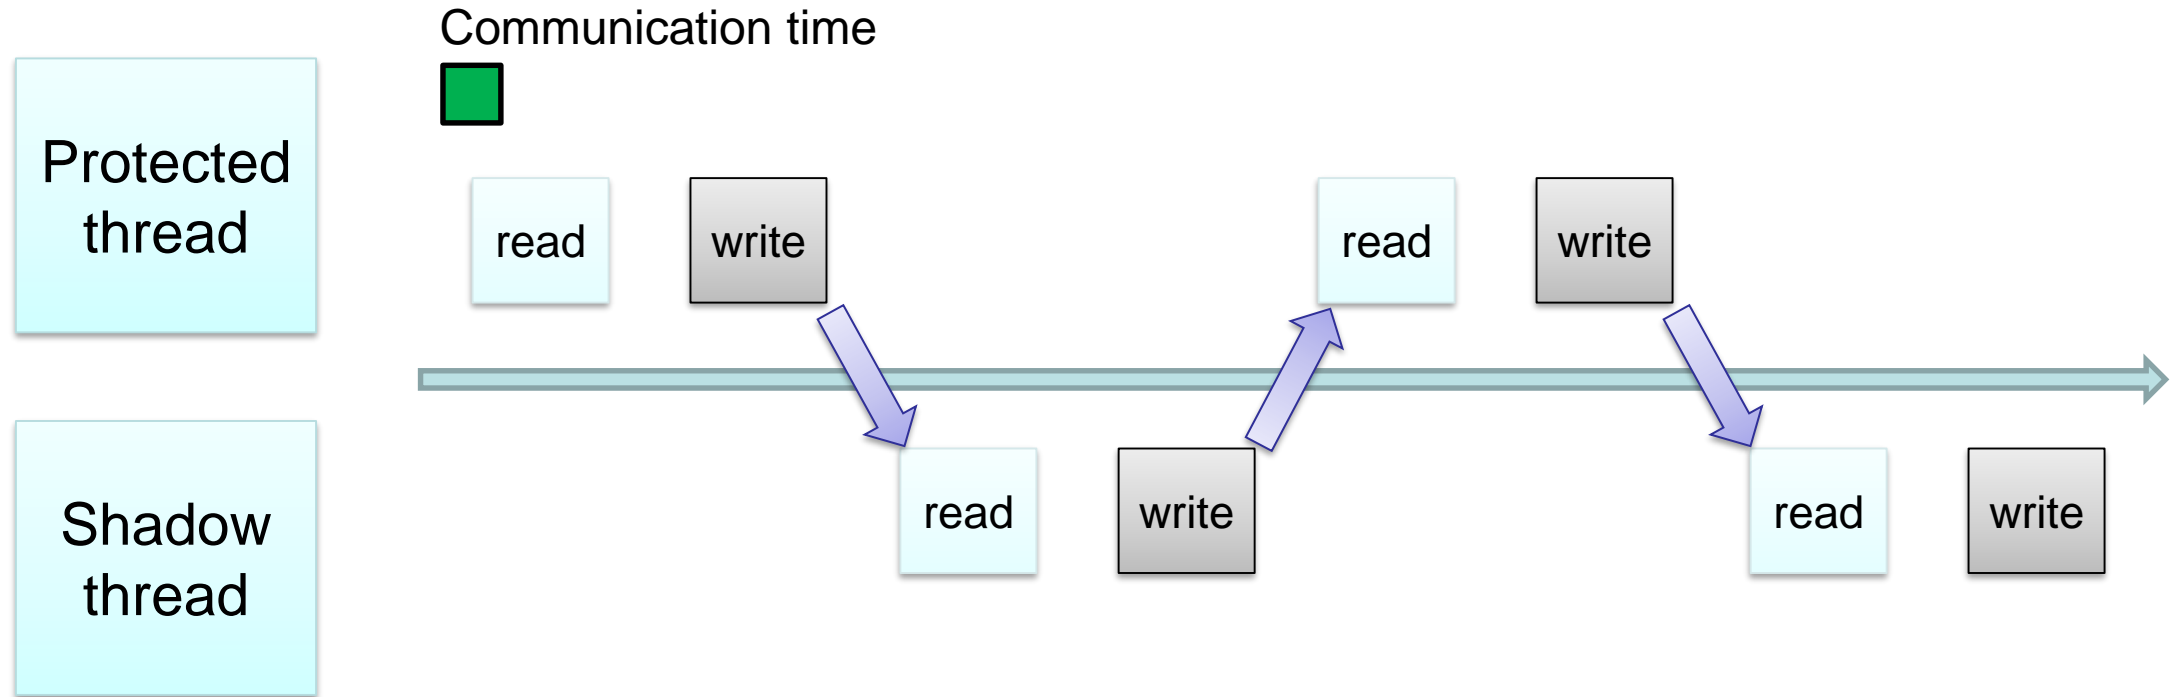
Closing HT-Side Channels on SGX with Contrived Data Races

❑ Co-location test with Contrived Data Races



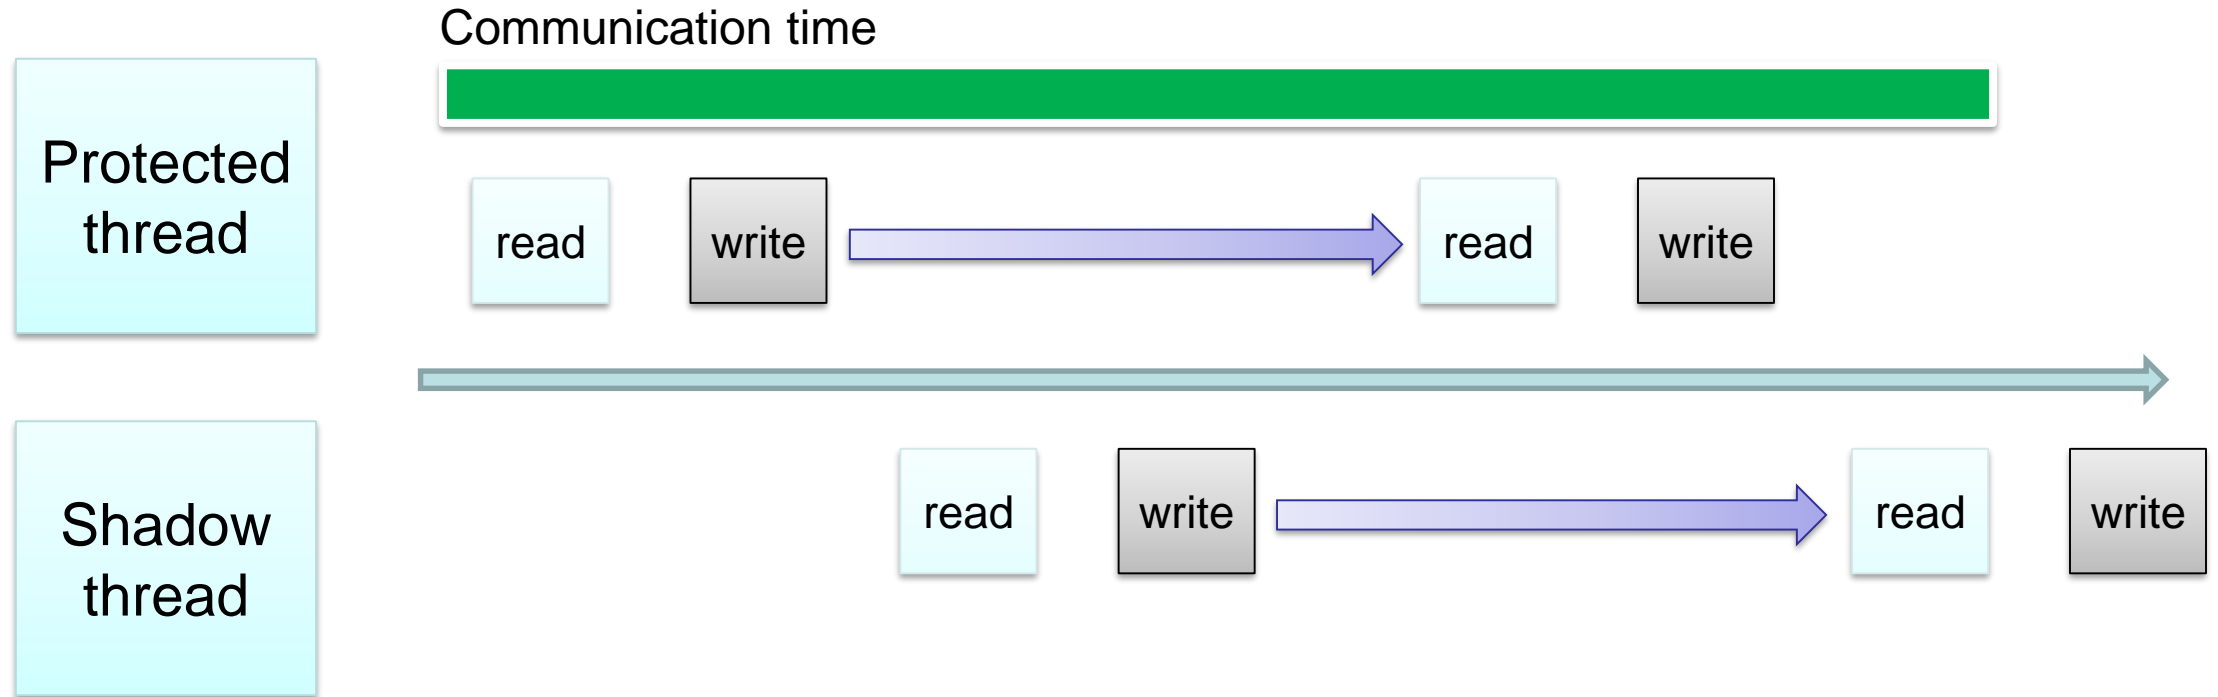
- ❑ Co-located: Both threads observe data races with high probability
- ❑ Otherwise: At least one observe data races with low probability

Closing HT-Side Channels on SGX with Contrived Data Races



- When co-located, communication time < execution time
- Each thread read the value written by the other thread with very **high** probability.

Closing HT-Side Channels on SGX with Contrived Data Races



- When **not** co-located, communication time $>$ execution time
- Each thread read the value written by the other thread with very **low** probability.

Closing HT-Side Channels on SGX with Contrived Data Races

Use of CMOV instructions

Thread T_0

```
1 <initialization>:
2   mov $colocation_count, %rdx
3   xor %rcx, %rcx
4   ; co-location test counter
5 <synchronization>:
6   ... ; acquire lock 0
7   .sync0:
8   mov %rdx, (sync_addr1)
9   cmp %rdx, (sync_addr0)
10  je .sync1
11  jmp .sync0
12  .sync1:
13  mfence
14  mov $0, (sync_addr0)
15 <initialize a round>:
16  mov $begin0, %rsi
17  mov $1, %rbx
18  mfence
19  mov $addr_v, %r8
20 <co-location test>:
21  .L0:
22  <load>:
23  mov (%r8), %rax
24  <store>:
25  mov %rsi, (%r8)
26  <update counter>:
27  mov $0, %r10
28  mov $0, %r11
29  cmp $end0, %rax
30  ; a data race happens?
```

```
31  cmovl %rbx, %r10
32  sub %rax, %r9
33  cmp $1, %r9
34  ; continuous number?
35  cmova %r11, %r10
36  add %r10, %rcx
37  shl $b_count, %rbx
38  ; bit length of $count
39  mov %rax, %r9
40  ; record the last number
41  <padding instructions 0>:
42  nop
43  nop
44  :
45  nop
46  mov (%r8), %rax
47  mov (%r8), %rax
48  :
49  mov (%r8), %rax
50  dec %rsi
51  cmp $end0, %rsi
52  jne .L0
53  ; finish 1 co-location test
54  <all rounds finished?>:
55  ... ; release lock 1
56  dec %rdx
57  cmp $0, %rdx
58  jne .sync0
```

Thread T_1

```
1 <initialization>:
2   mov $colocation_count, %rdx
3   xor %rcx, %rcx
4   ; co-location test counter
5 <synchronization>:
6   ... ; release lock 0
7   .sync2:
8   mov %rdx, (sync_addr0)
9   cmp %rdx, (sync_addr1)
10  je .sync3
11  jmp .sync2
12  .sync3:
13  mfence
14  mov $0, (sync_addr1)
15 <initialize a round>:
16  mov $begin1, %rsi
17  mov $1, %rbx
18  mfence
19  mov $addr_v, %r8
20 <co-location test>:
21  .L2:
22  <load>:
23  mov (%r8), %rax
24  <update counter>:
25  mov $0, %r10
26  mov $0, %r11
27  cmp $end0, %rax
28  ; a data race happens?
29  cmovg %rbx, %r10
30  sub %rax, %r9
```

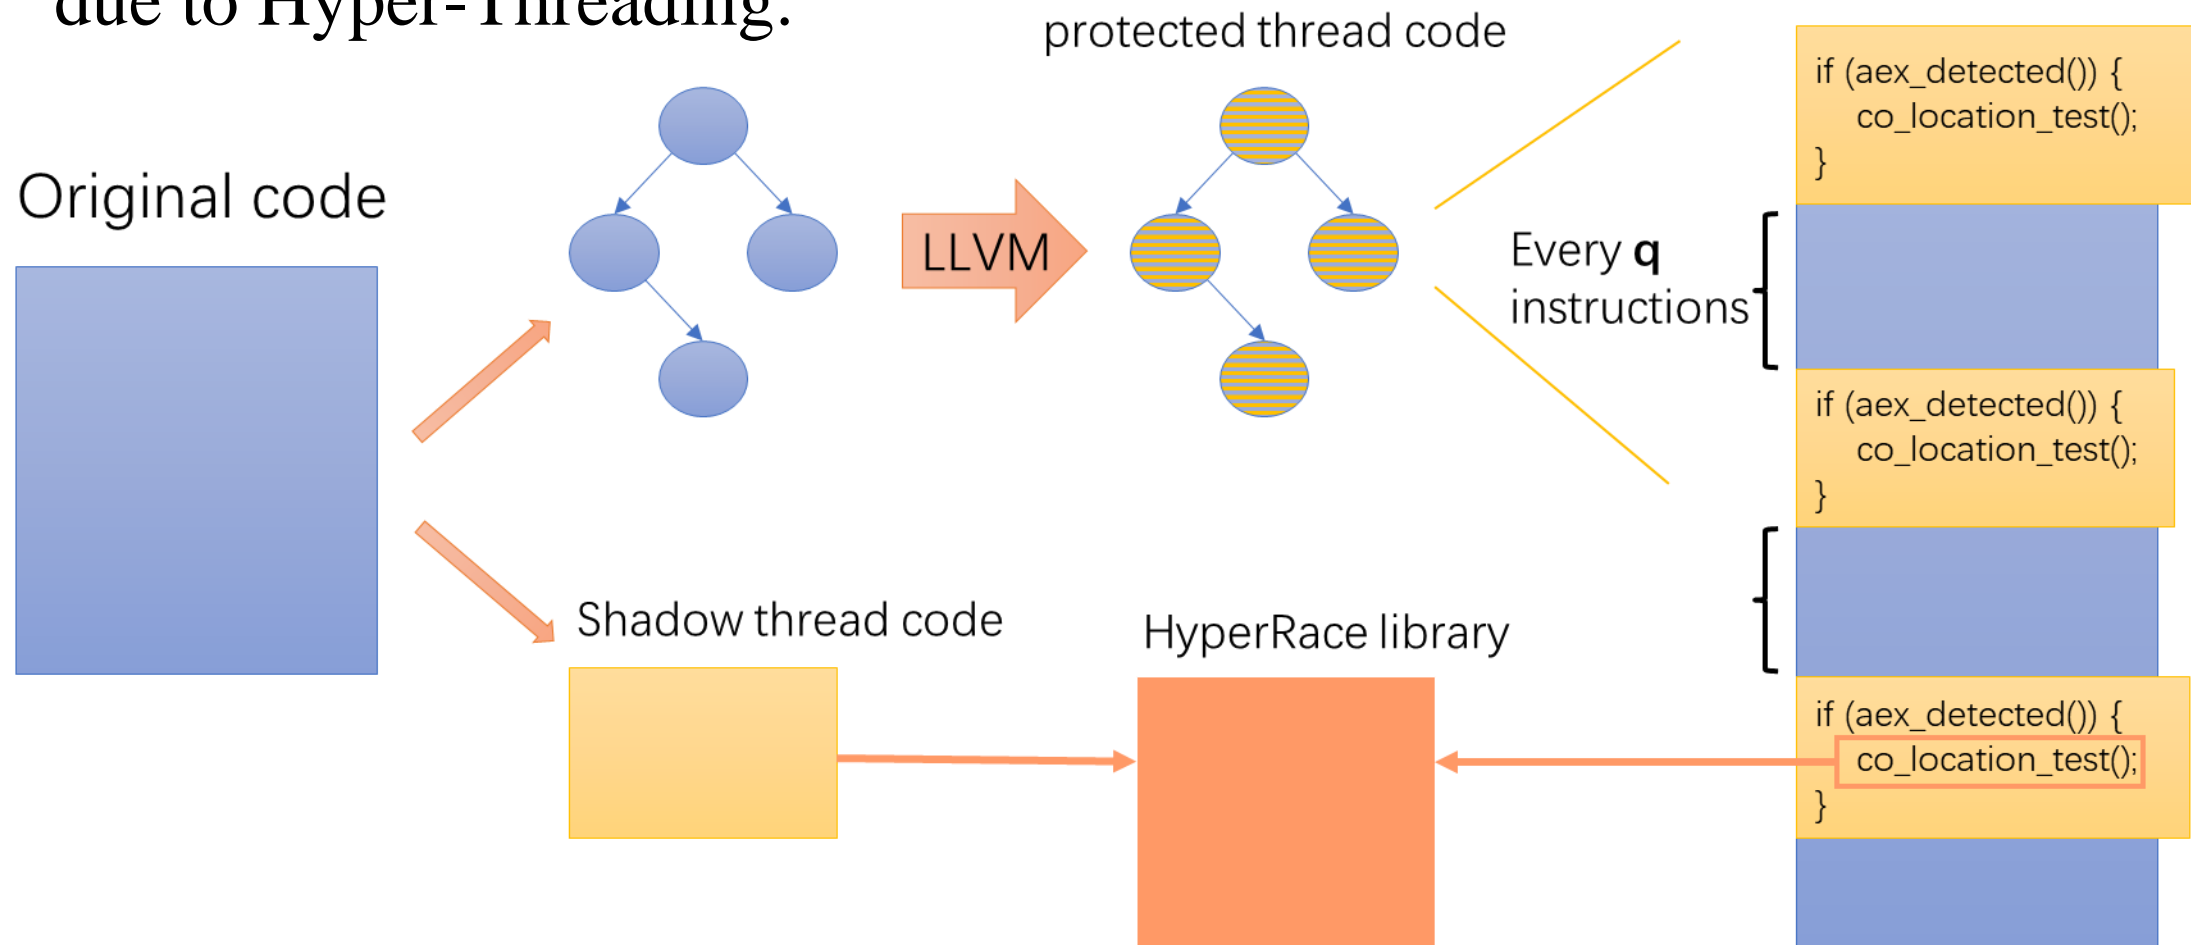
```
31  cmp $1, %r9
32  ; continuous number?
33  cmova %r11, %r10
34  add %r10, %rcx
35  shl $b_count, %rbx
36  ; bit length of $count
37  mov %rax, %r9
38  ; record the last number
39  <store>:
40  mov %rsi, (%r8)
41  <padding instructions 1>:
42  mov (%r8), %rax
43  lfence
44  mov (%r8), %rax
45  lfence
46  mov (%r8), %rax
47  lfence
48  mov (%r8), %rax
49  lfence
50  mov (%r8), %rax
51  lfence
52  dec %rsi
53  cmp $end1, %rsi
54  jne .L2
55  ; finish 1 co-location test
56  <all rounds finished?>:
57  ... ; acquire lock 1
58  dec %rdx
59  cmp $0, %rdx
60  jne .sync2
```

Hypothesis Test based security model

Different padding instruction patterns

Closing HT-Side Channels on SGX with Contrived Data Races

- HyperRace: An LLVM based tool to eradicate all side-channel threats due to Hyper-Threading.



Conclusion

- ❑ The SGX design opens up many side channels.
- ❑ These side channels can be combined
 - To make the attack stealthy and hard to detect
 - To achieve fine-grained observation
- ❑ The attacker can even reduce the noises by controlling the SW/HW environment.
- ❑ The side channel threats against SGX can not be ignored.

- ❑ How to design future TEEs?
 - HW/SW co-design?
 - Real world implications

Questions?



中国科学院 信息工程研究所
INSTITUTE OF INFORMATION ENGINEERING, CAS

Thank you