

Avoiding Leakage and Synchronization Attacks through Enclave-Side Preemption Control

Marcus Völp, Adam Lackorzynski*, Jérémie Decouchant, Vincent Rahli, Francisco Rocha, and Paulo Esteves-Veríssimo

University of Luxembourg Snt
CritiX Lab
Luxembourg
<name>.<surname>@uni.lu

*Kernkonzept GmbH and
TU Dresden – Operating-systems group
Dresden, Germany
adam.lackorzynski@kernkonzept.com

The functionality/code size dilemma

- application scenarios require the system to implement a certain set of functionalities
- implementing these functionalities comes at the cost of a certain minimal amount of code
 - even if development time and costs don't matter; and
 - even if you only use high-class developers
- correlation of code size and complexity to vulnerabilities
 - Chou et al., “*An Empirical Study of Operating Systems Errors*”, SOSP 2001
 - Asadollah et al., “*A Study of Concurrency Bugs in an Open Source Software*”, OSS 2016

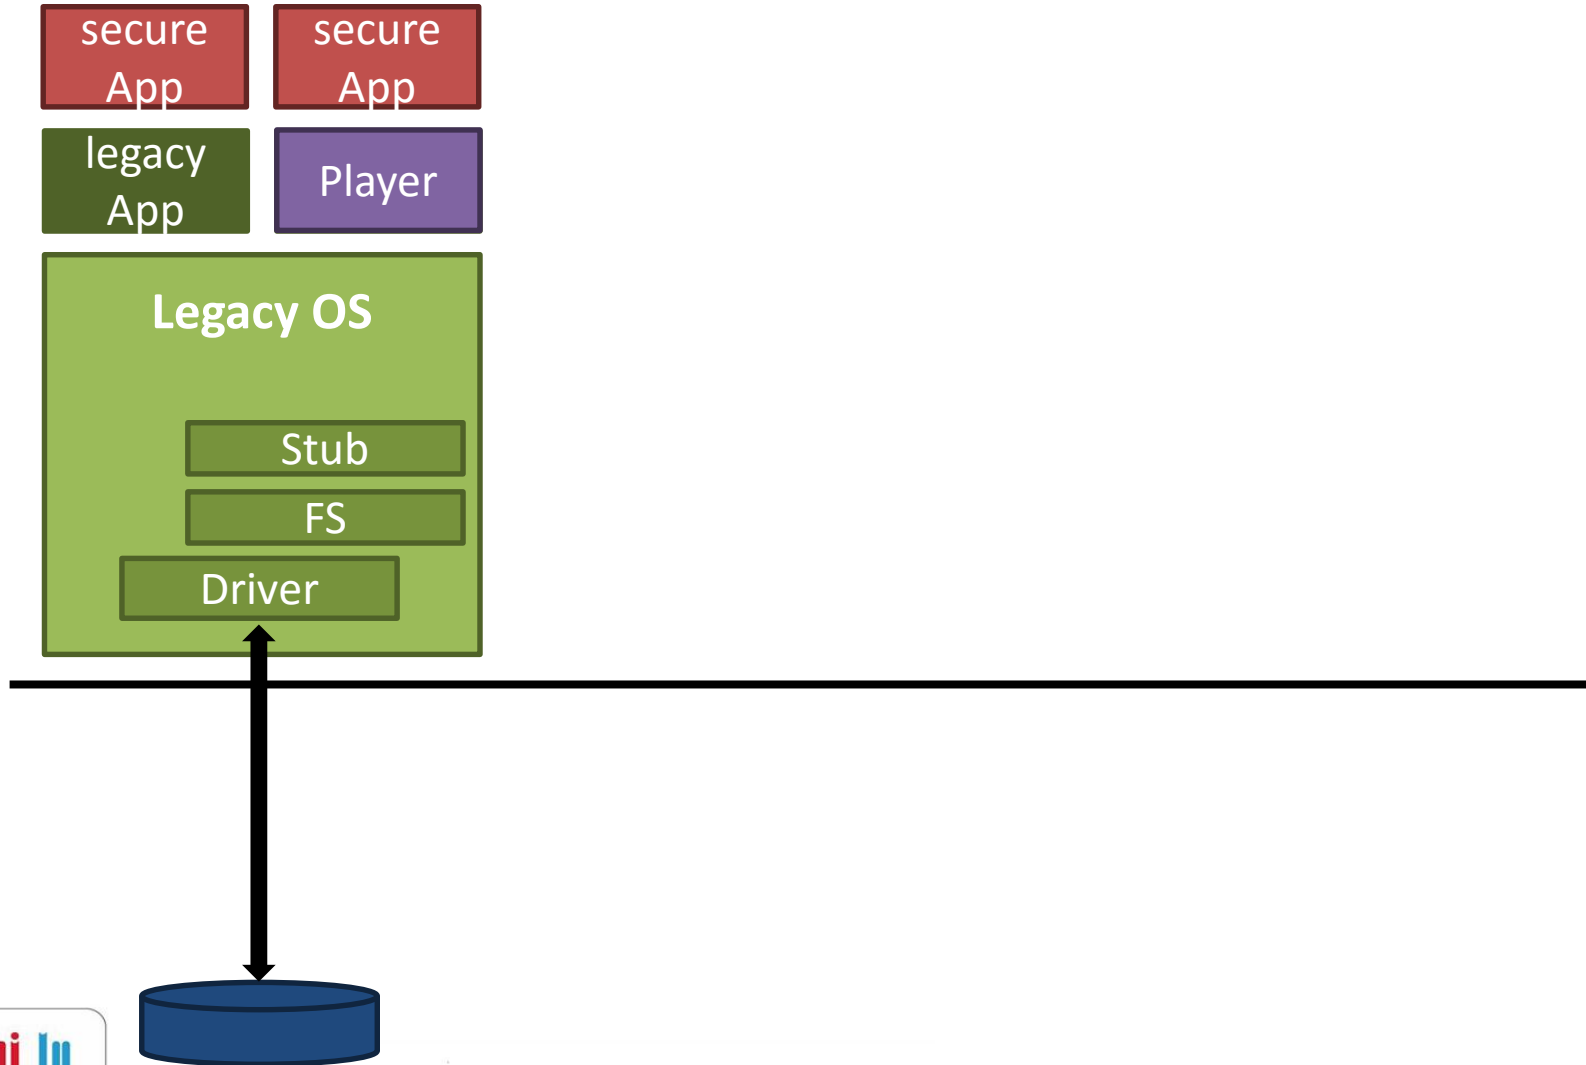
The functionality/code size dilemma

- application scenarios require the system to implement a certain set of functionalities
- implementing these functionalities comes at the cost of a certain minimal amount of code
 - even if d
 - even if y
- correlation

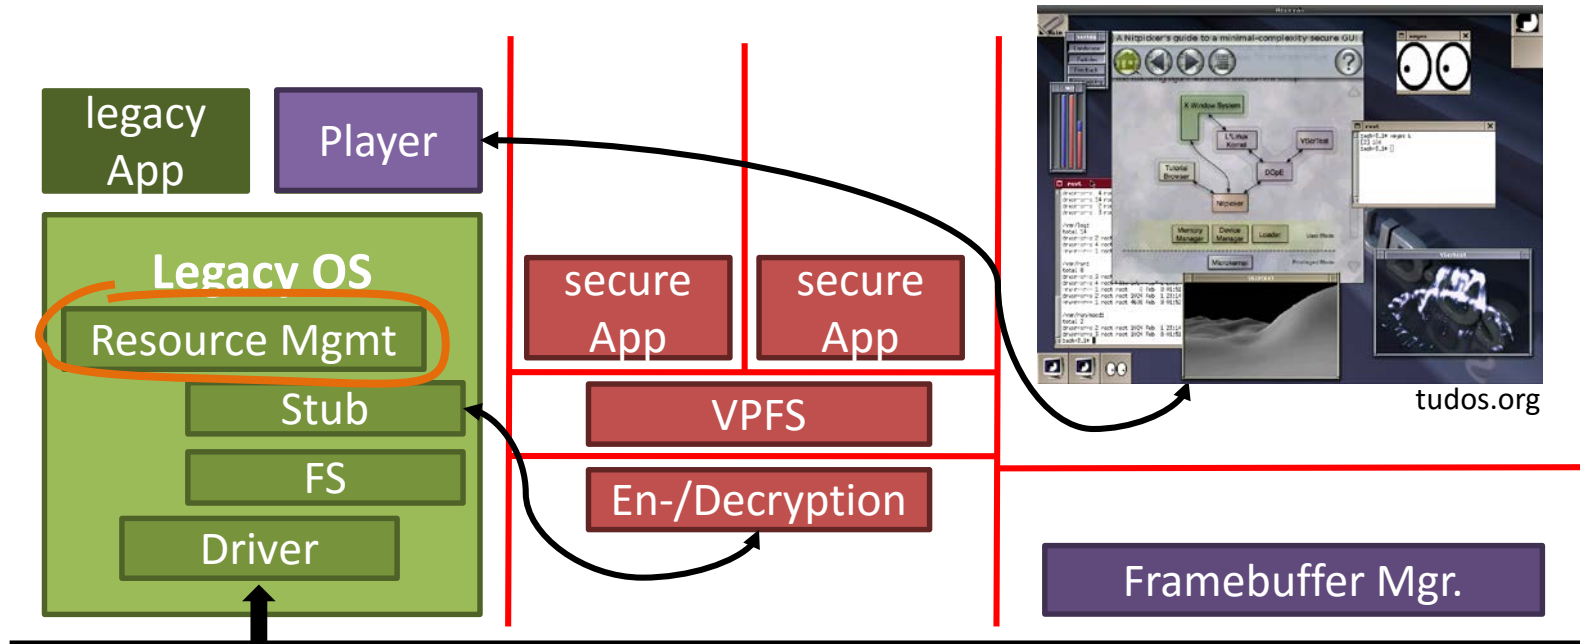
• RTOS	ca. 5 KLOC	} a 5-13 PY formal verification
• Microkernel	10 – 15 KLOC	
• Legacy OS	15 – 50 MLOC	

 - Chou et al., “An Empirical Study of Operating Systems Errors”, SOSP 2001
 - Asadollah et al., “A Study of Concurrency Bugs in an Open Source Software”, OSS 2016

Intransitive trust



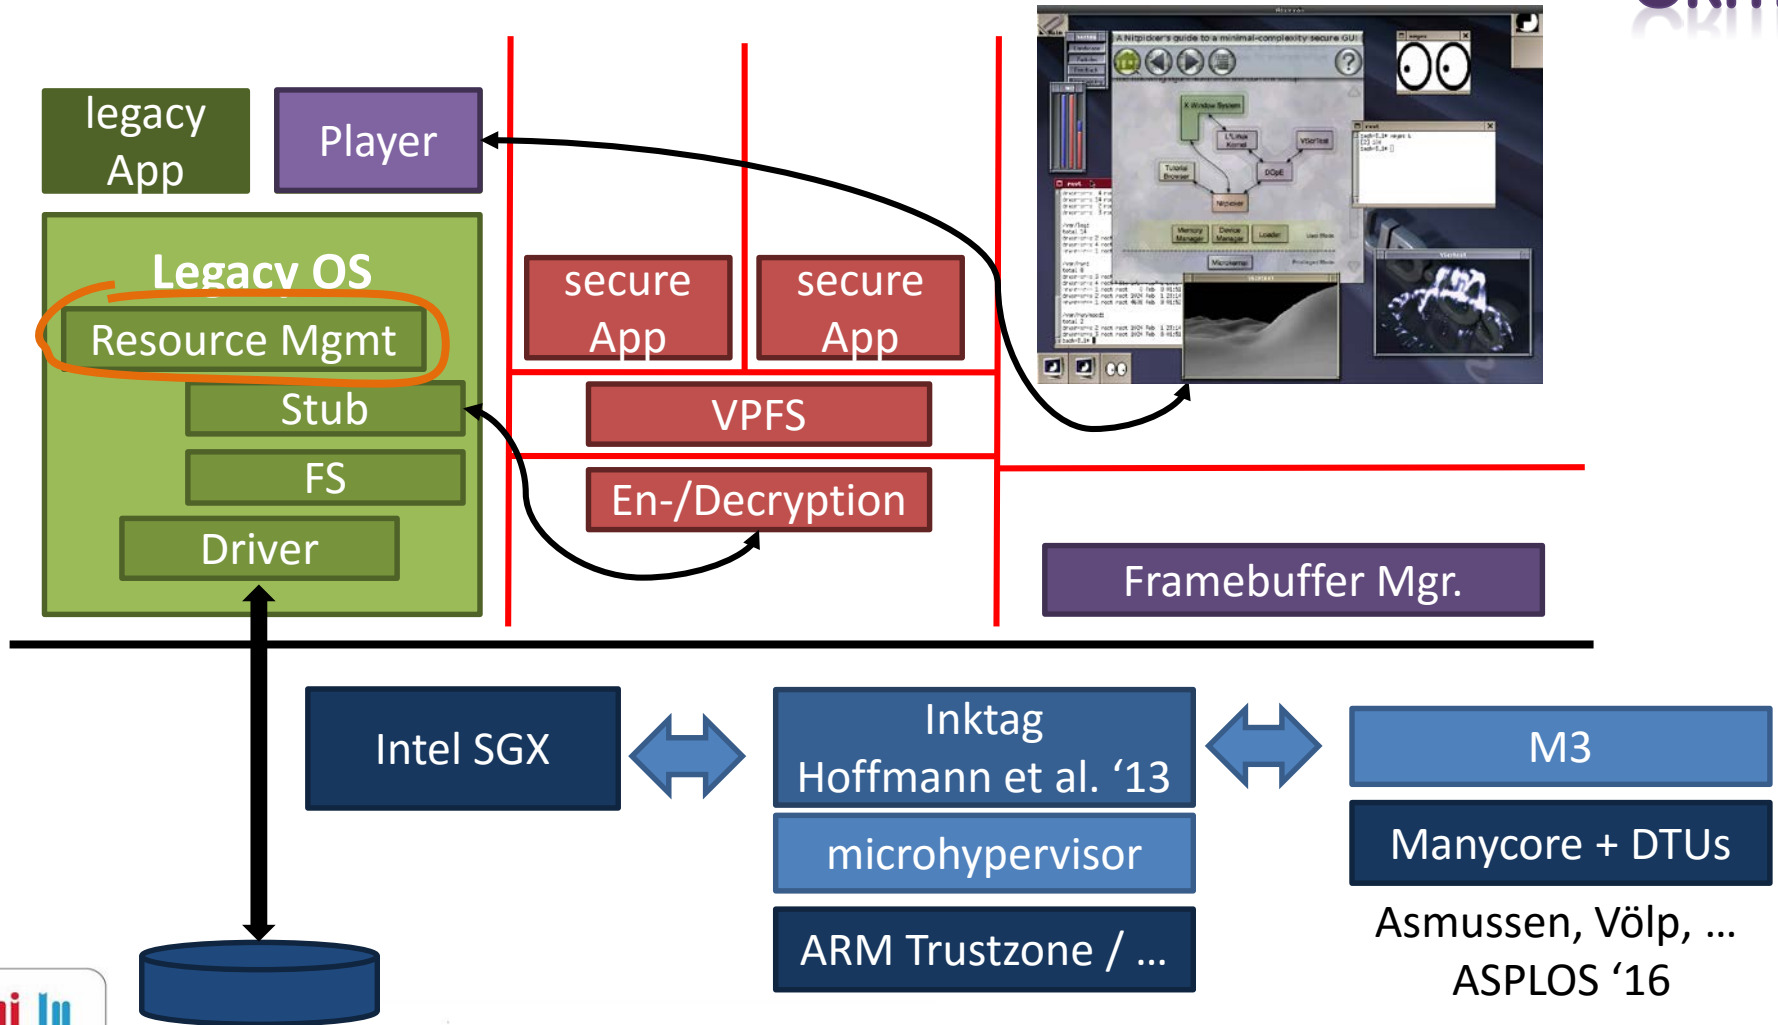
Intransitive trust



- Weinhold et al., “jVPFS: Adding Robustness to a Secure Stacked File System with Untrusted Local Storage Components”, USENIX ATC, 2011
- Singaravelu et al., “Reducing TCB Complexity for Security-Sensitive Applications: Three Case Studies”, Eurosys, 2006
- ...

Asmussen, Völp, ...
ASPLOS '16

Intransitive trust



SGX Vulnerabilities

AsyncShock: Exploiting Synchronisation Bugs in Intel SGX Enclaves

Nico Weichbrodt, Anil Kurmust, Peter Pietzuch*, and Rüdiger Kapitza
 TU Braunschweig, †IBM Research Zurich, *Imperial College London
 {weichbr,kapitza}@ibr.cs.tu-bs.de, kur@zurich.ibm.com, prp@imperial.ac.uk

Abstract. Intel's Software Guard Extensions (SGX) provide a new hardware-based trusted execution environment on Intel CPUs using *secure enclaves* that are resilient to accesses by privileged code and physical attackers. Original promise to protect sensitive data in the cloud. By exploiting synchronization bugs in enclave access control, we present a new attack. We present a multithreaded application that only manipulates enclave code. The application is implemented on Skylake CPUs and TOCTOU.

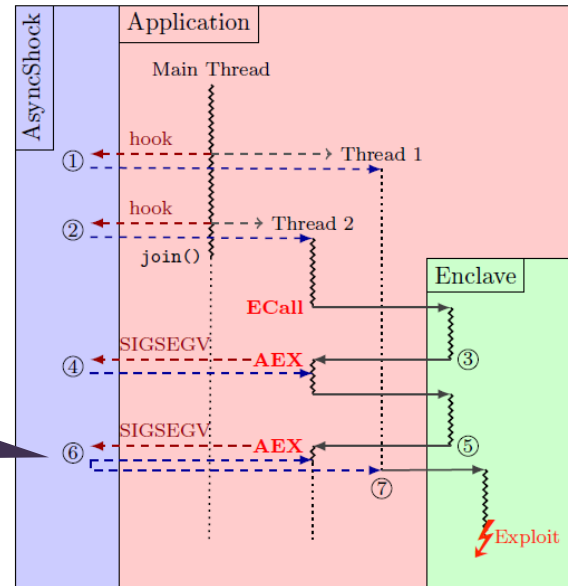
Keyword: synchronization

```

1 char *glob_str_ptr;
2
3 int other_functions(const char *c) { /* do other things */ }
4
5 int puts(const char * c) {
6     printf("%s", c);
7     return 0;
8 }
9
10 struct my_func_ptr {
11     int (*my_puts) (const char *);
12     char desc[8];
13 } my_func_ptr;
14
15 void ecall(const char *c) {
16     glob_str_ptr = c;
17 }
18
19 void ecall(const char *c) {
20     struct my_func_ptr mfp;
21     mfp.my_puts = puts;
22     if (glob_str_ptr != NULL)
23         mfp.desc = glob_str_ptr;
24     glob_str_ptr = NULL;
25     free(mfp);
26     free(mfp);
27     glob_str_ptr = NULL;
28 }
29 free(mfp);
30 }
    
```

Fine grain preemption control to widen the window of vulnerability of synchronization bugs

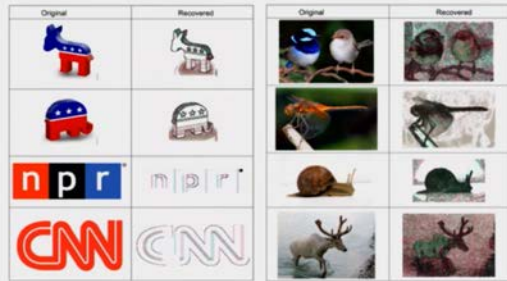
Source: AsyncShock



SGX Vulnerabilities

36th IEEE Symposium on Security and Privacy

libjpeg results



Controlled-Channel Attacks: Deterministic Side Channels for Untrusted Operating Systems

Yuanzhong Xu
The University of Texas at Austin
yxu@cs.utexas.edu

Weidong Cui
Microsoft Research
wdcui@microsoft.com

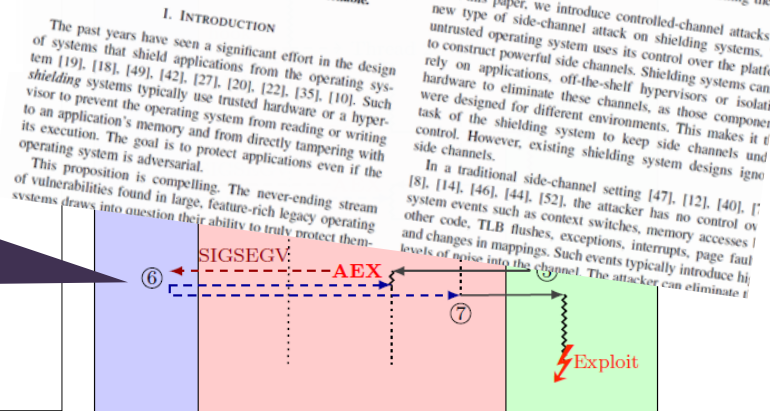
Marcus Peinado
Microsoft Research
marcus@peinado.com

Fine grain preemption control to widen the window of vulnerability for side-channel attacks

```

13 } my_func_ptr,
14
15 void ecal
16     glob_s
17 }
18
19 void ecal
20     struct
21     mfp->p
22     if (g
23     me
24     glo
25     mfp
26     fre
27     glob_str_ptr = NULL;
28 }
29 free(mfp);
30 }
    
```

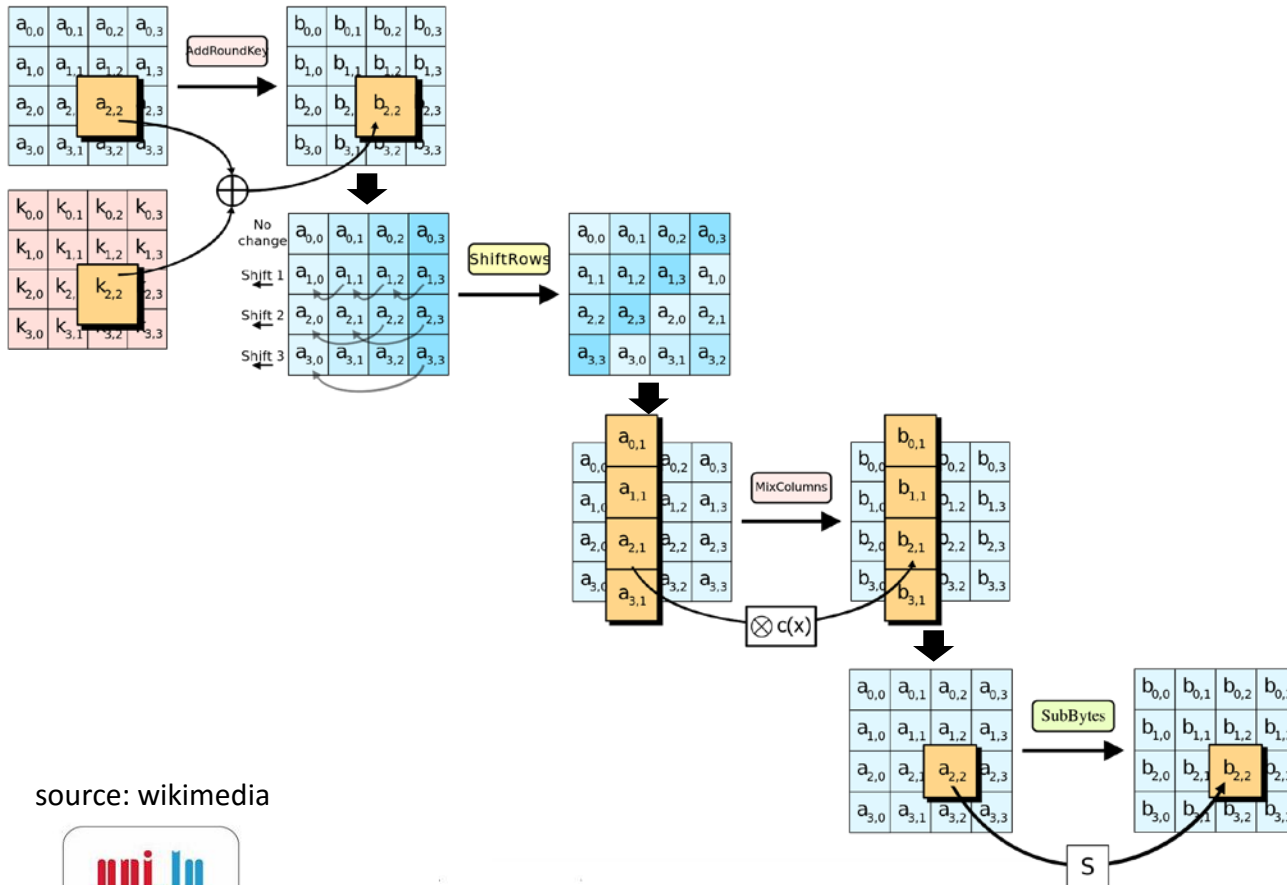
Fine grain preemption control to widen the window of vulnerability of synchronization bugs



SGX Vulnerabilities

- Running Example:

Osvik et al., "Cache Attacks and Countermeasures: the Case of AES", CT-RSA 2006



in-memory tables T_i

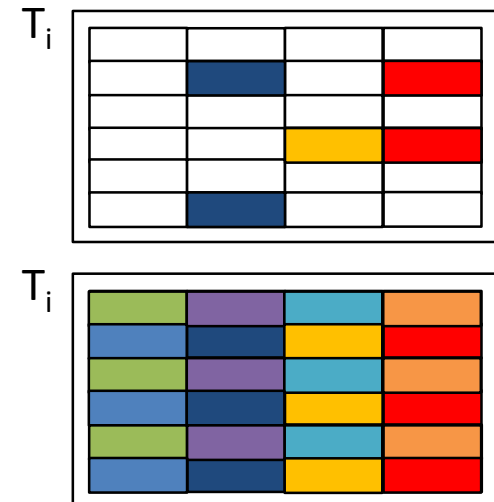
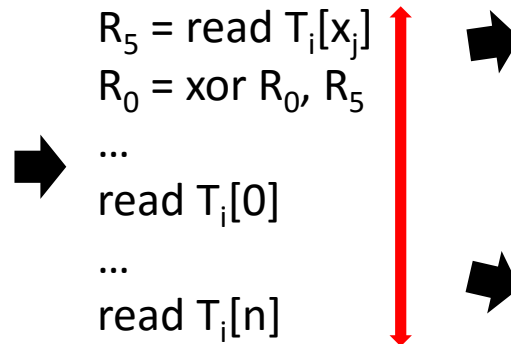
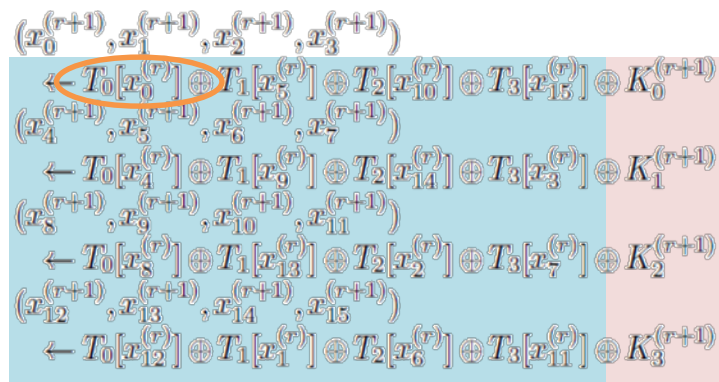
source: wikimedia



SGX Vulnerabilities

- Running Example:

Osvik et al., "Cache Attacks and Countermeasures: the Case of AES", CT-RSA 2006



SGX Vulnerabilities

- Running Example:

Osvik et al., "Cache Attacks and Countermeasures: the Case of AES", CT-RSA 2006

```

(x0(r+1), x1(r+1), x2(r+1), x3(r+1))
← T0[x0(r)] ⊕ T1[x5(r)] ⊕ T2[x10(r)] ⊕ T3[x15(r)] ⊕ K0(r+1)
(x4(r+1), x5(r+1), x6(r+1), x7(r+1))
← T0[x4(r)] ⊕ T1[x9(r)] ⊕ T2[x14(r)] ⊕ T3[x3(r)] ⊕ K1(r+1)
(x8(r+1), x9(r+1), x10(r+1), x11(r+1))
← T0[x8(r)] ⊕ T1[x13(r)] ⊕ T2[x2(r)] ⊕ T3[x7(r)] ⊕ K2(r+1)
(x12(r+1), x13(r+1), x14(r+1), x15(r+1))
← T0[x12(r)] ⊕ T1[x1(r)] ⊕ T2[x6(r)] ⊕ T3[x11(r)] ⊕ K3(r+1)
    
```



```

R6 = read Ti[0]
cmp 0, xj
R5 = cmov R6
R0 = xor R0, R5
...
    
```



T_i

Green	Purple	Light Blue	Orange
Blue	Dark Blue	Yellow	Red
Green	Purple	Light Blue	Orange
Blue	Dark Blue	Yellow	Red



T_i

Green	Purple	Light Blue	Orange
Blue	Dark Blue	Yellow	Red
Green	Purple	Light Blue	Orange
Blue	Dark Blue	Yellow	Red

low indistinguishable data access pattern embedded into low indistinguishable control flow

SGX Vulnerabilities

- Running Example:

Osvik et al., "Cache Attacks and Countermeasures: the Case of AES", CT-RSA 2006

```

(x0(r+1), x1(r+1), x2(r+1), x3(r+1))
← T0[x0(r)] ⊕ T1[x5(r)] ⊕ T2[x10(r)] ⊕ T3[x15(r)] ⊕ K0(r+1)
(x4(r+1), x5(r+1), x6(r+1), x7(r+1))
← T0[x4(r)] ⊕ T1[x9(r)] ⊕ T2[x14(r)] ⊕ T3[x3(r)] ⊕ K1(r+1)
(x8(r+1), x9(r+1), x10(r+1), x11(r+1))
← T0[x8(r)] ⊕ T1[x13(r)] ⊕ T2[x2(r)] ⊕ T3[x7(r)] ⊕ K2(r+1)
(x12(r+1), x13(r+1), x14(r+1), x15(r+1))
← T0[x12(r)] ⊕ T1[x1(r)] ⊕ T2[x6(r)] ⊕ T3[x11(r)] ⊕ K3(r+1)
    
```

disable preemptions

R₅ = read T_i[x_j]

R₀ = xor R₀, R₅


...

read T_i[0]

...

read T_i[n]

enable preemptions



This talk

Re-investigate delayed-preemption:

- How can we allow user-level applications (in enclaves) to disable preemptions without being able to monopolizing the system?
- How can we prevent solicited exits through which the management OS could regain control?
- How can we translate delayed-preemption to Intel SGX?

Towards Scalable Multiprocessor Virtual Machines

Volkmar Uhlig Joshua LeVasseur Espen Skoglund Uwe Dannowski
System Architecture Group
Universität Karlsruhe
contact@14ka.org

Abstract

A multiprocessor virtual machine benefits its guest operating system in supporting scalable job throughput and request latency—useful properties in server consolidation where servers require several of the system processors for steady state or to handle load bursts.

Typical operating systems, optimized for multiprocessor systems in their use of spin-locks for critical sections, can defeat flexible virtual machine scheduling due to lock-holder preemption and misbalanced load. The virtual machine must assist the guest operating system to avoid lock-holder preemption and to schedule jobs with knowledge of asymmetric processor allocation. We want to support a virtual machine environment with flexible scheduling policies, while maximizing guest performance.

This paper presents solutions to avoid lock-holder preemption for both fully virtualized and paravirtualized environments. Experiments show that we can nearly eliminate the effects of lock-holder preemption. Furthermore, the paper presents a scheduler feedback mechanism that despite the presence of asymmetric processor allocation achieves optimal and fair load balancing in the guest operating system.

1 Introduction

A recent trend in server consolidation has been to provide virtual machines that can be safely multiplexed on

of guests, such that they only ever access a fraction of the physical processors, or alternatively time-multiplex guests across a set of physical processors to, e.g., accommodate for spikes in guest OS workloads. It can also map guest operating systems to virtual processors (which can exceed the number of physical processors), and migrate between physical processors without notifying the guest operating systems. This allows for, e.g., migration to other machine configurations or hot-swapping of CPUs without adequate support from the guest operating system. It is important to recognize that guest operating system offers much more flexibility than schemes where one can only configure a virtual machine to either have an arbitrary share of a single processor [7, 24], or have uniform shares over multiple physical processors [10, 24].

Isolating commodity operating systems within virtual machines can defeat the assumptions of the guest operating system. Where the guest operating system expects constant resource configurations, critical timing behavior, and unrestrained access to the platform, the virtual machine provides illusionary access as it sees fit. Several methods exist to attempt to satisfy (a subset of) the assumptions of the guest operating system. The solutions may focus on the issues of instruction set emulation, such as trapping on system instructions [22], or they may focus on the behavior of the guest operating system algorithms, such as dynamic allocation of physical memory [25].

This paper presents solutions to two problems that arise with scheduling of virtual machines which provide

This talk

disable preemptions

$R_5 = \text{read } T_i[x_j]$

$R_0 = \text{xor } R_0, R_5$

...

read $T_i[0]$

...

read $T_i[n]$

enable preemptions



← How can we prevent solicited exits in sensitive code?



How can we make sure the enclave enables preemptions again?

This talk

disable preemptions ←

prepare ←

if preempted

goto retry

$R_5 = \text{read } T_i[x_j]$

$R_0 = \text{xor } R_0, R_5$

...

read $T_i[0]$

...

read $T_i[n]$

enable preemptions ←

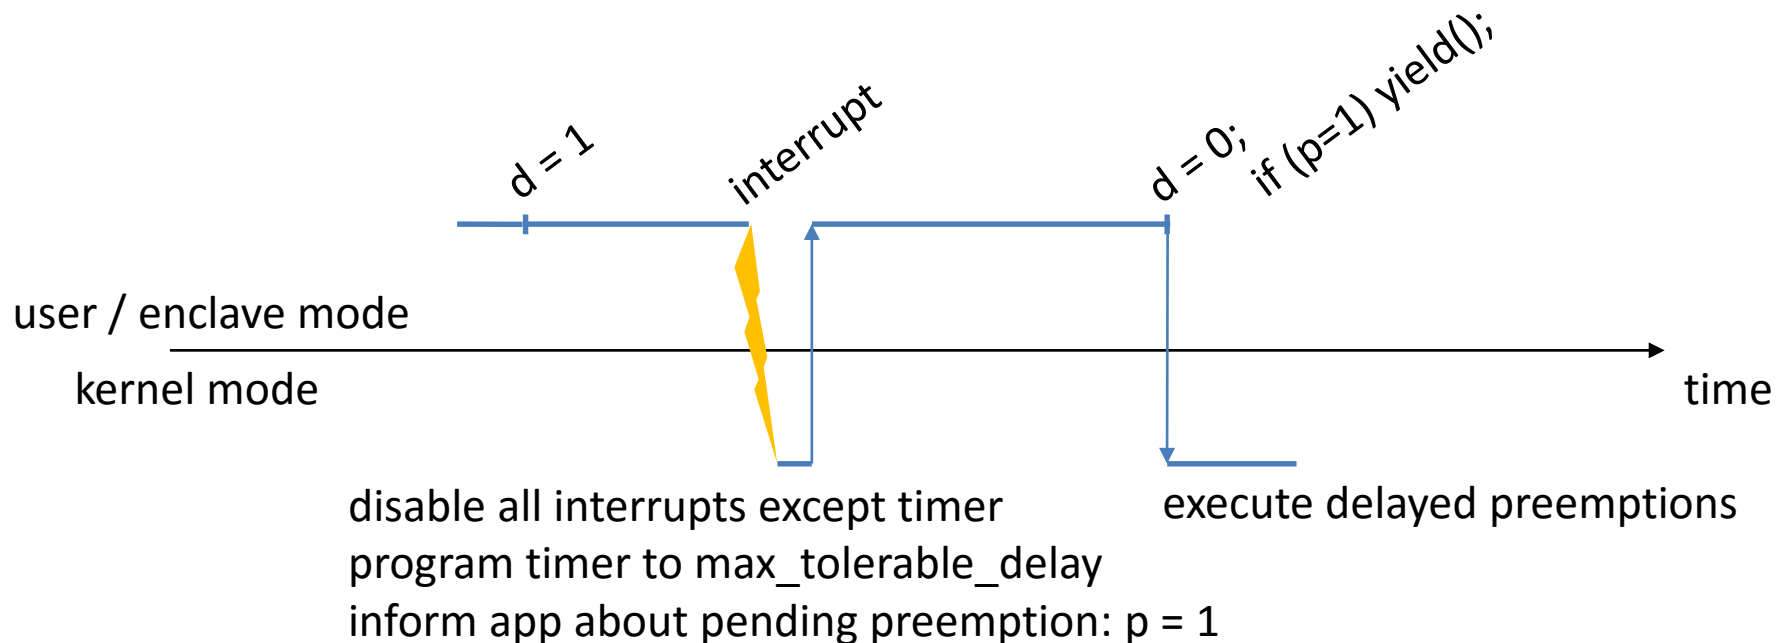


How can we prevent solicited exits in sensitive code?

How can we make sure the enclave enables preemptions again?

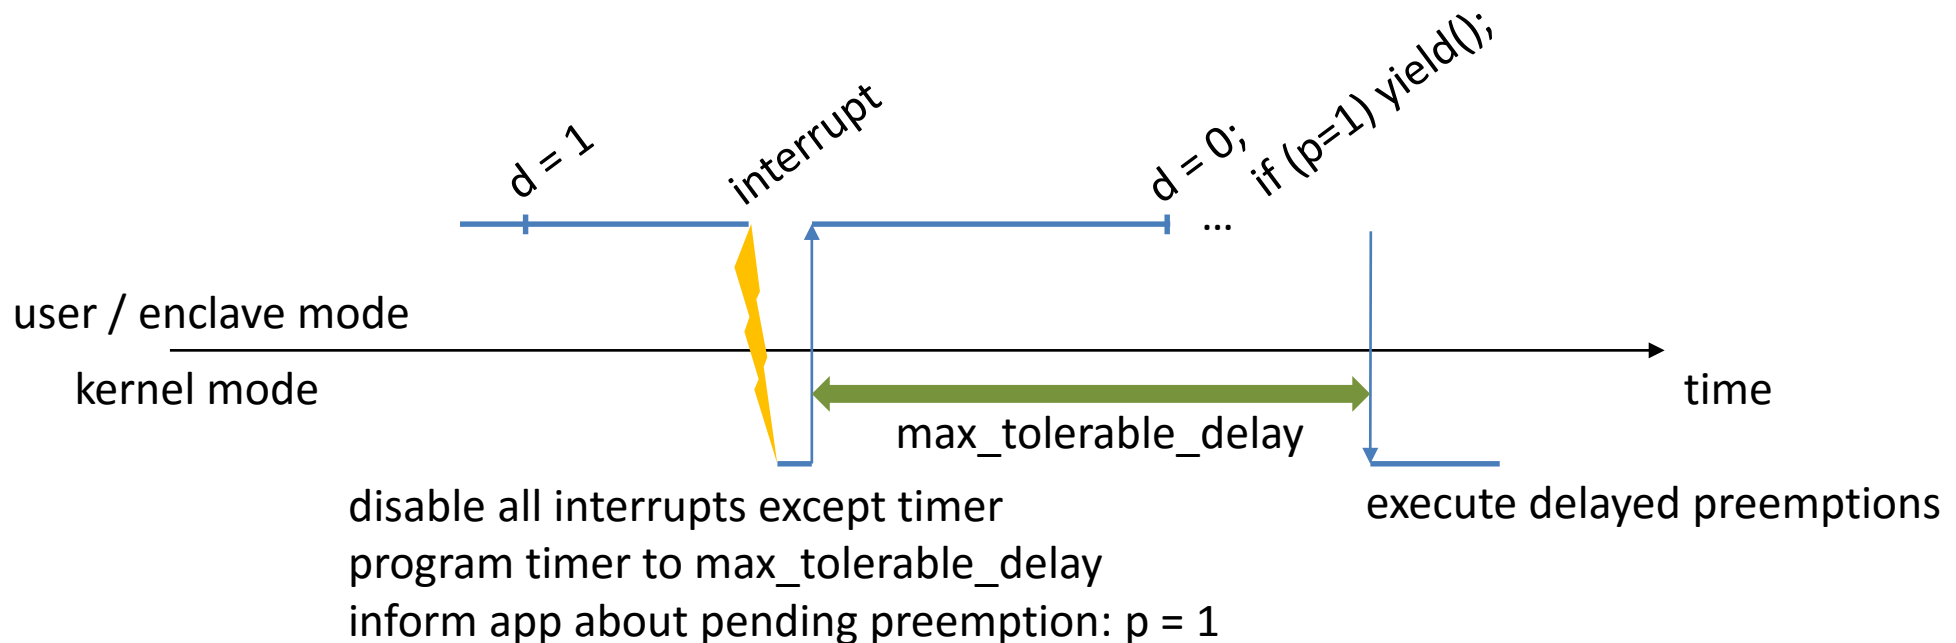
Delayed Preemption

- ... in a Trusted-Trustworthy Hypervisor




Delayed Preemption

- ... in a Trusted-Trustworthy Hypervisor



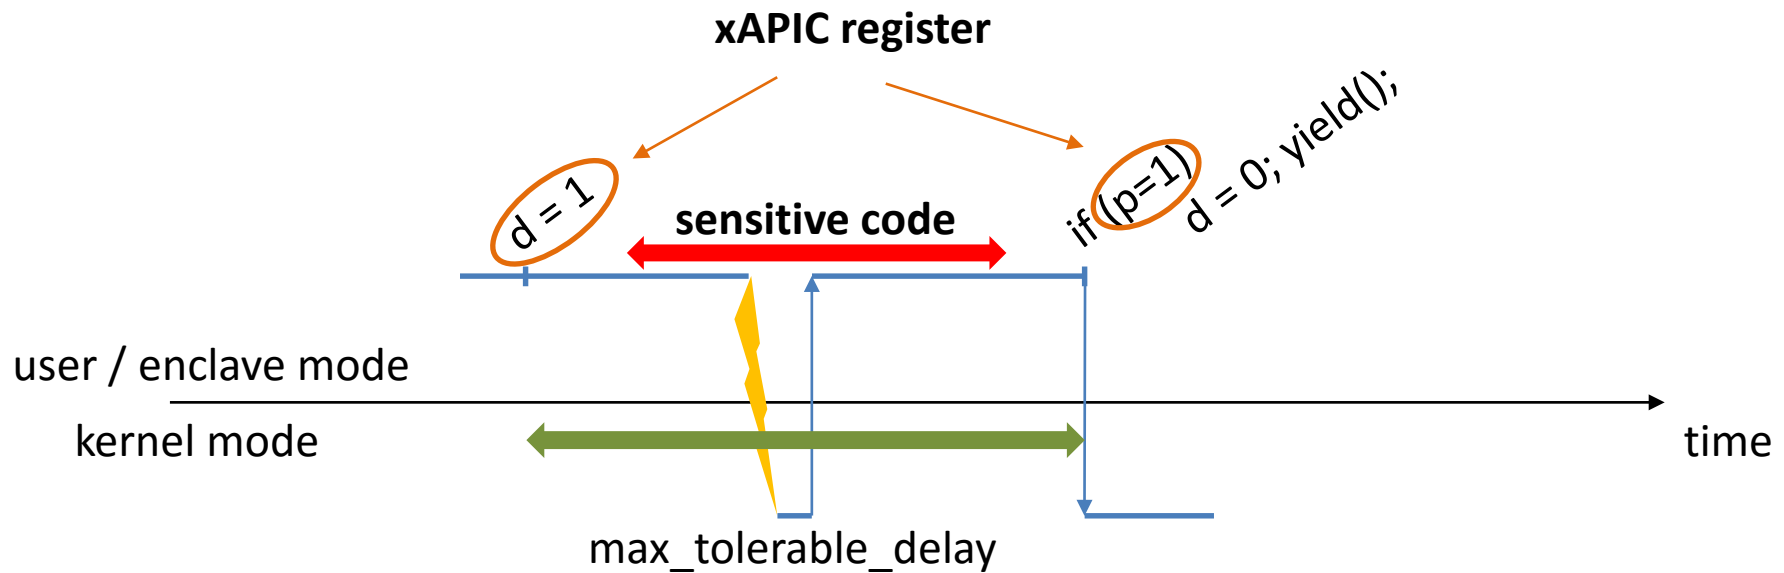
Delayed Preemption

- ... in a Trusted-Trustworthy Hypervisor



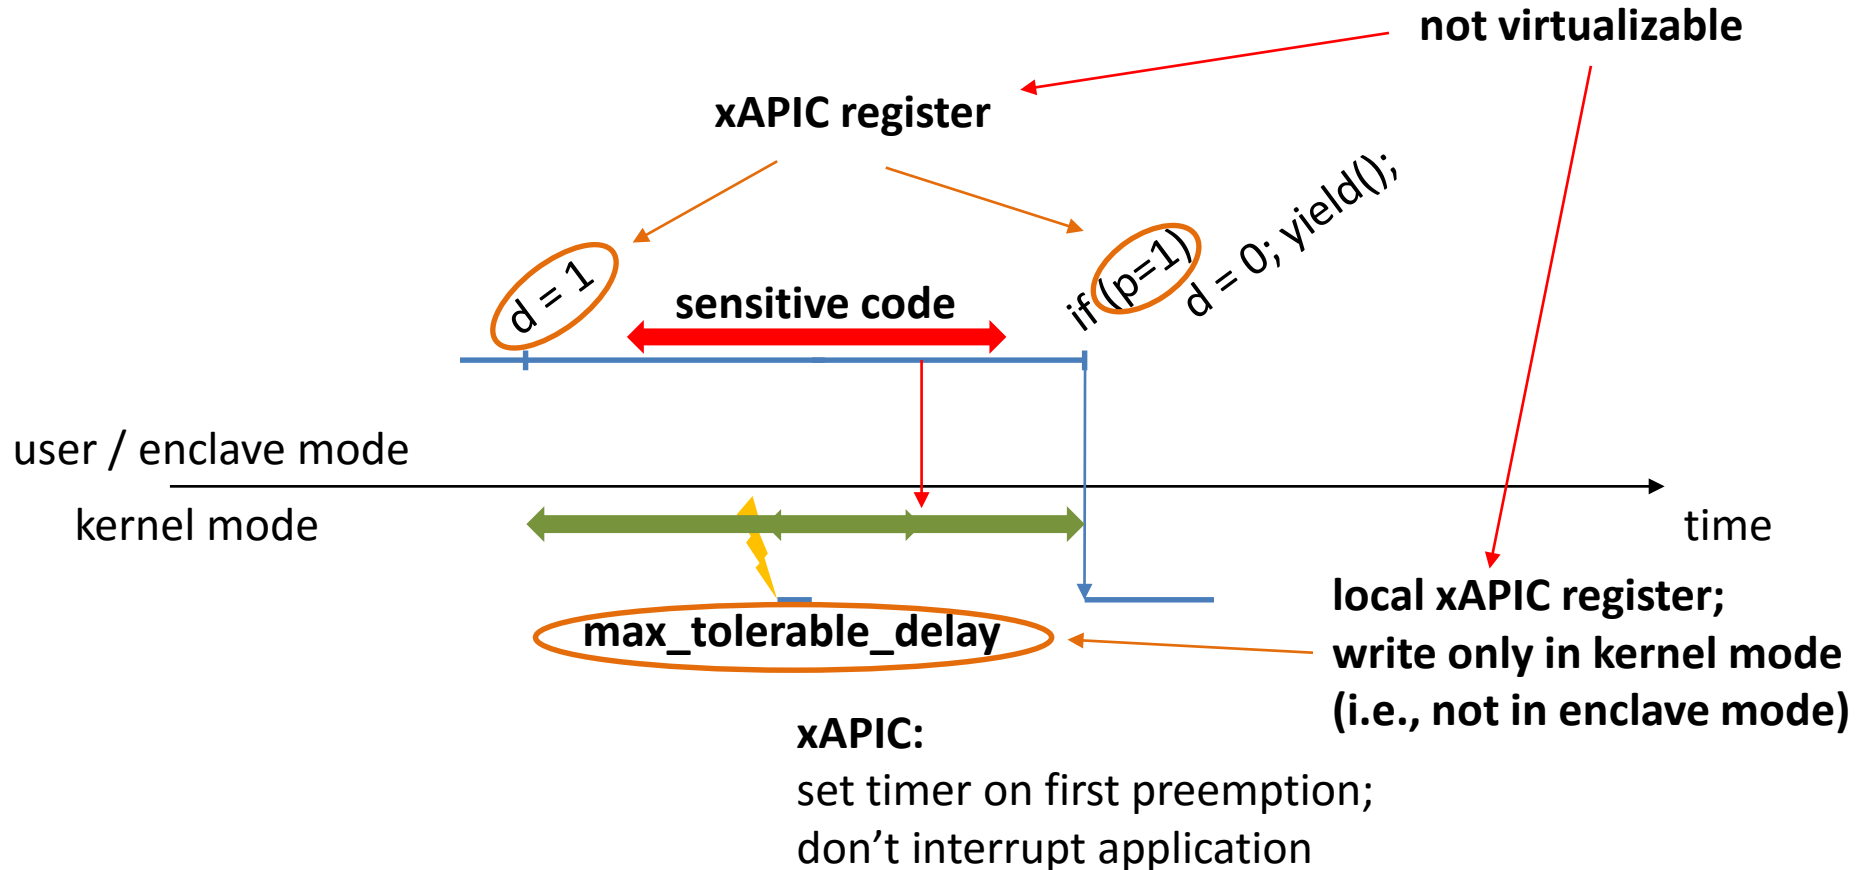
Delayed Preemption

- ... in SGX



Delayed Preemption

- ... in SGX



Solicited Exits

disable preemptions

prepare

if preempted

goto retry

$R_5 = \text{read } T_i[x_j]$

$R_0 = \text{xor } R_0, R_5$

...

read $T_i[0]$

...

read $T_i[n]$

enable preemptions



Solicited Exits


```
retry:
xApic.d = 1;
-----
prepare
if (p = 1)
    goto retry
R5 = read Ti[xj]
R0 = xor R0, R5
...
read Ti[0]
...
read Ti[n]
-----
xApic.d = 0
// if (xApic.p = 1) -> AEX ✓ max_tolerable_delay
```

✓
Trigger all such exits during non-sensitive prepare phase;
Set p flag to make code aware of these exits;
Context switch p flag as part of enclave state

How to prevent solicited exits in sensitive code?
• data / instruction page-faults
• lazy FPU context switch
• power management
• device virtualization

Solicited Exits

- data / instruction TLB



retry:

```
d = 1;  
//prepare  
call pg(f)  
call pg(g)  
or $0, [pg(T0)]  
or $0, [pg(T1)]  
or $0, [pg(T2)]  
or $0, [pg(T3)]  
if (p = 1)
```

goto retry

← set p-flag on instruction / data pagefault

Recall: cross-CPU page-table changes require IPIs to shutdown TLBs

Solicited Exits

```

retry:
xApic.d = 1;
-----
prepare
if (p = 1)
    goto retry
R5 = read Ti[xj]
R0 = xor R0, R5
...
read Ti[0]
...
read Ti[n]
-----
xApic.d = 0
// if (xApic.p = 1) -> AEX
    
```

max_tolerable_delay



How to prevent solicited exits in sensitive code?

- data / instruction page-faults
- lazy FPU context switch
- power management
- device virtualization

} access pages / FPU;
 } check p-flag
 } report power state
 } access device MMIO / ports;
 } check p-flag

check
 max_tolerable_delay >
 WCET(prepare + sensitive)
 of current power state

Concurrency Bugs

- Cannot fix concurrency bugs by delaying preemptions
- Avoid widening the window of vulnerability

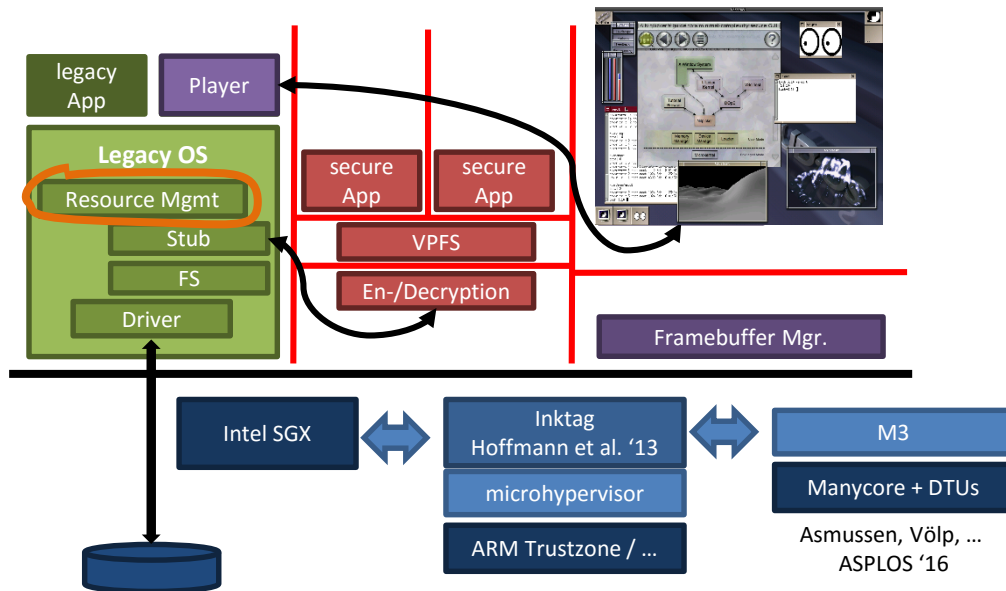
disable preemptions
free object
invalidate pointer
enable preemptions

||

disable preemptions
if (pointer)
 use object
enable preemptions

This talk in one slide...

intransitive trust: enabler for TCB reduction



CritiX Lab (Critical and Extreme Security and Dependability)

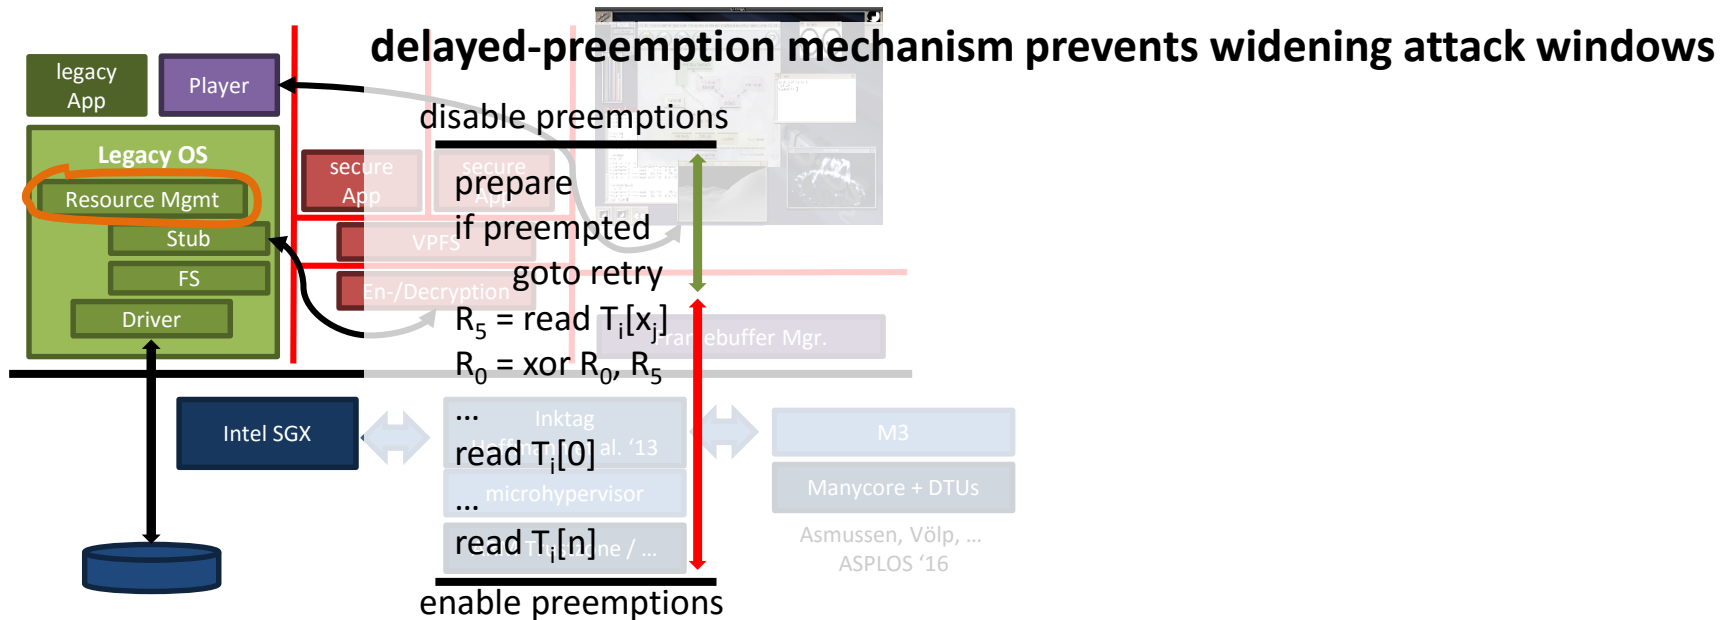
Interdisciplinary Centre for Security, Reliability and Trust - University of Luxembourg

PEARL Grant FNR/P14/8149128 – Paulo Esteves-Veríssimo

We are hiring bright post-docs and research associates!

This talk in one slide...

intransitive trust: enabler for TCB reduction



CritiX Lab (Critical and Extreme Security and Dependability)

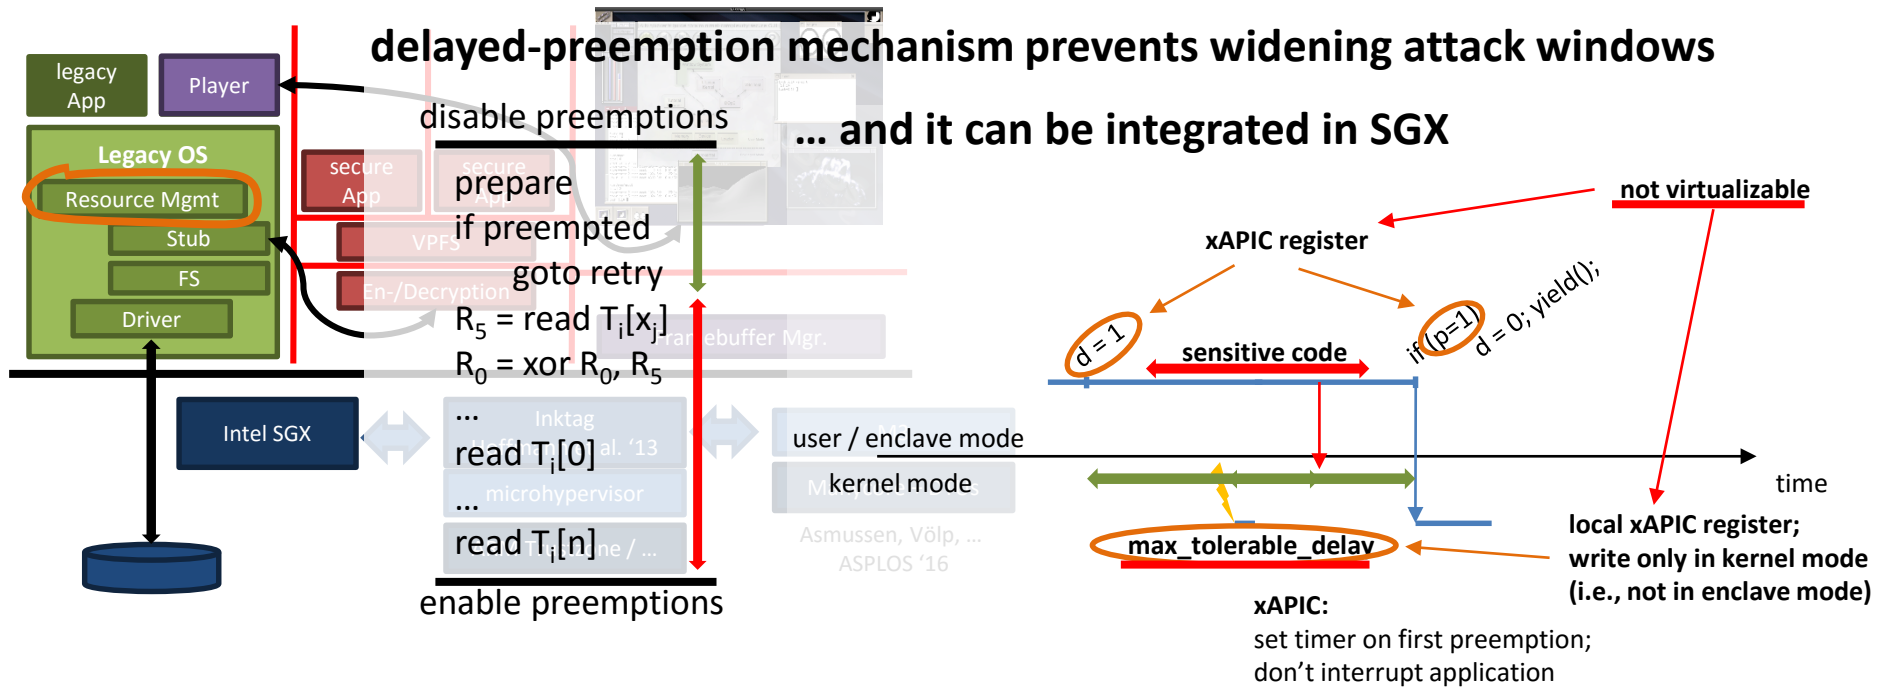
Interdisciplinary Centre for Security, Reliability and Trust - University of Luxembourg

PEARL Grant FNR/P14/8149128 – Paulo Esteves-Veríssimo

We are hiring bright post-docs and research associates!

This talk in one slide...

intransitive trust: enabler for TCB reduction



CritiX Lab (Critical and Extreme Security and Dependability)

Interdisciplinary Centre for Security, Reliability and Trust - University of Luxembourg

PEARL Grant FNR/P14/8149128 – Paulo Esteves-Veríssimo

We are hiring bright post-docs and research associates!