DESIGN, AUTOMATION & TEST IN EUROPE
01 – 05 February 2021 · virtual conference
The European Event for Electronic System Design & Test

Microarchitectural Timing Channels and their Prevention on an Open-Source 64-bit RISC-V Core

Nils Wistoff ETH Zurich
Moritz Schneider ETH Zurich
Frank K. Gürkaynak ETH Zurich
Luca Benini ETH Zurich and University of Bologna
Gernot Heiser UNSW Sydney and Data61 CSIRO

Security Model

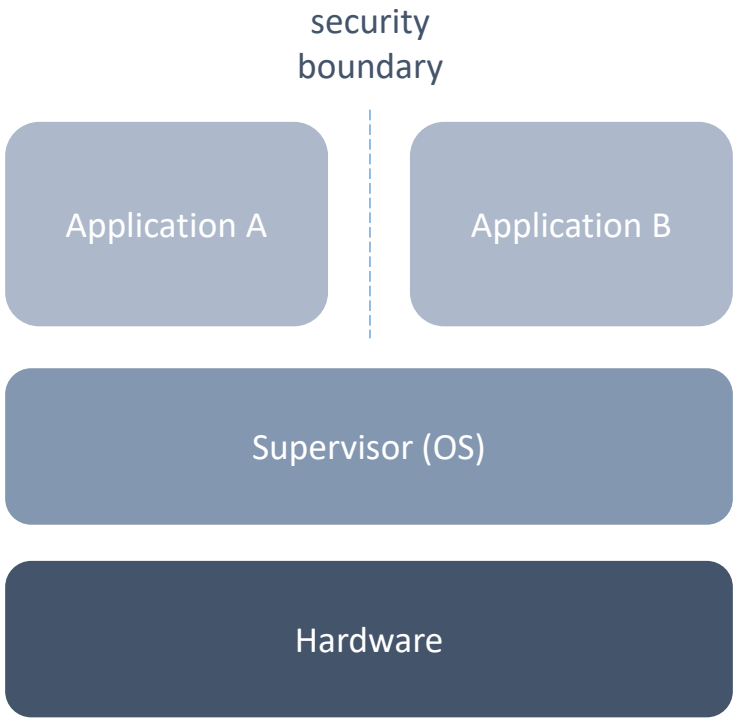

Microarchitectural Timing Channel

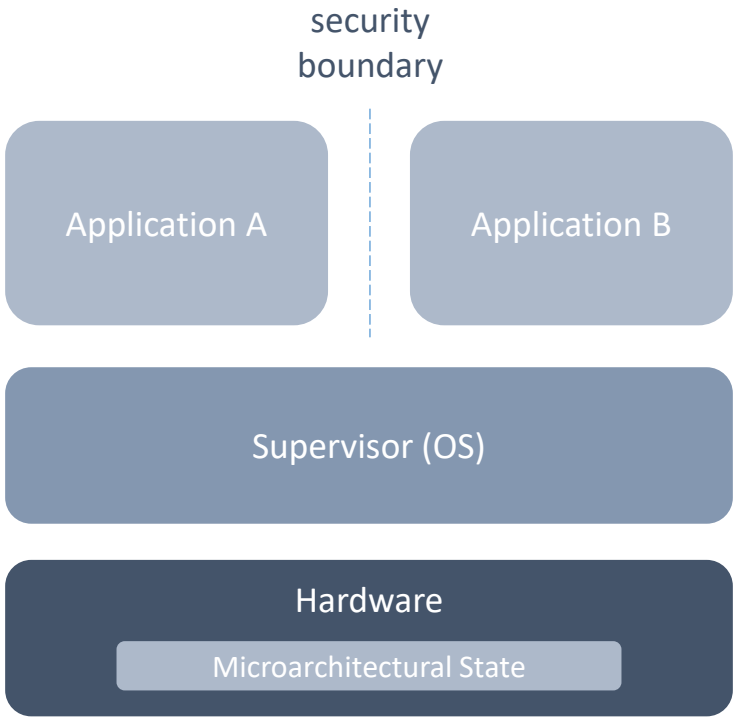

Microarchitectural Timing Channel

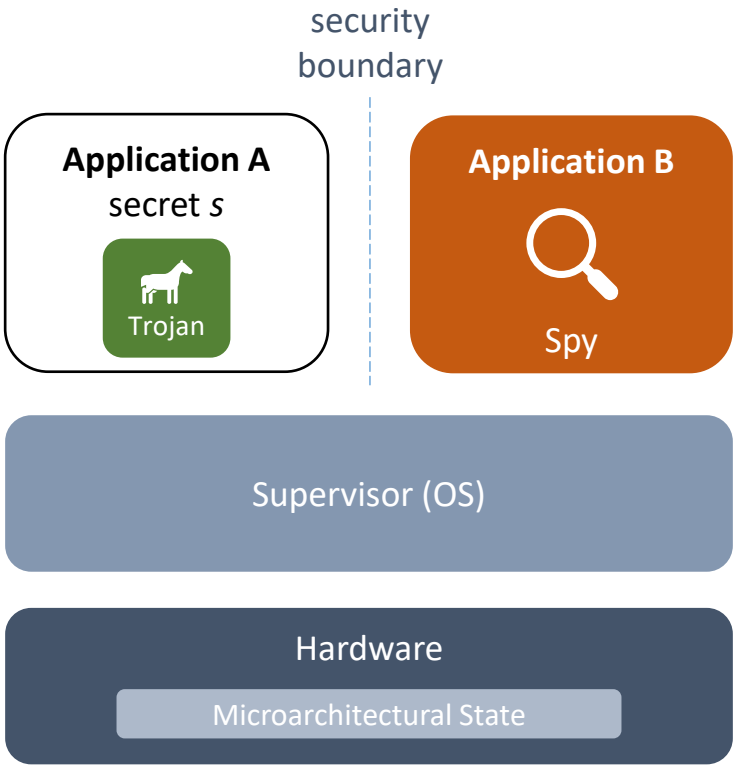

Microarchitectural Timing Channel

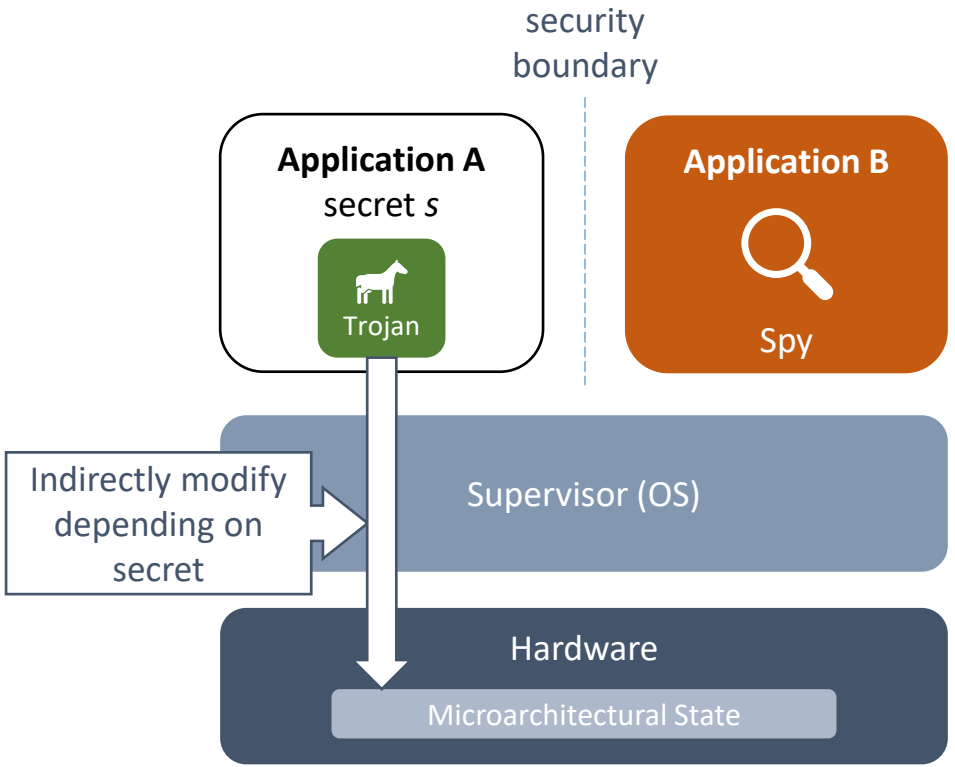

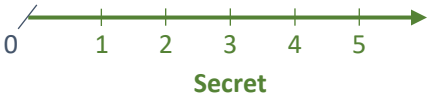

Microarchitectural Timing Channel

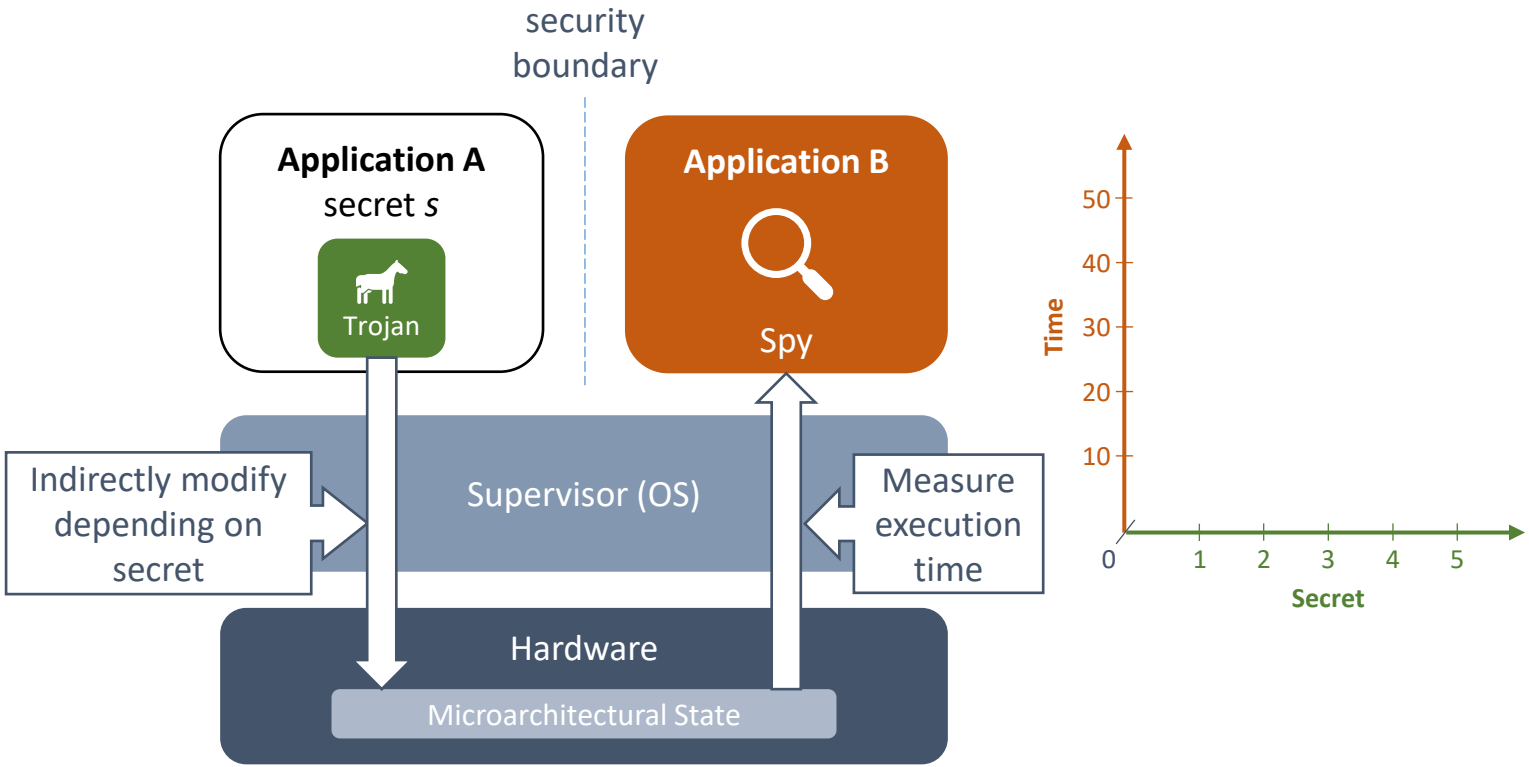

Microarchitectural Timing Channel

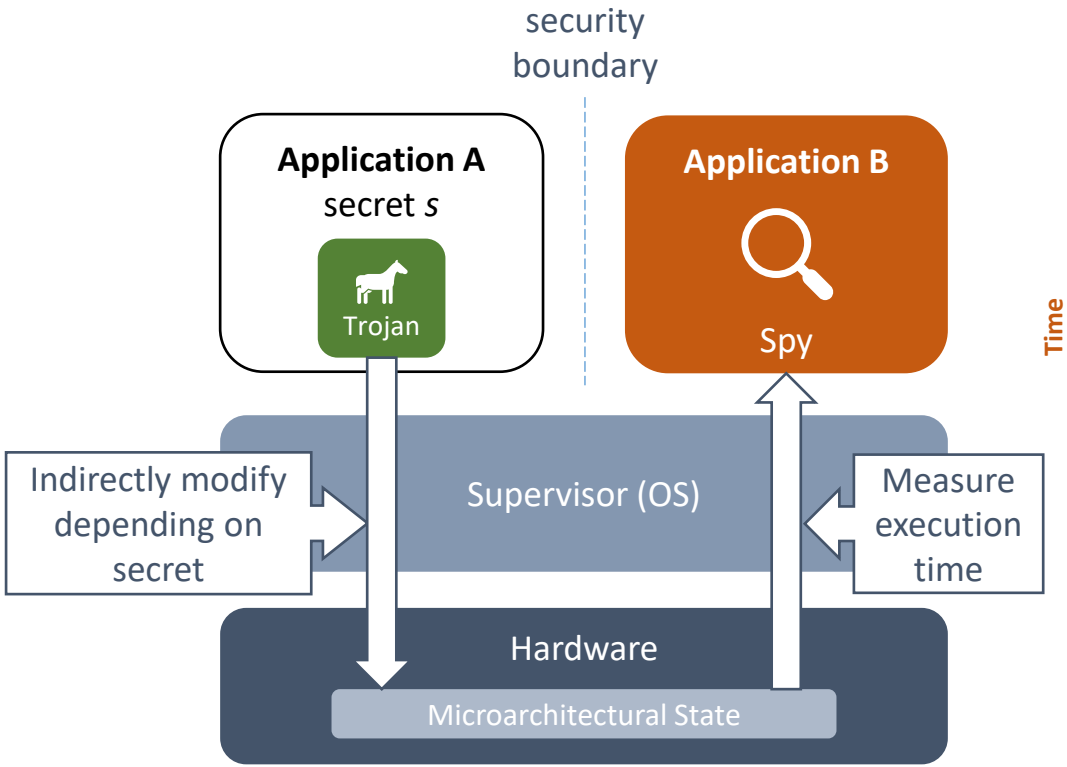

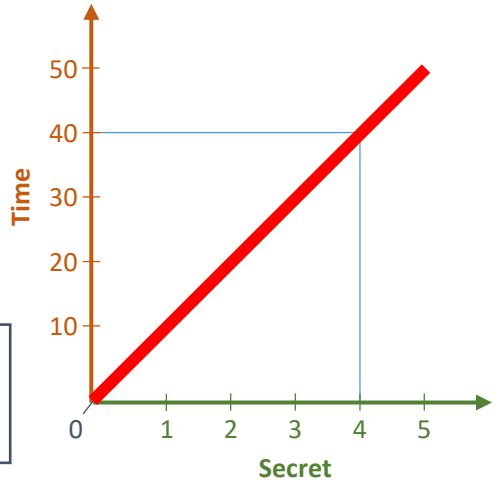

Microarchitectural Timing Channel

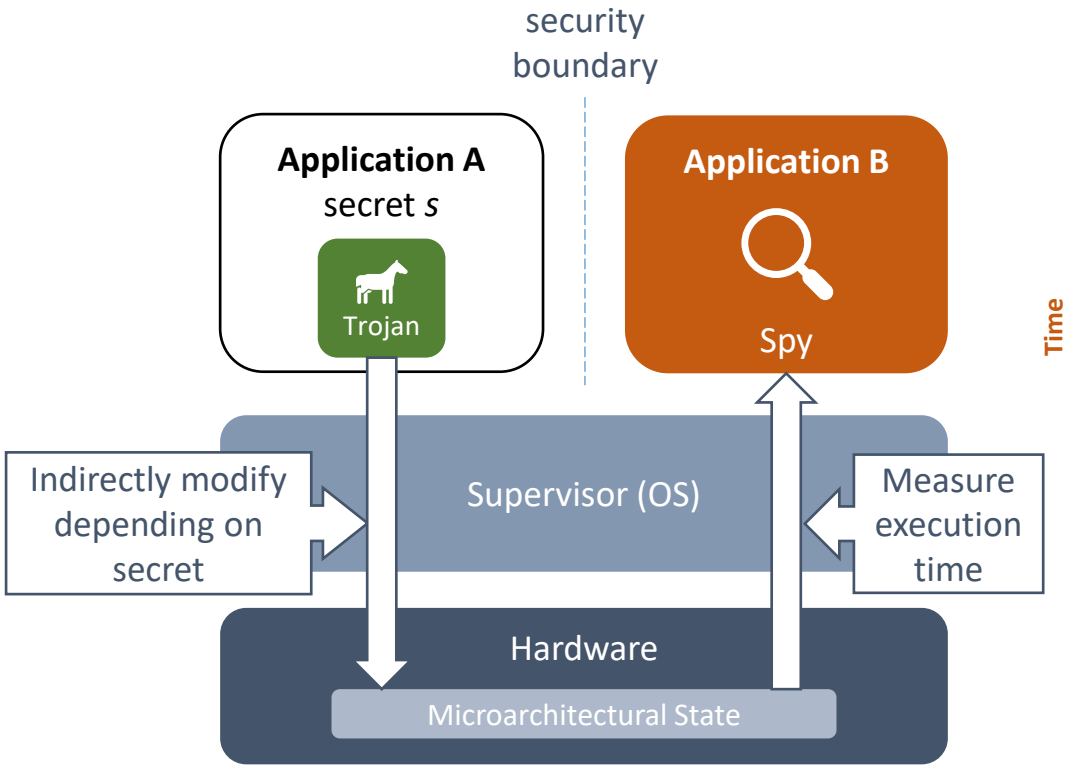

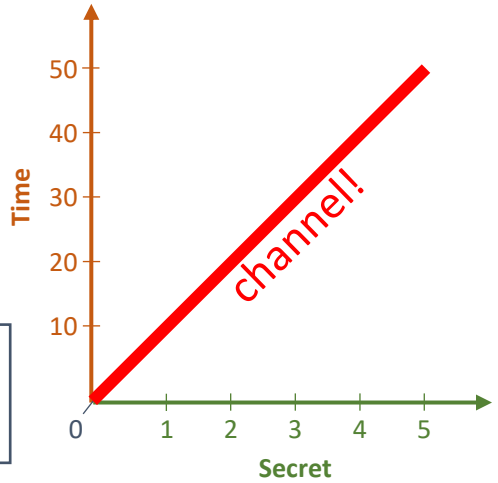

Microarchitectural Timing Channel

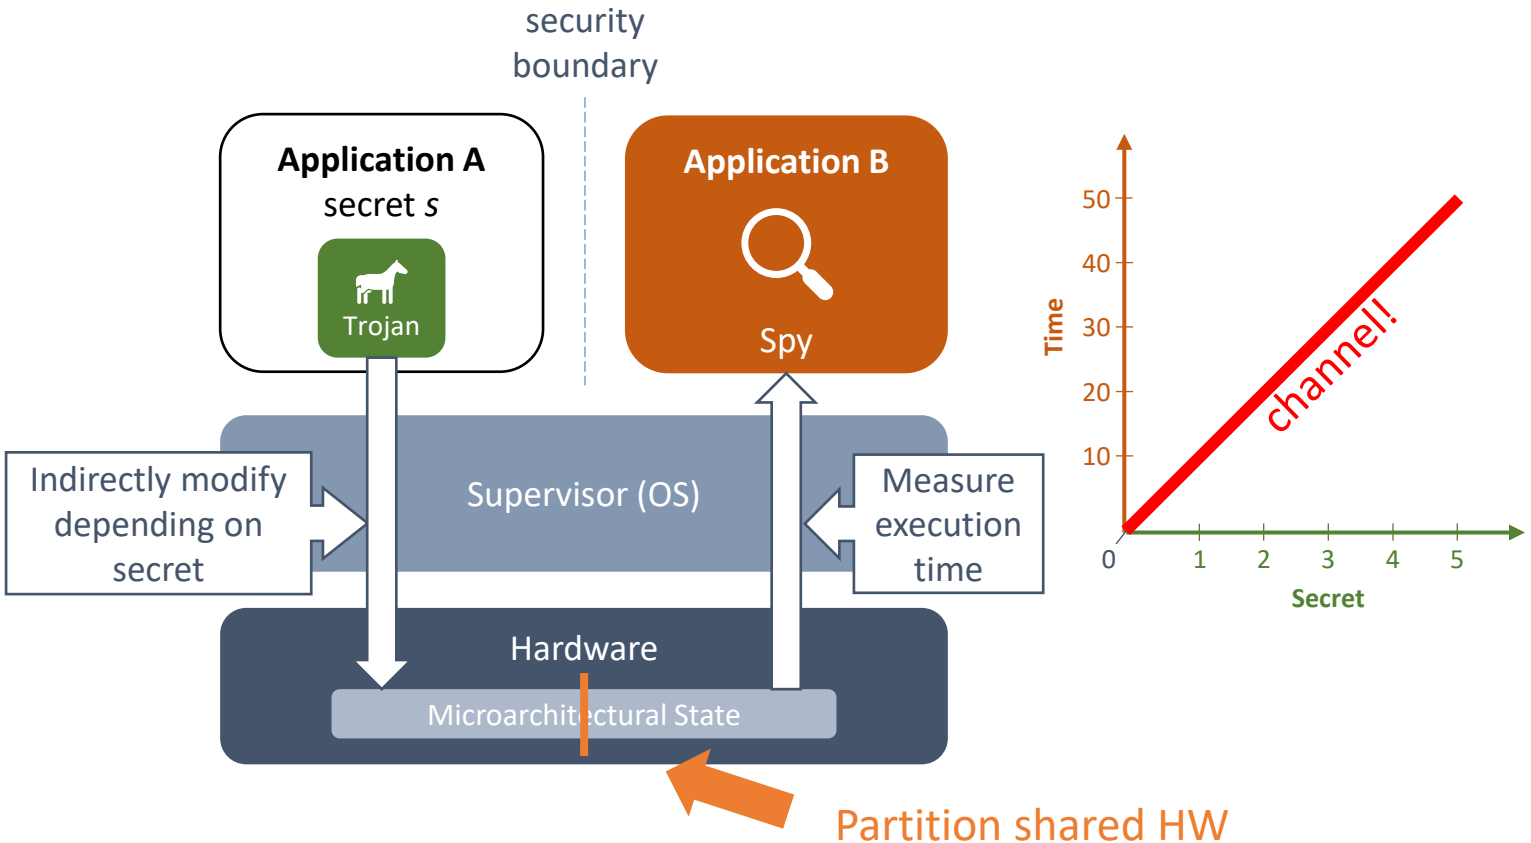

Microarchitectural Timing Channel

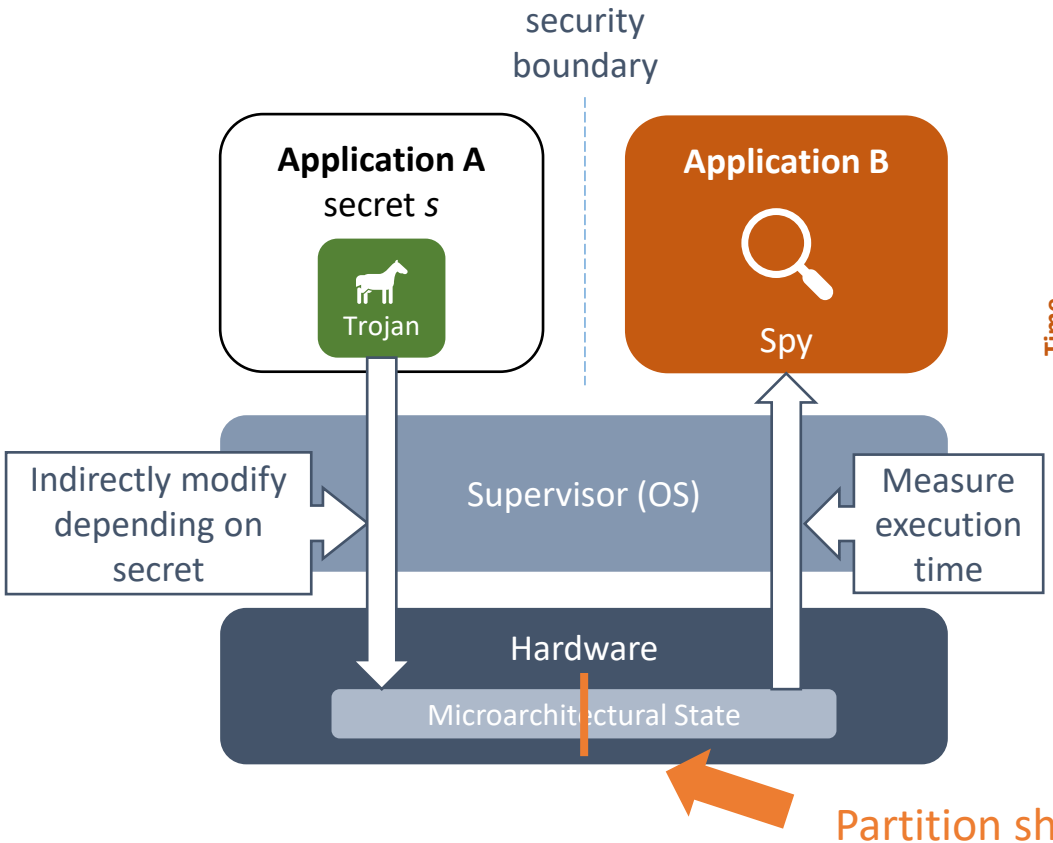

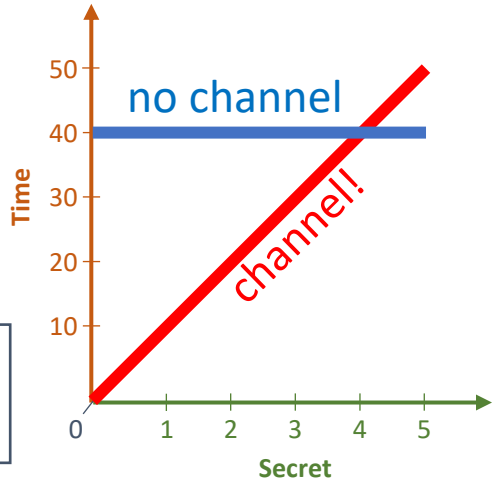

Evaluation Platform

Hardware platform

CVA6 RV64GC core [2]
on FPGA

Evaluation Platform

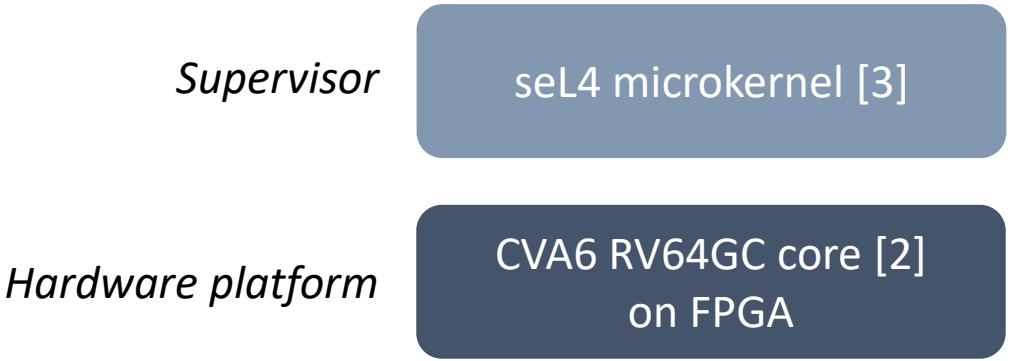

Evaluation Platform

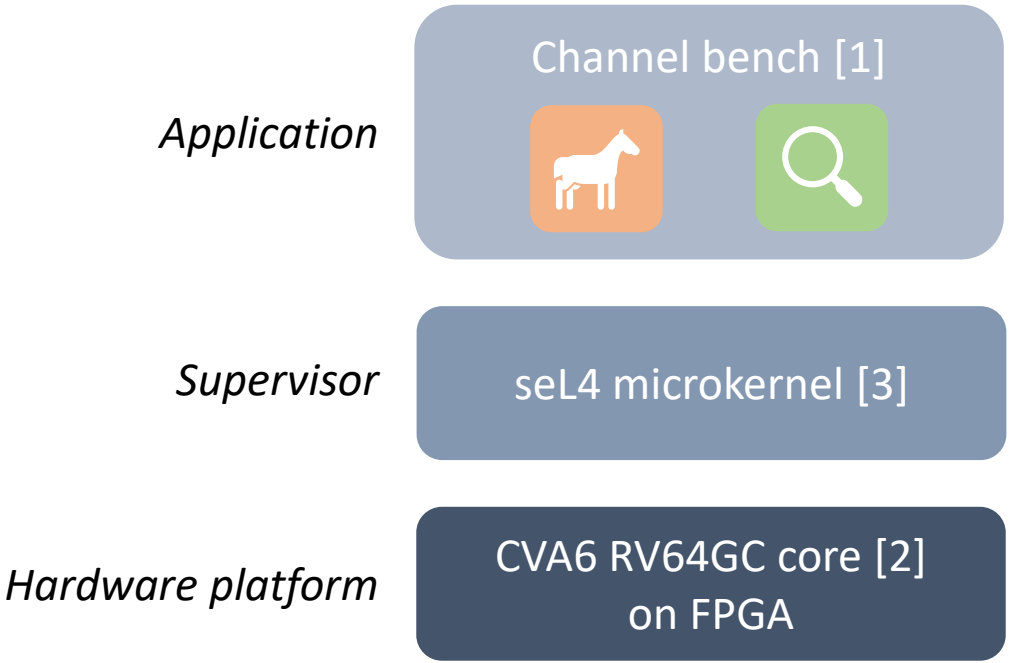

Channel Matrix: L1 D$

$N = 10^6$

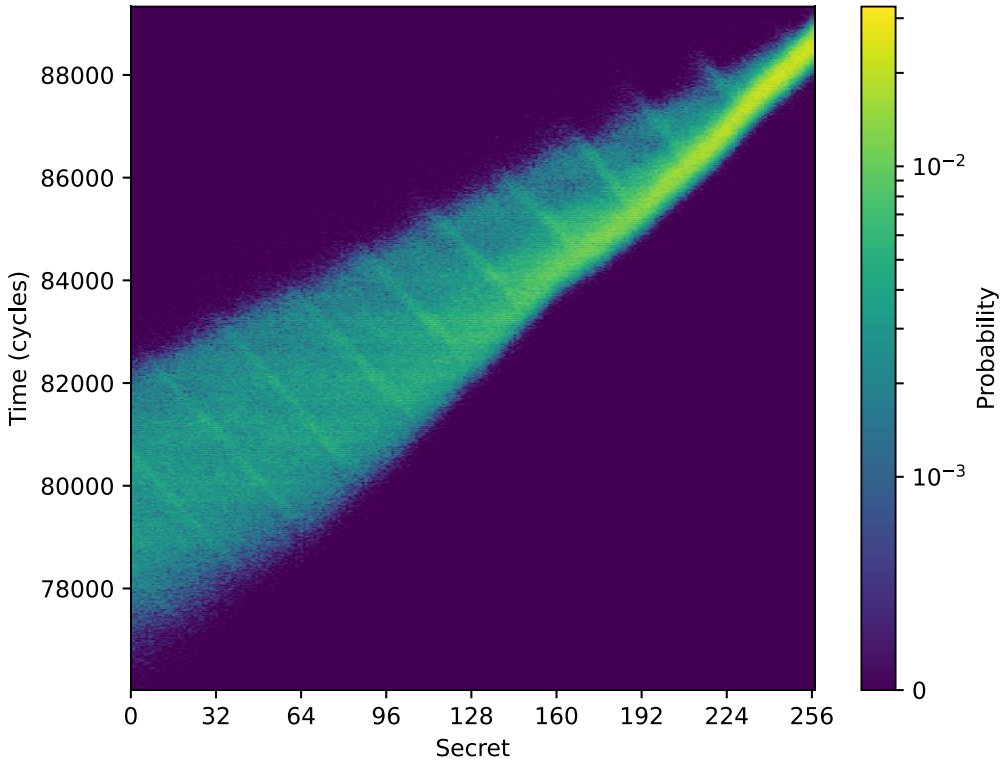

Channel Matrix: L1 D$

$N = 10^6$

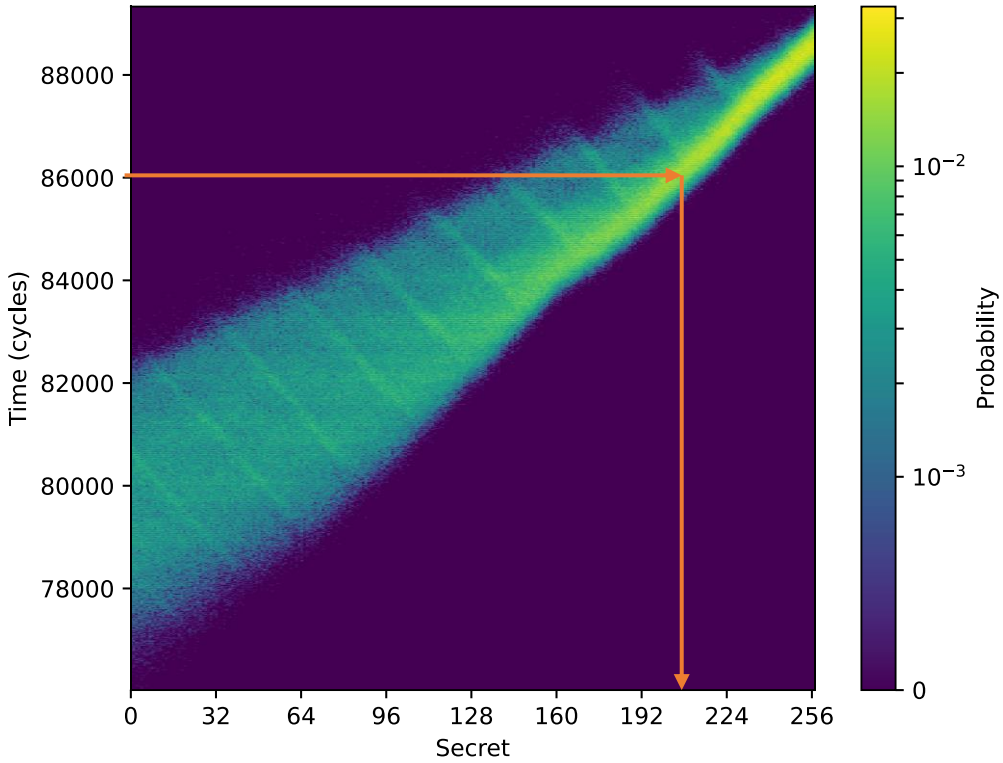

Channel Matrix: L1 D$

$N = 10^6$

$M = 1667.3$ mb

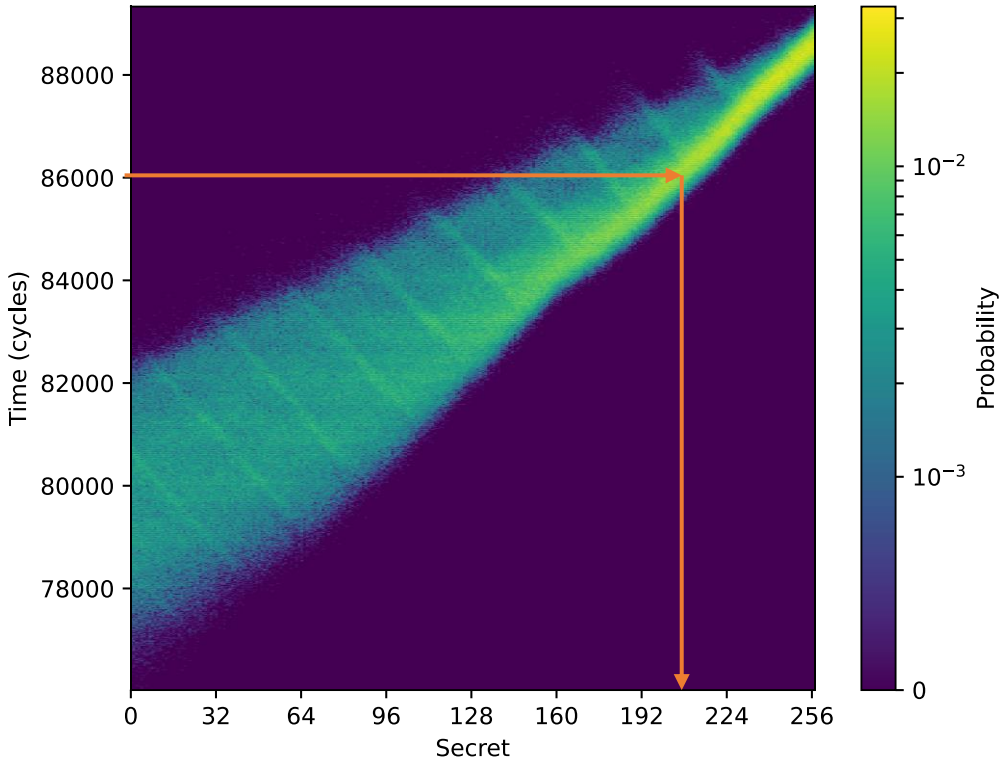

Channel Matrix: L1 D$

N = 10^6

M = 1667.3 mb

M_0 = 0.5 mb

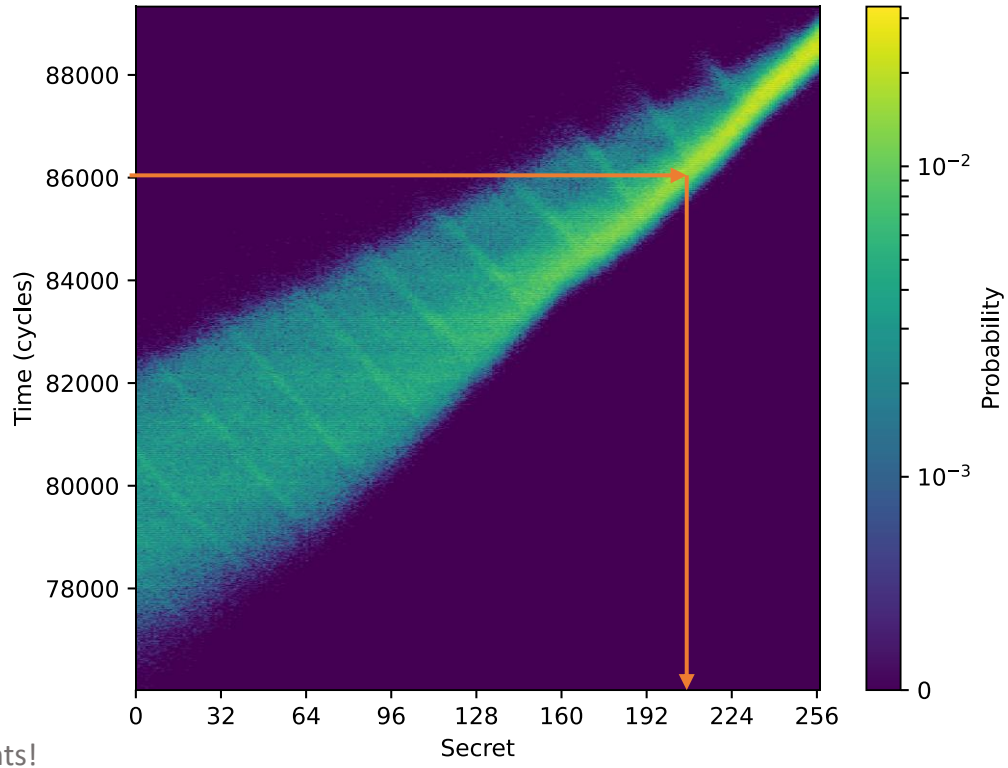

M_0 varies between Measurements!

Software Mitigation: L1 D$ Channel

Unmitigated

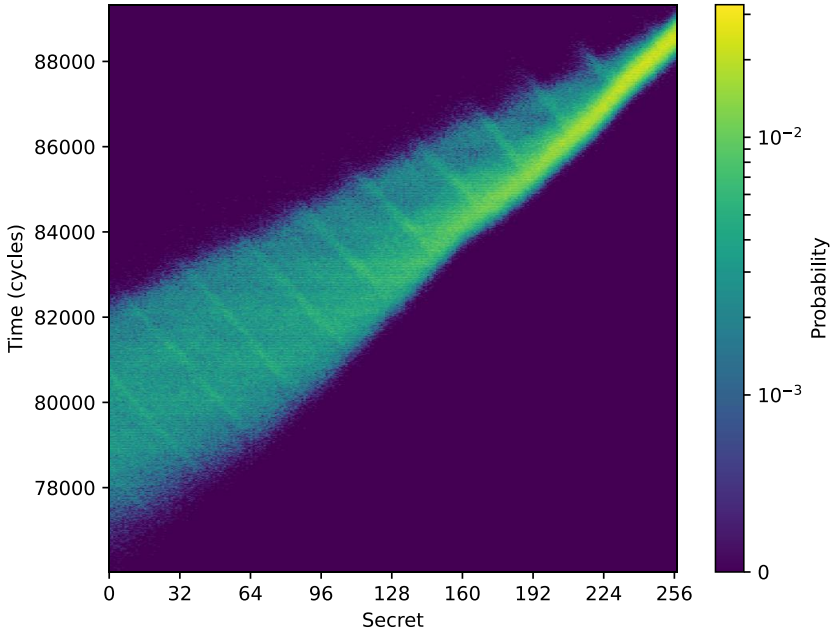

$N = 10^6$, $M = 1667.3$ mb, $M_0 = 0.5$ mb

Double L1 D$ prime on context switch

Reduced Range

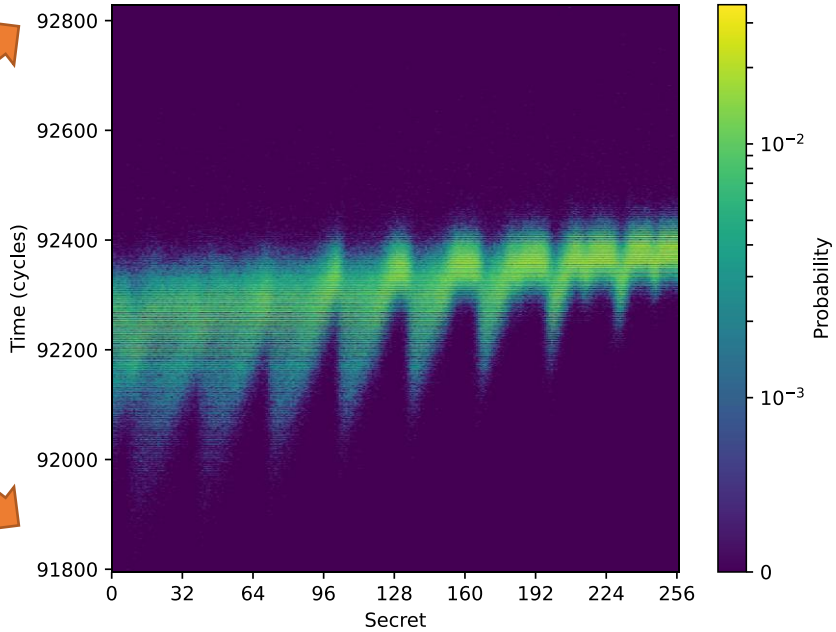

$N = 10^6$, $M = 515.7$ mb, $M_0 = 1.1$ mb

Software Mitigation: L1 D$ Channel

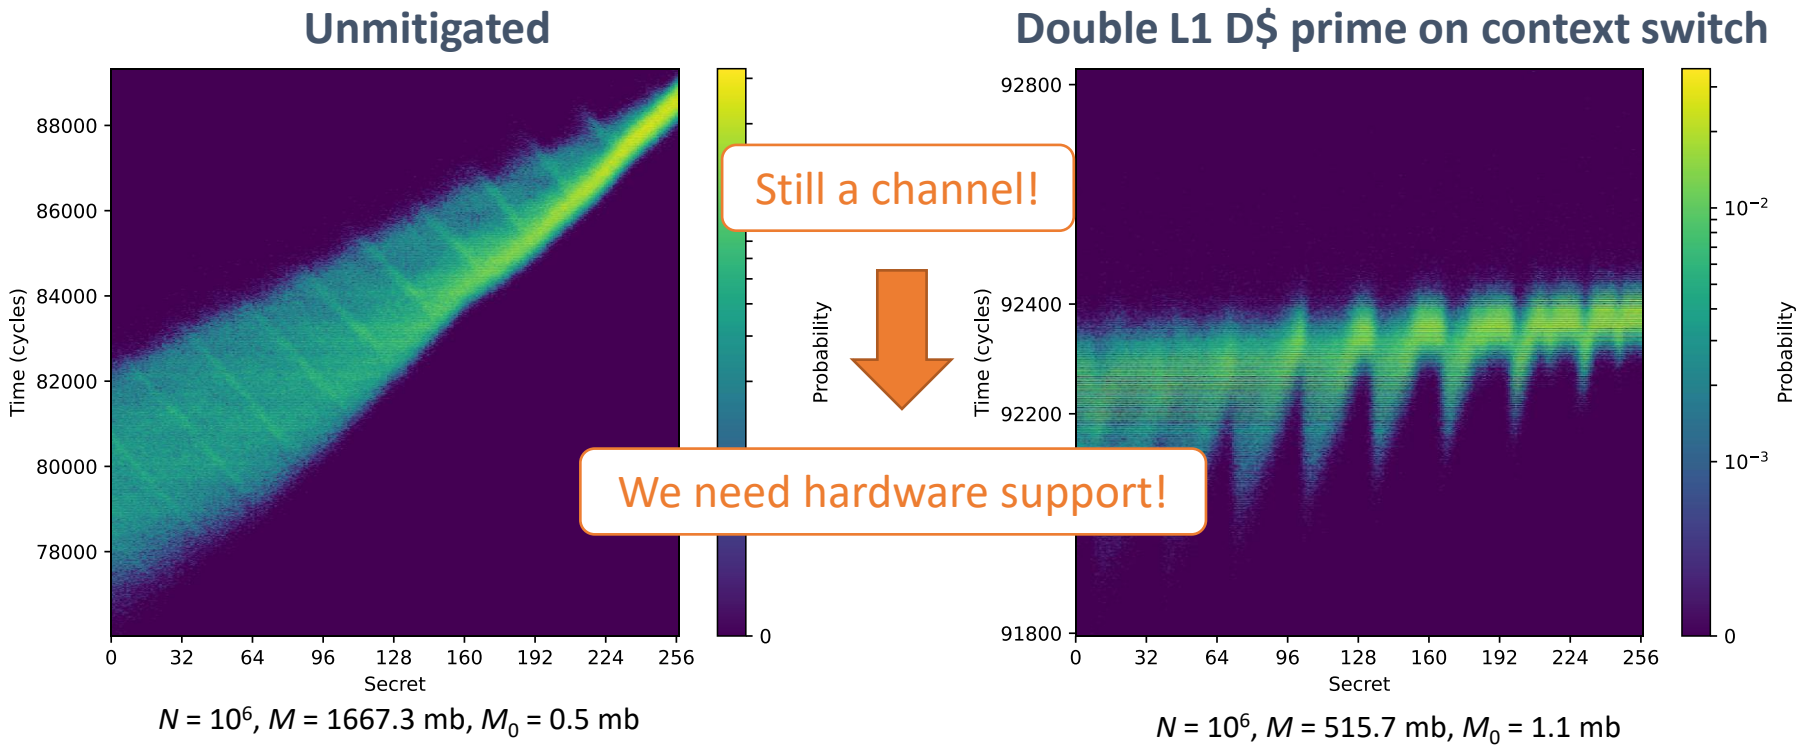

$N = 10^6$, $M = 1667.3$ mb, $M_0 = 0.5$ mb

$N = 10^6$, $M = 515.7$ mb, $M_0 = 1.1$ mb

Temporal Fence Instruction (`fence.t`)

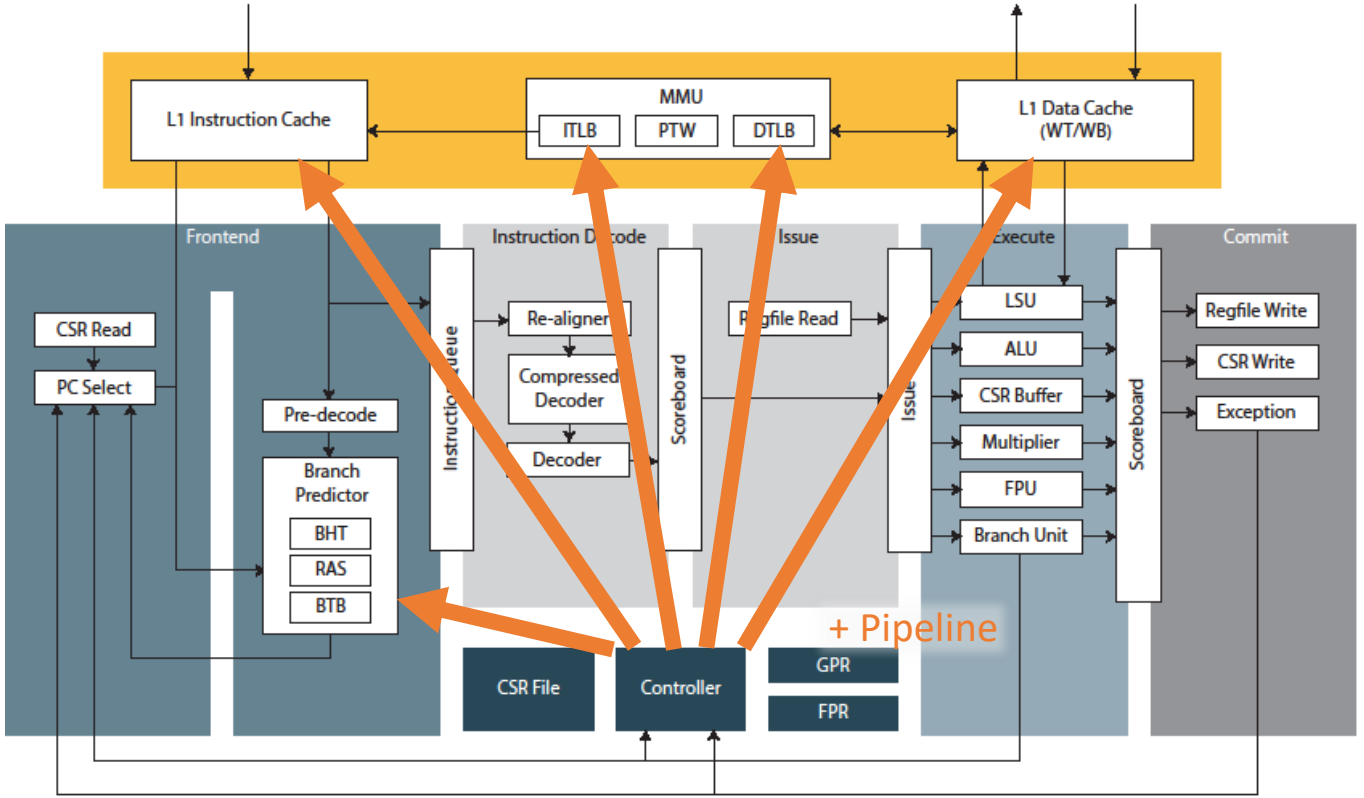

`fence.t`: L1 D$ Channel

Unmitigated

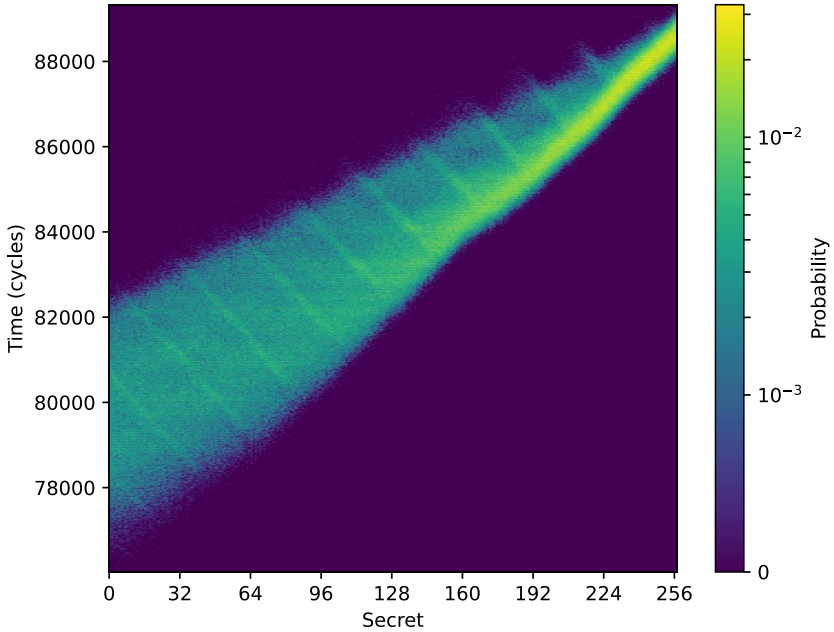

$N = 10^6$, $M = 1667.3$ mb, $M_0 = 0.5$ mb

Flush targeted components on context switch

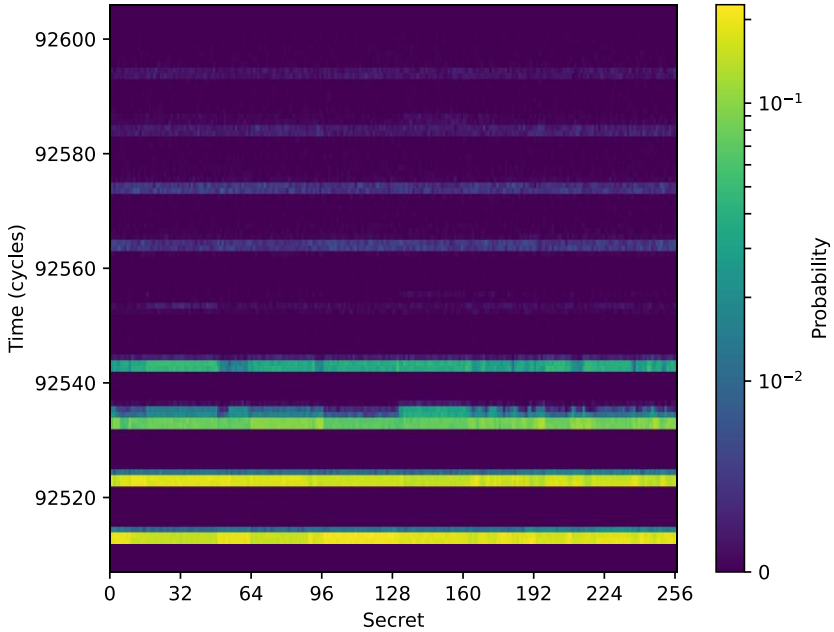

$N = 10^6$, $M = 7.7$ mb, $M_0 = 1.4$ mb

fence.t: L1 D$ Channel

Unmitigated

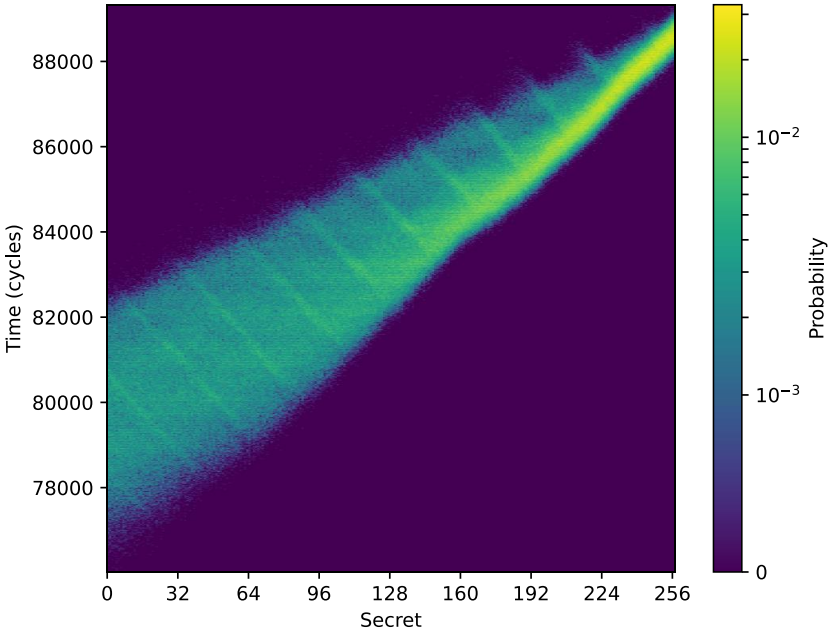

$N = 10^6$, $M = 1667.3$ mb, $M_0 = 0.5$ mb

Flush targeted components on context switch

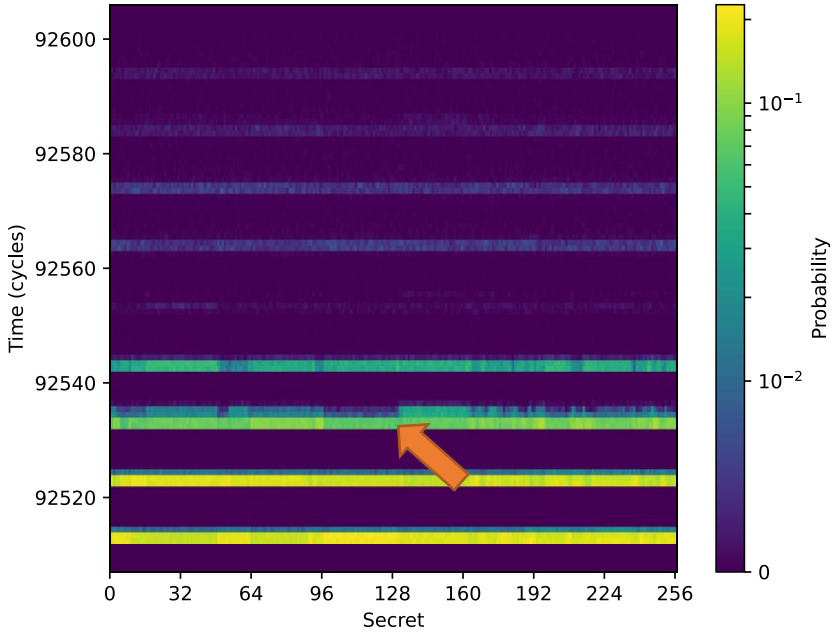

$N = 10^6$, $M = 7.7$ mb, $M_0 = 1.4$ mb

Vulnerable 2nd Order State-Holding Components

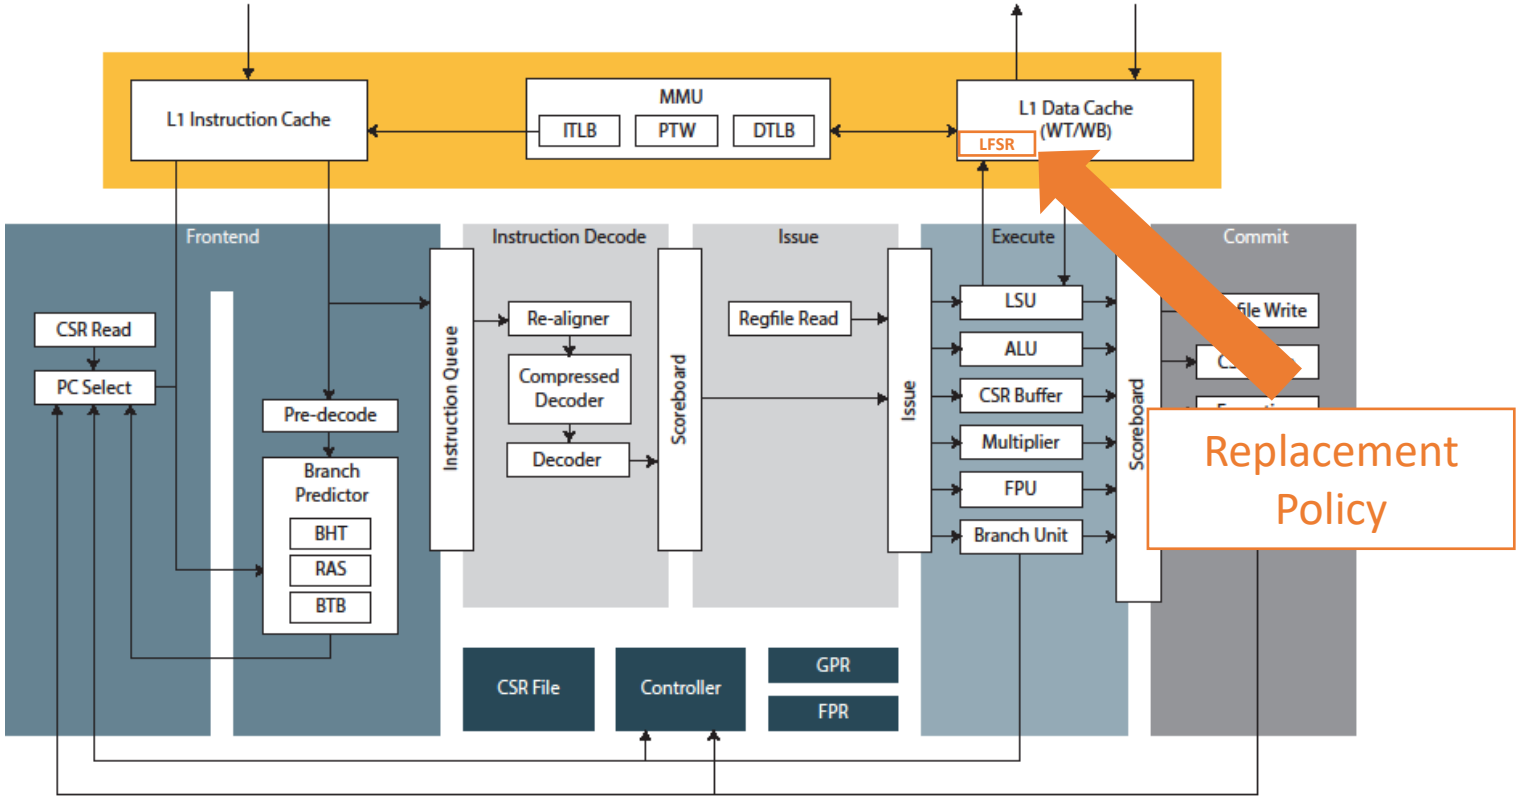

Vulnerable 2nd Order State-Holding Components

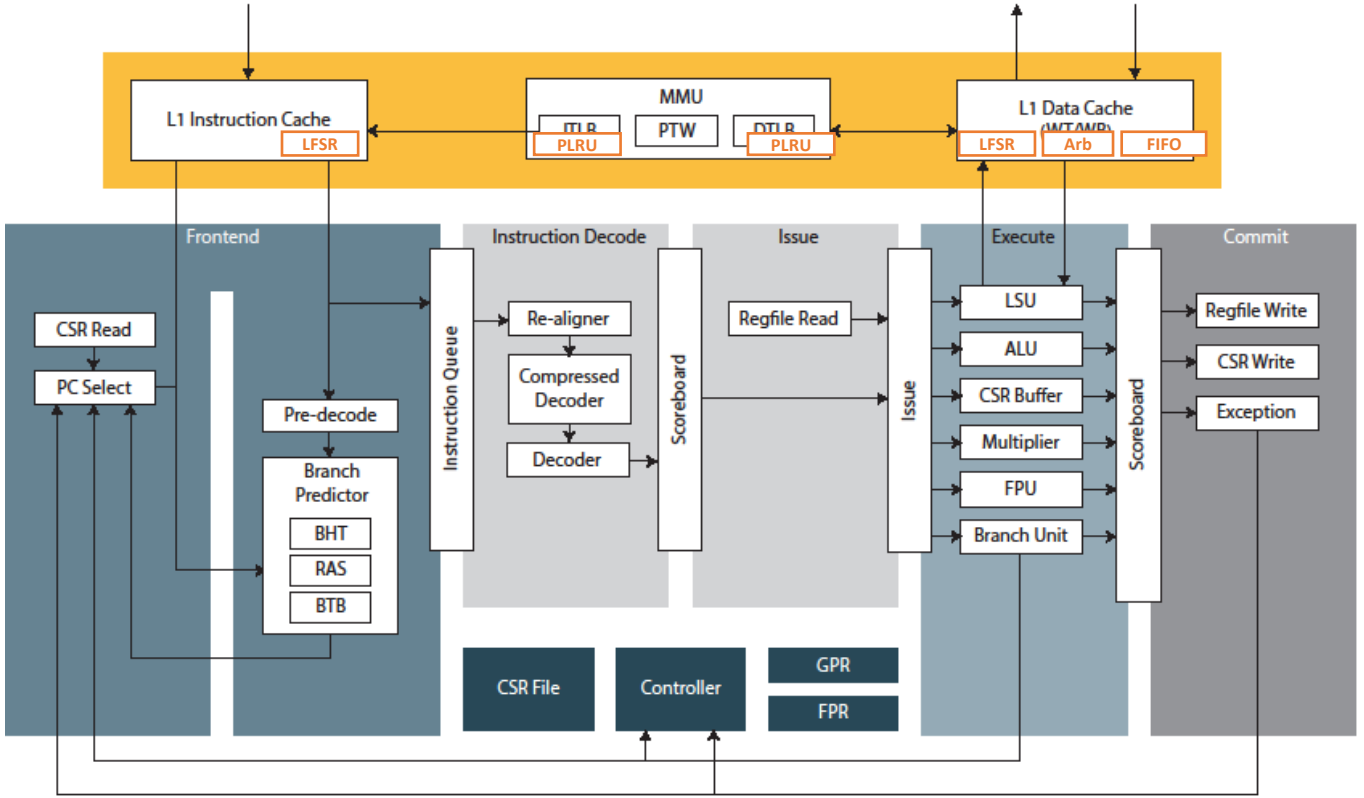

All Channels are Closed!

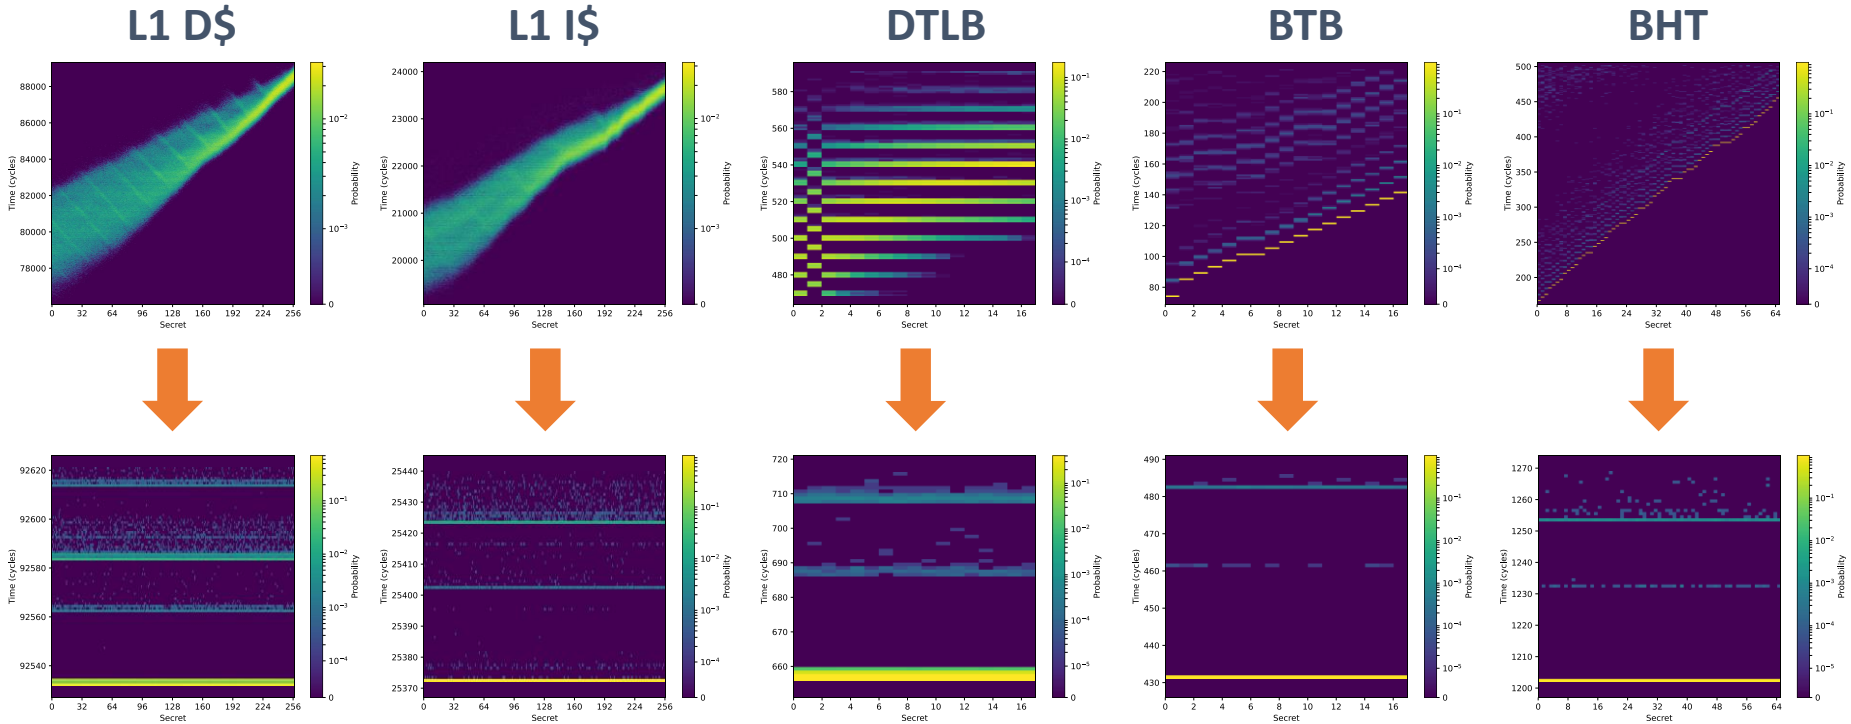

Costs

Context Switch Latency

Unmitigated		D$ Software Flush		HW Flush
Hot	Cold	Single	Double	
430 (±7.0)	1,180 (±1.0)	12,099 (±52)	51,876 (±256)	1,502 (±0.9)

320 cycles overhead per context switch

Clk @1GHz, CS @1KHz: **+ 0.032%**

Costs

Context Switch Latency

Unmitigated		D$ Software Flush		HW Flush
Hot	Cold	Single	Double	
430 (±7.0)	1,180 (±1.0)	12,099 (±52)	51,876 (±256)	1,502 (±0.9)

320 cycles overhead per context switch

Clk @1GHz, CS @1KHz: **+ 0.032%**

Hardware Costs

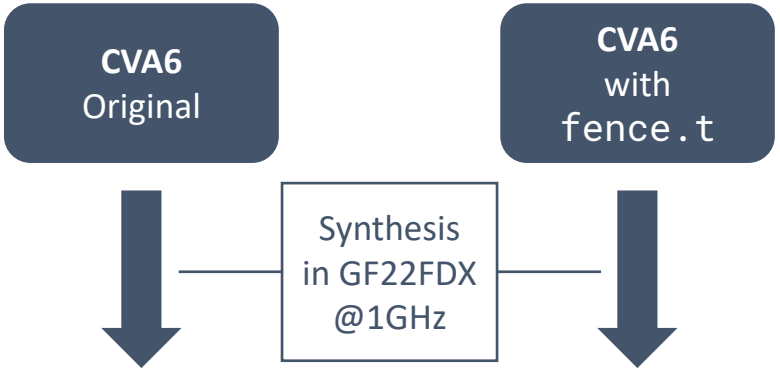

No notable area difference!

Conclusion

- We measure timing channels on an in-order RISC-V core (CVA6)
- We show that SW alone cannot solve the problem!
- Solution: Enable OS to flush microarchitectural state
  - We propose a temporal fence (`fence.t`) instruction
  - Closes all evaluated channels at negligible costs
- Need to flush *all* μArch state with possible timing impact!
- Future work
  - Evaluate performance with *write-back* L1 D$
  - Develop systematic approach to identify vulnerable μArch state

Sources

[1] Qian Ge, Yuval Yarom, Tom Chothia, and Gernot Heiser: “Time Protection: The Missing OS Abstraction”, EuroSys, 2019

[2] Florian Zaruba and Luca Benini: “The Cost of Application-Class Processing: Energy and Performance Analysis of a Linux-Ready 1.7-GHz 64-Bit RISC-V Core in 22-nm FDSOI Technology”, IEEE Trans. on VLSI Systems 27, 2019

[3] Gerwin Klein, June Andronick, Kevin Elphistone, Toby Murray, Thomas Sewell, Rafal Kolanski, and Gernot Heiser: “Comprehensive Formal Verification of an OS Microkernel”, ACM Trans. Comp. Syst. 32, 2014

DESIGN, AUTOMATION & TEST IN EUROPE
01 – 05 February 2021 · virtual conference
The European Event for Electronic System Design & Test

Microarchitectural Timing Channels and their Prevention on an Open-Source 64-bit RISC-V Core

Nils Wistoff ETH Zurich
Moritz Schneider ETH Zurich
Frank K. Gürkaynak ETH Zurich
Luca Benini ETH Zurich and University of Bologna
Gernot Heiser UNSW Sydney and Data61 CSIRO

Temporal Fence Instruction (`fence.t`)

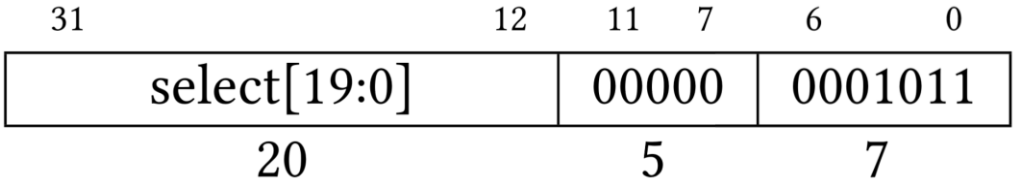

Temporal Fence Instruction (`fence.t`)

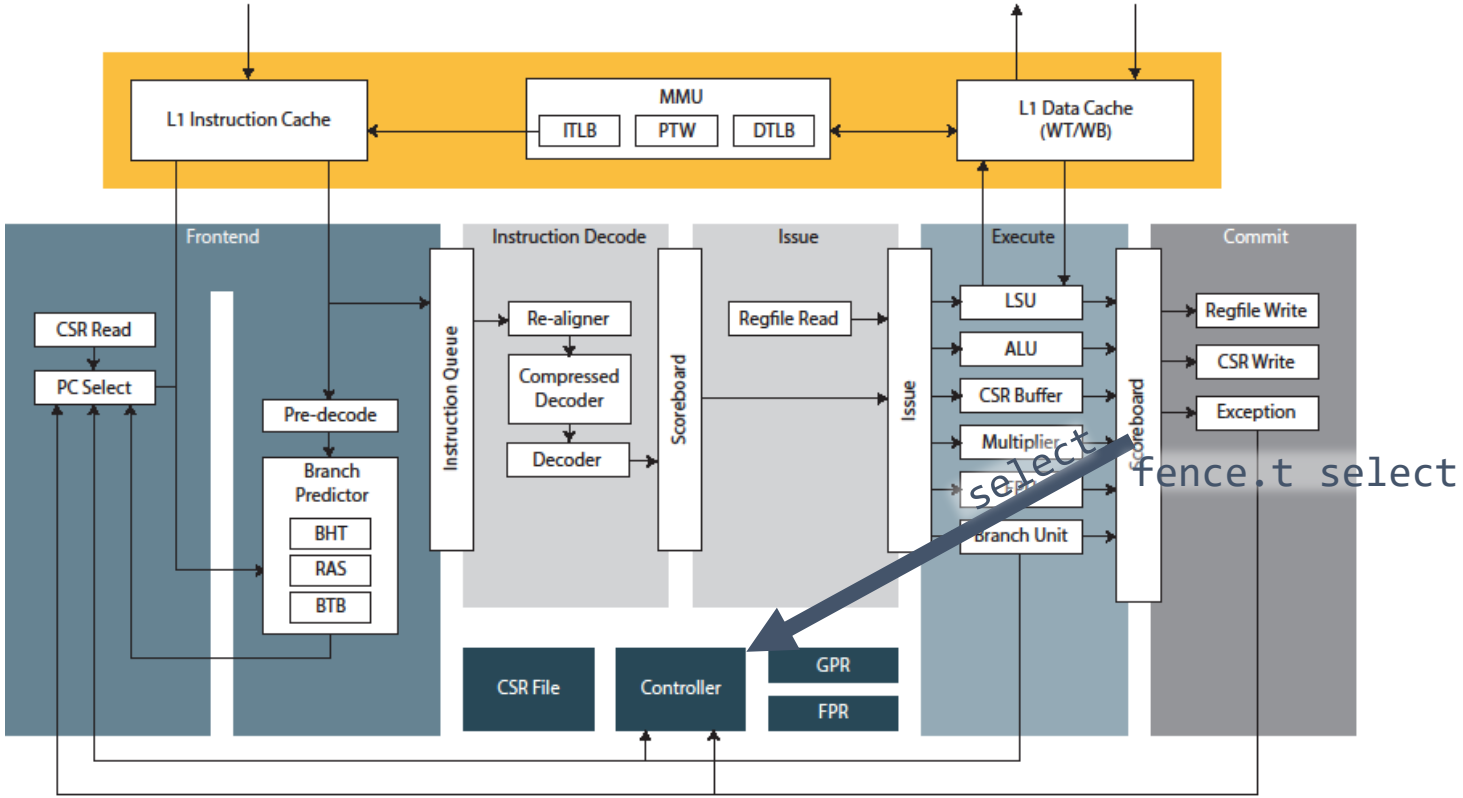

Covert Channel

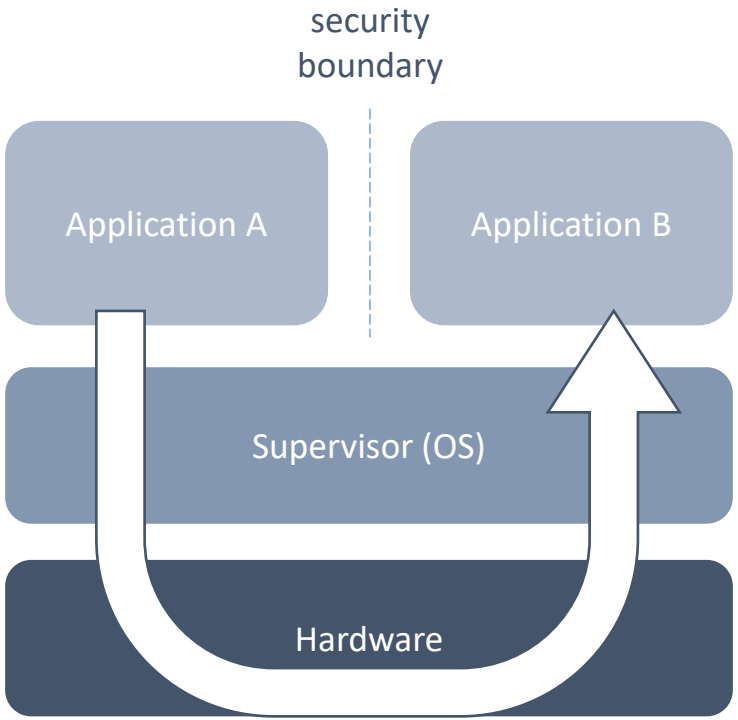

Evaluation Platform