Hardware Support for Time Protection

seL4 Summit 2024, Sydney, Australia, 2024-10-17

Nils Wistoff nwistoff@iis.ee.ethz.ch
Gernot Heiser gernot@unsw.edu.au
Luca Benini lbenini@iis.ee.ethz.ch

PULP Platform
Open Source Hardware, the way it should be!

@pulp_platform
pulp-platform.org
youtube.com/pulp_platform

Team of 100 people in ETH Zürich – University of Bologna

- **Research on open-source energy-efficient computing**

PULP open-source hardware ecosystem

RISC-V Cores and Vector Units

RI5CY *CV32E*	Zero R *Ibex*	Snitch	Spatz	Ariane *CVA6*	Ara
RV32	RV32	RV32	RVV	RV64	RVV

Peripherals

JTAG, SPI, UART, I2S, DMA, GPIO

Interconnects

LIC, HCI, APB, FlooNoC, AXI4

Platforms

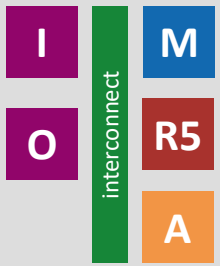

Single core
- PULPino, PULPissimo
- Cheshire

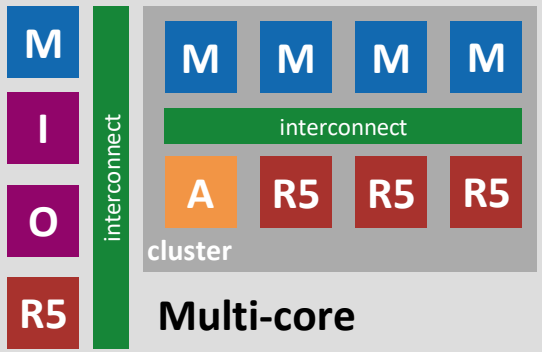

Multi-core
- OpenPULP
- ControlPULP

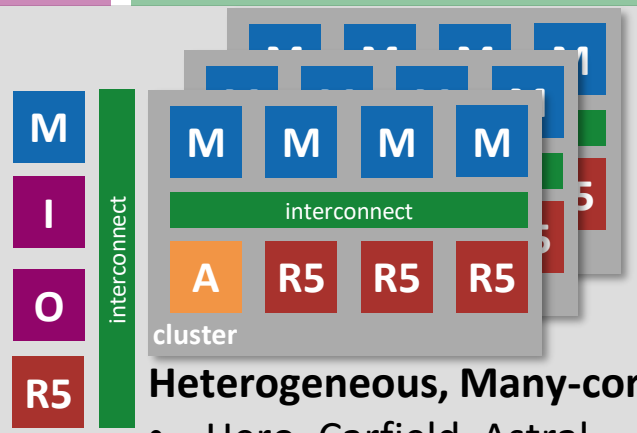

Heterogeneous, Many-core
- Hero, Carfield, Astral
- Occamy, Mempool

IOT → **HPC**

Accelerators and ISA extensions

XpulpNN, XpulpTNN	ITA (Transformers)	RBE, NEUREKA (QNNs)	FFT (DSP)	(FP-Tensor)

github.com/pulp-platform

In 11 years PULP team has designed more than 60 chips

Spectre: Exploiting timing channels to leak data [1]

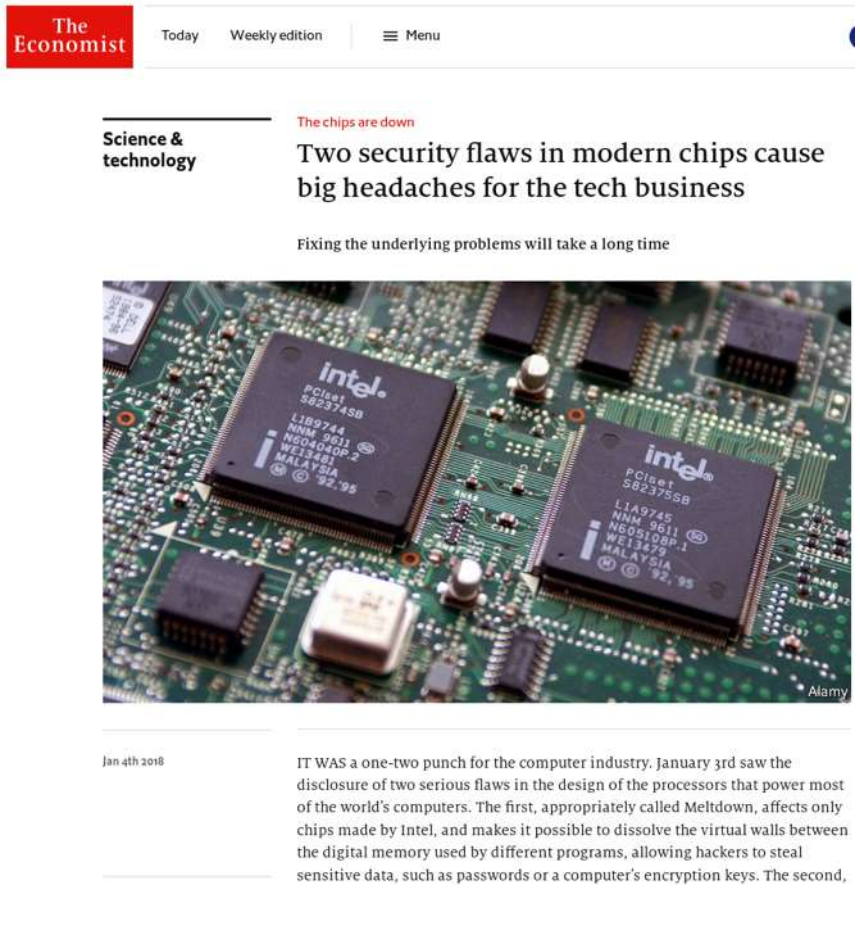
The Economist

Today Weekly edition Menu

Science & technology

The chips are down

Two security flaws in modern chips cause big headaches for the tech business

Fixing the underlying problems will take a long time

Jan 4th 2018

IT WAS a one-two punch for the computer industry. January 3rd saw the disclosure of two serious flaws in the design of the processors that power most of the world's computers. The first, appropriately called Meltdown, affects only chips made by Intel, and makes it possible to dissolve the virtual walls between the digital memory used by different programs, allowing hackers to steal sensitive data, such as passwords or a computer's encryption keys. The second,


ANDY GREENBERG SECURITY 01.03.2018 03.00 PM

A Critical Intel Flaw Breaks Basic Security for Most Computers

A Google-led team of researchers has found a critical chip flaw that developers are scrambling to patch in millions of computers.

SPECTRE

Speculative Execution + **Timing Channel**

theguardian

news opinion sport arts lifestyle more

Data and computer security

Meltdown and Spectre: 'worst ever' CPU bugs affect virtually all computers

Everything from smartphones and PCs to cloud computing affected by major security flaw found in Intel and other processors – and fix could slow devices

Spectre and Meltdown processor security flaws – explained

Samuel Gibbs

THE VERGE TECH REVIEWS SCIENCE CREATORS ENTERTAINMENT VIDEO MORE

Intel's processors have a security bug and the fix could slow down PCs

By Tom Warren | @tomwarren | Jan 3, 2018, 8:45am EST

[1] Kocher et al., *Spectre Attacks: Exploiting Speculative Execution*, IEEE S&P 2019

Timing channel

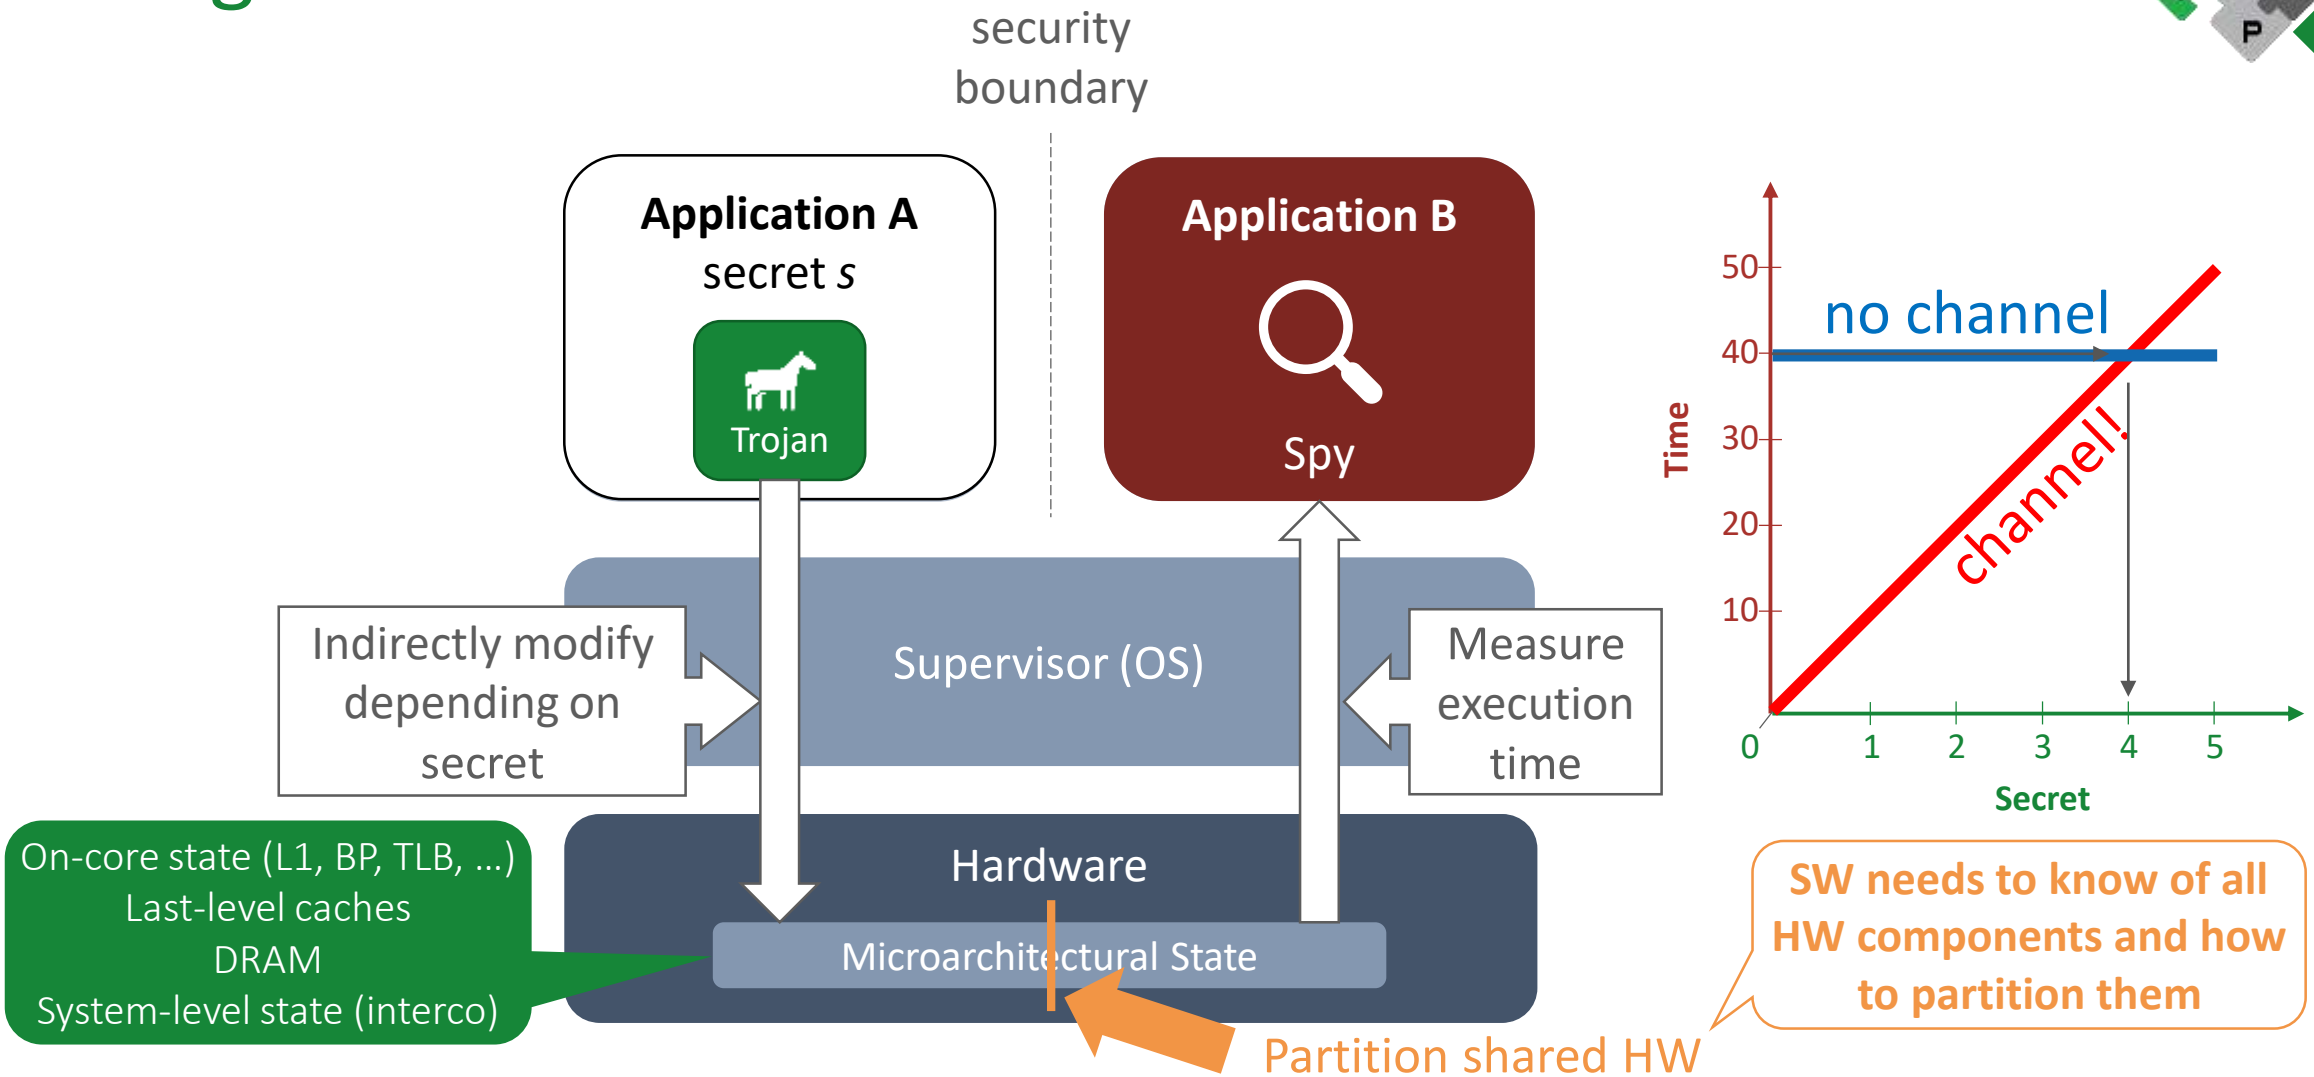

Spatial and temporal partitioning

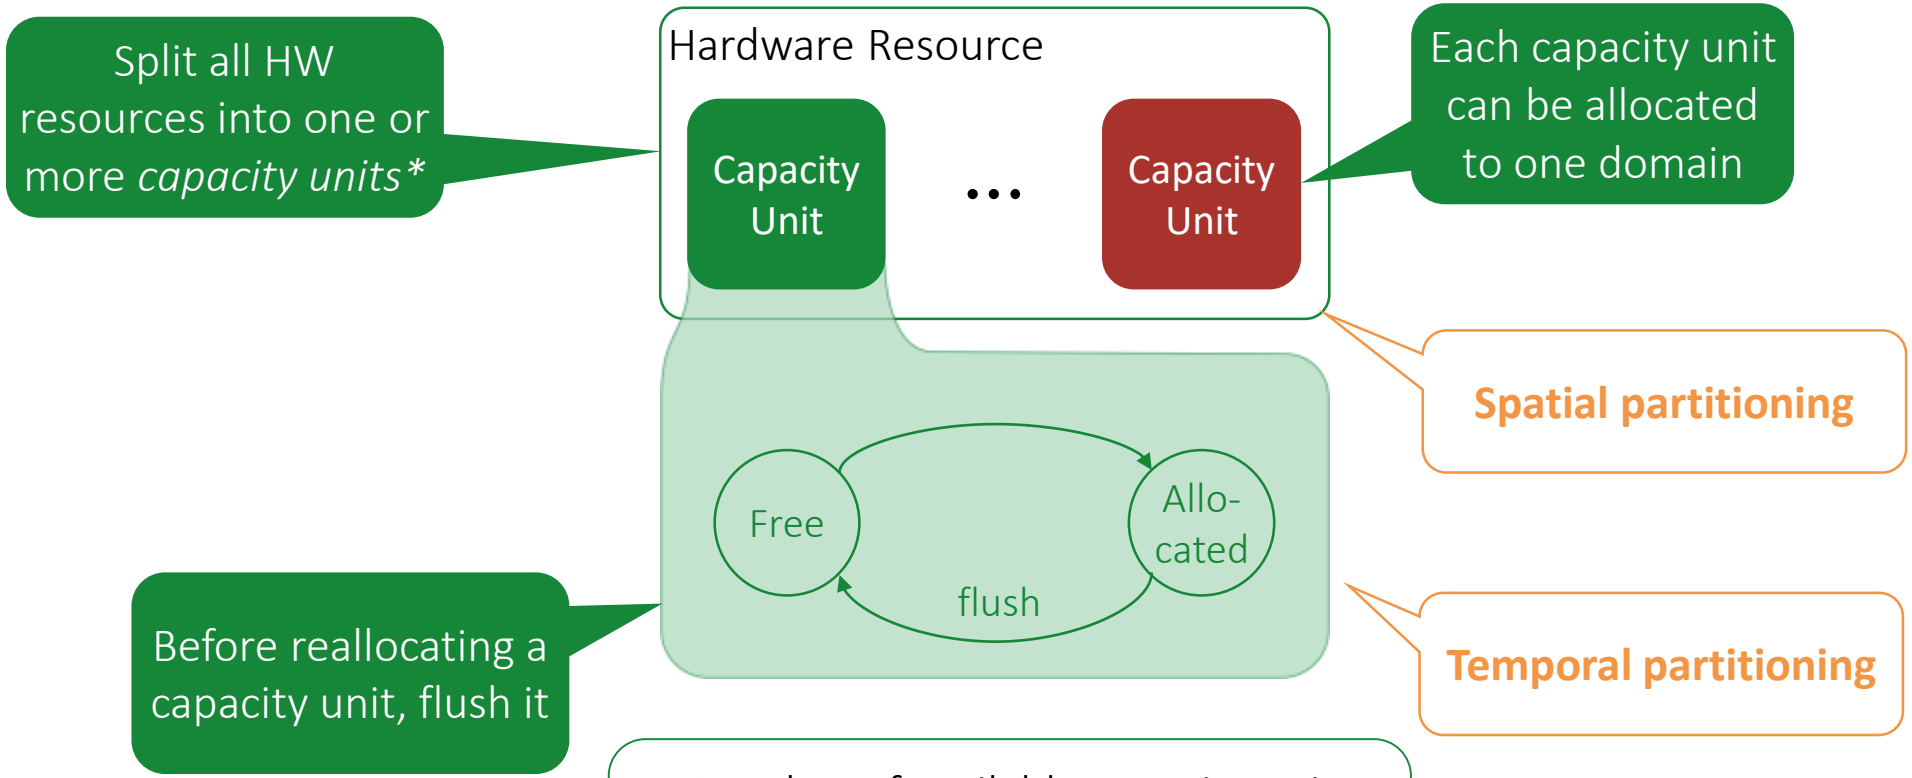

- Number of available capacity units given by HW
- Allocation policy determined by OS

** Terminology adopted from the RISC-V CBQRI extension*

Partitioning a system on chip

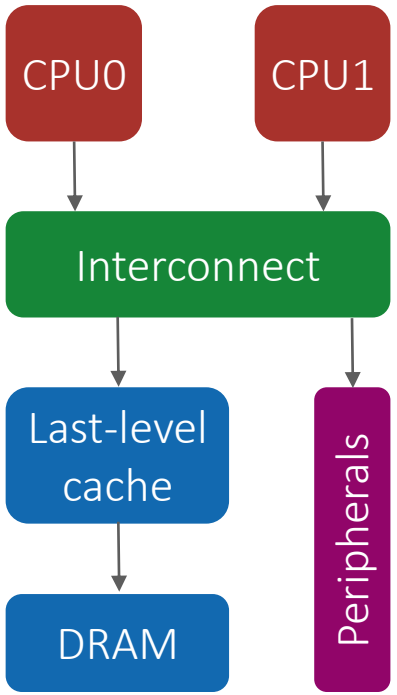

Partitioning the CPU

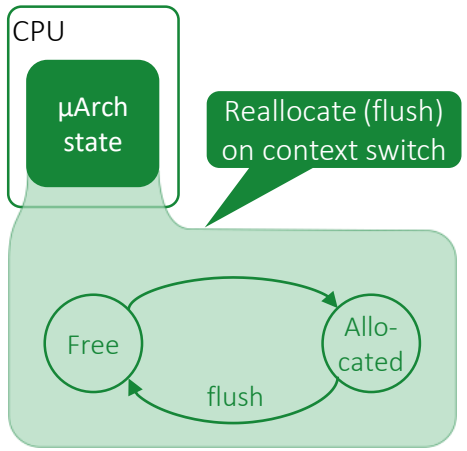

CVA6: 64-bit, application-class, 6-stage, in-order RISC-V core

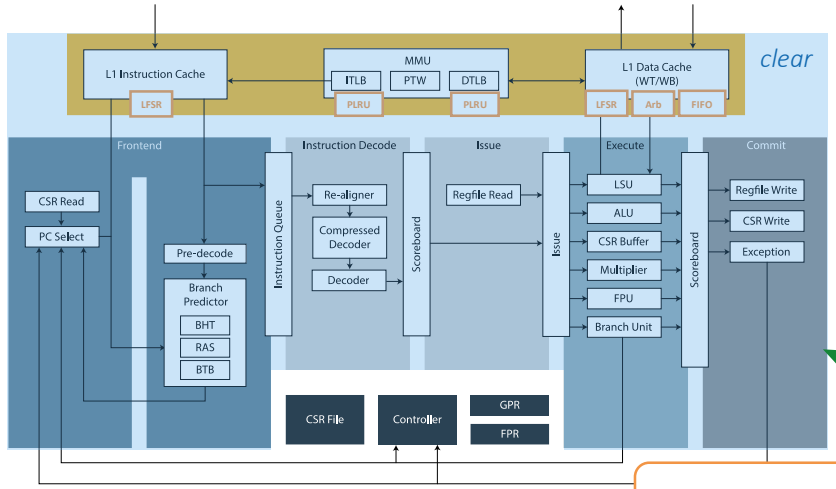

GROUP OPENHW PROVEN PROCESSOR IP

`fence.t` [TCOMP'23]

Negligible HW costs

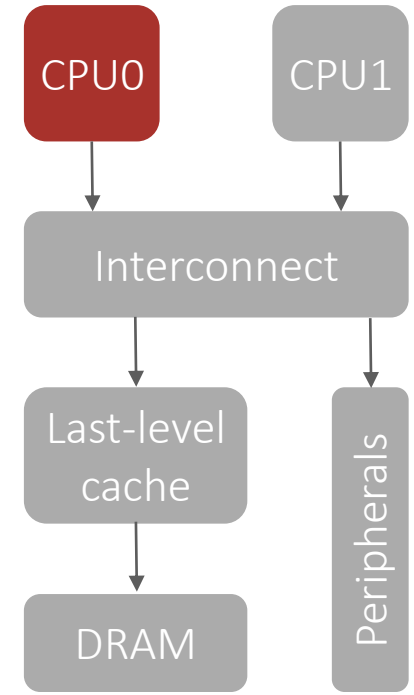

seL4 Security. Performance. Proof.

L1D Unmitigated

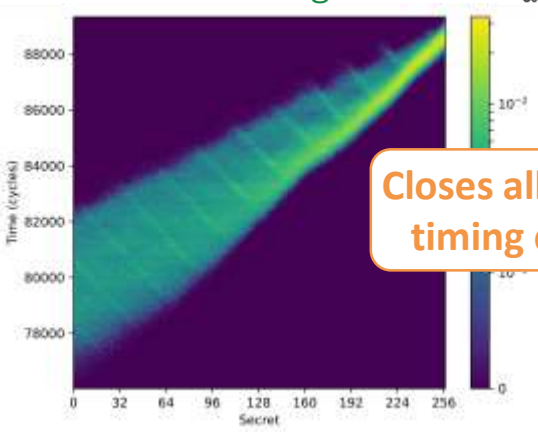

M = **1667.3 mb**, M_0 = 0.5 mb

Closes all observed timing channels

L1D fence.t

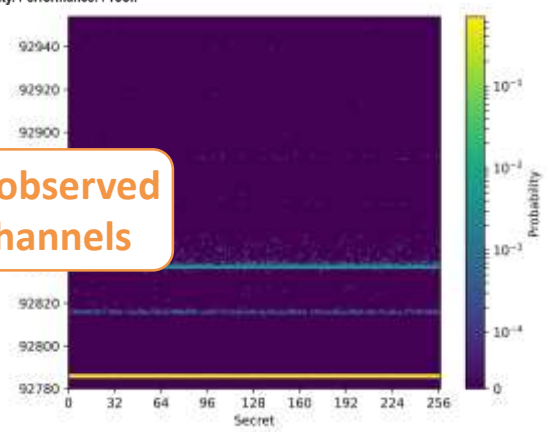

M = 21.7 mb, M_0 = 27.8 mb

Splash-2 Benchmark Overhead (%)

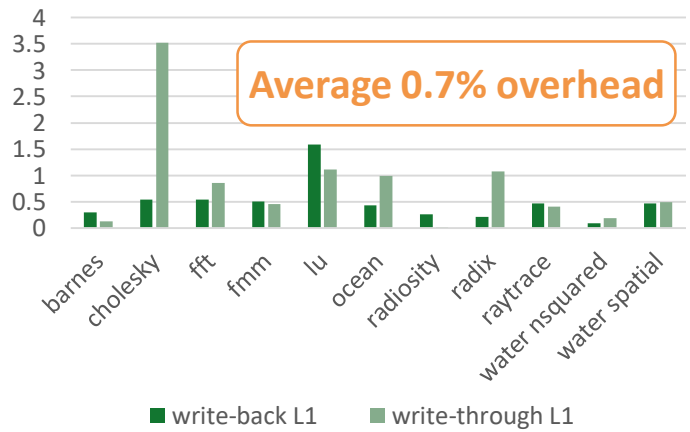

Average 0.7% overhead

Partitioning the CPU

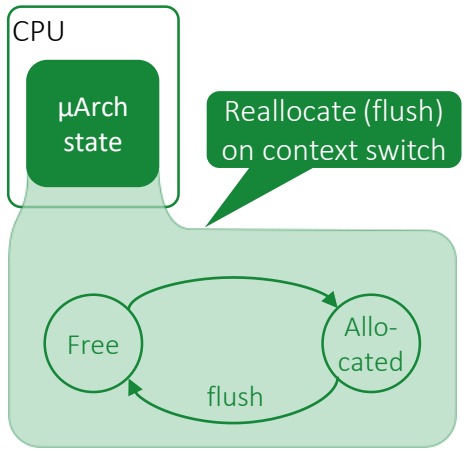

OpenC910: 64-bit, application-class, **12-stage**, **superscalar**, **out-of-order** RISC-V core

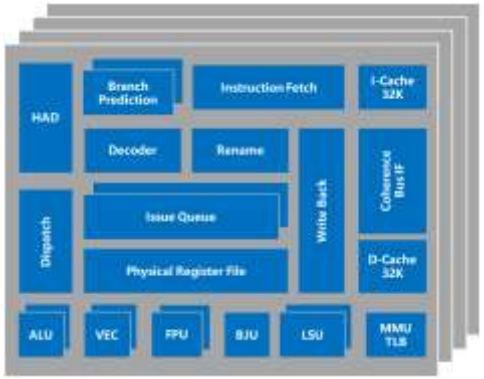

Chen et al., ISCA'20

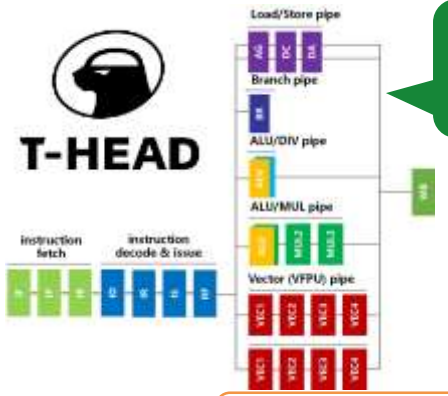

`fence.t.s` [ApplePies'24]

Negligible HW costs

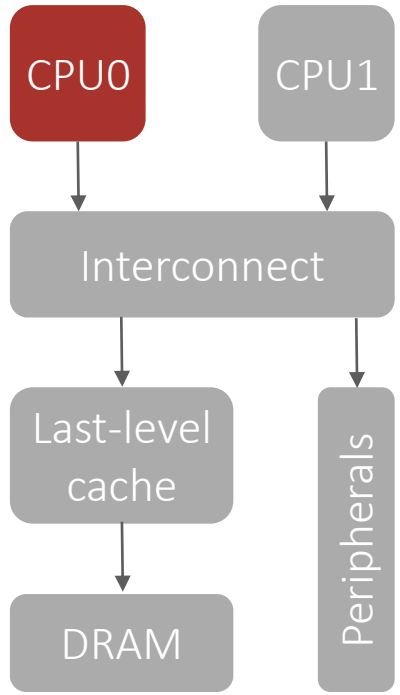

L1D Unmitigated

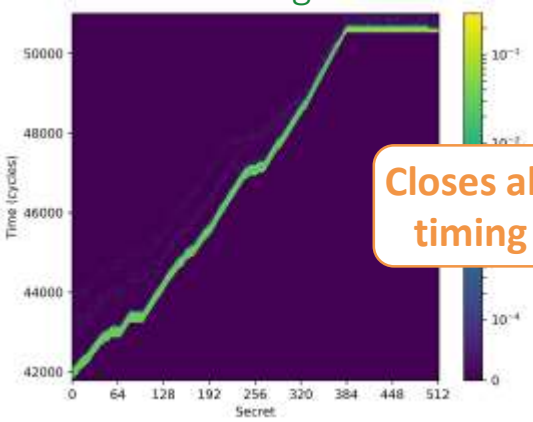

M = **4283 mb**, M_0 = 0.7 mb

L1D fence.t.s

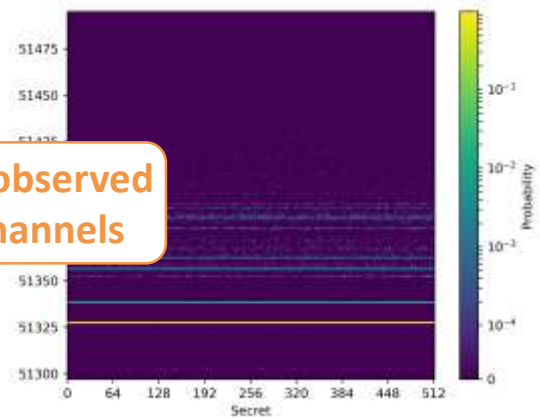

M = 64.3 mb, M_0 = 71.0 mb

Closes all observed timing channels

Splash-2 Benchmark Overhead (%)

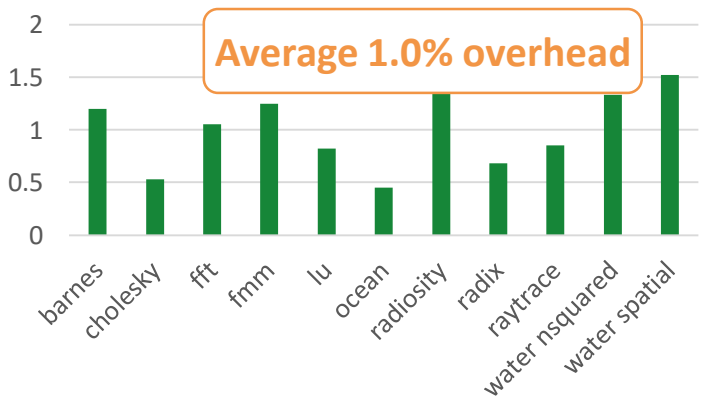

Partitioning the interconnect

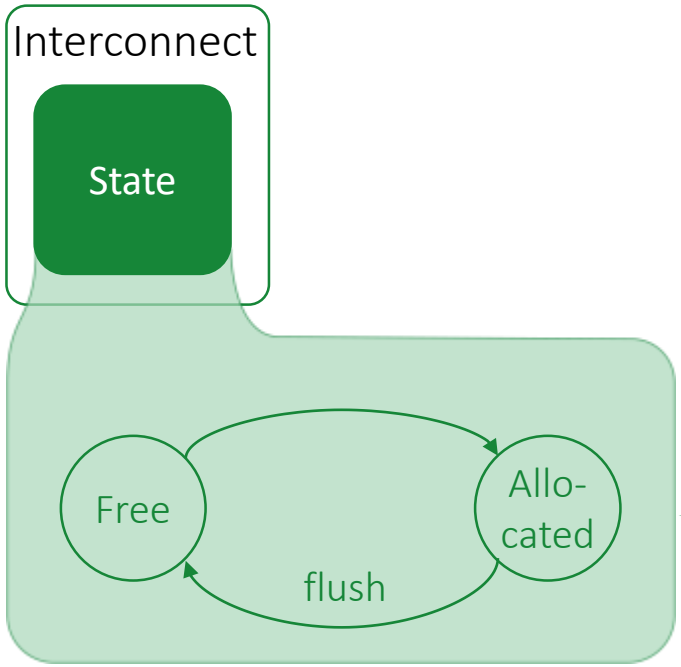

Can be avoided through coscheduling

Challenges:

- Cores compete for **bandwidth**
- Interconnect contains **state** (e.g. arbiters)

Flush on context switch

Work-in-progress

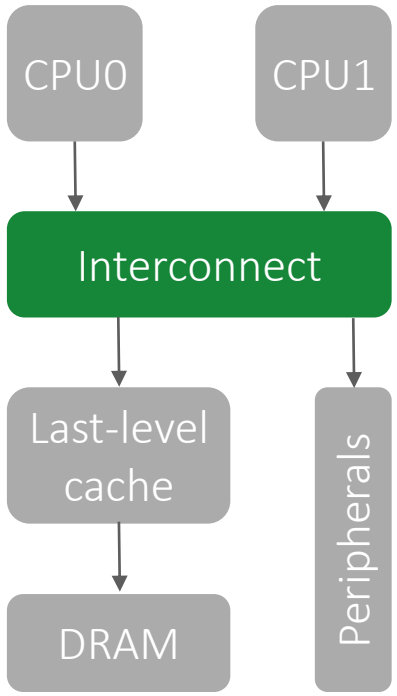

Partitioning the last-level cache

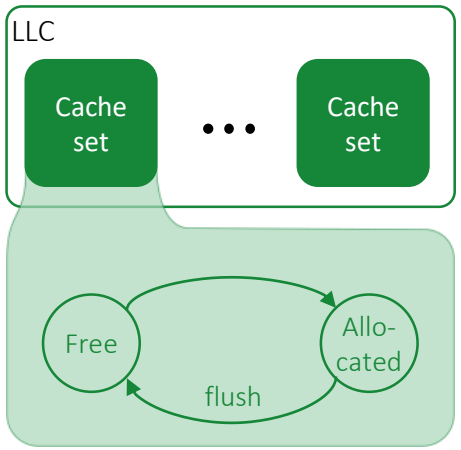

- **Cache controller** state comprises a **single capacity unit**
- HW prototype implemented, system integration/evaluation ongoing

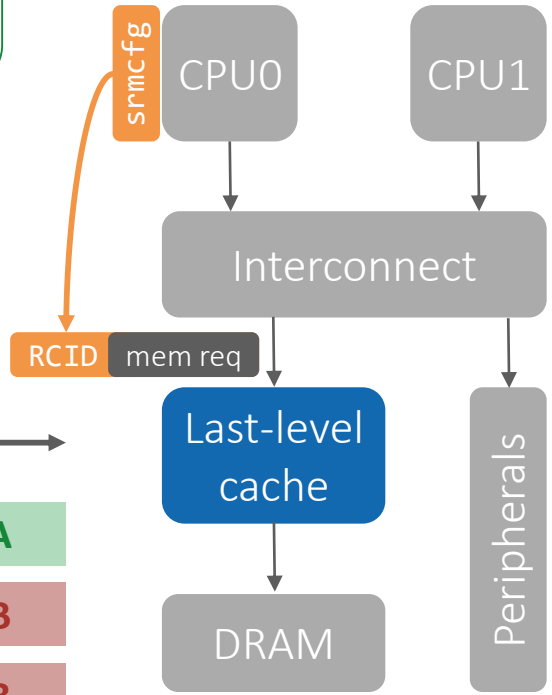

Augment set index with RCID*

Temporal partitioning

Spatial partitioning

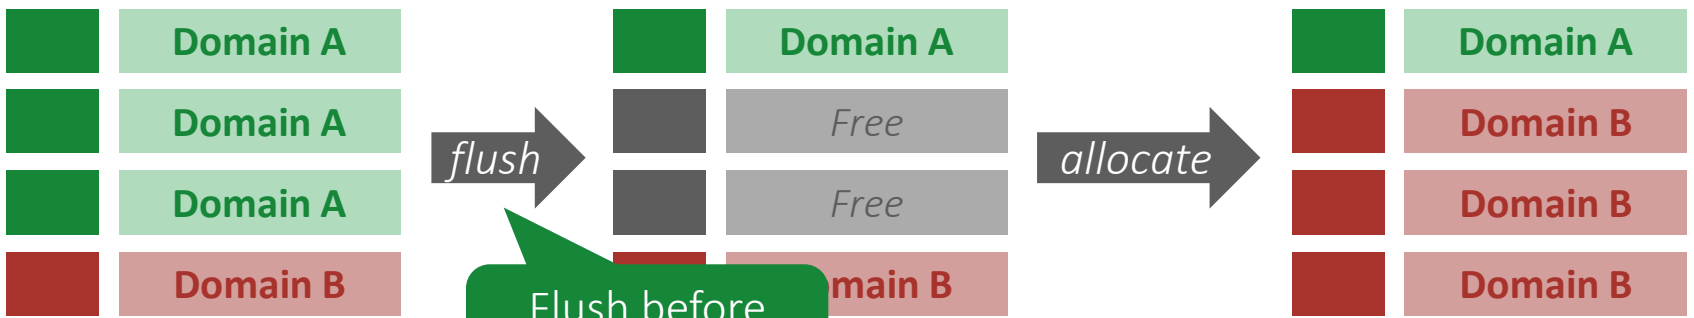

Flush before reallocating

** Terminology adopted from the RISC-V CBQRI extension*

Partitioning DRAM

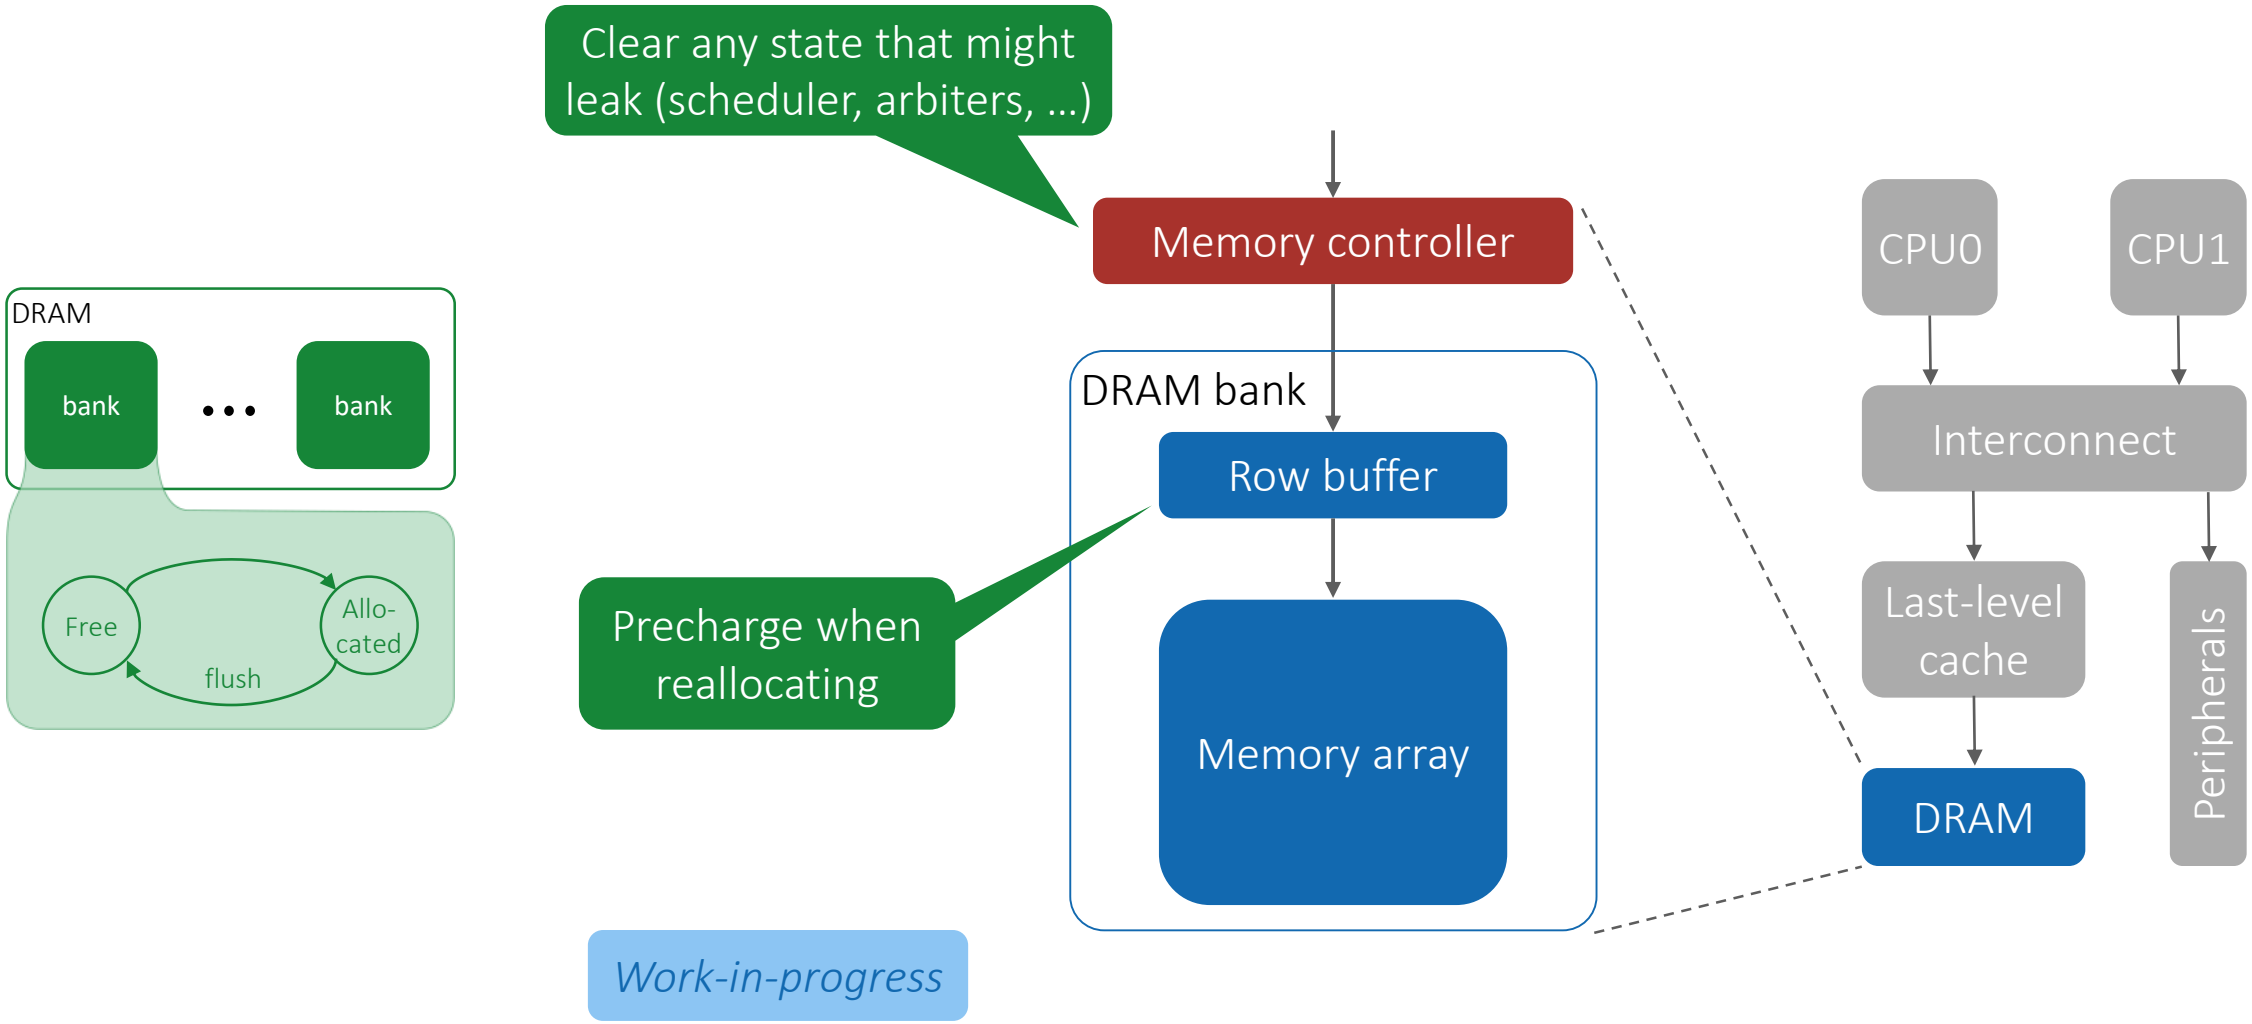

AutoCC: A formal tool to find timing channels [MICRO'23]

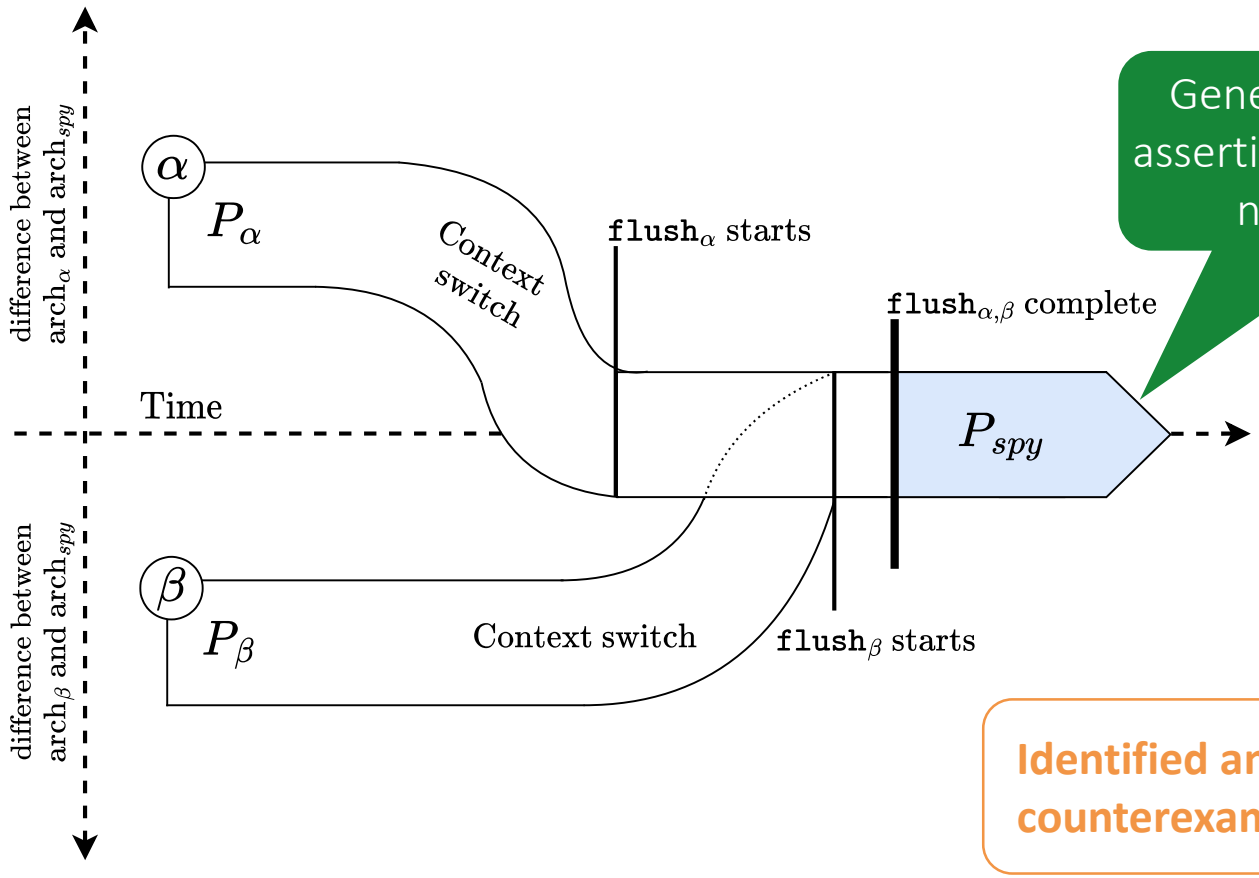

Generate SystemVerilog assertions to formally verify non-interference

Identified and fixed three counterexamples in CVA6

`github.com/morenes/AutoCC`

PRINCETON UNIVERSITY

ETH zürich

UNSW SYDNEY

Carfield: Automotive platform for MCS

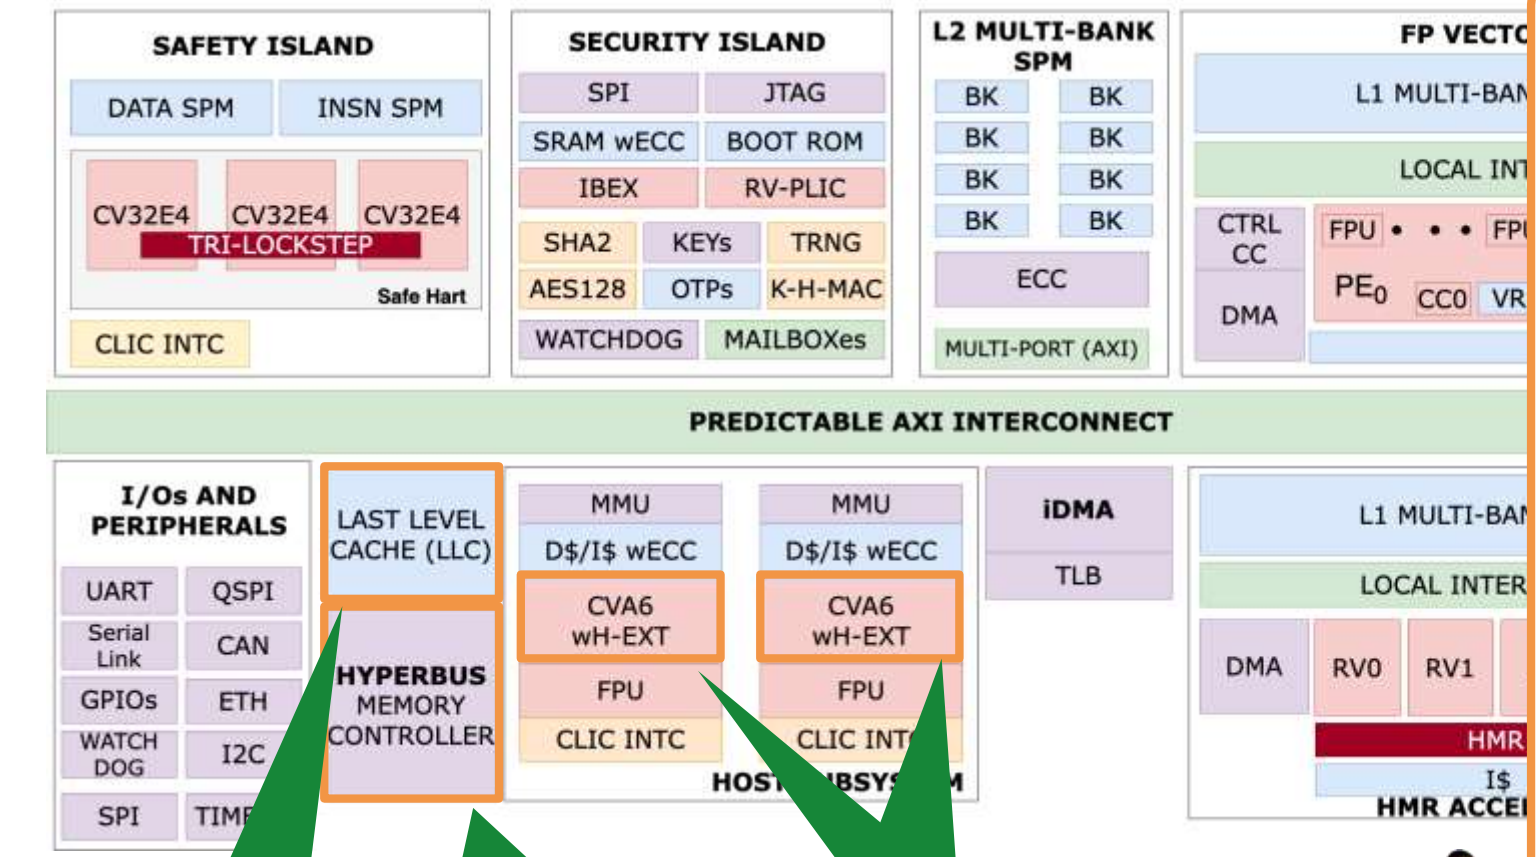

Taped-out in Intel 16

(16.0mm², 600MHz@0.8V)

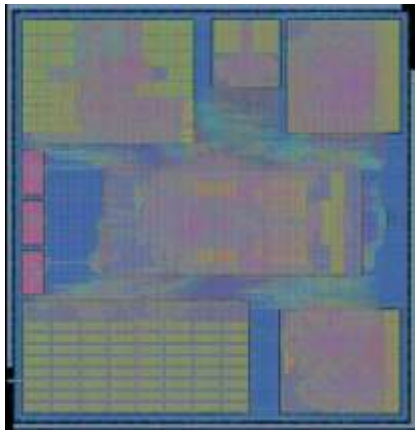

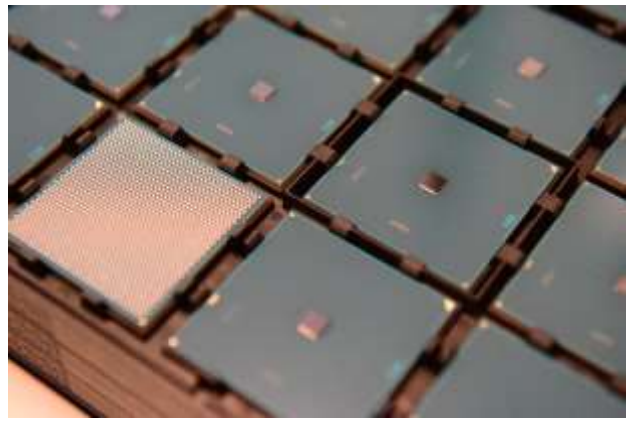

Conclusion

- Closing timing channels requires **HW/SW co-design**.
- By exposing HW partitioning mechanisms, all **on-core timing channels can be closed** at **minimal performance** impact (1.0%) and **negligible hardware overhead**.
- Proven **empirically** on seL4, using **formal** verification, and in **silicon**!
- DRAM and interconnect still under investigation
  - further need for HW/SW co-design expected.
- To be specified in RISC-V

RISC-V®

Microarchitecture Side Channels Special Interest Group
Timing Fences Task Group

PULP Platform

Open Source Hardware, the way it should be!

Nils Wistoff nwistoff@iis.ee.ethz.ch
Gernot Heiser gernot@unsw.edu.au
Luca Benini lbenini@iis.ee.ethz.ch

Institut für Integrierte Systeme – ETH Zürich
Gloriastrasse 35
Zürich, Switzerland

DEI – Università di Bologna
Viale del Risorgimento 2
Bologna, Italy

ETH zürich

ALMA MATER STUDIORUM UNIVERSITÀ DI BOLOGNA

@pulp_platform

pulp-platform.org

youtube.com/pulp_platform