

CPU Port Contention Without SMT

Thomas Rokicki - Univ Rennes, CNRS, IRISA

Clémentine Maurice - Univ Lille, CNRS, Inria

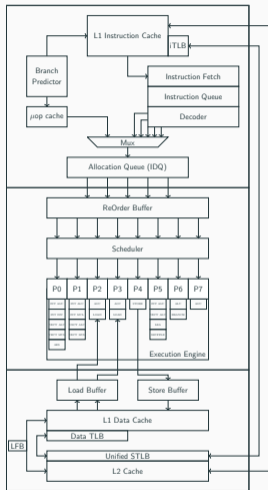
Michael Schwarz - CISPA Helmholtz Center for Information Security

ESORICS 2022 - 24/09/2022

Background: Microarchitectural attacks

Exploit subtle timing differences caused by the microarchitecture.

Cache attacks are the most famous, but most microarchitectural optimizations are targeted.

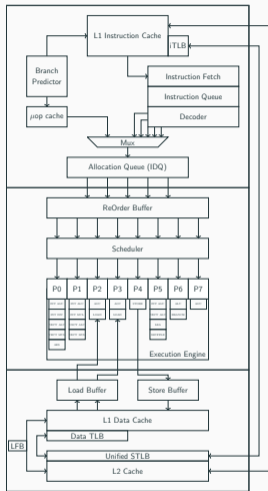



Background: Microarchitectural attacks

Exploit subtle timing differences caused by the microarchitecture.

Cache attacks are the most famous, but most microarchitectural optimizations are targeted.


Here: CPU Ports





Simultaneous computation (Also called Hyper-Threading).

- Physical cores are shared in several (often 2) logical cores.
- Abstraction at the OS level.

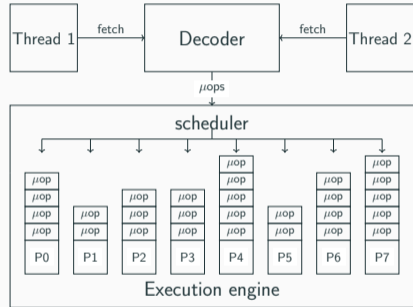


Simultaneous computation (Also called Hyper-Threading).

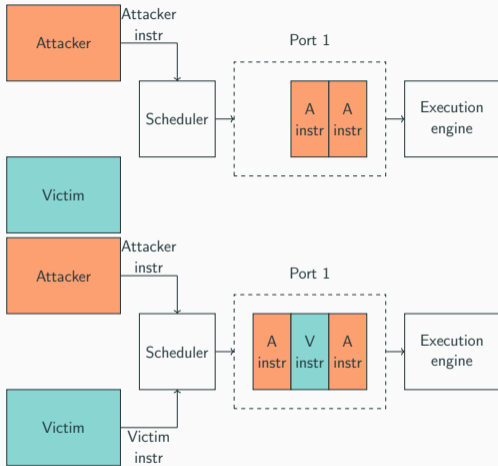
- Physical cores are shared in several (often 2) logical cores.
- Abstraction at the OS level.
- **Hardware resources are shared between logical cores.**

Background: Execution pipeline

- Instructions are decomposed in micro-operations (μops) to optimize Out-of-Order computation.
- The decomposition of instructions into μops is deterministic.
- μops are dispatched to specialized execution units through **CPU ports**.



Background: Port contention¹



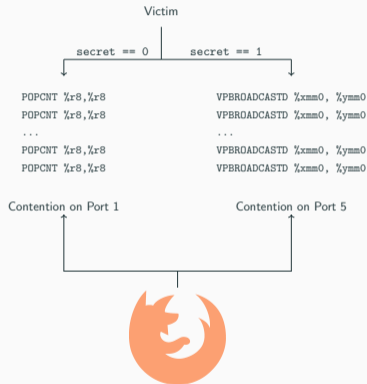
No Contention All the attacker instructions are executed in a row, **fast execution time**.

Contention Attacker instructions are delayed, **slow execution time**.

¹Aldaya et al. , Port Contention for Fun and Profit, S&P, 2019

Background: Port contention in web browsers²

- Port contention is also implementable in browsers with WebAssembly.
- Huge threat surface but restricted environment.
- We proposed a framework to determine the port usage of WebAssembly instructions.
- Spatial resolution on par with Prime+Probe.



²Rokicki et al. , Port contention goes portable, AsiaCCS, 2021

Port contention attacks rely on the attacker and victim sharing a hardware component. They are highly dependent on SMT.

Countermeasures to SMT-attacks are starting to appear:

Port contention attacks rely on the attacker and victim sharing a hardware component. They are highly dependent on SMT.

Countermeasures to SMT-attacks are starting to appear:

Disable SMT (RedHat)

Port contention attacks rely on the attacker and victim sharing a hardware component. They are highly dependent on SMT.

Countermeasures to SMT-attacks are starting to appear:
Disable SMT (RedHat) Dynamic Sharing ³.

³Taram et al. , SecSMT: Securing SMT Processors against Contention-Based Covert Channels, USENIX 22

Port contention attacks rely on the attacker and victim sharing a hardware component. They are highly dependent on SMT.

Countermeasures to SMT-attacks are starting to appear:
Disable SMT (RedHat) Dynamic Sharing ³.

Can we create port contention without SMT?

³Taram et al. , SecSMT: Securing SMT Processors against Contention-Based Covert Channels, USENIX 22

We introduce **Sequential Port Contention**.

We introduce **Sequential Port Contention**.

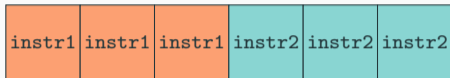
Exploit parallelism at instruction level.

We introduce **Sequential Port Contention**.


Exploit parallelism at instruction level.

Creates contention on ports and exploits it without SMT.

Port Contention Without SMT - Concept

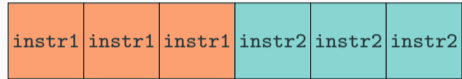


(a) Grouped



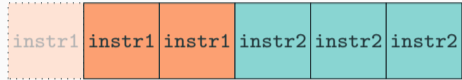
(b) Interleaved

Port Contention Without SMT - Concept




Cycle 0

Port Contention Without SMT - Concept



Cycle 1

Port Contention Without SMT - Concept



Cycle 2

Port Contention Without SMT - Concept



Cycle 3

Port Contention Without SMT - Concept



Cycle 4

Port Contention Without SMT - Concept

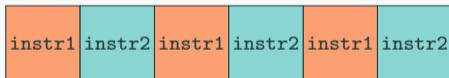


Cycle 5


Execution is never parallelized

Port Contention Without SMT - Concept

Instructions use the same ports



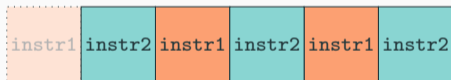
Instructions use different ports



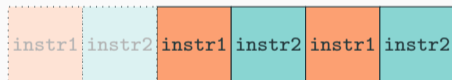
Cycle 0

Port Contention Without SMT - Concept

Instructions use the same ports



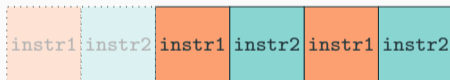
Instructions use different ports



Cycle 1

Port Contention Without SMT - Concept

Instructions use the same ports




Instructions use different ports




Cycle 2

Port Contention Without SMT - Concept

Instructions use the same ports



Instructions use different ports



Cycle 3


Port Contention Without SMT - Concept

Instructions use the same ports



The instructions create contention at the port level:
slower execution

Instructions use different ports



The execution is parallelized at port level:
faster execution




Figure 8: Execution time for grouped vs. interleaved with different ports.

We compute the ratio $\rho_{grouped/interleaved}$.

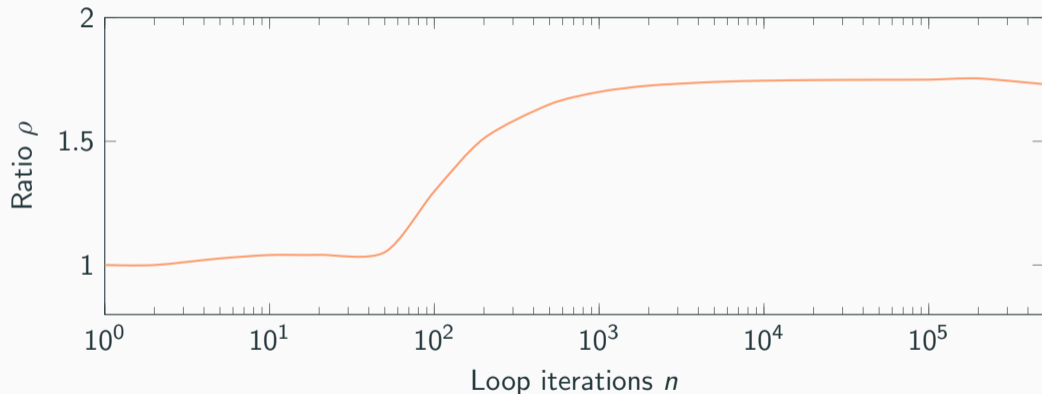




Figure 9: $\rho_{grouped/interleaved}$.

Same idea but in WebAssembly.

Same idea but in WebAssembly.



Same idea but in WebAssembly.



Ratio is lower due to lowest-resolution timers / less control on port usage.

Allows for way more portability!

- CPU generations bring changes to the microarchitecture.



Application to fingerprint - Idea

- CPU generations bring changes to the microarchitecture.
- Instructions can have different port usages between generations.



Application to fingerprint - Idea

- CPU generations bring changes to the microarchitecture.
- Instructions can have different port usages between generations.
- If we can determine the port usage of these instructions from the web, we can guess the generation!



Application to fingerprint - Idea

- CPU generations bring changes to the microarchitecture.
- Instructions can have different port usages between generations.
- If we can determine the port usage of these instructions from the web, we can guess the generation!
- Consolidate software attributes for fingerprinting⁴.



⁴Trampert et al. , Browser-based CPU Fingerprinting, Esorics 2022

We need to find **distinguishers**, *i.e.*, pairs of instructions that:

- Exhibit different contention on different generations;
- Exhibit similar contention on different CPUs of the same generation.





We need to find **distinguishers**, *i.e.*, pairs of instructions that:

- Exhibit different contention on different generations;
- Exhibit similar contention on different CPUs of the same generation.

Problem: We do not know how our WebAssembly instructions are translated.



We need to find **distinguishers**, *i.e.*, pairs of instructions that:

- Exhibit different contention on different generations;
- Exhibit similar contention on different CPUs of the same generation.

Problem: We do not know how our WebAssembly instructions are translated.

We built a **framework**, testing 458 pairs of instructions for distinguishers and found **30**.

- Once we have these distinguishers, we create *generation fingerprints*, ie the behavior of the distinguishers for a given generation.



- Once we have these distinguishers, we create *generation fingerprints*, ie the behavior of the distinguishers for a given generation.
- We use it to train a k -NN model to classify unknown CPUs.



- Once we have these distinguishers, we create *generation fingerprints*, ie the behavior of the distinguishers for a given generation.
- We use it to train a k -NN model to classify unknown CPUs.
- We created a website to get these fingerprints:
<https://fp-cpu-gen.github.io/fp-cpu-gen>
Feel free to try and send us results!



Application to fingerprint - Results

SKL	1	0	0	0	0	0
CFL	0.056	0.94	0	0	0	0
HSL	0	0	1	0	0	0
SNB	0	0	0	1	0	0
TGL	0	0.17	0	0	0.83	0
ZEN	0	0	0	0	0	1
	SKL	CFL	HSL	SNB	TGL	ZEN

- Evaluation on 50 different CPUs, spanning 13 generations.
- Includes Intel CPUs and AMD.
- 92% accuracy.
- Highly stable and resistant to noise.

Software Diversification : JavaScript engine can reorder code while keeping functionality or add fences to the output.

Performance Counters Detection : We can detect when sequential port contention occurs by measuring backend-bound execution.



- Threat surface extension for port contention.

- Threat surface extension for port contention.
- Applications to browser fingerprinting.

- Threat surface extension for port contention.
- Applications to browser fingerprinting.
- Highly resistant to noise.

- Threat surface extension for port contention.
- Applications to browser fingerprinting.
- Highly resistant to noise.
- Maybe other SMT attacks can be leveraged with instruction-level parallelism?

Questions?

Contact me here: `thomas.rokicki@irisa.fr`

Feel free to read the paper for more technical details!

Find the code here:

`https://github.com/MIAOUS-group/port-contention-without-smt`



Credits for images:

- [Vecteezy.com](https://www.vecteezy.com/)
- [Veryicon.com](https://www.veryicon.com/)