Formal methods to protect against Microarchitectural attacks

PEPR Sécurité – Winter school



Lesly-Ann Daniel – KU Leuven

Who am I?

2018–2021

Phd Student

- Symbolic Binary-Level Code Analysis for Security
- CEA List & Université Côte d'Azur
- Sébastien Bardin and Tamara Rezk

2021

Postdoc

- Hardware /Software co-Designs for Microarchitectural Security
- KU Leuven (Belgium)
- Frank Piessens

Outline

1. Microarchitectural side-channel attacks

- What are microarchitectural side-channel attacks?
- How can formal methods help to protect against them?

2. Spectre attacks

- Hardware optimizations can leak secrets!
- Model the microarchitecture with formal methods?

3. Mind the gap: model <> HW

• HW/SW contracts to help bridging the gap

PART 1

Microarchitectural side-channel attacks



How formal methods can help you protect your secrets from the vagaries of time

What are side-channels?

Programs manipulate secret data

Critical software is prevalent:

- Secure communications
- Online banking
- Protect health data



Their security relies on cryptography:

- Mathematical guarantees
- Verified implementations (no bugs, functional)
- But what about their execution in the physical world?









Timing and microarchitectural attacks can be run remotely [1]

[1] Remote Timing Attacks Are Practical, David Brumley and Dan Boneh at USENIX 2003



 $\begin{array}{c} 0000 \rightarrow 1s \\ 1000 \rightarrow 1s \\ 2000 \rightarrow 1s \\ 3000 \rightarrow 1s \\ 4000 \rightarrow 2s \\ 5000 \rightarrow 1s \end{array}$

. . .



 $\begin{array}{c} 0000 \rightarrow 1s \\ 1000 \rightarrow 1s \\ 2000 \rightarrow 1s \\ 3000 \rightarrow 1s \\ 4000 \rightarrow 2s \\ 5000 \rightarrow 1s \end{array}$

. . .



 $4000 \rightarrow 2s$ $4100 \rightarrow 2s$ $4200 \rightarrow 2s$ $4300 \rightarrow 3s$ $4400 \rightarrow 2s$ $4500 \rightarrow 2s$

...



 $4000 \rightarrow 2s$ $4100 \rightarrow 2s$ $4200 \rightarrow 2s$ $4300 \rightarrow 3s$ $4400 \rightarrow 2s$ $4500 \rightarrow 2s$

...



 $4000 \rightarrow 2s$ $4100 \rightarrow 2s$ $4200 \rightarrow 2s$ $4300 \rightarrow 3s$ $4400 \rightarrow 2s$ $4500 \rightarrow 2s$

...



Attack Complexity: from 10^4 to 10×4



Countermeasure

```
bool check_pin(char* guess) {
  good = true;
  for (i=0; i<4; i++)
   good &= guess[i] == pin[i];
  return good;
}</pre>
```

Make timing independent of secret Remove secret-dependent branch!





trace' → secret

Control-flow leaks

- end-to-end timing
- different resource consumption
- branch predictor state
- instruction cache
- instruction prefetcher
- micro-op cache
- ...

Memory accesses leak

- caches

- ...

- data pre-fetchers
- load/store dependencies

x = tab[secret]

Cache





Memory accesses leak

- caches

- ...

- data pre-fetchers
- load/store dependencies

x = tab[secret]







Memory accesses leak

- caches

- ...

- data pre-fetchers
- load/store dependencies



Variable time instructions leak

- divisions

- ...

- multiplication
- depends on microarchitecture



Why does it matter?

Timing Attacks on Implementations of Diffie-Hellman, RSA, DSS, and Other Systems

Paul C. Kocher

Cryptography Research, Inc. 607 Market Street, 5th Floor, San Francisco, CA 94105, USA. E-mail: paul@cryptography.com.

Abstract. By carefully measuring the amount of time required to perform private key operations, attackers may be able to find fixed Diffie-Hellman exponents, factor RSA keys, and break other cryptosystems. Against a vulnerable system, the attack is computationally inexpensive and often requires only known ciphertext. Actual systems are potentially at risk, including cryptographic tokens, network-based cryptosystems, and other applications where attackers can make reasonably accurate timing measurements. Techniques for preventing the attack for RSA and Diffie-Hellman are presented. Some cryptosystems will need to be revised to protect against the attack, and new protocols and algorithms may need to incorporate measures to prevent timing attacks.

Cache-timing attacks on AES

Daniel J. Bernstein *

Department of Mathematics, Statistics, and Computer Science (M/C 249) The University of Illinois at Chicago Chicago, IL 60607-7045 djb@cr.vp.to

Abstract. This paper demonstrates complete AES key recovery from known-plaintext timings of a network server on another computer. This attack should be blamed on the AES design, not on the particular AES library used by the server; it is extremely difficult to write constant-time high-speed AES software for common general-purpose computers. This paper discusses several of the obstacles in detail.

Solution? Constant-time programming!

Write programs with:

- No secret-dependent branches
- No secret-dependent memory accesses



Already used in many cryptographic implementations

Constant-time is not easy to implement



Compilers can break constant-time!



Compilers can break constant-time!



How to formally define constant-time?

Constant-time programming, formally?

Side-channel observations produced by program executions must be independent from secret input

Constant-time programming, formally?



How do we formalize program executions?

System model

Small asm
language(Values)
$$v \in \mathbb{V}$$
(Registers) $\mathbf{x} \in \mathbb{R}$ (Labels) $\ell \in \mathbb{L}$ $\langle exp \rangle ::= v \mid \mathbf{x}$
 $\langle inst \rangle ::= add $\mathbf{x} \langle exp \rangle \langle exp \rangle \mid mul \mathbf{x} \langle exp \rangle \langle exp \rangle$
 $\mid load $\mathbf{x} \langle exp \rangle \mid store \langle exp \rangle \langle exp \rangle$
 $\mid beqz \langle exp \rangle \mid \ell \mid jmp \langle exp \rangle$
(Program) $P : \mathbb{L} \rightarrow \langle inst \rangle$$$

Configurations
$$\sigma = \langle r, m \rangle$$
 where $\begin{cases} r : \mathbb{R} \to \mathbb{V} & (\text{Register map}) \\ m : \mathbb{V} \to \mathbb{V} & (\text{Memory}) \end{cases}$

System model

Expression evaluation
$$[\![e]\!]_r$$
 =

$$[\![e]\!]_r = v$$

 $\sigma \rightarrow \sigma'$ Instruction evaluation

$$\begin{array}{ll} \overset{\text{ADD}}{\underline{\ell} = r(\texttt{pc})} & P[\ell] = \texttt{add } \texttt{x} \ e_1 \ e_2 & v = \llbracket e_1 \rrbracket_r + \llbracket e_2 \rrbracket_r & r' = r[\texttt{x} \mapsto v][\texttt{pc} \mapsto \ell + 1] \\ & \langle m, r \rangle \to \langle m, r' \rangle \end{array}$$

Constant-time programming, formally?

Side-channel observations produced by program executions must be independent from secret input

How do we define side-channel observations?

Define side-channel observations

Semantics instrumented with observations

$$\sigma \xrightarrow{o} \sigma'$$
 with $o \in \mathcal{O}$ (Set of observations)

Constant-time observation mode (or leakage model)

- Program counter is observable
- Memory addresses are observable

$$\mathcal{O} = \{\bullet, \textbf{load} \ a, \textbf{store} \ a, \textbf{pc} \ \ell\}$$

Other observation modes are possible
Additions leak an atomic leakage

$$\frac{\substack{\text{ADD}}{\ell = r(\texttt{pc})} \qquad P[\ell] = \texttt{add } \texttt{x} \ e_1 \ e_2 \qquad v = \llbracket e_1 \rrbracket_r + \llbracket e_2 \rrbracket_r \qquad r' = r[\texttt{x} \mapsto v][\texttt{pc} \mapsto \ell + 1]}{\langle m, r \rangle \xrightarrow{\bullet} \langle m, r' \rangle}$$

Loads leak their address

Define side-channel observations

Control-flow instruction leak their target

$$\frac{P[r(\texttt{pc})] = \texttt{beqz} \ e \ \ell \quad \llbracket e \rrbracket_r = 0 \qquad r' = r[\texttt{pc} \mapsto \ell]}{\langle m, r \rangle \xrightarrow{pc \ \ell} \langle m, r' \rangle}$$

$$\frac{P[r(pc)] = beqz \ e \ \ell}{\langle m, r \rangle} \frac{pc \ \ell'}{\langle m, r' \rangle} \langle m, r' \rangle$$

Constant-time programming, formally?

Side-channel observations produced by program executions must

be independent from secret input

What does it mean to be independent from secret input?

Define security

Define public/secrets

Partition state into *public (low) / secret (high)* registers and memory

Low-equivalence relation $\sigma \sim_L \sigma'$

Two configurations are *low-equivalent* if they have the same public values

Define security

Definition 1 (Security). A program P is secure, if and only if: For any pair of initial configurations σ_0 , σ'_0 , if $\sigma_0 \sim_L \sigma'_0$ and $\sigma_0 \stackrel{o}{\rightarrow} {}^n \sigma_n$ then $\sigma'_0 \stackrel{o'}{\rightarrow} {}^n \sigma'_n$ and o = o'



Define security

Definition 1 (Security). A program P is secure, if and only if: For any pair of initial configurations σ_0 , σ'_0 , if $\sigma_0 \sim_L \sigma'_0$ and $\sigma_0 \stackrel{o}{\rightarrow} {}^n \sigma_n$ then $\sigma'_0 \stackrel{o'}{\rightarrow} {}^n \sigma'_n$ and o = o'

Property relating 2 execution traces (2-hypersafety)

Now how do we verify CT?

Several approaches

A Systematic Evaluation of Automated Tools for Side-Channel Vulnerabilities Detection in Cryptographic Libraries

Antoine Geimer Univ. Lille, CNRS, Inria Univ. Rennes, CNRS, IRISA Lille, France

Lesly-Ann Daniel KU Leuven, imec-DistriNet Leuven, Belgium Mathéo Vergnolle Université Paris-Saclay, CEA, List Gif-sur-Yvettes, France Frédéric Recoules Université Paris-Saclay, CEA, List Gif-sur-Yvettes, France

Sébastien Bardin Université Paris-Saclay, CEA, List Gif-sur-Yvettes, France Clémentine Maurice Univ. Lille, CNRS, Inria Lille, France

Static

- Type systems
- Abstract interpretation
- Symbolic execution

Dynamic

- Record and compare observations
- Statistical tests
- Fuzzing
- Dynamic symbolic execution

Several approaches

A Systematic Evaluation of Automated Tools for Side-Channel Vulnerabilities Detection in Cryptographic Libraries

Antoine Geimer Univ. Lille, CNRS, Inria Univ. Rennes, CNRS, IRISA Lille, France

Lesly-Ann Daniel KU Leuven, imec-DistriNet Leuven, Belgium Mathéo Vergnolle Université Paris-Saclay, CEA, List Gif-sur-Yvettes, France Frédéric Recoules Université Paris-Saclay, CEA, List Gif-sur-Yvettes, France

Sébastien Bardin Université Paris-Saclay, CEA, List Gif-sur-Yvettes, France Clémentine Maurice Univ. Lille, CNRS, Inria Lille, France

Static

- Type systems
- Abstract interpretation
- Symbolic execution

Dynamic

- Record and compare observations
- Statistical tests
- Fuzzing
- Dynamic symbolic execution

```
foo(public p, secret s) {
    c := p * s - 48;
    if(c = 0) error();
    else return s/c;
}
```

Can error be reached?

[1] James C. King. Symbolic execution and program testing, Communications of the ACM, 1976
 [2] Cristian Cadar and Sen Koushik. Symbolic execution for software testing: three decades later. Communications of the ACM, 2013
 46

```
foo (public p, secret s) {
    c := p * s - 48;
    if(c = 0) error();
    else return s/c;
}
```

Can error be reached?

[1] James C. King. Symbolic execution and program testing, Communications of the ACM, 1976 [2] Cristian Cadar and Sen Koushik. Symbolic execution for software testing: three decades later. Communications of the ACM, 2013 ⁴⁷

Symbolic store

$$\begin{array}{ccc} p & \mapsto & p \\ s & \mapsto & s \end{array}$$

```
foo(public p, secret s) {
    c := p * s - 48;
    if(c = 0) error();
    else return s/c;
}
```

Can error be reached?

Symbolic store

$$p \mapsto p$$

s $\mapsto s$
c $\mapsto p \times s - 48$

[1] James C. King. Symbolic execution and program testing, Communications of the ACM, 1976
 [2] Cristian Cadar and Sen Koushik. Symbolic execution for software testing: three decades later. Communications of the ACM, 2013
 ⁴⁸



Can error be reached?





Path predicate

[1] James C. King. Symbolic execution and program testing, Communications of the ACM, 1976

[2] Cristian Cadar and Sen Koushik. Symbolic execution for software testing: three decades later. Communications of the ACM, 2013



[1] James C. King. Symbolic execution and program testing, Communications of the ACM, 1976

[2] Cristian Cadar and Sen Koushik. Symbolic execution for software testing: three decades later. Communications of the ACM, 2013

CT is a 2-hypersafety!

Definition 1 (Security). A program P is secure, if and only if: For any pair of initial configurations σ_0 , σ'_0 , if $\sigma_0 \sim_L \sigma'_0$ and $\sigma_0 \stackrel{o}{\rightarrow} {}^n \sigma_n$ then $\sigma'_0 \stackrel{o'}{\rightarrow} {}^n \sigma'_n$ and o = o'

Property relating 2 execution traces (2-hypersafety)

Verification techniques/tools for safety do not apply

Secure Information Flow by Self-Composition*

Gilles Barthe¹ Pedro R. D'Argenio² Tamara Rezk (corresponding author) ³

Key idea: Turn a 2-hypersafety property of a program **P** to a safety property of a self-composed program **P;P'**

Can re-use verification techniques/tools for safety!

```
foo(public p, secret s) {
    c := p * s - 48;
    if(c = 0) error();
    else return s/c;
}
```

Can c = 0 depend on s?



Symbolic Execution
Formula
$$F(p, s)$$

 $c = p \times s - 48$





Symbolic Execution
Formula
$$F(p, s)$$

 $c = p \times s - 48$

$$F(p, s, p', s')$$

$$p = p' \wedge \frac{c}{c'} = p \times s - 48}{c'} \wedge c = 0 \neq c' = 0$$



Beyond self-composition: Optimization for SE

Limitations:

- Whole formula is duplicated F(p, s, p', s')
- High number of queries to the solver

Many techniques to optimize self-composed programs... Parallel SC, Product programs, Lazy SC, etc.

Beyond self-composition: Optimization for SE

Relational Symbolic Execution

Gian Pietro Farina*1, Stephen Chong^{†2} and Marco Gaboardi^{‡1}

¹University at Buffalo, SUNY ²Harvard University

- 2 execution in 1 SE instance
- Maximize sharing
- Spare queries

BINSEC/REL: Efficient Relational Symbolic Execution for Constant-Time at Binary-Level

Lesly-Ann Daniel*, Sébastien Bardin*, Tamara Rezk[†]

* CEA, List, Université Paris-Saclay, France † INRIA Sophia-Antipolis, INDES Project, France

 $les ly-ann.daniel @\,cea.fr,\ sebastien.bardin @\,cea.fr,\ tamara.rezk@inria.fr$

- RelSE for CT
- Optimization for binary-level

Formalization and theorems



And concretely?

BINSEC/REL: Efficient Relational Symbolic Execution for Constant-Time at Binary-Level

Lesly-Ann Daniel*, Sébastien Bardin*, Tamara Rezk[†]



Binsec/Rel <u>https://binsec.github.io/</u>

CT-analysis of cryptographic primitives

Preservation of constant-time by compilers

11 compiler versions

- 5 versions of clang for x86
- 5 versions of gcc for x86
- 1 version of gcc for ARM

Programs

- Analyze 34 small programs
- Total: 4148 binaries



• Optimization level O1 ... O3

Individual optimizations

Optimization setups



https://github.com/binsec/rel_bench/tree /main/properties_vs_compilers/ct

Fully reproducible build: Nix virtual env

• X86-cmov-converter, if-conversion



Clang adds secret dependent memory access

	1 sort2:
	2 esi := load (in+0)
1 void sort2(132* out, 132* in) {	3 edi := load (in+4)
2 a0 = load in[0]	4 cmp esi edi
3 a1 = load in[1]	5 edi := cmovle esi
4 a = select (a0 < a1) a0 a1	6 store (out+0) edi
5 store a out[0]	7 ecx := in+0
6 b1 = load in[1]	edx := in+4
7 b0 = load in[0]	9 edx := cmovge ecx
8	10 ecx := load edx
<pre>9 store b out[1] }</pre>	11 store (out+4) ecx
LIVM-IR	
	clang-9 –m32 –O3 –march=i686



- Constant-Time = de facto standard against microarchitectural SCA
- We can formalize CT as a 2-hypersafety
- There are tools to verify crypto primitives / find bugs
- We can find cool bugs introduced by compilers *LLVM analysis is not sufficient!*





Binsec/Rel

PART 2

Spectre Attacks



Or why is my code still leaking and what can I do about it?

Spectres are haunting our code

Spectre Attacks: Exploiting Speculative Execution

Paul Kocher¹, Jann Horn², Anders Fogh³, Daniel Genkin⁴, Daniel Gruss⁵, Werner Haas⁶, Mike Hamburg⁷, Moritz Lipp⁵, Stefan Mangard⁵, Thomas Prescher⁶, Michael Schwarz⁵, Yuval Yarom⁸



- Wrong speculation = transient executions
- Transient executions are reverted at architectural level
- But *not the microarchitectural state* (e.g. cache)



2018

Idea. Force victim to encode secret data in microarchitecture during transient execution & recover them with microarchitectural attacks

Is this code secure?

Leaks x to the microarchitectural state (e.g. load, or branch instr.)

Secure iff mysecret does not flow to leak



ISA (sequential) execution

Conditional bound check ensures idx is in bounds

x only contains public data





Actual (speculative) execution

Branch condition can be (mis)predicted



Can I exploit that to leak(mysecret) ?



 \bullet

- **Trains** branch predictor to predict true
- **Run** victim with idx = len
 - Branch is mispredicted to true



- Oob access to mysecret
- Transient execution leak(mysecret)
- **Extract mysecret** from microarch.


Many variants of Spectre



- 1. Misspeculation leads to transient execution Many sources of speculation
- 2. Transient execution leaks secret via side-channel

Many side-channel vectors (timing, caches, buffers, etc.)

Many sources of speculation \Leftrightarrow many variants of Spectre [1]

- Spectre-PHT: conditional branch
- Spectre-BTB: indirect branch
- Sprectre-RSB: return address
- Spectre-STL: memory dependencies
- etc. (see [2] for the most recent list)

[1] Canella, Claudio, et al. "A systematic evaluation of transient execution attacks and defenses." USENIX Security (2019)
 [2] Randal, Allison. "This is how you lose the transient execution war." arXiv (2023).

Countermeasures?

How to protect against Spectre?

In Software?



- Speculation barriers (fence)
- Load hardening
- Retpolines
- etc.

Full software solution
Variant-specific
Can be costly

In Hardware?



Microarchitectural partitioning, Invisible speculation, OISA, STT, SPT, ConTExT, etc.

Better performance
Comprehensive (but not always)
Adoption is harder

Fences to block speculative execution

Branch is mispredicted to true



- **fence** stalls until branch is resolved
- Rollback before leak(mysecret)

Transiently execution only until fence

Speculative Load Hardening



Branch is mispredicted to true



- Oob access to mysecret
- x = 0 if branch is mispredicted

leak(0) 🗸

But where should I insert protections? Don't worry kid, formal methods can help you!

Speculative Constant-Time (SCT)

Constant-Time Foundations for the New Spectre Era

Sunjay Cauligi[†] Craig Disselkoen[†] Klaus v. Gleissenthall[†] Dean Tullsen[†] Deian Stefan[†] Tamara Rezk^{*} Gilles Barthe^{**}

[†]UC San Diego, USA *****INRIA Sophia Antipolis, France *****MPI for Security and Privacy, Germany *****IMDEA Software Institute, Spain

Idea: Security in the constant-time observation mode on a *speculative semantics*

Many flavors of microarchitectural semantics / ways to define security (see [1])

[1] Cauligi, S., Disselkoen, C., Moghimi, D., Barthe, G., & Stefan, D. (2022, May). SoK: Practical foundations for software Spectre defenses. *SP'22*

Why is that hard?

Problem. Formalize microarchitectural semantics with predictions and out-of-order execution

Challenge. Microarchitectural features are complex, often undocumented

Goals. Find suitable abstraction to reason about Spectre

- Capture all variants of Spectre
- Keep it simple

First, how does my microarchitecture work?

Fetch/Decode

- In-order
- Get instruction from memory
- Decode instruction
- Fill reorder buffer (ROB)
- Predict branches



- Out-of-order
- Execute instructions from ROB
- Depends on available operands/execution units
- Rollback incorrect pred.

- In-order
- Commits oldest instruction from ROB

Retire

• Write result in register file/memory







84















Reorder Buffer (ROB)











Now, how do we formalize that?

Small asm
language(Values)
$$v \in \mathbb{V}$$
(Registers) $x \in \mathbb{R}$ (Labels) $\ell \in \mathbb{L}$
 $\langle exp \rangle ::= v \mid x$
 $\langle inst \rangle ::= add x \langle exp \rangle \langle exp \rangle \mid mul x \langle exp \rangle \langle exp \rangle$
 $\mid load x \langle exp \rangle \mid store \langle exp \rangle \langle exp \rangle$
 $\mid beqz \langle exp \rangle \ell \mid jmp \langle exp \rangle \mid fence$
(Program) $P : \mathbb{L} \rightarrow \langle inst \rangle$

Configurations
$$\sigma = \langle r, m, buf \rangle$$
 where $\begin{cases} r : \mathbb{R} \to \mathbb{V} & (\text{Register map}) \\ m : \mathbb{V} \to \mathbb{V} & (\text{Memory}) \\ buf : \langle inst_{rob} \rangle \ list & (\text{Reorder buffer}) \end{cases}$

Semantics instrumented with observations and attacker directives

$$\sigma \xrightarrow[d]{o} \sigma' \text{ with } \begin{cases} o \in \mathcal{O} & \text{(Observation)} \\ d \in \mathcal{D} & \text{(Directive)} \end{cases}$$

Attacker directives

Model attacker ability to influence scheduling / predictions

 $\mathcal{D} = \{ \text{fetch}, \text{execute } i, \text{retire} \}$

Semantics instrumented with observations and attacker directives

$$\sigma \xrightarrow[d]{o} \sigma' \text{ with } \begin{cases} o \in \mathcal{O} & \text{(Observation)} \\ \mathbf{d} \in \mathcal{D} & \text{(Directive)} \end{cases}$$

Constant-time observation mode (or leakage model)

- Program counter is observable (also commit and rollback)
- Memory addresses are observable

$$\mathcal{O} = \{\bullet, \text{load } a, \text{store } a, pc \ \ell, \text{commit}, \text{rollback}\}$$

Example: add instruction

$$\begin{array}{l} \mbox{FETCH-ADD} \\ \ell = \llbracket {\tt pc} \rrbracket_{apl(buf,r)} & P[\ell] = {\tt add} \ge e_1 \ e_2 & buf' = buf \cdot ({\tt add} \ge e_1 \ e_2 @ \varepsilon) \cdot ({\tt pc} \leftarrow \ell + 1 @ \varepsilon) \\ \\ \hline & \langle m,r,buf \rangle \xrightarrow[fetch]{} \langle m,r,buf' \rangle \end{array}$$

$$\frac{|buf| = i \quad \text{fence } \notin buf \quad inst = \text{add } x \ e_1 \ e_2 @\varepsilon \quad r' = apl(buf, r) \quad v = \llbracket e_1 \rrbracket_{r'} + \llbracket e_2 \rrbracket_{r'} \quad inst' = x \leftarrow v @\varepsilon}{\langle m, r, buf \cdot inst \cdot buf' \rangle}$$

$$\frac{\langle m, r, buf \cdot inst \cdot buf' \rangle}{\overset{\bullet}{\text{execute } i}} \langle m, r, buf \cdot inst' \cdot buf' \rangle$$

$$RETIRE$$

$$\frac{inst = \mathbf{x} \leftarrow v @\varepsilon}{\langle m, r, inst \cdot buf' \rangle} \xrightarrow[retire]{retire} r' = r[\mathbf{x} \mapsto v]$$

Example: branches

$$\begin{array}{c|c} \mbox{FETCH-BRANCH-TAKEN} \\ \hline \ell = \llbracket {\tt pc} \rrbracket_{apl(buf,r)} & P[\ell] = {\tt beqz} \ e \ \ell' & buf' = buf \cdot ({\tt pc} \leftarrow \ell' @ \ell) \\ \hline & \langle m,r,buf \rangle \xrightarrow[fetch \ true]{} \langle m,r,buf' \rangle \end{array}$$

$$\begin{array}{l} \begin{array}{l} \text{EXECUTE-COMMIT-BRANCH-TAKEN} \\ |buf| = i & \texttt{fence} \not\in buf & inst = \texttt{pc} \leftarrow \ell' @\ell & P[\ell] = \texttt{beqz} \ e \ \ell' & \llbracket e \rrbracket_{apl(buf,r)} = 0 & inst' = \texttt{pc} \leftarrow \ell' @\varepsilon \\ \\ & & \langle m, r, buf \cdot inst \cdot buf' \rangle \xrightarrow[execute i]{} \langle m, r, buf \cdot inst' \cdot buf' \rangle \end{array}$$

Define security



Now how do we verify SCT?

Modelling speculative semantics

Litmus tests (328 instrutions):

- Sequential semantics
 → 14 paths
- Speculative semantics
 → 37M paths



Modelling all transient paths *explicitly* is intractable We need to be smarter

ReISE for architectural semantics



ReISE for Spectre-PHT (naive)



RelSE for Spectre-PHT (but let's be smarter)



And concretely?
Verify/optimize Spectre protections

- Find gadgets in crypto [1,2]
- Find attacks combining Spectre variants [2,3]
- Insert Spectre protections smartly [4,5]
- Type system to protect crypto against Spectre [5]
- Find gadgets in the Linux kernel [6]

[1] Cauligi, Sunjay, et al. "Constant-time foundations for the new spectre era." PLDI'20

[2] Daniel, Lesly-Ann, Sébastien Bardin, and Tamara Rezk. "Hunting the haunter-efficient relational symbolic execution for spectre with haunted relse." *NDSS' 21*

[3] Fabian, Xaver, Marco Guarnieri, and Marco Patrignani. "Automatic Detection of Speculative Execution Combinations." *CCS'22*

[4] Vassena, Marco, et al. "Automatically eliminating speculative leaks from cryptographic code with blade." POPL'21

[5] Shivakumar, Basavesh Ammanaghatta, et al. "Typing High-Speed Cryptography against Spectre v1." SP'23

[6] Johannesmeyer, Brian, et al. "Kasper: scanning for generalized transient execution gadgets in the linux kernel." NDSS'22

Recap

• Constant-Time is vulnerable against Spectre



• New programming model: SCT Speculative ooo semantics

$$\sigma \xrightarrow[d]{o} \sigma' \text{ with } \begin{cases} o \in \mathcal{O} & \text{(Observation)} \\ d \in \mathcal{D} & \text{(Directive)} \end{cases}$$

- Harder: need clever tricks to avoid complexity
- Yet, formal methods can help optimizing protections and detect bugs!



PART 3

Fill the gap between models and hardware



Or how can we get sound hardware abstractions that can be leveraged by software?

First problem: gap model <> HW



Software Security Property (e.g. CT/SCT)



Second problem: many HW defenses



Second problem: how do we know they work?



Hardware-Software Contracts for Secure Speculation

Marco Guarnieri^{*}, Boris Köpf[†], Jan Reineke[‡], and Pepe Vila^{*} **IMDEA Software Institute* [†]*Microsoft Research* [‡]*Saarland University*



HW/SW contracts for side-channel-free programs

Definition. Contracts specify which program execution a side-channel adversary can distinguish

Goals.

- Capture security guarantees of hardware defenses
- Abstracts away hardware details
- Distribute security obligations between software/hardware
- Basis for secure programming

Contract. labeled deterministic semantics

Define a trace of observation produced during execution

Observer mode

- Constant-time (ct)
 - Control-flow + memory accesses
- Architectural observer (arch)
 - Leaks values of loads

$$\sigma_0 \stackrel{l_1}{\rightharpoonup} \sigma_1 \stackrel{l_2}{\frown} \dots \stackrel{l_n}{\frown} \sigma_n$$

$$\llbracket p \rrbracket (\sigma_0) = l_1 l_2 \dots l_n$$

Execution mode

- Sequential (seq)
 - In-order execution
- Speculative (spec)
 - Always mispredict branches

Hardware world



Adversary Model

Projections of μ

Hardware observation trace

 $\{p\}(\sigma)$

Close the gap HW <> contract

Definition 1 ($\{\cdot\} \vdash [\![\cdot]\!]$). A hardware semantics $\{\cdot\}$ satisfies a contract $[\![\cdot]\!]$ if, for all programs p and all initial architectural states σ, σ' , if $[\![p]\!](\sigma) = [\![p]\!](\sigma')$, then $\{\![p]\!](\sigma) = \{\![p]\!](\sigma')$.

States are indistinguishable in contract semantics Then they should be indistinguishable on HW

End-to-end guarantees

Program noninterference w.r.t to contract

Definition 3 $(p \vdash NI(\pi, \llbracket \cdot \rrbracket))$. Program *p* is *non-interferent* w.r.t. contract $\llbracket \cdot \rrbracket$ and policy π if for all initial architectural states $\sigma, \sigma': \sigma \simeq_{\pi} \sigma' \Rightarrow \llbracket p \rrbracket(\sigma) = \llbracket p \rrbracket(\sigma')$.

Proposition 2. If $p \vdash NI(\pi, \llbracket \cdot \rrbracket)$ and $\{ \cdot \} \vdash \llbracket \cdot \rrbracket$, then $p \vdash NI(\pi, \{ \cdot \})$.

Program security w.r.t. contract gives HW-security on any HW satisfying the contract

And concretely?

Formally study HW countermeasures

- seq: disable all speculation
- loadDelay: delaying all speculative loads
- tt: taint speculative load values and delay computations



Comparison of hardware countermeasures

PROSPECT: Provably Secure Speculation for the Constant-Time Policy

Lesly-Ann Daniel¹, Marton Bognar¹, Job Noorman¹, Sébastien Bardin², Tamara Rezk³ and Frank Piessens¹

¹imec-DistriNet, KU Leuven, 3001 Leuven, Belgium ²CEA, List, Université Paris Saclay, France ³INRIA, Université Côte d'Azur, Sophia Antipolis, France

- Track and protect secrets during speculative execution
- CT program in ISA semantics \Rightarrow secure on HW semantics
- Proof based on contract framework

Revizor: Testing Black-Box CPUs against Speculation Contracts

Oleksii Oleksenko* Hide and Seek with Spectres: Efficient discovery of Boris Köpf Microsoft Research **Christof Fetzer** Cambridge, UK TU Dresden speculative information leaks with random testing Dresden, Germany Oleksii Oleksenko Marco Guarnieri Boris Köpf Mark Silberstein Microsoft Research IMDEA Software Institute Microsoft Research Technion

- Test CPU against contracts
 - Generate pairs of programs indistinguishable wrt. contract
 - Execute them on CPU, check if they differ
- Rediscover existing Spectre variants
- Discover two new variants
 - Zero-dividend-injection
 - String-comparison overrun (repe, repne)

Specification and Verification of Side-channel Security for Open-source Processors via Leakage Contracts		
Zilong Wang	Gideon Mohr	Klaus von Gleissenthall
IMDEA Software Institute	Saarland University	Vrije Universiteit Amsterdam
Jan Reine	eke Marco	o Guarnieri
Saarland Univ	versity IMDEA So	oftware Institute

- Verify RTL processor designs against contract (ISA level)
- Applied on 3 RISC-V processors leaking CF, MEM, variable-time instr.
- Small in-order processors, no speculative execution

Contract-Aware Secure Compilation

Marco Guarnieri IMDEA Software Institute

Marco Patrignani Stanford University CISPA Helmholz Center for Information Security

- Source code shouldn't be tailored to specific HW guarantees
- Contract-Aware Secure COmpilation (CASCO)
 - Compiler parametric wrt. HW/SW contract
 - Make compilers aware of HW security guarantees
 - Leverage these to produce secure code
- (Still theoretical)

Recap

- Gap between model and hardware
- Hard to reason about HW defenses
- Contract can help formalizing HW leakage and guarantees
- Strong formal basis to reduce the gap! With already strong concrete results







Conclusion

- Concrete HW execution leak information
 - HW optimizations do not care for security
- Formal methods can help
 - Formalize observations & define secure programming models
 - Find bugs / prove that SW is secure
- Still a gap between HW-models
 - HW-SW contracts can help reduce it!
 - Opens exciting research directions!





Backup

Credits











Beyond self-composition: Optimization for symbolic execution

```
foo(public p, secret s) {
    c := p * s - 48;
    if(c = 0) error();
    else return s/c;
}
```



```
foo(public p, secret s) {
    c := p * s - 48;
    if(c = 0) error();
    else return s/c;
}
```







```
foo(public p, secret s) {
    c := p * s - 48;
    if(c = 0) error();
    else return s/c;
}
```



Relational formula:
$$F(p, s, s')$$

 $c = p \times s - 48$
 $c' = p \times s' - 48 \wedge c = 0 \neq c' = 0$
Sharing

```
foo(public p, secret s) {
    c := p * s - 48;
    if(c = 0) error();
    else return s/c;
}
```



Relational formula:
$$F(p, s, s')$$

 $c = p \times s - 48$
 $c' = p \times s' - 48 \wedge c = 0 \neq c' = 0$
Sharing



```
foo(public p, secret s) {
    c := p - 48;
    if(c = 0) error();
    else return s/c;
}
```

Symbolic store

p
$$\mapsto$$

s $\mapsto < s \mid s' >$
c \mapsto

[1] "Shadow of a doubt", Palikareva, Kuchta, and Cadar 2016[2] "Relational Symbolic Execution", Farina, Chong, and Gaboardi 2017

Better approach: Relational SE



Symbolic store

$$p \mapsto$$

s
$$\mapsto < s \mid s' >$$

c
$$\mapsto$$

[1] "Shadow of a doubt", Palikareva, Kuchta, and Cadar 2016

[2] "Relational Symbolic Execution", Farina, Chong, and Gaboardi 2017

Better approach: Relational SE



[1] "Shadow of a doubt", Palikareva, Kuchta, and Cadar 2016

[2] "Relational Symbolic Execution", Farina, Chong, and Gaboardi 2017



Spectre-STL: Loads can speculatively bypass prior stores

Sequential execution



Spectre-STL: Loads can speculatively bypass prior stores

Sequential execution + Transient Executions



Spectre-STL: Loads can speculatively bypass prior stores

Sequential execution + Transient Executions



Spectre-STL: Loads can speculatively bypass prior stores

Sequential execution + Transient Executions


ReISE for architectural semantics



ReISE for Spectre-PHT (naive)



ReISE for Spectre-STL (but let's be smarter)



147

RelSE for Spectre-STL (but let's be smarter)

