Hardware-software security contracts

Principled foundations for building secure microarchitectures

Marco Guarnieri IMDEA Software Institute

Based on joint work with Boris Köpf @ Microsoft Research Jan Reineke @ Saarland University José F. Morales, Pepe Vila, Andrés Sánchez @ IMDEA Software Institute Marco Patrignani @ CISPA

Contacts: marco.guarnieri@imdea.org @MarcoGuarnier1



institute software



Security Programming languages • Formal methods





Outline

- Introduction
- Speculative execution attacks
- Hardware countermeasures and limitations
- Hardware/software security contracts for secure speculation
- Next steps & challenges











Attacks exploit µarchitectural side-effects invisible at ISA level



FORESHADOW



ISA: Benefits

High-level language



Instruction set architecture (ISA)

Microarchitecture





ISA: Benefits

High-level language



Instruction set architecture (ISA)

Microarchitecture



Can program independently of microarchitecture





ISA: Benefits

High-level language



Instruction set architecture (ISA)

Microarchitecture



Can program independently of microarchitecture



Can implement **arbitrary** optimizations as long as ISA semantics are obeyed



High-level language



Instruction set architecture (ISA)

Microarchitecture





High-level language



Instruction set architecture (ISA)

Microarchitecture





High-level language



Instruction set architecture (ISA)

Microarchitecture



No guarantees about timing and side channels

Can implement arbitrary insecure optimizations as long as ISA semantics are obeyed



High-level language



Instruction set architecture (ISA)



No guarantees about timing and side channels

Can implement arbitrary insecure optimizations as long as ISA semantics are obeyed



High-level language



Instruction set architecture (ISA)



Impossible to program securely cryptographic algorithms? sandboxing untrusted code?

No guarantees about timing and side channels

Can implement arbitrary insecure optimizations as long as ISA semantics are obeyed





Hw-Sw contract = ISA + X

Succinctly captures possible information leakage



Can program securely on top contract independently of microarchitecture

Hw-Sw contract = ISA + X

Succinctly captures possible information leakage



Can program securely on top contract independently of microarchitecture

Hw-Sw contract = ISA + X

as long as contract is obeyed

Succinctly captures possible information leakage

Can implement arbitrary insecure optimizations





In this talk Proof-of-concept of HW/SW contracts for secure speculation



Speculative execution attacks





Almost *all* modern *CPUs* are *affected*

P. Kocher, J. Horn, A. Fogh, D. Genkin, D. Gruss, W. Haas, M. Hamburg, M. Lipp, S. Mangard, T. Prescher, M. Schwarz, Y. Yarom — Spectre Attacks:



Size of array A if (x < A size) y = B[A[x]]

Size of array A

if $(\mathbf{x} < \mathbf{A} \text{ size})$ $y = \mathbf{B}[\mathbf{A}[\mathbf{x}]]$

Size of array A

if $(\mathbf{x} < \mathbf{A} \text{ size})$



Size of array A

if (x < A size) y = B[A[x]]

Prediction based on **branch** history & program structure





Size of array A if (x < A size) y = B[A[x]]

Prediction based on **branch** history & program structure









Size of array A

if (x < A size) y = B[A[x]]

Wrong predicton? Rollback changes!

Architectural (ISA) state

Microarchitectural state

Prediction based on **branch** history & program structure









void f(int x) if (x < A_size) y = B[A[x]]</pre>

void f(int x) if (x < A_size) y = B[A[x]]</pre>



void f(int x) if (x < A size) y = B[A[x]]</pre>



















2) Prepare cache







2) Prepare cache

3) Run with x = 128







2) Prepare cache

3) Run with x = 128







2) Prepare cache

3) Run with x = 128






1) Train branch predictor

2) Prepare cache

3) Run with x = 128







1) Train branch predictor

2) Prepare cache

3) Run with *x* = 128







1) Train branch predictor

2) Prepare cache

3) Run with x = 128

4) Extract from cache





Countermeasures

Hardware countermeasures InvisiSpec: Making Speculative Execution Invisible in the Cache Hierarchy

Mengjia Yan[†], Jiho Choi[†], Dimitrios Skarlatos, Adam Morrison^{*}, Christopher W. Fletcher, and Josep Torrellas University of Illinois at Urbana-Champaign *Tel Aviv University {myan8, jchoi42, skarlat2}@illinois.edu, mad@cs.tau.ac.il, {cwfletch, torrella}@illinois.edu

Hardware countermeasures InvisiSpec: Making Speculative Execution Invisible in the Cache Hierarchy

Mengjia Yan[†], Jiho Choi[†], Dimitrios Skarlatos, Adam Morrison^{*}, Christopher W. Fletcher, and Josep Torrellas University of Illinois at Urbana-Champaign *Tel Aviv University {myan8, jchoi42, skarlat2}@illinois.edu, mad@cs.tau.ac.il, {cwfletch, torrella}@illinois.edu

CleanupSpec: An "Undo" Approach to Safe Speculation Moinuddin K. Qureshi moin@gaiccu.cua Georgia Institute of Technology 14

Hardware countermeasures InvisiSpec: Making Speculative Execution Invisible in the Cache Hierarchy

Mengjia Yan[†], Jiho Choi[†], Dimitrios Skarlatos, Adam Morrison^{*}, Christopher W. Fletcher, and Josep Torrellas University of Illinois at Urbana-Champaign *Tel Aviv University {myan8, jchoi42, skarlat2}@illinois.edu, mad@cs.tau.ac.il, {cwfletch, torrella}@illinois.edu

CleanupSpec: An "Undo" Approach to Safe Speculation Moinuddin K. Qureshi moin@gaiccu.ccc Georgia Institute of Technology 14





NDA: Preventing Speculative Execution Attacks at Their Source Efficient Invisible Speculative Execution through Norwegian University of Science and stefanos.kaxiras@it.uu.se Trondheim, Norway magnus.sjalander@ntnu.no Christos Sakalis Alexandra Jimborean Uppsala University Uppsala, Sweden Uppsala University hristos.sakalis@it.uu.se alexandra.jimborean@it.uu.se





if (x < A_size) y = A[x] z = B[y] end





if (x < A_size) y = A[x] z = B[y] end

1. if (x < A size)2. $\mathbf{y} = \mathbf{A}[\mathbf{x}]$ z = B[y]3. end 4.



Delay loads until they can be retired [Sakalis et al., ISCA'19]

Delay loads until they cannot be squashed [Sakalis et al., ISCA'19]





1. if (x < A size) $y = \mathbf{A}[\mathbf{x}]$ 2. z = B[y]3. end 4.



Delay loads until they can be retired [Sakalis et al., ISCA'19]

Delay loads until they cannot be squashed [Sakalis et al., ISCA'19]

Taint speculatively loaded data + delay tainted loads [STT and NDA, MICRO'19]







y = A[x] if (x < A_size) z = B[y] end

1. $\mathbf{y} = \mathbf{A}[\mathbf{x}]$ 2. if $(\mathbf{x} < \mathbf{A} \text{ size})$ $\boldsymbol{z} = \boldsymbol{B}[\boldsymbol{y}]$ 3. 4. end



Delay loads until they can be retired [Sakalis et al., ISCA'19]

Delay loads until they cannot be squashed [Sakalis et al., ISCA'19]





1. $\mathbf{y} = \mathbf{A}[\mathbf{x}]$ 2. if $(\mathbf{x} < \mathbf{A} \text{ size})$ z = B[y]3. end 4.



Delay loads until they can be retired [Sakalis et al., ISCA'19]

Delay loads until they cannot be squashed [Sakalis et al., ISCA'19]

Taint speculatively loaded data + delay tainted loads [STT and NDA, MICRO'19]











torrella@illinois.edu

How can we capture security guarantees?







Hardware/software security contracts

Guarnieri, Köpf, Reineke, Vila — Hardware-software contracts for secure speculation — IEEE S&P 2021 https://arxiv.org/abs/2006.03841 19



Contracts for side-channel-free systems

Contracts for side-channel-free systems

Contracts specify which **program executions** a side-channel adversary can distinguish



Contracts for side-channel-free systems

• Capture security guarantees of HW

• Basis for secure programming

Contracts specify which **program executions** a side-channel adversary can distinguish

Goals



Contracts

Contracts

SA extended with observations







Contract traces: $[\equiv](p, \sigma)$

- 21



Contract traces: $[m](p, \sigma)$







Contract traces: $[\mathbb{E}](p, \sigma)$



Hardware traces: (p, σ)





Contract traces: $[\mathbb{E}](p, \sigma)$

Attacker observes sequences of *µarch states*



Hardware traces: (p, σ)











Contracts for secure speculation
Contract = Execution Mode · Observer Mode



Contract = Execution Mode · Observer Mode

How are programs executed?



Contract = Execution Mode · Observer Mode

How are programs executed?

What is visible about the execution?



Contract = Execution Mode · Observer Mode

seq — sequential executionspec — mispredict branch instrs.





Contract = Execution Mode · Observer Mode

pc — only program counter arch — ct + loaded values

ct — pc + addr. of loads and stores











Leaks all data accessed nonspeculatively

Seq-arch

Spec-arch

Leaks "everything"



Leaks all data accessed nonspeculatively

Seq-arch

Spec-arch

Leaks "everything"



1. if (x < A_size) 2. $\mathbf{y} = \mathbf{A}[\mathbf{x}]$ 3. $\boldsymbol{z} = \boldsymbol{B}[\boldsymbol{y}]$

4. end







1. if (x < A_size) 2. $\mathbf{y} = \mathbf{A}[\mathbf{x}]$ 3. $\boldsymbol{z} = \boldsymbol{B}[\boldsymbol{y}]$ 4. end

x < A size







1. if (x < A size)2. $\mathbf{y} = \mathbf{A}[\mathbf{x}]$ $\boldsymbol{z} = \boldsymbol{B}[\boldsymbol{y}]$ 3. end 4.

x < A size









1. if (x < A size)2. $\mathbf{y} = \mathbf{A}[\mathbf{x}]$ $\boldsymbol{z} = \boldsymbol{B}[\boldsymbol{y}]$ 3. end 4.

x < A size





load A+x



1. if (x < A size) 2. $y = \mathbf{A}[\mathbf{x}]$ $\boldsymbol{z} = \boldsymbol{B}[\boldsymbol{y}]$ 3. end 4.

x < A size





load **B+A**[**x**]



1. if (x < A_size) 2. $\mathbf{y} = \mathbf{A}[\mathbf{x}]$ 3. $\boldsymbol{z} = \boldsymbol{B}[\boldsymbol{y}]$

4. end







1. if (x < A_size) 2. $\mathbf{y} = \mathbf{A}[\mathbf{x}]$ 3. $\boldsymbol{z} = \boldsymbol{B}[\boldsymbol{y}]$ 4. end

x < A size







1. if (x < A size)2. $\mathbf{y} = \mathbf{A}[\mathbf{x}]$ 3. $\boldsymbol{z} = \boldsymbol{B}[\boldsymbol{y}]$ end 4.

x < A size









if (x < A_size) y = A[x] z = B[y] end

x < A_size







if (x < A_size) y = A[x] z = B[y] end

x < A size



load B+A[x],B[A[x]]



1. if (x < A_size) 2. $\mathbf{y} = \mathbf{A}[\mathbf{x}]$ 3. $\boldsymbol{z} = \boldsymbol{B}[\boldsymbol{y}]$ 4. end





1. if (x < A_size) 2. $\mathbf{y} = \mathbf{A}[\mathbf{x}]$ 3. $\boldsymbol{z} = \boldsymbol{B}[\boldsymbol{y}]$ 4. end

x > A size





1. if (x < A size)2. $\mathbf{y} = \mathbf{A}[\mathbf{x}]$ 3. $\boldsymbol{z} = \boldsymbol{B}[\boldsymbol{y}]$ 4. end









1. if (x < A size)2. $\mathbf{y} = \mathbf{A}[\mathbf{x}]$ $\boldsymbol{z} = \boldsymbol{B}[\boldsymbol{y}]$ 3. end 4.

x > A size





start pc 2





1. if (x < A size)2. $\mathbf{y} = \mathbf{A}[\mathbf{x}]$ $\boldsymbol{z} = \boldsymbol{B}[\boldsymbol{y}]$ 3. end 4.









1. if (x < A size)2. $\mathbf{y} = \mathbf{A}[\mathbf{x}]$ $\boldsymbol{z} = \boldsymbol{B}[\boldsymbol{y}]$ 3. end 4.









1. if (x < A size) 2. $\mathbf{y} = \mathbf{A}[\mathbf{x}]$ $\boldsymbol{z} = \boldsymbol{B}[\boldsymbol{y}]$ 3. end 4.









1. if (x < A size) 2. $\mathbf{y} = \mathbf{A}[\mathbf{x}]$ $\boldsymbol{z} = \boldsymbol{B}[\boldsymbol{y}]$ 3. end 4.

x > A size





load **B+A**[**x**]





1. if (x < A size) 2. $\mathbf{y} = \mathbf{A}[\mathbf{x}]$ $\boldsymbol{z} = \boldsymbol{B}[\boldsymbol{y}]$ 3. end 4.

x > A size





rollback pc 4





1. if (x < A size)2. $y = \mathbf{A}[\mathbf{x}]$ z = B[y]3. end 4.









Hardware countermeasures



See paper for: processor operational semantics and security proofs





Speculative and out-of-order executions

See paper for: processor operational semantics and security proofs





Speculative and out-of-order executions

3-stage pipeline (fetch, execute, retire)

See paper for: processor operational semantics and security proofs





See paper for: processor operational semantics and security proofs

Speculative and out-of-order executions

3-stage pipeline (fetch, execute, retire)

Parametric in **branch predictor** and memory hierarchy


A simple processor

Modeled as **operational semantics** \Rightarrow describing how processor's state changes (inspired by [Cauligi et al. 2019])

See paper for: processor operational semantics and security proofs

Speculative and out-of-order executions

Parametric in branch predictor and memory hierarchy









if (x < A_size) z = A[x] y = B[z]</pre>





if (x < A_size) z = A[x] y = B[z]</pre>





if (x < A_size) z = A[x] y = B[z]</pre>









if (x < A_size) z = A[x] y = B[z]</pre>

Satisfies spec-ct



Disabling speculative execution





Disabling speculative execution





Instructions are executed sequentially















Delaying loads until all sources of speculation are resolved







Security guarantees?

if (x < A_size) z = A[x] y = B[z]</pre>

if (x < A_size) z = A[x] y = B[z]</pre>

if (x < A_size) z = A[x] y = B[z]</pre>

A[x] and B[z] delayed until x < A size is resolved</pre>



if (x < A_size) z = A[x] y = B[z]</pre>

A[x] and B[z] delayed until x < A size is resolved</pre>







z = A[x] if(x < A_size) if(z==0) skip</pre>

z = A[x] if(x < A size) if(z==0) skip</pre>

z = A[x] if(x < A size) if(z==0) skip</pre>

if (z==0) is not delayed



z = A[x] if(x < A size) if(z==0) skip</pre>

if (z==0) is not delayed

Program speculatively leaks **A**[**x**]







Satisfies seq-arch

Satisfies seq/spec-ct/pc











Taint speculatively loaded data





Taint speculatively loaded data

Propagate taint through computation





Taint speculatively loaded data

Propagate taint through computation

Delay tainted operations







Taint speculatively loaded data

Security guarantees?

Delay tainted operations



if (x < A size) z = A[x]y = B[z]



if (x < A size) z = A[x]y = B[z]



if (x < A size) z = A[x]y = B[z]

A[**x**] tainted as **unsafe B**[**z**] **delayed** until **A**[**x**] is safe





if (x < A size) z = A[x]y = B[z]

A[**x**] tainted as **unsafe B**[**z**] **delayed** until **A**[**x**] is safe









z = A[x]if (x < A_size) y = B[z]



z = A[x]if (x < A size) $\mathbf{y} = \mathbf{B}[\mathbf{z}]$



z = A[x]if (x < A size) y = B[z]

A[**x**] tagged as **safe B**[*z*] not delayed





z = A[x]if (x < A size) $\boldsymbol{y} = \boldsymbol{B}[\boldsymbol{z}]$

A[**x**] tagged as **safe B**[*z*] not delayed

Program speculatively leaks **A**[**x**] ----






Hardware taint-tracking [Yu et al. 2019, Weisse et al. 2019]





Hardware taint-tracking [Yu et al. 2019, Weisse et al. 2019]



Satisfies seq-arch

Satisfies **spec-ct**





See paper for: models and security proofs

Seq-ct Seq/spiec-ct/pc Spec-ct



Vanilla OoO CPU + SPEC. EXEC

In-order CPU (no specExec)

OoO CPU+load delay







See paper for: models and security proofs

Seq-ct Seq/spiec-ct/pc Spec-ct



Vanilla OoO CPU + SPEC. EXEC

In-order CPU (no specExec)

OoO CPU+load delay









Vanilla OoO CPU + SPEC. EXEC

In-order CPU (no specExec)

OoO CPU+load delay







See paper for: models and security proofs



Vanilla OoO CPU + SPEC. EXEC

In-order CPU (no specExec)

OoO CPU+load delay







See paper for: models and security proofs



Vanilla OoO CPU + SPEC. EXEC

In-order CPU (no specExec)

OoO CPU+load delay







Secure programming

Two flavors of secure programming



Constant-time



Sandboxing

Two flavors of secure programming





Sandboxing

Two flavors of secure programming



Constant-time



Checking secure pro	
	Cor
Seq-ct	
Seq-arch	
Spec-ct	
	l

See paper for: additional security checks, proofs, automation





See paper for: additional security checks, proofs, automation





Hardware

Hw-Sw security contracts

Modeling and specifications

Software



Hardware

Hw-Sw security contracts

Modeling and specifications

Software



flow leaks + execution modes for branch speculation

So far: toy ISA + observation modes for cache/control-



flow leaks + execution modes for branch speculation

Challenge: scale to real-world ISAs

So far: toy ISA + observation modes for cache/control-



flow leaks + execution modes for branch speculation

Challenge: scale to real-world ISAs

extensions to capture leaks

So far: toy ISA + observation modes for cache/control-

What do we need? Formal models of ISAs + adequate



Hardware

Hw-Sw security contracts

Modeling and specifications

Software



Hardware

Hw-Sw security contracts

Modeling and specifications

Software



So far: manual security proofs on toy CPU/ISA





So far: manual security proofs on toy CPU/ISA

Challenge: automation for realistic CPUs





So far: manual security proofs on toy CPU/ISA

Challenge: automation for realistic CPUs

What do we need? Automated verification and testing techniques for contract satisfaction



49



Hardware

Hw-Sw security contracts

Modeling and specifications

Software



Hardware

Hw-Sw security contracts

Modeling and specifications

Software



tools targeting "fixed contracts"

So far: contracts as secure programming foundations +



tools targeting "fixed contracts"

Challenge: support for large classes of contracts

So far: contracts as secure programming foundations +



tools targeting "fixed contracts"

Challenge: support for large classes of contracts

and secure compilation passes



So far: contracts as secure programming foundations +

What do we need? Contract-aware program analyses





Conclusions

Contracts



Contract ISA extended with observations

Contract traces: $[m](p, \sigma)$



Contract satisfaction

Hardware \blacksquare satisfies contract \blacksquare if for all programs p and arch. states σ , σ' : if $\square(p,\sigma) = \square(p,\sigma')$ then $\square(p,\sigma) = \square(p,\sigma')$

Security guarantees



Two flavors of secure programming



Constant-time

Contracts



Contract ISA extended with observations



Contract traces



Security gua

TT

Seq-arch⁻

Spec-arch
















Knows what *data* is *sensitive*



Software

Knows what *data* is *sensitive*

Hardware

Controls microarchitectural state



Contracts as basis for HW/SW co-design C \mathbf{O} n a C Hardware





Knows what **data** is **sensitive**

Controls microarchitectural state



Contracts @ Software



Contracts @ Software

Program analysis



Guarnieri et al., Spectector: Principled detection of speculative leaks, S&P 2020



Contracts @ Software

Program analysis



Guarnieri et al., Spectector: Principled detection of speculative leaks, S&P 2020



Patrignani and Guarnieri, Exorcising Spectre with secure compilers, CCS 2021

Contracts @ Hardware

Contracts @ Hardware

Verification



Tools for checking *compliance with contracts*

(Partially) automate contract satisfaction proofs

Contracts @ Hardware

Verification



Tools for checking *compliance with contracts*

(Partially) automate contract satisfaction proofs

Testing/Fuzzing

Contracts as test oracles

Black box: see Oleksenko et al., Revizor: Fuzzing for Leaks in Black-box CPUs https://arxiv.org/ abs/2105.06872

White box: WIP!





 $\langle m, a, buf, cs, bp, sc \rangle$







 $\langle m, a, buf, cs, bp, sc \rangle$





 $\langle m, a, buf, cs, bp, sc \rangle$





 $\langle m, a, buf, cs, bp, sc \rangle$





 $\langle m, a, buf, cs, bp, sc \rangle$ (metadata)









 $\langle m, a, buf, cs, bp, sc \rangle$ Scheduler predictor (metadata)







Semantics: Describe how configurations evolve



Semantics: Describe how configurations evolve

 $\langle m, a, buf, cs, bp \rangle \stackrel{d}{\Rightarrow} \langle m', a', buf', cs', bp' \rangle$ $d = next(sc) \qquad sc' = update(sc, buf'\downarrow)$ $\langle m, a, buf, cs, bp, sc \rangle \Rightarrow \langle m', a', buf', cs', bp', sc' \rangle$



Semantics: Describe how configurations evolve

 $\langle m, a, buf, cs, bp \rangle \stackrel{d}{\Rightarrow} \langle m', a', buf', cs', bp' \rangle$ $d = next(sc) \qquad sc' = update(sc, buf'\downarrow)$ $\langle m, a, buf, cs, bp, sc \rangle \Rightarrow \langle m', a', buf', cs', bp', sc' \rangle$



Current state



Semantics: Describe how configurations evolve

 $\langle m, a, buf, cs, bp, sc \rangle \Rightarrow \langle m', a', buf', cs', bp', sc' \rangle$

Current state

 $\langle m, a, buf, cs, bp \rangle \stackrel{d}{\Rightarrow} \langle m', a', buf', cs', bp' \rangle$ $d = next(sc) \qquad sc' = update(sc, buf'\downarrow)$



Semantics: Describe how configurations evolve

Directive from scheduler fetch, execute i, retire

 $\langle m, a, buf, cs, bp \rangle \stackrel{d}{\Rightarrow} \langle m', a', buf', cs', bp' \rangle$ $d = next(sc) \qquad sc' = update(sc, buf'\downarrow)$

 $\langle m, a, buf, cs, bp, sc \rangle \Rightarrow \langle m', a', buf', cs', bp', sc' \rangle$





Semantics: Describe how configurations evolve

Directive from scheduler fetch, execute i, retire

Semantics for *pipeline* stages Next state 60

 $\begin{cases} \langle m, a, buf, cs, bp \rangle \stackrel{d}{\Rightarrow} \langle m', a', buf', cs', bp' \rangle \\ d = next(sc) \qquad sc' = update(sc, buf'\downarrow) \end{cases}$ $\langle m, a, buf, cs, bp, sc \rangle \Rightarrow \langle m', a', buf', cs', bp', sc' \rangle$





















































Fetch-Branch-Hit

 $\langle m, a, buf, cs, bp \rangle \xrightarrow{\text{fetch}} \langle m, a, buf \cdot \mathbf{pc} \leftarrow \ell' @a'(\mathbf{pc}), cs', bp \rangle$

a' = apl(buf, a) $|buf| < \mathbf{W}$ $a'(\mathbf{pc}) \neq \bot$ $p(a'(\mathbf{pc})) = \mathbf{beqz} \ x, \ell$ $\ell' = predict(bp, a'(\mathbf{pc}))$ $access(cs, a'(\mathbf{pc})) = Hit$ $update(cs, a'(\mathbf{pc})) = cs'$

applying reorder buffer to register state Fetch-Branch-Hit $a' = apl(buf, a) |buf| < \mathbf{w} a'(\mathbf{pc}) \neq \bot$ $p(a'(\mathbf{pc})) = \mathbf{beqz} x, \ell \qquad \ell' = predict(bp, a'(\mathbf{pc}))$

 $\langle m, a, buf, cs, bp \rangle \xrightarrow{\text{fetch}} \langle m, a, buf \cdot \mathbf{pc} \leftarrow \ell' @a'(\mathbf{pc}), cs', bp \rangle$

 $access(cs, a'(\mathbf{pc})) = Hit$ $update(cs, a'(\mathbf{pc})) = cs'$

Rules capture effect of directives - Fetch reorder buffer applying reorder buffer is not full

applying reorder buffer to register state FETCH-BRANCH-HIT a' = apl(buf, a) $p(a'(\mathbf{pc})) = \mathbf{beqz} \ x, \ell$ $access(cs, a'(\mathbf{pc})) = \text{Hit}$

 $\langle m, a, buf, cs, bp \rangle \xrightarrow{\text{fetch}} \langle m, a, buf \cdot \mathbf{pc} \leftarrow \ell' @a'(\mathbf{pc}), cs', bp \rangle$

 $a'(\mathbf{pc}) \neq \bot$ $|buf| < \mathbf{w}$ $p(a'(\mathbf{pc})) = \mathbf{beqz} \ x, \ell$ $\ell' = predict(bp, a'(\mathbf{pc}))$ $access(cs, a'(\mathbf{pc})) = Hit$ $update(cs, a'(\mathbf{pc})) = cs'$

Rules capture effect of directives - Fetch reorder buffer applying reorder buffer is not full to register state $\sum_{\substack{a = apl(buf, a) \\ p(a'(\mathbf{pc})) = \mathbf{beqz} \ x, \ell \\ access(cs, a'(\mathbf{pc})) = \mathrm{Hit} \\ ext{ buf } < \mathbf{w} \qquad a'(\mathbf{pc}) \neq \bot \\ \ell' = predict(bp, a'(\mathbf{pc})) \\ update(cs, c') \\ \ell' = predict(bp, a'(\mathbf{pc})) \\ \ell' = predict($ $update(cs, a'(\mathbf{pc})) = cs'$ $\langle m, a, buf, cs, bp \rangle \xrightarrow{\text{fetch}} \langle m, a, buf \cdot \mathbf{pc} \leftarrow \ell' @a'(\mathbf{pc}), cs', bp \rangle$ instruction is a branch

applying reorder buffer to register state Fetch-Branch-Hit a' = apl(buf, a) $access(cs, a'(\mathbf{pc})) = Hit$

 $\langle m, a, buf, cs, bp \rangle \xrightarrow{\text{fetch}} \langle m, a, buf \cdot \mathbf{pc} \leftarrow \ell' @a'(\mathbf{pc}), cs', bp \rangle$

instruction is a branch





applying reorder buffer to register state Fetch-Branch-Hit $access(cs, a'(\mathbf{pc})) = Hit$

instruction is a branch



applying reorder buffer to register state Fetch-Branch-Hit

instruction is a branch

branch predictor reorder buffer says ℓ' is next is not full $\Rightarrow a' = apl(buf, a) \qquad |buf| < \mathbf{w} \qquad a'(\mathbf{pc}) \neq \bot$ $p(a'(\mathbf{pc})) = \mathbf{beqz} \ x, \ell \qquad \ell' = predict(bp, a'(\mathbf{pc}))$ $access(cs, a'(\mathbf{pc})) = Hit$ $update(cs, a'(\mathbf{pc})) = cs'$ $\langle m, a, buf, cs, bp \rangle \xrightarrow{\text{fetch}} \langle m, a, buf \cdot \mathbf{pc} \leftarrow \ell' @a'(\mathbf{pc}), cs', bp \rangle$ Cache hit + updating add change of pc cache state to reorder buffer 61



RETIRE-ASSIGNMENT

 $buf = x \leftarrow v@\varepsilon \cdot buf'$ $v \in Vals$ $\langle m, a, buf, cs, bp \rangle \xrightarrow{\text{retire}} \langle m, a[x \mapsto v], buf', cs, bp \rangle$
Rules capture effect of directive - Retire

result of instruction at head of reorder buffer is resolved



RETIRE-ASSIGNMENT $buf = x \leftarrow v @ \varepsilon \cdot buf'$ $v \in Vals$ $\langle m, a, buf, cs, bp \rangle \xrightarrow{\text{retire}} \langle m, a[x \mapsto v], buf', cs, bp \rangle$

Rules capture effect of directive - Retire

result of instruction at head of reorder buffer is resolved

RETIRE-ASSIGNMENT

 $buf = x \leftarrow v @\varepsilon \cdot buf'$ $v \in Vals$

 $\langle m, a, buf, cs, bp \rangle \xrightarrow{\text{retire}} \langle m, a[x \mapsto v], buf', cs, bp \rangle$

apply change to registers and remove entry from reorder buffer



Eager load delay (Sakalis et al. 2019)

STEP-EAGER-DELAY

 $\langle m, a, buf, cs, bp \rangle \stackrel{d}{\Rightarrow} \langle m', a', buf', cs',$ d = next(sc) $sc' = update(sc, buf'\downarrow)$ $buf|_i =$ load x, e $\forall \mathbf{pc} \leftarrow \ell @ \ell' \in buf [0.$

 $\langle m, a, buf, cs, bp, sc \rangle \Rightarrow_{\text{loadDelay}} \langle m', a', buf',$

STEP-OTHERS

 $\langle m, a, buf, cs, bp \rangle \xrightarrow{d} \langle m', a', buf', cs',$ d = next(sc) sc' = update(sc, bl)

 $d \in \{\text{fetch}, \text{retire}\} \lor (d = \text{execute } i \land buf|_i \neq \text{load } x, e)$

 $\langle m, a, buf, cs, bp, sc \rangle \Rightarrow_{\text{loadDelay}} \langle m', a', buf', cs', bp', sc' \rangle$

$$bp'\rangle$$

 $d = execute i$
 $(i-1]. \ell' = \varepsilon$
 $cs', bp', sc'\rangle$

$$\left| bp' \right\rangle$$

 $uf' \downarrow$)

Eager load delay [Sakalis et al. 2019]

STEP-EAGER-DELAY

 $\langle m, a, buf, cs, bp \rangle \stackrel{d}{\Rightarrow} \langle m', a', buf', cs', bp' \rangle$ d = next(sc) $sc' = update(sc, buf'\downarrow)$ d = execute i $buf|_i = \text{load } x, e$ $\forall \mathbf{pc} \leftarrow \ell @ \ell' \in buf[0..i-1]. \ \ell' = \varepsilon$ $\langle m, a, buf, cs, bp, sc \rangle \Rightarrow_{\text{loadDelay}} \langle m', a', buf', cs', bp', sc' \rangle$

STEP-OTHERS

 $\langle m, a, buf, cs, bp \rangle \stackrel{d}{\Rightarrow} \langle m', a', buf', cs', bp' \rangle$ d = next(sc) $sc' = update(sc, buf' \downarrow)$

 $d \in \{\text{fetch}, \text{retire}\} \lor (d = \text{execute } i \land buf|_i \neq \text{load } x, e)$

 $\langle m, a, buf, cs, bp, sc \rangle \Rightarrow_{\text{loadDelay}} \langle m', a', buf', cs', bp', sc' \rangle$

Loads are executed only if prior branch instructions are resolved



Eager load delay [Sakalis et al. 2019]

STEP-EAGER-DELAY

 $\langle m, a, buf, cs, bp \rangle \stackrel{d}{\Rightarrow} \langle m', a', buf', cs', bp' \rangle$ d = next(sc) $sc' = update(sc, buf'\downarrow)$ d = execute i $buf|_i = \text{load } x, e$ $\forall \mathbf{pc} \leftarrow \ell @ \ell' \in buf[0..i-1]. \ \ell' = \varepsilon$ $\langle m, a, buf, cs, bp, sc \rangle \Rightarrow_{\text{loadDelay}} \langle m', a', buf', cs', bp', sc' \rangle$

STEP-OTHERS

 $\langle m, a, buf, cs, bp \rangle \stackrel{d}{\Rightarrow} \langle m', a', buf', cs', bp' \rangle$ d = next(sc) $sc' = update(sc, buf'\downarrow)$

 $d \in \{\text{fetch}, \text{retire}\} \lor (d = \text{execute } i \land buf|_i \neq \text{load } x, e)$

 $\langle m, a, buf, cs, bp, sc \rangle \Rightarrow_{\text{loadDelay}} \langle m', a', buf', cs', bp', sc' \rangle$

Loads are executed only if prior branch instructions are resolved

Everything else works as before





Hw-level taint-tracking [Yu et al. 2019, Weisse et al. 2019]

STEP

$$d = next(sc) \qquad buf_{ul} = unlbl(buf,d)$$

$$\langle m, a, buf_{ul}, cs, bp \rangle \stackrel{d}{\Rightarrow} \langle m', a', buf'_{ul}, cs', bp' \rangle$$

$$sc' = update(sc, buf' \downarrow) \qquad buf' = lbl(buf'_{ul}, buf, d)$$

$$\langle m, a, buf, cs, bp, sc \rangle \Rightarrow_{tt} \langle m', a'buf', cs', bp', sc' \rangle$$



Hw-level taint-tracking (Yu et al. 2019, Weisse et al. 2019)



$$buf_{ul} = unlbl(buf, d)$$

$$\stackrel{d}{\Rightarrow} \langle m', a', buf'_{ul}, cs', bp' \rangle$$

$$buf' = lbl(buf'_{ul}, buf, d)$$

$$\Rightarrow_{tt} \langle m', a'buf', cs', bp', sc' \rangle$$
For are labelled as safe/uns



Hw-level taint-tracking M

STEP d = next(sc) $\langle m, a, buf_{ul}, cs, bp \rangle$ $sc' = update(sc, buf'\downarrow)$ $\langle m, a, buf, cs, bp, sc \rangle =$ Entries in the *reorder buffer* are labelled as *safe/unsa*





Hw-level taint-tracking Yu et al. 2019, Weisse et al. 2019] Derive unlabeled buffer (hides information tagged as unsafe) $buf_{ul} = unlbl(buf, d)$ $\langle m, a, buf_{ul}, cs, bp \rangle \xrightarrow{d} \langle m', a', buf_{ul}', cs', bp' \rangle$ $sc' = update(sc, buf'\downarrow)$ $buf' = lbl(buf'_{ul}, buf, d)$ $\langle m, a, buf, cs, bp, sc \rangle \Rightarrow_{tt} \langle m', a'buf', cs', bp', sc' \rangle$ Entries in the *reorder buffer* are labelled as *safe/unsafe*



Computation on *unlabeled buffer* STEP

Entries in the *reorder buffer* are labelled as *safe/unsafe*

Hw-level taint-tracking (Yu et al. 2019, Weisse et al. 2019) Derive unlabeled buffer (hides information tagged as unsafe) d = next(sc) $buf_{ul} = unlbl(buf, d)$ $\langle m, a, buf_{ul}, cs, bp \rangle \stackrel{d}{\Rightarrow} \langle m', a', buf'_{ul}, cs', bp' \rangle$ $sc' = update(sc, buf'\downarrow)$ $buf' = lbl(buf'_{ul}, buf, d)$ $\langle m, a, buf, cs, bp, sc \rangle \Rightarrow_{tt} \langle m', a'buf', cs', bp', sc' \rangle$ **Update** labels









See paper for: security definitions

 Key hardware data structure for out-of-order and speculative execution

- Key hardware data structure for out-of-order and speculative execution
- Keeps track of "in-flight instructions"

- Key hardware data structure for out-of-order and speculative execution
- Keeps track of "in-flight instructions"
- Example:

Entry Instruction Control Dep. O: c ← x < A size 1: beqz c, END 2: 3:





- Key hardware data structure for out-of-order and speculative execution
- Keeps track of "in-flight instructions"
- Example:

Entry Instruction Control Dep. 0: c ← x < A size 1: beqz c, END 2: 3:

 $C \leftarrow X < A size$ beqz c, END L1: load t, \mathbf{A} + x load y, **B** + t END:

Speculative Instruction Fetch

Entry	Instruction	Control Dep.
0:	c ← x < A_size	_
1:	beqz c, END	_
2:	load t, \mathbf{A} + x	1
3:		_
	•••	



- Key hardware data structure for out-of-order and speculative execution
- Keeps track of "in-flight instructions"
- Example:

. . .

. . .

Entry Instruction Control Dep. O: c ← x < A size 1: beqz c, END 2: load t, **A** + x 3:

 $C \leftarrow X < A size$ beqz c, END L1: load t, \mathbf{A} + x load y, **B** + t END:

Speculative Instruction Fetch

Entry	Instruction	Control Dep.
0:	c ← x < A_size	_
1:	beqz c, <i>END</i>	_
2:	load t, $\mathbf{A} + \mathbf{x}$	1
3:	load y, B + t	1

. . .



- Key hardware data structure for out-of-order and speculative execution
- Keeps track of "in-flight instructions"
- Example:

Entry Instruction	Control Dep.	Evaluate	Entry	Instruction	Control Dep.
0: c ← x < A_size	_	x < A_size	0:	c ← 0	_
1: beqz c, END	_		1:	beqz c, <i>END</i>	_
2: load t, \mathbf{A} + x	1		2:	load t, $\mathbf{A} + \mathbf{x}$	1
3: load y, B + t	1		3:	load y, B + t	1
	•••			• • •	





- Key hardware data structure for out-of-order and speculative execution
- Keeps track of "in-flight instructions"
- Example: Entry Instruction Control Dep. **0:** c ← 0 1: beqz c, END 2: load t, \mathbf{A} + x 3: load y, **B** + t

. . .

. . .



Rollback	Entry	Instruction	Control Dep.
mis-speculation	0:	c ← 0	_
	1:	beqz c, END	_
	2:	—	_
	3:	_	_
		•••	

. . .



- Key hardware data structure for out-of-order and speculative execution
- Keeps track of "in-flight instructions"
- Example: Entry Instruction Control Dep. **0:** c ← 0 1: beqz c, END 2: 3:

. . .

. . .

Retire





. . .



Speculative execution attacks 101



Size of array A if (x < A size) y = B[A[x]]

Size of array A

if $(\mathbf{x} < \mathbf{A} \text{ size})$ $y = \mathbf{B}[\mathbf{A}[\mathbf{x}]]$

Size of array A

if $(\mathbf{x} < \mathbf{A} \text{ size})$



Size of array A

if (x < A size) y = B[A[x]]

Prediction based on **branch** history & program structure





Size of array A if (x < A size) y = B[A[x]]

Prediction based on **branch** history & program structure









Size of array A

if (x < A size) y = B[A[x]]

Wrong predicton? Rollback changes!

Architectural (ISA) state

Microarchitectural state

Prediction based on **branch** history & program structure









void f(int x) if (x < A_size) y = B[A[x]]</pre>

void f(int x) if (x < A_size) y = B[A[x]]</pre>



void f(int x) if (x < A size) y = B[A[x]]</pre>













1) Train branch predictor







1) Train branch predictor

2) Prepare cache







1) Train branch predictor

2) Prepare cache

3) Run with x = 128






2) Prepare cache







2) Prepare cache







2) Prepare cache







2) Prepare cache







2) Prepare cache

3) Run with x = 128

4) Extract from cache



Secure programming — foundations

Secure programming — foundations

Program p is **non-interferent** wrt contract and policy π if for all arch. states σ , σ' : if $\sigma \approx_{\pi} \sigma'$ then $mathbb{(}p, \sigma) = mathbb{(}p, \sigma')$

Program p is *non-interferent* wrt contract []] and policy π if for all arch. states σ , σ' : if $\sigma \approx_{\pi} \sigma'$ then $\equiv (p, \sigma) = \equiv (p, \sigma')$



Program p is **non-interferent** wrt contract []] and policy π if for all arch. states σ , σ' : if $\sigma \approx_{\pi} \sigma'$ then $\equiv (p, \sigma) = \equiv (p, \sigma')$



Secure programming — foundations Specify secret data Program p is **non-interferent** wrt contract []] and policy π if for all arch. states σ , σ' : if $\sigma \approx_{\pi} \sigma'$ then $\exists (p, \sigma) = \exists (p, \sigma')$



Program p is **non-interferent** wrt contract []] and policy π if for all arch. states σ , σ' : if $\sigma \approx_{\pi} \sigma'$ then $\equiv (p, \sigma) = \equiv (p, \sigma')$

If p is **non-interferent** wrt contract $[\mathbf{m}]$ and policy π , and hardware \blacksquare satisfies \blacksquare , then p is *non-interferent* wrt hardware \square and policy π



Program p is **non-interferent** wrt contract []] and policy π if for all arch. states σ , σ' : if $\sigma \approx_{\pi} \sigma'$ then $\equiv (p, \sigma) = \equiv (p, \sigma')$









Program p is **non-interferent** wrt contract []] and policy π if for all arch. states σ , σ' : if $\sigma \approx_{\pi} \sigma'$ then $\equiv (p, \sigma) = \equiv (p, \sigma')$

If p is **non-interferent** wrt contract $[\blacksquare]$ and policy π , and hardware satisfies 🗐 then p is *non-interferent* wrt hardware \square and policy π





Program p is **non-interferent** wrt contract []] and policy π if for all arch. states σ , σ' : if $\sigma \approx_{\pi} \sigma'$ then $\equiv (p, \sigma) = \equiv (p, \sigma')$

If p is **non-interferent** wrt contract $[\mathbf{m}]$ and policy π , and hardware \blacksquare satisfies \blacksquare , then p is *non-interferent* wrt hardware \square and policy π



Constant-time

Constant-time

Traditional CT wrt policy $\pi \equiv$ non-interference wrt seq-ct and π



Control-flow and memory accesses do not depend on secrets

Traditional CT wrt policy $\pi \equiv$ non-interference wrt seq-ct and π



General CT wrt π and \equiv non-interference wrt \equiv and π



Control-flow and memory accesses do not depend on secrets

Traditional CT wrt policy $\pi \equiv$ non-interference wrt seq-ct and π

Sandboxing

Sandboxing

Traditional SB wrt policy $\pi \equiv$ non-interference wrt seq-arch and π





Programs never access high memory locations (out-of-sandbox)

Traditional SB wrt policy $\pi \equiv$ non-interference wrt seq-arch and π





General SB wrt π and $\equiv \equiv$ Traditional SB wrt π + non-interference wrt π and \equiv

Programs never access high memory locations (out-of-sandbox)

Traditional SB wrt policy $\pi \equiv$ non-interference wrt seq-arch and π





Checking secure programming				
	Constant-time	Sandboxing		
$[\cdot]_{ct}^{seq}$	Traditional CT	Traditional SB		
[[•]] ^{seq} arch	+ NI wrt [[•]] seq arch	Traditional SB		
$\llbracket \bullet \rrbracket_{ct}^{spec}$	+ Speculative NI [Guarnieri et al., S&P'20]	+ weak spec. N		



Checking secure programming				
		Constant-time	Sandboxing	
	$\llbracket \cdot \rrbracket_{ct}^{seq}$	Traditional CT	Traditional SB	
	[[•]] ^{seq} arch	+ NI wrt [[•]] ^{seq} arch	Traditional SB	
	$\llbracket \cdot \rrbracket_{ct}^{spec}$	+ speculative NI [Guarnieri et al., S&P'20]	+ weak spec. N	

