

## Beyond constant-time programming Hardware-software co-designs for microarchitectural security

Lesly-Ann Daniel

Chalmers Security & Privacy Seminar - November 20<sup>th</sup> 2023

## Programs handle secret data...



### ... which can affect timing/microarchitecture



#### ... and leak via side-channel attacks



### Solution? Constant-time programming!



#### What can influence execution time/microarchitecture?

## Control Flow if secret then foo() else bar() secret

#### What can influence execution time/microarchitecture?





#### What can influence execution time/microarchitecture?





## Solution? Constant-time programming!



#### Leaky instructions

- Control-Flow
- Memory accesses

- Full software countermeasure
- Plenty of analysis tools: ctgrind, MicroWalk, Binsec/Rel
- De facto standard for crypto: BearSSL, Libsodium, HACL\*, etc.

#### Constant-time is not perfect





# 

#### In Software?

Constant-time Programming Speculative Constant-time

#### In Hardware?

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





#### Hardware-software co-design

Best performance/security tradeoff



#### HW/SW Collaboration for End-to-End Security

#### **PROSPECT: Provably Secure Speculation for the Constant-Time Policy**

Lesly-Ann Daniel<sup>1</sup>, Marton Bognar<sup>1</sup>, Job Noorman<sup>1</sup>, Sébastien Bardin<sup>2</sup>, Tamara Rezk<sup>3</sup> and Frank Piessens<sup>1</sup>

<sup>1</sup>imec-DistriNet, KU Le <sup>2</sup>CEA, List, Univer <sup>3</sup>INRIA, Université Côte d

#### Abstract

We propose PROSPECT, a generic formal processor model providing provably secure speculation for the constant-time



Architectural Mimicry: Innovative Instructions to Efficiently Address Control-Flow Leakage in Data-Oblivious Programs

Hans Winderix imec-DistriNet KU Leuven Marton Bognar imec-DistriNet KU Leuven Job Noorman imec-DistriNet KU Leuven

Lesly-Ann Daniel imec-DistriNet KU Leuven

Frank Piessens imec-DistriNet KU Leuven

Abstract—The control flow of a program can often be observed through side-channel attacks. Hence, when control flow depends on secrets, attackers can learn information about these secrets. Widely used software-based countermeasures ensure that attacker-observable aspects of the control flow do not depend on secrets, relying on techniques like *dummy execution* (for balancing code) or *conditional execution* (for linearizing code). In the current state-of-practice, the primitives to implement these techniques have to be found in an existing i tion set architecture (ISA) that was not designed a p provide them, leading to performance, security, and por issues. To counter these issues, this paper proposes ligh hardware extensions for supporting these technique principled way. We propose (1) a powel bordware benefits compared to other approaches [5], [6]. Second, *linearization* [7]–[10] ensures that control flow does not depend on program secrets at all.

Balancing and linearization are important ingredients in state-of-practice software-based countermeasures (such as *constant-time programming* [11]), as well as in recent research prototypes [5]–[7], [10], [12]. They are based on techniques like *dummy execution* (i.e., using architectural no-ops

45th IEEE Symposium on Security and Privacy

#### **PROSPECT: Provably Secure Speculation for the Constant-Time Policy**

Lesly-Ann Daniel<sup>1</sup>, Marton Bognar<sup>1</sup>, Job Noorman<sup>1</sup>, Sébastien Bardin<sup>2</sup>, Tamara Rezk<sup>3</sup> and Frank Piessens<sup>1</sup>

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

#### Constant-time is vulnerable to Spectre



#### Constant-time is vulnerable to Spectre



#### Constant-time is vulnerable to Spectre



## How can I protect my code?

#### **Constant-Time Foundations for the New Spectre Era**

Sunjay Cauligi<sup>†</sup> Craig Disselkoen<sup>†</sup> Klaus v. Gleissenthall<sup>†</sup> Dean Tullsen<sup>†</sup> Deian Stefan<sup>†</sup> Tamara Rezk<sup>\*</sup> Gilles Barthe<sup>\*\*</sup>

<sup>†</sup>UC San Diego, USA **\***INRIA Sophia Antipolis, France **\***MPI for Security and Privacy, Germany **\***IMDEA Software Institute, Spain

#### **Speculative constant-time**

- Hard to reason about
- New speculation mechanisms?



#### We need Secure Speculation for Constant-Time!



Developers should not care about speculations



Hardware shall not speculatively leak secrets



But still be efficient and enable speculation



#### Hardware defense:

Secure speculation for constant-time!

## Hardware Secrecy Tracking



Software side	Hardware side	
Label secrets	Track security labels	
Constant-time program	Secrets do not speculatively flow to insecure instructions	
ConTExT: A Generic Approach for Mitigating Spectre	SpectreGuard: An Efficient Data-centric Defense Mechanism against Spectre Attacks	
Michael Schwarz <sup>1</sup> , Moritz Lipp <sup>1</sup> , Claudio Canella <sup>1</sup> , Robert Schilling <sup>1,2</sup> , Florian Kargl <sup>1</sup> , Daniel Gruss <sup>1</sup> <sup>1</sup> Graz University of Technology	Jacob Fustos Farzad Farshchi Heechul Yun University of Kansas University of Kansas	
Speculative Execution Without Compromising Privacy		
Rutvik Choudhary UIUC, USA Christopher W. Fletcher	Jiyong Yu UIUC, USA Adam Morrison	

```
char array[len]
char mysecret
1: if (idx < len)
2: x = array[idx]
3: leak(x)
```

secret char mysecret	
	ſ
li (lax < len)	
2: $x = array[idx]$	
3: leak(x)	

Developer marks secrets

	<pre>char array[len]</pre>	Developer marks secrets
	<pre>secret char mysecret</pre>	Speculative execution 🌔
1:	if (idx < len)	
2:	x = array[idx]	
3:	<pre>leak(x)</pre>	





## How do I know that my defense works?



## How do I know that my defense works?





#### Adapt HW/SW contract framework to account for

- All existing speculation mechanisms (Spectre, LVI)
- Futuristic speculation mechanisms (value prediction)
- Declassification

## Our contributions

- ProSpeCT: Formal processor model with HST
  - Proof: constant-time programs do not leak secrets
  - Allows for declassification
  - Generic: all Spectre variants / LVI
- First to consider (Load) Value Speculation
  - Novel insight: sometimes need to rollback *correct* speculations for security
- Implementation in a RISC-V microarchitecture
  - First synthesizable implementation
  - Evaluation: hardware cost, performance, annotations





#### ProSpeCT: Generic formal processor model for HST

Semantics of out-of-order speculative processor with HST

- → Abstract microarchitectural context
- → Functions *update*, *predict*, *next*

All public values are leaked / influence predictions

**Declassify** = write secrets to public memory

 $\rightarrow$  Beware unintentional declassification

Attacker observations/influence

Generic/Powerful predictors

### ProSpeCT: Generic formal processor model for HST

Semantics of out-of-order speculative processor with HST

- Abstract microarchitectural context  $\rightarrow$
- Functions *update*, *predict*, *next*  $\rightarrow$

All public values are leaked / influence predictions

**Declassify** = write secrets to public memory

Beware unintentional declassification  $\rightarrow$ 

Security proof

Constant-time programs (ISA semantics)

do not leak secrets (micro-arch. semantics)

Attacker observations/influence

Generic/Powerful predictors

#### Load Prediction: Rollback correct executions?

1: x = load mysecret y = x + 4

#### Load Prediction: Rollback correct executions?



#### Load Prediction: Rollback correct executions?



Implicit resolution-based channel!
### Load Prediction: Rollback correct executions?



Solution: always rollback when value is secret

### Implementation

#### **Prototype RISC-V implementation**

On top of Proteus modular RISC-V processor

- Branch target prediction
- Conservative approach
- 2 secret regions defined by CSRs





github.com/proteus-core/prospect

### Evaluation

#### **Hardware Cost**

#### Synthesized on FPGA

- LUTs: +17%
- Registers: +6%
- Critical path: +2%

#### **Annotation burden**

- 4 primitives (HACL\*)
- Annotate secret
- Ensure no secrets spilled
- Stack public in 3/4 cases
- $\leq 1h / primitive$

### Evaluation

#### **Performance overhead** (benchmark from [1])

Speculation/Crypto	25/75	50/50	75/25	90/10
Precise (Key)	0%	0%	0%	0%
Conservative (All)	10%	25%	36%	45%

No overhead in software for constant-time code when secrets are precisely annotated



[1] Jacob Fustos, Farzad Farshchi, and Heechul Yun. "SpectreGuard: An Efficient Data-Centric Defense Mechanism against Spectre Attacks". In: DAC. 2019





#### Software informs hardware about secrets



#### Strong security guarantees

End-to-end security for constant-time programs



#### Low overhead

No software overhead for constant-time code



#### Architectural Mimicry: Innovative Instructions to Efficiently Address Control-Flow Leakage in Data-Oblivious Programs

Hans Winderix imec-DistriNet KU Leuven Marton Bognar imec-DistriNet KU Leuven

Job Noorman imec-DistriNet KU Leuven Lesly-Ann Daniel imec-DistriNet KU Leuven Frank Piessens imec-DistriNet KU Leuven

### Secrets can leak via control-flow



### Secrets can leak via control-flow



C Conservative leakage model / high-end platforms

### Secrets can leak via control-flow



Conservative leakage model / high-end platforms

Liberal leakage model / small microcontrollers

### Control-Flow Leakage Mitigations?

#### Linearization (CT)

```
c = (secret ≠ 0)
add v1 a a
add v1 v1 8
add v2 a 4
v = select c v1 v2
C + L
```

		E	Bal	anc	ing	
if ( <mark>s</mark>	ecre	t)				
•	add	V	а	а		
	add	V	V	8		
else						
	add	V	а	4		
	add	V	V	0		(L)

## Control-Flow Leakage Mitigations?

Linearization (CT)	Balancing
$c = (secret \neq 0)$ add v1 a a add v1 v1 8 add v2 a 4	<pre>if (secret)     add v a a     add v v 8 else</pre>
v = select c v1 v2 C + L	add v a 4 add v v 0

- Portability (microcontrollers ≠ servers)
- **Issues** Performance (extra instructions, registers)
  - Security (no security guarantees)

### Goal



- Portability: linearization  $\bigcirc + \bigcirc$  and balancing  $\bigcirc$
- Performance: improve over std. linearization/balancing
- Security: leakage contract drives secure software development

### Contributions

- Mimic Execution
  - HW primitive for mimicking instructions
- Architectural Mimicry (AMi)
  - Instructions to control mimic execution
- Programming models
  - Secure/correct balancing/linearization with AMi
- Implementation in RISC-V
  - Evaluation: hardware cost, performance



### Mimic execution



**5** qualifiers

to control mimic execution



# AMi for Linearization C

#### Insecure Code

bnz <mark>secret</mark> else	jmp secret≠0
add v a a	add
add v v 8	add
j end	jmp
else:	
add v a 4	add
end:	

# AMi for Linearization (C)



Insecure Code			
bnz secret else add v a a add v v 8 j end	jmp secret≠0 add add jmp		
else: add v a 4 end:	add		

#### Linearization

a.bnz secret else	jmp
add v a a	add
add v v 8	add
else:	
a.beqz secret end	jmp
add v a 4	add
end:	

# AMi for Linearization (C)



Insecure Code			
bnz secret else add v a a add v v 8 i end	jmp secret≠0 add add imn		
else: add v a 4 end:	add		

#### Linearization

a.bnz secret else	jmp
add v a a	add add
else:	uuu
a.beqz secret end	jmp
add v a 4	add

when secret  $\neq 0$ 

# AMi for Linearization (C)



Insecure Code			
bnz secret else jm add v a a add v v 8 j end	p secret≠0 add add imp		
else: add v a 4 end:	add		

#### Linearization

a.bnz secret else	jmp
add v a a	add add
else:	
a.beqz secret end	jmp
add v a 4 end:	add

when secret = 0

### Maintain Correctness?



#### **Correctness:**

No effect on live state in mimic mode



### Maintain Correctness?



in mimic mode



 $c \neq 0$ 

### Maintain Correctness?



### **Enforce Security?**



### **Enforce Security?**



**Security violation!** 

### **Enforce Security?**



+correct assuming a is not live

### Formalization

- Operational ISA semantics for AMi
  - Instrumented with leakage
- Definition well-behaved activating regions
  - Proof (under conditions) activating regions are well-behaved
  - Nested and recursive activation
- Definitions correct / secure programming model
  - Linearization
  - Balancing

### Implementation

#### **Prototype 32-bit RISC-V implementation on Proteus**

- In order & out-of-order
- Mimic instr = no update register file
- Mode-independent stalling
- Constant-time branch (no prediction)



• 2-bit for instruction qualifiers + 3 CSR for store processor mode

## Evaluation

#### Benchmark

- 11 programs from [1]
- 4 configurations:

Balancing (B) with/wo AMi

Linearization (L) Molnar/AMi



- Security: tests
- Hardware: max 26% LUTS/FF, 0% CP
- Binary size: +19% → +0% (L)\*
  - Performance: +48%  $\rightarrow$  +19% (L)\*

#### \*No change for balancing

[1] H. Winderix, J. T. Mühlberg, and F. Piessens, "Compiler-assisted hardening of embedded software against interrupt latency side-channel attacks," in EuroS&P, 2021.

### Summary



#### Principled linearization and balancing



#### Security-oriented ISA extension

Can be leveraged to write side-channels free sw



#### Accelerate CT code

-60% overhead of linearized code



### A step back



#### **RISC-V** open standard ISA

 $\rightarrow$  HW-SW co-design for security

	1	
1		2

- Proteus: extensible RISC-V processor
- Security extensions
  - ProSpeCT
  - AMi
  - CHERI
  - ...

### Hardware-Software Contracts for Secure Speculation

Marco Guarnieri<sup>\*</sup>, Boris Köpf<sup>†</sup>, Jan Reineke<sup>‡</sup>, and Pepe Vila<sup>\*</sup> \**IMDEA Software Institute* <sup>†</sup>*Microsoft Research* <sup>‡</sup>*Saarland University* 



### How to ease adoption of HW-SW co-designs?

#### Need infrastructure around HW-SW contracts

Compilation support (LLVM / Jasmin)

Support larger code, more realistic performance evaluation

• Validate HW implementation (fuzzing / verification)

Reduce gap between model and implementation

### Conclusion

#### ProSpeCT

#### **Architectural Mimicry**



Software informs hardware about secrets

Strong security guarantees *End-to-end security for constant-time programs* 



#### Low overhead

No software overhead for constant-time code





#### Principled linearization and balancing



Can be leveraged to write side-channels free sw



#### Accelerate CT code -60% overhead of linearized code



#### **Goal:** end-to-end security & performance





Insecure Code		
bnz secret else	jmp	
add v a a add v v 8 j end	add add jmp	
add v a 4	add	
end:		



Insecure Code		AMi Balancing	
onz <mark>secret</mark> else	jmp	bnz secret else	jmp
add v a a add v v 8 j end	add add jmp	add v a a add v v 8 j end	add add jmp
else: add v a 4 end:	add	else: add v a 4 m.add v v 8 j end end:	add add jmp

### Hardware Costs

	LUT	Flip-Flops	<b>Critical path</b>
In-order	+19.5%	+22.9%	+0.6%
Out-of-order	+26.3%	+9.2%	-1.0%

#### Hardware cost overhead of AMi (synthesized on an FPGA)


	Baseline size (bytes)	Balanced		Linearized	
		No AMi	AMi	Molnar	AMi
Min	132	+0%	+0%	+8%	-6%
Max	500	+41%	+41%	+92%	<b>+2%</b>
Mean	321	+8%	+7%	+19%	+0%

Binary size overhead compared to insecure baseline

## **Execution time**

	Balanced (in order)		Linearized (in order)		Linearized (ooo)	
	No AMi	AMi	Molnar	AMi	Molnar	AMi
Min	+6%	+6%	+9%	-11%	-5%	-11%
Max	+143%	+143%	+275%	+69%	+233%	+77%
Mean	+59%	+59%	+57%	+ <b>2</b> 4%	+48%	+19%

Execution time overhead compared to insecure baseline (cycles)