

ProSpeCT: Provably Secure Speculation for the Constant-Time Policy

September 9th 2023

TASER Workshop

Lesly-Ann DanielMarton BognarJob NoormanSébastien BardinTamara RezkFrank PiessensKU LeuvenKU LeuvenCEA ListINRIAKU Leuven

I want to protect my secrets...

Easy: Constant-Time Programming!

De facto standard for crypto

... even against Spectre attacks



... even against Spectre attacks



Predict branch taken



x = mysecret

Leaks mysecret to microarchitecture!

How can I protect my code?

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

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 ¹ , Moritz Lipp ¹ , Claudio Canella ¹ , Robert Schilling ^{1,2} , Florian Kargl ¹ , Daniel Gruss ¹ ¹ Graz University of Technology ² Knc ¹ Carter Carbu Speculative Privacy Tracking	Jacob Fustos Farzad Farshchi Heechul Yun University of Kansas University of Kansas				
Speculative Execution Without Compromising Privacy					
Rutvik Choudhary UIUC, USA	Jiyong Yu UIUC, USA				
Christopher W. Fletcher UIUC, USA	Adam Morrison Tel Aviv University, Israel 7				

How do I know that my defense works?



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



	<pre>char array[len]</pre>
	<pre>secret char mysecret</pre>
1:	if (idx < len)
2:	x = array[idx]
3:	<pre>leak(x)</pre>

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>	





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

ProSpeCT: Generic formal processor model for HST

Semantics of out-of-order speculative processor with HST

All public values are leaked & predictions can depend on any public value

Declassify = write secrets to public memory



Security proof

Constant-time programs do not leak secrets

Implementation

Prototype RISC-V implementation

On top of Proteus modular RISC-V processor

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

https://github.com/proteus-core/prospect



Limited Hardware Cost

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

Evaluation

4 primitives (HACL*)

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

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



Conclusion



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

Icons made by Freepik, Vectors Market, monkik from www.flaticon.com





github.com/proteus-core/prospect