

A classic locked-room mystery. Eve was in the false branch of a conditional the whole time, *how could she do it*?

Mozilla Research | DePaul University | U. California San Diego

<ロ> (四) (四) (三) (三) (三) (三)

3 January 2018



A day out at the Tate Modern

The Code That Never <u>Ran</u>

Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Introduction Spectre Optimizations Simplified Spectre Results Experiments Conclusions

◆□ ▶ ◆□ ▶ ◆ 三 ▶ ◆ 三 ▶ ● ○ ○ ○ ○

3 January 2018



The Code That Never Ran

Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Introduction Spectre Optimizations Simplified Spectre Results Experiments Conclusions

▲□▶ ▲□▶ ▲目▶ ▲目▶ ▲□▶ ▲□▶

3 January 2018







Alan Jeffrey @asaieffrev

OMG a timing back channel based on speculative evaluation and caching. Three issues that are often brushed under the carpet by formal models.

Heather Adkins @argvee

Now with full details. CPU bugs Kudos @tehih googleprojectzero.blogspot.co.uk/2018/01/ readin

5:12 PM · 03 Jan 18

I View Tweet activity

6 Retweets 23 Likes





The Code That Never Ran

Craig Disselkoen. Radha Jagadeesan, Alan Jeffrey, James Rielv

Introduction Spectre Simplified Spectre Results

Spectre



Attacks bypass run-time security checks.

Can bypass array bounds checks, and read whole process memory.

Can be exploited from JS, so evil.ad.com can read your bank.com data.

Attacks *speculative evaluation* hardware optimization.

The Code That Never Ran

Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Introduction Spectre Optimizations Simplified Spectre Results Experiments Conclusions

A lie we tell programmers: "computers execute instructions one after the other."

$$x := x + 1; y := 1$$

has execution:

$$\begin{array}{c} \mathsf{R} x 1 \longrightarrow \mathsf{W} x 2 \longrightarrow \mathsf{W} y 1 \end{array}$$

The Code That Never Ran

Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Introduction

Spectre

Optimizations

Simplified Spectre

Results

Experiments

Conclusions

・ロト・日本・モー・モー シックショ

A lie we tell programmers: "computers execute instructions one after the other."

$$x := x + 1; y := 1$$

has execution where W y 1 might happen first:

$$\begin{array}{c} \mathsf{R} x 1 \longrightarrow \mathsf{W} x 2 \\ \mathsf{W} y 1 \end{array}$$

The Code That Never Ran

Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Introduction

Spectre

Optimizations

Simplified Spectre

Results

Experiments

Conclusions

A lie we tell programmers: "computers execute instructions one after the other."

$$x := x + 1; y := 1$$

has execution:

$$\begin{array}{c} \mathsf{R} x 1 \longrightarrow \mathsf{W} x 2 \\ \mathsf{W} y 1 \end{array}$$

Shared-memory concurrency leaks the abstraction

The Code That Never Ran

Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Introduction

Spectre

Optimizations

Simplified Spectre

Results

Experiments

Conclusions

A lie we tell programmers: "computers execute instructions one after the other."

$$x := x + 1; y := 1$$

has execution:

$$\begin{array}{c} \mathsf{R} x 1 \longrightarrow \mathsf{W} x 2 \\ \mathsf{W} y 1 \end{array}$$

Shared-memory concurrency leaks the abstraction

Resulted in entire research area: weak memory models (e.g. Pugh et al.; C11)

The Code That Never Ran

Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Introduction

Spectre

Optimizations Simplified Spectre Results

Experiments

Conclusions

Another lie we tell programmers: "only one branch of an if is executed."

if
$$(x) \{ y := 1; z := 1 \}$$
 else $\{ y := 2; z := 1 \}$

has execution:



The Code That Never Ran

Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Introduction

Spectre

Optimizations

Simplified Spectre

Results

Experiments

Conclusions

▲□▶ ▲圖▶ ▲目▶ ▲目▶ 目 めんぐ

Another lie we tell programmers: "only one branch of an if is executed."

if
$$(x) \{ y := 1; z := 1 \}$$
 else $\{ y := 2; z := 1 \}$

has execution where W z 1 might happen before W y 1:



The Code That Never Ran

Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Introduction Spectre

Optimizations Simplified Spectre Results Experiments Conclusions

▲□▶ ▲圖▶ ▲≣▶ ▲≣▶ = ● のへで

Another lie we tell programmers: "only one branch of an if is executed."

if
$$(x) \{ y := 1; z := 1 \}$$
 else $\{ y := 2; z := 1 \}$

has execution where W y 2 might happen, then get rolled back:



The Code That Never Ran

Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Introduction Spectre Optimizations Simplified Spectre Results Experiments Conclusions

Optimizations in hardware and compilers

Another lie we tell programmers: "only one branch of an if is executed."

if
$$(x) \{ y := 1; z := 1 \}$$
 else $\{ y := 2; z := 1 \}$

has execution where W z 1 might happen first:



The Code That Never Ran

Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Introduction Spectre Optimizations Simplified Spectre Results Experiments Conclusions

▲□▶ ▲圖▶ ▲≣▶ ▲≣▶ ■ の∢ぐ

Optimizations in hardware and compilers

Another lie we tell programmers: "only one branch of an if is executed."

if
$$(x) \{ y := 1; z := 1 \}$$
 else $\{ y := 2; z := 1 \}$

has execution:



No language-level model for this!

As weak memory models are to OOO, so what is to speculation?

The Code That Never Ran

Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Introduction Spectre Optimizations Simplified Spectre Results Experiments Conclusions

・ロト ・母 ・ ・ キャ ・ ゆ ・ うへの

Imagine a SECRET, protected by a run-time security check:

if canRead(SECRET) { \dots use SECRET \dots } else { \dots }

For attacker code canRead(SECRET) is always false



Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Introduction Spectre

Optimizations

Simplified Spectre

Results

Experiments

Conclusions

Imagine a SECRET, protected by a run-time security check:

```
\mathsf{if} \mathsf{canRead}(\mathsf{SECRET}) \{ \dots \mathsf{use} \ \mathsf{SECRET} \dots \} \mathsf{else} \{ \dots \}
```

For attacker code canRead(SECRET) is always false, e.g.



is an execution of if y { if canRead(SECRET) { x := SECRET } else { x := 2 } }.



Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Introduction Spectre Optimizations Simplified Spectre Results Experiments Conclusions

▲□▶ ▲圖▶ ▲≣▶ ▲≣▶ 三三 のへで

Imagine a SECRET, protected by a run-time security check:

```
\mathsf{if} \mathsf{canRead}(\mathsf{SECRET}) \{ \dots \mathsf{use} \ \mathsf{SECRET} \dots \} \mathsf{else} \{ \dots \}
```

For attacker code canRead(SECRET) is always false, e.g.



is an execution of if y { if canRead(SECRET) { x := SECRET } else { x := 2 } }.

Attacker goal: learn if SECRET is 0 or 1.



Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Introduction Spectre Optimizations Simplified Spectre Results Experiments Conclusions

A very simplified Spectre attack:

if canRead(SECRET) { a[SECRET]:= 1 }
else if touched (a[0]) { x:= 0 }
else if touched (a[1]) { x:= 1 }

with execution

Information flow from SECRET to x, if there's an implementation of "magic".

The Code That Never Ran

Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Introduction Spectre Optimizations Simplified Spectre Results Experiments Conclusions

▲□▶ ▲圖▶ ▲圖▶ ▲圖▶ ■ のへで

A very simplified Spectre attack:

if canRead(SECRET) { a[SECRET]:= 1 }
else if touched (a[0]) { x:= 0 }
else if touched (a[1]) { x:= 1 }

with execution

Information flow from SECRET to x, if there's an implementation of "magic".

Narrator: there was one.



Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Introduction Spectre Optimizations Simplified Spectre Results Experiments Conclusions Formalization of pretty pictures as partially ordered multisets (Gisher, 1988).

Compositional semantics based on weak memory models (e.g. C11).

Examples modeling Spectre, Spectre mitigations, PRIME+ABORT attack on transactional memory...

The Code That Never Ran

Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Introduction Spectre Optimizations Simplified Spectre **Results** Experiments Conclusions

・ロト ・四ト ・ヨト ・ヨー うらぐ

Formalization of pretty pictures as partially ordered multisets (Gisher, 1988).

Compositional semantics based on weak memory models (e.g. C11).

Examples modeling Spectre, Spectre mitigations, PRIME+ABORT attack on transactional memory... and a new family of attacks on compiler optimizations.

The Code That Never Ran

Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Introduction Spectre Optimizations Simplified Spectre **Results** Experiments Conclusions

Modeling an attack on compiler optimizations

An attacker running two threads (initially x = y = 0):

$$y:=x$$
 || if $(y==0) \{ x:=1 \}$
else if $(canRead(SECRET)) \{ x:=SECRET \}$
else $\{ x:=1; z:=1 \}$

If SECRET is 1, there is an execution:



If SECRET is 2, there is no execution (due to cyclic dependency):



The Code That Never Ran

Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Introduction Spectre Optimizations Simplified Spectre **Results** Experiments Conclusions

▲□▶ ▲□▶ ▲目▶ ▲目▶ 目 のへで

Implementing attacks on compiler optimizations

Spectre and Prime+Abort are implemented.

Can we implement the attacks on compiler optimizations?



Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Introduction Spectre Optimizations Simplified Spectre Results Experiments

Conclusions

▲□▶ ▲□▶ ▲三▶ ▲三▶ 三三 - のへで

Implementing attacks on compiler optimizations

Spectre and Prime+Abort are implemented.

Can we implement the attacks on compiler optimizations?

Yes

The Code That Never Ran

Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Introduction Spectre Optimizations Simplified Spectre Results Experiments

・ロト ・四ト ・ヨト ・ヨー うへぐ

Implementing attacks on compiler optimizations

Spectre and Prime+Abort are implemented.

Can we implement the attacks on compiler optimizations?

Yes, under unrealistic assumptions:

- SECRET is a constant known at compile-time
- canRead(SECRET) is a run-time check

The Code That Never Ran

Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Introduction Spectre Optimizations Simplified Spectre Results Experiments Conclusions

▲□▶ ▲□▶ ▲目▶ ▲目▶ 目 のへの

Main attacker thread: x:=1; if (canRead(SECRET)) { x:= SECRET; } r:=y;



Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Introduction

Spectre

Optimizations

Simplified Spectre

Results

Experiments

Conclusions

```
Main attacker thread: x := 1; if (canRead(SECRET)) { x := SECRET; } r := y;
```

```
When SECRET \neq 1, gcc generates:

mov canReadSecret(%rip), %eax

mov $1, ×(%rip)

test %eax, %eax

je label1

mov $0, ×(%rip)

label1:

mov y(%rip), %eax
```

```
When SECRET = 1, gcc generates:
```

```
mov canReadSecret(%rip), %eax
mov y(%rip), %eax
mov $1, x(%rip)
```

The Code That Never Ran

```
Craig Disselkoen,
Radha Jagadeesan,
Alan Jeffrey,
James Riely
```

Introduction Spectre Optimizations Simplified Spectre Results Experiments Conclusions

Main attacker thread: x := 1; if (canRead(SECRET)) { x := SECRET; } r := y;



The Code That Never Ran

Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Introduction Spectre Optimizations Simplified Spectre Results Experiments Conclusions

▲□▶ ▲□▶ ▲目▶ ▲目▶ 目 のへで

Main attacker thread: x := 1; if (canRead(SECRET)) { x := SECRET; } r := y;



The Code That Never Ran

Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Main attacker thread: x := 1; if (canRead(SECRET)) { x := SECRET; } r := y;



Forwarding thread x := y allows attacker to spot the reordering

・ロト ・四ト ・ヨト ・ヨー うへぐ

The Code

That Never Ran Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Small delay between write x and read y: increases probability of round trip

gcc will reorder across 30 straight-line instructions

Repeat to leak multiple bits, error correction

Bitwise accuracy 99.99% at 300Kbps

The Code That Never Ran

Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Introduction Spectre Optimizations Simplified Spectre Results Experiments

・ロト・日本・モート モー うんの

Implementing an attack on dead store elimination

A similar attack targets dead store elimination

Works on clang + gcc

Bitwise accuracy 99.99% at 1.2Mbps



Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Introduction Spectre Optimizations

Simplified Spectre

Results

Experiments

Conclusions

◆□▶ ◆圖▶ ◆臣▶ ◆臣▶ 臣 - ∽��?

Contributions

A compositional model of program execution that includes speculation.

Examples including existing information flow attacks on branch prediction and transactional memory, and new attacks on optimizing compilers.

Experimental evidence that the new attacks can be carried out, but only against compile-time secrets.

(Phew, we failed to mount attacks on JIT compilers.)

https://github.com/chicago-relaxed-memory/spec-eval

The Code That Never Ran

Craig Disselkoen, Radha Jagadeesan, Alan Jeffrey, James Riely

Introduction Spectre Optimizations Simplified Spectre Results Experiments Conclusions