

Exorcising Spectres with Secure Compilers



Marco Patrignani^{1,2}

Marco Guarnieri³



CISPA
HELMHOLTZ CENTER FOR
INFORMATION SECURITY

2



Stanford
University

3

institute
imdea

Contributions & Outline

1. **Formal framework** for assessing security of Spectre v1 compiler countermeasures

Contributions & Outline

1. **Formal framework** for assessing security of Spectre v1 compiler countermeasures

2. **Proofs of security** and insecurity of existing Spectre v1 compiler countermeasures

Contributions & Outline

1. **Formal framework** for assessing security of Spectre v1 compiler countermeasures
 - Based on recent **secure compilation** theory

2. **Proofs of security** and insecurity of existing Spectre v1 compiler countermeasures

Contributions & Outline

1. **Formal framework** for assessing security of Spectre v1 compiler countermeasures
 - Based on recent **secure compilation** theory
 - Preservation of SNI: **semantic** notion of security

2. **Proofs of security** and insecurity of existing Spectre v1 compiler countermeasures

Guarnieri et al. S&P'19

Contributions & Outline

1. **Formal framework** for assessing security of Spectre v1 compiler countermeasures
 - Based on recent **secure compilation** theory
 - Preservation of SNI: **semantic** notion of security
 - **Source** semantics is standard
2. **Proofs of security** and insecurity of existing Spectre v1 compiler countermeasures

Guarnieri et al. S&P'19

Contributions & Outline

1. **Formal framework** for assessing security of Spectre v1 compiler countermeasures
 - Based on recent **secure compilation** theory
 - Preservation of SNI: **semantic** notion of security
 - **Source** semantics is standard
 - **Target** semantics is speculative
2. **Proofs of security** and insecurity of existing Spectre v1 compiler countermeasures

Guarnieri et al. S&P'19

Contributions & Outline

1. **Formal framework** for assessing security of Spectre v1 compiler countermeasures
 - Based on recent **secure compilation** theory
 - Preservation of SNI: **semantic** notion of security
 - **Source** semantics is standard
 - **Target** semantics is speculative
2. **Proofs of security** and insecurity of existing Spectre v1 compiler countermeasures
 - **Secure**: Intel ICC, SLH(ish), SSLH
 - **Insecure**: MSVC, Interprocedural SLH

Guarnieri et al. S&P'19

Contributions & Outline

1. Formal framework for assessing security of Spectre v1 compiler countermeasures

- Based on recent **secure compilation** theory
- Preservation of SNI: **semantic** notion of security
- **Source** semantics is standard

Guarnieri et al. S&P'19

(1) **Target** semantics is speculative

2. Proofs of security and insecurity of existing Spectre v1 compiler countermeasures

- **Secure**: Intel ICC, SLH(ish), SSLH
- **Insecure**: MSVC, Interprocedural SLH

Contributions & Outline

1. Formal framework for assessing security of Spectre v1 compiler countermeasures

- Based on recent **secure compilation** theory

(2) Preservation of SNI: **semantic** notion of security

Guarnieri et al. S&P'19

- **Source** semantics is standard

(1) **Target** semantics is speculative

2. Proofs of security and insecurity of existing Spectre v1 compiler countermeasures

- **Secure**: Intel ICC, SLH(ish), SSLH
- **Insecure**: MSVC, Interprocedural SLH

Contributions & Outline

(3) **Formal framework** for assessing security of Spectre v1 compiler countermeasures

- Based on recent **secure compilation** theory

(2) Preservation of SNI: **semantic** notion of security

Guarnieri et al. S&P'19

- **Source** semantics is standard

(1) **Target** semantics is speculative

2. **Proofs of security** and insecurity of existing Spectre v1 compiler countermeasures

- **Secure**: Intel ICC, SLH(ish), SSLH
- **Insecure**: MSVC, Interprocedural SLH

Contributions & Outline

(3) **Formal framework** for assessing security of Spectre v1 compiler countermeasures

- Based on recent **secure compilation** theory

(2) Preservation of SNI: **semantic** notion of security

Guarnieri et al. S&P'19

- **Source** semantics is standard

(1) **Target** semantics is speculative

2. **Proofs of security** and insecurity of existing Spectre v1 compiler countermeasures

(4) **Secure**: Intel ICC, SLH(ish), SSLH

- Insecure: MSVC, Interprocedural SLH

Contributions & Outline

(3) **Formal framework** for assessing security of Spectre v1 compiler countermeasures

- Based on recent **secure compilation** theory

(2) Observation of SNI: **semantic** notion of security

Guarnieri et al. S&P'19

- **Source** semantics is standard

(1) **Target** semantics is speculative

2. **Proofs of security** and insecurity of existing Spectre v1 compiler countermeasures

(4) **Secure**: Intel ICC, SLH(ish), SSLH

(5) **Secure**: MSVC, Interprocedural SLH

Speculative Semantics & SNI

```
void f (int x) ↦ if(x < A.size) {y = B[A[x]]}
```

```
run 1: A.size = 16, A[128] = 3
```

call f 128

Speculative Semantics & SNI

```
void f (int x) ↦ if(x < A.size) {y = B[A[x]]}
```

```
run 1: A.size = 16, A[128] = 3
```


call f 128



if (128 < 16) { y = B[A[128]] }


Speculative Semantics & SNI

```
void f (int x) ↦ if(x < A.size) {y = B[A[x]]}  
run 1: A.size = 16, A[128] = 3
```




Speculative Semantics & SNI

```
void f (int x)  $\mapsto$  if(x < A.size) {y = B[A[x]]}  
run 1: A.size = 16, A[128] = 3
```




Speculative Semantics & SNI

`void f (int x) ↦ if(x < A.size) {y = B[A[x]]}`
run 1: A.size = 16, A[128] = 3




Speculative Semantics & SNI

`void f (int x) ↦ if(x < A.size) {y = B[A[x]]}`
run 1: A.size = 16, A[128] = 3




Speculative Semantics & SNI

`void f (int x) ↦ if(x < A.size) {y = B[A[x]]}`
run 1: A.size = 16, A[128] = 3



Speculative Semantics & SNI

```
void f (int x)  $\mapsto$  if(x < A.size) {y = B[A[x]]}  
run 1: A.size = 16, A[128] = 3
```



trace 1: rd A[128]


rd B[3]

Speculative Semantics & SNI

`void f (int x) ↦ if(x < A.size) {y = B[A[x]]}`

run 1: A.size = 16, A[128] = 3

run 2: A[128] = 7 different H values



trace 1: rd A[128]


rd B[3]

Speculative Semantics & SNI

`void f (int x) ↦ if(x < A.size) {y = B[A[x]]}`

run 1: A.size = 16, A[128] = 3

run 2: A[128] = 7 different H values




trace 1: rd A[128] rd B[3]
 rd A[128]

Speculative Semantics & SNI

`void f (int x) ↦ if(x < A.size) {y = B[A[x]]}`

run 1: A.size = 16, A[128] = 3

run 2: A[128] = 7 different H values




trace 1: rd A[128] rd B[3]
 rd A[128] rd B[7]

Speculative Semantics & SNI

`void f (int x) ↦ if(x < A.size) {y = B[A[x]]}`

run 1: `A.size = 16, A[128] = 3`

run 2: `A[128] = 7` different `H` values



trace 1: `rd A[128]`

trace 2: `rd A[128]`

`rd B[3]` different traces

`rd B[7]` ⇒ SNI violation

Speculative Semantics & SNI

A program is **SNI** ($\vdash P : \text{SNI}$) if, given two runs from low-equivalent states:

- assuming the non-speculative traces are low-equivalent
- then the **speculative traces are also low-equivalent**

trace 1: rd A[128]

trace 2: rd A[128]

rd B[3] different traces


rd B[7] \Rightarrow SNI violation

Speculative Semantics & SNI

`void f (int x) ↦ if(x < A.size) {y = B[A[x]]}`

run 1: `A.size = 16, A[128] = 3`

run 2: `A[128] = 7` different `H` values



trace 1: `rd A[128]`

trace 2: `rd A[128]`

`rd B[3]` different traces

`rd B[7]` ⇒ SNI violation

Problems Problems Problems ...

Problem: Proving compiler preserves SNI is hard

Problems Problems Problems ...

Problem: Proving compiler preserves SNI is hard

Solution: overapproximate SNI with a novel property: speculative safety (SS)

Speculative Safety (SS): Taint Tracking

`void f (int x) ↦ if(x < A.size) {y = B[A[x]]}`

only 1 run needed: `A.size=16, A[128]=3`

integrity lattice: $S \subset U$ $S \sqcap U = S$ U does not flow to S

call f 128
pc : S

Speculative Safety (*SS*): Taint Tracking

`void f (int x) ↦ if(x < A.size) {y = B[A[x]]}`

only 1 run needed: `A.size=16, A[128]=3`

integrity lattice: $S \subset U$ $S \sqcap U = S$ U does not flow to S



Speculative Safety (SS): Taint Tracking

`void f (int x) \mapsto if(x < A.size) {y = B[A[x]]}`

only 1 run needed: A.size=16, A[128]=3

integrity lattice: $S \subset U$ $S \sqcap U = S$ U does not flow to S




Speculative Safety (SS): Taint Tracking

`void f (int x) \mapsto if(x < A.size) {y = B[A[x]]}`

only 1 run needed: A.size=16, A[128]=3

integrity lattice: $S \subset U$ $S \sqcap U = S$ U does not flow to S




Speculative Safety (*SS*): Taint Tracking

`void f (int x) ↦ if(x < A.size) {y = B[A[x]]}`

only 1 run needed: `A.size=16, A[128]=3`

integrity lattice: $S \subset U$ $S \sqcap U = S$ U does not flow to S




Speculative Safety (*SS*): Taint Tracking

`void f (int x) ↦ if(x < A.size) {y = B[A[x]]}`

only 1 run needed: `A.size=16, A[128]=3`

integrity lattice: $S \subset U$ $S \sqcap U = S$ U does not flow to S




Speculative Safety (SS): Taint Tracking

`void f (int x) \mapsto if(x < A.size) {y = B[A[x]]}`

only 1 run needed: A.size=16, A[128]=3

integrity lattice: $S \subset U$ $S \sqcap U = S$ U does not flow to S



rd A[128] : S

rd B[3] : U

Speculative Safety (*SS*): Taint Tracking

```
void f (int x)  $\mapsto$  if(x < A.size) {y = B[A[x]]}  
only 1 run needed: A.size=16, A[128]=3
```

A program is **SS** ($\vdash \mathbf{P} : \mathbf{SS}$) if its traces do not contain **U** actions

call f 128
pc : S

→ if (128 < 16) { y = B[A[128]] }
pc : S

rd A[128] : S

rd B[3] : U

Speculative Safety (SS): Taint Tracking

`void f (int x) \mapsto if(x < A.size) {y = B[A[x]]}`

only 1 run needed: A.size=16, A[128]=3

integrity lattice: $S \subset U$ $S \sqcap U = S$ U does not flow to S



rd A[128] : S

rd B[3] : U

SS and SNI

SS overapproximates SNI, so:

- in the **target**: $\forall P \vdash P : \mathbf{SS} \Rightarrow P : \mathbf{SNI}$

SS and SNI

SS overapproximates SNI, so:

- in the **target**: $\forall P \vdash P : \mathbf{SS} \Rightarrow P : \mathbf{SNI}$
- in the **source**: $\forall P \vdash P : \mathbf{SS} \iff P : \mathbf{SNI}$
(recall, no speculative execution in **source**)

SS-Preserving Compiler: RSSC & RSSP

$\llbracket \cdot \rrbracket : \text{RSSP} \stackrel{\text{def}}{=} \text{if } \forall A.A [P] : SS \text{ then } \forall A.A [\llbracket P \rrbracket] : SS$

SS-Preserving Compiler: RSSC & RSSP

$\llbracket \cdot \rrbracket : \text{RSSP} \stackrel{\text{def}}{=} \text{if } \forall A.A [P] : SS \text{ then } \forall A.A [\llbracket P \rrbracket] : SS$

$\llbracket \cdot \rrbracket : \text{RSSC} \stackrel{\text{def}}{=} \text{if } \forall A.A [\llbracket P \rrbracket] \rightsquigarrow \mathbf{m} \text{ then } \exists A.A [P] \rightsquigarrow \mathbf{m} \approx \mathbf{m}$

$\approx =$ same traces, plus **S** actions in **m**

SS-Preserving Compiler: RSSC & RSSP

$\llbracket \cdot \rrbracket : \text{RSSP} \stackrel{\text{def}}{=} \text{if } \forall A.A [P] : SS \text{ then } \forall A.A [\llbracket P \rrbracket] : SS$

$\llbracket \cdot \rrbracket : \text{RSSC} \stackrel{\text{def}}{=} \text{if } \forall A.A [\llbracket P \rrbracket] \rightsquigarrow \mathbf{m} \text{ then } \exists A.A [P] \rightsquigarrow \mathbf{m} \approx \mathbf{m}$

$\approx =$ same traces, plus **S** actions in **m**

- \forall attackers: explicit attacker model
robustness

SS-Preserving Compiler: RSSC & RSSP

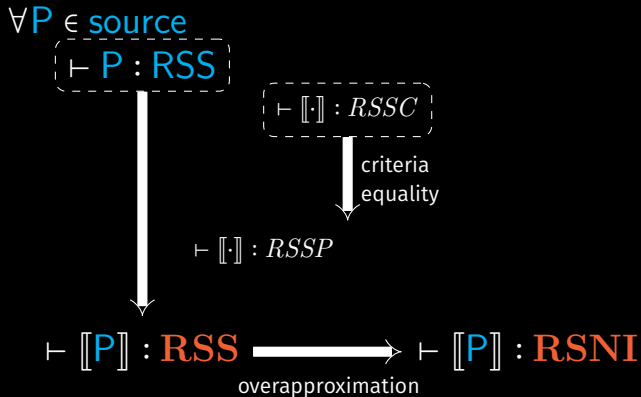
$\llbracket \cdot \rrbracket : \text{RSSP} \stackrel{\text{def}}{=} \text{if } \forall A.A [P] : SS \text{ then } \forall A.A [\llbracket P \rrbracket] : SS$

$\llbracket \cdot \rrbracket : \text{RSSC} \stackrel{\text{def}}{=} \text{if } \forall A.A [\llbracket P \rrbracket] \rightsquigarrow m \text{ then } \exists A.A [P] \rightsquigarrow m \approx m$


$\approx =$ same traces, plus **S** actions in **m**

- \forall attackers: explicit attacker model
robustness
- Proof: RSSC & RSSP are **equivalent**
RSSC : clear security guarantees
RSSP : simpler proofs

Secure Compilation Framework for Spectre

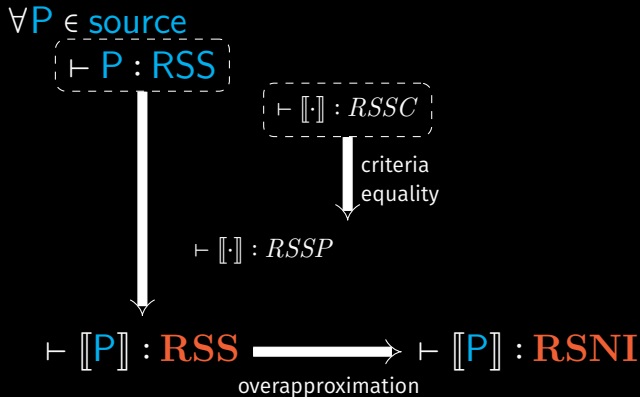


Secure Compilation Framework for Spectre



- dashed premises are already discharged

Secure Compilation Framework for Spectre



- dashed premises are already discharged
- to show security: **simply prove** RSSP

Security Spectrum

- 2 notions of SS and SNI (thus 2 **targets**):
 - **strong(+)**: no speculative leaks

Security Spectrum

- 2 notions of SS and SNI (thus 2 **targets**):
 - **strong**(+): no speculative leaks
 - **weak**(-): allows speculative leaks of data retrieved non-speculatively

Security Spectrum

- 2 notions of SS and SNI (thus 2 **targets**):
 - **strong(+)**: no speculative leaks
 - **weak(-)**: allows speculative leaks of data retrieved non-speculatively

```
1 void get (int y)
2   if (y < size) then
3     temp = B[A[y]*512]
```

Violates + and -

```
1 void get (int y)
2   x = A[y];
3   if (y < size) then
4     temp = B[x];
```

Violates +, Satisfies -

RSSC **for** lfence


```
void f(int x) ↦ if(x < A.size){y = B[A[x]]} // A.size=16, A[128]=3  
[[·]] = void f(int x) ↦ if(x < A.size){lfence; y = B[A[x]]}
```

call f 128

pc: 5


RSSC **for** lfence

```
void f(int x) ↦ if(x < A.size){y = B[A[x]]} // A.size=16, A[128]=3  
[[·]] = void f(int x) ↦ if(x < A.size){lfence; y = B[A[x]]}
```




RSSC **for** lfence

```
void f(int x) ↦ if(x < A.size){y = B[A[x]]} // A.size=16, A[128]=3  
[[·]] = void f(int x) ↦ if(x < A.size){lfence; y = B[A[x]]}
```



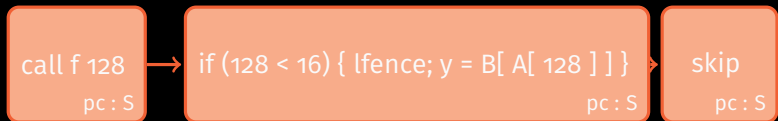
RSSC **for** lfence

```
void f(int x) ↦ if(x < A.size){y = B[A[x]]} // A.size=16, A[128]=3  
[[·]] = void f(int x) ↦ if(x < A.size){lfence; y = B[A[x]]}
```



RSSC **for** lfence

```
void f(int x) ↦ if(x < A.size){y = B[A[x]]} // A.size=16, A[128]=3  
[[·]] = void f(int x) ↦ if(x < A.size){lfence; y = B[A[x]]}
```



RSSC **for** SLH

```
void f(int x) ↦ if(x < A.size){y = B[A[x]]} // A.size=16, A[128]=3  
[[·]] = void f(int x) ↦ if(x < A.size){y = B[mask(A[x])]}
```

call f 128
pc : S

RSSC **for** SLH

```
void f(int x) ↦ if(x < A.size){y = B[A[x]]} // A.size=16, A[128]=3  
[[·]] = void f(int x) ↦ if(x < A.size){y = B[mask(A[x])]}
```

call f 128
pc : S



if (128 < 16) { y = B[mask(A[128])] }

pc : S

RSSC **for** SLH

```
void f(int x) ↦ if(x < A.size){y = B[A[x]]} // A.size=16, A[128]=3  
[[·]] = void f(int x) ↦ if(x < A.size){y = B[mask(A[x])]}
```

call f 128
pc : S




if (128 < 16) { y = B[mask(A[128])] }

pc : S


RSSC for SLH

```
void f(int x) ↦ if(x < A.size){y = B[A[x]]} // A.size=16, A[128]=3  
[[·]] = void f(int x) ↦ if(x < A.size){y = B[mask(A[x])]}
```




RSSC for SLH

```
void f(int x) ↦ if(x < A.size){y = B[A[x]]} // A.size=16, A[128]=3  
[[·]] = void f(int x) ↦ if(x < A.size){y = B[mask(A[x])]}
```




RSSC for SLH

```
void f(int x) ↦ if(x < A.size){y = B[A[x]]} // A.size=16, A[128]=3  
[[·]] = void f(int x) ↦ if(x < A.size){y = B[mask(A[x])]}
```




RSSC for SLH

```
void f(int x) ↦ if(x < A.size){y = B[A[x]]} // A.size=16, A[128]=3  
[[·]] = void f(int x) ↦ if(x < A.size){y = B[mask(A[x])]}
```



RSSC for SLH

`void f(int x) ↦ if(x < A.size){y = B[A[x]]}` // A.size=16, A[128]=3
`[·] = void f(int x) ↦ if(x < A.size){y = B[mask(A[x])]}`



SLH preserves SS- (and thus SNI-) but
not SS+ (and thus not SNI+)

Framework benefits: **fine-grained
analysis** of countermeasures security

rd B[0] : S

Insecurity Results

- MSVC is Insecure
- Non-interprocedural SLH is insecure

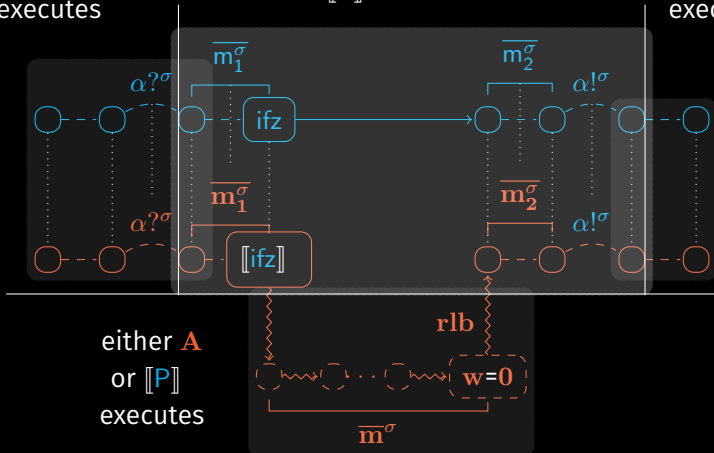
Both omit speculation barriers

Proofs Insight

$\langle\langle A \rangle\rangle / A$
executes

$P / \llbracket P \rrbracket$ executes

$\langle\langle A \rangle\rangle / A$
executes



either A
or $\llbracket P \rrbracket$
executes