

# Binsec/Rel



## Binsec/Rel Symbolic Binary Analyzer for Security

*Application to Constant-Time & Secret-Erasure*

Lesly-Ann Daniel, KU Leuven



Sébastien Bardin, CEA List

Tamara Rezk, INRIA

# Timing and Microarchitectural Attacks

## Timing and microarchitectural attacks:


Execution time & microarchitectural state **depends on secret** data



First timing attack in **1996** by Paul Kocher: full recovery of **RSA encryption key**

# Protect software with constant-time programming

**Constant-Time.** Execution time / changes to microarchitectural state must be independent from secret input



Already used in many cryptographic implementations

# What can influence execution time/microarchitecture?

## Control Flow

```
if secret
```

```
then foo()
```



```
else bar()
```



secret



~~secret~~

# What can influence execution time/microarchitecture?

## Control Flow

```
if secret
```

```
then foo()
```

```
else bar()
```



secret




~~secret~~

## Memory Accesses

```
x = buf[secret]
```

### Cache



# What can influence execution time/microarchitecture?

## Control Flow

```
if secret
```

```
then foo()
```

```
else bar()
```



secret




~~secret~~

## Memory Accesses

```
x = buf[secret]
```

Cache

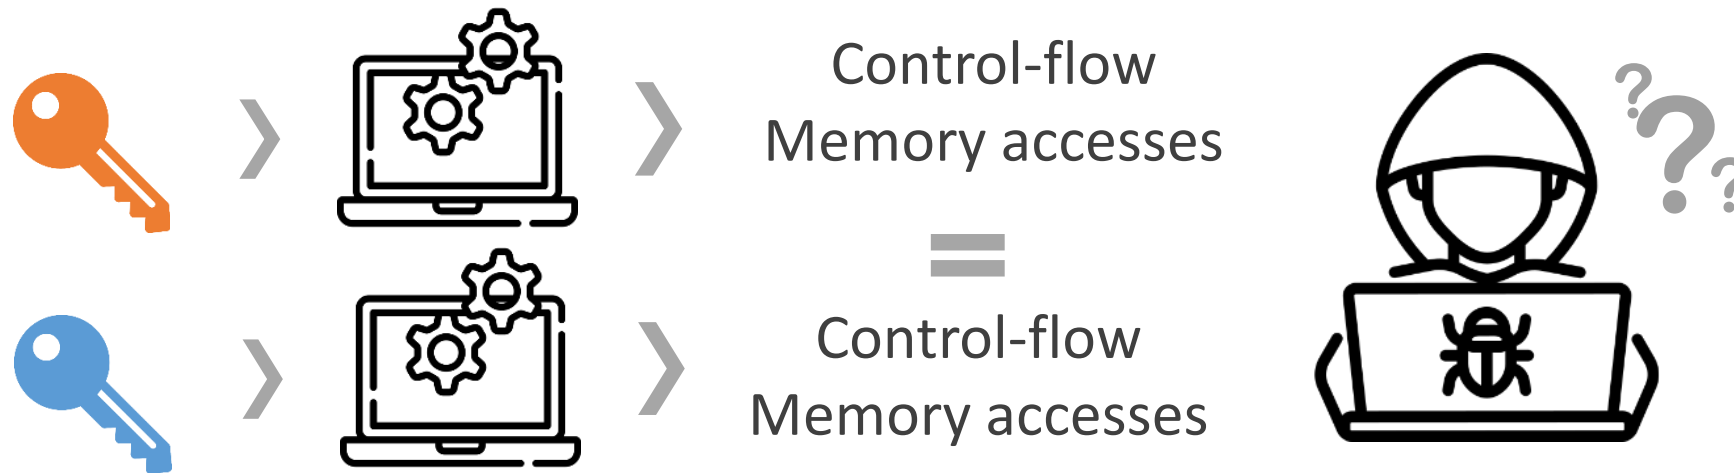


secret




# Protect software with constant-time programming

**Constant-Time.** **Control-flow** and **memory accesses** must be independent from secret input



# Protect software with constant-time programming

**Constant-Time.** **Control-flow** and **memory accesses** must be independent from secret input



Property relating *2 execution traces* (2-hypersafety)



# Constant-time is not easy to implement

```
uint32_t select(uint32_t x, uint32_t y, bool secret) {  
    if (secret) return x;  
    else return y;  
}
```



```
uint32_t ct_select(uint32_t x, uint32_t y, bool secret) {  
    signed b = 0 - secret;  
    return (x & b) | (y & ~b);  
}
```



# Compilers can break constant-time!

```
uint32_t ct_select(uint32_t x, uint32_t y, bool secret) {  
    signed b = 0 - secret;  
    return (x & b) | (y & ~b);  
}
```



```
public ct_select_u32_v4  
ct_select_u32_v4 proc near  
  
var_14= dword ptr -14h  
var_D= byte ptr -0Dh  
var_C= dword ptr -0Ch  
var_8= dword ptr -8  
arg_0= dword ptr 4  
arg_4= dword ptr 8  
arg_8= byte ptr 0Ch  
  
push esi  
sub esp, 10h  
mov al, [esp+14h+arg_8]  
mov ecx, [esp+14h+arg_4]  
mov edx, [esp+14h+arg_0]  
mov [esp+14h+var_8], edx  
mov [esp+14h+var_C], ecx  
and al, 1  
mov [esp+14h+var_D], al  
mov al, [esp+14h+var_D]  
and al, 1  
movzx ecx, al  
mov edx, 0  
sub edx, ecx  
mov [esp+14h+var_14], edx  
mov ecx, [esp+14h+var_8]  
and ecx, [esp+14h+var_14]  
mov edx, [esp+14h+var_C]  
mov esi, [esp+14h+var_14]  
xor esi, 0FFFFFFFFh  
and esi, edx  
or esi, ecx  
mov eax, esi  
add esp, 10h  
pop esi  
retn  
ct_select_u32_v4 endp
```



clang-3.0 -O0


clang-3.0 -O3

```
public ct_select_u32_v4  
ct_select_u32_v4 proc near  
  
arg_0= byte ptr 4  
arg_4= byte ptr 8  
arg_8= byte ptr 0Ch  
  
mov al, [esp+arg_8]  
test al, al  
jz short loc_804842F  
  
lea eax, [esp+arg_0]  
mov eax, [eax]  
retn  
  
loc_804842F:  
lea eax, [esp+arg_4]  
mov eax, [eax]  
retn  
ct_select_u32_v4 endp
```



Simon, Laurent, David Chisnall, and Ross Anderson. "What you get is what you C: Controlling side effects in mainstream C compilers."  
2018 IEEE European Symposium on Security and Privacy (EuroS&P).

# Automated program verification



## Perfect verification tool:

- Reject only **insecure** programs
- Accept only **secure** programs
- Always **terminate**
- Be fully **automatic**



## Not possible:

Non trivial semantic properties of programs are **undecidable**  
*Rice Theorem (1951)*

# Automated program verification



## Perfect verification tool:

- Reject only **insecure** programs
- Accept only **secure** programs **up to a given bound**
- Always **terminate**
- Be fully **automatic**

## Symbolic Execution (SE)



The KeY Project




BINSEC



# Challenges: SE for constant-time analysis

Property of 2 executions



**ReISE**

SE for pairs of traces with **sharing**



Not necessarily preserved by compilers



**Binary-analysis**

*Reason explicitly about memory*

**Does not scale** 😞

# Many verification tools for constant-time but...

	Target	Bounded-Verif	Bug-Finding
CT-SC [1] & CT-AI [2]	C	✓ <sup>+</sup>	✗
Casym [4] & CT-Verif [3]	LLVM	✓ <sup>+</sup>	✗
CacheAudit [5]	Binary	✓ <sup>+</sup>	✗
CacheD [6]	Binary	✗	✓

C/LLVM analysis might miss constant-time violations 😞

<sup>+</sup> Full proof

[1] J. Bacelar Almeida, M. Barbosa, J. S. Pinto, and B. Vieira, “Formal verification of side-channel countermeasures using self-composition,” in Science of Computer Programming, 2013

[2] S. Blazy, D. Pichardie, and A. Trieu, “Verifying Constant-Time Implementations by Abstract Interpretation,” in ESORICS, 2017

[3] J. B. Almeida, M. Barbosa, G. Barthe, F. Dupressoir, and M. Emmi, “Verifying Constant-Time Implementations..” in USENIX, 2016

[4] R. Brotzman, S. Liu, D. Zhang, G. Tan, and M. Kandemir, “CaSym: Cache aware symbolic execution for side channel detection and mitigation,” in IEEE SP, 2019

[5] G. Doychev and B. Köpf, “Rigorous analysis of software countermeasures against cache attacks,” in PLDI, 2017

[6] S. Wang, P. Wang, X. Liu, D. Zhang, and D. Wu, “CacheD: Identifying cache-based timing channels in production software,” in USENIX, 2017

# Many verification tools for constant-time but...

	Target	Bounded-Verif	Bug-Finding
CT-SC [1] & CT-AI [2]	C	✓ <sup>+</sup>	✗
Casym [4] & CT-Verif [3]	LLVM	✓ <sup>+</sup>	✗
CacheAudit [5]	Binary	✓ <sup>+</sup>	✗
CacheD [6]	Binary	✗	✓
Binsec/Rel	Binary	✓	✓

C/LLVM analysis might miss constant-time violations 😞

+ Full proof

[1] J. Bacelar Almeida, M. Barbosa, J. S. Pinto, and B. Vieira, "Formal verification of side-channel countermeasures using self-composition," in Science of Computer Programming, 2013

[2] S. Blazy, D. Pichardie, and A. Trieu, "Verifying Constant-Time Implementations by Abstract Interpretation," in ESORICS, 2017

[3] J. B. Almeida, M. Barbosa, G. Barthe, F. Dupressoir, and M. Emmi, "Verifying Constant-Time Implementations.," in USENIX, 2016

[4] R. Brotzman, S. Liu, D. Zhang, G. Tan, and M. Kandemir, "CaSym: Cache aware symbolic execution for side channel detection and mitigation," in IEEE SP, 2019

[5] G. Doychev and B. Köpf, "Rigorous analysis of software countermeasures against cache attacks," in PLDI, 2017

[6] S. Wang, P. Wang, X. Liu, D. Zhang, and D. Wu, "CacheD: Identifying cache-based timing channels in production software," in USENIX, 2017

# Contributions

- **Optimizations**: symbolic execution for constant-time + secret-erasure
- **Implementation** in an open source tools



- **Application** to cryptographic primitives
  - Violations introduced by compilers from verified llvm code



Background: SE for constant-time

# Symbolic Execution [1,2]

```
foo(public p, secret s) {  
  c := p * s - 48;  
  if(c = 0) error();  
  else return s/c;  
}
```

Can **error** be reached?

[1] James C. King. *Symbolic execution and program testing*, Communications of the ACM, 1976

[2] Cristian Cadar and Sen Koushik. *Symbolic execution for software testing: three decades later*. Communications of the ACM, 2013

# Symbolic Execution [1,2]

```
foo(public p, secret s) {  
  c := p * s - 48;  
  if(c = 0) error();  
  else return s/c;  
}
```



**Symbolic store**

$p \mapsto p$   
 $s \mapsto s$

Can **error** be reached?

[1] James C. King. *Symbolic execution and program testing*, Communications of the ACM, 1976

[2] Cristian Cadar and Sen Koushik. *Symbolic execution for software testing: three decades later*. Communications of the ACM, 2013

# Symbolic Execution [1,2]

```
foo(public p, secret s) {  
  c := p * s - 48;  
  if(c = 0) error();  
  else return s/c;  
}
```



**Symbolic store**

$p \mapsto p$

$s \mapsto s$

$c \mapsto p \times s - 48$

Can **error** be reached?

[1] James C. King. *Symbolic execution and program testing*, Communications of the ACM, 1976

[2] Cristian Cadar and Sen Koushik. *Symbolic execution for software testing: three decades later*. Communications of the ACM, 2013

# Symbolic Execution [1,2]

```
foo(public p, secret s) {  
  c := p * s - 48;  
  if (c = 0) error();  
  else return s/c;  
}
```

Can **error** be reached?




Symbolic store

$p \mapsto p$

$s \mapsto s$

$c \mapsto p \times s - 48$

Path predicate



Formula  $F(p, s)$

$c = p \times s - 48 \wedge c = 0$

[1] James C. King. *Symbolic execution and program testing*, Communications of the ACM, 1976

[2] Cristian Cadar and Sen Koushik. *Symbolic execution for software testing: three decades later*. Communications of the ACM, 2013

# Symbolic Execution [1,2]

```
foo(public p, secret s) {  
  c := p * s - 48;  
  if(c = 0) error();  
  else return s/c;  
}
```

Can **error** be reached?

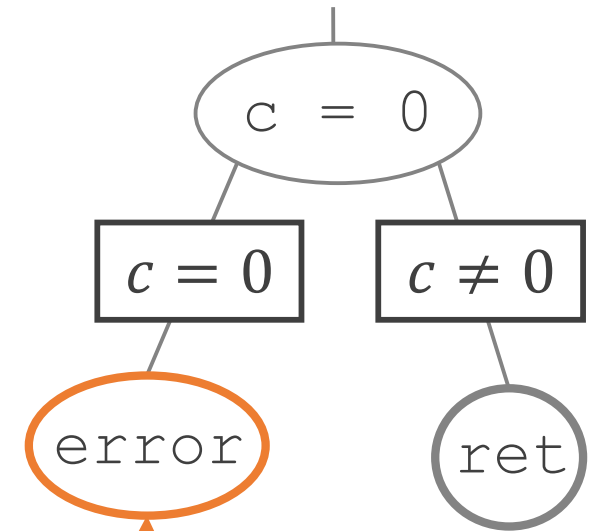
Symbolic store

$p \mapsto p$


$s \mapsto s$

$c \mapsto p \times s - 48$

Path predicate



SMT-Solver



Formula  $F(p, s)$

$c = p \times s - 48 \wedge c = 0$

[1] James C. King. *Symbolic execution and program testing*, Communications of the ACM, 1976

[2] Cristian Cadar and Sen Koushik. *Symbolic execution for software testing: three decades later*. Communications of the ACM, 2013

# SE for constant-time via self-composition [1]

```
foo(public p, secret s) {  
  c := p * s - 48;  
  if(c = 0) error();  
  else return s/c;  
}
```

Can  $c = 0$  depend on  $s$ ?



Symbolic Execution

Formula  $F(p, s)$

$c = p \times s - 48 \wedge c = 0$

# SE for constant-time via self-composition [1]

```
foo(public p, secret s) {  
  c := p * s - 48;  
  if(c = 0) error();  
  else return s/c;  
}
```



Symbolic Execution

Formula  $F(p, s)$

$$c = p \times s - 48 \wedge c = 0$$


Can  $c = 0$  depend on  $s$ ?

Self-composition:  $F(p, s, p', s')$

$$p = p' \wedge \begin{matrix} c = p \times s - 48 \\ c' = p' \times s' - 48 \end{matrix} \wedge c = 0 \neq c' = 0$$



SMT-Solver



$p = 6, s = 8$   
 $p' = 6, s' = 1$





# SE for constant-time via self-composition

## Limitations

- Whole formula is duplicated  $F(p, s, p', s')$
- High number of insecurity queries to the solver

*Relational Symbolic Execution to overcome these limitation*

# Better approach: Relational SE [1,2]

```
foo(public p, secret s) {  
  c := p * s - 48;  
  if(c = 0) error();  
  else return s/c;  
}
```



## Symbolic store

p  $\mapsto$   $\langle p \rangle$

s  $\mapsto$   $\langle s \mid s' \rangle$

c  $\mapsto$   $\langle p \times s - 48 \mid p \times s' - 48 \rangle$

Sharing in SE 

[1] "Shadow of a doubt", Palikareva, Kuchta, and Cadar 2016

[2] "Relational Symbolic Execution", Farina, Chong, and Gaboardi 2017

# Better approach: Relational SE [1,2]

```
foo(public p, secret s) {  
  c := p * s - 48;  
  if(c = 0) error();  
  else return s/c;  
}
```



Symbolic store

$p \mapsto \langle p \rangle$

$s \mapsto \langle s \mid s' \rangle$

$c \mapsto \langle p \times s - 48 \mid p \times s' - 48 \rangle$


Sharing in SE

Relational formula:  $F(p, s, s')$

$c = p \times s - 48$   
 $c' = p \times s' - 48$   
 $\wedge c = 0 \neq c' = 0$



SMT-Solver



$p = 6$   
 $s = 8$   $s' = 1$

[1] "Shadow of a doubt", Palikareva, Kuchta, and Cadar 2016

[2] "Relational Symbolic Execution", Farina, Chong, and Gaboardi 2017

# Better approach: Relational SE [1,2]

```
foo(public p, secret s) {  
  c := p - 48;  
  if(c = 0) error();  
  else return s/c;  
}
```



## Symbolic store

p  $\mapsto$   $\langle p \rangle$   
s  $\mapsto$   $\langle s \mid s' \rangle$   
c  $\mapsto$   $\langle p - 48 \rangle$

[1] "Shadow of a doubt", Palikareva, Kuchta, and Cadar 2016

[2] "Relational Symbolic Execution", Farina, Chong, and Gaboardi 2017

# Better approach: Relational SE [1,2]


```
foo(public p, secret s) {  
  c := p - 48;  
  if(c = 0) error();  
  else return s/c;  
}
```





## Symbolic store

p  $\mapsto$   $\langle p \rangle$   
s  $\mapsto$   $\langle s \mid s' \rangle$   
c  $\mapsto$   $\langle p - 48 \rangle$

Spared query !



Sharing in SE   
Secret tracking 

[1] "Shadow of a doubt", Palikareva, Kuchta, and Cadar 2016

[2] "Relational Symbolic Execution", Farina, Chong, and Gaboardi 2017

# Limitations of ReSE

## Problem:

- Memory = symbolic array  $\langle \mu \mid \mu' \rangle$
- Duplicate load operations  $\langle \text{select } \mu (\text{esp} - 4) \mid \text{select } \mu' (\text{esp} - 4) \rangle$
- Many loads in binary code ☹️

*ReSE is inefficient at binary-level*  
*ReSE cannot efficiently model speculations*

# Binary-level RelSE

## On-the-fly read-over-write

- Relational expressions in memory
- Builds on *read-over-write* [1]
- Simplify loads on-the-fly

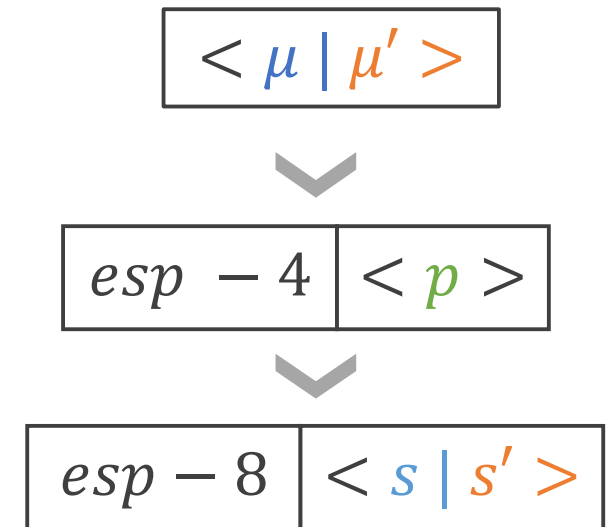
[1] Farinier B, David R, Bardin S, Lemerre M. *Arrays Made Simpler: An Efficient, Scalable and Thorough Preprocessing*. LPAR 2018

# Binary-level ReSE

## On-the-fly read-over-write

- Relational expressions in memory
- Builds on *read-over-write* [1]
- Simplify loads on-the-fly

## Memory as the history of stores.



[1] Farinier B, David R, Bardin S, Lemerre M. *Arrays Made Simpler: An Efficient, Scalable and Thorough Preprocessing*. LPAR 2018



# Binary-level RelSE


## On-the-fly read-over-write

- Relational expressions in memory
- Builds on *read-over-write* [1]
- Simplify loads on-the-fly

### Example.

load  $esp-4$  returns  $\langle p \rangle$  instead of  
 $\langle \text{select } \mu(esp-4) \mid \text{select } \mu'(esp-4) \rangle$

## Memory as the history of stores.



[1] Farinier B, David R, Bardin S, Lemerre M. *Arrays Made Simpler: An Efficient, Scalable and Thorough Preprocessing*. LPAR 2018

# Dedicated optimizations for constant-time

## Untainting

Use solver response to transform

$\langle a \mid a' \rangle$  to  $\langle a \rangle$

- Better sharing
- Better secret tracking

## Fault-Packing

Pack queries along basic-blocks

- Reduces number of queries
- Useful for constant-time analysis (many queries)


# Implementation



# Binsec Framework




## Binary



X86-32 (64 incoming)  
ARMv7/AARCH64/AMD64  
RISC-V 32 (64 incoming)



## IR



## Analysis



Symbolic execution  
Backward-bounded SE  
Relational SE  
Abstract interpretation  
Concrete interpretation


## Helpers



Loader for ELF/PE  
Build & simplify formulas  
[...]



## SMT-Solver




Boolector  
Bitwuzla  
z3, cvc4, yices

<https://binsec.github.io/>


# Binsec/Rel

Binary



Easily specify secrets using dedicated stubs

Initial memory configuration



Concretize esp, .data, canaries, ...  
Default = symbolic

## Binsec/Rel

Exploration module	Insecurity module
<ul style="list-style-type: none"><li>• Updates sym. state</li><li>• Build path predicate</li><li>• Check satisfiability</li></ul>	<ul style="list-style-type: none"><li>• Build insecurity queries<ul style="list-style-type: none"><li>• Constant-time</li><li>• Secret-erasure</li></ul></li><li>• Ensure unsatisfiability</li></ul>



- 1 Only if exhaustive exploration
- 2 Violation + counterexample (concrete input)
- 3 No violations but non-exhaustive exploration

# Limitations

## Bounded-verification

- Loop & recursion by unrolling
- Bounded enumeration of jump targets



Can miss violations so Binsec/Rel reports “Unknown”

## Implementation

- No dynamic libraries
- No dynamically allocated memory
- No syscall stubs
- No floating-point instructions

**Keep in mind:** when you *concretize* something (e.g. input size, initial memory, etc.) it might lead to unexplored behavior & *missed violations*

# Experimental evaluation

[https://github.com/binsec/rel\\_bench](https://github.com/binsec/rel_bench)

# Ablation study: Binsec/Rel vs. vanilla ReISE

	Instructions	Instructions / sec	Time	Timeouts
ReISE	349k	6.2	15h47	13
Binsec/Rel	23M	4429	1h26	0

*Total on 338 cryptographic samples (secure & insecure)  
Timeout set to 1h*

Binsec/Rel **700×** faster than ReISE  
**No timeouts** even on large programs (e.g. donna)



# Preservation of constant-time by compilers

Prior *manual* study on constant-time bugs introduced by compilers [1]

- We *automate* this study with Binsec/Rel
- We *extend* this study:
  - 34 functions
  - i386 / i686 / ARM architectures
  - 6 gcc + 6 clang version
  - 4 optimization level
  - impact of `-x86-cmov-converter` & `if-conversion`

Total
4148 binaries

- `clang backend passes` introduce violations in programs deemed *secure* by `llvm` analyzers
- clang use of `cmov` can introduce *secret-dependent memory accesses*
- `gcc` optimizations tend to preserve CT (`if-conversion` can even make secure non-ct source)
- Depend on multiple factors, hard to predict: compiler-support remains the best option



[1] “What you get is what you C”, Simon, Chisnall, and Anderson 2018

# Conclusion

# Conclusion

- Dedicated **optimizations** for RelSE at binary-level
  - Sharing between pairs of executions
- Open source tool **Binsec/Rel**
  - Bug-finding & bounded-verification of constant-time at binary-level
- Analysis of crypto libraries at binary-level
  - Constant-time **llvm** may yield **vulnerable binary**



# Extensions

- Binsec/Rel for **secret erasure**
  - Framework to check preservation of secret-erasure by compilers  
17 scrubbing functions / 10 compilers / 4 opt. level + DSE pass / total = 1156 binaries  
Open source & easy to extend on [https://github.com/binsec/rel\\_bench](https://github.com/binsec/rel_bench)
- Binsec/Haunted to find **Spectre-PHT/STL** vulnerabilities



<https://github.com/binsec/haunted>

# Future of Binsec/Rel

- Binsec/Rel not really maintained but...
- Binsec team is working on integrating Binsec/Rel in Binsec
  - Better (relational) symbolic execution engine
  - Better maintenance
  - Tutorials
- Any feedback is welcome:
  - [sebastien.bardin@cea.fr](mailto:sebastien.bardin@cea.fr)
  - [frederic.recoules@cea.fr](mailto:frederic.recoules@cea.fr)



<https://binsec.github.io/>

Backup

# Extension: Secret-erasure

# Secret-erasure

```
void scrub(char * buf, size_t size){
    memset(buf, 0, size );
}

int critical_function () {
    char secret [SIZE];
    read_secret(secret, SIZE);
    process_secret(secret, SIZE); // computation on secret
    scrub(secret, SIZE); // erase secret from memory
    return 0;
}
```





# Secret-erasure


```
void scrub(char * buf, size_t size){
    memset(buf, 0, size );
}

int critical_function () {
    char secret [SIZE];
    read_secret(secret, SIZE);
    process_secret(secret, SIZE); // computation on secret
    scrub(secret, SIZE); // erase secret from memory
    return 0;
}
```



gcc **-O2**

Dead store elimination pass  
removes memset call



- Crucial for **cryptographic** code
- Property of **2 executions**
- Not always preserved by **compilers**

# Generalizing Binary-level RelSE

- Binary-level RelSE **parametric** in the **leakage model**
  - *Symbolic leakage predicate* instantiated according to leakage model
  - For IF properties restricting to *pairs of traces following same path*

$$\frac{\mathbb{P}[l] = \mathbf{halt} \quad \boxed{\tilde{\lambda}_{\perp}(\pi, \hat{\mu})}}{(l, \rho, \hat{\mu}, \pi) \rightsquigarrow (l, \rho, \hat{\mu}, \pi)}$$

- New leakage model + property for **capturing secret-erasure**
  - *Leaks value of all store operations that are not overwritten*
  - *Forbids secret dependent control-flow*
- Adaptation of **Binsec/Rel** to **secret-erasure**

# Application: Secret-Erasure

New framework to check secret-erasure

*Easily extensible with new compilers and new scrubbing functions*

- We analyze **17 scrubbing functions**
- 5 versions of clang & 5 versions of gcc
- 4 optimization levels + DSE pass



Total

1156 binaries

- Dedicated secure scrubbing functions (e.g. `memset_s`) are secure
- Disabling DSE sometimes works but is *not always sufficient*
- **Volatile function pointers** can introduce additional **register spilling** that might **break secret-erasure** with `gcc -O2` and `gcc -O3`



# Extension: Spectre



Haunted ReISE: detect Spectre vulnerabilities



# Spectre-PHT

## Spectre-PHT

Exploits conditional branch predictor

```
if idx < size {  
    v = tab[idx]  
    leak(v)  
}
```

- `idx` is attacker controlled
- content of `tab` is public
- `leak(v)` encodes `v` to cache

## Sequential execution

- Conditional bound check ensures `idx` is in bounds
- `v` contains public data

# Spectre-PHT

## Spectre-PHT

Exploits conditional branch predictor

```
if idx < size {  
    v = tab[idx]  
    leak(v)  
}
```

- `idx` is attacker controlled
- content of `tab` is public
- `leak(v)` encodes `v` to cache

## Sequential execution

- Conditional bound check ensures `idx` is in bounds
- `v` contains public data

## Transient Execution

- Conditional is misspeculated
- Out-of-bound array access  
→ load secret data in `v`
- `v` is leaked to the cache



# Spectre-STL

**Spectre-STL:** Loads can speculatively bypass prior stores

## Sequential execution

```
store a s  
store a p  
store b q  
v = load a  
leak(v)
```

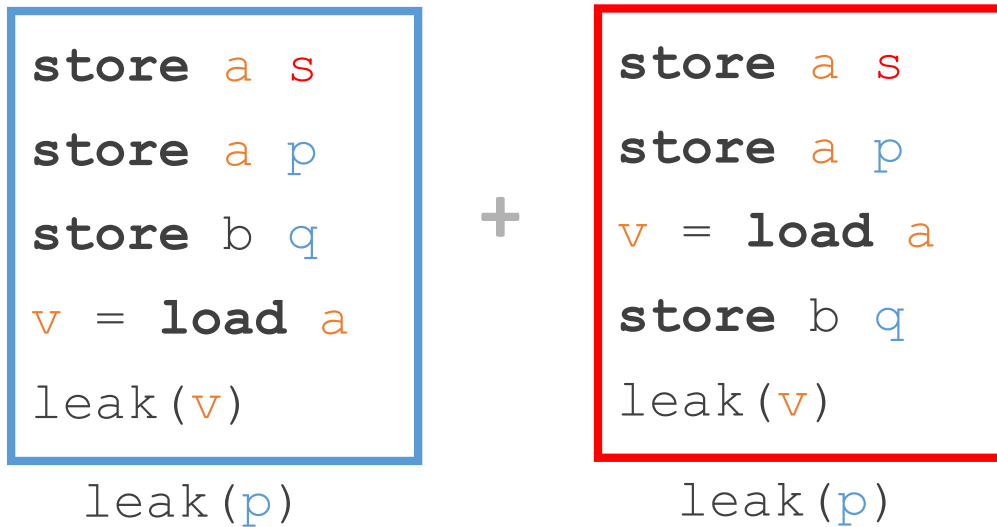
```
leak(p)
```

- where `s` is secret, `p` and `q` are public
- where `a`  $\neq$  `b`
- `leak(v)` encodes `v` to cache

# Spectre-STL

**Spectre-STL:** Loads can speculatively bypass prior stores

**Sequential execution** + **Transient Executions**



- where `s` is secret, `p` and `q` are public
- where `a`  $\neq$  `b`
- `leak(v)` encodes `v` to cache



# Spectre-STL

**Spectre-STL:** Loads can speculatively bypass prior stores

**Sequential execution** + **Transient Executions**

```
store a s
store a p
store b q
v = load a
leak(v)
```

leak(p)

+

```
store a s
store a p
v = load a
store b q
leak(v)
```

leak(p)

+

```
store a s
v = load a
store a p
store b q
leak(v)
```

leak(s)

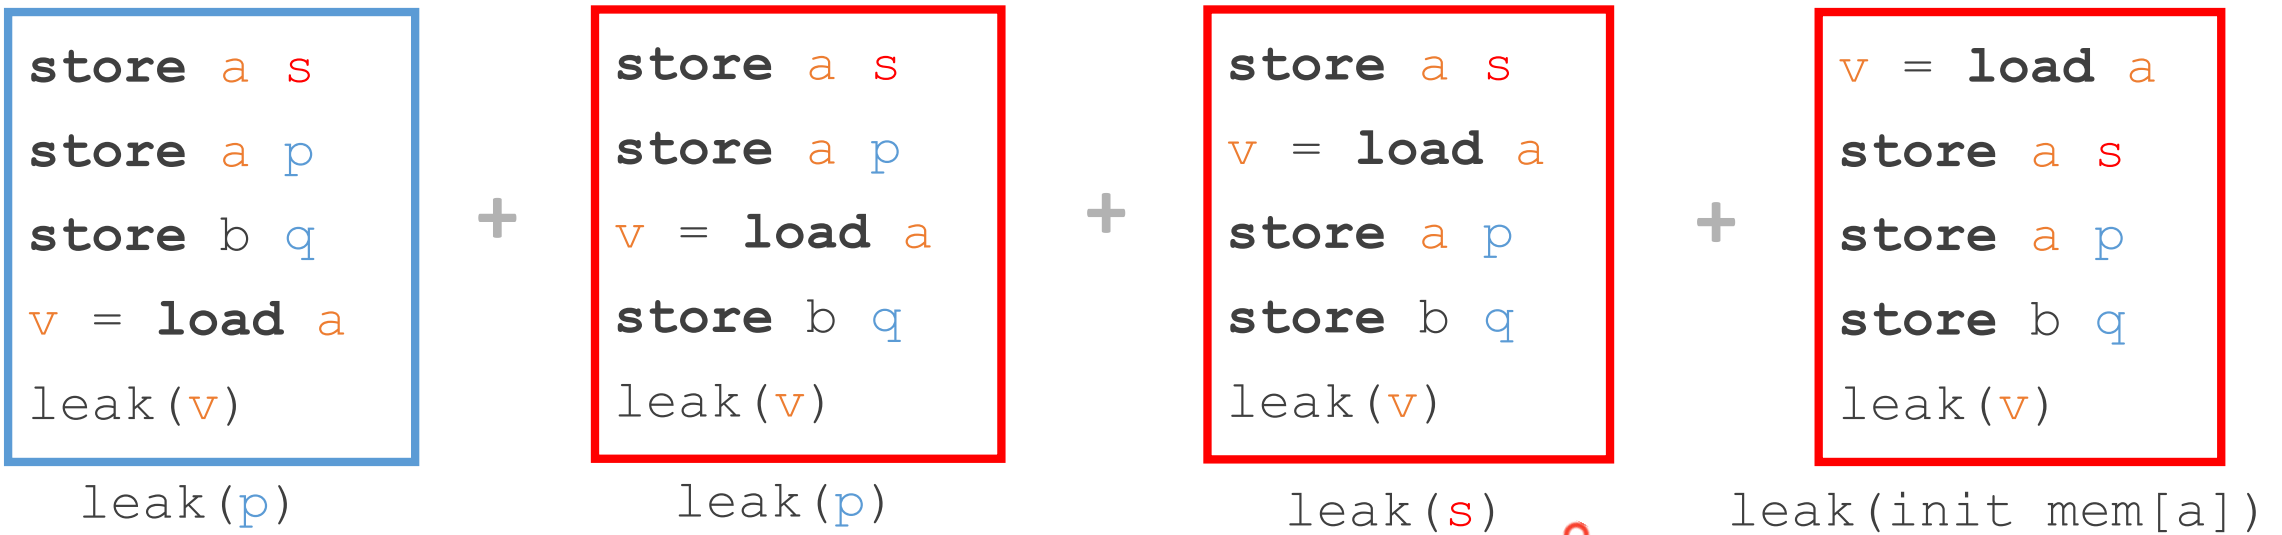


- where **s** is secret, **p** and **q** are public
- where **a**  $\neq$  **b**
- leak(v) encodes v to cache

# Spectre-STL

**Spectre-STL:** Loads can speculatively bypass prior stores

**Sequential execution** + **Transient Executions**




- where `s` is secret, `p` and `q` are public
- where `a`  $\neq$  `b`
- `leak(v)` encodes `v` to cache




# Constant-time verification in the Spectre era

## Not easy to write constant-time programs

- Sequence of **instructions** executed
  - First timing attacks by Paul Kocher, 1996
- **Memory** accesses
  - Cache attacks, 2005
- **Processors optimizations**
  - Spectre attacks, 2018



## Multiple failure points



We need efficient **automated verification tools** that take into account **speculation mechanisms in processors**

# Modelling speculative semantics

## Litmus tests (328 instructions):

- Sequential semantics  
→ **14 paths**
- Speculative semantics (Spectre-STL)  
→ **37M paths**



*Modelling all transient paths **explicitly** is intractable*

# No efficient verification tools for Spectre 😞

	Target	Spectre-PHT	Spectre-STL
<b>KLEESpectre [1]</b>	LLVM	😊	-
<b>SpecuSym [2]</b>	LLVM	😊	-
<b>FASS [3]</b>	Binary	😞	-
<b>Spectector [4]</b>	Binary	😞	-
<b>Pitchfork [5]</b>	Binary	😐	😞

## Legend

- 😊 Good perms. on crypto
- 😐 Good on small programs  
Limited perms. On crypto
- 😞 Limited to small programs

LLVM analysis might  
miss violations 😐

- [1] G. Wang, et al “KLEESpectre: Detecting Information Leakage through Speculative Cache Attacks via Symbolic Execution”, ACM Trans. Softw. Eng. Methodol., vol. 29, no. 3, 2020.
- [2] S. Guo, Y. Chen, P. Li, Y. Cheng, H. Wang, M. Wu, and Z. Zuo, “SpecuSym: Speculative Symbolic Execution for Cache Timing Leak Detection”, in ICSE 2020 Technical Papers, 2020.
- [3] K. Cheang, C. Rasmussen, S. A. Seshia, and P. Subramanyan, “A Formal Approach to Secure Speculation”, in CSF, 2019.
- [4] M. Guarnieri, B. Köpf, J. F. Morales, J. Reineke, and A. Sánchez, “Spectector: Principled Detection of Speculative Information Flows”, in S&P, 2020
- [5] S. Cauligi, C. Disselkoe, K. von Gleissenthall, D. M. Tullsen, D. Stefan, T. Rezk, and G. Barthe, “Constant-Time Foundations for the New Spectre Era”, in PLDI, 2020.

# No efficient verification tools for Spectre ?

	Target	Spectre-PHT	Spectre-STL
<b>KLEESpectre [1]</b>	LLVM	😊	-
<b>SpecuSym [2]</b>	LLVM	😊	-
<b>FASS [3]</b>	Binary	😞	-
<b>Spectector [4]</b>	Binary	😞	-
<b>Pitchfork [5]</b>	Binary	😐	😞
<b>Binsec/Haunted</b>	<b>Binary</b>	😊	😐

## Legend

- 😊 Good perms. on crypto
- 😐 Good on small programs  
Limited perms. On crypto
- 😞 Limited to small programs


LLVM analysis might  
miss violations 😐

- [1] G. Wang, et al “KLEESpectre: Detecting Information Leakage through Speculative Cache Attacks via Symbolic Execution”, ACM Trans. Softw. Eng. Methodol., vol. 29, no. 3, 2020.
- [2] S. Guo, Y. Chen, P. Li, Y. Cheng, H. Wang, M. Wu, and Z. Zuo, “SpecuSym: Speculative Symbolic Execution for Cache Timing Leak Detection”, in ICSE 2020 Technical Papers, 2020.
- [3] K. Cheang, C. Rasmussen, S. A. Seshia, and P. Subramanyan, “A Formal Approach to Secure Speculation”, in CSF, 2019.
- [4] M. Guarnieri, B. Köpf, J. F. Morales, J. Reineke, and A. Sánchez, “Spectector: Principled Detection of Speculative Information Flows”, in S&P, 2020
- [5] S. Cauligi, C. Disselkoe, K. von Gleissenthall, D. M. Tullsen, D. Stefan, T. Rezk, and G. Barthe, “Constant-Time Foundations for the New Spectre Era”, in PLDI, 2020.

# Haunted ReISE

# Explicit ReSE for Spectre PHT


**Symbolic execution** with sequential semantics





# Explicit ReSE for Spectre PHT

**Spectre-PHT.** Conditional branches can be executed speculatively



## Explicit ReSE.

Fork execution into 4 at conditionals:

- 2 sequential branches
- 2 transient branches

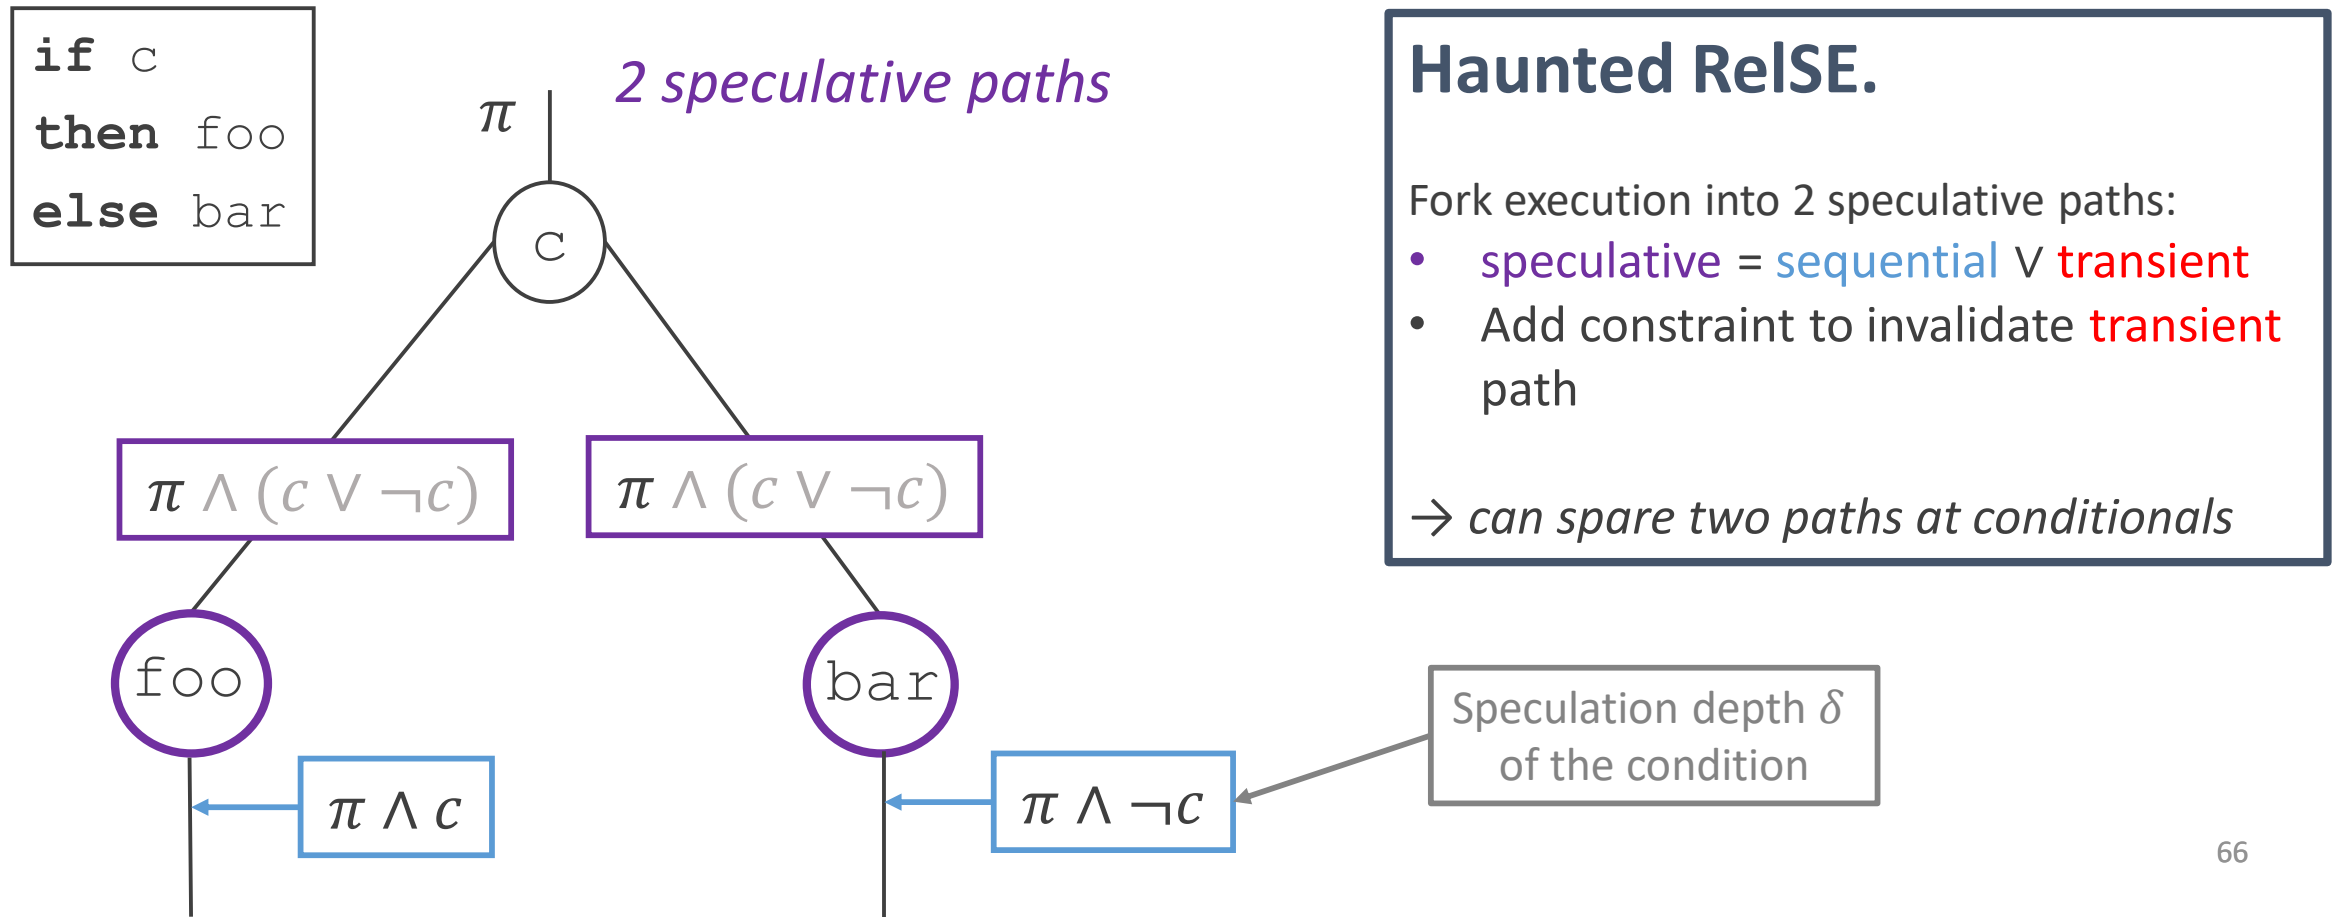
On sequential and transient branches:

- Verify no secret can leak.

(e.g. KLEESpectre)

# Haunted ReISE for Spectre PHT

**Spectre-PHT.** Conditional branches can be executed speculatively




# Explicit ReSE for Spectre-STL

```
store a s  
store a p  
store b q  
v = load a
```

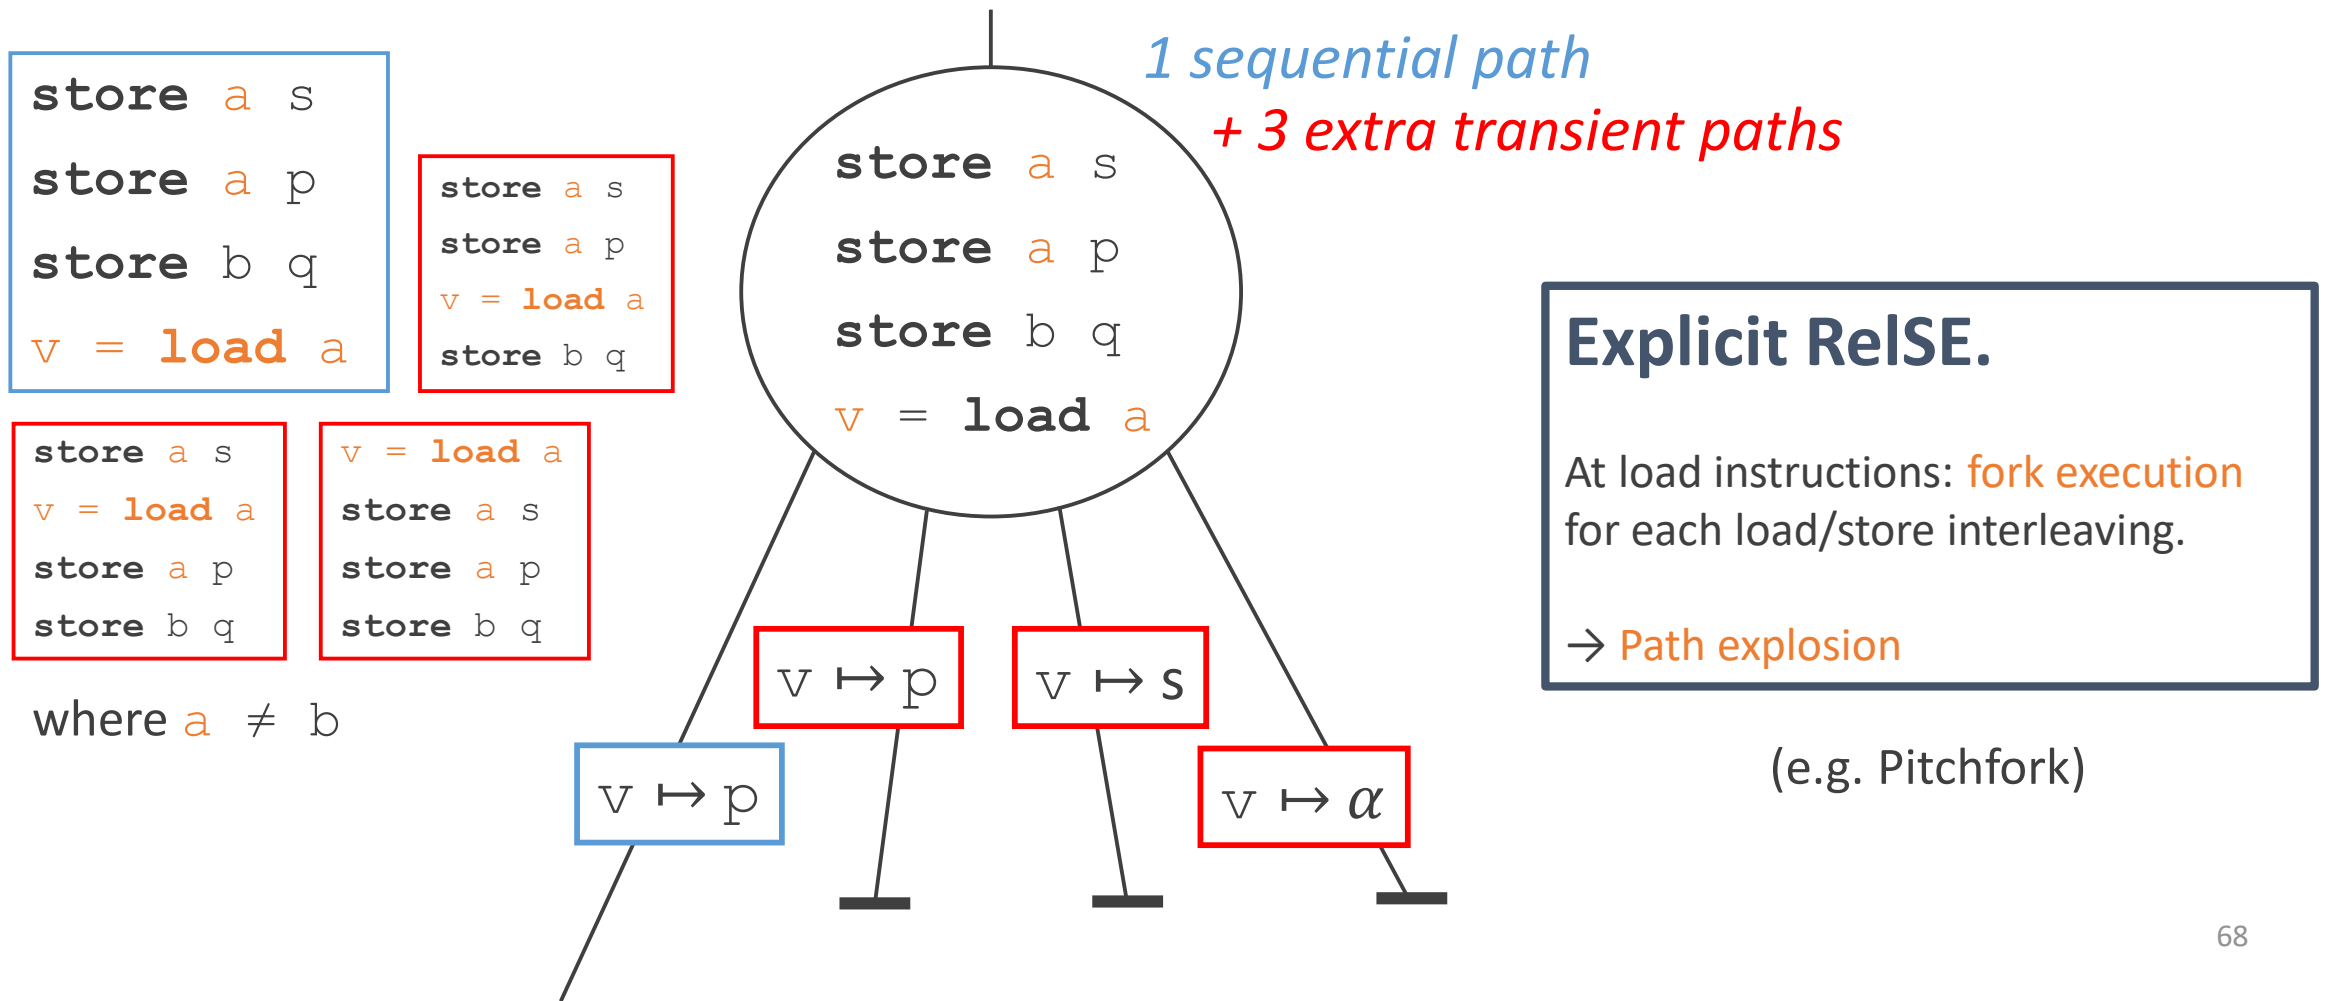
where  $a \neq b$

```
v ↦ p
```



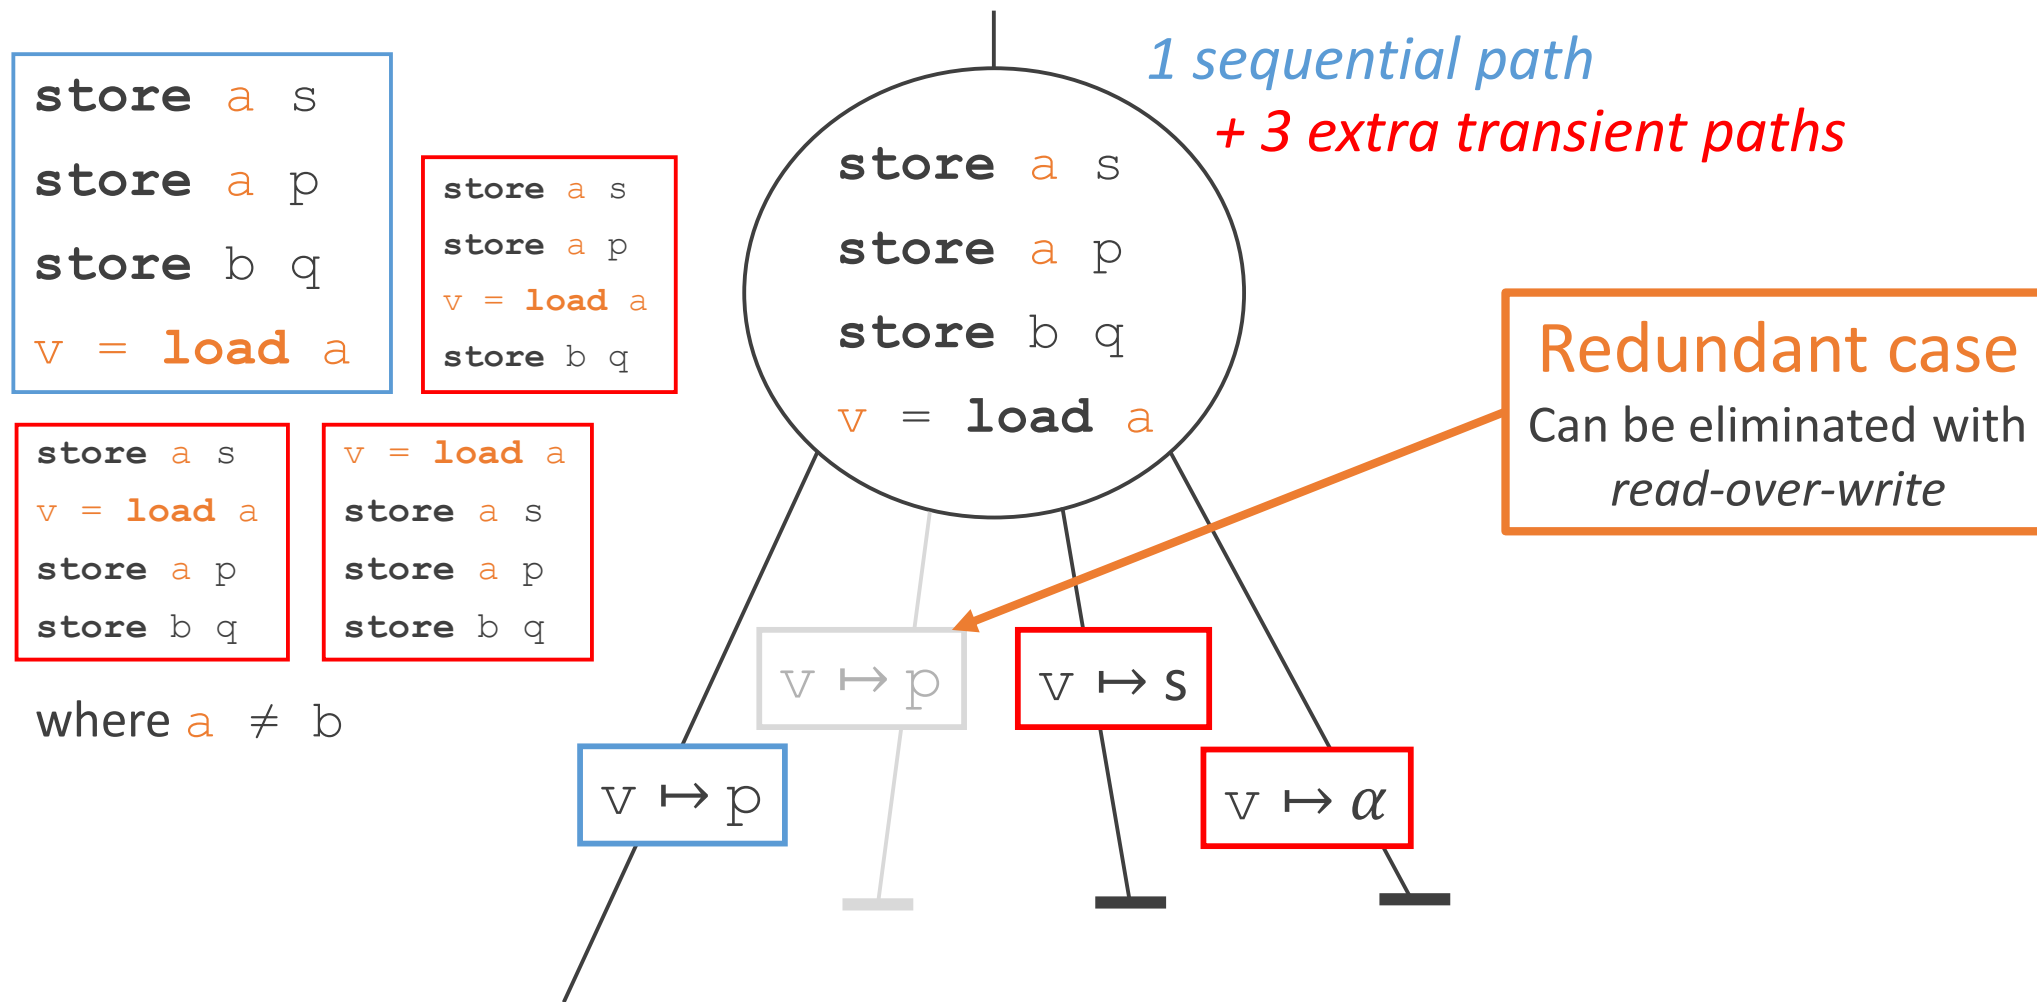
# Explicit RelSE for Spectre-STL

**Spectre-STL.** Loads can speculatively bypass prior stores



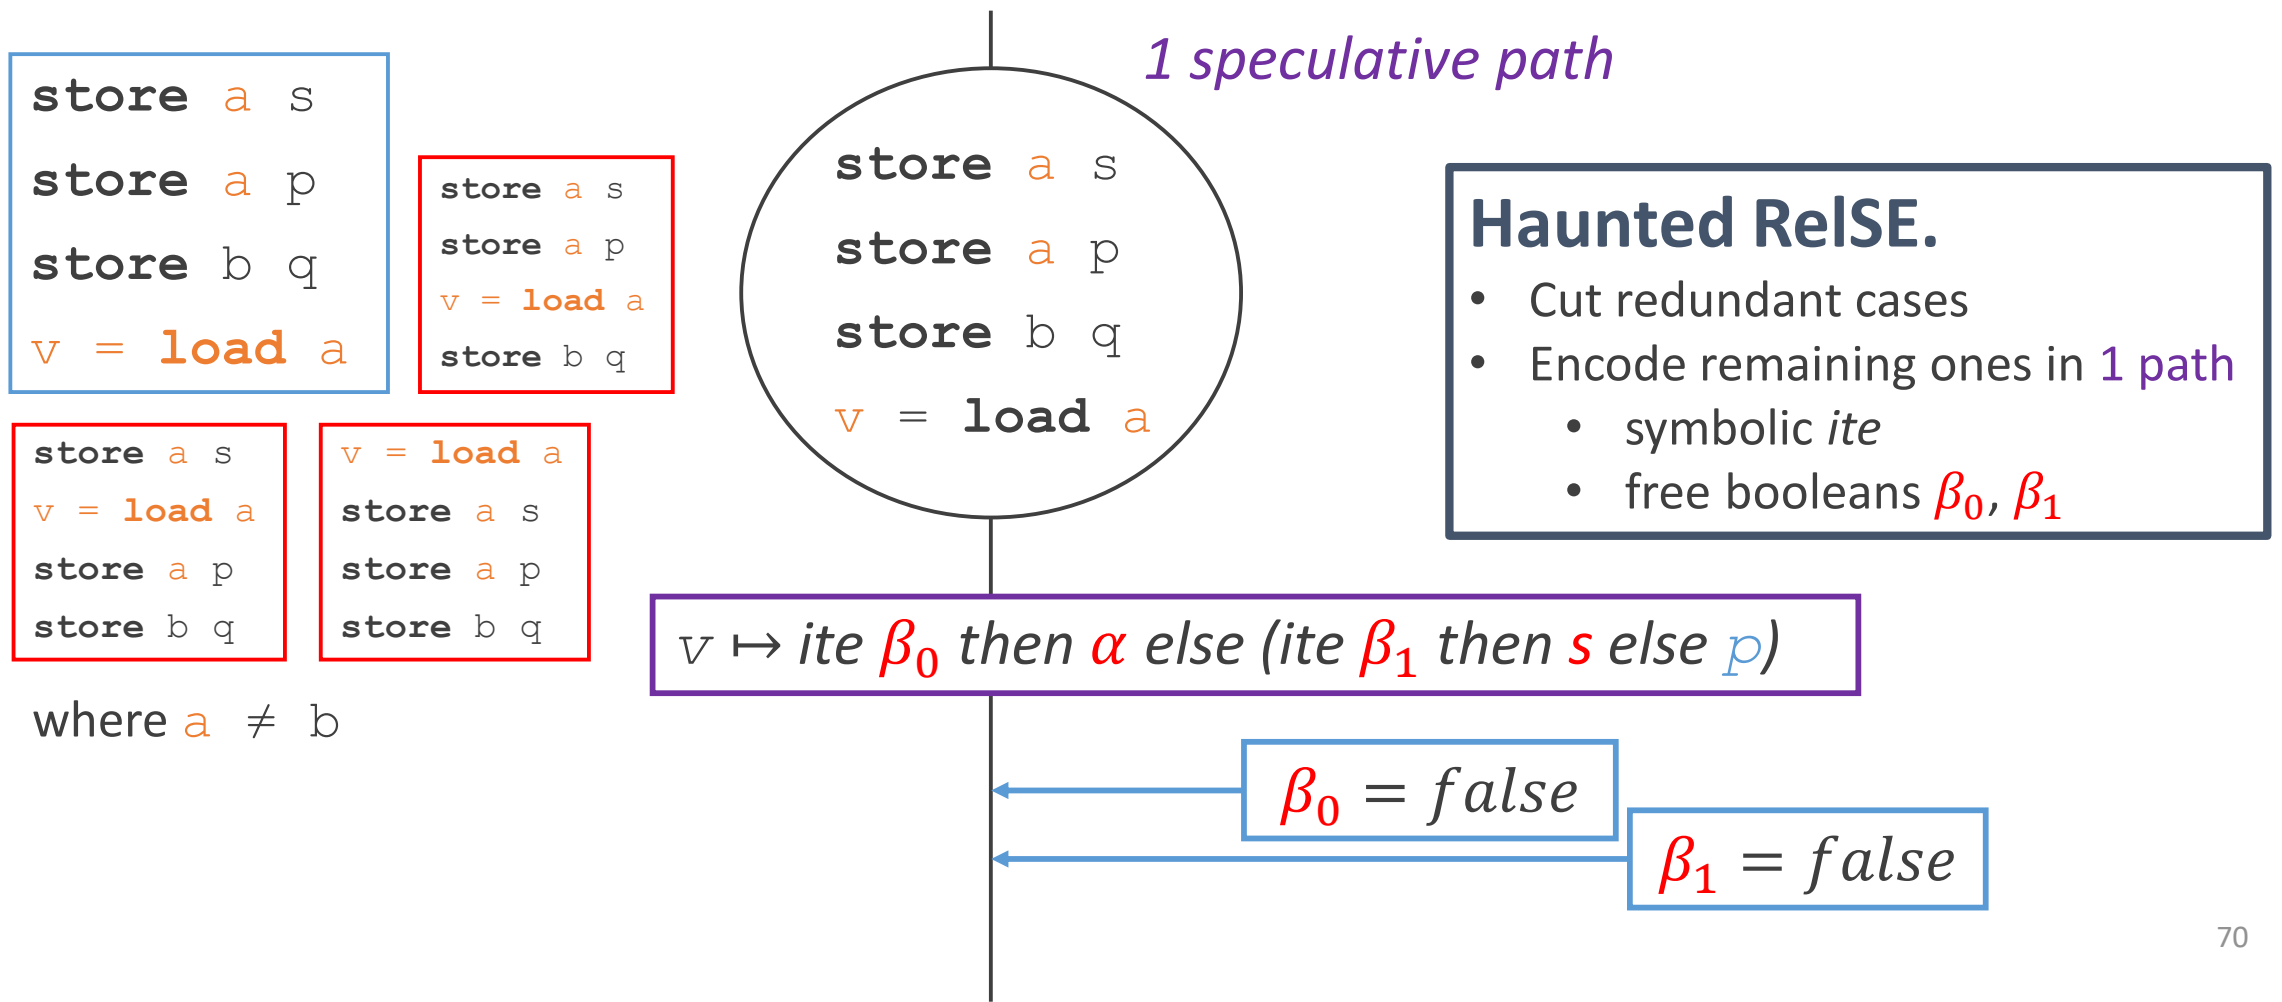
# Haunted ReSE for Spectre-STL

**Spectre-STL.** Loads can speculatively bypass prior stores



# Haunted RelSE for Spectre-STL

**Spectre-STL.** Loads can speculatively bypass prior stores



# Experimental evaluation



<https://github.com/binsec/haunted>

# Experimental evaluation

## Benchmark

**Litmus tests:** Spectre-PHT = Paul Kocher standard, Spectre-STL = **new** set of litmus tests

**Cryptographic primitives:** tea, donna, Libsodium secretbox, OpenSSL ssl3-digest-record & mee-cdc-decrypt

### Effective on real code?

→ Spectre-PHT 😊 & Spectre-STL 😞

### Haunted RelSE vs. Explicit RelSE?

→ Spectre-PHT: ≈ or ↗ & Spectre-STL: *always* ↗

### Comparison against KLEESpectre & Pitchfork

→ Spectre-PHT: ≈ or ↗ & Spectre-STL: *always* ↗

PHT	STL
<b>Litmus:</b>	
Paths: 1546 → 370	Paths: 93M → 42
Time: 3h → 15s	Coverage: 2k → 17k
<b>Libsodium + OpenSSL:</b>	Timeouts: 15 → 8
Coverage: 2273 → 8634	Bugs: 22 → 148
<b>Total:</b>	
Timeouts: 5 → 1	



# Weakness of index-masking countermeasure

+ Position independent code

# Weakness of Spectre-PHT countermeasure

Program vulnerable to Spectre-PHT

```
if (idx < size) { // size = 256  
  
    v = tab[idx]  
    leak(v)  
  
}
```

# Weakness of Spectre-PHT countermeasure

## Index masking countermeasure

```
if (idx < size) { // size = 256
    idx = idx & (0xff)
    v = tab[idx]
    leak(v)
}
```

# Weakness of Spectre-PHT countermeasure

## Index masking countermeasure

```
if (idx < size) { // size = 256
    idx = idx & (0xff)
    v = tab[idx]
    leak(v)
}
```



## Compiled version with gcc -O0 -m32

```
store  @idx (idx & 0xff)
eax = load @idx
al = [@tab + eax]
leak (al)
```

- Store + load masked index
- Store may be **bypassed** with **Spectre-STL !**

# Weakness of Spectre-PHT countermeasure

## Index masking countermeasure

```
if (idx < size) { // size = 256
    idx = idx & (0xff)
    v = tab[idx]
    leak(v)
}
```



## Compiled version with gcc -O0 -m32

```
store @idx (idx & 0xff)
eax = load @idx
al = [@tab + eax]
leak (al)
```

- Store + load masked index
- Store may be **bypassed** with **Spectre-STL !**

## Verified mitigations:

- Enable optimizations (depends on compiler choices)
- Explicitly put masked index in a register

```
register uint32_t ridx asm ("eax");
```