


Spectector: Principled detection of speculative information flows


Jan Reineke @  UNIVERSITÄT
DES
SAARLANDES

Joint work with

Marco Guarnieri, Jose Morales, Andres Sanchez @ IMDEA Software, Madrid
Boris Köpf @ Microsoft Research, Cambridge, UK

Supported by Intel Strategic Research Alliance (ISRA)

“Information Flow Tracking across the Hardware-Software Boundary”



SPECTRE

Exploits **speculative execution** to leak sensitive information

Almost all modern processors are affected

Countermeasures

Long Term: Co-Design of Software and Hardware countermeasures

Short and Mid Term: Software countermeasures

In particular: Compiler-level countermeasures

- ✓ *Example: insert “fences” to selectively terminate speculative execution*
- ✓ *Implemented in major compilers (Microsoft Visual C++, Intel ICC, Clang)*

PROBLEM SOLVED ?

Compiler-level countermeasures

Spectre Mitigations in Microsoft's C/C++ Compiler

Paul Kocher
February 13, 2018

“The countermeasure [...] is conceptually straightforward
but **challenging in practice**”

“compiler [...] produces **unsafe code** when the
static analyzer is unable to determine whether
a code pattern will be exploitable”

“there is **no guarantee** that all possible instances of
[Spectre] will be instrumented”

Bottom line: No guarantees!

Goals

1. Introduce **semantic notion of security**
against **speculative execution attacks**

2. Static analysis to **detect vulnerability**
or to **prove security**

Outline

1. Speculative execution attacks
2. Speculative non-interference
3. Spectector: Detecting speculative leaks
4. Challenges

1. Speculative execution attacks


Background: Speculative execution

- Predict instructions' outcomes and speculatively continue execution
- Rollback changes if speculation was wrong


Only architectural (ISA, “logical”) state,
not microarchitectural state

Background: Branch prediction


Size of array **A**

```
if ( x < A_size )   
  y = B[A[x]]
```


Predictions based on
branch history &
program structure



Spectre V1




What is in **A**[128]?




```
void f(int x)
  if (x < A_size)
    y = B[A[x]]
```

1a) Training




Cache state

Spectre V1



What is in **A**[128]?



```
void f(int x)  
  if (x < A_size)  
    y = B[A[x]]
```

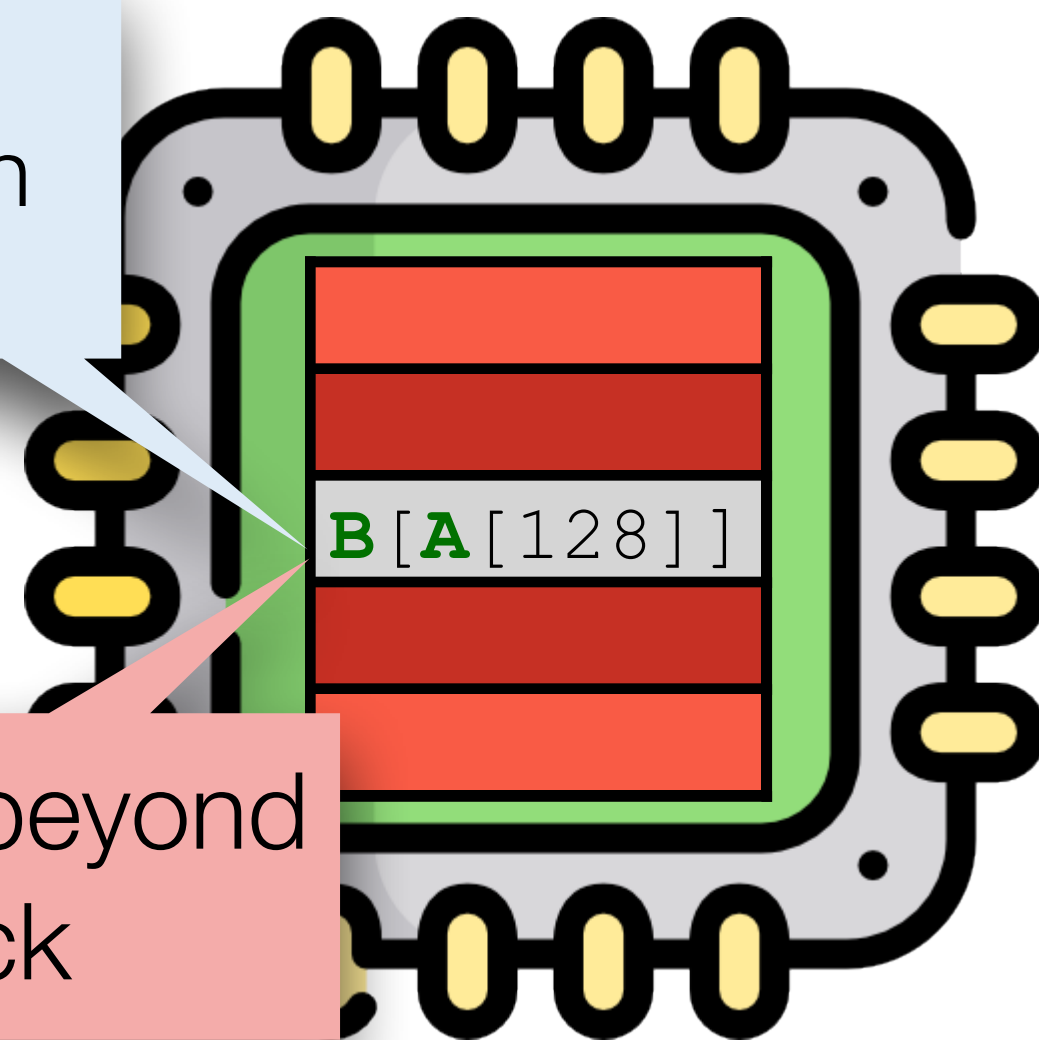
1a) Training  `f(0); f(1); f(2); ...`

1b) Prepare cache

2) Run `f(128)`

3) Extract from cache

Address depends on **A**[128]



Persistent beyond rollback

Cache state

2. Speculative non-interference

Generalizing the Spectre V1 example

1a) Training  `f(0); f(1); f(2); ...`

1b) Prepare cache

2) Run `f(128)`

3) Extract from cache

Attacker

Victim

Attacker

Generalizing the Spectre V1 example

1) Prepares microarchitectural state

} Attacker

Victim {

2) Leaks information into microarchitectural state

3) Extracts information from microarchitecture

} Attacker

Speculative non-interference

Extended
with policies

Program \mathbf{P} is **speculatively non-interferent** if

Informally:

Leakage of \mathbf{P} in
non-speculative
execution

?

=

Leakage of \mathbf{P} in
speculative
execution

For all program states \mathbf{s} and \mathbf{s}' :

More formally:

$$\begin{aligned} & \mathbf{P}_{\text{non-spec}}(\mathbf{s}) = \mathbf{P}_{\text{non-spec}}(\mathbf{s}') \\ \Rightarrow & \mathbf{P}_{\text{spec}}(\mathbf{s}) = \mathbf{P}_{\text{spec}}(\mathbf{s}') \end{aligned}$$

How to capture leakage into microarchitectural state?

Non-speculative semantics


Speculative semantics

+

Attacker/Observer model

μAssembly + Non-speculative semantics

```
if (x < A_size)  
  y = B[A[x]]
```



```
rax <- A_size
```


```
rcx <- x
```

```
jmp rcx ≥ rax, END
```

```
L1: load rax, A + rcx
```

```
load rax, B + rax
```

```
END:
```



Non-speculative semantics: Inference Rules

Expression evaluation

$$\llbracket n \rrbracket(a) = n \quad \llbracket x \rrbracket(a) = a(x) \quad \llbracket \ominus e \rrbracket(a) = \ominus \llbracket e \rrbracket(a) \quad \llbracket e_1 \otimes e_2 \rrbracket(a) = \llbracket e_1 \rrbracket(a) \otimes \llbracket e_2 \rrbracket(a)$$

Instruction evaluation

SKIP

$$\frac{p(a(\mathbf{pc})) = \mathbf{skip}}{\langle m, a \rangle \rightarrow \langle m, a[\mathbf{pc} \mapsto a(\mathbf{pc}) + 1] \rangle}$$

BARRIER

$$\frac{p(a(\mathbf{pc})) = \mathbf{spbarr}}{\langle m, a \rangle \rightarrow \langle m, a[\mathbf{pc} \mapsto a(\mathbf{pc}) + 1] \rangle}$$

ASSIGN

$$\frac{p(a(\mathbf{pc})) = x \leftarrow e \quad x \neq \mathbf{pc}}{\langle m, a \rangle \rightarrow \langle m, a[\mathbf{pc} \mapsto a(\mathbf{pc}) + 1, x \mapsto \llbracket e \rrbracket(a)] \rangle}$$

CONDITIONALUPDATE-SAT

$$\frac{p(a(\mathbf{pc})) = x \xleftarrow{e'} e \quad \llbracket e' \rrbracket(a) = 0 \quad x \neq \mathbf{pc}}{\langle m, a \rangle \rightarrow \langle m, a[\mathbf{pc} \mapsto a(\mathbf{pc}) + 1, x \mapsto \llbracket e \rrbracket(a)] \rangle}$$

CONDITIONALUPDATE-UNSAT

$$\frac{p(a(\mathbf{pc})) = x \xleftarrow{e'} e \quad \llbracket e' \rrbracket(a) \neq 0 \quad x \neq \mathbf{pc}}{\langle m, a \rangle \rightarrow \langle m, a[\mathbf{pc} \mapsto a(\mathbf{pc}) + 1] \rangle}$$

TERMINATE

$$\frac{p(a(\mathbf{pc})) = \perp}{\langle m, a \rangle \rightarrow \langle m, a[\mathbf{pc} \mapsto \perp] \rangle}$$

LOAD

$$\frac{p(a(\mathbf{pc})) = \mathbf{load} \ x, e \quad x \neq \mathbf{pc} \quad n = \llbracket e \rrbracket(a)}{\langle m, a \rangle \xrightarrow{\mathbf{load} \ n} \langle m, a[\mathbf{pc} \mapsto a(\mathbf{pc}) + 1, x \mapsto m(n)] \rangle}$$

STORE

$$\frac{p(a(\mathbf{pc})) = \mathbf{store} \ x, e \quad n = \llbracket e \rrbracket(a)}{\langle m, a \rangle \xrightarrow{\mathbf{store} \ n} \langle m[n \mapsto a(x)], a[\mathbf{pc} \mapsto a(\mathbf{pc}) + 1] \rangle}$$

BEQZ-SAT

$$\frac{p(a(\mathbf{pc})) = \mathbf{beqz} \ x, \ell \quad a(x) = 0}{\langle m, a \rangle \xrightarrow{\mathbf{pc} \ \ell} \langle m, a[\mathbf{pc} \mapsto \ell] \rangle}$$

BEQZ-UNSAT

$$\frac{p(a(\mathbf{pc})) = \mathbf{beqz} \ x, \ell \quad a(x) \neq 0}{\langle m, a \rangle \xrightarrow{\mathbf{pc} \ a(\mathbf{pc})+1} \langle m, a[\mathbf{pc} \mapsto a(\mathbf{pc}) + 1] \rangle}$$

JMP

$$\frac{p(a(\mathbf{pc})) = \mathbf{jmp} \ e \quad \ell = \llbracket e \rrbracket(a)}{\langle m, a \rangle \xrightarrow{\mathbf{pc} \ \ell} \langle m, a[\mathbf{pc} \mapsto \ell] \rangle}$$

Speculative semantics

```
rax <- A_size
```


```
rcx <- x
```

```
jmp rcx ≥ rax, END
```

```
L1: load rax, A + rcx
```

```
load rax, B + rax
```

```
END:
```



Starts **speculative transactions** upon branches

Committed upon correct speculation

Rolled back upon misspeculation

Prediction Oracle  determines branch prediction + length of speculative window

Observer model: Leakage into μ architectural state

```
rax <- A_size
```

```
rcx <- x
```



```
jmp rcx  $\geq$  rax, END
```

```
L1: load rax, A + rcx
```

```
load rax, B + rax
```

```
END:
```

```
load B+A [x]
```



Attacker can observe:

- locations of **memory accesses**
- **branch/jump** targets
- **start/end** speculative execution


Inspired by “constant-time” programming requirements

No need for detailed model of memory hierarchy:

- **possibly pessimistic**
- **more robust**

Reasoning about arbitrary prediction oracles

Speculative semantics
+
Prediction oracle



Always-mispredict
speculative semantics

Always-mispredict speculative semantics

```
rax ←- A_size
```


```
rcx ←- x
```

```
jmp rcx ≥ rax, END
```

```
L1: load rax, A + rcx
```

```
load rax, B + rax
```

```
END:
```



Always mispredict branch instructions' outcomes

Fixed speculative window

Rollback of every transaction

Always-mispredict speculative semantics: Inference Rules

SE-NOBRANCH

$$\begin{array}{c}
 p(\sigma(\mathbf{pc})) \neq \mathbf{beqz} \ x, \ell \quad \sigma \xrightarrow{\tau}_s \sigma' \quad \mathit{enabled}'(s) \\
 s' = \begin{cases} \mathit{decr}'(s) & \text{if } p(\sigma(\mathbf{pc})) \neq \mathbf{spbarr} \\ \mathit{zeroes}'(s) & \text{otherwise} \end{cases} \\
 \hline
 \langle \mathit{ctr}, \sigma, s \rangle \xRightarrow{\tau}_s \langle \mathit{ctr}, \sigma', s' \rangle
 \end{array}$$

SE-BRANCH-SYMB


$$\begin{array}{c}
 p(\sigma(\mathbf{pc})) = \mathbf{beqz} \ x, \ell'' \quad \mathit{enabled}'(s) \\
 \sigma \xrightarrow{\mathit{symPc}(se).\mathbf{pc} \ \ell'}_s \sigma' \quad \ell = \begin{cases} \sigma(\mathbf{pc}) + 1 & \text{if } \ell' \neq \sigma(\mathbf{pc}) + 1 \\ \ell'' & \text{if } \ell' = \sigma(\mathbf{pc}) + 1 \end{cases} \\
 s' = s \cdot \langle \sigma, \mathit{ctr}, \min(w, \mathit{wndw}(s) - 1), \ell \rangle \quad \mathit{id} = \mathit{ctr} \\
 \hline
 \langle \mathit{ctr}, \sigma, s \rangle \xRightarrow{\mathit{symPc}(se).\mathbf{start} \ \mathit{id}.\mathbf{pc} \ \ell}_s \langle \mathit{ctr} + 1, \sigma[\mathbf{pc} \mapsto \ell], s' \rangle
 \end{array}$$

SE-ROLLBACK

$$\begin{array}{c}
 \sigma' \xrightarrow{\tau}_s \sigma'' \\
 \hline
 \langle \mathit{ctr}, \sigma, s \cdot \langle \sigma', \mathit{id}, 0, \ell \rangle \rangle \xRightarrow{\mathit{rollback} \ \mathit{id}.\mathbf{pc} \ \sigma''(\mathbf{pc})}_s \langle \mathit{ctr}, \sigma'', s \rangle
 \end{array}$$

Always-mispredict leaks maximally

Speculative semantics
+
Prediction oracle



Always-mispredict
speculative semantics

For all program states s and s' :

$$P_{\text{spec}}(s) = P_{\text{spec}}(s')$$

$$\Leftrightarrow \forall O: P_{\text{spec},O}(s) = P_{\text{spec},O}(s')$$

Recap: Speculative non-interference

Program \mathbf{P} is **speculatively non-interferent** if

For all program states \mathbf{s} and \mathbf{s}' :

$$\mathbf{P}_{\text{non-spec}}(\mathbf{s}) = \mathbf{P}_{\text{non-spec}}(\mathbf{s}')$$

$$\Rightarrow \mathbf{P}_{\text{spec}}(\mathbf{s}) = \mathbf{P}_{\text{spec}}(\mathbf{s}')$$

Speculative non-interference: Example

```
rax <- A_size
```

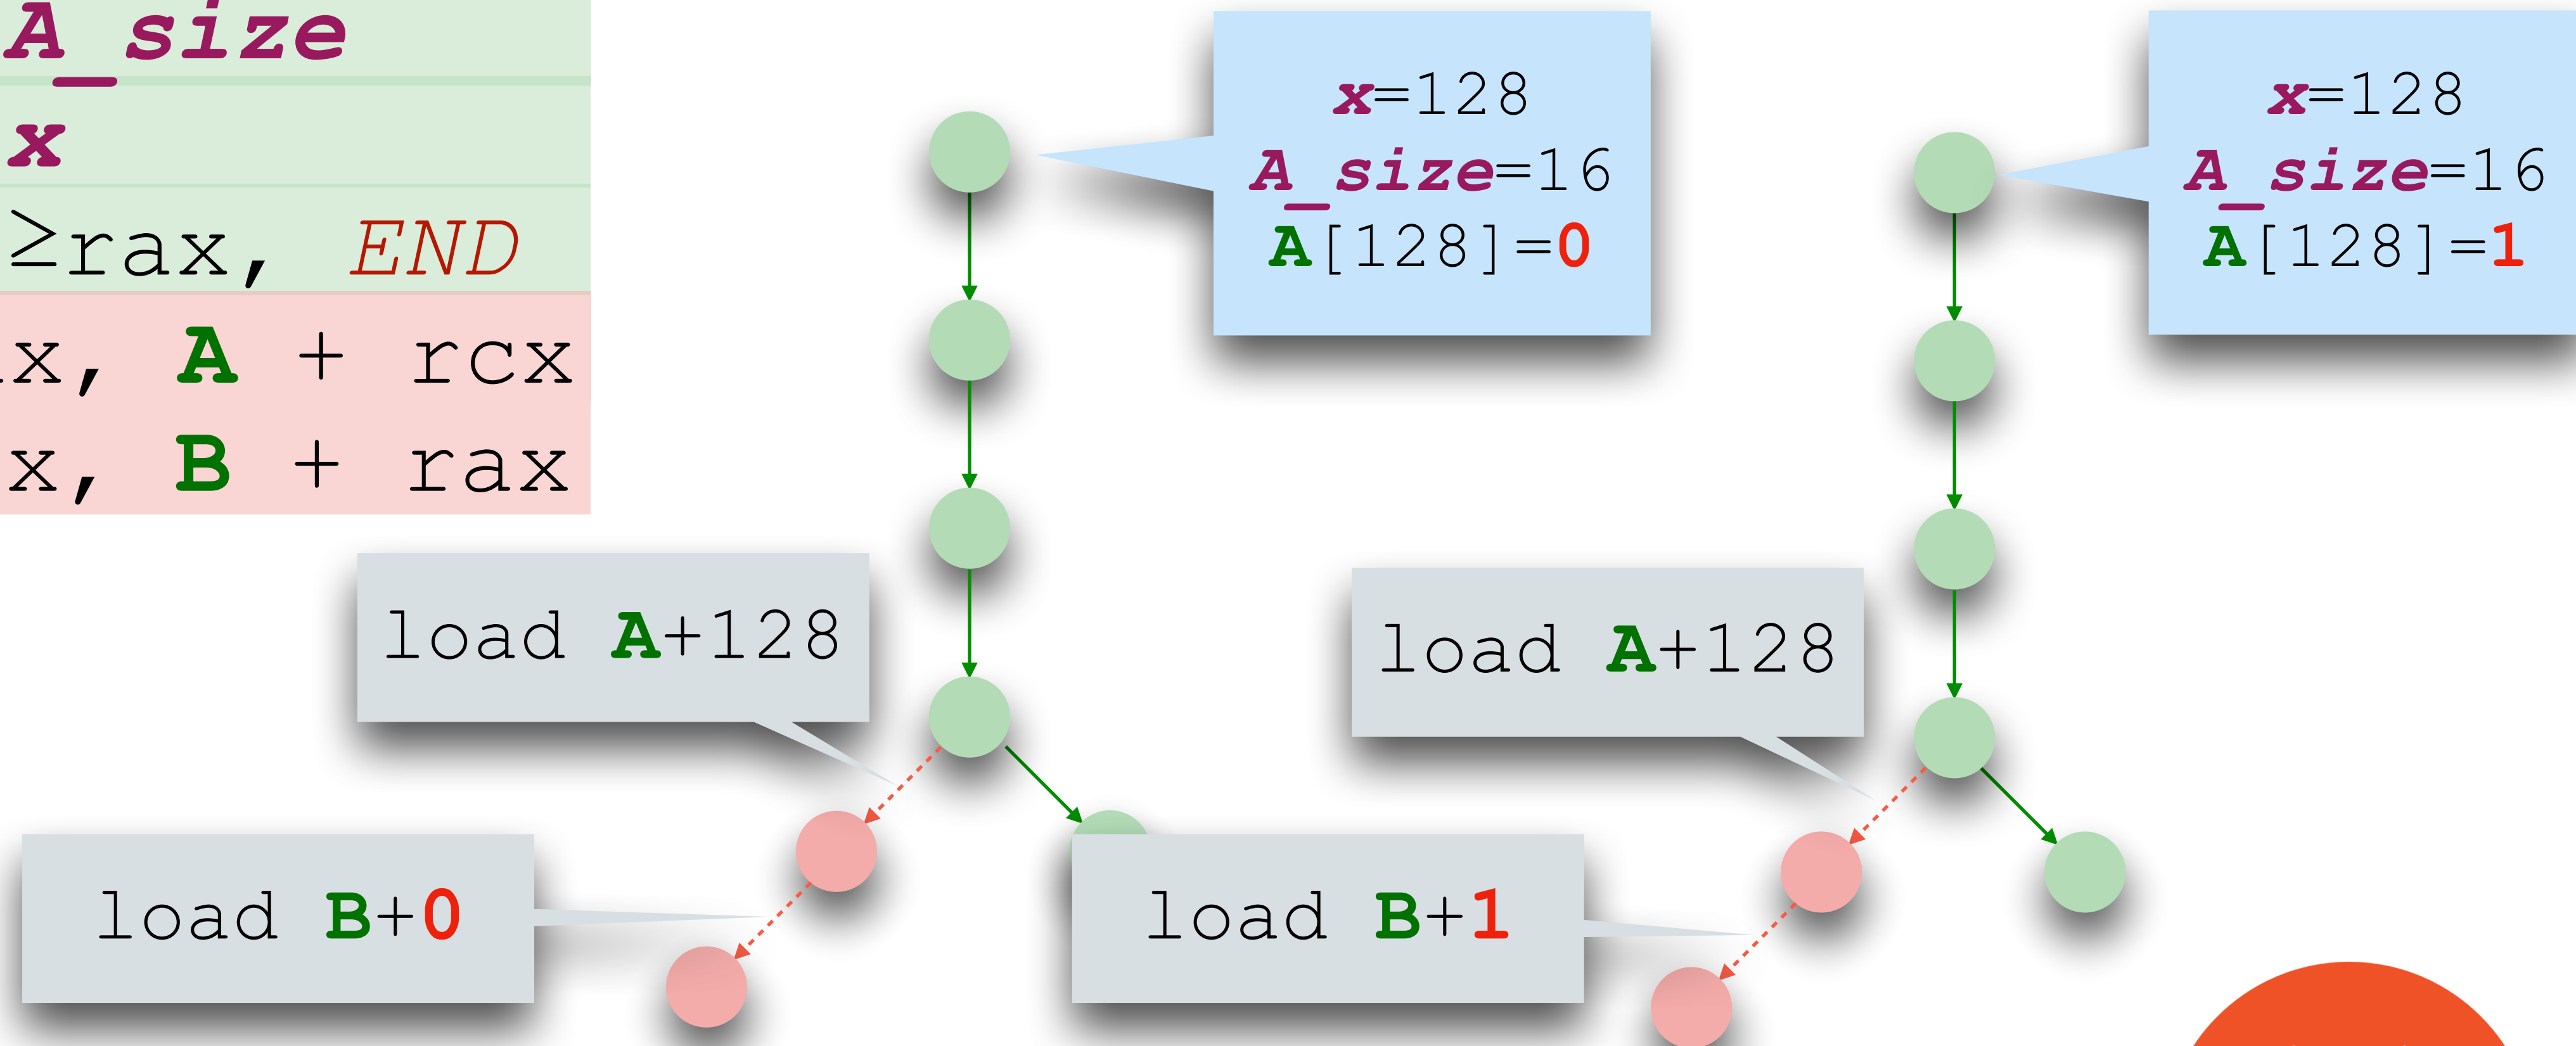
```
rcx <- x
```

```
jmp rcx ≥ rax, END
```

```
L1: load rax, A + rcx
```

```
load rax, B + rax
```

END:




3. Spectector: Detecting speculative leaks

Spectector: Detecting speculative leaks




```
rax <- A_size  
rcx <- x  
jmp rcx ≥ rax, END  
L1: load rax, A + rcx  
load rax, B + rax  
END:
```

Symbolic
execution



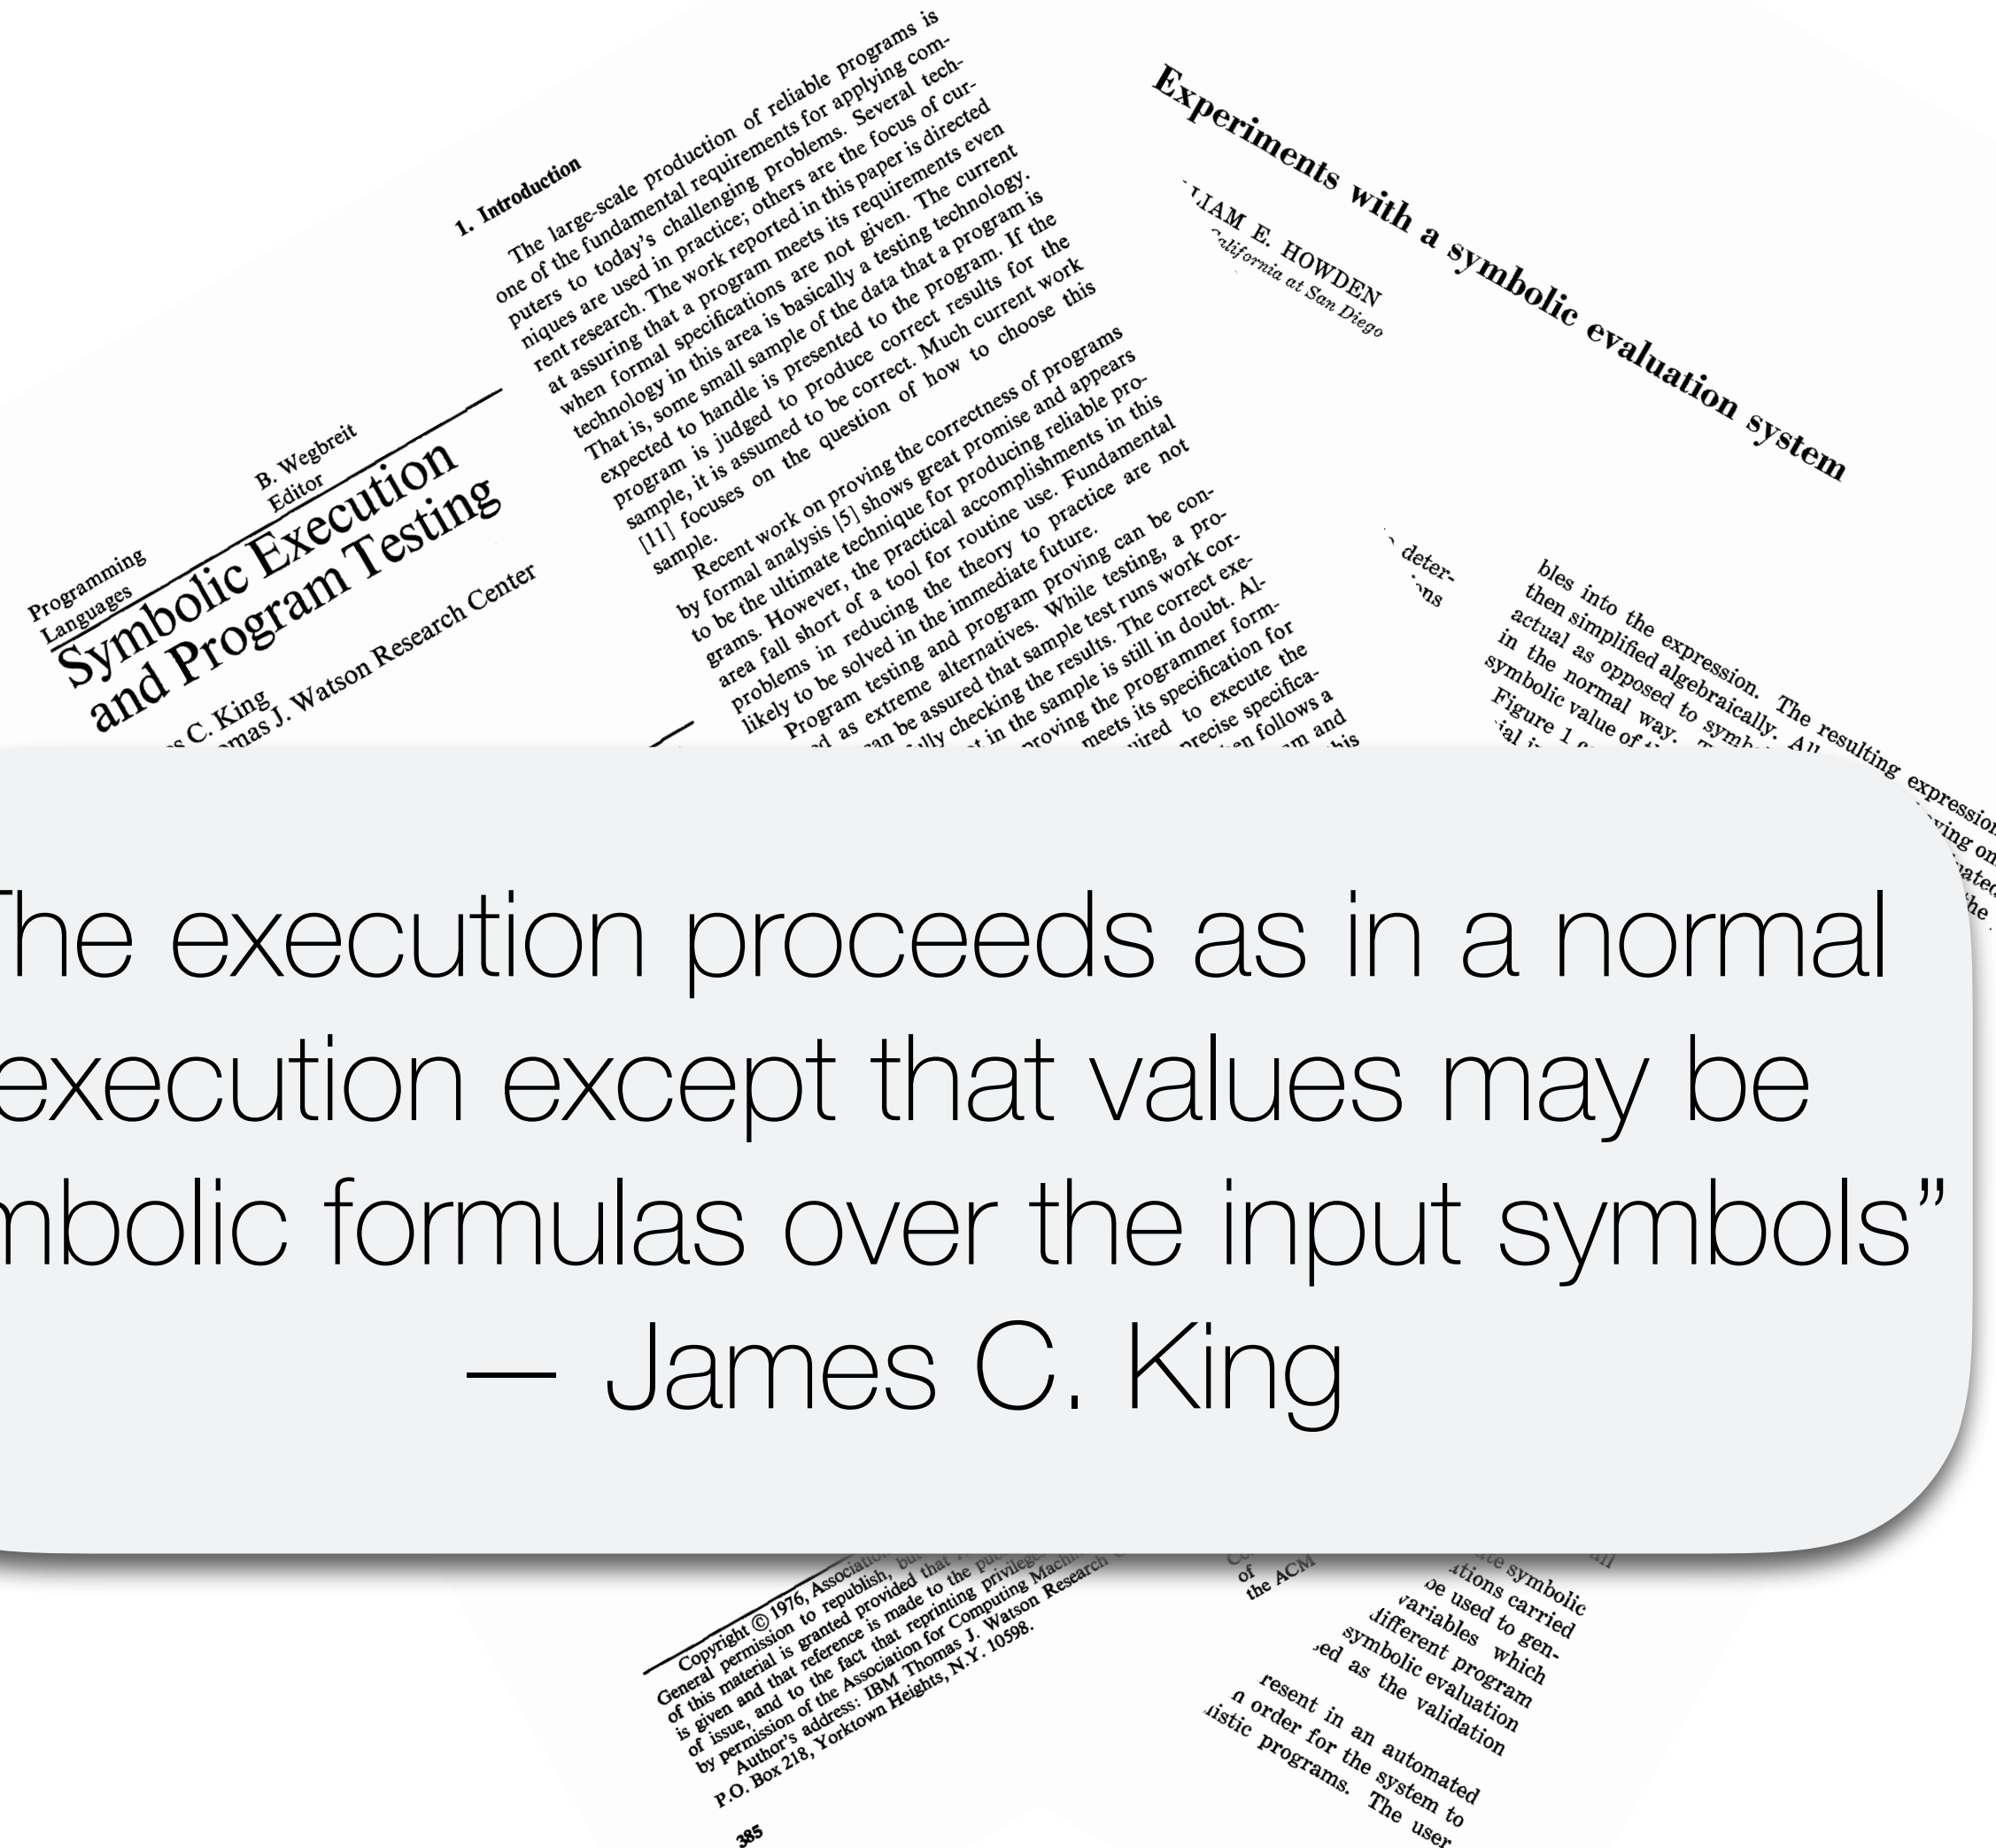
Detect leaks



Symbolic execution

- Program analysis technique
- Execute programs over symbolic values
 - Explore all paths, each with its own path constraint
 - Each path represents all possible executions satisfying the constraints
 - Branch and jump instructions: fork paths and update path constraint

“The execution proceeds as in a normal execution except that values may be symbolic formulas over the input symbols”
— James C. King



Detecting speculative leaks



```
rax <- A_s  
rcx <- x  
jmp rcx >= ra  
L1: load rax,  
load rax,
```

END:

```
For each  $\tau \in \text{sym-traces}(P)$   
  if MemLeak( $\tau$ ) then  
    return INSECURE  
  if CtrlLeak( $\tau$ ) then  
    return INSECURE  
  return SECURE
```







Memory leaks

Speculative memory accesses **must** depend only on

- Non-sensitive information (determined by policy), *or*
- be determined by non-speculative observations



$$\boxed{pathCnd(\tau)} \wedge \boxed{obsEqv(\tau |_{non-spec})} \wedge \boxed{\neg obsEqv(\tau |_{spec})}$$






Memory leaks

```

rax <- A_size
rcx <- x
jmp rcx >= rax, END
L1: load rax, A + rcx
    load rax, B + rax
END:

```



Policy
x, *A_size*, *A*, *B*
 are public

$\tau =$ start; pc L1; load A+x; load B+A[x]; rollback; pc END

$$pathCnd(\tau) \wedge obsEqv(\tau |_{non-spec}) \wedge \neg obsEqv(\tau |_{spec})$$

S_1	\models	$x_1 \geq A_size_1$	pc END	\parallel	$A_1 + x_1$	\nparallel	\vee	$B_1 + A_1[x_1]$	\nparallel
S_2	\models	$x_2 \geq A_size_2$	pc END	\parallel	$A_2 + x_2$	\nparallel	\vee	$B_2 + A_2[x_2]$	\nparallel

$$x_1 = x_2 \wedge A_size_1 = A_size_2 \wedge A_1 = A_2 \wedge B_1 = B_2$$

Always true!



Experimental results

No countermeasures

Automated insertions of fences

Speculative load hardening

15 Spectre variants from Paul Kocher

ICC

CLANG

UNP

FEN

UNP

FEN

SLH

-00 -00

2 -00 -02

Summary

Performance

- Leak (except #05)
 - Conf
 - Prog
 - Bu
 - Prog
- Programs ~20-200 lines of assembly code from Kocher
 - Analysis terminates in less than 30 sec
 - Except for example #05 (< 2 min)

01
02
03
04
05
06
07
08
09
10
11
12
13
14
15

if

y

4. Challenges

Scalable analysis

Goal:

Analysis of large, security-critical applications:

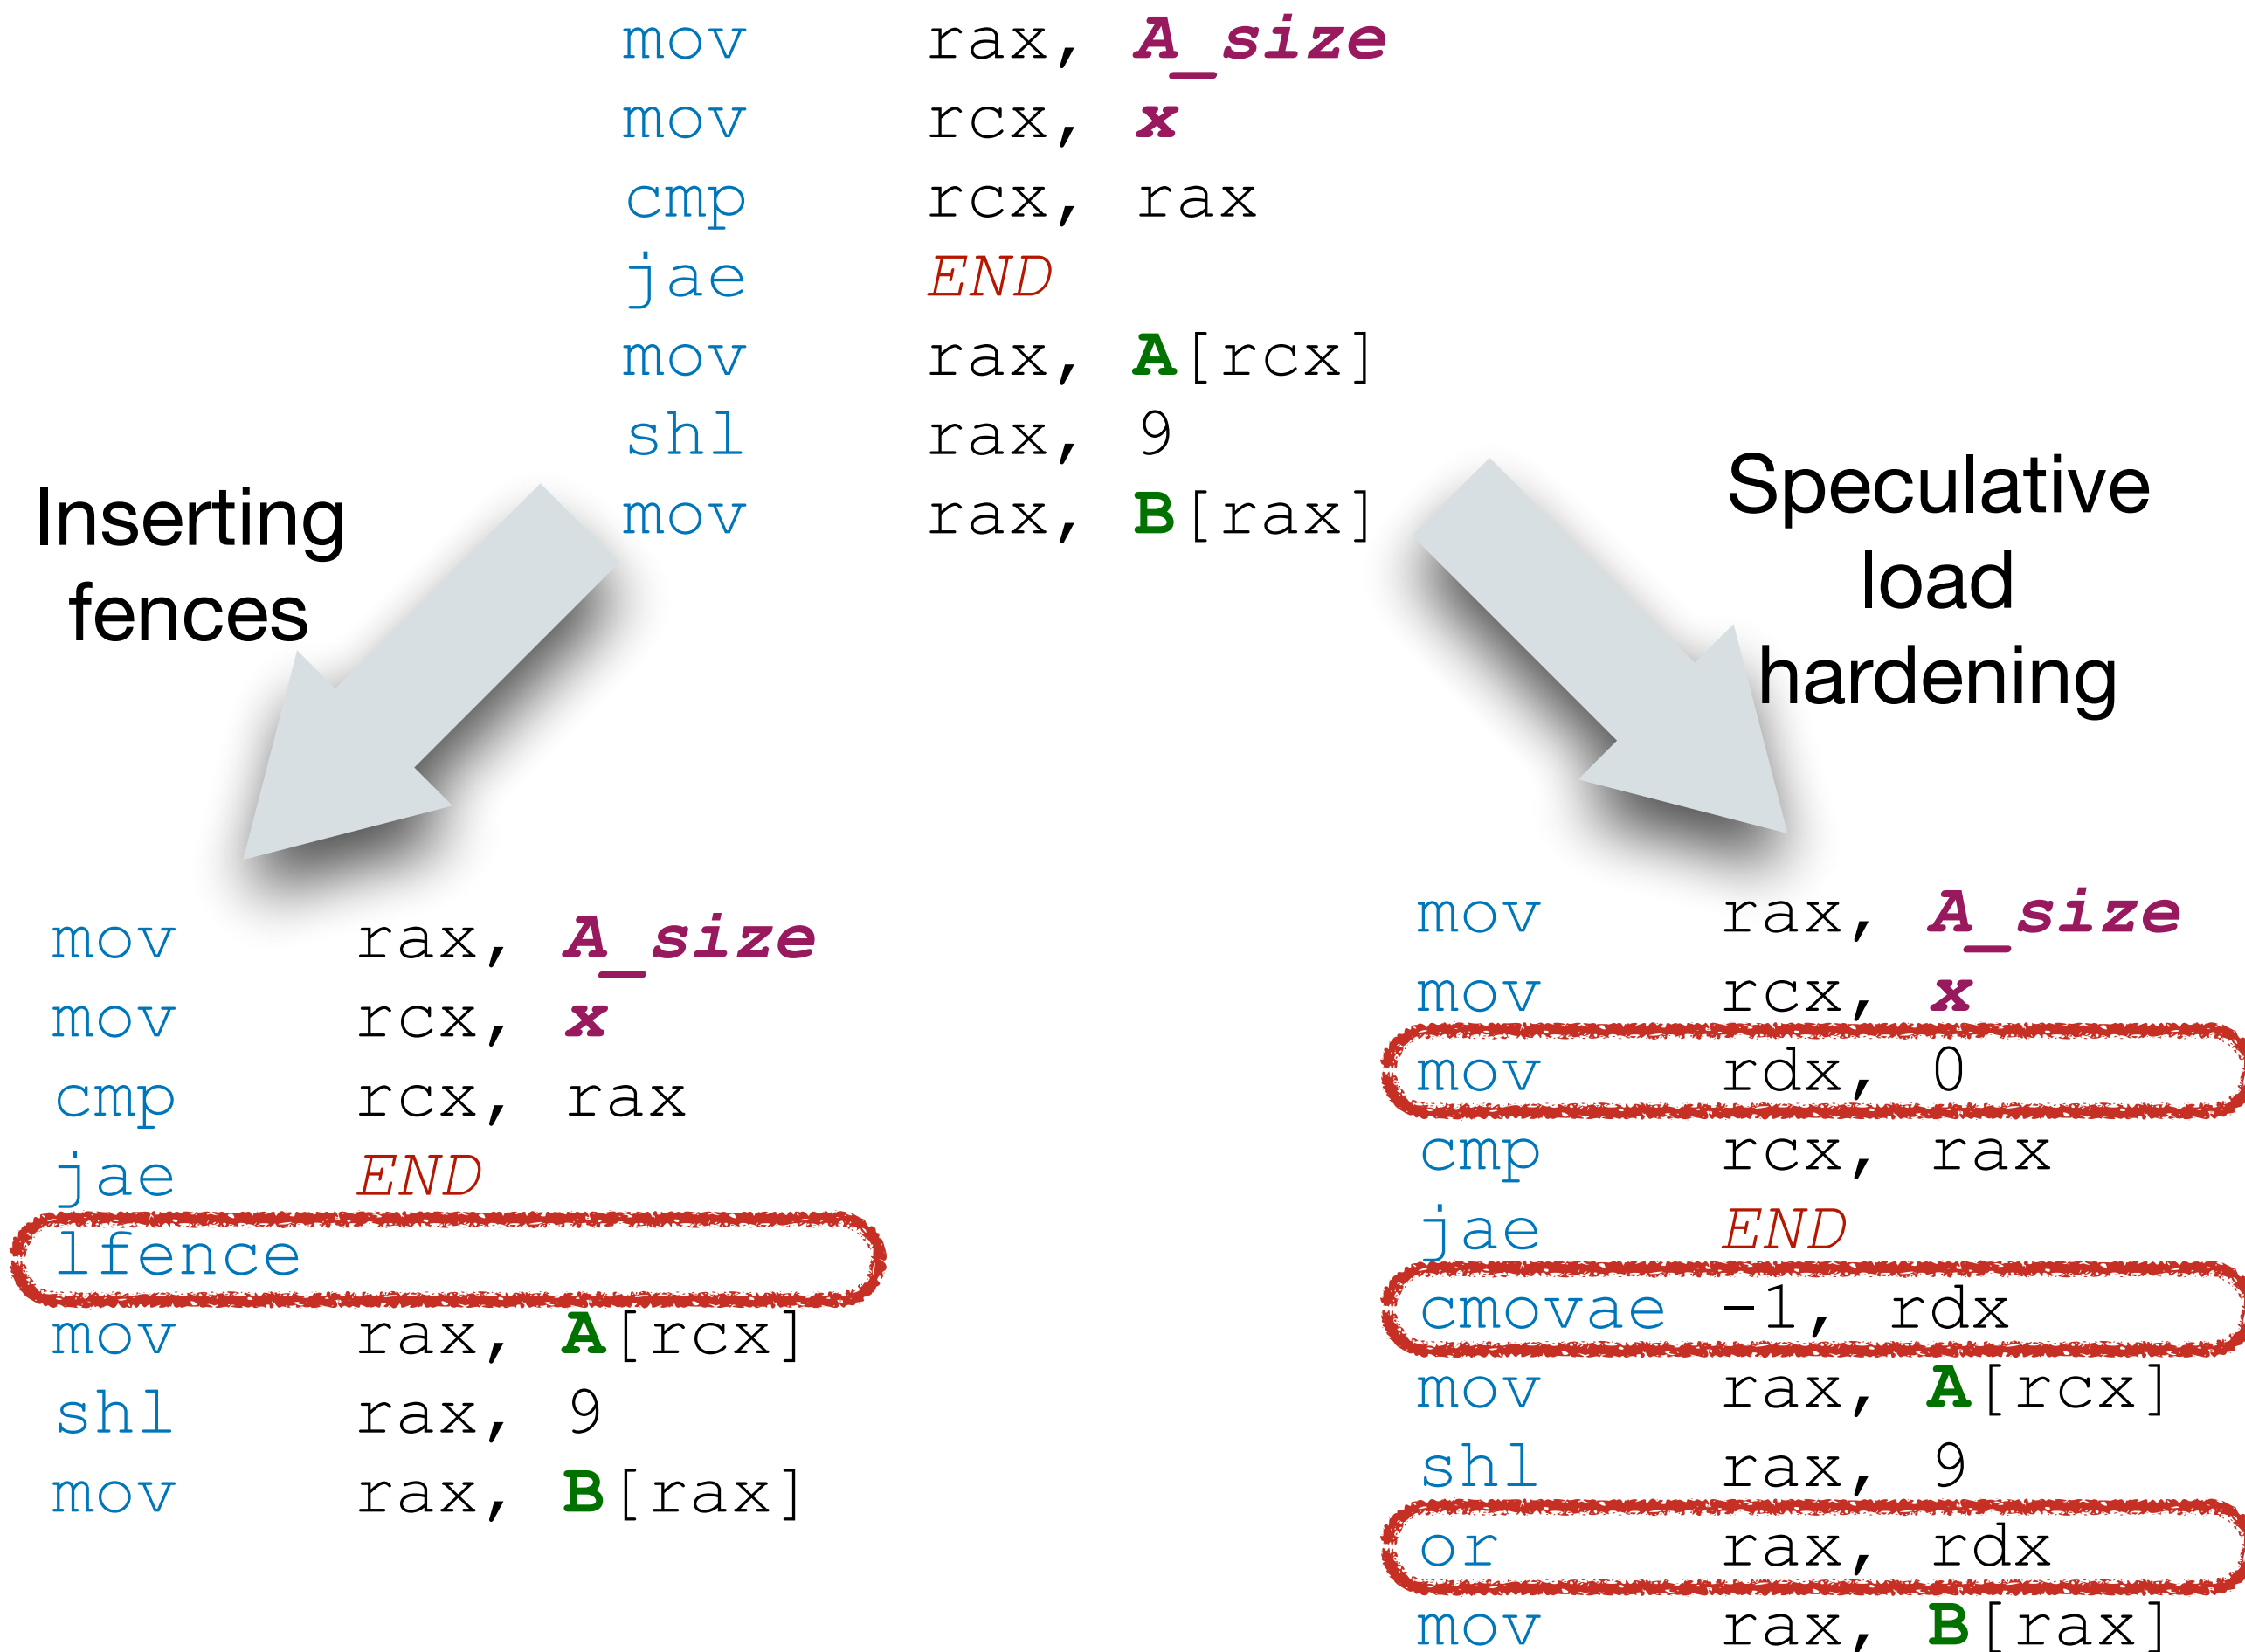
- Intel SGX SDK
- Xen hypervisor
- microkernels



Need: Scalable analysis of speculative non-interference

- Exploit “locality” of speculative execution
- Develop scalable abstractions

Verifying compiler-level countermeasures



How can we **verify** such countermeasures?

A sound HW/SW security contract

Instruction-set architecture: to weak for security guarantees

HW/SW security contract

Microarchitecture: not available publicly, and too detailed for analysis

Find out more in the paper:
<https://arxiv.org/abs/1812.08639>

To appear in: *IEEE Symposium on Security & Privacy, 2020*

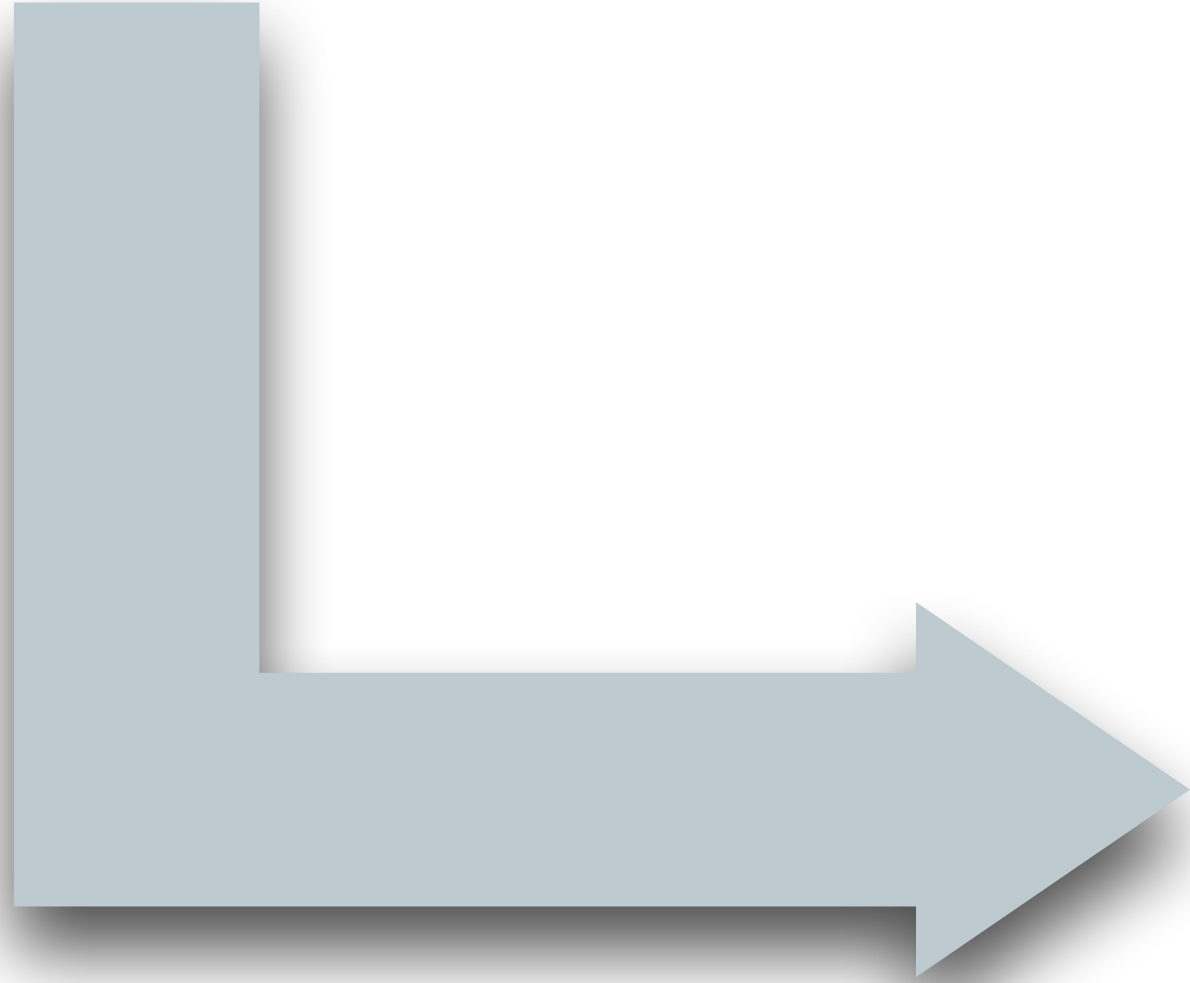
I am looking for PhD students and postdocs!
Feel free to get in touch by email.

Thank you for your attention!

Backup

Example #01 - SLH

```
if (x < A_size)
  y = B[A[x]*512]
```

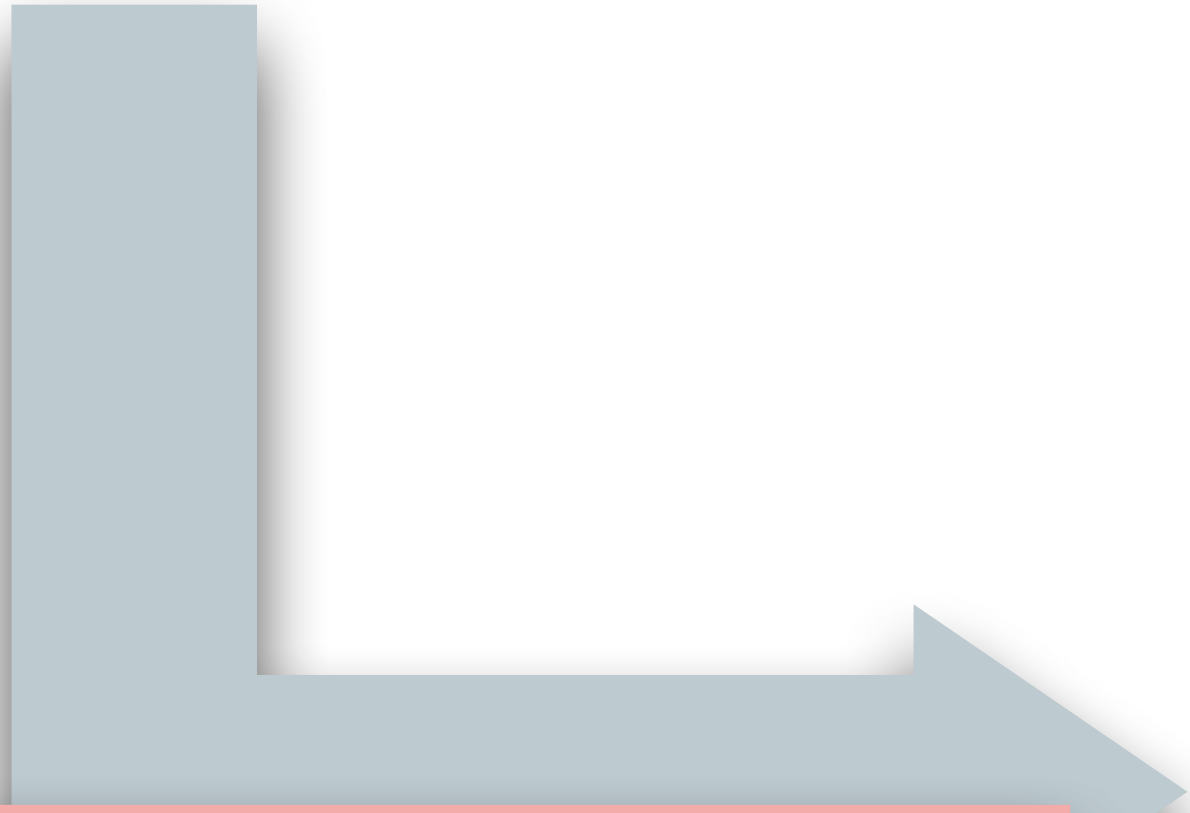


```
mov    rax, A_size
mov    rcx, x
mov    rdx, 0
cmp    rcx, rax
jae    END
cmovae -1, rdx
mov    rax, A[rcx]
shl   rax, 9
or    rax, rdx
mov    rax, B[rax]
```

rax is -1 whenever $x \geq A_size$
We can prove security

Example #10 - SLH

```
if (x < A_size)
  if (A[x] == 0)
    y = B[0]
```



Leaks `A[x] == 0` via
control-flow
We detect the leak!

```
mov    rax, A_size
mov    rcx, x
mov    rdx, 0
cmp    rcx, rax
jae    END
cmovae -1, rdx
mov    rax, A[rcx]
jne    rax, END
cmovne -1, rdx
mov    rax, [B]
```

Example #08 - FEN

$y = B[A[x < A_size ? (x+1) : 0] * 512]$



`lfence` is unnecessary

```
mov    rax, A_size
mov    rcx, x
lea    rcx, [rcx+1]
xor    rdx, rdx
cmp    rcx, rax
cmovae rdx, rcx
mov    rax, A[rdx]
shl   rax, 9
lfence
mov    rax, B[rax]
```