SECURE COMPILATION: SOFTWARE FAULT ISOLATION AND INFORMATION FLOW PRESERVATION

# F. Besson, S. Blazy, A. Dang, T. Jensen, P. Wilke

Celtique/Inria/Univ Rennes

GDR MFSEC, MARS 2021

What is the expected guarantee?

Semantic preservation

If  $beh(S) \neq \emptyset$  Then  $beh(T) \subseteq beh(S)$ .

If source is deterministic, target has same behaviour.
 If source has undefined behaviour, all bets are off.
 Beware: aggressive optimisations exploit undefined behaviours<sup>1</sup>.

Formal verification: CompCert, Vellum, CakeML

<sup>&</sup>lt;sup>1</sup>Undefined behavior: what happened to my code?, Wang et al. [2012]

Hyp1: My compiler has no bug (e.g., LLVM)

Hyp2 : My program has no UB (e.g., Linux kernel)

Functional properties are preserved.

 $\Rightarrow$  I can reason at source level!

# SECURITY PROPERTIES OF TARGET CODE?

#### Compilers may enhance security

- Shadow stack
- Canaries
- Security instrumentation

Compilers may also break security counter-measures<sup>1</sup>

- Introduction of jump breaks CT-programming
- Associativity of xor breaks masking
- CSE breaks Fault-Injection protection
- (Dead) code removal breaks CFI; breaks safe erasure
- $\Rightarrow$  Security people do not trust compilers.

<sup>1</sup>The Correctness-Security Gap in Compiler Optimization, D'Silva et al. [2015]

A secure compiler inserts security counter-measures (in the source) and preserves them (in the assembly).

Attackers get a **disadvantage** at attacking the target.

**Research Agenda** 

- Define classes of properties/attackers.
- Revisit/patch existing compiler passes.

## Security Enhancement: Software Fault Isolation<sup>1</sup> Property: Integrity of a host running untrusted code

Security Preservation: Information Flow Preservation<sup>2</sup> Property: Preservation of lifetime of secrets (secure erasure)

<sup>1</sup>Compiling Sandboxes: Formally Verified Software Fault Isolation, ESOP 2019 <sup>2</sup>Information-Flow Preservation in Compiler Optimisations, CSF 2019

# **SOFTWARE FAULT ISOLATION**

# SOFTWARE FAULT ISOLATION (SFI)

A trusted host wishes to run untrusted guest plugins

### Full speed & Full security

- Speed: native code, same address space
- Security: strong isolation
  - Code: calls limited to host API
  - Data: memory accesses limited to sandbox



"Run safely binary code of untrusted origin"

A modified Compiler ( $\notin$  TCB) masks memory accesses.

Binary verifier ( $\in$  TCB) checks masking is correct.

Ex: (P)NaCl (Google Chrome) [S&P'09, USENIX Sec'10, CACM'10] A compiler:

$$C \rightarrow$$
 CLANG  $\rightarrow$  LLVM  $\rightarrow$  SFI  $\rightarrow$  binary

A verified verifier [RockSalt, PLDI'12]

verifier : binary  $\rightarrow \mathbb{B}$ 

# PORTABLE SOFTWARE FAULT ISOLATION [KROLL ET AL., CSF'14]

$$C \rightarrow [Frontend] \rightarrow Cm \rightarrow [SFI] \rightarrow Cm' \rightarrow [Backend] \rightarrow Asm$$

# Property (SFI Security)

A program P is **SFI-secure** if all its memory accesses are within the sandbox memory region.

Property (Safety)

A program P is safe if all its behaviours are defined i.e. not stuck

Transfer of security from *Cm*<sup>'</sup> down-to *Asm*.

Let *P* be a program that is both **SFI-secure** and **safe**. Let  $B \in behave(ccomp(P))$ . By semantic preservation,  $B \in behave(P)$  (*P* is **safe**) *B* is a secure behaviour (*P* is **SFI-secure**).

8 / 34

# **OUR WORK: FULLY VERIFIED SFI WITHIN COMPCERT**

# Machined-checked proof of SFI-security and Safety

- Security: see [Kroll et al.]
- **Safety:** Re-design of the SFI transformation

#### Reduced TCB (no axiom for)

- Sandboxing memory accesses Low-level pointer arithmetic
- Control-flow integrity
   Trampoline indirect function calls

#### Other features

- Support for multi-threading
- Trusted Runtime

# SANDBOX RELOCATION





<i>msr</i> (A)       A    G    I    U    V	<i>msr</i> (A)       A    G    I    U    V
--	--



msk(A)

TA	G	Т	U	V	
----	---	---	---	---	--

sb

 T
 A
 G
 O
 O
 ...
 T
 A
 G
 F
 F
 F

A A & OxFFF (A & OxFFF)|TAGOOO



msk(A)

TA	G	Т	U	V
----	---	---	---	---



 $msk(A) = (A\&(2^k-1))|\&sb$ 



 $msk(A) = (A\&(2^k-1))|\&sb$ 

## Masking pointer arithmetic has no C semantics

SecurityVacuously trueSafetyVacuously false

#### Pointers are compiled into their numeric value

```
sfi(&sb) = tag × 2<sup>k</sup>
...
sfi(*(e)) = *(sfi(e)&msk + &sb)
sb[2^k]= {5;...};
long foo(sp:int,bar:int -> int -> unit){
    sp1=sp + 8 ;
    *bar(sp1,*(&sb),sp);
    return(*(sp&msk + &sb));
}
```

# **EXPERIMENTS WITH COMPCERT BENCHMARKS**



# COMPARISON WITH (P)NACL



■ GCCSFI/CLANGSFI very competitive

COMPCERTSFI average overhead 9% (removing outliers)

 $\Rightarrow$  Optimisations improve SFI

# Our source SFI pass deviates existing binary instrumentations:

- Masking without bitwise pointer arithmetics
- Source level control-flow integrity

### Compilers can be used for Security

PriceGuarantee only holds for safe programsLimitationCompilers only preserve observable behavioursBut, security is not always reducible to safety.

# **INFORMATION FLOW PRESERVATION**

Our Information-Flow Preservation property aims at protecting against:

- Data remanence
- Lifetime extension
- Increased information leakage
- Duplication of information

# Dead Store Elimination (DSE) is not secure<sup>1</sup>



<sup>1</sup>Dead Store Elimination (Still) Considered Harmful, Yang et al. [2017]

Code motion is not secure.



Common Expression Elimination is not secure.



Register Allocation is not secure.



# FORMAL DEFINITION OF IFP

## Trace based execution model

### Memory states: data observable by attackers



- Attackers know the code
- Attackers observe n bits in the trace



- Attackers know the code
- Attackers observe n bits in the trace



- Attackers know the code
- Attackers observe n bits in the trace



- Attackers know the code
- Attackers observe n bits in the trace



- Attackers know the code
- Attackers observe *n* bits in the trace



- Attackers know the code
- Attackers observe *n* bits in the trace



## **RATIONALE FOR HIERARCHY OF ATTACKERS**



equally insecure for a strong attacker

## **RATIONALE FOR HIERARCHY OF ATTACKERS**



- equally insecure for a strong attacker
- p1 is secure for the 1-bit attacker

# ATTACKER KNOWLEDGE<sup>1</sup>

- Attackers try to guess the initial memory used
- Possible initial memories matching its observations



<sup>1</sup>Gradual Release: Unifying Declassification, Encryption and Key Release Policies, Askarov and Sabelfeld [2007]

# ATTACKER KNOWLEDGE<sup>1</sup>

Attackers try to guess the initial memory usedPossible initial memories matching its observations



<sup>1</sup>Gradual Release: Unifying Declassification, Encryption and Key Release Policies, Askarov and Sabelfeld [2007]

# IFP TRANSFORMATION (1/2)

#### Intuition

Any information that can be learned with a trace observation of the transformed program can also be learned with the source program



# IFP TRANSFORMATION (1/2)

#### Intuition

Any information that can be learned with a trace observation of the transformed program can also be learned with the source program



> Source program  $p_1$ Transformed program  $p_2$





For any execution from the same initial memory  $m_{\rm o}$ 



For attackers with any observation capabilities



Exists lockstep pairings of observations from  $t_2$  to  $t_1$ 



For any observation  $o_2$  of size n on the trace  $t_2$ 





# TRANSLATION VALIDATION FOR REGIS-TER ALLOCATION

# **REGISTER ALLOCATION**

### Introduce spilling of values in the stack

- Usually not IFP:
  - Duplication on both stack and registers
  - Erasure may not be applied to both locations

#### Example with a 2-register machine:



# **REGISTER ALLOCATION**

### Introduce spilling of values in the stack

- Usually not IFP:
  - Duplication on both stack and registers
  - Erasure may not be applied to both locations

#### Example with a 2-register machine:



# VALIDATION AND PATCHING TOOLCHAIN

- Validator verifies the sufficient condition
- Detected leakage are patched





$$\begin{array}{rrrr} k & \leftarrow & r1 \\ t & \leftarrow & r2 \\ salt & \leftarrow & stack\_salt \end{array}$$













## PATCHING LEAKAGE

#### Leakage are patched with constant values



$$k \leftarrow r^2$$
salt  $\leftarrow$  stack\_salt
$$\bullet \leftarrow stack_k$$

- Observation points are placed at function calls and returns
- On the verified compiler CompCert<sup>1</sup>
- We measure the impact of patching on the programs
- Correctness is ensured by CompCert original validator

<sup>1</sup>Formal Certification of a Compiler Back-end, Leroy [2006]

## MEASURING IMPACT OF PATCHING



## MEASURING IMPACT OF PATCHING



# **Related work and Conclusion**

# Securing a compiler transformation<sup>12</sup>

- preserve programs that do not leak
- does not differentiate between degrees of leakage

Preservation of side-channel countermeasures<sup>3</sup> <sup>4</sup>

- framework to preserve security properties
- different leakage model
- use a 2-simulation property

<sup>1</sup>Securing a Compiler Transformation, Deng and Namjoshi [2016]
 <sup>2</sup>Securing the SSA Transform, Deng and Namjoshi [2017]
 <sup>3</sup>Secure Compilation of Side-Channel Countermeasures, Barthe et al. [2018]
 <sup>4</sup>Formal verification of a constant-time preserving C compiler, Barthe et al. [2020]

General purpose compilers are not designed for security

They aim at preserving **observable** behaviours

- $\blacksquare$  Software Fault Isolation  $\checkmark$
- Information Flow X

In theory, compiler may not preserve information flow In practice, they do break security of the source code

The **best** compilers are the **least** secure!  $\Rightarrow$  Optimisations need to be carefully reviewed

An opportunity for secure (and verified) compilation?