# High System-Code Security with Low Overhead

Jonas Wagner, Volodymyr Kuznetsov, George Candea, and Johannes Kinder



École Polytechnique Fédérale de Lausanne



Royal Holloway, University of London

### High System-Code Security?

Today's software is dangerous.

Example: OpenSSL Overflow in ssl/t1\_lib.c:3997 →

OpenSSL contains 53,073 memory accesses. How to protect them all?



### Protect *all* dangerous operations using sanity checks \langle Checks are automatically added at compile time \] ✓ No source code modification is needed



AddressSanitizer

\*p = 42;

if (!isValidAddress(p)) { reportError(p); abort(); \*p = 42;

### Problem: Sanity checks cause high performance overhead

:

ΤοοΙ	Ανς
AddressSanitizer (memory errors)	
SoftBound/CETS (full memory safety)	
UndefinedBehaviorSanitizer (integer overflows, type errors)	
Assertions, code contracts,	

#### g. Overhead

73%

116%

71%

#### depends



### Problem: Sanity checks cause high performance overhead

### People use checks heavily for testing, but disable them in production

### Goal: checks in production



# Insight: Checks are not all equal

Most of the *overhead* comes from a few expensive checks

Checks in hot code, each executed many times

Most of the *protection* comes from many cheap checks

Checks in cold code

Lets users choose their *overhead budget* (e.g., 5%) Automatically identifies sanity checks in software Analyzes the cost of every check Selects as many checks as fit in the user's budget

Most of the overhead comes from a few expensive checks Most of the protection comes from many cheap checks



## ASAP Insight & Results



		مر می و بر می مان این این این می و بر و می مسلم این می می ماند این می و بر و می و بر و می و بر و می و مسلم این می می و بر و می و بر و می و بر و می و بر و می	
ocks a	re verv	expensi	ive
-0/ h			
alop	efation	5	
Refer	ie prote	ction	
ome	ς		
(S			
30%	40%	50%	60%

#### **Overhead**

### Outline

Introduction: What is ASAP?

Design

Key Algorithms

Results

Conclusion

# Design

ASAP is built into the compiler
✓ Easy to use (set CC and CFLAGS)
✓ Compatible (parallel compilation, ...)

Compiler (LLVM)	
Identify sanity checks	P m che

Profiler: neasure eck costs

Optimizer: select maximum set of checks

#### Users use ASAP like a regular compiler that adds checks



### ASAP stores intermediate compiler output



ASAP generates a program variant with profiling



#### Users use ASAP like a regular compiler that adds checks

### instrumentation. Users run this to measure check costs.

#### Users use ASAP like a regular compiler that adds checks



ASAP generates a program variant with profiling



#### ASAP uses costs & budget to generate an optimized program



### instrumentation. Users run this to measure check costs.



### Outline

Introduction: What is ASAP?

Design

Key Algorithms

Results

Conclusion



#### Budget

ASAP







#### Can users trust ASAP to select checks that use the least CPU cycles?

#### If ASAP says you're 87% protected, what does this mean?

...
if (!isValidAddress(p)) {
 reportError(p);
 abort();
}
\*p = 42;
...

```
prof1++;
if (!isValidAddress(p)) {
    prof2++;
    reportError(p);
    abort();
}
prof3++;
*p = 42;
```

#### 1. Add profiling counters

```
prof1++;
if (!isValidAddress(p)) {
    prof2++;
    reportError(p);
    abort();
}
prof3++;
*p = 42;
```

- 1. Add profiling counters
- 2. Identify check instructions

```
prof1++;
if (!isValidAddress(p)) {
    prof2++;
    reportError(p);
    abort();
}
prof3++;
*p = 42;
```

- 1. Add profiling counters
- 2. Identify check instructions
- 3. Use static model

of cycles per instruction

```
prof1++;
if (!isValidAddress(p)) {
    prof2++;
    reportError(p);
    abort();
}
prof3++;
*p = 42;
...
```

#### Precise cost in CPU cycles

- 1. Add profiling counters
- 2. Identify check instructions
- 3. Use static model of cycles per instruction
- 4. Compute cost for check in CPU cycles:





#### If ASAP says you're 87% protected, what does this mean?

ASAP quantifies protection using the *sanity level* 

We would like to know the effective protection level

Methodology: measure how many bugs/vulnerabilities are effectively prevented

ASAP quantifies protection using the *sanity level* 

We would like to know the effective protection level

Methodology: measure how many bugs/vulnerabilities are effectively prevented

#### **Experiment 1**

Source code of Python 2.7

Bug: line that has received a patch between version 2.7.1 and 2.7.8

#### Sanity level 87% ≈ 91% protection ★

ASAP quantifies protection using the *sanity level* 

We would like to know the *effective protection level* 

Methodology: measure how many bugs/vulnerabilities are effectively prevented







#### All of these are in cold code



#### **Experiment 3**

Vulnerabilities from CVE DB

Analyze 145 vulnerabilities from 2014

- Memory errors
- Open source
- Patch available
- Error location known

83% of these are in cold code

#### Checks in cold code provide real value



### Outline

Introduction: What is ASAP?

Design

Key Algorithms

Results

Conclusion

Overhead for: SPEC benchmarks AddressSanitizer

# 100/0







Overhead for: SPEC benchmarks AddressSanitizer

95%







Overhead for: SPEC benchmarks AddressSanitizer

90%





Overhead for: SPEC benchmarks AddressSanitizer

80%







### Conclusion

#### Run-time checks deliver strong protection at high cost



Most of the overhead comes from a few expensive checks Most of the protection comes from many cheap checks

**Protect your software!** dslab.epfl.ch/proj/asap



### Backup slides

# CHECK ALL THE ACCESSES

### pay all the overhead?



; <label>:0 %5 = **lshr** i64 %4, 3 %9 = icmp eq i8 %8, 0

## Sanity Check in LLVM

; <label>:18 %19 = load i32\* %3, align 4

```
%1 = load i32* %fmap_i_ptr, align 4
%2 = zext i32 %1 to i64
%3 = getelementptr inbounds i32* %eclass, i64 %2
%4 = ptrtoint i32* %3 to i64
%6 = add i64 %5, 17592186044416
%7 = inttoptr i64 %6 to i8
%8 = load i8* %7, align 1
br i1 %9, label %18, label %10
           ; <label>:10
             %11 = ptrtoint i32* %3 to i64
             %12 = and i64 %11, 7
             %13 = add i64 %12, 3
             %14 = trunc i64 %13 to i8
             %15 = jcmp slt i8 %14, %8
             br i1 %15, label %18, label %16
                     ; <label>:16
                       %17 = ptrtoint i32* %3 to i64
                       call void @__asan_report_load4(i64 %17) #3
                       call void asm sideeffect "", ""() #3
                       unreachable
```

