

# **Shadowclone: Thwarting and Detecting DOP Attacks with Stack Layout Randomization and Canary**

EECS 583

Yunjie Pan, Shibo Chen, Cheng Chi, Yifan Guan

# Shadowclone

- Motivation & Background
- Methodology
- Implementation
- Evaluation
- Demo

# Shadowclone

- Motivation & Background
- Methodology
- Implementation
- Evaluation
- Demo

# DOP Example

Modeled after FTP server

Uses a stack buffer overflow vulnerability to control a few stack variables



```

                                     Affected Variables
1 struct server{ int * cur_max, total, typ;} *srv;
2 int connect_limit = MAXCONN; int *size, *type;
3 char buf[MAXLEN]; Buffer Overflow
4 size = &buf[8]; type = &buf[12];
5 ...
6 while(connect_limit--){ Gadget Dispatcher
7     readData(sockfd, buf); // stack bof
8     if(*type == NONE) break;
9     if(*type == STREAM) { // condition
10         *size = *(srv->cur_max); // dereference
11     } else { Gadgets
12         srv->typ = *type; // assignment
13         srv->total += *size; // addition
14     }
15     ...
16 }
```

# Simple Example

## Simple Stack Overflow

If successful, var1 and var 2 will be changed to 583

```
1 #include <stdio.h>
2 #include <stdlib.h>
3
4 void func(){
5     double var1 = 483.0;
6     int var2 = 483;
7     char var3 = 'e';
8     char buff[4];
9     char* buff_ptr = &buff[0];
10    size_t size = 1024;
11    getline(&buff_ptr, &size, stdin);
12    printf("var1: %f\n", var1);
13    printf("var2: %i\n", var2);
14    printf("var3: %c\n", var3);
15    printf("buff: %s\n", buff);
16    printf("This is function func.\n");
17 }
18
19 int main(){
20    printf("Calling func:\n");
21    func();
22    printf("func returned.\n");
23 }
```

Affected Variables

Buffer Overflow

# Prior work - Smokestack

Randomizes the order of stack variables during runtime with P-BOX

- + Much harder to deliver DOP attacks
- + Negligible memory overhead
- Runtime performance overhead
- Cannot detect attacks when happening

# Goals

- Reduce runtime overhead by compile time randomization
- Detect attacks when happening

# Threat Model

- CFI (Control Flow Integrity) defenses deployed
- Stack buffer overflow vulnerability
- Attackers cannot see the code, but can learn gradually




# Shadowclone

- Motivation & Background
- Methodology
- Implementation
- Evaluation
- Demo

# Shadowclone

- Generate compile-time randomized clones of vulnerable functions
- Insert compile-time random canary into stack variables and check before the function returns
- Randomly select copy to execute in run time



```
1 #include <stdio.h>
2 #include <stdlib.h>
3
4 void func(){
5     double var1 = 483.0;
6     int var2 = 483;
7     char var3 = 'e';
8     char buff[4];
9     char* buff_ptr = &buff[0];
10    size_t size = 1024;
11    getline(&buff_ptr, &size, stdin);
12    printf("var1: %f\n", var1);
13    printf("var2: %i\n", var2);
14    printf("var3: %c\n", var3);
15    printf("buff: %s\n", buff);
16    printf("This is function func.\n");
17 }
18
19 int main(){
20    printf("Calling func:\n");
21    func();
22    printf("func returned.\n");
23 }
```

Affected  
Variables

Buffer Overflow

Randomized  
Stack order



```
1 void func1(){
2     int var2 = 483;
3     uint32_t canary = 1092384;
4     double var1 = 483.0;
5     char buff[4];
6     char var3 = 'e';
7     char* buff_ptr = &buff[0];
8     size_t size = 1024;
9     getline(&buff_ptr, &size, stdin);
10    ...
11    if (canary != 1092384){
12        exit(1);
13    }
14 }
15
16 ...
17
18 void func_wrapper(){
19     int fp_index = rand() % 3;
20     if (fp_index == 0){
21         func0();
22     } else if (fp_index == 1){
23         func1();
24     } else {
25         func2();
26     }
27 }
```

Canary Var

Hard-coded canary check  
(cmp embedded constant)

Randomly select  
clone to execute

# Shadowclone

- Motivation & Background
- Methodology
- Implementation
- Evaluation
- Demo

# Implementation

1. Find all concrete functions (except main and *syscall*)
2. Find all *alloca* instructions
3. Clone a function  $\min(\textit{threshold}, \text{num}(\textit{alloca})!)$  times
4. Randomize the order of stack variables
5. Insert the canary and checks
6. Convert original function to randomly select a clone in run-time

# Randomize the order of stack variables

Generate a random ordering (a configuration)

If this configuration already exists:

*Continue*

Else:

*Apply this config to one of the clones*

Repeat until all clones have been randomized

# Insert canary and checks

1. Randomly select an insertion point and insert a 32-bit canary
2. Generate a random number and store it to the location of our canary
3. Insert a *compare-and-branch* duo before each *return* instruction  
(branch to the exception handler if compromise detected)

```

1 define void @func() #0 {
2   %1 = alloca double, align 8
3   %2 = alloca i32, align 4
4   %3 = alloca i8, align 1
5   %4 = alloca [4 x i8], align 1
6   %5 = alloca i8*, align 8
7   %6 = alloca i64, align 8
8   store double 4.830000e+02, double*
9   store i32 483, i32* %2, align 4
10  store i8 101, i8* %3, align 1
11  %7 = getelementptr inbounds [4 x i8
12  store i8* %7, i8** %5, align 8
13  store i64 1024, i64* %6, align 8
14  %8 = load %struct.__sFILE*, %struct
15  %9 = call i64 @getline(i8** %5, i64
16  ret void
17 }

```

```

1 define dso_local void @func.1() #0 {
2 entry:
3   %buff = alloca [4 x i8], align 1
4   %var1 = alloca double, align 8
5   %canary = alloca i32
6   store i32 780689205, i32* %canary
7   %var2 = alloca i32, align 4
8   %buff_ptr = alloca i8*, align 8
9   %var3 = alloca i8, align 1
10  %size = alloca i64, align 8
11  store double 4.830000e+02, double* %var1, align 8
12  store i32 483, i32* %var2, align 4
13  store i8 101, i8* %var3, align 1
14  %arrayidx = getelementptr inbounds [4 x i8], [4 x i8]* %buff
15  store i8* %arrayidx, i8** %buff_ptr, align 8
16  store i64 1024, i64* %size, align 8
17  %0 = load %struct._IO_FILE*, %struct._IO_FILE** @stdin, align
18  %call = call i64 @getline(i8** %buff_ptr, i64* %size, %struc
19  %1 = load i32, i32* %canary
20  %2 = icmp eq i32 %1, 780689205
21  br i1 %2, label %3, label %func_exit
22
23 3:
24  ret void
25
26 func_exit:
27  call void @detect_breach()
28  br label %3
29 }

```

Canary Var

Hard Coded CMP



# Run-time Selection

get\_rand() is defined in our run-time library

Generates a i32 random number with RDRAND instruction

```
1 define dso_local void @func() #0 {
2   rand_bb:
3     %0 = call i32 @get_rand()
4     %1 = icmp eq i32 %0, 0
5     br i1 %1, label %func_func.1, label %ctrl0
6
7   func_func.1:                                ; preds = %rand_bb
8     call void @func.1()
9     ret void
10
11  func_func.2:                                ; preds = %ctrl0
12    call void @func.2()
13    ret void
14
15  func_func.3:                                ; preds = %ctrl1
16    call void @func.3()
17    ret void
18
19  func_func.4:                                ; preds = %ctrl1
20    call void @func.4()
21    ret void
22
23  ctrl0:                                       ; preds = %rand_bb
24    %2 = icmp eq i32 %0, 1
25    br i1 %2, label %func_func.2, label %ctrl1
26
27  ctrl1:                                       ; preds = %ctrl0
28    %3 = icmp eq i32 %0, 2
29    br i1 %3, label %func_func.3, label %func_func.4
30 }
```

Branch based on random number

# Shadowclone

- Motivation & Background
- Methodology
- Implementation
- Evaluation
- Demo

# Experiment Setup

## ***Platform:***

Xeon Gold 6126, Ubuntu 18.04 Linux, 256GB of memory

## ***Benchmarks:***

Three in-House testcases (*big\_array*, *wc*, and *compress*)

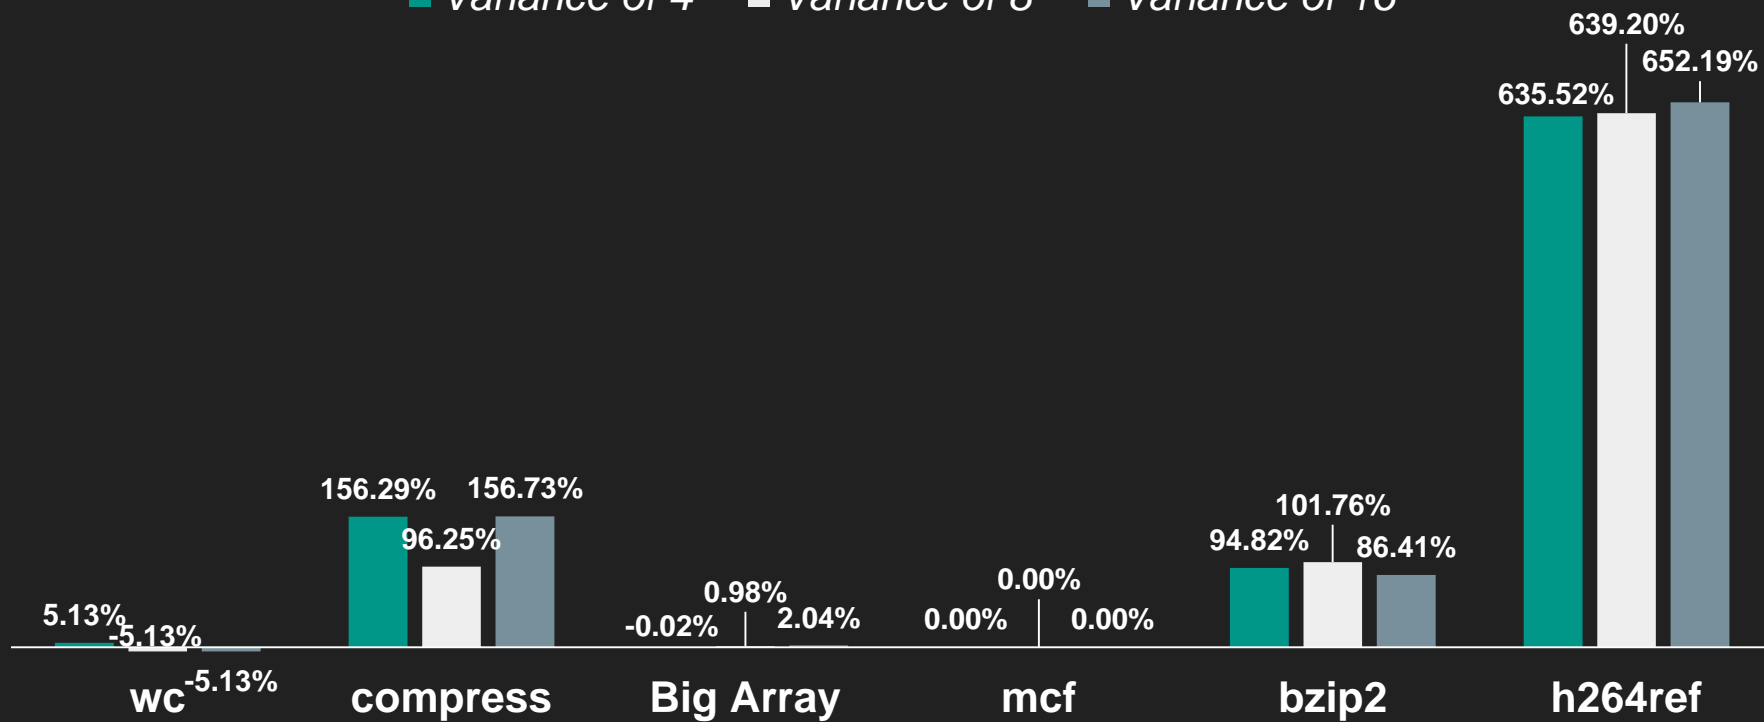
Three Spec06 benchmarks (*bzip2*, *mcf*, and *h264ref*)

## ***Source of random numbers:***


RDRAND

# Performance Overhead

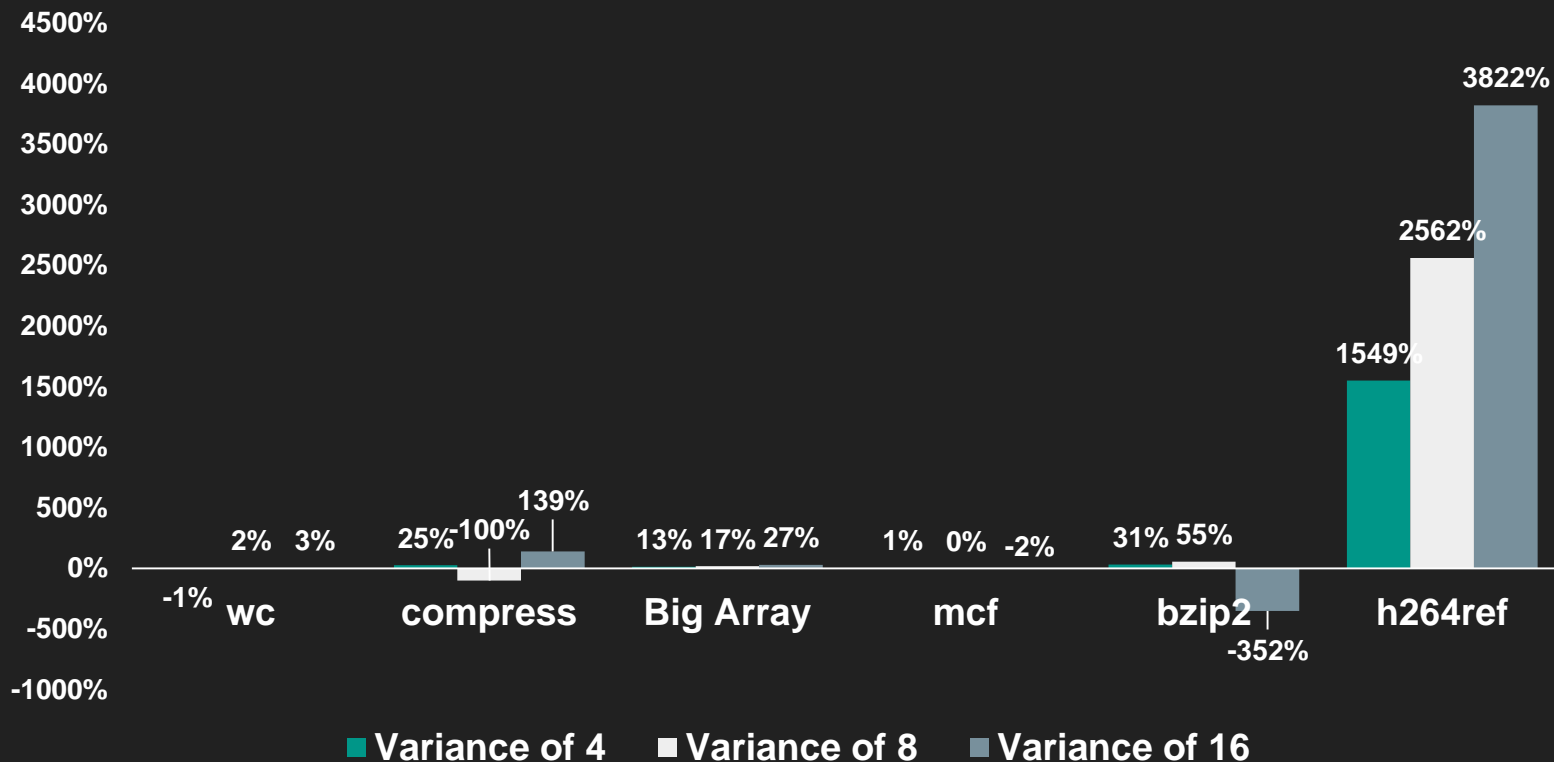
■ Variance of 4   ■ Variance of 8   ■ Variance of 16



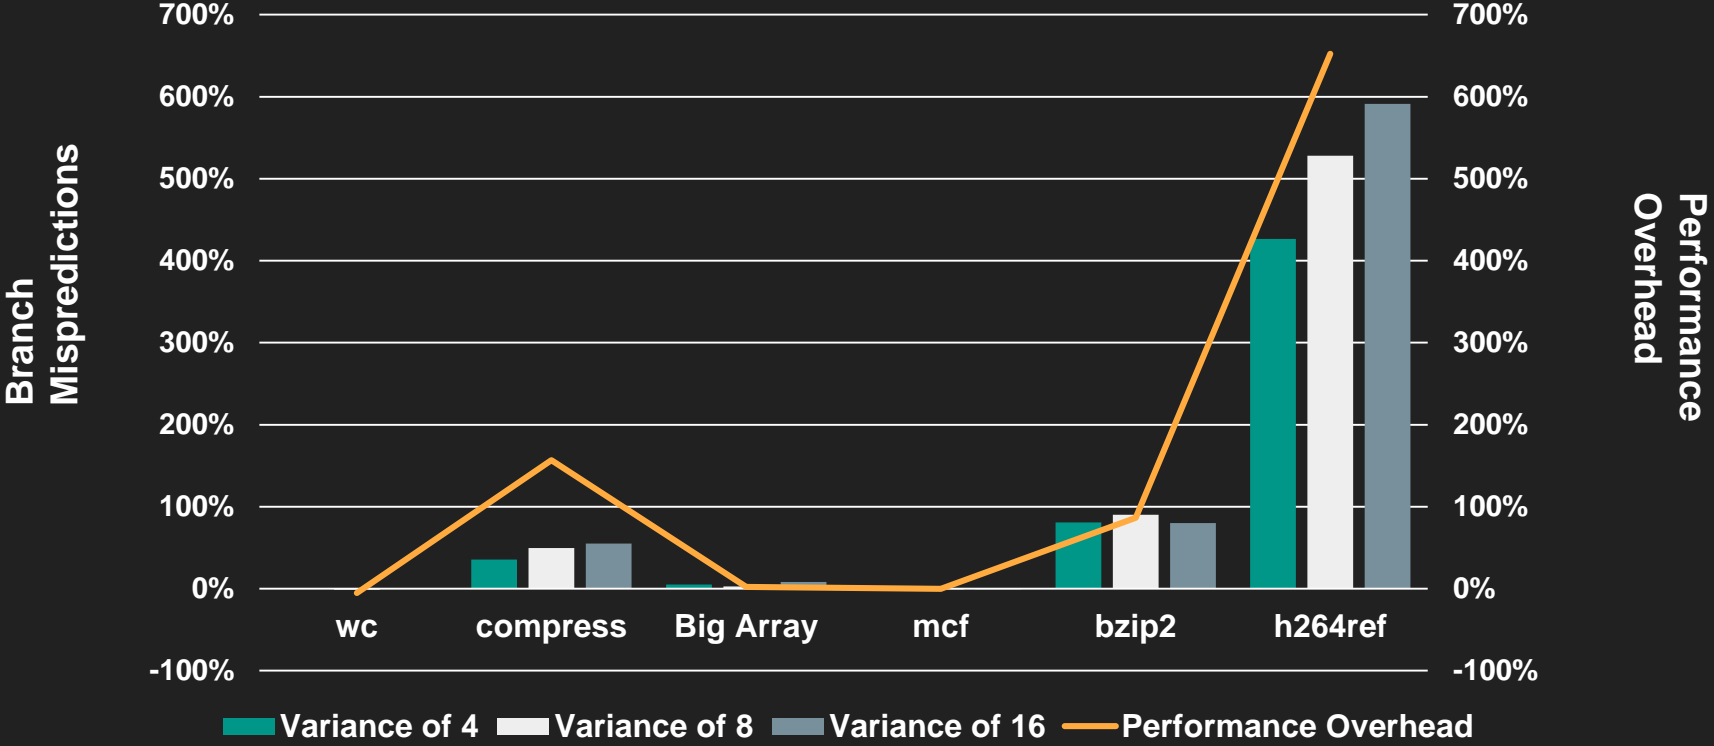
# Spatial Locality: *Code Size (in KB)*



# Spatial Locality: # of I-Cache Misses



# Temporal Locality & Speculation: # of Branch Mispredictions



# Security Analysis

- The attacker learns quickly
  - Learns about any configuration after this very configuration has been run only once
- The attacker doesn't trigger any exception by accident

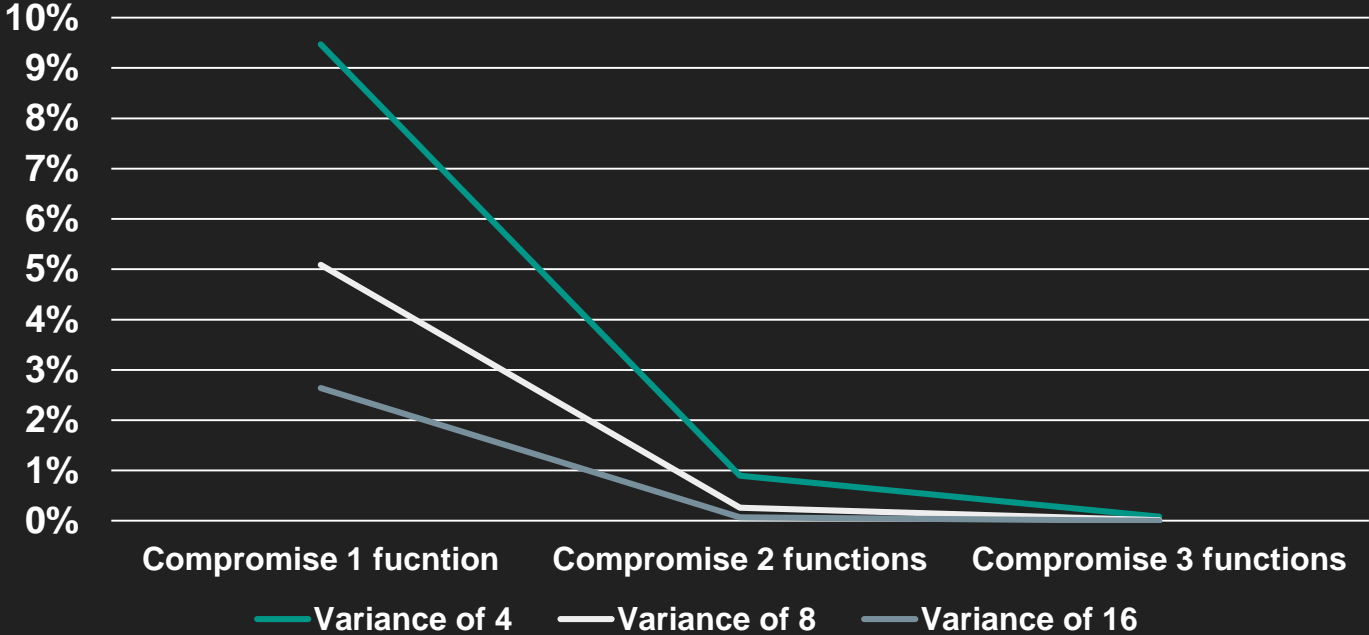
## **Metrics:**

*What's the chance for an attacker to successfully compromise our system without being detected?*



# Security Analysis

Probability of Attackers Successfully Deliver Attack w/o Being Detected



# Shadowclone

- Motivation & Background
- Methodology
- Implementation
- Evaluation
- Demo

# Conclusion

- Shadowclone can efficiently thwart and detect DOP attacks.
- Shadowclone has low performance overhead when running small programs. Its performance deteriorates as the size of program gets larger and the program gets more function calls.

Question?

# Probability of Attackers Successfully Deliver Attack w/o Being Detected

- $P(\text{attacker succeeds without being attacked}) = \sum_{k=1}^{\infty} P(\text{the first } (k - 1) \text{ times failed and without being attacked}) * P(\text{the } k\text{th time succeeds and not being detected})$
- $P(\text{the first time succeed}) = 1/N!$        $P(\text{the } k\text{th time succeed}) = 1/M$ 
  - (N is the average number of stack variables in a function, M is the number of clones)
  - N = 10 in the benchmarks we analyzed
- $P(\text{an attack would be detected}) = 1/2$