

# CS 858: Software Security

## Offensive and Defensive Approaches

### **Attacks: memory corruption**

Meng Xu (*University of Waterloo*)

Fall 2022

# Outline

- 1 Introduction
- 2 Intuition
- 3 Spatial safety
- 4 Temporal safety
- 5 Countermeasures

## Definition: memory

**Q:** What is “memory” in memory corruption?

## Definition: memory

**Q:** What is “memory” in memory corruption?

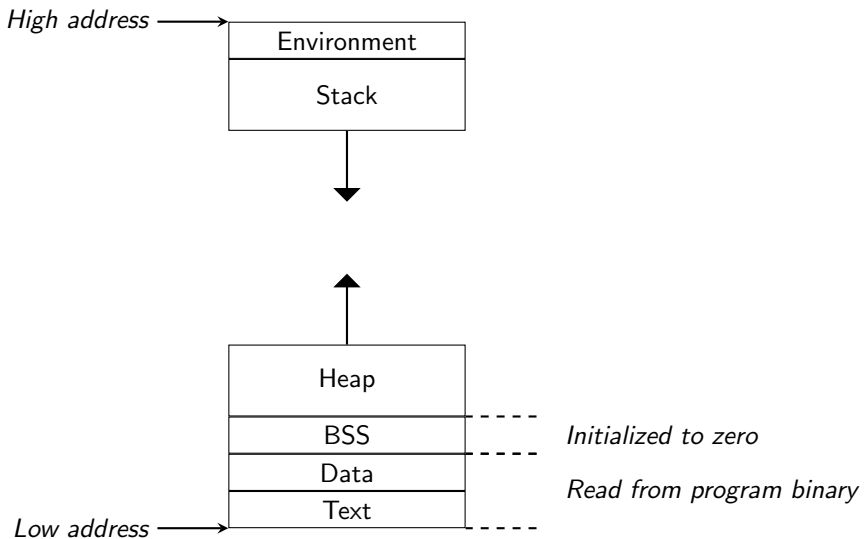
**A:** Three types of memory in system level:

- Stack
- Heap
- Global (a.k.a., static)

# Example

```
1 #include <stdlib.h>
2
3 // this is in the data section
4 const char *HELLO = "hello";
5
6 // this is in the BSS section
7 long counter;
8
9 void foo() {
10     // this is in the stack memory
11     int val;
12
13     // the msg pointer is in the stack memory
14     // the msg content is in the heap memory
15     char *msg = malloc(120);
16
17     // msg content is explicitly freed here
18     free(msg);
19
20     // the val and msg pointer is implicitly freed here
21 }
22
23 // the global memory is only destroyed on program exit
```

# Memory layout (Linux x86-64 convention)



# Stack layout (Linux x86-64 convention)

```
1 long foo(  
2     long a, long b, long c,  
3     long d, long e, long f,  
4     long g, long h)  
5 {  
6     long xx = a * b * c;  
7     long yy = d + e + f;  
8     long zz = bar(xx, yy, g + h);  
9     return zz + 20;  
10 }
```

High address

RBP + 24

h

RBP + 16

g

RBP + 8

return address

RBP

saved rbp

RBP - 8

xx

RBP - 16

yy

RBP - 24

zz

Low address

Argument a to f passed by registers.

# Heap layout (GNU C library implementation)

Refer to the article from [Azeria Labs](#).



# Memory layout

**Q:** What about stacks and heap in multi-threaded programs?

# Memory layout

**Q:** What about stacks and heap in multi-threaded programs?

- Each thread has its own stack
- All threads in the same process share the heap and global data

# For exploitation of memory errors

Smashing The Stack For Fun And Profit

How2Heap — Educational Heap Exploitation

# Outline

- 1 Introduction
- 2 Intuition**
- 3 Spatial safety
- 4 Temporal safety
- 5 Countermeasures

# A quick recap

This presentation is about **memory corruption**, a.k.a.,

- **memory errors**, or
- **violations of memory safety properties**, or
- **unsafe programs**

## Definition: safety

**Q:** What is “safety” in memory safety?

## Definition: safety

**Q:** What is “safety” in memory safety?

**Observation 1:** At runtime, memory is a pool of **objects**

## Definition: safety

**Q:** What is “safety” in memory safety?

**Observation 1:** At runtime, memory is a pool of **objects**

**Observation 2:** Each **object** has known and limited **size** and **lifetime**



## Definition: safety

**Q:** What is “safety” in memory safety?

**Observation 1:** At runtime, memory is a pool of **objects**

**Observation 2:** Each **object** has known and limited **size** and **lifetime**

**Observation 3:** Once allocated, the size of an **object** never changes

## Definition: safety

**Q:** What is “safety” in memory safety?

**Observation 1:** At runtime, memory is a pool of **objects**

**Observation 2:** Each **object** has known and limited **size** and **lifetime**

**Observation 3:** Once allocated, the size of an **object** never changes

**Observation 4:** A memory access is always **object-oriented**, i.e.

- Memory read: (object\_id, offset, length)
- Memory write: (object\_id, offset, length, value)

## Definition: safety

**Q:** What is “safety” in memory safety?

**Observation 1:** At runtime, memory is a pool of **objects**

**Observation 2:** Each **object** has known and limited **size** and **lifetime**

**Observation 3:** Once allocated, the size of an **object** never changes

**Observation 4:** A memory access is always **object-oriented**, i.e.

- Memory read: (object\_id, offset, length)
- Memory write: (object\_id, offset, length, value)

Wait..., in C/C++, pointers are just 32/64-bit integers. I can do:  
`int *p = 0xdeadbeef; int v = *p;` Which object I refer to here?

# Definition: safety

Q: What is “safety” in memory safety?

At any point of time during the program execution,  
for any object in memory, we know its  
(object\_id, size [int], alive [bool])

At the same time, for each memory access, we know:

- Memory read: (object\_id, offset [int], length [int])
- Memory write: (object\_id, offset [int], length [int], \_)

# Outline

- 1 Introduction
- 2 Intuition
- 3 Spatial safety**
- 4 Temporal safety
- 5 Countermeasures

## Definition: spatial safety

At any point of time during the program execution,  
for any object in memory, we know its  
(object\_id, size [int], alive [bool])

At the same time, for each memory access, we know:

- Memory read: (object\_id, offset [int], length [int])
- Memory write: (object\_id, offset [int], length [int], \_)

## Definition: spatial safety

At any point of time during the program execution,  
for any object in memory, we know its  
(**object\_id**, **size [int]**, alive [bool])

At the same time, for each memory access, we know:

- Memory read: (**object\_id**, **offset [int]**, **length [int]**)
- Memory write: (**object\_id**, **offset [int]**, **length [int]**, **\_**)

It is a violation of spatial safety if:

- $\text{offset} + \text{length} \geq \text{size}$  or
- $\text{offset} < 0$

# Example: spatial safety violations

```
1 int foo(int x) {  
2     int arr[16] = {0};  
3     return arr[x];  
4 }
```



# Example: spatial safety violations

```
1 int foo(int x) {  
2     int arr[16] = {0};  
3     return arr[x];  
4 }
```

```
1 long foo() {  
2     int a = 0;  
3     return *(long *)(&a);  
4 }
```

# Definition: NULL-pointer dereference

```
1 int foo(int *p) {  
2     // it is possible that p == NULL  
3     return *p + 42;  
4 }
```

## Definition: NULL-pointer dereference

```
1 int foo(int *p) {  
2     // it is possible that p == NULL  
3     return *p + 42;  
4 }
```

NULL-pointer dereference is sometimes considered as **undefined behavior** — meaning, its behavior is not given in the C language specification, although most operating systems chooses to panic the program on such behavior.

## Definition: NULL-pointer dereference

At any point of time during the program execution,  
for any object in memory, we know its  
(**object\_id**  $\neq$  0, size [int], alive [bool])

At the same time, for each memory access, we know:

- Memory read: (**object\_id**, offset [int], length [int])
- Memory write: (**object\_id**, offset [int], length [int], \_)

## Definition: NULL-pointer dereference

At any point of time during the program execution,  
for any object in memory, we know its  
(**object\_id**  $\neq$  0, size [int], alive [bool])

At the same time, for each memory access, we know:

- Memory read: (**object\_id**, offset [int], length [int])
- Memory write: (**object\_id**, offset [int], length [int], \_)

It is a NULL-pointer dereference if

- `object_id == 0`

# Outline

- 1 Introduction
- 2 Intuition
- 3 Spatial safety
- 4 Temporal safety**
- 5 Countermeasures

## Definition: temporal safety

At any point of time during the program execution,  
for any object in memory, we know its  
(**object\_id**, size [int], **alive** [bool])

At the same time, for each memory access, we know:

- Memory read: (**object\_id**, offset [int], length [int])
- Memory write: (**object\_id**, offset [int], length [int], \_)
- Memory free: (**object\_id**)

## Definition: temporal safety

At any point of time during the program execution,  
for any object in memory, we know its  
(**object\_id**, size [int], **alive** [bool])

At the same time, for each memory access, we know:

- Memory read: (**object\_id**, offset [int], length [int])
- Memory write: (**object\_id**, offset [int], length [int], \_)
- Memory free: (**object\_id**)

It is a violation of temporal safety if:

- !alive



# Example: temporal safety violations

```
1 int foo() {  
2     int *p = malloc(sizeof(int));  
3     *p = 42;  
4     free(p);  
5     return *p;  
6 }
```

# Example: temporal safety violations

```
1 int foo() {
2     int *p = malloc(sizeof(int));
3     *p = 42;
4     free(p);
5     return *p;
6 }
```

```
1 int *ptr;
2
3 void foo() {
4     int p = 100;
5     ptr = &p;
6 }
7 int bar() {
8     return *ptr;
9 }
```

# Example: temporal safety violations

```
1 int foo() {
2     int *p = malloc(sizeof(int));
3     *p = 42;
4     free(p);
5     return *p;
6 }
```

```
1 int *ptr;
2
3 void foo() {
4     int p = 100;
5     ptr = &p;
6 }
7 int bar() {
8     return *ptr;
9 }
```

```
1 int foo() {
2     int *p = malloc(sizeof(int));
3     *p = 42;
4     free(p);
5     free(p);
6     return *p;
7 }
```

## Definition: temporal safety (revisited)

At any point of time during the program execution,  
for any object in memory, we know its  
(**object\_id**, size [int], **status** [alloc|init|dead])

At the same time, for each memory access, we know:

- Memory read: (**object\_id**, offset [int], length [int])
- Memory write: (**object\_id**, offset [int], length [int], \_)
- Memory free: (**object\_id**)

## Definition: temporal safety (revisited)

At any point of time during the program execution,  
for any object in memory, we know its  
(**object\_id**, size [int], **status** [alloc|init|dead])

At the same time, for each memory access, we know:

- Memory read: (**object\_id**, offset [int], length [int])
- Memory write: (**object\_id**, offset [int], length [int], \_)
- Memory free: (**object\_id**)

It is a violation of temporal safety if:

- Read: **status** != **init**
- Write: **status** == **dead**
- Free: **status** == **dead**

# Example: temporal safety violations

```
1 int foo() {  
2   int p;  
3   return p;  
4   // what is the value returned?  
5 }
```

# Example: temporal safety violations

```
1 int foo() {
2   int p;
3   return p;
4   // what is the value returned?
5 }
```

```
1 int foo() {
2   int *p = malloc(sizeof(int));
3   return *p;
4   // what is the value returned?
5 }
```

## Definition: memory leak

At any point of time during the program execution,  
for any object in memory, we know its  
(**object\_id**, size [int], **status** [alloc|init|dead])

At the same time, for each memory access, we know:

- Memory read: (**object\_id**, offset [int], length [int])
- Memory write: (**object\_id**, offset [int], length [int], \_)
- Memory free: (**object\_id**)



## Definition: memory leak

At any point of time during the program execution,  
for any object in memory, we know its  
(**object\_id**, size [int], **status** [alloc|init|dead])

At the same time, for each memory access, we know:

- Memory read: (**object\_id**, offset [int], length [int])
- Memory write: (**object\_id**, offset [int], length [int], \_)
- Memory free: (**object\_id**)

It is a memory leak if **exists one object\_id** whose:

- status != dead

# Example: memory leak

```
1 int foo() {  
2     int *p = malloc(sizeof(int));  
3     int *q = malloc(sizeof(int));  
4     *p = 42;  
5     free(q);  
6     return *p;  
7 }
```

# Outline

- 1 Introduction
- 2 Intuition
- 3 Spatial safety
- 4 Temporal safety
- 5 Countermeasures**

# Detecting memory errors

- Static analysis
- Dynamic analysis

# What is so hard about static analysis?

# What is so hard about static analysis?

```
1 void foo(bool cond) {
2     char *p = malloc(sizeof(char) * 16);
3     char[4] q;
4
5     char *r;
6     if (cond) {
7         r = p;
8     } else {
9         r = q;
10    }
11    memcpy(r, "HELLO", 6);
12
13    free(p);
14 }
```

It is possible that one pointer may points to multiple locations.

# What is so hard about static analysis?

```
1 struct S {
2     char *field;
3 }
4
5 void foo() {
6     char *p = malloc(sizeof(char) * 16);
7     struct *s = malloc(sizeof(struct S));
8     s->field = p;
9     free(s);
10    free(p);
11 }
12
13 void bar(struct S* s) {
14     free(s->field);
15 }
```

It is possible that one location may have two aliased pointers.

## Preventing memory errors

- (Selective) hardening
- Use a safer language (Java / Rust / Modern C++)



# Why not harden everything?

There is actually a very simple way of preventing memory errors **completely!**

At any point of time during the program execution,  
for any object in memory, we track its  
(**object\_id**, size [int], status [alloc|init|dead])

At the same time, for each memory access, we check:

- Memory read: (**object\_id**, offset [int], length [int])
- Memory write: (**object\_id**, offset [int], length [int], \_)

# Why not harden everything?

There is actually a very simple way of preventing memory errors **completely!**

At any point of time during the program execution,  
**for any object** in memory, we **track** its  
(**object\_id**, size [int], status [alloc|init|dead])

At the same time, **for each memory access**, we **check**:

- Memory read: (**object\_id**, offset [int], length [int])
- Memory write: (**object\_id**, offset [int], length [int], \_)

This is essentially what is implemented in AddressSanitizer and MemorySanitizer. The result? **Over 100% performance overhead...**

# Zero-cost abstraction

# Zero-cost abstraction

```
1 int foo(int *arr, size_t len) {  
2     int sum = 0;  
3     for(size_t i = 0; i < len; i++) {  
4         // memory access check at each access  
5         sum += arr[i];  
6     }  
7     return sum;  
8 }
```

# Zero-cost abstraction

```
1 int foo(int *arr, size_t len) {
2     int sum = 0;
3     for(size_t i = 0; i < len; i++) {
4         // memory access check at each access
5         sum += arr[i];
6     }
7     return sum;
8 }
```

```
1 fn foo(arr: &Vec<i32>) -> i32 {
2     let mut sum = 0;
3     arr.iter().map(
4         // no need to check memory access here
5         |e| sum += e
6     );
7     return sum;
8 }
```

〈 **End** 〉