

## Machine-Level Programming V: Advanced Topics

CSci 2021: Machine Architecture and Organization  
March 6th, 2020  
**Your instructor:** Stephen McCamant

Based on slides originally by:  
Randy Bryant, Dave O'Hallaron

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

1

# Today

- Memory Layout
  - Unions
  - Buffer Overflow
    - Vulnerability
    - Protection

2

## x86-64 Linux Memory Layout

- **Stack**
    - Runtime stack (default 8MB soft limit)
    - E. g., local variables
  - **Heap**
    - Dynamically allocated as needed
    - When you call `malloc()`, `calloc()`, C++ `new`
  - **Data**
    - Statically (compiler-)allocated data
    - E.g., global vars, **static** vars, string constants
  - **Text / Shared Libraries**
    - Executable machine instructions
    - Read-only

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

**Hex Address** 400000  
000000

The diagram illustrates the memory layout of a process. It consists of several sections stacked vertically:

- Stack**: The top section, colored light blue, contains a downward-pointing arrow.
- Shared Libraries**: A green section below the stack.
- Heap**: A green section below the shared libraries.
- Data**: A pink section below the heap.
- Text**: The bottom section, colored yellow.

A brace on the right side of the diagram spans from the bottom of the Stack section to the top of the Text section, labeled "8M".

3

## Memory Allocation Example

```

char big_array[1L<<24]; /* 16 MB */
char huge_array[1L<<31]; /* 2 GB */

int global = 0;

int useless() { return 0; }

int main ()
{
    void *p1, *p2, *p3, *p4;
    int local = 0;
    p1 = malloc(1L << 28); /* 256 MB */
    p2 = malloc(1L << 8); /* 256 B */
    p3 = malloc(1L << 32); /* 4 GB */
    p4 = malloc(1L << 8); /* 256 B */
/* Some print statements ... */
}

```

*Where does everything go?*

*not drawn to scale*

The diagram illustrates the memory layout of a process. It consists of several horizontal sections separated by thin lines. From top to bottom, the sections are labeled: **Stack** (purple), **Shared Libraries** (green), **Heap** (light blue), **Data** (pink), and **Text** (yellow). A black arrow points downwards from the **Stack** section towards the **Shared Libraries** section. Another black arrow points upwards from the **Text** section towards the **Heap** section.

4

## x86-64 Example Addresses

*address range*  $\sim 2^{47}$

```
local          0x000007ffe4d3be87c  
p1           0x000007f7262a1e010  
p3           0x000007f7162a1d010  
p4           0x0000000008359d120  
p2           0x0000000008359d010  
  
big_array    0x00000000080601060  
huge_array   0x00000000000601060  
main()        0x000000000040060c  
useless()     0x0000000000400590
```

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

The diagram shows a vertical stack of memory regions:

- Stack**: Top-most region, light purple.
- Heap**: Second region from top, light green.
- Heap**: Third region from top, light green.
- Data**: Fourth region from top, pink.
- Text**: Bottom-most region, yellow.

Three arrows originate from the bottom of the second **Heap** section and point to the top of the third **Heap** section, indicating a recursive or circular reference between them.

5

Today

- Memory Layout
  - **Unions**
  - Buffer Overflow
    - Vulnerability
    - Protection

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

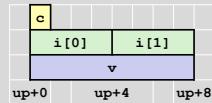
6

## Union Allocation

- Allocate according to largest element
- Can only use one field at a time

```
union U1 {
    char c;
    int i[2];
    double v;
} *up;
```

```
struct S1 {
    char c;
    int i[2];
    double v;
} *sp;
```

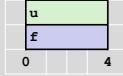


Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

7

## Using Union to Access Bit Patterns

```
typedef union {
    float f;
    unsigned u;
} bit_float_t;
```



```
float bit2float(unsigned u)
{
    bit_float_t arg;
    arg.u = u;
    return arg.f;
}
```

Same as `(float) u`?

```
unsigned float2bit(float f)
{
    bit_float_t arg;
    arg.f = f;
    return arg.u;
}
```

Same as `(unsigned) f`?

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

8

## Byte Ordering Revisited

### Idea

- Short/long/quad words stored in memory as 2/4/8 consecutive bytes
- Which byte is most (least) significant?
- Can cause problems when exchanging binary data between machines

### Big Endian

- Most significant byte has lowest address
- Sparc

### LittleEndian

- Least significant byte has lowest address
- Intel x86, ARM Android and iOS

### Bi Endian

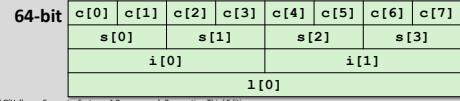
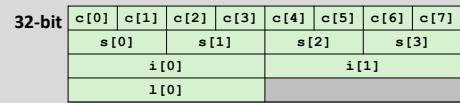
- Can be configured either way
- ARM

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

9

## Byte Ordering Example

```
union {
    unsigned char c[8];
    unsigned short s[4];
    unsigned int i[2];
    unsigned long l[1];
} dw;
```



Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

10

## Byte Ordering Example (Cont.).

```
int j;
for (j = 0; j < 8; j++)
    dw.c[j] = 0xf0 + j;

printf("Characters 0-7 == "
    "[0x%x,0x%x,0x%x,0x%x,0x%x,0x%x,0x%x]\n",
    dw.c[0], dw.c[1], dw.c[2], dw.c[3],
    dw.c[4], dw.c[5], dw.c[6], dw.c[7]);

printf("Shorts 0-3 == [0x%x,0x%x,0x%x,0x%x]\n",
    dw.s[0], dw.s[1], dw.s[2], dw.s[3]);

printf("Ints 0-1 == [0x%lx,0x%lx]\n",
    dw.i[0], dw.i[1]);

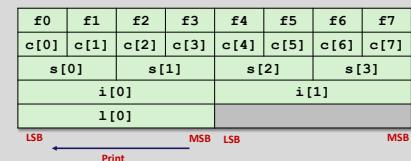
printf("Long 0 == [0x%lx]\n",
    dw.l[0]);
```

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

11

## Byte Ordering on IA32

### Little Endian



### Output:

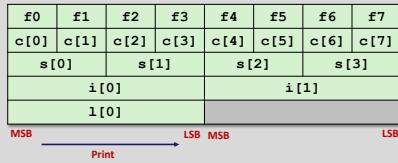
```
Characters 0-7 == [0xf0,0xf1,0xf2,0xf3,0xf4,0xf5,0xf6,0xf7]
Shorts 0-3 == [0xf1f0,0xf3f2,0xf5f4,0xf7f6]
Ints 0-1 == [0xf3f2f1f0,0xf7f6f5f4]
Long 0 == [0xf3f2f1f0]
```

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

12

## Byte Ordering on Sun

### Big Endian



### Output on Sun:

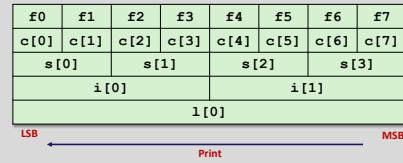
```
Characters 0-7 == [0xf0, 0xf1, 0xf2, 0xf3, 0xf4, 0xf5, 0xf6, 0xf7]
Shorts    0-3 == [0x0f10, 0xf2f3, 0xf4f5, 0xf6f7]
Ints      0-1 == [0x0f0f1f2f3, 0xf4f5f6f7]
Long       0 == [0xf0f1f2f3]
```

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

13

## Byte Ordering on x86-64

### Little Endian



### Output on x86-64:

```
Characters 0-7 == [0xf0, 0xf1, 0xf2, 0xf3, 0xf4, 0xf5, 0xf6, 0xf7]
Shorts    0-3 == [0x1f0, 0xf2f3, 0xf4f5, 0xf6f7]
Ints      0-1 == [0xf3f2f1f0, 0xe7f6f5f4]
Long       0 == [0xe7f6f5f4f3f2f1f0]
```

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

14

## Summary of Compound Types in C

### Arrays

- Contiguous allocation of memory
- Aligned to satisfy every element's alignment requirement
- Pointer to first element
- No bounds checking

### Structures

- Allocate bytes in order declared
- Pad in middle and at end to satisfy alignment

### Unions

- Overlay declarations
- Way to circumvent type system

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

15

## Today

### Memory Layout

### Unions

### Buffer Overflow

- Vulnerability
- Protection

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

16

## Recall: Memory Referencing Bug Example

```
typedef struct {
    int a[2];
    double d;
} struct_t;

double fun(int i) {
    volatile struct_t s;
    s.d = 3.14;
    s.a[i] = 1073741824; /* Possibly out of bounds */
    return s.d;
}
```

```
fun(0) → 3.14
fun(1) → 3.14
fun(2) → 3.1399998664856
fun(3) → 2.00000061035156
fun(4) → 3.14
fun(6) → Segmentation fault
```

- Result is system specific

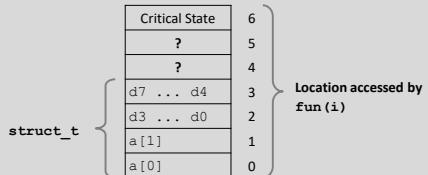
Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

17

## Memory Referencing Bug Example

typedef struct {	fun(0) → 3.14
int a[2];	fun(1) → 3.14
double d;	fun(2) → 3.1399998664856
struct_t;	fun(4) → 3.14
	fun(6) → Segmentation fault

### Explanation:



Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

18

## Such problems are a BIG deal

- Generally called a “buffer overflow”
  - when exceeding the memory size allocated for an array
- Why a big deal?
  - One of the most common technical causes of security vulnerabilities
- Most common form
  - Unchecked lengths on string inputs
  - Particularly for bounded character arrays on the stack
    - sometimes referred to as stack smashing

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

19

## String Library Code

### ■ Implementation of old standard C function gets()

```
/* Get string from stdin */
char *gets(char *dest)
{
    int c = getchar();
    char *p = dest;
    while (c != EOF && c != '\n') {
        *p++ = c;
        c = getchar();
    }
    *p = '\0';
    return dest;
}
```

▪ Bad design: no way to specify limit on number of characters to read

### ■ Similar problems with other library functions

- `strcpy`, `strcat`: Copy strings of arbitrary length
- `scanf`, `fscanf`, `sscanf`, when given `%s` conversion specification

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

20

## Vulnerable Buffer Code

```
/* Echo Line */
void echo()
{
    char buf[4]; /* Way too small! */
    gets(buf);
    puts(buf);
}

void call_echo() {
    echo();
}
```

**btw, how big  
is big enough?**

```
unix> ./bufdemo-nsp
Type a string: 012345678901234567890123
012345678901234567890123

unix> ./bufdemo-nsp
Type a string: 0123456789012345678901234
Segmentation Fault
```

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

21

## Buffer Overflow Disassembly

### echo:

```
00000000004006cf <echo>:
4006cf: 48 83 ec 18           sub    $0x18,%rsp
4006d3: 48 89 e7             mov    %rsp,%rdi
4006d6: e8 a5 ff ff ff     callq  400680 <gets>
4006db: 48 89 e7             mov    %rsp,%rdi
4006de: e8 3d fe ff ff     callq  400520 <puts@plt>
4006e3: 48 83 c4 18           add    $0x18,%rsp
4006e7: c3                   retq
```

### call\_echo:

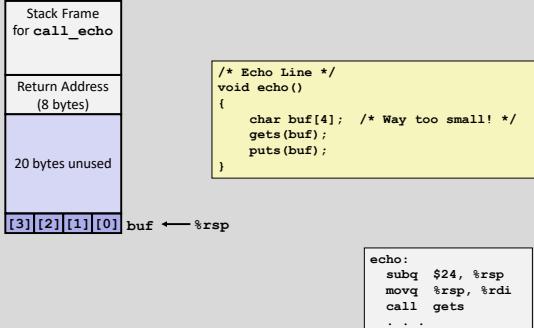
```
4006e8: 48 83 ec 08           sub    $0x8,%rsp
4006ec: b8 00 00 00 00         mov    $0x0,%eax
4006f1: e8 d9 ff ff ff     callq  4006cf <echo>
4006f6: 48 83 c4 08           add    $0x8,%rsp
4006fa: c3                   retq
```

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

22

## Buffer Overflow Stack

### Before call to gets



Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

23

## Buffer Overflow Stack Example

### Before call to gets

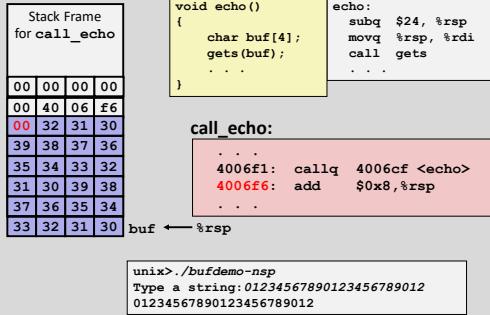
<pre>Stack Frame for call_echo</pre> <pre>00 00 00 00 00 40 06 f6</pre> <p>20 bytes unused</p> <pre>[3][2][1][0] buf ← %rsp</pre>	<pre>echo:</pre> <pre>subq \$24, %rsp movq %rsp, %rdi call gets ...</pre>
---	---

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

24

## Buffer Overflow Stack Example #1

After call to gets



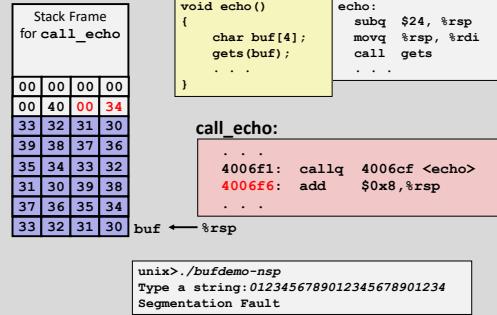
Overflowed buffer, but did not corrupt state

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

25

## Buffer Overflow Stack Example #2

After call to gets



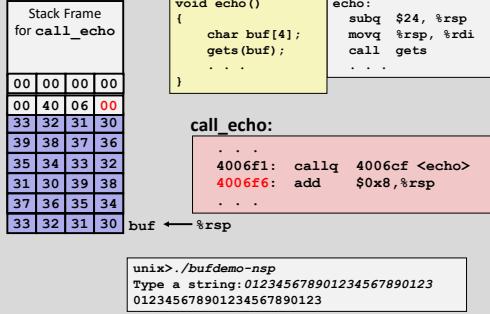
Overflowed buffer and corrupted return pointer

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

26

## Buffer Overflow Stack Example #3

After call to gets



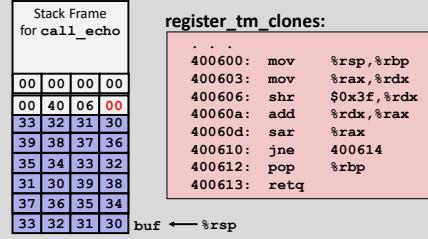
Overflowed buffer, corrupted return pointer, but program seems to work!

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

27

## Buffer Overflow Stack Example #3 Explained

After call to gets

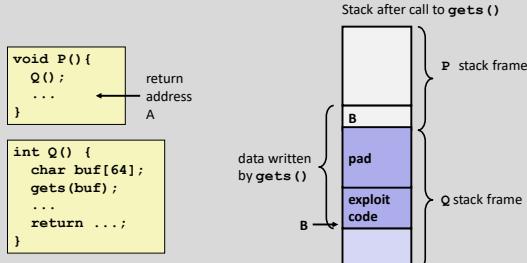


"Returns" to unrelated code  
Lots of things happen, without modifying critical state  
Eventually executes retq back to main

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

28

## Code Injection Attacks



- Input string contains byte representation of executable code
- Overwrite return address A with address of buffer B
- When Q executes ret, will jump to exploit code

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

29

## Exploits Based on Buffer Overflows

- **Buffer overflow bugs can allow remote machines to execute arbitrary code on victim machines**

- **Distressingly common in real programs**

- Programmers keep making the same mistakes ☹
- Recent measures make these attacks much more difficult

- **Examples across the decades**

- Original "Internet worm" (1988)
- "IM wars" (1999)
- Twilight hack on Wii (2000s)
- ... and many, many more

- **You will try out some techniques in lab**

- Hopefully to convince you to never leave such holes in your programs!!

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

30

## Example: the original Internet worm (1988)

- Exploited a few vulnerabilities to spread
  - Early versions of the finger server (fingerd) used `gets()` to read the argument sent by the client:
    - `finger drob@cs.cmu.edu`
  - Worm attacked fingerd server by sending phony argument:
    - `finger "exploit-code padding new-return-address"`
  - exploit code: executed a root shell on the victim machine with a direct TCP connection to the attacker.
- Once on a machine, scanned for other machines to attack
  - invaded ~6000 computers in hours (10% of the Internet ☺)
    - see June 1989 article in *Comm. of the ACM*
  - the young author of the worm was prosecuted
  - and CERT was formed

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

31

## Discussion Break: Unknown Addresses?

- Basic attack requires attacker to know address B of buffer
- Is an attack still possible if B is variable?
- E.g. what if attacker only knows B +/- 30?

35

## Discussion Break: Unknown Addresses?

- Basic attack requires attacker to know address B of buffer
- Is an attack still possible if B is variable?
- E.g. what if attacker only knows B +/- 30?

### Some possible attack strategies:

- Try attack repeatedly
- “NOP sled”: (0x90 is one-byte no-operation in x86)

NOP Exploit Code

36

## OK, what to do about buffer overflow attacks

- Avoid overflow vulnerabilities
- Employ system-level protections
- Have compiler use “stack canaries”
- Lets talk about each...

41

## 1. Avoid Overflow Vulnerabilities in Code (!)

```
/* Echo Line */
void echo()
{
    char buf[4]; /* Way too small! */
    fgets(buf, 4, stdin);
    puts(buf);
}
```

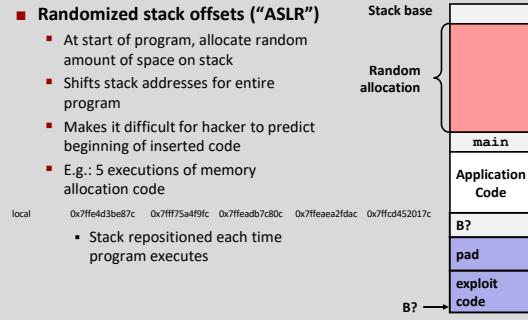
- For example, use library routines that limit string lengths
  - `fgets` instead of `gets`
  - `strncpy` instead of `strcpy`
  - Don't use `scanf` with `%s` conversion specification
    - Use `fgets` to read the string
    - Or use `%ns` where `n` is a suitable integer

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

42

## 2. System-Level Protections can help

- Randomized stack offsets (“ASLR”)
  - At start of program, allocate random amount of space on stack
  - Shifts stack addresses for entire program
  - Makes it difficult for hacker to predict beginning of inserted code
  - E.g.: 5 executions of memory allocation code
    - Stack repositioned each time program executes

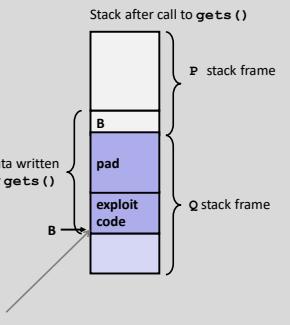


43

## 2. System-Level Protections can help

### ■ Nonexecutable code segments

- In traditional x86, can mark region of memory as either “read-only” or “writeable”
  - Can execute anything readable
- X86-64 era CPUs added explicit “(non-)execute” permission
- Stack marked as non-executable



Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

44

## 3. Stack Canaries can help

### ■ Idea

- Place special value (“canary”) on stack just beyond buffer
- Check for corruption before exiting function

### ■ GCC Implementation

- `-fstack-protector`
- Now commonly enabled by default

```
unix>./bufdemo -sp
Type a string: 0123456
0123456

unix>./bufdemo -sp
Type a string: 01234567
*** stack smashing detected ***
```

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

45

## Protected Buffer Disassembly

### echo:

```
40072f: sub    $0x18,%rsp
400733: mov    %fs:0x28,%rax
40073c: mov    %rax,0xb(%rsp)
400741: xor    %eax,%eax
400743: mov    %rsp,%rdi
400746: callq  4006e0 <gets>
40074b: mov    %rsp,%rdi
40074e: callq  400570 <puts@plt>
400753: mov    0x8(%rsp),%rax
400758: xor    %fs:0x28,%rax
400761: je    400768 <echo+0x39>
400763: callq  400580 <_stack_chk_fail@plt>
400768: add    $0x18,%rsp
40076c: retq
```

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

46

## Setting Up Canary

### Before call to gets

```
/* Echo Line */
void echo()
{
    char buf[4]; /* Way too small! */
    gets(buf);
    puts(buf);
}
```



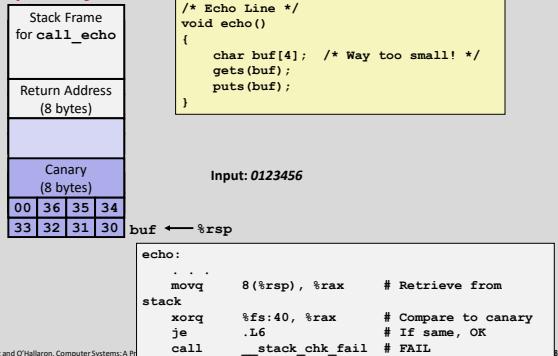
```
echo:
    ...
    movq    %fs:40, %rax # Get canary
    movq    %rax, 8(%rsp) # Place on stack
    xorl    %eax, %eax # Erase canary
    ...
```

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

47

## Checking Canary

### After call to gets



Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

## Return-Oriented Programming Attacks

### ■ Challenge (for hackers)

- Stack randomization makes it hard to predict buffer location
- Marking stack nonexecutable makes it hard to insert binary code

### ■ Alternative Strategy

- Use existing code
  - E.g., library code from stdlib
- String together fragments to achieve overall desired outcome
- Does not on its own overcome stack canaries

### ■ Construct program from gadgets

- Sequence of instructions ending in `ret`
  - Encoded by single byte `0xc3`
- Code positions fixed from run to run
- Code is executable

Bryant and O'Hallaron, Computer Systems: A Programmer's Perspective, Third Edition

48

## Gadget Example #1

```
long ab_plus_c
(long a, long b, long c)
{
    return a*b + c;
}
```

```
00000000004004d0 <ab_plus_c>:
4004d0: 48 0f af fe imul %rsi,%rdi
4004d4: 48 8d 04 17 lea (%rdi,%rdx,1),%rax
4004d8: c3 retq
```

$\text{rax} \leftarrow \text{rdi} + \text{rdx}$   
Gadget address = 0x4004d4

- Use tail end of existing functions

## Gadget Example #2

```
void setval(unsigned *p) {
    *p = 3347663060u;
}
```

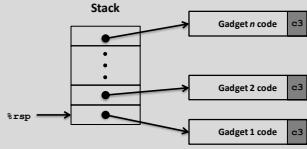
```
<setval>:
4004d9: c7 07 d4 48 89 c7 movl $0xc78948d4,(%rdi)
4004df: c3 retq
```

Encodes `movq %rax, %rdi`

$\text{rdi} \leftarrow \text{rax}$   
Gadget address = 0x4004dc

- Repurpose instruction bytes

## ROP Execution



- Trigger with `ret` instruction
  - Will start executing Gadget 1
  - Final `ret` in each gadget will start next one