

18-600 Foundations of Computer Systems

Lecture 7: "Machine-Level Programming III: Data & Program"

John Shen & Zhiyi Yu
September 21, 2016

- Required Reading Assignment:
 - Chapter 3 of CS:APP (3rd edition) by Randy Bryant & Dave O'Hallaron
- Assignments for This Week:
 - ❖ Lab 2



Today

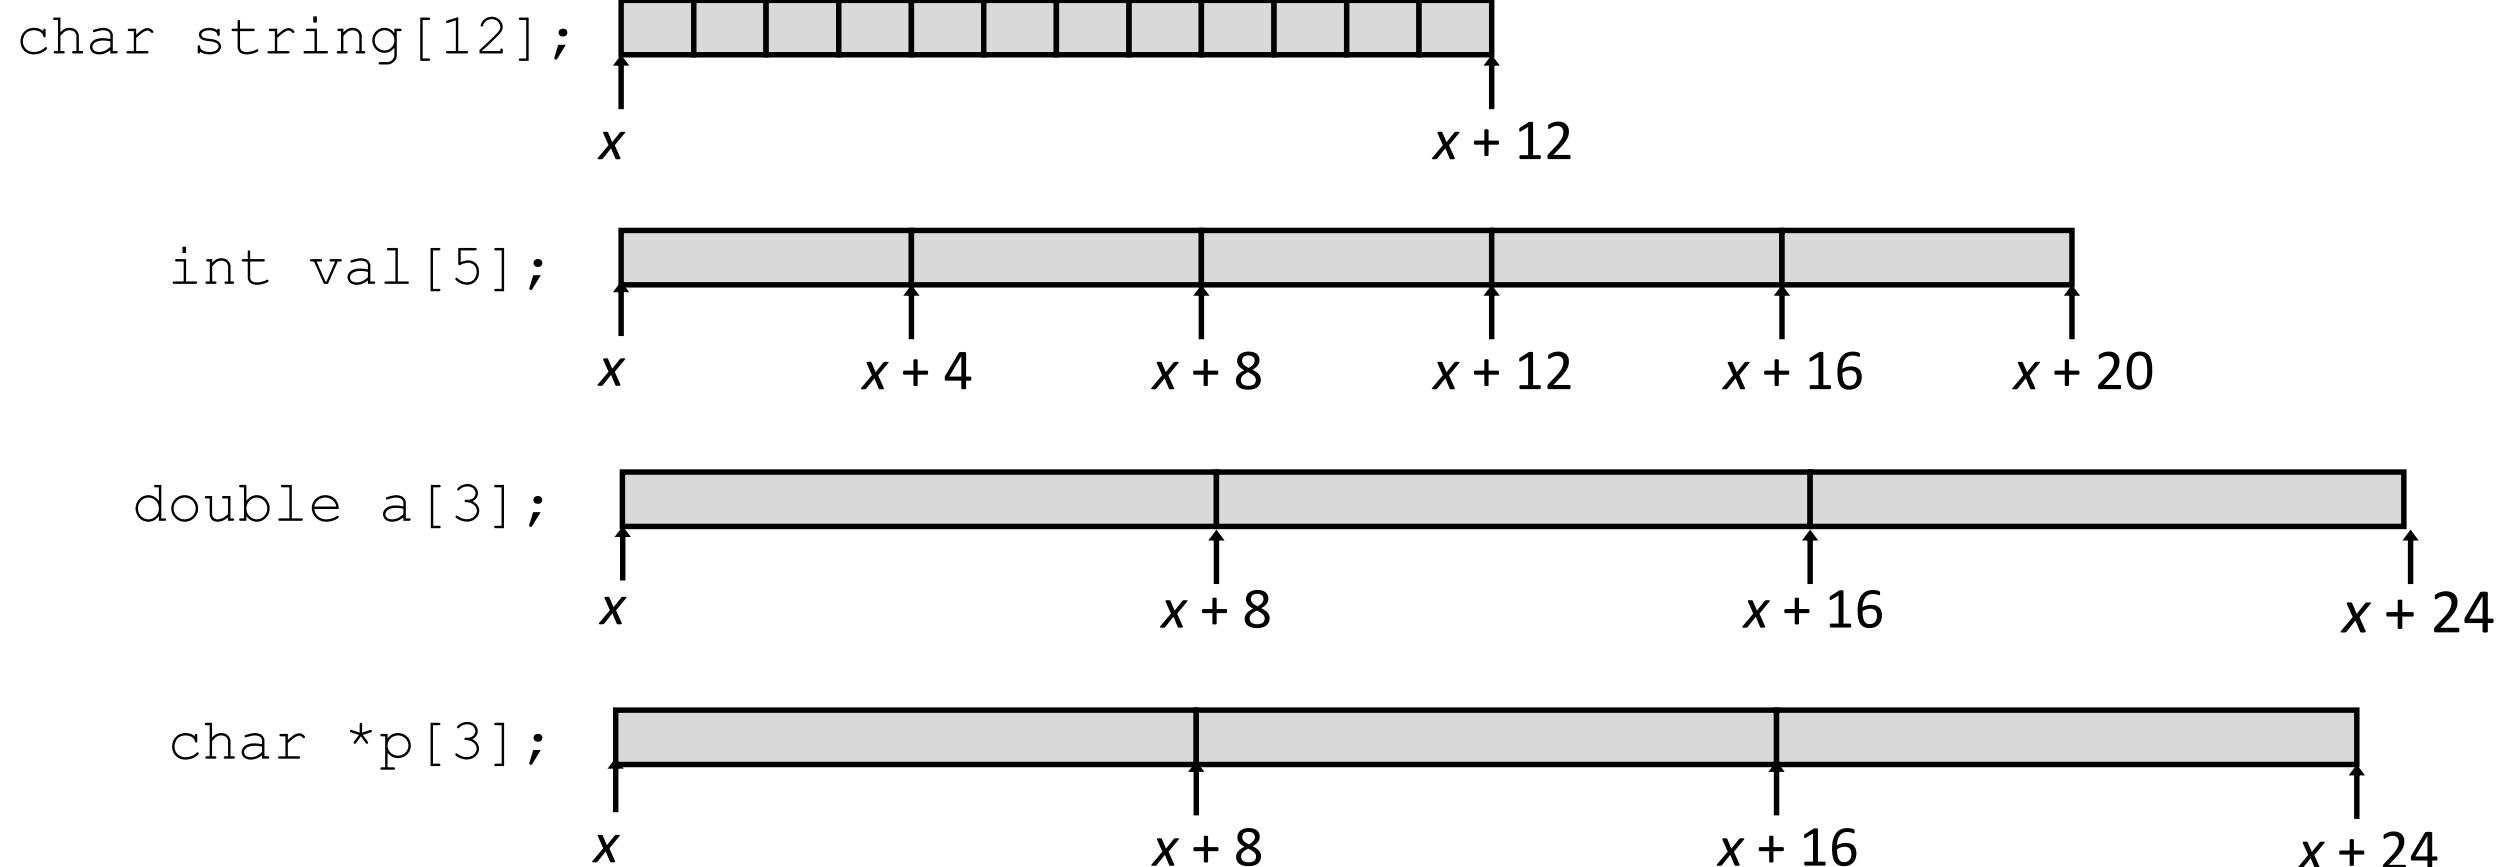
- Data
 - Arrays
 - Structures
 - Unions
 - Floating Point and SIMD operations
- Buffer Overflow
 - Memory Layout
 - Vulnerability and Protection

Array Allocation

- Basic Principle

$T \mathbf{A}[L];$

- Array of data type T and length L
- Contiguously allocated region of $L * \text{sizeof}(T)$ bytes in memory

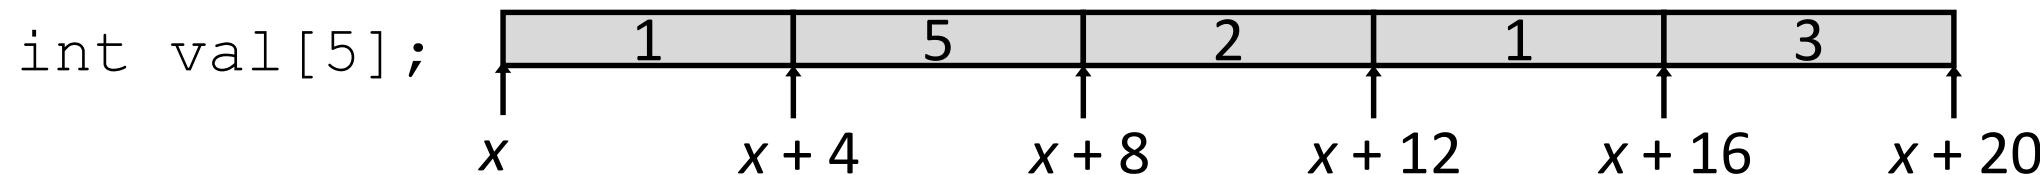


Array Access

- Basic Principle

$T \mathbf{A}[L]$;

- Array of data type T and length L
- Identifier **A** can be used as a pointer to array element 0: Type T^*



• Reference	Type	Value
<code>val[4]</code>	<code>int</code>	3
<code>val</code>	<code>int *</code>	x
<code>val+1</code>	<code>int *</code>	$x + 4$
<code>&val[2]</code>	<code>int *</code>	$x + 8$
<code>val[5]</code>	<code>int</code>	??
<code>* (val+1)</code>	<code>int</code>	5
<code>val + i</code>	<code>int *</code>	$x + 4i$

Multidimensional (Nested) Arrays

- Declaration

$T \text{ } \mathbf{A}[R][C];$

- 2D array of data type T
- R rows, C columns
- Type T element requires K bytes

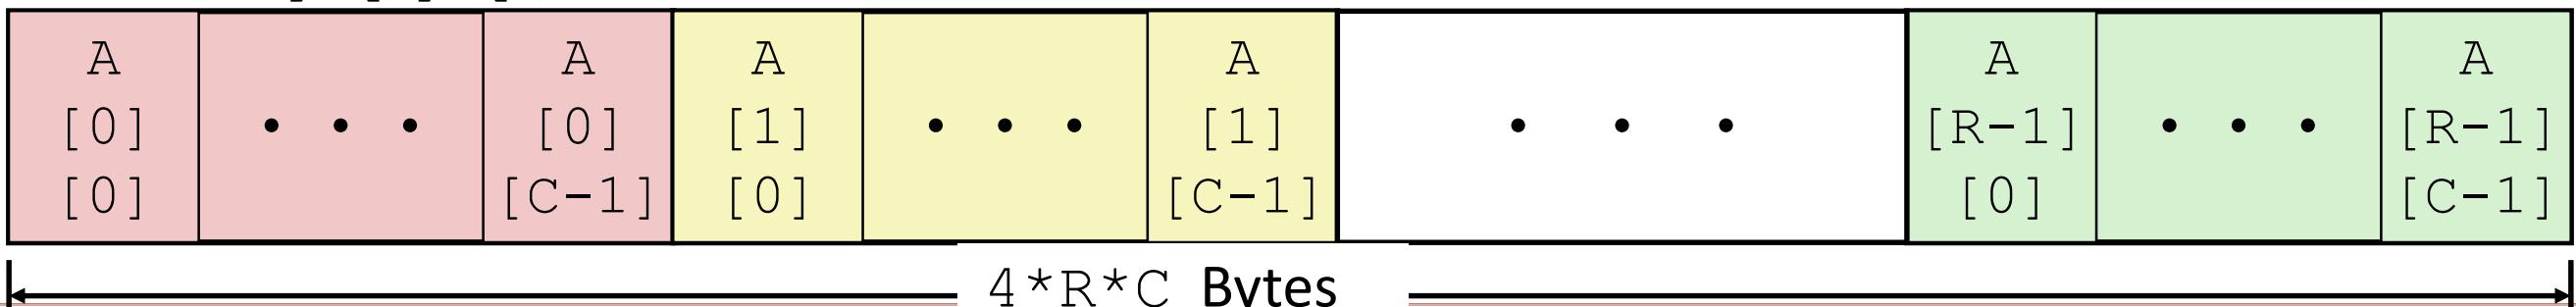
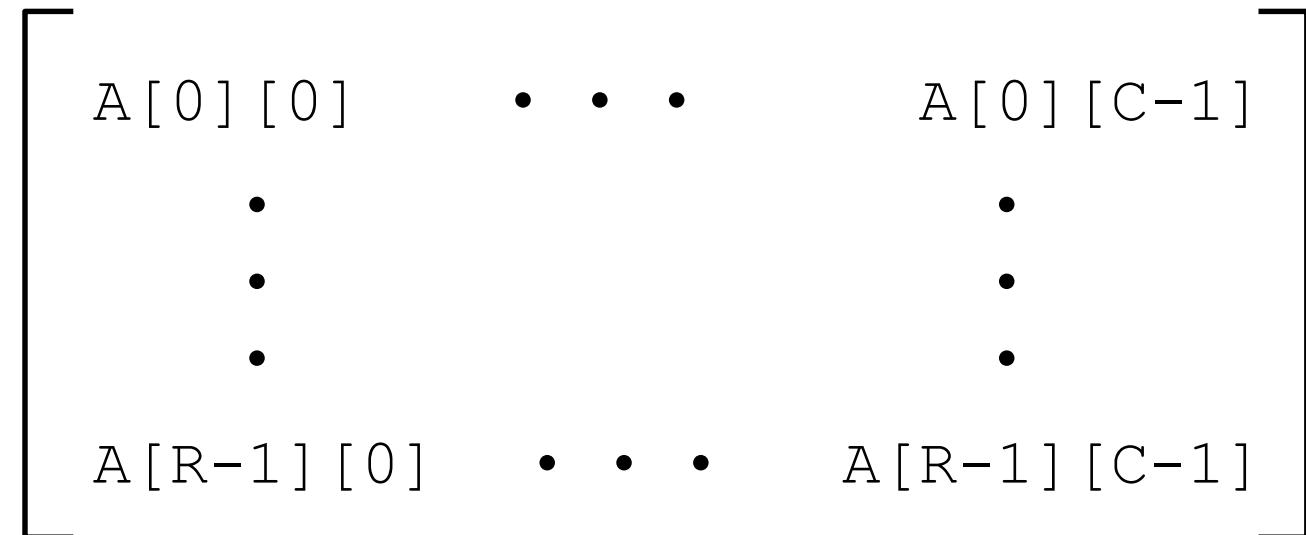
- Array Size

- $R * C * K$ bytes

- Arrangement

- Row-Major Ordering

```
int A[R][C];
```



Nested Array Access

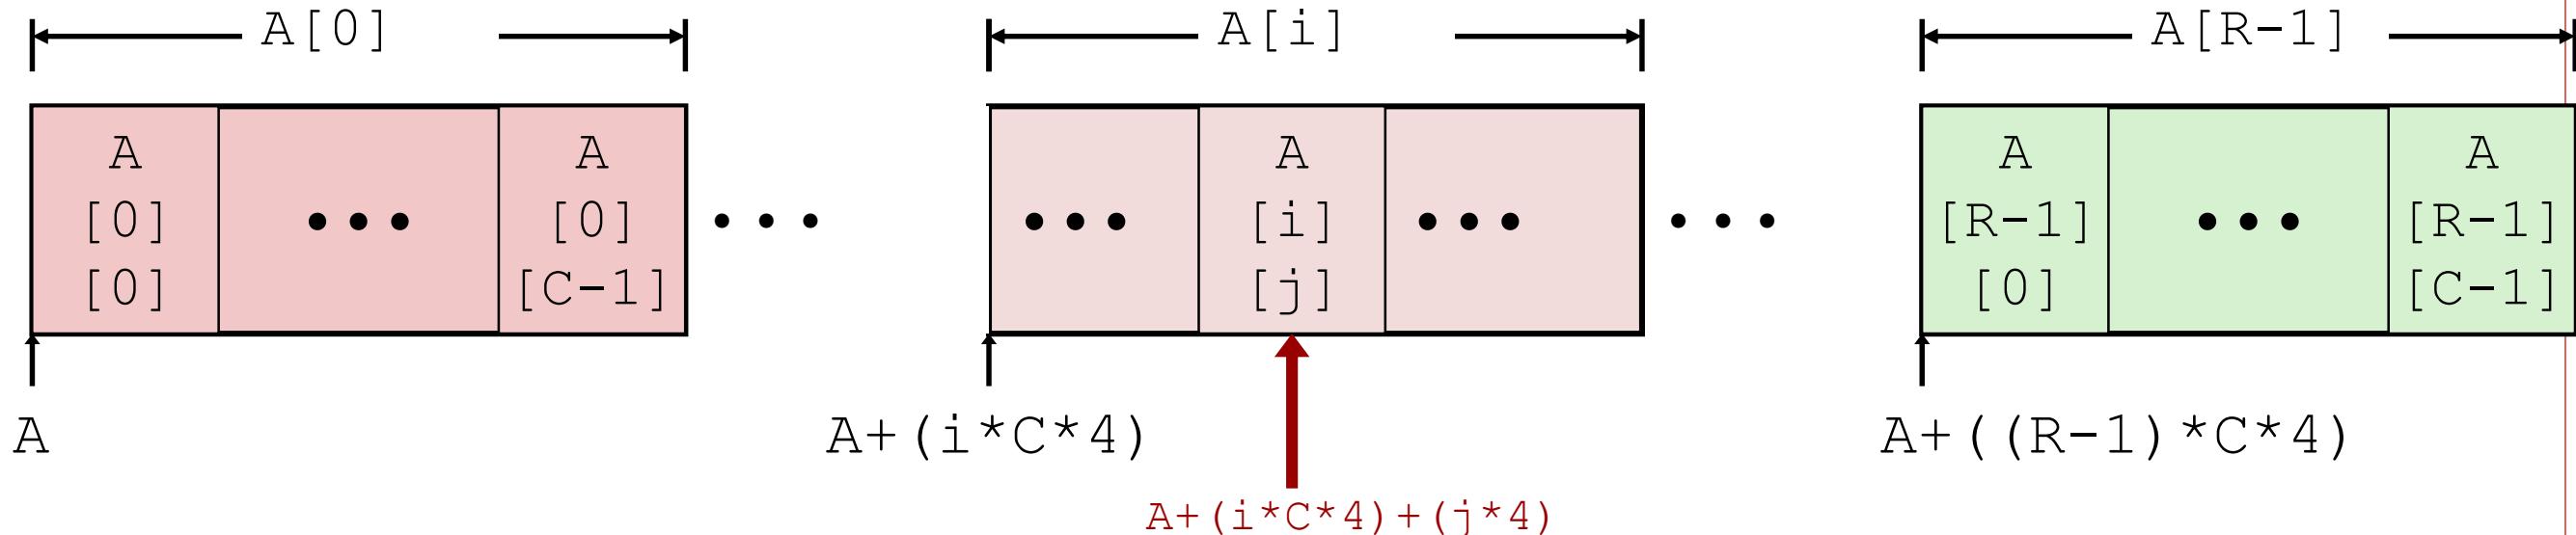
Row Vectors

- $\mathbf{A}[i]$ is array of C elements
- Each element of type T requires K bytes
- Starting address $\mathbf{A} + i * (C * K)$

• Array Elements

- $\mathbf{A}[i][j]$ is element of type T , which requires K bytes
- Address $\mathbf{A} + i * (C * K) + j * K = \mathbf{A} + (i * C + j) * K$

```
int A[R][C];
```



16 X 16 Matrix Access

■ Array Elements

- Address $\mathbf{A} + i * (C * K) + j * K$
- $C = 16, K = 4$

```
/* Get element a[i][j] */
int fix_ele(fix_matrix a, size_t i, size_t j) {
    return a[i][j];
}

# a in %rdi, i in %rsi, j in %rdx
salq    $6, %rsi           # 64*i
addq    %rsi, %rdi         # a + 64*i
movl    (%rdi,%rdx,4), %eax # M[a + 64*i + 4*j]
ret
```

n X n Matrix Access

■ Array Elements

- Address $\mathbf{A} + i * (C * K) + j * K$
- C = n, K = 4
- Must perform integer multiplication

```
/* Get element a[i][j] */
int var_ele(size_t n, int a[n][n], size_t i, size_t j) {
    return a[i][j];
}
```

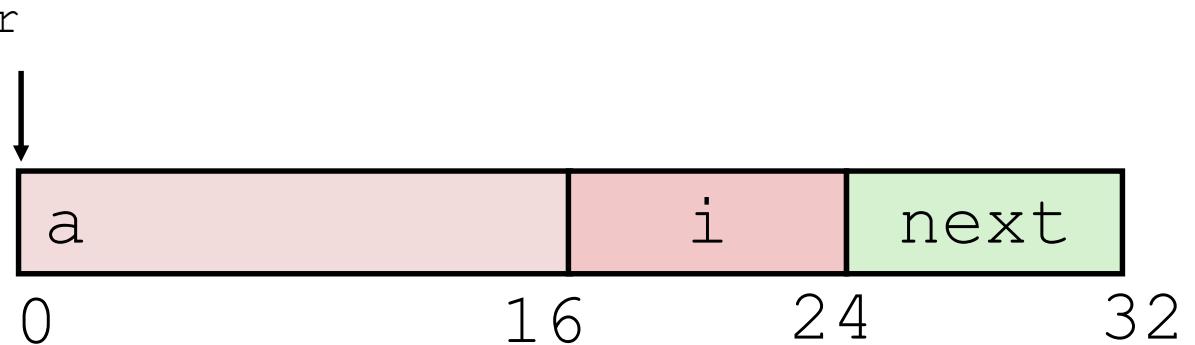
```
# n in %rdi, a in %rsi, i in %rdx, j in %rcx
imulq    %rdx, %rdi          # n*i
leaq     (%rsi,%rdi,4), %rax # a + 4*n*i
movl     (%rax,%rcx,4), %eax # a + 4*n*i + 4*j
ret
```

Today

- Data
 - Arrays
 - Structures
 - Unions
 - Floating Point and SIMD operations
- Buffer Overflow
 - Memory Layout
 - Vulnerability and Protection

Structure Representation

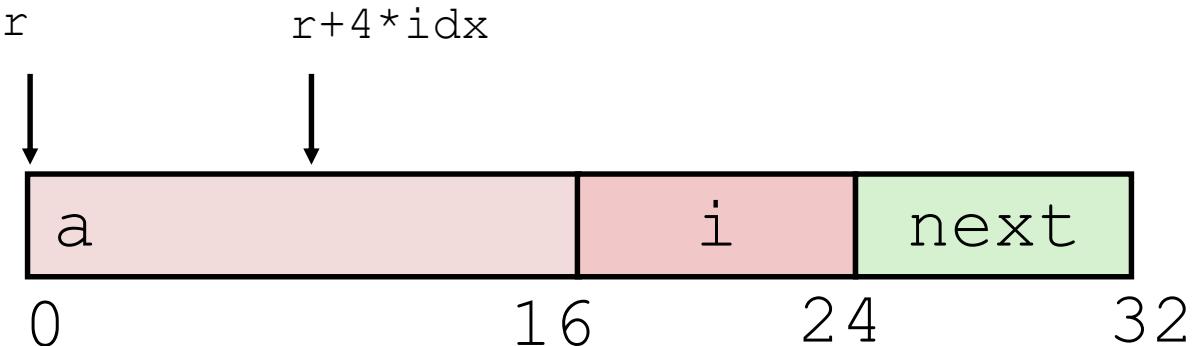
```
struct rec {  
    int a[4];  
    size_t i;  
    struct rec *next;  
};
```



- Structure represented as block of memory
 - **Big enough to hold all of the fields**
- Fields ordered according to declaration
 - **Even if another ordering could yield a more compact representation**
- Compiler determines overall size + positions of fields
 - **Machine-level program has no understanding of the structures in the source code**

Generating Pointer to Structure Member

```
struct rec {
    int a[4];
    size_t i;
    struct rec *next;
};
```



- Generating Pointer to Array Element
 - Offset of each structure member determined at compile time
 - Compute as `r + 4*idx`

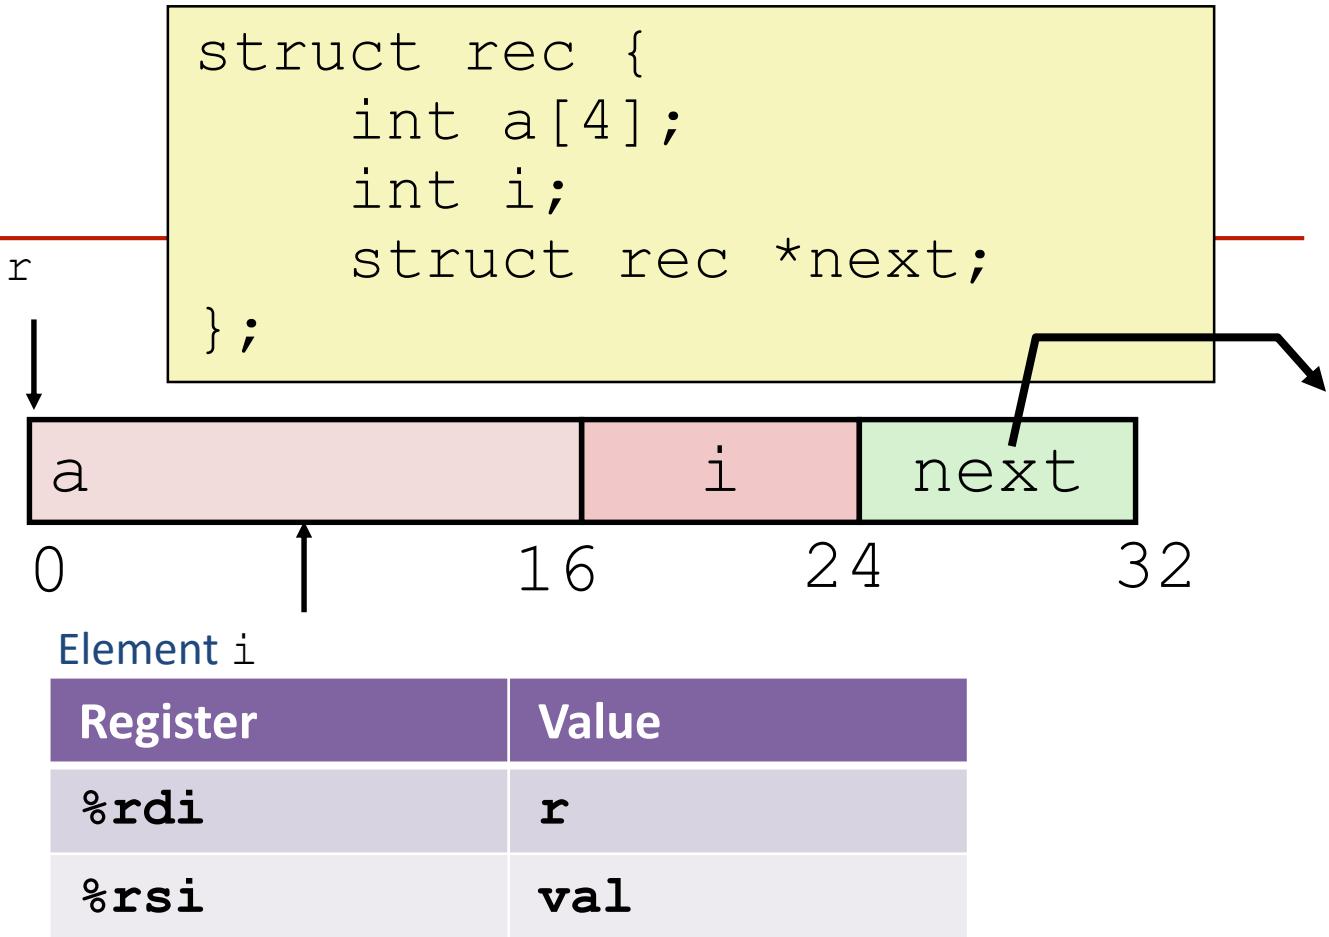
```
int *get_ap
(struct rec *r, size_t idx)
{
    return &r->a[idx];
}
```

```
# r in %rdi, idx in %rsi
leaq (%rdi,%rsi,4), %rax
ret
```

Following Linked List

- C Code

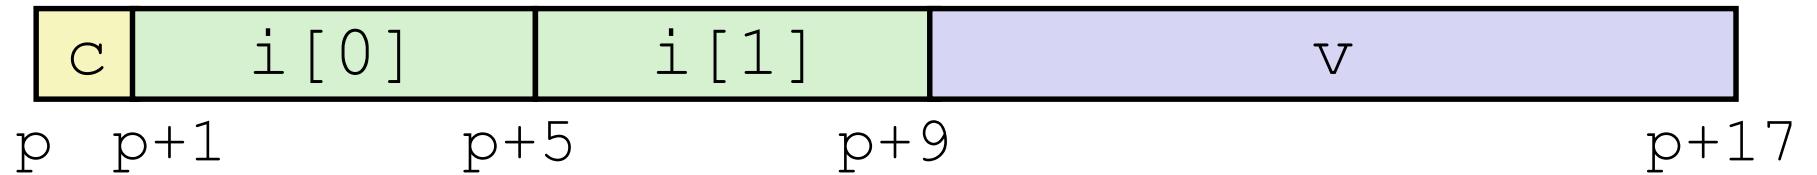
```
void set_val
    (struct rec *r, int val)
{
    while (r) {
        int i = r->i;
        r->a[i] = val;
        r = r->next;
    }
}
```



```
.L11:                                # loop:
    movslq 16(%rdi), %rax          #     i = M[r+16]
    movl    %esi, (%rdi,%rax,4)   #     M[r+4*i] = val
    movq    24(%rdi), %rdi        #     r = M[r+24]
    testq   %rdi, %rdi           #     Test r
    jne     .L11                  #     if !=0 goto loop
```

Structures & Alignment

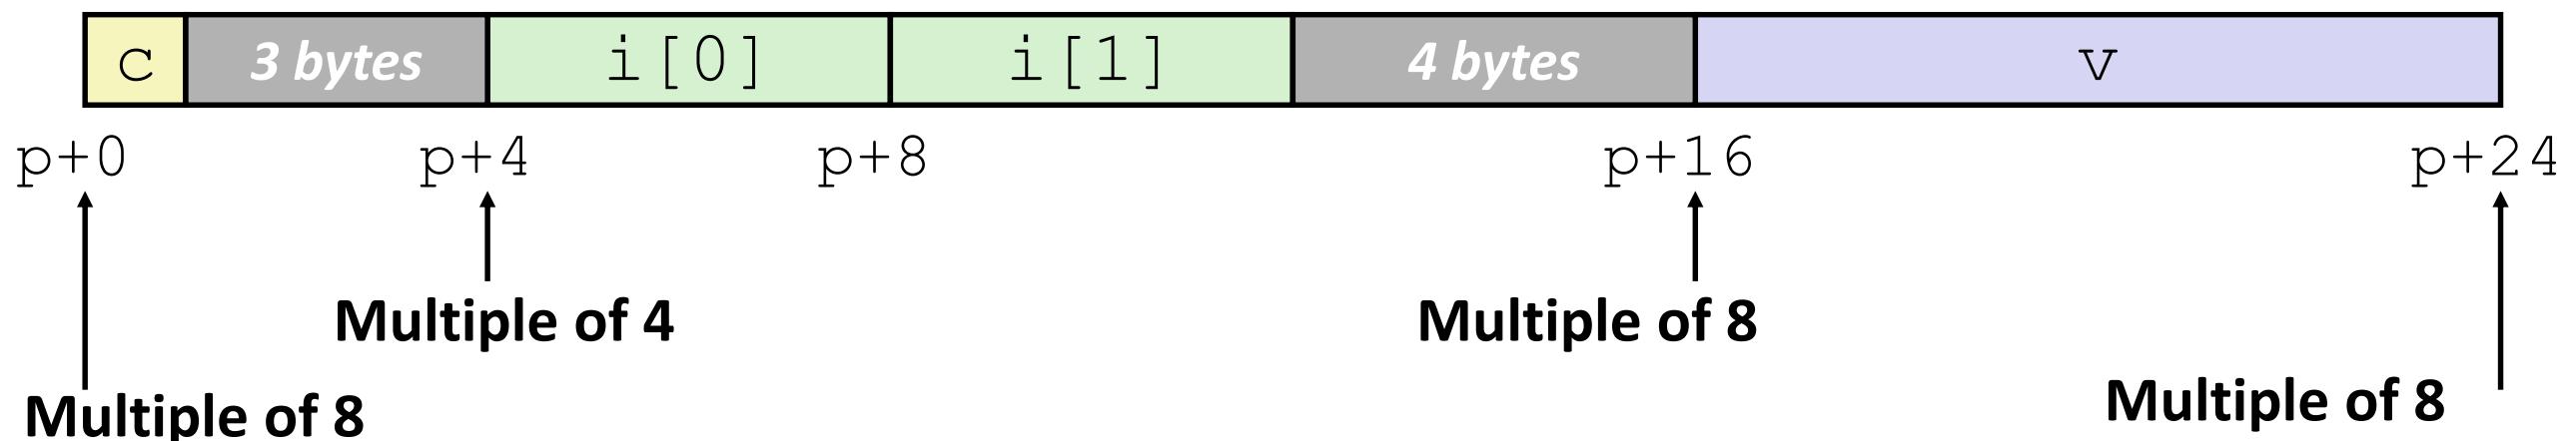
- Unaligned Data



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

- Aligned Data

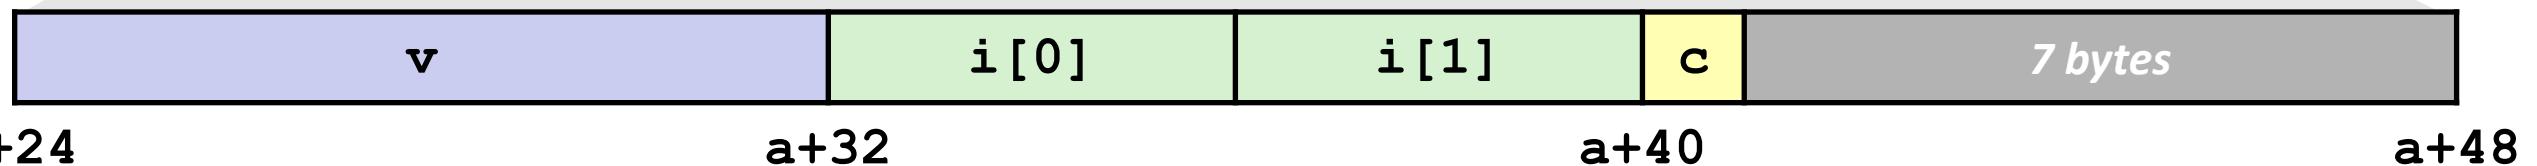
- Primitive data type requires K bytes
- Address must be multiple of K
- Required on some machines; advised on x86-64
- Motivation: Inefficient to load/store datum that spans word boundaries



Arrays of Structures

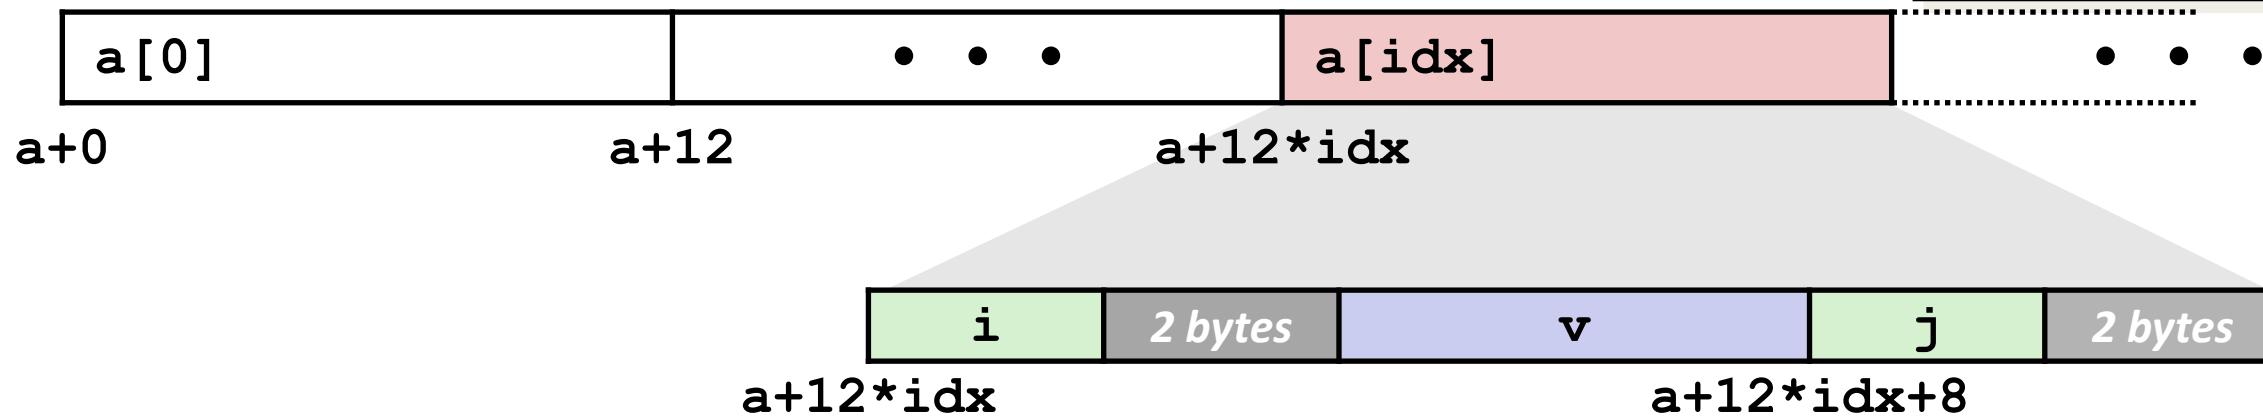
- Overall structure length multiple of K
- Satisfy alignment requirement for every element

```
struct S2 {
    double v;
    int i[2];
    char c;
} a[10];
```



Accessing Array Elements

- Compute array offset $12 * \text{idx}$
 - `sizeof(S3)`, including alignment spacers
- Element **j** is at offset 8 within structure
- Assembler gives offset **a+8**
 - Resolved during linking



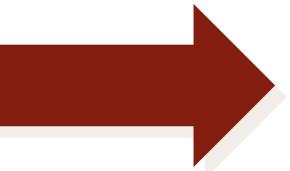
```
short get_j(int idx)
{
    return a[idx].j;
}
```

```
# %rdi = idx
leaq (%rdi,%rdi,2),%rax # 3*idx
movzwl a+8(%rax,4),%eax
```

Saving Space

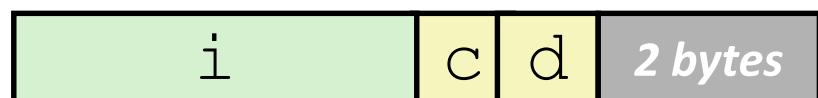
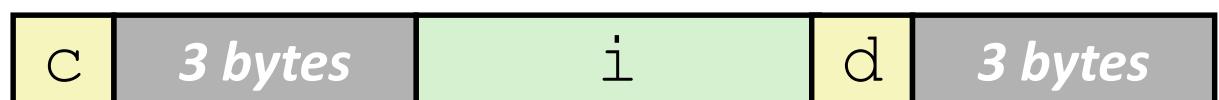
- Put large data types first

```
struct S4 {  
    char c;  
    int i;  
    char d;  
} *p;
```



```
struct S5 {  
    int i;  
    char c;  
    char d;  
} *p;
```

- Effect (K=4)



Today

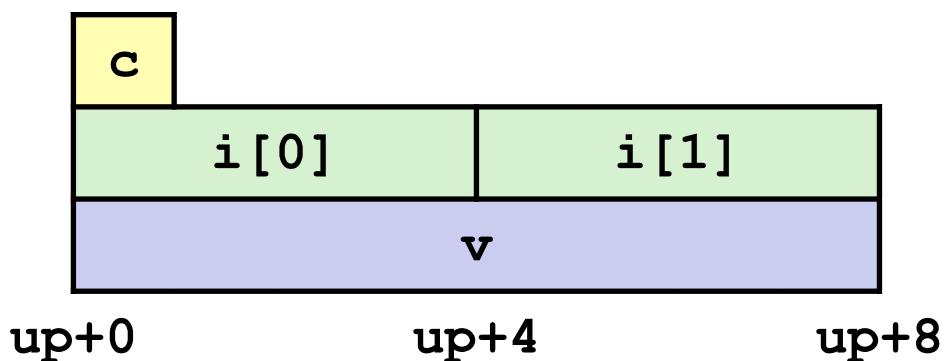
- Data
 - Arrays
 - Structures
 - Unions
 - Floating Point and SIMD operations
- Buffer Overflow
 - Memory Layout
 - Vulnerability and Protection

Union Allocation

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

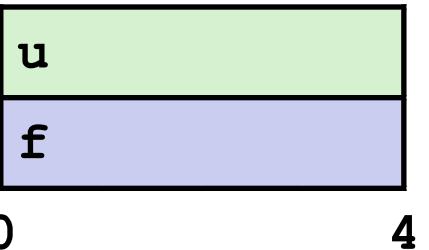
```
union U1 {
    char c;
    int i[2];
    double v;
} *up;
```

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



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;  
}
```

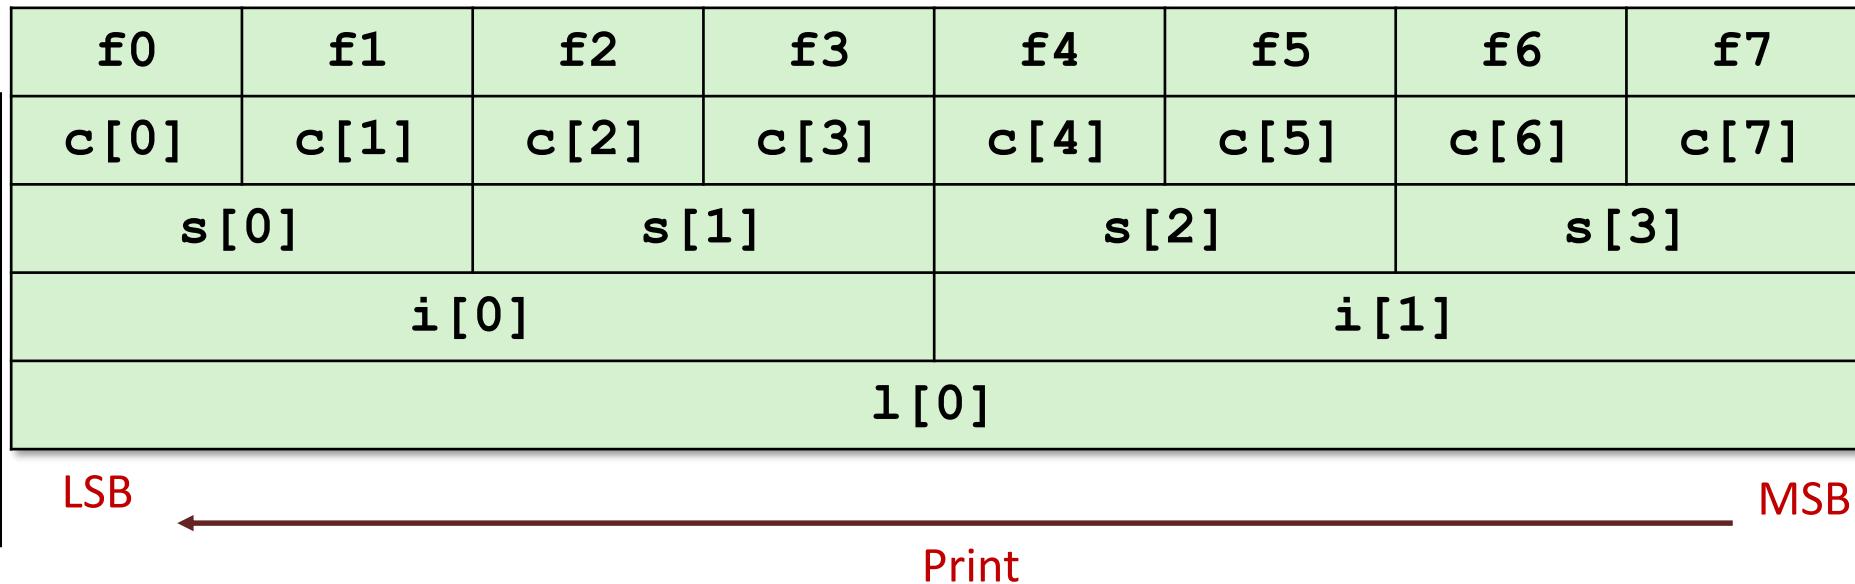
```
unsigned float2bit(float f) {  
    bit_float_t arg;  
    arg.f = f;  
    return arg.u;  
}
```

Same as (float) u ?

Same as (unsigned) f ?

Byte Ordering Example on x86-64

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



Output on x86-64 (little endian):

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 ==	[0xf7f6f5f4f3f2f1f0]

Today

- Data
 - Arrays
 - Structures
 - Unions
 - Floating Point and SIMD operation
- Buffer Overflow
 - Memory Layout
 - Vulnerability and Protection

Background

- History
 - x87 FP
 - Legacy, very ugly
 - SSE FP
 - Supported by Shark machines
 - Special case use of vector instructions
 - AVX FP
 - Newest version
 - Similar to SSE
 - Documented in book

Programming with SSE3

XMM Registers

- 16 total, each 16 bytes
- 16 single-byte integers



- 8 16-bit integers



- 4 32-bit integers



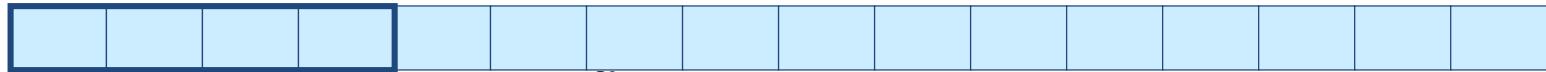
- 4 single-precision floats



- 2 double-precision floats



- 1 single-precision float

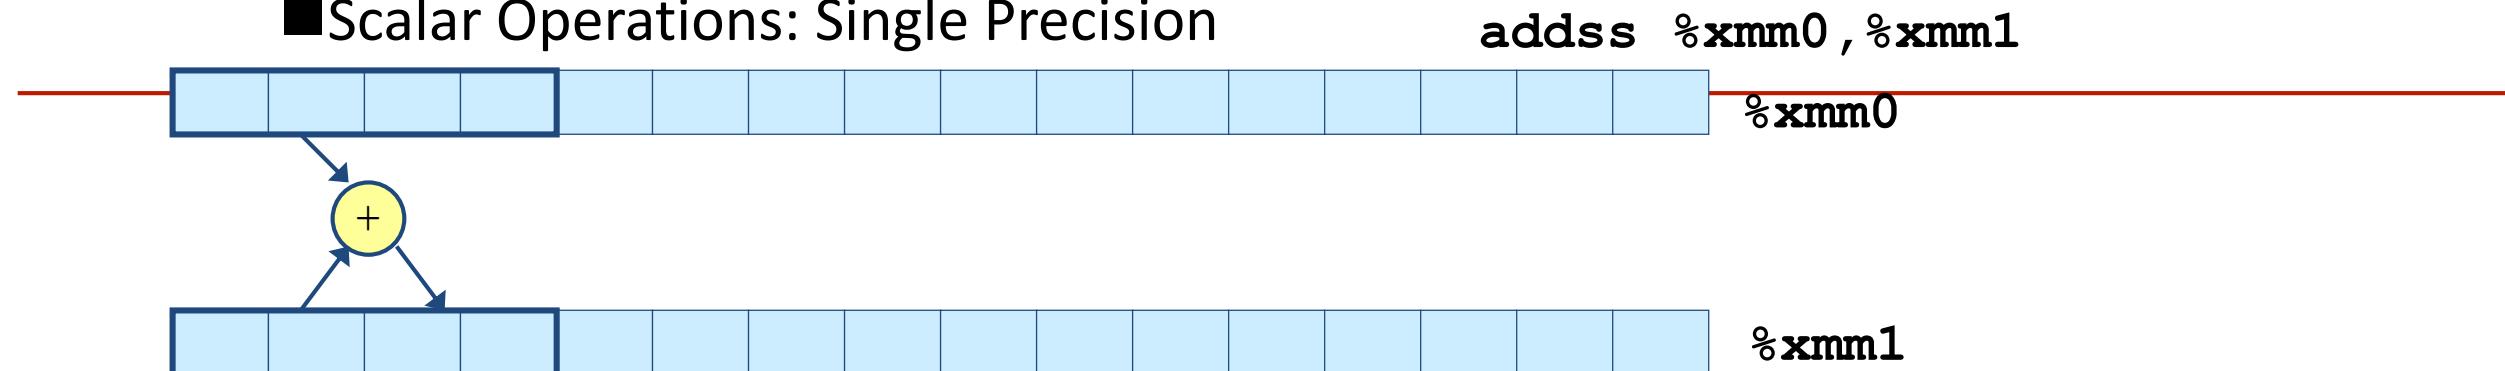


- 1 double-precision float

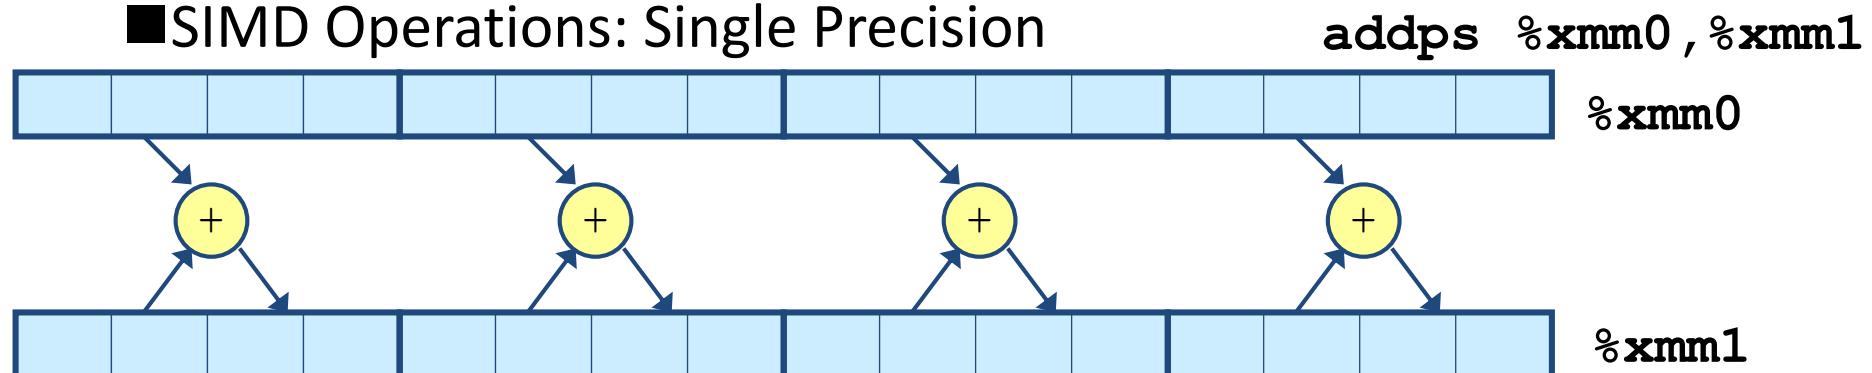


Scalar & SIMD Operations

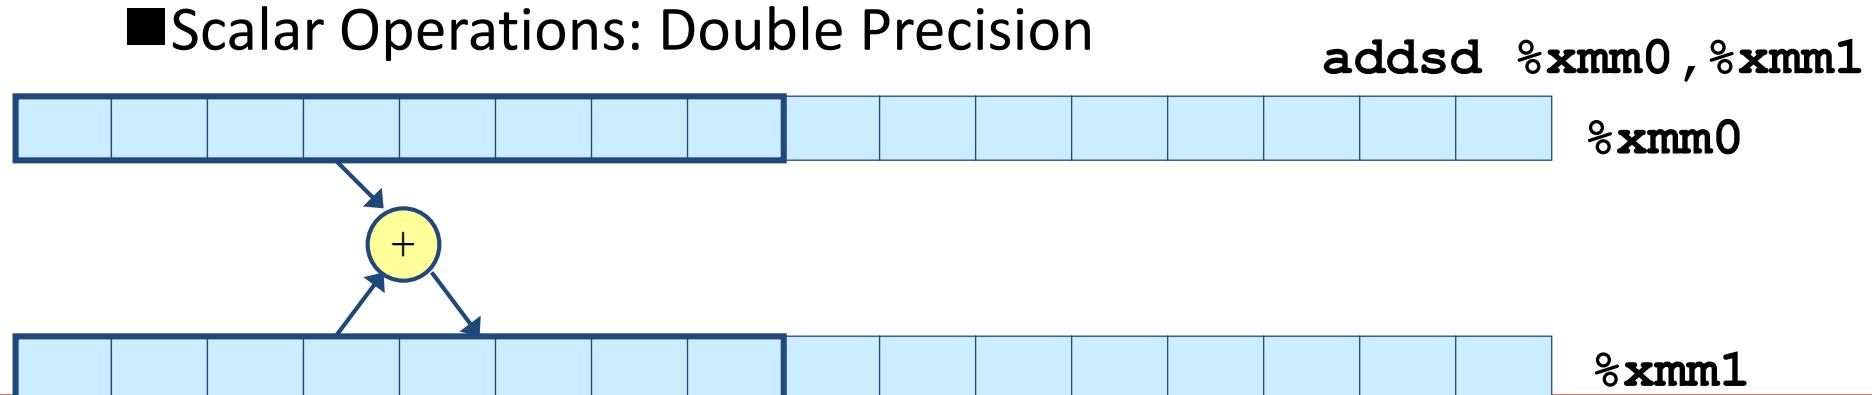
■ Scalar Operations: Single Precision



■ SIMD Operations: Single Precision



■ Scalar Operations: Double Precision



FP Basics

- Arguments passed in `%xmm0, %xmm1, ...`
- Result returned in `%xmm0`
- All XMM registers caller-saved

```
float fadd(float x, float y)
{
    return x + y;
}
```

```
# x in %xmm0, y in %xmm1
addss    %xmm1, %xmm0
ret
```

```
double dadd(double x, double y)
{
    return x + y;
}
```

```
# x in %xmm0, y in %xmm1
addsd    %xmm1, %xmm0
ret
```

FP Memory Referencing

- Integer (and pointer) arguments passed in regular registers
- FP values passed in XMM registers
- Different mov instructions to move between XMM registers, and between memory and XMM registers

```
double dincr(double *p, double v)
{
    double x = *p;
    *p = x + v;
    return x;
}
```

```
# p in %rdi, v in %xmm0
movapd  %xmm0, %xmm1    # Copy v
movsd   (%rdi), %xmm0  # x = *p
addsd   %xmm0, %xmm1    # t = x + v
movsd   %xmm1, (%rdi)  # *p = t
ret
```

Today

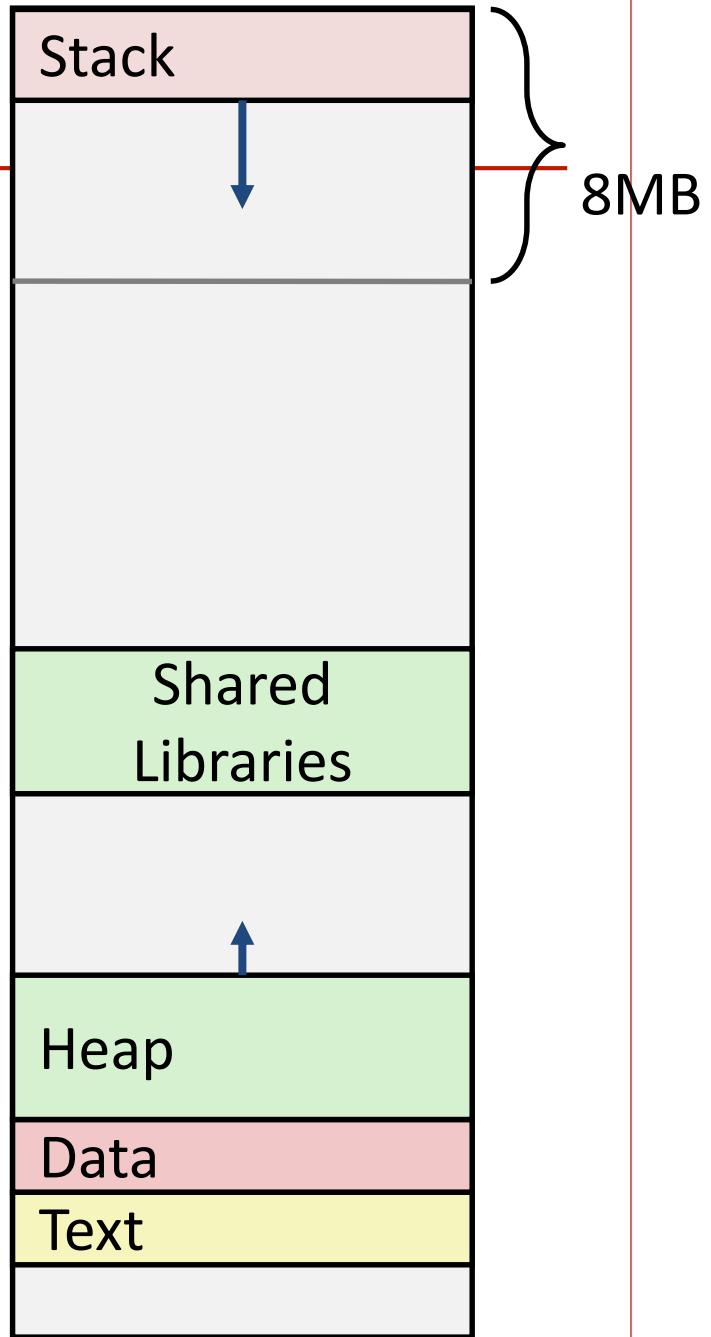
- Data
 - Arrays
 - Structures
 - Unions
 - Floating Point and SIMD operations
- Buffer Overflow
 - Memory Layout
 - Vulnerability and Protection

x86-64 Linux Memory Layout

00007FFFFFFFFF

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

Hex Address

400000
000000

Memory Allocation Example

```

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

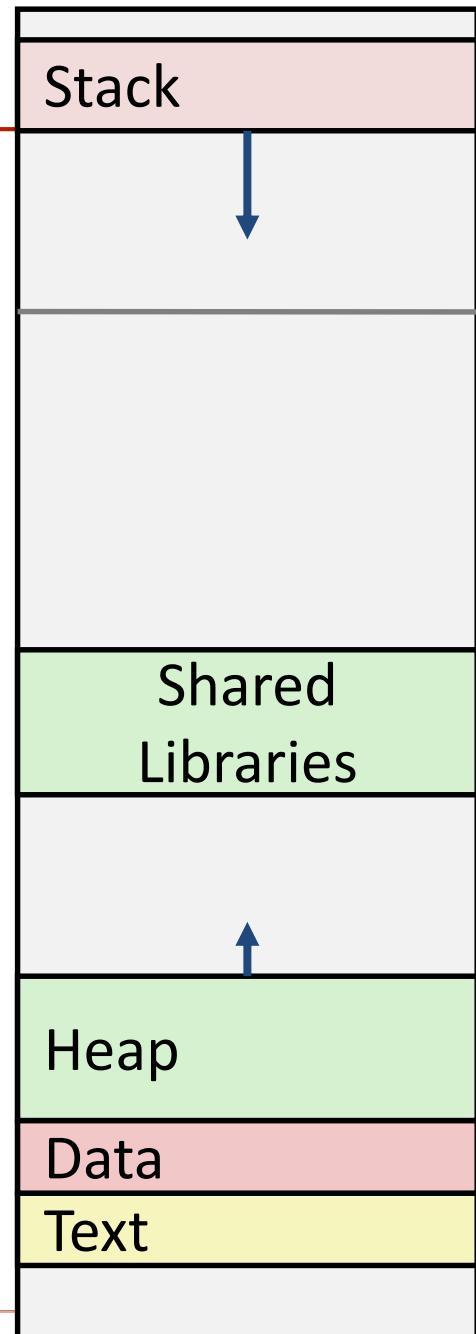
int global = 0;

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 ... */
}

```

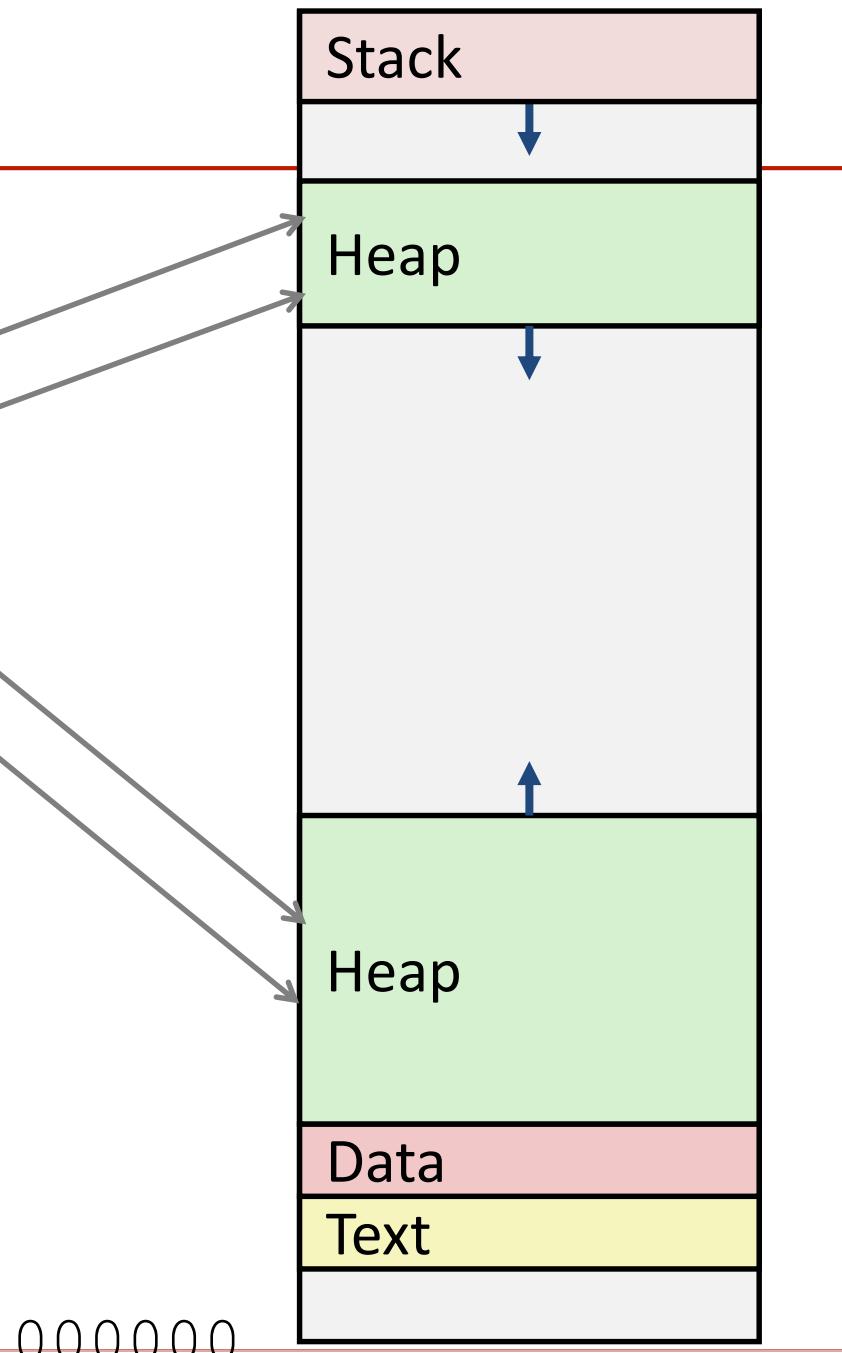
Where does everything go?



x86-64 Example Addresses

address range $\sim 2^{47}$

local	0x00007ffe4d3be87c
p1	0x00007f7262a1e010
p3	0x00007f7162a1d010
p4	0x000000008359d120
p2	0x000000008359d010
big_array	0x0000000080601060
huge_array	0x0000000000601060
main()	0x000000000040060c
useless()	0x0000000000400590



Today

- Data
 - Arrays
 - Structures
 - Unions
 - Floating Point and SIMD operations
- Buffer Overflow
 - Memory Layout
 - Vulnerability and protection

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
```

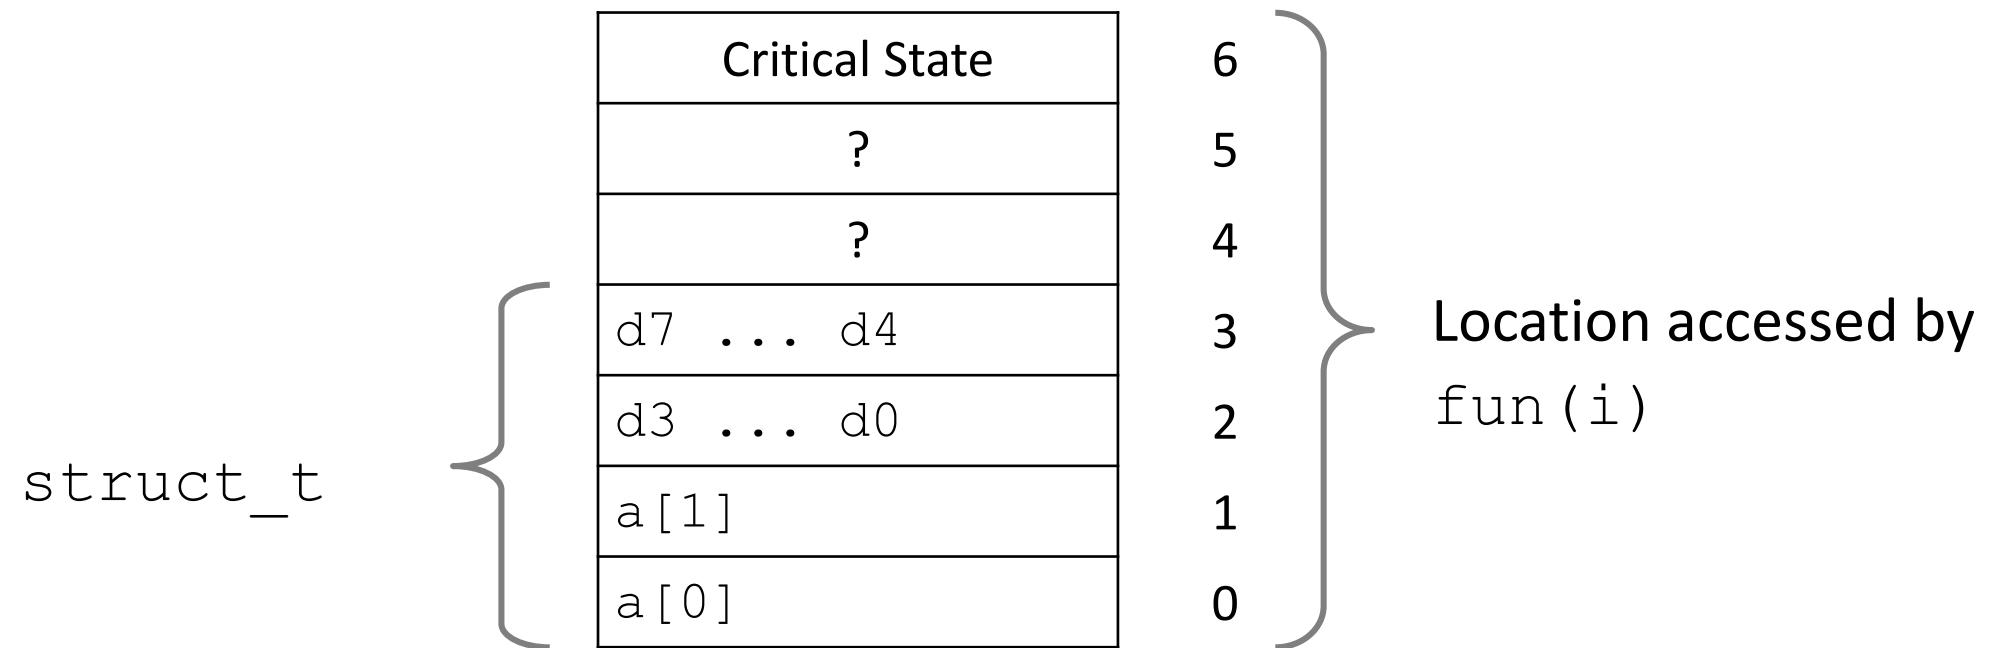
- Result is system specific

Memory Referencing Bug Example

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

Explanation:

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



Such problems are a BIG deal

- Generally called a “buffer overflow”
 - when exceeding the memory size allocated for an array
- Why a big deal?
 - It’s the #1 technical cause of security vulnerabilities
 - #1 overall cause is social engineering / user ignorance
- Most common form
 - Unchecked lengths on string inputs
 - Particularly for bounded character arrays on the stack
 - sometimes referred to as stack smashing

String Library Code

- Implementation of Unix 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;
}
```

- 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

Vulnerable Buffer Code

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

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

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

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

Buffer Overflow Disassembly

echo:

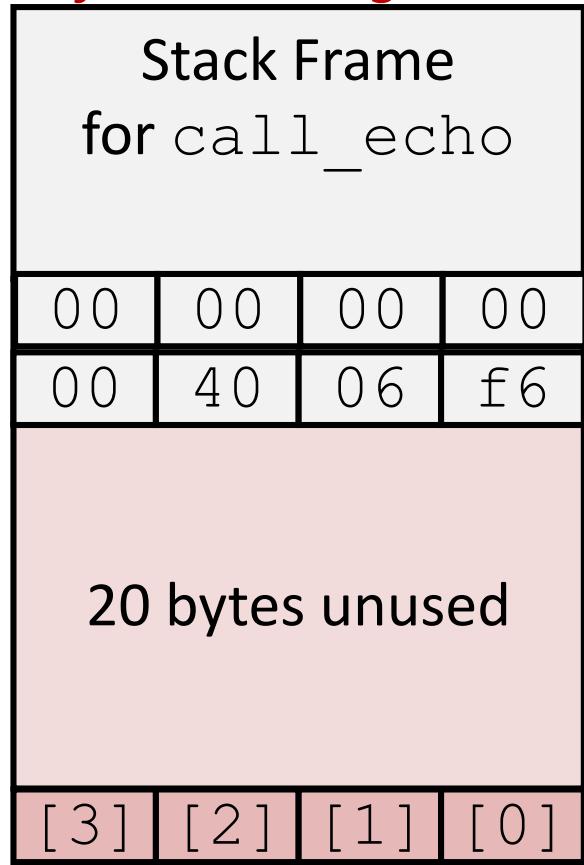
```
00000000004006cf <echo>:
4006cf: 48 83 ec 18           sub    $0x18,%rsp
4006d3: 48 89 e7             mov    %rsp,%rdi
4006d6: e8 a5 ff ff ff      callq   400680 <gets>
4006db: 48 89 e7             mov    %rsp,%rdi
4006de: e8 3d fe ff ff      callq   400520 <puts@plt>
4006e3: 48 83 c4 18        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

Buffer Overflow Stack Example

Before call to gets



```
void echo ()
{
    char buf[4];
    gets(buf);
    ...
}
```

```
echo:
subq $24, %rsp
movq %rsp, %rdi
call gets
...
```

```
call_echo:
...
4006f1: callq 4006cf <echo>
4006f6: add    $0x8, %rsp
...
```

buf ← %rsp

Buffer Overflow Stack Example #1

After call to gets

Stack Frame for call_echo			
00	00	00	00
00	40	06	f6
00	32	31	30
39	38	37	36
35	34	33	32
31	30	39	38
37	36	35	34
33	32	31	30

Overflowed buffer, but did not
corrupt state

```
void echo()
{
    char buf[4];
    gets(buf);
    ...
}
```

call_echo:

```
...
4006f1: callq 4006cf <echo>
4006f6: add    $0x8,%rsp
...
```

buf ← %rsp

```
unix>./bufdemo-nsp
Type a string: 01234567890123456789012
01234567890123456789012
```

Buffer Overflow Stack Example #2

After call to gets

Stack Frame for call_echo			
00	00	00	00
00	40	00	34
33	32	31	30
39	38	37	36
35	34	33	32
31	30	39	38
37	36	35	34
33	32	31	30

```
void echo()
{
    char buf[4];
    gets(buf);
    ...
}
```

```
echo:
    subq $24, %rsp
    movq %rsp, %rdi
    call gets
    ...
call_echo:
```

```
...
4006f1: callq 4006cf <echo>
4006f6: add    $0x8, %rsp
...
buf ← %rsp
```

Overflowed buffer and
corrupted return pointer

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

Buffer Overflow Stack Example #3

After call to gets

Stack Frame for call_echo			
00	00	00	00
00	40	06	00
33	32	31	30
39	38	37	36
35	34	33	32
31	30	39	38
37	36	35	34
33	32	31	30

buf ← %rsp

```
void echo()
{
    char buf[4];
    gets(buf);
    ...
}
```

```
echo:
    subq $24, %rsp
    movq %rsp, %rdi
    call gets
    ...
call_echo:
```

```
...
4006f1: callq 4006cf <echo>
4006f6: add    $0x8, %rsp
...
```

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

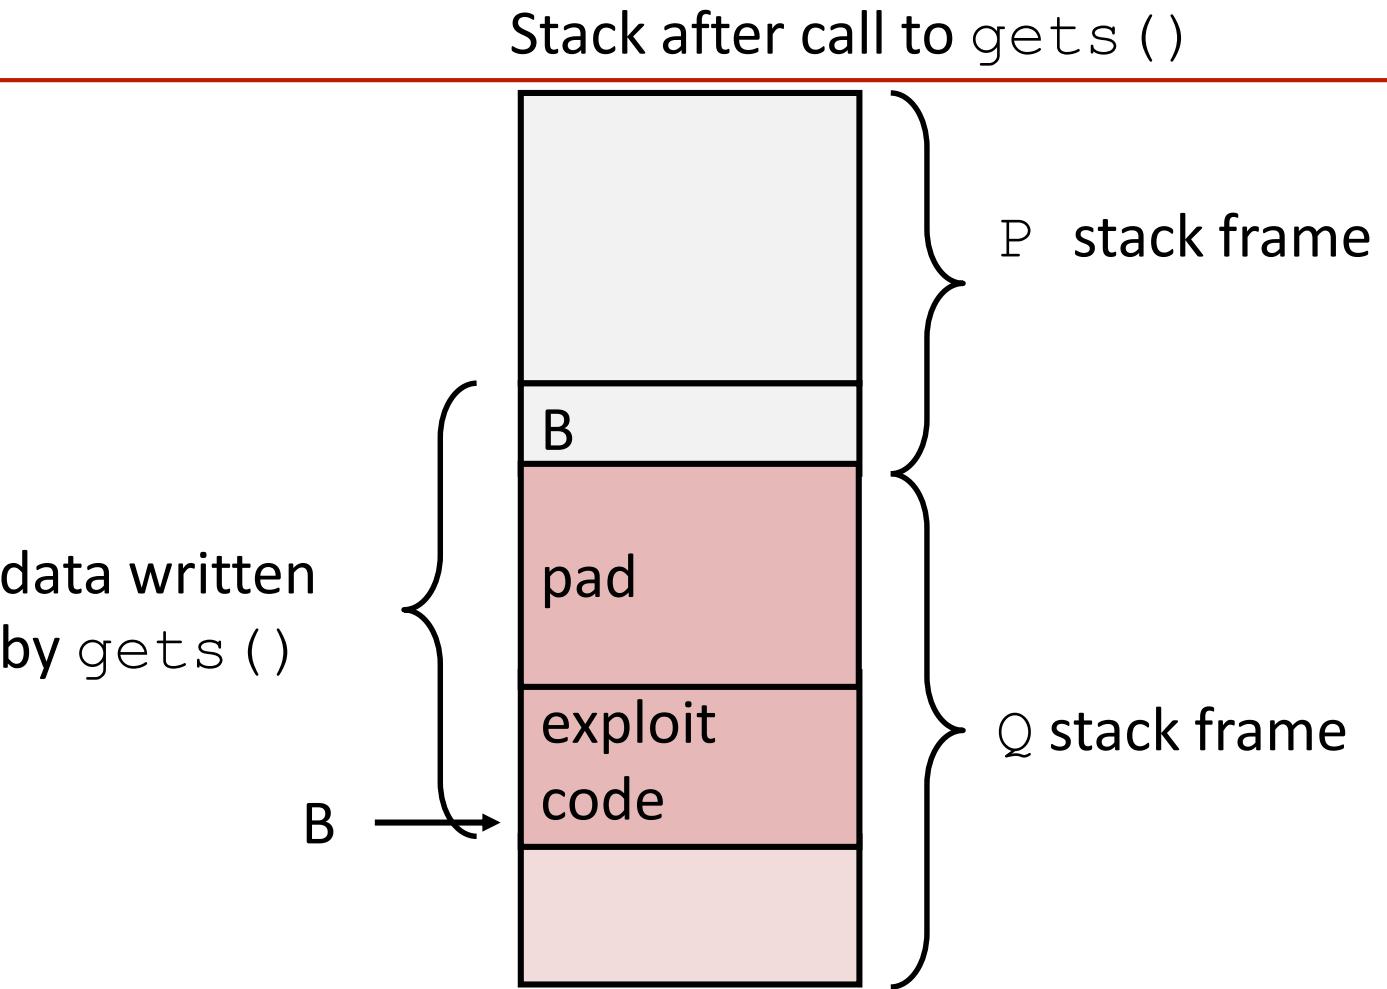
But “Returns” to unrelated code

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

Code Injection Attacks

```
void P() {
    Q();
    ...
}
int Q() {
    char buf[64];
    gets(buf);
    ...
    return ...;
}
```

return address
A



- 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

What to do about buffer overflow attacks

- Avoid overflow vulnerabilities
- Employ system-level protections
- Have compiler use “stack canaries”

1. Avoid Overflow Vulnerabilities in Code (!)

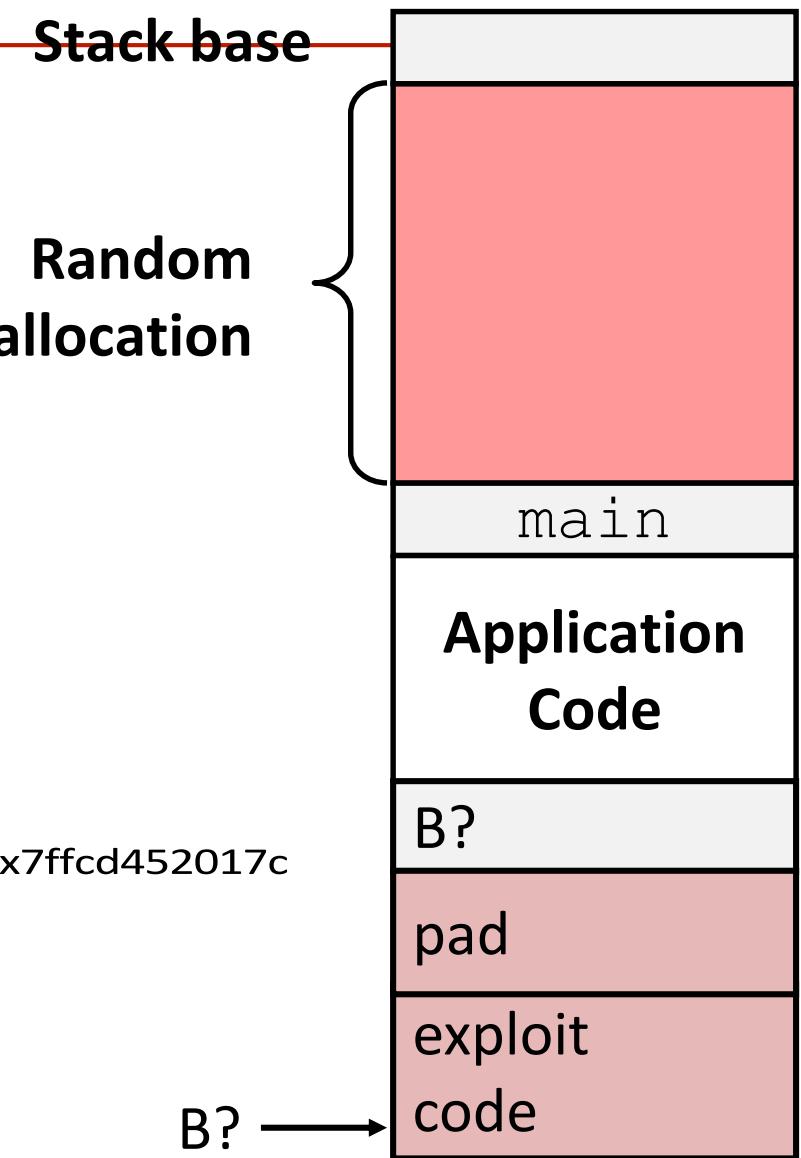
```
/* Echo Line */
void echo()
{
    char buf[4]; /* Way too small! */
    fgets(buf, 4, stdin);
    puts(buf);
}
```

- 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

2. System-Level Protections can help

- Randomized stack offsets
 - 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

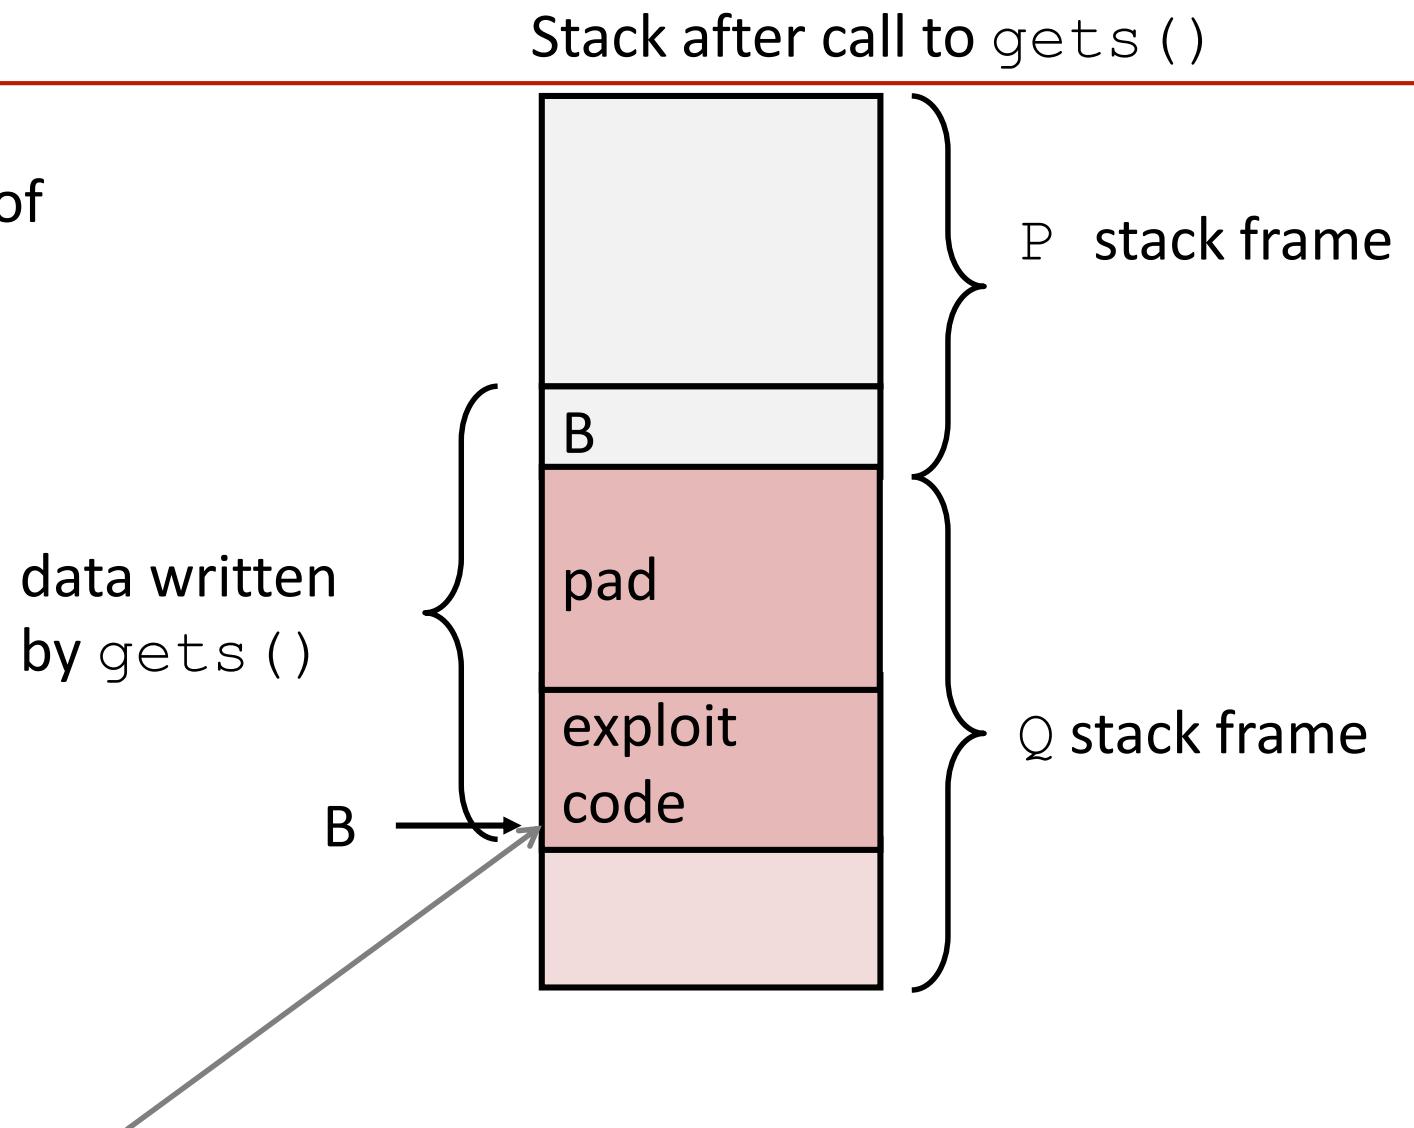
local 0x7ffe4d3be87c 0x7fff75a4f9fc 0x7ffeadb7c80c 0x7ffeaea2fdac 0x7ffcd452017c



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 added explicit “execute” permission
- Stack marked as non-executable



Any attempt to execute this code will fail

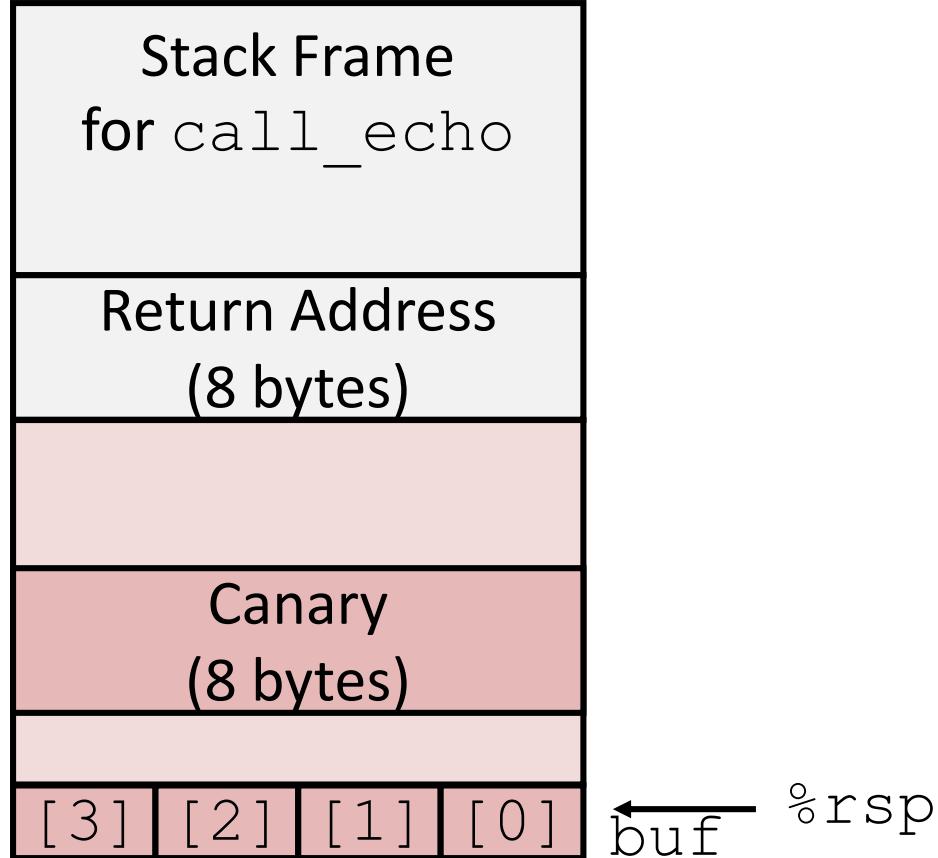
3. Stack Canaries can help

- Idea
 - Place special value (“canary”) on stack just beyond buffer
 - Check for corruption before exiting function
- GCC Implementation
 - **-fstack-protector**
 - Now the default (disabled earlier)

```
unix>./bufdemo-sp
Type a string: 0123456
0123456
unix>./bufdemo-sp
Type a string: 01234567
*** stack smashing detected ***
```

Setting Up Canary

Before call to gets

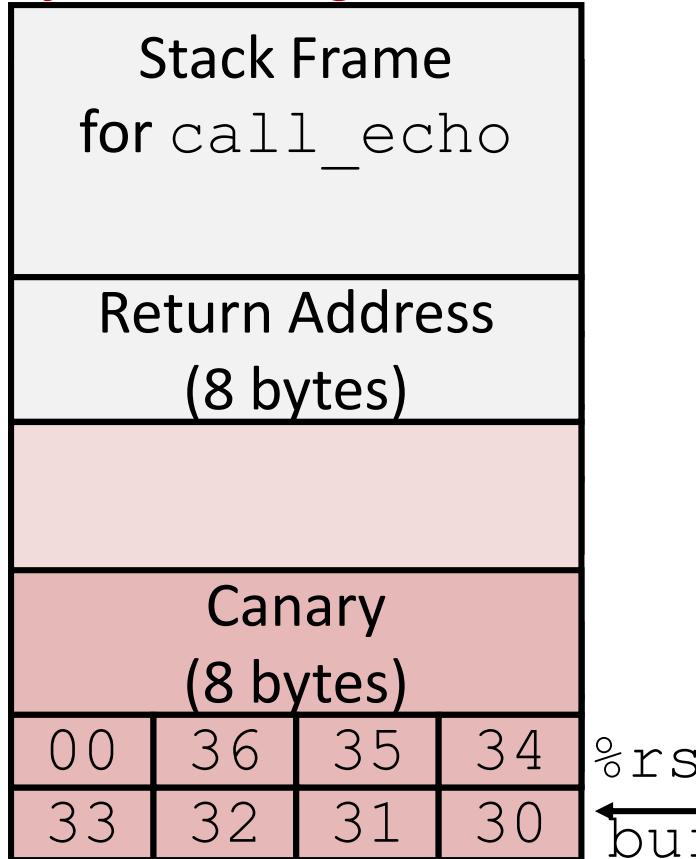


```
/* Echo Line */
void echo()
{
    char buf[4]; /* Way too small! */
    gets(buf);
    puts(buf);
}
```

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

Checking Canary

After call to gets



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

Input: 0123456

```
echo:
    . . .
    movq    8(%rsp), %rax      # Retrieve from stack
    xorq    %fs:40, %rax      # Compare to canary
    je     .L6                  # If same, OK
    call   __stack_chk_fail    # FAIL
.L6:
    . . .
```

Protected Buffer Disassembly

echo:

```
40072f: sub    $0x18,%rsp
400733: mov    %fs:0x28,%rax
40073c: mov    %rax,0x8(%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
```

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 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

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
```

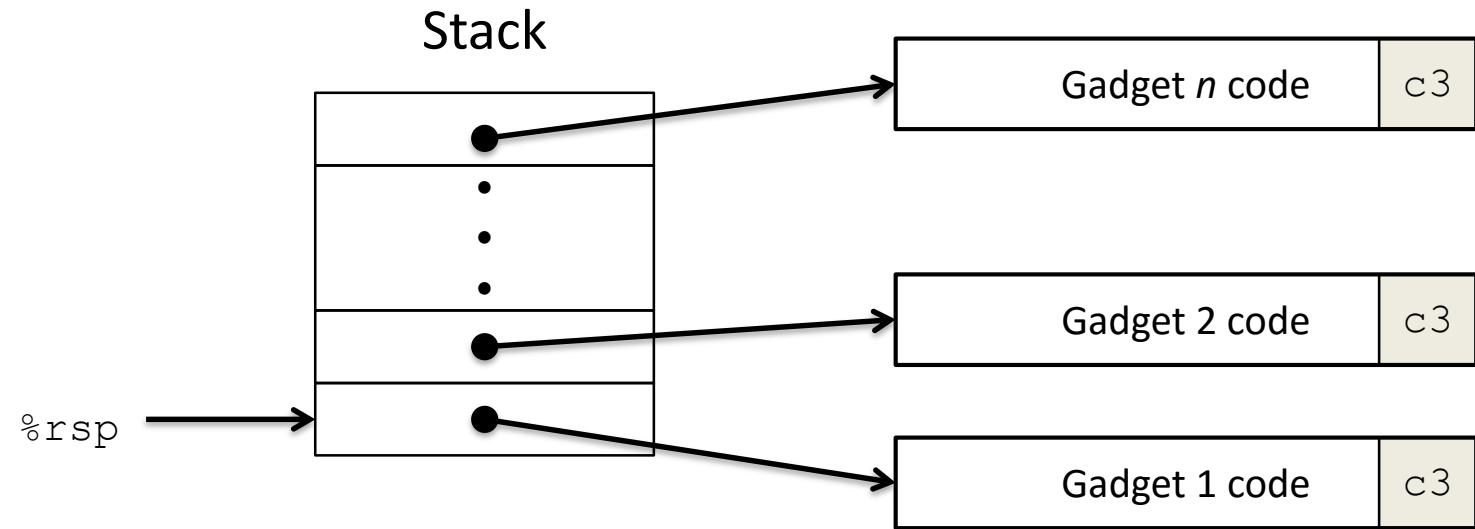
Encodes `movq %rax, %rdi`

`rdi ← rax`

Gadget address = 0x4004dc

- Repurpose byte codes

ROP Execution



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

Summary

- Arrays
 - Elements packed into contiguous region of memory
 - Use index arithmetic to locate individual elements
- Structures
 - Elements packed into single region of memory
 - Access using offsets determined by compiler
 - Possible require internal and external padding to ensure alignment
- Unions
 - Overlay declarations
 - Way to circumvent type system
- Floating Point
 - Data held and operated on in XMM registers
- “Buffer overflow” exceeds the memory size allocated for an array

18-600 Foundations of Computer Systems

Lecture 8: "Processor Architecture and Design"

John P. Shen & Zhiyi Yu

September 26, 2016

Next Time ...

- Required Reading Assignment:
 - Chapter 4 of CS:APP (3rd edition) by Randy Bryant & Dave O'Hallaron.
- Recommended Reference:
 - ❖ Chapters 1 and 2 of Shen and Lipasti (SnL).



Electrical & Computer
ENGINEERING