

18-600 Foundations of Computer Systems

Lecture 6:

"Machine-Level Programming III: Loops, Procedures, the Stack, and More"

September 18, 2017

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



Today

- Control
 - **Conditional Branches**
 - Switch Statements
- Procedures
 - Stack Structure
 - Passing control & data
 - Managing local data
 - Illustration of Recursion
- Buffer Overflow (Attacks)
- SSE, SIMD, FP

Jumping

- jX Instructions
 - Jump to different part of code depending on condition codes

| jX | Condition | Description |
|-----|--------------------------------------|---------------------------|
| jmp | 1 | Unconditional |
| je | ZF | Equal / Zero |
| jne | $\sim ZF$ | Not Equal / Not Zero |
| js | SF | Negative |
| jns | $\sim SF$ | Nonnegative |
| jg | $\sim (SF \wedge OF) \ \& \ \sim ZF$ | Greater (Signed) |
| jge | $\sim (SF \wedge OF)$ | Greater or Equal (Signed) |
| jl | $(SF \wedge OF)$ | Less (Signed) |
| jle | $(SF \wedge OF) \ \ ZF$ | Less or Equal (Signed) |
| ja | $\sim CF \ \& \ \sim ZF$ | Above (unsigned) |
| jb | CF | Below (unsigned) |

Conditional Branch Example (Using Branch)

- Generation

```
shark> gcc -Og -S -fno-if-conversion control.c
```

```
long absdiff
(long x, long y)
{
    long result;
    if (x > y)
        result = x-y;
    else
        result = y-x;
    return result;
}
```

```
absdiff:
    cmpq    %rsi, %rdi    # x:y
    jle    .L4
    movq   %rdi, %rax
    subq   %rsi, %rax
    ret
.L4:     # x <= y
    movq   %rsi, %rax
    subq   %rdi, %rax
    ret
```

| Register | Use(s) |
|----------|-------------------|
| %rdi | Argument x |
| %rsi | Argument y |
| %rax | Return value |

General Conditional Expression Translation (Using Branches)

C Code

```
val = Test ? Then_Expr : Else_Expr;
```

```
val = x > y ? x - y : y - x;
```

Goto Version

```
n_test = !Test;  
if (n_test) goto Else;  
val = Then_Expr;  
goto Done;  
Else:  
    val = Else_Expr;  
Done:  
    . . .
```

- Create separate code regions for then & else expressions
- Execute appropriate one

Using Conditional Moves

- Conditional Move Instructions
 - Instruction supports:
if (Test) Dest \leftarrow Src
 - Supported in post-1995 x86 processors
 - GCC tries to use them
 - But, only when known to be safe
- Why?
 - Branches are very disruptive to instruction flow through pipelines
 - Conditional moves do not require control transfer
- Only make sense when both conditional calculations are simple and safe

C Code

```
val = Test  
  ? Then_Expr  
  : Else_Expr ;
```

Goto Version

```
result = Then_Expr ;  
eval = Else_Expr ;  
nt = !Test ;  
if (nt) result = eval ;  
return result ;
```

Conditional Move Example

```

long absdiff
(long x, long y)
{
    long result;
    if (x > y)
        result = x-y;
    else
        result = y-x;
    return result;
}

```

| Register | Use(s) |
|-------------------|-------------------------|
| <code>%rdi</code> | Argument <code>x</code> |
| <code>%rsi</code> | Argument <code>y</code> |
| <code>%rax</code> | Return value |

```

absdiff:
movq    %rdi, %rax    # x
subq    %rsi, %rax    # result = x-y
movq    %rsi, %rdx
subq    %rdi, %rdx    # eval = y-x
cmpq    %rsi, %rdi    # x:y
cmovle  %rdx, %rax    # if <=, result = eval
ret

```

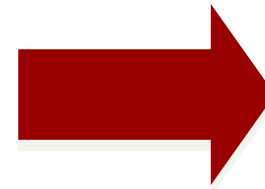
Today

- Control
 - Conditional branches
 - **Loops**
 - Switch Statements
- Procedures
 - Stack Structure
 - Passing control & data
 - Managing local data
 - Illustration of Recursion
- Buffer Overflow (Attacks)
- SSE, SIMD, FP

"While" Translation #1 (Jump to Middle)

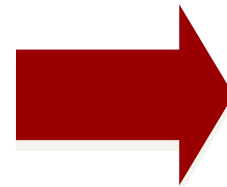
- "Jump-to-middle" translation; Used with `-Og`

```
while (Test)
    Body
```



```
goto test;
loop:
    Body
test:
    if (Test)
        goto loop;
done:
```

```
long pcount_while
(unsigned long x) {
    long result = 0;
    while (x) {
        result += x & 0x1;
        x >>= 1;
    }
    return result;
}
```



```
long pcount_goto_jtm
(unsigned long x) {
    long result = 0;
    goto test;
loop:
    result += x & 0x1;
    x >>= 1;
test:
    if(x) goto loop;
    return result;
}
```

"While" Translation #2 (Do while)

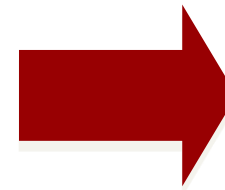
- "Do-while" conversion; Used with `-O1`

```
while (Test)
    Body
```



```
if (!Test)
    goto done;
loop:
    Body
    if (Test)
        goto loop;
done:
```

```
long pcount_while
(unsigned long x) {
    long result = 0;
    while (x) {
        result += x & 0x1;
        x >>= 1;
    }
    return result;
}
```

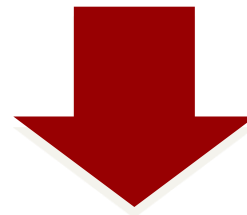


```
long pcount_goto_dw
(unsigned long x) {
    long result = 0;
    if (!x) goto done;
loop:
    result += x & 0x1;
    x >>= 1;
    if(x) goto loop;
done:
    return result;
}
```

"For" Loop → While Loop

General Form

```
for (Init; Test; Update )  
    Body
```

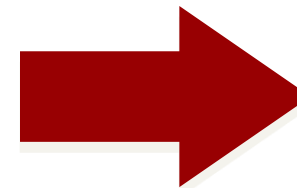


While Version

```
Init ;  
while (Test) {  
    Body  
    Update ;  
}
```

"For" Loop → While Loop (example)

```
#define WSIZE 8*sizeof(int)
long pcount_for
(unsigned long x)
{
    size_t i;
    long result = 0;
    for (i = 0; i < WSIZE; i++)
    {
        unsigned bit =
            (x >> i) & 0x1;
        result += bit;
    }
    return result;
}
```



```
long pcount_for_while
(unsigned long x)
{
    size_t i;
    long result = 0;
    i = 0;
    while (i < WSIZE)
    {
        unsigned bit =
            (x >> i) & 0x1;
        result += bit;
        i++;
    }
    return result;
}
```

"For" Loop Do-While Conversion

```

long pcount_for
(unsigned long x)
{
    size_t i;
    long result = 0;
    for (i = 0; i < WSIZE; i++)
    {
        unsigned bit =
            (x >> i) & 0x1;
        result += bit;
    }
    return result;
}

```

- Initial test can be optimized away

```

long pcount_for_goto_dw
(unsigned long x) {
    size_t i;
    long result = 0;
    i = 0;
    if (!(i < WSIZE))
    goto done;
loop:
    {
        unsigned bit =
            (x >> i) & 0x1;
        result += bit;
    }
    i++;
    if (i < WSIZE)
        goto loop;
done:
    return result;
}

```

Init

~~!Test~~

Body

Update

Test

Today

- Control
 - Control: Condition codes
 - Conditional branches
 - Loops
 - **Switch Statements**
- Procedures
 - Stack Structure
 - Passing control & data
 - Managing local data
 - Illustration of Recursion
- Buffer Overflow (Attacks)
- SSE, SIMD, FP

Switch Statement Example

```
long switch_eg
(long x, long y, long z)
{
    long w = 1;
    switch(x) {
    case 1:
        w = y*z;
        break;
    case 2:
        w = y/z;
        /* Fall Through */
    case 3:
        w += z;
        break;
    case 5:
    case 6:
        w -= z;
        break;
    default:
        w = 2;
    }
    return w;
}
```

- Multiple case labels
 - Here: 5 & 6
- Fall through cases
 - Here: 2
- Missing cases
 - Here: 4

Jump Table Structure

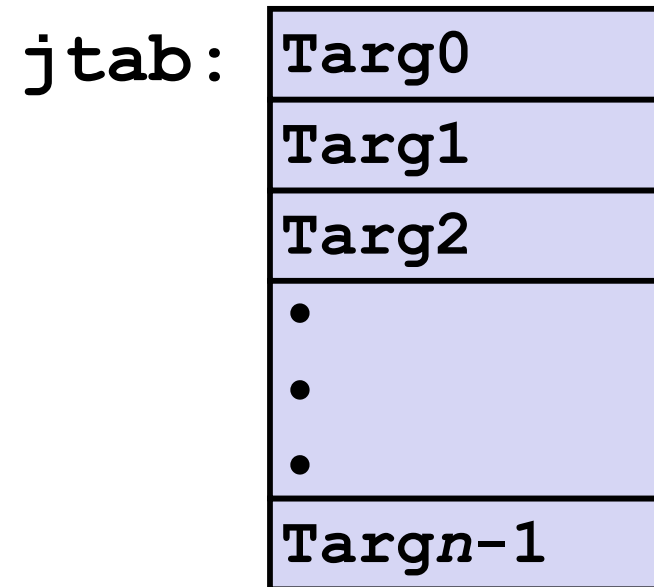
Switch Form

```
switch (x) {
  case val_0:
    Block 0
  case val_1:
    Block 1
    • • •
  case val_n-1:
    Block n-1
}
```

Translation (Extended C)

```
goto *JTab[x];
```

Jump Table



Jump Targets

Targ0:

Code Block
0

Targ1:

Code Block
1

Targ2:

Code Block
2

•

•

•

Targn-1:

Code Block
n-1

Switch Statement Example

```
long switch_eg(long x, long y, long z)
{
    long w = 1;
    switch(x) {
        . . .
    }
    return w;
}
```

Setup:

```
switch_eg:
    movq    %rdx, %rcx
    cmpq    $6, %rdi    # x:6
    ja     .L8          # Use default
    jmp     *.L4(, %rdi, 8)
```

*Indirect
jump* →

What range of values takes default?

| Register | Use(s) |
|----------|-------------------|
| %rdi | Argument x |
| %rsi | Argument y |
| %rdx | Argument z |
| %rax | Return value |

Jump table

```
.section .rodata
    .align 8
.L4:
    .quad .L8 # x = 0
    .quad .L3 # x = 1
    .quad .L5 # x = 2
    .quad .L9 # x = 3
    .quad .L8 # x = 4
    .quad .L7 # x = 5
    .quad .L7 # x = 6
```

Assembly Setup Explanation

- Table Structure
 - Each target requires 8 bytes
 - Base address at `.L4`
- Jumping
 - **Direct:** `jmp .L8`
 - Jump target is denoted by label `.L8`
 - **Indirect:** `jmp *.L4(, %rdi, 8)`
 - Start of jump table: `.L4`
 - Must scale by factor of 8 (addresses are 8 bytes)
 - Fetch target from effective Address `.L4 + x*8`
 - Only for $0 \leq x \leq 6$

Jump table

```
.section      .rodata
    .align 8
.L4:
    .quad     .L8    # x = 0
    .quad     .L3    # x = 1
    .quad     .L5    # x = 2
    .quad     .L9    # x = 3
    .quad     .L8    # x = 4
    .quad     .L7    # x = 5
    .quad     .L7    # x = 6
```

Jump Table

```

.section    .rodata
.align 8
.L4:
.quad      .L8    # x = 0
.quad      .L3    # x = 1
.quad      .L5    # x = 2
.quad      .L9    # x = 3
.quad      .L8    # x = 4
.quad      .L7    # x = 5
.quad      .L7    # x = 6

```

```

switch(x) {
case 1:      // .L3
    w = y*z;
    break;
case 2:      // .L5
    w = y/z;
    /* Fall Through */
case 3:      // .L9
    w += z;
    break;
case 5:
case 6:      // .L7
    w -= z;
    break;
default:    // .L8
    w = 2;
}

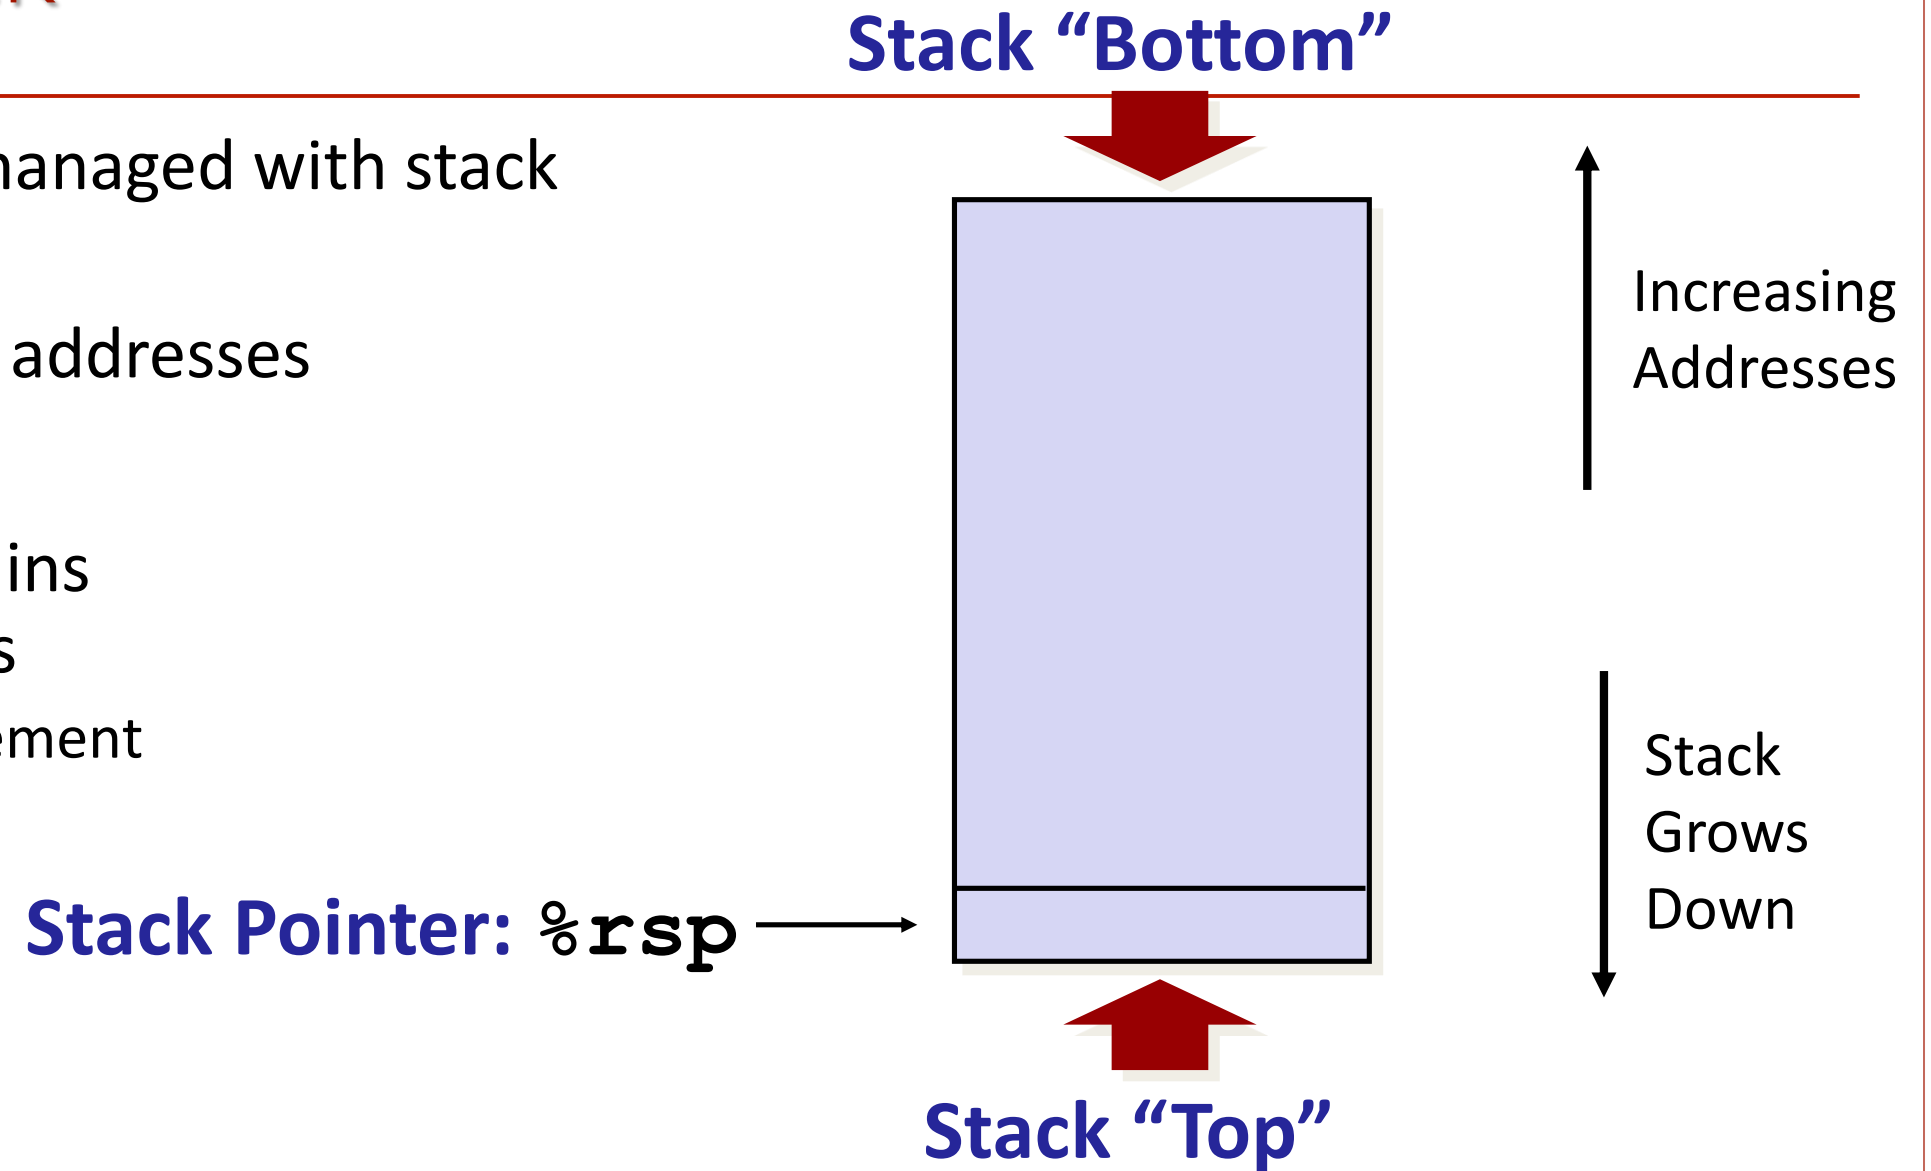
```

Today

- Control
 - Control: Condition codes
 - Conditional branches
 - Loops
 - Switch Statements
- Procedures
 - **Stack Structure**
 - Passing control & data
 - Managing local data
 - Illustration of Recursion
- Buffer Overflow (Attacks)
- SSE, SIMD, FP

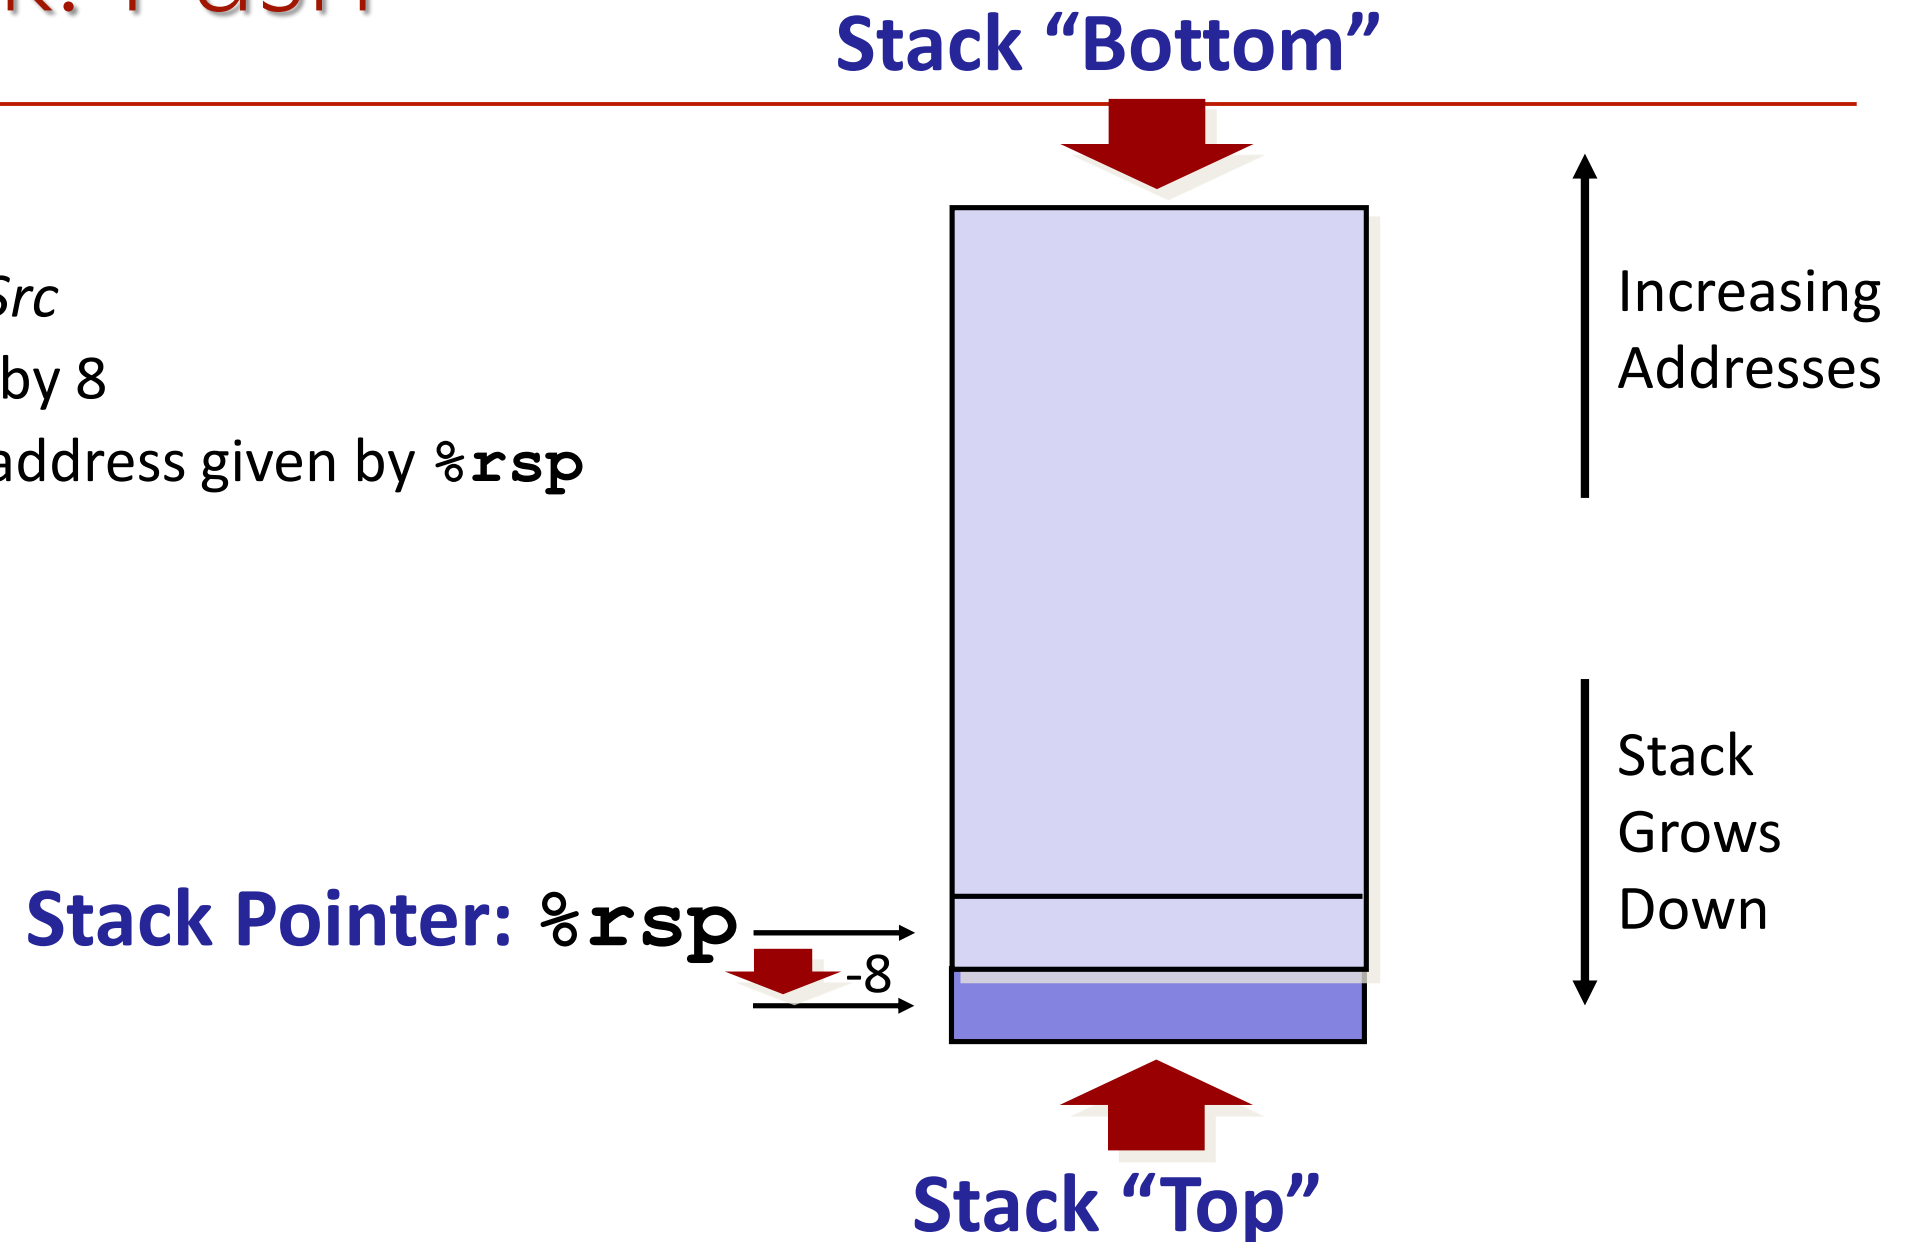
x86-64 Stack

- Region of memory managed with stack discipline
- Grows toward lower addresses
- Register `%rsp` contains lowest stack address
 - address of "top" element



x86-64 Stack: Push

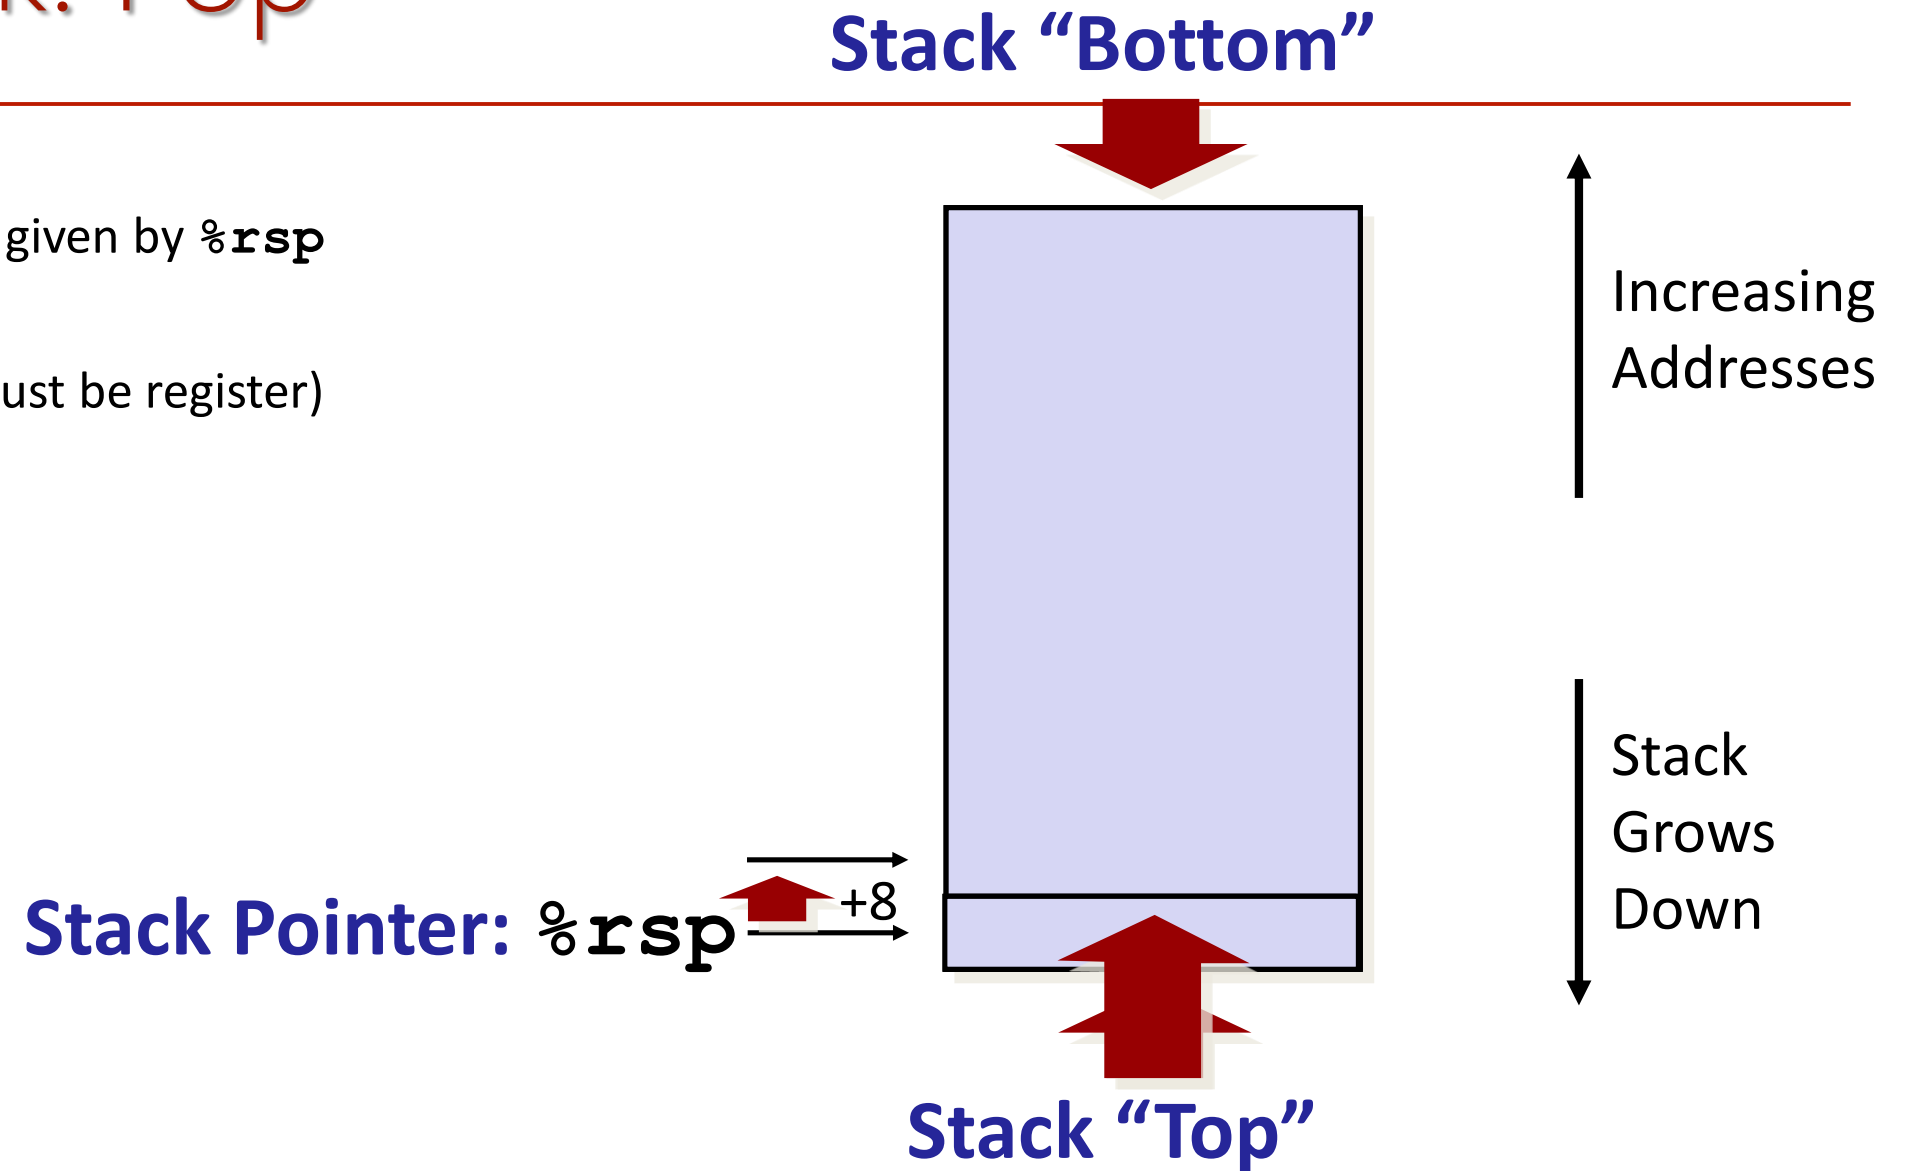
- **pushq Src**
 - Fetch operand at *Src*
 - Decrement `%rsp` by 8
 - Write operand at address given by `%rsp`



x86-64 Stack: Pop

■ `popq Dest`

- Read value at address given by `%rsp`
- Increment `%rsp` by 8
- Store value at `Dest` (must be register)



Today

- Control
 - Control: Condition codes
 - Conditional branches
 - Loops
 - Switch Statements
- Procedures
 - Stack Structure
 - **Passing control & data**
 - Managing local data
 - Illustration of Recursion
 - Buffer Overflow (Attacks)
 - SSE, SIMD, FP

Code Examples

```
void multstore
(long x, long y, long *dest) {
    long t = mult2(x, y);
    *dest = t;
}
```

```
00000000000400540 <multstore>:
400540: push    %rbx          # Save %rbx
400541: mov     %rdx,%rbx     # Save dest
400544: callq  400550 <mult2> # mult2(x,y)
400549: mov     %rax,(%rbx)   # Save at dest
40054c: pop     %rbx          # Restore %rbx
40054d: retq                    # Return
```

```
long mult2
(long a, long b)
{
    long s = a * b;
    return s;
}
```

```
00000000000400550 <mult2>:
400550: mov     %rdi,%rax     # a
400553: imul   %rsi,%rax     # a * b
400557: retq                    # Return
```

Procedure Control Flow

- Use stack to support procedure call and return
- **Procedure call: `call label`**
 - Push return address on stack
 - Jump to *label*
- Return address:
 - Address of the next instruction right after call
 - Example from disassembly
- **Procedure return: `ret`**
 - Pop address from stack
 - Jump to address

Control Flow Example

```

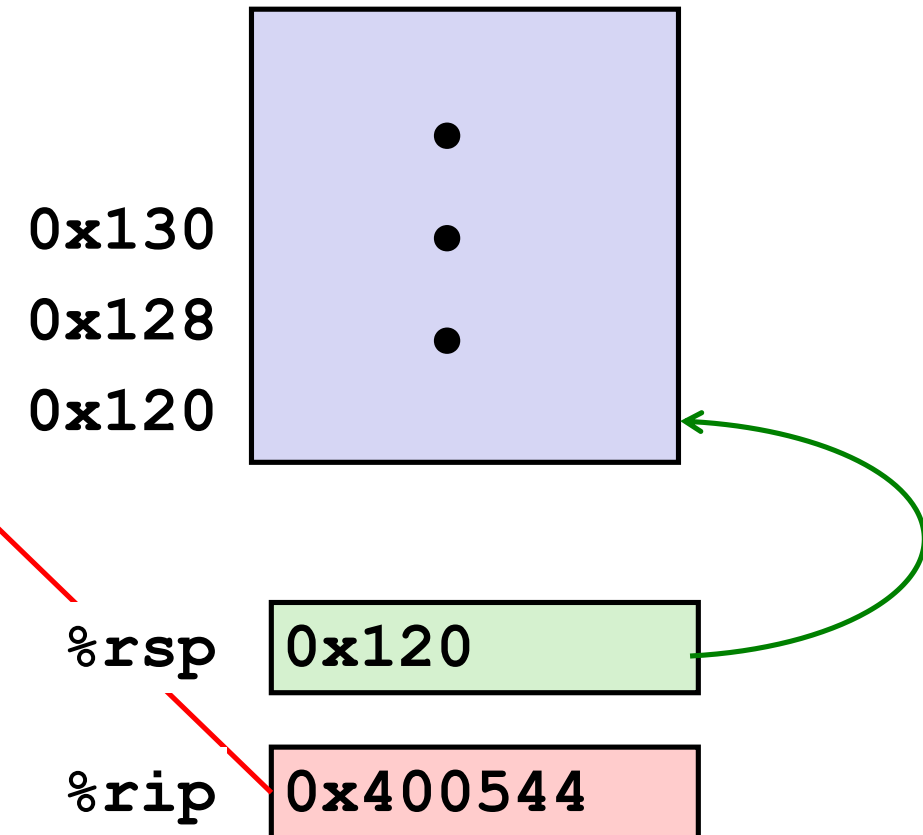
00000000000400540 <multstore>:
•
•
400544: callq  400550 <mult2>
400549: mov   %rax, (%rbx)
•
•

```

```

00000000000400550 <mult2>:
400550: mov   %rdi, %rax
•
•
400557: retq

```



Control Flow Example

```

00000000000400540 <multstore>:
.
.
400544: callq  400550 <mult2>
400549: mov   %rax, (%rbx)
.
.

```

```

00000000000400550 <mult2>:
400550: mov   %rdi, %rax
.
.
400557: retq

```

0x130

0x128

0x120

0x118 0x400549

%rsp 0x118

%rip 0x400550

Control Flow Example

```

00000000000400540 <multstore>:
.
.
400544: callq  400550 <mult2>
400549: mov   %rax, (%rbx)
.
.

```

```

00000000000400550 <mult2>:
400550: mov   %rdi, %rax
.
.
400557: retq

```

0x130

0x128

0x120

0x118

0x400549

%rsp

0x118

%rip

0x400557

Control Flow Example

```
00000000000400540 <multstore>:
•
•
400544: callq 400550 <mult2>
400549: mov  %rax, (%rbx)
•
•
```

```
00000000000400550 <mult2>:
400550: mov  %rdi, %rax
•
•
400557: retq
```

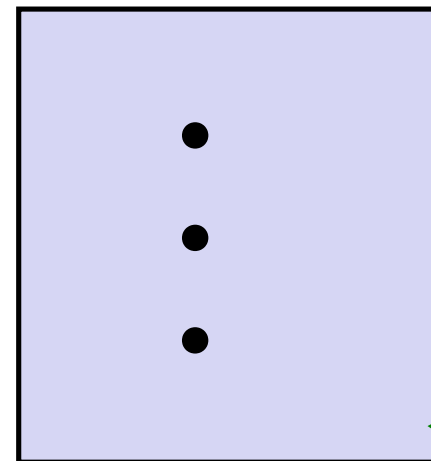
0x130

0x128

0x120

%rsp

%rip



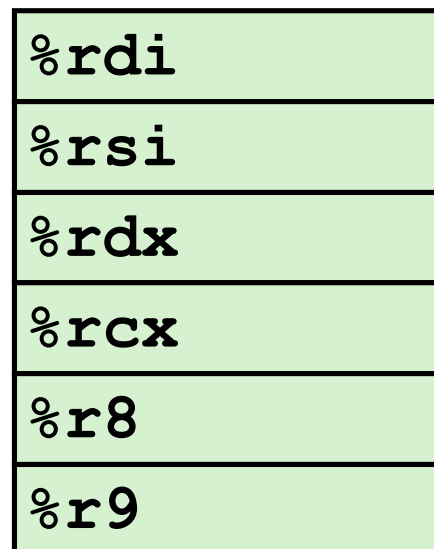
0x120

0x400549

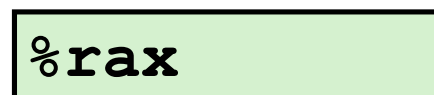
Procedure Data Flow

Registers

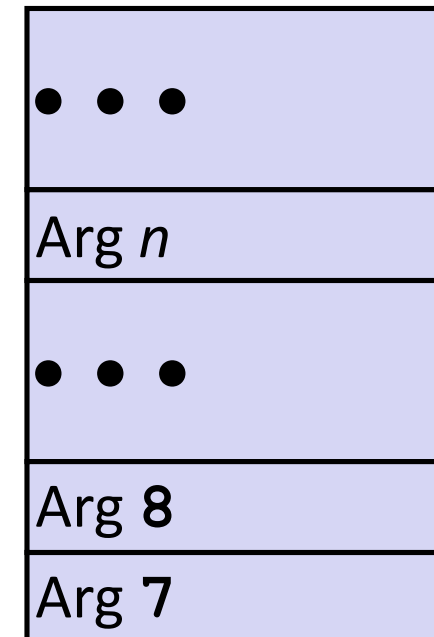
- First 6 arguments



- Return value



Stack



- Only allocate stack space when needed

Data Flow Examples

```
void multstore
(long x, long y, long *dest) {
    long t = mult2(x, y);
    *dest = t;
}
```

```
0000000000400540 <multstore>:
    # x in %rdi, y in %rsi, dest in %rdx
    ...
400541: mov     %rdx,%rbx        # Save dest
400544: callq  400550 <mult2>    # mult2(x,y)
    # t in %rax
400549: mov     %rax,(%rbx)     # Save at dest
    ...
```

```
long mult2
(long a, long b)
{
    long s = a * b;
    return s;
}
```

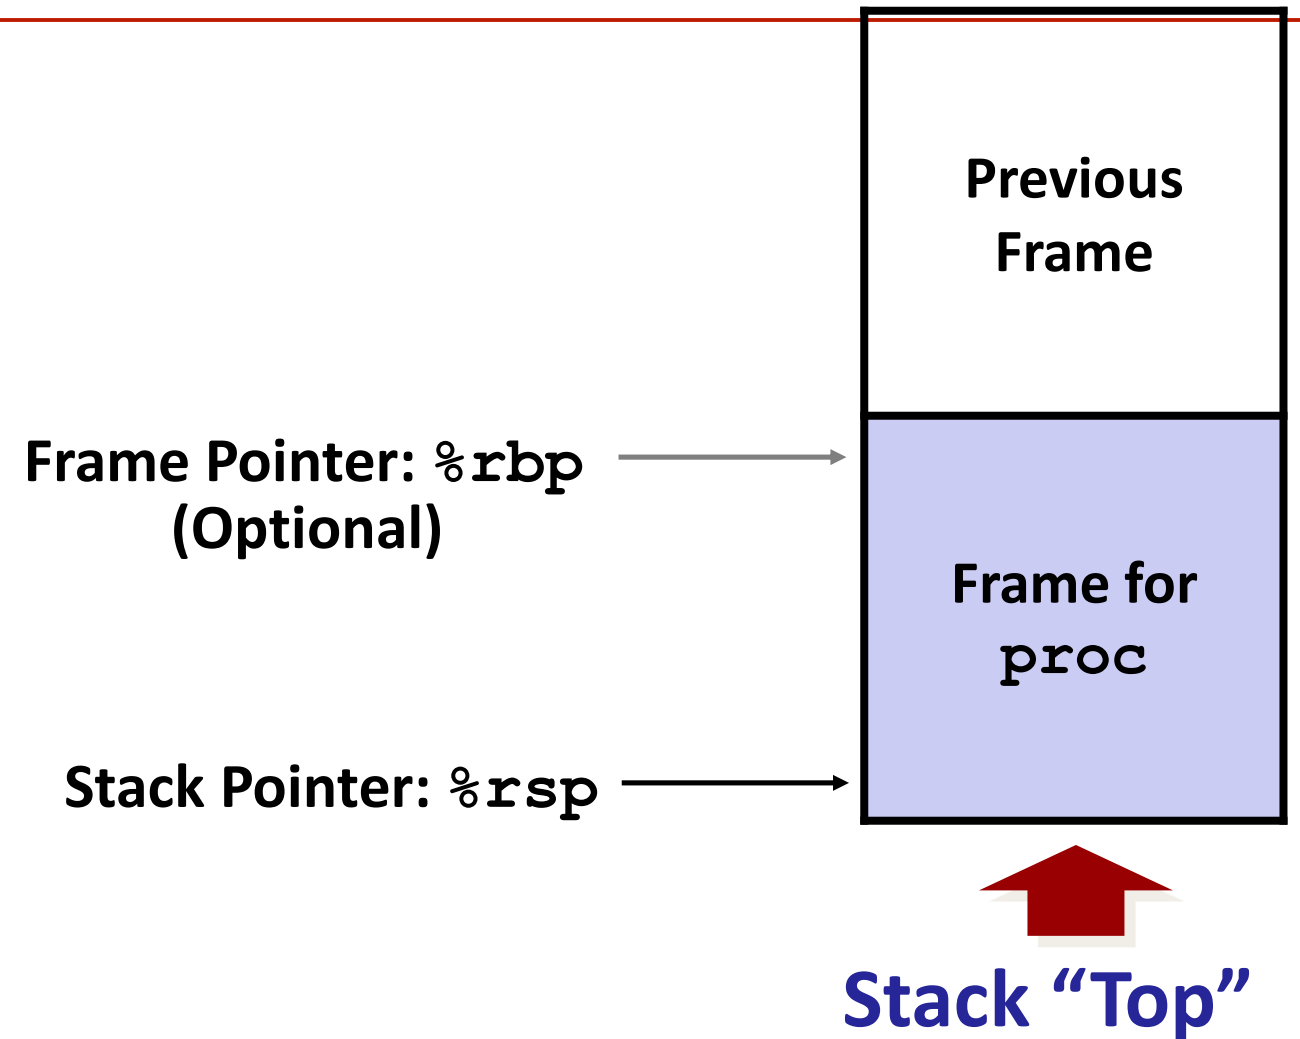
```
0000000000400550 <mult2>:
    # a in %rdi, b in %rsi
400550: mov     %rdi,%rax      # a
400553: imul   %rsi,%rax      # a * b
    # s in %rax
400557: retq                               # Return
```


Today

- Control
 - Control: Condition codes
 - Conditional branches
 - Loops
 - Switch Statements
- Procedures
 - Stack Structure
 - Passing control & data
 - **Managing local data**
 - Illustration of Recursion
- Buffer Overflow (Attacks)
- SSE, SIMD, FP

Stack Frames

- Contents
 - Return information
 - Local storage (if needed)
 - Temporary space (if needed)
- Management
 - Space allocated when enter procedure
 - "Set-up" code
 - Includes push by `call` instruction
 - Deallocated when return
 - "Finish" code
 - Includes pop by `ret` instruction



Example: `incr`

```
long incr(long *p, long val) {
    long x = *p;
    long y = x + val;
    *p = y;
    return x;
}
```

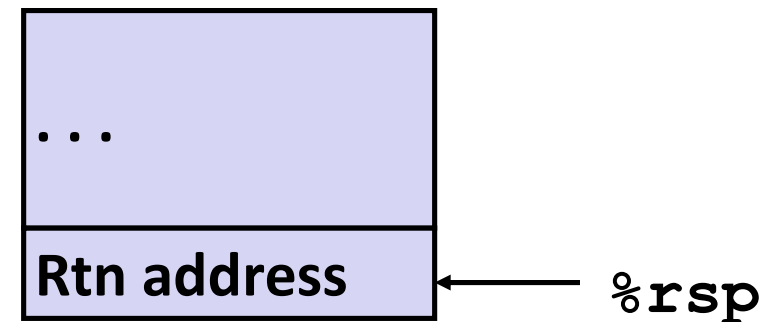
```
incr:
    movq    (%rdi), %rax
    addq    %rax, %rsi
    movq    %rsi, (%rdi)
    ret
```

| Register | Use(s) |
|-------------------|--|
| <code>%rdi</code> | Argument <code>p</code> |
| <code>%rsi</code> | Argument <code>val</code> , <code>y</code> |
| <code>%rax</code> | <code>x</code> , Return value |

Example: Calling `incr` #1

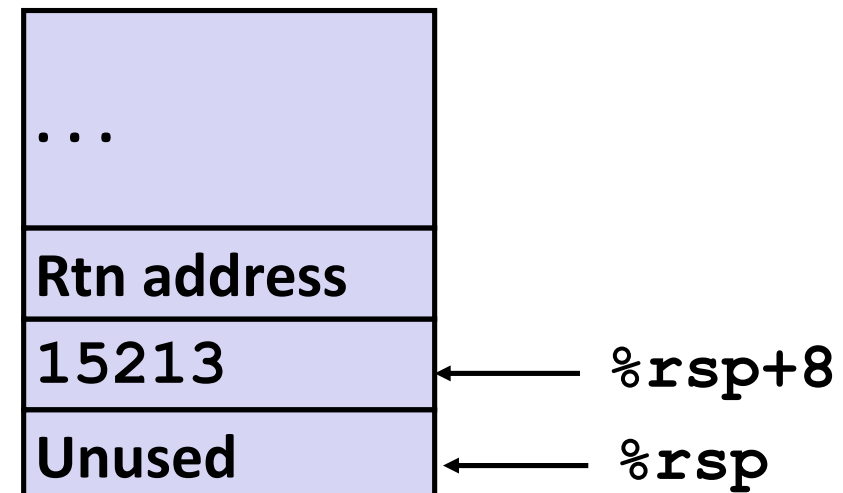
```
long call_incr() {
    long v1 = 15213;
    long v2 = incr(&v1, 3000);
    return v1+v2;
}
```

Initial Stack Structure



```
call_incr:
    subq    $16, %rsp
    movq    $15213, 8(%rsp)
    movl    $3000, %esi
    leaq    8(%rsp), %rdi
    call    incr
    addq    8(%rsp), %rax
    addq    $16, %rsp
    ret
```

Resulting Stack Structure

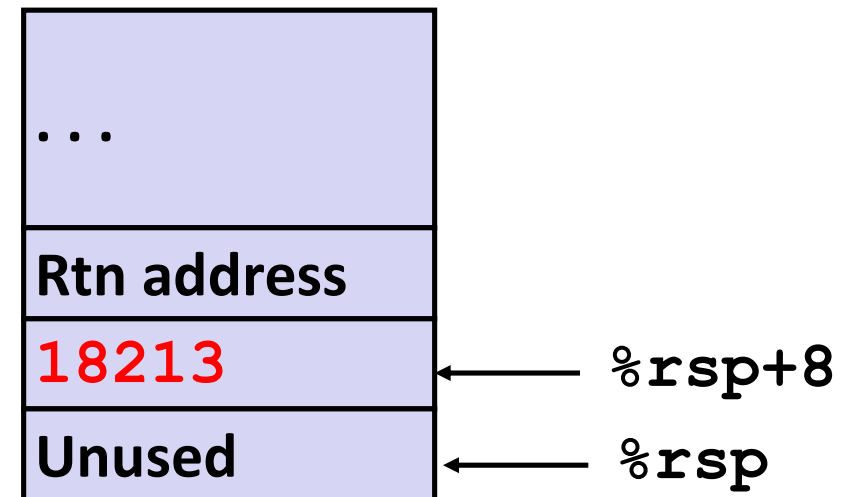


Example: Calling `incr` #2

```
long call_incr() {
    long v1 = 15213;
    long v2 = incr(&v1, 3000);
    return v1+v2;
}
```

```
call_incr:
    subq    $16, %rsp
    movq    $15213, 8(%rsp)
    movl    $3000, %esi
    leaq    8(%rsp), %rdi
    call    incr
    addq    8(%rsp), %rax
    addq    $16, %rsp
    ret
```

Stack Structure



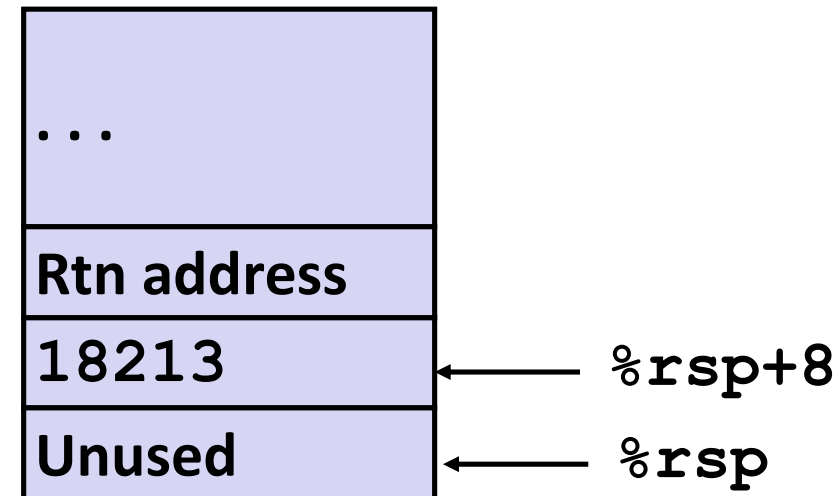
| Register | Use(s) |
|-------------------|----------------------|
| <code>%rdi</code> | <code>&v1</code> |
| <code>%rsi</code> | 3000 |

Example: Calling `incr` #3

```
long call_incr() {
    long v1 = 15213;
    long v2 = incr(&v1, 3000);
    return v1+v2;
}
```

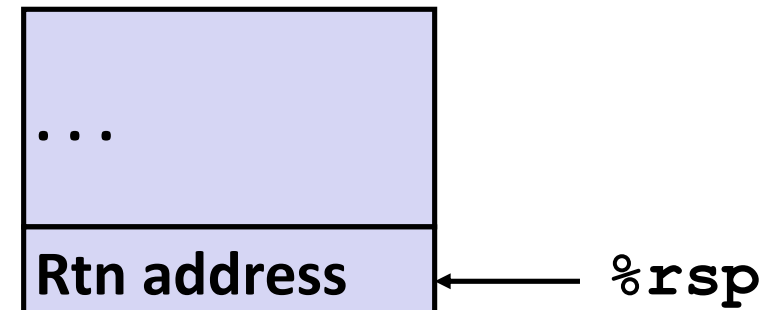
```
call_incr:
    subq    $16, %rsp
    movq    $15213, 8(%rsp)
    movl    $3000, %esi
    leaq    8(%rsp), %rdi
    call    incr
    addq    8(%rsp), %rax
    addq    $16, %rsp
    ret
```

Stack Structure



| Register | Use(s) |
|-------------------|--------------|
| <code>%rax</code> | Return value |

Updated Stack Structure

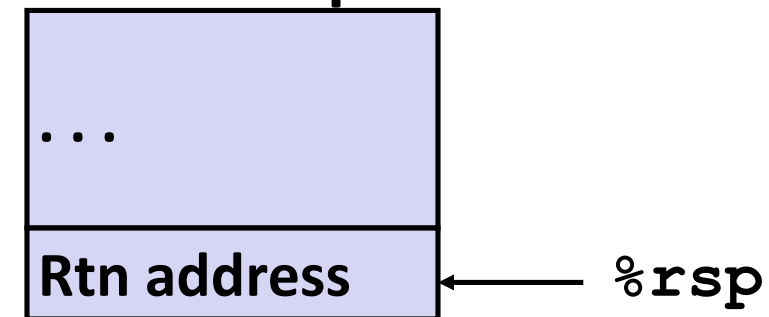


Example: Calling `incr` #4

```
long call_incr() {
    long v1 = 15213;
    long v2 = incr(&v1, 3000);
    return v1+v2;
}
```

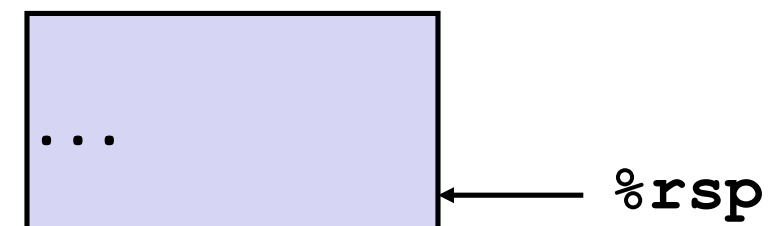
```
call_incr:
    subq    $16, %rsp
    movq    $15213, 8(%rsp)
    movl    $3000, %esi
    leaq    8(%rsp), %rdi
    call   incr
    addq    8(%rsp), %rax
    addq    $16, %rsp
    ret
```

Updated Stack Structure



| Register | Use(s) |
|-------------------|--------------|
| <code>%rax</code> | Return value |

Final Stack Structure



Register Saving Conventions

- When procedure **yoo** calls **who**:
 - **yoo** is the *caller*; **who** is the *callee*
- Can register be used for temporary storage?

```
yoo:
  . . .
  movq $15213, %rdx
  call who
  addq %rdx, %rax
  . . .
  ret
```

```
who:
  . . .
  subq $18213, %rdx
  . . .
  ret
```

- Contents of register **%rdx** overwritten by **who**. This could be trouble.

■ Conventions

- *“Caller Saved”*: Caller saves temporary values in its frame before the call
- *“Callee Saved”*: Callee saves temporary values in its frame before using, and restores them before returning to caller

x86-64 Linux Register Usage #1

- **%rax**
 - Return value
 - Also caller-saved
 - Can be modified by procedure
- **%rdi, ..., %r9**
 - Arguments
 - Also caller-saved
 - Can be modified by procedure
- **%r10, %r11**
 - Caller-saved
 - Can be modified by procedure

Return value

%rax

Arguments

%rdi

%rsi

%rdx

%rcx

%r8

%r9

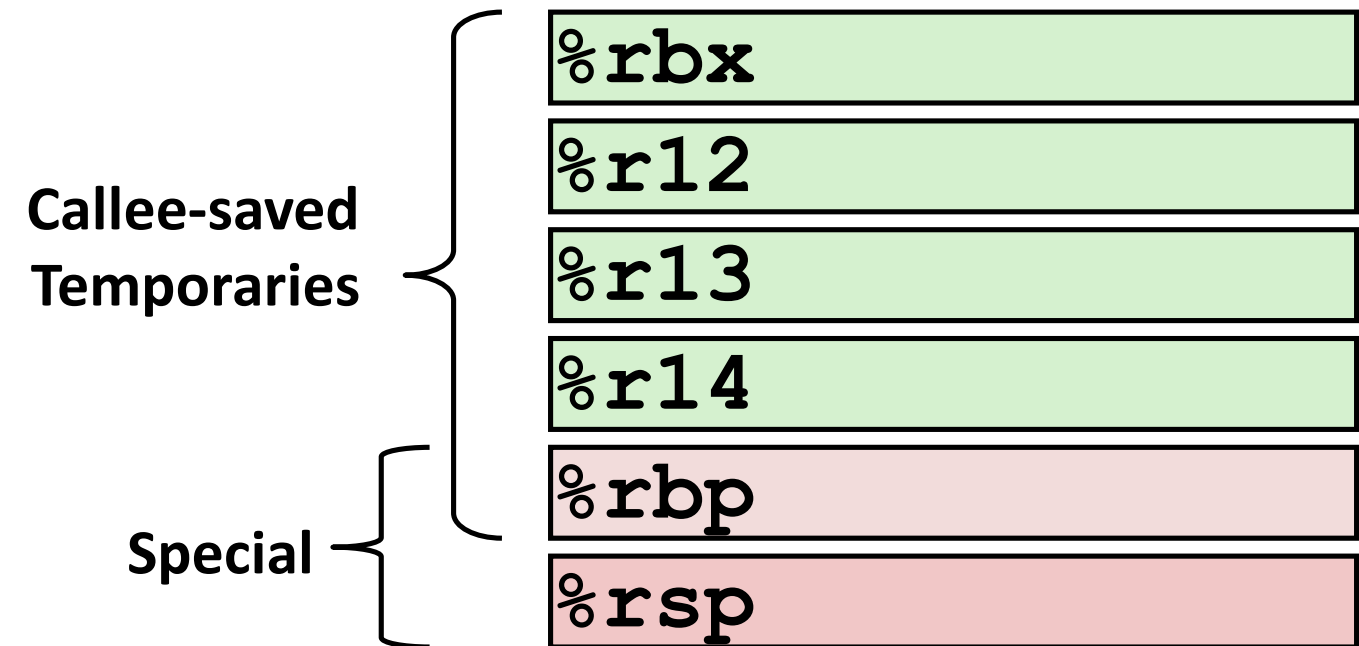
Caller-saved
temporaries

%r10

%r11

x86-64 Linux Register Usage #2

- **%rbx, %r12, %r13, %r14**
 - Callee-saved
 - Callee must save & restore
- **%rbp**
 - Callee-saved
 - Callee must save & restore
 - May be used as frame pointer
 - Can mix & match
- **%rsp**
 - Special form of callee save
 - Restored to original value upon exit from procedure



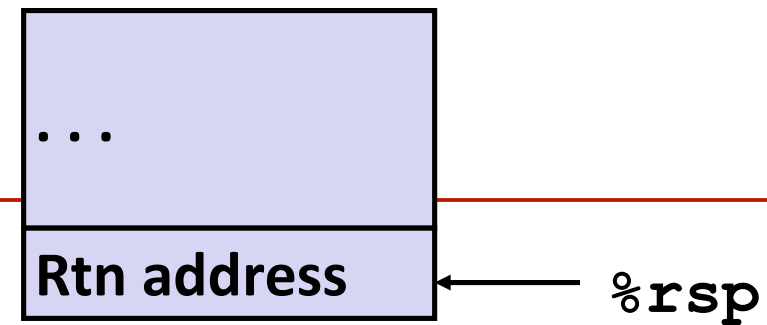
Callee-Saved Example

```
long call_incr2(long x) {
    long v1 = 15213;
    long v2 = incr(&v1, 3000);
    return x+v2;
}
```

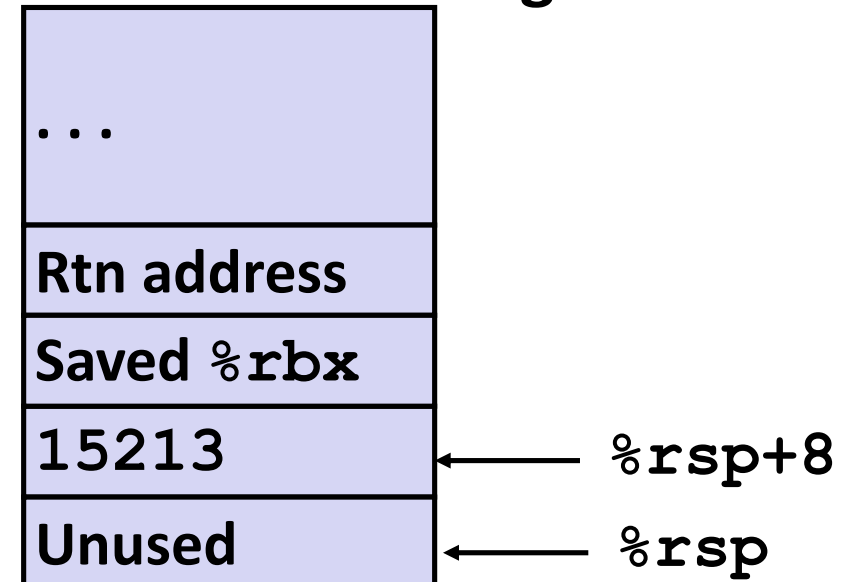
call_incr2:

```
pushq    %rbx
subq     $16, %rsp
movq     %rdi, %rbx
movq     $15213, 8(%rsp)
movl     $3000, %esi
leaq    8(%rsp), %rdi
call     incr
addq     %rbx, %rax
addq     $16, %rsp
popq     %rbx
ret
```

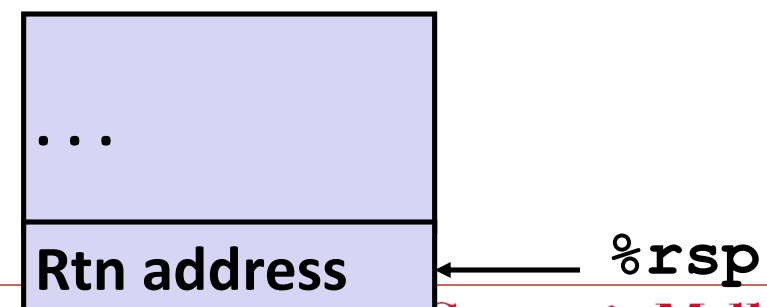
Initial Stack Structure



Resulting Stack Structure



Pre-return Stack Structure



Today

- Control
 - Control: Condition codes
 - Conditional branches
 - Loops
 - Switch Statements
- Procedures
 - Stack Structure
 - Passing control & data
 - Managing local data
 - **Illustration of Recursion**
- Buffer Overflow (Attacks)

Recursive Function Terminal Case

```

/* Recursive popcount */
long pcount_r(unsigned long x) {
    if (x == 0)
        return 0;
    else
        return (x & 1)
            + pcount_r(x >> 1);
}

```

```

pcount_r:
    movl    $0, %eax
    testq   %rdi, %rdi
    je     .L6
    pushq   %rbx
    movq    %rdi, %rbx
    andl    $1, %ebx
    shrq    %rdi
    call    pcount_r
    addq    %rbx, %rax
    popq    %rbx
.L6:
    rep; ret

```

| Register | Use(s) | Type |
|----------|--------------|--------------|
| %rdi | x | Argument |
| %rax | Return value | Return value |

Recursive Function Register Save

```

/* Recursive popcount */
long pcount_r(unsigned long x) {
    if (x == 0)
        return 0;
    else
        return (x & 1)
            + pcount_r(x >> 1);
}

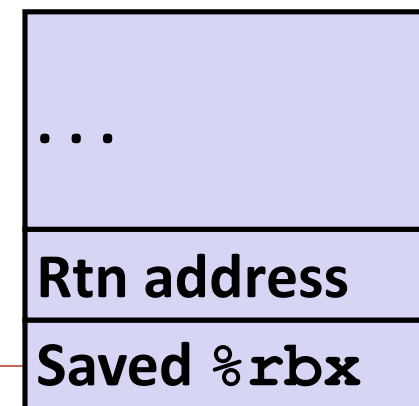
```

```

pcount_r:
    movl    $0, %eax
    testq   %rdi, %rdi
    je     .L6
    pushq   %rbx
    movq   %rdi, %rbx
    andl   $1, %ebx
    shrq   %rdi
    call   pcount_r
    addq   %rbx, %rax
    popq   %rbx
.L6:
    rep; ret

```

| Register | Use(s) | Type |
|----------|--------|----------|
| %rdi | x | Argument |



Recursive Function Call Setup

```

/* Recursive popcount */
long pcount_r(unsigned long x) {
    if (x == 0)
        return 0;
    else
        return (x & 1)
            + pcount_r(x >> 1);
}

```

```

pcount_r:
    movl    $0, %eax
    testq   %rdi, %rdi
    je      .L6
    pushq   %rbx
    movq    %rdi, %rbx
    andl    $1, %ebx
    shrq    %rdi
    call    pcount_r
    addq    %rbx, %rax
    popq    %rbx
.L6:
    rep; ret

```

| Register | Use(s) | Type |
|----------|--------|---------------|
| %rdi | x >> 1 | Rec. argument |
| %rbx | x & 1 | Callee-saved |

Recursive Function Call

```

/* Recursive popcount */
long pcount_r(unsigned long x) {
    if (x == 0)
        return 0;
    else
        return (x & 1)
            + pcount_r(x >> 1);
}

```

```

pcount_r:
    movl    $0, %eax
    testq   %rdi, %rdi
    je     .L6
    pushq   %rbx
    movq    %rdi, %rbx
    andl    $1, %ebx
    shrq    %rdi
    call    pcount_r
    addq    %rbx, %rax
    popq    %rbx
.L6:
    rep; ret

```

| Register | Use(s) | Type |
|----------|-----------------------------|--------------|
| %rbx | x & 1 | Callee-saved |
| %rax | Recursive call return value | |

Recursive Function Result

```

/* Recursive popcount */
long pcount_r(unsigned long x) {
    if (x == 0)
        return 0;
    else
        return (x & 1)
            + pcount_r(x >> 1);
}

```

| Register | Use(s) | Type |
|-------------------|------------------------|--------------|
| <code>%rbx</code> | <code>x & 1</code> | Callee-saved |
| <code>%rax</code> | Return value | |

```

pcount_r:
    movl    $0, %eax
    testq   %rdi, %rdi
    je     .L6
    pushq   %rbx
    movq    %rdi, %rbx
    andl    $1, %ebx
    shrq    %rdi
    call    pcount_r
    addq    %rbx, %rax
    popq    %rbx
.L6:
    rep; ret

```

Recursive Function Completion

```

/* Recursive popcount */
long pcount_r(unsigned long x) {
    if (x == 0)
        return 0;
    else
        return (x & 1)
            + pcount_r(x >> 1);
}

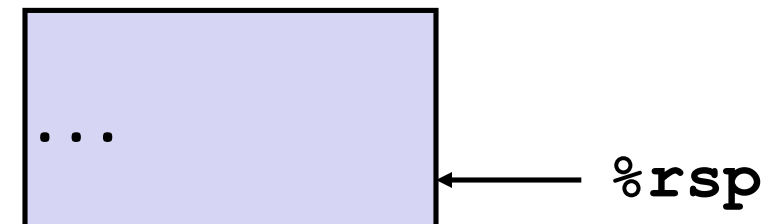
```

| Register | Use(s) | Type |
|-------------------|--------------|--------------|
| <code>%rax</code> | Return value | Return value |

```

pcount_r:
    movl    $0, %eax
    testq   %rdi, %rdi
    je     .L6
    pushq  %rbx
    movq   %rdi, %rbx
    andl   $1, %ebx
    shrq   %rdi
    call   pcount_r
    addq   %rbx, %rax
    popq   %rbx
.L6:
    rep; ret

```



Today

- Control
 - Control: Condition codes
 - Conditional branches
 - Loops
 - Switch Statements
- Procedures
 - Stack Structure
 - Passing control & data
 - Managing local data
 - Illustration of Recursion
- Buffer Overflow (Attacks)
- SSE, SIMD, FP

x86-64 Linux Memory Layout

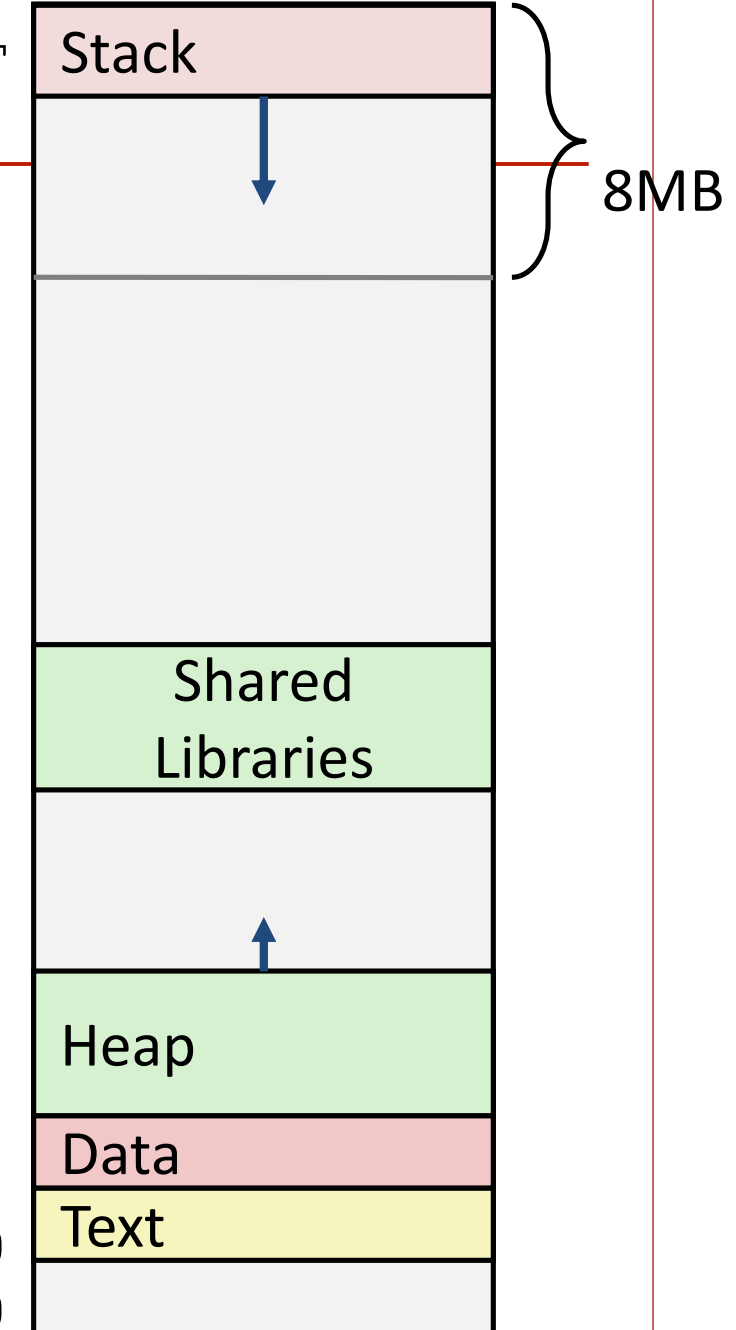
00007FFFFFFF

- 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 */

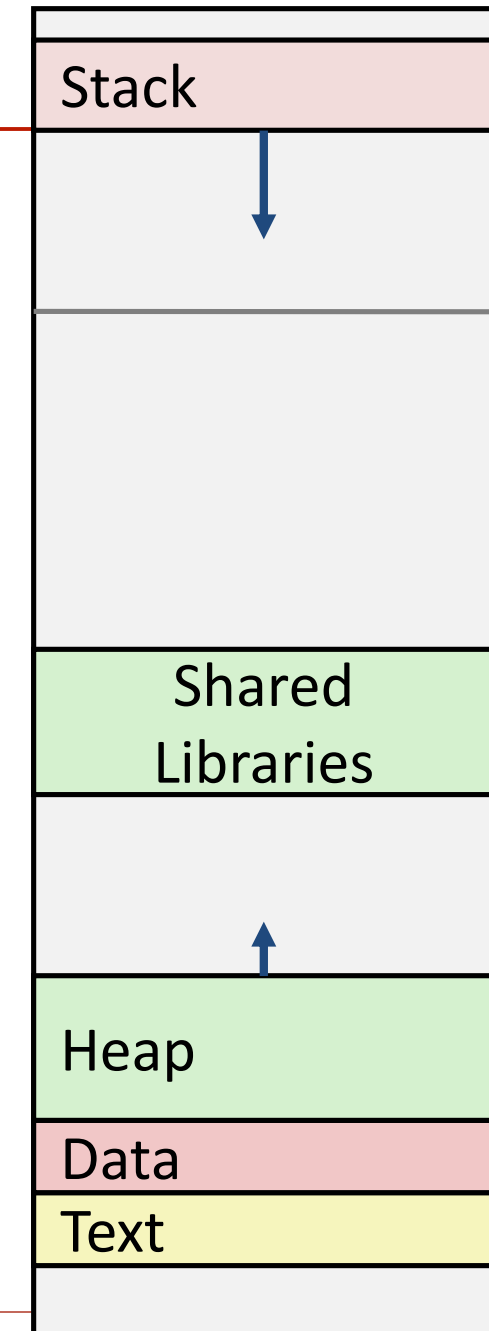
int global = 0;

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 ... */
}

```

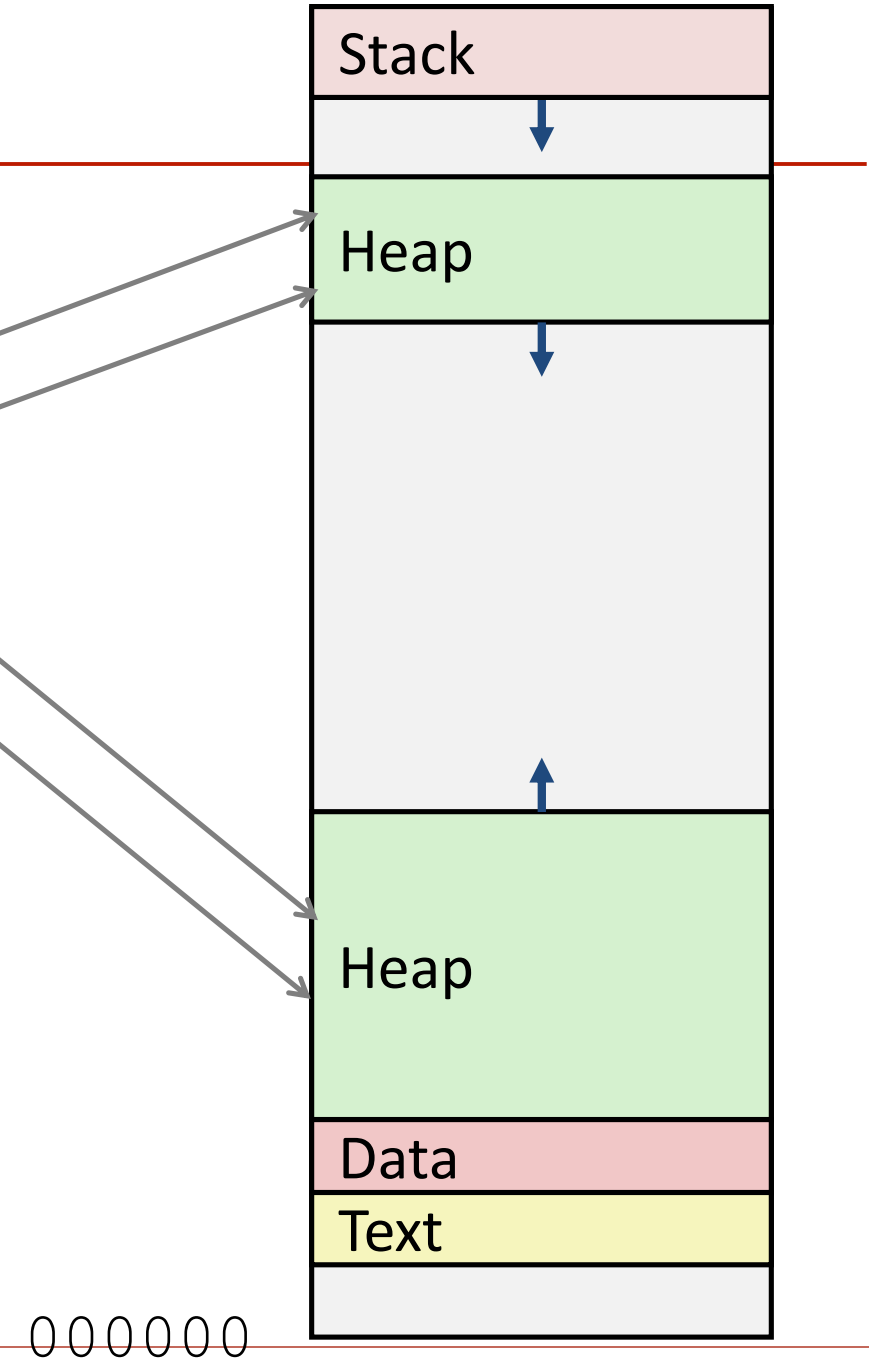
Where does everything go?



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

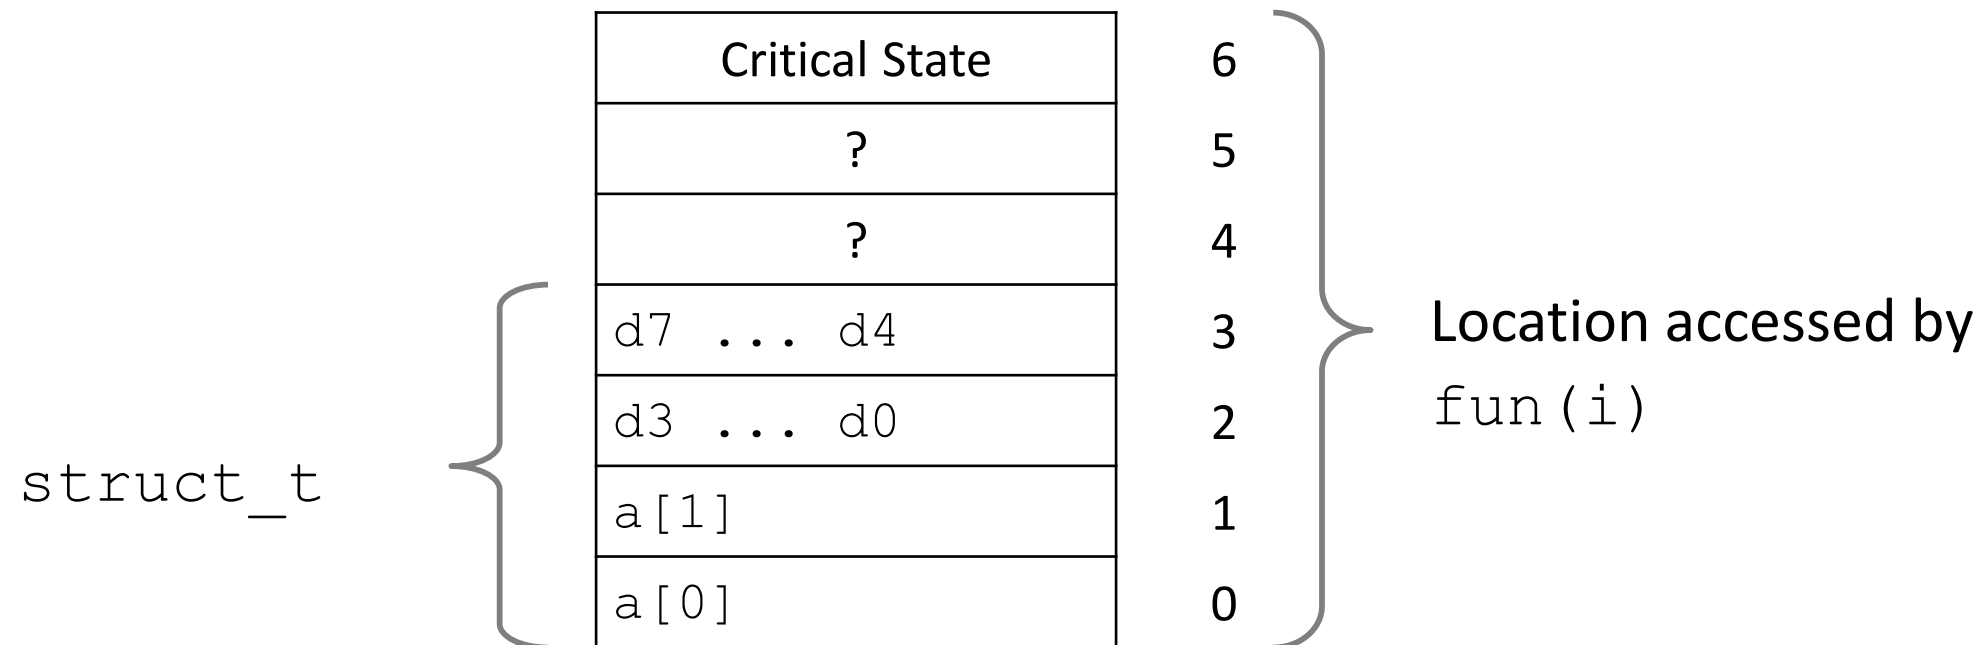
- Control
 - Control: Condition codes
 - Conditional branches
 - Loops
 - Switch Statements
- Procedures
 - Stack Structure
 - Passing control & data
 - Managing local data
 - Illustration of Recursion
- Buffer Overflow (Attacks)
- SSE, SIMD, FP

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);
}
```

← btw, how big
is big enough?

```
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:

```

000000000004006cf <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:

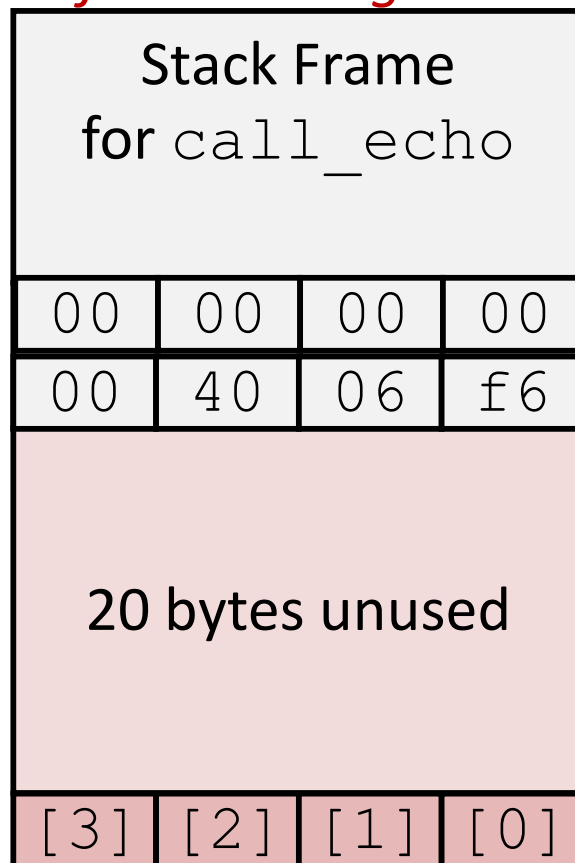
```

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 |

buf ← %rsp

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

call_echo:

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

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

Overflowed buffer, but did not corrupt state

```
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 |

buf ← %rsp

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

call_echo:

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

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

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:

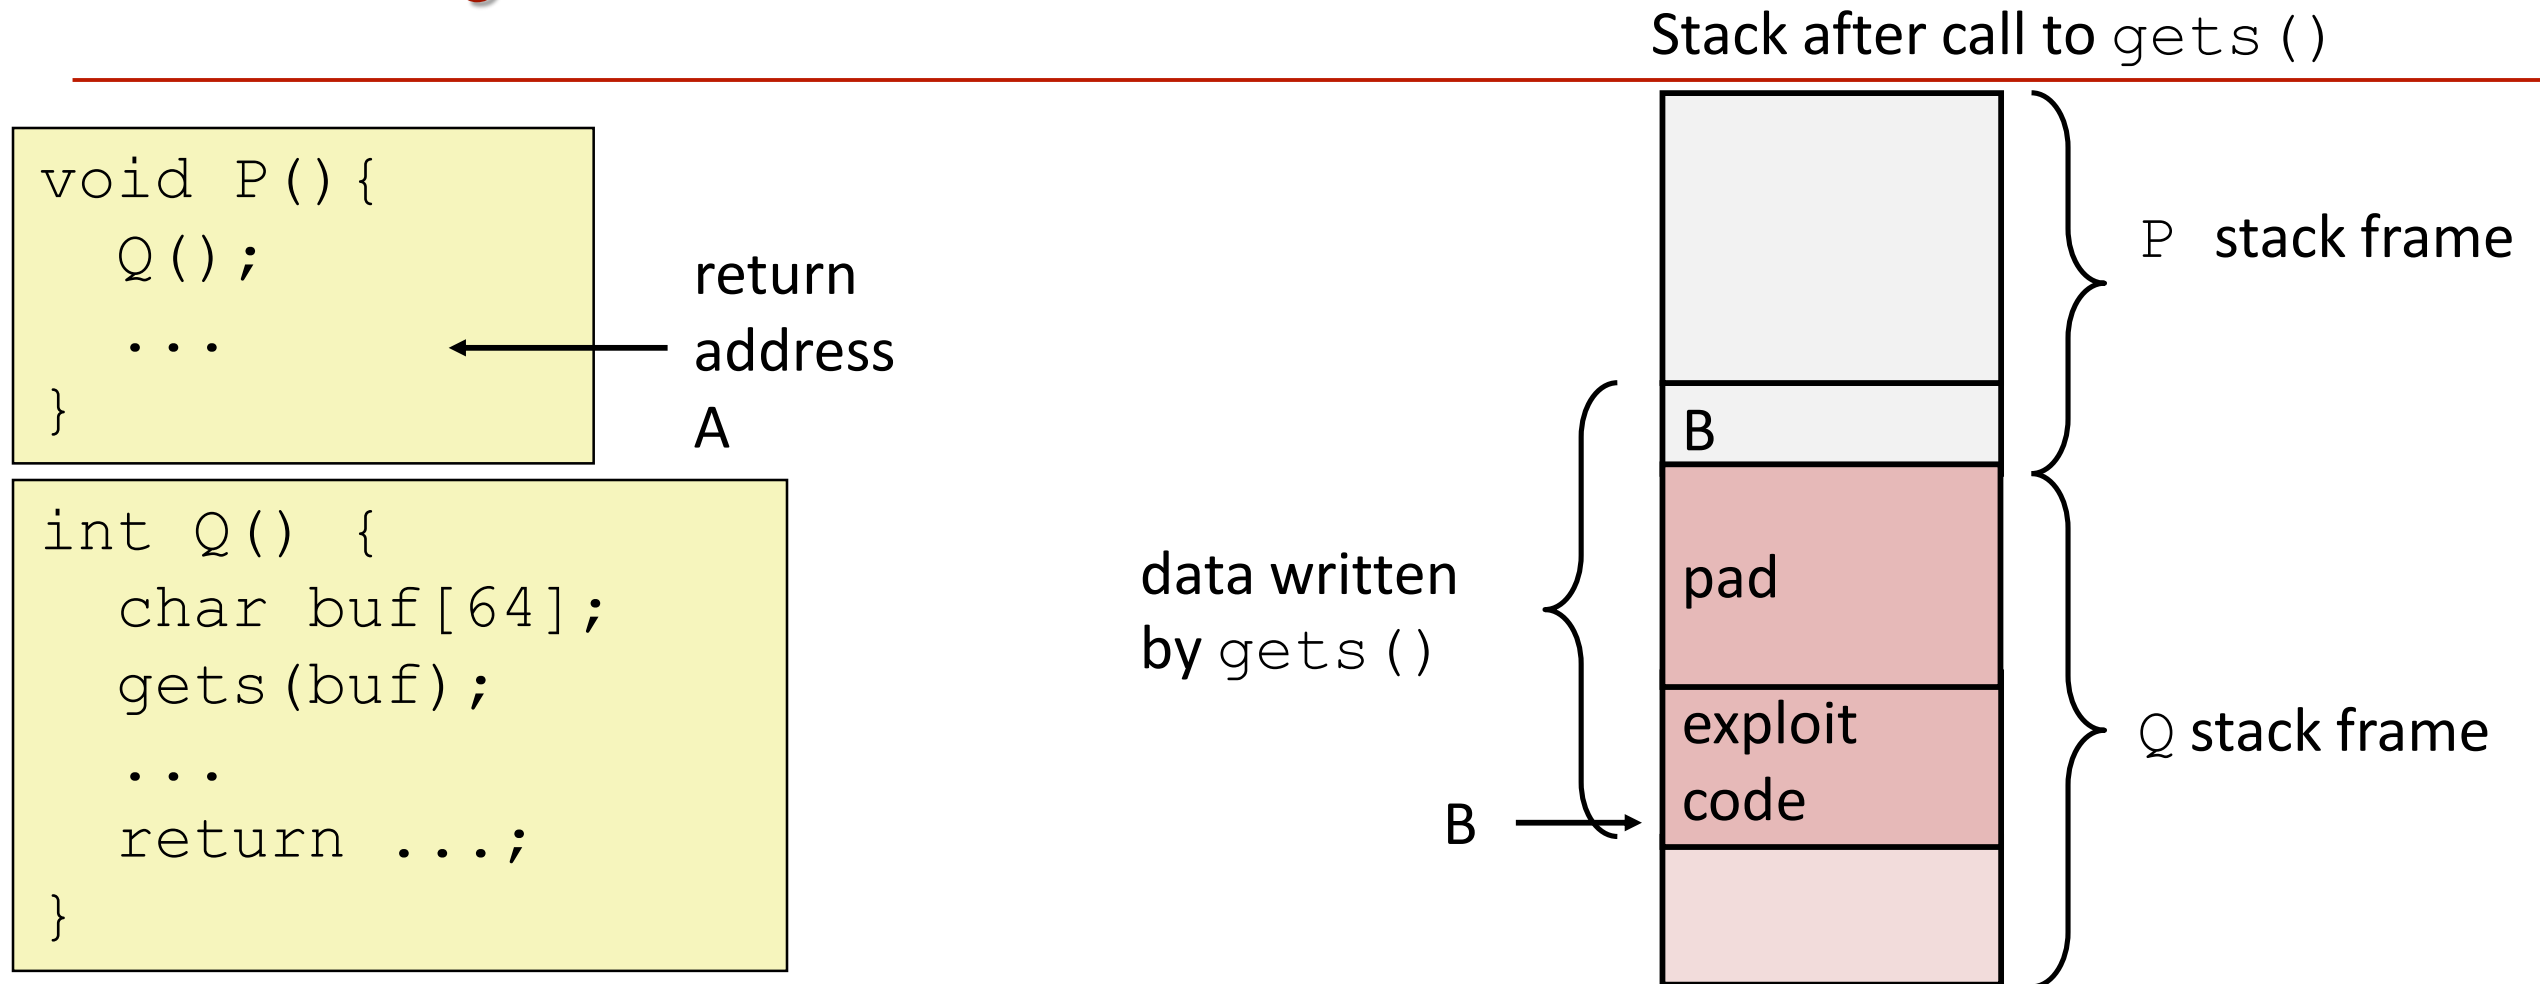
```
. . .
4006f1:    callq    4006cf <echo>
4006f6:    add     $0x8, %rsp
. . .
```

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



- 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 (!)

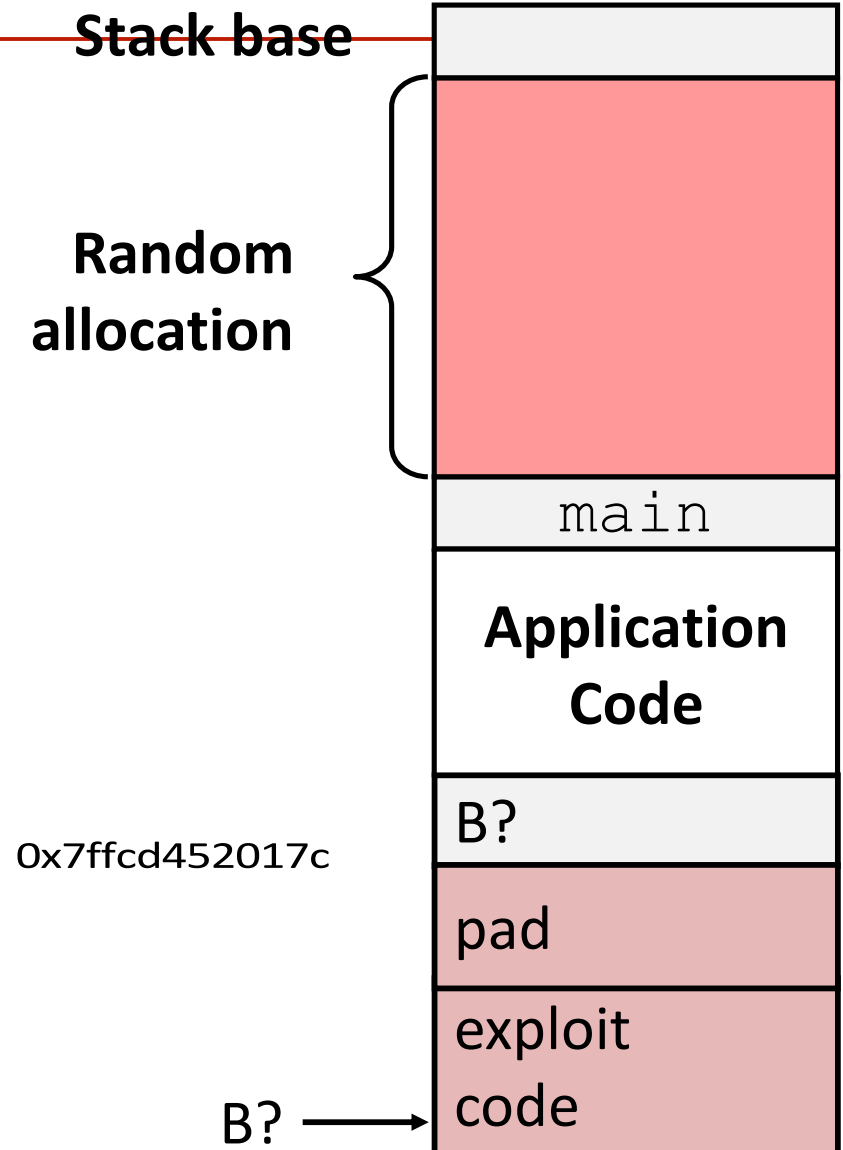
```
/* Echo Line */
void echo()
{
    char buf[4]; /* Way too small! */
    fgets(buf, 4, stdin);
    puts(buf);
}
```

- 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

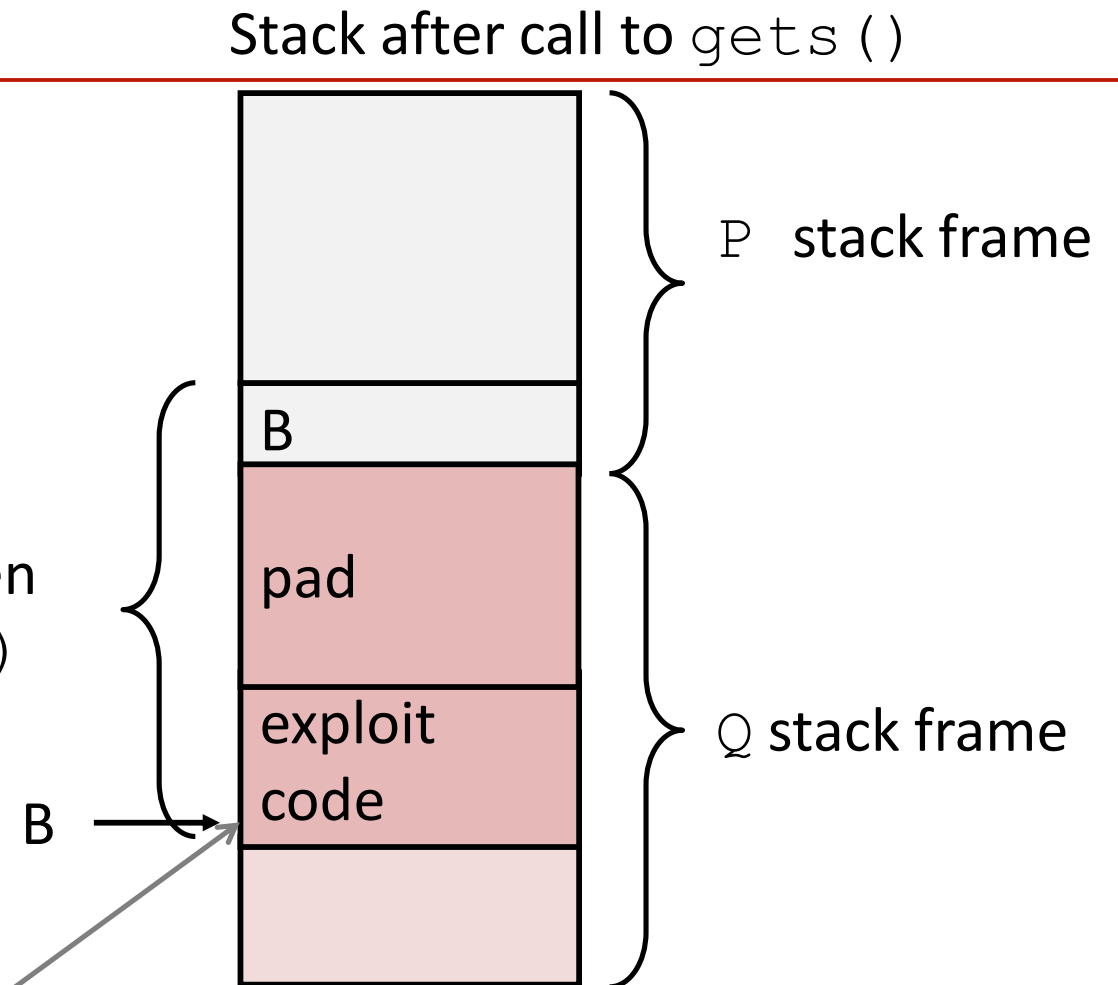
local 0x7ffe4d3be87c 0x7ff75a4f9fc 0x7ffeadb7c80c 0x7ffeaea2fdac 0x7ffcd452017c



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

data written
by `gets()`



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 ***
```

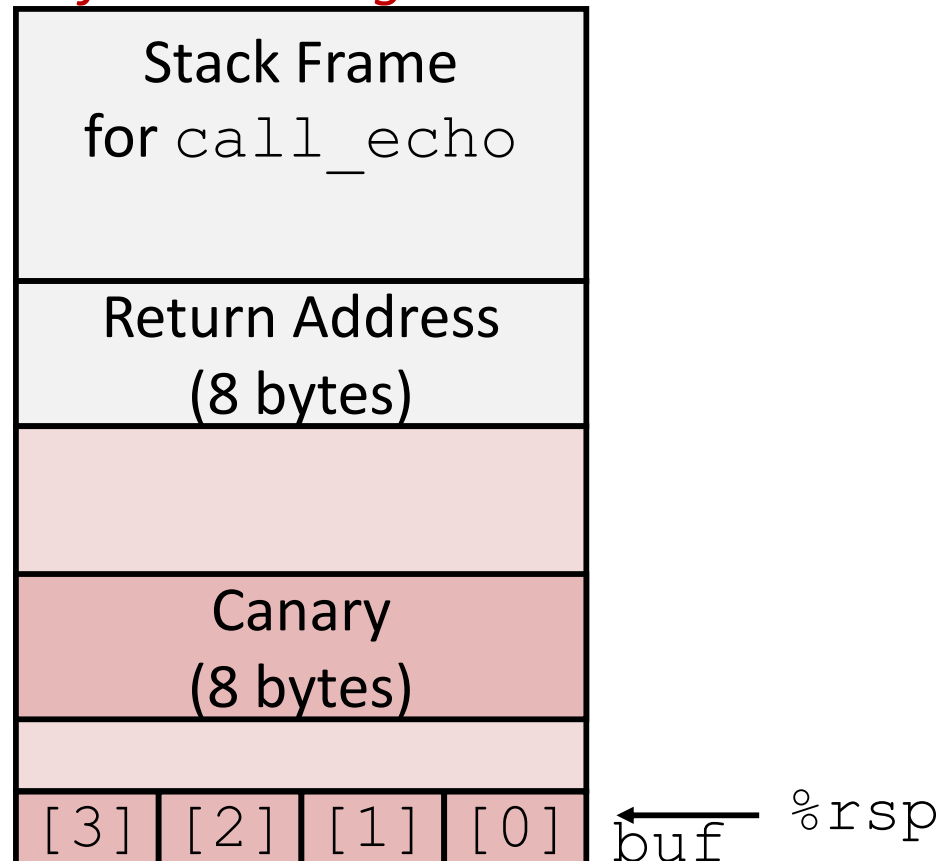
Protected Buffer Disassembly

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
```

Setting Up Canary

Before call to gets

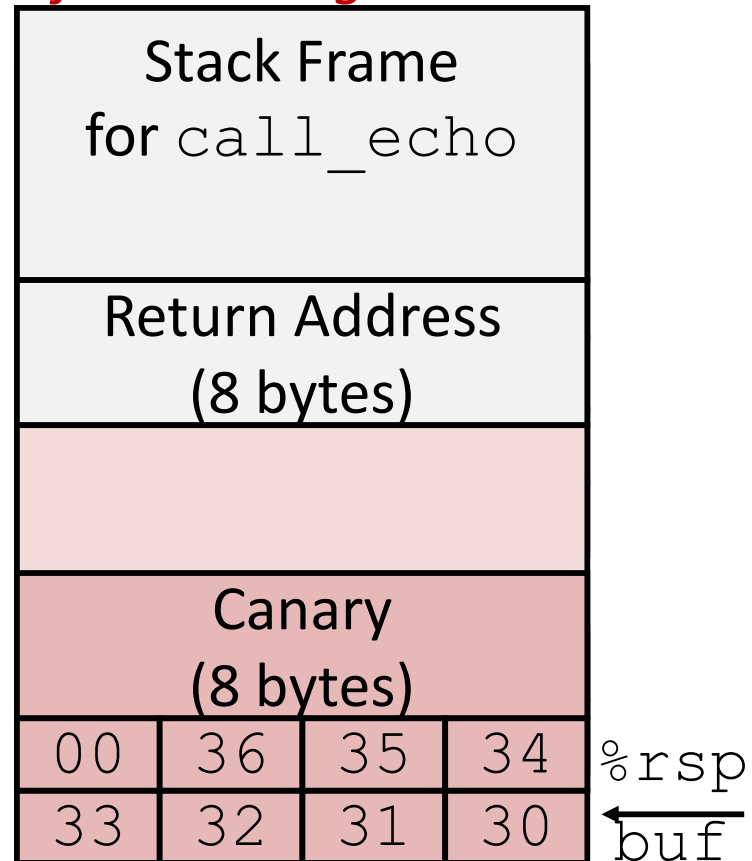


```
/* Echo Line */
void echo()
{
    char buf[4]; /* Way too small! */
    gets(buf);
    puts(buf);
}
```

```
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:
    . . .

```

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
```

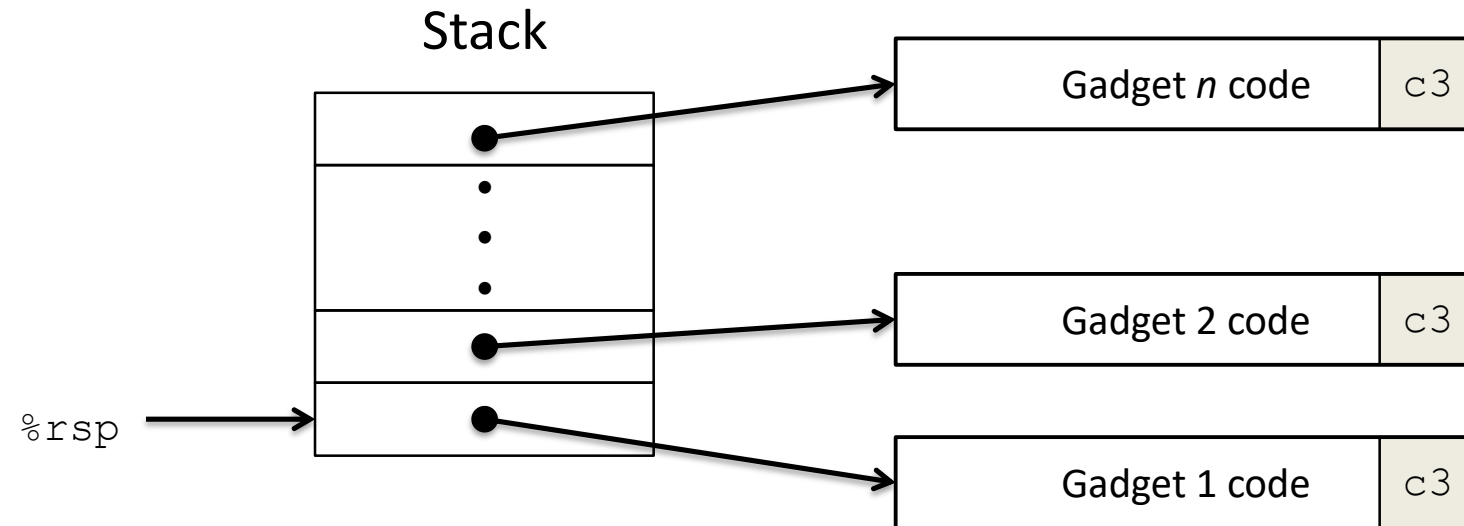
Encodes `movq %rax, %rdi`

`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

Today

- Control
 - Control: Condition codes
 - Conditional branches
 - Loops
 - Switch Statements
- Procedures
 - Stack Structure
 - Passing control & data
 - Managing local data
 - Illustration of Recursion
- Buffer Overflow (Attacks)
- **SSE, SIMD, FP**

Programming with SSE3

XMM Registers

- 16 total, each 16 bytes

- 16 single-byte integers



- 8 16-bit integers



- 4 32-bit integers



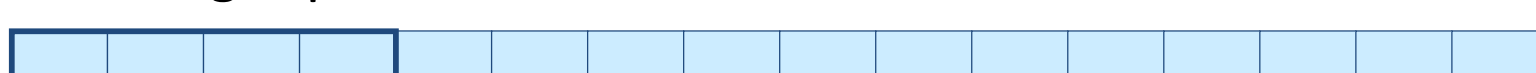
- 4 single-precision floats



- 2 double-precision floats



- 1 single-precision float



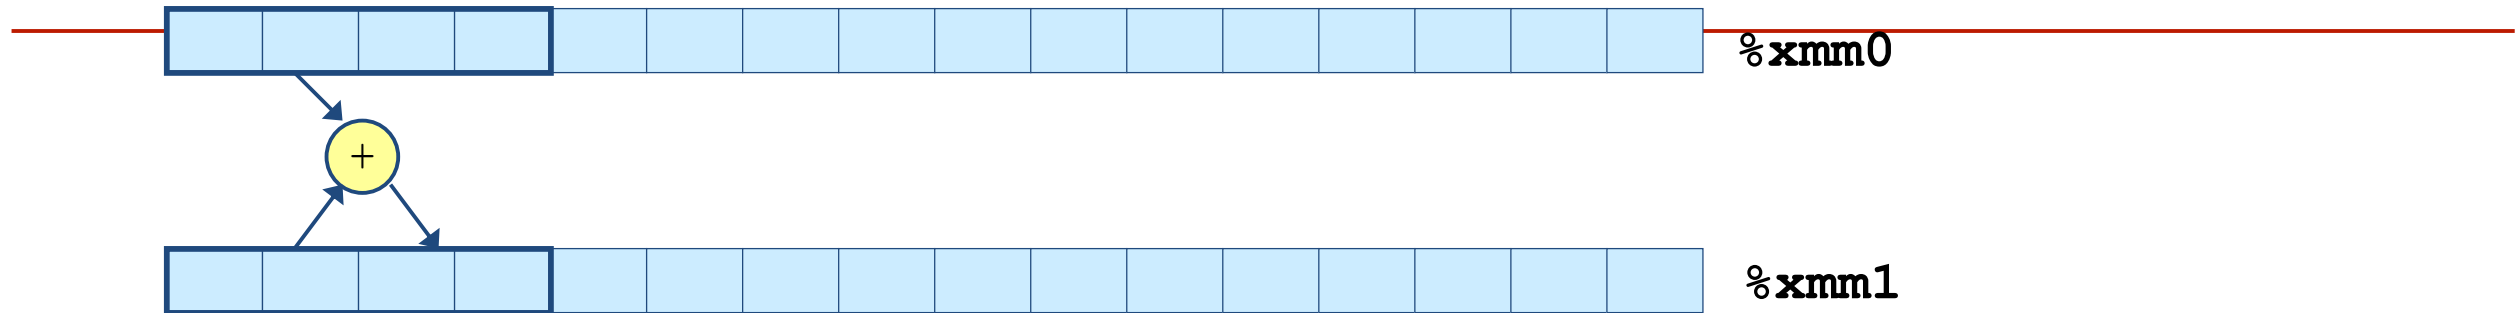
- 1 double-precision float



Scalar & SIMD Operations

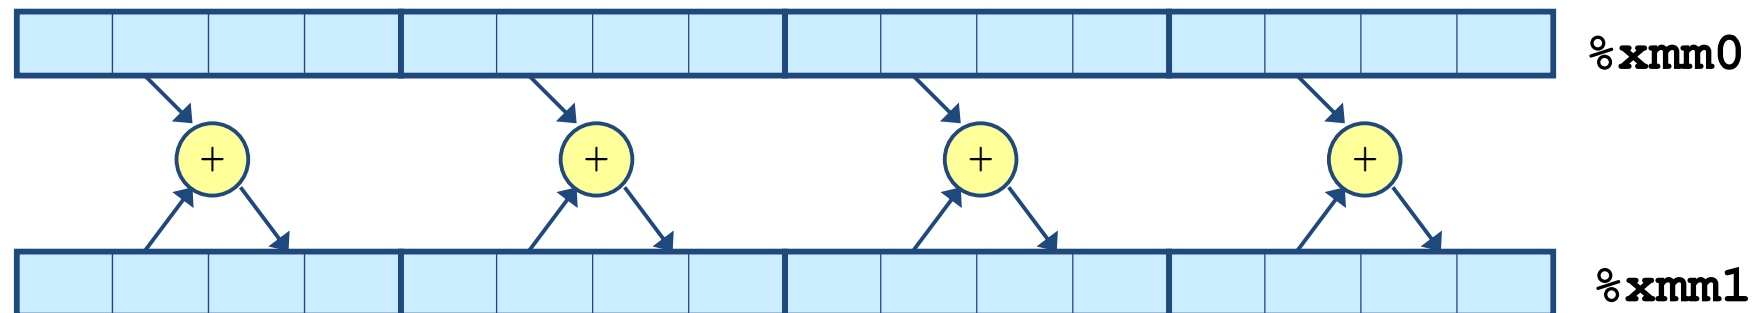
■ Scalar Operations: Single Precision

```
addss %xmm0, %xmm1
```



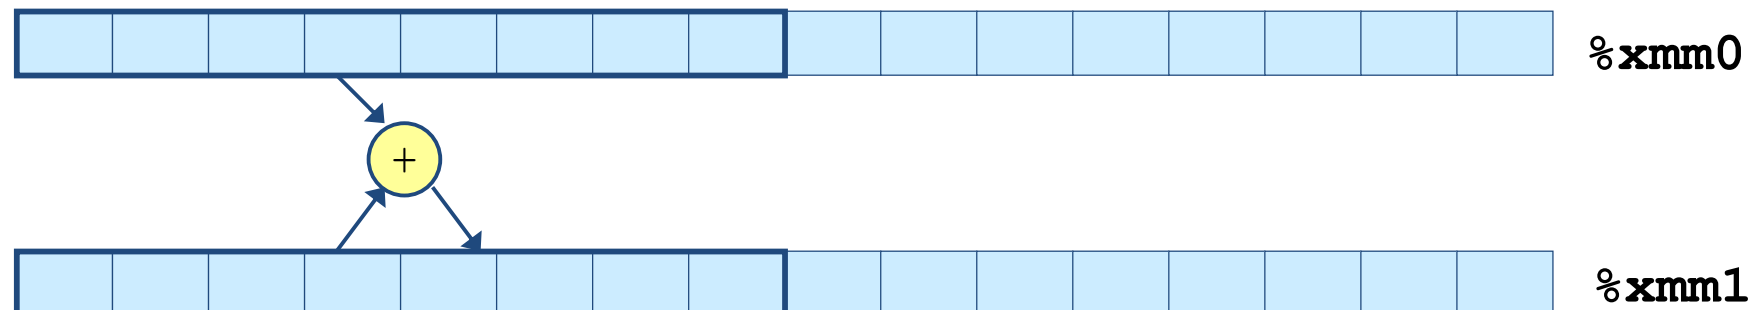
■ SIMD Operations: Single Precision

```
addps %xmm0, %xmm1
```



■ Scalar Operations: Double Precision

```
addsd %xmm0, %xmm1
```



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
```

Summarizing

- Control
 - Jmp instructions implement “goto” like syntax
 - Conditional jmp instructions implement “if” like syntax
 - An “if” that jumps back is a “while loop” that can implement all other loops
- Procedures
 - Registers are used for variable storage and to orchestrate other things
 - Register discipline is important, e.g. return values, caller- and callee-saved, etc
 - Recursion is not a special case
- Buffer Overflow (Attacks)
 - Result from poor coding practices
- SSE, SIMD, FP
 - Specialized instructions implement specialized models

18-600 Foundations of Computer Systems

Lecture 7: "Processor Architecture and Design"

John P. Shen

September 20, 2017

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).

