


Calling Conventions

Hakim Weatherspoon

CS 3410

Computer Science

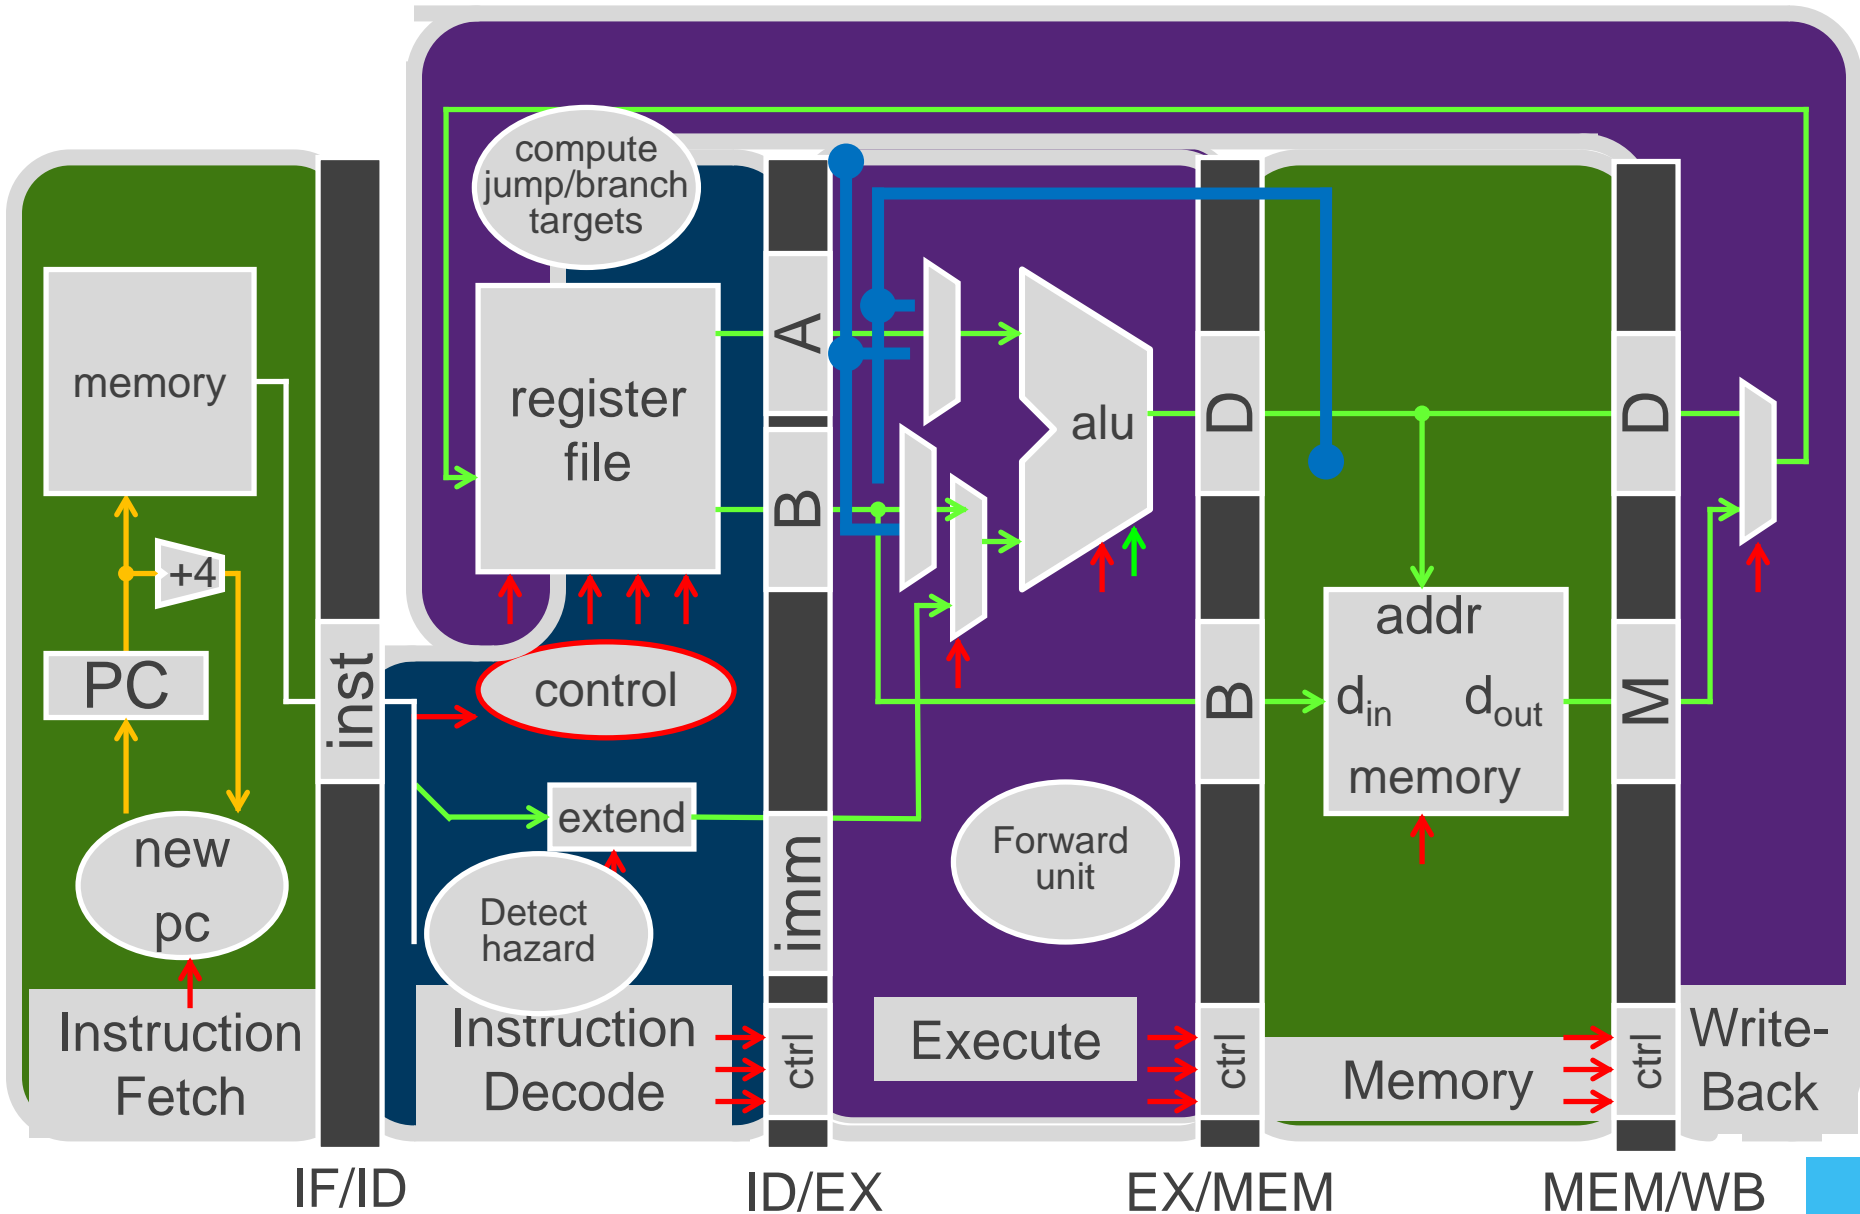
Cornell University



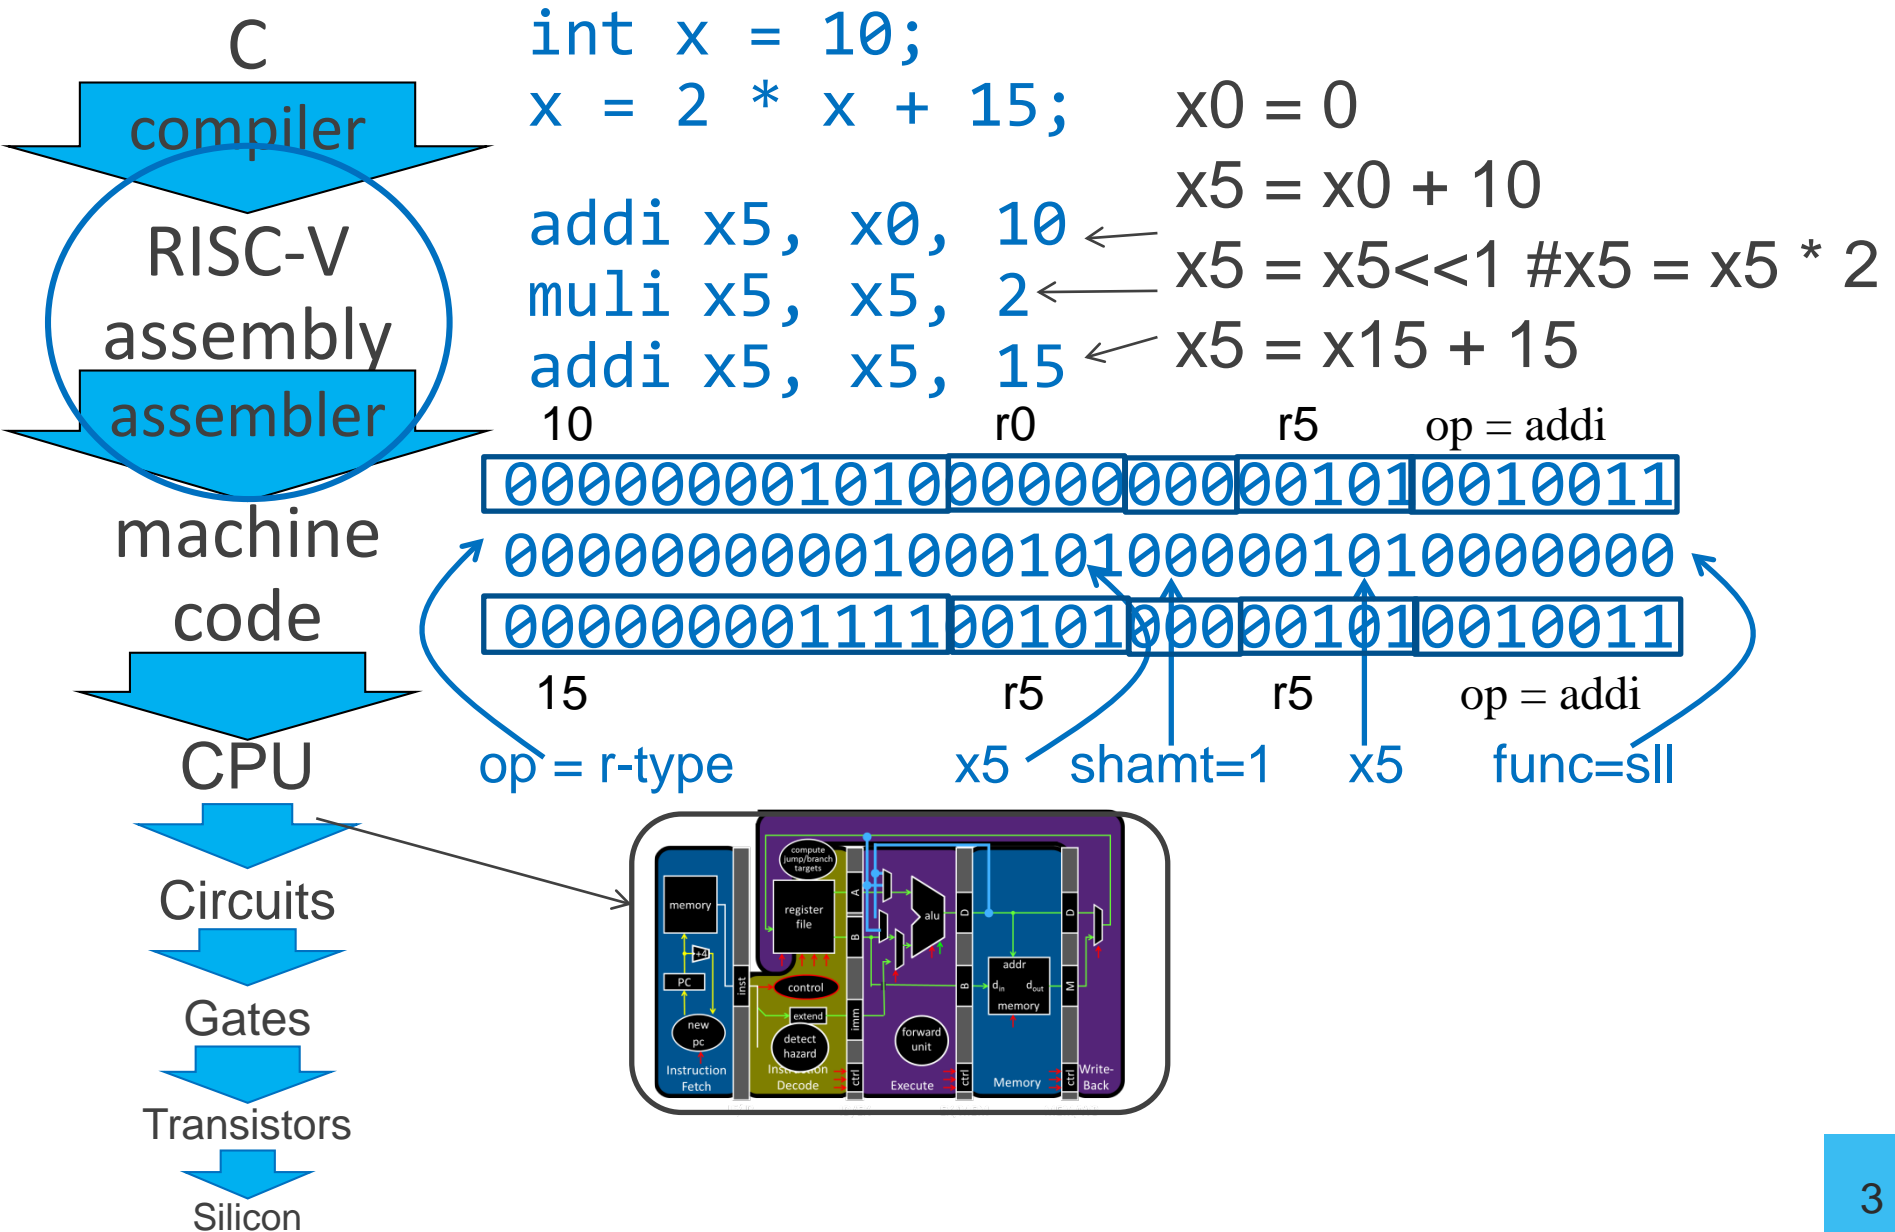
Cornell CIS
COMPUTING AND INFORMATION SCIENCE

[Weatherspoon, Bala, Bracy, McKee and Sierer]


Big Picture: Where are we going?



Big Picture: Where are we going?



Big Picture: Where are we going?



Goals for this week

Calling Convention for Procedure Calls

Enable code to be reused by allowing code snippets to be invoked

Will need a way to

- call the routine (i.e. transfer control to procedure)
- pass arguments
 - fixed length, variable length, recursively
- return to the caller
 - Putting results in a place where caller can find them
- Manage register

Calling Convention for Procedure Calls

Transfer Control

- Caller → Routine
- Routine → Caller

Pass Arguments to and from the routine

- fixed length, variable length, recursively
- Get return value back to the caller

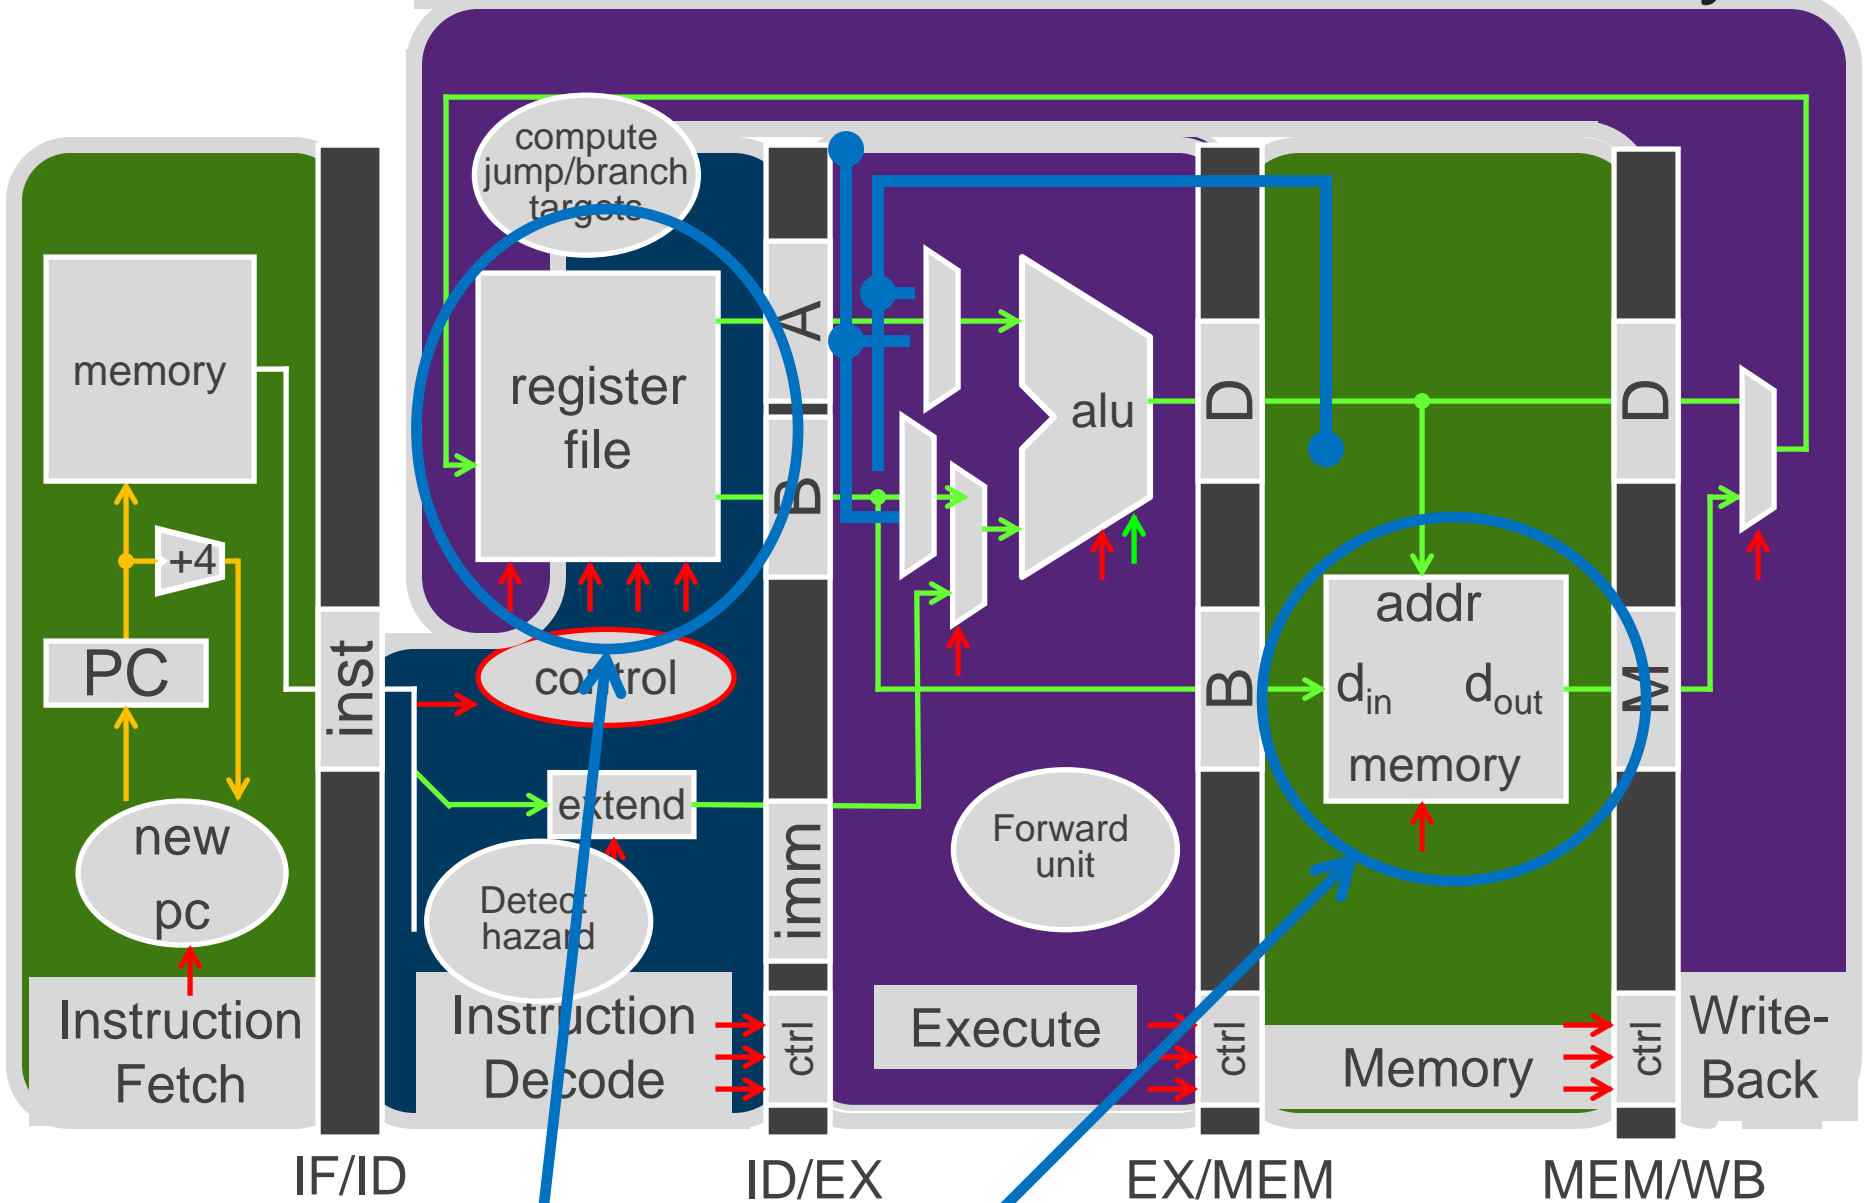
Manage Registers

- Allow each routine to use registers
- Prevent routines from clobbering each others' data

What is a Convention?

Warning: There is no one true RISC-V calling convention.
lecture != book != gcc != spim != web

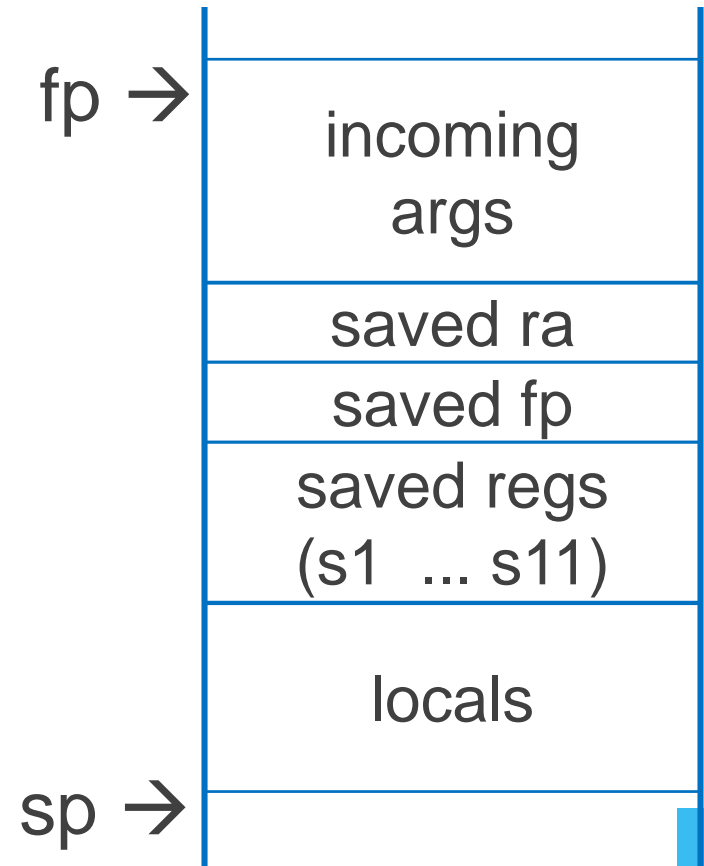
Cheat Sheet and Mental Model for Today



How do we share registers and use memory when making procedure calls?

Cheat Sheet and Mental Model for Today

- **first eight** arg words passed in a0, a1, ... , a7
- remaining arg words passed **in parent's stack frame**
- return value (if any) in a0, a1
- stack frame at sp
 - contains ra (clobbered on JAL to sub-functions)
 - contains local vars (possibly clobbered by sub-functions)
 - contains space for incoming args
- **callee save regs are preserved**
- **caller save regs are not**
- **Global data accessed via \$gp**



RISC-V Register

- Return address: x1 (ra)
- Stack pointer: x2 (sp)
- Frame pointer: x8 (fp/s0)
- First eight arguments: x10-x17 (a0-a7)
- Return result: x10-x11 (a0-a1)
- Callee-save free regs: x18-x27 (s2-s11)
- Caller-save free regs: x5-x7, x28-x31 (t0-t6)
- Global pointer: x3 (gp)
- Thread pointer: x4 (tp)

RISC-V Register Conventions

x0	zero	zero	x15	a5	function arguments
x1	ra	return address	x16	a6	
x2	sp	stack pointer	x17	a7	
x3	gp	global data pointer	x18	s2	saved (callee save)
x4	tp	thread pointer	x19	s3	
x5	t0	temps (caller save)	x20	s4	
x6	t1				
x7	t2				
x8	s0/fp	frame pointer	x22	s6	
x9	s1	saved (callee save)	x23	s7	
x10	a0	function args or return values	x24	s7	
x11	a1				
x12	a2	function arguments	x25	s9	
x13	a3				
x14	a4				
			x26	s10	temps (caller save)
			x27	s11	
			x28	t3	
			x29	t4	
			x30	t5	
			x31	t6	

Calling Convention for Procedure Calls

Transfer Control

- Caller → Routine
- Routine → Caller

Pass Arguments to and from the routine

- fixed length, variable length, recursively
- Get return value back to the caller

Manage Registers

- Allow each routine to use registers
- Prevent routines from clobbering each others' data

What is a Convention?


Warning: There is no one true RISC-V calling convention.
lecture != book != gcc != spim != web

How does a function call work?

```
int main (int argc, char* argv[ ]) {  
    int n = 9;  
    int result = myfn(n);  
}
```

```
int myfn(int n) {  
    int f = 1;  
    int i = 1;  
    int j = n - 1;  
    while(j >= 0) {  
        f *= i;  
        i++;  
        j = n - i;  
    }  
    return f;  
}
```

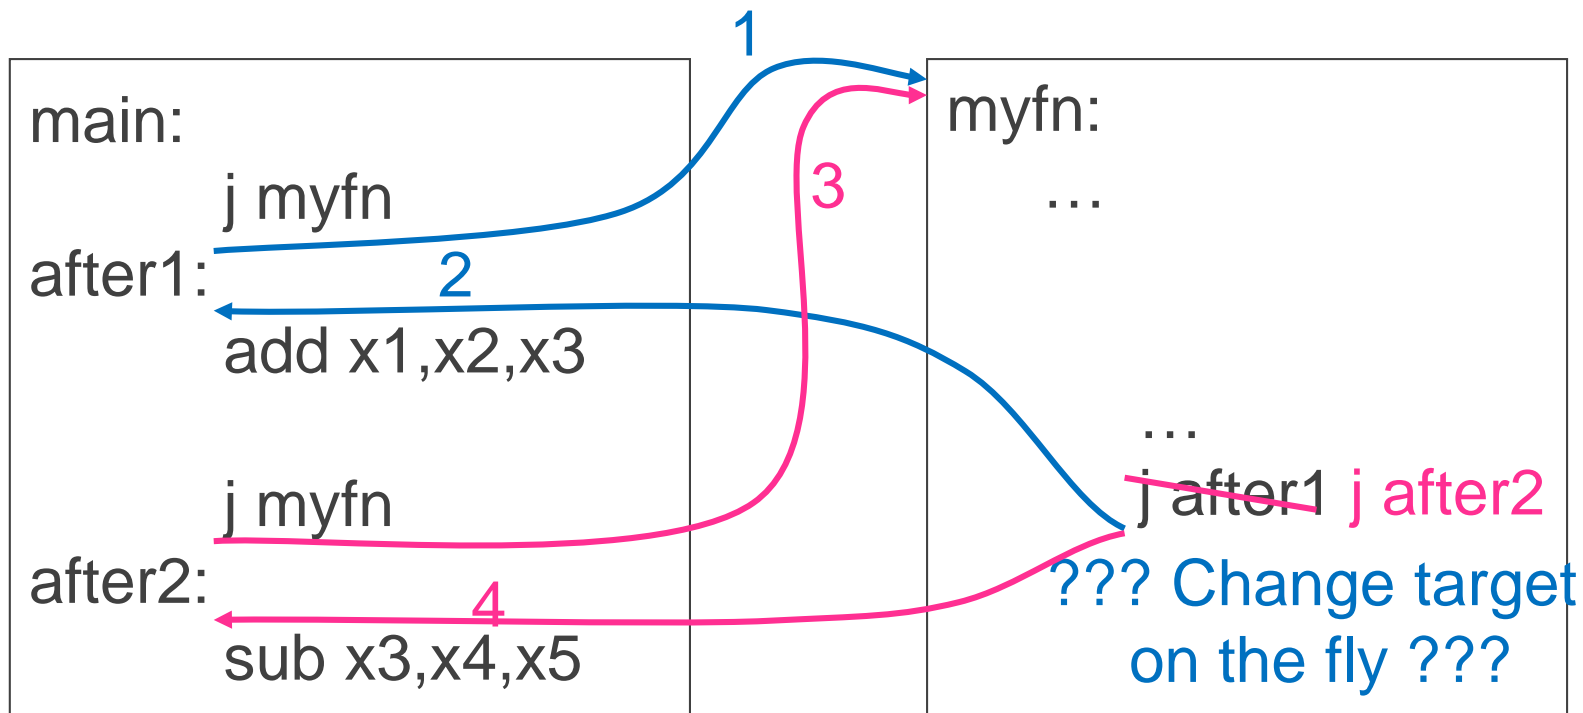
Jumps are not enough



Jumps to the callee

Jumps back

Jumps are not enough



Jumps to the callee

Jumps back

What about multiple sites?

Takeaway1: Need Jump And Link
JAL (Jump And Link) instruction moves a new value into the PC, and simultaneously saves the old value in register x1 (aka \$ra or return address)

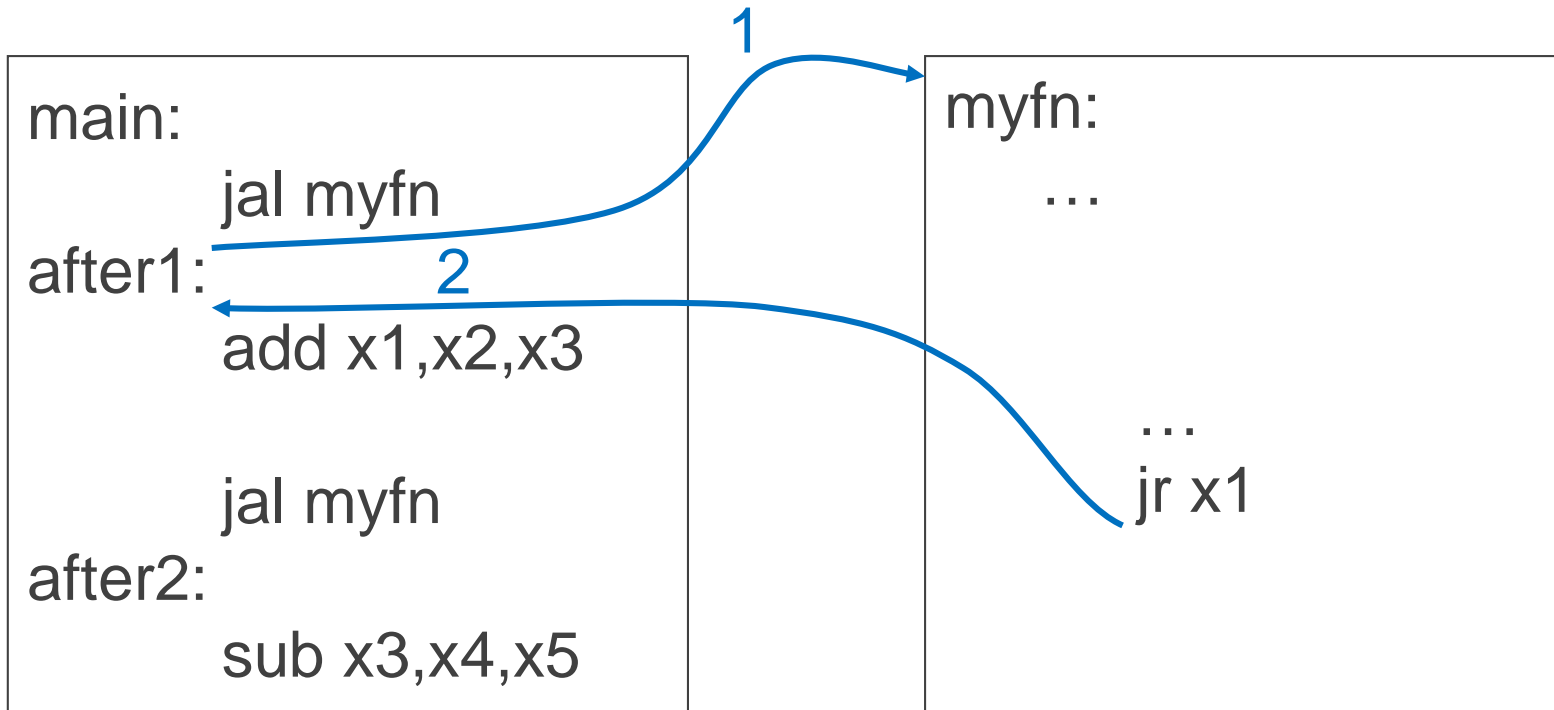
Thus, can get back from the subroutine to the instruction immediately following the jump by transferring control back to PC in register x1

Jump-and-Link / Jump Register

First call

x1

after1



JAL saves the PC in register \$31

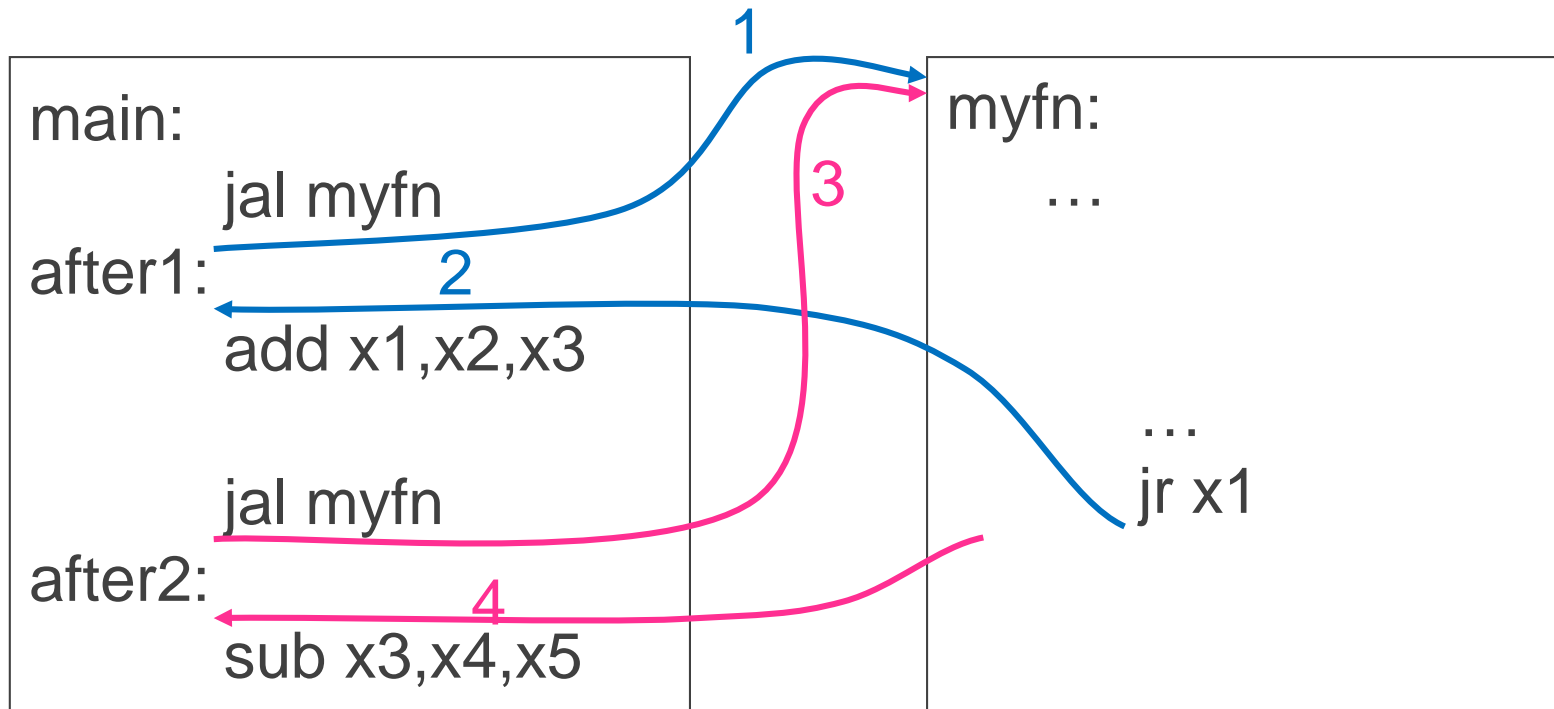
Subroutine returns by jumping to \$31

Jump-and-Link / Jump Register

Second call

x1

after2



JAL saves the PC in register x1

Subroutine returns by jumping to x1

What happens for recursive invocations?

JAL / JR for Recursion?

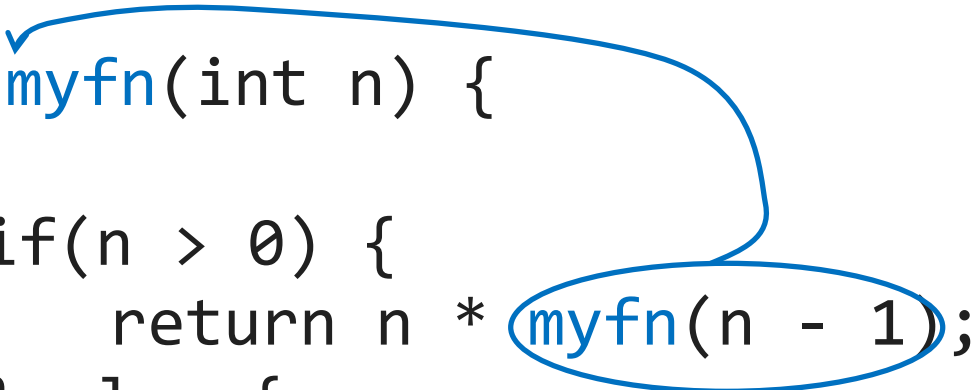
```
int main (int argc, char* argv[ ]) {  
    int n = 9;  
    int result = myfn(n);  
}
```

```
int myfn(int n) {  
    int f = 1;  
    int i = 1;  
    int j = n - 1;  
    while(j >= 0) {  
        f *= i;  
        i++;  
        j = n - i;  
    }  
    return f;  
}
```

JAL / JR for Recursion?

```
int main (int argc, char* argv[ ]) {  
    int n = 9;  
    int result = myfn(n);  
}
```

```
int myfn(int n) {  
    if(n > 0) {  
        return n * myfn(n - 1);  
    } else {  
        return 1;  
    }  
}
```

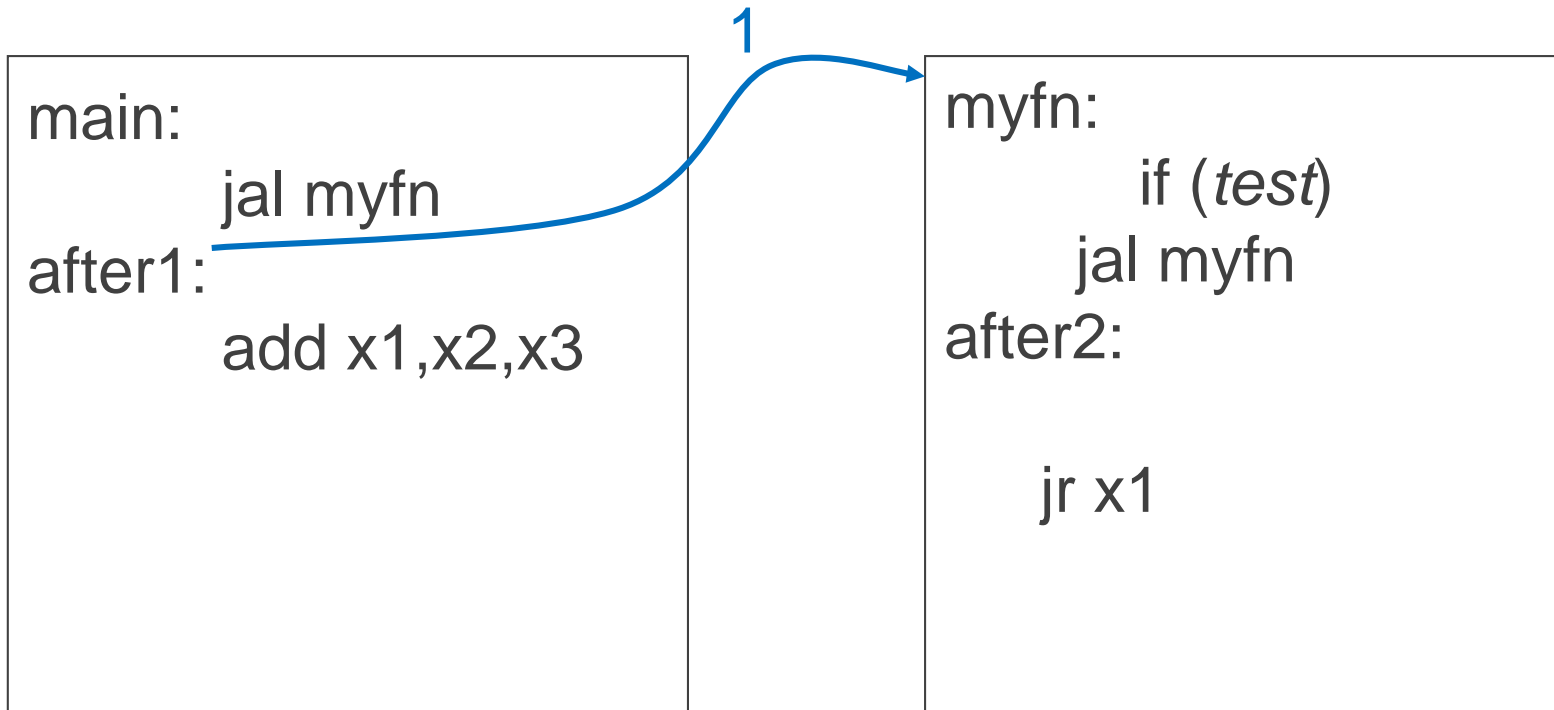


JAL / JR for Recursion?

First call

x1

after1

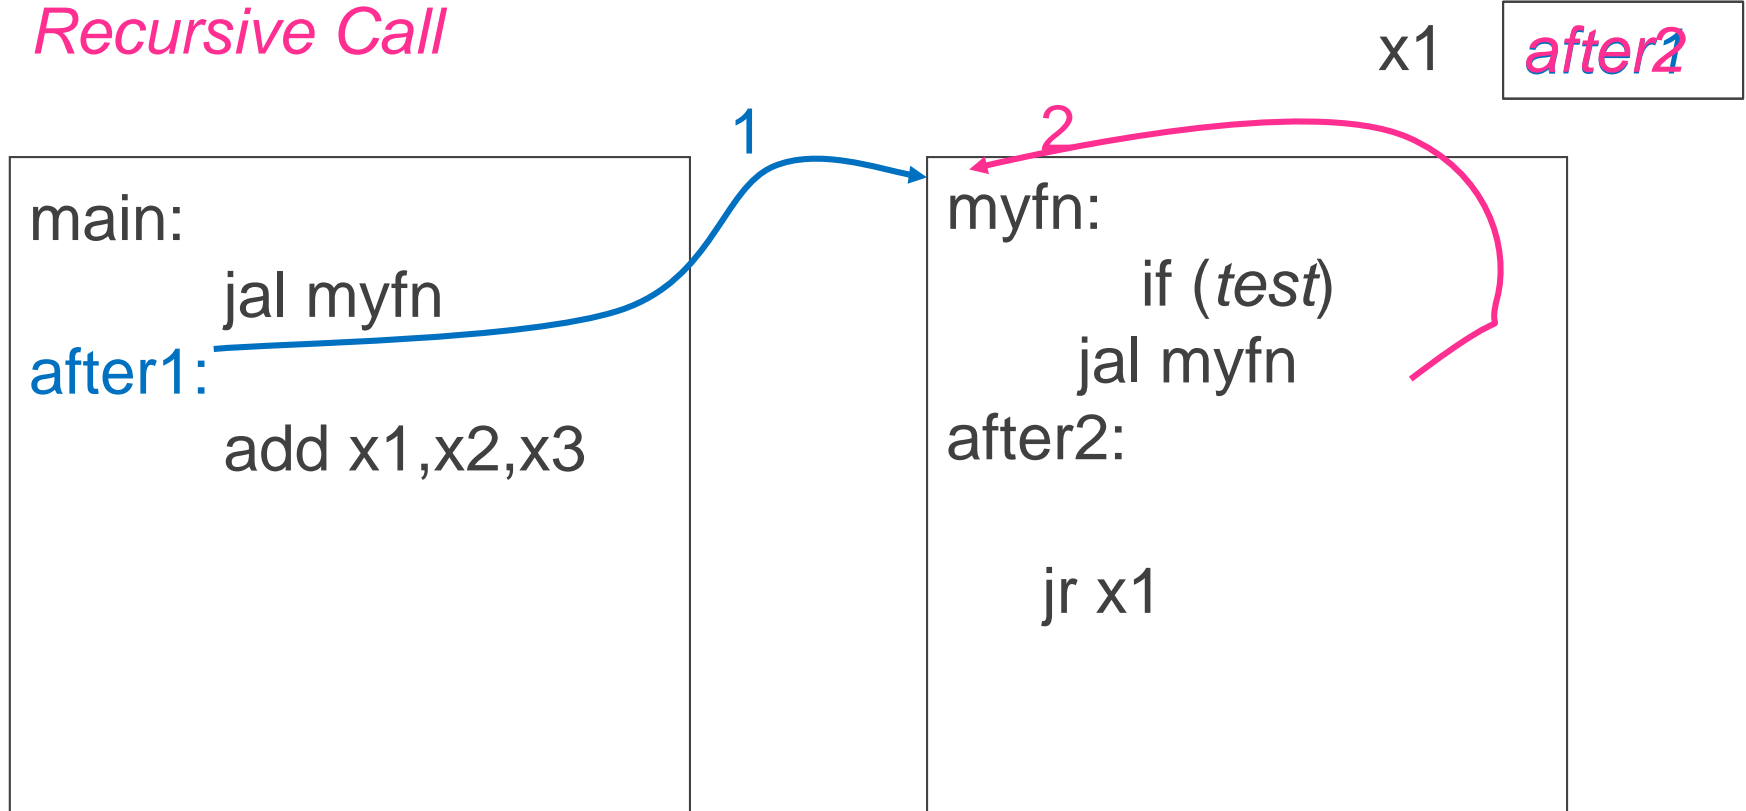


Problems with recursion:

- overwrites contents of x1

JAL / JR for Recursion?

Recursive Call

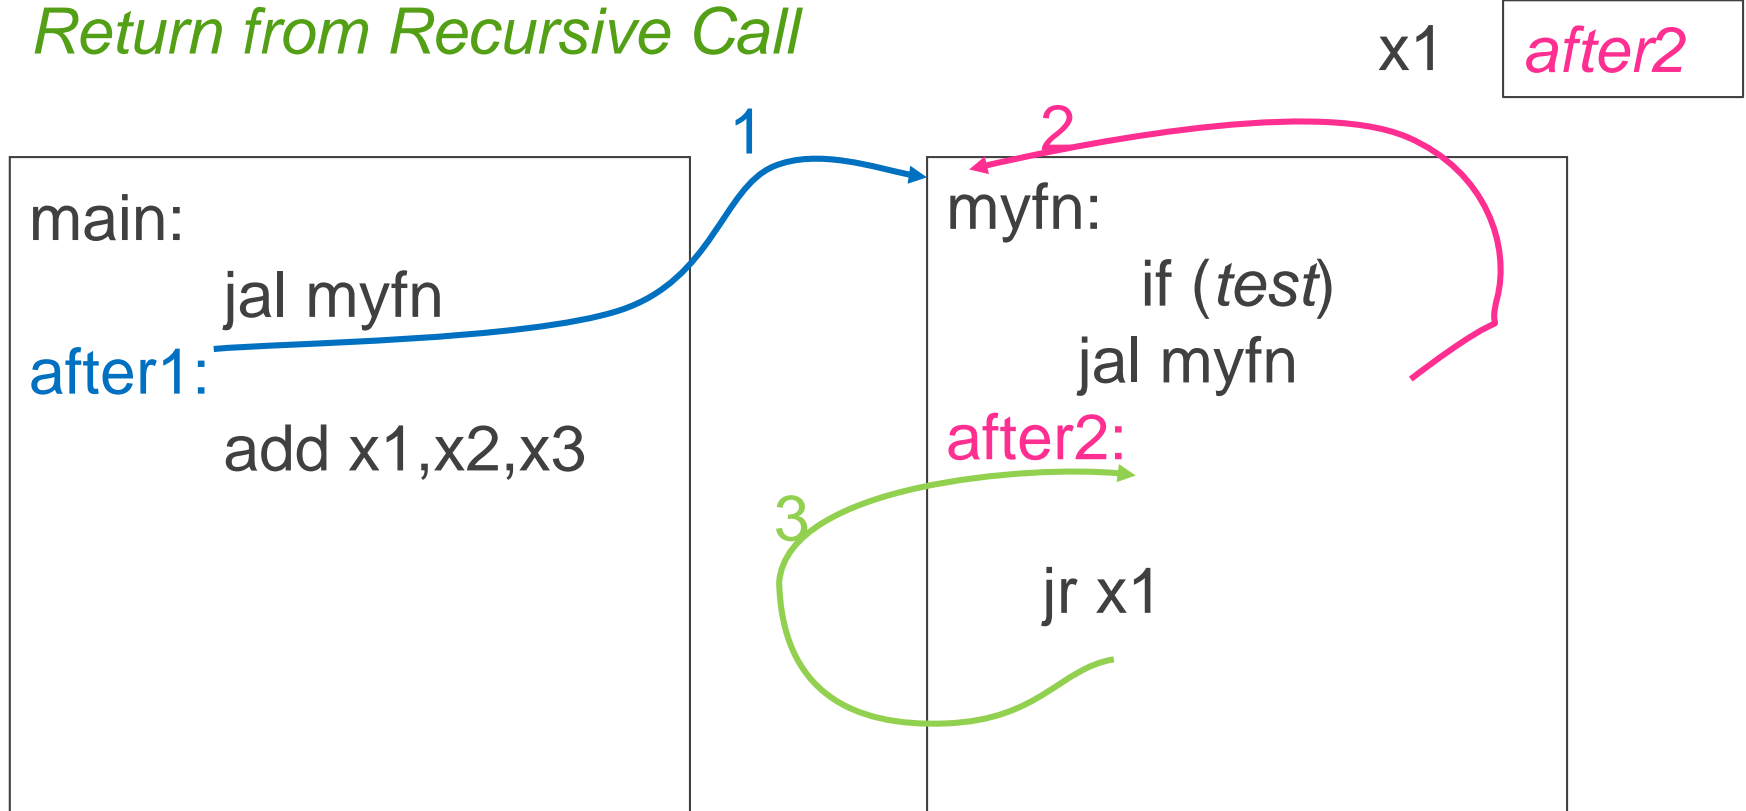


Problems with recursion:

- overwrites contents of x1

JAL / JR for Recursion?

Return from Recursive Call

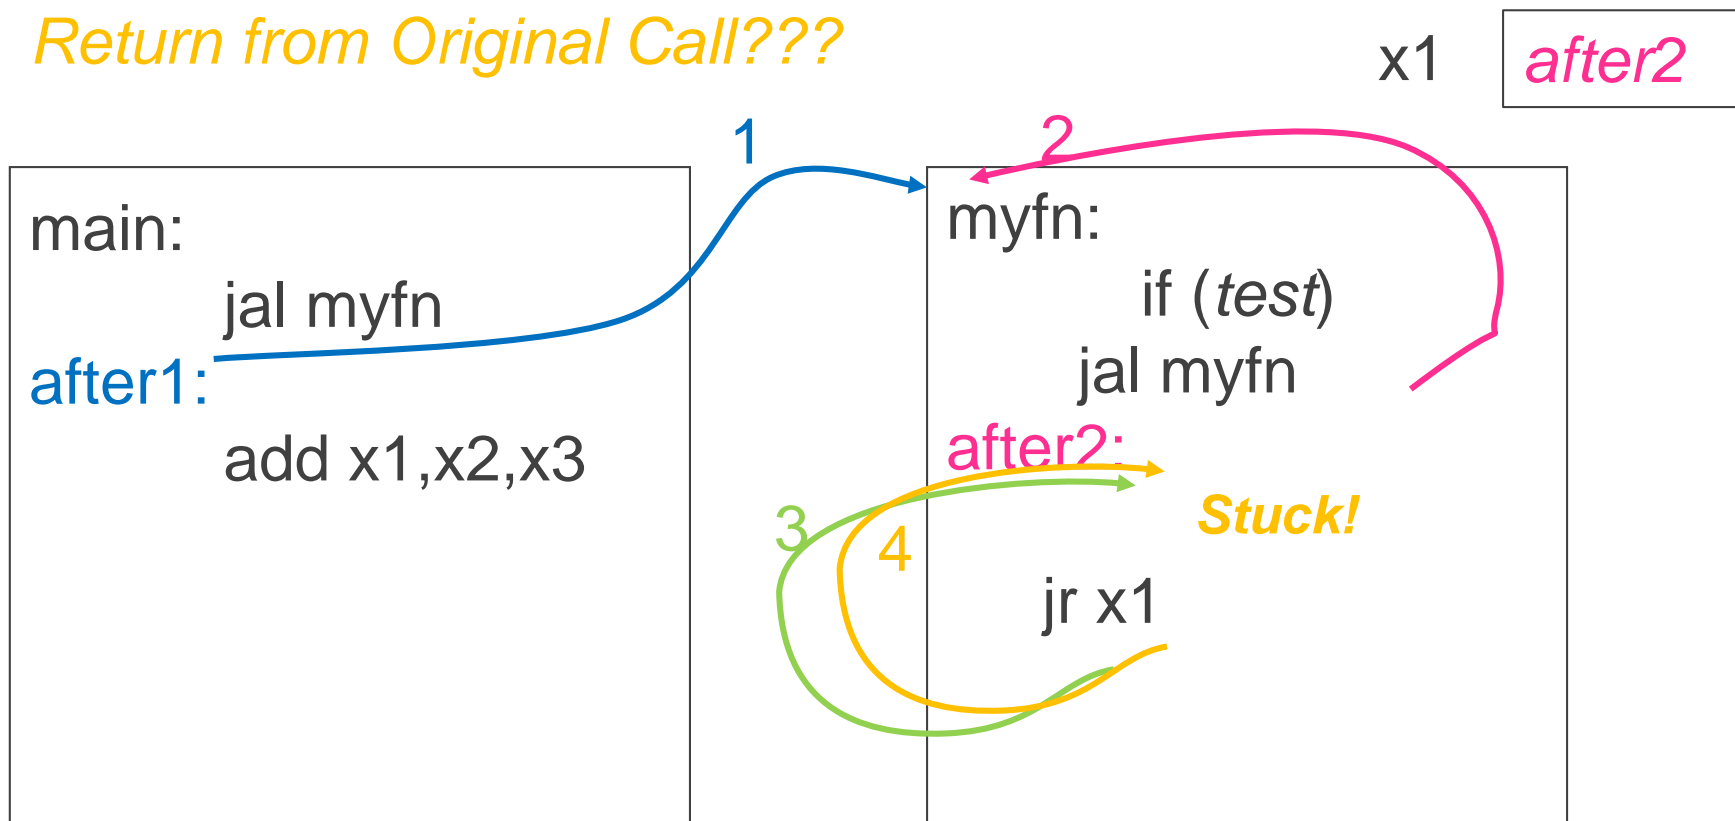


Problems with recursion:

- overwrites contents of x1

JAL / JR for Recursion?

Return from Original Call???

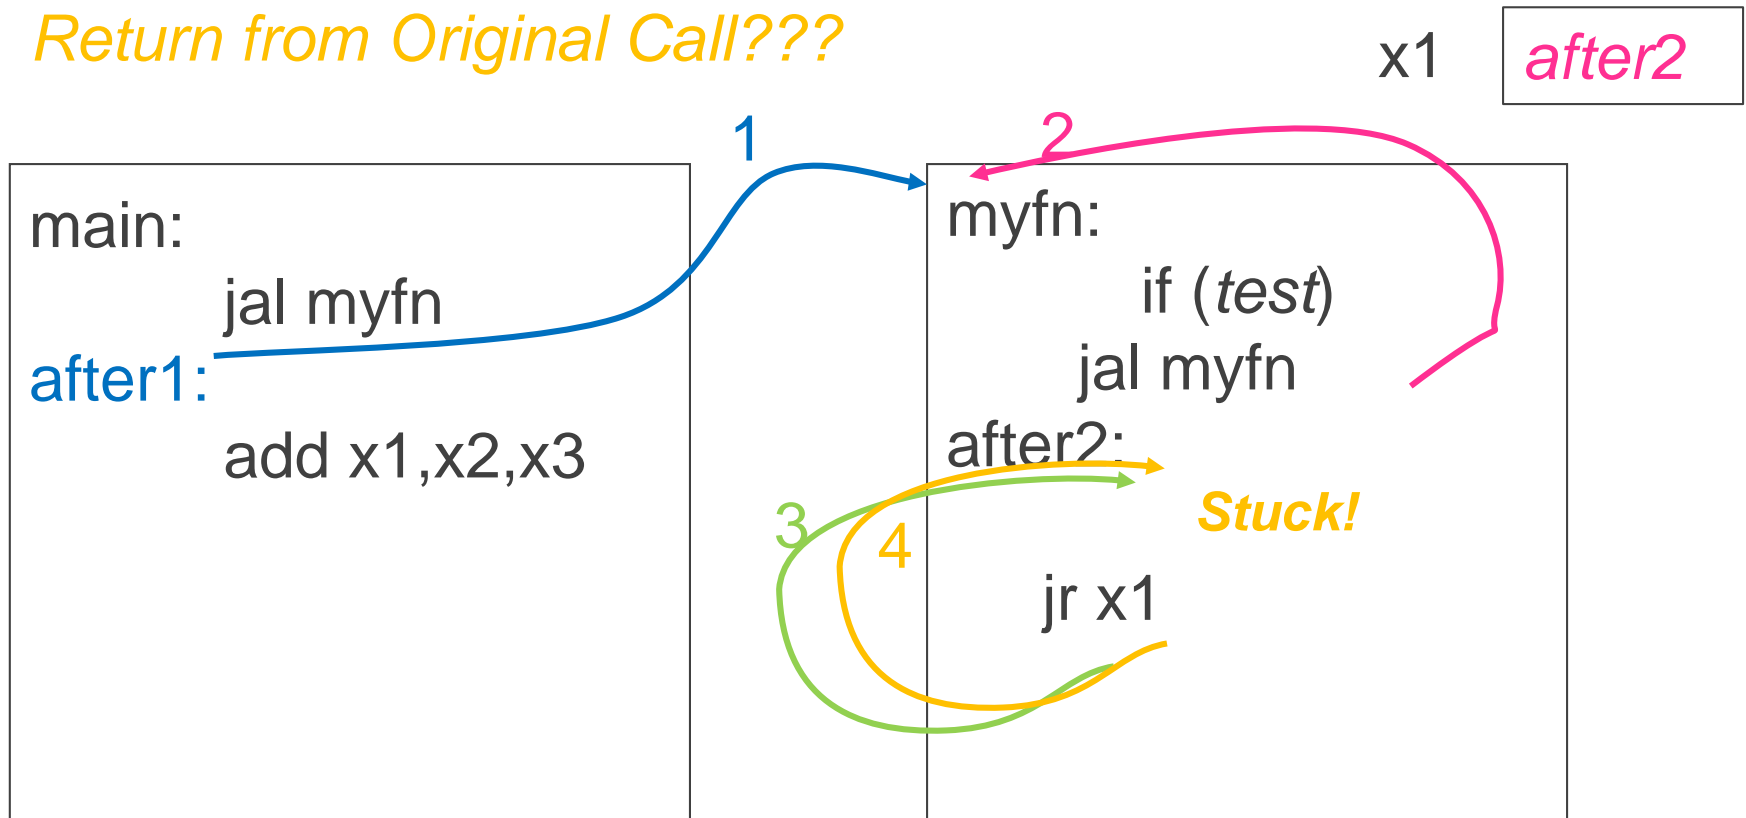


Problems with recursion:

- overwrites contents of x1

JAL / JR for Recursion?

Return from Original Call???



Problems with recursion:

overwrites contents of x1

Need a way to save and restore register contents

Need a “Call Stack”

Call stack

- contains activation records (aka stack frames)

Each activation record contains


- the return address for that invocation
- the local variables for that procedure

A **stack pointer (sp)** keeps track of the top of the stack

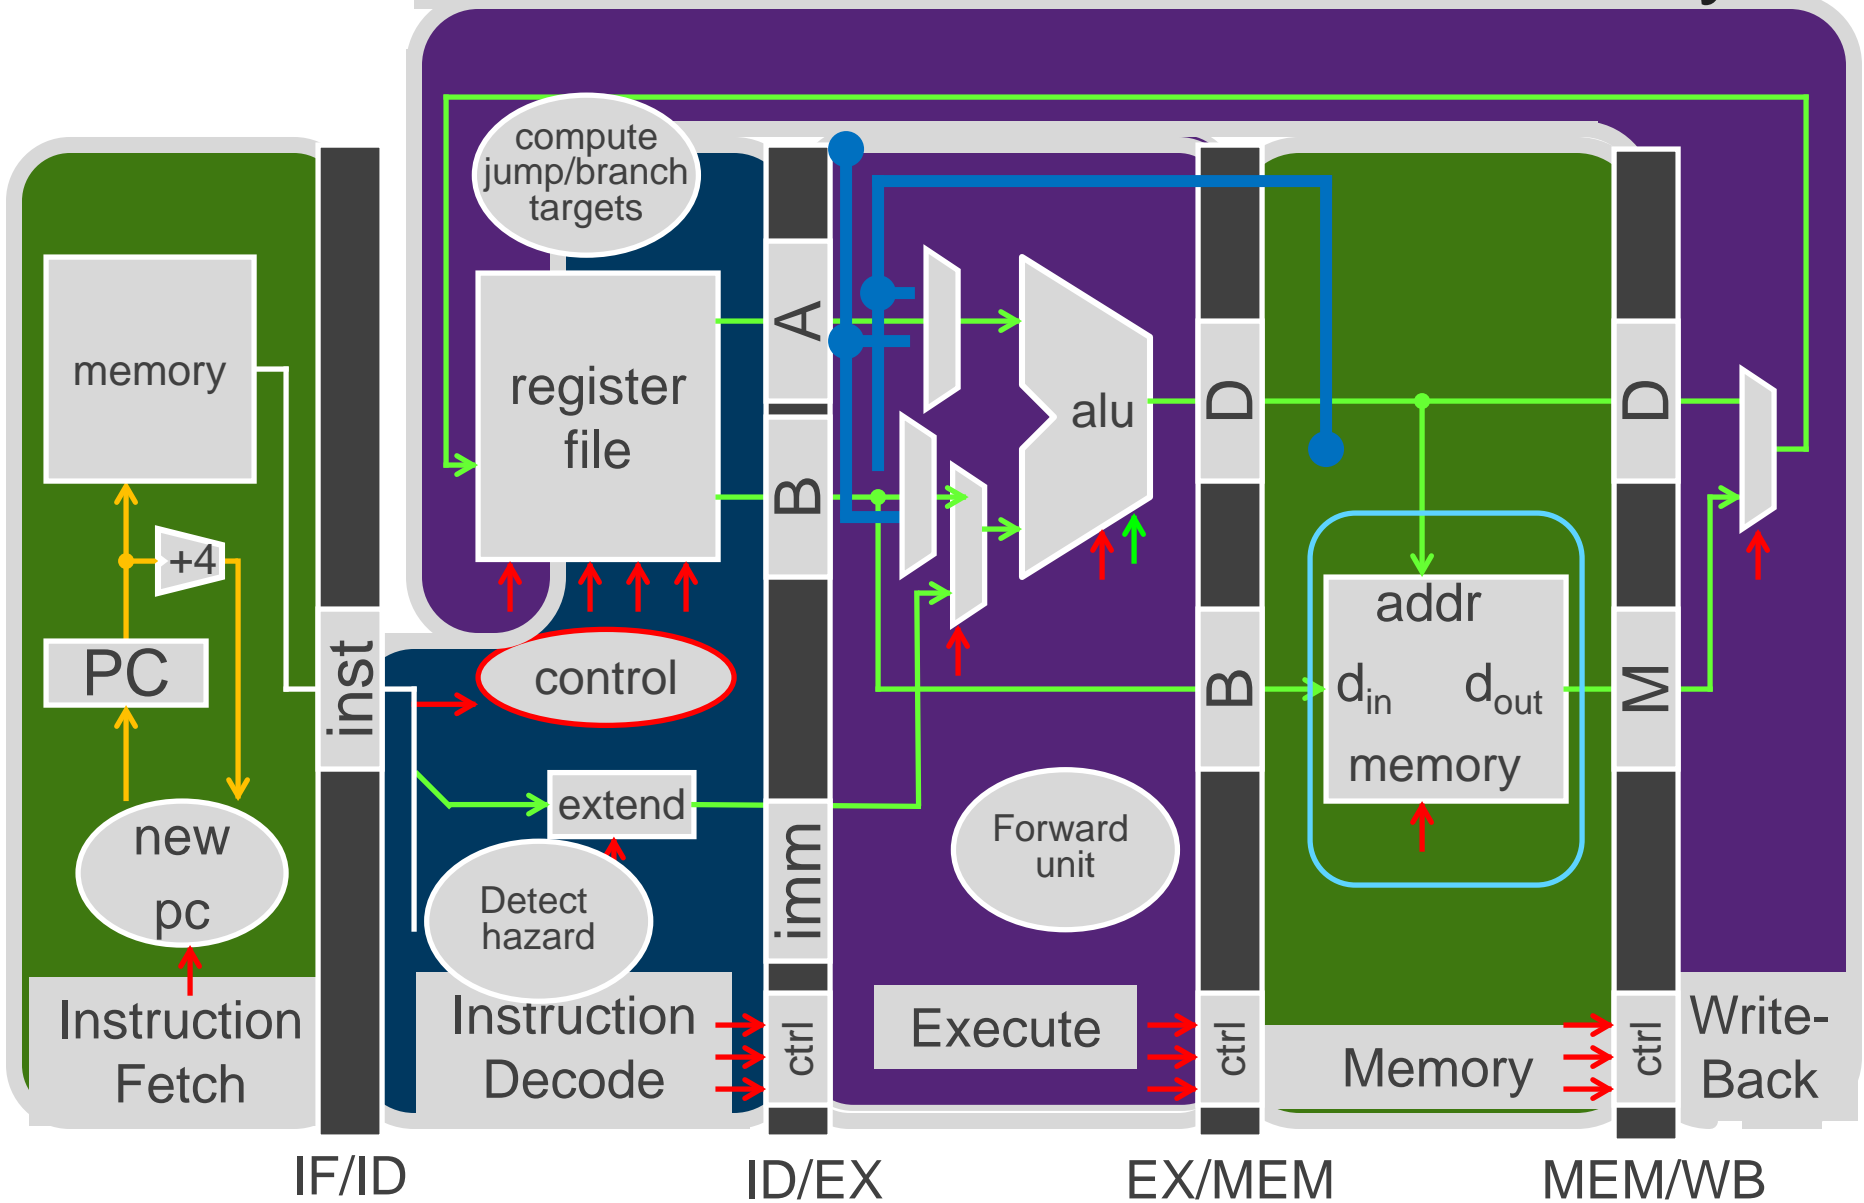
- dedicated register (**x2**) on the RISC-V

Manipulated by **push/pop** operations

- **push**: move sp down, store
- **pop**: load, move sp up



Cheat Sheet and Mental Model for Today



Need a “Call Stack”

Call stack

- contains activation records (aka stack frames)

Each activation record contains

- the return address for that invocation
- the local variables for that procedure

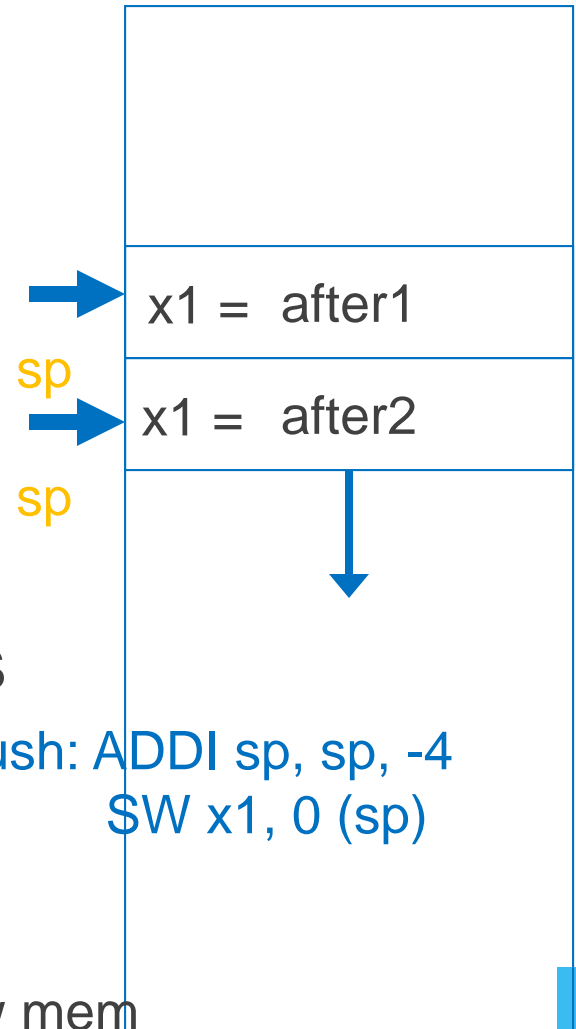
A **stack pointer (sp)** keeps track of the top of the stack

- dedicated register (**x2**) on the RISC-V

Manipulated by **push/pop** operations

- **push**: move sp down, store
- **pop**: load, move sp up

high mem



low mem

Need a “Call Stack”

Call stack

- contains activation records (aka stack frames)

Each activation record contains

- the return address for that invocation
- the local variables for that procedure

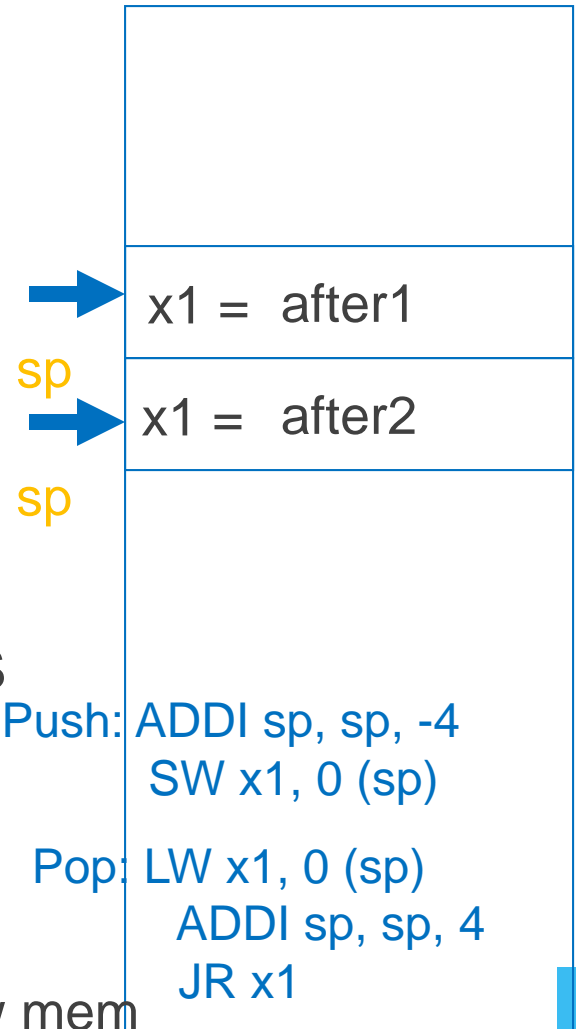
A **stack pointer (sp)** keeps track of the top of the stack

- dedicated register (**x2**) on the RISC-V

Manipulated by **push/pop** operations

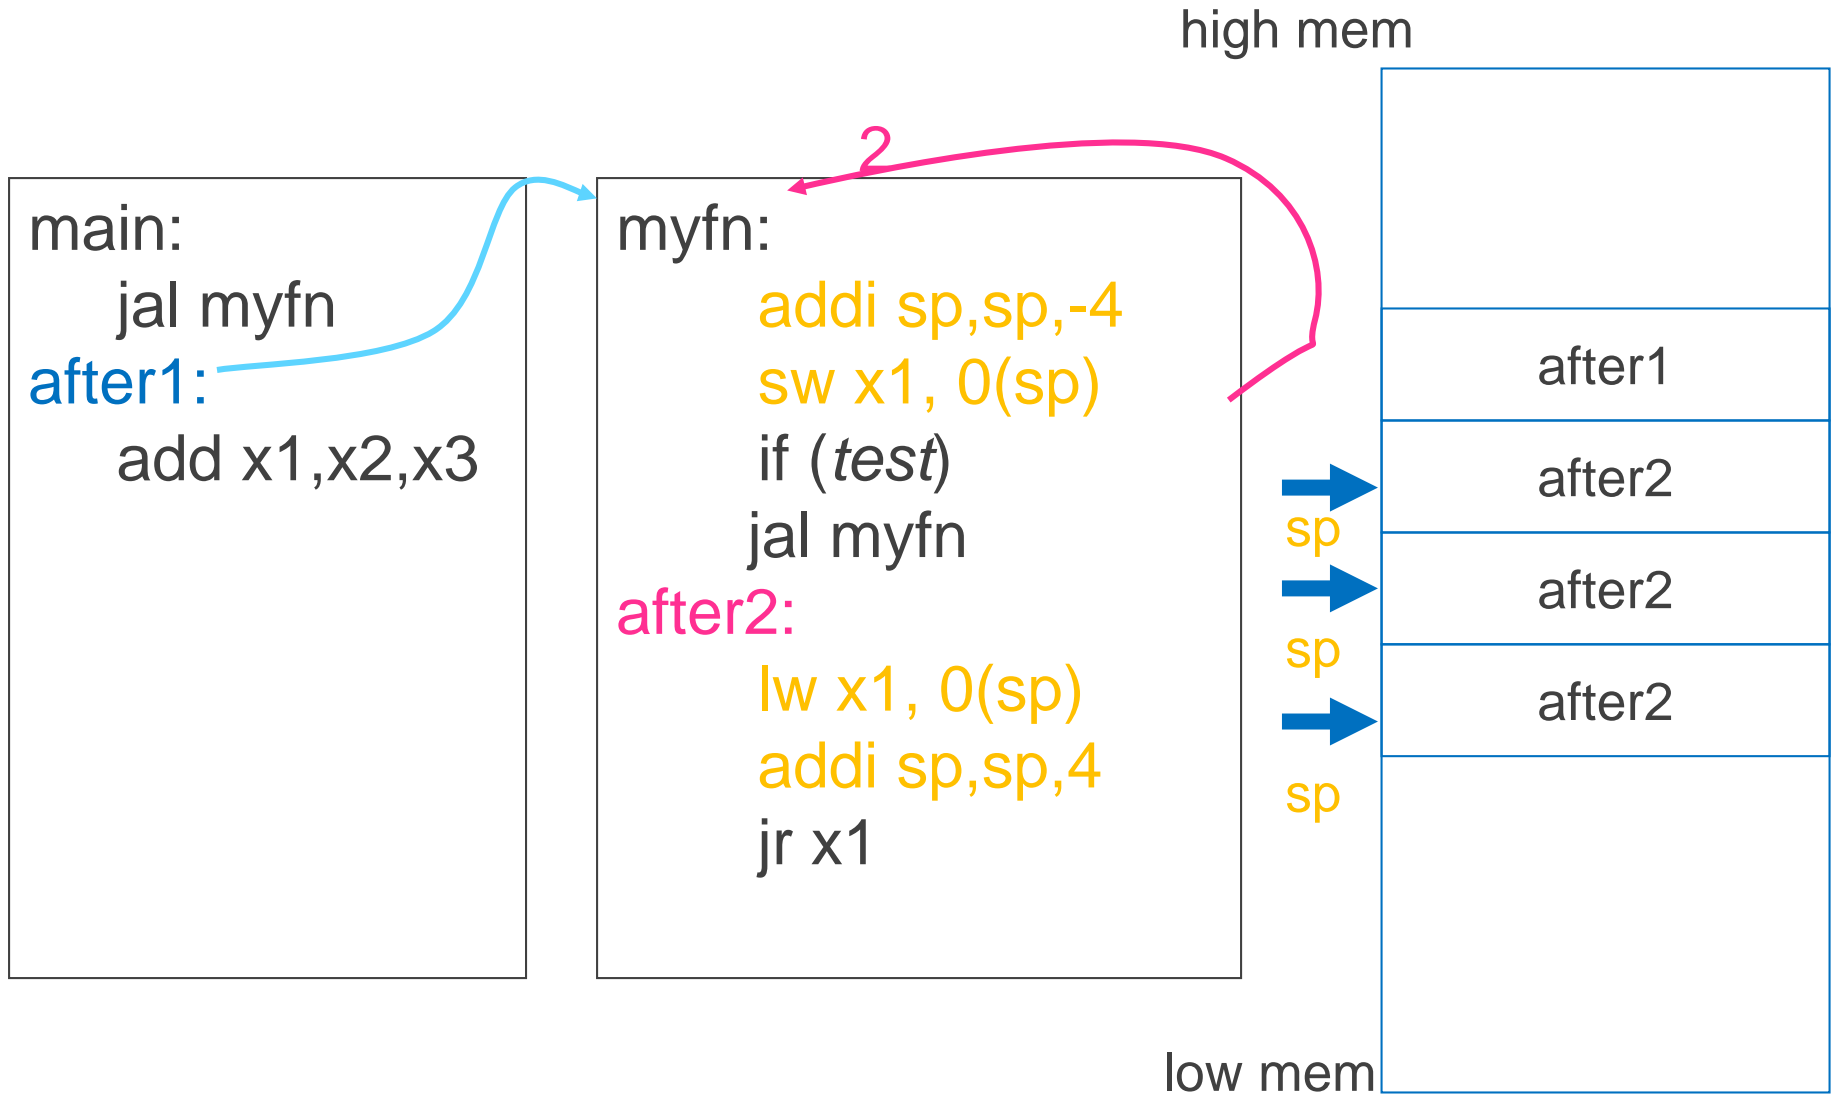
- **push**: move sp down, store
- **pop**: load, move sp up

high mem




low mem

Need a "Call Stack"



Stack used to save and restore contents of x1

Need a "Call Stack"



Stack used to save and restore contents of x1

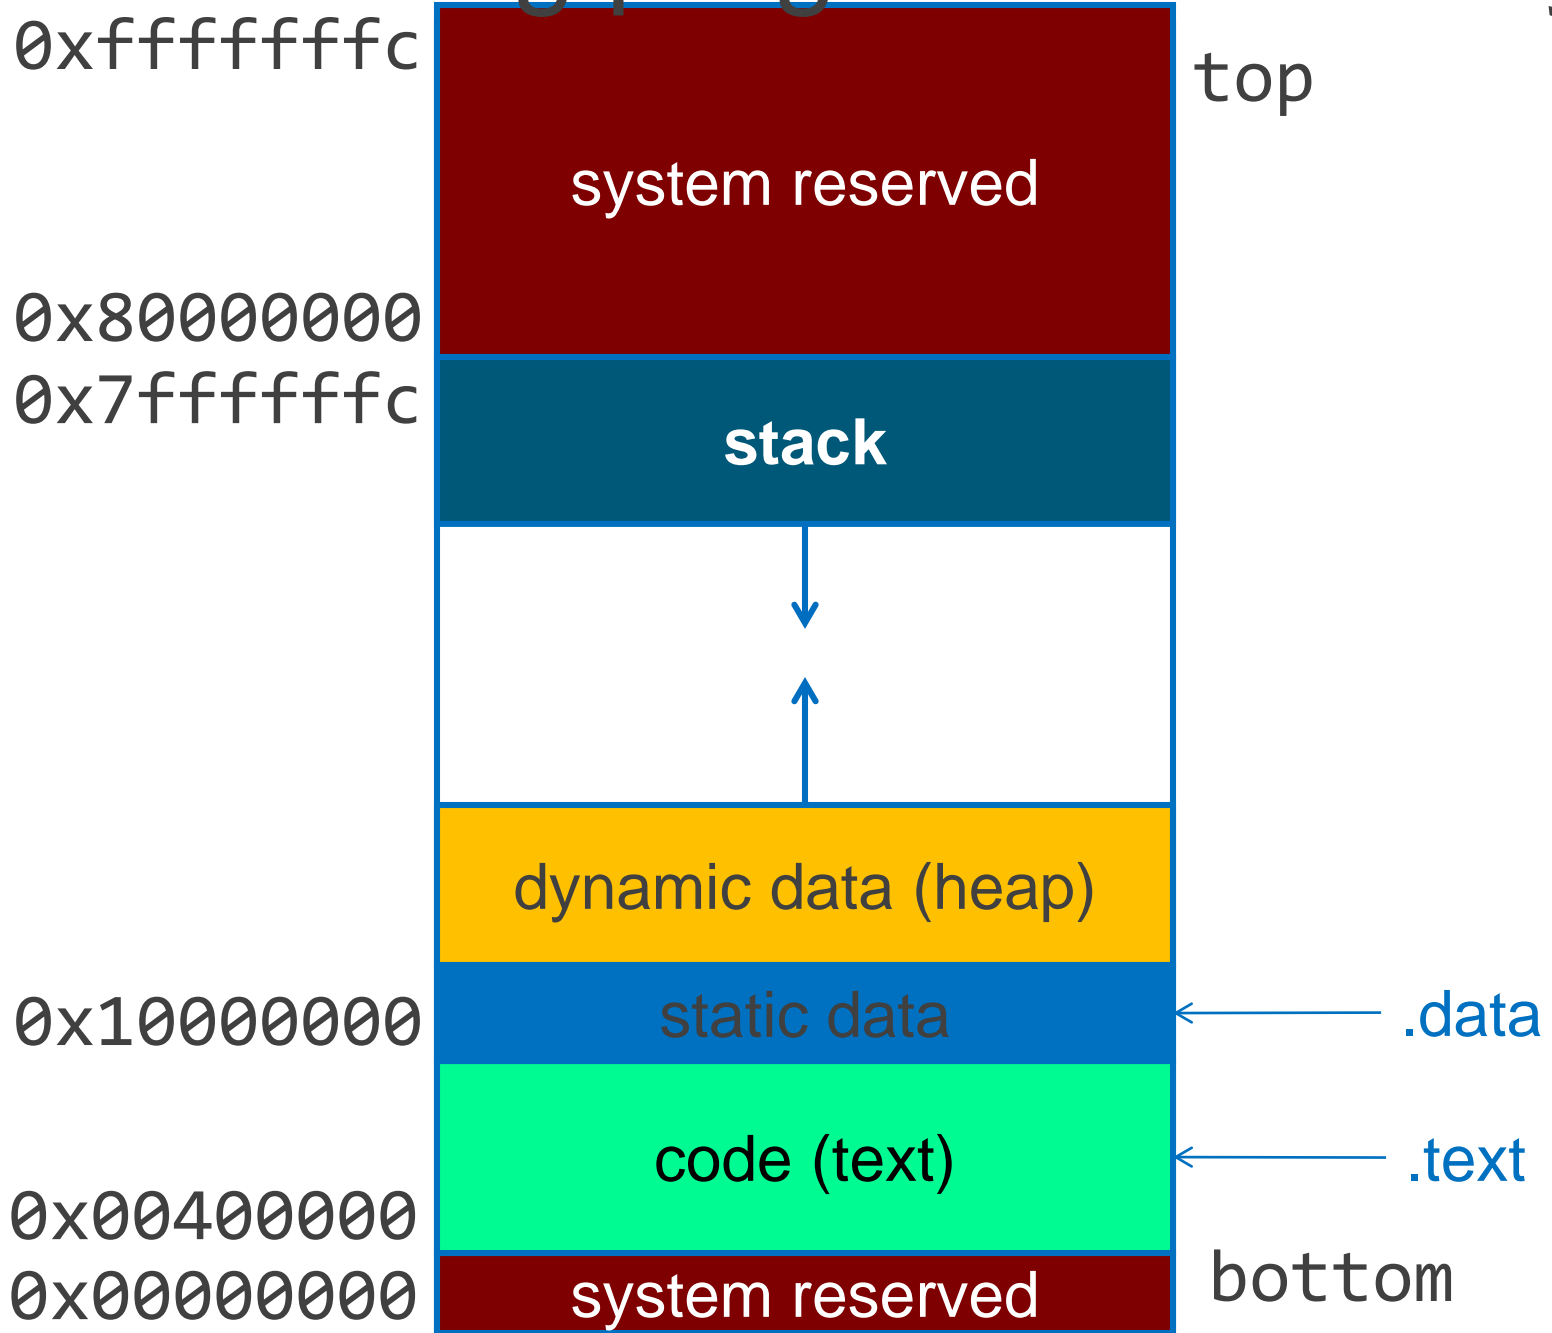
Stack Growth

(Call) Stacks start at a high address in memory

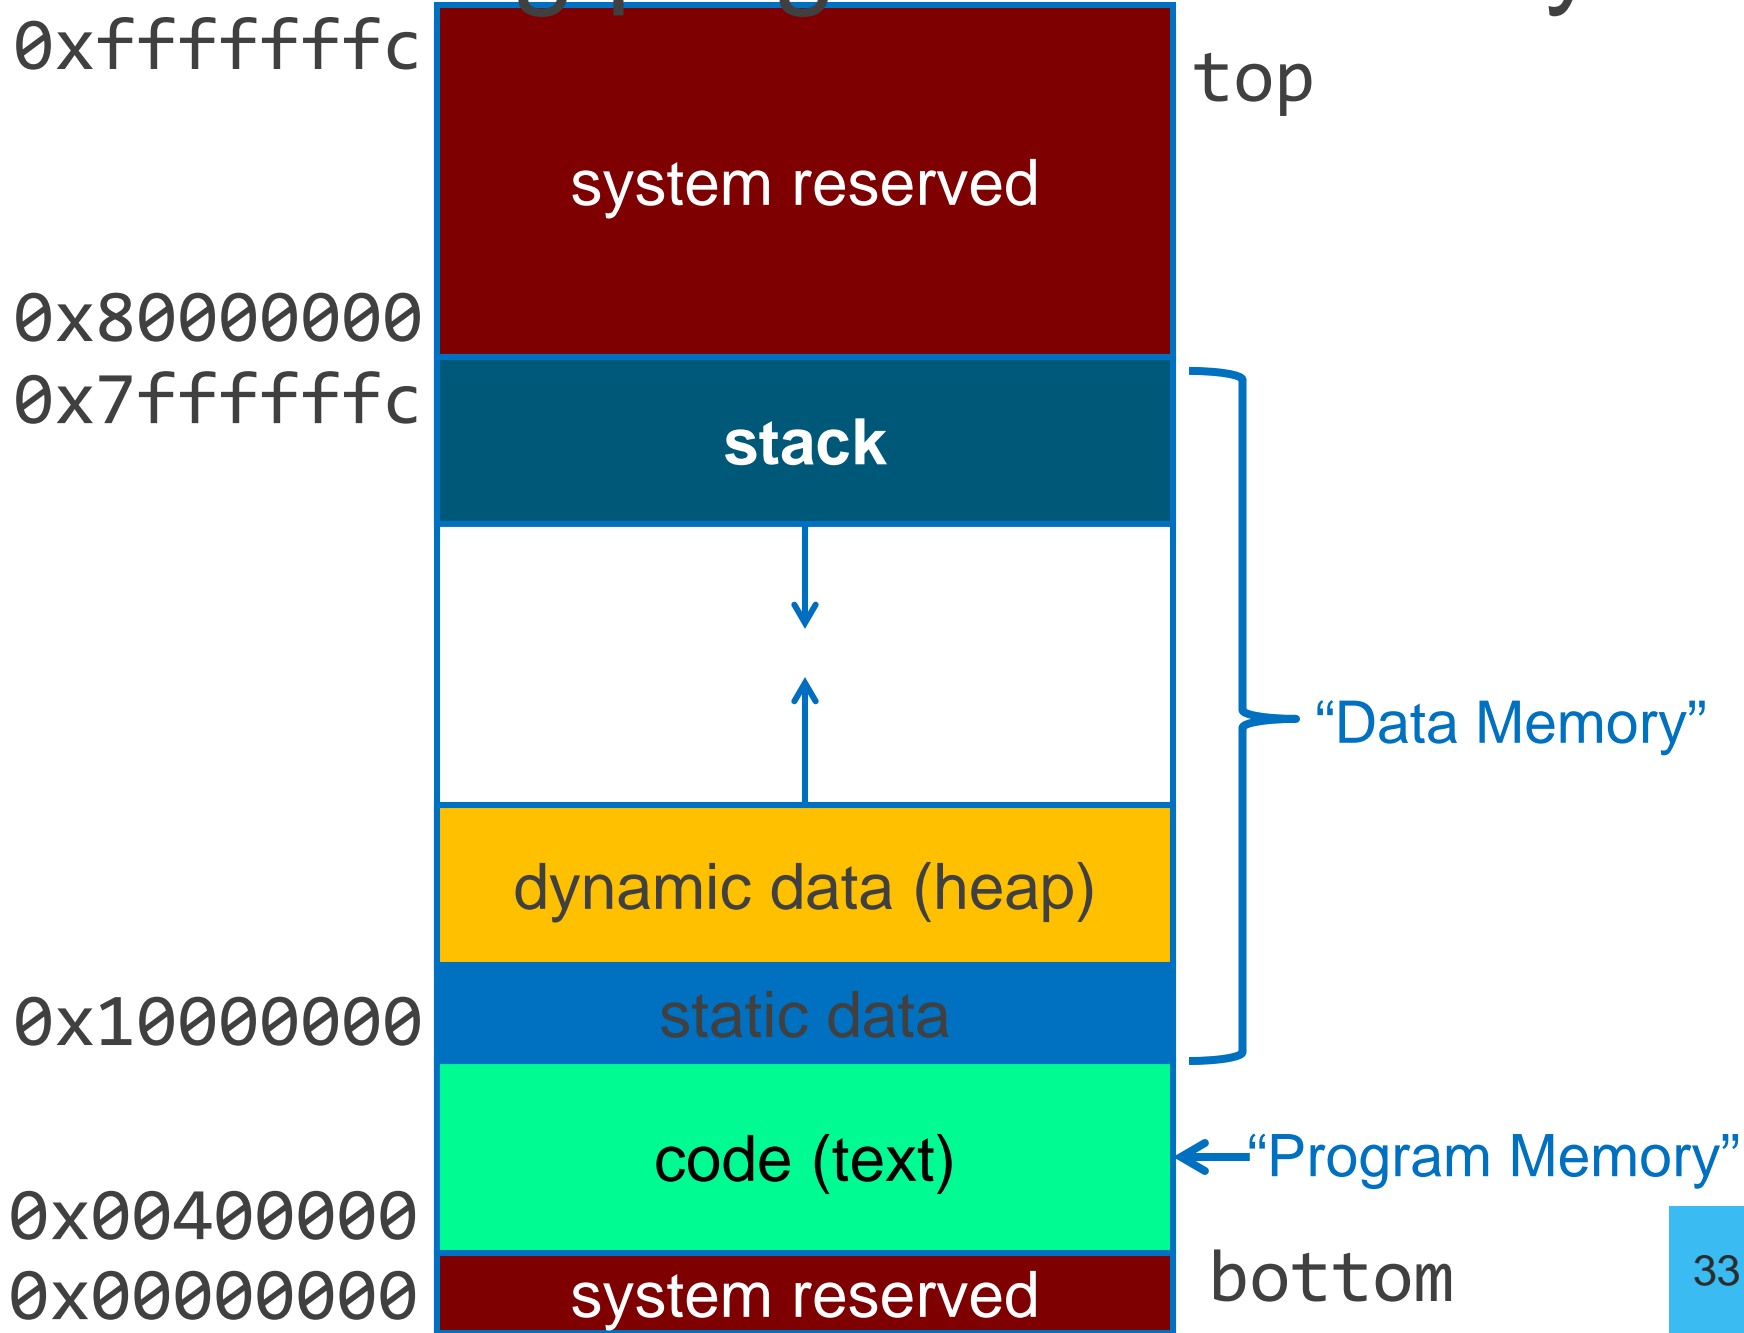
Stacks grow down as frames are pushed on

- Note: data region starts at a low address and grows up
- The growth potential of stacks and data region are not artificially limited

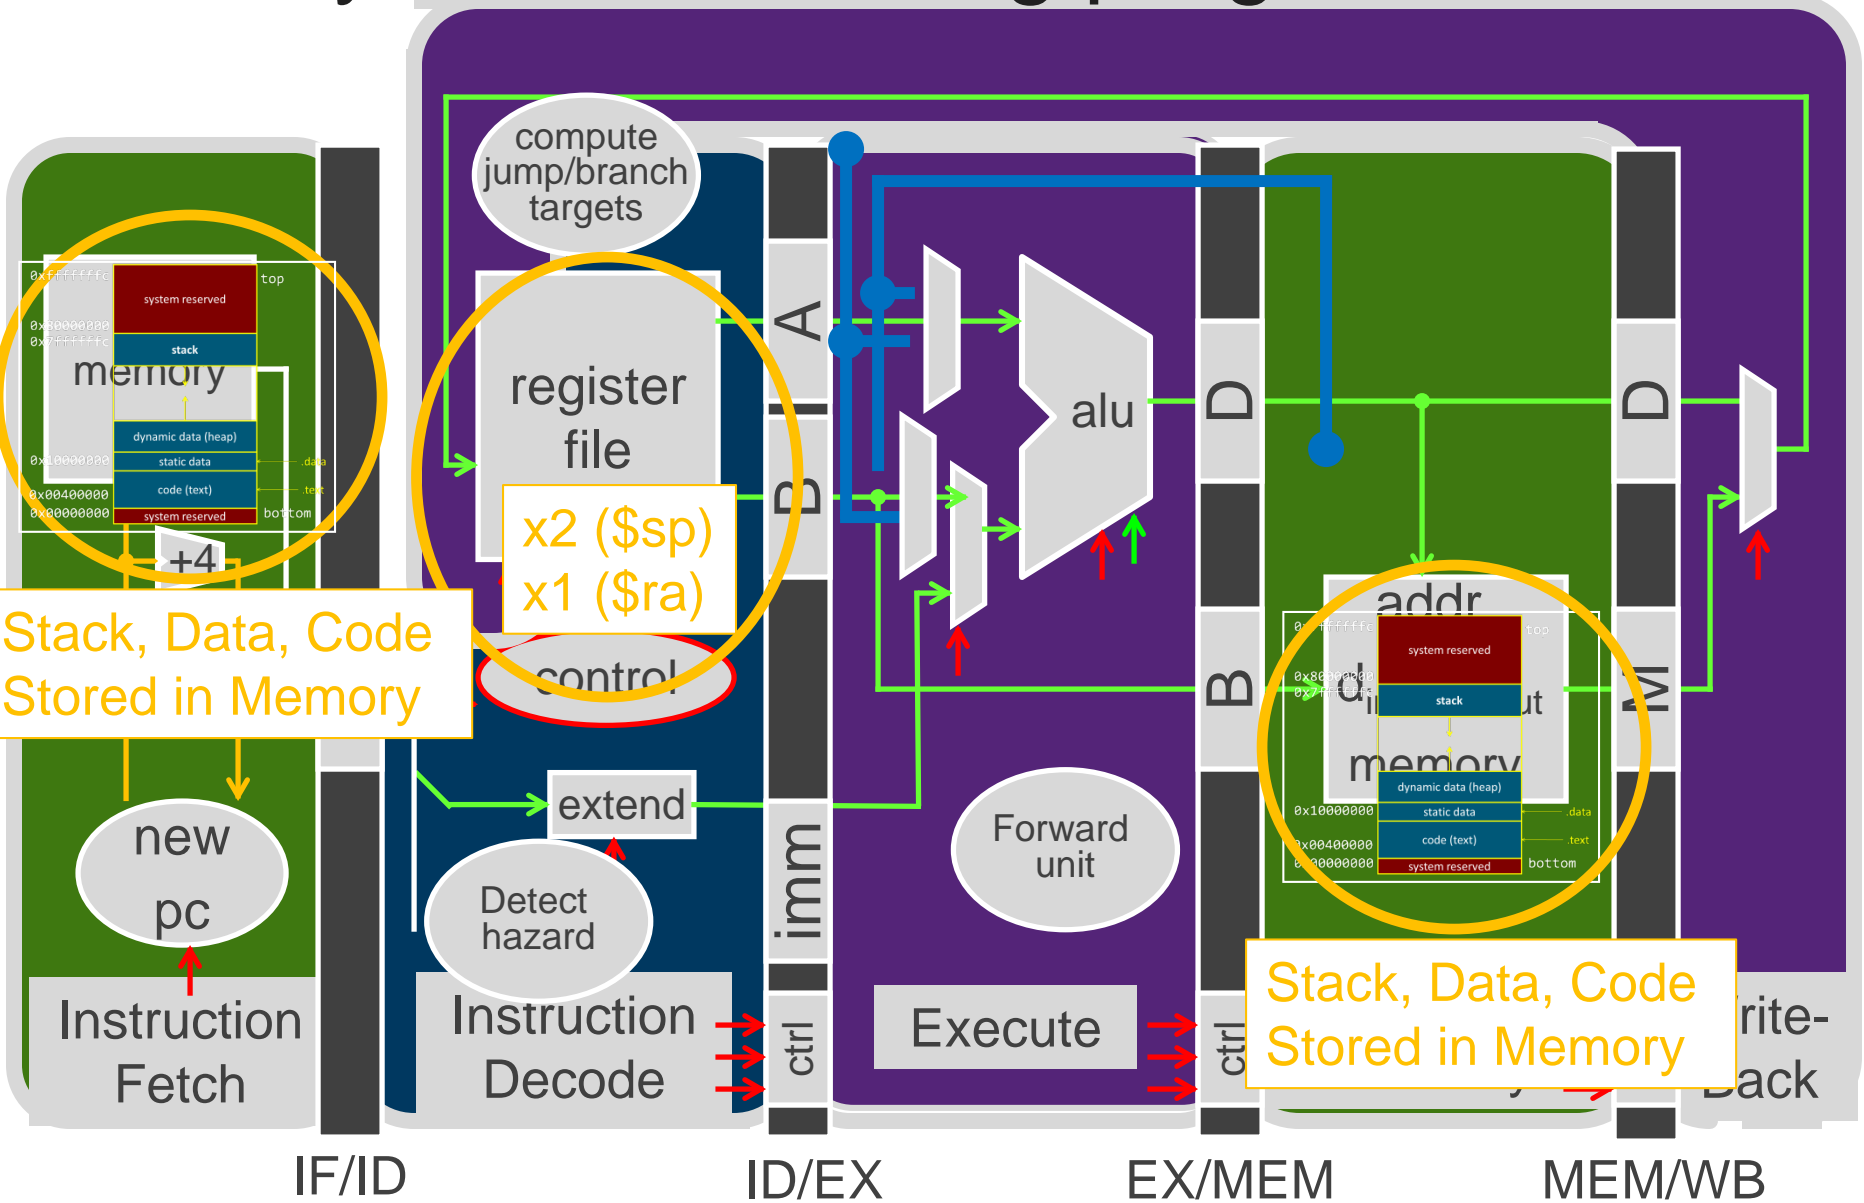
An executing program in memory



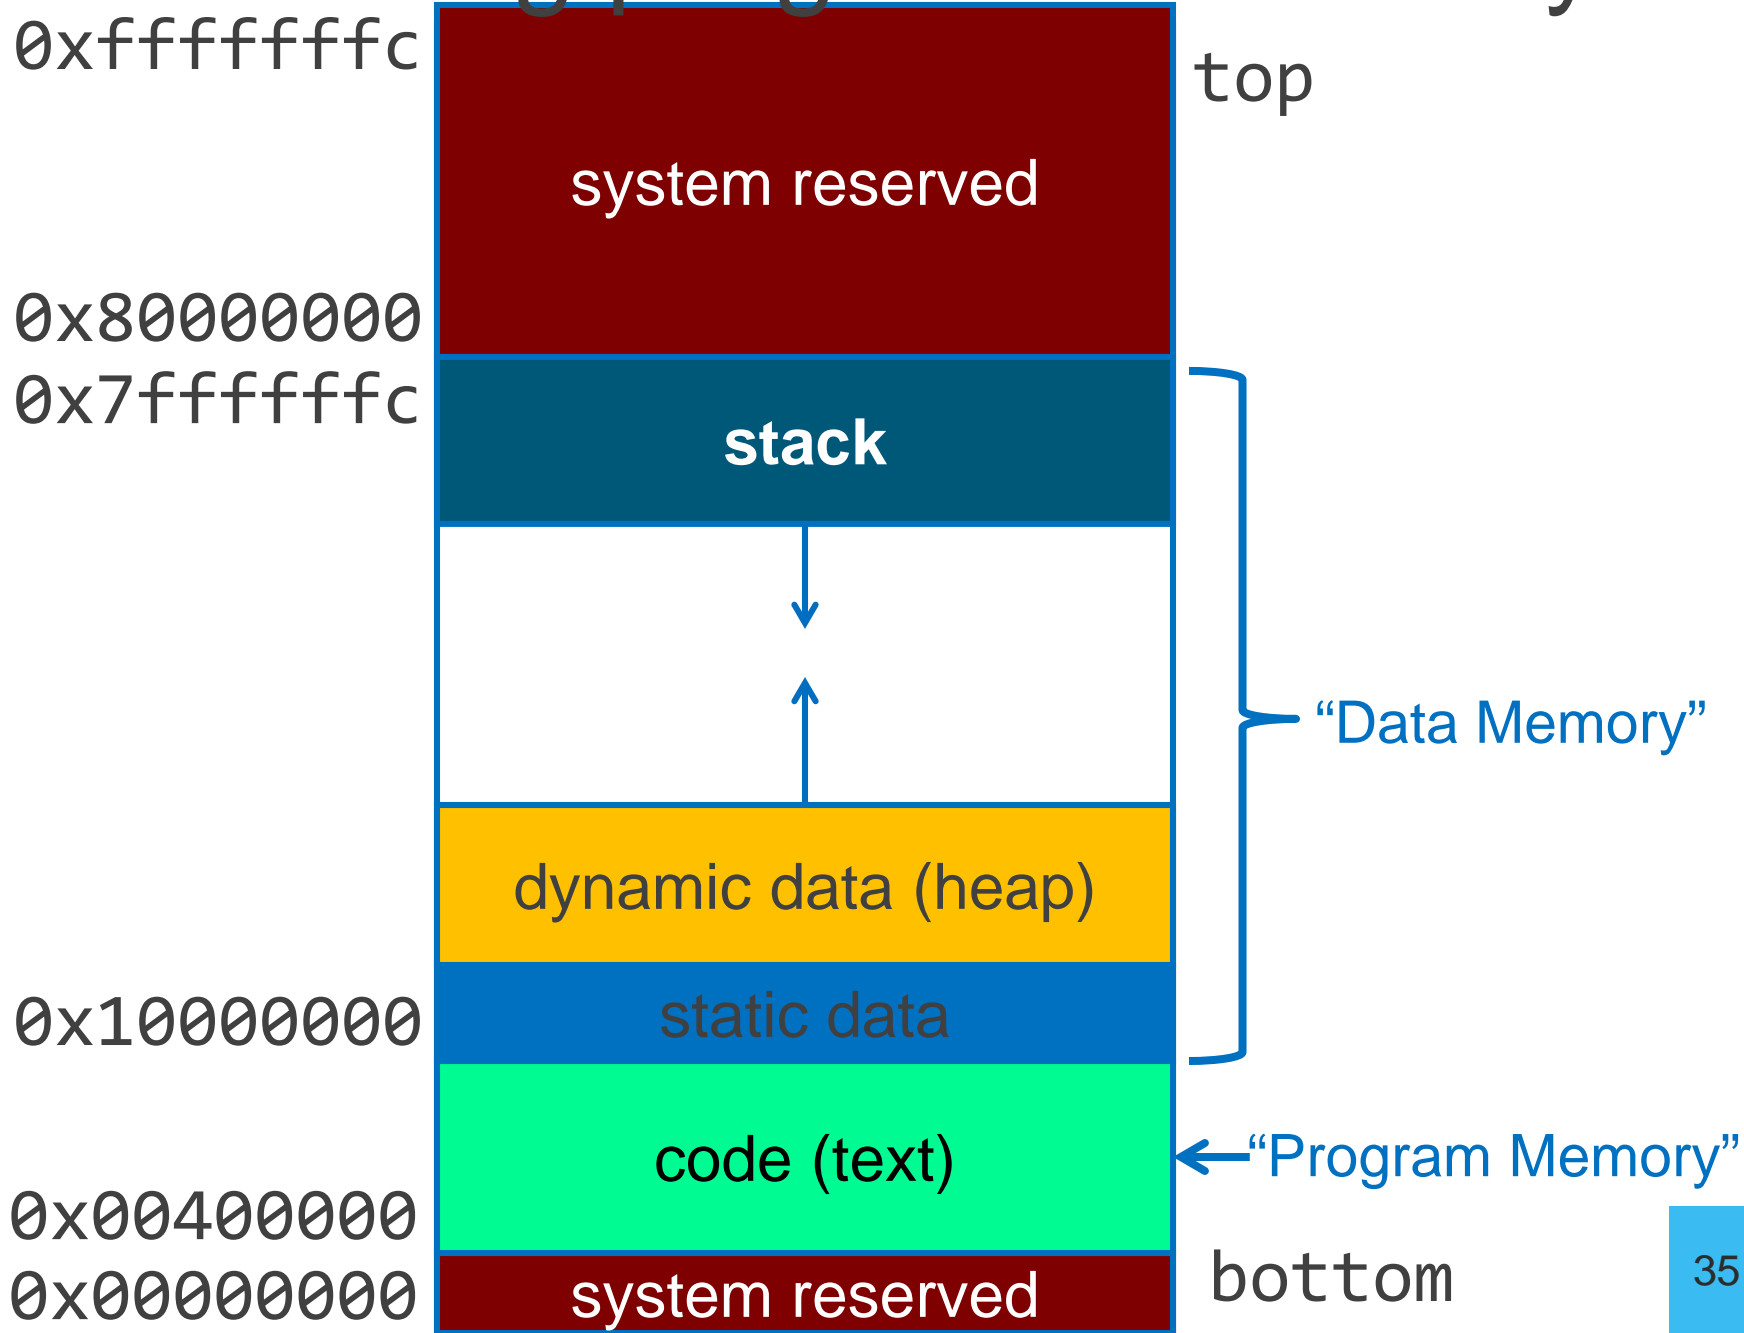
An executing program in memory



Anatomy of an executing program



An executing program in memory



The Stack

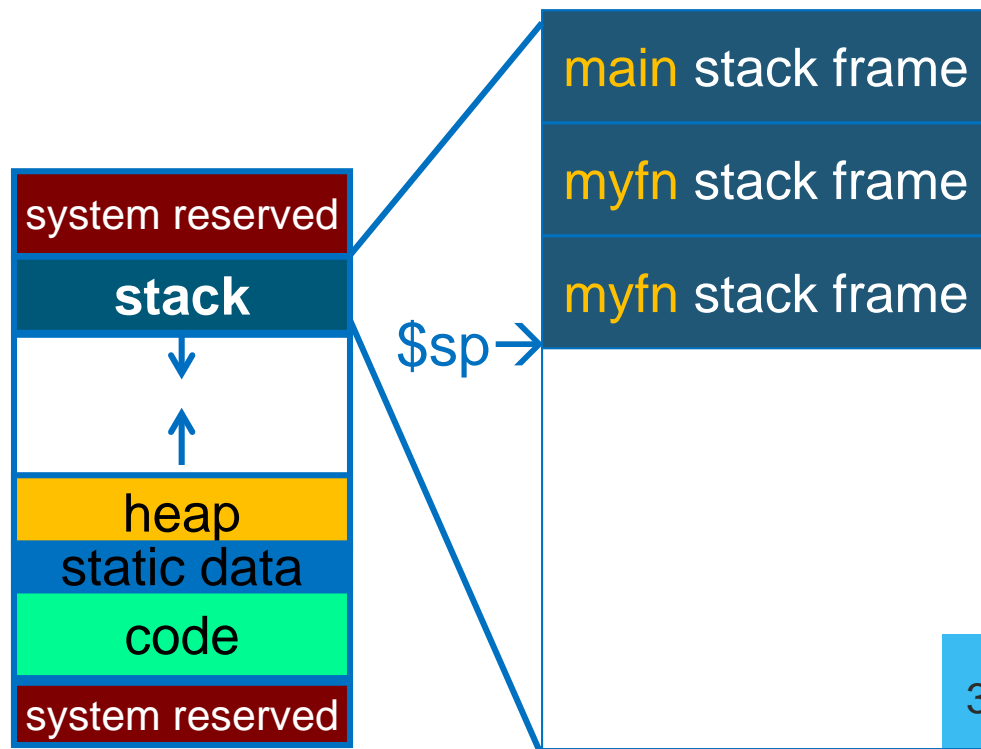
Stack contains stack frames (aka “activation records”)

- 1 stack frame per dynamic function
- Exists only for the duration of function
- Grows down, “top” of stack is `sp`, `x2`
- Example: `lw x5, 0(sp)` puts word at top of stack into `x5`

Each stack frame contains:

- Local variables, return address (later), register backups (later)

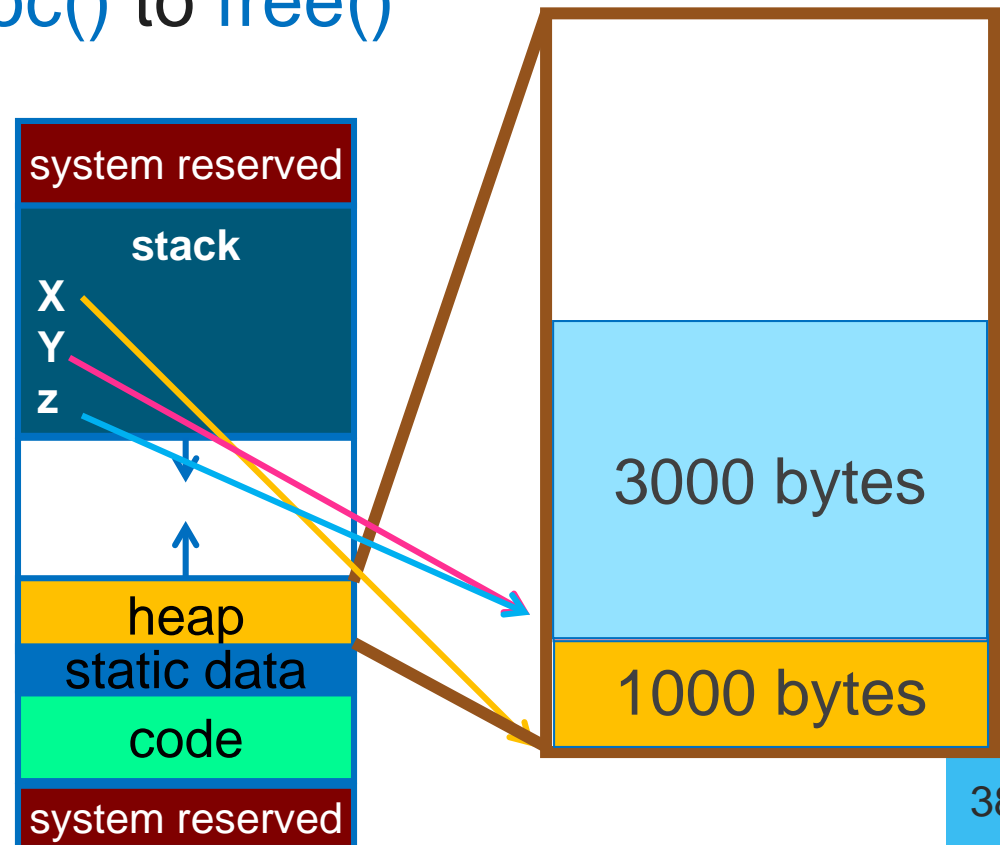
```
int main(...) {  
    ...  
    myfn(x);  
}  
int myfn(int n) {  
    ...  
    myfn();  
}
```



The Heap

- Heap holds dynamically allocated memory
- Program must maintain pointers to anything allocated
 - Example: if x5 holds x
 - lw x6, 0(x5) gets first word x points to
- Data exists from `malloc()` to `free()`

```
void some_function() {  
    int *x = malloc(1000);  
    int *y = malloc(2000);  
    free(y);  
    int *z = malloc(3000);  
}
```

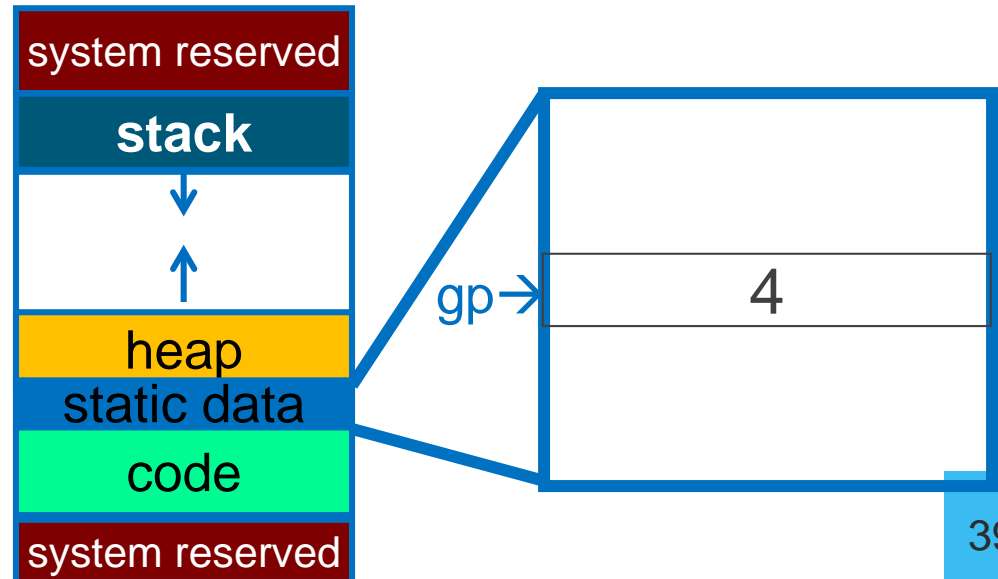


Data Segment

Data segment contains global variables

- Exist for all time, accessible to all routines
- Accessed w/global pointer
 - `gp, x3`, points to middle of segment
 - Example: `lw x5, 0(gp)` gets middle-most word (here, `max_players`)

```
int max_players = 4;  
  
int main(...) {  
    ...  
}
```



Globals and Locals

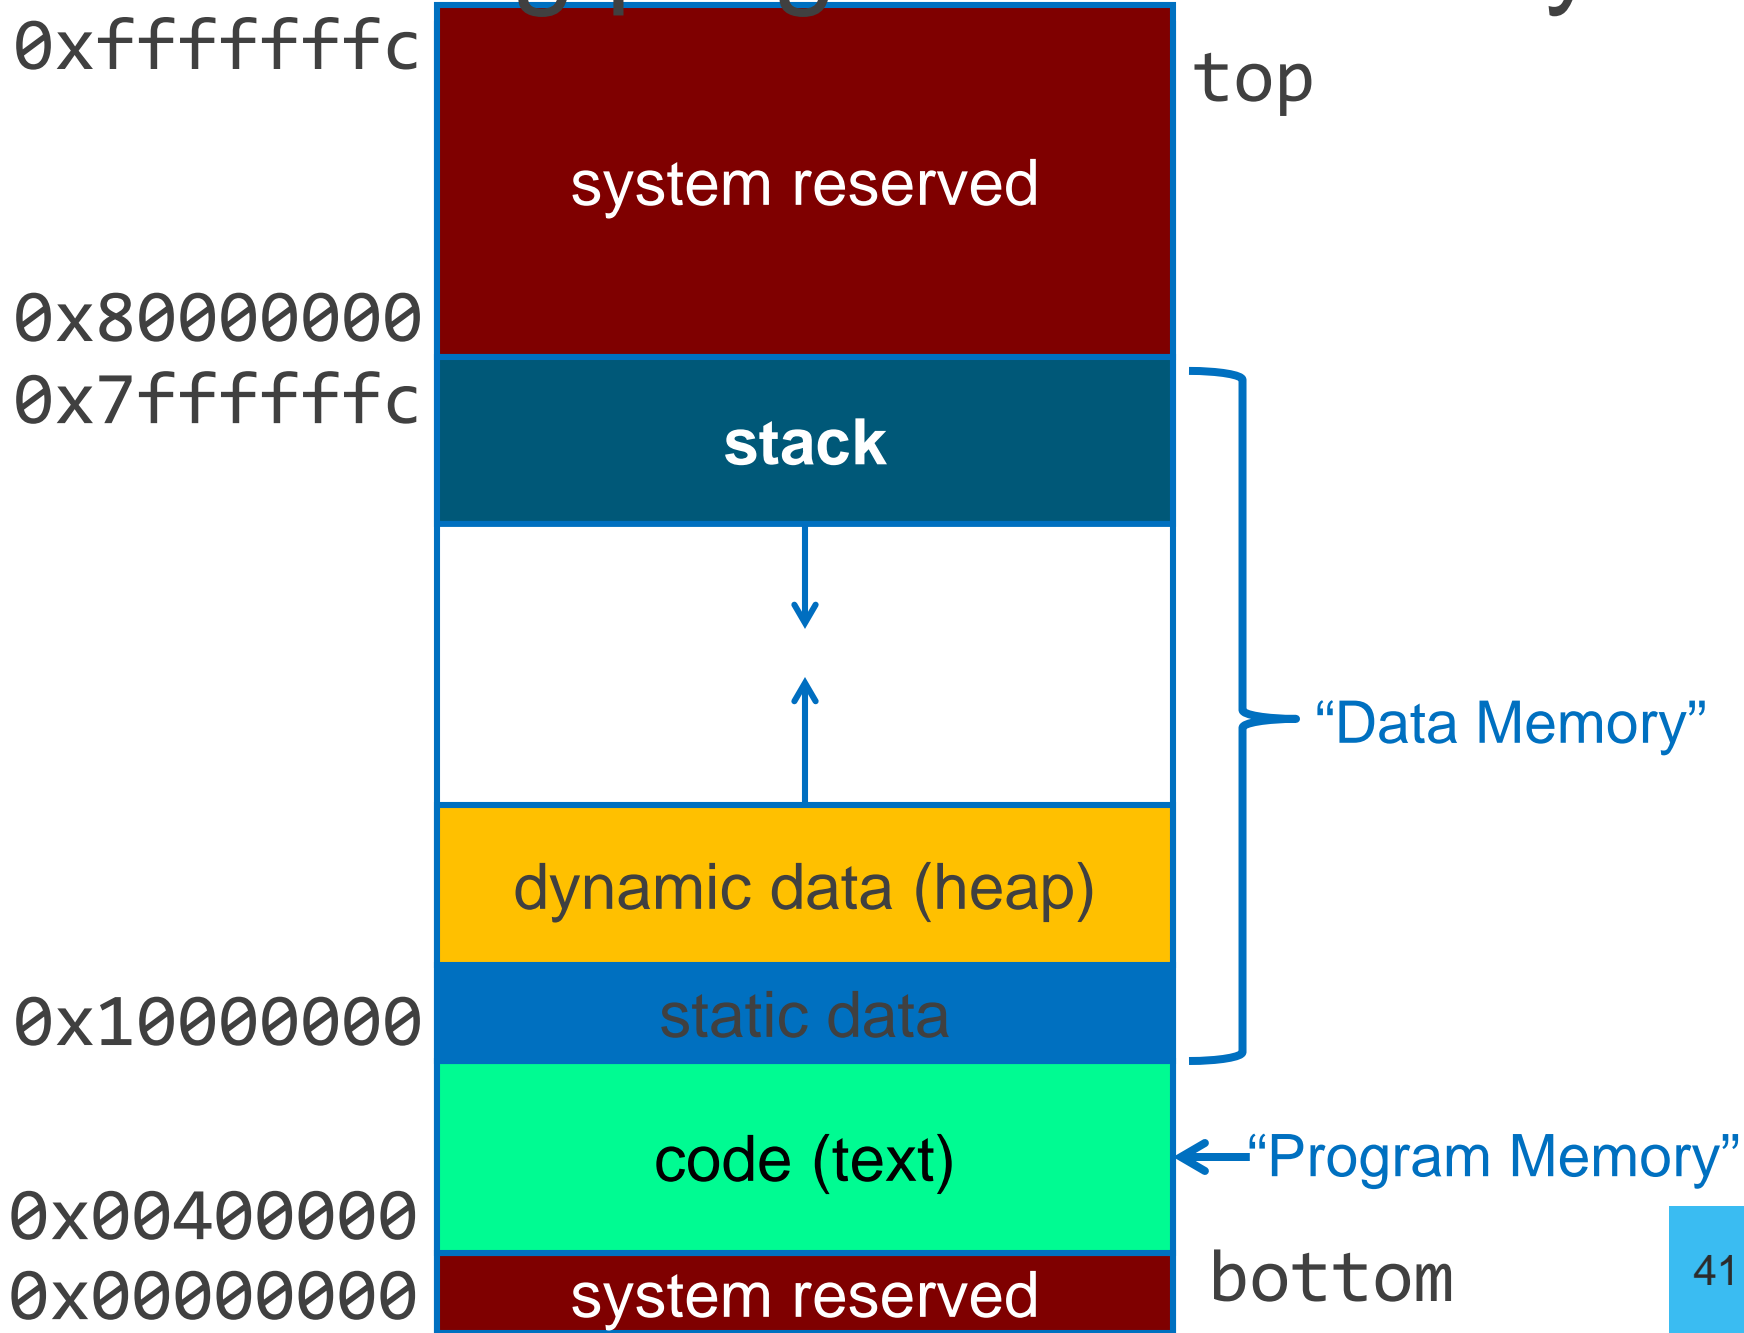
Variables	Visibility	Lifetime	Location
Function-Local			
Global			
Dynamic			

```
int n = 100;
int main (int argc, char* argv[ ]) {
    int i, m = n, sum = 0;
    int* A = malloc(4*m + 4);
    for (i = 1; i <= m; i++) {
        sum += i; A[i] = sum; }
    printf ("Sum 1 to %d is %d\n", n, sum);
}
```

Where is **main** ?

- (A) Stack
- (B) Heap
- (C) Global Data
- (D) Text

An executing program in memory



Globals and Locals

Variables	Visibility	Lifetime	Location
Function-Local <code>i, m, sum, A</code>	w/in function	function invocation	stack
Global <code>n, str</code>	whole program	program execution	.data
Dynamic <code>*A</code>	Anywhere that has a pointer	b/w malloc and free	heap

```
int n = 100;
int main (int argc, char* argv[ ]) {
    int i, m = n, sum = 0;
    int* A = malloc(4*m + 4);
    for (i = 1; i <= m; i++) {
        sum += i; A[i] = sum; }
    printf ("Sum 1 to %d is %d\n", n, sum);
}
```

Takeaway2: Need a Call Stack

JAL (Jump And Link) instruction moves a new value into the PC, and simultaneously saves the old value in register x1 (aka ra or return address). Thus, can get back from the subroutine to the instruction immediately following the jump by transferring control back to PC in register x1.

Need a Call Stack to return to correct calling procedure. To maintain a stack, need to store an **activation record** (aka a “stack frame”) in memory. Stacks keep track of the correct return address by storing the contents of x1 in memory (the stack).

Calling Convention for Procedure Calls

~~Transfer Control~~

- ~~• Caller → Routine~~
- ~~• Routine → Caller~~

Pass Arguments to and from the routine


- fixed length, variable length, recursively
- Get return value back to the caller

Manage Registers

- Allow each routine to use registers
- Prevent routines from clobbering each others' data

Next Goal






Need consistent way of passing arguments and getting the result of a subroutine invocation



Arguments & Return Values

Need consistent way of passing arguments and getting the result of a subroutine invocation

Given a procedure signature, need to know where arguments should be placed

- `int min(int a, int b);` 
- `int subf(int a, int b, int c, int d, int e, int f, int g, int h, int i);` 
- `int isalpha(char c);` stack?
- `int treesort(struct Tree *root);` 
- `struct Node *createNode();` 
- `struct Node mynode();` 

Too many combinations of char, short, int, void *, struct, etc.

- RISC-V treats char, short, int and void * identically

Simple Argument Passing (1-8 args)

```
main() {  
    int x = myfn(6, 7);  
    x = x + 2;  
}
```

```
main:  
    li x10, 6  
    li x11, 7  
    jal myfn  
    addi x5, x10, 2
```

First eight arguments:

passed in registers x10-x17

- aka \$a0, \$a1, ..., \$a7

Returned result:

passed back in a register

- Specifically, x10, aka a0
- And x11, aka a1

Note: This is *not* the entire story for 1-8 arguments.
Please see *the Full Story* slides.

Conventions so far:


- args passed in `$a0, $a1, ..., $a7`
- return value (if any) in `$a0, $a1`
- stack frame at `$sp`
 - contains `$ra` (clobbered on JAL to sub-functions)

Q: What about argument lists?

Many Arguments (8+ args)

```
main() {  
    myfn(0,1,2,...,7,8,9);  
    ...  
}
```

```
main:  
    li x10, 0  
    li x11, 1  
    ...  
    li x17, 7  
    li x5, 8  
    sw x5, -8(x2)  
    li x5, 9  
    sw x5, -4(x2)  
    jal myfn
```



First eight arguments:

passed in registers x10-x17

- aka `a0, a1, ..., a7`

Subsequent arguments:

"spill" onto the stack


Args passed in child's
stack frame

Note: This is *not* the entire story for 9+ args.
Please see *the Full Story* slides.

Many Arguments (8+ args)

```
main() {  
    myfn(0,1,2,...,7,8,9);  
    ...  
}
```

```
main:  
    li a0, 0  
    li a1, 1  
    ...  
    li a7, 7  
    li t0, 8  
    sw t0, -8(sp)  
    li t0, 9  
    sw t0, -4(sp)  
    jal myfn
```



First eight arguments:

passed in registers x10-x17

- aka `a0, a1, ..., a7`

Subsequent arguments:

"spill" onto the stack


Args passed in child's
stack frame

Note: This is *not* the entire story for 9+ args.
Please see *the Full Story* slides.

Argument Passing: *the Full Story*

```
main() {  
    myfn(0,1,2,...,7,8,9);  
    ...  
}
```

```
main:  
    li a0, 0  
    li a1, 1  
    ...  
    li a7, 7  
    li t0, 8  
    sw t0, -8(x2)  
    li t0, 9  
    sw t0, -4(x2)  
    jal myfn
```



Arguments 1-8:
passed in x10-x17
room on stack

Arguments 9+:
placed on stack

Args passed in
child's stack frame

Pros of Argument Passing Convention

- Consistent way of passing arguments to and from subroutines
- Creates single location for all arguments
 - Caller makes room for a0-a7 on stack
 - Callee must copy values from a0-a7 to stack
 - callee may treat all args as an array in memory
 - Particularly helpful for functions w/ variable length inputs:
`printf("Scores: %d %d %d\n", 1, 2, 3);`
- Aside: not a bad place to store inputs if callee needs to call a function (your input cannot stay in \$a0 if you need to call another function!)

iClicker Question

Which is a true statement about the arguments to the function

```
void sub(int a, int b, int c, int d, int e, int f,  
int g, int h, int i);
```

- A. Arguments a-i are all passed in registers.
- B. Arguments a-i are all stored on the stack.
- C. Only i is stored on the stack,
but space is allocated for all 9 arguments.
- D. Only a-h are stored on the stack,
but space is allocated for all 9 arguments.


iClicker Question

Which is a true statement about the arguments to the function

```
void sub(int a, int b, int c, int d, int e, int f,  
int g, int h, int i);
```

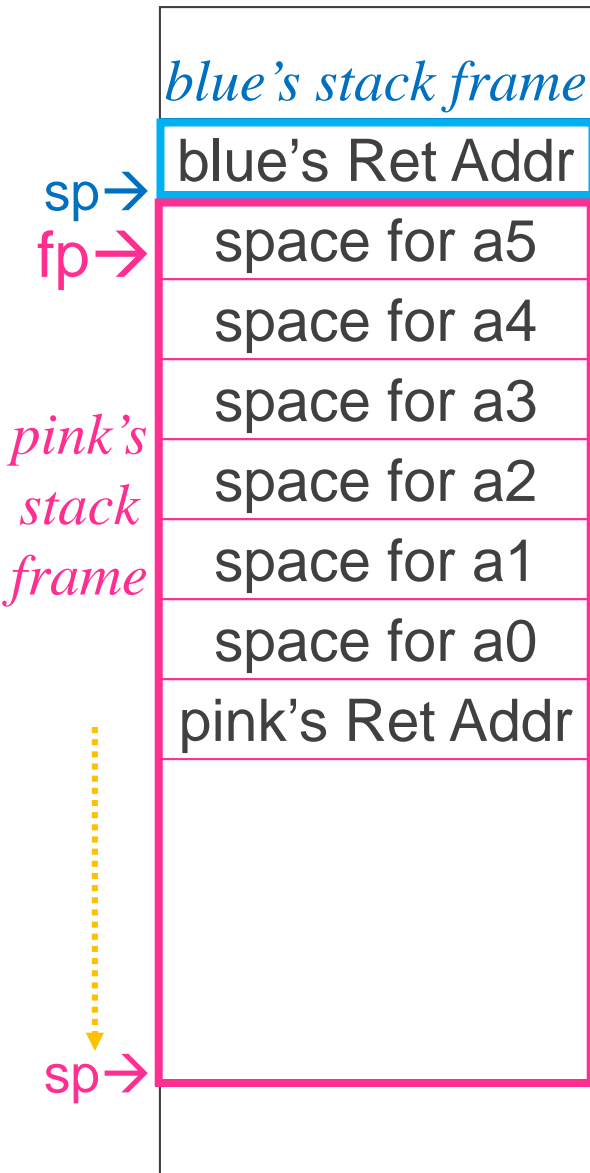
- A. Arguments a-i are all passed in registers.
- B. Arguments a-i are all stored on the stack.
- C. Only i is stored on the stack, but space is allocated for all 9 arguments.
- D. Only a-h are stored on the stack, but space is allocated for all 9 arguments.

Frame Layout & the Frame Pointer



```
blue() {  
    pink(0,1,2,3,4,5);  
}
```

Frame Layout & the Frame Pointer



Notice

- Pink's arguments are on pink's stack
 - **sp** changes as functions call other functions, complicates accesses
- Convenient to keep pointer to bottom of stack == **frame pointer**
x8, aka fp (also known as s0)
can be used to restore **sp** on exit

```
blue() {  
    pink(0,1,2,3,4,5);  
}  
pink(int a, int b, int c, int d, int e, int f) {  
    ...  
}
```

Conventions so far

- **first eight** arg words passed in \$a0, \$a1, ..., \$a7
- Space for args **in child's stack frame**
- return value (if any) in \$a0, \$a1
- stack frame (\$fp/\$s0 to \$sp) contains:
 - \$ra (clobbered on JAL to sub-functions)
 - space for 8 arguments to Callees
 - arguments 9+ to Callees

RISCV Register Conventions so far:

x0	zero	zero	x16		
x1	ra	Return address	x17		
x2	sp	Stack pointer	x18	s2	Saved registers
x3			x19		
x4			x20		
x5	t0	Temporary registers	x21		
x6	t1		x22		
x7	t2		x23		
x8	s0/fp	Saved register or framepointer	x24		
x9	s1	Saved register	x25		
x10	a0	Function args or return values	x26		
x11	a1		x27		
x12	a2	Function args	x28	t3	Temporary registers
x13	a3		x29		
x14	a4		x30		
			x31		

Globals and Locals

Global variables are allocated in the “data” region of the program

- Exist for all time, accessible to all routines


Local variables are allocated within the stack frame

- Exist solely for the duration of the stack frame

Dangling pointers are pointers into a destroyed stack frame

- C lets you create these, Java does not

- `int *foo() { int a; return &a; }`



Return the address of `a`,
But `a` is stored on stack, so will be removed
when call returns and point will be invalid

Global and Locals

How does a function load global data?

- global variables are just above 0x10000000

Convention: *global pointer*

- *x3* is *gp* (pointer into *middle* of global data section)
gp = 0x10000800
- Access most global data using LW at gp +/- offset
LW t0, 0x800(gp)
LW t1, 0x7FF(gp)

Anatomy of an executing program

0xfffffffffc

top

system reserved

0x80000000

0x7fffffff

stack



dynamic data (heap)

\$gp

0x10000000

static data

code (text)

0x00400000

0x00000000

system reserved

bottom

Frame Pointer

It is often cumbersome to keep track of location of data on the stack

- The offsets change as new values are pushed onto and popped off of the stack

Keep a pointer to the bottom of the top stack frame

- Simplifies the task of referring to items on the stack

A frame pointer, **x8**, aka **fp/s0**

- Value of sp upon procedure entry
- Can be used to restore sp on exit

Conventions so far

- first eight arg words passed in a0-a7
- Space for args in child's stack frame
- return value (if any) in a0, a1
- stack frame (fp/s0 to sp) contains:
 - ra (clobbered on JALs)
 - space for 8 arguments
 - arguments 9+
- global data accessed via gp

Calling Convention for Procedure Calls

~~Transfer Control~~

- ~~• Caller → Routine~~
- ~~• Routine → Caller~~

~~Pass Arguments to and from the routine~~


- ~~• fixed length, variable length, recursively~~
- ~~• Get return value back to the caller~~

Manage Registers

- Allow each routine to use registers
- Prevent routines from clobbering each others' data

Next Goal

What convention should we use to share use of registers across procedure calls?



Register Management

Functions:

- Are compiled in isolation
- Make use of general purpose registers
- Call other functions in the middle of their execution
 - These functions also use general purpose registers!
 - No way to coordinate between caller & callee

→ Need a convention for register management

Register Usage

Suppose a routine would like to store a value in a register

Two options: *callee-save* and *caller-save*

Callee-save:

- Assume that one of the callers is already using that register to hold a value of interest
- Save the previous contents of the register on procedure entry, restore just before procedure return
- E.g. \$ra, \$fp/\$s0, \$s1-\$s11, \$gp, \$tp
- Also, \$sp

Caller-save:

- Assume that a caller can clobber any one of the registers
- **Save** the previous contents of the register **before** proc call
- **Restore after** the call
- E.g. \$a0-a7, \$t0-\$t6

RISCV calling convention supports both

Caller-saved

Registers that the caller cares about: t0... t9

About to call a function?

- Need value in a t-register *after* function returns?
 - **save** it to the stack **before** fn call
 - **restore** it from the stack after fn returns
- Don't need value? → do nothing

Suppose:
t0 holds x
t1 holds y
t2 holds z

Where do we save and restore?

Functions

- Can freely use these registers
- Must assume that their contents are destroyed by other functions

```
void myfn(int a) {  
    int x = 10;  
    int y = max(x, a);  
    int z = some_fn(y);  
    return (z + y);  
}
```

Callee-saved

Registers a function intends to use: s0... s9

About to use an s-register? You **MUST**:

- Save the current value on the stack **before** using
- Restore the old value from the stack before fn returns

Suppose:

s1 holds x

s2 holds y

s3 holds z

Functions

Where do we save and restore?

- Must save these registers before using them
- May assume that their contents are preserved even across fn calls

```
void myfn(int a) {  
    int x = 10;  
    int y = max(x, a);  
    int z = some_fn(y);  
    return (z + y);  
}
```

Caller-Saved Registers in Practice

```
main:
```

```
...
```

```
[use x5 & x6]
```

```
...
```

```
addi x2, x2, -8
```

```
sw x6, 4(x2)
```

```
sw x5, 0(x2)
```

```
jal myfn
```

```
lw x6, 4(x2)
```

```
lw x5, 0(x2)
```

```
addi x2, x2, 8
```

```
...
```

```
[use x5 & x6]
```

Assume the registers are free for the taking, use with no overhead

Since subroutines will do the same, must protect values needed later:

Save before fn call

Restore after fn call

Notice: Good registers to use if you don't call too many functions or if the values don't matter later on anyway.

Caller-Saved Registers in Practice

main:

...

[use \$t0 & \$t1]

...

addi \$sp, \$sp, -8

sw \$t1, 4(\$sp)

sw \$t0, 0(\$sp)

jal myfn

lw \$t1, 4(\$sp)

lw \$t0, 0(\$sp)

addi \$sp, \$sp, 8

...

[use \$t0 & \$t1]

Assume the registers are free for the taking, use with no overhead

Since subroutines will do the same, must protect values needed later:

Save before fn call

Restore after fn call

Notice: Good registers to use if you don't call too many functions or if the values don't matter later on anyway.

Callee-Saved Registers in Practice

main:

```
addi x2, x2, -16
```

```
sw x1, 12(x2)
```

```
sw x8, 8(x2)
```

```
sw x18, 4(x2)
```

```
sw x9, 0(x2)
```

```
addi x8, x2, 12
```

...

```
[use x9 and x18]
```

...

```
lw x1, 12(x2)
```

```
lw x8, 8(x2)
```

```
lw x18, 4(x2)
```

```
lw x9, 0(x2)
```

```
addi x2, x2, 16
```

```
jr x1
```

Assume caller is using the registers

Save on entry

Restore on exit

Notice: Good registers to use if you make a lot of function calls and need values that are preserved across all of them.

Also, good if caller is actually using the registers, otherwise the save and restores are wasted. But hard to know this.

Callee-Saved Registers in Practice

main:

```
addi $sp, $sp, -16
```

```
sw $ra, 12($sp)
```

```
sw $fp, 8($sp)
```

```
sw $s2, 4($sp)
```

```
sw $s1, 0($sp)
```

```
addi $fp, $sp, 12
```

...

```
[use $s1 and $s2]
```

...

```
lw $ra, 12($sp)
```

```
lw $fp, 8($sp)
```

```
lw $s2, 4($sp)
```

```
lw $s1, 0($sp)
```

```
addi $sp, $sp, 16
```

```
jr $ra
```

Assume caller is using the registers

Save on entry

Restore on exit

Notice: Good registers to use if you make a lot of function calls and need values that are preserved across all of them.

Also, good if caller is actually using the registers, otherwise the save and restores are wasted. But hard to know this.

Clicker Question

```
int foo() {
    int a = 0;
    int b = 12;
    int c = 1;

    while(b + c > 0) {
        int e = b + bar(c);
        c = b + e;

        int d = c + baz(b);
        a = d - e;
    }

    return a;
}
```

You are a compiler. Do you choose to put **a** in a:

- (A) Caller-saved register (t)
- (B) Callee-saved register (s)
- (C) Depends on where we put the other variables in this fn
- (D) Both are equally valid

Clicker Question

```
int foo() {
    int a = 0;
    int b = 12;
    int c = 1;

    while(b + c > 0) {
        int e = b + bar(c);
        c = b + e;

        int d = c + baz(b);
        a = d - e;
    }

    return a;
}
```

You are a compiler. Do you choose to put **a** in a:

- (A) Caller-saved register (t)
- (B) Callee-saved register (s)
- (C) Depends on where we put the other variables in this fn
- (D) Both are equally valid

Repeat but assume that foo is recursive (bar/baz → foo)

Clicker Question

```
int foo() {
    int a = 0;
    int b = 12;
    int c = 1;

    while(b + c > 0) {
        int e = b + bar(c);
        c = b + e;


        int d = c + baz(b);
        a = d - e;
    }

    return a;
}
```

You are a compiler. Do you choose to put **b** in a:

- (A) Caller-saved register (t)
- (B) Callee-saved register (s)
- (C) Depends on where we put the other variables in this fn
- (D) Both are equally valid

Frame Layout on Stack



Assume a function uses two callee-save registers.

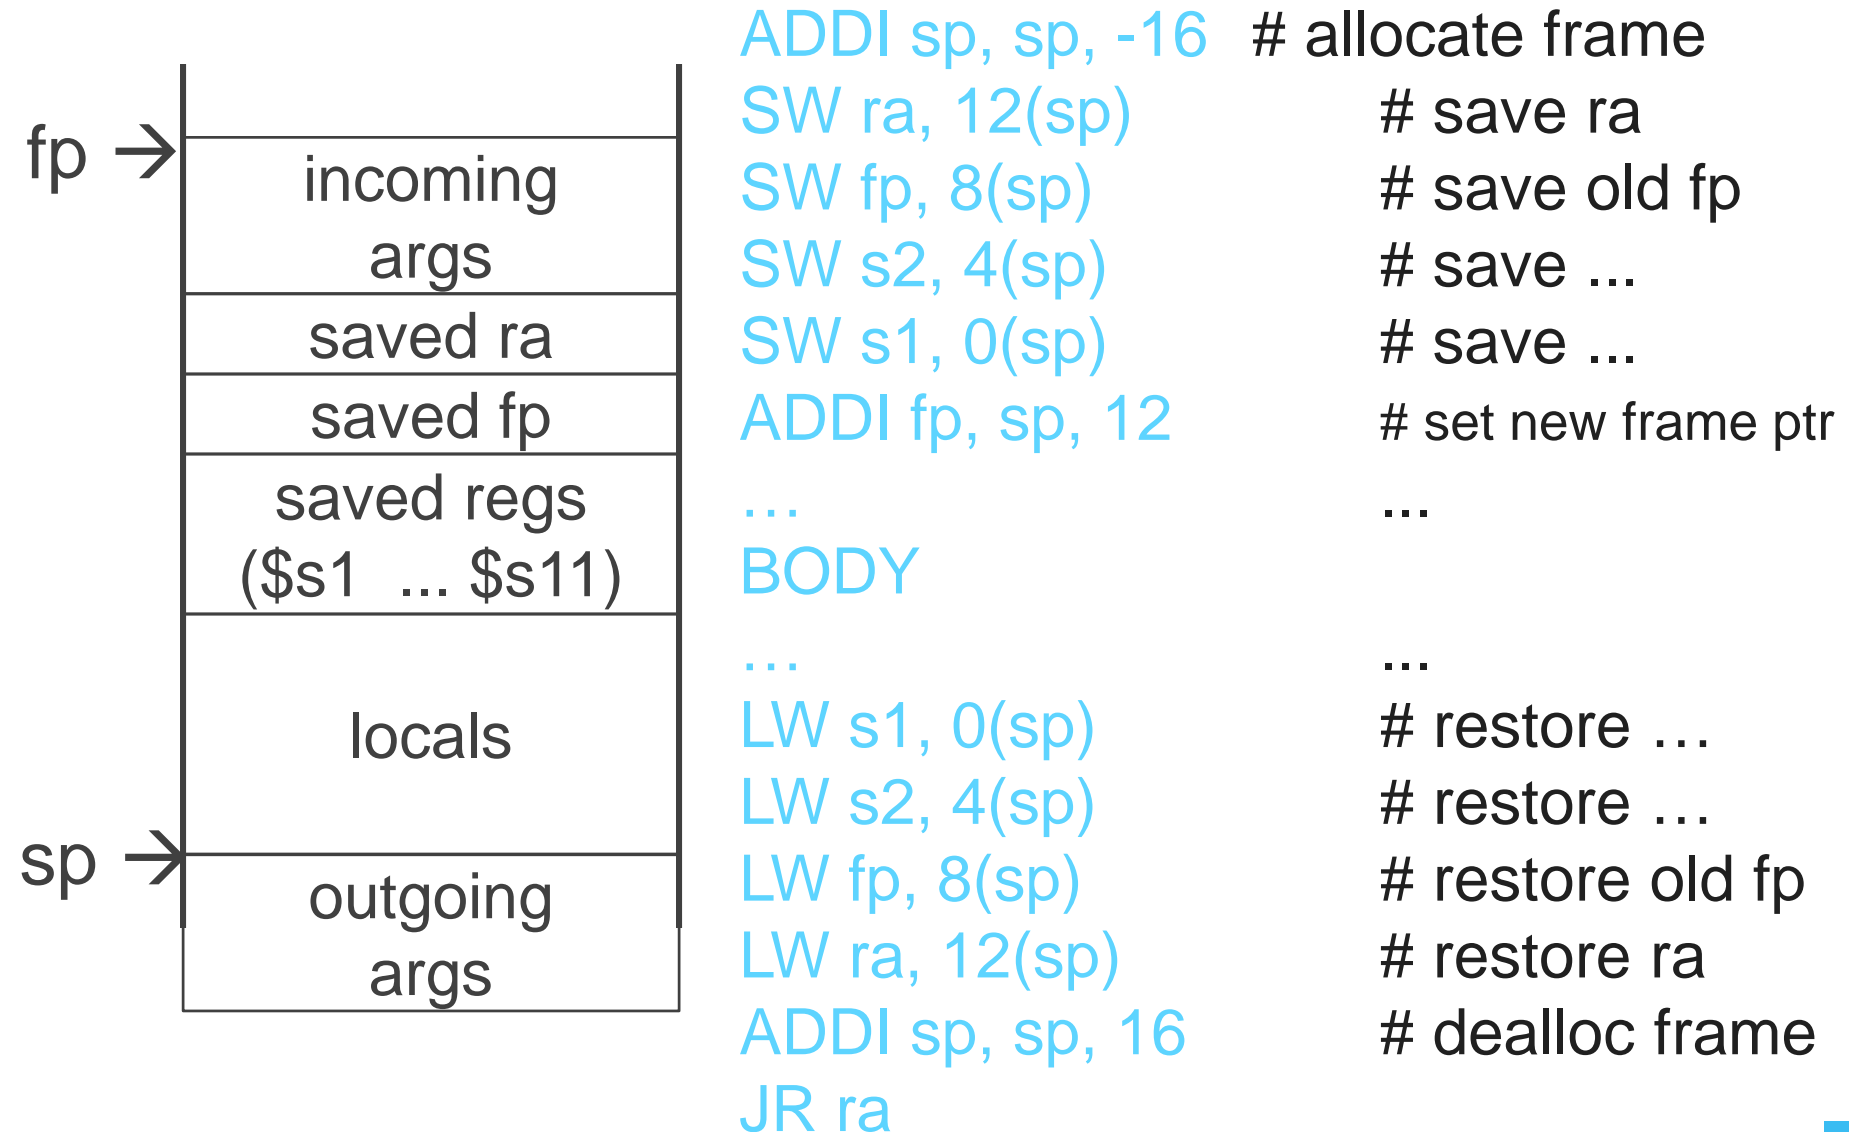
How do we allocate a stack frame?

How large is the stack frame?

What should be stored in the stack frame?

Where should everything be stored?

Frame Layout on Stack



Frame Layout on Stack

fp→
blue's
stack
frame
sp→

blue's ra

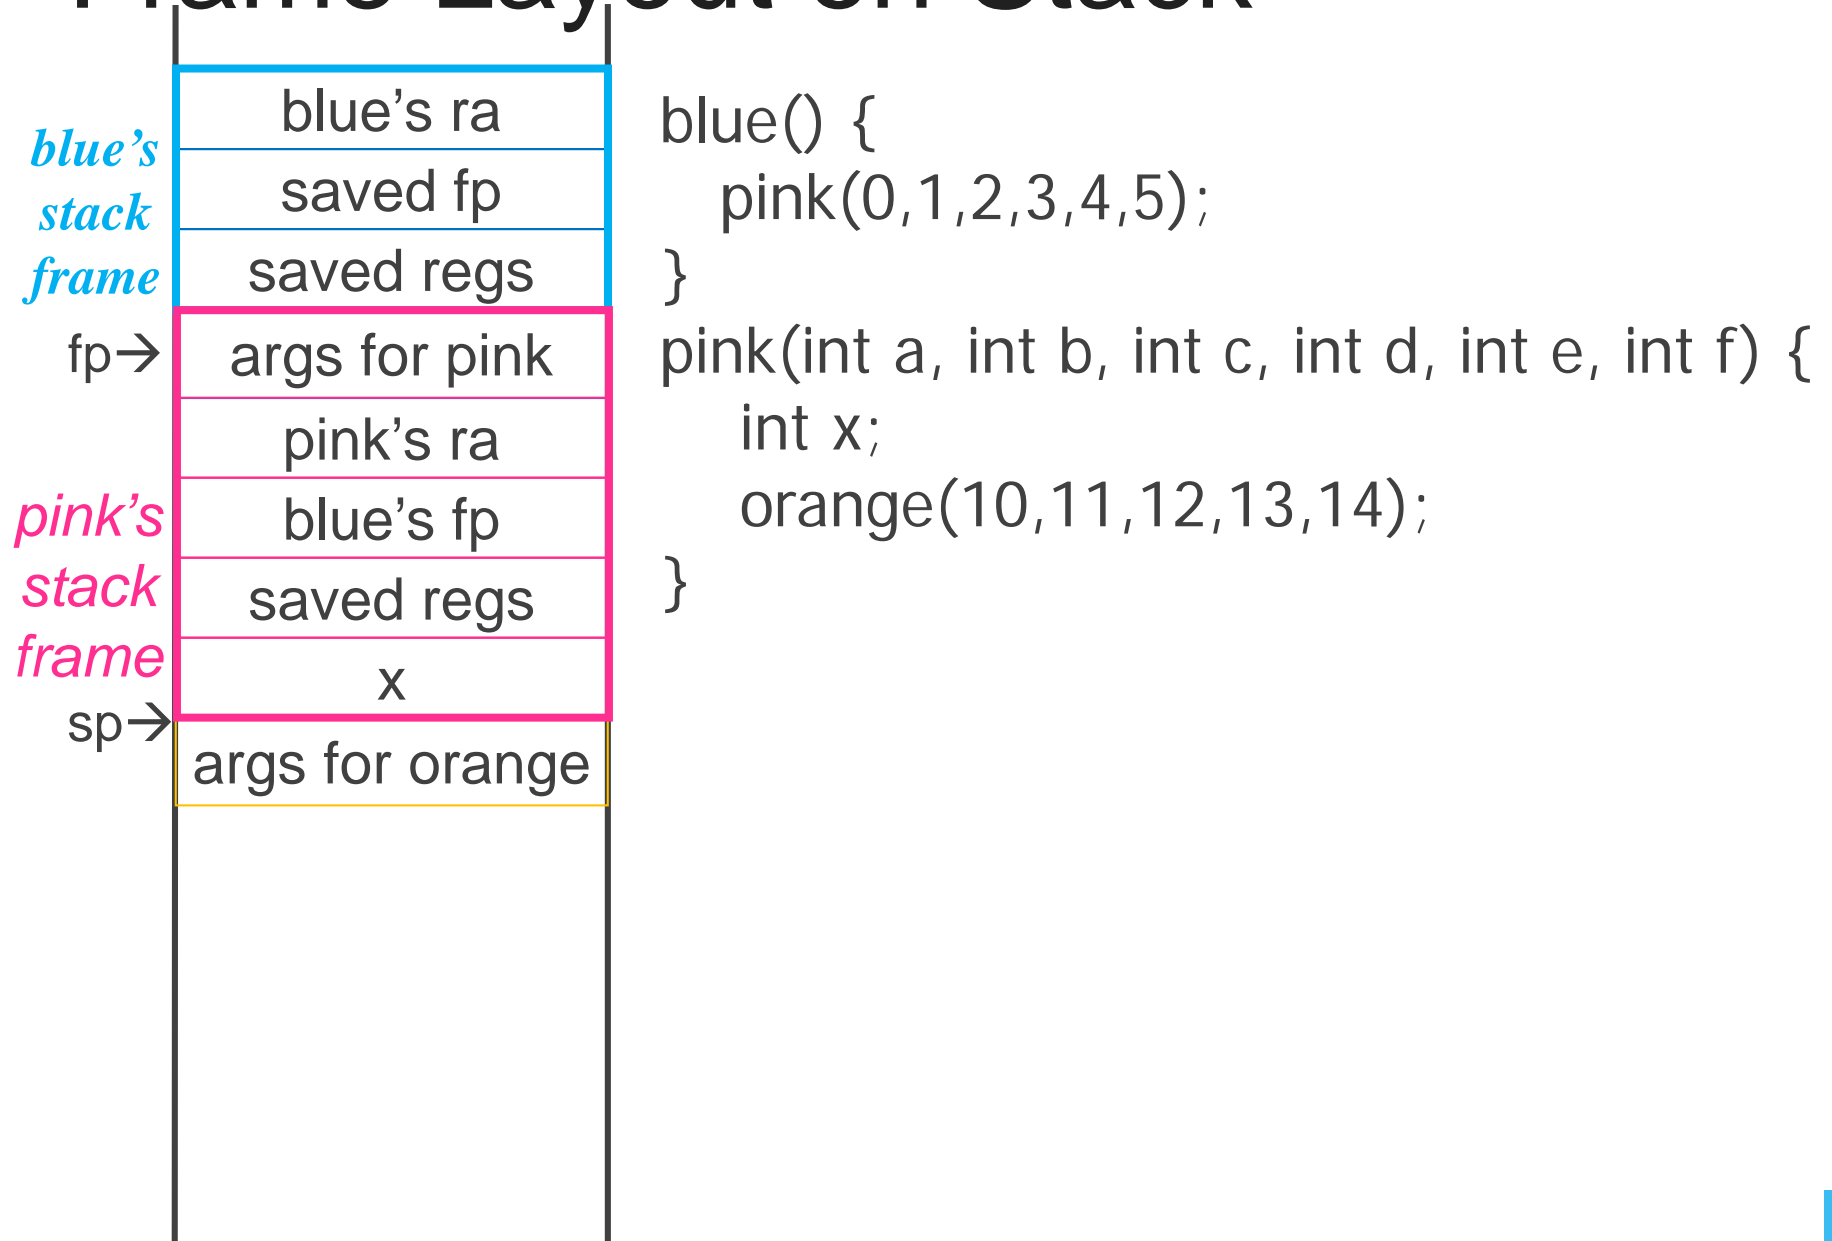
saved fp

saved regs

args for pink

```
blue() {  
    pink(0,1,2,3,4,5);  
}
```

Frame Layout on Stack



Frame Layout on Stack



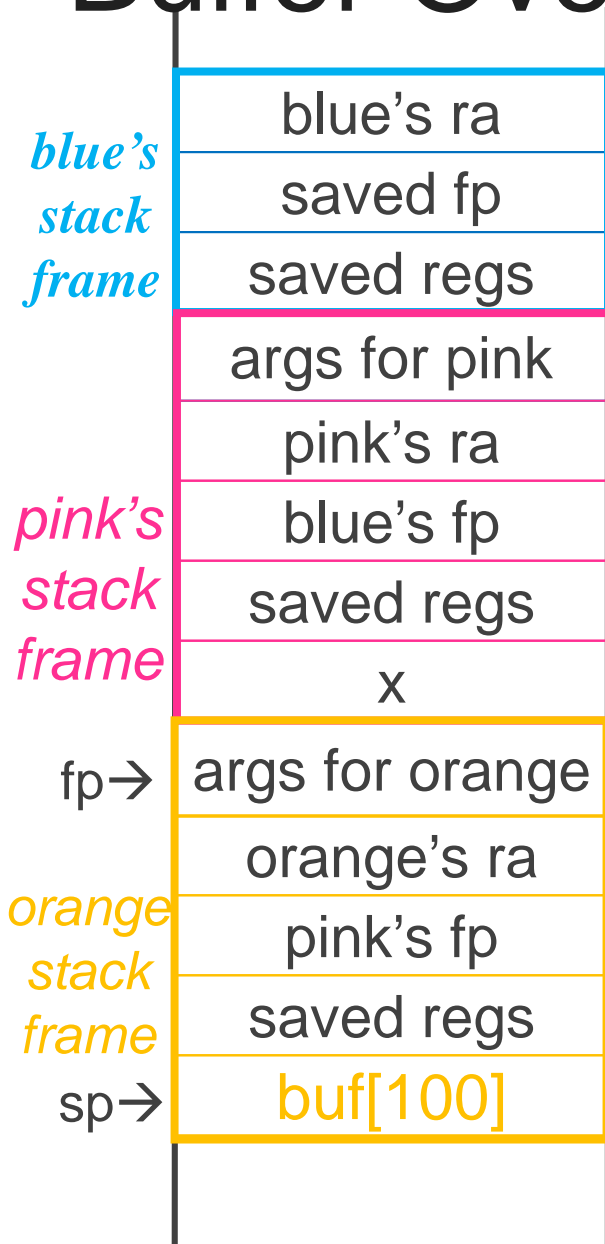
```
blue() {  
    pink(0,1,2,3,4,5);  
}
```

```
pink(int a, int b, int c, int d, int e, int f) {  
    int x;  
    orange(10,11,12,13,14);  
}
```

```
orange(int a, int b, int c, int, d, int e) {  
    char buf[100];  
    gets(buf);    // no bounds check!  
}
```

What happens if more than 100 bytes is written to buf?

Buffer Overflow



```
blue() {  
    pink(0,1,2,3,4,5);  
}
```

```
pink(int a, int b, int c, int d, int e, int f) {  
    int x;  
    orange(10,11,12,13,14);  
}
```

```
orange(int a, int b, int c, int, d, int e) {  
    char buf[100];  
    gets(buf);    // no bounds check!  
}
```

What happens if more than 100 bytes is written to buf?

RISCV Register Recap

Return address: x1 (ra)

Stack pointer: x2 (sp)

Frame pointer: x8 (fp/s0)

First four arguments: x10-x17 (a0-a7)

Return result: x10-x11 (a0-a1)

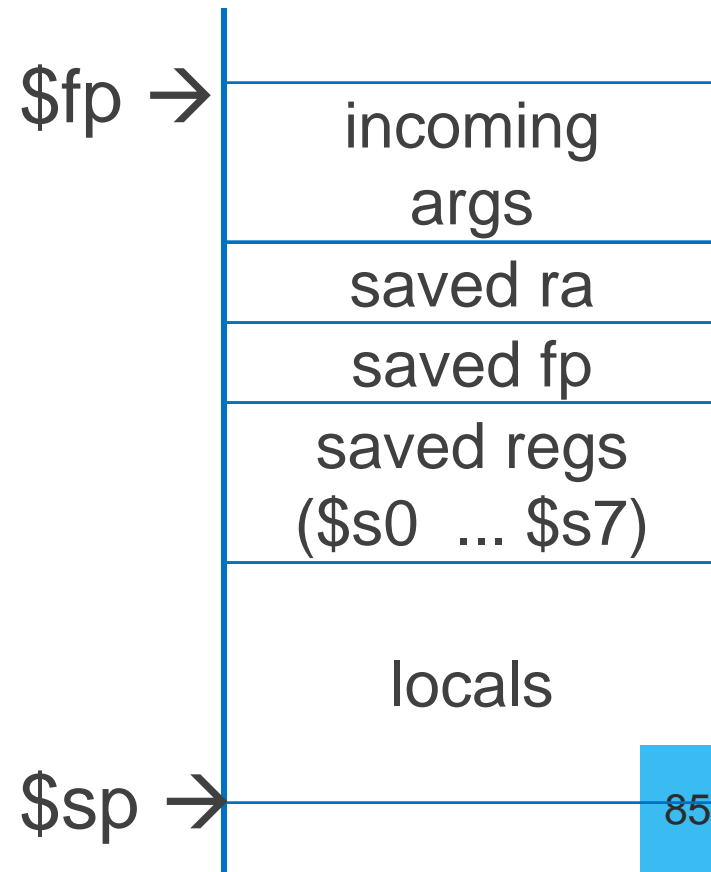
Callee-save free regs: x9, x18-x27 (s1-s11)

Caller-save (temp) free regs: x5-x7, x28-x31 (t0-t6)

Global pointer: x3 (gp)

Convention Summary

- first eight arg words passed in \$a0-\$a7
- Space for args in child's stack frame
- return value (if any) in \$a0, \$a1
- stack frame (\$fp to \$sp) contains:
 - \$ra (clobbered on JALs)
 - local variables
 - space for 8 arguments to Callees
 - arguments 9+ to Callees
- callee save regs: preserved
- caller save regs: not preserved
- global data accessed via \$gp



Activity #1: Calling Convention Example

```
int test(int a, int b) {  
    int tmp = (a&b)+(a|b);  
    int s = sum(tmp,1,2,3,4,5,6,7,8);  
    int u = sum(s,tmp,b,a,b,a);  
    return u + a + b;  
}
```

Correct Order:

1. Body First
2. Determine stack frame size
3. Complete Prologue/Epilogue

Activity #1: Calling Convention Example

```
int test(int a, int b) {  
    int tmp = (a&b)+(a|b);  
    int s = sum(tmp,1,2,3,4,5,6,7,8);  
    int u = sum(s,tmp,b,a,b,a);  
    return u + a + b;  
}
```

test:

Prologue

```
MOVE s1, a0  
MOVE s2, a1  
AND t0, a0, a1  
OR t1, a0, a1  
ADD t0, t0, t1  
MOVE a0, t0  
LI a1, 1  
LI a2, 2  
...  
LI a7, 7  
LI t1, 8  
SW t1, -4(sp)  
SW t0, 0(sp)  
JAL sum
```

LW t0, 0(sp)

```
MOVE a0, a0 # s  
MOVE a1, t0 # tmp  
MOVE a2, s2 # b  
MOVE a3, s1 # a  
MOVE a4, s2 # b  
MOVE a5, s1 # a  
JAL sum
```

add u (a0) and a (s1)

```
ADD a0, a0, s1
```

```
ADD a0, a0, s2
```

a0 = u + a + b

Epilogue

Activity #1: Calling Convention Example

```
int test(int a, int b) {  
    int tmp = (a&b)+(a|b);  
    int s =sum(tmp,1,2,3,4,5,6,7,8);  
    int u = sum(s,tmp,b,a,b,a);  
    return u + a + b;  
}
```

How many bytes do we need to allocate for the stack frame?

a) 24

b) 28

c) 36

d) 40

e) 48

test:

Prologue

```
MOVE s1, a0  
MOVE s2, a1  
AND t0, a0, a1  
OR t1, a0, a1  
ADD t0, t0, t1  
MOVE a0, t0  
LI a1, 1  
LI a2, 2  
...  
LI a7, 7  
LI t1, 8  
SW t1, -4(sp)
```

SW t0, 0(sp)

JAL sum

LW t0, 0(sp)

```
MOVE a0, a0 # s  
MOVE a1, t0 # tmp  
MOVE a2, s2 # b  
MOVE a3, s1 # a  
MOVE a4, s2 # b  
MOVE a5, s1 # a  
JAL sum
```

add u (v0) and a (s1)

```
ADD a0, a0, s1
```


```
ADD a0, a0, s2
```

a0 = u + a + b

Epilogue

Activity #1: Calling Convention Example

```
int test(int a, int b) {  
    int tmp = (a&b)+(a|b);  
    int s =sum(tmp,1,2,3,4,5,6,7,8);  
    int u = sum(s,tmp,b,a,b,a);  
    return u + a + b;  
}
```



test:

Prologue

```
MOVE s1, a0  
MOVE s2, a1  
AND t0, a0, a1  
OR t1, a0, a1  
ADD t0, t0, t1  
MOVE a0, t0  
LI a1, 1  
LI a2, 2  
...  
LI a7, 7  
LI t1, 8  
SW t1, -4(sp)  
  
SW t0, 0(sp)  
JAL sum
```

LW t0, 0(sp)

```
MOVE a0, v0 # s  
MOVE a1, t0 # tmp  
MOVE a2, s2 # b  
MOVE a3, s1 # a  
MOVE a4, s2 # b  
MOVE a5, s1 # a  
JAL sum
```

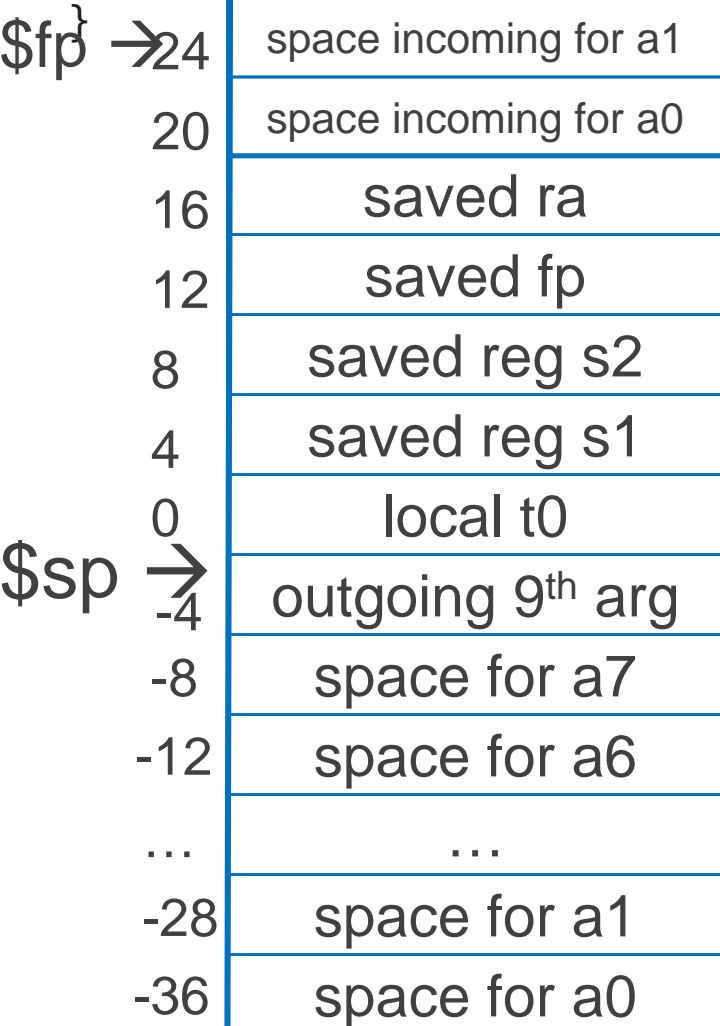
add u (a0) and a (s1)

```
ADD a0, a0, s1  
ADD a0, a0, s2  
# a0 = u + a + b
```

Epilogue

Activity #1: Calling Convention Example

```
int test(int a, int b) {
    int tmp = (a&b)+(a|b);
    int s =sum(tmp,1,2,3,4,5,6,7,8);
    int u = sum(s,tmp,b,a,b,a);
    return u + a + b;
}
```



test:

Prologue

```
MOVE s1, a0
MOVE s2, a1
AND t0, a0, a1
OR t1, a0, a1
ADD t0, t0, t1
MOVE a0, t0
LI a1, 1
LI a2, 2
...
LI a7, 7
LI t1, 8
SW t1, -4(sp)

SW t0, 0(sp)
JAL sum
```

LW t0, 0(sp)

```
MOVE a0, a0 # s
MOVE a1, t0 # tmp
MOVE a2, s2 # b
MOVE a3, s1 # a
MOVE a4, s2 # b
MOVE a5, s1 # a
JAL sum
```

add u (a0) and a (s1)

```
ADD a0, a0, s1
ADD a0, a0, s2
```

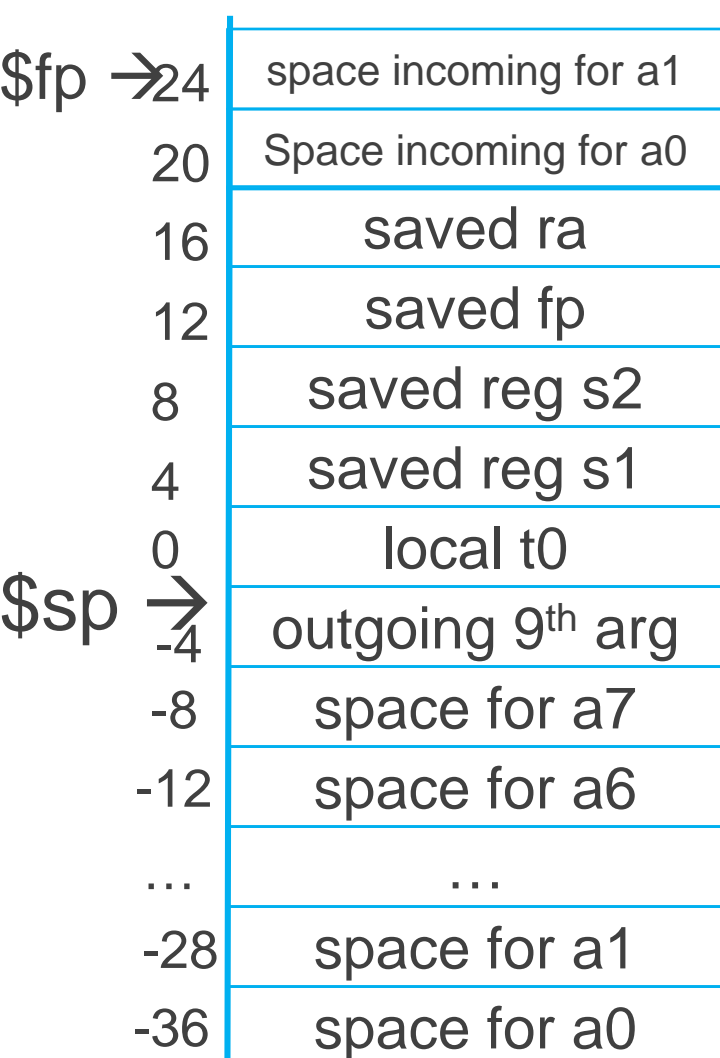
a0 = u + a + b

Epilogue

Activity #2: Calling Convention Example:

Prologue, Epilogue

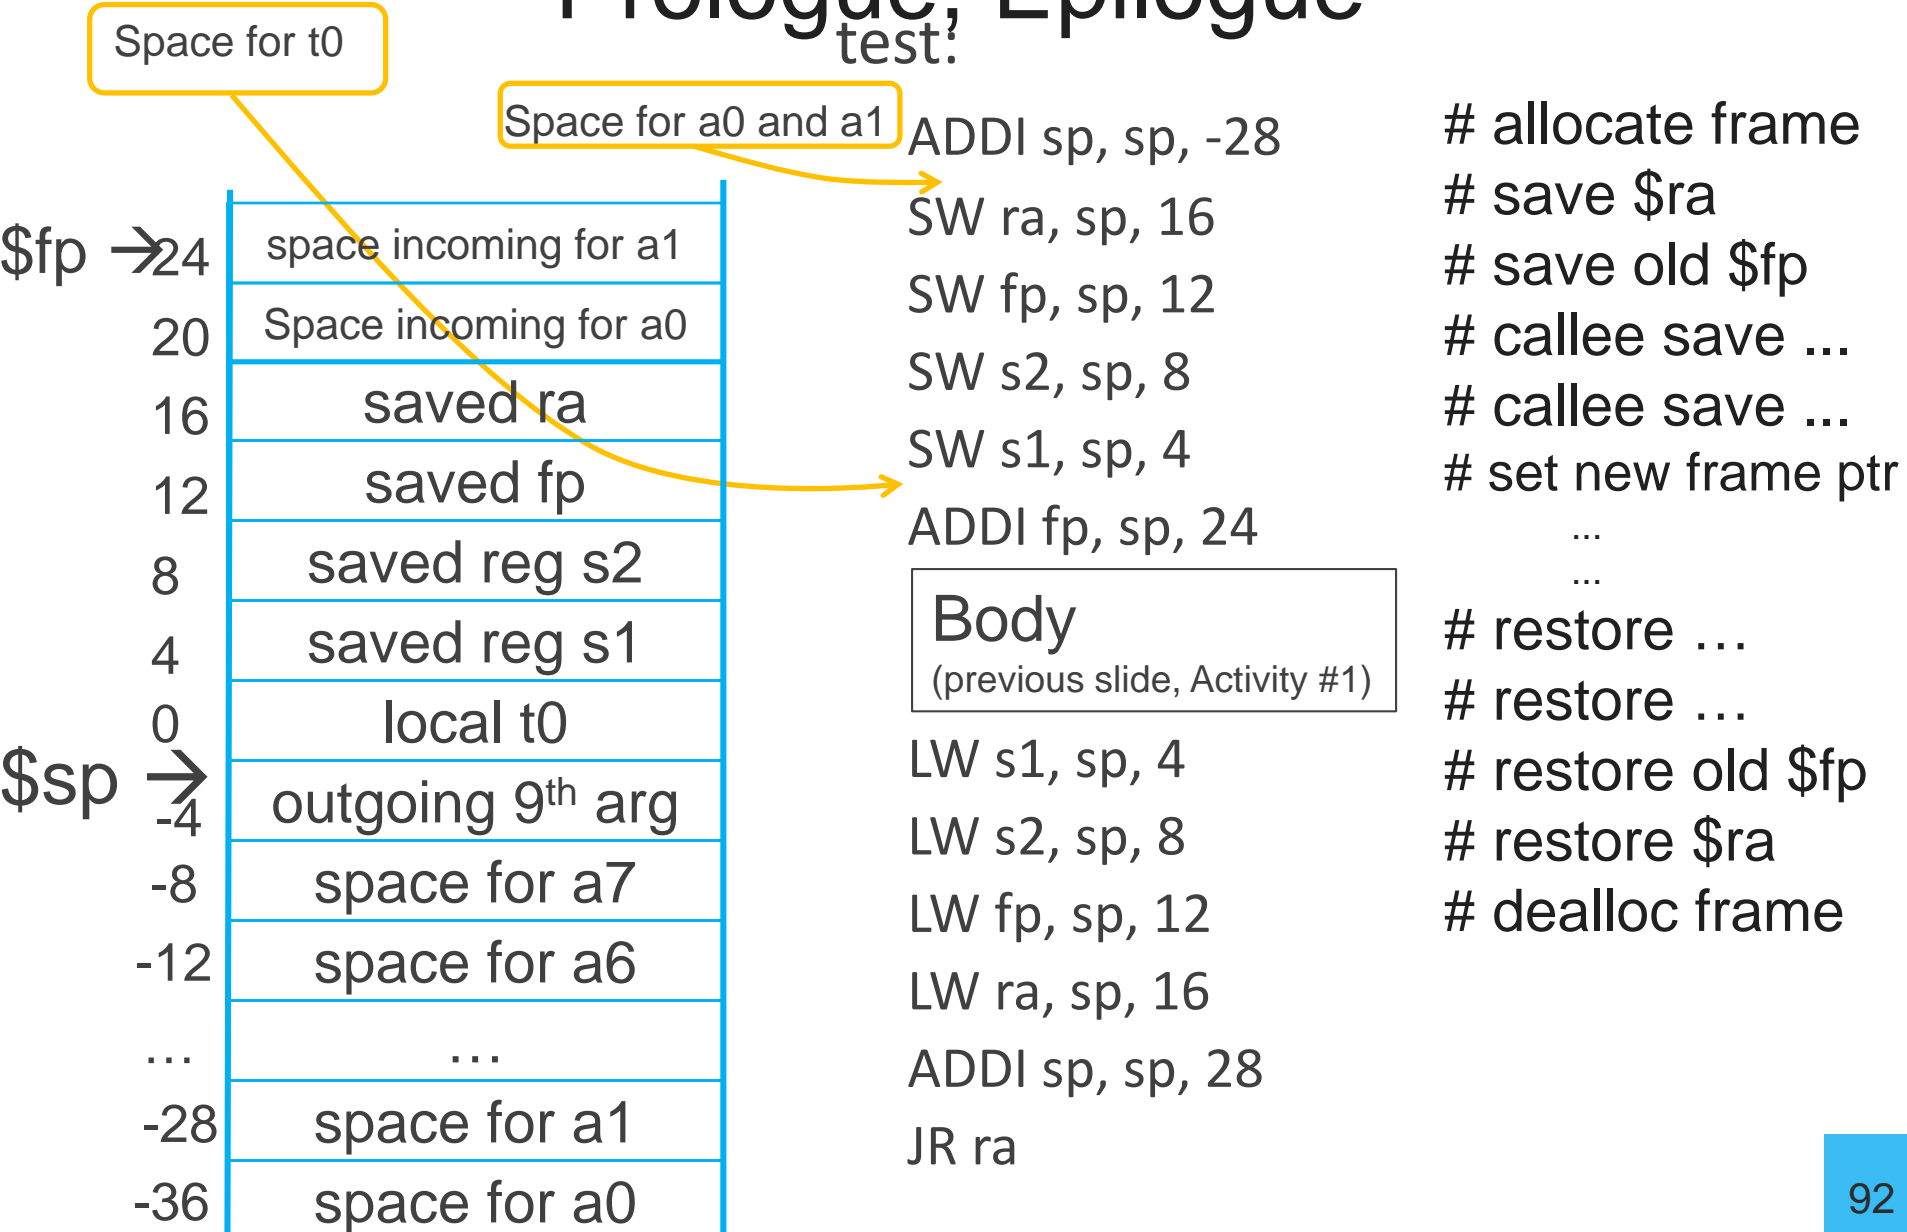
test:



```


# allocate frame
# save $ra
# save old $fp
# callee save ...
# callee save ...
# set new frame ptr
•   ...
•   ...
# restore ...
# restore ...
# restore old $fp
# restore $ra
# dealloc frame
    
```

Activity #2: Calling Convention Example: Prologue, Epilogue



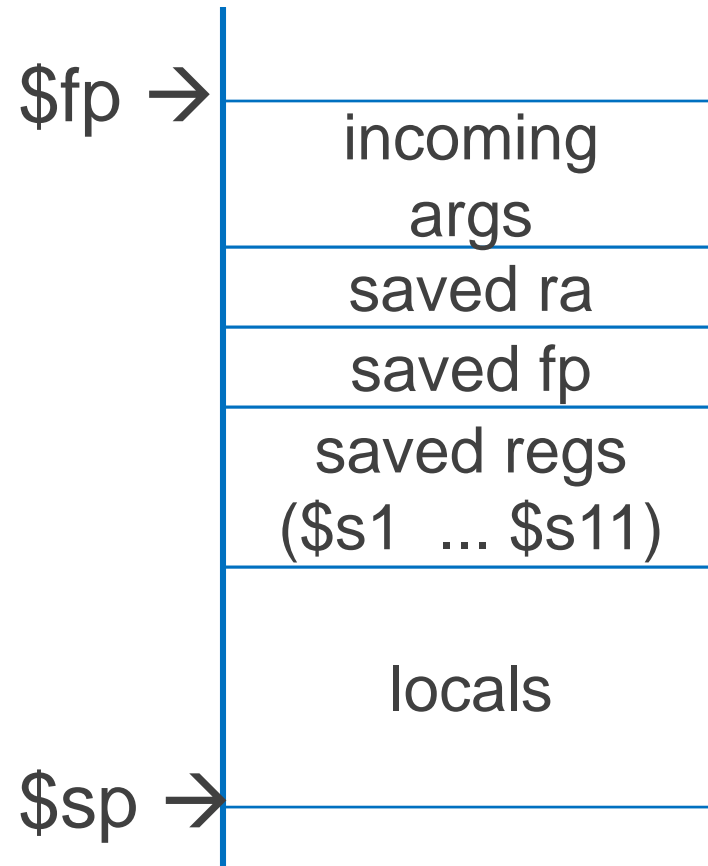
Next Goal

Can we optimize the assembly code at all?



Minimum stack size for a standard function?

Minimum stack size for a standard function?



Minimum stack size for a standard function?

Leaf function does not invoke any other functions

```
int f(int x, int y) { return (x+y); }
```

Optimizations?


No saved regs (or locals)

No incoming args

Don't push \$ra


No frame at all?

Maybe.

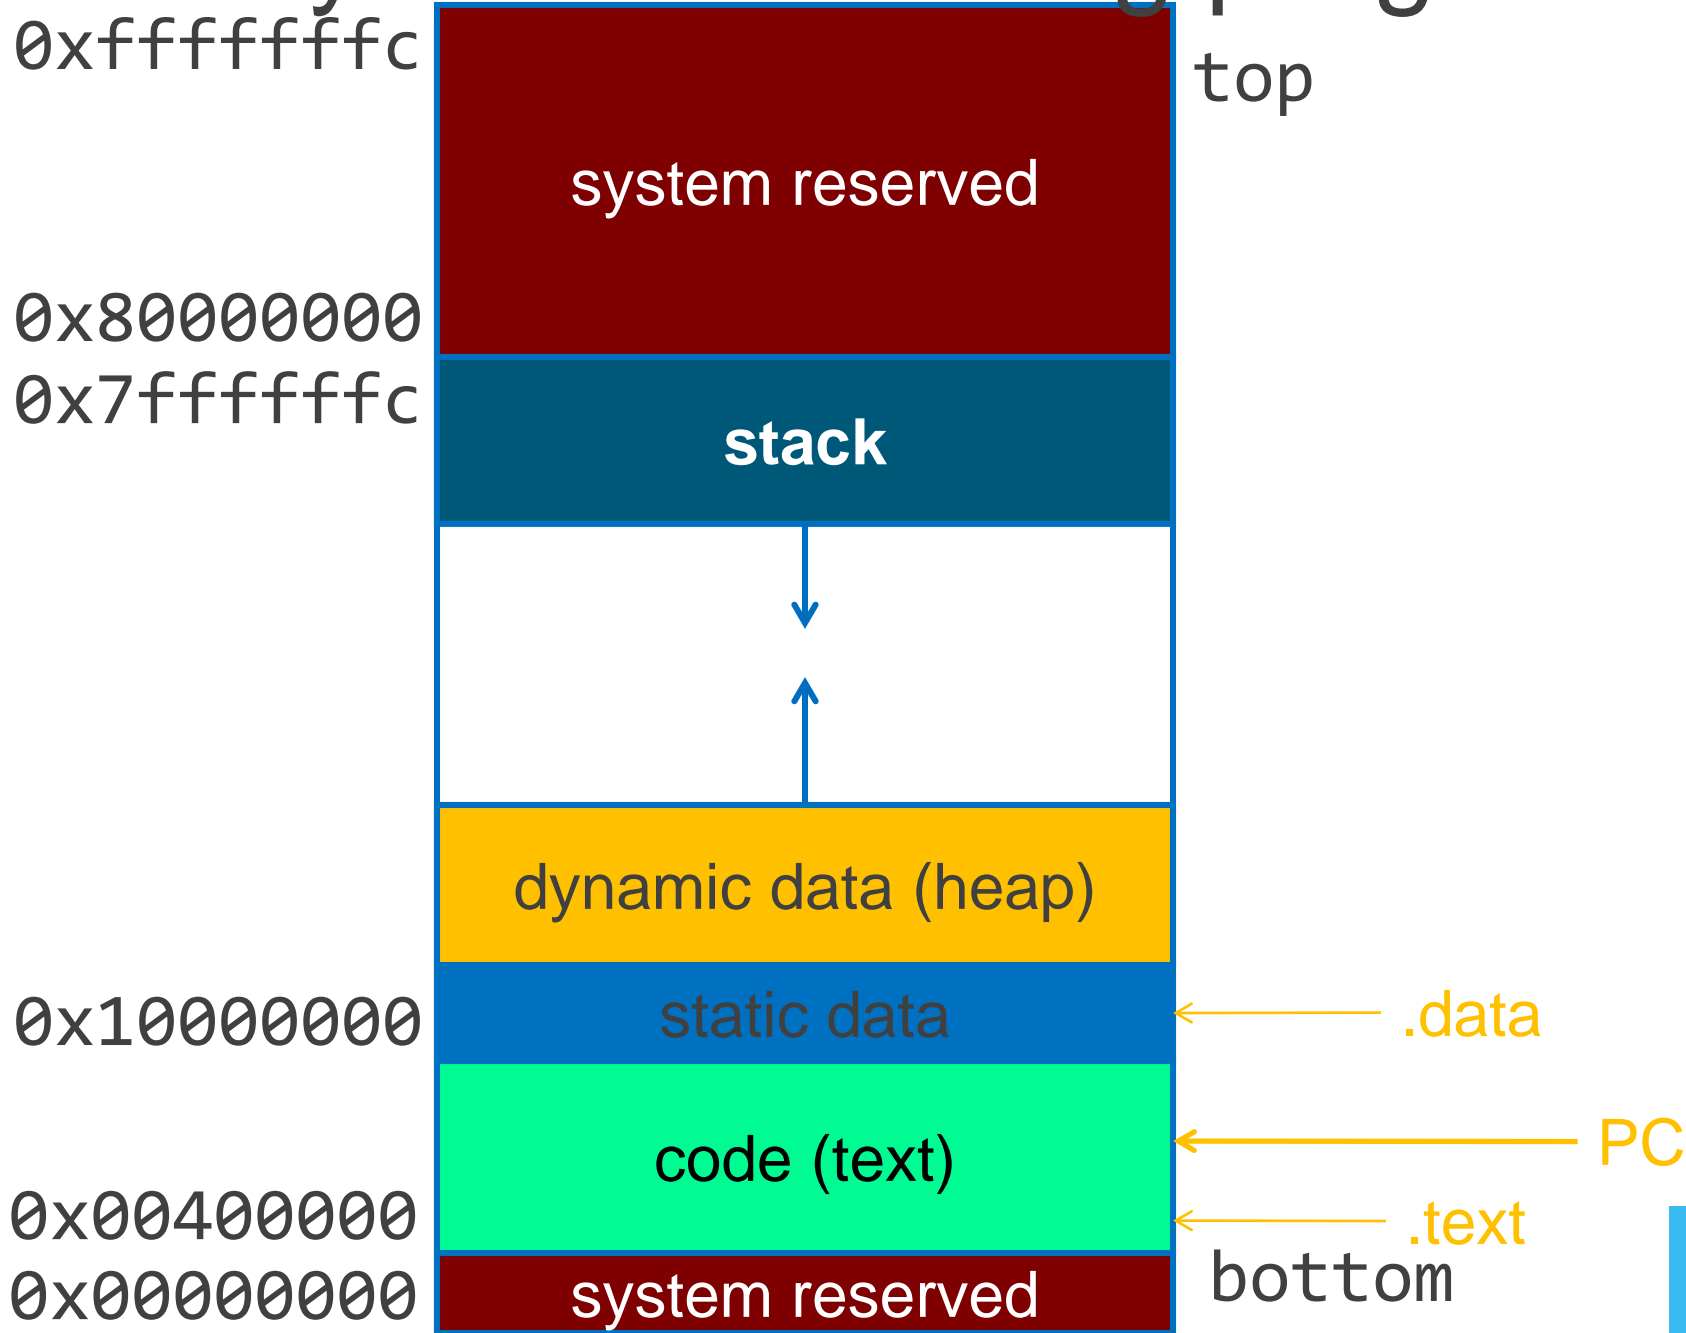


Next Goal

Given a running program (a process), how do we know what is going on (what function is executing, what arguments were passed to where, where is the stack and current stack frame, where is the code and data, etc)?



Anatomy of an executing program



Activity #4: Debugging

```
init():          0x400000  
printf(s, ...): 0x4002B4  
vnorm(a,b):     0x40107C  
main(a,b):      0x4010A0  
pi:             0x10000000  
str1:           0x10000004
```

```
CPU:  
$pc=0x004003C0  
$sp=0x7FFFFFFAC  
$ra=0x00401090
```

0x00000000
0x0040010c
0x7FFFFFF4
0x00000000
0x00000000
0x00000000
0x00000000
0x004010c4
0x7FFFFFFDC
0x00000000
0x00000000
0x00000015
0x10000004
0x00401090

0x7FFFFFFB0

What func is running?

Who called it?

Has it called anything?

Will it?

Args?

Stack depth?

Call trace?

Activity #4: Debugging

init(): 0x400000
printf(s, ...): 0x4002B4
vnorm(a,b): 0x40107C
main(a,b): 0x4010A0
pi: 0x10000000
str1: 0x10000004

The image shows a GDB memory dump with the following structure:

- CPU:** A box containing register values: `$pc=0x004003C0`, `$sp=0x7FFFFFFAC`, and `$ra=0x00401090`. The label `main` is positioned below this box.
- Stack:** A list of memory addresses and their contents, with labels on the right indicating stack frames:
 - `0x00000000` (a0)
 - `0x0040010c` (ra)
 - `0x7FFFFFFF4` (fp)
 - `0x00000000` (a3)
 - `0x00000000` (a2)
 - `0x00000000` (a1)
 - `0x00000000` (a0)
 - `0x004010c4` (ra)
 - `0x7FFFFFFDC` (fp)
 - `0x00000000` (a3)
 - `0x00000000` (a2)
 - `0x00000015` (a1)
 - `0x10000004` (a0)
 - `0x00401090` (ra)
 - `0x7FFFFFFC4` (fp)

Arrows indicate the call trace: `main` points to `$ra=0x00401090`, `vnorm` points to `0x004010c4`, and `printf` points to `0x00401090`.

What func is running? **printf**

Who called it? **vnorm**

Has it called anything? **no**

Will it? **no** b/c no space for outgoing args

Args? **Str1** and **0x15**

Stack depth? **4**

Call trace? **printf, vnorm, main, init**

Recap

- How to write and Debug a RISC-V program using calling convention
- **First eight** arg words passed in a0, a1, ..., a7
- Space for args passed **in child's stack frame**
- return value (if any) in a0, a1
- stack frame (fp/s0 to sp) contains:
 - ra (clobbered on JAL to sub-functions)
 - fp
 - **local vars** (possibly clobbered by sub-functions)
 - **Contains space for incoming args**
- **callee** save regs are **preserved**
- **caller** save regs are **not**
- Global data accessed via **gp**

