

CS 370: INTRODUCTION TO SECURITY
05.25: SOFTWARE SECURITY PRELIM.

Tu/Th 4:00 – 5:50 PM

Sanghyun Hong

sanghyun.hong@oregonstate.edu



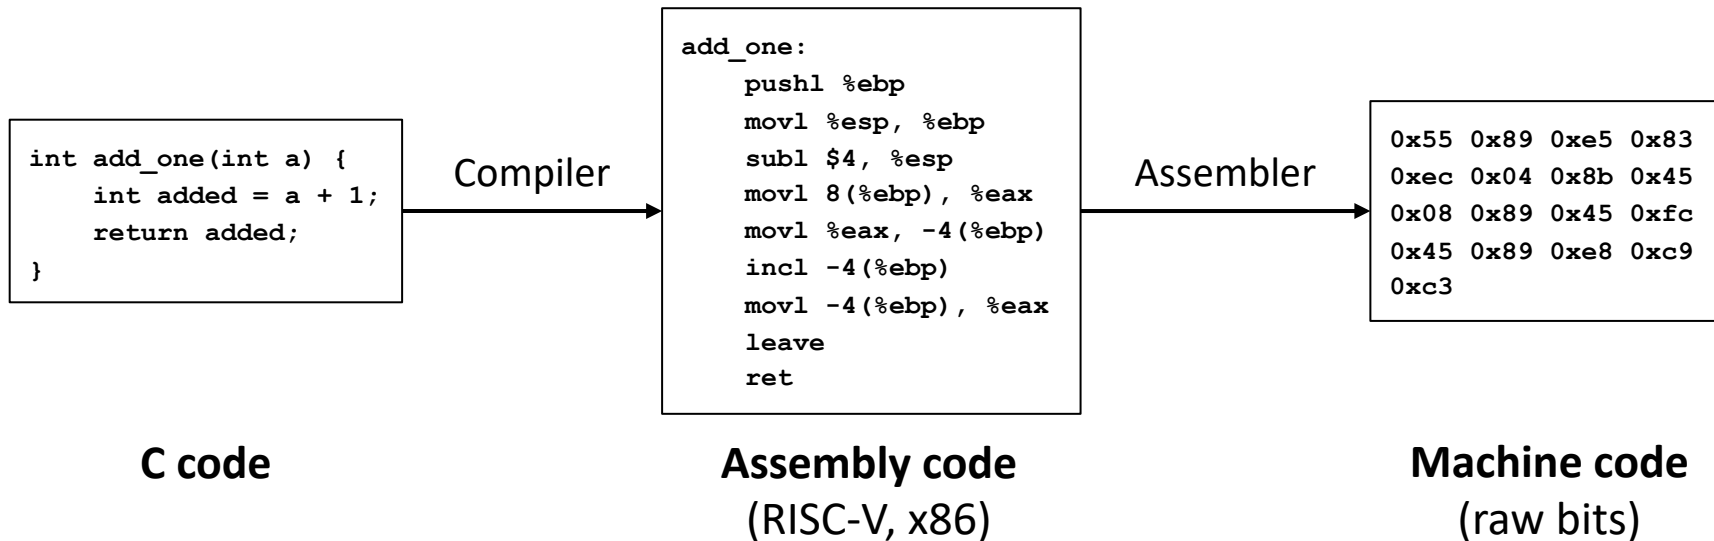
Oregon State
University

SAIL
Secure AI Systems Lab

TOPICS FOR TODAY

- Preliminaries (x86 assembly and call stack)
 - C program
 - Memory layout
 - x86 architecture
 - Stack layout
 - Calling convention
 - x86 calling convention design
 - x86 calling convention example

RUNNING A C PROGRAM: COMPILER AND ASSEMBLER



RUNNING A C PROGRAM: LINKER AND LOADER


- To run a C program:
 - Compiler : Converts C code into assembly code (RISC-V, x86)
 - Assembler : Converts assembly code into machine code (raw bits)
 - Linker : Deals with dependencies and libraries (learn more in CS444)
 - Loader : Sets up memory space and runs the machine code

TOPICS FOR TODAY

- Preliminaries (x86 assembly and call stack)
 - C program
 - Memory layout
 - x86 architecture
 - Stack layout
 - Calling convention
 - x86 calling convention design
 - x86 calling convention example


MEMORY LAYOUT

- C memory layout
 - At runtime, the loader tells an OS to give your program a big blob of memory
 - On a 32-bit system, the memory has 32-bit addresses
 - On a 64-bit system, the memory has 64-bit addresses
 - ex. the “solve” server is the 64-bit system
 - In this lecture slides, we consider a 32-bit system
 - Each address refers to 1 byte, which means you have 2^{32} bytes of memory



MEMORY LAYOUT

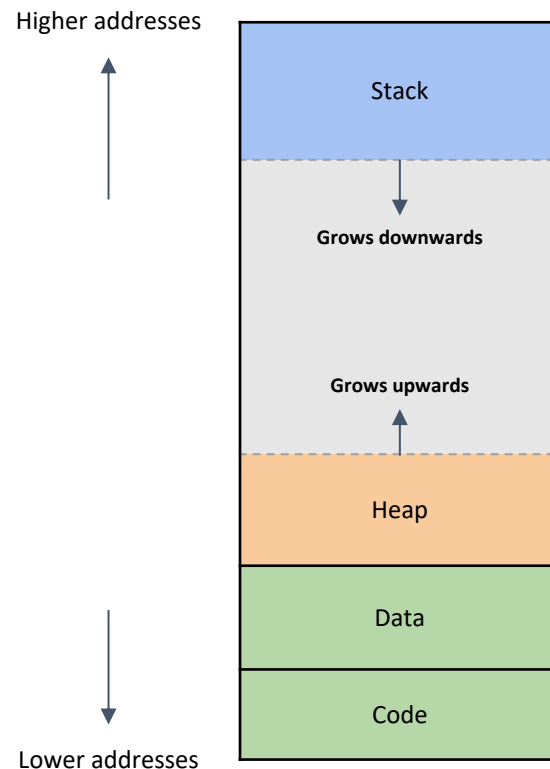
- C memory layout
 - Drawn vertically for ease of drawing
 - But memory is just a long array of bytes



MEMORY LAYOUT: X86

- Process has 4 segments

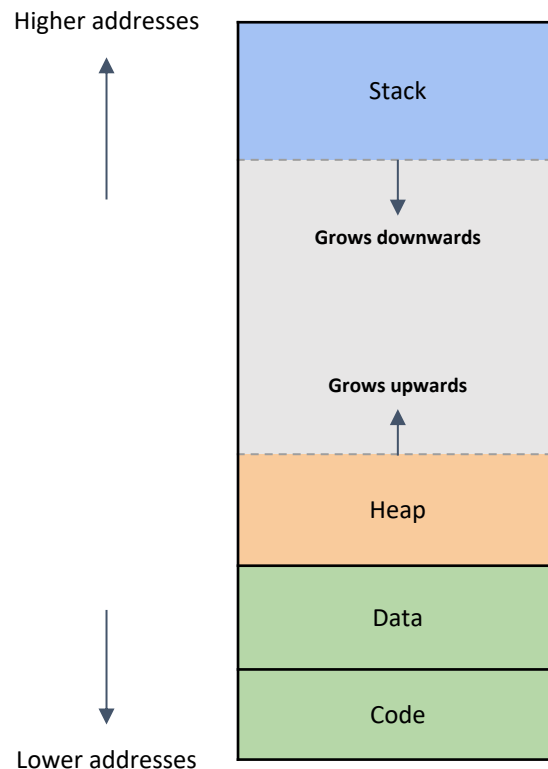
- Code (or text)
 - The program code itself
- Data
 - Static variables
 - Allocated when the program is started
- Heap
 - Dynamically allocated memory using *malloc* and *free*
 - Heap grows upwards
- Stack:
 - Local variables and stack frames
 - Stack grows downwards



MEMORY LAYOUT: X86

- Registers

- A quickly accessible location **on the CPU**
- Use names (ebp, esp, eip), not addresses
 - Memory: addresses are 32-bit numbers
- This is different from the memory layout

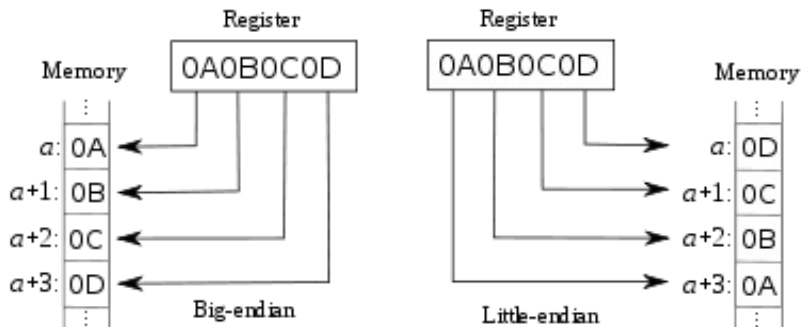


TOPICS FOR TODAY

- Preliminaries (x86 assembly and call stack)
 - C program
 - Memory layout
 - x86 architecture
 - Stack layout
 - Calling convention
 - x86 calling convention design
 - x86 calling convention example

x86 ARCHITECTURE: PRELIMINARIES

- x86 architecture
 - Most commonly used architecture
 - Use **little-endian**
 - The LSB is placed at the first/lowest memory address



- Support **variable-length instructions**
 - If assembled into machine code, instructions can be anywhere from 1 to 16 bytes long
 - Some other architectures could support fixed-length instructions (e.g., RISC-V; 4-byte)

X86 ARCHITECTURE: REGISTERS

- x86 registers
 - A quickly accessible location (separately) **on the CPU**
 - 8 main general-purpose registers:
 - EAX, EBX, ECX, EDX, ESI, EDI: General-purpose
 - ESP: Stack pointer
 - EBP: Base pointer
 - Instruction pointer register: EIP


X86 ARCHITECTURE: REGISTERS

- x86 registers
 - A quickly accessible location (separately) **on the CPU**
 - 8 main general-purpose registers:
 - EAX, EBX, ECX, EDX, ESI, EDI: General-purpose
 - ESP: Stack pointer
 - EBP: Base pointer
 - Instruction pointer register: EIP
- Syntax
 - Register references are preceded with a percent sign % (e.g., %eax, %esp, %edi)
 - Immediates are preceded with a dollar sign \$ (e.g., \$1, \$161, \$0x4)
 - Memory references use parentheses and can have immediate offsets
 - e.g., 8(%esp) dereferences memory 8 bytes above the address contained in ESP

X86 ARCHITECTURE: ASSEMBLY

- x86 assembly

- Instructions are composed of an opcode and zero or more operands.



- Pseudocode: **EBX = EBX + 0x8**

- The destination comes last

- The **add** instruction has two operands; and the destination is an input


- This instruction uses a register and an immediate

X86 ARCHITECTURE: ASSEMBLY

- x86 assembly

- Instructions are composed of an opcode and zero or more operands.

- `xorl 4(%esi) %eax`



- Pseudocode: $EAX = EAX \oplus *(ESI + 4)$

- This is a memory reference:

- The value at 4 bytes above the address in ESI is dereferenced
 - XOR'd with EAX
 - Stored back into EAX

TOPICS FOR TODAY

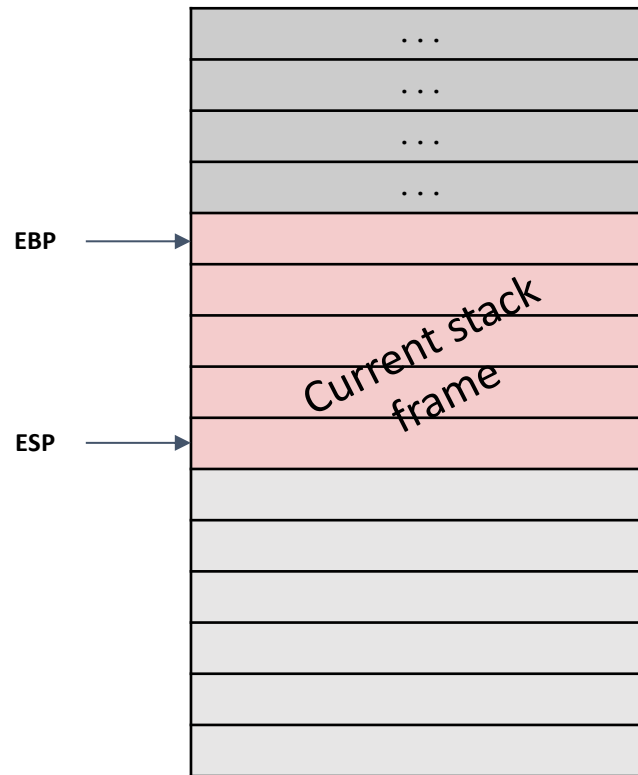
- Preliminaries (x86 assembly and call stack)
 - C program
 - Memory layout
 - x86 architecture
 - Stack layout
 - Calling convention
 - x86 calling convention design
 - x86 calling convention example

STACK LAYOUT

- Stack frames
 - If code calls a function:
 - Memory space is made on the stack for local variables
 - The space is known as the stack frame for the function
 - The stack frame will be free-ed once the function returns
 - The stack makes extra space by growing down
 - The stack starts at higher addresses
 - Every time your code calls a function, it grows down
 - Note:
 - Data on the stack, e.g., a string, is still stored from lowest address to highest address.
 - “Growing down” only happens when extra memory needs to be allocated.

STACK LAYOUT

- Stack frames
 - To keep track of the current stack frame
 - Store two pointers in registers
 - The EBP (base pointer) points to the top of the current stack frame
 - The ESP (stack pointer) points to the bottom of the current stack frame



STACK LAYOUT

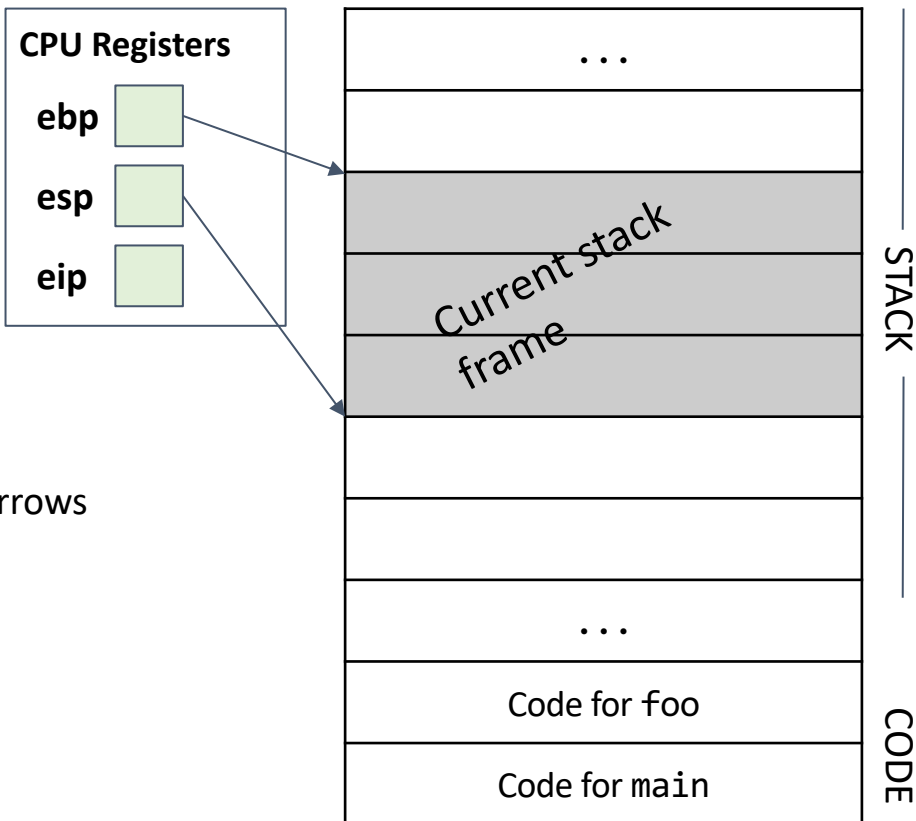
- Stack frames

- To keep track of the current stack frame

- Store two pointers in registers
 - The EBP (base pointer) points to the top of the current stack frame
 - The ESP (stack pointer) points to the bottom of the current stack frame

- Store

- The **ebp** and **esp** registers are drawn as arrows



STACK LAYOUT

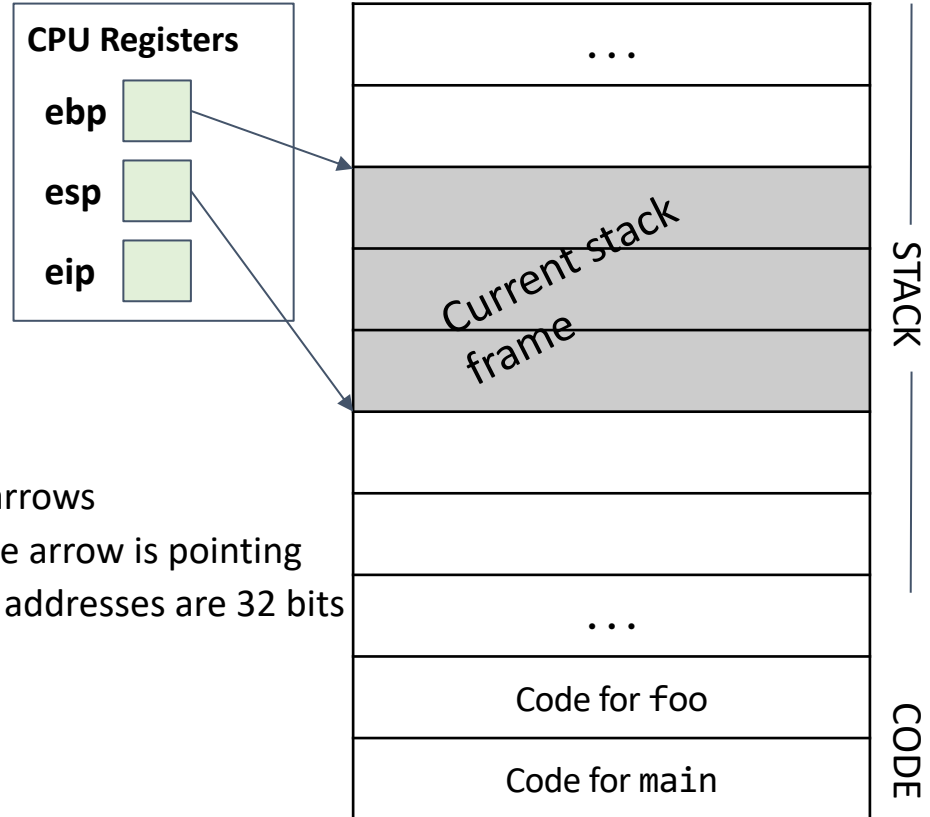
- Stack frames

- To keep track of the current stack frame

- Store two pointers in registers
 - The EBP (base pointer) points to the top of the current stack frame
 - The ESP (stack pointer) points to the bottom of the current stack frame

- Store (pointers)

- The **ebp** and **esp** registers are drawn as arrows
 - They are storing the address of where the arrow is pointing
 - This works as registers store 32 bits, and addresses are 32 bits



STACK LAYOUT

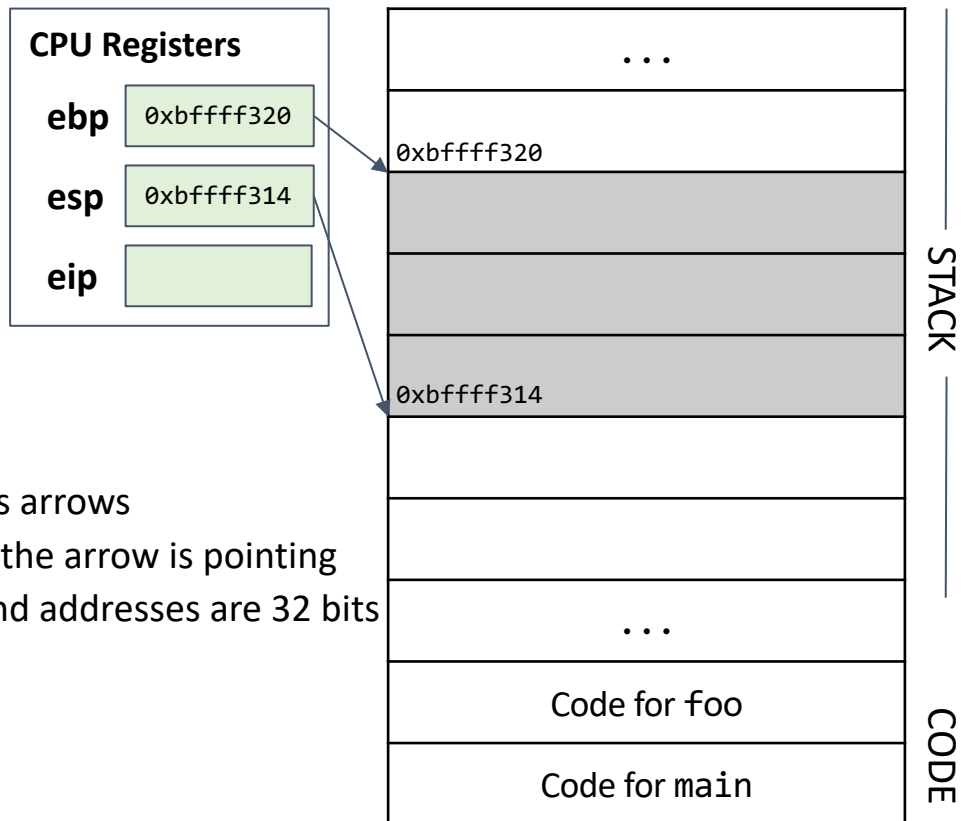
- Stack frames

- To keep track of the current stack frame

- Store two pointers in registers
 - The EBP (base pointer) points to the top of the current stack frame
 - The ESP (stack pointer) points to the bottom of the current stack frame

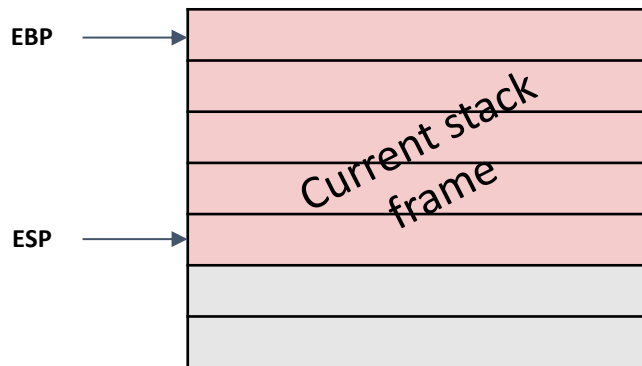
- Store (pointers)

- The **ebp** and **esp** registers are drawn as arrows
 - They are storing the address of where the arrow is pointing
 - This works as registers store 32 bits, and addresses are 32 bits



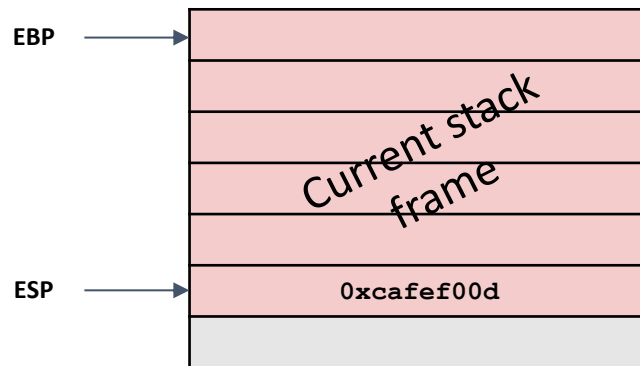
STACK LAYOUT

- Push and pop
 - The **push** instruction adds an element to the stack
 - Decrement ESP to allocate more memory on the stack
 - Save the new value on the lowest value address of the stack



EAX = 0xcafef00d
EBX = ...

Before **push %eax**




EAX = 0xcafef00d
EBX = ...

After **push %eax**


STACK LAYOUT

- Push and pop
 - The **pop** instruction removes an element from the stack
 - Load the value from the lowest value address on the stack and store it in a register
 - Increment ESP to deallocate the memory on the stack



EAX = `0x00000000`
EBX = ...

Before `pop %eax`



EAX = `0xcafef00d`
EBX = ...

After `pop %eax`

STACK LAYOUT

- Storing convention
 - Local variables are always allocated on the stack
 - Individual variables within a stack frame are stored with the first variable at the highest address
 - Members of a struct are stored with the first member at the lowest address
 - Global variables (not on the stack) are stored with the first variable at the lowest address

STACK LAYOUT

- Storing convention

- Local variables are always allocated on the stack
- Individual variables within a stack frame are stored with the first variable at the highest address
- Members of a struct are stored with the first member at the lowest address
- Global variables (not on the stack) are stored with the first variable at the lowest address


```
struct foo {  
    long long f1; // 8 bytes  
    int f2;      // 4 bytes  
    int f3;      // 4 bytes  
};
```

```
void func(void) {  
    int a;      // 4 bytes  
    struct foo b;  
    int c;      // 4 bytes  
}
```

Higher addresses



Lower addresses

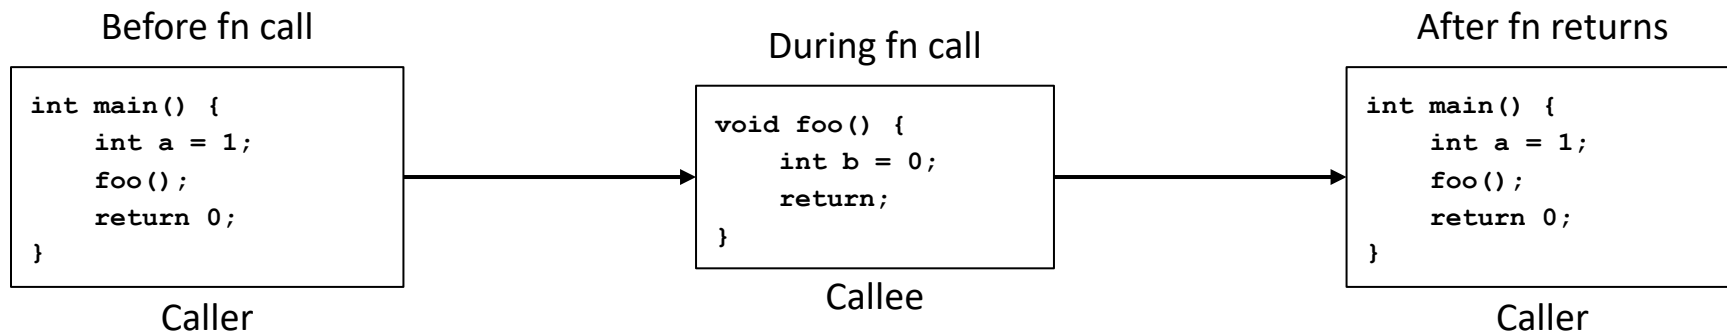


← 4 bytes →

TOPICS FOR TODAY

- Preliminaries (x86 assembly and call stack)
 - C program
 - Memory layout
 - x86 architecture
 - Stack layout
 - Calling convention
 - x86 calling convention design
 - x86 calling convention example

CALLING CONVENTION: FUNCTION CALLS



The **caller** function (**main**)
calls the **callee** function (**foo**)


The callee function executes and then
returns control to the caller function

CALLING CONVENTION

- x86 convention
 - A way for functions to call other functions (i.e., know what state the processor will return in)
 - How to pass **arguments**
 - Arguments are pushed onto the stack in reverse order
 - **func(val1, val2, val3)** will place **val3** at the highest memory address, then **val2**, then **val1**
 - How to receive **return values**
 - Return values are passed in **EAX**
 - Which registers are **caller-saved** or **callee-saved**
 - **Callee-saved**: The callee must not change the value of the register when it returns
 - **Caller-saved** : The callee may overwrite the register without saving or restoring it


CALLING CONVENTION

- x86 convention
 - The values in the caller-saved registers to stay unchanged when calling a function (i.e., If the function returns, the value in these registers should stay the same)
 - What if the function wants to change the values in these registers?
 - Before calling the function: write these values on the stack
 - After the function returns: move the values from the stack back to the registers




CALLING CONVENTION

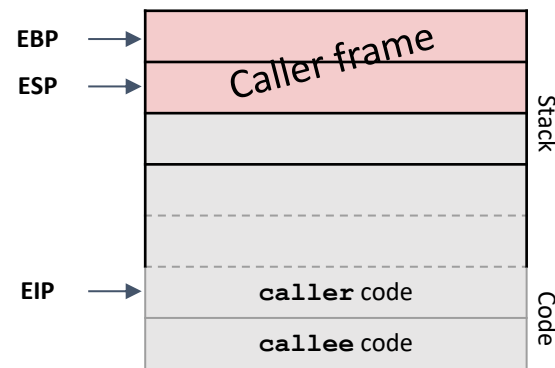
- Calling a function in x86
 - Call:
 - The ESP and EBP need to shift to create a new stack frame
 - The EIP must move to the callee's code
 - Return:
 - The ESP, EBP, and EIP must return to their old values



Before fn call



During fn call





After fn call

X86 FUNCTION CALL DESIGN



- Stack and registers

- If code calls a function, space is made on the stack for local variables
- The space goes away once the function returns
- The stack starts at higher addresses and grows down
- Registers are 32-bit (or 4-byte, 1-word) units of memory located on CPU



X86 FUNCTION CALL DESIGN



- Word and code segment
 - The **code segment** contains raw bytes that represent assembly instructions
 - Each row of the diagram is **1 word** = 4 bytes = 32 bits
 - Addresses increase as you move up the diagram



X86 FUNCTION CALL DESIGN

- Stack frames

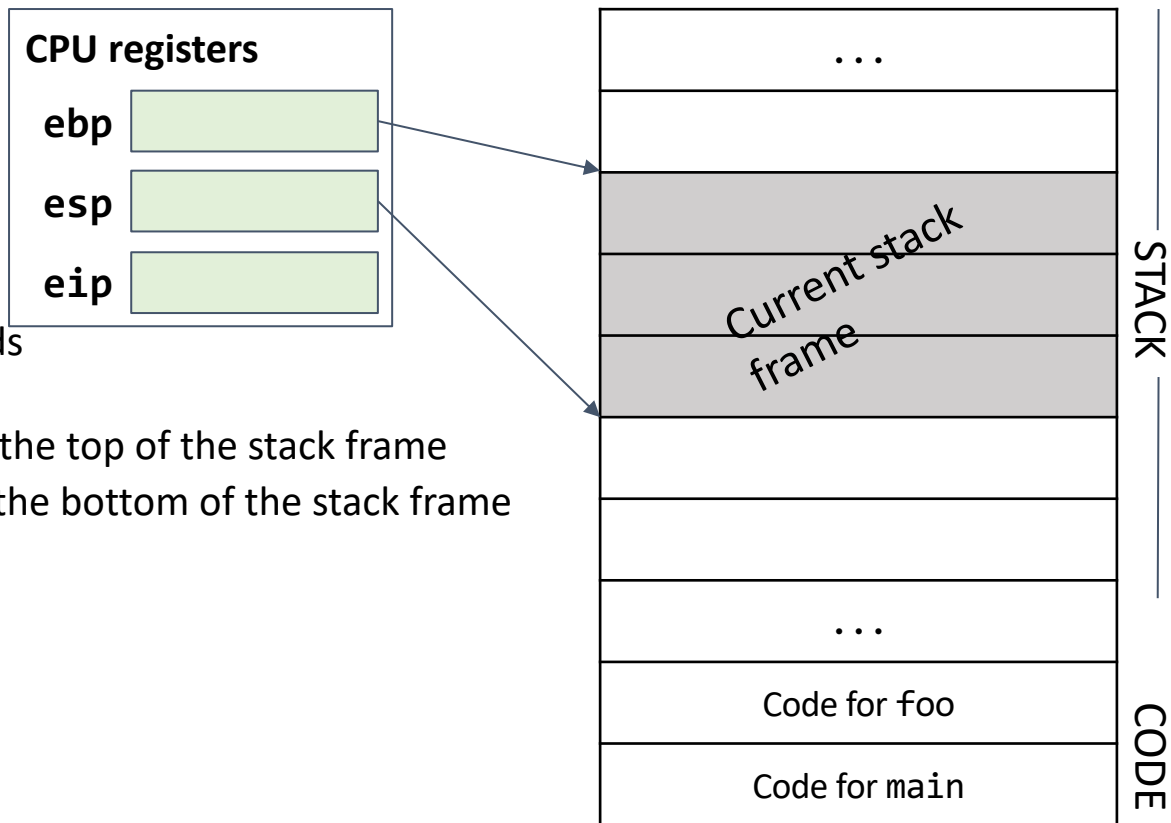
- Use two pointers to tell us which part of the stack is being used by the current function
- This is called a **stack frame**
- One stack frame corresponds to one function being called



X86 FUNCTION CALL DESIGN

- Stack frames

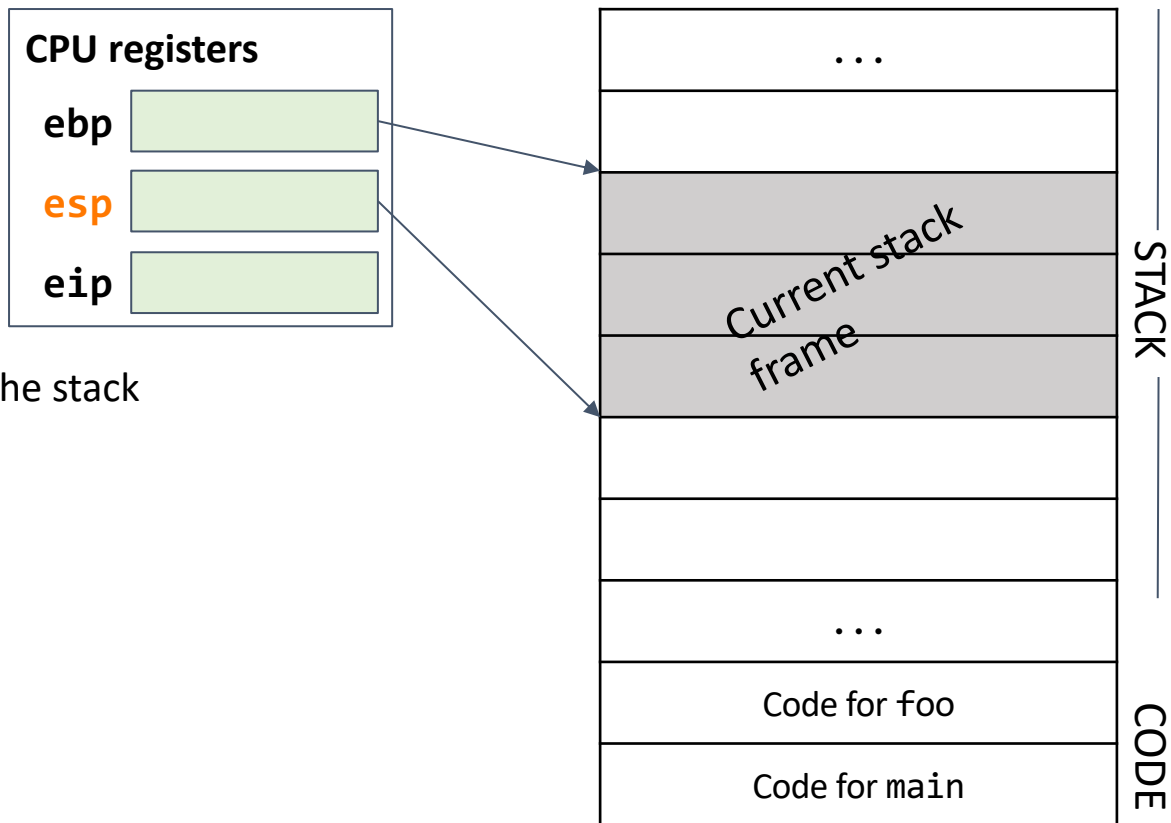
- Use **two pointers** to tell us which part of the stack is being used by the current function
- This is called a **stack frame**
- One stack frame corresponds to one function being called
- The **ebp** register is used for the top of the stack frame
- The **esp** register is used for the bottom of the stack frame



X86 FUNCTION CALL DESIGN

- ESP

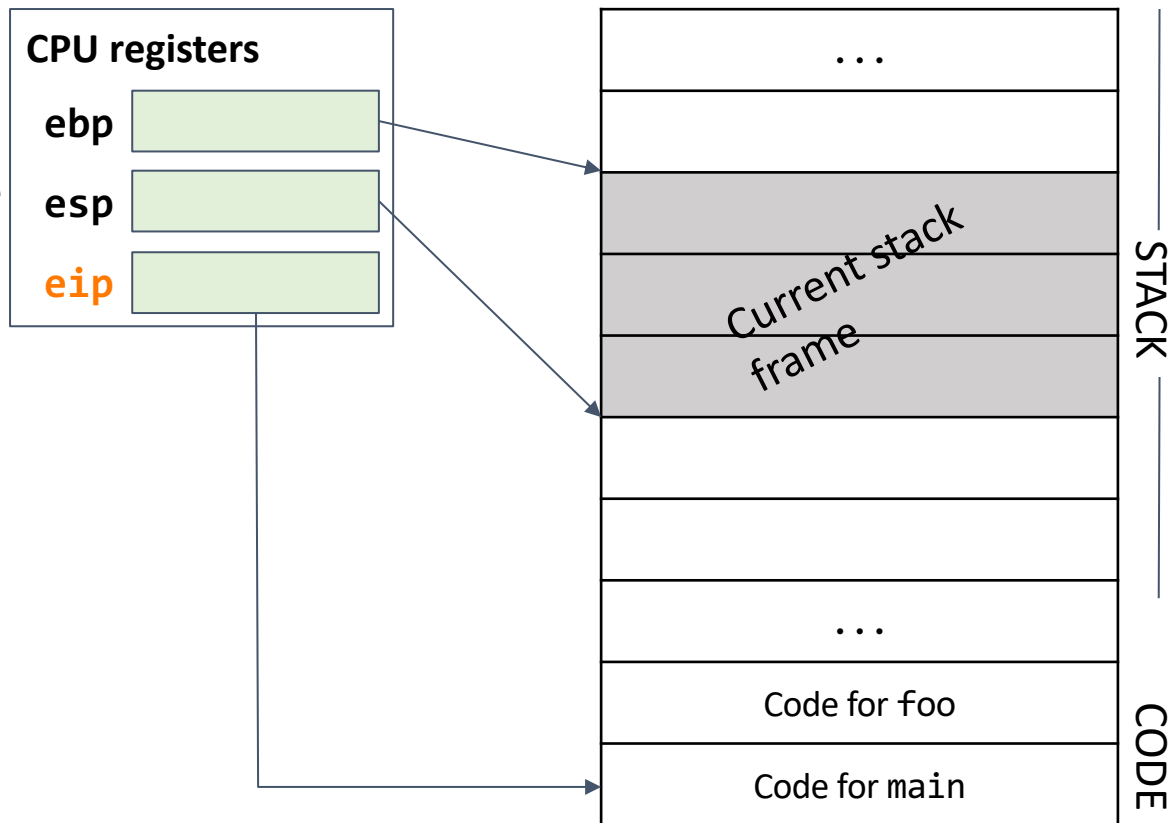
- **esp** also denotes the current lowest value on the stack
- Everything below **esp** is **undefined**
- If we push a value onto the stack, **esp** must adjust to match the lowest value on the stack



X86 FUNCTION CALL DESIGN

- EIP

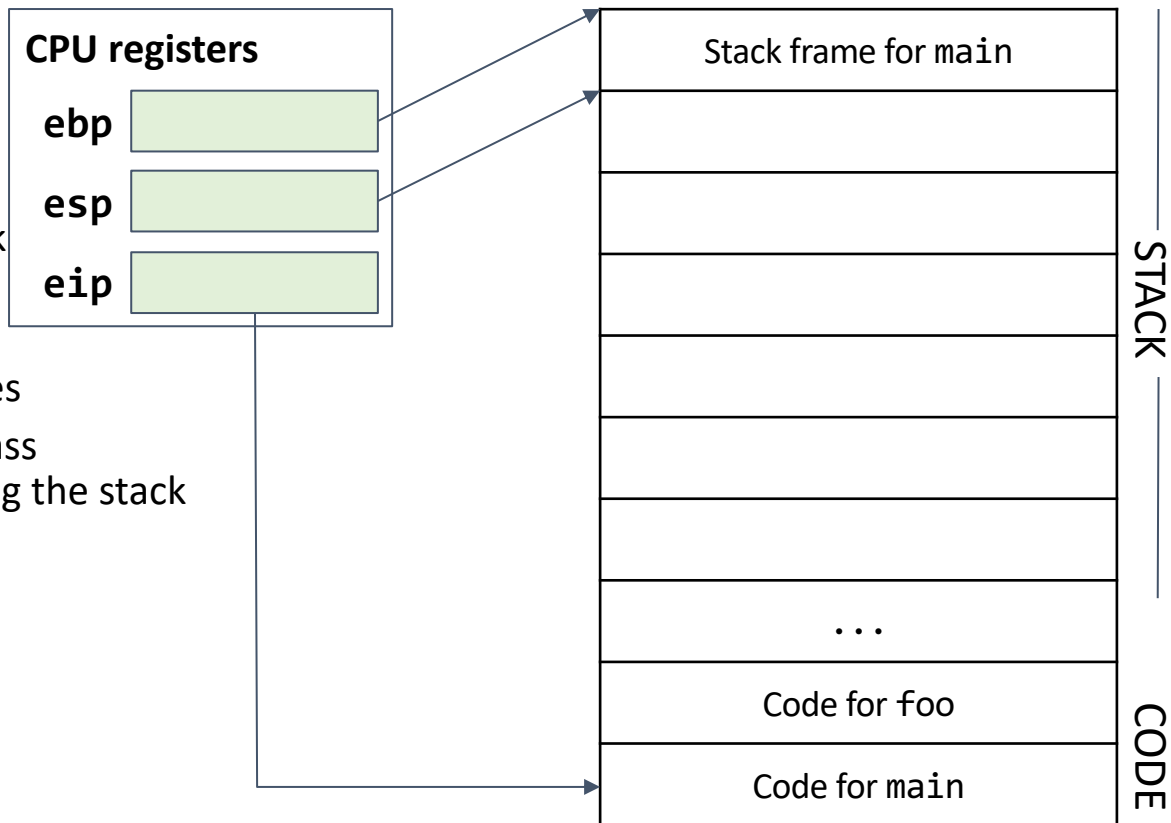
- To keep track of what step we're at in the instructions
- Use the **eip** register to store a pointer to the current instruction



X86 FUNCTION CALL DESIGN

- Stack design

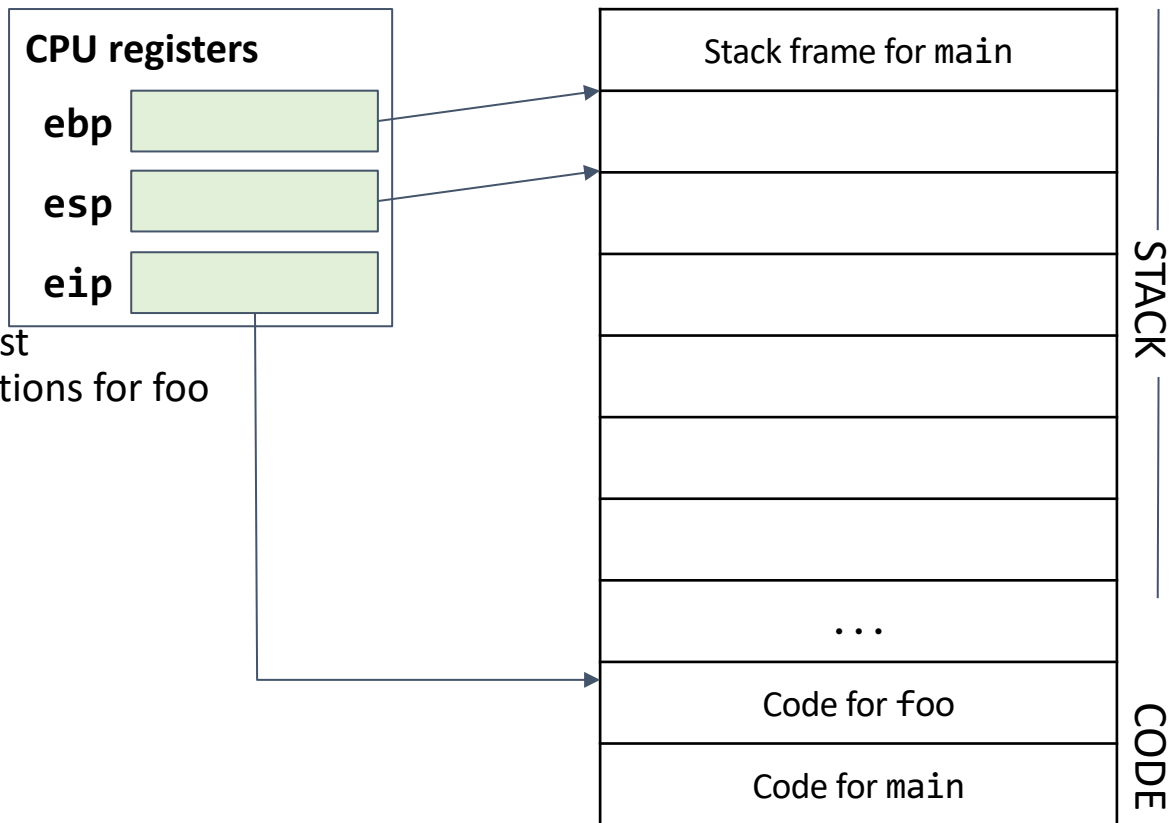
- Every time we call a func., a new stack frame must be created
- If the func returns, the stack frame must be discarded
- Each stack frame needs to have space for local variables
- Require to design how to pass arguments to functions using the stack



X86 FUNCTION CALL DESIGN

- Stack design

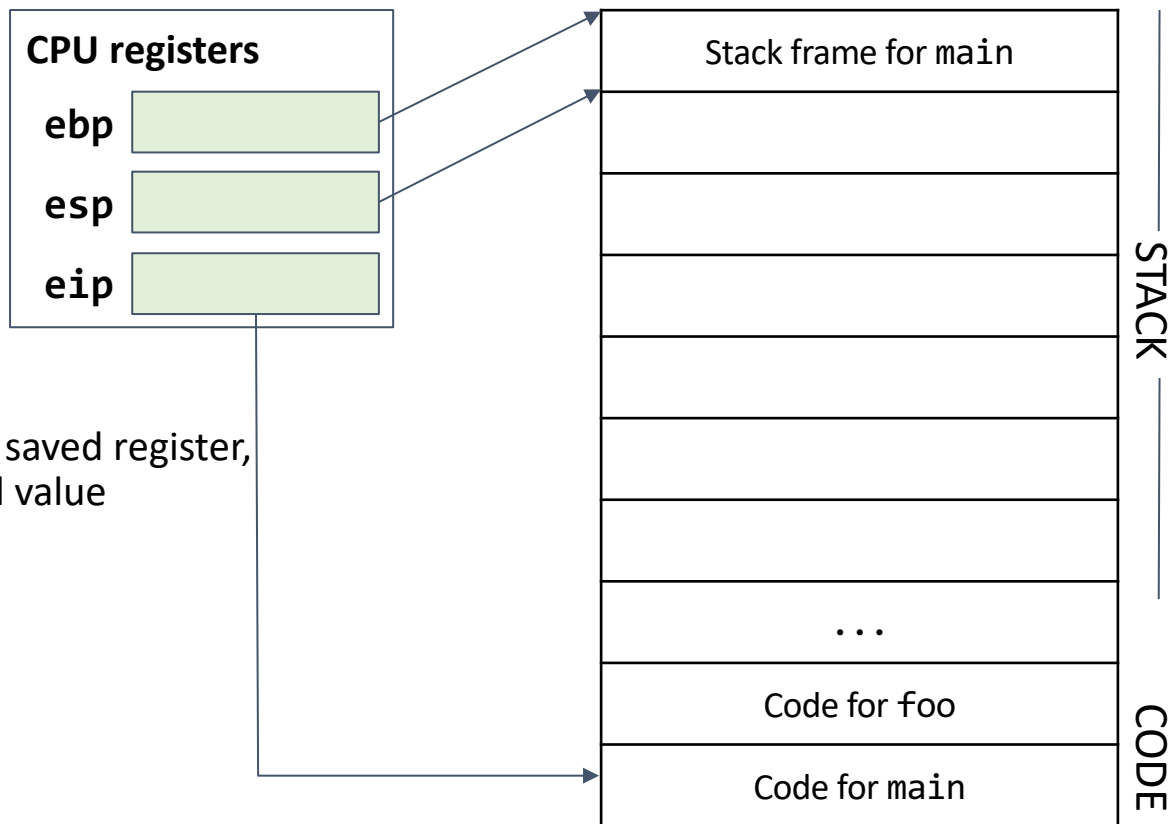
- Example: **foo** called
- The **ebp** and **esp** registers should adjust to give us a stack frame for **foo** with the correct size
- The **eip** register should adjust to let us execute the instructions for **foo**



X86 FUNCTION CALL DESIGN

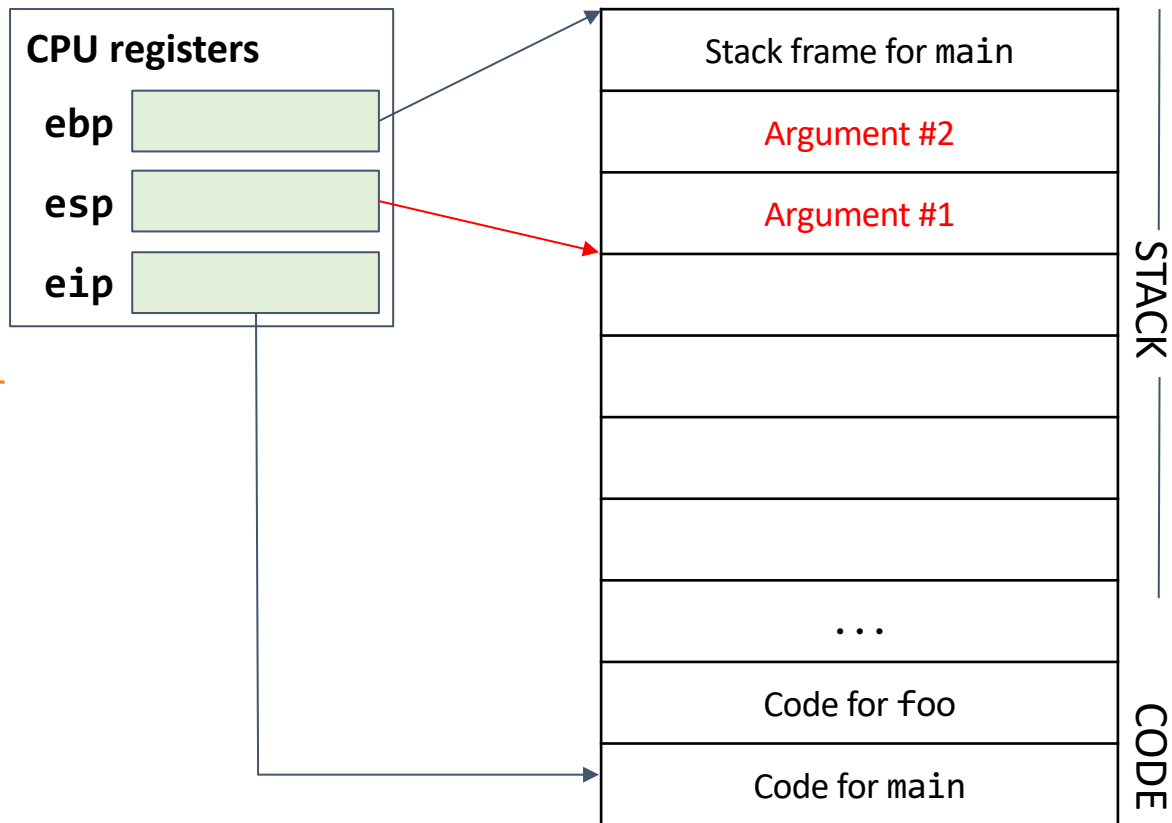
- Stack design

- Example: **foo** returns
- The stack should look exactly like it did before **foo** was called
- Require to design **how to pass arguments** to functions using the stack
- Rule: if we ever overwrite a saved register, we should remember its old value by putting it on the stack



X86 FUNCTION CALL DESIGN

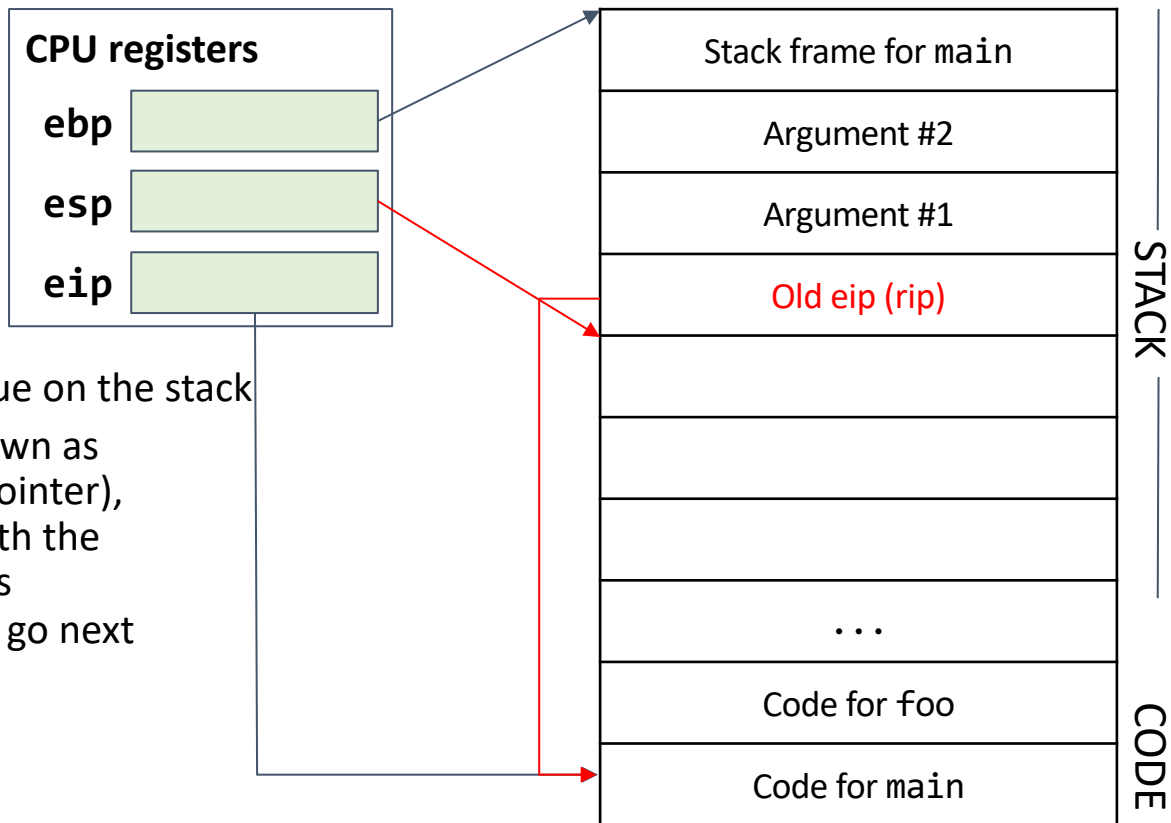
- Store arguments
 - Push the arguments onto the stack
 - Remember to adjust **esp** to point to the new lowest value on the stack
 - Arguments are added to the stack **in reverse order**



x86 FUNCTION CALL DESIGN

- Remember **eip**

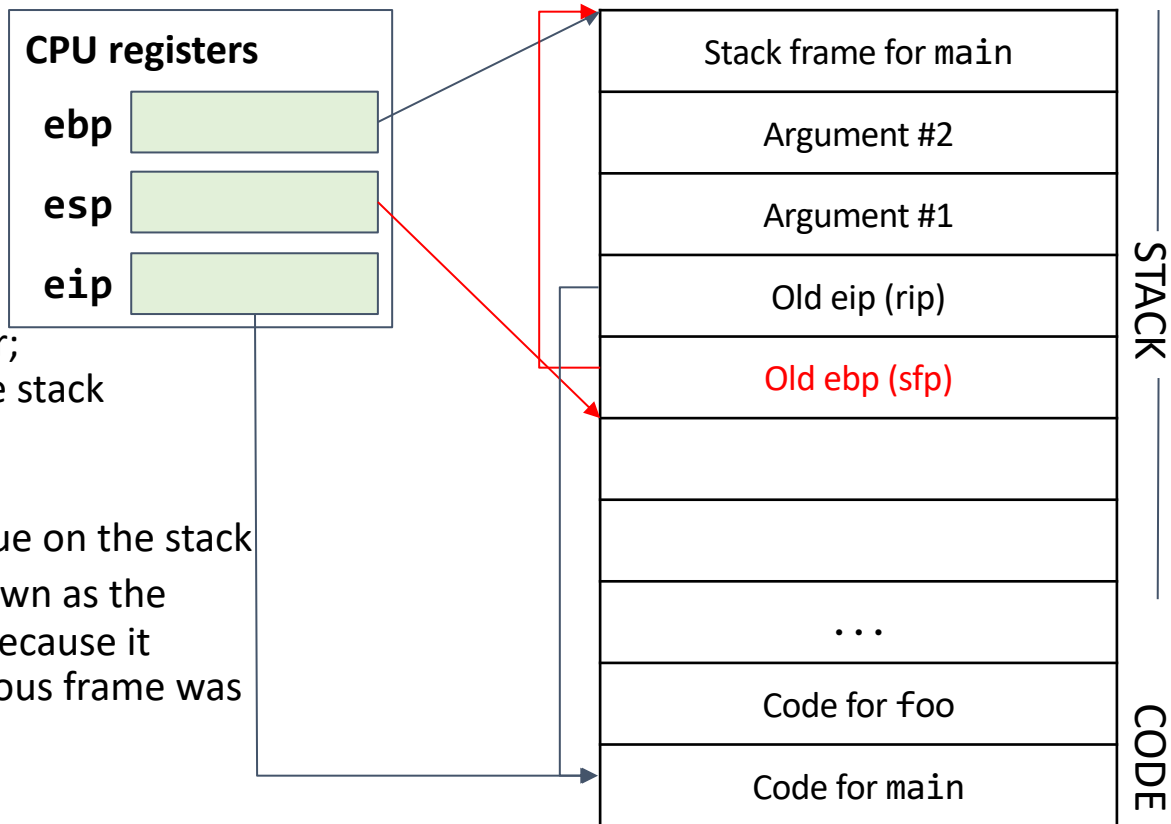
- Push the current value of **eip** on the stack
- This tells us what code to execute next after the function returns
- Remember to adjust **esp** to point to the new lowest value on the stack
- This value is sometimes known as the **rip** (return instruction pointer), because if we're finished with the function, this pointer tells us where in the instructions to go next



X86 FUNCTION CALL DESIGN

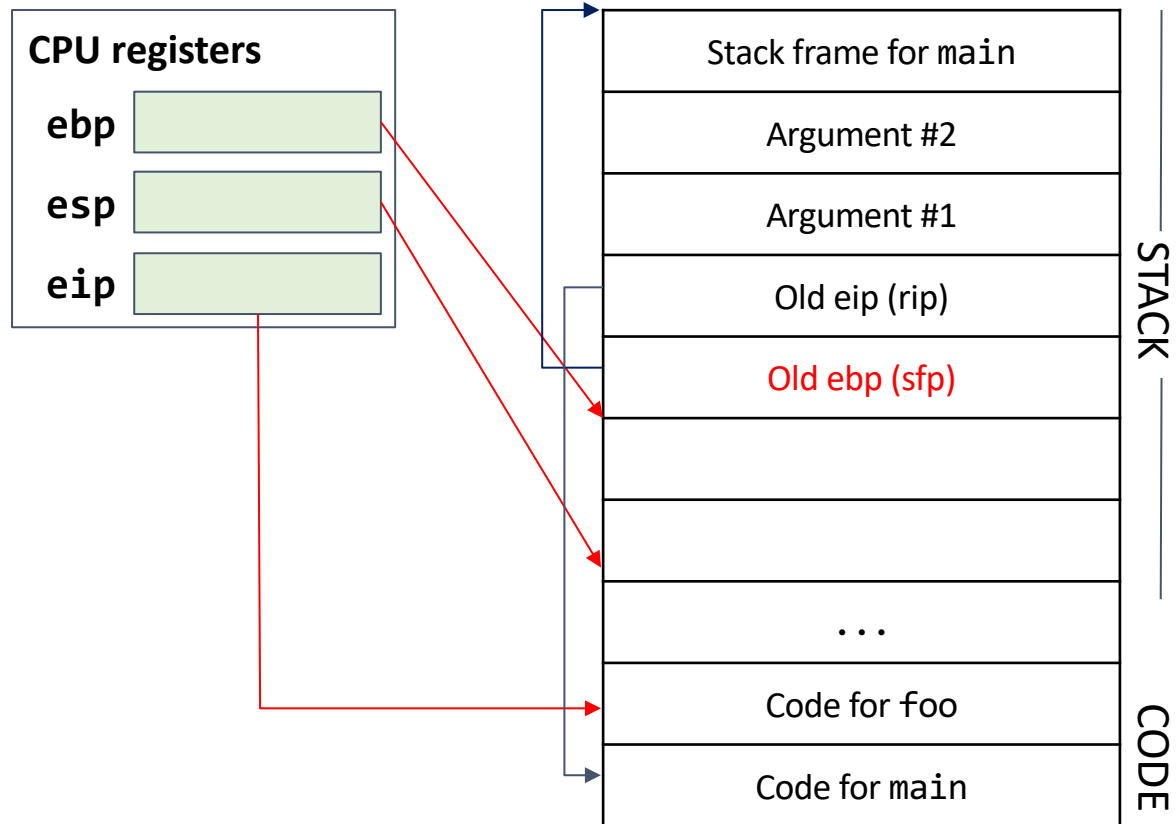
- Remember **ebp**

- Push the current value of **ebp** on the stack.
- This will let us restore the top of the previous stack frame when we return
- Note: **ebp** is a saved register; we store its old value on the stack before overwriting it
- Remember to adjust **esp** to point to the new lowest value on the stack
- This value is sometimes known as the **sfp** (saved frame pointer), because it reminds us where the previous frame was



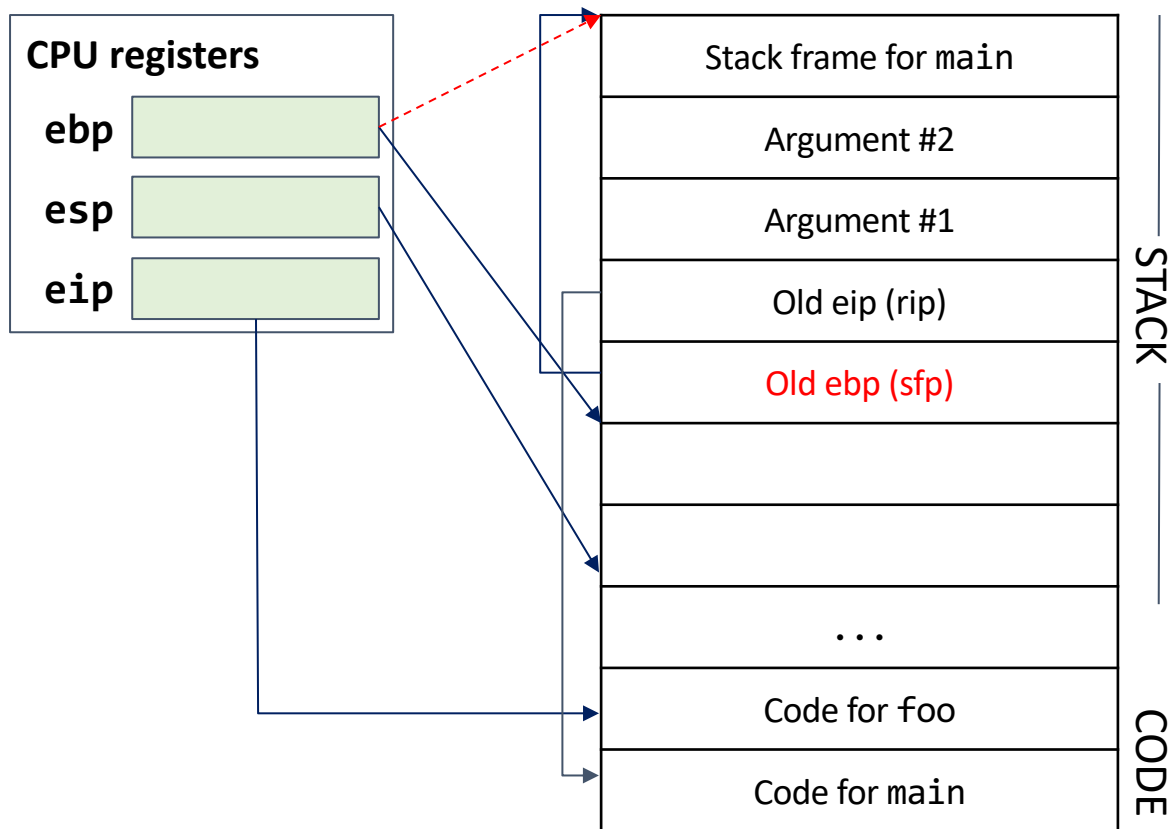
x86 FUNCTION CALL DESIGN

- Adjust the stack frame
 - Update all 3 registers
 - We can safely do this as we've just saved the old values of **ebp** and **eip**
 - Note: **esp** will always be the bottom of the stack, so there's no need to save it



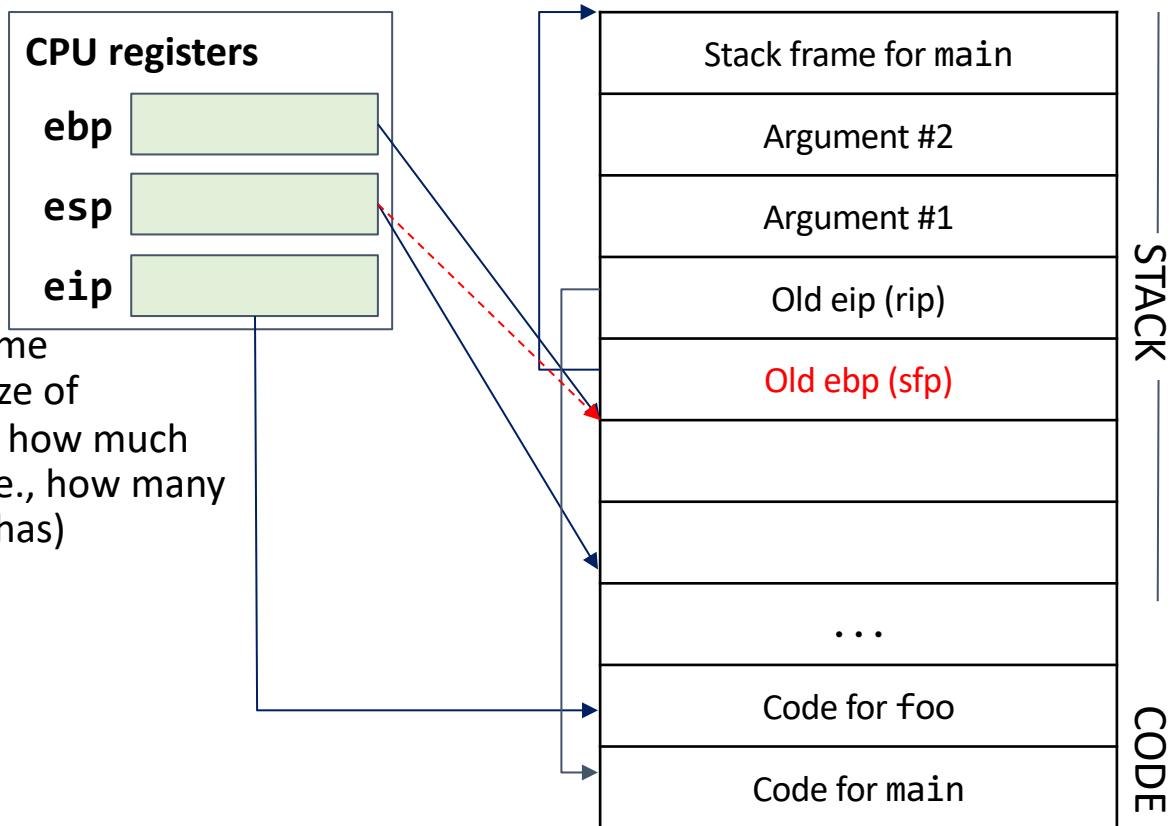
X86 FUNCTION CALL DESIGN

- Adjust the stack frame
 - Update all 3 registers
 - **ebp** now points to the top of the current stack frame, which is always the **sfp**



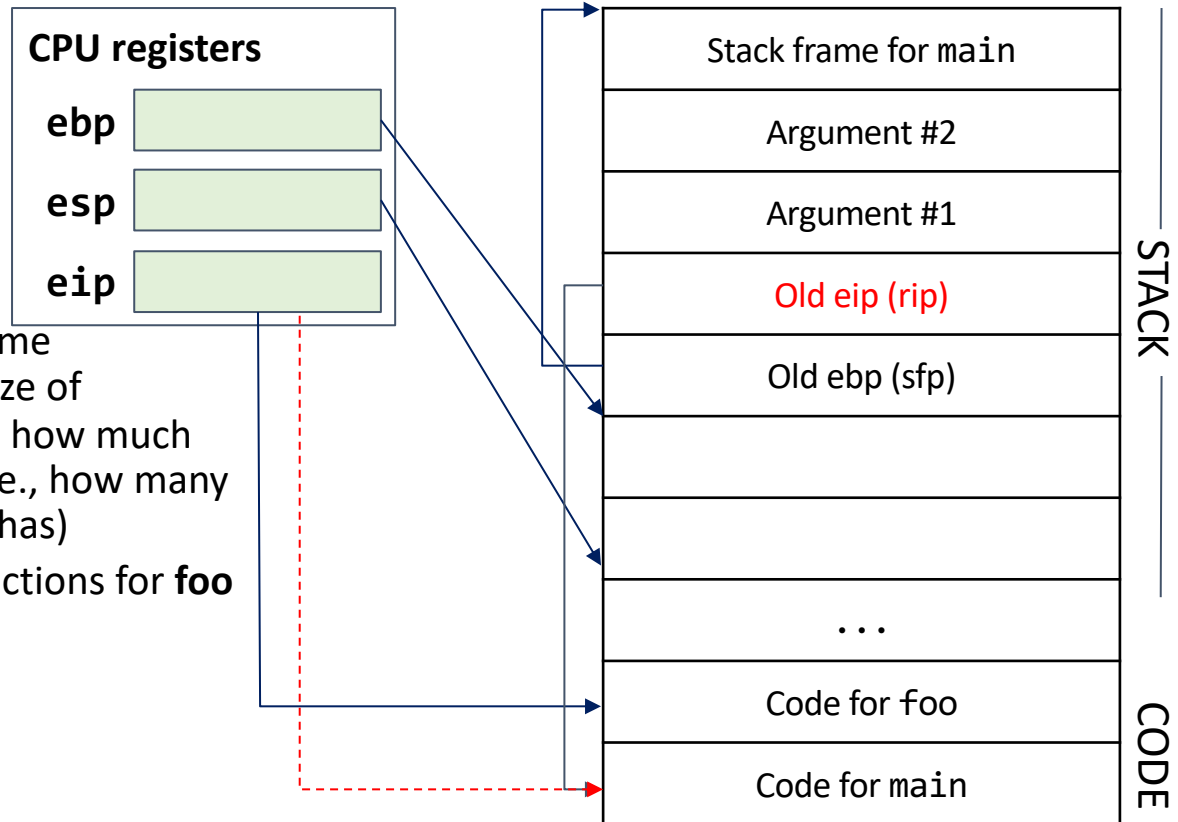
X86 FUNCTION CALL DESIGN

- Adjust the stack frame
 - Update all 3 registers
 - **ebp** now points to the top of the current stack frame, which is always the **sfp**
 - **esp** now points to the bottom of the current stack frame (the compiler decides the size of the stack frame by checking how much space the function needs, i.e., how many local variables the function has)



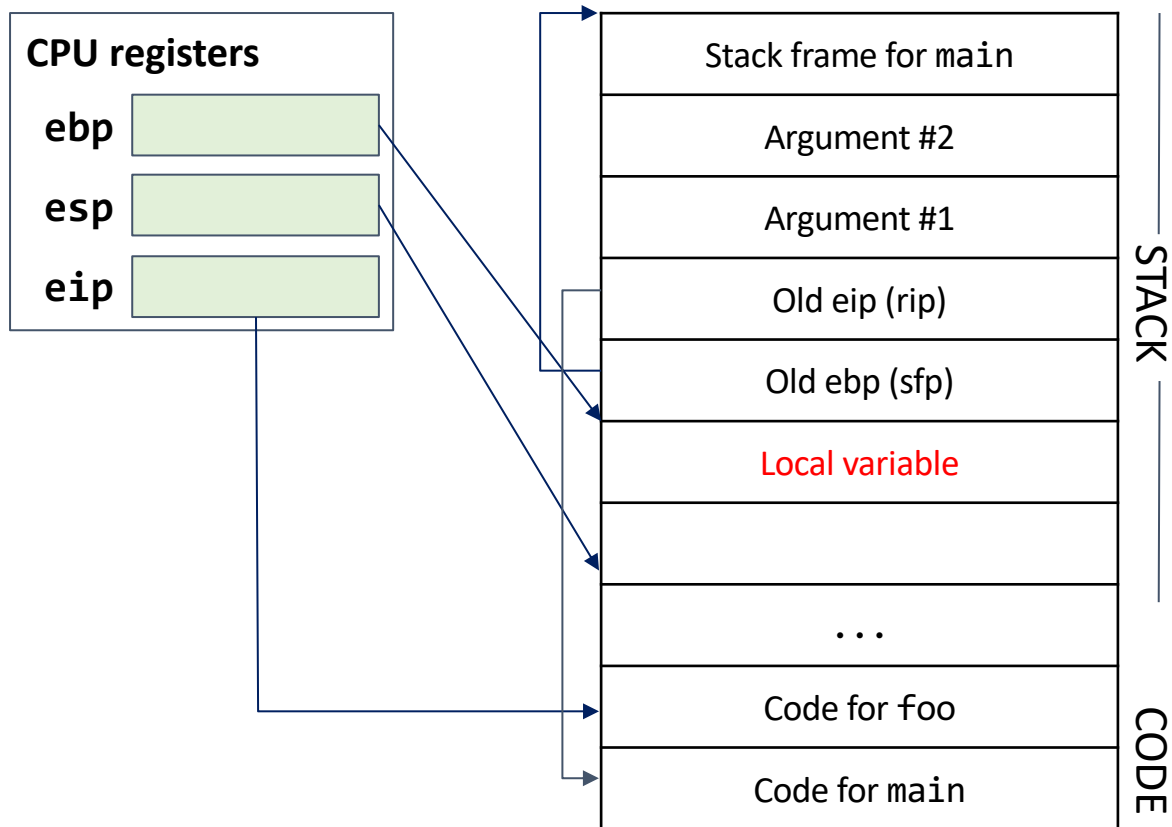
X86 FUNCTION CALL DESIGN

- Adjust the stack frame
 - Update all 3 registers
 - **ebp** now points to the top of the current stack frame, which is always the **sfp**
 - **esp** now points to the bottom of the current stack frame (the compiler decides the size of the stack frame by checking how much space the function needs, i.e., how many local variables the function has)
 - **eip** now points to the instructions for **foo**



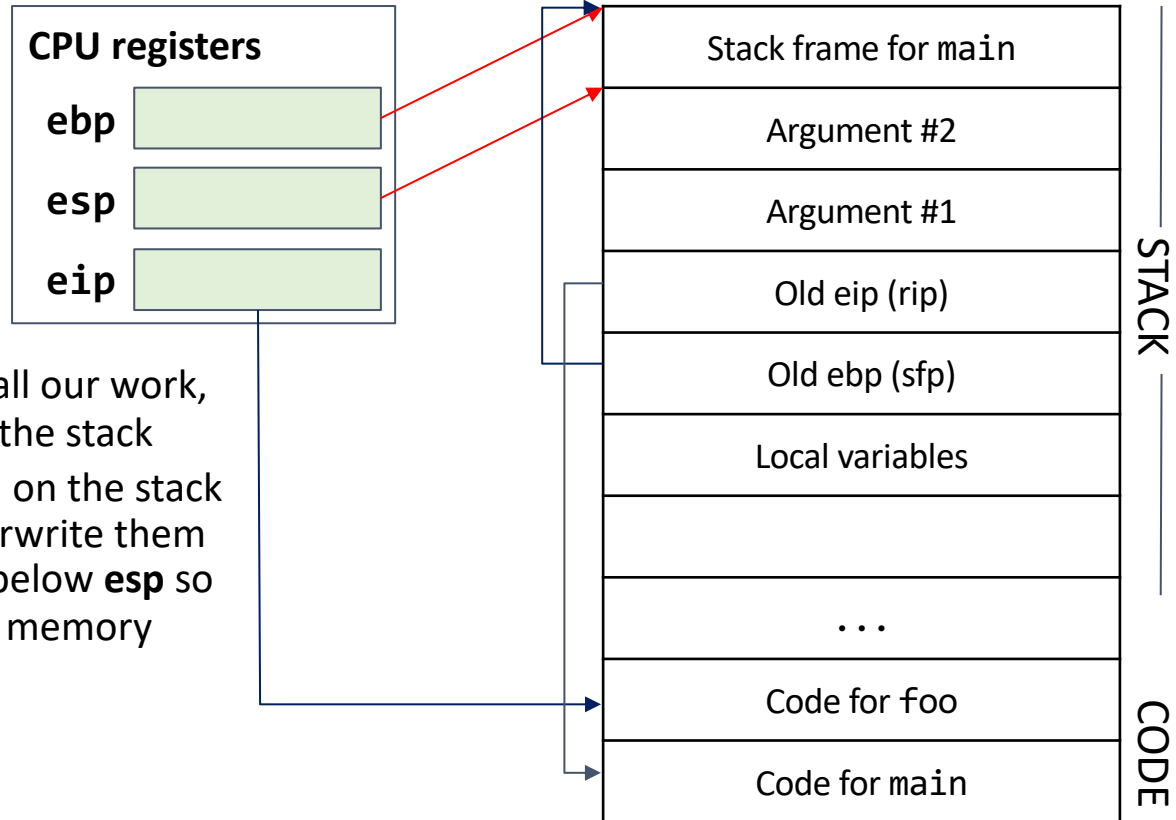
X86 FUNCTION CALL DESIGN

- Run the function
 - Now the stack frame is ready to do whatever the function instructions are
 - Any local variables will be stored to the stack now



X86 FUNCTION CALL DESIGN

- Return from the function
 - Put all 3 registers back where they were before
 - Use the addresses stored in **rip** and **sfp** to restore **eip** and **ebp** to their old values
 - **esp** naturally moves back to its old place as we undo all our work, which is popping values off the stack
 - Note: the values we pushed on the stack are still there (we don't overwrite them to save time), but they are below **esp** so they cannot be accessed by memory



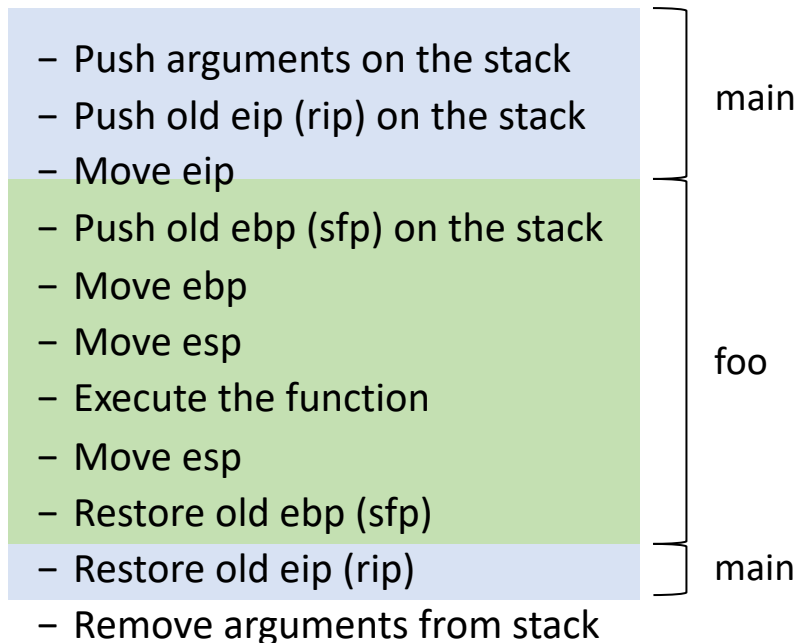
X86 FUNCTION CALL DESIGN

- Steps of a function call

- Push arguments on the stack
- Push old eip (rip) on the stack

- Push old ebp (sfp) on the stack
- Adjust the stack frame

- Execute the function
- Restore everything



TOPICS FOR TODAY

- Preliminaries (x86 assembly and call stack)
 - C program
 - Memory layout
 - x86 architecture
 - Stack layout
 - Calling convention
 - x86 calling convention design
 - x86 calling convention example

X86 FUNCTION CALL

```
void caller(void) {  
    callee(1, 2);  
}
```

```
int callee(int a, int b) {  
    int local;  
    return 42;  
}
```

- Illustration

- The code above snippets are the C functions
- On the right, the code compiled into x86 assembly

caller:

```
...  
push $2  
push $1  
call callee  
add $8, %esp  
...
```

callee:

```
push %ebp  
mov %esp, %ebp  
sub $4, %esp  
  
mov $42, %eax  
  
mov %ebp, %esp  
pop %ebp  
ret
```

X86 FUNCTION CALL

```
void caller(void) {  
    callee(1, 2);  
}
```

```
int callee(int a, int b) {  
    int local;  
    return 42;  
}
```

- Illustration

- The code above snippets are the C functions
- On the right, the code compiled into x86 assembly
- The instruction just executed in red
- The **EIP** points to the address of the next instruction

```
caller:  
    ...  
EIP → push $2  
       push $1  
       call callee  
       add $8, %esp  
       ...  
  
callee:  
    push %ebp  
    mov %esp, %ebp  
    sub $4, %esp  
  
    mov $42, %eax  
  
    mov %ebp, %esp  
    pop %ebp  
    ret
```

X86 FUNCTION CALL

```
void caller(void) {  
    callee(1, 2);  
}
```

```
int callee(int a, int b) {  
    int local;  
    return 42;  
}
```

- Illustration

- The code above snippets are the C functions
- On the right, the code compiled into x86 assembly
- The instruction just executed in red
- The **EIP** points to the address of the **next instruction**
- The below is the diagram of the stack (each row represents a word, 4-byte)

```
caller:  
    ...  
EIP → push $2  
       push $1  
       call callee  
       add $8, %esp  
       ...
```

```
callee:  
    push %ebp  
    mov %esp, %ebp  
    sub $4, %esp  
  
    mov $42, %eax  
  
    mov %ebp, %esp  
    pop %ebp  
    ret
```



X86 FUNCTION CALL


```
void caller(void) {  
    callee(1, 2);  
}
```

```
int callee(int a, int b) {  
    int local;  
    return 42;  
}
```

• Illustration

- Push the arguments to the stack
 - The **push** instruction decrements the ESP to make space on the stack
 - The arguments are pushed in reverse order

```
caller:  
    ...  
    push $2  
EIP → push $1  
       call callee  
       add $8, %esp  
    ...
```



```
callee:  
    push %ebp  
    mov %esp, %ebp  
    sub $4, %esp  
  
    mov $42, %eax  
  
    mov %ebp, %esp  
    pop %ebp  
    ret
```

X86 FUNCTION CALL

```
void caller(void) {  
    callee(1, 2);  
}
```

```
int callee(int a, int b) {  
    int local;  
    return 42;  
}
```

• Illustration

- Push the arguments to the stack
 - The **push** instruction decrements the ESP to make space on the stack
 - The arguments are pushed in reverse order

caller:

...

push \$2

push \$1

EIP → call callee

add \$8, %esp

...

callee:

push %ebp

mov %esp, %ebp

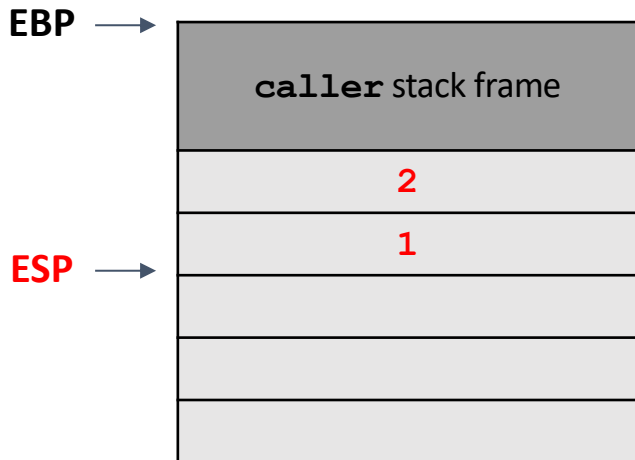
sub \$4, %esp

mov \$42, %eax

mov %ebp, %esp

pop %ebp

ret



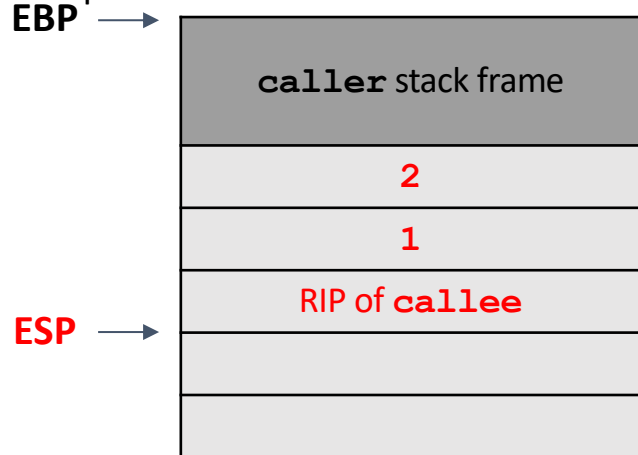
X86 FUNCTION CALL

```
void caller(void) {  
    callee(1, 2);  
}
```

```
int callee(int a, int b) {  
    int local;  
    return 42;  
}
```

• Illustration

- Push old EIP (RIP) on the stack
- Move EIP
 - The call instruction does 2 things
 - It first pushes the current value of EIP on the stack
 - The saved EIP value on the stack is called the RIP
 - It also changes EIP to point to the instructions of the callee



```
caller:  
    ...  
    push $2  
    push $1  
    call callee  
    add $8, %esp  
    ...
```

```
callee:  
    push %ebp  
    mov %esp, %ebp  
    sub $4, %esp  
  
    mov $42, %eax  
  
    mov %ebp, %esp  
    pop %ebp  
    ret
```


X86 FUNCTION CALL


```
void caller(void) {  
    callee(1, 2);  
}
```

```
int callee(int a, int b) {  
    int local;  
    return 42;  
}
```

• Illustration

- The next 3 steps set up a stack frame for the callee function
- These instructions are sometimes called the **function prologue** because they appear at the start of every function

```
caller:  
    ...  
    push $2  
    push $1  
    call callee  
    add $8, %esp  
    ...
```



```
callee:  
    push %ebp  
    mov %esp, %ebp  
    sub $4, %esp  
    mov $42, %eax  
    mov %ebp, %esp  
    pop %ebp  
    ret
```

Function prologue

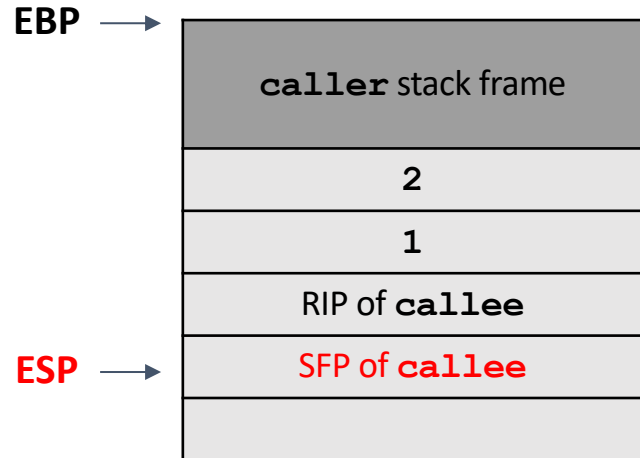
X86 FUNCTION CALL

```
void caller(void) {  
    callee(1, 2);  
}
```

```
int callee(int a, int b) {  
    int local;  
    return 42;  
}
```

• Illustration

- Push old EBP (SFP) on the stack
 - Restore the value of the EBP when returning, so we push the current value of the EBP on the stack
 - The saved value of the EBP on the stack is called the SFP



```
caller:  
    ...  
    push $2  
    push $1  
    call callee  
    add $8, %esp  
    ...
```

```
callee:  
    push %ebp  
    mov %esp, %ebp  
    sub $4, %esp  
  
    mov $42, %eax  
  
    mov %ebp, %esp  
    pop %ebp  
    ret
```

X86 FUNCTION CALL

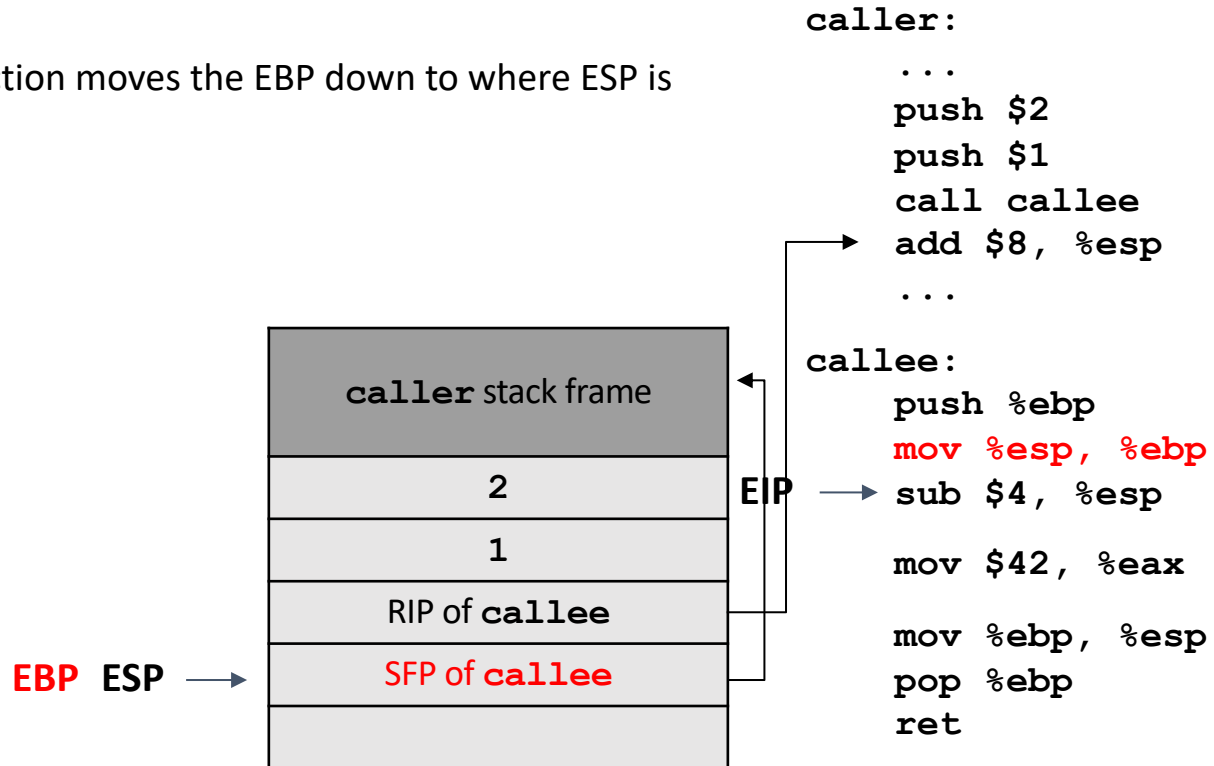
```
void caller(void) {  
    callee(1, 2);  
}
```

```
int callee(int a, int b) {  
    int local;  
    return 42;  
}
```

• Illustration

– Move EBP

- The instruction moves the EBP down to where ESP is



X86 FUNCTION CALL

```
void caller(void) {  
    callee(1, 2);  
}
```

```
int callee(int a, int b) {  
    int local;  
    return 42;  
}
```

• Illustration

– Move ESP


- The instruction moves the ESP down to create a new stack frame

caller:

```
...  
push $2  
push $1  
call callee  
add $8, %esp  
...
```

callee:

```
push %ebp  
mov %esp, %ebp  
sub $4, %esp  
mov $42, %eax  
mov %ebp, %esp  
pop %ebp  
ret
```



X86 FUNCTION CALL

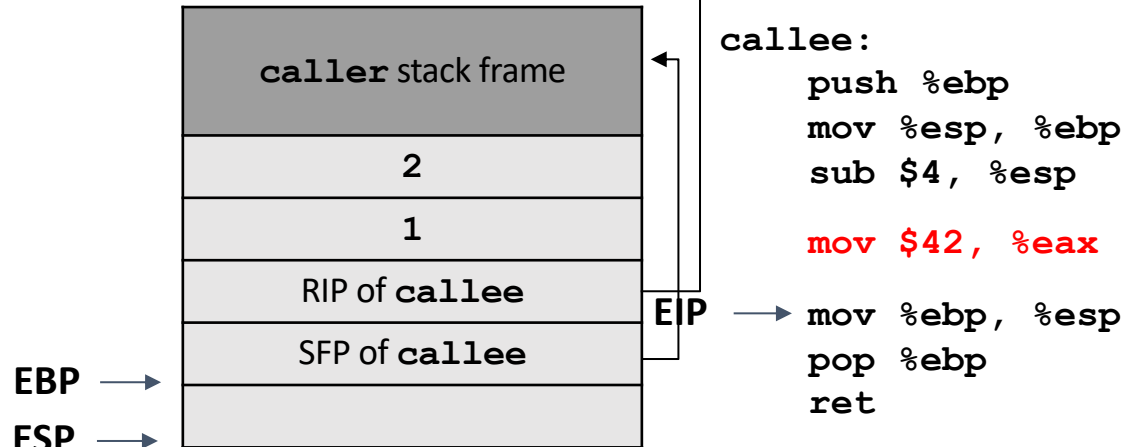
```
void caller(void) {  
    callee(1, 2);  
}
```

```
int callee(int a, int b) {  
    int local;  
    return 42;  
}
```

• Illustration

– Run the function

- The stack frame is set up
- The function can run
- This function just returns 42, so we put 42 in the EAX register



X86 FUNCTION CALL

```
void caller(void) {  
    callee(1, 2);  
}
```

```
int callee(int a, int b) {  
    int local;  
    return 42;  
}
```

• Illustration

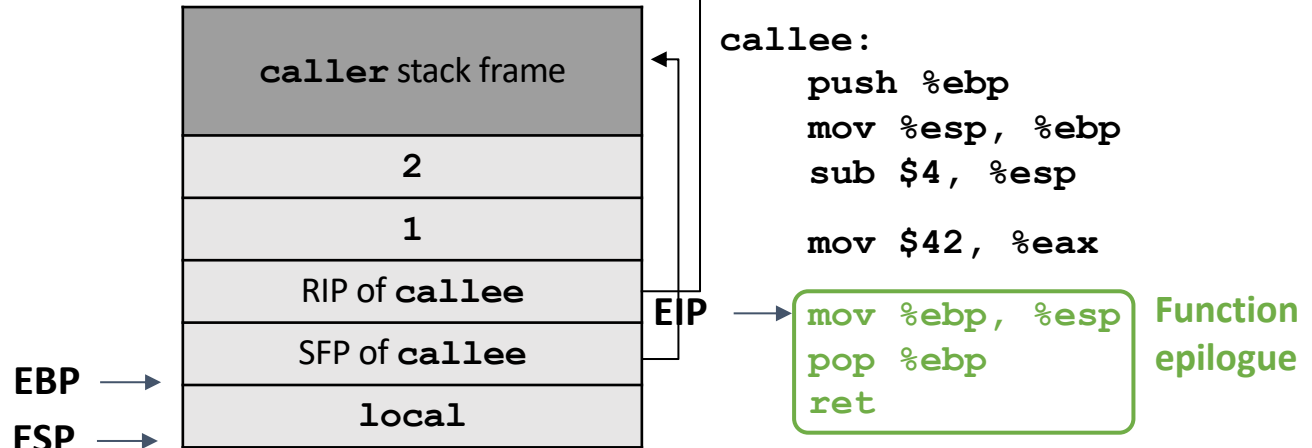
- The next 3 steps restore the caller's stack frame
- These instructions are sometimes called the **function epilogue**, because they appear at the end of every function
- Sometimes the **mov** and **pop** instructions are replaced with the **leave** and **ret** instruction

caller:

```
...  
push $2  
push $1  
call callee  
add $8, %esp  
...
```

callee:

```
push %ebp  
mov %esp, %ebp  
sub $4, %esp  
  
mov $42, %eax
```



X86 FUNCTION CALL

```
void caller(void) {  
    callee(1, 2);  
}
```

```
int callee(int a, int b) {  
    int local;  
    return 42;  
}
```

• Illustration

– Move ESP

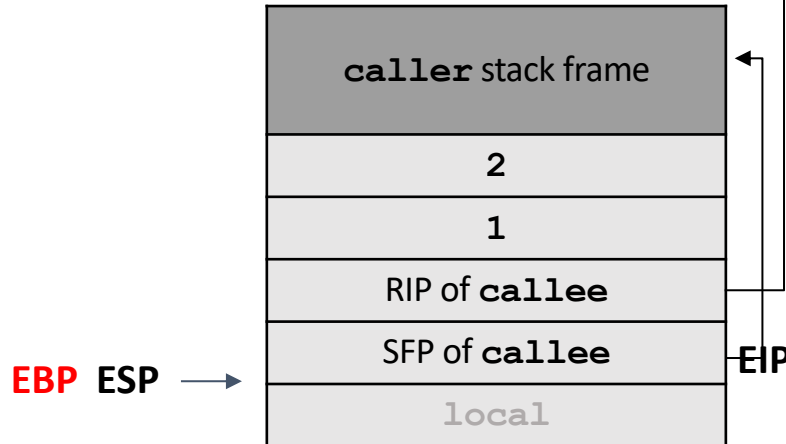
- This instruction moves the ESP up to where the EBP is located
- This effectively deletes the space allocated for the callee stack frame

caller:

```
...  
push $2  
push $1  
call callee  
add $8, %esp  
...
```

callee:

```
push %ebp  
mov %esp, %ebp  
sub $4, %esp  
  
mov $42, %eax  
  
mov %ebp, %esp  
pop %ebp  
ret
```



X86 FUNCTION CALL

```
void caller(void) {  
    callee(1, 2);  
}
```

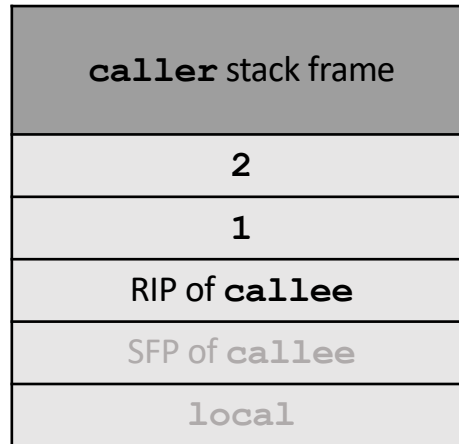
```
int callee(int a, int b) {  
    int local;  
    return 42;  
}
```

• Illustration

- Pop (restore) old EBP (SFP)
 - The **pop** instruction puts the SFP (saved EBP) back in EBP
 - It also increments ESP to delete the popped SFP from the stack

EBP →

ESP →



caller:

```
...  
push $2  
push $1  
call callee  
add $8, %esp  
...
```

callee:

```
push %ebp  
mov %esp, %ebp  
sub $4, %esp  
  
mov $42, %eax  
  
mov %ebp, %esp  
pop %ebp  
ret
```


X86 FUNCTION CALL

```
void caller(void) {  
    callee(1, 2);  
}
```


```
int callee(int a, int b) {  
    int local;  
    return 42;  
}
```

• Illustration

- Pop (restore) old EBP (SFP)
- The **ret** instruction acts like **pop %eip**
- It puts the next value on the stack (the RIP) into the EIP, which returns program execution to the caller
- It increases ESP to delete the popped RIP from the stack

```
caller:  
    ...  
    push $2  
    push $1  
    call callee  
    add $8, %esp  
    ...
```

```
callee:  
    push %ebp  
    mov %esp, %ebp  
    sub $4, %esp  
  
    mov $42, %eax  
  
    mov %ebp, %esp  
    pop %ebp  
    ret
```



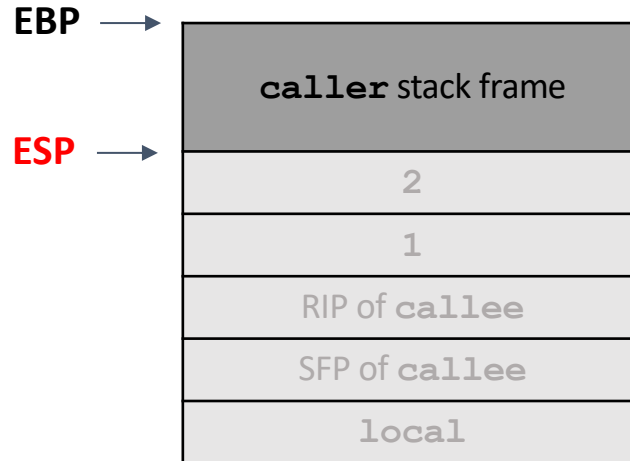
X86 FUNCTION CALL

```
void caller(void) {  
    callee(1, 2);  
}
```

```
int callee(int a, int b) {  
    int local;  
    return 42;  
}
```

• Illustration

- Remove arguments from stack
 - Back in the caller, we increment ESP to delete the arguments from the stack
 - The stack has returned to its original state before the function call



```
caller:  
    ...  
    push $2  
    push $1  
    call callee  
    add $8, %esp  
    ...  
EIP →
```

```
callee:  
    push %ebp  
    mov %esp, %ebp  
    sub $4, %esp  
  
    mov $42, %eax  
  
    mov %ebp, %esp  
    pop %ebp  
    ret
```

TOPICS FOR TODAY

- Preliminaries (x86 assembly and call stack)
 - C program
 - Memory layout
 - x86 architecture
 - Stack layout
 - Calling convention
 - x86 calling convention design
 - x86 calling convention example

Thank You!

Tu/Th 4:00 – 5:50 PM

Sanghyun Hong

sanghyun.hong@oregonstate.edu



Oregon State
University

SAIL
Secure AI Systems Lab