

CS 344: OPERATING SYSTEMS I

01.18: PART I - PROCESS

M/W 12:00 – 1:50 PM (LINC #200)

Sanghyun Hong

sanghyun.hong@oregonstate.edu



Oregon State
University

SAIL
Secure AI Systems Lab

NOTICE

- Announcements
 - Begin office hours
 - Time and locations: available on Canvas
 - Other times: at Discord server

Office Hours						
Time	Mon	Tue	Wed	Thu	Fri	
10:00 AM		Eunjin 10 AM - 1 PM (Zoom)			Eunjin 10 - 12:30 PM (Zoom)	
10:30 AM						
11:00 AM						
11:30 AM						
12:00 PM	[Grey]		[Grey]			
12:30 PM						
1:00 PM	Eunjin 2 - 6:30 PM (Zoom)	Radhika 1 - 4:30 PM (Zoom)		Radhika 1 - 4:30 PM (In-person)	Radhika 1:30 - 3 PM (Zoom)	
1:30 PM						
2:00 PM						
2:30 PM						
3:00 PM					Sanghyun 3 - 4:30 PM (Zoom)	
3:30 PM						
4:00 PM						
4:30 PM						
5:00 PM						
5:30 PM						
6:00 PM						

NOTICE – CONT'D

- Announcements
 - Begin office hours
 - Time and locations: available on Canvas
 - Other times: at Discord server
 - Notes
 - Discord: allow us a few hours to answer questions (2 TAs for 135+ students)
 - Discord: post questions to corresponding channels (e.g., #assignment-1 for the assignment 1)
 - Discord: feel free to DM instructor or TAs (Sanghyun, Radhika, or Eunjin)
 - All: help others, when you already know answers
(*do not share your code with others)

NOTICE – CONT'D

- Deadlines
 - (Passed) Syllabus quiz
 - (1/23 11:59 PM) Programming assignment 1
 - (1/30 11:59 PM) Midterm quiz 1

TOPICS FOR TODAY

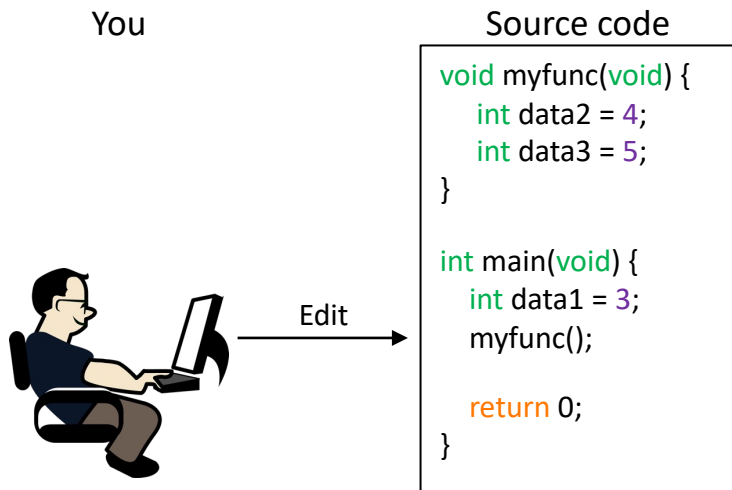
- Part I: Process
 - Provide abstraction
 - What is a program?
 - What is a process?
 - How does OS run a program?
 - Offer standard libraries
 - How do we run (or stop) a process?
 - How does OS manage the process(es) we ran?
 - Manage resources
 - (Note) We will talk about this in the “scheduling” class

PROVIDE ABSTRACTION: A PROGRAM

- (Computer) Program
 - **Definition:** a set of instructions for an OS to execute
 - **An example program for Linux computer**

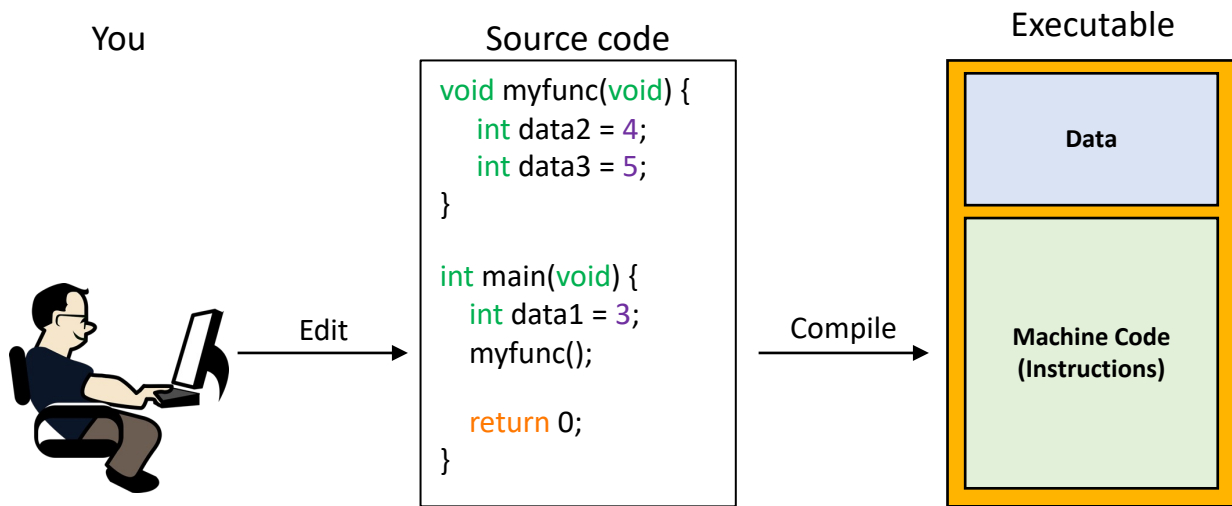
PROVIDE ABSTRACTION: A PROGRAM

- (Computer) Program
 - **Definition:** a set of instructions for an OS to execute
 - **An example program for Linux computer**



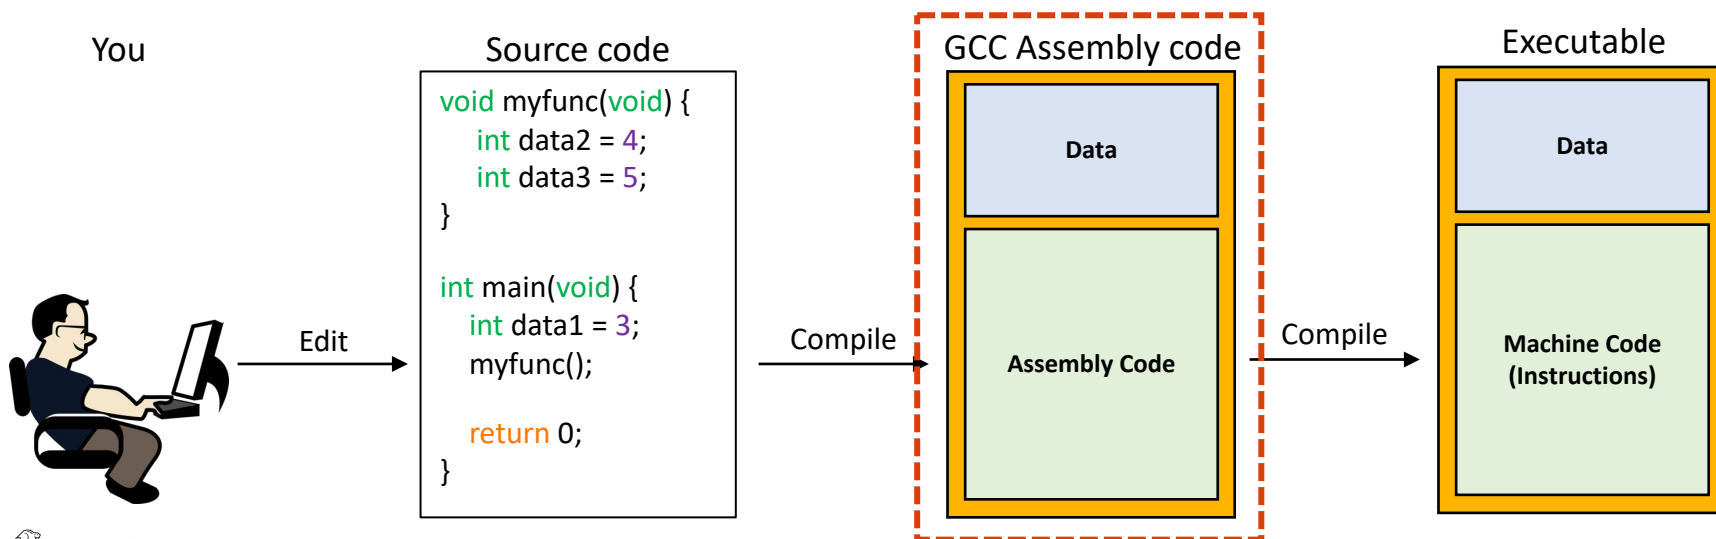
PROVIDE ABSTRACTION: A PROGRAM

- (Computer) Program
 - **Definition:** a set of instructions for an OS to execute
 - **An example program for Linux computer**



EXAMPLE: C COMPILATION WITH GCC

- GCC compilation
 - It converts source code to **assembly** code (`$ gcc -c -S <filename.c>`)
 - It then converts the assembly code to **instructions** (`$ gcc -c <filename.s> -o <filename.o>; gcc -o <filename.o> -o filename`)



EXAMPLE: C COMPILATION WITH GCC

- GCC compilation

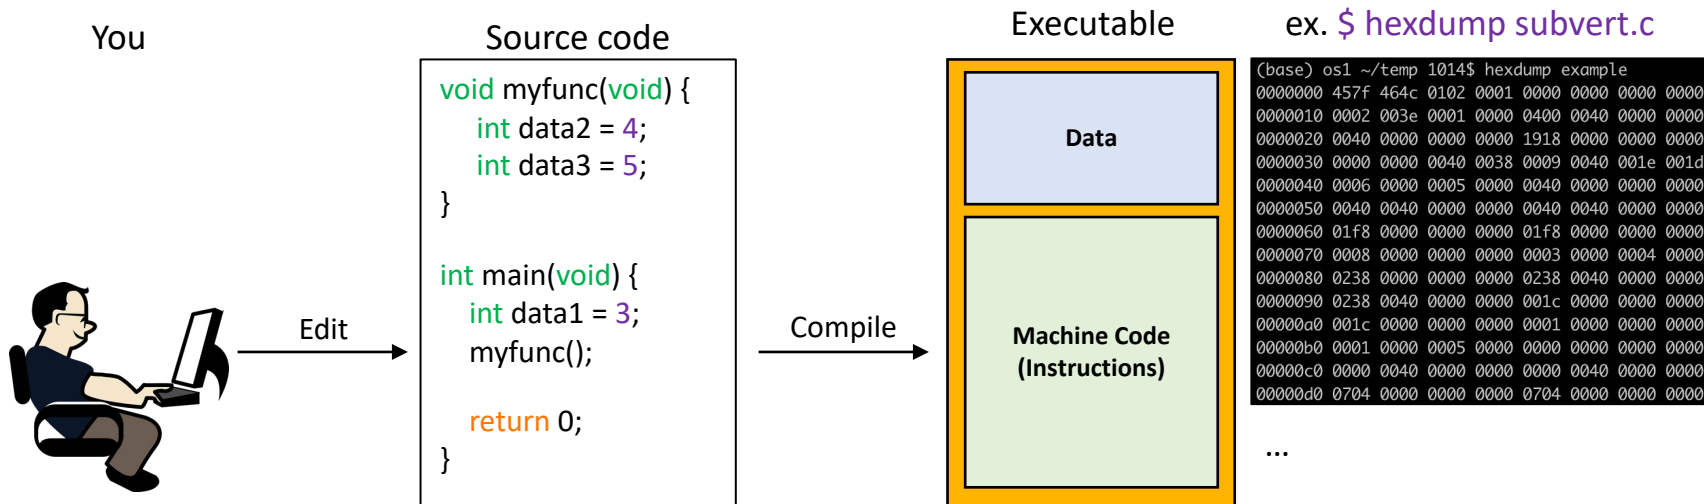
- It converts source code to **assembly** code (`$ gcc -c -S <filename.c>`)

```
.file "example.c"
.text
.globl myfunc
.type myfunc, @function
myfunc:
.LFB0:
.cfi_startproc
pushq %rbp
.cfi_def_cfa_offset 16
.cfi_offset 6, -16
movq %rsp, %rbp
.cfi_def_cfa_register 6
movl $4, -4(%rbp)
movl $5, -8(%rbp)
popq %rbp
.cfi_def_cfa 7, 8
ret
.cfi_endproc
.LFE0:
.size myfunc, .-myfunc
.globl main
.type main, @function
main:
example.s
```

```
.size myfunc, .-myfunc
.globl main
.type main, @function
main:
.LFB1:
.cfi_startproc
pushq %rbp
.cfi_def_cfa_offset 16
.cfi_offset 6, -16
movq %rsp, %rbp
.cfi_def_cfa_register 6
subq $16, %rsp
movl $3, -4(%rbp)
call myfunc
movl $0, %eax
leave
.cfi_def_cfa 7, 8
ret
.cfi_endproc
.LFE1:
.size main, .-main
.ident "GCC: (GNU) 4.8.5 20150623 (Red Hat 4.8.5-44)"
section .note.GNU-stack,"",@progbits
example.s
```


PROVIDE ABSTRACTION: A PROGRAM

- (Computer) Program
 - **Definition:** a set of instructions for an OS to execute
 - **An example program for Linux computer**



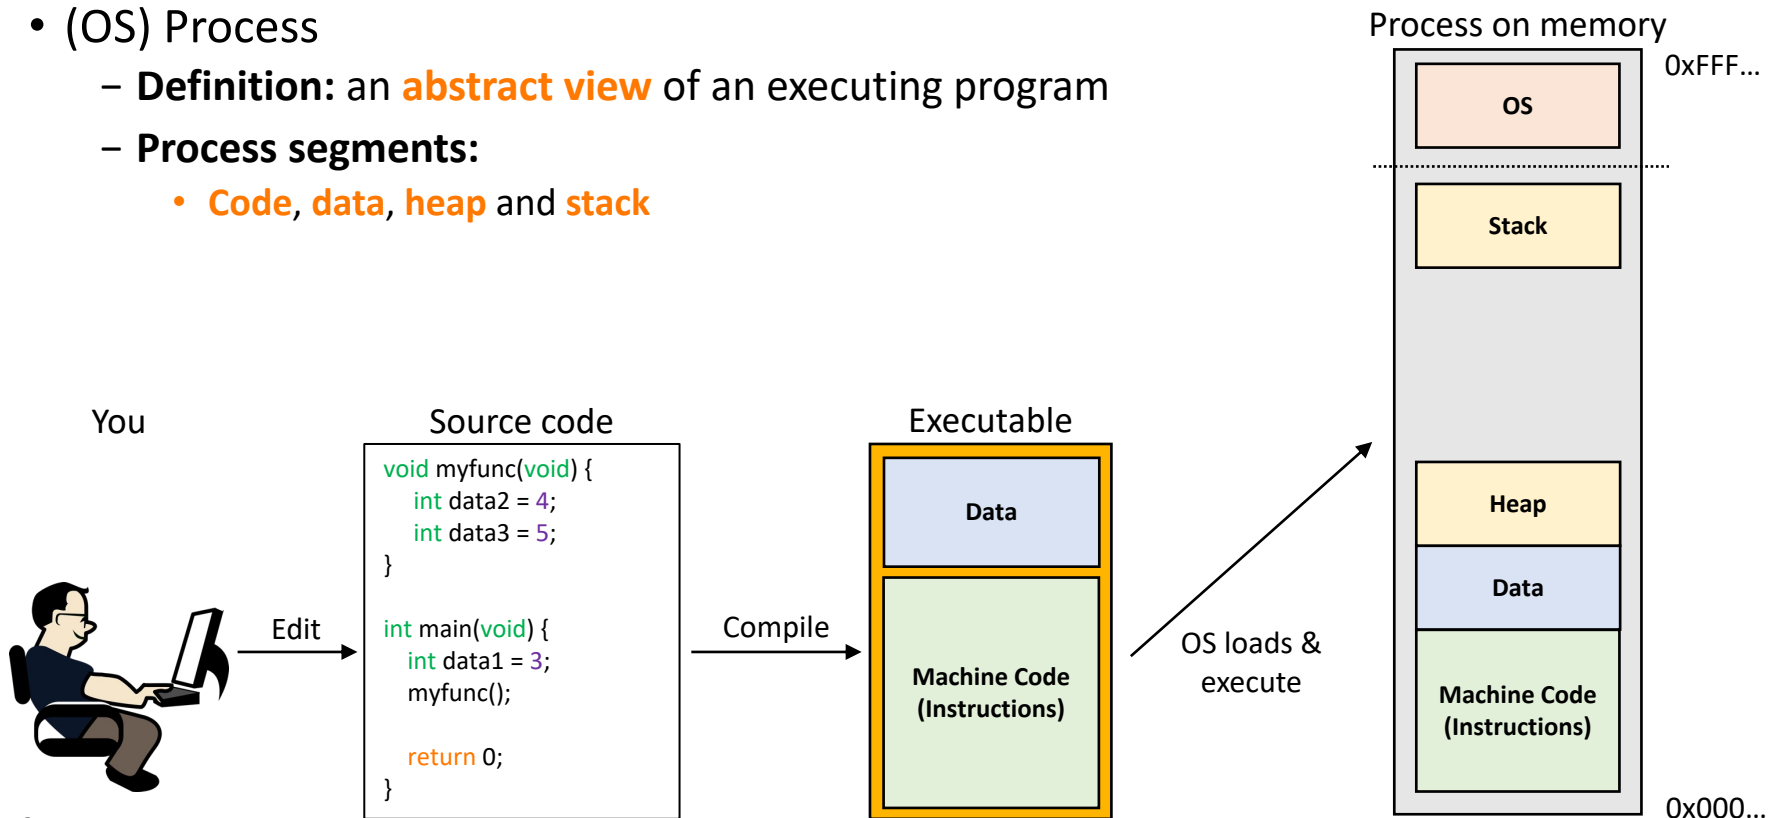
PROVIDE ABSTRACTION: A PROCESS

- (OS) Process
 - **Definition:** an **abstract view** of an executing program
 - **Process segments:**
 - **Code, data, heap** and **stack**



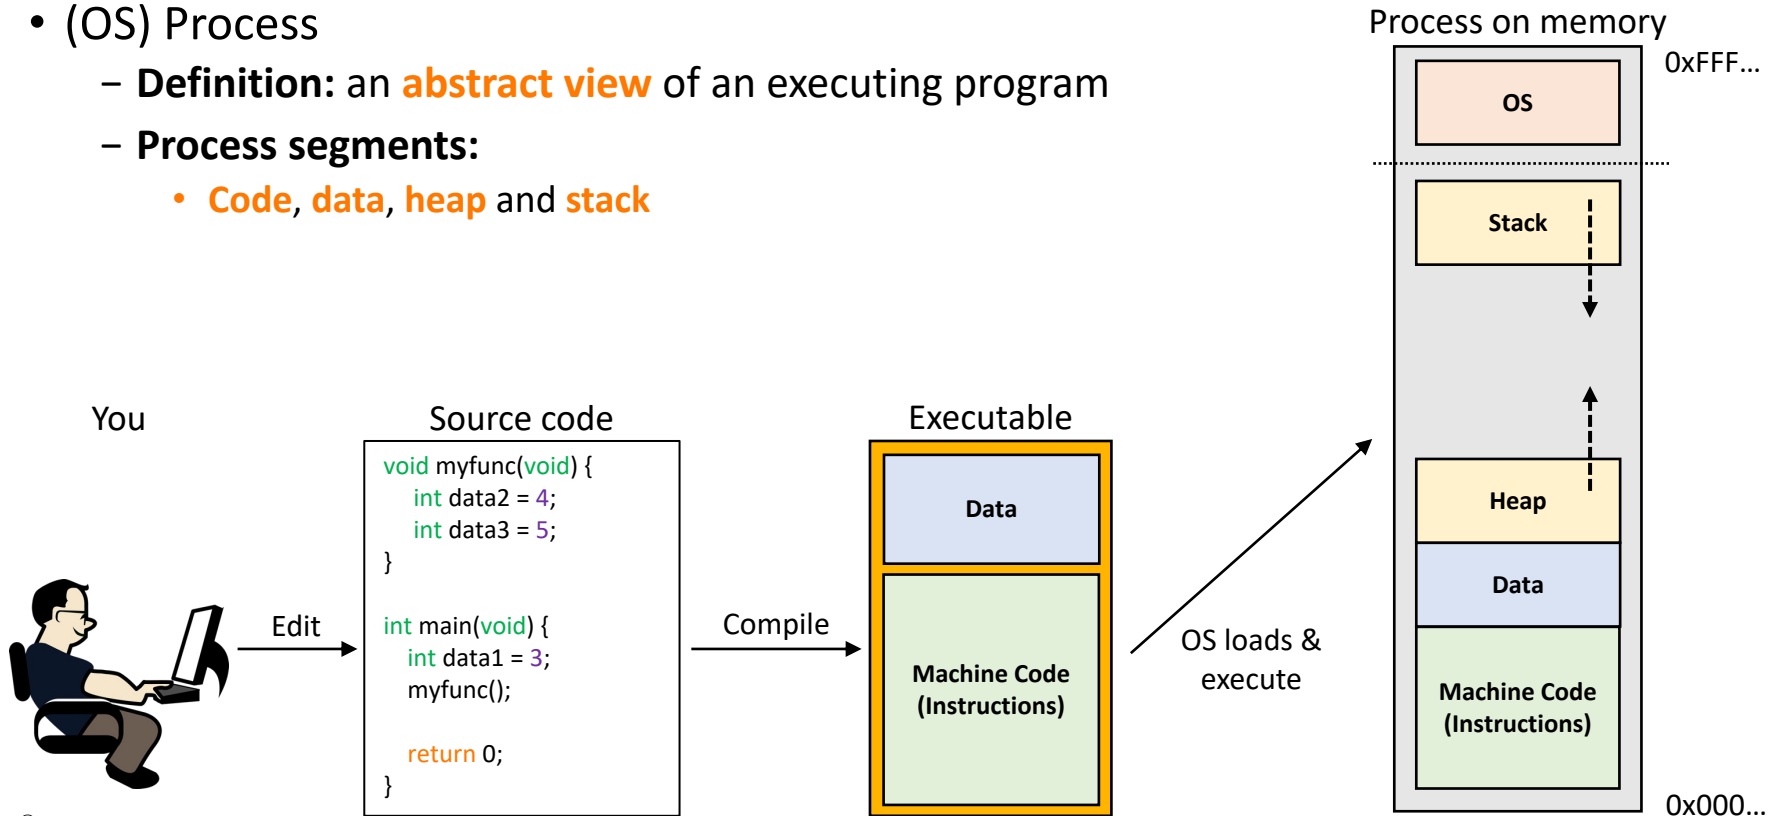
PROVIDE ABSTRACTION: A PROCESS

- (OS) Process
 - **Definition:** an **abstract view** of an executing program
 - **Process segments:**
 - **Code, data, heap** and **stack**



PROVIDE ABSTRACTION: A PROCESS

- (OS) Process
 - **Definition:** an **abstract view** of an executing program
 - **Process segments:**
 - **Code, data, heap** and **stack**



PROVIDE ABSTRACTION: HOW OS DEFINES A PROCESS?

- (Linux) has the process context

- Code

- Program counter
 - Instruction pointer

- Stack and heap

- Stack pointer
 - Heap pointer

- Running context

- Process state (ID, ...)
 - Execution flags
 - CPU # to run
 - (OS II) Scheduling policy
 - (OS II) Mem. virtualization

– ...


Process Context: A set of information that OS requires to run a process on a CPU, different from CPU vendors (ex. In Linux, it's defined as *task_struct*, [Link](#))

```
728 struct task_struct {
729 #ifdef CONFIG_THREAD_INFO_IN_TASK
730     /*
731      * For reasons of header soup (see current_thread_info()), this
732      * must be the first element of task_struct.
733      */
734     struct thread_info      thread_info;
735 #endif
736     unsigned int            __state;
737
738 #ifdef CONFIG_PREEMPT_RT
739     /* saved state for "spinlock sleepers" */
740     unsigned int            saved_state;
741 #endif
742
743     /*
744      * This begins the randomizable portion of task_struct. Only
745      * scheduling-critical items should be added above here.
746      */
747     randomized_struct_fields_start
748
749     void                    *stack;
750     refcount_t              usage;
751     /* Per task flags (PF_*), defined further below: */
752     unsigned int            flags;
753     unsigned int            ptrace;
754
755     struct sched_info       sched_info;
756
757     struct list_head        tasks;
758 #ifdef CONFIG_SMP
759     struct plist_node        pushable_tasks;
760     struct rb_node          pushable_dl_tasks;
761 #endif
762
763     struct mm_struct         *mm;
764     struct mm_struct         *active_mm;
765
766     /* Per-thread vma caching: */
767     struct vmacache         vmacache;
768
769 #ifdef SPLIT_RSS_COUNTING
770     struct task_rss_stat     rss_stat;
771 #endif
772
773     int                     exit_state;
774     int                     exit_code;
775     int                     exit_signal;
776     /* The signal sent when the parent dies: */
777     int                     pdeath_signal;
778     /* JOBCTL_*, siglock protected: */
779     unsigned long           jobctl;
780
781     /* Used for emulating ABI behavior of previous Linux versions: */
782     unsigned int            personality;

```

PROVIDE ABSTRACTION: HOW OS LOADS A PROCESS?

- (OS) Process
 - **Definition:** an **abstract view** of an executing program
 - **Load a process:**
 - **Code:** OS loads the instructions to “code” segments
 - **Data :** OS loads the data (such as static vars) to “data” segments
 - **Stack** and **heap:** OS creates those mem. spaces
 - (Ready) OS sets the program counter (PC) to the first code location



PROVIDE ABSTRACTION: HOW OS RUNS A PROCESS?

- OS makes the CPU run the machine code
 - Example: IBM machines
 - Submit a **punch card** that have a set of **instructions**
 - Machine **reads** instructions line by line and **do sth.**

Punch card

Example punch holes	instructions
● ○ ● ○ ○ ○ ○	// load 8
○ ● ● ○ ○ ● ○	// load 5
● ● ● ○ ○ ○ ○	// add 8 and 5
...	
...	

PROVIDE ABSTRACTION: HOW OS RUNS A PROCESS? – CONT'D

- OS makes the CPU run the machine code
 - Example: IBM machines
 - Submit a **punch card** that have a set of **instructions**
 - Machine **reads** instructions line by line and **do sth.**
 - Modern computers
 - Machine := a processor (CPU)
 - Instructions := instructions (100+ for Intel CPUs)
 - Punch card := a process in memory
 - Operates := execute the instructions

Punch card

Example punch holes	instructions	
● ○ ● ○ ○ ○ ○	// load 8	←---
○ ● ● ○ ○ ● ○	// load 5	
● ● ● ○ ○ ○ ○	// add 8 and 5	
...		
...		

Memory

Example instructions	operations	
0x11 0x12 0x05 0x00	// load 5 to r12	←---
0x08 0x12 0x08 0x00	// add r12 and 8	
0x12 0xF9 0xFF 0xF4	// store r12	
...		
...		
...		

PROVIDE ABSTRACTION: HOW OS RUNS A PROCESS? – CONT'D

- OS makes the CPU run the machine code
 - Example: IBM machines
 - Submit a **punch card** that have a set of **instructions**
 - Machine **reads** instructions line by line and **do sth.**
 - Modern computers
 - Machine := a processor (CPU)
 - Instructions := instructions (100+ for Intel CPUs)
 - Punch card := a process in memory
 - Operates := execute the instructions

The program counter (PC) in a CPU is always holding the memory address where the next instruction to execute is

Punch card


Example punch holes	instructions
● ○ ● ○ ○ ○ ○	// load 8
○ ● ● ○ ○ ● ○	// load 5
● ● ● ○ ○ ○ ○	// add 8 and 5
...	
...	

Memory

Example instructions	operations
0x11 0x12 0x05 0x00	// load 5 to r12
0x08 0x12 0x08 0x00	// add r12 and 8
0x12 0xF9 0xFF 0xF4	// store r12
...	
...	
...	


PROVIDE ABSTRACTION: HOW OS LOADS/RUNS A PROCESS?

- (OS) Process
 - **Definition:** an **abstract view** of an executing program
 - **Load a process:**
 - **Code:** OS loads the instructions to “code” segments
 - **Data :** OS loads the data (such as static vars) to “data” segments
 - **Stack** and **heap:** OS creates those mem. spaces
 - **(Ready) OS** sets the **program counter (PC)** to the first code location



PROVIDE ABSTRACTION: STACK VS. HEAP

- Stack vs. heap
 - **Definition:** Both are the **areas of memory**
 - Stack
 - **OS controls** the memory allocations (size)
 - Store data in Last in first out (**LIFO**) manner
 - Stack mostly holds data initialized within a function




PROVIDE ABSTRACTION: STACK VS. HEAP

- Stack vs. heap
 - **Definition:** Both are the **areas of memory**
 - Stack
 - **OS controls** the memory allocations (size)
 - Store data in Last in first out (**LIFO**) manner
 - Stack mostly holds data initialized within a function

```
void myfunc(void) {  
    int data2 = 4;  
    int data3 = 5;  
}
```

```
int main(void) {  
    int data1 = 3;  
    myfunc();  
  
    return 0;  
}
```

←----- Run




PROVIDE ABSTRACTION: STACK VS. HEAP

- Stack vs. heap
 - **Definition:** Both are the **areas of memory**
 - Stack
 - **OS controls** the memory allocations (size)
 - Store data in Last in first out (**LIFO**) manner
 - Stack mostly holds data initialized within a function

```
void myfunc(void) {  
    int data2 = 4;  
    int data3 = 5;  
}  
  
int main(void) {  
    int data1 = 3;  
    myfunc();  
  
    return 0;  
}
```


←----- Run



PROVIDE ABSTRACTION: STACK VS. HEAP

- Stack vs. heap
 - **Definition:** Both are the **areas of memory**
 - Stack
 - **OS controls** the memory allocations (size)
 - Store data in Last in first out (**LIFO**) manner
 - Stack mostly holds data initialized within a function

```
void myfunc(void) {  
    int data2 = 4;  
    int data3 = 5;    ←----- Run  
}  
  
int main(void) {  
    int data1 = 3;  
    myfunc();  
  
    return 0;  
}
```




PROVIDE ABSTRACTION: STACK VS. HEAP

- Stack vs. heap
 - **Definition:** Both are the **areas of memory**
 - Stack
 - **OS controls** the memory allocations (size)
 - Store data in Last in first out (**LIFO**) manner
 - Stack mostly holds data initialized within a function

```
void myfunc(void) {  
    int data2 = 4;  
    int data3 = 5;  
}
```

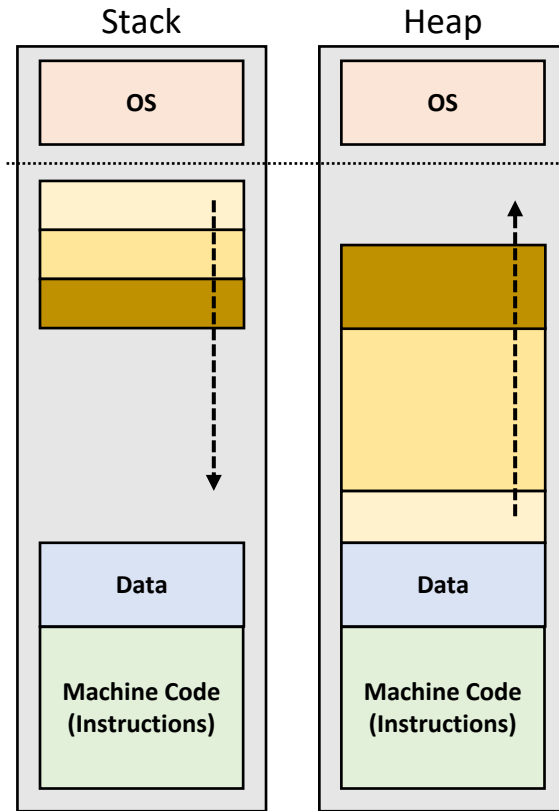
```
int main(void) {  
    int data1 = 3;  
    myfunc();  
  
    return 0;  
}
```

←----- Run



PROVIDE ABSTRACTION: STACK VS. HEAP

- Stack vs. heap
 - **Definition:** Both are the **areas of memory**
 - Heap
 - **User allocates** the memory with a specific size
 - **OS finds an empty space** and then place the mem.
 - Mem. fragmentation (also **mem. leak!**) can occur



PROVIDE ABSTRACTION: STACK VS. HEAP

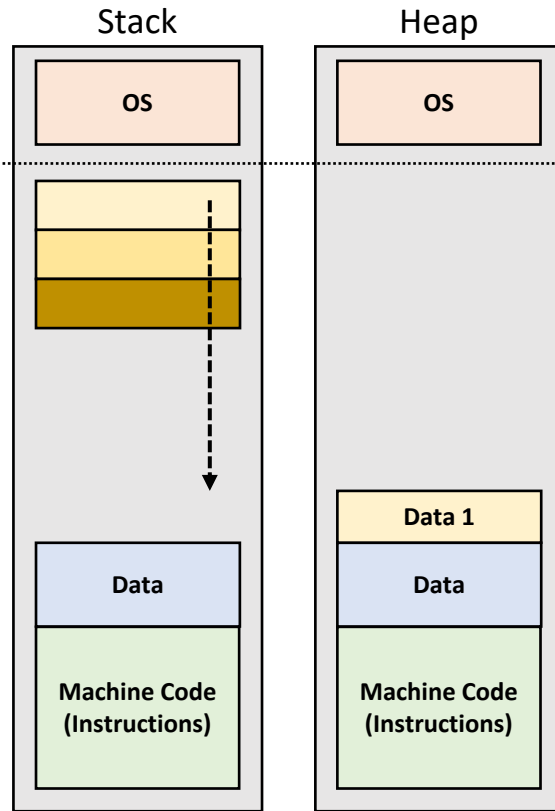
- Stack vs. heap

- **Definition:** Both are the **areas of memory**
- Heap
 - **User allocates** the memory with a specific size
 - **OS finds an empty space** and then place the mem.
 - Mem. fragmentation (also **mem. leak!**) can occur

```
void myfunc(void) {  
    char *data2 = (char *) malloc(5);  
    char *data3 = (char *) malloc(2);  
    free(data2);  
}
```

```
int main(void) {  
    char *data1 = (char *) malloc(1);  
    myfunc();  
  
    return 0;
```

←---- Run



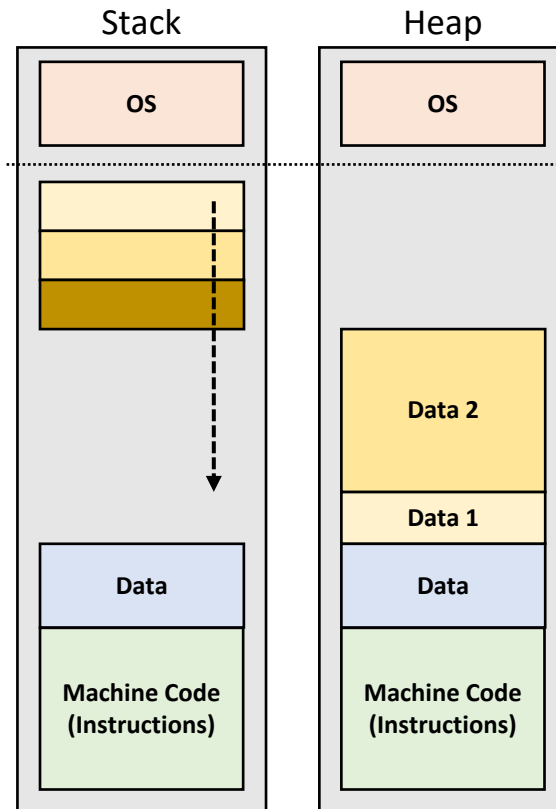
PROVIDE ABSTRACTION: STACK VS. HEAP

- Stack vs. heap

- **Definition:** Both are the **areas of memory**
- Heap
 - **User allocates** the memory with a specific size
 - **OS finds an empty space** and then place the mem.
 - Mem. fragmentation (also **mem. leak!**) can occur

```
void myfunc(void) {  
    char *data2 = (char *) malloc(5);  
    char *data3 = (char *) malloc(2);  
    free(data2);  
}  
  
int main(void) {  
    char *data1 = (char *) malloc(1);  
    myfunc();  
  
    return 0;  
}
```

←----- Run

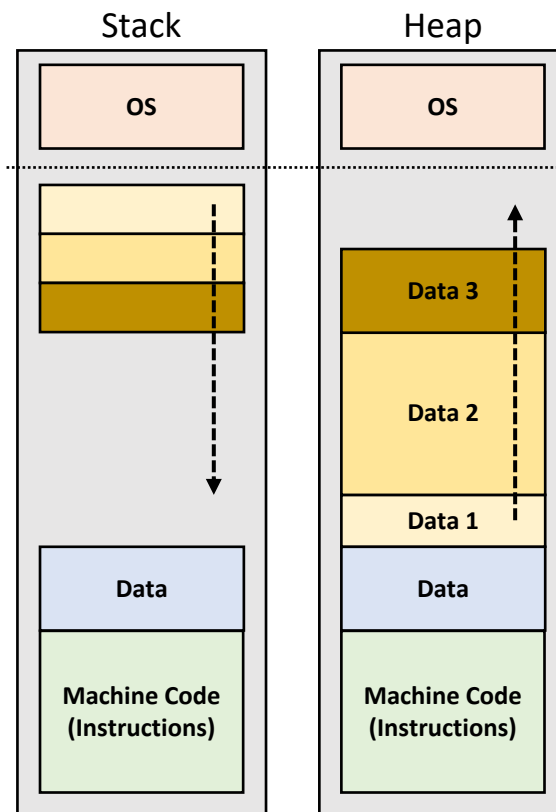


PROVIDE ABSTRACTION: STACK VS. HEAP

- Stack vs. heap
 - **Definition:** Both are the **areas of memory**
 - Heap
 - **User allocates** the memory with a specific size
 - **OS finds an empty space** and then place the mem.
 - Mem. fragmentation (also **mem. leak!**) can occur

```
void myfunc(void) {  
    char *data2 = (char *) malloc(5);  
    char *data3 = (char *) malloc(2);  
    free(data2);  
}  
  
int main(void) {  
    char *data1 = (char *) malloc(1);  
    myfunc();  
  
    return 0;  
}
```

←---- Run



PROVIDE ABSTRACTION: STACK VS. HEAP

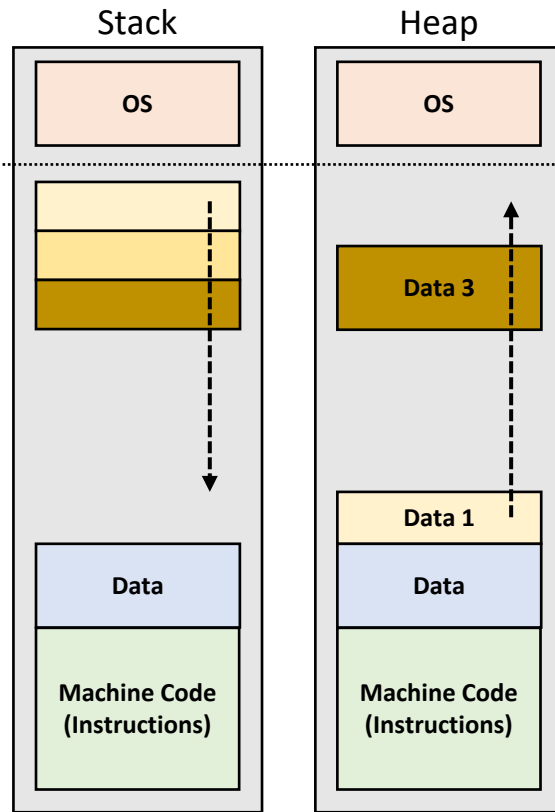
- Stack vs. heap

- **Definition:** Both are the **areas of memory**
- Heap
 - **User allocates** the memory with a specific size
 - **OS finds an empty space** and then place the mem.
 - Mem. fragmentation (also **mem. leak!**) can occur

```
void myfunc(void) {  
    char *data2 = (char *) malloc(5);  
    char *data3 = (char *) malloc(2);  
    free(data2);  
}
```

←----- Run

```
int main(void) {  
    char *data1 = (char *) malloc(1);  
    myfunc();  
  
    return 0;  
}
```



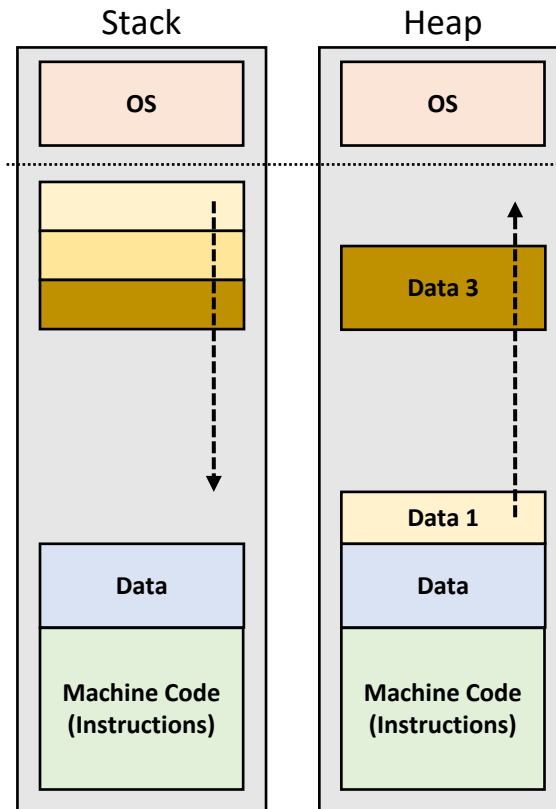
PROVIDE ABSTRACTION: STACK VS. HEAP

- Stack vs. heap
 - **Definition:** Both are the **areas of memory**
 - Heap
 - **User allocates** the memory with a specific size
 - **OS finds an empty space** and then place the mem.
 - Mem. fragmentation (also **mem. leak!**) can occur

```
void myfunc(void) {  
    char *data2 = (char *) malloc(5);  
    char *data3 = (char *) malloc(2);  
    free(data2);  
}
```

```
int main(void) {  
    char *data1 = (char *) malloc(1);  
    myfunc();  
  
    return 0;
```

←---- Run



TOPICS FOR TODAY

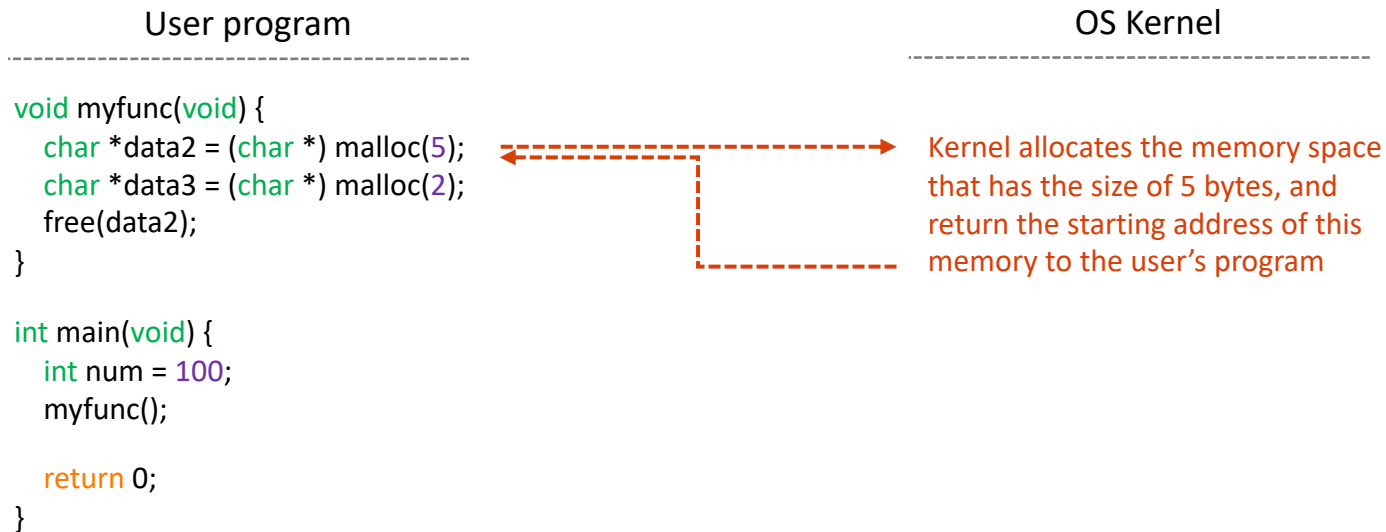
- Part I: Process
 - Provide abstraction
 - What is a program?
 - What is a process?
 - How does OS run a program?
 - Offer standard libraries
 - How do we run (or stop) a process?
 - How does OS manage the process(es) we ran?
 - Manage resources
 - (Note) We will talk about this in the “scheduling” class

OFFERS STANDARD INTERFACE

- How do we run a process?
 - Double click an icon
 - Type `./<program name>` in the terminal
 - ...


OFFERS STANDARD INTERFACE: SYSTEM CALL

- System call
 - **Definition:** a user-level function call to request a service from the OS
 - **Example:** when we allocate memory with “`malloc()`”




OFFERS STANDARD INTERFACE: SYSTEM CALL

- OS offers a set of system calls
 - To create/terminate a process
 - To open/read/write/close a file
 - To request/release a device (such as display, mouse, etc.)
 - To request/modify system information
 - To initiate/close networking
 - To set the security properties
 - ...



OFFERS STANDARD INTERFACE: **FORK** SYSTEM CALL

- fork() system call
 - **Operation:**
 - Create a new process that is an exact copy of the calling process
 - Return the process ID (PID) of a new process (and if it's in child, returns 0)



OFFERS STANDARD INTERFACE: FORK SYSTEM CALL

- `fork()` sample code in C

```
#include <stdio.h>
#include <sys/types.h>
#include <unistd.h>
```

```
int main(void) {
    int number = 10;
    pid_t pid;
    switch (pid = fork()) {
        case -1:
            perror ("fork");
            exit (1);
        case 0:
            number++;
            printf("I am a child process [%d]!", number);
            break;
        default:
            number--;
            printf("I am a parent process [%d]!", number);
            break;
    }

    printf("I will be executed by both");
    return 0;
}
```

Parent process
(pid = child's PID)

Child process
(pid = 0)

Execution result (sample):

I am a child process [11]!
I will be executed by both
I am a parent process [9]!
I will be executed by both

```
switch (pid = fork()) {
    case -1:
        perror ("fork");
        exit (1);
    case 0:
        number++;
        printf("I am a child process [%d]!", number);
        break;
    default:
        number--;
        printf("I am a parent process [%d]!", number);
        break;
}

printf("I will be executed by both");
return 0;
}
```

OFFERS STANDARD INTERFACE: FORK SYSTEM CALL

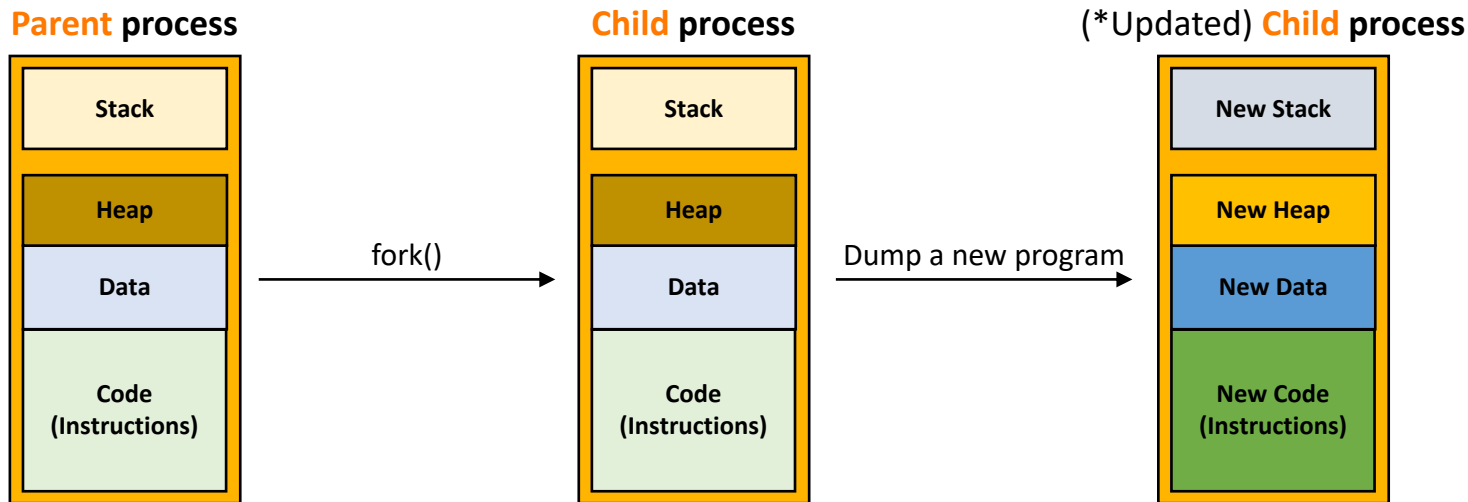
- fork() system call
 - **Operation:**
 - Create a new process that is an exact copy of the calling process
 - Return the process ID (PID) of a new process (and if it's in child, returns 0)
- Other system calls
 - exec(program to run):
 - Create a new process with fork() and dump the program to run into it
 - Return 0 if exec() is successful; otherwise, it returns the corresponding error
 - wait(status) or wait(PID):
 - Make the current process wait until the status (of a process, PID) changes
 - Returns the PID of the process that changes the status; otherwise, -1
 - exit() or kill():
 - Terminate the process with the given PID

OFFERS STANDARD INTERFACE: EXEC SYSTEM CALL

- `exec()` system call

- **Operation:**

- Create a new process with `fork()` and dump the program to run into it
 - Return 0 if `exec()` is successful; otherwise, it returns the corresponding error



OFFERS STANDARD INTERFACE: WHAT IF WE DO FORK INFINITELY?

- **fork() bomb** ([link](#))
 - A DoS attack that a process continuously fork() to deplete available system resources
 - **Consequence:** resource starvation
 - **Defense:** limit the number of processes a user can create ([check with \\$ ulimit -u](#))
- **Take-aways**
 - An attacker can exploit the standard interfaces for achieving adversarial goals
 - We should consider the worst-cases when designing/offering such interfaces
 - Defense mechanisms should also be offered to defeat such attacks


OFFER STANDARD INTERFACE: HOW OS MANAGES PROCESSES?

- Possible scenarios

- S1: Recursively fork()



- S2: Multiple fork()s from a process




What Would Be the Best Data Structure to Manage Processes?

OFFER STANDARD INTERFACE: HOW OS MANAGES PROCESSES?

- **fork() tree**

- OS manages processes with a tree
- Use (`$ pstree`) command to see the tree!
- Root of the fork() tree (in Linux)
 - PID=0: **Sched** (swapper) process
 - PID=1: **Init** process



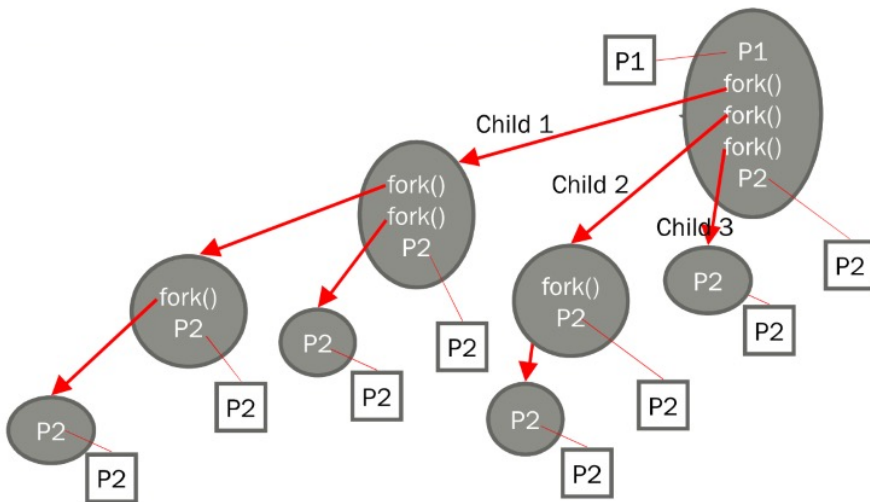
OFFER STANDARD INTERFACE: HOW OS MANAGES PROCESSES?

- **fork() tree**

- OS manages processes with a tree
- Use (`$ pstree`) command to see the tree!
- Root of the fork() tree (in Linux)
 - PID=0: **Sched** (swapper) process
 - PID=1: **Init** process

- Properties

- User processes always have a parent
- If we kill the parent, all the child processes will be killed, too (an exception, any process launched by `$ nohup` or `$ disown`)
- PIDs allocated by OS increases as we fork() more



TOPICS COVERED TODAY

- Part I: Process
 - Provide abstraction
 - What is a program?
 - What is a process?
 - How does OS run a program?
 - Offer standard libraries
 - How do we run (or stop) a process?
 - How does OS manage the process(es) we ran?
 - Manage resources
 - (Note) We will talk about this in the “scheduling” class

Thank You!

M/W 12:00 – 1:50 PM (LINC #200)

Sanghyun Hong

sanghyun.hong@oregonstate.edu



Oregon State
University

SAIL
Secure AI Systems Lab