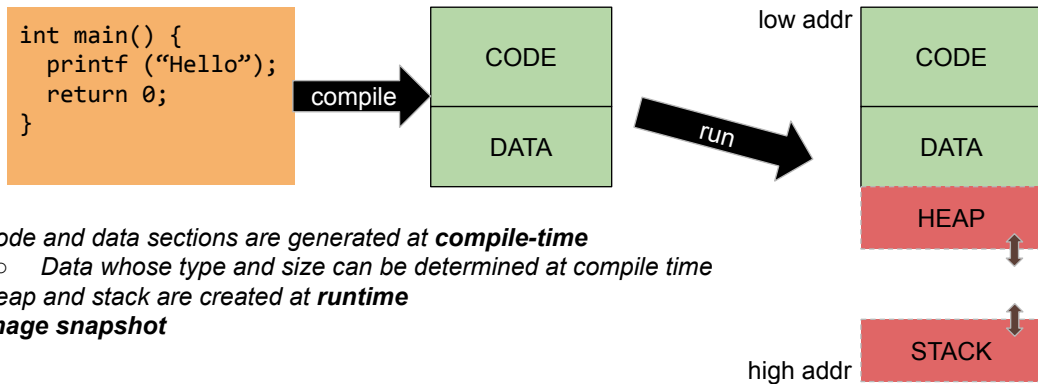


The slide features two decorative L-shaped lines made of thin brown lines. One is in the upper-left quadrant, and the other is in the lower-right quadrant, framing the central text.

# Unix Process Management

Process  $\neq$  Program

# Source $\Rightarrow$ Binary Executable $\Rightarrow$ Process (Image)



- Code and data sections are generated at **compile-time**
  - Data whose type and size can be determined at compile time
- Heap and stack are created at **runtime**
- **Image snapshot**

3

## Data Section vs. Heap vs. Stack

- Data section
  - Global variables (lifetime is NOT tied to any function, but tied to the entire program)
- Heap
  - Dynamically allocated memory (`malloc()`, Java `new`, C++ `new`)
- Stack
  - **Local** variables (lifetime is tied to the function lifetime)
  - C **fixed length** arrays
  - **Function Parameters**
  - **Return addresses**

4

# Memory Allocation: ([PollEv.com](http://PollEv.com))

```
float here;

int main() {
    int val;
    float *numbers;
    numbers = calloc (100, sizeof(float));
    for (.....) {

    }
    free (numbers);
    return 0;
}
```

Data section, stack, or heap?

- here is allocated in \_\_\_\_\_
- val is allocated in \_\_\_\_\_
- numbers is allocated in \_\_\_\_\_
- The 400 bytes from calloc() is allocated in \_\_\_\_\_

Assume 4-byte floats

5

# fork()

- Create a new (child) process using the **current copy** of the parent image
  - Parent and child processes are like **twins!**
- **At the time of fork()**
  - Code section: **exact duplicate** of each other
  - Heap section: **exact duplicate** of each other (most likely)
- **After return from fork()**
  - *Data (or Stack) section may differ by a few bytes*
- Thereafter, the two images are **independent/unrelated**
  - Parent and child **share NOTHING** (*not actually TRUE, further explanation later*)



7

# Fork()

```
/* parent */
int main() {
    printf ("Begin\n");
    fork();
    printf ("PID = %d\n, getpid());
    printf ("End\n");
    return 0;
}
```

# Fork()

```
/* parent */
int main() {
    printf ("Begin\n");
    fork();
    printf ("PID = %d\n, getpid());
    printf ("End\n");
    return 0;
}
```

# Fork()

a new process is created, DUPLICATING the parent process image

```
/* parent */
int main() {
    printf ("Begin\n");
    fork();
    printf ("PID = %d\n, getpid());
    printf ("End\n");
    return 0;
}
```

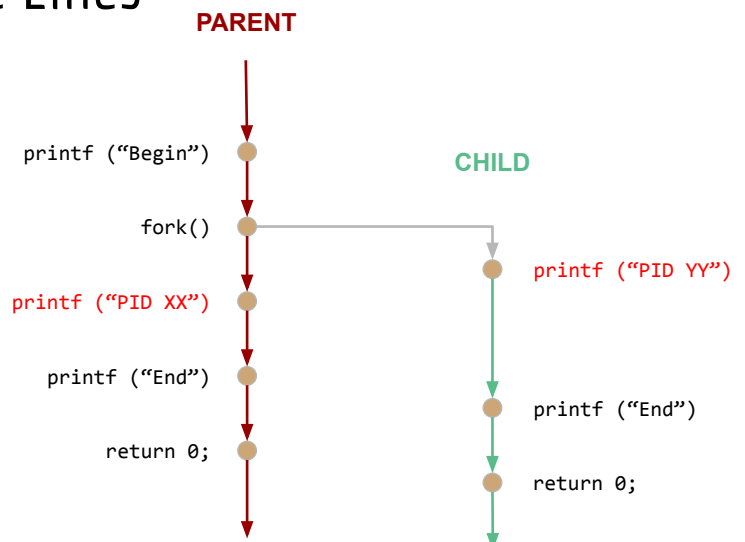
```
/* child */
int main() {
    printf ("Begin\n");
    fork();
    printf ("PID = %d\n, getpid());
    printf ("End\n");
    return 0;
}
```

At this point, both parent and child will run *independently* competing for the same CPU

10

# Parent & Child Time Lines


```
/* parent */
int main() {
    printf ("Begin\n");
    fork();
    printf ("PID = %d\n, getpid());
    printf ("End\n");
    return 0;
}
```



11

# Parent & Child Time Lines (Horizontal)

```
/* parent */
int main() {
    printf ("Begin\n");
    fork();
    printf ("PID = %d\n, getpid());
    printf ("End\n");
    return 0;
}
```



12

## Review

1. fork() creates a new process (child) by copying the current image (parent)
  - a. The child image is **EXACT DUPLICATE** of the parent image
2. **No** code/data/heap/stack are shared between parent and child
  - a. When one modifies its data/heap/stack the other won't see it
3. Both processes (parent & child) will **run independently** and compete for the same CPU(s) on your system
  - a. We never know the relative order of execution **ACROSS** the two processes
  - b. We only know the relative order of execution **WITHIN** each process
4. **Both processes** will resume execution to the **next statement after fork()**

13

True / False?

*The OS inspects each instruction of a program before it runs on the CPU*

14

## Fork()

```
/* parent */
int main() {
    printf ("Begin\n");
    pid_t who = fork();
    if (who == 0)
        printf ("Mug %d\n", getpid());
    else {
        printf ("Cup %d\n", who);
        printf ("Bowl %d\n", getpid());
    }

    printf ("End\n");
    return 0;
}
```

15

# Fork()

```
/* parent */
int main() {
    printf ("Begin\n");
    pid_t who = fork();
    if (who == 0)
        printf ("Mug %d\n", getpid());
    else {
        printf ("Cup %d\n", who);
        printf ("Bowl %d\n", getpid());
    }

    printf ("End\n");
    return 0;
}
```

fork() return value:

- ZERO (in child process)
- child PID (in parent process)

16

# Fork()

```
/* parent */
int main() {
    printf ("Begin\n");
    pid_t who = fork();
    if (who == 0)
        printf ("Mug %d\n", getpid());
    else {
        printf ("Cup %d\n", who);
        printf ("Bowl %d\n", getpid());
    }

    printf ("End\n");
    return 0;
}
```

```
/* child */
int main() {
    printf ("Begin\n");
    pid_t who = fork();
    if (who == 0)
        printf ("Mug %d\n", getpid());
    else {
        printf ("Cup %d\n", who);
        printf ("Bowl %d\n", getpid());
    }

    printf ("End\n");
    return 0;
}
```

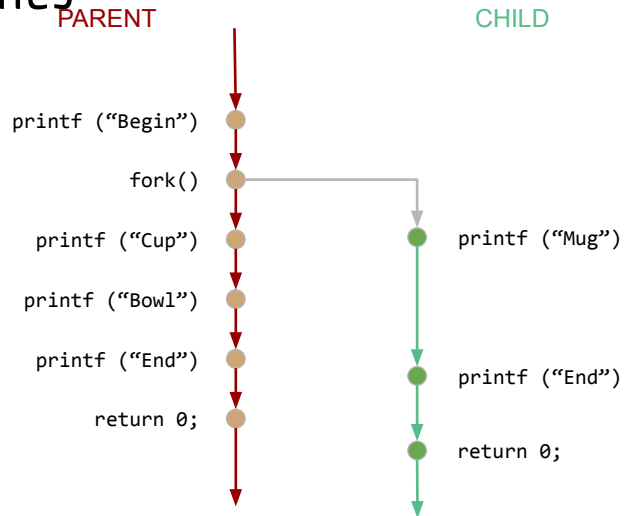
17



# Parent & Child Time lines

```
/* parent */
int main() {
    printf ("Begin\n");
    pid_t who = fork();
    if (who == 0)
        printf ("Mug %d\n, getpid());
    else {
        printf ("Cup %d\n, who);
        printf ("Bowl %d\n", getpid());
    }

    printf ("End\n");
    return 0;
}
```



18

## exit() & wait()

- **exit(N)**: terminate and report its status number (N) to parent
  - Every process in Unix/Linux (except init) has a parent
  - exit() should be called by a "child" process
  - Automatically called when returning from main()
    - `return 71;` translates to `exit(71);`
- **wait()**: wait for a child to terminate, and accept its status
  - wait() should be called by a "parent" process who spawns child processes

19

# Parent-Child Handshake: `exit()` $\leftrightarrow$ `wait()`

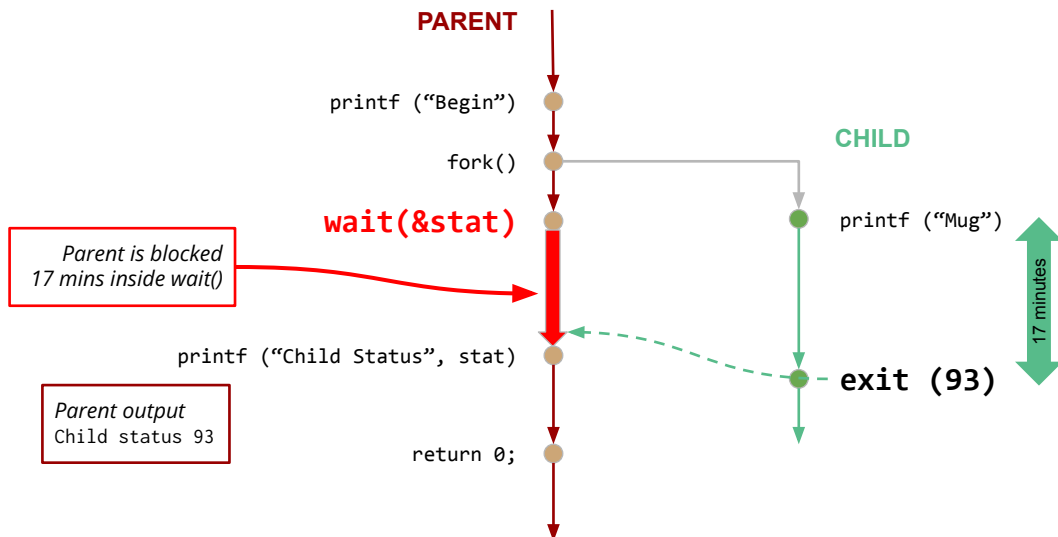
```
/* parent */
int main() {
    pid_t who = fork();
    if (who == 0) {
        printf ("Mug %d\n", getpid());
        exit (93);
    }
    else {
        int status;
        who = wait (&status);
        printf ("Child status is %d\n",
            WEXITSTATUS(status));
    }
    return 0;
}
```

parent output "Child status is 93"

```
/* child */
int main() {
    pid_t who = fork();
    if (who == 0) {
        printf ("Mug %d\n", getpid());
        exit (93); // any number you like
    }
    else {
        int status;
        who = wait (&status);
        printf ("Child status is %d\n",
            WEXITSTATUS(status));
    }
    return 0;
}
```

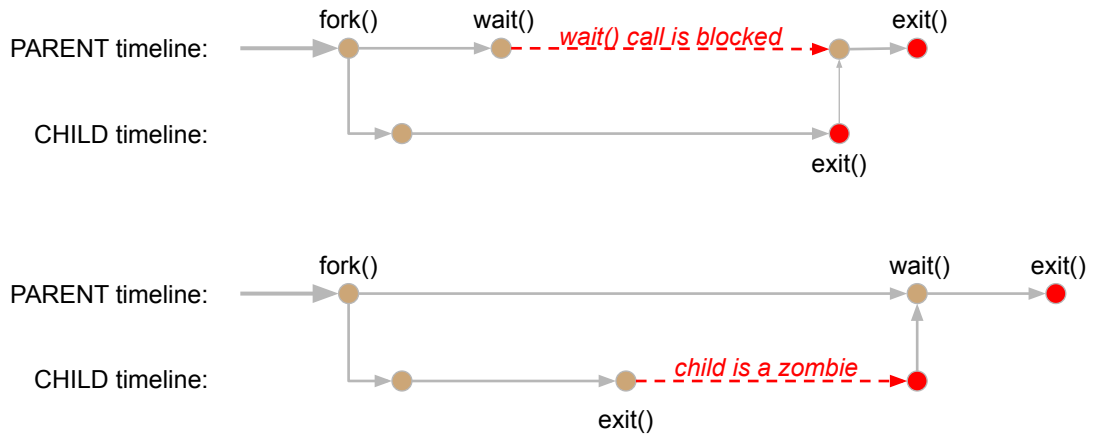
20

# Parent & Child Time lines



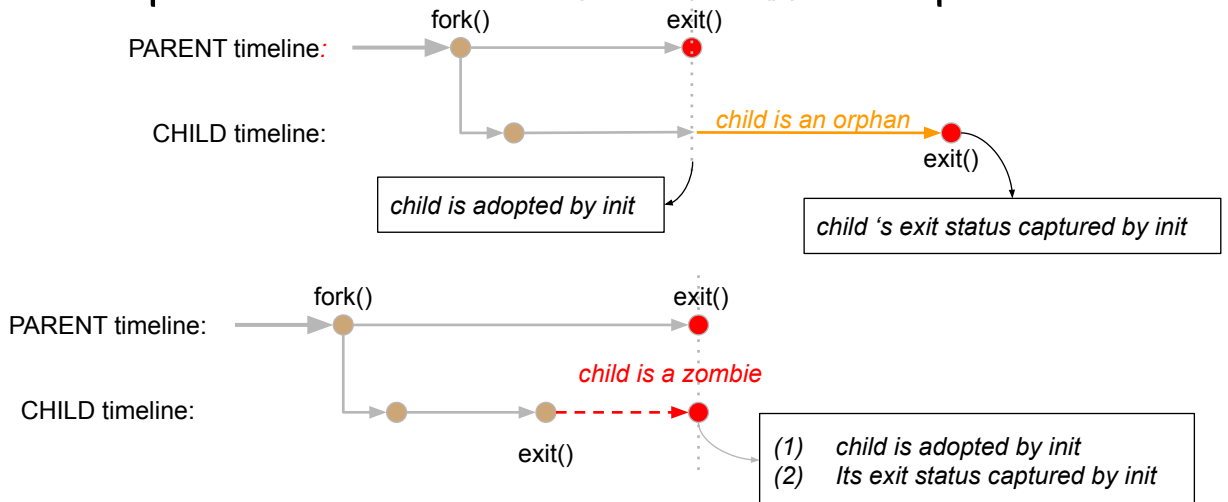
21

# Zombies (<defunct> processes)



22

# "Irresponsible" Parents (no wait()) & Orphans



23

# exec\*(): load an **external** executable

- **Replace** the current process image with a new binary executable
  - Continue running from the “main()” of the new executable
- The new binary executable does NOT entail a new process
  - The current process is the “home” of the new binary executable
- The current process image **stays intact** if the replacement executable cannot be loaded
  - Continue running from the “next” statement in the current process image

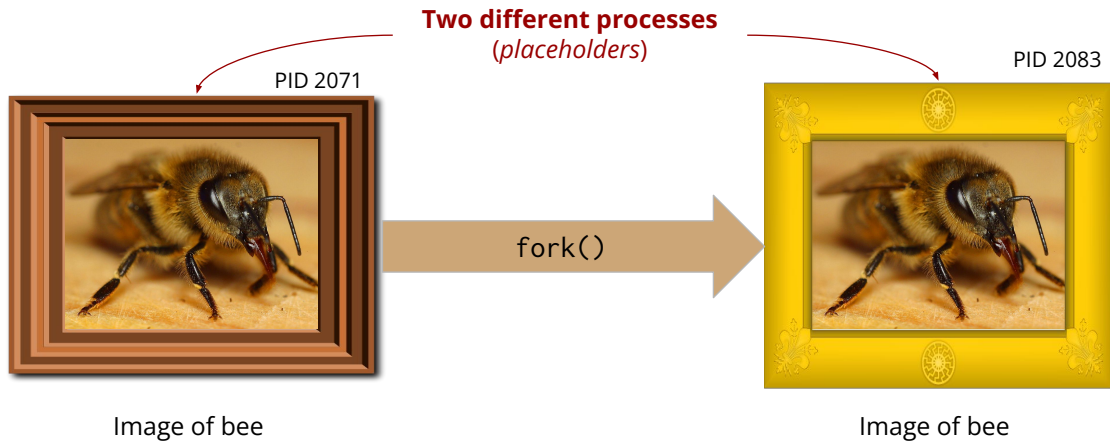
24

## Photo Album $\Rightarrow$ Frame + Image



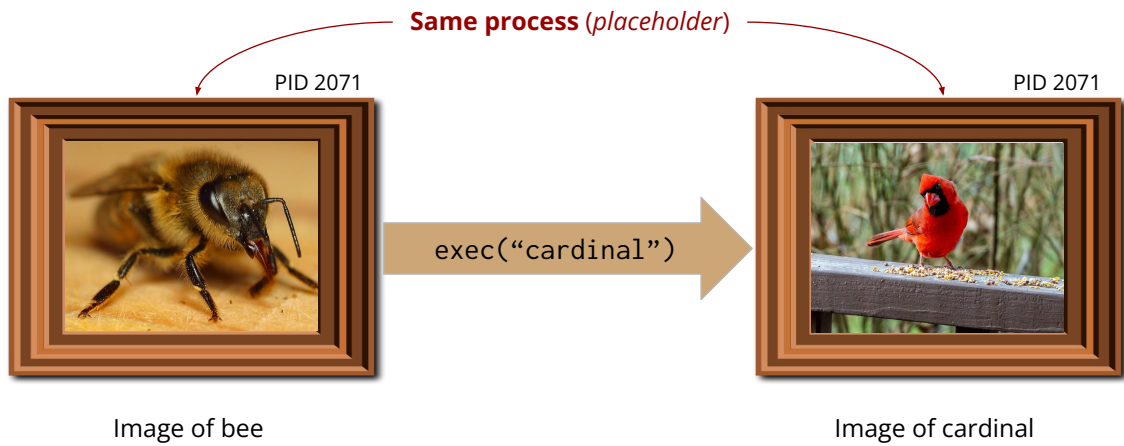
25

# fork()



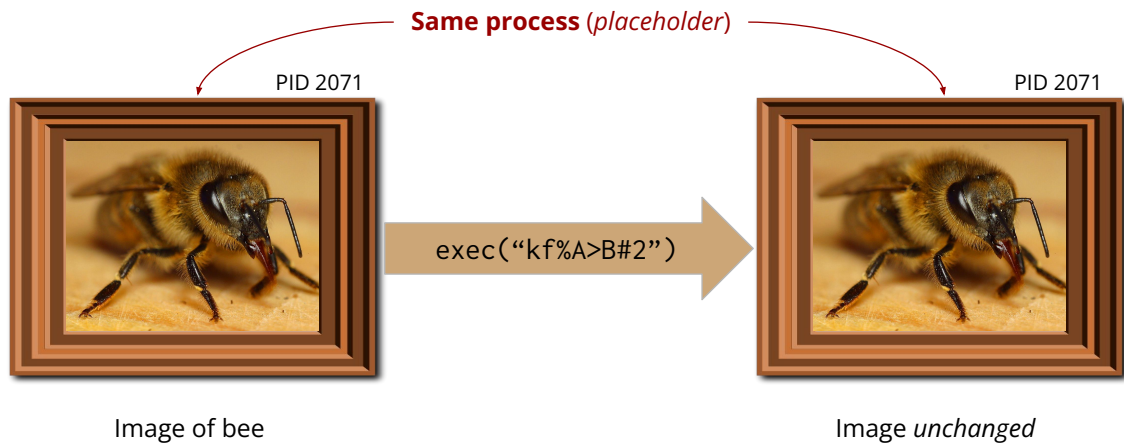
26

# exec()



27

# Exec() failure



28

## `exec*()` variants

- `execl()/exec1p()`: the "list" variant
  - (command line) arguments are supplied to the new binary executable using a list
- `execv()/execvp()`: the "vector" variant
  - (command line) arguments are supplied to the new binary executable using an array
- the "p" suffix: use the PATH environment variable to search for the new binary executable
- The FIRST argument to `exec*` is the **location** of the new binary executable
- The second (and remaining arguments) are arguments passed to the new binary executable

29

# exec\*() demo exec\_cal.c (EOS)

30

## Using exec\*()

```
int main() {  
    printf ("Begin\n");  
    execl ("/usr/bin/cal", "Venus", NULL);  
    printf ("End\n");  
    while (1) {}  
    return 0;  
}
```

when "/usr/bin/cal" **replaces** the process image

- "End" will never get printed!
- the infinite while-loop will never run

```
int main() {  
    printf ("Begin\n");  
    execl ("/usr/bin/california", "Venus", NULL);  
    printf ("End\n");  
    return 0;  
}
```

"/usr/bin/california" was not found  
the current process image was NOT replaced  
**"End" will be printed!**

31

# exec1() vs. execv(): passing arguments

```
// myprog.c
void main(int argc, char*argv[])
{
    for (int k = 0; k < argc; k++)
        printf("Arg-%d: |%s|\n", argv[k]);
}
```

```
exec1 ("/path/to/myprog", "23", NULL);
Arg-0: |23|
```

*argc: 1*

```
exec1 ("/path/to/myprog", "23", "and me", NULL);
Arg-0: |23|
Arg-1: |and me|
```

*argc: 2*

```
char* args[] = {"23", NULL};
execv ("/path/to/myprog", args);
Arg-0: |23|
```

*argc: 1*

```
char* args[] = {"23", "and me", NULL};
execv ("/path/to/myprog", args);
Arg-0: |23|
Arg-1: |and me|
```

*argc: 2*

32

```
ls -l -a -R
```

```
exec1("/usr/bin/ls", "GVSU", "-l", "-a", "-R", NULL);
```

```
execlp("ls", "GVSU", "-l", "-a", "-R", NULL);
```

```
char* arr[] = {"-l", "-a", "-R", NULL};
execv("/usr/bin/ls", arr);
```

```
char* arr[] = {"-l", "-a", "-R", NULL};
execvp("ls", arr);
```

33



# Time To Disclose the Truth

- **False statement:** Parent and child processes DO NOT share any contents of code/data/heap/stack
- Facts
  - Code section of both parent and child are exact copy of each other (**shareable**)
  - Heap section are very likely to be exact copy of each other (**shareable**)
  - Data (or stack) section may differ in a few bytes (**when the return value of fork() is saved to a global (or local) variable**)

34

# COW (Copy-on-Write)

- Code section (always read-only) can be shared
- Data, Heap, and Stack sections of parent and child can be shared IF these sections are used in **read-only** fashion
- If either parent/child **attempts to modify/write** data, heap, or stack, must be allocated **its own copy** (the OS does it for you)
- *Class discussion:* how to enforce COW?

35

# Program Execution

```
int main() { // Program A
    while (1);
    return 0;
}
```

- Is program A **running on the CPU** at all time?
  - Why?
  - Why not?

```
int main() { // Program B
    int num;
    while (1) {
        printf ("Number? ");
        scanf ("%d", &num);
        printf ("[%d]\n", num);
    }
    return 0;
}
```

- Is program B **running on the CPU** at all time?
  - Why?
  - Why not?
- How is it different from the execution of Program A?

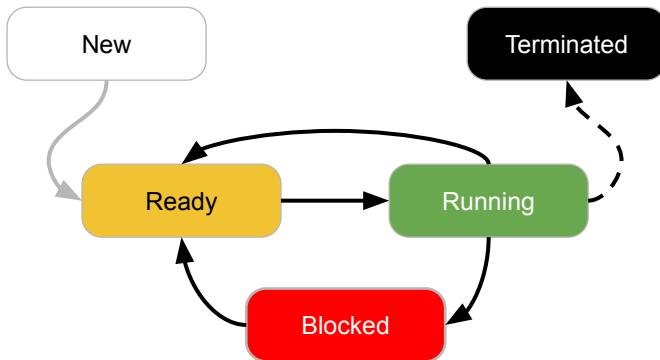
38

Sharing (Virtualize) the CPU?

Create N virtual CPUs out of ONE physical CPU?

39

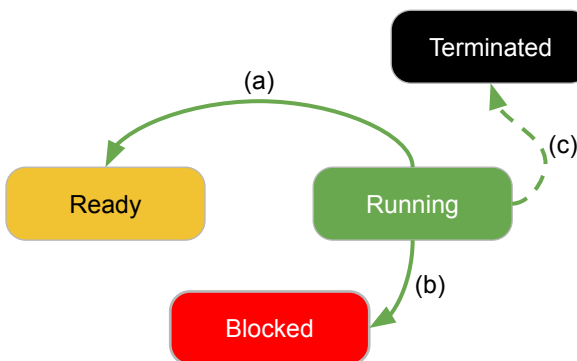
# Process State Transition Diagram



- New → Rdy: the process just created, ready to use the CPU
- Rdy → Run: dispatched by OS to use the CPU
- Run → Term: the process exits
- Run → Rdy: process time slice expired
- Run → Blk: the process issued a blocking system call (read(), sleep(), wait()) for child, wait for mouse click, ....
- Blk → Rdy: the blocking syscall completed, the process is ready to use the CPU again

40

# Process “Eviction” from CPU



- The current process is evicted from CPU
  - The CPU becomes vacant
- What triggers the eviction?
- What is the nature of eviction?
  - Abrupt or gradual?
- For (a) and (b), what needs to be done to **properly resume** an evicted process?

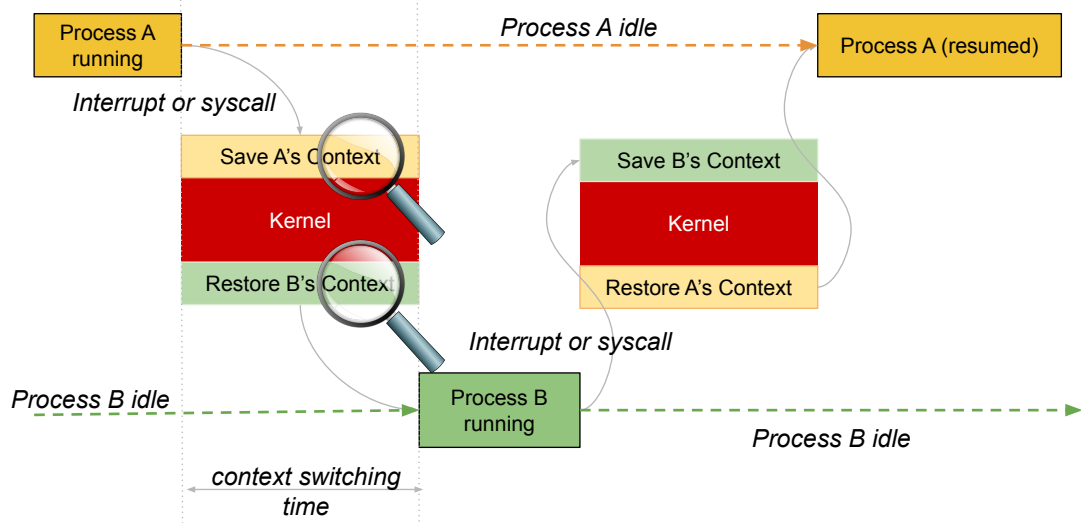
41

# How to share the CPU?

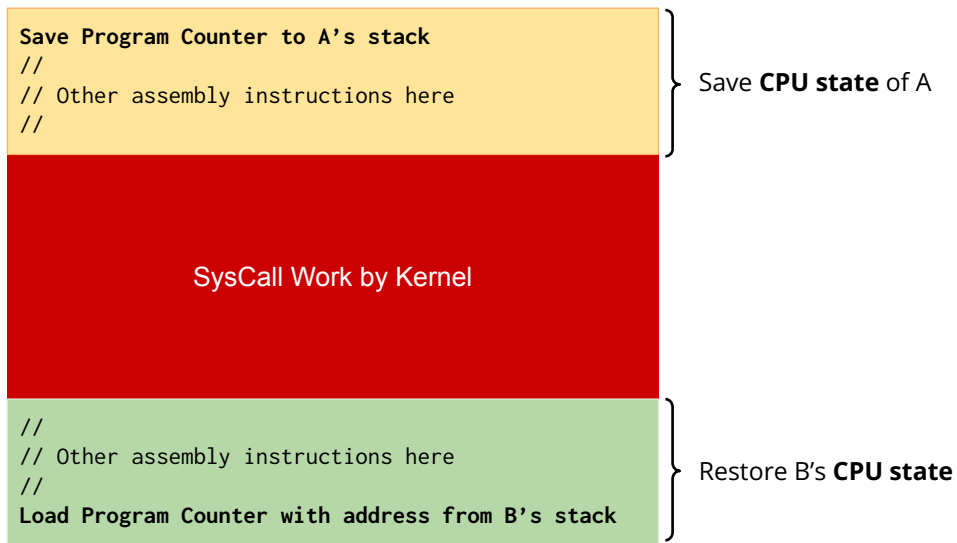
- When a program is **running**, its data can be found in
  - RAM (data section, heap section, stack section)
  - and also in \_\_\_\_\_
- What are potential problems in sharing the CPU?
- Solution?
- How do we share (and avoid conflict):
  - A classroom in a campus building?
  - A CPU?

42

## Context Switching



43



44

## Process Control Block (PCB)

- PCB is an OS data structure for storing important information of a process
  - Current context: CPU registers, Program Counter, control registers, ...
  - Scheduling-related information: priority, wait time, total CPU time, ....
  - Memory-related info: total memory, segment size, page tables,
  - I/O related info
  - Accounting info: CPU time used, memory used (real mem, virtual mem), ...

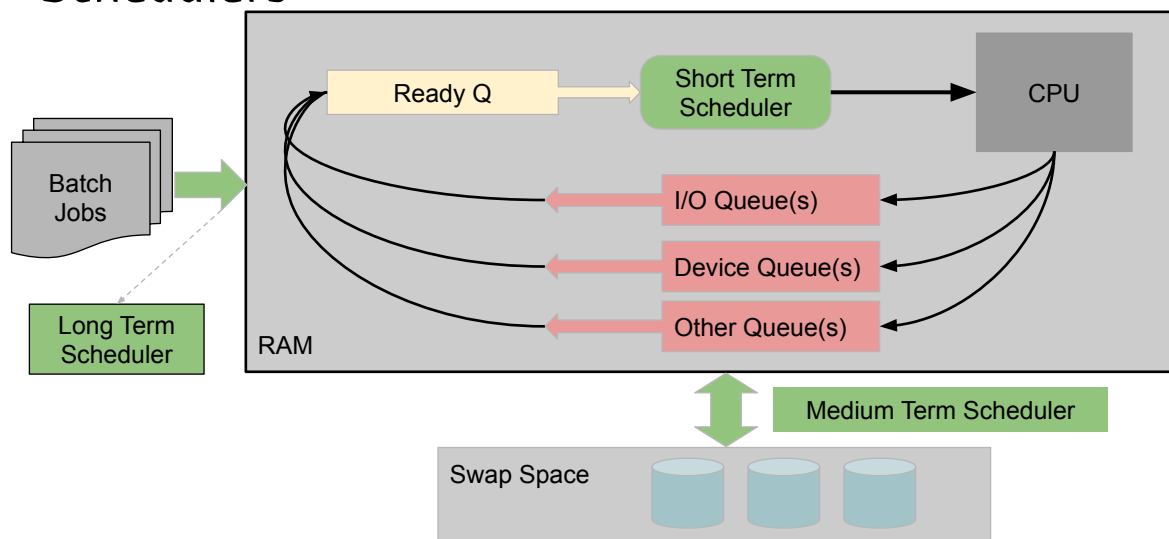
45

# Process Mgmt ↔ Processor Mgmt

- Process(or) Scheduler: dispatch a process to the CPU from the ready Q
- Queues for managing processes
  - Ready Queue: keeps all the processes who are in Ready state
  - Device Queue(s) or I/O Queue(s): keeps all the processes who are in Blocked state, waiting for I/O completion to/from a particular device
- Schedulers
  - Short-Term or CPU Scheduler: dispatches a process from the Ready Q to use the CPU
  - Medium-Term Scheduler: decides which process gets swapped-in / swapped-out (between RAM and swap disk)
  - Long-Term Scheduler (in a batch system): dispatch batch jobs to enter the system

46

## Schedulers



47