CS 344: Operating Systems I
01.25: Scheduling (101)

M/W 12:00 – 1:50 PM (LINC #200)

Sanghyun Hong
sanghyun.hong@oregonstate.edu

SAIL
Secure AI Systems Lab

Notice

- Announcement
  - Chat GPT (#random channel)
  - Misc.: we are here to help!

Notice

- Deadlines
  - ~~(1/23 11:59 PM) Programming assignment 1~~
  - (1/30 11:59 PM) Midterm quiz 1
  - (2/06 11:59 PM) Programming assignment 2

Recap

- Part I: Threads
  - Provide abstraction
    - What is a thread?
    - How is it different from a process?
    - How does OS run threads?
  - Offer standard libraries
    - How do we create/run/kill a thread?
    - How does OS manage the thread(s) we ran?
  - Manage resources
    - (Note) We will talk about this in the “scheduling” and “synchronization” classes

Topics for today

- Part I: Scheduling
  - Provide abstraction
    - What is scheduling?
    - What does OS achieve by scheduling?
  - Offer standard libraries
    - (Note) We will talk about this more in the “synchronization” lecture
  - Manage resources
    - What happens during scheduling?
    - How OS performs scheduling?
    - How OS implements this scheduling?

Problem: multiple program, yet limited processors

- Your Chrome browser:
  - Open multiple websites (tabs)
    - Tab 1: Open Canvas website
    - Tab 2: Stack Overflow
    - Tab 3: Discord website
    - … (many more 10+)
  - 4-8 CPUs (Processors)

How Can OS Address This Problem?

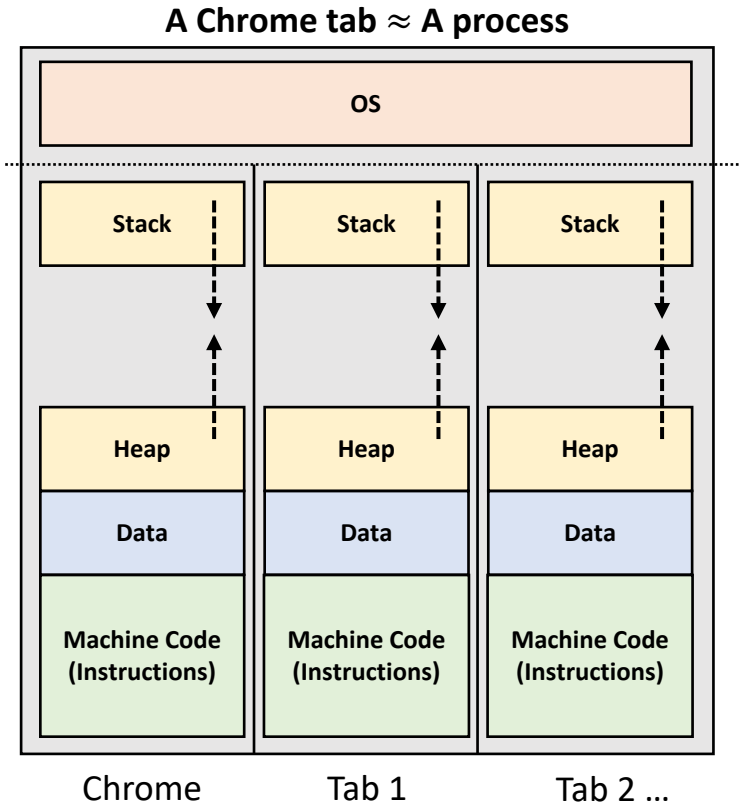

Provide abstraction: scheduling

- (Process/thread) scheduling:
  - **Definition:** the action of assigning resources to perform tasks
  - **Example:** your Chrome browser
    - An OS assigns each tab (process) to one of the processors
    - The OS takes over the processor and assigns to another process
    - … (continues)

Provide abstraction: scheduling – cont'd

- Goal:
  - Generate illusion
  - **Illusion:**
    - Make you feel that you're running 100+ processes at the same time
    - But in truth, it's not

Topics for today

- Part I: Scheduling
  - Provide abstraction
    - What is scheduling?
    - What does OS achieve by scheduling?
  - Offer standard libraries
    - (Note) We will talk about this more in the “synchronization” lecture
  - Manage resources
    - What happens during scheduling?
    - How OS performs scheduling?
    - How OS implements this scheduling?

Preliminaries on terminology

- **Definitions:**
  - Multiprogramming vs. multi-processing vs. multi-threading
    - Multi-programming: multiple jobs (or processes)
    - Multi-processing: multiple processors (CPUs)
    - Multi-threading: multiple threads

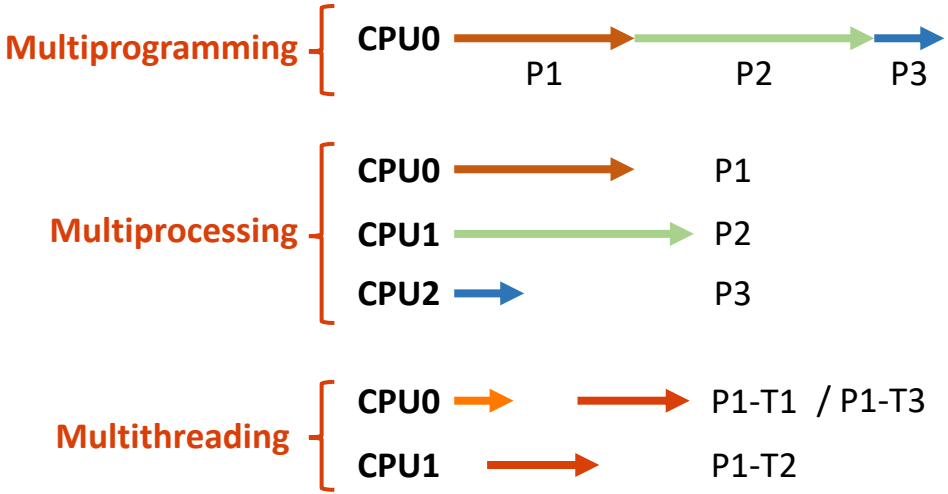

Manage resources: what happens during scheduling?

- **Context switch**
  - **Definition:** OS stores the current process's status and loads the new process's one
  - **Informal:** OS takes a CPU from one process and gives it to another

Manage resources: what happens during scheduling?

- **Context switch**
  - **Definition:** OS stores the current process's status and loads the new process's one
  - **Informal:** OS takes a CPU from one process and gives it to another

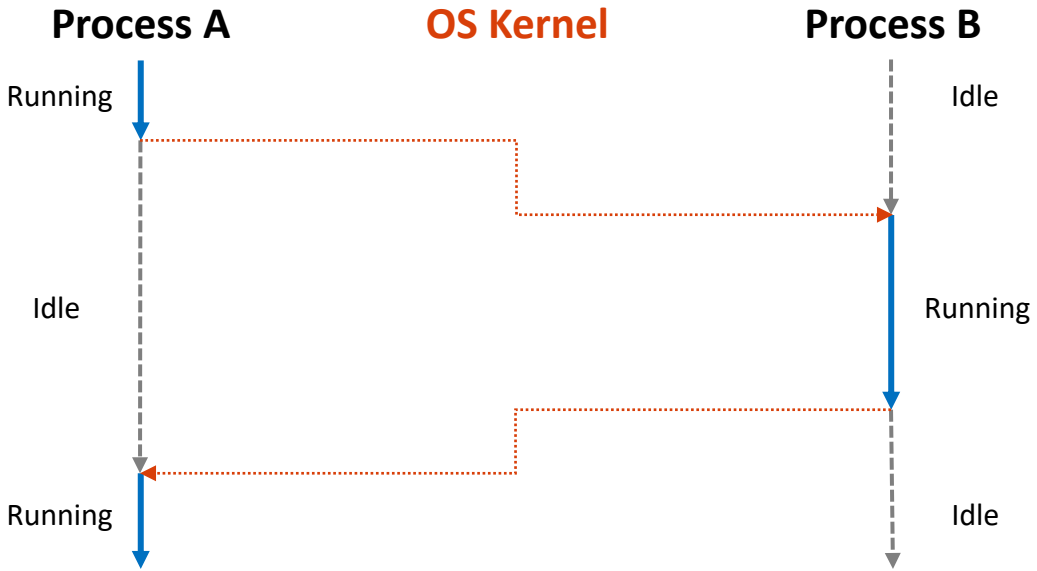

Manage resources: what happens during scheduling?

- **Context switch**
  - **Definition:** OS stores the current process's status and loads the new process's one
  - **Informal:** OS takes a CPU from one process and gives it to another

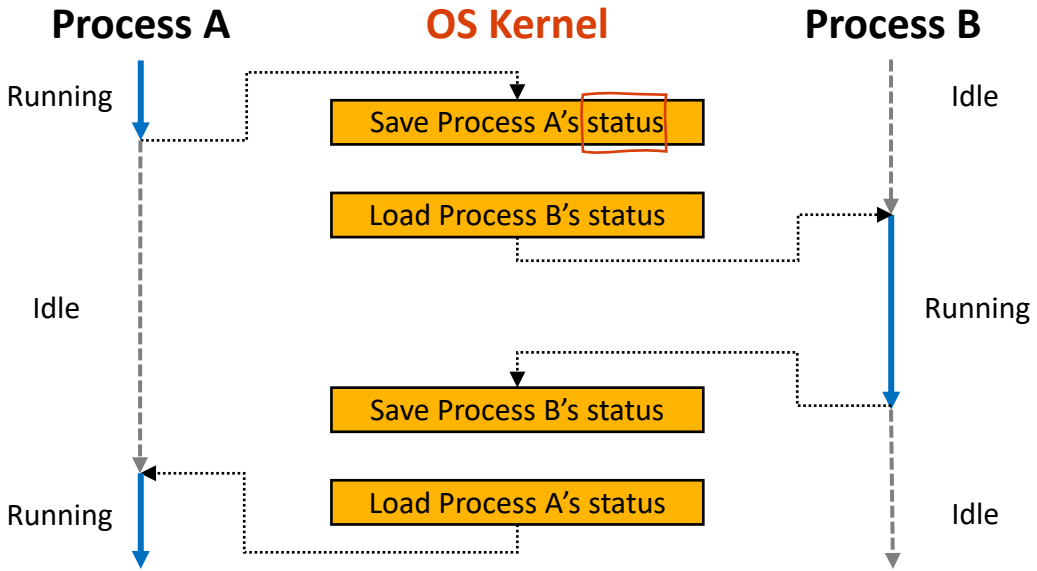

Manage resources: what happens during scheduling?

- **Context switch**
  - **Definition:** OS stores the current process's status and loads the new process's one
  - **Informal:** OS takes a CPU from one process and gives it to another

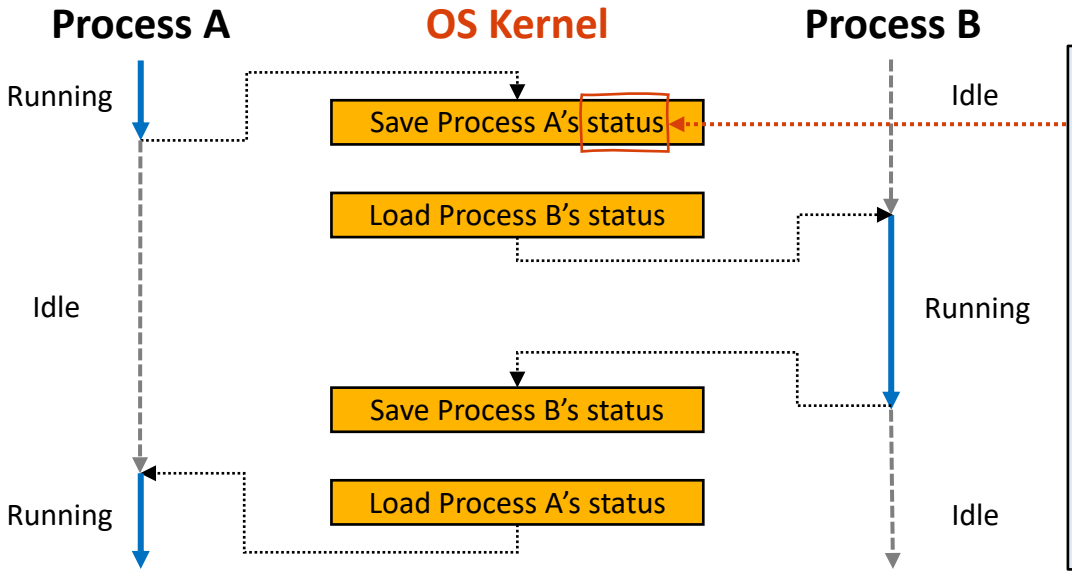

Recall: Process control block

A structure in OS that contains a set of information required to run a process on a CPU. Recall that Linux has *task_struct*.

- CPU#
- Program counter
- Instruction pointer
- Heap/stack pointer
- Process state [!]
- …

Revisit: process context

- (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;
```

```
852         struct sched_info               sched_info;
853
854         struct list_head                tasks;
855 #ifdef CONFIG_SMP
856         struct plist_node               pushable_tasks;
857         struct rb_node                  pushable_dl_tasks;
858 #endif
859
860         struct mm_struct                *mm;
861         struct mm_struct                *active_mm;
862
863         /* Per-thread vma caching: */
864         struct vmacache                 vmacache;
865
866 #ifdef SPLIT_RSS_COUNTING
867         struct task_rss_stat            rss_stat;
868 #endif
869         int                             exit_state;
870         int                             exit_code;
871         int                             exit_signal;
872         /* The signal sent when the parent dies: */
873         int                             pdeath_signal;
874         /* JOBCTL_*, siglock protected: */
875         unsigned long                   jobctl;
876
877         /* Used for emulating ABI behavior of previous Linux versions: */
878         unsigned int                    personality;
```

Manage resources: what happens during scheduling?

- **Context switch**
  - **Definition:** OS stores the current process's status and loads the new process's one
  - **Informal:** OS takes a CPU from one process and gives it to another

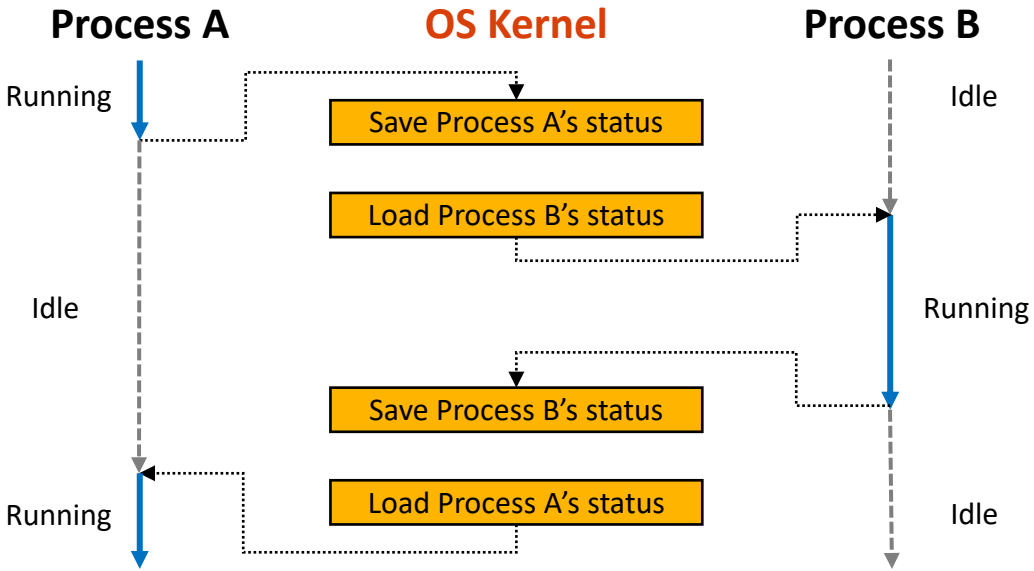

Revisit: Process Context to a CPU

- (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)

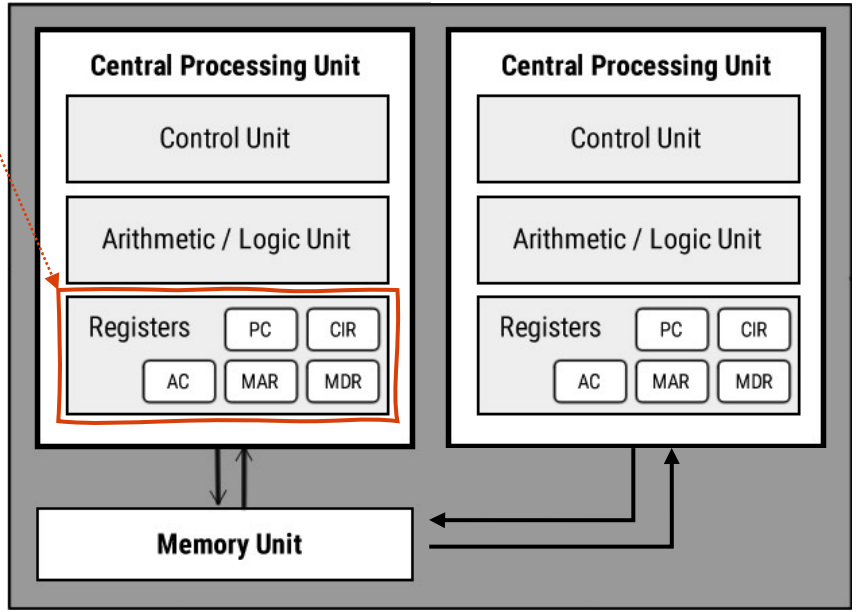

Manage resources: what happens during scheduling?

- **Context switch**
  - **Definition:** OS stores the current process's status and loads the new process's one
  - **Informal:** OS takes a CPU from one process and gives it to another

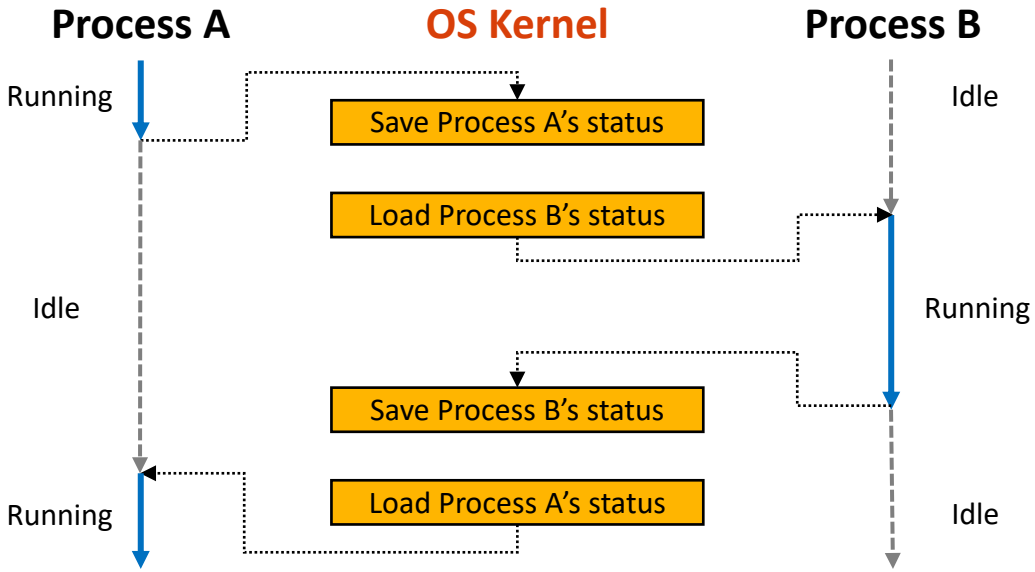

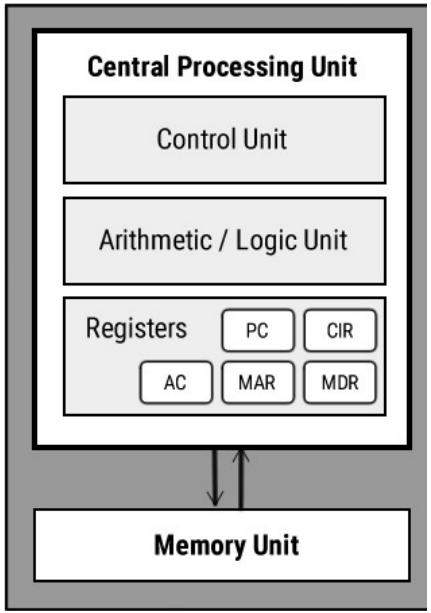

Manage resources: what happens during scheduling?

- **Context switch**
  - **Definition:** OS stores the current process's status and loads the new process's one
  - **Informal:** OS takes a CPU from one process and gives it to another

- No free lunch
  - Context switching takes $\sim 5\ \mu s$ on average
  - OS typically runs 100+ processes
  - Too many context switching makes a system unable to respond…

Manage resources: what happens during scheduling?

- A process can have **five states**:
  - **New:** a process (or thread) is being created (by fork())
  - **Ready:** the process is waiting to run
  - **Running:** the process is running on a CPU(or CPUs)
  - **Waiting:** the process is waiting for some events to occur (*e.g.*, a data loaded from storage)
  - **Terminated:** the process has finished execution; waiting for removal

Manage resources: what happens during scheduling?

- A process can have **five states**:
  - **New:** a process (or thread) is being created (by fork())
  - **Ready:** the process is waiting to run
  - **Running:** the process is running on a CPU(or CPUs)
  - **Waiting:** the process is waiting for some events to occur (*e.g.*, a data loaded from storage)
  - **Terminated:** the process has finished execution; waiting for removal
- State transition (life cycle):

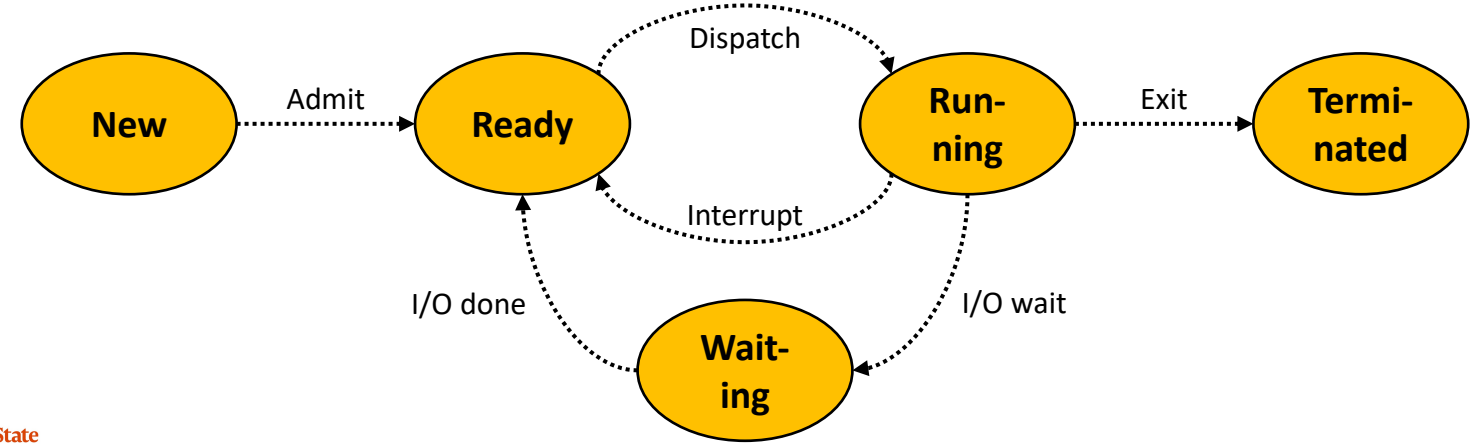

Manage resources: what happens during scheduling?

- A process can have **five states**:
  - **New:** a process (or thread) is being created (by fork())
  - **Ready:** the process is waiting to run
  - **Running:** the process is running on a CPU(or CPUs)
  - **Waiting:** the process is waiting for some events to occur (*e.g.*, a data loaded from storage)
  - **Terminated:** the process has finished execution; waiting for removal
- State transition (life cycle):

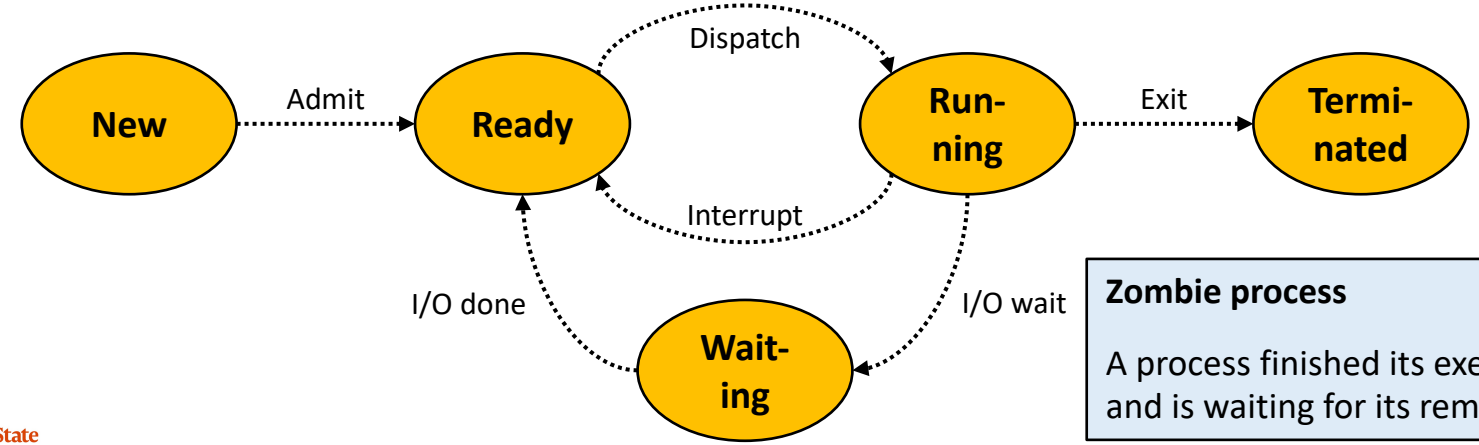

Zombie process

A process finished its execution and is waiting for its removal

Topics for today

- Part I: Scheduling
  - Provide abstraction
    - What is scheduling?
    - What does OS achieve by scheduling?
  - Offer standard libraries
    - (Note) We will talk about this more in the “synchronization” lecture
  - Manage resources
    - What happens during scheduling?
    - How OS performs scheduling?
    - How OS implements this scheduling?

Manage resources: how OS performs scheduling?

- **Scheduling**
  - **Definition:** an OS activity that schedules processes in different states
  - **Note:** OS implements queues to hold multiple processes in the same state

- **Illustration (single CPU)**

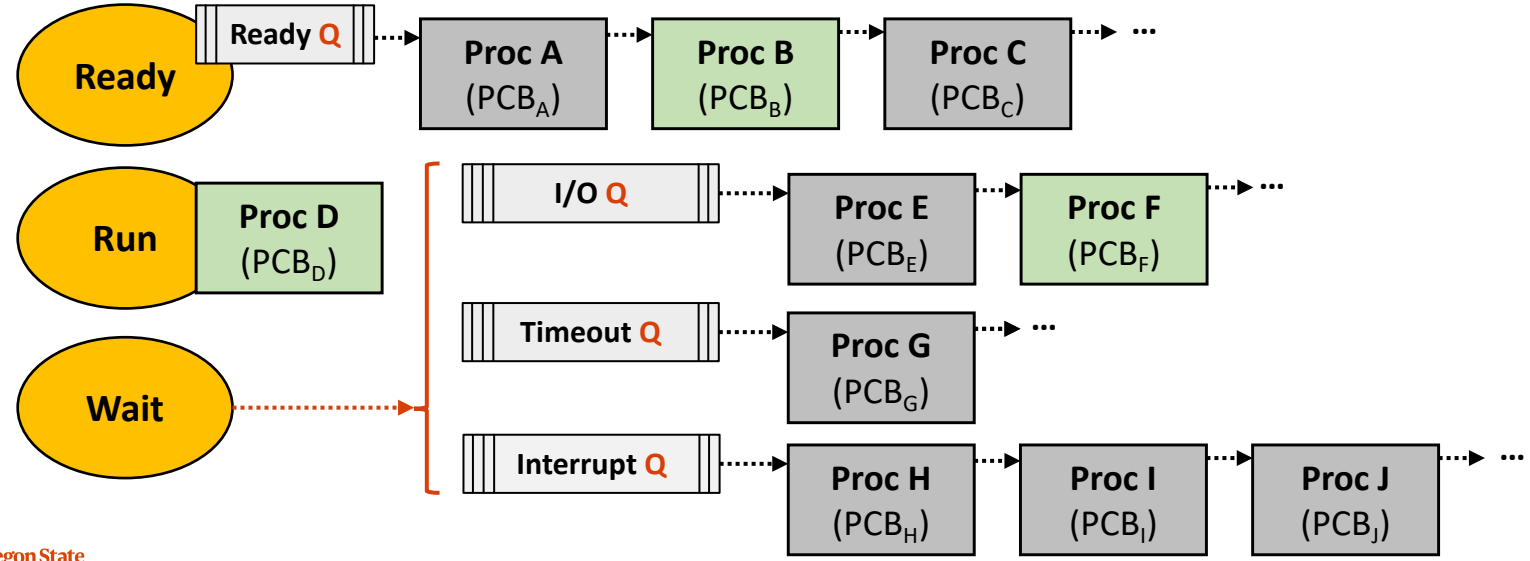

Manage resources: how OS performs scheduling?

- **Scheduling**
  - **Definition:** an OS activity that schedules processes in different st
  - **Note:** OS implements queues to hold multiple processes in the sa

Illustrated Example

1. Kicks out Proc D (timeout)
2. Runs Proc B
3. Puts Proc F in the ready Q (I/O has done, in this case)

- **Illustration (single CPU)**

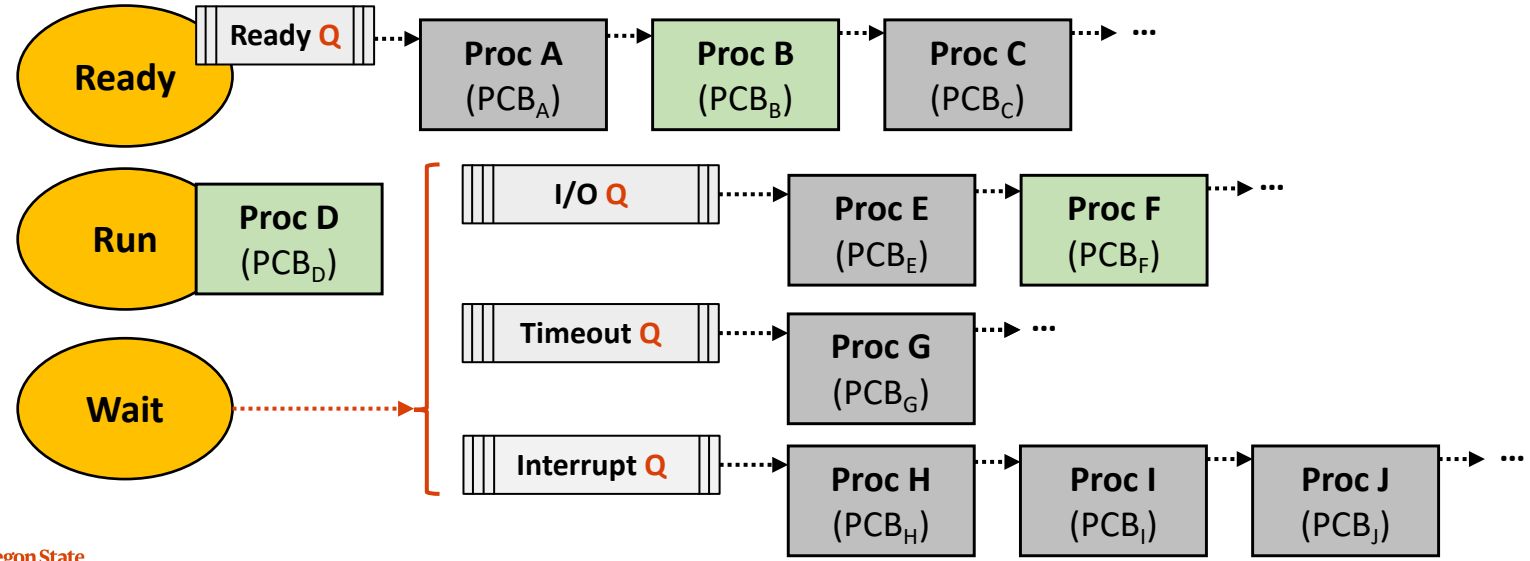

Scheduling example: high-level view

- 3 Processes in Chrome:
  - **P1:** Download movies
  - **P2:** Open Canvas
  - **P3:** Search StackOverflow
- Example
  - **New** : P1
  - **Ready**:
  - **Run** :
  - **Wait** :
  - **Term..**:

  - **Scenario:** open a website for downloading movies

Scheduling example: high-level view

- 3 Processes in Chrome:
  - **P1:** Download movies
  - **P2:** Open Canvas
  - **P3:** Search StackOverflow
- Example
  - **New** :
  - **Ready**:
  - **Run** :
  - **Wait** :
  - **Term..**:

  - **Scenario:** the website opened and open two other websites

Scheduling example: high-level view

- 3 Processes in Chrome:
  - **P1:** Download movies
  - **P2:** Open Canvas
  - **P3:** Search StackOverflow
- Example
  - **New** :
  - **Ready**:
  - **Run** :
  - **Wait** :
  - **Term..**:

  - **Scenario:** downloads started, and you focus on Canvas

Scheduling example: high-level view

- 3 Processes in Chrome:
  - **P1:** Download movies
  - **P2:** Open Canvas
  - **P3:** Search StackOverflow
- Example
  - **New** :
  - **Ready**:
  - **Run** :
  - **Wait** :
  - **Term..**:

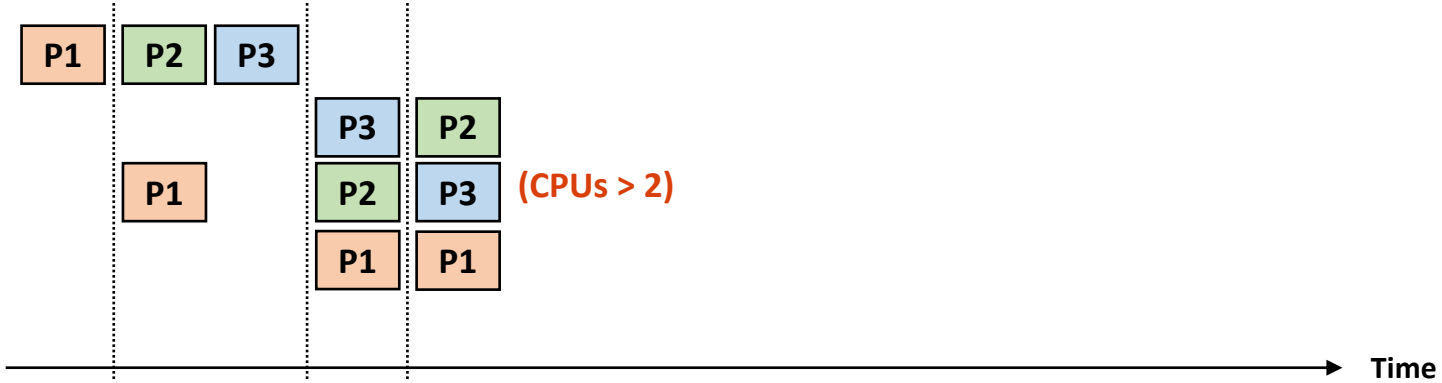

  - **Scenario:** while downloading, you start searching StackOverflow

Scheduling example: high-level view

- 3 Processes in Chrome:
  - **P1:** Download movies
  - **P2:** Open Canvas
  - **P3:** Search StackOverflow
- Example
  - **New** :
  - **Ready**:
  - **Run** :
  - **Wait** :
  - **Term..**:

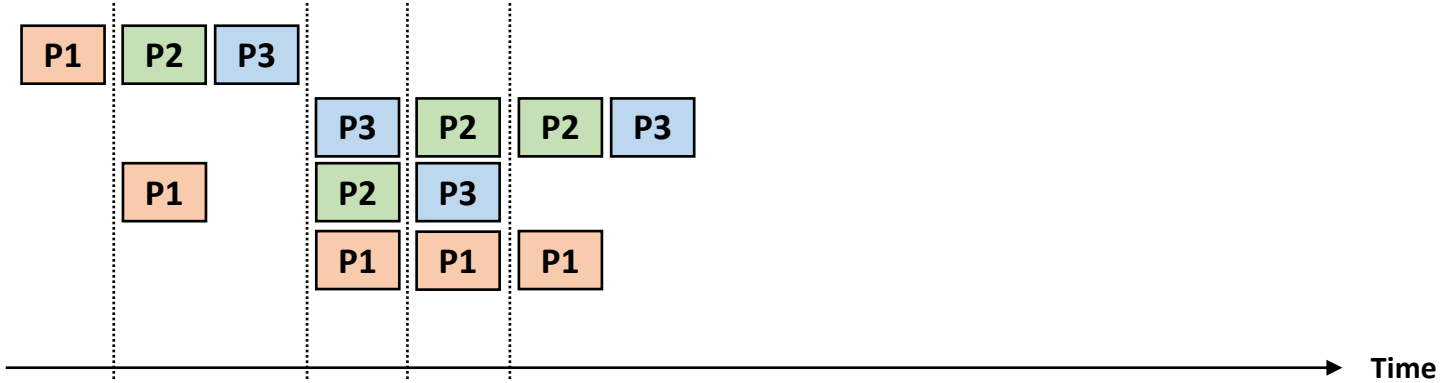

  - **Scenario:** downloading movies are done

Scheduling example: high-level view

- 3 Processes in Chrome:
  - **P1:** Download movies
  - **P2:** Open Canvas
  - **P3:** Search StackOverflow
- Example
  - **New** :
  - **Ready**:
  - **Run** :
  - **Wait** :
  - **Term..**:

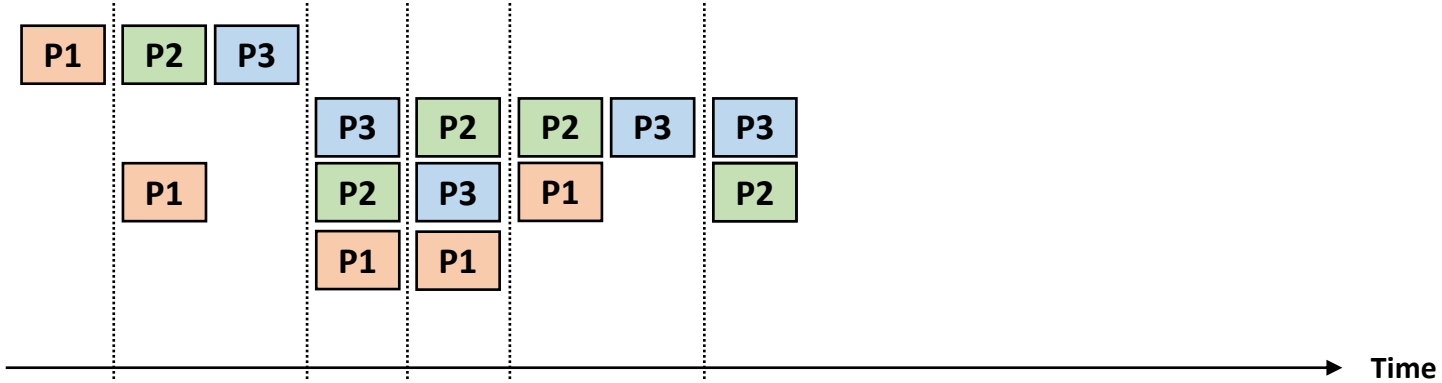

  - **Scenario:** close the download tab, and keep looking at Canvas

Scheduling example: high-level view

- 3 Processes in Chrome:
  - **P1:** Download movies
  - **P2:** Open Canvas
  - **P3:** Search StackOverflow
- Example
  - **New** :
  - **Ready**:
  - **Run** :
  - **Wait** :
  - **Term..**:
  - **Scenario:** close the other two tabs to go to bed

Linked list: OS keeps a list of processes for each state by using the queue structure. Based on a policy, OS searches the queue and find the process with the highest priority

Zombie process(es): A list of processes whose job has already terminated (exited) but has not removed from the process list. One case is that a parent process wait for a child's exit status.

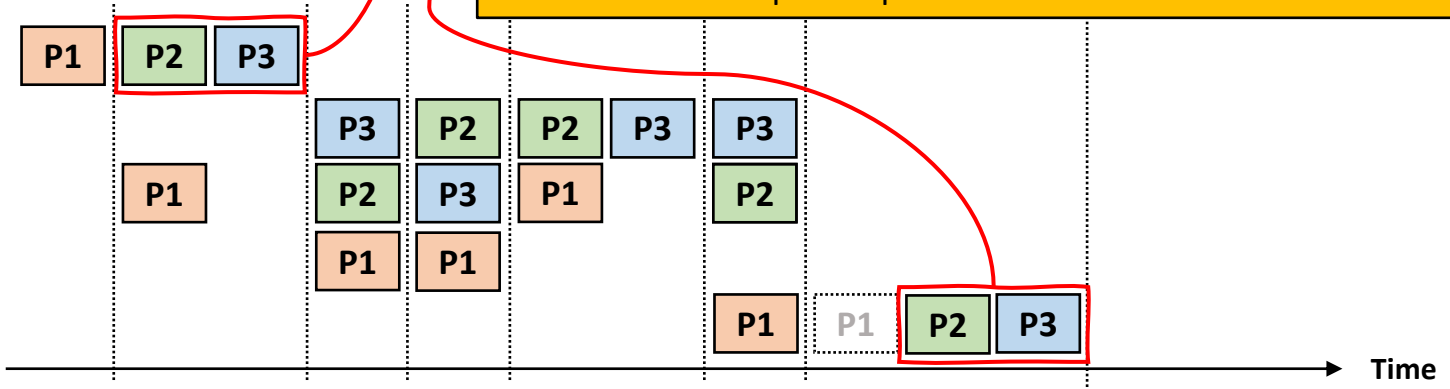

Topics for today

- Part I: Scheduling
  - Provide abstraction
    - What is scheduling?
    - What does OS achieve by scheduling?
  - Offer standard libraries
    - (Note) We will talk about this more in the “synchronization” lecture
  - Manage resources
    - What happens during scheduling?
    - How OS performs scheduling?
    - How OS implements this scheduling?

Manage resources: how OS implements scheduling?

- **(OS) Scheduler:**
  - **Definition:** An OS task (process) that manages the process scheduling activity

Manage resources: how OS implements scheduling?

- **(OS) Scheduler:**
  - **Definition:** An OS task (process) that manages the process scheduling activity

- **Implementation**

```
while ( <some condition,
          but eventually will be infinite>) {

   RunProcess( curProc );
   newProc = chooseNextProc();
   saveCurrentProc( curProc );
   LoadNextState( newProc );

}
```

  - It is also a process (an *infinite* loop)
  - The scheduler process terminates if we *stop* (turn-off) a computer
  - Example mechanisms that trigger scheduling, *e.g.*, yield and interrupt

Manage resources: how OS implements scheduling?

- **Problem:** how to choose a next process?
  - FIFO (first come, first served)
  - LIFO (last come, first served)
  - Shortest-job first (do a short-period job first)
  - Priority-based (do an important job first)
  - … (It's an open-problem; You'll learn in OS II)

Process Scheduling

- Multiple objectives:
  - **Fairness**: no monopoly
  - **Priority**: consider importance of a process
  - **Deadlines**: a process should be done before/by the time *T*
  - **Throughput**: maximize the number of tasks done
  - **Efficiency**: minimize the scheduling overheads
  - …

- **No silver bullet:**
  - Depends on the objectives (ex. NASA's Perseverance)
  - Oftentimes, the objectives are in conflicts (ex. Priority vs. Fairness)
  - … (You'll learn in OS II)

Topics for today

- Part I: Scheduling
  - Provide abstraction
    - What is scheduling?
    - What does OS achieve by scheduling?
  - Offer standard libraries
    - (Note) We will talk about this more in the “synchronization” lecture
  - Manage resources
    - What happens during scheduling?
    - How OS performs scheduling?
    - How OS implements this scheduling?

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