CS 1550: Scheduling

Jonathan Misurda
jmisurda@cs.pitt.edu
http://www.cs.pitt.edu/~jmisurda/

Multiprogramming

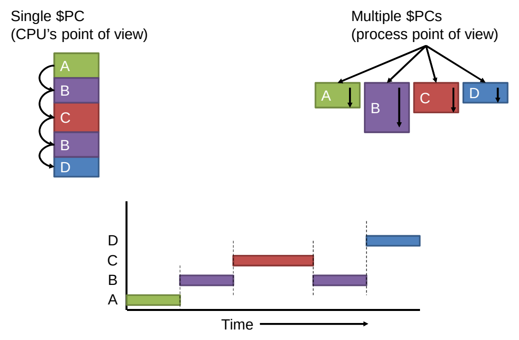

Process

A running program and its associated data

Life Cycle of a Process

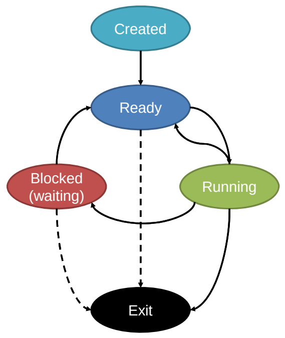

Process Table

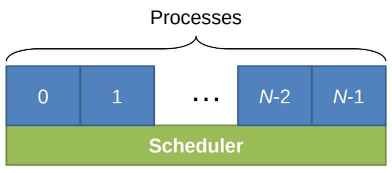

Process Table Entry

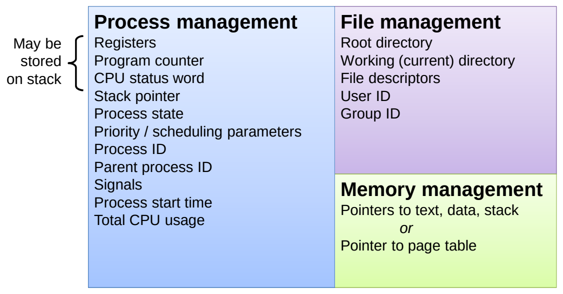

May be stored on stack

Process management	File management
Registers	Root directory
Program counter	Working (current) directory
CPU status word	File descriptors
Stack pointer	User ID
Process state	Group ID
Priority / scheduling parameters	
Process ID	**Memory management**
Parent process ID	Pointers to text, data, stack
Signals	*or*
Process start time	Pointer to page table
Total CPU usage	

Thread

A stream of instructions and their associated state

Processes and Threads

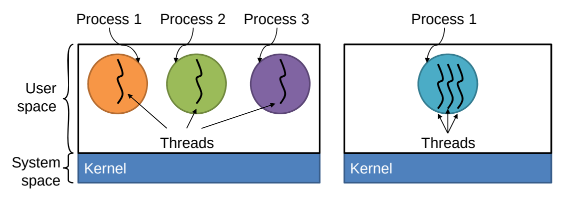

Thread State

Per process items
Address space
Open files
Child processes
Signals & handlers
Accounting info
Global variables

Per thread items	Per thread items	Per thread items
Program counter	Program counter	Program counter
Registers	Registers	Registers
Stack & stack pointer	Stack & stack pointer	Stack & stack pointer
State	State	State

Threading

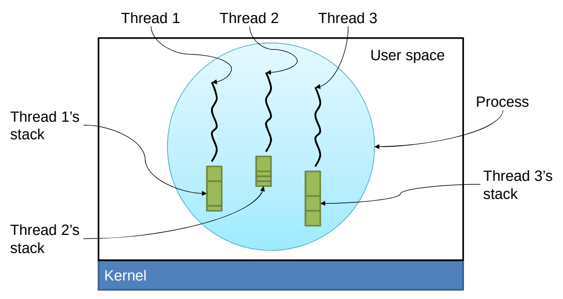

Multithreading in Action

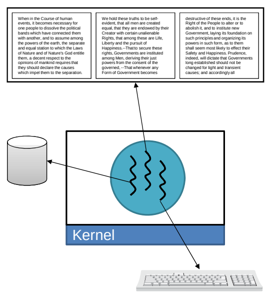

Multithreaded Webserver

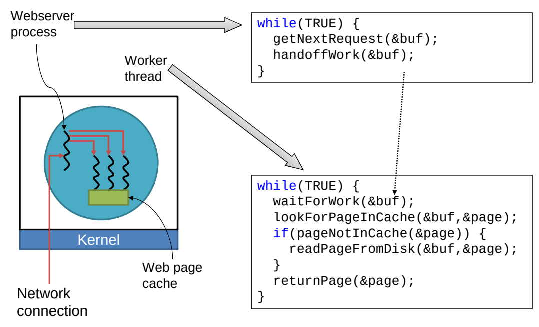

```
while(TRUE) {
  getNextRequest(&buf);
  handoffWork(&buf);
}
```

```
while(TRUE) {
  waitForWork(&buf);
  lookForPageInCache(&buf,&page);
  if(pageNotInCache(&page)) {
    readPageFromDisk(&buf,&page);
  }
  returnPage(&page);
}
```

User Threads vs. Kernel Threads

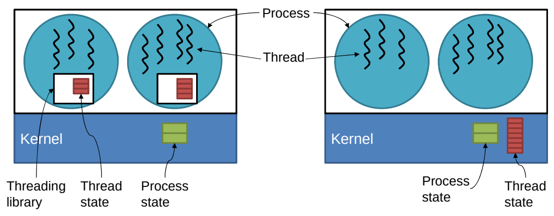

Scheduling

How to choose which of the Ready processes/threads gets to Run next

CPU Bound vs. I/O Bound

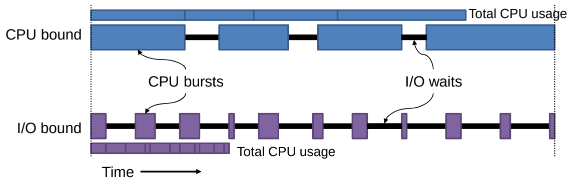

When to Schedule

- Process Creation
- Process Exit
- Blocked
- I/O Interrupt
- Clock Interrupts

Three-Level Scheduling

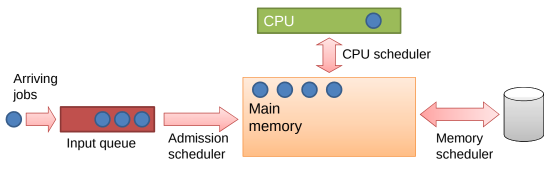

Fairness

Comparable processes get comparable service

Throughput

Number of jobs completed per unit time

Turnaround Time

Time from job submission to job completion

Average Turnaround Time

Average of all turnaround times for a set of jobs

Batch Scheduling

Non-interactive jobs that can be run "overnight"

First Come, First Served

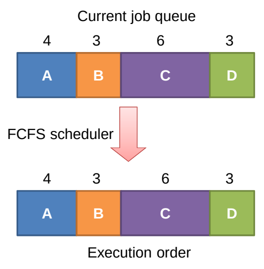

Shortest Job First (SJF)

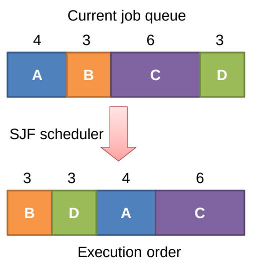

Interactive scheduling

Impatient users waiting

Round Robin Scheduling

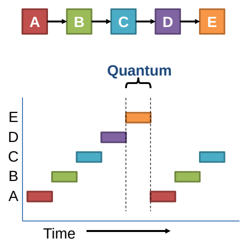

Priority Scheduling

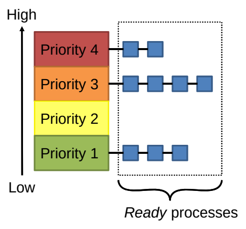

Other Scheduling Algorithms

- **Shortest Process Next**
  - *SJF applied to Interactive Systems*
- **Guaranteed Scheduling**
  - *N processes get 1/N of the CPU Time*
- **Lottery Scheduling**
  - *Give out tickets, pull one at random, winner runs*
- **Fair Share**
  - *N users get 1/N CPU time*

Earliest Deadline First (EDF)

Real-time: How you do homework

Mechanism

The way something is done (e.g., an algorithm)

Policy

The rules a particular mechanism should follow (i.e., the parameters of an algorithm)

Scheduling User Threads

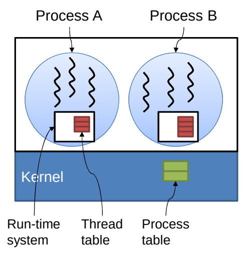

Scheduling Kernel Threads

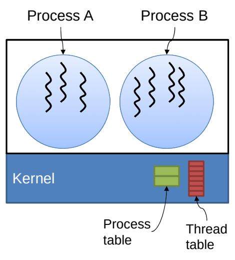
Process A
Process B
Kernel
Process table
Thread table