

# CPU Scheduling

CS 256/456

Dept. of Computer Science, University of Rochester

1/26/2005

CSC 256/456 - Spring 2005

1

## Recap of the Last Class

- Process
  - Process concept
  - A process's image in the computer
  - Operations on processes
- Thread
  - Thread concept
  - Compared with process
    - less context switch overhead
    - more efficient synchronization between threads
  - User/kernel threads

1/26/2005

CSC 256/456 - Spring 2005

2

## User/Kernel Threads

- Thread switching
  - save the registers (including the SP) of the old thread;
  - restore the registers of the new thread;
  - set the new PC appropriately for the new thread;
- Can be done without the help from the OS.
- User threads
  - Thread data structure is in user-mode memory
  - scheduling done at user mode
- Kernel threads
  - Thread data structure is in kernel memory
  - scheduling done by the OS kernel

1/26/2005

CSC 256/456 - Spring 2005

3

## User/Kernel Threads (cont.)

- Benefits of user threads
  - lightweight - less context switching overhead
  - more efficient synchronization??
  - flexibility - allow application-controlled scheduling
- Problems of user threads
  - can't user more than one processor
  - oblivious to kernel events, e.g., all threads in a process are put to wait when only one of them does I/O


1/26/2005

CSC 256/456 - Spring 2005

4

## Mixed User/Kernel Threads

- M user threads run on N kernel threads ( $M \geq N$ )
  - $N=1$ : pure user threads
  - $M=N$ : pure kernel threads
  - $M > N > 1$ : mixed model



1/26/2005

CSC 256/456 - Spring 2005

5

## Solaris/Linux Threads

- Solaris
  - supports mixed model
  - provides application interface to control the underlying kernel threads
- Linux
  - No standard user threads on Linux
  - Processes are treated similarly with threads (both called tasks)
  - Processes are tasks with exclusive address space
  - Tasks can also share the address space, open files, ...

1/26/2005

CSC 256/456 - Spring 2005

6

## Pthreads

- Different OS has its own thread package with different Application Programming Interfaces  $\Rightarrow$  **poor portability**.
- Pthreads
  - A POSIX standard API for thread management and synchronization.
  - API specifies behavior of the thread library, not the implementation.
  - Commonly supported in UNIX operating systems.

1/26/2005

CSC 256/456 - Spring 2005

7

## CPU Scheduling

- Selects from among the processes/threads that are ready to execute, and allocates the CPU to it.
- CPU scheduling may take place at:
  1. Hardware interrupt/software exception.
  2. System calls.
- *Nonpreemptive*:
  - Scheduling only when the current process terminates or not able to run further
- *Preemptive*:
  - Scheduling can occur at any opportunity possible

1/26/2005

CSC 256/456 - Spring 2005

8

## Scheduling Criteria

- Minimize turnaround time - amount of time to execute a particular process
- Maximize throughput - # of processes that complete their execution per time unit
- Maximize CPU utilization - the proportion of the CPU that is not idle
- Minimize response time - amount of time it takes from when a request was submitted until the first response is produced (interactivity)
- Fairness: avoid starvation

1/26/2005


CSC 256/456 - Spring 2005

9

## First-Come, First-Served (FCFS) Scheduling

Process	CPU Time
$P_1$	24
$P_2$	3
$P_3$	3

- Suppose that the processes arrive in the order:  $P_1, P_2, P_3$   
The schedule is:



- Turnaround time for  $P_1 = 24$ ;  $P_2 = 27$ ;  $P_3 = 30$
- Average turnaround time:  $(24 + 27 + 30)/3 = 27$

1/26/2005


CSC 256/456 - Spring 2005

10

## FCFS Scheduling (Cont.)

Suppose that the processes arrive in the order  $P_2, P_3, P_1$ .

- The schedule is:



- Turnaround time for  $P_1 = 30$ ;  $P_2 = 3$ ;  $P_3 = 6$
- Average turnaround time:  $(30 + 3 + 6)/3 = 13$
- Much better than previous case.
- Short process delayed by long process: *Convoy effect*

1/26/2005

CSC 256/456 - Spring 2005

11

## Shortest-Job-First (SJR) Scheduling

- Associate with each process the length of its CPU time. Use these lengths to schedule the process with the shortest CPU time.
- SJR is optimal - gives minimum average turnaround time for a given set of processes.
- Problem:
  - don't know the process CPU time ahead of time.
- Two variations:
  - nonpreemptive - once CPU given to the process it cannot be taken away until it completes.
  - preemptive - if a new process arrives with CPU time less than remaining time of current executing process, preempt.

1/26/2005


CSC 256/456 - Spring 2005

12

### Example of Non-Preemptive SJF

Process	Arrival Time	CPU Time
$P_1$	0.0	7
$P_2$	2.0	4
$P_3$	4.0	1
$P_4$	5.0	4

- SJF (non-preemptive)



- Average turnaround time =  $(7 + 10 + 4 + 11)/4 = 8$

1/26/2005


CSC 256/456 - Spring 2005

13

### Example of Preemptive SJF

Process	Arrival Time	CPU Time
$P_1$	0.0	7
$P_2$	2.0	4
$P_3$	4.0	1
$P_4$	5.0	4

- SJF (preemptive)



- Average turnaround time =  $(16 + 5 + 1 + 6)/4 = 7$

1/26/2005

CSC 256/456 - Spring 2005

14

### Priority Scheduling

- A priority number (integer) is associated with each process
- The CPU is allocated to the process with the highest priority.
  - Preemptive
  - nonpreemptive
- SJF is a priority scheduling where priority is the predicted CPU time.
- Problem: **Starvation** - low priority processes may never execute.
- Solution: **Aging** - as time progresses increase the priority of the process.

1/26/2005

CSC 256/456 - Spring 2005

15

### Round Robin (RR)

- Each process gets a fixed unit of CPU time (*time quantum*), usually 10-100 milliseconds. After this time has elapsed, the process is preempted and added to the end of the ready queue.
- If there are  $n$  processes in the ready queue and the time quantum is  $q$ , then each process gets  $1/n$  of the CPU time in chunks of at most  $q$  time units at once. No process waits more than  $(n-1)q$  time units.
- Performance
  - $q$  small  $\Rightarrow$  fair, starvation-free, better interactivity
  - $q$  large  $\Rightarrow$  FIFO
  - $q$  must be large with respect to context switch cost, otherwise overhead is too high.

1/26/2005

CSC 256/456 - Spring 2005

16

## Cost of Context Switch

- Direct overhead of context switch
  - saving old contexts, restoring new contexts, ... ..
- Indirect overhead of context switch
  - caching, memory management overhead

1/26/2005


CSC 256/456 - Spring 2005

17

## Example of RR with Quantum = 20

Process	CPU Time
$P_1$	53
$P_2$	17
$P_3$	68
$P_4$	24

- The schedule is:



- Typically, higher average turnaround than SJF, but better *response*.

1/26/2005

CSC 256/456 - Spring 2005

18

## Multilevel Queue

- Ready queue is partitioned into separate queues:
  - foreground (interactive)
  - background (batch)
- Each queue has its own scheduling algorithm,
  - foreground - RR
  - background - FCFS
- Scheduling must be done between the queues.
  - Fixed priority scheduling; (i.e., serve all from foreground then from background). Possibility of starvation.
  - Time slice - each queue gets a certain amount of CPU time which it can schedule amongst its processes; e.g.,
    - 80% to foreground in RR
    - 20% to background in FCFS

1/26/2005

CSC 256/456 - Spring 2005

19

## Real-Time Scheduling

- *Hard real-time systems* - required to complete a critical task within a guaranteed amount of time.
- *Soft real-time computing* - requires that critical processes receive priority over less fortunate ones.
- EDF - Earliest Deadline First Scheduling.

1/26/2005

CSC 256/456 - Spring 2005

20

## Linux Task Scheduling

- Linux uses two task-scheduling classes:
  - A time-sharing class for fair preemptive scheduling between multiple tasks.
  - A real-time class that conforms to POSIX real-time standard. (FIFO/RR)
- For time-sharing tasks, Linux uses a prioritized, credit based algorithm.
  - Scheduling is prioritized based on the initial credit
  - Credit of the running task decrements by one at every clock tick
  - Recrediting when no runnable tasks have any credit

$$\text{credits} = \frac{\text{credits}}{2} + \text{priority}$$

- This crediting system automatically prioritizes interactive or I/O-bound tasks.

1/26/2005

CSC 256/456 - Spring 2005

21

## Disclaimer

- Parts of the lecture slides contain original work of Abraham Silberschatz, Peter B. Galvin, Greg Gagne, Andrew S. Tanenbaum, and Gary Nutt. The slides are intended for the sole purpose of instruction of operating systems at the University of Rochester. All copyrighted materials belong to their original owner(s).

1/26/2005

CSC 256/456 - Spring 2005

22