Fall 2015 COMP 3511 Homework Assignment #2 Solution

Handout Date: Oct 5, 2015 Due Date: Oct 19, 2015

Name: _____ ID: _____
E-Mail: _____ Section: _____

Please read the following instructions carefully before answering the questions:

- You should finish the homework assignment **individually**.
- There are a total of **4** questions.
- When you write your answers, please try to be precise and concise.
- Fill in your name, student ID, email at the top of each page.
- Please fill in your answers in the space provided, or you can type your answers in the MS Word file.

- **Homework Collection: the hardcopy** is required and the homework is collected in **collection box #16** (for **L1**) and **collection box #17** (for **L2**). The collection boxes are located outside **Room 4210**, near **Lift 21** (there are labels on the boxes).

1. (20 points) Multiple choices

1) A ____ provides an API for creating and managing threads.

A) set of system calls
B) multicore system
C) thread library
D) multithreading model
Answer: C

2) Which of the following would be an acceptable signal handling scheme for a multithreaded program?

A) Deliver the signal to the thread to which the signal applies.
B) Deliver the signal to every thread in the process.
C) Deliver the signal to only certain threads in the process.
D) All of the above
Answer: D

3) The ____ multithreading model multiplexes many user-level threads to a smaller or equal number of kernel threads.

A) many-to-one model
B) one-to-one model
C) many-to-many model

D) many-to-some model
Answer: C

4) _____ involves distributing tasks across multiple computing cores.

A) Concurrency
B) Task parallelism
C) Data parallelism
D) Parallelism
Answer: B

5) According to Amdahl's Law, what is the speedup gain for an application that is 40% parallel and we run it on a machine with 4 processing cores?

A) 0.7
B) 1.82
C) 0.55
D) 1.43
Answer: D

6) ____ is the number of processes that are completed per time unit.

A) CPU utilization
B) Response time
C) Turnaround time
D) Throughput
Answer: D

7) The ____ scheduling algorithm is designed especially for time-sharing systems.

A) SJF
B) FCFS
C) RR
D) Multilevel queue
Answer: C

8) Which of the following is true of multilevel queue scheduling

A) Processes can move between queues.
B) Each queue has its own scheduling algorithm.
C) A queue cannot have absolute priority over lower-priority queues.
D) It is the most general CPU-scheduling algorithm.
Answer: B

9) _____ allows a thread to run on only one processor.

A) Processor affinity
B) Processor set
C) NUMA
D) Load balancing
Answer: A

10) A significant problem with priority scheduling algorithms is _____.

A) complexity
B) starvation
C) determining the length of the next CPU burst
D) determining the length of the time quantum
Answer: B

2. [15 points] Threads

1) (5 points) Please briefly explain why a (busy) database server should not run as a single-threaded process.

Answer: If a database server runs as a single-threaded process, only one client can be serviced at a time. This either results in potentially enormous waiting times for a busy server or a large number of processes have to be created with significant redundancies (program and data).

2) (5 points) Please briefly describe the difference between the fork() and clone() Linux system calls

Answer: The fork() system call duplicates the entire process. The clone() system call behaves similarly except that, instead of creating a copy of the process, it creates a separate process that shares the address space of the calling process based on the flags specified in the clone() parameters..

3) (5 points) What is the advantage of **deferred cancellation** over **asynchronous cancellation**?

Answer: In the asynchronous cancellation, the thread is immediately cancelled in response to a cancellation request. There could be potential danger if the thread is in the middle of a data update. In the deferred cancellation, the thread polls whether or not it should terminate, thus the thread can be made to cancel at a safer time.

3. [25 points] Please answer the following questions in a few sentences

1) (5 points) What is processor affinity? What are hard affinity and soft affinity?

Answer: Processor affinity implies that a process has an affinity for the processor on which it is currently running. The hard affinity allows a process to specify a subset of processors on which it may run. If the operating system only attempts to keep a process running on the same processors, but not guaranteeing that, this situation is known as soft affinity.

2) (10 points) Consider a system running ten I/O-bound tasks and one CPU-bound task. Assume that the I/O-bound tasks issue an I/O operation once for every millisecond of CPU computing and that each I/O operation takes 10 milliseconds to complete. Also assume that the context-switching overhead is 0.1 millisecond and that all processes are long-running tasks. Calculate the CPU utilization for a round-robin scheduler when:

   a) The time quantum is 1 millisecond

   b) The time quantum is 10 milliseconds

Answer:

a. The time quantum is 1millisecond: Irrespective of which process is scheduled, the scheduler incurs a 0.1 millisecond context-switching cost for every context-switch. This results in a CPU utilization of 1/1.1 * 100 = 91%.

b. The time quantum is 10 milliseconds: The I/O-bound tasks incur a context switch after using up only 1 millisecond of the time quantum. The time required to cycle through all the processes is therefore 10*1.1 + 10.1 (as each I/O-bound task executes for 1 millisecond and then incur the context switch task, whereas the CPU-bound task executes for 10 milliseconds before incurring a context switch). The CPU utilization is therefore 20/21.1 * 100 = 94%.

3) (5 points) Consider the scheduling algorithms, FCFS, SJF, RR, SRTF and priority scheduling, which can result in starvation and why?

Answer: Any priority-based scheduling algorithm can result in starvation, in which if there are always higher priority processes in the system. These include SJF, SRTF and priority scheduling.

4) (5 points) Please describe the usage of process-contention scope and system-contention scope.

Answer: The process-contention scope or PCS dictates how the thread library schedules user-level threads to run on an available LWP. This is typically done according to priority. The system-contention scope or CSC is used by the kernel to determine which kernel-level thread to schedule onto a CPU.

4. [40 points] CPU Scheduling

1) (20 points) Given the arrival time and CPU-burst of 5 processes shown in the following diagram:

Process	Arrival Time (ms)	Burst Time (ms)
P1	0	8
P2	2	2
P3	6	15
P4	12	4
P5	19	13

Suppose the OS uses a 3-level feedback queue to schedule the above 5 processes. Round-Robin scheduling strategy is used for the queue with the highest priority and the queue with the second highest priority, but the time quantum used in these two queues is different. First-come-first serve scheduling strategy is used for the queue with the lowest priority. The scheduling is **preemptive**.

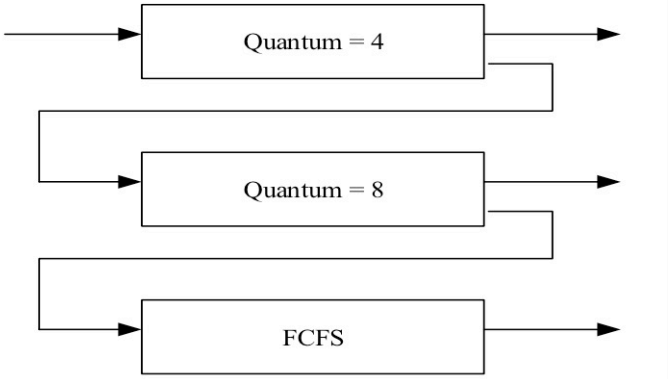

a) (10 points) Construct a Gantt chart depicting the scheduling for the set of processes specified in the above diagram using this 3-level feedback queue.

Answer:

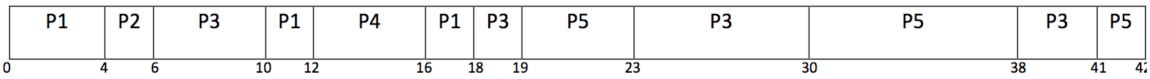

b) (10 points) Calculate the average waiting time for the schedule constructed in a).

Answer:

Waiting time for each process:

P1: 6 + 4 = 10,

P2: 2,

P3: 8 + 4 + 8 = 20,

P4: 0,

P5: 7 + 3 = 10

Average waiting time: (10 + 2 + 20 + 0 + 10) / 5 = 8.4

2) (20 points) Consider the following set of processes, with the length of the CPU burst time given in milliseconds:

Process	Arrival Time(ms)	Burst Time(ms)
P1	0	30
P2	19	15
P3	42	24
P4	27	10
P5	3	59

a) (10 points) Draw three Gantt charts that illustrate the execution of these processes using the scheduling algorithms listed below:
(i) Shortest-Job-First
(ii) Shortest-Remaining-Time-First (Preemptive Shortest-Job-First)
(iii) Round-Roubin (quantum = 8)

Answer:
(i) Shortest-Job-First

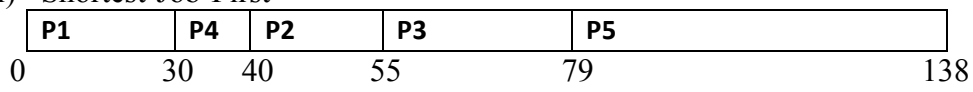

(ii) Shortest-Remaining-Time-First

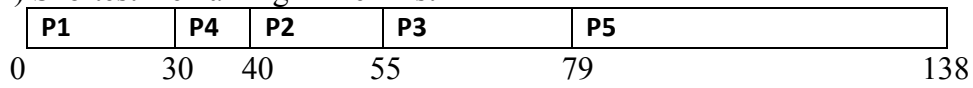

(iii) RR (quantum = 8)

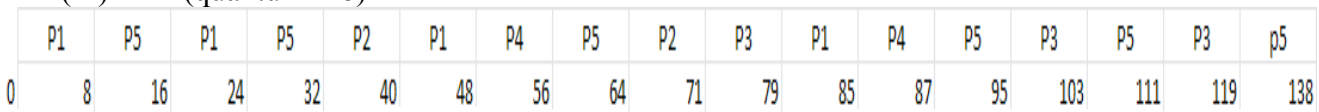

b) (5 points) What is the turnaround time of each process for each of the scheduling algorithms in part a?

Turnaround time	P1	P2	P3	P4	P5

SJF					
SRJF					
RR					

Answer:

Turnaround time	P1	P2	P3	P4	P5
SJF	30	36	37	13	135
SRJF	30	36	37	13	135
RR	85	52	77	60	135

c) (5 points) What is the waiting time of each process for each of these scheduling algorithms in part a?

Waiting time	P1	P2	P3	P4	P5
SJF					
SRJF					
RR					

Answer:

Waiting time	P1	P2	P3	P4	P5
SJF	0	21	13	3	76
SRJF	0	21	13	3	76
RR	55	37	53	50	76