

Lecture

- In the lectures in week 10 we will finish Chapter 4 (Mutithreaded Programming). We will start with Chapter 5 (Process Scheduling). Examples will be shown for the simulation of Process Scheduling. Solaris process scheduling administration will be illustrated using `dispadmin`.
- Note, as usual, that you find even more exercises including solutions here :
<http://codex.cs.yale.edu/avi/os-book/OS9/practice-exer-dir/index.html>

Prepare for the Tutorial Sessions in week 11, 2018: all exercises not discussed so far (which includes very likely all exercises from the weekly notes in week 10). In addition:

- 5.1 Why is it important for the scheduler to distinguish I/O-bound programs from CPU-bound programs?
- 5.2 Discuss how the following pairs of scheduling criteria conflict in certain settings.
 - a. CPU utilization and response time
 - b. Average turnaround time and maximum waiting time
 - c. I/O device utilization and CPU utilization
- 5.4 In this chapter, we discussed possible race conditions on various kernel data structures. Most scheduling algorithms maintain a run queue, which lists processes eligible to run on a processor. On multicore systems, there are two general options: (1) each processing core has its own run queue, or (2) a single run queue is shared by all processing cores. What are the advantages and disadvantages of each of these approaches?
- 5.5 Consider the exponential average formula used to predict the length of the next CPU burst. What are the implications of assigning the following values to the parameters used by the algorithm?
 - a. $\alpha = 0$ and $\tau_0 = 100$ milliseconds
 - b. $\alpha = 0.99$ and $\tau_0 = 10$ milliseconds

- 5.7 Consider the following set of processes, with the length of the CPU burst time given in milliseconds:

Process	Burst Time	Priority
P ₁	2	2
P ₂	1	1
P ₃	8	4
P ₄	4	2
P ₅	5	3

The processes are assumed to have arrived in the order P₁, P₂, P₃, P₄, P₅ all at time 0.

- a. Draw four Gantt charts that illustrate the execution of these processes using the following scheduling algorithms: FCFS, SJF, nonpreemptive priority (a smaller priority number implies a higher priority), and RR (quantum = 1).
 - b. What is the turnaround time of each process for each of the scheduling algorithms in part a?
 - c. What is the waiting time of each process for each of these scheduling algorithms?
 - d. Which of the algorithms results in the minimum average waiting time (over all processes)?
- 5.8 The following processes are being scheduled using a preemptive, round-robin scheduling algorithm. Each process is assigned a numerical priority, with a higher number indicating a higher relative priority. In addition to the processes listed below, the system also has an idle task (which consumes no CPU resources and is identified as P_{idle}). This task has priority 0 and is scheduled whenever the system has no other available processes to run. The length of a time quantum is 10 units. If a process is preempted by a higher-priority process, the preempted process is placed at the end of the queue.

Thread	Priority	Burst	Arrival
P_1	40	20	0
P_2	30	25	25
P_3	30	25	30
P_4	35	15	60
P_5	5	10	100
P_6	10	10	105

- a. Show the scheduling order of the processes using a Gantt chart.
 - b. What is the turnaround time for each process?
 - c. What is the waiting time for each process?
 - d. What is the CPU utilization rate?
- 5.10 Which of the following scheduling algorithms could result in starvation?
- a. First-come, first-served
 - b. Shortest job first
 - c. Round robin
 - d. Priority
- 5.11 Consider a variant of the RR scheduling algorithm where the entries in the ready queue are pointers to the PCBs.
- a. What would be the effect of putting two pointers to the same process in the ready queue?

- b. What would be two major advantages and disadvantages of this scheme?
 - c. How would you modify the basic RR algorithm to achieve the same effect without the duplicate pointers?
- 5.12 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. Describe is the CPU utilization for a round-robin scheduler when:
- a. The time quantum is 1 millisecond
 - b. The time quantum is 10 milliseconds
- 5.13 Consider a system implementing multilevel queue scheduling. What strategy can a computer user employ to maximize the amount of CPU time allocated to the user's process?
- 5.14 Consider a preemptive priority scheduling algorithm based on dynamically changing priorities. Larger priority numbers imply higher priority. When a process is waiting for the CPU (in the ready queue, but not running), its priority changes at a rate α ; when it is running, its priority changes at a rate β . All processes are given a priority of 0 when they enter the ready queue. The parameters α and β can be set to give many different scheduling algorithms.
- a. What is the algorithm that results from $\beta > \alpha > 0$?
 - b. What is the algorithm that results from $\alpha < \beta < 0$?
- 5.16 Explain the differences in how much the following scheduling algorithms discriminate in favor of short processes:
- a. FCFS
 - b. RR
 - c. Multilevel feedback queues
- 5.19 Assume that two tasks A and B are running on a Linux system. The nice values of A and B are -5 and +5, respectively. Using the CFS scheduler as a guide, describe how the respective values of vruntime vary between the two processes given each of the following scenarios:
- a. Both A and B are CPU-bound.
 - b. A is I/O-bound, and B is CPU-bound.
 - c. A is CPU-bound, and B is I/O-bound.