

IA Operating Systems: Processes, Scheduling, Interrupts, System Calls

Lent 2024

Last Updated: Friday 23rd February, 2024 23:47

- 1 For each of the following, indicate if the statement is true or false, and explain why:
 - (a) Preemptive schedulers require hardware support.
 - (b) A context switch can be implemented by a flip-flop.
 - (c) System calls are an optional extra in modern operating systems.
 - (d) Shortest job first (SJF) is an optimal scheduling algorithm.
 - (e) Round-robin scheduling can suffer from the so-called ‘convoy effect’.

- 2
 - (a) Operating systems need to be able to prevent applications from crashing or locking up the system, or from interfering with other applications. Which three kinds of hardware support do we require to accomplish this?
 - (b) How do applications request that the operating system perform tasks on their behalf?
 - (c) What could we do if we did not have the requisite hardware support?

- 3 Process scheduling can be *preemptive* or *non-preemptive*. Compare and contrast these approaches, commenting on issues of simplicity, fairness, performance and required hardware support.

4 (a) Describe how the CPU is allocated to processes if static priority scheduling is used. Be sure to consider the various possibilities available in the case of a tie.

(b) “All scheduling algorithms are essentially priority scheduling algorithms.”

Discuss this statement with reference to the first-come first-served (FCFS), shortest job first (SJF), shortest remaining time first (SRTF) and round-robin (RR) scheduling algorithms.

(c) What is the major problem with static priority scheduling and how may it be addressed?

(d) Why do many CPU scheduling algorithms try to favour I/O intensive jobs?

5 An operating system uses a single queue round-robin scheduling algorithm for all processes. You are told that a *quantum* of three time units is used.

(a) What can you infer about the scheduling algorithm?

(b) Why is this sort of algorithm suitable for a multi-user operating system?

(c) The following processes are to be scheduled by the operating system.

Process	Creation Time	Required Computing Time
P_1	0	9
P_2	1	4
P_3	7	2

None of the processes ever blocks. New processes are added to the tail of the queue and do not disrupt the currently running process. Assuming context switches are instantaneous, determine the *response time* for each process.

(d) Give one advantage and one disadvantage of using a small quantum.

- 6 (a) Describe with the aid of a diagram the life-cycle of a process. You should describe each of the states that it can be in, and the reasons it moves between these states.
- (b) What information does the operating system keep in the process control block?
- (c) Give one advantage and one disadvantage of non-preemptive scheduling.
- (d) What steps does the operating system take when an interrupt occurs?
- (e) What problems could occur if a system experienced a very high interrupt load?
- (f) What information do the shortest job first (SJF) and shortest remaining time first (SRTF) algorithms require about each job or process? How can this information be obtained?

7 *System calls* are part of most modern operating systems.

- (a) What is the purpose of a system call?
- (b) What mechanism is typically used to implement system calls?

8 Past paper questions

- [y2013p2q3](#)
- [y2012p2q3](#)
- [y2010p2q3](#) [not (c)]
- [y2015p2q3](#) [not (c)]
- [y2014p2q3](#)
- [y2011p2q3](#) [not (a)(iv), (a)(v), (b)]
- [y2007p1q8](#) (a)–(c)