

COMPUTER SCIENCE TRIPOS Part IA – 2013 – Paper 2

3 Operating Systems (IML)

Consider a single processor system in which multiple processes are running.

- (a) What does it mean for a process to be I/O bound? What does it mean for it to be CPU bound? [2 marks]
- (b) What is the difference between preemptive and non-preemptive scheduling? Which one requires specific hardware support and what is that hardware support? [3 marks]
- (c) Two processes, *A* and *B*, have the following sequential execution patterns:

A: [cpu 4 ms; I/O 2 ms; cpu 4 ms; I/O 2 ms; cpu 4 ms]

B: [cpu 1 ms; I/O 2 ms; cpu 1 ms; I/O 2 ms; cpu 1 ms]

I/O operations for the two processes do not interfere with each other and are blocking.

- (i) If the processes are run consecutively one after another, what is the elapsed time for all to complete? [2 marks]
- (ii) Sketch the execution pattern under non-preemptive scheduling and determine the total elapsed time for completion. You may assume that processes are scheduled in the order in which they become ready to run and that in the event of a tie *A* has priority over *B*. You may further assume that the scheduler and context switches take negligible time. [6 marks]
- (iii) Repeat (c)(ii) but for a preemptive scheduler that operates on a time slice of 2 ms, that is, no process can run for more than 2 ms at a time (unless no other process is runnable). [6 marks]
- (iv) Is there any evidence from your results for (c)(ii) and (c)(iii) of a significant advantage for either scheduling method? [1 mark]