UNIVERSITY OF EDINBURGH

COLLEGE OF SCIENCE AND ENGINEERING

SCHOOL OF INFORMATICS

OPERATING SYSTEMS

Wednesday 28th April 2010

09:30 to 11:30

Year 3 Courses

Convener: K. Kalorloti
External Examiners: K. Eder, A. Frisch

INSTRUCTIONS TO CANDIDATES

Answer any TWO questions.

All questions carry equal weight.

CALCULATORS MAY NOT BE USED IN THIS EXAMINATION

1. (a) Give **brief** (one or two sentence) definitions of the following four terms in memory management: *virtual memory*, *page table*, *secondary storage*, *segment register*. [*8 marks*]

   (b) Describe in detail what happens in a demand-paged system when a user process accesses memory that is currently paged out. [*8 marks*]

   (c) The IBM System/370 family (and its successors) has an instruction
   `MVC dst(len) src`
   which moves *len* bytes (from 1 to 256) from memory location *src* to memory location *dst*. The source and destination may overlap – the instruction behaves as if it copies one byte at a time, storing the nth byte before fetching the $(n+1)$th byte.

   Suppose that an S/370 system is running with a page size of 4K, and that for some user process the page from (virtual address) 0 to 4095 is in memory, and the page from 4096 to 8191 is paged out. The user process executes the instruction
   `MVC 4088(100) 4090`

   i. Explain, with reference to your answer to part (b), what problem arises. [*5 marks*]

   ii. Propose ways in which the hardware could handle the problem. [*4 marks*]

2. This question is about disk i/o and disk scheduling.

   (a) Explain the translations that occur between a user program requesting some data from an open file, and the disk scheduler being called to schedule a real disk i/o request. [*4 marks*]

   (b) Explain briefly how the three scheduling algorithms SSTF, SCAN and C-SCAN work. [*9 marks*]

   (c) Modern disks often present their structure to the OS in terms of *logical block addresses* (LBAs), a simple sequential numbering of blocks, rather than in terms of the real cylinder, head and sector (CHS) addresses assumed by the scheduling algorithms. If the algorithms of part (b) are still to work, what assumption needs to be true about the relationship between the LBA and CHS addresses? [*3 marks*]

   (d) Modern disks are usually internally formatted with some spare sectors, which are swapped in to replace any sectors that fail during the lifetime of the disk. This re-mapping is internal to the disk. What problem could this cause for disk scheduling, and how would you design the disk format to avoid it? [*4 marks*]

   (e) None of the algorithms in part (b) is *fair* – for each of them, there are situations in which a process may get slower service than another, simply because of its pattern of i/o. Give examples of such situations for each of SSTF and SCAN. [*5 marks*]

3. (a) What is a pre-emptive process scheduler? Describe the *round-robin* (RR) and *feedback* (*priority queue*) approaches. [*8 marks*]

(b) Suppose that the usual description of the RR algorithm is modified so that a single process can have two (or more) independent entries in the algorithm's data structures, so that a process can be doubly, triply, etc. scheduled.

i. What is the effect of having two entries for the same process? Does it depend on how they are placed? [*5 marks*]

ii. Under what circumstances would you consider it appropriate to give a process multiple entries? [*5 marks*]

(c) In a scheduling system known as *preemptive priority scheduling*, the priority of each process is a (positive or negative) number that changes dynamically, at a rate α when it is on the ready queue, and at a rate β when it is actually running. Processes are initialized with priority 0, and reset to priority 0 when they return from a blocked state. (Positive α means increasing priority.) What happens if the algorithm is run with the following settings? Explain your answers.

i. $\beta > \alpha > 0$ [*3 marks*]

ii. $\beta < \alpha < 0$ [*4 marks*]