UNIVERSITY OF EDINBURGH

FACULTY OF SCIENCE AND ENGINEERING

MASTER OF SCIENCE IN COMPUTER SCIENCE

OPERATING SYSTEMS

Wednesday 19 April 2000

2:00 to 3:30

Examiners:

I. Watson (External)

R.N. Ibbett (Chairman)

INSTRUCTIONS TO CANDIDATES

Answer any TWO questions.

All questions carry equal weight.

1. (a) What is a *semaphore*? Define the *P* and *V* operations on a semaphore. How can P and V operations be used to ensure mutual exclusion of processes from a critical region? [*4 marks*]

   (b) What is meant by the term *busy waiting* in relation to semaphores and how can it arise? How is it avoided in operating systems? [*4 marks*]

   (c) How might semaphore facilities be implemented by the kernel of an operating system? Describe the data structures that would be needed and how they would be used. [*6 marks*]

   (d) What are the necessary conditions for *deadlock* to arise? Give an example of deadlock caused by the use of semaphores by *three* processes. [*5 marks*]

   (e) How could the kernel detect the presence of deadlock arising from use of semaphores? Describe a scheme using the data structures you have suggested in part (c). [*6 marks*]

2. (a) List objectives that the designer of a CPU scheduling strategy should aim to meet in each of the following contexts:

   i. a single-user workstation
   ii. a multiple-user interactive server
   iii. a background job-stream server
   iv. a real-time system

   [*6 marks*]

   (b) Briefly describe the operation of the following scheduling schemes and the advantages and disadvantages that each has:

   i. first-come-first-served
   ii. round-robin
   iii. highest response ratio
   iv. preemptive priority

   [*8 marks*]

   (c) What is meant by the terms *hard* and *soft* deadlines for scheduling in real-time systems? Give examples where the need for each arises. Describe scheduling schemes which are appropriate where hard and soft deadlines arise. [*8 marks*]

   (d) What is meant by the term *priority inversion* and how can the problems it causes be alleviated? [*3 marks*]

3. (a) Describe what is meant by virtual memory *paging*. What problems does paging overcome and how does it do so? What is a *page table* and what does it consist of? Describe the purpose of each component of a page table entry. [*7 marks*]

   (b) What is meant by the terms *page fault* and *demand paging*? Describe the sequence of processing stages, both hardware and software, that occur when a virtual page address is accessed, in particular when the page is not present in physical main memory. [*7 marks*]

   (c) What are *multi-level* page tables, how do they work and what problems do they overcome? Describe the organisation of the Sun's Sparc architecture page tables. [*4 marks*]

   (d) What is an *inverted* page table and how does it work? What is meant by the term *hardware hashing*? [*4 marks*]

   (e) What is meant by a *zero-level* paging scheme? [*3 marks*]