Where an algorithm is asked for, you may write in any suitable pseudocode. Correct syntax for any computer language is not expected.

Answer 3 questions.

1.

a) What information is kept in a *Process Control Block* (PCB) and why? [6]

b) Of that information, what must be saved or updated when a process is context switched? [4]

c) Consider a simple priority-based scheduling system. Priority is determined by the fraction of its time quantum that a process last used. Thus, a CPU bound process using its entire time quantum would get a priority of 1, whereas a process using a fifth of its time quantum would get a priority of 5.

i) Describe in detail how such a scheme might work.

ii) Give a critique of the scheme and suggest an alternative.

iii) How would you test whether your scheme was, in fact, better than the scheme proposed above? [15]

d) If the simple scheme above is extended by making it preemptive, what is the average wait time for a process if the order and duration is as follows, and the timeslice is 100ms. You may ignore context-switching time and assume that a process always starts at priority level 1. State any other assumptions you make.

	P_1	P_2	P_3	P_4
Arrival time	0	55	100	300
Duration /ms	510	20	280	10
CPU burst /ms	120	5	50	10
I/O burst/ms	20	20	30	0

[8]

[TURN OVER]

2.

a) State Little's Law. [2]

b) A paging disk has a queue associated with it. Some measurements have been done on this queue. At 1ms intervals, the length of the queue was found to be:

1 2 1 1 2 3 4 5 4 3 4 4 3 2 2 3 3 3 4 5 6

i) What is the arrival rate?

ii) What is the throughput?

iii) Why is there a difference between these values?

iv) What is the average residence time?

State any assumptions you make. [8]

c) If the hardware components of a 64-bit machine's memory management system include a Translation Lookaside Buffer (TLB), main memory, and disk, explain how one would implement a paging system for the machine. [15]

d) The following timings apply to the system you have described in part (c):

i) A TLB lookup takes 10ns.

ii) Main memory access time is 100ns.

iii) 75% of pages referred to are in the TLB. 10% of those that are not in the TLB cause page faults.

iv) A page fault is serviced by the disk described in part (b).

What is the effective memory access time? [8]

3.

You have been commissioned by a publisher to produce a web-based course on distributed systems. The intended market for this course are humanities graduates who have had no previous distributed systems training.

Describe what you would put into the course and justify its relevance. Also, say how each section relates to the sections before and after it, and say how the course as a whole attempts to achieve the aims of the project as you interpret them. [33]

[CONTINUED]

4.

a) Consider the following program segments for two different processes executing concurrently. `x` is a shared variable, initialised to zero and assignment is atomic.

Process 1	Process 2
`x = x + 1`	`x = x * 2`

i) What are the possible interleavings?

ii) What are the possible resulting values for `x` and which is the least likely?

Hint: consider how this might compile into machine code. [6]

b) If processes share variables, there is a need to use concurrency control primitives to define critical sections inside which the shared variables may be manipulated. What three criteria should be satisfied by such concurrency control primitives? [6]

c) Analyse and explain the operation of the following primitives and say whether or not it meets the criteria you identified in part (b) above and/or whether it has other problems. You may assume that the code for each primitive is executed atomically. [10]

```
boolean b = FALSE;

lock()
{ while (b)
    no_op
  b = TRUE;
}

unlock()
{ b = FALSE; }
```

d) In the following system there are four processes. Process P_1 places data into a bounded buffer BB_1. Process P_2 places different data into an unbounded buffer UB_1. Process P_3 extracts a datum from BB_1 and another from UB_1, combines them in some way and puts the result into a bounded buffer BB_2. Process P_4 extracts data from BB_2 and prints it on the screen. Diagrammatically, this can be seen as follows:

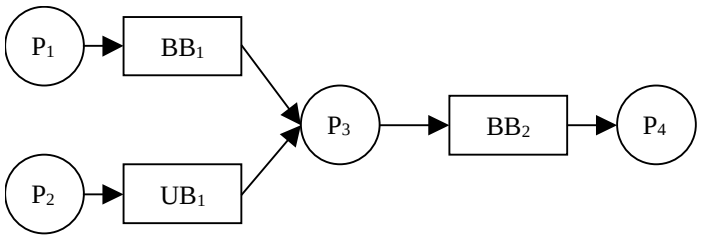

Write pseudocode for all of the processes using any concurrency control or synchronisation primitives you wish. Make clear which variables are shared and between which processes. [11]

[TURN OVER]

5.

a) At Bentham University, it is necessary to transfer certain information between departments and the central administration via a campus-wide network. This information must be protected against accidental or willful damage and it must be possible to authenticate the originator. (Examples of such information are salary approvals, expense payments requests, or examinations marks transfer.) Identify the security threats in such systems and propose suitable mechanisms to counter them. [18]

b) After the system is installed, it is noted that some actions (e.g. overtime requests, expense payment requests) must be approved by Heads of Department (HoD). However, it is well known that the Bentham HoDs spend a large amount of time at conferences around the world. Consequently, it is necessary to allow the HoDs some means of logging in to a Bentham site, and approving (or not) the actions whilst they are away. However, the links between the machines they use remotely and Bentham must be considered insecure as must the remote machines themselves. Devise a mechanism that would allow HoDs to log in, and state its limitations and any assumptions you make. [15]

[END OF PAPER]