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. You have been commissioned by a publisher to produce two one hour videos which describe the fundamental differences between centralised and distributed systems. The intended market for these videos is computer science graduates and working programmers who have had no previous distributed systems training.

   Describe what you would put into each of the videos. To do this, describe what the major subject headings will be and, for each one, justify its relevance and how it relates to the sections before and after it. [33]

2. The 'thugs' and the 'hippies' are two groups of people who live near each other. They must cross a river to get to the nearest pub. A single boat is used to cross the river, but it only seats three people, and must always carry a full load. In order to guarantee the safety of the hippies, you cannot put one hippy and two thugs in the same boat (because the thugs would overpower and rob the hippy), but all other combinations are legal.

   Two procedures are needed: `HippyArrives()` and `ThugArrives()`, called by a hippy or thug when they arrive at the river bank. The procedures must arrange the arriving hippies and thugs into safe boatloads. Once the boat is full, then the hippies and thugs who make up the boatload must be transported across the river (some thread calls `RowTheBoat()`) and the corresponding `HippyArrives()` and `ThugArrives()` procedure calls must return.

   a) Define a *critical region*. [2]

   b) What is the shared resource in this case? [1]

   c) Using any concurrency control primitives you choose, implement a solution to this problem, and show that it operates correctly and that there is no starvation, no deadlock, and no undue waiting (i.e. as soon as there is a boatload of people, it should cross the river.) [30]

[TURN OVER]

3.

a)

i) Why cannot one use simple paging as a mechanism for memory management in machines with a 64 bit architecture? [3]

ii) Briefly, describe and explain an alternative mechanism which might be more appropriate. [5]

b) What is a *Translation Lookaside Buffer (TLB)* and why might one be needed in the system you described in a) (ii)? Explain what the various implementation options are for a TLB, describing what the advantages and disadvantages of each are. [8]

c)

i) What is Belady's anomaly? [2]

ii) Show that the least recently used (LRU) page replacement algorithm cannot suffer from Belady's anomaly. [4]

iii) LRU is not regarded as a practical page replacement algorithm. Why not? [2]

iv) Briefly, describe the implementation of an approximation to LRU. [4]

d) Suppose we have a demand-paged virtual memory (16 bit words) where the address is arranged as below:

Page number	Offset
4 bits	12 bits

The most recent references are cached in a TLB (access time 10ns) and there is a 95% hit rate for memory references. In case of a miss, the page table held in main memory (access time 200ns) must be consulted. 1% of the TLB misses cause page faults which must be serviced from a paging disk whose latency is 5ms and whose transfer rate is 2 MB/s. What is the effective memory access time? State any assumptions you make. [5]

[CONTINUED]

4.

a) Briefly define the terms *Public Key Cryptography (PKC)* and *Secret Key Cryptography (SKC).* [8]

b) It has been decided that, at a well-known college of London University, all exam marks will be entered into spreadsheets by the departments which set the examination, then sent to the Registry over the college network where they will be placed in individual student's (computerised) records. The exam marks must be protected against accidental or willful damage at all stages and it must be possible to authenticate the originator of any information. Identify the security threats in such a system and propose suitable mechanisms to counter them. [25]

5.

a) My local supermarket has a policy which states that if there is (on average) more than one person waiting in a queue they will open another checkout. It takes an average of 5 minutes for each person to pay for their goods

i) The following graph shows the number of tills open as the rate at which people complete their shopping increases. Derive values for the points labeled *a, b, c, d, e*. State any assumptions you make. [4]

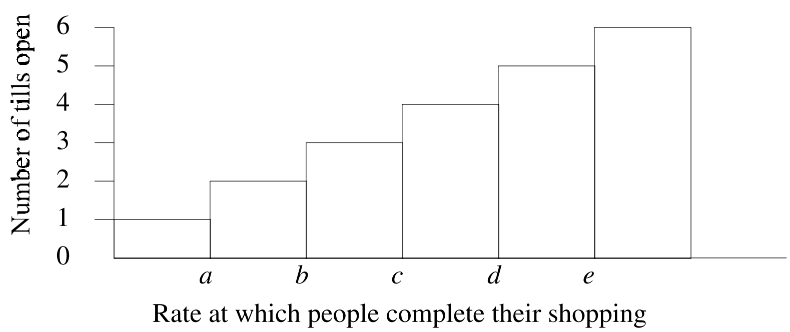

ii) Sketch the graph which shows the average wait time per shopper as the rate at which people complete their shopping increases, marking values in for salient points. Remember that the number of checkouts open changes as the rate at which people finish their shopping increases. [13]

[Question 5 cont. over page]

[Question 5 cont.]

b) In the following system the visit counts for different types of traffic at different devices has been determined.

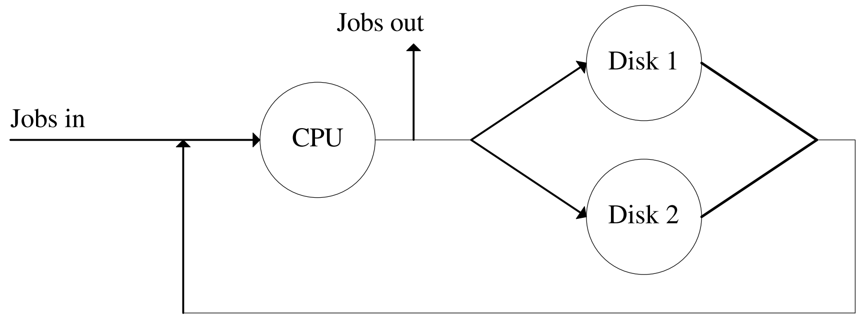

The results are as follows:

Traffic type	Arrival rate (jobs/hour)
Compilation	**60**
Execution	**1200**
Edits	**300**

The visit counts and service time have been calculated for each type of traffic at each device.

Traffic Type	CPU		Disk 1		Disk 2	
	S /s	V	S /s	V	S /s	V
Compilation	**0.9**	**?**	**1**	**10**	**0.5**	**10**
Execution	**0.04**	**?**	**0.02**	**20**	**0.5**	**5**
Edits	**0.01**	**?**	**0.01**	**10**	**0.02**	**10**

i) Fill in the questionmarks in the above table. [3]

ii) Determine the utilisation of each device. State any assumptions you make. [7]

iii) Which device is the bottleneck device? [1]

iv) Sketch a graph which shows system throughput as the overall rate of jobs coming into the system increases (but the same proportions are retained). Think time is 10s. [5]

[END OF PAPER]