Computer Science Department

Dummy Cover Sheet

M.Sc. / Coll Dip

E1 Real Time Computing 1997-98

2 hours 30 minutes

Computer Science Department

*Answer **TWO** questions from each section. All questions carry equal credit.*

The use of electronic calculators is permitted.

Part C, Real-Time Computing

*Answer **TWO** questions from this section.*

Question C1

An automatic video system is to be used to transmit lectures live to other institutions by using a pan and tilt camera with variable zoom to image the podium of a lecture and to track lecturers in close-up as they move about. The tracking is to be performed automatically by processing images from the camera as well as transmitting them to remote sites.

(a) Describe the main characteristics of a real-time computing system and explain why the computer system required to perform the tracking above is 'real-time'. Discuss whether it is a *hard* real-time computing system, or a *soft* real-time computing system.

[**12 marks]**

(b) A high resolution, digital, colour camera transmitting 1536 x 2048 images at 24 bits/pixel in RGB format (ie: in red, green and blue colour channels, each encoded at 8 bits/pixel) at 30 frames/sec is used. The tracking algorithm involves computing the colour and motion characteristics of pixels in successive images and requires approximately 20 arithmetic operations per pixel. Estimate:

(i) the amount of data for each image frame in MBytes,

(ii) the data rate in Mbytes/sec,

(iii) the processing rate required.

Comment on the implications of the processing rates you have estimated:

(1) for building such a system from readily available processors,

(2) for building such a system for mass marketing.

In each case, suggest what you would do to implement the system and, if necessary, what steps you would take to make a solution easier to achieve.

[**8 marks]**

[TURN OVER]

Question C2

(a) Explain the difference between a pipelined multiprocessor system and a parallel multiprocessor system, and describe the advantages and disadvantages of each. Describe what is meant by an 'MIMD' parallel system and a 'processor farm' and explain why a processor farm is often used to implement an MIMD system.

[12 marks]

(b) Describe how a processor farm is implemented and explain why architectures based on pipelines of processor farms are often useful in the implementation of computationally demanding real-time systems. Illustrate your answer by reference to a system that has been broken down into a three-stage pipeline with computational times of 10, 15, and 20 ms respectively. Indicate what problems might frustrate your efforts, and how you might mitigate their effects.

[**8 marks]**

Question C3

(a) Explain the difference between *independent*, *co-operating* and *competing* tasks in a real-time computing system. Show, by writing appropriate pseudo-code, how you would use a semaphore to ensure that two tasks A and B can only gain mutually exclusive access to a memory M. Explain what the functions you have used do and indicate briefly what problems you might expect to encounter in using a semaphore in this way.

[**12 marks]**

(b) In a particular real-time system, tasks may be *active*, *runnable* or *suspended*. Describe what these terms mean, and use them to explain why two tasks A and B cannot, in general, safely share access to the same code for a module M. Describe briefly how you would implement M so as to ensure that the code could safely be shared, *without* use of a semaphore or other access control mechanism such as a monitor.

[8 marks]

[CONTINUED]

Question C4

(a) Define what is meant when the tasks in a real-time system are said to be *schedulable* and explain the difference between *cyclic*, *round-robin* and *priority* scheduling strategies in a multi-tasking, real-time system. Explain how the implementation of round-robin and priority scheduling strategies in a real-time system may differ from the implementation of similar strategies on a conventional computer system such as a UNIX workstation.

[12 marks]

(b) According to the rate monotonic algorithm, n tasks are schedulable on a single processor if the overall CPU utilisation is less than $n\left(\sqrt[n]{2}-1\right)$.

Explain how priorities are determined by the rate monotonic algorithm and describe any limitations on the application of the formula. Explain what is meant by CPU utilisation and why, according to the above formula, a single task can be scheduled if its utilisation is less than 100%, but two tasks only if their utilisation is less than approximately 83%, and three if their utilisation is less than approximately 78%. Discuss whether these are necessary or sufficient conditions, illustrating your answer by reference to a system with three tasks, A, B and C with periods of 20ms, 60ms and 80ms respectively and run-times of 10ms, 21ms and 8 ms respectively.

[8 marks]

[End of paper]