Computer Science Department

Dummy Cover Sheet

B.Sc. Degree; B. Eng. Degree 1998

B229 Real Time Systems 1997-98

2 hours 30 minutes

*Part A of this paper is **compulsory**. You should attempt **all of question 1** in part A and any **THREE** questions from part B. All questions carry equal credit.*

The use of electronic calculators is permitted.

Part A

Compulsory question

Question 1

A prestigious university wishes to install an automatic video system in one of its lecture theatres so that lectures by eminent staff may be transmitted live to other institutions. To create the most favourable impact, it has been decided to use two camera systems to image the theatre podium, one a fixed camera capable of imaging the whole podium at once, the other a pan and tilt camera with variable zoom so that lecturers may be tracked in close-up. The tracking is to be performed automatically by processing images from both cameras 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 is 'real-time'.

[9 marks]

(b) If both cameras are high resolution, digital, colour cameras transmitting 768 x 1024 images at 24 bits/pixel in RGB format (ie: in red, green and blue colour channels, each encoded at 8 bits/pixel) at 60 frames/sec, estimate: (i) the amount of data for each image frame in MBytes, (ii) the data rate in Mbytes/sec for each camera, (iii) the bandwidth of the cables linking the cameras to the computer system, in GHz., if the data is transmitted serially on a single line connected to each camera. Explain your working.

[4 marks]

[Question 1 cont. over page]

[TURN OVER]

[Question 1 cont.]

(c) Tracking is based on two algorithms. The first requires computing the colour characteristics of pixels in successive images from both cameras. The second requires computing an additional cross-correlation of selected pixels of the desired colour in images from the moveable camera with a template obtained from the preceding image in the same camera. The first algorithm requires at most 10 arithmetic operations per pixel, whilst the second, the cross-correlation, on average, has to be carried out on 5% of the pixels, over a search region of 15 x 20 pixels. Estimate the processing rates required for each algorithm.

[6 marks]

(d) Comment on the implications of the processing rates you have estimated:

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

(ii) 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.

[6 marks]

[CONTINUED]

Part B

Attempt THREE questions from this section.

Question 2

(a) Define the z-transform and show how it may be used to represent: (i) filtering of a discrete, regularly sampled signal x[n] to produce an output y[n], (ii) recursive filtering of a similar signal, and (iii), show how the latter may be represented as a feedback loop.

[7 marks]

(b) A signal is to be processed by averaging it uniformly over every five samples, subsequently differentiating the smoothed signal by differencing suitable samples, and the result recombined with the original signal. Describe how you would choose an appropriate discrete differentiation operator, write down the z-transform for each of the processes, draw a diagram showing how they should be combined in series and parallel, and calculate the z-transform of the overall system transfer function.

[8 marks]

(c) Explain how the average in (b) above may be carried out recursively, and use the z-transform to prove that your recursive implementation is correct. Use your result to suggest a suitable recursive filter if the uniform average is replaced by one with relative weights 1,2,2,2,1 as sketched below and use the z-transform to prove that your filter is correct. Generalise your result to a filter of length n, with relative weights 1,2,2,…..2,2,1 and again use the z-transform to prove that your generalisation is correct.

[10 marks]

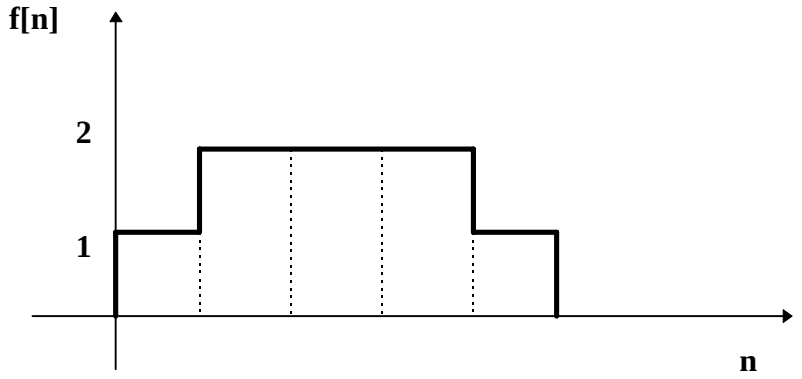

[TURN OVER]

Question 3

(a) Explain the difference between a pipelined multiprocessor system and a parallel multiprocessor system, and describe the advantages and disadvantages of each.

[8 marks]

(b) Explain the difference between a shared memory parallel architecture and a distributed memory parallel architecture, describe briefly how they are usually implemented, and discuss the advantages and disadvantages of each. Explain what is meant by an MIMD parallel system and describe why 'processor farms' are often used to implement MIMD parallel systems. Describe what a processor farm is, how it is implemented, and discuss what factors are important in the implementation of MIMD systems by means of processor farms.

[10 marks]

(c) Describe 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 25 ms respectively. Indicate what problems might frustrate your efforts, and how you might mitigate their effects.

[7 marks]

[CONTINUED]

Question 4

(a) (i) Explain the difference between independent, co-operating and competing tasks in a real-time computing system. (ii) Define what is meant by a semaphore and describe how it is manipulated.

[8 marks]

(b) Show, by writing appropriate pieces of 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. Indicate briefly what problems you might expect to encounter in using a semaphore in this way and what might be done to solve them.

[9 marks]

(c) 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.

[8 marks]

[TURN OVER]

Question 5

(a) Explain what is meant when a real-time computing system or task is described as 'hard' or 'soft'. Define what is meant when the tasks in a real-time system are said to be *schedulable* and describe how the characteristics of hard and soft real-time tasks lead to different approaches to satisfying their schedulability requirements.

[7 marks]

(b) 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.

[9 marks]

(c) 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 82.8%. Discuss whether either of these are necessary or sufficient conditions, illustrating your answer by reference to a system with two tasks, A and B, with periods of 20ms and 60ms respectively and run-times of 12ms and 21ms respectively.

[9 marks]

[CONTINUED]

Question 6

(a) Define what is meant by the reliability of a real-time computing system and explain how duplication of hardware and N-version programming can enhance the reliability of a real-time computing system. Include in your answer an explanation of why different procedures have to be adopted for hardware and software components. Explain carefully how many duplicates and versions are needed (i) to detect errors, and (ii) to ensure that the results computed are correct. Comment on use of the word 'ensure' in (ii).

[10 marks]

(b) Describe what is meant by 'fault-tolerant' operation of a real-time computing system and explain the difference between 'spatial' and 'temporal' fault-tolerance. Explain why some real-time systems may always exhibit a certain amount of temporal fault-tolerance.

[6 marks]

(c) Explain carefully why software is tested, and describe what is meant by (i) exhaustive, (ii) statistical and (iii) regression testing. Discuss why the testing of real-time system software is often more difficult than testing software developed for conventional computer systems.

[9 marks]

[End of paper]