COMPUTER SCIENCE TRIPOS Part II

Wednesday 7 June 2000 1.30 to 4.30

Paper 8

Answer **five** *questions.*
Submit the answers in five **separate** *bundles, each with its own cover sheet. On each cover sheet, write the numbers of* **all** *attempted questions, and circle the number of the question attached.*
Write on **one** *side of the paper only.*

1 Distributed Systems

In a large-scale distributed system, what advantages might be gained by using role-based access control (RBAC)? [2 marks]

A server which is an object manager implements RBAC by issuing principals with persistent capabilities which include the fields "role" and "object-name".

Describe how a signature field in the capability can be used to ensure its integrity. Define the protection the signature does and does not enforce. [6 marks]

How can such a persistent capability be made principal-specific? To what extent does this improve protection? [4 marks]

Discuss the desirability or otherwise of the persistence of capabilities and the use of persistent principal names. [4 marks]

Discuss revocation for the capability schemes suggested above. [4 marks]

2 Specification and Verification I

Describe how Floyd–Hoare logic can be used to define the semantics of programming languages (axiomatic semantics). Give advantages and disadvantages of this approach. [5 marks]

Describe the general idea of how program verification may be mechanised using verification conditions. Discuss what can be done completely automatically and what requires manual assistance. [5 marks]

Describe the difference between deep and shallow semantic embeddings of one language in another and compare and contrast the benefits of each. What is the point of embedding Floyd–Hoare logic in higher order logic? [5 marks]

Discuss possible commercial uses of formal software verification. How would you justify its use? [5 marks]

3 Digital Communication II

What information does a host require to transmit an IP datagram to a given destination IP address? [3 marks]

Describe how the host determines this information. How does this change in the presence of subnets? [3 marks]

Routing protocols can be classified as either *link-state* or *vector distance*.

(a) Briefly describe the operation of each of these schemes. [4 marks]

(b) In which case does the overall state converge more rapidly after a change? Justify your answer. [2 marks]

(c) Which scheme is more scalable? Justify your answer. [2 marks]

Sketch a design for a transport protocol providing reliable data transmission over IP multicast. What are the main issues here? Which changes within the network could lead to a more efficient implementation? [6 marks]

4 Comparative Architectures

Outline the design of a simple dynamic branch prediction cache which would provide a high degree of accuracy for branches that exhibit a strong bias in one or other direction (for example, loop-closing branches). [5 marks]

Give an example of a simple code sequence containing branches that would be mispredicted by your design, and describe how it would behave. [5 marks]

By storing a record of a branch's recent behaviour it is possible to devise prediction caches which are able to predict accurately branches whose behaviour follows some short regular pattern.

Revise your prediction cache design to make use of a four-bit "local history" for each entry. [5 marks]

How would this predictor perform with a branch exhibiting the following repeating sequences?
(`T` = taken, `N` = not-taken)

(a) `NTTNNTTNNTTN...`

(b) `NTTNNTNTTNNT...`

(c) `TNTTTTTNTTTT...`

[5 marks]

5 Business Studies

What are the differences between *profit and loss* and *cash flow* statements? [5 marks]

What are the differences between *debt* and *equity* finance? [5 marks]

What is an *option* and how might it be valued? [5 marks]

Comment on the current prices of high-tech stocks. [5 marks]

6 Information Retrieval

Simple natural language words are very basic index keys for retrieval systems. Since the topics for which users seek to retrieve documents are often complex, it is frequently claimed that simply looking for matches on any subset of a set of submitted single-word keys is not a discriminating enough approach to retrieval.

Describe *three* different ways in which the relations between words can be defined so as to form compound index terms. You should describe in detail how each is implemented and what its advantages and disadvantages are by comparison with (*a*) single terms, and (*b*) the other two types of compound.
[6 marks for each method]

Which would you implement in a general-purpose system, and why? [2 marks]

7 Artificial Intelligence

This question considers the monkey-and-bananas problem, in which there is a monkey in a room with some bananas hanging out of reach from the ceiling, but a box is available that will enable the monkey to reach the bananas if he climbs onto it. Initially the monkey is at location A, the bananas at B and the box at C. The monkey and box have height x, but if the monkey climbs onto the box he will have height y, the same as the bananas. The actions available to the monkey include Go from one place to another, Push an object from one place to another, Climb onto an object, and Grasp an object. Grasping results in holding the object if the monkey and object are in the same place at the same height.

(*a*) Write the initial state description using a representation of your choice.
[4 marks]

(*b*) Write definitions of the four actions, providing at least some obvious preconditions, additions and deletions. [10 marks]

(*c*) Suppose the monkey wants to fool the observers, who have gone to lunch, by grabbing the bananas but leaving the box in its original place. Write this as a goal (but not assuming the box is necessarily at location C) in the language of situation calculus. [3 marks]

(*d*) If the box is filled with bricks, its position will remain the same when the Push operator is applied. Is this an example of the frame problem or the circumscription problem? Justify your answer. [3 marks]

8 Neural Computing

(*a*) (*i*) Define *generalisation* in neural networks that learn from training data, and then are tested on new data. Why should not *all* the available data be used in the training set? [4 marks]

(*ii*) Draw a simple connectivity diagram that illustrates the idea of lateral inhibition in a competitive neural network. [2 marks]

(*iii*) With another diagram showing plots for input and output, illustrate how lateral inhibition in such a competitive network sharpens any input signal by effectively amplifying its first derivative. [2 marks]

(*iv*) What class of multi-layer neural network can be used to overcome the mathematical difficulties caused by intrinsic non-orthogonality of representation in many sensory and control systems? [2 marks]

(*b*) The study of neurological trauma to the brain gives clues about its modular organisation and specialisations of function, which may reveal some computational principles.

(*i*) What two fundamental principles of brain function did Karl Lashley's neurological research seem to reveal? [2 marks]

(*ii*) What are generally the differences between recovery prospects after a sudden brain trauma (in which all the damage is done at once), *versus* the same damage done more gradually (e.g. by a growing tumour)? [2 marks]

(*iii*) What mechanism might explain this difference? [2 marks]

(*iv*) Comment on its possible computational significance in terms of fault-tolerance, circuit adaptability and flexibility. [2 marks]

(*v*) Describe *two* different types of language-related disorders that can result from trauma to Broca's area or Wernicke's area, and comment on the computational inferences we might draw concerning language processing and linguistic representation. [2 marks]

9 Natural Language Processing

> The house was beautiful – Kim and Sandy especially liked the windows.
> They wanted to buy it, so they rang the estate agent.

The preceding narrative contains several types of anaphoric relation between pronouns or definite noun phrases and discourse antecedents. Identify them. [4 marks]

Describe how they might be recovered using either statistical or knowledge-based techniques. [6 marks]

What are the relative advantages and disadvantages of each approach? [10 marks]

10 Numerical Analysis II

A Riemann integral over $[a, b]$ is defined by

$$\int_a^b f(x)\,dx = \lim_{\substack{n\to\infty \\ \Delta\xi\to 0}} \sum_{i=1}^{n} (\xi_i - \xi_{i-1})\, f(x_i)$$

Explain the terms *Riemann sum* and *mesh norm*. [4 marks]

With respect to an integral over $[-1, 1]$ which of the following are *not* Riemann sums? Give explanations.

(a) $0.2f(-0.9) + 0.8f(-0.1) + 0.8f(+0.1) + 0.2f(+0.9)$

(b) $0.8f(-0.9) + 0.2f(-0.1) + 0.2f(+0.1) + 0.8f(+0.9)$

(c) $0.7f(-0.6) + 0.3f(-0.4) + 0.3f(+0.4) + 0.7f(+0.6)$

(d) $0.5f(-0.7) + 0.8f(0) + 0.5f(+0.7)$

(e) $0.3f(-0.7) + 1.0f(+0.1) + 0.7f(+0.7)$

[5 marks]

Suppose $\mathbf{R}$ is a rule that integrates constants exactly over $[-1, 1]$, and $f(x)$ is bounded and Riemann-integrable over $[a, b]$. Write down a formula for the composite rule $(n \times \mathbf{R})f$ and prove that

$$\lim_{n\to\infty} (n \times \mathbf{R})f = \int_a^b f(x)\,dx$$ [6 marks]

Which of the examples (a) to (e) converge in composite form? [2 marks]

Does the rule

$$-0.5f(-1) + 1.5f(-0.4) + 1.5f(+0.4) - 0.5f(+1)$$

converge in composite form? Comment on its suitability for this purpose.
[3 marks]

11 Information Theory and Coding

(*a*) If a continuous signal is discretely sampled by multiplying it with a sequence of uniformly-spaced Dirac delta functions, having frequency f_s, what happens to the Fourier spectrum of the signal? [3 marks]

(*b*) What is the conditional probability $p(x|y)$, the probability of event x given that event y has occurred, provided that we know the following?

$p(x)$, the unconditional probability of event x
$p(y)$, the unconditional probability of event y
$p(y|x)$, the probability of event y given that event x has occurred

[3 marks]

(*c*) Consider a binary symmetric communication channel, whose input source is the alphabet $X = \{0, 1\}$ with probabilities $\{0.5, 0.5\}$; whose output alphabet is $Y = \{0, 1\}$; and whose channel matrix is

$$\begin{pmatrix} 1-\epsilon & \epsilon \\ \epsilon & 1-\epsilon \end{pmatrix}$$

where ϵ is the probability of transmission error.

(*i*) What is the entropy of the source, $H(X)$? [1 mark]

(*ii*) What is the probability distribution of the outputs, $p(Y)$, and the entropy of this output distribution, $H(Y)$? [3 marks]

(*iii*) What is the joint probability distribution for the source and the output, $p(X, Y)$, and what is the joint entropy, $H(X, Y)$? [3 marks]

(*iv*) What is the mutual information of this channel, $I(X; Y)$? [2 marks]

(*v*) How many values are there for ϵ for which the mutual information of this channel is maximal? What are those values, and what then is the capacity of such a channel in bits? [3 marks]

(*vi*) For what value of ϵ is the capacity of this channel minimal? What is the channel capacity in that case? [2 marks]

12 Computer Vision

Explain the Bayesian approach to solving problems in computer vision. Explain the notion of an *Inverse Problem* and how computer vision can be regarded thereby in a formal sense as *inverse graphics*. Write down Bayes' rule in general form, and explain the interpretation of its terms as:

- probability of the image, given the object
- probability of the object, given the image

What is the role of the "prior?"

Discuss and illustrate the Bayesian approach in terms of 3D surface reconstruction, given the reflectance data in an image.

[20 marks]

13 Specification and Verification II

Describe the syntax and semantics of Computation Tree Logic (CTL). [4 marks]

Write down CTL formulae expressing the following properties:

(*a*) *if* req *goes high then it will stay high until* ack *goes high and then go low on the next cycle* [2 marks]

(*b*) *if ever* req *is high and* started *is low then eventually* error *will become permanently high* [2 marks]

Briefly describe the main ideas underlying model checking. What is symbolic model checking? [4 + 4 marks]

Compare and contrast the use of a model checker and a theorem prover for formal verification. [4 marks]

END OF PAPER