COMPUTER SCIENCE TRIPOS Part II

Wednesday 4 June 2014 1.30 to 4.30 pm

COMPUTER SCIENCE 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.*

You may not start to read the questions printed on the subsequent pages of this question paper until instructed that you may do so by the Invigilator

STATIONERY REQUIREMENTS	SPECIAL REQUIREMENTS
Script paper	*Approved calculator permitted*
Blue cover sheets	
Tags	
Rough work pad	

1 Advanced Graphics

(*a*) From a uniform knot vector, derive the basis functions of a uniform quadratic B-spline. That is, derive $N_{1,3}$ from the knot vector and also show how $N_{i,3}$ is related to $N_{1,3}$ for arbitrary i. [7 marks]

(*b*) It is known that the uniform quadratic B-spline curve is continuous in its first derivative but that it is not guaranteed continuous in its second derivative. Prove that $N_{1,3}$ is discontinuous in its second derivative at one or more points. [4 marks]

(*c*) The Chaikin corner-cutting subdivision scheme is related to the uniform quadratic B-spline in that the limit curve of the Chaikin scheme is the uniform quadratic B-spline curve generated from the original control points, $(P_0^0, P_1^0, P_2^0, \ldots)$.

Given a set of control points, $(P_0^n, P_1^n, P_2^n, \ldots)$, at subdivision level n, the Chaikin scheme generates a new set of control points at level $n+1$ by two rules:

$$\begin{aligned} P_{2i}^{n+1} &= \tfrac{3}{4}P_i^n + \tfrac{1}{4}P_{i+1}^n \\ P_{2i+1}^{n+1} &= \tfrac{1}{4}P_i^n + \tfrac{3}{4}P_{i+1}^n \end{aligned}$$

Consider the sequence $P_i^0, P_{2i}^1, P_{4i}^2, \ldots, P_{2^n i}^n, \ldots$. From your answer to part (*a*), or otherwise, determine $\lim_{n\to\infty} P_{2^n i}^n$ in terms of the original control points. [4 marks]

(*d*) The univariate Chaikin curve scheme described in part (*c*) can be generalised to a bivariate scheme that generates surfaces. Explain how it can be generalised to generate surfaces from an arbitrary mesh of control points, paying attention to both the regular and the extraordinary cases. [5 marks]

2 Artificial Intelligence II

Consider a *reinforcement learning* problem having states $\{s_1, \ldots, s_n\}$, actions $\{a_1, \ldots, a_m\}$, reward function $R(s, a)$ and next state function $S(s, a)$.

(*a*) Give a general definition of *discounted cumulative reward*, a *policy*, and an *optimal policy* for a problem of this kind. [5 marks]

(*b*) Give a detailed derivation of the *Q-learning algorithm*. [5 marks]

(*c*) In the reinforcement learning problem shown in the diagram, states are positions on a grid and actions are `down` and `right`. The initial state is s_1. The only way an agent can receive a (non-zero) reward is by moving into one of two special positions, one of which has reward -50 and the other 150.

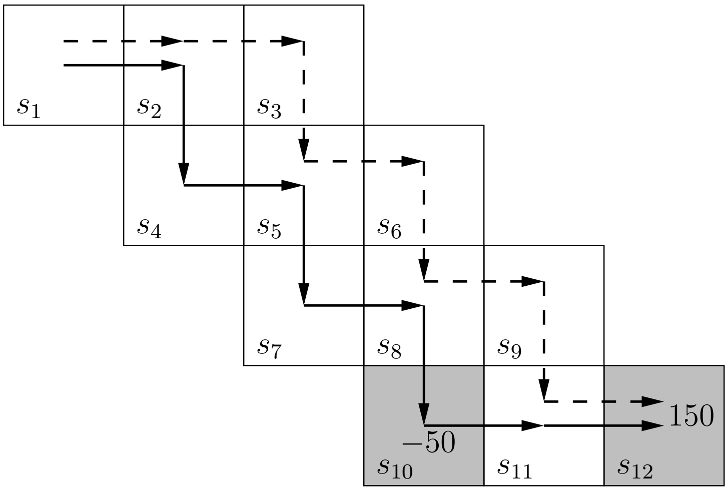

A possible sequence of actions (sequence 1) is shown by solid arrows, and another (sequence 2) by dashed arrows. Assume that all Q values are initialised at 0. Explain how the Q values are modified by the Q-learning algorithm if sequence 1 is used *once*, followed by *two* uses of sequence 2, and then *one* final use of sequence 1. [10 marks]

3 Comparative Architectures

(*a*) Throughout the 1990s mainstream microprocessors were developed with ever deeper pipelines. Since then manufacturers have scaled back to more moderate pipeline depths.

(*i*) What were the benefits from implementing deep pipelines and why were they scaled back? [4 marks]

(*ii*) How do pipelines that support in-order and out-of-order execution differ in their microarchitectural components? [4 marks]

(*b*) Modern high-performance processors incorporate a dynamic branch predictor to avoid stalling when branches are fetched.

(*i*) What is a tournament branch predictor and why might it outperform either a global or local branch predictor alone? [4 marks]

(*ii*) You develop a new branch predictor that is significantly more accurate than existing designs. However, its complexity means that it takes several cycles to produce a prediction. How can you make use of this predictor without always introducing a pipeline bubble? [4 marks]

(*iii*) If you were designing an out-of-order core, why might you decide not to allow predicated execution? [4 marks]

4 Computer Systems Modelling

Consider a birth death process $X(t)$ for $t \geq 0$ with states $0, 1, 2, \ldots$ and where the birth rate is λ_j in states $j = 0, 1, \ldots$ and the death rate is μ_j in states $j = 1, 2, \ldots$

(*a*) Draw the state space diagram for the birth death process. [2 marks]

(*b*) Derive the Chapman-Kolmogorov differential equations for the birth death process in terms of $P_j(t) = \mathbb{P}(X(t) = j)$ for $j = 0, 1, \ldots$ and $t \geq 0$. [4 marks]

(*c*) State the detailed balance conditions for an equilibrium distribution $p_j = P_j(t)$ for $j = 0, 1, \ldots$ to solve the Chapman-Kolmogorov equations with $\frac{dP_j(t)}{dt} = 0$ and determine a further condition to ensure the existence of the equilibrium distribution. Derive an expression for equilibrium distribution p_j when your further condition holds. [4 marks]

(*d*) Describe a birth death process model for an $M/M/K$ system in the limit as the number of servers $K \to \infty$. Draw the state space diagram, give the birth and death rates and derive the equilibrium distribution stating whether there are any conditions for its existence. [5 marks]

(*e*) Now consider the Chapman-Kolmogorov equations derived in part (*b*) in the special case of a pure birth process with constant birth rates $\lambda_j = \lambda$ for $j = 0, 1, \ldots$ and zero death rates $\mu_j = 0$ for $j = 1, 2, \ldots$ Suppose that the process starts in state 0 at time $t = 0$ (that is $P_0(0) = 1$ and $P_j(0) = 0$ for $j = 1, 2, \ldots$). Thus $X(t)$ is the number of events in a Poisson process of rate λ. Determine the Chapman-Kolmogorov differential equations for the time-dependent solution $P_j(t)$ for $j = 0, 1, 2, \ldots$ and $t \geq 0$ in this case and solve for the explicit solutions $P_j(t)$ obeying the initial conditions. [5 marks]

5 Computer Vision

(*a*) Present five experimental observations about human vision that support the thesis that "vision is graphics": what we see is explicable only partly by the optical image itself, but is more strongly determined by top-down knowledge, model-building and inference processes. [5 marks]

(*b*) Consider the following pair of (6×6) filter kernels:

-1	-1	2	2	-1	-1
-1	-3	4	4	-3	-1
-1	-4	5	5	-4	-1
-1	-4	5	5	-4	-1
-1	-3	4	4	-3	-1
-1	-1	2	2	-1	-1

1	-1	-1	1	1	-1
1	-2	-3	3	2	-1
1	-3	-4	4	3	-1
1	-3	-4	4	3	-1
1	-2	-3	3	2	-1
1	-1	-1	1	1	-1

(*i*) Why do these two kernels constitute a quadrature pair? [2 marks]

(*ii*) To what kinds of image structure, and which orientations, are these detector kernels most sensitive? [2 marks]

(*iii*) How would these kernels be applied directly to an image for filtering or feature extraction? [2 marks]

(*iv*) How could their respective Fourier Transforms alternatively be applied to an image, to achieve the same effect as in (*iii*) but faster? [2 marks]

(*v*) What is the "DC" response of each of the kernels, and what is the significance of this? [2 marks]

(*vi*) How could these kernels be combined to locate facial features? [2 marks]

(*c*) What information about the shape and orientation of an object can be inferred, and how, from the extraction of texture descriptors; and what is the role of prior assumptions in making such inferences from texture? [3 marks]

6 Digital Signal Processing

(*a*) Consider a causal, order-2 digital filter with real-valued infinite impulse response sequence $h_0, h_1, h_2, \ldots$

(*i*) What is the z-transform $H(z)$ of this filter's impulse response? [2 marks]

(*ii*) Express $H(z)$ in terms of the locations c_1, c_2 of its two zeros and the locations d_1, d_2 of its two poles in $\mathbb{C}$. [4 marks]

(*iii*) Give a necessary condition for c_1, c_2, d_1, d_2 to ensure that $\{h_n\}$ has only real values. [4 marks]

(*iv*) If we operate that filter at sampling frequency f_s, what will its amplitude gain at frequency f be? [2 marks]

(*b*) A *notch filter* aims to suppress a single frequency f_c. One way of designing an order-2 notch filter, as in part (*a*), involves placing the zeros directly onto the unit circle, and the poles right next to them inside the unit circle, at distance $0 < \alpha < 1$ from 0:

$$c_1 = \mathrm{e}^{\mathrm{j}\omega}, \quad d_1 = \alpha \cdot c_1, \quad c_2 = \mathrm{e}^{-\mathrm{j}\omega}, \quad d_2 = \alpha \cdot c_2, \quad \text{with} \quad \omega = 2\pi f_\text{c}/f_\text{s}$$

(*i*) What is the z-transform of the impulse response of the resulting filter, written as a fraction of two polynomials of z^{-1}? [4 marks]

(*ii*) The *OxyMax* is a medical device designed in the United States. It processes a heart-beat signal with a sampling rate of $f_\text{s} = 600$ Hz. It contains the following C function, which implements a notch filter, as in part (*b*)(*i*), to suppress in the input signal interference from the North American power grid at $f_\text{c} = 60$ Hz:

```
double mains_notch(double sample) {
  static double x[4], y[4];
  static int n = 0;
  x[n&3] = sample;
  y[n&3] = sample + x[(n-1)&3] * b1 + x[(n-2)&3]
                  - y[(n-1)&3] * a1 - y[(n-2)&3] * a2;
  return y[n++&3];
}
```

The U.S. version initializes the constants used with `b1` $= -2\cos(\pi/5)$, `a1` $=$ `b1` $\times\, 0.9$ and `a2` $= 0.81$. What changed constant(s) will instead suppress the power-grid frequency at $f_\text{c} = 50$ Hz for the European version? [4 marks]

7 E-Commerce

(*a*) Define “Fair market value”. [2 marks]

(*b*) For a fair market do the participants need to be anonymous before the transaction is agreed? [4 marks]

(*c*) Are anonymity and reputation contradictory? [4 marks]

(*d*) Are fully anonymous online markets desirable? Justify your answer. [10 marks]

8 Temporal Logic and Model Checking

From Wikipedia: "Tic-tac-toe (or Noughts and Crosses, `X`s and `O`s) is a paper-and-pencil game for two players, `X` and `O`, who take turns marking the spaces in a 3×3 grid. The player who succeeds in placing three respective marks in a horizontal, vertical, or diagonal row wins the game." For example, `X` is the first player in both example games shown below; the first game is won by the `X`, the second is drawn.

```
 | |X   O| |X   O| |X   O| |X   O| |X   O| |X   O| |X
-|-|-   -|-|-   -|-|-   -|-|-   -|-|-   -|-|-   -|-|-
 | |     | |     | |     |O|     |O|     |O|O    |O|O
-|-|-   -|-|-   -|-|-   -|-|-   -|-|-   -|-|-   -|-|-
 | |     | |    X| |    X| |    X| |X   X| |X   X|X|X

 | |     | |O    | |O   O| |O   O|X|O   O|X|O   O|X|O   O|X|O   O|X|O
-|-|-   -|-|-   -|-|-   -|-|-   -|-|-   -|-|-   -|-|-   -|-|-   -|-|-
 |X|     |X|     |X|     |X|     |X|     |X|    X|X|    X|X|O   X|X|O
-|-|-   -|-|-   -|-|-   -|-|-   -|-|-   -|-|-   -|-|-   -|-|-   -|-|-
 | |     | |     | |X    | |X    | |X    |O|X    |O|X    |O|X   X|O|X
```

(*a*) This part of the question asks you to define a Kripke structure $M = (S, S_0, R, L)$ to model Tic-tac-toe. Assume the set AP consists of atomic propositions `Start`(p) and `Has`(i, v). `Start`(p) means player p starts, where $p \in \{0, 1\}$ represents a player: 0, 1 represent `O`, `X`, respectively. `Has`(i, v) means space i contains value v, where $i \in \{1, \ldots, 9\}$ names a grid space and $v \in \{0, 1, 2\}$ represents the state of a space: 0, 1, 2 represent `O`, `X`, empty-space, respectively.

(*i*) Specify a suitable representation S of states. [2 marks]

(*ii*) Specify the set of initial states S_0. [2 marks]

(*iii*) Specify a transition relation R to model the moves in the game. [6 marks]

(*iv*) Specify a labelling function L to define which atomic propositions hold in each state. [2 marks]

(*b*) In a suitable temporal logic, which you should name, devise and explain a formula ψ such that $M \models \psi$ if and only if the first player can always win or draw, no matter how the second player plays. [8 marks]

9 Information Retrieval

In the inverted index of an information retrieval system, dictionary terms can be represented using different data structures.

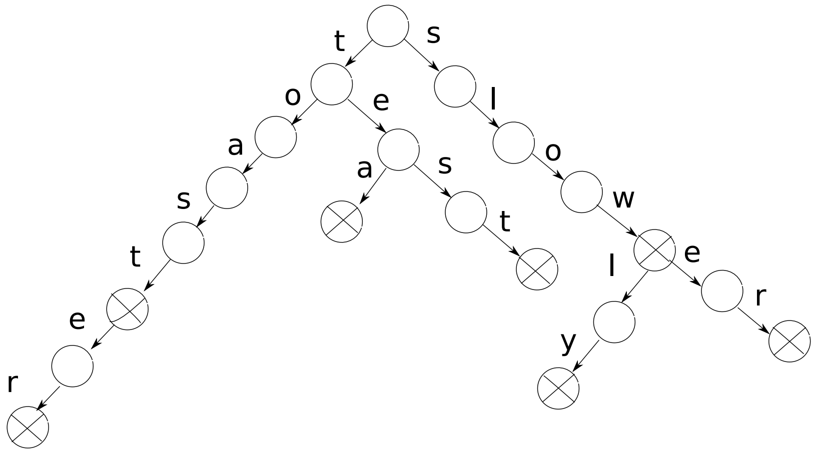

(*a*) Consider the *trie* in the figure above, which encodes several dictionary terms.

(*i*) List the terms contained in this trie. [2 marks]

(*ii*) Explain how terms are looked up in a trie. [2 marks]

(*b*) Alternatively, we could store the terms in a *binary search tree.*

(*i*) Draw the binary search tree with minimal depth that stores the dictionary terms from the figure above. [3 marks]

(*ii*) Compare the worst-case time complexity of dictionary lookup for a binary tree and a trie. What are the conditions where the binary tree is preferable to a trie? [3 marks]

(*c*) Next consider a *radix tree*, a space-optimised trie data structure where each internal node with only one child is merged with its child. (An internal node is one not associated with a term, and thus not pointing to any data.)

(*i*) Draw the radix tree containing the dictionary terms from the figure above. [2 marks]

(*ii*) Give an algorithm for insertion of a new index term $t = t(0) \dots t(k)$ into a radix tree. Use examples to illustrate your algorithm. You may use pseudocode as long as you clearly explain your thoughts. [8 marks]

10 Principles of Communications

(*a*) In distributed routing in networked systems such as the Internet, there may be intermittent faults which could cause the choice of routes to osciliate between one path and another, leading to undesirable consequences.

(*i*) What problems would rapid route change cause? [5 marks]

(*ii*) How might a feedback control system be used to damp such oscillations? [5 marks]

(*b*) A network provider deploys Explicit Notification (ECN) capability in their routers. End systems running TCP can take advantage of this to trigger Congestion Control behaviour in TCP when receiving packets with ECN marks. A suggestion that the TCP congestion control scheme to date, using Additive Increase and Multiplicative Decrease (AIMD) could be replaced by a proportional-integral-derivative controller (PID controller) is made.

Describe how such a controller could operate to adjust TCP's congestion control window in response to ECN marked packets in qualitative terms. Explain any assumptions that might have to be made about the bottleneck router marking packets and what the advantages of a PID controller over AIMD might then accrue. [10 marks]

11 Security II

(*a*) Let $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ be a private-key encryption scheme that operates on fixed-length messages from $\mathcal{M} = \{0,1\}^m$. Briefly explain a game that a user $\mathcal{U}$ of Π must be able to win against any polynomial-time adversary $\mathcal{A}$ with probability $\frac{1}{2} - \epsilon$ (where is ϵ is "negligible" with growing key length) for Π to be able to claim to offer "indistinguishable multiple encryptions under chosen-plaintext attack" (CPA security). [8 marks]

(*b*) Let $\Pi = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ be a private-key encryption scheme that offers CPA security and operates on fixed-length messages $M \in \mathcal{M} = \{0,1\}^m$ with keys $K \in \mathcal{K} = \{0,1\}^\ell$. We use it to construct a new encryption scheme $\Pi' = (\mathsf{Gen}, \mathsf{Enc}', \mathsf{Dec}')$. In which of the following cases is Π' also CPA secure? Explain your answer. [2 marks each]

(*i*) $\mathsf{Enc}'_K(M) = \mathsf{Enc}_K(M \oplus 1^m)$

(*ii*) $\mathsf{Enc}'_K(M) = \mathsf{Enc}_K(M) || \mathrm{LSB}(M)$

(*iii*) $\mathsf{Enc}'_K(M) = \mathsf{Enc}_K(M) || \mathrm{LSB}(K)$

(*iv*) $\mathsf{Enc}'_K(M) = \mathsf{Enc}_{0^\ell}(M)$

[*Note:* LSB outputs the least significant bit of its input word, || is concatenation.]

(*c*) While reviewing an implementation of AES-CBC, you discover that it simply uses the last ciphertext block from the previously encrypted message as the IV value C_0 for encrypting the next message. The implementation's author argues that as long as the IV of the very first message was chosen uniformly at random, all resulting subsequent ciphertext blocks will also be distributed uniformly at random, and therefore make good IVs. Why is this construction nevertheless not CPA secure? [4 marks]

12 System-on-Chip Design

Programmed I/O, also known as memory-mapped I/O, uses a processor's load and store instructions to read and write peripheral registers.

(*a*) How can programmed I/O adversely interact with compiler optimisations and how can this be avoided? [2 marks]

(*b*) How can programmed I/O adversely interact with cache subsystems and how can this be avoided? [2 marks]

(*c*) How might the memory addresses for progammed I/O typically be decided and how are these decisions typically embodied in hardware and in software? [3 marks]

(*d*) Give two code fragments that both implement a pair of registers in a peripheral device that, perversely, when one register is read (e.g. by programmed I/O) it returns the value last written to the other.

(*i*) The first fragment should be for a physical implementation (use synthesisable RTL/SystemC or a schematic diagram), [3 marks]

(*ii*) The second fragment should be for a transactional model. [3 marks]

(*e*) Estimate, to the nearest order of magnitude, how many instructions are needed by the simulator when modelling a single programmed I/O write to a device register using the following two modelling styles:

(*i*) the system bus and the peripheral device modelled as a gate-level net list; [4 marks]

(*ii*) the system bus and the peripheral device hand-coded as a cycle-accurate model. [3 marks]

You may disregard the modelling of the initiating CPU core and answer only for the peripheral device connected by a simple bus. Greater credit will be awarded for arguments supporting your answers than for the figures themselves. State assumptions as to whether the modelling language is interpreted, compiled or otherwise in each case.

13 Topical Issues

Retailers anticipate using Bluetooth Low Energy (BLE) beacons distributed around a store to provide location-based marketing information. The beacons regularly send their unique identifier via BLE advertisements. A smartphone that receives an identifier can then look up the related information to display.

(*a*) State and explain three optimisations in the BLE technology that would allow the beacons to operate on a single coin cell for many months or years. [6 marks]

(*b*) If the smartphone is given a spatial map of the beacons, it can also provide customer tracking around the store.

(*i*) Explain how signal fingerprinting could be used to track customers in this context. Explain how fingerprints are formed and include an example matching metric in your answer. [4 marks]

(*ii*) Give three advantages of BLE fingerprinting over the more traditional WiFi fingerprinting. [3 marks]

(*iii*) Discuss the practical issues you would expect to see in a BLE fingerprint system for retail. [7 marks]

END OF PAPER