COMPUTER SCIENCE TRIPOS Part II

Thursday 4 June 2009 1.30 to 4.30

COMPUTER SCIENCE Paper 9

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
Script paper
Blue cover sheets
Tags

SPECIAL REQUIREMENTS
Approved calculator permitted

1 Additional Topics

(*a*) "The future that ubicomp has been attempting to build is not our own future, but 1989's future—yesterday's tomorrows." (Bell & Dourish, 2007).

Explain what is meant by this quotation, using examples where appropriate. [5 marks]

(*b*) A company wishes to have a camera in a video-conferencing suite to track the current speaker at any time. To do so they propose distributing a set of microphones around the room, each connected to a single computer via a set of sound cards. They intend to use a Time Difference of Arrival (TDOA) procedure to locate the speaker.

(*i*) What is the minimum number of microphones needed, and how would you distribute them spatially to get the best results? [2 marks]

(*ii*) Each microphone input is sampled simultaneously for a time window of length T. Give *two* factors that affect the choice of T. [3 marks]

(*iii*) Given matching windows of samples for two microphones, how would you compute the time differences of the signals? [2 marks]

(*iv*) How does the sampling rate for the signals affect the accuracy of each TDOA estimate? Assuming the audio signal of interest has a frequency of 22kHz, estimate the largest error that could be associated with a single TDOA estimate. [5 marks]

(*v*) Suggest *two* other error sources likely to contribute to the ultimate error in position. What would you expect to be the dominant error source? [3 marks]

2 Advanced Systems Topics

(*a*) Distributed storage systems can typically be divided into *network attached storage* (NAS) and *storage area networks* (SAN).

(*i*) Describe with the aid of a diagram the operation of a typical NAS system. Use as an example the access of a file by a client system. [3 marks]

(*ii*) Describe with the aid of a diagram the operation of a typical SAN system. Use as an example the access of a file by a client system. [3 marks]

(*iii*) Which would be more suitable for a high-performance database system? Justify your answer. [2 marks]

(*b*) Database systems often use *indexes* in order to accelerate certain operations.

(*i*) What exactly is an index used for? [1 mark]

(*ii*) Sequential indexes can be either sparse or dense. Give *two* advantages of sparse indexes and *two* advantages of dense indexes. [4 marks]

(*iii*) The B+-tree is a commonly used data-structure for implementing indexes. Sketch the structure of a B+-tree, and describe how lookup and insertion occur. [4 marks]

(*iv*) A related data-structure is the B-tree. What are the differences between B-trees and B+-trees? [2 marks]

(*v*) Why are B+-trees typically preferred over B-trees in database systems? [1 mark]

3 Bioinformatics

(*a*) Compute the global alignment between the two strings s1 = ACCGTT and s2 = AGTTCA, considering the following scoring parameters: +1 for match, −1 for mismatch, and −1 for a gap.

(*i*) What is the maximum similarity score between the two sequences s1 and s2? [2 marks]

(*ii*) Find an alignment with this similarity score. [2 marks]

(*iii*) Is the alignment you found unique, or are there multiple alignments achieving the maximum similarity score? [1 mark]

(*b*) Discuss the complexity of the Sankoff parsimony algorithm. [4 marks]

(*c*) Discuss the main differences between K-means, Superparamagnetic and Markov clustering algorithms. [7 marks]

(*d*) Discuss the utility of the Gillespie algorithm in system biology. [4 marks]

4 Computer Systems Modelling

Consider the birth death process model for a $M/M/1$ queue with arrival rate $\lambda > 0$, service rate $\mu > 0$ such that the traffic intensity $\rho = \lambda/\mu < 1$. Let p_k for $k = 0, 1, 2, \ldots$ be the equilibrium distribution for the number of jobs, k, in the queue.

(*a*) By considering transitions into and out of a given state i construct the *global balance* conditions for the equilibrium distribution. [5 marks]

(*b*) By considering the transitions between a pair of adjacent states i and $i + 1$ construct the *detailed balance* conditions for the equilibrium distribution. [5 marks]

(*c*) Solve the detailed balance conditions to derive the equilibrium distribution when $\rho < 1$. [5 marks]

(*d*) Show that your solution for the equilibrium distribution derived in part (*c*) also solves your global balance conditions given in part (*a*). [5 marks]

5 Computer Vision

(*a*) Consider the following two contrasting kinds of filter kernels, A and B:

A

0	-1	-1	-1	-1	0
-1	-2	-3	-3	-2	-1
-1	-3	12	12	-3	-1
-1	-3	12	12	-3	-1
-1	-2	-3	-3	-2	-1
0	-1	-1	-1	-1	0

B

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*) Given that the sum of all taps in each kernel is zero, how will each filter respond to an image region having only uniform brightness? [1 mark]

(*ii*) Which filter has much broader bandwidth in spatial frequency? [1 mark]

(*iii*) Categorise each filter as being either essentially isotropic or anisotropic, and explain the significance of these terms. [2 marks]

(*iv*) Which filter is better described as an oriented edge detector? What orientation of edges is it best able to detect? [2 marks]

(*v*) Apply the terms "$\nabla^2 G_\sigma(x, y)$" or "Imaginary part of a 2D Gabor wavelet" as you think most appropriate to each filter. [2 marks]

(*b*) What is accomplished by the lateral signal flows within both plexiform layers of the mammalian retina, in terms of image processing and coding? [4 marks]

(*c*) Consider the task of "anti-spoofing" in the automated visual recognition of iris patterns: how to confirm that an image is acquired from a living iris instead of from a spoofing artefact like a photograph or a fake printed contact lens. How would you implement strategies (*i*)–(*ii*), and detect properties (*iii*)–(*iv*)?

(*i*) 3D shape description using inferences from stereo, shading information, or structured light and the fact that a real iris is planar whereas a printed contact lens on the cornea floats on a spherical surface. [2 marks]

(*ii*) Motion detection tuned for radial changes in pupil size with iris pattern stretching and light-driven deformations. [2 marks]

(*iii*) Lambertian properties and dynamic specular reflections. [2 marks]

(*iv*) Photonic properties of living tissue compared with artificial objects, including reflectance maps and wavelength-dependent absorption spectra. [2 marks]

6 Denotational Semantics

(*a*) (*i*) Define the notion of *contextual equivalence* in PCF. [2 marks]

(*ii*) Consider the following two closed PCF terms of type $nat \to nat \to nat$.

$$\begin{aligned} F \stackrel{\text{def}}{=} \mathbf{fix}\big(\ & \mathbf{fn}\, f : nat \to nat \to nat.\ \mathbf{fn}\, x : nat.\ \mathbf{fn}\ y : nat. \\ & \quad \mathbf{if\ zero}(x)\ \mathbf{then\ 0} \\ & \quad \mathbf{else\ if\ zero}(y)\ \mathbf{then\ 0} \\ & \quad \mathbf{else\ succ}\big(\, f\,(\mathbf{pred}\ x)\,(\mathbf{pred}\ y)\,\big)\ \big) \end{aligned}$$

$$\begin{aligned} G \stackrel{\text{def}}{=} \mathbf{fix}\big(\ & \mathbf{fn}\, g : nat \to nat \to nat.\ \mathbf{fn}\, x : nat.\ \mathbf{fn}\ y : nat. \\ & \quad \mathbf{if\ zero}(y)\ \mathbf{then\ 0} \\ & \quad \mathbf{else\ if\ zero}(x)\ \mathbf{then\ 0} \\ & \quad \mathbf{else\ succ}\big(\, g\,(\mathbf{pred}\ x)\,(\mathbf{pred}\ y)\,\big)\ \big) \end{aligned}$$

Are F and G contextually equivalent? Justify your answer. [5 marks]

(*b*) (*i*) Define a closed PCF term $H : (nat \to nat \to nat) \to nat \to nat \to nat$ such that $[\![\mathbf{fix}(H)]\!] \in \big(\mathbb{N}_\perp \to (\mathbb{N}_\perp \to \mathbb{N}_\perp)\big)$ satisfies

$$[\![\mathbf{fix}(H)]\!]\,(i)\,(j) \;=\; \max(i, j)$$

for all $i, j \in \mathbb{N}$. [4 marks]

(*ii*) Let

$$S \stackrel{\text{def}}{=} \{\, f \in \big(\mathbb{N}_\perp \to (\mathbb{N}_\perp \to \mathbb{N}_\perp)\big) \mid f(x)(y) = f(y)(x) \text{ for all } x, y \in \mathbb{N}_\perp \,\}$$

Show that the subset $S \subseteq \big(\mathbb{N}_\perp \to (\mathbb{N}_\perp \to \mathbb{N}_\perp)\big)$ is admissible. [4 marks]

(*iii*) Show that

$$[\![\mathbf{fix}(H)]\!]\,(x)\,(y) \;=\; [\![\mathbf{fix}(H)]\!]\,(y)\,(x)$$

for all $x, y \in \mathbb{N}_\perp$. [5 marks]

[Hint: Use Scott's Fixed-Point Induction Principle.]

7 Digital Communication II

(*a*) Discuss the trade-off in state and the trade-off in protocol headers that is exhibited by protocols for cell switching (e.g. ATM) and packet switching (e.g. IP). [10 marks]

(*b*) What properties contribute to scaling in the use of hierarchy in allocation and treatment of identifiers such as the Internet's IP addresses and ATM's VCIs? [10 marks]

8 System-on-Chip Design

Unless otherwise stated, use any appropriate combination of text, diagrams, SystemC and/or Register Transfer Language (RTL) code in your answers.

(*a*) Give a programming model for a simple DMA (direct memory access) controller with one control/status register and three operand registers: for block length, source address and destination address.
[The DMA controller, when active, becomes a bus master and copies a block of data from one area to another.] [4 marks]

(*b*) Give the precise condition for your DMA controller to generate an interrupt. [2 marks]

(*c*) Sketch the implementation (in C and/or assembly language) for a general-purpose block-copy function that uses the DMA controller. Include the steps of enabling and disabling interrupts. Assume all copies are an integer number of words and are word-aligned. [4 marks]

(*d*) Give a programming model for an audio sub-system that includes a modified DMA controller that is dedicated to copying audio samples to and from a sound codec (ADC (Analogue/Digital Converter) and DAC (Digital/Analogue Converter) pair). [6 marks]

(*e*) How must the buffering in the audio-subsystem of part (*d*) be dimensioned and how are interrupts used to maintain continuous streams of audio at the correct sample rate? [4 marks]

9 Information Theory and Coding

(*a*) Calculate the entropy in bits for each of the following random variables:

(*i*) Pixel values in an image whose possible grey values are all the integers from 0 to 255 with uniform probability. [1 mark]

(*ii*) Humans grouped by whether they are, or are not, mammals. [1 mark]

(*iii*) Gender in a tri-sexual insect population whose three genders occur with probabilities 1/4, 1/4, and 1/2. [2 marks]

(*iv*) A population of persons classified by whether they are older, or not older, than the population's median age. [1 mark]

(*b*) Let $p(x)$ and $q(x)$ be two discrete probability distributions.

(*i*) What is the *Kullback–Leibler distance* (*KL*) between these distributions? [2 marks]

(*ii*) If we have devised an optimally compact code for the random variable described by $q(x)$, what does the *KL* tell us about the effectiveness of our code if the probability distribution is $p(x)$ instead of $q(x)$? [1 mark]

(*iii*) Which axiom of distance metrics is violated by this distance? [1 mark]

(*iv*) What happens to this metric if there are some forbidden values of x for which $p(x) = 0$, and other values of x for which $q(x) = 0$? [1 mark]

(*c*) Consider a continuous real signal whose bandwidth extends from 0 to 1 KHz. We wish to represent a 1.0 second interval of it exactly, using just a finite list of numbers obtained by sampling the signal at discrete, periodic, points in time.

(*i*) What is the length of the shortest list of such discrete samples needed to reconstruct this interval of the signal from them completely? [2 marks]

(*ii*) Name, define, and sketch a plot of the function used to reconstruct the continuous signal from its samples, by superimposing shifted copies of this function weighted by the discrete samples. [3 marks]

(*d*) Explain why data can be compressed by encoding it into transforms (such as the DCT, Fourier or Gabor) that result in coefficients that have a more narrow, peaked, distribution than the original data. Without going into details about particular transforms, explain why the coefficients obtained have distributions with less entropy than the original signal or image, and why this enables compression. [5 marks]

10 Optimising Compilers

Given a flowgraph G representing a function f having parameters $(x_1, \ldots, x_k)$, we are interested in when and how its parameters are used, in particular when we can guarantee that one or more of its variables are read during its execution. You may assume G has conditionals, assignment and arithmetic operators but no function calls.

(*a*) Representing dataflow properties as the set of variables that must inevitably be read when execution starts at node n, give dataflow equations for the sets *mustuse*(n). [4 marks]

(*b*) Now, supposing this analysis is not precise enough for our needs, we refine the dataflow properties to be *mustuseoneof*(n)—the set of sets of variables so that, for example, $\{\{v_1, v_2\}, \{v_3\}\}$ means that "every path starting at n must either read v_3 or read both v_1 and v_2; no other possibility exists". Give dataflow equations for *mustuseoneof*(n). [Hint: you may need to generalise the traditional $(\cdots \setminus \mathit{kill}_n \cup \mathit{gen}_n)$ to e.g. $\mathit{gen}_n(\mathit{kill}_n(\cdots))$.] [4 marks]

(*c*) Assuming that n_0 is the entry node of G and that the values of its arguments have already been stored in its parameters $(x_1, \ldots, x_k)$, relate *mustuseoneof*(n_0) to the strictness function $f^\sharp$ corresponding to f. For this purpose you may now assume G to be appropriately restricted for a lazy functional language (e.g. assignments only corresponding to initialisations). Your answer should include

(*i*) a map from sets of sets of variables to Boolean functions, discussing whether this is injective;

(*ii*) the difference between $\{\}$ and $\{\{\}\}$;

(*iii*) remarking on any way in which the correspondence is not perfect.

[8 marks]

(*d*) Now suppose we add function calls, each summarised by a *mustuseoneof* set, how might we deal with the dataflow of (say) `z = h(x-1,y+1)` assuming lazy semantics? Suppose f now has a recursive call, for example

```
f(x,y) = if x=0 then y else f(x-1,y+1)
```

explain what you expect the *mustuseoneof* set for the body of f to be and how this might be achieved by iterating from a given initial set. [4 marks]

11 Quantum Computing

(*a*) The SWAP-gate is a two-qubit quantum gate implementing the unitary operation $S: |\psi\rangle|\phi\rangle \mapsto |\phi\rangle|\psi\rangle$, for all arbitrary states $|\psi\rangle$ and $|\phi\rangle$ in $\mathbb{C}^2$.

(*i*) Write the matrix for S, in the computational basis. [3 marks]

(*ii*) Consider the quantum circuit and inputs depicted below. Write the final (output) state in terms of $(|\Psi\rangle + S|\Psi\rangle)$ and $(|\Psi\rangle - S|\Psi\rangle)$. [2 marks]

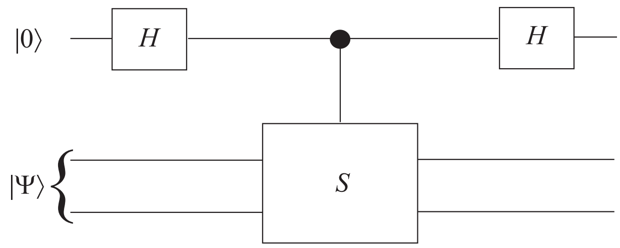

(*iii*) Let $p_0(|\Psi\rangle)$ denote the probability of getting outcome "0" if, after the final Hadamard gate in the above circuit with input $|0\rangle|\Psi\rangle$, you measure the top qubit with respect to the computational basis. Compute $p_0(|B\rangle)$ for each Bell state $|B\rangle$. [2 marks]

(*iv*) Suppose you are given two qubits, whose states are denoted $|\psi\rangle$ and $|\phi\rangle$, and you are promised that, with equal probability, $|\psi\rangle$ and $|\phi\rangle$ are either *identical* or *orthogonal*. (1) Describe how to use the above quantum circuit and measurement to help you decide which *case* – identical or orthogonal – you are in; i.e., specify the input state for the circuit and say which case you should guess for each measurement outcome. (2) Which case can be decided correctly with certainty? (3) Assuming you were in this case (but, of course, did not know it), what is the probability of getting the measurement outcome that allows you to be certain? [3 marks]

(*b*) Let D denote the quantum Fourier transform (QFT) on $\mathbb{C}^M$.

(*i*) Write the matrix for the two-qubit QFT, in the computational basis. [3 marks]

(*ii*) Derive a simplified expression for D^2 in terms of outer products of the computational basis elements. [5 marks]

(*iii*) Suppose that the final two steps in a quantum algorithm on an M-dimensional quantum system are (1) to apply D^{-1} and (2) to measure the entire quantum system with respect to the computational basis. Suppose you perform all but the second-to-last step of the algorithm correctly, applying D instead of D^{-1} in the second-to-last step. Is your final measurement outcome still meaningful? Explain. [2 marks]

12 Digital Signal Processing

(*a*) Make the following statements correct by changing one word or number. (Negating the sentence is not sufficient.)

(*i*) In the stopband, a filter design approximates a gain of -1. [1 mark]

(*ii*) For infinite sequences the z-transform always converges across the entire complex plane. [1 mark]

(*iii*) The Barlett window is the product of a rectangular window and a raised cosine function. [1 mark]

(*iv*) Multiplying two complex variables can be implemented with two real-valued multiplications and five real-valued additions. [1 mark]

(*v*) As a continuous signal is sampled, its Fourier spectrum becomes non-linear. [1 mark]

(*b*) Briefly explain

(*i*) the zigzag ordering of DCT coefficients in JPEG; [3 marks]

(*ii*) the difference between I-frames, P-frames and B-frames in MPEG; [3 marks]

(*iii*) the relationship between YCrCb and RGB colour coordinates. [4 marks]

(*c*) A 300 Hz sine wave is sampled at 1000 Hz. This discrete sequence is then multiplied, sample by sample, with the discrete sequence

$$\ldots, 0, +1, 0, -1, 0, +1, 0, -1, 0, +1, 0, -1, \ldots$$

Which frequencies appear in the Fourier transform of the result?

[5 marks]

13 Specification and Verification II

(*a*) What is the "false-implies-everything" problem? [2 marks]

(*b*) What is the difference between temporal and data abstraction? [2 marks]

(*c*) How do models of combinational and sequential devices differ? [2 marks]

(*d*) How is the rising edge of a signal modelled in higher order logic? [2 marks]

(*e*) Write down a formula that asserts that if a signal s has the value $\mathtt{a}$ then at all later times it also has the value $\mathtt{a}$. [2 marks]

(*f*) Give *two* properties of transistors that are not modelled by the simple switch model. [2 marks]

(*g*) What are advantages of using binary decision diagrams to represent Boolean formulae? [2 marks]

(*h*) What is the difference between linear and branching time? [2 marks]

(*i*) What are the relative expressive powers of LTL and CTL? [2 marks]

(*j*) How do the Verilog and VHDL simulation cycles differ? [2 marks]

14 Topics in Concurrency

(*a*) (*i*) Describe the modal μ-calculus and its semantics. [4 marks]

(*ii*) Describe how to express maximum fixed points $\nu Y.A$ in terms of minimum fixed points. [1 mark]

(*b*) (*i*) Describe an algorithm to determine whether a state in a finite-state transition system satisfies an assertion in the modal μ-calculus. [4 marks]

(*ii*) Explain briefly why the algorithm always terminates. [3 marks]

(*iii*) Use the algorithm to determine whether or not the state s in the labelled transition system below satisfies the assertion $[a]\nu Y.(\langle b\rangle T \wedge [b]Y)$, where T stands for "true".

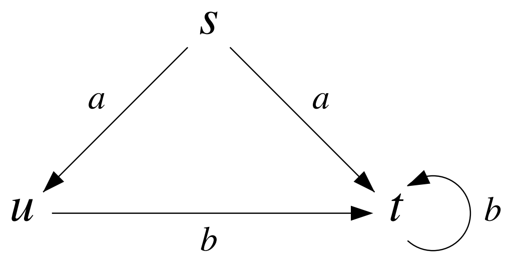

[6 marks]

(*iv*) Describe, without proof, the meaning of the assertion from the modal μ-calculus in part (*b*)(*iii*) above. [2 marks]

15 Types

(*a*) What is meant by *beta-reduction*, *beta-conversion* and *beta-normal forms* for the polymorphic lambda calculus (PLC)? Explain why typeable PLC expressions are beta-convertible to beta-normal forms that are unique up to alpha-conversion. Is the same true for untypeable PLC expressions? (Any general properties of PLC you use should be clearly stated, but need not be proved.) [10 marks]

(*b*) Let τ be the PLC type $\forall\beta((\alpha \to \beta) \to \beta)$, where α and β are distinct type variables. Give closed PLC beta-normal forms I and J with the following properties:

(*i*) I has type $\forall\alpha(\alpha \to \tau)$

(*ii*) J has type $\forall\alpha(\tau \to \alpha)$

(*iii*) $\Lambda\alpha(\lambda x : \alpha(J\, \alpha\, (I\, \alpha\, x)))$ has beta-normal form $\Lambda\alpha(\lambda x : \alpha(x))$

Justify your answers by giving proofs of typing and beta-conversion. [8 marks]

What is the beta-normal form of $\Lambda\alpha(\lambda y : \tau(I\, \alpha\, (J\, \alpha\, y)))$? [2 marks]

END OF PAPER