COMPUTER SCIENCE TRIPOS Part II

Thursday 9 June 2011 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 Computer Systems Modelling

(*a*) Suppose that the random variable X has an exponential distribution with parameter $\lambda > 0$.

(*i*) What are the probability density function, $f_X(x)$, and the probability distribution function, $F_X(x)$, for the random variable X? [2 marks]

(*ii*) Derive the mean and variance of the random variable X and determine its coefficient of variation. [3 marks]

(*iii*) Show that the random variable X obeys the *memoryless property*

$$P(X > t + s | X > t) = P(X > s)$$

for all $s, t > 0$. [2 marks]

(*iv*) Use the inverse transform method to derive a method to simulate random variables, X_i, indexed by $i = 1, 2, \ldots$ from an exponential distribution with parameter $\lambda > 0$ given a sequence of pseudo-random values U_i from the uniform distribution $U(0, 1)$. [3 marks]

(*b*) Suppose that $X_1, X_2, \ldots$ is a sequence of independent and identically distributed random variables with each random variable X_i having a marginal distribution that is an exponential distribution with parameter $\lambda > 0$. Let $S_n = \sum_{i=1}^{n} X_i$ where n is a positive integer and let the random variable $N(t)$ for $t > 0$ be the number of events in a Poisson Process with parameter $\lambda > 0$ that occur in the time interval $(0, t)$.

(*i*) State the probability distribution of $N(t)$. [2 marks]

(*ii*) State a relation between S_n and $N(t)$. [2 marks]

(*iii*) Derive the probability density of the random variable S_n. [6 marks]

2 Computer Vision

(*a*) Name **four** Gestalt laws of perceptual organisation. What is the main theme of Gestalt Psychology, and in what way is it relevant to automated object recognition in images and video sequences? [4 marks]

(*b*) Many computer vision algorithms such as SIFT (scale-invariant feature transform) seek to detect and analyse features at multiple scales of analysis.

(*i*) What is "scale-space" and what is a "scale-space fingerprint"? [2 marks]

(*ii*) Briefly describe how SIFT achieves scale and rotation invariance. [3 marks]

(*iii*) Show that the difference-of-Gaussian operation DoG given by

$$DoG(x, y, \sigma) = (G(x, y, k\sigma) - G(x, y, \sigma)) * I(x, y)$$

$$G(x, y, \sigma) = \frac{1}{2\pi\sigma^2} e^{-(x^2+y^2)/2\sigma^2}$$

on image data I as used by SIFT is an approximation of the scale-normalised Laplacian of Gaussian given by $\sigma^2 \nabla^2 G$. You may assume that

$$\frac{\partial G}{\partial \sigma} = \sigma \nabla^2 G$$

Comment on a good choice for the value of k for practical applications of SIFT. [6 marks]

(*c*) The curvature map of a closed image contour is a bandlimited signal and can be described by its zero-crossings. Explain how this property can be exploited for 2D shape description. What properties make such a representation suitable for classification and recognition of shapes? [5 marks]

3 Denotational Semantics

(*a*) A PCF type τ is said to be *finite* (resp. *infinite)* if the domain $[\![\tau]\!]$ is finite (resp. infinite). An element d of the domain $[\![\tau]\!]$ is said to be *definable* whenever there exists a closed PCF term $M : \tau$ such that $[\![M]\!] = d$.

Indicate whether the following statements are true or false. Provide an argument for each answer. You may use standard results provided that you state them clearly.

(*i*) For all finite PCF types τ, every element of the domain $[\![\tau]\!]$ is definable. [5 marks]

(*ii*) For all infinite PCF types τ, every element of the domain $[\![\tau]\!]$ is definable. [5 marks]

(*b*) Consider the following two statements for PCF terms M_1 and M_2 for which the typings $\Gamma \vdash M_1 : \tau$ and $\Gamma \vdash M_2 : \tau$ hold for some type environment Γ and type τ.

(1) For all PCF contexts $\mathcal{C}[-]$ for which $\mathcal{C}[M_1] : bool$ and $\mathcal{C}[M_2] : bool$,

$$\mathcal{C}[M_1] \Downarrow_{bool} \iff \mathcal{C}[M_2] \Downarrow_{bool}$$

where, for $M : \tau$, the notation $M \Downarrow_\tau$ stands for the existence of a value $V : \tau$ for which $M \Downarrow_\tau V$.

(2) For all PCF contexts $\mathcal{C}[-]$ for which $\mathcal{C}[M_1] : bool$ and $\mathcal{C}[M_2] : bool$,

$$\mathcal{C}[M_1] \Downarrow_{bool} \mathbf{true} \iff \mathcal{C}[M_2] \Downarrow_{bool} \mathbf{true}$$

(*i*) Show that (1) implies (2). [5 marks]

(*ii*) Define the notion of contextual equivalence in PCF and show that (2) implies that M_1 and M_2 are contextually equivalent. [5 marks]

4 Digital Signal Processing

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

(*i*) A square-summable sequence is also called power signal. [1 mark]

(*ii*) Adding together sine waves of the same frequency always results in another sine wave of the same phase. [1 mark]

(*iii*) The Fourier transform of a Dirac comb is a Dirac impulse. [1 mark]

(*iv*) 60 dBm = 1 W [1 mark]

(*b*) Before resampling a digital image at a quarter of its original resolution, you want to apply an anti-aliasing low-pass filter.

(*i*) If you apply a 1-dimensional filter with impulse response $\{h_n\}$ both horizontally and vertically to image pixels $I_{x,y}$, what are the resulting filtered pixel values $\tilde{I}_{x,y}$? [4 marks]

(*ii*) What would be the discrete impulse response $\{h_n\}$ of an ideal low-pass filter for this application, if its length were of no concern? [4 marks]

(*iii*) You decide to truncate the impulse response $\{h_n\}$ at its second zero-crossing on each side, resulting in a new impulse response $\{\bar{h}_n\}$. In the frequency domain, this results in $\bar{H} = H * T$ for what function T? [4 marks]

(*iv*) In order to make the frequency-domain response of your filter $\{\bar{h}_n\}$ smoother, you convolve it in the frequency domain with a rectangular pulse, the width of which is twice the distance between the zero crossings of T. What does the resulting time-domain impulse response $\{\breve{h}_n\}$ look like? [4 marks]

5 Information Retrieval

Consider the following text, which gives the input to three different information extraction systems dealing with texts about job succession events.

> Last Monday Sunny Simmers resigned his position as CEO of Sparkling Inc., the well-known cleaning supply manufacturer, following recent corruption allegations. This move has been expected for some time now. He is succeeded by his brother, Sandy Simmers.

The tasks of the three systems are as follows:

- System 1 determines all person and organisation names.
- System 2 determines the `employed-by` relationship, which holds between employees and employers.
- System 3 fills templates about job successions, in the style of MUC templates (as in the following figure).

START-JOB-EVENT
Person:
Position:
Company:
Start Date:

END-JOB-EVENT
Person:
Position:
Company:
End Date:

(*a*) For each system, give a description of how it could plausibly solve its task. What is the most difficult problem each system will encounter? Use examples from the text. [4 marks per system]

(*b*) Define each of the following concepts, giving an example in the context of the above text if applicable:

(*i*) sure-fire rule; [2 marks]

(*ii*) lexico-semantic template; [2 marks]

(*iii*) inference rule; [2 marks]

(*iv*) template learning. [2 marks]

6 Mobile and Sensor Systems

This question concerns Wireless Sensor Networks (WSNs).

(*a*) Explain the advantages of using event-based concurrency versus traditional concurrency through processes. Also explain how TinyOS implements event-based concurrency. [4 marks]

(*b*) In most WSNs, devices are battery operated and the use of the radio interface needs to be limited to preserve energy. Describe the salient characteristics of duty cycling through basic preamble sampling. [2 marks]

(*c*) Describe XMAC and how it optimises over basic preamble sampling approaches. [3 marks]

(*d*) Consider the WSN depicted in the picture below. Assume MintRoute is used as the routing protocol. Describe how MintRoute works and explain how the numbers on the edges in the network depicted below are generated by the protocol. What is the path that the packets would follow from source S to destination D if MintRoute were applied?

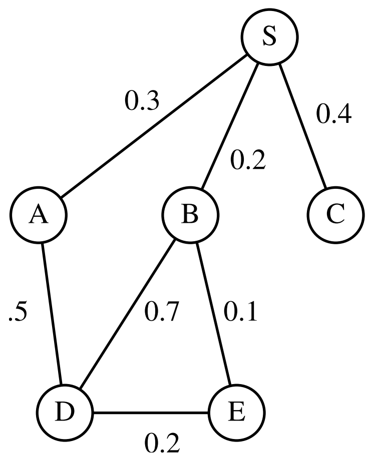

[5 marks]

(*e*) Explain why wireless reprogramming is important in WSNs. Describe the sensor network reprogramming described in the lectures and explain the role of the parameters which govern its behaviour. Also, what are the limitations of the approach when applied to mobile sensor networks? [6 marks]

7 Natural Language Processing

(*a*) Explain what is meant by *subsumption* and *unification* when applied to feature structures. [5 marks]

(*b*) The feature structures below represent a grammar rule and two lexical entries. The rule is intended to allow for adjectival and adverbial modification in phrases such as *big dog* and *suddenly jumps*. Give lexical entries for *big* and *suddenly* and show how the rule applies when these phrases are parsed. [7 marks]

$$\text{modifier–modified rule:} \begin{bmatrix} \textsc{mother} & \boxed{1} \\ \textsc{dtr1} & \begin{bmatrix} \textsc{mod} & \boxed{1} \end{bmatrix} \\ \textsc{dtr2} & \boxed{1} \begin{bmatrix} \textsc{obj} & \textbf{filled} \end{bmatrix} \end{bmatrix}$$

$$\text{dog:} \begin{bmatrix} \textsc{head} & \begin{bmatrix} \textsc{cat} & \textbf{noun} \\ \textsc{agr} & \textbf{sg} \end{bmatrix} \\ \textsc{obj} & \textbf{filled} \\ \textsc{subj} & \textbf{filled} \\ \textsc{mod} & \textbf{filled} \end{bmatrix} \qquad \text{jumps:} \begin{bmatrix} \textsc{head} & \begin{bmatrix} \textsc{cat} & \textbf{verb} \\ \textsc{agr} & \textbf{sg} \end{bmatrix} \\ \textsc{obj} & \textbf{filled} \\ \textsc{subj} & \begin{bmatrix} \textsc{head} & \begin{bmatrix} \textsc{cat} & \textbf{noun} \end{bmatrix} \end{bmatrix} \\ \textsc{mod} & \textbf{filled} \end{bmatrix}$$

(*c*) In the verb–object rule discussed in the lectures, the value of the HEAD feature of the mother is reentrant with the HEAD feature of DTR1, MOTHER SUBJ is reentrant with DTR1 SUBJ, and DTR1 OBJ is reentrant with DTR2. Explain why the pattern of reentrancies in the modifier–modified rule differs from that of the verb–object rule. [6 marks]

(*d*) Why is the value of DTR2 OBJ specified as **filled** in the modifier–modified rule? [2 marks]

8 Optimising Compilers

(*a*) Explain what is meant by "instruction scheduling" in a compiler. Indicate what properties of architectures make it beneficial. Are some architectures too simple, or too complex, for it to be useful? [3 marks]

(*b*) Sketch an algorithm that schedules instructions within a basic block for an architecture of your choice on which scheduling is useful. Indicate its time complexity $O(f(n))$ stating to what 'n' refers. What additional complications arise in choosing which instruction to emit first from a basic block? [5 marks]

(*c*) Indicate how your algorithm could be modified to deal with a *static multiple-issue* VLIW-style architecture. The input basic block contains simple instructions, but the output basic block contains wide instructions consisting of two simple instructions (either of which may be NOP). These two instructions execute in parallel.

Now explain how your algorithm can schedule the following sequence of instructions to execute in parallel, or detail any adjustments necessary to permit this.

```
st  r1,4(r8)
ld  r2,0(r8)
```

Explain how *alias analysis* or *points-to analysis* might provide information to the instruction scheduler to enable a wider range of load/store pairs to be re-scheduled. [6 marks]

(*d*) Briefly summarise a source-level approach to alias or points-to analysis, for example Andersen's algorithm. Your summary should emphasise key choices, for example any data-structures used; detailed algorithms are not required. [6 marks]

9 Principles of Communications

(*a*) Control Theory allows us to find various useful properties of a system such as stability. Draw a picture of a generic control system, explaining the functions of *feedback* and the design goals for the *controller*. [5 marks]

(*b*) The Transmission Control Protocol of the Internet uses a feedback controller and responds to absence or presence of congestion signals by *Additive Increase Multiplicative Decrease* of a sending window.

(*i*) Describe the operation of this scheme. [10 marks]

(*ii*) Explain qualitatively why the scheme is necessary, but perhaps not sufficient, to create a stable set of traffic flows in the Internet. [5 marks]

10 System-on-Chip Design

(*a*) What is an *instruction set simulator* (ISS) and what is typically the best way for it to achieve high performance? [2 marks]

(*b*) What performance bottleneck can typically arise when CPU-intensive software is run on a model of a complete system on chip (SoC)? How can this be avoided and at what costs to modelling accuracy? [4 marks]

(*c*) What problems can arise when operating system device drivers are run on an ISS that has been optimised according to your answers to parts (*a*) and (*b*) above? [4 marks]

(*d*) Two processors on a SoC have separate address spaces. What does this mean? Describe a simple hardware mechanism for sending non-trivial amounts of data between processors in separate address spaces. Interrupts should be used. [4 marks]

(*e*) Briefly describe both a *high-level* and a *mid-level* model of your answer to part (*d*) where the mid-level model requires minimal or no modification to the device driver firmware but is not a net-level model and where the high-level model dispenses with much of the device driver code and many or all of the interrupts. [3 marks each]

11 Temporal Logic and Model Checking

(*a*) Explain the difference between static and dynamic methods for property checking. Is LTL or CTL a better property language for dynamic checking? Justify your answer. [4 marks]

(*b*) Outline an algorithm for model checking LTL properties of the form `G` **p**, where **p** is an atomic property. [4 marks]

(*c*) What is the difference between explicit state and symbolic model checking? [4 marks]

(*d*) Who is the person credited with introducing temporal logic as a property specification language? What were the main contributions of Clarke & Emerson, and of McMillan, to model checking? Why were they significant? [4 marks]

(*e*) Briefly explain the terms *Boolean layer*, *Temporal layer*, *Verification layer* and *Modelling layer* in connection with PSL. [4 marks]

12 Topical Issues

(*a*) Describe the Binary Exponential Backoff (BEB) procedure found in 802.11 (WiFi) and explain why it is necessary. Give **two** shortcomings of the BEB approach. [4 marks]

(*b*) Describe **two** modifications to the BEB algorithm that could be made by a malicious user attempting to increase their chances of gaining access to the wireless channel. At least one of your modifications, if used by all wireless stations, should enable the provision of Quality of Service and you should explain how it would do so. [6 marks]

(*c*) Discuss the behaviour of the BEB procedure in heavily congested networks. Propose an alternative procedure, which may be based on BEB, that has improved behaviour in such networks. Take care to explain how it outperforms the default technique. Discuss how your solution behaves in lightly congested networks. [10 marks]

13 Topics in Concurrency

This question concerns CTL$^-$, a variant of CTL in which assertions are of the form:

$$A := T \mid A_0 \vee A_1 \mid A_0 \wedge A_1 \mid \neg A \mid \langle \lambda \rangle A \mid \langle - \rangle A \mid \mathbf{AG}\, A \mid \mathbf{EG}\, A$$

where λ ranges over action names. Recall that a *path* π from state s is a maximal sequence of states $\pi = (\pi_0, \pi_1, \pi_2, \ldots)$ such that $s = \pi_0$ and $\pi_i \stackrel{\cdot}{\longrightarrow} \pi_{i+1}$ for all i. The logical connectives have the standard interpretation and T represents "true". The interpretation of the modalities is:

$s \models \langle \lambda \rangle A$	iff	there exists s' such that $s \stackrel{\lambda}{\longrightarrow} s'$ and $s' \models A$
$s \models \langle - \rangle A$	iff	there exists s' such that $s \stackrel{\cdot}{\longrightarrow} s'$ and $s' \models A$
$s \models \mathbf{EG}\, A$	iff	for some path π from s, we have $\pi_i \models A$ for all i
$s \models \mathbf{AG}\, A$	iff	for all paths π from s, we have $\pi_i \models A$ for all i

(*a*) What is the interpretation of the CTL modality $\mathbf{E}[A_0\ \mathbf{U}\ A_1]$? How can it be used to express the CTL$^-$ modality $\mathbf{AG}\, A$? [4 marks]

(*b*) Consider the following three formulae:

$$A_1 : \quad \mathbf{AG}\, \langle a \rangle T \qquad A_2 : \quad \neg\, \mathbf{AG} \neg \langle b \rangle T \qquad A_3 : \quad \neg\, \mathbf{EG}\, \langle b \rangle T$$

(*i*) For each of the following two transition systems, state which of A_1, A_2 and A_3 are satisfied in the initial state.

P_1 : 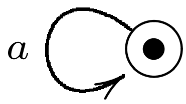

P_2 : 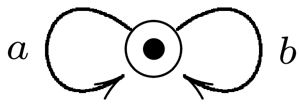

[4 marks]

(*ii*) Draw a transition system with an initial state that satisfies $A_1 \wedge A_2 \wedge A_3$. [3 marks]

(*c*) Give a modal-μ formula that corresponds to the CTL$^-$ formula $\mathbf{AG}\, \langle a \rangle T$. [2 marks]

(*d*) Let the function φ on sets of states of a transition system be defined as:

$$\varphi(X) \stackrel{\text{def}}{=} \langle a \rangle T \vee \langle - \rangle X$$

Show by induction on $n \geq 1$ that

$$s \models \varphi^n(\emptyset) \quad \text{iff} \quad \text{there exists } m \leq n \text{ and states } s_1, \ldots, s_m, s' \text{ such that } s = s_1 \stackrel{\cdot}{\longrightarrow} \ldots \stackrel{\cdot}{\longrightarrow} s_m \stackrel{a}{\longrightarrow} s'$$

Deduce that $s \models \mu X. \langle A \rangle T \vee \langle - \rangle X$ in a finite-state transition system if, and only if, $s \models \neg\, \mathbf{AG} \neg \langle a \rangle T$. [7 marks]

14 Types

Let x range over a set of identifiers and α range over a set of type variables. Now suppose we have a set of types, τ, and a set of type schemes, σ, given by

$$\begin{aligned}\tau &::= \alpha \quad | \quad \tau \to \tau \quad | \quad \tau\ list \\ \sigma &::= \forall \alpha_1, \ldots, \alpha_n\,(\tau)\end{aligned}$$

and a language of terms, M, given by

$$\begin{aligned}M \ ::= \ & x \ \ | \ \ \lambda x\,(M) \ \ | \ \ M\ M \ \ | \ \ \texttt{let}\ x = M\ \texttt{in}\ M \ \ | \ \ \texttt{nil} \ \ | \ \ M :: M \\ | \ \ & \texttt{case}\ M\ \texttt{of}\ \texttt{nil} \Rightarrow M \mid x :: x \Rightarrow M\end{aligned}$$

(*a*) Define the relation of specialisation, $\tau \prec \sigma$, between types τ and type schemes σ. [3 marks]

(*b*) Give the ML-like type inference rules for judgements of the form $\Gamma \vdash M : \tau$, and explain why λ-bound variables cannot be used polymorphically within a function abstraction, while `let`-bound variables can within a local declaration.

Hint: Consider the terms $\lambda f(f\ f)$ and $\texttt{let}\ f = \lambda x(x)\ \texttt{in}\ f\ f$.

[8 marks]

(*c*) Briefly explain what is meant by *capture-avoiding substitution* for type schemes. [2 marks]

(*d*) Prove that for all τ, all σ and all substitutions for type schemes S, if $\tau \prec \sigma$ holds, then also $S(\tau) \prec S(\sigma)$.

Hint: Use the following property of simultaneous substitution:

$$(\tau[\tau_1/\alpha_1, \ldots, \tau_n/\alpha_n])[\vec{\tau'}/\vec{\alpha'}] = \tau[\vec{\tau'}/\vec{\alpha'}][\tau_1[\vec{\tau'}/\vec{\alpha'}]/\alpha_1, \ldots, \tau_n[\vec{\tau'}/\vec{\alpha'}]/\alpha_n]$$

which holds, provided that for each i, α_i is distinct from the type variables $\vec{\alpha'}$, and α_i does not occur in type schemes $\vec{\tau'}$.

[7 marks]

END OF PAPER