COMPUTER SCIENCE TRIPOS Part II

Wednesday 6 June 2007 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.*

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	*None*
Blue cover sheets	
Tags	

1 Advanced Systems Topics

(*a*) Mutexes are usually built from *atomic* processor instructions. What does it mean for a processor instruction to be atomic, and how can this property be implemented in a cache-coherent multiprocessor system? [4 marks]

(*b*) Provide pseudocode for a simple multi-reader spinlock, including all four operations supported by this type of lock. Describe any atomic operations that your pseudocode uses. [8 marks]

(*c*) Why might the simple multi-reader spinlock scale poorly? Sketch a more scalable design assuming that read-only critical sections are vastly more frequent than critical sections that modify shared state. [4 marks]

(*d*) Is it possible to implement mutual exclusion in a multiprocessor system that provides only atomic load and store instructions? If so, why do modern processors provide read–modify–write instructions? [4 marks]

2 Natural Language Processing

(*a*) In German, the third person singular present inflection of weak verbs is generally formed by adding the 't' to the stem. Exceptions to this rule include verbs with stems that end in 't' or 'd' which are formed by adding 'et' instead of 't'. The following table gives some examples (ˆ is used as the affix marker in the underlying form).

stem	surface	underlying
kauf	kauft	kaufˆt
arbeit	arbeitet	arbeitˆt

Draw a finite state transducer (FST) that relates surface and underlying forms according to this pattern. (Only the inflected forms should be accepted by the transducer since the stems by themselves do not correspond to words.) Explain the notation that you use and outline how the FST could be used in morphological analysis and generation. [14 marks]

(*b*) The past participle of these verbs is the same as the third person singular present, but with 'ge' before the stem. Assume that the underlying form of the past participle is treated as having the artificial suffix 'ˆP', as indicated below.

stem	surface	underlying
kauf	gekauft	kaufˆP
arbeit	gearbeitet	arbeitˆP

Modify the FST that you gave in answer to part (*a*) above to allow for the past participle as well as the third person singular present. What does this example illustrate about limitations of the FST approach? [6 marks]

3 Comparative Architectures

(*a*) Why is multi-processing more important than ever? [3 marks]

(*b*) Name the *two* basic forms of parallelism that are exploited by all new high-performance processors. [2 marks]

(*c*) Which form(s) of parallelism is/are exploited by the following four styles of multi-processing: multi-core processors, coarse and fine interleaved multi-threading, and simultaneous multi-threading? Explain these terms and say how good each is at hiding memory latency. [6 marks]

(*d*) How does compiler and operating system support need to differ for each form and style? [4 marks]

(*e*) Thinking of the complete flow from high-level language design through to hardware execution, what tends to limit the available parallelism in a single thread? Suggest various ways of increasing it. [5 marks]

4 Digital Communication II

Routing in the Internet is distributed, adaptive, and above all *opportunistic.*

(*a*) Describe, using simple topologies, the operation of *distance vector* and *link state* algorithms. [5 marks each]

(*b*) End-to-end reliability, flow control, and congestion avoidance are provided through TCP. Describe the impact on throughput for an end-to-end TCP connection of a link failure, followed by the provision of a new end-to-end path by the convergence of the distributed route computation. [5 marks]

(*c*) How would the same failure and repair affect a Voice-over-IP session in terms of packet deliveries? [2 marks]

(*d*) How could one mitigate the impact of dynamic single path routing on so-called realtime flows, and what would the impact (overhead) of this mitigation be? [3 marks]

5 Distributed Systems

(*a*) Define the characteristics of middleware based on the following interaction paradigms:

(*i*) request-reply;

(*ii*) one-to-one message passing.

Discuss any limitations imposed by the paradigms and mention the properties of the application domains for which each is particularly suited. [9 marks]

(*b*) Web Services standards have been specified relatively recently. What does the Web Services paradigm offer? [3 marks]

(*c*) (*i*) Define the publish/subscribe communication paradigm. [2 marks]

(*ii*) What does this paradigm provide that is lacking in the three described above in parts (*a*) and (*b*)? [3 marks]

(*iii*) What are the shortcomings of publish/subscribe? [3 marks]

6 Computer Vision

(*a*) Briefly define each of the following concepts as it relates to vision:

(*i*) "signal-to-symbol converter"; [2 marks]

(*ii*) "inverse graphics"; [2 marks]

(*iii*) blurred Laplacian operator; [2 marks]

(*iv*) volumetric coordinates; [2 marks]

(*v*) correspondence problem. [2 marks]

(*b*) Give *three* examples of methodologies or tools used in Computer Vision in which Fourier analysis plays a role, either to solve a problem, or to make a computation more efficient, or to elucidate how and why a procedure works. For each of your examples, clarify the benefit offered by the Fourier perspective or implementation. [6 marks]

(*c*) When designing a pattern classifier, what roles are played by within-class variability and between-class variability? Which one is helpful and which one is undesirable? How should the definition or selection of features reflect these two kinds of variability? Illustrate these points in the context of face recognition. [4 marks]

7 Information Theory and Coding

(*a*) Suppose that X is a random variable whose entropy $H(X)$ is 8 bits. Suppose that $Y(X)$ is a deterministic function that takes on a different value for each value of X.

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

(*ii*) What is $H(Y|X)$, the conditional entropy of Y given X? [1 mark]

(*iii*) What is $H(X|Y)$, the conditional entropy of X given Y? [1 mark]

(*iv*) What is $H(X, Y)$, the joint entropy of X and Y? [1 mark]

(*v*) Suppose now that the deterministic function $Y(X)$ is not invertible; in other words, different values of X may correspond to the same value of $Y(X)$. In that case, what could you say about $H(Y)$? [2 marks]

(*vi*) In that case, what could you say about $H(X|Y)$? [1 mark]

(*b*) A continuous real-valued signal has a bandwidth limited to ±W Hertz.

(*i*) In a duration of time T, at most how many regularly-spaced samples are needed in order for the signal to be specified completely at all points within T? State the theorem that is the basis of your answer. [2 marks]

(*ii*) What about the signal values in between the points that are sampled – how can anything be known about those unobserved values? [1 mark]

(*c*) Write down the general functional form for a 1-D Gabor wavelet, and explain how particular choices for the values of its parameters would turn it into either the Fourier basis or the delta function sampling basis, as two special cases. [3 marks]

(*d*) Show that the set of all Gabor wavelets is closed under convolution. That is, show that the convolution of any two Gabor wavelets is also a Gabor wavelet. Comment on how this property relates to the fact that these wavelets are also closed under multiplication, and that they are also self-Fourier. [3 marks]

(*e*) We wish to compute the Fourier Transform of a data sequence of 1,024 samples:

(*i*) Approximately how many multiplications would be needed if the Fourier integral expressions were to be computed literally (as written mathematically) and without a clever algorithm? [2 marks]

(*ii*) Approximately how many multiplications would be needed if an FFT algorithm were used? [2 marks]

8 Optimising Compilers

(*a*) Summarise the core ideas of *Common Sub-Expression Elimination* (CSE) optimisation based on dataflow analysis. Your explanation should include dataflow equations (including auxiliary functions used) and sketches of how (*i*) to compute dataflow solutions and (*ii*) these solutions drive the transformation itself. [8 marks]

(*b*) Consider *P*, the following SSA-form intermediate code, resulting from CSE (treat `READ A` as an atomic instruction writing a run-time-determined value to `A`):

```
1: READ A
2: READ B
   ...
6: Z := A+B
7: C := Z
8: D := Z
9: PRINT C+D
```

(*i*) Instruction 6 was inserted by CSE. Leaving the other instructions alone, where else might it have been reasonably inserted? Indicate any conditions your answer places on the unspecified instructions 3 to 5, recalling that *P* is in SSA form. [3 marks]

(*ii*) Registers are often in short supply. One heuristic for positioning instruction 6 is to minimise the sum of the lengths (in lines) of all live ranges of program variables (here `A`, `B`, `C`, `D`, `Z`). Calculate the values of this heuristic for the given program, and also for program *Q* obtained by following *P* with `10: PRINT A-B`. (For this purpose, it is sufficient to assume instructions 3 to 5 are simple "no-operations" reading and writing no variables.) [4 marks]

(*iii*) Now calculate the minimal values of this heuristic based on swapping instruction 6 with one of instructions 3 to 5 for the two programs. Hence, or otherwise, give a one-sentence refinement to your answer in part (*a*). [5 marks]

9 Artificial Intelligence II

An agent can exist in a state $s \in S$ and can move between states by performing actions, the outcome of which might be uncertain.

(a) Explain what is meant by a *Utility Function* within this context. [2 marks]

(b) Give a definition of *Maximum Expected Utility* and describe the way in which it can be used to decide which action to perform next. [3 marks]

(c) What difficulties might you expect to have to overcome in practice in order to implement such a scheme? [3 marks]

(d) Explain why it makes sense to use a utility function in the design of an agent, even though it can be argued that real agents (such as humans) appear not to do this, but rather to act on the basis of *preferences.* [4 marks]

(e) As well as actions allowing an agent to move between states, an agent might be capable of performing actions that allow it to discover more about its environment. Give a full derivation of the *Value of Perfect Information*, and explain how this idea can be used as the basis for an agent that can gather further information in a way that takes account of the potential cost of performing such actions. [8 marks]

10 Digital Signal Processing

(*a*) The DAUB4 wavelet transform involves a pair of 4-point FIR filters.

(*i*) Explain the properties that these filters are designed to have and provide a system of equations that defines the two impulse responses accordingly. [8 marks]

(*ii*) Explain briefly how this filter pair is used in the wavelet transform. [4 marks]

(*b*) Consider a digital radio designed to receive all signals in the frequency range 90–105 MHz. Its antenna amplifier includes a bandpass filter that eliminates any signals outside this frequency range. The filtered antenna signal is directly fed into an analogue-to-digital converter, such that all subsequent demodulation steps can be performed in software.

(*i*) What is the lowest sampling frequency that can be used without risking loss of information due to aliasing? Explain briefly why. [5 marks]

(*ii*) If the resulting discrete sequence were turned into a continuous baseband signal through sinc interpolation, what relationship would there be between the spectra of the input and output signal? In particular, what would a 94 MHz sine-wave antenna signal be converted into? [3 marks]

11 Security

(*a*) In cryptography, what do we mean when we describe a hash function as a "pseudorandom function"? [4 marks]

(*b*) A hash function h is said to be *collision resistant* if it is hard to find $m_1 \neq m_2$ such that $h(m_1) = h(m_2)$ and *preimage resistant* if given a random y it is hard to find X such that $h(X) = y$.

Describe how hash functions can be used in the following applications, in each case indicating whether the function needs to be collision-resistant or preimage-resistant.

(*i*) digital certificates generated by a trustworthy CA for use in SSL/TLS; [4 marks]

(*ii*) the exchange of electronically-signed contracts by companies; [4 marks]

(*iii*) computing fingerprints on "known good" files by anti-virus software; [4 marks]

(*iv*) hash chains of digital coins for use in an electronic cash system. [4 marks]

12 Quantum Computing

(*a*) Give a schematic circuit diagram for Grover's algorithm. [4 marks]

(*b*) Suppose we apply Grover's algorithm to a four-qubit register in which exactly one of the states is marked.

(*i*) Writing O for the oracle matrix, write out the Grover iterate in matrix form. [4 marks]

(*ii*) What are the probabilities of measuring the marked state after applying the Grover iterate *once*, *twice* and *three* times? [3 marks each]

(*iii*) If you continue to increase the number of times the Grover iterate is applied, will the probability continue to increase? Justify your answer. [3 marks]

13 Bioinformatics

(*a*) Describe with *one* example the difference between Hamming and Edit distances. [2 marks]

(*b*) Discuss the Smith–Waterman algorithm. What is the complexity and the relationship with the problem of finding the longest common subsequences? [5 marks]

(*c*) Describe the Banded algorithm for local alignment and its complexity. [5 marks]

(*d*) Describe the four Russian speedup algorithm. [8 marks]

14 Human–Computer Interaction

In designing a university-wide student information system for Cambridge, it is necessary to meet the usability requirements of various groups, including (amongst others):

(*i*) staff working in college tutorial offices, and

(*ii*) people applying online for admission as graduate students.

(*a*) In what respects would the usability requirements of *these two* groups be likely to differ? [2 marks each]

(*b*) Which usability analysis methods would provide most insight into typical usage scenarios for each of these two groups? For *each* method, explain the reason why it would be appropriate, and also the analysis procedures that should be followed. [5 marks each]

(*c*) Drawing on either heuristic evaluation, cognitive dimensions of notations, or relevant components of both, suggest *three* specific interaction features that might address the needs of these different users. [6 marks]

15 Denotational Semantics

(*a*) Suppose that $(D, \sqsubseteq_D)$ and $(E, \sqsubseteq_E)$ are cpos.

(*i*) What properties does a function $f : D \to E$ need to satisfy in order to be continuous? [2 marks]

(*ii*) Assume also that $(C, \sqsubseteq_C)$ is a cpo and that $g : C \times D \to E$ is a continuous function. Let $g^* : C \to (D \to E)$ be defined by $g^*(c) = \lambda d \in D.\, g(c, d)$. Prove that g^* is continuous. You may refer to general facts about least upper bounds in product and function cpos provided that you state them clearly. [6 marks]

(*b*) Let $\mathbf{2} = (\{\bot, \top\}, \bot \sqsubseteq \top)$ be the unique domain with two elements.

(*i*) Draw a diagram which represents the elements of the function domain $\mathbf{2} \to \mathbf{2}$ and shows their ordering; [1 mark]

(*ii*) Any set X can be considered as a flat domain $X_\bot$ by adding a bottom element. Show that the strict continuous functions $X_\bot \to \mathbf{2}$ are in 1-1 correspondence with the subsets of X. [2 marks]

(*c*) Define what is meant by an admissible subset of a domain D. [2 marks]

(*d*) State the principle of Scott induction and prove its validity. [4 marks]

(*e*) Suppose that D is a domain and $f : D \times D \to D$ is a continuous function satisfying the property $\forall d, e \in D.\ f(d, e) = f(e, d)$. Let $g : D \times D \to D \times D$ be defined by $g(d_1, d_2) = (f(d_1, f(d_1, d_2)), f(f(d_1, d_2), d_2))$. Let $(u_1, u_2) = \mathit{fix}(g)$. Show that $u_1 = u_2$ using Scott induction. [3 marks]

16 Specification and Verification I

(*a*) What is the difference between the methods of refinement and *post hoc* verification? List *one* benefit of using refinement besides creating correct-by-construction programs. [4 marks]

(*b*) What features are needed in a wide-spectrum language to support refinement? Illustrate your answer by describing constructs found in such a language that are not present in programming languages. [4 marks]

(*c*) What are the "Laws of Programming" in relation to refinement? Illustrate your answer with a concrete example of a law and contrast your law with a related axiom or rule of Floyd–Hoare Logic. [4 marks]

(*d*) What does it mean for refinement to be *monotonic*? Why is monotonicity important? [4 marks]

(*e*) Exhibit a refinement of `[Y, X=0` $\wedge$ `Y=2`$\times$`X]` to `X:=0;Y:=0`. [4 marks]

END OF PAPER