UNIVERSITY OF CAMBRIDGE

COMPUTER SCIENCE TRIPOS Part IB

Tuesday 31 May 2016 1.30 to 4.30

COMPUTER SCIENCE Paper 4

*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 Artificial Intelligence I

(*a*) A planning problem has start state $\{\neg A, \neg B, \neg C, \neg D, \neg E\}$ and goal $\{D, E\}$. The following actions are available:

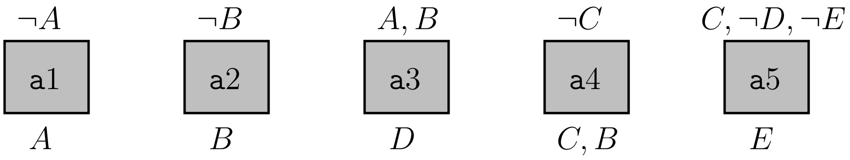

Give a detailed explanation of how the *Partial Order Planning Algorithm* can be used to solve this problem. [8 marks]

(*b*) In order to include heuristics in your planning system you decide also to construct the planning graph for the problem in Part (*a*). Provide a description of the first two state levels and the first action level in the planning graph, and explain why each item is included in each level. [4 marks]

(*c*) You now wish to add a further action level and a further state level to the graph. Explain what new items will be added to each new level and why. You need not mention items that are identical to those in earlier levels, only those that are added now for the first time. [3 marks]

(*d*) Provide an example of how each of the five possible kinds of mutex link will appear in the graph constructed in Part (*b*) and Part (*c*). [5 marks]

2 Artificial Intelligence I

(*a*) Provide a detailed description of the *Iterative Deepening A* (IDA*)* algorithm. Your answer should include a clear statement of the algorithm in pseudo-code, and a general description of how it works. [8 marks]

(*b*) Explain how the IDA* algorithm searches the following search tree. Numbers between nodes denote the cost of the path between those nodes, and numbers on the nodes denote the value of the heuristic for that node. [8 marks]

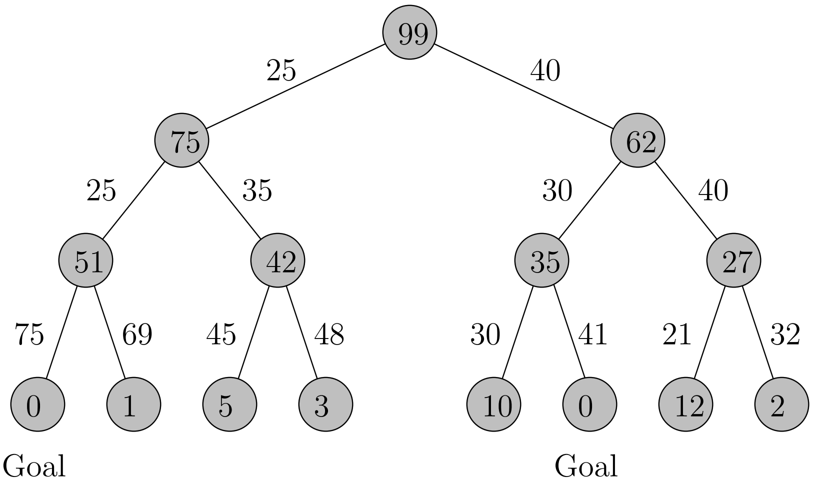

(*c*) Give *two* reasons why the IDA* algorithm might prove unsuitable as a solution to a search problem, in each case giving a brief explanation of why this is the case, and suggesting a potential solution. [4 marks]

3 Computer Graphics and Image Processing

A program is required to draw an arc from $(0, 1)$ to $(1, 0)$ of the circle centred at the origin with unit radius.

(*a*) One approach would be to draw a segment of the cubic Overhauser curve defined by $(-1, 0)$, $(0, 1)$, $(1, 0)$ and $(0, -1)$.

(*i*) Explain how a segment of an Overhauser curve in general can be represented as an Hermite cubic and so as a Bézier cubic. [4 marks]

(*ii*) Derive the formula for the resulting Bézier curve, $\mathbf{P}(t)$. [3 marks]

(*iii*) Calculate the coordinates of $\mathbf{P}(\frac{1}{2})$. How large is the error? [*Hint*: $\sqrt{2} \approx 1.414$.] [3 marks]

(*b*) Calculate revised control points for the Bézier curve so that it models the circular arc more accurately. [4 marks]

(*c*) Describe in outline an alternative way of efficiently drawing the arc by calculating the pixels that lie on it directly. [6 marks]

4 Computer Graphics and Image Processing

(*a*) Describe in detail the Cohen-Sutherland algorithm to clip a straight line segment against a rectangle. [8 marks]

(*b*) Extend the algorithm from part (*a*) to clip a line against a three-dimensional viewing frustrum. [6 marks]

(*c*) Describe how to clip a Bézier curve against a screen rectangle. [6 marks]

5 Databases

(*a*) Define the concept of a *functional dependency.* [3 marks]

(*b*) Suppose that relation R has m attributes. Give an upper bound on the number of functional dependencies that R could satisfy (including trivial dependencies). [3 marks]

(*c*) Let $R(A, B, C, D, E)$ be a relational scheme with the following dependencies.

$$\begin{array}{rcl} A & \rightarrow & C \\ B, C & \rightarrow & D \\ A & \rightarrow & E \\ B, D & \rightarrow & C \\ C & \rightarrow & E \\ E & \rightarrow & D \\ E & \rightarrow & B \end{array}$$

Which, if any, of these dependencies are redundant? [4 marks]

(*d*) Suppose $R(A,\ B,\ C)$ is a relational schema with functional dependency $A \rightarrow B$. What can you deduce about the results of $\pi_{A,B}(R) \bowtie_A \pi_{A,C}(R)$? Justify your answer. [3 marks]

(*e*) Suppose $R(A,\ B,\ C)$ is a relational schema. In addition, you know that the following is always true in any correct database instance.

$$R = \pi_{A,B}(R) \bowtie_A \pi_{A,C}(R).$$

What can you deduce about the dependencies between attributes A, B, and C? Prove any of your claims. [7 marks]

6 Databases

This question deals with the variety of approaches to database design.

(*a*) What is meant by the term *on-line transaction processing* (OLTP)? [3 marks]

(*b*) What is meant by the term *on-line analytic processing* (OLAP)? [3 marks]

(*c*) Compare and contrast the approach to schema design for OLTP and OLAP databases. [3 marks]

(*d*) Compare OLAP with the so-called NoSQL approach to database design. [3 marks]

(*e*) Give an example of a set of requirements whose solution would need to combine OLAP, OLTP and NoSQL. Describe an architecture integrating these elements in the system design. [8 marks]

7 Economics, Law and Ethics

(*a*) Describe three ways in which information goods and services markets differ from the market for coal or for potatoes. [6 marks]

(*b*) What are the usual effects of these differences on the structure of such markets? [4 marks]

(*c*) You are the CEO of a car company considering adoption of Android as the platform for the entertainment, navigation and related systems in your next generation of vehicles. Should the app store be run by your company or by Google, and how should the safety case for apps be established? [5 marks]

(*d*) You are an academic advising the Secretary of State for Transport on how vehicle app stores should be regulated. What would your advice be, both in terms of the UK public interest and the likely effects of EU regulation? [5 marks]

8 Security I

(*a*) Block ciphers usually process 64 or 128-bit blocks at a time. To illustrate how their modes of operation work, we can use instead a pseudo-random permutation that operates on the 26 letters of the English alphabet:

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25
m	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z
$E_K(m)$	D	G	W	X	T	E	R	L	Y	Z	O	J	N	S	I	Q	P	C	U	H	B	V	F	A	M	K

As the XOR operation is not defined on the set $\{\texttt{A}, \ldots, \texttt{Z}\}$, we replace it here during encryption with modulo-26 addition (e.g., $\texttt{C} \oplus \texttt{D} = \texttt{F}$ and $\texttt{Y} \oplus \texttt{C} = \texttt{A}$).

(*i*) Decrypt the following ciphertexts, which were encrypted using

(A) Electronic codebook mode: `UOMHDJT` [2 marks]

(B) Cipher feedback mode: `RVPHTUH` [4 marks]

(C) Output feedback mode: `LNMSUUY` [4 marks]

(*ii*) Determine the CBC-MAC for the message `TRIPOS`. [4 marks]

(*b*) Consider another small pseudo-random permutation, this time defined over the set of decimal digits $\{0, 1, 2, \ldots, 9\}$, using modulo-10 addition instead of XOR (e.g., $7 \oplus 3 = 0$).

(*i*) You have intercepted the message `100` with appended CBC-MAC block `4`. The message represents an amount of money to be paid to you and can be of variable length. Use this information to generate a message that represents a much larger number, and provide a valid CBC-MAC digit, without knowing the pseudo-random permutation or key that the recipient will use to verify it. [4 marks]

(*ii*) What mistake did the designer of the communication system attacked in part (*b*)(*i*) make (leaving aside the tiny block size), and how can this be fixed? [2 marks]

9 Security I

(*a*) Briefly explain *return-oriented programming*: what kind of software vulnerability and countermeasure does this class of attacks target, how does it work, and under what conditions is it applicable? [6 marks]

(*b*) Identify and fix a potential vulnerability in the following C function: [2 marks]

```
#include <stdlib.h>
void *bitmalloc(size_t bits) {
  return malloc((bits + 7)/8);
}
```

(*c*) On a Linux file server, you find this file:

```
$ ls -l
-rw----r-- 1 frank students  13593 May 31 14:55 question.tex
```

User `frank` is a member of group `students`.

(*i*) Based on the POSIX access-control settings shown, illustrate how the server's operating system will authorize access (if-statement pseudo code). [3 marks]

(*ii*) What does an equivalent Windows NTFS access-control list look like? [3 marks]

(*iii*) Does the Windows GUI for manipulating NTFS access-control lists allow users to enter this configuration? [2 marks]

(*d*) Give an example of how POSIX file-system access control can be used to provide the equivalent of password protection for parts of the file space. In particular, show how user `alice` can set up a directory `papers` such that only those members of group `committee` (which includes `alice`) who know the secret string "`SEL-4sB3`" can read a file `restricted.pdf`. Show the setup either as a sequence of shell commands that `alice` can use to create it, or in the form of the metadata of the files and directories involved (as `ls -l` would output it). [4 marks]

END OF PAPER