UNIVERSITY OF CAMBRIDGE

COMPUTER SCIENCE TRIPOS Part IA

Tuesday 31 May 2016 1.30 to 4.30

COMPUTER SCIENCE Paper 2

*Answer **one** question from each of Sections A, B and C, and **two** questions from Section D.*

*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

SECTION A

1 Digital Electronics

(*a*) Consider a 4-input Boolean function that outputs a binary 1 whenever an odd number of its inputs are binary 1.

(*i*) Using Boolean logic or otherwise, show how the above function can be implemented using only 2-input XOR gates. [4 marks]

(*ii*) Show how the above function may alternatively be implemented using one 4-input decoder, and a minimum number of 4-input NOR and 4-input NAND gates. [3 marks]

(*b*) Consider the following Boolean expression

$$F = \overline{B}.\overline{C} + \overline{A}.B.C + A.C.\overline{D}$$

(*i*) Show that F can be represented by the following Product of Sums (POS) form

$$F = (\overline{B} + C).(\overline{A} + \overline{C} + \overline{D}).(A + B + \overline{C})$$

[3 marks]

(*ii*) Show how F can be implemented in a 2-level form using OR gates followed by an AND gate. Remember to indicate any NOT gates required, since only uncomplemented input variables are available. [2 marks]

(*c*) Consider your implementation in part (*b*)(*ii*).

(*i*) Assume that the gates have finite propagation delay. Describe in detail what happens at the output F when the inputs $\{A, B, C, D\}$ change from $\{1, 1, 0, 1\}$ to $\{1, 1, 1, 1\}$. [4 marks]

(*ii*) Using a Karnaugh map or otherwise, determine the other single input variable change that will give rise to a similar problem to that observed in part (*c*)(*i*). [2 marks]

(*iii*) Using a Karnaugh map or otherwise, determine a modified POS expression for F that will eliminate the problems observed in parts (*c*)(*i*) and (*c*)(*ii*). [2 marks]

2 Digital Electronics

(*a*) Briefly describe what is meant by synchronous logic. Show how a Master-Slave D-type Flip-Flop may be constructed from two transparent D-latches and describe its operation with the help of a timing diagram. [7 marks]

(*b*) With the use of appropriate diagrams, briefly explain the operation of Moore and Mealy finite state machines, paying particular regard to their differences. [4 marks]

(*c*) A two-bit synchronous binary Up/Down (U/D) counter is capable of either up-counting (e.g., $0, 1, 2, 3, 0, \ldots$) or down-counting (e.g., $3, 2, 1, 0, 3, \ldots$) and randomly changes between these two modes of operation. It has outputs X and Y, where X is the Most Significant Bit (MSB).

The U/D counter is connected to a count Direction Detection System (DDS) that has two outputs, namely C_U and C_D, where C_U is required to give a binary 1 pulse when the U/D counter up-counts and C_D is required to give a binary 1 pulse when the U/D counter down-counts, otherwise the two outputs are both to remain at binary 0.

Assume that the count DDS has two state registers, each implemented as a D-type Flip-Flop (FF). The next state outputs of each FF, namely Q_X and Q_Y are given by the current inputs to the DDS, i.e., the U/D counter output bits X and Y respectively. Also assume that the FFs in the count DDS are clocked at a much higher rate than the U/D counter.

(*i*) Draw a state diagram for the count DDS, where the arcs connecting the states show the bits X and Y input to the count DDS, and also the output signals C_U and C_D. [4 marks]

(*ii*) Determine the combinational logic required in the count DDS to generate C_U and C_D from the inputs X and Y, and from the two FF outputs, namely Q_X and Q_Y. [5 marks]

SECTION B

3 Operating Systems

Consider a simple operating system where live processes are either *running*, *ready* to run, or *blocked* on an event, before they *exit*.

(*a*) State four conditions under which the operating system will try to schedule processes. [4 marks]

(*b*) OS schedulers are said to be *pre-emptive* or *non-preemptive*. State the principal problem with non-preemptive schedulers. Explain how pre-emptive schedulers solve this problem. [2 marks]

(*c*) Explain why pre-emptive schedulers are more complex to implement. [2 marks]

(*d*) Consider four CPU-bound processes arriving to a Shortest Remaining Time First (SRTF) scheduler as follows:

Process	**Arrival Time**	**Duration**
P_1	0	8
P_2	3	3
P_3	5	1
P_4	9	4

(*i*) Give a schedule computed by the SRTF scheduler, and compute the average waiting time across all four processes. Your answer should be explicit about the state of each process and the *ready* queue at all times. Clearly state any assumptions you make. [6 marks]

(*ii*) Explain why SRTF is difficult to implement in practice, and propose how to address this difficulty. [2 marks]

(*iii*) Assume that jobs are no longer CPU-bound but also perform blocking and non-blocking IO. Discuss how this can affect the effectiveness and fairness of the SRTF scheduler. [4 marks]

4 Operating Systems

(*a*) What is stored in a *directory entry* and what is stored in a *file entry*? [2 marks]

(*b*) Consider a file system that uses a fixed size 128 byte directory entry structure, a fixed size 32 byte file entry structure, and manages storage via the "chaining in a map" technique, on a disk that uses 4 kByte blocks. Assume that the map of blocks is permanently in memory and need not be read from disk. Assume further that the file system contains 2^{20} files.

(*i*) Assume this file system is structured as a *file system with a single level directory*. Explain how many blocks must be read in order to open and read a 16 byte file named `/a/b/c/d/e`. [3 marks]

(*ii*) Instead, assume this file system is structured as *a simple hierarchical file system with directories that can contain up to 256 sub-directories and 1024 files*. Give upper and lower bounds on how many blocks must be read in order to open and read a 16 byte file named `/a/b/c/d/e` where `/` denotes the path separator. State any assumptions you make in computing your bounds. [10 marks]

(*c*) Unix uses a more complex file entry structure than assumed above, known as an *i-node*. In addition to file metadata, the i-node stores 12 direct blocks, plus single, double and triple indirect blocks (each of which has 512 entries). Explain what purpose is served by each of these four types of block. [3 marks]

(*d*) Using Unix i-nodes as above, compute the maximum size file that can be stored. Assume a block is 4 kBytes in size. [2 marks]

SECTION C

5 Software and Interface Design

The following design brief was prepared by the client for one of the 2016 Cambridge group design projects.

We would like you to make a hazard warning app on the smartwatch Pebble. When a user takes their Pebble for a jog, he/she can record GPS location and details of any potential hazard they come across on the road, such as broken pavement, dog dirt, flooding. After the run, they can pair their Pebble to a computer program or mobile app, and upload the data via an interactive map. This allows other Pebble users to download the data, and receive an alarm when they are near the hazard. The relatively constrained set of controls on the Pebble, and the constraint that the jogger will want to work quickly before cooling down, means that you will need to provide an intelligent text entry method.

(*a*) For each of the following software project phases, suggest a design model or representation that would be a helpful aid in the design process. For each of these, sketch an example to show what this model looks like, based on some part of the above design brief.

(*i*) Inception phase

(*ii*) Elaboration phase

(*iii*) Construction phase

(*iv*) Transition phase

[12 marks]

(*b*) For each of the sketched examples in part (*a*), describe how the design work so far could be evaluated before proceeding to the next phase. [4 marks]

(*c*) Choose two of the above design models, representations or evaluation methods, and explain how they would be done differently if the project plan was following an agile rather than spiral project management approach. [4 marks]

6 Software and Interface Design

Imagine you have been commissioned to design a system that will help students learn how to test and debug software in a typical classroom situation.

(*a*) Explain the difference between testing and debugging. [2 marks]

(*b*) Name *four* different approaches to testing, giving a brief definition of *each*. [8 marks]

(*c*) Prepare a preliminary design, as suited to the inception phase of a project, that could be presented to the client who has commissioned this educational testing application. You should include *two different kinds of diagram*, so that the client understands the overall structure of the proposed user interaction, and also the kinds of data that will be processed and stored when the system is operational. [10 marks]

SECTION D

7 Discrete Mathematics

You may use standard results provided that you mention them clearly.

(*a*) (*i*) State a sufficient condition on a pair of positive integers a and b so that the following holds:

$$\forall\, x, y \in \mathbb{Z}.\ \big(x \equiv y \pmod{a} \wedge x \equiv y \pmod{b}\big) \iff x \equiv y \pmod{ab}$$

[2 marks]

(*ii*) Recall that, for a positive integer m, we let $\mathbb{Z}_m = \{n \in \mathbb{N} \mid 0 \leq n < m\}$ and that, for an integer k, we write $[k]_m$ for the unique element of $\mathbb{Z}_m$ such that $k \equiv [k]_m \pmod{m}$.

Let a and b be positive integers and let k and l be integers such that $k\,a + l\,b = 1$. Consider the functions $f : \mathbb{Z}_{ab} \to \mathbb{Z}_a \times \mathbb{Z}_b$ and $g : \mathbb{Z}_a \times \mathbb{Z}_b \to \mathbb{Z}_{ab}$ given by

$$f(n) = \big(\, [n]_a\,,\ [n]_b \,\big) \quad , \quad g(x, y) = [k\,a\,(y - x) + x]_{ab}$$

Prove either that $g \circ f = \mathrm{id}_{\mathbb{Z}_{ab}}$ or that $f \circ g = \mathrm{id}_{\mathbb{Z}_a \times \mathbb{Z}_b}$. [8 marks]

(*b*) Let T^* denote the reflexive-transitive closure of a relation T on a set A.

For relations R and S on a set A, prove that if $\mathrm{id}_A \subseteq (R \cap S)$ then $(R \cup S)^* = (R \circ S)^*$.

Note: You may alternatively consider T^* to be defined as either

$$\bigcup_{n \in \mathbb{N}} T^{\circ n} \quad \text{, where } T^{\circ 0} = \mathrm{id}_A \text{ and } T^{\circ(n+1)} = T \circ T^{\circ n}$$

or as

$$\bigcap \{\, R \subseteq A \times A \mid (T \cup \mathrm{id}_A) \subseteq R \ \wedge\ R \circ R \subseteq R \,\}$$

or as inductively given by the rules

$$\frac{}{(x, y)}\ \big((x, y) \in T\big) \qquad \frac{}{(x, x)}\ (x \in A) \qquad \frac{(x, y)\ \ (y, z)}{(x, z)}\ (x, y, z \in A)$$

[10 marks]

8 Discrete Mathematics

(*a*) (*i*) Calculate two integers x and y satisfying $177\,x + 78\,y = 3$. [3 marks]

(*ii*) Describe all the integer pairs (x, y) that satisfy the above equation. [3 marks]

(*b*) Let $\mathbb{N}_{\geq 5} = \{n \in \mathbb{N} \mid n \geq 5\}$. Prove that: $\forall\, n \in \mathbb{N}_{\geq 5}.\, 2^n > n^2$. [6 marks]

(*c*) Let $\mathrm{C}(X, Y) = \{\, S \subseteq X \mid S \cong Y \,\}$.

(*i*) For finite sets X and Y, what is the cardinality of $\mathrm{C}(X, Y)$ in terms of that of X and Y? [2 marks]

(*ii*) For elements a and b, and sets A and B such that $\{a, b\} \cap (A \cup B) = \emptyset$, consider the functions

- $f : \mathrm{C}\big(A \cup \{a\}, B \cup \{b\}\big) \longrightarrow \mathrm{C}(A, B) \uplus \mathrm{C}\big(A, B \cup \{b\}\big)$ given by

$$f(S) = \begin{cases} \big(0, S \setminus \{a\}\big) & , a \in S \\ (1, S) & , a \notin S \end{cases}$$

and

- $g : \mathrm{C}(A, B) \uplus \mathrm{C}\big(A, B \cup \{b\}\big) \longrightarrow \mathrm{C}\big(A \cup \{a\}, B \cup \{b\}\big)$ given by

$$g(i, S) = \begin{cases} S \cup \{a\} & , i = 0 \\ S & , i = 1 \end{cases}$$

Prove either that $g \circ f = \mathrm{id}$ or that $f \circ g = \mathrm{id}$. [6 marks]

9 Discrete Mathematics

(a) Let p and m be positive integers such that $p > m$.

(i) Prove that $\gcd(p, m) = \gcd(p, p - m)$. [3 marks]

(ii) Without using the Fundamental Theorem of Arithmetic, prove that if $\gcd(p, m) = 1$ then $p \mid \binom{p}{m}$. You may use any other standard results provided that you state them clearly. [3 marks]

(b) Let A^* denote the set of strings over a set A.

For a function $h : X \to Y$, let $\mathrm{map}_h : X^* \to Y^*$ be the function inductively defined by

$$\begin{array}{rcll} \mathrm{map}_h(\varepsilon) & = & \varepsilon & \\ \mathrm{map}_h(x\,\omega) & = & \big(h(x)\big)\,\big(\mathrm{map}_h(\omega)\big) & \quad (x \in X, \omega \in X^*) \end{array}$$

Prove that, for functions $f : A \to B$ and $g : B \to C$,

$$\mathrm{map}_g \circ \mathrm{map}_f \;=\; \mathrm{map}_{g \circ f}$$

Note: You may use the following Principle of Structural Induction for properties $P(\omega)$ of strings $\omega \in A^*$:

$$\big(P(\varepsilon) \wedge \forall\, \omega \in A^*.\, P(\omega) \Rightarrow \forall\, a \in A.\, P(a\,\omega)\big) \implies \forall \omega \in A^*.\, P(\omega)$$

[6 marks]

(c) We say that a relation $T \subseteq A \times B$ is a *total cover* whenever $\mathrm{id}_A \subseteq T^{\mathrm{op}} \circ T$ and $\mathrm{id}_B \subseteq T \circ T^{\mathrm{op}}$. (Recall that $T^{\mathrm{op}} \subseteq B \times A$ denotes the opposite, or dual, of the relation $T \subseteq A \times B$.)

For a relation $R \subseteq \{1, \ldots, m\} \times \{1, \ldots, n\}$ $(m, n \in \mathbb{N})$, we define a new relation $\stackrel{R}{\rightsquigarrow}$ between strings over a set X as follows: for all $u, v \in X^*$,

$$\begin{array}{rcl} u \stackrel{R}{\rightsquigarrow} v & \iff & R \text{ is a total cover and} \\ & & u = a_1 \ldots a_m,\ v = b_1 \ldots b_n, \text{ and } a_i = b_j \text{ for all } (i, j) \in R \end{array}$$

(i) Prove that for $R = \mathrm{id}_{\{1,\ldots,m\}}$, we have that $u \stackrel{R}{\rightsquigarrow} u$ for all $u = a_1 \ldots a_m$.

(ii) Prove that $u \stackrel{R}{\rightsquigarrow} v$ implies $v \stackrel{R^{\mathrm{op}}}{\rightsquigarrow} u$.

(iii) Prove that $u \stackrel{R}{\rightsquigarrow} v$ and $v \stackrel{S}{\rightsquigarrow} w$ imply $u \stackrel{S \circ R}{\rightsquigarrow} w$.

(iv) Prove that the further relation $\sim$ on X^* defined by

$$u \sim v \iff \exists\, R.\, u \stackrel{R}{\rightsquigarrow} v$$

is an equivalence relation.

[8 marks]

10 Discrete Mathematics

(*a*) Let $\Sigma = \{a, b, c\}$. Consider each of the subsets of Σ^* defined by the following groups of axioms and rules, and for each prove or disprove that $\#_a(u) \geq \#_b(u)$ for all $u \in \Sigma^*$, where $\#_x(u)$ is the number of occurrences of the symbol x in the string u.

(*i*) $\dfrac{}{\epsilon} \quad , \quad \dfrac{u}{aaub} \quad , \quad \dfrac{u,v}{ucv}$ for all $u, v \in \Sigma^*$

(*ii*) $\dfrac{}{a} \quad , \quad \dfrac{u}{au} \quad , \quad \dfrac{u}{ubc}$ for all $u \in \Sigma^*$

[6 marks]

(*b*) For each of the subsets in part (*a*), indicate with justification whether they are regular languages.
Note: Complete proofs are not necessary but you should clearly outline any proof strategy. [10 marks]

(*c*) For two regular expressions r and s and an alphabet Σ, define r & s to match a string in Σ^* if both r and s do. Given Kleene's Theorem, sketch a proof that the set of strings matched by r & s is a regular language for any regular expressions r and s. [4 marks]

END OF PAPER