SUPERVISOR TO ATTACH
PROCESSING LABEL HERE

Write your **student number** in the boxes above. **Letter**

Algorithmics (HESS)

Question and Answer Book

VCE Examination – Friday 8 November 2024

- Reading time is **15 minutes**: 3.00 pm to 3.15 pm
- Writing time is **2 hours**: 3.15 pm to 5.15 pm

Approved materials

- One scientific calculator

Materials supplied

- Question and Answer Book of 32 pages
- Multiple-Choice Answer Sheet

Instructions

- Follow the instructions on your Multiple-Choice Answer Sheet.
- At the end of the examination, place your Multiple-Choice Answer Sheet inside the front cover of this book.

Students are **not** permitted to bring mobile phones and/or any unauthorised electronic devices into the examination room.

Contents

VICTORIAN CURRICULUM AND ASSESSMENT AUTHORITY

VICTORIA State Government

Section A – Multiple-choice questions

Instructions

- Answer **all** questions in pencil on your Multiple-Choice Answer Sheet.
- Choose the response that is **correct** or that **best answers** the question.
- A correct answer scores 1; an incorrect answer scores 0.
- Marks will **not** be deducted for incorrect answers.
- No marks will be given if more than one answer is completed for any question.

Use the Master Theorem to solve recurrence relations of the form shown below.

$$T(n) = \begin{cases} aT\left(\frac{n}{b}\right) + kn^c & \text{if } n > 1 \\ d & \text{if } n = 1 \end{cases} \quad \text{where } a > 0,\ b > 1,\ c \geq 0,\ d \geq 0,\ k > 0$$

$$\text{and its solution } T(n) = \begin{cases} O(n^c) & \text{if } a < b^c \\ O(n^c \log n) & \text{if } a = b^c \\ O(n^{\log_b a}) & \text{if } a > b^c \end{cases}$$

Question 1

A garden nursery is performing a stock inventory of all its plants. The nursery owner wishes to be able to look up the name of any individual plant and receive as output the total number of that plant currently in stock.

Which one of the following abstract data types (ADTs) is the most suitable for modelling this situation?

A. array

B. dictionary

C. list

D. priority queue

Use the following information to answer Questions 2 and 3.

The graph G is shown below.

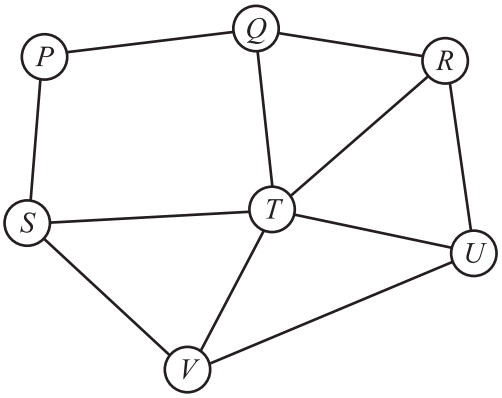

Question 2

Which of the following categories of graphs is G a member of?

A. directed graphs

B. complete graphs

C. connected graphs

D. trees

Question 3

A depth-first search algorithm is performed on G starting from Node P.

The order in which the algorithm visits nodes could be

A. P, Q, R, S, T, U, V

B. P, Q, R, T, S, U, V

C. P, Q, S, R, T, V, U

D. P, S, V, U, T, Q, R

Question 4

The algorithm that is most fit-for-purpose to find the solution to the all-pairs shortest path problem is

A. the Bellman-Ford algorithm.

B. Dijkstra's algorithm.

C. the Floyd-Warshall algorithm.

D. Prim's algorithm.

Question 5

The algorithm *ShortestPath* (G, u, v) has three inputs:

- a graph, G,
- two nodes, u and v.

The algorithm returns the shortest distance from u to v in G. Nodes in G are sequentially labelled using integers, starting from 0.

Which of the following algorithms has the least time complexity?

A.
```
Algorithm A(G, v):
  Foreach node u in G Do
    ShortestPath(G, u, v)
```

B.
```
Algorithm B(G):
  Foreach node u in G Do
    Foreach node v in G Do
      ShortestPath(G, u, v)
```

C.
```
Algorithm C(G, v):
  u ← 0
  While u < v Do
    ShortestPath(G, u, v)
    Increment u by 1
```

D.
```
Algorithm D(G, u, v):
  If u in G Do
    If v in G Do
      ShortestPath(G, u, v)
```

Question 6

The mergesort algorithm is an example of which algorithm design pattern?

A. brute-force search

B. divide and conquer

C. dynamic programming

D. greedy

Question 7

Which of the following properties is **not** a required element of a solution to the travelling salesman problem?

A. The path must include every edge in the graph.

B. The path must include every node in the graph.

C. The path must have the minimum possible total cost.

D. The path must be closed.

Question 8

Which one of the following is the signature specification for the 'in' operation of a set?

A. set $\rightarrow$ element

B. set $\times$ element $\rightarrow$ set

C. set $\times$ element $\rightarrow$ element

D. set $\times$ element $\rightarrow$ Boolean

Question 9

Consider the following algorithm.

```
Algorithm A(n):
  a ← 0
  b ← 10
  While n > 0 Do
    a ← b - a
    b ← 2n + b
    n ← n - 1
  Return a + b
```

What value is returned when $A(4)$ is run?

A. 30

B. 42

C. 44

D. 48

Question 10

Consider the following algorithm.

```
Algorithm B(x, y):
  If x ≤ 0 Do
    Return x
  If y ≤ 0 Do
    Return x
  Return B(x - y, y - 1)
```

What is the value returned by $B(30, 4)$?

A. 10

B. 14

C. 18

D. 20

Question 11

Consider the worst-case running time of four algorithms:

Algorithm $P - O(\log(n))$

Algorithm $Q - O(n)$

Algorithm $R - O(n^{\log_3 2})$

Algorithm $S - O\left(\sqrt{n}\right)$

The order of these algorithms when sorted from best to worst running time for large n is:

A. P, Q, R, S

B. P, S, R, Q

C. Q, R, S, P

D. S, R, Q, P

Question 12

The recurrence relation for $T(x)$ given below describes the running time of an algorithm for an input of size x.

$$T(x) = \begin{cases} T\left(\frac{x}{2}\right) + 2x & \text{if } x > 1 \\ 2 & \text{if } x = 1 \end{cases}$$

What is the tightest Big-O time complexity of the algorithm?

A. $O(\log(x))$

B. $O\left(\sqrt{x}\right)$

C. $O(x)$

D. $O(x \log(x))$

Question 13

Consider the following sorted list of eleven integers.

[1, 4, 5, 9, 9, 10, 18, 19, 20, 25, 27]

The binary search algorithm is called on the list to search for the value 9.

What is the second comparison value used by the algorithm?

A. 5

B. 9

C. 10

D. 20

Question 14

What is the difference between strong AI and weak AI?

A. Strong AI learns patterns of inputs and outputs, while weak AI understands semantics and uses reasoning to deduce actions in new situations.

B. Strong AI uses support vector machines (SVMs) to understand semantics of language, while weak AI uses decision trees to understand semantics of language.

C. Strong AI understands semantics and uses reasoning to deduce actions in new situations, while weak AI learns patterns of inputs and outputs.

D. Strong AI uses neural networks to understand semantics of language, while weak AI uses a binary classifier to understand semantics of language.

Question 15

Why is the undecidability of the Halting Problem important for our understanding of the limits of computation?

A. The Halting Problem was an early problem to be proved as undecidable, and subsequent computationally difficult problems can be shown to be undecidable from it.

B. If we do not know a problem halts, then it is difficult to understand what the limits of computation are.

C. The Halting Problem shows us that the only computational problems that can be solved are those that can be solved on a Turing machine.

D. If a problem sometimes halts, or sometimes goes into an infinite loop, then it is difficult for us to understand how long it takes for it to run.

Question 16

Which one of the following statements is **false** about SVMs?

A. SVMs can classify higher dimensional data.

B. SVMs use backward propagation to determine their parameters.

C. SVMs can classify data that can be separated by a straight line.

D. SVMs find a decision boundary that is equally distant from the closest points on each side of the boundary.

Question 17

In 2014, a company began developing an algorithm to review resumés to help filter job applications. The algorithm was trained using resumés that had been previously submitted to the company. This data had significantly more applications from males. It was discovered the resulting algorithm gave lower ratings to applicants who attended some all-women colleges and whose resumés included the word 'women's'

What kind of ethical issue is this an example of?

A. accountability

B. bias

C. machine ethics

D. transparency

Question 18

A problem Y is in the NP-Hard time computational complexity class if

A. solutions to Y can be verified in polynomial time.

B. Y can be reduced to all problems in P in polynomial time.

C. Y can be reduced to all problems in NP in polynomial time.

D. all problems in NP can be reduced to Y in polynomial time.

Question 19

In the change-making problem, the variables are:

- an amount A,
- a set of coin denominations $d_1, d_2, \ldots, d_n$, and
- a set of coin quantities $x_1, x_2, \ldots, x_n$.

The value that is optimised by the change-making problem is

A. $\sum_{i=1}^{n} x_i$

B. $\sum_{i=1}^{n} d_i$

C. $\sum_{i=1}^{n} d_i x_i$

D. $\sum_{i=1}^{n} \frac{d_i}{x_i}$

Question 20

If the A* algorithm is run with $h(n) = 0$ then it is equivalent to

A. the Bellman-Ford algorithm.

B. breadth-first search.

C. Dijkstra's algorithm.

D. the Floyd-Warshall algorithm.

End of Section A

Section B

Instructions

- Answer **all** questions in the spaces provided.
- Write your responses in English.

Use the Master Theorem to solve recurrence relations of the form shown below.

$$T(n) = \begin{cases} aT\left(\frac{n}{b}\right) + kn^c & \text{if } n > 1 \\ d & \text{if } n = 1 \end{cases} \quad \text{where } a > 0,\ b > 1,\ c \geq 0,\ d \geq 0,\ k > 0$$

$$\text{and its solution } T(n) = \begin{cases} O(n^c) & \text{if } a < b^c \\ O(n^c \log n) & \text{if } a = b^c \\ O(n^{\log_b a}) & \text{if } a > b^c \end{cases}$$

Question 1 (3 marks)

Write a definition for a directed acyclic graph.

Question 2 (3 marks)

Explain the characteristics of the stack abstract data type (ADT) and give an example application of a stack.

Question 3 (5 marks)

At a local dry-cleaning store, customers drop off their dirty clothes to be cleaned. The store offers three different types of cleaning:

1. Same-day cleaning – a priority cleaning that is completed in a matter of hours.
2. Standard cleaning – the standard option that usually takes a couple of days.
3. Deep cleaning – an option for heavily soiled items that takes about four days.

The following information is stored about each job:

- a job identification number
- the type of cleaning required (same-day, standard, deep)
- the phone number of the customer
- the date and time that the clothing can be picked up by the customer.

The store can only process one job at a time. It attempts to process jobs on a first-come, first-served basis, but has to prioritise same-day cleaning jobs when they arrive.

Describe a combination of ADTs that the dry-cleaning store could use to store the information about each job that would allow the business to easily identify information about each job and the order in which jobs should be processed. Justify why each ADT is required and what aspect of the situation each ADT is modelling.

Question 4 (5 marks)

An electrician has been tasked with installing electrical cabling in a new building.
The electrician must connect all the switches and power points to a central power source using the smallest total length of cabling possible.

a. Explain how the features of this problem could be represented using the graph ADT. 2 marks

b. The electrician considers using one of Prim's algorithm, Dijkstra's algorithm or the Bellman-Ford algorithm to assist in finding the shortest total length of cabling required to connect all the switches and power points to the central power source.

What algorithm would be most suitable for the electrician to use to solve this problem? Justify your answer by comparing it to the other listed options. 3 marks

Question 5 (3 marks)

Consider the following neural network.

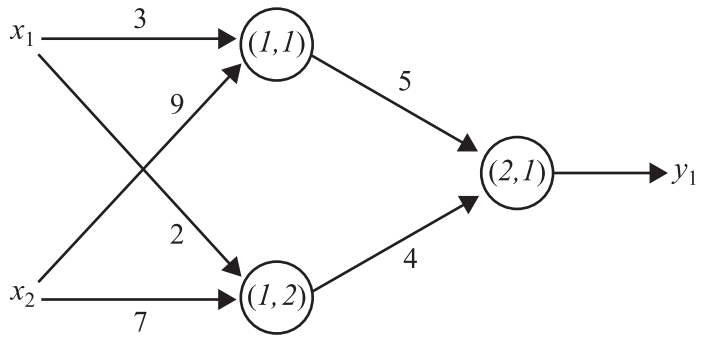

For a given node k, let a_k denote the total node activation or input value and let o_k denote the output value.

a. Nodes $(1,1)$, $(1,2)$ and $(2,1)$ all use a linear transfer function, where for a given node, k

$$o_k = a_k$$

Calculate the output value, y_1, for the input $x_1 = 1$ and $x_2 = 4$. 2 marks

b. The neural network from **part a** is modified so that node $(2,1)$ uses the following transfer function:

$$o_k = \begin{cases} 1, a_k \geq 50 \\ 0, a_k < 50 \end{cases}$$

The same inputs $x_1 = 1$ and $x_2 = 4$ are given to the neural network.

Describe how the output of the neural network would differ from that of the output calculated in **part a**. 1 mark

Do not write in this area.

Question 6 (6 marks)

The grid below contains six regions, each denoted by a letter from the alphabet from A to F. All pairs of adjacent regions have a property called their 'accessibility', which is calculated as the number of shared borders between those two regions. For example:

- The accessibility of regions B and A is 2.
- The accessibility of regions B and C is 0.

F	F	F	F	D
F	B	D	D	D
E	B	A	D	D
E	B	A	C	C
E	E	C	C	C

a. Draw an undirected graph that represents this grid. Include a node to represent each region and a weighted edge for all node pairs with non-zero accessibility. Label each edge with the accessibility of its adjacent nodes. 2 marks

Question 6 continues on the next page.

b. Given a grid, g, where $g(x, y)$ returns the region label of the grid square in column x, row y, the following algorithm creates a graph from this grid.

```
Algorithm regionGraph(g):
  Let r be the number of rows in g
  Let c be the number of columns in g
  Create an empty graph H
  Foreach x in {1, ..., c} Do
    Foreach y in {1, ..., r} Do
      If g(x, y) is not a node in H Do
        Create a node called g(x, y) in H
  Foreach pair of nodes u, v in H Do
    Create an edge (u, v) in H with a weight of 0
  Foreach x in {1, ..., (c - 1)} Do
    Foreach y in {1, ..., (r - 1)} Do
      If g(x, y) ≠ g(x + 1, y) Do
        Add 1 to the weight of edge (g(x, y), g(x + 1, y)) in H
      If g(x, y) ≠ g(x, y + 1) Do
        Add 1 to the weight of edge (g(x, y), g(x, y + 1)) in H
  Foreach edge e in H Do
    If weight of e = 0 Do
      Remove e from H
```

For the purposes of this problem, the following operations can be performed in constant time:

- creating an empty graph
- creating a new node in a graph
- checking if a node exists in a graph
- creating a new edge or removing an edge in a graph
- accessing or changing the weight of an edge in a graph.

Analyse this algorithm to determine its tightest Big-O time complexity with respect to:

- r, the number of rows,
- c, the number of columns, and
- n, the number of regions in the grid.

4 marks

Question 7 (4 marks)

Jane and Aniket are starting an online clothing company that employs just-in-time manufacturing and sends customer orders directly to a factory, which manufactures them to the customer's requirements. Jane has studied artificial intelligence (AI) and wants to embed AI into the website to allow users to upload or take photos of themselves and virtually try on the available clothing designs and colours. The AI will automatically estimate clothing sizes and suggest designs based on the past browsing and purchase history of the user and other similar users.

a. Describe two potential risks or harms that could arise from this use of AI to fit clothing onto real photos of users. 2 marks

- _____
- _____

b. In the first six months of the company's operation, 5000 users sign up and, on average, each user tries on a few pieces of clothing. The company uses this data to train a neural network with 64 layers and hundreds of millions of parameters.

Describe the problem that would result from training a neural network of this size with this amount of training data. 2 marks

_____

_____

_____

_____

_____

_____

Question 8 (4 marks)

A small convenience store has a fixed amount of shelf space. This shelf space is divided into displays of similar items, such as types of drinks or snacks. These sections are fixed in size as they contain a group of similar items from the same brand.

From historical sales data, the owner has a reliable understanding of what profit to expect from different types of stock. She is considering what she should display on the shelves to make the greatest profit. A sample of the display options is shown below.

Display	**Size** (shelf sections)	**Expected profit of display** ($/fortnight)
Cheepy Cheep chips	3	1000
White Nugget premium chips	4	1400
Farmer's Bounty vegetable chips	3	800

a. What type of computational problem is the store owner's problem analogous to? 1 mark

b. Describe a heuristic method that could be used to optimise the store's shelves to be most profitable. 3 marks

Question 9 (7 marks)

Isabelle and Christopher are organising the painting of a new picket fence. There are five painters available and each painter takes 1 unit of time to paint 1 linear metre of fence picket. The fence is made up of n vertical pickets of length $\{L_1, L_2, L_3, ..., L_n\}$, where $n \geq 20$.

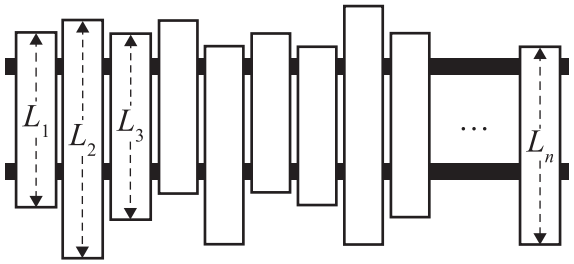

Isabelle and Christopher wish to find the minimum time required to get the fence painted, given that a picket can only have one painter painting it and any given painter will only paint a set of adjacent pickets. For example, an individual painter might paint pickets 1, 2 and 3, but they will not paint pickets 1, 3 and 5.

a. Write pseudocode for a brute-force algorithm that will return the minimum time required to paint the fence. 5 marks

b. A picket fence for a house might have 50 pickets, while a picket fence for a large school might have 50 000 pickets.

Discuss whether the brute-force algorithm described in **part a** could feasibly be applied to this range of picket fence sizes. 2 marks

Question 10 (3 marks)

Compare the backtracking algorithm design and brute-force search algorithm design strategies to explain how backtracking improves brute-force search.

Question 11 (3 marks)

The algorithm *equalDegreeNodes* takes one input, a connected graph G with at least two nodes, and identifies and returns a pair of nodes that have the same degree.

```
   Algorithm equalDegreeNodes(G):
1    Let V be the set of nodes in G
2    Foreach d in {|V| - 1, ..., 1} Do
3      Foreach u in V Do
4        Foreach v in V with u ≠ v Do
5          If degree(u) = d and degree(v) = d Do
6            Return (u, v)
7    Return None
```

Use an argument by contradiction to show that the algorithm *equalDegreeNodes* will always return a pair of nodes for all simple, undirected, connected graphs with at least two nodes.

Do not write in this area.

Question 12 (9 marks)

In the puzzle game of Oriento, there are nine round dials arranged in a 3 × 3 grid. Each dial has an arrow that indicates the dial's current direction, and these arrows point either up, down, left or right. An example grid is shown below.

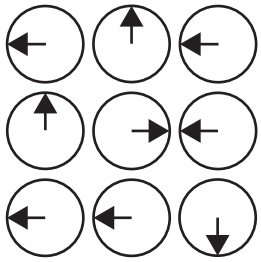

The dials can be turned in 90-degree increments, either clockwise or anticlockwise. When a dial is manually turned, any adjacent dials above, below, left or right of that dial also turn by the same amount.

For example, if the middle row, left dial is turned 90 degrees clockwise in the above grid, it would result in the grid shown below.

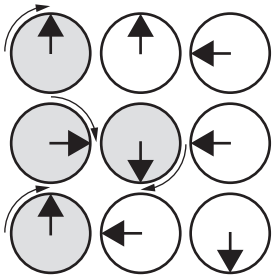

The goal of the game is to orient all the arrows to point in the same direction as quickly as possible. All possible starting arrangements of the grid can be solved to point in any of the four directions.

a. Describe an ADT that could be used to store information about a single Oriento grid and describe how features of the grid map to features of the ADT. 2 marks

Question 12 continues on the next page.

The following information is also provided.

There are 262 144 possible Oriento grids. A graph, G, is created with this many nodes, each representing one possible state of the grid. For each pair of nodes, u and v, where node v's grid can be obtained by manually turning one dial in node u's grid, a directed edge (u, v) is created with the following two properties:

- dial: the dial to be turned
- turn: the amount of turn (−90, 90 or 180 degrees clockwise).

There are four specially labelled solution nodes, Up, Down, Left and Right, that correspond to the nodes that have all of the dials oriented in one of the four directions.

b. Sami would like to find out how to solve the game by manually turning the fewest number of dials possible.

Describe an algorithm that takes two inputs, the graph G and a starting node s, and efficiently returns the solution node that can be reached from s by turning the fewest dials. In the case of a tie, either solution can be returned. No pseudocode is required. 2 marks

c. In contrast, Zoni would like to find out how to solve the game by performing the least amount of total rotation. The total amount of rotation is the sum of the number of 90-degree turns manually performed on all nine of the dials.

Describe an algorithm that takes two inputs, the graph G and a starting node s, and efficiently returns the solution node that can be reached from s by performing the least amount of total rotation of the dials. No pseudocode is required. 3 marks

d. Zoni creates a program that solves the Oriento game and builds it into a robot that can solve the puzzle.

With reference to one particular conception of artificial intelligence, explain whether Zoni's robot possesses artificial intelligence. 2 marks

Question 13 (8 marks)

The algorithm *CD* draws a size *s* diamond shape composed of cubes. Below is a sample run of algorithm *CD* with inputs 1 to 5.

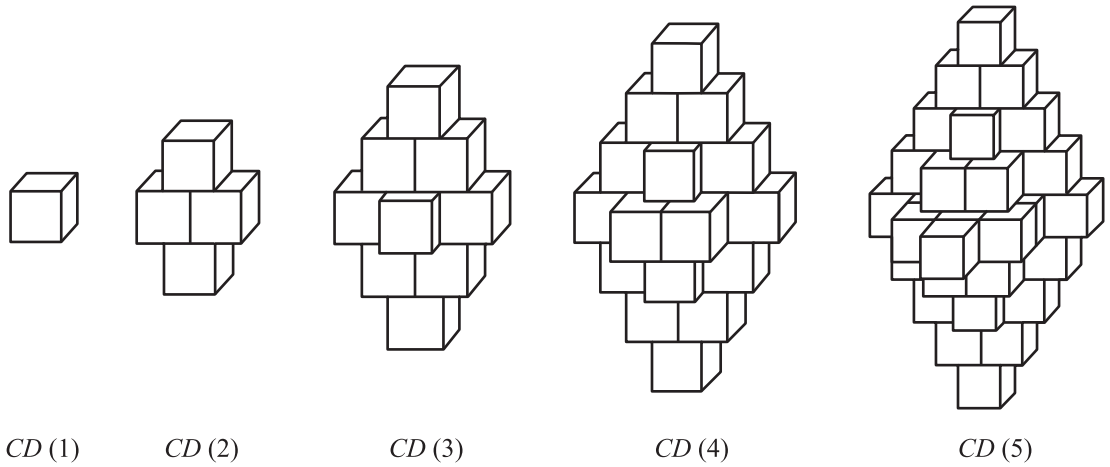

CD (1) *CD* (2) *CD* (3) *CD* (4) *CD* (5)

The algorithm *CD* draws the diamond one level at a time from top to bottom and uses these two functions:

- *CalculateNumberOfCubes*(*level, column*), which computes the number of cubes in the given vertical level and horizontal column of the diamond
- *DrawCubes*(*level, column, n*), which draws *n* cubes in the position corresponding to the given vertical level and horizontal column of the diamond, with those cubes drawn centred to form a diamond shape. This function has a constant runtime regardless of its inputs.

```
   Algorithm CD(s):
1    Foreach level in {0, ..., (s - 1)} Do
2      Foreach column in {0, ..., level} Do
3        n ← CalculateNumberOfCubes(level, column)
4        DrawCubes(level, column, n)
5    Foreach level in {(s - 2), ..., 0} Do
6      Foreach column in {0, ..., level} Do
7        n ← CalculateNumberOfCubes(level, column)
8        DrawCubes(level, column, n)
```

```
   Algorithm CalculateNumberOfCubes(level, column):
1    ans ← 1
2    If (column > level - column) Do
3      column ← level - column
4    Foreach i in {0, ..., (column - 1)} Do
5      ans ← ans × (level - i)
6      ans ← ans ÷ (i + 1)
7    Return ans
```

Do not write in this area.

a. Analyse the algorithm CD and find its tightest Big-O time complexity with respect to the size of the diamond, s. 4 marks

Question 13 continues on the next page.

b. Each level of the diamond is composed of a two-dimensional grid of cubes. For those levels in the top half of the diamond, the number of columns in each level grows linearly with the height of the diamond and the number of cubes in each column is the sum of the number of cubes in the two adjacent columns in the level above it.

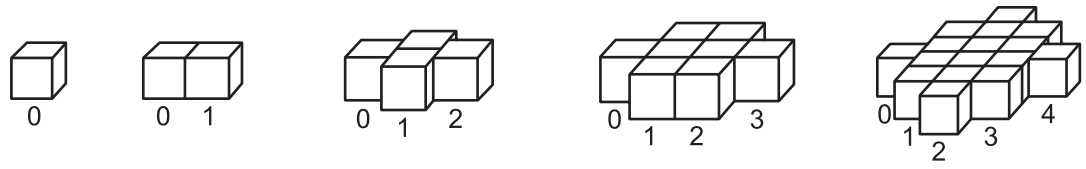

Level 1 Level 2 Level 3 Level 4 Level 5

A new function is created, *DrawLevel*(*level*, *columns*), which takes two inputs, the level and an array of column sizes, and draws all of the cubes in the level. This function has a runtime linearly proportional to the length of *columns*.

Apply the dynamic programming design pattern to design an algorithm with improved efficiency for drawing a diamond of cubes. Write your algorithm in pseudocode. Your algorithm should use the *DrawLevel* function. 4 marks

Do not write in this area.

Question 14 (5 marks)

Mirabelle and Tiffany are currently studying computational theory.

a. While studying, they find the following statement in an old textbook they found in the library.

'If the computational complexity classes P = NP, then we can solve many problems that are thought intractable.'

Do you agree with this statement? Explain your reasoning. 2 marks

b. Later, in the same book, they find the following statement.

'NP-Hard problems are the same as NP-Complete problems.'

Do you agree with this statement? Explain your reasoning and provide an example that supports your claim. 3 marks

Question 15 (2 marks)

Lila has an algorithm that recursively reduces its search space by half as it searches through the input data. She expects it to have a logarithmic running time. However, after testing the algorithm, it was determined that its running time grows quadratically with the size of the input.

Give a possible explanation for this result.

Question 16 (10 marks)

In the computer game Griddo, players move a character around a map. The map is a two-dimensional, square grid. Each square of the grid has a type of terrain.

- Field squares allow for normal movement.
- Forest squares can be moved through, but only at half speed.
- Mountain squares cannot be moved through.

During the game, both field and forest squares may be temporarily modified with rain, halving the normal movement speed through that square.

The location of fields, forests and mountains does not change during the game and, hence, it is possible to pre-compute the shortest paths between locations on the map without the presence of rain.

An example of a small map is shown below. Maps in the actual game can have millions of squares.

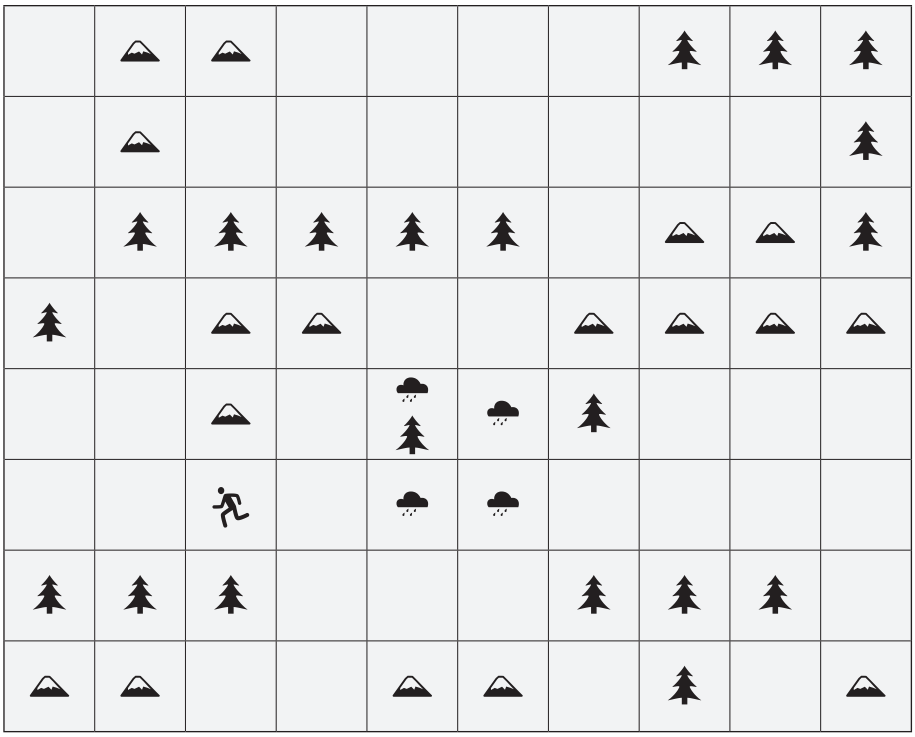

Key	
	Field
	Forest
	Mountain
	Rain
	Character

Characters can move up, down, left or right to adjacent squares around the grid.

During the game, players can instruct their character to move to a location. When this is done, the game needs to determine the path that the character will take to get there and begin moving the character. It needs to both determine the path quickly and also find a fast path for the character to take to get to the destination.

- Describe how the computer should represent this problem.
- Describe an algorithmic approach for determining a fast path for the character to take to travel to a target location on the map.
- Explain why your algorithmic approach would be the best choice on the basis of whether it would determine the path quickly and accurately.

Question 16 continues on the next page.

End of examination questions

This page is blank.