

STUDENT NUMBER Letter

ALGORITHMICS (HESS)

Written examination

Tuesday 9 November 2021

Reading time: 2.00 pm to 2.15 pm (15 minutes)

Writing time: 2.15 pm to 4.15 pm (2 hours)

QUESTION AND ANSWER BOOK

Structure of book

<i>Section</i>	<i>Number of questions</i>	<i>Number of questions to be answered</i>	<i>Number of marks</i>
A	20	20	20
B	11	11	80
			Total 100

- Students are permitted to bring into the examination room: pens, pencils, highlighters, erasers, sharpeners, rulers and one scientific calculator.
- Students are NOT permitted to bring into the examination room: blank sheets of paper and/or correction fluid/tape.

Materials supplied

- Question and answer book of 28 pages
- Answer sheet for multiple-choice questions

Instructions

- Write your **student number** in the space provided above on this page.
- Check that your **name** and **student number** as printed on your answer sheet for multiple-choice questions are correct, **and** sign your name in the space provided to verify this.
- All written responses must be in English.

At the end of the examination

- Place the answer sheet for multiple-choice questions inside the front cover of this book.

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

SECTION A – Multiple-choice questions**Instructions for Section A**

Answer **all** questions in pencil on the answer sheet provided for multiple-choice questions.

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.

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

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

Question 1

An abstract data type (ADT) has the following partial signature specification.

$$\text{Operation_1: } \textit{element} \times Y \rightarrow Y$$

The ADT Y could be

- A. a list ADT, where Operation_1 is the ‘add’ operation.
- B. an array ADT, where Operation_1 is the ‘lookup’ operation.
- C. a queue ADT, where Operation_1 is the ‘dequeue’ operation.
- D. a stack ADT, where Operation_1 is the ‘pop’ operation.

Question 2

Consider the following algorithm.

```

Algorithm numberbox(n)
  a ← 1
  b ← 1
  For each digit in n Do
    If digit is even Do
      a ← a + b
      b ← b + 1
    If digit is odd Do
      a ← a × b
      b ← b - 1
  Return a

```

What does the algorithm numberbox output when $n = 2021$?

- A. 3
- B. 11
- C. 12
- D. 28

Question 3

At one Victorian school, a student is allowed to enrol in advanced studies in a subject if they satisfy all of the following requirements:

- They are not studying a language outside of school.
- They are achieving high marks in the subject.
- They are not failing any other subjects.

Which one of the following conditions would be suitable for testing whether a student can study an advanced subject?

- A. NOT (studying a language outside of school AND failing other subjects) AND (getting high marks in the subject)
- B. NOT (studying a language outside of school OR failing other subjects) AND (getting high marks in the subject)
- C. NOT (studying a language outside of school OR failing other subjects) OR (getting high marks in the subject)
- D. NOT studying a language outside of school OR NOT failing other subjects OR getting high marks in the subject

Use the following information to answer Questions 4 and 5.

Consider the following algorithm that searches for a node s in a graph G , starting from node u .

```

Algorithm search( $G, u, s$ ) :
  If  $u = s$  Do
    Return True
  For each neighbour  $v$  of  $u$  Do
    result  $\leftarrow$  search( $G, v, s$ )
    If result = True Do
      Return True
  Return False

```

Question 4

This algorithm is an example of

- A. breadth-first search.
- B. depth-first search.
- C. best-first search.
- D. topological search.

Question 5

This algorithm, as written, would experience problems because

- A. it does not consider whether a node has already been visited.
- B. it returns both True and False, leading to a contradiction.
- C. if node u is equal to node s , the algorithm always returns True and does not complete the search.
- D. it uses recursion, whereas search algorithms work correctly only when written using iteration.

Question 6

In the PageRank algorithm, the score for each node in each iteration is calculated using the rule

$$\text{PR}(v) = \frac{1-d}{N} + d \cdot \sum_{u \in U} \frac{\text{PR}(u)}{\text{out-degree}(u)}$$

where U is the set of incoming neighbours of v .

The parameter d represents

- A. 0.15
- B. 0.85
- C. the probability that an explorer of the graph will randomly choose a new node
- D. the probability that an explorer of the graph will follow an edge from their current location

Question 7

A student finds that a weighted, connected graph G has two different minimal spanning trees, T_1 and T_2 . Which one of the following statements is necessarily true?

- A. The graph G must contain a cycle.
- B. The student has made a mistake as each graph has a unique minimal spanning tree.
- C. The weights of all edges in the graph G must be the same.
- D. The graph G is acyclic.

Question 8

Which one of the following statements is true?

- A. All connected graphs are trees.
- B. All trees are connected graphs.
- C. All connected graphs are cyclic graphs.
- D. All cyclic graphs are connected graphs.

Question 9

Which one of the following recurrence relations describes the worst case for the number of comparisons made by the quicksort algorithm on an input of n elements?

- A. $T(n) = 2T\left(\frac{n}{2}\right) + O(n)$
- B. $T(n) = 2T(n) + O(n)$
- C. $T(n) = T(n-1) + O(n)$
- D. $T(n) = T(n^2) + O(n)$

Use the following information to answer Questions 10 and 11.

Consider a country where the only coins in circulation are 1 cent, 3 cent and 4 cent coins. A citizen would like to determine the fewest coins required to pay for a chocolate bar. The following algorithm has been provided for this purpose. The last line of the algorithm, which makes three recursive calls to itself, is incomplete.

Input: x , the price of the chocolate bar given in cents

Algorithm minCoins(x):

If $x < 5$ **Then**

If $x = 2$ **Then**

Return 2

Else

Return 1

Else

Return _____

Question 10

Which one of the following correctly completes the algorithm above?

- A. $\text{minimum}(\text{minCoins}(x-1)+1, \text{minCoins}(x-3)+1, \text{minCoins}(x-4)+1)$
- B. $\text{minimum}(\text{minCoins}(x+1)-1, \text{minCoins}(x+3)-1, \text{minCoins}(x+4)-1)$
- C. $\text{minimum}(\text{minCoins}(x-1)-1, \text{minCoins}(x-1)-3, \text{minCoins}(x-1)-4)$
- D. $\text{minimum}(\text{minCoins}(x-1)+1, \text{minCoins}(x-1)+3, \text{minCoins}(x-1)+4)$

Question 11

Why is it infeasible to run this algorithm on large values of x ?


- A. This approach is greedy and would not return an optimal value when the value of x gets large.
- B. Recursive algorithms will necessarily take longer to execute than iterative algorithms.
- C. A very large number of repeated recursive calls would mean the algorithm is unlikely to terminate in a reasonable amount of time.
- D. The minimum function is likely to be slow, so calling it repeatedly will mean the algorithm is unlikely to terminate in a reasonable amount of time.

Question 12


Consider a Turing machine with the following state table.

State	Cell		
	Blank	0	1
A	1, left, halt	X	1, left, B
B	0, left, halt	0, left, B	1, left, A
halt			

The machine commences in state A with the tape below, on the cell indicated by the arrow.



The final state of the tape should be as follows, with the machine in the halt state.



What should the transition **X** be in the state table above to achieve this?

- A. 0, left, A
- B. 1, left, A
- C. 0, left, B
- D. 1, left, B

Question 13

Consider the recurrence relation

$$T(n) = \begin{cases} 5T\left(\frac{n}{2}\right) + 2n & \text{if } n > 1 \\ 2 & \text{if } n = 1 \end{cases}$$

$T(n)$ describes a function that is

- A. $O(n)$
- B. $O(n \log n)$
- C. $O(n^{\log_2 5})$
- D. $O(n^2)$

Question 14

An algorithm is known to have a running time that is $O(n^4)$ and $\Omega(n^3)$ for an input of size n .

Which one of the following statements about this algorithm is necessarily correct?

- A. The running time of the algorithm is $O(n^2)$.
- B. The running time of the algorithm is $O(n^3)$.
- C. The running time of the algorithm is $\Omega(n^2)$.
- D. The running time of the algorithm is $\Omega(n^4)$.

Use the following information to answer Questions 15 and 16.

Consider the following algorithm.

```

Algorithm direction(n, count)
    If count > 1
        If n is divisible by 2
            Return direction(n + 7, count - 1)
        Else
            Return direction((n-1)/2, count - 1)
    If the remainder of dividing n by 4 is 0 Return 'n'
    If the remainder of dividing n by 4 is 1 Return 'e'
    If the remainder of dividing n by 4 is 2 Return 's'
    If the remainder of dividing n by 4 is 3 Return 'w'
  
```

Question 15

What does the algorithm `direction` return with inputs $n = 14$ and $count = 6$?

- A. 'n'
- B. 'e'
- C. 's'
- D. 'w'

Question 16

This algorithm incorporates the concept of

- A. modularisation.
- B. recursion.
- C. iteration.
- D. divide and conquer.

Use the following information to answer Questions 17 and 18.

A school is using the following combination of ADTs to store information about its classes:

- a dictionary, where each key is a teacher's name and each value is an array of classes that the teacher teaches

$$\left\{ \begin{array}{l} \text{Ms Algo : ['7 English A', '9 History C', '12 English D']} \\ \text{Mr Rithmics : ['8 Maths B', '10 Maths C', '12 Further Maths E']} \end{array} \right\}$$

- an associated dictionary that stores the students in each class

$$\left\{ \begin{array}{l} \text{'7 English A': ['Robert Prim', 'Lester Ford', ...]} \\ \text{'8 Maths B': ['Richard Bellman', 'Edsger Dijkstra' ...]} \end{array} \right\}$$

Question 17

Which one of the following operations could **not** be performed in $O(1)$ time?

- Set the first item in the array for 'Ms Algo' to '7 English D'.
- Determine whether a given teacher teaches at the school.
- Determine which maths class a given student is enrolled in.
- Determine if there is a class called '7 History G' at the school.

Question 18

An algorithm iterates over each class to find the number of classes with 25 or more students.

The time complexity of this algorithm is

- $O(c)$
- $O(s)$
- $O(c + s)$
- $O(c \times s)$

Question 19

A Turing machine can solve all

- mathematical questions.
- mathematical questions that are computable.
- mathematical questions about natural numbers.
- mathematical questions except for the Halting Problem.

Question 20

In DNA computing, a part of the computation is performed by

- genetically engineering cells to perform calculations.
- simulating brain neurons using DNA.
- using algorithms that mutate themselves.
- mixing DNA to generate the entire solution space.

DO NOT WRITE IN THIS AREA

CONTINUES OVER PAGE

TURN OVER

SECTION B

Instructions for Section B

Answer **all** questions in the spaces provided.


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

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

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

Question 1 (5 marks)

Consider the following graph with edge weights as shown.



- a. What is the distance of the C-D-E-G path? 1 mark

- b. Write down the degree of node G. 1 mark

DO NOT WRITE IN THIS AREA

c. Consider running Dijkstra's algorithm, starting from node A, to find the shortest distance to all other nodes.

i. What property of this graph would make Dijkstra's algorithm unsuitable?

1 mark

ii. Find a node for which the distance from node A that Dijkstra's algorithm returns is incorrect. State the distance found by Dijkstra's algorithm and the true shortest distance, respectively, in your response.

2 marks

DO NOT WRITE IN THIS AREA

Question 2 (5 marks)

Alex is developing a computer simulation to model a chemical reaction. In the simulation, the reactor is initially empty and the reaction starts when one or more chemicals are added. More chemicals may be added while the reaction is in progress. A chemical reaction is characterised by what chemicals are added, how much of them are added and when they are added.

- a. To store a description of a chemical reaction for simulation, Alex needs the following information about each addition into the reactor:
- the name of the chemical added
 - the time the chemical was added, in seconds, after the reaction begins
 - the amount of the chemical added, in milligrams

Describe a combination of abstract data types (ADTs) that Alex could use to store a description of a chemical reaction. Justify your answer.

3 marks

- b. Alex intends to use the combination of ADTs described in **part a.** often and decides to define the combination of ADTs as a new ADT called ChemEvents.

For the ChemEvents ADT, as described in **part a.**, write the signature specification of the following operations:

- Add a new chemical to the reaction.
- Look up what chemical was added to the reaction at a given time.

2 marks

DO NOT WRITE IN THIS AREA

CONTINUES OVER PAGE

Question 3 (10 marks)


HexaQuesta is a game played on a board that is divided into hexagonal regions. There are three kinds of borders between regions:

- Normal borders take one move to cross.
- River borders take two moves to cross.
- Mountain borders cannot be crossed.

Game boards may contain thousands of hexagonal regions.

The objective of the game is to travel across the board from the start, S, to the finish, X, with the fewest number of moves possible.

An example of a HexaQuesta board, showing the start, S, and the finish, X



Samarth would like to write a program that determines how to win the game.

- a. Describe an ADT that appropriately models a board of the HexaQuesta game. 2 marks

DO NOT WRITE IN THIS AREA

- b. Write an algorithm, in pseudocode, that Samarth could use to find a winning path from the start to the finish on any board.

5 marks

- c. Samarth would like to find the best starting region on a given board. He decides that the best starting region is the one with the lowest average number of moves to travel to every other region on the board.

Describe how Samarth could efficiently find the best starting region on the board and derive a running time for your solution. Do **not** provide pseudocode in your answer.


3 marks

Question 4 (14 marks)

A delivery company, Pilliga Postage, wants to introduce automated robots that will navigate its warehouses to store and retrieve items. Shown below is a two-dimensional floor plan of one of its warehouses, where the lines represent the paths that a robot can move along, and the intersections are the only positions where the robot can stop and change direction.

Multiple robots will move around each warehouse at any time. In a given warehouse, every robot will be fitted with a navigation system that allows it to track the movements of robots.

An example of a warehouse layout



- a. Pilliga Postage has decided to use a weighted, directed graph to model the warehouse floor.

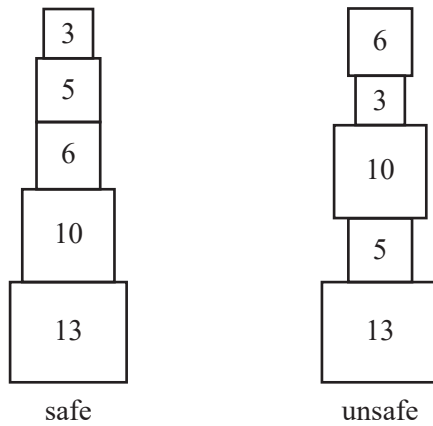
Explain what could be represented by the weights and the direction in Pilliga Postage’s model. 2 marks

- b. A robot is required to find a path from one location to another in its warehouse. Pilliga Postage has decided to use a breadth-first search to find a path from one location to another.


Will this algorithm return a valid path and, if so, will it be a shortest path? Justify your answers.

2 marks

Pilliga Postage would also like to program the robots with a safety feature that will allow the robot to correct unsafe piles of boxes. A pile of boxes is deemed unsafe if there is a larger box above a smaller box anywhere in the pile. In the example shown below, the number on each box indicates the size of the box.



The robots will be able to correct unsafe piles by performing a series of box flips until the pile is safe. A flip involves the robot picking up any number of boxes at the top of the pile, flipping them and placing them back on the pile upside down.



- c. Explain why the stack ADT is not an appropriate data type to model a pile of boxes. State and justify an alternative ADT that would be appropriate.

3 marks

DO NOT WRITE IN THIS AREA

- d. Write pseudocode for a recursive `check_safe` algorithm that the robots will use to determine whether a pile is safe using the ADT stated in **part c**. The algorithm should take the current pile of boxes as input and output True if the pile is safe and False otherwise. 3 marks

- e. An optimal sequence of box flips is one that makes the pile safe with the minimum number of flips.
Describe a combination of ADTs that could be used to model the problem of finding an optimal sequence of box flips and describe how a single-source shortest path algorithm could be applied to this model to find an optimal sequence of box flips. Do **not** describe a specific single-source shortest path algorithm. 4 marks

Question 5 (4 marks)

Consider the inner loop of the Bellman-Ford algorithm.

```

For all  $v$  in  $V$  Do
    distance[ $i, v$ ]  $\leftarrow$  distance[ $i-1, v$ ]
For each edge  $(u, v)$  in  $E$  Do
    distance[ $i, v$ ]  $\leftarrow$  min(distance[ $i, v$ ], distance[ $i-1, u$ ] + weight( $u, v$ ))
  
```

In this algorithm, $\text{distance}[i, v]$ stores the shortest known distance from s to v in the i th iteration of the outer loop of the algorithm.

- a. Write the missing word in the statement below that describes what $\text{distance}[i, v]$ represents.

1 mark

$\text{distance}[i, v]$ represents the shortest path from s to v with at most i _____.

- b. Given that $\text{distance}[i-1, v]$ is correct, demonstrate that the distance calculated for $\text{distance}[i, v]$ is correct.

3 marks

Question 6 (4 marks)

Eddie has an idea for a business based on editing photographs of people's eyes. Eddie has a number of standard algorithms that he can use on photographs to produce artistic effects, such as highlighting, shading, grey-scaling and so on. Eddie is also interested in allowing users to supply their own code to produce their own individual effects if they wish. Eddie will use the following process:

- Step 1 – The user sends their code and a photograph of their eyes to a dedicated email address.
- Step 2 – The code is tested to ensure it will terminate on the given input. If not, the code is deleted and the user is notified.
- Step 3 – Only a code that has passed the test in Step 2, and is thus guaranteed to terminate on the given input, will be run on Eddie's servers to produce the desired effect.

a. Explain why Step 2 in Eddie's process will not work.

2 marks

b. Explain how Eddie can improve his process while still allowing users to submit their own code.

2 marks

DO NOT WRITE IN THIS AREA

Question 7 (4 marks)

Consider the following algorithm

```

Algorithm Sort(List)
  If List is empty
    Return []
  Else
    SortedList = []
    While List is not empty
      Item = head(List)
      List = tail(List)
      NewSortedList = insert(Item, SortedList)
      SortedList = NewSortedList
    Return SortedList

```

where `insert(Item, List)` does an item-by-item search of the list `List` until it finds the right place and inserts `Item` appropriately. For example, `insert(4, [1, 2, 3, 5])` will compare 4 to every item in the list and return the result `[1, 2, 3, 4, 5]`.

State the best case and worst case time complexity of this algorithm, providing an example input for each.

Question 8 (10 marks)

The decision version of the travelling salesman problem asks: given a weighted, undirected graph and a source node, does there exist a path starting and finishing at the source node that visits all remaining nodes exactly once and has a total weight of less than X (where X is a given integer)?

- a. Explain why the decision version of the travelling salesman problem is in NP. 2 marks

- b. Describe a greedy approach to solving the travelling salesman problem. 2 marks

DO NOT WRITE IN THIS AREA

- c. An alternative approach to solving the travelling salesman problem is to use a simulated annealing algorithm. The greedy approach described in **part b.** can be used to generate an initial candidate solution for the simulated annealing algorithm. High-level pseudocode for the simulated annealing algorithm is provided below.

```

Generate initial candidate solution, S
Repeat until a terminating condition is reached:
    Generate a new candidate solution, S_new
    Set a temperature, T
    If the new candidate solution is accepted:
        Set S = S_new
Return S

```

- i. Explain how the algorithm would determine whether a new candidate solution is accepted.

3 marks

- ii. State a terminating condition that could be used for this algorithm.

1 mark

- iii. Describe **one** advantage and **one** limitation of using a simulated annealing algorithm to solve the travelling salesman problem.

2 marks

Question 9 (10 marks)

A Melbourne property developer would like to buy a stretch of land to build some townhouses. She has identified a street on which she would like to build and has spoken to each of the property owners on the street to determine how much profit she would make by building on their land.

She has stored this information in an array, with each of the elements in the array representing the profit (or loss) for the purchase, in thousands of dollars.

For example, the array $P = [80, -120, 50]$ would represent three houses, where buying the first would give the developer \$80 000 in profit, buying the second would result in a \$120 000 loss and buying the third would give a \$50 000 profit.

The developer would like to solve the problem of determining the greatest profit she could make by buying a single, continuous stretch of land.

Three examples of continuous stretches of land with the greatest profit are provided below.

$P = [80, -120, 50]$: greatest profit is 80 thousand. ([80, -120, 50])

$P = [20, -40, 90, 20, -50, 70]$: greatest profit is 130 thousand. ([20, -40, 90, 20, -50, 70])

$P = [25, -10, -20, 50, 40, -30]$: greatest profit is 90 thousand. ([25, -10, -20, 50, 40, -30])

- a. Describe a brute-force approach for solving this problem. Do **not** provide pseudocode in your answer.

3 marks

- b. Discuss whether a divide and conquer approach could be used to solve this problem. Justify your answer. 3 marks

- c. Write pseudocode for a dynamic programming algorithm that would solve this problem. 4 marks

DO NOT WRITE IN THIS AREA

Question 10 (4 marks)

Zion would like to create a neural-network-based algorithm that would automatically tag any photographs of him on social media. To do so, the algorithm will need to be able to classify whether a photograph shows ‘Zion’ or ‘not Zion’ and tag only the former.

- a. Describe a motivation for using a neural network rather than a traditional algorithm for classifying photographs.

2 marks

- b. Zion trains his neural network using a set of 100 photographs that includes both photographs of him and photographs of other people. After training his model, he tests its accuracy by renaming the photographs as ‘Photo 1’, ‘Photo 2’ and so on, and running the model on the de-identified photographs.

Describe **two** problems with Zion’s current approach to training and testing a neural network. 2 marks

Question 11 (10 marks)

Consider the following discussion between two Algorithmics students, Raya and Manu, covering the topics of computability and John Searle's Chinese Room Argument.

Raya: 'I find it difficult to explain what an undecidable problem is.'

Manu: 'It's easy! An undecidable problem is one where there is no instance of the problem that can be solved.'

- a. Discuss why Manu's explanation is incorrect and provide an improved explanation of what an undecidable problem is.

3 marks

Raya: 'And then of course there are intractable problems.'

Manu: 'Intractable problems? What are they?'

Raya: 'They are problems for which there is no way of checking a solution quickly and so it is impossible to know if they are decidable or not.'

- b. Explain why Raya's definition is incorrect and provide an improved definition of intractable problems.

3 marks

Raya: ‘Do you remember much about John Searle’s Chinese Room Argument?’

Manu: ‘I just remember that the motivation for John Searle’s Chinese Room Argument was the idea of strong AI.’

- c. Explain what ‘strong AI’ means in the context of the debates about artificial intelligence. 2 marks

Raya: ‘There seem to be many people who think Searle is wrong! For example, those who put forward the Systems Reply argue that Searle is wrong to only consider the brain of the person in the room. Instead, he should be considering the whole room as a “system” that does understand Chinese.’

- d. Describe a standard counterargument to the Systems Reply. 2 marks

DO NOT WRITE IN THIS AREA