

STUDENT NUMBER Letter

ALGORITHMICS (HESS)

Written examination

Monday 6 November 2017

Reading time: 11.45 am to 12.00 noon (15 minutes)

Writing time: 12.00 noon to 2.00 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	16	16	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 26 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.

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

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

For extremely large values of n , which one of the following statements in relation to the time complexity of an algorithm is correct?

- A. The best case running time of an algorithm that runs in $O(n)$ time can be faster than the worst case running time of an algorithm that runs in $O(\log n)$ time.
- B. The worst case running time of an algorithm that runs in $O(n)$ time can be faster than the best case running time of an algorithm that runs in $O(\log n)$ time.
- C. An algorithm that runs in $O(n)$ time is always slower than an algorithm that runs in $O(\log n)$ time.
- D. An algorithm that runs in $\Omega(n)$ time is always slower than an algorithm that runs in $O(n^2)$ time.

Question 2

One difference between the Bellman-Ford and Dijkstra's shortest path algorithms is that

- A. Dijkstra's algorithm revisits edges of nodes.
- B. the Bellman-Ford algorithm revisits edges of nodes.
- C. Dijkstra's algorithm will not underestimate shortest path lengths, whereas the Bellman-Ford algorithm will.
- D. the Bellman-Ford algorithm fails when computed on graphs with negative cycles, whereas Dijkstra's algorithm does not.

Question 3

Consider the following pseudocode for the algorithm $f(a, b)$.

```

Algorithm f (a, b)
  Inputs: a, b; both integers  $\geq 0$ 
  Output: an integer  $\geq 0$ 

  if (a < b)
    return a
  else
    return f(a - b, a)
  endif

```

What does the algorithm f return when $a = 2$ and $b = 3$?

- A. 0
- B. 1
- C. 2
- D. 3

Use the following information to answer Questions 4 and 5.

A large, straight shipping canal uses a special system to allow ships to pass through safely. For the system to work, only one ship can pass through the canal at a time, regardless of which direction that ship comes from. If a ship is currently passing through the canal and another ship arrives, the second ship must wait until the first ship has reached the other side before it can proceed. If there are ships waiting on either side of the canal at the same time, the ship that arrived earliest is the next to proceed.

Question 4

Which abstract data type (ADT) best models the behaviour of ships that want to pass through the canal?

- A. stack
- B. queue
- C. dictionary
- D. associative array

Question 5


The method currently used to manage ships when they arrive at the canal does not account for the case when two ships arrive, one on each side, at exactly the same time. This could be solved by implementing a heuristic to select which ship will go first.

To minimise the time it takes for ships to pass through the canal, the most relevant data each ship could provide as input to a heuristic is its

- A. width.
- B. speed.
- C. length.
- D. weight.

Question 6

Consider the following graph representation of pseudocode.



Starting at C, in which order would each of the nodes be first examined if they were traversed using depth-first search? (Alphabetical order is used when there is more than one option.)

- A. A, AND, B, OR, C
- B. A, AND, OR, B, C
- C. C, OR, B, AND, A
- D. C, B, A, AND, OR

Question 7

A special type of Turing machine, M , terminates correctly with output O when run with a finite input tape, T . Which one of the following changes can be made to T to guarantee that when M is run with T , it still terminates with the identical answer O ?

- A. shortening the tape
- B. lengthening the tape
- C. adding new symbols to non-blank cells on the tape
- D. adding existing symbols to non-blank cells on the tape

Question 8

Given finite time, a problem that is classed as undecidable is known to have

- A. an algorithm that returns a valid solution.
- B. an algorithm that checks if a solution is valid.
- C. no single algorithm that will return a true or false solution.
- D. more than one algorithm that will return a true or false solution.

Question 9

Church and Turing defined the computability of a function by a human and made the assumption that

- A. a finite amount of space and time was available to the human.
- B. the algorithm can recover from any error a human would make.
- C. the human performing the calculation would make no mistakes.
- D. the human would only need infinite space to compute the function, but not infinite time.

Question 10

A professor is working on a problem and has thought of an algorithm to verify her solution in $O(n)$ time.

Without knowing anything further about the professor's algorithm, which one of the following statements is definitely true?

- A. The problem can be solved by an algorithm in quadratic time.
- B. The problem can be solved by an algorithm in linear time.
- C. The problem is in class NP.
- D. The problem is in class P.

Question 11

Claire is thinking of a way to solve a problem that is demonstrating combinatorial characteristics. She is choosing between a brute-force or a greedy approach.

To find a solution with large input in reasonable time, which approach should Claire choose and for which reason?

- A. brute-force, because it will always find an optimal solution
- B. greedy, because it will always find an optimal solution
- C. brute-force, because it may find an optimal solution
- D. greedy, because it may find an optimal solution

Question 12

A graph algorithm will need to work in conditions where the edges and nodes already traversed cannot be checked. Cycles can exist, all edges have a weight of 0 and no pre-computation can be undertaken on the graph.

Which one of the following will guarantee that a solution will be found?

- A. best-first search
- B. depth-first search
- C. breadth-first search
- D. meta-heuristic search

Question 13

A farmer is breeding rabbits that reproduce at a monthly rate that follows a Fibonacci series.

Assume that all rabbits that are born will live.

If the farmer wants to know how many rabbits she will have in two years' time, the solving strategy that gives the smallest guaranteed Big-O time complexity is

- A. simulated annealing.
- B. dynamic programming.
- C. tree search by backtracking.
- D. recursive divide-and-conquer.

Question 14

Consider the following functions:

$$f(n) = 4n^2 + 6$$

$$g(n) = 4n^2 + \log n$$

$$h(n) = f(n) + g(n)$$

$$j(n) = \frac{f(n)}{g(n)}$$

What is the best Big-O time complexity description of the asymptotic growth rate of $h(n)$ and $j(n)$, respectively?

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

Question 15

A student makes the following statements about the Floyd-Warshall algorithm on a graph with n nodes:

- I Its running time is in the P complexity class.
- II In the best case, the running time is $\Theta(n^2)$.
- III In the worst case, the running time is $O(n^3)$.
- IV In the average case, the running time is $O(n \log n)$.
- V It must run in $\Omega(n)$ time because every node must be inspected at least once.

Which group of statements is correct?

- A. II and IV
- B. III and IV
- C. I, II and V
- D. I, III and V

Question 16

What values of a , b and c in the recurrence $T(n) = \begin{cases} aT\left(\frac{n}{b}\right) + kn^c & \text{if } n > 1 \\ 42 & \text{if } n = 1 \end{cases}$

would give the running time $T(n)$ as $O(n^2 \log n)$?

- A. $a = 9$, $b = 3$, $c = 2$
- B. $a = 9$, $b = 2$, $c = 3$
- C. $a = 27$, $b = 3$, $c = 2$
- D. $a = 27$, $b = 2$, $c = 3$

Question 17

The owner of a business has taken on a number of contracts that require her to hire new employees. These employees will work in small teams, with a manager who will provide daily reports to the owner. As the project evolves, certain tasks are added that have varying levels of importance. The owner has decided to expand her employee system to reflect the inclusion of the new teams.

Which combination of ADTs would be the most appropriate to represent the reporting structure and the tasks within each team, respectively?

- A. disconnected graph, priority queue
- B. directed graph, priority queue
- C. weighted graph, stack
- D. decision tree, stack

Question 18

The time complexity of the best known algorithm to solve four different computational problems is given below:

Problem 1: $O(2^n)$

Problem 2: $O(n^2)$

Problem 3: $O(\log n)$

Problem 4: $O(n \log n)$

According to Cobham's thesis, which problem is not known to be feasible?

- A. Problem 1
- B. Problem 2
- C. Problem 3
- D. Problem 4

Question 19

Which one of the following did Searle **not** explicitly state regarding his original Chinese Room Argument?

- A. The Chinese characters can be interpreted and replicated.
- B. The entity inside the room needs to understand the language to produce an output.
- C. The language used by the entity is fixed and every statement is known in advance and in its entirety.
- D. The entity inside the room contains instructions that, if given an input, can produce an appropriate output.

Question 20

One benefit of using DNA computing as an alternative method of computation is that it allows for

- A. large amounts of parallel processing.
- B. the self-correction of errors during computation.
- C. results that can be easily interpreted when compared to that of a binary machine.
- D. large sequential computing that can be computed quickly, typically in the order of milliseconds.

THIS PAGE IS BLANK

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 (3 marks)

A local artist is planning to create a very large mural made out of small pieces of coloured glass that are joined together. Her requirements are that each piece of glass cannot touch the side of any other piece of glass if they are the same colour.

There are nine colours of glass and the glass pieces are all six-sided polygons.

Outline an algorithm the artist could use to create a mural that meets her criteria in a reasonable amount of time. As part of your description, identify an appropriate abstract data type (ADT) that could be used, referring to its specifications.

Question 2 (6 marks)

Simon is documenting a newly discovered language, where words vary in the number of letters and vowels ('a', 'e', 'i', 'o', 'u'). This language marks each vowel in one of three ways. For example, 'a' can be marked as 'ā', 'ă' and 'â'.

Simon would like to produce an algorithm, `generateWords`, that he can use to discover potential new words. When given an input word, `word`, without markings shown on any vowels, `generateWords(word)` would produce as output all of the potential words that might exist with the vowel markings included.

Example with one vowel

Input: cat

Output: cāt, căt, cât

Example with two vowels

Input: cate

Output: cātē, cātĕ, cātē, cătĕ, cătĕ, cătĕ, cătĕ, cătĕ, cătĕ

Simon will use this list to ask speakers of the language if any of the generated words are real words in their language.

- a. Describe an algorithm design pattern that Simon could use in the `generateWords` algorithm to solve his problem.

3 marks


- b. Simon would also like an algorithm that takes as input a set of words without vowel markings, returning as output all of the potential words based on vowel combinations, as well as all of the potential sentences based on word order, where each word can appear in any position in the sentence.

Explain how modifying the algorithm to generate all possible sentences will affect the solvability of Simon's problem.

3 marks

Question 3 (6 marks)

A locksmith has forgotten the combination to a lock that he owns. The combination lock has three dials, where each dial has 10 possible numbers. All three numbers must be correct for the lock to open. An example is shown below.



Write an algorithm that will find the correct combination. Assume that the following functions exist:

- `set(x, y)` – sets the value of dial x to y
 - x is the dial from left to right, where $0 \leq x \leq 2$
 - y is the number, where $0 \leq y \leq 9$
- `unlock()` – will return `True` if the combination is correct
- `reset()` – will reset all dials to zero

Question 4 (4 marks)

A software engineer has designed a program that uses a number of different algorithms with varying complexity to solve a larger problem.

Discuss how variation in complexity can affect the overall running time of the program if the algorithms are in sequence compared to if the algorithms are nested, with one calling another. In your response, use examples of time complexity.

Question 5 (3 marks)

Describe the Halting Problem as defined by the Church-Turing Thesis. In your response, refer to the decidability of a problem and how the Halting Problem affects the ability of a program to be verified.

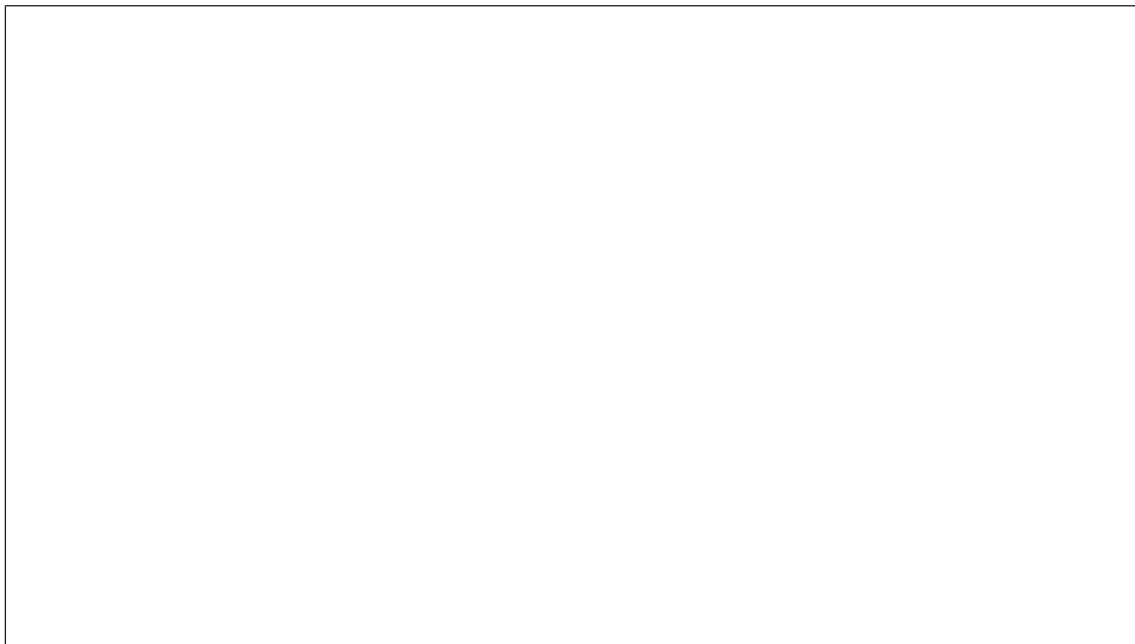
Question 7 (8 marks)

Emerging technologies may help overcome some of the limits of current computing systems.

- a. Describe **one** limit that exists in current models of computing. 2 marks

- b. Describe **one** way a neural network could be used to overcome current limits of computation. 3 marks

- c. Draw a diagram of a basic neural network represented as a graph, making sure to indicate the purpose of each edge and node. 3 marks



Question 8 (10 marks)

Harry owns a childcare centre. He urgently needs to know the median age of all children at the centre in order to qualify for new funding.

The children are split into two rooms, each with one childcare specialist. The first room contains n children and the second room contains $n + 1$ children. The specialist in each room knows the order, by age, of all of the children in that room and can instantly answer the question, ‘What is the age of the k^{th} child in your room?’

- a. Write a brute-force algorithm that would return the median age of all of the children in the childcare centre and would require Harry to ask the specialists a total of $2n + 1$ times for the age of some child in their room.

5 marks

- b. Harry has been given an unfinished draft of a different algorithm that will also give him the median age of all of the children when called as $\text{Search}(0, n - 1, 0, n)$. Before Harry completes the algorithm, he wants an estimation of the total number of questions he would ask the specialists.

```

Algorithm Search(L1, R1, L2, R2)
    Input L1: lowest rank of child in room 1 being considered
    Input R1: highest rank of child in room 1 being considered
    Input L2: lowest rank of child in room 2 being considered
    Input R2: highest rank of child in room 2 being considered

    Let n1 = R1 - L1 + 1
    Let n2 = R2 - L2 + 1

    Ask specialist in Room 1 for age of child L1 + n1/2 and call that m1.
    Ask specialist in Room 2 for age of child L2 + n2/2 and call that m2.

    if n1 equals 0
        return m2
    if n2 equals 0
        return m1
    if m1 > m2
        return Search(L1, L1 + n1/2, L2 + n2/2, R2)
    else
        return Search(L1 + n1/2, R1, L2, L2 + n2/2)
    endif

```

Write a recurrence relation that reflects the running time of the algorithm as it is written. 2 marks

- c. Assume that the completed `Search` algorithm has the same time complexity as the draft. If Harry's aim is to ask as few questions as possible of his specialists, would you recommend a brute-force algorithm or the `Search` algorithm? Justify your recommendation. 3 marks

Question 9 (5 marks)

Consider a connected, weighted, undirected graph, G , whose smallest edge weight is w . There is only one edge (u, v) in G that has weight w .

Outline a proof by contradiction that shows that (u, v) must be in any minimal spanning tree of G .

Question 10 (4 marks)

Joseph is having a birthday party for his little sister and he wants to make sure she wins the prize in a game that the children will play. In the game, all the children at the party, N , will sit in a circle and pass the prize to their left M times. After M passes, the child holding the prize passes it to the child to the left if $N > 1$, then leaves the circle, and the prize is passed M times again. Joseph wants to choose M so that his sister is the last child left in the circle and wins the prize.

Joseph does not know yet how many children will be coming to the party.

He decides to write an algorithm that will let him input M , N , P and S , where S is the starting location of his sister and P is the starting location of the prize ($0 \leq P < N$). The algorithm will return true if his sister will win and false if any other child wins.

Algorithm GetWinner (M, N, P, S)

Input: M , integer, number of prize passes to the left

Input: N , integer, number of children at the start of the game

Input: P , integer, starting location of the prize

Input: S , integer, starting location of the sister

Output: true if S is the only player left, false otherwise

Let children = $[0, 1, \dots, N-1]$

Let current = P

repeat $N - 1$ times

 Let current = remainder of $(\text{current} + M) / (\text{number of children})$

 Remove current from children

end repeat

if children contains only S

 return true

else

 return false

endif

- a. Joseph writes a brute-force search algorithm that calls algorithm GetWinner on all possible combinations of M , N , P and S to find winning values for his sister.

Assuming that the algorithm GetWinner could be implemented in $O(N)$ time, what is the best case running time of Joseph's brute-force search algorithm? Why?

2 marks

- b. Assuming that the algorithm GetWinner could be implemented in $O(N)$ time, is the worst case running time of Joseph's brute-force search algorithm greater than $O(N^2)$? Why or why not?

2 marks

Question 11 (6 marks)

A planetary space station is populated by workers and their families, where each family lives in its own housing dome. Domes are connected to one another via an extensive tunnel system that can run for multiple kilometres, and contains many entry and exit points to other tunnels.

An outbreak of a virus is rapidly spreading through the space station. The virus is spread between people if they live in the same dome. At least one person from each dome needs to be vaccinated against the virus to halt its spread.

Scientists have developed a vaccine that works when sprayed from the ceiling of a housing dome. They have a large map of the area with the location of all domes clearly marked, sorted by ascending distance from their laboratory. They propose to use the following algorithm to carry out the vaccinations.

Algorithm Vaccinate(L)

Input: L, a list of domes sorted by distance from the laboratory.

while L is empty

 Deliver vaccinations in the first dome in L.

 Remove the last element of L.

- a. The proposed algorithm has errors.

How many housing domes will be vaccinated using this algorithm as it is written?

1 mark

- b. Write a corrected version of the algorithm that will allow the vaccination to be delivered to each dome in order of the dome's distance from the laboratory.

2 marks

- c. After the scientists have vaccinated the people in some domes using a correct algorithm, they discover that vaccinated people can spread the vaccine to others without the spray being needed. This means that any vaccinated person can vaccinate everyone in their nearest neighbouring domes by simply touching one member of each neighbouring dome.

What ADT would be most appropriate to model housing domes and neighbouring domes? Assuming that a dome can vaccinate exactly all of its neighbouring domes in one day, what type of algorithm could be used with your selected ADT to compute the maximum number of days after scientists stop spraying until everyone in the space station is vaccinated? Briefly explain how this type of algorithm would work.

3 marks

Question 12 (5 marks)

BattleForFlags is a computer game where pairs of players compete against one another in virtual combat. Each player in a pair has a particular set of skills, with the most effective pairs containing players who have complementary skills. Skill levels are represented by whole integers, where larger integers indicate higher ability in that skill. A complementary rating of a skill can therefore be determined by summing the absolute differences between the skills of two players. The complementary rating (CR) can be represented as

$$CR(A, B) = \sum_{k=0}^{n-1} |A[k] - B[k]|, n > 0$$

where n is the total number of skills, A and B are arrays of player skill values and k is the index value in those arrays. Assume all players have the same sets of skills.

For example, Caro is trying to decide if she should team up with Max or Nik. These are their skills.

	k	Caro	Max	Nik
Speed	0	10	6	10
Accuracy	1	4	8	9
Endurance	2	5	10	4

Caro has a CR of 13 between herself and Max, and a CR of 6 with Nik. She should, therefore, pair up with Max.

BattleForFlags is a very popular computer game, with hundreds of thousands of players. Players can add each other as friends, allowing them to keep track of people they enjoy playing with. It would be useful for the game to have an algorithm that suggests to a player the most complementary player that is available to play with them in a game, based on their friends, the friends of their friends, and so on. A player is unavailable if they are already competing in a game. If the selected player is unavailable, the next best player who is not in a game is selected.

- a. A player who has just installed the game will have no friends.

Suggest an additional condition to include in the algorithm that will allow new players without friends to receive complementary player suggestions. 2 marks

- b. Explain how a combination of tree search and any other algorithm design patterns could be used in the algorithm to return the best available complementary player in reasonable time. As part of your explanation, refer to any assumptions that must be made about how player data is represented.

3 marks

Question 13 (3 marks)

Two friends, June and David, are having a conversation about the power of algorithmic thinking in real-world scenarios.

June challenges David with the following: ‘Taking into account that there are approximately 50 000 cities in the world and that they can be sorted in alphabetical order, choose any city in the world and don’t tell me its name. Allow me to ask you “yes” or “no” questions about its name in relation to other cities. I’m sure that I can find out which city you’ve chosen in fewer than 16 questions!’

State and describe an algorithm design pattern that June can use to guarantee that she will select David’s chosen city using fewer than 16 questions.

Question 14 (4 marks)

A post office is testing a new self-driving delivery van. The van starts its day in a depot, where it receives a list of requests to transport packages from the depot to their final destinations. Assume the self-driving van is of a sufficiently large size to fit all of the packages that need to be delivered on any given day, that each location will only have one package delivered to it and that the van will always have enough time to deliver all of its packages.

- a. Describe an algorithm that will quickly find a route for the self-driving van on a single day. 2 marks

- b. Describe an algorithm that will find the shortest route for the self-driving van on a single day. 2 marks

Question 15 (4 marks)

Hard drives from a large storage device have been installed incorrectly, as shown in Figure 1. The data cannot be recovered in its entirety unless the drives are reordered as shown in Figure 2. To further complicate the issue, the system itself can only deactivate one drive out of eight, temporarily, to be moved to an adjacent, vacant drive bay. Bays above, below, to the left or to the right are configured to receive swapped drives.

HDD1	HDD2	HDD5
HDD4	HDD3	
HDD6	HDD7	HDD8

Figure 1

HDD1	HDD2	HDD3
HDD4	HDD5	HDD6
HDD7	HDD8	

Figure 2

Before a hard drive can be moved to the vacant drive bay, the manager wants to confirm that he can successfully move the drives in Figure 1 to match the order in Figure 2.

Explain a method the manager could use to check if the hard drives in Figure 1 can be reordered to match Figure 2.

Question 16 (3 marks)

In a simulation, a shark is hunting a fish to eat for dinner. The shark and the fish are both contained in a square 5×5 grid made of cells. Both the shark and the fish can move up, down, left and right to adjacent cells in the grid. They cannot move diagonally.

Assume that the shark starts in the cell in the bottom-left corner and the fish starts in the cell in the upper-right corner of the grid. Assume also that each time the shark moves into a cell adjacent to its current position in the grid, the fish will randomly move into a cell adjacent to its current position in the grid.

Explain how the minimax algorithm can be used by the shark to maximise the likelihood of it catching the fish.
