

STUDENT NUMBER Letter

ALGORITHMICS (HESS)

Written examination

Monday 31 October 2016

Reading time: 3.00 pm to 3.15 pm (15 minutes)

Writing time: 3.15 pm to 5.15 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	17	17	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 30 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.

Question 1

Consider the following pseudocode, where a and n are positive integers.

```

b = -1
c = a
i = 0
while(i is less than or equal to n)
    c = c * b
    i = i + 1
endwhile

```

Which one of the following statements correctly describes the value of c after the algorithm is executed?

- A. c is undefined
- B. c is equal to -1
- C. c is equal to $-a$ if n is odd
- D. c is equal to $-a$ if n is even

Question 2

Which one of the following statements about Floyd-Warshall's algorithm running on a graph with V nodes and E edges is correct?

- A. The recursive version finds the transitive closure of a graph in $O(3^V)$ time.
- B. The iterative (dynamic programming) version finds the shortest path between all pairs of nodes in $O(V^3)$ time.
- C. The iterative (dynamic programming) version finds the shortest path between all pairs of nodes in $O(3^E)$ time.
- D. The iterative (dynamic programming) version always finds a minimal spanning tree rooted at every node in $O(V^3)$ time.

Question 3

Consider the following three parameters of a function:

- *row*, which can take the values 1 or 2
- *column*, which can take the values 'a' or 'b' or 'c'
- *sound*, which can take the values True or False

Which one of the following statements is **true**?

- A. Pair-wise testing of the function requires six test cases.
- B. Pair-wise testing of the function requires 12 test cases.
- C. Exhaustive testing of the function requires six test cases.
- D. Exhaustive black-box testing of the function requires $7!$ test cases.

Question 4

Consider the following function $f(x, y)$ that takes two integers as input.

```

f(x, y)
begin
  if (x < 0 or y < 0)
    if (x > y)
      a = y
    else
      a = x
    endif
  else
    if (x > y)
      a = x
    else
      a = y
    endif
  endif

  return a
end

```

Which one of the following sets of test cases covers all paths of the function?

- A. $\{(x=-3, y=4), (x=-2, y=6), (x=3, y=6), (x=3, y=1)\}$
- B. $\{(x=-3, y=-4), (x=-2, y=6), (x=3, y=6), (x=3, y=1)\}$
- C. $\{(x=-3, y=-4), (x=-2, y=6), (x=3, y=6), (x=3, y=3)\}$
- D. $\{(x=-3, y=-4), (x=-2, y=6), (x=3, y=-6), (x=3, y=1)\}$

Question 5

When considering different types of algorithms, which one of the following statements is **true**?

- A. Divide and conquer algorithms are always faster than greedy algorithms for the same problem.
- B. Greedy algorithms are always faster than divide and conquer algorithms for the same problem.
- C. Greedy algorithms give good approximate answers to problems, but never the best possible answer.
- D. Brute-force algorithms can never be faster than a well-designed greedy algorithm for the same problem.


Question 6

A large number of people will be completing an online survey. Each survey response will be stored in the order it is submitted.

Which abstract data type (ADT) would be the most appropriate for storing the survey responses?

- A. graph
- B. queue
- C. stack
- D. dictionary

Use the following information to answer Questions 7 and 8.



Question 7

A depth-first search traversal of the graph above will visit the nodes in which one of the following orders? (Alphabetical order is used when there is more than one option.)

- A. ABDGEFCH
- B. ABDEFCHG
- C. ABEGFCHD
- D. ABDFCHEG

Question 8

A breadth-first search traversal of the graph above will visit the nodes in which one of the following orders? (Alphabetical order is used when there is more than one option.)

- A. ABEGFCHD
- B. ABDGFECH
- C. ABDGEFCH
- D. AGBDEFCH

Question 9

A connected, undirected graph with distinct edge weights has maximum edge weight e_{\max} and minimum edge weight e_{\min} .

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

- A. e_{\max} is not in any minimal spanning tree.
- B. Every minimal spanning tree of the graph must contain e_{\min} .
- C. Prim's algorithm will generate a unique minimal spanning tree.
- D. If e_{\max} is in a minimal spanning tree, its removal will disconnect the graph.

Question 10

Dijkstra's single-source shortest path algorithm in an undirected graph reports distances from the source to each node.

These distances

- A. are the shortest possible distances to every destination node.
- B. are never the shortest possible distances when negative edge weights are present.
- C. may be the shortest possible distances when negative edge weights are present.
- D. may not always be the shortest possible distances when all edge weights are positive.

Question 11

Consider the following four algorithms, operating on a graph with V nodes and E edges:

1. Floyd-Warshall's algorithm for transitive closure
2. Bellman-Ford's algorithm for the single-source shortest path problem
3. depth-first traversal algorithm
4. Dijkstra's algorithm for the single-source shortest path problem

The time complexities of these algorithms, in order, are

- A. $O(V^3)$, $O(VE)$, $O(V^2)$, $O(V+E)$
- B. $O(V+E)$, $O(V^3)$, $O(VE)$, $O(V^2)$
- C. $O(V^3)$, $O(VE)$, $O(V+E)$, $O(V^2)$
- D. $O(VE)$, $O(V+E)$, $O(V^3)$, $O(V^2)$

Question 12

```

Algorithm A
Input: two nodes X, Y
Output: "yes" if successful, "no" otherwise

Let L = an empty list
Append X to L

while(L is not empty)
  Let Z = first node in L
  if(Z = Y) then
    return "yes"
  else
    Remove Z from L
    for every neighbour N connected to Z
      if(N.visited = False
        and N is not in L)
        Append N to the end of L
      endif
    endfor
    Z.visited = True
  endif
endwhile

return "no"

```

The algorithm above uses

- A. breadth-first search.
- B. depth-first search.
- C. mergesort.
- D. minimax.

Question 13

After k iterations of Bellman-Ford's algorithm in a graph with n nodes ($1 < k < n - 1$), which statement is **true** about the shortest paths from the source node to every other node?

- A. The shortest paths with at most k edges have been found.
- B. The shortest paths with at most $k - 1$ edges have been found.
- C. The shortest paths with at most $k + 1$ edges have been found.
- D. The shortest paths for any number of edges can change as the algorithm has not completed.

Question 14

The PageRank algorithm uses two attributes to calculate the popularity of a web page.

Which one of the following statements best describes these attributes for calculating a web page's PageRank?

- A. the number of outbound and inbound links of a web page
- B. the probability of a random direct landing on a web page and the size of the web page
- C. the number of outbound links of a web page and the probability of a random direct landing on the web page
- D. the number of inbound links of a web page and the probability of a random direct landing on the web page

Use the following information to answer Questions 15 and 16.

Blackbox Inc. is trying to sell new timetabling software to a school.

The following data shows running times (in microseconds) for how the new timetabling software's algorithm compares with a standard baseline, where n is a measure of the size of the timetabling problem.

n	Baseline	New algorithm
1	5	3
10	105	53
100	1998	103
1000	29902	152
10000	398636	202

Question 15

Which Big-O expression most closely describes the running time of the baseline?

- A. $O(n)$
- B. $O(n^2)$
- C. $O(\log n)$
- D. $O(n \log n)$

Question 16

Which Big-O expression most closely describes the running time of the new algorithm?

- A. $O(n)$
- B. $O(n^2)$
- C. $O(\log n)$
- D. $O(n \log n)$


Question 17

In the context of Algorithmics, a decision problem is

- A. a problem with only two possible solutions.
- B. a problem with more than two possible solutions.
- C. a problem for which all known algorithms do not terminate.
- D. undecidable according to the Church-Turing thesis because it does not terminate on a Turing machine.

Use the following information to answer Questions 18 and 19.

A Turing machine is configured with the instructions represented in the state diagram below.



Each edge is labelled $i/j : k$, where:

- i is the input
- j is the output
- k is the direction the head moves (L = left, R = right) after the output.


The machine begins in state q_0 .

The machine is given the following tape. For this machine, the tape remains stationary while the head moves.

		0	1	0	0	1	0	1		
--	--	---	---	---	---	---	---	---	--	--

Question 18

The Turing machine is run with the tape. When the Turing machine halts, the appearance of the tape and the position of the head are as follows. The arrow shows the ending point of the head.



Which one of the following best represents the position of the head directly before the Turing machine started?

- A.

		0	1	0	0	1	0	1		
									↑	
- B.


		0	1	0	0	1	0	1		
								↑		
- C.

		0	1	0	0	1	0	1		
↑										
- D.

		0	1	0	0	1	0	1		
										↑

Question 19

The arrow below shows the starting point of the head before the machine is run again.



When the Turing machine halts, the number of steps that the head will have moved is

- A. 6
 B. 7
 C. 9
 D. 10

Question 20

A student begins to write the following signature for a queue ADT.

```
name         queue
import      element, boolean
operations  empty : → queue
              isEmpty : queue → boolean
              peek : queue → element
              enqueue : queue × element → queue
              dequeue :
```

Which one of the following is the correct representation of the dequeue operation?

- A. dequeue : queue → queue
- B. dequeue : queue → boolean
- C. dequeue : element → queue
- D. dequeue : queue → element

SECTION B**Instructions for Section B**

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

Question 1 (3 marks)

Explain how randomised heuristics can help overcome the soft limits of computation. Use an example as part of your explanation.

Question 2 (3 marks)

A laboratory is working on a 3D-printed structure that uses a new experimental material. The material is made up of six separate components and needs to be structured in such a way that no component connects to another component of the same type. If components of the same type do connect, the structure will collapse. The connections between the components can be treated as a graph.

Describe an approach that could be used in reasonable time to ensure that structures of any size would not collapse.

Question 3 (3 marks)

A company makes synthetic diamonds to sell to customers who use them in high-powered lasers. The company has received complaints from its customers about some of the diamonds being faulty, making the lasers unsafe to use.

On investigating, a technician notices that the faulty diamonds weigh a few micrograms less than the diamonds that are not faulty and that each batch of diamonds is producing exactly one faulty diamond. As the diamonds are produced in very large batches, it would not be feasible to simply weigh each diamond until the faulty one is found. The manager would like to put in place a method to find faulty diamonds.

State an appropriate algorithm design pattern that the manager could use to successfully find faulty diamonds and explain how that algorithm will solve the problem efficiently.

Question 4 (2 marks)

When data is transferred across computer networks, it is first broken up into packets. Computer network traffic is normally processed in the order that packets arrive at each device along the path between communicating devices.

Packets may be pieces of email, web content, voice or video. While some traffic, such as email or web content, can withstand delays in delivery, others, such as voice and video, cannot have delays; these packets cannot wait at each device for other traffic to be processed ahead of them.

Describe a standard abstract data type (ADT) that could be used to manage the packets arriving at a computer.


Question 5 (2 marks)

An undirected graph, G , is possibly disconnected.

How can a depth-first search algorithm be modified to check for disconnectedness?

Question 6 (6 marks)

The following graph represents links between web pages.



The PageRank of Page *A* is given by

$$PR(A) = \frac{(1-d)}{N} + d \left(\frac{PR(B)}{L(B)} + \frac{PR(C)}{L(C)} + \frac{PR(D)}{L(D)} \right)$$


where $PR(x)$ is the PageRank of Page x , N is the number of pages in this network and $L(x)$ is the number of outgoing links from Page x .

- a.** Explain the purpose of d in the PageRank. 2 marks

- b.** What does $\frac{(1-d)}{N}$ represent in the PageRank? 1 mark

- c.** What does $d \left(\frac{PR(B)}{L(B)} + \frac{PR(C)}{L(C)} + \frac{PR(D)}{L(D)} \right)$ represent in the PageRank? 1 mark


- d. A new page, E , is added to the graph as a node, shown below.



Explain how the PageRank would include node E if there are no outbound links from Page E . 2 marks

Question 7 (3 marks)

A student runs Bellman-Ford’s single-source shortest path algorithm on the following directed graph using node *A* as the source. After nine iterations, she notes the distance from *A* to each of the other nodes. She then runs a tenth iteration of the algorithm and notes the distance from *A* to each of the other nodes.



- a. Which nodes will show a change in distance from source node *A* between the ninth and tenth iterations?

1 mark


- b. Explain why some nodes have remained the same distance from source node *A* while others have a new distance.

2 marks


Question 8 (4 marks)

An artist wishes to construct a sculpture. The sculpture will be made with numerous L-shaped blocks consisting of four cubes stuck together. These blocks can only be placed on top of each other using their connectors, shown below.

View from top



View from side



The artist would like to join a collection of blocks together so that, upon looking down on the sculpture, every cell of an 8×8 square of cells is covered, where each block would cover four cells.

Describe two different approaches that could be used by the artist to find a possible arrangement of blocks.

1. _____

2. _____

Question 9 (4 marks)

Describe the decision version of the travelling salesman problem and explain why it is a Non-deterministic Polynomial-time (NP) problem.

Question 10 (4 marks)

The Master Theorem provides a general solution to recurrence relations of the form

$$T(n) = aT\left(\frac{n}{b}\right) + f(n), \text{ where } a > 1, b > 1$$

If $f(n) \in O(n^d)$, the above recurrence relation has the solution

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

Consider the following version of mergesort, where a list is divided into three sub-lists of approximately equal size. Each of the lists is sorted recursively into sub-lists and then the sorted sub-lists are merged into a single sorted list.

Assume that the sub-lists can be merged in linear time.

Write a recurrence relation for the time complexity of this version of mergesort. State the appropriate case of the Master Theorem for this version of mergesort, including the values of a , b and d , and use it to solve the time complexity.

Question 11 (5 marks)

Consider the following algorithm that might be used to solve a problem where solutions can be randomly generated.

```

soln = generate random solution
temperature = 1
min_temperature = 0.01
n_iterations = 100
while temperature > min_temperature
  for i = 1 to n_iterations
    soln_new = generate neighbouring solution of soln
    if cost(soln) >= cost(soln_new)
      soln = soln_new
    else
      prob = (random 0 to 100)/100
      if e(cost(soln)-cost(soln_new))/temperature > prob
        soln = soln_new
      endif
    endif
  endfor
  temperature = cooling_factor * temperature
endwhile

```

- a. State the range of valid values for `cooling_factor`, so that at least 200 random solutions are generated and the algorithm terminates. 2 marks

- b. Given that `soln_new` is generated in the neighbourhood of `soln`, why is it a good idea to sometimes replace `soln` with `soln_new` when `cost(soln) < cost(soln_new)`? 1 mark


- c. Give **one** example of a problem where a version of the algorithm above is likely to give an acceptable solution. Describe a possible `cost(soln)` for that problem. 2 marks

Question 12 (3 marks)

Consider a game for two players, Player A and Player B, which uses three piles of stones. Each player takes a turn to remove as many stones as they wish from one of the piles. The object of the game is to make an opponent pick up the last stone.

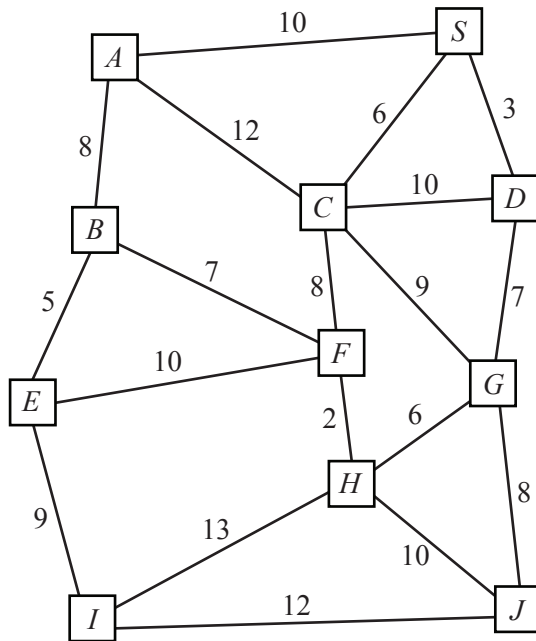
Suppose that during the game the three piles of stones consist of one stone, two stones and one stone, and it is Player A's turn.

Using the minimax algorithm, complete the following game tree until it determines the move Player A should make to guarantee that they win.



Question 13 (7 marks)

Below is a graph representation of a possible way in which a collection of computers can be connected. Each computer is labelled with a letter and is a node in the graph. Cables that are used to connect the computers are shown as edges and the length of each cable is given as an edge weight.




The collection of computers needs to be connected with cables such that the following conditions are met:

- Condition 1: There are no cycles.
- Condition 2: The shortest length of cabling is used from *S*, the source, to every other computer while the total cable length for the whole network is the smallest possible length.

a. Draw the graph produced by Prim's algorithm and indicate the condition(s) that the graph meets. 2 marks


Condition(s) met _____



- b. Draw the graph produced by Dijkstra's algorithm and indicate the condition(s) that the graph meets.

2 marks

Condition(s) met _____



- c. Is there a modification to Dijkstra's algorithm that will allow for both Condition 1 and Condition 2 to be met? Explain your answer.

3 marks

Question 14 (8 marks)

Consider the following table of distances between schools.

School 1	School 2	Distance
<i>A</i>	<i>B</i>	10 km
<i>A</i>	<i>C</i>	8 km
<i>B</i>	<i>C</i>	5 km
<i>C</i>	<i>E</i>	3 km
<i>B</i>	<i>D</i>	1 km
<i>D</i>	<i>E</i>	9 km

- a. Draw a graph that would represent this data and indicate how the schools and distances are represented in the graph drawn.

3 marks

How the schools are represented _____

How the distances are represented _____

- b.** Gerty wants to walk from School *A* to School *D*, beginning at School *A*, and will always choose to walk to the closest school.

What distance does Gerty cover if she never revisits a school?

1 mark

- c.** There is a shorter path for Gerty to get from School *A* to School *D*.

What is that path?

1 mark

- d.** Ned is interested in connecting all of these schools with fibre-optic cable for easy videoconferencing.

Assuming that each school requires a length of cable equal to the distance between the schools, what is the shortest total length of cable required to connect all of the schools?

1 mark

- e.** Sam has to walk between School *A* and some other school. For every school he passes, he stops and gives the principal at that school one apple. It costs Sam \$2 for each kilometre travelled and \$1 for each apple that he gives out.

How would the problem be modelled so that a single-source shortest path algorithm could be used to determine the lowest cost for Sam to travel to each school starting at School *A*?

2 marks

Question 15 (8 marks)

Wally the washer uses a specific approach to wash a basket of clothes. He begins by dividing the clothes into two smaller baskets and washes the first basket of clothes using his specific approach. Once he has washed all of the clothes in that basket, he then washes the second basket of clothes using the same approach.

- a. Write the pseudocode for an algorithm, `washClothes (basket)`, that describes Wally’s recursive washing system.

2 marks

When the washing has been completed, Wally has to match up each pair of washed socks.

- b. Write the pseudocode for an algorithm, `findMatch (sock, sockList)`, that will look through a list of socks, `sockList`, to return a matching pair for a given sock, `sock`. All socks are unique pairs.

3 marks

Question 16 (9 marks)

To revise her studies, Trudi plans to do six test papers in a row without sleep. Trudi is very particular about her study environment and insists that there be only three stacks of tests: incomplete tests, complete tests and marked tests. She may also have a single test on her desk.

Trudi begins with the first test on the stack of incomplete tests, does the test and then places it on top of the stack of complete tests. Once all of the tests are complete, she then marks them by taking a test from the top of the complete stack, marking it and placing it on the stack of marked tests.

- a. Assuming each stack of tests is modelled as a stack ADT with the usual operations, write the pseudocode for an algorithm that best represents Trudi's study session. 4 marks

- b. How many pop operations will be executed for Trudi's six tests? 1 mark

- c. In general, if there are n tests, how many pop operations will be executed? 1 mark
