

CSE1305 Algorithms & Data Structures

Resit Analysis Exam

05 April 2022, 09:00–11:00

Examiners:

Examiner responsible: Joana Gonçalves and Ivo van Kreveld

Examination reviewer: Stefan Hugtenburg

Parts of the examination and determination of the grade:

Exam part	Number of questions	Question specifics	Grade (%)	Grade (points)
Multiple-choice	20 questions (equal weights)	One correct answer per question	50%	$5 \cdot \frac{\text{score}}{20}$
Open questions	4 questions (different weights)	Multiple parts	50%	$5 \cdot \frac{\text{score}}{20}$

Use of information sources and aids:

- A hand-written double-sided A4 cheat sheet can be used during the exam.
- No other materials may be used, including but not limited to books, lecture slides in any form, or devices such as laptops and phones.
- Scrap paper sheets are provided at the beginning of the exam. Additional scrap paper can be requested.

General instructions:

- Solve the exam on your own. Any form of collaboration is prohibited.
- You cannot leave the examination room during the first 30 minutes.
- If you are eligible for extra time, place the “Verklaring Tentamentijd Verlenging” on your desk, together with your student card.

Instructions for writing down your answers:

- You should answer the questions using the provided answer sheets.
- Write your name and student number on every sheet of paper.
- **For multiple-choice questions**, mark the answers on this exam paper first, and copy them to the answer form after revising.
- **For open questions**, provide all requested information and always give an explanation. Avoid irrelevant data, it could lead to deductions.
- **For proofs**, make sure your proof is properly structured and sufficiently explained. Statements or steps without justification could lead to point deductions.

This page is intentionally left blank.

Multiple-choice questions (50%, 20 points)

- (1 point) Given two functions $f(n)$ and $g(n)$ where $f(n)$ is $\Omega(g(n))$. Which of the following statements about asymptotic algorithmic complexities must be **true**?
 - $g(n)$ is $\mathcal{O}(f(n))$.
 - $g(n)$ is $\Theta(f(n))$.
 - $g(n)$ is **not** $\Omega(f(n))$.
 - $f(n) > g(n)$ for some value of n .
- (1 point) The Java methods `sumA` and `sumB` both calculate the sum of the integer elements in array `arr`.

```
1 public static int sumA(int[] arr, int low, int high) {
2     if(low == high) {
3         return arr[low];
4     }
5     return sumA(arr, low, high - 1) + arr[high];
6 }
7
8 public static int sumB(int[] arr, int low, int high) {
9     if(low == high) {
10        return arr[low];
11    }
12    int mid = (high + low) / 2;
13    return sumB(arr, low, mid) + sumB(arr, mid + 1, high);
14 }
```

The problem size is the difference between `high` and `low`. What can we say about the tightest worst-case time and space complexity of methods `sumA` and `sumB`, expressed in big-Oh notation?

- Method `sumA` has larger time complexity and larger space complexity than method `sumB`.
 - Method `sumA` has larger time complexity than method `sumB`, space complexities are the same for both methods.
 - Method `sumA` has larger space complexity than `sumB`, time complexities are the same for both methods.
 - The methods `sumA` and `sumB` have the same time and space complexities.
- (1 point) Consider implementations of a stack and a queue data structures using a singly-linked list. Where do you need to add and remove elements to get a correctly working data structure with the optimal time complexities for both methods?

A. stack : add to head, remove from head	queue : add to head, remove from tail
B. stack : add to head, remove from head	queue : add to tail, remove from head
C. stack : add to tail, remove from tail	queue : add to head, remove from tail
D. stack : add to tail, remove from tail	queue : add to tail, remove from head
 - (1 point) Consider a **sorted** sequence of integers elements. The task is to insert a new element such that the sequence is kept sorted. What is the time complexity of the most time-efficient way to do this, using an array-based implementation and a doubly-linked list (DLL) implementation?

A. Array : $\mathcal{O}(\log(n))$	DLL : $\mathcal{O}(\log(n))$
B. Array : $\mathcal{O}(\log(n))$	DLL : $\mathcal{O}(n)$
C. Array : $\mathcal{O}(n)$	DLL : $\mathcal{O}(\log(n))$
D. Array : $\mathcal{O}(n)$	DLL : $\mathcal{O}(n)$


5. (1 point) Consider a dynamic array, where elements are inserted or removed always at the end. There can be arbitrary insertions and deletions in no particular order. When and how could we resize the array to make sure that the amortized time complexities of adding and removing one element are both $\mathcal{O}(1)$? Here n denotes the number of elements in the array, and C is the capacity of the array.
- A. When $n = C$, increase capacity to $n + 10$. When $n = C - 20$, decrease capacity to $n + 10$.
 - B. When $n = C$, increase capacity to $n + 10$. When $n = C - 10$, decrease capacity to n .
 - C. When $n = C$, increase capacity to $2 \cdot n$. When $n = \frac{1}{2} \cdot C$, decrease capacity to n .
 - D. When $n = C$, increase capacity to $2 \cdot n$. When $n = \frac{1}{2} \cdot C$, decrease capacity to $\frac{3}{2} \cdot n$.
6. (1 point) Consider the following Java implementation of an iterator `ADSLISTIterator` below.

```
1 public class ADSListIterator<V> implements Iterator<V> {
2     LinkedList<V> list;
3     int index;
4
5     public ADSListIterator(LinkedList<V> list) {
6         this.list = list;
7         index = 0;
8     }
9
10    public boolean hasNext() {
11        return index < list.size();
12    }
13
14    public V next() throws NoSuchElementException {
15        if (!hasNext()) throw new NoSuchElementException();
16        return list.get(index++);
17    }
18 }
```

Is this a lazy or snapshot iterator? And what is the time complexity of the `next` method?

- A. It is a lazy iterator. The time complexity of the `next` method is $\mathcal{O}(1)$.
 - B. It is a lazy iterator. The time complexity of the `next` method is $\mathcal{O}(n)$.
 - C. It is a snapshot iterator. The time complexity of the `next` method is $\mathcal{O}(1)$.
 - D. It is a snapshot iterator. The time complexity of the `next` method is $\mathcal{O}(n)$.
7. (1 point) Consider a tree storing integer elements, in which the element of a parent node is always larger than or equal to the elements of its children, and the element of the left child is also always larger than the element of the right child. Which of the following tree traversals visits the nodes of the tree in non-increasing order?
- A. Preorder traversal
 - B. Inorder traversal
 - C. Postorder traversal
 - D. Breadth-first traversal


8. (1 point) Consider the following heap:



Which nodes would have a different element after removing the minimum element from the heap?

- A. The nodes currently containing 3, 6, 7 and 11.
 - B. The nodes currently containing 3, 6, 7 and 9.
 - C. The nodes currently containing 3, 6 and 10.
 - D. The nodes currently containing 3, 6, 7 and 12.
9. (1 point) Consider a sorting algorithm which first enqueues all elements of a sequence into a priority queue, and then dequeues all the elements from that priority queue. The priority queue uses an unsorted list implementation. What sorting algorithm does this description correspond to?
- A. This is a stable version of insertion sort.
 - B. This is an unstable version of insertion sort.
 - C. This is a stable version of selection sort.
 - D. This is an unstable version of selection sort.
10. (1 point) Consider the array (10,4,2,6,14,12,16,18,8). What is the content of the array after **two** iterations of the **in-place** heap sort algorithm to sort the array in **decreasing** order? Note: in case you need to apply heapify, you should apply the bottom-up algorithm with the $O(n)$ time complexity.
- A. (18,16,14,12,10,8,6,4,2)
 - B. (2,4,10,6,14,12,16,18,8)
 - C. (4,6,10,8,14,12,16,18,2)
 - D. (6,8,10,18,14,12,16,4,2)
11. (1 point) Consider a modified version of merge sort, which uses selection sort in the base case. How should the base case be defined such that this variant of the merge sort algorithm has a time complexity of $\mathcal{O}(n \cdot \log(n))$? If there are multiple correct options, choose the most general one.
- A. Immediately go into the base case, because selection sort is $\mathcal{O}(n \cdot \log(n))$.
 - B. Go into the base case after k recursive calls (k is a constant greater than 0).
 - C. Go into the base case when we have a list of size k or smaller (k is a constant greater than 0).
 - D. Go into the base case when we have a list of size 2 or 1.
12. (1 point) Say we have a list of n integers, with values ranging from 0 to 9. In the initial list, it is guaranteed that the integers from 0 to 4 are earlier in the list than the integers from 5 to 9 (for each integer i , if $i < 5$ then there exists no integer j earlier in the list such that $j > 4$). What is the fastest sorting algorithm to sort this list (for an arbitrarily large n)?
- A. Bucket sort
 - B. Insertion sort
 - C. Merge sort
 - D. Quick sort


13. (1 point) What is the expected time complexity of the *get*, *put* and *remove* methods with a hashmap with a load factor of $\lambda = 2$, for both separate chaining and linear probing? You may assume that the hash function is perfect, so each hash value in the range $[0, N - 1]$ will have an equal chance to be hashed to.
- Separate chaining: $\mathcal{O}(1)$, linear probing: $\mathcal{O}(n)$.
 - Separate chaining: $\mathcal{O}(n)$, linear probing: $\mathcal{O}(n)$.
 - Separate chaining: $\mathcal{O}(1)$, linear probing is not able to handle load factors above 1.
 - Separate chaining: $\mathcal{O}(n)$, linear probing is not able to handle load factors above 1.
14. (1 point) Consider the following AVL tree:



When deleting entry **19** from the tree, what trinode restructuring would occur? If the node of the entry to be deleted has non-null children, the node will be replaced with the in-order predecessor (i.e. the maximal entry in the left child of the entry).

- A single rotation trinode restructuring involving nodes 3, 8 and 10.
 - A double rotation trinode restructuring involving node 3, 8 and 10.
 - A single rotation trinode restructuring involving node 8, 10 and 18.
 - A double rotation trinode restructuring involving nodes 8, 10 and 18.
15. (1 point) Given a corresponding (2,4) tree and red-black tree. What is **false**?
- The depth of the leaf nodes in the (2,4) tree is equal to the black depth of the leaf nodes in the red-black tree.
 - The (2,4) tree has as many nodes as the red-black tree has black nodes.
 - The (2,4) tree has as many 3-nodes and 4-nodes as the red-black tree has red nodes.
 - The set of entries in the root of the (2,4) tree contains the entry found at the root of the red-black tree.
16. (1 point) Consider that we need a data structure to keep the scoreboard of a game, with entries stored as (key, value) pairs. We would like to be able to do frequent lookups of the top k players, and to do frequent lookups of all players with scores within a given range s . Note that the number of players with scores in a range s is typically much smaller than the size of the scoreboard. Which of the following data structures would be the most time-efficient for each of these operations?
- | | |
|--------------------------------------------------|-------------------------------------------------|
| A. Top k: Hashtable and Heap | Score in range s: Hashtable |
| B. Top k: Hashtable and Heap | Score in range s: Heap |
| C. Top k: Heap and AVL tree | Score in range s: Heap |
| D. Top k: Heap and AVL tree | Score in range s: AVL tree |
17. (1 point) Consider a graph G with n vertices and m edges implemented using an **adjacency map** and a list L of k vertices $[x_1, \dots, x_k]$. The task is to determine whether this list of vertices forms a path i.e. whether for all $i \in [1, \dots, k - 1]$, the vertices x_i and x_{i+1} are adjacent. What is the tightest upper bound on the time complexity of this operation?
- $\mathcal{O}(k)$
 - $\mathcal{O}(k \cdot n)$
 - $\mathcal{O}(n \cdot m)$
 - $\mathcal{O}(\sum_{v \in L} \deg(v))$

18. (1 point) Consider a directed acyclic graph G and a topological sorting of that graph where node a appears before node b . What must be **true**?
- A. There is a directed path from a to b .
 - B. There is no directed path from b to a .
 - C. If node b has no incoming edges, then a also had no incoming edges.
 - D. If there is a directed path from c to a and a directed path from c to b , the directed path from c to a has less edges than the directed path from c to b .
19. (1 point) Consider the following weighted graph:



- If we calculate the shortest path between A and F in this graph using Dijkstra's algorithm, how many times do we need to update the distance that is stored for node B (including the update from ∞ to an initial value)?
- A. 0 times
 - B. 1 time
 - C. 2 times
 - D. 3 times
20. (1 point) While applying Kruskal's algorithm to a given graph, we iterate over the edges in a sorted order. What happens if we encounter an edge between 2 vertices of the same cluster?
- A. We add the edge to the minimum spanning tree, nothing else.
 - B. We add the edge to the minimum spanning tree, and we remove the largest edge from the other edges that now form a cycle.
 - C. We do **not** add the edge to the minimum spanning tree, nothing else.
 - D. We do **not** add the edge to the minimum spanning tree, and the algorithm terminates because the minimum spanning tree is now complete.

Open questions (50%, 20 points)

21. (5 points) Consider the following Java implementation of method `methodX`. Note that `remove` corresponds to the dequeue operation and `add` corresponds to the enqueue operation.

```

1  public static int methodX(Queue<Integer> q, int x) {
2      int n = q.size();
3      if(n == 1) {
4          if(q.peek() == x) {
5              return 1;
6          } else {
7              return 0;
8          }
9      }
10     int sum = 0;
11     for(int i = 0; i < n; i++) {
12         int y = q.remove();
13         sum += methodX(q, x);
14         q.add(y);
15     }
16     return sum;
17 }

```

In this method, the size of the input n is the number of elements in queue q .

Give the recurrence equation expressing the tightest worst-case **time** complexity of `methodX` as a function of n . Define all variables and constants, and explain which lines of code contribute to each term of the recurrence equation. Make sure that each line is mentioned in at least one of the terms.

22. (5 points) Derive the closed form of the following recurrence equation:

$$\begin{aligned}
 T(5) &= c_0 \\
 T(n) &= c_1 + c_2 \cdot n + T(n-1) \quad \text{if } n > 5
 \end{aligned}$$

To answer this question, you should either:

- derive the closed form solution by repeatedly unfolding the recurrence equation, or
- guess the closed form and prove the correctness of your solution by induction.

Note the base case is for $n = 5$.

23. Prove that $\frac{1}{2} \cdot n^2 + n - 10$ is $\Omega(n^2)$.
- (2 points) State in detail the mathematical conditions that should be proven.
 - (2 points) Prove that these conditions hold. Explain all the steps in your proof.
24. In a construction market, wooden planks of different lengths are being sold. You are tasked to design the database that stores information about these planks. Each type of plank will be an item, which has two properties: length and price. Your database should be able to perform the following three operations:
- A new type of plank is being sold: given length and price, add the respective item to the database.
 - A type of plank is not sold anymore: given the length, remove the respective item from the database.
 - A customer wants to buy a type of plank: given a length, retrieve the price from the database.
- (3 points) What data structure would be most **time-efficient** for this database of plank type items? Explain your answer by giving the time complexities of the three operations for that data structure. Also explain how it stores the properties of each item.

- (b) (3 points) You have noticed that customers sometimes want a plank of custom length, which is not being sold. In those cases, you can sell them a plank that is slightly longer and they can saw it themselves to get the correct length. Thus, the third operation is replaced with:
- A customer wants to buy a type of plank: given a length, retrieve the length and price of the smallest plank that is at least as large as the given length.

What data structure is most **time-efficient** for this new operation? Explain why the data structure you chose in question 24.a cannot do this operation as efficiently. Describe how this operation can now be efficiently performed in the new data structure (i.e. explain the algorithm). You do not need to give or explain the time complexity.