

Major (COL 106)

Total marks: 34

2 September 2020, 9:30 AM - 11:30 AM.

Pages: 5

Answer all questions

IMPORTANT NOTE:

- **Accompanying Moodle Quiz** ends sharply at 11:45 AM. No Extensions will be provided at any cost. **We recommend that you fill in the quiz as you complete your answers here, instead of doing it all at once in the end of the exam. Moodle can get overloaded, and we will not make any allowances.** Note that you can restart the Moodle quiz if you lose the connection. Your already filled-in answers will remain saved. You can also go back and forth to questions on quiz to update your answers. Note that the quiz will auto-submit at the end of allotted time.
- **Answersheets have to be scanned to PDF** and uploaded on Gradescope before 3:00 PM. It is also necessary that proper pages are mapped to question answers. No extensions or relaxations will be provided. **We will evaluate only those pages mapped correctly. No submissions through emails will be accepted.**

Guidelines

- Read the complete guidelines posted on Piazza carefully, and follow them. *We have reproduced only a critical subset of them here, but you are expected to follow all.*
- All answers have to be handwritten (by yourself), using either blue or black pen.
- **You must begin answer to each question in a new page.**
- Each question is followed by the **maximum number of pages** allowed for its answer. We will not evaluate answers beyond that.
- Your submission **must contain** the hand-written honor-code (given in Piazza and reproduced below for convenience), signed to declare that you understand and have followed it.
- Your submission **also must contain** the scan of your id-card (list of allowed ids are given in the detailed guidelines on Piazza/Google doc). We will reject all submissions which do not come with id and honor-code.
- **A Moodle Quiz** submission should accompany your answersheet submission on Gradescope. We will evaluate the submission only if the quiz answers are fully consistent with your answersheet.

- Moodle quiz asks for each page of the question: (a) the first word and the last word of the first and last lines of the page, and (b) the first word and the last word of a given set of line numbers. **In the response you should type the line-specific two words separated by a space in a single line in the response.** If you have only one word in the line, write that word both times.
- Make appropriate and meaningful assumptions that may be necessary to answer questions, and mention them explicitly in your answer.
- We are aware of some websites which may provide you the final solution for a given question. But their use is prohibited. So if you just write the final solution, or replicate the answer from those websites, no marks will be awarded. Further, this will be considered as plagiarism, attracting disciplinary action.

Honor-code

I, <your name>, acknowledge the IITD Honor Code, and confirm that the submission below is entirely my own except where I have cited the source clearly. I have not:

1. used the services of any person, organization or discussions during the examination in preparation of this open-book exam answers.
2. given or received assistance in either accessing the question paper, or providing specimen answer/hint/diagram/model/code to or from other candidates submitting this open-book exam.

I am fully aware of the fact that if found to have used unfair and disallowed practices in the examination, I am liable for strict action which could include receiving a failure grade in this course, as well as other actions as deemed fit by the institute disciplinary committee.

Entry number:

Mobile number:

Signature:

Date:

1. (4 points) A chess tournament takes place between n players: $\{1, \dots, n\}$, where each participant plays with every other participant exactly once. The results of the tournament are being maintained in a matrix M of size $n \times n$. Initially, the matrix contains all 0 entries. For a match between player i and player j , we fill the matrix M based on the following criteria:

- If player i defeats player j then $M[i, j] = 1$ and $M[j, i] = 0$
- If the match is a draw then both $M[i, j] = 1$ and $M[j, i] = 1$
- If the two players fight during a match, then both $M[i, j] = 0$ and $M[j, i] = 0$

A player is declared the *winner* of the tournament if he/she properly wins all $n - 1$ matches without any draw. Given the matrix M , design an $O(n)$ algorithm (with pseudocode) to find the winner of the tournament. If there is no winner, then your algorithm should declare “*tournament wasted*”. **[Max Pages for Answer: 2]**

2. (3 points) Suppose you are enrolled in an elective course: “COL NAN”. This is a dummy course in which random marks are given to students based on certain criteria. This time the instructor has decided to distribute the marks based on your entry number. The instructor randomly orders the elements of the set of natural numbers from 1 to 200, $S := \{1, \dots, 200\}$. Suppose $A := (a_0, a_1, \dots, a_{199})$ is an ordering of S . If t is the last two digits of your entry number, you are awarded $a_{t+10}/2$ marks. Suppose the ordering A is an array representation of the *max-heap* on S such that $a_0 = 200$. What could be the maximum and minimum possible marks that you can get? Give proper reasoning for the same. **[Max Pages for Answer: 2]**

OR

Consider the hash function $h(x) = x \bmod 13$. Insert the following sequence of numbers in the the hash table of size 13 using quadratic probing: 24, 26, 49, 39, 36, 18, 14. Show all the steps and the entries where collision (if any) happens. **[Max Pages for Answer: 2]**

0	1	2	3	4	5	6	7	8	9	10	11	12

3. (4 points) Suppose $f: \mathbb{R} \rightarrow \mathbb{R}$ is a *convex* function (e.g., $f(x) = e^x$, $f(x) = x^2$, $f(x) = |x|$, etc.). You are given a set $S := \{x_1, x_2, \dots, x_n\}$ of n real numbers. Assume that $x_{i+1} \geq x_i$ for every $1 \leq i \leq n - 1$. Design an $O(\log n)$ time algorithm (with pseudocode) to find the value of the following objective function:

$$\Phi(S) := \min_{i \in \{1, \dots, n\}} f(x_i)$$

Assume that for each $i \in \{1, \dots, n\}$, computing $f(x_i)$ takes $O(1)$ time. **[Max Pages for Answer: 1]**

4. (6 points) A *Neat City* is a city with some buildings and roads in between buildings such that one can go from any building to any other via roads. *Arun* is an architect who is in charge of its plan. He went to his office and showed the plan to a board of members. *Raima*, *Arun's* greatest competitor in the office has pointed out that the city *Arun* planned is not a *Neat* one, as it has buildings that have no road connectivity between them.
- (a) (3 points) To *Arun's* surprise, *Raima* suggested to remove all the roads *Arun* planned, and add all those which were not there in *Arun's* plan. *Raima* is pretty sure that the city according to the new plan would indeed be a *Neat* one. Do you think *Raima* is right? You should provide a proof to justify your answer.
- (b) (3 points) On top of that *Raima* also shares her concern about minimizing the construction and maintenance cost of roads: Between any two buildings B_1 and B_2 , only one route suffices. *Raima* needs an algorithm that can tell her whether there is some road that need not be constructed for a given plan of *Neat City*. Design such an algorithm that requires time linear in the number of buildings and roads in the city. **[Max Pages for Answer: 3]**
5. (4 points) *Lata* has recently done a course on data structures and algorithm, and learnt about Quick-sort. One fine afternoon she came up with a new idea of sorting, and named it *MIXED-SORT*. *Lata* has one array of size n , and intends to use Quick-sort initially, but upon calling Quick-sort on a sub-array with fewer than k elements, *Lata* makes it simply return without sorting the sub-array. After the top-level call to Quick-sort returns, *Lata* runs insertion sort on the entire array to finish the sorting process. What would be the time complexity of this *MIXED-SORT* – derive based on your knowledge of the time complexity of Quick-sort and insertion-sort? Under what conditions can *MIXED-SORT* perform better than Quick-sort – derive these conditions. **[Max Pages for Answer: 2]**
6. (5 points) In a distributed computing system, suppose you are the coordinator, and you have K employees, each in charge of one machine coding up sorting algorithms to run on own machine. *Rajesh*, one of your employees has implemented a quick-sort and observed a surprising pattern when executed: The pivot creates the worst split for first two times, and a best split after that, and again two worst split followed by a best one, and so on.
- (a) (3 points) *Rajesh* is quite concerned as he thinks that his quick-sort will not work in $O(n \log n)$ due to the occurrences of those bad splits. Do you agree with him? Answer with a clear proof. Your justification should include calculation of the time-complexity of his quick-sort.
- (b) (2 points) Once all the employees are done with their work, they report to you, and send you their sorted lists. Now coordinator –that is, you– should write an efficient ($O(N \log K)$) strategy (pseudo-code) to combine all their results and prepare a single sorted list. You have to clearly state which data-structure(s) you plan to use, write the pseudo-code clearly and analyze the time-complexity of your strategy.

[Max Pages for Answer: 3]

7. (4 points) Consider the following array of numbers [8, 1, 15, 20, 13, 25, 27]. Insert these numbers sequentially from the beginning into a **Binary Search Tree** without performing any rotations. On the resulting Binary Search Tree, can you colour the nodes using red and black colors, to form a valid Red Black tree? If yes, show a valid coloring. If no, explain why. **[Max Pages for Answer: 2]**
8. (4 points) *Mohit* had been tasked with setting up radio stations in a newly developed township. After setting up the towers, he realized that some of the towers had been placed too close to each other, resulting in interference among the signals. He does not have the budget to re-install the stations at new locations. So he decided to split the spectrum bandwidth into two channels — *Channel-1* and *Channel-2*, with non-overlapping frequencies, and assign different channels to interfering stations. He has the list of n stations, and a list of k pairs of stations which interfere with each other. Can you help *Mohit* assign a channel to each radio station so that there is no interference among the radio stations? If it is not possible to perform such an assignment, you must let *Mohit* know. Write the pseudo-code for the proposed algorithm to perform the task in $O(n + k)$ time. **[Max Pages for Answer: 2]**

END OF THE QUESTION PAPER
