

NAME: _____

CMPSCI 611
Advanced Algorithms
First Midterm Exam Fall 2017

A. McGregor

11 October 2017

DIRECTIONS:

- Do not turn over the page until you are told to do so.
- This is a *closed book exam*. No communicating with other students, or looking at notes, or using electronic devices. You may ask the professor or TA to clarify the meaning of a question but do so in a way that causes minimal disruption.
- If you finish early, you may leave early but do so as quietly as possible. The exam script should be given to the professor.
- There are five questions. All carry the same number of marks but some questions may be easier than others. Don't spend too long on a problem if you're stuck – you may find that there are other easier questions.
- The front and back of the pages can be used for solutions. There are also a blank page at the end that can be used. If you are using these pages, clearly indicate which question you're answering. Remember that the best answers are those that are clear and concise (and of course correct). Most questions can be answered in two or three sentences.

1	/10
2	/10
3	/10
4	/8+2
5	/10
Total	/48+2

Question 1. For each of the following statements, indicate whether they are TRUE or FALSE by circling the appropriate answer. No justification is required.

1. If the running time $T(n)$ of an algorithm satisfies $T(n) = O(n^2)$ and $T(n) = \Omega(n^2)$ then it also satisfies $T(n) = \Theta(n^2)$.

TRUE FALSE

2. If a subset system (E, \mathcal{I}) doesn't satisfy the cardinality property, there exist subsets $A, B \in \mathcal{I}$ such that $|B| > |A|$ and $A + e \notin \mathcal{I}$ for all $e \in B - A$.

TRUE FALSE

3. If the complex number z is a 4th root of unity, then it is also an 8th root of unity.

TRUE FALSE

4. Consider an undirected graph where all edges have weight at least 1. If there are exactly two edges with weight 1 then any minimum spanning tree includes both of these edges.

TRUE FALSE

5. If E consists of n elements and \mathcal{I} consists of 2^n subsets of E , then (E, \mathcal{I}) is a matroid.

TRUE FALSE

Question 2. No justification required for your answers.


1. Consider the following undirected graph. The number on an edge indicates its weight.

Weight of minimum spanning tree =

Weight of maximum weight matching =

Length of shortest path between c and d =

Diameter of the graph =



For diameter and shortest paths, you should treat the weight as the length of the edge.

2. What is the adjacency matrix of the graph considered above? You should ignore the weights on the edges.

3. Suppose that a divide and conquer algorithm solves a problem of size n by first solving three instances of size $n/2$ and then taking $O(n)$ additional basic steps. If the time to create the three instances is $O(n^2)$, what is the running time, $T(n)$, of the algorithm.

4. If $T(n) = T(n-1) + n$ for $n > 1$ and $T(1) = 1$, then $T(n) = \Theta(n^x)$ for what value of x ?

Question 3. As part of an ongoing study, scientists in Iceland have been logging the noon temperature over the last n days. Let $T[i]$ be the temperature on day i for $i = 1, 2, \dots, n$. Temperatures can be positive or negative.

1. Design an algorithm that finds, given a value of r , finds the set of r consecutive days with the maximum average temperature. Analyze the running time of your algorithm but no proof of correctness is required. The running time should not depend on r .

2. Consider the following algorithm that claims to determine whether there is a set of consecutive days whose average temperature is exactly 0.

Step 1. Let $C[1] = T[1]$ and for $i = 2, \dots, n$, let $C[i] = T[i] + C[i - 1]$

Step 2. Sort C by increasing value using MergeSort and let C' be the resulting array.

Step 3. If there exists i such that $C'[i] = C'[i+1]$, output “yes”. Output “no” otherwise.

- (a) What is the running time of this algorithm?

- (b) Is the algorithm correct? Either provide an example where the algorithm returns the wrong answer or provide a brief proof of correctness.

3. Rather than releasing all the data exactly, the scientists release a rough approximation of their reading. We say list L is a k -bucket histogram if there is *at most* $k - 1$ values of i such that $L[i] \neq L[i + 1]$. For example, $L = [6, 6, 6, 2, 2, 3, 3, 3, 3, 2, 2, 4]$ is a 5-bucket histogram of length 12. Design a dynamic programming algorithm that finds the k -bucket histogram L of length n that minimizes the following quantity:

$$\text{error}(L, T) = \sum_{i=1}^n (L[i] - T[i])^2$$

Hint: You may assume that for any set of real numbers S , $\sum_{e \in S} (e - x)^2$ is minimized by setting x to be $\text{average}(S) = \sum_{e \in S} e / |S|$. You may wish to consider the value $D[j, k']$ that is the minimum possible value of $\sum_{i=1}^j (L[i] - T[i])^2$ where L is a k' -bucket histogram.

(a) Briefly describe your algorithm.

(b) Briefly explain why your algorithm returns the correct answer.

(c) Analyze the running time of your algorithm.

Question 4. Let $\mu(n)$ be the time to multiply two $n \times n$ matrices together. For simplicity, you may ignore any dependence on the size of the entries of the matrices. Given a matrix A , the square of the matrix is defined as $A^2 = A \times A$ and, more generally, for $p \geq 2$, the p th power is defined as $A^p = A \times A \times \dots \times A$ where the product consists of p copies of A . The faster the algorithms you design, the more marks you will get. Your answers should include a *brief description* of your algorithms and an analysis of the running time (perhaps in terms of $\mu(n)$). You *do not need* to prove correctness.

1. Design an algorithm for computing the p th power of a matrix A where p is a power of two. You may assume the existence of a subroutine to multiply matrices together.

2. Design an algorithm for computing the p th power of a matrix D where $D[i, j] = 0$ if $i \neq j$ and $D[i, i] \in \{-1, 1\}$ for all i . The algorithm only needs to output the non-zero entries of D^p . In this part of the question, you *may not assume* p is a power of two.

3. **Extra Credit.** Suppose you have developed an algorithm that can find the square of a matrix in $O(n^2)$ time. Using this algorithm as a subroutine, find an algorithm that multiplies two $n \times n$ different matrices together in $O(n^2)$ time.

Question 5. A thief breaks into a house and finds n objects that could be stolen. The i th item is valued at $v_i > 0$ dollars and weighs $w_i \in \{1, 2, \dots, W\}$ kilograms. The thief can carry at most W kilograms where W is some positive integer.

1. Let $D[i, j]$ be the maximum total value that can be stolen when the thief can carry at most j kilograms and may pick any subset of the first i items. Write a formula for $D[i, j]$ assuming $D[i', j']$ has already been computed for all $i' < i$ and $j' \leq j$.

$$D[i, j] =$$

2. Based on your above answer, write a dynamic program for computing the maximum total value of a set of objects that the thief can carry. You do not need to argue correctness but do need to state the running time.

3. Consider the greedy algorithm that orders the objects by decreasing value and adds each object to the collection to be stolen if it does not violate the weight constraint. Give an example with $W = 10$ and three objects with different values such that this greedy does *not* maximize the total value stolen.

$$w_1 = \quad v_1 = \quad w_2 = \quad v_2 = \quad w_3 = \quad v_3 =$$

4. Suppose that the objects are *divisible*, i.e., it is possible to take any fraction of each object. For example, if one of the objects was milk, you can take all the milk, none of the milk, or any amount in between. If you take a f th fraction of object i the weight is $f \times w_i$ and the value is $f \times v_i$. Let $r_i = v_i/w_i$. Design a greedy algorithm to find the maximum total value that can be stolen. You should not assume the items are initially sorted in any particular way.

(a) Briefly describe your algorithm.

(b) What is the running time?

(c) Prove the correctness. **Hint:** Let f_i be the fraction of item i that your algorithm takes and consider another solution where g_i is the fraction of item i that is taken. If the solutions are different, show that the solution with the g_i fractions is not optimal. For simplicity, you may assume all the v_i/w_i values are different. Your analysis does not need to involve matroids.

