

CSE 201

Homework Assignment 7

Solutions

1. Recall the coin changing problem: Given denominations $d = (d_1, d_2, \dots, d_n)$ and an amount N to be paid, determine the number of coins in each denomination necessary to disburse N units using the fewest possible coins. Assume that there is an unlimited supply of coins in each denomination.
 - a. Write pseudo-code for a greedy algorithm that attempts to solve this problem. (Recall that the greedy strategy doesn't necessarily produce an optimal solution to this problem. Whether it does or not depends on the denomination set d .) Your algorithm will take the array d as input and return an array G as output, where $G[i]$ is the number of coins of type i to be disbursed. Assume the denominations are arranged by increasing value $d_1 < d_2 < \dots < d_n$, so your algorithm will step through array d in reverse order. Also assume that $d_1 = 1$ so any amount can be paid.

Solution:

GreedyCoinChange(d, N) (Pre: $1 = d[1] < d[2] < \dots < d[n]$)

1. $n = \text{length}[d]$
2. $G[1 \dots n] = (0, 0, \dots, 0)$
3. $\text{sum} = 0$
4. $i = n$
5. while $\text{sum} < N$
6. while $\text{sum} + d[i] \leq N$
7. $\text{sum} += d[i]$
8. $G[i]++$
9. $i--$
10. return G

- b. Let $d_i = b^{i-1}$ for some integer $b > 1$, and $1 \leq i \leq n$, i.e. $d = (1, b, b^2, \dots, b^{n-1})$. Does the greedy strategy always produce an optimal solution in this case? Either prove that it does, or give a counter-example.

Theorem

The greedy strategy produces an optimal solution with this denomination set for any amount N .

Proof:

Let $N \geq 0$, and suppose that $x = (x_1, x_2, \dots, x_n)$ is an optimal disbursement of N units, where x_i is the number of coins of type i to be paid ($1 \leq i \leq n$). Let $g = (g_1, g_2, \dots, g_n)$ be the solution produced by the greedy strategy. To show that g is optimal, it suffices to show that $g_i = x_i$ for $1 \leq i \leq n$. Since both solutions disburse N units, we have $\sum_{i=1}^n x_i b^{i-1} = N = \sum_{i=1}^n g_i b^{i-1}$, i.e.

$$x_1 + x_2 b + x_3 b^2 + \dots + x_n b^{n-1} = g_1 + g_2 b + g_3 b^2 + \dots + g_n b^{n-1}.$$

Reducing this equation mod b yields the congruence $x_1 \equiv g_1 \pmod{b}$. Observe $0 \leq x_1 < b$, since otherwise it would be possible to replace b coins of type 1 by 1 coin of type 2, contradicting that x is optimal. Likewise $0 \leq g_1 < b$, since the greedy strategy guarantees that as many coins as possible of type 2 (value b) will be disbursed before any coins of type 1 are. These inequalities,

along with the preceding congruence, imply $x_1 = g_1$. Cancel x_1 from both sides of the above equation and divide through by b to obtain

$$x_2 + x_3b + \cdots + x_nb^{n-2} = g_2 + g_3b + \cdots + g_nb^{n-2}.$$

A similar argument shows that $x_2 = g_2$. Continuing in this manner, we can show $g_i = x_i$ for $1 \leq i \leq n$, and hence the greedy solution g is in fact optimal, as required. ■

- c. Let $d_1 = 1$ and $2d_i \leq d_{i+1}$ for $1 \leq i \leq n-1$. Does the greedy strategy always produce an optimal solution in this case? Either prove that it does, or give a counter example.

Counter Example:

Let $n = 3$, $d = (1, 10, 25)$, and $N = 30$. Observe that $2 \cdot 1 \leq 10$, and $2 \cdot 10 \leq 25$. The greedy solution in this case is $g = (5, 0, 1)$, which uses 6 coins, while the optimal solution $x = (0, 3, 0)$ uses only 3 coins. ■

2. **Activity Scheduling Problem:** Consider n activities $\{1, 2, \dots, n\}$ with start times $s_1, \dots, s_n$ and finish times $f_1, \dots, f_n$, that must use the same resource (such as lectures in a lecture hall, or jobs on a machine.) At any time only one activity can be scheduled. Two activities i and j are *compatible* if their time intervals $[s_i, f_i]$ and $[s_j, f_j]$ have non-overlapping interiors. Your objective is to determine a set of compatible activities of maximum possible size. For each of the greedy strategies below, determine whether or not it provides a correct solution to all instances of the problem. If your answer is yes, state and prove a theorem establishing the correctness of the proposed strategy. If your answer is no, provide a counterexample (i.e. specific start and end times) showing that the strategy can fail to find an optimal solution.

- a. Order the activities by increasing total duration. Schedule activities with the shortest duration first, satisfying the compatibility constraint. If there is a tie, choose the one that starts first.

Counter-Example:

Activity:	1	2	3
Start time:	2	0	3
Finish time:	4	3	6

The greedy strategy yields the schedule $\{1\}$, while the optimal schedule is $\{2, 3\}$.

- b. Order the activities by increasing start time. Schedule the activities with the earliest start times first, satisfying the compatibility constraint. If there is a tie, choose the one having shortest duration.

Counter-Example:

Activity:	1	2	3
Start time:	0	1	2
Finish time:	3	2	4

The greedy strategy yields the schedule $\{1\}$, while the optimal schedule is $\{2, 3\}$.

- c. Order the activities by increasing finish times. Schedule the activities with the earliest finish times first, satisfying the compatibility constraint. If there is a tie, pick one arbitrarily.

Theorem

This strategy maximizes the number of scheduled activities.

Proof:

Let $A = \{1, 2, \dots, n\}$ denote set of activities, let $s(1), s(2), \dots, s(n)$ be the start times, and let $f(1), f(2), \dots, f(n)$ be the finish times. Assume that A is already ordered by finish times: $f(1) \leq f(2) \leq \dots \leq f(n)$. Observe two activities i and j with $i < j$ are compatible iff $f(i) \leq s(j)$. We call a subset of A *feasible* if all its activities are mutually compatible. We call a feasible set *optimal* if it has maximum cardinality amongst all feasible subsets of A .

Let $X = \{i_1 < i_2 < \dots < i_k\} \subseteq A$ be the subset obtained by the greedy strategy described above, and let $Y = \{j_1 < j_2 < \dots < j_l\} \subseteq A$ be any optimal subset, so in particular $l \geq k$. Clearly the greedy strategy guarantees X is feasible. We must show that X is also optimal, i.e. $l = k$.

If $X = Y$ there is nothing to prove, so suppose $X \neq Y$. Let t be the first index $1 \leq t \leq k$ such that $i_t \neq j_t$. In other words $i_r = j_r$ for all $r < t$, and $i_t \neq j_t$. (Such an index t must exist for otherwise $i_r = j_r$ for all $1 \leq r \leq k$, and hence $X \subseteq Y$. Since $X \neq Y$ there must be an element $j_{k+1} \in Y - X$ that is compatible with all the activities in X . But the greedy strategy should then have added j_{k+1} to X . Since it did not, an index t with the above properties must exist.) Both activities i_t and j_t must therefore be compatible with $\{i_1, i_2, \dots, i_{t-1}\} = \{j_1, j_2, \dots, j_{t-1}\}$. Since i_t was chosen by the greedy strategy we must have $f(i_t) \leq f(j_t)$. The compatibility constraint implies that $f(j_t) \leq s(j_{t+1})$, and therefore $f(i_t) \leq s(j_{t+1})$. It follows that i_t is also compatible with j_{t+1} , and hence also with all the activities in the set $\{j_{t+1}, \dots, j_l\}$.

Let $Y_1 = Y - \{j_t\} \cup \{i_t\}$. The above remarks show that Y_1 is a feasible set, and since $|Y_1| = |Y|$, Y_1 is optimal as well. But Y_1 has one more activity in common with X than Y does. If $X = Y_1$, then X is optimal and we are done. If not, we repeat the above process with Y_1 in place of Y , to obtain another optimal set Y_2 having one more activity in common with X than Y_1 . Continuing in this fashion we construct a sequence of optimal sets, each having more in common with X than the previous. Eventually we must reach an optimal set Y_m that is identical to X , showing that X is itself optimal as claimed. ■

3. Let m be an integer in the set $\{0, 1, 2, 3, 4, 5, 6, 7, 8, 9\}$, and consider the following problem: determine m by asking 3-way questions, i.e. questions with at most 3 possible answers. For instance, one could ask which of 3 specific subsets m belongs to.
 - a. Give a decision tree argument showing that at least 3 such questions are necessary in worst case. In other words, prove that no correct algorithm can solve this problem by asking only 2 questions in worst case.

Proof:


Any algorithm that solves this problem can be represented by ternary decision tree, since each question can have at most 3 answers. Since there are 10 possible verdicts, the height of such a tree must satisfy $h \geq \lceil \log_3(10) \rceil = 3$. Therefore no such algorithm can ask less than 3 questions in worst case. ■

- b. Design an algorithm that will solve this problem by asking 3 such questions in worst case. Express your algorithm as a decision tree.

Solution:

Each internal node below represents a question asking whether m belongs to one of 3 possible subsets of $\{0, 1, 2, 3, 4, 5, 6, 7, 8, 9\}$. For instance $0123 | 456 | 789$ represents the question: "does m belong to $\{0, 1, 2, 3\}$, to $\{4, 5, 6\}$, or to $\{7, 8, 9\}$?" Verdicts are placed in brackets "[]".

Decision tree:



4. **Bar Weighing Problem**

Assume we are given 12 gold bars numbered 1 to 12 where 11 bars are pure gold and one is counterfeit: either gold-plated lead (which is heavier than gold), or gold-plated tin (lighter than gold). The problem is to find the counterfeit bar and what metal it is made of using only a balance scale.



Any number of bars can be placed on each side of the scale, and each use of the scale produces one of three outcomes: either the left side is heavier, or the two sides are the same weight, or the right side is heavier.

- a. Give a decision tree lower bound for the (worst case) number of weighing's that must be performed by any algorithm solving this problem.

Theorem

Any algorithm that solves this problem must perform at least 3 weighing's in worst case.


Proof: Since there are 3 possible outcomes to each weighing, the decision tree for this problem is a ternary tree. Observe that there are 24 distinct verdicts, which may be represented as: $\{1L, 2L, 3L, \dots, 12L, 1H, 2H, 3H, \dots, 12H\}$. For instance, $8L$ is the conclusion that bar 8 is light (i.e. tin) and all other bars are genuine, while $11H$ represents the conclusion that bar 11 is heavy (lead) and all others are pure gold.

The height h of a decision tree representing any algorithm for this problem must therefore satisfy $h \geq \lceil \log_3(24) \rceil = 3$. Therefore at least 3 weighing's must be performed (in worst case) to identify the counterfeit bar. ■

- b. Design an algorithm that solves this problem with (worst case) number of weighing's equal to the lower bound you found in (a). Present your algorithm by drawing a decision tree, rather than pseudo-code.

Solution:

The following decision tree represents one algorithm (among many) that identifies the counterfeit bar in exactly 3 weighing's.



- c. Alter the problem slightly to allow the possibility that all 12 bars are pure gold. Thus, there is one additional possible verdict: "all gold". Make a minor change to your algorithm in part (b) so that it gives a correct answer to this more general problem.

Solution:

Observe that if all bars are pure gold, then all weighing's (of equal numbers of bars) will balance. Therefore, by adding a single leaf to the above tree labeled AllGold as the middle child of the 1:12 node, we obtain an algorithm that solves this problem in the more general setting.

5. **Water Jug Problem** (Problem 8-4: page 206 of CLRS 3rd edition)

Suppose that you are given n red and n blue water jugs, all of different shapes and sizes. All red jugs hold different amounts of water, as do the blue ones. Moreover, for every red jug, there is a blue jug that holds the same amount of water, and vice versa.



It is your task to find a grouping of the jugs into pairs of red and blue jugs that hold the same volume of water. To do so, you may perform the following operation: pick a pair of jugs, one red, one blue, fill the red jug with water, and then pour the water into the blue jug. This operation will tell you if the two jugs hold the same amount of water, and if not, which one holds more water. Assume that such a comparison takes one unit of time. Your goal is to find an algorithm that solves this problem. Remember that you may not directly compare two red jugs or two blue jugs.

- a. Describe an algorithm that uses $\Theta(n^2)$ comparisons (in worst case) to group the jugs into pairs.

Solution:

Call the red jugs $R_1, R_2, \dots, R_n$ and the blue jugs $B_1, B_2, \dots, B_n$. Begin by comparing R_1 to each of the n blue jugs until a match is found. At most $n - 1$ comparisons are necessary to determine a match, since if no match has been found during $n - 1$ comparisons, the last blue jug must match, and no additional comparison is necessary. Next compare R_2 to the $n - 1$ unmatched blue jugs. This time at most $n - 2$ comparisons are necessary to determine a match. Next compare R_3 to at most $n - 3$ blue jugs. In general, one compares R_i to at most $n - i$ blue jugs. The last red jug R_n doesn't need to be compared to any blue jugs since there is only one possibility left at that point. Once this process is complete, a matching of red and blue jugs has been attained. The total number of comparisons performed is at most

$$\sum_{i=1}^{n-1} (n - i) = \sum_{i=1}^{n-1} i = \frac{n(n-1)}{2} = \Theta(n^2)$$

■

- b. Prove a lower bound of $\lceil \log_3(n!) \rceil$ for the worst case, and $\log_3(n!)$ for the average case number of comparisons to be performed by any algorithm that solves this problem.

Proof:

Each comparison described above has 3 possible outcomes, and hence any algorithm solving this problem is represented by a ternary tree. There is a total of $n!$ ways to match the n red jugs to the n blue jugs, so this problem has $n!$ possible verdicts. Any algorithm solving this problem can be represented by a decision tree whose height satisfies $h \geq \lceil \log_3(n!) \rceil$ and whose average height satisfies $a \geq \log_3(n!)$. These quantities therefore serve as lower bounds on the worst case and average case number of comparisons, respectively.

■

- c. Write pseudo-code for an algorithm that solves this problem and that runs in (average case) time $\Theta(n \log n)$. (Hint: modify the algorithms Partition() and Quicksort().)

Solution:

The input for our algorithm will be two arrays $A[1 \dots n]$ and $B[1 \dots n]$ each containing the same set of n distinct numbers, in different orders. Our goal is to simultaneously sort these two arrays without ever comparing two elements in the same array. The desired matching is then given by $A[i] == B[i]$ for $1 \leq i \leq n$. The following recursive algorithms Match() and MatchPartition() are based on Quicksort() and Partition(), respectively. Both Match(A, B, p, r) and its subroutine MatchPartition(A, B, p, r) have as precondition that $A[p \dots r]$ and $B[p \dots r]$ contain the same $r - p + 1$ distinct elements.

Match(A, B, p, r)

1. if $p < r$
2. $q = \text{MatchPartition}(A, B, p, r)$
3. Match($A, B, p, q - 1$)
4. Match($A, B, q + 1, r$)

The subroutine MatchPartition(A, B, q, r) re-arranges the subarrays $A[p \dots r]$ and $B[p \dots r]$, returning an index q such that $A[q] == B[q]$ and such that both

$$A[p \dots (q - 1)] < A[q] < A[(q + 1) \dots r],$$

and

$$B[p \dots (q - 1)] < B[q] < B[(q + 1) \dots r].$$

The precondition on $A[p \dots r]$ and $B[p \dots r]$ implies, with the above inequalities, that the subarray pairs $A[p \dots (q - 1)]$, $B[p \dots (q - 1)]$ and $A[(q + 1) \dots r]$, $B[(q + 1) \dots r]$ satisfy the very same condition, namely that each pair contains the same distinct elements. This guarantees that the recursive calls to Match() on lines 3 and 4 will be valid.

MatchPartition(A, B, p, r)

1. $i = p - 1$
2. for $j = p$ to $r - 1$
3. if $A[j] == B[r]$
4. $A[j] \leftrightarrow A[r]$ (swap)
5. if $A[j] < B[r]$
6. $i++$
7. $A[j] \leftrightarrow A[i]$
8. $i = p - 1$
9. for $j = p$ to $r - 1$
10. if $B[j] < A[r]$
11. $i++$
12. $B[j] \leftrightarrow B[i]$
13. $A[r] \leftrightarrow A[i + 1]$
14. $B[r] \leftrightarrow B[i + 1]$
15. return $(i + 1)$

Lines 1-7 partition $A[p \dots r]$ around the pivot $B[r]$, along the way finding the pivot in $A[p \dots r]$ and placing it at $A[r]$. At this point we have $A[p \dots i] < B[r] < A[(i + 1) \dots (r - 1)]$ and

$A[r] == B[r]$. Lines 8-12 similarly partition $B[p \cdots r]$ around $A[r]$. Lines 13 and 14 correctly place the pivot in both arrays, and line 15 returns the pivot index. Observe that the comparisons on lines 3, 5 and 10 are always between elements of different arrays. Thus $\text{MatchPartition}()$ performs $3(n-1)$ such comparisons on a subarray of length n .

The analysis of the (average case) runtime of $\text{Match}(A, B, 1, n)$ is almost identical to that of $\text{Quicksort}(A, 1, n)$. (See the lecture notes on the average case runtime of $\text{Quicksort}()$.) The recurrence for $\text{Match}(A, B, 1, n)$ is

$$t(n) = 3(n-1) + \left(\frac{2}{n}\right) \cdot \sum_{q=1}^{n-1} t(q)$$

One checks that the effect of the factor 3 in the first term is to multiply the solution itself by 3. Thus $t(n) = 3(-4n + 2(n+1)H_n) \sim 6n \ln(n) = \Theta(n \log n)$. ■

6. Show that at least $\binom{n}{2}$ “adjacency” questions are necessary to determine whether a graph G on n vertices is acyclic. (Hint: use the following adversary strategy. Answer yes to any edge probe, unless that answer would prove the existence of a cycle.)

Proof: First observe that the result is trivial when $n = 2$ since in that case, there is only one potential edge to probe. We may therefore assume without loss of generality that $n \geq 3$. Consider any algorithm for this problem, and run it against the following adversary simulating an input graph G with n vertices. The daemon’s strategy is to answer yes to any edge probe, unless that answer would prove the existence of a cycle. More precisely, the daemon maintains two edge sets A and B , where initially B is empty and A contains all $\binom{n}{2}$ edges in K_n . He then performs the following algorithm whenever the algorithm probes an edge e .

Probe(e)

1. if $B + e$ is acyclic
2. $B = B + e$
3. return Yes
4. else
5. $A = A - e$
6. return No

Observe that at all times we have $B \subseteq A$, and the set $A - B$ consists of exactly those edges in K_n not yet been probed. Also note, both A and B are consistent with the daemon’s entire sequence of answers, since when the answer Yes is given, the edge is added to B and not removed from A , while if No is given, the edge is removed from A and not added to B . The following invariants are maintained over any sequence of edge probes:

- (a) The subgraph B is acyclic. This is obvious from the daemon’s strategy.
- (b) If B is disconnected, then A contains all edges of K_n that join distinct connected components of B . **Proof:** Let e be an edge joining two components of B . Thus, if e were added to B , it would not create a cycle in B . Such an edge could not have been removed from A , hence $e \in A$.
- (c) The subgraph A is connected. **Proof:** Assume, to get a contradiction, that at some point an edge e is probed and removed from A , causing it to become disconnected. At the time of its removal, e does not form a cycle with the edges of A (since otherwise its removal would not disconnect

A.) Since $B \subseteq A$, e does not form a cycle with the edges of B either. But then e should have been added to B rather than removed from A . This contradiction shows that no such edge e exists, and therefore A is connected.


- (d) If $A \neq B$, then A contains a cycle. **Proof:** Assume, to get a contradiction, that A is acyclic. Then being connected, A is a tree. Since $B \subseteq A$ and $A \neq B$, we see B must be disconnected. (Another property of trees: removing an edge from a tree results in a disconnected graph.) Let $e \in A - B$. Then e joins two components of B . Since A is acyclic, it can contain no other edges joining the same two components. But since $n \geq 3$, at least one other edge e' joining those two components exists in K_n . By (b) above, e' also belongs to A , contradicting our earlier observation. This contradiction shows that A must contain a cycle.

Now suppose the algorithm halts and returns a verdict (acyclic/not acyclic) after probing fewer than $\binom{n}{2}$ edges. Then some edge of K_n was not probed, hence $A \neq B$. Remark (d) implies that A is not acyclic, and by (a) B is acyclic. Since the algorithm cannot distinguish between A and B , it cannot be correct. If the algorithm says G contains a cycle, then the daemon can claim $G = B$, while if the algorithm says G is acyclic, the daemon claims $G = A$. Thus, any correct algorithm solving this problem must probe all $\binom{n}{2}$ edges of K_n . ■

7. Let $b = x_1x_2x_3x_4x_5$ be a bit string of length 5, i.e. $x_i \in \{0, 1\}$ for $1 \leq i \leq 5$. Consider the following problem. Determine whether or not b contains the substring 111. Restrict attention to those algorithms whose only operation is to peek at a bit. Obviously 5 peeks are sufficient. A decision tree argument provides the (useless) fact that at least one peek is necessary.
- a. Design an algorithm that solves the problem using only 4 peeks in worst case. Express your algorithm as a decision tree.

Solution:

The following decision tree is just one of many correct solutions. Each internal node represents the index of the bit being peeked, while the leaves indicate verdicts. A branch to the left signifies that the peek saw a 0, while a branch to the right indicates the peek detected a 1.



- b. Use an adversary argument to show that 4 peeks are necessary in general.

Proof:

Consider any algorithm purporting to solve this problem, and start it on an unspecified bit string b of length 5. The daemon's strategy is as follows. The first time one of the bits $\{b_2, b_4\}$ are peeked, answer 0, the second time a bit in this set is peeked, answer 1; any time one of the bits $\{b_1, b_3, b_5\}$ is peeked, answer 1. To keep track, the daemon maintains two bit strings x and y ,

where initially $x = 11111$ and $y = 00000$. Each time the answer 0 is given, the corresponding bit in x is flipped to 0, and if the answer 1 is given, the corresponding bit in y is flipped to 1. We may assume the algorithm never peeks at the same bit twice (otherwise the daemon just repeats his previous answer.) Expressing the daemon's strategy in pseudo-code we have:

Peek(i)

1. if $i \in \{1, 3, 5\}$
2. $y_i = 1$ and answer 1
3. else if $i == 2$
4. if $x_4 == 0$ // b_4 was already peeked and answered 0
5. $y_2 = 1$ and answer 1
6. else // this is our first peek at $\{b_2, b_4\}$
7. $x_2 = 0$ and answer 0
8. else // $i == 4$
9. if $x_2 == 0$ // b_2 was already peeked and answered 0
10. $y_4 = 1$ and answer 1
11. else // this is our first peek at $\{b_2, b_4\}$
12. $x_4 = 0$ and answer 0

Note that after any set of bit peeks, both strings x and y are consistent with the daemon's sequence of answers: if 0 was returned for the i^{th} bit, then x_i was flipped to 0 and y_i remained 0, while if 1 was returned for the i^{th} bit, then y_i was flipped to 1 and x_i remained 1. Notice also that any sequence of bit peeks results in at most one 0 being returned, and hence x can only be one of the strings 11111, 10111 or 11101, all of which contain 111.

Now suppose the algorithm halts and returns an answer (yes/no) after performing at most 3 peeks. In this case, y cannot contain the substring 111. (The only way y can contain three 1's is if bits 2 and 4 were never peeked. But then y is 10101, which does not contain 111.) We conclude that the algorithm cannot be correct, since after only three peeks, there will necessarily exist a bit string consistent with the daemon's sequence of answers, but which contradicts the algorithm's verdict. Therefore, any correct algorithm for this problem must perform at least 4 bit peeks in worst case. ■