

CSE 201

Final Exam Review Problems

Solutions to some of the more difficult problems

2. Suppose we are given 4 gold bars (labeled 1, 2, 3, 4), one of which *may* be counterfeit: gold-plated tin (lighter than gold) or gold-plated lead (heavier than gold). Again the problem is to determine which bar, if any, is counterfeit and what it is made of. The only tool at your disposal is a balance scale, each use of which produces one of three outcomes: tilt left, balance, or tilt right.
- c. Find an adversary argument that proves 3 weighings are necessary (in worst case), and therefore the algorithm you found in (b) is best possible. (Hint: study the adversary argument for the min-max problem discussed in the handout to gain some insight into this problem. Further hint: put some marks on the 4 bars and design an adversary strategy that, on each weighing, removes the fewest possible marks, then show that if the balance scale is only used 2 times, not enough marks will be removed.)

Proof:

We assume that any algorithm for this problem always places an equal number of bars on the left and right sides of the scale, since the greater number of bars will always outweigh the lesser. When such an algorithm is run against the following adversary, the daemon first marks each bar with both a + sign (meaning the bar is possibly heavy), and a - sign (meaning it may be light). As the algorithm progresses, the daemon will remove certain marks. Each bar is therefore always in one of 4 states.

$\pm$: a candidate both for heavy and light

+: a candidate for heavy, but not light

-: a candidate for light, but not heavy

N (no mark): neither heavy nor light, known to be pure gold

Each answer that the daemon gives to a weighing (*left*, *balance* or *right*), implies that some marks should be erased. If for instance, a probe is answered by stating that the scale tilted left, then all - signs on the left and all + signs on the right must be deleted. Tilting left also implies that all bars not on the scale are genuine, and hence their marks should be erased as well. If the daemon says the scale tilted right, then all + signs on the left, all - signs on the right, and any marks on bars not on the scale should be removed. If the daemon says that the scale balances, then all bars on the scale are genuine, so their marks must be erased.

The daemon's strategy is to always answer in such a way that the minimum number of marks are removed. Specifically, let $L, R \subseteq \{1, 2, 3, 4\}$ denote the sets of bars placed on the left and right sides of the scale, respectively. Define

$L_+ = \#$ of + marks on the left side of the scale

$L_- = \#$ of - marks on the left side of the scale

$R_+ = \#$ of + marks on the right side of the scale

$R_- = \#$ of - marks on the right side of the scale

$a = L_+ + R_-$

$b = L_- + R_+$

$c = \#$ of marks on bars not placed on the scale

Thus the daemon's strategy is to find $\min(a + b, a + c, b + c)$, remove that number of marks and answer accordingly. Observe that both $a + b \leq a + c$ and $a + b \leq b + c$ are true if and only if $b \leq c$ and $a \leq c$, which is equivalent to $c = \max(a, b, c)$. Similarly, $a + c \leq a + b$ and $a + c \leq b + c$ are

true iff $b = \max(a, b, c)$, and also $b + c \leq a + b$ and $b + c \leq a + c$ hold iff $a = \max(a, b, c)$. When the algorithm places bars on the scale, the daemon executes the following algorithm.

```

Weigh(L,R)  pre: |L| == |R|
a.  compute the quantities  $a, b$  and  $c$  defined above
b.   $m = \max(a, b, c)$ 
c.  if  $a == m$ 
d.      remove  $b + c$  marks:  $-$  on left,  $+$  on right, all marks not on the scale
e.      return "left"
f.  else if  $b == m$ 
g.      remove  $a + c$  marks:  $+$  on left,  $-$  on right, all marks not on the scale
h.      return "right"
i.  else //  $c == m$ 
j.      remove  $a + b$  marks: all marks on the scale
k.      return "balance"

```

Now suppose the algorithm halts and returns a verdict after using the scale at most 2 times. Since there are only 4 bars, and since $|L| = |R|$, only two kinds of weighings can take place: 2 bars on each side, or 1 bar on each side. Since the scale is used two times, we have 4 cases to consider, each with a number of subcases.

<u>1st weighing</u>	<u>2nd weighing</u>
(1) 2 bars on each side	2 bars on each side
(2) 2 bars on each side	1 bar on each side
(3) 1 bar on each side	2 bars on each side
(4) 1 bar on each side	1 bar on each side

Case (1): From the initial set of states $\{\pm, \pm, \pm, \pm\}$, the first weighing removes 2 +'s and 2 -'s, leaving the states $\{+, +, -, -\}$. The second weighing breaks into two subcases.

- (1.1) Place $++$ on one side and $--$ on the other, symbolized by $++ | --$. Note it does not matter whether the algorithm places $++$ on the left and $--$ on the right, or the other way around, since the daemon's strategy removes the same number of marks (namely 0) in both cases. The daemon simply tilts in the direction that confirms the existing marks, leaving the bars in states $\{+, +, -, -\}$.
- (1.2) Place $+ - | + -$. In this case we have $a = b = 2$ and $c = 0$, so the daemon removes 2 marks leaving $\{+, -, N, N\}$.

Case (2): Again the initial states $\{\pm, \pm, \pm, \pm\}$ become $\{+, +, -, -\}$ after the first weighing, giving 3 subcases for the second.

- (2.1) Place $++$. In this case $a = 1, b = 1$ and $c = 2$, so the daemon returns "balance" and removes 2 marks, leaving $\{-, -, N, N\}$.
- (2.2) Place $--$. Again the daemon removes 2 marks and leaves $\{+, +, N, N\}$.
- (2.3) Place $+ | -$. Recall the notation $+ | -$ does not specify which side has the $+$ and which has the $-$, since the same number of marks are removed. In this case the daemon removes 2 marks, leaving the states $\{+, -, N, N\}$.

Case (3): After the first weighing with 1 bar on each side, the initial states $\{\pm, \pm, \pm, \pm\}$ become $\{\pm, \pm, N, N\}$, leaving 2 subcases:

- (3.1) Place $\pm \pm | NN$. In this case $a = b = 2$ and $c = 0$, so the daemon removes 2 marks leaving either $\{+, +, N, N\}$ or $\{-, -, N, N\}$.

(3.2) Place $\pm N \mid \pm N$. Again 2 marks are removed leaving $\{+, -, N, N\}$.

Case (4): The first weighing again transforms initial states $\{\pm, \pm, \pm, \pm\}$ into $\{\pm, \pm, N, N\}$, leaving 3 subcases:

(4.1) Place $\pm \mid \pm$. The daemon removes 2 marks leaving $\{+, -, N, N\}$.

(4.2) Place $\pm \mid N$. The daemon returns "balance", and erases $\pm$ on one bar leaving $\{\pm, N, N, N\}$.

(4.3) Place $N \mid N$. The daemon again returns "balance", and removes 0 marks leaving the bars in states $\{\pm, \pm, N, N\}$.

Notice that in all cases at least two marks remain. In most cases, these marks are on different bars, but in subcase (4.2) they are on the same bar. In any case other than (4.2), if the algorithm says a particular bar is heavy or light, the daemon can produce another valid candidate as the counterfeit bar. In case (4.2), if the algorithm identifies the bar marked $\pm$ as say heavy, the daemon can claim that bar to be light. A similar contradiction arises if the algorithm identifies the $\pm$ bar as light. In all cases then, the daemon has a valid contradiction to the algorithm's verdict, and therefore, no correct algorithm can perform fewer than 3 weighings. ■

3. Recall the coin changing problem again. Given denominations $d = (d_1, d_2, \dots, d_n)$ and an amount N , 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. Prove that the greedy strategy works for any amount N with the coin system $d = (1, 5, 10, 25)$.

Proof:

Let $N \geq 0$, let $x = (x_1, x_2, x_3, x_4)$ be the optimal solution, and let $g = (g_1, g_2, g_3, g_4)$ be the solution obtained by the greedy strategy. We will show that $x = g$, whence g is also optimal. Since both solutions disburse the amount N , we have

$$(1) \quad x_1 + 5x_2 + 10x_3 + 25x_4 = N = g_1 + 5g_2 + 10g_3 + 25g_4,$$

which implies $x_1 \equiv g_1 \pmod{5}$. Observe that $0 \leq x_1 < 5$ since otherwise it would be possible to replace 5 coins of type 1 (pennies) by 1 coin of type 2 (nickel), contradicting that x is an optimal solution. Likewise $0 \leq g_1 < 5$, since the greedy strategy guarantees that as many nickels as possible will be disbursed before any pennies are. These inequalities, along with the preceding congruence, imply $x_1 = g_1$.

Cancel x_1 from (1) and divide through by 5 to get $x_2 + 2x_3 + 5x_4 = g_2 + 2g_3 + 5g_4$, which we write as

$$(2) \quad (x_2 - g_2) + 2(x_3 - g_3) + 5(x_4 - g_4) = 0.$$

The fact that x is optimal means $\sum_{i=1}^4 x_i \leq \sum_{i=1}^4 g_i$, and since $x_1 = g_1$ we have

$$(3) \quad \sum_{i=2}^4 (x_i - g_i) \leq 0$$

Note that $0 \leq x_2 < 2$ since any two nickels can be replaced by one dime, and $0 \leq g_2 < 2$ since the greedy strategy disburses the maximum number of dimes before any nickels. Therefore

$$(4) \quad -1 \leq (x_2 - g_2) \leq 1,$$

Our goal is to show $x_2 = g_2$. Assume, to get a contradiction, that $x_2 \neq g_2$, so that $x_2 - g_2 \neq 0$. Inequality (4) now implies $x_2 - g_2 = \pm 1$. Now the very first action of the greedy strategy is to pay as many quarters

as possible. Thus $g_4 \geq x_4$, and hence $g_4 - x_4 \geq 0$. If it were the case that $g_4 - x_4 = 0$, then (2) would imply $2(x_3 - g_3) = \pm 1$, which is impossible since the left side is even while the right is odd. Therefore we conclude

$$(5) \quad (g_4 - x_4) \geq 1.$$

Equation (2) implies $(x_3 - g_3) = \frac{1}{2}(g_2 - x_2) + \frac{5}{2}(g_4 - x_4)$. Combining this with inequality (3) gives

$$\begin{aligned} 0 &\geq (x_2 - g_2) + (x_3 - g_3) + (x_4 - g_4) \\ &= (x_2 - g_2) + \left[\frac{1}{2}(g_2 - x_2) + \frac{5}{2}(g_4 - x_4) \right] + (x_4 - g_4) \\ &= \frac{1}{2}(x_2 - g_2) + \frac{3}{2}(g_4 - x_4) \\ &\geq -\frac{1}{2} + \frac{3}{2} \quad (\text{by inequalities (4) and (5)}) \\ &= 1, \end{aligned}$$

i.e. $1 \leq 0$, which is absurd. Therefore our assumption was false, and hence $x_2 = g_2$.

Now equation (2) reduces to $2(x_3 - g_3) + 5(x_4 - g_4) = 0$, which we write as

$$(6) \quad x_3 - g_3 = \frac{5}{2}(g_4 - x_4).$$

Inequality (3), together with $x_2 = g_2$, and equation (6) implies

$$\begin{aligned} 0 &\geq (x_3 - g_3) + (x_4 - g_4) \\ &= \frac{5}{2}(g_4 - x_4) + (x_4 - g_4) \\ &= \frac{3}{2}(g_4 - x_4). \end{aligned}$$

Therefore $g_4 - x_4 \leq 0$, so that $g_4 \leq x_4$. But as noted previously $g_4 \geq x_4$ (no solution disburses more quarters than the greedy solution.) Hence $g_4 = x_4$, and equation (6) now gives $g_3 = x_3$. This completes the proof that $x = g$. We conclude that the greedy strategy is optimal for this denomination set. ■

4. **Scheduling to Minimize Average Completion Time:** (This is problem 16-2a on page 402 of CLRS.) Suppose you are given a set $S = \{a_1, a_2, \dots, a_n\}$ of tasks, where task a_i requires p_i units of processing time to complete, once it has started. You have one computer on which to run these tasks, and the computer can run only one task at a time. Let c_i be the *completion time* of task a_i , that is, the time at which task a_i completes processing. Your goal is to minimize the average completion time, that is to minimize the quantity $(1/n) \sum_{i=1}^n c_i$. For example, suppose there are two tasks, a_1 and a_2 , with $p_1 = 3$ and $p_2 = 5$. Consider the schedule in which a_2 runs first, followed by a_1 . Then $c_2 = 5$, $c_1 = 8$, and the average completion time is $(5 + 8)/2 = 6.5$. Another schedule runs a_1 first, followed by a_2 . In that case $c_1 = 3$, $c_2 = 8$, and the average completion time is $(3 + 8)/2 = 5.5$. Thus the schedule (a_1, a_2) has a better average completion time as compared to (a_2, a_1) . Observe that a schedule is nothing more than a permutation of the set $S = \{a_1, a_2, \dots, a_n\}$ of tasks.

Give a greedy algorithm that schedules the tasks so as to minimize the average completion time. Each task must run non-preemptively, that is, once task a_i is started, it runs continuously for p_i units of time. Prove that your algorithm minimizes the average completion time, and state its running time.

Solution:

We adopt the following greedy strategy: sort S by increasing processing times p_i , schedule those tasks with shortest processing times first, and those with longer processing times later. To put it another way, determine a permutation of the indices $(1, 2, \dots, n)$ which sorts the array of processing times $(p_1, \dots, p_n)$, then apply that same permutation to the array of tasks $(a_1, \dots, a_n)$ to obtain an optimal schedule. The running time of this algorithm is obviously the running time of the sort used, which can be chosen to be $\theta(n \log(n))$ in worst and average cases.

Theorem

The above greedy strategy results in a schedule which minimizes the average completion time of the tasks $a_1, a_2, \dots, a_n$.

Proof:

To simplify notation, let us assume from the start that the tasks $a_1, a_2, \dots, a_n$ are already sorted by increasing processing times, i.e we assume that $p_1 \leq p_2 \leq \dots \leq p_n$. Let us also assume that processing begins at time 0, so that the completion times are:

$$\begin{aligned} c_1 &= p_1 \\ c_2 &= p_1 + p_2 \\ &\vdots \\ c_n &= p_1 + p_2 + \dots + p_n \end{aligned}$$

and in general $c_k = \sum_{i=1}^k p_i$. Thus the average completion time resulting from the above strategy is

$$A = (1/n) \sum_{k=1}^n c_k = (1/n) \sum_{k=1}^n \sum_{i=1}^k p_i = (1/n) \sum_{k=1}^n (n - k + 1) p_k$$

More Generally if σ is any permutation of the set $\{1, 2, \dots, n\}$, then the average completion time of the schedule $a_{\sigma(1)}, a_{\sigma(2)}, \dots, a_{\sigma(n)}$ is given by $A(\sigma) = (1/n) \sum_{k=1}^n (n - k + 1) p_{\sigma(k)}$. Given such a permutation σ , we must show that $A = A(\text{identity}) \leq A(\sigma)$. Let i be the least index such that $\sigma(i) \neq i$. Since i is the smallest such index, we must have $i < \sigma(i)$. Let $j = \sigma(i)$ and $l = \sigma^{-1}(i)$, so that $\sigma(l) = i$. Let τ be the permutation obtained from σ by swapping i with j . In other words

$$\tau(k) = \begin{cases} i & \text{for } k = i \\ j & \text{for } k = l \\ \sigma(k) & \text{for all other } k \end{cases}$$

Thus

$$\begin{aligned} A(\sigma) - A(\tau) &= (1/n) \sum_{k=1}^n (n - k + 1) p_{\sigma(k)} - (1/n) \sum_{k=1}^n (n - k + 1) p_{\tau(k)} \\ &= (1/n) [(n - i + 1) p_j + (n - l + 1) p_i - (n - i + 1) p_i - (n - l + 1) p_j] \\ &= (1/n) [(l - i) p_j + (i - l) p_i] \\ &= (1/n) (l - i) (p_j - p_i) \end{aligned}$$

Now recall $i < \sigma(i) = j$ whence $p_i \leq p_j$, and $\sigma(l) = i \neq l$ so that $i < l$. Hence $A(\sigma) - A(\tau) \geq 0$, and therefore $A(\sigma) \geq A(\tau)$. In other words the schedule $a_{\tau(1)}, a_{\tau(2)}, \dots, a_{\tau(n)}$ has a lower average completion time than does the schedule $a_{\sigma(1)}, a_{\sigma(2)}, \dots, a_{\sigma(n)}$. But the permutation τ has one more term in common with the identity permutation than σ does. That is $\tau(k) = k$ for $1 \leq k \leq i$, while $\sigma(k) = k$ for $1 \leq k < i$

and $\sigma(i) \neq i$. If τ equals the identity permutation (i.e. $\tau(k) = k$ for all k), then we are done. If not, then repeat the above procedure with τ in place of σ to obtain a permutation τ_1 with one more term in common with the identity permutation than τ , and giving a schedule with lower average completion time. Continuing in this fashion we can construct a sequence of permutations $\sigma, \tau, \tau_1, \tau_2, \dots$ ending in the identity permutation giving schedules with descending average completion times: $A(\sigma) \geq A(\tau) \geq A(\tau_1) \geq A(\tau_2) \geq \dots \geq A(\text{identity}) = A$. Thus $A(\sigma) \geq A$ as required. ■