CSE 201
Final Exam Review
Solutions to Selected Problems

1. Suppose $T(n)$ satisfies the recurrence $T(n) = 3T(n/4) + F(n)$, where $F(n)$ itself satisfies the recurrence $F(n) = 5F(n/9) + n^{3/4}$. Find a tight asymptotic bound for $T(n)$. Be sure to fully justify each use of the Master Theorem. (Hint: $\log_9(5) < 3/4 < \log_4(3)$.)

   Solution:
   First observe

   $$5^2 = 25 < 27 = 3^3 \;\Rightarrow\; 5 < 3^{3/2} = 9^{3/4} \;\Rightarrow\; \log_9(5) < 3/4$$

   and

   $$2^3 = 8 < 9 = 3^2 \;\Rightarrow\; 4^{3/4} = 2^{3/2} < 3 \;\Rightarrow\; 3/4 < \log_4(3).$$

   To determine $F(n)$, let $\epsilon = \frac{3}{4} - \log_9(5)$. Then $\epsilon > 0$ and $\frac{3}{4} = \log_9(5) + \epsilon$, giving $n^{3/4} = \Omega\left(n^{3/4}\right) = \Omega\left(n^{\log_9(5)+\epsilon}\right)$. Also, for any c in the range $\frac{5}{9^{3/4}} \le c < 1$ we have

   $$5\left(\frac{n}{9}\right)^{3/4} = \frac{5}{9^{3/4}} \cdot n^{3/4} \le cn^{3/4}$$

   for all $n \ge 1$, establishing the regularity condition. Case 3 yields $F(n) = \Theta\left(n^{3/4}\right)$.

   We may now write the recurrence for $T(n)$ as $T(n) = 3T(n/4) + n^{3/4}$. Let $\epsilon = \log_4(3) - \frac{3}{4}$. It follows that $\epsilon > 0$ and $\frac{3}{4} = \log_4(3) - \epsilon$, whence $n^{3/4} = O\left(n^{3/4}\right) = O\left(n^{\log_4(3)-\epsilon}\right)$. Case 1 of the Master Theorem gives us $T(n) = \Theta\left(n^{\log_4(3)}\right)$. ■

2. Recall that $\{0,1\}^*$ denotes the set of all bit-strings of any finite length. A language L is a collection of bit-strings, i.e. a subset $L \subseteq \{0,1\}^*$. Let $A(x)$ be an algorithm whose input is a bit-string $x \in \{0,1\}^*$, and whose output is 0 or 1.

   a. Define what it means for a language L to be *accepted* by A.
      Solution:
      Algorithm A *accepts* L if and only if $L = \{x \in \{0,1\}^* \mid A(x) = 1\}$. (Note $x \in \{0,1\}^* - L$ does not imply that $A(x) = 0$, since it is possible that $A(x)$ does not halt.)

   b. Define what it means for a language L to be *decided* by A.
      Solution:
      Algorithm A *decides* the language L if and only if $A(x) = 1$ for all $x \in L$, and $A(x) = 0$ for all $x \in \{0,1\}^* - L$. In particular, A must always halt. Note also: $A\ decides\ L \;\Rightarrow\; A\ accepts\ L$ but the converse is false: $A\ accepts\ L \;\not\Rightarrow\; A\ decides\ L$.

4. Recall the decision problems Hamiltonian Cycle (HC) and Travelling Salesman (TSP).

   HC: Given a graph *G*, determine whether or not *G* contains a Hamiltonian cycle (a cycle that visits every vertex in *G*).
   TSP: Given a complete graph K_n, a weight function $d: E(K_n) \to \mathbb{R}$, and a bound $b \geq 0$, determine whether or not K_n contains a Hamiltonian cycle of total weight no more than *b*.

   Recall also the mapping f:HC → TSP that takes instances of HC to instances of TSP, defined as follows. Given a graph *G* with $|V(G)| = n$, identify $V(G)$ with $V(K_n)$, define $d: E(K_n) \to \mathbb{R}$ by

$$d(u,v) = \begin{cases} 1 & \text{if } \{u,v\} \in E(G) \\ 2 & \text{if } \{u,v\} \notin E(G), \end{cases}$$

   and let $b = n$.
   a. Prove that if G is a Yes instance of HC, then $f(G)$ is a Yes instance of TSP.
   b. Prove that if $f(G)$ is a Yes instance of TSP, then G is a Yes instance of HC.
   c. Explain how $f(G)$ can be computed in polynomial time. (Make some assumption as to how G will be represented, such as adjacency-list, adjacency-matrix, or incidence-matrix.)

   Solution to (a) and (b): See pages 3 and 4 of the handout on NP-completeness.

   Solution to (c):
   Let $n = |V(G)|$ and $m = |E(G)|$. Label the vertices of G (and of K_n) using the positive integers from 1 to n, so $V(G) = V(K_n) = \{1, 2, \ldots, n\}$. Represent G by a 1-dimensional array of adjacency lists $\text{adj}[1 \cdots n]$, and represent the distance function d by a 2-dimensional array of integers $d[1 \cdots n; 1 \cdots n]$. Begin by initializing $d[u,v] = 2$ for all $1 \leq u \leq n$ and $1 \leq v \leq n$, which takes $\Theta(n^2)$ time. Then, for each u in the range $1 \leq u \leq n$ and each $v \in \text{adj}[u]$, set $d[u,v] = 1$. Since the total length of all of the adjacency lists is $2m$, and $m \leq n^2$, this step takes time $O(n^2)$. Return the triple (n, d, n) at total cost $\Theta(n^2)$. ■

5. 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.

   a. Use a decision tree argument to prove that at least 2 weighings must be performed (in worst case) by any algorithm that solves this problem. Carefully enumerate the set of possible verdicts.

      Proof:
      Any algorithm for this problem can be represented by a ternary decision tree. Since there are 9 possible verdicts: {1L, 2L, 3L, 4L, 1H, 2H, 3H, 4H, AllGold}, at least $\lceil \log_3(9) \rceil = 2$ weighings are necessary. ■

   b. Determine an algorithm that solves this problem using 3 weighings (in worst case). Express your algorithm as a decision tree.

      Solution:
      As usual a left branch indicates that the scale tilted left, right branch means tilt right, and middle branch means the scale balanced. There are of course many correct solutions, one of which is pictured below.

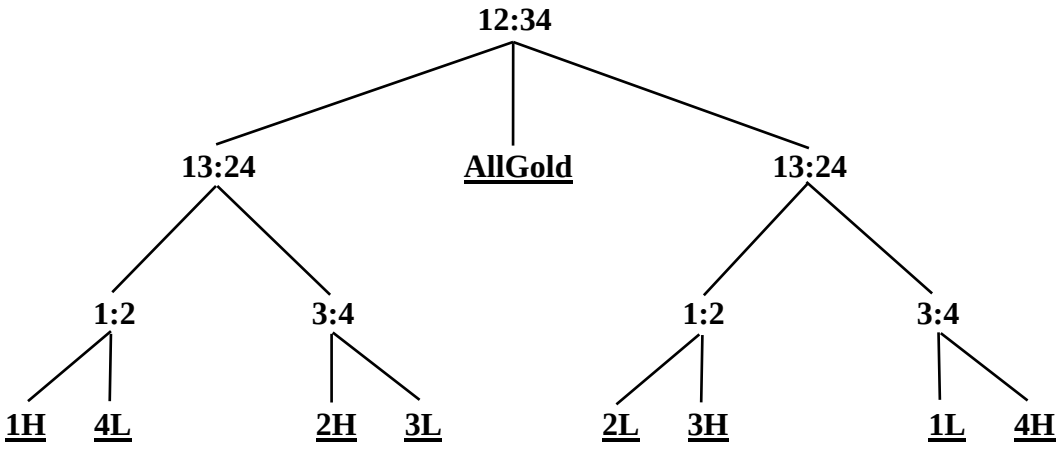

Note that although this is a ternary tree, only one node (the root) actually has 3 children. The question naturally arises as to whether this can be done with just 2 weighings, as the decision tree lower bound in (a) would suggest. Part (c) below shows this is not possible. ■

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 class 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

$$L_+ = \#\text{ of } + \text{ marks on the left side of the scale}$$
$$L_- = \#\text{ of } - \text{ marks on the left side of the scale}$$

$R_+ =$ # 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 $a+b \le a+c$ and $a+b \le b+c$ are true if and only if $b \le c$ and $a \le c$, which is equivalent to $c = \max(a,b,c)$. Similarly, $a+c \le a+b$ and $a+c \le b+c$ are true iff $b = \max(a,b,c)$, and also $b+c \le a+b$ and $b+c \le a+c$ 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.

	1st weighing	2nd weighing
(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 $\pm N\,|\,\pm N$. 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 $\pm\,|\,\pm$. The daemon removes 2 marks leaving $\{+,-,N,N\}$.
(4.2) Place $\pm\,|\,N$. The daemon returns "balance", and erases $\pm$ on one bar leaving $\{\pm,N,N,N\}$.
(4.3) Place $N\,|\,N$. 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. ■

Remark The case by case analysis above shows that it *is* possible to identify the counterfeit bar in just 2 weighings by designing an algorithm that utilizes case (4.2). To decide if the counterfeit is heavy or light still takes one more weighing though. The algorithm represented by the decision tree below utilizes this strategy.

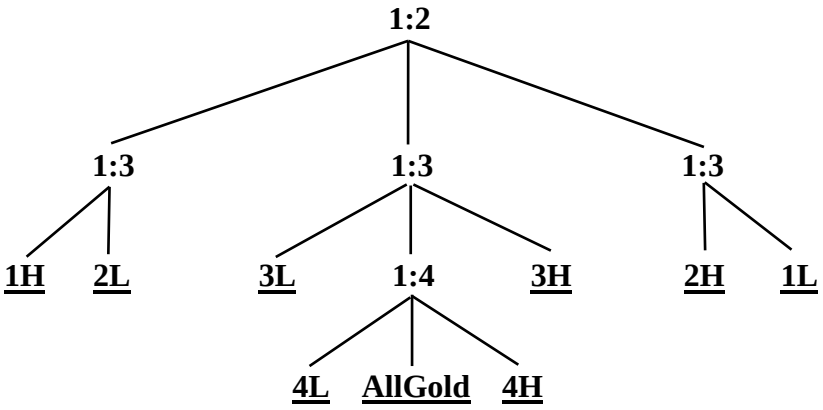

Which algorithm one chooses may depend on prior beliefs as to which bars are good or bad. For instance, if you think it is highly likely that a certain bar is genuine, then calling that bar number 4, the above algorithm will produce a good average case run time. If you believe it likely that there is no counterfeit bar, then the solution given in (b) is preferred. The following exercise should be the last word on the bar weighing problem.

Additional Exercise
Suppose we are given $n \geq 3$ gold bars, where again one bar may be counterfeit, light or heavy. Again we have a balance scale whose use produces one of three outcomes. Since there are $2n + 1$ possible verdicts, we immediately have the decision tree lower bound of $\lceil \log_3(2n+1) \rceil$ for the minimum number of weighings to find the counterfeit and its type (light or heavy), or show that there is no counterfeit.

a. Give an adversary argument showing that $\lceil \log_3(2n+3) \rceil$ weighings are necessary in general. (Observe that the two lower bounds are different only when $n = \frac{1}{2}(3^k - 1)$ for some $k \geq 2$.)
b. Design an algorithm that solves this problem using the scale at most $\lceil \log_3(2n+3) \rceil$ times.

6. (This is Problem 34.1-6 page 1061 of CLRS, see pages 1057-58 for definitions.) Show that the class *P*, viewed as a set of languages, is closed under union, intersection, concatenation, complement, and Kleene star. That is, if $L_1, L_2 \in P$, then $L_1 \cup L_2 \in P, L_1 \cap L_2 \in P, L_1 L_2 \in P, \overline{L_1} \in P$, and $L_1^* \in P$.

Proof:
Given $L_1, L_2 \in P$, there exist polynomial time algorithms A_1 and A_2 that decide the languages L_1 and L_2, respectively. (This is the definition of *P* on page 1059 of CLRS.) Specifically, there exist numbers k_1 and k_2 such that for any $x \in \{0,1\}^*$, $A_1(x)$ halts in time $O(n^{k_1})$ (where $n = |x|$) returning "yes" if $x \in L_1$, and "no" if $x \notin L_1$. Similarly, $A_2(x)$ halts in time $O(n^{k_2})$ and decides whether $x \in L_2$.

We must show there exist algorithms that decide each of the languages $L_1 \cup L_2$, $L_1 \cap L_2$, $L_1 L_2$, $\overline{L_1}$ and L_1^* in polynomial time. In fact, Theorem 34.2 (bottom of page 1059) says it is sufficient to find polynomial time algorithms that merely accept these languages. Recall that an algorithm A accepts L in polynomial time if whenever $x \in L$, $A(x)$ halts and returns "yes" in time $O(n^k)$. If $x \notin L$ then $A(x)$ may halt and return 0, or it may not halt.

$L_1 \cup L_2$:
Given $x \in \{0,1\}^*$, compute both $A_1(x)$ and $A_2(x)$. If either algorithm returns yes, then return "yes" for $x \in L_1 \cup L_2$, otherwise return "no". The runtime is $O(n^{k_1}) + O(n^{k_2}) = O\left(n^{\max(k_1,k_2)}\right)$, where as before $n = |x|$, the length of the input string x.

$L_1 \cap L_2$:
Given $x \in \{0,1\}^*$, compute both $A_1(x)$ and $A_2(x)$. If both algorithms returns yes, then return "yes" for $x \in L_1 \cap L_2$, otherwise return "no". Again the runtime is $O(n^{k_1}) + O(n^{k_2}) = O\left(n^{\max(k_1,k_2)}\right)$.

$L_1 L_2$:
Given $x \in \{0,1\}^*$, iterate over all $n+1$ factorizations $x = yz$ into a concatenation of two substrings *y* and *z*. For each such factorization, compute both $A_1(y)$ and $A_2(z)$. If on any iteration, both computations return yes, halt and return "yes" for $x \in L_1 L_2$. If the end of the loop is reached without returning yes, then return "no". Since $|x| = |y| + |z|$, we must have $|y| = j$ and $|z| = n - j$ for some integer *j* in the range $0 \leq j \leq n$. Iteration *j* runs in time $O(j^{k_1}) + O((n-j)^{k_2})$, so the whole procedure runs in time

$$\begin{aligned}
&\sum_{j=0}^{n}\left(O(j^{k_1}) + O((n-j)^{k_2})\right) \\
&\quad = \sum_{j=0}^{n} O(j^{k_1}) + \sum_{j=0}^{n} O((n-j)^{k_2}) \\
&\quad = O\left(\sum_{j=0}^{n} j^{k_1}\right) + O\left(\sum_{j=0}^{n} j^{k_2}\right) \\
&\quad = O(n^{k_1+1}) + O(n^{k_2+1}) \qquad \text{(Ex 4 on p.4 of asymptotic handout)} \\
&\quad = O\left(n^{\max(k_1+1,k_2+1)}\right) \\
&\quad = O\left(n^{\max(k_1,k_2)+1}\right).
\end{aligned}$$

$\overline{L_1}$:
Given $x \in \{0,1\}^*$, compute $A_1(x)$, and if yes is returned, then return "no" for $\overline{L_1}$. If no is returned, then return "yes" for $\overline{L_1}$. This algorithm obviously runs in time $O(n^{k_1})$.

L_1^*:
To simplify notation, we write L for the language L_1. Since $L \in P$, there exists and algorithm A That decides $x \in L^*$, and halts in time $O(n^k)$, where $n = |x|$. Our goal is to write an algorithm Kleene(x) that decides $x \in L^* = \bigcup_{i=0}^{\infty} L^i$ in polynomial time. Since this algorithm is somewhat complicated, we express it in pseudo-code, rather than using a verbal description as above. We will write $x[i{:}j]$ to denote the substring of x from index i to index j, inclusive. Indices range from 1 to n, so $x = x[1]x[2]\cdots x[n] = x[1{:}n]$. Our goal is to fill a 2-dimensional table $T[i,j]$, where $1 \le i \le j \le n$, with values true an false, indicating exactly which substrings of x belong to L^*. Specifically, we assign $T[i,j] =$ true iff $x[i{:}j] \in L^*$. Once this is done, we return the Boolean value $T[1,n]$ as the answer to $x \in L^*$.

```
Kleene(x)   pre: A(x) decides x ∈ L
1.  n = |x|
2.  if n == 0
3.      return true
4.  for l = 1 to n
5.      for i = 1 to n − l + 1
6.          j = i + l − 1
7.          if A(x[i:j])
8.              T[i,j] = true
9.          else
10.             for m = i to j − 1
11.                 if T[i,m] and T[m + 1,j]
12.                     T[i,j] = true
13.                 else
14.                     T[i,j] = false
15. return T[1,n]
```

The correctness of this algorithm relies on the fact that $L^* = \{\epsilon\} \cup L \cup L^*L^*$, which is readily verified. Thus every $x \in L^*$ satisfies (at least) one of the following conditions:

(1) $x = \epsilon$,
(2) $x \in L$, or
(3) $x = yz$ for some $y, z \in L^*$.

Line 3 returns true if $x = \epsilon$, in which case (1) holds and we need not consider any substrings of x. Line 8 sets $T[i,j] =$ true iff the substring $x[i{:}j] \in L$, which is condition (2). Loop 10-14 checks condition (3), setting $T[i,j] =$ true iff $x[i{:}j] = x[i{:}m]x[m+1{:}j]$, where the substrings $x[i{:}m]$ and $x[m+1{:}j]$ were previously found to belong to L^*.

Note that the above algorithm is a classic instance of Dynamic Programming, similar to the Matrix Chain Multiplication problem. It remains only to show the algorithm runs in polynomial time. Observe the call to $A(x[i{:}j])$ is inside two nested for loops, each of which executes at most n times. Since $x[i{:}j]$ is a substring of x, $A(x[i{:}j])$ runs in time $O(n^k)$. Also, within the double loop is another for loop that executes at most n times. Therefore Kleene(x) runs in time

$$O\Big(n^2(O(n^k)+n)\Big) \subseteq O(n^{k+2}),$$

as required. ■

7. Recall the coin changing problem again. Given denominations $d = (d_1, d_2, \ldots, d_n)$ 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) $$x_1 + 5x_2 + 10x_3 + 25x_4 = N = g_1 + 5g_2 + 10g_3 + 25g_4,$$

which implies $x_1 \equiv g_1 (mod\ 5)$. 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) $$(x_2 - g_2) + 2(x_3 - g_3) + 5(x_4 - g_4) = 0.$$

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

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

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

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

$$\begin{aligned} 0 &\geq (x_2 - g_2) + (x_3 - g_3) + (x_4 - g_4) \\ &= (x_2 - g_2) + \left[\tfrac{1}{2}(g_2 - x_2) + \tfrac{5}{2}(g_4 - x_4)\right] + (x_4 - g_4) \end{aligned}$$

$$= \frac{1}{2}(x_2 - g_2) + \frac{3}{2}(g_4 - x_4)$$
$$\geq -\frac{1}{2} + \frac{3}{2} \qquad \text{(by inequalities (4) and (5))}$$
$$= 1,$$

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

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

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

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. ■

8. **Scheduling to Minimize Average Completion Time:** (This is problem 16-2a on page 402 of CLRS.) Suppose you are given a set $S = \{a_1, a_2, \ldots, a_n\}$ 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 $p_2 = 5$, and 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 $(5 + 8)/2 = 6.5$.

   Give an 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 must run continuously for p_i units of time. Prove that your algorithm minimizes the average completion time, and state the running time of your algorithm.

   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,…, *n*) which sorts the array of processing times $(p_1, \ldots, p_n)$, then apply that same permutation to the array of tasks $(a_1, \ldots, a_n)$ 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 $a_1, a_2, \ldots, a_n$.

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

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

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, …, *n*}, then the average completion time of the schedule $a_{\sigma(1)}, a_{\sigma(2)}, \ldots, a_{\sigma(n)}$ 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 $A = A(\text{identity}) \le A(\sigma)$. Let *i* be the least index such that $\sigma(i) \ne i$. 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 $p_i \le p_j$, and $\sigma(l) = i \ne l$ so that $i < l$. Hence $A(\sigma) - A(\tau) \ge 0$, and therefore $A(\sigma) \ge A(\tau)$. In other words the schedule $a_{\tau(1)}, a_{\tau(2)}, \ldots, a_{\tau(n)}$ has a lower average completion time than does the schedule $a_{\sigma(1)}, a_{\sigma(2)}, \ldots, a_{\sigma(n)}$. But the permutation τ has one more term in common with the identity permutation than σ does. That is $\tau(k) = k$ for $1 \le k \le i$, while $\sigma(k) = k$ for $1 \le k < i$ and $\sigma(i) \ne i$. 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 $\sigma, \tau, \tau_1, \tau_2, \ldots$ending in the identity permutation giving schedules with descending average completion times: $A(\sigma) \ge A(\tau) \ge A(\tau_1) \ge A(\tau_2) \ge \cdots \ge A(\text{identity}) = A$. Thus $A(\sigma) \ge A$ as required. ■