

# CS 373: Combinatorial Algorithms, Spring 2001

## Homework 1 (due Thursday, February 1, 2001 at 11:59:59 p.m.)

Name:		
Net ID:	Alias:	U <sup>3</sup> / <sub>4</sub> 1

Name:		
Net ID:	Alias:	U <sup>3</sup> / <sub>4</sub> 1

Name:		
Net ID:	Alias:	U <sup>3</sup> / <sub>4</sub> 1

---

Starting with Homework 1, homeworks may be done in teams of up to three people. Each team turns in just one solution, and every member of a team gets the same grade. Since 1-unit graduate students are required to solve problems that are worth extra credit for other students, **1-unit grad students may not be on the same team as 3/4-unit grad students or undergraduates.**

Neatly print your name(s), NetID(s), and the alias(es) you used for Homework 0 in the boxes above. Please also tell us whether you are an undergraduate, 3/4-unit grad student, or 1-unit grad student by circling U, <sup>3</sup>/<sub>4</sub>, or 1, respectively. Staple this sheet to the top of your homework.

---

### Required Problems

1. Suppose you are a simple shopkeeper living in a country with  $n$  different types of coins, with values  $1 = c[1] < c[2] < \dots < c[n]$ . (In the U.S., for example,  $n = 6$  and the values are 1, 5, 10, 25, 50 and 100 cents.) Your beloved and belevolent dictator, El Generalissimo, has decreed that whenever you give a customer change, you must use the smallest possible number of coins, so as not to wear out the image of El Generalissimo lovingly engraved on each coin by servants of the Royal Treasury.
  - (a) In the United States, there is a simple greedy algorithm that always results in the smallest number of coins: subtract the largest coin and recursively give change for the remainder. El Generalissimo does not approve of American capitalist greed. Show that there is a set of coin values for which the greedy algorithm does *not* always give the smallest possible of coins.
  - (b) Describe and analyze a dynamic programming algorithm to determine, given a target amount  $A$  and a sorted array  $c[1..n]$  of coin values, the smallest number of coins needed to make  $A$  cents in change. You can assume that  $c[1] = 1$ , so that it is possible to make change for any amount  $A$ .

2. Consider the following sorting algorithm:

```

STUPIDSORT( $A[0..n-1]$ ) :
  if  $n = 2$  and  $A[0] > A[1]$ 
    swap  $A[0] \leftrightarrow A[1]$ 
  else if  $n > 2$ 
     $m \leftarrow \lceil 2n/3 \rceil$ 
    STUPIDSORT( $A[0..m-1]$ )
    STUPIDSORT( $A[n-m..n-1]$ )
    STUPIDSORT( $A[0..m-1]$ )

```

- (a) Prove that STUPIDSORT actually sorts its input.
- (b) Would the algorithm still sort correctly if we replaced the line  $m \leftarrow \lceil 2n/3 \rceil$  with  $m \leftarrow \lfloor 2n/3 \rfloor$ ? Justify your answer.
- (c) State a recurrence (including the base case(s)) for the number of comparisons executed by STUPIDSORT.
- (d) Solve the recurrence, and prove that your solution is correct. [Hint: Ignore the ceiling.] Does the algorithm deserve its name?
- \*(e) Show that the number of *swaps* executed by STUPIDSORT is at most  $\binom{n}{2}$ .
3. The following randomized algorithm selects the  $r$ th smallest element in an unsorted array  $A[1..n]$ . For example, to find the smallest element, you would call RANDOMSELECT( $A, 1$ ); to find the median element, you would call RANDOMSELECT( $A, \lfloor n/2 \rfloor$ ). Recall from lecture that PARTITION splits the array into three parts by comparing the pivot element  $A[p]$  to every other element of the array, using  $n - 1$  comparisons altogether, and returns the new index of the pivot element.

```

RANDOMSELECT( $A[1..n], r$ ) :
   $p \leftarrow \text{RANDOM}(1, n)$ 
   $k \leftarrow \text{PARTITION}(A[1..n], p)$ 
  if  $r < k$ 
    return RANDOMSELECT( $A[1..k-1], r$ )
  else if  $r > k$ 
    return RANDOMSELECT( $A[k+1..n], r-k$ )
  else
    return  $A[k]$ 

```

- (a) State a recurrence for the expected running time of RANDOMSELECT, as a function of  $n$  and  $r$ .
- (b) What is the *exact* probability that RANDOMSELECT compares the  $i$ th smallest and  $j$ th smallest elements in the input array? The correct answer is a simple function of  $i$ ,  $j$ , and  $r$ . [Hint: Check your answer by trying a few small examples.]
- (c) Show that for any  $n$  and  $r$ , the expected running time of RANDOMSELECT is  $\Theta(n)$ . You can use either the recurrence from part (a) or the probabilities from part (b). For extra credit, find the exact expected number of comparisons, as a function of  $n$  and  $r$ .
- (d) What is the expected number of times that RANDOMSELECT calls itself recursively?

4. What excitement! The Champaign Spinners and the Urbana Dreamweavers have advanced to meet each other in the World Series of Basketweaving! The World Champions will be decided by a best-of- $2n-1$  series of head-to-head weaving matches, and the first to win  $n$  matches will take home the coveted Golden Basket (for example, a best-of-7 series requiring four match wins, but we will keep the generalized case). We know that for any given match there is a constant probability  $p$  that Champaign will win, and a subsequent probability  $q = 1 - p$  that Urbana will win.

Let  $P(i, j)$  be the probability that Champaign will win the series given that they still need  $i$  more victories, whereas Urbana needs  $j$  more victories for the championship.  $P(0, j) = 1$ ,  $1 \leq j \leq n$ , because Champaign needs no more victories to win.  $P(i, 0) = 0$ ,  $1 \leq i \leq n$ , as Champaign cannot possibly win if Urbana already has.  $P(0, 0)$  is meaningless. Champaign wins any particular match with probability  $p$  and loses with probability  $q$ , so

$$P(i, j) = p \cdot P(i - 1, j) + q \cdot P(i, j - 1)$$

for any  $i \geq 1$  and  $j \geq 1$ .

Create and analyze an  $O(n^2)$ -time dynamic programming algorithm that takes the parameters  $n$ ,  $p$  and  $q$  and returns the probability that Champaign will win the series (that is, calculate  $P(n, n)$ ).

5. The traditional Devonian/Cornish drinking song “The Barley Mow” has the following pseudolyrics<sup>1</sup>, where  $container[i]$  is the name of a container that holds  $2^i$  ounces of beer.<sup>2</sup>

```

BARLEYMOW( $n$ ):
  “Here’s a health to the barley-mow, my brave boys,”
  “Here’s a health to the barley-mow!”

  “We’ll drink it out of the jolly brown bowl,”
  “Here’s a health to the barley-mow!”
  “Here’s a health to the barley-mow, my brave boys,”
  “Here’s a health to the barley-mow!”

  for  $i \leftarrow 1$  to  $n$ 
    “We’ll drink it out of the  $container[i]$ , boys,”
    “Here’s a health to the barley-mow!”
    for  $j \leftarrow i$  downto 1
      “The  $container[j]$ ,”
      “And the jolly brown bowl!”
      “Here’s a health to the barley-mow!”
    “Here’s a health to the barley-mow, my brave boys,”
    “Here’s a health to the barley-mow!”

```

- (a) Suppose each container name  $container[i]$  is a single word, and you can sing four words a second. How long would it take you to sing BARLEYMOW( $n$ )? (Give a tight asymptotic bound.)
- (b) If you want to sing this song for  $n > 20$ , you’ll have to make up your own container names, and to avoid repetition, these names will get progressively longer as  $n$  increases<sup>3</sup>. Suppose  $container[n]$  has  $\Theta(\log n)$  syllables, and you can sing six syllables per second. Now how long would it take you to sing BARLEYMOW( $n$ )? (Give a tight asymptotic bound.)
- (c) Suppose each time you mention the name of a container, you drink the corresponding amount of beer: one ounce for the jolly brown bowl, and  $2^i$  ounces for each  $container[i]$ . Assuming for purposes of this problem that you are at least 21 years old, *exactly* how many ounces of beer would you drink if you sang BARLEYMOW( $n$ )? (Give an *exact* answer, not just an asymptotic bound.)

<sup>1</sup>Pseudolyrics are to lyrics as pseudocode is to code.

<sup>2</sup>One version of the song uses the following containers: nipperkin, gill pot, half-pint, pint, quart, pottle, gallon, half-anker, anker, firkin, half-barrel, barrel, hogshead, pipe, well, river, and ocean. Every container in this list is twice as big as its predecessor, except that a firkin is actually 2.25 ankers, and the last three units are just silly.

<sup>3</sup>“We’ll drink it out of the hemisemidemiyottapint, boys!”

6. [This problem is required only for graduate students taking CS 373 for a full unit; anyone else can submit a solution for extra credit.]


Suppose we want to display a paragraph of text on a computer screen. The text consists of  $n$  words, where the  $i$ th word is  $p_i$  pixels wide. We want to break the paragraph into several lines, each exactly  $P$  pixels long. Depending on which words we put on each line, we will need to insert different amounts of white space between the words. The paragraph should be fully justified, meaning that the first word on each line starts at its leftmost pixel, and *except for the last line*, the last character on each line ends at its rightmost pixel. There must be at least one pixel of whitespace between any two words on the same line.

Define the *slop* of a paragraph layout as the sum over all lines, *except the last*, of the cube of the number of extra white-space pixels in each line (not counting the one pixel required between every adjacent pair of words). Specifically, if a line contains words  $i$  through  $j$ , then the amount of extra white space on that line is  $P - j + i - \sum_{k=i}^j p_k$ . Describe a dynamic programming algorithm to print the paragraph with minimum slop.

## Practice Problems


1. Give an  $O(n^2)$  algorithm to find the longest increasing subsequence of a sequence of numbers. The elements of the subsequence need not be adjacent in the sequence. For example, the sequence  $\langle 1, 5, 3, 2, 4 \rangle$  has longest increasing subsequence  $\langle 1, 3, 4 \rangle$ .
2. You are at a political convention with  $n$  delegates. Each delegate is a member of exactly one political party. It is impossible to tell which political party a delegate belongs to. However, you can check whether any two delegates are in the *same* party or not by introducing them to each other. (Members of the same party always greet each other with smiles and friendly handshakes; members of different parties always greet each other with angry stares and insults.)
  - (a) Suppose a majority (more than half) of the delegates are from the same political party. Give an efficient algorithm that identifies a member of the majority party.
  - (b) Suppose exactly  $k$  political parties are represented at the convention and one party has a *plurality*: more delegates belong to that party than to any other. Present a practical procedure to pick a person from the plurality party as parsimoniously as possible. (Please.)
3. Give an algorithm that finds the *second* smallest of  $n$  elements in at most  $n + \lceil \lg n \rceil - 2$  comparisons. [Hint: divide and conquer to find the smallest; where is the second smallest?]
4. Some graphics hardware includes support for an operation called *blit*, or **block transfer**, which quickly copies a rectangular chunk of a pixelmap (a two-dimensional array of pixel values) from one location to another. This is a two-dimensional version of the standard C library function `memcpy()`.

Suppose we want to rotate an  $n \times n$  pixelmap  $90^\circ$  clockwise. One way to do this is to split the pixelmap into four  $n/2 \times n/2$  blocks, move each block to its proper position using a sequence of five blits, and then recursively rotate each block. Alternately, we can first recursively rotate the blocks and blit them into place afterwards.



Two algorithms for rotating a pixelmap.  
 Black arrows indicate blitting the blocks into place.  
 White arrows indicate recursively rotating the blocks.

The following sequence of pictures shows the first algorithm (blit then recurse) in action.



In the following questions, assume  $n$  is a power of two.

- (a) Prove that both versions of the algorithm are correct. [Hint: If you exploit all the available symmetries, your proof will only be a half of a page long.]
  - (b) *Exactly* how many blits does the algorithm perform?
  - (c) What is the algorithm's running time if a  $k \times k$  blit takes  $O(k^2)$  time?
  - (d) What if a  $k \times k$  blit takes only  $O(k)$  time?
5. A company is planning a party for its employees. The employees in the company are organized into a strict hierarchy, that is, a tree with the company president at the root. The organizers of the party have assigned a real number to each employee measuring how 'fun' the employee is. In order to keep things social, there is one restriction on the guest list: an employee cannot attend the party if their immediate supervisor is present. On the other hand, the president of the company *must* attend the party, even though she has a negative fun rating; it's her company, after all. Give an algorithm that makes a guest list for the party that maximizes the sum of the 'fun' ratings of the guests.
  6. Suppose you have a subroutine that can find the median of a set of  $n$  items (*i.e.*, the  $\lfloor n/2 \rfloor$  smallest) in  $O(n)$  time. Give an algorithm to find the  $k$ th biggest element (for arbitrary  $k$ ) in  $O(n)$  time.
  7. You're walking along the beach and you stub your toe on something in the sand. You dig around it and find that it is a treasure chest full of gold bricks of different (integral) weight. Your knapsack can only carry up to weight  $n$  before it breaks apart. You want to put as much in it as possible without going over, but you *cannot* break the gold bricks up.
    - (a) Suppose that the gold bricks have the weights  $1, 2, 4, 8, \dots, 2^k$ ,  $k \geq 1$ . Describe and prove correct a greedy algorithm that fills the knapsack as much as possible without going over.
    - (b) Give a set of 3 weight values for which the greedy algorithm does *not* yield an optimal solution and show why.
    - (c) Give a dynamic programming algorithm that yields an optimal solution for an arbitrary set of gold brick values.