
CS 573: Graduate Algorithms, Fall 2008 Homework 4

Due at 11:59:59pm, Wednesday, October 31, 2008

- Groups of up to three students may submit a single, common solution. Please neatly print (or typeset) the full name, NetID, and the HWO alias (if any) of every group member on the first page of your submission.
- Unless a problem explicitly states otherwise, you can assume the existence of a function RANDOM(k), which returns an integer uniformly distributed in the range $\{1, 2, ..., k\}$ in O(1) time; the argument k must be a positive integer. For example, RANDOM(2) simulates a fair coin flip, and RANDOM(1) always returns 1.
- 1. Death knocks on your door one cold blustery morning and challenges you to a game. Death knows that you are an algorithms student, so instead of the traditional game of chess, Death presents you with a complete binary tree with 4ⁿ leaves, each colored either black or white. There is a token at the root of the tree. To play the game, you and Death will take turns moving the token from its current node to one of its children. The game will end after 2n moves, when the token lands on a leaf. If the final leaf is black, you die; if it's white, you will live forever. You move first, so Death gets the last turn.



You can decide whether it's worth playing or not as follows. Imagine that the nodes at even levels (where it's your turn) are OR gates, the nodes at odd levels (where it's Death's turn) are AND gates. Each gate gets its input from its children and passes its output to its parent. White and black leaves stand represent TRUE and FALSE inputs, respectively. If the output at the top of the tree is TRUE, then you can win and live forever! If the output at the top of the tree is FALSE, you should challenge Death to a game of Twister instead.

- (a) Describe and analyze a deterministic algorithm to determine whether or not you can win. [Hint: This is easy!]
- (b) Unfortunately, Death won't let you even look at every node in the tree. Describe and analyze a *randomized* algorithm that determines whether you can win in $O(3^n)$ expected time. [Hint: Consider the case n = 1.]
- (c) [Extra credit] Describe and analyze a randomized algorithm that determines whether you can win in $O(c^n)$ expected time, for some constant c < 3. [Hint: You may not need to change your algorithm at all.]

2. Consider the following randomized algorithm for choosing the largest bolt. Draw a bolt uniformly at random from the set of *n* bolts, and draw a nut uniformly at random from the set of *n* nuts. If the bolt is smaller than the nut, discard the bolt, draw a new bolt uniformly at random from the unchosen bolts, and repeat. Otherwise, discard the nut, draw a new nut uniformly at random from the unchosen nuts, and repeat. Stop either when every nut has been discarded, or every bolt except the one in your hand has been discarded.

What is the *exact* expected number of nut-bolt tests performed by this algorithm? Prove your answer is correct. [Hint: What is the expected number of unchosen nuts and bolts when the algorithm terminates?]

- 3. (a) Prove that the expected number of proper descendants of any node in a treap is *exactly* equal to the expected depth of that node.
 - (b) Why doesn't the Chernoff-bound argument for depth imply that, with high probability, *every* node in a treap has $O(\log n)$ descendants? The conclusion is obviously bogus—every n-node treap has one node with exactly n descendants!—but what is the flaw in the argument?
 - (c) What is the expected number of leaves in an *n*-node treap? [Hint: What is the probability that in an *n*-node treap, the node with kth smallest search key is a leaf?]

- (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 ith smallest and jth 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).
- *(d) [Extra Credit] Find the exact expected number of comparisons executed by RANDOMSELECT, as a function of n and r.

- 5. A *meldable priority queue* stores a set of keys from some totally-ordered universe (such as the integers) and supports the following operations:
 - MakeQueue: Return a new priority queue containing the empty set.
 - FINDMIN(Q): Return the smallest element of Q (if any).
 - DeleteMin(Q): Remove the smallest element in Q (if any).
 - INSERT(Q, x): Insert element x into Q, if it is not already there.
 - DecreaseKey(Q, x, y): Replace an element $x \in Q$ with a smaller key y. (If y > x, the operation fails.) The input is a pointer directly to the node in Q containing x.
 - Delete the element $x \in Q$. The input is a pointer directly to the node in Q containing x.
 - Meld (Q_1, Q_2) : Return a new priority queue containing all the elements of Q_1 and Q_2 ; this operation destroys Q_1 and Q_2 .

A simple way to implement such a data structure is to use a heap-ordered binary tree, where each node stores a key, along with pointers to its parent and two children. Meld can be implemented using the following randomized algorithm:

- (a) Prove that for *any* heap-ordered binary trees Q_1 and Q_2 (*not* just those constructed by the operations listed above), the expected running time of $Meld(Q_1, Q_2)$ is $O(\log n)$, where n is the total number of nodes in both trees. [Hint: How long is a random root-to-leaf path in an n-node binary tree if each left/right choice is made with equal probability?]
- (b) Prove that $Meld(Q_1, Q_2)$ runs in $O(\log n)$ time with high probability. [Hint: You don't need Chernoff bounds, but you might use the identity $\binom{c^k}{b} \leq (ce)^k$.]
- (c) Show that each of the other meldable priority queue operations can be implemented with at most one call to Meld and O(1) additional time. (This implies that every operation takes $O(\log n)$ time with high probability.)

*6. [Extra credit] In the usual theoretical presentation of treaps, the priorities are random real numbers chosen uniformly from the interval [0, 1], but in practice, computers only have access to random bits. This problem asks you to analyze a modification of treaps that takes this limitation into account.

Suppose the priority of a node ν is abstractly represented as an infinite sequence $\pi_{\nu}[1..\infty]$ of random bits, which is interpreted as the rational number

$$priority(v) = \sum_{i=1}^{\infty} \pi_{v}[i] \cdot 2^{-i}.$$

However, only a finite number ℓ_{ν} of these bits are actually known at any given time. When a node ν is first created, *none* of the priority bits are known: $\ell_{\nu} = 0$. We generate (or 'reveal') new random bits only when they are necessary to compare priorities. The following algorithm compares the priorities of any two nodes in O(1) expected time:

```
 \begin{array}{|c|c|} \hline \text{LargerPriority}(\nu,w) : \\ \hline \text{for } i \leftarrow 1 \text{ to } \infty \\ & \text{if } i > \ell_{\nu} \\ & \ell_{\nu} \leftarrow i; \quad \pi_{\nu}[i] \leftarrow \text{RandomBit} \\ & \text{if } i > \ell_{w} \\ & \ell_{w} \leftarrow i; \quad \pi_{w}[i] \leftarrow \text{RandomBit} \\ & \text{if } \pi_{\nu}[i] > \pi_{w}[i] \\ & \text{return } \nu \\ & \text{else if } \pi_{\nu}[i] < \pi_{w}[i] \\ & \text{return } w \\ \hline \end{array}
```

Suppose we insert n items one at a time into an initially empty treap. Let $L = \sum_{\nu} \ell_{\nu}$ denote the total number of random bits generated by calls to LARGER PRIORITY during these insertions.

- (a) Prove that $E[L] = \Theta(n)$.
- (b) Prove that $E[\ell_v] = \Theta(1)$ for any node v. [Hint: This is equivalent to part (a). Why?]
- (c) Prove that $E[\ell_{\text{root}}] = \Theta(\log n)$. [Hint: Why doesn't this contradict part (b)?]