

Problem Set 5

MIT students: This problem set is due in lecture on *Monday, October 22*.

SMA students: This problem set is due after the video-conferencing session on *Wednesday, October 24*.

Reading: Chapters 12 and 13

Both exercises and problems should be solved, but *only the problems* should be turned in. Exercises are intended to help you master the course material. Even though you should not turn in the exercise solutions, you are responsible for material covered by the exercises.

Mark the top of each sheet with your name, the course number, the problem number, your recitation instructor and time, the date, and the names of any students with whom you collaborated.

MIT students: Each problem should be done on a separate sheet (or sheets) of three-hole punched paper.

SMA students: Each problem should be done on a separate sheet (or sheets) of two-hole punched paper.

You will often be called upon to “give an algorithm” to solve a certain problem. Your write-up should take the form of a short essay. A topic paragraph should summarize the problem you are solving and what your results are. The body of your essay should provide the following:

1. A description of the algorithm in English and, if helpful, pseudocode.
2. At least one worked example or diagram to show more precisely how your algorithm works.
3. A proof (or indication) of the correctness of the algorithm.
4. An analysis of the running time of the algorithm.

Remember, your goal is to communicate. Graders will be instructed to take off points for convoluted and obtuse descriptions.

Exercise 5-1. Do exercise 12.2-9 on page 260 of CLRS.

Exercise 5-2. Do exercise 12.4-3 on page 268 of CLRS.




Figure 1: A treap. Each node x is labeled with $key[x] : priority[x]$. For example, the root has key G and priority 4.

Exercise 5-3. Do exercise 13.2-4 on page 279 of CLRS.

Exercise 5-4. Do exercise 13.4-7 on page 295 of CLRS.

Problem 5-1. Treaps

If we insert a set of n items into a binary search tree, the resulting tree may be horribly unbalanced, leading to long search times. On the other hand, we know that randomly built binary search trees tend to be balanced. Therefore, a strategy that, on average, builds a balanced tree for a fixed set of items is to randomly permute the items and then insert them in that order into the tree.

But, what if we do not have all the items at once? If we receive the items one at a time, can we still randomly build a binary search tree out of them?

We will examine a data structure that answers this question in the affirmative. A *treap* is a binary search tree with a modified way of ordering the nodes. Figure 1 shows an example. As usual, each node x in the tree has a key value $key[x]$. In addition, we assign $priority[x]$, which is a random number chosen independently for each node. We assume that all priorities are distinct and also that all keys are distinct. The nodes of the treap are ordered so that the keys obey the binary-search-tree property and the priorities obey the min-heap order property:

- If v is a left child of u , then $key[v] < key[u]$.
- If v is a right child of u , then $key[v] > key[u]$.
- If v is a child of u , then $priority[v] > priority[u]$.

(This combination of properties is why the tree is called a “treap;” it has features of both a binary search tree and a heap.)

It helps to think of treaps in the following way. Suppose that we insert nodes x_1, x_2, \dots, x_n , with associated keys, into a treap. Then the resulting treap is the tree that would have been formed if the nodes had been inserted into a normal binary search tree in the order given by their (randomly chosen) priorities, i.e., $priority[x_i] < priority[x_j]$ means that x_i was inserted before x_j .

- (a) Explain why a treap on n nodes is equivalent to a randomly built binary search tree on n nodes.
- (b) Conclude that the expected time to search for a value in the treap is $\Theta(\lg n)$.

Let us see how to insert a new node into an existing treap. The first thing we do is assign to the new node a random priority. Then we call the insertion algorithm, which we call TREAP-INSERT, whose operation is illustrated in Figure 2.

- (c) Explain how TREAP-INSERT works. Explain the idea in English and give pseudocode. (*Hint: Execute the usual binary-search-tree insertion procedure and then perform rotations to restore the min-heap order property.*)
- (d) Show that the expected running time of TREAP-INSERT is $\Theta(\lg n)$.




Figure 2: The operation of TREAP-INSERT. (a) The original treap, prior to insertion. (b) The treap after inserting a node with key C and priority 25. (c)–(d) Intermediate stages when inserting a node with key D and priority 9. (e) The treap after the insertion of parts (c) and (d) is done. (f) The treap after inserting a node with key F and priority 2.

Problem 5-2. Join operation on red-black trees

The *join* operation takes two dynamic sets S_1 and S_2 and an element x such that for any $x_1 \in S_1$ and $x_2 \in S_2$, we have $key[x_1] \leq key[x] \leq key[x_2]$. It returns a set $S = S_1 \cup \{x\} \cup S_2$. In this problem, we investigate how to implement the join operation on red-black trees.

- (a) Given a red-black tree T , we store its black-height as the field $bh[T]$. Argue that this field can be maintained by RB-INSERT and RB-DELETE (as given in the textbook, on pages 280 and 288 respectively) without requiring extra storage in the nodes of the tree and without increasing the asymptotic running times. Show that while descending through T , we can determine the black-height of each node we visit in $O(1)$ time per node visited.

We wish to implement the operation RB-JOIN(T_1, x, T_2), which may destroy T_1 and T_2 and returns a red-black tree $T = T_1 \cup \{x\} \cup T_2$. Let n be the total number of nodes in T_1 and T_2 .

- (b) Assume that $bh[T_1] \geq bh[T_2]$. Describe an $O(\lg n)$ -time algorithm that finds a black node y in T_1 with the largest key from among those nodes whose black-height is $bh[T_2]$.
- (c) Let T_y be the subtree rooted at y . Describe how $T_y \cup \{x\} \cup T_2$ can replace T_y in $O(1)$ time without destroying the binary-search-tree property.

Consider the following red-black properties:

- every node is either red or black
- every leaf is black
- for each node, all paths from the node to descendant leaves contain the same number of black nodes

- (d) What color should we make x so that the above red-black properties are maintained?

Consider the following red-black properties:

- the root is black
- if a node is red, then both its children are black

- (e) Describe how the above two properties can be enforced in $O(\lg n)$ time.
- (f) Argue that the running time of RB-JOIN is $O(\lg n)$.