

Treap

Nicolas Camenisch

Januar 2019

1 What is a treap?

A treap (conjunction of the words **tree** and **heap**) is a binary search tree that uses randomization and the binary heap property to maintain balance with a high probability. It maintains a dynamic set of ordered keys and allows binary searches among these keys.^[1]

Each node of a treap contains two values: A key (ordered using binary search tree properties) and a priority (ordered using heap properties). The heap order can either be maintained using a max heap or a min heap. In all following examples a max heap will be used.

The structure of a treap follows these rules:^[2]

1. The nodes in the left subtree of the root will have key values that are less than the key of the root. $\forall x \in \text{root.left}, x.\text{key} < \text{root.key}$
2. The nodes in the right subtree of the root will have key values that are greater than the key of the root. $\forall x \in \text{root.right}, x.\text{key} > \text{root.key}$
3. All priorities are unique.
4. The priority value of the root is greater than the priority values of its children. $\forall x \in \text{root.left} \cup \text{root.right}, x.\text{priority} < \text{root.priority}$

These rules ensure, that there is only one unique and valid structure for each possible set of keys and priorities.

2 Operations

Treaps support a variety of operations. These operations can either be implemented using tree rotations (Figure 1), or using merge and split. Because the latter supports more operations, we're going to focus on merge and spit.

2.1 Merge & Split

We start out by defining the struct for our treap and the corresponding nodes (Listing 1). The Treap struct is very basic and only contains a pointer to the root node and some functions. A Node consists of two pointers pointing to its left and right child, the key and priority values, and a size variable that keeps track of the size of both subtrees. During initialization the priority gets assigned a random integer.




Figure 1: A tree rotation is an operation on a binary tree that changes the structure without interfering with the order of the elements. A tree rotation moves one node up in the tree and one node down.^{[3][4]}

```

1  template <typename T>
2  struct Treap {
3      struct Node {
4          int priority;
5          T key;
6          size_t size = 0;
7          std::unique_ptr<Node> l, r;
8
9          Node(T key): priority(rand()), key(key) {}
10     };
11
12     std::unique_ptr<Node> root;
13
14     void insert(const T& key);
15     void erase(const T& key);
16     T at(size_t index);
17     size_t size() const {
18         return (root == NULL ? 0 : root->size + 1);
19     }
20
21 private:
22     static std::pair<std::unique_ptr<Node>, std::unique_ptr<Node>> split(std::
23         ↪ unique_ptr<Node> current, const T& key);
24     static std::unique_ptr<Node> merge(std::unique_ptr<Node> l, std::unique_ptr<Node> r
25         ↪ );
26 };

```

Listing 1: The Treap struct used in the following implementation.

2.1.1 Split

The split function (Listing 2) splits a treap into two separate treaps, those smaller or equal the key x , and those larger than the key x . The following implementation is in-place destructive, but there also exist nondestructive algorithms.

```

25  static std::pair<std::unique_ptr<Node>, std::unique_ptr<Node>> split(std::unique_ptr<
26     ↪ Node> current, const T& key) {
27     std::unique_ptr<Node> l, r;
28
29     if (current == NULL)

```

```

29     return {std::move(l), std::move(r)};
30     if (key < current->key) {
31         r = std::move(current);
32         tie(l, r->l) = split(std::move(r->l), key);
33         r->size -= (l == NULL ? 0 : l->size + 1);
34     } else {
35         l = std::move(current);
36         tie(l->r, r) = split(std::move(l->r), key);
37         l->size -= (r == NULL ? 0 : r->size + 1);
38     }
39
40     return {std::move(l), std::move(r)};
41 }

```

Listing 2: The split function.^[5]

2.1.2 Merge

The merge function (Listing 3) merges two treaps into one treap. For this to work the greatest key in the first treap has to be less than or equal the smallest value in the second treap. This is always true if the first and second tree both are the result of a split operation. Like the split function, this implementation is in-place destructive, but there also exist nondestructive versions.

```

42 static std::unique_ptr<Node> merge(std::unique_ptr<Node> l, std::unique_ptr<Node> r) {
43     if (l == NULL || r == NULL)
44         return (l != NULL) ? std::move(l) : std::move(r);
45
46     if (l->priority > r->priority) {
47         l->size += r->size + 1;
48         l->r = merge(std::move(l->r), std::move(r));
49         return l;
50     } else {
51         r->size += l->size + 1;
52         r->l = merge(std::move(l), std::move(r->l));
53         return r;
54     }
55 }

```

Listing 3: The merge function.^[5]

2.2 Basic operation

Based on the split and merge functions we can implement three different basic operations: search, insert and delete. Because the search operation is a standard binary search in a binary tree (ignoring the priorities), we'll skip over it. Since the following operations only modify the treap by calling the merge and split functions, the heap invariant never gets violated.

2.2.1 Insert

To insert a node N , you first search the correct position to insert it into. This means a position that doesn't violate either the BST or max heap properties. This can be done by splitting the treap into two treaps l and r , one where all values are less than or equal the key of N (l), and one where the key is greater (r). Finally you merge l , N and r to form a treap that contains the new node N (Listing 4).

```

56 void insert(const T& key) {
57     auto node = std::make_unique<Node>(key);
58
59     if (root == NULL) {
60         root = std::move(node);
61     } else {
62         auto [l, r] = split(std::move(root), key);
63         root = merge(merge(std::move(l), std::move(node)), std::move(r));
64     }
65 }

```

Listing 4: The insert function.^[5]

2.2.2 Delete

To remove a key k , you have to split the treap into three separate treaps: one containing only keys that are smaller than k (l), equal to k (eq) and greater than k (r). You then merge the left and the right subtree of eq to remove the node of highest priority with key k . Finally you merge all three treaps back together to form a treap where the node with key k is missing (Listing 5).

```

66 void erase(const T& key) {
67     if (root == NULL)
68         return;
69
70     std::unique_ptr<Node> l, r, eq;
71
72     tie(l, r) = split(std::move(root), key);
73     tie(l, eq) = split(std::move(l), key - 1);
74
75     eq = merge(std::move(eq->l), std::move(eq->r)); // Remove node with max priority
76     root = merge(merge(std::move(l), std::move(eq)), std::move(r));
77 }

```

Listing 5: The erase function.^[5]

2.3 Indexing

A treap provides random access (Listing 6) in $O(\log n)$ time. This is possible by storing the size of the subtrees of each node and then walking down the tree until the specified index is reached.

```

78 T at(size_t index) {
79     Node* current = root.get();
80
81     while (current != NULL) {
82         size_t leftSize = (current->l != NULL ? current->l->size+1 : 0);
83
84         if (index < leftSize) {
85             current = current->l.get();
86         } else if (index > leftSize) {
87             current = current->r.get();
88             index -= leftSize + 1;
89         } else {
90             return current->key;

```

```

91     }
92 }
93
94     throw std::out_of_range("");
95 }

```

Listing 6: The at function.

2.4 Bulk Operations

Treaps also support several efficient bulk operations: union, intersection and set difference. The time complexity of these bulk operations is $O(m \log \frac{n}{m})$, where m and n are the treap sizes and $m \leq n$.

3 Analysis

Due to the random nature of a treap the expected depth of a node in a tree is less than $2 \ln n = O(\log n)$, but the worst case depth is $O(n)$.^[6] Both the split and merge operation have a running time proportional to the height of the treap because they only traverse the path from the root down to the node being split / the node where it was split. Because the expected depth of any node is $O(\log n)$, the expected run time of split and merge is $O(\log n)$.

Operation	Average	Worst case
Space	$O(n)$	$O(n)$
Search	$O(\log n)$	$O(n)$
Insert	$O(\log n)$	$O(n)$
Delete	$O(\log n)$	$O(n)$
Indexing	$O(\log n)$	$O(n)$

Table 1: Time complexity of a treap in big O notation.

4 Motivation

During a normal programming contest you have 5 hours of time to solve three to four tasks. Due this tightly limited amount of time it is important to be able to implement algorithms fast and without bugs. When it comes to implementing a balanced BST, there are many options to choose from: 2-3 tree, red-black tree, scapegoat tree and many others. But by far the simplest to implement is the treap. The implementation of all other BSTs is way to complex and most of the times, the performance improvement due the guaranteed maximal height of $\log n$ isn't even necessary.^[7]

If, for example, you wanted to maintain an ordered set of elements, you could use the `std::set` provided by the C++ standard library. But since a `std::set` only provides "random access" in $O(n)$ time, you sometimes have to implement your own set. In that case a treap is the simplest way to implement an ordered set with random access in $O(\log n)$ time.

References

[1] Wikipedia contributors. Treap, . URL <https://en.wikipedia.org/wiki/Treap>.

- [2] Apple Juice Teaching. Data structures: Treaps explained. URL <https://www.youtube.com/watch?v=6podLUYinH8>.
- [3] Wikipedia contributors. Tree rotation, . URL https://en.wikipedia.org/wiki/Tree_rotation.
- [4] Ramasamy. File:tree rotation.png. URL https://commons.wikimedia.org/wiki/File:Tree_rotation.png.
- [5] ABDELKARIM0. Treap simple implementation. URL <https://acmcairoscience.wordpress.com/2016/03/31/treap-simple-implementation/>.
- [6] Margaret Reid-Miller. Lecture 16 — treaps; augmented bsts. URL <http://www.cs.cmu.edu/afs/cs/academic/class/15210-s12/www/lectures/lecture16.pdf>.
- [7] Paul Shved. Treap: simple balanced search tree. URL <http://coldattic.info/post/4/>.