

# Splay Trees

Thomas A. Anastasio

1 November 1999

## 1 Introduction

The average binary search tree supports insertion, deletion, and find operations in  $O(\lg n)$  time. However, there are worst case trees for which these operations are in  $O(n)$ . One approach to avoiding the worst case behavior is to require that the tree always remain balanced. The AVL tree is one such balanced tree. Balanced trees guarantee  $O(\lg n)$  time per operation but do so at the cost of a high constant of proportionality. AVL trees adjust their structure after an insertion or deletion operation if necessary to maintain balance. They do not adjust their structure in response to a find operation.

Splay trees are “self-adjusting” binary search trees in which the adjustment is done on the basis of the history of accesses rather than in an attempt to maintain an explicit balance. After access, a tree node is moved to the root by a sequence of rotations called “splaying.” No balance information need be stored in the nodes and, in general, splay trees are not balanced. The most recently accessed nodes in a splay tree will tend to be close to the root. Accessing these nodes is less expensive than accessing nodes further down the tree. Because of data locality effects, having the most recently accessed nodes near or at the root can be a big gain.

The major point about splay trees is that the amortized performance of a sequence of  $m$  operations is  $O(m \lg n)$ .  $n$  is the total number of insertions (the maximum size attainable by the tree during the sequence of  $m$  operations). Any individual operation in the sequence could be as costly as  $O(n)$ ; there is no guarantee that a particular operation in the sequence will be efficient. Compare this with the performance of a sequence of operations in an AVL tree. Since each operation on the AVL tree is in  $O(\lg n)$ , the sequence of  $m$  operations will be in  $O(m \lg n)$ . Splay trees have the same asymptotic performance over a sequence of operations as do balanced trees. In the splay tree, some operations in the sequence may be inefficient, but other operations will be super-efficient to make up for them.

So what’s a sequence of operations? It’s any sequence of `insert`, `remove`, or `find` operations performed in building and manipulating a splay tree from scratch. There are some constraints, of course. For example, the  $k$ -th deletion must have been preceded by at least  $k$  insertions.

## 2 Tree Operations

The `insert` operation in a splay tree works by first doing a normal binary-search tree insertion, then splaying the inserted node. This puts the inserted node at the root. If the insertion fails because a duplicate element is already in the tree, the node holding the duplicate is splayed.

The `find` operation searches the splay tree for the element sought. If found, that node is splayed to the root. If not found, the last node encountered on the search path (the parent of the non-existent element) is splayed to the root.

The `remove` operation, as in regular binary-search trees, removes the current element. In a splay tree, `remove` involves three splaying operations. The current node is splayed to the root, producing a new splay tree  $T$ . Let  $T_L$  and  $T_R$  be the left and right subtrees of  $T$ . Disconnect these subtrees from  $T$ , then splay the largest element in  $T_L$  and the smallest element in  $T_R$ , after which  $T_L$  will have no right subtree and  $T_R$  will have no left subtree. Now, form a new tree  $T$  by making  $T_R$  the right subtree of  $T_L$  (or vice-versa, if  $T_L$  is empty). This new tree is the original tree after the deletion.

## 3 Overview of the Splaying Operation

A binary tree is *splayed* at a node  $X$  by traversing the tree from  $X$  to the root, performing sequences of single rotations along the way. The following splay operation is repetitively applied, beginning at node  $X$ , until  $X$  is at the root.

1. If  $X$  is root, do nothing.
2. If  $X$  has no grandparent (*i.e.*, the parent of  $X$  is root), rotate  $X$  about its parent. This will make  $X$  be root. See Figure 1.




Figure 1: Node Has No Grandparent

3. If X has a grandparent:

- (a) If X and its parent are both left-children or both right-children, rotate the parent about the grandparent, then rotate X about its parent. See Figure 2.




Figure 2: Node and Parent are Same Type

- (b) If  $X$  and its parent are opposite-type children (one is left, the other is right) rotate  $X$  about its parent, then rotate  $X$  about its new parent (*i.e.*, its former grandparent). See Figure 3.




Figure 3: Node and Parent are Opposite Type

## 4 Examples

Figure 4 shows the splay tree resulting from insertion of the Integers 1..4 in order. Note that the sequential insertions resulted in the same type of long tree that would have been obtained with ordinary binary search trees. In ordinary binary search trees, each insertion in the sequence of four insertions would have taken  $O(n)$  time. Thus, the entire sequence of operations would have taken  $O(n^2)$ . This is not the case with the splay tree. Notice that each insertion was done just below the root because splaying moved the newly inserted node to the root. Thus, this particular sequence of four insertions was done in  $O(1)$  per operation, for a total of  $O(n)$  for the entire sequence.

The last tree shown in Figure 4 results from performing a `find` operation on the node with element 1.

## 5 Implementation of Splay Trees

Splay trees are implemented as extensions of ordinary binary search trees. The efficiency of the splaying operation depends heavily on the ability to find the

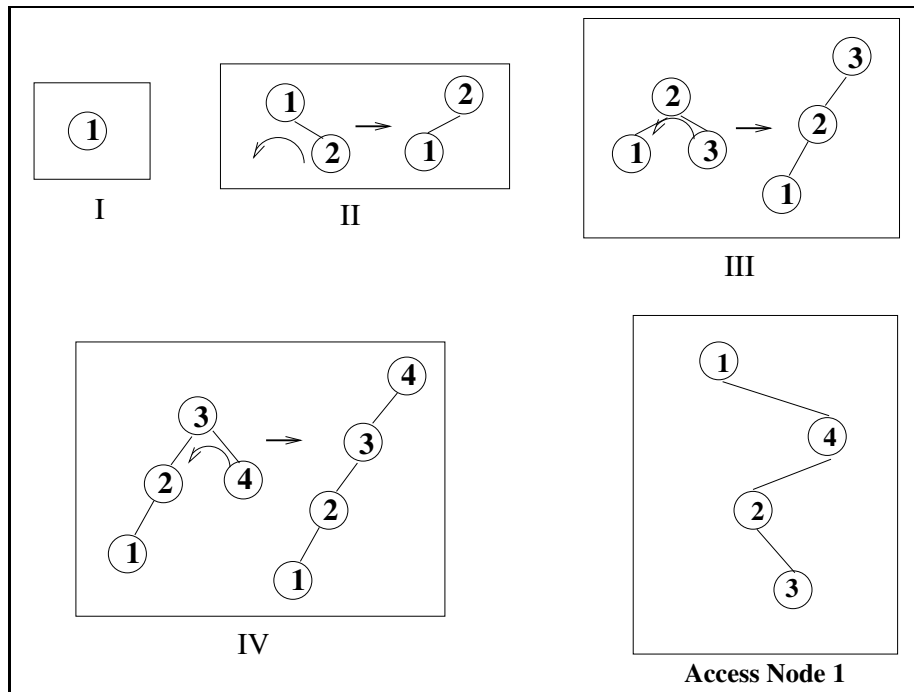


Figure 4: Insertion in Order

parent of a node. This can be done in  $O(1)$  time (at the cost of  $O(n)$  additional storage) by having each node store a reference to its parent. Alternatively, parents can be stored on a stack during recursive descent of the tree. This complicates the code somewhat but requires only  $O(\text{depth})$  additional storage.

Splaying takes place at a particular node, so a `splay` method can be added as a member function to the node class.

### 5.1 The find Operation

The `find` operation searches for a node with a given key, starting from a given node, usually the root of the tree. This is the same type of search as is done in an ordinary binary search tree. In a splay tree, the search is followed by a splay operation. If the search succeeds, the node containing the key is splayed. If the search fails, the last node visited is splayed.

### 5.2 The insert Operation

Insertion in a splay tree is done in precisely the same way as in an ordinary binary search tree. This insertion method does not allow duplicate elements to be inserted. If the element to be inserted is already in the tree, then the

insertion is not done, but the node holding the duplicate element is splayed. If the insertion succeeds, then the newly-inserted node is splayed.

### 5.3 The remove Operation

The `remove` operation for a splay tree differs markedly from that in an ordinary BST. In fact, it's a lot simpler. If the item to be deleted is not in the tree, the node last visited in the search for it is splayed and we are done. On the other hand, if the item is in the tree, its node is splayed and then deletion occurs as described below. In either event, a `find` for the item will result in the appropriate node being splayed to the root.

Let's look at the case in which the item is in the tree. The splay that resulted from the `find` operation puts that item's node at the root. In general, the node has both left and right children. Do the following operations:

1. Disconnect the left and right subtrees from the root. Call these  $T_L$  and  $T_R$ , respectively. Discard the root node.
2. Splay the maximum item in  $T_L$ . This will result in the root of  $T_L$  having no right child. (The largest element is at the root, so the root has no larger elements - no right subtree.)
3. Splay the minimum item in  $T_R$ . This will result in the root of  $T_R$  having no left child. (The smallest element is at the root, so the root has no smaller elements - no left subtree.)
4. All the elements in  $T_L$  are smaller than the elements in  $T_R$ . Make  $T_R$  be the right subtree of the root of  $T_L$  (or vice-versa).

Notice that there are three splays done.

1. The splay done when searching for the item in the tree.
2. The splay done in the left subtree  $T_L$ .
3. The splay done in the right subtree  $T_R$ .

There are no "cases" to worry about as in the ordinary BST deletion method.

## 6 Performance Considerations

The `insert` operation uses ordinary BST insertion, then splays the inserted node to the root. Clearly the insertion portion of this operation has the same asymptotic behavior as that for BSTs, namely  $O(\text{depth})$ . What about the splaying portion? If the work required at each step of the splay is independent of the number of nodes in the tree, the splay portion will be in  $O(1)$  per step. The number of steps is bounded by the depth of the tree, so the splaying portion would also be in  $O(\text{depth})$ .

Well, is the splaying portion independent of the number of nodes in the tree? Look at what's involved:

1. Find the parent of the node and perhaps its parent too.
2. Determine if the node and its parent are the same type (left or right).
3. Rearrange the local structure of the tree.

Determining the type and rearranging the structure are local operations involving only a few nodes (the node, its parent, its grandparent). These operations must be in  $O(1)$ .

The cost of determining the parent of a node depends on how the parent is found. If each node stores a reference to its parent (or if a stack of parents is maintained), the cost is in  $O(1)$ , keeping the splay operation in  $O(\text{depth})$ .

### 6.1 Amortized Cost

For the average BST,  $d = O(\lg n)$ , so on average, operations on BSTs are in  $O(\lg n)$ . A sequence of  $m$  operations (`insert`, `remove`, and `find`) will cost  $O(m \lg n)$ .

The situation with splay trees is different. The structure of the tree is changed with each operation, so the notion of average is not relevant. A technique known as “amortized analysis” is used to show that over a sequence of  $m$  operations (starting with an empty tree) the cost of the sequence will be  $O(m \lg n)$ . This is the same cost as that for BSTs so overall performance is the same. What’s different is that the splay tree does not guarantee  $O(\lg n)$  performance for *each* operation.

## 7 A Final Example

Figure 5 shows a splay tree before and after node ‘‘X’’ is accessed by an `find` operation.




Figure 5: A Splaying Example