

Binary search trees provide $O(\log N)$ search times provided that the nodes are distributed in a reasonably “balanced” manner. Unfortunately, that is not always the case and performing a sequence of deletions and insertions can often exacerbate the problem.

When a BST becomes badly unbalanced, the search behavior can degenerate to that of a sorted linked list, $O(N)$.

There are a number of strategies for dealing with this problem; most involve adding some sort of restructuring to the insert/delete algorithms.

That can be effective only if the restructuring reduces the average depth of a node from the root of the BST, and if the cost of the restructuring is, on average, $O(\log N)$.


We will examine one such restructuring algorithm...

AVL tree*: a binary search tree in which the heights of the left and right subtrees of the root differ by at most 1, and in which the left and right subtrees are themselves AVL trees.

Each AVL tree node has an associated balance factor indicating the relative heights of its subtrees (left-higher, equal, right-higher). Normally, this adds one data element to each tree node and an enumerated type is used.

How effective is this? The height of an AVL tree with N nodes never exceeds $1.44 \log N$ and is typically much closer to $\log N$.


***G. M. Adelson-Velskii and E. M. Landis, 1962.**




This is an AVL tree...

and this is not.

Consider inserting the value 45 into the AVL tree:




The result would be unbalanced at the node containing 25:




The unbalance is repaired by applying one of two types of “rotation” to the unbalanced subtree...

The subtree rooted at 25 is right-higher.



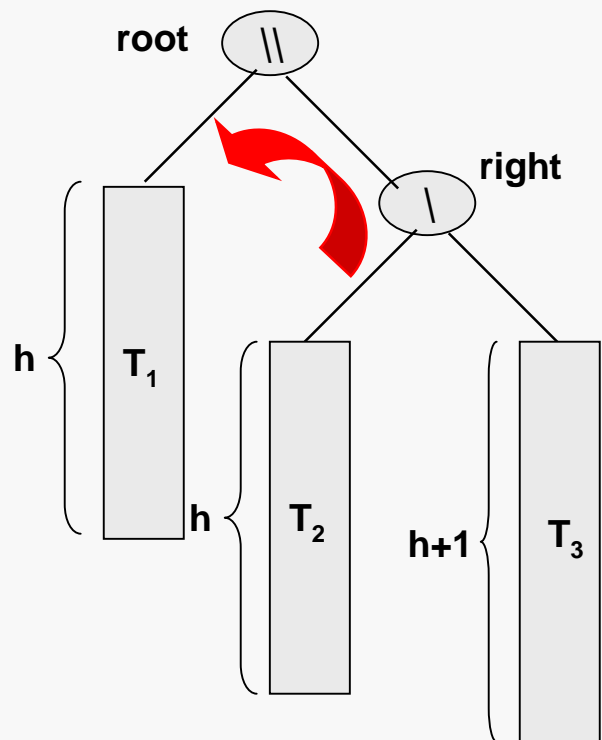
We restructure the subtree, resulting in a balanced subtree:



The transformation is relatively simple, requiring only a few operations, and results in a subtree that has equal balance.

There are two unbalance cases to consider, each defined by the state of the subtree that just received a new node. For simplicity, assume for now that the insertion was to the right subtree (of the subtree).

Let `root` be the root of the newly unbalanced subtree, and suppose that its right subtree is now right-higher:




In this case, the subtree rooted at `right` was previously equally balanced (why?) and the subtree rooted at `root` was previously right-higher (why?).

The height labels follow from those observations.


Balance can be restored by “rotating” the values so that `right` becomes the subtree root node and `root` becomes the left child.

The manipulation just described is known as a “left rotation” and the result is:



That covers the case where the right subtree has become right-higher... the case where the left subtree has become left-higher is analogous and solved by a right rotation.

Now suppose that the right subtree has become left-higher:




The insertion occurred in the left subtree of the right subtree of root.

In this case, the left subtree of the right subtree (rooted at right-left) may be either left-higher or right-higher, but not balanced (why?).


Surprisingly (perhaps), this case is more difficult. The unbalance cannot be removed by performing a single left or right rotation.

Applying a single right rotation to the subtree rooted at `right` produces...



...a subtree rooted at `right-left` that is now right-higher...

Now, applying a single left rotation to the subtree rooted at `root` produces...



...a balanced subtree.

The case where the left subtree of `root` is right-higher is handled similarly (by a double rotation).

An AVL tree implementation can be derived from the `BinNodeT` and `BSTreeT` classes seen earlier.

The `AVLNodeT` class must add a data member for the balance factor, and accessor and mutator member functions for that data member:

```
#include "BinNodeT.h"
#include "Enums.h"

template <class Data> class AVLNodeT : public BinNodeT<Data> {
protected:
    BFactor Balance;

public:
    AVLNodeT();
    AVLNodeT(Data newData, AVLNodeT<Data>* newLeft = NULL,
              AVLNodeT<Data>* newRight = NULL);
    void setBalance(BFactor newBFactor);
    BFactor getBalance() const;
    ~AVLNodeT();
};
```

QTP: Must the `BinNodeT` class be modified as well?

The AVL tree interface:

```
template <class Data> class AVLTreeT : public BSTreeT<Data> {
protected:

    bool AVLInsertHelper(Data D, BinNodeT<Data>* sRoot, bool& Taller);
    void RightBalance(BinNodeT<Data>* sRoot);
    void LeftBalance(BinNodeT<Data>* sRoot);
    void RotateLeft(BinNodeT<Data>* sRoot);
    void RotateRight(BinNodeT<Data>* sRoot);

    // delete helper function prototypes . . .

public:
    AVLTreeT();
    AVLTreeT(Data newData);

    bool Insert(Data D);
    bool Delete(Data D);

    ~AVLTreeT();
};
```

QTP: Must the `BSTreeT` class be modified as well?

```
void AVLTreeT<Data>::RotateLeft(BinNodeT<Data>* sRoot) {  
  
    if ( (sRoot == NULL) || (sRoot->getRight() == NULL) ) return;  
  
    AVLNodeT<Data>* Temp = new AVLNodeT<Data>(sRoot->getData());  
    if (Temp == NULL) return;  
  
    Temp->setLeft(sRoot->getLeft());  
    sRoot->setLeft(Temp);  
    Temp->setRight(sRoot->getRight()->getLeft());  
    BinNodeT<Data>* toDelete = sRoot->getRight();  
    sRoot->setData(toDelete->getData());  
    sRoot->setRight(toDelete->getRight());  
    delete toDelete;  
}
```

QTP: Is everything that should be copied taken into account?

```
template <class Data>
void AVLTreeT<Data>::RightBalance(BinNodeT<Data>* sRoot) {

    if ( (sRoot == NULL) || (sRoot->getRight() == NULL) ) return;

    BinNodeT<Data>* rightSubTreeRoot = sRoot->getRight();

    switch ( sRoot->getRight()->getBalance() ) {

    case RHIGHER:    sRoot->setBalance(EQUALHT);
                    rightSubTreeRoot->setBalance(EQUALHT);
                    RotateLeft(sRoot);
                    break;

    // . . . continues . . .
```

```
// . . . continued . . .

case LHIGHER:  BinNodeT<Data>* leftSubtreeRightSubTree =
                rightSubTreeRoot->getLeft();

                switch ( leftSubtreeRightSubTree->getBalance() ) {

case EQUALHT:  sRoot->setBalance(EQUALHT);
                rightSubTreeRoot->setBalance(EQUALHT);
                break;

case LHIGHER:  sRoot->setBalance(EQUALHT);
                rightSubTreeRoot->setBalance(RHIGHER);
                break;

case RHIGHER:  sRoot->setBalance(LHIGHER);
                rightSubTreeRoot->setBalance(EQUALHT);
                break;

                }
                leftSubtreeRightSubTree->setBalance(EQUALHT);
                RotateRight(rightSubTreeRoot);
                RotateLeft(sRoot);
                break;
        }
}
```

```
template <class Data>
bool AVLTreeT<Data>::Insert(Data D) {

    if (Root == NULL) {
        // insert node
        AVLNodeT<Data>* Temp = new AVLNodeT<Data>(D);
        if (Temp == NULL) return false;
        Root = Current = Temp;
        return true;
    }
    bool Taller;
    return AVLInsertHelper(D, Root, Taller);
}
```

```
template <class Data>
bool AVLTreeT<Data>::AVLInsertHelper(Data D, BinNodeT<Data>* sRoot,
                                     bool& Taller) {

    if ( (D < sRoot->getData()) && (sRoot->getLeft() == NULL) ) {

        AVLNodeT<Data>* Temp = new AVLNodeT<Data>(D);
        if (Temp == NULL) return false;

        sRoot->setLeft(Temp);

        if (sRoot->getBalance() == RHIGHER) {
            sRoot->setBalance(EQUALHT);
            Taller = false;
        }
        else if (sRoot->getBalance() == EQUALHT) {
            sRoot->setBalance(LHIGHER);
            Taller = true;
        }
        return true; // completed insertion
    }

    // . . . continues . . .
}
```

Inserting a new left leaf at the current subtree root node.

Correct the balance factor at the subtree root, and notify its parent whether the subtree grew taller as a result of the insertion.

```
// . . . continued . . .  
  
    if ( (D >= sRoot->getData()) && (sRoot->getRight() == NULL) ) {  
  
        // code is analogous to first case . . .  
  
        return true; // completed insertion  
    }  
  
// . . . continues . . .
```

**Inserting a new right leaf
at the current subtree
root node.**

```
// . . . continued . . .
```

```
if (D < sRoot->getData()) {  
    bool Success = AVLInsertHelper(D, sRoot->getLeft(), Taller);
```

Insertion occurs in the left subtree of sRoot.

```
    // analogous to succeeding case . . .  
    return Success;
```

If left subtree just grew . . .

```
}
```

```
// . . . continues . . .
```

AVL Insert Helper: Recurse to Right Subtree

```
// . . . continued . . .
```

```
else {
```

```
    bool Success = AVLInsertHelper(D, sRoot->getRight(), Taller);
```

```
    if (Taller) {
```

```
        switch ( sRoot->getBalance() ) {
```

```
            case RHIGHER:  RightBalance(sRoot);  
                          Taller = false;  
                          break;
```

```
            case EQUALHT:  sRoot->setBalance(RHIGHER);  
                          break;
```

```
            case LHIGHER:  sRoot->setBalance(EQUALHT);  
                          Taller = false;  
                          break;
```

```
        }
```

```
    }
```

```
    return Success;
```

```
}
```

```
}
```

Insertion occurs in the right subtree of sRoot.

See if subtree just grew. . .

sRoot was right-higher. . .

sRoot was balanced. . .

sRoot was left-higher. . .

Deleting a node from an AVL tree can also create an imbalance that must be corrected.

The effects of deletion are potentially more complex than those of insertion.

The basic idea remains the same: delete the node, track changes in balance factors as the recursion backs out, and apply rotations as needed to restore AVL balance at each node along the path followed down during the deletion.

However, rebalancing after a deletion may require applying single or double rotations at more than one point along that path.

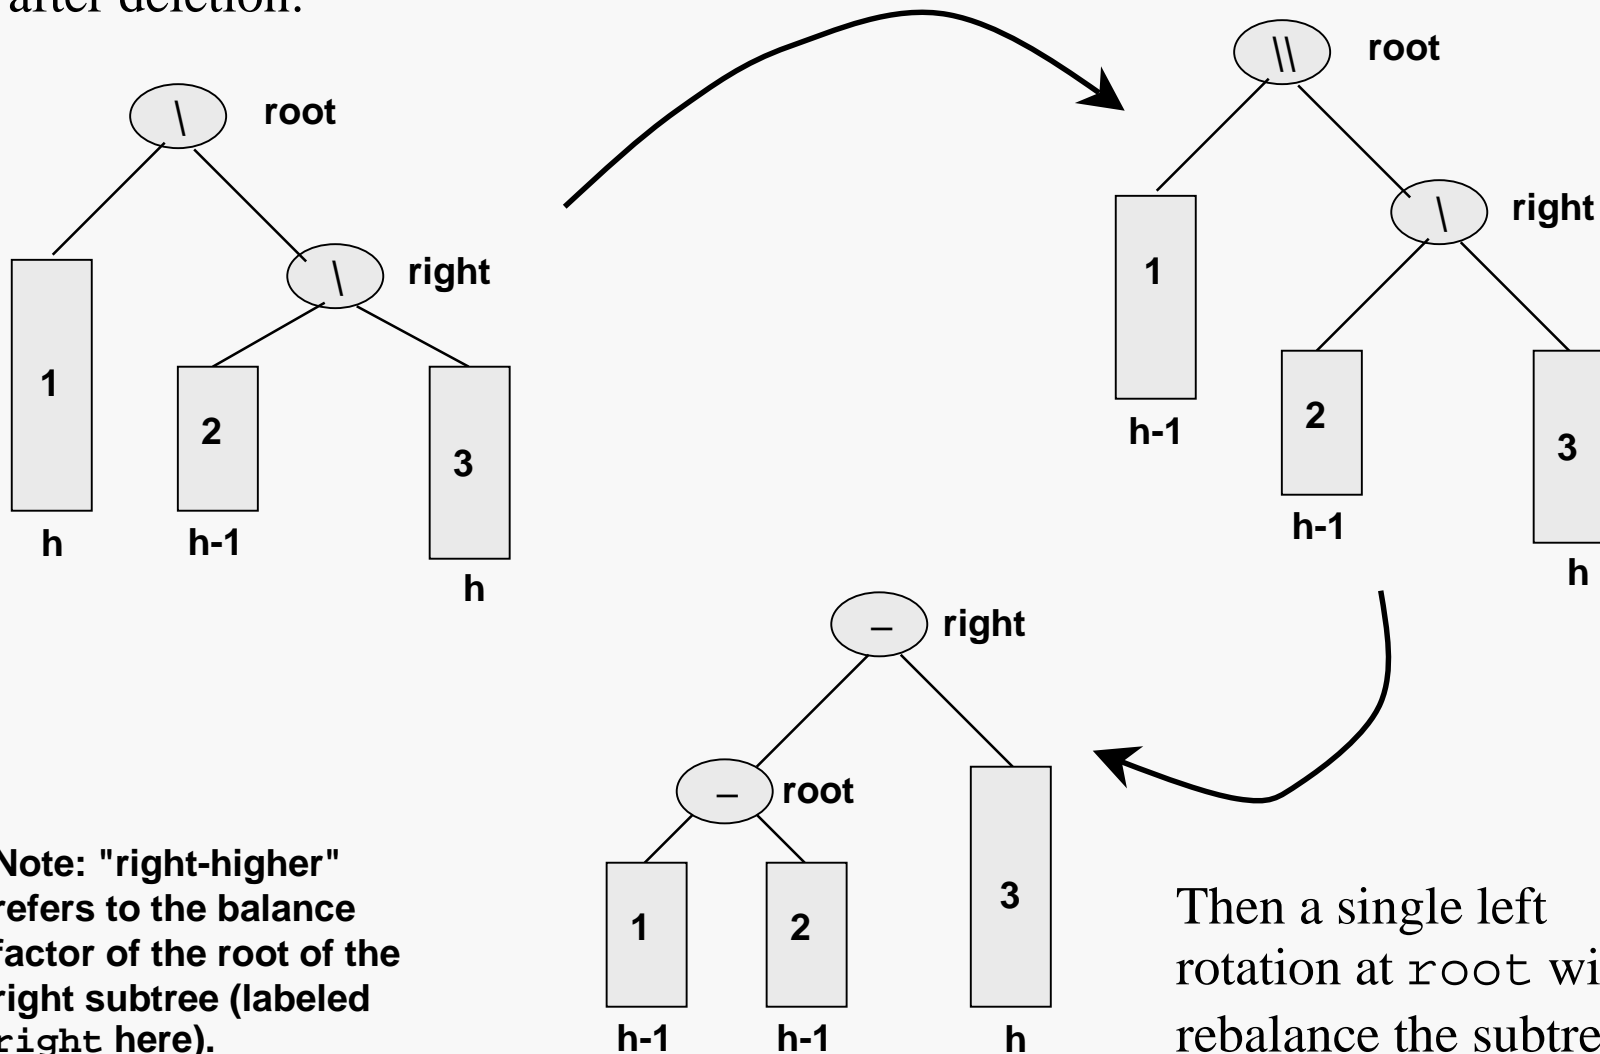
As usual, there are cases...

Here, we will make the following assumptions:

- the lowest imbalance occurs at the node `root` (a subtree root)
- the deletion occurred in the left subtree of `root`

AVL Deletion Case: right-higher


Suppose we have the subtree on the left prior to deletion and that on the right after deletion:



Note: "right-higher" refers to the balance factor of the root of the right subtree (labeled **right** here).


Then a single left rotation at **root** will rebalance the subtree.

Suppose we the right subtree root has balance factor equal-height:



The difference is the resulting balance factor at the old subtree root node, `root`, which depends upon the original balance factor of the node `right`.

If the right subtree root was left-higher, we have the following situation:




Deleting a node from the left subtree of root now will cause root to become double right higher.


As you should expect, the resulting imbalance can be cured by first applying a right rotation at the node *right*, and then applying a left rotation at the node *root*.

However, we must be careful because the balance factors will depend upon the original balance factor at the node labeled *right-left*...

If the right-left subtree root was also left-higher, we obtain:



If the right-left subtree root was right-higher, we obtain:



And, finally, if the right-left subtree root was equal-height, we'd obtain a tree where all three of the labeled nodes have equal-height.

We have considered a number of distinct deletion cases, assuming that the deletion occurred in the left subtree of the imbalanced node.

There are an equal number of entirely similar, symmetric cases for the assumption the deletion was in the right subtree of the imbalanced node.

Drawing diagrams helps...

This discussion also has some logical implications for how insertion is handled in an AVL tree. The determination of the balance factors in the tree, following the rotations, involves similar logic in both cases.

Turning the previous discussion into an implementation involves a considerable amount of work.

At the top level, the public delete function must take into account whether the deletion shortened either subtree of the tree root node. If that was the case, then it may be necessary to perform a rotation at the tree root itself.

Thus, the helper function(s) must be recursive and must indicate whether a subtree was shortened; this may be accomplished by use of a `bool` parameter, as was done for the insertion functions.

Deletion of the tree root node itself is a special case, and should take into account whether the tree root node had more than one subtree. If not, the root pointer should simply be retargeted to the appropriate subtree and no imbalance occurs.

If the tree root node has two subtrees, its value should be swapped with the minimum value in the right subtree and then the corresponding node should be deleted from the right subtree.

```
template <class Data>
bool AVLTreeT<Data>::Delete(Data D) {

    if (Root == NULL) return false;

    if (Root->getData() == D) {                // delete tree root
        BinNodeT<Data>* toDelete = Root;
        if (Root->getLeft() == NULL) {        // no left subtree
            Root = Root->getRight();
            delete toDelete;
            return true;
        }
        if (Root->getRight() == NULL) {       // left but no right subtree
            Root = Root->getLeft();
            delete toDelete;
            return true;
        }
    }

    // . . . continues . . .
}
```

```
// . . . continued . . . still deleting the root node
bool Shorter;

BinNodeT<Data>* Look = Root->getRight();
while (Look->getLeft() != NULL)
    Look = Look->getLeft();

Root->setData(Look->getData());
Look->setData(D);

if (Look == Root->getRight()) {
    Root->setRight(Look->getRight());
    delete Look;
    Shorter = true;
}
else {

    DeleteRightMinimum(Root->getRight(), Shorter);
}
// . . . continues . . .
```

Find minimum value in right subtree...

...swap with root value...

Case where the minimum was a child of the root is a special case (in my implementation).

...delete swap node from subtree...

```
// . . . continued . . . deleted root node, need to rebalance?
```

```
if (Shorter) {  
    switch ( Root->getBalance() ) {  
  
        case EQUALHT:  Root->setBalance(LHIGHER);  
                       break;  
  
        case LHIGHER:  // Root is unbalanced  
                       LeftDelBalance(Root, Shorter);  
                       break;  
  
        case RHIGHER:  Root->setBalance(EQUALHT);  
                       // sRoot's subtree DID get shorter  
                       break;  
  
    }  
}  
return true;  
}  
// . . . continues . . .
```

The right subtree just got shorter...

Logic is different than in insertion, and so is this function... or is it?

Public AVL Delete Function: Deletion Below Root

AVL Trees 32

```
// . . . continued . . . not deleting the root node
```

```
bool Shorter;
```

```
AVLDeleteHelper(D, Root, Shorter);
```

```
if (D < Root->getData()) {
```

```
    if (Shorter) {
```

```
        switch ( Root->getBalance() ) {
```

```
            case EQUALHT: Root->setBalance(RHIGHER);  
                        break;
```

```
            case RHIGHER: // Root is unbalanced  
                        RightDelBalance(Root, Shorter);  
                        break;
```

```
            case LHIGHER: Root->setBalance(EQUALHT);  
                        // sRoot's subtree DID get shorter  
                        break;
```

```
        }
```

```
    }
```

```
}
```

```
else {
```

```
    // similar (!) to if-clause
```

```
}
```

```
return true;
```

```
}
```

Value is not in root, so we must find and delete a lower node...

If left subtree became shorter, rebalance...

...otherwise if right subtree became shorter, rebalance...

Deleting the Right Minimum

```
template <class Data>
bool AVLTreeT<Data>::DeleteRightMinimum(BinNodeT<Data>* sRoot,
                                         bool& Shorter) {

    if (sRoot == NULL || sRoot->getLeft() == NULL)
        return false;

    BinNodeT<Data>* ParentOfMinimum = sRoot;

    if (ParentOfMinimum->getLeft()->getLeft() != NULL) {
        DeleteRightMinimum(ParentOfMinimum->getLeft(), Shorter);
    }
    else {
        BinNodeT<Data>* toDelete = ParentOfMinimum->getLeft();
        ParentOfMinimum->setLeft(toDelete->getRight());
        delete toDelete;
        Shorter = true;
    }
    if (Shorter) {
        // fun stuff . . .
    }
    return true;
}
```

I assume here that sRoot does NOT contain the minimum value...

Recurse to the left until it's time to stop....

Must check for imbalance in the parent of the node just deleted...

The AVL Delete function uses a recursive helper function for the case where the deleted node is not the tree root.

That function, `AVLDeleteHelper()`, employs similar search logic to the `DeleteHelper()` function for the BST. However, once the target node is found, the usual swap is performed and `DeleteRightMinimum()` is called.

The AVL Delete code also use different top-level balance functions than those used for insertion.

Logically the functions are very similar... the primary differences lie in the readjustment of balance factors.

The left and right rotation logic is **NOT** changed for deletion, so those functions are shared.

In testing the AVL implementation, I found a few useful ideas:

- Test with small cases first, and vary them to cover all the different scenarios you can think of.
- Work out the solutions manually as well. If you don't understand how to solve the problem by hand, you'll NEVER program a correct solution.
- Test insertion first, thoroughly.
- Then build an initial, correct AVL tree and test deletion, thoroughly.
- Then test alternating combinations of insertions and deletions.
- Modify your tree print function to show the balance factors:

