


CSCB63 – Design and Analysis of Data Structures

Anya Tafliovich¹

¹based on notes by Anna Bretscher and Albert Lai

our dear friend the BST (Binary Search Tree)



Q. What did we use it for?

A. Sets (keys only), ordered dictionaries (key/value pairs)

Q. Can things go wrong?

A. Yes, actual worst-case complexity of, say, search is $\mathcal{O}(n)$, not $\mathcal{O}(\log n)$.

BST — when things go wrong

Example: show a sequence of values that, when inserted into an initially empty BST, creates a “bad” BST.

Insert 5, 4, 3, 2, 1.

Q. What is the complexity of search in this tree?

A. $\mathcal{O}(n)$, basically a linked list: not good enough.

solution — balanced trees

Idea: Maintain a Binary Search Tree that always stays balanced.


Several ways to accomplish more-or-less the same result:

- Red-Black trees
- AVL trees — G. Abelson-Velsky and E. Landis
- B-trees
- Splay trees

AVL tree

- stores key/value pairs in all nodes (both leaf and internal)
- has a property relating the keys stored in a subtree to the key stored in the parent node (ordering)
- maintains the height (number of nodes/edges on a root-to-leaf path) of $\mathcal{O}(\log n)$
 - balance factor = height(left subtree) – height(right subtree)
 - maintain balance factor of ± 1 or 0 for all nodes

AVL tree



AVL Tree

- maintains the height (number of nodes/edges on a root-to-leaf path) of $\mathcal{O}(\log n)$
 - balance factor = height(left subtree) – height(right subtree)
 - maintain balance factor of ± 1 or 0 for all nodes

AVL tree operations

Operations of an ordered dictionary:

- `search(k, T)`: return the value corresponding to key k in the tree T
 - special scenario if $k \notin T$
- `insert(k, v, T)`: insert the new key/value pair k/v into the tree T
 - special scenario if $k \in T$
- `delete(k, T)`: delete the key/value pair with key k from the tree T
 - special scenario if $k \notin T$

AVL search

Q. How should we implement `search(k, T)`?

A. Same as BST!

AVL insert

Q. How should we implement `insert(k, v, T)`?

Q. Can we take the same approach as with `search`?



A. No! The tree may become unbalanced as a result of insertion.

For example:

Insert key 60 into the example tree above.

AVL insert

Insert scenario:



Insert key 6.



Q. What does the new tree look like?

Q. What are the new balance factors?

A. Tree on the right above.

AVL insert

Another insert scenario:



Insert key 35.

Q. What does the new tree look like? What are the new balance factors?


A. Tree on the right above.

Q. What is the problem?

A. The balance property is broken.

AVL insert



Solve the problem:



- perform a single rotation: counter-clockwise
- update balance factors

AVL insert

Another insert scenario:



Insert key 45.

Q. What does the new tree look like? What are the new balance factors?



A. Tree on the right above.

Q. What is the problem?

A. The balance property is broken.

AVL insert


Solve the problem:



- perform a single rotation: clockwise
- note: node with key 62 gets a new parent!
- update balance factors

AVL insert

Another insert scenario:




Insert key 46.

Q. What is the problem? Can we fix it by one of the methods above?

A. The balance property is broken. No! We need a double rotation.

AVL insert

Solve the problem:



Q. How do we know that we need a double rotation?

A. Have change in sign of a balance factor.

AVL insert

Important observations before we develop the complete algorithm:

- do we need to perform only 1 rotation per insert?
 - no, sometimes need 2 to rebalance

- how can we be sure this won't end up being $\mathcal{O}(n)$?
 - Intuitively: Because we only update on one leaf-to-root path, the number of updates is $\mathcal{O}(\log n)$ and each update is $\mathcal{O}(1)$

 - Formally:
 - prove updates only on root-to-leaf path
 - prove height of AVL tree is $\mathcal{O}(\log n)$

AVL rebalancing

For each node v on the new-node-to-root path:


```
if height(v.left) - height(v.right) > 1:
    let x = v.left
    if height(x.left) >= height(x.right):
        single rotation: clockwise
    else:
        double rotation: counter-clockwise then clockwise

else if height(v.right) - height(v.left) > 1:
    let x = v.right
    if height(x.left) <= height(x.right):
        single rotation: counter-clockwise
    else:
        double rotation: clockwise then counter-clockwise

else:
    no rotation
```

AVL rebalancing: rotation clockwise

```
if height(v.left) - height(v.right) > 1:  
  let x = v.left  
  if height(x.left) >= height(x.right):  
    single rotation: clockwise
```




Q. How do we know x exists?

Q. How do we know the result is a BST?

AVL rebalancing: rotation counter-clockwise

```
if height(v.right) - height(v.left) > 1:  
  let x = v.right  
  if height(x.left) <= height(x.right):  
    single rotation: counter-clockwise
```




Q. How do we know x exists?

Q. How do we know the result is a BST?

AVL rebalancing: double rotation

```


if height(v.right) - height(v.left) > 1:
    let x = v.right
    ...
    // height(x.left) > height(x.right):
        let w = x.left
        double rotation: clockwise then counter-clockwise
    
```



AVL rebalancing: double rotation

```

if height(v.left) - height(v.right) > 1:
    let x = v.left
    ...
    // height(x.left) < height(x.right):
        let w = x.right
        double rotation: counter-clockwise then clockwise
    
```



AVL insert

- find the node to become parent of new node
 - complexity: $\Theta(\log n)$
- put new node there
 - complexity: $\Theta(1)$
- rebalance and update heights if needed
 - complexity: $\Theta(\log n)$

This means we need to store, for each node, either the height of the subtree rooted at it or its balancing factor.

We will prove that the height of the AVL tree is $\mathcal{O}(\log n)$ shortly.