

Unit Outline

Unit #6: AVL Trees

CPSC 221: Algorithms and Data Structures

Will Evans and Jan Manuch

2016W1

- ▶ Binary search trees
- ▶ Balance implies shallow (shallow is good)
- ▶ How to achieve balance
- ▶ Single and double rotations
- ▶ AVL tree implementation

2 / 29

Learning Goals

- ▶ Compare and contrast balanced/unbalanced trees.
- ▶ Describe and apply rotation to a BST to achieve a balanced tree.
- ▶ Recognize balanced binary search trees (among other tree types you recognize, e.g., heaps, general binary trees, general BSTs).


Dictionary ADT Implementations

Worst Case time	insert	find	delete (after find)
Linked list			
Unsorted array			
Sorted array			
Hash table			

3 / 29

4 / 29

Binary Search in a Sorted List



```
int bSearch(int A[], int key, int i, int j) {
    if (j < i) return -1;
    int m = (i + j) / 2;
    if (key < A[m]) return bSearch(A, key, i, m-1);
    else if (key > A[m]) return bSearch(A, key, m+1, j);
    else return m;
}
```

5 / 29


Binary Search Tree as Dictionary Data Structure

Binary tree property

- ▶ each node has ≤ 2 children

Search tree property


- ▶ all keys in left subtree smaller than node's key
- ▶ all keys in right subtree larger than node's key



Result: easy to find any given key

6 / 29

In-, Pre-, Post-Order Traversal



In-order: 2, 5, 7, 9, 10, 15, 17, 20, 30

Pre-order:

Post-order:

Beauty is Only $O(\log n)$ Deep

Binary Search Trees are fast if they're shallow.

Know any shallow trees?

- ▶ perfectly complete
- ▶ perfectly complete except the last level (like a heap)
- ▶ anything else?


What matters here?

Siblings should have about the same height.

7 / 29

8 / 29

Balance



$$\text{balance}(x) = \text{height}(x.\text{left}) - \text{height}(x.\text{right})$$

$$\text{height}(\text{NULL}) = -1.$$

If for all nodes x ,


- ▶ $\text{balance}(x) = 0$ then perfectly balanced.
- ▶ $|\text{balance}(x)|$ is small then balanced enough.
- ▶ $-1 \leq \text{balance}(x) \leq 1$ then tree height $\leq c \lg n$ where $c < 2$.

9 / 29

AVL (Adelson-Velsky and Landis) Tree

Binary tree property

- ▶ each node has ≤ 2 children



Search tree property

- ▶ all keys in left subtree smaller than node's key
- ▶ all keys in right subtree larger than node's key


Balance property

- ▶ For all nodes x , $-1 \leq \text{balance}(x) \leq 1$

Result: height is $\Theta(\log n)$.


10 / 29

Is this an AVL tree?



11 / 29


An AVL Tree



12 / 29

How Do We Stay Balanced?

Suppose we start with a balanced search tree (an AVL tree),




It's no longer an AVL tree. What can we do?

ROTATE!


13 / 29

Rotation Animation




14 / 29

Rotation Animation




14 / 29

Rotation Animation




14 / 29

Rotation Animation




14 / 29

Rotation Animation



14 / 29


Single Rotation



15 / 29

Single Rotation

rotateRight is shown. There's also a symmetric rotateLeft.




After rotation, subtree's height is the same as before insert.
So heights of ancestors don't change.
So?

16 / 29

Double Rotation

Start with




and insert 6

A single rotation won't fix this.

DOUBLE ROTATE!

Double Rotation




17 / 29

18 / 29

Double Rotation

doubleRotateLeft is shown. There's also a symmetric doubleRotateRight.



insert occurred here or here

Either x or y increased to height $h - 1$ after insert.
After rotation, subtree's height is the same as before insert.
So height of ancestors doesn't change.

Insert Algorithm

1. Find location for new key.
2. Add new leaf node with new key.
3. Go up tree from new leaf searching for imbalance.
4. At lowest unbalanced ancestor:

Case LL: rotateRight

Case RR: rotateLeft

Case LR: doubleRotateRight

Case RL: doubleRotateLeft


The case names are the first two steps on the path from the unbalanced ancestor to the new leaf.

19 / 29

20 / 29


Insert: No Imbalance

Insert(3)



Insert: Imbalance Case RR


Insert(33)




21 / 29

22 / 29

Case RR: rotateLeft



Single Rotation Code




```
void rotateLeft( Node *&a ) {
    Node * b = a->right;
    a->right = b->left;
    b->left = a;
    updateHeight(a);
    updateHeight(b);
    a = b;
}
```

23 / 29


24 / 29

Insert: Imbalance Case RL

Insert(18)




Case RL: doubleRotateLeft



25 / 29

26 / 29

Double Rotation Code




```
void doubleRotateLeft( Node *&a ) {
    rotateRight(a->right);
    rotateLeft(a);
}
```

27 / 29

Delete

1. Delete as for general binary search tree. (This way we reduce the problem to deleting a node with 0 or 1 child.)
2. Go up tree from deleted node searching for imbalance (and fixing heights).
3. Fix **all** unbalanced ancestors (bottom-up)



28 / 29

Thinking about AVL trees

Observations

- ▶ AVL trees are binary search trees that allow only slight imbalance
- ▶ Worst-case $O(\log n)$ time for find, insert, and delete
- ▶ Elements (even siblings) may be scattered in memory

Realities

- ▶ For large data sets, disk accesses dominate runtime

Could we have perfect balance if we relax binary tree restriction?