

CSE373: Data Structures and Algorithms

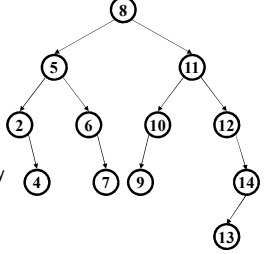
## AVL Trees

Steve Tanimoto  
Winter 2016

This lecture material represents the work of multiple instructors at the University of Washington. Thank you to all who have contributed!

### Review: Binary Search Tree (BST)

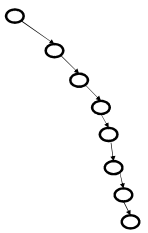
- **Structure property** (binary tree)
  - Each node has  $\leq 2$  children
  - Result: keeps operations simple
- **Order property**
  - All keys in left subtree smaller than node's key
  - All keys in right subtree larger than node's key
  - Result: easy to find any given key



Winter 2016 CSE373: Data Structures & Algorithms 2

### BST: Efficiency of Operations?

- Problem: operations may be inefficient if BST is unbalanced.
- Find, insert, delete
  - $O(n)$  in the worst case
- BuildTree
  - $O(n^2)$  in the worst case



Winter 2016 CSE373: Data Structures & Algorithms 3

### How can we make a BST efficient?

*Observation*

- BST: the shallower the better!

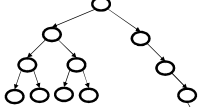
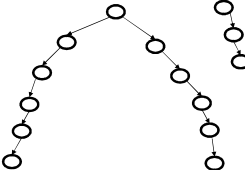
*Solution:* Require and maintain a **Balance Condition** that

1. Ensures depth is always  $O(\log n)$  - strong enough!
2. Is efficient to maintain - not too strong!

- When we **build** the tree, make sure it's balanced.
- **BUT...** Balancing a tree **only** at build time is insufficient because sequences of operations can eventually transform our carefully balanced tree into the **dreaded list** ☹️
- So, we also need to also **keep** the tree balanced as we perform operations.

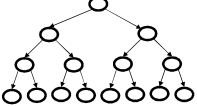
Winter 2016 CSE373: Data Structures & Algorithms 4

### Potential Balance Conditions

1. Left and right subtrees of the **root** have equal number of nodes
  - Too weak!**  
Height mismatch example:
2. Left and right subtrees of the **root** have equal height
  - Too weak!**  
Double chain example:

Winter 2016 CSE373: Data Structures & Algorithms 5

### Potential Balance Conditions

3. Left and right subtrees of **every node** have equal number of nodes
  - Too strong!**  
Only perfect trees ( $2^n - 1$  nodes)
4. Left and right subtrees of **every node** have equal height
  - Too strong!**  
Only perfect trees ( $2^n - 1$  nodes)

Winter 2016 CSE373: Data Structures & Algorithms 6

### The AVL Balance Condition

Left and right subtrees of *every node* have heights differing by at most 1

Definition:  $balance(node) = height(node.left) - height(node.right)$

AVL property: for every node  $x$ ,  $-1 \leq balance(x) \leq 1$

- Ensures small depth
  - Will prove this by showing that an AVL tree of height  $h$  must have a number of nodes exponential in  $h$  (i.e. height must be logarithmic in number of nodes)
- Efficient to maintain
  - Using single and double rotations

### The AVL Tree Data Structure

An AVL tree is a self-balancing binary search tree.


Structural properties

- Binary tree property (same as BST)
  - Order property (same as for BST)
1. Balance property:  
balance of every node is between -1 and 1

Result: Worst-case depth is  $O(\log n)$


- Named after inventors Adelson-Velskii and Landis (AVL)
  - First invented in 1962

Is this an AVL tree?



Yes! Because the left and right subtrees of every node have heights differing by at most 1


Is this an AVL tree?



Nope! The left and right subtrees of some nodes (e.g. 1, 4, 6) have heights that differ by more than 1

### The shallowness bound

- Let  $S(h)$  = the minimum number of nodes in an AVL tree of height  $h$
- $S(h)$  grows exponentially in  $h$ .
    - (Can be proved, but we will not do it in class.)
  - Therefore, a tree with  $n$  nodes has a logarithmic height
  - Thus FIND can be done in  $O(\log n)$  time.
  - We will also see that INSERT and DELETE can be done in  $O(\log n)$  time, while maintaining the AVL property.



### Implementing AVL Trees

Node structure

Tree operations

(We'll want to be sure these operate in  $O(\log n)$  worst case time.)

### An AVL Tree

Track height at all times!

Winter 2016 CSE373: Data Structures & Algorithms 13

### AVL tree operations

- **AVL find:**
  - Same as BST find
- **AVL insert:**
  - First BST insert, then check balance and potentially "fix" the AVL tree
  - Four different imbalance cases
- **AVL delete:**
  - The "easy way" is lazy deletion
  - Otherwise, do the deletion and then check for several imbalance cases (we will skip this)

Winter 2016 CSE373: Data Structures & Algorithms 14

### Insert: detect potential imbalance

1. Insert the new node as in a BST (a new leaf)
2. For each node on the path from the root to the new leaf, the insertion may (or may not) have changed the node's height
3. So after insertion in a subtree, detect height imbalance and perform a *rotation* to restore balance at that node

All the action is in defining the correct rotations to restore balance

Fact that an implementation can ignore:

- There must be a deepest element that is imbalanced after the insert (all descendants still balanced)
- After rebalancing this deepest node, every node is balanced
- So at most one node needs to be rebalanced

Winter 2016 CSE373: Data Structures & Algorithms 15

### Case #1: Example

Insert(6)  
Insert(3)  
Insert(1)

Third insertion violates balance property

- happens to be at the root

What is the only way to fix this?

Winter 2016 CSE373: Data Structures & Algorithms 16

### Fix: Apply "Single Rotation"

- **Single rotation:** The basic operation we'll use to rebalance
  - Move child of unbalanced node into parent position
  - Parent becomes the "other" child (always okay in a BST!)
  - Other subtrees move in only way BST allows (next slide)

AVL Property violated at node 6

Child's new-height = old-height-before-insert

Winter 2016 CSE373: Data Structures & Algorithms 17

### The example generalized

- Insertion into left-left grandchild causes an imbalance
  - 1 of 4 possible imbalance causes (other 3 coming up!)
- Creates an imbalance in the AVL tree (specifically **a** is imbalanced)

Winter 2016 CSE373: Data Structures & Algorithms 18

### The general left-left case

- So we rotate at **a**
  - Move child of unbalanced node into parent position
  - Parent becomes the "other" child
  - Other sub-trees move in the only way BST allows:
    - using BST facts:  $X < b < Y < a < Z$

- A single rotation restores balance at the node
  - To same height as before insertion, so ancestors now balanced

Winter 2016 CSE373: Data Structures & Algorithms 19

### Another example: insert (16)

Winter 2016 CSE373: Data Structures & Algorithms 20

### The general right-right case

- Mirror image to left-left case, so you rotate the other way
  - Exact same concept, but need different code

Winter 2016 CSE373: Data Structures & Algorithms 21

### Two cases to go

Unfortunately, single rotations are not enough for insertions in the left-right subtree or the right-left subtree

Simple example: `insert(1), insert(6), insert(3)`

- First wrong idea: single rotation like we did for left-left

Winter 2016 CSE373: Data Structures & Algorithms 22

### Two cases to go

Unfortunately, single rotations are not enough for insertions in the left-right subtree or the right-left subtree

Simple example: `insert(1), insert(6), insert(3)`

- Second wrong idea: single rotation on the child of the unbalanced node

Winter 2016 CSE373: Data Structures & Algorithms 23

### Sometimes two wrongs make a right ☺

- First idea violated the order property
- Second idea didn't fix balance
- But if we do both single rotations, starting with the second, it works! (And not just for this example.)
- Double rotation:
  - Rotate problematic child and grandchild
  - Then rotate between self and new child

Winter 2016 CSE373: Data Structures & Algorithms 24

### The general right-left case

Winter 2016 CSE373: Data Structures & Algorithms 25

### Comments

- Like in the left-left and right-right cases, the height of the subtree after rebalancing is the same as before the insert
  - So no ancestor in the tree will need rebalancing
- Does not have to be implemented as two rotations; can just do:
  - Move *c* to grandparent's position
  - Put *a*, *b*, *X*, *U*, *V*, and *Z* in the only legal positions for a BST

Easier to remember than you may think:  
 Move *c* to grandparent's position  
 Put *a*, *b*, *X*, *U*, *V*, and *Z* in the only legal positions for a BST

Winter 2016 CSE373: Data Structures & Algorithms 26

### The last case: left-right

- Mirror image of right-left
  - Again, no new concepts, only new code to write

Winter 2016 CSE373: Data Structures & Algorithms 27

### Insert, summarized

- Insert as in a BST
- Check back up path for imbalance, which will be 1 of 4 cases:
  - Node's left-left grandchild is too tall
  - Node's left-right grandchild is too tall
  - Node's right-left grandchild is too tall
  - Node's right-right grandchild is too tall
- Only one case occurs because tree was balanced before insert
- After the appropriate single or double rotation, the smallest-unbalanced subtree has the same height as before the insertion
  - So all ancestors are now balanced

Winter 2016 CSE373: Data Structures & Algorithms 28

### Now efficiency

- Worst-case complexity of `find`:  $O(\log n)$ 
  - Tree is balanced
- Worst-case complexity of `insert`:  $O(\log n)$ 
  - Tree starts balanced
  - A rotation is  $O(1)$  and there's an  $O(\log n)$  path to root
  - Tree ends balanced
- Worst-case complexity of `buildTree`:  $O(n \log n)$

Takes some more rotation action to handle `delete`...

Winter 2016 CSE373: Data Structures & Algorithms 29

### Pros and Cons of AVL Trees

Arguments for AVL trees:

- All operations logarithmic worst-case because trees are *always* balanced
- Height balancing adds no more than a constant factor to the speed of `insert` and `delete`

Arguments against AVL trees:

- Difficult to program & debug [but done once in a library!]
- More space for height field
- Asymptotically faster but rebalancing takes a little time
- If "amortized" logarithmic time is enough, use "splay trees."

Winter 2016 CSE373: Data Structures & Algorithms 30