

AVL Trees

CSE 332 Spring 2020

Instructor: Hannah C. Tang

Teaching Assistants:

Annie Mao

Howard Xiao

Lucy Jiang

Chris Choi

Jade Watkins

Nathan Lipiarski

Ethan Knutson

Khushi Chaudhari

Richard Jiang

Announcements

- ❖ Exercise 5 due Monday – we hope you've started!

Learning Objectives

- ❖ Be able to implement and prove runtimes for BST operations
- ❖ Articulate the differences between BSTs and AVLs, and make intelligent tradeoffs between them
- ❖ Be able to implement and prove runtimes for AVL find() and add()

Lecture Outline

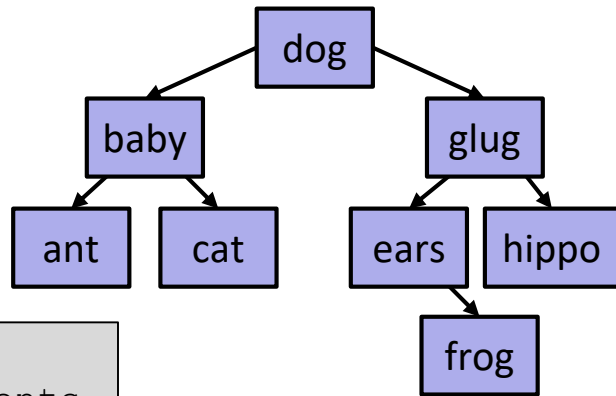
- ❖ Binary Search Trees as Dictionary/Set Data Structures (cont.)
 - Add
 - Remove

- ❖ AVL Tree
 - *(Proving the AVL tree's height bound)*
 - Find
 - Add

Binary Search Trees: Add

- ❖ Where does the new item belong?
- ❖ How do we use BST invariants to ensure the leaf is added correctly?

new leaf!



```

BSTNode add(BSTNode t, Item i) {
    // Implement by putting statements
    // in the correct order

```

D B C A } both will work!
D C B A }

A	B	C	D
<pre>return t;</pre>	<pre>if (k < i.key) { t.left = add(t.left, i); }</pre>	<pre>if (k > i.key) { t.right = add(t.right, i); }</pre>	<pre>if (t == null) { return new BSTNode(i); }</pre>

Lecture Outline

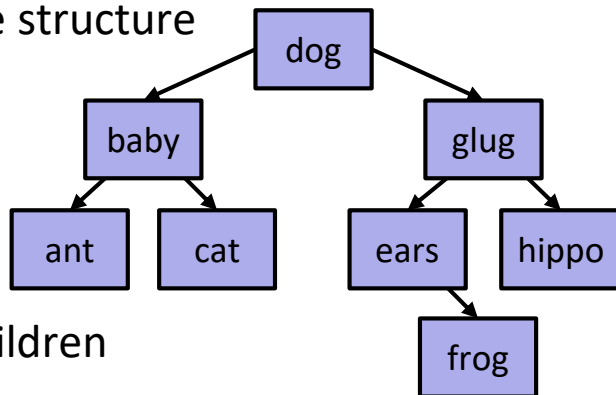
- ❖ Binary Search Trees as Dictionary/Set Data Structures (cont.)
 - Add
 - **Remove**

- ❖ AVL Tree
 - *(Proving the AVL tree's height bound)*
 - Find
 - Add

Binary Search Trees: Remove

- ❖ Removing an item disrupts the tree structure

- **find** the node to be removed
- Remove it
- “Fix” the tree so that it is still a BST



- ❖ 3 cases based on the number of children

1. Node has no children
2. Node has one child
3. Node has two children

- ❖ In each case, we must maintain the **BST Ordering!**


Reminder: a dictionary maps *keys* to *values*;
an *item* or *data* refers to the (key, value) pair

BST Remove: Case #1: Leaf

- ❖ Remove the node with the key **hippo**

```
BSTNode remove(BSTNode n) {
```

```
}
```




BST Remove: Case #2: One Child

- ❖ Remove the node with the key **ears**
 - What does the BST invariant say about the descendant's keys?

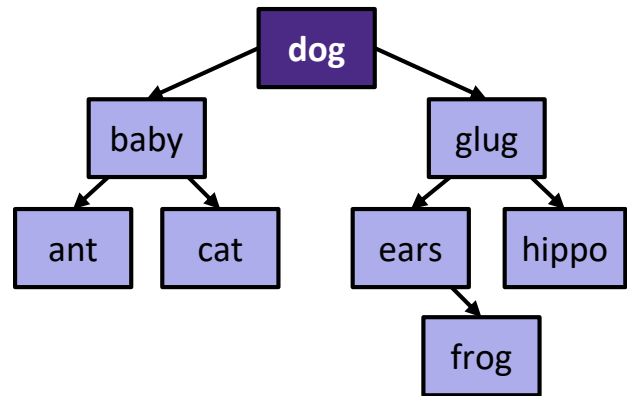
```
BSTNode remove(BSTNode n) {
```

```
}
```



BST Remove: Case #3: Two Children

- ❖ Remove the node with the key **dog**
- ❖ The replacement node's key:
 - Must be $>$ than all keys in left subtree
 - Must be $<$ than all keys in right subtree



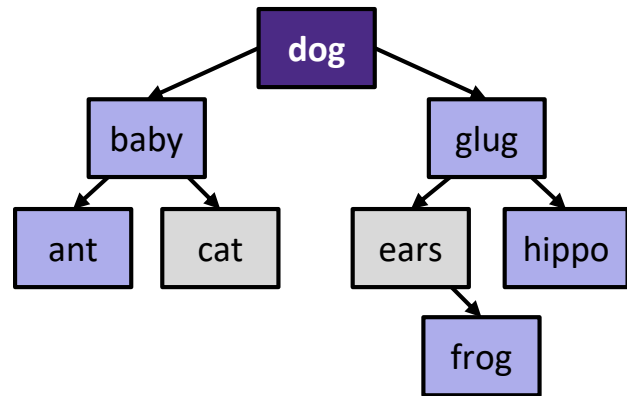
in-order traversal: ant, baby, cat, dog, ears, frog, glug, hippo
("flattened tree")

BST Remove: Case #3: Two Children

❖ Remove the node with the key **dog**

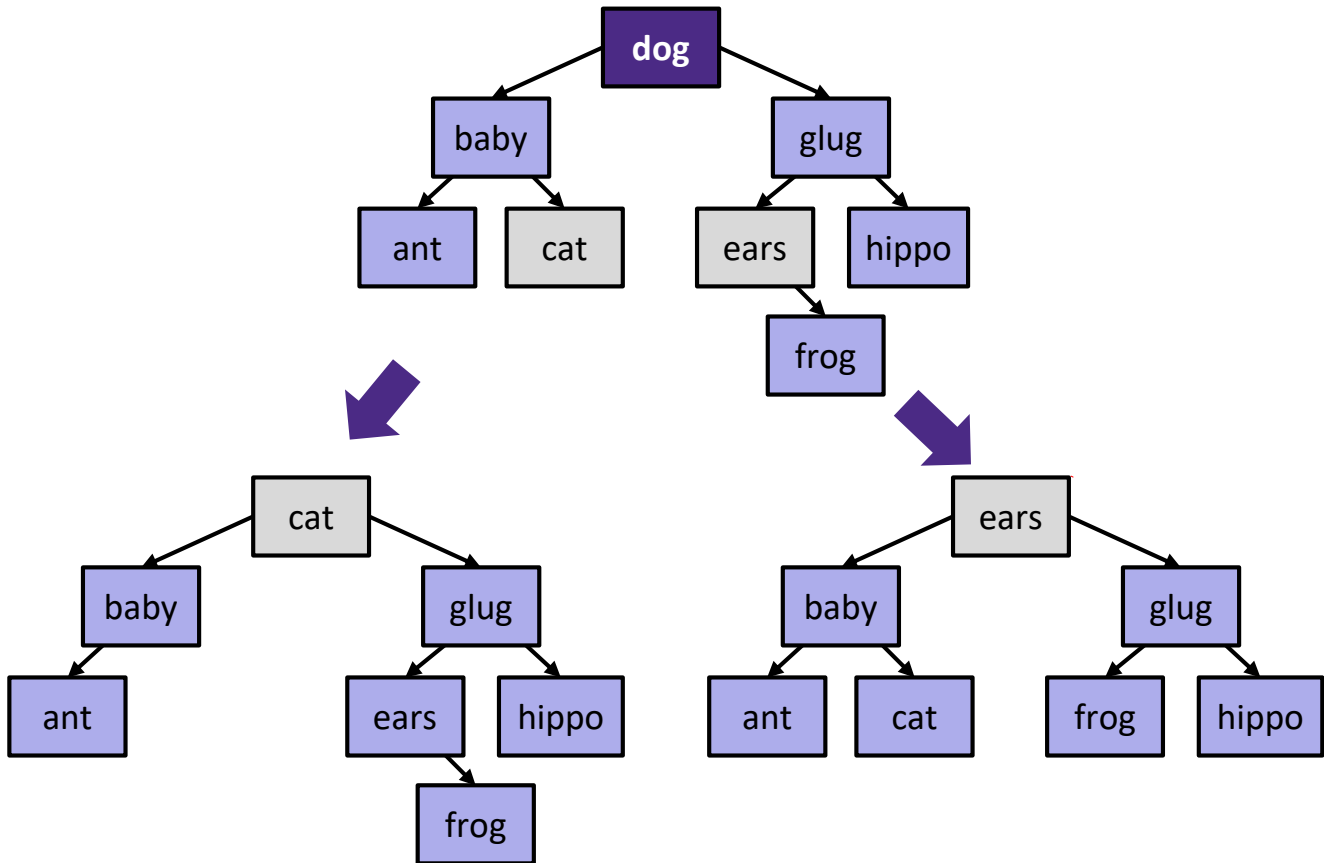
❖ The replacement node's key:

- Must be $>$ than all keys in left subtree: **predecessor** (**cat**)
- Must be $<$ than all keys in right subtree: **successor** (**ears**)



❖ The predecessor or successor have either 0 or 1 children


BST Remove: Case #3: Two Children



Aside: Finding the largest (or smallest) node

- ❖ The predecessor is the largest item in the left subtree
- ❖ The successor is the smallest item in the right subtree
- ❖ How do you find the largest (and smallest) item in a tree?
 - Remember that subtrees are trees too

```
BSTNode largest(BSTNode n) {  
    while (n.right != null) {  
        n = n.right;  
    }  
    return n;  
}
```



BST Summary

- ❖ Binary Search Trees implement both the Set and Dictionary ADTs
- ❖ Binary Search Trees are *recursively defined*

	LinkedList Dictionary, Worst Case	BST Dictionary, Average Case <i>Randomized</i>	BST Dictionary, Worst Case
Find	$\Theta(N)$	$\Theta(h)$ aka $\Theta(\log N)$	$\Theta(h)$ aka $\Theta(N)$
Add	$\Theta(N)$	$\Theta(h)$ aka $\Theta(\log N)$	$\Theta(h)$ aka $\Theta(N)$
Remove	$\Theta(N)$	$\Theta(h)$ aka $\Theta(\log N)$	$\Theta(h)$ aka $\Theta(N)$




Lecture Outline

- ❖ Binary Search Trees as Dictionary/Set Data Structures (cont.)
 - Add
 - Remove

- ❖ **AVL Tree**
 - *(Proving the AVL tree's height bound)*
 - Find
 - Add

buildTree for Binary Search Trees

- ❖ We had **buildHeap**, so let's consider **buildTree**
- ❖ Insert keys 1, 2, 3, 4, 5, 6, 7, 8, 9 into an empty BST
 - If inserted in given order, what is the resultant tree?
 - What is the big-O runtime for sorted input?
 - Is inserting in the reverse order any better?



buildTree Runtime: $O(n^2)$

(not a happy place)

Balancing a BST

Observation

- ❖ BST: the shallower the better!
- ❖ For a BST with n items inserted in arbitrary order
 - ~~Average~~ ^{Randomized} height is $O(\log n)$ – see text for proof
 - Worst case height is $O(n)$
- ❖ Simple cases such as inserting in order lead to worst-case!


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

- ❖ ensures depth is always $O(\log n)$
- ❖ is easy to maintain

Potential BST Balance Conditions


- ❖ Left and right subtrees of the *root* have equal number of nodes

Too weak!
Height mismatch example:




- ❖ Left and right subtrees of the *root* have equal *height*

Too weak!
Double chain example:



The AVL Balance Condition (1 of 2)

- ❖ Left and right subtrees of the *root* have equal number of nodes
- ❖ Left and right subtrees of the *root* have equal *height*



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

The AVL Balance Condition (2 of 2)

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*, $-1 \leq \text{balance}(x) \leq 1$

Results:

- ❖ Ensures shallow depth: $h \in O(\log n)$
 - Will prove this by showing that an AVL tree of height h must have a number of nodes *exponential* in h
- ❖ Efficient to maintain using rotations


The AVL Tree Data Structure (1 of 2)

❖ Structural properties


- Binary tree property (0, 1, or 2 children)
- Heights of left and right subtrees for every node differ by at most 1

❖ Ordering property

- Same as for BST




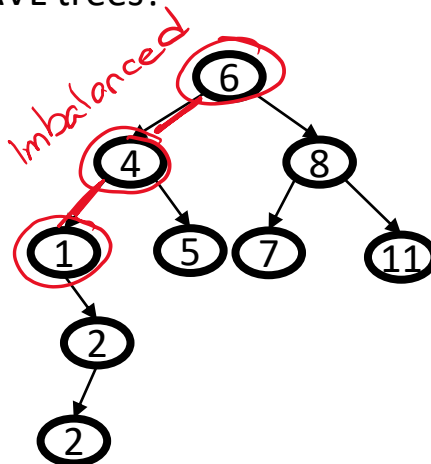

The AVL Tree Data Structure (2 of 2)




Poll Everywhere

pollev.com/cse332

❖ Are the following trees AVL trees?




- A. No / No / No
- B. Yes / No / No
- C. Yes / Yes / No
- D. Yes / Yes / Yes
- E. Yes / No / Yes
- F. I'm not sure ...



Height of an AVL Tree? (1 of 2)

- ❖ The “best case” AVL tree is a perfect tree




- ❖ What does the “worst case” AVL tree look like? *Spindlet!*
- ❖ Let $S(h)$ be the minimum # of nodes in a tree of height h :
 - Where $S(-1) = 0$ and $S(0) = 1$


Minimal AVL Tree (height = 0)




Minimal AVL Tree (height = 1)



Minimal AVL Tree (height = 2)



Minimal AVL Tree (height = 3)



Height of an AVL Tree? (2 of 2)

- ❖ Using the AVL balance property, we can determine the minimum number of nodes in an AVL tree of height h

- ❖ Let $\mathbf{S}(h)$ be the minimum # of nodes in a tree of height h

$$\rightarrow \mathbf{S}(h) = \mathbf{S}(h-1) + \mathbf{S}(h-2) + 1$$

where $\mathbf{S}(-1) = 0$ and $\mathbf{S}(0) = 1$

plus the root node

- ❖ Solution of Recurrence: $\mathbf{S}(h) \approx 1.62^h$

$$n \approx 1.62^h$$
$$\log_{1.62} n \approx h$$



Lecture Outline

- ❖ Binary Search Trees as Dictionary/Set Data Structures (cont.)
 - Add
 - Remove

- ❖ AVL Tree
 - *(Proving the AVL tree's height bound)*
 - Find
 - Add

Before We Prove It

- ❖ Good intuition from plots comparing:
 - $S(h)$ computed directly from the definition
 - $((1+\sqrt{5})/2)^h$
- ❖ $S(h)$ is always bigger, up to trees with huge numbers of nodes
 - Graphs aren't proofs, so let's prove it




The Proof Outline

Let $S(h)$ = the minimum number of nodes in an AVL tree of height h

- If we can prove that $S(h)$ grows exponentially in h , then a tree with n nodes has a logarithmic height

❖ Step 1: Define $S(h)$ inductively using AVL property

- $S(-1)=0, S(0)=1, S(1)=2$
- For $h \geq 1, S(h) = 1+S(h-1)+S(h-2)$




❖ Step 2: Show this recurrence grows really fast

- Similar to Fibonacci numbers
- Can prove for all $h, S(h) > \phi^h - 1$ where ϕ is the golden ratio, $(1+\sqrt{5})/2$, about 1.62
- Growing faster than 1.6^h is “plenty exponential”

The Golden Ratio

$$\phi = \frac{1 + \sqrt{5}}{2} \approx 1.62$$



This is a special number

- Aside: Since the Renaissance, many artists and architects have proportioned their work (e.g., length:height) to approximate the *golden ratio*: If $(a+b)/a = a/b$, then $a = \phi b$
- We will need one special arithmetic fact about ϕ :

$$\begin{aligned}
 \phi^2 &= ((1+5^{1/2})/2)^2 \\
 &= (1 + 2*5^{1/2} + 5) / 4 \\
 &= (6 + 2*5^{1/2}) / 4 \\
 &= (3 + 5^{1/2}) / 2 \\
 &= 1 + (1 + 5^{1/2}) / 2 \\
 &= 1 + \phi
 \end{aligned}$$

The Proof

$$S(-1)=0, S(0)=1, S(1)=2$$

$$\text{For } h \geq 1, S(h) = 1+S(h-1)+S(h-2)$$

Theorem: For all $h \geq 0$, $S(h) > \phi^h - 1$

Proof: By induction on h

Base cases:

$$S(0) = 1 > \phi^0 - 1 = 0$$

$$S(1) = 2 > \phi^1 - 1 \approx 0.62$$

Inductive case ($k > 1$):

Show $S(k+1) > \phi^{k+1} - 1$ assuming $S(k) > \phi^k - 1$ and $S(k-1) > \phi^{k-1} - 1$

$$S(k+1) = 1 + S(k) + S(k-1)$$

by definition of S

$$> 1 + \phi^k - 1 + \phi^{k-1} - 1$$

by induction

$$= \phi^k + \phi^{k-1} - 1$$

by arithmetic ($1-1=0$)

$$= \phi^{k-1} (\phi + 1) - 1$$

by arithmetic (factor ϕ^{k-1})

$$= \phi^{k-1} \phi^2 - 1$$

by special property of ϕ

$$= \phi^{k+1} - 1$$

by arithmetic (add exponents)

Lecture Outline

- ❖ Binary Search Trees as Dictionary/Set Data Structures (cont.)
 - Add
 - Remove

- ❖ AVL Tree
 - *(Proving the AVL tree's height bound)*
 - **Find**
 - Add

AVL Find


- ❖ Surprise! You already know this one

Same implementation, but better runtime!

And Now for Some Bad News ...

- ❖ find() is $O(\log n)$!
- ❖ But as we add() and remove elements(), we need to:
 - Track heights
 - Detect imbalance
 - Restore balance

*Is this tree AVL-balanced?
How about after insert(8)?*

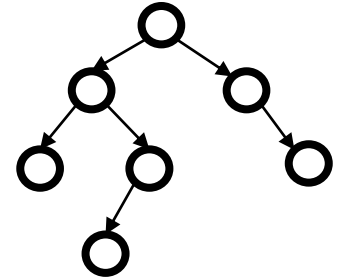


Lecture Outline

- ❖ Binary Search Trees as Dictionary/Set Data Structures (cont.)
 - Add
 - Remove

- ❖ AVL Tree
 - *(Proving the AVL tree's height bound)*
 - Find
 - **Add**

AVL add(): Overall Approach




- ❖ Our overall algorithm looks like:
 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
 - Detect height imbalance and perform a *rotation* to restore balance

- ❖ Fact that makes it a bit easier:
 - Imbalances only occur along the path from the new leaf to the root
 - There must be a deepest element that is unbalanced
 - After rebalancing this deepest node, every node above it is also rebalanced
 - Therefore, at most one node needs to be rebalanced

AVL add(): Cases

- ❖ Let b be the node where an imbalance occurs
- ❖ There are four cases to consider. The insertion is in the:
 1. left subtree of the left child of b
 2. right subtree of the left child of b
 3. left subtree of the right child of b
 4. right subtree of the right child of b



Case #1: Example

add(6)


add(3)

add(1)

The insertion is in the:

1. left subtree of the left child of b
2. right subtree of the left child of b
3. left subtree of the right child of b
4. right subtree of the right child of b


- ❖ Last add() violates balance property
- ❖ What is the only way to fix this?




Case #1 Fix: Apply “Single Rotation”

❖ *Single rotation:*

- Move child of unbalanced node into parent position
- Parent becomes the “other” child



Case #1: Pseudocode




```
void RotateWithLeftChild(Node root) {
    Node temp = root.left
    root.left = temp.right
    temp.right = root
    root.height = max(root.right.height(),
                      root.left.height()) + 1
    temp.height = max(temp.right.height(),
                      temp.left.height()) + 1
    root = temp
}
```

RotateWithLeftChild rotates the tree clockwise

Case #1: Why It Works (1 of 2)

Oval: a node in the tree
Triangle: a subtree

- ❖ Node is imbalanced due to insertion *somewhere* in **left-left grandchild**
- ❖ First we did the insertion, which would make **b** imbalanced




Case #1: Why It Works (2 of 2)

❖ So we rotate at b , maintaining BST order: $X < a < Y < b < Z$

❖ Result:


- A single rotation restores balance at the formerly-imbalanced node
- Height is same as before insertion, so ancestors now balanced



Case #1: Another Example: add(16)

The insertion is in the:


1. left subtree of the left child of b
2. right subtree of the left child of b
3. left subtree of the right child of b
4. right subtree of the right child of b



Case #1: Another Example: add(16)

The insertion is in the:


1. left subtree of the left child of b
2. right subtree of the left child of b
3. left subtree of the right child of b
4. right subtree of the right child of b



Case #1 ≈ Case #4

- The insertion is in the:
1. left subtree of the left child of *b*
 2. right subtree of the left child of *b*
 3. left subtree of the right child of *b*
 4. right subtree of the right child of *b*

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



RotateWithRightChild rotates the tree counter-clockwise