

CSE 250

Data Structures

Dr. Eric Mikida
epmikida@buffalo.edu
208 Capen Hall

Red-Black Trees

BST Operations

Operation	Runtime
<code>find</code>	$O(d)$
<code>insert</code>	$O(d)$
<code>remove</code>	$O(d)$

What is the runtime in terms of n ? $O(n)$

$$\log(n) \leq d \leq n$$

AVL Trees

An **AVL tree** (**Adelson-**V**elsky and **L**andis) is a ***BST*** where every subtree is depth-balanced**

Remember: Tree depth = height(root)

Balanced: $|\text{height}(\text{root.left}) - \text{height}(\text{root.right})| \leq 1$

AVL Trees

Define $\text{balance}(v) = \text{height}(v.\text{right}) - \text{height}(v.\text{Left})$

Goal: Maintaining $\text{balance}(v) \in \{-1, 0, 1\}$

- $\text{balance}(v) = 0 \rightarrow$ "**v** is balanced"
- $\text{balance}(v) = -1 \rightarrow$ "**v** is left-heavy"
- $\text{balance}(v) = 1 \rightarrow$ "**v** is right-heavy"

AVL Trees - Depth Bounds

Question: Does the AVL property result in any guarantees about depth?


YES! Depth balance forces a maximum possible depth of $\log(n)$

AVL Trees - Enforcing the Depth Bound

Key Observations:


- Adding a node to an AVL tree can increase subtree height by at most 1
- Removing a node can decrease subtree height by at most 1
- Both of these modifications only affect ancestors
- A rotation maintains ordering, and changes tree height by at most +/-1

Enforcing the AVL Constraint: Case 1




How can we fix this?

Enforcing the AVL Constraint: Case 1



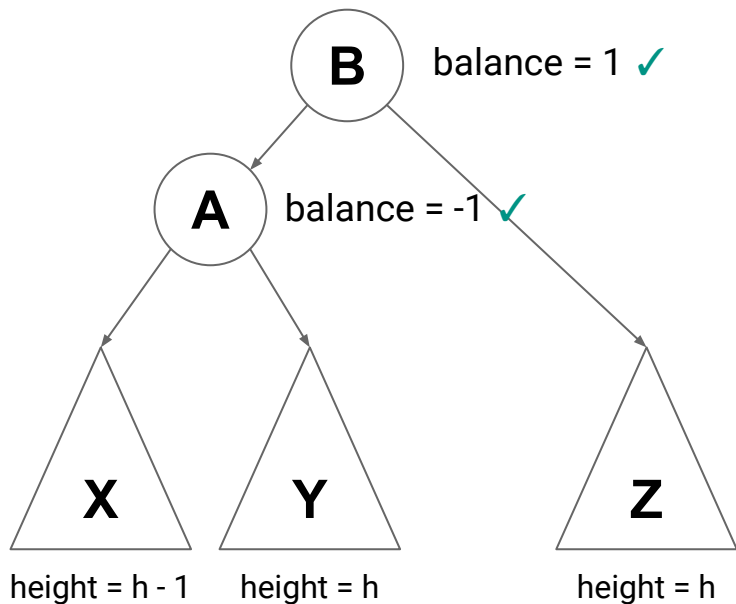
How can we fix this? `rotate(A,B)`

Enforcing the AVL Constraint: Case 2



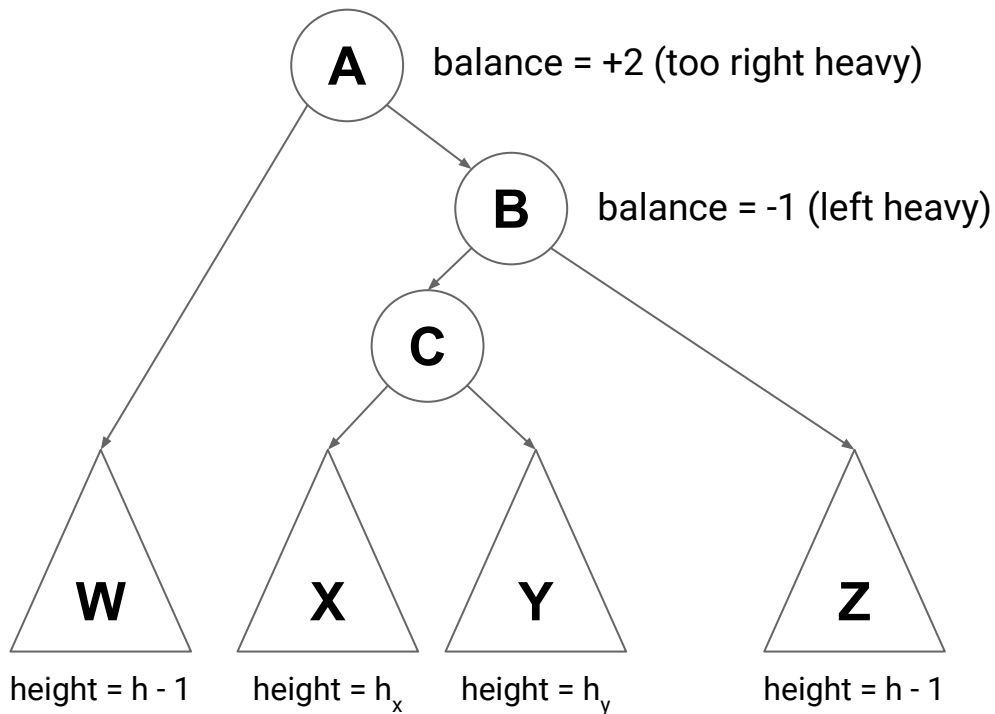
How can we fix this?

Enforcing the AVL Constraint: Case 2



How can we fix this? `rotate(A,B)`

Enforcing the AVL Constraint: Case 3




How can we fix this?

Height of **C** we know must be h

Therefore At least one of h_x or h_y must be $h - 1$

The other can also be $h - 2$, or $h - 1$

Enforcing the AVL Constraint: Case 3



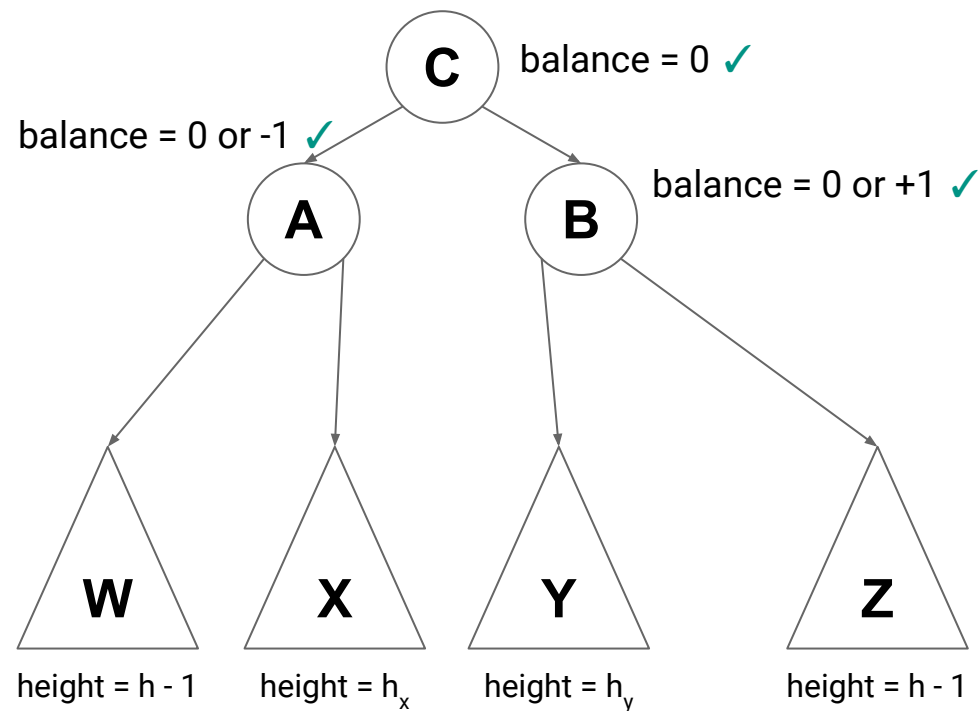
How can we fix this?
Rotate right first: `rotate(B, C)`

Height of **C** we know must be h

Therefore At least one of h_x or h_y must be $h - 1$

The other can also be $h - 2$, or $h - 1$

Enforcing the AVL Constraint: Case 3



How can we fix this?
Rotate right first: `rotate(B, C)`
Then right left: `rotate(A, C)`

Height of **C** we know must be h

Therefore At least one of h_x or h_y must be $h - 1$

The other can also be $h - 2$, or $h - 1$

Enforcing the AVL Constraint

- If too right heavy (balance == +2)
 - If right child is right heavy (balance == +1) or balanced (balance == 0)
 - rotate left around the root
 - If right child is left heavy (balance == -1)
 - rotate right around root of right child, then rotate left around root
- If too left heavy (balance == -2)
 - Same as above but flipped

Therefore if we have a balance factor that is off, but all children are AVL trees, we can fix the balance factor in at most 2 rotations

Inserting Records

To insert a record into an AVL Tree:

1. Find the insertion point (remember it is a BST) $O(d) = O(\log n)$
2. Insert the new leaf and set balance factor to 0 $O(1)$
3. Trace path back up to root and update balance factors $O(d) = O(\log n)$
 - a. If a balance factor becomes +/-2 then rotate to fix $O(1)$

Inserting Records

```
def insert[K, V](key: K, value: V, root: AVLNode[K, V]): Unit = {
```

```
  var node = findInsertionPoint(key, root)
  node._key = key; node._value = value
  node._isLeftHeavy = node._isRightHeavy = false
```

Find insertion point and create the new
leaf $O(d) = O(\log n)$

```
  while(node._parent.isDefined){ ←
```

$O(d) = O(\log n)$ iterations

```
    if(node._parent._left == node){
      if(node._parent._isRightHeavy){
        node._parent._isRightHeavy = false; return
      } else if(node._parent._isLeftHeavy) {
        if(node._isLeftHeavy){ node._parent.rotateRight() }
        else { node._parent.rotateLeftRight() }
        return
      } else {
        node._parent._isLeftHeavy = true
      }
    } else { /* symmetric to above */ }
```

$O(1)$ per iteration

```
    node = node._parent
```

```
  }
```

```
}
```


AVL Tree

What was our initial goal?

AVL Tree

What was our initial goal? **To constrain the depth of the tree**

AVL Tree

What was our initial goal? **To constrain the depth of the tree**

How did we accomplish it?

AVL Tree

What was our initial goal? **To constrain the depth of the tree**

How did we accomplish it? **By keeping the tree balanced
(subtree heights within 1 of each other)**

AVL Tree

What was our initial goal? **To constrain the depth of the tree**

How did we accomplish it? **By keeping the tree balanced
(subtree heights within 1 of each other)**

This approach is indirect, and a bit more restrictive than it has to be

Maintaining Balance - Another Approach


Enforcing height-balance is too strict (May do “unnecessary” rotations)

Weaker (and more direct) restriction:

- Balance the depth of EmptyTree nodes
- If ***a***, ***b*** are EmptyTree nodes, then enforce that for all ***a***, ***b***:
 - $\text{depth}(\mathbf{a}) \geq (\text{depth}(\mathbf{b}) \div 2)$
 - or
 - $\text{depth}(\mathbf{b}) \geq (\text{depth}(\mathbf{a}) \div 2)$

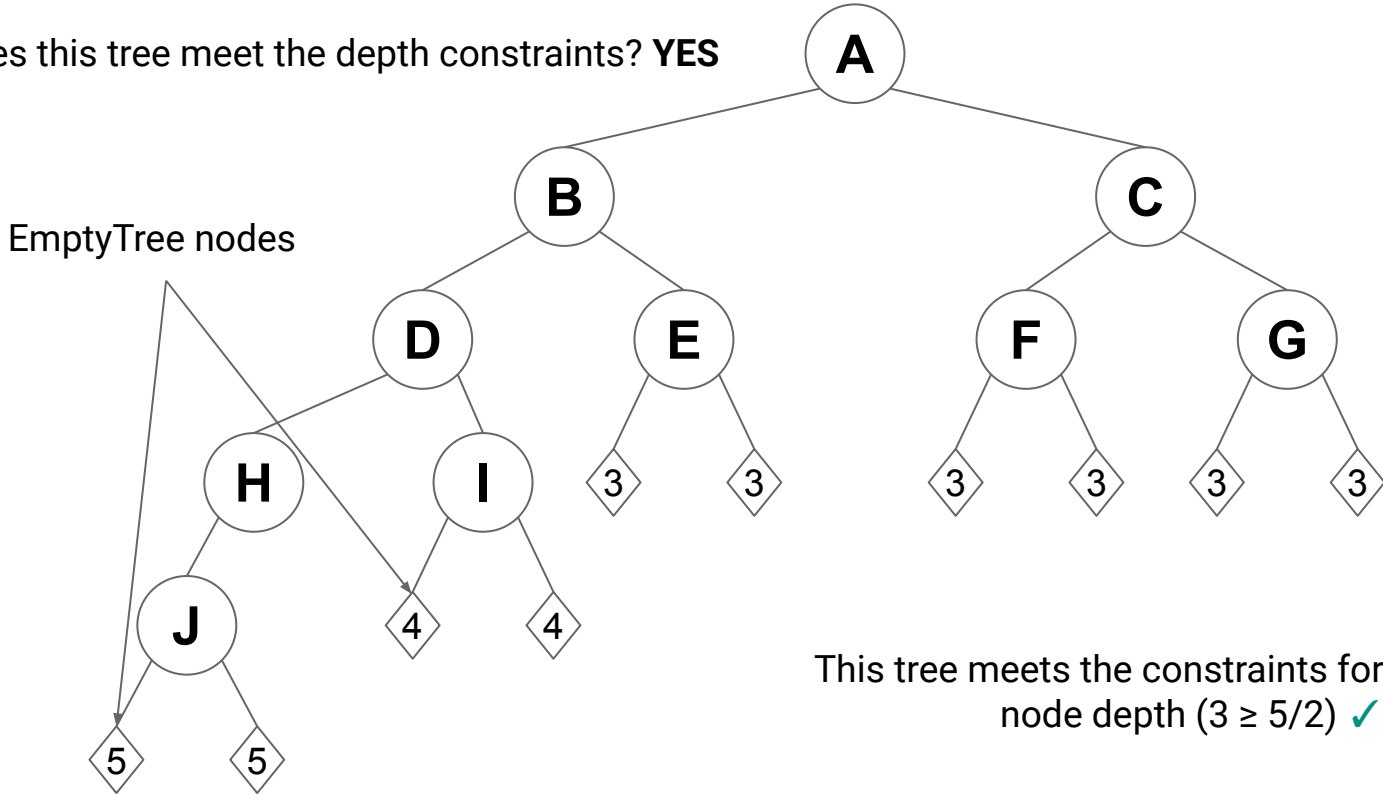
Depth Balancing

Does this tree meet the depth constraints?



Depth Balancing


Does this tree meet the depth constraints? **YES**



This tree meets the constraints for EmptyTree node depth ($3 \geq 5/2$) ✓

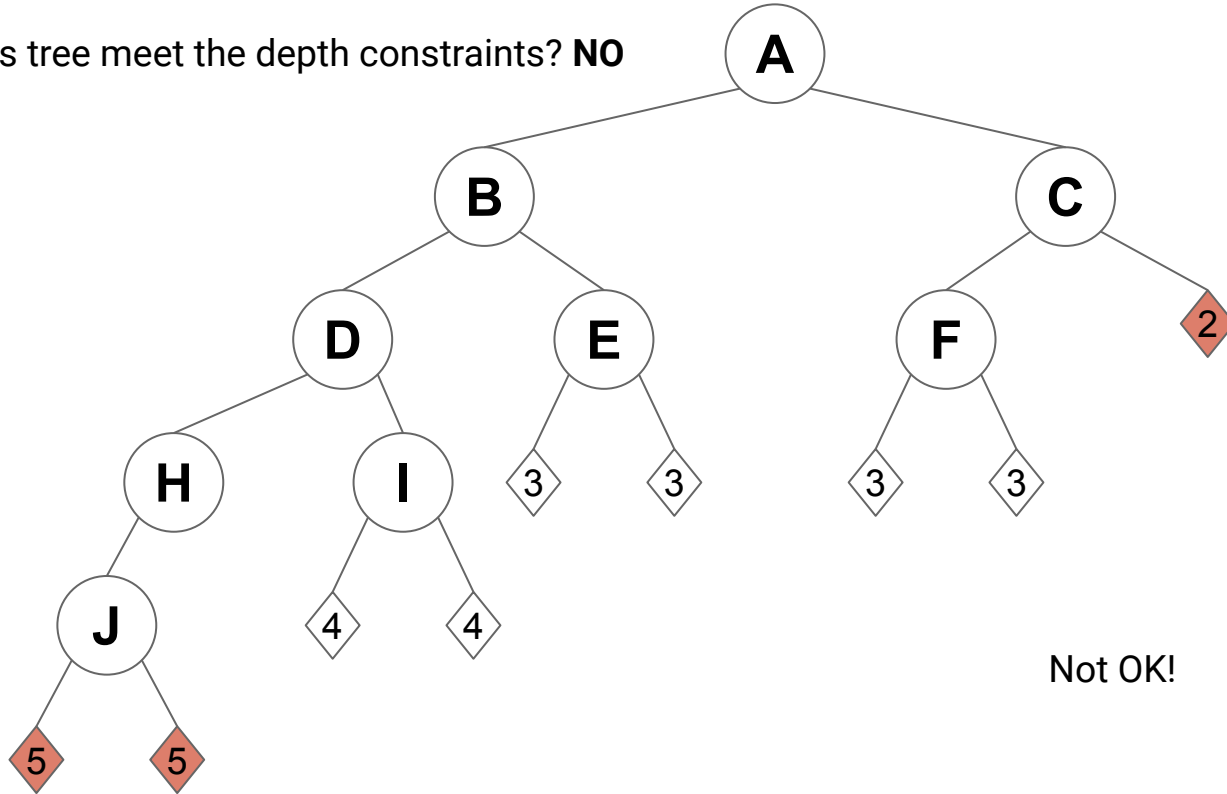
Depth Balancing

Does this tree meet the depth constraints?




Depth Balancing

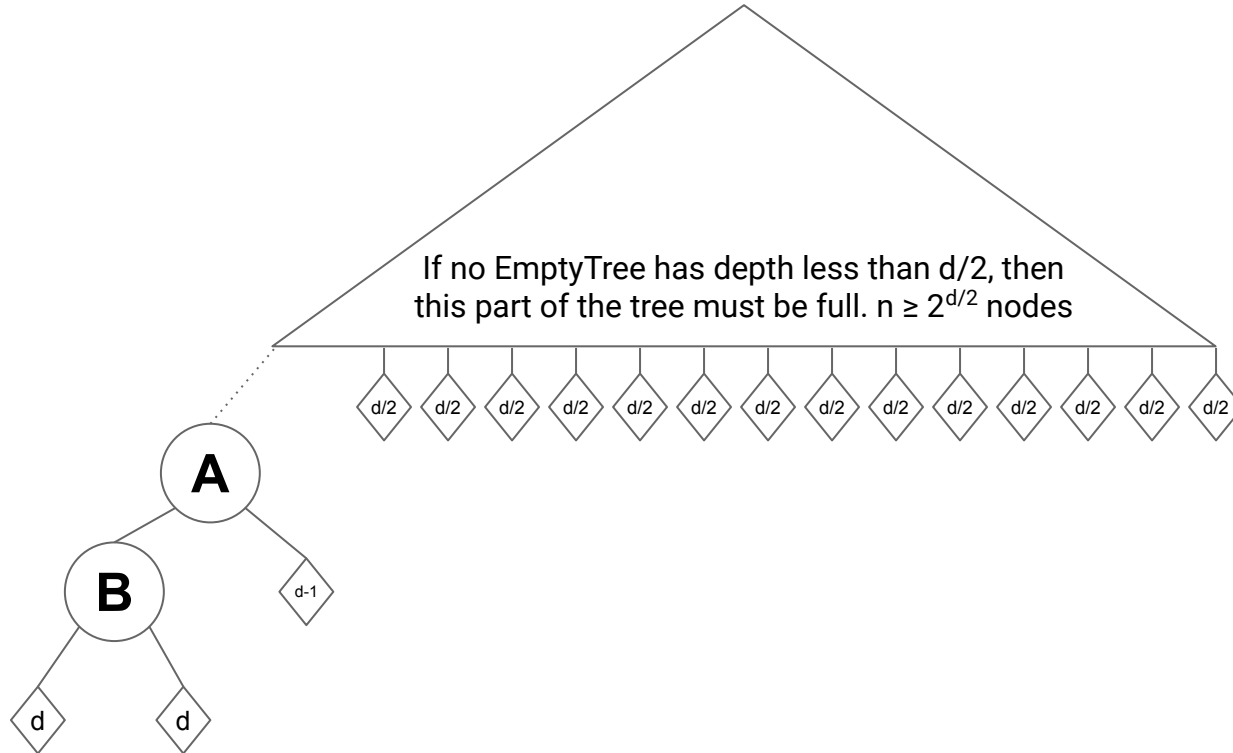
Does this tree meet the depth constraints? **NO**



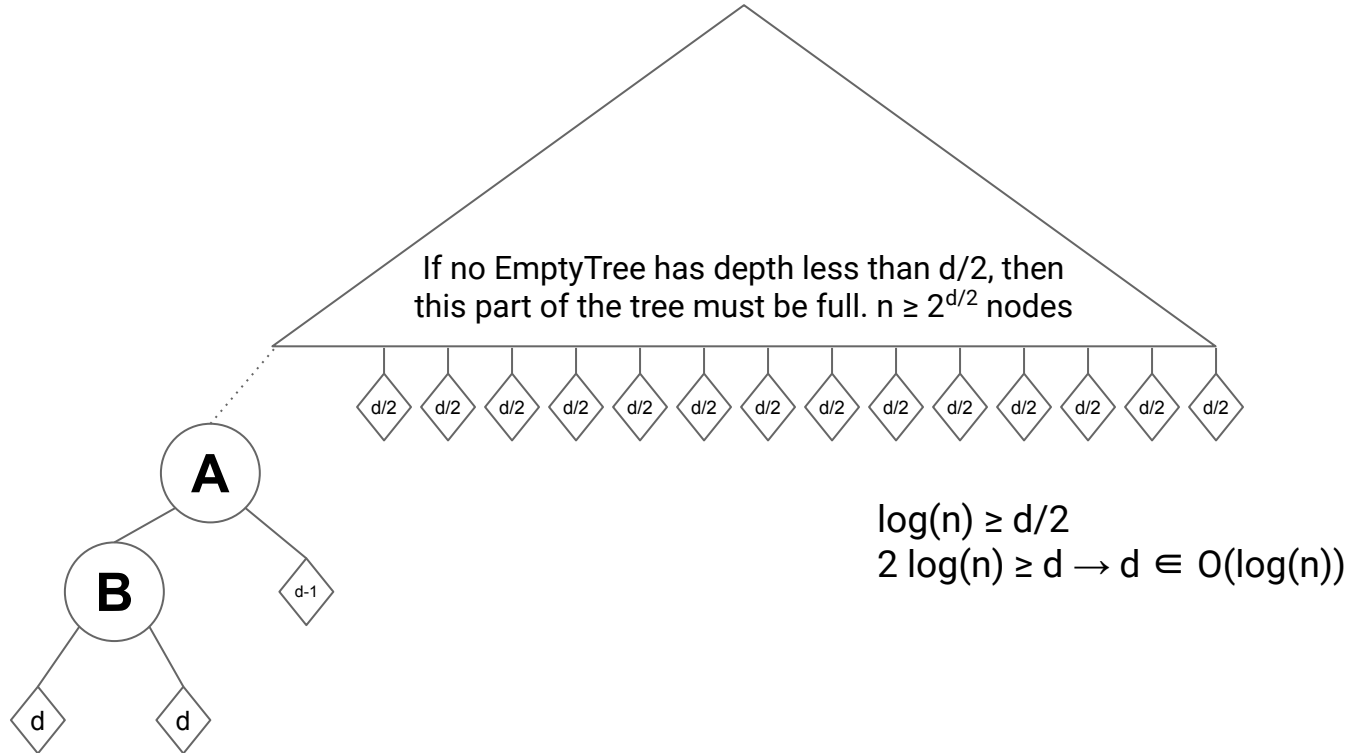
Depth Balancing



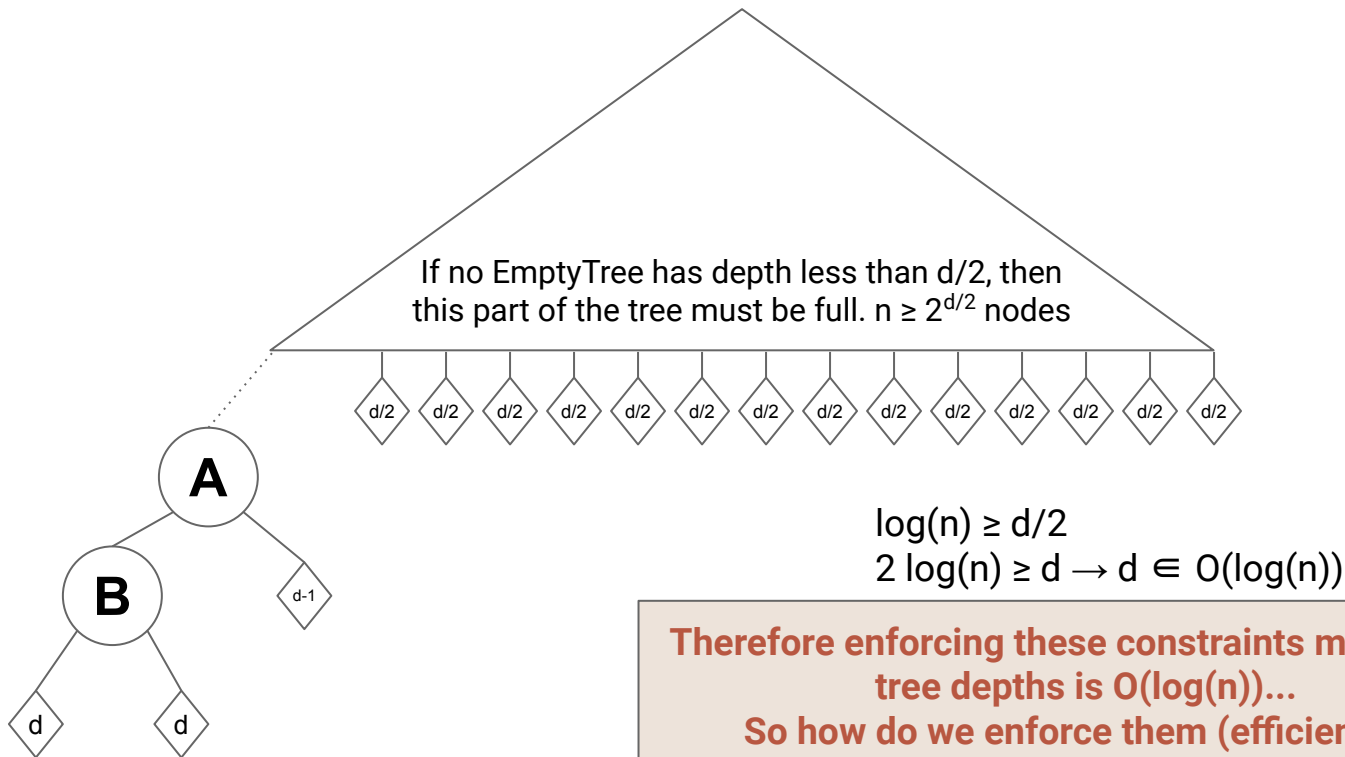
Depth Balancing



Depth Balancing



Depth Balancing




Red-Black Trees

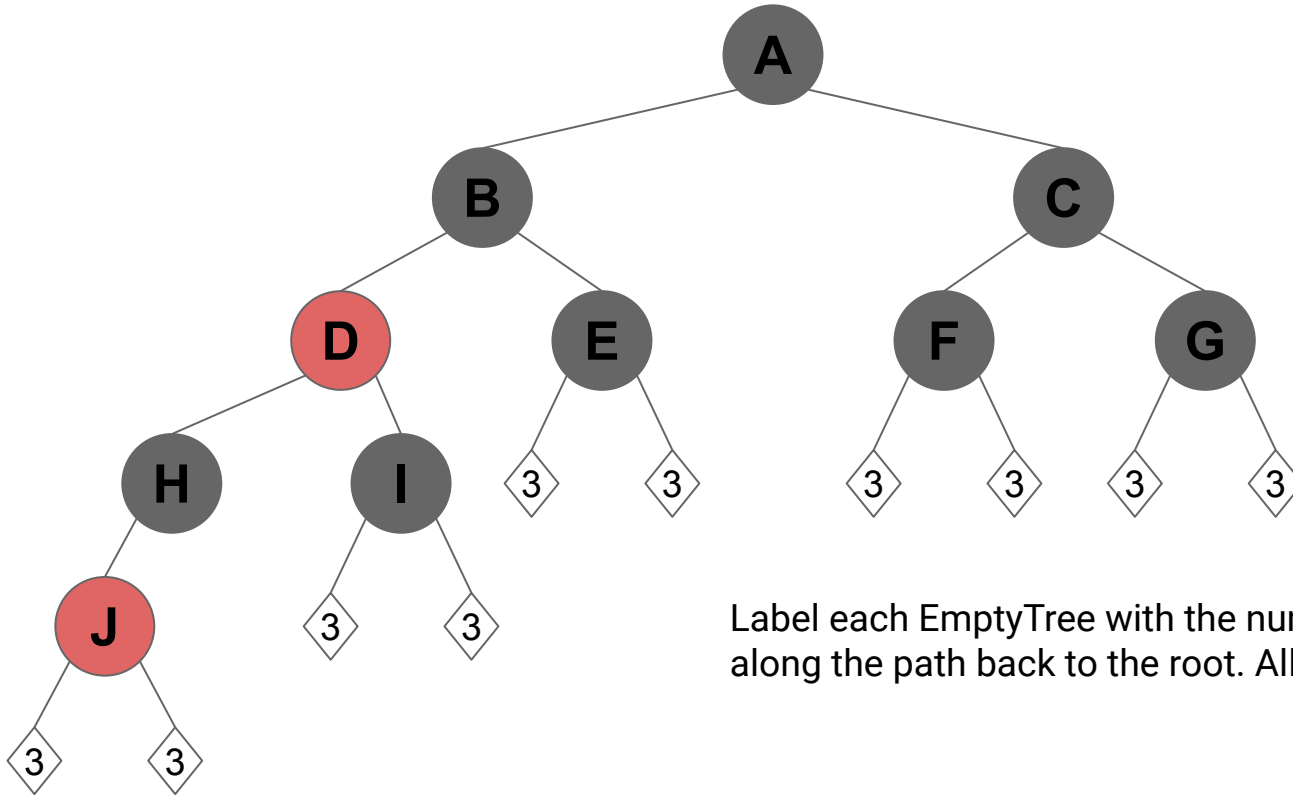
To Enforce the Depth Constraint on EmptyTree nodes:

1. Color each node red or black
 - a. The # of black nodes from each EmptyTree node to root must be same
 - b. The parent of a red node must always be black
2. On insertion (or deletion)
 - a. Inserted nodes are red (won't break 1a)
 - b. Repair violations of 1b by rotating and/or recoloring
 - i. Make sure repairs don't break 1a

Red-Black Trees

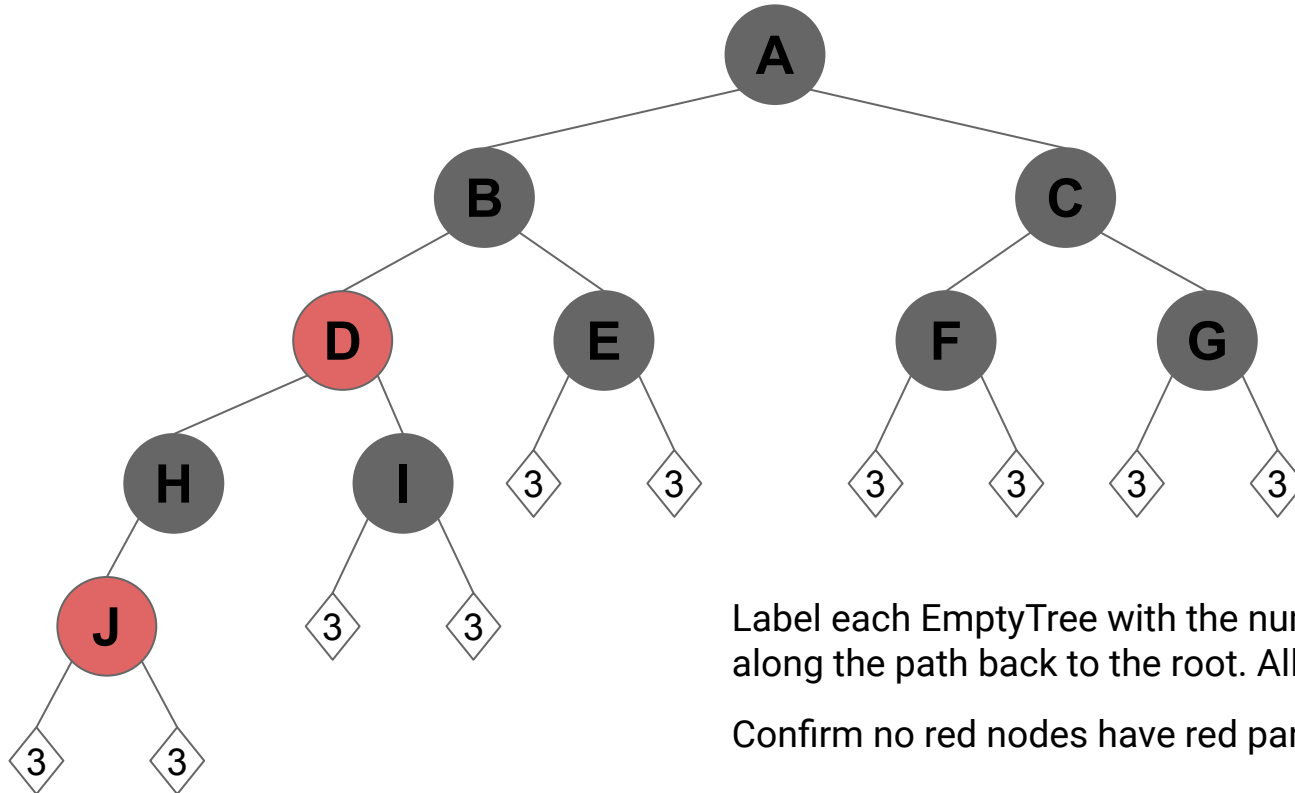


Red-Black Trees



Label each EmptyTree with the number of black nodes along the path back to the root. All 3 in this case ✓

Red-Black Trees



Label each EmptyTree with the number of black nodes along the path back to the root. All 3 in this case ✓

Confirm no red nodes have red parents ✓

Red-Black Trees

How does this coloring relate to our depth constraint?

Red-Black Trees

Assume we have a valid Red-Black tree with X black nodes from on each path from EmptyTree to root

What is the shallowest possible depth of an EmptyTree node?

Red-Black Trees

Assume we have a valid Red-Black tree with X black nodes from on each path from EmptyTree to root

What is the shallowest possible depth of an EmptyTree node?

X black nodes in a row = X

Red-Black Trees

Assume we have a valid Red-Black tree with X black nodes from on each path from EmptyTree to root

What is the shallowest possible depth of an EmptyTree node?

X black nodes in a row = X

What is the deepest possible depth of an EmptyTree node?

Red-Black Trees

Assume we have a valid Red-Black tree with X black nodes from on each path from EmptyTree to root

What is the shallowest possible depth of an EmptyTree node?

X black nodes in a row = X

What is the deepest possible depth of an EmptyTree node?

X black nodes with 1 red node between each one = $2X$

Red-Black Trees

Now we have:

1. If we color nodes red and black with the rules described, then the shallowest EmptyTree will be at least half the depth of the deepest
2. If the shallowest EmptyTree is at least half the depth of the deepest then the depth of our tree is $O(\log(n))$

Red-Black Trees

Now we have:

1. If we color nodes red and black with the rules described, then the shallowest EmptyTree will be at least half the depth of the deepest
2. If the shallowest EmptyTree is at least half the depth of the deepest then the depth of our tree is $O(\log(n))$

So how do we build/color our tree?


Red-Black Tree

After insertion or deletion, what situations can we encounter?

Red-Black Tree

After insertion or deletion, what situations can we encounter?

Case 1a: Our root is red, we're all good! ✓




Triangles represent **valid**
Red-Black tree fragments

Red-Black Tree

After insertion or deletion, what situations can we encounter?

Case 1b: Our root is black, we're all good! ✓




Triangles represent **valid**
Red-Black tree fragments

Red-Black Tree

After insertion or deletion, what situations can we encounter?

Case 2: The node we are checking is red...




Triangles represent **valid**
Red-Black tree fragments

Red-Black Tree

After insertion or deletion, what situations can we encounter?

Case 2: The node we are checking is red...
and its parent is black. We are all good! ✓




Triangles represent **valid**
Red-Black tree fragments

Red-Black Tree

After insertion or deletion, what situations can we encounter?


Case 3: The node we are checking is red... and it's parent is red. Now we have to fix the tree.



Red-Black Tree

After insertion or deletion, what situations can we encounter?

Case 3a: The node we are checking is red...
and its parent is red. That node's parent is
black and its sibling is red...




Red-Black Tree

After insertion or deletion, what situations can we encounter?

Case 3a: The node we are checking is red...
and it's parent is red. That node's parent is
black and it's sibling is red...

Recolor B,C,D. Are we all good?

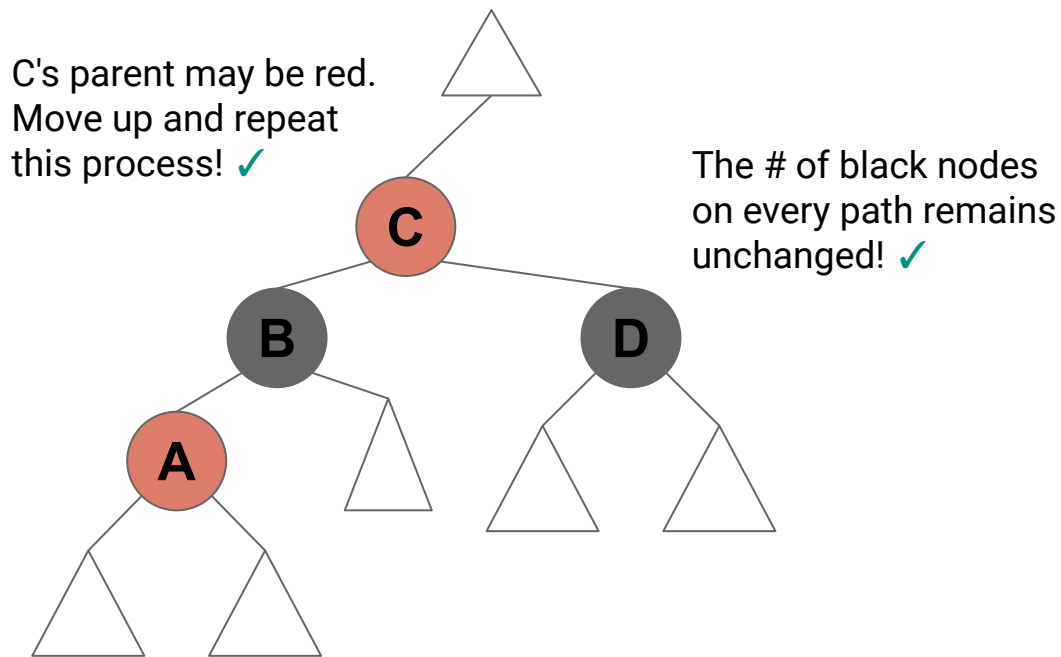


Red-Black Tree

After insertion or deletion, what situations can we encounter?

Case 3a: The node we are checking is red... and its parent is red. That node's parent is black and its sibling is red...

Recolor B,C,D. Are we all good?




Red-Black Tree

After insertion or deletion, what situations can we encounter?

Case 3a: The node we are checking is red... and its parent is red. That node's parent is black and its sibling is red...

Recolor B,C,D. Are we all good?


Note: This also works if A is right child of B and/or B is right child of C



Red-Black Tree

After insertion or deletion, what situations can we encounter?

Case 3b: The node we are checking is red...
and its parent is red. That node's parent is
black and its sibling is black...




Red-Black Tree

After insertion or deletion, what situations can we encounter?

Case 3b: The node we are checking is red...
and its parent is red. That node's parent is
black and its sibling is black...

Rotate(B,C)




Red-Black Tree

After insertion or deletion, what situations can we encounter?

Case 3b: The node we are checking is red... and its parent is red. That node's parent is black and its sibling is black...

Rotate(B,C)

1 less black node to root for this part of the tree...




Same # of black nodes to the root from this part of tree

Red-Black Tree

After insertion or deletion, what situations can we encounter?

Case 3b: The node we are checking is red... and its parent is red. That node's parent is black and its sibling is black...


Rotate(B,C)
Recolor(B,C)



Red-Black Tree

After insertion or deletion, what situations can we encounter?

Case 3c: The node we are checking is red... and its parent is red. That node's parent is black and its sibling is black...but A is the right child of B




Red-Black Tree

After insertion or deletion, what situations can we encounter?

Case 3c: The node we are checking is red... and it's parent is red. That node's parent is black and it's sibling is black...but A is the right child of B

Rotate(B,A) now we are back to **3b**



Red-Black Tree

Note: Each insertion creates at most one red-red parent-child conflict

- $O(1)$ time to recolor/rotate to repair the parent-child conflict
- May create a red-red conflict in grandparent
 - Up to $d/2 = O(\log(n))$ repairs required, but each repair is $O(1)$
- **Insertion therefore remains $O(\log(n))$**

Note: Each deletion removes at most one black node (red doesn't matter)

- $O(1)$ time to recolor/rotate to preserve black-depth
- May require recoloring (grand-)parent from black to red
 - Up to $d = O(\log(n))$ repairs required
- **Deletion therefore remains $O(\log(n))$**