


CSE 332

Data Structures and Parallelism

AVL Trees



Outline

- 1 Introducing AVL Trees
- 2 Tree Representation in Code
- 3 How Does an AVL Tree Work?
- 4 Why Does an AVL Tree Work?
- 5 AVL Tree Examples

Left and right subtrees **recursively** have heights differing by at most one.

Definition (balance)

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


Definition (AVL Balance Property)

An AVL tree is balanced when:

$$\text{For every node } n, \text{balance}(n) \leq 1$$

- This ensures a small depth
- It's relatively easy to maintain

AVL Tree



Structure Property:
0, 1, or 2 children

BST Property:
Keys in Left Subtree are smaller
Keys in Right Subtree are larger

AVL Balance Property:
Left and Right subtrees have heights
that differ by at most one.

That is, **all AVL Trees are BSTs**, but the reverse is not true.

AVL Trees rule out **unbalanced BSTs**.

Node Class?

```
class Node {  
    Data data;  
    Node left;  
    Node right;  
}
```

This Definition Leads to Redundant Code

```
1 boolean find(Node current, int data) {  
2     if (current == null) {  
3         return false;  
4     }  
5     else if (current.data == data) {  
6         return true;  
7     }  
8  
9     if (current.data < data) {  
10        return find(current.left, data);  
11    }  
12    else {  
13        return find(current.right, data);  
14    }  
15 }
```

But that's what we've been writing! Why is it ugly?

- It's redundant
- The left and right cases are **the same**, why write them twice?
- It's not idiomatic (e.g., the right abstraction would allow us to write the two cases found vs. not found)

Node Class?

```
class Node {  
    Data data;  
    Node left;  
    Node right;  
}
```

A Bad Fix

```
1 boolean find(Node current, int data) {  
2     if (current == null) {  
3         return false;  
4     }  
5  
6     if (current.data == data) {  
7         return true;  
8     }  
9     else {  
10        Node next = null;  
11        if (current.data < data) { next = current.left; }  
12        else { next = current.right; }  
13        return find(next, data);  
14    }  
15 }
```

How is This Code?

```
int a0 = 0;  
int a1 = 0;  
int a2 = 0;  
  
for (int i = 0; i < 3; i++) {  
    if (i == 0) { a0 = i; }  
    else if (i == 1) { a1 = i; }  
    else { a2 = i; }  
}
```

This course is about **making the right data abstractions**. This is a perfect example of where we could improve.

Keep an **array** of children!

Node Class?

```
class Node {
    Data data;
    Node[] children;
}
```

Is This Really Any Better?

```
1 boolean find(Node current, int data) {
2     if (current == null) {
3         return false;
4     }
5     else if (current.data == data) {
6         return true;
7     }
8
9     int next = current.data < data ? 0 : 1;
10    return find(current.children[next], data)
11        ;
12 }
```

Actually, yes! How do I get “the other child” in each of these versions?


```
1 Node getOtherChild(Node me, Node child1) {
2     if (me.left == child1) { return me.right; }
3     else { return me.left; }
4 }
```

vs.

```
1 Node getOtherChild(Node me, int child1) {
2     return me.children[1 - child1];
3 }
```

Since operations on binary trees are **almost always symmetric**, this is a big deal for complicated operations. Keep this in mind.


Worst Case



When we insert 3, we violate the AVL Balance condition. What to do?


There's only one tree with the BST Property and the Balance Property:

FIXING The Worst Case



This “fix” is called a rotation. We’re “rotating” the child node “up”:


Rotation



This is the only fundamental of AVL Trees!

You can either look at this as “the only way to correctly rearrange the subtrees” or it’s helpful to think of it as gravity.

Rotation




The Code

```
1 void rotate(Node current) {  
2     Node child = current.right;  
3     current.right = child.left;  
4     child.left = current;  
5  
6     child.height = child.updateHeight();  
7     current.height = current.updateHeight();  
8  
9     current = child;  
10 }
```


Inserting 16

Is the result an AVL tree? If not, how do we fix it?



This is just the same rotation in the other direction!


Rotation



The Code

```
1 void rotate(Node current) {  
2     Node child = current.left;  
3     current.left = child.right;  
4     child.right = current;  
5  
6     child.height = child.updateHeight();  
7     current.height = current.updateHeight();  
8  
9     current = child;  
10 }
```

We Want...




Cases We've Handled



Cases To Handle




Second Case




When we insert 2, we violate the AVL Balance condition. What to do?

There's only one tree with the BST Property and the Balance Property:


FIXING The Second Case




Double Rotation




$\text{fix}(a)$



First, we rotate b.




$\text{rotate}(a)$



And now it's balanced!

Now, we're back to the **line** case.

Double Rotation



Left Double Rotation Code

```
1 void leftDoubleRotation(Node current) {  
2     rotation(current.right, RIGHT);  
3     rotation(current, LEFT);  
4 }
```

Expectations and Tips

- For any one bug, debug for **30 minutes**, then stop.
- For any one exercise, you should be spending drastically different times.
- Most exercises are **direct applications** of lecture. They are not 311 problems.
- Partners: There is no “formal” way of saying “my partner isn’t doing enough work”, but we DO factor that information in if you let us know

AVL Operations

- `find(x)` is identical to BST `find`
- `insert(x)` by (1) doing a BST `insert`, and (2) fixing the tree with either a rotation or a double rotation
- `delete(x)` by either a similar method to `insert`—or doing lazy `delete`




AVL Fields

- We've seen that the code is very redundant if we use `left` and `right` fields; so, we should use a `children` array
- We've seen quick access to `height` is very important; so, it should be a field


Okay, so does it work?

We must **guarantee** that the AVL property gives us a small enough tree. Our approach: Find a big **lower bound** on the number of nodes necessary to make a tree with height h .

What is the **smallest** number of nodes to get a height h AVL Tree?



What is the **smallest** number of nodes to get a height h AVL Tree?



The general number of nodes to get a height of h is:

$$f(h) = f(h-2) + f(h-1) + 1$$

We break down where each term comes from. We want a tree that has the **smallest** number of nodes where each branch has the AVL Balance condition.

- $f(h-1)$: To force the height to be h , we take the smallest tree of height $h-1$ as one of the children
- $f(h-2)$: We are allowed to have the branches differ by one; so, we can get a smaller number of nodes by using $f(h-2)$
- $+1$ comes from the root node to join together the two branches

So, now we solve our recurrence. How?

Ratio Between Terms

A good way of solving a recurrence that we expect to be of the form X^n is to look at the ratio between terms. If $\frac{f(h+1)}{f(h)} > X$, then

$$f(h+1) > Xf(h) > X(X(f(h-1))) > \dots > X^h$$

So, we evaluate these ratios and see the following:

OUTPUT

```
>> 2.0
>> 2.0
>> 1.75
>> 1.7142857142857142
>> 1.6666666666666667
>> 1.65
>> 1.6363636363636365
>> 1.6296296296296295
>> 1.625
>> 1.6223776223776223
>> 1.6206896551724137
>> 1.6196808510638299
>> 1.619047619047619
>> 1.618661257606491
>> 1.618421052631579
>> ...
```

In this case, we see that $f(h)$ pretty quickly converges to $\phi(1.618\dots)$. Before trying to prove this closed form, we should look at a few examples:

- $f(0) = 1$ vs. $(\phi)^0 = 1$
- $f(1) = 2$ vs. $(\phi)^1 = \phi$

We want to show that $f(h) >$ some closed form, but looking at the first base case, $1 \not> 1$. So, we'll prove $f(h) > \phi^h - 1$ instead.

Induction Proof

- Base Cases: Note that $f(0) = 1 > 1 - 1 = 0$ and $f(1) = 2 > \phi - 1 \approx 0.618$
- Induction Hypothesis: Suppose that $f(h) > \phi^h - 1$ for all $0 \leq h \leq k$ for some $k \geq 1$.
- Induction Step:

$$\begin{aligned}
 f(k+1) &\geq f(k) + f(k-1) + 1 \\
 &> (\phi^k - 1) + (\phi^{k-1} - 1) + 1 \quad [\text{By IH}] \\
 &= \phi^{k-1}(\phi + 1) + 1 - 2 \\
 &= \phi^{k+1} - 1 \quad [\text{By } \phi]
 \end{aligned}$$

In the step labeled “by ϕ ”, we use the property $\phi^2 = \phi + 1$.

So, since $n \geq f(h) > \phi^h - 1$, taking \lg of both sides gives us:

$$\lg(n) > \lg(\phi^h - 1) \approx \lg(\phi^h) = h \lg(\phi)$$

So, $h \in \mathcal{O}(\lg n)$.

- Worst-case complexity of `find`:
- Worst-case complexity of `insert`:
 - Tree starts balanced
 - A rotation is $\mathcal{O}(1)$ and there's an $\mathcal{O}(\lg n)$ path to root
 - (Same complexity even without one-rotation-is-enough fact)
 - Tree ends balanced
- Worst-case complexity of `buildTree`:
- Worst-case complexity of `delete`: (requires more rotations)
- Worst-case complexity of `lazyDelete`:

So, since $n \geq f(h) > \phi^h - 1$, taking \lg of both sides gives us:

$$\lg(n) > \lg(\phi^h - 1) \approx \lg(\phi^h) = h \lg(\phi)$$

So, $h \in \mathcal{O}(\lg n)$.


- Worst-case complexity of `find`: $\mathcal{O}(\lg n)$
- Worst-case complexity of `insert`: $\mathcal{O}(\lg n)$
 - Tree starts balanced
 - A rotation is $\mathcal{O}(1)$ and there's an $\mathcal{O}(\lg n)$ path to root
 - (Same complexity even without one-rotation-is-enough fact)
 - Tree ends balanced
- Worst-case complexity of `buildTree`: $\mathcal{O}(n \lg n)$
- Worst-case complexity of `delete`: (requires more rotations) $\mathcal{O}(\lg n)$
- Worst-case complexity of `lazyDelete`: $\mathcal{O}(1)$

Pros of 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

Cons of AVL trees




- Difficult to program & debug
- More space for height field
- Asymptotically faster but rebalancing takes a little time
- Most large searches are done in database-like systems on disk and use other structures (e.g., B-trees, our next data structure)

Example (Insert a, b, e, c, d into an AVL Tree)


Example (Which Rotation?)




- Which insertions would cause a **single rotation**?




Example (Which Rotation?)






- Which insertions would cause a **double rotation**?




Example (Which Rotation?)




- Which insertions would cause no rotation?




Example (Insert 3, 33, 18, 32)



Example (Insert 3, 33, 18, 32)



Example (Insert 3, 33, 18, 32)



Example (Insert 3, 33, 18, 32)

