

CSE 332: Data Structures and Parallelism

Fall 2022
Richard Anderson
Lecture 7: Адельсон-Вельский
Ландис деревья


Announcements*

- 10/12: AVL Trees
- 10/14: AVL Trees
- 10/17: B-Trees
- 10/19: B-Trees
- 10/21: Hashing I
- 10/24: Hashing II
- 10/26: Sorting I
- 10/28: Sorting II
- 10/31: Sorting III
- 11/02: Splay Trees
- 11/04: Midterm

*Subject to change

Binary Search Tree Data Structure

- Structural property
 - each node has ≤ 2 children
- Order property
 - all keys in left subtree smaller than root's key
 - all keys in right subtree larger than root's key
- Find / Insert
 - Compare with node value to go left or right
 - Runtime $O(\text{height})$
- Works great, unless tree is unbalanced



Balanced binary trees

- Binary tree with guarantee on depths of leaves
- $O(\log n)$ insert and delete
- Many flavors
 - Red-black trees
 - Self-adjusting binary trees
 - 2-3 trees
 - AVL Trees
- Issues
 - Ensure height is bounded by $c \log n$
 - Maintain this bound on insert and delete

Imbalance is a real issue

$$\log_2 1,000 \approx 10$$

$$\log_2 1,000,000 \approx 20$$

$$\log_2 1,000,000,000 \approx 30$$

Almost sorted input is a realistic case


More on binary search trees

What does the following function do:

```
LV(Node root, int low, int high) {
    if (root == null)
        return;
    if (low <= root.key)
        LV(root.left, low, high);
    if (low <= root.key && high >= root.key)
        Print(root.key);
    if (high >= root.key)
        LV(root.right, low, high);
}
```

What is the run time?

- Tree of height h
- A total of k nodes are output



10/12/2022

CSE 332

7

AVL Trees

- Developed in 1962 by Soviet mathematicians Georgy Adelson-Velsky and Evgeny Landis (Адельсон-Вельский, Ландис)
- Structural property on tree guarantees depth $O(\log n)$
- Rebalance operation to ensure property
- Practical
- First published balanced search tree

AN ALGORITHM FOR THE ORGANIZATION OF INFORMATION

G. M. ADEL'SON-VEL'SKIĬ AND E. M. LANDIS

An algorithm is proposed in which both the search and the recording are carried out in $C \lg N$ operations, where N is the number of information elements which have entered at a given moment.


10/12/2022

CSE 332

8

AVL Tree overview

- Balance condition
- Depth bound
- Rotations to rebalance the tree



10/12/2022

CSE 332

9

The AVL Tree Data Structure

Structural properties


1. Binary tree property
2. Balance: $\text{left.height} - \text{right.height}$
3. Balance property: balance of every node is between -1 and 1

Result:

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

Ordering property

- Same as for BST




10/12/2022

CSE 332

10

An AVL tree?




10/12/2022

CSE 332

11

An AVL tree?



10/12/2022

CSE 332

12

Height of an AVL Tree?

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


Let $S(h)$ be the minimum # of nodes in an AVL tree of height h , then:

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

where $S(0) = 1$ and $S(1) = 2$

Solution of Recurrence: $S(h) \approx 1.62^h$

$$n \geq \phi^h \implies \log_{\phi} n \geq h$$



10/12/2022

CSE 332

13

Let $S(h)$ be the minimum # of nodes in an AVL tree of height h , then:

$$S(h) = S(h-1) + S(h-2) + 1 \quad \text{where } S(0) = 1 \text{ and } S(1) = 2$$

h

Minimal AVL Tree

$S(h)$

10/12/2022

CSE 332

14

A simpler bound: $S_k \geq 2^{k/2}$

$$S_0 = 1, S_1 = 2, S_k = 1 + S_{k-1} + S_{k-2}$$

1, 2, 4, 7, 12, 20, 33, 54, ...

Growing faster than the Fibonacci sequence

Observation: $S_k \geq 2 S_{k-2}$

10/12/2022

CSE 332

15

The Golden Ratio


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

This is a special number

- Golden ratio: If $(a+b)/a = a/b$, then $a = \phi b$

- We will need one special arithmetic fact about ϕ :

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



10/12/2022

CSE 332

16

$$S(0)=1, S(1)=2, S(2)=4$$

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

The Proof

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 $S(k-1) > \phi^{k-1} - 1$

$$S(k+1) = 1 + S(k) + S(k-1) \quad \text{by definition of } S$$

$$> 1 + \phi^k - 1 + \phi^{k-1} - 1 \quad \text{by induction}$$

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

$$= \phi^{k-1} (\phi + 1) - 1 \quad \text{by arithmetic (factor } \phi^{k-1} \text{)}$$

$$= \phi^{k-1} \phi^2 - 1 \quad \text{by special property of } \phi$$

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

10/12/2022

CSE 332

17


Good news

Proof means that if we have an AVL tree, then **find** is $O(\log n)$

– Recall logarithms of different bases > 1 differ by only a constant factor

But as we insert and delete elements, we need to:

1. Track balance
2. Detect imbalance
3. Restore balance



Is this AVL tree balanced?


How about after insert(30)?

10/12/2022

CSE 332

18

An AVL Tree



10/12/2022

CSE 332

19

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 have several imbalance cases (next lecture, maybe)

10/12/2022

CSE 332

20

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 recursive 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

Facts 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


10/12/2022

CSE 332

21

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?

10/12/2022


CSE 332

22

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 here




10/12/2022

CSE 332

23

The example generalized

- Node imbalanced due to insertion *somewhere* in **left-left grandchild** increasing height
 - 1 of 4 possible imbalance causes (other three coming)
- First we did the insertion, which would make **a** imbalanced




10/12/2022

CSE 332

24

The general left-left case

- Node imbalanced due to insertion *somewhere* in **left-left grandchild**
 - 1 of 4 possible imbalance causes (other three coming)
- So we rotate at **a**, 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

10/12/2022

CSE 332

25

Another example: insert (16)




10/12/2022

CSE 332

26

Another example: insert (16)




10/12/2022

CSE 332

27

The general right-right case

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



10/12/2022

CSE 332


28

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



10/12/2022

CSE 332


29

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




10/12/2022

CSE 332

30

Sometimes two wrongs make a right ☺

- First idea violated the BST 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:
 1. Rotate problematic child and grandchild
 2. Then rotate between self and new child




10/12/2022

CSE 332

31

The general right-left case




10/12/2022

CSE 332

32

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:



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


10/12/2022

CSE 332

33

The last case: left-right

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



10/12/2022

CSE 332

34

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

10/12/2022

CSE 332

35