



CS 225

Data Structures

March 19 – AVL Applications

Brad Solomon

Informal Early Feedback Reminder

CS 225 All SP21: Data Structures (Evans, C)

[Dashboard](#) / [My courses](#) / [CS 225 All SP21](#)



Announcements



Lab Attendance



Informal Early Feedback



Lab Informal Early Feedback



Learning Objectives

- Review Big O in the contexts of an AVL tree
- Formalize relationship between n and h in an AVL tree
- Prove h has an upper bound of $O(\log n)$
- Wrap up balanced binary trees

AVL Tree Analysis

For AVL tree of height h , we know:


find runs in: _____.

insert runs in: _____.

remove runs in: _____.


We will argue that: h is _____.

AVL Tree Analysis




The height of the tree, $f(n)$, will always be less than $c \times g(n)$ for all values where $n > k$.

AVL Tree Analysis



$f(n)$ = "Tree height given nodes"



$f^{-1}(h)$ = "Nodes in tree given height"

The number of nodes in the tree, $f^{-1}(h)$, will always be greater than $c \times g^{-1}(h)$ for all values where $n > k$.

Plan of Action

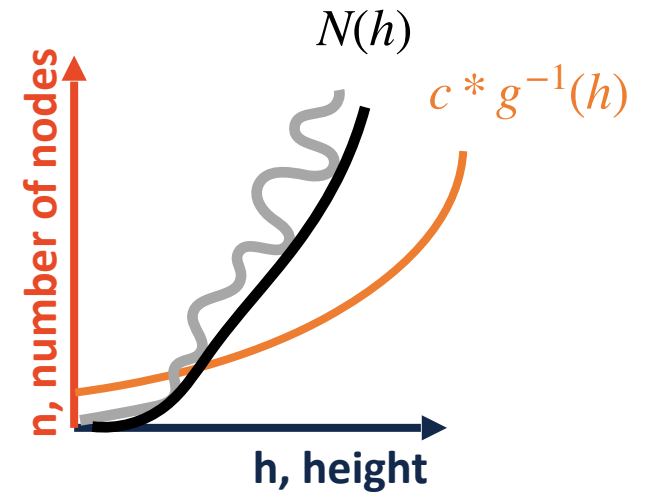
Since our goal is to find the lower bound on n given h , we can begin by defining a function given h which describes the smallest number of nodes in an AVL tree of height h :

$N(h)$ = minimum number of nodes in an AVL tree of height h

Simplify the Recurrence

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

$$N(h) \geq N(h) - 1$$



State a Theorem

Theorem: An AVL tree of height h has at least _____.

Proof by Induction:

I. Consider an AVL tree and let h denote its height.

II. Base Case: _____

An AVL tree of height _____ has at least _____ nodes.

Prove a Theorem

III. Base Case: _____

An AVL tree of height _____ has at least _____ nodes.

Prove a Theorem

IV. Induction Case: _____

If for all heights $i < h$, $N(i) \geq 2^{i/2}$

then we must show for height h that $N(h) \geq 2^{h/2}$

Prove a Theorem

V. Using a proof by induction, we have shown that:

...and inverting:

AVL Runtime Proof

An upper-bound on the height of an AVL tree is **$O(\lg(n))$** :

$N(h)$:= Minimum # of nodes in an AVL tree of height h

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

$$> 1 + 2^{h-1/2} + 2^{h-2/2}$$

$$> 2 \times 2^{h-2/2} = 2^{h-2/2+1} = 2^{h/2}$$

Theorem #1:

Every AVL tree of height h has at least $2^{h/2}$ nodes.



AVL Runtime Proof

An upper-bound on the height of an AVL tree is **$O(\lg(n))$** :

$$\# \text{ of nodes } (n) \geq N(h) > 2^{h/2}$$

$$n > 2^{h/2}$$

$$\lg(n) > h/2$$

$$2 \times \lg(n) > h$$

$$h < 2 \times \lg(n) \quad , \text{ for } h \geq 1$$

Proved: The maximum number of nodes in an AVL tree of height h is less than $2 \times \lg(n)$.

Summary of Balanced BST

AVL Trees

- Max height: $1.44 * \lg(n)$
- Rotations:

Summary of Balanced BST

AVL Trees

- Max height: $1.44 * \lg(n)$
- Rotations:
 - Zero rotations on find
 - One rotation on insert
 - $O(h) == O(\lg(n))$ rotations on remove

Red-Black Trees

- Max height: $2 * \lg(n)$
- Constant number of rotations on insert (max 2), remove (max 3).

Red-Black Trees in C++

C++ provides us a balanced BST as part of the standard library:

```
std::map<K, V> map;
```

Modifiers		Lookup	
<code>clear</code>	clears the contents (public member function)	<code>count</code>	returns the number of elements matching specific key (public member function)
<code>insert</code>	inserts elements or node (public member function)	<code>find</code>	finds element with specific key (public member function)
<code>insert_or_assign</code> (C++17)	inserts an element or assigns (public member function)	<code>contains</code> (C++20)	checks if the container contains element with specific key (public member function)
<code>emplace</code> (C++11)	constructs element in-place (public member function)	<code>equal_range</code>	returns range of elements matching a specific key (public member function)
<code>emplace_hint</code> (C++11)	constructs elements in-place (public member function)	<code>lower_bound</code>	returns an iterator to the first element <i>not less</i> than the given key (public member function)
<code>try_emplace</code> (C++17)	inserts in-place if the key does not exist, does nothing if the key exists (public member function)	<code>upper_bound</code>	returns an iterator to the first element <i>greater</i> than the given key (public member function)
<code>erase</code>	erases elements (public member function)		
<code>swap</code>	swaps the contents (public member function)		
<code>extract</code> (C++17)	extracts nodes from the container (public member function)		
<code>merge</code> (C++17)	splices nodes from another container (public member function)		



Red-Black Trees in C++

```
V & std::map<K, V>::operator[] ( const K & )
```

```
void std::map<K, V>::erase ( const K & )
```



Why Balanced BST?

Summary of Balanced BST

Pros:

- Running Time:
 - Improvement Over:
- Great for specific applications:

Trees in the Real World

Q: Can we always fit our data in main memory?

Q: Where else can we keep our data?


We assume constant time memory access, but the constant factor can be limiting in real world settings!

Memory Hierarchy (Speed of access)

(Measured in 2011 at <https://gist.github.com/hellerbarde/2843375>)

	Time x1 billion
L1 cache reference	0.5 seconds
Branch mispredict	5 seconds
L2 cache reference	7 seconds
Mutex lock/unlock	25 seconds
Main memory reference	100 seconds
Compress 1K bytes	50 minutes
Send 2K bytes over 1 Gbps network	5.5 hours
SSD random read	1.7 days
Read 1 MB sequentially from memory	2.9 days
Read 1 MB sequentially from SSD	11.6 days
Disk seek	16.5 weeks
Read 1 MB sequentially from disk	7.8 months
Above two together	1 year
Send packet CA->Netherlands->CA	4.8 years

AVLs in the Cloud



BTree Motivations

Knowing that we have large seek times for data, we want to:

BTree (of order m)

A **BTree** of order m is an m -way tree:

- All keys within a node are ordered
- All nodes contain no more than **$m-1$** keys.



BTree in the Real World

-3	8	23	25	31	42	43	55
----	---	----	----	----	----	----	----

$m=9$

Goal: Minimize the number of reads!

Build a tree that uses _____ / node
[1 network packet]
[1 disk block]