

Programming Abstractions

CS106B

Cynthia Lee

Topics:

- Monday:
 - › Priority Queue implemented as a binary heap
 - Enqueue (“bubble up”)
 - Dequeue (“trickle down”)
- Today: **Map** implemented as a Binary Search Tree (BST)
 - › Starting with a dream: binary search in a linked list?
 - › How our dream provided the inspiration for the BST
 - › BST insert
 - › Big-O analysis of BST
- Friday: BST balance issues



Binary Search Trees

Implementing the **Map interface** with Binary Search Trees

Implementing Map interface with a Binary Search Tree (BST)

- Binary Search Trees is one option for implementing Map
 - › C++'s **Standard Template Library (STL)** uses a Red-Black tree (a type of BST) for their map
 - › Stanford library also uses a BST
- Another Map implementation is a hash table
 - › We will talk about this later!
 - › This is what Stanford's HashMap uses

Binary Search in a Linked List?

Exploring a good idea, finding way to make it work

Imagine storing sorted data in an array

0	1	2	3	4	5	6	7	8	9	10
2	7	8	13	25	29	33	51	89	90	95

- How long does it take us to find?
 - › **Binary search!**
 - › **$O(\log n)$** : awesome!
- But slow to insert
 - › **Move everyone over to make space**
 - › **$O(n)$** : not terrible, but pretty bad compared to $\log(n)$ or $O(1)$
 - › In contrast, linked list stores its nodes scattered all over memory, so it doesn't have to move things over


Q. Can we do binary search on a linked list to implement a quick insert?

A. No.

- The nodes are spread all over memory, and we must follow “next” pointers one at a time to navigate (the treasure hunt).
- **Therefore cannot jump right to the middle.**
- **Therefore cannot do binary search.**
- **Find is $O(N)$:** not terrible, but pretty bad compared to $O(\log n)$ or $O(1)$

Binary Search Tree can be thought of roughly like a linked list that has pointers to the middle, again and again (recursively), to form a tree structure

An Idealized Binary Search Tree



TreeMap


This is basically the same as Stanford Map

tree-map.h

```
template <typename Key, typename Value>
class TreeMap {
public:
    TreeMap();
    ~TreeMap();

    bool isEmpty() const;
    int size() const { return count; }
    bool containsKey(const Key& key) const;
    void put(const Key& key, const Value& value);
    Value get(const Key& key) const;
    Value& operator[](const Key& key);
    ...
}
```

tree-map.h






// class TreeMap continued...

private:

```
struct node {  
    Key key;  
    Value value;  
    node *left, *right;  
};
```

```
int size;  
node *root;
```

```
};
```





Example: insert 2


BST put()

Pretty simple!

- If $\text{key} > \text{node's key}$
 - › Go right!
- If key < node's key
 - › Go left!
- If there is nothing currently in the direction you are going, that's where you end up
- Example: `put(23)`



BST put()



Insert: 22, 9, 34, 18, 3

How many of these result in the same tree structure as above?

22, 34, 9, 18, 3

22, 18, 9, 3, 34

22, 9, 3, 18, 34

- A. None of these
- B. 1 of these
- C. 2 of these
- D. All of these

Question about our put() algorithm:

Pretty simple!

- If key > node's key
 - › Go right!
- If key < node's key
 - › Go left!

Q. What do we do if the key is equal to the node's key?

Stanford Map example:

```
Map<int, string> mymap;  
mymap.put(5, "five");  
mymap.put(5, "cinco");           // what should happen?  
cout << mymap.get(5) << endl;   // what should print?
```

Performance analysis:

What is the WORST CASE cost for doing containsKey() in BST?

- A. $O(1)$
- B. $O(\log n)$
- C. $O(n)$
- D. $O(n \log n)$
- E. $O(n^2)$

What is the worst case cost for doing
containsKey() in BST *if the BST is balanced?*

$O(\log N)$ —awesome!

BSTs are great when balanced

BSTs are bad when unbalanced

...and Balance depends on order of insert of elements...

...but user controls this, not “us” (author of the Map class)

...no way for “us” (author of Map class) to ensure our Map
doesn't perform terribly! ☹ ☹

Ok, so, long-chain BSTs are bad, should we worry about it? [math puzzle time]

One way to create a bad BST is to insert the elements in *decreasing* order: 34, 22, 9, 3

That's not the only way...

How many **distinctly structured** BSTs are there that exhibit the worst case height (height equals number of nodes) for a tree with the 4 nodes listed above?

A. 2-3

B. 4-5

C. 6-7

D. 8-9

E. More than 9

$$2^{n-1}$$

Bonus question: general formula for any BST of size n ?

Extra bonus question (CS109): what is this as a fraction of all trees (i.e., probability of worst-case tree).

Midterm Questions?

Are you nervous? What can I do to help?