

CSE 143

Binary Search Trees

5/26/2004

(c) 2001-4, University of Washington

21-1

Costliness of *contains*

- Review: in a binary tree, *contains* is $O(N)$
- *contains* may be a frequent operation in an application
- Can we do better than $O(N)$?
- Turn to list searching for inspiration...
 - Why was binary search so much better than linear search?
 - Can we apply the same idea to trees?

5/26/2004

(c) 2001-4, University of Washington

21-2

Binary Search Trees

- Idea: order the nodes in the tree so that, given that a node contains a value v ,
 - All nodes in its left subtree contain values $< v$
 - All nodes in its right subtree contain values $> v$
- A binary tree with these properties is called a *binary search tree* (BST)
- Notes:
 - Can also define a BST using \geq and \leq instead of $>$, $<$
This implies there could be duplicate values in the tree
 - In Java, if the values are not primitive types, they must implement interface comparable (i.e., provide compareTo)


5/26/2004

(c) 2001-4, University of Washington

21-3

Examples(?)

- Are these are binary search trees? Why or why not?



5/26/2004

(c) 2001-4, University of Washington

21-4

Implementing a Set with a BST

- Can exploit properties of BSTs to have fast, divide-and-conquer implementations of Set's add and contains operations
- TreeSet!
- A TreeSet can be represented by a pointer to the root node of a binary search tree, or null of no elements yet

```
public class SimpleTreeSet implements Set {
    private BTNode root; // root node, or null if none
    public SimpleTreeSet() { root = null; }
    // size as for BinTree
    ...
}
```

5/26/2004

(c) 2001-4, University of Washington

21-5

contains for a BST

- For a general binary tree, contains had to search both subtrees
 - Like linear search
- With BSTs, need to only search one subtree
 - All small elements to the left, all large elements to the right
 - Search either left or right subtree, based on comparison between elem and value at root of tree
- Like binary search

5/26/2004

(c) 2001-4, University of Washington

21-6

Code for *contains* (in TreeSet)

```
/** Return whether elem is in set */
public boolean contains(Object elem) {
    return subtreeContains(root, (Comparable)elem);
}

// Return whether elem is in (sub-)tree with root r
private boolean subtreeContains(BTNode r, Comparable elem) {
    if (r == null) {
        return _____;
    } else {
        int comp = elem.compareTo(r.item);
        if (comp == 0) { return _____; } // found it!
        else if (comp < 0) { return _____; } // search left
        else { return _____; } // search right
    }
}
```


5/26/2004

(c) 2001-4, University of Washington


21-7

Examples

contains(6)



contains(10)



5/26/2004

(c) 2001-4, University of Washington

21-8

Cost of BST *contains*

- Work done at each node:
- Number of nodes visited (depth of recursion):
- Total cost:

5/26/2004

(c) 2001-4, University of Washington

21-9

add

- Must preserve BST invariant: insert new element in correct place in BST
- Two base cases
 - Tree is empty: create new node which becomes the root of the tree
 - If node contains the value, found it; suppress duplicate add
- Recursive case
 - Compare value to current node's value
 - If value < current node's value, add to left subtree recursively
 - Otherwise, add to right subtree recursively

5/26/2004

(c) 2001-4, University of Washington

21-10

Example

- Add 8, 10, 5, 1, 7, 11 to an initially empty BST, in that order:

5/26/2004

(c) 2001-4, University of Washington

21-11

Example (2)

- What if we change the order in which the numbers are added?
- Add 1, 5, 7, 8, 10, 11 to a BST, in that order (following the algorithm):

5/26/2004

(c) 2001-4, University of Washington

21-12

Code for *add* (in TreeSet)

```
// instance variable
private boolean treeChanged; // true if addToSubtree changes the tree, false if not
// (hack since addToSubtree can only return one value)

/** Ensure that elem is in the set. Return true if elem was added, false otherwise. */
public boolean add(Object elem) {
    treeChanged = false;
    root = addToSubtree(root, (Comparable)elem); // add elem to tree
    return treeChanged;
}

/** Add elem to tree rooted at r. Return (possibly new) tree containing elem, and set
 * treeChanged = true if the node was actually added */
private BTreeNode addToSubtree(BTreeNode r, Comparable elem) {
    ...
}
```

5/26/2004

(c) 2001-4, University of Washington

21-13

Code for *addToSubtree*

```
/** Add elem to tree rooted at r. Return (possibly new) tree containing elem, or throw
DuplicateAdded if elem already was in tree */
private BTreeNode addToSubtree(BTreeNode r, Comparable elem) throws DuplicateAdded {
    if (r == null) { // adding to empty tree
        treeChanged = true;
        return new BTreeNode(elem, null, null);
    }
    int comp = elem.compareTo(r.item);
    if (comp == 0) { return; } // elem already in tree
    if (comp < 0) { // add to left subtree
        r.left = addToSubtree(r.left, elem);
    } else { // comp > 0 // add to right subtree
        r.right = addToSubtree(r.right, elem);
    }
    return r; // this tree has been modified to contain elem
}
```

5/26/2004

(c) 2001-4, University of Washington

21-14

Cost of *add*

- Cost at each node:
- How many recursive calls?
 - Proportional to height of tree
- Best case?
- Worst case?

5/26/2004

(c) 2001-4, University of Washington

21-15

A Challenge: iterator

- How to return an iterator that traverses the sorted set in order?
 - Need to iterate through the items in the BST, from smallest to largest
- Problem: how to keep track of position in tree where iteration is currently suspended
 - Need to be able to implement next(), which advances to the correct next node in the tree
- Solution: keep track of a path from the root to the current node
 - Still some tricky code to find the correct next node in the tree

5/26/2004

(c) 2001-4, University of Washington

21-16

Another Challenge: *remove*

- Algorithm: find the node containing the element value being removed, and remove that node from the tree
- Removing a leaf node is easy: replace with an empty tree
- Removing a node with only one non-empty subtree is easy: replace with that subtree
- How to remove a node that has two non-empty subtrees?
 - Need to pick a new element to be the new root node, and adjust at least one of the subtrees
 - E.g., remove the largest element of the left subtree (will be one of the easy cases described above), make that the new root

5/26/2004

(c) 2001-4, University of Washington

21-17

Analysis of Binary Search Tree Operations

- Cost of operations is proportional to height of tree
- Best case: tree is *balanced*
 - Depth of all leaf nodes is roughly the same
 - Height of a balanced tree with n nodes is $\sim \log_2 n$
- If tree is unbalanced, height can be as bad as the number of nodes in the tree
 - Tree becomes just a linear list

5/26/2004

(c) 2001-4, University of Washington

21-18

Summary

- A binary search tree is a good general implementation of a set, if the elements can be ordered
 - Both contains and add benefit from divide-and-conquer strategy
 - No sliding needed for add
 - Good properties depend on the tree being roughly balanced
- Not covered (or, why take a data structures course?)
 - How are other operations implemented (e.g. iterator, remove)?
 - Can you keep the tree balanced as items are added and removed?

5/26/2004

(c) 2001-4, University of Washington

21-19