

# CSCI-1200 Data Structures — Fall 2014

## Lecture 17 – Trees, Part II

### Review from Lecture 16

- Binary Trees, Binary Search Trees, & Balanced Trees
- STL `set` container class (like STL `map`, but without the pairs!)
- Finding the smallest element in a BST.

### Today's Lecture

- Overview of the `ds_set` implementation: `begin`, `find`, and `destroy_tree`
- In-order, pre-order, and post-order traversal; Breadth-first and depth-first tree search
- Other `ds_set` operations: `insert`, and `erase`

### 17.1 `ds_set` and Binary Search Tree Implementation

- A partial implementation of a set using a binary search tree is in the code attached. We will continue to study this implementation in tomorrow's lab & the next lecture.
- The increment and decrement operations for iterators have been omitted from this implementation. Next lecture we will discuss a couple strategies for adding these operations.
- We will use this as the basis both for understanding an initial selection of tree algorithms and for thinking about how standard library sets really work.

### 17.2 `ds_set`: Class Overview

- There is two auxiliary classes, `TreeNode` and `tree_iterator`. All three classes are templated.
- The only member variables of the `ds_set` class are the root and the size (number of tree nodes).
- The iterator class is declared internally, and is effectively a wrapper on the `TreeNode` pointers.
  - Note that `operator*` returns a `const` reference because the keys can't change.
  - The increment and decrement operators are missing (we'll fill this in next lecture!).
- The main public member functions just call a private (and often recursive) member function (passing the root node) that does all of the work.
- Because the class stores and manages dynamically allocated memory, a copy constructor, `operator=`, and destructor must be provided.

### 17.3 Exercises

1. Provide the implementation of the member function `ds_set<T>::begin`. This is essentially the problem of finding the node in the tree that stores the smallest value.

2. Write a recursive version of the member function `ds_set::find`.

3. Write the `destroy_tree` private helper function.

## 17.4 In-order, Pre-Order, Post-Order Traversal

- One of the fundamental tree operations is “traversing” the nodes in the tree and doing something at each node. The “doing something”, which is often just printing, is referred to generically as “visiting” the node.
- There are three general orders in which binary trees are traversed: pre-order, in-order and post-order.
- In order to explain these, let’s first draw an “exactly balanced” binary search tree with the elements 1-7:
  - What is the *in-order traversal* of this tree? Hint: it is monotonically increasing, which is always true for an in-order traversal of a binary search tree!
  - What is the *post-order traversal* of this tree? Hint, it ends with “4” and the 3rd element printed is “2”.
  - What is the *pre-order traversal* of this tree? Hint, the last element is the same as the last element of the in-order traversal (but that is not true in general! why not?)
- What is the traversal order of the `destroy_tree` function we wrote earlier?

- Now let's write code to print out the elements in a binary tree in each of these three orders. These functions are easy to write recursively, and the code for the three functions looks amazingly similar. Here's the code for an in-order traversal to print the contents of a tree:

```
void print_in_order(ostream& ostr, const TreeNode<T>* p) {
    if (p) {
        print_in_order(ostr, p->left);
        ostr << p->value << "\n";
        print_in_order(ostr, p->right);
    }
}
```

How would you modify this code to perform pre-order and post-order traversals?

## 17.5 Depth-first vs. Breadth-first Search

- We should also discuss two other important tree traversal terms related to problem solving and searching.
  - In a *depth-first* search, we greedily follow links down into the tree, and don't backtrack until we have hit a leaf.
 

When we hit a leaf we step back out, but only to the last decision point and then proceed to the next leaf. This search method will quickly investigate leaf nodes, but if it has made "incorrect" branch decision early in the search, it will take a long time to work back to that point and go down the "right" branch.
  - In a *breadth-first* search, the nodes are visited with priority based on their distance from the root, with nodes closer to the root visited first.
 

In other words, we visit the nodes by level, first the root (level 0), then all children of the root (level 1), then all nodes 2 links from the root (level 2), etc.

If there are multiple solution nodes, this search method will find the solution node with the shortest path to the root node.

However, the breadth-first search method is memory-intensive, because the implementation must store all nodes at the current level – and the worst case number of nodes on each level doubles as we progress down the tree!
- Both depth-first and breadth-first will eventually visit all elements in the tree.
- Note: The ordering of elements visited by depth-first and breadth-first is not fully specified.
  - In-order, pre-order, and post-order are all *examples* of depth-first tree traversals.
  - What is a breadth-first traversal of the elements in our sample binary search tree above?

## 17.6 General-Purpose Breadth-First Search/Tree Traversal

- Write an algorithm to print the nodes in the tree one tier at a time, that is, in a *breadth-first* manner.

- What is the best/average/worst-case running time of this algorithm? What is the best/average/worst-case memory usage of this algorithm? Give a specific example tree that illustrates each case.

## 17.7 Insert

- Move left and right down the tree based on comparing keys. The goal is to find the location to do an insert *that preserves the binary search tree ordering property*.
- We will always inserting at an empty (NULL) pointer location. Why does this work? Is there always a place to put the new item? Is there ever more than one place to put the new item?
- IMPORTANT NOTE: Passing pointers by reference ensures that the new node is truly inserted into the tree. This is subtle but important.
- Note how the return value pair is constructed.

**Exercise:** How does the order that the nodes are inserted affect the final tree structure? Give an ordering that produces a balanced tree and an insertion ordering that produces a highly unbalanced tree.

## 17.8 Erase

First we need to find the node to remove. Once it is found, the actual removal is easy if the node has no children or only one child. It is harder if there are two children:

- Find the node with the greatest value in the left subtree or the node with the smallest value in the right subtree.
- The value in this node may be safely moved into the current node because of the tree ordering.
- Then we recursively apply erase to remove that node — which is guaranteed to have at most one child.

**Exercise:** Write a recursive version of erase.

**Exercise:** How does the order that nodes are deleted affect the tree structure? Starting with a mostly balanced tree, give an erase ordering that yields an unbalanced tree.

```

// -----
// TREE NODE CLASS
template <class T>
class TreeNode {
public:
    TreeNode() : left(NULL), right(NULL) {}
    TreeNode(const T& init) : value(init), left(NULL), right(NULL) {}
    T value;
    TreeNode* left;
    TreeNode* right;
};

// -----
// TREE NODE ITERATOR CLASS
template <class T>
class tree_iterator {
public:
    tree_iterator() : ptr_(NULL) {}
    tree_iterator(TreeNode<T>* p) : ptr_(p) {}
    tree_iterator(const tree_iterator& old) : ptr_(old.ptr_) {}
    ~tree_iterator() {}
    tree_iterator& operator=(const tree_iterator& old) { ptr_ = old.ptr_; return *this; }
    // operator* gives constant access to the value at the pointer
    const T& operator*() const { return ptr_->value; }
    // comparisons operators are straightforward
    bool operator==(const tree_iterator& rgt) { return ptr_ == rgt.ptr_; }
    bool operator!=(const tree_iterator& rgt) { return ptr_ != rgt.ptr_; }
    // increment & decrement will be discussed in Lecture 18

private:
    // representation
    TreeNode<T>* ptr_;
};

// -----
// DS_SET CLASS
template <class T>
class ds_set {
public:
    ds_set() : root_(NULL), size_(0) {}
    ds_set(const ds_set<T>& old) : size_(old.size_) { root_ = this->copy_tree(old.root_); }
    ~ds_set() { this->destroy_tree(root_); }
    ds_set& operator=(const ds_set<T>& old) {
        if (old != *this) {
            this->destroy_tree(root_);
            root_ = this->copy_tree(old.root_);
            size_ = old.size_;
        }
        return *this;
    }
    typedef tree_iterator<T> iterator;
    int size() const { return size_; }
    bool operator==(const ds_set<T>& old) const { return (old.root_ == this->root_); }

    // FIND, INSERT & ERASE
    iterator find(const T& key_value) { return find(key_value, root_); }
    std::pair< iterator, bool > insert(T const& key_value) { return insert(key_value, root_); }
    int erase(T const& key_value) { return erase(key_value, root_); }

    // OUTPUT & PRINTING
    friend std::ostream& operator<< (std::ostream& ostr, const ds_set<T>& s) {
        s.print_in_order(ostr, s.root_);
        return ostr;
    }
    void print_as_sideways_tree(std::ostream& ostr) const { print_as_sideways_tree(ostr, root_, 0); }
};

```

```

// ITERATORS
iterator begin() const {
    /* Implemented in Lecture 17 */
}
iterator end() const { return iterator(NULL); }

private:
// REPRESENTATION
TreeNode<T>* root_;
int size_;

// PRIVATE HELPER FUNCTIONS
TreeNode<T>* copy_tree(TreeNode<T>* old_root) { /* Implemented in Lab 10 */ }
void destroy_tree(TreeNode<T>* p) { /* Implemented in Lecture 17 */ }

iterator find(const T& key_value, TreeNode<T>* p) {
    /* Implemented in Lecture 17 */
}

std::pair<iterator,bool> insert(const T& key_value, TreeNode<T>* &p) {
    if (!p) {
        p = new TreeNode<T>(key_value);
        this->size_++;
        return std::pair<iterator,bool>(iterator(p), true);
    }
    else if (key_value < p->value)
        return insert(key_value, p->left);
    else if (key_value > p->value)
        return insert(key_value, p->right);
    else
        return std::pair<iterator,bool>(iterator(p), false);
}

int erase(T const& key_value, TreeNode<T>* &p) {
    /* Implemented in Lecture 17 */
}

void print_in_order(std::ostream& ostr, const TreeNode<T>* p) const {
    if (p) {
        print_in_order(ostr, p->left);
        ostr << p->value << "\n";
        print_in_order(ostr, p->right);
    }
}

void print_as_sideways_tree(std::ostream& ostr, const TreeNode<T>* p, int depth) const {
    if (p) {
        print_as_sideways_tree(ostr, p->right, depth+1);
        for (int i=0; i<depth; ++i) ostr << "  ";
        ostr << p->value << "\n";
        print_as_sideways_tree(ostr, p->left, depth+1);
    }
}
};

```