

CSCI-1200 Computer Science II — Fall 2008

Lecture 18 – Trees, Part III

Review from Lecture 17

- implementing cs2set operations: insert, destroy, printing, erase

Today's Lecture

- Adding parent pointers, insert & erase w/ parent pointers
- Increment and decrement operations on iterators
- Limitations of our implementation, brief intro to red-black trees
- HW6 contest results

18.1 Tree Height, Longest & Shortest Paths

- (From last lecture...) The height of a node in a tree is the length of the longest path down the tree from that node to a leaf node. The height of a leaf is therefore 0. We will think of the height of a null pointer as -1.
- The height of the tree is the height of the root node, and therefore if the tree is empty the height will be -1.
- Last time we wrote a simple recursive algorithm to calculate the height of a tree. What was the running time of this algorithm?

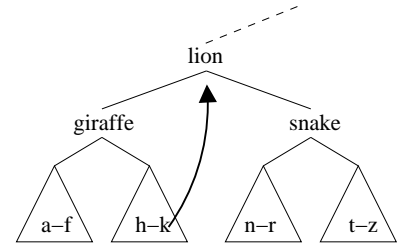
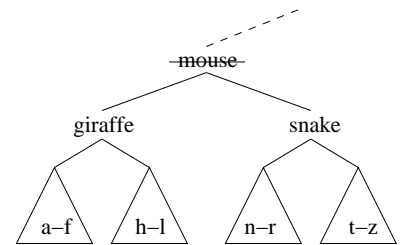
- Now let's write a function to calculate the shortest path to a NULL child pointer. What is the running time of this algorithm? Can we do better?

18.2 Tree Iterators (review and with more detail)

- The increment operator should change the iterator's pointer to point to the next `TreeNode` in an in-order traversal — the “in-order successor” — while the decrement operator should change the iterator's pointer to point to the “in-order predecessor”.
- Unlike the situation with lists and vectors, these predecessors and successors are not necessarily “nearby” (either in physical memory or by following a link) in the tree, as examples we draw in class will illustrate.
- There are two common solution approaches:
 - Each iterator maintains a stack of pointers representing the path down the tree to the current node.
 - Each node stores a parent pointer. Only the root node has a null parent pointer.
- Let's go with the parent pointer method. Most of the changes are simple, but we need to revise the implementation of several member functions, including `insert` and `erase`, to correctly adjust parent pointers. See the attached source code.

18.3 Erase (now with parent pointers)

- (Review from last lecture...) 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.



18.4 Exercise

Rewrite erase, now with parent pointers.

18.5 Exercise

- Although iterator increment looks expensive in the worst case for a single application of `operator++`, it is fairly easy to show that iterating through a tree storing n nodes requires $O(n)$ operations overall.
- Implement an algorithm for finding the in-order successor of a node. Use this to implement the iterator increment function.

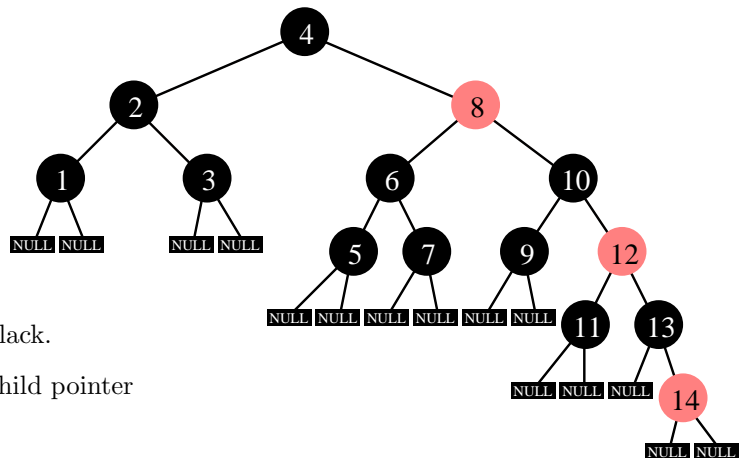
18.6 Limitations of Our BST Implementation

- The efficiency of the main insert, find and erase algorithms depends on the height of the tree.
- The best-case and average-case heights of a binary search tree storing n nodes are both $O(\log n)$. The worst-case, which often can happen in practice, is $O(n)$.
- Developing more sophisticated algorithms to avoid the worst-case behavior will be covered in Data Structures and Algorithms.

18.7 Red-Black Trees

In addition to the binary search tree properties, the following red-black tree properties are maintained throughout all modifications to the data structure:

1. Each node is either red or black.
2. The NULL child pointers are black.
3. Both children of every red node are black.
Thus, the parent of a red node must also be black.
4. All paths from a particular node to a NULL child pointer contain the same number of black nodes.



What tree does our `cs2set` implementation produce if we insert the numbers 1-14 *in order*?

The tree at the right is the result using a red-black tree. Notice how the tree is still quite balanced.

Visit this link for an animation of the sequential insertion and rebalancing:

http://babbage.clarku.edu/~achou/cs160fall03/examples/bst_animation/RedBlackTree-Example.html

```

// -----
// TREE NODE CLASS
template <class T>
class TreeNode {
public:
    TreeNode() : left(NULL), right(NULL), parent(NULL) {}
    TreeNode(const T& init) : value(init), left(NULL), right(NULL), parent(NULL) {}
    T value;
    TreeNode* left;
    TreeNode* right;
    TreeNode* parent; // to allow implementation of iterator increment & decrement
};

// -----
// 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
    friend bool operator==(const tree_iterator& lft, const tree_iterator& rgt)
    { return lft.ptr_ == rgt.ptr_; }
    friend bool operator!=(const tree_iterator& lft, const tree_iterator& rgt)
    { return lft.ptr_ != rgt.ptr_; }
    // increment & decrement operators
    tree_iterator<T> & operator++() { /* implemented in Lecture 18 */

        return *this;
    }
    tree_iterator<T> operator++(int) { tree_iterator<T> temp(*this); ++(*this); return temp; }
    tree_iterator<T> & operator--() { /* implemented in Lab 11 */ }
    tree_iterator<T> operator--(int) { tree_iterator<T> temp(*this); --(*this); return temp; }
private:
    // representation
    TreeNode<T>* ptr_;
};

// -----
// CS2 SET CLASS
template <class T>
class cs2set {
public:
    cs2set() : root_(NULL), size_(0) {}
    cs2set(const cs2set<T>& old) : size_(old.size_) { root_ = this->copy_tree(old.root_,NULL); }
    ~cs2set() { this->destroy_tree(root_); root_ = NULL; }
    cs2set& operator=(const cs2set<T>& old) { /* implementation omitted */ }

    typedef tree_iterator<T> iterator;
    int size() const { return size_; }
    bool operator==(const cs2set<T>& old) const { return (old.root_ == this->root_); }

    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_, NULL); }
    int erase(T const& key_value) { return erase(key_value, root_); }
};

```

```

// MAKE SURE THE DATA STRUCTURE'S CHILD & PARENT POINTERS ARE CONSISTENT
bool sanity_check() const {
    if (root_ == NULL) return true;
    if (root_>parent != NULL) return false;
    return sanity_check(root_);
}
iterator begin() const {
    if (!root_) return iterator(NULL);
    TreeNode<T>* p = root_;
    while (p->left) p = p->left;
    return iterator(p);
}
iterator end() const { return iterator(NULL); }

private:
// REPRESENTATION
TreeNode<T>* root_;
int size_;
// now with parent pointers, need to pass the current node as the parent pointer for the child
TreeNode<T>* copy_tree(TreeNode<T>* old_root, TreeNode<T>* the_parent) {
    if (old_root == NULL) return NULL;
    TreeNode<T> *answer = new TreeNode<T>();
    answer->value = old_root->value;
    answer->left = copy_tree(old_root->left, answer);
    answer->right = copy_tree(old_root->right, answer);
    answer->parent = the_parent; // link up the parent
    return answer;
}
void destroy_tree(TreeNode<T>* p) { /* implementation omitted */ }
iterator find(const T& key_value, TreeNode<T>* p) { /* implementation omitted */ }
// now with parent pointers, need to pass the current node as the parent pointer for the child
std::pair<iterator, bool> insert(const T& key_value, TreeNode<T>* p, TreeNode<T>* the_parent) {
    if (!p) {
        p = new TreeNode<T>(key_value);
        p->parent = the_parent; // link up the parent
        this->size_++;
        return std::pair<iterator, bool>(iterator(p), true);
    }
    else if (key_value < p->value)
        return insert(key_value, p->left, p);
    else if (key_value > p->value)
        return insert(key_value, p->right, p);
    else
        return std::pair<iterator, bool>(iterator(p), false);
}

int erase(T const& key_value, TreeNode<T>* &p) { /* Implemented in Lecture 18 */

}

bool sanity_check(TreeNode<T>* p) const {
    if (p == NULL) return true;
    if (p->left != NULL && p->left->parent != p) { return false; }
    if (p->right != NULL && p->right->parent != p) { return false; }
    return sanity_check(p->left) && sanity_check(p->right);
}
};

```