

CS062

DATA STRUCTURES AND ADVANCED PROGRAMMING

20: Left-Leaning Red-Black Trees



Alexandra Papoutsaki
she/her/hers

Lecture 20: Left-leaning Red-Black Trees


- ▶ Introduction
- ▶ Elementary red-black BST operations
- ▶ Insertion
- ▶ Mathematical analysis
- ▶ Historical context

Left-leaning red-black BSTs correspond 1-1 with 2-3 trees


- ▶ Start with standard BSTs which are made up of 2-nodes.
- ▶ Add extra information to encode 3-nodes. We will introduce two types of links.
- ▶ **Red links**: bind together two 2-nodes to represent a 3-node.
 - ▶ Specifically, 3-nodes are represented as two 2-nodes connected by a single red link that leans left (one of the 2-nodes is the left child of the other).
- ▶ **Black links**: bind together the 2-3 tree.
- ▶ Advantage: Can use BST code with minimal modification.

Left-leaning red-black BSTs correspond 1-1 with 2-3 trees


red-black BST





horizontal red links



2-3 tree



3-node



1-1 correspondence between red-black BSTs and 2-3 trees

Search

- ▶ Exactly the same as for elementary BSTs (we ignore the color).
 - ▶ But runs faster because of better balance.

```
public Value get(Key key) {
    if (key == null) throw new IllegalArgumentException("argument to get() is null");
    return get(root, key);
}

// value associated with the given key in subtree rooted at x; null if no such key
private Value get(Node x, Key key) {
    while (x != null) {
        int cmp = key.compareTo(x.key);
        if (cmp < 0) x = x.left;
        else if (cmp > 0) x = x.right;
        else return x.val;
    }
    return null;
}
```

- ▶ Operations such as floor, iteration, rank, selection are also identical.

Representation


- ▶ Each node is pointed to by one node, its parent. We can use this to encode the color of the links in nodes.
- ▶ True if the link from the parent is red and false if it is black. Null links are black.

```
private static final boolean RED    = true;
private static final boolean BLACK = false;

private Node root;    // root of the BST


// BST helper node data type
private class Node {
    private Key key;           // key
    private Value val;        // associated data
    private Node left, right; // links to left and right subtrees
    private boolean color;    // color of parent link
    private int size;         // subtree count

    private boolean isRed(Node x) {
        if (x == null) return false;
        return x.color == RED;
    }
}
```




Practice Time

- ▶ Which of the following are legal LLRB trees?



Answer



- ▶ Which of the following are legal LLRB trees?
- ▶ iii and iv
- ▶ i is not balanced and ii is also not in symmetrical order



Lecture 20: Left-leaning Red-Black Trees


- ▶ Introduction
- ▶ Elementary red-black BST operations
- ▶ Insertion
- ▶ Mathematical analysis
- ▶ Historical context

Left rotation: Orient a (temporarily) right-leaning red link to lean left





```

Node rotateLeft(Node h)
{
    Node x = h.right;
    h.right = x.left;
    x.left = h;
    x.color = h.color;
    h.color = RED;
    x.N = h.N;
    h.N = 1 + size(h.left)
        + size(h.right);
    return x;
}
    
```




Left rotate (right link of h)

Right rotation: Orient a left-leaning red link to a (temporarily) lean right




```
Node rotateRight(Node h)
{
    Node x = h.left;
    h.left = x.right;
    x.right = h;
    x.color = h.color;
    h.color = RED;
    x.N = h.N;
    h.N = 1 + size(h.left)
        + size(h.right);
    return x;
}
```




Right rotate (left link of h)

Color flip: Recolor to split a (temporary) 4-node



```
void flipColors(Node h)
{
    h.color = RED;
    h.left.color = BLACK;
    h.right.color = BLACK;
}
```



Flipping colors to split a 4-node

Lecture 20: Left-leaning Red-Black Trees

- ▶ Introduction
- ▶ Elementary red-black BST operations
- ▶ **Insertion**
- ▶ Mathematical analysis
- ▶ Historical context

Basic strategy: Maintain 1-1 correspondence with 2-3 trees


- ▶ During internal operations, maintain:
 - ▶ symmetric order
 - ▶ perfect black balance.
- ▶ But we might violate color invariants. For example:
 - ▶ Right-leaning red link.
 - ▶ Two red children (temporary 4-node).
 - ▶ Left-left red (temporary 4-node).
 - ▶ Left-right red (temporary 4-node).
- ▶ To restore color invariant we will be performing **rotations** and **color flips**.

Insertion into a LLRB

- ▶ Do standard BST insertion and color the new link red.
- ▶ Repeat until color invariants restored:
 - ▶ Both children red? Flip colors.
 - ▶ Right link red? Rotate left.
 - ▶ Two left reds in a row? Rotate right.

Red-black BST construction demo

red-black BST



Implementation


- ▶ Only three cases:
 - ▶ Right child red; left child black: rotate left.
 - ▶ Left child red; left-left grandchild red: rotate right.
 - ▶ Both children red: flip colors.

```
// insert the key-value pair in the subtree rooted at h
private Node put(Node h, Key key, Value val) {
    if (h == null) return new Node(key, val, RED, 1);

    int cmp = key.compareTo(h.key);
    if (cmp < 0) h.left = put(h.left, key, val);
    else if (cmp > 0) h.right = put(h.right, key, val);
    else
        h.val = val;

    // fix-up any right-leaning links
    if (isRed(h.right) && !isRed(h.left)) h = rotateLeft(h);
    if (isRed(h.left) && isRed(h.left.left)) h = rotateRight(h);
    if (isRed(h.left) && isRed(h.right)) flipColors(h);
    h.size = size(h.left) + size(h.right) + 1;

    return h;
}
```



Visualization of insertion into a LLRB tree

- ▶ 255 insertions in ascending order.

Visualization of insertion into a LLRB tree

- ▶ 255 insertions in descending order.

Visualization of insertion into a LLRB tree


- ▶ 255 insertions in random order.

Practice Time - Worksheet #20

- ▶ Draw the LLRB tree that results when you insert the keys 10, 18, 7, 15, 16, 30, 25, 40, 60 in that order in an initially empty tree.

ANSWER - Worksheet #20

- ▶ Draw the LLRB tree that results when you insert the keys 10, 18, 7, 15, 16, 30, 25, 40, 60 in that order in an initially empty tree.



Lecture 20: Left-leaning Red-Black Trees

- ▶ Introduction
- ▶ Elementary red-black BST operations
- ▶ Insertion
- ▶ **Mathematical analysis**
- ▶ Historical context

Balance in LLRB trees

- ▶ Height of LLRB trees is $\leq 2 \log n$ in the worst case.
- ▶ Worst case is a 2-3 tree that is all 2-nodes except that the left-most path is made up of 3-nodes.
- ▶ All ordered operations (min, max, floor, ceiling) etc. are also $O(\log n)$.

Lecture 20: Left-leaning Red-Black Trees

- ▶ Introduction
- ▶ Elementary red-black BST operations
- ▶ Insertion
- ▶ Mathematical analysis
- ▶ **Historical context**

Red-black trees

- ▶ A dichromatic framework for balanced trees. [Guibas and Sedgwick, 1978].
- ▶ Why red-black? Xerox PARC had a laser printer and red and black had the best contrast...
- ▶ Left-leaning red-black trees [Sedgwick, 2008]
 - ▶ Inspired by difficulties in proper implementation of RB BSTs.
- ▶ RB BSTs have been involved in lawsuit because of improper implementation.

Balanced trees in the wild

- ▶ Red-black trees are widely used as system dictionaries.
 - ▶ e.g., Java: `java.util.TreeMap` and `java.util.TreeSet`.
- ▶ Other balanced BSTs: AVL, splay, randomized.
- ▶ 2-3 search trees are a subset of b-trees.
 - ▶ See recommended textbook for more.
 - ▶ B-trees are widely used for file systems and databases.

Readings:

- ▶ Recommended Textbook: Chapter 3.3 (Pages 424-447)
- ▶ Website:
 - ▶ <https://algs4.cs.princeton.edu/33balanced/>
- ▶ Visualization:
 - ▶ <https://algs4.cs.princeton.edu/GrowingTree/> (for LLRB trees)

Worksheet:

- ▶ [Lecture 20 worksheet](#)

Problem 1

- ▶ Draw the left-leaning red-black BST that results when you insert items with the keys E, A, S, Y, Q, U, T, I, O, N in that order into an initially empty tree.

ANSWER 1

- ▶ Draw the left-leaning red-black BST that results when you insert items with the keys E, A, S, Y, Q, U, T, I, O, N in that order into an initially empty tree.

