

# **Lecture 7: Red Black Trees (1997)**

**Steven Skiena**


Department of Computer Science  
State University of New York  
Stony Brook, NY 11794-4400

<http://www.cs.sunysb.edu/~skiena>


*Describe a Red-Black tree with the largest and smallest ratio of red nodes.*

---

To minimize the ratio of red-black nodes, make all black (possible for  $n = 2^k - 1$ )



To maximize the ratio of red nodes, interleave with red nodes as *real* leaves



$$\#red/n \leq n/2 + n/8 + n/32 = .656n$$


$$\#black/n \leq n/4 + n/16 + n/64 = .328n$$

$$n - n \sum_{i=1}^{\max} (1/4)^i = 2n/3$$

# Rotations

---

The basic restructuring step for binary search trees are left and right rotation:



1. Rotation is a local operation changing  $O(1)$  pointers.
2. An in-order search tree before a rotation *stays* an in-order search tree.

3. In a rotation, one subtree gets one level closer to the root and one subtree one level further from the root.

LEFT-ROTATE( $T, x$ )

$y \leftarrow \text{right}[x]$  (\* Set  $y$ \*)

$\text{right}[x] \leftarrow \text{left}[y]$  (\* Turn  $y$ 's left into  $x$ 's right\*)

if  $\text{left}[y] = \text{NIL}$

    then  $p[\text{left}[y]] \leftarrow x$

$p[y] \leftarrow p[x]$  (\* Link  $x$ 's parent to  $y$  \*)

if  $p[x] = \text{NIL}$

    then  $\text{root}[T] \leftarrow y$

    else if  $x = \text{left}[p[x]]$

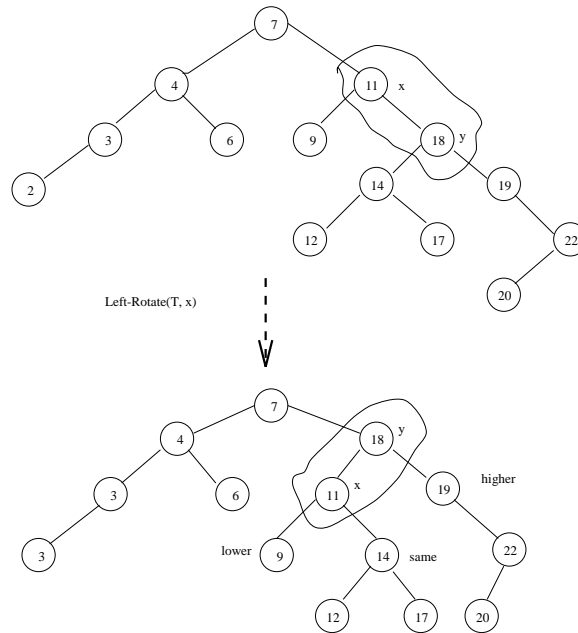
        then  $\text{left}[p[x]] \leftarrow y$

        else  $\text{right}[p[x]] \leftarrow y$

$\text{left}[y] \leftarrow x$

$p[x] \leftarrow y$

Note the in-order property is preserved.



## Red-Black Insertion


---

Since red-black trees have  $\Theta(\lg n)$  height, if we can preserve all properties of such trees under insertion/deletion, we have a balanced tree!

Suppose we just did a regular insertion. Under what conditions does it stay a red-black tree?

Since every insertion take places at a leaf, we will change a black NIL pointer to a node with two black NIL pointers.



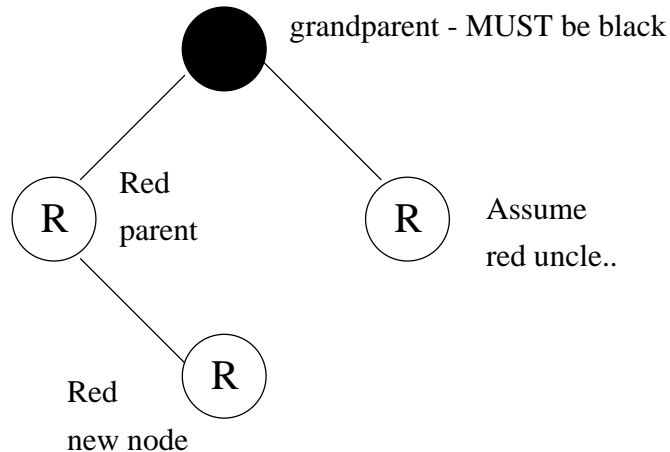


To preserve the black height of the tree, the new node must be *red*. If its *new* parent is black, we can stop, otherwise we must restructure!


# How can we fix two reds in a row?

---

It depends upon our uncle's color:



If our uncle is red, reversing our relatives' color either solves the problem or pushes it higher!



Note that after the recoloring:

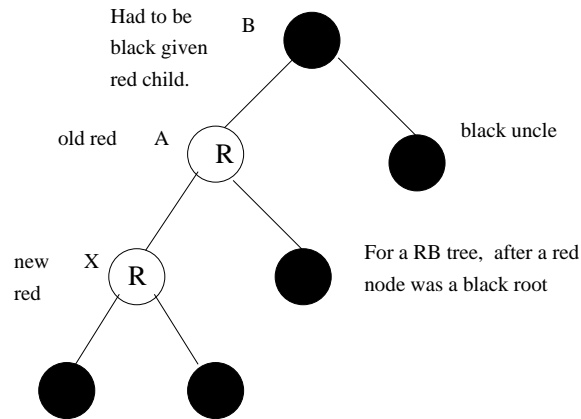
1. The black height is unchanged.
2. The shape of the tree is unchanged.
3. We are done if our great-grandparent is black.

If we get all the way to the root, recall we can always color a red-black tree's root black. We always will, so initially it *was* black, and so this process terminates.

# The Case of the Black Uncle


---

If our uncle was black, observe that all the nodes around us have to be black:



Left as RB trees by our color change or are nil

Solution - rotate right about B:



Since the root of the subtree is now black with the same black-height as before, we have restored the colors and can stop!

A double rotation can be required to set things up depending upon the left-right turn sequence, but the principle is the same.

## DOUBLE ROTATION ILLUSTRATION

# Pseudocode and Figures

---




# Deletion from Red-Black Trees


---

Recall the three cases for deletion from a binary tree:


Case (a) The node to be deleted was a leaf;



Case (b) The node to be deleted had one child;



Case (c) relabel to node as its successor and delete the successor.



## Deletion Color Cases

---

Suppose the node we remove was *red*, do we still have a red-black tree?

*Yes!* No two reds will be together, and the black height for each leaf stays the same.

However, if the dead node  $y$  was black, we must give each of its descendants another black ancestor. If an appropriate node is red, we can simply color it black otherwise we must restructure.

Case (a) black NIL becomes “double black”;

Case (b) red  $\beta$  becomes black and black  $\beta$  becomes “double black”;

Case (c) red  $\beta$  becomes black and black  $\beta$  becomes “double

black”.

Our goal will be to recolor and restructure the tree so as to get rid of the “double black” node.

In setting up any case analysis, we must be sure that:

1. All possible cases are covered.
2. No case is covered twice.

In the case analysis for red-black trees, the breakdown is:

Case 1: The double black node  $x$  has a red brother.

Case 2:  $x$  has a black brother and two black nephews.

Case 3:  $x$  has a black brother, and its left nephew is red and its right nephew is black.

Case 4:  $x$  has a black brother, and its right nephew is red (left nephew can be any color).

## Conclusion

---

Red-Black trees let us implement all dictionary operations in  $O(\log n)$ . Further, in no case are more than 3 rotations done to rebalance. Certain very advanced data structures have data stored at nodes which requires a lot of work to adjust after a rotation — red-black trees ensure it won't happen often.

*Example:* Each node represents the endpoint of a line, and is augmented with a list of segments in its subtree which it intersects.

We will not study such complicated structures, however.