



Presentation for use with the textbook *Data Structures and Algorithms in Java, 6th edition*, by M. T. Goodrich, R. Tamassia, and M. H. Goldwasser, Wiley, 2014

Red-Black Trees




From (2,4) to Red-Black Trees

- ◆ A red-black tree is a representation of a (2,4) tree by means of a binary tree whose nodes are colored **red** or **black**
- ◆ In comparison with its associated (2,4) tree, a red-black tree has
 - same logarithmic time performance
 - simpler implementation with a single node type



Red-Black Trees

- ◆ A red-black tree can also be defined as a binary search tree that satisfies the following properties:
 - **Root Property:** the root is black
 - **External Property:** every leaf is black
 - **Internal Property:** the children of a red node are black
 - **Depth Property:** all the leaves have the same black depth



© 2014 Goodrich, Tamassia, Goldwasser

Red-Black Trees

3

Height of a Red-Black Tree

- ◆ **Theorem:** A red-black tree storing n items has height $O(\log n)$
 Proof:
 - The height of a red-black tree is at most twice the height of its associated (2,4) tree, which is $O(\log n)$
- ◆ The search algorithm for a binary search tree is the same as that for a binary search tree
- ◆ By the above theorem, searching in a red-black tree takes $O(\log n)$ time


© 2014 Goodrich, Tamassia, Goldwasser

Red-Black Trees

4

Insertion

- ◆ To insert (k, o) , we execute the insertion algorithm for binary search trees and color **red** the newly inserted node z unless it is the root
 - We preserve the root, external, and depth properties
 - If the parent v of z is black, we also preserve the internal property and we are done
 - Else (v is red) we have a **double red** (i.e., a violation of the internal property), which requires a reorganization of the tree
- ◆ Example where the insertion of 4 causes a double red:



© 2014 Goodrich, Tamassia, Goldwasser

Red-Black Trees


5

Remedying a Double Red

- ◆ Consider a double red with child z and parent v , and let w be the sibling of v


Case 1: w is black

- The double red is an incorrect replacement of a 4-node
- **Restructuring:** we change the 4-node replacement



Case 2: w is red

- The double red corresponds to an overflow
- **Recoloring:** we perform the equivalent of a **split**




© 2014 Goodrich, Tamassia, Goldwasser

Red-Black Trees

6

Restructuring

- ◆ A restructuring remedies a child-parent double red when the parent red node has a black sibling
- ◆ It is equivalent to restoring the correct replacement of a 4-node
- ◆ The internal property is restored and the other properties are preserved




© 2014 Goodrich, Tamassia, Goldwasser

Red-Black Trees

7

Restructuring (cont.)

- ◆ There are four restructuring configurations depending on whether the double red nodes are left or right children




© 2014 Goodrich, Tamassia, Goldwasser

Red-Black Trees

8

Recoloring

- ◆ A recoloring remedies a child-parent double red when the parent red node has a red sibling
- ◆ The parent v and its sibling w become black and the grandparent u becomes red, unless it is the root
- ◆ It is equivalent to performing a split on a 5-node
- ◆ The double red violation may propagate to the grandparent u



© 2014 Goodrich, Tamassia, Goldwasser

Red-Black Trees

9

Analysis of Insertion

Algorithm *insert(k, o)*

1. We search for key k to locate the insertion node z
2. We add the new entry (k, o) at node z and color z red
3. **while** *doubleRed*(z)
 - if** *isBlack*(*sibling*(*parent*(z)))
 - $z \leftarrow$ *restructure*(z)
 - return**
 - else** { *sibling*(*parent*(z)) is red }
 - $z \leftarrow$ *recolor*(z)

- ◆ Recall that a red-black tree has $O(\log n)$ height
- ◆ Step 1 takes $O(\log n)$ time because we visit $O(\log n)$ nodes
- ◆ Step 2 takes $O(1)$ time
- ◆ Step 3 takes $O(\log n)$ time because we perform
 - $O(\log n)$ recolorings, each taking $O(1)$ time, and
 - at most one restructuring taking $O(1)$ time
- ◆ Thus, an insertion in a red-black tree takes $O(\log n)$ time


© 2014 Goodrich, Tamassia, Goldwasser

Red-Black Trees

10

Deletion

- ◆ To perform operation `remove(k)`, we first execute the deletion algorithm for binary search trees
- ◆ Let v be the internal node removed, w the external node removed, and r the sibling of w
 - If either v or r was red, we color r black and we are done
 - Else (v and r were both black) we color r **double black**, which is a violation of the internal property requiring a reorganization of the tree
- ◆ Example where the deletion of 8 causes a double black:



© 2014 Goodrich, Tamassia, Goldwasser

Red-Black Trees

11

Remedying a Double Black

- ◆ The algorithm for remedying a double black node w with sibling y considers three cases
 - Case 1:** y is black and has a red child
 - We perform a **restructuring**, equivalent to a **transfer**, and we are done
 - Case 2:** y is black and its children are both black
 - We perform a **recoloring**, equivalent to a **fusion**, which may propagate up the double black violation
 - Case 3:** y is red
 - We perform an **adjustment**, equivalent to choosing a different representation of a 3-node, after which either Case 1 or Case 2 applies
- ◆ Deletion in a red-black tree takes $O(\log n)$ time

© 2014 Goodrich, Tamassia, Goldwasser

Red-Black Trees

12

Red-Black Tree Reorganization

Insertion		remedy double red
Red-black tree action	(2,4) tree action	result
restructuring	change of 4-node representation	double red removed
recoloring	split	double red removed or propagated up
Deletion		remedy double black
Red-black tree action	(2,4) tree action	result
restructuring	transfer	double black removed
recoloring	fusion	double black removed or propagated up
adjustment	change of 3-node representation	restructuring or recoloring follows

© 2014 Goodrich, Tamassia, Goldwasser

Red-Black Trees

13

Java Implementation

```

1  /** An implementation of a sorted map using a red-black tree. */
2  public class RBTreeMap<K,V> extends TreeMap<K,V> {
3  /** Constructs an empty map using the natural ordering of keys. */
4  public RBTreeMap() { super(); }
5  /** Constructs an empty map using the given comparator to order keys. */
6  public RBTreeMap(Comparator<K> comp) { super(comp); }
7  // we use the inherited aux field with convention that 0=black and 1=red
8  // (note that new leaves will be black by default, as aux=0)
9  private boolean isBlack(Position<Entry<K,V>> p) { return tree.getAux(p)==0;}
10 private boolean isRed(Position<Entry<K,V>> p) { return tree.getAux(p)==1; }
11 private void makeBlack(Position<Entry<K,V>> p) { tree.setAux(p, 0); }
12 private void makeRed(Position<Entry<K,V>> p) { tree.setAux(p, 1); }
13 private void setColor(Position<Entry<K,V>> p, boolean toRed) {
14     tree.setAux(p, toRed ? 1 : 0);
15 }
16 /** Overrides the TreeMap rebalancing hook that is called after an insertion. */
17 protected void rebalanceInsert(Position<Entry<K,V>> p) {
18     if (!isRoot(p)) {
19         makeRed(p); // the new internal node is initially colored red
20         resolveRed(p); // but this may cause a double-red problem
21     }
22 }

```

© 2014 Goodrich, Tamassia, Goldwasser

Red-Black Trees

14

Java Implementation, 2

```

23  /** Remedies potential double-red violation above red position p. */
24  private void resolveRed(Position<Entry<K,V>> p) {
25      Position<Entry<K,V>> parent,uncle,middle,grand; // used in case analysis
26      parent = parent(p);
27      if (isRed(parent)) { // double-red problem exists
28          uncle = sibling(parent);
29          if (isBlack(uncle)) { // Case 1: misshapen 4-node
30              middle = restructure(p); // do trinode restructuring
31              makeBlack(middle);
32              makeRed(left(middle));
33              makeRed(right(middle));
34          } else { // Case 2: overfull 5-node
35              makeBlack(parent); // perform recoloring
36              makeBlack(uncle);
37              grand = parent(parent);
38              if (!isRoot(grand)) {
39                  makeRed(grand); // grandparent becomes red
40                  resolveRed(grand); // recur at red grandparent
41              }
42          }
43      }
44  }

```

© 2014 Goodrich, Tamassia, Goldwasser

Red-Black Trees

15

Java Implementation, 3

```

45  /** Overrides the TreeMap rebalancing hook that is called after a deletion. */
46  protected void rebalanceDelete(Position<Entry<K,V>> p) {
47      if (isRed(p)) // deleted parent was black
48          makeBlack(p); // so this restores black depth
49      else if (!isRoot(p)) {
50          Position<Entry<K,V>> sib = sibling(p);
51          if (isInternal(sib) && (isBlack(sib) || isInternal(left(sib))))
52              remedyDoubleBlack(p); // sib's subtree has nonzero black height
53      }
54  }
55

```

© 2014 Goodrich, Tamassia, Goldwasser

Red-Black Trees

16

Java Implementation, 4

```

56  /** Remedies a presumed double-black violation at the given (nonroot) position. */
57  private void remedyDoubleBlack(Position<Entry<K,V>> p) {
58      Position<Entry<K,V>> z = parent(p);
59      Position<Entry<K,V>> y = sibling(p);
60      if (isBlack(y)) {
61          if (isRed(left(y)) || isRed(right(y))) { // Case 1: trinode restructuring
62              Position<Entry<K,V>> x = (isRed(left(y)) ? left(y) : right(y));
63              Position<Entry<K,V>> middle = restructure(x);
64              setColor(middle, isRed(z)); // root of restructured subtree gets z's old color
65              makeBlack(left(middle));
66              makeBlack(right(middle));
67          } else { // Case 2: recoloring
68              makeRed(y);
69              if (isRed(z))
70                  makeBlack(z); // problem is resolved
71              else if (!isRoot(z))
72                  remedyDoubleBlack(z); // propagate the problem
73          }
74      } else { // Case 3: reorient 3-node
75          rotate(y);
76          makeBlack(y);
77          makeRed(z);
78          remedyDoubleBlack(p); // restart the process at p
79      }
80  }
81  }

```