

Algorithms and Data Structures (II)

Gabriel Istrate

April 1, 2020

- AVL trees.
- Splay trees
- Red-black trees

- Want: data structure to support INSERT, DELETE, SEARCH in $O(\log n)$ time.
- Binary search trees: insert, delete, search.
- But **complexity bound not met unless trees balanced**.
- Can rebalance.

Today: three **self-adjusting trees**, finally meet the $O(\log n)$ bound: **AVL trees, splay trees, red-black-trees**.

■ Binary search trees

- ▶ embody the *divide-and-conquer* search strategy
- ▶ **SEARCH**, **INSERT**, **MIN**, and **MAX** are $O(h)$, where h is the *height of the tree*
- ▶ in general, $h(n) = \Omega(\log n)$ and $h(n) = O(n)$
- ▶ *randomization* can be used to make the worst-case scenario $h(n) = n$ highly unlikely

■ Binary search trees

- ▶ embody the *divide-and-conquer* search strategy
- ▶ **SEARCH**, **INSERT**, **MIN**, and **MAX** are $O(h)$, where h is the *height of the tree*
- ▶ in general, $h(n) = \Omega(\log n)$ and $h(n) = O(n)$
- ▶ *randomization* can be used to make the worst-case scenario $h(n) = n$ highly unlikely

■ Problem

- ▶ worst-case scenario is unlikely but still possible
- ▶ simply bad cases are even more probable

- Why study all this stuff ?
- Linked list: **search linear**.
- **Balanced** binary trees: **search logarithmic**.
- **For frequent searches it pays**.
- Advantage: as long as trees “approximately balanced”.
- But: operations (inserts/deletes) can destroy balance.

Self-balancing trees

If insertion/deletion unbalances the tree, **rebalance it**.

- AVL trees: more strictly balanced than R-B trees. Better for **lookup intensive** programs.
- For an **insert intensive tasks**, use a Red-Black tree.
- Simplicity of implementation: splay trees > red-black trees > AVL trees.
- Splay trees: only $O(\log n)$ amortized.
- Splay trees: suitable for cases where there are **large number of nodes but only few of them are accessed frequently**.
- Splay trees: more *memory-efficient* than AVL trees, because they do not need to store balance information in the nodes.
- AVL trees: more useful in multithreaded environments with lots of lookups, because lookups in an AVL tree can be done in parallel.
- Benchmarking: **AVL trees more than 20% faster than R-B trees** in "real-life" benchmarkis

- treaps
- T -trees
- tango trees

... many other ! (But this is an introduction)

Tango trees: A type of binary search tree proposed by Erik D. Demaine, Dion Harmon, John Iacono, and Mihai Pătrașcu in 2004.

They work by partitioning a BST into a set of preferred paths, which are themselves stored in auxiliary trees (so the tango tree is represented as a tree of trees)


- **red-black trees:**

- 1 **Java:** `java.util.TreeMap` , `java.util.TreeSet` .
- 2 **C++ STL:** `map`, `multimap`, `multiset`.
- 3 **Linux kernel:** completely fair scheduler, `linux/rbtree.h`

- **Splay trees:** typically used in the implementation of caches, memory allocators, routers, garbage collectors, data compression, etc.

- **Implementations** of AVL trees, RB-trees, splay trees: *not standardized*. STL provides only **minimal set of containers**.


- Balancedness condition #1: left and right subtrees of the root have the same height. **too weak.**
- Balancedness condition #2: left and right subtrees of **every node** have the same height. **too strong.**
- AVL (Adelson-Velskii and Landis) trees: binary search trees that verify the following **balancedness condition**: for every node v **the left and right subtrees of v have height differing by at most one.**
- When a tree violates rule #3 a repair is done.
- **The repair is done during insertions, as soon as rule #3 is violated.**
- The repair is accomplished via **"single" and "double" rotations.**

Single Rotation


Suppose an item is added at the bottom of subtree X , thus causing an imbalance at $k2$. Then pull $k1$ up. Note that after the rotation, the height of the tree is the same as it was before the insertion.

- **Balance factor of a node:** difference of heights between left and right subtrees.
- **AVL trees:** each node balance factor 0 or ± 1 .
- After single rotations, the new height of the entire subtree is exactly the same as the height of the original subtree prior to the insertion of the new data item that caused X to grow.
- Thus no further updating of heights on the path to the root is needed, and consequently no further rotations are needed.


Single rotations: another example




Suppose an item is added at the bottom of subtree X, thus causing an imbalance at $k2$. Then pull $k1$ up. Note that after the rotation, the height of the tree is the same as it was before the insertion.




```
void SingleRotate1 (AvlNode *&k2) const {
    AvlNode *k1 = k2->left;
    k2->left = k1->right;
    k1->right = k2;
    k2 = k1;
}
```

Double Rotation

Suppose an item is added below k_2 . This causes an imbalance at k_2 . This causes an imbalance at k_3 . Then pull k_2 up. Note that after the rotation, the height of the tree is the same as it was before the insertion.


Another Double Rotation

Suppose an item is added below k_2 . This causes an imbalance at k_3 . Then pull k_2 up. Note that after the rotation, the height of the tree is the same as it was before the insertion.



```
void DoubleRotate1 (AvlNode *&k3) const {
    SingleRotate1( k3->left );
    SingleRotate1( k3 );
}
```

An Example



Imbalance at node 8 solved with double rotation.


- Recognizing which rotation you have to use is the hardest part.
 - 1 Find the imbalanced node.
 - 2 Go down two nodes towards the newly inserted node.
 - 3 If the path is straight, use single rotation.
 - 4 If the path zig-zags, use double rotation.

- Use `deleteByCopying()` to delete a node. This allows reducing the problem of deleting a node with two descendants to deleting a node with at most one descendant.
- After a node has been deleted, balance factors updated from the parent of the deleted node to the root.
- For each node whose balance becomes ± 2 , a single or double rotation has to be performed to restore balance of the tree.
- Deletion: at most $O(\log n)$ rotations.
- Deletion might improve balance factor of its parent.
- It may also worsen the balance factor of its grandparent.

- As with the single rotations, double rotations restore the height of the subtree to what it was before the insertion.
- This guarantees that all rebalancing and height updating is complete.
- AVL trees maintain balance of binary search trees while they are being created via insertions of data.
- An alternative approach is to have trees that readjust themselves when data is accessed, making often accessed data items move to the top of the tree (splay trees).


- Invented by Sleator and Tarjan (1985).
- to splay ~ to spread out.
- Self-balancing binary trees, simpler to implement than AVL, red-black trees.
- Additional property: recently accessed elements quick to access.
- Insertion, lookup, removal: $O(\log n)$ **amortized time**.
- That roughly means that the average price per operation in a long sequence of operations is $O(\log n)$.
- Fundamental operation: **splaying**. Rearranging the tree such that certain elements brought at the top of the tree.

- When a node x is accessed, a **splaying operation** performed to bring it to the top.
- Composed of a sequence of **splaying steps**.
- Each splaying step brings x closer to the root.
- Steps depend on:
 - Whether z is left or right child of its parent p .
 - Whether p is root or not, and
 - Whether p is left or right child of **its** parent g .
- **Three types** of splaying steps.




“zig”: basically rotation.

Second case: p not the root, x, p both left or both right children





“Zigzig”





“Zigzag”

- First case: rotation.
- All cases: actually two mirror-image cases (only one shown in picture).
- Advantages: more accessed nodes closer to root. Useful for implementing caches, garbage collection.
- Disadvantages: random access worse than for other balanced BST.
- Particularly bad: access elements in sorted order.











■ *Red-black-tree property*




■ *Red-black-tree property*

- 1 every node is either **red** or **black**




■ *Red-black-tree property*

- 1 every node is either **red** or **black**
- 2 the root is **black**




■ *Red-black-tree property*

- 1 every node is either **red** or **black**
- 2 the root is **black**
- 3 every (NIL) leaf is **black**



■ *Red-black-tree property*


- 1 every node is either **red** or **black**
- 2 the root is **black**
- 3 every (NIL) leaf is **black**
- 4 if a node is **red**, then both its children are **black**




■ Red-black-tree property

- 1 every node is either **red** or **black**
- 2 the root is **black**
- 3 every (NIL) leaf is **black**
- 4 if a node is **red**, then both its children are **black**
- 5 for every node x , each path from x to its descendant leaves has the same number of **black** nodes $bh(x)$ (the *black-height* of x)

■ *Implementation*




■ *Implementation*



- ▶ we use a common "sentinel" node to represent leaf nodes

■ *Implementation*



- ▶ we use a common “sentinel” node to represent leaf nodes
- ▶ the sentinel is also the parent of the root node

- *Implementation*

- ▶ T represents the tree, which consists of a set of *nodes*

■ *Implementation*

- ▶ T represents the tree, which consists of a set of *nodes*
- ▶ $T.root$ is the root node of tree T

■ *Implementation*


- ▶ T represents the tree, which consists of a set of *nodes*
- ▶ $T.root$ is the root node of tree T
- ▶ $T.nil$ is the “sentinel” node of tree T

■ Implementation

- ▶ T represents the tree, which consists of a set of *nodes*
- ▶ $T.root$ is the root node of tree T
- ▶ $T.nil$ is the “sentinel” node of tree T

Nodes

- ▶ $x.parent$ is the parent of node x
- ▶ $x.key$ is the key stored in node x
- ▶ $x.left$ is the left child of node x
- ▶ $x.right$ is the right child of node x




■ Implementation

- ▶ T represents the tree, which consists of a set of *nodes*
- ▶ $T.root$ is the root node of tree T
- ▶ $T.nil$ is the "sentinel" node of tree T

Nodes

- ▶ $x.parent$ is the parent of node x
- ▶ $x.key$ is the key stored in node x
- ▶ $x.left$ is the left child of node x
- ▶ $x.right$ is the right child of node x
- ▶ $x.color \in \{RED, BLACK\}$ is the color of node x



Height of a Red-Black Tree

Lemma: the height $h(x)$ of a red-black tree with $n = \text{size}(x)$ internal nodes is at most $2 \log(n + 1)$.

Lemma: the height $h(x)$ of a red-black tree with $n = \text{size}(x)$ internal nodes is at most $2 \log(n + 1)$.

Proof:

① prove that $\forall x : \text{size}(x) \geq 2^{bh(x)} - 1$ by induction:

① **base case:** x is a leaf, so $\text{size}(x) = 0$ and $bh(x) = 0$

Lemma: the height $h(x)$ of a red-black tree with $n = \text{size}(x)$ internal nodes is at most $2 \log(n + 1)$.

Proof:

- 1 prove that $\forall x : \text{size}(x) \geq 2^{bh(x)} - 1$ by induction:
 - 1 **base case:** x is a leaf, so $\text{size}(x) = 0$ and $bh(x) = 0$
 - 2 **induction step:** consider y_1, y_2 , and x such that $y_1.\text{parent} = y_2.\text{parent} = x$, and assume (induction) that $\text{size}(y_1) \geq 2^{bh(y_1)} - 1$ and $\text{size}(y_2) \geq 2^{bh(y_2)} - 1$; prove that $\text{size}(x) \geq 2^{bh(x)} - 1$

Lemma: the height $h(x)$ of a red-black tree with $n = \text{size}(x)$ internal nodes is at most $2 \log(n + 1)$.

Proof:

① prove that $\forall x : \text{size}(x) \geq 2^{bh(x)} - 1$ by induction:

① **base case:** x is a leaf, so $\text{size}(x) = 0$ and $bh(x) = 0$

② **induction step:** consider y_1, y_2 , and x such that $y_1.\text{parent} = y_2.\text{parent} = x$, and assume (induction) that $\text{size}(y_1) \geq 2^{bh(y_1)} - 1$ and $\text{size}(y_2) \geq 2^{bh(y_2)} - 1$;
prove that $\text{size}(x) \geq 2^{bh(x)} - 1$

proof:

$$\text{size}(x) = \text{size}(y_1) + \text{size}(y_2) + 1 \geq (2^{bh(y_1)} - 1) + (2^{bh(y_2)} - 1) + 1$$

Lemma: the height $h(x)$ of a red-black tree with $n = \text{size}(x)$ internal nodes is at most $2 \log(n + 1)$.

Proof:

① prove that $\forall x : \text{size}(x) \geq 2^{bh(x)} - 1$ by induction:

① **base case:** x is a leaf, so $\text{size}(x) = 0$ and $bh(x) = 0$

② **induction step:** consider y_1, y_2 , and x such that $y_1.\text{parent} = y_2.\text{parent} = x$, and assume (induction) that $\text{size}(y_1) \geq 2^{bh(y_1)} - 1$ and $\text{size}(y_2) \geq 2^{bh(y_2)} - 1$;
prove that $\text{size}(x) \geq 2^{bh(x)} - 1$

proof:

$$\text{size}(x) = \text{size}(y_1) + \text{size}(y_2) + 1 \geq (2^{bh(y_1)} - 1) + (2^{bh(y_2)} - 1) + 1$$

since

$$bh(x) = \begin{cases} bh(y) & \text{if } \text{color}(y) = \text{RED} \\ bh(y) + 1 & \text{if } \text{color}(y) = \text{BLACK} \end{cases}$$

Lemma: the height $h(x)$ of a red-black tree with $n = \text{size}(x)$ internal nodes is at most $2 \log(n + 1)$.

Proof:

① prove that $\forall x : \text{size}(x) \geq 2^{bh(x)} - 1$ by induction:

① **base case:** x is a leaf, so $\text{size}(x) = 0$ and $bh(x) = 0$

② **induction step:** consider y_1, y_2 , and x such that $y_1.\text{parent} = y_2.\text{parent} = x$, and assume (induction) that $\text{size}(y_1) \geq 2^{bh(y_1)} - 1$ and $\text{size}(y_2) \geq 2^{bh(y_2)} - 1$; prove that $\text{size}(x) \geq 2^{bh(x)} - 1$

proof:

$$\text{size}(x) = \text{size}(y_1) + \text{size}(y_2) + 1 \geq (2^{bh(y_1)} - 1) + (2^{bh(y_2)} - 1) + 1$$

since

$$bh(x) = \begin{cases} bh(y) & \text{if } \text{color}(y) = \text{RED} \\ bh(y) + 1 & \text{if } \text{color}(y) = \text{BLACK} \end{cases}$$

$$\text{size}(x) \geq (2^{bh(x)-1} - 1) + (2^{bh(x)-1} - 1) + 1$$

Lemma: the height $h(x)$ of a red-black tree with $n = \text{size}(x)$ internal nodes is at most $2 \log(n + 1)$.

Proof:

① prove that $\forall x : \text{size}(x) \geq 2^{bh(x)} - 1$ by induction:

① **base case:** x is a leaf, so $\text{size}(x) = 0$ and $bh(x) = 0$

② **induction step:** consider y_1, y_2 , and x such that $y_1.\text{parent} = y_2.\text{parent} = x$, and assume (induction) that $\text{size}(y_1) \geq 2^{bh(y_1)} - 1$ and $\text{size}(y_2) \geq 2^{bh(y_2)} - 1$;
prove that $\text{size}(x) \geq 2^{bh(x)} - 1$

proof:

$$\text{size}(x) = \text{size}(y_1) + \text{size}(y_2) + 1 \geq (2^{bh(y_1)} - 1) + (2^{bh(y_2)} - 1) + 1$$

since

$$bh(x) = \begin{cases} bh(y) & \text{if } \text{color}(y) = \text{RED} \\ bh(y) + 1 & \text{if } \text{color}(y) = \text{BLACK} \end{cases}$$

$$\text{size}(x) \geq (2^{bh(x)-1} - 1) + (2^{bh(x)-1} - 1) + 1 = 2^{bh(x)} - 1$$

1 $size(x) \geq 2^{bh(x)} - 1$

- 1 $size(x) \geq 2^{bh(x)} - 1$
- 2 Since every red node has black children, in every path from x to a leaf node, at least half the nodes are black

- 1 $size(x) \geq 2^{bh(x)} - 1$
- 2 Since every red node has black children, in every path from x to a leaf node, at least half the nodes are black, thus $bh(x) \geq h(x)/2$

- 1 $size(x) \geq 2^{bh(x)} - 1$
- 2 Since every red node has black children, in every path from x to a leaf node, at least half the nodes are black, thus $bh(x) \geq h(x)/2$
- 3 From steps 1 and 2, $n = size(x) \geq 2^{h(x)/2} - 1$

- 1 $size(x) \geq 2^{bh(x)} - 1$
- 2 Since every red node has black children, in every path from x to a leaf node, at least half the nodes are black, thus $bh(x) \geq h(x)/2$
- 3 From steps 1 and 2, $n = size(x) \geq 2^{h(x)/2} - 1$, therefore

$$h \leq 2 \log(n + 1)$$

1 $size(x) \geq 2^{bh(x)} - 1$

2 Since every red node has black children, in every path from x to a leaf node, at least half the nodes are black, thus $bh(x) \geq h(x)/2$

3 From steps 1 and 2, $n = size(x) \geq 2^{h(x)/2} - 1$, therefore


$$h \leq 2 \log(n + 1)$$


■ A red-black tree works as a binary search tree for search, etc.


■ So, the complexity of those operations is $T(n) = O(h)$, that is

$$T(n) = O(\log n)$$


▶ which is also the *worst-case* complexity







■ $x = \text{RIGHT-ROTATE}(x)$



■ $x = \text{RIGHT-ROTATE}(x)$

■ $x = \text{LEFT-ROTATE}(x)$

- **RB-INSERT**(T, z) works as in a binary search tree


- **RB-INSERT**(T, z) works as in a binary search tree
- Except that it must preserve the *red-black-tree property*


- **RB-INSERT**(T, z) works as in a binary search tree
- Except that it must preserve the *red-black-tree property*
 - ① every node is either **red** or **black**
 - ② the root is **black**
 - ③ every (NIL) leaf is **black**
 - ④ if a node is **red**, then both its children are **black**
 - ⑤ for every node x , each path from x to its descendant leaves has the same number of **black** nodes $bh(x)$ (the *black-height* of x)


- **RB-INSERT**(T, z) works as in a binary search tree
- Except that it must preserve the *red-black-tree property*
 - 1 every node is either **red** or **black**
 - 2 the root is **black**
 - 3 every (NIL) leaf is **black**
 - 4 if a node is **red**, then both its children are **black**
 - 5 for every node x , each path from x to its descendant leaves has the same number of **black** nodes $bh(x)$ (the *black-height* of x)
- *General strategy*


- **RB-INSERT**(T, z) works as in a binary search tree
- Except that it must preserve the *red-black-tree property*
 - ① every node is either **red** or **black**
 - ② the root is **black**
 - ③ every (NIL) leaf is **black**
 - ④ if a node is **red**, then both its children are **black**
 - ⑤ for every node x , each path from x to its descendant leaves has the same number of **black** nodes $bh(x)$ (the *black-height* of x)
- *General strategy*
 - ① insert z as in a binary search tree
 - ② color z **red** so as to preserve property 5
 - ③ *fix the tree* to correct possible violations of property 4


```
RB-INSERT(T, z) 1 y = T.nil
                   2 x = T.root
                   3 while x ≠ T.nil
                   4     y = x
                   5     if z.key < x.key
                   6       x = x.left
                   7     else x = x.right
                   8 z.parent = y
                   9 if y == T.nil
                  10   T.root = z
                  11 else if z.key < y.key
                  12   y.left = z
                  13   else y.right = z
                  14 z.left = z.right = T.nil
                  15 z.color = RED
                  16 RB-INSERT-FIXUP(T, z)
```
















- z's father is **black**, so no fixup needed















































- A **black** node can become **red** and transfer its **black** color to its two children





- A **black** node can become **red** and transfer its **black** color to its two children
- This may cause other **red-red** conflicts, so we iterate...





- A **black** node can become **red** and transfer its **black** color to its two children
- This may cause other **red-red** conflicts, so we iterate...
- The root can change to **black** without causing conflicts


































- An *in-line red-red* conflicts can be resolved with a rotation plus a color switch




































- A zig-zag **red-red** conflicts can be resolved with a rotation to turn it into an *in-line* conflict, and then a rotation plus a color switch


Recap on Deletion in Binary Trees

Recap on Deletion in Binary Trees




Recap on Deletion in Binary Trees

1. z has no children




Recap on Deletion in Binary Trees




1. z has no children
 - ▶ simply remove z

Recap on Deletion in Binary Trees




1. z has no children
 - ▶ simply remove z

Recap on Deletion in Binary Trees




1. z has no children
 - ▶ simply remove z
2. z has one child

Recap on Deletion in Binary Trees




1. z has no children
 - ▶ simply remove z
2. z has one child
 - ▶ remove z

Recap on Deletion in Binary Trees




1. z has no children
 - ▶ simply remove z
2. z has one child
 - ▶ remove z
 - ▶ connect $z.parent$ to $z.right$

Recap on Deletion in Binary Trees




1. z has no children
 - ▶ simply remove z
2. z has one child
 - ▶ remove z
 - ▶ connect $z.parent$ to $z.right$

Recap on Deletion in Binary Trees




1. z has no children
 - ▶ simply remove z
2. z has one child
 - ▶ remove z
 - ▶ connect $z.parent$ to $z.right$
3. z has two children

Recap on Deletion in Binary Trees





1. z has no children
 - ▶ simply remove z
2. z has one child
 - ▶ remove z
 - ▶ connect $z.parent$ to $z.right$
3. z has two children
 - ▶ replace z with $y = \mathbf{TREE-SUCCESSOR}(z)$
 - ▶ remove y (1 child!)


Recap on Deletion in Binary Trees





1. z has no children
 - ▶ simply remove z
2. z has one child
 - ▶ remove z
 - ▶ connect $z.parent$ to $z.right$
3. z has two children
 - ▶ replace z with $y = \mathbf{TREE-SUCCESSOR}(z)$
 - ▶ remove y (1 child!)
 - ▶ connect $y.parent$ to $y.right$














- A deleting a **red leaf** does not require any adjustment





- A deleting a **red leaf** does not require any adjustment
 - ▶ the deletion does not affect the black height (property 5)





- A deleting a **red leaf** does not require any adjustment
 - ▶ the deletion does not affect the black height (property 5)
 - ▶ no two red nodes become adjacent (property 4)













- Deleting a **black** node changes the balance of black-height in a subtree x



- Deleting a **black** node changes the balance of black-height in a subtree x
 - ▶ **RB-DELETE-FIXUP**(T, x) fixes the tree after a deletion



- Deleting a **black** node changes the balance of black-height in a subtree x
 - ▶ **RB-DELETE-FIXUP**(T, x) fixes the tree after a deletion
 - ▶ in this simple case: $x.color = \text{BLACK}$


- y is the spliced node ($y = z$ if z has zero or one child)
 - ▶ if y is **red**, then no fixup is necessary
 - ▶ so, here we assume that y is **black**


- y is the spliced node ($y = z$ if z has zero or one child)
 - ▶ if y is **red**, then no fixup is necessary
 - ▶ so, here we assume that y is **black**
- x is either y 's only child or $T.nil$
 - ▶ y was spliced out, so y can not have two children
 - ▶ $x = T.nil$ iff y has no (key-bearing) children


- y is the spliced node ($y = z$ if z has zero or one child)
 - ▶ if y is **red**, then no fixup is necessary
 - ▶ so, here we assume that y is **black**
- x is either y 's only child or $T.nil$
 - ▶ y was spliced out, so y can not have two children
 - ▶ $x = T.nil$ iff y has no (key-bearing) children
- **Problem 1:** $y = T.root$ and x is **red**
 - ▶ violates red-black property ?? ($root$ must be **black**)

- y is the spliced node ($y = z$ if z has zero or one child)
 - ▶ if y is **red**, then no fixup is necessary
 - ▶ so, here we assume that y is **black**
- x is either y 's only child or $T.nil$
 - ▶ y was spliced out, so y can not have two children
 - ▶ $x = T.nil$ iff y has no (key-bearing) children
- **Problem 1:** $y = T.root$ and x is **red**
 - ▶ violates red-black property ?? ($root$ must be **black**)
- **Problem 2:** both x and $y.parent$ are **red**
 - ▶ violates red-black property 4 (no adjacent red nodes)


- y is the spliced node ($y = z$ if z has zero or one child)
 - ▶ if y is **red**, then no fixup is necessary
 - ▶ so, here we assume that y is **black**
- x is either y 's only child or $T.nil$
 - ▶ y was spliced out, so y can not have two children
 - ▶ $x = T.nil$ iff y has no (key-bearing) children
- **Problem 1:** $y = T.root$ and x is **red**
 - ▶ violates red-black property ?? (*root* must be **black**)
- **Problem 2:** both x and $y.parent$ are **red**
 - ▶ violates red-black property 4 (no adjacent red nodes)
- **Problem 3:** we are removing y , which is black
 - ▶ violates red-black property 5 (same *black height* for all paths)









- x carries an *additional **black** weight*
 - ▶ the fixup algorithm pushes it up towards to root




- x carries an *additional **black** weight*
 - ▶ the fixup algorithm pushes it up towards to root




- x carries an *additional **black** weight*
 - ▶ the fixup algorithm pushes it up towards to root





- x carries an *additional **black** weight*
 - ▶ the fixup algorithm pushes it up towards to root





- x carries an *additional **black** weight*
 - ▶ the fixup algorithm pushes it up towards to root
- The *additional **black** weight* can be discarded if it reaches the *root*, otherwise...
















- The *additional black weight* can also stop as soon as it reaches a **red** node, which will absorb the extra **black** color






















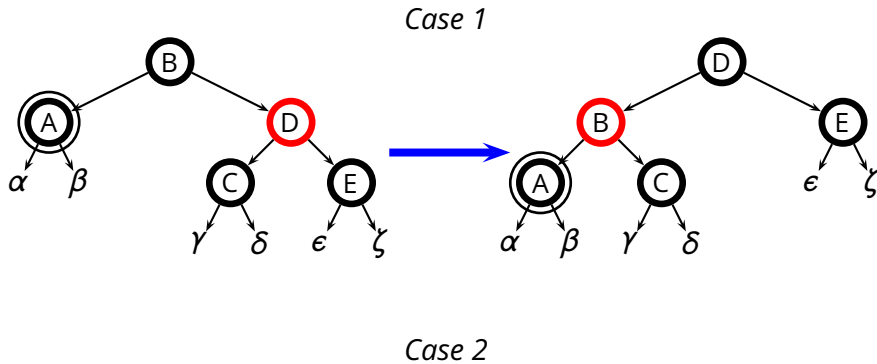
- In other cases where we can not push the additional black color up, we can apply appropriate rotations and color transfers that preserve all other red-black properties



Case 1

Case 1











Case 1





Case 2





Case 3


Case 3








Case 3



Case 4




```

RB-DELETE-FIXUP( $T, x$ ) 1  while  $x \neq T.root \wedge x.color = BLACK$ 
2      if  $x == x.parent.left$ 
3           $w = x.parent.right$ 
4          if  $w.color == RED$ 
5              case 1...
6          if  $w.left.color == BLACK \wedge w.right.color = BLACK$ 
7               $w.color = RED$                                 // case 2
8               $x = x.parent$ 
9          else if  $w.right.color == BLACK$ 
10             case 3...
11             case 4...
12          else same as above, exchanging right and left
13   $x.color = BLACK$ 

```