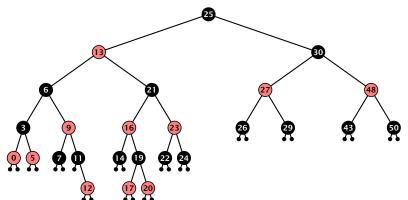
Definition 1

A red black tree is a balanced binary search tree in which each internal node has two children. Each internal node has a color, such that

- 1. The root is black.
- 2. All leaf nodes are black.
- 3. For each node, all paths to descendant leaves contain the same number of black nodes.
- 4. If a node is red then both its children are black.

The null-pointers in a binary search tree are replaced by pointers to special null-vertices, that do not carry any object-data

Red Black Trees: Example



Lemma 2

A red-black tree with n internal nodes has height at most $\mathcal{O}(\log n)$.

Definition 3

The black height bh(v) of a node v in a red black tree is the number of black nodes on a path from v to a leaf vertex (not counting v).

We first show:

Lemma 4

A sub-tree of black height $\mathrm{bh}(v)$ in a red black tree contains at least $2^{\mathrm{bh}(v)}-1$ internal vertices.

Proof of Lemma 4.

Induction on the height of v.

base case (height(v) = 0)

- If height(v) (maximum distance btw. v and a node in the sub-tree rooted at v) is 0 then v is a leaf.
- ightharpoonup The black height of v is 0.
- ► The sub-tree rooted at v contains $0 = 2^{bh(v)} 1$ inner vertices.

Proof (cont.)

induction step

- Supose v is a node with height(v) > 0.
- ightharpoonup v has two children with strictly smaller height.
- ► These children (c_1 , c_2) either have $bh(c_i) = bh(v)$ or $bh(c_i) = bh(v) 1$.
- **By** induction hypothesis both sub-trees contain at least $2^{\text{bh}(v)-1}-1$ internal vertices.
- ► Then T_v contains at least $2(2^{\text{bh}(v)-1}-1)+1 \ge 2^{\text{bh}(v)}-1$ vertices.



Proof of Lemma 2.

Let h denote the height of the red-black tree, and let P denote a path from the root to the furthest leaf.

At least half of the node on P must be black, since a red node must be followed by a black node.

Hence, the black height of the root is at least h/2.

The tree contains at least $2^{h/2}-1$ internal vertices. Hence, $2^{h/2}-1 \le n$.

Hence, $h \le 2\log(n+1) = \mathcal{O}(\log n)$.



Definition 1

A red black tree is a balanced binary search tree in which each internal node has two children. Each internal node has a color, such that

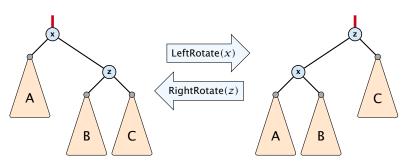
- 1. The root is black.
- 2. All leaf nodes are black.
- 3. For each node, all paths to descendant leaves contain the same number of black nodes.
- 4. If a node is red then both its children are black.

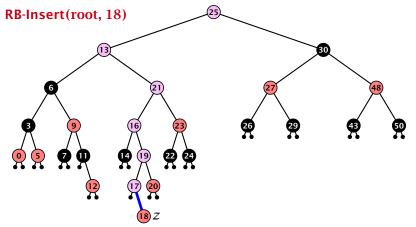
The null-pointers in a binary search tree are replaced by pointers to special null-vertices, that do not carry any object-data.

We need to adapt the insert and delete operations so that the red black properties are maintained.

Rotations

The properties will be maintained through rotations:





Insert:

- first make a normal insert into a binary search tree
- then fix red-black properties



26/41

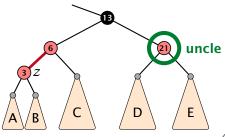
Invariant of the fix-up algorithm:

- z is a red node
- the black-height property is fulfilled at every node
- the only violation of red-black properties occurs at z and parent[z]
 - either both of them are red (most important case)
 - or the parent does not exist (violation since root must be black)

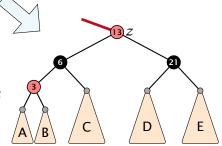
If z has a parent but no grand-parent we could simply color the parent/root black; however this case never happens.

```
Algorithm 10 InsertFix(z)
 1: while parent[z] \neq null and col[parent[z]] = red do
         if parent[z] = left[gp[z]] then z in left subtree of grandparent
 2:
 3:
               uncle \leftarrow right[grandparent[z]]
               if col[uncle] = red then
 4:
                                                                 Case 1: uncle red
                    col[p[z]] \leftarrow black; col[u] \leftarrow black;
 5:
                    col[gp[z]] \leftarrow red; z \leftarrow grandparent[z];
 6:
 7:
              else
                                                               Case 2: uncle black
                    if z = right[parent[z]] then
 8:
                                                                  2a: z right child
                         z \leftarrow p[z]; LeftRotate(z);
 9:
                    col[p[z]] \leftarrow black; col[gp[z]] \leftarrow red; 2b: z left child
10:
11:
                    RightRotate(gp[z]);
12:
         else same as then-clause but right and left exchanged
13: \operatorname{col}(\operatorname{root}[T]) \leftarrow \operatorname{black};
```

Case 1: Red Uncle



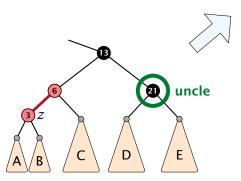
- 1. recolour
- 2. move z to grand-parent
- 3. invariant is fulfilled for new z
- 4. you made progress

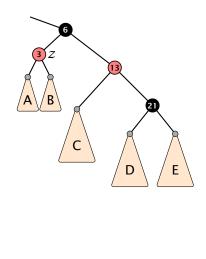


7.2 Red Black Trees

Case 2b: Black uncle and z is left child

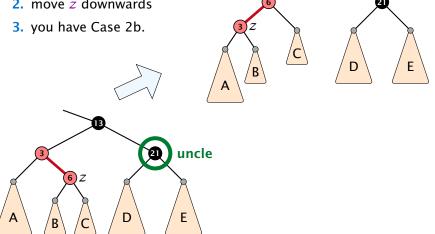
- 1. rotate around grandparent
- re-colour to ensure that black height property holds
- 3. you have a red black tree





Case 2a: Black uncle and z is right child

- 1. rotate around parent
- 2. move z downwards



Running time:

- ▶ Only Case 1 may repeat; but only h/2 many steps, where h is the height of the tree.
- Case 2a → Case 2b → red-black tree
- Case 2b → red-black tree

Performing Case 1 at most $\mathcal{O}(\log n)$ times and every other case at most once, we get a red-black tree. Hence $\mathcal{O}(\log n)$ re-colorings and at most 2 rotations.

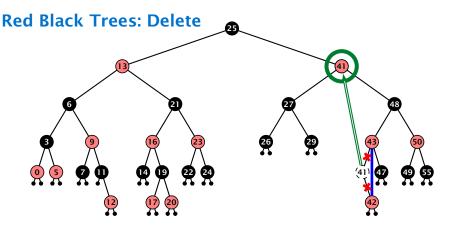
Red Black Trees: Delete

First do a standard delete.

If the spliced out node x was red everything is fine.

If it was black there may be the following problems.

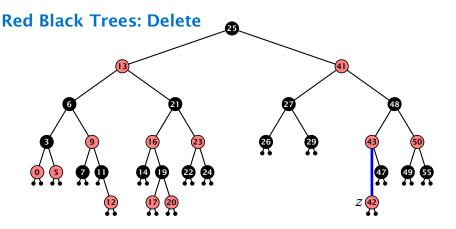
- ightharpoonup Parent and child of x were red; two adjacent red vertices.
- If you delete the root, the root may now be red.
- Every path from an ancestor of x to a descendant leaf of x changes the number of black nodes. Black height property might be violated.



Case 3:

Element has two children

- do normal delete
- when replacing content by content of successor, don't change color of node



Delete:

- deleting black node messes up black-height property
- ightharpoonup if z is red, we can simply color it black and everything is fine
- the problem is if z is black (e.g. a dummy-leaf); we call a fix-up procedure to fix the problem.

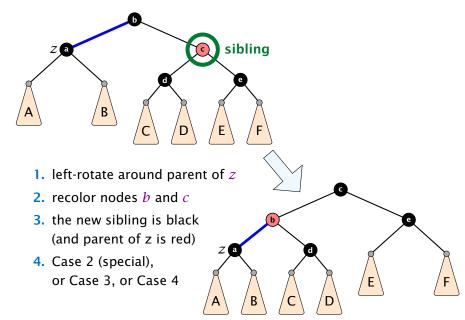
Red Black Trees: Delete

Invariant of the fix-up algorithm

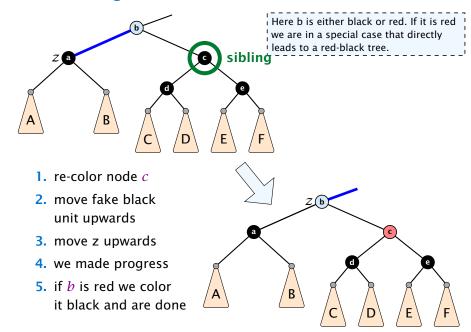
- ► the node z is black
- if we "assign" a fake black unit to the edge from z to its parent then the black-height property is fulfilled

Goal: make rotations in such a way that you at some point can remove the fake black unit from the edge.

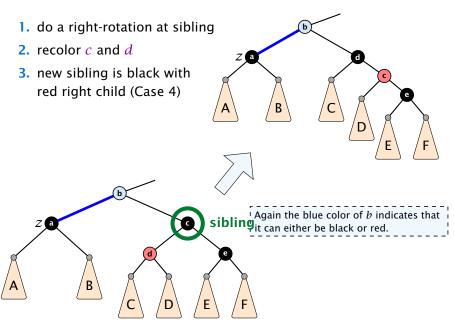
Case 1: Sibling of z is red



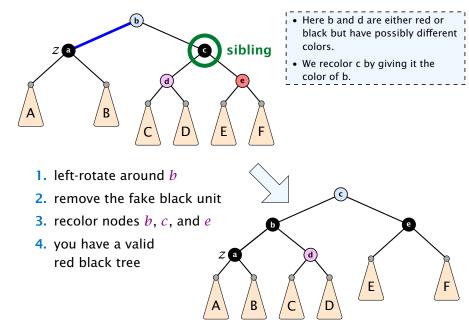
Case 2: Sibling is black with two black children



Case 3: Sibling black with one black child to the right



Case 4: Sibling is black with red right child



Running time:

- only Case 2 can repeat; but only h many steps, where h is the height of the tree
- Case 1 → Case 2 (special) → red black tree Case 1 → Case 3 → Case 4 → red black tree Case 1 → Case 4 → red black tree
- Case 3 → Case 4 → red black tree
- Case 4 → red black tree

Performing Case 2 at most $\mathcal{O}(\log n)$ times and every other step at most once, we get a red black tree. Hence, $\mathcal{O}(\log n)$ re-colorings and at most 3 rotations.

Red-Black Trees

Bibliography

[CLRS90] Thomas H. Cormen, Charles E. Leiserson, Ron L. Rivest, Clifford Stein: Introduction to Algorithms (3rd ed.), MIT Press and McGraw-Hill, 2009

Red black trees are covered in detail in Chapter 13 of [CLRS90].