

Red-Black Trees

- What is a red-black tree?
 - node color: red or black
 - $nil[T]$ and black height
- Subtree rotation
- Node insertion
- Node deletion

Red-black trees: Overview

- ♦ Red-black trees are a variation of binary search trees to ensure that the tree is *balanced*.
 - » Height is $O(\lg n)$, where n is the number of nodes.
- ♦ Operations take $O(\lg n)$ time in the **worst case**.
- ♦ A red-black tree is normally not perfectly balanced, but satisfying:

The length of the longest path from a node to a leaf is less than two times of the length of the shortest path from that node to a leaf.


Red-black Tree

- ◆ Every node in a red-black tree is associated with a bit: the attribute *color*, which is either **red** or **black**.
- ◆ All other attributes of BSTs are inherited:
 - » *key*, *left*, *right*, and *p*.
- ◆ If a child or the parent of a node does not exist, the corresponding pointer field of the node contains the value *nil*.
- ◆ Sentinel - *nil[T]*, representing all the *virtual nil* nodes.
 - A node, if it has only one child, a virtual nil child will be created. If it has no children (i.e., it is a leaf node), two virtual nil children will be created.
 - For the tree root, a virtual nil parent will be created.


Red-black Properties

1. Every node is either **red** or **black**.
2. The **root** is **black**.
3. Every *virtual* **node** (*nil*) is **black**.
4. If a node is **red**, then both its children are **black**.
5. For each node, all paths from the node to descendant leaves contain the same number of **black** nodes.


Red-black Tree – Example



Red-black Tree – Example



Red-black Tree – Example



Red-black Properties

1. Every node is either **red** or **black**.
2. The **root** is **black**.
3. Every *virtual* **leaf** (*nil*) is **black**.
4. If a node is **red**, then both its children are **black**.
5. For each node, all paths from the node to descendant leaves contain the same number of **black** nodes.

Height of a Red-black Tree

- ◆ Height of a node:

- » $h(x)$ = number of edges in a longest path to a leaf.

- ◆ Black-height of a node x , $bh(x)$:

- » $bh(x)$ = number of black nodes (including $nil[T]$) on the path from x to leaf, not counting x .

- ◆ Black-height of a red-black tree is the black-height of its root.

- » By Property 5, black height is well defined.

Height of a Red-black Tree

♦ Example:

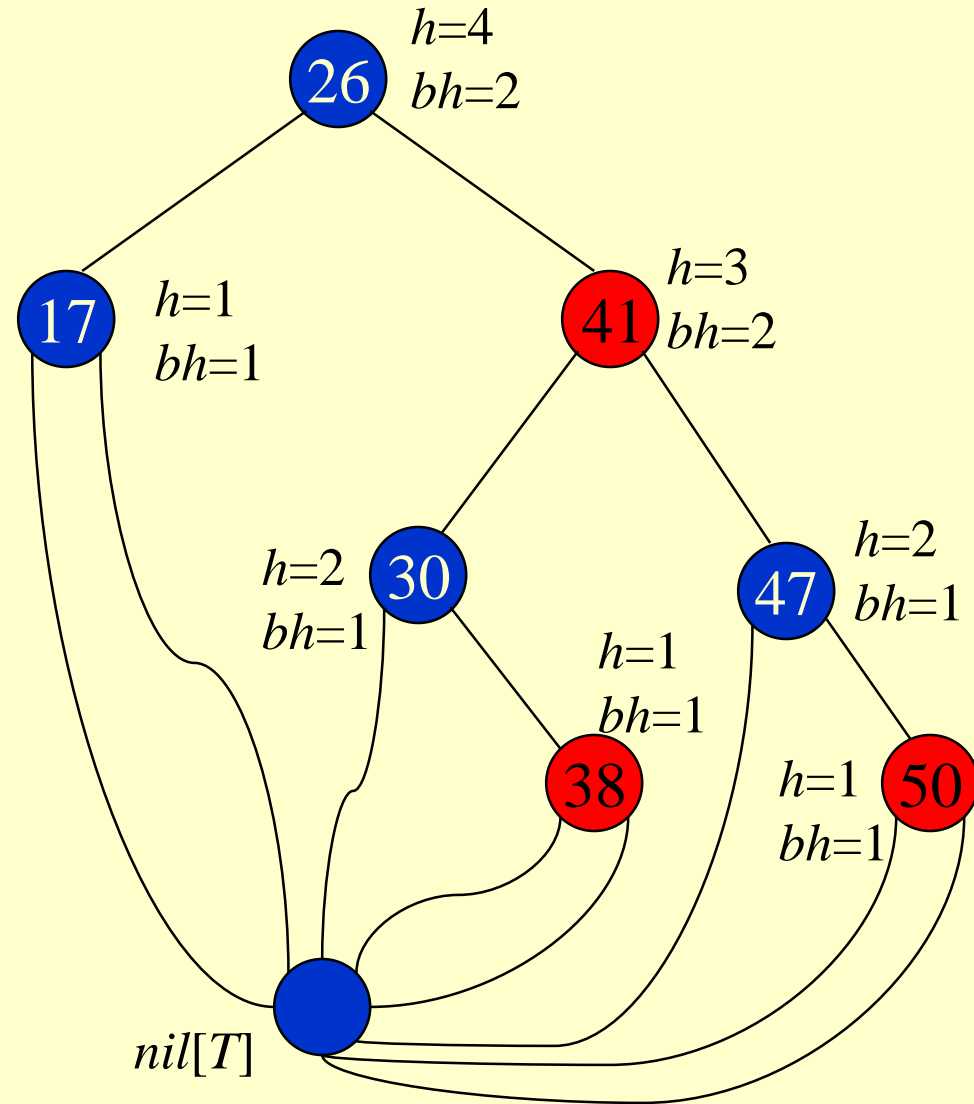
♦ Height of a node:

$h(x)$ = # of edges in a longest path to a leaf.

♦ Black-height of a node
 $bh(x)$ = # of black nodes on path from x to leaf, not counting x .

♦ How are they related?

» $bh(x) \leq h(x) \leq 2bh(x)$



Lemma “RB Height”

Consider a node x in an RB tree: The longest descending path from x to a leaf has length $h(x)$, which is at most twice the length of the shortest descending path from x to a leaf.

Proof:

black nodes on any path from $x = bh(x)$ (prop 5)
 \leq # nodes on shortest path from x , $s(x)$. (prop 1)

But, there are no consecutive red (prop 4),
and we end with black (prop 3), so $h(x) \leq 2 bh(x)$.

Thus, $h(x) \leq 2s(x)$. QED


Bound on RB Tree Height

- ♦ **Lemma:** The subtree rooted at any node x has $\geq 2^{bh(x)} - 1$ internal nodes.
- ♦ **Proof:** By induction on height of x , $h(x)$.
 - » **Base Case:** Height $h(x) = 0 \Rightarrow x$ is a leaf $\Rightarrow bh(x) = 0$. Subtree has $2^0 - 1 = 0$ internal nodes.
 - » **Induction Step:** Assume that for any node with height $< h$ the lemma holds.

Consider node x with $h(x) = h > 0$ and $bh(x) = b$.

- Each child of x has height at most $h - 1$ and black-height either b (child is red) or $b - 1$ (child is black).
- By ind. hyp., each child has $\geq 2^{bh(x)-1} - 1$ internal nodes.
- Subtree rooted at x has $\geq 2(2^{bh(x)-1} - 1) + 1$
 $= 2^{bh(x)} - 1$ internal nodes. (The $+1$ is for x itself.)

Bound on RB Tree Height



Bound on RB Tree Height

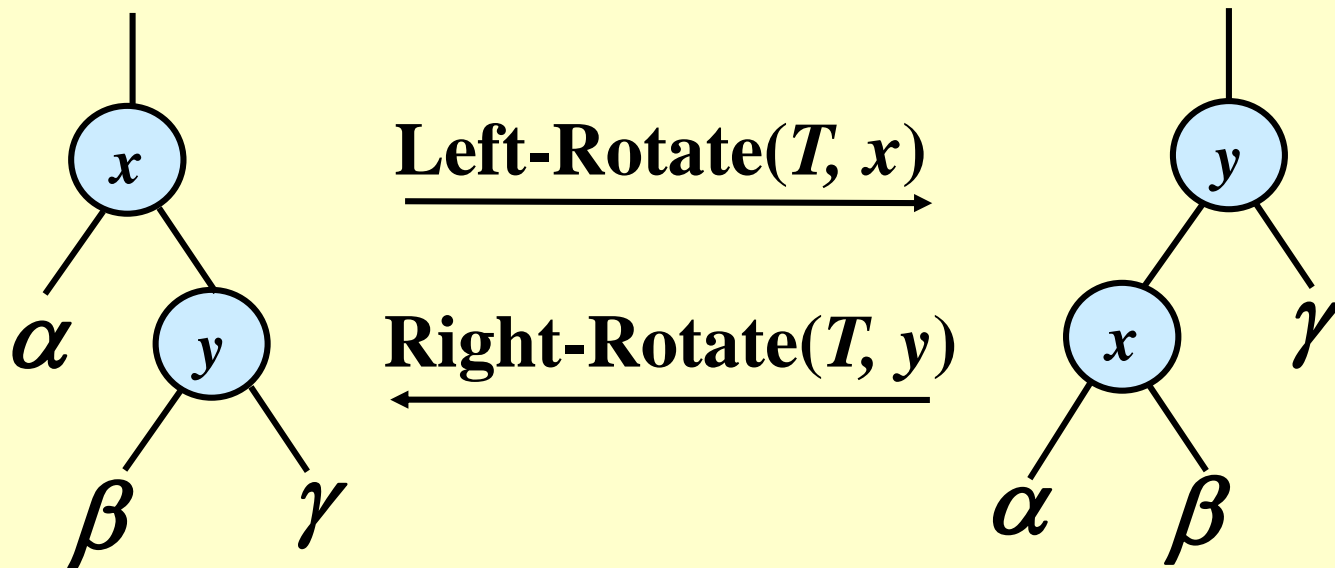
- ♦ Lemma: The subtree rooted at any node x has $\geq 2^{bh(x)} - 1$ internal nodes.
- ♦ **Lemma 13.1:** A red-black tree with n internal nodes has height at most $2\lg(n+1)$.
- ♦ **Proof:**
 - » By the above lemma, $n \geq 2^{bh} - 1$,
 - » and since $bh \geq h/2$, we have $n \geq 2^{h/2} - 1$.
 - » $\Rightarrow h \leq 2\lg(n + 1)$.

Operations on RB Trees

- ◆ All operations can be performed in $O(\lg n)$ time.
- ◆ The query operations, which don't modify the tree, are performed in exactly the same way as they are in binary search trees.
- ◆ Insertion and Deletion are not straightforward.

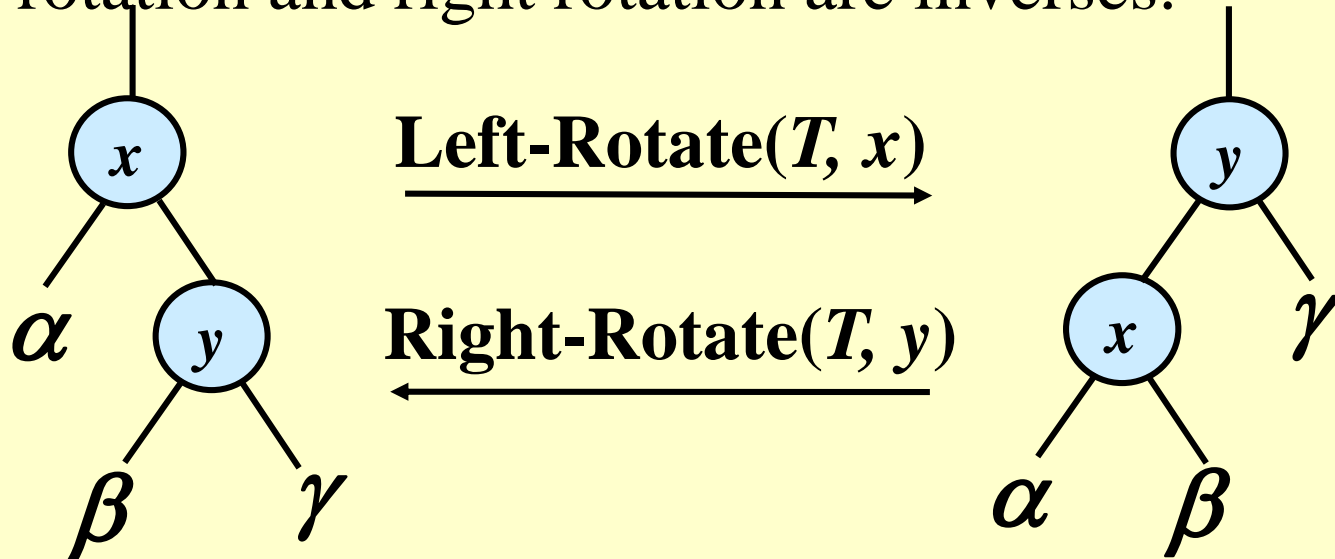
Why?

Rotations



Rotations


- ◆ Rotations are the basic **tree-restructuring** operation for almost all *balanced* search trees.
- ◆ Rotation takes a red-black-tree and a node as the input,
- ◆ Change pointers to change the local structure, and
- ◆ Won't violate the binary-search-tree property.
- ◆ Left rotation and right rotation are inverses.



Left Rotation – Pseudo-code

Left-Rotate (T, x)

1. $y \leftarrow \text{right}[x]$ // Set y .
2. $\text{right}[x] \leftarrow \text{left}[y]$ // Turn y 's left subtree β into x 's right subtree.
3. **if** $\text{left}[y] \neq \text{nil}[T]$
4. **then** $p[\text{left}[y]] \leftarrow x$ // Set x to be the parent of $\text{left}[y] = \beta$.
5. $p[y] \leftarrow p[x]$ // Link x 's parent to y .
6. **if** $p[x] = \text{nil}[T]$ // If x is the root.
7. **then** $\text{root}[T] \leftarrow y$
8. **else if** $x = \text{left}[p[x]]$
9. **then** $\text{left}[p[x]] \leftarrow y$
10. **else** $\text{right}[p[x]] \leftarrow y$
11. $\text{left}[y] \leftarrow x$ // Put x as y 's left child.
12. $p[x] \leftarrow y$



Rotation

- ◆ The pseudo-code for Left-Rotate assumes that
 - » $right[x] \neq nil[T]$, and
 - » root's parent is $nil[T]$.
- ◆ Left Rotation on x , makes x the left child of y , and the left subtree of y into the right subtree of x .
- ◆ Pseudocode for Right-Rotate is symmetric: exchange *left* and *right* accordingly.
- ◆ **Time:** $O(1)$ for both Left-Rotate and Right-Rotate, since a constant number of pointers are modified.

Reminder: Red-black Properties

1. Every node is either **red** or **black**.
2. The **root** is **black**.
3. Every **leaf** (*nil*) is **black**.
4. If a node is **red**, then both its children are **black**.
5. For each node, all paths from the node to descendant leaves contain the same number of **black** nodes.

Insertion in RB Trees

- ◆ Insertion must preserve all red-black properties.
- ◆ Should an inserted node be colored **Red**? **Black**?
- ◆ **Basic steps:**
 - » Use Tree-Insert from BST (slightly modified) to insert a node z into T .
 - Procedure **RB-Insert(z)**.
 - » Color the node z red.
 - » Fix the modified tree by re-coloring nodes and performing rotation to preserve RB tree property.
 - Procedure **RB-Insert-Fixup**.

Insertion

RB-Insert(T, z)

```
1.   $y \leftarrow nil[T]$ 
2.   $x \leftarrow root[T]$ 
3.  while  $x \neq nil[T]$ 
4.    do  $y \leftarrow x$ 
5.      if  $key[z] < key[x]$ 
6.        then  $x \leftarrow left[x]$ 
7.        else  $x \leftarrow right[x]$ 
8.   $p[z] \leftarrow y$ 
9.  if  $y = nil[T]$ 
10.   then  $root[T] \leftarrow z$ 
11.   else if  $key[z] < key[y]$ 
12.     then  $left[y] \leftarrow z$ 
13.     else  $right[y] \leftarrow z$ 
```

RB-Insert(T, z) Contd.

```
14.  $left[z] \leftarrow nil[T]$ 
15.  $right[z] \leftarrow nil[T]$ 
16.  $color[z] \leftarrow RED$ 
17. RB-Insert-Fixup( $T, z$ )
```

How does it differ from the Tree-Insert procedure of BSTs?

Which of the RB properties might be violated?


Fix the violations by calling RB-Insert-Fixup.

Insertion – Fixup


- ◆ Problem: we may have a pair of consecutive reds where we did the insertion.
- ◆ Solution: rotate it up the tree and away...
Six cases have to be handled:

Case 1:


$p[z]$ is the
left child



Case 2:




Case 3:




Case 4:


$p[z]$ is the
right child



Case 5:



Case 6:




Insertion – Fixup


RB-Insert-Fixup (T, z)

1. **while** $color[p[z]] = \text{RED}$
2. **do if** $p[z] = \text{left}[p[p[z]]]$ //for cases 1 – 3
 then $y \leftarrow \text{right}[p[p[z]]]$
3. **if** $color[y] = \text{RED}$
4. **then** $color[p[z]] \leftarrow \text{BLACK}$ // Case 1
5. $color[y] \leftarrow \text{BLACK}$ // Case 1
6. $color[p[p[z]]] \leftarrow \text{RED}$ // Case 1
7. $z \leftarrow p[p[z]]$ // Case 1

z 's parent is the left child
of its own parent



Case 1:





← Change this node to red to
keep the number of black
nodes not increased

Insertion – Fixup

RB-Insert-Fixup(T, z) (Contd.)

9. **else if** $z = \text{right}[p[z]]$ // $\text{color}[y] \neq \text{RED}$
10. **then** $z \leftarrow p[z]$ // Case 2
11. $\text{LEFT-ROTATE}(T, z)$ // Case 2
12. $\text{color}[p[z]] \leftarrow \text{BLACK}$ // Case 3
13. $\text{color}[p[p[z]]] \leftarrow \text{RED}$ // Case 3
14. $\text{RIGHT-ROTATE}(T, p[p[z]])$ // Case 3
15. **else** (if $p[z] = \text{right}[p[p[z]]]$) (for cases 4 – 6, same
16. as 3-14 with “right” and “left” exchanged)
17. $\text{color}[\text{root}[T]] \leftarrow \text{BLACK}$



Correctness

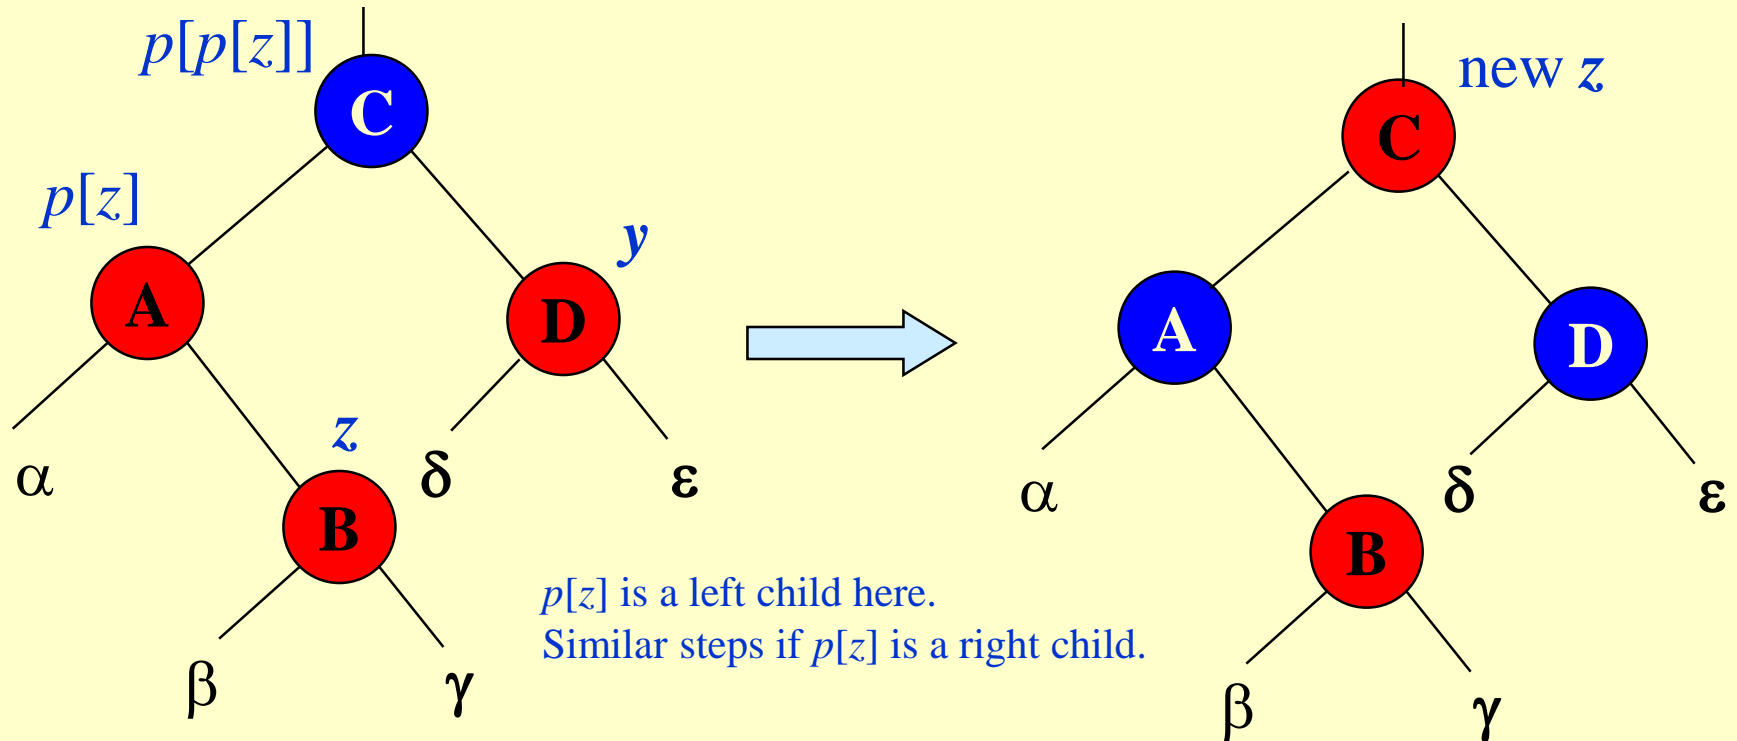
Loop invariant:

- ◆ At the start of each iteration of the **while** loop,
 - » z is red.
 - » If $p[z]$ is the root, then $p[z]$ is black.
 - » There is at most one red-black violation:
 - Property 2: z is a red root, or
 - Property 4: z and $p[z]$ are both red.

Correctness – Contd.

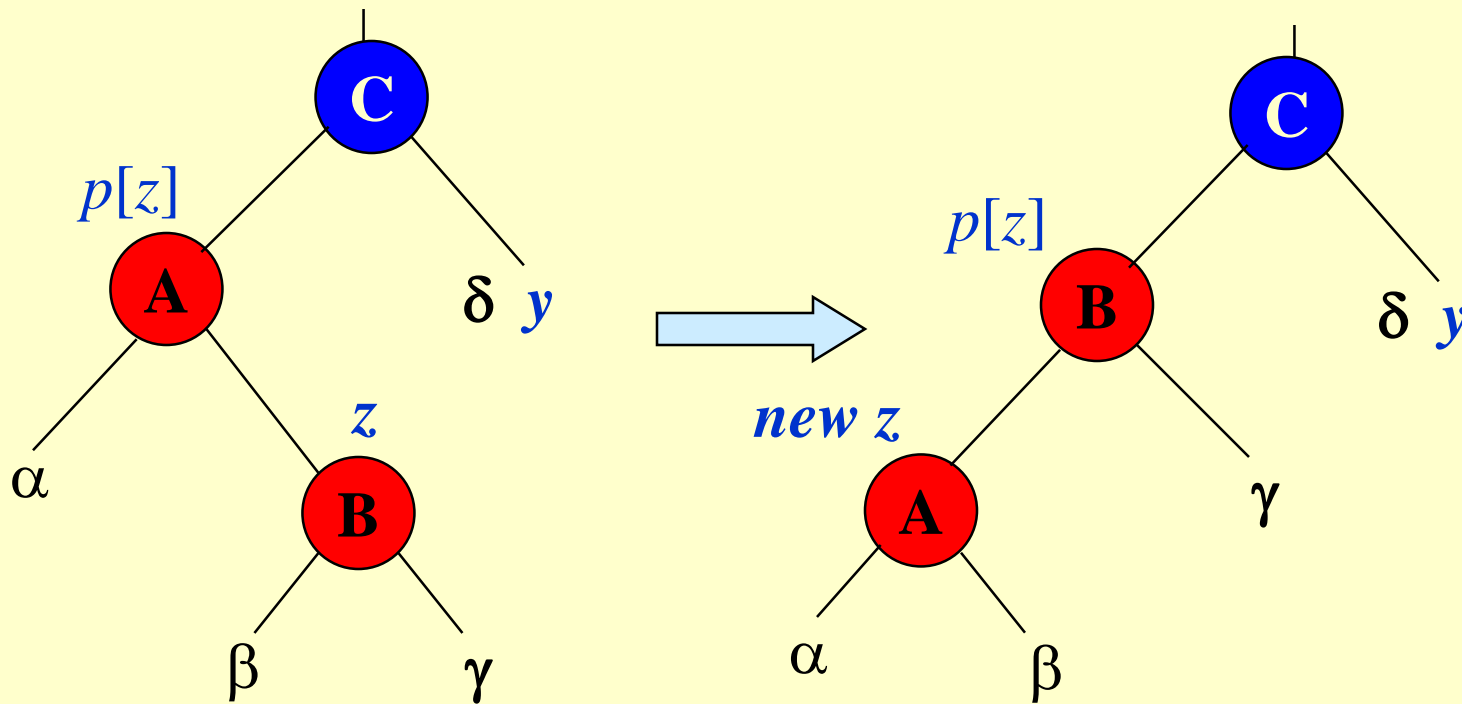
- ♦ **Initialization:** OK.
- ♦ **Termination:** The loop terminates only if $p[z]$ is black.
Hence, property 4 is OK.
The last line ensures property 2 always holds.
- ♦ **Maintenance:** We drop out when z is the root (since then $p[z]$ is sentinel $nil[T]$, which is black). When we start the loop body, the only violation is of property 4.
 - » There are 6 cases, 3 of which are symmetric to the other 3.
We consider cases in which $p[z]$ is a left child.
 - » Let y be z 's uncle ($p[z]$'s sibling).

Case 1 – uncle y is red




- ♦ $p[p[z]]$ (z 's grandparent) must be black, since z and $p[z]$ are both red and there are no other violations of property 4.
- ♦ Make $p[z]$ and y black \Rightarrow now z and $p[z]$ are not both red. But property 5 might now be violated.
- ♦ Make $p[p[z]]$ red \Rightarrow restores property 5.
- ♦ The next iteration has $p[p[z]]$ as the new z (i.e., z moves up 2 levels).

Case 2 – y is black, z is a right child



- ◆ Left rotate around $p[z]$, $p[z]$ and z switch roles \Rightarrow now z is a left child, and both z and $p[z]$ are red.
- ◆ Takes us immediately to case 3.

Case 3 – y is black, z is a left child



- ◆ Make $p[z]$ black and $p[p[z]]$ red.
- ◆ Then right rotate on $p[p[z]]$. Ensures property 4 is maintained.
- ◆ No longer have 2 reds in a row.
- ◆ $p[z]$ is now black \Rightarrow no more iterations.

Algorithm Analysis

- ◆ $O(\lg n)$ time to get through RB-Insert up to the call of RB-Insert-Fixup.
- ◆ Within **RB-Insert-Fixup**:
 - » Each iteration takes $O(1)$ time.
 - » Each iteration but the last **moves z up 2 levels**.
 - » $O(\lg n)$ levels $\Rightarrow O(\lg n)$ time.
 - » Thus, insertion in a red-black tree takes $O(\lg n)$ time.
 - » Note: there are at most 2 rotations overall.

Deletion

- ◆ Deletion, like insertion, should preserve all the RB properties.
- ◆ The properties that may be violated depends on the color of the deleted node.
 - » Red – OK. Why?
 - » Black?
- ◆ Steps:
 - » Do regular BST deletion.
 - » Fix any violations of RB properties that may be caused by a deletion.

Deletion

RB-Delete(T, z)

1. **if** $left[z] = nil[T]$ or $right[z] = nil[T]$
2. **then** $y \leftarrow z$
3. **else** $y \leftarrow \text{TREE-SUCCESSOR}(z)$
4. **if** $left[y] \neq nil[T]$
5. **then** $x \leftarrow left[y]$
6. **else** $x \leftarrow right[y]$
7. $p[x] \leftarrow p[y]$ // Do this, even if x is $nil[T]$

Deletion

RB-Delete (T, z) (Contd.)

```
8. if  $p[y] = \text{nil}[T]$ 
9.   then  $\text{root}[T] \leftarrow x$ 
10.  else if  $y = \text{left}[p[y]]$  (*if  $y$  is a left child.*)
11.    then  $\text{left}[p[y]] \leftarrow x$ 
12.    else  $\text{right}[p[y]] \leftarrow x$  (*if  $y$  is a right
13.  if  $y \neq z$                                 child.*)
14.  then  $\text{key}[z] \leftarrow \text{key}[y]$ 
15.                                copy  $y$ 's satellite data
    into  $z$ 
16. if  $\text{color}[y] = \text{BLACK}$ 
17.  then RB-Delete-Fixup( $T, x$ )
18. return  $y$ 
```

The node passed to the fixup routine is the only child of the spliced up node, or the sentinel.




RB Properties Violation

- ◆ If y is black, we could have violations of red-black properties:

- » Prop. 1. OK.

- » Prop. 2. If y is the root and x is red, then the root has become red.



- » Prop. 3. OK.


- » Prop. 4. Violation if $p[y]$ and x are both red.

- » Prop. 5. Any path containing y now has 1 fewer black node.

RB Properties Violation

♦ Prop. 5. Any path containing y now has 1 fewer black node.

- » Correct by giving x an “extra black.”
- » Add 1 to the count of black nodes on paths containing x .
- » Now property 5 is OK, but property 1 is not.
- » x is either **doubly black** (if $color[x] = \text{BLACK}$) or **red & black** (if $color[x] = \text{RED}$).
- » The attribute $color[x]$ is still either RED or BLACK. No new values for $color$ attribute.
- » In other words, the extra blackness on a node is by virtue of “ x pointing to the node”. (If a node is pointed to by x , it has an extra black.)



♦ Remove the violations by calling RB-Delete-Fixup.


Deletion – Fixup

RB-Delete-Fixup(T, x)


1. **while** $x \neq \text{root}[T]$ and $\text{color}[x] = \text{BLACK}$
2. **do if** $x = \text{left}[p[x]]$ //for cases 1 - 4
3. **then** $w \leftarrow \text{right}[p[x]]$
4. **if** $\text{color}[w] = \text{RED}$ // Case 1
5. **then** $\text{color}[w] \leftarrow \text{BLACK}$ // Case 1
6. $\text{color}[p[x]] \leftarrow \text{RED}$ // Case 1
7. LEFT-ROTATE($T, p[x]$) // Case 1
8. $w \leftarrow \text{right}[p[x]]$ // Case 1

not necessary

Case 1:



left rotation
→




RB-Delete-Fixup(T, x) (Contd.)


/ x is still $\text{left}[p[x]]$ */*

9. **if** $\text{color}[\text{left}[w]] = \text{BLACK}$ and $\text{color}[\text{right}[w]] = \text{BLACK}$
10. **then** $\text{color}[w] \leftarrow \text{RED}$ *// Case 2*
11. $x \leftarrow p[x]$ *// Case 2*
12. **else if** $\text{color}[\text{right}[w]] = \text{BLACK}$ *// Case 3*
13. **then** $\text{color}[\text{left}[w]] \leftarrow \text{BLACK}$ *// Case 3*
14. $\text{color}[w] \leftarrow \text{RED}$ *// Case 3*
15. $\text{RIGHT-ROTATE}(T, w)$ *// Case 3*
16. $w \leftarrow \text{right}[p[x]]$ *// Case 3*

Case 2:



Case 3:




RB-Delete-Fixup(T, x) (Contd.)

/ x is still $left[p[x]]$ */*

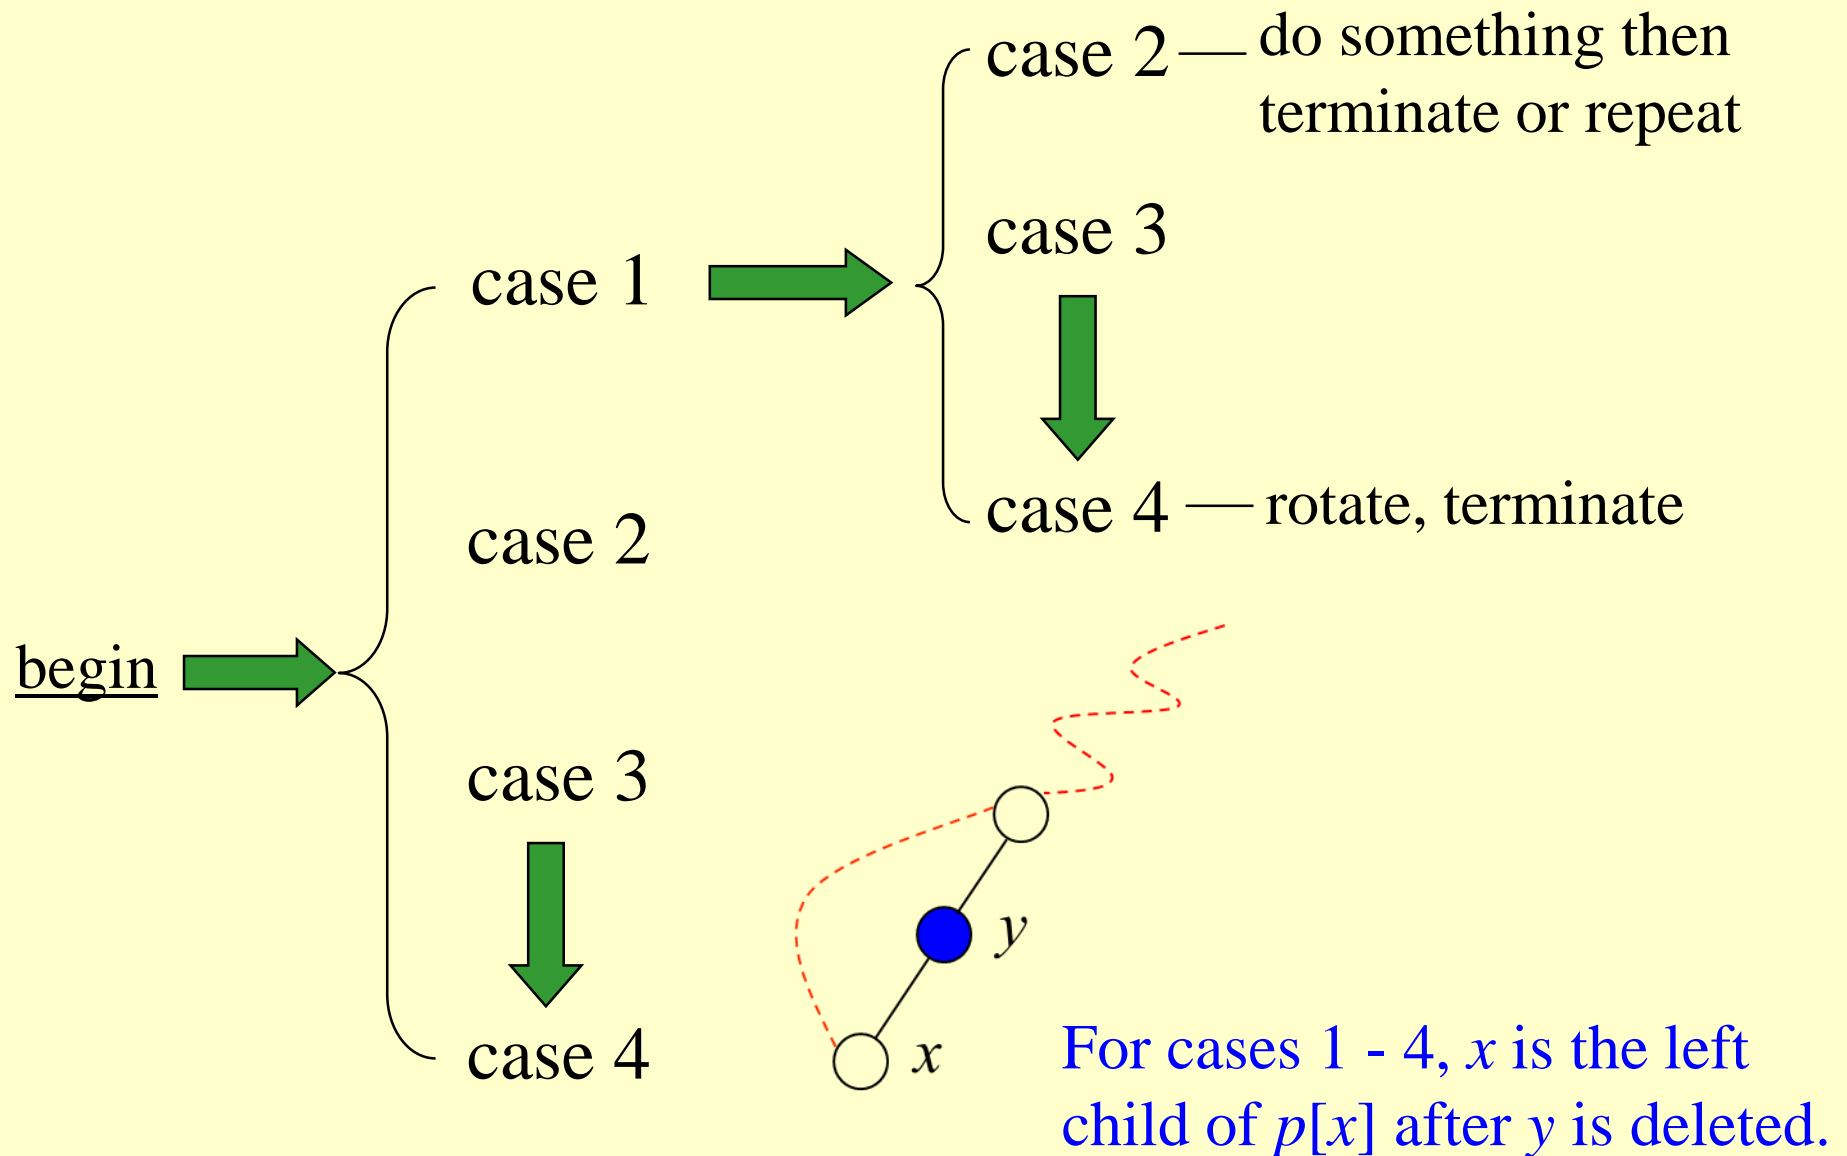
17. $color[w] \leftarrow color[p[x]]$ // Case 4
 18. $color[p[x]] \leftarrow \text{BLACK}$ // Case 4
 19. $color[right[w]] \leftarrow \text{BLACK}$ // Case 4
 20. $\text{LEFT-ROTATE}(T, p[x])$ // Case 4
 21. $x \leftarrow root[T]$ // Case 4
 22. **else** (for cases 5 – 8, same as lines 3 - 21 with “right” and “left”
exchanged)
 23. $color[x] \leftarrow \text{BLACK}$
- to go out the while-loop


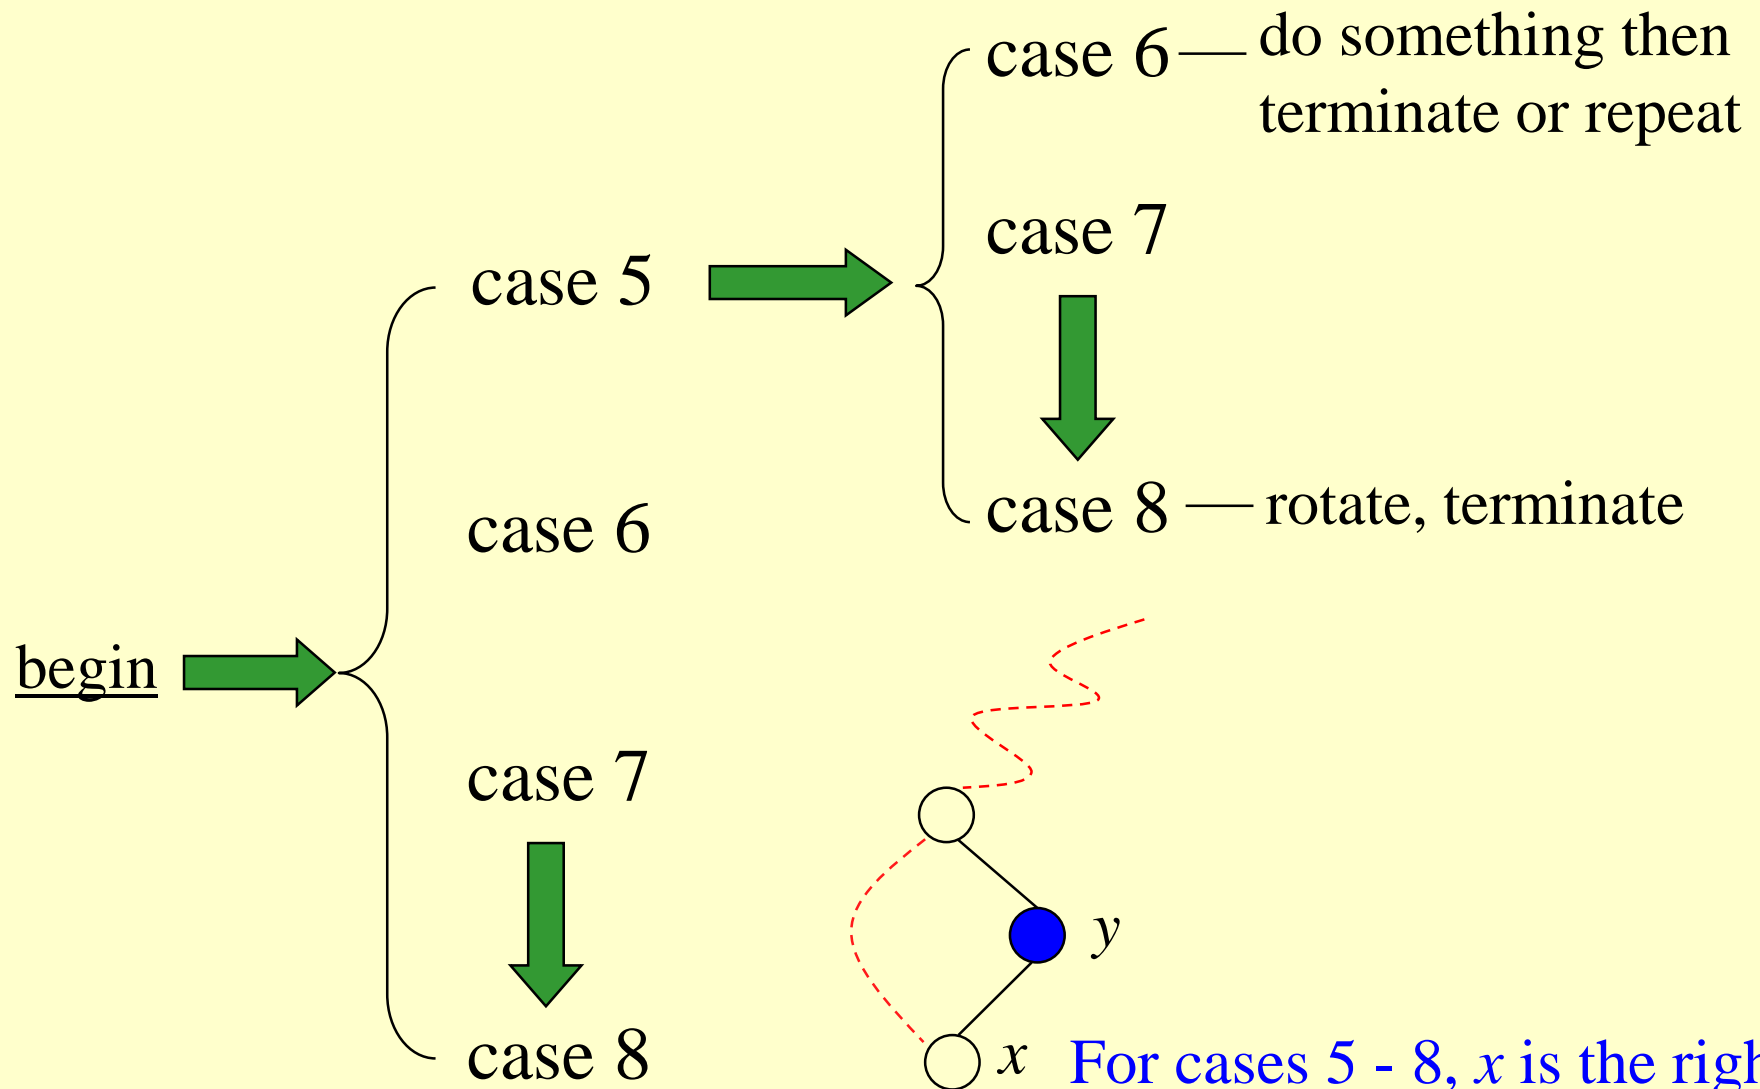
Case 4:



Deletion – Fixup

- ◆ **Idea:** Move the extra black (represented by x) up the tree until
- ◆ x points to a red node (this node is considered to be a red & black node since “ x points to” means an extra black) \Rightarrow turn it into a black node,
- ◆ x points to the root \Rightarrow just remove the extra black, or
- ◆ We can do certain rotations and recoloring and finish.
- ◆ 8 cases in all, 4 of which are symmetric to the other. (4 cases for the situation that x is the left child of $p[x]$; 4 cases for the situation that x is the right child of $p[x]$.)
- ◆ Within the **while** loop:
 - » x always points to a nonroot doubly black node.
 - » w is x 's sibling.
 - » w cannot be $nil[T]$. Otherwise, it would violate property 5 at $p[x]$.

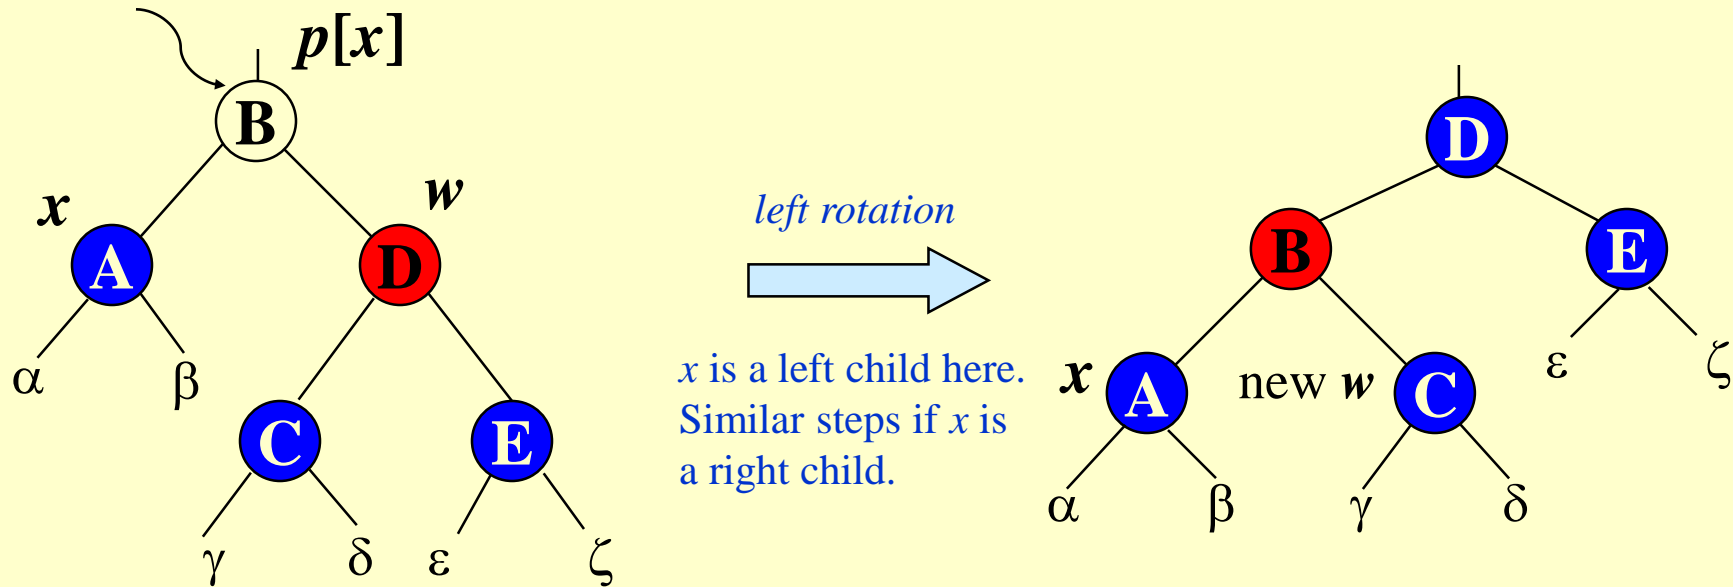




For cases 5 - 8, x is the right child of $p[x]$ after y is deleted.


Case 1 – w is red

B must be black.




- ◆ w must have black children.
- ◆ Make w black and $p[x]$ red (because w is red $p[x]$ cannot be red).
- ◆ Then left rotate on $p[x]$.
- ◆ New sibling of x was a child of w before rotation \Rightarrow it must be black.
- ◆ Go immediately to case 2, case 3, or case 4.

Case 2 – w is black, both w 's children are black




- ◆ Take 1 black off x (\Rightarrow singly black) and 1 black off w (\Rightarrow red).
- ◆ Move that black to $p[x]$.
- ◆ Do the next iteration with $p[x]$ as the new x .
- ◆ If entered this case from case 1, then $p[x]$ was red \Rightarrow new x is red & black \Rightarrow color attribute of new x is RED \Rightarrow loop terminates. Then new x is made black in the last line of the algorithm.

Case 3 – w is black, w 's left child is red, w 's right child is black



- ◆ Make w red and w 's left child black.
- ◆ Then right rotate on w .
- ◆ New sibling w of x is black with a red right child \Rightarrow case 4.

Case 4 – w is black, w 's right child is red



- ◆ Make w be $p[x]$'s color (c).
- ◆ Make $p[x]$ black and w 's right child black.
- ◆ Then left rotate on $p[x]$.
- ◆ Remove the extra black on x ($\Rightarrow x$ is now singly black) without violating any red-black properties.
- ◆ All done. Setting x to root (see line 21 in the algorithm) causes the loop to terminate.

Analysis

- ◆ $O(\lg n)$ time to get through RB-Delete up to the call of RB-Delete-Fixup.
- ◆ Within RB-Delete-Fixup:
 - » Case 2 is the only case in which more iterations occur.
 - x moves up 1 level.
 - Hence, $O(\lg n)$ iterations.
 - » Each of cases 1, 3, and 4 has 1 rotation $\Rightarrow \leq 3$ rotations in all.
 - » Hence, $O(\lg n)$ time.

Hysteresis : or the value of lazyness

- ◆ The red nodes give us some slack – we don't have to keep the tree perfectly balanced.
- ◆ The colors make the analysis and code much easier than some other types of balanced trees.
- ◆ Still, these aren't free – balancing costs some time on insertion and deletion.