

Week 9: Balanced Trees

In the worst case, all BST operations runs in $O(N)$ time. This case happens when the tree has a linear structure. We want to maintain additional properties on a BST so that it is balanced. Perfect balance is hard to achieve, so we usually relax the balance properties. As long as we can make sure that the height of the BST is always $O(\log N)$, we are happy!

Readings

- Required
 - [Weiss] ch19.4 and ch19.5 – 19.5.1
 - <http://www.ddj.com/print/documentID=14585>
- Fun
 - <http://www.seanet.com/users/arsen/avl/tree.html>
 - <http://www.cs.washington.edu/homes/sds/rb.html>
(don't take this too seriously)
- Exercise
 - [Weiss] 19.5

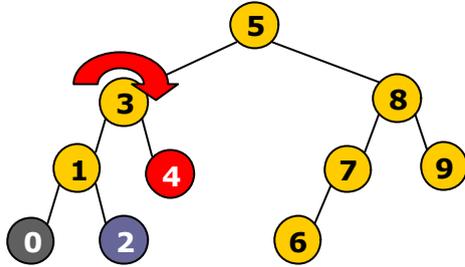
nus.soc.cs1102b.week9

2

Rotate Operation

The rotate operation is an important operation for maintaining the balance of a BST.

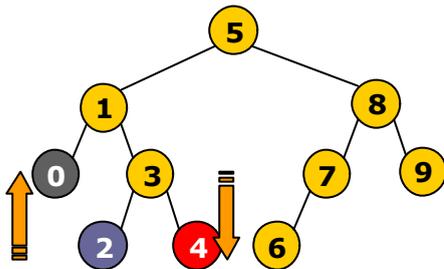
Rotate Right at 3



nus.soc.cs1102b.week9

7

After Rotate Right at 3

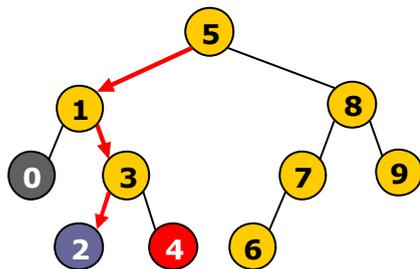


nus.soc.cs1102b.week9

8

Rotation changes the heights of nodes. In this example, the depth of node 4 increases by 1. The depth of node 0 decreases by 1. The depth of node 2 remains unchanged.

After Rotate Right at 3



nus.soc.cs1102b.week9

9

Right rotation modifies the following pointers.

Rotate Right at x

```

rotateRight(x)
l = x.left
if l is empty
  return
x.left = l.right
l.right = x
p = x.parent
if x is a left child
  p.left = l
else
  p.right = l
  
```

nus.soc.cs1102b.week9 10

The pseudocode on the right shows how we rotate right at x. The red arrows are the pointers after modification.

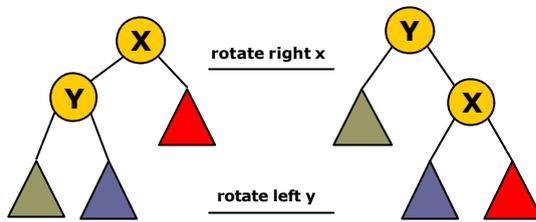
Rotate Left at 5

nus.soc.cs1102b.week9 11

After Rotate Left at 5

nus.soc.cs1102b.week9 12

Rotation Summary



nus.soc.cs1102b.week9

13

AVL Tree

15

AVL Tree Properties

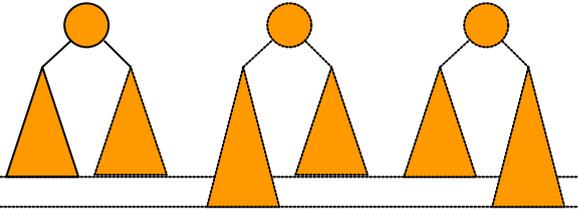
- Difference in height between left and right subtree is at most one.

$$|H_l - H_r| \leq 1$$

nus.soc.cs1102b.week9

16

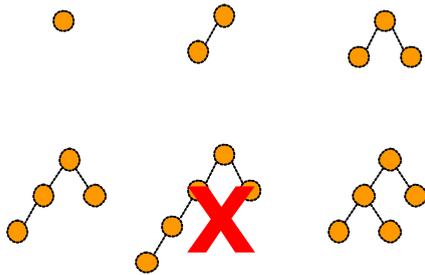
AVL Tree Properties



nus.soc.cs1102b.week9

17

AVL Tree Example



nus.soc.cs1102b.week9

18

Height of an AVL Tree

- Minimal AVL trees: AVL Tree with fewest possible number of nodes
- Minimal AVL trees with height 0



- Minimal AVL trees with height 1



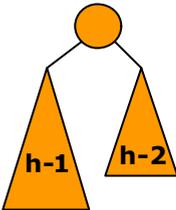
nus.soc.cs1102b.week9

19

The last line comes from the fact that $n(h-1) > n(h-2)$.

Height of an AVL Tree

- $n(h)$: number of nodes in a minimal AVL tree with height h .
- $n(0) = 1$
- $n(1) = 2$
- $n(h) = 1 + n(h-1) + n(h-2)$
 $> 2n(h-2)$



nus.soc.cs1102b.week9 20

Height of an AVL Tree

$$\begin{aligned}
 n(h) &> 2n(h-2) \\
 &> 2 * 2n(h-4) \\
 &> 4n(h-4) \\
 &> 4 * 2n(h-6) \\
 &> 8n(h-6) \\
 &\vdots \\
 &> 2^i n(h-2i)
 \end{aligned}$$

nus.soc.cs1102b.week9 21

Verify that when h is odd, the height is also $O(\log N)$

Height of an AVL Tree

$$n(h) > 2^i n(h-2i)$$

when $h - 2i = 0, i = h/2$

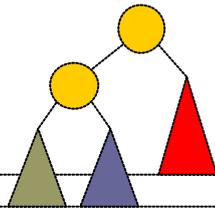
$$\begin{aligned}
 n(h) &> 2^{h/2} \\
 h &< 2 \log n(h) \\
 h &< 2 \log N \\
 h &= O(\log N)
 \end{aligned}$$

nus.soc.cs1102b.week9 22

AVL Tree Insertion

23

Idea



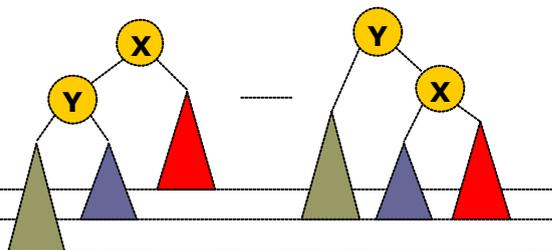
nus.soc.cs1102b.week9

24

We can insert into the red subtree, blue subtree, or gray subtree. Insertion in red subtree never violates the AVL tree property. But insertion into blue and gray subtree *may* cause a violation.

To insert into an AVL tree, we insert the node as usual. After insertion, travel from new node back to the root. At each node, checks if $|H_l - H_r| = 1$. If violation occurs, rotate the tree based on the following cases.

Case 1: Outside

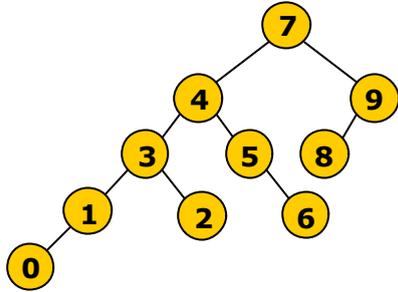


rotate right around X

nus.soc.cs1102b.week9

25

Example: Insert Outside

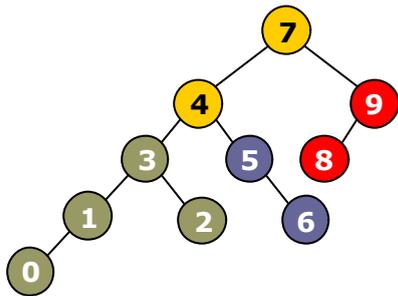


nus.soc.cs1102b.week9

26

0 is the new inserted node. In the first pass, we move down the tree just like insertion into a BST.

Example: Insert Outside

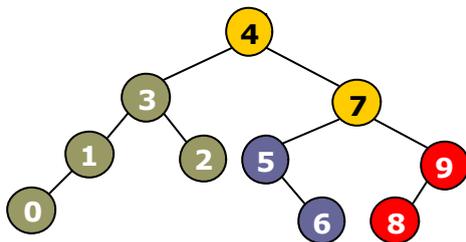


nus.soc.cs1102b.week9

27

Then, we climb our way back up. On our way towards the root, we check if the current subtree violates the AVL Tree properties. In this example, the AVL Tree property is violated at the root 7.

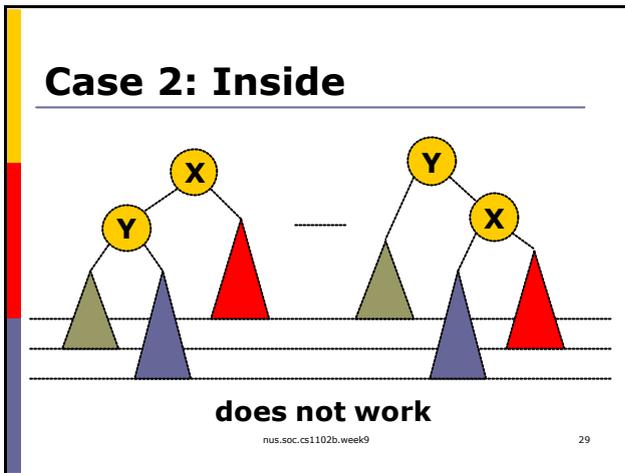
Example: Insert Outside



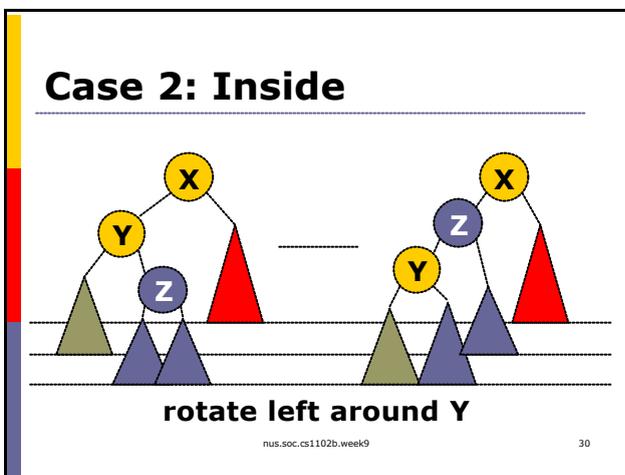
nus.soc.cs1102b.week9

28

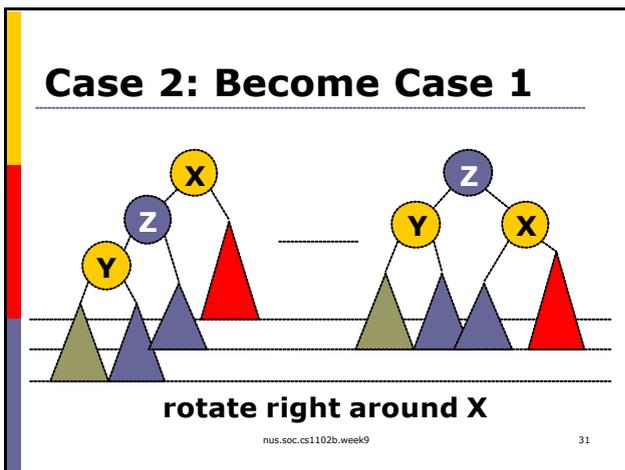
We therefore perform a single right rotation at 7 to fix the tree.



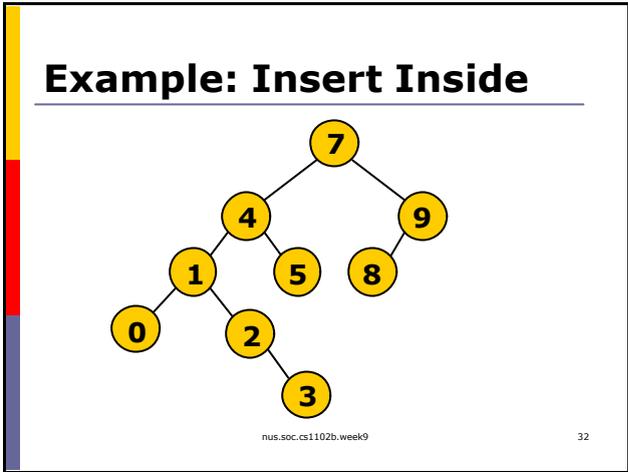
Single rotation does not work if the new node that causes violation belongs in the blue sub-tree. We need double rotations.



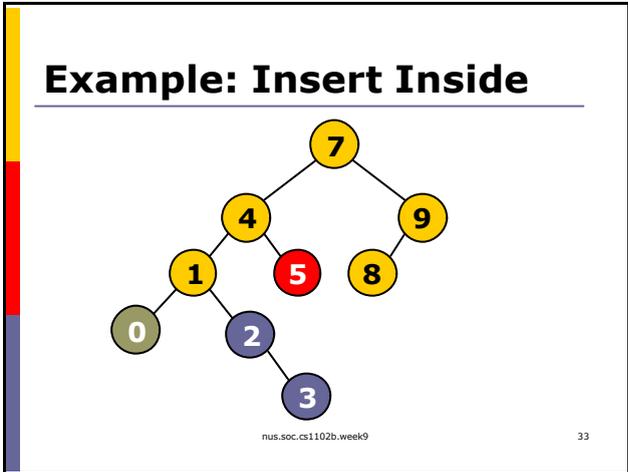
We first rotate left around Y. We have reduced our problem to the previous one: Case 1.



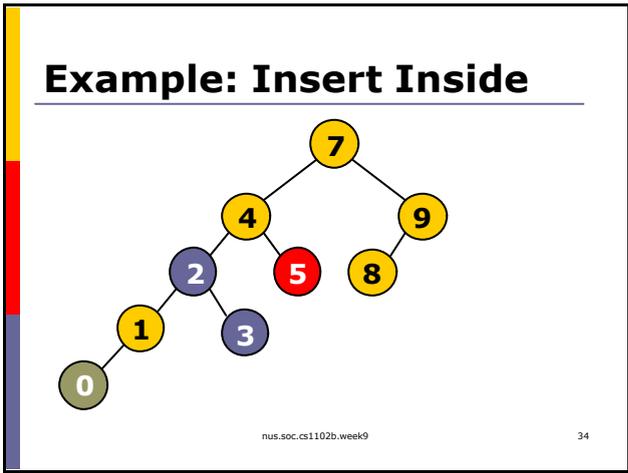
Now, we do a rotation around X to fix the problem.



In this example, we insert 3.

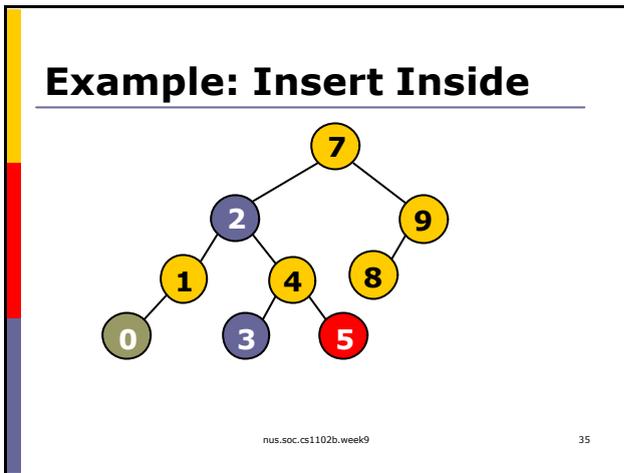


The AVL Tree property is violated at 4. We do a double rotation.



First, we rotate left at 1. We get the tree shown in the left.

Then, we rotate right around 4.

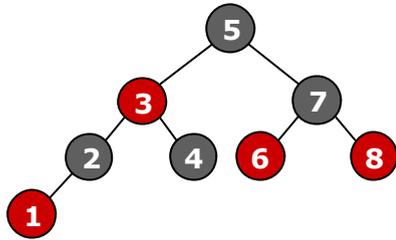


- ### Summary
- Insert outside: Single Rotation
 - Insert inside: Double Rotation
 - Two passes needed: first pass down to insert, second pass up to fix.
- nus.soc.cs1102b.week9 36

Red-Black Tree

38

A Red-Black Tree



nus.soc.cs1102b.week9

39

A Red-Black Trees has 4 properties.

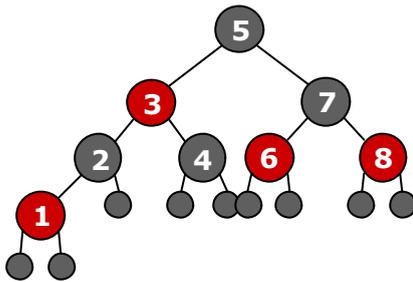
(i) All nodes must be colored red or black.

(ii) The root must be black

(iii) If a parent is red, its children must be black

(iv) All paths from the root to the null pointer must contain the same number of black nodes.

Dummy Leaves



nus.soc.cs1102b.week9

41

We add dummy black leaves to simplify implementation of red-black trees.

Properties of RB Tree

- B: Number of black nodes along a path
- N: Number of nodes

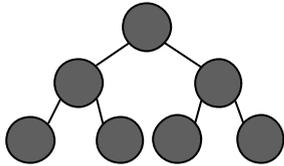
$$N \geq 2^B - 1$$
$$B \leq \log_2(N+1)$$

nus.soc.cs1102b.week9

42

To get the first inequality, we consider the smallest possible red-black tree with B black nodes along the path from root to null pointer. Such a red-black tree must contain only black nodes, and is a complete binary tree.

Complete Binary Tree



$$= 2^B - 1$$

nus.soc.cs1102b.week9

43

Recall how we compute the number of nodes in a complete binary tree (See Week 8: Trees). The number of nodes is $2^B - 1$. Note that we do not consider the dummy leaves here.

Height of an RB Tree

□ h: Height of an RB Tree

$$h \leq 2B$$

$$h \leq 2 \log_2 (N+1)$$

$$h = O(\log N)$$

nus.soc.cs1102b.week9

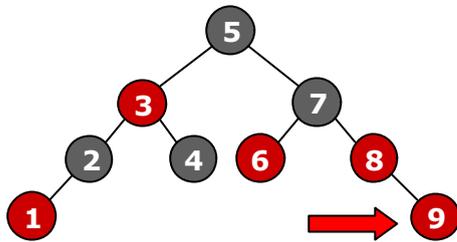
44

For a general red-black tree, since we cannot have two consecutive red nodes, the maximum height must be $2B$. From $B = \log(N+1)$, we get $h = 2 \log(N+1)$.

Red-Black Tree Insertion

45

Insertion into a RBT



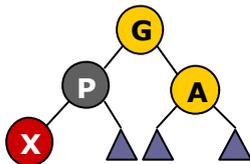
nus.soc.cs1102b.week9

46

We always color the new node red, because if the new node is black, we violate rule number 4.

We insert as usual, then travel up the tree, trying to fix the tree if violation of red-black tree rules occurs.

Case 0: No Problemo

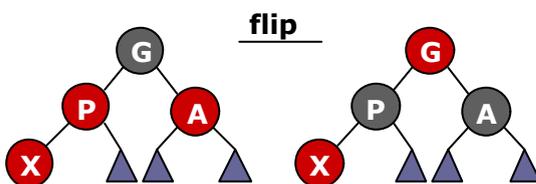


nus.soc.cs1102b.week9

47

If the parent of X is black, we are done. Note that in this case, we do not care if X is left or right child of its parent. Nor do we care about the colors of the grandparent or uncle.

Case 1: Color Flip



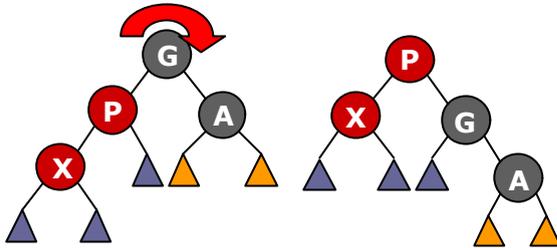
nus.soc.cs1102b.week9

48

If the parent is red, we look at the color of its uncle. If the uncle is red, we just do a color flip. Note that in this case, we do not care if X is left or right child of its parent either. If the parent of G is black, then we are done. Otherwise, parent of G is also red, and we have violated rule 3. We must continue to fix the tree.

Note that even though I draw X as a leaf, this case may happen in general when X is an internal node.

Case 2: Rotate

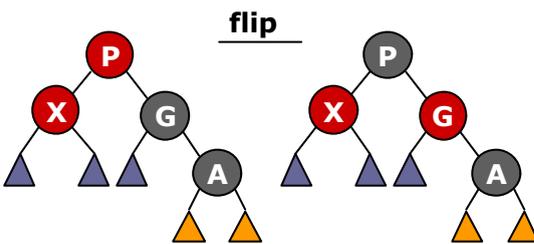


nus.soc.cs1102b.week9

49

If X's uncle is black, and X is a left child, we do a single right rotation around its grandparent, and then do a color flip. After this, node that the parent is black, so we do not need to fix the tree any further (unlike case 1).

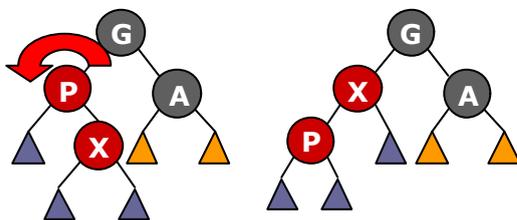
Case 2: Color Flip



nus.soc.cs1102b.week9

50

Case 3: Rotate 1



nus.soc.cs1102b.week9

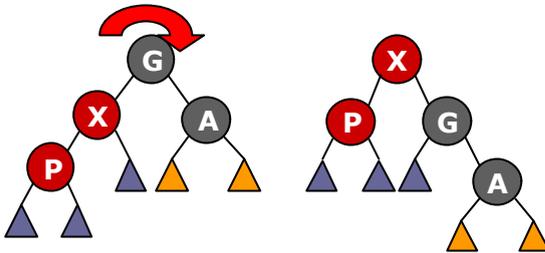
51

If X's uncle is black, and X is a right child, we need two rotations and a color flip.

Just like case 2, after two rotations and one color flip, we do not need to fix the tree any further.

Note that after the first rotation, we have reduce the problem to the previous one (case 2).

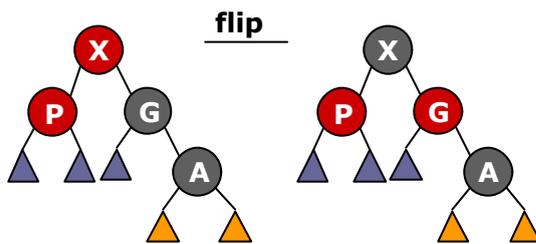
Case 3: Rotate 2



nus.soc.cs1102b.week9

52

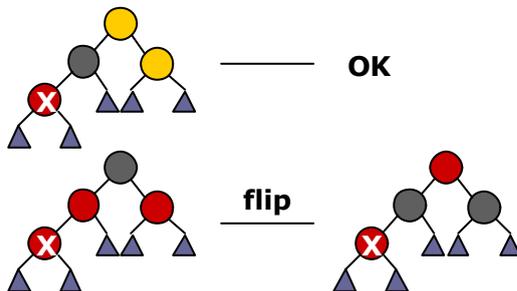
Case 3: Color Flip



nus.soc.cs1102b.week9

53

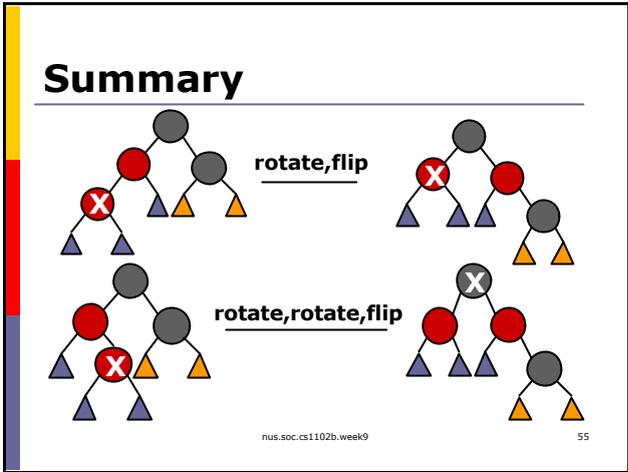
Summary



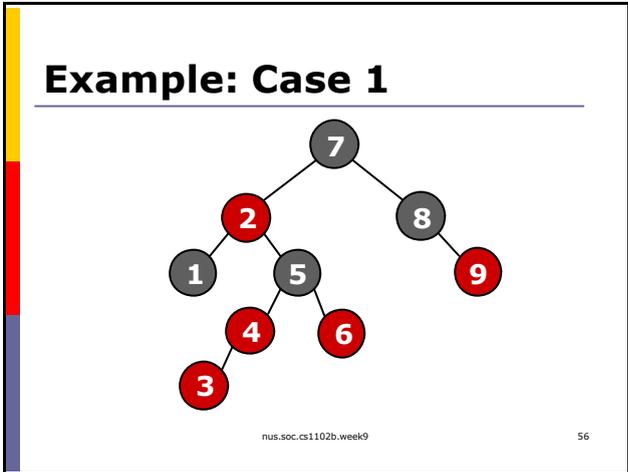
nus.soc.cs1102b.week9

54

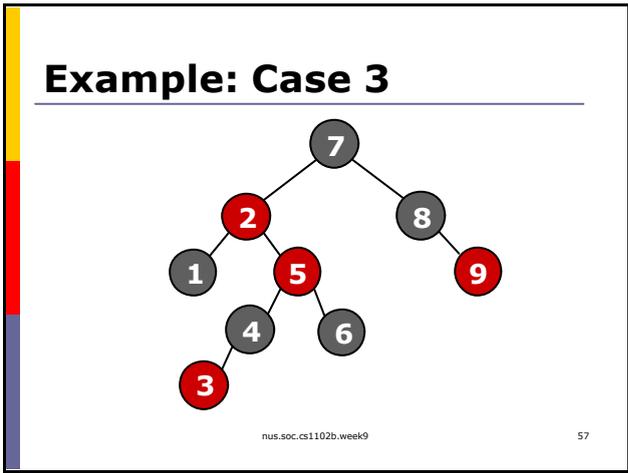
In summary, we insert, mess up the tree, and then use a bunch of rotations and color flips to correct the tree.



In this example, we insert 3. We color 3 as red.

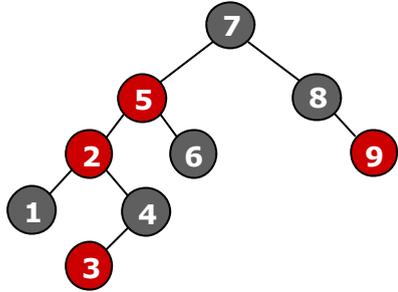


We flip the color. But we still have two consecutive red nodes in 2 and 5.



We move up the tree to fix it again. Note the position of 2,5,7 and 8. This is case 3. Do a double rotation and then a color flip.

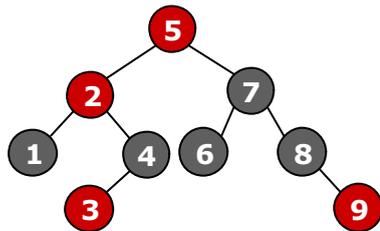
Example: Rotate



nus.soc.cs1102b.week9

58

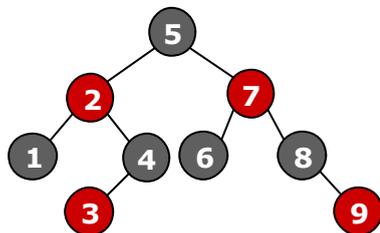
Example: Rotate



nus.soc.cs1102b.week9

59

Example: Flip

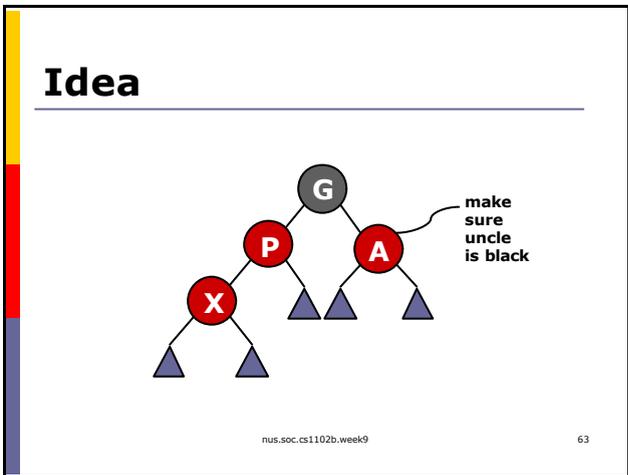


nus.soc.cs1102b.week9

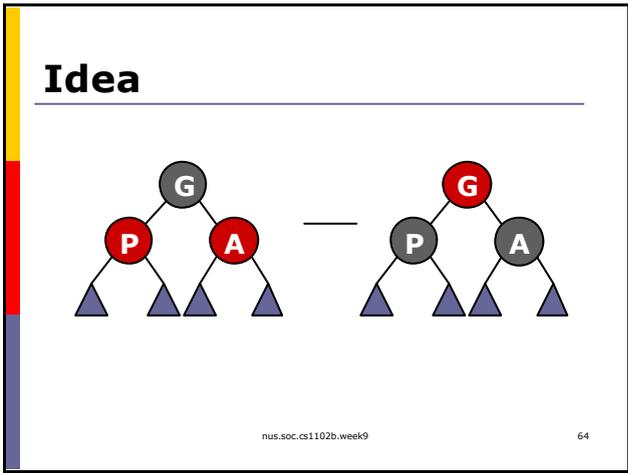
60

Top-Down Insertion

The rest of these materials are not covered in the lecture, and are optional.

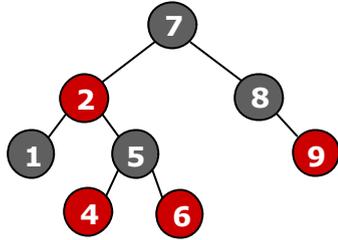


Using the insertion routine described previously, we need two passes to insert into a red-black tree: first pass down to insert, second pass up to rotate. But we only need the second pass because in case 1, we change the color of the grandparent to red. If we make sure that case 1 never happens (i.e., the uncle is always black), then we do not need to insert in two passes.



During the first pass of insert, when we traverse down the tree, everytime there is a potential that parent and uncle are both red, (that is, we see a black node with two red children), we do a color flip. Now, the grandparent is red, and this may violate rule 3 if the great-grandparent is red as well. But we can be sure that granduncle is black (WHY?)

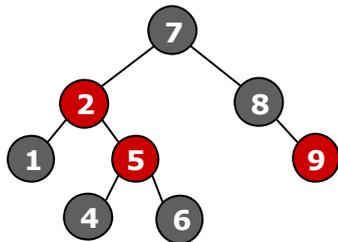
Insert 3



nus.soc.cs1102b.week9

65

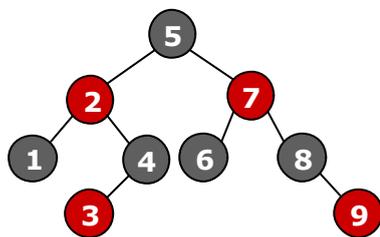
Insert 3



nus.soc.cs1102b.week9

66

Insert 3



nus.soc.cs1102b.week9

67

Top-Down Deletion

68

delete key

```
X = root
while X is not empty
  if X.data == key
    delete X
  else if X.data < key
    X = X.left
  else
    X = X.right
```

nus.soc.cs1102b.week9

69

Just like top-down insertion, we can do top-down deletion. The idea here is that we want to make sure the node to be deleted is red. (WHY?). As we traverse down the tree to find the node to be deleted, we color the current node as red. This of course will mess up the tree. Therefore, we have to fix it.

Idea: Make X Red

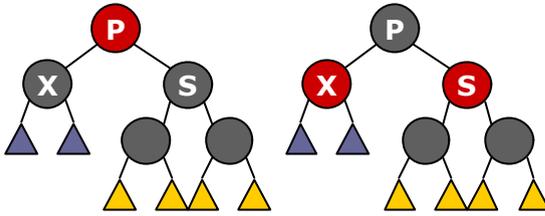
```
X = root
while X is not empty
  if X.data == key
    delete X
  else if X.data < key
    X = X.left
  else
    X = X.right
```

make X red
fix the tree

nus.soc.cs1102b.week9

70

Case 1.1

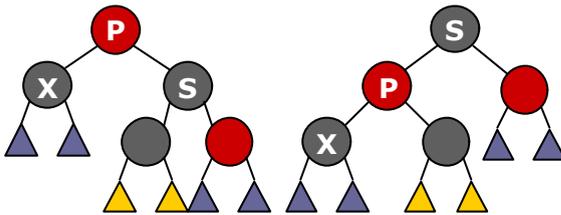


nus.soc.cs1102b.week9

71

If both children of X are black, and both children of X's sibling are black as well, we just do a color flip to color X red and keep the RB-Tree properties.

Case 1.2

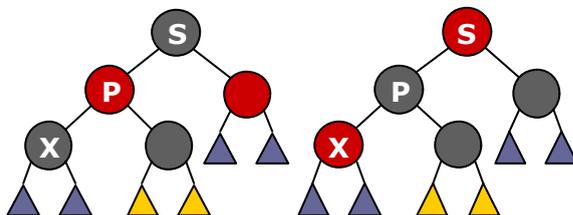


nus.soc.cs1102b.week9

72

If both children of X are black, and X's right niece (children of X's sibling) is red, we need to perform a single rotation, follow by a color flip to color X as red and maintain the RB Tree properties.

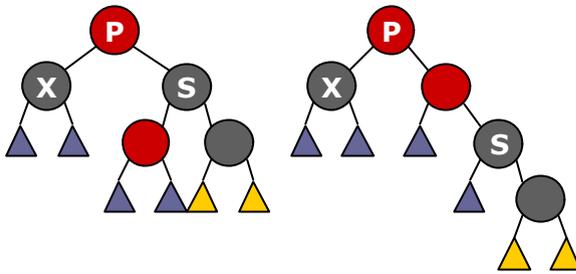
Case 1.2



nus.soc.cs1102b.week9

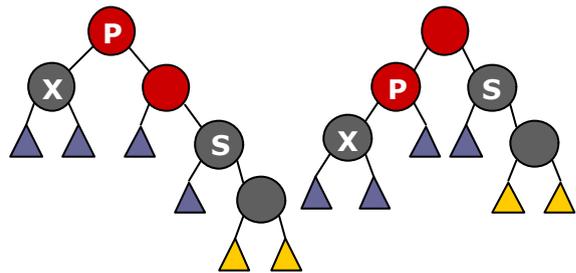
73

Case 1.3

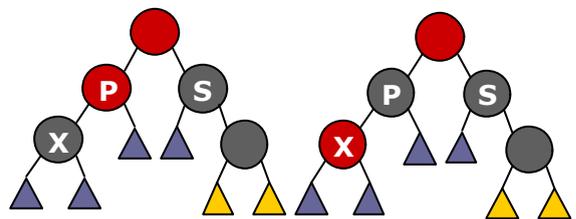


If both children of X are black, and X's left niece (children of X's sibling) is red, we need to perform a double rotation, follow by a color flip to color X as red and maintain the RB Tree properties.

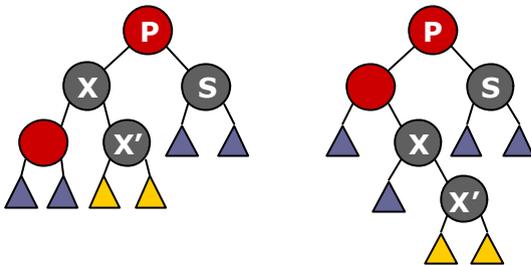
Case 1.3



Case 1.3



Case 2



Now consider the case if X has a red child. We do nothing, and hope that the next step of traversal brings up to the red node. If so, we are OK. Otherwise, we have step onto a black node X'. We do a rotation and a color flip. But X' is still black. Now we check which case (1.1,1.2,1.3 or 2) applies again to make X' red.

Case 2

