


Balanced Search Trees

CLRS ch 13, chap 18



Next: no discuss data structures that build approximately balanced binary trees

- Red black trees \neq binary
- 2-3-4 Trees, B-trees, ¹⁹⁷⁰ 23 \neq not binary
- AVL trees (1962, 1st one)
- Skip Lists, Treaps \leftarrow 1996 most recent

RED-BLACK trees - all operations in $O(\log n)$ time

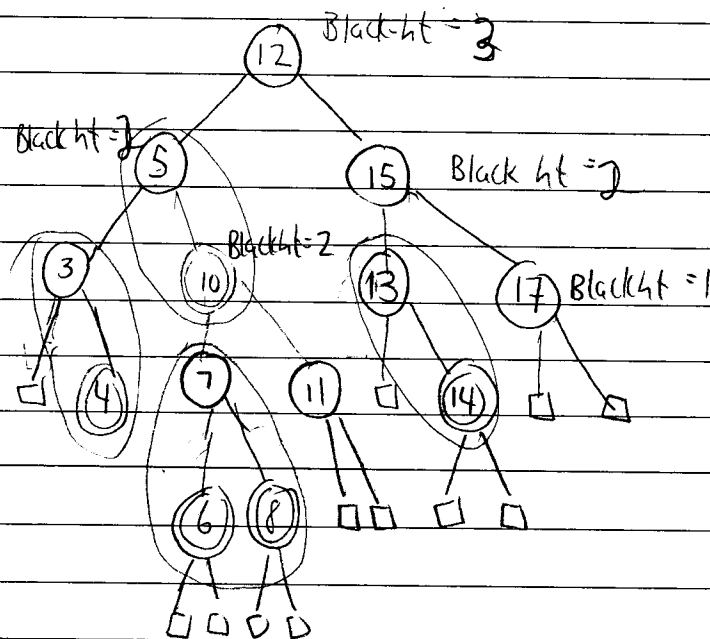
A BST is a red-black tree if it satisfies the following properties

1. Each node is red or black
2. Root is black
3. Leaves (NIL) are black
4. Parent of every red node is black (so no 2 consec. reds)
5. \forall nodes, all paths from node to leaf had same # of black nodes

Black-ht(x) = # black nodes on any path from but not including x to a leaf

All internal nodes have exactly 2 children, leaves no children

Example

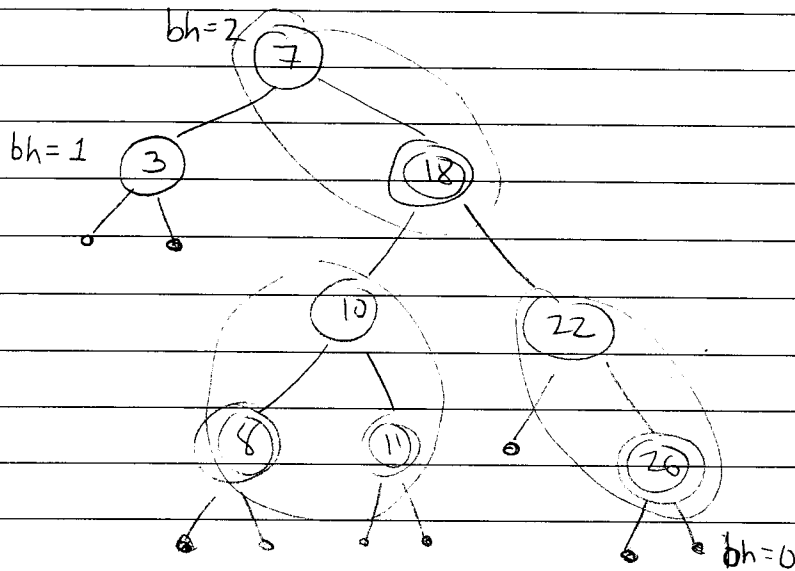


← skip this example

★ the properties force the tree to have $\log n$ ht, and make properties easy to maintain

Tricky to maintain properties as we do updates.

Ex.



Valid binary search tree

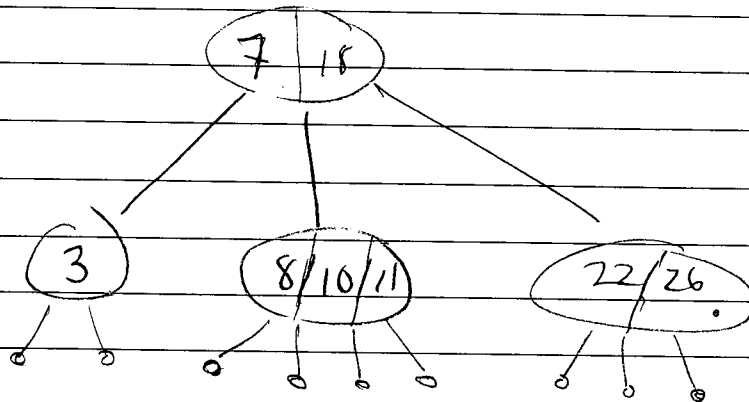
in-order traversal is sorted \odot is red

Theorem Any red-black tree with n nodes has $ht \leq 2 \ln(n+1) = O(\log n)$

PF in book by induction

PF here more intuition.

Take red nodes + merge with their parents
 in our example we get:



This is a 2-3-4 tree. Each parent has 2, 3 or 4 kids
 (Foreshadowing of what is to come)

- In a 2-3-4 tree all leaves have same depth ($= \text{bh of root}$)
- Every internal node has between 2 - 4 children

Number of leaves in a binary tree with
 n internal nodes is $n+1$

Also $n+1$ leaves in the 2-3-4 tree. (because we didn't change the
 # of leaves)

A 2-3-4 tree of height h' has between $2^{h'}$ and
 $4^{h'}$ leaves.

$$\text{so } 2^{h'} \leq n+1 \leq 4^{h'}$$

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

Now relate h (ht of Red-black tree) to h' .

Using property 4, we know $h \leq 2h'$

$$h' \leq h \leq 2h'$$

$$\text{so } h \leq 2h' \leq 2 \log(n+1).$$

Later we'll see how to manipulate 2-3-4 trees directly

Today we'll do red-black trees.

Recap: we know red-black trees always balanced

If we maintain red-black tree properties, all operations will be efficient - $O(\text{ht}) = O(\log n)$.

Ordinary insert/delete on BSTs will not satisfy red-black properties. Updates must therefore modify tree to preserve red-black properties.

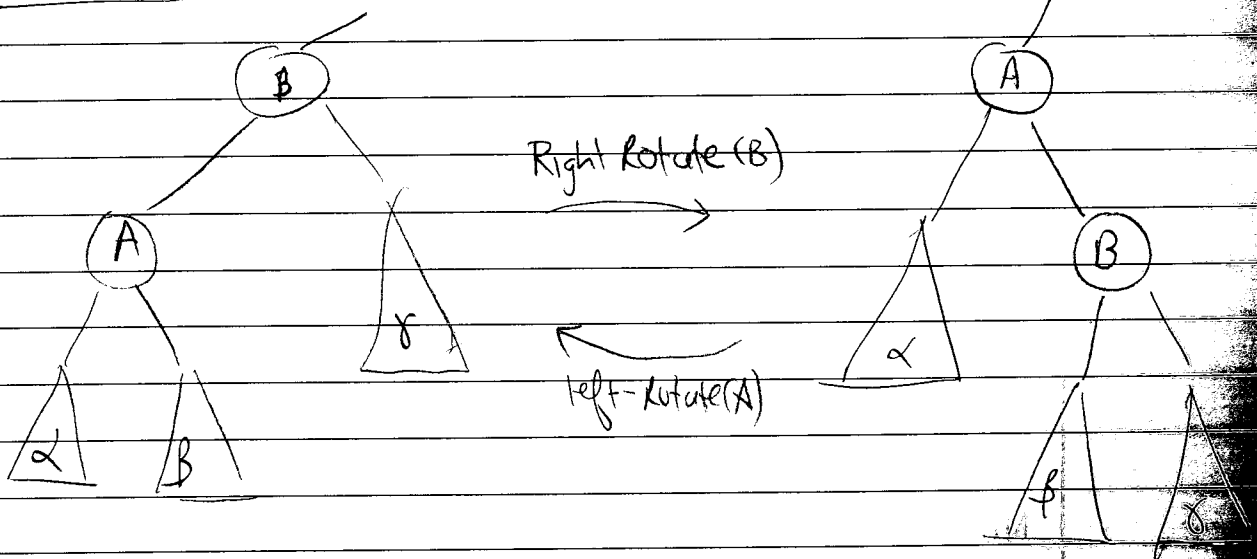
Types of operations we'll do

BST operation

Color changes

change links via rotations (this is the only way we'll change links)

Rotations



check: it preserves BST property.

i.e. \forall nodes v , nodes in left subtree all $< v$

$a \in A \leq b \in B \leq c \in X$
 in left tree $a \leq A \leq b \leq B \leq c$
 same in rt tree.

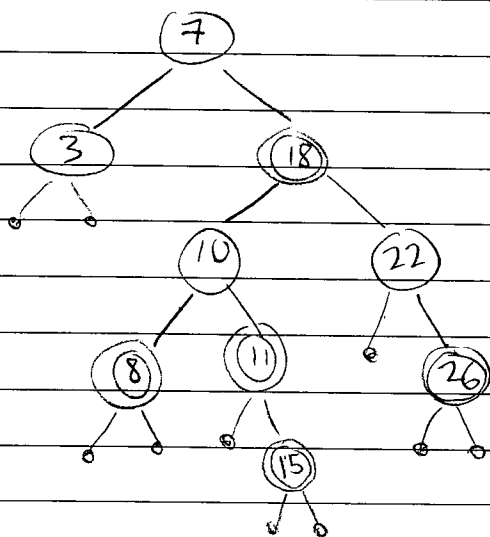
- Reversible operation,
- So Left rotate (A) v opposite
- Constant time to perform a rotation

INSERTION

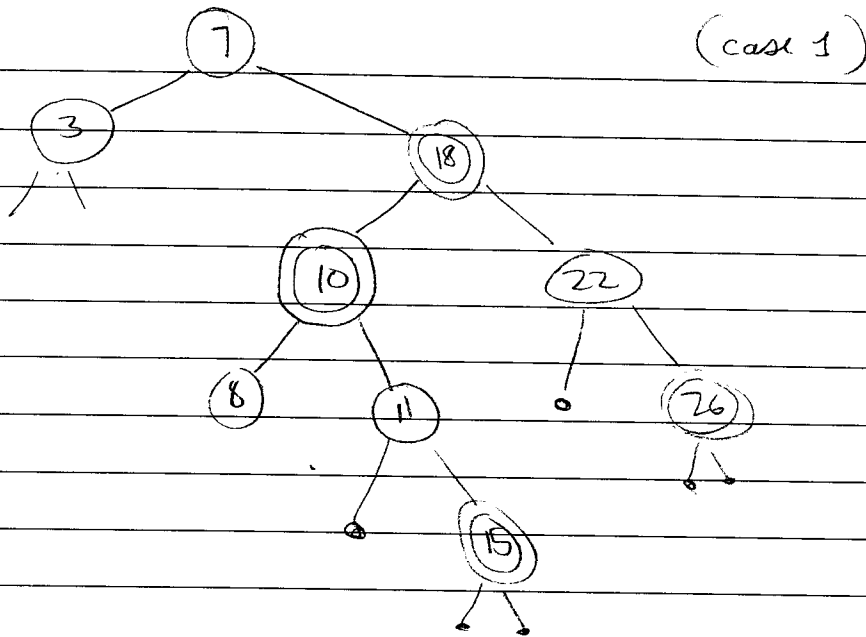
(RB-insert)

- Idea:
- Tree insert (x)
 - Make x red
 - Problem if x's parent is red (Prop. 4 may be violated)
 - But bh property still holds
 - Move violation of prop. 4 up the tree via recoloring until we can fix violation using a rotation (plus recoloring)

Example insert 15




violate prop 4.
 can't make it black since it screws up
 bh property.
 Try to recolor
 Look at grandparent of 15 -
 it is black with 2 red children
 turn it into red with 2 black
 children.



violation moved
up the tree —
between 10 + 18

Look at grandparent of 10.
Can't do same trick since grandparent has
one red, one black child
Instead do a rotation —
rotate 18 to the right

Rotate - Rt (18)



still a violation
between 10 + 18
but now it is straight
(7-10-18)
not zig-zagged

Now rotate left (7) + recoloring

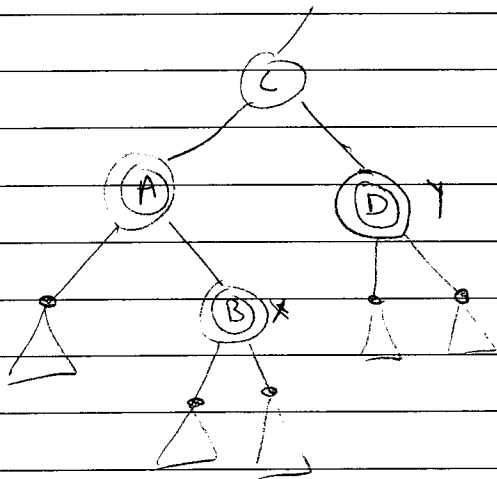
3 cases



means this subtree has a black root

all \triangle have same black ht


Case A1 (Case A.1)



key D, node y
since A, B, D all red,
 \triangle 's all have same black ht

in 2-3-4 tree this has
5 children which is bad

make A, D black c red



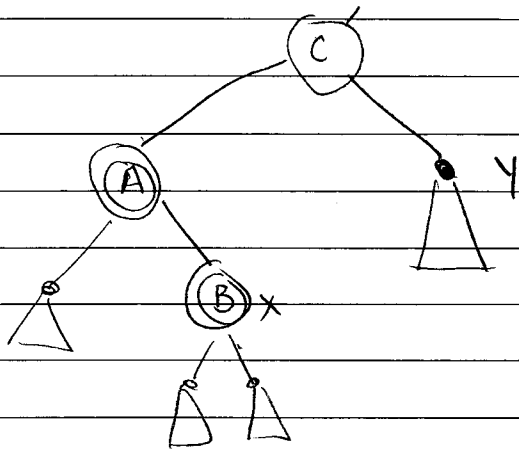
NEW $x = c$

Now c may violate prop. 4
c is higher up.

Continue with $x = c$.

Note x can be on left or rt of A -
doesn't matter.


Case A.2



x is right child of A

2-3-4,
only 4 children.


Do left-rotate (A)



b4
Zigzag between X
and grandparent
NOW zigzag

Case B3 ↑ picture is as above

Do Right-Rotate (c), then recolor



NO violations!
(check for hw)

Runtime

Cases 2 + 3 are terminal

(Case 3 - done)

(Case 2 \rightarrow immediately case 3)

Case 1 repeats only $\log n$ times

So total time is $O(\log n)$

Also - only $O(1)$ rotations. Most changes are recoloring. So very fast in practice.

Also can do other stuff (^{search} queries) during the recoloring

Rotation more expensive since nodes have to be locked during rotations

DELETION - more complicated.

SEE BOOK

END
lecture
2

Tutorial 2: Deletion in Red/Black Trees

Middle Q: 2-3 trees

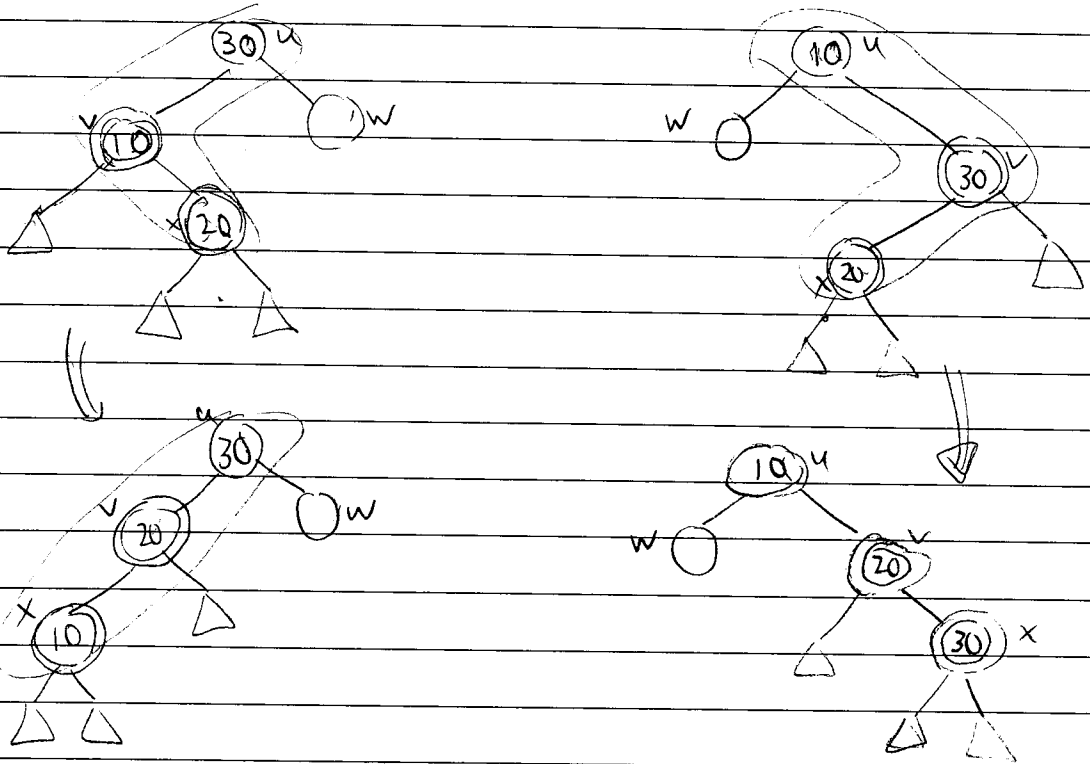
modified Red/Black

only RT red. so no recoloring?




Another new

1. grandparent of x has a black child, w . 4 cases






Ely Z.

Restructure (1 or 2 rotations) to get




2. grandparent of x has 2 red children. Recolor.






Eliminates violation or moves violation up tree ← equivalent to a split op



Example Red Black Insert (14)



corresponds to
this 234 →
red



Deletion in Red-Black Trees


Do Binary search to find item u to be removed

If u does not have an external child

find inmed successor v , swap u & v

+ then remove u (which now has an 'external child')

← same as
BSF delete



- If v was red (then r black) - done
- If v black and r red, make r black - done
- OW v black, r black. Make r double black
(creates underflow in corresp. 2-3-4 tree)




To remedy double-black, there are 3 cases

Case 1

Sibling y of r is Black + has red child z . Eliminates violation

rotation
recoloring

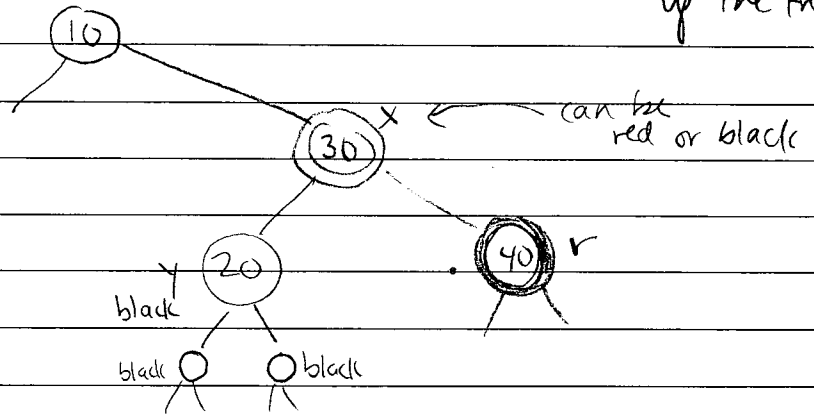
x can be
black or red
 y black
 z red
 r dbl black



Case 2 Sibling y of r is black, both children of y are black

eliminates dbl black or propagates it up the tree


recoloring



color r black


color y red

color x black if it was red, or dbl black if it was black




Case 3 Sibling y of r is red. Turns it into case 1 or case 2.

rotate
recolor



OR



now we're in case 1 or case 2