

## Balanced Search: The Problem

Search trees important?

Insertion/deletion fast (on every operation, unlike hash table, which has to expand from time to time).

Range queries, sorting (unlike hash tables)

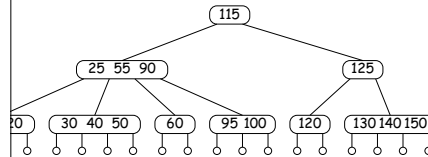
Performance from binary search tree requires remaining balanced (balanced by some constant  $> 1$  at each node).

Worst case, that tree be "bushy"

Leaf nodes (most inner nodes with one child) perform like linked list

That heights of any two subtrees of a node always differ by no more than some constant factor  $C > 0$ .

## Example of Direct Approach: B-Trees



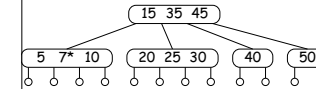
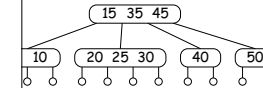
A B-tree grows/shrinks only at the root, then the two sides are roughly the same height.

Each node, except the root, has from  $\lceil M/2 \rceil$  to  $M$  children, and one child is empty between each two children.

$M$  can range from 2 to  $M$  children (in a non-empty tree).

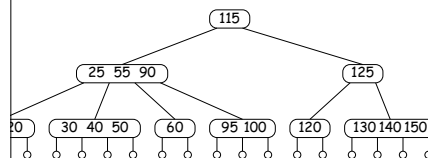
When a node is full, add to nodes just above the (empty) leaves: split overfull node into two, moving one key up to its parent.

## Inserting in (2, 4) tree (Simple Case)



## CS61B Lecture #29

## Example of Direct Approach: B-Trees



A B-tree is an  $M$ -ary search tree,  $M > 2$ .

For  $M = 4$ , non-root nodes have at least 2 nodes, so we use (2, 4) (or 2-4) trees.

Search-tree property:

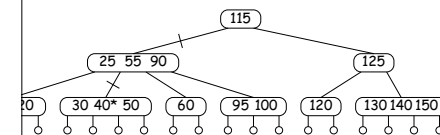
Keys are sorted in each node.

Keys in subtrees to the left of a key,  $K$ , are  $< K$ , and all to the right are  $> K$ .

Empty slots at the bottom of tree are all empty (don't really exist) and are removed from root.

A B-tree is a simple generalization of binary search.

## Example Order 4 B-tree ((2,4) Tree)



We show path when finding 40.

For each side of each child pointer in path bracket 40.

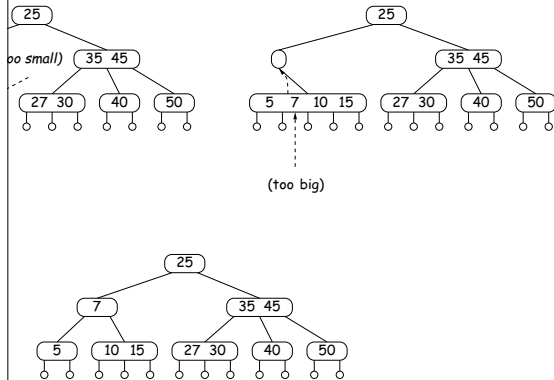
Each node has at least 2 children, and all leaves (little circles) are full, so height must be  $O(\lg N)$ .

B-trees are stored on secondary storage, and the order is much bigger

Depends on the size of a disk sector, page, or other convenient unit.

## Deleting Keys from B-tree

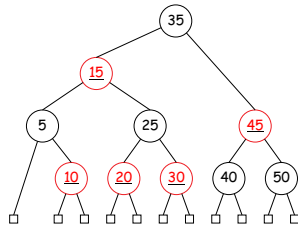
from last tree.



7:40 2021

CS61B: Lecture #29 8

## Red-Black Tree Constraints



(conceptually) colored red or black.

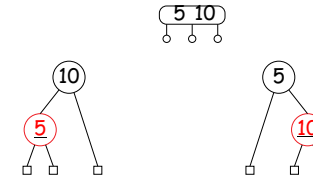
- Node contains no data (as for B-trees) and is black.
- Every leaf has same number of black ancestors.
- Every internal node has two children.
- Every red node has two black children.
- Nodes 5, 15, 25, 35, and 45 guarantee  $O(\lg N)$  searches.

7:40 2021

CS61B: Lecture #29 10

## Constraints: Left-Leaning Red-Black Trees

(2,4) or (2,3) tree with three children may be represented in different ways in a red-black tree:

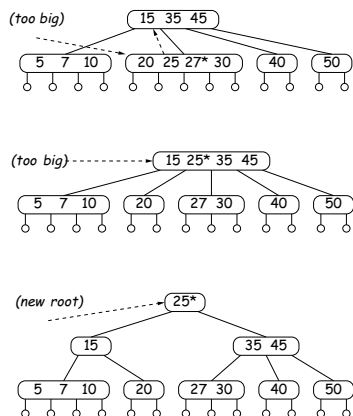


- This considerably simplifies insertion and deletion in a red-black tree by choosing the option on the left.
- Under this constraint, there is a one-to-one relationship between (2,4) trees and red-black trees.
- Left-leaning red-black trees are called *left-leaning red-black trees*.
- For simplification, let's restrict ourselves to red-black trees that correspond to (2,3) trees (whose nodes have no more than two children), so that no red-black node has two red children.

7:40 2021

CS61B: Lecture #29 12

## Inserting in B-Tree (Splitting)



7:40 2021

CS61B: Lecture #29 7

## Red-Black Trees

A *red-black tree* is a binary search tree with additional constraints that prevent it from becoming unbalanced. Searching in a red-black tree is always  $O(\lg N)$ .

Java's `TreeSet` and `TreeMap` types.

When keys are inserted or deleted, the tree is *rotated* and *recoloring* to restore balance.

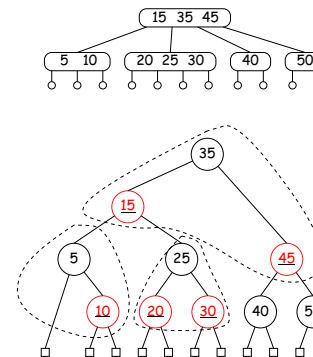
7:40 2021

CS61B: Lecture #29 9

## Red-Black Trees and (2,4) Trees

A (2,4) tree corresponds to a (2,4) tree, and the operations on (2,4) trees correspond to those on the other.

Each (2,4) tree corresponds to a cluster of 1-3 red-black trees, where the top node is black and any others are red.



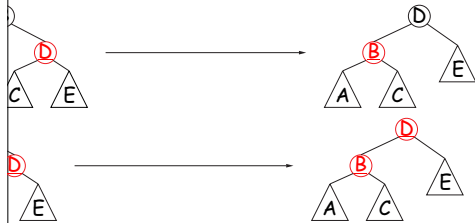
7:40 2021

CS61B: Lecture #29 11

## Rotations and Recolorings

cases, we'll augment the general rotation algorithms with coloring.

the color from the original root to the new root, and color the root red. Examples:



these changes the number of black nodes along any path from root and the leaves.

## The Algorithm (Sedgwick)

Binary-tree type RBTREE: basically ordinary BST nodes

the same as for ordinary BSTs, but we add some fixups to preserve the red-black properties.

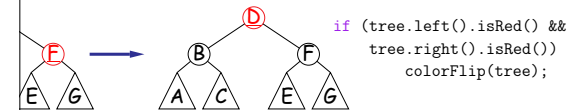
```

insert(RBTREE tree, KeyType key) {
    if (tree == null)
        return new RBTREE(key, null, null, RED);
    cmp = key.compareTo(tree.label());
    if (cmp < 0) tree.setLeft(insert(tree.left(), key));
    else tree.setRight(insert(tree.right(), key));
}
    
```

fixup(tree); // Only line that's all new!

## Fixing Up the Tree (II)

break up 4-nodes into 3-nodes or 2-nodes.



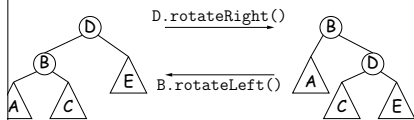
As a result of other fixups, or of insertion into the empty tree, the root may end up red, so color the root black after the rest of the fixups are finished. (Not part of the fixup function; the end).

## Red-Black Insertion and Rotations

Just as for binary tree (color red except when tree is empty).

(and recolor) to restore red-black property, and thus preserve binary tree property, but changes balance.

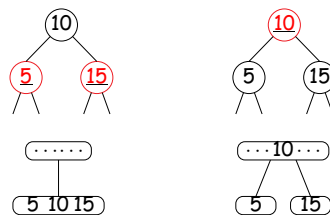
Insertion *preserves* binary tree property, but changes balance.



## Splitting by Recoloring

Insertions will temporarily create nodes with too many children, so we split them up.

Recoloring allows us to split nodes. We'll call it *colorFlip*:

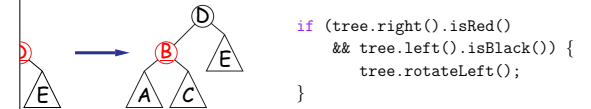


colorFlip joins the parent node, splitting the original.

## Fixing Up the Tree

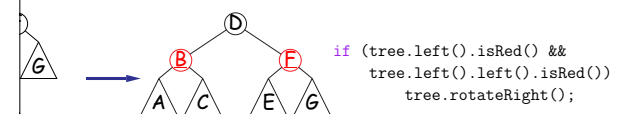
After breaking up the BST, we restore the left-leaning red-black tree and limit ourselves to red-black trees that correspond to the original by applying the following (in order) to each node:

Convert right-leaning trees to left-leaning:



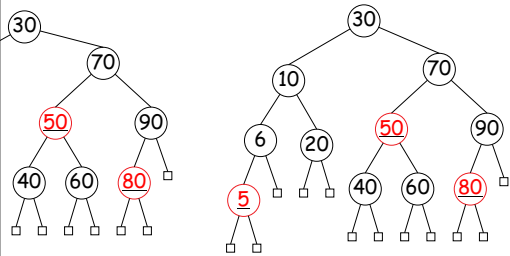
Node B will be red, so that both B and D end up red. This is a 4-node.

Convert linked red nodes into a normal 4-node (temporarily).



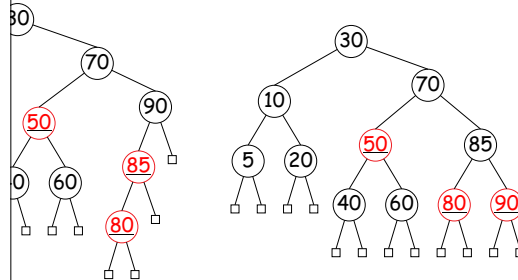
### Insertion Example (II)

..., let's insert 6, leading to the tree on the left. This is a 4-node, so apply Fixup 1:



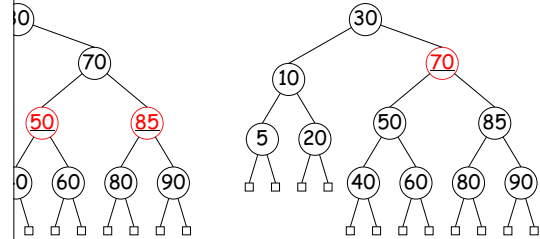
### Insertion Example (IIIa)

step 2.



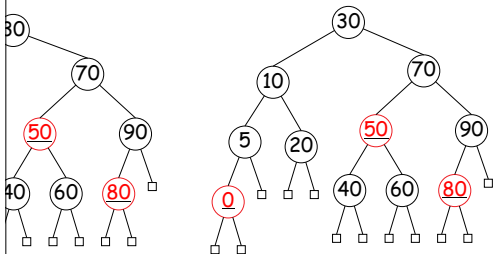
### Insertion Example (IIIc)

another 4-node, so apply fixup 3 again.



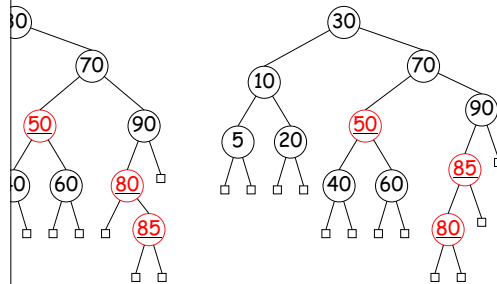
### of Left-Leaning 2-3 Red-Black Insertion

initial tree on left. No fixups needed.



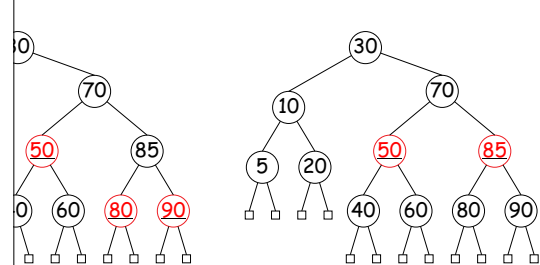
### Insertion Example (III)

or inserting 85. We need fixup 1 first.



### Insertion Example (IIIb)

is a 4-node, so apply fixup 3.



### Insertion Example (IIId)

a right-leaning tree, so apply fixup 1.

