

Augmenting Red-black trees

Suppose we want to add the following two operations

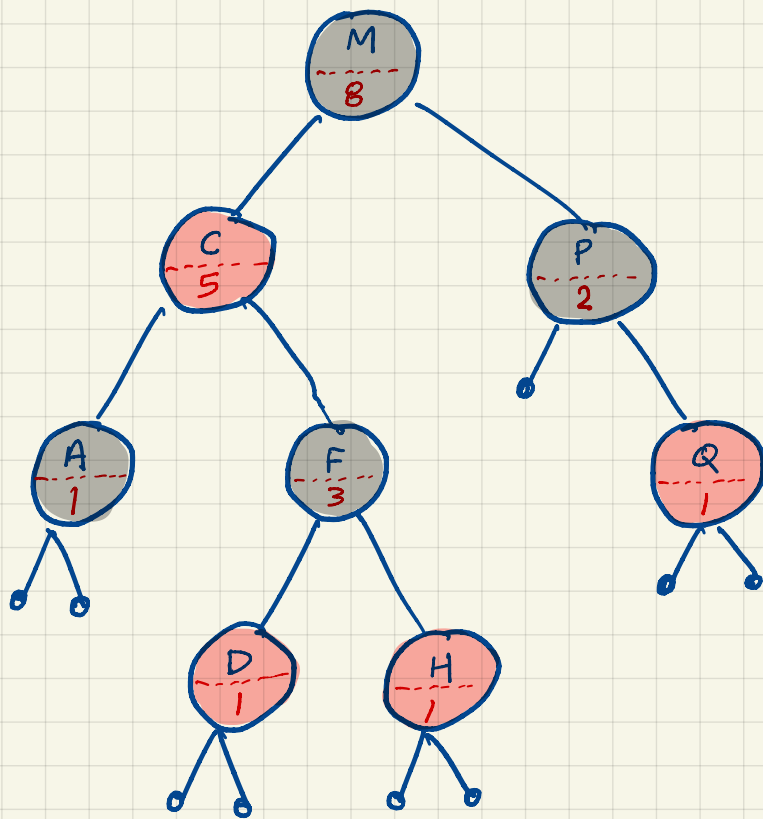
$\text{Select}(x, i)$: returns node with i th smallest key in x 's subtree

$\text{rank}(T, x)$: returns rank of x in the inorder traversal of T .

IDEA: Augment each node x with $\text{size}[x]$

$\text{size}[x] = \#$ internal nodes in x 's subtree (including x)

$\text{size}[\text{NIL}[T]] = 0$



ACDFH

Note:

$$\text{Size}[x] = \text{Size}[\text{left}[x]] + \text{Size}[\text{right}[x]] + 1$$

$\text{Size}[\text{left}[x]] + 1 = \text{rank of } x \text{ in its subtree}$

Select(x, i)

$$r \leftarrow \text{Size}[\text{left}[x]] + 1$$

if $i = r$

then return x

else if $i < r$

then return Select($\text{left}[x], i$)

else return Select($\text{right}[x], i - r$)

Rank(T, x)

$r \leftarrow \text{size}[\text{left}[x]] + 1$ \blacktriangleright rank of x in its subtree

$y \leftarrow x$

while $y \neq \text{root}[T]$

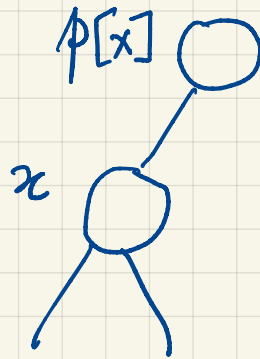
do if $y = \text{right}[p[y]]$ \blacktriangleright y is a right child

then $r \leftarrow r + \text{size}[\text{left}[p[y]]] + 1$

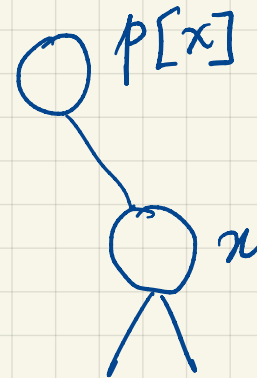
$y \leftarrow p[y]$

return r

$r = \text{size}[\text{left}[x]] + 1$

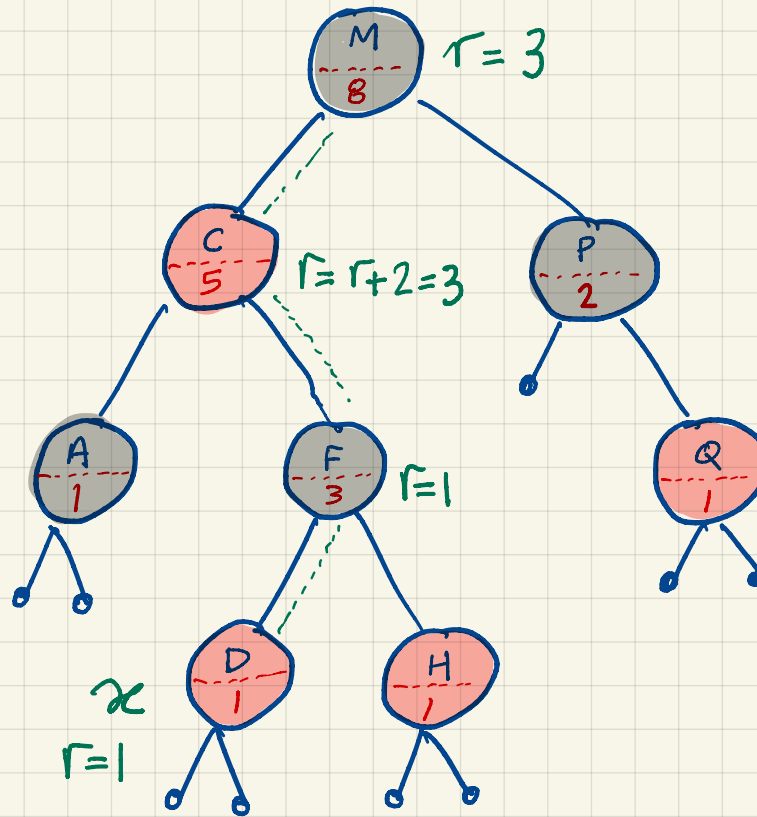


$r =$ rank of x in x 's subtree
is also rank of x in $p[x]$'s subtree



rank of x in y 's subtree
 $= r +$ rank of $p[x]$

Example:



ACDFHMPQ
↑
3

Augmenting Red-black trees in general

Let f be a field that augments a Red-black tree.

Suppose $f[x]$ can be computed using only the information in

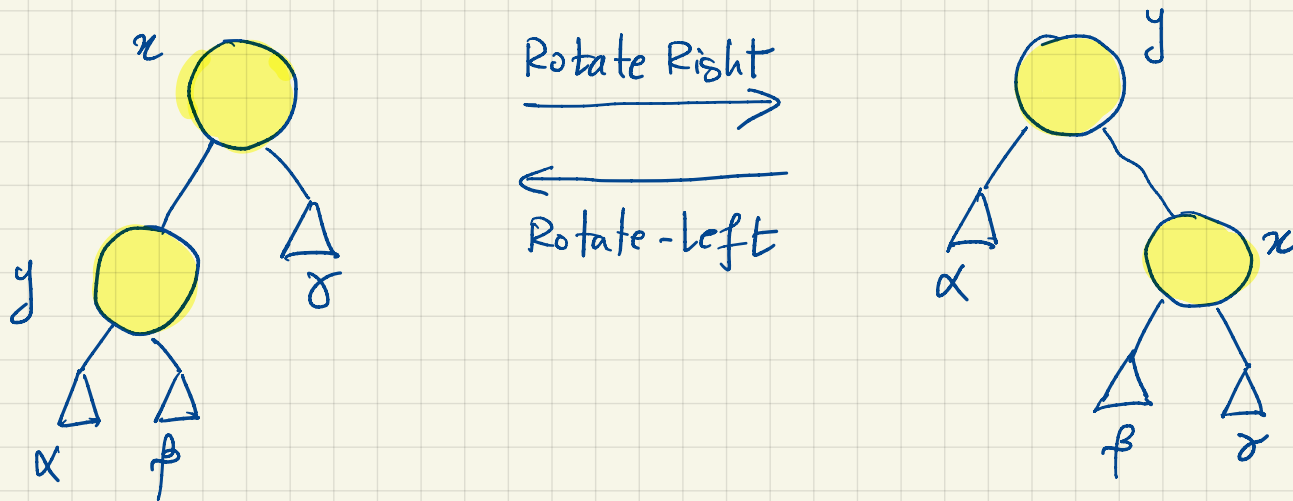
- x itself
- $\text{left}[x]$
- $\text{right}[x]$

including $f[\text{left}[x]]$ and $f[\text{right}[x]]$

Then one can maintain f in all nodes during insertion & deletion without affecting the $O(\log n)$ asymptotic performance.

Proof:

- Changing $f[x]$ only affects the field in nodes on the path from root $[T]$ to x ($O(\log n)$ computations)
- Rotations require local changes only.

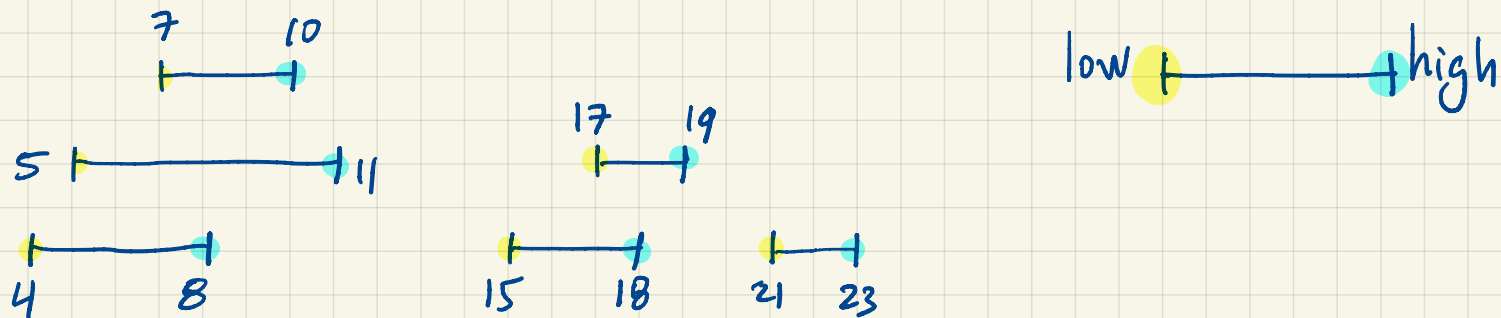


Only two nodes require changes in a rotation, and we have a constant number of rotations per operation.

Application: Interval trees

- Maintain a dynamic set of intervals
- Support operation: Find an interval in the set that overlaps a given query interval.

Example:

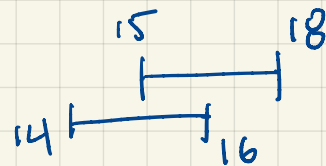


Query:

$[14, 16] \longrightarrow [15, 18]$

$[16, 19] \longrightarrow [15, 18] \text{ or } [17, 19]$

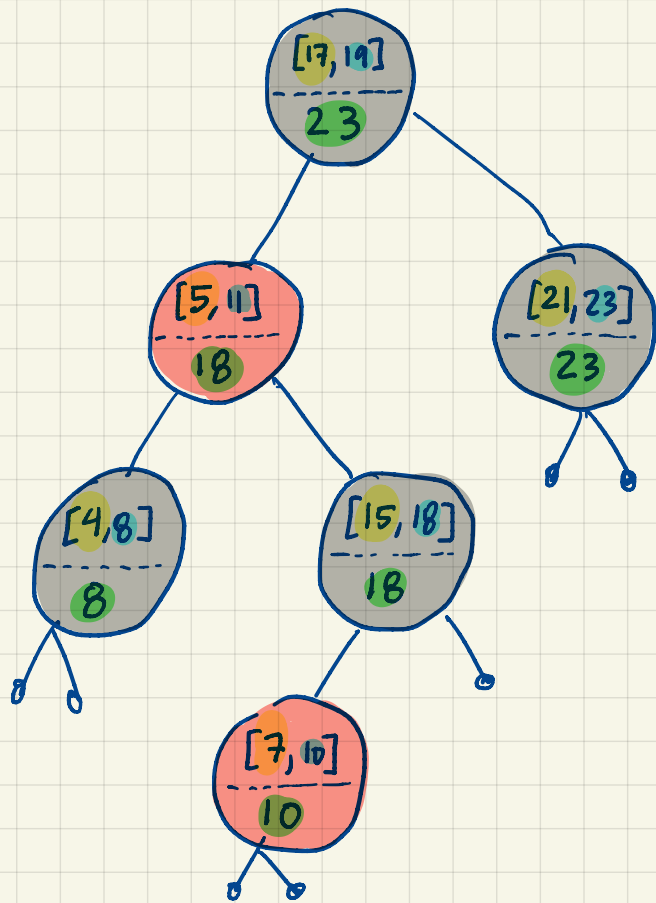
$[12, 14] \longrightarrow \text{NIL}[T]$



— Red-black tree with $\text{key}[x] = \text{low}[\text{int}[x]]$

So the binary search tree is sorted on the low end points of intervals

— Augment with $\text{max}[x] = \max \left\{ \begin{array}{l} \text{high}[\text{int}[x]] \\ \text{max}[\text{left}[x]] \\ \text{max}[\text{right}[x]] \end{array} \right. , \text{max}[\text{Nil}[T]] = 0$



BST on \bullet

\bullet is maximum of all \bullet in subtree

Interval-Search (T, i) \triangleright i is interval

$x \leftarrow \text{root}[T]$

while $x \neq \text{NIL}[T]$ and $i \cap \text{int}[x] = \emptyset$ (no overlap)

do if $\text{left}[x] \neq \text{NIL}[T]$ and $\text{max}[\text{left}[x]] \geq \text{low}[i]$

then $x \leftarrow \text{left}[x]$

else $x \leftarrow \text{right}[x]$

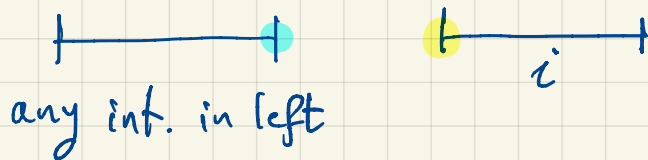
return x

Prove every move is safe

Case 1: Go right.

$\text{left}[x] = \text{NIL}[T]$ or

$\text{max}[\text{left}[x]] < \text{low}[i]$



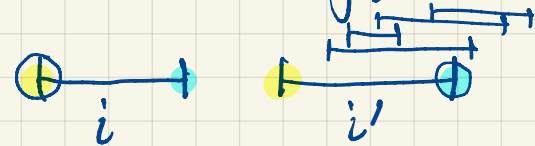
so there is no overlap in left.

Case 2: Go left.

If no overlap in left subtree

$\text{low}[i] \leq \text{max}[\text{left}[x]] = \text{high}[i']$

for some i' in left subtree



But BST \Rightarrow $\text{high}[i] <$ any
low in right subtree

So no overlap in right either.