

Predecessor search data structures

Given n integers in $[0, U-1]$, store them in a data structure, supporting

- symmetric /
- pred(x): return the largest number that is at most x .
 - succ(x): return the smallest number that is at least x .
 - ins(x): insert x to the set of integers ($x \in [0, U]$)
 - del(x): delete x .


Easy sol: sorted array + binary search (no ins/del)
or BST $O(\log n)$ time.

Today: $O(\log \log U)$ time.

Application: sort n integers faster than $O(n \log n)$

- build the data structure by inserting n integers $O(n \log \log U)$ (when U is subexponential in n)
- find the max in $O(n)$ time
- keep querying the pred of max-1, and deleting the old max. $O(n \log \log U)$

Van Emde Boas tree (vEB tree)



divide $[0, u-1]$ into \sqrt{u} blocks of size \sqrt{u}

store: min, max

$(i=0, \dots, \sqrt{u}-1)$ children $[i]$: a recursive vEB for numbers in block i over \sqrt{u}

aux: a recursive vEB tree for the \sqrt{u} blocks (the inserted "numbers" are the nonempty blocks).

(min/max do not go into the recursion)

pred(x):

if $x \geq \max$
return max

if $x < \min$
return NULL

$h = \lfloor x / \sqrt{u} \rfloor$, $l = x \bmod \sqrt{u}$.

if $l \geq \text{children}[h].\text{min}$

recursively find $p = \text{pred}(l)$ in $\text{children}[h]$

return $h: \sqrt{l} + p$

recursively find $i = \text{pred}(h)$ in aux

return $\text{children}[i].\text{max}$

Time: only 1 recursive call, reduce l to \sqrt{l} .

$O(\log \log l)$ levels of recursion.

$O(\log \log l)$ time.

$\text{ins}(x)$:

if empty or $\text{min} = \text{max}$

update min, max , return

if $x < \text{min}$, swap $x \leftrightarrow \text{min}$

// ins old min
instead

if $x > \text{max}$, swap $x \leftrightarrow \text{max}$

$h = \lfloor x / \sqrt{l} \rfloor$, $l = x \bmod \sqrt{l}$

if $\text{children}[h]$ empty

$\text{ins}(h)$ in aux

$\text{ins}(l)$ in $\text{children}[h]$

Time: never actually make two recursive calls
if "ins(h) in aux" is called,
then children[h] was empty, and
"ins(l) in children[h]" takes $O(1)$ time.
same analysis as pred, $O(\log \log U)$ time.

del(x) is similar.

All operations in $O(\log \log U)$ time.

Space: $O(U)$ (too large!)

binary trie

write each integer as a $\log U$ -bit binary string

binary trie:

- complete binary tree with height = $\log U$
- left child = bit 0, right child = bit 1
- each leaf = a $\log U$ -bit integer

pred(x):


- binary search to find the lowest ancestor p of x in the trie

- if p has no right child

• return $p.\text{left.max}$

- if p has no left child

• return $p.\text{right.min.pred.}$



Binary search takes $O(\log \log U)$ time.

Space: $O(n \log U)$

Each ins/del may need to update the ptrs in at most $O(\log U)$ ancestors, $O(\log U)$ time.