

Advanced Data Structures

Skip Lists
+ Scapegoat
Trees



Recap:

- HW1: partially done,
posted by Wed.
- Sub on Friday & next
Wednesday
(no class next Monday)

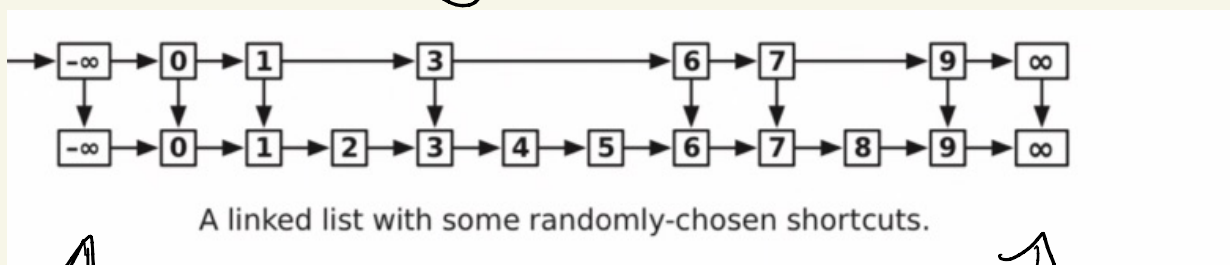
Next: Skip Lists

(Bill Pugh, 1990)

An alternative to balanced binary search trees.

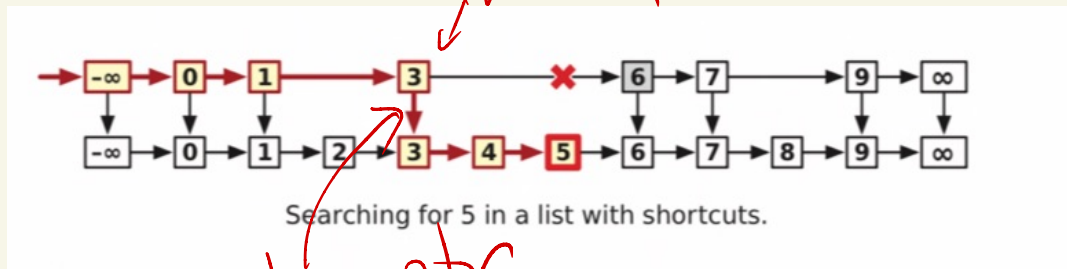
Essentially, just a sorted list where we add shortcuts - but to speed up, we'll duplicate some elements.

For each item, duplicate with probability $\frac{1}{2}$:



↑
plus some sentinel nodes

Searching:



Scan in top list.

If found, great!

Otherwise:

Look until next element is too large

Follow down ptr
& scan lower list

Some probability!

Expectation: $\sum_{\text{values possible}} (\text{value}) (\text{prob of value})$

Ex: 6 sided dice

$$E[\text{value}] = \frac{1}{6} \cdot 1 + \frac{1}{6} \cdot 2 + \frac{1}{6} \cdot 3 + \dots + \frac{1}{6} \cdot 6 \\ \underline{\underline{= 3.5}}$$

Each node is copied with prob = $\frac{1}{2}$,
 $E[\# \text{ nodes in top}] = \frac{n}{2}$

Goal: Bound expected # of comparisons
in top is $n/2$ # comp.


Prob [a node is follow by k
without duplicates]

$$= \underbrace{\frac{1}{2} \cdot \frac{1}{2} \cdots \frac{1}{2}}_k = \frac{1}{2^k}$$

So: Expected [# comparisons
in lower list]

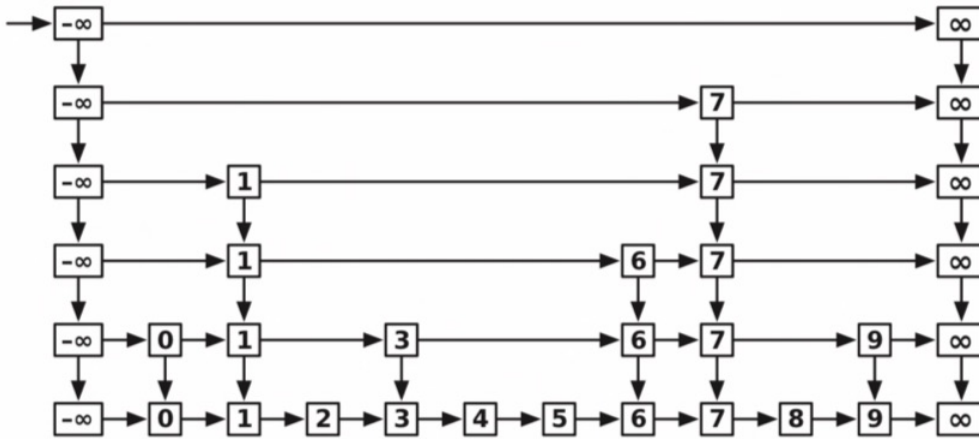
$$= \underline{1} + 1 \cdot \frac{1}{2} + 1 \cdot \frac{1}{2^2} + \dots + \frac{1}{2^k} = \sum_{k=0}^n 2^{-k} \leq 2$$

\Rightarrow expect $\frac{n}{2} + 2$ comparisons



Searching for 5 in a list with shortcuts.

What next? Rearchse!



A skip list is a linked list with recursive random shortcuts.

To search!

SKIPLISTFIND(x, L):

$v \leftarrow L$

while ($v \neq \text{NULL}$ and $\text{key}(v) \neq x$)


 if $\text{key}(\text{right}(v)) > x$

$v \leftarrow \text{down}(v)$

 else

$v \leftarrow \text{right}(v)$

return v



Searching for 5 in a skip list.

How many levels?

$$\begin{aligned} \text{Well, } E[\text{size at level } i] \\ = \frac{1}{2} E[\text{size at level } i-1] \end{aligned}$$

So (intuitively):

$O(\log n)$ runtime

Each time we add a level,

$E[\# \text{ searches}]$
goes down by $\frac{1}{2}$.

More formally?


See posted notes!

(Assumes some probability...)

Binary Search Trees:

What is the "best" one?

Recap:



Search:

start at root
if $v == \text{target}$
return yes
else if $v < \text{target}$
recurse left


Insert:

else
recurse right

while (v has children)

if $x \leq v$
else go left
go right

Delete:



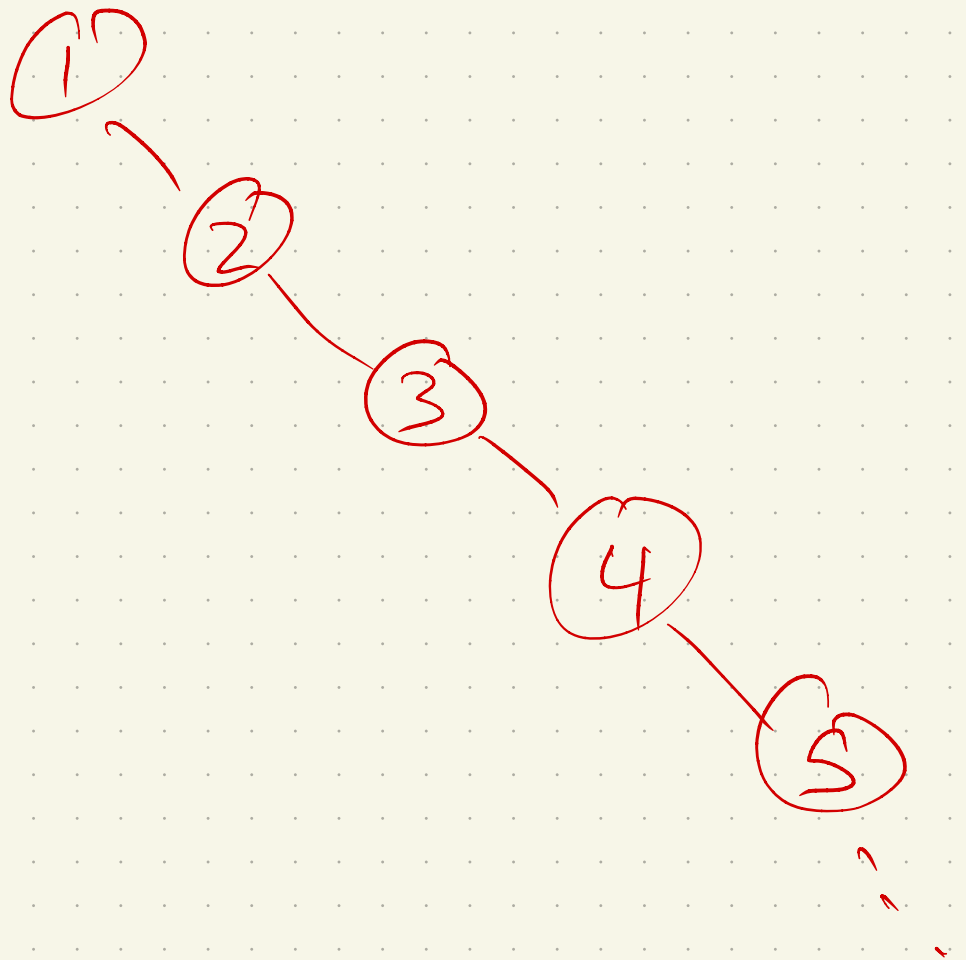
~ Find next node on in order traversal

Data Structures Class

- "Vanilla" BSTs (no rotations or balancing)


Runtime: $O(n)$

How can it get this bad?



BSTrees : balancing

Rotation/pivot : $\leftarrow \text{height}(y) = \text{height}(x) + 1$



unbalanced: left (or right)

- AVL trees \leftarrow is too big
- Red-Black trees \leftarrow want $|h(l(v)) - h(r(v))| \leq 1$

~~Today~~ : - Scapegoat Trees

This week - Splay Trees

Terminology I'll assume:

- search key
- node
- left/right child, parent
- internal/leaf node
- root
- ancestor/descendant
- preorder, inorder, postorder

Recap:

- Height(v): distance to
↳ furthest leaf in v 's
subtree
- Depth(v): distance from
 v to the root
- Size(v): # of nodes in
 v 's subtree

Scapegoat Trees:

[Anderson '89, Galperin-Rivest '93]

Supports amortized $O(\log n)$.

Basic idea:

- Standard BST search
- Delete: mark "deleted" node.
When tree is half dirty, rebuild into perfect tree.

Runtime:

Claim: rebuild a perfect tree in linear time

$\Rightarrow O(\log n)$ amortized time

And insert:

Standard insert

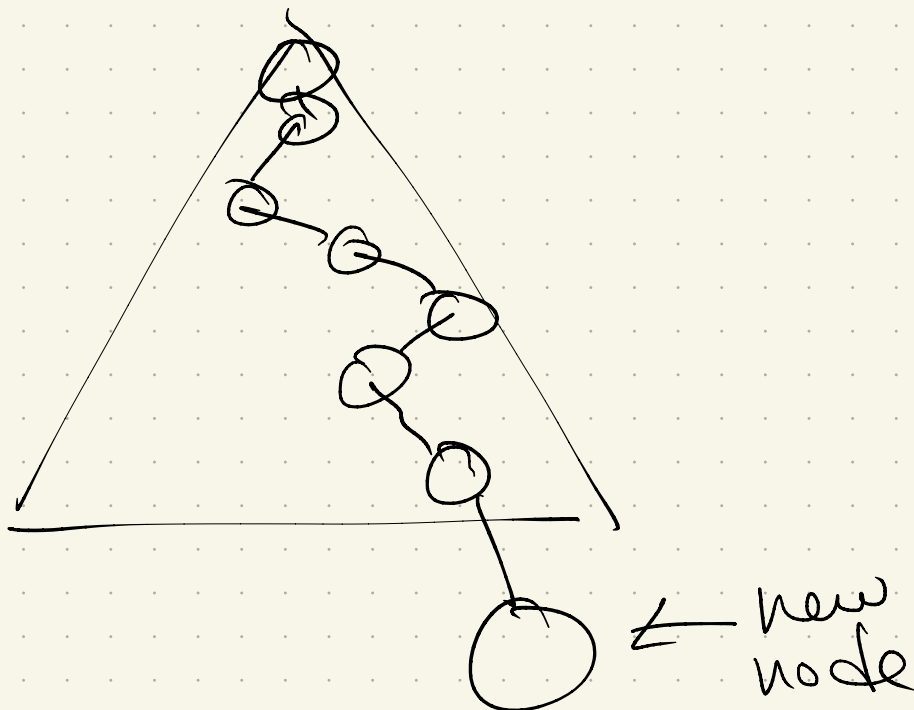
But: if imbalanced,
rebuild a subtree
containing new leaf

Dfn: Fix any $\alpha > 2$.

A node is imbalanced

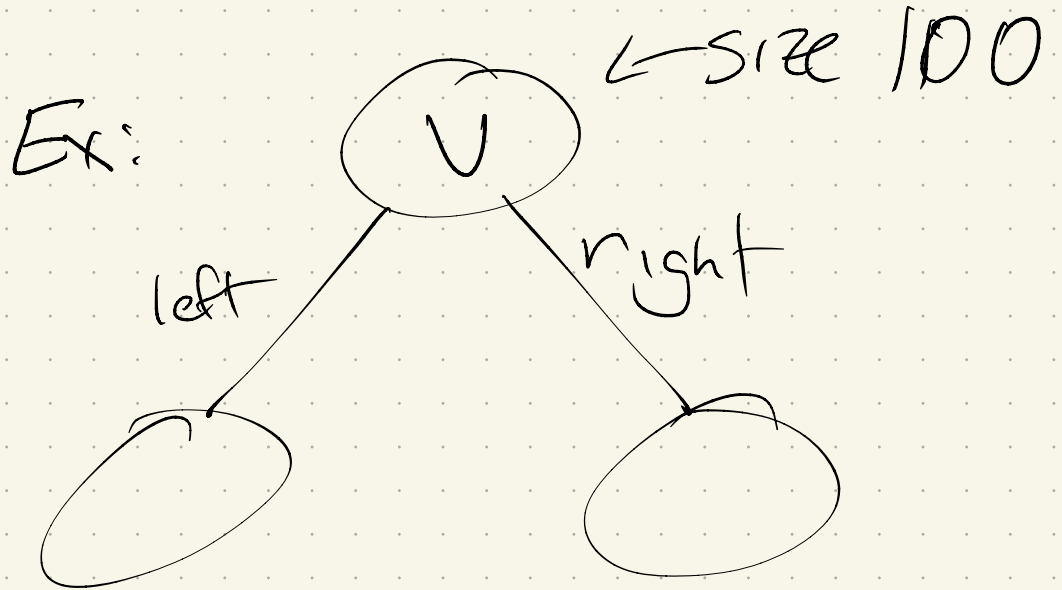
if $\text{height}(v) > \alpha \lg(\text{size}(v))$

So here:



Let : $I(v) =$

$$\max \{ 0, | \text{size}(\text{left}(v)) - \text{size}(\text{right}(v)) - 1 | \}$$



Lemma:

Just before rebuilding at v ,
 $I(v) = \Omega(n)$

proof:

If imbalanced, $h(v) > \alpha(\lg \text{size}(v))$
(by defn of imbalanced)
but $\text{left}(v)$ & $\text{right}(v)$
were not imbalanced:

$$h(\text{left}(v)) \leq$$

$$h(\text{right}(v)) \leq$$

Wlog:

Assume insert on left; so:



Some intense math:

So: takeaway

$$I(v) = \Omega(\text{size}(v))$$


This means $\sim \text{size}(v)$ insertions
since the last rebuilding.

So: rebuild! How?

Several ways to do this
in $O(\text{size}(v))$ time.

(HW question!)

Claim: ≤ 1 tree rebuild
for each insertion



Find runtime:

Find:

Delete:

Insert: