

Prof. Erik Demaine & Dr. André Schulz
TA: Aleksandar Zlateski

Topics: comparison, graph, integer, geometric,
string, succinct, cache-efficient DSS

Administration:

- Signup sheet
- requirements: homework, scribing, project
- listeners welcome
- problem session (starting ~ week 3)
- scribe for today

TODAY: Dynamic Optimality I

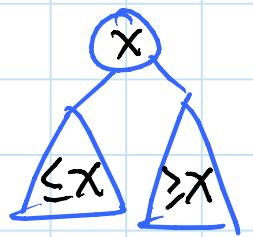
- binary search trees
- analytic bounds
- splay trees
- geometric view
- greedy algorithm

Q: is there one best binary search tree (BST)?

BST: comparison data structure

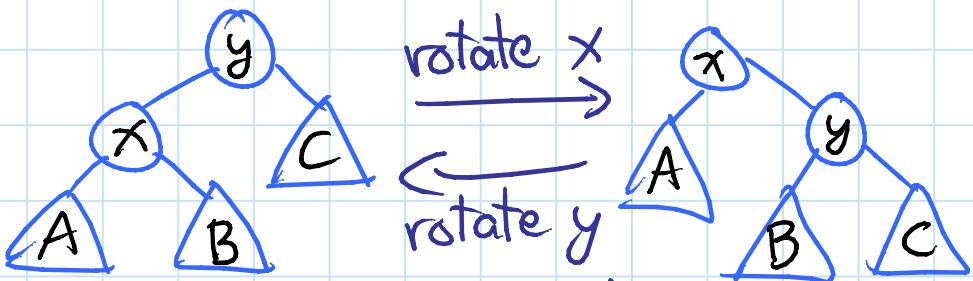
supporting search

(& predecessor/successor, insert/delete)



Also a model of computation (for DSSs)

- data must be stored in a BST
- unit-cost operations:
 - walk left, right, or up (parent)
 - rotate this node & its parent



(- create/destroy leaf)

⇒ search cost = length of root-to-node path

DSSs in this model:

- vanilla BST (no rotations)

- AVL trees

- red-black trees (B-trees)

- weight-balanced trees, treaps

- splay trees

- Tango trees

- Greedy

$O(\lg n)$
/ op.

} focus
here

Is $O(\lg n)$ search optimal?

- depends on sequence of searches
- say we're storing keys $\{1, 2, \dots, n\}$ & search for x_1, x_2, \dots, x_m

Sequential access property:

$1, 2, \dots, n \Rightarrow O(1)$ amortized / op.

[in-order traversal in any BST]

Dynamic finger property:

$|x_i - x_{i-1}| = k \Rightarrow O(\lg k)$ / op. possible

[think level-linked B-trees ~ but BST] *

Entropy bound / static optimality: best possible without rotation

k appears p_k fraction of the time $\Rightarrow O\left(\sum_{k=1}^n p_k \lg \frac{1}{p_k}\right)$ / op.

[store x_i at height $\leq \lg \frac{1}{p_k} + 1$]

[Jacomo - SWAT 2000]

Working-set property:

if t_i distinct keys accessed since last access to x_i , then $O(\lg t_i)$ possible

[intuition: store most recent higher up] *

\Rightarrow if all $x_i \in S$ then $O(\lg |S|)$ / op. possible

[form BST on S , put rest below]

* = hard to do with BST, but possible!

Unified property: [Iacono - SODA 2001]

if t_{ij} distinct keys accessed in x_i, \dots, x_j
then x_j costs $O(\lg \min_i [\underbrace{|x_i - x_j|}_{\text{space}} + \underbrace{t_{ij} + 2}_{\text{time}}])$

"fast if close to something recent" *

- e.g. $1, \frac{n}{2}, 2, \frac{n}{2}+1, 3, \frac{n}{2}+3, \dots \Rightarrow O(1)/\text{op.}$
- Implies both working set & unified
- possible on pointer machine [Iacono: Badiou, Cole, Demaine, Iacono - Algorithmica 2007]
- possible on BST up to additive $O(\lg \lg n)$ [Bose, Douieb, Dujmović, Howat - arXiv 2009]
- OPEN: possible on a BST?

Dynamic optimality / $O(1)$ -competitive:

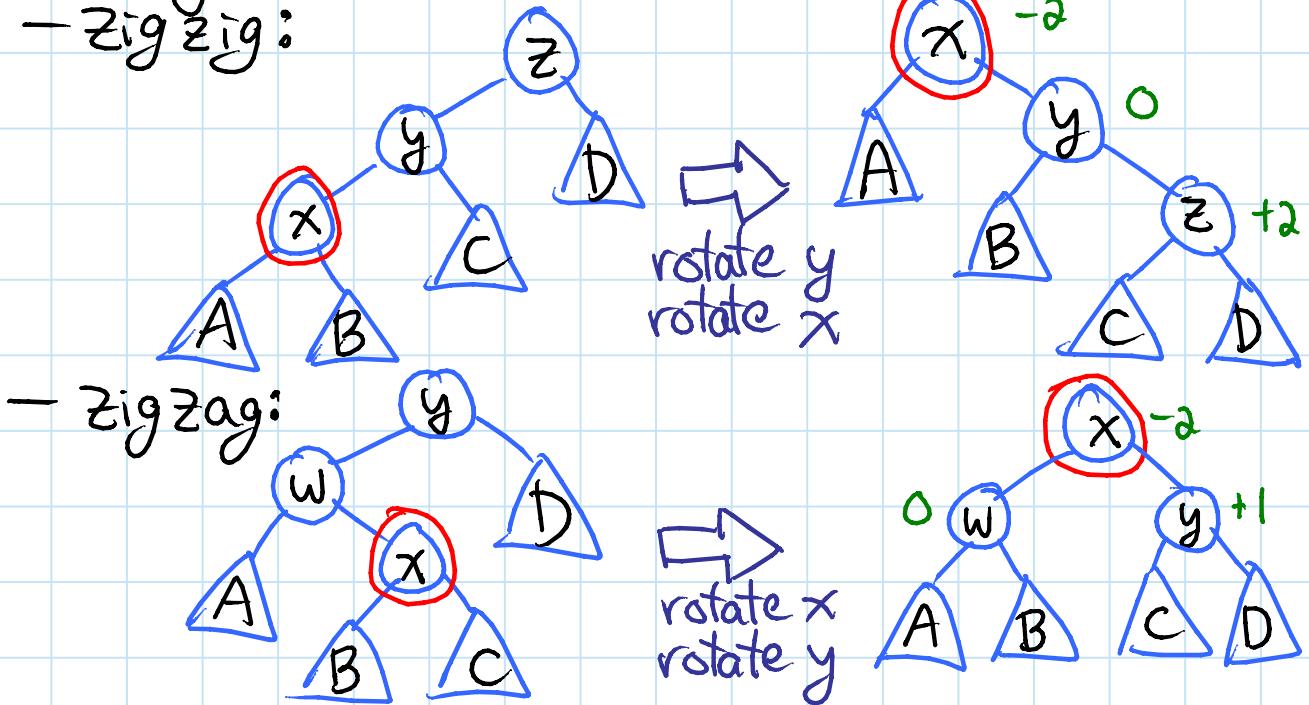
total cost = $\underbrace{O(\text{OPT})}$

min. cost of any BST on this access sequence

- OPEN: is this possible for any BST?
(online)
for any pointer-machine DS?

Splay trees: [Sleator & Tarjan - JACM 1985]

- binary search for x
- modify the path:
 - zigzag:



- at the end, possible single rotation to put x at root
- key feature: at most half the nodes on the path go down in the tree

Performance: (amortized)

- has working-set property
- has dynamic-finger property

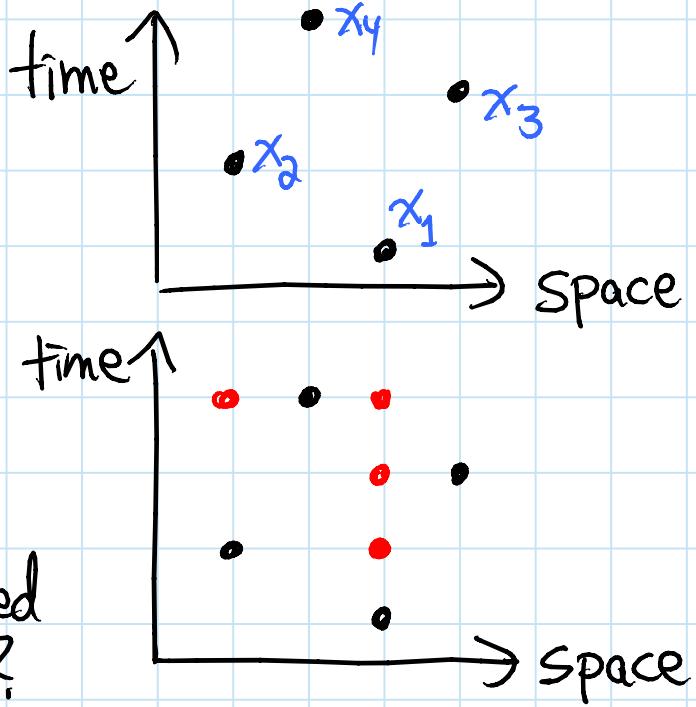
[Sleator & Tarjan]
[Cole - SICOMP 2000]

- **CONJECTURE**: has unified property
- **CONJECTURE**: dynamically optimal

Geometric view:

[Demaine, Harmon, Iacono,
Kane, Pătrașcu - SODA 2009]

access sequence
→ point set
 $\{(x_i, i)\}$

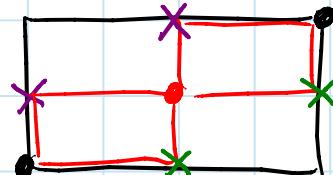


BST execution
→ point set:
which nodes touched
during search(x_i)?

Theorem: point set is a valid BST execution
 \Leftrightarrow Arborally Satisfied Set (ASS)

↳ rectangle spanned by two points
in set, not on horizontal/vertical line,
contains another point

- in fact must have another point
on a rectangle
side incident
to either corner:



Corollary: OPT = smallest ASS containing input

OPEN: complexity? $O(1)$ -approximation?

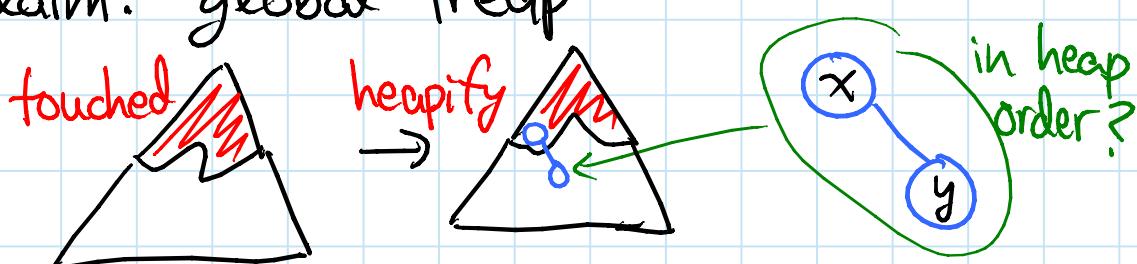
Proof of Theorem:

(\Rightarrow) let $a = \text{lca of } x_i \& x_j$ (changes over time)

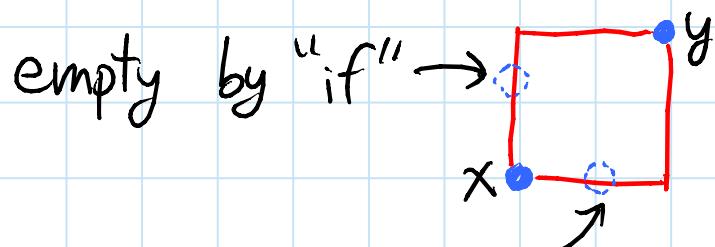
- so $x_i \leq a \leq x_j$ & a is ancestor of $x_i \& x_j$
- if $a \neq x_i$ just before $\text{search}(x_i)$
then $(a, i) \in \text{execution}$
- if $a \neq x_j$ just before $\text{search}(x_j)$
then $(a, j) \in \text{execution}$
- else: a changes from x_i to x_j
between times i & j
 $\Rightarrow x_j$ rotated before $\text{search}(x_j)$

\Rightarrow get third point in $(x_i, i) \rightarrow (x_j, j)$ rectangle
in all cases

- (\Leftarrow) define tree at all times to be treap:
 BST & heap ordered by next-touch-time
- note: next-touch-time has some ties,
 So this is not uniquely defined
 - when we reach time i , nodes to touch
 form a connected subtree at the top
 (by heap-order property)
 - these nodes get new next-touch-time
 - re-arrange into local treap
 (this still may be ambiguous — break ties
 arbitrarily — but still restricts global choice)
 - claim: global treap

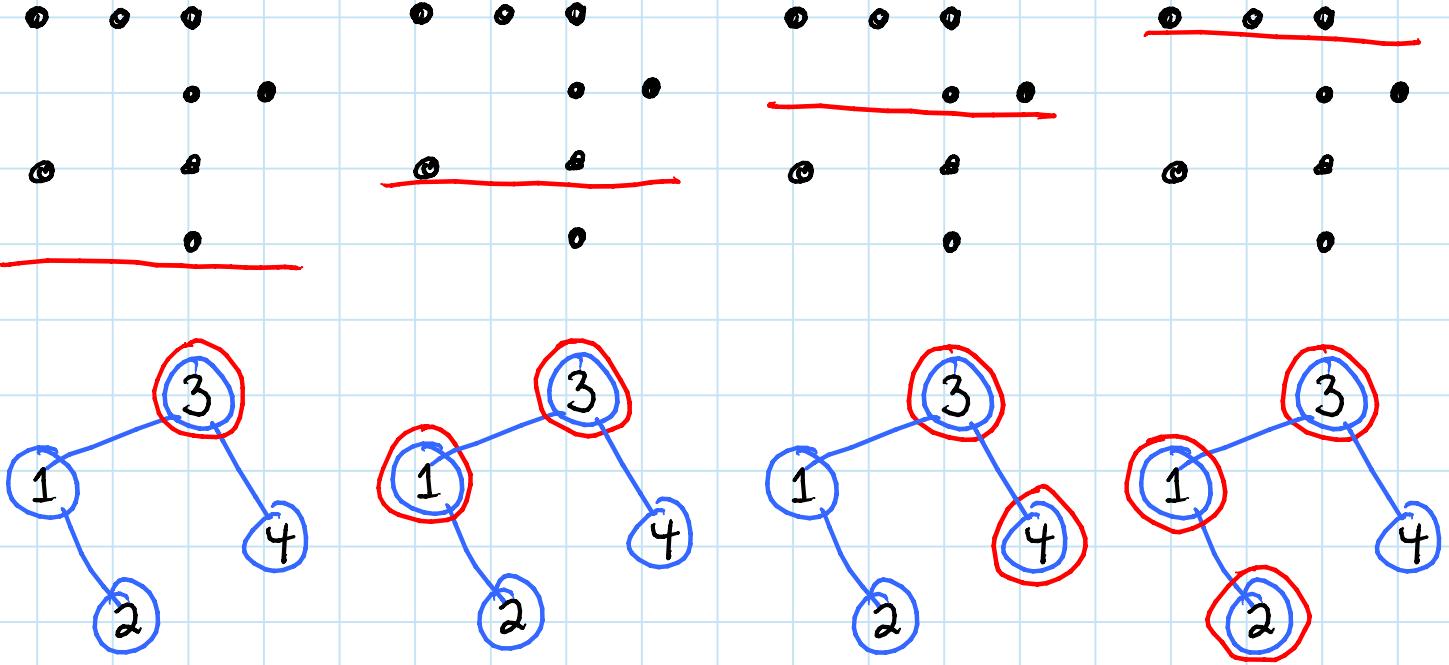


if y to be touched sooner than x
 then $(x, \text{next-touch}) \rightarrow (y, \text{next-touch})$
 is an unsatisfied rectangle:
 (according to 2nd definition of ASS)



leftmost such point would be right child
 of x after $\text{search}(x_i)$, not y \square

Simple example:



Greedy algorithm: [Lucas 1988; Munro 2000]

- consider point set one row at a time
- add the necessary points on that row
- in tree view: re-arrange root-to- x path optimally for future searches

CONJECTURE: Greedy = $O(OPT)$

or even: = $OPT + O(m)$

- seems obvious... "just" need to show you needn't stray from the access path

So what?

Theorem: online ASS algorithm

→ online BST (with $O(1)$ slowdown)

Corollary: Greedy is actually an online BST!

- Conjecture \Rightarrow dynamically optimal