

Week 7 Lecture Notes

Topics: Fixed universe successor
Van Emde Boas data structure
Amortized analysis
Computation Geometry

Fixed-universe successor problem

Goal: Maintain a dynamic subset S of size n of the universe $U = \{0, 1, \dots, u-1\}$ of size u subject to these operations.

INSERT ($x \in U \setminus S$): Add x to S

DELETE ($x \in S$): Remove x from S

SUCCESSOR ($x \in U$): Find the next element in S larger than any element y of the universe U

PREDECESSOR ($x \in U$): Find the previous element in S smaller than x .

Solutions to fixed-universe successor problem

Goal: Maintain a dynamic subset S of size n of the universe $U = \{0, 1, \dots, u-1\}$ of size u , subject to **INSERT, DELETE, SUCCESSOR, PREDECESSOR**

- Balanced search trees can implement operations in $O(lgn)$ time, without fixed-universe assumptions.
- In 1975, Peter van Emde Boas solved this problem in $O(\lg \lg u)$ time per operation.
 - If u is only polynomial in n , that is, $u = O(n^c)$, then $O(\lg \lg n)$ time per operation - exponential speedup

$O(\lg \lg u)$?!

Where could a bound of $O(\lg \lg u)$ arise?

- Binary search over $O(\lg u)$ things
- $T(u) = T(\sqrt{u}) + O(1)$

$$\begin{aligned} T'(\lg u) &= T'((\lg u)/2) + O(1) \\ &= O(\lg \lg u). \end{aligned}$$

1). Starting point: Bit Vector

Bit vector v stores, for each $\alpha \in U$

$$v_\alpha = \begin{cases} 1 & \text{if } \alpha \in S \\ 0 & \text{if } \alpha \notin S \end{cases}$$

Example:

$$U = 16, n = 4, S = \{1, 9, 10, 15\}$$

0	1	0	0	0	0	0	0	0	1	1	1	0	0	0	0	1
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	

Insert/Delete run in $O(1)$ time

Successor/Predecessor run in $O(\sqrt{u})$ worst-case time

2. Split universe into widgets

Carve universe of size u into \sqrt{u} widgets.

$w_0, w_1, \dots, w_{\sqrt{u}-1}$ each of size \sqrt{u}

Example:

$$w_0 \quad u = 16, \sqrt{u} = 4$$

0	1	0	0	0	0	0	0	0	0	1	1	0	0	0	1
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15

w_0 represents $0, 1, \dots, \sqrt{u}-1 \in U$;

w_i represents $\sqrt{u}, \sqrt{u}+1, \dots, 2\sqrt{u}-1 \in U$;

:

w_i represents $i\sqrt{u}, i\sqrt{u}+1, \dots, (i+1)\sqrt{u}-1 \in U$;

:

$w_{\sqrt{u}-1}$ represents $u-\sqrt{u}, u-\sqrt{u}+1, \dots, u-1 \in U$.

Define $\text{high}(x) \geq 0$ and $\text{low}(x) \geq 0$

so that $x = \text{high}(x) \sqrt{u} + \text{low}(x)$.

1	0	0	1
$\text{high}(x)$		$\text{low}(x)$	
$= 1$			

That is, if we write $x \in U$ in binary,

$\text{high}(x)$ is the high-order half of the bits, and

$\text{low}(x)$ is the low-order half of the bits.

For $x \in U$, $\text{high}(x)$ is the index of widget containing x and $\text{low}(x)$ is the index of x within that widget.

w_0	w_1	w_2	w_3
0 1 0 0	0 0 0 0	0 1 1 0	0 0 0 1
0 1 2 3	4 5 6 7	8 9 10 11	12 13 14 15

INSERT(x)

1. insert x into widget $w_{\text{high}(x)}$ at position $\text{low}(x)$
2. mark $w_{\text{high}(x)}$ as nonempty

Running time: $T(u) = O(1)$

SUCCESSOR(x)

1. look for successor of x within widget $w_{\text{high}(x)}$
 2. starting after position $\text{low}(x)$
 3. if successor found
 4. then return it
 5. else find smallest $i > \text{high}(x)$
 6. for which w_i is non-empty
 7. return smallest element in w_i
- $\} O(\sqrt{u})$
- $\} O(\sqrt{u})$
- $\} O(\sqrt{u})$

Running time: $T(u) = O(\sqrt{u})$

Revelation

SUCCESSOR (α)

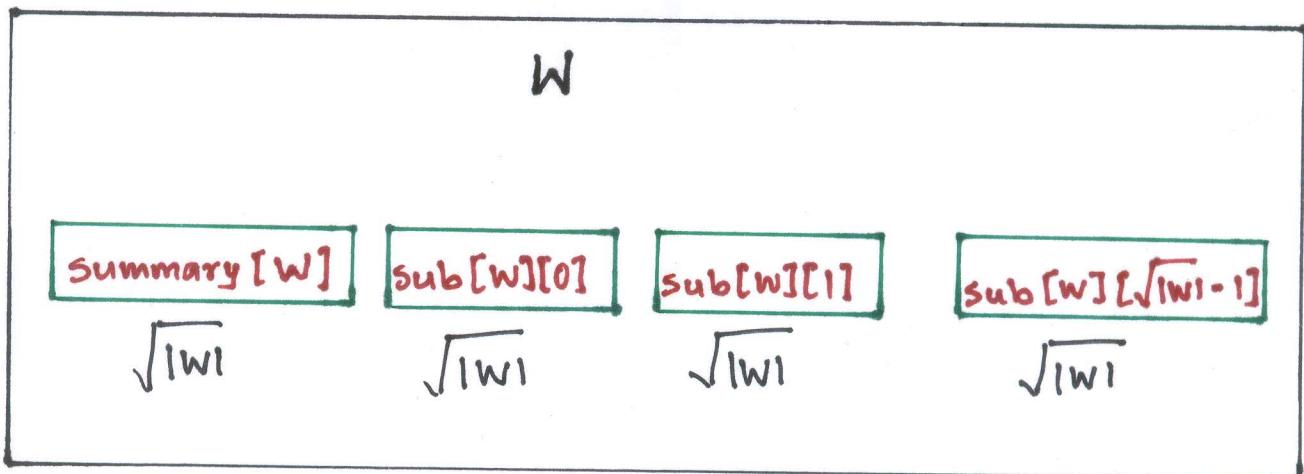
1. look for successor of α within widget $W_{high(\alpha)}$
 2. starting after position $low(\alpha)$
 3. if successor found
 4. then return it
 5. else find smallest $i > high(\alpha)$ for which W_i is nonempty
 6. return smallest element in W_i
- } recursive successor
} recursive successor
} recursive successor

3. Recursion

Represent universe by widget of size u

Recursively split each widget W of size $|W|$ into $\sqrt{|W|}$ subwidgets $\underline{\text{sub}[W][0]}, \underline{\text{sub}[W][1]}, \dots, \underline{\text{sub}[W][\sqrt{|W|}-1]}$ each of size $\sqrt{|W|}$

Store a summary widget $\text{summary}[W]$ of size $\sqrt{|W|}$ representing which subwidgets are non-empty.



Define $\text{high}(n) \geq 0$ and $\text{low}(n) \geq 0$

so that $a = \text{high}(n)\sqrt{|W|} + \text{low}(n)$

INSERT(x, w)

1. if $\text{sub}[w][\text{high}(n)]$ is empty
2. then $\text{INSERT}(\text{high}(n), \text{summary}[w])$
3. $\text{INSERT}(\text{low}(n), \text{sub}[w][\text{high}(n)])$

Running time: $T(u) = 2T(\sqrt{u}) + O(1)$

$$\begin{aligned} T'(\lg u) &= 2T'((\lg u)/2) + O(1) \\ &= O(\underline{\lg u}) \end{aligned}$$

SUCCESSOR(x, w)

$j \leftarrow \text{SUCCESSOR}(\text{low}(n), \text{sub}[w][\text{high}(n)]) \} T(\sqrt{u})$

if $j < \infty$

then return $\text{high}(n)\sqrt{|W|} + j$

else $i \leftarrow \text{SUCCESSOR}(\text{high}(n), \text{summary}[w]) \} T(\sqrt{u})$

$j \leftarrow \text{SUCCESSOR}(-\infty, \text{sub}[w][i]) \} T(\sqrt{u})$

return $i\sqrt{|W|} + j$

Running time: $T(u) = 3T(\sqrt{u}) + O(1)$

$$\begin{aligned} T'(\lg u) &= 3T'((\lg u)/2) + O(1) \\ &= O((\lg u)^{\lg 3}) \end{aligned}$$

Improvements

Need to reduce INSERT and SUCCESSOR down to 1 recursive call each.

- 1 call: $T(u) = 1T(\sqrt{u}) + O(1)$
 $= O(\lg \lg n)$
- 2 calls: $T(u) = 2T(\sqrt{u}) + O(1)$
 $= O(\log n)$
- 3 calls: $T(u) = 3T(\sqrt{u}) + O(1)$
 $= O((\lg n)^{\lg 3})$

We are closer to this goal than it may seem!

Recursive calls in successor

If x has a successor within $\text{sub}[w][\text{high}(x)]$, then there is only 1 recursive call to SUCCESSOR. Otherwise, there are 3 recursive calls.

- SUCCESSOR ($\text{low}(x), \text{sub}[w][\text{high}(x)]$)
discovers that $\text{sub}[w][\text{high}(x)]$ has no successor
- SUCCESSOR ($\text{high}(x), \text{summary}[W]$)
finds next non-empty subwidget $\text{sub}[w][i]$
- SUCCESSOR ($-\infty, \text{sub}[w][i]$)
finds smallest element in subwidget
 $\text{sub}[w][i]$

Reducing recursive calls in successor

If α has no successor within $\text{sub}[w][\text{high}(\alpha)]$, there are 3 recursive calls:

- $\text{SUCCESSOR}(\underline{\text{low}(\alpha)}, \underline{\text{sub}[w][\text{high}(\alpha)]})$ discovers that $\underline{\text{sub}[w][\text{high}(\alpha)]}$ hasn't successor
 - Could be determined using the maximum value in the subwidget $\underline{\text{sub}[w][\text{high}(\alpha)]}$
- $\text{SUCCESSOR}(\underline{\text{high}(\alpha)}, \underline{\text{summary}[w]})$ finds next nonempty subwidget $\underline{\text{sub}[w][i]}$
- $\text{SUCCESSOR}(-\infty, \underline{\text{sub}[w][i]})$
finds minimum element in subwidget $\underline{\text{sub}[w][i]}$.

Improved Successor

INSERT (α, W)

if $\text{sub}[W][\text{high}(\alpha)]$ is empty

then $\text{INSERT}(\text{high}(\alpha), \text{summary}[W])$

$\text{INSERT}([\text{low}(\alpha), \text{sub}[W][\text{high}(\alpha)])]$

if $\alpha < \min[W]$, then $\min[W] \leftarrow \alpha$ } new
augmentation

if $\alpha > \max[W]$, then $\max[W] \leftarrow \alpha$

Running Time: $T(u) = 2T(\sqrt{u}) + O(1)$

$$T'(\lg u) = 2T'((\lg u)/2) + O(1)$$
$$= O(\lg u)$$

SUCCESSOR (α, W)

if $\text{low}(\alpha) < \max[\text{sub}[W][\text{high}(\alpha)]]$

then $j \leftarrow \text{SUCCESSOR}([\text{low}(\alpha), \text{sub}[W][\text{high}(\alpha)])$ } $T(\sqrt{u})$

return $\text{high}(\alpha) \sqrt{|W|} + j$

else $i \leftarrow \text{SUCCESSOR}(\text{high}(\alpha), \text{summary}[W])$ } $T(\sqrt{u})$

$j \leftarrow \min[\text{sub}[W][i]]$

return $i \sqrt{|W|} + j$

Running Time: $T(u) = T(\sqrt{u}) + O(1)$

$$= O(\lg \lg u)$$

Recursive calls in insert

If $\text{sub}[w][\text{high}(n)]$ is already in $\text{summary}[w]$,
then there is only 1 recursive call to INSERT .

Otherwise, there are 2 recursive calls:

- $\text{INSERT}(\text{high}(n), \text{summary}[w])$
- $\text{INSERT}(\text{low}(n), \text{sub}[w][\text{high}(n)])$

Idea:

We know that $\text{sub}[w][\text{high}(n)]$ is empty.

Avoid second recursive call by specially
storing a widget containing just 1 element.

Specifically, do not store min recursively.

Improved insert

INSERT(x, w)

1. if $x < \min[w]$ then exchange $x \leftrightarrow \min[w]$
2. if $\text{sub}[w][\text{high}(x)]$ is nonempty, that is
 $\min[\text{sub}[w][\text{high}(x)]] \neq \text{NIL}$
3. then $\text{INSERT}(\text{low}(x), \text{sub}[w][\text{high}(x)])$
4. else $\min[\text{sub}[w][\text{high}(x)]] \leftarrow \text{low}(x)$
5. $\text{INSERT}(\text{high}(x), \text{summary}[w])$
6. if $x > \max[w]$ then $\max[w] \leftarrow x$

Running Time: $T(u) = 1 T(\sqrt{u}) + O(1)$
 $= O(\lg \lg u)$

SUCCESSOR(x, w)

1. if $x < \min[w]$ then return $\min[w]$
2. if $\text{low}(x) < \max[\text{sub}[w][\text{high}(x)]]$
3. then $j \leftarrow \text{SUCCESSOR}(\text{low}(x), \text{sub}[w][\text{high}(x)])$
4. return $\text{high}(x), \sqrt{|w|} + j$
5. else $j \leftarrow \text{SUCCESSOR}(\text{high}(x), \text{summary}[w])$
6. $j \leftarrow \min[\text{sub}[w][\epsilon]]$
7. return $i, \sqrt{|w|} + j$

Running Time: $T(u) = 1 T(\sqrt{u}) + O(1)$
 $= O(\lg \lg u)$

Deletion

$\text{DELETE}(x, w)$

1. if $\min[w] = \text{NIL}$ or $x < \min[w]$ then return
2. if $x = \min[w]$
3. then $i \leftarrow \min[\text{summary}[w]]$
4. $x \leftarrow i\sqrt{|w|} + \min[\text{sub}[w][i]]$
5. $\min[w] \leftarrow x$
6. $\text{DELETE}(\text{low}(x), \text{sub}[w][\text{high}(x)])$
7. if $\text{sub}[w][\text{high}(x)]$ is now empty, that is
 $\min[\text{sub}[w][\text{high}(x)]] = \text{NIL}$
8. then $\text{DELETE}(\text{high}(x), \text{summary}[w])$

(in this case, the first recursive call was cheap)

Amortized Algorithms

How large should a hash table be?

Goal:

Make the table as small as possible, but large enough so that it won't overflow (or otherwise become inefficient).

Problem:

What if we don't know the proper size in advance?

Solution:

Dynamic Tables

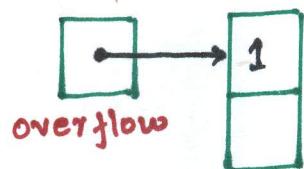
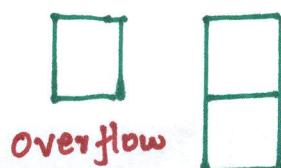
Idea:

Whenever the table overflows, "grow" it by allocating (via `malloc` or `new`) a new, larger table. Move all items from the old table into the new one, and free the storage for the old table.

Example of a dynamic table

1. INSERT

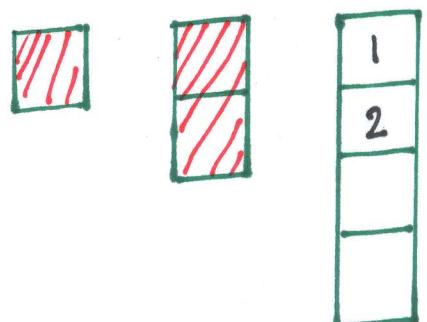
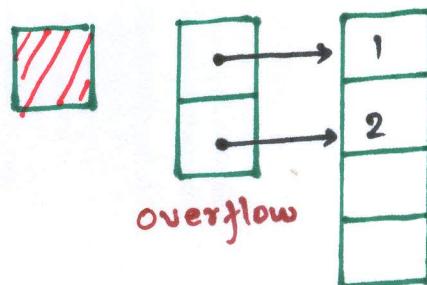
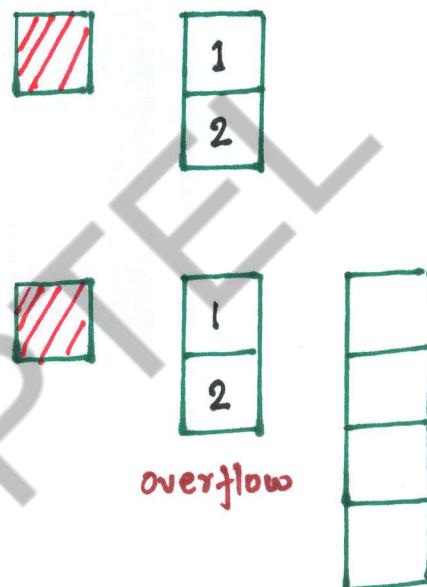
2. INSERT



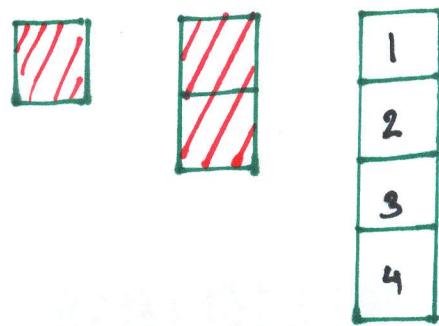
1. INSERT

2. INSERT

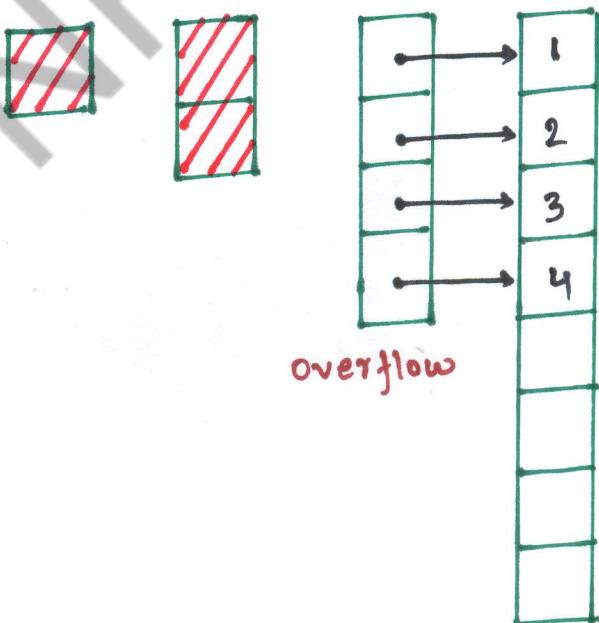
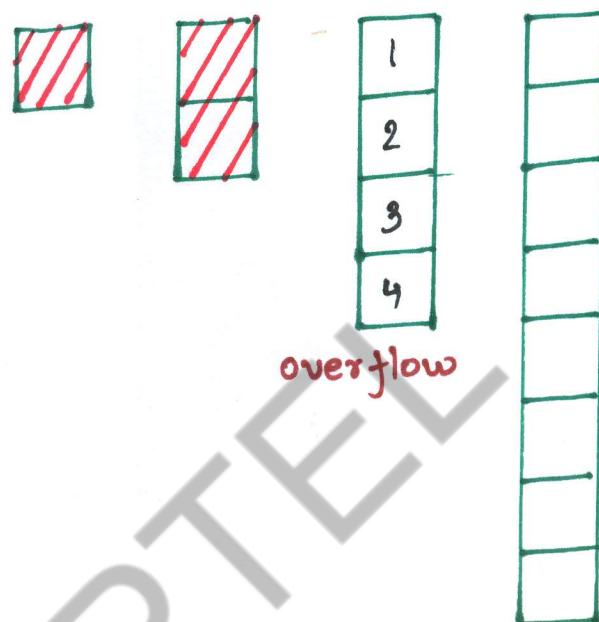
3. INSERT



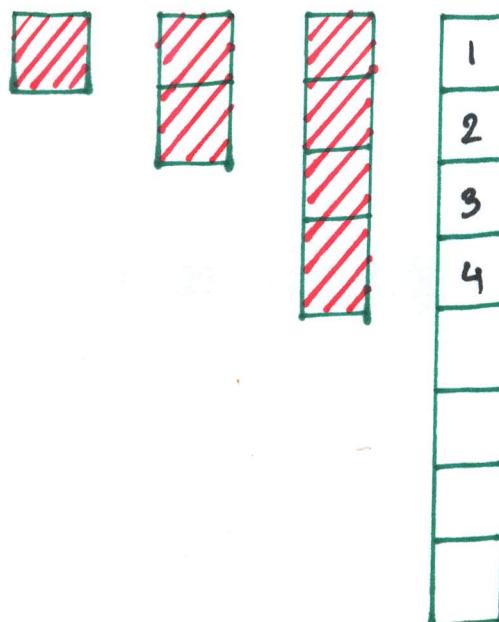
1. INSERT
2. INSERT
3. INSERT
4. INSERT



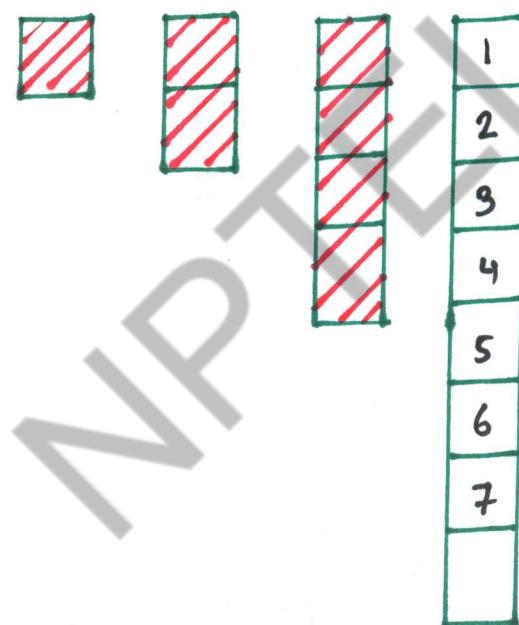
1. INSERT
2. INSERT
3. INSERT
4. INSERT
5. INSERT



1. INSERT
2. INSERT
3. INSERT
4. INSERT
5. INSERT



1. INSERT
2. INSERT
3. INSERT
4. INSERT
5. INSERT
6. INSERT
7. INSERT



Worst-case analysis

Consider a sequence of n insertions. The worst-case time to execute one insertion is $\Theta(n)$.

Therefore, the worst case time for n insertions is:

$$n \cdot \Theta(n) = \Theta(n^2).$$

WRONG! In fact, the worst case cost for n insertions is only $\Theta(n) \ll \Theta(n^2)$

Let us see why?

Tighter analysis

Let c_i = the cost of i^{th} insertion

$$= \begin{cases} i, & \text{if } i-1 \text{ is an exact power of 2;} \\ 1, & \text{otherwise.} \end{cases}$$

i	1	2	3	4	5	6	7	8	9	10
size_i	1	2	4	4	8	8	8	16	16	
c_i	1	2	3	1	5	1	1	9	1	

Tighter analysis (continued)

i	1	2	3	4	5	6	7	8	9	10
size _i	1	2	4	4	8	8	8	8	16	16
c _i	1	1	1	1	1	1	1	1	1	1
	1	2		4					8	

$$\begin{aligned}
 \text{Cost of } n \text{ operations} &= \sum_{i=1}^n c_i \\
 &\leq n + \sum_{j=0}^{\lfloor \lg(n-1) \rfloor} 2^j \\
 &\leq 3n \\
 &= \Theta(n)
 \end{aligned}$$

Thus, the average cost of each dynamic-table operation is:

$$\frac{\Theta(n)}{n} = \Theta(1)$$

Amortized Analysis

An **amortized analysis** is any strategy for analyzing a sequence of operations to show that the average cost per operation is small, even though a single operation within the sequence might be expensive.

Even though we are taking averages, however, probability is not involved!

- An amortized analysis guarantees the average performance of each operation in the worst case.

Types of amortized analyses

Three common amortization arguments:

the **aggregate** method,

the **accounting** method,

the **potential** method.

We have just seen an aggregate analysis.

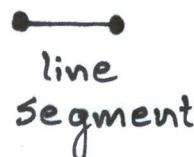
The aggregate method, though simple, lacks the precision of other two methods.

In particular, the accounting and potential methods allow a specific **amortized cost** to be allocated to each operation.

Computational Geometry

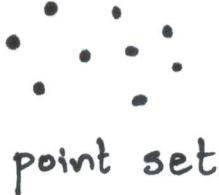
- Algorithms for solving "geometric problems" in 2D and higher.
- Fundamental Objects

point



line

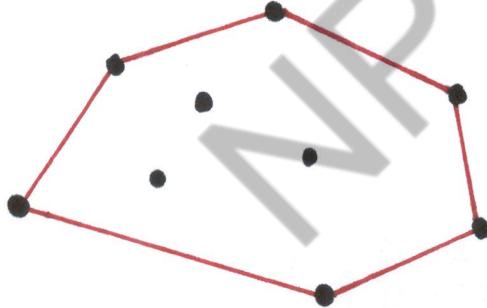
- Basic Structures



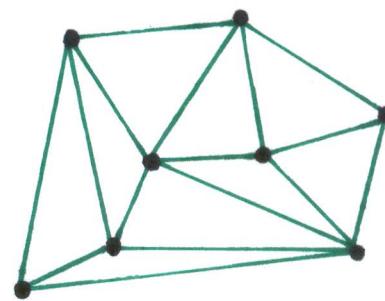
point set



polygon



Convex hull



triangulation.

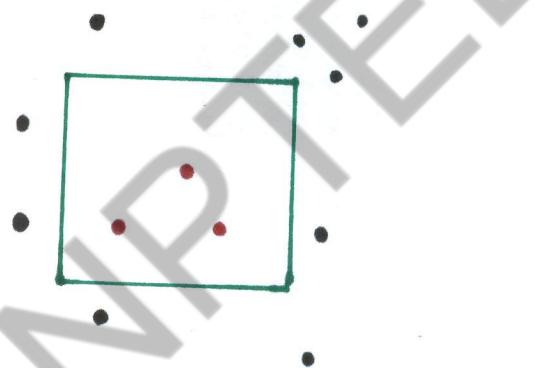
Orthogonal Range Searching

Input: n points in d dimensions.

- E.g. representing a database of n records each with d numeric fields.

Query: Axis aligned box (in 2D, a rectangle)

- Reports on the points inside the box:
 - Are there any points?
 - How many are there?
 - List the points.



Goal: Preprocessing points into data structure to support fast queries.

- Primary goal: Static data Structure
- In 1D we will also obtain a dynamic data structure supporting insert and delete.

1D range searching

In 1D, the query is an interval



First solution using ideas we know:

- Interval trees
 - Represent each point x by the interval $[x, x]$.
 - Obtain a dynamic structure that can list K answers in a query in $O(K \lg n)$ time.

Second solution using ideas we know:

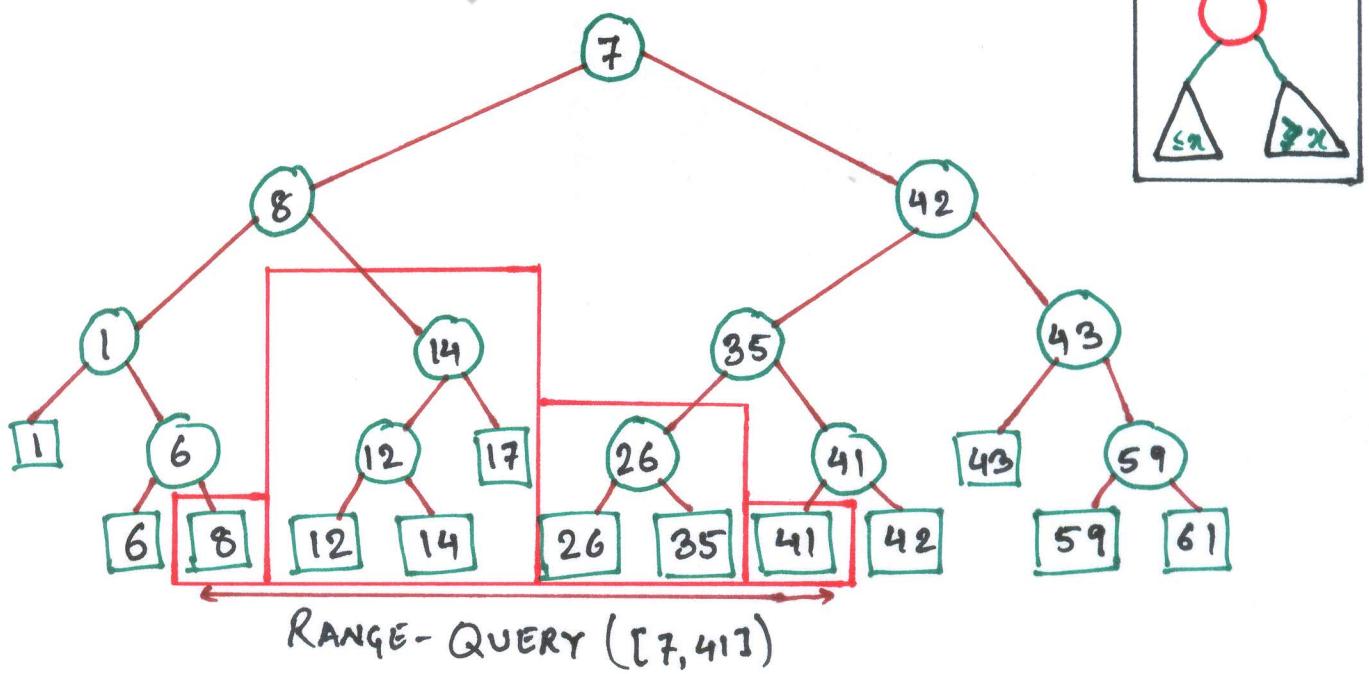
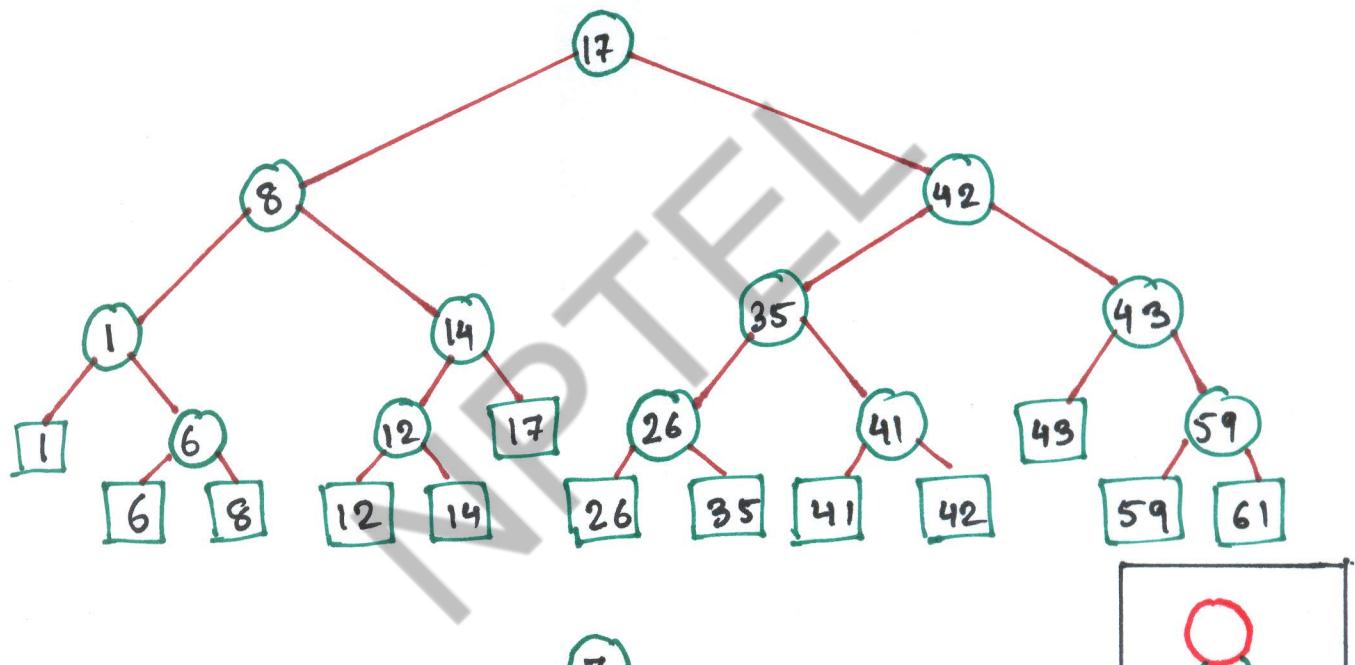
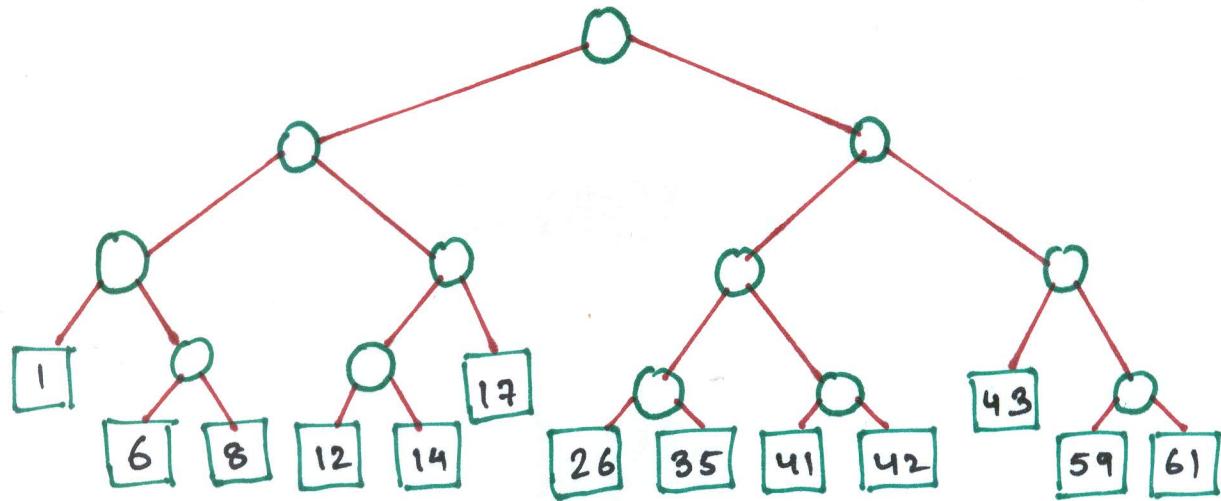
- Sort the points and store them in an array
 - Solve query by binary search on endpoints
 - Obtain a static structure that can list K answers in a query in $O(k + \log n)$ time.

Goal: Obtain a dynamic structure that can list K answers in a query in $O(k + \log n)$ time.

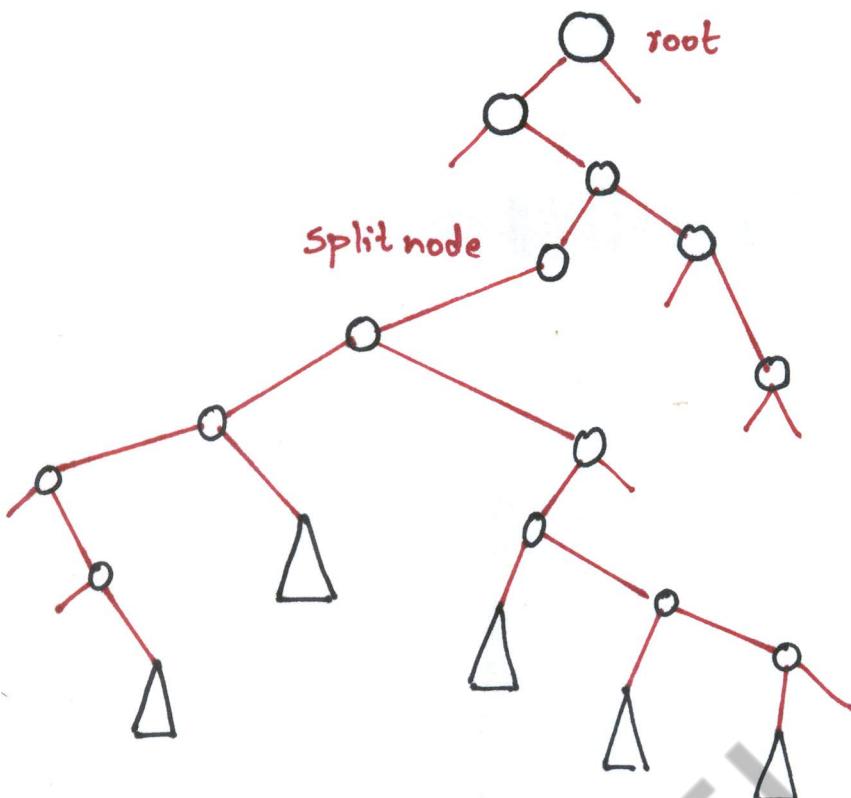
New solution that extends to higher dimensions:

- Balanced binary search tree
 - New organization principle:
 - store points in the leaves of the tree
 - Internal nodes store copies of the leaves to satisfy binary search property.
 - Node x stores in $\text{key}[x]$ the maximum key of any leaf in the left subtree of x .

Example of a 1D range tree



General 1D range query



Pseudocode, part 1: Find the split node

1D-RANGE-QUERY ($T, [x_1, x_2]$)

1. $w \leftarrow \text{root}[T]$
2. while w is not a leaf and ($x_2 \leq \text{key}[w]$ or $\text{key}[w] < x_1$)
 3. do if $x_2 \leq \text{key}[w]$
 4. then $w \leftarrow \text{left}[w]$
 5. else $w \leftarrow \text{right}[w]$
 - w is now the split node

traverse left and right from ' w ' and report relevant subtrees.

Pseudo code, part 2: Traverse left and right from Split Node.

1D-RANGE-QUERY ($T, [x_1, x_2]$)

[find the split node]

► w is now the split node

1. if w is leaf

2. then output leaf w if $x_1 \leq \text{key}[w] \leq x_2$

3. else $v \leftarrow \text{left}[w]$ ► left traversal

4. while v is not a leaf

5. do if $x_1 \leq \text{key}[v]$

6. then output the subtree rooted at $\text{right}[v]$

7. $v \leftarrow \text{left}[v]$

8. else $v \leftarrow \text{right}[v]$

9. output the leaf v if $x_1 \leq \text{key}[v] \leq x_2$

[symmetrically for right traversal]

Analysis of 1D-range query

Query Time: Answer to range query represented by $O(\log n)$ subtrees found in $O(\log n)$ time.

Thus

- Can test for points in interval in $O(\log n)$ time
- Can count points in interval in $O(\log n)$ time if we augment the tree with subtree sizes.
- Can report the first K points in $O(K + \log n)$ time.

Space: $O(n)$

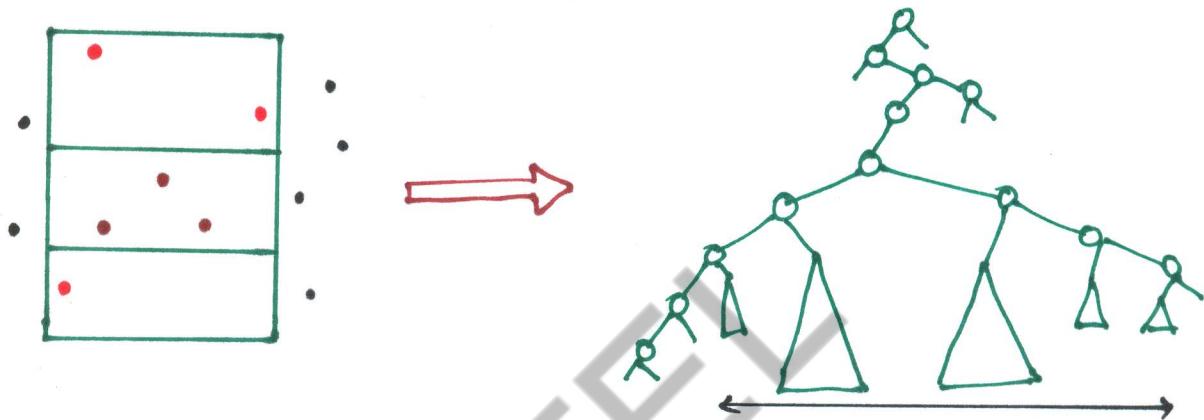
Preprocessing time: $O(n \log n)$

2D range trees

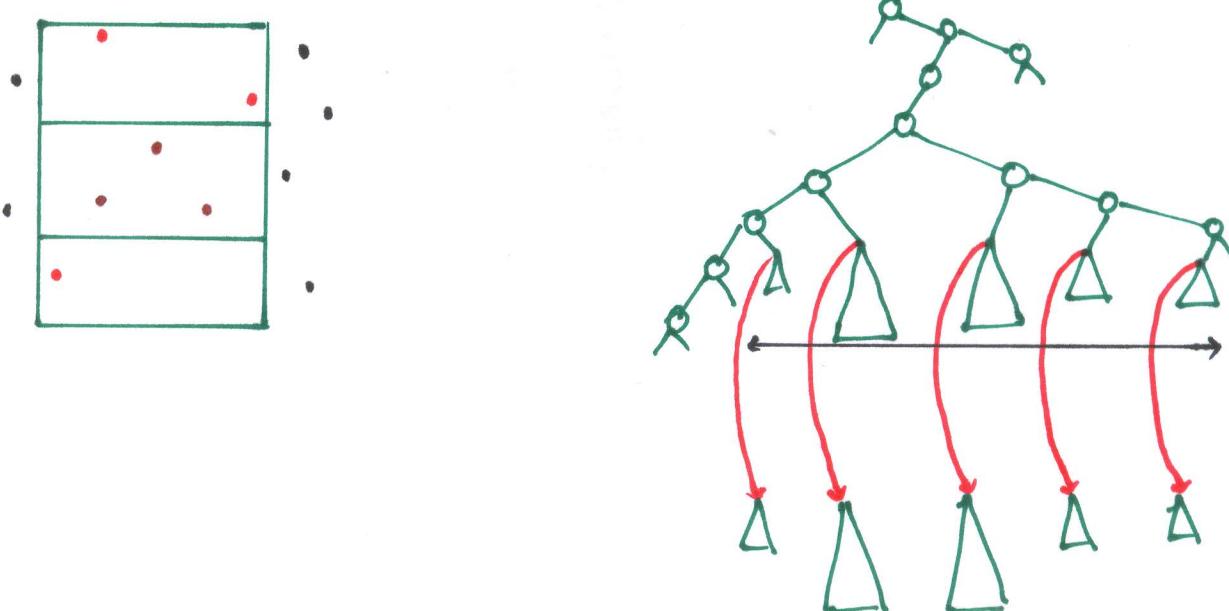
Store a primary 1D range tree for all the points based on x-coordinate

Thus in $O(\log n)$ time we can find $O(\log n)$ subtrees representing the points with proper x-coordinates.

How to restrict to points with proper y-coordinates?



Idea: In primary 1D range tree of x-coordinate, every node stores a secondary 1D range tree based on y-coordinate for all points in the subtree of the node
Recursively search within each.



Analysis of 2D range tree

Query time: In $O(\lg^2 n) = O((\lg n)^2)$ time, we can represent answer to range query by $O(\lg^2 n)$ subtrees. Total cost for reporting K points: $O(K + (\lg n)^2)$.

Space: The secondary trees at each level of the primary tree together store a copy of the points. Also, each point is present in each secondary tree along the path from the leaf to the root. Either way, we obtain that space is $O(n \lg n)$.

Preprocessing time: $O(n \lg n)$

d-dimensional range trees

Each node of secondary y-structure stores a tertiary z-structure representing the points in the subtree rooted at the node, etc.

Query time: $O(K + \log^d n)$ to report K points

Space: $O(n \log^{d-1} n)$

Preprocessing time: $O(n \log^{d-1} n)$

Best Data Structure to date:

Query time: $O(K + \log^{d-1} n)$ to report K points.

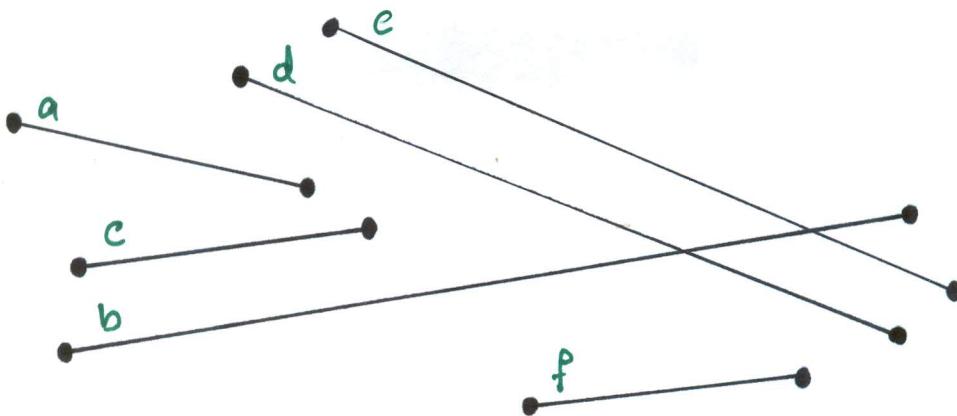
Space: $O(n(\lg n / \log \lg n)^{d-1})$

Preprocessing time: $O(n \log^{d-1} n)$

Line-segment intersection

Given n line segments, does any pair intersect?

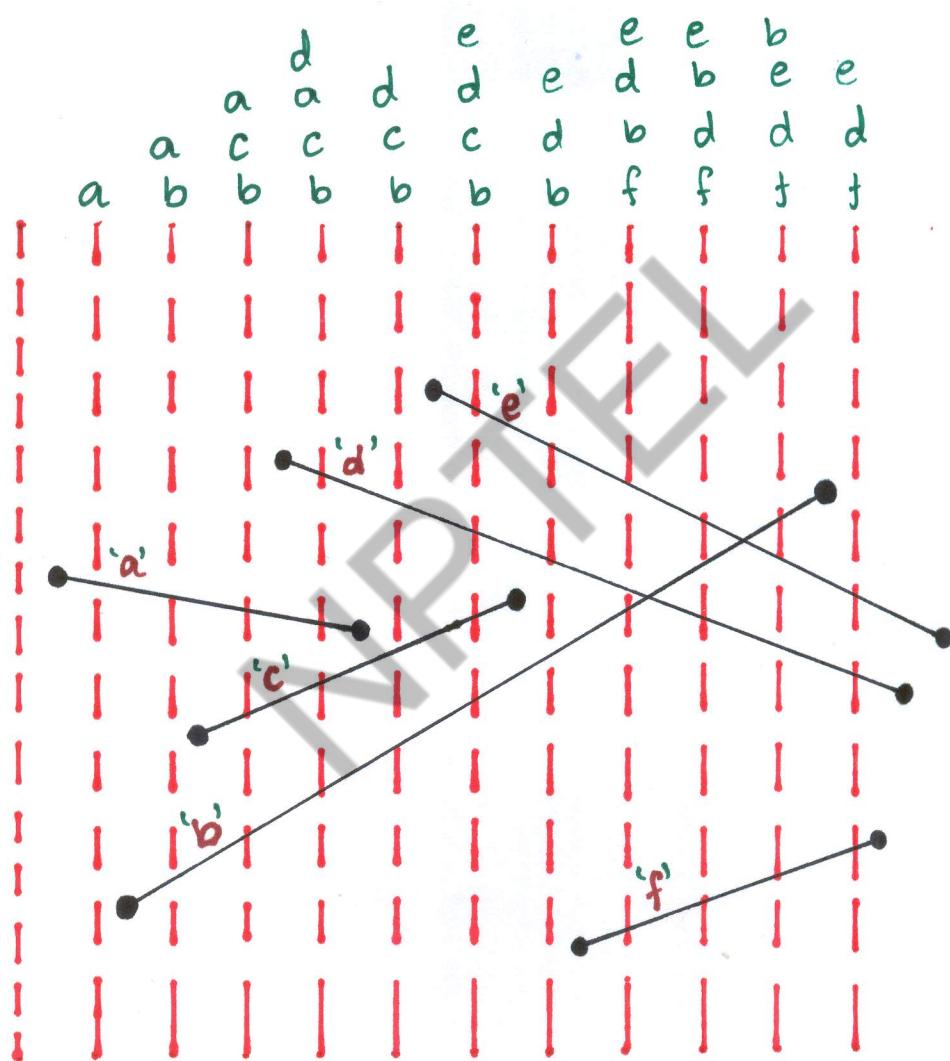
Obvious algorithm: $O(n^2)$



Sweep-line algorithm

- Sweep a vertical line from left to right
(conceptually replacing x -coordinate with time)
- Maintain dynamic set S of segments that intersect the sweep line, ordered (tentatively) by y -co-ordinate of intersection.
- Order changes when
 - new segment is encountered,
 - existing segment finishes, or
 - two segments cross
- Key event points are therefore segment endpoints.

} segment
endpoints



Sweep-line algorithm

Process event points in order by sorting segment endpoints by x -coordinates and looping through:

- For a left endpoint of segment s :
 - Add segment s to dynamic set S
 - Check for intersection between s and its neighbours in S
- For a right endpoint of segment s :
 - Remove segment s from dynamic set S
 - Check for intersection between s and the neighbours of s in S .

Analysis

Use red black tree to store dynamic set S .

Total running time: $O(n \log n)$

Correctness

Theorem: If there is an intersection, the algorithm finds it.

Proof: Let X be the leftmost intersection point.

Assume for simplicity that

- only two segments s_1, s_2 pass through X , and
- no two points have the same x -co-ordinates.

At some point before we reach X , s_1 and s_2 become consecutive in order of S .

Either initially consecutive when s_1 and s_2 inserted or became consecutive when another deleted.