

Geometric Data Structures for Range

Searching and Windowing

Balanced Binary Search Tree for 1-D Range Searching

Given: A set $P := \{p_1, \dots, p_n\} \subseteq \mathbb{R}$

Task: Process P into a data structure that supports range queries of the type: Report all points of P that lie in a query interval $[x, x']$

⇒ Balanced binary search tree T

- Leaves of T store points of P

- Inner vertices store splitting values

Inner vertex v stores x_v such that

• Left subtree of v contains all $p \in P$ with $p \leq x_v$

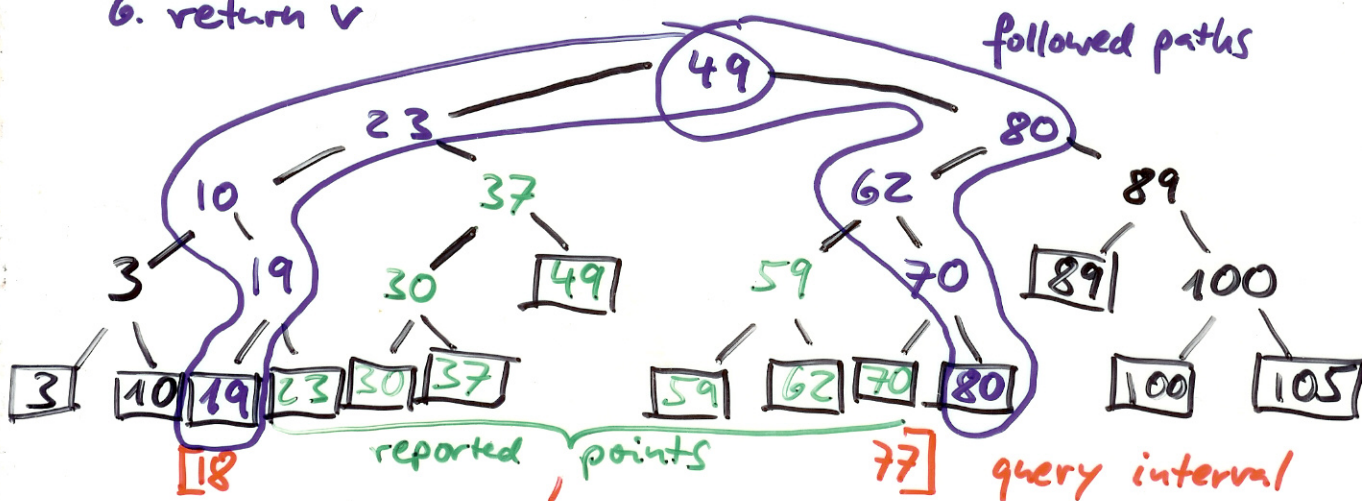
• Right subtree of v contains all $p \in P$ with $p > x_v$

Find Split Node (T, x, x'):

Input: T and $x \leq x'$

Output: The node v where the paths to x and x' split, or the leaf where both paths end

1. $v := \text{root}(T)$
2. while v is no leaf and $(x' \leq x_v \text{ or } x_v < x)$ do
3. if $x' \leq x_v$ then
4. $v := \text{left-child}(v)$
5. else $v := \text{right-child}(v)$
6. return v



Algorithm 1D-Range-Query ($\mathcal{T}, [x, x']$)

Input: A binary search tree \mathcal{T} , a range $[x, x']$

Output: All points stored in leaves of \mathcal{T} that lie in $[x, x']$

1. $v_{split} := \text{FindSplitNode}(\mathcal{T}, x, x')$
2. if v_{split} is a leaf then
3. Check if point stored at v_{split} must be reported
4. else // Follow path to x and report points in subtrees right of path
5. $v := \text{left_child}(v_{split})$
6. while v is no leaf do
7. if $x \leq x_v$ then
8. $\text{Report_Subtree}(\text{right_child}(v))$
9. $v := \text{left_child}(v)$
10. else $v := \text{right_child}(v)$
11. Check if point stored at leaf v must be reported
12. Similarly follow path to x' and report points in subtrees left of path

Analysis:


- Preprocessing:

$O(n \log n)$ construction time

$O(n)$ storage

- Query:

$O(k + \log n)$ time to report all k points in query range



$P(v) := \text{Canonical subset of } v$

$O(\log n)$ subtrees



\Rightarrow Union of canonical subsets = output

KD-Trees for 2-D Range Searching

Given: A set $P := \{p_1, \dots, p_n\} \subseteq \mathbb{R}^2$

Task: Process P into a data structure that allows fast 2D range queries: Report all points in P that lie in the query rectangle $[x, x'] \times [y, y']$

\Rightarrow Recursively split P into two sets of same size, alternatingly along a vertical or horizontal line



- Sort P by x - and by y -coordinate in advance
- Use these two sorted lists to find median
- Pass sorted lists into recursive calls

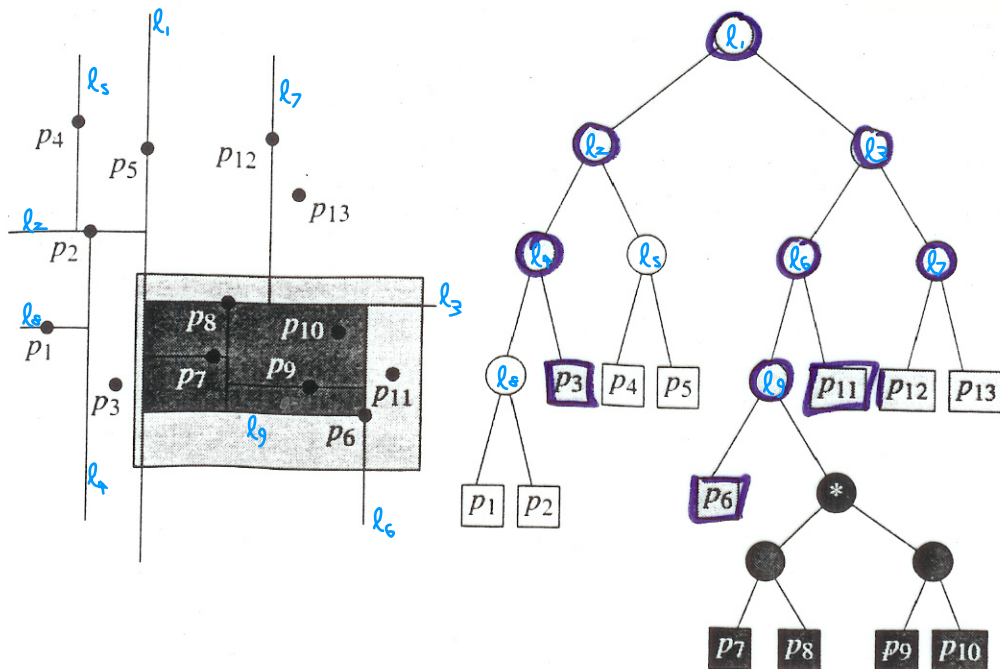
$$\Rightarrow T(n) = \begin{cases} O(1) & , n=1 \\ O(n) + 2T(\lfloor n/2 \rfloor) & , n > 1 \end{cases} = O(n \log n) \text{ time}$$

Algorithm BUILDKDTREE($P, depth$)

Input. A set of points P and the current depth $depth$.

Output. The root of a kd-tree storing P .

1. if P contains only one point
2. then return a leaf storing this point
3. else if $depth$ is even
4. then Split P into two subsets with a vertical line ℓ through the median x -coordinate of the points in P . Let P_1 be the set of points to the left of ℓ or on ℓ , and let P_2 be the set of points to the right of ℓ .
5. else Split P into two subsets with a horizontal line ℓ through the median y -coordinate of the points in P . Let P_1 be the set of points below ℓ or on ℓ , and let P_2 be the set of points above ℓ .
6. $v_{\text{left}} \leftarrow \text{BUILDKDTREE}(P_1, depth + 1)$
7. $v_{\text{right}} \leftarrow \text{BUILDKDTREE}(P_2, depth + 1)$
8. Create a node v storing ℓ , make v_{left} the left child of v , and make v_{right} the right child of v .
9. return v



$lc(v) = \text{left-child}(v)$
 $rc(v) = \text{right-child}(v)$

Algorithm SEARCHKDTREE(v, R)

Input. The root of (a subtree of) a kd-tree, and a range R .

Output. All points at leaves below v that lie in the range.

1. if v is a leaf
2. then Report the point stored at v if it lies in R .
3. else if $\text{region}(lc(v))$ is fully contained in R
4. then REPORTSUBTREE($lc(v)$)
5. else if $\text{region}(lc(v))$ intersects R
6. then SEARCHKDTREE($lc(v), R$)
7. if $\text{region}(rc(v))$ is fully contained in R
8. then REPORTSUBTREE($rc(v)$)
9. else if $\text{region}(rc(v))$ intersects R
10. then SEARCHKDTREE($rc(v), R$)

$\text{region}(lc(v))$
 $= \text{region}(v) \cap \ell(v)^{\text{left}}$
 \rightarrow compute on the fly


Theorem: A kd-tree for a set P of n points can be constructed in $O(n \log n)$ time and uses $O(n)$ space. A rectangular range query can be answered in $O(\sqrt{n} + k)$ time; $k = \#$ reported points. (Generalization to d dimensions: Also $O(n)$ storage, $O(n \log n)$ construction time, but $O(n^{1-1/d} + k)$ query time.)

Proof Sketch:

- Sum of $\#$ of visited vertices in Report Subtree = $O(k)$
- $\#$ visited vertices that are not in one of the reported subtrees = $O(\# \text{ regions}(v) \text{ intersected by a line})$

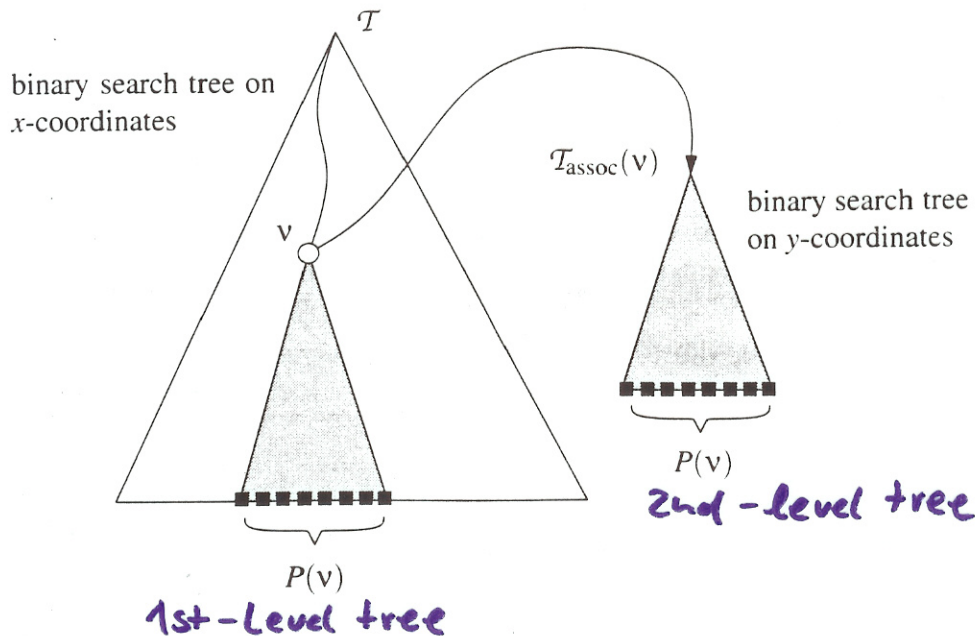
$\rightarrow Q(n) := \#$ intersected regions in kd-tree of n points whose root contains vertical splitting line

$$\rightarrow Q(n) = 2 + 2Q\left(\frac{n}{4}\right), n > 1 \Rightarrow Q(n) = O(\sqrt{n})$$



Range Trees in 2D

Task: Answer 2D rectangular range query within $O(\log^2 n + k)$ query time



Algorithm BUILD2DRANGETREE(P)

Input. A set P of points in the plane.

Output. The root of a 2-dimensional range tree.

1. Construct the associated structure: Build a binary search tree T_{assoc} on the set P_y of y-coordinates of the points in P . Store at the leaves of T_{assoc} not just the y-coordinate of the points in P_y , but the points themselves.
2. **if** P contains only one point
3. **then** Create a leaf v storing this point, and make T_{assoc} the associated structure of v .
4. **else** Split P into two subsets; one subset P_{left} contains the points with x -coordinate less than or equal to x_{mid} , the median x -coordinate, and the other subset P_{right} contains the points with x -coordinate larger than x_{mid} .
5. $v_{\text{left}} \leftarrow \text{BUILD2DRANGETREE}(P_{\text{left}})$
6. $v_{\text{right}} \leftarrow \text{BUILD2DRANGETREE}(P_{\text{right}})$
7. Create a node v storing x_{mid} , make v_{left} the left child of v , make v_{right} the right child of v , and make T_{assoc} the associated structure of v .
8. **return** v

Lemma: $O(n \log n)$ storage; $O(n \log n)$ construction time

Proof: • $O(n)$ storage at every level of $\mathcal{T} \rightarrow O(n \log n)$ altogether
• Presort P into two lists, sorted by x - or y -coordinate
 \rightarrow Construct search trees in linear time
 \rightarrow Construction time proportional to storage □

Algorithm 2DRANGEQUERY($\mathcal{T}, [x : x'] \times [y : y']$)

Input. A 2-dimensional range tree \mathcal{T} and a range $[x : x'] \times [y : y']$.

Output. All points in \mathcal{T} that lie in the range.

1. $v_{\text{split}} \leftarrow \text{FINDSPLITNODE}(\mathcal{T}, x, x')$
2. **if** v_{split} is a leaf
3. **then** Check if the point stored at v_{split} must be reported.
4. **else** (* Follow the path to x and call 1DRANGEQUERY on the subtrees right of the path. *)
5. $v \leftarrow lc(v_{\text{split}})$
6. **while** v is not a leaf
7. **do if** $x \leq x_v$
8. **then** 1DRANGEQUERY($\mathcal{T}_{\text{assoc}}(rc(v)), [y : y']$)
9. $v \leftarrow lc(v)$
10. **else** $v \leftarrow rc(v)$
11. Check if the point stored at v must be reported.
12. Similarly, follow the path from $rc(v_{\text{split}})$ to x' , call 1DRANGEQUERY with the range $[y : y']$ on the associated structures of subtrees left of the path, and check if the point stored at the leaf where the path ends must be reported.

Theorem: A range tree storing n points in the plane can be constructed in $O(n \log n)$ time and storage. A rectangular range query can be answered in $O(\log^2 n + k)$ time; $k = \#$ reported points.

Proof: Query time $= \sum_{\text{visited } v} O(\log n + k_v) = O\left(\underbrace{\sum_{\text{visited } v} \log n}_{= \log^2 n} + \underbrace{\sum_{\text{visited } v} k_v}_{= k}\right)$
Search paths in \mathcal{T} have length $\log^2 n$

Generalizing Range Trees to Higher Dimensions

- d Levels of trees; one level per coordinate
 - 1st level: balanced bin. search tree on 1st coordinate
 - 2nd Level: For each v in 1st level construct $(d-1)$ -dim. range tree for points in $P(v)$ restricted to last $(d-1)$ coordinates
 - ⋮
- Query recursively as before $\rightarrow O(\log^d n + k)$ time

Theorem: A range tree storing n points in \mathbb{R}^d can be constructed in $O(n \log^{d-1} n)$ time and storage. A rectangular range query can be answered in $O(\log^d n + k)$ time.

Proof: $T_d(n) = O(n \log n) + O(\log n) \cdot T_{d-1}(n)$ construction time
construct bal. bin. search tree #levels construction time per level
 $T_2(n) = O(n \log n) \Rightarrow T_d(n) = O(n \log^{d-1} n)$


Query time $Q_d(n) = O(\log n) + O(\log n) \cdot Q_{d-1}(n)$
1st level # subtrees
 $Q_2(n) = O(\log^2 n) \Rightarrow Q_d(n) = O(\log^d n)$

Theorem: Using fractional cascading the query time of a d -dimensional range tree can be sped up to $O(\log^{d-1} n + k)$; thus to $O(\log n + k)$ for $d=2$.

Windowing Problem

Given: A set S of n line segments in the plane
(non-intersecting)

Task: Process S into a data structure such that the following windowing query can be answered efficiently:
Report all segments in S that intersect a given query window $W := [x, x'] \times [y, y']$




Segments having at least one endpoint in W can be found by range queries in range trees

→ $O(\log n + k)$ time (with fractional cascading)

Subproblem I (for horizontal segments)

Process a set of horizontal line segments s.t. segments intersecting a vertical query segment can be reported efficiently.

→ Consider query line instead of segment



Subproblem II (1 dimensional):

Given: A set $I = \{[x_n, x_n'], \dots, [x_n, x_n']\}$ of intervals in \mathbb{R}

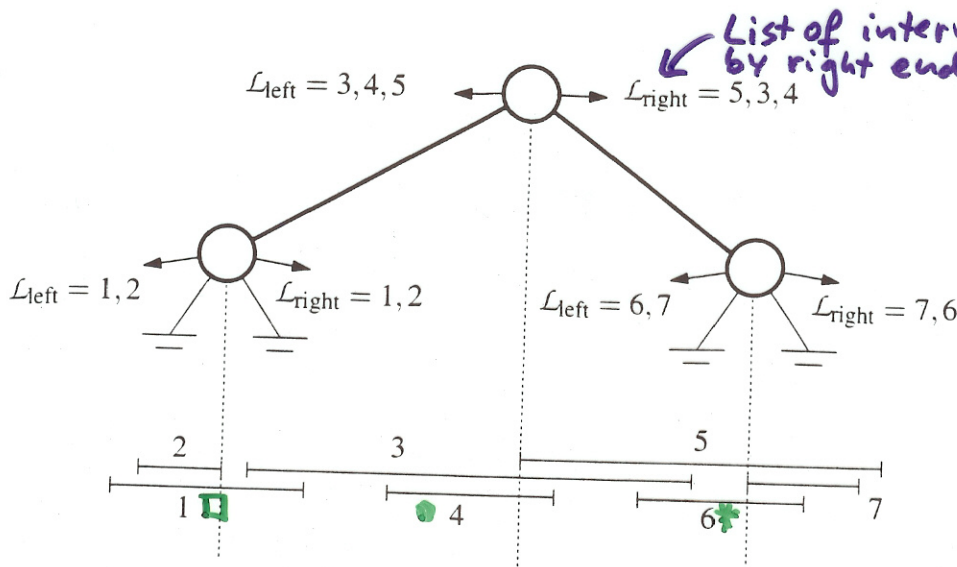

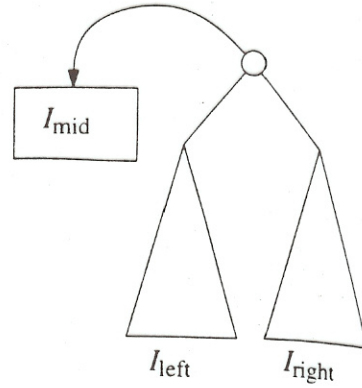
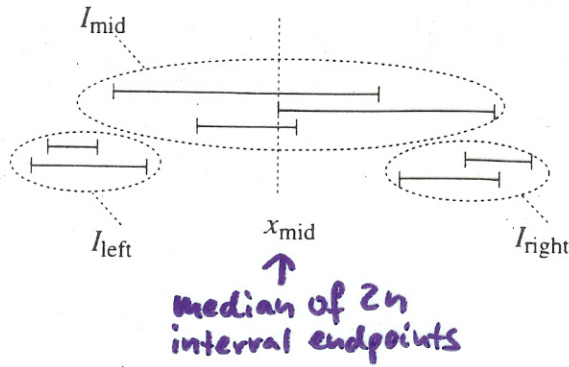
Task: Process I into a data structure which supports queries of the type: Report all intervals that contain a query point

→ Interval trees

→ Segment trees

Interval Trees

$$I = I_{\text{left}} \cup I_{\text{mid}} \cup I_{\text{right}}$$



Interval Tree

Lemma: An interval tree on a set of n intervals uses $O(n)$ storage and has depth $O(\log n)$.

Proof: Each interval is stored in a set I_{mid} only once $\rightarrow O(n)$ storage.

Algorithm CONSTRUCTINTERVALTREE(I)

Input. A set I of intervals on the real line.

Output. The root of an interval tree for I .

1. if $I = \emptyset$
2. then return an empty leaf
3. else Create a node v . Compute x_{mid} , the median of the set of interval endpoints, and store x_{mid} with v .
4. Compute I_{mid} and construct two sorted lists for I_{mid} : a list $\mathcal{L}_{\text{left}}(v)$ sorted on left endpoint and a list $\mathcal{L}_{\text{right}}(v)$ sorted on right endpoint. Store these two lists at v .
5. $lc(v) \leftarrow \text{CONSTRUCTINTERVALTREE}(I_{\text{left}})$
6. $rc(v) \leftarrow \text{CONSTRUCTINTERVALTREE}(I_{\text{right}})$
7. return v

Time analysis: $O(\sqrt{n} + |I_{\text{mid}}| \cdot \log |I_{\text{mid}}|)$ per vertex $\Rightarrow O(n \log n)$

II

Algorithm QUERYINTERVALTREE(v, q_x)

Input. The root v of an interval tree and a query point q_x .

Output. All intervals that contain q_x .

1. **if** v is not a leaf
2. **then if** $q_x < x_{\text{mid}}(v)$
3. **then** Walk along the list $\mathcal{L}_{\text{left}}(v)$, starting at the interval with the leftmost endpoint, reporting all the intervals that contain q_x . Stop as soon as an interval does not contain q_x .
4. QUERYINTERVALTREE($lc(v), q_x$)
5. **else** Walk along the list $\mathcal{L}_{\text{right}}(v)$, starting at the interval with the rightmost endpoint, reporting all the intervals that contain q_x . Stop as soon as an interval does not contain q_x .
6. QUERYINTERVALTREE($rc(v), q_x$)

Time analysis:

- $O(1+k_v)$ time at vertex v ; $k_v := \# \text{intervals reported at } v$
 - visit ≤ 1 node at any depth
- $\leadsto O(\log n + k)$

Theorem: An interval tree for a set of n intervals ~~in the plane~~ can be constructed in $O(n \log n)$ time and uses $O(n)$ storage. All intervals that contain a query point can be reported in $O(\log n + k)$ time; $k = \# \text{reported intervals}$.

Segment Trees

Let p_1, p_2, \dots, p_m the sorted list of distinct interval endpoints of I



→ Consider partitioning into elementary intervals

$(-\infty, p_1), [p_1, p_1], (p_1, p_2), [p_2, p_2], \dots, (p_{m-1}, p_m), [p_m, p_m], (p_m, \infty)$

- Balanced binary search tree T with leaves corresponding to elementary intervals
- $\text{Int}(\mu) :=$ elementary interval corresponding to leaf μ
- $\text{Int}(v) :=$ union of $\text{Int}(\mu)$ of all leaves in subtree rooted at v
- Each node or leaf v stores
 - $\text{Int}(v)$
 - the canonical subset $I(v) \subseteq I$:

$$I(v) := \{[x, x'] \in I \mid \text{Int}(v) \subseteq [x, x'] \text{ and } \text{Int}(\text{parent}(v)) \not\subseteq [x, x']\}$$

Store intervals "as high as possible".



Lemma: A segment tree on n intervals uses $O(n \log n)$ storage

Proof idea: Any interval is stored in a set $I(v)$ for at most two nodes at the same level of T .

Algorithm QUERYSEGMENTTREE(v, q_x)

Input. The root of a (subtree of a) segment tree and a query point q_x .

Output. All intervals in the tree containing q_x .

1. Report all the intervals in $I(v)$.
2. **if** v is not a leaf
3. **then if** $q_x \in \text{Int}(lc(v))$
4. **then** QUERYSEGMENTTREE($lc(v), q_x$)
5. **else** QUERYSEGMENTTREE($rc(v), q_x$)

Time analysis: • Visit one node per level } $O(\log n + k)$ time
• Spend $O(1 + kv)$ per node v

Segment-Tree-Construction:

- 1) Sort interval endpoints of $I \rightarrow$ elementary intervals } $O(n \log n)$
- 2) Construct balanced bin. search tree on elem. intervals } $O(n)$
- 3) Determine $\text{Int}(v)$ bottom-up
- 4) Compute canonical subsets by incrementally inserting the intervals $[x, x'] \in I$ into \mathcal{T} , using **Insert Segment Tree**

Algorithm INSERTSEGMENTTREE($v, [x : x']$)

Input. The root of a (subtree of a) segment tree and an interval.

Output. The interval will be stored in the subtree.

1. **if** $\text{Int}(v) \subseteq [x : x']$
2. **then** store $[x : x']$ at v
3. **else if** $\text{Int}(lc(v)) \cap [x : x'] \neq \emptyset$
4. **then** INSERTSEGMENTTREE($lc(v), [x : x']$)
5. **if** $\text{Int}(rc(v)) \cap [x : x'] \neq \emptyset$
6. **then** INSERTSEGMENTTREE($rc(v), [x : x']$)

Time analysis: • Spend constant time per node

- If we don't store $[x, x']$ at v , then $x \in \text{Int}(v)$ or $x' \in \text{Int}(v)$
- Each interval stored \leq twice at each level.

At most one node per level whose interval contains x (similar for x')
 \rightarrow Visit ≤ 4 nodes per level $\Rightarrow O(\log n) \Rightarrow O(n \log n)$ together.

Theorem: A **Segment tree** for a set of n intervals can be built in $O(n \log n)$ time and uses $O(n \log n)$ storage.

All intervals that contain a query point can be reported in $O(\log n + k)$ time.

2D Windowing Revisited


Given: A set S of n disjoint segments in the plane

Task: Process S into a data structure such that all segments intersecting a vertical query segment $q := q_x \times [q_y, q'_y]$ can be reported efficiently

- Build segment tree \mathcal{T} based on x -intervals of segments in S
 \rightarrow each $\text{Int}(v) \cong \text{Int}(v) \times (-\infty, \infty)$ vertical slab
- $I(v) \subseteq S(v)$ canonical subset of segments spanning vertical slab

Analysis: Store $S(v)$ in binary search tree $\mathcal{T}(v)$ based on vertical order of segments at v

- Storage $O(n \log n)$
- Bottom-up construction maintaining vertical order of segments
 $\rightarrow O(n \log n)$ time together



Query algorithm:

- Search regularly for q_x in \mathcal{T}
- In every visited node v report segments in $\mathcal{T}(v)$ between q_y and q'_y (1D range query)
 $\rightarrow O(\log n + kv)$ time for $\mathcal{T}(v) \rightarrow O(\log^2 n + k)$ altogether

Theorem: Let S be a set of (interior-) disjoint segments in \mathbb{R}^2 . The segments intersecting a vertical query segment (or an axis-parallel rectangular query window) can be reported in $O(\log^2 n + k)$ time, with $O(n \log n)$ preprocessing time and $O(n \log n)$ storage.