

CS 240 – Data Structures and Data Management

Module 8: Range-Searching in Dictionaries for Points

Olga Veksler

Based on lecture notes by many previous cs240 instructors

David R. Cheriton School of Computer Science, University of Waterloo

Winter 2025

Outline

8 Range-Searching in Dictionaries for Points

- Range Searches
- Multi-Dimensional Data
- Quadtrees
- kd-Trees
- Range Trees

Outline

8 Range-Searching in Dictionaries for Points

- Range Searches
- Multi-Dimensional Data
- Quadtrees
- kd-Trees
- Range Trees

Range searches

- So far: *search*(k) looks for *one* specific item.
- New operation **range-search**: look for *all* items that fall within a given range.
 - ▶ Input: A **range**, i.e., an interval $Q = (x, x')$
It may be open or closed at the ends.
 - ▶ Want: Report all KVPs in the dictionary whose key k satisfies $k \in Q$

Example:

5	10	11	17	19	33	45	51	55	59
---	----	----	----	----	----	----	----	----	----

range-search((18,45]) should return {19, 33, 45}

Range searches

- So far: *search*(k) looks for *one* specific item.
- New operation **range-search**: look for *all* items that fall within a given range.
 - ▶ Input: A **range**, i.e., an interval $Q = (x, x')$
It may be open or closed at the ends.
 - ▶ Want: Report all KVPs in the dictionary whose key k satisfies $k \in Q$

Example:

5	10	11	17	19	33	45	51	55	59
---	----	----	----	----	----	----	----	----	----

range-search((18,45]) should return {19, 33, 45}

- As usual n denotes the number of input-items.
- Let s be the **output-size**, i.e., the number of items in the range.
- We need $\Omega(s)$ time simply to report the items.
- Note that sometimes $s = 0$ and sometimes $s = n$; we therefore keep it as a separate parameter when analyzing the run-time.

Typical run-time: $O(\log n + s)$.

Outline


8 Range-Searching in Dictionaries for Points

- Range Searches
- **Multi-Dimensional Data**
- Quadtrees
- kd-Trees
- Range Trees

Multi-Dimensional Data

Range searches are of special interest for **multi-dimensional data**.

Example: flights that leave between 9am and noon, and cost \$400-\$600



- Each item has d **aspects** (coordinates): $(x_0, x_1, \dots, x_{d-1})$ so corresponds to a point in d -dimensional space
- We concentrate on $d = 2$, i.e., points in Euclidean plane
- (Orthogonal) **d -dimensional range search**: Given a **query rectangle** $Q = [x_1, x'_1] \times \dots \times [x_d, x'_d]$, find all points that lie within Q .

Multi-dimensional Range Search

The time for range searches depends on how the points are stored.

- Could store a 1-dimensional dictionary (where the key is some combination of the aspects.)
Problem: Range search on one aspect is not straightforward
- Could use one dictionary for each aspect
Problem: inefficient, wastes space
- **Better idea:** Design new data structures specifically for points.
 - ▶ Quadtrees
 - ▶ kd-trees
 - ▶ range-trees
- **Assumption:** Points are in **general position**:
 - ▶ No two points on a horizontal line.
 - ▶ No two points on a vertical line.

This simplifies presentation; data structures can be generalized.

Outline

8 Range-Searching in Dictionaries for Points

- Range Searches
- Multi-Dimensional Data
- **Quadtrees**
- kd-Trees
- Range Trees

Quadtrees

We have n points $P = \{(x_0, y_0), (x_1, y_1), \dots, (x_{n-1}, y_{n-1})\}$ in the plane.


Find a **bounding box** $R = [0, 2^k) \times [0, 2^k)$: a square containing all points.

- Assume (after translation) that all coordinates are non-negative.
- Find max-coordinate in P , use the smallest k such that it is $< 2^k$.

Structure (and also how to *build* the quadtree that stores P):



- Root r of the quadtree is associated with region R
- If R contains 0 or 1 points, then root r is a leaf that stores point.
- Else *split*: Partition R into four equal subsquares (**quadrants**)
 $R_{NE}, R_{NW}, R_{SW}, R_{SE}$
- Partition P into sets $P_{NE}, P_{NW}, P_{SW}, P_{SE}$ of points in these regions.
 - ▶ **Convention**: Points on split lines belong to right/top side
- Recursively build tree T_i for points P_i in region R_i and make them children of the root.

Quadtree example





$([0, 16] \times [0, 16])$

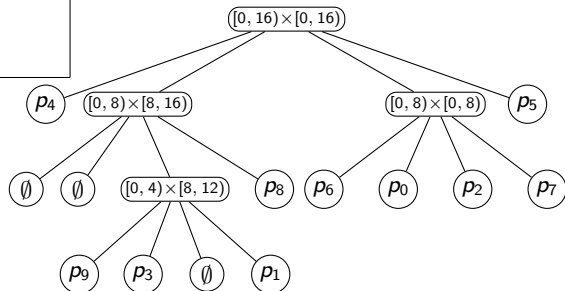

Quadtree example



Quadtree example




Quadtree example



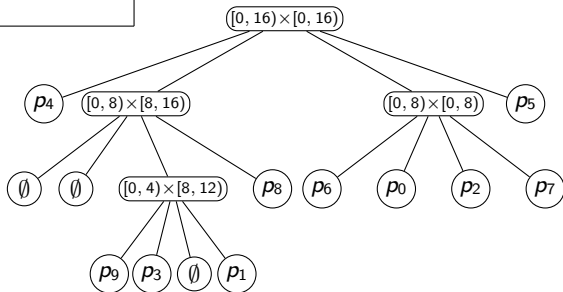
Quadtree Dictionary Operations

- *search*: Analogous to binary search trees and tries
- *insert*:
 - ▶ Search for the point
 - ▶ Split the leaf while there are two points in one region
- *delete*:
 - ▶ Search for the point
 - ▶ Remove the point
 - ▶ If its parent has only one point left: delete parent (and recursively all ancestors that have only one point left)


Quadtree Insert example



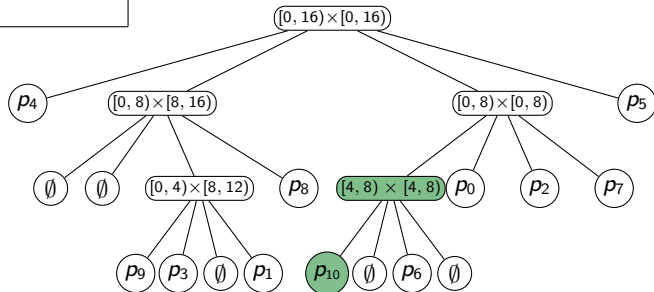
insert(p_{10})



Quadtree Insert example



insert(p_{10})



Quadtree Range Search

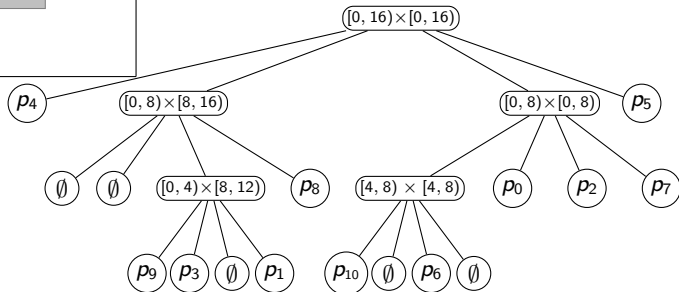

QTree::range-search($r \leftarrow \text{root}$, Q)

r : The root of a quadtree, Q : Query-rectangle


1. $R \leftarrow$ region associated with node r
2. **if** ($R \subseteq Q$) **then** // inside node, stop searching
 report all points below r and **return**
3. **else if** ($R \cap Q$ is empty) **then return** // outside node, stop searching
 // boundary node, recurse
4. **if** (r is a leaf) **then**
5. $p \leftarrow$ point stored at r
6. **if** p is not NULL and in Q **then** report it and **return**
7. **else return**
8. **for** each child v of r **do** *QTree::range-search*(v , Q)

Note: We assume here that each node of the quadtree stores the associated square. Alternatively, these could be re-computed during the search (space-time tradeoff).

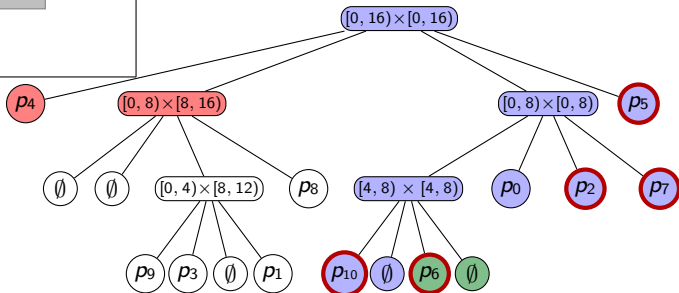
Quadtree range search example



Quadtree range search example




- Green: Search stopped due to $R \subseteq Q$.
- Red: Search stopped due to $R \cap Q = \emptyset$.
- Blue: Must continue search in children / evaluate.



Quadtree Analysis

Crucial for analysis: what is the height of a quadtree?

- Can have very large height for bad distributions of points.
- Even with $n = 3$ points, the height could be arbitrarily large.



- There exists a (weaker) bound that depends on the **spread factor** of points P :

$$\frac{\text{sidelength of } R}{\text{minimum distance between points in } P}$$

- Can show: height h of quadtree is in $\Theta(\log(\text{spread factor}))$
- Complexity to build initial tree: $\Theta(nh)$ worst-case
- Complexity of range search: $\Theta(nh)$ worst-case even if the answer is \emptyset

Quadtrees in other dimensions

- Quad-tree of 1-dimensional points:

“Points:”	0	9	12	14	24	26	28
(in base-2)	00000	01001	01100	01110	11000	11010	11100

Quadtrees in other dimensions

- Quad-tree of 1-dimensional points:


“Points:”	0	9	12	14	24	26	28
(in base-2)	00000	01001	01100	01110	11000	11010	11100

$[0,32)$

Quadtrees in other dimensions

- Quad-tree of 1-dimensional points:


“Points:”	0	9	12	14	24	26	28
(in base-2)	00000	01001	01100	01110	11000	11010	11100



Quadtrees in other dimensions

- Quad-tree of 1-dimensional points:

“Points:” 0 9 12 14 24 26 28
(in base-2) 00000 01001 01100 01110 11000 11010 11100




Same as a pruned trie

- Quadtrees also easily generalize to higher dimensions (split into octants → octrees, etc.) but are rarely used beyond dimension 3.

Quadtree summary

- Very easy to compute and handle
- No complicated arithmetic, only divisions by 2 (bit-shift!) if the width/height of bounding box R is a power of 2
- Space potentially wasteful, but good if points are well-distributed
- Variation: We could stop splitting earlier and allow up to K points in a leaf (for some fixed bound K).
- Variation: Use quad-tree to store pixelated images.



Outline

8 Range-Searching in Dictionaries for Points


- Range Searches
- Multi-Dimensional Data
- Quadtrees
- **kd-Trees**
- Range Trees

kd-trees


- We have n points $P = \{(x_0, y_0), (x_1, y_1), \dots, (x_{n-1}, y_{n-1})\}$
- Quadtrees split square into quadrants regardless of where points are
- (Point-based) kd-tree idea: Split the region at upper median of coordinates (\rightsquigarrow roughly half of the point are in each subtree)
- Each node of the kd-tree keeps track of a **splitting line** in one dimension (2D: either vertical or horizontal)
- **Convention:** Points on split lines belong to right/top side
- Continue splitting, switching between vertical and horizontal lines, until every point is in a separate region

(There are alternatives, e.g., split by the dimension that has better aspect ratios for the resulting regions.)

kd-tree example





kd-tree example




\mathbb{R}^2
$x < p_8.x?$

kd-tree example

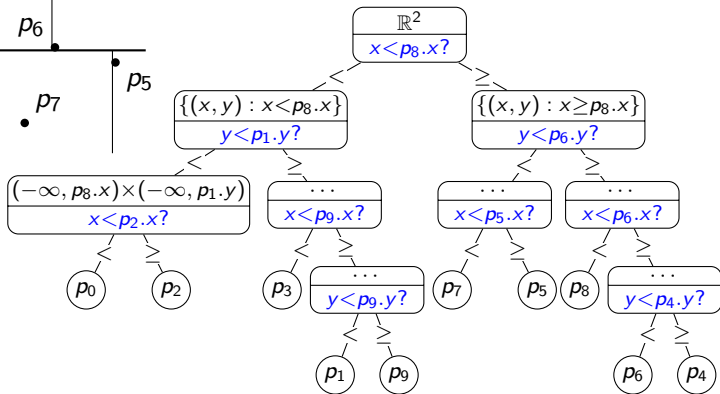



kd-tree example



For ease of drawing, we will usually not show the associated regions.

kd-tree example



For ease of drawing, we will usually not show the associated regions.

Constructing kd-trees

Build kd-tree with initial split by x on points P :

- If $|P| \leq 1$ create a leaf and return.
- Else $X := \text{randomized-quick-select}(P, \lfloor \frac{n}{2} \rfloor)$ (select by x -coordinate)
- Partition P by x -coordinate into $P_{x < X}$ and $P_{x \geq X}$
 - ▶ $\lfloor \frac{n}{2} \rfloor$ points on one side and $\lceil \frac{n}{2} \rceil$ points on the other.
(Recall: Points in general position.)
- Create left subtree recursively (splitting by y) for points $P_{x < X}$.
- Create right subtree recursively (splitting by y) for points $P_{x \geq X}$.

Building with initial y -split symmetric.

Constructing kd-trees

Run-time:

- Find X and partition P in $\Theta(n)$ expected time using *randomized-quick-select*.
- Both subtrees have $\approx n/2$ points.

$$T^{\text{exp}}(n) = 2T^{\text{exp}}(n/2) + O(n) \quad (\text{sloppy recurrence})$$

This resolves to $\Theta(n \log n)$ expected time.

- This can be reduced to $\Theta(n \log n)$ *worst-case* time by pre-sorting (no details).

Height: $h(1) = 0$, $h(n) \leq h(\lceil n/2 \rceil) + 1$.

- This resolves to $O(\log n)$ (specifically $\lceil \log n \rceil$).
- This is tight (binary tree with n leaves)

Space: All interior nodes have exactly two children.

- Therefore have $n - 1$ interior nodes.
- Space is $\Theta(n)$.

kd-tree Dictionary Operations

- *search* (for single point): as in binary search tree using indicated coordinate
- *insert*: search, insert as new leaf.
- *delete*: search, remove leaf.

Problem: After insert or delete, the split might no longer be at exact median and the height is no longer guaranteed to be $\lceil \log_2 n \rceil$.

We can maintain $O(\log n)$ height by occasionally re-building entire subtrees. (No details.) But *range-search* will be slower.

kd-trees do not handle insertion/deletion well.

kd-tree Range Search

- Range search is *exactly* as for quad-trees, except that there are only two children and leaves always store points.



```
kdTree::range-search( $r \leftarrow \text{root}$ ,  $Q$ )
```

r : The root of a kd-tree, Q : Query-rectangle

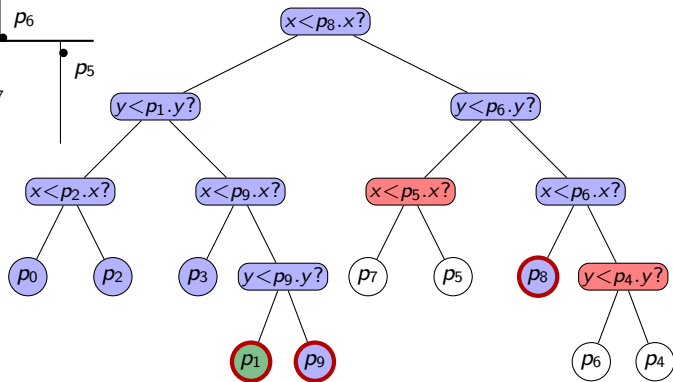

1. $R \leftarrow$ region associated with node r
2. **if** ($R \subseteq Q$) **then** report all points below r ; **return**
3. **if** ($R \cap Q$ is empty) **then return**
4. **if** (r is a leaf) **then**
5. $p \leftarrow$ point stored at r
6. **if** p is in Q **return** p
7. **else return**
8. **for** each child v of r **do** *kdTree::range-search*(v , Q)

- We assume again that each node stores its associated region.
- To save space, we could instead pass the region as a parameter and compute the region for each child using the splitting line.

kd-tree: Range Search Example



kd-tree: Range Search Example



Red: Search stopped due to $R \cap Q = \emptyset$. Green: Search stopped due to $R \subseteq Q$.

kd-tree: Range Search Complexity

- We spend $O(1)$ time at each visited node, except in line 2.
- All calls to line 2 together take $O(s)$ time (recall: s is the output-size)
- **Observe:** # visited nodes is $O(\beta(n))$
where $\beta(n)$ is the number of “boundary” nodes (blue):
 - ▶ *kdTree::range-search* was called.
 - ▶ Neither $R \subseteq Q$ nor $R \cap Q = \emptyset$
- **Can show:** $\beta(n)$ satisfies the following recurrence relation:

$$\beta(n) \leq 2\beta(n/4) + O(1)$$

- This implies $\beta(n) \in O(\sqrt{n})$
- Therefore, the complexity of range search in kd-trees is $O(s + \sqrt{n})$

kd-tree: Higher Dimensions

- kd-trees for d -dimensional space:
 - ▶ At the root the point set is partitioned based on the first coordinate
 - ▶ At the subtrees of the root the partition is based on the second coordinate
 - ▶ At depth $d - 1$ the partition is based on the last coordinate
 - ▶ At depth d we start all over again, partitioning on first coordinate
- **Storage:** $O(n)$
- **Height:** $O(\log n)$
- **Construction time:** $O(n \log n)$
- **Range search time:** $O(s + n^{1-1/d})$

This assumes that d is a constant.

Outline

8 Range-Searching in Dictionaries for Points

- Range Searches
- Multi-Dimensional Data
- Quadtrees
- kd-Trees
- **Range Trees**

Towards Range Trees

- Both Quadtrees and kd-trees are intuitive and simple.
- But: both may be very slow for range searches.
- Quadtrees are also potentially wasteful in space.

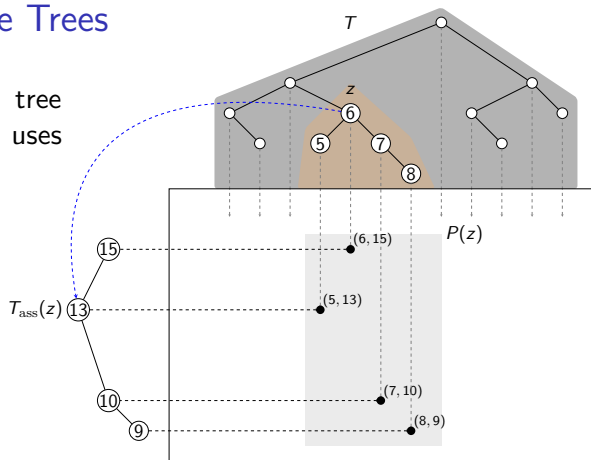
New idea: **Range trees**

- **Tree of trees** (a *multi-level* data structure)
 - ▶ So far, nodes in our trees stored a key-value pair and references to children and (maybe) the parent
 - ▶ But we can store much more in a node!
 - ▶ Here: Each node stores in another binary search tree (!)
- They are wasteful in space, but permit much faster range search.

2-dimensional Range Trees

Primary structure:


Balanced binary search tree T that stores P and uses x -coordinates as keys.




Every node z of T stores an **associate structure** $T_{\text{ass}}(z)$:

- Let $P(z)$ be all points in subtree of z in T (including point at z)
- $T_{\text{ass}}(z)$ stores $P(z)$ in a balanced binary search tree, using the y -coordinates as key
- Note: Point of z is not necessarily the root of $T_{\text{ass}}(z)$


Bigger example



Bigger example




Bigger example



Not all associate trees are shown.

Range Tree Space Analysis

- Primary tree T uses $O(n)$ space.
- How many nodes do all associate trees together have?



- ▶ point of a is only in associate tree $T_{\text{ass}}(a)$
- ▶ point of b is in associate trees $T_{\text{ass}}(a), T_{\text{ass}}(b)$
- ▶ point of c is in associate trees $T_{\text{ass}}(a), T_{\text{ass}}(b), T_{\text{ass}}(c)$
- ▶ **Key insight:** point of z is in associate tree $T_{\text{ass}}(u)$ if and only if u is an ancestor of z in T
- ▶ So every point belongs to $O(\log n)$ associate trees.
- ▶ So all associate trees together use $O(n \log n)$ space.

- A range-tree with n points uses $O(n \log n)$ space.

This is tight for some primary trees.

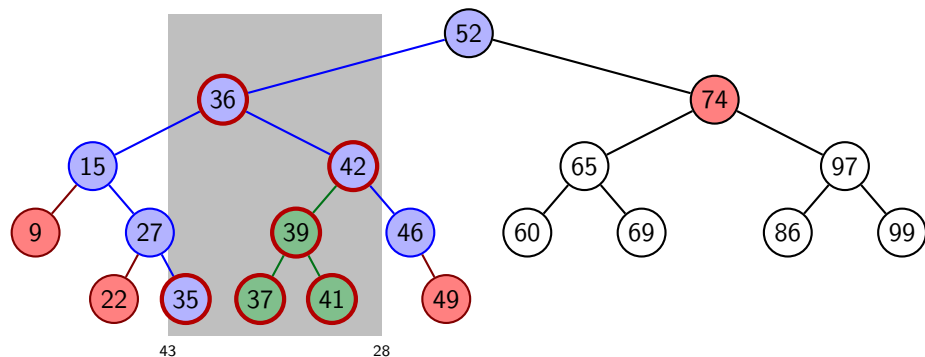
Range Trees Operations

- *search*: search by x -coordinate in T
- *insert*: First, insert point by x -coordinate into T .
Then, walk back up to the root and insert the point by y -coordinate in *all* associate trees $T_{\text{ass}}(z)$ of nodes z on path to the root.
- *delete*: analogous to insertion
- **Problem**: We want the binary search trees to be balanced.
 - ▶ This makes *insert/delete* very slow if we use AVL-trees. (A rotation at v changes $P(v)$ and hence requires a re-build of $T_{\text{ass}}(v)$.)
 - ▶ **Solution**: Completely rebuild highly unbalanced subtrees (no details)
 - ▶ Run-time for *insert/delete* becomes $O(\log^2 n)$ amortized.

Range Trees Operations


- *search*: search by x -coordinate in T
- *insert*: First, insert point by x -coordinate into T .
Then, walk back up to the root and insert the point by y -coordinate in *all* associate trees $T_{\text{ass}}(z)$ of nodes z on path to the root.
- *delete*: analogous to insertion
- **Problem**: We want the binary search trees to be balanced.
 - ▶ This makes *insert/delete* very slow if we use AVL-trees.
(A rotation at v changes $P(v)$ and hence requires a re-build of $T_{\text{ass}}(v)$.)
 - ▶ **Solution**: Completely rebuild highly unbalanced subtrees (no details)
 - ▶ Run-time for *insert/delete* becomes $O(\log^2 n)$ amortized.
- *range-search*: search by x -range in T .
Among found points, search by y -range in some associated trees.
- Must understand first: How to do (1-dimensional) range search in binary search tree?

BST Range Search



- Search for left boundary x_1 : this gives path P_1
- Search for right boundary x_2 : this gives path P_2
- Three types of nodes: outside, on, or between the paths.
- This classification will be crucial later!

BST Range Search re-phrased




- **boundary nodes:** nodes in P_1 or P_2
 - ▶ For each boundary node, test whether it is in the range.
- **outside nodes:** nodes that are left of P_1 or right of P_2
 - ▶ These are *not* in the range, we do not visit them.
- **inside nodes:** nodes that are right of P_1 and left of P_2
 - ▶ We keep a list of the topmost inside nodes.
 - ▶ All descendants of such a node are *in* the range.
For a 1d range search, report them.

BST Range Search analysis

Assume that the binary search tree is balanced:


- Search for path P_1 : $O(\log n)$
- Search for path P_2 : $O(\log n)$
- $O(\log n)$ boundary nodes
- We spend $O(1)$ time on each.



BST Range Search analysis

Assume that the binary search tree is balanced:

- Search for path P_1 : $O(\log n)$
- Search for path P_2 : $O(\log n)$
- $O(\log n)$ boundary nodes
- We spend $O(1)$ time on each.



- We spend $O(1)$ time per **topmost inside node** v .
 - ▶ They are children of boundary nodes, so this takes $O(\log n)$ time.
- For 1d range search, also report the descendants of v .
 - ▶ We have $\sum_z \text{topmost inside } \#\{\text{descendants of } z\} \leq s$ since subtrees of topmost inside nodes are disjoint. So this takes time $O(s)$ overall.


Run-time for 1d range search: $O(\log n + s)$. This is no faster overall, but topmost inside nodes will be important for 2d range search.

Range Trees: Range Search


Range search for $Q = [x_1, x_2] \times [y_1, y_2]$ is a two stage process:

- Perform a range search (on the x -coordinates) for the interval $[x_1, x_2]$ in primary tree T ($BST::range-search(T, x_1, x_2)$)
- Get **boundary** and **topmost inside** nodes as before.
- For every **boundary node**, test to see if the corresponding point is within the region Q .
- For every **topmost inside node** v :
 - ▶ Let $P(z)$ be the points in the subtree of z in T .
 - ▶ We know that all x -coordinates of points in $P(z)$ are within range.
 - ▶ Recall: $P(z)$ is stored in $T_{ass}(z)$.
 - ▶ To find points in $P(z)$ where the y -coordinates are within range as well, perform a range search in $T_{ass}(z)$: $BST::range-search(T_{ass}(z), y_1, y_2)$


Range tree range search example




Range tree range search example




Range tree range search example




Range tree range search example



Range tree range search example



Range tree range search example



Range Trees: Range Search Run-time

- $O(\log n)$ time to find boundary and topmost inside nodes in primary tree.
- There are $O(\log n)$ such nodes.
- $O(\log n + s_z)$ time for each topmost inside node z , where s_z is the number of points in $T_{\text{ass}}(z)$ that are reported
- Two topmost inside nodes have no common point in their trees
 \Rightarrow every point is reported in at most one associate structure
 $\Rightarrow \sum_z \text{topmost inside } s_z \leq s$

Time for range search in range-tree is proportional to

$$\sum_{z \text{ topmost inside}} (\log n + s_z) \in O(\log^2 n + s)$$

(There are ways to make this even faster. No details.)

Range Trees: Higher Dimensions


- Range trees can be generalized to d -dimensional space.

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

Construction time $O(n (\log n)^d)$

Range search time $O(s + (\log n)^d)$

(Note: d is considered to be a constant.)




Range Trees: Higher Dimensions

- Range trees can be generalized to d -dimensional space.

Space	$O(n(\log n)^{d-1})$	kd-trees: $O(n)$
Construction time	$O(n(\log n)^d)$	kd-trees: $O(n \log n)$
Range search time	$O(s + (\log n)^d)$	kd-trees: $O(s + n^{1-1/d})$

(Note: d is considered to be a constant.)


- Space/time trade-off compared to kd-trees.



Range search data structures summary


- Quadtrees

- ▶ simple (also for dynamic set of points)
- ▶ work well only if points evenly distributed
- ▶ wastes space for higher dimensions




- kd-trees

- ▶ linear space
- ▶ range search time $O(\sqrt{n} + s)$
- ▶ inserts/deletes destroy balance and range search time (no simple fix)



- range-trees

- ▶ range search time $O(\log^2 n + s)$
- ▶ wastes some space
- ▶ inserts/deletes destroy balance (can fix this with occasional rebuild)



Convention: Points on split lines belong to right/top side.