Computational Geometry

4. Range search

Content

4.1 Motivation & problem definition

4.2 Multi-dimensional search

4.3 kd-trees

4.4 Range trees

4.1 Motivation

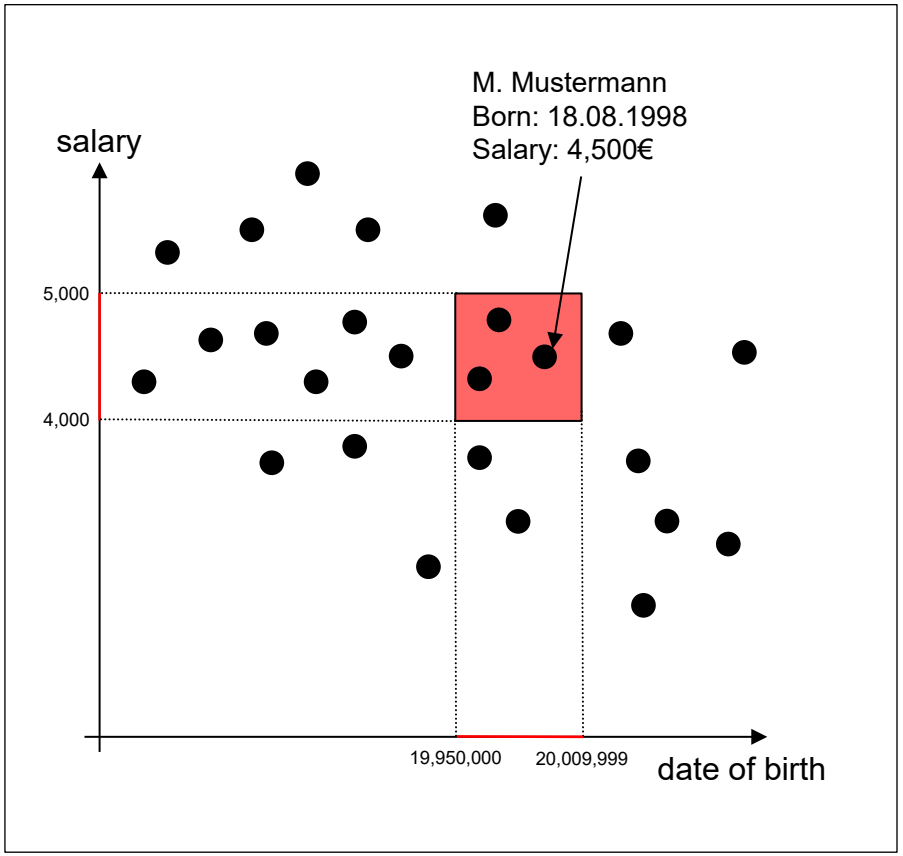
- Data base with names, dates of birth, salary of employees, number of children, etc. of a company.
- Query: Find all employees
  1. born between 1995 and 2000,
  2. earning between 4,000€ and 5,000€,
  3. having 1-3 kids.
- Solution: Compute a single number from the birthday:

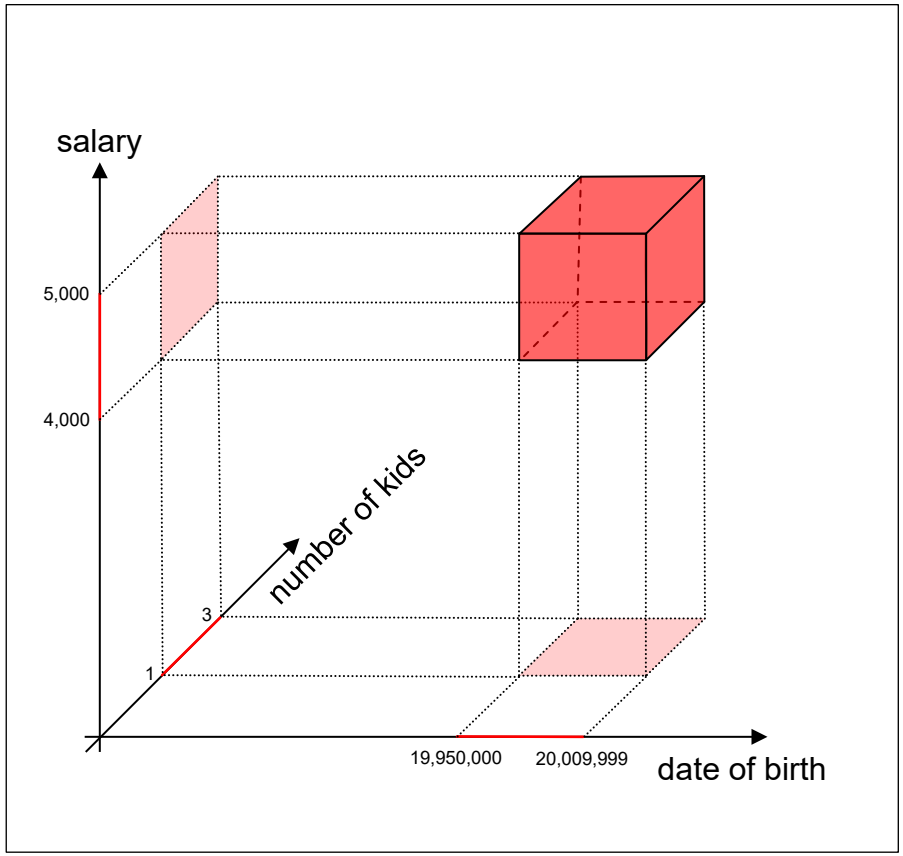
  $$10{,}000 \times \mathit{Year} + 100 \times \mathit{Month} + \mathit{Day}$$

- This yields an order for the birthdays that can be used with the income and number of kids in a *3d-range-query*.

4.1 Motivation

- Geometric interpretation





Source: de Berg et al.

4.1 Problem

More examples for different dimensions

- $k = 1$
  - Enumerate all elements in a sequence of keys between a and b.
  - ► Compare with iso-oriented line intersections.
- $k = 2$
  - Enumerate all cities in a square of 100 km edge length centered in Konstanz.
- $k \geq 3$
  - Visualize or render all objects in a k-dimensional scene inside the viewing frustum.

4.1 Problem

- Input: Set $P = \{p_1, \ldots, p_n\}$ of n points in a k-dimensional space (kd-space) and a k-dimensional *axis-aligned query range* D.
- Output: All points in $P \cap D$.
- Goal
  - Multiple queries for the same point set P with different query ranges D should be computed as fast as possible.
- Idea
  - Store P in a data structure, that supports range-queries as efficiently as possible.

4.2 Multi-dimensional search

Naïve approach

- Test all points sequentially, if they lie in the query range $D = [x_l, x_r] \times [y_l, y_r]$.
- Run time: $O(n)$ for every range search.
- Inefficient, if the search result for every range-query contains only a small, constant number of points.
- An output-size sensitive algorithm would be better, where the run time depends on the number of points in the query range.

4.2 Multi-dimensional search

Quad-tree approach (grid-method)

- Place regular grid over the set of points P.
- Consider only points in grid cells, that intersect the query range.
- Efficient, if the points are *uniformly* distributed.
- Inefficient, if the points are concentrated in one grid area and the grid has many empty grid cells.

4.2 Multi-dimensional search

BSP-tree approach

- Partition the point set P along arbitrary (hyper-)planes.
- Inefficient, because the intersection of a rectangular query range with arbitrary half-spaces is an arbitrary convex polygon (convex polytope in kd).
- Efficient, if the (hyper-)planes are axis-aligned:

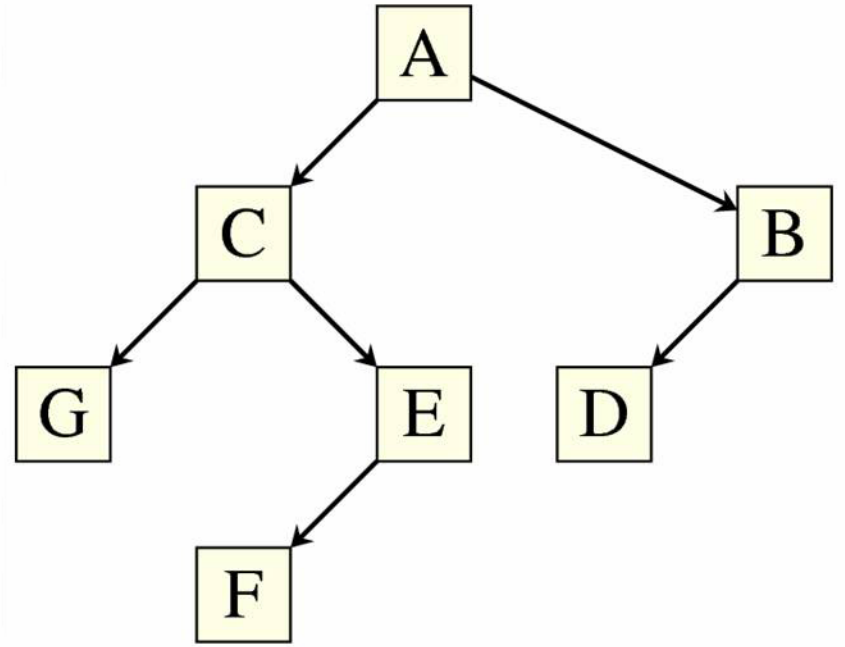
two-dimensional trees (kd-trees).

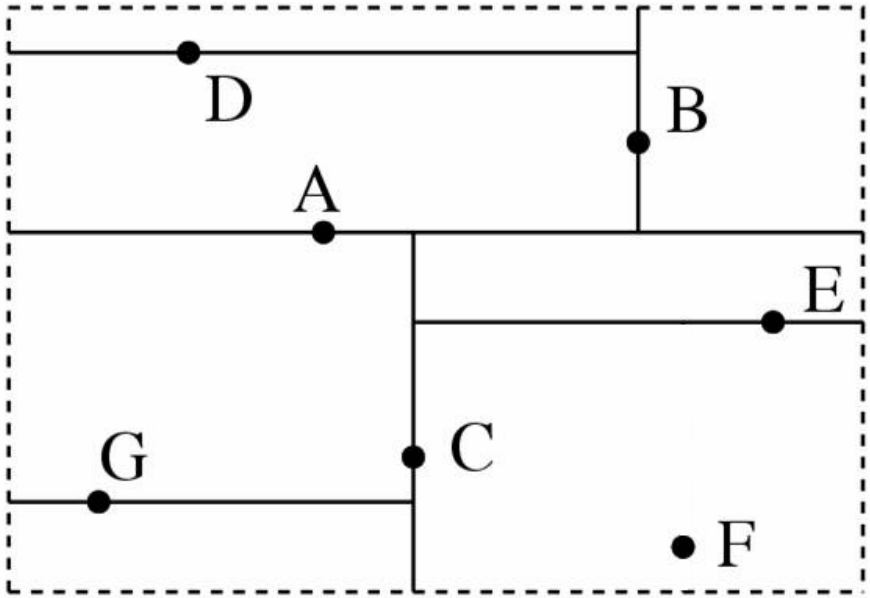
4.3 kd-trees

- In the sequel only two-dimensional range search, i.e., $k = 2$.
- Extension to higher dimensions is then straight-forward.
- Idea
  - Partition 2d search space in a way similar to a binary search tree for one-dimensional search spaces.
  - Use alternating the x- und y-coordinates as keys.
  - Construct a binary tree, representing a partition of the plane:
    - The knots correspond to the n points in 2d.
    - On an even level of the tree (level=depth+1) use the x-coordinate as key, otherwise the y-coordinate.

4.3 kd-trees

- A 2d-search tree and the resulting partition of the plane.





4.3 kd-trees

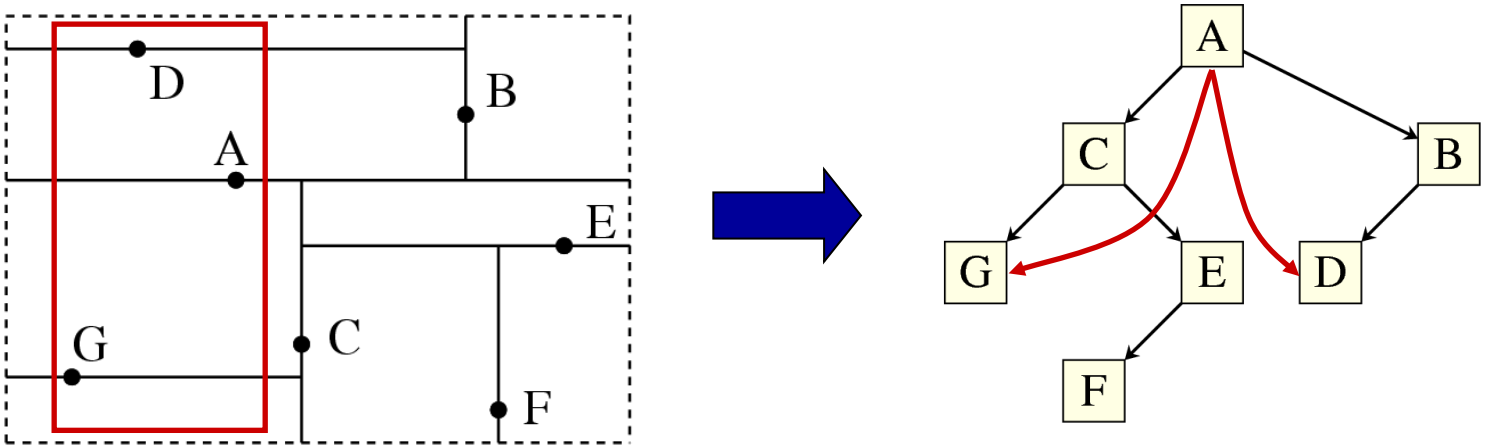
Properties

- For a knot in the tree v_x representing point p_x, whose key is its x-coordinate (even level), we have:
  - The left sub-tree of v_x contains the points left of p_x.
  - The right sub-tree of v_x contains the points right of p_x.
  - The sub-trees of children of v_x are determined by y-coordinates.
- For a knot in the tree v_y representing point p_y, whose key is its y-coordinate (odd level), we have:
  - The left sub-tree of v_y contains the points below of p_y.
  - The right sub-tree of v_y contains the points above of p_y.
  - The sub-trees of children of v_y are determined by x-coordinates.

4.3 kd-trees

Search in a 2d-tree

- The search is analogous to the usual search in a binary search tree.
- **Attention:** For every visited knot in the tree, the sub-trees of the knot are determined by the x-coordinate **or** the y-coordinate.
- **Attention:** For a range search both sub-trees must be considered.
- Example:



4.3 kd-trees

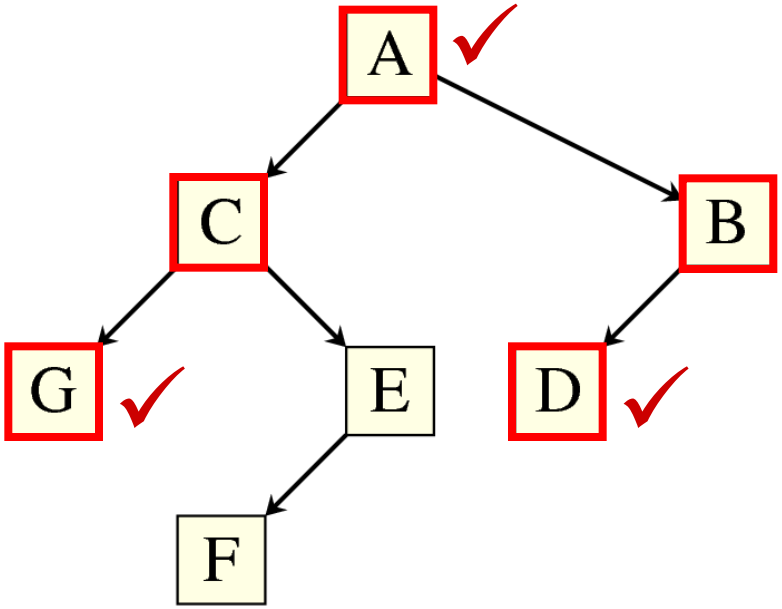
Algorithm 1: RangeSearch(knot k, direction d, query range D)

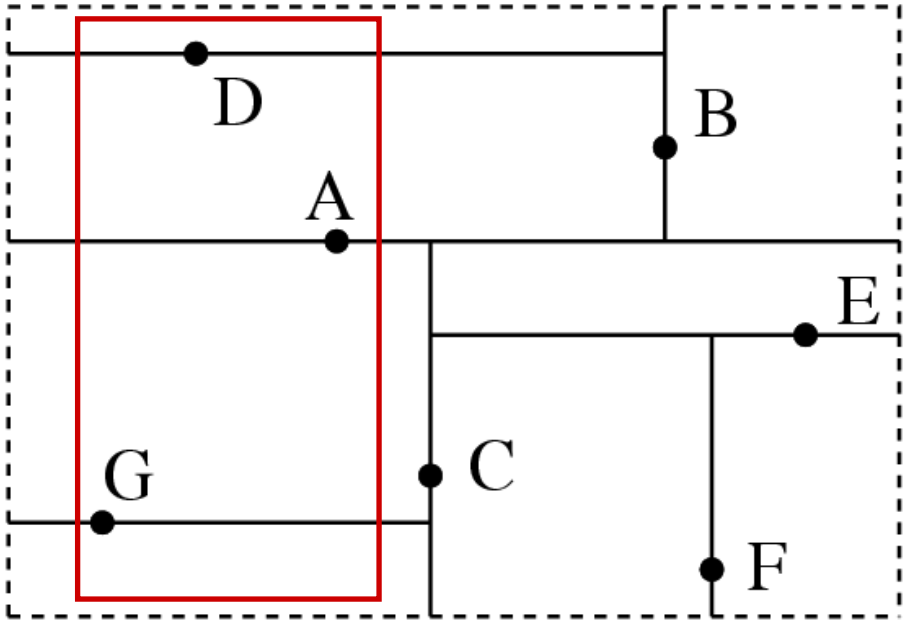
```
 1: if (k ≠ NULL) then {
 2:   if (d == vertical) then {
 3:     (l,r) = (D.y1,D.y2);
 4:     coord = k.y;
 5:     dNew  = horizontal;
 6:   } else {
 7:     (l,r) = (D.x1,D.x2);
 8:     coord = k.x;
 9:     dNew  = vertical;
10:   }
11:   if (k ∈ D    ) then Add k to the output;
12:   if (l < coord) then RangeSearch(k.left ,dNew,D);
13:   if (r > coord) then RangeSearch(k.right,dNew,D);
14: }
```

Call: `RangeSearch(root,vertical,D)`

4.3 kd-trees

- Example





4.3 kd-trees

Run time

- If the 2d-tree is balanced, the height of the tree is $O(\log n)$.
- If the 2d-tree is unbalanced, a path in the tree might degenerate to a linear list resulting in a linear run time.

- Goal: The construction of the tree must guarantee the balanced-ness.
- Solution: Partition the points at the median of the sequence of x- and y-coordinates.

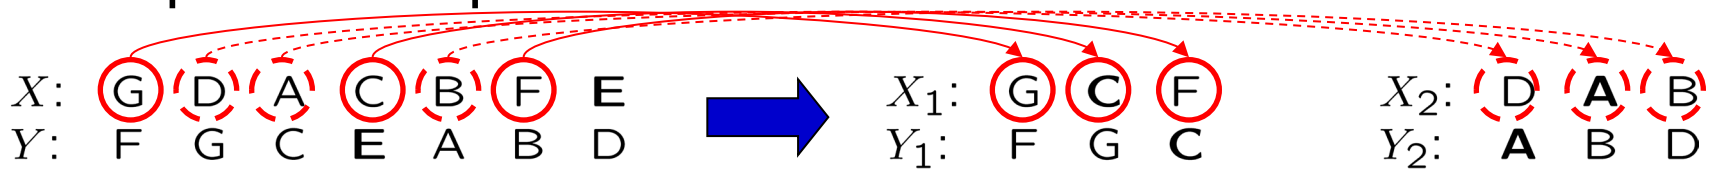
4.3 kd-trees

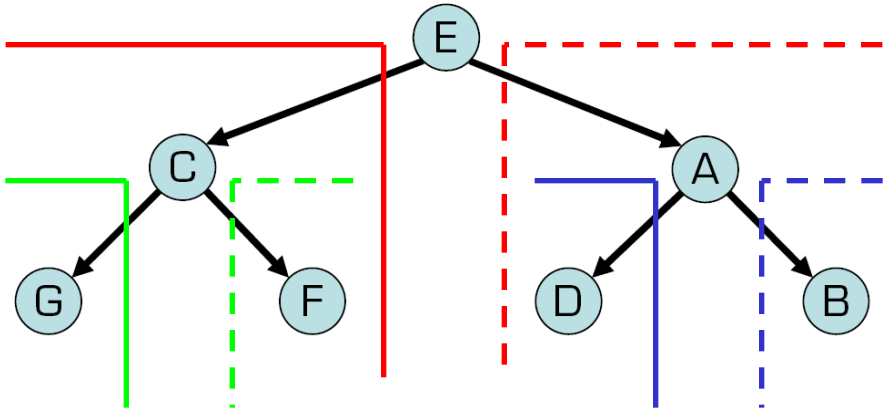
Partitioning at the median

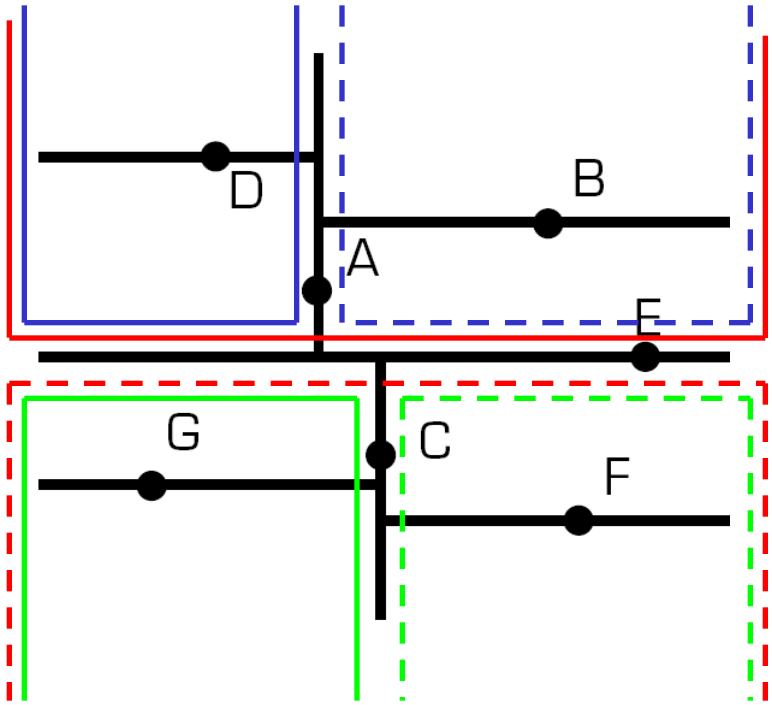
- Sort all points both by x- and y-coordinates
  - Two sorted sequences X and Y.
- Subdivide Y at its median and take it as the root of the tree
  - Two sub-sequences Y_1 and Y_2.
- Partition X into two sequences X_1 and X_2, such that X_1 contains the same points as Y_1 and X_2 the same points as Y_2.
- Partition X_1 and X_2 recursively at the median and subdivide Y_1 and Y_2 accordingly as above, until these sequences contain only one point.
  - These points are the leafs of the tree.

4.3 kd-trees

- Example for the partition







4.3 kd-trees

- Global, pre-sorted sequences X and Y.

Algorithm 2: ConstructBalanced2DTree

Input: left index l, right index r, knot k, direction d

```
 1: if (l ≤ r) then {
 2:    m = ⌈(l+r)/2⌉;
 3:    if (d == vertical) then {
 4:       k.value = Y[m];
 5:       PartitionField(X,l,r,m);
 6:    } else {
 7:       k.value = X[m];
 8:       PartitionField(Y,l,r,m);
 9:    }
10:    ConstructBalanced2dTree(l   ,m-1,k.left ,!d);
11:    ConstructBalanced2dTree(m+1,r   ,k.right,!d);
12: }
```

Call: `ConstructBalanced2DTree(1,n,root,vertical)`

4.3 kd-trees

Run time

- Construction
  - Sorting of sequences: $O(n \log n)$.
  - Partitioning of the sequences: $O(n)$.
  - Recursive calls: $T(n) = T\left(\left\lceil \frac{n-1}{2} \right\rceil\right) + T\left(\left\lfloor \frac{n-1}{2} \right\rfloor\right) + O(n)$
  - Solution of the recursion: $O(n \log n)$.

- Range search in a balanced 2d-tree
  - Run time for R points in the query range D: $O(\sqrt{n} + R)$.
  - Proof: Black board.
  - For small R this is much faster than testing all points.

4.3 kd-trees

Higher dimensions

- Use the same algorithm!
- **For the construction** of the tree cycle through the k dimensions and partition the coordinate-sorted lists one by one.
- **For the search** also cycle through the k dimensions one after the other.
- **Run time and memory** in k dimensions
  - Construction: $O(k \cdot n \log n)$
  - Range search: $O(kn^{1-\frac{1}{k}} + R)$
  - Memory: $O(n)$ (A kd-tree is a binary tree with n leaves.)

4.4 Range trees

- For large n and small R the query time of a kd-tree is relatively large.

- Principle of kd-trees:

  Do comparisons of x- and y-coordinates **in alternation**.

- Principle of range-trees:

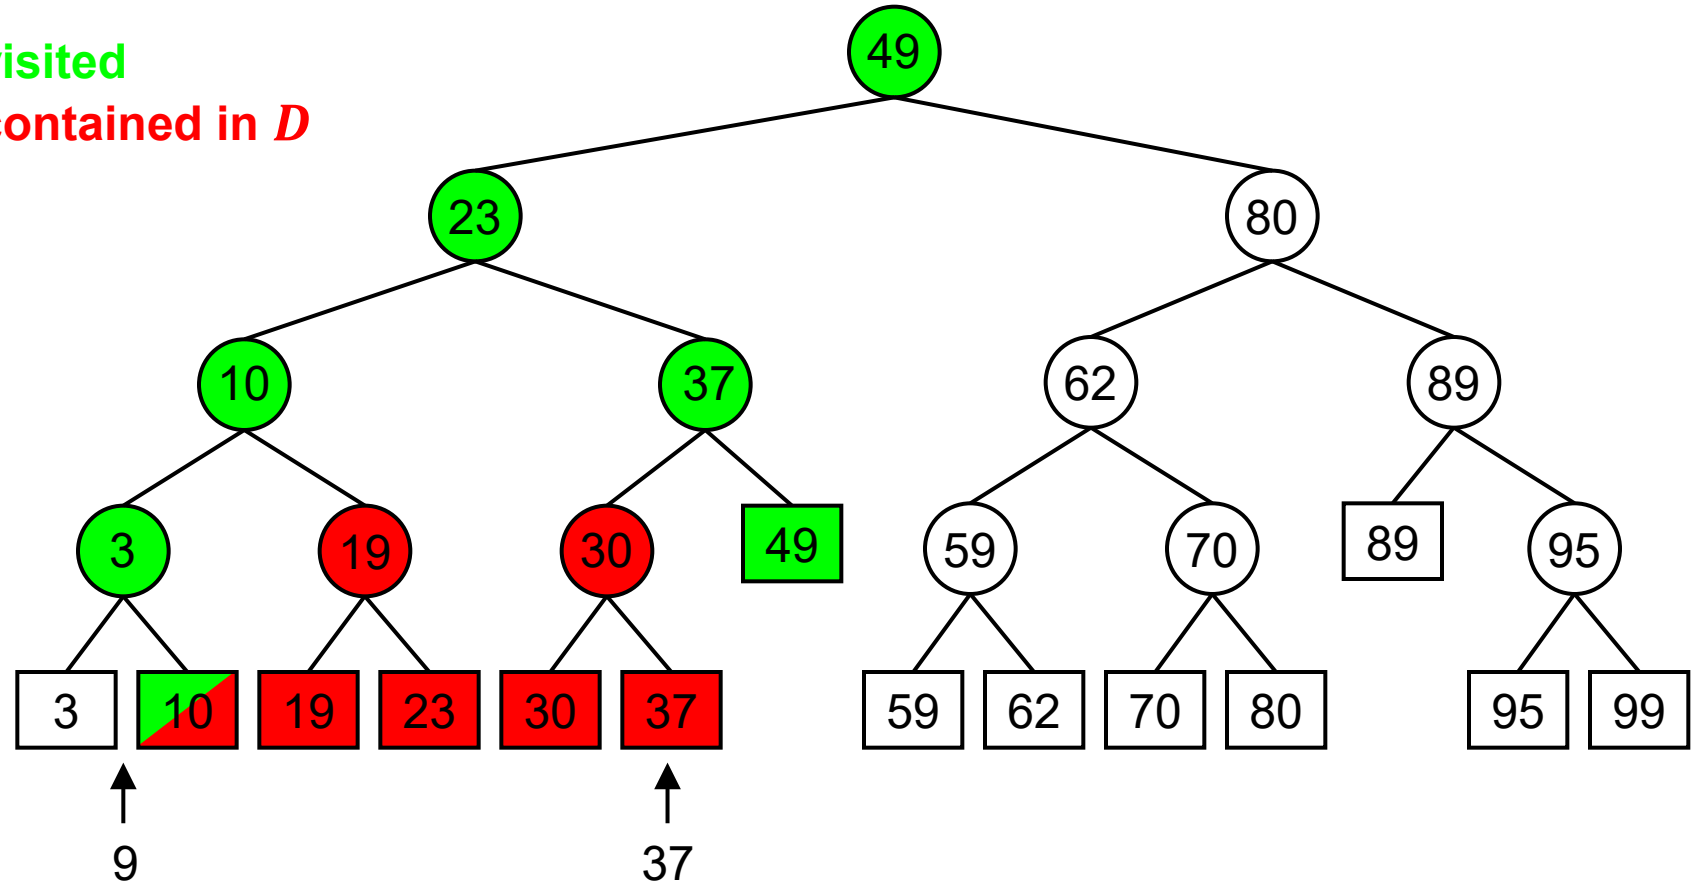
  Do comparisons of x- and y-coordinates **in succession**.

4.4 Range trees

- First the x-coordinates:
  - First level data structure: Use a balanced binary tree for the x-coordinates of the points.
  - Query: Use a one-dimensional range query to find all points with x-coordinates in $[x_l, x_r]$.
  - This yields a list of candidates that lie potentially in the query range D.
- Second the y-coordinates:
  - Second level data structure: Use a balanced binary tree for the y-coordinates of the candidates.
  - Query: Use a one-dimensional range query to find all candidates with y-coordinates in $[y_l, y_r]$.

4.4.1 1d range search

Example (1d range query): Search for $D = [9, 37]$



4.4.1 1d range search

1d-Range-Search: Search for all leaves in a balanced, binary search tree whose key lies in $[x_l, x_r]$.

Idea: Search for x_l and x_r in the tree simultaneously:

1. Find the node v_{split}, where the paths to x_l and x_r split.
2. From v_{split} follow the path to x_l.
   - At a node where this path turns left, report all leaves of the right sub-tree.
3. From v_{split} follow the path to x_r.
   - At a node where this path turns right, report all leaves of the left sub-tree.

4.4.1 1d range search

Algorithm 3: FindSplitNode(tree T, left bound x_l, right bound x_r)

Output: Node where paths to x_l and x_r split, or the leaf where both paths end.

```
1: v = Root(T);
2: while (v is not a leaf and (x_l > x_v or x_r ≤ x_v)) do {
3:   if (x_r ≤ x_v) then v = LeftChild (v);
4:   else               v = RightChild(v);
5: }
6: return v;
```

4.4.1 1d range search

Algorithm 4: 1dRangeSearch(tree T, left bound x_l, right bound x_r)

Output: All points in T that lie in the interval $[x_l, x_r]$

```
 1: v_split = FindSplitNode(T, x_l, x_r);
 2: if (v_split is a leaf) Return the point in v_split, if necessary;
 3: else {
 4:    v = LeftChild(v_split);
 5:    while (v is not a leaf) do {
 6:       if(x_l ≤ x_v) then {
 7:             ReportSubtree(RightChild(v));
 8:             v = LeftChild (v);
 9:       } else v = RightChild(v);
10:    }
11:    Return the point in v, if necessary;
12:    // Analogous for RightChild(v_split) and x_r;
13:    v = RightChild(v_split);
 ⋮           ⋮
 ⋮    }
```

4.4.1 1d range search

Proposition 1

A 1d range search can be computed with $O(n)$ memory in $O(n \log n)$ for the pre-processing and $O(R + \log n)$ for the range search, where R is the number of points, that lie in the query range.

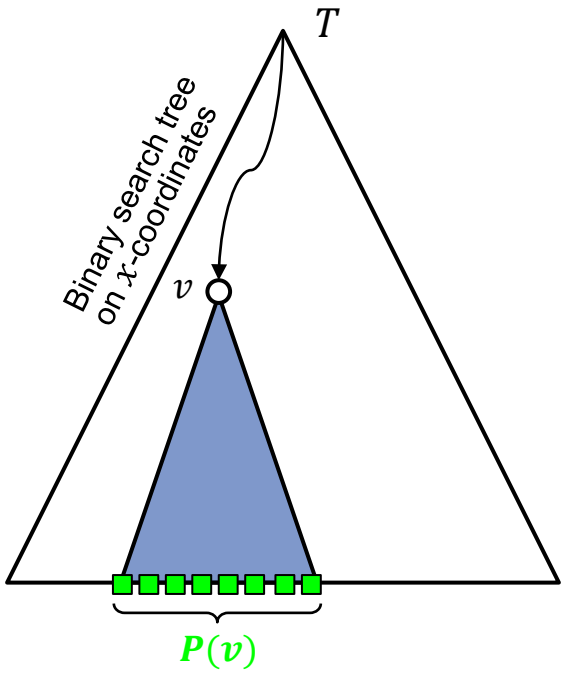
Proof:

- **Pre-processing:** Store the 1d-point sequence in a balanced, binary search tree (e.g. AVL-tree)
  - Memory: $O(n)$
  - Construction time: $O(n \log n)$
- **Range query**:
  - To compute v_{split} and to follow only two paths to x_l and x_r takes $O(\log n)$, because the time spend in each node is $O(1)$ and the tree is balanced.
  - `ReportSubtree` is linear in the reported points which is in total $O(R)$.
  - Total: $O(R + \log n)$. □

4.4 Range trees

2d Range Search: Search for all leaves in a range tree whose key lies in $[x_l, x_r] \times [y_l, y_r]$.

- First level data structure: Balanced, binary search tree T for the x-coordinates of the points.
- Every sub-tree rooted at a node v represents a subset of the points, the so-called **canonical subset $P(v)$**.

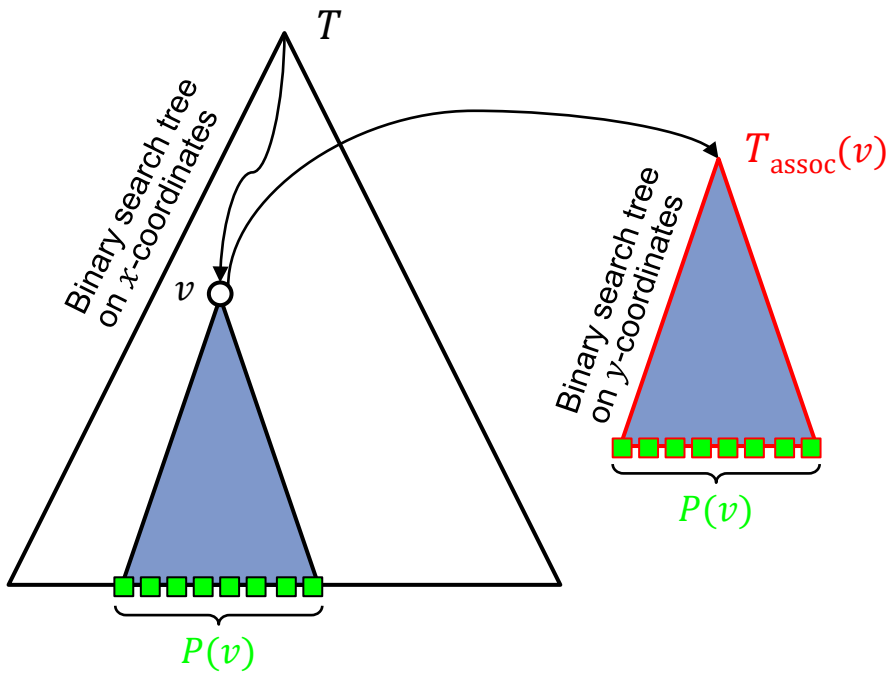


- For the range query we need only those points in $P(v)$ whose y-coordinates are in $[y_l, y_r]$.
- Second level data structure: Balanced, binary search tree for the y-coordinates of the points in the canonical subsets.

4.4 Range trees

Range tree (2d)

- Balanced, binary search tree T for the x-coordinates of the points.
- At every inner node v store a balanced, binary search tree for the y-coordinates of the points in $P(v)$.
  - This is the so-called associated structure $T_{\text{assoc}}(v)$.
  - The leaves of $T_{\text{assoc}}(v)$ hold the points.



4.4 Range trees

Proposition 2

A range tree for n 2d points uses $O(n \log n)$ memory and can be constructed in $O(n \log n)$.

Proof:

- **Memory**
  - At each level every point is stored in exactly one associated structure, which uses linear memory.
  - Thus, at every level $O(n)$ memory is used in total for all associated structures.
  - Because there are $O(\log n)$ levels, $O(n \log n)$ of memory is used.
- **Construction**
  - To build a range-tree pre-compute two lists of the points sorted by x- and y-coordinates.
  - From these two lists build the search trees bottom up in $O(n \log n)$.

4.4 Range trees

Algorithm 5: 2dRangeSearch(tree T, query range $[x_l,x_r]\times[y_l,y_r]$)

Output: All point in T that lie in the query range $[x_l,x_r]\times[y_l,y_r]$

```
 1: v_split = FindSplitNode(T,x_l,x_r);
 2: if (v_split is a leaf) Return the point in v_split, if necessary;
 3: else {
 4:    v = LeftChild(v_split);
 5:    Return the point in v, if necessary;
 6:    while (v is not a leaf) do {
 7:       if(x_l≤x_v) then {
 8:             1dRangeSearch(T_assoc(RightChild(v)),y_l,y_r);
 9:             v = LeftChild (v);
10:       } else v = RightChild(v);
11:    }
12:    // Analogous for RightChild(v_split) and x_r;
 ⋮
 ⋮  }
```

4.4 Range trees

Proposition 3

A query in a range tree of n 2d points takes $O(R + \log^2 n)$ time, where R is the number of points, that lie in the query range.

Proof:

- Every call of **`1DRangeSearch`** at a node v takes $O(R_v + \log n)$, where R_v is the number of points in $D \cap P(v)$, i.e. $\Sigma\, R_v = R$, see page 28.
- Hence, the total time spend in **`1DRangeSearch`** is

$$\Sigma_v\, O(R_v + \log n),$$

  where the summation is over all nodes v that are visited on the search paths to x_l and x_r.
- Because the search paths to x_l and x_r have length $O(\log n)$, the total run time is $O(R + \log^2 n)$.

4.4 Range trees

Higher dimensions

- Construct a k-dimensional range-tree recursively, i.e., the second level data structure is a $(k-1)$-dimensional range-tree.
- Run time and memory in k dimensions
  - Construction: $O\left(n \log^{k-1} n\right)$
  - Range search: $O\left(R + \log^{k} n\right)$
  - Memory: $O\left(n \log^{k-1} n\right)$

4.4 Range trees

Comparison of memory and runtimes

	kd-tree		Range-tree	
	2d	kd	2d	kd
Memory	$O(n)$	$O(n)$	$O(n \log n)$	$O(n \log^{k-1} n)$
Construction	$O(n \log n)$	$O(kn \log n)$	$O(n \log n)$	$O(n \log^{k-1} n)$
Range query	$O(R + \sqrt{n})$	$O(R + kn^{1-\frac{1}{k}})$	$O(R + \log^2 n)$	$O(R + \log^k n)$