




<https://algs4.cs.princeton.edu>

GEOMETRIC APPLICATIONS OF BSTs


- ▶ *1d range search*
- ▶ *line segment intersection*
- ▶ *kd trees*

Overview

This lecture. Intersections among **geometric objects**.



2d orthogonal range search







line segment intersection

Applications. CAD, games, movies, virtual reality, databases, GIS,

Efficient solutions. **Binary search trees** (and extensions).

Overview


This lecture. Only the tip of the iceberg.



**Computer Science 451
Computational Geometry**

[Bernard Chazelle](#)

Princeton University
Computer Science
Department



GEOMETRIC APPLICATIONS OF BSTs

- ▶ *1d range search*
- ▶ *line segment intersection*
- ▶ *kd trees*



<https://algs4.cs.princeton.edu>

1d range search


Extension of ordered symbol table.

- Insert key–value pair.
- Search for key k .
- Delete key k .
- **Range search:** find all keys between k_1 and k_2 .
- **Range count:** number of keys between k_1 and k_2 .

Application. Database queries.

Geometric interpretation.

- Keys are point on a **line**.
- Find/count points in a given **1d interval**.



```
insert B      B
insert D      B D
insert A      A B D
insert I      A B D I
insert H      A B D H I
insert F      A B D F H I
insert P      A B D F H I P
search G to K H I
count G to K  2
```



Suppose that the keys are stored in a sorted array. What is the running time for **range count** as a function of n and R ?

n = number of keys
 R = number of matching keys

- A. $\log R$
- B. $\log n$
- C. $\log n + R$
- D. $n + R$

1d range search: elementary implementations

Ordered array. Slow insert; fast range search.

Unordered list. Slow insert; slow range search.

order of growth of running time for 1d range search

data structure	insert	range count	range search
ordered array	n	$\log n$	$R + \log n$
unordered list	n	n	n
goal	$\log n$	$\log n$	$R + \log n$

n = number of keys


R = number of keys that match

1d range search: BST implementation

1d range search. Find all keys between l_0 and h_i .

- Recursively find all keys in left subtree (if any could fall in range).
- Check key in current node.
- Recursively find all keys in right subtree (if any could fall in range).

range search [D...Q]



Proposition. Running time proportional to $R + \log n$.

Pf. Nodes examined = search path to l_0 + search path to h_i + matches.

1d range search: summary of performance

Ordered array. Slow insert; fast range search.


Unordered list. Slow insert; slow range search.

BST. Fast insert; fast range search/count.

order of growth of running time for 1d range search

data structure	insert	range count	range search
ordered array	n	$\log n$	$R + \log n$
unordered list	n	n	n
goal	$\log n$	$\log n$	$R + \log n$

use `rank()` function
(see precept)



n = number of keys
 R = number of keys that match


INTERVAL STABBING QUERIES



Goal. Given n intervals (*left, right*), support queries of the form “how many intervals contain x ?”

Performance requirement $\log n$ per query.

Non-degeneracy assumption. Assume no two intervals contain an endpoint in common, and that no query is equal to an endpoint.



GEOMETRIC APPLICATIONS OF BSTs


- ▶ *1d range search*
- ▶ *line segment intersection*
- ▶ *kd trees*



<https://algs4.cs.princeton.edu>

Orthogonal line segment intersection

Given n horizontal and vertical line segments, find all intersections.



Quadratic algorithm. Check all pairs of line segments for intersection.


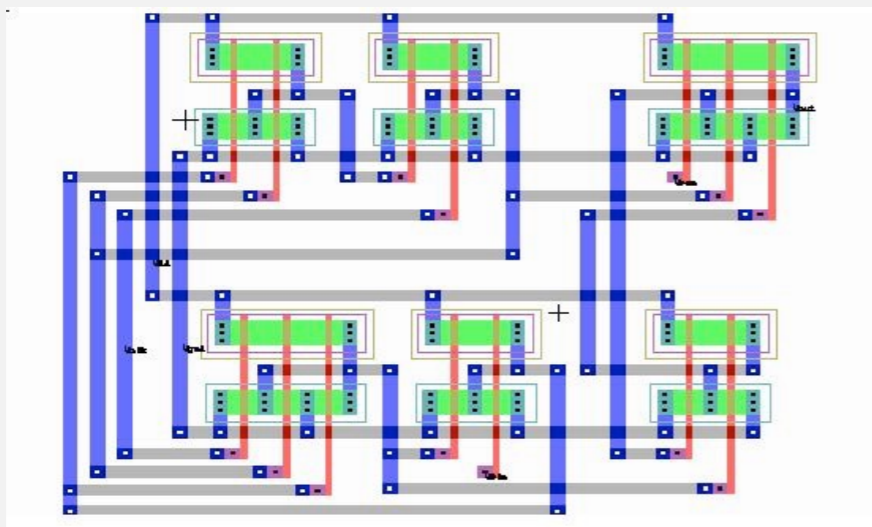
Microprocessors and geometry

Early 1970s. microprocessor design became a **geometric** problem.

- Very Large Scale Integration (VLSI).
- Computer-Aided Design (CAD).


Design-rule checking.

- Certain wires cannot intersect.
- Certain spacing needed between different types of wires.
- Debugging = line segment (or rectangle) intersection.




Algorithms and Moore's law

Moore's law (1965). Transistor count doubles every 2 years.



Gordon Moore



http://commons.wikimedia.org/wiki/File%3ATransistor_Count_and_Moore's_Law_-_2011.svg

Algorithms and Moore's law

Sustaining Moore's law.

- Problem size doubles every 2 years. ← problem size = transistor count
- Processing power doubles every 2 years. ← get to use faster computer
- How much \$ do I need to get the job done with a quadratic algorithm?

$$T_n = a n^2 \quad \text{running time today}$$


$$\begin{aligned} T_{2n} &= (a/2) (2n)^2 \\ &= 2 T_n \end{aligned} \quad \text{running time in 2 years}$$

running time	1970	1972	1974	2000
n	$\$ x$	$\$ x$	$\$ x$	$\$ x$
$n \log n$	$\$ x$	$\$ x$	$\$ x$	$\$ x$
n^2	$\$ x$	$\$ 2x$	$\$ 4x$	$\$ 2^{15}x$

Bottom line. Linearithmic algorithm is **necessary** to sustain Moore's Law.

Orthogonal line segment intersection: sweep-line algorithm


Non-degeneracy assumption. All x - and y -coordinates are distinct.




Orthogonal line segment intersection: sweep-line algorithm

Sweep vertical line from left to right.

- x -coordinates define events.
- h -segment (left endpoint): insert y -coordinate into BST.



non-degeneracy assumption: all x - and y -coordinates are distinct




y -coordinates


Orthogonal line segment intersection: sweep-line algorithm

Sweep vertical line from left to right.

- x -coordinates define events.
- h -segment (left endpoint): insert y -coordinate into BST.
- h -segment (right endpoint): remove y -coordinate from BST.



non-degeneracy assumption: all x - and y -coordinates are distinct




y -coordinates


Orthogonal line segment intersection: sweep-line algorithm

Sweep vertical line from left to right.

- x -coordinates define events.
- h -segment (left endpoint): insert y -coordinate into BST.
- h -segment (right endpoint): remove y -coordinate from BST.
- v -segment: range search for interval of y -endpoints.



non-degeneracy assumption: all x - and y -coordinates are distinct



y -coordinates

Orthogonal line segment intersection: sweep-line analysis

Proposition. The sweep-line algorithm takes time proportional to $n \log n + R$ to find all R intersections among n orthogonal line segments.

Pf.

- Put x -coordinates on a PQ (or sort). $\longleftarrow n \log n$
- Insert y -coordinates into BST. $\longleftarrow n \log n$
- Delete y -coordinates from BST. $\longleftarrow n \log n$
- Range searches in BST. $\longleftarrow n \log n + R$

Bottom line. Sweep line reduces 2d orthogonal line segment intersection search to 1d range search.

Sweep-line algorithm: context

The **sweep-line algorithm** is a key technique in computational geometry.




Geometric intersection.

- General line segment intersection.
- Axis-aligned rectangle intersection.
- ...



More problems.

- Andrew's algorithm for convex hull.
- Fortune's algorithm Voronoi diagram.
- Scanline algorithm for rendering computer graphics.
- ...





<https://algs4.cs.princeton.edu>

GEOMETRIC APPLICATIONS OF BSTs

- ▶ *1d range search*
- ▶ *line segment intersection*
- ▶ *kd trees*

2-d orthogonal range search

Extension of ordered symbol-table to 2d keys.


- Insert a 2d key.
- Search for a 2d key.
- **Range search:** find all keys that lie in a 2d range.
- **Range count:** number of keys that lie in a 2d range.

Applications. Networking, circuit design, databases, ...

Geometric interpretation.

- Keys are point in the **plane**.
- Find/count points in a given **$h-v$ rectangle**

↑
rectangle is axis-aligned



Space-partitioning trees


Use a **tree** to represent a recursive subdivision of 2d space.

Grid. Divide space uniformly into squares.


Quadtree. Recursively divide space into four quadrants.

2d tree. Recursively divide space into two halfplanes.


BSP tree. Recursively divide space into two regions.




Grid



Quadtree



2d tree




BSP tree


Space-partitioning trees: applications

Applications.


- Ray tracing.
- **2d range search.**
- Flight simulators.
- N-body simulation.
- Collision detection.
- Astronomical databases.
- **Nearest neighbor search.**
- Adaptive mesh generation.
- Accelerate rendering in Doom.
- Hidden surface removal and shadow casting.




Grid



Quadtree



2d tree


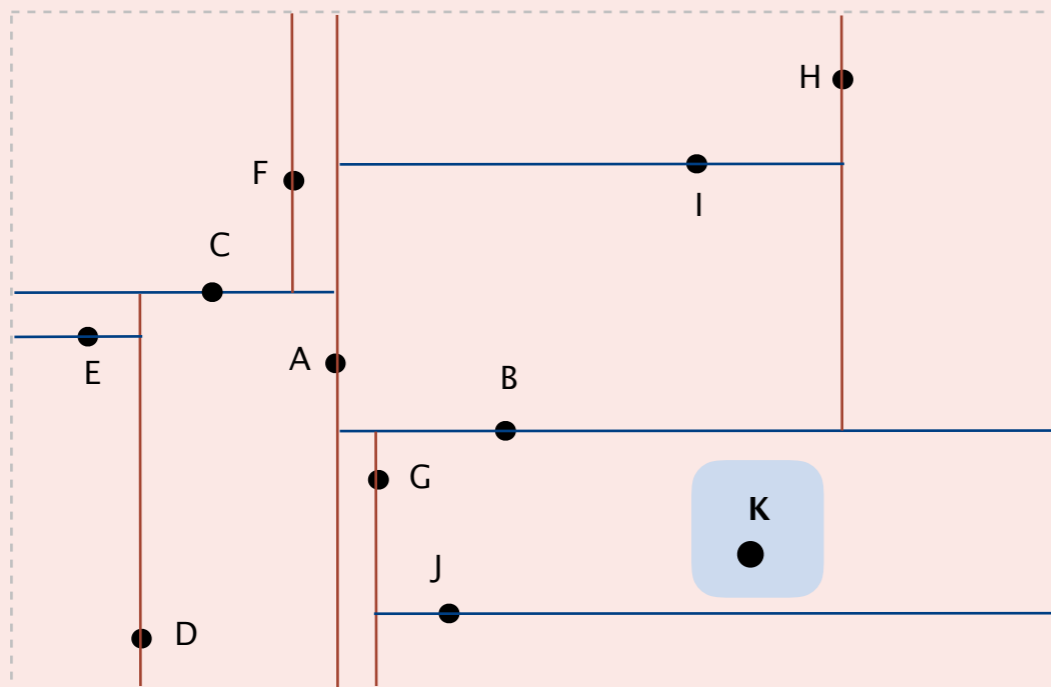


BSP tree



Where would point K be inserted in the 2d tree below?





- A. Left child of G.
- B. Left child of J.
- C. Right child of J.
- D. Right child of I.



2d tree implementation

Data structure. BST, but alternate using x - and y -coordinates as key.



- Search gives rectangle containing point.
- Insert further subdivides the plane.



2d tree demo: range search

Goal. Find all points in a query axis-aligned rectangle.


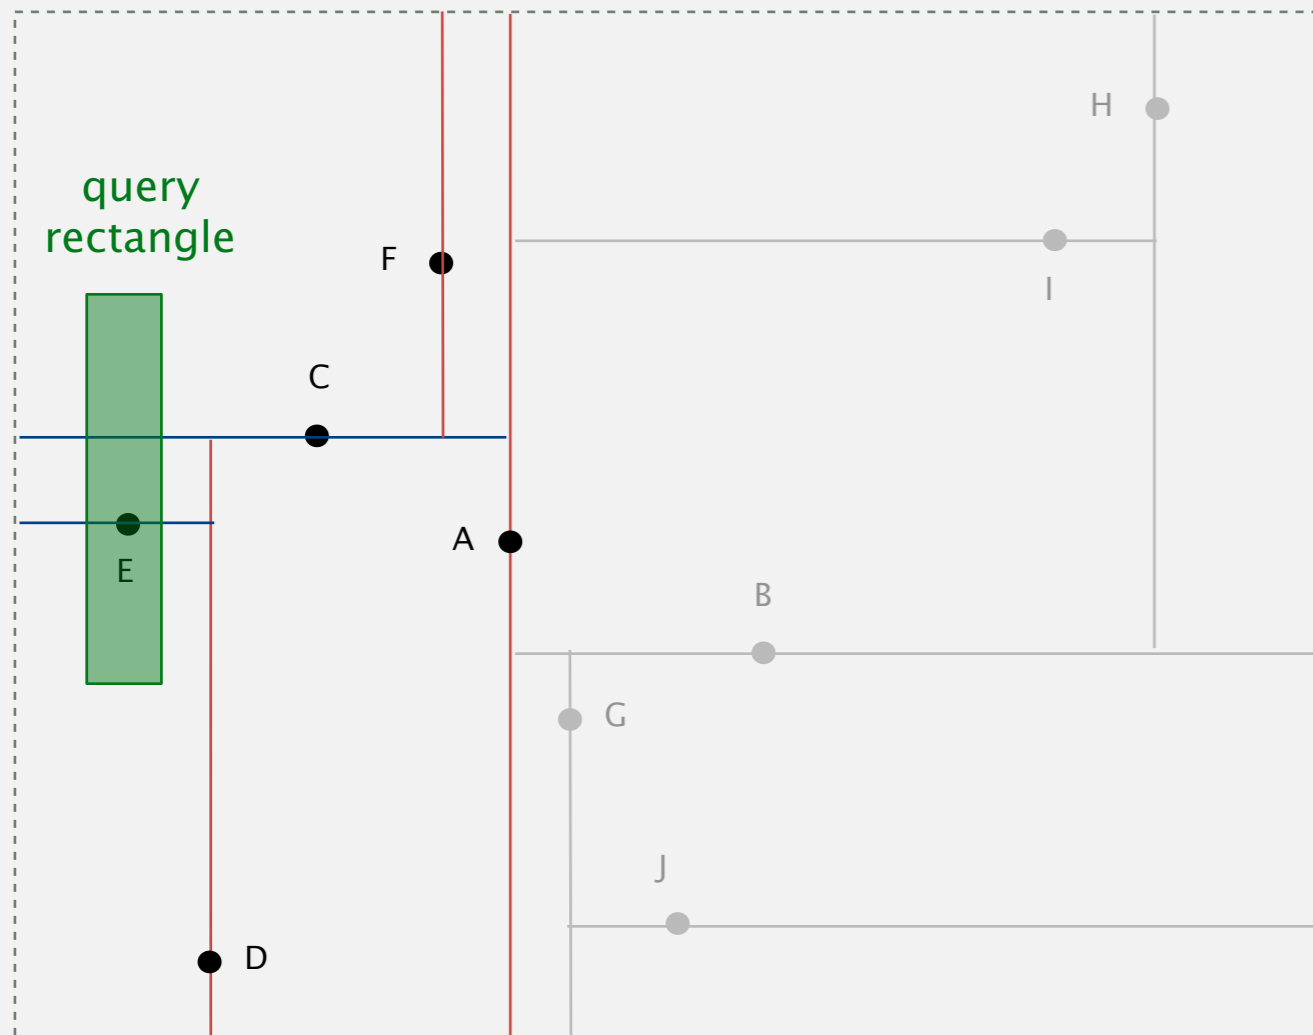
- Check if point in node lies in given rectangle.
- Recursively search left/bottom (if any could fall in rectangle).
- Recursively search right/top (if any could fall in rectangle).



2d tree demo: range search

Goal. Find all points in a query axis-aligned rectangle.

- Check if point in node lies in given rectangle.
- Recursively search left/bottom (if any could fall in rectangle).
- Recursively search right/top (if any could fall in rectangle).







Suppose we explore the right/top subtree before the left/bottom subtree in range search. What effect would it have on typical inputs?

- A.** Returns wrong answer.
- B.** Explores more nodes.
- C.** Both A and B.
- D.** Neither A nor B.

Range search in a 2d tree analysis



Typical case. $R + \log n$.

Worst case (assuming tree is balanced). $R + \sqrt{n}$.





2d tree demo: nearest neighbor

Goal. Find closest point to query point.



2d tree demo: nearest neighbor

- Check distance from point in node to query point.
- Recursively search left/bottom (if it could contain a closer point).
- Recursively search right/top (if it could contain a closer point).
- Organize method so that it begins by searching for query point.



nearest neighbor = E



Suppose we always explore the left/bottom subtree before the right/top subtree in nearest-neighbor search. What effect will it have on typical inputs?

- A.** Returns wrong answer.
- B.** Explores more nodes.
- C.** Both A and B.
- D.** Neither A nor B.




Which of the following is the worst case for nearest-neighbor search?


A.




C.



B.





D.



Nearest neighbor search in a 2d tree analysis

Typical case. $\log n$.

Worst case (even if tree is balanced). n .



nearest neighbor = E

Flocking birds

Q. Which “natural algorithm” do starlings, migrating geese, starlings, cranes, bait balls of fish, and flashing fireflies use to flock?



<http://www.youtube.com/watch?v=XH-groCeKbE>

Flocking boids [Craig Reynolds, 1986]

Boids. Three simple rules lead to complex emergent flocking behavior:



- Collision avoidance: point away from **k-nearest** boids.
- Flock centering: point towards the center of mass of **k-nearest** boids.
- Velocity matching: update velocity to the average of **k-nearest** boids.



Kd tree

Kd tree. Recursively partition k -dimensional space into 2 halfspaces.

Implementation. BST, but cycle through dimensions ala 2d trees.



Efficient, simple data structure for processing k -dimensional data.

- Widely used.
- Adapts well to high-dimensional and clustered data.
- Discovered by an undergrad in an algorithms class!




Jon Bentley

N-body simulation

Goal. Simulate the motion of n particles, mutually affected by gravity.

Brute force. For each pair of particles, compute force: $F = \frac{G m_1 m_2}{r^2}$

Running time. Time per step is n^2 .



http://www.youtube.com/watch?v=ua7YIN4eL_w

Appel's algorithm for n-body simulation

Key idea. Suppose particle is far, far away from cluster of particles.

- Treat cluster of particles as a single aggregate particle.
- Compute force between particle and **center of mass** of aggregate.



Appel's algorithm for n-body simulation

- Build 3d-tree with n particles as nodes.
- Store center-of-mass of subtree in each node.
- To compute total force acting on a particle, traverse tree, but stop as soon as distance from particle to subdivision is sufficiently large.

SIAM J. SCI. STAT. COMPUT.
Vol. 6, No. 1, January 1985

© 1985 Society for Industrial and Applied Mathematics
008


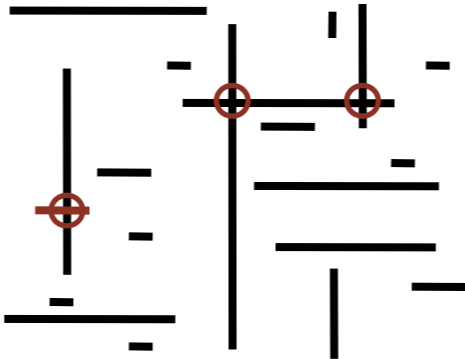
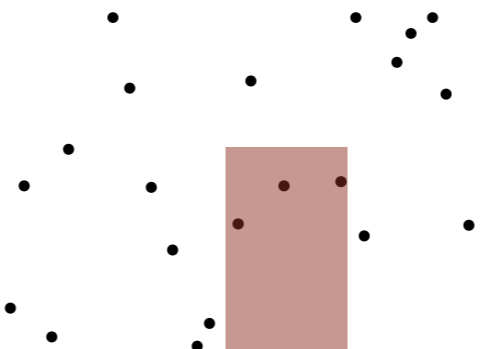
AN EFFICIENT PROGRAM FOR MANY-BODY SIMULATION*

ANDREW W. APPEL†

Abstract. The simulation of N particles interacting in a gravitational force field is useful in astrophysics, but such simulations become costly for large N . Representing the universe as a tree structure with the particles at the leaves and internal nodes labeled with the centers of mass of their descendants allows several simultaneous attacks on the computation time required by the problem. These approaches range from algorithmic changes (replacing an $O(N^2)$ algorithm with an algorithm whose time-complexity is believed to be $O(N \log N)$) to data structure modifications, code-tuning, and hardware modifications. The changes reduced the running time of a large problem ($N = 10,000$) by a factor of four hundred. This paper describes both the particular program and the methodology underlying such speedups.

Impact. Running time per step is $n \log n \Rightarrow$ enables new research.

Geometric applications of BSTs

problem	example	solution
1d range search		<i>binary search tree</i>
2d orthogonal line segment intersection		<i>sweep line reduces problem to 1d range search</i>
2d range search kd range search		<i>2d tree kd tree</i>