

Exercises and hand-in

Advanced database technology

April 15, 2004

Today's exercises consider three of the multi-dimensional index structures covered in G UW:

- The grid file.
- The kd-tree.
- The R-tree.

These structures are often used in practice as they are simple and typically work well. The aim of these exercises is to investigate in what situations (i.e., for what kind of data) these index structures behave poorly. For any set of data there are many possible choices for the index (which grid lines to use for grid files, and which rectangles to use for R-trees). We seek data that is bad *no matter* how these choices are made. To simplify, we will consider the 2-dimensional case, and static data structures (no updates).

1. First we consider nearest neighbor queries for a point set that is distributed *almost* exactly along a circle.
 - How will each of the index structures represent such a set? How much internal memory is needed?
 - How will they cope with a query for the nearest neighbor of a point almost exactly at the center of the circle?¹
2. We then consider range queries on a point set that is distributed almost exactly along a horizontal line.
 - How will each of the index structures represent such a set? How much internal memory is needed?
 - How will they cope with a query for the points in a rectangle that covers all of the points with respect to x-coordinate, but is just above the horizontal line and contains just a tiny fraction of the points?
 - What is the best (worst case) search time we could hope for?²

The following exercise is to be handed in at the latest April 29 at 10.00 AM.

Problem 3 from the ADBT exam, June 2003. (See News section on home page.) Note: You will probably need knowledge of the April 22 lecture for solving this problem in a good way. Hand in before the lecture on that day, or earlier in our mail boxes in the information office.

¹A much better solution in this case is to use the planar point location data structure you saw in the lecture. This is left as a home exercise.

²ITU has some of the world's leading researchers in 2-dimensional range searching. They have devised very good worst-case efficient data structures for internal memory.