A Plane-Sweep Algorithm for the All-Nearest-Neighbors Problem for a Set of Convex Planar Objects

Thorsten Graf and Klaus Hinrichs

Institut für numerische und instrumentelle Mathematik
-INFORMATIK-
Westfälische Wilhelms-Universität Münster
Einsteinstr. 62
D-48149 Münster

Abstract. We present a plane-sweep algorithm that solves the all - nearest - neighbors problem with respect to an arbitrary Minkowski-metric d_t ($1 \leq t \leq \infty$) for a set of non-intersecting planar compact convex objects, such as points, line segments, circular arcs and convex polygons. The algorithm also applies if we replace the condition of disjointness by the weaker condition that the objects in the configuration are *diagonal-disjoint*. For configurations of points, line segments or disks the algorithm runs in asymptotically optimal time $O(n \log n)$. For a configuration of n convex polygons with a total of N edges it finds nearest neighbors with respect to the Euclidean L^2-metric in time $O(n \log N)$ if each polygon is given by its vertices in cyclic order.

1 Introduction

We consider the two-dimensional *all-nearest-neighbors problem* (ANN problem) for a set of non-intersecting convex objects with respect to arbitrary L^t-metrics d_t ($1 \leq t \leq \infty$):

> Given a set S of n non-intersecting compact convex objects in the plane, find a nearest neighbor of each with respect to d_t.

The Minkowski-distance d_t of two objects $s_1, s_2 \in S$ is defined as follows:

$$d_t(s_1, s_2) := \min\{(|p.x - q.x|^t + |p.y - q.y|^t)^{\frac{1}{t}} : p \in s_1, q \in s_2\}, \quad 1 \leq t < \infty$$
$$d_\infty(s_1, s_2) := \min\{\max\{|p.x - q.x|, |p.y - q.y|\} : p \in s_1, q \in s_2\}$$

If S consists of n points $\Omega(n \log n)$ is a well known tight lower bound for this problem in the algebraic decision tree model of computation. Optimal algorithms for this simple version of the problem are given in [9, 11, 12]; a survey is given in [8]. However, all these algorithms do not apply to sets of more complex objects, e.g. convex polygons.

For sets of line segments and circular arcs in the plane a nearest neighbor with respect to the Euclidean metric can be found in optimal time $O(n \log n)$

by first computing the Voronoi diagram ([7, 13]) and then extracting nearest neighbors from the Voronoi diagram in linear time. For a set of n polygons with a total of N edges this technique costs $\Theta(N \log N)$ time which is worse than the $O(n \log N)$ time complexity of the algorithm presented in this paper. Furthermore, the Voronoi diagram approach cannot be applied to solve the all-nearest-neighbors problem for general convex objects with respect to an arbitrary Minkowski metric.

[2] presents a plane-sweep algorithm to solve the closest-pair problem for a set of convex planar objects with respect to an arbitrary Minkowski-metric. This algorithm tests whether a new object encountered during the left-to-right sweep forms a new closest pair with any of those objects seen so far. Among these objects only those intersecting the δ-slice to the left of the sweep-line have to be considered, where δ denotes the minimal object distance found so far. A total ordering on these *active* objects is determined by their intersections with the δ-slice. During the sweep the algorithm maintains the set of active objects and the minimal distance δ between any pair of active objects which become neighbors with respect to this total ordering. New neighbor pairs are obtained either by encountering a new object or by deactivating an object which does not intersect the δ-slice any more. It seems to be surprising that this algorithm finds the correct result by just testing pairs of active objects which become neighbors during the sweep. [2] proves that this is sufficient for both an intersection free configuration and a configuration with intersecting objects, a more intuitive proof is given in [1]. The technique used in the above algorithm cannot be applied to the ANN problem for convex objects. Since we search for a nearest neighbor of each object in our configuration, we do not have a global δ for all objects but an individual δ for each object. This makes the technique used in [2] unsuitable for the ANN problem since the deactivation events do not occur in a predefined order and therefore have to be rearranged dynamically. Even worse, the individual δ-values increase the local nature of the problem, and therefore it is not necessarily true that at any time during the left-to-right sweep or the right-to-left sweep an object and any of its nearest neighbors become neighbors in the y-table, i.e. with respect to the total ordering defined in [2]. It is not difficult to construct a set of points for which this approach does not work with respect to the L^2-metric.

The algorithm PSANN (=Plane Sweep All Nearest Neighbors) presented in this paper uses four sweeps, from left to right, from right to left, from top to bottom and from bottom to top. The algorithm also works correctly if we weaken the condition of disjointness to the conditon that the objects in our configuration are *diagonal-disjoint*, which means that for any object the x-diagonal connecting certain x-extremal and the y-diagonal connecting certain y-extremal points do not intersect any other object of S. For configurations consisting of n line segments, circular disks or convex polygons whose number of edges is treated as a constant, PSANN runs in asymptotically optimal time $O(n \log n)$. Let S be a set of n convex polygons with a total of N edges. If each polygon is given by its vertices in cyclic order then PSANN finds nearest neighbors with respect to the Euclidean metric in time $O(n \log N)$. This runtime is achieved by employing