On Neighbors in Geometric Permutations*

Micha Sharir† Shakhar Smorodinsky‡

November 6, 2001

Abstract

We introduce a new notion of 'neighbors' in geometric permutations. We conjecture that the maximum number of neighbors in a set of n pairwise disjoint convex bodies in $\mathbb{R}^d$ is $O(n)$, and we settle this conjecture for $d = 2$. We show that if the set of pairs of neighbors in a set S is of size N, then S admits at most $O(N^{d-1})$ geometric permutations. Hence we obtain an alternative proof of a linear upper bound on the number of geometric permutations for any finite family of pairwise disjoint convex bodies in the plane.

1 Background and Motivation

Let A be a finite family of pairwise disjoint convex bodies in $\mathbb{R}^d$. A line ℓ that intersects every member of A (called a *transversal* of A) induces two linear orderings on A—the orders in which the members of A are met by ℓ, corresponding to the two orientations of ℓ. Katchalski et al. [2] were the first to study such orderings and called them *geometric permutations*. We refer to [2, 3, 4, 6, 7, 10] for the recent study of this concept, its applications and its generalizations.

Let $g_d(n)$ denote the maximum number of geometric permutations in $\mathbb{R}^d$, where the maximum is taken over all such families A of size n. A major open problem in transversal theory is to give sharp asymptotic bounds on $g_d(n)$. Edelsbrunner and Sharir [4] have shown that $g_2(n) = 2n - 2$. Katchalski et al. [6] showed that $g_d(n) = \Omega(n^{d-1})$. The only known general upper bound on $g_d(n)$ is $O(n^{2d-2})$ and is due to Wenger [5]. Hence, for $d \geq 3$ there still exists a wide gap between the known upper and lower bounds. Recently, Smorodinsky et al. [8, 9], obtained the sharp upper and lower bounds of $\Theta(n^{d-1})$ on the number of geometric permutations, in the special case where A consists of pairwise disjoint balls in $\mathbb{R}^d$. (For the case of congruent balls, or nearly congruent the number of geometric permutations is only $O(1)$ [11].) This result was followed by an extension of the same bound

*Work on this paper has been supported by NSF Grants CCR-97-32101 and CCR-00-98246, by a grant from the U.S.-Israeli Binational Science Foundation, by a grant from the Israeli Academy of Sciences for a Center of Excellence in Geometric Computing at Tel Aviv University, and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University.

†School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel, and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA; sharir@math.tau.ac.il.

‡School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel; smoro@tau.ac.il. The research by this author was done while the author was a Ph.D. student under the supervision of Micha Sharir.

to the case of "fat" convex bodies, by Katz and Varadarajan [10]. It has been conjectured that $g_d(n)$ is $O(n^{d-1})$.

1.1 Separation Sets and Neighbors

Wenger [5] introduced the notion of *separation set* which was later generalized in [9]:

Definition 1.1 *Let S be a family of pairwise disjoint convex sets in $\mathbb{R}^d$, and let P be a set of hyperplanes in $\mathbb{R}^d$ passing through the origin. We say that P is a* separation set *for S if for each pair $s_i, s_j \in S$ there exists a hyperplane h, parallel to a hyperplane in P, such that s_i and s_j are contained in different open half-spaces bounded by h.*

For a family S of pairwise disjoint convex sets, let $GP(S)$ denote the number of geometric permutations of S.

Lemma 1.2 *(see [5, 8]) Let S be a collection of pairwise disjoint convex sets in $\mathbb{R}^d$ and let P be a separation set for S. Then $GP(S) = O(|P|^{d-1})$.*

Proof (sketch): Consider the arrangement $\mathcal{A}(P)$ of $|P|$ great spheres in the unit sphere $\mathbb{S}^{d-1}$, each consisting of all orientations parallel to some hyperplane in P. It is easy to show that, for each cell C of $\mathcal{A}(P)$, there is at most one possible geometric permutation that is induced by lines with orientation in C, and this implies the asserted bound. $\square$

Since any set of n pairwise disjoint convex bodies admits a separation set of size $\binom{n}{2}$, one obtains a trivial upper bound of $O(n^{2d-2})$ on the number of geometric permutations in d-space.

It was shown in [9] that, if S is a collection of pairwise disjoint balls in $\mathbb{R}^d$, then there exist a separation set for S of size $O(n)$, where the constant of proportionality depends on the dimension d. Hence, combined with Lemma 1.2, one immediately gets an $O(n^{d-1})$ upper bound on the number of geometric permutations for S. The fact that, in any dimension, balls can be separated with only $O(n)$ hyperplane directions seems to depend on fatness. Indeed, Katz et al. [10] showed later that any set of pairwise disjoint α-fat convex bodies in $\mathbb{R}^d$ (i.e., the ratio between the circumradius and inradius of any input body is at most α, where α is a constant), admits a separation set of size $O(n)$, where the constant of proportionality depends on α and on the dimension d.

The problem with the notion of separation set is that, already in three dimensions, one cannot hope to get a separation set of linear size for general convex bodies. In fact, there exist collections S of n pairwise disjoint convex bodies in $\mathbb{R}^3$, for arbitrarily large n, such that any separation set for S is of size $\Omega(n^2)$. For example, one can take S to be the collection of Voronoi cells of the Voronoi diagram of a set consisting of $n/2$ points on the line $l_1 : x = 0; z = 1$ and of $n/2$ points on the line $l_2 : y = 0; z = 0$. It is easily seen that every pair of points (p, q), such that $p \in l_1$ and $q \in l_2$, have touching Voronoi cells, and that the separating facets of these pairs of cells have different orientations. Shrink the bodies in S by a small amount, to make them pairwise disjoint. Hence any separation set for S must contain a distinct plane for each of these pairs of cells, so its size is $\Theta(n^2)$. Therefore, Lemma 1.2 is useless in the general case.

In this paper, we introduce a weaker notion of separation, and show that it can be used, in a manner similar to that in Lemma 1.2, to derive bounds on the number of geometric permutations.

Definition 1.3 *Let S be a family of n pairwise disjoint compact convex sets in $\mathbb{R}^d$. Two objects $a, b \in S$ are called* neighbors *if there exists a tranversal line l of S such that a and b are consecutive elements of the geometric permutation induced by l.*

Denote by $N(S)$ the set of all neighbors in S. Note that if $GP(S) = 0$ then $N(S) = \emptyset$ and if $GP(S) \geq 1$ then $|N(S)| \geq n - 1$.

Conjecture 1.4 *If S is a family of pairwise disjoint convex bodies in $\mathbb{R}^d$ then $|N(S)| = O(n)$.*

In Section 2 we establish the conjecture for the planar case.

It can be shown that, in the above example involving Voronoi cells, we have $GP(S) \geq 1$, and there are indeed only $O(n)$ neighbors. The following lemma shows that, if the conjecture is true, it leads to sharp bounds on the number of geometric permutations in the general case.

Lemma 1.5 *Let S be a family of n pairwise disjoint compact convex bodies in $\mathbb{R}^d$. Then $GP(S) = O(|N(S)|^{d-1})$*

Proof: For each pair of bodies $a, b \in N(S)$, let $h_{a,b}$ be any hyperplane separating a and b, and let P be the set of these $|N(S)|$ hyperplanes. We proceed as in the proof of Lemma 1.2, by considering the arrangement $\mathcal{A}(P)$ of the great spheres in $\mathbb{S}^{d-1}$ associated with the hyperplanes in P. Fix a (full-dimensional) cell C of $\mathcal{A}(P)$ and let ℓ be an oriented line (transversal to S) with orientation in C. ℓ induces a geometric permutation $\pi = (\pi_1, \pi_2, \ldots, \pi_n)$ on S. For each $i = 1, \ldots, n-1$, the elements π_i, π_{i+1} are neighbors (by definition), so there exists a hyperplane $h_i \in P$ that separates them. h_i corresponds to one of the great spheres in the arrangement $\mathcal{A}(P)$, with C lying in one of the two corresponding hemispheres. This means that every oriented line that intersects both π_i and π_{i+1} and has orientation in C, must intersect these two sets in a fixed order. This implies that in any geometric permutation induced by a line transversal with orientation in C, π_i must precede π_{i+1} , for each $i = 1, \ldots, n-1$. It is easily seen that exactly one geometric permutation (which is π) has this property. We conclude that there is at most one geometric permutation that can be induced by line transversals with orientation in C. Hence the number of geometric permutations is at most the number of cells in $\mathcal{A}(P)$, which is $O(|N(S)|^{d-1})$. $\square$

2 Linear Bound on the Number of Neighbors in the Plane

In this section we prove our main result, showing that our conjecture is true in the plane.

Theorem 2.1 *Let S be a set of n pairwise disjoint compact convex bodies in the plane. Then the number of neighbor pairs in S is $O(n)$.*

Proof: We define a graph G on S, whose edges connect each of the neighbor pairs. We will show that G is a *quasi-planar* graph, i.e., G can be drawn in the plane so that no three edges with distinct endpoints are pairwise crossing. The theorem then follows from the result of Agarwal et al. [1], that quasi-planar graphs have linear size.

We draw G as follows. For each object $a \in S$ we fix a point v_a inside a; these are the vertices of (the drawing of) G. If (a, b) are neighbors, we choose one transversal line ℓ of S along which $\ell \cap a$ and $\ell \cap b$ appear consecutively. Let q_a and q_b be the two nearest endpoints of these two respective intervals. Then we draw the edge ab of G as the polygonal arc $v_a q_a q_b v_b$. See Figure 1.

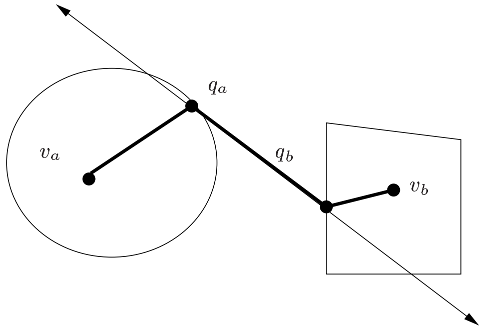

Figure 1: Drawing an edge of G.

Let a, b, c, d, e, f be six distinct objects in S, such that the edges ab, cd, ef are pairwise crossing. Any of these edges, say ab, is drawn as a polygonal arc, whose portion outside $a \cup b$ is a straight segment s_{ab}, connecting a point $q_a \in \partial a$ to a point $q_b \in \partial b$, such that s_{ab} is contained in a line transversal ℓ_{ab} of S, along which a and b are neighbors, and so that the relative interior of s_{ab} is disjoint from a and b (and from all other objects in S as well). Similar properties hold for the edges cd and ef. Note that the crossing between, say ab and cd must be at an interior point of s_{ab} and of s_{cd}, and similarly for the two other crossings.

The removal of s_{ab} from ℓ_{ab} partitions this line into two rays, one of which meets a and is denoted as ρ_a, and the other meets b, and is denoted as ρ_b. A similar notation holds for the two other lines.

Note that the lines ℓ_{ab}, ℓ_{cd}, ℓ_{ef} must be distinct transversals of S. Let Δ be the triangle whose vertices are the crossing points of the segments s_{ab}, s_{cd}, s_{ef}. Order the six objects a, b, c, d, e, f according to the counterclockwise angular order of the endpoints $q_a, q_b, q_c, q_d, q_e, q_f$ about (any point of) Δ. Without loss of generality, assume that this circular order is (a, c, e, b, d, f); note that the three pairs a and b, c and d, and e and f must be 'antipodal' elements of this order. See Figure 2.

Since each of the lines ℓ_{ab}, ℓ_{cd}, ℓ_{ef} is transversal to S, it meets each of the six sets a, b, c, d, e, f. Consider, say, the line ℓ_{ab}. Since the relative interior of the contained segment s_{ab} is disjoint from all sets in S, each of the four other objects c, d, e, f meets ℓ_{ab} at one (and only one) of the two complementary rays ρ_a, ρ_b.

We make the following simplification of the configuration, by replacing each of these six objects by a straight segment. Consider one of the objects, say a. Pick a point $q_a' \in a \cap \ell_{cd}$, and a point $q_a'' \in a \cap \ell_{ef}$. Then $q_a' \in \rho_c \cup \rho_d$ and $q_a'' \in \rho_e \cup \rho_f$. It is impossible that $q_a' \in \rho_d$

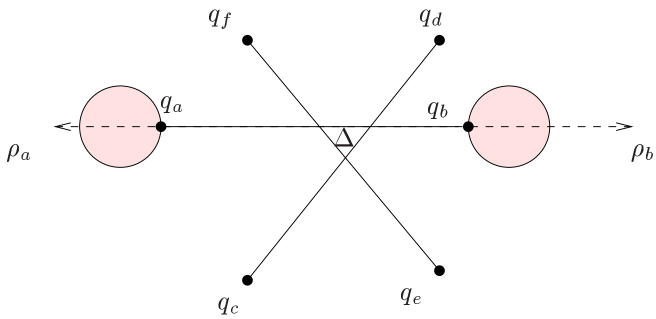

Figure 2: The structure of a triple edge-crossing in G.

and $q''_a \in \rho_e$, for then the segment s_{ab} would not be disjoint from a, as is easily seen (see Figure 2). Hence there are three possible cases (the mnemonics refer to the positioning of q_a in the triple $\{q_a, q'_a, q''_a\}$):

(i) $q'_a \in \rho_c$ and $q''_a \in \rho_f$. In this case we refer to a as a *middle* object, and replace it by the straight segment $q'_a q''_a$; see Figure 3(a).

(ii) $q'_a \in \rho_c$ and $q''_a \in \rho_e$. In this case we refer to a as a *clockwise* object, and replace it by the straight segment $q_a q''_a$; see Figure 3(b).

(iii) $q'_a \in \rho_d$ and $q''_a \in \rho_f$. In this case we refer to a as a *counterclockwise* object, and replace it by the straight segment $q_a q'_a$; see Figure 3(c).

Similar 'straightenings' are applied to the five other objects. Each straightened object is contained in the corresponding original object, and thus the resulting six segments are pairwise disjoint. Note that the straightened a need not contain the point q_a (in case (i)). In this case we replace q_a by the intersection of the straightened a with ℓ_{ab}. This step, applied to each of the objects if needed, may cause some of the segments s_{ab}, s_{cd}, s_{ef} to expand, but it is easily checked that the relative interiors of the new segments are still all disjoint from all the new six objects; see Figure 3(a).

Lemma 2.2 *If one of the six objects is middle or clockwise (resp., middle or counterclockwise) then the object immediately succeeding (resp., preceding) it in the counterclockwise circular order about Δ must be clockwise (resp., counterclockwise).*

Proof: Without loss of generality, assume that the first object is a and that it is middle or clockwise. The succeeding object, c, must meet ℓ_{ab} within ρ_b, or else a and c would have to intersect (because any segment connecting q_c to a point on ρ_a has to cross $q_a q'_a \subseteq a$). Then c must be clockwise. $\square$

Corollary 2.3 *Either all six objects are clockwise or all are counterclockwise.*

Without loss of generality, we may assume that all objects are clockwise. Start at the a-endpoint $v_1 = q_a$ of s_{ab}, and draw through it a line that is parallel to ℓ_{ef} and intersects

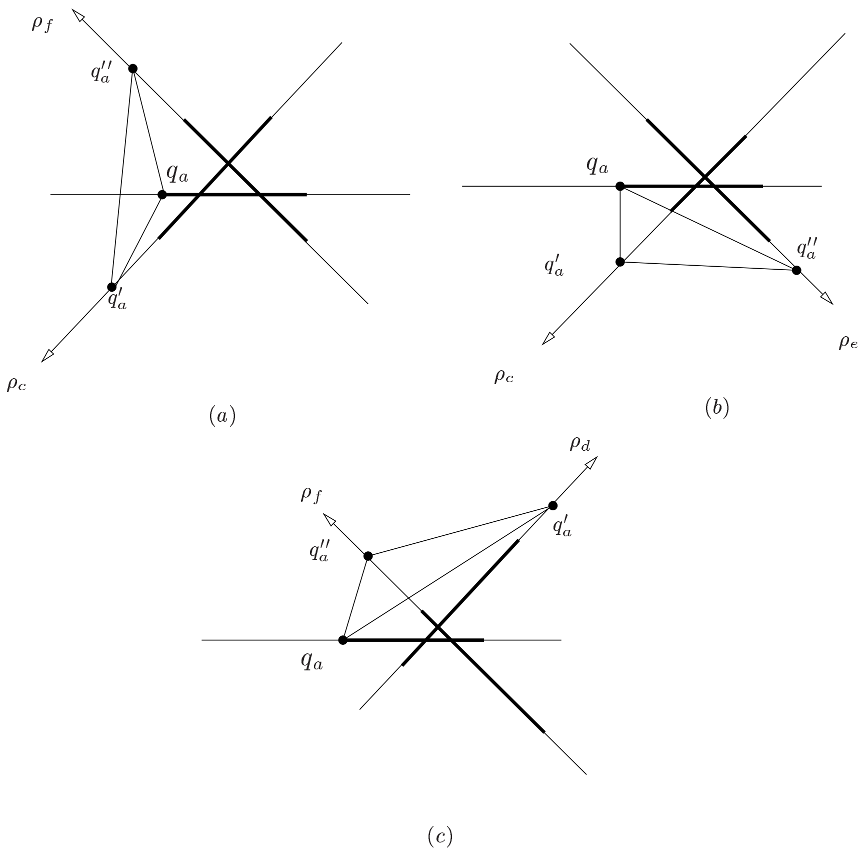

Figure 3: The three types of objects and their straightening.

ℓ_{cd} at a point v_2, which lies counterclockwise from v_1 with respect to Δ. Then draw from v_2 a line parallel to ℓ_{ab} which intersects ℓ_{ef} at a point v_3, again, lying counterclockwise from v_2. Keep 'turning' like this to trace a closed hexagon H. (It is not hard to show that this process does indeed always produce a closed hexagon.) See Figure 4.

We claim that the vertices of H satisfy

$$v_1 = q_a, \quad v_2 \in \rho_c, \quad v_3 \in \rho_e, \quad v_4 \in \rho_b, \quad v_5 \in \rho_d, \quad v_6 \in \rho_f.$$

Indeed, since the segment a is clockwise, it meets ρ_e, and thus must lie counterclockwise to the ray $\overline{q_a v_2}$. Since $a \cap \ell_{cd} \in \rho_c$, we have $v_2 \in \rho_c$ as well. A similar argument, proceeding inductively counterclockwise around H, holds for all other vertices.

Note that the segment f, being clockwise, meets ρ_a and ρ_c. Hence its orientation is on one hand counterclockwise to that of $\overline{v_6 q_a} = \overline{v_6 v_1}$ (or else it would not have met ρ_c), and is on the other hand clockwise to that of $\overline{v_6 v_1}$ (since $v_6 \in \rho_f$ and f meets ρ_a). This contradiction implies that the graph G cannot have three pairwise crossing edges, and thus G is quasi-planar and has linear size. $\square$

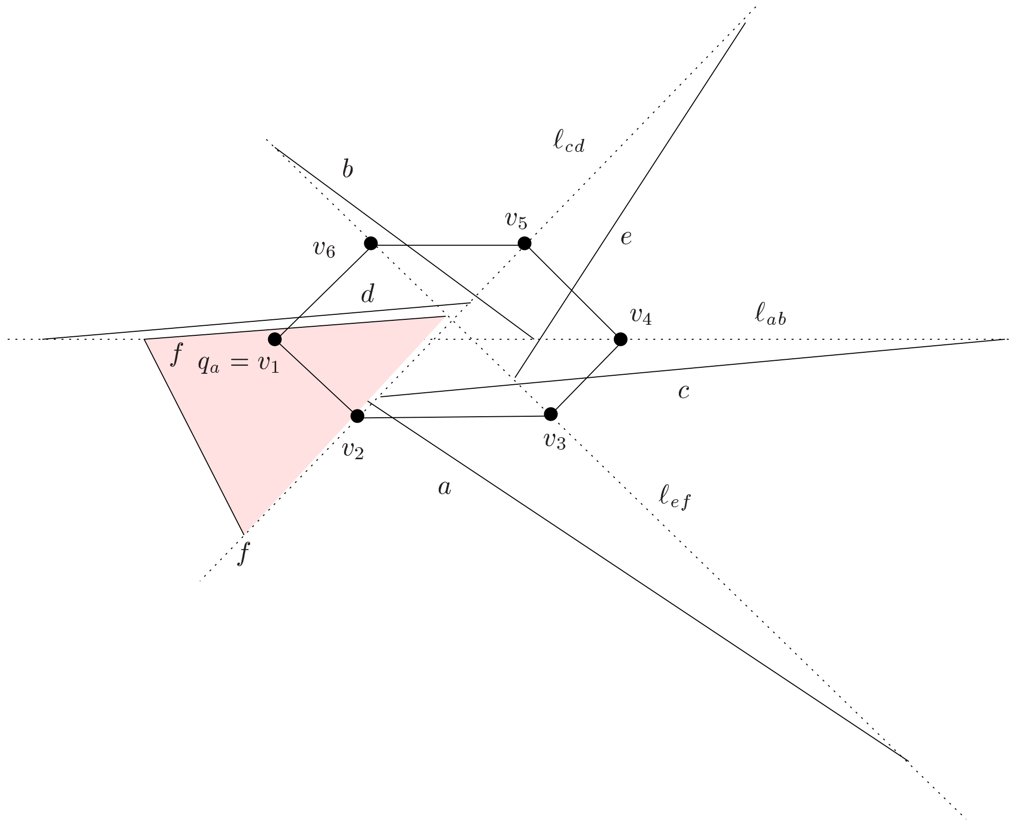

Figure 4: The hexagon H and the final contradiction.

3 Conclusion

Theorem 2.1 combined with Lemma 1.5, yields an alternate proof that $g_2(n) = O(n)$ (albeit weaker than the $2n-2$ bound in [4]). The main open problem is to establish the conjecture in higher dimensions, in particular for $d = 3$.

Interesting by its self, is the use of the quasi-planarity result in the proof of Theorem ref-main. As far as we are aware, this is the first application of it.

References

[1] P. Agarwal, B. Aronov, J. Pach, R. Pollack and M. Sharir, Quasi-planar graphs have a linear number of edges, *Combinatorica* 17:1–9, 1997.

[2] M. Katchalski, T. Lewis, and A. Liu. Geometric permutations and common transversals. *Discrete Comput. Geom.*, 1:371–377, 1986.

[3] J.E. Goodman and R. Pollack. Hadwiger's transversal theorem in higher dimensions. *J. Amer. Math. Soc.*, 1(2):301–309, 1988.

[4] H. Edelsbrunner and M. Sharir. The maximum number of ways to stab n convex non-intersecting sets in the plane is $2n-2$. *Discrete Comput. Geom.*, 5(1):35–42, 1990.

[5] R. Wenger. Upper bounds on geometric permutations for convex sets. *Discrete Comput. Geom.*, 5:27–33, 1990.

[6] M. Katchalski, T. Lewis, and A. Liu. The different ways of stabbing disjoint convex sets. *Discrete Comput. Geom.*, 7:197–206, 1992.

[7] J.E. Goodman, R. Pollack and R. Wenger. Geometric transversal theory. In *New Trends in Discrete and Computational Geometry*, (J. Pach, ed.), Springer Verlag, Heidelberg, 1993, pp. 163–198.

[8] S. Smorodinsky. Geometric Permutations and Common Transversals. Master's thesis, Tel-Aviv University, School of Mathematical Sciences, Tel-Aviv, Israel, July 1998.

[9] S. Smorodinsky, J.S.B Mitchell, and M. Sharir. Sharp bounds on geometric permutations of pairwise disjoint balls in $\mathbb{R}^d$. *Discrete Comput. Geom.*, 23(2):247–259, 2000.

[10] M.J. Katz and K.R. Varadarajan. A tight bound on the number of geometric permutations of convex fat objects in $\mathbb{R}^d$. *Proc. 17th Annu. ACM Sympos. Comput. Geom*, 2001 :249–251.

[11] Y. Zhou and S. Suri. Shape sensitive geometric permutations. *Proc. 12th ACM-SIAM Sympos. Discrete. Algorithms*, 2001 :234–243.