Sharp Bounds on Geometric Permutations of Pairwise Disjoint Balls in $\mathbb{R}^{d*}$

S. Smorodinsky,[1] J. S. B. Mitchell,[2] and M. Sharir[1,3]

[1]School of Mathematical Sciences, Tel Aviv University,
Tel Aviv 69978, Israel
{shakhar,sharir}@math.tau.ac.il

[2]Department of Applied Mathematics and Statistics,
State University of New York,
Stony Brook, NY 11794, USA
jsbm@ams.sunysb.edu
http://www.ams.sunysb.edu/~jsbm/

[3]Courant Institute of Mathematical Sciences,
New York University,
New York, NY 10012, USA

Abstract. We prove that the maximum number of geometric permutations, induced by line transversals to a collection of n pairwise disjoint balls in $\mathbb{R}^d$, is $\Theta(n^{d-1})$. This improves substantially the upper bound of $O(n^{2d-2})$ known for general convex sets [9].

We show that the maximum number of geometric permutations of a sufficiently large collection of pairwise disjoint unit disks in the plane is two, improving the previous upper bound of three given in [5].

1. Introduction

Let $\mathcal{A}$ be a family of bodies in $\mathbb{R}^d$. A line ℓ is said to be a *transversal* for $\mathcal{A}$ if it intersects every member of $\mathcal{A}$. If $\mathcal{A}$ consists of pairwise disjoint convex bodies, then a line transversal for $\mathcal{A}$ induces two linear orderings on $\mathcal{A}$—the orders in which the members of $\mathcal{A}$ are met by ℓ, corresponding to the two orientations of ℓ. Katchalski et al. [4] were the first to study such orderings and called them *geometric permutations*.

Let $g_d(n)$ denote the maximum number of geometric permutations, where the maximum is taken over all such families $\mathcal{A}$ of size n. The following asymptotic bounds are known on $g_d(n)$.

1. $g_2(n) = 2n - 2$ (see [3]).
2. $g_d(n) = \Omega(n^{d-1})$ (see [6]).
3. $g_d(n) = O(n^{2d-2})$ (see [9]).

It has been conjectured that $g_d(n)$ is $O(n^{d-1})$.

In 1985 Katchalski et al. [7] constructed a family of $n \geq 4$ pairwise disjoint convex sets in $\mathbb{R}^2$ that has $2n-2$ geometric permutations; see Fig. 1. Five years later, Edelsbrunner and Sharir [3] showed that $2n-2$ is the maximum possible number of geometric permutations in the plane.

In 1992 Katchalski et al. [6] generalized their lower bound construction and showed that there exist families $\mathcal{A}$ of n pairwise disjoint convex sets in $\mathbb{R}^d$ admitting $\Omega(n^{d-1})$

Fig. 1. An example of n convex sets admitting $2n - 2$ geometric permutations (taken from [7]).

geometric permutations. In this work we show that this lower-bound holds even for collections of balls in d dimensions.

The only known general upper bound for the number of geometric permutations of collections of pairwise disjoint convex sets in $\mathbb{R}^d$ is $O(n^{2d-2})$, and is due to Wenger [9]. Hence, for $d \geq 3$ there still exists a wide gap between the known upper and lower bounds.

In this work we close this gap for families of balls. Specifically, we show that the maximum number of geometric permutations for collections of n pairwise disjoint balls in $\mathbb{R}^d$ is $\Theta(n^{d-1})$. The upper bound is a consequence of the following main technical result of the paper: A collection $\mathcal{A}$ of n pairwise disjoint balls in $\mathbb{R}^d$ admits a family H of at most $c_d n$ hyperplanes, such that every pair of balls in $\mathcal{A}$ is separated by a hyperplane in H.

An interesting special case of the problem is when all the balls in $\mathcal{A}$ are congruent (say, unit balls). It is conjectured that if $|\mathcal{A}|$ is sufficiently large, then it can have only $O(1)$ geometric permutations, with the bound depending on d. This conjecture is open for $d \geq 3$. In the plane it is known [5] that a collection of n pairwise disjoint translates of any convex body admits at most three geometric permutations, assuming n is sufficiently large, and this bound is tight in the worst case.

In this paper we show that the maximum number of geometric permutations for sufficiently many pairwise disjoint unit disks in the plane is only two. An independent (and different) proof of the same bound has recently been obtained by Katchalski and Asinowski (private communication; see [2]).

2.   Geometric Preliminaries

In this section we develop a few technical concepts that we use in subsequent sections. We also generalize the notion of a *separation set* used by Wenger (see [9]).

Definition 2.1.   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 2.2 (see [8] and [9]).   *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.   Consider the arrangement, $A(P)$, of the great spheres in $\mathcal{S}^{d-1}$ associated with the hyperplanes in P (i.e., each of these great spheres consists of all orientations parallel to some hyperplane in P). The arrangement partitions the unit sphere into $O(|P|^{d-1})$ connected components (i.e., open $(d-1)$-dimensional cells of the arrangement). Fix a connected component C and a pair of sets $s_i, s_j \in S$. There exists a hyperplane H that separates s_i and s_j; H corresponds to one of the great spheres in the arrangement

$A(P)$, with C lying in one of the two corresponding hemispheres. This means that every oriented line that intersects both s_i and s_j with orientation in C, must intersect the sets in a fixed order. This is true for every pair of sets in S. Therefore every line transversal to S with orientation in C induces the same order on S. Thus the number of geometric permutations for S is at most the number of connected components in the arrangement $A(P)$, implying the asserted bound. (Note that we can ignore lower-dimensional faces of $A(P)$, because no line with orientation in such a face can be a transversal of S.)   □

Corollary 2.3 (see [9]). $g_d(n) = O(n^{2d-2})$.

Proof.   For any set S of n pairwise disjoint convex sets in $\mathbb{R}^d$ there exists a separation set of size $\binom{n}{2}$: separate each pair of sets in S by a different hyperplane. The corollary is now immediate from Lemma 2.2.   □

3.   Geometric Permutations of Pairwise Disjoint Unit Disks

In this section we study the case of families of pairwise disjoint unit disks in the plane. We show that, for n sufficiently large, such a family admits at most two geometric permutations. The case of pairwise disjoint unit disks is a special case of disjoint translates of a convex body. Geometric permutations of disjoint translates of a convex set were studied by Katchalski et al. in [5] and [6], where the following result was proved.

Theorem 3.1.   *The maximum number of geometric permutations for finite families of disjoint translates of a convex body in $\mathbb{R}^2$ is three.*

For a construction that achieves this bound (for unit disks), see Fig. 2. Asinowski and Katchalski (see [2]) gave an example of a family of more than three translates admitting three geometric permutations; see Fig. 3. This construction can be generalized to any number of translates of the same fixed body.

However, we show that if the convex body s is a disk, there exists a constant c such that for any set S of more than c pairwise disjoint translates of s, the number of geometric

Fig. 2.   An example of three unit disks admitting three geometric permutations.

Fig. 3.   An example of five translates admitting three geometric permutations.

permutations for S is at most two. This proof is inspired by that of [8], but is simpler. Asinowski and Katchalski (see [2]) have an independent proof for this result. We begin the proof with some observations:

Lemma 3.2.   *Let S be a collection of n pairwise disjoint unit disks in the plane, such that S admits at least one line transversal. There exist two disks such that the distance between their centers is at least $\pi n/4 - 2$.*

Proof.   If ℓ is a line transversal of S, then all disks in S are contained in a strip of width 4 about ℓ. Denote by t the distance between the center of the first disk crossed by ℓ and the center of the last disk crossed by ℓ. Since the sum of areas of the disks is less than the area of the smallest rectangle bounding the disks with an edge in direction ℓ, we have $(t + 2) \cdot 4 \geq n\pi$, which implies the above inequality. $\square$

Claim 3.3.   *Let a and b be two disjoint unit disks in the plane, and let w_{ab} denote the double wedge that contains a and b and is bounded by the two separating common tangents of a and b. Then the set of orientations of the line transversals to a and b consists of a pair of antipodal arcs on the circle of orientations delimited by the orientations of these two separating tangents and consisting of orientations of lines that cross both wedges of w_{ab}. (See Fig. 4 for an illustration.)*

As a result of Lemma 3.2 and Claim 3.3, and using the fact that the angle between the two separating tangents is proportional to $\sin^{-1}(1/t)$, where t is the distance between the first and last stabbed disks, we can assume that all orientations of line transversals to

Fig. 4.   A transversal ℓ for two disks must cross the double wedge defined by the two common tangents.

a collection S of pairwise disjoint unit disks lie in an ε-neighborhood (more precisely, in the union of two antipodal ε-neighborhoods) on the unit circle of directions, for $\varepsilon = O(1/n)$. We can thus make ε as small as we wish by increasing the number of disks.

Let ε be the angle between the two separating common tangents of the farthest pair of disks in S. All line transversal orientations should lie in the union I of the two antipodal ε-neighborhoods of orientations mentioned above. Moreover, a line transversal to S is, in particular, a transversal to all pairs of disks in S and therefore must lie in the intersection of all of the corresponding unions of antipodal intervals, one pair for each pair of disks in S. The only case in which each of the ε-neighborhoods in I can be broken into more than one connected component is that in which a pair of disks is close enough so that their common tangents define an antipodal pair of intervals of admissible orientations whose sizes are close to π and such that they both cross the two antipodal intervals of I. See Fig. 5 for an illustration. In that case, simple trigonometry shows that the two corresponding disks must be very close to one another (the distance between their centers is at most $2 + O(\varepsilon^2) = 2 + O(1/n^2)$), and that the two disks in such a pair must be adjacent in any geometric permutation. Moreover, the line connecting the centers of such a pair of disks must be almost perpendicular to any line transversal of S. This implies that a disk can belong to at most one such pair. Also, pairs of disks that are not as close to each other appear in the same order in any geometric permutation.

Fig. 5. The intersection of the antipodal arcs I (corresponding to disks a and b) with another pair of antipodal arcs I' (corresponding to disks a' and b') can break each interval of I into two connected components.

We use the following lemma, which was proved in [5]:

Lemma 3.4. *Let $S = \{a, b, c, d\}$ be a set of four pairwise disjoint convex bodies in the plane. If (a, b, c, d) is a geometric permutation of S, then (b, a, d, c) cannot be a geometric permutation of S.*

It follows immediately that any two geometric permutations differ in a single pair of consecutive disks, whose order is interchanged between the two permutations. This in turn implies that there exists at most one pair (a, b) of disks such that any two geometric permutations are obtained from each other by interchanging a and b, implying that S admits at most two geometric permutations. To see this last claim, assume to the contrary that there are two pairs (a, b) and (c, d) that consist of four distinct disks such that there are at least three geometric permutations in which, without loss of generality, both pairs appear in the following possible orders:

- $\pi_1 = (\ldots, a, b, \ldots, c, d, \ldots)$,
- $\pi_2 = (\ldots, b, a, \ldots, c, d, \ldots)$,
- $\pi_3 = (\ldots, a, b, \ldots, d, c, \ldots)$;

that is, π_2 is obtained from π_1 by interchanging a and b and π_3 is obtained from π_1 by interchanging c and d. However, this is a contradiction since π_2 and π_3 are two geometric permutations that contradict Lemma 3.4.

We have proved:

Theorem 3.5. *There exists a constant c such that any family of more than c pairwise disjoint congruent disks in the plane admits at most two geometric permutations.*

Remark 3.6. As we will show in Section 4.2, this result does not hold when the disks are not congruent. In this case there are examples with $\Omega(n)$ geometric permutations. However, in these examples the ratio between the largest and the smallest radii is unbounded (the ratio depends on n). We conjecture that Theorem 3.5 continues to hold when this ratio is bounded. (Recall that the theorem may fail for translates of a fixed convex set, as depicted in Fig. 3.)

4. Geometric Permutations of Pairwise Disjoint Balls in $\mathbb{R}^d$

4.1. Upper Bounds

Let S be a given set of n pairwise disjoint (closed) balls in $\mathbb{R}^d$. We prove that $g_d(S) = O(n^{d-1})$. The main step of the proof is to show that S admits a separation set of size $O(n)$. As a matter of fact, we prove the stronger result that there exists a set H of $O(n)$ hyperplanes such that each pair of balls in S is separated by a hyperplane in H, rather than a hyperplane parallel to one in H.

Fig. 6. The construction of $h_{i,k}$.

Let $S = \{B_1, \ldots, B_n\}$ be a set of n pairwise disjoint balls in $\mathbb{R}^d$; ball B_i has radius r_i and center b_i. We assume, without loss of generality, that $r_1 > r_2 > \cdots > r_n$. (If several balls have the same radius, we slightly increase their radii, making them all distinct and keeping the balls disjoint. This can only increase $g_d(S)$.)

Let $\mathcal{S}_{d-1}$ be the unit sphere of directions. Let $\mathcal{C} = \{C_1, \ldots, C_K\}$ be a covering of $\mathcal{S}_{d-1}$ by a set of K spherical patches of diameter δ, where δ is chosen so that the angle θ between any pair of unit vectors $\hat{u}, \hat{v} \in C_k$ is at most $\sin^{-1}((\sqrt{3} - 1)/2) \approx 0.375$ (or about $21.5°$). Each set C_k determines a convex cone $C_k(p)$ with respect to any given apex point p; this is the union of all rays emanating from p and having orientations in C_k. Note that we can always cover $\mathcal{S}_{d-1}$ with a *constant* number (depending on dimension) of sets C_k; i.e., K is a constant, depending (exponentially) on d.

We construct a set H of $O(n)$ hyperplanes as follows. Consider a ball B_i and a set C_k of directions, which define a cone, $C_k(b_i)$, with apex at b_i. If $C_k(b_i)$ contains the center of at least one ball that is larger than B_i, then we let B_j ($j < i$) be that ball with center $b_j \in C_k(b_i)$ closest to b_i, and we define $h_{i,k}$ to be the hyperplane supporting B_i, orthogonal to the vector $b_j - b_i$ and separating b_i and b_j; see Fig. 6. Clearly, $h_{i,k}$ separates B_i from B_j. We let H be the set of all such hyperplanes $h_{i,k}$; since K is a constant depending on dimension, $|H| = O(n)$, for any fixed dimension d.

Theorem 4.1. *H is a separating set for S.*

Proof. We must show that for every choice of B_i, and $j < i$, there is a hyperplane in H that separates B_i from B_j.

Our proof is by induction on i. The base of the induction is the trivial claim that H contains hyperplanes separating B_1 from each ball that has larger radius (there are none). We now make the following induction hypothesis (on i): *H contains a hyperplane separating B_i from each B_j with $j < i$*.

Suppose the hypothesis holds for all $i' \leq i$, and consider ball $B = B_{i+1}$. Without loss of generality, we can assume that $r_{i+1} = 1$ and b_{i+1} is the origin, O. Consider an

arbitrary $B' = B_j$, with $j < i + 1$, radius $r' = r_j > 1$, and center $v = b_j$ lying in a cone $C = C_k(b_{i+1})$, for some $k \in \{1, \dots, K\}$.

By the construction of H, since C contains the center of a larger ball, we know that there exists a hyperplane $h = h_{i+1,k} \in H$ separating B from some ball, B'', with radius $r'' > 1$ and center $u \in C$. (In fact, by construction, h is supporting B and is orthogonal to u.) Our goal is to show that H contains a hyperplane separating B from B'. If $B' = B''$, we are done. So, we assume that B' and B'' are distinct.

By the induction hypothesis, there exists a hyperplane $h' \in H$ that separates B' from B'' (since each has radius larger than that of B). If h already separates B' from B, then we are done. So we assume that it does not, which means that B' intersects h.

We let θ be the angle between u and v. We let ρ denote the ray containing u with endpoint at the origin. We let $p = h \cap B$ denote the point on ρ where h supports B, and we let p' denote the point on ρ, further from p, at distance $|v - p|$ from p. Finally, we let θ' denote the angle between vector $v - p$ and ρ. See Fig. 7 for an illustration.

Fig. 7. Illustration of the notation in the proof of Theorem 4.1. (The dotted loop surrounding B' is meant to convey the fact that B' is assumed to cross h, even though, for clarity, we have not drawn it large enough to do so.)

We will need the following technical lemma:

Lemma 4.2. $2 \sin(\theta'/2) \leq \cos \theta'$.

Proof. Referring to Fig. 7, we need to show that $|vp'| \leq |pp''|$, where p'' is the foot of the perpendicular from v to ρ. It is easily seen that this can be rewritten as

$$\frac{|v| \sin \theta}{\cos(\theta'/2)} \leq |v| \cos \theta - 1.$$

Since θ is acute and $|v| > 2$, it follows that $\angle Ovp < \theta$ and hence $\theta' < 2\theta$. We thus have

$$\frac{|v| \sin \theta}{\cos(\theta'/2)} \leq |v| \tan \theta,$$

so it suffices to show that $|v| \tan \theta \leq |v| \cos \theta - 1$; since $|v| > 2$, it suffices to show that $\cos \theta - \tan \theta > \frac{1}{2}$, or that $1 - \sin^2 \theta - \sin \theta \geq \frac{1}{2} \cos \theta$. By construction, we have $\sin \theta \leq (\sqrt{3} - 1)/2$, which implies that $1 - \sin^2 \theta - \sin \theta \geq \frac{1}{2}$, thus completing the proof of the lemma. □

Note that Lemma 4.2 trivially implies that $\theta' \leq \pi/4$.

First, we claim that B' intersects ρ in an interval that lies after u (i.e., an interval of points that are farther from the origin than is the point u); thus, h' separates the origin (and B'') from B'. We argue as follows. Since $\theta' \leq \pi/4$, we know that point v is at least as close to ray ρ as it is to hyperplane h; thus, B' intersects ray ρ. By Lemma 4.2, v is in fact closer to point p' than to any point on h; thus, B' contains point p'. Now, by construction of H, $|u| \leq |v|$, which implies that $|u - p| = |u| - 1 \leq |v| - 1 \leq |v - p| = |p' - p|$. Thus, ray ρ intersects B' after B''. Since h' separates B' and B'', ray ρ must intersect B before B'' before h' before B'.

Second, we claim that h' does not intersect B; thus, h' separates B from B'. To see this claim, consider for each $q \in B$ the ray ρ_q that is parallel to ρ, with apex q. Since B'' is larger than B, each ray ρ_q must intersect B''. Now ray ρ intersects B before B'' before h', so, by continuity, each ray ρ_q must also intersect B before B'' before h'. This shows that h' cannot intersect B, since every point $q \in B$ is the apex of a ray that intersects h' only after passing through B'' (which is disjoint from h').

Since we have shown that h' separates B and B', this completes the induction step and thus concludes the proof of the theorem. □

As a result of Lemma 2.2 and Theorem 4.1 we have:

Theorem 4.3. *The number of geometric permutations of a set of n pairwise disjoint balls in $\mathbb{R}^d$ is $O(n^{d-1})$.*

Remark 4.4. For general pairwise disjoint convex sets in $\mathbb{R}^3$, the size of a separating set can be $\Theta(n^2)$. For example, in the standard construction of a Voronoi diagram in $\mathbb{R}^3$

with $\Theta(n^2)$ complexity, one needs $\Theta(n^2)$ different plane orientations to separate all pairs of cells. Hence the current proof of Theorem 4.3 does not extend to families of general convex sets.

4.2. *Lower Bounds*

In this subsection we prove that the lower bound of $\Omega(n^{d-1})$ for $g_d(n)$ can be attained by a family of n pairwise disjoint balls in $\mathbb{R}^d$.

Theorem 4.5. *There exists a set S of n pairwise disjoint balls in $\mathbb{R}^d$ that admits $\Omega(n^{d-1})$ geometric permutations.*

Proof. For simplicity we start with $\mathbb{R}^2$. We construct a family of n pairwise disjoint disks in the plane admitting $\lfloor n/2 \rfloor$ geometric permutations. We then extend this construction to obtain a family of n pairwise disjoint balls in $\mathbb{R}^d$ admitting $\Omega(n^{d-1})$ geometric permutations. The constructions are inspired by those of Agarwal et al. [1].

Consider a line ℓ through the origin O in $\mathbb{R}^2$. Let ε be a small quantity to be fixed later and let $R_1 > 0$. Place two disks of radius R_1 tangent to ℓ at O from the opposite sides of ℓ, and then move them apart slightly, perpendicular to ℓ, such that any line that passes through O and makes an angle larger than ε with ℓ is a transversal for this pair of disks. Denote the resulting two disks by D_l^1 and D_l^2. Note that no line transversal for $\{D_l^1, D_l^2\}$ is parallel to ℓ. Note that the distance ρ, between D_l^1 and D_l^2 can be made arbitrarily large if R_1 is sufficiently large, while ε is kept fixed.

Assume, without loss of generality, that n is even. Consider a set of $n/2$ distinct lines $\{l_1, \ldots, l_{n/2}\}$ passing through the origin. We construct the set $S = \{D_{l_i}^1, D_{l_i}^2 \mid i = 1, \ldots, n/2\}$ by induction. Suppose we have already constructed the set of pairwise disjoint disks $S_j = \{D_{l_k}^1, D_{l_k}^2 \mid k = 1, \ldots, j\}$, for some $j \geq 0$ We choose R_{j+1} sufficiently large such that it is possible to place two disks $D_{l_{j+1}}^1, D_{l_{j+1}}^2$ of radius R_{j+1} outside the convex hull of the union of the disks in S_j, so that the line connecting their centers passes through O and is orthogonal to l_{j+1}, as above, and so that every line that passes through O and makes an angle larger than ε with line l_{j+1} is a transversal for $\{D_{l_{j+1}}^1, D_{l_{j+1}}^2\}$ and no line transversal for $\{D_{l_{j+1}}^1, D_{l_{j+1}}^2\}$ is parallel to l_{j+1}. We proceed in this manner until the whole set S is constructed. We claim that if 2ε is smaller than the smallest angle between any pair l_i, l_j of our lines, then $GP(S)$ consists of $n/2$ geometric permutations. Indeed, fix one of the lines l_i. No line that makes an angle less than ε with line l_i is a transversal for $\{D_{l_i}^1, D_{l_i}^2\}$. Hence no line with orientation that lies in an ε-neighborhood of the orientation α_i of l_i (or in a neighborhood of $\alpha_i + \pi$) is a transversal for S. Moreover, every line through the origin whose orientation is in the complement of the union of all ε-neighborhoods of α_i, and of $\alpha_i + \pi$ for $i = 1, \ldots, n/2$, must be a transversal for S since it intersects $D_{l_i}^1$ and $D_{l_i}^2$ for every $i = 1, \ldots, n/2$. Thus taking ε to be sufficiently small, as prescribed above, we get n pairwise disjoint intervals on the circle of directions $\mathcal{S}_1$ and any orientation in the

Fig. 8.   The set S_2 consists of four pairwise disjoint disks and admits two geometric permutations.

complement of the union of these intervals is the orientation of some line transversal for S. This complement consists of n connected components that come in $n/2$ antipodal pairs, and any pair of nonantipodal components gives rise to different geometric permutations, since the order of at least one pair $D_{l_i}^1$, $D_{l_i}^2$ is reversed. It is easily seen that orientations in the same component of the complement give rise to the same geometric permutation. We thus get $n/2$ different geometric permutations. See Fig. 8 for an illustration.

We modify the above construction, to get a set S of n pairwise disjoint balls in $\mathbb{R}^d$ admitting $\Omega(n^{d-1})$ geometric permutations, in the following manner. Instead of lines we place $n/2$ distinct hyperplanes $\{\pi_1, \ldots, \pi_{n/2}\}$ passing through the origin, such that no d hyperplanes pass through the same line. The pairs of balls are constructed in an analogous manner; i.e., for each hyperplane π_i in turn, we place two congruent balls outside the convex hull of all the balls $\{B_j^1, B_j^2 \mid j = 1, \ldots, i-1\}$ such that every line that passes through the origin and makes an angle larger than ε with hyperplane π_i is a transversal for $\{B_i^1, B_i^2\}$ and no line parallel to hyperplane π_i is a transversal for this pair of balls. Let $\mathcal{S}_{d-1}$ denote the (unit-radius) sphere of directions in $\mathbb{R}^d$, centered at the origin; each direction in d-space can be represented as a point on $\mathcal{S}_{d-1}$. The locus of orientations of lines parallel to π_i is a great $(d-2)$-sphere on $\mathcal{S}_{d-1}$, and the $n/2$ great spheres $\{C_1, \ldots, C_{n/2}\}$, in general position, corresponding to the $n/2$ hyperplanes $\{\pi_1, \ldots, \pi_{n/2}\}$ induce an arrangement $\mathcal{A}$ on $\mathcal{S}_{d-1}$ with $\Theta(n^{d-1})$ cells. To finish the argument it is sufficient to show that each such cell contains at least one line orientation corresponding to a transversal for S. Indeed, any line passing through O and making an angle larger than ε with each hyperplane π_i is a transversal for S. Therefore every point on $\mathcal{S}_{d-1}$ outside of each of the bands of half-width ε centered around each of the $n/2$ great spheres C_i is the orientation of a line transversal for S through O. Picking ε small enough (as a function of n) we can assure that the bands are sufficiently narrow so that every cell of $\mathcal{A}$ contains a point outside of the bands. It is easy to see that two lines with orientations that belong to different nonantipodal cells in this arrangement induce two different linear orderings on the balls, since at least two balls B_i^1, B_i^2 change their order in these two orderings. This completes the construction of a family of balls admitting $\Omega(n^{d-1})$ linear orderings. Since each geometric permutation is counted at most twice in these linear orderings, we have $\Omega(n^{d-1})$ geometric permutations. $\square$

5. Open Problems

In conclusion, we mention a few open problems suggested by the results of this paper:

1. Although Remark 4.4 shows that separation sets may have quadratic size in general, we conjecture that $O(n)$ orientations suffice to separate all pairs of sets that are adjacent in some geometric permutation of the given collection, for any collection of n pairwise disjoint convex sets in $\mathbb{R}^d$. If true, this conjecture would imply an $O(n^{d-1})$ bound on the maximum number of geometric permutations for families of pairwise disjoint convex sets in $\mathbb{R}^d$.
2. If the balls in S are all congruent, we conjecture (as noted in the Introduction) that the number of geometric permutations of S is $O(1)$. As far as we know, this problem is open already for $d = 3$.
3. Is $O(n^{d-1})$ also a bound on the maximum number of geometric permutations for families of pairwise disjoint convex sets in $\mathbb{R}^d$?

References

1. P. K. Agarwal, B. Aronov, and M. Sharir. Line traversals of balls and smallest enclosing cylinders in three dimensions. *Discrete Comput. Geom.*, 21:373–388, 1999.
2. A. Asinowski. Common transversals and geometric permutations. Master's thesis, Technion IIT, Haifa, 1998.
3. H. Edelsbrunner and M. Sharir. The maximum number of ways to stab n convex nonintersecting sets in the plane is $2n - 2$. *Discrete Comput. Geom.*, 5(1):35–42, 1990.
4. M. Katchalski, T. Lewis, and A. Liu. Geometric permutations and common transversals. *Discrete Comput. Geom.*, 1:371–377, 1986.
5. M. Katchalski, T. Lewis, and A. Liu. Geometric permutations of disjoint translates of convex sets. *Discrete Math.*, 65:249–259, 1987.
6. M. Katchalski, T. Lewis, and A. Liu. The different ways of stabbing disjoint convex sets. *Discrete Comput. Geom.*, 7:197–206, 1992.
7. M. Katchalski, T. Lewis, and J. Zaks. Geometric permutations for convex sets. *Discrete Math.*, 54:271–284, 1985.
8. S. Smorodinsky. Geometric permutations and common transversals. Master's thesis, School of Mathematical Sciences, Tel-Aviv University, Tel-Aviv, July 1998.
9. R. Wenger. Upper bounds on geometric permutations for convex sets. *Discrete Comput. Geom.*, 5:27–33, 1990.