Improved Bounds for Geometric Permutations

Natan Rubin
Tel Aviv University
rubinnat@post.tau.ac.il

Haim Kaplan
Tel Aviv University
haimk@post.tau.ac.il

Micha Sharir
Tel Aviv University
michas@post.tau.ac.il

***Abstract*—We show that the number of geometric permutations of an arbitrary collection of n pairwise disjoint convex sets in $\mathbb{R}^d$, for $d \geq 3$, is $O(n^{2d-3} \log n)$, improving Wenger's 20 years old bound of $O(n^{2d-2})$.**

***Keywords*-geometric permutations; line transversals; convex sets; arrangements**

I. Introduction

Let $\mathcal{K}$ be a collection of n convex sets in $\mathbb{R}^d$. A line ℓ is a *transversal* of $\mathcal{K}$ if it intersects all the sets in $\mathcal{K}$. If the objects in $\mathcal{K}$ are *pairwise disjoint*, an oriented line transversal meets them in a well-defined order, called a *geometric permutation*. The study of geometric permutations plays a central role in geometric transversal theory; see [7], [20] for compherensive surveys.

Previous work. In 1985, Katchalski et al. [10] initiated the study of the maximum possible number $g_d(n)$ of geometric permutations induced by a set $\mathcal{K}$ of n pairwise disjoint convex objects in $\mathbb{R}^d$. For $d = 2$, their analysis, combined with that in [5], yields $g_2(n) = 2n - 2$. Wenger [19] proved in 1990 that $g_d(n) = O(n^{2d-2})$ in any dimension $d \geq 3$. In 1992, Katchalski et al. [11] showed that $g_d(n) = \Omega(n^{d-1})$ for any $d \geq 3$. Since then, closing (or even reducing) the fairly large gap between these upper and lower bounds on $g_d(n)$, in any dimension $d \geq 3$, has remained one of the major long standing open problems in geometric transversal theory.

Improved bounds were obtained for several special cases. Smorodinsky et al. [17] derived a tight upper bound of $\Theta(n^{d-1})$ on the number of geometric permutations induced by an arbitrary collection of n pairwise disjoint balls in $\mathbb{R}^d$. Katz and Varadarajan [13] generalized this result to arbitrary collections of n pairwise disjoint *fat* convex bodies. Other recent works [3], [8], [12], [21] show that the maximum possible number of geometric permutations induced by pairwise disjoint *unit* balls (or, more generally, balls of bounded size disparity) is constant in any dimension.

Work by Haim Kaplan has been supported by Grant 2006/204 from the U.S.-Israel Binational Science Foundation and by Grant 975/06 from the Israel Science Fund. Work by Micha Sharir was partially supported by NSF Grant CCF-08-30272, by Grant 2006/194 from the U.S.-Israel Binational Science Foundation, by grant 338/09 from the Israel Science Fund, and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University. This work is part of the first author's Ph.D. Dissertation prepared under the supervision of the other two authors.

Other studies bound the number of geometric permutations induced by arbitrary collections of pairwise disjoint convex sets, whose realizing transversal lines belong to some restricted subfamily of lines in $\mathbb{R}^d$. For example, Aronov and Smorodinsky [2] derive a tight bound of $\Theta(n^{d-1})$ on the maximum number of geometric permutations realized by lines that pass through a fixed point in $\mathbb{R}^d$. A recent paper [9] by the authors studies line transversals of arbitrary convex polyhedra in $\mathbb{R}^3$ and derives (as a byproduct) an improved upper bound of $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$, on the number of geometric permutations realized by lines which pass through a fixed line in $\mathbb{R}^3$.

The space of line transversals. Lines in $\mathbb{R}^d$ have $2d-2$ degrees of freedom, and are naturally represented in a real projective space (so-called the *Grassmannian manifold*; see [7]). However, for the purpose of combinatorial analysis, we can represent them (with the exclusion of some "negligible" subset which we may ignore) by points in the real Euclidean space $\mathbb{R}^{2d-2}$; see [7] for more details.

Let $\mathcal{K}$ be a collection of n convex sets in $\mathbb{R}^d$, not necessarily pairwise disjoint. The *transversal space* $\mathcal{T}(\mathcal{K})$ of $\mathcal{K}$ is the set in $\mathbb{R}^{2d-2}$ of all (points representing) the transversal lines of $\mathcal{K}$.

If the sets of $\mathcal{K}$ are pairwise disjoint, any two lines in the same connected component of $\mathcal{T}(\mathcal{K})$ induce the same geometric permutation, so the number of geometric permutations is upper bounded by the number of components of $\mathcal{T}(\mathcal{K})$. In two dimensions, the converse property also holds [6].

The situation becomes considerably more complicated already in $\mathbb{R}^3$: There exist collections of four (pairwise disjoint) convex sets whose transversal space consists of an arbitrarily large number of connected components [6], [9]. Thus, the shape of $\mathcal{T}(\mathcal{K})$ depends on the shape of the sets in $\mathcal{K}$, and may grow out of control if we do not impose any restrictions on the sets of $\mathcal{K}$.

In three dimensions, if the sets in $\mathcal{K}$ have *constant description complexity* (i.e., each set can be described as a Boolean combination of a constant number of polynomial equalities and inequalities of constant maximum degree), the analysis of Koltun and Sharir [14] yields an improved bound of $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$, on the combinatorial complexity, and thus also on the number of connected components, of $\mathcal{T}(\mathcal{K})$. (If $\mathcal{K}$ is a collection of n triangles in $\mathbb{R}^3$, an improved bound of $O(n^3 \log n)$ holds; see [1].) Hence, this also serves as an upper bound on the number of geometric permutations induced by any such collection $\mathcal{K}$. The strength (and beauty) of Wenger's analysis is that it does not make any assumptions whatsoever on the shape of the sets in $\mathcal{K}$ (other than being convex and pairwise disjoint).

Our results. We first show that the number of geometric permutations admitted by *any* collection of n pairwise disjoint convex sets in $\mathbb{R}^3$ is $O(n^3 \log n)$, thus improving Wenger's previous upper bound on $g_3(n)$ roughly by a factor of n. Our approach can be generalized to higher dimensions, and yields an improved upper bound of $O(n^{2d-3} \log n)$ on $g_d(n)$, for any $d \geq 3$. (Our bound is also a slight improvement of the bound $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$, of [14] for the case where the sets in $\mathcal{K}$ have constant description complexity.)

Here is a brief overview of our solution in $\mathbb{R}^3$. Following the approach of Wenger [19], we represent the directions of transversal lines by points on the unit 2-sphere $\mathbb{S}^2$, separate every pair of objects in $\mathcal{K}$ by a plane, and associate with each such plane the great circle on $\mathbb{S}^2$ parallel to it. We then consider the arrangement $\mathcal{A}$ of the resulting $\binom{n}{2}$ great circles on $\mathbb{S}^2$, which consists of $O(n^4)$ 2-faces. The crucial observation made in [19] is that all transversal lines, whose directions belong to the same 2-face of $\mathcal{A}$, stab the sets of $\mathcal{K}$ in the same order (if the face contains such directions at all). Hence, the number of geometric permutations is upper bounded by the total number of 2-faces of $\mathcal{A}$, implying that $g_3(n) = O(n^4)$.

We improve this bound by showing that the actual number of faces which contain at least one direction of a transversal line (so-called *permutation faces*) is only $O(n^3 \log n)$. Moreover, we show that the overall number of edges and vertices on the boundaries of these faces is also at most $O(n^3 \log n)$.

The analysis proceeds in two steps. First, we use a direct geometric analysis to show that the number of vertices whose four incident faces are all permutation faces is $O(n^3)$. We refer to such vertices as *popular vertices*. Informally, we associate with each popular vertex v (with the possible exception of $O(n^3)$ "degenerate" ones) the intersection line λ_v of the two separating planes h, h' that correspond to the two circles incident to v, and show that λ_v stabs exactly $n-4$ sets of $\mathcal{K}$ (all but the sets in the two pairs separated by h and h', respectively). We then apply, within each of the $\binom{n}{2}$ separating planes, the linear bound of [5] on the number of geometric permutations in $\mathbb{R}^2$, combined with a simple application of the Clarkson-Shor probabilistic analysis technique [4], and thereby obtain the overall $O(n^3)$ asserted bound on the number of popular vertices.

We then use this bound to analyze the overall number of vertices incident to permutation faces. This is achieved by a refined (and simplified) variant of the charging scheme of Tagansky [18].

The analysis can be extended to any dimension $d \geq 4$, but its technical details become somewhat more involved.

The paper is organized as follows. We first derive the nearly-cubic upper bound on $g_3(n)$. To this end, we begin in Section II by introducing some notations and the infrastructure, and then establish this bound in Section III. In Section IV, we extend the analysis to any dimension $d \geq 4$.

II. Preliminaries

The setup in $\mathbb{R}^3$. Let $\mathcal{K} = \{K_1, \ldots, K_n\}$ be a collection of n arbitrary pairwise disjoint convex sets in $\mathbb{R}^3$. We may also assume, without loss of generality, that the elements of $\mathcal{K}$ are *compact*; see the full version [15] for the easy argument. For each $1 \leq i < j \leq n$ we fix some plane h_{ij} which strictly separates K_i and K_j. We orient h_{ij} so that K_i lies in the open negative halfspace h_{ij}^- that it bounds, and K_j lies in the open positive halfspace h_{ij}^+. We represent directions of (oriented) lines in $\mathbb{R}^3$ by points on the unit 2-sphere $\mathbb{S}^2$. Without loss of generality we may assume that the planes h_{ij} are in *general position*, meaning that every triple of them intersect at a single point, and no four meet at a common point.

Each separating plane h_{ij} induces a great circle C_{ij} on $\mathbb{S}^2$, formed by the intersection of $\mathbb{S}^2$ with the plane parallel to h_{ij} through the origin. C_{ij}^+ (resp., C_{ij}^-) consists of the directions of lines which cross h_{ij} from h_{ij}^- to h_{ij}^+ (resp., from h_{ij}^+ to h_{ij}^-). Note that lines whose directions lie in C_{ij} cannot stab both K_i and K_j. Thus, any oriented common transversal line of K_i and K_j intersects K_j after (resp., before) K_i if and only if its direction lies in C_{ij}^+ (resp., C_{ij}^-).

Put $\mathcal{C}(\mathcal{K}) = \{C_{ij} \mid 1 \leq i < j \leq n\}$, and consider the arrangement $\mathcal{A}(\mathcal{K})$ of the $\binom{n}{2}$ great circles of $\mathcal{C}(\mathcal{K})$. The assumption that the planes h_{ij} are in general position is easily seen to imply that the circles in $\mathcal{C}(\mathcal{K})$ are also in general position, in the sense that no pair of them coincide and no three have a common point. Each 2-face f of $\mathcal{A}(\mathcal{K})$ induces a relation $\prec_f$ on $\mathcal{K}$, in which $K_i \prec_f K_j$ (resp., $K_j \prec_f K_i$) if $f \subseteq C_{ij}^+$ (resp., $f \subseteq C_{ij}^-$). Clearly, the direction of each oriented line transversal λ of $\mathcal{K}$ belongs to the unique 2-face f of $\mathcal{A}(\mathcal{K})$ whose relation $\prec_f$ coincides with the order in which λ visits the sets of $\mathcal{K}$. In particular, the number of geometric permutations is bounded by the number of 2-faces of $\mathcal{A}(\mathcal{K})$, which is $O(n^4)$.

This is the way in which Wenger established this upper bound (in three dimensions) 20 years ago [19]. Moreover, this approach can be extended to any dimension $d \geq 3$, and yields the upper bound $O(n^{2d-2})$ on $g_d(n)$; see [19] and Section IV below. The main weakness of this argument (as follows from the analysis in this paper) is that most faces of $\mathcal{A}(\mathcal{K})$ do not induce a geometric permutation of $\mathcal{K}$. Specifically, for some faces f the relation $\prec_f$ might have cycles, in which case f clearly cannot contain the direction of a transversal of $\mathcal{K}$.

More definitions. We need a few more notations. We call a 2-face of $\mathcal{A}(\mathcal{K})$ a *permutation face* if there is at least one line transversal of $\mathcal{K}$ whose direction belongs to f. Note, however, that the directions of the line transversals of $\mathcal{K}$ within a fixed permutation face f is only a subset of f, which need not even be connected; see, e.g., a construction in [6] and the introduction.

Each pair of great circles of $\mathcal{C}(\mathcal{K})$ intersect at exactly two antipodal points of $\mathbb{S}^2$. By the general position assumption, all the circles are distinct, and each vertex v of $\mathcal{A}(\mathcal{K})$ is incident to exactly two great circles. Hence, each vertex is incident to exactly four (distinct) faces of $\mathcal{A}(\mathcal{K})$. This implies that (for $|\mathcal{K}| \geq 3$) the number of permutation faces in $\mathcal{A}(\mathcal{K})$ is at most proportional to the overall number of their vertices. It is this latter quantity that we proceed to bound.

We say that vertex v in $\mathcal{A}(\mathcal{K})$ is *regular* if the two great circles $C_{ij}, C_{k\ell}$ incident to v are defined by four *distinct* sets of $\mathcal{K}$; otherwise, when only three of the indices i, j, k, ℓ are distinct, we call v a *degenerate* vertex. Clearly, the number of degenerate vertices is $O(n^3)$, so it suffices to bound the number of regular vertices of permutation faces.

In the forthcoming analysis we will use subcollections $\mathcal{K}'$ of $\mathcal{K}$, typically obtained by removing one set, say K_q, from $\mathcal{K}$. Doing so eliminates all separating planes h_{iq}, for $i = 1, \ldots, q-1$, and h_{qi}, for $i = q+1, \ldots, n$. Accordingly, the corresponding circles C_{iq}, C_{qi} are also eliminated from $\mathcal{C}(\mathcal{K}')$, and $\mathcal{A}(\mathcal{K}')$ is constructed only from the remaining circles. In particular, a regular vertex v of $\mathcal{A}(\mathcal{K})$, formed by the intersection of C_{ij} and $C_{k\ell}$, remains a vertex of $\mathcal{A}(\mathcal{K}')$ if and only if $q \neq i, j, k, \ell$. An edge (resp., face) of $\mathcal{A}(\mathcal{K}')$ may contain several edges (resp., faces) of $\mathcal{A}(\mathcal{K})$. Note that if f' is a face of $\mathcal{A}(\mathcal{K}')$ which contains a permutation face f of $\mathcal{A}(\mathcal{K})$ then f' is a permutation face in $\mathcal{A}(\mathcal{K}')$; the permutation that it induces is the permutation of f with K_q removed.

III. The Number of Geometric Permutations in $\mathbb{R}^3$

Popular vertices and edges. We say that an edge e of $\mathcal{A}(\mathcal{K})$ is *popular* if its two incident faces are both permutation faces. We say that a vertex v of $\mathcal{A}(\mathcal{K})$ is *popular* if its four incident faces are all permutation faces. We establish the upper bound $O(n^3)$ on the number of popular vertices, using a direct geometric argument. The analysis then proceeds by applying two charging schemes. The first scheme results in a recurrence which expresses the number of popular edges in terms of the number of popular vertices. The second scheme leads to a recurrence which expresses the number of vertices of permutation faces in terms of the number of popular edges. The solutions of both recurrences are nearly cubic. Naive (and simpler) implementation of both schemes incurs an extra logarithmic factor in each recurrence, resulting in the overall bound $g_3(n) = O(n^3 \log^2 n)$. With a more careful analysis of the second scheme, we are able to eliminate one of these factors, and thus obtain the bound $g_3(n) = O(n^3 \log n)$.

A. The number of popular vertices

For a regular vertex v of $\mathcal{A}(\mathcal{K})$, formed by the intersection of $C_{ij}, C_{k\ell} \in \mathcal{C}(\mathcal{K})$, we denote by $\mathcal{K}_v$ the collection $\{K_i, K_j, K_k, K_\ell\}$ of the four sets defining (the circles meeting at) v.

Lemma III.1. *Let v be a regular popular vertex of $\mathcal{A}(\mathcal{K})$, incident to $C_{ij}, C_{k\ell} \in \mathcal{C}(\mathcal{K})$.*
(i) Each pair of sets $K_a \in \mathcal{K}_v$ and $K_b \in \mathcal{K} \setminus \mathcal{K}_v$ appear in the same order in all four permutations induced by the faces incident to v.
(ii) The elements of each pair K_i, K_j and K_k, K_ℓ are consecutive in all four permutations induced by the faces incident to v.

Proof: Any two distinct faces f, g incident to v are separated only by one or two great circles from $\{C_{ij}, C_{k\ell}\}$, so the orders $\prec_f$ and $\prec_g$ may disagree only over the pairs (K_i, K_j) and (K_k, K_ℓ). As a matter of fact, the four permutations are obtained from each other only by swapping K_i and K_j and/or swapping K_k and K_ℓ. This is easily seen to imply both parts of the lemma. ∎

Lemma III.2. *Let v be a regular popular vertex in $\mathcal{A}(\mathcal{K})$, incident to $C_{ij}, C_{k\ell} \in \mathcal{C}(\mathcal{K})$. Then the line $\lambda_v = h_{ij} \cap h_{k\ell}$ stabs all the $n-4$ sets in $\mathcal{K} \setminus \mathcal{K}_v$, and misses all four sets in $\mathcal{K}_v$.*

Proof: By definition, λ_v misses every set $K \in \mathcal{K}_v$, because it is contained in a plane separating K from another set in $\mathcal{K}_v$. Hence, it suffices to show that λ_v is a transversal of $\mathcal{K} \setminus \mathcal{K}_v$.

To show this, we fix a set $K_a \in \mathcal{K} \setminus \mathcal{K}_v$ and show that each of the four dihedral wedges determined by h_{ij} and $h_{k\ell}$ meets K_a. The convexity of K_a then implies that λ_v intersects K_a; see Figure 1 (left).

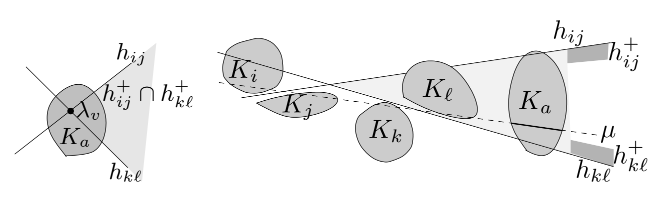

Figure 1. Left: K_a must cross $\lambda_v = h_{ij} \cap h_{k\ell}$ since it meets each of the four incident wedges (one of which is highlighted). Right: The transversal line μ crosses K_a after K_i, K_j, K_k, K_ℓ, so the segment $K_a \cap \mu$ (highlighted) is contained in $h^+_{ij} \cap h^+_{k\ell}$.

Lemma III.1 implies that K_a lies at the same position in each of the four permutations induced by the faces incident to v. Without loss of generality, assume that the consecutive pair K_i, K_j appears in these permutations before the consecutive pair K_k, K_ℓ. Then either K_a precedes both pairs in all four permutations, or appears in between them, or succeeds both of them. In what follows we assume that K_a succeeds both pairs in all the permutations, but similar arguments handle the other two cases too.

Consider the permutation π_1 induced by the face f_1 incident to v and lying in $C_{ij}^+ \cap C_{k\ell}^+$, and let μ be a line transversal which induces π_1. Since the direction of μ lies in $C_{ij}^+ \cap C_{k\ell}^+$, it follows that μ crosses h_{ij} from the side containing K_i to the side containing K_j, and it crosses $h_{k\ell}$ from the side containing K_k to the side containing K_ℓ. Hence K_i precedes K_j and K_k precedes K_ℓ in π_1. Moreover, μ crosses h_{ij} in between its intersections with K_i and K_j, and it crosses $h_{k\ell}$ in between its intersections with K_k and K_ℓ. Thus, $\mu \cap K_a$ lies in $h_{ij}^+ \cap h_{k\ell}^+$; see Figure 1 (right). That is, K_a intersects the dihedral wedge $h_{ij}^+ \cap h_{k\ell}^+$. Fully symmetric arguments, applied to the permutations induced by the three other faces f_2, f_3, f_4 incident to v, show that K_a intersects each of the three other dihedral wedges determined by h_{ij} and $h_{k\ell}$, which, as argued above, implies that λ_v stabs K_a. Slightly modified variants of this argument (with different correspondences between the wedges around λ_v and the faces around v) handle the cases where K_a precedes both pairs K_i, K_j and K_k, K_ℓ in all four permutations, or appears in between these pairs. ■

Theorem III.3. *Let $\mathcal{K}$ be a collection of n pairwise disjoint compact convex sets in $\mathbb{R}^3$. Then the number of popular vertices in $\mathcal{A}(\mathcal{K})$ is $O(n^3)$.*

Proof: As is easily seen, each popular vertex must be regular; see the full version [15] for more details. Let v be a regular popular vertex in $\mathcal{A}(\mathcal{K})$, incident to $C_{ij}, C_{k\ell} \in \mathcal{C}(\mathcal{K})$, and let $\lambda_v = h_{ij} \cap h_{k\ell}$ be the line considered in Lemma III.2. Put $K_q^* = K_q \cap h_{ij}$, for each index $q \neq i, j$, and denote by $\mathcal{K}^*$ the collection of these $n-2$ planar cross-sections within h_{ij}. Clearly, all sets in $\mathcal{K}^*$ are pairwise disjoint, compact, and convex.

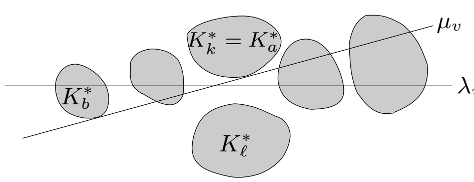

Figure 2. View inside h_{ij}: The line $\lambda_v = h_{ij} \cap h_{k\ell}$ misses K_k^*, K_ℓ^* but stabs all other sets in $\mathcal{K}^*$. The line μ_v is tangent to $K_a^* = K_k^*$ and to K_b^*, so it misses only K_ℓ^*.

By Lemma III.2, λ_v lies in h_{ij}, stabs all the sets in $\mathcal{K}^* \setminus \{K_k^*, K_\ell^*\}$ (so they are all nonempty) and misses the two sets K_k^*, K_ℓ^*. (As can be easily verified, both of K_k^*, K_ℓ^* are also nonempty, although our analysis does not rely on this property.)

Translate λ_v within h_{ij} until it becomes tangent to some set $K_a^* \in \mathcal{K}^*$, and then rotate the resulting line around K_a^*, say counterclockwise, keeping it tangent to that set, until it becomes tangent to another set $K_b^* \in \mathcal{K}^* \setminus \{K_a^*\}$. The sets $K_k^*, K_\ell^*, K_a^*, K_b^*$ need not all be distinct, so the resulting extremal tangent μ_v misses *at most* two sets of $\mathcal{K}^*$ and intersects all the other sets; see Figure 2.

We charge λ_v to μ_v, and argue that each extremal line μ in h_{ij}, which is tangent to two sets of $\mathcal{K}^*$ and misses at most two other sets of $\mathcal{K}^*$, is charged in this manner at most twice. Indeed, by the general position assumption, μ lies in a single plane h_{ij}. Within that plane, if μ misses two sets of $\mathcal{K}^*$ then these must be the sets K_k^*, K_ℓ^*. If μ misses only one set of $\mathcal{K}^*$ then this set must be one of the sets K_k^*, K_ℓ^*, and the other set is one of the two sets μ is tangent to. Finally, if μ does not miss any set of $\mathcal{K}^*$ then K_k^*, K_ℓ^* are the two sets μ is tangent to. Hence μ determines at most two quadruples K_i, K_j, K_k, K_ℓ whose lines λ_v can charge μ, and the claim follows.

It therefore suffices to bound the number of extremal lines μ charged in this manner. This can be done using the Clarkson-Shor technique [4], by observing that each such line μ is defined by two sets of $\mathcal{K}^*$ (those it is tangent to; any such pair of sets determine four common tangents) and is "in conflict" with at most two other sets (those that it misses). Thus, the Clarkson-Shor technique implies that the number of lines μ_v is $O\left(L_0(n/2)\right)$, where $L_0(r)$ is the (expected) number of extremal lines which are transversals to a (random) sample of r sets of $\mathcal{K}^*$. Edelsbrunner and Sharir [5] establish an upper bound of $O(r)$ on the complexity of the space of line transversals to a collection of r pairwise-disjoint compact convex sets in the plane, implying that $L_0(r) = O(r)$. Hence the number of charged lines μ in a single plane h_{ij} is $O(n)$, for a total of $O\left(\binom{n}{2} \cdot n\right) = O(n^3)$. Since, as noted above, each line is charged at most twice in its plane, this also bounds the number of popular vertices. ■

B. The number of popular edges

Define an *edge border* in $\mathcal{A}(\mathcal{K})$ to be a pair (v, Q), where v is a vertex of $\mathcal{A}(\mathcal{K})$, incident to two great circles $C_{ij}, C_{k\ell}$, and Q is one of the four open hemispheres $C_{ij}^-, C_{ij}^+, C_{k\ell}^-, C_{k\ell}^+$ determined by one of these circles. See Figure 3 (left). Note that Q determines a unique edge e of $\mathcal{A}(\mathcal{K})$ which is incident to v and is contained in Q. If, in addition, e is a popular edge, we say that (v, Q) is a *popular edge border*. For the purpose of the analysis, we will also refer to (v, Q) as a 0-*level edge border*.

One useful feature of the border notation is that if (v, Q) is an edge border in $\mathcal{A}(\mathcal{K})$ and $\mathcal{K}'$ is a subcollection of $\mathcal{K}$ so that v is still a vertex of $\mathcal{A}(\mathcal{K}')$, then (v, Q) is also an edge border in $\mathcal{A}(\mathcal{K}')$. The edge e' of $\mathcal{A}(\mathcal{K}')$ associated with (v, Q) in $\mathcal{A}(\mathcal{K}')$ either is equal to e, or strictly contains e (in the latter case both e and e' have v as a common endpoint).

If an edge border (v, Q), which is not a 0-level edge border, becomes a 0-level edge border after removing from $\mathcal{K}$ some single set $K_a \in \mathcal{K}$, we call it a 1-*level edge border*. In this case we say that (v, Q) is *in conflict* with K_a. Note

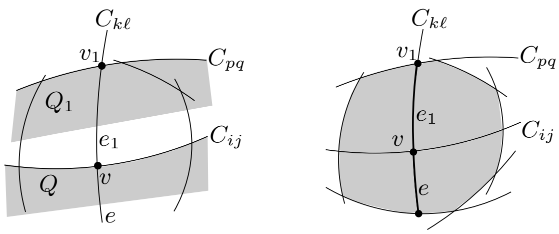

Figure 3. Left: Charging a 0-level edge border (v, Q) to a 1-level edge border (v_1, Q_1). Right: If the edges e, e_1 are both popular then v is a popular vertex.

that the set K_a, whose removal makes (v, Q) a 0-level edge border, need not be unique.

Clearly, to bound the number of popular edges it suffices to bound the number of 0-level edge borders, which is twice the number of popular edges in $\mathcal{A}(\mathcal{K})$ (each edge is counted once at each of its endpoints).

Since each vertex of $\mathcal{A}(\mathcal{K})$ participates in exactly four edge borders, the number of edge borders which are incident to a degenerate vertex is $O(n^3)$. We bound the number of remaining 0-level edge borders using the following charging scheme.

Let (v, Q) be a 0-level edge border, where v is incident to C_{ij} and $C_{k\ell}$, so that $Q = C_{ij}^+$, say. Let e be the popular edge associated with (v, Q). Trace $C_{k\ell}$ from v away from e (into C_{ij}^-), and let v_1 be the next encountered vertex. Let e_1 be the edge connecting v and v_1. Let C_{pq} be the other circle incident to v_1 and assume, without loss of generality, that v lies in C_{pq}^+. See Figure 3 (left). Note that, assuming $|\mathcal{K}| \geq 3$, we have $C_{pq} \neq C_{ij}$ (i.e., v_1 is not antipodal to v), because otherwise C_{ij} would have intersected only $C_{k\ell}$. One of the following cases must arise:

(i) v_1 is degenerate.

(ii) The edge e_1 is also popular, so v is a popular vertex; see Figure 3 (right).

(iii) e_1 is not popular. Since $C_{pq} \neq C_{ij}$, one of i, j, say i, is different from both p and q. This (and the fact that $i \neq k, \ell$) implies that removing K_i from $\mathcal{K}$ also removes C_{ij} from $\mathcal{A}$, keeps v_1 intact, and makes the appropriate extension of e reach (and terminate at) v_1, thereby making (v_1, Q_1) a 0-level edge border in $\mathcal{A}(\mathcal{K} \setminus \{K_i\})$, where $Q_1 = C_{pq}^+$. See Figure 3 (left).

In case (i) we charge (v, Q) to v_1. The number of degenerate vertices is $O(n^3)$ and each of them can be charged only $O(1)$ times in this manner. Hence, the number of 0-level edge borders that fall into this subcase is $O(n^3)$.

In case (ii) we can charge (v, Q) to v. Since a popular vertex participates in exactly four 0-level edge borders, the number of 0-level edge borders that fall into this subcase is $O(n^3)$, by Theorem III.3.

In case (iii) we charge (v, Q) to the 1-level edge border (v_1, Q_1). Note that (v_1, Q_1) is charged in this manner only by (v, Q).

Let us denote by $E_0(\mathcal{K})$ (resp., $E_1(\mathcal{K})$) the number of 0-level edge borders (resp., 1-level edge borders) in $\mathcal{A}(\mathcal{K})$. Then we have the following recurrence:

$$E_0(\mathcal{K}) \leq E_1(\mathcal{K}) + O(n^3). \tag{1}$$

To solve this recurrence, we apply the technique of Tagansky [18]. Specifically, we remove from $\mathcal{K}$ a randomly chosen set $K \in \mathcal{K}$, and denote by $\mathcal{R}$ the collection of the $n-1$ remaining sets. A 0-level edge border (v, Q) in $\mathcal{A}(\mathcal{K})$, where v is an intersection point of C_{ij} and $C_{k\ell}$ and is regular, appears as a 0-level edge border in $\mathcal{A}(\mathcal{R})$ if and only if K is different from each of the four sets K_i, K_j, K_k, K_ℓ defining v, which happens with probability $\frac{n-4}{n}$. A 1-level edge border (v, Q) in $\mathcal{A}(\mathcal{K})$ becomes a 0-level edge border in $\mathcal{A}(\mathcal{R})$ if and only if K is in conflict with (v, Q), which happens with probability *at least* $\frac{1}{n}$. No other edge border in $\mathcal{A}(\mathcal{K})$ can appear as a 0-level edge border in $\mathcal{A}(\mathcal{R})$. Hence, we obtain

$$\mathbf{E}\{E_0(\mathcal{R})\} \geq \frac{n-4}{n} E_0(\mathcal{K}) + \frac{1}{n} E_1(\mathcal{K}), \tag{2}$$

where $\mathbf{E}$ denotes expectation with respect to the random sample $\mathcal{R}$, as constructed above. Combining (1) and (2) yields

$$\frac{1}{n} E_0(\mathcal{K}) \leq \frac{1}{n} E_1(\mathcal{K}) + O(n^2) \leq \mathbf{E}\{E_0(\mathcal{R})\} - \frac{n-4}{n} E_0(\mathcal{K}) + O(n^2).$$

Denoting by $E_0(n)$ the maximum number of 0-level edge borders in $\mathcal{A}(\mathcal{K})$, for any collection $\mathcal{K}$ of size n with the assumed properties, we get the recurrence

$$\frac{n-3}{n} E_0(n) \leq E_0(n-1) + O(n^2),$$

whose solution is easily seen to be $E_0(n) = O(n^3 \log n)$.

C. The number of permutation faces

Finally, we bound the number of vertices of permutation faces using the bound on popular edges just derived. This will also serve as an upper bound on the number of permutation faces, and thus also on $g_3(n)$. We present the analysis in two stages. The first stage derives the slightly weaker upper bound $O(n^3 \log^2 n)$, but is considerably simpler. The second stage involves a more careful examination of the possible charging scenarios, and leads to a sharper recurrence, whose soution is only $O(n^3 \log n)$.

Each vertex v is incident to exactly four faces of $\mathcal{A}(\mathcal{K})$, so we need to count v with multiplicity of at most 4—once for each permutation face incident to v. For this we extend the notion of borders as follows. The two great circles passing through v partition $\mathbb{S}^2$ into four wedges, or rather slices. Each such slice R contains a unique face f incident to v, and defines, together with v, a *border* (v, R). We call f the face *associated* with (v, R). We call (v, R) a *popular border*, or a 0-*level border*, if the face associated with (v, R) is a permutation face. It thus suffices to bound the number of 0-level borders in $\mathcal{A}(\mathcal{K})$.

If (v, R) is a border in $\mathcal{A}(\mathcal{K})$ with an associated face f, and $\mathcal{K}'$ is a subcollection of $\mathcal{K}$, so that v is still a vertex of $\mathcal{A}(\mathcal{K}')$, then (v, R) is also a border in $\mathcal{A}(\mathcal{K}')$, except that the face f' of $\mathcal{A}(\mathcal{K}')$ associated with (v, R) may be different from f (or, more precisely, properly contain f).

If a border (v, R), which is not a 0-level border in $\mathcal{A}(\mathcal{K})$, becomes a 0-level border after removing from $\mathcal{K}$ some set K, we call it a 1-*level border*. The set K is said to be *in conflict* with (v, R). Note that K cannot be one of the (at most) four sets defining v, and that a 1-level border may be in conflict with more than one set of $\mathcal{K}$. See Figure 4 (left).

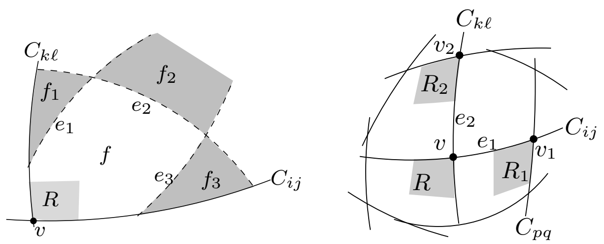

Figure 4. Left: A non-permutation face f, associated with the 1-level border (v, R), is separated from permutation faces f_1, f_2, f_3 by the respective edges $e_1 \subset C_{p_1q_1}, e_2 \subset C_{p_2q_2}, e_3 \subset C_{p_3q_3}$. If none of $p_1, q_1, p_2, q_2, p_3, q_3$ belongs to $\{i, j, k, \ell\}$ then (v, R) is a 1-level border in conflict with each of $K_{p_1}, K_{q_1}, K_{p_2}, K_{q_2}, K_{p_3}, K_{q_3}$. Right: Charging a 0-level border (v, R) to the two 1-level borders (v_1, R_1), (v_2, R_2), along the two edges e_1, e_2 emanating from v away from R.

We bound the number of 0-level borders using a charging scheme similar to that in Section III-B. Let (v, R) be a 0-level border, and let f be the permutation face associated with it. Note that the number of borders incident to degenerate vertices is $O(n^3)$. We may therefore assume that v is regular, and let C_{ij} and $C_{k\ell}$ denote the two great circles incident to v (so i, j, k, ℓ are all distinct). Without loss of generality, assume that $R = C_{ij}^+ \cap C_{k\ell}^+$.

Let e_1 and e_2 be the two edges incident to v and emanating from it away from R, where $e_1 \subset C_{ij} \cap C_{k\ell}^-$ and $e_2 \subset C_{k\ell} \cap C_{ij}^-$; see Figure 4 (right). Let v_1 (resp., v_2) be the other endpoint of e_1 (resp., of e_2).

Our charging scheme is based on the following case analysis:

(i) If one of the two edges incident to v and bounding R is popular, we charge (v, R) to this edge. Since the number of popular edges is $O(n^3 \log n)$, and each of them is charged by at most four 0-level borders (twice for each of its endpoints), the number of 0-level borders that fall into this subcase is also $O(n^3 \log n)$.

(ii) If no edge incident to v and bounding R is popular, we charge (v, R) to two 1-level borders, one incident to v_1 and one to v_2. Specifically, consider v_1, say, and let C_{pq} be the circle whose intersection with C_{ij} forms v_1, and assume, again without loss of generality, that v lies in C_{pq}^+. We then charge (v, R) to (v_1, R_1), where $R_1 = C_{pq}^+ \cap C_{ij}^+$. Let f_1 be the face of $\mathcal{A}(\mathcal{K})$ associated with (v_1, R_1) (this is the face whose boundary we trace from v to v_1 along e_1, and it is also incident to v). Since the edge incident to f, f_1 (and to v) is not popular, f_1 is not a permutation face. Clearly, one of the indices k, ℓ, say k, is different from both p, q. Thus, removing K_k keeps v_1 as a vertex in the new spherical arrangement, and makes $C_{k\ell}$ disappear, so both faces f, f_1 fuse into a single larger permutation face contained in R_1. Hence, (v_1, R_1) is a 1-level border which is in conflict with K_k. A fully symmmetic argument applies to v_2. We say that the 1-level borders (v_1, R_1) and (v_2, R_2), which we charge, are the *neighbors* of (v, R) in $\mathcal{A}(\mathcal{K})$.

Note that each 1-level border (v', R') is charged by at most two 0-level borders in this manner (at most once along each of the two edges incident to v' and bounding the face associated with the border).

Let $V_0(\mathcal{K})$ and $V_1(\mathcal{K})$ denote, respectively, the number of 0-level borders and the number of 1-level borders in $\mathcal{A}(\mathcal{K})$ (where we also include degenerate vertices in both counts). Then we have the following recurrence:

$$V_0(\mathcal{K}) \le V_1(\mathcal{K}) + O(n^3 \log n). \tag{3}$$

Indeed, each 0-level border which falls into case (ii) charges two 1-level borders, and each 1-level border is charged at most twice. The number of all other 0-level borders is $O(n^3 \log n)$, as argued above. Combining this inequality with the random sampling technique of Tagansky [18], as in Section III-B, results in the recurrence

$$\frac{n-3}{n} V_0(n) \le V_0(n-1) + O(n^2 \log n),$$

where $V_0(n)$ is the maximum value of $V_0(\mathcal{K})$, over all collections $\mathcal{K}$ of n pairwise disjoint compact convex sets in $\mathbb{R}^3$. The solution of this recurrence is $V_0(n) = O(n^3 \log^2 n)$, which yields the same upper bound on the number of geometric permutations induced by $\mathcal{K}$.

An improved bound. We next improve the bound by replacing the recurrence (3) by a refined recurrence. Let (v, R) be a 1-level border which is in conflict with $w \ge 1$ sets of $\mathcal{K}$. Then (v, R) becomes a 0-level border in $\mathcal{A}(K \setminus \{K\})$, after removing a random set $K \in \mathcal{K}$, with probability exactly $\frac{w}{n}$. Namely, this happens if and only if K is one of the w sets in conflict with (v, R). We refer to w as the *weight* of (v, R).

In the refined setting, $V_1(\mathcal{K})$ counts the total weight of all the 1-level borders in $\mathcal{A}(\mathcal{K})$, so now the contribution of each 1-level border to $V_1(\mathcal{K})$ is equal to its weight. By an appropriate adaptation of the argument in Section III-B, we obtain the following *equality*:

$$\mathbf{E}\{V_0(\mathcal{R})\} = \frac{n-4}{n} V_0(\mathcal{K}) + \frac{1}{n} V_1(\mathcal{K}), \tag{4}$$

where $\mathcal{R}$ denotes a random sample of $n-1$ sets of $\mathcal{K}$. This follows by noting that the probability of a 1-level border of weight w to be counted in $V_0(\mathcal{R})$ is $\frac{w}{n}$, and it contributes w to $V_1(\mathcal{K})$.

In the refined charging scheme, each 1-level border (v, R) of weight $w \geq 1$ gets a supply of w units of charge, which it can give to its charging neighboring 0-level borders. Hence, as long as the number of these charging 0-level borders, which is at most two, does not exceed w, (v, R) can pay each of its neighbors 1 unit. Hence, the only problematic case is when $w = 1$ and (v, R) is charged twice. The following technical lemma (whose somewhat involved proof appears in the full version [15]) takes care of this case.

Lemma III.4. *The number of* 1*-level borders having weight* 1 *and charged by two* 0*-level borders is* $O(n^3 \log n)$.

If a 1-level border (v, R) has only one neighboring 0-level border (v', R') then (v', R') can receive one unit of charge from (v, R), regardless of what the weight of (v, R) is. Similarly, if (v, R) has weight at least 2, and it has two neighboring 0-level borders, each of these 0-level borders can receive one unit of charge from (v, R). The number of remaining 1-level borders, namely the 1-level borders of weight 1 with two neighboring 0-level borders, is $O(n^3 \log n)$, by Lemma III.4.

This implies

$$2V_0(\mathcal{K}) \leq V_1(\mathcal{K}) + O(n^3 \log n).$$

Combining this with (4) we get

$$\frac{2}{n} V_0(\mathcal{K}) \leq \frac{1}{n} V_1(\mathcal{K}) + O(n^2 \log n)$$

$$\leq \mathbf{E}\{V_0(\mathcal{R})\} - \frac{n-4}{n} V_0(\mathcal{K}) + O(n^2 \log n),$$

or
$$\frac{n-2}{n} V_0(\mathcal{K}) \leq \mathbf{E}\{V_0(\mathcal{R})\} + O(n^2 \log n).$$

Replacing $V_0(\mathcal{K})$, $V_0(\mathcal{R})$ by their respective maximum values $V_0(n)$, $V_0(n-1)$, we thus obtain the recurrence

$$\frac{n-2}{n} V_0(n) \leq V_0(n-1) + O(n^2 \log n),$$

whose solution is easily seen to be $V_0(n) = O(n^3 \log n)$. That is, we have:

Theorem III.5. *Any collection* $\mathcal{K}$ *of* n *pairwise disjoint convex sets in* $\mathbb{R}^3$ *admits at most* $O(n^3 \log n)$ *geometric permutations.*

IV. Geometric Permutations in Higher Dimensions

In this section we generalize Theorem III.5 by showing that the number of geometric permutations induced by any collection $\mathcal{K} = \{K_1, \ldots, K_n\}$ of n pairwise disjoint convex sets in $\mathbb{R}^d$ is $O(n^{2d-3} \log n)$, for any $d \geq 3$.

Setup. As in the three-dimensional case, we may assume that the sets of $\mathcal{K}$ are also compact. For each $1 \leq i < j \leq n$ we fix some hyperplane h_{ij} which strictly separates K_i and K_j. We orient h_{ij} so that K_i lies in the negative open halfspace h_{ij}^- that it bounds, and K_j lies in the positive open halfspace h_{ij}^+. We represent directions of lines in $\mathbb{R}^d$ by points on the unit $(d-1)$-sphere $\mathbb{S}^{d-1}$. We may assume that the separating hyperplanes h_{ij} are in *general position*, so that every d of them intersect in a unique point, and no $d+1$ of them have a point in common.

Each separating hyperplane h_{ij} induces a great $(d-2)$-sphere C_{ij} on $\mathbb{S}^{d-1}$, which is the locus of the directions of all lines parallel to h_{ij}. C_{ij} partitions $\mathbb{S}^{d-1}$ into two open hemispheres C_{ij}^+, C_{ij}^-, so that C_{ij}^+ (resp., C_{ij}^-), consists of the directions of lines which cross h_{ij} from h_{ij}^- to h_{ij}^+ (resp., from h_{ij}^+ to h_{ij}^-). Any oriented common transversal line of K_i and K_j visits K_j after (resp., before) K_i if and only if its direction lies in C_{ij}^+ (resp., in C_{ij}^-).

Put $\mathcal{C}(\mathcal{K}) = \{C_{ij} \mid 1 \leq i < j \leq n\}$, and consider the arrangement $\mathcal{A}(\mathcal{K})$ of these $\binom{n}{2}$ $(d-2)$-spheres on $\mathbb{S}^{d-1}$. It partitions $\mathbb{S}^{d-1}$ into relatively open cells of dimensions $0, 1, \ldots, d-1$; we refer to an s-dimensional cell of $\mathcal{A}(\mathcal{K})$ simply as an s-cell. Since the hyperplanes h_{ij} are in general position, the $(d-2)$-spheres of $\mathcal{C}(\mathcal{K})$ are also in general position, in the sense that the intersection of any s distinct spheres of $\mathcal{C}(\mathcal{K})$, for $1 \leq s \leq d-1$, is a $(d-s-1)$-sphere, and the intersection of any d distinct spheres of $\mathcal{C}(\mathcal{K})$ is empty. Each $(d-1)$-cell f of $\mathcal{A}(\mathcal{K})$ induces a relation $\prec_f$ on $\mathcal{K}$, in which $K_i \prec_f K_j$ (resp., $K_j \prec_f K_i$) if $f \subseteq C_{ij}^+$ (resp., $f \subseteq C_{ij}^-$). The direction of each oriented line transversal λ of $\mathcal{K}$ belongs to the unique $(d-1)$-cell f of $\mathcal{A}(\mathcal{K})$ whose relation $\prec_f$ coincides with the linear order in which λ visits the sets of $\mathcal{K}$. In particular, as noted by Wenger [19], the number of geometric permutations is bounded by the number of $(d-1)$-cells of $\mathcal{A}(\mathcal{K})$, which is $O(n^{2d-2})$.

We call a $(d-1)$-cell f of $\mathcal{A}(\mathcal{K})$ a *permutation cell* if there is at least one line transversal of $\mathcal{K}$ whose direction belongs to f. As in the three-dimensional case, we improve the above bound by showing that the number of permutation cells in $\mathcal{A}(\mathcal{K})$ is $O(n^{2d-3} \log n)$, which also bounds the number of geometric permutations induced by $\mathcal{K}$.

We refer to 0-cells in $\mathcal{A}(\mathcal{K})$ as *vertices*, and to 1-cells as *edges*. We say that a vertex v of $\mathcal{A}(\mathcal{K})$ is *regular* if the $d-1$ $(d-2)$-spheres of $\mathcal{C}(\mathcal{K})$ that are incident to v are defined by $2d-2$ distinct sets of $\mathcal{K}$; otherwise v is a *degenerate* vertex. Clearly, the number of degenerate vertices is $O(n^{2d-3})$, so it suffices to bound the number of regular vertices of permutation cells.

As in the three-dimensional case, we will also consider subcollections $\mathcal{K}'$ of $\mathcal{K}$, typically obtained by removing one set, say K_q, from $\mathcal{K}$. Doing so eliminates all separating hyperplanes h_{iq}, h_{qi}, as well as all the corresponding $(d-2)$-spheres C_{iq}, C_{qi}, and $\mathcal{A}(\mathcal{K}')$ is constructed only from the remaining spheres. In particular, a vertex[1] v of the intersection $C_{i_1 j_1} \cap C_{i_2 j_2} \cap \cdots \cap C_{i_{d-1} j_{d-1}}$ of $\mathcal{A}(\mathcal{K})$ remains a vertex of $\mathcal{A}(\mathcal{K}')$ if and only if $q \notin \{i_1, j_1, \ldots, i_{d-1}, j_{d-1}\}$.

[1]As in the three-dimensional case, the intersection consists of two antipodal points, so there are two choices for v.

A cell of $\mathcal{A}(\mathcal{K}')$, of any dimension $s \geq 1$, may contain several cells of $\mathcal{A}(\mathcal{K})$. As before, if f' is a $(d-1)$-cell of $\mathcal{A}(\mathcal{K}')$ which contains a permutation cell f of $\mathcal{A}(\mathcal{K})$ then f' is a permutation cell in $\mathcal{A}(\mathcal{K}')$; the permutation that it induces is the permutation of f with K_q removed.

Each s-cell f of $\mathcal{A}(\mathcal{K})$ is incident to 2^{d-s-1} $(d-1)$-cells of $\mathcal{A}(\mathcal{K})$. If all these cells are permutation cells, f is called *popular*. In particular, a popular vertex (resp., edge) is incident to 2^{d-1} (resp., 2^{d-2}) permutation cells, and a popular $(d-1)$-cell is a permutation cell.

Overview. We show that the number of popular vertices is $O(n^{2d-3})$ by a straightforward generalization of the analysis in Section III-A. The analysis then proceeds by applying, for each $1 \leq s \leq d-1$, a charging scheme, which expresses the number of popular s-cells in terms of the number of popular $(s-1)$-cells (and degenerate vertices). A naive charging scheme produces a recurrence whose solution incurs an additional logarithmic factor for each s, resulting in the weaker bound $g_d(n) = O(n^{2d-3}\log^{d-1} n)$. A more careful analysis, as in the three-dimensional case, leads to refined recurrences, whose solution yields the improved bound $g_d(n) = O(n^{2d-3}\log n)$. (We lose a logarithmic factor only when passing from vertices to edges, as in the three-dimensional case.)

A. The number of popular vertices

For a regular vertex $v \in \bigcap_{q=1}^{d-1} C_{i_q j_q}$ of $\mathcal{A}(\mathcal{K})$, we denote by $\mathcal{K}_v$ the collection $\{K_{i_q}, K_{j_q} \mid 1 \leq q \leq d-1\}$ of the $2d-2$ sets defining v.

Lemma IV.1. *Let $v \in \bigcap_{q=1}^{d-1} C_{i_q j_q}$ be a regular popular vertex of $\mathcal{A}(\mathcal{K})$.*
(i) Each pair of sets $K_a \in \mathcal{K}_v$ and $K_b \in \mathcal{K} \setminus \mathcal{K}_v$ appear in the same order in all the 2^{d-1} permutations induced by the $(d-1)$-cells incident to v.
(ii) The elements of each pair $K_{i_q}, K_{j_q} \in \mathcal{K}_v$, for $1 \leq q \leq d-1$, are consecutive in all these 2^{d-1} permutations.

Proof: Each pair of distinct $(d-1)$-cells f, g incident to v are separated by at most $d-1$ $(d-2)$-spheres from $\{C_{i_1 j_1}, \ldots, C_{i_{d-1} j_{d-1}}\}$, and only by these spheres. Hence the orders $\prec_f$ and $\prec_g$ may disagree only over the pairs (K_{i_q}, K_{j_q}), for $1 \leq q \leq d-1$. This is easily seen to imply both parts of the lemma. ∎

Lemma IV.2. *Let $v \in \bigcap_{q=1}^{d-1} C_{i_q j_q}$ be a regular popular vertex in $\mathcal{A}(\mathcal{K})$. Then the line $\lambda_v = \bigcap_{q=1}^{d-1} h_{i_q j_q}$ stabs all the $n-2d+2$ sets in $\mathcal{K} \setminus \mathcal{K}_v$, and misses all the $2d-2$ sets in $\mathcal{K}_v$.*

Proof: By definition, λ_v misses every set $K \in \mathcal{K}_v$, because it is contained in a hyperplane separating K from another set in $\mathcal{K}_v$. Hence, it suffices to show that λ_v is a transversal of $\mathcal{K} \setminus \mathcal{K}_v$.

To show this, we fix a set $K_a \in \mathcal{K} \setminus \mathcal{K}_v$ and show that each of the 2^{d-1} wedges determined by $\{h_{i_q j_q} \mid 1 \leq q \leq d-1\}$ meets K_a. Each of these wedges is the intersection of $d-1$ halfspaces, where the q-th halfspace is either $h^+_{i_q j_q}$ or $h^-_{i_q j_q}$, for $q = 1, \ldots, d-1$. All these wedges have λ_q on their boundary, and the convexity of K_a then implies, exactly as in the three-dimensional case, that λ_v intersects K_a.

For specificity, we show that K_a intersects the wedge $\bigcap_{q=1}^{d-1} h^+_{i_q j_q}$; the proof for the other wedges is essentially the same. Lemma IV.1 implies that K_a lies at the same position in each of the 2^{d-1} permutations induced by the cells incident to v. For each index q, if K_{i_q}, K_{j_q} appear before K_a (resp., after K_a) in all permutations induced by the cells incident to v, put $C_q = C^+_{i_q j_q}$ (resp., $C_q = C^-_{i_q j_q}$).

Let f be the cell incident to v and contained in $\bigcap_{q=1}^{d-1} C_q$, and let μ_f be a transversal line stabbing $\mathcal{K}$ in the order $\prec_f$ (so its direction lies in f). By the choice of f and by our assumption, we have either $K_{i_q} \prec_f K_{j_q} \prec_f K_a$, or $K_a \prec_f K_{j_q} \prec_f K_{i_q}$. This implies in the former case that μ_f visits K_a after crossing $h_{i_q j_q}$ from $h^-_{i_q j_q}$ (the side containing K_{i_q}) to $h^+_{i_q j_q}$ (the side containing K_{j_q}). In the latter case, μ_f first visits K_a and then crosses $h_{i_q j_q}$ from $h^+_{i_q j_q}$ to $h^-_{i_q j_q}$. Thus, in either case, the segment $\lambda_f \cap K_a$ lies in $h^+_{i_q j_q}$, and this holds for every $1 \leq q \leq d-1$. Hence $\lambda_f \cap K_a \subset \bigcap_{q=1}^{d-1} h^+_{i_q j_q}$, and the claim follows. ∎

Theorem IV.3. *Let $\mathcal{K}$ be a collection of n pairwise disjoint compact convex sets in $\mathbb{R}^d$. Then the number of popular vertices in $\mathcal{A}(\mathcal{K})$ is $O(n^{2d-3})$.*

Proof: As in the three-dimensional case, it is easily checked that a popular vertex must be regular. Let $v \in \bigcap_{q=1}^{d-1} C_{i_q j_q}$ be a (regular) popular vertex in $\mathcal{A}(\mathcal{K})$, and let $\lambda_v = \bigcap_{q=1}^{d-1} h_{i_q j_q}$ be the intersection line of the corresponding hyperplanes. Consider the plane $H = \bigcap_{q=1}^{d-2} h_{i_q j_q}$, put $K^*_a = K_a \cap H$, for each index $a \notin \{i_q, j_q \mid 1 \leq q \leq d-2\}$, and denote by $\mathcal{K}^*$ the collection of these $n-2d+4$ planar cross-sections. Clearly, all sets in $\mathcal{K}^*$ are pairwise disjoint, compact, and convex.

By Lemma IV.2, λ_v lies in H, stabs all the sets in $\mathcal{K}^* \setminus \{K^*_{i_{d-1}}, K^*_{j_{d-1}}\}$, and misses the two sets $K^*_{i_{d-1}}, K^*_{j_{d-1}}$. As in Theorem III.3, we charge λ_v to an extremal line $\mu = \mu_v$ within H which is tangent to two sets of $\mathcal{K}^*$, and misses only the sets among $K^*_{i_{d-1}}, K^*_{j_{d-1}}$ that it does not touch. As in the preceding analysis, each extremal line μ of this kind is charged at most twice. Applying the Clarkson-Shor analysis [4], similarly to Theorem III.3, the number of lines μ, charged within H, is $O(n)$. Summing over all possible choices of the 2-planes H, i.e., all choices of $d-2$ of the hyperplanes h_{ij}, the number of lines λ_v, and thus the number of popular vertices, is $O(n \cdot n^{2d-4}) = O(n^{2d-3})$. ∎

B. The number of permutation cells

We next generalize the analysis of Section III-C to higher dimensions. We first extend the notion of borders. Let v be a vertex of $\mathcal{A}(\mathcal{K})$, so that $v \in \bigcap_{1 \leq q \leq d-1} C_{i_q j_q}$. For any subset

J of $\{1, \ldots, d-1\}$, let $R \subseteq \mathbb{S}^{d-1}$ be a connected component of $\mathbb{S}^{d-1} \setminus \bigcup_{q \in J} C_{i_q j_q}$. Equivalently, it is the intersection of $|J|$ hemispheres, where the q-th hemisphere, for $q \in J$, is either $C^+_{i_q j_q}$ or $C^-_{i_q j_q}$. Note that there are $2^{|J|}$ such regions (for any fixed J). We call (v, R) an *s-border*, where $s = |J|$. Given v and R, for $s \geq 1$, there is a unique s-dimensional cell f of $\mathcal{A}(\mathcal{K})$ which is incident to v and is contained in the interior of R. This cell f lies in the intersection of the interior of R with $\bigcap_{q \in J^c} C_{i_q j_q}$, where $J^c = \{1, \ldots, d-1\} \setminus J$. We refer to f as the s-cell of $\mathcal{A}(\mathcal{K})$ associated with (v, R). For $s = 0$ we define the s-cell of $\mathcal{A}(\mathcal{K})$ associated with (v, R) to be v itself, and for $s = d-1$ we define the s-cell of $\mathcal{A}(\mathcal{K})$ associated with (v, R) to be the unique $(d-1)$-cell incident to v and contained in R. The reader is invited to check that, for $d = 3$, a 0-border, in the new definition, is a vertex of $\mathcal{A}(\mathcal{K})$, a 1-border is an edge border, and a 2-border is what we simply called a border.

Let (v, R) be an s-border and let f be the s-cell associated with (v, R). If f is a popular cell, we say that (v, R) is a *0-level s-border* of $\mathcal{A}(\mathcal{K})$. An s-border (v, R) is a *1-level s-border* in $\mathcal{A}(\mathcal{K})$ if it is not a 0-level s-border, but becomes such a border after removing from $\mathcal{K}$ some single set K. In this case we say that K is *in conflict* with (v, R). As in the three-dimensional case, K need not be unique.

For each $t = 0, 1$ and $0 \leq s \leq d-1$, let $N_t^{(s)}(\mathcal{K})$ be the number of t-level s-borders in $\mathcal{A}(\mathcal{K})$, and let $N_t^{(s)}(n)$ denote the maximum value of $N_t^{(s)}(\mathcal{K})$, over all collections $\mathcal{K}$ of n pairwise disjoint compact convex sets in $\mathbb{R}^d$.

Note that $N_0^{(0)}(\mathcal{K})$ is the number of popular vertices in $\mathcal{A}(\mathcal{K})$, so we have $N_0^{(0)}(n) = O(n^{2d-3})$. The term $N_0^{(d-1)}(\mathcal{K})$ counts the overall number of vertices incident to permutation cells, which is an upper bound on the number of cells, and thus also on the number of geometric permutations of $\mathcal{K}$.

For each $1 \leq s \leq d-1$, we apply a charging scheme, which results in a recurrence which expresses $N_0^{(s)}(\mathcal{K})$ in terms of $N_0^{(s-1)}(\mathcal{K})$ and $N_1^{(s)}(\mathcal{K})$.

Fix $1 \leq s \leq d-1$. Let (v, R) be a 0-level s-border in $\mathcal{A}(\mathcal{K})$, and let f be the popular s-cell associated with (v, R). Let $C_{i_1 j_1}, C_{i_2 j_2}, \ldots, C_{i_{d-1} j_{d-1}}$ be the $d-1$ $(d-2)$-spheres of $\mathcal{C}(\mathcal{K})$ incident to v, and assume, with no loss of generality, that $R = \bigcap_{q=1}^{s} C^+_{i_q j_q}$. Moreover, we may assume that v is regular, since the number of s-borders incident to degenerate vertices is clearly $O(n^{2d-3})$.

For each $1 \leq q \leq s$ there exists a unique edge e_q^+ of f which is incident to v and not contained in $C_{i_q j_q}$. Indeed, by construction, f lies in the intersection s-sphere $\bigcap_{q=s+1}^{d-1} C_{i_q j_q}$ (for $s = d-1$, this is the entire $\mathbb{S}^{d-1}$), and each edge of f incident to v is formed by further intersecting this sphere with $s-1$ additional spheres from $C_{i_1 j_1}, \ldots, C_{i_s j_s}$. The claim follows since only one side of the resulting intersection circle lies (near v) in the closure of R. Let e_q^- denote the other edge of $\mathcal{A}(\mathcal{K})$ which is incident to v and lies on the same intersection circle γ as e_q^+, so e_q^- emanates from v away from R. Let v_q denote the other endpoint of e_q^-, and let g_q denote the (unique) $(s-1)$-cell which bounds f, lies in $C_{i_q j_q} \cap \left(\bigcap_{t=s+1}^{d-1} C_{i_t j_t}\right)$, is incident to v and is contained in $R_q = \bigcap_{1 \leq t \leq s,\, t \neq q} C^+_{i_t j_t}$. Also, g_q is the $(s-1)$-cell associated with (v, R_q). See Figure 5. There are two possible cases:

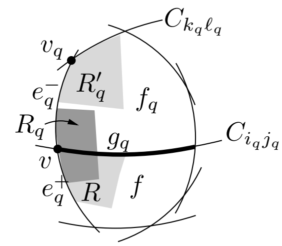

Figure 5. Charging a 0-level s-border along the edge e_q.

(i) If one of the $(s-1)$-cells $g_1, g_2, \ldots, g_s$, say g_s, is popular, we charge (v, R) to the 0-level $(s-1)$-border (v, R_s), noting, as above, that g_s is the $(s-1)$-cell associated with this border. By construction, each 0-level $(s-1)$-border (v, R) is charged at most $2(d-s)$ times in this manner, once from each s-border associated with an s-cell which is bounded by the $(s-1)$-cell associated with this border. Hence, the number of 0-level s-borders falling into subcase (i) is $O\left(N_0^{(s-1)}(\mathcal{K})\right)$.

(ii) None of the $(s-1)$-cells $g_1, g_2, \ldots, g_s$ is popular. For each $1 \leq q \leq s$, let $C_{k_q \ell_q}$ be the additional great sphere incident to v_q, and suppose, for specificity, that $v \in C^+_{k_q \ell_q}$. The vertex v_q participates in the 1-level s-border (v_q, R'_q), where $R'_q = C^+_{k_q \ell_q} \cap \left(\bigcap_{1 \leq t \leq s,\, t \neq q} C^+_{i_t j_t}\right)$.

Since g_q is not popular, (v_q, R'_q) is not a 0-level s-border. Let f_q be the s-cell associated with (v_q, R'_q). Clearly, at least one of i_q, j_q does not belong to $\{k_q, \ell_q\}$; say it is i_q. Thus, and since v is regular, removing K_{i_q} keeps v_q (and hence (v_q, R'_q)) intact, and makes f and f_q fuse into a larger s-cell f' containing both of them. Clearly, f' is the cell associated with (v_q, R'_q) in $\mathcal{A}(\mathcal{K} \setminus \{K_{i_q}\})$, and it is popular there because $f \subset f'$ was popular in $\mathcal{A}(\mathcal{K})$. We say that the borders (v, R), (v_q, R'_q) are *neighbors* in $\mathcal{A}(\mathcal{K})$.

We then charge (v, R) to its s neighboring 1-level s-borders (v_q, R'_q), for $q = 1, \ldots, s$. Note that each 1-level s-border (v, R) is charged at most s times, once along each of the s edges, incident to v, of the s-cell associated with it. We thus obtain the following recurrence.

$$N_0^{(s)}(\mathcal{K}) \leq N_1^{(s)}(\mathcal{K}) + O\left(N_0^{(s-1)}(\mathcal{K}) + n^{2d-3}\right). \quad (5)$$

Applying Tagansky's technique: The simpler variant. We prove that $N_0^{(s)}(n) = O(n^{2d-3} \log^s n)$ by induction on s. For the base case $s = 0$, we have $N_0^{(0)}(n) = O(n^{2d-3})$ by Theorem IV.3. Consider a fixed $s \geq 1$ and assume that the bound holds for $s-1$, so (5) becomes

$$N_0^{(s)}(\mathcal{K}) \le N_1^{(s)}(\mathcal{K}) + O(n^{2d-3}\log^{s-1} n). \qquad (6)$$

Let $\mathcal{R}$ be a random sample of $n-1$ sets of $\mathcal{K}$, obtained by removing a random set K from $\mathcal{K}$. Arguing exactly as in the three-dimensional case, the expected number of 0-level popular s-borders in $\mathcal{A}(\mathcal{R})$ satisfies

$$\mathbf{E}\{N_0^{(s)}(\mathcal{R})\} \ge \frac{n-2d+2}{n} N_0^{(s)}(\mathcal{K}) + \frac{1}{n} N_1^{(s)}(\mathcal{K}). \qquad (7)$$

Combining this inequality with (6), we get

$$\frac{1}{n} N_0^{(s)}(\mathcal{K}) \le \frac{1}{n} N_1^{(s)}(\mathcal{K}) + O(n^{2d-4}\log^{s-1} n) \le$$
$$\mathbf{E}\left\{N_0^{(s)}(\mathcal{R})\right\} - \frac{n-2d+2}{n} N_0^{(s)}(\mathcal{K}) + O(n^{2d-4}\log^{s-1} n),$$

or

$$\frac{n-2d+3}{n} N_0^{(s)}(\mathcal{K}) \le \mathbf{E}\left\{N_0^{(s)}(\mathcal{R})\right\} + O(n^{2d-4}\log^{s-1} n).$$

Replacing $N_0^{(s)}(\mathcal{K})$ and $N_0^{(s)}(\mathcal{R})$ by their respective maximum possible values $N_0^{(s)}(n)$ and $N_0^{(s)}(n-1)$, we get the recurrence

$$\frac{n-2d+3}{n} N_0^{(s)}(n) \le N_0^{(s)}(n-1) + O(n^{2d-4}\log^{s-1} n),$$

whose solution is easily seen to be $N_0^{(s)}(n) = O(n^{2d-3}\log^{s} n)$. This establishes the induction step and thus proves the asserted bound. In particular, we have so far

$$g_d(n) = O(n^{2d-3}\log^{d-1} n).$$

Improved bounds for $s \ge 2$. In the full version [15], we extend the refined analysis in Section III-C to higher dimensions, and show that

$$N_0^{(s)}(\mathcal{K}) = O(n^{2d-3}\log n), \qquad (8)$$

for any $1 \le s \le d-1$, by establishing a sharper variant of (5). We thus obtain the following main result.

Theorem IV.4. *Any collection $\mathcal{K}$ of n pairwise disjoint convex sets in $\mathbb{R}^d$, for any $d \ge 3$, admits at most $O(n^{2d-3}\log n)$ geometric permutations.*

V. Discussion

Although the improvement presented in this paper is significant, especially since no progress was made on the problem during the past 20 years, it is far from satisfactory, since we strongly believe (and tend to conjecture) that the correct upper bounds are close to $O(n^{d-1})$, for any $d \ge 3$. Improving further the bounds is the main open problem left by this study. A modest subgoal is to get rid of the logarithmic factor in our bounds, and show, e.g., that $g_3(n) = O(n^3)$.

References

[1] P. K. Agarwal, On stabbing lines for convex polyhedra in 3D, *Comput. Geom. Theory Appls.* 4 (1994), 177–189.

[2] B. Aronov and S. Smorodinsky, Geometric permutations induced by line transversals through a fixed point, *Discrete Comput. Geom.* 34 (2005), 285–294.

[3] O. Cheong, X. Goaoc and H. Na, Geometric permutations of disjoint unit spheres, *Comput. Geom. Theory Appls.*, 30 (2005), 253–270.

[4] K. Clarkson and P. Shor, Applications of random sampling in computational geometry, II, *Discrete Comput. Geom.* 4 (1989), 387–421.

[5] 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 (1990), 35–42.

[6] J. E. Goodman, R. Pollack and R. Wenger, On the connected components of the space of line transversals to a family of convex sets, *Discrete Comput. Geom.* 13 (1995), 469–476.

[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, Berlin, 1993, pp. 163–198.

[8] Y. Huang, J. Xu and D. Z. Chen, Geometric permutations of high-dimensional spheres, in *Proc. 12th Annu. ACM-SIAM Sympos. Discrete Algo.*, 2001, pp. 342–343.

[9] H. Kaplan, N. Rubin and M. Sharir, Line transversals of convex polyhedra in $\mathbb{R}^3$, *Proc. 19th Annu. ACM-SIAM Symp. on Discrete Algorithms* (2009), 170–179. Also *SIAM. J. Comput., to appear.*

[10] M. Katchalski, T. Lewis and J. Zaks, Geometric permutations for convex sets, *Discrete Math.* 54 (1985), 271–284.

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

[12] M. Katchalski, S. Suri and Y. Zhou, A constant bound for geometric permutations of disjoint unit balls, *Discrete Comput. Geom.* 29(2) (2003), 161–173.

[13] M. J. Katz and K. R. Varadarajan, A tight bound on the number of geometric permutations of convex fat objects in $\mathbb{R}^d$, *Discrete Comput. Geom.* 26(4) (2001), 543–548.

[14] V. Koltun and M. Sharir, The partition technique for overlays of envelopes, *SIAM J. Comput.* 32(4) (2003), 841–863.

[15] N. Rubin, H. Kaplan and M. Sharir, Improved bounds for geometric permutations, `http://www.cs.tau.ac.il/~rubinnat/PermsF.pdf`.

[16] M. Sharir and P. K. Agarwal, *Davenport-Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, New York, 1995.

[17] 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) (2000), 247–259.

[18] B. Tagansky, A new technique for analyzing substructures in arrangements of piecewise linear surfaces, *Discrete Comput. Geom.* 16(4) (1996), 455–479.

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

[20] R. Wenger, Helly-type theorems and geometric transversals, in *Handbook of Discrete and Computational Geometry*, 2nd Edition (J.E. Goodman and J. O'Rourke, Eds.), Chapman & Hall/CRC Press, 2004, pp. 73–96.

[21] Y. Zhou and S. Suri, Geometric permutations of balls with bounded size disparity, *Comput. Geom. Theory Appls.* 26(1) (2003), 3–20.