

LINE TRANSVERSALS OF CONVEX POLYHEDRA IN \mathbb{R}^3 *

HAIM KAPLAN[†], NATAN RUBIN[†], AND MICHA SHARIR[‡]

Abstract. We establish a bound of $O(n^2k^{1+\varepsilon})$, for any $\varepsilon > 0$, on the combinatorial complexity of the set \mathcal{T} of line transversals of a collection \mathcal{P} of k convex polyhedra in \mathbb{R}^3 with a total of n facets, and we present a randomized algorithm which computes the boundary of \mathcal{T} in comparable expected time. Thus, when $k \ll n$, the new bounds on the complexity (and construction cost) of \mathcal{T} improve upon the previously best known bounds, which are nearly cubic in n . To obtain the above result, we study the set \mathcal{T}_{ℓ_0} of line transversals which emanate from a fixed line ℓ_0 , establish an almost tight bound of $O(nk^{1+\varepsilon})$ on the complexity of \mathcal{T}_{ℓ_0} , and provide a randomized algorithm which computes \mathcal{T}_{ℓ_0} in comparable expected time. Slightly improved combinatorial bounds for the complexity of \mathcal{T}_{ℓ_0} and comparable improvements in the cost of constructing this set are established for two special cases, both assuming that the polyhedra of \mathcal{P} are pairwise disjoint: the case where ℓ_0 is disjoint from the polyhedra of \mathcal{P} , and the case where the polyhedra of \mathcal{P} are unbounded in a direction parallel to ℓ_0 . Our result is related to the problem of bounding the number of geometric permutations of a collection \mathcal{C} of k pairwise-disjoint convex sets in \mathbb{R}^3 , namely, the number of distinct orders in which the line transversals of \mathcal{C} visit its members. We obtain a new partial result on this problem.

Key words. lines in space, convex polyhedra, line transversals, combinatorial complexity, extremal stabbing lines

AMS subject classifications. 52B05, 52B10, 52C35, 52C45, 65D18, 68Q25, 68U05

DOI. 10.1137/080744694

1. Introduction.

Line transversals—a brief background. In this paper we study the combinatorial complexity of the set of line transversals of a collection of convex polyhedra in \mathbb{R}^3 . This is a special case of the general study of line transversals to a collection of convex sets in \mathbb{R}^d , for any $d \geq 2$, a topic that has been extensively studied for several decades; see the survey papers [18, 28, 34].

Let \mathcal{P} be a family of k convex sets in \mathbb{R}^3 . A line ℓ is a *transversal* of \mathcal{P} if it intersects every member of \mathcal{P} . The set of all line transversals of \mathcal{P} is called the *transversal space* (or *stabbing region*) of \mathcal{P} and is denoted by $\mathcal{T}(\mathcal{P})$. Lines in \mathbb{R}^3 have four degrees of freedom, and they can be represented as points in an appropriate projective 4-space, called the Grassmannian manifold [11, 32]. For simplicity (and with some care), we will represent them as points in \mathbb{R}^4 . Thus $\mathcal{T}(\mathcal{P})$ is a subset of \mathbb{R}^4 , and we define its *combinatorial complexity* as the total number of topological faces, of all dimensions,

*Received by the editors December 22, 2008; accepted for publication (in revised form) April 29, 2010; published electronically July 29, 2010. A preliminary version of this paper appeared in *Proceedings of the 20th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2009, pp. 170–179. This paper is part of the second author’s Ph.D. dissertation, prepared under the supervision of the first and third authors.

<http://www.siam.org/journals/sicomp/39-7/74469.html>

[†]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel (haimk@post.tau.ac.il, rubinnat@post.tau.ac.il). The work of these authors was supported by grant 975/06 from the Israel Science Foundation and by grant 2006-204 from the United States–Israel Binational Science Foundation.

[‡]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel, and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012 (michas@post.tau.ac.il). This author’s work was partially supported by NSF grants CCF-05-14079 and CCF-08-30272, by grant 2006-194 from the United States–Israeli Binational Science Foundation, by grants 155/05 and 338/09 from the Israel Science Foundation, by a grant from the AFIRST joint French-Israeli program, and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University.

forming the boundary of $\mathcal{T}(\mathcal{P})$. (If the combinatorial complexity of the individual members of \mathcal{P} is not constant, as is the case in this paper, the complexity of the individual faces of $\partial\mathcal{T}(\mathcal{P})$ need not be constant either, and we need to take this into account in the complexity bound of $\mathcal{T}(\mathcal{P})$. A more detailed definition and analysis for the case at hand are given later in the paper.)

A standard reduction leads to a representation of $\mathcal{T}(\mathcal{P})$ as a region (a “sandwich” region) in \mathbb{R}^4 , enclosed between the lower envelope and the upper envelope of two collections of surfaces, describing upper and lower tangencies to each member of \mathcal{P} ; see, for example, [2, 24] for a description of this reduction, a special case of which is described below. When \mathcal{P} is a set of convex polyhedra, the surface of line tangents to any $P \in \mathcal{P}$ is decomposed into patches, each representing tangents to P at a fixed edge of P . The boundary vertices of $\mathcal{T}(\mathcal{P})$ correspond to *extremal stabbing lines*, which are transversals of \mathcal{P} that (i) are tangent to some polyhedra of \mathcal{P} at specific edges or vertices, and (ii) cannot be continuously moved while continuing to touch the same edges and vertices of those polyhedra. As we will later note, the worst-case combinatorial complexity of $\mathcal{T}(\mathcal{P})$ is bounded by the maximum number of vertices of $\mathcal{T}(\mathcal{P})$, so it suffices to bound the latter quantity.

Any upper bound on the maximal *combinatorial complexity* of $\mathcal{T}(\mathcal{P})$ also serves as a natural upper bound on the maximal number of connected components of $\mathcal{T}(\mathcal{P})$ (and, a posteriori, also on the number of geometric permutations of \mathcal{P} , as defined in the abstract, assuming that its elements are pairwise disjoint).

In the planar case, the complexity of $\mathcal{T}(\mathcal{P})$, when the elements of \mathcal{P} are pairwise disjoint, is $O(k)$ (see, e.g., [15]), but can be $\Omega(n)$ otherwise, where n is the total description complexity of the objects of \mathcal{T} , with an almost matching slightly super-linear upper bound (in n); see [30, section 8.2.1] and [14] (e.g., when the objects of \mathcal{P} are convex polygons, and n is the overall number of their edges, the complexity of $\mathcal{T}(\mathcal{P})$ is $O(n\alpha(n))$, where $\alpha(n)$ is the slowly growing inverse Ackermann function). In the three-dimensional case, the complexity of $\mathcal{T}(\mathcal{P})$ depends on n , even if the elements of \mathcal{P} are pairwise disjoint (see, e.g., [9]).

Consider first the case where \mathcal{P} is a set of convex polyhedra, with a total of n facets. The first algorithms for computing $\mathcal{T}(\mathcal{P})$ in this case run in time about $O(n^4)$ [8, 25]. Pellegrini and Shor [27] establish an upper bound of $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$, on the complexity of $\mathcal{T}(\mathcal{P})$ and describe an algorithm for computing the boundary vertices of $\mathcal{T}(\mathcal{P})$, with a comparable running time. Agarwal [1] improves the upper bound for the complexity of $\mathcal{T}(\mathcal{P})$ to $O(n^3 \log n)$.

When the sets in \mathcal{P} are semialgebraic of *constant description complexity*¹ (for example, if the sets in \mathcal{P} are balls, or tetrahedra), the complexity of $\mathcal{T}(\mathcal{P})$ is $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$, as follows from the general and more recent result of Koltun and Sharir [24] on the complexity of sandwich regions of trivariate functions.² In \mathbb{R}^3 there are almost matching lower-bound constructions of n triangles in \mathbb{R}^3 , showing that the complexity of $\mathcal{T}(\mathcal{P})$ can be $\Omega(n^3)$; see [2, 26]. However, for the number of connected components of $\mathcal{T}(\mathcal{P})$, the best known lower bounds are $\Omega(n^2)$, or $\Omega(n^{d-1})$ in d dimensions [31]. Narrowing this gap, even for restricted families of objects, is an intriguing open problem already for $d = 3$ dimensions.

These are the best general known bounds on the complexity of $\mathcal{T}(\mathcal{P})$, but there are some improved bounds in restricted cases: Aronov and Smorodinsky [7] proved

¹That is, each set is a semialgebraic set defined by a Boolean combination of a constant number of polynomial equalities and inequalities of constant maximum degree.

²It is indeed stated in [24] as a corollary.

that when restricting the transversals to pass through a fixed point, the transversal space has a maximum of $\Theta(k^{d-1})$ components for any collection \mathcal{P} of k (not necessarily pairwise disjoint) convex sets in \mathbb{R}^d . Brönnimann et al. [10] gave a complete description of the transversal space of k segments in \mathbb{R}^3 . In this case the transversal space consists of a maximum of k connected components.

If the objects in \mathcal{P} are pairwise disjoint, a line transversal meets them in a well-defined order, called a *geometric permutation*. Another line of research, initiated by Katchalski, Lewis, and Liu [21], studies the maximum possible number $g_d(k)$ of geometric permutations of a set \mathcal{P} of k pairwise disjoint convex objects in \mathbb{R}^d . The known bounds on $g_d(k)$ in the general case are as follows: $g_2(k) = 2k - 2$ (see [15]), $g_d(k) = O(k^{2d-2})$, $d \geq 3$ (see [33]), and $g_d(k) = \Omega(k^{d-1})$, $d \geq 3$ (see [22, 31]). A tight bound of $\Theta(k^{d-1})$, for all $d \geq 1$, is achieved in several special cases such as pairwise disjoint balls [31] and families of fat objects [23].

Consider now the special case studied in this paper where \mathcal{P} is a collection of k convex polyhedra in \mathbb{R}^3 with a total of n facets. One of the most relevant predecessors of this paper is a recent paper of Brönnimann et al. [9], who prove that the entire arrangement \mathcal{A}^* of surfaces (in parametric 4-space) describing upper and lower tangencies to the individual polyhedra of \mathcal{P} has complexity $O(n^2k^2)$, and this bound is tight in the worst case. They also describe a deterministic algorithm which computes the combinatorial description of \mathcal{A}^* in comparable time. Each cell in \mathcal{A}^* corresponds to a maximal connected set of lines which stab the same subset of \mathcal{P} . In particular, $\mathcal{T}(\mathcal{P})$ is equal to the union of all the cells in $\mathcal{A}^*(\mathcal{P})$ whose stabbed subset is the entire \mathcal{P} .

Efrat et al. [16] assume that the polyhedra in \mathcal{P} are pairwise disjoint and consider the restricted space \mathcal{L} of lines that pass through a fixed line ℓ_0 . They prove that the set of all lines in \mathcal{L} which stab at least one polyhedron in \mathcal{P} (alternatively, the set of all lines in \mathcal{L} which miss all the polyhedra of \mathcal{P}) has combinatorial complexity $O(nk^2)$, and this bound is tight in the worst case. When the polyhedra in \mathcal{P} are unbounded in the ℓ_0 -negative direction, the bound improves to $O(nk2^{\alpha(k)})$; see [16] for more details.

Other studies consider *ray shooting* amid such a collection of polyhedra and obtain performance bounds that depend on both k and n ; see [4, 20]. The latter study [20], by the present authors, considers the scenario where the rays along which one shoots are constrained to lie on lines that pass through a fixed line ℓ_0 and describes a data structure which answers ray-shooting queries in polylogarithmic time and requires storage which is near-linear in n .

Our contribution. We consider an arbitrary collection \mathcal{P} of k convex polyhedra in \mathbb{R}^3 , with a total of n facets, and derive an upper bound of $O(n^2k^{1+\varepsilon})$, for any $\varepsilon > 0$, on the complexity of $\mathcal{T}(\mathcal{P})$. We provide a randomized (Las Vegas) algorithm for computing a description of the boundary of $\mathcal{T}(\mathcal{P})$, with comparable expected running time. We also present a lower bound construction of such a collection \mathcal{P} for which $\mathcal{T}(\mathcal{P})$ has complexity $\Omega(n^2 + nk^2)$.

To achieve the general upper bound on the complexity of $\mathcal{T}(\mathcal{P})$, we focus on the restricted case where we consider only lines which pass through a fixed line ℓ_0 , and we study the resulting stabbing region $\mathcal{T}_{\ell_0}(\mathcal{P}) := \mathcal{T}(\mathcal{P}) \cap \mathcal{L}$, where $\mathcal{L} = \mathcal{L}_{\ell_0}$ is the space of these restricted lines. Unlike the general case of lines in 3-space, which have four degrees of freedom, lines in \mathcal{L} have only three degrees of freedom, and we represent them as points in an appropriately parametrized three-dimensional space. The overall bound is obtained by repeating this analysis for all the $O(n)$ lines ℓ_0 which contain polyhedra edges.

The *combinatorial complexity* of $\mathcal{T}_{\ell_0}(\mathcal{P})$ is dominated by the number of *vertices* of $\mathcal{T}_{\ell_0}(\mathcal{P})$, which are formed by intersections of triples of tangency surfaces and hence represent extremal stabbing lines which pass through ℓ_0 and are tangent to three distinct polyhedra of \mathcal{P} (as we will argue later, the number of all other features of $\mathcal{T}_{\ell_0}(\mathcal{P})$ is proportional to nk plus the number of these vertices). We show that the combinatorial complexity of $\mathcal{T}_{\ell_0}(\mathcal{P})$ is $O(nk^{1+\varepsilon})$, for any $\varepsilon > 0$, and that the boundary representation of $\mathcal{T}_{\ell_0}(\mathcal{P})$ can be computed in comparable randomized expected time. To appreciate this bound, we note that the standard representation of $\mathcal{T}(\mathcal{P})$ as a sandwich region between a lower envelope and an upper envelope also holds in the case of $\mathcal{T}_{\ell_0}(\mathcal{P})$, except that here the envelopes are of *bivariate* functions. There are only k functions in each collection, where each function represents an upper tangency or a lower tangency to some fixed polyhedron in \mathcal{P} , but their graphs do not have constant description complexity—each graph is partitioned into patches, each representing tangency at some fixed edge of the respective polyhedron. We can thus regard the sandwich region as being formed by a total of $O(n)$ partially defined bivariate functions, each now of constant description complexity, so by the results of [3, 24], the complexity of the stabbing region is $O(n^{2+\varepsilon})$ for any $\varepsilon > 0$. Our contribution is thus in making this bound depend also on k , so that it becomes only *linear* in n ; this is a significant improvement when $k \ll n$.

We also consider a pair of restricted instances of the problem, both of which assume that the polyhedra of \mathcal{P} are pairwise disjoint. In the first case, when the polyhedra of \mathcal{P} are disjoint from ℓ_0 , our general analysis easily implies that the complexity of the stabbing region is only $O((nk + k^3)\beta_4(k))$, where³ $\beta_4(k) = 2^{\alpha(k)}$ and $\alpha(k)$ is the extremely slowly growing inverse Ackermann function. In the second case, when all the polyhedra in \mathcal{P} are unbounded in a direction parallel to ℓ_0 , we improve the upper bound on the stabbing region to $O(nk\beta_4(k))$. In this case, the sandwich region degenerates to the region above the upper envelope of the lower tangency functions or, symmetrically, to the region below the lower envelope of the upper tangency functions. The improved bound then follows by showing that the complexity of a single envelope, rather than of a sandwich region between two envelopes, is only $O(nk\beta_4(k))$. (This bound holds regardless of whether the polyhedra are unbounded or not, but it requires them to be pairwise disjoint.) In both special instances, the stabbing region within \mathcal{L} can be computed in deterministic time, asymptotically close to its worst-case complexity.

We also show that, for any collection \mathcal{P} of k pairwise disjoint convex objects in \mathbb{R}^3 , one of which is a line, the number of geometric permutations of \mathcal{P} is $O(k^3)$. A naive bound on this number is $O(k^4)$ (which is the general bound mentioned above [33]). In the special case where \mathcal{P} is a pairwise disjoint collection of convex polyhedra (all disjoint from ℓ_0), the number of geometric permutations of \mathcal{P} induced by lines in \mathcal{L} (i.e., the number of geometric permutations of $\mathcal{P} \cup \{\ell_0\}$) is $O(\min\{k^3, nk^{1+\varepsilon}\})$, for any $\varepsilon > 0$, where the second term follows from the fact that the number of geometric permutations is always upper bounded by the complexity of $\mathcal{T}(\mathcal{P})$. We also apply our general $O(nk^{1+\varepsilon})$ bound on the complexity of $\mathcal{T}_{\ell_0}(\mathcal{P})$ to obtain an upper bound of $O(k^{3+\varepsilon})$, for any $\varepsilon > 0$, on the number of geometric permutations of a collection \mathcal{P} of k pairwise disjoint convex objects, induced by lines passing through some fixed line ℓ_0 . The advantage of this seemingly weaker bound is that it does not require ℓ_0 to be disjoint from all of the objects of \mathcal{P} . Still, the only known lower bound on this

³The reason for the index 4 is that $\beta_4(k) = \Theta(\lambda_4(k)/k)$, where $\lambda_4(k)$ is the maximum length of Davenport–Schinzel sequences of order 4 on k symbols; see [30] and below.




FIG. 1. Representing an oriented line $\ell \in \mathcal{L}$. The plane $\Pi_{\theta(\ell)}$ contains ℓ , $(\theta(\ell), \varphi(\ell))$ are the spherical coordinates of the orientation vector of ℓ , and $z(\ell) = \ell \cap \ell_0$ is the z -intercept of ℓ .

quantity is $\Omega(k^2)$, and it would be very interesting to show that this is also an upper bound in the special case of collections of pairwise disjoint convex objects in \mathbb{R}^3 , one of which is a line. See section 7 for details.

2. Preliminaries. Let $\mathcal{P} = \{P_1, \dots, P_k\}$ be a collection of k convex polyhedra in \mathbb{R}^3 with a total of n facets, and let ℓ_0 be a fixed line. Without loss of generality, we take ℓ_0 , for the time being, to be the z -axis.⁴

Let $\mathcal{L} = \mathcal{L}_{\ell_0}$ denote the space of all lines that pass through ℓ_0 (other than ℓ_0 itself). Lines in \mathcal{L} have three degrees of freedom, and we represent each (directed) line $\ell \in \mathcal{L}$ by a triple $(\theta(\ell), \varphi(\ell), z(\ell))$, where $z(\ell)$ is the z -coordinate of the intercept of ℓ at ℓ_0 and $(\theta(\ell), \varphi(\ell))$ are the spherical coordinates of the orientation of ℓ . Clearly, all lines ℓ with $\theta(\ell) = \theta$ lie in the plane through ℓ_0 at xy -orientation θ ; we denote this plane by Π_θ . See Figure 1. The intersection of Π_θ with a polyhedron P is the polygon $P(\theta) = P \cap \Pi_\theta$ (if it is not empty).

We define, for each polyhedron $P \in \mathcal{P}$, a pair of (partial) bivariate functions σ_P^+ and σ_P^- over the $\theta\varphi$ -domain, so that $\sigma_P^-(\theta, \varphi)$ (resp., $\sigma_P^+(\theta, \varphi)$) is the z -intercept (at ℓ_0) of the line whose orientation has spherical coordinates (θ, φ) and which is tangent to P from below (resp., above).

For each $P \in \mathcal{P}$, the graphs of σ_P^+ and σ_P^- are (θ, φ) -monotone surfaces, representing tangents to the upper and lower portions of ∂P , respectively. With an appropriate reparametrization (e.g., replacing θ with $\tan \frac{\theta}{2}$ and φ with $\cot \frac{\varphi}{2}$), each surface σ_P^- (resp., σ_P^+) consists of monotone semialgebraic surface patches, each of which is a graph of a partially defined function, representing lower (resp., upper) tangents to P at a fixed edge of its lower (resp., upper) boundary. For simplicity of presentation, we continue to use the coordinates θ, φ .

The set $\mathcal{T}_{\ell_0}(\mathcal{P})$ of all transversals to \mathcal{P} (in \mathcal{L}) is then the *sandwich region*

$$(2.1) \quad \left\{ (\theta, \varphi, z) \in \mathcal{L} \mid \max_{P \in \mathcal{P}} \sigma_P^-(\theta, \varphi) \leq z \leq \min_{P \in \mathcal{P}} \sigma_P^+(\theta, \varphi) \right\}$$

between the upper envelope $E_U = \max_{P \in \mathcal{P}} \sigma_P^-$ of the functions σ_P^- and the lower envelope $E_L = \min_{P \in \mathcal{P}} \sigma_P^+$ of the functions σ_P^+ .

Following [9, 27] (see also the introduction), we define an *extremal line* to be a line $\ell \in \mathcal{L}$ tangent to some polyhedra of \mathcal{P} at a set A of respective vertices and

⁴Later, in section 3, we will use a separate coordinate system attached to each edge of every polyhedron in \mathcal{P} .

edges, so that ℓ cannot be continuously moved (within \mathcal{L}) while remaining transversal to the elements of A . An *extremal stabbing line* is an extremal line, which is also a transversal of \mathcal{P} . Every vertex of $\mathcal{T}_{\ell_0}(\mathcal{P})$ corresponds to an extremal stabbing line; see, e.g., [27]. Assuming ℓ_0 is not coplanar with any facet or edge of any polyhedron of \mathcal{P} , the converse statement also holds. Namely, each extremal stabbing line corresponds to a vertex of $\mathcal{T}_{\ell_0}(\mathcal{P})$; see a discussion of this issue in the end of section 3.

The following theorem was proven by Brönnimann et al. [9]; see also [20].

THEOREM 2.1. *Let P and Q be two convex polyhedra in \mathbb{R}^3 , having n_P and n_Q facets, respectively. Then the arrangement of the tangency surfaces σ_P^+ , σ_P^- , σ_Q^+ , and σ_Q^- has combinatorial complexity $O(n_P + n_Q)$, and it can be computed in $O((n_P + n_Q) \log(n_P + n_Q))$ time. In particular, the combinatorial complexity of $\mathcal{T}_{\ell_0}(\{P, Q\})$ is $O(n_P + n_Q)$, and it can be computed in (deterministic) $O((n_P + n_Q) \log(n_P + n_Q))$ time.*

Henceforth, we assume that the polyhedra of \mathcal{P} and ℓ_0 are in *general position*. In particular, at most three edges, or a vertex and an edge of the given polyhedra, can admit a common transversal through ℓ_0 , no edge of any polyhedron is coplanar with ℓ_0 or with an edge of another polyhedron, and no facet of any polyhedron contains a vertex of another polyhedron. In addition, we assume that lines containing polyhedra edges are in general position, in the sense that at most three of these lines can lie in a common *regulus* (see [32] and below). In particular, any four of these lines admit at most two common transversal lines [32]. This assumption involves no loss of generality, because, as we will show toward the end of section 3, the bounds obtained under this assumption continue to hold also for collections of polyhedra which are not in general position.

We say that a vertex or an edge ξ of $\partial\mathcal{T}_{\ell_0}(\mathcal{P})$ is *defined* by a set of polyhedra $\mathcal{P}' \subset \mathcal{P}$ if \mathcal{P}' is a minimal set of polyhedra such that ξ is present in $\partial\mathcal{T}_{\ell_0}(\mathcal{P}')$. Assuming general position, each vertex v of $\partial\mathcal{T}_{\ell_0}(\mathcal{P})$ is defined by a unique set of between one and three polyhedra. By Theorem 2.1, any two polyhedra $P, Q \in \mathcal{P}$, having, respectively, n_P and n_Q facets, define $O(n_P + n_Q)$ features of $\partial\mathcal{T}_{\ell_0}(\mathcal{P})$, each of which is the locus of lines tangent to P and Q at a fixed number (at most three) specific boundary edges and vertices. Summing over all pairs of polyhedra in \mathcal{P} , we obtain a bound of $O(nk)$ on the number of features of this kind defined by at most two polyhedra. Any other feature is incident to at least one vertex of $\mathcal{T}_{\ell_0}(\mathcal{P})$ and can be charged to it, so that each vertex is charged by only $O(1)$ incident features of $\mathcal{T}_{\ell_0}(\mathcal{P})$.

Therefore, it is sufficient to bound the number of vertices of $\partial\mathcal{T}_{\ell_0}(\mathcal{P})$ which are defined by three polyhedra; each such vertex is an extremal stabbing line (in \mathcal{L}) which is a common tangent to three polyhedra of \mathcal{P} , at (the relative interiors of) three edges, one of each polyhedron. The additional complexity $O(nk)$ will be subsumed in the bound that we will obtain for the latter kind of vertices.

Let P and Q be a pair of polyhedra of \mathcal{P} , and let ζ be a boundary edge of $\mathcal{T}_{\ell_0}(\{P, Q\})$, contained in the common intersection of two semialgebraic patches $\sigma_e, \sigma_{e'}$, where σ_e (resp., $\sigma_{e'}$) is contained in $\sigma_P^+ \cup \sigma_P^-$ (resp., $\sigma_Q^+ \cup \sigma_Q^-$) and represents oriented lines (in \mathcal{L}) tangent to P (resp., Q) at e (resp., e'). That is, ζ represents (i.e., is the trace of) a maximal connected set of lines that are tangent to P at e and to Q at e' . Note that ζ is a connected portion of the locus of lines that pass through three fixed lines, namely, ℓ_0 , and the two lines supporting e and e' , respectively. In general position, this locus is commonly referred to as a *regulus*, whose lines have only one degree of freedom and trace (a portion of) a ruled surface in \mathbb{R}^3 , which is either a hyperbolic paraboloid or a 1-sheeted hyperboloid; see [11, 32] for more details on reguli.

In the restricted context of this paper, a *regulus* denotes a maximal connected set of oriented lines in \mathcal{L} that are tangent to two fixed polyhedra of \mathcal{P} at two fixed edges, one of each polyhedron. In particular, each regulus represents some boundary edge of $\mathcal{T}_{\ell_0}(\{P, Q\})$ for some pair of distinct polyhedra $P, Q \in \mathcal{P}$. Theorem 2.1, together with the follow-up discussion, imply that the polyhedra of \mathcal{P} define a total of $O(nk)$ such reguli.

Transversals parallel to a fixed plane. Let h be a fixed plane in \mathbb{R}^3 . Denote by \mathcal{L}_h the space of lines passing through ℓ_0 and parallel to h . Clearly, lines in \mathcal{L}_h have only two degrees of freedom, and if h is generic (i.e., it is not parallel to any extremal line in \mathcal{L} that is tangent to three polyhedra in \mathcal{P}), any extremal stabbing line within \mathcal{L}_h is tangent to at most two polyhedra at a corresponding pair of edges. We establish the following lemma, which we need in our analysis.

LEMMA 2.2. *Let h be a fixed plane in \mathbb{R}^3 . Then the number of extremal stabbing lines of \mathcal{P} within the space \mathcal{L}_h , as defined above, is $O(n\beta_4(k)) = O(n \cdot 2^{\alpha(k)})$.*

Proof. Assume first that h is not parallel to ℓ_0 . Then lines in \mathcal{L}_h can be parametrized by their two coordinates (θ, z) , where the third coordinate $\varphi = \varphi_h(\theta)$ depends only on θ . Hence, using (2.1), the stabbing region of \mathcal{P} within \mathcal{L}_h is the set

$$\left\{ (\theta, z) \mid \max_{P \in \mathcal{P}} \sigma_P^-(\theta, \varphi_h(\theta)) \leq z \leq \min_{P \in \mathcal{P}} \sigma_P^+(\theta, \varphi_h(\theta)) \right\}.$$

That is, the transversal space within \mathcal{L}_h is the sandwich region between the lower envelope of the k univariate functions $\sigma_P^+(\theta, \varphi_h(\theta))$ and the upper envelope of the k univariate functions $\sigma_P^-(\theta, \varphi_h(\theta))$ for $P \in \mathcal{P}$. The graphs of these functions are (connected and) piecewise smooth, where each piece represents lines tangent to some specific polyhedron at some specific edge of its boundary. The overall number of these pieces (subarcs) is thus $O(n)$, and any pair of them intersect at most twice. Indeed, any such intersection represents a line ℓ in \mathcal{L} , parallel to h , which intersects the two polyhedra edges e_1 and e_2 corresponding to those subarcs. Hence, ℓ is a line that passes through ℓ_0 , the two lines containing e_1 and e_2 , and the line at infinity in h . Assuming general position, there are at most two such lines ℓ [32] (see also [10]).

Consider next the case where h is parallel to ℓ_0 . We can assume that h contains ℓ_0 , and thus equals to a plane Π_{θ_h} , for some fixed $0 \leq \theta_h \leq 2\pi$. That is, the transversal space within \mathcal{L}_h is the sandwich region between the lower envelope of the k univariate functions (of φ) $\sigma_P^+(\theta_h, \varphi)$ and the upper envelope of the k functions $\sigma_P^-(\theta_h, \varphi)$ for $P \in \mathcal{P}$. Again, the graphs of those functions are connected and piecewise smooth, where each piece represents lines tangent to some specific polygon $P(\theta_h)$ at some specific vertex of its boundary; in this case any pair of these pieces intersect at most once—the intersection corresponds to the unique line in Π_{θ_h} that passes through the two respective vertices. See [20] for a similar analysis. The overall number of these pieces is, as above, $O(n)$.

In both cases, the lemma follows from the upper bound on the complexity of the sandwich region defined by $O(k)$ piecewise-smooth connected curves in \mathbb{R}^2 , which are composed of a total of $O(n)$ algebraic arcs, each pair of which intersect at most twice. Recalling our definition of $\beta_4(k)$, this bound is $O(\frac{n}{k} \lambda_4(k)) = O(n\beta_4(k))$, where $\lambda_4(k) = \Theta(k \cdot 2^{\alpha(k)})$ is the maximum length of Davenport–Schinzel sequences of order 4 on k symbols; see [30, Theorem 1.4]. \square

Separating convex bodies in \mathbb{R}^3 . We denote by \mathbb{S}^2 the unit sphere in 3-space centered at the origin. For each oriented line ℓ in 3-space, we denote its orientation by $\vec{d}(\ell)$ and represent it as a point on \mathbb{S}^2 , with spherical coordinates $(\theta(\ell), \varphi(\ell))$. For

a plane $h \subset \mathbb{R}^3$, we denote by $C_h \subset \mathbb{S}^2$ the great circle obtained by intersecting \mathbb{S}^2 with the plane parallel to h through the origin; equivalently, C_h is the locus of all orientations on \mathbb{S}^2 that are parallel to h . The following (easy) lemma has been proven by Wenger [33].

LEMMA 2.3. *Let P and Q be a pair of disjoint convex bodies in \mathbb{R}^3 , and let h be a plane which separates them. Then C_h partitions \mathbb{S}^2 into a pair of hemispheres \mathbb{S}_h^+ and \mathbb{S}_h^- such that, for any (oriented) common transversal ℓ of P and Q , ℓ stabs P before (resp., after) Q if and only if $\vec{d}(\ell) \in \mathbb{S}_h^+$ (resp., $\vec{d}(\ell) \in \mathbb{S}_h^-$).*

3. The complexity of $\mathcal{T}_{\ell_0}(\mathcal{P})$. We first establish Theorem 3.1 that bounds the complexity of $\mathcal{T}_{\ell_0}(\mathcal{P})$. In order to get the bound on the complexity of $\mathcal{T}_{\ell_0}(\mathcal{P})$ to depend on k , we reduce the global problem, involving a sandwich region in the three-dimensional space \mathcal{L} , to a collection of two-dimensional problems. The difficulty is that a naive reduction of this sort yields subproblems in which the relevant portion of $\mathcal{T}_{\ell_0}(\mathcal{P})$ is *not* a sandwich region. Our solution uses a more involved approach to ensure that the resulting subproblems do have a sandwich structure, but this requires a more careful and somewhat intricate analysis.

THEOREM 3.1. *Let \mathcal{P} be a set of k convex polyhedra in \mathbb{R}^3 with a total of n facets, and let ℓ_0 be a fixed line. Then the space $\mathcal{T}_{\ell_0}(\mathcal{P})$ of line transversals to \mathcal{P} in $\mathcal{L} = \mathcal{L}_{\ell_0}$ has combinatorial complexity $O(nk^{1+\varepsilon})$ for any $\varepsilon > 0$.*

The bound is almost tight; below we show how to construct collections of k convex polyhedra in general position, with a total of n facets, for arbitrarily large values of k and n , and a line ℓ_0 in general position with respect to \mathcal{P} for which the number of extremal stabbing lines, as above, is $\Omega(nk)$.

In addition, the proof given below assumes general position of ℓ_0 and \mathcal{P} , as defined in section 2. Later, toward the end of this section, we will explain how to extend this theorem to degenerate configurations.

Proof. For each polyhedron $P \in \mathcal{P}$ and each connected component C of $\ell_0 \setminus P$, we choose a single plane that separates C from P (there are at most two such components, and therefore at most two corresponding planes are chosen; if $\ell_0 \cap P = \emptyset$, we choose a single separating plane which is parallel to ℓ_0). See Figure 2 for an illustration. Let H be the resulting set of at most $2k$ separating planes. These planes intersect ℓ_0 in at most $2k$ points, partitioning it into a collection \mathcal{I} of up to $2k + 1$ open “atomic” intervals, so that, for each interval $I \in \mathcal{I}$ and for each polyhedron $P \in \mathcal{P}$, either I is fully contained in P or I is disjoint from P .

Let C_H denote the collection of the great circles $C_h \subset \mathbb{S}^2$, for $h \in H$, as defined in the previous section, and let $\mathcal{A}(C_H)$ denote the arrangement that they form on the sphere \mathbb{S}^2 . The construction and Lemma 2.3 imply that for each cell D of $\mathcal{A}(C_H)$




FIG. 2. We separate the components C_1 and C_2 of $\ell_0 \setminus P$ by the planes h_1 and h_2 , respectively, each containing a facet of P .




FIG. 3. Representing lines in $\mathcal{L}[e_0]$: A side view (left), and the cross-section within the plane $\tilde{\Pi}_\theta$ (right).

and for each $I \in \mathcal{I}$, there exists a partition of \mathcal{P} into three subsets $\mathcal{P}_0 = \mathcal{P}_0(D, I)$, $\mathcal{P}^- = \mathcal{P}^-(D, I)$, and $\mathcal{P}^+ = \mathcal{P}^+(D, I)$ with the following property. For any oriented stabbing line $\ell \in \mathcal{L}$ which emanates from a point on I and has orientation in D , the point $q = \ell \cap \ell_0$ lies in P if $P \in \mathcal{P}_0$, before $\ell \cap P$ along (the oriented) ℓ if $P \in \mathcal{P}^-$, and after $\ell \cap P$ along (the oriented) ℓ if $P \in \mathcal{P}^+$.

As mentioned above, it is sufficient to bound the number of extremal stabbing lines to \mathcal{P} in \mathcal{L} that are tangent to some three polyhedra in \mathcal{P} at three respective edges. We now fix an edge e_0 on the boundary of some polyhedron $P_0 \in \mathcal{P}$, and we denote by $\mathcal{L}[e_0]$ the set of oriented lines which pass through ℓ_0 and are tangent to P_0 at (the relative interior of) e_0 . (Assuming general position of \mathcal{P} and ℓ_0 , we can exclude extremal stabbing lines that pass through one of the endpoints of e_0 or overlap a facet of P_0 . As argued in section 2 (see Theorem 2.1) the number of such “degenerate” extremal stabbers is only $O(nk)$.) We parametrize these lines by the appropriately rotated spherical coordinate system, in which the line ℓ_{e_0} supporting e_0 is the z -axis (with a fixed assigned orientation), and the xz -plane supports one of the facets of P_0 incident to e_0 . It is easy to check that, in general position, this is a unique parametrization. Thus, in the new system, which, for convenience, we continue to denote by (θ, φ) , the angle θ encodes the orientation of a plane $\tilde{\Pi}_\theta$ containing e_0 and rotating about it. However, we are only interested in values of θ for which $\tilde{\Pi}_\theta$ is tangent to P_0 at e_0 . Thus, if θ_0 denotes π minus the dihedral angle of P_0 at e_0 , then we can restrict θ to lie in the union of the two antipodal angular ranges $(0, \theta_0)$ and $(\pi, \theta_0 + \pi)$. For simplicity, and with no loss of generality, we consider only lines in $\mathcal{L}[e_0]$ whose orientation (θ, φ) satisfies $0 < \theta < \theta_0$. See Figure 3 for an illustration.

For a fixed $0 < \theta < \theta_0$, we intersect each $P \in \mathcal{P}$ with the plane $\tilde{\Pi}_\theta$, thus obtaining a set $\tilde{\mathcal{P}}(\theta)$ of (possibly empty) convex polygons of the form $\tilde{P}(\theta) := P \cap \tilde{\Pi}_\theta$. In particular, we have $\tilde{P}_0(\theta) \equiv e_0$ for all $\theta \in (0, \theta_0)$.

Let $\ell_0(\theta)$ denote the point $\ell_0 \cap \tilde{\Pi}_\theta$. By the general position assumption, e_0 is not coplanar with ℓ_0 , so the point $\ell_0(\theta)$ is well defined, except for the unique orientation $\theta_{e_0}^*$, if it exists, at which $\tilde{\Pi}_\theta$ is parallel to ℓ_0 .⁵ Clearly, an oriented line ℓ in $\mathcal{L}[e_0]$ with $\theta(\ell) = \theta$ must pass through $\ell_0(\theta)$. If $\theta_{e_0}^*$ exists, then we split $(0, \theta_0)$ into the two subintervals $(0, \theta_{e_0}^*)$ and $(\theta_{e_0}^*, \theta_0)$, and we apply the ongoing analysis to each of them separately. (By the definition of $\theta_{e_0}^*$, we have $\theta(\ell) \neq \theta_{e_0}^*$ for all $\ell \in \mathcal{L}[e_0]$.) For simplicity of presentation and with no loss of generality, we assume that $\theta_{e_0}^*$ does not

⁵By saying that $\theta_{e_0}^*$ exists, we mean that it lies in $(0, \theta_0)$. Note also that, by the general position, this orientation cannot be 0 or θ_0 .

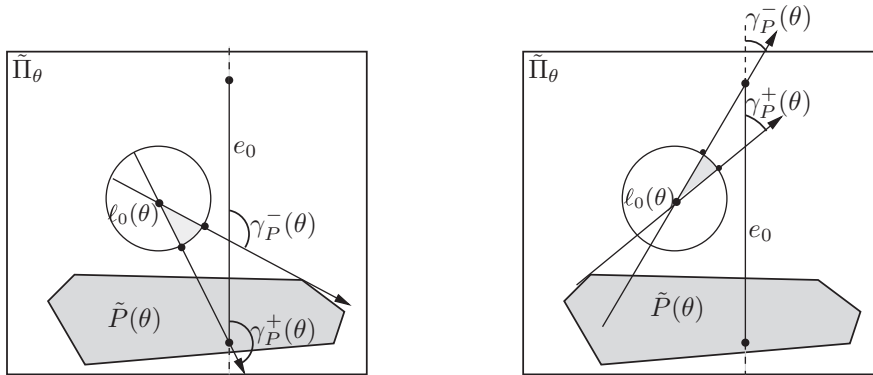


FIG. 4. The interval $(\gamma_P^-(\theta), \gamma_P^+(\theta))$ for $P \in \mathcal{P}^+$ (left) and for $P \in \mathcal{P}^-$ (right).

exist, and we deal with the entire original interval $(0, \theta_0)$. Thus, we can assume that, for all $\theta \in (0, \theta_0)$, the point $\ell_0(\theta)$ remains within $\tilde{\Pi}_\theta$ on the same side of ℓ_{e_0} .

Now with all these preparations, the portion under consideration of $\mathcal{L}[e_0]$, which, for simplicity of presentation, we continue to denote by $\mathcal{L}[e_0]$, is represented as the region⁶

$$\{(\theta, \varphi) \mid 0 < \theta < \theta_0, \varphi^-(\theta) < \varphi < \varphi^+(\theta)\},$$

where $\varphi^-(\theta), \varphi^+(\theta)$ are the φ -coordinates of the respective lines passing through $\ell_0(\theta)$ and through each of the endpoints of e_0 . See Figure 3 (right) for an illustration.

Next we fix a cell D of $\mathcal{A}(C_H)$ (defined at the beginning of the proof) and consider the subset $\mathcal{L}[e_0, D]$ of those lines $\ell \in \mathcal{L}[e_0]$ with orientation in D . (If $\mathcal{L}[e_0, D]$ is empty, then we ignore D .)

For each polyhedron $P \in \mathcal{P}$, we define two functions $\gamma_P^-(\theta), \gamma_P^+(\theta)$, for $\theta \in (0, \theta_0)$, as follows. Let $I \in \mathcal{I}$ be the segment which contains $\ell_0(\theta)$. Then we define $\gamma_P^-(\theta)$ (resp., $\gamma_P^+(\theta)$) to be the minimum (resp., maximum) value of $\varphi \in (\varphi^-(\theta), \varphi^+(\theta))$ such that (a) the line ℓ with representation (θ, φ) intersects P , and (b) the order of $\ell_0 \cap \ell = \ell_0(\theta)$ and $P \cap \ell$ along ℓ is as prescribed by D and I ; see Figure 4. It is easily seen that, for any fixed θ , the range of values of φ for which the line ℓ with local coordinates (θ, φ) satisfies (a) and (b) is an angular interval, which we denote by $T_P^+(\theta)$ (resp., $T_P^-(\theta)$) for $P \in \mathcal{P}^+(D, I)$ (resp., $P \in \mathcal{P}^-(D, I)$). The interval coincides with the full range $(\varphi^-(\theta), \varphi^+(\theta))$ when $P \in \mathcal{P}_0(D, I)$ and is empty when $\tilde{P}(\theta) = \emptyset$. In the latter case, $\gamma_P^-(\theta)$ and $\gamma_P^+(\theta)$ are undefined, but we use the convention of artificially defining them as $\gamma_P^-(\theta) = \varphi^+(\theta)$ and $\gamma_P^+(\theta) = \varphi^-(\theta)$. The definition is tailored to the task of expressing the stabbing region in $\mathcal{L}[e_0, D]$ as a *sandwich region*; see Lemma 3.2.

It follows that, for $\theta \in (0, \theta_0)$, the line ℓ having representation $(\theta, \gamma_P^-(\theta))$ or $(\theta, \gamma_P^+(\theta))$ is either (a) a tangent to $\tilde{P}(\theta)$ passing through $\ell_0(\theta)$ in the plane $\tilde{\Pi}_\theta$, or (b) a line which connects $\ell_0(\theta)$ with one of the endpoints of $e_0 = \tilde{P}_0(\theta)$. The graph of each of the functions γ_P^-, γ_P^+ is connected and piecewise smooth, composed of maximal subarcs, so that each arc represents either lines in $\mathcal{L}[e_0]$ that are tangent to P at some fixed edge, or lines tangent to P_0 at one of the endpoints of e_0 . Since, by construction, all lines in $\mathcal{L}[e_0]$ have $\varphi \in (\varphi^-(\theta), \varphi^+(\theta))$, the arcs of the latter type are redundant,

⁶The inequalities are sharp since lines in $\mathcal{L}[e_0]$ can neither overlap a facet of P_0 incident to e_0 nor pass through any of the endpoints of e_0 ; see the definition of $\mathcal{L}[e_0]$.

in the sense that they lie on the boundary of $\mathcal{L}[e_0]$, and can thus be safely ignored in the subsequent analysis. (Alternatively, whatever happens along these arcs involves only two polyhedra and is therefore already accounted for in Theorem 2.1.)

Now consider arcs of the first kind. An endpoint of such an arc of, say, γ_P^+ , occurs at some θ for which one of the following conditions holds:

- (i) $\tilde{P}(\theta)$ is a singleton point, namely, a vertex v of P . The line $\ell^+(\theta)$ represented by $(\theta, \gamma_P^+(\theta))$ is tangent to P at v .
- (ii) $\ell^+(\theta)$ (as defined in (i)) changes the edge of P it is tangent to. This can happen either when $\ell^+(\theta)$ passes through a vertex of P (while continuing to be tangent to P) or when it overlaps with a facet of P .
- (iii) $\ell_0(\theta)$ coincides with an endpoint of an interval in \mathcal{I} which lies on ∂P . At that moment, $\ell^+(\theta)$ overlaps a facet of P .
- (iv) The line $\ell^+(\theta)$ passes through a vertex of e_0 and, at the same time, is tangent to $\tilde{P}(\theta)$ in $\tilde{\Pi}_\theta$.

As was shown in [9], at each one of the events (i)–(iv), $(\theta, \gamma_P^+(\theta))$ represents an extremal stabbing line to $\{P_0, P\}$, and, therefore, it is a vertex of $\mathcal{T}_{\ell_0}(\{P_0, P\})$.

Note that in this analysis we do not “start” or “stop” an arc at orientations at which it enters or leaves D , so a single arc can straddle several regions D ; an exception is case (iii), where the crossing point into/from D occurs at a boundary great circle induced by P itself. Moreover, the arc does not even have to meet D at all. Each subarc ζ of γ_P^+ , defined by some fixed edge e of P , represents a regulus defined by e_0 and e . Namely, ζ represents a maximal connected set of (oriented) lines in \mathcal{L} that are tangent to P_0 at e_0 and to P at e .

In particular, it follows that any pair of arcs, corresponding to a pair of edges e, e' of distinct respective polyhedra P, P' , intersect at most twice, since these intersections represent lines that pass through four fixed lines: ℓ_0 , and the lines containing the edges e_0, e , and e' ; see [32].

The discussion following Theorem 2.1 implies that the overall number of distinct arcs that comprise the graphs of the functions γ_P^+, γ_P^- , over all possible edges e_0 , polyhedra P , and cells D is $O(nk)$ (where, as noted, a single arc may traverse several cells $D \in \mathcal{D}$, but then we count it only once in this bound). Indeed, for a given pair of distinct polyhedra $P_1, P_2 \in \mathcal{P}$, the regulus of lines tangent to P_1 , at a fixed edge e_1 , and to P_2 , at a fixed edge e_2 , can show up (with different parametrizations) only in the subspaces $\mathcal{L}[e_1], \mathcal{L}[e_2]$ of \mathcal{L} (although it can appear in the subsets $\mathcal{L}[e_1, D]$ of $\mathcal{L}[e_1]$ for several cells D and similarly for e_2). This, combined with Theorem 2.1, implies the claim.

LEMMA 3.2. *Let e_0, D , and P be as above, and let ℓ be an oriented line in $\mathcal{L}[e_0, D]$, with (the $\mathcal{L}[e_0]$ -local) spherical coordinates (θ, φ) . Then ℓ stabs P if and only if*

$$(3.1) \quad \gamma_P^-(\theta) \leq \varphi \leq \gamma_P^+(\theta).$$

Proof. If $\ell_0(\theta) \notin P$, then ℓ stabs P if and only if φ lies in one of the two angular intervals $T_P^-(\theta), T_P^+(\theta)$. (Note that the choice of $\ell \in \mathcal{L}[e_0]$ trivially ensures that φ lies in the interval $(\varphi^-(\theta), \varphi^+(\theta))$.) However, since $\ell \in \mathcal{L}[e_0, D]$, φ can belong only to one of these intervals (that is, to $T_P^-(\theta)$ if P belongs to the corresponding set \mathcal{P}^- or to $T_P^+(\theta)$ if $P \in \mathcal{P}^+$). The claim now follows by definition.

If $\ell_0(\theta) \in P$, then ℓ stabs P , and by definition, (3.1) becomes $\varphi^-(\theta) \leq \varphi \leq \varphi^+(\theta)$, which holds trivially. If $\tilde{P}(\theta) = \emptyset$, the inequality becomes vacuous because of our convention, so the lemma holds in this case too. \square




FIG. 5. The sandwich region in $\mathcal{L}[e_0, D]$, as given by (3.2), overlaid with (the shaded) D and drawn on the unit sphere \mathbb{S}^2 . Note that $\gamma_{P_1}^-$ and $\gamma_{P_2}^+$ appear on the boundary of the sandwich region, but not within D , so their arcs are not counted in $R[e_0, D]$.

It follows from Lemma 3.2 that for a fixed D and e_0 , a line $\ell \in \mathcal{L}[e_0, D]$, with coordinates (θ, φ) , is a transversal to \mathcal{P} if and only if

$$(3.2) \quad \max_{P \in \mathcal{P}} \gamma_P^-(\theta) \leq \varphi \leq \min_{P \in \mathcal{P}} \gamma_P^+(\theta).$$

COROLLARY 3.3. *Let e_0 and D be as above, and let ℓ be an oriented extremal stabbing line to \mathcal{P} in $\mathcal{L}[e_0, D]$ with coordinates (θ, φ) . Then the point (θ, φ) is a vertex on the boundary of the sandwich region given by (3.2).*

Note that we do not guarantee that the great circular arc $\{\theta\} \times (\gamma_P^-(\theta), \gamma_P^+(\theta))$ is contained or even meets D at all. Still, since D is spherically convex (and contained in a hemisphere), the arc intersects D in a connected subarc, so the portion of the sandwich region within D is also θ -monotone (see Figure 5).

Hence, the number of extremal stabbing lines in $\mathcal{L}[e_0, D]$ is upper bounded by the complexity of the sandwich region given by (3.2). Since this region is formed by the graphs of $O(k)$ functions, consisting of some number $\Gamma[e_0, D]$ of connected subarcs, each pair of which intersect at most twice, it follows from [30, Theorem 1.4] (as in the proof of Lemma 2.2) that the complexity of the sandwich region is $O(\Gamma[e_0, D]\beta_4(k))$.

In other words, if we fix D and sum this bound over all edges e_0 of the polyhedra of \mathcal{P} , we get an overall bound of $O(nk\beta_4(k))$. Unfortunately, multiplying this bound by the number $O(k^2)$ of cells D yields an upper bound of $O(nk^3\beta_4(k))$, which is much too large.⁷ (Here we face the difficulty that an arc may appear in several cells D , so we do not have a good bound on $\sum_{e_0, D} \Gamma[e_0, D]$.)

In order to keep the overall bound close to $O(nk)$, we note that we can replace $\Gamma[e_0, D]$ by the number $R[e_0, D]$ of those arcs that actually show up on the boundary of the sandwich region (3.2) at points which represent lines in $\mathcal{L}[e_0, D]$. That is, we exclude arcs which either do not appear on the boundary of the sandwich region at all or appear there but only at (lines represented by) orientations outside D . Our goal is to bound $\sum_{e_0, D} R[e_0, D]$. See Figure 5 for an illustration.

To recap, the preceding analysis implies that the number $N[e_0, D]$ of extremal stabbing lines in $\mathcal{L}[e_0, D]$ satisfies

$$(3.3) \quad N[e_0, D] = O(R[e_0, D]\beta_4(k)).$$

Put

$$N = \sum_{e_0, D} N[e_0, D] \quad \text{and} \quad R = \sum_{e_0, D} R[e_0, D],$$

⁷It is an easy exercise, omitted here, to reduce this bound to $O(nk^2)$, but we are after a stronger bound.

where the sums extend over all choices of edges e_0 of the polyhedra in \mathcal{P} and all cells D of $\mathcal{A}(C_H)$. Denote by $N(n, k)$ (resp., $R(n, k)$) the maximum value of N (resp., R) over all possible choices of a collection \mathcal{P} of k convex polyhedra with a total of n facets and of a fixed line ℓ_0 .

By (3.3), we have

$$(3.4) \quad N(n, k) = O(R(n, k)\beta_4(k)).$$

We next derive a recurrence formula for $N(n, k)$ as follows. For any line $\ell \in \mathcal{L}$, define the *depth* of ℓ to be the number of polyhedra of \mathcal{P} which are *not stabbed* by ℓ . For any integer $t \geq 0$, denote by $N_{\leq t}(n, k)$ the number of extremal lines in \mathcal{L} whose depth is at most t . Since each extremal line in \mathcal{L} is defined by at most three polyhedra, the standard probabilistic argument of Clarkson and Shor [12] implies that

$$N_{\leq t}(n, k) = \mathbf{E} \left\{ O \left(t^3 N \left(n_{\mathcal{R}}, \frac{k}{t} \right) \right) \right\},$$

where the expectation is with respect to a random sample \mathcal{R} of $\frac{k}{t}$ polyhedra of \mathcal{P} and where $n_{\mathcal{R}}$ is the random variable equal to the number of facets in the sampled polyhedra. Clearly, the expected value of $n_{\mathcal{R}}$ is $\frac{n}{t}$. To simplify the presentation, we will rewrite the Clarkson–Shor bound as

$$(3.5) \quad N_{\leq t}(n, k) = O \left(t^3 N \left(\frac{n}{t}, \frac{k}{t} \right) \right).$$

The justification of this bound is similar to those used in earlier works, such as in [17], and will become clearer later in the analysis.

Let r be a regulus of lines passing through ℓ_0 that are tangent to P_0 at e_0 , and to some other polyhedron P , at a fixed edge e of its boundary. According to section 2, r corresponds to an edge of $\mathcal{T}_{\ell_0}(\{P_0, P\})$, and there is a total of $O(nk)$ such edges over all choices of e_0 and P . Since lines in r have one degree of freedom, r is represented by a connected subarc of a constant-degree curve, within the (θ, φ) -parametric space $\mathcal{L}[e_0]$ (or, rather, within an algebraically reparametrized version of it). To recursively bound $N(n, k)$, we distinguish between two possible cases.

(i) There is at most one cell $D \in \mathcal{A}(C_H)$ such that r contains an extremal stabbing line $\ell \in \mathcal{L}[e_0, D]$. Then r contributes at most one unit to the sum $\sum_{D \in \mathcal{A}(C_H)} R[e_0, D]$. The total contribution in this case is $O(nk)$.

(ii) There are $m > 1$ cells $D \in \mathcal{A}(C_H)$ such that r contains an extremal stabbing line in $\mathcal{L}[e_0, D]$. Each such cell D contains a connected portion \tilde{r} of $r \cap D$, one of whose endpoints, call it a , belongs to the boundary of D and the other, call it b , represents an extremal stabbing line in $\mathcal{L}[e_0, D]$.

Choose a positive threshold t . If the depth of (the line represented by) a is less than t , we charge the appearance of r in D to a . Since a is contained in the boundary of D , there is a plane $h \in H$ such that a represents a line in the space \mathcal{L}_h of lines which pass through ℓ_0 and are parallel to h . Moreover, since r is a common tangent to P_0 and P and lines in \mathcal{L}_h have two degrees of freedom, a represents an extremal line at depth at most t , within \mathcal{L}_h . By Lemma 2.2, the number of extremal lines at depth 0 within \mathcal{L}_h is $O(n\beta_4(k))$. (More precisely, if h is defined by a polytope Q which is disjoint from ℓ_0 , then no line in \mathcal{L}_h has depth 0, and the above bound holds for lines at depth one. Nevertheless, the asymptotic bound of Clarkson and Shor continues to hold in this degenerate situation.) By applying the Clarkson–Shor




FIG. 6. A lower-bound construction, viewed from above (left). (The actual gaps between the pairs of parallel “plates” around ℓ_0 , relative to the width of the plates, are much smaller than depicted.) If \mathcal{P} consists of a single polygon, which is coplanar with ℓ_0 , then every pair of distinct vertices of P define an extremal stabbing line passing through them (right).

argument [12], as above, and using the fact that an extremal stabbing line in \mathcal{L}_h is defined by at most two polyhedra, we conclude that the number of such extremal lines at depth at most t , within \mathcal{L}_h , is $O(t^2 \cdot \mathbf{E}\{n_{\mathcal{R}}\beta_4(k/t)\})$, where $n_{\mathcal{R}}$ is the number of facets in a random sample of k/t polyhedra of \mathcal{P} . Since $\mathbf{E}\{n_{\mathcal{R}}\} = n/t$, the bound is $O(tn\beta_4(k/t))$. Since $|H| = O(k)$, the overall number of lines a , as above, of depth smaller than t is $O(tnk\beta_4(k/t)) = O(tnk\beta_4(k))$. Moreover, any such line a is charged at most a constant number of times, over all possible choices of e_0 and D .

If the depth of (the line represented by) a is at least t , we walk from b to a along \tilde{r} and collect at least t extremal lines at depth at most t contained in $\mathcal{L}[e_0, D]$. (These are lines on the regulus r which are tangent to a third polyhedron, one of at least t polyhedra that a misses.) We charge the appearance of r in D to those lines. Clearly, each line is charged in this manner only a constant number of times, over all choices of e_0 and D . Using (3.5) and taking case (i) also into account, we get

$$R(n, k) = O\left(nk + tnk\beta_4(k) + \frac{1}{t} \cdot t^3 N\left(\frac{n}{t}, \frac{k}{t}\right)\right).$$

Combining this with (3.4), we finally get a recurrence for $N(n, k)$:

$$(3.6) \quad N(n, k) \leq ctnk\beta_4^2(k) + ct^2 N\left(\frac{n}{t}, \frac{k}{t}\right)\beta_4(k)$$

for an appropriate constant c .

The recurrence terminates when k becomes smaller than t . In this case, we use the fact, which was proved in [9], that the complexity of the arrangement, within \mathcal{L} , of the upper and lower tangency surfaces σ_P^+, σ_P^- of all polyhedra is bounded by $O(nk^2) = O(nt^2)$, which thus also serves as an upper bound for $N(n, k)$.

Using standard analysis, such as that in [6, 17], it follows that the solution of the recurrence (3.6) is $N(n, k) = O(nk^{1+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε (the choice of ε affects the choice of t). This completes⁸ the proof of the upper bound in Theorem 3.1.

Lower bound. A lower bound construction for the complexity of $\mathcal{T}_{\ell_0}(\mathcal{P})$ is depicted in Figure 6 (left). We use k pairs of polyhedra, each of which is a thin and long vertical

⁸The fact that the bound that we derive is linear in n justifies the replacement of $n_{\mathcal{R}}$ by its expectation $\frac{n}{t}$ in the bound (3.5). Although this deserves a more formal argument, we omit the details, which are similar to those in earlier studies [17, 29].

plate, whose vertical edges and facets are all parallel to ℓ_0 . The polyhedra in each pair are parallel to each other and are situated symmetrically around ℓ_0 . In addition, we include in \mathcal{P} one drum-like polyhedral prism P_0 , whose axis is horizontal (i.e., orthogonal to ℓ_0) and which has $n \gg k$ long and narrow facets. Overall, we have $\Theta(k)$ polyhedra with a total of $\Theta(n)$ facets.

The polyhedra in $\mathcal{P} \setminus \{P_0\}$ are arranged so that as we rotate a line of \mathcal{L} around ℓ_0 by varying its θ -component while keeping its z - and φ -components within some reasonable range, we obtain $\Theta(k)$ distinct geometric permutations of the polyhedra in $\mathcal{P} \setminus \{P_0\}$ (see Figure 6 (left) and [11, 13] for similar constructions). With an appropriate choice of the layout, we can construct, for each of these permutations and for each edge e of P_0 , a transversal line in \mathcal{L} that realizes the permutation and is tangent to P_0 at e and to two of the other polyhedra at two respective vertical edges. We thus obtain a lower bound of $\Omega(nk)$ on the complexity of $\mathcal{T}_{\ell_0}(\mathcal{P})$. (The construction is not in general position, but can be transformed into general position by a small perturbation of its polyhedra.) \square

Handling degeneracies. The proof of Theorem 3.1 assumed the general position of ℓ_0 and the polyhedra of \mathcal{P} .

Degenerate configurations of $\mathcal{P} \cup \{\ell_0\}$ lead to degeneracies in the arrangement of the tangency surfaces $\{\sigma_P^+, \sigma_P^-\}_{P \in \mathcal{P}}$ within the three-dimensional space whose points represent lines of \mathcal{L} . For example, a single vertex or edge of $\mathcal{T}_{\ell_0}(\mathcal{P})$ may be contained in the tangency surfaces of any number of polyhedra, and any pair of edges or 2-faces of the tangency surfaces may be tangent to each other or overlap (partially or completely). In addition, any polyhedron $P \in \mathcal{P}$ may contain faces of dimensions one and two which are coplanar with ℓ_0 , in which case the the stabbing region $\mathcal{T}_{\ell_0}(\{P\})$ of P contains z -parallel boundary faces which are not contained in σ_P^+, σ_P^- . Similarly, for any polyhedron P_0 and one of its edges e_0 , we may encounter similar degeneracies in the arrangement of the tangency curves $\{\gamma_P^+, \gamma_P^-\}_{P \in \mathcal{P} \setminus \{P_0\}}$ (defined in the proof of Theorem 3.1) within the local two-dimensional space $\mathcal{L}[e_0] \subset \mathcal{L}$ whose points represent lines of \mathcal{L} tangent to P_0 at e_0 . To extend our analysis to degenerate configurations of $\mathcal{P} \cup \{\ell_0\}$, we apply a small perturbation of the polyhedra of \mathcal{P} , as described, e.g., in [5, 6] and [30, Theorem 7.3.1], which can only increase the number of faces (of dimensions zero, one, and two) bounding the full-dimensional components (cells) of $\mathcal{T}_{\ell_0}(\mathcal{P})$. To prevent loss of the lower-dimensional isolated components of $\mathcal{T}_{\ell_0}(\mathcal{P})$ (which are not possible in general position), we slightly enlarge the polyhedra of \mathcal{P} and then apply the perturbation. With some care, this expansion can be performed without changing the boundary topology of the full-dimensional components of $\mathcal{T}_{\ell_0}(\mathcal{P})$; see [5] for a similar argument. At the end, the elements of $\mathcal{P} \cup \{\ell_0\}$ are in *general position*, and the complexity of the transversal space $\mathcal{T}_{\ell_0}(\mathcal{P})$ is no smaller than that of the unperturbed collection, namely, $O(nk^{1+\varepsilon})$, for any $\varepsilon > 0$.

Theorem 3.1 implies the bound of $O(nk^{1+\varepsilon})$, for any $\varepsilon > 0$, on the number of extremal stabbing lines in \mathcal{L} that correspond to the boundary vertices of $\mathcal{T}_{\ell_0}(\mathcal{P})$. In degenerate settings, however, \mathcal{L} may contain additional extremal stabbing lines which do not correspond to these vertices. Each of these lines $\ell \in \mathcal{L}$ is tangent to some polyhedron $P \in \mathcal{P}$ but is represented on neither of the $\theta\kappa$ -monotone surfaces σ_P^+, σ_P^- . This situation arises when ℓ is tangent to P at its facet or edge f , which is coplanar with ℓ_0 , and ℓ is contained in the unique plane Π_{θ^*} , for some $0 \leq \theta^* \leq 2\pi$, containing both ℓ_0 and f . See Figure 6 (right). In our particular parametrization of \mathcal{L} , the corresponding surfaces σ_P^+ and σ_P^- contain only the points representing lines $\ell \in \mathcal{L}$ that are tangent to $P \cap \Pi_{\theta^*}$ in the plane Π_{θ^*} . In particular, such lines $\ell \in \mathcal{L}$ that overlap the relative interior of $f = P \cap \Pi_{\theta^*}$ are not represented on $\sigma_P^+ \cup \sigma_P^-$;

see also [20]. (These lines form a relatively open z -parallel 2-face on the boundary of $\mathcal{T}_{\ell_0}(\{P\})$.) As noticed in [9], in such a case $\mathcal{T}_{\ell_0}(\mathcal{P})$ may contain $\Omega(n^2)$ extremal stabbing lines of this kind. (This construction applies even when \mathcal{P} contains a single polyhedron P .)

4. The combinatorial complexity of $\mathcal{T}(\mathcal{P})$. In this section we establish the following theorem.

THEOREM 4.1. *Let \mathcal{P} be a collection of k convex polyhedra in \mathbb{R}^3 with a total of n facets. Then the set $\mathcal{T}(\mathcal{P})$ of line transversals of \mathcal{P} has complexity $O(n^2 k^{1+\varepsilon})$ for any $\varepsilon > 0$.*

Proof. Following [24, 27], we represent the transversal space $\mathcal{T}(P)$ as a sandwich region in \mathbb{R}^4 . To do so, we represent each line as a point in \mathbb{R}^4 . For example, we can represent each line (not orthogonal to the x -axis) $\{(x, y, z) \mid y = ax + b, z = cx + d\}$ by the point $(a, b, c, d) \in \mathbb{R}^4$. Similar to section 2, each polyhedron $P \in \mathcal{P}$ defines two abc -monotone partially defined (hyper-)surfaces σ_P^+, σ_P^- in 4-space, consisting, respectively, of points representing upper and lower tangents to P . In this four-dimensional parametrization, $\mathcal{T}(\mathcal{P})$ is the sandwich region lying below the lower envelope of $\{\sigma_P^+\}_{P \in \mathcal{P}}$ and above the upper envelope of $\{\sigma_P^-\}_{P \in \mathcal{P}}$; see, e.g., [24]. As in sections 2 and 3, we may assume with no loss of generality that the polyhedra of \mathcal{P} are in general position; otherwise we can enforce general position by slightly perturbing (and enlarging) the polyhedra so that the number of boundary faces of $\mathcal{T}_{\ell_0}(P)$ can only increase. Moreover, arguing as in section 2 (see also [24, 27]), the worst-case complexity of $\mathcal{T}(\mathcal{P})$ is dominated by the maximum possible number of its vertices, each of which corresponds to an extremal stabbing line of \mathcal{P} . We now proceed to bound the number of these lines.

Let then e_0 be a fixed edge of the boundary of some polyhedron $P_0 \in \mathcal{P}$, let ℓ_0 be the line which contains e_0 , and let \mathcal{L} be the space of lines passing through ℓ_0 . Clearly, any extremal stabbing line to \mathcal{P} that is tangent to P_0 at the relative interior of e_0 corresponds to a boundary vertex of $\mathcal{T}_{\ell_0}(\mathcal{P} \setminus \{P_0\})$, unless it either passes through an endpoint of e_0 or overlaps a facet of P_0 . Using Theorem 3.1 and summing over $O(n)$ possible choices of e_0 , we immediately derive the upper bound of $O(n^2 k^{1+\varepsilon})$, for any $\varepsilon > 0$, on the number of all extremal stabbing lines that are tangent to at least one polyhedron at the relative interior of an edge of it. The number of remaining extremal stabbing lines is $O(n^2)$ because each of them passes through two polyhedra vertices or overlaps two polyhedra facets.⁹ Theorem 4.1 now follows from the combination of all these bounds. \square

Remarks. (1) Unlike the restricted case analyzed in section 3, here we do not have a matching lower bound. The best lower bound is the trivial $\Omega(n^2)$ bound; see [9]. A different lower bound of $\Omega(nk^2)$ can be shown by slightly modifying the “standard” example of k pairwise-disjoint convex bodies in \mathbb{R}^3 which admit $\Omega(k^2)$ geometric permutations (realized by lines passing) through the origin; e.g., see [7]. In our construction, we take each convex set to be a convex polyhedron with $O(1)$ facets and denote the resulting collection by \mathcal{P}' . Then we place a sufficiently small polyhedron P_0 near the origin, where P_0 is a sufficiently flat pyramid whose base is an n -gon lying in some generic plane through the origin. For small enough P_0 , there are $\Omega(k^2)$ geometric permutations through each edge of the base of P_0 . Moreover, by our construction, “almost” all lines passing through edges of the base of P_0 are tangent to

⁹Recall that, because of the general position of \mathcal{P} , no polyhedron facet can be coplanar with a vertex of another polyhedron.

P_0 . (The exceptional lines are nearly parallel to the base of P_0 and can be ignored in the construction.) This easily implies that, for $\mathcal{P} = \mathcal{P}' \cup \{P_0\}$, the space $\mathcal{T}(\mathcal{P})$ has $\Omega(nk^2)$ complexity. Closing the gap between the resulting lower bound $\Omega(n^2 + nk^2)$ and the upper bound in Theorem 4.1 remains an interesting and challenging open problem.

(2) Note also that, unlike the case considered in section 3, here we do not have to distinguish between vertices of $\mathcal{T}(\mathcal{P})$ and extremal stabbing lines which are not vertices (as in the degenerate configuration depicted in Figure 6 (right) and discussed at the end of the previous section) because the analysis of Theorem 4.1 takes care of both kinds of lines.

5. Constructing $\mathcal{T}(\mathcal{P})$. Let us first describe an efficient randomized algorithm for the restricted space $\mathcal{T}_{\ell_0}(\mathcal{P})$.

THEOREM 5.1. *Let \mathcal{P} be a set of k convex polyhedra with a total of n facets, and let ℓ_0 be a fixed line. Then one can compute (the boundary representation of) the stabbing region $\mathcal{T}_{\ell_0}(\mathcal{P})$ of \mathcal{P} , within the set \mathcal{L} of lines passing through ℓ_0 , in $O(nk^{1+\varepsilon} \log n)$ randomized expected time for any $\varepsilon > 0$.*

Proof. As before, we assume general position of $\mathcal{P} \cup \{\ell_0\}$. This assumption involves no loss of generality. Extending the algorithm to handle degenerate inputs can be achieved by symbolic perturbation of the tangency surfaces σ_P^+, σ_P^- , which closely follows the one described in section 3 (see [20] for a similar treatment of degenerate configurations). Moreover, we assume, without loss of generality, that ℓ_0 is the z -axis. Recall that for each $P \in \mathcal{P}$, the graphs of σ_P^+ and σ_P^- are (θ, φ) -monotone surfaces which are comprised of monotone semialgebraic surface patches, each of which is a graph of a partially defined function of constant description complexity, representing upper (resp., lower) tangents to P at a fixed edge of its upper (resp., lower) boundary. For each edge e of P , we denote its corresponding tangency function (and its graph) by σ_e . The domain of σ_e is the region $D_e \cup \bar{D}_e$, where \bar{D}_e is obtained from D_e by mapping¹⁰ (θ, φ) to $(\pi + \theta, \pi - \varphi)$ and

$$(5.1) \quad D_e = \{(\theta, \varphi) \mid \theta_e^- \leq \theta \leq \theta_e^+, \tau_{f^-}(\theta) \leq \varphi \leq \tau_{f^+}(\theta)\},$$

where θ_e^- and θ_e^+ are the θ -coordinates of the two (clockwise and counterclockwise) endpoints of e , f^- and f^+ are the two facets of P incident to e , and τ_{f^-} (resp., τ_{f^+}) is the locus of all orientations (θ, φ) of lines parallel to f^- (resp., f^+); since, as is easily checked, the sets τ_{f^-}, τ_{f^+} are θ -monotone, the functional notation (5.1), as well as the determination of which of the adjacent facets is f^- and which is f^+ , are well defined. If e is a silhouette edge of P , i.e., it admits a supporting plane of P parallel to ℓ_0 , then we consider it as a pair of identical copies, one of which belongs to the upper portion of ∂P and the other to the lower portion of ∂P . In each of these cases, one of τ_{f^+} or τ_{f^-} , as appropriate, is set to $+\infty$ or $-\infty$, respectively.

As a preparatory step, we use the algorithm of Theorem 2.1 to construct $\mathcal{T}_{\ell_0}(\{P, Q\})$, for each pair of polyhedra $P, Q \in \mathcal{P}$, in a total of $O(nk \log n)$ time. In particular, for each edge e of the upper (resp., lower) boundary of some polyhedron P and for each $Q \in \mathcal{P} \setminus \{P\}$, the above algorithm computes the intersection of σ_Q^+ and of σ_Q^- with the (two-dimensional) patch σ_e representing upper (resp., lower) tangents to P at e (recall that silhouette edges are treated as both upper and lower edges).

The algorithm for constructing $\mathcal{T}_{\ell_0}(\mathcal{P})$ then proceeds as follows. First, we choose a fixed random permutation π of the polyhedra of \mathcal{P} . Next, we fix a polyhedron $P \in \mathcal{P}$

¹⁰That is, if an oriented line $\ell \in \mathcal{L}$ has orientation $(\theta, \varphi) \in D_e$, then the oppositely oriented copy of ℓ has orientation $(\theta + \pi, \pi - \varphi) \in \bar{D}_e$.

and a boundary edge e of P . Recall that σ_e consists of two connected patches of a constant-degree algebraic surface (in an appropriate reparametrization); one patch is defined over the domain D_e given in (5.1), and the other is defined over the symmetric domain \tilde{D}_e . Without loss of generality, we consider the portion of σ_e defined over D_e . To obtain $\mathcal{T}_{\ell_0}(\mathcal{P})$, it is sufficient to compute $\sigma_e \cap \mathcal{T}_{\ell_0}(\mathcal{P} \setminus \{P\})$ for all possible choices of P and e as above; the union of all these two-dimensional patches constitutes the boundary of $\mathcal{T}_{\ell_0}(\mathcal{P})$. We follow a standard randomized divide-and-conquer approach; see [30]. Let \mathcal{P}_R be the set of the first $\lfloor \frac{k-1}{2} \rfloor$ elements of $\mathcal{P} \setminus \{P\}$ in the permutation π , and let \mathcal{P}_B be the set of the $\lceil \frac{k-1}{2} \rceil$ remaining elements of $\mathcal{P} \setminus \{P\}$. We think of the elements of \mathcal{P}_R as “red” and of those of \mathcal{P}_B as “blue.” We start by recursively computing the “red region” $\mathcal{R} := \sigma_e \cap \mathcal{T}_{\ell_0}(\mathcal{P}_R)$ and the “blue region” $\mathcal{B} := \sigma_e \cap \mathcal{T}_{\ell_0}(\mathcal{P}_B)$ within σ_e . Then $\sigma_e \cap \mathcal{T}_{\ell_0}(\mathcal{P} \setminus \{P\}) = \mathcal{R} \cap \mathcal{B}$. We denote this intersection by \mathcal{G} and refer to it as the “green region.” Let n_R , n_B , and n_G denote the total number of vertices and edges on the boundary of \mathcal{R} , \mathcal{B} , and \mathcal{G} , respectively. Then \mathcal{G} can be computed in $O((n_R + n_B + n_G) \log(n_R + n_B + n_G))$ time using (an appropriate variant of) a standard planar sweep algorithm; see, e.g., [30]. We repeat this procedure for all possible choices of P and e . The above algorithm has expected running time $O(N \log N \log k)$, where N is the overall expected number of vertices and edges which are constructed by the algorithm at any recursive step, over all choices of P and e .

We next establish an upper bound on N . Note that any edge of $\mathcal{T}_{\ell_0}(\mathcal{P})$ which does not have any incident vertex and which appears in one of the regions constructed by the algorithm appears in $\mathcal{T}_{\ell_0}(\{P, Q\})$, for some pair of polyhedra $P, Q \in \mathcal{P}$, so, according to Theorem 2.1, the overall number of such edges is $O(nk)$. The appearance of any other edge can be charged to that of an incident vertex so that each vertex is charged at most a constant number of times, assuming general position. (The argument is somewhat subtle because an edge can be split many times into subedges. However, if we charge each newly formed subedge to a *new* delimiting vertex, we ensure that every vertex is charged only $O(1)$ times, as claimed.) Therefore, it remains to bound the expected number of vertices which appear in some region \mathcal{G} throughout the execution of the algorithm (over all P and e). Clearly, the overall number of vertices (i.e., breakpoints) of the curves $\sigma_e \cap \sigma_Q^+$, $\sigma_e \cap \sigma_Q^-$, over all choices of P , $e \in \partial P$, and $Q \in \mathcal{P} \setminus \{P\}$, is $O(nk)$, since those vertices appear in the respective regions $\mathcal{T}_{\ell_0}(\{P, Q\})$. Hence, it suffices to bound the expected number of vertices defined by triples of polyhedra. For a fixed choice of P and e , any such vertex v belongs to the intersection of $\sigma_e \cap \sigma_Q^+$ (or $\sigma_e \cap \sigma_Q^-$) and $\sigma_e \cap \sigma_R^+$ (or $\sigma_e \cap \sigma_R^-$) for some pair of distinct polyhedra $Q, R \in \mathcal{P} \setminus \{P\}$. Let \mathcal{K}_v denote the set of polyhedra which are not stabbed by the extremal line ℓ_v corresponding to v . Thus, $|\mathcal{K}_v|$ is equal to the depth of ℓ_v . As is easily verified, a necessary condition for v to appear in \mathcal{G} , at some recursive step, is that no polyhedron of \mathcal{K}_v appears between Q and R in π . We say that v has depth t if ℓ_v has depth t . Restricting π to the $t+2$ polyhedra Q, R , and the t polyhedra that ℓ_v misses, it follows that the probability that v is constructed by the algorithm, when processing σ_e , is at most

$$\frac{2(t+1)!}{(t+2)!} = \frac{2}{t+2}.$$

Note that v may also arise when processing the respective edges of Q and of R which define v , so we actually need to multiply the corresponding probabilities and expectations by 3.

Let $N_{\leq t}$ denote the number of extremal lines at depth at most t , summed over all P and e . Then the overall number of extremal lines at depth t , for $0 \leq t \leq k - 3$, is $N_{\leq t} - N_{\leq t-1}$, where $N_{\leq -1} = 0$ and the expected number of vertices at depth t , which appear in some region \mathcal{G} throughout the algorithm, again over all P and e , is at most

$$\frac{6(N_{\leq t} - N_{\leq t-1})}{t + 2}.$$

Summing over all depths $0 \leq t \leq k - 3$, we get

$$(5.2) \quad N \leq 6 \sum_{0 \leq t \leq k-3} \frac{(N_{\leq t} - N_{\leq t-1})}{t + 2} + O(nk).$$

Rearranging the sum, we get

$$(5.3) \quad \begin{aligned} N &\leq 6 \sum_{0 \leq t \leq k-4} N_{\leq t} \left(\frac{1}{t + 2} - \frac{1}{t + 3} \right) + \frac{6}{k - 1} N_{\leq k-3} + O(nk) \\ &= 6 \sum_{0 \leq t \leq k-4} \frac{N_{\leq t}}{(t + 2)(t + 3)} + \frac{6}{k - 1} N_{\leq k-3} + O(nk). \end{aligned}$$

By Theorem 3.1, for any $\varepsilon > 0$, $N_{\leq 0} = O(nk^{1+\varepsilon})$. Hence, by the Clarkson–Shor probabilistic argument [12], using (3.5), $N_{\leq t} = O(tnk^{1+\varepsilon})$ for all $1 \leq t \leq k - 3$; see [6] for a similar argument. Plugging this into (5.3), we obtain that $N = O(nk^{1+\varepsilon})$ for any $\varepsilon > 0$. (The substitution yields a harmonic series which adds a factor $O(\log k)$ to the bound, but this factor is “swallowed” by $k^{1+\varepsilon}$, by slightly increasing ε , still keeping it arbitrary small.) Our algorithm thus runs in $O(nk^{1+\varepsilon} \log n)$ randomized expected time for any $\varepsilon > 0$. (Again, the $O(\log k)$ factor yielded by the divide-and-conquer process is subsumed in the factor $k^{1+\varepsilon}$.) \square

The following theorem follows from Theorems 5.1 and 4.1.

THEOREM 5.2. *Let \mathcal{P} be a set of k convex polyhedra with a total of n facets. Then one can compute (the boundary representation of) the stabbing region $\mathcal{T}(\mathcal{P})$ in $O(n^2 k^{1+\varepsilon} \log n)$ randomized expected time for any $\varepsilon > 0$.*

Remark. Technically, the randomized procedure of Theorem 5.1 (repeated to each of the $O(n)$ lines passing through the polyhedra edges) yields most of the boundary features of $\mathcal{T}(\mathcal{P})$ but not all of them. For example, it excludes features that involve tangencies at the polyhedra vertices only, or some features involving tangencies where the line overlaps a facet, and so on. These extra features can be constructed by additional ad hoc procedures, within the same running time bound. We also omit details about gluing the features yielded by Theorem 5.1 to each other to obtain a global representation of $\mathcal{T}(\mathcal{P})$.

6. Extensions. Assume that \mathcal{P} is a collection of k pairwise disjoint convex polyhedra in \mathbb{R}^3 , with a total of n facets. As in the preceding sections, we denote by \mathcal{L} the space of lines passing through a fixed line ℓ_0 and assume with no loss of generality that \mathcal{P} and ℓ_0 are in general position. In this section, we establish slightly improved upper bounds on the complexity of the stabbing region $\mathcal{T}_{\ell_0}(\mathcal{P})$ in the following two cases: the case where all polyhedra of \mathcal{P} are disjoint from ℓ_0 , and the case where the polyhedra of \mathcal{P} are all unbounded in one of the two directions parallel to ℓ_0 . In both cases, we also provide deterministic algorithms for efficiently computing $\mathcal{T}_{\ell_0}(\mathcal{P})$.

6.1. ℓ_0 is disjoint from all polyhedra in \mathcal{P} . When the polyhedra of \mathcal{P} are disjoint from each other and from ℓ_0 , we establish an improved bound on the complexity of the restricted space $\mathcal{T}_{\ell_0}(\mathcal{P})$ (provided $k \ll \sqrt{n}$).

THEOREM 6.1. *Let \mathcal{P} be a set of k pairwise disjoint convex polyhedra with a total of n facets, and let ℓ_0 be a fixed line disjoint from all polyhedra of \mathcal{P} . Then the complexity of the stabbing region $\mathcal{T}_{\ell_0}(\mathcal{P})$ of \mathcal{P} in the space \mathcal{L} of lines passing through ℓ_0 is $O((nk + k^3)\beta_4(k))$. Moreover, one can compute (the boundary representation of) $\mathcal{T}_{\ell_0}(\mathcal{P})$ in $O((nk + k^3)(\log n + \alpha(k) \log k))$ deterministic time.*

Proof. To bound the number of extremal stabbing lines in \mathcal{L} , we return to the proof of Theorem 3.1. Recall that, in the current context where ℓ_0 is disjoint from all the polyhedra of \mathcal{P} , H is a collection of k planes, each separating a polyhedron $P \in \mathcal{P}$ from ℓ_0 . Clearly, we can choose H such that all of its planes contain ℓ_0 .

Let e_0 be a boundary edge of some polyhedron $P_0 \in \mathcal{P}$, and let D be a cell in the arrangement $\mathcal{A}(C_H)$, as defined above. Recall that $R[e_0, D]$ is defined to be the number of arcs that actually show up on the boundary of the sandwich region (3.2), at points which represent lines in $\mathcal{L}[e_0, D]$. Then $R[e_0, D]$ is bounded by the number of reguli of lines (in \mathcal{L}) that are tangent to P_0 at e_0 and to another polyhedron $P \in \mathcal{P}$ at some fixed boundary edge e , and contain a line with direction in D . (Here we relax the requirement that such a regulus contains a *stabbing* line with direction in D .) For each such regulus r , we denote by r^* the locus, on \mathbb{S}^2 , of the orientations of the lines in r . Recall that we consider only those portions of the reguli consisting of lines that pass through the respective edges e_0 and e .

Recall that the number of extremal stabbing lines (in \mathcal{L}) is

$$O\left(\sum_{P_0 \in \mathcal{P}, e_0 \in \partial P} \sum_{D \in \mathcal{A}(C_H)} R[e_0, D] \beta_4(k)\right).$$

As noted in section 3, there is a total of $O(nk)$ reguli over all choices of e_0 . Hence, the sum

$$(6.1) \quad \sum_{P_0 \in \mathcal{P}, e_0 \in \partial P} \sum_{D \in \mathcal{A}(C_H)} R[e_0, D]$$

is bounded by $O(nk)$ plus a constant times the number of crossings of reguli r and boundaries of cells $D \in \mathcal{A}(C_H)$. Any such intersection corresponds to a distinct (oriented) line, *contained* in some plane $h \in H$ and tangent to a pair of polyhedra of \mathcal{P} . (By the choice of H , a line in \mathcal{L} is parallel to a plane $h \in H$ if and only if it lies in h .) Since the polyhedra of \mathcal{P} are pairwise disjoint, the number of such lines is $O(k^2)$ for a fixed $h \in H$. Since $|H| = k$, we get that the sum in (6.1) is $O(nk + k^3)$. Hence, the number of extremal stabbing lines in \mathcal{L} , and thus the complexity of the stabbing region $\mathcal{T}_{\ell_0}(\mathcal{P})$, is $O((nk + k^3)\beta_4(k))$. Note that this bound constitutes an improvement over Theorem 3.1 when $k = O(\sqrt{n})$.

To compute the extremal stabbing lines in \mathcal{L} , we first compute the set of all reguli r , in $O(nk \log n)$ overall time, using the algorithm of Theorem 2.1. Then, for each such regulus r , we trace r^* through the arrangement $\mathcal{A}(C_H)$ in $O((nk + k^3) \log k)$ total time. To do so, we prepare $\mathcal{A}(C_H)$ for point location queries and then locate, for each regulus r , the cell of $\mathcal{A}(C_H)$ that contains the orientation of an “endpoint” line of r . This takes $O(nk \log k)$ time. Then, for each plane $h \in H$, we intersect the polyhedra of \mathcal{P} with h and find the $O(k^2)$ common tangent lines (in h) to pairs of these intersections. This takes $O(n + k^2 \log n)$ time for each h for a total of $O(nk + k^3 \log n)$ time. Then we take each such common tangent and locate the arc of $\mathcal{A}(C_H)$ (on C_h) which contains

its orientation. Combining the outputs of this and the preceding step, we obtain the sets $R[e_0, D]$, for all edges e_0 and cells D , in overall $O((nk + k^3) \log n)$ time (see the preceding analysis for justification of this procedure). Finally, we compute, for each choice of P_0 , e_0 , and D , the vertices of the sandwich region, defined in the proof of Lemma 3.1, in $O(R[e_0, D]\beta_3(k) \log k) = O(R[e_0, D]\alpha(k) \log k)$ time using the algorithm of Hershberger [19]; see also [30, Theorem 6.5]. Summing over all choices of P_0 , e_0 , and $D \in \mathcal{A}(C_H)$, the total construction cost is $O((nk + k^3)\alpha(k) \log k)$ time. We can thus compute the extremal stabbing lines in \mathcal{L} in $O((nk + k^3)(\log n + \alpha(k) \log k))$ total time.

Computing the other features of $\mathcal{T}_{\ell_0}(\mathcal{P})$, such as edges without incident vertices, can also be done within this time bound (see section 2 for details). This completes the proof of the theorem. \square

6.2. The polyhedra of \mathcal{P} are unbounded in a direction parallel to ℓ_0 .

For this subsection, we represent lines in \mathcal{L} as in section 2, and we establish the following result, which holds for more general collections of convex polyhedra.

THEOREM 6.2. *Let \mathcal{P} be a collection of k pairwise disjoint convex polyhedra in \mathbb{R}^3 having a total of n facets, and let ℓ_0 be a fixed line. The number of vertices of the upper envelope E_U of the partially defined functions σ_P^- , for $P \in \mathcal{P}$ (as defined in section 2), is $O(nk\beta_4(k))$. Moreover, there is a deterministic algorithm which computes E_U in $O(nk(\log n + \alpha(k) \log k))$ time.*

Proof. As in section 2, we assume that ℓ_0 is the z -axis. Fix a polyhedron P of \mathcal{P} , and consider the graph of the function $z = \sigma_P^-(\theta, \varphi)$, which represents lower tangents to P in \mathcal{L} . We denote this graph also as σ_P^- and assume, without loss of generality, that σ_P^- is not empty. For each $Q \in \mathcal{P} \setminus \{P\}$, let σ_{PQ}^- denote the (relatively open) portion of σ_P^- that lies below σ_Q^- (in the z -direction); that is, σ_Q^- is higher than σ_P^- over this portion and thus “prevents” σ_P^- from attaining E_U at these points. Hence, the portion of σ_P^- that appears on E_U is the complement of the union of the regions σ_{PQ}^- for $Q \in \mathcal{P} \setminus \{P\}$.

We fix P and Q , and we study in more detail the structure of σ_{PQ}^- . Fix some value θ_0 of θ , and consider the cross sections of σ_P^- and σ_Q^- at θ_0 . Any line ℓ with $\theta = \theta_0$ lies in the plane Π_{θ_0} . Any point on the intersection curve $\gamma_{PQ} = \sigma_P^- \cap \sigma_Q^-$, with $\theta = \theta_0$, corresponds to a line in Π_{θ_0} , which is a common lower tangent to $P(\theta_0)$ and $Q(\theta_0)$. Since these two convex polygons are disjoint, they have at most two common lower tangents.

Moreover, the only way in which they can have two common lower tangents is when one of $P(\theta_0)$, $Q(\theta_0)$ lies “fully above” the other. Formally, this happens if and only if (i) the vertical projections of $P(\theta_0)$ and $Q(\theta_0)$ (on the axis of Π_{θ_0} orthogonal to ℓ_0) are nested within each other, and (ii) the polygon with the larger projection, say, $Q(\theta_0)$, lies above the other polygon $P(\theta_0)$. We then say that Q *overshadows* P (at θ_0) and denote this as $P \prec Q$ (at θ_0). Clearly, for any fixed θ_0 , this is a partial order; see Figure 7 for an illustration.

Suppose then that $P \prec Q$ at θ_0 . Then, as noted, γ_{PQ} meets $\theta = \theta_0$ at two points (θ_0, φ_1) , (θ_0, φ_2) , with $\varphi_1 < \varphi_2$. Moreover, as is easily checked, σ_{PQ}^- is the portion of σ_P^- (strictly) between φ_1 and φ_2 . This is a bad scenario because then (the cross-section at $\theta = \theta_0$ of) σ_{PQ}^- is a “strip,” and the union of such “strips” can have quadratic complexity; see Figure 8 (left) for an illustration.

Before proceeding to address this issue, we note that all the other cases are “good”: If neither of $P(\theta_0)$, $Q(\theta_0)$ overshadows the other, then they have at most one common




FIG. 7. Two polygons $P_1(\theta)$ and $P_2(\theta)$ can have two common lower tangents only when one of them, P_2 , overshadows P_1 .




FIG. 8. In the bad scenario (left), to be avoided, the regions σ_{PQ}^- may appear as “strips” within σ_P^- (as defined in the text), and their union may have quadratic complexity. In the good scenario (right), we consider only regions σ_{PQ}^- that appear as “half-planes” (or as pairs of disjoint half-planes) within σ_P^- (as defined in the text), in which case the complement of their union (shown unshaded) behaves like a sandwich region and has nearly linear complexity.

lower tangent line and then σ_{PQ}^- , at θ_0 , is either the portion $\varphi > \varphi_0$ or the portion $\varphi < \varphi_0$, where (θ_0, φ_0) is the intersection point of γ_{PQ} with $\theta = \theta_0$ (if it exists at all; if not, σ_{PQ}^- is either empty at $\theta = \theta_0$ or consists of the entire range of φ). If, on the other hand, $Q \prec P$ at θ_0 , then, arguing as above, σ_{PQ}^- consists of the portion of σ_P^- (strictly) above φ_2 and of the portion below φ_1 . In either case, the union of these good regions σ_{PQ}^- is the union of “half-planes,” where each half-plane consists of those points whose $\theta\varphi$ -projections lie above some θ -monotone curve (*upper half-planes*) or below such a curve (*lower half-planes*); here “above” and “below” are with respect to the φ -direction. Thus, for the “good” interactions involving P , the complement of the union is a (two-dimensional) sandwich region between the upper envelope of the (curves bounding the) lower half-planes and the lower envelope of the (curves bounding the) upper half-planes. This property will be crucial for the analysis of the number of vertices of E_U . Note that, when $Q \prec P$, σ_{PQ}^- is the union of two half-planes, so we may have up to $2(k-1)$ half-planes which form the sandwich region.

Our strategy is thus as follows. Let v be a vertex of the upper envelope E_U , incident to three surfaces, which we denote as $\sigma_P^-, \sigma_Q^-, \sigma_R^-$, corresponding to three respective polyhedra $P, Q, R \in \mathcal{P}$. We claim that there exists at least one surface, say,




FIG. 9. A critical line tangent to three polyhedra of \mathcal{P} from below. Here $P_1 \prec P_2 \prec P_3$, so the overshadowing relationships are “good” for P_3 .

$\sigma_{\bar{P}}$, such that neither of the relations $P \prec Q, P \prec R$ holds (at the θ -coordinate of v). Indeed, if this does not hold for $\sigma_{\bar{P}}$, then, say, $P \prec Q$. If $Q \prec R$, then, by transitivity, we also have $P \prec R$, so R has the desired property (using the antisymmetry of \prec). If, on the other hand, $Q \not\prec R$, then Q has the desired property. See Figure 9 for an illustration.

Suppose, without loss of generality, that P has this property; that is, neither of the relations $P \prec Q, P \prec R$ holds. By the preceding observations, both $\sigma_{\bar{P}Q}$ and $\sigma_{\bar{P}R}$ are good at (the vicinity of) v , so v appears as a vertex of the sandwich region of the good regions $\sigma_{\bar{P}S}$ for $S \in \mathcal{P} \setminus \{P\}$.

It therefore suffices to bound the overall complexity of the sandwich regions of the good portions $\sigma_{\bar{P}Q}$, within the respective surfaces $\sigma_{\bar{P}}$.

Let us study in detail the structure of a good portion $\sigma_{\bar{P}Q}$ within $\sigma_{\bar{P}}$. It is a “partially defined” half-plane, bounded by two extreme θ -values $\theta_1 < \theta_2$, and lying above or below a θ -monotone portion of the curve γ_{PQ} . If $Q \not\prec P$, then γ_{PQ} itself is θ -monotone. If $Q \prec P$, γ_{PQ} has two θ -monotone arcs, and, as already mentioned, we consider the portions of $\sigma_{\bar{P}Q}$ that lie above the top arc and below the bottom arc, respectively, as two separate half-planes.

At each of the delimiting orientations θ_1, θ_2 , the overshadowing relation between $P(\theta)$ and $Q(\theta)$ starts or stops holding. When this happens, the vertical projections of $P(\theta)$ and $Q(\theta)$ have a common endpoint. Alternatively, if we project the polyhedra of \mathcal{P} onto some plane h_0 orthogonal to ℓ_0 , then θ_1 and θ_2 are the θ -coordinates of intersection points between the projected *silhouettes* of P and Q . As is easily checked, the total number of such intersection points, over all pairs P, Q , is only $O(nk)$. (That is, in an arrangement of k convex polygons in the plane, with a total of n edges, there are at most nk points where a pair of polygon boundaries cross each other.)

With an appropriate reparametrization of θ and φ , each curve γ_{PQ} is piecewise algebraic. Each piece corresponds to a regulus of lower tangent lines of P, Q which touch P, Q at two fixed respective edges e, e' (and pass through ℓ_0). A breakpoint of γ_{PQ} corresponds to situations where one of the incident edges e, e' changes. As already argued in the proof of Theorem 3.1 (see also Theorem 2.1), the number of such breakpoints is $O(nk)$.

To recap, on each surface $\sigma_{\bar{P}}$, for $P \in \mathcal{P}$, we need to bound the complexity of a sandwich region between the lower envelope of up to $k - 1$ θ -monotone curves and

the upper envelope of up to $k - 1$ other θ -monotone curves. Each curve consists of several arcs, each representing tangency with some fixed pair of edges of two of the polyhedra, and the total number of these arcs is $O(nk)$ over all surfaces σ_P^- .

We next observe that any pair of these arcs, say, $\gamma' \subset \gamma_{PQ}, \gamma'' \subset \gamma_{PR}$, on the same surface σ_P , intersect at most twice. Indeed, any such intersection point p represents a line that passes through four lines, namely, ℓ_0 , the line containing the edge e of P touched by the line ℓ_p represented by p , and the lines containing the edges of Q and R , corresponding to γ', γ'' , respectively, touched by ℓ_p . Since at most two lines can touch four distinct lines (not all lying on a common regulus; see [32]), the claim follows.

It therefore follows, using similar arguments to those given above, that the overall complexity of all the sandwich regions, over all surfaces σ_P , is $O(n\lambda_4(k)) = O(nk\beta_4(k))$. This completes the proof of the upper bound on the complexity of E_U .

The above analysis immediately implies the existence of a deterministic algorithm for computing the vertices of $\mathcal{T}_{\ell_0}(\mathcal{P})$ in $O(nk(\log n + \alpha(k) \log k))$ time. Indeed, we can obtain the curves γ_{PQ} , for every pair of distinct polyhedra $P, Q \in \mathcal{P}$, using the algorithm of Theorem 2.1. We then compute the good portion of each γ_{PQ} for each pair of distinct polyhedra $P, Q \in \mathcal{P}$. Finally, we compute, for each $P \in \mathcal{P}$, the corresponding sandwich region bounded by the good portions of the curves γ_{PQ} for $Q \in \mathcal{P} \setminus \{P\}$. As can be easily verified, the first two steps take $O(nk \log n)$ time, and the last step takes a total of $O(nk\alpha(k) \log k + nk \log n)$ time, using Hershberger’s algorithm [19], where the second term is the cost of sorting arc endpoints for each sandwich region separately. The other features of E_U can also be computed (deterministically) within this time bound. This completes the proof of the theorem. \square

COROLLARY 6.3. *If all the polyhedra of \mathcal{P} are unbounded in one of the two directions parallel to ℓ_0 , then the complexity of $\mathcal{T}_{\ell_0}(\mathcal{P})$ is $O(nk\beta_4(k))$, and there is a deterministic algorithm which computes (the boundary representation of) $\mathcal{T}_{\ell_0}(\mathcal{P})$ in $O(nk(\log n + \alpha(k) \log k))$ time.*

Proof. Suppose that the direction in which the polyhedra are unbounded is the positive z -direction. In this case the functions σ_P^+ are all undefined, and the stabbing region is the region above the upper envelope E_U of the functions σ_P^- for $P \in \mathcal{P}$. The claim is then immediate from Theorem 6.2. \square

Remarks. It would be nice to extend Theorem 6.2 to the sandwich region between the two envelopes E_U and E_L , as in (2.1). The difficulty in obtaining such an extension is that the overshadowing relation is not a necessary condition for two cross-sections $P(\theta_0), Q(\theta_0)$ to have two common tangents which are, say, lower tangents to $P(\theta_0)$ and upper tangents to $Q(\theta_0)$; see Figure 10. Hence, the reduction to two-dimensional sandwich regions, as in the proof of Theorem 6.2, does not carry over to the “mixed” case. The analysis in section 3 is the way we have managed to overcome this problem,




FIG. 10. *The overshadowing relation is not necessary for the cross-sections $P(\theta), Q(\theta)$ to have two common tangents.*

at the cost of a slight degradation in the bound and a more involved analysis. See also a discussion in section 8.

7. A note on geometric permutations. In this section we apply the machinery developed in this paper to establish a number of results on geometric permutations.¹¹

THEOREM 7.1. *Let \mathcal{P} be a collection of k pairwise disjoint convex objects in \mathbb{R}^3 , one of which is a line ℓ_0 . Then the number of geometric permutations induced by \mathcal{P} is $O(k^3)$.*

Proof. For each convex set $C \in \mathcal{P} \setminus \{\ell_0\}$, let h_C be a plane passing through ℓ_0 and disjoint from C . The collection of the planes h_C partitions \mathbb{R}^3 into (up to) $2(k - 1)$ wedges, with the following property, which extends the discussion in the preceding section: For each of these wedges W and for any directed line ℓ , which is a transversal of \mathcal{P} so that its forward ray ℓ^+ from ℓ_0 is contained in W , the set \mathcal{P}^+ of the objects of $\mathcal{P} \setminus \{\ell_0\}$ that are crossed by ℓ^+ is the same for all such lines ℓ (it depends only on W and not on the choice of ℓ). Clearly, this also holds for the complementary set \mathcal{P}^- of objects crossed by the backward ray ℓ^- of ℓ . Denote by \mathcal{W} the set of these wedges.

Next, for each pair of objects $C, C' \in \mathcal{P} \setminus \{\ell_0\}$, choose an arbitrary plane $h_{C,C'}$ which separates C and C' ; for simplicity, we may assume that $h_{C,C'}$ is not parallel to ℓ_0 . Let $z_{C,C'}$ denote the point $\ell_0 \cap h_{C,C'}$. We may also assume that all the points $z_{C,C'}$ are distinct. These points partition ℓ_0 into $\binom{k-1}{2} + 1$ intervals, and we denote by \mathcal{I} the set of these intervals; see Figure 11.

Each directed line ℓ which passes through ℓ_0 can be labeled by the pair $(W, I) \in \mathcal{W} \times \mathcal{I}$, where W is the wedge containing the forward ray of ℓ and I is the interval containing the intercept of ℓ with ℓ_0 . We claim that, for each fixed pair (W, I) , all transversals of \mathcal{P} labeled by (W, I) (if only exist) generate the same geometric permutation. Indeed, let ℓ be such a transversal. As argued above, the sets $\mathcal{P}^+, \mathcal{P}^-$, consisting of those objects of $\mathcal{P} \setminus \{\ell_0\}$ crossed by the respective forward and backward rays ℓ^+, ℓ^- of ℓ , are fixed and are independent of ℓ . Moreover, we can sort, in a unique manner independent of ℓ , the elements of \mathcal{P}^+ along ℓ^+ , as follows. For any pair of sets $C, C' \in \mathcal{P}^+$, consider the separating plane $h_{C,C'}$. By construction, I lies fully on one side of $h_{C,C'}$, say, the side containing C . But then C must precede C' along ℓ^+ , for otherwise ℓ^+ would have to cross $h_{C,C'}$ twice, once from a point in ℓ_0 (i.e., in I) to a point in C' and then back to a point in C . See Figure 11. Hence the order of any two elements of \mathcal{P}^+ along ℓ^+ is fixed for all lines ℓ labeled by (W, I) ,




FIG. 11. Geometric permutations through ℓ_0 : the interval I lies on one side of $h_{C,C'}$.

¹¹The proof of Theorem 7.1 uses a technique reminiscent of, and borrowing some ideas from, the analysis in [7].

so the suffix of the permutation induced by ℓ , from ℓ_0 on, is fixed too. A symmetric argument shows that the prefix of the permutation, up to ℓ_0 , is also fixed, so the permutation itself is fixed for all lines labeled by (W, I) . Since the number of labels is $O(k^3)$, the theorem follows. \square

Remarks. (a) Clearly, the proof shows that the number of geometric permutations is at most $O(Hk)$, where H is the smallest cardinality of a set of planes which separate \mathcal{P} (that is, each pair of objects of \mathcal{P} is separated by some plane in the set). Unfortunately, there are constructions of sets of k pairwise disjoint polyhedra for which $H = \Theta(k^2)$. For example, consider the hyperbolic paraboloid $z = xy$, and draw on it two sets, L_1, L_2 , of generating lines, where L_1 consists of the lines $x = i, z = iy$, for $i = 1, \dots, k/2$, and L_2 consists of the lines $y = j, z = jx$, for $j = 1, \dots, k/2$. Now keep L_1 intact, and shift L_2 upward by some $\varepsilon > 0$. It is then easy to check that if ε is chosen sufficiently small, for each pair of lines $\ell_1 \in L_1, \ell_2 \in L_2$, any plane that separates ℓ_1 and ℓ_2 must intersect all the other lines in $L_1 \cup L_2$, so $H \geq k^2/4$.

(b) Note that the proof of Theorem 7.1 fails if we replace ℓ_0 by a *line segment*, because it may then be impossible to separate the sets of \mathcal{P} from ℓ_0 by planes passing through ℓ_0 .

Geometric permutations induced by lines passing through a fixed line. We next establish a slightly inferior bound for a slightly more general setup, as follows.

THEOREM 7.2. *Let \mathcal{P} be a collection of k pairwise disjoint convex objects in \mathbb{R}^3 , and let ℓ_0 be a line in \mathbb{R}^3 (not necessarily disjoint from the elements of \mathcal{P}). Then the number of geometric permutations of \mathcal{P} induced by lines passing through ℓ_0 is $O(k^{3+\varepsilon})$ for any $\varepsilon > 0$.*

In particular, this implies a near-cubic upper bound on the number of geometric permutations induced by a collection \mathcal{P} of k pairwise disjoint convex objects, one of which is a *line segment*.

Proof. For each pair of objects $C, C' \in \mathcal{P}$ we select a pair of parallel planes $h_{C, C'}$ and $h_{C', C}$ separating them, where $h_{C, C'}$ is closer to C and $h_{C', C}$ is closer to C' . Let $H_{C, C'}$ denote the half-space bounded by $h_{C, C'}$ and containing C , and let $H_{C', C}$ denote the half-space bounded by $h_{C', C}$ and containing C' . For each $C \in \mathcal{P}$ let P_C be the polyhedron $\bigcap_{C' \in \mathcal{P} \setminus \{C\}} H_{C, C'}$. Then $C \subseteq P_C$, for each $C \in \mathcal{P}$, and the collection $\tilde{\mathcal{P}} = \{P_C \mid C \in \mathcal{P}\}$ consists of a total of k pairwise disjoint convex polyhedra, having a total of $O(k^2)$ facets. Plugging this into Theorem 3.1, we get that the combinatorial complexity of $\mathcal{T}_{\ell_0}(\tilde{\mathcal{P}})$ is $O(k^{3+\varepsilon})$, for any $\varepsilon > 0$, which also bounds the number of geometric permutations of $\tilde{\mathcal{P}}$ induced by lines passing through ℓ_0 . The theorem then follows since each transversal of \mathcal{P} is also a transversal of $\tilde{\mathcal{P}}$. \square

8. Conclusion. In this paper we obtained an improved bound, close to n^2k , on the combinatorial complexity of the set $\mathcal{T}(\mathcal{P})$ of line transversals to a collection \mathcal{P} of k convex polyhedra in \mathbb{R}^3 with a total of n facets, and we showed how to compute $\mathcal{T}(\mathcal{P})$ in comparable randomized expected time. We reduce this general problem to the restricted instance in which line transversals of \mathcal{P} are constrained to pass through some fixed line ℓ_0 . Specializing to this restricted instance, we obtain *nearly tight* bounds, close to nk , on the maximum complexity of the stabbing region $\mathcal{T}_{\ell_0}(\mathcal{P})$ within the space \mathcal{L} of lines that pass through ℓ_0 . Our analysis combines the classical charging scheme previously used to bound various substructures in arrangements of surfaces in \mathbb{R}^d , multiple examples of which can be found in [30], with (also classical) methods previously used to study geometric permutations of pairwise disjoint convex sets, as in [7, 33]. An advantage of our main analysis is that it does not assume that the polyhedra of \mathcal{P} are pairwise disjoint.

There are several challenging open problems for further research, including the following:

(i) Close the gap between the lower and the upper bounds for the complexity of the stabbing region $\mathcal{T}(\mathcal{P})$, as given in section 4.

(ii) Devise a deterministic algorithm for computing the stabbing region $\mathcal{T}(\mathcal{P})$, which runs in time close to $O(n^2k)$. In particular, do so for the restricted case of $\mathcal{T}_{\ell_0}(\mathcal{P})$, with running time close to $O(nk)$.

Finally, we mention the following general open problem: We are given two collections \mathcal{F}, \mathcal{G} , each of k bivariate functions, whose graphs consist of a total of $O(n)$ patches of constant description complexity, and the complexity of the arrangement of the graphs of any pair or triple of functions in $\mathcal{F} \cup \mathcal{G}$ is at most proportional to the overall number of their patches. Is it true that the complexity of the sandwich region between the upper envelope of \mathcal{F} and the lower envelope of \mathcal{G} is $O(nk^{1+\varepsilon})$ for any $\varepsilon > 0$? The lack of such a result forced us to go into the intricate analysis of section 3. We note that the complexity of the *overlay* of the minimization diagrams (cf. [24]) of \mathcal{F} and of \mathcal{G} can be $\Omega(n^2)$. Nevertheless, one can easily derive an upper bound of $O(nk^2)$ on the complexity of the sandwich region, which is only linear in n , for any fixed k .¹² This shows that the two problems of the overlay of minimization diagrams, and of sandwich regions, are significantly different in the above context.

Acknowledgments. We thank the anonymous referees for valuable suggestions that helped us to improve the presentation.

REFERENCES

- [1] P. K. AGARWAL, *On stabbing lines for convex polyhedra in 3D*, *Comput. Geom.*, 4 (1994), pp. 177–189.
- [2] P. K. AGARWAL, B. ARONOV, AND M. SHARIR, *Line transversals of balls and smallest enclosing cylinders in three dimensions*, *Discrete Comput. Geom.*, 21 (1999), pp. 373–388.
- [3] P. K. AGARWAL, O. SCHWARZKOPF, AND M. SHARIR, *The overlay of lower envelopes in 3-space and its applications*, *Discrete Comput. Geom.*, 15 (1996), pp. 1–13.
- [4] P. K. AGARWAL AND M. SHARIR, *Ray shooting amidst convex polyhedra and polyhedral terrains in three dimensions*, *SIAM J. Comput.*, 25 (1996), pp. 100–116.
- [5] B. ARONOV AND M. SHARIR, *Triangles in space, or building (and analyzing) castles in the air*, *Combinatorica*, 10 (1990), pp. 137–173.
- [6] B. ARONOV, M. SHARIR, AND B. TAGANSKY, *The union of convex polyhedra in three dimensions*, *SIAM J. Comput.*, 26 (1997), pp. 1670–1688.
- [7] B. ARONOV AND S. SMORODINSKY, *Geometric permutations induced by line transversals through a fixed point*, *Discrete Comput. Geom.*, 34 (2005), pp. 285–294.
- [8] D. AVIS AND R. WENGER, *Polyhedral line transversals in space*, *Discrete Comput. Geom.*, 3 (1988), pp. 257–265.
- [9] H. BRÖNNIMANN, O. DEVILLERS, V. DUJMOVIĆ, H. EVERETT, M. GLISSE, X. GOAOC, S. LAZARD, H.-S. NA, AND S. WHITESIDES, *Lines and free line segments tangent to arbitrary three-dimensional convex polyhedra*, *SIAM J. Comput.*, 37 (2007), pp. 522–551.
- [10] H. BRÖNNIMANN, H. EVERETT, S. LAZARD, F. SOTTILE, AND S. WHITESIDES, *Transversals to line segments in \mathbb{R}^3* , *Discrete Comput. Geom.*, 35 (2005), pp. 381–390.
- [11] B. CHAZELLE, H. EDELSBRUNNER, L. GUIBAS, M. SHARIR, AND J. STOLFI, *Lines in space: Combinatorics and algorithms*, *Algorithmica*, 15 (1996), pp. 428–447.
- [12] K. CLARKSON AND P. SHOR, *Applications of random sampling in computational geometry, II*, *Discrete Comput. Geom.*, 4 (1989), pp. 387–421.
- [13] R. COLE AND M. SHARIR, *Visibility problems for polyhedral terrains*, *J. Symbolic Comput.*, 7 (1989), pp. 11–30.

¹²Each triple of patches of three different surfaces define at most a constant number of sandwich vertices. So if n_i , $1 \leq i \leq k$, is the number of patches of surface i , then the total complexity of the sandwich region (or of the entire arrangement, for that matter) is $O(\sum_{i \neq j \neq \ell} (n_i + n_j + n_\ell)) = O(nk^2)$.

- [14] D. DOBKIN AND D. SOUVAINE, *Computational geometry in a curved world*, *Algorithmica*, 5 (1989), pp. 421–457.
- [15] 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), pp. 35–42.
- [16] A. EFRAT, L. J. GUIBAS, O. A. HALL-HOLT, AND L. ZHANG, *On incremental rendering of silhouette maps of a polyhedral scene*, *Comput. Geom.*, 38 (2007), pp. 129–138.
- [17] E. EZRA AND M. SHARIR, *A single cell in an arrangement of convex polyhedra in \mathbb{R}^3* , *Discrete Comput. Geom.*, 37 (2007), pp. 21–41.
- [18] 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.
- [19] J. HERSHBERGER, *Finding the upper envelope of n line segments in $O(n \log n)$ time*, *Inform. Process. Lett.*, 33 (1989), pp. 169–174.
- [20] H. KAPLAN, N. RUBIN, AND M. SHARIR, *Linear data structures for fast ray shooting amidst convex polyhedra*, *Algorithmica*, 55 (2009), pp. 283–310.
- [21] M. KATCHALSKI, T. LEWIS, AND J. ZAKS, *Geometric permutations for convex sets*, *Discrete Math.*, 54 (1985), pp. 271–284.
- [22] M. KATCHALSKI, T. LEWIS, AND A. LIU, *The different ways of stabbing disjoint convex sets*, *Discrete Comput. Geom.*, 7 (1992), pp. 197–206.
- [23] 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 (2001), pp. 543–548.
- [24] V. KOLTUN AND M. SHARIR, *The partition technique for overlays of envelopes*, *SIAM J. Comput.*, 32 (2003), pp. 841–863.
- [25] M. MCKENNA AND J. O’ROURKE, *Arrangements of lines in 3-space: A data structure with applications*, in *Proceedings of the 4th ACM Symposium on Computational Geometry*, 1988, pp. 371–380.
- [26] M. PELLEGRINI, *Lower bounds for line stabbing in 3-space*, *Comput. Geom.*, 3 (1993), pp. 53–58.
- [27] M. PELLEGRINI AND P. W. SHOR, *Finding stabbing lines in 3-space*, *Discrete Comput. Geom.*, 8 (1992), pp. 191–208.
- [28] M. PELLEGRINI, *Ray shooting and lines in space*, in *Handbook of Discrete and Computational Geometry*, 2nd ed., J. E. Goodman and J. O’Rourke, eds., Chapman and Hall/CRC Press, London, 2004, pp. 839–856.
- [29] M. SHARIR, *Almost tight upper bounds for lower envelopes in higher dimensions*, *Discrete Comput. Geom.*, 12 (1994), pp. 327–345.
- [30] M. SHARIR AND P. K. AGARWAL, *Davenport-Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, New York, 1995.
- [31] 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 (2000), pp. 247–259.
- [32] D. SOMMERVILLE, *Analytical Geometry in Three Dimensions*, Cambridge University Press, Cambridge, 1951.
- [33] R. WENGER, *Upper bounds on geometric permutations for convex sets*, *Discrete Comput. Geom.*, 5 (1990), pp. 27–33.
- [34] R. WENGER, *Helly-type theorems and geometric transversals*, in *Handbook of Discrete and Computational Geometry*, 2nd ed., J. E. Goodman and J. O’Rourke, eds., Chapman & Hall/CRC Press, London, 2004, pp. 73–96.