

Line Transversals of Convex Polyhedra in \mathbb{R}^{3*}

Haim Kaplan[†]

Natan Rubin[‡]

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 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 .

1 Introduction

Line transversals—a brief background. In this paper we study the combinatorial complexity of the set of line transversals of a collection \mathcal{P} of k convex polyhedra in \mathbb{R}^3 with a total of n facets. 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 [15, 26, 31].

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 $\mathcal{T}(\mathcal{P})$. The *combinatorial complexity* of $\mathcal{T}(\mathcal{P})$ is defined as the total number of topological faces, of all dimensions, on the boundary of $\mathcal{T}(\mathcal{P})$. If the objects in \mathcal{P} are pairwise disjoint, a line transversal meets them in a well-defined order, called a *geometric permutation*.

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, 22] for a description of this reduction, a special case of which is described later in this paper. When \mathcal{P} is a set of convex polyhedra, the case considered in this paper, 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} , 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} , 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., [13]), but can be $\Omega(n)$ otherwise, where n is the total description complexity of the objects of \mathcal{T} , with a matching slightly super-linear upper bound (in n). (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 3-dimensional

*Work by Haim Kaplan and Natan Rubin has been partially supported by Grant 975/06 from the Israel Science Fund. and by grant 2006204 from the U.S.-Israeli Binational Science Foundation. Work by Micha Sharir was partially supported by NSF Grant CCF-05-14079, by a grant from the U.S.-Israeli Binational Science Foundation, by grant 155/05 from the Israel Science Fund, Israeli Academy of Sciences, by a grant from the AFIRST joint French-Israeli program, and by the Hermann Minkowski-MINERVA Center for Geometry at Tel Aviv University.

[†]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel. E-mail: haimk@post.tau.ac.il

[‡]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel. E-mail: rubinnat@post.tau.ac.il

[§]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel, and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA. E-mail: michas@post.tau.ac.il

case, the complexity of $\mathcal{T}(\mathcal{P})$ depends on n , even if the elements of \mathcal{P} are pairwise disjoint. The first algorithms for computing $\mathcal{T}(\mathcal{P})$, where \mathcal{P} is a set of convex polyhedra, with a total of n facets, run in time about $O(n^4)$ [7, 23]. Pellegrini and Shor [25] establish an upper bound of $O(n^{3+\varepsilon})$ on the complexity of $\mathcal{T}(\mathcal{P})$. They also 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 semi-algebraic 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 [22] on the complexity of sandwich regions of trivariate functions.² In \mathbb{R}^3 there are almost matching lower-bound constructions, showing that the complexity of $\mathcal{T}(\mathcal{P})$ can be $\Omega(n^3)$; see [2, 24]. Those lower bounds are constructed using collections of n triangles in \mathbb{R}^3 . 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 [28]. Narrowing this gap, even for restricted families of objects, is an intriguing open problem, already for $d = 3$.

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 [6] proved 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. [9] 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 $\Theta(k)$ connected components.

Another line of research initiated by Katchalski et al. [19] studies the number $g_d(k)$ of geometric permutations of a set \mathcal{P} of pairwise disjoint convex objects. The known bounds on $g_d(k)$ in the general case are: $g_2(k) = 2k - 2$; see [13], $g_d(k) = O(k^{2d-2})$, $d \geq 3$; see [30], and $g_d(k) = \Omega(k^{d-1})$, $d \geq 3$; see [20, 28]. Improved bounds are known in several special cases such as pairwise disjoint balls [28], and families of fat objects [21].

One of the most relevant predecessors of this paper is a recent paper of Brönnimann et al. [8], who prove that the entire arrangement $\mathcal{A}^*(\mathcal{P})$ of surfaces describing upper and lower tangencies to the individual polyhedra of \mathcal{P} has complexity $O(n^2 k^2)$, and this bound is tight in the worst case. Each cell in $\mathcal{A}^*(\mathcal{P})$

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. [14] 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 [14] for more details.

A closely related problem is *ray shooting* amid such a collection of polyhedra. A number of studies of the problem obtain performance bounds that depend on both k and n ; see [4, 17]. The latter study [17], 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 time 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 $O(n^2 k^{1+\varepsilon})$, 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 only consider lines which pass through a fixed line ℓ_0 , and 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 3-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 $\mathcal{T}_{\ell_0}(\mathcal{P})$ can be measured in terms of the number of its *vertices*, each representing an extremal stabbing line, which passes through (the polyhedron edge contained in) ℓ_0 . 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

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

²It is indeed stated in [22] as a corollary.

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, 22], 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 $O((nk + k^3)\beta_4(k))$, where $\beta_4(k) = 2^{\alpha(k)}$ is an extremely slowly growing function³, expressed in terms of the inverse Ackermann function $\alpha(k)$. 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 the number of geometric permutations of \mathcal{P} induced by lines in \mathcal{L} is $O(\min\{k^3, nk^{1+\varepsilon}\})$. A naive bound on this number is $O(k^4)$ (which is Wenger’s general bound mentioned above [30]). The second term in this bound follows from the fact that the number of geometric permutations is always upper bounded by the complexity of $\mathcal{T}(P)$. The first term is interesting because it depends only on k (and improves Wenger’s bound). Still, the only known lower bound

on this quantity is $\Omega(k^2)$, even for the restricted case of lines in \mathcal{L} , 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 5 for details.

Due to lack of space, we omit proofs of some lemmas and theorems. They can be found in the full version [18].

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. We further assume that the polyhedra of \mathcal{P} and ℓ_0 are in general position,⁴ and handle degenerate configurations in the full version; see [18]. Without loss of generality, we take ℓ_0 , for the time being, to be the z -axis.

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 (left). The intersection of Π_θ with a polyhedron P is the polygon $P(\theta) = P \cap \Pi_\theta$ (if it is not empty).




Figure 1: (Left) 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 . (Right) 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 .

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).

³The reason for the index 4 is that $\beta_4(k) = \Theta\left(\frac{\lambda_4(k)}{k}\right)$, where $\lambda_4(k)$ is the maximum length of Davenport-Schinzel sequences of order 4 on k symbols; see [27] and below.

⁴In particular, we assume that at most three edges, or a vertex and an edge of the given polyhedra, admit a common transversal through ℓ_0 , and that no edge of any polyhedron is coplanar with ℓ_0 or with a facet of another polyhedron.

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 re-parametrization (e.g., replacing θ with $\tan \frac{\theta}{2}$, and φ with $\cot \varphi$), each surface σ_P^- (resp., σ_P^+) consists of monotone semi-algebraic 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.

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 [8, 25] (see also the introduction), we define an *extremal line* to be a line ℓ tangent to some polyhedra of \mathcal{P} at a set A of respective vertices and 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, and vice versa; see, e.g., [25]. The following theorem was proven by Brönnimann et al. [8].

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. This implies that there are $O(n_P + n_Q)$ pairs of edges, one of P and one of Q admitting common tangent lines to P, Q through them which belong to \mathcal{L} . In particular, the combinatorial complexity of $\mathcal{T}_{\ell_0}(\{P, Q\})$ is $O(n_P + n_Q)$, and it can be computed in $O((n_P + n_Q) \log(n_P + n_Q))$ time.*

We say that a vertex or an edge ξ of $\mathcal{T}_{\ell_0}(\mathcal{P})$ 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 $\mathcal{T}_{\ell_0}(\mathcal{P}')$. Assuming general position of $\mathcal{P} \cup \{\ell_0\}$, each vertex v of $\mathcal{T}_{\ell_0}(\mathcal{P})$ 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 $\mathcal{T}_{\ell_0}(\mathcal{P})$, each of which is the locus of lines tangent to P and Q at two specific boundary edges. 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

can be charged to a vertex of $\mathcal{T}_{\ell_0}(\mathcal{P})$, so that no vertex is charged more than $O(1)$ times. Therefore, it is sufficient to bound the number of vertices of $\mathcal{T}_{\ell_0}(\mathcal{P})$ 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.

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 semi-algebraic 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 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 [10, 29] 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 $P, Q \in \mathcal{P}$. Lemma 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, 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; the proof is omitted in this version.

LEMMA 2.1. *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 $O(n\beta_4(k)) = O(n \cdot 2^{\alpha(k)})$.*

Informally, to prove Lemma 2.1, we represent the stabbing region in \mathcal{L}_h as a sandwich region of $2k$ univariate functions, whose graphs are composed of a total of $O(n)$ algebraic arcs.

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 and containing 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 [30].

LEMMA 2.2. *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 and Construction of $\mathcal{T}_{\ell_0}(\mathcal{P})$

We first establish Theorem 3.1 which 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 3-dimensional space \mathcal{L} , to a collection of 2-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 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 number of extremal stabbing lines to \mathcal{P} in the set \mathcal{L} of lines passing through ℓ_0 is $O(nk^{1+\varepsilon})$, for any $\varepsilon > 0$. One can construct collections of k convex polyhedra in general position (together with ℓ_0), with a total of n facets, for arbitrarily large values of k and n , for which the complexity of the stabbing region is $\Omega(nk)$.*

Proof. For each polyhedron $P \in \mathcal{P}$, we separate each connected component of $\ell_0 \setminus P$ from P by a plane (there are at most two such components, and therefore at most two corresponding planes; if $\ell_0 \cap P = \emptyset$, there is a single separating plane which is parallel to ℓ_0). See Figure 1 (right) 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 2.2 imply that for each cell D of $\mathcal{A}(C_H)$, 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 properties. 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 $P \in \mathcal{P}^+$.

We now fix an edge e_0 on the boundary of some polyhedron $P_0 \in \mathcal{P}$, and denote by $\mathcal{L}[e_0]$ the set of *oriented* lines which pass through ℓ_0 and are tangent to P_0 at (the 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 the number of such “degenerate” 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 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, we only consider lines in $\mathcal{L}[e_0]$ whose orientation (θ, φ) satisfies $0 < \theta < \theta_0$. See Figure 2 for an illustration.




Figure 2: Representing lines in $\mathcal{L}[e_0]$: A side view (above), and the cross-section within the plane $\tilde{\Pi}_\theta$ (below).

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 co-planar 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 , which we ignore here.

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 region⁵

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

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 2 (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 $\ell \in W$ 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 $\gamma_P^-(\theta)$ ($\gamma_P^+(\theta)$) 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 ;⁶ see Figure 3 (right). Although the definition is fairly straightforward, it provides the central tool for expressing the stabbing region in $\mathcal{L}[e_0, D]$ as a *sandwich region*; see Lemma 3.2.




Figure 3: The interval $(\gamma_P^-(\theta), \gamma_P^+(\theta))$ when $P \in \mathcal{P}^+$ (left), and $P \in \mathcal{P}^-$ (right).

It follows that, for $\theta \in (0, \theta_0)$, the line ℓ having representation $(\theta, \gamma_P^-(\theta))$ ($(\theta, \gamma_P^+(\theta))$) is one of: (a) a

⁵The 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 $\mathcal{L}[e_0]$.

⁶Namely, $\ell_0(\theta)$ is contained in P (and thus $\tilde{P}(\theta)$), if $P \in \mathcal{P}_0(D, I)$; and P (and, therefore, $\tilde{P}(\theta) \cap \ell$) lies before (resp., after) $\ell_0(\theta)$ along ℓ , if $P \in \mathcal{P}^-(D, I)$ (resp., $P \in \mathcal{P}^+(D, I)$).

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 $e_0 = \tilde{P}_0(\theta)$. With an appropriate, standard re-parametrization of θ, φ , as above, the graph of each of the functions γ_P^-, γ_P^+ is piecewise algebraic, composed of algebraic arcs of constant maximum degree. Each arc either represents lines in $\mathcal{L}[e_0]$ that are tangent to P at some fixed edge (each such arc is a trace, on \mathbb{S}^2 , of a *regulus* of lines, defined in Section 2), or it represents lines tangent to P_0 at one of the endpoints of e_0 . Since by assumption all lines in $\mathcal{L}[e_0]$ have $\varphi \in (\varphi^-(\theta), \varphi^+(\theta))$, the arcs of the latter type are redundant and, therefore, can be safely ignored in the subsequent analysis.

Formally, we have the following lemma, whose proof is omitted.

LEMMA 3.1. *Let P be a polyhedron of $\mathcal{P} \setminus \{P_0\}$. Then each algebraic subarc ζ of γ_P^\pm , defined by some fixed edge e of P , is (a trace of) 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 .*

Now the discussion following Theorem 2.1 implies that the overall number of distinct algebraic arcs that comprise the graphs of the functions γ_P^\pm , over all possible edges e_0 , polyhedra P , and cells D , is $O(nk)$ (where a single arc may traverse several cells $D \in \mathcal{D}$). 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 parameterizations) 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 .

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 (local) spherical coordinates (θ, φ) . Then ℓ stabs P if and only if*

$$(3.2) \quad \gamma_P^-(\theta) \leq \varphi \leq \gamma_P^+(\theta).$$

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

$$(3.3) \quad \max_{P \in \mathcal{P}} \gamma_P^-(\theta) \leq \varphi \leq \min_{P \in \mathcal{P}} \gamma_P^+(\theta).$$

COROLLARY 3.1. *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) corresponds to a vertex on the boundary of the sandwich region given by (3.3).*

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 (3.3). This region is formed by the graphs

of $O(k)$ functions, consisting of some number, $\Gamma[e_0, D]$, of algebraic arcs. A pair of such arcs, corresponding to a pair of distinct 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 [29]. It follows from [27, Theorem 1.4] 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. (Here we face the difficulty that arcs may appear in several cells 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 (3.3), 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 4 for an illustration.




Figure 4: The sandwich region in $\mathcal{L}[e_0, D]$, as given by (3.3), and D , 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 $\mathcal{R}[e_0, D]$.

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

$$(3.4) \quad N[e_0, D] = O(R[e_0, D]\beta_4(k)).$$

Put $N = \sum_{e_0, D} N[e_0, D]$ and $R = \sum_{e_0, D} R[e_0, D]$, 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 (3.4), we have $N(n, k) = O(R(n, k)\beta_4(k))$. Now the last step of the analysis is deriving a recurrence

formula for $N(n, k)$ which we obtain by proving that

$$(3.5) \quad R(n, k) = O\left(nk + tnk\beta_4(k) + \frac{1}{t} \cdot t^3 N\left(\frac{n}{t}, \frac{k}{t}\right)\right)$$

for any integer threshold $t > 0$. The solution of the resulting recurrence is easily seen to imply the desired bound. We derive (3.5) using a standard slide-and-charge scheme; e.g., see [16, 27]. Specifically, we say that the level of a line $\ell \in \mathcal{L}$ is j if ℓ stabs all but j polyhedra of \mathcal{P} . Now, fix a polyhedron P_0 and an edge e_0 of ∂P_0 . As argued above, there is a total of $O(nk)$ arcs counted by $R[e_0, D]$, for at least one cell $D \in \mathcal{D}$, over all choices of P_0 and e_0 . Let ζ be an arc counted in $R[e_0, D]$, for more than one cell $D \in \mathcal{D}$. Then each such cell D contains a connected portion $\tilde{\zeta}$ of $\zeta \cap D$, 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 $\mathcal{L}[e_0, D]$. We distinguish between two possible cases.

(i) The depth of (the line represented by) a is less than t . In this case we charge the appearance of ζ in D to a , which represents an extremal line in \mathcal{L}_h (see Section 2), at level $\leq t$, where h is one of the planes such that the great circle C_h appears on the boundary of D . To bound the number of extremal lines charged in this manner, we combine Lemma 2.1 with the general randomized technique of Clarkson and Shor [11] in \mathcal{L}_h , and get that the number of extremal lines at level $\leq t$ in \mathcal{L}_h is bounded by $O(nt\beta_4(k))$, and the number of extremal lines at level $\leq t$ in all $O(k)$ planes \mathcal{L}_h , $h \in H$ is $O(nkt\beta_4(k))$.

(ii) The depth of (the line represented by) a is at least t . In this case we walk from b to a along $\tilde{\zeta}$, and collect at least t extremal lines at depth $\leq t$ contained in $\mathcal{L}[e_0, D]$. We charge the appearance of ζ in D to those lines. To bound the number of extremal lines charged in this manner, we also apply the Clarkson and Shor technique (this time by sampling k/t out of the k polyhedra in the entire space) to deduce that the number of extremal lines at level $\leq t$ is $O(t^3 N(\frac{n}{t}, \frac{k}{t}))$.

In both cases, (i) and (ii), each extremal line is charged at most a constant number of times, over all choices of e_0 and D , so we eventually obtain (3.5).

Lower bound. A lower bound construction for the complexity of $\mathcal{T}_{\ell_0}(\mathcal{P})$ is depicted in Figure 5. We use k pairs of polyhedra, each of which is a thin and long vertical plate, whose vertical edges and facets are all parallel to ℓ_0 . The polyhedra in each pair are parallel to each other, and 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 $\Theta(n)$




Figure 5: A lower-bound construction, viewed from above. (The actual gaps between the pairs of parallel “plates” around ℓ_0 , relative to the width of the plates, are much smaller than depicted.)

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 5, and [10, 12] 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 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.)

This completes the proof of Theorem 3.1.

THEOREM 3.2. (Algorithm for constructing $\mathcal{T}_{\ell_0}(\mathcal{P})$)
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 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. Without any loss of generality, we assume that ℓ_0 is the z -axis. 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 .

The algorithm for constructing $\mathcal{T}_{\ell_0}(\mathcal{P})$ proceeds as follows. First, we choose a fixed random permutation π of the polyhedra of \mathcal{P} . Next, we fix a polyhedron $P \in \mathcal{P}$, and a boundary edge e of P . 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 2-dimensional patches constitutes the boundary of $\mathcal{T}_{\ell_0}(\mathcal{P})$. We follow a standard randomized divide-and-conquer approach; see [27]. 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 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., [27]. 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 . By a standard reduction, it suffices to bound the expected contribution to N of vertices defined by triples of polyhedra; see Section 2 for details. 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 $Q, R \in \mathcal{P} \setminus \{P\}$. We say that v has depth t if its corresponding line ℓ_v has depth t . As is easily verified the probability that a vertex v at level t is constructed by the algorithm, when processing σ_e , is at most $\frac{2}{t+2}$.

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

$$(3.6) \quad N \leq 6 \sum_{0 \leq t \leq k-3} \frac{(N_{\leq t} - N_{\leq t-1})}{t+2} + O(nk).$$

Rearranging the sum, we get

$$(3.7) \quad N \leq 6 \sum_{0 \leq t \leq k-4} N_{\leq t} \left(\frac{1}{t+2} - \frac{1}{t+3} \right) + \frac{6}{k-1} N_{\leq k-3} + O(nk) = 6 \sum_{0 \leq t \leq k-4} \frac{N_{\leq t}}{(t+2)(t+3)} + \frac{6}{k-1} N_{\leq k-3} + O(nk).$$

By Theorem 3.1, for any $\varepsilon > 0$, $N_{\leq 0} = O(nk^{1+\varepsilon})$. Hence, by the Clarkson-Shor probabilistic

argument [11], $N_{\leq t} = O(tnk^{1+\varepsilon})$, for all $1 \leq t \leq k-3$; see [5] for a similar argument. Plugging this into (3.7), we obtain that $N = O(nk^{1+\varepsilon})$, for any $\varepsilon > 0$.

4 The Complexity and Construction of $\mathcal{T}(\mathcal{P})$

THEOREM 4.1. (The general $\mathcal{T}(\mathcal{P})$) *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 $O(n^2k^{1+\varepsilon})$, for any $\varepsilon > 0$.*

Proof. For the sake of simplicity, we only bound the number of extremal stabbing lines in $\mathcal{T}(\mathcal{P})$, which correspond to boundary vertices of the stabbing region, within an appropriate parametric space of lines; see [25] and Section 2 for a standard justification.

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 is also an extremal stabbing line ℓ to $\mathcal{P} \setminus \{P_0\}$ in \mathcal{L} , unless it passes through an endpoint of e_0 , or it overlaps a facet of P_0 incident to e_0 . Now using Theorem 3.1, and summing over $O(n)$ possible choices of e_0 , we immediately derive an $O(n^2k^{1+\varepsilon})$ upper bound, for any $\varepsilon > 0$, on the number of extremal stabbing lines that are tangent to at least one polyhedron at the relative interior of one of its edges, and not overlapping any of its facets.

To complete our proof, we next show an upper bound of $O(n^2k)$ on the number of remaining extremal stabbing lines, each of which, assuming general position of $\mathcal{P} \cup \{\ell_0\}$, is defined by two or fewer polyhedra of \mathcal{P} . For the sake of simplicity, we derive the above bound only for extremal stabbing lines defined by pairs of polyhedra. Fix any pair of distinct polyhedra $P, Q \in \mathcal{P}$, with n_P and n_Q facets, respectively. As shown by Brönnimann et al. [8], the number of such extremal stabbing lines defined by P and Q is $O((n_P+n_Q)^2)$. The desired bound of $O(n^2k)$ is now obtained by summing over all choices of P and Q .

Lower bounds. 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 [8]. A different lower bound of $\Omega(nk^2)$ is described in the full version. Closing the gap between the resulting lower bound $\Omega(n^2+nk^2)$ and the upper bound in Theorem 4.1 remains a major challenging open problem.

The following theorem follows from Theorem 3.2 and Theorem 4.1.

THEOREM 4.2. (Algorithm for constructing $\mathcal{T}(\mathcal{P})$) *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 $O(n^2k^{1+\varepsilon} \log n)$ randomized expected time, for any $\varepsilon > 0$.

5 Extensions

In the full version [18] we prove slightly improved upper bounds on the complexity of the stabbing region $\mathcal{T}_{\ell_0}(\mathcal{P})$, in the following two cases: (i) all polyhedra of \mathcal{P} are disjoint from ℓ_0 , and (ii) 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 $\mathcal{T}_{\ell_0}(\mathcal{P})$. The first result is established in Theorem 5.1 and the second is a corollary of Theorem 5.2.

THEOREM 5.1. *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 $\mathcal{T}_{\ell_0}(\mathcal{P})$ in $O((nk+k^3)(\log n + \alpha(k) \log k))$ deterministic time.*

THEOREM 5.2. *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.*

Using our techniques we also derive the following result about geometric permutations.

THEOREM 5.3. *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 $O(k^3)$.*

6 Conclusion

In this paper we obtained an improved bound 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 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 on the 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 successfully combines the slide-and-charge scheme from [16] with the method of separating planes, that was previously used in [30] to study geometric permutations; also see [6].

There are several challenging open problems for further research, including the problem of closing 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.

Acknowledgements. The authors thank the anonymous SODA referees for helpful comments.

References

- [1] P. K. Agarwal, On stabbing lines for convex polyhedra in 3D, *Comput. Geom. Theory Appl.* 4 (1994), 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(3) (1999), 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), 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), 100–116.
- [5] B. Aronov, M. Sharir and B. Tagansky, The union of convex polyhedra in three dimensions, *SIAM J. Comput.* 26 (1997), 1670–1688.
- [6] B. Aronov and S. Smorodinsky, Geometric permutations induced by line transversals through a fixed point, *Discrete Comput. Geom.* 34 (2005), 285–294.
- [7] D. Avis and R. Wenger, Polyhedral line transversals in space, *Discrete Comput. Geom.* 3 (1988), 257–265.
- [8] H. Brönnimann, O. Devillers, V. Dujmovic, 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), 522–551.
- [9] H. Brönnimann, H. Everett, S. Lazard, F. Sottile and S. Whitesides, Transversals to line segments in \mathbb{R}^3 , *Discrete Comput. Geom.* 35 (2005), 381–390.
- [10] B. Chazelle, H. Edelsbrunner, L. Guibas, M. Sharir and J. Stolfi, Lines in space: Combinatorics and algorithms, *Algorithmica* 15 (1996), 428–447.
- [11] K. Clarkson and P. Shor, Applications of random sampling in computational geometry, II, *Discrete Comput. Geom.* 4 (1989), 387–421.
- [12] R. Cole and M. Sharir, Visibility problems for polyhedral terrains, *J. Symb. Comput.* 7 (1989), 11–30.
- [13] H. Edelsbrunner and M. Sharir, The maximum number of ways to stab n convex nonintersecting sets in the plane is $2n - 2$, *Discrete Comput. Geom.* 5 (1990), 35–42.
- [14] A. Efrat, L. J. Guibas, O. A. Hall-Holt, and L. Zhang, On incremental rendering of silhouette maps of a polyhedral scene, *Comput. Geom. Theory Appl.* 38 (2007), 129–138.
- [15] 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.
- [16] D. Halperin and M. Sharir, New bounds for lower envelopes in three dimensions, with applications to visibility in terrains, *Discrete Comput. Geom.* 12 (1994), 313–326.
- [17] H. Kaplan, N. Rubinfeld and M. Sharir, Linear data structures for fast ray shooting amidst convex polyhedra, *Algorithmica*, to appear. Also in *Proc. 15th Europ. Sympos. Algo.*, 2007, 287–298.
- [18] H. Kaplan, N. Rubinfeld and M. Sharir, Line transversals of convex polyhedra in \mathbb{R}^3 , <http://www.cs.tau.ac.il/~rubinnat/TransPoly.pdf>. Also see [arXiv:0807.1221](https://arxiv.org/abs/0807.1221).
- [19] M. Katchalski, T. Lewis and J. Zaks, Geometric permutations for convex sets, *Discrete Math.* 54 (1985), 271–284.
- [20] M. Katchalski, T. Lewis and A. Liu, The different ways of stabbing disjoint convex sets, *Discrete Comput. Geom.* 7 (1992), 197–206.
- [21] M. J. Katz and K. R. Varadarajan, A tight bound on the number of geometric permutations of convex fat objects in \mathbb{R}^d , *Discrete Comput. Geom.* 26(4) (2001), 543–548.
- [22] V. Koltun and M. Sharir, The partition technique for overlays of envelopes, *SIAM J. Comput.* 32(4) (2003), 841–863.
- [23] M. McKenna and J. O’Rourke, Arrangements of lines in 3-space: a data structure with applications. *Proc. 3rd ACM Sympos. Comput. Geom.*, 1988, 371–380.
- [24] M. Pellegrini, Lower bounds for line stabbing in 3-space, *Comput. Geom. Theory Appl.* 3 (1993), 53–58.
- [25] M. Pellegrini, and P. W. Shor, Finding stabbing lines in 3-space, *Discrete Comput. Geom.* 8 (1992), 191–208.
- [26] M. Pellegrini, Ray shooting and lines in space, in *Handbook of Discrete and Computational Geometry*, 2nd Edition, (J. E. Goodman and Joseph O’Rourke, Eds.), Chapman and Hall/CRC Press, 2004, pp. 839–856.
- [27] M. Sharir and P. K. Agarwal, *Davenport-Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, New York, 1995.
- [28] S. Smorodinsky, J. S. B. Mitchell and M. Sharir, Sharp bounds on geometric permutations of pairwise disjoint balls in \mathbb{R}^d , *Discrete Comput. Geom.* 23(2) (2000), 247–259.
- [29] D. Sommerville, *Analytical Geometry in Three Dimensions*, Cambridge University Press, Cambridge, 1951.
- [30] R. Wenger, Upper bounds on geometric permutations for convex sets, *Discrete Comput. Geom.* 5 (1990), 27–33.
- [31] R. Wenger, Helly-type theorems and geometric transversals, in *Handbook of Discrete and Computational Geometry*, 2nd Edition (J.E. Goodman and J. O’Rourke, Eds.), Chapman & Hall/CRC Press, 2004, pp. 73–96.