

Lines and free line segments Tangent to Arbitrary Three-dimensional Convex Polyhedra

Hervé Brönnimann, Olivier Devillers, Vida Dujmovic, Hazel Everett, Marc Glisse, Xavier Goaoc, Sylvain Lazard, Hyeon-Suk Na, Sue Whitesides

► To cite this version:

Hervé Brönnimann, Olivier Devillers, Vida Dujmovic, Hazel Everett, Marc Glisse, et al.. Lines and free line segments Tangent to Arbitrary Three-dimensional Convex Polyhedra. SIAM Journal on Computing, Society for Industrial and Applied Mathematics, 2007, 37 (2), pp.522-551. <10.1137/S0097539705447116>. <inria-00103916>

HAL Id: inria-00103916

<https://hal.inria.fr/inria-00103916>

Submitted on 19 Nov 2007

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

LINES AND FREE LINE SEGMENTS TANGENT TO ARBITRARY THREE-DIMENSIONAL CONVEX POLYHEDRA *

HERVÉ BRÖNNIMANN[†], OLIVIER DEVILLERS[‡], VIDA DUJMOVIĆ[§], HAZEL EVERETT[¶], MARC
GLISSE[¶], XAVIER GOAOC[¶], SYLVAIN LAZARD[¶], HYEON-SUK NA^{||}, AND SUE WHITESIDES^{**}

Abstract. Motivated by visibility problems in three dimensions, we investigate the complexity and construction of the set of tangent lines in a scene of three-dimensional polyhedra. We prove that the set of lines tangent to four possibly intersecting convex polyhedra in \mathbb{R}^3 with a total of n edges consists of $\Theta(n^2)$ connected components in the worst case. In the generic case, each connected component is a single line, but our result still holds for arbitrarily degenerate scenes. More generally, we show that a set of k possibly intersecting convex polyhedra with a total of n edges admits, in the worst case, $\Theta(n^2k^2)$ connected components of maximal free line segments tangent to at least four polytopes. Furthermore, these bounds also hold for possibly occluded lines rather than maximal free line segments.

Finally, we present a $O(n^2k^2 \log n)$ time and $O(nk^2)$ space algorithm that, given a scene of k possibly intersecting convex polyhedra, computes all the *minimal* free line segments that are tangent to any four of the polytopes and are isolated transversals to the set of edges they intersect; in particular, we compute at least one line segment per connected component of tangent lines.

Key words. Computational geometry, 3D visibility, visibility complex, visual events.

AMS subject classifications. 65D18, 68U05.

1. Introduction. Computing visibility relations in a 3D environment is a problem central to computer graphics and engineering tasks such as radio propagation simulation and fast prototyping. Examples of visibility computations include determining the view from a given point, and computing the umbra and penumbra cast by a light source. In many applications, visibility computations are well-known to account for a significant portion of the total computation cost. Consequently a large body of research is devoted to speeding up visibility computations through the use of data structures (see [14] for a survey).

One such structure, the visibility complex [16, 23], encodes visibility relations by partitioning the set of maximal free line segments. The size of this partition is intimately related to the number of maximal free line segments tangent to four objects in the scene; for a scene of n triangles in \mathbb{R}^3 , the complex can have size $\Theta(n^4)$ in the worst case [16], even when the triangles form a terrain (see [11] or Figure 1.1). The complex is thus potentially enormous, which has hindered its application in practice. However, there is evidence, both theoretical and practical, that this estimation is pessimistic. The lower bound examples, which are carefully designed to exhibit the worst-case behavior, are unrealistic in practice. For realistic scenes, Durand *et al.* [15] observe a quadratic growth rate, albeit for rather small scenes. For random scenes, Devillers *et al.* [12] prove that the expected size of the visibility complex is much smaller; for uniformly distributed unit balls the expected size is linear and for polygons or polyhedra of bounded aspect ratio and similar size it is at most quadratic. Also, in 2D, while the worst-case complexity of the visibility complex is quadratic, experimental re-

*This research was initiated at the McGill-INRIA Workshop on Computational Geometry in Computer Graphics, February 9-15, 2002, co-organized by H. Everett, S. Lazard, and S. Whitesides, and held at the Bellairs Research Institute of McGill University. This research was supported by the McGill-ISA collaborative INRIA project, NSERC, FCAR and NSF. A preliminary version of this paper appeared in Proceedings of the 20th Annual ACM Symposium on Computational Geometry, New-York, USA, 2004, pp. 46–55.

[†]Polytechnic University, Brooklyn, NY 11201, USA. hbr@poly.edu.

[‡]INRIA Sophia-Antipolis, Sophia-Antipolis, France. olivier.Devillers@inria.fr.

[§]School of Computer Science, Carleton University, Ottawa, Canada. vida@scs.carleton.ca.

[¶]LORIA - INRIA Lorraine, University Nancy 2, Nancy, France. Firstname.Name@loria.fr.

^{||}School of Computing, Soongsil University, Seoul, South Korea. hsnaa@computing.ssu.ac.kr.

^{**}School of Computer Science, McGill University, Montréal, Canada. sue@cs.mcgill.ca.




FIG. 1.1. A terrain of size n with $\Omega(n^4)$ maximal free line segments tangent in four points.

sults strongly suggest that the size of the visibility complex of a scene consisting of scattered triangles is linear [10].

While these results are encouraging, most scenes are not random. In fact, most scenes have a lot of structure which we can exploit; a scene is typically represented by many triangles which form a much smaller number of convex patches. In particular, if a scene consists of k disjoint convex polyhedra with a total of n edges, then under a strong general position assumption, the number of maximal free line segments tangent to four of the polyhedra is at most $O(n^2k^2)$; this follows directly from the bound proved in [17] on the number of combinatorial changes of the silhouette map viewed from a point moving along a straight line, and was also later proved in [8]. We present in this paper a generalization of these results. After preliminary definitions, we give a detailed account of our results and then present related previous work.

Preliminary definitions. We consider a scene that consists of a finite number of polytopes, not necessarily disjoint, not necessarily fully dimensional, and in arbitrary position. The definitions below are standard, yet carefully phrased in a way that remains valid in those situations.

A *polytope* is the convex hull of a point set. A plane is *tangent* to a polytope if it intersects the polytope and bounds a closed half-space that contains the polytope. A face, an edge, or a vertex of a polytope in \mathbb{R}^3 is the 2, 1 or 0-dimensional intersection of the polytope with a tangent plane. Note that, with this usual definition of polytopes, edges and faces are closed and they are not subdivided in any way.

A line or segment is *tangent* to a polytope (whether or not the latter is fully dimensional) if it intersects the polytope and is contained in a tangent plane. In a given plane, a line is tangent to a polygon if it intersects the polygon and bounds a closed half-plane that contains the polygon. With these definitions, given a polygon in a plane π , and a line contained in π that intersects the relative interior of this polygon, the line is tangent to the polygon when considered as a polytope in \mathbb{R}^3 , but not tangent to the polygon in the plane π .

The set of lines in \mathbb{R}^3 has a natural topological structure, namely, that of Plücker space [25]. The set of lines tangent to at least four polytopes is a subspace, whose *connected components* correspond to lines that can be continuously moved one into the other while remaining tangent to at least four polytopes.¹ A line or line segment is *free* if it is tangent to each polytope that its relative interior intersects;² otherwise it is *occluded*. A free line segment is a *maximal free line segment* if it is not properly contained into another free line segment. The space of

¹The set of polytopes to which the line is tangent might change during the motion.

²When the polytopes are fully dimensional, a segment is free if it does not intersect the interior of any of them. Our definition ensures that a segment is free also when it intersects and is coplanar with a two-dimensional polytope. The endpoints of a free segment may also lie on the boundary of a polytope.

line segments also has a natural topological structure and the *connected components* of maximal free line segments tangent to at least four among the k polytopes are defined similarly as for lines.

A *support vertex* of a line is a polytope vertex that lies on the line. A *support edge* of a line is a polytope edge that intersects the line but has no endpoint on it (a support edge intersects the line at only one point of its relative interior). A *support* of a line is one of its support vertices or support edges. The supports of a segment are defined similarly. Notice that it follows from the definition of polytopes that any line has at most two supports in any given polytope.

A line is *isolated with respect to* a set of edges and vertices if the line cannot be moved continuously while remaining a common transversal to these edges and vertices. Furthermore, we say that a set \mathcal{S} of edges and vertices *admits an isolated transversal* if these edges and vertices admit a common transversal that is isolated with respect to \mathcal{S} . Finally, a line is *isolated* if it is isolated with respect to a set of some, and hence all, of its supports.

Our results. In this paper we present two types of results, combinatorial bounds and algorithms.

Combinatorial bounds. We generalize the result of [8, 17] in two ways. First, we consider polytopes that may **intersect**. We show that among k polytopes of total complexity n , the number of lines tangent to any four of them is in the worst case either infinite or $\Theta(n^2k^2)$. The most surprising aspect of this result is that the bound (which is tight) is the same whether the polytopes intersect or not. This is in sharp contrast to the 2D case, where the number of tangents of two convex polygons is always 4 if disjoint, and could be linear in the size of the polygons if they intersect. Second, we consider polytopes in **arbitrary position**: we drop all general position assumptions. The polytopes may intersect in any way; they may overlap or coincide. They may degenerate to polygons, segments or points. While four polytopes in general position (as defined in [8]) admit a finite number of common tangents, four polytopes in arbitrary position may admit an infinite number of common tangents which can be partitioned into connected components.

Our main results are, more precisely, the following.

THEOREM 1.1. *Given k polytopes in \mathbb{R}^3 with n edges in total, there are, in the worst case, $\Theta(n^2k^2)$ connected components of maximal free line segments tangent to at least four of the polytopes. This bound also holds for connected components of possibly occluded lines tangent to at least four of the polytopes.*

These results improve the trivial bound of $O(n^4)$. Note that, when $k \neq 4$, neither of the two results stated in Theorem 1.1 implies the other since a line tangent to at least four among k polytopes may contain many, but does not necessarily contain any, maximal free line segments tangent to four polytopes.

When $k = 4$, Theorem 1.1 implies that there are $\Theta(n^2)$ connected components of lines tangent to the four polytopes, an improvement on the previously known upper bound of $O(n^3 \log n)$ which follows from the same bound on the complexity of the set of line transversals to a set of polyhedra (here four) with n edges in total [1]. Moreover, we prove a tighter bound when one of the four polytopes has few edges.

THEOREM 1.2. *Given 3 polytopes with n edges in total and one polytope with m edges, there are, in the worst case, $\Theta(mn)$ connected components of lines tangent to the four polytopes.*

We also prove the following result which is more powerful, though more technical, than Theorem 1.1. Whereas Theorem 1.1 bounds the number of connected components of tangents, Theorem 1.3 bounds the number of isolated tangents with some notion of multiplicity. For example, the line in Figure 1.2 is counted $\binom{k}{2}$ times which is the number of minimal sets




FIG. 1.2. A line tangent at a vertex of each of k polytopes.

of vertices that admit that line as an isolated transversal. Although neither theorem implies the other, we will prove in Proposition 3.4 that the upper bound of Theorem 1.1 is easily proved using Theorem 1.3.

THEOREM 1.3. *Given k polytopes in \mathbb{R}^3 with n edges in total, there are, in the worst case, $\Theta(n^2k^2)$ minimal sets of open edges and vertices, chosen from some of the polytopes, that admit a possibly occluded isolated transversal that is tangent to these polytopes.*

Algorithm. We now turn our attention to the computation of all free segments that are isolated transversals to their set of supports and tangent to the corresponding polytopes. Durand *et al.* [16] proposed an algorithm for this problem with worst-case time complexity $O((n^3 + p) \log n)$ where p is the output size; this algorithm, based on a double-sweep, has proven to be difficult to implement. Durand *et al.* also presented an algorithm with $\Theta(n^5)$ worst-case time complexity that incorporates interesting heuristics leading to reasonable performance in practice [15]. We present an algorithm that uses, in the worst case, $O(n^2k^2 \log n)$ time and $O(nk^2)$ space, is readily implementable, and uses only simple data structures. The polytopes may intersect and be in arbitrary position. A preliminary version of this algorithm was described for disjoint convex polyhedra in X. Goaoc's Ph.D. thesis [19].

THEOREM 1.4. *Given k polytopes in \mathbb{R}^3 with n edges in total, we can compute, in $O(n^2k^2 \log n)$ time and $O(nk)$ space, all the possibly occluded lines that are isolated transversals to their set of supports and tangent to the corresponding polytopes. We can also compute, in $O(n^2k^2 \log n)$ time and $O(nk^2)$ space, all the minimal free segments that are isolated transversals to their set of supports and tangent to the corresponding polytopes.*

It should be noted that our algorithm does not provide the endpoints (possibly at infinity) of the maximal free segments. Computing the endpoints of each such segment can be done by shooting rays in $O(\log^2 n)$ time per ray using $O((nk)^{2+\epsilon})$ preprocessing time and storage [3]. Such ray-shooting data structures are not, however, readily implementable. Alternatively, each ray-shooting query can be answered in $O(k \log n)$ time after $O(n \log n)$ preprocessing time and using additional $O(n)$ space by applying the Dobkin-Kirkpatrick hierarchy on each polytope [13].

To emphasize the importance of considering intersecting polytopes, observe that computer graphics scenes often contain non-convex objects. These objects, however, can be decomposed into sets of convex polyhedra. Notice that simply decomposing these objects into convex polyhedra with disjoint interiors may induce a scene of much higher complexity than a decomposition into intersecting polytopes. Moreover, the decomposition of a polyhedron into interior-disjoint polytopes may introduce new tangents which were not present in the original scene; indeed a line tangent to two polytopes along a shared face is not tangent to their union.

	Worst-case	Expected
free lines to a polyhedron	$\Theta(n^4)$ (trivial)	
free lines above a polyhedral terrain	$O(n^3 2^{c\sqrt{\log n}})$ [20, 22]	
free lines among disjoint homothetic polytopes	$\Omega(n^3)$ [4]	
free lines among unit balls	$\Omega(n^2)$ [12], $O(n^{3+\epsilon})$ [2]	$\Theta(n)$ [12]
max. free segments above a polyhedral terrain	$\Theta(n^4)$ [11]	
isolated maximal free segments among k generic disjoint convex polyhedra	$\Theta(n^2 k^2)$ [17, 8]	
max. free segments among unit balls	$\Omega(n^2)$ [12], $O(n^4)$	$\Theta(n)$ [12]

TABLE 1.1

Published bounds on the complexity of the set of free lines or maximal free line segments among objects of total complexity n . The expected complexities are given for the uniform distribution of the balls centers.

The importance of considering polytopes in arbitrary position comes from the fact that graphics scenes are full of degeneracies both in the sense that four polytopes may admit infinitely many tangents and that polytopes may share edges or faces. There may actually be more connected components of tangents when the objects are in degenerate position; this is, for instance, the case for line segments [9]. Also, we could not find a perturbation argument that guarantees the preservation of all (or at least a constant fraction of) the connected components of tangents and we do not believe that finding such a perturbation is a simple matter.

Related results. Previous results on this topic include those that bound the complexity of sets of free lines or free line segments among different sets of objects. They are summarized in Table 1.1.

Recently, Agarwal *et al.* [2] proved that the set of free lines among n unit balls has complexity $O(n^{3+\epsilon})$. Devillers *et al.* showed a simple bound of $\Omega(n^2)$ [12] for this problem, and Koltun recently sketched a bound of $\Omega(n^3)$ (personal communication, 2004).

The complexity of the set of free line segments among n balls is trivially $O(n^4)$. Devillers and Ramos showed that the set of free line segments can have complexity $\Omega(n^3)$ (personal communication 2001, see also [12]). When the balls are unit size, the $\Omega(n^2)$ lower bound for the set of free lines holds. A lower bound of $\Omega(n^4)$ that applies to either case was recently sketched by Glisse (personal communication, 2004).

We mention two results for polyhedral environments. Halperin and Sharir [20] and Pellegrini [22] proved that, in a polyhedral terrain with n edges, the set of free lines has near-cubic complexity. De Berg, Everett and Guibas [4] showed a $\Omega(n^3)$ lower bound on the complexity of the set of free lines (and thus free segments) among n disjoint homothetic convex polyhedra.

This paper is organized as follows. We prove the upper bounds of Theorems 1.1, 1.2, and 1.3 in Sections 2 and 3, and the lower bounds in Section 4. In section 5, we present our algorithm for computing free segments.

2. Main lemma. We prove in this section a lemma which is fundamental for the proofs of the upper bounds of Theorems 1.1, 1.2, and 1.3. Consider four polytopes \mathbf{P} , \mathbf{Q} , \mathbf{R} , and \mathbf{S} in \mathbb{R}^3 , with p , q , r , and $s \geq 1$ edges, respectively, and let e be an edge of \mathbf{S} .

MAIN LEMMA. *There are $O(p+q+r)$ isolated lines intersecting e and tangent to \mathbf{P} , \mathbf{Q} , \mathbf{R} and \mathbf{S} excluding those that lie in planes that contain e and are tangent to all four polytopes.*

The proof of the Main Lemma is rather complicated because it handles polytopes which may intersect as well as all the degenerate cases. To assist the reader, we first give an overview




FIG. 2.1. Plane Π_t contains edge e and intersects polytopes \mathbf{P} , \mathbf{Q} , and \mathbf{R} in polygons P_t , Q_t , and R_t .

of the proof. We then state preliminaries and definitions in Section 2.2. In Sections 2.3 and 2.4, we bound the number of so-called “generic tangent lines”. In Section 2.5, we bound the number of “non-generic tangent lines”. Finally, in Section 2.6, we pull these results together to conclude the proof of the Main Lemma.

2.1. Proof overview. The proof is inspired by a method which was, to our knowledge, first used in [6] (and later in [5, 17, 8]). We present here an overview of the proof in which we do not address most of the problems arising from degeneracies. In particular, some definitions and remarks will require more elaboration in the context of the complete proof.

We sweep the space with a plane Π_t rotating about the line containing e . The sweep plane intersects the three polytopes \mathbf{P} , \mathbf{Q} , and \mathbf{R} in three, possibly degenerate or empty, convex polygons denoted P_t , Q_t , and R_t , respectively (see Figure 2.1). During the sweep, we track the *bitangents*, that is, the lines tangent to P_t and Q_t , or to Q_t and R_t , in Π_t . As the sweep plane rotates, the three polygons deform and the bitangents move accordingly. Every time two bitangents become aligned during the sweep, the common line they form is tangent to \mathbf{P} , \mathbf{Q} , and \mathbf{R} .

In any given instance of the sweep plane Π_t , we consider the pairs of bitangents (one involving P_t and Q_t , and the other Q_t and R_t) that share a vertex of Q_t (see Figure 2.1). The isolated lines intersecting e and tangent to \mathbf{P} , \mathbf{Q} , \mathbf{R} and \mathbf{S} are isolated transversals with respect to a tuple of supports that consists of e and the supports of two such bitangents. We consider all *candidate* such tuples of supports as the sweep plane rotates.

Such a tuple induced by an instance of the sweep plane changes as the plane rotates only when a support of a bitangent changes. We define *critical planes* in such a way that the supports of the bitangents do not change as the sweep plane rotates between two consecutive critical planes. As the sweep plane rotates, the supports of a bitangent change if a support starts or ceases to be swept, or if, during its motion, the bitangent becomes tangent to one of the polygons along an edge of that polygon (see Figure 2.2). In the latter case, this means that the bitangent crosses a face or contains an edge of one of the polytopes. We thus define two types of critical planes: an instance of the sweep plane is critical if it contains a vertex of one of the polytopes, or if it contains a line that lies in the plane containing a face of one of the polytopes, and is tangent to another of the polytopes (see Figures 2.2 and 2.3). We will show that the number of critical planes is $O(p + q + r)$.

When the polytopes intersect there may exist a linear number of bitangents in an instance of the sweep plane (two intersecting convex polygons may admit a linear number of bitangents, as is the case for two regular n -gons where one is a rotation of the other about its center). Thus there can be a linear number of candidate tuples induced by any instance of the sweep plane, and the linear number of critical planes leads to a quadratic bound on the




FIG. 2.2. A bitangent to P_t and Q_t is tangent to P_t along an edge. The plane Π_t is F -critical.

total number of distinct candidate tuples. In the detailed proof of the lemma, we amortize the count of candidate tuples over all the critical planes to get a linear bound on the number of distinct candidate tuples and thus on the number of isolated lines intersecting e and tangent to \mathbf{P} , \mathbf{Q} , \mathbf{R} and \mathbf{S} ; this bound will however not hold for those isolated lines that lie in planes that contain e and are tangent to all four polytopes. Indeed, the number of such isolated tangent lines can be quadratic, in degenerate cases; for instance, four polytopes such that a plane contains edge e and a face of linear complexity from each other polytope may admit in this plane a quadratic number of such isolated tangent lines (one through each of a quadratic number of pairs of vertices).

2.2. Preliminaries and definitions. We can assume without loss of generality that \mathbf{P} , \mathbf{Q} , \mathbf{R} and \mathbf{S} have *non-empty interior*. Indeed, since the set of isolated tangent lines to the four polytopes is zero-dimensional, there is always room to extend any polytope with empty interior in such a way that none of the original isolated tangent lines are lost.

We say that a line *properly* intersects a polygon if it intersects its relative interior. In the sequel, we use this definition only when the line and polygon are coplanar. Notice that a line that contains a segment is tangent to the segment as well as properly intersects it.

Let l_e be the line containing e and let Π_t denote the sweep plane parameterized by $t \in [0, \pi]$ such that Π_t contains the line l_e for all t and $\Pi_0 = \Pi_\pi$. Each plane Π_t intersects the three polytopes \mathbf{P} , \mathbf{Q} , and \mathbf{R} in three, possibly degenerate or empty, convex polygons, P_t , Q_t , and R_t , respectively (see Figure 2.1).

For any t , a *bitangent* to polygons P_t and Q_t is a line tangent to P_t and Q_t in Π_t (the line may intersect the polygon R_t in any way, possibly not at all). For any t , let a (P_t, Q_t) -tuple be the unordered set of all supports in \mathbf{P} and \mathbf{Q} of one of the bitangents to polygons P_t and Q_t . Note that a support in \mathbf{P} may be identical to a support in \mathbf{Q} , in which case the (P_t, Q_t) -tuple does not contain duplicates. Also note that a (P_t, Q_t) -tuple consists of exactly one support in \mathbf{P} and one support in \mathbf{Q} (possibly identical) except when the corresponding bitangent is tangent to \mathbf{P} (or \mathbf{Q}) along a face (either intersecting the face properly or containing one of its edges); then the (P_t, Q_t) -tuple contains two supports in \mathbf{P} (or \mathbf{Q}) instead of one. A \mathbf{PQ} -tuple is a set of edges and vertices that is a (P_t, Q_t) -tuple for some t . We define similarly the (Q_t, R_t) -tuples and \mathbf{QR} -tuples.

We say that a (P_t, Q_t) -tuple is *maximal for some t* if it is not contained in any other




FIG. 2.3. Plane Π_t is *F-critical*: it contains a line that lies in a plane Ψ containing a face of \mathbf{P} such that the line is tangent to $\mathbf{Q} \cap \Psi$ at a point not on l_e .

(P_t, Q_t) -tuple, for the same t . Note that a (P_t, Q_t) -tuple is non-maximal for some t if and only if all its supports intersect Π_t in one and the same point, and P_t and Q_t are not equal to one and the same point (see Figure 2.5(b)).

For any t , let a (P_t, Q_t, R_t) -tuple be the union of a (P_t, Q_t) -tuple and a (Q_t, R_t) -tuple that share at least one support in \mathbf{Q} . A (P_t, Q_t, R_t) -tuple is maximal for some t if it is not contained in any other (P_t, Q_t, R_t) -tuple, for the same t . A **PQR-tuple** is a set of edges and vertices that is a (P_t, Q_t, R_t) -tuple for some t . Note that a **PQR-tuple** typically consists of three supports, one from each polytope, and consists, in all cases, of at most two supports in \mathbf{P} , at most three supports in \mathbf{Q} , and at most two supports in \mathbf{R} .

A line intersecting e and tangent to \mathbf{P} , \mathbf{Q} , \mathbf{R} and \mathbf{S} is called a *generic tangent line* if and only if it intersects \mathbf{S} only on e and is tangent to P_t , Q_t , and R_t in some plane Π_t . Otherwise it is called a *non-generic tangent line*. A non-generic tangent line properly intersects a face of \mathbf{S} or properly intersects P_t , Q_t , or R_t in some plane Π_t . In the latter case P_t , Q_t , or R_t is a face or an edge of \mathbf{P} , \mathbf{Q} , or \mathbf{R} lying in Π_t ; thus a non-generic tangent line is (in both cases) tangent to \mathbf{P} , \mathbf{Q} , \mathbf{R} and \mathbf{S} in a plane containing a face or two edges of these polytopes, a degenerate situation.

In the following three subsections, we bound the number of generic and non-generic tangent lines. It is helpful to keep in mind that, as observed earlier, two convex polygons in a plane Π_t (such as P_t and Q_t) may admit a linear number of tangents if they intersect.

2.3. Generic tangent lines.

LEMMA 2.1. *The set of supports in \mathbf{P} , \mathbf{Q} , and \mathbf{R} of a generic tangent line is a **PQR-tuple**.*

Proof. Any generic tangent line ℓ is tangent in Π_t to P_t , Q_t , and R_t for some value t . Thus the set of supports of ℓ in \mathbf{P} and \mathbf{Q} (resp. in \mathbf{Q} and \mathbf{R}) is a (P_t, Q_t) -tuple (resp. a (Q_t, R_t) -tuple). Moreover the (P_t, Q_t) -tuple and the (Q_t, R_t) -tuple contain the same supports in \mathbf{Q} , and thus their union is a (P_t, Q_t, R_t) -tuple, hence a **PQR-tuple**. \square

We now define the *critical planes* Π_t in such a way that, as we will later prove, the set of (P_t, Q_t, R_t) -tuples is invariant for t ranging strictly between two consecutive critical values. We introduce two types of critical planes: the *V-critical* and *F-critical* planes.

A plane Π_t is *V-critical* if it contains a vertex of \mathbf{P} , \mathbf{Q} , or \mathbf{R} , not on l_e . (The constraint that the vertex does *not* lie on l_e ensures that the number of V-critical planes is finite even in degenerate configurations.) A plane Π_t is *F-critical* relative to an ordered pair of polytopes (\mathbf{P}, \mathbf{Q}) if (see Figure 2.3) it contains a line ℓ such that





FIG. 2.4. Plane Π_{t^*} contains a line m such that (i) m lies in a plane $\Psi \neq \Pi_{t^*}$ containing a face of \mathbf{P} , and (ii) m is tangent to polygon $\mathbf{P} \cap \Psi$ at some point not on l_e ; however m is not tangent to $\mathbf{Q} \cap \Psi$. If the definition of F -critical planes was not considering such plane Π_{t^*} to be F -critical then Lemma 2.3 would not hold. Indeed the set u of supports of line $\Pi_{t^*-\epsilon} \cap \Psi$ is a maximal (P_t, Q_t) -tuple for some but not all t in any open neighborhood of t^* , and, although Π_{t^*} is V -critical, there exists no V -critical event (t^*, v) such that u contains v or an edge with endpoint v .

- (i) ℓ lies in a plane $\Psi \neq \Pi_t$ containing a face of \mathbf{P} , and
- (ii) ℓ is tangent in Ψ to polygon $\mathbf{Q} \cap \Psi$ or $\mathbf{P} \cap \Psi$, at some point not on l_e .

For simplicity, we do not require that ℓ is tangent to \mathbf{P} ; this leads to overestimating the number of common tangents to \mathbf{P} , \mathbf{Q} , \mathbf{R} , and \mathbf{S} but only by an asymptotically negligible amount. Note that not all lines in Ψ tangent to \mathbf{Q} are tangent to the polygon $\mathbf{Q} \cap \Psi$ when that polygon is a face or edge of \mathbf{Q} lying in Ψ . Note also that we define Π_t to be F -critical when ℓ is tangent to $\mathbf{P} \cap \Psi$ at some point not on l_e only for handling the very degenerate case where $\mathbf{Q} \cap \Psi$ is an edge of \mathbf{Q} and there exists a line in Ψ that properly intersects $\mathbf{Q} \cap \Psi$ and is tangent to $\mathbf{P} \cap \Psi$ along an edge that has an endpoint on l_e (see Figure 2.4). Note finally that if $\ell \in \Pi_t$ satisfies (i) and is tangent, in Ψ , to $\mathbf{P} \cap \Psi$ at some point not on l_e then polytope \mathbf{Q} plays no role and thus Π_t is F -critical relative to (\mathbf{P}, \mathbf{Q}) for all polytopes \mathbf{Q} .

F -critical planes relative to (\mathbf{Q}, \mathbf{P}) , (\mathbf{Q}, \mathbf{R}) , and (\mathbf{R}, \mathbf{Q}) are defined similarly. A plane Π_t is F -critical if it is F -critical relative to pairs of polytopes (\mathbf{P}, \mathbf{Q}) , (\mathbf{Q}, \mathbf{P}) , (\mathbf{Q}, \mathbf{R}) , or (\mathbf{R}, \mathbf{Q}) .

The values of t corresponding to critical planes Π_t are called *critical values*. We call V -critical and F -critical events the ordered pairs (t, o) where t is a critical value and o is a vertex or line depending on the type of critical event. In a V -critical event, o is a vertex of \mathbf{P} , \mathbf{Q} , or \mathbf{R} that belongs to $\Pi_t \setminus l_e$. In an F -critical event, o is a line lying in some plane Π_t and satisfying Conditions (i-ii) above. A *critical event* is a V -critical or F -critical event.

FIG. 2.5. Lines through x in Π_t and tangent to P_t and Q_t .

LEMMA 2.2. *There are at most $\frac{2}{3}(p+q+r)$ V-critical events and $\frac{8}{3}(p+2q+r)$ F-critical events.*

Proof. The number of V-critical events is at most the total number of vertices of \mathbf{P} , \mathbf{Q} , and \mathbf{R} , and hence is less than two thirds the total number of edges of \mathbf{P} , \mathbf{Q} , and \mathbf{R} . We now count the number of F-critical events relative to polytopes (\mathbf{P}, \mathbf{Q}) . Let Ψ be a plane containing a face of \mathbf{P} , and suppose that for some plane Π_t , line $\ell = \Pi_t \cap \Psi$ satisfies Conditions (i-ii). Plane Ψ does not contain l_e because otherwise both l_e and ℓ lie in the two distinct planes Ψ and Π_t , so $\ell = l_e$ but then ℓ cannot satisfy Condition (ii). Furthermore ℓ and l_e intersect or are parallel since they both lie in Π_t . Thus if $\Psi \cap l_e$ is a point then ℓ contains it, and otherwise $\Psi \cap l_e = \emptyset$ and ℓ is parallel to l_e .

If $\Psi \cap l_e$ is a point, there are at most four candidates for a line ℓ in plane Ψ going through $\Psi \cap l_e$ and tangent to $\mathbf{Q} \cap \Psi$ or $\mathbf{P} \cap \Psi$ at some point not on l_e . Likewise, if $\Psi \cap l_e$ is empty, there are at most four candidates for a line ℓ in plane Ψ that is parallel to l_e and tangent to $\mathbf{Q} \cap \Psi$ or $\mathbf{P} \cap \Psi$. In either case, each candidate line is contained in a unique plane Π_t , for $t \in [0, \pi]$, since $\ell \neq l_e$ (ℓ contains a point not on l_e). Hence, a face of \mathbf{P} generates at most four F-critical events relative to (\mathbf{P}, \mathbf{Q}) . Therefore the number of critical events relative to (\mathbf{P}, \mathbf{Q}) is at most $\frac{8}{3}p$ since the number of faces of a polytope is at most two thirds the number of its edges. Hence the number of critical events relative to (\mathbf{P}, \mathbf{Q}) , (\mathbf{Q}, \mathbf{P}) , (\mathbf{Q}, \mathbf{R}) and (\mathbf{R}, \mathbf{Q}) is at most $\frac{8}{3}(p+2q+r)$. \square

The following lemma states that the critical planes have the desired property. Let u_e be the set of supports of l_e in \mathbf{P} and \mathbf{Q} and let u denote some (P_t, Q_t) -tuple.

LEMMA 2.3. *Let t^* be the endpoint of a maximal interval³ throughout which $u \neq u_e$ is a maximal (P_t, Q_t) -tuple. Then t^* is a critical value. Moreover, there exists a V-critical event (t^*, v) or a F-critical event (t^*, m) such that u contains v or an edge with endpoint v , or u is contained in the set of supports of m .*

The proof of this lemma is rather long and intricate; we postpone it to Section 2.4. Note that, as stated, this lemma only applies under the assumptions that u is maximal and distinct from u_e . These assumptions are made in order to simplify the proof of Lemma 2.3; we don't suggest that the lemma is false without them.

LEMMA 2.4. *Any edge or vertex of \mathbf{P} or \mathbf{Q} is in at most 2 \mathbf{PQ} -tuples that are maximal (P_t, Q_t) -tuples for all t in any given non-empty interval³ of $\mathbb{R}/\pi\mathbb{Z}$.*

Proof. Let \tilde{t} be an element of a non-empty interval I of $\mathbb{R}/\pi\mathbb{Z}$ and x be an edge or vertex of \mathbf{P} or \mathbf{Q} . If x does not intersect $\Pi_{\tilde{t}}$ then no $(P_{\tilde{t}}, Q_{\tilde{t}})$ -tuple contains x . If x intersects $\Pi_{\tilde{t}}$ in one point then there are, in general, at most two lines in $\Pi_{\tilde{t}}$ going through x and tangent to $P_{\tilde{t}}$ and $Q_{\tilde{t}}$ (see Figure 2.5(a)); in all cases there are at most 3 $(P_{\tilde{t}}, Q_{\tilde{t}})$ -tuples containing x (see Figure 2.5(b)), however at most 2 of them are maximal. If x intersects $\Pi_{\tilde{t}}$ in more than one point, x is an edge lying in $\Pi_{\tilde{t}}$. Then any line in $\Pi_{\tilde{t}}$ intersecting x and tangent to $P_{\tilde{t}}$ and $Q_{\tilde{t}}$ contains an endpoint of x and thus x belongs to no $(P_{\tilde{t}}, Q_{\tilde{t}})$ -tuple.

³Such an interval could be open or closed, a single point or an interval of positive length.

Hence at most 2 **PQ**-tuples contain x and are maximal (P_t, Q_t) -tuples for $t = \tilde{t}$, and thus at most 2 **PQ**-tuples contain x and are maximal (P_t, Q_t) -tuples for all t in I . \square

LEMMA 2.5. *There are at most $O(p + q + r)$ **PQR**-tuples.*

Proof. In order to count the number of distinct (P_t, Q_t, R_t) -tuples, we charge each maximal (P_t, Q_t, R_t) -tuple to a critical event. We then show that each critical event is charged at most a constant number of times. It then follows from Lemma 2.2 that there are $O(p + q + r)$ distinct maximal (P_t, Q_t, R_t) -tuples. A maximal (P_t, Q_t, R_t) -tuple consists of at most two supports in **P**, at most three supports in **Q**, and at most two supports in **R**, and thus contains at most $(2^2 - 1)(2^3 - 1)(2^2 - 1)$ distinct subsets with at least one support in each of **P**, **Q** and **R**. Each maximal (P_t, Q_t, R_t) -tuple thus contains at most a constant number of distinct (P_t, Q_t, R_t) -tuples, which implies the result.

Let s be a maximal (P_t, Q_t, R_t) -tuple and let I be any maximal connected subset of $\mathbb{R}/\pi\mathbb{Z}$ such that s is a maximal (P_t, Q_t, R_t) -tuple for all $t \in I$. Let u be a maximal (P_t, Q_t) -tuple and u' a maximal (Q_t, R_t) -tuple such that the union of u and u' is s and such that u and u' share at least one support in **Q**.

First, suppose that $I = \mathbb{R}/\pi\mathbb{Z}$. Then u is a maximal (P_t, Q_t) -tuple for all $t \in \mathbb{R}/\pi\mathbb{Z}$. Thus each support in u intersects Π_t for all $t \in \mathbb{R}/\pi\mathbb{Z}$ and thus intersects l_e ; moreover each support in u intersects Π_t only on l_e for all $t \in \mathbb{R}/\pi\mathbb{Z}$ except possibly for one value of t . Since **P** and **Q** have non-empty interior, $P_t \cup Q_t$ is not reduced to a point for all t in some interval of positive length. For all t in such an interval, since u is maximal, the union of the supports in u intersects Π_t in at least two distinct points. These at least two distinct points lie on l_e for some values of t by the above argument. Thus, for these values of t , l_e is the only line in Π_t whose set of supports contains u . Hence u is the set of supports of l_e . The same property holds for v and thus s is also the set of supports of l_e . We can thus assume in the following that $I \neq \mathbb{R}/\pi\mathbb{Z}$, and only count the maximal (P_t, Q_t, R_t) -tuples that are not the set of supports of l_e .

Interval I is thus a non-empty interval of $\mathbb{R}/\pi\mathbb{Z}$; it can be open or closed, a single point or an interval of positive length. Let w_0 and w_1 denote the endpoints of $I \neq \mathbb{R}/\pi\mathbb{Z}$.

If s contains a vertex v , or an edge with endpoint v , such that v lies in $\Pi_{w_i} \setminus l_e$, for $i = 0$ or 1 , then we charge s to the V-critical event (w_i, v) . Otherwise, we charge s to an F-critical event (w_i, m) where m is a line in Π_{w_i} whose set of supports contains u or u' . Such a V-critical or F-critical event exists by Lemma 2.3.

We now prove that each critical event is charged by at most a constant number of distinct maximal (P_t, Q_t, R_t) -tuples. As mentioned before, that will imply the result.

Consider a V-critical event (t^*, v) that is charged by a maximal (P_t, Q_t, R_t) -tuple s . By the charging scheme, s contains a support x that is v or an edge with endpoint v , and s is a maximal (P_t, Q_t, R_t) -tuple for all t in at least one of three intervals, $\{t^*\}$ and two open intervals having t^* as endpoint; denote these intervals by I_1, I_2, I_3 .

By Lemma 2.4, at most 2 **PQ**-tuples contain x and are maximal (P_t, Q_t) -tuples for all t in I_i . Moreover, each of these **PQ**-tuples contains at most 2 supports in **Q**, and each of these supports belongs to at most 2 **QR**-tuples that are maximal (Q_t, R_t) -tuples for all t in I_i . Thus at most 8 **PQR**-tuples contain x and are maximal (P_t, Q_t, R_t) -tuples for all t in I_i , for each $i = 1, \dots, 3$. Hence any V-critical event (t^*, v) is charged by at most 24 distinct maximal (P_t, Q_t, R_t) -tuples.

Consider now an F-critical event (t^*, m) that is charged by a maximal (P_t, Q_t, R_t) -tuple s , and define as before u and u' . By the charging scheme, the set of supports of m contains u or u' (or both); suppose without loss of generality that it contains u . The set of supports of m contains at most two supports in **P** and at most two supports in **Q**. Since u contains at least one support in **P** and at least one support in **Q**, there are at most 3^2 choices for u .

By the charging scheme, s is a maximal (P_t, Q_t, R_t) -tuple for all t in at least one of 3 intervals, $\{t^*\}$ and two open intervals having t^* as endpoint; denote by I_1, I_2, I_3 these intervals. It follows from Lemma 2.4 that, for each support x of \mathbf{Q} in u , at most 2 \mathbf{QR} -tuples contain x and are maximal (Q_t, R_t) -tuples for all t in I_i . There are at most 3^2 choices for u (as shown above), 2 for x , 3 for i and 2 for the \mathbf{QR} -tuples containing x . Hence any F-critical event (t^*, m) is charged by at most $2^2 \times 3^3$ distinct maximal (P_t, Q_t, R_t) -tuples.

Therefore each critical event is charged by at most a constant number of distinct maximal (P_t, Q_t, R_t) -tuples, which concludes the proof. \square

COROLLARY 2.6. *There are at most $O(p+q)$ \mathbf{PQ} -tuples.*

Proof. Replace \mathbf{R} by a copy of \mathbf{Q} in Lemma 2.5. Any \mathbf{PQ} -tuple is also a \mathbf{PQQ} -tuple, and there are at most $O(p+q+q) = O(p+q)$ of these. \square

PROPOSITION 2.7. *There are $O(p+q+r)$ isolated generic tangent lines.*

Proof. A generic tangent line is transversal to e and to the edges and vertices of a \mathbf{PQR} -tuple, by definition and Lemma 2.1. An isolated generic tangent line is thus an isolated transversal with respect to a set of edges and vertices that consists of a \mathbf{PQR} -tuple and either edge e or one or both of its endpoints. The number of such sets is four times the number of \mathbf{PQR} -tuples, which is in $O(p+q+r)$ by Lemma 2.5. The result follows since each such set consists of at most eight edges and vertices (at most two supports from each of the four polytopes) and thus admits at most eight isolated transversals [9]. \square

2.4. Proof of Lemma 2.3. Recall that u_e denotes the set of supports of l_e in \mathbf{P} and \mathbf{Q} , and that Lemma 2.3 states the following.

Let t^ be the endpoint of a maximal interval throughout which $u \neq u_e$ is a maximal (P_t, Q_t) -tuple. Then t^* is a critical value. Moreover, there exists a V-critical event (t^*, v) or a F-critical event (t^*, m) such that u contains v or an edge with endpoint v , or u is contained in the set of supports of m .*

We can assume that u contains no vertex v and no edge with endpoint v , such that v lies on $\Pi_{t^*} \setminus l_e$ because otherwise (t^*, v) is a V-critical event such that u contains v or an edge with endpoint v , which concludes the proof.

We prove a series of lemmas that yields Lemma 2.3. Indeed, we prove the existence of a line m in Π_{t^*} whose set of supports contains u (Lemma 2.10) such that (i) m lies in a plane $\Psi \neq \Pi_{t^*}$ containing a face of \mathbf{P} (Lemma 2.11), and (ii) m is tangent in Ψ to polygon $\mathbf{Q} \cap \Psi$ or $\mathbf{P} \cap \Psi$, at some point not on l_e (Lemma 2.12). This proves that Π_{t^*} contains a line m whose set of supports contains u and such that (t^*, m) is an F-critical event, which concludes the proof.

By hypothesis, for any sufficiently small open neighborhood \mathcal{N} of t^* whose endpoints are denoted by t_0 and t_1 , u is not a maximal (P_t, Q_t) -tuple for some $t \in \mathcal{N}$ and u is a maximal (P_t, Q_t) -tuple for $t = t^*$ or for all $t \in (t^*, t_1)$ (or by symmetry for all $t \in (t_0, t^*)$).

We only consider in the following supports in \mathbf{P} and in \mathbf{Q} ; polytope \mathbf{R} plays no role. We start by proving two preliminary lemmas.

LEMMA 2.8. *Each support in u intersects Π_t in exactly one point (possibly on l_e), for all t in any sufficiently small open neighborhood \mathcal{N} of t^* .*

Moreover, the union of all supports in u intersects Π_t in at least two distinct points for all $t \neq t^$ in \mathcal{N} . This property also holds for $t = t^*$ if u is a maximal (P_{t^*}, Q_{t^*}) -tuple.*

Proof. Since u is a (P_t, Q_t) -tuple for some t in every open neighborhood of t^* , each support in u intersects Π_t for some t in every open neighborhood of t^* . It thus follows from the assumption that u contains no vertex v and no edge with endpoint v , such that v lies on $\Pi_{t^*} \setminus l_e$, that each support in u intersects Π_t for all t in any sufficiently small open neighborhood \mathcal{N} of t^* . It follows that each support in u either lies in l_e or intersects Π_t in exactly one point for all $t \in \mathcal{N}$. However, no edge of u lies in l_e because otherwise, if x denotes such an edge of, say, \mathbf{P} , then any line tangent to P_t in Π_t and intersecting x contains an endpoint of x which is

a vertex of \mathbf{P} ; thus, by definition, u does not contain x but one of its endpoints. Hence each support of u intersects Π_t in exactly one point for all $t \in \mathcal{N}$.

We now prove that the union of the supports in u intersects Π_t in at least two distinct points for any $t \in \mathcal{N}$ such that u is a maximal (P_t, Q_t) -tuple. Suppose for a contradiction that the union of the supports in u intersects Π_t in one single point v for some $t \in \mathcal{N}$ such that u is a maximal (P_t, Q_t) -tuple. Then polygons P_t and Q_t are both reduced to point v because otherwise u is not maximal (otherwise, a line in Π_t tangent to P_t and Q_t at v can be rotated about v until it becomes tangent to P_t or Q_t at some other points). Thus $v = P_t = Q_t$ is a vertex of \mathbf{P} and of \mathbf{Q} because the polytopes have non-empty interior. Hence $u = \{v\}$ because each support in u contains v . It follows that v lies on l_e since each support in u intersects Π_t for all $t \in \mathcal{N}$. Moreover, since P_t and Q_t are both reduced to point $v = l_e \cap \mathbf{P} = l_e \cap \mathbf{Q}$, the set u_e of supports of l_e is u , contradicting the hypotheses of Lemma 2.3.

Thus, if u is a maximal (P_t, Q_t) -tuple for all $t \in (t^*, t_1)$, the union of the supports in u intersects Π_t in at least two distinct points for all $t \in (t^*, t_1)$ and thus for all $t \neq t^*$ in any sufficiently small open neighborhood of t^* . Also, if u is a maximal (P_t, Q_t) -tuple for $t = t^*$, the union of the supports in u intersects Π_t in at least two distinct points for $t = t^*$ and thus for all t in any sufficiently small open neighborhood of t^* . \square

LEMMA 2.9. *If u is a maximal (P_{t^*}, Q_{t^*}) -tuple then u consists of at least three supports.*

Proof. Note that it follows from Lemma 2.8 that u contains at least two supports. Suppose for a contradiction that u consists of only two supports. By Lemma 2.8, they intersect Π_t in exactly two distinct points for all t in any sufficiently small open neighborhood \mathcal{N} of t^* . Thus there exists for all $t \in \mathcal{N}$ a unique line m_t in Π_t whose set of supports contains u ; moreover m_t is continuous in terms of t . Since u is a maximal (P_{t^*}, Q_{t^*}) -tuple, the set of supports of m_{t^*} is u . Thus, for all t in any sufficiently small \mathcal{N} , the set of supports of m_t is u . Thus the set of supports of m_t is invariant for $t \in \mathcal{N}$ and since m_{t^*} is tangent to P_{t^*} and Q_{t^*} , line m_t is tangent to P_t and Q_t for all $t \in \mathcal{N}$.

Hence, for all $t \in \mathcal{N}$, line m_t , whose set of supports is u , is tangent to P_t and Q_t in Π_t . Thus u is a maximal (P_t, Q_t) -tuple for all $t \in \mathcal{N}$. Moreover, m_t is the unique line in Π_t whose set of supports contains u , thus u is a maximal (P_t, Q_t) -tuple for all $t \in \mathcal{N}$, contradicting the hypotheses of the lemma. \square

LEMMA 2.10. *There exists a line m in Π_{t^*} whose set of supports contains u that is tangent to P_{t^*} and Q_{t^*} along an edge of one of them, say of P_{t^*} .*

Proof. Consider first the case where u is a maximal (P_{t^*}, Q_{t^*}) -tuple. There exists in Π_{t^*} a line m tangent to P_{t^*} and Q_{t^*} whose set of supports is u . By Lemma 2.9, the set u of supports of m contains at least three supports, and hence at least two supports in \mathbf{P} (or in \mathbf{Q}). Furthermore, the supports of m in one polytope intersect Π_{t^*} in distinct points (by definition of supports). Thus m intersects P_{t^*} (or Q_{t^*}) in at least two distinct points and is tangent to P_{t^*} and Q_{t^*} . The result follows since P_{t^*} (and Q_{t^*}) is convex.

Consider now the case where u is a maximal (P_t, Q_t) -tuple for all $t \in (t^*, t_1)$. Then, for all $t \in (t^*, t_1)$, there exists a line in Π_t tangent to P_t and Q_t and whose set of supports is u . Moreover, by Lemma 2.8, this line is unique for each $t \in (t^*, t_1)$ and varies continuously in terms of $t \in (t^*, t_1)$. When t tends to t^* , the line tends to a line m_{t^*} in Π_{t^*} which is tangent to P_{t^*} and Q_{t^*} and whose set of supports contains u . If its set of supports strictly contains u then m_{t^*} is tangent to P_{t^*} and Q_{t^*} along an edge of one of them because the polygons are convex, and hence we can choose $m = m_{t^*}$ to complete the proof. Otherwise, u is a (P_{t^*}, Q_{t^*}) -tuple.

We can suppose that u is a non-maximal (P_{t^*}, Q_{t^*}) -tuple since we already treated the case where u is maximal. There exists in Π_{t^*} a line tangent to P_{t^*} and Q_{t^*} whose set of supports is u . Since u is non-maximal this line is tangent to P_{t^*} and Q_{t^*} at a shared vertex, and can be rotated about this vertex in Π_{t^*} until it becomes tangent to P_{t^*} and Q_{t^*} at some other points,




FIG. 2.6. Line m is tangent to \mathbf{P} along a face in plane $\Psi \neq \Pi_{t^*}$.

which must occur because u is non-maximal; let m denote the resulting line. The set of supports of m contains u and m is tangent to P_{t^*} and Q_{t^*} along an edge of one of them because the polygons are convex. \square

LEMMA 2.11. *Line m lies in a plane $\Psi \neq \Pi_{t^*}$ containing a face of \mathbf{P} .*

Proof. By Lemma 2.10, m contains an edge of P_{t^*} ; see Figure 2.6. This edge either intersects the relative interior of some face of \mathbf{P} in which case we take Ψ to be the plane containing that face, or it is an edge of \mathbf{P} in which case we take Ψ to be a plane, different from Π_{t^*} , containing one of the two faces of \mathbf{P} incident to that edge. \square

Let m_t be the line $\Psi \cap \Pi_t$ for all t in any sufficiently small open neighborhood \mathcal{N} of t^* ; line m_t is well defined since $\Psi \cap \Pi_{t^*}$ is line m by Lemmas 2.10 and 2.11.

LEMMA 2.12. *Line m is tangent to $\mathbf{P} \cap \Psi$ or to $\mathbf{Q} \cap \Psi$, at some point not on l_e .*

Proof. We assume for a contradiction that line m does not satisfy the lemma, i.e., m is not tangent to $\mathbf{P} \cap \Psi$ or to $\mathbf{Q} \cap \Psi$ at any point other than on l_e . We prove that the set of supports of m is u and is a maximal (P_t, Q_t) -tuple for all t in any sufficiently small neighborhood of t^* , contradicting the hypotheses of Lemma 2.3 and thus proving Lemma 2.12.

Since m is tangent to \mathbf{Q} (by Lemma 2.10), m is tangent to $\mathbf{Q} \cap \Psi$ only on l_e (see Figure 2.7(a)), or m properly intersects $\mathbf{Q} \cap \Psi$ which is then a face or an edge of \mathbf{Q} (see Figure 2.7(b))⁴. Similarly m is tangent to $\mathbf{P} \cap \Psi$ only on l_e , or m properly intersects it; however $\mathbf{P} \cap \Psi$ is necessarily a face of \mathbf{P} by Lemma 2.11.

The following Lemmas 2.13 and 2.14 imply that the set of supports of m_t is invariant and equal to u for all t in any sufficiently small open neighborhood \mathcal{N} of t^* . Moreover, since m_t varies continuously with t and $m = m_{t^*}$ is tangent to P_{t^*} and Q_{t^*} (by Lemma 2.10), line m_t is tangent to P_t and Q_t for all $t \in \mathcal{N}$. Hence u is a (P_t, Q_t) -tuple for all $t \in \mathcal{N}$. We now prove that u is a maximal (P_t, Q_t) -tuple for all $t \in \mathcal{N}$.

As we have seen before, $m = m_{t^*}$ is tangent to \mathbf{P} in at least two points (by Lemma 2.10), thus m_{t^*} intersects its supports in at least two distinct points. Moreover the set of supports of m_{t^*} is u . Thus there is a unique line in Π_{t^*} whose set of supports contains u . Hence u is a maximal (P_{t^*}, Q_{t^*}) -tuple.

By Lemma 2.8, m_t is the unique line in Π_t whose set of supports contains u for all $t \neq t^*$ in \mathcal{N} . Thus u is a maximal (P_t, Q_t) -tuple for all $t \neq t^*$ in \mathcal{N} .

Hence u is a maximal (P_t, Q_t) -tuple for all $t \in \mathcal{N}$, contradicting the hypotheses of Lemma 2.3 and thus concluding the proof of Lemma 2.12. \square

LEMMA 2.13. *The set of supports of m_t is u for some t in any sufficiently small open neighborhood \mathcal{N} of t^* .*

Proof. We first prove that the supports in u are supports of m_t for all $t \in \mathcal{N}$. A support vertex in u lies on l_e by Lemma 2.8 and thus lies in Π_t for all t . A support vertex in u also

⁴Note that in these two situations, two edges of two distinct polytopes are then coplanar (in the first case an edge of \mathbf{Q} and e are coplanar, and in the later case a face of \mathbf{P} is coplanar with a face or an edge of \mathbf{Q}). Hence proving this lemma is straightforward under some general position assumption that excludes such situations.




FIG. 2.7. m is tangent to \mathbf{P} along a face in Ψ and (a) to $\mathbf{Q} \cap \Psi$ only on l_e or (b) to \mathbf{Q} along a face in Ψ .

lies on m by Lemma 2.10 and thus lies in plane Ψ by Lemma 2.11. Hence, for all $t \in \mathcal{N}$, the support vertices in u lie on m_t , and thus are supports of m_t .

In order to prove that the support edges in u are supports of m_t , it is sufficient (by Lemma 2.10) to prove that the support edges of m are supports of m_t . The support edges of m in \mathbf{P} lie in plane Ψ (see Figure 2.7(b)) because Ψ contains m and a face of \mathbf{P} (indeed if m intersects an edge of \mathbf{P} not in Ψ then m contains one of its endpoints, and thus the edge is not a support). Thus all the support edges of m lie in Ψ and m contains none of their endpoints (by definition). Since m_t lies in Ψ for all t and $m_t^* = m$, line m_t intersects all the support edges of m and contains none of their endpoints for all t in any sufficiently small open neighborhood \mathcal{N} of t^* . Hence the support edges of m in \mathbf{P} are supports of m_t for all $t \in \mathcal{N}$.

Consider the case where $\mathbf{Q} \cap \Psi$ is a face or an edge of \mathbf{Q} . Similarly as for \mathbf{P} , the support edges of m in \mathbf{Q} lie in plane Ψ , and thus are supports of m_t for all $t \in \mathcal{N}$.

Consider now the case where m is tangent to $\mathbf{Q} \cap \Psi$ only on l_e at, say, point v (see Figure 2.7(a)). Then v lies in Ψ (since $m \subset \Psi$ by Lemma 2.11) and also lies in Π_t for all t (since $l_e \subset \Pi_t$ for all t). Hence m_t contains v for all $t \in \mathcal{N}$. Moreover, m_t is tangent to $\mathbf{Q} \cap \Psi$ only at v for all t in any sufficiently small open neighborhood \mathcal{N} of t^* . Hence the set of supports of m_t in \mathbf{Q} is invariant for all $t \in \mathcal{N}$.

We have so far proved that the set of supports of m_t contains u for all $t \in \mathcal{N}$.

We now prove that the set of supports of m_t is u for some $t \in \mathcal{N}$. Consider first the case where u is a maximal (P_{t^*}, Q_{t^*}) -tuple. Then, by Lemma 2.8, the union of the supports in u intersects Π_{t^*} in at least two distinct points, thus $m_{t^*} = m$ is the only line in Π_{t^*} whose set of supports contains u . Moreover, since u is a (P_{t^*}, Q_{t^*}) -tuple, there exists a line in Π_{t^*} whose set of supports is u . Hence the set of supports of m_{t^*} is u .

Consider now the case where u is a maximal (P_t, Q_t) -tuple for all $t \in (t^*, t_1)$. By Lemma 2.8, for all $t \in (t^*, t_1)$, the union of the supports in u intersects Π_t in at least two distinct points, thus m_t is the only line in Π_t whose set of supports contains u . For all $t \in (t^*, t_1)$, since u is a (P_t, Q_t) -tuple there exists a line in Π_t whose set of supports is u . Hence the set of supports of m_t is u for all $t \in (t^*, t_1)$. \square

LEMMA 2.14. *The set of supports of m_t is invariant for t ranging in any sufficiently small open neighborhood \mathcal{N} of t^* .*

Proof. First if $m = l_e$ then $m_t = l_e$ for all $t \in \mathcal{N}$ because Ψ contains $m = l_e$ (by Lemma 2.11) and Π_t contains l_e for all t (by definition). Thus the set of supports of m_t is invariant for all $t \in \mathcal{N}$. We now assume that $m \neq l_e$.

Line m is tangent to polygon P_{t^*} along an edge by Lemma 2.10. Thus m is tangent to \mathbf{P} in at least two points. Hence, since $\mathbf{P} \cap \Psi$ is a face of \mathbf{P} and m lies in Ψ , either m properly intersects $\mathbf{P} \cap \Psi$ or m is tangent to $\mathbf{P} \cap \Psi$ along one of its edges. In the later case, the edge does not lie in l_e since $m \neq l_e$, thus m is tangent to $\mathbf{P} \cap \Psi$ at some point not on l_e , contradicting

our assumptions. Hence m properly intersects the face of \mathbf{P} in Ψ .

It follows that, if m contains a vertex of \mathbf{P} , then this vertex is an endpoint of a support edge of m_t for all t in any sufficiently small open neighborhood of t^* (indeed m_t lies in Ψ and tends to m when t tends to t^*). By Lemma 2.13, the set of supports of m_t is u for some t in any sufficiently small open neighborhood of t^* . Hence, if m contains a vertex of \mathbf{P} , this vertex is an endpoint of a support edge in u . By assumption u contains no edge with endpoint on $\Pi_{t^*} \setminus l_e$, thus m contains no vertex of \mathbf{P} except possibly on l_e (since m lies in Π_{t^*}). It thus follows that the set of supports of m_t in \mathbf{P} is invariant for t ranging in any sufficiently small open neighborhood of t^* (since $m_t \subset \Psi$ tends to m when t tends to t^* and all supports of m lie in Ψ).

Now consider the case where m properly intersects $\mathbf{Q} \cap \Psi$ which is a face or an edge of \mathbf{Q} . Similarly as for \mathbf{P} , m contains no vertex of \mathbf{Q} except possibly on l_e and thus the set of supports of m_t in \mathbf{Q} is invariant for t ranging in any sufficiently small open neighborhood of t^* .

Finally, consider the case where m is tangent to $\mathbf{Q} \cap \Psi$ only on l_e . Then, as in the proof of Lemma 2.13, the set of supports of m_t in \mathbf{Q} is invariant for all t ranging in any sufficiently small open neighborhood of t^* , which concludes the proof. \square

2.5. Non-generic tangent lines. We count here the number of non-generic tangent lines. Note that, as mentioned before, there are no such lines under some adequate general position assumption.

PROPOSITION 2.15. *There are at most $O(p+q+r)$ isolated non-generic tangent lines except possibly for those that lie in planes that contain e and are tangent to all four polytopes.*

Proof. An isolated non-generic tangent line lies in plane Π_t for some t and contains (at least) two distinct points, each of which is a vertex of \mathbf{P} , \mathbf{Q} , \mathbf{R} , or \mathbf{S} , or a point of tangency between the line and one of the polygons P_t , Q_t , and R_t ; indeed, otherwise the line can be moved in Π_t while keeping the same supports.

We count first the isolated non-generic tangent lines that contain two distinct points of tangency with two of the polygons P_t , Q_t , and R_t in Π_t for some t . Consider such a line ℓ tangent to, say, P_t and Q_t in Π_t . Line ℓ is non-generic and thus properly intersects a face of \mathbf{S} or a face or an edge of \mathbf{R} lying in Π_t . If ℓ properly intersects a face of \mathbf{S} or a face or an edge of \mathbf{R} lying in Π_t but not entirely contained in l_e , then Π_t is one of the at most four planes tangent to \mathbf{R} or \mathbf{S} . There are $O(p+q)$ lines tangent to P_t and Q_t in two distinct points in each of these planes and thus $O(p+q)$ such lines in total. Otherwise, Π_t intersects each of \mathbf{R} and \mathbf{S} in an edge contained in l_e . The supports of ℓ are thus the union of a \mathbf{PQ} -tuple, and of, in each of \mathbf{R} and \mathbf{S} , the edge lying in l_e or one (or both) of its endpoint. It follows that at most a constant number of such isolated non-generic tangent lines contain a given \mathbf{PQ} -tuple in its set of supports. Hence the number of such lines is at most the number of \mathbf{PQ} -tuples, which is in $O(p+q)$ by Corollary 2.6. It follows that there are at most $O(p+q+r)$ isolated non-generic tangent lines that contain two distinct points of tangency with two of the polygons P_t , Q_t , and R_t in Π_t for some t . We obtain similarly that there are at most $O(p+q+r)$ isolated non-generic tangent lines that contain two distinct points of tangency with only one the polygons P_t , Q_t , and R_t .

We now count the isolated non-generic tangent lines that contain a unique vertex of \mathbf{P} , \mathbf{Q} , \mathbf{R} , or \mathbf{S} and a unique point of tangency with the polygons P_t , Q_t , and R_t in Π_t for some t . Each vertex v of \mathbf{P} , \mathbf{Q} , \mathbf{R} , or \mathbf{S} that does not lie on l_e is contained in a unique plane Π_t and there are, in that plane, at most six lines through v and tangent to P_t , Q_t , or R_t . There are thus $O(p+q+r)$ such lines in total. Consider now a line ℓ through a vertex v on l_e and tangent to P_t at $w \neq v$ in Π_t for some t . We can suppose that each of Q_t and R_t is either tangent to ℓ at w or is properly intersected by ℓ ; indeed otherwise ℓ is tangent to two polygons in two distinct

points. If Q_t (or R_t) is a face of \mathbf{Q} (resp. \mathbf{R}) or an edge not contained in l_e then Π_t is one of the at most two planes tangent to \mathbf{Q} (resp. \mathbf{R}) and, in each of these planes, there are at most two lines through v and tangent to P_t . If Q_t (or R_t) is tangent to ℓ at w such that the support edges of ℓ in \mathbf{P} and in \mathbf{Q} (resp. \mathbf{R}) are not collinear then ℓ goes through a vertex of \mathbf{P} , \mathbf{Q} , \mathbf{R} , or \mathbf{S} that lies on l_e , and through a vertex of the intersection of two of these polytopes. There are at most eight vertices of \mathbf{P} , \mathbf{Q} , \mathbf{R} , and \mathbf{S} on l_e and $O(p+q+r)$ vertices on the intersection of two of these polytopes. There are thus $O(p+q+r)$ such lines in total. Otherwise, Q_t (and R_t) is an edge contained in l_e or is tangent to ℓ at w such that the support edges of ℓ in \mathbf{P} and in \mathbf{Q} (resp. \mathbf{R}) are collinear; then ℓ is not isolated.

We finally bound the number of isolated non-generic tangent lines that contain no point of tangency with the polygons P_t , Q_t , and R_t in Π_t for any t (and thus contain at least two vertices of \mathbf{P} , \mathbf{Q} , \mathbf{R} , and \mathbf{S}). Consider such a line ℓ that lies in plane Π_t for some t . Line ℓ is tangent to \mathbf{P} , \mathbf{Q} , and \mathbf{R} and thus properly intersect P_t , Q_t , and R_t in plane Π_t which is tangent to \mathbf{P} , \mathbf{Q} , and \mathbf{R} . If plane Π_t is not tangent to \mathbf{S} , ℓ goes through an endpoint of e (since ℓ is tangent to \mathbf{S}) and there are $O(p+q+r)$ such lines ℓ that go through an endpoint of e and at least another vertex of \mathbf{P} , \mathbf{Q} , or \mathbf{R} . If plane Π_t is tangent to \mathbf{S} , line ℓ lies in a plane Π_t tangent to \mathbf{P} , \mathbf{Q} , \mathbf{R} , and \mathbf{S} , which concludes the proof. \square

Note that there can be $\Omega(n^2)$ isolated non-generic tangent lines that lie in a plane tangent to all four polytopes. Consider, for instance, four polytopes that admit a common tangent plane containing edge e , an edge e' of \mathbf{P} , and two faces of \mathbf{Q} and \mathbf{R} of linear complexity such that all the lines through a vertex of each face intersect e and e' . All these lines are isolated non-generic tangent lines.

2.6. Proof of the Main Lemma. Proposition 2.7, which handles the isolated generic tangent lines, and Proposition 2.15, which handles the isolated non-generic tangent lines, directly yield the Main Lemma.

3. Upper bounds. We prove in this section the upper bounds of Theorems 1.1, 1.2, and 1.3. The lower bounds are proved in Section 4. Consider k pairwise distinct polytopes $\mathbf{P}_1, \dots, \mathbf{P}_k$ with n_1, \dots, n_k edges, respectively, and n edges in total.

LEMMA 3.1. *For any edge e of \mathbf{P}_i , there are $O(n_j + n_l + n_m)$ sets of open edges, chosen from \mathbf{P}_i , \mathbf{P}_j , \mathbf{P}_l , and \mathbf{P}_m , that admit an isolated transversal that intersects e and is tangent to these four polytopes.*

Proof. Any isolated transversal to a set of edges is isolated with respect to the set of all its supports. It is thus sufficient to bound the number of sets of open edges, chosen from \mathbf{P}_i , \mathbf{P}_j , \mathbf{P}_l , and \mathbf{P}_m , that are intersected by an isolated line that intersects e and is tangent to these four polytopes. The Main Lemma states that there are $O(n_j + n_l + n_m)$ isolated lines intersecting e and tangent to \mathbf{P}_i , \mathbf{P}_j , \mathbf{P}_l , and \mathbf{P}_m , excluding those that lie in planes that contain e and are tangent to all four polytopes. Any of these $O(n_j + n_l + n_m)$ isolated lines intersects at most two open edges in any polytope. Thus there are $O(n_j + n_l + n_m)$ sets of open edges (chosen from \mathbf{P}_i , \mathbf{P}_j , \mathbf{P}_l , and \mathbf{P}_m) that are intersected by one of these isolated lines. Now consider any isolated line that lies in a plane that contains e and is tangent to all four polytopes. This plane contains all the open edges that are intersected by the isolated line. Thus these edges (and any subset of them) admit no isolated transversal. \square

LEMMA 3.2. *A minimal set of open edges and vertices that admit an isolated transversal consists of (i) two vertices, (ii) one vertex and one or two edges, or (iii) two, three, or four edges.*

Proof. Consider a minimal set of open edges and vertices that admits an isolated transversal. The elements are necessarily distinct because the set is minimal. If the set contains two vertices, it contains no other element since the two vertices admit a unique transversal.

Suppose now that the set contains one vertex. None of the open edges contain the vertex because otherwise such an edge would be redundant. Thus, the vertex and any segment define either a line, and thus admit an isolated transversal, or they define a plane. If none of the other edges intersect that plane in a unique point, the vertex and all open edges admit zero or infinitely many common transversals, a contradiction. Thus there exists an edge that intersects the plane in a unique point. Hence, the vertex and two open edges admit a unique transversal, and the minimal set contains no other element.

Suppose finally that the set only contains open edges. The characterization of the transversals to a set of line segments [9] shows that either two, three or four of these line segments admit at most two transversals, or that the set of common transversals to all the open line segments can be parameterized by an open set of parameters in \mathbb{R}^2 , \mathbb{R} or $\mathbb{R}/\pi\mathbb{Z}$. In the latter case, the edges admit no isolated transversal, a contradiction. Hence, the minimal set of edges consists of two, three or four edges. (Note that two or three edges may admit an isolated transversal if that transversal contains one or two of the edges.) \square

We can now prove the upper bound of Theorem 1.3.

PROPOSITION 3.3. *There are $O(n^2k^2)$ minimal sets of open edges and vertices, chosen from some polytopes, that admit an isolated transversal that is tangent to these polytopes.*

Proof. We bound the number of minimal sets depending of their type according to Lemma 3.2. First, there are $O(n^2)$ pairs of vertices, pairs of edges, and sets of one vertex and one edge. Hence, at most $O(n^2)$ such pairs admit an isolated transversal.

Consider a minimal set of one vertex and two open edges, chosen from some polytopes, that admit an isolated transversal that is tangent to these polytopes. The open edges do not contain the vertex because otherwise they admit no isolated transversal. Thus the vertex and each edge define a plane. For each of the $O(n^2)$ planes defined by a vertex and an open edge not containing it, there are $O(k)$ lines in that plane that are tangent to one of the polytopes at some point other than the vertex. Hence there are $O(n^2k)$ sets of one vertex and two edges, chosen from some polytopes, that admit an isolated transversal that is tangent to these polytopes.

It is straightforward to show that three open edges admit an isolated transversal only if the line containing one of the edges intersects the two other edges. Since any line intersects at most two open edges in any of the k polytopes, there are $O(nk^2)$ sets of three open edges that admit an isolated transversal.

Consider now the case of four edges, chosen from at most three polytopes, that admit an isolated transversal that is tangent to these polytopes. The two edges chosen from the same polytope belong to the same face, and the isolated transversal lies in the plane containing that face. Each of the two other open edges intersects that plane in one point, because otherwise the four open edges admit zero or infinitely many transversals. For each of the $O(n)$ planes containing a face of one of the polytopes, and each of the $O(n)$ edges intersecting that plane in exactly one point, there are at most $2k$ lines in that plane that contain this point and are tangent to one of the k polytopes at some other point. Hence there are $O(n^2k)$ sets of four open edges, chosen from at most three polytopes, that admit an isolated transversal that is tangent to these polytopes.

We finally bound the number of sets of four edges, no two chosen from the same polytope. By Lemma 3.1 and by summing over all n edges e of the polytopes, the number T of sets of four open edges, chosen from four polytopes, that admits an isolated transversal that is tangent to these four polytopes satisfies

$$T \leq n \sum_{j < l < m} C(n_j + n_l + n_m),$$

where C is some constant. Since each n_i , $1 \leq i \leq k$, appears $\binom{k-1}{2}$ times in the sum, it follows that

$$T \leq Cn \sum_{1 \leq i \leq k} n_i \binom{k-1}{2} = Cn^2 \binom{k-1}{2}$$

so T is in $O(n^2k^2)$ as claimed. \square

The above result implies the following upper bounds and in particular those of Theorem 1.1.

PROPOSITION 3.4. *There are $O(n^2k^2)$ connected components of maximal free line segments tangent to at least four of the polytopes. This bound also holds for connected components of possibly occluded lines tangent to at least four of the polytopes. Furthermore, the same bound holds for isolated such segments or lines.*

Proof. We prove the proposition for possibly occluded lines tangent to at least four of the polytopes; the proof is similar for maximal free line segments. By Proposition 3.3, there are $O(n^2k^2)$ minimal sets of open edges and vertices, chosen from some polytopes, that admit an isolated transversal that is tangent to these polytopes. The bound on the number of connected components thus follows from the fact that any connected component of lines tangent to four polytopes contains an isolated line. Indeed, any non-isolated line can be moved while keeping the same set of supports until (at the limit) the line intersects a new edge or vertex. During the motion, the line remains tangent to all four polytopes since it keeps the same supports (except at the limit); if the line has more than one degree of freedom, this can be repeated until the line becomes isolated. \square

We now prove the upper bound of Theorem 1.2. We start by two preliminary lemmas.

LEMMA 3.5. *Four possibly intersecting convex polygons in \mathbb{R}^2 admit at most a constant number of connected components of line transversals.*

Proof. Consider the usual geometric transform where a line in \mathbb{R}^2 with equation $y = ax + b$ is mapped to the point $(-a, b)$ in the dual space (see e.g. [24, §8.2.1]). The transversals to a convex polygon are mapped to a region bounded from above by a convex x -monotone curve and from below by a concave x -monotone curve; such a region is called stabbing region, and the curves are referred to as the upper and lower boundaries of the stabbing region. The transversals to four polygons are mapped to the intersection of four stabbing regions. There exists no transversal of a given slope if and only if the lower boundary of a stabbing region lies above the upper boundary of another stabbing region at that slope. Two such boundaries intersect in at most two points, and thus the transversals to four polygons form at most a constant number of connected components of transversals. \square

As in Section 2, let \mathbf{P} , \mathbf{Q} , \mathbf{R} , and \mathbf{S} be four polytopes in \mathbb{R}^3 , with p , q , r , and $s \geq 1$ edges, respectively, and let e be a closed edge of \mathbf{S} .

LEMMA 3.6. *There are $O(p + q + r)$ connected components of lines intersecting e and tangent to \mathbf{P} , \mathbf{Q} , \mathbf{R} and \mathbf{S} .*

Proof. As in the proof of Proposition 3.4, any connected component of lines intersecting e and tangent to \mathbf{P} , \mathbf{Q} , \mathbf{R} , and \mathbf{S} contains an isolated line. The Main Lemma thus yields that there are $O(p + q + r)$ connected components of lines intersecting e and tangent to \mathbf{P} , \mathbf{Q} , \mathbf{R} and \mathbf{S} except for the components that only contain isolated lines that lie in planes that contain e and are tangent to all four polytopes.

We show that there are at most a constant number of connected components of lines intersecting e and tangent to \mathbf{P} , \mathbf{Q} , \mathbf{R} and \mathbf{S} that lie in planes that contain e and are tangent to all four polytopes. There may be infinitely many such planes that intersect \mathbf{P} , \mathbf{Q} , \mathbf{R} and \mathbf{S} only on l_e but all the lines tangent to the four polytopes in all these planes belong to the same connected component. Besides these planes there are at most two planes containing e




FIG. 4.1. Lower bound examples for Lemmas 4.1 and 4.2.

and tangent to all four polytopes. In any such plane, the lines tangent to the four polytopes are the transversals to the four polygons that are the faces, edges, or vertices of \mathbf{P} , \mathbf{Q} , \mathbf{R} , and \mathbf{S} lying in the plane. Lemma 3.5 thus yields the result. \square

We can now prove the upper bound of Theorem 1.2.

PROPOSITION 3.7. *Given 3 polytopes with n edges in total and one polytope with m edges, there are $O(mn)$ connected components of lines tangent to the four polytopes.*

Proof. Let \mathbf{S} denote the polytope with m edges. First, if \mathbf{S} consists of a single point, it is straightforward to show that there are $O(n)$ connected components of lines tangent to the four polytopes. Otherwise, by summing over all the edges of \mathbf{S} , Proposition 3.6 yields that the number of connected components of lines tangent to the four polytopes is $O(mn)$. \square

4. Lower bounds. We provide in this section the lower-bound examples needed for Theorems 1.1, 1.2, and 1.3. The following proposition proves the lower bound of Theorem 1.2.

LEMMA 4.1. *There exist four disjoint polytopes of complexity n such that the number of common tangent lines is finite and $\Omega(n^2)$. There also exist two polytopes of complexity n and two polytopes of complexity m such that the number of common tangent lines is finite and $\Omega(mn)$.*

Proof. We consider four planar regular polygons P , Q , R , and S , each with n vertices, embedded in \mathbb{R}^3 . P is centered at the origin and parallel to the yz -plane, Q is obtained from P by a rotation of angle $\frac{\pi}{n}$ about the x -axis, and R and S are obtained from P and Q , respectively, by a translation of length 1 in the positive x -direction (see Figure 4.1). We transform the polygons P and Q into the polytopes \mathbf{P} and \mathbf{Q} by adding a vertex at coordinates $(\varepsilon, 0, 0)$. Similarly, we transform the polygons R and S into the polytopes \mathbf{R} and \mathbf{S} by adding a vertex at coordinates $(1 + \varepsilon, 0, 0)$.

For ε sufficiently small, the lines tangent to \mathbf{P} , \mathbf{Q} , \mathbf{R} and \mathbf{S} are the lines through a vertex of $P \cap Q$ and a vertex of $R \cap S$. Since $P \cap Q$ and $R \cap S$ have $2n$ vertices each, there are $4n^2$ tangent lines. Now, moving \mathbf{P} and \mathbf{S} by 2ε in the x direction ensures the disjointness of the polytopes while preserving the existence of the tangents if ε is small enough.

Replacing R and S in the above construction by regular polygons each with m vertices yields the $\Omega(mn)$ lower bound in the case of two polytopes of complexity n and two polytopes of complexity m . \square

We now prove the lower bounds of Theorems 1.1 and 1.3. The following proposition directly yields these bounds since the number of isolated tangents to any four of the polytopes

is less or equal to the number of sets of open edges and vertices in at most four polytopes that admit an isolated transversal that is tangent to these polytopes.

LEMMA 4.2. *There exist k disjoint polytopes of total complexity n such that the number of maximal free line segments tangent to four of them is finite and $\Omega(n^2k^2)$. Moreover these segments lie in pairwise distinct lines.*

Proof. The lower bound example is similar to the one with four polyhedra. For simplicity suppose that n and k are such that $\frac{n}{k}$ and $\frac{k}{4}$ are integers. We first take a $\frac{n}{k}$ -regular polygon A_1 in the plane $x = 0$. Next we consider a copy, B_0 , of A_1 scaled by a factor of $(1 + \varepsilon)$, and on each edge of B_0 we place $\frac{k}{4}$ points. Polygon B_i , $1 \leq i \leq \frac{k}{4}$, is constructed by taking the i^{th} point on each edge of B_0 . If ε is small enough, the intersection points of A_1 and B_i are outside the other polygons B_j for $1 \leq j \leq \frac{k}{4}$ and $i \neq j$. Now the A_i , for $2 \leq i \leq \frac{k}{4}$, are constructed as copies of A_1 scaled by a factor $1 + \frac{i}{k}\varepsilon$ (see Figure 4.1). For the moment, all polygons lie in plane $x = 0$. We now construct 4 families of $\frac{k}{4}$ polygons each:

- P_i is a copy of A_i translated by $i\varepsilon$ in the negative x direction
- Q_i is a copy of B_i translated by $i\varepsilon$ in the positive x direction
- R_i is a copy of B_i translated by $1 - i\varepsilon$ in the positive x direction
- S_i is a copy of A_i translated by $1 + i\varepsilon$ in the positive x direction

Any choice of four polygons, one in each family P_i , Q_j , R_l and S_m , reproduces the quadratic example of Lemma 4.1 with polygons of size $\frac{n}{k}$ and thus with total number of tangents larger than $(\frac{k}{4})^4 4 (\frac{n}{k})^2 = \frac{n^2k^2}{4}$. Furthermore the lines tangent to P_i , Q_j , R_l and S_m are only occluded by $P_{i'}$ and $S_{m'}$ for $i' > i$ and $m' > m$, that is, beyond the portion of the tangents containing the contact points. The k polygons can be transformed into k convex polyhedra as in Lemma 4.1. \square

5. Algorithm. Using the sweep-plane algorithm outlined in Section 2.1, we can compute in $O(n^2k^2 \log n)$ time all minimal sets of open edges and vertices, chosen from some of the polytopes, that admit a possibly occluded isolated transversal that is tangent to these polytopes. Now, for some of these lines, the segment joining the contact points with the polytopes is free. We can use standard, but complicated, ray-shooting data structures in order to determine which of these $O(n^2k^2)$ segments are free; this can be done in $O(\log^2 n)$ -time per query using $O((nk)^{2+\varepsilon})$ preprocessing time and storage [3].

We present in this section a solution that uses $O(n^2k^2 \log n)$ time and $O(nk^2)$ space. We adapt the algorithm outlined in Section 2.1 to directly compute the minimal sets of edges and vertices admitting an isolated line transversal that contains a free segment tangent to their respective polytopes. Our algorithm has better time and space complexities than the previously mentioned approach, and is readily implementable. Moreover, the space complexity drops to $O(nk)$ if no occlusion is taken into account. Precisely, we prove the following theorem which is more powerful, though more technical, than Theorem 1.4 and directly yields it.

THEOREM 5.1. *Given k polytopes in \mathbb{R}^3 with n edges in total, we can compute in $O(n^2k^2 \log n)$ time and $O(nk)$ space all the minimal sets of open edges and vertices, chosen from some of the polytopes, that admit an isolated, possibly occluded, line transversal tangent to these polytopes. We can also compute, in $O(n^2k^2 \log n)$ time and $O(nk^2)$ space, all the minimal sets of open edges and vertices that admit an isolated line transversal containing a maximal free segment that is tangent to these polytopes. Furthermore, the algorithm reports which of the transversals contains such a free line segment.*

For ease of presentation, we describe a simplified version of the algorithm in which we assume that the polytopes are in generic position; see Section 5.2 for details. Using the same techniques as in Section 2, it is straightforward though tedious to generalize the algorithm for arbitrary situations. We also only detail the algorithm for the case of minimal sets of

four edges, no two chosen from the same polytope; the other sets of at most four edges and vertices can be computed similarly.

5.1. Algorithm overview and data structures. The input to our algorithm is a set of possibly intersecting polytopes structured in a standard way so that classic incidence queries can be performed in constant time (see, for instance, [7, §9.1]).

We consider each polytope edge, e , in turn and sweep a plane around it between its two incident faces. During the sweep we create and maintain the following objects.

Combinatorial polygons. The sweep plane intersects each polytope in a (possibly empty) convex polygon whose vertices correspond to polytope edges. For each of these polygons, we maintain the set of vertices, each represented by its corresponding polytope edge, in a data structure that admits logarithmic-time vertex insertion, deletion and look-up operations, as well as ray-shooting queries. This can be done with a balanced binary search tree (see [21, §7.9.1]).

Combinatorial bitangents. The algorithm keeps track of the lines contained in the sweep plane and tangent to two polygons. The polytopes properly intersected by such a bitangent between its two supports are its *blockers*. A bitangent is represented by (pointers to) its two supports and a set of its blockers, ordered by polytope index, stored in a balanced binary search tree.

Polytope edges. We associate with each polytope edge a list of pointers to the combinatorial bitangents it supports in the current sweep plane.

Critical events. The sweep stops at critical events at which time combinatorial polygons and bitangents are updated. In addition to the V- and F-critical events defined in Section 2.3, we introduce the following two new types of events at which the set of blockers of some combinatorial bitangents may change. A *T-critical* event occurs whenever three bitangents, supported by a **PQR**-tuple, become aligned (see Figure 5.1b). An *I-critical* event occurs when the sweep plane contains a point of intersection between an edge and a face of two (distinct) polytopes (see Figure 5.2).

Each event is represented by a data structure providing pointers to the primitives that define it: a vertex for a V-event, a bitangent and a face for a F-event, three bitangents for a T-event, and a face and an edge for a I-event. In addition, for a T-event, we store a bit of information specifying which of the line transversals to l_e and the three support edges defines the T-event. Note that the critical value of each critical event can be computed in constant time from the information associated with the event; it thus does not need to be explicitly stored.

Finally, critical events are sorted in the order in which they appear during the sweep and stored in an *event queue* supporting insertion and deletion in logarithmic time.

5.2. Generic position assumption. Our generic position assumption is that *the ordered set of events does not change under any arbitrarily small perturbation of the input polytopes*. This assumption corresponds to (i) the events are generic, and (ii) no two events occur in the same sweep plane, except for F- and I-critical events induced by the same pair of edge and face. The genericity of the events is ensured by (but not characterized by) the following geometric conditions:

V-critical events: no vertex lies on a line containing another edge,

F-critical events: no two edges in two distinct polytopes are coplanar,

I-critical events: if an edge intersects a face of another polytope, it does so properly and not on a line containing another edge,

T-critical events: any four lines containing polytope edges admit zero or two transversals.




FIG. 5.1. (a) The sweep plane in which the combinatorial bitangent with support edges e_1 and e_2 is created. (b) The sweep plane at a T-critical event induced by the three bitangents with support edges in e_1 , e_2 , and e_3 . (c-d) A line ℓ that defines an F-critical event. (d) The F-event defined by ℓ occurs simultaneously with an I-critical event.

5.3. Initialization. For each new sweep, we initialize the event queue and construct the combinatorial polygons and combinatorial bitangents as follows.

Combinatorial polygons. Computing the combinatorial polygons in the initial sweep plane can easily be done in $O(n)$ time.

Combinatorial bitangents. The bitangent lines to two polygons P and Q in the initial sweep plane through a given vertex of P can be computed by a binary search on Q in $O(\log n)$ time. The blockers of a given bitangent can be found using one ray-shooting query per combinatorial polygon, for a total time of $O(k \log n)$. Altogether, the $O(nk)$ combinatorial bitangents can thus be computed in $O(nk^2 \log n)$ time.

Event queue. There are $O(n)$ V-critical events and $O(nk)$ I-critical events, since an edge intersects a polytope in at most two faces. The $O(nk)$ edge-face intersection points are computed and stored once before the beginning of the first sweep; this computation can be done by using brute force in $O(n^2)$ time, and with $O(nk)$ space, since it is done once for all the sweeps. For each sweep, all the V- and I-critical events can then be inserted in $O(nk \log n)$ time. For each of the $O(nk)$ combinatorial bitangents, we also insert F- and T-critical events in $O(k \log n)$ time as explained in Section 5.4 (Lemma 5.2). In total, initializing the event queue takes $O(nk^2 \log n)$ time per sweep.

Thus, initializing all the combinatorial polygons, bitangents, and the event queue can be done in $O(nk^2 \log n)$ time per sweep plus $O(n^2)$ time overhead for a total of $O(n^2 k^2 \log n)$ time as announced in Theorem 5.1.

5.4. Updating the event queue. Every time a new combinatorial bitangent is created, we compute and insert into the queue new F- and T-events as described below. Let e_1 and e_2 denote the two support edges of a new combinatorial bitangent. Let Π_{t_0} denote the critical plane at which the new combinatorial bitangent is created.

New T-critical events. See Figure 5.1a-b. Consider all the bitangents having e_1 as support edge and compute the set of support edges (distinct from e_1 and e_2) of all these bitangents.

Compute the intersection of this set with the similar set for e_2 ; this can be done in $O(k \log k)$ time by ordering the edges by their indices. For each edge e_3 in that set, insert a T-event for each line transversal to l_e, e_1, e_2 , and e_3 if the transversal is tangent to the three polytopes containing e_1, e_2 , and e_3 ; this test can be done in constant time. Each of the at most k insertions into the event queue takes $O(\log n)$. Thus computing and inserting the new T-critical events takes $O(k \log n)$ time per new bitangent.

New F-critical events. Consider in turn each of the four faces incident to one of the two support edges. Let e_1 and f denote the considered edge and face. We compute a candidate F-event, in constant time, as follows. Compute the line ℓ (if any) that lies in the plane Ψ containing f and goes through l_e and e_2 (see Figure 5.1c). If ℓ is tangent to the polytope containing e_2 , ℓ defines an F-event. We reject this F-event if ℓ does not intersect e_1 (in such a case, the edge e_1 does not intersect the sweep plane at the F-event and thus the combinatorial bitangent to e_1 and e_2 would have been deleted at some V-event before the F-event). We also discard this F-event if it occurs at the critical value t_0 where the (considered) bitangent is created (that is Π_{t_0} contains ℓ); we discard such F-events because when a bitangent is created at an F-event, we do not re-insert the same F-event into the queue. We thus retain at most four F-events, at most one for each of the four faces incident to one of the two support edges. If no F-event is retained, the bitangent will be deleted at a V-critical event and no new F-critical event is created. If more than one F-event is retained, we need only keep the first one, since, as we shall see in Section 5.5.2, the combinatorial bitangent will be deleted at the first of these events.

Again, let f denote the face incident to edge e_1 that induces that F-critical event. If the other support edge, e_2 , intersects face f (see Figure 5.1d), then this event will be treated as an I-critical event and again we create no new F-event. Otherwise, we insert the F-event into the queue in $O(\log n)$ time. We thus get the following lemma.

LEMMA 5.2. *Each time a combinatorial bitangent is created, the event queue can be updated in $O(k \log n)$ time.*


5.5. Processing events.

5.5.1. V-critical events. Let v denote the vertex that induces a V-critical event. As the sweep plane reaches v , all edges incident to v start or cease to be swept; we call the former *starting* edges and the latter *terminating* edges. Let \mathbf{Q} denote the polytope to which v belongs and let Π_{t_0} be the sweep plane containing v . When processing a V-event, we perform the following operations.

Create and delete combinatorial bitangents. Suppose first that the critical plane through v properly intersects \mathbf{Q} . Consider in turn each combinatorial bitangent supported by a terminating edge, e_t , incident to v and let h denote the other support edge of this bitangent. We check all starting edges incident to v to find the edge e_s such that the line in $\Pi_{t_0+\varepsilon}$ through e_s and h is tangent to \mathbf{Q} for $\varepsilon > 0$ arbitrarily small. We create a new combinatorial bitangent and delete the old one; in fact, we simply replace e_t by e_s in the combinatorial bitangent, create a pointer from edge e_s to the bitangent, and update the event queue. After handling the last bitangent supported by edge e_t , delete all the pointers from e_t to the bitangents.

The critical plane through v contains $O(k)$ bitangents through v , thus, by continuity, at most $O(k)$ combinatorial bitangents are deleted and created. Each deletion and creation takes linear time in the degree of v plus $O(k \log n)$ time for updating the event queue (Lemma 5.2). Hence, since the sum of the degrees of the vertices is $O(n)$, this step takes $O(nk^2 \log n)$ time in total for all non-extremal V-events.

Suppose now that the critical plane through v is tangent to \mathbf{Q} and that all edges incident to v are starting. For each edge not incident to v , we can decide in constant time whether

FIG. 5.2. *I-critical event.*

it supports a bitangent through v in the critical plane through v . If so, we check, for each edge incident to v , if the line in plane $\Pi_{t_0+\epsilon}$ that goes through these two edges is tangent to \mathbf{Q} for $\epsilon > 0$ arbitrarily small. If so, we create a new combinatorial bitangent. By continuity, $O(k)$ bitangents are created in total time $O(n + kd)$ where d is the degree of v . For each of these newly created bitangents, we compute its set of blockers in (brute force) $O(n)$ time and update the event queue in $O(k \log n)$ time (Lemma 5.2). This takes $O(nk \log n)$ time per event, hence $O(nk^2 \log n)$ time per sweep since there are at most two sweep planes tangent to any polytope.

Finally, if all edges incident to v are terminating, we delete all the $O(k)$ bitangents supported by these edges; for each bitangent, deleting its blockers and the pointer from the edge not incident to v can be done in $O(k)$ time. Hence, this takes $O(k^2)$ time per critical event and $O(k^3)$ time per sweep.

Update the combinatorial polygon associated with \mathbf{Q} . This takes $O(\log n)$ time per polytope edge incident to v , thus $O(n \log n)$ time in total for all V-events.

Hence, processing all V-events takes $O(nk^2 \log n)$ time per sweep.

5.5.2. F-critical events. We process an F-critical event as follows. Let b and f denote the bitangent and face associated with the event. Let e_1 and e_2 denote the two support edges of b such that e_1 is the edge that belong to f (see Figure 5.1c-d). By construction of F-events (see Section 5.4), e_2 does not intersect face f (see Figure 5.1c), thus the bitangent b is deleted and a new combinatorial bitangent is created.

Bitangent b is removed from the lists of bitangents supported by e_1 and e_2 in $O(k)$ time. The support edges of the new bitangent are e_2 and the edge $e'_1 \neq e_1$ of f that is intersected by the line in the plane Ψ (containing f) through l_e and e_2 (see Figure 5.1c). This edge e'_1 is also one of the two edges adjacent to e_1 in its combinatorial polygon. Edge e'_1 can thus be computed in $O(\log n)$ time. As usual, the new bitangent is added to the lists of bitangents supported by e'_1 and e_2 . We then compute all the blockers of this new bitangent by performing one ray-shooting query per combinatorial polygon, for a total time of $O(k \log n)$. We finally update the event queue in $O(k \log n)$ time (Lemma 5.2).

There are $O(k)$ F-events associated to each polytope face, thus $O(nk)$ F-events per sweep. Hence, the total time complexity for processing all F-events is $O(nk^2 \log n)$ per sweep.

5.5.3. I-critical events. An I-event is associated with a face f of some polytope \mathbf{P} and an edge e_1 of some other polytope \mathbf{Q} . Let p denote the point of intersection between f and e_1 . The sweep plane, Π_{t_0} , that contains p intersects the two polytopes \mathbf{P} and \mathbf{Q} in two polygons P_{t_0} and Q_{t_0} . See Figure 5.2. Point p lies on an edge of P_{t_0} ; the two endpoints of this edge are

the intersection of two edges of \mathbf{P} , say e_0 and e_2 . These two polytope edges can be computed in $O(\log n)$ time using the combinatorial polygon associated with \mathbf{P} .

Create or delete combinatorial bitangents. If the two polygons P_{i_0} and Q_{i_0} are tangent at p (see Figure 5.2a), the two combinatorial bitangents whose pairs of support edges are (e_0, e_1) and (e_1, e_2) are either created or deleted at the I-event. If these bitangents appear in the list of bitangents having edge e_1 as support, we remove them from the list and delete them; this can be done brute force in $O(k)$ time. Otherwise we create these two combinatorial bitangents. We compute their set of occluders in $O(k \log n)$ time by intersecting the bitangents with all the polytopes using their associated combinatorial polygons. Finally, we update the event queue in $O(k \log n)$ time.

Update sets of blockers. Consider now each of the $O(k)$ bitangents having e_1 as a support edge except for the two bitangents that might have just been created. We update its set of blockers as follows. First, note that only polytope \mathbf{P} may have to be added to, or removed from, the set of blockers. Two situations occur: either the geometric bitangent segment joining the two support edges in Π_{i_0} properly intersects polygon P_{i_0} , or not. In the first case (e.g., segment pq in Figure 5.2), polytope \mathbf{P} was and remains a blocker of the bitangent. In the second case (e.g., segment pr in Figure 5.2), \mathbf{P} has to be either removed from, or added to, the set of blockers. This can be done in $O(k \log k)$ time by searching for \mathbf{P} in the set (recall that polytopes are ordered by their index in a binary search tree).

Processing an I-event thus takes $O(k \log n)$ time. Since any polytope edge intersects any other polytope in at most two points, there are $O(nk)$ I-events which can be processed in $O(nk^2 \log n)$ time in total per sweep.

5.5.4. T-critical events. Suppose that on the line transversal to e_1, e_2, e_3 and l_e (the one associated to the T-event) edges e_1, e_2, e_3 are met in that order at points p_1, p_2, p_3 . Let \mathbf{Q}_i be the polytope containing e_i , $1 \leq i \leq 3$.

Update sets of blockers. Update the occluder set for the bitangent with support edges e_1 and e_3 by either removing \mathbf{Q}_2 (if it appears in the set) or adding \mathbf{Q}_2 (if it does not appear in the set); this can be done in $O(\log n)$ time.

Output. First determine if the segment $p_1 p_3$ is unoccluded by checking if the set of blockers of the bitangent with support edges e_1 and e_3 is empty or reduced to \mathbf{Q}_2 . If so and if the segment intersects the reference edge e , then it is a free segment transversal to the four edges e, e_1, e_2, e_3 . In order to report each such transversal exactly once, we report it only if the reference edge e is smaller than e_2 for some global ordering of all edges. This can be done in constant time.

There are $O(nk^2)$ T-critical events per sweep (see the proof of Proposition 3.3), thus all the T-events can be processed in $O(nk^2 \log n)$ time per sweep.

5.6. Complexity. Note first that we assume a model of computation in which bounded-degree algebraic polynomials may be evaluated in constant time. See [18] for a detailed description of the predicates concerning line transversals that are used in this algorithm.

In this model of computation, we have described a $\Theta(n^2 k^2 \log n)$ -time algorithm for computing all the minimal sets of edges, no two chosen from the same polytope, that admit an isolated line transversal containing a free segment that is tangent to all these polytopes. As mentioned earlier, the sweep-plane algorithm can be easily modified to report all types of minimal support sets.

The space used by the algorithm is $\Theta(nk^2)$ in the worst case. To see this, first notice that storing the combinatorial polygons and the V-, F- and I-critical events uses $O(nk)$ space.

There are also $O(nk)$ combinatorial bitangents in any sweep plane. Storing the combinatorial bitangents thus requires $\Theta(nk^2)$ space since, in the worst case, $\Theta(nk)$ of them may be intersected by $\Theta(k)$ polytopes. Furthermore, there may be $\Theta(nk^2)$ T-events in the queue since each of the $\Theta(nk)$ bitangents may share a support with $\Theta(k)$ other bitangents. This yields the bounds of Theorem 1.4 for computing minimal free segments.

Notice that, with a slight modification to the algorithm, and no increase in the time complexity, we can reduce the storage requirement of the T-events to $O(nk)$. To do this we maintain the bitangents sorted by polar angle around each vertex of the combinatorial polygons, which can easily be done since the cyclic ordering changes only at T-critical events or when a bitangent is created or deleted. Since two bitangents become aligned only when they are neighbors in this cyclic ordering, we only need to maintain the T-events for pairs of consecutive bitangents and there can only be $O(nk)$ of these at any one time.

Finally, the bounds of Theorem 1.4 that concern the computation of potentially occluded isolated lines tangent to polytopes are obtained by noticing that we need not maintain the sets of blockers of the bitangents which reduces the space requirements for the combinatorial bitangents to $O(nk)$.

6. Conclusion. We have presented a tight bound on the *number* of (connected components of) lines and maximal free line segments that are tangent to at least four among k possibly intersecting polytopes in arbitrary position. A problem that we leave open is to prove that the same bound holds for the *combinatorial complexity* of the set of all maximal free line segments among k polytopes.

We have also shown how to compute in near-optimal worst-case time all the *minimal* free line segments that are isolated transversals to their set of supports and tangent to the corresponding polytopes. We believe that our algorithm can also be made to report all connected sets of minimal free segments that are transversal to the same set of edges. A problem that we have not solved, however, is to compute in the same time and space complexities, respectively, the polytopes supporting the endpoints of the corresponding *maximal* free line segments.

REFERENCES

- [1] P. K. AGARWAL, *On stabbing lines for convex polyhedra in 3D*, Computational Geometry: Theory and Applications, 4:4 (1994), pp. 177–189.
- [2] P. K. AGARWAL, B. ARONOV, V. KOLTUN, AND M. SHARIR, *Lines avoiding unit balls in three dimensions*, Discrete and Computational Geometry, 34:2 (2005), pp.231–250.
- [3] P. K. AGARWAL AND M. SHARIR, *Ray shooting among convex polytopes in 3D*, SIAM Journal on Computing, 25:1 (1996), pp. 100–116.
- [4] M. DE BERG, H. EVERETT, AND L.J. GUIBAS, *The union of moving polygonal pseudodiscs – Combinatorial bounds and applications*, Computational Geometry: Theory and Applications, 11 (1998), pp. 69–82.
- [5] M. DE BERG, D. HALPERIN, M. OVERMARS, AND M. VAN KREVELD, *Sparse arrangements and the number of views of polyhedral scenes*, Internat. J. Comput. Geom. Appl. 7:3 (1997), pp. 175–195.
- [6] M. BERN, D.P. DOBKIN, D. EPPSTEIN, AND R. GROSSMAN, *Visibility with a moving point of view*, Algorithmica 11 (1994), pp. 360–378.
- [7] J. D. BOISSONNAT AND M. YVINEC, *Algorithmic Geometry*, Cambridge University Press, 1998.
- [8] H. BRÖNNIMANN, O. DEVILLERS, V. DUJMOVIĆ, H. EVERETT, M. GLISSE, X. GOAOC, S. LAZARD, H.-S. NA, AND S. WHITESIDES, *On the Number of Lines Tangent to Four Convex Polyhedra*, in Proceedings of the Canadian Conference on Computational Geometry, Lethbridge, Canada, 2002, pp. 113–117.
- [9] H. BRÖNNIMANN, H. EVERETT, S. LAZARD, F. SOTTILE, AND S. WHITESIDES, *Transversals to line segments in three-dimensional space*, Discrete and Computational Geometry, 34:3 (2005), pp. 381–390.
- [10] F. S. CHO AND D. FORSYTH, *Interactive ray tracing with the visibility complex*, Computers and Graphics, Special Issue on Visibility - Techniques and Applications, 23:5 (1999), pp. 703–717.
- [11] R. COLE AND M. SHARIR, *Visibility problems for polyhedral terrains*, Journal of Symbolic Computation, 7:1 (1989), pp. 11–30.

- [12] O. DEVILLERS, V. DUJMOVIĆ, H. EVERETT, X. GOAOC, S. LAZARD, H.-S. NA, AND S. PETITJEAN, *The expected number of 3D visibility events is linear*, SIAM Journal on Computing, 32:6 (2003), pp. 1586–1620.
- [13] D. P. DOBKIN AND D. G. KIRKPATRICK, *Fast detection of polyhedral intersection*, Theoretical Computer Science, 27:3(1983), pp. 241–253.
- [14] F. DURAND, *A multidisciplinary survey of visibility*, ACM Siggraph course notes, Visibility, Problems, Techniques, and Applications, 2000.
- [15] F. DURAND, G. DRETTAKIS, AND C. PUECH, *The visibility skeleton: a powerful and efficient multi-purpose global visibility tool*, in Comput. Graphics Proc. Ann. Conf. Series (SIGGRAPH) 31 (1997), pp. 89–100.
- [16] F. DURAND, G. DRETTAKIS, AND C. PUECH, *The 3D visibility complex*, ACM Transactions on Graphics, 21:2 (2002), pp. 176–206.
- [17] A. EFRAT, L. GUIBAS, O. HALL-HOLT, AND L. ZHANG, *On Incremental Rendering of Silhouette Maps of a Polyhedral Scene*, in Proceedings of the 11th ACM-SIAM Symposium on Discrete Algorithms, 2000, pp. 910–917.
- [18] H. EVERETT, S. LAZARD, W. LENHART, J. REDBURN, AND L. ZHANG, *Predicates for line transversals in 3D*, in Proceedings of the 18th Canadian Conference on Computational Geometry, 2006, pp. 43–46.
- [19] X. GOAOC, *Structures de visibilité globales : taille, calcul et dégénérescences*, PhD Thesis, Université Nancy 2, Nancy, France, May 2004.
- [20] D. HALPERIN AND M. SHARIR, *New Bounds for Lower Envelopes in Three Dimensions, with Applications to Visibility in Terrains*, Discrete and Computational Geometry 12 (1994), pp. 313–326.
- [21] J. O’ROURKE, *Computational Geometry in C*, 2nd edition, Cambridge University Press, 1998.
- [22] M. PELLEGRINI, *On lines missing polyhedral sets in 3-space*, Discrete and Computational Geometry 12 (1994), pp. 203–221.
- [23] M. POCCHIOLA AND G. VEGTER, *The visibility complex*, Internat. J. Comput. Geom. Appl., 6:3 (1996), pp. 279–308.
- [24] M. SHARIR AND P. K. AGARWAL, *Davenport-Schinzle sequences and their geometric applications*, Cambridge university press, 1995.
- [25] J. STOLFI, *Oriented Projective Geometry: A Framework for Geometric Computations*, Academic Press, New York, NY, 1991.