On the Number of Lines Tangent to Four Convex Polyhedra*

H. Brönnimann† O. Devillers‖ V. Dujmović‡ H. Everett§ M. Glisse¶
X. Goaoc§ S. Lazard§ H.-S. Na§ S. Whitesides‡

May 3, 2002

Abstract

We prove that, under a certain general position assumption, the number of lines tangent to four bounded disjoint convex polyhedra of total complexity n in $\mathbb{R}^3$ is $O(n^2)$. Under the same assumption, we show that a set of k bounded disjoint convex polyhedra has at most $O(n^2k^2)$ lines that are tangent to four of these polyhedra (possibly occluded by some remaining ones).

1 Introduction

The running time of many solutions to basic visibility problems arising in graphics applications depends on the number of visibility events in a scene. A visibility event is a combinatorial change in the view of a moving observer; such an event occurs when a viewing direction becomes tangent to some objects. For convex objects in $\mathbb{R}^3$, a viewing direction typically cannot be tangent to more than four of the objects. In a scene consisting of n triangles, the number of lines tangent to four of them is $O(n^4)$, and it is not hard to see that $\Omega(n^4)$ is also a lower bound.

Although graphics scenes are ultimately treated in hardware simply as sets of triangles, they usually consist of highly structured objects such as polyhedra. This additional structure may be exploited to increase the efficiency of solutions to visibility problems.

This paper studies the number of visibility events in a scene of convex polyhedra in $\mathbb{R}^3$ as a function of both the complexity of the scene and of the number of polyhedra. We assume that the polyhedra are bounded, convex, disjoint, and satisfy a certain general position property. The visibility events we count take the form of lines tangent to four of the polyhedra (see Figure 1).

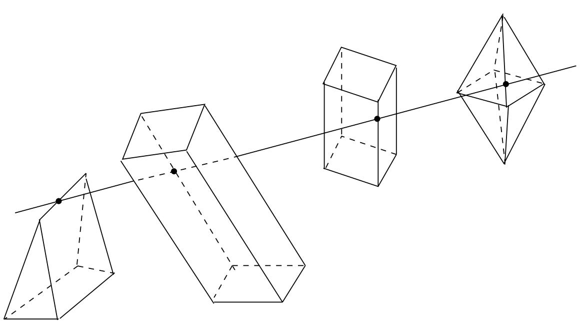

Figure 1: A line tangent to four convex polyhedra.

In order to define our notion of general position for polyhedra, we first define this notion for sets of lines. The following well-known fact is essential here and throughout the paper.

Fact: Four lines in $\mathbb{R}^3$ admit either 0, 1, 2, or an infinite number of line transversals (see [1]).

We say that four lines are in *general position* if they admit at most 2 transversals. See Appendix A for a characterization of this general position assumption.

*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 INRIA, NSERC, FCAR and NSF.

†Polytechnic University, hbr@poly.edu

‡McGill University, {vida,sue}@mcgill.ca

§INRIA Lorraine, {everett,goaoc,lazard,na}@loria.fr

¶École Normale Supérieure, marc.glisse@ens.fr

‖INRIA Sophia-Antipolis, Olivier.Devillers@inria.fr

We say that a set of convex polyhedra is in *general position* if every subset of four edges, no two of which belong to the same polyhedron, are supported by four lines in general position.

Our main result is an upper bound on the number of lines tangent to four convex polyhedra in $\mathbb{R}^3$.

Theorem 1 *Given four bounded, disjoint convex polyhedra in general position in $\mathbb{R}^3$ with n edges in total, the number of lines tangent to all four polyhedra is $O(n^2)$. If one of the polyhedra has constant size, then this number reduces to $O(n)$.*

When the number k of polyhedra is greater than 4, a direct application of the above theorem yields an upper bound of $O(n^2k^4)$, which can be improved as stated in the following theorem.

Theorem 2 *Given k bounded, disjoint convex polyhedra having a total of n edges and lying in general position in $\mathbb{R}^3$, the number of lines tangent to four of the polyhedra (possibly intersecting some remaining ones) is $O(n^2k^2)$.*

2 Proof of Theorem 1

The following lemma directly yields Theorem 1.

Lemma 3 *Given four bounded, disjoint convex polyhedra in general position in $\mathbb{R}^3$, the number of lines intersecting a given edge e of one of the polyhedra and tangent to the three other polyhedra, P, Q, and R, having p, q, and r edges, respectively, is $O(p+q+r)$.*

Proof: The proof proceeds as follows. We rotate a sweep plane about the line l supporting e. We show that some critical events occur during the sweep, that between two such events there are at most a constant number of tangent lines, and that the number of such events is $O(p+q+r)$.

The sweep plane intersects all three polyhedra P, Q and R throughout some closed, maximum interval of rotation about l. Let the sweep plane Π_t be parameterized by $t \in [0,1]$, where $t=0$ and $t=1$ correspond to the endpoints of this rotation interval[1]. Plane Π_0 (resp. Π_1) intersects P, Q, and R in three possibly degenerate convex polygons. At most one of these is precisely an edge or a face of one of the polyhedra; otherwise, one could choose three coplanar edges from distinct polyhedra (edge e and two others), leading to a contradiction of the general position assumption. Thus a line in Π_0 (resp. Π_1) that is tangent to P, Q and R has to be tangent to two polygons, so Π_0 (resp. Π_1) contains at most four lines tangent to P, Q and R.

Now consider the planes Π_t with $t \in (0,1)$ in the open rotation interval. Each such plane Π_t intersects the three polyhedra in three non-degenerate, disjoint convex polygons, $P(t)$, $Q(t)$ and $R(t)$ (see Figure 2). Each tangent to P, Q, R that intersects e has to be contained in some plane Π_t, and must be tangent to all three polygons $P(t)$, $Q(t)$ and $R(t)$ in that plane. The number of such tangents is thus equal to the number of times a bitangent to the polygons $P(t)$ and $Q(t)$ coincides with a bitangent to the polygons $Q(t)$ and $R(t)$, for all $t \in (0,1)$. (A bitangent is a line tangent to two polygons; its contact points are called support vertices.)

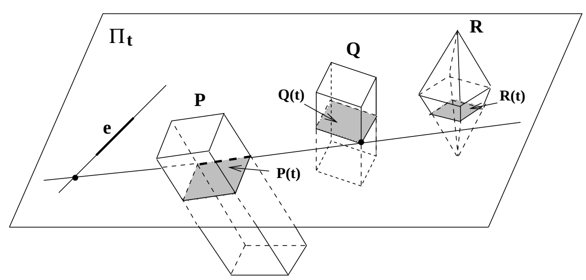

Figure 2: Plane Π_t contains edge e and intersects polyhedra P, Q and R in polygons $P(t)$, $Q(t)$, and $R(t)$, respectively. In this figure Π_t is in critical F-position.

We define a set of *critical positions* of the plane Π_t which contains, in particular, those positions for which the set of polyhedral edges containing support vertices of bitangents in Π_{t^-} is not equal

[1] When the line l intersects each of the polyhedra P, Q and R, the plane Π_t sweeps the whole space $\mathbb{R}^3$ and then Π_0 is equal to Π_1 which can be any plane containing l.

to the corresponding set of of edges for Π_{t^+}. Positions $\Pi_{t=0}$ and $\Pi_{t=1}$ are regarded as critical. Note that in the open interval between consecutive critical positions, the support vertices of the bitangents to polygons $P(t)$ and $Q(t)$, and of the bitangents to $Q(t)$ and $R(t)$, stay on the same polyhedron edges. It follows that between two consecutive critical positions Π_τ and $\Pi_{\tau'}$, the planes Π_t $(t \in [\tau, \tau'] \cap (0,1))^2$ contain at most 16 lines tangent to P, Q, R and intersecting e. Indeed, for each of the (at most four) bitangents between $P(t)$ and $Q(t)$, at most two bitangents between $Q(t)$ and $R(t)$ leave $Q(t)$ on the same side, and each of those (at most) eight pairs of bitangents corresponds to a quadruplet of edges which is met by at most two lines (by **Fact**).

We now characterize two types of critical positions, *critical V-positions* and *critical F-positions*.

Plane Π_t is at a *critical V-position* if it goes through a vertex v of P, Q or R. There can be at most $p+q+r$ such positions, since the number of vertices of a polyhedron is less than the number of its edges. As the plane rotates between two consecutive critical V-positions, the vertices of the polygons $P(t)$, $Q(t)$, and $R(t)$ stay on the same polyhedron edges.

The other kind of critical position of Π_t occurs when a support vertex of a bitangent in Π_t jumps from one vertex of a polygon in Π_t to another vertex of the polygon (see Figure 2). Then the bitangent (i) contains an edge of one of the polygons $P(t), Q(t)$ or $R(t)$, (ii) is tangent, in Π_t, to another of these polygons, and (iii) intersects (or is parallel to) the line l, since l also lies in Π_t. Thus the bitangent is a line m that (i) goes through one of the faces of the polyhedra P, Q or R, (ii) is tangent to another of these polyhedra and (iii) intersects (or is parallel to) the line l. The *critical F-positions* are the positions for which Π_t contains such a line m.

When a plane Ψ supporting a face of P, Q or R (say P) is tangent to another of these polyhedra (say Q), we potentially have an infinite number of critical F-positions corresponding to the lines in Ψ that intersect (or are parallel to) the line l and that pass through $P \cap \Psi$ and through $Q \cap \Psi$. In this case we consider as critical F-positions only the positions that correspond to the two extremal such lines. This is sufficient because it is still true that between any two consecutive critical positions, the support vertices of the bitangents stay on the same polyhedron edges, although some of these bitangents may have more than two support vertices. This does not change the bound of 16 on the number of tangent lines between two consecutive critical positions.

A plane Ψ supporting a face of P intersects Q and R in convex polygons, and either contains the line l or intersects it in at most one point.

In the first case, $\Psi \cap \Pi_t$ contains l; also, $\Psi \cap \Pi_t$ contains m. Thus either Ψ and Π_t are distinct, and thus $l = m$, or $\Psi = \Pi_t$, and $t = 0$ or 1 since only Π_0 and Π_1, among all Π_t, can contain a face of the polyhedra. The plane $\Psi = \Pi_0$ or Π_1 has already been considered. In the other case, $l = m$ is a bitangent in every Π_t, $t \in (0,1)$. However, the polyhedron edges containing the support vertices for bitangent $l = m$ do not change as t varies, so $l = m$ does not by itself give rise to any critical F-positions.

In the latter case, in plane Ψ, at most four lines are tangent to Q or R and intersect (or are parallel to) the line l. Thus a face of P generates at most four critical F-positions. A similar argument holds for the planes supporting the faces of Q and R. Thus the total number of critical F-positions, for a given edge e, is at most four times the number of faces of P, Q and R, that is, at most $4\,(p+q+r)$, since the number of faces of a polyhedron is at most the number of its edges.

Since there are at most $5\,(p+q+r)$ critical positions (at most $4\,(p+q+r)$ F-positions and at most $(p+q+r)$ V-positions), and Π_0 and Π_1 contain at most 8 tangent lines, there can be at most $5 \cdot 16\,(p+q+r) + 8$ lines tangent to P, Q, R and intersecting e. This proves the lemma. □

[2] As we will soon see, any such lines (tangent to P, Q, R and intersecting e) lying in a critical plane distinct from Π_0 and Π_1, are effectively counted with the lines arising from the adjacent open interval(s); the tangent lines lying in the planes Π_0 and Π_1 may not be tangent to all polygons $P(t)$, $Q(t)$ and $R(t)$, and thus may not be counted with the lines arising from the adjacent open interval(s).

3 Proof of Theorem 2

Given k polyhedra $P_1,\ldots,P_k$, each with n_i edges, we choose an edge e of P_i and count the number of tangents to P_j, P_ℓ, and P_m intersecting e. From Lemma 3, we know that this number is no more than $C(n_j + n_\ell + n_m)$, where C is some constant. Since there are n edges in the scene, the number of tangents T to four polyhedra is therefore at most

$$\begin{aligned} T &\leq n \sum_{j\neq\ell\neq m} C\,(n_j + n_\ell + n_m) \\ &\leq 3\,C\,n \sum_{j\neq\ell\neq m} n_j \\ &\leq 3\,C\,n \sum_{\ell\neq m}\sum_{j} n_j \\ &= 3\,C\,n \binom{k}{2} n, \end{aligned}$$

which is $O(n^2k^2)$ as claimed.

4 Discussion and open problems

We have presented bounds for the number of lines tangent to four polyhedra. The proofs are inspired by a method which was (to our knowledge) first used for sparse arrangements and the number of views of k convex polyhedra [3] (see also Schiffenbauer's survey [4]). We conclude with a few open problems.

There is a gap between the upper bound we prove for the number of lines tangent to four polyhedra, that is $O(n^2)$, and the best known lower bound which is $\Omega(n)$. It would be interesting to come up with better constructions for the lower bound. The bound can also be stated as $O(nm)$ where m is the minimum number of edges of one of the polyhedra. Are other bounds possible, e.g. in terms of the maximum number of edges on the silhouette of one of the polyhedra?

In addition, there is an easy example of $\Omega(n)$ convex polyhedra admitting $\Omega(n^4)$ lines tangent to four of them. We know of no non-trivial lower bound that takes into account the number of objects (in contrast with other quantities, e.g. the number of views of k convex polyhedra [2]).

Finally, the bound of Theorem 2 is only for lines tangent to four polyhedra among k. What we really would like to bound is the number of tangent lines such that the shortest line segment spanning all four points of contact does not intersect another polyhedron (non-occluded visibility events). Is it possible to get a better bound for this number?

Acknowledgements. We thank the other workshop participants for stimulating conversations and for creating a productive atmosphere.

References

[1] M. Hohmeyer and S. Teller, "Determining the lines through four lines," *Journal of Graphics Tools*, vol 4, no. 3, pp. 11-12, 1999.

[2] B. Aronov, H. Brönnimann, D. Halperin, and R. Schiffenbauer, "On the number of views of polyhedral scenes," in Discrete and Computational Geometry, Proc. of JCDCG 2000, J. Akiyama, M. Kano, and M. Urabe (Eds.), Springer-Verlag, Lecture Notes in Comp. Sci. series (LNCS), vol. 2098, 2001, pp. 81-90.

[3] M. de Berg, D. Halperin, M. Overmars, and M. van Kreveld, "Sparse arrangements and the number of views of polyhedral scenes," in *Internat. J. Comput. Geom. Appl.*, vol. 7, no. 3, 1997, pp. 175-195.

[4] R. Schiffenbauer, "A Survey of Aspect Graphs", Technical Report TR-CIS-2001-01, Dept. of Computer and Information Science, Polytechnic University, 2001.

A Characterization of non-general position configurations

In this appendix we prove the following elementary lemma in order to give the reader a better understanding of the impact that the general position assumption has on the sets of convex polyhedra that we allow.

Lemma 4 *Four lines are not in general position if and only if (i) they belong to the same family of generators of a hyperbolic paraboloid, (ii) two of the lines are coplanar, and the two other lines intersect in that plane, (iii) three of the lines are coplanar and the fourth line intersects that plane, or (iv) at least three of the four lines are concurrent. (If two of the lines coincide, then (i) or (ii) applies.)*

Proof: We consider four pairwise distinct lines l_1, l_2, l_3 and l_4. We first assume that they are pairwise distinct because, as we will see in the proof, three pairwise distinct lines have either 0 or infinitely many transversals, and so do two distinct lines.

We distinguish the cases where (1) l_1, l_2 and l_3 are concurrent, (2) l_1, l_2 and l_3 are coplanar, (3) l_1 and l_2 are parallel, (4) l_1 and l_2 intersect, and (5) l_1, l_2 and l_3 are pairwise skew.

Case 1. See Figure 3. There are infinitely many lines that go through the intersection point of l_1, l_2 and l_3, and meet the fourth line (case (iv)).

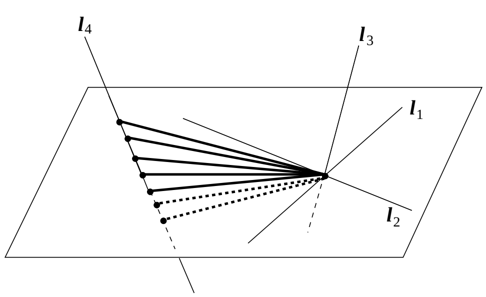

Figure 3: l_1, l_2 and l_3 are concurrent (case 1).

Case 2. See Figure 4. l_1, l_2 and l_3 are contained in some plane Π, and any line that meets the three of them is also in that plane. If l_4 meets

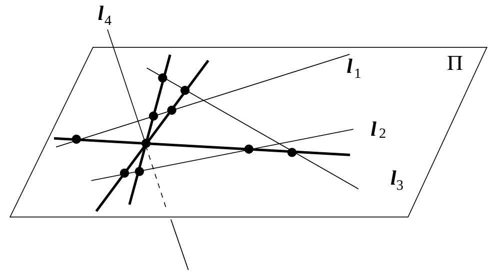

Figure 4: l_1, l_2 and l_3 are coplanar (case 2).

Π (case (iii)), then there is an infinite number of lines meeting the four lines, otherwise there is no such line.

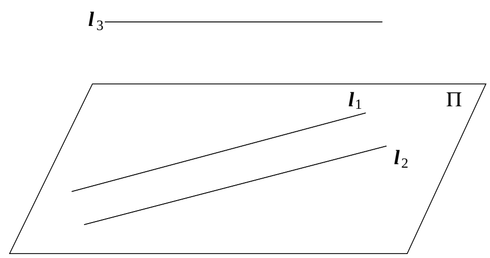

Figure 5: l_3 is parallel to Π (case 3).

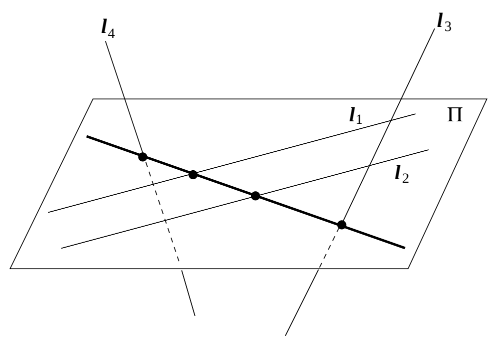

Figure 6: l_3 and l_4 meet Π at distinct points (case 3).

Case 3. Let Π denote the plane containing l_1 and l_2. All the lines that meet l_1 and l_2 both are contained in Π. Thus, if l_3 or l_4 is parallel to Π there is no line that pierces the four of them (see Figure 5); if l_3 and l_4 meet Π in distinct points, only the line joining these two points meets the four lines, provided that it is not parallel to l_1

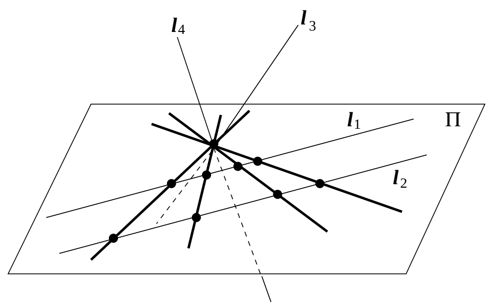

Figure 7: l_3 and l_4 meet Π at the same point (case 3).

and l_2 (see Figure 6); if l_3 and l_4 meet Π in the same point (case (ii)), then infinitely many lines go through that point and meet l_1 and l_2 (see Figure 7). Note that if l_3 or l_4 is contained in Π then we are in case (2).

Case 4. l_1 and l_2 intersect in some point A, let Π be the plane containing l_1 and l_2, and Ψ be the plane containing A and l_3. We need to distinguish between the case where l_3 is parallel to Π and the case where l_3 meets Π in some point (l_3 is not contained in Π because otherwise we would be in case (1) or (2)).

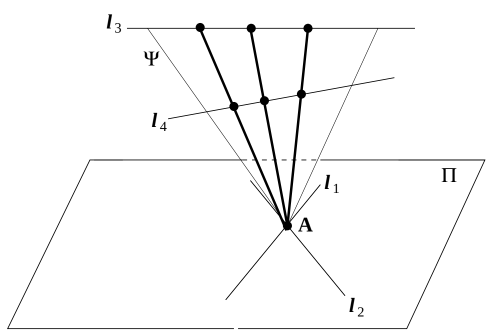

Figure 8: l_4 is contained in Ψ (case 4).

In the former case, the lines that meet l_1, l_2 and l_3 are the lines of Ψ passing through A. Thus, if l_4 is contained in Ψ (case (ii)), then there are infinitely many lines that intersect the four lines: all those that go through A and meet the two coplanar lines l_3 and l_4 (see Figure 8); otherwise, if l_4 is parallel to Ψ then no line meets the four lines; finally, if l_4 meets Ψ in a point B, which is distinct from A (otherwise we would be

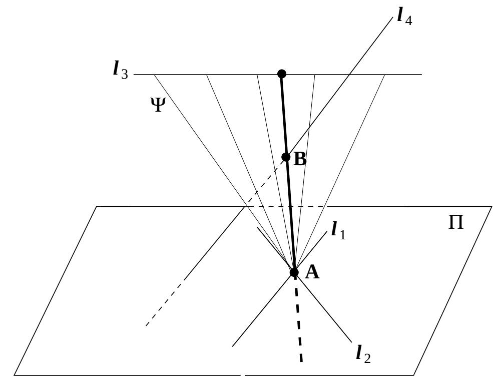

Figure 9: l_4 meets Ψ in a point B (case 4).

in case (1)), then only the line passing through A and B may meet the four lines, provided that it is not parallel to l_3 (see Figure 9).

In the latter case, let B be the intersection point of l_3 and Π. The lines that meet l_1, l_2 and l_3 are the lines in Π that go through B, and those in Ψ that go through A. If l_4 is contained in Π, l_1, l_2 and l_4 are coplanar and we are in case (1) or (2). If l_4 is contained in Ψ (case (ii)), then there are infinitely many lines that intersect the four lines (see Figure 10). If l_4 meets neither Π nor Ψ then no line simultaneously intersects the four lines. If l_4 passes through A we are in case (1). If l_4 passes through B (case (ii)), there are infinitely many lines in Π that meet the four lines (see Figure 11). If l_4 meets only Π or Ψ in a point distinct from A and B, exactly one line meets the four lines. If l_4 meets both Π and Ψ in two points distinct from A and B, exactly two lines meet the four lines (see Figure 12). If l_4 meets both Π and Ψ in the same point, which is distinct from A and B, then exactly one line meet the four lines.

Case 5. The three pairwise skew lines l_1, l_2 and l_3 lie on a degree-two ruled surface, a hyperbolic paraboloid (saddle) $\mathcal{Q}$ (see Figure 13). There are two families of lines that generate $\mathcal{Q}$, one that contains l_1, l_2 and l_3, and the other that consists of the lines that intersect each of l_1, l_2 and l_3. If the fourth line l_4 meets $\mathcal{Q}$ in at most two points, each of these points determines a line in the second family. Then there are at most two lines m_1 and m_2 that meet l_1, l_2, l_3 and l_4.

If l_4 lies on $\mathcal{Q}$, two cases can arise. If l_4 belongs to the same family as l_1, l_2 and l_3 (case (i)), then all the lines in the second family meet the four lines, and there is thus infinitely many lines that intersect the four lines. If l_4 belongs to the other family of generators of $\mathcal{Q}$, then only one line meets the four lines, namely l_4. □

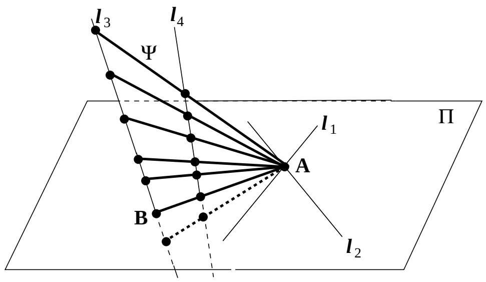

Figure 10: l_4 is contained in Ψ (case 5).

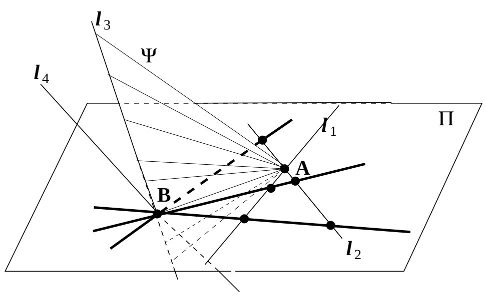

Figure 11: l_4 passes through B (case 5).

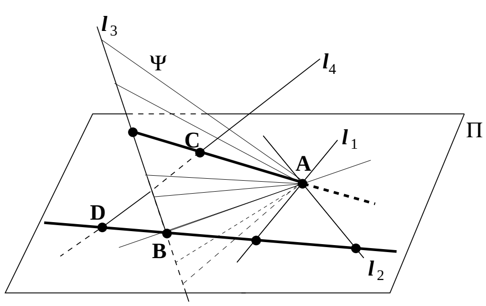

Figure 12: l_4 meets Π and Ψ in distinct points, different from A and B (case 5).

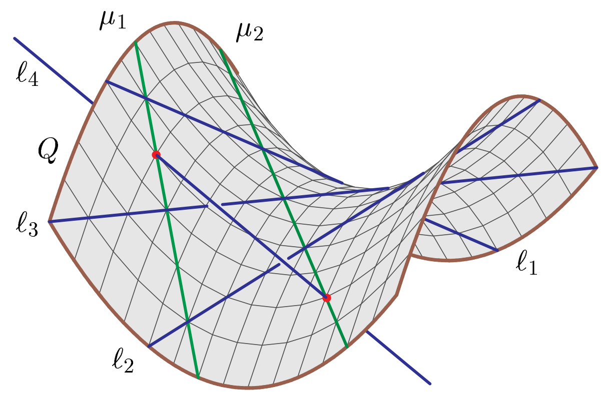

Figure 13: The two lines meeting four general lines in space (from F. Sottile, `http://www.maths.univ-rennes1.fr/~raag01/surveys/ERAG/S4/1.html`).