

Problèmes de Géométrie Algorithmique sur les Droites et les Quadriques en Trois Dimensions

MÉMOIRE

présenté et soutenu publiquement le 24/09/2007

pour l'obtention de

Habilitation de l'Université Nancy 2
(Spécialité Informatique)

par

Sylvain Lazard

Composition du jury

Président : Jean-Daniel Boissonnat, Directeur de recherche INRIA

Rapporteurs : Claude Puech, Professeur à l'université de Paris 11
Günter Rote, Professeur à l'université libre de Berlin

Examineurs : Said Benachour, Professeur à l'université de Nancy 2
Michel Pocchiola, Maître de conférence à l'École Normale Supérieure

Autres rapporteurs

Micha Sharir, Professeur à l'université de Tel Aviv

Table des matières

Table des figures	vii
Liste des tableaux	xi
I Introduction	1
1 Introduction	3
1.1 Context	3
1.2 Overview and methodology	4
2 3D visibility and lines in space	11
2.1 Introduction	11
2.2 Summary of contributions	13
3 Geometric computing with low-degree algebraic surfaces	29
3.1 Introduction	29
3.2 Summary of contributions	30
4 Conclusion	51
II Propriétés des droites et segments de \mathbb{R}^3 et problèmes de visibilité tridimensionnelle	53
5 Common tangents to spheres in \mathbb{R}^3	55
5.1 Introduction	55
5.2 Preliminaries	56
5.3 Affinely independent centers	58

5.4	Coplanar centers	58
5.5	Collinear centers	63
5.6	Conclusion	64
6	Transversals to line segments in three-dimensional space	67
6.1	Introduction	67
6.2	Our results	68
6.3	Proof of Theorem 6.1	69
6.4	Algorithmic considerations and conclusion	72
7	Lines tangent to four triangles in three-dimensional space	75
7.1	Introduction	75
7.2	A construction with 62 tangents	77
7.3	Upper bound for disjoint triangles in general position	79
7.4	Upper bounds on the number of components	81
7.5	Random triangles	81
8	Predicates for line transversals in 3D	85
8.1	Introduction	85
8.2	Computing lines through four lines	86
8.3	Predicates	89
8.4	Experiments	96
9	Lines and free line segments tangent to arbitrary three-dimensional convex polyhedra	99
9.1	Introduction	99
9.2	Main lemma	103
9.3	Upper bounds	113
9.4	Lower bounds	116
9.5	Algorithm	117
9.6	Conclusion	122
10	Towards an implementation of the 3D visibility skeleton	123
10.1	Introduction	123
10.2	The algorithm	123
10.3	Implementation issues and technical details	124

11 The expected number of 3D visibility events is linear	127
11.1 Introduction	127
11.2 Our model and results	129
11.3 The expected number of $T4$ -segments is at most linear	130
11.4 The expected number of $T4$ -segments is at least linear	141
11.5 The expected size of the visibility complex is linear	142
11.6 Worst-case lower bound	143
11.7 Generalizations	144
11.8 Conclusion	147
11.9 Appendix A. Volume of the intersection of a 3D hippodrome with a ball . .	148
11.10 Appendix B. Volume of K	150
11.11 Appendix C. Volume of the intersection of two spherical shells	155
12 An upper bound on the average size of silhouettes	159
12.1 Introduction	159
12.2 Definitions	160
12.3 Main results	161
12.4 Conclusion	170
13 Between umbra and penumbra	171
13.1 Introduction	171
13.2 Preliminaries	172
13.3 The penumbra boundary	173
13.4 Upper bounds	175
13.5 Lower bounds	178
13.6 Conclusion	184
III Géométrie algorithmique non linéaire sur les quadriques en trois dimensions	185
14 Near-optimal parameterization of the intersection of quadrics : I. The generic algorithm	187
14.1 Introduction	187

14.2	Notation and preliminaries	190
14.3	Levin's pencil method	191
14.4	Generic algorithm	193
14.5	Canonical forms and proof of Theorem 14.3	197
14.6	Optimality of the parameterizations	199
14.7	Near-optimality in the smooth quartic case	201
14.8	Examples	207
14.9	Conclusion	208
14.10	Appendix : The parameterizations of Table 14.3 are faithful	210
15 Near-optimal parameterization of the intersection of quadrics : II. A classification of pencils		213
15.1	Introduction	213
15.2	Classification of pencils of quadrics over the complexes	214
15.3	Classification of regular pencils of $\mathbb{P}^3(\mathbb{R})$ over the reals	220
15.4	Classification of singular pencils of $\mathbb{P}^3(\mathbb{R})$ over the reals	230
15.5	Classifying degenerate intersections	233
15.6	Examples	236
15.7	Conclusion	241
16 Near-optimal parameterization of the intersection of quadrics : III. Parameterizing singular intersections		243
16.1	Introduction	243
16.2	Parameterizing degenerate intersections : regular pencils	244
16.3	Parameterizing degenerate intersections : singular pencils	255
16.4	Examples	258
16.5	Conclusion	259
16.6	Appendix A : A primer on Galois theory	261
16.7	Appendix B : Rational canonical form for the case of four lines forming a skew quadrilateral	262
16.8	Appendix C : Examples in all cases	263
17 Near-optimal parameterization of the intersection of quadrics : IV. An efficient and exact implementation		269
17.1	Introduction	269

17.2 Preliminaries	271
17.3 Algorithm description	271
17.4 Height of output coefficients : smooth quartics	274
17.5 Height of output coefficients : singular intersections	277
17.6 Implementation	284
17.7 Experimental results	286
17.8 Examples	288
17.9 Conclusion	292
18 The Voronoi diagram of three lines	295
18.1 Introduction	295
18.2 Structure of the trisector	297
18.3 Properties of the Voronoi diagram	306
18.4 Topology of the Voronoi diagram	309
18.5 Configurations of three lines whose trisector contains a line	311
18.6 Algorithms	311
18.7 Conclusion	315
18.8. Appendix : Maple-sheet computations	315
Bibliographie	321

Table des figures

1.1	A chess set modeled entirely with quadrics.	6
2.1	Two examples of quadruples of spheres with infinitely many common tangents.	17
2.2	Two views of a hyperboloid of one sheet containing four line segments and their four connected components of transversals.	18
2.3	Four segments having three connected components of transversals.	19
2.4	Triangles with many common tangents	20
2.5	A line tangent at a vertex of each of k polytopes.	23
2.6	Four connected components of umbra cast on a plane by a segment light source in the presence of two triangles.	28
3.1	A gallery of intersections of quadrics.	32
3.2	Intersection of an elliptic cylinder and a hyperboloid of one sheet in a smooth quartic.	40
3.3	Voronoi diagram of three lines	48
5.1	Two examples of quadruples of spheres with infinitely many common tangents.	56
6.1	Two views of a hyperboloid of one sheet containing four line segments and their four connected components of transversals.	70
6.2	Four segments having three connected components of transversals.	71
6.3	Four coplanar segments having four connected components of transversals.	72
7.1	Hyperbolic paraboloid spanned by three lines.	76
7.2	Configuration in plane π_4	77
7.3	Conics in the planes π_i	78
7.4	Configuration in plane π_4	79
7.5	Stabbing and non-stabbing configurations.	80
7.6	Configurations with 26 and 25 contributing triples.	80
7.7	Lines stabbing two triangles.	81
7.8	Triangles with many common tangents.	82
8.1	(a) : Transversal ℓ intersects segment pq only if $(\ell \odot op) (\ell \odot oq) \leq 0$. (b-c) : An illustration for the proof of Lemma 8.10.	90
8.2	Planes P_1 and P_2 such that $P_1 < P_2$	94
9.1	A terrain of size n with $\Omega(n^4)$ maximal free line segments tangent in four points.	100

9.2	A line tangent at a vertex of each of k polytopes.	102
9.3	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	104
9.4	A bitangent to P_t and Q_t is tangent to P_t along an edge. The plane Π_t is F-critical.	104
9.5	Plane Π_t is F-critical.	106
9.6	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	107
9.7	Lines through x in Π_t and tangent to P_t and Q_t	108
9.8	Line m is tangent to \mathbf{P} along a face in plane $\Psi \neq \Pi_{t^*}$	110
9.9	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	111
9.10	Lower bound examples for Lemmas 9.27 and 9.28.	116
9.11	The sweep plane in which the combinatorial bitangent with support edges e_1 and e_2 is created.	119
9.12	I-critical event.	121
10.1	One position of the sweep plane and the view inside the sweep plane.	124
10.2	The 2D visibility skeleton for Figure 10.1(b).	124
11.1	$\mathcal{H}(p_i, p_l, 2)$ and $\mathcal{H}(t_i, t_l, 1)$ are shown shaded.	131
11.2	For the sketch of the proof of Lemma 11.14 ($v \in (0, 1)$).	136
11.3	For the sketch of the proof of Lemma 11.17.	138
11.4	For the proof of Lemma 11.19.	139
11.5	Quadratic view from the origin [DR01].	143
11.6	For the proof of Lemma 11.28.	145
11.7	Illustration for the case of polygons of bounded aspect ratio.	147
11.8	For the definition of t and C	148
11.9	For the definition of a, a', b, b'	149
11.10	The height from p_l to \mathcal{P} is greater than $ t_i t_l \cos \theta - 1$	150
11.11	The part E of $\mathcal{H}(p_i, p_l, 1)$ outside F	151
11.12	G , a part of ∂F enclosing $E \cap \partial F$	152
11.13	The region H and a plane Π	152
11.14	Section of H by a plane Π intersecting segment $p_i p_l$ at p	152
11.15	The length of the circular arcs ab	153
11.16	For the computing a bound on $\Delta\theta$	154
11.17	For the proof of Proposition 11.37.	155
11.18	For the proof of Proposition 11.37, case 1.	156
12.1	A worst-case linear silhouette of a polyhedron approximating a cylinder.	160
12.2	Three different notions of silhouettes.	160
12.3	Two half lanterns of Schwarz.	163
12.4	Length and dihedral angle of an edge and set of directions for which e is on the silhouette.	163
12.5	Construction of Ω	165
12.6	For the proofs of Lemma 12.4 and Corollary 12.5.	165

12.7	For the proof of Lemma 12.4.	167
12.8	Projection of ∂T_i , $f_n(\partial T_i)$ and D^-	167
12.9	For the proof of Lemma 12.6.	168
13.1	$\Omega(k^2)$ lower bound.	174
13.2	$\Omega(k\alpha(k))$ lower bound.	174
13.3	$\Omega(n\alpha(k))$ lower bound.	175
13.4	$\Omega(mk)$ lower bound.	176
13.5	Two triangles and a segment light source that cast 4 connected components of umbra on a plane.	178
13.6	The four connected components of umbra rendered with the ray tracer OpenRT.	179
13.7	Superset of the shadow arrangement on plane Π	179
13.8	Views in the sweep plane with bitangents that define the umbra.	180
13.9	The four connected components of umbra and the four lines used in the proof of Theorem 13.7.	181
13.10	$\Omega(n)$ lower bound.	181
13.11	$\Omega(nk^2)$ lower bound.	182
13.12	$\Omega(k^4)$ lower bound.	182
13.13	$\Omega(n^2k^3)$ lower bound.	183
13.14	$\Omega(nk^5)$ lower bound.	183
15.1	A gallery of intersections of quadrics.	222
17.1	Examples of intersections of quadrics.	272
17.2	Evolution of the height of $\Delta(u, v)$ when the intersection is a smooth quartic.	277
17.3	Observed height of the parameterization of the cubic in the cubic and tangent line case.	281
17.4	Computation time for the cubic and tangent line case.	281
17.5	Evolution of execution time in the smooth quartic case as a function of the size of the input for very large input sizes.	286
17.6	Evolution of execution time in the smooth quartic case as a function of the input size.	287
17.7	Computation time for 120 pairs of quadrics covering all intersection cases.	288
17.8	Three CSG models made entirely of quadrics.	289
17.9	Further examples of intersection.	289
18.1	Voronoi diagram of 3 lines ℓ_1 , ℓ_2 , and ℓ_3 in general position : (a) Voronoi 2D face of ℓ_1 and ℓ_2 , <i>i.e.</i> , set of points equidistant to ℓ_1 and ℓ_2 and closer to them than to ℓ_3 . (b) Orthogonal projection of a 2D face on a plane \mathcal{P} with coordinate system (X, Y) ; the plane's normal is parallel to the common perpendicular of ℓ_1 and ℓ_2 and the X and Y -axes are parallel to the two bisector lines (in \mathcal{P}) of the projection of ℓ_1 and ℓ_2 on \mathcal{P} . The 2D face is bounded by four branches of a non-singular quartic.	296
18.2	Three lines in general position.	298
18.3	The parallelepiped formed by ℓ_1 , ℓ_2 , and ℓ_3 and the associated frame (C, w_1, w_2, w_3) of positive orientation.	307

18.4	(a) Projection of the two-dimensional Voronoi cell V_{12} onto the XY -plane. (b) Vertical ordering of the sheets of the connected components of the two-dimensional Voronoi diagram cells above each region induced by the projection of the trisector and the silhouette curves of the bisectors ; the ordering over the small cell in the middle is $T_{13} < T_{13} < T_{23} < T_{23}$ (<i>i.e.</i> , a vertical line over that cell intersects twice T_{13} and twice T_{23} in that order).	308
18.5	Separating the two components of a two-dimensional Voronoi cell.	312

Liste des tableaux

2.1	Bounds on the complexity of the set of free lines or maximal free line segments among objects of total complexity n	24
2.2	Lower bounds on the number of connected components and upper bounds on the complexity of the umbra cast on a plane in the presence of k polytopes of total complexity $O(n)$	27
2.3	Bounds on the complexity of the union of umbra and penumbra cast on a plane by a m -gon light source in the presence of k polytopes of total complexity $O(n)$	27
3.1	Correspondence between quadric inertias and Euclidean types.	35
3.2	Parameterizations of canonical simple ruled quadrics [Lev76].	36
3.3	Parameterization of projective quadrics of inertia different from $(3, 1)$	38
3.4	Classification of pencils in the case where $\mathcal{D}(\lambda, \mu)$ does not identically vanish.	42
3.5	Classification of pencils in the case where $\mathcal{D}(\lambda, \mu)$ identically vanishes.	43
3.6	Ring of definition of the projective coordinates of the parameterization of each component of the intersection and optimality.	45
7.1	Four triangles with 62 common tangents.	77
7.2	Number of triangles with a given number of tangents, out of 5 000 000 randomly constructed triangles.	82
7.3	Four triangles with 40 common tangents.	82
8.1	Percentages of failure of the degree 168 and degree 3 predicates using double-precision floating-point interval-arithmetic, for ϵ varying from 10^{-12} to 10^{-2}	96
9.1	Published bounds on the complexity of the set of free lines or maximal free line segments among objects of total complexity n	103
11.1	Known bounds on the complexity of the set of lines, free lines or maximal free line segments tangent to 4 amongst n objects.	128
13.1	Lower bounds on the number of connected components and upper bounds on the complexity of the umbra cast on a plane by k polytopes of total complexity $O(n)$	172
14.1	Correspondence between quadric inertias and Euclidean types.	191
14.2	Parameterizations of canonical simple ruled quadrics [Lev76].	192
14.3	Parameterization of projective quadrics of inertia different from $(3, 1)$	193

15.1	Classification of pencils by Segre symbol in the case where $\mathcal{D}(\lambda, \mu)$ does not identically vanish.	217
15.2	Classification of pencils by Segre symbol in the case where $\mathcal{D}(\lambda, \mu) \equiv 0$	217
15.3	Classification of pencils by Segre symbol in the case where $\mathcal{D}(\lambda, \mu) \equiv 0$ and $\mathcal{D}_3(\lambda, \mu) \equiv 0$	217
15.4	Classification of pencils in the case where $\mathcal{D}(\lambda, \mu)$ does not identically vanish.	218
15.5	Classification of pencils in the case where $\mathcal{D}(\lambda, \mu)$ identically vanishes.	219
16.1	Ring of definition of the projective coordinates of the parameterization of each component of the intersection and optimality.	245
16.2	Exhaustive list of examples when the intersection is 0- or 1-dimensional over \mathbb{C}	263
17.1	Asymptotic heights of parameterizations in major cases, when the determinantal equation has a unique multiple root.	278
18.1	For the proof of the Main Lemma.	302
18.2	About the proof of the Main Lemma.	316
18.3	About the proof of the Main Lemma.	317
18.4	For the proof of Lemma 18.8.	318
18.5	For the proof of Lemma 18.10.	319

Première partie

Introduction

Chapitre 1

Introduction

1.1 Context

Even the youngest child playing with his building blocks appreciates the inherent difficulty of three-dimensional geometry. Unfortunately, or luckily, depending on how you lean, geometry and geometric models are ubiquitous in science and technology and geometric computing is present in many fields, such as computer-aided design, manufacturing, computer graphics, robotics, molecular biology, geographic information systems, astrophysics, computer vision, metrology and many others.

Since the 1970's, computational geometry has been dedicated to supplying a solid foundation for the study of geometric algorithms which are relevant to all these applications. Computational geometry has traditionally treated linear objects such as line segments and polygons in the plane, and point sets and polyhedra in higher-dimensional space, occasionally, and more recently, venturing into the world of non-linear curves such as circles and ellipses. The methodological experience and the know-how accumulated over the past thirty years is enormous.

For many applications, particularly in the fields of computer graphics and solid modeling, it is necessary to manipulate more general objects such as complex curves and surfaces given in either implicit or parametric form. Typically such objects are handled by approximating them by simple objects such as triangles. This approach is extremely important and it has been used in almost all of the usable software existing in industry today. It does, however, have some disadvantages. Using a tessellated form in place of its exact geometry may introduce spurious numerical errors (the famous gap between the wing and the body of the aircraft), not to mention that thousands if not hundreds of thousands of triangles could be needed to adequately represent the object. Moreover, the curved objects that need to be considered are not necessarily everyday three-dimensional objects, but also abstract mathematical objects that are not linear, that may live in high-dimensional space, and whose geometry we do not control. For example, the set of lines in 3D that are tangent to three polyhedra, an object at the core of visibility issues, span a ruled quadratic surface and the lines tangent to a sphere correspond, in projective five-dimensional space, to the intersection of two quadratic hypersurfaces.

Classical computational geometry has little to offer when it comes to handling the exact geometry of curved objects. Among the difficulties involved in computing with curved objects are, the extension (if not complete rework) of even the most basic data structures and algorithms

(such as Voronoi diagrams, sweep algorithms, and arrangements), the massive intrusion of algebraic issues and therefore the need for efficient algebraic software, the explosion of the number of degenerate cases (in geometric applications, they are often the norm, not the exception) and therefore the increased difficulty of writing robust software, and the need to incorporate arithmetic complexity, the assumption that operations take constant time (real RAM) not being appropriate in this context.

1.2 Overview and methodology

Over the last few years, I have focused my research activities on *effective geometric computing* dedicated to *non-linear geometric objects*. The objects that we consider are not necessarily everyday three-dimensional curved objects, but also abstract mathematical objects that are not linear and may be defined in higher dimension (such as the set of line transversals to three lines). My contributions on this subject are centered on, firstly, the structural, combinatorial, and algorithmic properties of geometric structures on lines and line segments in three dimensions in the context of three-dimensional visibility problems and, secondly, on the algorithmics for geometric curved objects and, in particular, quadrics.

Concerning the development of algorithms, I focus on algorithms that account for the *exact geometry* of the objects, in particular when they are curved. Furthermore, *effectiveness* is a key issue. Effective algorithms should be *robust* and *efficient*. By robust we mean algorithms that do not crash on degenerate inputs and always output topologically consistent data. By efficient we mean algorithms that run reasonably quickly on realistic data where performance is ascertained both experimentally and theoretically.

Meeting these computational objectives requires mathematical tools that are both *geometric* and *algebraic*. In particular, we need further knowledge of the basic geometry of lines and surfaces in a variety of spaces and dimensions as well as to adapt sophisticated algebraic methods, often computationally prohibitive in the most general setting, for use in solving seemingly simple geometric problems.

I present in this document my main contributions on the topics of *3D visibility and lines in space* and *geometric computing with low-degree algebraic surfaces*.

It should be stressed that these two directions of research are strongly related. A pragmatic illustration of that fact is that, in many of my contributions to 3D visibility and lines in space, the methodology and proof techniques are heavily based on the use of low-degree algebraic surfaces and, in particular, on quadrics (see Chapters 5, 6, 7, 13, and 18). More concretely, these connections are well illustrated by the following two examples. First, the line transversals to three lines in space span a ruled surface which is quadratic. Hence computing the lines through four lines, a basic operation in 3D visibility, is akin, in some sense, to intersecting two quadratic surfaces. Second, it is useful when studying lines in 3D to represent them as the points of some space. Unfortunately, the parameterization of the lines of projective 3-space cannot be as simple as for points and planes : it is known that there is no “algebraic model” for the space of lines that is itself a projective space [PW01, p. 143]. The best we can hope for is a parameterization of lines by the points of a quadratic hypersurface Ψ in projective 5-space (the so-called Plücker quadric). Now, each tangency condition in 3-dimensional space has an equivalent in 5-dimensional space : the lines meeting a given line correspond to the intersection

of a hyperplane of $\mathbb{P}^5(\mathbb{R})$ with Ψ and the lines tangent to a sphere correspond to the intersection of a hypersurface of degree 2 with Ψ . In this way, visibility problems in $\mathbb{P}^3(\mathbb{R})$ (or \mathbb{R}^3) can be formulated as problems on arrangements of low-degree surfaces in $\mathbb{P}^5(\mathbb{R})$. This shows that 3D visibility problems are inherently non-linear, even when dealing initially with polyhedral objects. Knowing how difficult problems on arrangements of low-degree surfaces already are in 3-dimensional space should give an idea as to why robust visibility computations present an enormous research challenge.

I present in the following subsections a brief description of the context and methodology of my work concerning curved objects, robustness, and efficiency. Then, in Chapters 2 and 3, I present a description of my contributions on the topics of *3D visibility and lines in space* and *geometric computing with low-degree algebraic surfaces* and include the corresponding papers in Parts II and III.

Curved objects

When confronted with curved objects, most applications in realistic rendering and solid modeling use polyhedral approximations to these surfaces and apply existing algorithms to manipulate these approximations. This approach leads invariably to very large data sets which has a negative impact on the algorithm's efficiency. In addition, the resulting computations are inaccurate and topological consistency cannot be guaranteed.

While research aimed at finding more and more efficient algorithms for dealing with triangulated meshes is essential due to the ever increasing complexity of such models, effective geometric computing may also be achieved by treating directly non-linear objects. Unfortunately, computational geometry provides efficient methods principally when the models can be described as large collections of linear objects. State-of-the-art algorithmic knowledge is largely helpless when it comes to robustly and efficiently manipulating even simple curved objects.

Most objects of interest in computer graphics and solid modeling are (semi-)algebraic, that is, they are defined by polynomial equations and inequations. Thus algebra plays a fundamental role in non-linear geometric computing. Many operations at the heart of algorithms dealing with curved primitives boil down to evaluating, manipulating and solving systems of polynomial equations.

Computer algebra and symbolic computation provide a possible approach to the realization of these operations. Well-known tools such as resultants, Gröbner bases and Sturm sequences have proved to be reliable means of attacking and solving simple geometric problems. Each problem can be formulated algebraically in many different ways and computer algebraists know from experience that the formulation can have a dramatic impact on the computational cost of the solving process. Also, the algebraic translation of a problem tends to obscure the information of geometrical nature. Finding formulations that are both geometrically meaningful and that lend themselves well to symbolic manipulation is always a difficult problem. Nevertheless, recent advances in computer algebra have opened new research alleys in the domain of geometric computing with curved surfaces. In particular, the performance of computer algebra systems has dramatically improved over the past few years and non-trivial problems can now be solved (see, for instance, Chapter 18 on Voronoi diagrams of lines). However, performing algebraic computations efficiently is more of an art than a science and reliance on these tools should enter only at the end of the pipeline, once the problem has been completely straightened




FIG. 1.1 – A chess set modeled entirely with quadrics (models courtesy of SGDL Systems, Inc.).

out from a geometric standpoint, as the following example illustrates.

Consider the problem of deciding whether four given spheres admit infinitely many common tangents. This problem can be directly translated into a polynomial system (of degree twelve) which can be solved, or, using the geometric characterization of Chapter 5, one can check whether the four spheres have aligned centers and, if so, check whether they admit one tangent (with a low-degree predicate). That the latter procedure will be computationally much more efficient than the former should be quite clear. Intuition also dictates that the smaller is the degree of the (class of) surfaces considered, the more the algorithms can take advantage of the geometry.

In general, the difficulties inherent in extending the computational geometry repertoire to curved objects are of several natures and mastering this challenge implies developing multidisciplinary research, involving *computational geometry*, *computer algebra*, *mathematics of curves and surfaces*, and *real and complex algebraic geometry*.

Indeed, compared to linear objects, handling even the simplest of curved objects is a major challenge and the intrusion of mathematics and computer algebra is often massive (see Chapter 5 on the characterization of degeneracies of lines tangent to spheres, Chapters 14, 15, 16 on the intersection of two quadrics and Chapter 18 on Voronoi diagrams of lines). Furthermore, basic algorithms have to be revisited and sometimes entirely rethought to suit the needs of curved primitives (see, again the work on quadrics, and Chapter 9 on computing line segments tangent to four among k polytopes). Robustness and degeneracy issues are also harder to handle in the world of curved objects than in the linear world and often require multidisciplinary tools (see Chapters 5 and 6 on the characterization of degeneracies of lines tangent to spheres or transversal to segments, Chapter 9 on lines tangent to arbitrary polytopes, and Chapters 15 and 16 on degenerate intersection of two quadrics).

Robustness

In geometric applications, degeneracies, which occur when the objects considered lie in special positions are, in practice, often the norm, not the exception. This is especially true in computer graphics and shape modeling. The example of Figure 1.1 can serve to illustrate this point. While the generic intersection between two quadrics is a non-planar curve of degree four, it is fairly clear that most if not all of the arcs outlining the boundary representation of the knight piece are line segments and conic arcs, and thus are planar. In other words, these curves are degenerate instances of surface intersections.

Completely and correctly accounting for degeneracies is a major challenge of most geometric software. The overall utility of an implementation may critically depend upon the correct treatment of special cases. Allowing for degenerate data vastly increases the number of special cases and thus the number of exceptional branchings and dedicated procedures in the software. One solution to this problem is to symbolically perturb the input, thereby symbolically moving the geometric problem away from a singular case and resolving degeneracies automatically (see the survey paper by Sugihara [Sug00]).

While handy, the symbolic perturbation approach is often inappropriate, for several reasons. First, its applicability is limited and it has been worked out in detail for only a small class of problems. A perturbation scheme has to move the input into generic position for every degenerate problem instance. Finding such a scheme can be extremely difficult and in any case requires recognizing that the input is degenerate to begin with, which is a substantial part of the work needed to fully treat degeneracies. Second, a symbolic perturbation may unnecessarily slow down computation. Indeed, solving a generic problem instance can involve manipulating arithmetic expressions of longer bit length, and thus be computationally more costly, than when the instance is singular. The work on the intersection of quadrics has shown precisely that (see Chapter 17). Since geometric inputs are often degenerate by design (aligned primitives, touching objects, etc), this can have a dramatic impact on the overall performance. Finally, when the input is degenerate, a geometric software applying symbolic perturbation does not necessarily solve the given problem instance. Indeed the output is obtained by a limit-taking process applied to the solution of the perturbed problem ; this limit may be structurally different than the solution to the given instance. For example, when intersecting two tangent objects, if the perturbation scheme moves them apart, the limit of the intersection of the perturbed instance will be the empty set instead of, say, a point. For applications where exactness is an issue, this might well be unacceptable.

By contrast, my work falls within the *paradigm of exact geometric computing*. Recall that a geometric object is really two things : a combinatorial structure (which, for instance, encodes the incidences between the constituting elements of the object) and a set of numerical quantities describing the embedding of the object in space, which one can usually consider to be algebraic. Since there are consistency constraints governing the relation between combinatorial information and numerical quantities, the numerical instability of geometric algorithms is intimately linked to this double nature of geometric objects. Exact geometric computing means performing computations in which numerical quantities are evaluated to sufficient precision in order for the underlying combinatorial structure to be mathematically exact.

The dependence of combinatorial decisions on numerical computation is encapsulated in the notion of *predicates*. Evaluating a geometric predicate consists in determining the sign of

some polynomial expression in the input numerical quantities ; a simple example is “does a given point lie to the left, to the right or on a given line?”. The *paradigm of exact geometric computing requires predicates to be evaluated correctly*, which ensures that the branchings made by the algorithm are correct and thus that the topological structure of the output is exact. This paradigm does not necessarily require, however, that the output numerical quantities are exact.

In full generality, solving a geometric problem in this paradigm amounts to

- characterize (mathematically and geometrically) the degeneracies, which can be intrinsic or algorithmic ;
- translate each geometric decision (and branching of the algorithm) into the determination of the signs of some algebraic expressions ;
- exactly and efficiently determine the sign of these algebraic expressions.

Note, however, that resolving each branching of the algorithm does not necessarily require evaluating algebraic expressions. Indeed, consider for instance the problem of determining the number of real roots of a univariate square-free polynomial. This can be done by evaluating pre-computed polynomials, such as Sturm sequences, or by dichotomy, using Descartes’ rule and Uspensky’s algorithm [RZ04]. It is known that the latter method is more efficient than the former, in particular for high-degree polynomials. Nonetheless, most simple geometric predicates are usually resolved by evaluating algebraic expressions and determining their signs.

Characterizing degenerate situations is crucial because undiscovered degeneracies usually result in fatal runtime error or combinatorially invalid output. Degeneracies are essentially of two different species : intrinsic degeneracies, which are inherent to the problem at hand and have to be handled by any algorithm that intends to robustly solve the problem, and algorithmic degeneracies, which are induced by algorithmic choices. A careful study of intrinsic degeneracies can lead to very efficient dedicated algorithms, as shown by our work on intersection of quadrics (see Chapters 14, 15, 16, and 17). Perhaps less intuitively, a fine analysis of non-intrinsic, algorithmic degeneracies can suggest structural modifications to the algorithm and induce improvements to its efficiency, as shown by the work of Angelier on 2D visibility complexes [Ang02].

Once a geometric predicate has been identified, it cannot in general be resolved using fixed-precision floating-point arithmetic : if a problem instance is nearly degenerate, then the value of the corresponding expression can be very small, possibly less than the rounding error in the floating-point evaluation of the expression. Hence the sign-evaluation may be incorrect, likely resulting, later, in a fatal error. Exact and efficient evaluation of predicates is usually performed with the use of exact (arbitrary-precision) arithmetic and filters. When the value of a polynomial expression is sufficiently far away from zero, filters can compute the exact sign of the expression without computing its value exactly (filters typically use interval arithmetic) ; otherwise the value of the expression is computed exactly.

The degrees of the polynomials expressing geometric predicates is a direct measure of the arithmetic bit length required for their exact evaluation. The degree is thus a measure of computational efficiency : the higher the degree of a predicate is, the more often the filters will fail and the more costly exact evaluation will be. Translating geometric decisions into low-degree predicates is thus important to strictly limit the arithmetic and computational demands of a robust implementation. However, measuring the computational efficiency of solving a predicate by the

degree of the polynomials involved is a model that seems appropriate only for simple predicates. Indeed, the above example about determining the number of real roots of a polynomial shows the limits of such model. Finally, it should be noted that deciding whether a predicate realizing a given decision has minimal degree or can be broken down into predicates of smaller degree seems to be, in general, a very difficult problem.

While simple enough for basic geometric primitives, the characterization of degenerate instances of a problem and its translation into low-degree predicates can be mathematically involved even when the primitives are simple. Consider for instance an environment made of spheres of arbitrary radii. From the point of view of 3D visibility, characterizing degeneracies requires detecting when quadruples of spheres admit infinitely many common tangents. While one may be content with such a high-level description, the “infinitely many common tangents” condition cannot be easily turned into (a set of) geometric predicates and detection will be both expensive and prone to numerical instability. Recently, we proved (using tools from complex projective geometry) that degenerate instances of quadruples of spheres are those with aligned centers and admitting at least one common tangent (see Chapter 5). Now, this condition has a much more geometric flavor and can be efficiently verified by low-degree geometric predicates.

My approach to geometric problems includes a systematic characterization of degeneracies (see Chapters 5, 6, 8, 9, 15, 16, and 17) and the design of low-degree predicates for making the corresponding geometric decisions (see Chapters 8). I however do not work on the design of arithmetic filters which is almost an independent field of research but use existing implementations such as the one developed in the kernel of the CGAL library.

Efficiency

In computational geometry an algorithm’s *efficiency* has been historically measured in terms of worst-case asymptotic complexity in the real RAM model of computation. In this model, time (and space) optimal algorithms are often the ultimate goal. Although worst-case optimal algorithms are clearly important, they are not always the most efficient in practice; quicksort is recognized as the fastest sorting algorithm in practice though only if it is implemented using the $\Theta(n^2)$ worst-case algorithm rather than the optimal $\Theta(n \log n)$ worst-case algorithm. Moreover, they can be hopelessly complicated; Chazelle’s linear-time algorithm for triangulating simple polygons [Cha91] is unanimously considered to be unimplementable [Ski97, pp. 355-357]. Nevertheless, we will always be interested in the worst-case complexity of an algorithm because it provides a starting point for further analysis.

If the size of the output of an algorithm can vary substantially, then the goal is to design algorithms that optimize the worst-case complexity expressed in terms of two parameters, input and output sizes, the so-called output-sensitive algorithms. An example of this is the problem of computing the convex hull of a set of points in 2D. Although the convex hull may contain all the input points it may also contain as few as three. An optimal worst-case algorithm computes the convex hull of n points in $\Theta(n \log n)$ time whereas there is a simple $\Theta(n \log h)$ -time algorithm where h is the output size [CSY97]. Clearly, the advantage of the output-sensitive approach is even greater when the output can be as large as $\Theta(n^4)$ which is the case when computing lines tangent to 4 amongst n objects in three dimensions.

Computational geometry problems, particularly in three and higher dimensions, are plagued by contorted worst-case configurations that simply never arise in practice. Data sets required to

induce an exponential-time performance of the simplex method for linear programming is a well-known example. When worst-case asymptotic complexity is not a relevant measure, another option is to do a probabilistic analysis of the algorithm in which one assumes a distribution on the inputs and derives the expected time complexity. This is often very difficult even for some simple cases such as uniform distributions (see, for instance, Chapter 11). Although such simple distributions are arguably irrelevant since real-world data is not uniformly random, a probabilistic analysis can give some insight beyond that given by a study of the complexity in the worst-case ; worst-case examples don't necessarily represent real-world data either !

A randomized algorithm is one which makes random choices. An analysis of a randomized algorithm yields an expected-time bound that holds for all inputs ; no assumption on the distribution of inputs is made. Randomized quicksort is the classic example. Randomized algorithms have known a tremendous success in computational geometry because they are generally incredibly simple to implement and are fast [Sei91]. Though randomized algorithms should perhaps not be used in critical real-time applications (such as landing a space shuttle), randomization is the way to go in graphics and solid-modeling applications.

Efficiency of algorithms can often be improved by first applying some simple algorithm if the data exhibits some special characteristics. One example is, again, in the quadric intersection algorithm ; degenerate cases are recognized and handled by specialized faster algorithms (see Chapters 15, 16, and 17). We advocate this approach. We do not, however, advocate the use of algorithms that solve the problem only in some special cases and crashes or outputs nonsense otherwise.

A judicious choice of data structures can have a huge impact on the efficiency of an algorithm. Hierarchical decompositions of space such as octrees is a classic example from graphics. When the problem involves answering many queries, this can be facilitated by first preprocessing the input into some data structure. Obviously, space becomes an important issue.

When the theoretical analysis of an algorithm shows that it is promising or when we fail to obtain an analysis because we simply don't have the right tools, the next step would be to implement the algorithm and analyze its running-time (and space) complexity on random and real data. This step has posed enormous problems to computational geometers partly because of the robustness issues discussed in the previous section. Many, possibly most, published geometric algorithms work only on non-degenerate input. Since typical graphics scenes contain degeneracies, it is critical to write algorithms that work on all inputs so that they may be tested.

To summarize, the objective is to design and implement algorithms, correct on all inputs, that require a reasonable amount of extra storage and that run quickly on realistic data. The performance will also be ascertained both theoretically and through experiments.

Chapitre 2

3D visibility and lines in space

2.1 Introduction

Three-dimensional visibility problems arose about 30 years ago, with the advent of the visualization of 3D objects. In three decades, the gap between theory and practice widened as the performance of hardware gradually supplemented the lack of theoretical understanding. Visibility computations are central in computer graphics applications and many problems are often considered solved in practice, though they do not necessarily have satisfying solutions. For instance, computing the limit of the umbra and penumbra cast by an area light source, identifying the set of blockers between any two polygons are typical examples of visibility problems that are essential for the realistic rendering of 3D scenes and that do not have any satisfying solutions : these problems are, in practice, approximated by many point-to-points visibility queries. To develop new algorithmic solutions for these questions, we need to improve the understanding of 3D visibility.

It should be stressed that, for solving most visibility problems, the graphics community has moved decades ago to hardware solutions such as dedicated z -buffer hardware or, more recently, GPUs (Graphics Processing Units). However, these solutions are mainly used for computing 2D images though very rarely for computing light-flow simulations in 3D scenes. The reason for this is that current hardware solutions are suited for dealing with two-dimensional sets of rays (for instance, a z -buffer can be used for efficiently computing the view from a point) but are not suited for general visibility queries which deal by essence with four-dimensional sets of rays (lines in 3D have 4 degrees of freedom).

My research objectives concerning 3D visibility are primarily the design and implementation of algorithmic solutions for 3D visibility problems. In particular, I am interested in two major challenges that are (i) to compute the limits of umbra and penumbra cast by an area light source and (ii) to solve surface-to-surface visibility queries, that is to answer queries of the type : are two elements of surfaces mutually (partially) visible ?

To reach these objectives, we first need a better understanding of the properties of sets of lines and line segments that are tangent or transversal to objects in a three-dimensional scene. This fundamental direction of research is of interest independently of any direct application and it has been the focus of my research activities on the subject since 2000. In particular, most of the work I present here focuses primarily on the structural and combinatorial properties of sets

of three-dimensional lines and line segments in the context of visibility problems.

This work also leads to some algorithmic solutions for 3D visibility problems. In particular, the study of the set of free line segments tangent to four among k polytopes leads to an algorithm for computing such segments (see Chapter 9); a brief description of an implementation of a preliminary version of this algorithm is given in Chapter 10. More generally, this work should also lead to an algorithm for computing the visibility skeleton, a data structure that encodes visibility information [DDP97]. Note also that the work on characterizing degenerate configurations of lines tangent to polytopes (see Chapters 6 and 9) and the work on predicates for computing such lines (see Chapter 8) is critical for any robust implementation of such algorithms. Finally, it should be mentioned that the work of Chapter 9 on the properties of lines tangent to polytopes also led to new combinatorial results on the umbra cast on a plane by polygonal light sources in the presence of obstacles (see Chapter 13). These results are critical to understand the structure of the umbra and penumbra cast by non-trivial light sources and they should ultimately impact the techniques for computing such objects exactly and efficiently.

I now present a panorama on 3D visibility and, in Section 2.2, give a description of my contributions on the topic.

Panorama

In a given scene, two points are visible if the line segment joining them is free of any obstacles in the scene. The study of visibility is thus essentially the study of free line segments. The set of free line segments is bounded by segments that are tangent to the obstacles; in 3D, a line is tangent to up to 4 objects in general position. To put it another way, consider a moving viewpoint in a 3D scene. The view will change when a new object appears (or disappears) from behind an obstacle. When this happens there is a line passing through the viewpoint, tangent to the obstacle and arriving at the new object (and actually tangent to it as well). It is thus imperative to fully understand tangencies between lines and obstacles in order to answer visibility queries.

As mentioned above, two major challenges in 3D visibility are (i) to compute the limits of umbra and penumbra cast by an area light source and (ii) to solve surface-to-surface visibility queries, that is to answer queries of the type: are two elements of surfaces mutually (partially) visible? Classical data structures based on spatial decomposition, such as octrees or binary space partition (BSP) trees, are of little help in that respect. These structures are, indeed, designed to solve ray-shooting or point-to-point visibility queries, *i.e.*, problems that are intrinsically of dimension two (the dimension of the set of lines going through a point), and they do that very well, whereas surface-to-surface visibility involves the set of lines intersecting two surfaces, which is four-dimensional.

One data structure that has been proposed for encoding 3D visibility information is the visibility complex, which, roughly, is a partition of the space of maximal non-occluded line segments into connected cells of segments that are tangent to exactly the same objects. The structure was introduced by Pocchiola and Vegter in 2D [PV96b] and by Durand, Drettakis, and Puech in 3D [Dur99, DDP02]. In 2D, the visibility complex has been extensively studied [PV96a, Riv97, HH02, AP03b] as well as its application for rendering [ORDP96, CF99]. In 3D, because this structure has cells of dimension up to four, Durand et al. introduced, for practical reasons, the visibility skeleton, the structure that consists of the cells of dimension

zero and one of the visibility complex [DDP97]. They proposed a proof-of-concept implementation showing that the skeleton leads to higher quality images in light simulation together with improved computation time compared to previous algorithms [DDP99].

Despite these positive results their pioneering approach suffered two major impediments : the poor performance of their algorithm, based on a systematic enumeration of all possibilities leading to worst-case complexity $\Theta(n^5)$ and observed complexity $\Theta(n^{2.5})$ thanks to the use of heuristics [Dur99], and the lack of robustness of their implementation which requires a great deal of time-consuming human intervention to remove degeneracies from the scenes. As a result, the largest scene they were able to handle had no more than 1,500 triangles [DDP99].

A few years later, Duguet and Drettakis proposed a practically robust implementation for computing a two-dimensional section of the visibility skeleton in which the lines supporting the maximal non-occluded line segments are restricted to be all concurrent (possibly at infinity) [DD02]. Their implementation was successfully tested on scenes with over 100,000 polygons showing that, in the restricted context of punctual lights or lights at infinity, the approach leads to higher quality images in light simulation together with improved computation time compared to previous algorithms.

There are two reasons why little research has been done on the topic of 3D global visibility. One is that the problems are formidable because, as hinted earlier, the problems are intrinsically non-linear, even when dealing with polyhedral data, and are fairly high-dimensional since the space of lines in 3D is of dimension four. The other reason is that the space requirements of global visibility data structures have always been considered huge due to theoretical complexity bounds : the size of the visibility skeleton is $\Theta(n^4)$ in the worst case for n triangles. Durand [Dur99] observed a much smaller empirical size of $\Theta(n^{2.4})$, which was unfortunately still much too large for real-world graphics applications. However the small size of the scenes they were able to handle seriously diminishes the significance of this experimental asymptotic complexity since it is not clear that it corresponds to the asymptotic behavior. I also present, in the next section, theoretical evidence supporting the idea that these bounds are, indeed, overly pessimistic for practical applications.

2.2 Summary of contributions

2.2.1 Introduction

I first present briefly my contributions, and their relevance, on 3D visibility and lines in space. I then present, in Sections 2.2.2 to 2.2.9, a more detailed, though concise, descriptions of these results.

First, I present some results on the *structural properties of lines that are tangent or transversal to four primitives* in three dimensions. Chapter 5 first presents a characterization of the configurations of four spheres in 3D that admit infinitely many common tangents. This characterization is complete and elementary : the spheres must have aligned centers and at least one common tangent. This elegant characterization settled a conjecture by T. Theobald and concluded a series of (seven) papers that investigated variants of this problem where the spheres are congruent or some spheres are replaced by lines.

Chapter 6 then presents a characterization of the set of common transversals to four (and, more generally, to n) line segments in three dimensions. We prove, in particular, that four segments in 3D admit at most 4 connected components of common transversals and possibly exactly 4 transversals. This result is fundamental and surprising since four lines in 3D admit at most 2 connected components of common transversals.

In Chapter 7, we investigate the number of common tangents to four triangles (a line is tangent to a triangle if it intersects its boundary). We show that the number of such common tangents can be, surprisingly, as high as 62. The corresponding construction consists of very thin triangles but we also ran experiments on five millions quadruples of random triangles; we obtained several quadruples with 40 common tangents and observed that quadruples that admit common tangents have 16 tangents or more with probability at least 15%. We also show an upper bound of 162 which also holds for arbitrary degenerate configurations where each tangent line may become a connected component of tangent lines. It should be noticed that while the lower bound of 62 is already surprisingly high, the gap between the lower and upper worst-case bounds remains large.

Related to these problems, Chapter 8 presents a study of various predicates concerning line transversals to lines and segments in 3D. In particular, we compute the degrees of standard methods of evaluating these predicates. We show that the degrees of some of these methods are surprisingly high (up to 168), which may explain why computing line transversals with finite-precision floating-point arithmetic is prone to error. Our results also suggest the need to explore alternatives to the standard methods of computing these quantities. Indeed, efficient predicates for solving such elementary queries is important for the development of robust and efficient algorithmic solutions to 3D visibility problems.

I then present several results on the *combinatorial properties of geometric structures* in the context of three-dimensional visibility. Chapters 9 and 11 focus on the complexity of the visibility complex, a data structure that encode visibility information, which is, roughly speaking, the partition of the space of maximal free line segments into connected components of segments that are tangent to the same objects. Chapter 12 focuses on the size of silhouettes of polyhedra and Chapter 13 on complexity of the umbra and penumbra cast by non-point light sources.

Chapter 9 presents new combinatorial bounds on the set of lines and free line segments tangent to possibly intersecting arbitrary convex polyhedra. In particular, we prove that k such polytopes with a total of n edges admit, in the worst case, $\Theta(n^2k^2)$ connected components of maximal free line segments tangent to at least four polytopes. This result is a significant improvement over the previous known trivial bound of $O(n^4)$ and over the same $O(n^2k^2)$ bound for the considerably easier case of disjoint convex polyhedra in algebraic general position [EGHHZ00, BDD⁺02]. This result also implies that the visibility complex admits $\Theta(n^2k^2)$ vertices, in the worst case. More generally, this bound also holds for the size of whole visibility complex but, if the proof is trivial for polytopes in generic position, it is more subtle (and unpublished) for polytopes in arbitrary position.

Chapter 11 focuses on the size of the visibility complex for uniformly distributed objects. We show that, amongst n uniformly distributed unit balls in \mathbb{R}^3 , the expected number of maximal non-occluded line segments tangent to four balls is linear. This bounds generalizes in various directions. In particular, the bound applies to the the expected size of the visibility complex of uniformly distributed unit balls. This also leads to a $O(n^2)$ bound on the expected size of

the visibility complex of random polyhedra or polygons (of total complexity $O(n)$) of bounded aspect ratio. These results significantly improve the best previously known bounds of $O(n^{8/3})$ for unit balls [DDP02].

The results of Chapters 9 and 11 show that taking into account the underlying structure of a scene provides new insight on the size of 3D visibility data structures. Furthermore, they give evidence that the previously known bounds (mentioned in the previous section) were overly pessimistic and that the storage requirement for these data structures is not necessarily prohibitive. There is thus some hope that these structure may be useful in real-world graphics applications.

Chapter 12 provides theoretical evidence supporting, for a large class of objects, the widely observed phenomenon that the size of the silhouette of a polyhedron is much smaller than the size of the whole polyhedron. We prove that the silhouette of a polyhedron that approximates a surface in a reasonable way has expected size $O(\sqrt{n})$ where the average is taken over all points of view and n is the complexity of the polyhedron; the surfaces may be non-convex and non-differentiable and they may have boundaries. Note that the average is taken over all viewpoints for a given surface, and not on a set of surfaces. This confirms a widely accepted belief in computer graphics and is the first complexity result for silhouettes of non-convex polyhedra.

Finally, Chapter 13 presents some combinatorial results on the umbra and penumbra cast by non-trivial light sources. A point is in the umbra if it does not see any part of any light source; it is in full light if it sees entirely all the light sources; otherwise, it is in the penumbra. While the common boundary of the penumbra and the full light is fairly well understood, less is known about the boundary of the umbra and computing such shadow boundaries is a difficult problem in the case of non-point light sources. We present several surprising lower bounds. In particular, a segment light source may cast on a plane, in the presence of two triangles, up to four connected components of umbra; also, replacing the two triangles by two disjoint fat polytopes of total size n induces $\Theta(n)$ connected components of umbra in the worst case. For polygonal light sources in a scene of k polytopes of total complexity n , we show that the umbra may admit $\Omega(n^2k^3 + nk^5)$ connected components and has complexity $O(n^3k^3)$. These results are the first non-trivial bounds on the size of the umbra and show that the umbra can be surprisingly complicated, even in the presence of disjoint fat obstacles. More generally, this work gives some insight on the structure of the umbra which may lead to new algorithms for computing them exactly and efficiently.

The above results focus on the structural and combinatorial properties of sets of lines or line segments that are tangent or transversal to primitives. Some of these results are also critical for the development of *robust and efficient algorithms* for 3D visibility. In particular, in Chapter 9, the study of the set of free line segments tangent to four among k polytopes leads to an algorithm for computing such segments (see Section 9.5). Chapter 10 presents a video that describes an implementation of a preliminary version of this algorithm. More generally, this algorithm can also compute the whole visibility skeleton of disjoint polytopes but its generalization to possibly intersecting polytopes or disjoint non-convex polyhedra is the subject of future research. Note also that the work on characterizing degenerate configurations of lines tangent to polytopes (see Chapters 6 and 9) and the work on predicates for computing such lines (see Chapter 8) is critical for any robust implementation of such algorithms.

Finally, I mention a few other results on 3D and 2D visibility problems that I do not include

in this document because they are only remotely connected to the results I mentioned so far. First, in a context of layer manufacturing in rapid prototyping, we have proposed a practical approximation solution to the problem of computing an optimal slicing strategy and drilling direction so as to minimize the volume of the regions that are inaccessible to a two and one-half degrees of freedom tool during the manufacturing process [LLB04, Lau05]. This strategy has been tested with success and is now used by CIRTES¹, a research company in rapid prototyping. We also presented preliminary results on the problem of effectively computing visibility graphs of spheres [MDC05] and solved several visibility problems in two dimensions. In particular, we proved theoretical bounds and made an experimental assessment of the size of the visibility complex of randomly distributed objects in the plane [ELPZ07] and presented an algorithm for parabolic ray-shooting [CEK⁺07].

I now present a more detailed descriptions of these results.

2.2.2 Common tangents to spheres in \mathbb{R}^3

The lines tangent to spheres has been persistently investigated in the recent years. Macdonald *et al.* [MPT01] proved that four unit spheres have at most 12 common tangents in general, and infinitely many common tangents if and only if the centers are aligned. The bound of 12 was independently obtained by Devillers *et al.* [DMPT03]. Examples show that, in the finite case, this bound is tight [DMPT03, MPT01], yet, according to Megyesi [Meg01], it drops to 8 in the case of unit spheres with coplanar but non-collinear centers. However, the upper bound of 12 remains valid when the spheres have arbitrary radii. Sottile and Theobald [ST02] proved that there are $3 \cdot 2^{n-1}$ complex common tangent lines to $2n - 2$ general spheres in \mathbb{R}^n , and that there exists a choice of spheres with all common tangents real.

Recently, progress has also been made in understanding the varieties of common tangents to spheres and transversals to lines. Theobald [The02] described the configurations of three lines and a sphere having infinitely many common tangents/transversals. Next, Megyesi *et al.* [MST03] characterized the families of two lines and two quadrics of $\mathbb{P}^3(\mathbb{C})$ with infinitely many tangents/transversals, and applied their results to the case of two lines and two spheres of \mathbb{R}^3 . Last, Megyesi and Sottile [MS05] classified the families of one line and three spheres of \mathbb{R}^3 with infinitely many tangents/transversals.

The question of characterizing the positions of four spheres of various radii with infinitely many common tangents remained open. Quoting Theobald [The02]: “*We conjecture that there does not exist any configuration with four balls of arbitrary radii, non-collinear centers and infinitely many common tangent lines.*” In Chapter 5, we confirm this expectation and prove the following.

Theorem 2.1. *Four distinct spheres in \mathbb{R}^3 have infinitely many real common tangent lines if and only if they have aligned centers and at least one real common tangent.*

More precisely we prove that four spheres with infinitely many common real tangents either intersect in a circle, possibly degenerating to a point, or each sphere has a circle of tangency with one and the same quadric of revolution with symmetry axis the line through all centers (see Figure 2.1); such a quadric is *unique* and can be a cone, a cylinder or a hyperboloid of one

¹<http://www.cirtes.fr>.

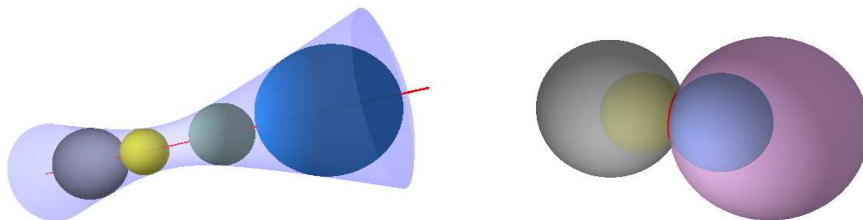


FIG. 2.1 – Two examples of quadruples of spheres with infinitely many common tangents.

sheet. Furthermore, the common tangents to the four spheres are exactly the common tangents to any three of them.

The idea of the proof is as follows. Considering a line parameterized by a point p and a direction v , we first characterize the lines tangent to four spheres as the solutions of three equations in p and v . We then eliminate p from that system of equations. This gives two curves¹ in the 2D projective space of directions, whose intersection contains all directions along which a common tangent line to the four spheres is observed. We then prove that the two curves intersect in a finite number of points.

The key idea behind the proof is that if the two curves, envisaged as complex projective curves, had a common component of positive dimension, this component would intersect, in complex space, the imaginary conic $|v|^2 = 0$ and we show that this is not the case.

It should be stressed that any solution to the problem of characterizing sets of four spheres with infinitely many tangent lines must be computational to some extent, because while we are interested in real lines, the “native” system of equations is over \mathbb{C} . Any understanding of the system should involve sensitivity to complex degeneracies. In our proof, computations flow towards revealing such complex degeneracies, but are short-circuited by use of reality assumptions.

2.2.3 Transversals to line segments in three-dimensional space

We address, in Chapter 6, the following basic question : What is the cardinality and geometry of the set of transversals to an arbitrary collection of n line segments in \mathbb{R}^3 ? Here a segment may be open, semi-open, or closed, and it may degenerate to a point ; segments may intersect or even overlap. Since a line in \mathbb{R}^3 has four degrees of freedom, it can intersect at most four lines or line segments in generic position. Conversely, it is well-known that four lines in generic position admit zero or two transversals ; moreover, four arbitrary lines in \mathbb{R}^3 admit zero, one, two, or infinitely many transversals [HCV52, p. 164]. In contrast, our work shows that four arbitrary line *segments* admit up to four or infinitely many transversals.

We completely describe the structure of the connected components of transversals to a collection of n line segments in \mathbb{R}^3 . Generically, the set of transversal to four segments consist of zero, one, or two lines. We catalog the non-generic cases and show that $n \geq 3$ arbitrary line

¹A cubic and a quartic when the centers are affinely independent, a conic and a sextic when the centers are coplanar with no three aligned.

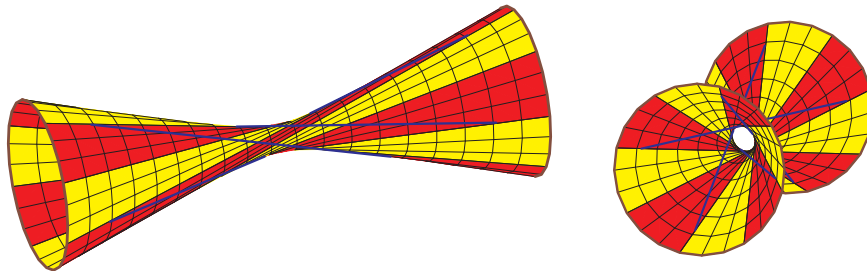


FIG. 2.2 – Two views of a hyperboloid of one sheet containing four line segments and their four connected components of transversals (corresponding to the shaded regions). The four segments are symmetric under rotation about the axis of the hyperboloid.

segments in \mathbb{R}^3 admit at most n connected components of line transversals, and that this bound can be achieved in certain configurations when the segments are coplanar, or they all lie on a hyperboloid of one sheet (see Figure 2.2). This implies a tight upper bound of n on the number of geometric permutations of line segments in \mathbb{R}^3 .

More precisely, our results are the following. We say that two transversals to a collection of line segments are in the same *connected component* if and only if one of the transversals can be continuously moved into the other while remaining a transversal in \mathbb{R}^3 to the collection of line segments. (For the sets of line transversals considered here, the notions of connected and path-connected components are equivalent since all sets are semi-algebraic.) Equivalently, the two points in line space (e.g., in Plücker space [PW01]) corresponding to the two transversals are in the same connected component of the set of points corresponding to all the transversals in \mathbb{R}^3 to the collection of line segments.

Our main result is the following theorem.

Theorem 2.2. *A collection of $n \geq 3$ arbitrary line segments in \mathbb{R}^3 admits any number from 0 to n of connected components of line transversals. More precisely, the set of line transversals consists of at most two isolated lines unless the segments lie in one of the following three configurations :*

1. *the n segments are all contained in lines of one ruling of (a) a hyperbolic paraboloid or (b) a hyperboloid of one sheet, or*
2. *they are all concurrent, or*
3. *they all lie in a plane, with the possible exception of a group of one or more segments that all meet that plane at the same point.*

In cases 1(a) and 2, the transversals form at most one connected component. In cases 1(b) and 3, the transversals can have any number from 0 to n of connected components. Moreover, in case 3, if all segments are not coplanar, this number is at most $n - 1$.

In cases 1–3, each connected component can consist of infinitely many lines or reduce to an isolated line. For example, three segments forming a triangle and a fourth segment intersecting the interior of the triangle in one point have exactly three transversals (Figure 2.3 shows a similar example with infinitely many transversals). Also, the four segments in Figure 2.2 can be shortened so that the four connected components of transversals reduce to four isolated transversals.

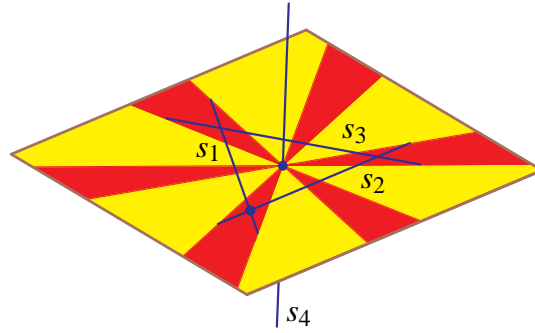


FIG. 2.3 – Four segments having three connected components of transversals.

We also show that a simple consequence of our theorem is the following bound on the number of geometric permutations of n segments in \mathbb{R}^3 .

Corollary 2.3. *A set of $n \geq 3$ pairwise disjoint segments in \mathbb{R}^3 admits up to n geometric permutations and this bound is tight.*

Finally, we also show that an $O(n \log n)$ -time algorithm for computing the transversals to n segments follows directly from the proof of the above theorem.

The idea of the proof is as follows. We first consider the segments' *supporting lines*, i.e., the lines containing the given segments (defined as vertical if the corresponding segment is a point). We prove Theorem 2.2 by studying the set of line transversals to three segments in the following three different cases, which cover all possibilities but are not exclusive : (i) three supporting lines are pairwise skew, (ii) two supporting lines are coplanar, and (iii) all the segments are coplanar.

For instance, in the first case and if, furthermore, all $n \geq 3$ supporting lines lie in the first ruling of a hyperboloid of one sheet (see Figure 2.2), the line transversals are the lines in the second ruling and they can be parameterized by points on an ellipse, for instance, by their intersection points with an ellipse lying on the hyperboloid of one sheet. Thus, the set of transversals to the n segments corresponds to the intersection of n intervals on this ellipse. This intersection can have any number of connected components from zero up to n , and any of these connected components may consist of an isolated point on the ellipse. The set of transversals can thus have any number of connected components from zero up to n , and any of these connected components may consist of an isolated transversal. Figure 2.2 shows two views of a configuration with $n = 4$ line segments having four connected components of transversals.

2.2.4 Lines tangent to four triangles in three-dimensional space

We study in Chapter 7 the basic question : What is the number of common tangent lines to four triangles ? We establish upper and lower bounds on this number. We show in particular that this number could be, surprisingly, as high as 62.

We say that a line is *tangent* to a triangle if it meets an edge of the triangle. Let $n(t_1, t_2, t_3, t_4)$ be number of lines tangent to four triangles $t_1, t_2, t_3,$ and t_4 in \mathbb{R}^3 . This number may be infinite if the lines supporting the edges of the different triangles are not in general position.

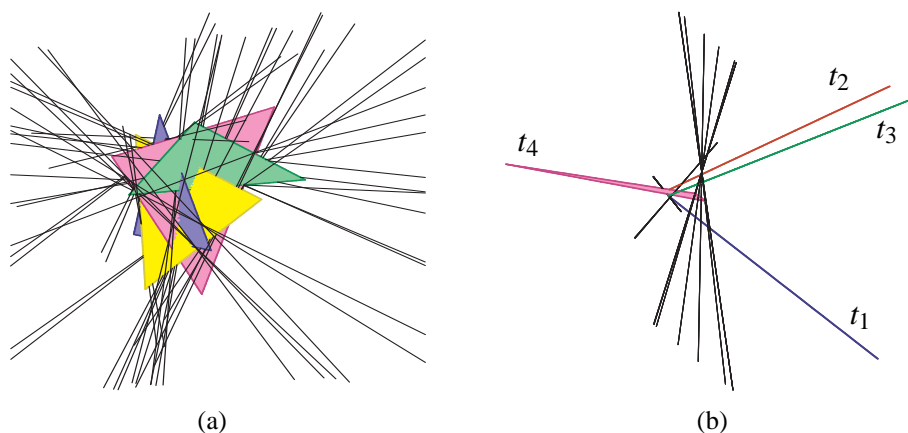


FIG. 2.4 – Triangles with many common tangents

Our first step is to consider the algebraic relaxation of this geometric problem in which we replace each edge of a triangle by the line in $\mathbb{P}^3(\mathbb{R})$ supporting it, and then ask for the set of lines in $\mathbb{P}^3(\mathbb{C})$ which meet one supporting line from each triangle. Since there are $3^4 = 81$ such quadruples of supporting lines, this is the disjunction of 81 instances of the classical problem of transversals to four given lines in $\mathbb{P}^3(\mathbb{C})$. As there are two such transversals to four given lines in general position, we expect that this algebraic relaxation has 162 solutions. We say that four triangles t_1, t_2, t_3, t_4 are in (algebraic) *general position* if each of the 81 quadruples of supporting lines have two transversals in $\mathbb{P}^3(\mathbb{C})$ and all 162 transversals are distinct. Let \mathcal{T} be the configuration space of all quadruples of triangles in \mathbb{R}^3 and $T \subset \mathcal{T}$ consist of those quadruples which are in general position. Thus if $(t_1, t_2, t_3, t_4) \in T$, the number $n(t_1, t_2, t_3, t_4)$ is finite and is at most 162.

Our primary interest is the number

$$N := \max\{n(t_1, t_2, t_3, t_4) \mid (t_1, t_2, t_3, t_4) \in T\}.$$

Our results about this number N are two-fold. First, we show that $N \geq 62$.

Theorem 2.4. *There are four disjoint triangles in T with 62 common tangent lines.*

The idea is to perturb a configuration of four lines in \mathbb{R}^3 with two real transversals. The triangles in our construction are very ‘thin’—the smallest angle among them measures about 10^{-11} degrees (see Figure 2.4(b)). We ran a computer search for ‘fatter’ triangles having many common tangents, checking the number of tangents to 5 million different quadruples of triangles over six months of CPU time. It appears that random quadruples of realistic triangles often have a fair number of common tangents (see Figure 2.4(a)). Several had as many as 40 common tangents, and quadruples that admit common tangents have 16 tangents or more with probability at least 15%.

We can improve the upper bound on N when the triangles are disjoint.

Theorem 2.5. *Four triangles in T admit at most 162 distinct common tangent lines. This number is at most 156 if the triangles are disjoint.*

We believe, however, that the upper bounds we give here are far from optimal. When the four triangles are not in general position, the number of tangent lines can be infinite. In this case, we may group these tangents by connected components : two line tangents are in the same component if one may move continuously between the two lines while staying tangent to the four triangles. Each quadruple of edges may induce up to four components of tangent lines (see Chapter 6), giving a trivial upper bound of 324. This may be improved.

Theorem 2.6. *Four triangles have at most 162 connected components of common tangents. If the triangles are disjoint, then this number is at most 156.*

Finally, we also prove the following result about the parity of the number of tangents.

Theorem 2.7. *If $(t_1, t_2, t_3, t_4) \in T$, then $n(t_1, t_2, t_3, t_4)$ is even.*

This result may not seem surprising as complex roots come in conjugate pairs. However, this usual argument does not apply because we seek tangents to triangles and not transversals to lines. Frequently, only one of two real transversals to a quadruple of supporting lines is tangent to the triangles.

2.2.5 Predicates for line transversals in 3D

Computing line transversals to lines or segments is an important operation in solving 3D visibility problems arising in computer graphics ; see Chapters 7, 9, 13 and [DD02, DDP97, DDP02, EGHZ00, PD90]. In Chapter 8, we study various predicates and their degrees concerning line transversals to lines and segments in 3D.

A predicate is a function that returns a value from a discrete set. Typically, geometric predicates answer questions of the type “Is a point inside, outside or on the boundary of a set?”. We consider predicates that are evaluated by boolean functions of more elementary predicates, the latter being functions that return the sign ($-$, 0 or $+$) of a multivariate polynomial whose arguments are a subset of the input parameters of the problem instance (see, for instance [BP00]). By *degree* of a procedure for evaluating a predicate, we mean the maximum degree in the input parameters among all polynomials used in the evaluation of the predicate by the procedure. In what follows we casually refer to this measure as the degree of the predicate. We are interested in the degree because it provides a measure of the number of bits required for an exact evaluation of our predicates when the input parameters are integers or floating-point numbers ; the number of bits required is then roughly the product of the degree with the number of bits used in representing each input value.

In Chapter 8, we first study the degree of standard predicates for determining the number of line transversals to four lines or four segments in 3D ; recall that four lines in \mathbb{R}^3 admit 0, 1, 2 or an infinite number of line transversals and that four segments admit up to 4 or an infinite number of line transversals (see Chapter 6). We also consider the predicate for determining whether a minimal (i.e., locally shortest) segment transversal to four line segments is occluded by a triangle. Finally, we study the predicate for comparing planes through two fixed points, each containing a third (rational) point or a line transversal to four segments or lines. This predicate arises in the rotating plane-sweep algorithm that computes the minimal free segments tangent to four among k convex polyhedra in 3D (see Chapters 9 and 10).

Our study shows that standard procedures for solving these predicates have high degree. In particular, we show that determining whether a minimal segment transversal to four line segments is occluded by a triangle can be evaluated by a predicate of degree 90 in the coordinates of the points defining the four segments and the triangle. Also, the predicate for comparing, in a rotational sweep, two planes, each defined by a line transversal, can be evaluated by a procedure of degree 168 in the coordinates of the points defining the input segments. These very high degrees may help explain why using fixed-precision floating-point arithmetic in implementations for solving 3D visibility problems are prone to errors when given real-world data.

2.2.6 Lines and free line segments tangent to arbitrary three-dimensional convex polyhedra

We investigate in Chapters 9 and 10 the complexity and construction of sets of tangent line segments in a scene of three-dimensional polyhedra. We present, in Chapter 9, two types of results. First we present some combinatorial bounds on the number of lines or free line segments tangent to four polytopes ; we then present an algorithm for computing such free line segments. In Chapter 10, we present a short description of a video illustrating our algorithm and its implementation in the case of disjoint polytopes. We now detail these results.

Combinatorial bounds. We show that among k polytopes, of total complexity n , that may *intersect*, 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 to two convex polygons is always 4 if disjoint, and could be linear in the size of the polygons if they intersect. Our result is also valid for 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 [BDD⁺02]) 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 2.8. *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.*

This result is a significant improvement over the previous known trivial bound of $O(n^4)$ and over the same $O(n^2k^2)$ bound for the considerably easier case of disjoint convex polyhedra in algebraic general position [EGHHZ00, BDD⁺02]. Note that, when $k > 4$, neither of the two results stated in Theorem 2.8 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 2.8 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

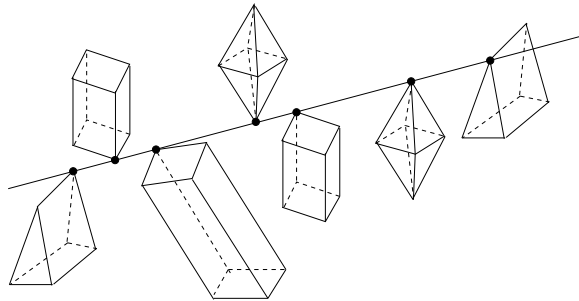


FIG. 2.5 – A line tangent at a vertex of each of k polytopes.

polyhedra (here four) with n edges in total [Aga94]. Moreover, we prove a tighter bound when one of the four polytopes has few edges.

Theorem 2.9. *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 2.8. Whereas Theorem 2.8 bounds the number of connected components of tangents, Theorem 2.10 bounds the number of isolated tangents with some notion of multiplicity. For example, a line through k vertices and tangent to the corresponding k polytopes (see Figure 2.5) is counted $\binom{k}{2}$ times which is the number of minimal sets of vertices that admit that line as an isolated transversal. Although neither theorem implies the other, we will prove that the upper bound of Theorem 2.8 is easily proved using Theorem 2.10.

Theorem 2.10. *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.* [DDP02] 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 [DDP97]. 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. Note that S. Hornus [Hor06] recently showed how the whole visibility complex of a set of disjoint convex polyhedra can be computed using this algorithm as a starting point. A preliminary version of this algorithm was described for disjoint convex polyhedra in X. Goaoc's Ph.D. thesis [Goa04] and implemented by L. Zhang; Chapter 10 presents a short description of a video illustrating our algorithm and its implementation.

Theorem 2.11. *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*

	Worst-case	Expected
free lines to a polyhedron	$\Theta(n^4)$ (trivial)	
free lines above a polyhedral terrain	$O(n^3 2^{c\sqrt{\log n}})$ [HS94, Pe194]	
free lines among disjoint homothetic polytopes	$\Omega(n^3)$ [dBEG98]	
free lines among unit balls	$\Omega(n^2)$ [§11], $O(n^{3+\epsilon})$ [AAS04]	$\Theta(n)$ [§11]
free lines among balls	$\Omega(n^3)$ [Gli07], $O(n^4)$	
max. free segments above a polyhedral terrain	$\Theta(n^4)$ [CS89b]	
isolated maximal free segments among k arbitrary polytopes	$\Theta(n^2 k^2)$ [§9]	
maximal free segments among unit balls	$\Theta(n^4)$ [Gli07]	$\Theta(n)$ [§11]

TAB. 2.1 – 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.

of supports and tangent to the corresponding polytopes. We can also compute, in $O(n^2 k^2 \log n)$ 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 [AS96]. 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 [DK83].

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.

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 (see Chapter 6). 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.

2.2.7 The expected number of visibility events is linear

The visibility complex, a partition of the set of maximal free line segments, has been proposed as a unified data structure encoding the visibility information of a scene [PV96b] and has been used for rendering purposes [DDP02]. Other related data structures include Pellegrini’s ray-shooting structure [Pel93], the aspect graph [PD90] and the visual hull [Lau94]; see [Dur00] for a recent survey.

One problem with these types of data structures which may prevent their application in practice is their potentially enormous size; the size of the visibility complex of a set of n triangles in \mathbb{R}^3 is $\Theta(n^4)$ in the worst case [DDP02], which is prohibitive even for scenes of relatively modest size. Worst-case examples are somewhat artificial and indeed Durand, Drettakis and Puech [DDP97] provide empirical evidence indicating that these worst-case upper bounds are largely pessimistic in practical situations; they observe a quadratic growth rate, albeit for rather small scenes. In 2D, while the worst-case complexity of the visibility complex is quadratic, experimental results strongly suggest that the size of the visibility complex of a scene consisting of scattered triangles is linear [CF99].

Our goal is to provide theoretical evidence to support these observations. To this end we investigate the *expected size* of the visibility complex, or equivalently, the expected number of visibility events, occurring in scenes in \mathbb{R}^3 . A visibility event is a combinatorial change in the view of a moving observer; such an event occurs when the viewing direction becomes tangent to some objects. For sets of convex objects in general position in \mathbb{R}^3 , the viewing direction can be tangent to at most four objects. Visibility events thus correspond to maximal non-occluded line segments tangent to at most four objects; combinatorially different visibility events correspond to the faces of the visibility complex.

In Chapter 11, we prove that the expected number of maximal non-occluded line segments tangent to four balls, amongst n uniformly distributed unit balls in \mathbb{R}^3 , is linear. This improves the previously known upper bound of $O(n^{8/3})$ by Durand et al. who proved the more general result that the expected number of (possibly occluded) lines tangent to four balls is $O(n^{8/3})$ for the same model [DDP02]. The intuition behind our proof is that, given a line segment tangent to four balls, the probability that this segment is not occluded by any other ball is the probability that a cylinder-like volume of radius 1 about the segment is free from the centers of the other balls. This probability decays roughly exponentially fast with the length of the segment, yielding the result. Using our techniques we then show a linear bound on the expected size of the visibility complex of n uniformly distributed unit balls in \mathbb{R}^3 . A simple computation then provides us with the same result for the Poisson distribution.

Our results generalize in the following ways. We show that, for certain types of visibility events, the linear bound also applies to balls of various but bounded radii, to polyhedral objects enclosed between two concentric balls of fixed radius, and even to non-fat objects such as polygons, enclosed between two concentric circles of fixed radius, whose centers and normals are uniformly distributed. For the remaining types of visibility events (namely those occurring close to the boundary of the scene – see Section 11.7.3 for the details), we prove only an $O(n^2)$ bound, which is still an improvement over the bound of $O(n^{8/3})$ by Durand et al. [DDP02].

Of course objects in graphics scenes are seldom distributed uniformly or according to a Poisson point process. We chose this model because it allows tractable proofs of theoretical results. This is important in a context where there are few rigorous results either theoretical or expe-

rimental. The same model, albeit with significant simplifying assumptions, has also been used to study the average complexity of ray shooting [SKHBS02, SKM98] and occlusion culling for 2D urban scenes [NFLYCO99]. It is interesting to note that Szirmay-Kalos et al. [SKHBS02], after establishing bounds on the average complexity of ray shooting in scenes consisting of unit balls distributed according to a Poisson point process, tested their algorithms on a small number of realistic scenes. The results they obtain are consistent with those predicted by the theoretical results thus providing some evidence that the model is helpful. No other model has been widely accepted by the graphics community and, in fact, generating meaningful random scenes usable for testing algorithms is a major problem. (Note that rather than attempting to generate random scenes, an alternative approach, which has been used to study the average complexity of ray shooting, is to fix the scene and randomly distribute the rays ; see, for example, [ABCC02].)

2.2.8 An upper bound on the average size of silhouettes

The silhouette of a polyhedron with respect to a given viewpoint is, roughly speaking, the set of edges incident to a front and a back face. Silhouettes arise in various problems in computer graphics such as hidden surface removal and shadow computations (see [Dug04, DD02, EGHHZ00] for some recent references) and algorithms to compute them efficiently have been well-studied (see the survey by Isenberg et al. [IFH⁺03]). They are important in shape recognition ; Sander et al. [SGG⁺00] claim that the silhouette “is one of the strongest visual cues of the shape of an object”.

It is a widely accepted fact that the silhouette of a polyhedron is usually much smaller than the whole polyhedron. Sander et al. [SGG⁺00], for instance, state the largely repeated claim that the silhouette of a mesh is often of size $\Theta(\sqrt{n})$ where n is the number of faces of the mesh. An experimental study by Kettner and Welzl [KW97] confirms this for a set of realistic objects. This experimental study was extended by McGuire [McG04] to a larger database of larger objects for which the observed size of the silhouette is approximately $n^{0.8}$.

There are few theoretical results supporting these observations. Kettner and Welzl [KW97] prove that a convex polyhedron that approximates a sphere with Hausdorff distance ϵ has $\Theta(1/\epsilon)$ edges, and a random orthographic projection of such a polytope has $\Theta(1/\sqrt{\epsilon})$ silhouette edges. Alt et al. [AGG03] give conditions under which it can be proved that the average silhouette of a *convex* polyhedron has size $O(\sqrt{n})$ and give additional conditions under which the worst-case size is provably sub-linear.

In Chapter 12, we study the average silhouette size of *non-convex* polyhedra. Convexity is a very strong assumption, which was crucial in the previous theoretical results. Here, rather, we assume that the polyhedron is a good approximation of some fixed (not necessarily convex) surface.

We provide theoretical evidence supporting, for a large class of objects, the widely observed phenomenon that the size of the silhouette of a polyhedron is much smaller than the size of the whole polyhedron. We prove that the silhouette of a polyhedron that approximates a surface in a reasonable way has expected size $O(\sqrt{n})$ where the average is taken over all points of view and n is the complexity of the polyhedron ; the surfaces may be non-convex and non-differentiable and they may have boundaries. Note that the average is taken over all viewpoints for a given surface, and not on a set of surfaces. This confirms a widely accepted belief in computer graphics and is the first complexity result for silhouettes of non-convex polyhedra.

Scene type	Lower bounds	Upper bounds
Segment light source		
2 triangles	4	$O(1)$
2 fat polytopes	$\Omega(n)$	$O(n)$
k polytopes	$\Omega(nk^2 + k^4)$	$O(nk^3)$
n-gon light source		
k polytopes	$\Omega(n^2k^3 + nk^5)$	$O(n^3k^3)$

TAB. 2.2 – Lower bounds on the number of connected components and upper bounds on the complexity of the umbra cast on a plane in the presence of k polytopes of total complexity $O(n)$.

Scene type	Lower bound	Upper bound
k polytopes	$\Omega(n\alpha(k) + km + k^2)$	$O(n\alpha(k) + km\alpha(k) + k^2)$

TAB. 2.3 – Bounds on the complexity of the union of umbra and penumbra cast on a plane by a m -gon light source in the presence of k polytopes of total complexity $O(n)$.

2.2.9 On the complexity of the umbra and penumbra

Shadows play a central role in human perception [MKK98, Wan92]. A wide variety of approaches have been considered for simulating and rendering shadows (see, for example, the surveys [Dur00, WPF90]) and many methods make extensive use of graphics hardware (see the survey [HLHS03]). Unfortunately, computing realistic shadows efficiently is a difficult problem, particularly in the case of non-point light sources. A part of this difficulty arises from the complicated internal structure that such shadows may have. We study, in Chapter 13, this structure.

A point is in the *umbra* if it does not see any part of the light source(s); it is in *full light* if it sees entirely all the light source(s); otherwise, it is in the *penumbra*. While the boundary between the penumbra and the full light is reasonably well-understood (see Section 13.3), less is known about the boundary of the umbra. Nevertheless, there is an extensive literature concerning the explicit computation of these shadow boundaries; see, for example, [DF94, DDP97, DDP99, DDP02, Hec92, NN83, SG94, Tel92].

In Chapter 13, we prove various bounds, summarized in Tables 2.2 and 2.3, on the complexity of the umbra and penumbra cast by a segment or polygonal light source on a plane in the presence of polygon or polytope (*i.e.* convex polyhedral) obstacles. In particular, we show that a single segment light source may cast, in the presence of two triangles, four connected components of umbra (see Figure 2.6). We prove that the umbra defined by one segment light source and two fat convex obstacles of total complexity n can have $\Omega(n)$ connected components. We also prove an $\Omega(nk^2 + k^4)$ lower bound on the maximum number of connected components of the umbra and a $O(nk^3)$ upper bound on its complexity in a scene consisting of a segment light source and k disjoint polytopes of total complexity n . Finally, we prove that the umbra cast on a plane by one (or more) polygonal light source(s) and k convex obstacles of total complexity n can have $\Omega(n^2k^3 + nk^5)$ connected components and has complexity $O(n^3k^3)$. These are the

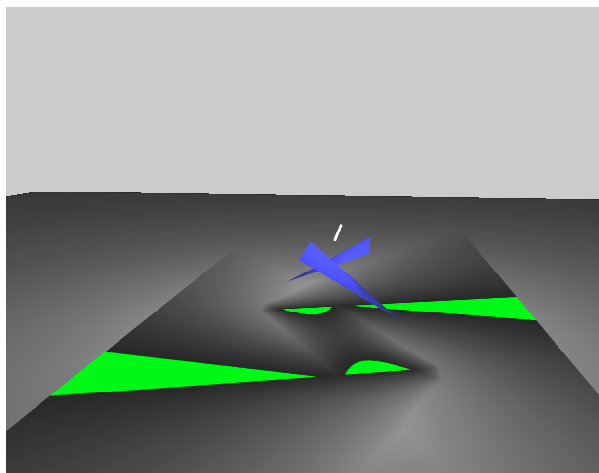


FIG. 2.6 – Four connected components of umbra cast on a plane by a segment light source in the presence of two triangles (the umbra is in light grey) ; courtesy of A. Dietrich.

first bounds on the size of the umbra in terms of both k and n .

These results are surprising in the sense that they show that the umbra cast by a single segment light source may have many connected components. The fact that the umbra may have four connected components in the case of two triangle obstacles comes as a total surprise. Our lower bounds of $\Omega(nk^2 + k^4)$ and $\Omega(n^2k^3 + nk^5)$ connected components, for k polytopes of total complexity n , is rather pathological in the sense that most of the obstacles are very long and thin. However, we also present a lower bound example of $\Omega(n)$ connected components with two fat polygons or polytopes of complexity $O(n)$. Concerning the upper bounds of $O(nk^3)$ and $O(n^3k^3)$, even though these bounds are not a priori tight, they substantially improve the only previously known bounds for this problem which were the trivial $O(n^4)$ and $O(n^6)$ upper bounds. It is also interesting to notice that the proof of the upper bound for a polygonal light source uses, as a key lemma, the Main Lemma of Chapter 9 on the number of lines intersecting a given line and tangent to three polytopes. Finally, it is interesting to point out that even for the simplest case of non-point light sources, obtaining tight bounds on the complexity of the umbra and understanding its structure is a very challenging problem.

These results show that the umbra, which is bounded by arcs of conics, is intrinsically much more intricate than the full light/penumbra boundary which is bounded by line segments and for which we prove that the worst-case complexity is in $\Omega(n\alpha(k) + km + k^2)$ and $O(n\alpha(k) + km\alpha(k) + k^2)$, where m is the complexity of the polygonal light source.

Chapitre 3

Geometric computing with low-degree algebraic surfaces

3.1 Introduction

Low-degree real algebraic surfaces such as quadrics allow for a good compromise between simplicity, flexibility, and modeling power and they play a leading role in the construction of accurate computer models of physical environments for simulation and prototyping purposes. As a consequence, geometric computing with curved surfaces has a long and rich history. Nevertheless, our recent work on the exact and optimal parameterization of intersections of quadrics (see Chapters 14, 15, 16, 17) has proved that, even on well-studied problems, startling improvements can be accomplished when the right mathematics are put into motion. Moreover, the recent advances in exact computation and the constant improvement of the available algebraic machinery have opened new alleys of research for performing, in this domain, computations with certified topology. In particular, the performance of computer algebra systems has dramatically improved over the past few years and we are now able to solve non-trivial problems (see Chapter 18). However, despite recent success, computer algebra systems remain computationally intensive and using them efficiently is more of an art than a science. Experience has also proved that having a thorough geometric understanding of the problem before resorting to such systems greatly increases the chances of success. Still, a lot of work remains to be done to obtain exact scalable geometric algorithms dealing with (low-degree) curved objects.

In the last few years, I focused my research activities on certified and effective computations with quadrics. I have mainly worked on the problems of computing the intersection of two quadrics and, more generally, on the problems of computing quadratic complexes, (*i.e.*, piecewise quadratic surfaces) in the context of boundary evaluation of solids. I have also worked on Voronoi diagrams of polyhedral objects. I present below a concise presentation of my contributions on the intersection of quadrics and on Voronoi diagrams of lines. I do not however discuss in this document my work on computing quadratic complexes, which is still ongoing and unpublished.

3.2 Summary of contributions

3.2.1 Introduction

My main achievement has been the completion of the first-ever exact, complete, near-optimal and usable algorithm and implementation for parameterizing the intersection of two quadrics in three-dimensional real projective space (see Chapters 14, 15, 16, and 17). This algorithm is a considerable breakthrough on a long-standing open question and, though there is still room for minor improvements, the problem is now largely closed.

More precisely, we presented (see Chapter 15) the first classification of pairs of quadrics based on the type of their intersection in real projective space (*e.g.*, smooth quartic, nodal quartic, a cubic and a line, two conics, etc.). Using this classification, we presented the first practical algorithm that correctly identifies, separates and parameterizes all the algebraic components of the intersection and gives all the relevant topological information (see Chapters 14 and 16). The parameterization is moreover simple in the sense that (i) the parameterization is rational whenever one exists (otherwise the intersection is a non-singular quartic and the parameterization involves the square root of a polynomial) and (ii) the number of square roots appearing in the coefficients of the parameterization is always minimal, except in a small number of well-identified cases, in which the coefficients may involve one extra square root. We have also implemented this algorithm in C++ and we showed that our implementation is extremely efficient in practice, on generic, degenerate as well as “real-life” data (see Chapter 17); it can routinely compute parameterizations of the intersection of quadrics with input coefficients having 50 digits in less than 50 milliseconds on a mainstream PC. Our code can be downloaded from the LORIA and INRIA web sites² and it can also be queried via a web interface³.

I have also recently started working on the problem of computing the medial axis or Voronoi diagram of polyhedra in 3D. Such a diagram is a partition of space into cells, each of which consists of the points closest to one particular object than to any other. Moreover, the set of points equidistant to two lines (or to a line and a point) is a quadric and the set of points equidistant to three lines is the intersection of two quadrics. Chapter 18 presents some very nice results characterizing the topology of the Voronoi diagrams of three lines. We show that the topology is invariant for lines in general position and we obtain a monotonicity property on the arcs of the diagram. We deduce a simple algorithm for sorting points along such an arc, which is presumably of great interest for future efficient algorithms for computing the medial axis of a polyhedron. The proof technique, which relies heavily upon modern tools of computer algebra, is also of great interest in its own right.

I now present, in Sections 3.2.2 and 3.2.3, a more detailed description of these results.

3.2.2 Intersection of two quadrics

I present here a detailed, though concise, overview of my contributions on the problem of intersecting two quadrics. I start by presenting a panorama of relevant past work on the subject. Then, after some preliminaries, I briefly recall Levin’s seminal algorithm for parameterizing

²<http://www.loria.fr>, <http://www.inria.fr>

³<http://www.loria.fr/equipes/vegas/qi>

the intersection of two quadrics and show how this algorithm introduces high-degree algebraic numbers which induce robustness problems. I then present my contributions.

Section 3.2.2.4 first presents a major improvement over Levin's algorithm that computes a *near-optimal* parameterization of the intersection of two quadrics in the generic case, that is when the intersection is a non-singular quartic (see Figure 3.1); the parameterization is near-optimal in the sense that, when the intersection is a non-singular quartic, the number of square roots appearing in the coefficients of the parameterization is minimal or minimal plus one. Then, Section 3.2.2.5 presents the first classification of pairs of quadrics based on the type of their intersection in real projective space (*e.g.*, smooth quartic, nodal quartic, a cubic and a line, two conics, etc.). Based on this classification, Section 3.2.2.6 presents near-optimal algorithms for all possible types of singular intersection (see Figure 3.1). This induces a practical algorithm that correctly identifies, separates and computes near-optimal parameterizations of all the algebraic components of the intersection, for all possible type of intersection. Finally, Section 3.2.2.7 presents an efficient exact implementation of this algorithm.

3.2.2.1 Panorama

The simplest of all the curved surfaces, quadrics (*i.e.*, algebraic surfaces of degree two), are fundamental geometric objects, arising in such diverse contexts as geometric modeling, statistical classification, pattern recognition, and computational geometry. Computing the intersection of two general quadrics is a fundamental problem and an exact parametric representation of the intersection is often desirable. For instance, it is at the basis of such complex geometric operations as computing convex hulls of quadric patches [HI95], arrangements of sets of quadrics [BHK⁺05, MTT05, SW06, Wol02], and boundary representations of quadric-based solid models [KCF⁺04, Sar83].

Until recently, the only known general method for computing a parametric representation of the intersection between two arbitrary quadrics was due to Levin [Lev76, Lev79]. It is based on an analysis of the pencil generated by the two quadrics, *i.e.* the set of linear combinations of the two quadrics.

Though useful, Levin's method has serious limitations. When the intersection is singular or reducible, a parameterization by rational functions is known to exist, but Levin's pencil method fails to find it and generates a parameterization that involves the square root of some polynomial. In addition, when a floating point representation of numbers is used, Levin's method sometimes outputs results that are topologically wrong and it may even fail to produce any parameterization at all and crash. On the other hand a correct implementation using exact arithmetic is essentially out of reach because Levin's method introduces algebraic numbers of fairly high degree. A good indication of this impracticality is that even for the simple generic example of Section 14.8.2, an exact parametric form output by Levin's algorithm (computed by hand with Maple) fills up over 100 megabytes of space !

Over the years, Levin's seminal work has been extended and refined in several different directions. Wilf and Manor [WM93] use a classification of quadric intersections by the Segre characteristic (see [Bro06]) to drive the parameterization of the intersection by the pencil method. Recently, Wang, Goldman and Tu [WGT03] further improved the method by making it capable of computing structural information on the intersection and its various connected components and able to produce a parameterization by rational functions when it exists. Whether

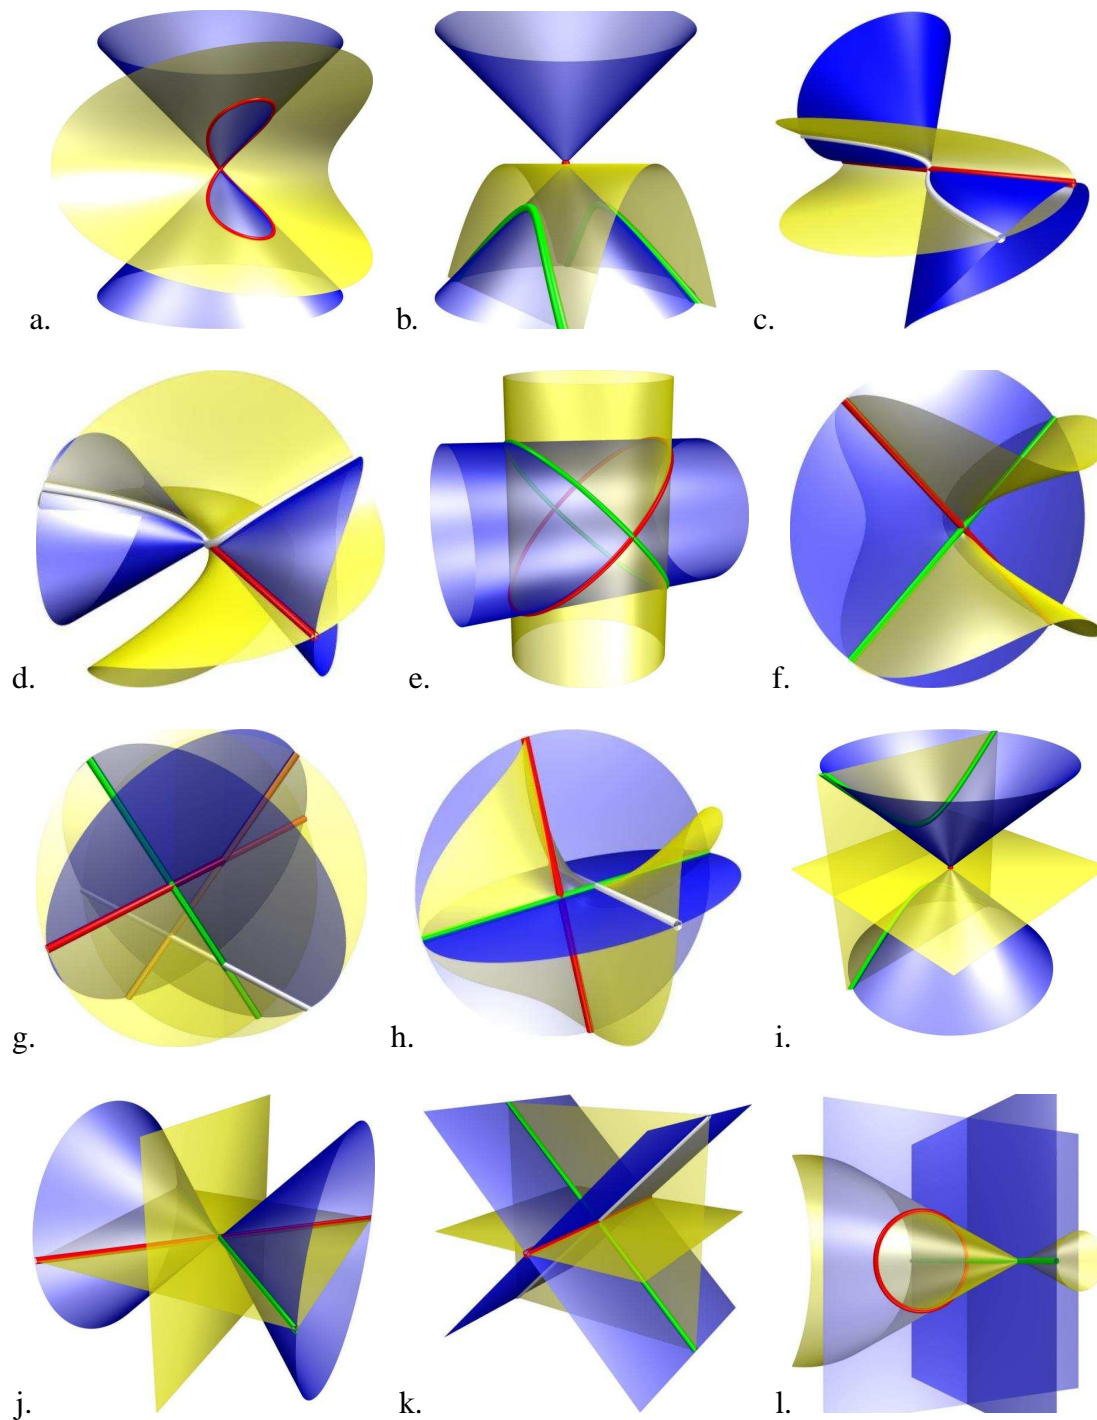


FIG. 3.1 – A gallery of intersections. a. Nodal quartic. b. Nodal quartic with isolated singular point. c. Cubic and secant line. d. Cubic and tangent line. e. Two secant conics. f. Two double lines. g. Four lines forming a skew quadrilateral. h. Two lines and a double line. i. Conic and two lines not crossing on the conic, the two lines being imaginary. j. Four concurrent lines, only two of which are real. k. Two lines and a double line, the three being concurrent. l. Conic and double line.

their refined algorithm is numerically robust is open to question.

Another method of algebraic flavor was introduced by Farouki, Neff and O'Connor [FNO89] when the intersection is degenerate. In such cases, using a combination of classical concepts (Segre characteristic) and algebraic tools (factorization of multivariate polynomials), the authors show that explicit information on the morphological type of the intersection curve can be reliably obtained. A notable feature of this method is that it can output an exact parameterization of the intersection in simple cases, when the input quadrics have rational coefficients. No implementation is however reported.

Rather than restricting the type of the intersection, others have sought to restrict the type of the input quadrics, taking advantage of the fact that geometric insights can then help compute the intersection curve [GM91, Mil87, MG95, SJ92, SJ94]. Specialized routines are devised to compute the intersection curve in each particular case. Even though such geometric approaches are numerically more stable than the algebraic ones, they are essentially limited to the class of so-called natural quadrics (i.e., the planes, right cones, circular cylinders and spheres) and planar intersections.

Perhaps the most interesting of the previously known algorithms for computing an explicit representation of the intersection of two arbitrary quadrics is the method of Wang, Joe and Goldman [WJG02]. This algebraic method is based on a birational mapping between the intersection curve and a plane cubic curve. The cubic curve is obtained by projection from a point lying on the intersection. Then the classification and parameterization of the intersection are obtained by invoking classical results on plane cubics. The authors claim that their algorithm is the first to produce a complete topological classification of the intersection (singularities, number and types of connected components, etc.). However, the computation of the center of projection uses (an enhanced version of) Levin's algorithm. Either floating point arithmetic is used and the point will in general not exactly lie on the curve, leading to possibly incorrect classification, or exact arithmetic is used and the parameterizations computed will involve algebraic numbers of very high degree, thereby limiting their practical value.

A related problem to computing the intersection of two quadrics is to classify pencils of quadrics based on the type of their intersection. Classifying pencils of quadrics over the complexes (that is based on the type of the intersection in $\mathbb{P}^3(\mathbb{C})$) was achieved by Segre in the nineteenth century [Seg83]. Its practical value is however limited since its proper interpretation lies in the complex domain (i.e. points on the intersection might be real or complex), whereas our concern is with real parts of the intersection.

Classifying pencils of quadrics over the reals, that is based on the type of the intersection in $\mathbb{P}^3(\mathbb{R})$, is an important problem of independent interest. Such a classification is also critical for parameterizing the intersection of two quadrics for two reasons. It induces some structural information on the intersection curve which can be used to drive an algorithm for computing a near-optimal parameterization of the intersection curve (see Chapter 16); it is also a prerequisite for proving the (near-)optimality of a parameterization algorithm (see Chapter 16). I now recall past work on the classification of pencils of quadrics over the reals.

In the context of the representation of the geometry of Boolean combinations of volumes bounded by quadric surfaces, Ocken, Schwartz, and Sharir [OSS87] showed in 1987 how two quadrics can be simultaneously diagonalized using a real projective transformation and used this diagonalization to parameterize the intersection of the quadrics. The analysis is however incomplete and some intersection morphologies are overlooked, leading to possible misclassifi-

cations. In particular, the cases when the characteristic polynomial of the pencil has two double roots, corresponding to such morphologies as a cubic and a secant line or four lines forming a skew quadrilateral, are missing.

The next result on the classification of pencil of quadrics based on the real type of the intersection was obtained in 2002 by Tu, Wang, and Wang [TWW02] who classified pencils in the generic case, that is when the intersection is a smooth quartic (in complex space) whose number of connected components (in real space) is two, one, or zero. Note that Wang and Krasauskas [WK04] also obtained results on the classification of pencils in the generic case when the pencil is furthermore restricted to be generated by two ellipsoids in affine space. Related results have also been obtained by Wang, Wang, and Kim on the separation of two ellipsoids in affine space [WWK01].

In September 2005, Tu, Wang, Mourrain, and Wang published a research report [TWMW05] presenting a classification of pencils very similar to the one of Chapter 15 (also published in Dupont's PhD Thesis in 2004 [Dup04]). They use the Canonical Pair Form Theorem of F. Uhlig as basic mathematical tool and refine the classification of pencils of quadrics over the complexes in exactly the same way as we do. There are however differences between the two approaches. First, we classify pencils using the inertia of the quadrics at the multiple roots of the characteristic polynomial, except for a small number of cases where simple geometric conditions allow to discriminate. By contrast, Tu et al. classify pencils using the inertia of the quadrics between the roots of the characteristic polynomial (plus the degree of the minimal polynomial of the characteristic polynomial in some cases), and rely on Puiseux expansion to deduce some information at the (multiple) roots. Second, the classification of Tu et al. is limited to non-degenerate pencils (i.e. pencils whose characteristic polynomial does not vanish identically), while ours covers all possible cases. Third, in addition to the enumeration of all real quadric intersection morphologies, we also provide algorithms for exactly and efficiently recovering the real projective type of the intersection of two arbitrary given quadrics.

3.2.2.2 Preliminaries

We call a *quadric* associated to a four-by-four symmetric real matrix S the set $Q_S = \{\mathbf{x} \in \mathbb{P}^3 \mid \mathbf{x}^T S \mathbf{x} = 0\}$, where $\mathbb{P}^n = \mathbb{P}(\mathbb{R})^n$ denotes the real projective space of dimension n . (Note that every matrix of the form αS , where $\alpha \in \mathbb{R} \setminus \{0\}$, represents the same quadric Q_S .) When the ambient space is \mathbb{R}^n instead of $\mathbb{P}(\mathbb{R})^n$, the quadric is simply Q_S minus its points at infinity.

Given a real symmetric matrix S of size $n + 1$, the upper left submatrix of size n , denoted S_u , is called the *principal submatrix* of S and the determinant of S_u the *principal subdeterminant* of S .

Matrix S being symmetric, all of its eigenvalues are real. Let σ^+ and σ^- be the numbers of positive and negative eigenvalues of S , respectively. The *rank* of S is the sum of σ^+ and σ^- . We define the *inertia* of S and Q_S as the pair $(\max(\sigma^+, \sigma^-), \min(\sigma^+, \sigma^-))$. (Note that it is more usual to define the inertia as the pair (σ^+, σ^-) , but our definition, in a sense, reflects the fact that Q_S and Q_{-S} are one and the same quadric.) Matrix S and quadric Q_S are called *singular* if the determinant of S is zero; otherwise they are called *nonsingular*.

The inertia of a quadric in \mathbb{P}^3 is a fundamental concept which somehow replaces the usual type of a quadric in \mathbb{R}^3 ; Table 3.1 recalls the correspondence between inertias in \mathbb{P}^3 and types in \mathbb{R}^3 .

Inertia of Q_S	Inertia of S_u	Euclidean canonical equation	Euclidean type of Q_S
(4, 0)	(3, 0)	$x^2 + y^2 + z^2 + 1$	\emptyset (imaginary ellipsoid)
(3, 1)	(3, 0)	$x^2 + y^2 + z^2 - 1$	ellipsoid
	(2, 1)	$x^2 + y^2 - z^2 + 1$	hyperboloid of two sheets
	(2, 0)	$x^2 + y^2 + z$	elliptic paraboloid
(3, 0)	(3, 0)	$x^2 + y^2 + z^2$	point
	(2, 0)	$x^2 + y^2 + 1$	\emptyset (imaginary elliptic cylinder)
(2, 2)	(2, 1)	$x^2 + y^2 - z^2 - 1$	hyperboloid of one sheet
	(1, 1)	$x^2 - y^2 + z$	hyperbolic paraboloid
(2, 1)	(2, 1)	$x^2 + y^2 - z^2$	cone
	(2, 0)	$x^2 + y^2 - 1$	elliptic cylinder
	(1, 1)	$x^2 - y^2 + 1$	hyperbolic cylinder
	(1, 0)	$x^2 + y$	parabolic cylinder
(2, 0)	(2, 0)	$x^2 + y^2$	line
	(1, 0)	$x^2 + 1$	\emptyset (imaginary parallel planes)
(1, 1)	(1, 1)	$x^2 - y^2$	intersecting planes
	(1, 0)	$x^2 - 1$	parallel planes
	(0, 0)	x	simple plane
(1, 0)	(1, 0)	x^2	double plane
	(0, 0)	1	\emptyset (double plane at infinity)

TAB. 3.1 – Correspondence between quadric inertias and Euclidean types.

In \mathbb{P}^3 , any quadric not of inertia (3, 1) is either a ruled surface or not a surface. Also, the quadrics of inertia (3, 1) are the only ones with a strictly negative determinant. The nonsingular quadrics are those of rank 4, i.e. those of inertia (4, 0), (3, 1) and (2, 2). Quadrics of inertia (4, 0) are however empty of real points. A quadric of rank 3 is called a *cone*. The cone is said to be *real* if its inertia is (2, 1). It is said to be *imaginary* otherwise, in which case its real projective locus is limited to its singular point. A quadric of rank 2 is a *pair of planes*. The pair of planes is real if its inertia is (1, 1). It is called imaginary if its inertia is (2, 0), in which case its real projective locus consists of its singular line, i.e. the line of intersection of the two planes. A quadric of inertia (1, 0) is called a *double plane* and is necessarily real.

Let S and T be two real symmetric matrices of the same size and let $R(\lambda, \mu) = \lambda S + \mu T$. The set

$$\{R(\lambda, \mu) \mid (\lambda, \mu) \in \mathbb{P}^1\}$$

is called the *pencil of matrices* generated by S and T . For the sake of simplicity, we sometimes write a member of the pencil $R(\lambda) = \lambda S - T$, $\lambda \in \overline{\mathbb{R}} = \mathbb{R} \cup \{\infty\}$. Associated to it is a *pencil of quadrics* $\{Q_{R(\lambda, \mu)} \mid (\lambda, \mu) \in \mathbb{P}^1\}$. Recall that the intersection of two distinct quadrics of a pencil is independent of the choice of the two quadrics. We call the binary form

$$\mathcal{D}(\lambda, \mu) = \det R(\lambda, \mu)$$

quadric	canonical equation ($a, b > 0$)	parameterization $\vec{X} = [x, y, z]$, $u, v \in \mathbb{R}$
simple plane	$x = 0$	$\vec{X}(u, v) = [0, u, v]$
double plane	$x^2 = 0$	$\vec{X}(u, v) = [0, u, v]$
parallel planes	$ax^2 = 1$	$\vec{X}(u, v) = [\frac{1}{\sqrt{a}}, u, v]$, $\vec{X}(u, v) = [-\frac{1}{\sqrt{a}}, u, v]$
intersecting planes	$ax^2 - by^2 = 0$	$\vec{X}(u, v) = [\frac{u}{\sqrt{a}}, \frac{u}{\sqrt{b}}, v]$, $\vec{X}(u, v) = [\frac{u}{\sqrt{a}}, -\frac{u}{\sqrt{b}}, v]$
hyperbolic paraboloid	$ax^2 - by^2 - z = 0$	$\vec{X}(u, v) = [\frac{u+v}{2\sqrt{a}}, \frac{u-v}{2\sqrt{b}}, uv]$
parabolic cylinder	$ax^2 - y = 0$	$\vec{X}(u, v) = [u, au^2, v]$
hyperbolic cylinder	$ax^2 - by^2 = 1$	$\vec{X}(u, v) = [\frac{1}{2\sqrt{a}}(u + \frac{1}{u}), \frac{1}{2\sqrt{b}}(u + \frac{1}{u}), v]$

TAB. 3.2 – Parameterizations of canonical simple ruled quadrics [Lev76].

the *characteristic polynomial* of the pencil.

3.2.2.3 Levin’s method

Levin’s quadric intersection algorithm [Lev76, Lev79] being seminal in almost all work on intersection of quadrics, including ours, we start by recalling the main steps of his algorithm for computing a parameterized representation of the intersection of two distinct implicit quadrics Q_S and Q_T of \mathbb{R}^3 . Starting from this short description, we then identify where this algorithm introduces high-degree algebraic numbers and why this is a problem.

The high-level idea behind Levin’s algorithm is this : if (say) Q_S is of some “good” type, then Q_S admits a parameterization which is linear in one of its parameters and plugging this parameterization in the implicit equation of Q_T yields a degree 2 equation in one of the parameters (instead of a degree 4 equation) which can be easily solved to get a parametric representation of $Q_S \cap Q_T$. When neither Q_S nor Q_T has a “good” type, then one can find a quadric Q_R of “good” type in the pencil generated by Q_S and Q_T , and we are back to the previous case replacing Q_S by Q_R .

The definition of a “good” type is embodied in Levin’s notion of simple ruled quadric⁴ and the existence of such a quadric Q_R is Levin’s key result :

Theorem 3.1 ([Lev76]). *The pencil generated by any two distinct quadrics contains at least one simple ruled quadric, i.e., one of the quadrics listed in Table 3.2, or the empty set.*

In more details, Levin’s method is as follows.

1. Find a simple ruled quadric in the pencil $\{Q_{R(\lambda)} = \lambda Q_S - T \mid \lambda \in \overline{\mathbb{R}}\}$ generated by Q_S and Q_T , or report an empty intersection. Since simple ruled quadrics have a vanishing principal subdeterminant, this is achieved by searching for a $\lambda_0 \in \overline{\mathbb{R}}$ such that $\det(R_u(\lambda_0)) = 0$ and $Q_R = Q_{R(\lambda_0)}$ is simple ruled ; by Theorem 3.1, such a quadric exists or the pencil contains the empty set. Assume, for the sake of simplicity, that the intersection is not empty and that Q_R and Q_S are distinct. Then $Q_S \cap Q_T = Q_S \cap Q_R$.

⁴In [Lev76, Lev79], Levin refers to these quadrics as to nonelliptic paras.

2. Determine the orthonormal transformation matrix P_u which sends R_u in diagonal form by computing the eigenvalues and the normalized eigenvectors of R_u . Deduce the transformation matrix P which sends Q_R into canonical form. In the orthonormal frame in which it is canonical, Q_R admits one of the parameterizations \vec{X} of Table 3.2.
3. Compute the matrix $S' = P^T S P$ of the quadric Q_S in the canonical frame of Q_R and consider the equation

$$\vec{X}^T S' \vec{X} = a(u)v^2 + b(u)v + c(u) = 0, \quad (3.1)$$

where \vec{X} has been augmented by a fourth coordinate set to 1. (The parameterizations of Table 3.2 are such that $a(u)$, $b(u)$ and $c(u)$ are polynomials of degree at most four in u .)

Solve (3.1) for v in terms of u and determine the corresponding domain of validity of u on which the solutions are defined, i.e., the set of u such that $\Delta(u) = b^2(u) - 4a(u)c(u) \geq 0$. Substituting v by its expression in terms of u in \vec{X} , we have a parameterization of $Q_S \cap Q_T = Q_S \cap Q_R$ in the orthonormal coordinate system in which Q_R is canonical.

4. Output $P\vec{X}(u)$, the parameterized equation of $Q_S \cap Q_T$ in the global coordinate frame, and the domain of $u \in \mathbb{R}$ on which it is valid.

This method is very nice and powerful since it gives an explicit representation of the intersection of two general quadrics. However, it is far from being ideal from the point of view of precision and robustness since it introduces non-rational numbers at several different places. Thus, if a floating point representation of numbers is used, the result may be wrong (geometrically and topologically) or, worse, the program may crash (especially in Step 1 when the type of the quadrics $Q_{R(\lambda_0)}$ are incorrectly computed). In theory, exact arithmetic would do, except that it would seriously slow down the computations. In practice, however, a correct implementation using exact arithmetic seems out of reach because of the high degree of the algebraic numbers involved.

Let us examine more closely the potential sources of numerical instability in Levin's algorithm.

- Step 1 : λ_0 is the root of a third degree polynomial with rational coefficients. In the worst case, it is thus expressed with nested radicals of depth two. Since determining if $Q_{R(\lambda_0)}$ is simple ruled involves computing its Euclidean type (not an easy task considering that $Q_{R(\lambda_0 - \varepsilon)}$ and $Q_{R(\lambda_0 + \varepsilon)}$ may be and often are of different types), this is probably the biggest source of non-robustness.
- Step 2 : Since Q_R is simple ruled, the characteristic polynomial of R_u is a degree three polynomial having zero as a root and whose coefficients are in the field extension $\mathbb{Q}(\lambda_0)$. Thus, the nonzero eigenvalues of R_u may involve nested radicals of depth three. Since the corresponding eigenvectors have to be normalized, the coefficients of the transformation matrix P are expressed with radicals of nesting depth four in the worst case. Since the coefficients of the parameterization \vec{X} of Q_R are expressed as square roots of the coefficients of the canonical equation Q_{PTRP} (as in Table 3.2), the coefficients of the parameterization of $Q_S \cap Q_T$ can involve *nested radicals of depth five* in the worst case.
- Step 3 : Computing the domain of \vec{X} amounts to solving the fourth degree equation $\Delta(u) = \overline{0}$ whose coefficients are nested radicals of worst-case depth five in \mathbb{Q} .

Note that this worst-case picture is the generic case. Indeed, given two arbitrary quadrics with rational coefficients, the polynomial $\det(R_u(\lambda))$ will generically have no rational root (a consequence of Hilbert's Irreducibility Theorem).

inertia of S	canonical equation ($a, b, c, d > 0$)	parameterization $\vec{X} = [x, y, z, w]$
(4, 0)	$ax^2 + by^2 + cz^2 + dw^2 = 0$	$Q_S = \emptyset$
(3, 0)	$ax^2 + by^2 + cz^2 = 0$	Q_S is point $(0, 0, 0, 1)$
(2, 2)	$ax^2 + by^2 - cz^2 - dw^2 = 0$	$\vec{X} = [\frac{ut+avs}{a}, \frac{us-bvt}{b}, \frac{ut-avs}{\sqrt{ac}}, \frac{us+bvt}{\sqrt{bd}}], (u, v), (s, t) \in \mathbb{P}^1$
(2, 1)	$ax^2 + by^2 - cz^2 = 0$	$\vec{X} = [uv, \frac{u^2-abv^2}{2b}, \frac{u^2+abv^2}{2\sqrt{bc}}, s], (u, v, s) \in \mathbb{P}^{*2}$
(2, 0)	$ax^2 + by^2 = 0$	$\vec{X} = [0, 0, u, v], (u, v) \in \mathbb{P}^1$
(1, 1)	$ax^2 - by^2 = 0$	$\vec{X}_1 = [u, \frac{\sqrt{ab}}{b}u, v, s], \vec{X}_2 = [u, -\frac{\sqrt{ab}}{b}u, v, s], (u, v, s) \in \mathbb{P}^2$
(1, 0)	$ax^2 = 0$	$\vec{X} = [0, u, v, s], (u, v, s) \in \mathbb{P}^2$

TAB. 3.3 – Parameterization of projective quadrics of inertia different from (3, 1). In the parameterization of projective cones, \mathbb{P}^{*2} stands for the 2-dimensional real quasi-projective space defined as the quotient of $\mathbb{R}^3 \setminus \{0, 0, 0\}$ by the equivalence relation \sim where $(x, y, z) \sim (y_1, y_2, y_3)$ iff $\exists \lambda \in \mathbb{R} \setminus \{0\}$ such that $(x, y, z) = (\lambda y_1, \lambda y_2, \lambda^2 y_3)$.

3.2.2.4 The generic algorithm

We present, in Chapter 14, a first but major improvement to Levin’s pencil method for computing parametric representations of the intersection of quadrics. This so-called “generic algorithm” removes most of the sources of radicals in Levin’s algorithm. We prove that this algorithm produces near-optimal parameterizations in the generic case, that is when the intersection is a non-singular quartic. It is however not optimal for all the possible types of intersection and will need later refinements (see Chapters 15 and 16). But it is sufficiently simple, robust and efficient to be of interest to many. We give here a short description of this algorithm and of the underlying ideas.

We start by introducing the projective framework underlying our approach and stating the main theorem on which the approach rests. From now on, all the input quadrics considered have equations with rational coefficients.

Key ideas. The first ingredient of our approach is to work not just over \mathbb{R}^3 but over the real projective space \mathbb{P}^3 . Recall that, in projective space, quadrics are characterized by their inertia (i.e., the number of positive and negative eigenvalues of the associated 4×4 matrix), while in Euclidean space they are characterized by their inertia and the inertia of their principal submatrix (see Table 3.1).

In our algorithm, quadrics of inertia *different* from (3, 1) (i.e., ruled quadrics) play the role of simple ruled quadrics in Levin’s method. In Table 3.3, we present a new set of faithful⁵ parameterizations of ruled projective quadrics that are both linear in one of their parameters and involve, in the worst case, a minimal number of square roots⁶.

⁵A parameterization is faithful if there is a one-to-one correspondence between the points of the quadric and the parameters.

⁶Note that there is necessarily a trade-off between the minimal degree of a parameterization in one of its parameters and the degree of its coefficient field. For instance, Wang, Joe and Goldman [WJG97] give parameterizations

Another key ingredient of our approach is encapsulated in the following theorem, which mirrors, in the projective setting, Levin's theorem on the existence of ruled quadrics in a pencil.

Theorem 3.2. *In a pencil generated by any two distinct quadrics, the set \mathcal{S} of quadrics of inertia different from $(3, 1)$ is not empty. Furthermore, if no quadric in \mathcal{S} has rational coefficients, then the intersection of the two initial quadrics is reduced to two distinct points.*

This theorem, which we prove using Uhlig's results [Uhl73, Uhl76] on simultaneous block diagonalizations of pairs of symmetric real matrices, generalizes Theorem 3.1. Indeed, it ensures that the two quadrics we end up intersecting have rational coefficients, except in one very specific situation where they intersect in only two points. This is how we remove the main source of nested radicals in Levin's algorithm.

The last basic ingredient of our approach is the use of Gauss reduction of quadratic forms for diagonalizing a symmetric matrix and computing the canonical form of the associated projective quadric, instead of the traditional eigenvalues/eigenvectors approach used by Levin. Since Gauss transformation is rational (the elements of the matrix P which sends S into canonical form are rational), this removes some layers of nested radicals from Levin's algorithm. Note, also, that there is no difficulty in parameterizing the reduced quadric $S' = P^T S P$ since, by Sylvester's Inertia Law, S and S' have the same inertia [Lam73].

Algorithm outline. We can now outline our generic algorithm.

Let $R(\lambda) = \lambda S - T$ be the pencil generated by the quadrics Q_S and Q_T of \mathbb{P}^3 and $\mathcal{D}(\lambda) = \det(R(\lambda))$ be the characteristic polynomial of the pencil. Recall that, although working in all cases, our generic algorithm is best designed when $\mathcal{D}(\lambda)$ is not identically zero and does not have any multiple root. In the other case, a better algorithm is described in Chapters 15 and 16. The outline of the intersection algorithm is as follows :

1. Find a quadric Q_R with rational coefficients in the pencil, such that $\det R > 0$ if possible or $\det R = 0$ otherwise. (If no such R exists, the intersection is reduced to two points, which we output.) If the inertia of R is $(4, 0)$, output empty intersection. Otherwise, proceed.

Assume for the sake of simplicity that $Q_S \neq Q_R$, in such a way that $Q_S \cap Q_R = Q_S \cap Q_T$.

2. If the inertia of R is not $(2, 2)$, apply Gauss reduction to R and compute a frame in which $P^T R P$ is diagonal.

If the inertia of R is $(2, 2)$, its parameterization of Table 3.3 contains in general two square roots but one can be eliminated as follows. First, find a rational point close enough to Q_R such that the quadric in the pencil through this point has the same inertia as Q_R . Replace Q_R by this quadric. Then use that rational point to compute a frame in which $P^T R P$ is the diagonal matrix $\text{diag}(1, 1, -1, -\delta)$, with $\delta \in \mathbb{Q}$.

In the local frame, Q_R can be described by one of the parameterizations \mathbf{X} of Table 3.3. Compute the parameterization $P\mathbf{X}$ of Q_R in the global frame.

3. Consider the equation

$$\Omega : (P\vec{X})^T S (P\vec{X}) = 0. \quad (3.2)$$

Equation Ω is of degree at most 2 in (at least) one of the parameters. Solve it for this parameter in terms of the other(s) and compute the domain of the solution.

of quadrics that have rational coefficients but are quadratic in all of their parameters.

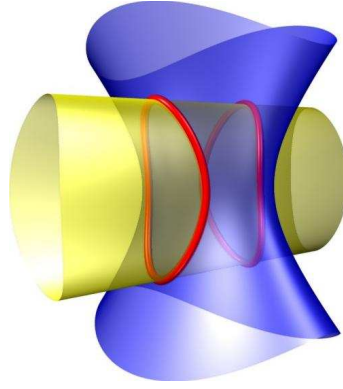


FIG. 3.2 – Intersection of an elliptic cylinder and a hyperboloid of one sheet in a smooth quartic.

4. Substitute this parameter in $P\vec{X}$, giving a parameterization of the intersection of Q_S and Q_T .

Example. It should be stressed that even though our algorithm produces an *exact* parameterization of the intersection of two quadrics, it tends toward producing “simple” parameterizations. It is, in particular, interesting to compare the exact parameterizations produced by our algorithm to the approximated parameterizations produced by other algorithms. We give here one example (see Chapters 14 and 17 for more).

Consider Example 4 from [WJG02], that is the intersection of an elliptic cylinder and a hyperboloid of one sheet (see Figure 3.2) of equations :

$$\begin{cases} 4x^2 + z^2 - w^2 = 0, \\ x^2 + 4y^2 - z^2 - w^2 = 0. \end{cases}$$

In [WJG02], the authors find the following parameterization for the curve of intersection :

$$\mathbf{X}(u) = \begin{pmatrix} 0.0 \\ 1131.3708u^3 - 5760.0u^2 + 10861.1602u - 8192.0 \\ -1600.0u^3 + 10861.1602u^2 - 21504.0u + 11585.2375 \\ 1600.0u^3 + 3620.2867u^2 + 5120.0u + 11585.2375 \end{pmatrix} \pm \begin{pmatrix} -80.0u + 1181.0193 \\ 0.0 \\ 0.0 \\ 0.0 \end{pmatrix} \sqrt{905.0967u^3 - 3328.0u^2 + 2896.3094u} \quad (3.3)$$

with u in the closure of \mathbb{R} and such that the square root is defined. The authors report a computation error on this example (measured as the maximum distance from a sequence of sample points on the curve to the input quadrics) of order $O(10^{-7})$. In comparison, our algorithm outputs the following exact and simple result (in less than 10 ms on a standard PC) :

$$\mathbf{X}(u) = \begin{pmatrix} 2u^3 - 6u \\ 7u^2 + 3 \\ 10u^2 - 6 \\ 2u^3 + 18u \end{pmatrix} \pm \begin{pmatrix} -2 \\ u \\ 2u \\ 2 \end{pmatrix} \sqrt{-3u^4 + 26u^2 - 3}.$$

Conclusion. This algorithm already represents a substantial improvement over Levin’s pencil method and its subsequent refinements. Indeed, we prove that, when the intersection is a smooth quartic (the generic case) the algorithm computes a parameterization which is optimal in the

number of radicals involved up to one possibly unnecessary square root. We also show that deciding (in all cases) whether this extra square root can be avoided is out of reach⁷, and that the parameterization is optimal in some cases. Moreover, for the first time, our algorithm enables to compute in practice an exact form of the parameterization of two arbitrary quadrics with rational coefficients.

Even though this algorithm focuses on the generic, smooth quartic case, it can also be used when the intersection is singular. Unfortunately, this does not always lead to a parameterization of the intersection that involves only rational functions. In particular, when the intersection C is a singular quartic, Ω is irreducible since C itself is, and solving Ω in one of its variables introduces the square root of a polynomial, while we know that there exists a parameterization of C with rational functions (since C is of genus zero).

Always computing parameterizations with rational functions when such parameterizations are known to exist will necessitate rethinking the basic philosophy of our algorithm. Essentially, while the idea of the generic algorithm is to use a rational quadric Q_R with *largest* rank as intermediate quadric for parameterizing the intersection, the refined method will instead use the rational quadric with *smallest* rank as intermediate quadric.

Proceeding that way will have the double benefit of always computing the simplest possible parameterizations and much better controlling the size of their coefficients. The price we pay is a multiplicity of cases and the need to write dedicated software for each (real projective) type of intersection. This is the subject of Sections 3.2.2.5 and 3.2.2.6 and of Chapters 15 and 16.

3.2.2.5 A classification of pencils

We present in Chapter 15 the first classification of pencils of quadrics based on the type of their intersection in real projective space. A summary of this classification is given in Tables 3.4 and 3.5. We also show how this classification can be used to compute efficiently the type of the real intersection. In particular, we show how computations with non-rational numbers can be avoided for detecting the type of the intersection when the input quadrics have rational coefficients.

As we shall see in Chapter 16, this classification is critical to our algorithm for parameterizing the intersection of two quadrics. Indeed, our algorithm (see Chapter 16) first determine, using this classification, the real type of the intersection and then uses the associated structural information on the intersection curve to drive the algorithm for computing a near-optimal parameterization of the intersection curve.

It should however be stressed that, even though the classification of pencils over the reals is presented here as an intermediate step in a more global process (i.e., parameterization of the intersection), this classification has an interest on its own. It is used, for instance, in the work of Chapter 18 on the characterization of the topology of Voronoi diagrams of three lines. It could also be used in a collision detection context to predict at which time stamps a collision between two moving objects will happen.

Our proof relies heavily on Uhlig's Canonical Pair Form Theorem [Uhl73, Uhl76] on simultaneous block diagonalizations of pairs of symmetric real matrices. The idea is as follows. For

⁷Indeed, deciding whether the extra square root can be avoided amounts, in general, to finding a rational point on a surface of degree 8.

Segre string	roots of $\mathcal{D}(\lambda, \mu)$ in \mathbb{C}	rank or inertia of $R(\lambda_1, \mu_1)$	rank or inertia of $R(\lambda_2, \mu_2)$	type of (λ_2, μ_2)	s	type of intersection in $\mathbb{P}^3(\mathbb{R})$
[1111]	4 simple roots					smooth quartic or \emptyset ; see [Fin37] & [TWW02] (or also Th. 14.5 & 14.25)
[112]	1 double root	(3, 0)		real		point
[112]	1 double root	(2, 1)		real	-	nodal quartic; isolated node
[112]	1 double root	(2, 1)		real	+	nodal quartic; convex sing.
[112]	1 double root	rank 3		complex		nodal quartic; concave sing.
[11(11)]	1 double root	(2, 0)		real	+	\emptyset
[11(11)]	1 double root	(2, 0)		real	-	two points
[11(11)]	1 double root	(1, 1)	(2, 1)	real	-	two non-secant conics
[11(11)]	1 double root	(1, 1)	(3, 0)	real	-	\emptyset
[11(11)]	1 double root	(1, 1)		real	+	two secant conics; convex sing.
[11(11)]	1 double root	rank 2		complex	-	conic
[11(11)]	1 double root	rank 2		complex	+	two secant conics; concave sing.
[13]	triple root	rank 3				cuspidal quartic
[1(21)]	triple root	(2, 0)				double point
[1(21)]	triple root	(1, 1)				two tangent conics
[1(111)]	triple root	rank 1	(2, 1)			double conic
[1(111)]	triple root	rank 1	(3, 0)			\emptyset
[4]	quadruple root	rank 3				cubic and tangent line
[(31)]	quadruple root	(1, 1)			-	conic
[(31)]	quadruple root	(1, 1)			+	conic and two lines crossing on the conic
[(22)]	quadruple root	(2, 0)				double line
[(22)]	quadruple root	(1, 1)			+	two single lines & a double line
[(211)]	quadruple root	rank 1			-	point
[(211)]	quadruple root	rank 1			+	two secant double lines
[(1111)]	quadruple root	rank 0				any smooth quadric of the pencil
[22]	2 double roots	rank 3	rank 3	real		cubic and secant line
[22]	2 double roots	rank 3	rank 3	complex		cubic and non-secant line
[2(11)]	2 double roots	(3, 0)	rank 2	real		point
[2(11)]	2 double roots	(2, 1)	rank 2	real	+	conic and two intersecting lines
[2(11)]	2 double roots	(2, 1)	rank 2	real	-	conic and point
[(11)(11)]	2 double roots	(2, 0)	(2, 0)	real		\emptyset
[(11)(11)]	2 double roots	(2, 0)	(1, 1)	real		two points
[(11)(11)]	2 double roots	(1, 1)	(2, 0)	real		two points
[(11)(11)]	2 double roots	(1, 1)	(1, 1)	real		four lines (skew quadrilateral)
[(11)(11)]	2 double roots	rank 2	rank 2	complex		two secant lines

TAB. 3.4 – Classification of pencils in the case where $\mathcal{D}(\lambda, \mu)$ does not identically vanish. (λ_1, μ_1) denotes a multiple root of $\mathcal{D}(\lambda, \mu)$ (if any) and (λ_2, μ_2) another root (not necessarily simple). If (λ_1, μ_1) is a double root then s denotes the sign of $\frac{\det(\lambda S + \mu T)}{(\mu_1 \lambda - \lambda_1 \mu)^2}$ at $(\lambda, \mu) = (\lambda_1, \mu_1)$; if (λ_1, μ_1) is a quadruple root then s denotes the sign of $\det(\lambda S + \mu T)$ for any $(\lambda, \mu) \neq (\lambda_1, \mu_1)$. When the characteristic polynomial has multiple roots, the additional simple roots are not indicated. The Segre characteristic is mentioned for clarity but is not needed for the classification.

Segre string	roots of $\mathcal{D}_3(\lambda, \mu)$ in \mathbb{C}	rank or inertia of $R(\lambda_1, \mu_1)$	inertia of $R(\lambda_2, \mu_2)$	type of (λ_2, μ_2)	type of intersection in $\mathbb{P}^3(\mathbb{R})$
[1{3}]	no common singular point				conic and double line
[111]	3 simple roots	(1, 1)	(1, 1)	real	four concurrent lines meeting at p point p point p two lines intersecting at p
[111]	3 simple roots	(2, 0)		real	
[111]	3 simple roots		(2, 0)	real	
[111]	3 simple roots			complex	
[12]	double root	(1, 1)			2 lines and a double line meeting at p double line
[12]	double root	(2, 0)			
[1(11)]	double root	rank 1	(1, 1)		two double lines meeting at p point p
[1(11)]	double root	rank 1	(2, 0)		
[3]	triple root	rank 2			a line and a triple line meeting at p
[(21)]	triple root	rank 1			a quadruple line
[(111)]	triple root	rank 0			any non-trivial quadric of the pencil
	$\mathcal{D}_3(\lambda, \mu) \equiv 0$				same as in Table 15.3

TAB. 3.5 – Classification of pencils in the case where $\mathcal{D}(\lambda, \mu)$ identically vanishes. In the bottom part, the quadrics of the pencil have a singular point p in common. $\mathcal{D}_3(\lambda, \mu)$ is the determinant of the 3×3 upper-left matrix of $R(\lambda, \mu)$ after a congruence transformation sending p to $(0, 0, 0, 1)$. The conic associated with a root of $\mathcal{D}_3(\lambda, \mu)$ corresponds to the 3×3 upper-left matrix or $R(\lambda, \mu)$. (λ_1, μ_1) denotes the multiple root of $\mathcal{D}_3(\lambda, \mu)$ (if any) and (λ_2, μ_2) another root. When $\mathcal{D}_3(\lambda, \mu)$ has a multiple root, the additional simple roots are not indicated. The Segre characteristic is mentioned for clarity but is not needed for the classification.

each possible type (real or complex) and multiplicity of the roots of the characteristic polynomial of the pencil, we compute, using Uhlig’s result, the canonical form of the two quadrics we intersect. The canonical form of these two quadrics are quadrics whose associated matrices are block diagonal and that are obtained from the two input quadrics by the same congruence transformation which preserves the roots (values and multiplicities) of the characteristic polynomial of the pencil. We then deduce from this canonical form a *normal form* of these two quadrics, over the reals, by rescaling and translating the roots to particularly simple values. This normal form is in a sense the “simplest pair” of quadrics having a given real intersection type. The normal pencil is equivalent by a real projective transformation to any pencil of quadrics with the same real and complex intersection type.

For instance, when the characteristic polynomial of the pencil has one double root and two simple real roots, the two quadrics, in their normal form, have equations $y^2 + az^2 + bw^2 = 0$ and $xy + cw^2 = 0$ with a, b, c equal to ± 1 . Their intersection is a nodal quartic in $\mathbb{P}^3(\mathbb{C})$ and, depending on the values of a, b, c , it is, in $\mathbb{P}^3(\mathbb{R})$, a point, a nodal quartic with an isolated node, or a nodal quartic with a convex singularity (see Table 3.4).

It should be noted that the real projective transformations that send the input quadrics into

normal form, if they preserve the real type of the intersection, may well involve irrational numbers. These transformations are thus not practical for *actually* computing the intersection of two quadrics. However, they serve perfectly their purpose for classifying pencils of quadrics depending on their intersection.

3.2.2.6 Parameterizing singular intersections

Our classification of pencils over the reals allows us to determine quickly and efficiently the real type of the intersection of two given quadrics. We can furthermore focus on singular intersections (*i.e.*, other than a smooth quartic) since the “generic algorithm” of Chapter 14 produces near-optimal parameterizations when the intersection is non-singular.

In Chapter 16, we design, for each real type of non-singular intersection, an algorithm for computing an optimal or near-optimal parameterization, that is a parameterization with at most one extra square root in the coefficients. When a parameterization is possibly not optimal, we show that testing if the parameterization is non-optimal and, if so, finding an optimal parameterization is equivalent to finding a rational point on a (possibly non-rational) conic. We also give, for every type of real intersection, worst-case examples where the maximum number of square roots is reached. A summary of these results is given in Table 3.6.

Our general philosophy for these refined algorithms is to modify the “generic algorithm” such that we use for intermediate quadric Q_R the rational quadric of the pencil of smallest rank (instead of one of largest rank). As we will see, in Chapter 16, this philosophy has the double advantage of (i) avoiding the square root of a polynomial, $\sqrt{\Delta}$, in all singular cases, and (ii) minimizing the number of radicals. As an additional benefit, it helps keep the size of the numbers involved in intermediate computations and in the final parameterizations to a minimum (see Chapter 17).

3.2.2.7 An efficient and exact implementation

We present, in Chapter 17, the first complete, exact, and efficient implementation of an algorithm for parameterizing the intersection of two arbitrary quadrics, given in implicit form, with integer coefficients. (Note that quadrics with rational or finite floating-point coefficients can be trivially converted to integer form.) This implementation is based on the parameterization method described in Chapters 14, 15, and 16.

Precisely, our implementation has the following features :

- it computes an exact parameterization of the intersection of two quadrics with integer coefficients of arbitrary size ;
- it places no restriction of any kind on the type of the intersection or the type of the input quadrics ;
- it correctly identifies, separates, and parameterizes all the algebraic components of the intersection and gives all the relevant topological information ;
- the parameterization is rational when one exists ; otherwise the intersection is a smooth quartic and the parameterization involves the square root of a polynomial ;
- the parameterization is either optimal in the degree of the extension of \mathbb{Z} on which its coefficients are defined or, in a small number of well-identified cases, involves one

Segre string	real type of intersection	worst case format of parameterization	worst-case optimality of parameterization
[1111]	nonsingular quartic (see part I)	$\mathbb{Q}(\sqrt{\delta})[\xi, \sqrt{\Delta}]$, $\Delta \in \mathbb{Q}(\sqrt{\delta})[\xi]$	rational point on degree-8 surface
[112]	point	\mathbb{Q}	optimal
	nodal quartic	$\mathbb{Q}(\sqrt{\delta})[\xi]$	rational point on conic
[11(11)]	two points	$\mathbb{Q}(\sqrt{\delta})$	optimal
	conic	$\mathbb{Q}(\sqrt{\delta}, \sqrt{\mu})[\xi]$, $\mu \in \mathbb{Q}(\sqrt{\delta})$	optimal if $\sqrt{\delta} \notin \mathbb{Q}$ rational point on conic if $\sqrt{\delta} \in \mathbb{Q}$
	two non-tangent conics	$\mathbb{Q}(\sqrt{\delta}, \sqrt{\delta'})[\xi]$	$\mathbb{Q}(\sqrt{\delta'})$ -rational point on $\mathbb{Q}(\sqrt{\delta'})$ -conic
[13]	cuspidal quartic	$\mathbb{Q}[\xi]$	optimal
[1(21)]	point	\mathbb{Q}	optimal
	two tangent conics	$\mathbb{Q}(\sqrt{\delta})[\xi]$	optimal
[1(111)]	double conic	$\mathbb{Q}(\sqrt{\delta})[\xi]$	rational point on conic
[4]	cubic and tangent line	$\mathbb{Q}[\xi]$	optimal
[(31)]	conic	$\mathbb{Q}[\xi]$	optimal
	conic and two lines crossing on the conic	$\mathbb{Q}(\sqrt{\delta})[\xi]$	optimal
[(22)]	double line	$\mathbb{Q}[\xi]$	optimal
	two simple skew lines cutting a double line	$\mathbb{Q}(\sqrt{\delta})[\xi]$	optimal
[(211)]	point	\mathbb{Q}	optimal
	two double concurrent lines	$\mathbb{Q}(\sqrt{\delta})[\xi]$	optimal
[22]	cubic and non-tangent line	$\mathbb{Q}[\xi]$	optimal
[2(11)]	point	\mathbb{Q}	optimal
	conic and point	$\mathbb{Q}(\sqrt{\delta})[\xi]$	rational point on conic
	conic and two lines not crossing on the conic	$\mathbb{Q}(\sqrt{\delta})[\xi]$	rational point on conic
[(11)(11)]	two points	$\mathbb{K}[\xi]$, degree (\mathbb{K}) = 4	optimal
	two skew lines	$\mathbb{K}[\xi]$, degree (\mathbb{K}) = 4	optimal
	four lines (skew quadrilateral)	$\mathbb{K}[\xi]$, degree (\mathbb{K}) = 4	optimal
[1{3}]	conic and double line	$\mathbb{Q}[\xi]$	optimal
[111]	point	\mathbb{Q}	optimal
	two concurrent lines	$\mathbb{K}[\xi]$, degree (\mathbb{K}) = 4	optimal
	four concurrent lines	$\mathbb{K}[\xi]$, degree (\mathbb{K}) = 4	optimal
[12]	double line	$\mathbb{Q}[\xi]$	optimal
	two simple and a double concurrent lines	$\mathbb{Q}(\sqrt{\delta})[\xi]$	optimal
[3]	concurrent simple and triple lines	$\mathbb{Q}[\xi]$	optimal
[1(11)]	point	\mathbb{Q}	optimal
	two concurrent double lines	$\mathbb{Q}(\sqrt{\delta})[\xi]$	optimal
[(21)]	quadruple line	$\mathbb{Q}[\xi]$	optimal
[11]	quadruple line	$\mathbb{Q}[\xi]$	optimal

TAB. 3.6 – Ring of definition of the projective coordinates of the parameterization of each component of the intersection and optimality, in all cases where the real part of the intersection is 0- or 1-dimensional. $\delta, \delta' \in \mathbb{Q}$.

- extra possibly unnecessary square root ;
- the implementation is carefully designed so that the size of the coefficients is kept small ;
- it is fast and efficient, and can routinely compute parameterizations of the intersection of quadrics with input coefficients having fifty digits in less than 50 milliseconds on a mainstream PC.

Our code can be downloaded from the LORIA and INRIA web sites⁸. The C++ implementation can also be queried via a web interface⁹.

We also prove theoretical bounds on the size of the output coefficients and compare those bounds to observed values. Roughly speaking, the height of an output coefficient is its logarithm with base the maximum (in absolute value) of the input quadric coefficients. For instance, we prove that, when the intersection is generic, the output integer coefficients have height at most $38 + 50h_p$ where h_p is the height of the coordinates of the rational point p chosen near the quadric Q_R during step 2 of the generic algorithm. We also observe that, for random inputs quadrics with coefficients with up to 10,000 digits, the height of the largest coefficients never exceeds 36. This partially follows from the fact that, in practice, the coordinates of p are bounded by some constant and thus h_p goes to zero when the input size goes to infinity. We however have no explanation of why the bound of 36 is reached instead of 38.

We also made interesting observations on the height of the coefficients for singular intersections that validate a key choice we made in the design of our algorithm. We observe, for instance, that when the intersection is a cubic and a line that are tangent, the observed height of the cubic coefficients tend to roughly 8 or 15 depending on whether the intermediate quadric Q_R used in the algorithm is chosen with minimal or maximal rank (that is with inertia $(2, 1)$ or $(2, 2)$). This shows that our choice of choosing an intermediate quadric Q_R with minimal rank is a good strategy regarding the height of the output coefficients.

We also give experimental results and performance evaluation on random and real data. For instance, we show that the intersection of two quadrics with random coefficients with up to 50 digits takes roughly 50 ms and that the intersections of pairs of quadrics in the chess set of Figure 1.1 take on average 3.4ms on a regular PC.

3.2.3 The Voronoi diagram of three lines

The Voronoi diagram of a set of disjoint objects is a decomposition of space into cells, one cell per object, such that the cell associated with an object consists of all points that are closer to that object than to any other object. In Chapter 18, we consider the Voronoi diagram of lines in \mathbb{R}^3 under the Euclidean metric.

Voronoi diagrams have been the subject of a tremendous amount of research. For points, these diagrams and their complexities are well understood and optimal algorithms as well as robust and efficient implementations exist for computing them in any dimension (see for instance [Aur91, AK99, BDP⁺02, BDS⁺92, CSY97, CS89a, For97, OBSC00, PT06, Sei81]). Nevertheless, some important problems remain and are addressed in recent papers. The same is true

⁸<http://www.loria.fr>, <http://www.inria.fr>

⁹<http://www.loria.fr/equipes/vegas/qi>

for segments and polygons in two dimensions [Kar04].

For lines, segments, and polyhedra in three dimensions much less is known. In particular, determining the combinatorial complexity of the Voronoi diagram of n lines or line segments in \mathbb{R}^3 is an outstanding open problem. The best known lower bound is $\Omega(n^2)$ and the best upper bound is $O(n^{3+\varepsilon})$ [Sha94]. It is conjectured that the complexity of such diagrams is near-quadratic. In the restricted case of a set of n lines with a fixed number, c , of possible orientations, Koltun and Sharir have shown an upper bound of $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$ [KS03].

There are few algorithms for computing exactly the Voronoi diagram of linear objects. Most of this work has been done in the context of computing the medial axis of a polyhedron, *i.e.*, the Voronoi diagram of the faces of the polyhedron [Cul00, Mil93]. Recently, some progress has been made on the related problem of computing arrangements of quadrics (each cell of the Voronoi diagram is a cell of such an arrangement) [BHK⁺05, KKM99, MTT05, SW06, SS97]. Finally, there have been many papers reporting algorithms for computing approximations of the Voronoi diagram (see for instance [DZ02, ER02, HCK⁺99, TT97]).

In Chapter 18, we address the fundamental problem of understanding the structure of the Voronoi diagram of three lines. A robust and effective implementation of Voronoi diagrams of three-dimensional linear objects requires a complete and thorough treatment of the base cases, that is the diagrams of three and four lines, points or planes. We also strongly believe that this is required in order to make progress on complexity issues, and in particular for proving tight worst-case bounds. We provide here a full and complete characterization of the geometry and topology of the elementary though difficult case of the Voronoi diagram of three lines in general position.

Our main result, which settles a conjecture of Koltun and Sharir [KS03], is the following (see Figure 3.3).

Theorem 3.3. *The topology of the Voronoi diagram of three pairwise skew lines that are not all parallel to a common plane is invariant. The trisector consists of four infinite branches of either a non-singular quartic¹⁰ or of a cubic and a line that do not intersect in $\mathbb{P}^3(\mathbb{R})$. Each cell of dimension two consists of two connected components on a hyperbolic paraboloid that are bounded, respectively, by three and one of the branches of the trisector.*

The proof technique, which relies heavily upon modern tools of computer algebra, is of interest in its own right. We also provide a geometric characterization of the configurations of three lines (which are pairwise skew and not all parallel to a common plane) whose trisector is not generic, that is consists of a cubic and a line.

The characterization of Theorem 3.3 yields some fundamental properties of the Voronoi diagram of three lines which are likely to be critical for the analysis of the complexity and the development of efficient algorithms for computing Voronoi diagrams and medial axes of lines or polyhedra. In particular, we obtain the following results.

Monotonicity Property. *Given three pairwise skew lines that are not all parallel to a common plane, there is a direction in which all four branches of the trisector are monotonic.*

Theorem 3.4. *Let p be a point that lies in (i) a two-dimensional cell or (ii) the one-dimensional cell of the Voronoi diagram of three pairwise skew lines that are not all parallel to a common plane. There are linear semi-algebraic tests for*

¹⁰By non-singular quartic, we mean an irreducible curve of degree four with no singular point in $\mathbb{P}^3(\mathbb{C})$.

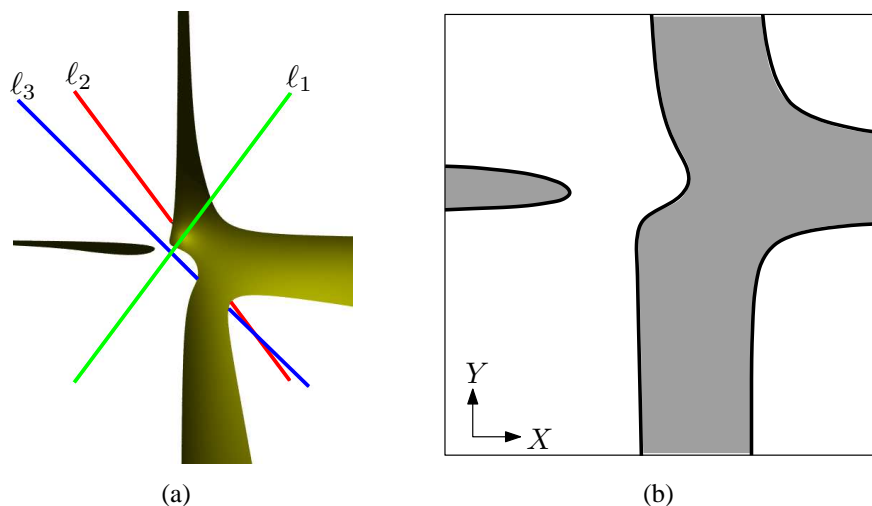


FIG. 3.3 – Voronoi diagram of three lines l_1 , l_2 , and l_3 in general position : (a) Voronoi 2D face of l_1 and l_2 , *i.e.*, set of points equidistant to l_1 and l_2 and closer to them than to l_3 . (b) Orthogonal projection of a 2D face on a plane \mathcal{P} with coordinate system (X, Y) ; the plane's normal is parallel to the common perpendicular of l_1 and l_2 and the X and Y -axes are parallel to the two bisector lines (in \mathcal{P}) of the projection of l_1 and l_2 on \mathcal{P} . The 2D face is bounded by four branches of a non-singular quartic.

(i) deciding on which connected components of the two-dimensional cell, point p lies,
or

(ii) deciding on which branches of the trisector, point p lies, respectively.

Furthermore, if the three lines are rational, these linear tests are rational.¹¹ Moreover, there is a linear semi-algebraic test for

(iii) ordering points on each branch of the trisector.

Notice that the tests (i) and (ii) enable us to answer queries of the form, given a point, determine in which connected component of which cell it lies. Notice also that tests (ii) and (iii) should be useful for computing the Voronoi diagram of n lines since it requires to locate the points equidistant to four lines on a Voronoi arc of three of these lines.

The idea of the proof of Theorem 3.3 is as follows. We consider three lines in *general position*, that is, pairwise skew and not all parallel to the same plane. The idea is to prove that the topology of the trisector is invariant by continuous deformation on the set of all triplets of three lines in general position and that this set is connected. The result then follows from the analysis of any example.

We show that the trisector is always homeomorphic to four lines that do not pairwise intersect. To prove this, we show that the trisector is always non-singular in $\mathbb{P}^3(\mathbb{R})$ and has four simple real points at infinity. To show that the trisector is always non-singular, we study the type of the intersection of two bisectors, which are hyperbolic paraboloids.

We use the classic result that the intersection of two quadrics is a non-singular quartic (in $\mathbb{P}^3(\mathbb{C})$) unless the characteristic equation of their pencil has (at least) a multiple root. In order

¹¹By rational linear test, we mean that the polynomials (whose signs determine the connected component) are of degree one in the coordinates of point p and have rational coefficients.

to determine when this equation has a multiple root, we determine when its discriminant Δ is zero.

This discriminant has several factors, some of which are trivially always positive. The remaining, so-called “*gros facteur*”, can be shown, using Safey’s software [RAG], to be never negative. This implies that it is zero only when all its partial derivatives are zero. We thus consider the system that consists of the *gros facteur* and all its partial derivatives, and compute its Gröbner basis. This gives three equations of degree six. We consider separately two components of solutions, one for which a (simple) polynomial F is zero, the other for which $F \neq 0$.

When $F \neq 0$, some manipulations and simplifications, which are interesting in their own rights, yield another Gröbner basis, with the same real roots, which consists of three equations of degree four. We show that one of these equations has no real root which implies that the system has no real root and thus that $\Delta = 0$ has no real root on the considered component. We can thus conclude that, in this case, the trisector is always a non-singular quartic in $\mathbb{P}^3(\mathbb{R})$. When $F = 0$, we show, by substituting $F = 0$ in Δ and by using the classification of the intersection of quadrics over the reals (see Chapter 15), that the trisector is a cubic and a line that do not intersect in $\mathbb{P}^3(\mathbb{R})$.

We can thus conclude that the trisector is always a non-singular quartic or a cubic and a line that do not intersect in real space and thus that the trisector is always non-singular in $\mathbb{P}^3(\mathbb{R})$.

We then prove that the trisector always contains four simple real points at infinity and thus that it is always homeomorphic to four lines that do not pairwise intersect. It follows that the topology of the Voronoi diagram is invariant by continuous deformation on any connected set of triplets of lines in general position. Next we prove that the set of triplets of lines in general position is connected which implies that the topology of the Voronoi diagram is constant. We finally determine the topology for (any) one arbitrary triplet of lines, which yields the result.

Chapitre 4

Conclusion

I have presented, in this document, my most significant recent work on the subject of non-linear computational geometry dealing with lines and quadrics in three dimensions. I have argued that the study of the fundamental properties and algorithmics of non-linear geometric primitives is a useful and necessary direction of research that is complementary to, and should be conducted along with, the more conventional research on geometric computing with linear discretized primitives.

The first part of this document presents a body of work on lines and line segments that are tangent or transversal to three-dimensional objects. Recent results are presented on their structural and combinatorial properties, as well as on the algorithmics for computing them. These results are applied to the problem of computing global data structures for three-dimensional visibility problems, which motivated the study of such tangents. This work leads me to believe that we can reasonably expect to obtain robust and certified algorithmic solutions to some difficult three-dimensional visibility problems such as the exact and efficient computation of limits of umbra and penumbra. It should be again stressed that such algorithms are not intended to compete in terms of speed with specialized graphics hardware, but rather to provide robust and exact solutions instead of approximations. It is indeed useful to know, from time-to-time, the correct (certified) answer, in particular, for the purpose of validating ad-hoc methods.

In the second part of this document, I present substantial progress on two well-known problems on geometric computing with simple curved primitives, namely, quadrics. The work on the intersection of two quadrics, aside from providing new theoretical results, demonstrates that a careful study of the geometry leads to fast and robust implementations. The work on the topology of Voronoi diagrams of three lines (a partition of space by quadrics patches) shows, in addition, that state-of-the-art computer algebra systems have reached the point where they can be used effectively for proving theorems on non-trivial geometric problems. However, we have seen that performing such algebraic computations is more of an art than a science, in particular, because the computations required for even seemingly simple non-linear geometric problems are often at the very limit of such systems. This body of work shows that significant progress can be made on various classical problems of non-linear computational geometry and that even non-discretized solutions can be extremely efficient in practice.

Despite much progress, there is much work left to do. Concerning three-dimensional visibility, algorithmic solutions for efficiently computing visibility structures, such as the visibility skeleton, are now in sight. This gives hope that the problem of efficiently computing limits of

umbra and penumbra is tractable, although the problem is still largely open. However, other major challenges such as solving exactly and effectively surface-to-surface visibility queries remains entirely open. Many problems also remain in geometric computing with low-degree algebraic surfaces. For instance, there is no robust and efficient implementation for computing arrangements, or parts of arrangements, of quadrics, or for the related problem of computing Voronoi diagrams of polyhedra. State-of-the-art methods are far from being efficient and/or usable and, though the algorithmic and arithmetic complexities of these problems are inherently high, there is certainly a lot of room for improvement, both on the theoretical and practical levels.

Deuxième partie

Propriétés des droites et segments de \mathbb{R}^3 et problèmes de visibilité tridimensionnelle

Chapitre 5

Common tangents to spheres in \mathbb{R}^3

Cet article a été publié dans *Discrete Computational Geometry* [BGLP06] ainsi que dans la thèse de X. Goaoc [Goa04].

Abstract

We prove that four spheres in \mathbb{R}^3 have infinitely many real common tangents if and only if they have aligned centers and at least one real common tangent.

5.1 Introduction

A major issue in geometric computing is to handle degenerate inputs properly in order to design robust algorithms. This often requires recognizing such an input to begin with. In 3D visibility problems, which are ubiquitous in computer graphics and image synthesis, objects with a set of common tangents of improper dimension constitute degenerate configurations, as detailed in the survey of Durand [Dur00]. In this paper, we determine all degenerate configurations of four distinct spheres, that is all configurations of four spheres with infinitely many common tangents.

The study of real lines tangent to basic geometric objects has been very active in recent years. This topic includes two closely related directions of research, namely the characterization of degenerate configurations and the enumeration of lines satisfying geometric constraints. Usually, these problems are approached by studying the degeneracies and counting the number of solutions of some specific polynomial system. The difficulty often resides in eliminating imaginary solutions, solutions at infinity, and components of positive dimension of solutions in order to retain only real affine solutions.

The case of lines tangent to spheres has been persistently investigated. Macdonald *et al.* [MPT01] proved that four unit spheres have at most 12 common tangents in general, and infinitely many common tangents if and only if the centers are aligned. The bound of 12 was independently obtained by Devillers *et al.* [DMPT03]. Examples show that, in the finite case, this bound is tight [DMPT03, MPT01], yet, according to Megyesi [Meg01], it drops to 8 in the case of unit spheres with coplanar but non-collinear centers. However, the upper bound of 12 remains valid when the spheres have arbitrary radii. Sottile and Theobald [ST02] proved that there are $3 \cdot 2^{n-1}$ complex common tangent lines to $2n - 2$ general spheres in \mathbb{R}^n , and that there exists a choice of spheres with all common tangents real.

Recently, progress has also been made in understanding the varieties of common tangents to spheres and transversals to lines. Theobald [The02] described the configurations of three lines and a sphere having infinitely many common tangents/transversals. Next, Megyesi *et al.* [MST03] characterized the families of two lines and two quadrics of $\mathbb{P}^3(\mathbb{C})$ with infinitely many tangents/transversals, and applied their results to the case of two lines and two spheres of \mathbb{R}^3 . Last, Megyesi and Sottile [MS05] classified the families of one line and three spheres of \mathbb{R}^3 with infinitely many tangents/transversals.

The question of characterizing the positions of four spheres of various radii with infinitely many common tangents remained open. Quoting Theobald [The02] : “*We conjecture that there does not exist any configuration*

with four balls of arbitrary radii, non-collinear centers and infinitely many common tangent lines.” In this paper, we confirm this expectation and prove

Theorem 5.1. *Four distinct spheres in \mathbb{R}^3 have infinitely many real common tangent lines if and only if they have aligned centers and at least one real common tangent.*

More precisely we prove that four spheres with infinitely many common real tangents either intersect in a circle, possibly degenerating to a point, or each sphere has a circle of tangency with one and the same quadric of revolution with symmetry axis the line through all centers (see Figure 5.1); such a quadric is *unique* and can be a cone, a cylinder or a hyperboloid of one sheet. Furthermore, the common tangents to the four spheres are exactly the common tangents to any three of them.

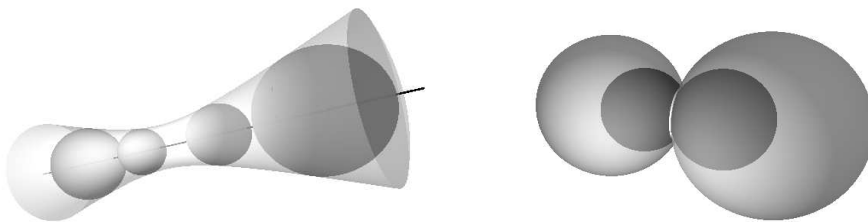


FIG. 5.1 – Two examples of quadruples of spheres with infinitely many common tangents.

After introducing some notations and preliminaries in Section 5.2, we treat the case of four spheres with affinely independent centers in Section 5.3. Next, we handle in Section 5.4 the more intricate case of spheres with coplanar centers, no three aligned. Section 5.5 ends the proof of Theorem 5.1 with the case of three aligned centers. We obtain, at the same time, the algebraic and semi-algebraic conditions on radii and mutual distances between centers, which characterize four spheres with infinitely many common real tangents.

5.2 Preliminaries

Notations

Our proofs use points and vectors from \mathbb{R}^n and from the real and complex projective spaces of dimension n , $\mathbb{P}^n(\mathbb{R})$ and $\mathbb{P}^n(\mathbb{C})$. We make no distinction between a point p and the vector from the origin of the frame to p . For more clarity, we denote an element of \mathbb{R}^n by (a_1, \dots, a_n) , and an element of $\mathbb{P}^n(\mathbb{R})$ or $\mathbb{P}^n(\mathbb{C})$ by $(a_1 : \dots : a_{n+1})$.

For any two vectors a, b of \mathbb{R}^n , $\mathbb{P}^n(\mathbb{R})$, or $\mathbb{P}^n(\mathbb{C})$, we denote by $a \cdot b$ their dot product, by $a \times b$ their cross product, and by $|a|^2$ the dot product $a \cdot a$ (note that $|a|^2$ is not the square of the norm of a when a has imaginary coordinates).

Let \mathcal{S}_i denote the sphere of \mathbb{R}^3 with center c_i and radius $r_i > 0$, for $i = 1, \dots, 4$, and (e_1, e_2, e_3) be an orthonormal frame of \mathbb{R}^3 . Without loss of generality, **we assume that c_1 is the origin of our frame**. The *axis* of a set of spheres with aligned centers is the line going through these centers.

Tangents to four spheres

We begin by reviewing the description of the common tangent lines to four spheres as solutions of a polynomial system, as in [MPT01]. We represent a line in \mathbb{R}^3 by its closest point to the origin $p \in \mathbb{R}^3$ and its direction vector $v \in \mathbb{P}^2(\mathbb{R})$. Let M denote the matrix $[c_2, c_3, c_4]^T$ and Φ_0 and $\Phi_2(v)$ be the vectors

$$\Phi_0 = \begin{pmatrix} |c_2|^2 + r_1^2 - r_2^2 \\ |c_3|^2 + r_1^2 - r_3^2 \\ |c_4|^2 + r_1^2 - r_4^2 \end{pmatrix}, \quad \Phi_2(v) = - \begin{pmatrix} (c_2 \cdot v)^2 \\ (c_3 \cdot v)^2 \\ (c_4 \cdot v)^2 \end{pmatrix}.$$

Lemma 5.2. *The lines tangent to the four spheres $\mathcal{S}_1, \dots, \mathcal{S}_4$ are the common solutions (p, v) in $\mathbb{R}^3 \times \mathbb{P}^2(\mathbb{R})$ of the equations*

$$p \cdot v = 0, \quad (5.1)$$

$$|p|^2 = r_1^2, \quad (5.2)$$

$$2|v|^2 Mp = \Phi_2(v) + |v|^2 \Phi_0. \quad (5.3)$$

Proof. A couple $(p, v) \in \mathbb{R}^3 \times \mathbb{P}^2(\mathbb{R})$ represents a line if and only if Equation (5.1) is satisfied. A line (p, v) is tangent to sphere \mathcal{S}_i if and only if its squared distance to c_i is r_i^2 that is, if and only if

$$|(c_i - p) \times v|^2 = r_i^2 |v|^2.$$

Expanding this equation yields

$$|c_i \times v|^2 + |p \times v|^2 - 2(c_i \times v) \cdot (p \times v) = r_i^2 |v|^2. \quad (5.4)$$

Applying to $(c_i \times v) \cdot (p \times v)$ the scalar triple product identity $a \cdot (b \times c) = b \cdot (c \times a)$, then the vector triple product identity $a \times (b \times c) = (a \cdot c)b - (a \cdot b)c$ and finally using Equation (5.1) we get

$$(p \times v) \cdot (c_i \times v) = c_i \cdot (v \times (p \times v)) = c_i \cdot ((v \cdot v)p - (v \cdot p)v) = |v|^2 c_i \cdot p.$$

Since p and v are orthogonal, $|p \times v|^2 = |p|^2 |v|^2$ and thus Equation (5.4) becomes

$$2|v|^2 c_i \cdot p = |c_i \times v|^2 + |v|^2 (|p|^2 - r_i^2).$$

As $|c_i \times v|^2 + (c_i \cdot v)^2 = |c_i|^2 |v|^2$, we finally get that

$$2|v|^2 c_i \cdot p = -(c_i \cdot v)^2 + |v|^2 (|c_i|^2 + |p|^2 - r_i^2). \quad (5.5)$$

Equation (5.5) for $i = 1$ is equivalent to Equation (5.2) since c_1 is the origin of the frame. It follows that the four equations (5.5) for $i = 1, \dots, 4$ are equivalent to the two equations (5.2) and (5.3). \square \square

The approach used to show that infinitely many tangent lines to spheres can only happen when the centers of the spheres are aligned is as follows. We eliminate p among the equations (5.1)-(5.3), giving two curves¹ in the 2D projective space of directions, whose intersection contains all directions along which a common tangent line to the four spheres is observed. We then prove that, when the centers are non-collinear, the two curves intersect in a finite number of points.

The key idea behind the proofs of Section 5.3 (affinely independent centers) and Section 5.4 (coplanar centers) is that if the two curves, envisaged as complex projective curves, had a common component of positive dimension, this component would intersect the imaginary conic $|v|^2 = 0$ and we show that this is not the case. Intersecting the curve with $|v|^2 = 0$ is inspired by the relation of the Grassmannian of lines in $\mathbb{P}^3(\mathbb{C})$ with the (p, v) coordinate system, well adapted to the representation of lines in the affine part $\mathbb{R}^3 \subset \mathbb{P}^3(\mathbb{R})$.

It should be stressed that any solution to the problem of characterizing sets of four spheres with infinitely many tangent lines must be computational to some extent, because while we are interested in real lines, the “native” system of equations is over \mathbb{C} . Any understanding of the system should involve sensitivity to complex degeneracies. In our proof, computations flow towards revealing such complex degeneracies, but are short-circuited by use of reality assumptions.

¹A cubic and a quartic when the centers are affinely independent, a conic and a sextic when the centers are coplanar with no three aligned.

5.3 Affinely independent centers

We first investigate the case of spheres with affinely independent centers.

Proposition 5.3. *Four spheres with affinely independent centers have at most twelve common tangent lines.*

Proof. First note that matrix M is invertible since the spheres have affinely independent centers. Considering (p, v) in $\mathbb{R}^3 \times \mathbb{P}^2(\mathbb{R})$, we have $|v|^2 \neq 0$ and thus Equations (5.1)-(5.3) are equivalent to the three equations

$$p = M^{-1} \left(\frac{\Phi_2(v)}{2|v|^2} + \frac{1}{2}\Phi_0 \right), \quad (5.6)$$

$$(M^{-1}(\Phi_2(v) + |v|^2\Phi_0)) \cdot v = 0, \quad (5.7)$$

$$|M^{-1}(\Phi_2(v) + |v|^2\Phi_0)|^2 = 4r_1^2|v|^4. \quad (5.8)$$

Equation (5.6) expresses the point p in terms of the direction vector v , proving that there is at most one line tangent to the four spheres with a given direction. The remaining equations are a cubic (5.7) and a quartic (5.8) in v , and their intersection represents the directions $v \in \mathbb{P}^2(\mathbb{R})$ along which there is a tangent to the four spheres. We want to prove that the cubic and the quartic intersect in at most 12 points in $\mathbb{P}^2(\mathbb{R})$. For that purpose we prove this property in $\mathbb{P}^2(\mathbb{C})$, by contradiction.

If the cubic and the quartic have in $\mathbb{P}^2(\mathbb{C})$ a common component of positive dimension, this component intersects the conic $|v|^2 = 0$; this is a property of any two curves in $\mathbb{P}^2(\mathbb{C})$ which does not dispute the fact that the real solutions of Equations (5.6)-(5.8) satisfy $|v|^2 \neq 0$. We now prove that the intersection in $\mathbb{P}^2(\mathbb{C})$ of the cubic (5.7), the quartic (5.8) and the conic $|v|^2 = 0$ is empty. This system simplifies to

$$\begin{cases} |v|^2 = 0, \\ (M^{-1}\Phi_2(v)) \cdot v = 0, \\ |M^{-1}\Phi_2(v)|^2 = 0. \end{cases}$$

The first two equations express the fact that $M^{-1}\Phi_2(v)$ is on the tangent at v to the smooth conic $|v|^2 = 0$, and the last that $M^{-1}\Phi_2(v)$ is itself on that conic. It follows that $M^{-1}\Phi_2(v)$ and v are one and the same projective point. Thus there exists $\mu \neq 0$ in \mathbb{C} such that

$$M^{-1}\Phi_2(v) = \mu v, \text{ that is } \Phi_2(v) = \mu Mv.$$

Expanding this last equality yields $-(c_i \cdot v)^2 = \mu c_i \cdot v$, for $i = 2, \dots, 4$, which implies that every term $c_i \cdot v$ is 0 or $-\mu$. This leads to

$$Mv = -\mu \begin{pmatrix} a_2 \\ a_3 \\ a_4 \end{pmatrix} \quad (5.9)$$

where each a_i is equal to 0 or 1. Let a denote the vector of the a_i . Plugging $v = \mu M^{-1}a$ in the equation of the conic $|v|^2 = 0$ yields

$$\mu^2 |M^{-1}a|^2 = 0.$$

The vector $M^{-1}a$ is real, thus $\mu = 0$ or $a = 0$. In both cases, Equation (5.9) implies $v = 0$. Thus there is no common solution in $\mathbb{P}^2(\mathbb{C})$ for the system of the conic, the cubic and the quartic, hence the cubic (5.7) and quartic (5.8) cannot intersect in a curve. By Bezout's Theorem, they intersect in at most 12 points, and since there is at most one line tangent to the four spheres with a given direction by Equation (5.6), this completes the proof. \square \square

5.4 Coplanar centers

We now treat the more intricate case of four spheres whose centers are coplanar but such that no three centers are aligned.

Proposition 5.4. *Four spheres with coplanar centers, no three aligned, have at most twelve common tangents.*

Let $(p, v) \in \mathbb{R}^3 \times \mathbb{P}^2(\mathbb{R})$ represent a line tangent to the four spheres $\mathcal{S}_1, \dots, \mathcal{S}_4$. By Lemma 5.2, (p, v) is solution of Equations (5.1)-(5.3). As in Section 5.3, we start by extracting from these equations two equations in v .

Without loss of generality, we may assume that the sphere centers span the plane (e_1, e_2) :

$$M = \begin{pmatrix} c_{21} & c_{22} & 0 \\ c_{31} & c_{32} & 0 \\ c_{41} & c_{42} & 0 \end{pmatrix}.$$

Let M_{12} be the 2×2 upper left sub-matrix of M , which is invertible since no three centers are aligned. For any vector a , let a_{12} be the vector that consists of the first two rows of a and a_3 be its third row.

Let us first assume that $v_3 \neq 0$. It follows from $p \cdot v = 0$ that

$$p_3 = -\frac{p_{12} \cdot v_{12}}{v_3}, \quad (5.10)$$

and p_{12} is characterized using Equation (5.3) :

$$2|v|^2 p_{12} = M_{12}^{-1} ((\Phi_2(v))_{12} + |v|^2 (\Phi_0)_{12}).$$

Let $\Psi_2(v) = M_{12}^{-1} (\Phi_2(v))_{12}$ and $\Psi_0 = M_{12}^{-1} (\Phi_0)_{12}$. As $\Phi_2(v)$ and $\Psi_2(v)$ do not depend on v_3 , we may write them as $\Phi_2(v_{12})$ and $\Psi_2(v_{12})$. Then

$$2|v|^2 p_{12} = \Psi_2(v_{12}) + |v|^2 \Psi_0. \quad (5.11)$$

Substituting the expression of p_3 from Equation (5.10) in Equation (5.2) gives

$$|p_{12}|^2 + \left(\frac{p_{12} \cdot v_{12}}{v_3} \right)^2 - r_1^2 = 0.$$

Then multiplying by $4|v|^4 v_3^2$ and substituting $2|v|^2 p_{12}$ by its expression from Equation (5.11) gives the following sextic equation in v :

$$v_3^2 |\Psi_2(v_{12}) + |v|^2 \Psi_0|^2 + ((\Psi_2(v_{12}) + |v|^2 \Psi_0) \cdot v_{12})^2 - 4|v|^4 v_3^2 r_1^2 = 0. \quad (5.12)$$

For any p, q in $\mathbb{P}^3(\mathbb{C})$, we have, by transposition :

$$(Mp) \cdot q = p \cdot (M^T q).$$

Let ω be a non-zero kernel vector of M^T . Then $(Mp) \cdot \omega = p \cdot (M^T \omega) = 0$. Substituting the expression of Mp from Equation (5.3), we obtain that v must be on the following conic :

$$\Phi_2(v_{12}) \cdot \omega + |v|^2 \Phi_0 \cdot \omega = 0. \quad (5.13)$$

Notice that Equations (5.12) and (5.13), obtained for $v_3 \neq 0$, are still valid for $v_3 = 0$ by continuity. We thus get the following lemma.

Lemma 5.5. *The direction $v \in \mathbb{P}^2(\mathbb{R})$ of a line tangent to the four spheres $\mathcal{S}_1, \dots, \mathcal{S}_4$ satisfies the sextic (5.12) and the conic (5.13).*

Lemma 5.6. *If the sextic (5.12) and the conic (5.13) admit a component of positive dimension of common solutions in $\mathbb{P}^2(\mathbb{C})$, then it intersects the conic $|v|^2 = 0$ and any point v in the intersection satisfies*

$$\exists \lambda \in \mathbb{C}, \quad \Psi_2(v_{12}) = \lambda v_{12} \quad (5.14)$$

$$\Phi_2(v_{12}) \cdot \omega = 0. \quad (5.15)$$

Proof. If Equations (5.12) and (5.13) share a component of positive dimension in $\mathbb{P}^2(\mathbb{C})$, then this component, seen as a curve of $\mathbb{P}^2(\mathbb{C})$, intersects the conic $|v|^2 = 0$. Let $v \in \mathbb{P}^2(\mathbb{C})$ be in this intersection. Then Equation (5.13) becomes Equation (5.15). Now, it follows from $|v|^2 = 0$ that $v_3^2 = -|v_{12}|^2$, and thus Equation (5.12) becomes

$$-|v_{12}|^2 |\Psi_2(v_{12})|^2 + (\Psi_2(v_{12}) \cdot v_{12})^2 = 0.$$

Since $|x|^2|y|^2 - (x \cdot y)^2 = \det(x, y)^2$ for any $x, y \in \mathbb{C}^2$, the equation is equivalent to $\det(v_{12}, \Psi_2(v_{12})) = 0$ which is equivalent to Equation (5.14) (v is on $|v|^2 = 0$ so we cannot have $v_{12} = 0$). \square \square

In the following we consider the centers $c_1 = 0, c_2, c_3, c_4$ as 2D points (i.e., we forget the third coordinate, which is 0). For any vector $x \in \mathbb{R}^2$ we denote by x^\perp its orthogonal vector obtained by a rotation of angle $\pi/2$.

Lemma 5.7. *If Equations (5.14) and (5.15) have a common solution v_{12} in $\mathbb{P}^1(\mathbb{C})$, it must satisfy $v_{12} = c_i^\perp$ and $v_{12} \cdot (c_j - c_k) = 0$, with $\{i, j, k\} = \{2, 3, 4\}$ (which implies that c_1, c_2, c_3, c_4 are the vertices of a trapezoid).*

Proof. From $M^T \omega = 0$ we get

$$M^T \omega = \begin{pmatrix} M_{12}^T & c_4 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} \omega_{12} \\ \omega_3 \end{pmatrix} = \begin{pmatrix} M_{12}^T \omega_{12} + \omega_3 c_4 \\ 0 \end{pmatrix} = 0.$$

Thus $\omega_{12} = -\omega_3 (M_{12}^T)^{-1} c_4$ and $\omega_3 \neq 0$ (otherwise, $\omega_{12} = 0$ thus $\omega = 0$ contradicting its definition). Now, we can write Equation (5.15) as $(\Phi_2(v_{12}))_{12} \cdot \omega_{12} - (c_4 \cdot v_{12})^2 \omega_3 = 0$, and substituting our expression of ω_{12} yields

$$-\omega_3 (\Phi_2(v_{12}))_{12} \cdot ((M_{12}^T)^{-1} c_4) - (c_4 \cdot v_{12})^2 \omega_3 = 0,$$

which simplifies, by transposition, into

$$(M_{12}^{-1} (\Phi_2(v_{12}))_{12}) \cdot c_4 + (c_4 \cdot v_{12})^2 = 0.$$

Hence, an equivalent expression for Equation (5.15) is :

$$\Psi_2(v_{12}) \cdot c_4 + (c_4 \cdot v_{12})^2 = 0. \quad (5.16)$$

Substituting $\Psi_2(v_{12}) = \lambda v_{12}$ from Equation (5.14) into (5.16) leads to

$$(c_4 \cdot v_{12})^2 = -\lambda c_4 \cdot v_{12}.$$

By a similar reasoning, we can express the conic (5.15) using c_2 or c_3 in expressions similar to Equation (5.16), and the above argument yields that :

$$(c_i \cdot v_{12})^2 = -\lambda c_i \cdot v_{12}, \quad i = 2, 3, 4. \quad (5.17)$$

If $c_i \cdot v_{12} \neq 0$ for $i = 2, 3$, and 4 then $(c_2 - c_3) \cdot v_{12} = (c_2 - c_4) \cdot v_{12} = 0$ and, since c_2, c_3 and c_4 are not aligned, $v_{12} = 0$ contradicting $v_{12} \in \mathbb{P}^1(\mathbb{C})$. Hence, v_{12} must be orthogonal to some c_i , $i \in \{2, 3, 4\}$. Since $v_{12} \in \mathbb{P}^1(\mathbb{C})$, we can assume that $v_{12} = c_i^\perp$. Since no three centers are aligned, v_{12} is orthogonal to neither c_j nor c_k , with $\{i, j, k\} = \{2, 3, 4\}$. Thus Equation (5.17) yields

$$-\lambda = c_j \cdot c_i^\perp = c_k \cdot c_i^\perp, \quad \text{and so} \quad c_i^\perp \cdot (c_j - c_k) = 0.$$

This means that the segments $c_1 c_i$ and $c_j c_k$ are parallel and thus the centers of the spheres are the vertices of a trapezoid. \square \square

Lemma 5.8. *If the sextic (5.12) and the conic (5.13) have a common component of positive dimension in $\mathbb{P}^2(\mathbb{C})$, Equations (5.14) and (5.15) have at least two distinct solutions in $\mathbb{P}^1(\mathbb{C})$.*

Proof. Assume that the sextic (5.12) and the conic (5.13) share a component of positive dimension. Then by Lemmas 5.6 and 5.7, Equations (5.14) and (5.15) admit a common solution $v_{12} = c_i^\perp$ for $i = 2, 3$, or 4. By relabeling if necessary, we can assume $v_{12} = c_4^\perp$. Suppose, for a contradiction, that c_4^\perp is the unique common solution of Equations (5.14) and (5.15).

By Lemma 5.6, any point in the intersection of the conic $|v|^2 = 0$ and the common component of the sextic (5.12) and the conic (5.13) satisfies Equations (5.14) and (5.15). Thus any such point satisfies $v_{12} = c_4^\perp$ and $|v|^2 = 0$, and is equal to one of the two points of coordinates $(c_4^\perp : \pm i|c_4|)$. Hence the common component contains at least one of these two points.

The common component of the sextic (5.12) and the conic (5.13) is either the conic itself or a line. In the latter case, the equation of the line is real because otherwise its conjugate is also contained in the conic and in the sextic (since their equations are real); the sextic then contains the conic, which corresponds to the first case. Hence the equation of the common component is real in both cases. Thus, since the common component contains one of the two points $(c_4^\perp : \pm i|c_4|)$, it also contains its conjugate, hence the two points.

We now discard the case where the common component is the conic by deriving a contradiction with our assumption that no three centers are collinear. If the conic is contained in the sextic, it meets $|v|^2 = 0$ in the two points $(c_4^\perp : \pm i|c_4|)$, which are therefore tangency points. This means that Equation (5.16), which is our conic mod $|v|^2 = 0$, has a double root at $v_{12} = c_4^\perp$. Since any degree-two polynomial in $v_{12} \in \mathbb{P}^1(\mathbb{C})$ that has c_4^\perp as double root is proportional to $(c_4 \cdot v_{12})^2$, we get that

$$\Psi_2(v_{12}) \cdot c_4 = \alpha(c_4 \cdot v_{12})^2$$

for some $\alpha \in \mathbb{C}$ and all $v_{12} \in \mathbb{P}^1(\mathbb{C})$. Computing $\det(M_{12})M_{12}^{-1}$ gives the matrix with columns $[-c_3^\perp c_2^\perp]$, thus our equation becomes

$$\Psi_2(v_{12}) \cdot c_4 = \frac{1}{\det(M_{12})} [(c_3^\perp \cdot c_4)(c_2 \cdot v_{12})^2 - (c_2^\perp \cdot c_4)(c_3 \cdot v_{12})^2] = \alpha(c_4 \cdot v_{12})^2.$$

Since the four centers form a trapezoid we have $c_4 = v(c_2 - c_3)$ for some $v \in \mathbb{R}^3$. Replacing c_4 by its expression and simplifying by factor $c_3^\perp \cdot c_2 = -c_2^\perp \cdot c_3$ yields

$$(c_2 \cdot v_{12})^2 - (c_3 \cdot v_{12})^2 = \kappa((c_2 - c_3) \cdot v_{12})^2,$$

for some $\kappa \in \mathbb{C}$. Writing $v_{12} = xc_2^\perp + yc_3^\perp$ we obtain

$$(c_2 \cdot c_3^\perp)^2(y^2 - x^2 - \kappa(x+y)^2) = 0$$

for all $(x, y) \in \mathbb{P}^1(\mathbb{C})$, which forces the proportionality of c_2 and c_3 and their alignment with c_1 . Thus, if no three centers are aligned the conic cannot be contained in the sextic.

Now we examine the second alternative, when the common component of the sextic (12) and the conic (13) is a line. This line contains the two points $(c_4^\perp : \pm i|c_4|)$ and thus contains the point $(c_4^\perp : v_3)$ for all $v_3 \in \mathbb{C}$. Thus all the coefficients of the sextic (5.12) viewed as an equation in v_3 with coefficients depending on $v_{12} = c_4^\perp$ must vanish. In particular the constant and the coefficient of v_3^2 minus $|c_4|^{-4}$ times the coefficient of v_3^6 both vanish and are equal to

$$\begin{aligned} \Psi_2(c_4^\perp) \cdot c_4^\perp + |c_4|^2 \Psi_0 \cdot c_4^\perp &= 0, \\ |\Psi_2(c_4^\perp)|^2 + 2|c_4|^2 \Psi_2(c_4^\perp) \cdot \Psi_0 &= 0. \end{aligned}$$

From the proof of Lemma 5.7, we know that $\Psi_2(c_4^\perp) = \lambda c_4^\perp$ with $\lambda = -c_2 \cdot c_4^\perp = -c_3 \cdot c_4^\perp$. Thus, the relations become

$$\begin{aligned} |c_4|^2(\lambda + \Psi_0 \cdot c_4^\perp) &= 0, \\ \lambda|c_4|^2(\lambda + 2\Psi_0 \cdot c_4^\perp) &= 0. \end{aligned}$$

Since no three centers are aligned, $\lambda \neq 0$ and $|c_4|^2 \neq 0$, and these two equations imply $\lambda = 0$, a contradiction. $\square \square$

Lemma 5.9. *The sextic (5.12) and the conic (5.13) cannot have a component of positive dimension of common solutions.*

Proof. Assume that the sextic (5.12) and the conic (5.13) have a common component of positive dimension. Lemmas 5.7 and 5.8 yield that Equations (5.14) and (5.15) then have at least two distinct solutions among $\{c_2^\perp, c_3^\perp, c_4^\perp\}$. By relabeling the centers, we may assume these solutions are c_2^\perp and c_3^\perp . Lemma 5.7 gives that

$$c_2^\perp \cdot (c_4 - c_3) = 0 \quad \text{and} \quad c_3^\perp \cdot (c_4 - c_2) = 0.$$

Thus, c_2 is proportional to $c_4 - c_3$, and c_3 is proportional to $c_4 - c_2$. Therefore, $c_2 + c_3 = c_4$ and the centers form a parallelogram. By translating our frame to the center of that parallelogram, we may assume that the centers are at $a = (a_1, a_2, 0)$, $b = (b_1, b_2, 0)$, $-a$ and $-b$, with corresponding radii r_i , $i = 1, \dots, 4$. On occasion, we abuse notation, and allow a and b to stand for (a_1, a_2) , respectively (b_1, b_2) .

Subtracting Equation (5.5) for $i = 1$ from its expression for $i = 3$ leads to

$$4(a \cdot p) = r_3^2 - r_1^2,$$

and the same operation for $i = 2$ and $i = 4$ yields

$$4(b \cdot p) = r_4^2 - r_2^2.$$

This shows that the first two coordinates p_{12} of p are determined by centers and radii alone, and remain constant. Thus, all the common tangents to the four spheres meet the line perpendicular to the plane of the centers in p_{12} .

A theorem in the preprint [MS05] addresses a situation of this nature and shows that the common tangents to three spheres which meet at the same time a fixed line cannot be infinitely many unless their three centers are collinear. We give here an independent proof which continues the above line of thought.

Recall that (12) and (13) were obtained from (1)-(3) by eliminating p . Operating "in reverse", it is easy to see that a one-dimensional component of solutions for (12) and (13) would produce a one-dimensional family of solutions for (1)-(3). We show now this cannot happen.

Rewriting (5.5) for the centers a and b gives

$$(a \cdot v)^2 = |v|^2 \left(|a|^2 + |p|^2 - \frac{1}{2}(r_1^2 + r_3^2) \right), \quad (5.18)$$

$$(b \cdot v)^2 = |v|^2 \left(|b|^2 + |p|^2 - \frac{1}{2}(r_2^2 + r_4^2) \right). \quad (5.19)$$

Let $\alpha = |a|^2 - \frac{1}{2}(r_1^2 + r_3^2)$ and $\beta = |b|^2 - \frac{1}{2}(r_2^2 + r_4^2)$. Subtracting (5.19) from (5.18) gives the conic

$$((a+b) \cdot v_{12}) ((a-b) \cdot v_{12}) = |v|^2 (\alpha - \beta). \quad (5.20)$$

Multiplying (5.18), (5.19), and v_3^2 together and dividing by $|v|^2$ gives

$$(a \cdot v_{12})^2 (\beta + |p_{12}|^2 + p_3^2) v_3^2 = (b \cdot v_{12})^2 (\alpha + |p_{12}|^2 + p_3^2) v_3^2,$$

or equivalently, using (5.10),

$$(|p_{12}|^2 v_3^2 + (p_{12} \cdot v_{12})^2) ((a+b) \cdot v_{12}) ((a-b) \cdot v_{12}) = v_3^2 (\alpha (b \cdot v_{12})^2 - \beta (a \cdot v_{12})^2). \quad (5.21)$$

For the conic (5.20) and the quartic (5.21) to have a common one-dimensional component, it is necessary that equality holds for any $v_{12} \in \mathbb{P}^1$ and some adequate value(s) for v_3 . Indeed, the projection $v \mapsto v_{12}$ of the common component cannot be constant, for with fixed v_{12} and (already known) fixed p_{12} , equations (5.10) and (5.18) (or (5.19)) would determine only a finite number of solutions v_3 .

Evaluating (5.20) and (5.21) at $v_{12} = (a+b)^\perp$, we find no possible value for v_3 , unless $\alpha = \beta$. Returning this necessary condition into (5.20) implies $v_{12} = (a \pm b)^\perp$ contradicting the fact that (5.20) and (5.21) holds for all $v_{12} \in \mathbb{P}^1$. \square \square

We now conclude on the case of spheres with coplanar centers.

Proof of Proposition 5.4. By Lemmas 5.5 and 5.9, there are finitely many directions along which the spheres have a common tangent. For each such direction v , a line tangent to the four spheres projects onto a plane orthogonal to v into a point that lies on the common intersection of the four circles obtained as the boundary of the projection of each sphere. There are thus at most two lines tangent to the four spheres per direction. Hence there are finitely many lines tangent to the four spheres. Now, the bound of 12 directly follows from the non-coplanar case (Proposition 5.3) by continuity. \square

5.5 Collinear centers

We conclude in this section the proof of Theorem 5.1. We first establish the following lemma.

Lemma 5.10. *The common tangents to three distinct spheres with collinear centers and no common intersection are, if any, the ruling(s) of a single quadric of revolution with symmetry axis the line through all centers. This quadric can be a cone, a cylinder or a hyperboloid of one sheet.*

Proof. Suppose that three distinct spheres with collinear centers admit a common tangent. Such a tangent is not orthogonal to the axis of the three spheres since they have no common intersection. Furthermore, such a tangent remains tangent after a rotation about this axis. Thus the common tangents to the three spheres are the rulings of a collection \mathcal{Q} of quadrics of revolution with symmetry axis the line through all centers (see Figure 5.1); these quadrics have to be cylinders, cones, or hyperboloids of one sheet. Assume for a contradiction that \mathcal{Q} consists of more than one quadric.

We take the line through the centers to be the y -axis in some (x, y) -plane. This plane intersects the quadrics of \mathcal{Q} into a collection \mathcal{C} of conics symmetric with respect to the y -axis which have equations of the following form :

$$x^2 + Ay^2 + By + C = 0, \quad A \leq 0, \quad B^2 - 4AC \leq 0. \quad (5.22)$$

The (x, y) -plane also intersects the three spheres into three circles, with centers $(0, \alpha_i)$ and radii r_i , $i = 1, \dots, 3$, that are tangent to the conics of \mathcal{C} . Since these conics and circles are symmetric with respect to the y -axis, two of them are tangent if and only if they intersect in exactly two points with same y -coordinate. Thus a conic (5.22) and a circle of center $(0, \alpha_i)$ and radius r_i are tangent if and only if

$$(x^2 + Ay^2 + By + C) - (x^2 + (y - \alpha_i)^2 - r_i^2) = 0$$

has a double solution in y , i.e. the discriminant vanishes :

$$\delta_i = (B + 2\alpha_i)^2 - 4(A - 1)(C + r_i^2 - \alpha_i^2) = 0. \quad (5.23)$$

For the three circles, this gives a system of three equations in the three indeterminates (A, B, C) . This system is linear in C (with a non-zero coefficient since $A \leq 0$) and thus has more than one solution only if the linear system in (A, B)

$$\begin{cases} \delta_1 - \delta_2 = ((\alpha_1^2 - \alpha_2^2) - (r_1^2 - r_2^2))A + (\alpha_1 - \alpha_2)B + r_1^2 - r_2^2 = 0 \\ \delta_1 - \delta_3 = ((\alpha_1^2 - \alpha_3^2) - (r_1^2 - r_3^2))A + (\alpha_1 - \alpha_3)B + r_1^2 - r_3^2 = 0 \end{cases}$$

does, that is only if the determinant of the coefficients of A and B , and the determinant of the constant coefficients and the coefficients of B both vanish. The sum of these determinants also vanishes and is equal to

$$\begin{vmatrix} \alpha_1^2 - \alpha_2^2 & \alpha_1 - \alpha_2 \\ \alpha_1^2 - \alpha_3^2 & \alpha_1 - \alpha_3 \end{vmatrix} = (\alpha_1 - \alpha_2)(\alpha_1 - \alpha_3)(\alpha_2 - \alpha_3).$$

Hence at least two centers are equal which implies that one sphere is strictly contained in another. The three spheres thus have no common tangent, a contradiction. \square

Remark 5.11. Actually solving the system (5.23), $i = 1, 2, 3$, yields, in terms of radii and oriented distances between centers $d_{ij} = \alpha_j - \alpha_i$:

$$A = \frac{1}{D} (r_1^2 d_{23} + r_2^2 d_{31} + r_3^2 d_{12}),$$

$$B^2 - 4AC = \frac{-1}{d_{23}d_{31}d_{12}D} (r_1 d_{23} + r_2 d_{31} + r_3 d_{12}) (r_1 d_{23} + r_2 d_{31} - r_3 d_{12})$$

$$(r_1 d_{23} - r_2 d_{31} + r_3 d_{12}) (-r_1 d_{23} + r_2 d_{31} + r_3 d_{12}),$$

where $D = d_{23}d_{31}d_{12} + r_1^2 d_{23} + r_2^2 d_{31} + r_3^2 d_{12}$.

We can now prove Theorem 5.1.

Proof of Theorem 5.1. Consider four distinct spheres with infinitely many real common tangents. By Propositions 5.3 and 5.4, the centers of at least three of the spheres are aligned.

If these three spheres intersect in a circle, their common tangents are the tangents to that circle in its plane. To be tangent to infinitely many of these lines, the fourth sphere has to contain that circle (and, if that circle is degenerate to a point, the four spheres must have the same tangent plane at this point). Thus all four spheres have aligned centers.

If the three spheres with aligned centers do not have a common intersection, then by Lemma 5.10 their common tangents are the rulings of a single quadric having their axis as axis of revolution. To be tangent to infinitely many lines contained in this quadric, the fourth sphere must have its center on the axis of the quadric (and adequate radius as determined below), hence the four spheres have aligned centers.

Conversely, four spheres with aligned centers and at least one common tangent have infinitely many common tangents, by symmetry of revolution. This concludes the proof of Theorem 5.1 and provides the finer geometric characterization stated in Section 5.1. \square

As shown above, four spheres with collinear centers and no common intersection admit infinitely many real common tangents if and only if there exists a conic (5.22) whose coefficients A, B, C satisfy Equation (5.23) for all $i = 1, \dots, 4$. These four equations admit a solution if and only if the relation obtained by eliminating A, B, C is satisfied. One can put the result in the permutation invariant form in terms of the oriented distances $d_{ij} = \alpha_j - \alpha_i$ and the radii r_k :

$$\sum_{k=1}^4 \frac{r_k^2}{\prod_{i \neq k} d_{ki}} = 0. \quad (5.24)$$

In order to obtain infinitely many *real* common tangents, the coefficients A, B, C must also satisfy the semi-algebraic conditions

$$A \leq 0, \quad B^2 - 4AC \leq 0 \quad (5.25)$$

noted in (5.22). A and $B^2 - 4AC$ can be obtained in terms of the d_{ij} and r_k by solving the system of equations, as illustrated after Lemma 7.

The case of four spheres intersecting in a common circle or tangent in a common point is a limit case of the situation above, and thereby subject to the same algebraic and semi-algebraic conditions.

Remark 5.12. When a configuration of four spheres is given in terms of the Cartesian coordinates of the four centers $c_i = (x_i, y_i, z_i)$ and the corresponding radii r_i , expressing the collinearity of the centers involves quadratic equations in their coordinates, and, in view of $\frac{d_{ij}}{d_{ik}} = \frac{x_j - x_i}{x_k - x_i} = \frac{y_j - y_i}{y_k - y_i} = \frac{z_j - z_i}{z_k - z_i}$ and $d_{ij}^2 = (x_j - x_i)^2 + (y_j - y_i)^2 + (z_j - z_i)^2$, testing Conditions (5.24) and (5.25) amounts to evaluating polynomials of degree at most five in the Cartesian coordinates and radii.

5.6 Conclusion

This paper answers a question left open for several years by characterizing the sets of four spheres of various radii with infinitely many common tangent lines. This completes the description of degeneracies for common tangents to spheres in \mathbb{R}^3 .

Some of our results generalize to the case of quadric surfaces. In a companion paper [BGLP03] we characterize the families of quadrics in $\mathbb{P}^3(\mathbb{C})$ whose common tangents sweep another quadric surface. The result of the present paper appears as a particular case obtained by considering *real* tangents to *real* spheres. Extending our characterization to quadruples of quadrics with infinitely many real common tangents remains an open problem.

Results of the kind proved in this paper have applications in the field of 3D visibility. Given a 3D scene, combinatorial changes appearing in the view of a moving observer occur when traversing special surfaces known as *visual event surfaces*. Such surfaces are swept by lines having prescribed contact with the objects of the scene. Various data structures based on visual events, like the visibility complex or the visibility skeleton [Dur00], have been proposed to speed up visibility computations. The 0-dimensional elements of these structures appear as discrete lines tangent to four objects. Failing to recognize that four objects admit infinitely many tangent lines leads to errors in the computations of these types of data structures. Hence, recognizing configurations of four objects with infinitely many tangent lines is crucial to the robustness of visibility computations. Our theorem settles the case of four spherical objects in \mathbb{R}^3 .

Chapitre 6

Transversals to line segments in three-dimensional space

Cet article a été publié dans *Discrete Computational Geometry* [BEL⁺05].

Abstract

We completely describe the structure of the connected components of transversals to a collection of n line segments in \mathbb{R}^3 . Generically, the set of transversal to four segments consist of zero or two lines. We catalog the non-generic cases and show that $n \geq 3$ arbitrary line segments in \mathbb{R}^3 admit at most n connected components of line transversals, and that this bound can be achieved in certain configurations when the segments are coplanar, or they all lie on a hyperboloid of one sheet. This implies a tight upper bound of n on the number of geometric permutations of line segments in \mathbb{R}^3 .

6.1 Introduction

A k -transversal to a family of convex sets in \mathbb{R}^d is an affine subspace of dimension k (e.g. a point, line, plane, or hyperplane) that intersects every member of the family. Goodman, Pollack, and Wenger [GPW93a] and Wenger [Wen98] provide two extensive surveys of the rich subject of geometric transversal theory. In this paper, we are interested in 1-transversals (also called line transversals, or simply transversals) to line segments. In \mathbb{R}^2 , this question was studied in the 1980's by Edelsbrunner et al. [EMP⁺82] : they proved that the set of transversals to n line segments has total description complexity $O(n)$ and can be computed in $O(n \log n)$ time ; moreover, it follows from their work that the set of transversals consists of up to n connected components (see Section 6.3.3). Here we study the subject in \mathbb{R}^3 .

We address the following basic question : What is the cardinality and geometry of the set of transversals to an arbitrary collection of n line segments in \mathbb{R}^3 ? Here a segment may be open, semi-open, or closed, and it may degenerate to a point ; segments may intersect or even overlap. Since a line in \mathbb{R}^3 has four degrees of freedom, it can intersect at most four lines or line segments in generic position. Conversely, it is well-known that four lines or line segments in generic position admit zero or two transversals ; moreover, four arbitrary lines in \mathbb{R}^3 admit zero, one, two, or infinitely many transversals [HCV52, p. 164]. In contrast, our work shows that four arbitrary line segments admit up to four or infinitely many transversals.

Our interest in line transversals to segments in \mathbb{R}^3 is motivated by visibility problems. In computer graphics and robotics, scenes are often represented as unions of not necessarily disjoint polygonal or polyhedral objects. The objects that can be seen in a particular direction from a moving viewpoint may change when the line of sight becomes tangent to one or more objects in the scene. Since the line of sight then becomes a transversal to a subset of the edges of the polygons and polyhedra representing the scene, questions about transversals to segments arise very naturally in this context.

As an example, the visibility complex [DDP02, PV96b] and its visibility skeleton [DDP97] are data structures that encode visibility information of a scene; an edge of these structures corresponds to a set of segments lying in line transversals to some k edges of the scene. Generically in \mathbb{R}^3 , k is equal to three. In degenerate configurations, however, k can be arbitrarily large. Such degenerate configurations can arise, for instance in architectural scenes, which frequently contain many coplanar edges. It is thus essential for computing these data structures to characterize and compute the transversals to k segments in \mathbb{R}^3 . Also, to bound the size of the visibility complex one needs to bound the number of connected components of transversals to k arbitrary line segments. The present paper establishes the actual bound.

As mentioned above, in the context of 3D visibility, lines tangent to objects are more relevant than transversals; lines tangent to a polygon or polyhedron along an edge happen to be transversals to this edge. (For bounds on the space of transversals to convex polyhedra in \mathbb{R}^3 see [PS92].) The literature related to lines tangent to objects falls into two categories. The one closest to our work deals with characterizing the degenerate configurations of curved objects with respect to tangent lines. MacDonald, Pach, and Theobald [MPT01] give a complete description of the set of lines tangent to four unit balls in \mathbb{R}^3 . Megyesi, Sottile, and Theobald [MST03] describe the set of lines meeting two lines and tangent to two spheres in \mathbb{R}^3 , or tangent to two quadrics in \mathbb{P}^3 . Megyesi and Sottile [MS05] describe the set of lines meeting one line and tangent to two or three spheres in \mathbb{R}^3 . A nice survey of these results can be found in [The97]. Very recently, in an as yet unpublished manuscript, Borcea, Goaoc, Lazard, and Petitjean completed this study by characterizing the set of lines tangent to four spheres in \mathbb{R}^3 .

The other category of results deals with lines tangent to k among n objects in \mathbb{R}^3 . For polyhedral objects, de Berg, Everett, and Guibas [dBEG98] showed a $\Omega(n^3)$ lower bound on the number of free (i.e., non-occluded by the interior of any object) lines tangent to four among n disjoint homothetic convex polyhedra. Brönnimann et al. [BDD⁺02] showed that, under a certain general position assumption, the number of lines tangent to four among k bounded disjoint convex polyhedra of total complexity n is $O(n^2k^2)$. For curved objects, Devillers et al. [DDE⁺03] presented a simple $\Omega(n^2)$ lower bound on the number of free maximal segments tangent to four among n unit balls, and give a bound of $\Omega(n^3)$ (due to Devillers and Ramos) for n arbitrarily sized balls. Agarwal, Aronov, and Sharir [AAS99] showed an upper bound of $O(n^{3+\epsilon})$ on the complexity of the space of line transversals to n balls; recently, with Koltun, they showed that the same upper bound holds for the complexity of the set of lines that do not intersect n balls [AAKS05]. Durand et al. [DDP02] showed an upper bound of $O(n^{8/3})$ on the expected number of possibly occluded lines tangent to four among n uniformly distributed unit balls. Under the same model, Devillers et al. [DDE⁺03] recently showed a bound of $\Theta(n)$ on the expected number of maximal free line segments tangent to four among n balls.

A topic closely related to line transversals is that of geometric permutations. A *geometric permutation* of pairwise disjoint convex objects in \mathbb{R}^d is an ordering of the objects (or its reverse) such that the objects are met in that order by a line transversal. Worst-case bounds for general convex objects are known: $2n - 2$ is tight in two dimensions [ES90], while in any dimension the best known bounds are $\Omega(n^{d-1})$ [KLL92] and $O(n^{2d-2})$ [Wen90]. The gap was closed for spheres by Smorodinsky et al. [SMS00], who showed that n spheres in \mathbb{R}^d admit up to $\Theta(n^{d-1})$ geometric permutations, and the same bound was also shown true for “fat” objects [KV01]. Recently, Cheong et al. [CGN05] improved the known bounds for congruent balls, by showing that n balls in \mathbb{R}^d of same radius admit at most two geometric permutations if $n \geq 9$, and at most three otherwise.

6.2 Our results

We say that two transversals to a collection of line segments are in the same *connected component* if and only if one of the transversals can be continuously moved into the other while remaining a transversal in \mathbb{R}^3 to the collection of line segments. (For the sets of line transversals considered here, the notions of connected and path-connected components are equivalent since all sets are semi-algebraic.) Equivalently, the two points in line space (e.g., in Plücker space [PW01]) corresponding to the two transversals are in the same connected component of the set of points corresponding to all the transversals in \mathbb{R}^3 to the collection of line segments.

Our main result is the following theorem.

Theorem 6.1. *A collection of $n \geq 3$ arbitrary line segments in \mathbb{R}^3 admits any number from 0 to n of connected components of line transversals. More precisely, the set of line transversals consists of at most two isolated lines unless the segments lie in one of the following three configurations :*

1. the n segments are all contained in lines of one ruling of (a) a hyperbolic paraboloid or (b) a hyperboloid of one sheet, or
2. they are all concurrent, or
3. they all lie in a plane, with the possible exception of a group of one or more segments that all meet that plane at the same point.

In cases 1(a) and 2, the transversals form at most one connected component. In cases 1(b) and 3, the transversals can have any number from 0 to n of connected components. Moreover, in case 3, if all segments are not coplanar, this number is at most $n - 1$.

In cases 1–3, each connected component can consist of infinitely many lines or reduce to an isolated line. For example, three segments forming a triangle and a fourth segment intersecting the interior of the triangle in one point have exactly three transversals (Figure 6.2b shows a similar example with infinitely many transversals). Also, the four segments in Figure 6.1 can be shortened so that the four connected components of transversals reduce to four isolated transversals.

A simple consequence of our theorem is the following bound on the number of geometric permutations of n segments in \mathbb{R}^3 .

Corollary 6.2. *A set of $n \geq 3$ pairwise disjoint segments in \mathbb{R}^3 admits up to n geometric permutations and this bound is tight.*

Proof. By the theorem above, n segments in \mathbb{R}^3 admit up to n connected components of line transversals. Within a connected component, the lines transversals must intersect the segments in the same order. Otherwise by continuity there would exist a line in that component where two objects would intersect somewhere on that line, a contradiction. Hence the upper bound. The lower bound is proved by the configuration of Figure 6.1 generalized to n segments : the n geometric permutations are all the permutations of the form $(i, i + 1, \dots, n, 1, \dots, i - 1)$ for $1 \leq i \leq n$. \square

Finally, as discussed in the conclusion, an $O(n \log n)$ -time algorithm for computing the transversals to n segments follows directly from the proof of Theorem 6.1.

6.3 Proof of Theorem 6.1

Every non-degenerate line segment is contained in its *supporting line*. We define the supporting line of a point to be the vertical line through that point. We prove Theorem 6.1 by considering the following three cases which cover all possibilities but are not exclusive.

1. Three supporting lines are pairwise skew.
2. Two supporting lines are coplanar.
3. All the segments are coplanar.

We can assume in what follows that *the supporting lines are pairwise distinct*. Indeed, if disjoint segments have the same supporting line ℓ , then ℓ is the only transversal to those segments, and so the set of transversals is either empty or consists of ℓ and the theorem is satisfied. If some non-disjoint segments have the same supporting line, then any transversal must meet the intersection of the segments. In that case, we can replace these overlapping segments by their common intersection and the theorem for the smaller collection will imply the result for the original collection.

6.3.1 Three supporting lines are pairwise skew

Three pairwise skew lines lie on a unique doubly-ruled hyperboloid, namely, a hyperbolic paraboloid or a hyperboloid of one sheet (see the discussion in [PW01, §3]). Furthermore, they are members of one ruling, say the “first” ruling, and their transversals are the lines in the “second” ruling that are not parallel to any of the three given skew lines.

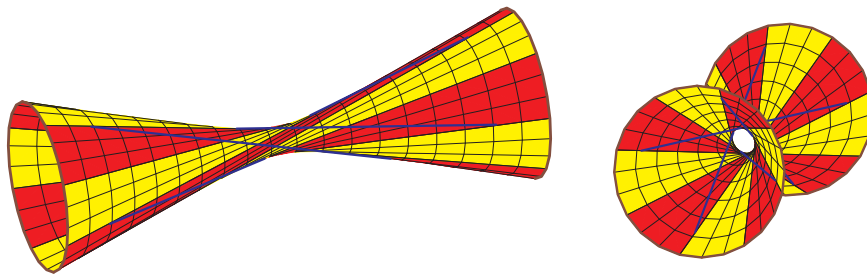


FIG. 6.1 – Two views of a hyperboloid of one sheet containing four line segments and their four connected components of transversals (corresponding to the shaded regions). The four segments are symmetric under rotation about the axis of the hyperboloid.

Consider first the case where there exists a fourth segment whose supporting line ℓ does not lie in the first ruling. Either ℓ is not contained in the hyperboloid or it lies in the second ruling. In both cases, there are at most two transversals to the four supporting lines, which are lines of the second ruling that meet or coincide with ℓ [HCV52, p. 164]. Thus there are at most two transversals to the n line segments.

Now suppose that all the $n \geq 3$ supporting lines of the segments s_i lie in the first ruling of a hyperbolic paraboloid. The lines in the second ruling can be parameterized by their intersection points with any line r of the first ruling. Thus the set of lines in the second ruling that meet a segment s_i corresponds to an interval on line r . Hence the set of transversals to the n segments corresponds to the intersection of n intervals on r , that is, to one interval on this line, and so the transversals form one connected component.

Consider finally the case where the $n \geq 3$ supporting lines lie in the first ruling of a hyperboloid of one sheet (see Figure 6.1). The lines in the second ruling can be parameterized by points on a circle, for instance, by their intersection points with a circle lying on the hyperboloid of one sheet. Thus the set of transversals to the n segments corresponds to the intersection of n intervals on this circle. This intersection can have any number of connected components from zero up to n , and any of these connected components may consist of an isolated point on the circle. The set of transversals can thus have any number of connected components from zero up to n , and any of these connected components may consist of an isolated transversal. Figure 6.1 shows two views of a configuration with $n = 4$ line segments having four connected components of transversals.

6.3.2 Two supporting lines are coplanar

Let ℓ_1 and ℓ_2 be two (distinct) coplanar supporting lines in a plane H . First consider the case where ℓ_1 and ℓ_2 are parallel. Then the transversals to the n segments all lie in H . If some segment does not intersect H then there are no transversals; otherwise, we can replace each segment by its intersection with H to obtain a set of coplanar segments, a configuration treated in Section 6.3.3.

Now suppose that ℓ_1 and ℓ_2 intersect at point p . Consider all the supporting lines not in H . If no such line exists then all segments are coplanar; see Section 6.3.3. If such lines exist and any one of them is parallel to H then all transversals to the n segments lie in the plane containing p and that line. We can again replace each segment by its intersection with that plane to obtain a set of coplanar segments, a configuration treated in Section 6.3.3.

We can now assume that there exists a supporting line not in H . Suppose that all the supporting lines not in H go through p . If all the segments lying in these supporting lines contain p then we may replace all these segments by the point p without changing the set of transversals to the n segments. Then all resulting segments are coplanar, a configuration treated in Section 6.3.3. Now if some segment s does not contain p then the only possible transversal to the n segments is the line containing s and p .

We can now assume that there exists a supporting line ℓ_3 intersecting H in exactly one point q distinct from p (see Figure 6.2(a)). Let K be the plane containing p and ℓ_3 . Any transversal to the lines ℓ_1 , ℓ_2 , and ℓ_3 lies in K and goes through p , or lies in H and goes through q .

If there exists a segment s that lies neither in H nor in K and goes through neither p nor q , then there are at most two transversals to the n segments, namely, at most one line in K through p and s and at most one line in H through q and s .

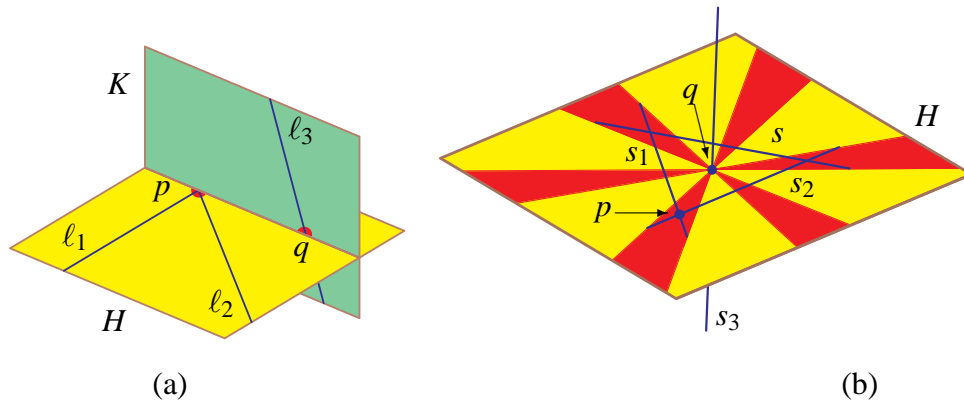


FIG. 6.2 – (a) Lines ℓ_1 and ℓ_2 intersect at point p , and line ℓ_3 intersects plane H in a point q distinct from p . (b) Four segments having three connected components of transversals.

We can thus assume that all segments lie in H or K or go through p or q . If there exists a segment s that goes through neither p nor q , it lies in H or K . If it lies in H then all the transversals to the n segments lie in H (see Figure 6.2(b)). Indeed, no line in K through p intersects s except possibly the line pq which also lies in H . We can again replace each segment by its intersection with H to obtain a set of coplanar segments; see Section 6.3.3. The case where s lies in K is similar.

We can now assume that all segments go through p or q (or both). Let n_p be the number of segments not containing p , and n_q be the number of segments not containing q . Note that $n_p + n_q \leq n$.

Among the lines in H through q , the transversals to the n segments are the transversals to the n_q segments not containing q . We can replace these n_q segments by their intersections with H to obtain a set of n_q coplanar segments in H . The transversals to these segments in H through q can form up to n_q connected components. Indeed, the lines in H through q can be parameterized by a point on a circle, for instance, by their polar angle in $\mathbb{R}/\pi\mathbb{Z}$. Thus the set of lines in H through q and through a segment in H corresponds to an interval of $\mathbb{R}/\pi\mathbb{Z}$. Hence the set of transversals to the n_q segments corresponds to the intersection of n_q intervals in $\mathbb{R}/\pi\mathbb{Z}$ which can have up to n_q connected components.

Similarly, the lines in K through p that are transversals to the n segments can form up to n_p connected components. Note furthermore that the line pq is a transversal to all segments and that the connected component of transversals that contains the line pq is counted twice. Hence there are at most $n_p + n_q - 1 \leq n - 1$ connected components of transversals to the n segments.

To see that the bound of $n - 1$ connected components is reached, first consider $\lfloor n/2 \rfloor$ lines in H through p , but not through q . Their transversals through q are all the lines in H through q , except for the lines that are parallel to any of the $\lfloor n/2 \rfloor$ given lines. This gives $\lfloor n/2 \rfloor$ connected components. Shrinking the $\lfloor n/2 \rfloor$ lines to sufficiently long segments still gives $\lfloor n/2 \rfloor$ connected components of transversals in H through q . The same construction with $\lceil n/2 \rceil$ line segments in plane K gives $\lceil n/2 \rceil$ connected components of transversals in K through p . This gives $n - 1$ connected components of transversals to the n segments since the component containing the line pq is counted twice. Figure 6.3(a) shows an example of four segments having three connected components of transversals.

6.3.3 All the segments are coplanar

Let H be the plane containing all the n segments. There exists a transversal not in H if and only if all segments are concurrent at a point p . In this case, the transversals consist of the lines through p together with the transversals lying in H . To see that they form only one connected component, notice that any transversal in H can be translated to p while remaining a transversal throughout the translation. We thus can assume in the following that all transversals lie in H , and we consider the problem in \mathbb{R}^2 .

We consider the usual geometric transform (see e.g. [EMP⁺82]) where a line in \mathbb{R}^2 with equation $y = ax + b$ is mapped to the point (a, b) in the dual space. The transversals to a segment are transformed to a double wedge; the double wedge degenerates to a line when the segment is a point. The apex of the double wedge is the dual of the line containing the segment.

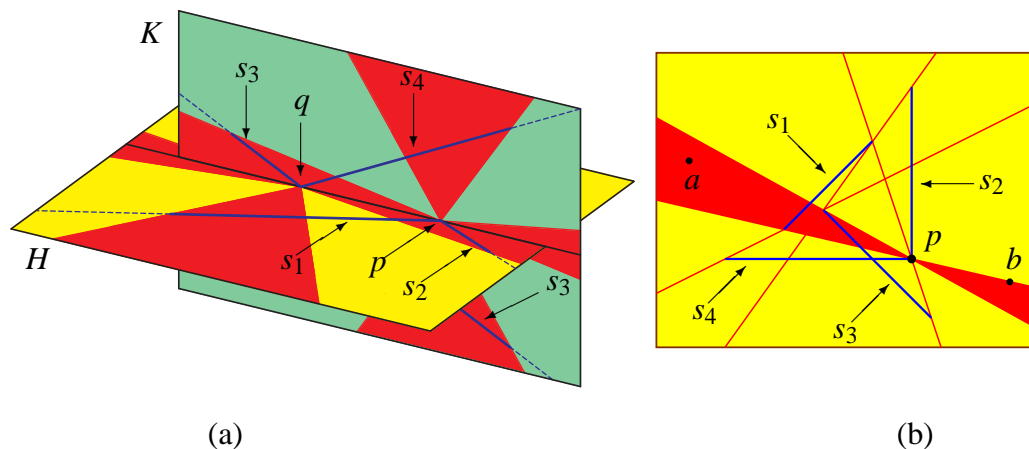


FIG. 6.3 – (a) Four segments having three connected components of transversals. (b) Four coplanar segments having four connected components of transversals.

A transversal to the n segments is represented in the dual by a point in the intersection of all the double wedges. There are at most $n + 1$ connected components of such points [EMP⁺82] (see also [Ede87a, Lemma 15.3]). Indeed, each double wedge consists of two wedges separated by the vertical line through the apex. The intersection of all the double wedges thus consists of at most $n + 1$ convex regions whose interiors are separated by at most n vertical lines.

Notice that if there are exactly $n + 1$ convex regions then two of these regions are connected at infinity by the dual of some vertical line, in which case the segments have a vertical transversal. Thus the number of connected components of transversals is at most n .

To see that this bound is sharp consider the configuration in Figure 6.3(b) of four segments having four components of transversals. Three of the components consist of isolated lines and one consists of a connected set of lines through p (shaded in the figure). Observe that the line segment ab meets the three isolated lines. Thus the set of transversals to the four initial segments and segment ab consists of the three previously mentioned isolated transversals, the line pb which is isolated, and a connected set of lines through p . This may be repeated for any number of additional segments, giving configurations of n coplanar line segments with n connected components of transversals.

6.4 Algorithmic considerations and conclusion

While algorithmic issues have not been the main concern of the paper, we note that the proof of Theorem 6.1 leads to an $O(n \log n)$ -time algorithm in the real RAM model of computation. First reduce in $O(n \log n)$ time the set of segments to the case of pairwise distinct supporting lines. Choose any three of these lines. Either they are pairwise skew or two of them are coplanar. If they are pairwise skew (see Section 6.3.1), their transversals, and hence the transversals to all n segments, lie in one ruling of a hyperboloid. Any segment that intersects the hyperboloid in at most two points admits at most two transversals that lie in that ruling. Checking whether these lines are transversals to the n segments can be done in linear time. Consider now the case of a segment that lies on the hyperboloid. Its set of transversals, lying in the ruling, can be parameterized in constant time by an interval on a line or a circle depending on the type of the hyperboloid. Computing the transversals to the n segments thus reduces in linear time to intersecting n intervals on a line or on a circle, which can be done in $O(n \log n)$ time. If two supporting lines are coplanar (see Section 6.3.2), computing the transversals to the n segments reduces in linear time to computing transversals to at most n segments in one or two planes, which can be done in $O(n \log n)$ time [EMP⁺82].

Finally, note that we did not consider in this paper, for simplicity, segments that can extend to lines or half-lines in \mathbb{R}^3 although our theorem holds in those situations as well. For example, in \mathbb{R}^3 , the transversals to $n \geq 3$ lines of one ruling of a hyperboloid of one sheet are all the lines of the other ruling with the exception of the lines

parallel to the n given lines. Thus, in \mathbb{R}^3 , the transversals form n connected components. Notice however that our theorem does not hold for lines in projective space \mathbb{P}^3 ; in this case, our proof directly yields that, if a set of lines admit infinitely many transversals, they form one connected component.

Acknowledgments

This research was initiated at the Second McGill-INRIA Workshop on Computational Geometry in Computer Graphics, February 7–14, 2003, co-organized by H. Everett, S. Lazard, and S. Whitesides, and held at the Belairs Research Institute of McGill University. We would like to thank the other participants of the workshop for useful discussions, and in particular Xavier Goaoc for pointing out the implication on the number of geometric permutations of segments in \mathbb{R}^3 .

Chapitre 7

Lines tangent to four triangles in three-dimensional space

Cet article a été accepté pour publication dans *Discrete Computational Geometry* [BDLS07].

Abstract

We investigate the lines tangent to four triangles in \mathbb{R}^3 . By a construction, there can be as many as 62 tangents. We show that there are at most 162 connected components of tangents, and at most 156 if the triangles are disjoint. In addition, if the triangles are in (algebraic) general position, then the number of tangents is finite and it is always even.

7.1 Introduction

Motivated by visibility problems, we investigate lines tangent to four triangles in \mathbb{R}^3 . In computer graphics and robotics, scenes are often represented as unions of not necessarily disjoint polygonal or polyhedral objects. The objects that can be seen in a particular direction from a moving viewpoint may change when the line of sight becomes tangent to one or more objects in the scene. Since this line of sight is tangent to a subset of the edges of the polygons and polyhedra representing the scene, we are also led to questions about lines tangent to segments and to polygons. Four polygons will typically have finitely many common tangents, while five or more will have none and three or fewer will have either none or infinitely many.

This paper follows a series of papers by the authors and their collaborators investigating such questions. The paper [BDD⁺07] investigated the lines of sight tangent to four convex polyhedra in a scene of k convex but not necessarily disjoint polyhedral objects, and proved that there could be up to but no more than $\Theta(n^2k^2)$ connected components of such lines. The same bound for the considerably easier case of disjoint convex polyhedra in algebraic general position was proved earlier [EGHHZ00, BDD⁺02]. The paper [BEL⁺05] offers a detailed study of transversals to n line segments in \mathbb{R}^3 and proved that although there are at most two such transversals for four segments in (algebraic) general position, there are at most n such connected components of transversals in any case. Dealing with curved objects in \mathbb{R}^3 , the paper [BGLP06] studies the tangent lines to four arbitrary spheres and [DDE⁺03] shows that there is a linear expected number of maximal non-occluded line segments tangent to four among n uniformly distributed unit balls.

Halperin and Sharir [HS94], and Pellegrini [Pel94], proved that, in a polyhedral terrain, the set of free lines with n edges has near-cubic complexity. De Berg, Everett and Guibas [dBEG98] 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. Recently, Agarwal et al. [AAKS05] proved that the set of free lines among n unit balls has complexity $O(n^{3+\epsilon})$. For related books and surveys, see [Ede87b, GPW93b, PW01, Wen98].

In this paper, we consider the case of four triangles in \mathbb{R}^3 , and establish lower and upper bounds on the number of tangent lines.

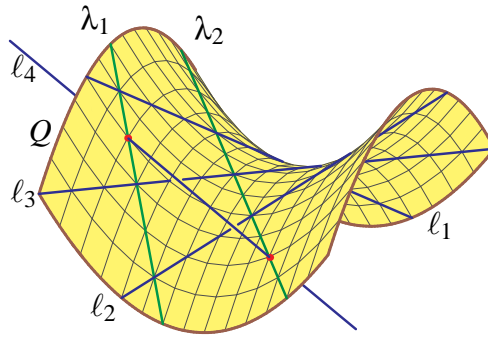


FIG. 7.1 – The lines l_1 , l_2 and l_3 span a hyperbolic paraboloid Q which meets line l_4 in two points. The two lines λ_1 and λ_2 are the transversals to the lines l_1 , l_2 , l_3 , and l_4 .

A *triangle* in \mathbb{R}^3 is the convex hull of three distinct (and non-collinear) points in \mathbb{R}^3 . A line is *tangent* to a triangle if it meets an edge of the triangle. Note that a line tangent to each of four triangles forming a scene corresponds to an unoccluded line of sight in that scene. If there are $k > 4$ triangles, then the bound $\Theta(k^4)$ of [BDD⁺07] stands (as the total number of edges is $n = 3k$ and one of the lower bound examples is made of triangles). We thus investigate the case of four triangles. Let $n(t_1, t_2, t_3, t_4)$ be number of lines tangent to four triangles t_1 , t_2 , t_3 , and t_4 in \mathbb{R}^3 . This number may be infinite if the lines supporting the edges of the different triangles are not in general position.

Our first step is to consider the algebraic relaxation of this geometric problem in which we replace each edge of a triangle by the line in \mathbb{CP}^3 supporting it, and then ask for the set of lines in \mathbb{CP}^3 which meet one supporting line from each triangle. Since there are $3^4 = 81$ such quadruples of supporting lines, this is the disjunction of 81 instances of the classical problem of transversals to four given lines in \mathbb{CP}^3 . As there are two such transversals to four given lines in general position, we expect that this algebraic relaxation has 162 solutions. We say that four triangles t_1, t_2, t_3, t_4 are in (algebraic) *general position* if each of the 81 quadruples of supporting lines have two transversals in \mathbb{CP}^3 and all 162 transversals are distinct. Let \mathcal{T} be the configuration space of all quadruples of triangles in \mathbb{R}^3 and $T \subset \mathcal{T}$ consist of those quadruples which are in general position. Thus if $(t_1, t_2, t_3, t_4) \in T$, the number $n(t_1, t_2, t_3, t_4)$ is finite and is at most 162.

Our primary interest is the number

$$N := \max\{n(t_1, t_2, t_3, t_4) \mid (t_1, t_2, t_3, t_4) \in T\}.$$

Our results about this number N are two-fold. First, we show that $N \geq 62$.

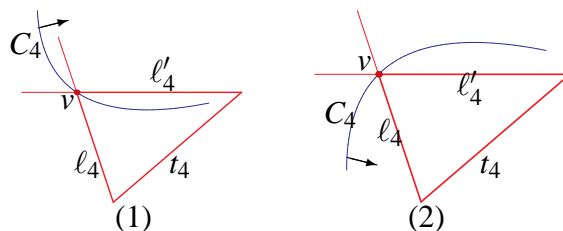
Theorem 7.1. *There are four disjoint triangles in T with 62 common tangent lines.*

The idea is to perturb a configuration of four lines in \mathbb{R}^3 with two real transversals, such as in Figure 7.1. The triangles in our construction are very ‘thin’—the smallest angle among them measures about 10^{-11} degrees. We ran a computer search for ‘fatter’ triangles having many common tangents, checking the number of tangents to 5 million different quadruples of triangles. It appears that random quadruples of realistic triangles often have a fair number of common tangents. Several had as many as 40 common tangents, and quadruples that admit common tangents have 16 tangents or more with probability at least 15%. This is discussed in Section 7.5.

We can improve the upper bound on N when the triangles are disjoint.

Theorem 7.2. *Four triangles in T admit at most 162 distinct common tangent lines. This number is at most 156 if the triangles are disjoint.*

We believe, however, that the upper bounds we give here are far from optimal. When the four triangles are not in general position, the number of tangent lines can be infinite. In this case, we may group these tangents by connected components: two line tangents are in the same component if one may move continuously between the two lines while staying tangent to the four triangles. Each quadruple of edges may induce up to four components of tangent lines [BEL⁺05], giving a trivial upper bound of 324. This may be improved.

FIG. 7.2 – Configuration in plane π_4 .

Theorem 7.3. *Four triangles have at most 162 connected components of common tangents. If the triangles are disjoint, then this number is at most 156.*

We have one more result which we do not prove in this paper, but is proved in the companion research report [BDLS05] and is relevant to mention here.

Theorem 7.4. *If $(t_1, t_2, t_3, t_4) \in T$, then $n(t_1, t_2, t_3, t_4)$ is even.*

This result may not seem surprising as complex roots come in conjugate pairs. However, this usual argument does not apply because we seek tangents to triangles and not transversals to lines. Frequently, only one of two real transversals to a quadruple of supporting lines is tangent to the triangles. The main new idea behind Theorem 7.4 is that such tangent lines essentially come in pairs.

In our proof of Theorem 7.4, we consider four *moving* triangles, and show that common tangents are created and destroyed in pairs, and so the parity of $n(t_1, t_2, t_3, t_4)$ does not change. There are two cases to consider. The first is when two real tangents which are transversal to the same four edges coalesce and become a pair of complex conjugate transversals; this is the ‘usual’ argument. The second case is when a real transversal to edges e_1, e_2, e_3 , and e_4 moves off of e_4 and is thus no longer tangent to the four triangles. In doing so, it must pass through a vertex v of e_4 . In this case, there is a real transversal to edges e_1, e_2, e_3 , and some *other* edge e'_4 meeting v which simultaneously moves off of e'_4 , also passing through the vertex v . Theorem 7.4 follows as there are triangles in T with no common tangents. We give a complete proof in the research report version [BDLS05].

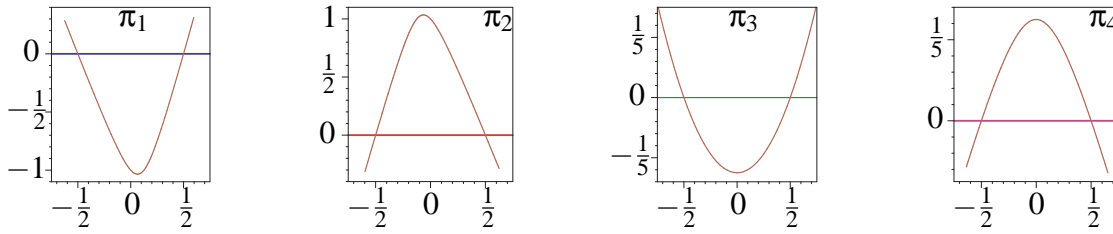
Theorems 7.1, 7.2, and 7.3 are proved in Sections 7.2, 7.3, and 7.4, respectively. Section 7.5 discusses our search for ‘fat’ triangles with many common tangents.

7.2 A construction with 62 tangents

Consider the four triangles whose vertices are given in Table 7.1.

t_1	(-10.5, 1, -10.5) (.5628568345479573470378601, 1, .5628568345479573470378601) (.56285683454726874605620706, .99999999999822994290647247, .56285683454726874605620706)
t_2	(-10.5, -1, 10.5) (1.394218989475, -1, -1.394218989475) (1.3942406911811439954597161, -1.0000237884694881275439271, -1.3942406911811439954597161)
t_3	(-9.5, -9.5, .25) (.685825, .685825, .25) (.69121730616063647303519136, .69121730616063647303519136, .26069756890079842876805653)
t_4	(9.5, 0, 0) (-.511, 0, 0) (-1.0873912730501133759642956, 0, -.51645811088049333541289247)

TAB. 7.1 – Four triangles with 62 common tangents.

FIG. 7.3 – Conics in the planes π_i .

Theorem 7.1'. *There are exactly 62 lines tangent to the four triangles of Table 7.1.*

This can be verified by a direct computation. Software is provided on this paper's web page[†]. More illuminating perhaps is our construction. The idea is to perturb a configuration of four lines in \mathbb{R}^3 with two transversals such as in Figure 7.1. The resulting triangles of Theorem 7.1' are very thin. In degrees, their smallest angles are

$$t_1: 6.482 \times 10^{-12}, \quad t_2: 8.103 \times 10^{-5}, \quad t_3: 4.253 \times 10^{-2}, \quad \text{and} \quad t_4: 2.793.$$

The construction. The lines given parametrically by

$$\ell_1 : (t, 1, t), \quad \ell_2 : (t, -1, -t), \quad \ell_3 : (t, t, \frac{1}{4}), \quad \text{and} \quad \ell_4 : (t, 0, 0),$$

have two transversals $\lambda_1 : (\frac{1}{2}, 2t, t)$ and $\lambda_2 : (-\frac{1}{2}, 2t, -t)$.

For each $i = 1, 2, 3, 4$, let Q_i be the hyperboloid spanned by the lines other than ℓ_i . For example, Q_3 has equation $z = xy$. The intersection of Q_i with a plane containing ℓ_i will be a conic which meets ℓ_i in two points (corresponding to the common transversals λ_1 and λ_2 at $t = \pm \frac{1}{2}$). We choose the plane π_i so that these two points lie in the same connected component of the conic. Here is one possible choice

$$\pi_1 : x = z, \quad \pi_2 : x = -z, \quad \pi_3 : x = y, \quad \text{and} \quad \pi_4 : y = 0.$$

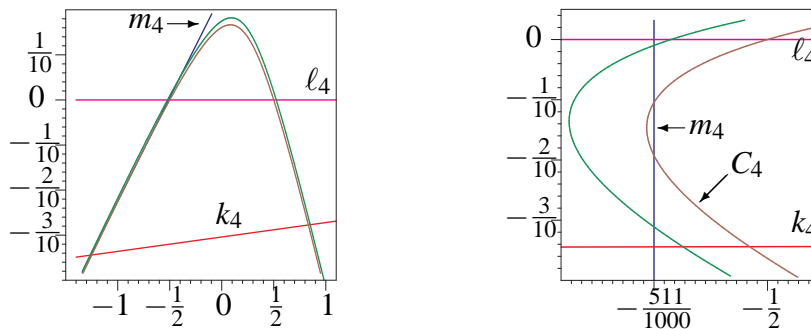
For each i , let C_i be the conic $\pi_i \cap Q_i$, shown in the plane π_i in Figure 7.3. Here, the horizontal coordinate is t , the parameter of the line ℓ_i , while the vertical coordinate is $y-1$ for π_1 , $y+1$ for π_2 , $z-\frac{1}{4}$ for π_3 , and z for π_4 .

For each $i = 1, \dots, 4$, rotate line ℓ_i in plane π_i very slightly about a point that is far from the conic C_i , obtaining a new line k_i in π_i which also meets C_i in two points. Consider now the transversals to $\ell_i \cup k_i$, for $i = 1, \dots, 4$. Because k_i is close to ℓ_i and there were two transversals to $\ell_1, \ell_2, \ell_3, \ell_4$, there will be two transversals to each of the 16 quadruples of lines obtained by choosing one of ℓ_i or k_i for $i = 1, \dots, 4$. By our choice of the point of rotation, all of these will meet ℓ_i and k_i in one of the two thin wedges they form. In this wedge, form a triangle by adding a third side so that the edges on ℓ_i and k_i contain all the points where the transversals meet the lines. The resulting triangles will then have at least 32 common tangents. We claim that by carefully choosing the third side (and tuning the rotations) we are able to get 30 additional tangents.

To begin, look at Figure 7.4 which displays the configuration in π_4 given by the four triangles from Table 7.1. Since the lines ℓ_i and k_i for $i = 1, 2$ are extremely close, the four conics given by transversals to them and to ℓ_3 cannot be resolved in these pictures. The same is true for the four conics given by k_3 , so that each of the apparent two conics are clusters of four nearby conics. The picture on the left is a view of this configuration in the coordinates for π_4 of Figure 7.3. It includes a secant line m_4 to the conics. We choose coordinates on the right so that m_4 is vertical, but do not change the coordinates on ℓ_4 . The horizontal scale has been accentuated to separate the two clusters of conics. The three lines, ℓ_4, k_4 , and m_4 form the triangle t_4 . Let its respective edges be e_4, f_4 , and g_4 . Each edge meets each of the eight conics in two points and these 48 points of intersection give 48 lines tangent to the four triangles.

This last assertion that the 16 lines transversal to m_4 and to $\ell_i \cup k_i$ for $i = 1, 2, 3$ meet the edges of the triangles t_1, t_2 , and t_3 needs justification. Consider for example the transversals to ℓ_1, ℓ_2 , and ℓ_3 . These form a ruling of the doubly-ruled quadric Q_4 and are parameterized by their point of intersection with ℓ_1 . The intersection of Q_4 with π_4 is the conic C_4 . Since the intersections of the conic C_4 with the segment g_4 supported on m_4 lie between its

[†]<http://www.math.tamu.edu/~sottile/stories/4triangles/index.html>

FIG. 7.4 – Configuration in plane π_4 .

intersections with ℓ_4 and k_4 , the corresponding transversals to ℓ_1, ℓ_2, ℓ_3 , and g_4 meet ℓ_1 between points of ℓ_1 met by common transversals to $\ell_4 \cup k_4$ and ℓ_1, ℓ_2 , and ℓ_3 . The same argument for the other lines and for all 8 conics justifies the assertion.

Naïvely, we would expect that this same construction (the third side cutting all eight conics in π_i) could work to select each of the remaining sides of the triangles g_3, g_2 , and g_1 , and that this would give four triangles having $32 + 16 + 16 + 16 + 16 = 96$ common tangents. Unfortunately this is not the case. In the earlier conference version of this paper [BDLS04], we gave a construction that we claimed would yield 88 common tangents. Attempting that construction using Maple revealed a flaw in the argument and the current construction of four triangles with 62 common tangents is the best we can accomplish.

In π_4 , the conics come in two clusters, depending upon whether or not they correspond to ℓ_3 or to k_3 . In order for the edge g_4 to cut all conics, the angle between ℓ_4 and k_4 has to be large, in fact significantly larger than the angle between ℓ_3 and k_3 . Thus in π_3 , the conics corresponding to ℓ_4 are quite far from the conics corresponding to k_4 , and the side g_3 can only be drawn to cut four of the conics, giving eight additional common tangents. Similarly, g_2 can only cut two conics, and g_1 only one. In this way, we arrive at four triangles having $32 + 16 + 8 + 4 + 2 = 62$ common tangents, which has been verified by computer. \square

7.3 Upper bound for disjoint triangles in general position

Four triangles in general position have at most 162 common tangents. If the triangles are disjoint, we slightly improve this upper bound to 156. Our method will be to show that not all $81 = 3^4$ quadruples of edges can give rise to a common tangent. Our proof follows that for the upper bound on the number of tangents to four polytopes [BDD⁺02], limiting the number of configurations for disjoint triangles in \mathbb{R}^3 . We divide the proof into two lemmas, which do not assume general position. The application of the lemmas to the proof of [BDD⁺02], however, requires the general position assumption.

In order for a tangent to meet an edge e , the plane it spans with e must meet one edge from each of the other triangles. A triple of edges, one from each of the other triangles, is *contributing* if there is a plane containing e which meets the three edges. We say that an edge e *stabs* a triangle t if its supporting line meets the interior of t .

Lemma 7.5. *Let e be an edge of some triangle. If e stabs exactly one of the other triangles, then there are at most 26 contributing triples of edges. If e stabs no other triangle, then there are at most 25 contributing triples.*

It is not hard to see that if e stabs at least two of the other triangles, then each of the $27 = 3^3$ triples of edges can be contributing.

Proof. Suppose that e is an edge of some triangle. Let $\pi(\alpha)$ be the pencil of planes containing e . (This is parameterized by the angle α .) For each edge f of another triangle t , there is an interval of angles α for which $\pi(\alpha)$ meets f . Figure 7.5 illustrates the two possible configurations for these intervals, which depend upon whether or not e stabs the triangle t . The intervals are labeled 1, 2, and 3 for the three edges of t . When e stabs t , these intervals cover the entire range of α and the picture is actually wrapped. Call this a *stabbing diagram*. When the supporting line of e does not meet t , these intervals do not cover the entire range of α , and there are two endpoints and one *interior*




FIG. 7.5 – Stabbing and non-stabbing configurations.

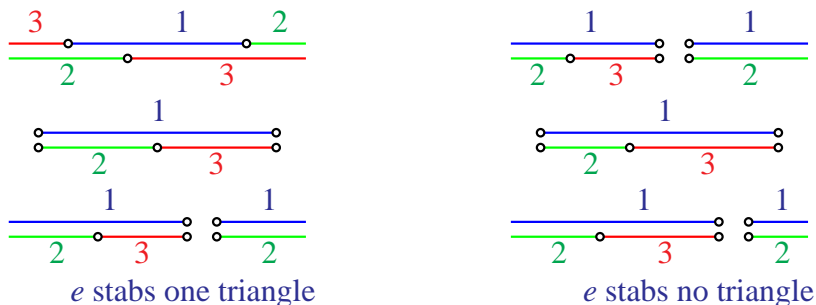


FIG. 7.6 – Configurations with 26 and 25 contributing triples.

vertex of the diagram. If the supporting line of e meets an edge of t , then the two endpoints of the non-stabbing diagram wrap around and coincide. Call either of these last two configurations a non-stabbing diagram.

To count contributing triples, we line up (overlay) diagrams from each of the three triangles not containing e and count how many of the 27 triples $\{1, 2, 3\}^3$, one from each triangle, occur at some value of α . For example, Figure 7.6 displays a configuration with 26 contributing triples (where e stabs a single triangle) and a configuration with 25 contributing triples (e stabs no other triangles). The configuration on the left is missing the triple $(2, 3, 3)$, while the configuration on the right is missing the triples $(2, 2, 3)$ and $(3, 3, 2)$.

These configurations are the best possible. Indeed, begin with two non-stabbing diagrams in which all 9 pairs of edges occur. (If only 8 pairs occurred, there would be at most 24 contributing triples.) The unique way to do this up to relabeling the edges is given by the lower two diagrams in either picture in Figure 7.6. These two diagrams divide the domain of α into six intervals (the two at the ends are wrapped). The five pairs involving 1 occur in two intervals, but four exceptional pairs $\{(2, 2), (2, 3), (3, 2), (3, 3)\}$ occur uniquely in different intervals.

Consider now a third diagram. An exceptional pair extends to three contributing triples only if all three sides in the third diagram meet the interval corresponding to that pair. If the third diagram is stabbing, then one of its three vertices lies in that interval—thus there is at least one triple which does not contribute. If the third diagram is non-stabbing, then either the middle vertex or else both endpoints must lie in that interval—thus there are at least two triples which do not contribute. \square

Lemma 7.6. *At most 78 quadruples of edges of four disjoint triangles can lead to a common tangent.*

Proof. First consider the maximum number of stabbing edges between two triangles. If the triangles are disjoint, then there are at most three stabbing edges; one triangle could have three edges stabbing the other. Indeed, if at least two supporting lines of a triangle t meet another triangle t' which is disjoint from t , then t lies entirely on one side of the plane supporting t' , and thus no supporting lines of t' can meet t . Figure 7.7(a) shows a configuration in which all three supporting lines of t stab t' .

Consider now the bipartite graph between 12 nodes representing the edges of the four triangles and 4 nodes representing the triangles. This graph has an arc between an edge e and a triangle t if the line supporting e stabs t . (We assume that e is not an edge of t .) We just showed that the edges of one triangle t can have at most three arcs incident on another triangle t' , and so this graph has at most 18 edges.

Let the weight of a triangle be the number of arcs emanating from its edges in this graph. As the graph has at most 18 arcs, at least one triangle has weight less than 5. We argue that there is a triangle of weight at most 3. This is immediate if the graph has 15 or fewer edges. On the other hand, this graph has more structure. If it has 18 edges, then all pairs of triangles are in the configuration of Figure 7.7(a), and so every triangle has weight a multiple of 3, which implies that some triangle has weight at most 3. If the graph has 17 edges, then there is exactly one pair

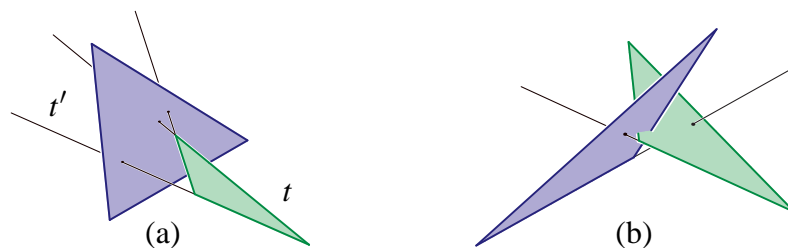


FIG. 7.7 – (a) Two disjoint triangles can have at most 3 stabbing lines. (b) Two intersecting triangles may have up to four.

of triangles with only two stabbing edges, and so the possible weights less than 5 are 0, 2, and 3. If the graph has 16 edges, then there is one pair with only one edge stabbing, or two pairs with two edges stabbing. There can be at most two triangles of weight 4, and again we conclude that there is a triangle with weight at most 3.

If a triangle has weight at most three, either all three edges stab a unique triangle, or else one edge stabs no triangles and another edge stabs at most one other triangle. We sum the number of contributing triples over the edges of this triangle. By Lemma 7.5, this sum will be at most $26+26+26=78$ if all three edges stab a unique triangle and at most $27+26+25=78$ if not. This proves the lemma. \square

Remark 7.7. *There exist four disjoint triangles whose bipartite graph has exactly 18 edges. Thus the previous argument cannot be improved without additional ideas. It is conceivable that further restrictions to the bipartite graph may exist, leading to a smaller upper bound.*

Remark 7.8. *This proof does not enable us to improve the bound when the triangles are not disjoint. Two intersecting triangles can induce up to four arcs (see Figure 7.7(b)) and thus the total number of arcs is bounded above by 24. The minimal weight of a triangle is then 6, and the edges of such a triangle could all have degree 2, which leads to no restrictions.*

7.4 Upper bounds on the number of components

Let \mathcal{F} and \mathcal{J} be the sets of quadruples of edges, one from each of four triangles, whose supporting lines have finitely and infinitely, respectively, many common transversals. Let $n_{\mathcal{F}}$ and $n_{\mathcal{J}}$ be the sum over all quadruples of edges in \mathcal{F} and \mathcal{J} , respectively, of the numbers of connected components of common transversals to each quadruple of edges. Note that the number of quadruples in \mathcal{F} and \mathcal{J} is $|\mathcal{F}| + |\mathcal{J}| = 81$.

Consider a connected component c of common transversals to a quadruple of edges $q \in \mathcal{J}$. The arguments of [BEL⁺05] show that c contains a line that meets a vertex of one of the four edges. That line is thus transversal to another quadruple q' of edges. Thus, the connected component c of common transversals to q is connected with a connected component c' of common transversals to q' . If $q' \in \mathcal{F}$ we charge the component $c \cup c'$ to c' . Otherwise q and q' are both in \mathcal{J} and the component $c \cup c'$ is counted twice. The number of connected components of tangents to four triangles is thus at most $n_{\mathcal{F}} + n_{\mathcal{J}}/2$.

Since any four lines admit at most two or infinitely many transversals, $n_{\mathcal{F}} \leq 2|\mathcal{F}|$. Also, any four segments admit at most four connected components of common transversals [BEL⁺05], thus $n_{\mathcal{J}} \leq 4|\mathcal{J}|$. Hence, the number of connected components of tangents to four triangles is at most $2|\mathcal{F}| + 2|\mathcal{J}| = 162$.

This still may overcount the number of connected components of tangents, but further analysis is very delicate. Such complicated arguments are not warranted as we have already obtained the upper bound of 162 common tangents to four triangles in T . As in Section 7.3, if the triangles are disjoint, then not all quadruples of edges can contribute, which lowers this bound to 156.

7.5 Random triangles

We proved Theorem 7.1 by exhibiting four triangles having 62 common tangents. We do not know if that is the best possible. Since the geometric problem of determining the tangents to four triangles is computationally

Number	0	2	4	6	8	10	12	14					
Frequency	1 515 706	331 443	646 150	403 679	637 202	327 159	358 312	238 913					
	16	18	20	22	24	26	28	30	32	34	36	38	40
	253 396	114 046	80 199	44 870	27 726	12 426	5 796	2 016	813	111	30	3	4

TAB. 7.2 – Number of triangles with a given number of tangents, out of 5 000 000 randomly constructed triangles.

Triangle	Vertices		
t_1	$(-4, -731, -336)$	$(297, -507, 978)$	$(824, -62, -359)$
t_2	$(531, -631, -820)$	$(-24, -716, 713)$	$(807, 377, 177)$
t_3	$(586, -205, 952)$	$(861, -774, 235)$	$(-450, 758, 161)$
t_4	$(330, -141, -908)$	$(942, -920, 651)$	$(-226, 489, 968)$

TAB. 7.3 – Four triangles with 40 common tangents.

feasible—it is the disjunction of 81 problems with algebraic degree 2 and simple inequalities on the solutions—we investigated it experimentally.

For this, we generated 5 000 000 quadruples of triangles whose vertices were points with integral coordinates chosen uniformly at random from the cube $[-1000, 1000]^3$. For each, we computed the number of tangents. The resulting frequencies are recorded in Table 7.2. This search consumed over six months of CPU time on 1.2GHz processors at the MSRI and a DEC Alpha machine at the University of Massachusetts in 2004. It is archived on the web page[†] accompanying this article.

In this search, we found four different quadruples of triangles with 40 common tangents, and none with more. Based on this *iiid* random model, we find that the probability that the four triangles have at least one tangent is around 69.7%, and that the expected number of tangents is somewhat around 6.325, with a standard deviation of about 12.93. The vertices of one are given in Table 7.3. These triangles are rather ‘fat’, in that none have very small angles. Contrast that to the triangles of our construction in Section 7.3. In Figure 7.8 we compare these two configurations of triangles. On the left is the configuration of triangles from Table 7.3, together with their 40 common tangents, while on the right is the configuration of triangles having 62 common tangents. The triangles are labeled in the second diagram, as they are hard to distinguish from the lines. As we remarked in Section 7.3, many of the lines are extremely close and cannot be easily distinguished; that is why one can only count eight lines in this picture.

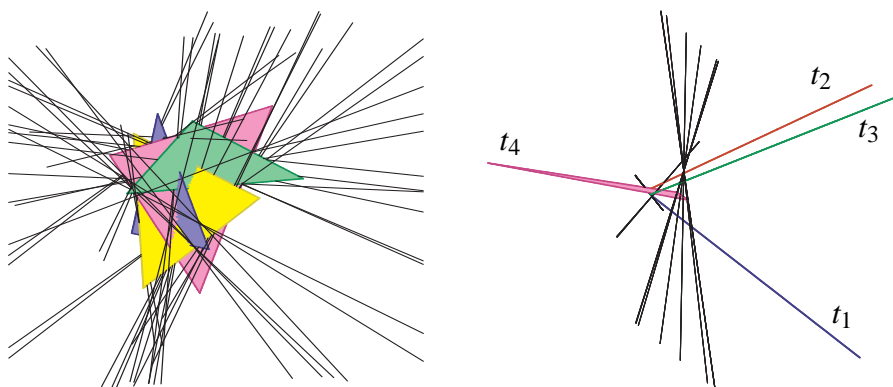


FIG. 7.8 – Triangles with many common tangents.

[†]www.math.tamu.edu/~sottile/stories/4triangles/index.html

Acknowledgments

Research on this paper was initiated at the Second McGill-INRIA Workshop on Computational Geometry in Computer Graphics, February 7–14, 2003, co-organized by H. Everett, S. Lazard, and S. Whitesides, and held at the Bellairs Research Institute of McGill University in Holetown, St. James, Barbados, West Indies. We would like to thank the other participants of the workshop for useful discussions.

Chapitre 8

Predicates for line transversals in 3D

Cet article a été soumis à *Computational Geometry : Theory and Applications* [ELL⁺06]. Une version courte est parue dans les proceedings de la *18th Canadian Conference on Computational Geometry*.

Abstract

In this paper we study various geometric predicates for determining the existence of and categorizing the configurations of lines in 3D that are transversal to lines or segments. We compute the degrees of standard procedures of evaluating these predicates. The degrees of some of these procedures are surprisingly high (up to 168), which may explain why computing line transversals with finite-precision floating-point arithmetic is prone to error. Our results suggest the need to explore alternatives to the standard methods of computing these quantities.

8.1 Introduction

Computing line transversals to lines or segments is an important operation in solving 3D visibility problems arising in computer graphics [BDD⁺07, DD02, DDP97, DDP02, EGHHZ00, PD90]. In this paper, we study various predicates and their degrees concerning line transversals to lines and segments in 3D.

A predicate is a function that returns a value from a discrete set. Typically, geometric predicates answer questions of the type “Is a point inside, outside or on the boundary of a set?”. We consider predicates that are evaluated by boolean functions of more elementary predicates, the latter being functions that return the sign ($-$, 0 or $+$) of a multivariate polynomial whose arguments are a subset of the input parameters of the problem instance (see, for instance [BP00]). By *degree* of a procedure for evaluating a predicate, we mean the maximum degree in the input parameters among all polynomials used in the evaluation of the predicate by the procedure. In what follows we casually refer to this measure as the degree of the predicate. We are interested in the degree because it provides a measure of the number of bits required for an exact evaluation of our predicates when the input parameters are integers or floating-point numbers; the number of bits required is then roughly the product of the degree with the number of bits used in representing each input value.

In this paper, we first study the degree of standard procedures for determining the number of line transversals to four lines or four segments in 3D; recall that four lines in \mathbb{R}^3 admit 0, 1, 2 or an infinite number of line transversals and that four segments admit up to 4 or an infinite number of line transversals [BEL⁺05]. We also consider the predicate for determining whether a minimal (i.e., locally shortest) segment transversal to four line segments is intersected by a triangle. These predicates are ubiquitous in 3D visibility problems. The latter predicate, for instance, can be used for determining whether two triangles see each other in a scene of triangles (that is, for determining whether there exists a segment joining the two triangles and that does not properly intersect any of the other triangles). Finally, we study the predicate for ordering planes through two fixed points, each plane containing a third rational point or a line transversal to four segments or lines. This predicate arises in the rotating plane-sweep algorithm that computes the minimal free segments tangent to four among k convex polyhedra in 3D [BDD⁺07].

Our study shows that standard procedures for solving these predicates have high degrees. We study, in particular, procedures that involve computing the Plücker coordinates of the line transversals involved in the predicates. Throughout the paper, the points defining input geometric primitives (which can be lines, segments, and triangles) are, by assumption, given by their Cartesian coordinates and the degrees of the procedures for evaluating predicates are expressed in these coordinates. We show that, for determining the number of transversals to four lines or four segments, such standard methods lead to procedures of degree 22 or 36, respectively. For determining whether a minimal segment transversal to four line segments is intersected by a triangle, we show that these methods lead to a procedure of degree 78. Also, for ordering, in a rotational sweep about a line, two planes, each defined by a line transversal to four lines, such methods lead to a procedure of degree 144. Furthermore, in some implementations, the Plücker coordinates of the relevant line transversals are computed in a way that the degrees of these procedures are even higher; for instance, the procedure for evaluating the latter predicate for ordering planes then become of degree 168 instead of 144. These very high degrees may help explain why using fixed-precision floating-point arithmetic in implementations for solving 3D visibility problems are prone to errors when given real-world data (see, for instance, [Gla07]).

The degrees we present are tight, that is, they correspond to the maximum degree of the polynomials to be evaluated, in the worst case, in the procedures we consider. It should be stressed that these degrees refer to polynomials used in specific evaluation procedures and we make no claim on the optimality of these procedures.

In the next section we describe a standard method used for computing the line transversals to four lines, which is common to all our predicates. In Section 8.3 we describe the predicates and their degrees. Some experimental results are presented in Section 8.4.

8.2 Computing lines through four lines

We describe here a method for computing the line transversals to four lines in real projective space \mathbb{P}^3 . This method is a variant, suggested by Devillers and Hall-Holt [DHH00] and also described in Redburn [Red03], of that by Hohmeyer and Teller [HT99]; note that, for evaluating predicates, the latter method is not appropriate because it uses singular value decomposition for which we only know of numerical methods and thus the line transversals cannot be computed exactly, when needed.

Each line can be described using Plücker coordinates (see [Sho98], for example, for a review of Plücker coordinates). If a line ℓ in \mathbb{R}^3 is represented by a direction vector \vec{u} and a point p in \mathbb{R}^3 then ℓ can be represented by the six-tuple $(\vec{u}, \vec{u} \times \vec{Op})$ in real projective space \mathbb{P}^5 , where O is any arbitrarily, fixed, origin and \times denote the cross product. The side product \odot of any two six-tuples $\ell = (a_1, a_2, a_3, a_4, a_5, a_6)$ and $k = (x_1, x_2, x_3, x_4, x_5, x_6)$ is $\ell \odot k = a_4x_1 + a_5x_2 + a_6x_3 + a_1x_4 + a_2x_5 + a_3x_6$. The fundamental importance of the side product lies in the fact that a six-tuple $k \in \mathbb{P}^5$ represents a line in 3D if and only if $k \odot k = 0$; this defines a quadric in \mathbb{P}^5 called the Plücker quadric. More generally, recall that two lines intersect in real projective space $\mathbb{P}^3(\mathbb{R})$ if and only if the side product of their Plücker coordinates is zero. Notice that this implies that there is a predicate for determining whether two lines intersect in $\mathbb{P}^3(\mathbb{R})$ which is of degree two in the Plücker coordinates of the lines and, if the lines are each defined by two points, of degree three in the Cartesian coordinates of these points.

Oriented lines of \mathbb{R}^3 , with direction vector \vec{u} and through a point p , can be represented similarly by a six-tuple $(\vec{u}, \vec{u} \times \vec{Op})$ in real oriented projective space (*i.e.*, the quotient of $\mathbb{R}^6 \setminus \{0\}$ by the equivalence relation induced by positive scaling). The sign (positive or negative) of the side operator of two oriented lines ℓ and k then determines on which “side” of ℓ , k lies; for instance, if op and oq are two lines oriented from o to p and from o to q and ℓ is an arbitrarily oriented line such that ℓ , p , q , and o are not coplanar, then $(\ell \odot op)(\ell \odot oq) \leq 0$ if and only if ℓ intersects segment pq (see Figure 8.1(a)).

Given four lines ℓ_1, \dots, ℓ_4 , our problem here is to compute all lines $k = (x_1, x_2, x_3, x_4, x_5, x_6) \in \mathbb{P}^5$ such that $k \odot \ell_i = 0$, for $1 \leq i \leq 4$, which can be written in the following form :

$$\begin{pmatrix} a_4 & a_5 & a_6 & a_1 & a_2 & a_3 \\ b_4 & b_5 & b_6 & b_1 & b_2 & b_3 \\ c_4 & c_5 & c_6 & c_1 & c_2 & c_3 \\ d_4 & d_5 & d_6 & d_1 & d_2 & d_3 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} \quad (8.1)$$

where the rows of the 4×6 matrix contain the Plücker coordinates of the four lines. This can be rewritten as

$$\begin{pmatrix} a_6 & a_1 & a_2 & a_3 \\ b_6 & b_1 & b_2 & b_3 \\ c_6 & c_1 & c_2 & c_3 \\ d_6 & d_1 & d_2 & d_3 \end{pmatrix} \begin{pmatrix} x_3 \\ x_4 \\ x_5 \\ x_6 \end{pmatrix} + \begin{pmatrix} a_4x_1 + a_5x_2 \\ b_4x_1 + b_5x_2 \\ c_4x_1 + c_5x_2 \\ d_4x_1 + d_5x_2 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}. \quad (8.2)$$

Let δ denote the determinant of the above 4×4 matrix. Assuming $\delta \neq 0$, we can solve the system for $x_3, x_4, x_5,$ and x_6 in terms of x_1 and x_2 . Applying Cramer's rule, we get

$$\begin{cases} x_3 = -(\alpha_1x_1 + \beta_1x_2)/\delta \\ x_4 = -(\alpha_2x_1 + \beta_2x_2)/\delta \\ x_5 = -(\alpha_3x_1 + \beta_3x_2)/\delta \\ x_6 = -(\alpha_4x_1 + \beta_4x_2)/\delta \end{cases}$$

where α_i (respectively β_i) is the determinant δ with the i^{th} column replaced by $(a_4, b_4, c_4, d_4)^T$ (respectively $(a_5, b_5, c_5, d_5)^T$). We rewrite this system as

$$\begin{cases} x_1 = -u\delta \\ x_2 = -v\delta \\ x_3 = \alpha_1u + \beta_1v \\ x_4 = \alpha_2u + \beta_2v \\ x_5 = \alpha_3u + \beta_3v \\ x_6 = \alpha_4u + \beta_4v \end{cases} \quad (8.3)$$

with $(u, v) \in \mathbb{P}^1$. Since k is a line, we have $k \odot k = 0$, which implies

$$x_1x_4 + x_2x_5 + x_3x_6 = 0.$$

Substituting in the expressions for $x_1 \dots x_6$, we get

$$Au^2 + Buv + Cv^2 = 0 \quad (8.4)$$

where

$$\begin{aligned} A &= \alpha_1\alpha_4 - \alpha_2\delta, \\ B &= \alpha_1\beta_4 + \beta_1\alpha_4 - \beta_2\delta - \alpha_3\delta, \\ C &= \beta_1\beta_4 - \beta_3\delta. \end{aligned}$$

Solving this degree-two equation in (u, v) and replacing in (8.3), we get (assuming that $A \neq 0$) that the Plücker coordinates of the transversal lines k are :

$$\begin{cases} x_1 = B\delta \mp \delta\sqrt{B^2 - 4AC} \\ x_2 = -2A\delta \\ x_3 = -B\alpha_1 + 2A\beta_1 \pm \alpha_1\sqrt{B^2 - 4AC} \\ x_4 = -B\alpha_2 + 2A\beta_2 \pm \alpha_2\sqrt{B^2 - 4AC} \\ x_5 = -B\alpha_3 + 2A\beta_3 \pm \alpha_3\sqrt{B^2 - 4AC} \\ x_6 = -B\alpha_4 + 2A\beta_4 \pm \alpha_4\sqrt{B^2 - 4AC}. \end{cases} \quad (8.5)$$

Lemma 8.1. *Consider four lines, given by the Cartesian coordinates of pairs of points, that admit finitely many line transversals in $\mathbb{P}^3(\mathbb{R})$. If the four lines are not parallel to a common plane, the Plücker coordinates of their transversals in $\mathbb{P}^3(\mathbb{R})$ can be written as $\phi_i + \varphi_i\sqrt{\Delta}$, $i = 1, \dots, 6$, where ϕ_i, φ_i , and Δ are polynomials of degree at most 17, 6, and 22, respectively, in the coordinates of the input points. Otherwise, the Plücker coordinates of the transversals can be written as polynomials of degree at most 19. Moreover, these bounds are, in the worst case, reached for three of the coordinates.*

Proof. The assumption that the four lines admit finitely many transversals in $\mathbb{P}^3(\mathbb{R})$ ensures that the 4×6 matrix of Plücker coordinates (in (8.1)) has rank 4. Consider first the case where the four input lines are not all parallel to a common plane. Then, the 4×3 matrix of the direction vectors of the four lines has rank 3. By the basis extension theorem, this matrix can be complemented by one of the other columns of the matrix of Plücker coordinates (of (8.1)) in order to get a 4×4 matrix of rank 4. We can thus assume, without loss of generality, that the 4×4 matrix of (8.2) has rank 4.

Since, by assumption, the four lines admit finitely many transversals in $\mathbb{P}^3(\mathbb{R})$, A, B , and C in (8.4) are not all zero. We compute the degree, in the coordinates of the input points, of the various polynomial terms in (8.5). For each input line ℓ_i , the first three and last three coordinates of its Plücker representation have degree 1 and 2, respectively. Hence δ , α_1 , and β_1 have degree 5 and α_i and β_i have degree 6 for $i = 2, 3, 4$. Hence, A, B , and C have degree 11 and the bounds on the degrees of ϕ_i, φ_i , and Δ follow. Note, in particular, that, if $A \neq 0$, these bounds are reached for $i = 4, 5, 6$.

Consider now the case where the four input lines are parallel to a common plane. Since the four lines admit finitely many transversals in $\mathbb{P}^3(\mathbb{R})$, they are not parallel. It follows that the 4×3 matrix of the direction vectors of the four lines has rank 2. Two vectors, say (a_i, b_i, c_i, d_i) for $i = 1, 2$, are thus linearly independent and, by the basis extension theorem, the corresponding 4×2 matrix can be complemented by two other columns (say, (a_i, b_i, c_i, d_i) for $i = 4, 5$) of the matrix of Plücker coordinates (of (8.1)) in order to define a 4×4 matrix of rank 4. As above, a straightforward computation gives the Plücker coordinates of the line transversal. We get

$$x_1 = \alpha_1 u, \quad x_2 = \alpha_2 u, \quad x_3 = -u\delta, \quad x_4 = \alpha_3 u + \beta_3 v, \quad x_5 = \alpha_4 u + \beta_4 v, \quad x_6 = -v\delta$$

where $(u, v) \in \mathbb{P}^1(\mathbb{R})$ is solution of the equation

$$A' u^2 + B' uv = 0 \quad \text{where} \quad A' = \alpha_1 \alpha_3 + \alpha_2 \alpha_4 \quad \text{and} \quad B' = \alpha_1 \beta_3 + \alpha_2 \beta_4 + \delta^2. \quad (8.6)$$

$\delta, \alpha_1, \alpha_2, \beta_3, \beta_4$ have degree 6 and α_3, α_4 have degree 7 (and $\beta_1 = \beta_2 = 0$) thus A' and B' have degree 13 and 12, respectively. Note that A' and B' are not both zero since there are finitely many transversals. The Plücker coordinates of the transversals can thus be written as polynomials of degree at most 19 and, for one of the transversals (the one not in the plane at infinity), this bound is reached for three coordinates (namely, x_4, x_5, x_6). \square

Lemma 8.2. *Consider four lines, given by the Cartesian coordinates of pairs of points, that admit finitely many line transversals in $\mathbb{P}^3(\mathbb{R})$. If the four lines are not parallel to a common plane, we can compute on each transversal two points whose homogeneous coordinates have the form $\phi_i + \varphi_i \sqrt{\Delta}$, $i = 1, \dots, 4$, where ϕ_i, φ_i , and Δ are polynomials of degree at most 17, 6, and 22, respectively, in the coordinates of the input points. Otherwise, we can compute on each transversal two points whose homogeneous coordinates are polynomials of degree at most 19. Moreover, these bounds are reached, in the worst case, for some coordinates.*

Proof. Denote by w_1 (resp. w_2) the vector of the first (resp. last) three coordinates of (x_1, \dots, x_6) , the Plücker coordinates of a line k , and let n denote any vector of \mathbb{R}^3 . Then, if the four-tuple $(w_2 \times n, w_1 \cdot n)$ is not equal to $(0, 0, 0, 0)$, it is a point (in homogeneous coordinates) on the line k (by Lagrange's triple product expansion formula). By considering the axis unit vectors for n , we get that the four-tuples $(0, x_6, -x_5, x_1)$, $(-x_6, 0, x_4, x_2)$, $(x_5, -x_4, 0, x_3)$ that are non-zero are points on the transversal lines k . Either five of the six Plücker coordinates of k are zero or at least two of these four-tuples are non-zero and thus are points on k . In the latter case, the result follows from Lemma 8.1. In the former case, two points with coordinates 0 or 1 can easily be computed on line k since the line is then one of the axis or a line at infinity defined by the directions orthogonal to one of the axis. \square

Remark 8.3. *In some implementations (for instance, the one of [Red03]), the 4×4 submatrix of the matrix of Plücker coordinates (see (8.1)) used for computing the line transversals is chosen, by default, as the leftmost submatrix whose determinant has degree 7 in the coordinates of the input points. In this case, the Plücker coordinates of the line transversals are written as $\phi_i + \varphi_i \sqrt{\Delta}$, $i = 1, \dots, 6$, where ϕ_i, φ_i , and Δ are polynomials of degree at most 20, 7, and 26, respectively, in the coordinates of the input points (and these bounds are reached). Similarly for the homogeneous coordinates of two points on the transversals.*

8.3 Predicates

8.3.1 Preliminaries

We start by two straightforward lemmas on the degree of predicates for determining the sign of simple algebraic numbers. If x is a polynomial expression in some variables, we denote by $\deg(x)$ the degree of x in these variables. This first lemma is trivial and its proof is omitted.

Lemma 8.4. *If a, b , and c are polynomial expressions of (input) rational numbers, the sign of $a + b\sqrt{c}$ can be determined by a predicate of degree $\max\{2\deg(a), 2\deg(b) + \deg(c)\}$.*

Lemma 8.5. *If $\alpha_i, \beta_i, \delta, \mu, i = 1, 2$, are polynomial expressions of (input) rational numbers, the sign of $\alpha_1 + \beta_1\sqrt{\delta} + (\alpha_2 + \beta_2\sqrt{\delta})\sqrt{\mu}$ can be obtained by a predicate of degree*

$$\max\{4\deg(\alpha_1), 4\deg(\beta_1) + 2\deg(\delta), 4\deg(\alpha_2) + 2\deg(\mu), 4\deg(\beta_2) + 2\deg(\delta) + 2\deg(\mu), \\ 2\deg(\alpha_1) + 2\deg(\beta_1) + \deg(\delta), 2\deg(\alpha_2) + 2\deg(\beta_2) + 2\deg(\mu) + \deg(\delta)\}.$$

Proof. The predicate is to evaluate the sign of an expression of the form $a + b\sqrt{\mu}$, where $a = \alpha_1 + \beta_1\sqrt{\delta}$, $b = \alpha_2 + \beta_2\sqrt{\delta}$, and $\alpha_i, \beta_i, \mu, \delta$ are rational. This can be done by evaluating the signs of a , b , and $a^2 - b^2\mu$. The first two signs can be obtained by directly applying Lemma 8.4. On the other hand, $a^2 - b^2\mu$ is equal to $A + B\sqrt{\delta}$ with $A = \alpha_1^2 + \beta_1^2\delta - \alpha_2^2\mu - \beta_2^2\mu\delta$ and $B = 2\alpha_1\beta_1 - 2\alpha_2\beta_2\mu$. The sign of $A + B\sqrt{\delta}$ can be determined by another application of Lemma 8.4, which gives the result. \square

8.3.2 Transversals to four lines

We consider first the predicate of determining whether four lines admit 0, 1, 2, or infinitely many line transversals in $\mathbb{P}^3(\mathbb{R})$ (that is lines in $\mathbb{P}^3(\mathbb{R})$ that intersect, in $\mathbb{P}^3(\mathbb{R})$, the four input lines). An evaluation of this predicate directly follows from the algorithm described in Section 8.2 for computing the line transversals. Recall that, in the sequel, all input points are, by assumption, given by their Cartesian coordinates.

Theorem 8.6. *Given four lines defined by pairs of points, there is a predicate of degree 22 in the coordinates of these points to determine whether the four lines admit 0, 1, 2, or infinitely many line transversals in $\mathbb{P}^3(\mathbb{R})$.*

Proof. We consider three cases. First, if the four lines are parallel, which can easily be determined by a predicate of degree 3, then they admit infinitely many line transversals in $\mathbb{P}^3(\mathbb{R})$. Second, if the four lines are not parallel but parallel to a common plane, which can easily be determined by a predicate of degree 3, then the four lines admit infinitely many transversals if Equation (8.6) is identically zero and, otherwise, 2 line transversals in $\mathbb{P}^3(\mathbb{R})$; this can thus be determined with a predicate of degree 13 (see the proof of Lemma 8.1). Finally, if the four lines are not parallel to a common plane, they admit infinitely many transversals if Equation 8.4 is identically zero and, otherwise, 0, 1, or 2 transversals depending on the sign of Δ (in Lemma 8.1) which is of degree 22 in the coordinates of the points defining the lines. \square

Note that if the leftmost (instead of the rightmost) 4×4 submatrix of the matrix of Plücker coordinates (in (8.1)) is used for computing line transversals (see Remark 8.3) then the procedure described in the above proof has degree 26 instead of 22.

All line transversals are defined in \mathbb{R}^3 except in the case where the four input lines are parallel to a common plane, in which case the intersection of this plane with the plane at infinity is a line transversal at infinity. Note also that, determining whether a line transversal in $\mathbb{P}^3(\mathbb{R})$ is transversal in \mathbb{R}^3 amounts to determining whether the transversal is parallel to one of the four input lines ℓ_i , that is if their direction vectors are collinear. This can be done, by Lemmas 8.1 and 8.4, by a predicate of degree 36 in the Cartesian coordinates of the points defining the input lines.

Note, however, that if the points defining the ℓ_i have rational coordinates and if the transversal is parallel to one of the ℓ_i , the Plücker coordinates of the transversal are rational; indeed, the multiplicative factor of the direction vectors is rational (since one of the coordinates of the direction vector of the transversal is rational, e.g.,

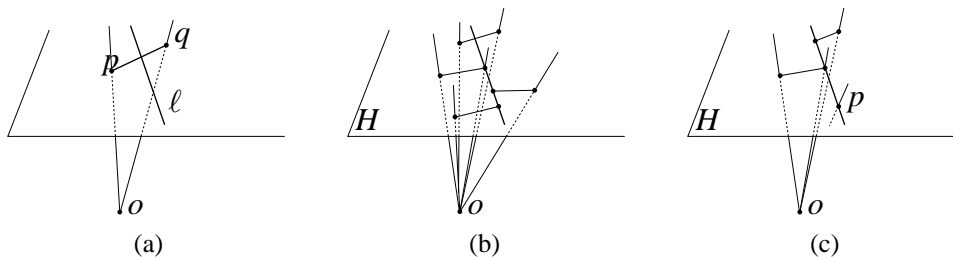


FIG. 8.1 – (a) : Transversal ℓ intersects segment pq only if $(\ell \odot op) (\ell \odot oq) \leq 0$. (b-c) : An illustration for the proof of Lemma 8.10.

x_2 in (8.5)) and thus all the coordinates of this direction vector are rational, which implies that Δ is a square in (8.5). Hence, deciding whether a transversal is parallel to one of the input lines ℓ_i can be done by first determining whether Δ is a square and, if so, testing whether the direction vectors are collinear. It thus follows from Lemma 8.1 that determining whether a transversal is parallel to one of the input lines ℓ_i can be done with a fixed-precision floating-point arithmetic using a number of bits roughly equal to 22 times the number of bits used in representing each input value. This should be compared to the degree 36 of the above procedure. In this paper we have restricted our attention to evaluation procedures for predicates that consist entirely of determining the signs of polynomial expressions in the input parameters. We see here an example of a predicate which may be more efficiently evaluated by a procedure which permits other operations, in this case, determining whether a rational number is a square. This provides an interesting example of a geometric predicate whose algebraic degree does not seem to be an entirely adequate measure of the number of bits needed for the computation.

8.3.3 Transversals to four segments

We consider here the predicate of determining how many transversals four segments of \mathbb{R}^3 admit. Recall that four segments may admit up to 4 or infinitely many line transversals [BEL⁺05]. In this section, we prove the following theorem.

Theorem 8.7. *Given four line segments, there is a predicate of degree 36 in the coordinates of their endpoints to determine whether those segments admit 0, 1, 2, 3, 4, or infinitely many line transversals.*

Note that if, the leftmost (instead of the rightmost) 4×4 submatrix of the matrix of Plücker coordinates (in (8.1)) is used for computing line transversals (see Remark 8.3) then the procedure described below for the predicate of Theorem 8.7 has degree 42 instead of 36.

We consider, in the following, the supporting lines of the four segments, that is, the lines containing the segments; in the case where one (or several) segment is reduced to a point, we consider as supporting line, any line through this point and parallel to at least another supporting line. We first consider the case where the four supporting lines admit finitely many transversals in \mathbb{P}^3 ; this can be determined by a predicate of degree 22, by Theorem 8.6.

Lemma 8.8. *Given four segments in \mathbb{R}^3 whose supporting lines admit finitely many line transversals in \mathbb{P}^3 , determining the number of transversals to the four segments can be done with a predicate of degree 36 in the coordinates of their endpoints.*

Proof. Let ℓ denote an (arbitrarily) oriented line, as well as its Plücker coordinates, that is transversals to the four lines; ℓ can be computed as described in Section 8.2. We consider the predicate of determining whether ℓ intersects each of the four segments, in turn. Let p and q denote the endpoints of one of these segments. For any two distinct points r and s , denote by rs the Plücker coordinates of the line rs oriented from r to s ; depending on the context, rs also denotes the line through r and s or the segment from r to s .

If a point o does not lie in the plane containing line ℓ and segment pq (see Figure 8.1(a)), then line ℓ intersects segment pq if and only if the oriented line ℓ is on opposite sides of the two oriented lines from o to p and from o to q , that is if $(\ell \odot op) (\ell \odot oq) \leq 0$ (recall that \odot denotes the side operator – see Section 8.2).

On the other hand, point o lies in a plane containing line ℓ and segment pq if and only if ℓ intersects (in \mathbb{P}^3) both lines op and oq , that is both side operators $\ell \odot op$ and $\ell \odot oq$ are zero. By choosing point o to be for instance $(1, 0, 0)$, $(0, 1, 0)$, $(0, 0, 1)$, or $(1, 1, 1)$, we ensure that one of these points will not be coplanar with ℓ and segment pq unless segment pq lies on ℓ .

Hence the predicate follows from the sign of side operators of the line transversal and of a line defined by two points, one of which with coordinates equal to 0 or 1. The degree of the Plücker coordinates of the line through these two points is thus 1 (in the coordinates of the input points). Hence, by Lemma 8.1, the predicate can be computed by determining the sign of polynomials of degree at most 20 if the input lines are parallel to a common plane and, otherwise, by determining the sign of expressions of the form $a + b\sqrt{c}$ where a , b and c have degree at most 18, 7, and 22, respectively; moreover, these bounds are reached. By Lemma 8.4, the predicate thus has degree 36, which concludes the proof. \square

We now consider the case where the four lines admit infinitely many transversals. Recall that, in \mathbb{P}^3 , four lines or line segments admit infinitely many transversals only if [BEL⁺05]:

1. they lie in one ruling of a hyperbolic paraboloid or a hyperboloid of one sheet,
2. they are all concurrent, or
3. they all lie in a plane, with the possible exception of a group of one or more that all meet that plane at the same point.

We treat the cases independently.

Lemma 8.9. *Given four segments in \mathbb{R}^3 whose supporting lines are pairwise skew and admit infinitely many line transversals, determining the number of their line transversals can be done with a predicate of degree at most 36 in the coordinates of their endpoints.*

Proof. When four lines are pairwise skew, their common transversals can be parameterized by their points of intersection with one of the lines; moreover, the set of common transversals to the four segments corresponds (through this parameterization) to up to four intervals on that line and the transversals that correspond to the endpoints of these intervals contain (at least) one endpoint of the segments [BEL⁺05]. We can compute and order all these interval endpoints and determine whether there exists a transversal (to the four segments) through each midpoint of two consecutive distinct interval endpoints. By construction and by [BEL⁺05], the four segments admit such a transversal if and only if they admit infinitely many transversals.

The set of interval endpoints, on, say, segment s_1 is a subset of the endpoints of s_1 and of the intersection points of s_1 with the planes containing s_2 and an endpoint of s_3 or s_4 and of the intersection points of s_1 with the planes containing s_3 and an endpoint of s_2 . The coordinates of these points can be trivially computed as rational expressions of degree 4 in the coordinates of the segment endpoints. The coordinates of the midpoints are thus rational expressions of degree at most 8.

The transversal to the four lines through (any) one of these midpoints intersects line ℓ_2 and lies in the plane containing line ℓ_3 and the considered midpoint; the coordinates of the intersection point between this plane and ℓ_2 are rational expressions of degree at most 19. Finally, determining whether a transversal (to the four lines) through two points whose coordinates are rational expressions of degree 8 and 19 is a transversal to each of the four segments can be done, as in the proof of Lemma 8.8, using side operators. Hence, we can decide whether the four segments admit infinitely many transversals with a predicate of degree at most 36 since the Plücker coordinates of the line transversal are of degree at most 35.

Now, if the four segments admit finitely many transversals, we can determine the number of transversals as follows. As mentioned above, the set of transversals can be parameterized by intervals on a line and the interval endpoints correspond to transversals that go through a segment endpoint. A transversal is isolated if and only if it corresponds to an interval that is reduced to a point. Thus, a transversal is isolated only if it goes through two distinct segment endpoints (the segments necessarily have distinct endpoints since, by assumption, their supporting lines are pairwise skew and thus no segment is reduced to a point). Determining whether the lines through two distinct endpoints intersect the other segments can easily be done, as described in the proof of Lemma 8.8, by

computing the sign of side operators which are here of degree 3 in the coordinates of the segment endpoints. \square

Lemma 8.10. *Given four segments in \mathbb{R}^3 whose supporting lines are not pairwise skew and admit infinitely many line transversals, determining the number of their line transversals can be done with a predicate of degree 7 in the coordinates of their endpoints.*

Proof. First, note that testing whether two segments intersect can be done using side operators with a predicate of degree 3. The four lines containing the segments are not pairwise skew and they admit infinitely many line transversals. Thus, they are all concurrent or they all lie in a plane H , with the possible exception of a group of one or more that all meet that plane at the same point [BEL⁺05]. Four cases may occur :

- (i) all four lines lie in a plane H ,
- (ii) three lines lie in a plane H and the fourth line intersects H in exactly one point,
- (iii) two lines lie in a plane H and two other lines intersect H in exactly one and the same point,
- (iv) three lines are concurrent but not coplanar.

Differentiating between these cases can be done by determining whether sets of four segment endpoints are coplanar (which is a predicate of degree 3). We study each case in turn.

Case (i). The four segments are coplanar. Any component of transversals contains a line through two distinct segment endpoints. Hence the four segments have finitely many transversals if and only if any line through two distinct endpoints that is a transversal to the four segments is an isolated transversal. This only occurs¹² (see Figure 8.1(b)) when the transversal goes through the endpoints of three segments such that the segment, whose endpoint is in between the two others, lies (in H) on the opposite side of the transversal than the two other segments. This can be tested by computing the sign of scalar products and side operators between the transversal and the lines through a point o not in H and the segment endpoints (see Figure 8.1(b)). This leads to a predicate of degree 4.

Case (ii). Three lines lie in a plane H . Testing whether the fourth segment intersects the plane H can easily be done by computing the point of intersection between H and the line containing the fourth segment, leading to a predicate of degree 3. If the fourth segment does not intersect plane H , the four segments have no transversal unless the first three segments are concurrent in which case the four segments have one or infinitely many transversals depending on whether the four lines supporting the segments are concurrent. Otherwise, let p denote the point of intersection. We assume that the three segments in H are not concurrent; otherwise the four segments have infinitely many transversals. Thus, any component of transversals contains a line through p and through a segment endpoint. Hence the four segments have finitely many transversals if and only if any line through p and a segment endpoint that is a transversal to the four segments is an isolated transversal. Testing whether such a line is a transversal to all segments can be done, as in the proof of Lemma 8.8, by computing the sign of side operators of the line transversal and of lines through a segment endpoint and a point o not in H ; the coordinates of point p are rational expressions of degree 4, thus the Plücker coordinates of the transversal have degree at most 6, which leads to a predicate of degree 7. Such a line transversal is isolated (see Figure 8.1(c)) if and only if¹³ the transversal goes through two endpoints of two distinct segments that lie on the same side (in plane H) of the transversal or not depending whether p is in between the two endpoints or not. This test can be done by computing the sign of scalar products and side operators between the transversal and the lines through a point o not in H and the segment endpoints (see Figure 8.1(c)). This test also leads to a predicate of degree 7. We can thus determine the number of isolated transversals with a predicate of degree 7.

Case (iii). Two lines lie in a plane H and two other lines intersect H in exactly one and the same point. (Note that there may be two instances of plane H for a given configuration.) This case can be treated similarly as Case (ii).

¹²For simplicity, we do not discuss here the case where the line transversal contains one of the four segments.

¹³We assume here for simplicity that the line transversal contains no segment.

Case (iv). Three lines are concurrent but not coplanar. If none of the three corresponding segments intersect, they have no common transversal. If only two segments intersect, the three segments have exactly one transversal; checking whether that transversal intersects the fourth segment can easily be done with a predicate of degree 3. Now, if the three segments intersect, then the four segments have infinitely many transversals if they are concurrent or if their supporting lines are not concurrent. Otherwise, if the four segments are not concurrent but their supporting lines are, the four segments then have a unique transversal. This can also be checked with a predicate of degree 3. \square

We can now conclude the proof of Theorem 8.7. By Theorem 8.6, we can determine with a predicate of degree 22 whether the four lines containing the four segments admit finitely many transversals in \mathbb{P}^3 . If the four lines admit finitely many transversals, then, by Lemma 8.8, determining the number of transversals to the four segments can be done with a predicate of degree 36. Assume now that the four lines admit infinitely many transversals. Note that determining whether the input lines are pairwise skew can easily be done with a predicate of degree 3. Thus, by Lemmas 8.9 and 8.10, determining whether the four segments admit 0, 1, 2, 3, 4, or infinitely many line transversals can be done by a predicate of degree at most 36. Hence, we can determine the number of transversals to four segments with a predicate of degree 36. \square

8.3.4 Transversals to four segments and a triangle

We consider here the predicate of determining whether a minimal segment transversal to four line segments is intersected by a triangle. Given a line transversal ℓ to a set S of segments, a triangle T *occludes* ℓ if ℓ intersects T and if there exist two segments in S whose intersections with ℓ lie on opposite sides of T . We describe a method for evaluating the predicate for determining whether a triangle occludes a transversal to a given set of line segments and establish its degree.

Theorem 8.11. *Let ℓ be a line transversal to four line segments that admit finitely many transversals and let T be a triangle. There is a predicate of degree 78 in the coordinates of the points defining the segments and the triangle to determine whether T occludes ℓ .*

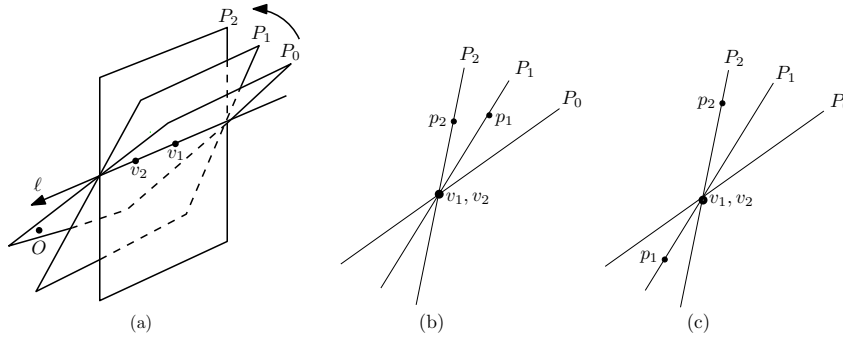
Proof. Let ℓ denote an oriented line transversal to segments s_1, \dots, s_4 , each defined by two points e_i and f_i , $i = 1, \dots, 4$, and let T be a triangle defined by three points p, q , and r . The Plücker coordinates of ℓ can be computed as described in Section 8.2. We only consider the case where the four lines containing segments s_i have finitely many transversals because, otherwise, since the four segments admit finitely many transversals, each transversal goes through at least one endpoint of the four segments and it is straightforward that the degree of the predicate is then much smaller.

We first determine whether ℓ intersects T by taking the side product of ℓ with each supporting line of T (oriented consistently); ℓ intersects T if and only if no two side products have opposite signs (*i.e.*, ± 1). Similarly as in the proof of Lemma 8.8, there is a predicate of degree 38 for determining the sign of these side operators.

Assuming that ℓ intersects T , we next find the point of intersection. By Lemma 8.2, ℓ can be represented parametrically in the form $\pi + \rho t$. We determine the value of t for which the determinant of $p, q, r, \pi + \rho t$ is equal to zero; denote this value of t by t_0 . This determinant has the form $a_0 + b_0 t_0$, where, by Lemma 8.2, a_0 and b_0 are polynomials of degree 22 if s_1, \dots, s_4 are parallel to a common plane or, otherwise, have the form $\phi + \phi\sqrt{\Delta}$ where ϕ, ϕ , and Δ have degree 20, 8, and 22, respectively, in the coordinates of p, q, r, e_i, f_i .

Now, for each segment s_i , we compute the point of intersection of s_i with ℓ in terms of the parameter t using the method similar to that of the previous section: choose a point o_i not in the plane determined by s_i and ℓ and compute the value t for which the determinant of $e_i, f_i, o_i, \pi + \rho t$ equals 0. Denote this value by t_i . Since o_i can be chosen with all coordinates equal to 0 or 1, we get, similarly as in the previous paragraph, that each of these determinants has the form $a_i + b_i t_i$ where a_i and b_i are polynomials of degree 21 if s_1, \dots, s_4 are parallel to a common plane or, otherwise, have the form $\phi + \phi\sqrt{\Delta}$ where ϕ, ϕ , and Δ have degree 19, 8, and 22, respectively.

Determining whether T occludes ℓ is now only a matter of determining whether t_0 lies between two of the values $t_i, i = 1, \dots, 4$, which requires only that we be able to compare t -values, that is, compute $\text{sign}(t_i - t_j)$. Observe that $t_i - t_j = \frac{a_j b_i - a_i b_j}{b_i b_j} < 0$, so $\text{sign}(t_i - t_j) = \text{sign}(a_j b_i - a_i b_j) \text{sign}(b_i) \text{sign}(b_j)$. It follows from the above characterization of the a_i and b_i that a product $a_i b_j$ is either a polynomial of degree 43 if s_1, \dots, s_4 are parallel to a common plane or, otherwise, has the form $\phi + \phi\sqrt{\Delta}$ where ϕ, ϕ , and Δ have degree at most 39, 28, and 22,

FIG. 8.2 – Planes P_1 and P_2 such that $P_1 < P_2$

respectively (and these bounds are reached in the worst case). Applying Lemma 8.4 yields a predicate of degree 78, which concludes the proof. \square

Note that, if the leftmost (instead of the rightmost) 4×4 submatrix of the matrix of Plücker coordinates (in (8.1)) is used for computing line transversals (see Remark 8.3) then the procedure described above for the predicate of Theorem 8.11 has degree 90 instead of 78.

8.3.5 Ordering planes through two fixed points, each containing a third (rational) point or a line transversal

Let ℓ be a line defined by two points v_1 and v_2 , and $\vec{\ell}$ be the line ℓ oriented in the direction $\overrightarrow{v_1 v_2}$.

We define an ordering of all the planes containing ℓ with respect to the oriented line $\vec{\ell}$ and a reference point O (not on ℓ). Let P_0 be the plane containing O and ℓ , and let P_1 and P_2 be two planes containing ℓ . We say that $P_1 < P_2$ if and only if P_1 is encountered strictly before P_2 when rotating counterclockwise about $\vec{\ell}$ a plane from P_0 (see Figure 8.2a).

Let p_i be any point on plane P_i but not on ℓ , for $i = 1, 2$, and let $D(p, q)$ denote the determinant of the four points (v_1, v_2, p, q) given in homogeneous coordinates.

Lemma 8.12. *With $\chi = D(O, p_1) \cdot D(O, p_2) \cdot D(p_1, p_2)$, we have :*

- (a) *If $\chi > 0$ then $P_1 > P_2$.*
- (b) *If $\chi < 0$ then $P_1 < P_2$.*
- (c) *If $\chi = 0$ then*
 - (i) *if $D(p_1, p_2) = 0$, then $P_1 = P_2$,*
 - (ii) *else if $D(O, p_1) = 0$, then $P_1 < P_2$,*
 - (iii) *else $P_1 > P_2$.*

Proof. Assume first that $D(O, p_1) \cdot D(O, p_2) > 0$, that is, that p_1 and p_2 lie strictly on the same side of the plane P_0 (see Figure 8.2b). Then the order of P_1 and P_2 is determined by the orientation of the four points (v_1, v_2, p_1, p_2) , that is by the sign of $D(p_1, p_2)$. It is then straightforward to notice that $P_1 > P_2$ if and only if $D(p_1, p_2) > 0$. Hence, if $\chi > 0$, then $P_1 > P_2$ and, if $\chi < 0$, then $P_1 < P_2$.

Suppose now that $D(O, p_1) \cdot D(O, p_2) < 0$, that is, that p_1 and p_2 lie strictly on opposite sides of the plane P_0 (see Figure 8.2c). The order of P_1 and P_2 is then still determined by the sign of $D(p_1, p_2)$. However, $P_1 > P_2$ if and only if $D(p_1, p_2) < 0$. Hence, we have in all cases that, if $\chi > 0$, then $P_1 > P_2$ and, if $\chi < 0$, then $P_1 < P_2$.

Suppose finally that $\chi = 0$. If $D(p_1, p_2) = 0$, then p_1 and p_2 are coplanar, and $P_1 = P_2$. Otherwise, if $D(O, p_1) = 0$, then $P_0 = P_1$ thus P_1 is smaller to all other planes (containing $\vec{\ell}$), and in particular $P_1 \leq P_2$. Furthermore, since $D(p_1, p_2) \neq 0$, $P_1 \neq P_2$ and thus $P_1 < P_2$. Otherwise, $D(O, p_2) = 0$ and we get similarly that $P_2 < P_1$. \square

Computing a point on a plane defined by ℓ and a line transversal. We want to order planes P_i that are defined by line ℓ and either a rational point not on ℓ , or by a line transversal to ℓ and three other lines. In the latter case, we consider a point on the line transversal (which is non-rational, in general; see Lemma 8.2). The following lemma tells us that, in general, such a plane P_i contains no rational points outside of ℓ , and that in the cases where it does contain such a rational point, the line transversal is then rational. Hence, if the points computed on the line transversal, as described in Lemma 8.2, are not rational, there is no need to search for simpler points on the plane (but not on ℓ).

Lemma 8.13. *The plane P containing a rational line ℓ and a line transversal to ℓ and three other segments, each determined by two rational points, contains in general no rational points except on ℓ . Furthermore, if plane P contains a rational point not on ℓ then the line transversal is rational.*

Proof. Suppose that the plane P contains a rational point p not on ℓ . Then the plane contains three (non-collinear) rational points, p and two points on ℓ , and thus P is a rational plane. This plane intersects the three other segments in three points, all of which are rational and lie on the transversal. So the transversal is a rational line which implies that the discriminant $B^2 - 4AC$ in Equation (8.5) is a square, which is not the case in general. \square

Comparing two planes. We want to order planes P_i that are defined by either line ℓ and another (input rational) point not on ℓ , or by line ℓ and a line transversal to ℓ and three other lines.

By Lemma 8.12, ordering such planes about ℓ amounts to computing the sign of determinants of four points (in homogeneous coordinates). Two of these points are input (affine rational) points on ℓ (v_1 and v_2) and each of the two other points is either an input (affine rational) point r_i , $i = 1, 2$, or is, by Lemma 8.2 (and Lemma 8.13), a point u_i whose homogeneous coordinates have degree at most 19 (in the coordinates of the input points) or a point of the form $p_i + q_i \sqrt{\Delta_i}$, $i = 1, 2$, where the Δ_i have degree 22 and where the p_i and q_i are points with homogeneous coordinates of degree at most 17 and 6, respectively. If the four points are all input points, then the determinant of the four points has degree 3 in their coordinates.

If only three of the four points are input points, then the determinant of the four points is either a polynomial of degree 22 or it has the form $D(p_1, r_1) + D(q_1, r_1) \sqrt{\Delta_1}$ where the degrees of the $D()$ are 20 and 9, respectively, in the coordinates of the input points. Hence, by Lemma 8.4, the sign of this expression can be determined with a predicate of degree 40.

Finally, if only two of the four points are input points, then the determinant has one of the following forms (depending on whether the quadruples of lines defining the transversals are parallel to a common plane); the degrees are given in terms of the coordinates of the input points :

- (i) $D(u_1, u_2)$ which is of degree 40.
- (ii) $D(u_1, p_1) + D(u_1, q_2) \sqrt{\Delta_1}$ where the $D()$ have degree 38 and 27, respectively.
- (iii) $D(p_1, p_2) + D(q_1, p_2) \sqrt{\Delta_1} + (D(p_1, q_2) + D(q_1, q_2) \sqrt{\Delta_1}) \sqrt{\Delta_2}$ where the $D()$ have degree 36, 25, 25, and 14, respectively.

Hence, by Lemma 8.5, the sign of these expressions can be determined with a predicate of degree at most 144 (and the bound is reached in the worst case). We thus get the following result.

Theorem 8.14. *Let ℓ be an oriented line defined by two points, let p_0 be a point not on ℓ , and let P_0 be the plane determined by ℓ and p_0 . Given two planes P_1, P_2 containing ℓ there is a predicate which determines the relative order of P_1 and P_2 about ℓ with respect to P_0 having the following degree in the coordinates of the input points :*

- (i) degree 3 if $P_i, i = 1, 2$ are each specified by a (input) point p_i ;
- (ii) degree 40 if P_1 is specified by a point p_1 and P_2 is determined by a line transversal to ℓ and three other lines ℓ_1, ℓ_2, ℓ_3 , each specified by two (input) points;
- (iii) degree 144 if $P_i, i = 1, 2$ are each determined by a line transversal to ℓ and three other lines $\ell_{i,1}, \ell_{i,2}, \ell_{i,3}$, each specified by two (input) points.

Remark 8.15. *Similarly as before, note that, if the leftmost (instead of the rightmost) 4×4 submatrix of the matrix of Plücker coordinates (in (8.1)) is used for computing line transversals (see Remark 8.3) then the predicates of Theorem 8.14 have degree 3, 46, and 168.*

predicates \ ϵ	10^{-12}	10^{-10}	10^{-8}	10^{-6}	10^{-4}	10^{-2}
degree 168	99.6%	50.4%	7.6%	0.8%	0.08%	0.008%
degree 3	99.5%	8.2%	0.08%	0.001%		

TAB. 8.1 – Percentages of failure of the degree 168 and degree 3 predicates using double-precision floating-point interval-arithmetic, for ϵ varying from 10^{-12} to 10^{-2} .

8.4 Experiments

In this section, we report the results of experiments that analyze the behavior of the predicate for ordering, in a rotational sweep about a line, two planes each defined by a line transversal to four lines, that is the predicate related to Theorem 8.14(iii). The degree of the procedure we use for evaluating this predicate is 168 because we use for computing line transversals to four lines the code of Redburn [Red03], which, as noted in Remarks 8.3 and 8.15, leads to degree 168 instead of 144 as in Theorem 8.14(iii).

The standard approach to comparing two such planes is to first evaluate the predicate using fixed-precision interval-arithmetic. This is very efficient but may fail when the sign of an expression cannot be successfully determined because the result of the evaluation of the expression is an interval that contains zero. If this happens, the answer to the predicate is then obtained by either evaluating exactly the expression (and thus its sign) using exact arithmetic or by increasing the precision of the interval arithmetic until either the result of the evaluation of the expression is an interval that does not contain zero or the separation bound is attained (see for instance [BFM⁺01, Mig82, Sch00, Yap97]); in both approaches the computation is much slower than when using fixed-precision interval-arithmetic. We are thus interested in determining how often the fixed-precision interval-arithmetic evaluation of our predicate fails.

To test our predicate, we generate pairs of planes, each defined by two lines, one chosen at random and common to the two planes, and the other defined as a transversal to the common line and to three other random lines. We are interested in evaluating our predicate in the case where the two planes are very close together, that is, when there is significant risk of producing an error when using finite-precision floating-point arithmetic.

We generate two sets of four lines. Each line of the first set is determined by two points, all of whose coordinates are double-precision floating-point numbers chosen uniformly at random from the interval $[-5000, 5000]$. The second set of lines is obtained by perturbing the points defining three of the lines of the first set; the fourth line is not perturbed and is thus common to the two sets. To perturb a point p , we translate it to a point chosen uniformly at random in a sphere centered at p , with radius ϵ .

We compute, for each of these two sets of four lines, a line transversal. If either set of four lines does not admit a transversal (which happens roughly 24% of the time), we throw out that data and start again. Otherwise, we choose a transversal in a consistent way for the two sets of four lines, that is, such that one transversal converges to the other when ϵ tends to zero. Each transversal, together with the common line, defines a plane.

For various values of ϵ , varying from 10^{-2} to 10^{-10} , we evaluate the predicate using double-precision floating-point interval arithmetic until we obtain 1000 pairs of planes for which the computation of the predicate fails. We measure the percentage of time that the computation fails. The results of these experiments are shown in Table 8.1.

We observe, as expected, that when ϵ is sufficiently small (10^{-10}), that is, when the two planes are often close enough to each other, the fixed-precision interval-arithmetic predicate fails with high probability and that this probability decreases as ϵ increases. When $\epsilon = 10^{-2}$, the probability of failure is close to zero. Finally, we have also observed that the predicate fails when the angle between the two planes is less than roughly 10^{-8} radians, which is, of course, independent of ϵ .

Note finally that the percentage of failure of the degree 168 predicate using fixed-precision interval-arithmetic is, as expected, high compared to lower-degree predicates. Table 8.1 also shows the failure rate for the degree 3 predicate related to Theorem 8.14(i). We use the same experimental scheme as above, that is, we chose at random three points that define a plane and perturb one of these points by at most ϵ .

All the experiments were made on a i686 machine with AMD Athlon 1.73 GHz CPU and 1 GB of main memory using the CGAL interval number type with double-precision floating-point numbers [CGA].

Acknowledgment

The authors gratefully acknowledge fruitful discussions on this topic with Olivier Devillers and Marc Glisse. We also especially thank Hervé Brönnimann for his comments on the paper and Günter Rote for pointing out an improvement, in Lemma 1, to the degrees appearing in the Plücker coordinates of line transversals.

Chapitre 9

Lines and free line segments tangent to arbitrary three-dimensional convex polyhedra

Cet article a été accepté pour publication dans *SIAM Journal on Computing* [BDD⁺07]. Des versions préliminaires sont parues dans les proceedings du *20th ACM Annual Symposium on Computational Geometry* [BDD⁺04], ceux de la *14th Canadian Conference on Computational Geometry* [BDD⁺02], et dans la thèse de X. Goaoc [Goa04].

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.

9.1 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 [Dur00] for a survey).

One such structure, the visibility complex [DDP02, PV96b], 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

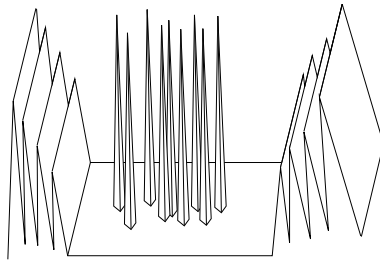


FIG. 9.1 – A terrain of size n with $\Omega(n^4)$ maximal free line segments tangent in four points.

the worst case [DDP02], even when the triangles form a terrain (see [CS89b] or Figure 9.1). 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.* [DDP97] observe a quadratic growth rate, albeit for rather small scenes. For random scenes, Devillers *et al.* [DDE⁺03] 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 results strongly suggest that the size of the visibility complex of a scene consisting of scattered triangles is linear [CF99].

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 [EGHHZ00] 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 [BDD⁺02]. 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 [Sto91]. 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.¹⁴ A line or line segment is *free* if it is tangent to each polytope that its relative interior intersects;¹⁵ 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 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.

¹⁴The set of polytopes to which the line is tangent might change during the motion.

¹⁵When 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.

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 [BDD⁺02, EGHZ00] 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 [BDD⁺02]) 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 9.1. *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 9.1 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 9.1 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 [Aga94]. Moreover, we prove a tighter bound when one of the four polytopes has few edges.

Theorem 9.2. *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 9.1. Whereas Theorem 9.1 bounds the number of connected components of tangents, Theorem 9.3 bounds the number of isolated tangents with some notion of multiplicity. For example, the line in Figure 9.2 is counted $\binom{k}{2}$ times which is the number of minimal sets of vertices that admit that line as an isolated transversal. Although neither theorem implies the other, we will prove in Proposition 9.23 that the upper bound of Theorem 9.1 is easily proved using Theorem 9.3.

Theorem 9.3. *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.* [DDP02] 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 [DDP97]. We present an algorithm that uses, in the worst case, $O(n^2k^2 \log n)$ time and $O(nk^2)$ space,

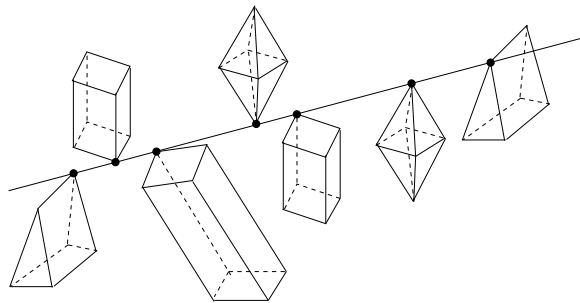


FIG. 9.2 – A line tangent at a vertex of each of k polytopes.

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 [Goa04].

Theorem 9.4. *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 [AS96]. 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 [DK83].

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.

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 [BEL⁺05]. 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 9.1.

Recently, Agarwal *et al.* [AAS04] 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 $\Omega(n^2)$ [DDE⁺03] 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 [DDE⁺03]). 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 [HS94] and Pellegrini [Pel94] proved that, in a polyhedral terrain with n edges, the set of free lines has near-cubic complexity. De Berg, Everett and Guibas [dBEG98] 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.

	Worst-case	Expected
free lines to a polyhedron	$\Theta(n^4)$ (trivial)	
free lines above a polyhedral terrain	$O(n^3 2^{c\sqrt{\log n}})$ [HS94, Pel94]	
free lines among disjoint homothetic polytopes	$\Omega(n^3)$ [dBEG98]	
free lines among unit balls	$\Omega(n^2)$ [DDE ⁺ 03], $O(n^{3+\varepsilon})$ [AAS04]	$\Theta(n)$ [DDE ⁺ 03]
max. free segments above a polyhedral terrain	$\Theta(n^4)$ [CS89b]	
isolated maximal free segments among k generic disjoint convex polyhedra	$\Theta(n^2 k^2)$ [EGHHZ00, BDD ⁺ 02]	
max. free segments among unit balls	$\Omega(n^2)$ [DDE ⁺ 03], $O(n^4)$	$\Theta(n)$ [DDE ⁺ 03]

TAB. 9.1 – 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.

This paper is organized as follows. We prove the upper bounds of Theorems 9.1, 9.2, and 9.3 in Sections 9.2 and 9.3, and the lower bounds in Section 9.4. In section 9.5, we present our algorithm for computing free segments.

9.2 Main lemma

We prove in this section a lemma which is fundamental for the proofs of the upper bounds of Theorems 9.1, 9.2, and 9.3. 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 of the proof. We then state preliminaries and definitions in Section 9.2.2. In Sections 9.2.3 and 9.2.4, we bound the number of so-called “generic tangent lines”. In Section 9.2.5, we bound the number of “non-generic tangent lines”. Finally, in Section 9.2.6, we pull these results together to conclude the proof of the Main Lemma.

9.2.1 Proof overview

The proof is inspired by a method which was, to our knowledge, first used in [BDEG94] (and later in [dBHOvK97, EGHHZ00, BDD⁺02]). 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 9.3). 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 9.3). 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 9.4). 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

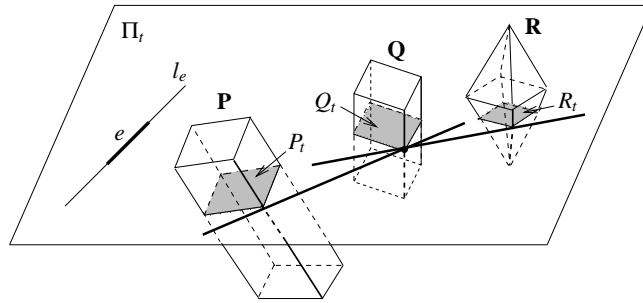


FIG. 9.3 – Plane Π_t contains edge e and intersects polytopes \mathbf{P} , \mathbf{Q} , and \mathbf{R} in polygons P_t , Q_t , and R_t .

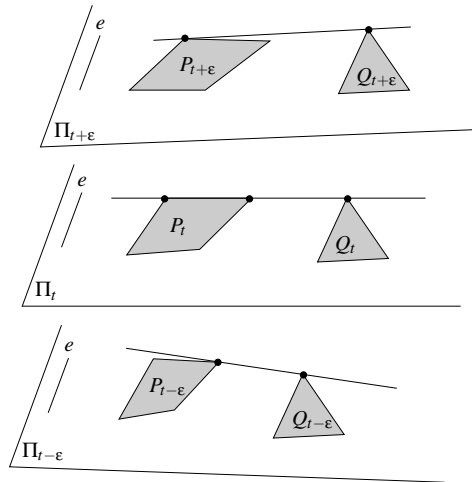


FIG. 9.4 – A bitangent to P_t and Q_t is tangent to P_t along an edge. The plane Π_t is F-critical.

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 9.4 and 9.5). 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 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 tangent 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).

9.2.2 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 9.3).

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 **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 **QR**-tuples.

We say that a (P_t, Q_t) -tuple is *maximal for some t* if it is not contained in any other (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 9.7(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.

9.2.3 Generic tangent lines

Lemma 9.5. *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 9.5) it contains a line ℓ such that

- (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

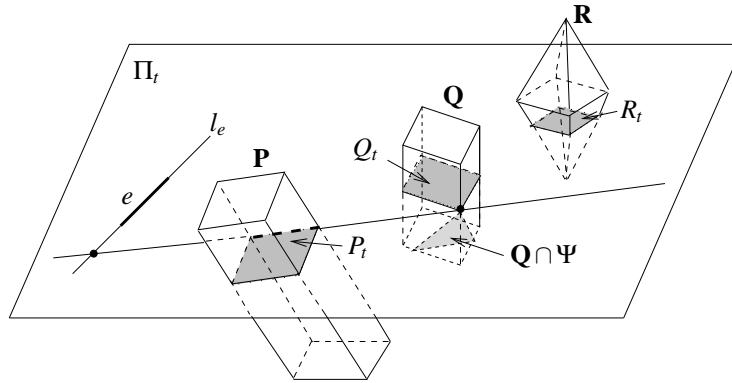


FIG. 9.5 – Plane Π_t is *F-critical* : it contains a line that lies in a plane Ψ containing a face of \mathbf{P} such that the line is tangent to $\mathbf{Q} \cap \Psi$ at a point not on l_e .

edge that has an endpoint on l_e (see Figure 9.6). 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.

Lemma 9.6. *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 9.7. *Let t^* be the endpoint of a maximal interval¹⁶ 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 9.2.4. 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 9.7 ; we don't suggest that the lemma is false without them.

¹⁶Such an interval could be open or closed, a single point or an interval of positive length.

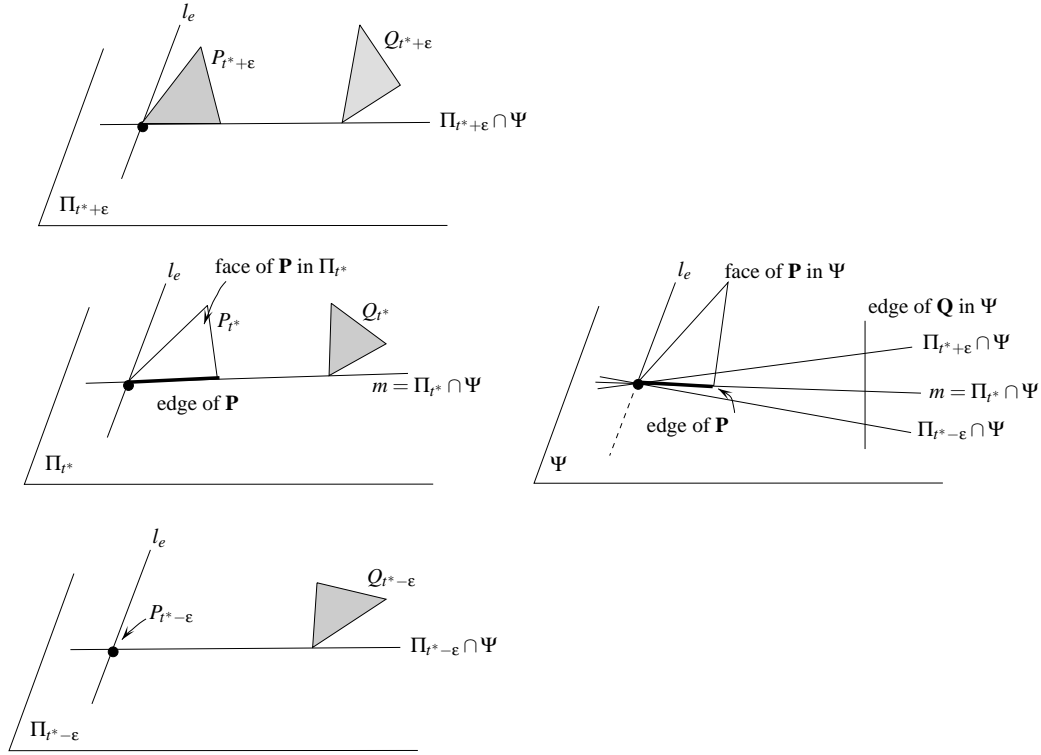


FIG. 9.6 – Plane Π_{t^*} contains a line m such that (i) m lies in a plane $\Psi \neq \Pi_{t^*}$ containing a face of \mathbf{P} , and (ii) m is tangent to polygon $\mathbf{P} \cap \Psi$ at some point not on l_e ; however m is not tangent to $\mathbf{Q} \cap \Psi$. If the definition of F-critical planes was not considering such plane Π_{t^*} to be F-critical then Lemma 9.7 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 .

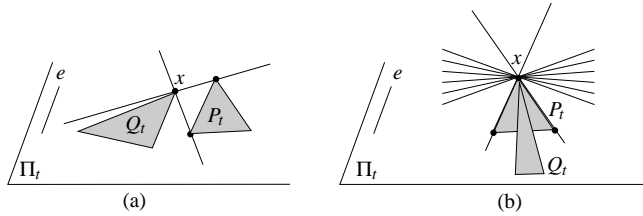
Lemma 9.8. Any edge or vertex of \mathbf{P} or \mathbf{Q} is in at most 2 \mathbf{PQ} -tuples that are maximal (P_t, Q_t) -tuples for all t in any given non-empty interval¹⁶ of $\mathbb{R}/\pi\mathbb{Z}$.

Proof. Let \tilde{t} be an element of a non-empty interval \mathcal{J} 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 9.7(a)); in all cases there are at most 3 $(P_{\tilde{t}}, Q_{\tilde{t}})$ -tuples containing x (see Figure 9.7(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.

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

Lemma 9.9. There are at most $O(p + q + r)$ \mathbf{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 9.6 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 \mathbf{P} , at most three supports in \mathbf{Q} , and at most two supports in \mathbf{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 \mathbf{P} , \mathbf{Q} and \mathbf{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.

FIG. 9.7 – Lines through x in Π_t and tangent to P_t and Q_t .

Let s be a maximal (P_t, Q_t, R_t) -tuple and let \mathcal{J} 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 $t \in \mathcal{J}$. 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 \mathbf{Q} .

First, suppose that $\mathcal{J} = \mathbb{R}/\pi\mathbb{Z}$. 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 \mathbf{P} and \mathbf{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 $\mathcal{J} \neq \mathbb{R}/\pi\mathbb{Z}$, and only count the maximal (P_t, Q_t, R_t) -tuples that are not the set of supports of l_e .

Interval \mathcal{J} 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 $\mathcal{J} \neq \mathbb{R}/\pi\mathbb{Z}$.

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 9.7.

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 $\mathcal{J}_1, \mathcal{J}_2, \mathcal{J}_3$.

By Lemma 9.8, at most 2 \mathbf{PQ} -tuples contain x and are maximal (P_t, Q_t) -tuples for all t in \mathcal{J}_i . Moreover, each of these \mathbf{PQ} -tuples contains at most 2 supports in \mathbf{Q} , and each of these supports belongs to at most 2 \mathbf{QR} -tuples that are maximal (Q_t, R_t) -tuples for all t in \mathcal{J}_i . Thus at most 8 \mathbf{PQR} -tuples contain x and are maximal (P_t, Q_t, R_t) -tuples for all t in \mathcal{J}_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 \mathbf{P} and at most two supports in \mathbf{Q} . Since u contains at least one support in \mathbf{P} and at least one support in \mathbf{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 $\mathcal{J}_1, \mathcal{J}_2, \mathcal{J}_3$ these intervals. It follows from Lemma 9.8 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 \mathcal{J}_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 9.10. *There are at most $O(p+q)$ \mathbf{PQ} -tuples.*

Proof. Replace \mathbf{R} by a copy of \mathbf{Q} in Lemma 9.9. Any \mathbf{PQ} -tuple is also a \mathbf{PQQ} -tuple, and there are at most

$O(p+q+r) = O(p+q)$ of these. □

Proposition 9.11. *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 **PQR**-tuple, by definition and Lemma 9.5. An isolated generic tangent line is thus an isolated transversal with respect to a set of edges and vertices that consists of a **PQR**-tuple and either edge e or one or both of its endpoints. The number of such sets is four times the number of **PQR**-tuples, which is in $O(p+q+r)$ by Lemma 9.9. 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 [BEL⁺05]. □

9.2.4 Proof of Lemma 9.7

Recall that u_e denotes the set of supports of l_e in **P** and **Q**, and that Lemma 9.7 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 9.7. Indeed, we prove the existence of a line m in Π_{t^*} whose set of supports contains u (Lemma 9.14) such that (i) m lies in a plane $\Psi \neq \Pi_{t^*}$ containing a face of **P** (Lemma 9.15), 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 9.16). 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 **P** and in **Q**; polytope **R** plays no role. We start by proving two preliminary lemmas.

Lemma 9.12. *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, **P**, then any line tangent to P_t in Π_t and intersecting x contains an endpoint of x which is a vertex of **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 **P** and of **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 9.7.

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,

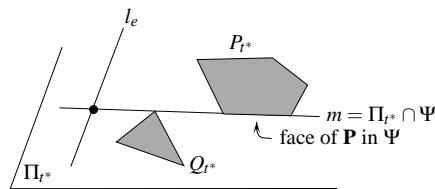


FIG. 9.8 – Line m is tangent to \mathbf{P} along a face in plane $\Psi \neq \Pi_{t^*}$.

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 9.13. *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 9.12 that u contains at least two supports. Suppose for a contradiction that u consists of only two supports. By Lemma 9.12, 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 9.14. *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 9.13, 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 9.12, 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, 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 9.15. *Line m lies in a plane $\Psi \neq \Pi_{t^*}$ containing a face of \mathbf{P} .*

Proof. By Lemma 9.14, m contains an edge of P_{t^*} ; see Figure 9.8. 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

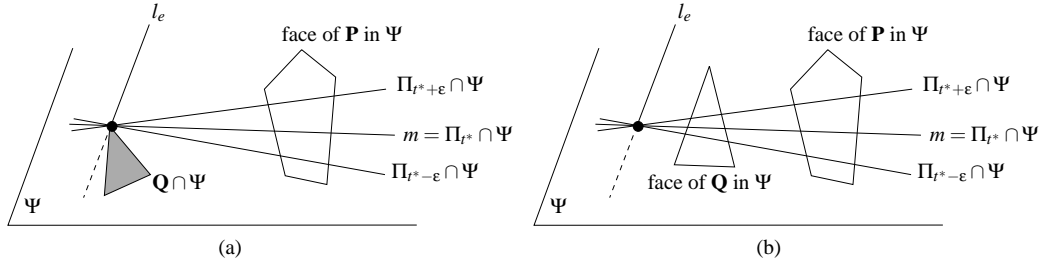


FIG. 9.9 – m is tangent to \mathbf{P} along a face in Ψ and (a) to $\mathbf{Q} \cap \Psi$ only on l_e or (b) to \mathbf{Q} along a face in Ψ .

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 9.14 and 9.15.

Lemma 9.16. *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 9.7 and thus proving Lemma 9.16.

Since m is tangent to \mathbf{Q} (by Lemma 9.14), m is tangent to $\mathbf{Q} \cap \Psi$ only on l_e (see Figure 9.9(a)), or m properly intersects $\mathbf{Q} \cap \Psi$ which is then a face or an edge of \mathbf{Q} (see Figure 9.9(b))¹⁷. 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 9.15.

The following Lemmas 9.17 and 9.18 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 9.14), 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 9.14), 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 9.12, 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 9.7 and thus concluding the proof of Lemma 9.16. \square

Lemma 9.17. *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 9.12 and thus lies in Π_t for all t . A support vertex in u also lies on m by Lemma 9.14 and thus lies in plane Ψ by Lemma 9.15. 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 9.14) 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 9.9(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 $t \in \mathcal{N}$.

¹⁷Note that in these two situations, two edges of two distinct polytopes are then coplanar (in the first case an edge of \mathbf{Q} and l_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.

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 9.9(a)). Then v lies in Ψ (since $m \subset \Psi$ by Lemma 9.15) 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 9.12, 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 9.12, 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 9.18. *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 9.15) 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 9.14. 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 9.17, 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 9.17, 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

9.2.5 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 9.19. *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 9.10. 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.

9.2.6 Proof of the Main Lemma

Proposition 9.11, which handles the isolated generic tangent lines, and Proposition 9.19, which handles the isolated non-generic tangent lines, directly yield the Main Lemma.

9.3 Upper bounds

We prove in this section the upper bounds of Theorems 9.1, 9.2, and 9.3. The lower bounds are proved in Section 9.4. 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 9.20. *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 9.21. *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 [BEL⁺05] 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 9.3.

Proposition 9.22. *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 9.21. 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 9.20 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 9.1.

Proposition 9.23. *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 9.22, 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 9.2. We start by two preliminary lemmas.

Lemma 9.24. *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. [SA95, §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 9.2, 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 9.25. *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 9.23, 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 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 9.24 thus yields the result. \square

We can now prove the upper bound of Theorem 9.2.

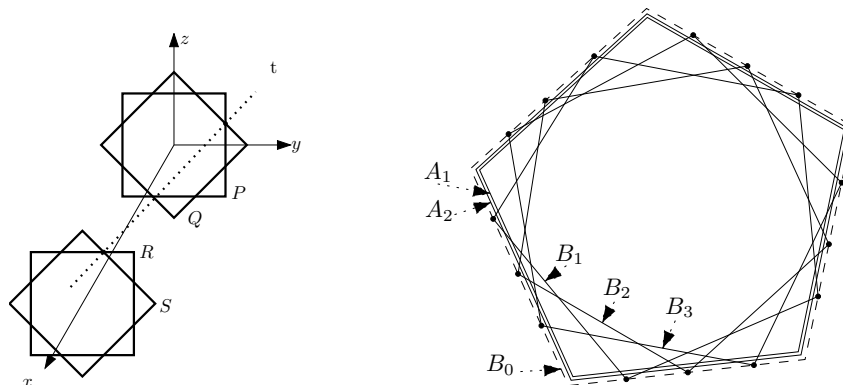


FIG. 9.10 – Lower bound examples for Lemmas 9.27 and 9.28.

Proposition 9.26. *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 9.25 yields that the number of connected components of lines tangent to the four polytopes is $O(mn)$. \square

9.4 Lower bounds

We provide in this section the lower-bound examples needed for Theorems 9.1, 9.2, and 9.3. The following proposition proves the lower bound of Theorem 9.2.

Lemma 9.27. *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 9.10). 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 9.1 and 9.3. 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 9.28. *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 $\Omega(n^2 k^2)$. 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 9.10). 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 9.27 with polygons of size $\frac{n}{k}$ and thus with total number of tangents larger than $(\frac{k}{4})^4 4 (\frac{n}{k})^2 = \frac{n^2 k^2}{4}$. 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 9.27. \square

9.5 Algorithm

Using the sweep-plane algorithm outlined in Section 2.1, we can compute in $O(n^2 k^2 \log n)$ 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 $O(n^2 k^2)$ segments are free; this can be done in $O(\log^2 n)$ -time per query using $O((nk)^{2+\varepsilon})$ preprocessing time and storage [AS96].

We present in this section a solution that uses $O(n^2 k^2 \log n)$ 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 9.4 and directly yields it.

Theorem 9.29. *Given k polytopes in \mathbb{R}^3 with n edges in total, we can compute in $O(n^2 k^2 \log n)$ 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 $O(n^2 k^2 \log n)$ 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 9.5.2 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.

9.5.1 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, [BY98, §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 [O'R98, §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 9.2.3, 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 9.11b). 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 9.12).

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.

9.5.2 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.

9.5.3 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

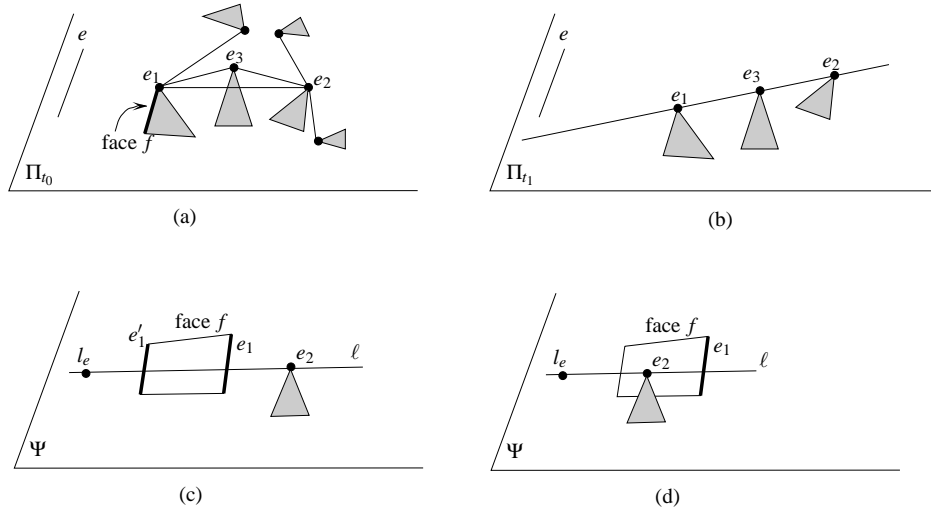


FIG. 9.11 – (a) The sweep plane in which the combinatorial bitangent with support edges e_1 and e_2 is created. (b) The sweep plane at a T-critical event induced by the three bitangents with support edges in e_1 , e_2 , and e_3 . (c-d) A line ℓ that defines an F-critical event. (d) The F-event defined by ℓ occurs simultaneously with an I-critical event.

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 9.5.4 (Lemma 9.30). 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 9.29.

9.5.4 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_0 denote the critical plane at which the new combinatorial bitangent is created.

New T-critical events. See Figure 9.11a-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 9.11c). 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_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 9.5.5.2, 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 9.11d), 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 9.30. *Each time a combinatorial bitangent is created, the event queue can be updated in $O(k \log n)$ time.*

9.5.5 Processing events

9.5.5.1 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 9.30). 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 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 $\Pi_{t_0+\varepsilon}$ that goes through these two edges is tangent to \mathbf{Q} for $\varepsilon > 0$ 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 9.30). 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.

9.5.5.2 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 9.11c-d). By construction of F-events (see Section 9.5.4), e_2 does not intersect face f (see Figure 9.11c), 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 9.11c). 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

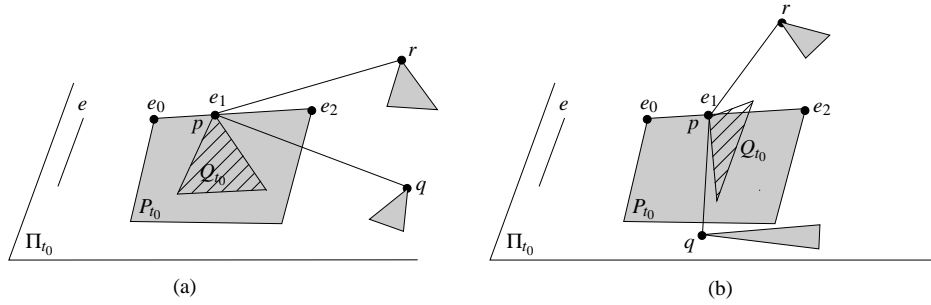


FIG. 9.12 – I-critical event.

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 9.30).

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.

9.5.5.3 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_{i_0} , that contains p intersects the two polytopes \mathbf{P} and \mathbf{Q} in two polygons P_{i_0} and Q_{i_0} . See Figure 9.12. Point p lies on an edge of P_{i_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 9.12a), 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 9.12), polytope \mathbf{P} was and remains a blocker of the bitangent. In the second case (e.g., segment pr in Figure 9.12), \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.

9.5.5.4 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 9.22), thus all the T-events can be processed in $O(nk^2 \log n)$ time per sweep.

9.5.6 Complexity

Note first that we assume a model of computation in which bounded-degree algebraic polynomials may be evaluated in constant time. See [ELL⁺06] 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 9.4 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 9.4 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 $O(nk)$.

9.6 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.

Chapitre 10

Towards an implementation of the 3D visibility skeleton

Cet article court présentant une vidéo va être publiée dans les proceedings du *23th ACM Annual Symposium on Computational Geometry* [ZELW07].

Abstract

In this note we describe the contents of a video illustrating an algorithm for computing the 3D visibility skeleton of a set of disjoint convex polytopes. The video can be found at <http://www.cs.mcgill.ca/~lzhang15/video/> with file name `socg07visidemo.mov`.

10.1 Introduction

The 3D visibility skeleton is a graph whose vertices correspond to the maximal free line segments that are transversal to four edges of at least three distinct polytopes and tangent to those polytopes ; its arcs correspond to sets of maximal free line segments that are tangent to three polytopes [DDP97]. The visibility skeleton has been used for visibility computations such as computing shadow boundaries [DDP99, DD02].

This video demonstrates a sweep plane algorithm for capturing the vertices of the 3D visibility skeleton of a set of polytopes in 3D [Goa04, BDD⁺07].

10.2 The algorithm

The input of the algorithm is a set of k disjoint convex polytopes in general position with n edges in total. The output of the algorithm is the set of $O(n^2k^2)$ vertices of the 3D visibility skeleton of the input polyhedra. The algorithm, which runs in $O(n^2k^2 \log n)$ time, can also be used to compute the arcs of the skeleton.

The algorithm performs a rotational plane sweep around each edge e of each polytope, sweeping from one incident face of that edge to the other incident face. The sweep plane intersects the polytopes in at most k disjoint convex polygons, which change their shape as the sweep plane rotates. Figure 10.1(a) shows one position of the sweep plane, drawn as a faint grid, as it rotates around edge e of polytope **C**. Polytopes **A** and **B** are intersected by the sweep plane. Polytope **C** lies above the plane, with edge e in the plane. Polytope **D** lies below the plane. Figure 10.1(b) shows the view inside the sweep plane. The polytopes **A** and **B** intersect the plane in convex polygons A and B , which support 4 bitangents. Figure 10.2 shows the 2D visibility skeleton corresponding to Figure 10.1(b). The circular cycle of directed arcs gives the ordering of the 4 bitangents around polygon A ; the cycle of the remaining directed arcs gives the ordering of the 4 bitangents around polygon B .

During the sweep, the algorithm maintains the 2D visibility skeleton of the intersected polytopes [PV96a]. The 2D visibility skeleton for the convex polygons in the initial sweep plane is computed and then used to determine

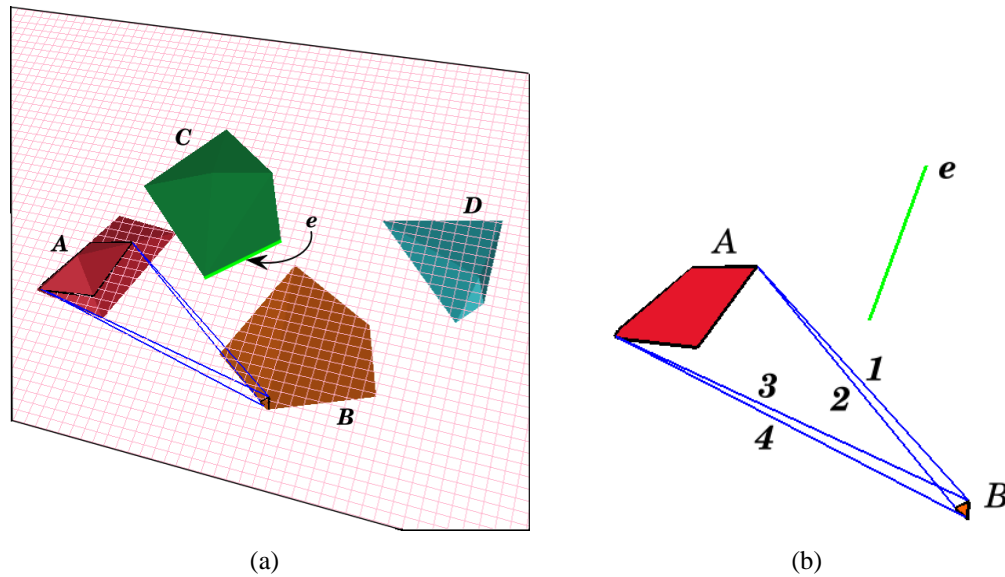


FIG. 10.1 – (a) One position of the sweep plane. (b) The view inside the sweep plane.

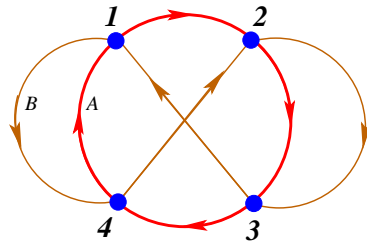


FIG. 10.2 – The 2D visibility skeleton for Figure 10.1(b).

the initial queue of critical events that will occur during the sweep. The sweep planes at which these critical events occur are called *event planes*. At a critical event, the 2D visibility skeleton may change its topology and the algorithm updates it, as well as the queue of critical events.

There are three types of critical events. A *V-event* occurs when the sweep plane encounters a polytope vertex that supports one or more bitangents in the sweep plane. A *T-event* occurs when two or three bitangents become colinear. An *F-event* occurs when a bitangent becomes colinear with a face of a polytope. There are $O(nk^2)$ events per sweep, and they can be computed and processed in $O(nk^2 \log n)$ time in the usual sweep algorithm paradigm.

The vertices of the 3D visibility skeleton are captured during the sweep, as they correspond to the *V*, *T*, *F*-events whose associated bitangents intersect the edge e that the sweep plane is rotating about. After the n sweeps, a description of the arcs of the 3D visibility skeleton can be computed, although the details are not illustrated in the video.

10.3 Implementation issues and technical details

Although the algorithm as described in [BDD⁺07] works for any set of possibly intersecting convex polytopes in any configuration, the current implementation requires that the polytopes satisfy certain general position assumptions¹⁸.

¹⁸The precise definition of our general position assumptions is straightforward but lengthy. It guarantees, for example, that each critical event corresponds to a unique position of the sweep plane.

The key predicate of the sweep algorithm compares two event planes, to order their occurrences in the sweep. A detailed study of this predicate and its degree is given in [ELL⁺06].

The algorithm was implemented in C++ using the *CGAL* [CGA] library. We used the *CORE* library [COR] to perform exact comparisons of algebraic numbers. To compute the 2D visibility skeleton we used the CGAL-based package due to Angelier and Pocchiola [AP03a], based on the Greedy Flip Algorithm [AP03b, PV96a].

The graphical output was produced using the *Geomview* software [Geo] through the interface provided by the *CGAL* library. We took snapshots of the Geomview window display, while rotating the viewpoint to provide a 3D view of the objects in the display window. Finally, we used *iMovie* [iMo] to assemble all the snapshots together into the final video. We used the *Audacity* [Aud] software for the audio.

Acknowledgments

We would like to thank Ethan Kim, Christophe Weibel, and Nathan Yu for their help.

Chapitre 11

The expected number of 3D visibility events is linear

Cet article est paru dans *SIAM Journal on Computing* [DDE⁺03] ainsi que dans la thèse de X. Goaoc [Goa04].

Abstract

In this paper, we show that, amongst n uniformly distributed unit balls in \mathbb{R}^3 , the expected number of maximal non-occluded line segments tangent to four balls is linear. Using our techniques we show a linear bound on the expected size of the visibility complex, a data structure encoding the visibility information of a scene, providing evidence that the storage requirement for this data structure is not necessarily prohibitive. These results significantly improve the best previously known bounds of $O(n^{8/3})$ [DDP02].

Our results generalize in various directions. We show that the linear bound on the expected number of maximal non-occluded line segments that are not too close to the boundary of the scene and tangent to four unit balls extends to balls of various but bounded radii, to polyhedra of bounded aspect ratio, and even to non-fat 3D objects such as polygons of bounded aspect ratio. We also prove that our results extend to other distributions such as the Poisson distribution. Finally, we indicate how our probabilistic analysis provides new insight on the expected size of other global visibility data structures, notably the aspect graph.

11.1 Introduction

Visibility computations are central in computer graphics applications. Computing the limits of the umbra and penumbra cast by an area light source, identifying the set of blockers between any two polygons and determining the view from a given point are examples of visibility queries that are essential for the realistic rendering of 3D scenes. In global illumination algorithms, where the flow of light in a scene is simulated according to the laws of geometrical optics, visibility computations are excessively costly. In fact, more than half of the overall computation time can routinely be spent on visibility queries in radiosity simulations [HSD94].

One approach to speeding up rendering is to store global visibility information in a data structure which can then be efficiently queried. The visibility complex, a partition of the set of maximal free line segments, has been proposed as a unified data structure encoding the visibility information of a scene [PV96b] and has been used for rendering purposes [DDP02]. Other related data structures include Pellegrini's ray-shooting structure [Pel93], the aspect graph [PD90] and the visual hull [Lau94]; see [Dur00] for a recent survey.

One problem with these types of data structures which may prevent their application in practice is their potentially enormous size; the size of the visibility complex of a set of n triangles in \mathbb{R}^3 is $\Theta(n^4)$ in the worst case [DDP02], which is prohibitive even for scenes of relatively modest size. Worst-case examples are somewhat artificial and indeed Durand, Drettakis and Puech [DDP97] provide empirical evidence indicating that these

	Worst-case	Expected
possibly occluded lines amongst unit balls	$\Theta(n^4)$	$O(n^{\frac{8}{3}})$ [DDP02]
free lines amongst unit balls	$\Omega(n^2) [\star], O(n^{3+\epsilon})$ [AAS99]	$\Theta(n) [\star]$
free lines amongst disjoint homothetic polytopes	$\Omega(n^3)$ [dBEG98]	
free segments amongst unit balls	$\Omega(n^2) [\star], O(n^4)$	$\Theta(n)[\star]$
free segments amongst arbitrary sized balls	$\Omega(n^3)$ [DR01], $O(n^4)$	
visibility complex of unit balls	$\Omega(n^2) [\star], O(n^4)$	$\Theta(n) [\star]$

TAB. 11.1 – Known bounds on the complexity of the set of lines, free lines or maximal free line segments tangent to 4 amongst n objects. The expected complexities are calculated for the uniform distribution. The results referenced by \star are established in this paper.

worst-case upper bounds are largely pessimistic in practical situations ; they observe a quadratic growth rate, albeit for rather small scenes. In 2D, while the worst-case complexity of the visibility complex is quadratic, experimental results strongly suggest that the size of the visibility complex of a scene consisting of scattered triangles is linear [CF99].

Our goal is to provide theoretical evidence to support these observations. To this end we investigate the *expected size* of the visibility complex, or equivalently, the expected number of visibility events, occurring in scenes in \mathbb{R}^3 . A visibility event is a combinatorial change in the view of a moving observer ; such an event occurs when the viewing direction becomes tangent to some objects. For sets of convex objects in general position in \mathbb{R}^3 , the viewing direction can be tangent to at most four objects. Visibility events thus correspond to maximal non-occluded line segments tangent to at most four objects ; combinatorially different visibility events correspond to the faces of the visibility complex.

In this paper we prove that the expected number of maximal non-occluded line segments tangent to four balls, amongst n uniformly distributed unit balls in \mathbb{R}^3 , is linear. This improves the previously known upper bound of $O(n^{8/3})$ by Durand et al. who proved the more general result that the expected number of (possibly occluded) lines tangent to four balls is $O(n^{8/3})$ for the same model [DDP02]. The intuition behind our proof is that, given a line segment tangent to four balls, the probability that this segment is not occluded by any other ball is the probability that a cylinder-like volume of radius 1 about the segment is free from the centers of the other balls. This probability decays roughly exponentially fast with the length of the segment, yielding the result. Using our techniques we then show a linear bound on the expected size of the visibility complex of n uniformly distributed unit balls in \mathbb{R}^3 . A simple computation then provides us with the same result for the Poisson distribution.

Our results generalize in the following ways. We show that, for certain types of visibility events, the linear bound also applies to balls of various but bounded radii, to polyhedral objects enclosed between two concentric balls of fixed radius, and even to non-fat objects such as polygons, enclosed between two concentric circles of fixed radius, whose centers and normals are uniformly distributed. For the remaining types of visibility events (namely those occurring close to the boundary of the scene – see Section 11.7.3 for the details), we prove only an $O(n^2)$ bound, which is still an improvement over the bound by Durand et al. [DDP02].

Of course objects in graphics scenes are seldom distributed uniformly or according to a Poisson point process. We chose this model because it allows tractable proofs of theoretical results. This is important in a context where there are few rigorous results either theoretical or experimental. The same model, albeit with significant simplifying assumptions, has also been used to study the average complexity of ray shooting [SKHBS02, SKM98] and occlusion culling for 2D urban scenes [NFLYCO99]. It is interesting to note that Szirmay-Kalos et al. [SKHBS02], after establishing bounds on the average complexity of ray shooting in scenes consisting of unit balls distributed according to a Poisson point process, tested their algorithms on a small number of realistic scenes. The results they obtain are consistent with those predicted by the theoretical results thus providing some evidence that the model is helpful. No other model has been widely accepted by the graphics community and, in fact, generating meaningful random scenes usable for testing algorithms is a major problem. (Note that rather than attempting to generate random scenes, an alternative approach, which has been used to study the average complexity of ray shooting, is to fix the scene and randomly distribute the rays ; see, for example, [ABCC02].)

Previous results on this topic include those that bound the number of lines and the number of free (i.e., non-occluded) lines amongst different sets of objects. They are summarized in Table 11.1. Agarwal, Aronov and Sha-

ir [AAS99] showed an upper bound of $O(n^{3+\epsilon})$ on the complexity of the space of line transversals of n balls by studying the lower envelope of a set of functions. A study of the upper envelope of the same set of functions yields the same upper bound on the number of free lines tangent to four balls [DR01]. Agarwal et al. [AAS99] also showed a lower bound on the complexity of the space of line transversals of n balls of $\Omega(n^3)$ for arbitrarily sized balls and $\Omega(n^2)$ for unit sized balls. De Berg, Everett and Guibas [dBEG98] showed a $\Omega(n^3)$ lower bound on the number of free lines (and thus free segments) tangent to four amongst n disjoint homothetic convex polyhedra. Recently, Devillers and Ramos [DR01] presented a simple $\Omega(n^3)$ lower bound on the number of free segments tangent to 4 amongst n arbitrarily sized balls, which also holds for non-intersecting balls. We also present a simple $\Omega(n^2)$ lower bound on the number of free segments tangent to 4 amongst n unit balls.

In the next section we carefully define the problem and state our main results. In Section 11.3 and Section 11.4 we prove the expected upper and lower linear bounds on the number of free segments tangent to four balls. In Section 11.5 we extend this result to the visibility complex. We present in Section 11.6 a $\Omega(n^2)$ worst-case lower bound. In Section 11.7 we discuss extensions of our results to some other models. We conclude in Section 11.8.

11.2 Our model and results

We first describe our objects and their distribution. Let $n \in \mathbb{N}$ and μ be a positive constant. A sample scene consists of n unit radius balls B_1, \dots, B_n whose centers p_1, \dots, p_n are independently chosen from the uniform distribution over a universal ball \mathcal{U} of radius R centered at O . Since we distribute the centers p_i over \mathcal{U} , the balls B_i may intersect each other and are contained in the ball, denoted \mathcal{U}^+ , whose radius is $R+1$ and whose center is that of \mathcal{U} .

We define the radius R of the universal ball \mathcal{U} to be a function of n satisfying

$$R^3 = n/\mu. \quad (11.1)$$

The constant μ reflects the density of the balls in the sense that the expected number of centers lying in any given solid of volume V in the universe is $\frac{3}{4\pi}\mu V$. (The model is interesting only if n is asymptotically proportional to R^3 . Indeed, if $\frac{n}{R^3}$ tends to infinity when n tends to infinity, then the universe gets entirely filled up with balls and visibility events only occur in $\mathcal{U}^+ \setminus \mathcal{U}$. Conversely, if $\frac{n}{R^3}$ tends to zero when n tends to infinity, then the balls get scattered so far apart that the probability that any four (or three) balls have a common tangent goes to zero.)

We now define the *visibility complex* of a set of objects [PV96b]. A *free* or *non-occluded* segment is a line segment that does not intersect the interior of any object. A free segment is maximal if it is not properly contained in another one. Thus, the endpoints of a maximal free segment are either on an object or at infinity. We say that two maximal free segments are similar if their endpoints lie on the same objects (possibly at infinity). The visibility complex of a collection of objects is roughly defined as the partition of the space of maximal free segments into connected components of similar segments¹⁹. Its faces have dimension between 0 and 4; when the objects are in adequate general position, a k -dimensional face corresponds to a connected set of similar maximal non-occluded line segments tangent to $4-k$ objects.

In order to bound the total number of faces of the visibility complex, we first bound the number of 0-faces. To do this, we count the *T4-segments*, which are the free segments tangent to 4 balls with endpoints on two of those balls. Since there is a one-to-one correspondence between 0-faces and *T4*-segments when the objects are in adequate general position, this yields a bound on the expected number of vertices of the visibility complex. Note that since the balls are contained in \mathcal{U}^+ , the *T4*-segments are also contained in \mathcal{U}^+ .

Our main result is the following.

Theorem 11.1. *The expected number of T4-segments amongst n uniformly distributed unit balls is $\Theta(n)$.*

We extend this result to the higher dimensional faces of the complex.

Theorem 11.2. *The expected size of the visibility complex of n uniformly distributed unit balls is $\Theta(n)$.*

¹⁹Formally, we consider the space of free segments quotiented by the equivalence relation that is the transitive and reflexive closure of the inclusion. In other words, two free segments are identified if they are both contained in the same maximal free segment. This allows the cells of the partition to be connected.

We also present an $\Omega(n^2)$ worst-case lower bound on the number of $T4$ -segments amongst n unit balls in \mathbb{R}^3 (see Proposition 11.27). In fact the lower bound holds for the number of k -faces of the visibility complex, for all k between 0 and 4.

11.3 The expected number of $T4$ -segments is at most linear

The general idea behind the proof of the upper bound of Theorem 11.1 is the following. For any ordered choice of four balls, we bound from above the probability that a line is tangent to these balls in the given order and is not occluded in between its contact points with the balls. Then we sum these probabilities over all ordered quadruples of balls and all potential tangent lines to these balls.

For any two points p and q , and positive real number α , let $\mathcal{H}(p, q, \alpha)$ denote the union of all the balls of radius α centered on the line segment pq (see Figure 11.1). We first show that a line is tangent to four balls B_i, B_j, B_k and B_l in that order only if p_j and p_k are in $\mathcal{H}(p_i, p_l, 2)$. Thus the volume of $\mathcal{H}(p_i, p_l, 2) \cap \mathcal{U}$ gives an upper bound on the probability that a line tangent to the four balls, in the given order, exists.

We next show that a segment tangent to four balls B_i, B_j, B_k and B_l in that order, at points t_i, t_j, t_k and t_l , respectively, is not occluded if and only if the centers of all remaining balls are outside or on the boundary of $\mathcal{H}(t_i, t_l, 1)$. The volume of $\mathcal{U} \setminus \mathcal{H}(t_i, t_l, 1)$ gives an upper bound on the probability that the tangent segment is not occluded. Thus, to get an upper bound on that probability, we need a lower bound on the volume of $\mathcal{H}(t_i, t_l, 1) \cap \mathcal{U}$.

To bound the probability that a $T4$ -segment exists, we integrate over the distance between p_i and p_l , and over the distance from p_i to the boundary of the universe \mathcal{U} . This integral is split into three parts covering the cases where

- (i) B_i and B_l are close to one another,
- (ii) at least one of B_i and B_l is entirely inside the universe,
- (iii) B_i and B_l are not close to one another and both are partially outside the universe.

In each case we over-estimate the volume of $\mathcal{H}(p_i, p_l, 2) \cap \mathcal{U}$ and under-estimate the volume of $\mathcal{H}(t_i, t_l, 1) \cap \mathcal{U}$. We apply the same general proof technique in each of the three cases. While Case (ii) illustrates the main idea behind the proof (Case (i) being a simplified version), extending this idea to Case (iii) is technically challenging because of the difficulties caused by the boundary of the universe.

11.3.1 Definitions

Let \mathcal{N} be the set of ordered 4-tuples (i, j, k, l) chosen from $\{1, 2, \dots, n\}$ such that i, j, k, l are pairwise distinct. In our model, the probability that four centers are collinear is zero, so we may assume that any set of four balls admits at most 12 real common tangent lines [DMPT01, MPT01]. Moreover, the real common tangent lines correspond to the real solutions of a degree 12 system of equations. For any set of four balls we order arbitrarily the 12 solutions of the associated system.

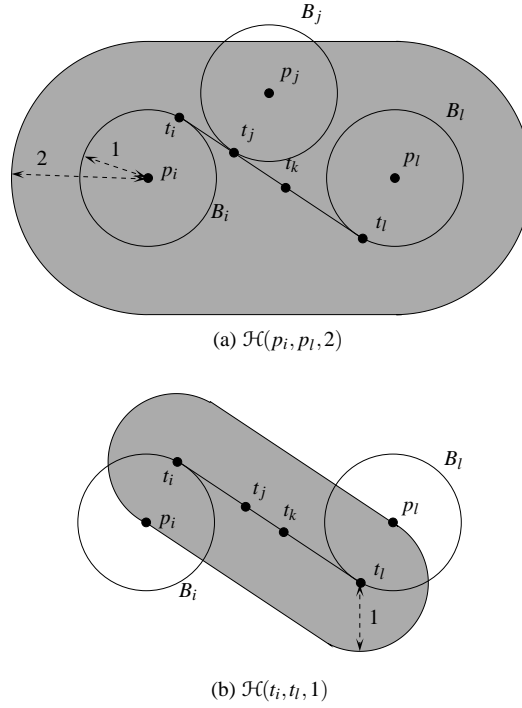
Given four balls B_i, B_j, B_k and B_l , we denote by $\mathcal{L}_{i,j,k,l}^\omega$, for ω in $\{1, \dots, 12\}$, the event that the ω^{th} solution of the system is real, that the corresponding real tangent line is tangent to the four balls B_i, B_j, B_k and B_l in that order, and that p_i is not closer than p_l to the boundary of \mathcal{U} . Whenever $\mathcal{L}_{i,j,k,l}^\omega$ occurs, we denote the points of tangency of that line on B_i, B_j, B_k, B_l by t_i, t_j, t_k, t_l , respectively. Let $\delta_{i,j,k,l}^\omega$ be the event that $\mathcal{L}_{i,j,k,l}^\omega$ occurs and the line segment $t_i t_l$ is not occluded. Notice that if $\delta_{i,j,k,l}^\omega$ occurs, the balls B_i, B_j, B_k, B_l define a $T4$ -segment, and that a $T4$ -segment corresponds to a unique $\delta_{i,j,k,l}^\omega$.

Let $\mathbf{x}_{i,l}$ be the random variable representing the distance from p_i to p_l , and \mathbf{y}_i (resp. \mathbf{y}_l) be the random variable denoting the distance from p_i (resp. p_l) to the boundary of the universe.

In the sequel, a random point p denotes a point chosen from the uniform distribution over \mathcal{U} .

11.3.2 The Proof

There is a one-to-one correspondence between the $T4$ -segments and the events $\delta_{i,j,k,l}^\omega$ that occur. We thus have the following straightforward lemma.

FIG. 11.1 – $\mathcal{H}(p_i, p_l, 2)$ and $\mathcal{H}(t_i, t_l, 1)$ are shown shaded.

Lemma 11.3. *The expected number of T4-segments amongst n uniformly distributed unit balls is*

$$\sum_{(i,j,k,l) \in \mathcal{N}} \sum_{\omega=1}^{12} \Pr(\delta_{i,j,k,l}^{\omega}).$$

We bound the probability $\Pr(\delta_{i,j,k,l}^{\omega})$ by integrating over the distance x between p_i and p_l , and over the distance y from p_i to the boundary of the universe \mathcal{U} . The integral is split into three parts covering the cases where (i) the balls B_i and B_l are close to one another, (ii) p_i is at distance at least 1 from the boundary of \mathcal{U} , and (iii) the balls B_i and B_l are not close to one another and p_i is at distance less than 1 from the boundary of \mathcal{U} . Note that in the last case, if $\delta_{i,j,k,l}^{\omega}$ occurs, then both ball centers p_i and p_l are within distance 1 from the boundary of \mathcal{U} . Two balls are considered close to one another if their centers are closer than some sufficiently large constant; for technical reasons which are embedded in the proof of Proposition 11.29, we actually define *close* to mean distance at most 6.

Lemma 11.4. $\Pr(\delta_{i,j,k,l}^{\omega}) \leq I_{x \leq 6} + I_{y \geq 1} + I_{x > 6, y < 1}$, where

$$\begin{aligned} I_{x \leq 6} &= \int_{x=0}^6 \Pr(\delta_{i,j,k,l}^{\omega} \mid \mathbf{x}_{i,l} = x) \cdot \Pr(x \leq \mathbf{x}_{i,l} < x + dx), \\ I_{y \geq 1} &= \int_{x=0}^{2R} \Pr(\delta_{i,j,k,l}^{\omega} \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i \geq 1) \cdot \Pr(x \leq \mathbf{x}_{i,l} < x + dx \mid \mathbf{y}_i \geq 1), \\ I_{x > 6, y < 1} &= \int_{x=6}^{2R} \int_{y=0}^1 \Pr(\delta_{i,j,k,l}^{\omega} \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i) \\ &\quad \cdot \Pr((x \leq \mathbf{x}_{i,l} < x + dx) \cap (\mathbf{y}_l \leq \mathbf{y}_i) \mid \mathbf{y}_i = y) \\ &\quad \cdot \Pr(y \leq \mathbf{y}_i < y + dy). \end{aligned}$$

Proof. By the Total Probability Theorem (see [Pap91]),

$$\Pr(\delta_{i,j,k,l}^{\omega}) = \int_{x=0}^{2R} \Pr(\delta_{i,j,k,l}^{\omega} \mid \mathbf{x}_{i,l} = x) \cdot \Pr(x \leq \mathbf{x}_{i,l} < x + dx).$$

The integral can be split at $x = 6$, giving $I_{x \leq 6}$. Then applying the Total Probability Theorem on what remains, we get

$$\int_{x=6}^{2R} \int_{y=0}^R \Pr(\delta_{i,j,k,l}^{\circ} \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y) \cdot \Pr((x \leq \mathbf{x}_{i,l} < x + dx) \mid \mathbf{y}_i = y) \cdot \Pr(y \leq \mathbf{y}_i < y + dy) \quad (11.2)$$

which can be split at $y = 1$. The part corresponding to y between 1 and R is equal to

$$\begin{aligned} & \int_{x=6}^{2R} \int_{y=1}^R \Pr(\delta_{i,j,k,l}^{\circ} \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_i \geq 1) \\ & \cdot \Pr((x \leq \mathbf{x}_{i,l} < x + dx) \mid \mathbf{y}_i = y, \mathbf{y}_i \geq 1) \cdot \Pr(y \leq \mathbf{y}_i < y + dy) \\ & \leq \int_{x=6}^{2R} \int_{y=0}^R \Pr(\delta_{i,j,k,l}^{\circ} \cap (x \leq \mathbf{x}_{i,l} < x + dx) \mid \mathbf{y}_i = y, \mathbf{y}_i \geq 1) \cdot \Pr(y \leq \mathbf{y}_i < y + dy). \end{aligned}$$

Applying the Total Probability Theorem again, we get

$$\int_{x=6}^{2R} \Pr(\delta_{i,j,k,l}^{\circ} \cap (x \leq \mathbf{x}_{i,l} < x + dx) \mid \mathbf{y}_i \geq 1)$$

which is less than $I_{y \geq 1}$. Consider now the part of (11.2) for y between 0 and 1. If $\mathbf{y}_l > \mathbf{y}_i$ then $\delta_{i,j,k,l}^{\circ}$ does not occur (by definition of $\mathcal{L}_{i,j,k,l}^{\circ}$), thus we have

$$\begin{aligned} & \Pr(\delta_{i,j,k,l}^{\circ} \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y) \cdot \Pr((x \leq \mathbf{x}_{i,l} < x + dx) \mid \mathbf{y}_i = y) \\ & = \Pr(\delta_{i,j,k,l}^{\circ} \cap (x \leq \mathbf{x}_{i,l} < x + dx) \mid \mathbf{y}_i = y) \\ & = \Pr(\delta_{i,j,k,l}^{\circ} \cap (x \leq \mathbf{x}_{i,l} < x + dx) \cap (\mathbf{y}_l \leq \mathbf{y}_i) \mid \mathbf{y}_i = y) \\ & = \Pr(\delta_{i,j,k,l}^{\circ} \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i) \\ & \cdot \Pr((x \leq \mathbf{x}_{i,l} < x + dx) \cap (\mathbf{y}_l \leq \mathbf{y}_i) \mid \mathbf{y}_i = y). \end{aligned}$$

Thus, the part of (11.2) for y between 0 and 1 is equal to $I_{x > 6, y < 1}$. \square

Let Ξ denote any of the following events : $(\mathbf{x}_{i,l} = x)$, $(\mathbf{x}_{i,l} = x, \mathbf{y}_i \geq 1)$, $(\mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i)$. The next three lemmas are used to bound $\Pr(\delta_{i,j,k,l}^{\circ} \mid \Xi)$ appearing in the three integrals $I_{x \leq 6}$, $I_{y \geq 1}$ and $I_{x > 6, y < 1}$.

Lemma 11.5. *If a line is tangent to four balls B_i, B_j, B_k, B_l in that order at t_i, t_j, t_k, t_l , respectively, then $p_j, p_k \in \mathcal{H}(p_i, p_l, 2)$. Also, the segment $t_i t_l$ is not occluded if and only if the interior of $\mathcal{H}(t_i, t_l, 1)$ does not contain the center of any other ball.*

Proof. Segment $t_i t_l$ is contained in $\mathcal{H}(p_i, p_l, 1)$. Since t_j and t_k belong to that segment, t_j and t_k are also in $\mathcal{H}(p_i, p_l, 1)$. Thus p_j, p_k are both in $\mathcal{H}(p_i, p_l, 2)$. See Figure 11.1 (a).

The segment $t_i t_l$ is occluded if and only if some ball B_γ , $\gamma \neq i, j, k, l$, properly intersects it, that is the center of B_γ lies in the interior of $\mathcal{H}(t_i, t_l, 1)$. See Figure 11.1 (b). \square

Lemma 11.6. $\Pr(p \in \mathcal{H}(p_i, p_l, 2) \mid \Xi) \leq \frac{(3x+8)}{R^3}$.

Proof.

$$\Pr(p \in \mathcal{H}(p_i, p_l, 2) \mid \Xi) = \frac{\text{Volume of } \mathcal{H}(p_i, p_l, 2) \cap \mathcal{U}}{\text{Volume of } \mathcal{U}} \mid \Xi \leq \frac{\text{Volume of } \mathcal{H}(p_i, p_l, 2)}{\text{Volume of } \mathcal{U}} \mid \Xi.$$

When Ξ occurs, $\mathbf{x}_{i,l} = x$ and the volumes of $\mathcal{H}(p_i, p_l, 2)$ and \mathcal{U} are $\frac{4\pi}{3}(3x+8)$ and $\frac{4\pi}{3}R^3$, respectively. Thus

$$\Pr(p \in \mathcal{H}(p_i, p_l, 2) \mid \Xi) \leq \frac{3x+8}{R^3}.$$

\square

Lemma 11.7. $\Pr(\delta_{i,j,k,l}^\circ | \Xi) \leq \frac{(3x+8)^2}{R^6} \cdot \Pr(p \notin \mathcal{H}(t_i, t_l, 1) | \mathcal{L}_{i,j,k,l}^\circ, \Xi)^{n-4}$.

Proof. If $\delta_{i,j,k,l}^\circ$ occurs, then $\mathcal{L}_{i,j,k,l}^\circ$ necessarily occurs, thus

$$\Pr(\delta_{i,j,k,l}^\circ | \Xi) = \Pr(\delta_{i,j,k,l}^\circ \cap \mathcal{L}_{i,j,k,l}^\circ | \Xi) = \Pr(\mathcal{L}_{i,j,k,l}^\circ | \Xi) \cdot \Pr(\delta_{i,j,k,l}^\circ | \mathcal{L}_{i,j,k,l}^\circ, \Xi).$$

By Lemma 11.5, $\Pr(\mathcal{L}_{i,j,k,l}^\circ | \Xi)$ is bounded by the probability that p_j and p_k belong to $\mathcal{H}(p_i, p_l, 2)$ given Ξ , and $\Pr(\delta_{i,j,k,l}^\circ | \mathcal{L}_{i,j,k,l}^\circ)$ is equal to the probability that for all $\gamma \neq i, j, k, l$, point p_γ is outside $\mathcal{H}(t_i, t_l, 1)$ given Ξ . Since all the points are independently and identically drawn from the uniform distribution over \mathcal{U} , Lemma 11.6 yields the result. \square

We consider the three integrals $I_{x \leq 6}$, $I_{y \geq 1}$ and $I_{x > 6, y < 1}$ in the following subsections, and prove that each is bounded by $O\left(\frac{1}{n^3}\right)$. This will complete the proof of the upper bound of Theorem 11.1 since, by Lemmas 11.3 and 11.4, the expected number of T_4 -segments is less than $12 \binom{n}{4} (I_{x \leq 6} + I_{y \geq 1} + I_{x > 6, y < 1})$.

B_i and B_l are close to one another

We prove here that $I_{x \leq 6}$ is $O\left(\frac{1}{n^3}\right)$. When B_i and B_l are close to one another, the probability that there exist two other balls, B_j and B_k , defining a line tangent to B_i, B_j, B_k, B_l in that order, is small enough that we do not need to consider occlusions in order to get the bound we want.

We first bound the term $\Pr(x \leq \mathbf{x}_{i,l} < x + dx)$ appearing in the integral $I_{x \leq 6}$.

Lemma 11.8. $\Pr(x \leq \mathbf{x}_{i,l} < x + dx) \leq \frac{3x^2}{R^3} dx$.

Proof. When p_i is given, p_l must belong to a spherical shell between two spheres of center p_i and radii x and $x + dx$. The probability $\Pr(x \leq \mathbf{x}_{i,l} < x + dx)$, if p_i is known, is exactly the volume of the part of the spherical shell inside \mathcal{U} divided by the volume of \mathcal{U} . The volume of the part of the spherical shell inside \mathcal{U} is bounded from above by the volume of the spherical shell which is $4\pi x^2 dx$. Since the volume of \mathcal{U} is $\frac{4}{3}\pi R^3$ we get the claimed bound. (The exact value of $\Pr(x \leq \mathbf{x}_{i,l} < x + dx)$ is actually given in [Mat99, San76] but the above approximate bound is enough for our purposes.) \square

Proposition 11.9. $I_{x \leq 6}$ is $O\left(\frac{1}{n^3}\right)$.

Proof. Recall that (see Lemma 11.4)

$$I_{x \leq 6} = \int_{x=0}^6 \Pr(\delta_{i,j,k,l}^\circ | \mathbf{x}_{i,l} = x) \cdot \Pr(x \leq \mathbf{x}_{i,l} < x + dx).$$

By Lemma 11.7,

$$\begin{aligned} \Pr(\delta_{i,j,k,l}^\circ | \mathbf{x}_{i,l} = x) &\leq \frac{(3x+8)^2}{R^6} \cdot \Pr(p \notin \mathcal{H}(t_i, t_l, 1) | \mathbf{x}_{i,l} = x, \mathcal{L}_{i,j,k,l}^\circ)^{n-4} \\ &\leq \frac{(3x+8)^2}{R^6}. \end{aligned}$$

It thus follows from Lemma 11.8 that

$$I_{x \leq 6} \leq \int_{x=0}^6 \frac{(3x+8)^2}{R^6} \cdot \frac{3x^2}{R^3} dx = \frac{\mu^3}{n^3} \int_{x=0}^6 3x^2 (3x+8)^2 dx = O\left(\frac{1}{n^3}\right).$$

\square

B_i is entirely inside \mathcal{U}

For the integral $I_{y \geq 1}$, occlusions must be taken into account. To this aim, we bound from below the volume of $\mathcal{H}(t_i, t_l, 1) \cap \mathcal{U}$ in the following lemma.

Lemma 11.10. *When $\mathcal{L}_{i,j,k,l}^\omega$ occurs and $\mathbf{y}_i \geq 1$, the volume of $\mathcal{H}(t_i, t_l, 1) \cap \mathcal{U}$ is greater than $\frac{\pi}{12} \mathbf{x}_{i,l}$.*

Proof. Let K be the ball having diameter $p_i t_i$. Note that K and p_l are both contained in \mathcal{U} and in $\mathcal{H}(t_i, t_l, 1)$. The convex hull of p_l and K is thus contained in $\mathcal{H}(t_i, t_l, 1) \cap \mathcal{U}$, and its volume is larger than half the volume of the ball K , $\frac{\pi}{12}$, plus the volume of a cone of apex p_l , of base a disk whose boundary is a great circle of K , and of height greater than $\mathbf{x}_{i,l} - 1$. The volume of that cone is at least $\frac{1}{3} \frac{\pi}{2^2} (\mathbf{x}_{i,l} - 1) = \frac{\pi}{12} \mathbf{x}_{i,l} - \frac{\pi}{12}$. \square

We now bound the probability that a tangent line segment $t_i t_l$ is not occluded by any of the other $n - 4$ balls, given that the line segment $t_i t_l$ exists and the ball B_i is entirely contained in \mathcal{U} .

Lemma 11.11. $\Pr\left(p \notin \mathcal{H}(t_i, t_l, 1) \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i \geq 1, \mathcal{L}_{i,j,k,l}^\omega\right)^{n-4} < 55 \exp\left(-\frac{\mu x}{16}\right)$.

Proof. First notice that

$$\begin{aligned} \Pr\left(p \notin \mathcal{H}(t_i, t_l, 1) \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i \geq 1, \mathcal{L}_{i,j,k,l}^\omega\right) \\ = 1 - \frac{\text{Volume of } \mathcal{H}(t_i, t_l, 1) \cap \mathcal{U}}{\text{Volume of } \mathcal{U}} \Big|_{\mathbf{x}_{i,l}=x, \mathbf{y}_i \geq 1, \mathcal{L}_{i,j,k,l}^\omega}. \end{aligned}$$

By Lemma 11.10, the volume of $\mathcal{H}(t_i, t_l, 1) \cap \mathcal{U}$ is bounded from below by $\frac{\pi}{12} x$. Since the volume of \mathcal{U} is $\frac{4}{3} \pi R^3$, we get

$$\Pr\left(p \notin \mathcal{H}(t_i, t_l, 1) \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i \geq 1, \mathcal{L}_{i,j,k,l}^\omega\right)^{n-4} < \left(1 - \frac{x}{16R^3}\right)^{n-4}.$$

For any $0 \leq t \leq 1$, we have $(1-t) \leq e^{-t}$ thus

$$(1-t)^{n-4} \leq e^{-t(n-4)} = e^{-tn} e^{4t} \leq e^4 e^{-tn} < 55 e^{-tn}.$$

Now $0 \leq x \leq 2R$ and $R \geq 1$ since B_i is entirely inside \mathcal{U} . Thus $0 \leq \frac{x}{16R^3} \leq \frac{1}{8R^2} \leq 1$ and

$$\begin{aligned} \Pr\left(p \notin \mathcal{H}(t_i, t_l, 1) \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i \geq 1, \mathcal{L}_{i,j,k,l}^\omega\right)^{n-4} &< 55 \exp\left(-\frac{nx}{16R^3}\right) \\ &= 55 \exp\left(-\frac{\mu x}{16}\right). \end{aligned}$$

\square

The following proposition now bounds the integral $I_{y \geq 1}$.

Proposition 11.12. $I_{y \geq 1}$ is $O\left(\frac{1}{n^3}\right)$.

Proof. Recall that

$$I_{y \geq 1} = \int_{x=0}^{2R} \Pr(\delta_{i,j,k,l}^\omega \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i \geq 1) \cdot \Pr(x \leq \mathbf{x}_{i,l} < x + dx \mid \mathbf{y}_i \geq 1).$$

By Lemmas 11.7 and 11.11 we have

$$\Pr(\delta_{i,j,k,l}^\omega \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i \geq 1) \leq \frac{(3x+8)^2}{R^6} \cdot 55 \exp\left(-\frac{\mu x}{16}\right).$$

Similarly as in Lemma 11.8 we have

$$\Pr(x \leq \mathbf{x}_{i,l} < x + dx \mid \mathbf{y}_i \geq 1) \leq \frac{3x^2}{R^3} dx.$$

Thus we get

$$\begin{aligned} I_{y \geq 1} &\leq \int_{x=0}^{2R} \frac{(3x+8)^2}{R^6} \cdot 55 \exp\left(-\frac{\mu x}{16}\right) \cdot \frac{3x^2}{R^3} dx \\ &\leq \frac{\mu^3}{n^3} \int_{x=0}^{+\infty} 3x^2 (3x+8)^2 \cdot 55 \exp\left(-\frac{\mu x}{16}\right) dx. \end{aligned}$$

Changing $\frac{\mu x}{16}$ by z we get integrals of the kind

$$\int_0^{\infty} z^r \exp(-z) dz$$

which is bounded by a constant and thus $I_{y \geq 1}$ is $O\left(\frac{1}{n^3}\right)$. \square

B_i and B_l are not close to one another and B_i is partially outside \mathcal{U}

The only remaining task is to bound the integral $I_{x > 6, y < 1}$. As in the previous case, we need to bound from below the volume of $\mathcal{H}(t_i, t_l, 1) \cap \mathcal{U}$. Here, however, the tangent $t_l t_i$ can be entirely outside \mathcal{U} , so the bound of Lemma 11.10 does not apply and a more intricate proof is needed. We need to distinguish two cases depending on the distance of segment $t_l t_i$ from O , the center of \mathcal{U} .

To this aim, we introduce two new types of events. For any $s \in \mathbb{R}$, let $\mathcal{F}_{i,j,k,l}^{\omega}(s)$ (resp. $\mathcal{N}_{i,j,k,l}^{\omega}(s)$) be the event that $\mathcal{L}_{i,j,k,l}^{\omega}$ occurs and the line segment $t_l t_i$ is at distance greater (resp. less) than $R+1-s$ from O . For reasons that will become clear in the proof of Lemma 11.15, we consider $s = y^{\frac{2}{3}}$.

The next five lemmas are used to bound the first term of the integral $I_{x > 6, y < 1}$.

Lemma 11.13. *For any random point p in \mathcal{U} , $\Pr(\delta_{i,j,k,l}^{\omega} \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i)$ is equal to*

$$\begin{aligned} &\Pr\left(\mathcal{F}_{i,j,k,l}^{\omega}(y^{\frac{2}{3}}) \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i\right) \\ &\cdot \Pr\left(p \notin \mathcal{H}(t_i, t_l, 1) \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i, \mathcal{F}_{i,j,k,l}^{\omega}(y^{\frac{2}{3}})\right)^{n-4} \\ &+ \Pr\left(\mathcal{N}_{i,j,k,l}^{\omega}(y^{\frac{2}{3}}) \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i\right) \\ &\cdot \Pr\left(p \notin \mathcal{H}(t_i, t_l, 1) \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i, \mathcal{N}_{i,j,k,l}^{\omega}(y^{\frac{2}{3}})\right)^{n-4}. \end{aligned}$$

Proof. $\delta_{i,j,k,l}^{\omega}$ implies $\mathcal{L}_{i,j,k,l}^{\omega}$ which can be split into $\mathcal{F}_{i,j,k,l}^{\omega}(y^{\frac{2}{3}})$, $\mathcal{N}_{i,j,k,l}^{\omega}(y^{\frac{2}{3}})$, and the event that $\mathcal{L}_{i,j,k,l}^{\omega}$ occurs and the line segment $t_l t_i$ is at distance exactly $R+1-y^{\frac{2}{3}}$ from O . This later event occurs with probability 0, thus

$$\begin{aligned} \Pr(\delta_{i,j,k,l}^{\omega} \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i) &= \\ &\Pr(\delta_{i,j,k,l}^{\omega} \cap \mathcal{F}_{i,j,k,l}^{\omega}(y^{\frac{2}{3}}) \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i) \\ &+ \Pr(\delta_{i,j,k,l}^{\omega} \cap \mathcal{N}_{i,j,k,l}^{\omega}(y^{\frac{2}{3}}) \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i), \end{aligned}$$

which can be expanded into

$$\begin{aligned} &\Pr(\mathcal{F}_{i,j,k,l}^{\omega}(y^{\frac{2}{3}}) \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i) \\ &\cdot \Pr(\delta_{i,j,k,l}^{\omega} \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i, \mathcal{F}_{i,j,k,l}^{\omega}(y^{\frac{2}{3}})) \\ &+ \Pr(\mathcal{N}_{i,j,k,l}^{\omega}(y^{\frac{2}{3}}) \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i) \\ &\cdot \Pr(\delta_{i,j,k,l}^{\omega} \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i, \mathcal{N}_{i,j,k,l}^{\omega}(y^{\frac{2}{3}})). \end{aligned}$$

When $\mathcal{F}_{i,j,k,l}^{\omega}(y^{\frac{2}{3}})$ occurs, the probability

$$\Pr(\delta_{i,j,k,l}^{\omega} \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i, \mathcal{F}_{i,j,k,l}^{\omega}(y^{\frac{2}{3}}))$$

is the probability that the tangent is not occluded, that is, p_γ does not belong to $\mathcal{H}(t_i, t_l, 1)$ for all the $n-4$ values of $\gamma \neq i, j, k, l$. The same argument holds for $\mathcal{N}_{i,j,k,l}^\omega(y^{\frac{2}{3}})$. Since the p_γ are independent, we get the result. \square

In order to bound the two terms in Lemma 11.13,

$$\Pr\left(p \notin \mathcal{H}(t_i, t_l, 1) \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i, \mathcal{F}_{i,j,k,l}^\omega(y^{\frac{2}{3}})\right)^{n-4} \text{ and}$$

$$\Pr\left(p \notin \mathcal{H}(t_i, t_l, 1) \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i, \mathcal{N}_{i,j,k,l}^\omega(y^{\frac{2}{3}})\right)^{n-4},$$

we need to bound the volume of $\mathcal{H}(t_i, t_l, 1) \cap \mathcal{U}$ from below.

Lemma 11.14. *When $\mathbf{x}_{i,l} \geq 6$, $\mathbf{y}_l \leq \mathbf{y}_i \leq 1$, $\mathcal{L}_{i,j,k,l}^\omega$ occurs and segment $t_i t_l$ is at distance less than $R+1-s$, $0 \leq s \leq 1$, from the center of \mathcal{U} , then the volume of $\mathcal{H}(t_i, t_l, 1) \cap \mathcal{U}$ is larger than $\frac{1}{6\sqrt{2}}(\mathbf{x}_{i,l}-5)s\sqrt{s}$.*

Proof. We give here the idea of the proof; full details can be found in Appendix A. Let t be the closest point on segment $t_i t_l$ from O , and D be a unit radius disk centered at t in a plane containing O , the center of \mathcal{U} . We define a quadrilateral with vertices a, b, a', b' such that a and a' are the closest and the farthest points, respectively, in $D \cap \mathcal{U}$ from O , and b and b' are the points of intersection of ∂D and the perpendicular bisector of segment aa' (see Figure 11.2). Let v be equal to $R+1$ minus the distance from O to segment $t_i t_l$. We prove that the convex hull of a, b, a', b' and p_l , which is included in $\mathcal{H}(t_i, t_l, 1) \cap \mathcal{U}$, has volume greater than $\frac{1}{6\sqrt{2}}(\mathbf{x}_{i,l}-5)\min(2\sqrt{2}, v\sqrt{v})$. It follows that, for any $0 \leq s \leq 1$, if segment $t_i t_l$ is at distance less than $R+1-s$ from O , then $v \geq s$ and the volume of $\mathcal{H}(t_i, t_l, 1) \cap \mathcal{U}$ is greater than $\frac{1}{6\sqrt{2}}(\mathbf{x}_{i,l}-5)s\sqrt{s}$. \square

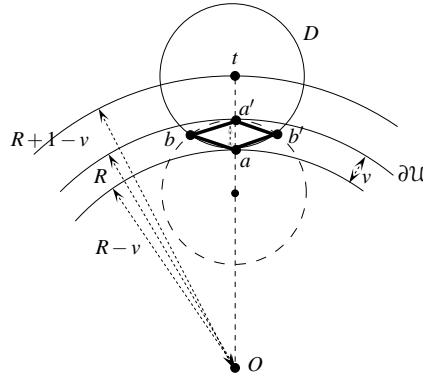


FIG. 11.2 – For the sketch of the proof of Lemma 11.14 ($v \in (0, 1)$).

Lemma 11.15. *For any random point p in \mathcal{U} , $x \geq 6$ and $0 \leq y \leq 1$,*

$$\Pr\left(p \notin \mathcal{H}(t_i, t_l, 1) \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i, \mathcal{F}_{i,j,k,l}^\omega(y^{\frac{2}{3}})\right)^{n-4}$$

$$< 55 \exp\left(-\frac{\mu(x-5)y^2}{8\sqrt{2}\pi}\right)$$

and

$$\Pr\left(p \notin \mathcal{H}(t_i, t_l, 1) \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i, \mathcal{N}_{i,j,k,l}^\omega(y^{\frac{2}{3}})\right)^{n-4}$$

$$< 55 \exp\left(-\frac{\mu(x-5)y}{8\sqrt{2}\pi}\right).$$

Furthermore, if $x \geq 6\sqrt{R}$ then

$$\Pr\left(p \notin \mathcal{H}(t_i, t_l, 1) \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i, \mathcal{N}_{i,j,k,l}^\omega(y^{\frac{2}{3}})\right)^{n-4} < 55 \exp\left(-\frac{\mu(x-5)}{8\sqrt{2}\pi}\right).$$

Proof. Let $\mathbf{x}_{i,l} = x, \mathbf{y}_i = y$ and suppose first that event $\mathcal{F}_{i,j,k,l}^\omega(y^{\frac{2}{3}})$ occurs. Since p_i is at distance $R - y$ from O , the segment $t_i t_l$ is at distance less than $R + 1 - y$ from O , and thus, by Lemma 11.14, the volume of $\mathcal{H}(t_i, t_l, 1) \cap \mathcal{U}$ is greater than $\frac{1}{6\sqrt{2}}(x-5)y\sqrt{y}$, which is bigger than $\frac{1}{6\sqrt{2}}(x-5)y^2$ since $0 \leq y \leq 1$ (we bound $y\sqrt{y}$ from below by y^2 only so that we can actually compute the integral \mathcal{J}_1 in the proof of Proposition 11.20). We now follow the proof of Lemma 11.11, except that the volume of $\mathcal{H}(t_i, t_l, 1) \cap \mathcal{U}$ is now bounded from below by $\frac{1}{6\sqrt{2}}(x-5)y^2$ instead of $\frac{\pi}{12}x$. We get

$$\Pr\left(p \notin \mathcal{H}(t_i, t_l, 1) \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i, \mathcal{F}_{i,j,k,l}^\omega(y^{\frac{2}{3}})\right)^{n-4} < 55 \exp\left(-\frac{\mu(x-5)y^2}{8\sqrt{2}\pi}\right).$$

When $\mathcal{N}_{i,j,k,l}^\omega(y^{\frac{2}{3}})$ occurs, the segment $t_i t_l$ is at distance less than $R + 1 - y^{\frac{2}{3}}$ from O , and thus, by Lemma 11.14, the volume of $\mathcal{H}(t_i, t_l, 1) \cap \mathcal{U}$ is bounded from below by $\frac{1}{6\sqrt{2}}(x-5)y^{\frac{2}{3}}\sqrt{y^{\frac{2}{3}}} = \frac{1}{6\sqrt{2}}(x-5)y$. Then, as before, we get

$$\Pr\left(p \notin \mathcal{H}(t_i, t_l, 1) \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i, \mathcal{N}_{i,j,k,l}^\omega(y^{\frac{2}{3}})\right)^{n-4} < 55 \exp\left(-\frac{\mu(x-5)y}{8\sqrt{2}\pi}\right).$$

Now, if $x \geq 6\sqrt{R}$, the length of the tangent $t_i t_l$ is at least $6\sqrt{R} - 2$. Since $x \geq 6, R > 3$ and a simple computation shows that $6\sqrt{R} - 2$ is bigger than $2\sqrt{2R+1}$ which is the length of the longest line segment that may entirely lie inside $\mathcal{U}^+ \setminus \mathcal{U}$. Thus $\text{dist}(O, t_i t_l) \leq R = R + 1 - s$ with $s = 1$ and, by Lemma 11.14, the volume of $\mathcal{H}(t_i, t_l, 1) \cap \mathcal{U}$ is greater than $\frac{1}{6\sqrt{2}}(x-5)$. Then, as before, we get

$$\Pr\left(p \notin \mathcal{H}(t_i, t_l, 1) \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i, \mathcal{N}_{i,j,k,l}^\omega(y^{\frac{2}{3}})\right)^{n-4} < 55 \exp\left(-\frac{\mu(x-5)}{8\sqrt{2}\pi}\right).$$

□

Lemma 11.16. $\Pr(\mathcal{N}_{i,j,k,l}^\omega(y^{\frac{2}{3}}) \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i) \leq \frac{(3x+8)^2}{R^6}$.

Proof. The event $\mathcal{N}_{i,j,k,l}^\omega(y^{\frac{2}{3}})$ occurs only if $\mathcal{L}_{i,j,k,l}^\omega$ occurs. The result thus follows since, by Lemmas 11.5 and 11.6, $\Pr(\mathcal{L}_{i,j,k,l}^\omega \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i) \leq \frac{(3x+8)^2}{R^6}$. □

Lemma 11.17. If $y < 1$, then

$$\Pr\left(\mathcal{F}_{i,j,k,l}^\omega(y^{\frac{2}{3}}) \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y\right) \leq 81\pi^2 \frac{(x+6)^2 y^2}{R^6}.$$

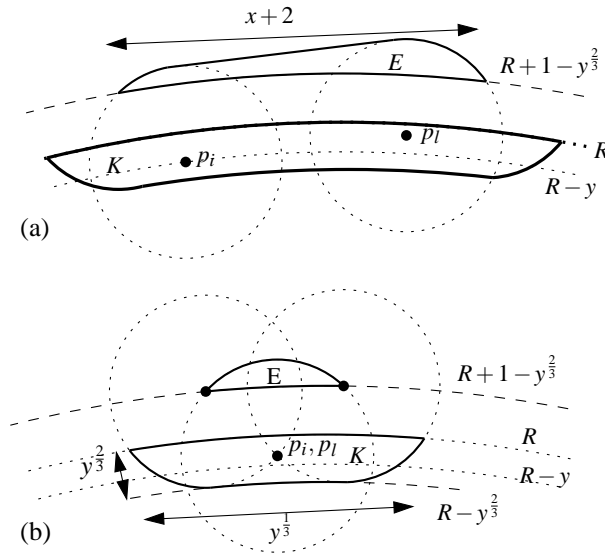


FIG. 11.3 – For the sketch of the proof of Lemma 11.17.

Proof. A “far” tangent $t_i t_l$ is at distance at least $R+1-y^{2/3}$ from the center O of \mathcal{U} . Such a segment also lies in $\mathcal{H}(p_i, p_l, 1)$. Let E be the part of $\mathcal{H}(p_i, p_l, 1)$ lying outside of the sphere of radius $R+1-y^{2/3}$ and center O . See Figure 11.3 (a). Now, both p_j and p_k must be in the region inside \mathcal{U} and within distance 1 from E . Denote this region by K . Then

$$\Pr\left(\mathcal{F}_{i,j,k,l}^{\text{co}}(y^{2/3}) \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i\right) \leq \left(\frac{\text{Volume of } K}{\text{Volume of } \mathcal{U}}\right)^2.$$

By Proposition 11.32, which we prove in Appendix B, the volume of K is bounded from above by $12\pi^2(x+6)y$, which yields the result. Here we give the intuition of the proof. Refer to Figure 11.3. First notice that the “length” of K is at most $x+4$. Since K is enclosed in between a sphere of radius R and one of radius $R-y^{2/3}$, its “height” is at most $y^{2/3}$. For the “width”, consider Figure 11.3 (b) which shows a cross-section of K taken with a plane through O and perpendicular to $p_i p_l$. The “width” of K is no more than 2 times the “width” of E . The “height” of E can be bounded by some constant times $y^{2/3}$; thus its “width” can be bounded by some constant times $\sqrt{y^{2/3}} = y^{1/3}$. Thus, intuitively, the volume of K is smaller than $(x+4)y^{2/3}y^{1/3} = (x+4)y$, up to a constant, and the result follows. \square

We now bound the two last terms of the integral $I_{x>6, y<1}$.

Lemma 11.18. $\Pr(y \leq \mathbf{y}_i < y+dy) \leq \frac{3dy}{R}$.

Proof. The event $(y \leq \mathbf{y}_i < y+dy)$ occurs only if p_i lies in the spherical shell delimited by the two spheres centered at O of radii $R-y$ and $R-y-dy$ whose volume is smaller than $4\pi R^2 dy$. Dividing by the volume of \mathcal{U} proves the result. \square

Lemma 11.19. For $6 \leq x \leq 2R$ and $y \leq 1$, we have

$$\Pr((x \leq \mathbf{x}_{i,l} < x+dx) \cap (\mathbf{y}_l \leq \mathbf{y}_i) \mid \mathbf{y}_i = y) \leq \frac{6xydx}{R^3}.$$

Proof. The probability $\Pr((x \leq \mathbf{x}_{i,l} < x+dx) \cap (\mathbf{y}_l \leq \mathbf{y}_i) \mid \mathbf{y}_i = y)$ is equal to the volume of the region (shown in grey in Figure 11.4) which is the intersection of the region in between the two spheres centered at p_i and of radii x and $x+dx$, and the region in between the two spheres centered at O and of radii R and $R-y$, divided by the volume

of \mathcal{U} . We prove in Proposition 11.37 in Appendix C that the volume of that region is at most $8\pi xy dx$. Roughly speaking, the volume bounded by the four spheres is at most $8\pi xy dx$ because, its “thickness” is dx , its “height” is y and its “radius” is x . Dividing by the volume of \mathcal{U} proves the result. \square

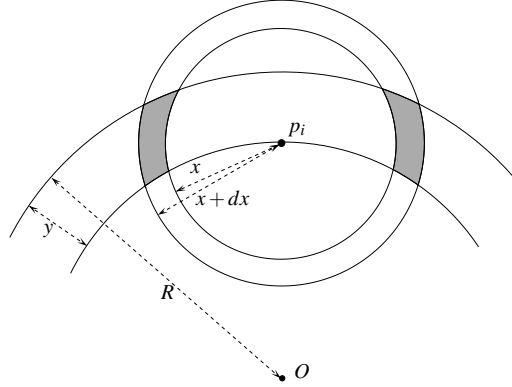


FIG. 11.4 – For the proof of Lemma 11.19.

We can now bound the integral $I_{x>6,y<1}$ of Lemma 11.4.

Proposition 11.20. $I_{x>6,y<1}$ is $O\left(\frac{1}{n^3}\right)$.

Proof. Recall that

$$\begin{aligned} I_{x>6,y<1} &= \int_{x=6}^{2R} \int_{y=0}^1 \Pr(\delta_{i,j,k,l}^\omega \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i) \\ &\quad \cdot \Pr((x \leq \mathbf{x}_{i,l} < x + dx) \cap (\mathbf{y}_l \leq \mathbf{y}_i) \mid \mathbf{y}_i = y) \\ &\quad \cdot \Pr(y \leq \mathbf{y}_i < y + dy). \end{aligned}$$

By Lemmas 11.18 and 11.19, we get

$$I_{x>6,y<1} \leq \int_{x=6}^{2R} \int_{y=0}^1 \Pr(\delta_{i,j,k,l}^\omega \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i) \cdot \frac{6xy dx}{R^3} \cdot \frac{3dy}{R}.$$

By Lemma 11.13, $\Pr(\delta_{i,j,k,l}^\omega \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i)$ is equal to

$$\begin{aligned} &\Pr\left(\mathcal{F}_{i,j,k,l}^\omega(y^{\frac{2}{3}}) \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i\right) \\ &\quad \cdot \Pr\left(p \notin \mathcal{H}(t_i, t_l, 1) \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i, \mathcal{F}_{i,j,k,l}^\omega(y^{\frac{2}{3}})\right)^{n-4} \\ &+ \Pr\left(\mathcal{N}_{i,j,k,l}^\omega(y^{\frac{2}{3}}) \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i\right) \\ &\quad \cdot \Pr\left(p \notin \mathcal{H}(t_i, t_l, 1) \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i, \mathcal{N}_{i,j,k,l}^\omega(y^{\frac{2}{3}})\right)^{n-4}. \end{aligned}$$

We split the integral at $x = 6\sqrt{R}$. When $x \geq 6\sqrt{R}$, the distance from O to the tangent $t_i t_l$ is less than R (see the proof of Lemma 11.15), which is less than $R + 1 - y^{\frac{2}{3}}$ for any y in $(0, 1)$. Thus, for any $x \geq 6\sqrt{R}$ and $y \in (0, 1)$, the probability $\Pr\left(\mathcal{F}_{i,j,k,l}^\omega(y^{\frac{2}{3}}) \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i\right)$ is equal to 0. It then follows from Lemmas 11.15, 11.16 and

11.17 that $I_{x>6,y<1} \leq J_1 + J_2 + J_3$ with

$$\begin{aligned} J_1 &= \int_{x=6}^{6\sqrt{R}} \int_{y=0}^1 81\pi^2 \frac{(x+6)^2 y^2}{R^6} \cdot 55 \exp\left(-\frac{\mu(x-5)y^2}{8\sqrt{2}\pi}\right) \cdot \frac{6xy dx}{R^3} \cdot \frac{3dy}{R}, \\ J_2 &= \int_{x=6}^{6\sqrt{R}} \int_{y=0}^1 \frac{(3x+8)^2}{R^6} \cdot 55 \exp\left(-\frac{\mu(x-5)y}{8\sqrt{2}\pi}\right) \cdot \frac{6xy dx}{R^3} \cdot \frac{3dy}{R}, \\ J_3 &= \int_{x=6\sqrt{R}}^{2R} \int_{y=0}^1 \frac{(3x+8)^2}{R^6} \cdot 55 \exp\left(-\frac{\mu(x-5)}{8\sqrt{2}\pi}\right) \cdot \frac{6xy dx}{R^3} \cdot \frac{3dy}{R}. \end{aligned}$$

Changing $\frac{\mu(x-5)}{8\sqrt{2}\pi}$ by z in the three integrals and y^2 by y' in J_1 , we get

$$\begin{aligned} J_1 &\leq \frac{K}{R^{10}} \sum_{u=0}^3 \int_{z=0}^{c\sqrt{R}} \int_{y'=0}^1 z^u y' \exp(-zy') dz dy', \\ J_2 &\leq \frac{K}{R^{10}} \sum_{u=0}^3 \int_{z=0}^{c\sqrt{R}} \int_{y=0}^1 z^u y \exp(-zy) dz dy, \\ J_3 &\leq \frac{K}{R^{10}} \sum_{u=0}^3 \int_{z=0}^{\infty} \int_{y=0}^1 z^u y \exp(-z) dz dy, \end{aligned}$$

where K and c are some positive constants.

Note first that J_3 is bounded from above by $\frac{K}{R^{10}} \sum_{u=0}^3 \int_{z=0}^{\infty} z^u \exp(-z) dz$. These integrals are bounded by a constant, thus J_3 is $O\left(\frac{1}{R^{10}}\right)$.

To bound the integrals J_1 and J_2 , we now compute the integral

$$\int_{z=0}^A \int_{y=0}^1 z^u y \exp(-zy) dz dy \quad (11.3)$$

for $u \in \{0, \dots, 3\}$ and $A > 0$, for example with Maple [Map]. For $u = 0$ it is equal to

$$\frac{\exp(-A) + A - 1}{A}. \quad (11.4)$$

For $u = 1$, the integral (11.3) is equal to

$$\exp(-A) + \ln A + Ei(1, A) + \gamma - 1 \quad (11.5)$$

where $Ei(1, A)$ denotes the exponential integral $\int_{t=1}^{\infty} \frac{\exp(-At)}{t} dt$ and γ denotes Euler's constant. Finally, for $u = 2$ or 3, the integral (11.3) is equal to

$$\exp(-A) P_1(A, u-1) + P_2(A, u-1) \quad (11.6)$$

where $P_i(A, u-1)$ denotes a polynomial of degree $u-1$ in A .

When A tends to ∞ , (11.4) tends to 1, (11.5) is equivalent to $\ln A$ (since $Ei(1, A)$ tends to 0) and (11.6) is equivalent to the leading monomial of $P_2(A, u-1)$ which is of degree $u-1 \leq 2$. This guarantees that for $A = c\sqrt{R}$ and $u \in \{0, \dots, 3\}$, the integral (11.3) is $O(R)$. It follows that J_1 and J_2 are $O\left(\frac{1}{R^9}\right)$.

Since $R^3 = n/\mu$, we get that $I_{x>6,y<1} \leq J_1 + J_2 + J_3 = O\left(\frac{1}{R^9}\right) = O\left(\frac{1}{n^3}\right)$. \square

We can now conclude the proof that the expected number of $T4$ -segments is $O(n)$, because, by Lemmas 11.3, 11.4, and Propositions 11.9, 11.12, and 11.20, the expected number of $T4$ -segments is smaller than

$$\sum_{(i,j,k,l) \in \mathcal{N}} \sum_{\omega=1}^{12} \left(O\left(\frac{1}{n^3}\right) + O\left(\frac{1}{n^3}\right) + O\left(\frac{1}{n^3}\right) \right) = O(n).$$

11.4 The expected number of $T4$ -segments is at least linear

In this section, we prove that the expected number of $T4$ -segments amongst n uniformly distributed unit balls is $\Omega(n)$. To do this, we bound from below the probability that four given balls have a given $T4$ -segment. The key step is to give a condition on the relative positions of four unit balls that guarantees that they have exactly twelve common tangent lines. We use here the notation as defined in Section 11.3.1.

Lemma 11.21. *Let e be a real number satisfying $\frac{4\sqrt{2}}{3} < e < 2$ and let the radius R of \mathcal{U} be strictly greater than e . There exists an $\varepsilon > 0$ such that for any point $p \in \mathcal{U}$, there exist three balls $\Gamma_1(p)$, $\Gamma_2(p)$, $\Gamma_3(p)$ of radius ε contained in \mathcal{U} and satisfying the following conditions :*

- p and the centers of the $\Gamma_i(p)$ form a regular tetrahedron with edges of length e , and
- for any triple of points (p_1, p_2, p_3) , p_i taken from $\Gamma_i(p)$, the four unit balls centered at p , p_1 , p_2 and p_3 have exactly 12 distinct tangent lines.

Proof. Macdonald, Pach and Theobald proved [MPT01, Lemma 3] that 4 unit balls centered on the vertices of a regular tetrahedron with edges of length e , $\frac{4\sqrt{2}}{3} < e < 2$, have exactly 12 distinct real common tangent lines. Moreover, these 12 tangent lines correspond to the 12 real roots of a system of equations of degree 12, thus each tangent line corresponds to a *simple* root of that system of equations. It thus follows that for any sufficiently small perturbation of the 4 ball centers, the 4 perturbed balls still have 12 real common tangent lines. Let $\varepsilon > 0$ be such that the 4 ball centers can move distance ε in any direction while keeping 12 distinct common tangents.

Now, for any point $p \in \mathcal{U}$, consider a regular tetrahedron with edge length e having p as a vertex and such that the other vertices are at distance at least ε from the boundary of \mathcal{U} ; for example, we can choose the other three vertices on a plane perpendicular to the segment Op . Let $\Gamma_1(p)$, $\Gamma_2(p)$, and $\Gamma_3(p)$ be the balls of radius ε centered at the vertices, distinct from p , of that tetrahedron. By the previous reasoning, for any $q \in \Gamma_1(p)$, $r \in \Gamma_2(p)$, and $s \in \Gamma_3(p)$, the four unit balls centered at p , q , r and s have exactly twelve tangents. \square

Now, by Lemma 11.3, the expected number of $T4$ -segments is

$$\sum_{(i,j,k,l) \in \mathcal{N}^{\omega=1}} \sum_{\omega=1}^{12} \Pr(\delta_{i,j,k,l}^{\omega}).$$

Thus we only need to bound from below the probability that the event $\delta_{i,j,k,l}^{\omega}$ occurs.

Lemma 11.22. $\Pr(\delta_{i,j,k,l}^{\omega})$ is $\Omega\left(\frac{1}{n^3}\right)$.

Proof. Assume that $n > 8\mu$ so that the radius $R = \sqrt[3]{n/\mu}$ of \mathcal{U} is larger than 2 and let $T(p)$ be the set $\Gamma_1(p) \times \Gamma_2(p) \times \Gamma_3(p)$ where $\Gamma_i(p)$ and e are defined as in Lemma 11.21. First, note that

$$\begin{aligned} \Pr(\delta_{i,j,k,l}^{\omega}) &\geq \Pr(\delta_{i,j,k,l}^{\omega} \cap (p_i, p_j, p_k) \in T(p_l)) \\ &= \Pr((p_i, p_j, p_k) \in T(p_l)) \cdot \Pr(\delta_{i,j,k,l}^{\omega} \mid (p_i, p_j, p_k) \in T(p_l)). \end{aligned}$$

Since $\Gamma_1(p_l)$, $\Gamma_2(p_l)$, and $\Gamma_3(p_l)$ are three balls of radius ε entirely contained in \mathcal{U} , we have

$$\Pr((p_i, p_j, p_k) \in T(p_l)) = \left(\frac{\frac{4}{3}\pi\varepsilon^3}{\frac{4}{3}\pi R^3}\right)^3 = \frac{\mu^3\varepsilon^9}{n^3}.$$

By Lemmas 11.5 and 11.21, the event $(\delta_{i,j,k,l}^{\omega} \mid (p_i, p_j, p_k) \in T(p_l))$ occurs if and only if the interior of $\mathcal{H}(t_i, t_l, 1) \cap \mathcal{U}$ does not contain the center of any ball. Note that the volume of $\mathcal{H}(t_i, t_l, 1) \cap \mathcal{U}$ is at most the volume of $\mathcal{H}(t_i, t_l, 1)$, which is at most $\frac{4}{3}\pi + \pi(2 + e + 2\varepsilon)$ since the length of $t_i t_l$ is at most $e + 2 + 2\varepsilon$. It follows that

$$\Pr(\delta_{i,j,k,l}^{\omega} \mid (p_i, p_j, p_k) \in T(p_l)) \geq \left(1 - \frac{\pi(\frac{4}{3} + 2 + e + 2\varepsilon)}{\text{Volume}(\mathcal{U})}\right)^{n-4}.$$

Since $e < 2$, we get, after some elementary calculations, that

$$\Pr(\delta_{i,j,k,l}^0 \mid (p_i, p_j, p_k) \in T(p_l)) \geq \left(1 - \frac{(6+2\varepsilon)\mu}{n}\right)^{n-4}. \quad (11.7)$$

We thus have

$$\Pr(\delta_{i,j,k,l}^\omega) \geq \frac{\mu^3 \varepsilon^9}{n^3} \left(1 - \frac{(6+2\varepsilon)\mu}{n}\right)^{n-4}.$$

Since $\left(1 - \frac{(6+2\varepsilon)\mu}{n}\right)^{n-4}$ tends to $e^{-(6+2\varepsilon)\mu}$ when n tends to infinity, we get

$$\Pr(\delta_{i,j,k,l}^\omega) = \Omega\left(\frac{1}{n^3}\right).$$

□

This completes the proof of the lower bound of Theorem 11.1 since the expected number of $T4$ -segments amongst n uniformly distributed unit balls is, by Lemmas 11.3 and 11.22,

$$\sum_{(i,j,k,l) \in \mathcal{N}^{\omega=1}} \sum_{\omega=1}^{12} \Pr(\delta_{i,j,k,l}^\omega) = \sum_{(i,j,k,l) \in \mathcal{N}^{\omega=1}} \sum_{\omega=1}^{12} \Omega\left(\frac{1}{n^3}\right) = \Omega(n).$$

11.5 The expected size of the visibility complex is linear

In this section we prove Theorem 11.2, that the expected size of the visibility complex of a set of n uniformly distributed unit balls is linear.

We say that the balls are in general position if any k -dimensional face of the visibility complex is a connected set of maximal free segments tangent to exactly $4 - k$ balls. We can assume that the balls are in general position since this occurs with probability 1. We give a bound on the expected number of k -faces, for $k = 0, \dots, 4$.

Lemma 11.23. *The expected number of 0-faces is $\Theta(n)$.*

Proof. A 0-face of the visibility complex is a maximal free line segment tangent to 4 balls. Each maximal free line segment tangent to 4 balls contains a $T4$ -segment and each $T4$ -segment is contained in one maximal free line segment. Thus, by Theorem 11.1, the expected number of 0-faces is linear. □

To deal with the faces of dimension $k \geq 1$, we divide them into two classes. A k -face is *open* if it is incident to at least one $(k-1)$ -face, otherwise it is *closed*. When the balls are in general position, the number of k -faces incident to a particular $(k-1)$ -face is constant. In the proof of the following lemmas, any constant can be used. However, for completeness, we will use the exact values, but without justifying them.

Lemma 11.24. *The expected number of 1-faces is $\Theta(n)$.*

Proof. Note that a 0-face corresponds to a maximal free segment tangent to 4 balls and it is incident to those 1-faces corresponding to free segments tangent to 3 amongst those 4 balls. So, a 0-face is incident to exactly six 1-faces, which implies that the number of open 1-faces is 6 times the number of 0-faces, and is thus $\Theta(n)$ by the previous lemma.

Proving that the expected number of closed 1-faces is $O(n)$ can be done in a way very similar to the proof of the upper bound in Theorem 11.1. The difference is that we consider now only three balls and thus in all proofs, we forget ball B_k . We have to consider only $\binom{n}{3}$ triples of balls instead of $\binom{n}{4}$ quadruples, but we remove from the integral the probability $\Pr(p_k \in \mathcal{H}(p_i, p_l, 2) \mid \mathbf{x}_{i,l} = x) \leq \frac{3x+8}{R^3}$. Since $\frac{n}{R^3} = \mu$, this amounts to dividing the terms over which we integrate by $\mu(3x+8)$ which does not change the general shape of the integrals (a polynomial multiplied by an exponential) which are convergent. Notice that B_i, B_j, B_l and ω now define a set of segments $t_i t_l$, rather than just a single segment. However, those segments define a closed 1-face only if none of them is occluded by one of the $n-3$ remaining balls. Any particular choice of a tangent $t_i t_l$ in the 1-face will give a relevant cylinder $\mathcal{H}(t_i, t_l, 1)$ to use in the proofs. □

Lemma 11.25. *The expected number of 2-faces is $\Theta(n)$.*

Proof. Since a 1-face has five incident 2-faces, the tight linear bound on the number of 1-faces gives a tight linear bound on the number of open 2-faces. The closed case is solved similarly to the proof of the upper bound in Theorem 11.1. We now consider $\binom{n}{2}$ pairs of balls B_i, B_l and we remove from the integrals the probability $\Pr(p_j, p_k \in \mathcal{H}(p_i, p_l, 2) | \mathbf{x}_{i,l} = x) \leq \left(\frac{3x+8}{R^3}\right)^2$ which gives an $O(n)$ bound on the number of closed 2-faces. \square

Lemma 11.26. *The expected numbers of 3-faces and 4-faces are $\Theta(n)$.*

Proof. A 3-face, corresponding to lines tangent to a ball, can only be closed if $n = 1$. The number of open 3-faces is linear by the fact that in general position a 2-face is incident to four 3-faces. The number of 4-faces is linear since a 3-face is incident to three 4-faces. \square

11.6 Worst-case lower bound

We provide here a $\Omega(n^2)$ lower bound on the number of k -faces in the visibility complex. Recall that for the case of n arbitrarily sized balls, Devillers and Ramos [DR01] presented a simple $\Omega(n^3)$ lower bound on the number of free segments tangent to 4 balls, which is also the number of vertices in the visibility complex. Their lower bound (see Figure 11.5) consists of (i) $\frac{n}{3}$ balls such that the view from the origin consists of $\frac{n}{3}$ disjoint disks centered on a circle, (ii) $\frac{n}{3}$ balls such that the view from the origin consists of $\frac{n}{3}$ disks whose boundaries are concentric circles intersecting (in projection) all the disks of (i), and (iii) $\frac{n}{3}$ tiny balls centered around the origin such that from any point on these $\frac{n}{3}$ tiny balls the view of the balls in (i) and (ii) is topologically invariant. Note that finding a $\Omega(n^3)$ lower bound on the number of free segments tangent to 4 balls, amongst n balls of bounded radii, is to the best of our knowledge, open.

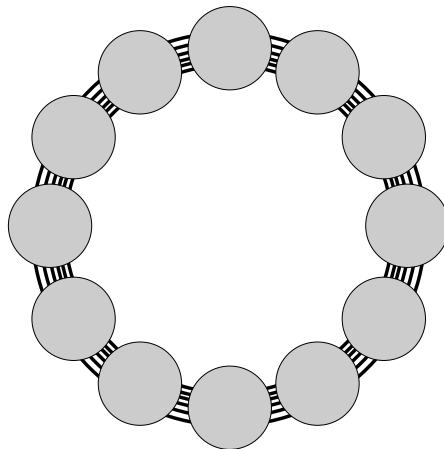


FIG. 11.5 – Quadratic view from the origin [DR01].

Proposition 11.27. *The number of k -faces in the visibility complex of n disjoint unit balls in \mathbb{R}^3 is $\Omega(n^2)$ for all k between 0 and 4.*

Proof. We first observe that the size of the visibility complex of n unit balls can trivially be quadratic by having the balls sparsely distributed in the space such that any pair of balls defines a closed 2-face.

Getting a quadratic number of free lines tangent to four balls amongst a set of n unit balls can be done by taking balls B_i centered at $(2i, 0, 0)$ for $1 \leq i \leq \frac{n}{2}$ and balls B'_j centered at $(2j, 10, 0)$ for $1 \leq j \leq \frac{n}{2}$. Then, for

any i and j , the line through the points $(2i + 1, 0, 1)$ and $(2j + 1, 10, 1)$ is free and can be moved down so that it comes into contact with the four balls B_i, B_{i+1}, B'_j and B'_{j+1} . This argument proves that the number of k -faces, for $0 \leq k \leq 2$, can be quadratic.

The free segment $(2i, 1, 0)(2j, 9, 0)$ belongs to the 4-face consisting of maximal free segments with endpoints on B_i and B'_j . Thus there is a quadratic number of 4-faces. The bound also applies to 3-faces by considering lines tangent to B_i and stabbing B'_j .

In the above construction, the balls can be pushed together (they will intersect) so that they fit inside a spherical universe of radius $\sqrt[3]{n/\mu}$ without changing the result. Note also that the above construction can be slightly perturbed to obtain the same result for a set of n unit balls, disjoint or not, with no 4 centers coplanar. \square

11.7 Generalizations

In this section we provide several generalizations of our results.

11.7.1 Poisson distribution

Consider a set of unit balls whose centers are drawn by a 3-D Poisson point process of parameter μ in the universe \mathcal{U} . By a *Poisson point process of parameter μ in \mathcal{U}* [GS92], we mean that we generate X random points inside \mathcal{U} so that

$$Pr(X = k) = \frac{(\mu \cdot \text{Volume}(\mathcal{U}))^k \cdot \exp(-\mu \cdot \text{Volume}(\mathcal{U}))}{k!} \quad (11.8)$$

and for any disjoint subsets M and M' of \mathcal{U} , the number of the points inside M and the number of points inside M' are independent random variables. Note that Equation (11.8) yields that the expected number of points inside \mathcal{U} is $\mu \cdot \text{Volume}(\mathcal{U}) = \frac{4\pi}{3}n$.

The following simple argument shows that our results extend to this distribution. Let X be the random variable representing the number of centers of unit balls generated by a Poisson point process with parameter μ in \mathcal{U} , and let Y be the random variable representing the number of $T4$ -segments amongst those balls. The expected number of $T4$ -segments is

$$E(Y) = \sum_{k=0}^{\infty} E(Y|X = k) \cdot Pr(X = k).$$

Theorem 11.1 gives $E(Y|X = k) = \Theta(k)$ and

$$Pr(X = k) = \frac{(\frac{4}{3}\pi n)^k \cdot \exp(-\frac{4}{3}\pi n)}{k!}.$$

Thus

$$\begin{aligned} E(Y) &= \Theta\left(\frac{4}{3}\pi n \exp(-\frac{4}{3}\pi n) \sum_{k=1}^{\infty} \frac{(\frac{4}{3}\pi n)^{k-1}}{(k-1)!}\right) \\ &= \Theta(n \exp(-\frac{4}{3}\pi n) \exp(\frac{4}{3}\pi n)) = \Theta(n). \end{aligned}$$

Therefore the expected number of $T4$ -segments amongst n balls whose centers are generated by a Poisson point process with parameter μ in \mathcal{U} is $\Theta(n)$. Similarly this bound extends to the expected size of the visibility complex.

We now investigate various models in which we change the shape of the universe or the nature of the objects.

11.7.2 Smooth convex universe

Our results can be generalized to the case where the universe is no longer a ball, but a homothet of a smooth convex set with homothety factor proportional to $\sqrt[3]{n}$. This can be achieved by considering the radius of curvature of the boundary of the universe, instead of R , in the proofs of the lemmas dealing with tangents outside the universe.

11.7.3 Other objects

Let r_{min} and r_{max} be two strictly positive real constants. In the following, we bound the expected number of $T4$ -segments amongst balls whose radii vary in the interval $[r_{min}, r_{max}]$, amongst polyhedra each enclosed between two concentric balls of radii r_{min} and r_{max} , and amongst polygons each enclosed between two concentric circles of radii r_{min} and r_{max} . The centers of the concentric balls or circles are called the centers of the polyhedra or polygons, respectively. In each case a $T4$ -segment is called *outer* if the centers of the two extremal objects it is tangent to are farther apart than $6r_{max}$ and are both at distance less than $2r_{max}$ from the boundary of \mathcal{U} . Otherwise the $T4$ -segment is called *inner*.

For these models, the proof of the $\Omega(n)$ lower bound on the expected number of $T4$ -segments (Section 11.4) generalizes directly because, for the kind of objects we consider, there always exist placements of four of them such that they admit at least one common tangent line with multiplicity one.

Balls of various radii

We have considered a model where all the balls have the same radius. If we allow the radii to vary in the interval $[r_{min}, r_{max}]$, then the proof of the linear upper bound on the expected number of inner $T4$ -segments generalizes almost immediately by considering the volumes $\mathcal{H}(p_i, p_l, 2r_{max})$ and $\mathcal{H}(t_i, t_l, r_{min})$ instead of $\mathcal{H}(p_i, p_l, 2)$ and $\mathcal{H}(t_i, t_l, 1)$.

Section 11.3.2 generalizes immediately to prove that the expected number of $T4$ -segments tangent to four balls B_i, B_j, B_k and B_l in that order such that p_i and p_l are closer to one another than $6r_{max}$ is $O(n)$. The only difficult task for extending Section 11.3.2 is the proof of the following analog of Lemma 11.10.

Lemma 11.28. *When $\mathbf{x}_{i,l} \geq 6r_{max}$, $\mathbf{y}_i \geq 2r_{max}$ and $\mathcal{L}_{i,j,k,l}^\omega$ occurs, the volume of $\mathcal{H}(t_i, t_l, r_{min}) \cap \mathcal{U}$ is greater than $\frac{\pi}{24} r_{min}^2 (\mathbf{x}_{i,l} - 6r_{max})$.*

Proof. The proof is similar to the proof of Lemma 11.10. Refer to Figure 11.6. Let m be the midpoint of segment $t_i t_l$ and K be the sphere of diameter r_{min} centered on the point c lying on segment $t_i p_i$ at distance $\frac{1}{2}r_{min}$ from t_i . The sphere K is entirely inside $\mathcal{H}(t_i, t_l, r_{min}) \cap \mathcal{U}$, m lies in $\mathcal{H}(t_i, t_l, r_{min})$ and a straightforward computation shows that m is in \mathcal{U} since t_i is in \mathcal{U} at distance at least r_{max} from its boundary and t_l is at distance at most r_{max} from \mathcal{U} . Thus $\mathcal{H}(t_i, t_l, r_{min}) \cap \mathcal{U}$ contains the convex hull of K and m which contains the cone of apex m , of base a disk whose boundary is a great circle of K , and of height the distance from m to the center c of K . Now

$$\begin{aligned} \mathbf{x}_{i,l} = |p_i p_l| &\leq |p_i c| + |cm| + |m t_l| + |t_l p_l| \\ &\leq r_{max} + |cm| + \frac{1}{2}|t_i t_l| + r_{max} \\ &\leq 2r_{max} + |cm| + \frac{1}{2}(\mathbf{x}_{i,l} + 2r_{max}). \end{aligned}$$

Thus $|cm| \geq \frac{1}{2}\mathbf{x}_{i,l} - 3r_{max}$ and the volume of the cone is at least $\frac{1}{3}\pi(\frac{r_{min}}{2})^2(\frac{1}{2}\mathbf{x}_{i,l} - 3r_{max}) = \frac{\pi}{24}r_{min}^2(\mathbf{x}_{i,l} - 6r_{max})$. \square

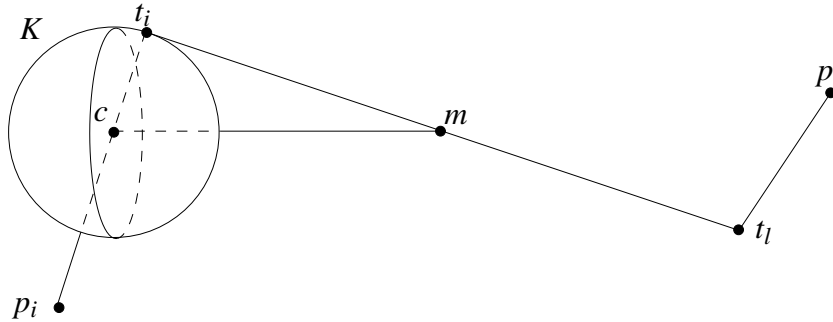


FIG. 11.6 – For the proof of Lemma 11.28.

The rest of Section 11.3.2 generalizes easily for proving that the expected number of $T4$ -segments tangent to four balls B_i, B_j, B_k and B_l in that order such that p_i and p_l are farther apart than $6r_{max}$ and p_i is farther than $2r_{max}$ from the boundary of \mathcal{U} , is $O(n)$. Hence the expected number of inner $T4$ -segments is $O(n)$.

Our proof cannot be extended to provide a linear upper bound on the expected number of outer $T4$ -segments. This is because, if balls B_i and B_l are of radius r_{max} then a line segment $t_i t_l$ tangent to B_i and B_l might be outside \mathcal{U} and at distance greater than r_{min} from its boundary. Then $\mathcal{H}(t_i, t_l, r_{min})$ does not intersect \mathcal{U} and we cannot bound $\mathcal{H}(t_i, t_l, r_{min}) \cap \mathcal{U}$ from below by a positive constant as in Lemma 11.14, which is crucial for the proof of Lemma 11.15 and thus for Proposition 11.20.

However, by not taking into account the occlusion in the proof of Proposition 11.20, we get that the expected number of outer $T4$ -segments is $O(n^2)$. Refer to the proof of Proposition 11.20 and consider $I_{x>6r_{max}, y<2r_{max}}$, the analog of $I_{x>6, y<1}$ for this case. The analogs of Lemmas 11.6 and 11.7 yield that

$$\Pr(\delta_{i,j,k,l}^{00} \mid \mathbf{x}_{i,l} = x, \mathbf{y}_i = y, \mathbf{y}_l \leq \mathbf{y}_i) \leq \frac{(3xr_{max}^2 + 8r_{max}^3)^2}{R^6}.$$

Lemma 11.18 still holds and we can easily prove the analog of Lemma 11.19. Both results imply that

$$\begin{aligned} I_{x>6r_{max}, y<2r_{max}} &\leq \int_{x=6r_{max}}^{2R} \int_{y=0}^{2r_{max}} \frac{(3xr_{max}^2 + 8r_{max}^3)^2}{R^6} \cdot \frac{6xy dx}{R^3} \cdot \frac{3 dy}{R} \\ &\in O\left(\frac{1}{R^6}\right) = O\left(\frac{1}{n^2}\right). \end{aligned}$$

Hence the expected number of inner $T4$ -segments is $O(n)$ and the expected number of outer $T4$ -segments is $O(n^2)$. This still improves the result of Durand et al. [DDP02] who proved a bound of $O(n^{8/3})$ for the same model.

In this section we have assumed that the sphere centers are uniformly distributed but we have made no assumption on the distribution of the radii of the spheres in the interval $[r_{min}, r_{max}]$, which are thus assumed to be worst case. The addition of some hypothesis on the radii distribution may yield better results on the number of outer $T4$ -segments.

Polyhedra of bounded aspect ratio

Consider polyhedra of constant complexity, each enclosed between two concentric balls of radii r_{min} and r_{max} whose centers are uniformly distributed in \mathcal{U} . In such a case, as for balls of various radii, the $O(n)$ bound on the expected number of inner $T4$ -segments immediately applies as well as the $O(n^2)$ bound on the expected number of outer $T4$ -segments.

Polygons of bounded aspect ratio

Our proof technique can also be generalized to non-fat 3D objects such as polygons. Consider polygons of constant complexity enclosed between two coplanar concentric circles of radii r_{min} and r_{max} , and whose centers and normals are independently chosen from the uniform distributions over \mathbb{R}^3 and \mathbb{S}^2 . Let T_1, \dots, T_n be such polygons with respective normals $\mathbf{n}_1, \dots, \mathbf{n}_n$ and centers p_1, \dots, p_n .

Four polygons T_i, T_j, T_k and T_l have a common tangent line that meet them in that order only if p_j and p_k lie in $\mathcal{H}(p_i, p_l, 2r_{max})$. This implies, as in Section 11.3.2, that the expected number of $T4$ -segments tangent to four polygons T_i, T_j, T_k and T_l in that order such that p_i and p_l are closer to one another than some constant, say $6r_{max}$, is $O(n)$.

When such a tangent, denoted $t_i t_l$, exists, it is not occluded only if, for any $\gamma \neq i, j, k, l$, point p_γ does not lie in the interior of $\mathcal{H}(t_i, t_l, r_{min} \cos \theta_\gamma)$ where θ_γ denotes the angle between \mathbf{n}_γ and the supporting line of $t_i t_l$ (see Figure 11.7 (a) and Lemma 11.5). Let γ be an integer distinct from i, j, k and l . By the Total Probability Theorem, the probability that T_γ does not occlude the tangent line segment $t_i t_l$ is bounded from above by

$$\int_{\theta=0}^{\pi/2} \Pr(p_\gamma \notin \mathcal{H}(t_i, t_l, r_{min} \cos \theta_\gamma) \mid \theta_\gamma = \theta) \cdot \Pr(\theta \leq \theta_\gamma < \theta + d\theta).$$

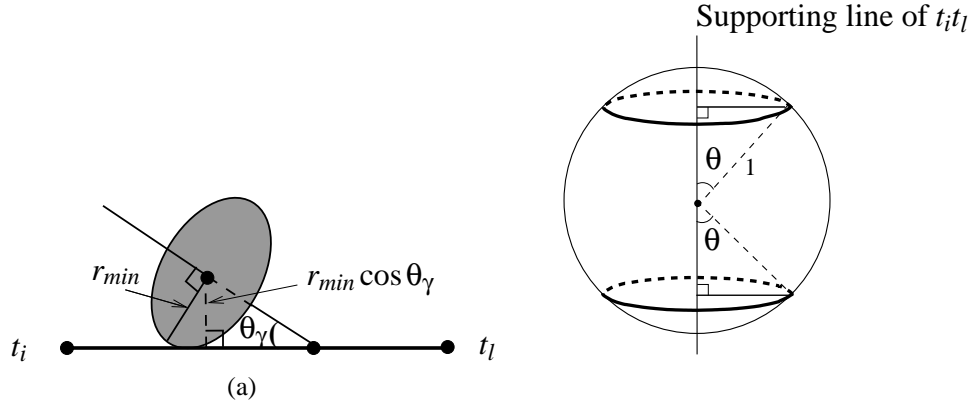


FIG. 11.7 – Illustration for the case of polygons of bounded aspect ratio.

Similarly as in Lemma 11.28, when the tangent $t_i t_l$ exists, $\mathbf{x}_{i,l} \geq 6r_{max}$ and $\mathbf{y}_i \geq 2r_{max}$, the volume of $\mathcal{H}(t_i, t_l, r_{min}) \cap \mathcal{U}$ is greater than $\frac{\pi}{24} (r_{min} \cos \theta_\gamma)^2 (\mathbf{x}_{i,l} - 6r_{max})$. Thus

$$\Pr(p_\gamma \notin \mathcal{H}(t_i, t_l, r_{min} \cos \theta_\gamma) \mid \theta_\gamma = \theta) \leq 1 - \frac{(r_{min} \cos \theta_\gamma)^2 (\mathbf{x}_{i,l} - 6r_{max})}{32R^3}.$$

The probability that θ_γ is in between θ and $\theta + d\theta$ is $\sin \theta d\theta$, which corresponds to twice the area of the spherical shell between the latitudes θ and $\theta + d\theta$ on the unit sphere, divided by the area of the unit sphere (see Figure 11.7 (b)). Thus when p_i is at distance greater than $6r_{max}$ from p_l and at distance greater than $2r_{max}$ from the boundary of \mathcal{U} , the probability that T_γ does not occlude the tangent line segment $t_i t_l$ is bounded from above by

$$\int_{\theta=0}^{\pi/2} \left(1 - \frac{(r_{min} \cos \theta_\gamma)^2 (\mathbf{x}_{i,l} - 6r_{max})}{32R^3} \right) \sin \theta d\theta = 1 - \frac{r_{min}^2 (\mathbf{x}_{i,l} - 6r_{max})}{96R^3}.$$

Then, similarly as in Lemma 11.11, the probability that the tangent line segment $t_i t_l$ is not occluded, when p_i is at distance greater than $6r_{max}$ from p_l and at distance greater than $2r_{max}$ from the boundary of \mathcal{U} , is at most

$$55 \exp\left(-\frac{\mu r_{min}^2 (\mathbf{x}_{i,l} - 6r_{max})}{96}\right).$$

We thus get the analog of Proposition 11.12 for the model considered here which implies that the expected number of $T4$ -segments tangent to four polygons T_i, T_j, T_k and T_l in that order such that p_i and p_l are farther apart than $6r_{max}$ and p_i is farther than $2r_{max}$ from the boundary of \mathcal{U} is $O(n)$.

We thus get that the expected number of inner $T4$ -segments is $O(n)$. Moreover, as for balls of various radii, the expected number of outer $T4$ -segments is $O(n^2)$.

11.8 Conclusion

In this paper, we proved that the expected number of $T4$ -segments amongst n uniformly distributed unit balls in \mathbb{R}^3 is $\Theta(n)$. We also proved that the expected size of the visibility complex of n uniformly distributed unit balls is $\Theta(n)$. Equivalently the expected number of combinatorially different visibility events amongst n uniformly distributed unit balls is $\Theta(n)$. We then proved that $\Theta(n)$ also bounds the expected number of $T4$ -segments occurring not too close to the boundary of the universe for various other models such as n uniformly distributed polyhedra, or polygons, of bounded aspect ratio and constant complexity. For these models, we also provided a $O(n^2)$ bound on the expected number of all the $T4$ -segments.

This paper is an attempt to analyze the average-case behavior of the size of visibility structures. The distribution models of scene objects investigated here are theoretical in nature since objects in graphics scenes are seldom distributed uniformly or by a Poisson process. However, our results are important in a context where there are

few rigorous results either theoretical or experimental. They provide theoretical ground to support the empirical evidence indicating that the worst-case upper bound on the number of visibility events is largely pessimistic in practical situations. As a consequence, there is reason to believe that an output-sensitive algorithm for computing all visibility events may work in practice.

Practitioners will be concerned about the size of the constant hidden in the Θ notation. We have calculated (in the proofs of Section 11.3) this constant to be no larger than $2^{16}\mu^3 + 2^{31}\mu + 2^{37}e^{-\mu/3}(\mu^2 + 1/\mu^2)$. Of course this is shocking. We suppose that the constant is actually much smaller. However estimating it in practice is a difficult problem which is still to be solved. After solving this problem, an interesting experiment will be to compare the number of visibility events occurring in a realistic graphic scene with the theoretical bound for uniformly distributed objects.

The results proved here also provide new insight on the complexity of other visibility structures. Consider for instance the aspect graph, a partition of viewpoint space into maximal connected regions by surfaces along which visibility events are observed. As explained in [Pet95], the complexity of the aspect graph is dominated by δ^m , where δ is the degree of the surface corresponding to lines “tangent” to 3 objects and m the dimension of the viewpoint space. For a scene composed of n disjoint spheres, δ is trivially $O(n^3)$, so the aspect graph has $O(n^6)$ orthographic views and $O(n^9)$ perspective views. However the results of this paper show that the expected value of δ is $\Theta(n)$ since the expected number of families of lines tangent to three objects (related to the 1-faces of the visibility complex) is linear and the degree of each family is bounded. It would thus be interesting to get a good bound on the expected value of δ^2 and δ^3 which is related to bounding the expected value of the square and the cube of the number of combinatorially different visibility events. Note that the former would also give the standard deviation of the expected number of combinatorially different visibility events. Similar observations hold for the polyhedral case.

11.9 Appendix A. Volume of the intersection of a 3D hippodrome with a ball

Recall that \mathcal{U} is a ball of radius R centered at O . Let B_i and B_l be two unit balls whose centers p_i and p_l are in \mathcal{U} , within distance 1 from its boundary, and distance $x \geq 6$ apart. Let $t_i t_l$ be a line segment tangent to B_i and B_l at its endpoints. The section is devoted to the proof of the following proposition which leads directly to Lemma 11.14.

Proposition 11.29. *For any $0 \leq s \leq 1$ such that segment $t_i t_l$ is at distance less than $R + 1 - s$ from O , the volume of $\mathcal{H}(t_i, t_l, 1) \cap \mathcal{U}$ is larger than $\frac{1}{6\sqrt{2}}(x - 5)s\sqrt{s}$.*

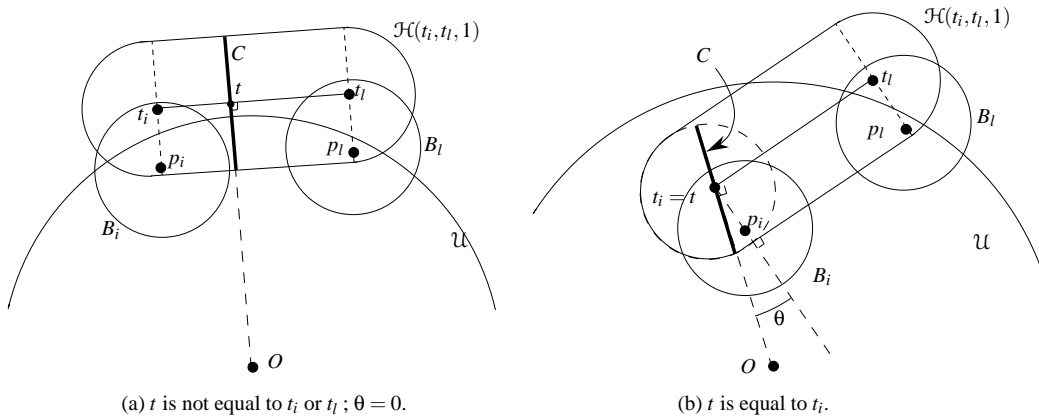
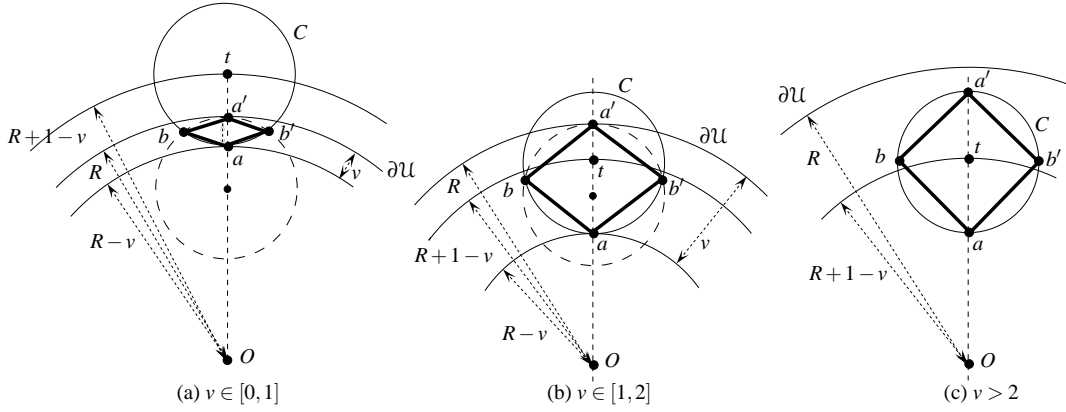


FIG. 11.8 – For the definition of t and C (C is shown from the side view).

We proceed as follows. Let v be such that the distance from O to the segment $t_i t_l$ is $R + 1 - v$, and let t be the point on segment $t_i t_l$ closest to O (see Figure 11.8). Assume without loss of generality that t is closer to t_i than to

FIG. 11.9 – For the definition of a, a', b, b' .

t_l . Let C (resp. D) be the unit radius circle (resp. disk) centered at t in the plane, denoted \mathcal{P} , containing the vectors $\vec{O}t$ and the cross product of $\vec{O}t$ and $\vec{t}_i t_l$. Let θ be the angle between the plane orthogonal to $t_i t_l$ and \mathcal{P} . We first prove the following lemma.

Lemma 11.30. *The volume of $\mathcal{H}(t_i, t_l, 1) \cap \mathcal{U}$ is greater than*

$$\frac{1}{3} \min\left(2, \frac{v\sqrt{v}}{\sqrt{2}}\right) \cdot \min\left(\frac{x-2}{2}, (x-2)\cos\theta - 1\right).$$

Proof. Let a denote the closest point on C from O , a' the farthest point in $D \cap \mathcal{U}$ from O , and b and b' the two points of intersection of C and the perpendicular bisector of segment aa' (see Figure 11.9).

The volume of $\mathcal{H}(t_i, t_l, 1) \cap \mathcal{U}$ is greater than the volume of the convex hull of a, b, a', b' and p_l because $\mathcal{H}(t_i, t_l, 1) \cap \mathcal{U}$ is convex and contains these five points. The volume of this polyhedron is equal to one third of the area of its base, the quadrilateral with vertices a, b, a', b' , times its height, the distance from p_l to the plane \mathcal{P} containing a, b, a', b' .

We first compute a lower bound on the area of the quadrilateral with vertices a, b, a', b' . If $v \leq 2$ (see Figure 11.9 (a) and (b)), then the length of aa' is equal to v , and a simple calculation gives that the length of bb' is equal to $2\sqrt{v - \frac{v^2}{4}} \geq \sqrt{2v}$. Thus the area of the quadrilateral a, b, a', b' is greater than $\frac{v\sqrt{v}}{\sqrt{2}}$. If $v > 2$ (see Figure 11.9 (c)), then C is entirely contained in \mathcal{U} and the area of the quadrilateral a, b, a', b' is equal to 2. Thus, the area of the quadrilateral is at least $\min(2, \frac{v\sqrt{v}}{\sqrt{2}})$.

The volume of the polyhedron is thus greater than $\frac{1}{3} \min(2, \frac{v\sqrt{v}}{\sqrt{2}})$ times the distance from p_l to the plane \mathcal{P} . We consider two cases.

First, suppose that t belongs to the interior of the segment $t_i t_l$ (see Figure 11.8 (a)). Then, the height is equal to the distance from t_l to t since p_l and a, b, a', b' belong, respectively, to the two planes, orthogonal to $t_i t_l$ and passing through t_l and t , respectively. Since t_i and t_l belong to B_i and B_l , they are at least distance $x - 2$ apart, thus t and t_l are at least distance $\frac{x-2}{2}$ apart. Thus, the height from p_l to \mathcal{P} is at least $\frac{x-2}{2}$.

Second, suppose that $t = t_i$ (see Figure 11.8 (b)); $t \neq t_l$ since we assumed that t is closer to t_i than to t_l . Refer to Figure 11.10. Let A and B be the orthogonal projections of p_l and t_l onto \mathcal{P} , respectively. Note that the lengths of Ap_l and Bt_l are the distances from p_l and t_l to the plane \mathcal{P} , respectively.

Considering the triangle $\triangle Ap_l t_l$ and that the distance between t_l and p_l is 1, we obtain that $|Ap_l| \geq |At_l| - |t_l p_l| = |At_l| - 1$, where $|ab|$ denotes the length of segment ab . Since $A \in \mathcal{P}$ and the length of Bt_l is the distance from t_l to the plane \mathcal{P} , the length of At_l is greater than that of Bt_l , thus $|Ap_l| \geq |Bt_l| - 1$.

To bound the length of Bt_l , we now consider the triangle $\triangle Bt_l t$. The angle $\angle Bt_l t$ is the angle between the normal of the plane \mathcal{P} and $t_i t_l$, that is, by definition, θ . So the length of Bt_l is the length of $t_i t_l$ times $\cos\theta$ and, since $|t_i t_l|$ is at least $x - 2$, $|Bt_l|$ is greater than $(x - 2)\cos\theta$. Thus the length of Ap_l is greater than $(x - 2)\cos\theta - 1$.

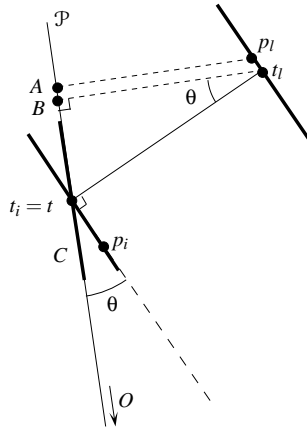


FIG. 11.10 – The height from p_l to \mathcal{P} is greater than $|t_i t_l| \cos \theta - 1$.

Hence the distance from p_l to the plane \mathcal{P} is greater than $\min(\frac{x-2}{2}, (x-2) \cos \theta - 1)$ and thus the volume of $\mathcal{H}(t_i, t_l, 1) \cap \mathcal{U}$ is greater than $\frac{1}{3} \min(2, \frac{v\sqrt{v}}{\sqrt{2}}) \cdot \min(\frac{x-2}{2}, (x-2) \cos \theta - 1)$. \square

The following lemma bounds $\cos \theta$.

Lemma 11.31. *The angle θ is such that $\cos \theta \geq \frac{\sqrt{7}}{4}$.*

Proof. Note first that this lemma is intuitively obvious. Indeed (see Figure 11.8 (b)), if x is sufficiently large and if t_i is the closest point on segment $t_i t_l$ to O , then the angle between the plane supporting C and the segment $t_i t_l$ is necessarily close to $\pi/2$, which implies that θ is close to 0. We now prove the lemma.

Refer to Figure 11.10 and consider the triangle $\triangle Ot_i t_l$. Let $|ab|$ denote the length of segment ab . Then the law of cosines yields

$$\begin{aligned} |Ot_l|^2 &= |Ot_i|^2 + |t_i t_l|^2 - 2 \cdot |Ot_i| \cdot |t_i t_l| \cdot \cos(\frac{\pi}{2} + \theta) \\ &= |Ot_i|^2 + |t_i t_l|^2 + 2 \cdot |Ot_i| \cdot |t_i t_l| \cdot \sin \theta \end{aligned}$$

which gives that

$$\sin \theta = \frac{|Ot_l|^2 - |Ot_i|^2 - |t_i t_l|^2}{2 \cdot |Ot_i| \cdot |t_i t_l|}.$$

The centers p_i and p_l of balls B_i and B_l are distance $x \geq 6$ apart and at distance less than 1 from the boundary of \mathcal{U} , so $|t_i t_l| \geq 4$, $|Ot_i| \geq R - 2$ and $|Ot_l| \leq R + 1$. Hence

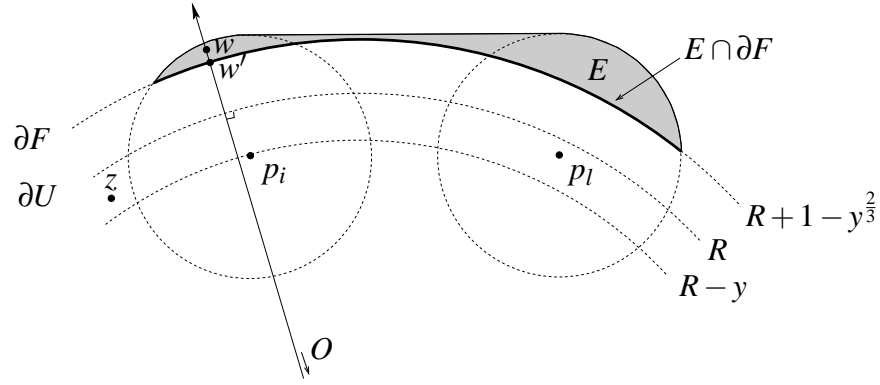
$$\sin \theta \leq \frac{(R+1)^2 - (R-2)^2 - 4^2}{2 \cdot (R-2) \cdot 4} \leq \frac{6(R-2)}{8(R-2)} = \frac{3}{4}.$$

Using $\cos \theta = \sqrt{1 - (\sin \theta)^2}$ proves that $\cos \theta \geq \frac{\sqrt{7}}{4}$. \square

We can now conclude the proof of Proposition 11.29. For any $0 \leq s \leq 1$, if segment $t_i t_l$ is at distance $R + 1 - v \leq R + 1 - s$ from the center of \mathcal{U} , then $v \geq s$. By Lemma 11.31, $(x-2) \cos \theta - 1 \geq \frac{x-5}{2}$ which means that $\min(\frac{x-2}{2}, (x-2) \cos \theta - 1) \geq \frac{x-5}{2}$. Thus Lemma 11.30 gives that the volume of $\mathcal{H}(t_i, t_l, 1) \cap \mathcal{U}$ is greater than $\frac{1}{6\sqrt{2}} (x-5) \min(2\sqrt{2}, v\sqrt{v}) \geq \frac{1}{6\sqrt{2}} (x-5) \min(2\sqrt{2}, s\sqrt{s}) = \frac{1}{6\sqrt{2}} s \sqrt{s} (x-5)$ since $s \leq 1$. Hence the volume of $\mathcal{H}(t_i, t_l, 1) \cap \mathcal{U}$ is greater than $\frac{1}{6\sqrt{2}} s \sqrt{s} (x-5)$.

11.10 Appendix B. Volume of K

Recall that \mathcal{U} is a ball of radius R centered at O and let p_i and p_l be two points in \mathcal{U} within distance 1 of its boundary and distance x apart. Let y be a real number such that $0 \leq y < 1$. Let F be the open ball with center O and

FIG. 11.11 – The part E of $\mathcal{H}(p_i, p_l, 1)$ outside F .

radius $R+1-y^{\frac{2}{3}}$ and ∂F its frontier. Let E be the part of $\mathcal{H}(p_i, p_l, 1)$ that is outside F and K be the intersection of \mathcal{U} with the union of all unit balls centered on points in E (see Figure 11.3). This section is devoted to the proof of the following proposition used in the proof of Lemma 11.17.

Proposition 11.32. *The volume of K is bounded from above by $12\pi^2(x+6)y$.*

Lemma 11.33. *If $z \in \mathcal{U}$ is at distance less than 1 from E , then z is at distance less than 1 from $E \cap \partial F$.*

Proof. Let $z \in \mathcal{U}$ and $w \in E$ be two points at distance less than 1 and refer to Figure 11.11. Let w' be the point of intersection of ∂F and the ray from O through w . For any ball B centered in \mathcal{U} , $B \setminus F$ lies in the cone of center O and base $B \cap \partial F$. Thus $E = \mathcal{H}(p_i, p_l, 1) \setminus F$ lies in the cone of center O and base $E \cap \partial F$. Hence the ray from O through w lies in this cone and $w' \in E \cap \partial F$. On the other hand, $|zw'| \leq |zw|$ since $z \in F$, $w' \in \partial F$ and w lies outside F on the ray from O through w' . Thus, since $w' \in E \cap \partial F$ and $|zw| < 1$ by hypothesis, the distance from z to $E \cap \partial F$ is less than 1. \square

The above lemma implies that K is the intersection of \mathcal{U} with the union of all unit balls centered on $E \cap \partial F$. To bound the volume of K , we enclose $E \cap \partial F$ in a subset of ∂F that will be easier to deal with.

Let $B(p)$ denote the ball of unit radius centered at p . Let $\pi(p)$ be the point that maximizes (under inclusion) the intersection $\partial F \cap B(q)$ for all q on the ray from O through p . A simple computation yields that the distance between $\pi(p)$ and O is

$$R_y = \sqrt{(R+1-y^{\frac{2}{3}})^2 - 1}.$$

Thus π is the orthogonal projection onto the sphere centered at O of radius R_y . Now let $\pi'(p)$ be the point that maximizes (under inclusion) the intersection $\partial F \cap B(q)$ for all q on the radius of \mathcal{U} through p (that is the part inside \mathcal{U} of the ray from O through p). Similarly, π' is the orthogonal projection onto the sphere centered at O of radius

$$R' = \min(R, R_y).$$

Let G be the union of the spherical caps $\partial F \cap B(\pi'(p))$ for all p on the segment from p_i to p_l (see Figure 11.12). Let H denote the points of \mathcal{U} at distance less than or equal to 1 from G (see Figure 11.13).

Lemma 11.34. $K \subseteq H$.

Proof. $E \cap \partial F$ is the union of $\partial F \cap B(p)$ for all p on the segment $p_i p_l$. Furthermore, for any such p , $\partial F \cap B(p) \subseteq \partial F \cap B(\pi'(p))$ by definition of π' since $p \in \mathcal{U}$. Thus $E \cap \partial F$ is contained in G .

By Lemma 11.33, K is the intersection of \mathcal{U} with the union of all unit balls centered on $E \cap \partial F$. Thus K is contained in H , the union of all unit balls centered in G . \square

To bound the volume of H from above, we first bound the area of its section by planes Π that contain O and are orthogonal to the plane, denoted (O, p_i, p_l) , containing O , p_i and p_l (see Figures 11.13 and 11.14).

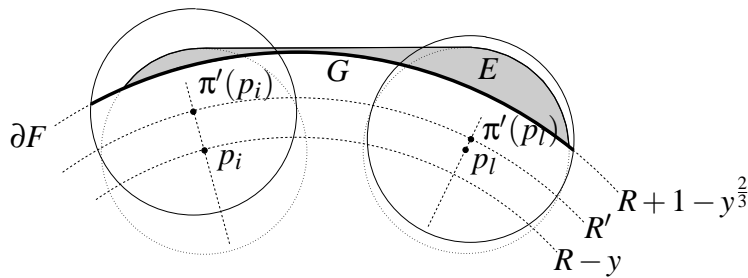


FIG. 11.12 – G , a part of ∂F enclosing $E \cap \partial F$.

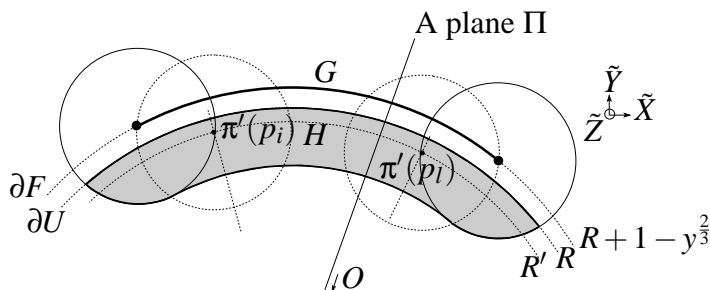


FIG. 11.13 – The region H and a plane Π .

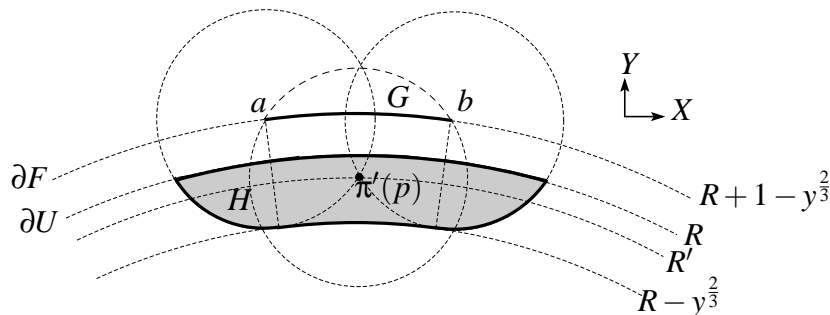
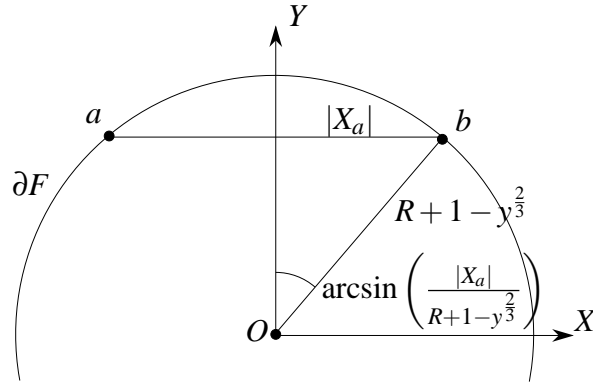


FIG. 11.14 – Section of H by a plane Π intersecting segment $p_i p_l$ at p .

Lemma 11.35. *The area of $\Pi \cap H$ is less than $12\pi y$.*

Proof. The section of G by a plane Π is a circular arc on ∂F . If Π intersects the segment $p_i p_l$, let p denote the point of intersection, then the circular arc is the intersection of ∂F and the disk $B(\pi'(p)) \cap \Pi$ (refer to Figures 11.12 and 11.14). Otherwise, the circular arc is the intersection of ∂F and the disk $B(\pi'(p_i)) \cap \Pi$ or $B(\pi'(p_l)) \cap \Pi$ (see Figure 11.12). The disk has radius 1 in the former case and radius less than one in the latter case. In both cases the center of the disk is at distance R' from O . Thus the length of the circular arc $G \cap \Pi$ is maximal if and only if Π intersects the segment $p_i p_l$. Thus the area of $\Pi \cap H$ is maximal if and only if Π intersects the segment $p_i p_l$. Hence we can assume that Π is such a plane. Let p denote its intersection with segment $p_i p_l$.

Let a and b denote the endpoints of $G \cap \Pi$ and refer to Figure 11.14. Points a and b are the intersection of ∂F and the circle in Π of radius 1 centered at $\pi'(p)$. The lines (Oa) and (Ob) split $\Pi \cap H$ into three parts, a left, a central and a right part. Symmetries with respect to the lines (Oa) and (Ob) send the left and right parts into the central one. Hence, the area of $\Pi \cap H$ is bounded by 3 times the area of its central part. This part is delimited by the two rays from O through a and b , and the two circles in Π with center O and radii R and $R - y^{2/3}$. So, if α denotes

FIG. 11.15 – The length of the circular arcs ab .

the length of the circular arc ab , the area A of the central part is

$$A = \frac{\alpha}{2\pi(R + 1 - y^{2/3})} \cdot \pi(R^2 - (R - y^{2/3})^2) = \alpha \frac{2Ry^{2/3} - y^{4/3}}{2(R + 1 - y^{2/3})} \leq \alpha y^{2/3}.$$

We now bound the length α of the arc ab . We choose an orthonormal frame $(\pi'(p), X, Y)$ in Π such that O has coordinates $(0, -R')$ (see Figure 11.14). Recall that a is one of the intersection points of the circle centered at $\pi'(p)$ of radius 1 and the circle centered at O of radius $R + 1 - y^{2/3}$. A simple computation yields that the coordinates (X_a, Y_a) of a are equal to

$$Y_a = \frac{(R + 1 - y^{2/3})^2 - 1 - R'^2}{2R'}, \quad |X_a| = \sqrt{1 - Y_a^2}.$$

If $R' = R$, then

$$Y_a = \frac{y^{4/3} + 2R - 2Ry^{2/3} - 2y^{2/3}}{2R} = 1 - y^{2/3} \left(1 + \frac{2 - y^{2/3}}{2R}\right) \geq 1 - 2y^{2/3}$$

which implies that

$$|X_a| \leq \sqrt{1 - (1 - 2y^{2/3})^2} = \sqrt{4y^{2/3} - 4y^{4/3}} \leq 2y^{1/3}.$$

Now if $R' \neq R$, then $(R + 1 - y^{2/3})^2 - 1 \leq R^2$ by definition. Expanding this inequality yields

$$\begin{aligned} y^{4/3} + 2R - 2Ry^{2/3} - 2y^{2/3} &\leq 0, \\ y^{2/3} &\geq \frac{y^{4/3} + 2R}{2(R + 1)} \geq \frac{R}{R + 1} \geq \frac{1}{2}. \end{aligned}$$

Thus $\sqrt{2y^{2/3}} \geq 1$ and since $|X_a| = \sqrt{1 - Y_a^2} \leq 1$ we get $|X_a| \leq \sqrt{2}y^{1/3}$. Hence, in both cases,

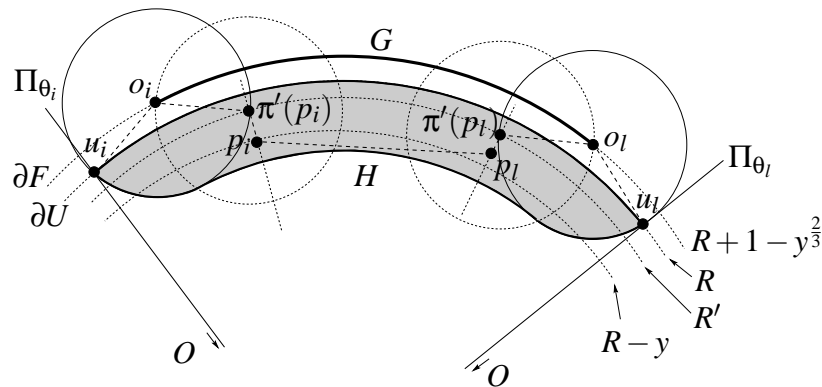
$$|X_a| \leq 2y^{1/3}.$$

Thus the length of the circular arc ab is (see Figure 11.15)

$$\alpha = (R + 1 - y^{2/3}) \cdot 2 \arcsin\left(\frac{|X_a|}{R + 1 - y^{2/3}}\right) \leq (R + 1 - y^{2/3}) \cdot 2 \arcsin\left(\frac{2y^{1/3}}{R + 1 - y^{2/3}}\right).$$

A straightforward computation shows that $\arcsin(x) - \pi x \leq 0$ for any $x \in [0, 1]$. Thus

$$\alpha \leq (R + 1 - y^{2/3}) \cdot 2\pi \frac{2y^{1/3}}{R + 1 - y^{2/3}} = 4\pi y^{1/3}.$$


 FIG. 11.16 – For the computing a bound on $\Delta\theta$.

Since the area A of the middle part is less than or equal to $\alpha y^{\frac{2}{3}}$,

$$A \leq 4\pi y^{\frac{1}{3}} y^{\frac{2}{3}} = 4\pi y.$$

This implies that the area of $\Pi \cap H$ is less than or equal to $12\pi y$. \square

Lemma 11.36. *The volume of H is bounded from above by $12\pi^2(x+6)y$.*

Proof. We express the volume of H by an integral using spherical coordinates (r, θ, ϕ) in an orthogonal frame $(O, \tilde{X}, \tilde{Y}, \tilde{Z})$ such that the plane $(O, \tilde{X}, \tilde{Y})$ contains p_i and p_l (see Figure 11.13). A plane $\theta = \text{constant}$ contains the \tilde{Z} -axis and thus is a plane Π . Let $1_H(r, \theta, \phi)$ denote the indicator function of H ; $1_H(r, \theta, \phi)$ is equal to 1 if the point of coordinates (r, θ, ϕ) belongs to H and to 0 otherwise. Then

$$\text{Volume of } H = \int_{\phi} \int_r \int_{\theta} 1_H(r, \theta, \phi) \cdot r^2 \sin \phi \, dr \, d\theta \, d\phi.$$

Since H is inside \mathcal{U} , $r \cdot 1_H(r, \theta, \phi) \leq R \cdot 1_H(r, \theta, \phi)$. Moreover $\sin \phi \leq 1$, thus

$$\text{Volume of } H \leq R \int_{\theta} \left(\int_{\phi} \int_r 1_H(r, \theta, \phi) \cdot r \, dr \, d\phi \right) d\theta.$$

The double integral in parentheses is equal to the area of the section of H by a plane $\Pi_{\theta} : \theta = \text{constant}$. By Lemma 11.35, this area is less than $12\pi y$, which is independent of θ . Moreover the area is equal to 0 when Π_{θ} does not intersect H . Let $\Delta\theta$ denote the angle between the two extreme planes Π_{θ} that intersect H . Thus we have

$$\text{Volume of } H \leq R \cdot 12\pi y \cdot \Delta\theta.$$

We now bound $\Delta\theta$. Refer to Figure 11.16. Let Π_{θ_i} and Π_{θ_l} be the two extreme planes that intersect H . Let u_i and u_l be the two points of intersection of H with Π_{θ_i} and Π_{θ_l} , respectively; u_i and u_l lie on $\partial\mathcal{U}$. Let o_i and o_l be the two points in G at distance 1 from u_i and u_l , respectively. $\pi'(p_i)$ and $\pi'(p_l)$ are at distance 1 from o_i and o_l , respectively.

The angle between the two extreme planes Π_{θ_i} and Π_{θ_l} is, as before,

$$\Delta\theta = 2 \arcsin \frac{|u_i u_l|/2}{R} \leq 2\pi \frac{|u_i u_l|/2}{R} = \pi \frac{|u_i u_l|}{R}.$$

Now we bound $|u_i u_l|$ by the length of the polygonal line shown in Figure 11.16.

$$\begin{aligned} |u_i u_l| &\leq |u_i o_i| + |o_i \pi'(p_i)| + |\pi'(p_i) p_i| + |p_i p_l| + |p_l \pi'(p_l)| + |\pi'(p_l) o_l| + |o_l u_l| \\ &= 1 + 1 + |\pi'(p_i) p_i| + x + |p_l \pi'(p_l)| + 1 + 1. \end{aligned}$$

We show that $|\pi'(p_i)p_i|$ and $|p_i\pi'(p_i)|$ are less than 1. $\pi'(p_i)$ is inside \mathcal{U} at distance less than 1 from ∂F which lies outside \mathcal{U} . Thus $\pi'(p_i)$ is inside \mathcal{U} at distance less than 1 from its frontier. Point p_i is also inside \mathcal{U} at distance less than 1 from its frontier. Since p_i and $\pi'(p_i)$ are on the same ray starting from O , they are at distance less than 1 apart. Similarly for $\pi'(p_l)$ and p_l . Hence

$$\Delta\theta \leq \pi \frac{|u_i u_l|}{R} \leq \pi \frac{x+6}{R}.$$

Therefore

$$\text{Volume of } H \leq R \cdot 12\pi y \cdot \Delta\theta \leq 12\pi^2 y(x+6).$$

□

Proposition 11.32 follows from Lemmas 11.34 and 11.36.

11.11 Appendix C. Volume of the intersection of two spherical shells

We prove in this section the following proposition used in the proof of Lemma 11.19.

Proposition 11.37. *Let $R > 0$, $x \in [6, 2R]$, $y \in [0, 1]$ and p be a point at distance $R - y$ from O . The volume of the intersection of the region in between the two spheres centered at p and of radii x and $x + dx$, and the region in between the two spheres centered at O and of radii R and $R - y$ (see Figure 11.17) is bounded from above by $8\pi xy dx$.*

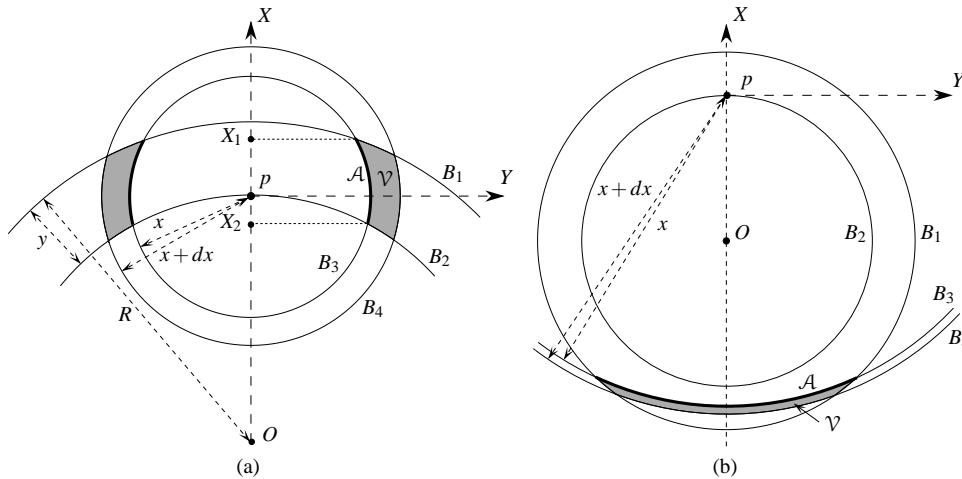


FIG. 11.17 – For the proof of Proposition 11.37.

Proof. Define the balls B_1 with center O and radius R , B_2 with center O and radius $R - y$, B_3 with center p and radius x and finally B_4 with center p and radius $x + dx$. Let \mathcal{V} denote the intersection of $(B_1 \setminus B_2)$ and $(B_4 \setminus B_3)$. We prove that the volume of \mathcal{V} is less than $8\pi xy dx$.

Since dx is infinitesimally small, the volume of \mathcal{V} is $\mathcal{A}dx$ where \mathcal{A} is the area of the intersection of the sphere ∂B_3 with $B_1 \setminus B_2$.

Let (p, X, Y, Z) be an orthogonal reference frame whose center is p and whose X -axis is oriented along \overrightarrow{Op} (see Figure 11.17). Notice that all spheres are centered on that axis. Let C_1 (resp. C_2, C_3) denote the circle that is

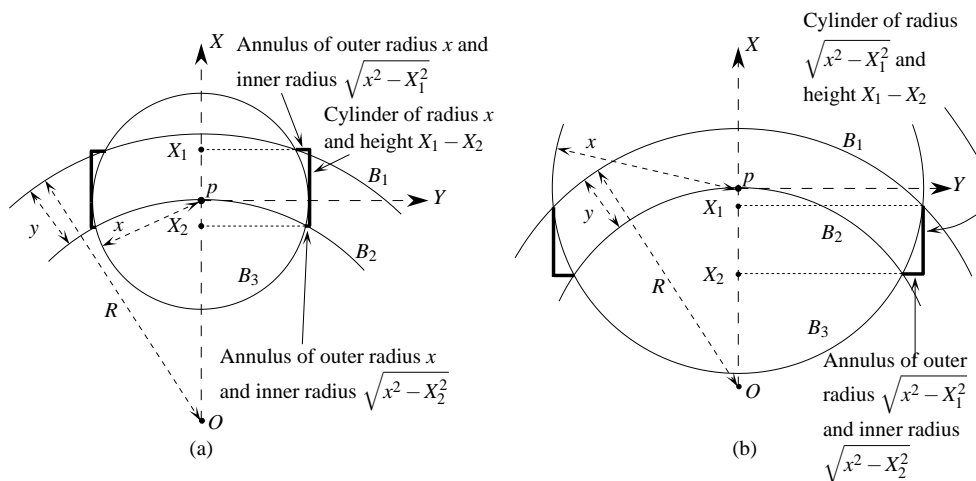


FIG. 11.18 – For the proof of Proposition 11.37, case 1.

the boundary of the intersection of B_1 (resp. B_2, B_3) and the plane (p, X, Y) in which Figure 11.17 is drawn. The equations of these circles are, in the frame (p, X, Y) ,

$$\begin{aligned} C_1 : (X + R - y)^2 + Y^2 &= R^2, \\ C_2 : (X + R - y)^2 + Y^2 &= (R - y)^2, \\ C_3 : X^2 + Y^2 &= x^2. \end{aligned}$$

Since C_3 is centered at a point on C_2 and has radius $x \geq 6 > 1 \geq y$, C_3 intersects or encloses C_1 and C_2 . In fact, C_3 intersects or encloses C_1 and C_2 in one of the three following ways.

Case 1 : If $6 \leq x \leq 2R - 2y$ then C_3 intersects both C_1 and C_2 (see Figure 11.17 (a)).

Case 2 : If $2R - 2y < x \leq 2R - y$ then C_3 intersects C_1 and encloses C_2 (see Figure 11.17 (b)).

Case 3 : If $2R - y < x$ then C_3 encloses both C_1 and C_2 . In that case, \mathcal{V} is empty and the volume is 0.

In the first case, let X_1 (resp. X_2) be the abscissa of the points of intersection of circles C_1 (resp. C_2) and C_3 . Note that $X_1 \geq X_2$ and their values can be computed directly from the equations of the circles C_1, C_2 and C_3 :

$$X_1 = \frac{R^2 - x^2 - (R - y)^2}{2R - 2y}, \quad X_2 = \frac{-x^2}{2R - 2y}.$$

Using the fact that $y \leq x \leq 2R - 2y$ we get

$$\begin{aligned} X_1 - X_2 &= \frac{y(2R - y)}{2R - 2y} = y \left(1 + \frac{y}{2R - 2y} \right) \leq 2y, \\ -X_1 - X_2 &= \frac{2x^2 - y(2R - y)}{2R - 2y} \leq 2x \frac{x}{2R - 2y} \leq 2x. \end{aligned}$$

We now bound from above the area \mathcal{A} of the surface $\partial B_3 \cap (B_1 \setminus B_2)$ by the area of a larger surface which depends on the sign of X_1 . If $X_1 \geq 0$, the surface consists of a cylinder of axis the X -axis, of radius x and height $X_1 - X_2$, and of two annuli in the planes $X = X_1$ and $X = X_2$, of inner radius $\sqrt{x^2 - X_1^2}$ and $\sqrt{x^2 - X_2^2}$, respectively, and outer radius x (see Figure 11.18 (a)). If $X_1 \leq 0$, the surface consists of a cylinder of axis the X -axis, of radius $\sqrt{x^2 - X_1^2}$ and height $X_1 - X_2$, and of an annulus in the plane $X = X_2$, of inner radius $\sqrt{x^2 - X_2^2}$ and outer radius $\sqrt{x^2 - X_1^2}$ (see Figure 11.18 (b)). In both cases that surface is larger than $\partial B_3 \cap (B_1 \setminus B_2)$ by convexity.

If $X_1 \geq 0$, the area of the cylinder is $2\pi x(X_1 - X_2) \leq 4\pi xy$ and the area of the annuli are $\pi x^2 - \pi(x^2 - X_i^2) = \pi X_i^2$, $i = 1, 2$. Since $X_1 \geq 0$, $X_1 \leq y \leq x$ and thus $\pi X_1^2 \leq \pi xy$. We also have from the expression of X_1 that $R^2 - x^2 - (R -$

$y)^2 \geq 0$ and thus $x^2 \leq y(2R - y)$. Thus

$$\pi X_2^2 = \pi \frac{x^2}{2R - 2y} \frac{x^2}{2R - 2y} \leq \pi x \frac{x}{2R - 2y} y \frac{2R - y}{2R - 2y} = \pi xy \frac{x}{2R - 2y} \left(1 + \frac{y}{2R - 2y}\right).$$

It thus follows from $y \leq x \leq 2R - 2y$ that $\pi X_2^2 \leq 2\pi xy$. Hence $\mathcal{A} \leq 7\pi xy$.

If $X_1 \leq 0$, the area of the cylinder is $2\pi\sqrt{x^2 - X_1^2}(X_1 - X_2) \leq 2\pi x(2y)$ and the area of the annulus is $\pi(x^2 - X_1^2) - \pi(x^2 - X_2^2) = \pi(X_1 - X_2)(-X_2 - X_1) \leq 4\pi xy$. Thus $\mathcal{A} \leq 8\pi xy$.

Consider now the second case $2R - 2y < x \leq 2R - y$ (see Figure 11.17 (b)). For a fixed value of y , \mathcal{A} is the area of a spherical cap whose perimeter and curvature decreases as x increases. Thus \mathcal{A} is a decreasing function of x . Since the bound $\mathcal{A} \leq 8\pi xy$ is valid for $x = 2R - 2y$ and $8\pi xy$ is an increasing function of x , $\mathcal{A} \leq 8\pi xy$ for any $x \geq 2R - 2y$. \square

Acknowledgments

The authors would like to thank Helmut Alt who led us to the lower bound proof of Section 11.4 and Luc Devroye for useful discussions.

Chapitre 12

An upper bound on the average size of silhouettes

Cet article a été soumis à *Discrete Computational Geometry* [GL06]. (Voir également le rapport de recherche [GL07].) Une version préliminaire a été publiée dans les proceedings du *22th ACM Annual Symposium on Computational Geometry* [Gli06].

Abstract

It is a widely observed phenomenon in computer graphics that the size of the silhouette of a polyhedron is much smaller than the size of the whole polyhedron. This paper provides, for the first time, theoretical evidence supporting this for a large class of objects, namely for polyhedra that approximate surfaces in some reasonable way; the surfaces may be non-convex and non-differentiable and they may have boundaries. We prove that such polyhedra have silhouettes of expected size $O(\sqrt{n})$ where the average is taken over all points of view and n is the complexity of the polyhedron.

12.1 Introduction

The silhouette of a polyhedron with respect to a given viewpoint is, roughly speaking, the set of edges incident to a front and a back face. Silhouettes arise in various problems in computer graphics such as hidden surface removal and shadow computations (see [Dug04, DD02, EGHZ00] for some recent references) and algorithms to compute them efficiently have been well-studied (see the survey by Isenberg et al. [IFH⁺03]). They are important in shape recognition; Sander et al. [SGG⁺00] claim that the silhouette “is one of the strongest visual cues of the shape of an object”.

It is a widely accepted fact that the silhouette of a polyhedron is usually much smaller than the whole polyhedron. Sander et al. [SGG⁺00], for instance, state the largely repeated claim that the silhouette of a mesh is often of size $\Theta(\sqrt{n})$ where n is the number of faces of the mesh. An experimental study by Kettner and Welzl [KW97] confirms this for a set of realistic objects. This experimental study was extended by McGuire [McG04] to a larger database of larger objects for which the observed size of the silhouette is approximately $n^{0.8}$.

There are few theoretical results supporting these observations. Kettner and Welzl [KW97] prove that a convex polyhedron that approximates a sphere with Hausdorff distance ϵ has $\Theta(1/\epsilon)$ edges, and a random orthographic projection of such a polytope has $\Theta(1/\sqrt{\epsilon})$ silhouette edges. Alt et al. [AGG03] give conditions under which it can be proved that the average silhouette of a *convex* polyhedron has size $O(\sqrt{n})$ and give additional conditions under which the worst-case size is provably sub-linear.

The goal of this paper is to study the average silhouette size of *non-convex* polyhedra. Convexity is a very strong assumption, which was crucial in the previous theoretical results. Here, rather, we assume that the polyhedron is a good approximation of some fixed (not necessarily convex) surface. Notice that it is very difficult to guarantee anything on the *worst-case* complexity of the silhouette of a polyhedron unless it approximates a strictly

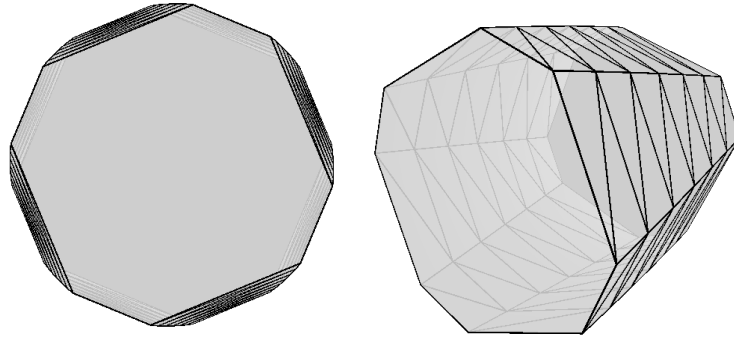


FIG. 12.1 – A worst-case linear silhouette (left) of a polyhedron approximating a cylinder.

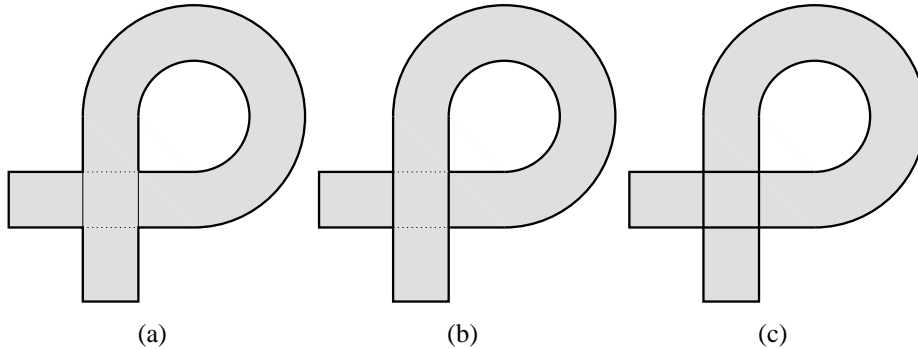


FIG. 12.2 – Three different notions of silhouettes : (a) outline of a solid, as cast by its shadow, (b) rim of an opaque object, and (c) rim of a transparent object.

convex surface. Alt et al. [AGG03] give an example of a polyhedral approximation of a section of a cylinder with worst-case silhouette size $\Theta(n)$ (see Figure 12.1). Moreover, their example can be modified in such a way that the surface is smooth, and its polyhedral approximation is as “nice” as one might hope (for instance, it can be required that the faces are fat triangles that all have almost the same size).

In this paper we prove an upper bound on the *expected size* of the silhouette for random viewpoints. We prove that the silhouette of a polyhedron that approximates a surface in a reasonable way has expected size $O(\sqrt{n})$. Note that the average is taken over all viewpoints for a given surface, and not on a set of surfaces.

In Section 12.2, we define precisely the notion of silhouette for polyhedra and general surfaces. We then present and prove our main result in Section 12.3 and conclude in Section 12.4.

12.2 Definitions

The term silhouette has been used in the literature to represent several different notions, depending on the application, reflecting such issues as : is the object considered opaque or transparent? Is occlusion taken into account? Is one interested by what the eye perceives, *i.e.*, a plane curve, or by the space curve which gave birth to it? In the area of photography, for instance, a silhouette (also called apparent boundary) is defined as an outline of a solid object, as cast by its shadow, that appears dark against a light background (Figure 12.2(a)). In the field of computer vision, by contrast, the silhouette (also called rim, profile or contour generator) is roughly defined as the curve on the surface that separates front face regions from the back ones, either for opaque (Figure 12.2(b)) or for transparent (Figure 12.2(c)) objects.

In this paper we prove an upper bound on the size of the transparent silhouette ; since such a silhouette contains the apparent boundary and the contour, our bounds also apply to all these types of silhouettes. In the rest of the paper the term silhouette will be used to mean transparent silhouette.

In the rest of this section we give a formal definition of silhouettes of polyhedra and then provide a definition for more general surfaces.

12.2.1 Polyhedra

The (transparent) *silhouette* of a polyhedron from a viewpoint (possibly at infinity) is the set of edges that are adjacent to a front face and a back face. A face is considered a front face if the angle between its normal vector and a vector from a point of the face to the viewpoint is acute, and a back face if that angle is larger than $\pi/2$. If the point of view is in the plane containing the face, we refer to the definition of silhouettes for the case of general surfaces. The normal vectors should point outwards, but what really matters is that the orientation is consistent for the two faces that share this edge, so this definition also applies to non-orientable (necessarily self-intersecting) polyhedra.

In this paper, we call complexity of a silhouette (of a polyhedron) its number of edges.

12.2.2 General surfaces

Our objective is to bound the size of the silhouette of a polyhedron. To achieve this goal, we need to relate the silhouette of the polyhedron to the silhouette of the surface it approximates, which means we need a definition of silhouettes that applies to a larger class of objects. Although this may seem unintuitive, we first define the silhouette as a set of rays, and then relate this to the more usual concept of a set of points on the surface.

Let S be a compact 2-manifold without boundary. It separates \mathbb{R}^3 in two non-empty open regions; call \mathcal{O} and \mathcal{O}' their closures (so $\mathcal{O} \cap \mathcal{O}' = S$ and $\mathcal{O} \cup \mathcal{O}' = \mathbb{R}^3$). Let V be a viewpoint not on S but possibly at infinity. The (transparent) *silhouette* of S from V is the set of rays R starting from V that are tangent to S in a non-crossing way (R may cross S elsewhere). More formally, we require that there exists an open segment u of R that contains a connected component of $R \cap S$ and is contained either in \mathcal{O} or \mathcal{O}' .

This definition defines a set of rays. The silhouette can also be seen as the trace of this set of rays on the surface. More precisely, for each ray R on the silhouette, we consider the closest point to V on each connected component of $R \cap S$ that satisfies the non-crossing property. This definition is consistent with the one given for the particular case of polyhedra, and is the one we will use in this paper.

For a given viewpoint at infinity, we define the (projected) *length* of the silhouette as the length (counted with multiplicity if several points have the same projection) of the projection of the silhouette, along the direction given by the viewpoint, on an orthogonal plane.

Remark. The definition of the silhouette can be extended to cases where S is not a 2-manifold, but an immersion of a compact 2-manifold. More precisely, we have a 2-manifold S' and an application $f : S' \rightarrow \mathbb{R}^3$ such that $S = f(S')$ and for any point on S' there exists a neighborhood U of that point such that U and $f(U)$ are homeomorphic. The local orientation is sufficient to decide whether R crosses S or not (note that more complicated things can happen than crossing or being tangent, even with smooth surfaces; for instance, the surface may ripple an infinite number of times in the neighborhood of a point, making it impossible to define on which side of S R is near the intersection point). This remark extends to the whole paper and, in particular, to Theorem 12.1. However, we do not give either a definition or a proof of this, as it would uselessly make everything more obscure.

12.3 Main results

Let S be a compact 2-manifold without boundary whose silhouettes have finite average length, $\text{silh}(S)$, where the average is taken over all viewpoints at infinity. Let P_n be a polyhedron with n triangular faces, that is homeomorphic to S through $f_n : P_n \rightarrow S$, such that :

1. the length of any edge of P_n is at least $\frac{\alpha}{\sqrt{n}}$ and
2. for any point x on P_n , $d(x, f_n(x)) < \frac{\beta h(x)}{\sqrt{n}}$ where $h(x)$ is the smallest height of the triangle(s) of P_n that contain(s) x ,

where α and β are two arbitrary positive numbers and $d(\cdot)$ denotes the Euclidean distance.

Theorem 12.1. *The expected complexity of the silhouette of P_n is $O(\sqrt{n})$, where the average is taken over all viewpoints at infinity. More precisely, for any n , the expected complexity is at most*

$$\left(15\beta + \frac{24}{\alpha} \text{silh}(S)\right) \sqrt{n}.$$

Note that the bound is valid for any n and any polyhedron P_n satisfying the above assumptions. Note also that the bound depends on S only by the average length of its silhouette.

We first clarify the meaning of the hypotheses on P_n and their implications. We then prove Theorem 12.1 in Section 12.3.2. We finally show in Section 12.3.3 how Theorem 12.1 can be generalized to surfaces with boundary and viewpoints at finite distance. In particular, we prove the following result.

Let S' be any compact two-manifold with boundary of finite length and whose silhouette has finite average length (taken over all viewpoints at infinity).

Theorem 12.2. *Any mesh P_n with n triangular faces that approximates S' according to Hypotheses 1 and 2 has a silhouette of expected complexity $O(\sqrt{n})$ when the viewpoint is chosen uniformly at random in a ball.*

12.3.1 Meaning of the hypotheses

Hypothesis 1 is here to avoid short edges. The main idea of the proof is to link the complexity of the silhouette to its length, and arbitrarily short edges would make this impossible. Now the $\frac{1}{\sqrt{n}}$ factor makes sense : intuitively, since the polyhedron has n faces, each face has area of order $\frac{1}{n}$, which means that the edges have length of order $\frac{1}{\sqrt{n}}$.

Hypothesis 2 is rather technical, and we discuss instead the meaning of the following two more intuitive hypotheses, which, together with Hypothesis 1, imply²⁰ Hypothesis 2.

3. The faces of P_n are fat.
4. For any x on P_n , $d(x, f_n(x)) < \frac{\gamma}{n}$, where γ is some positive constant.

Hypothesis 3 is quite natural. Hypothesis 4 ensures that P_n approximates S . Furthermore, the $\frac{1}{n}$ factor is reasonable ; indeed, in 2D, when considering a regular polygon with edge length $\Theta(\frac{1}{\sqrt{n}})$ inscribed in a circle of radius 1, the maximal distance between a point on the polygon and the circle is $\Theta(\frac{1}{n})$. The situation is the same in 3D. Basically it means that the error when approximating the surface with a plane is of the second order.

Our hypotheses (1-3-4 or 1-2) ensure that the homeomorphism f_n has good properties, that is that, roughly speaking, the polyhedron can be obtained by only a small perturbation of the surface while keeping the normal vectors in approximately the same directions. This is crucial for our proof since otherwise, for example, a cylinder can be approximated by a lantern of Schwarz [Sch90] (see Figure 12.3(a)) whose silhouette has expected complexity $\Theta(n)$ and unbounded length.

Notice that the existence of polyhedra with arbitrarily large number of edges that approximate the surface according to these hypotheses is a constraint on the surface. Not every surface admits such an approximation (think of the neighborhood of 0 in the surface defined by $z = (x^2 + y^2)^{1/8}$ as shown in Figure 12.3(b)). However, the class of surfaces for which such approximations exist is quite large. It includes, in particular, smooth surfaces and polyhedra with fat faces.

12.3.2 Proof of Theorem 12.1

We consider a point of view chosen randomly at infinity. We call l_e the length of an edge e of polyhedron P_n and θ_e the exterior dihedral angle associated to e (see Figure 12.4).

Let T_e denote the union of the two triangles adjacent to edge e (including e but not the other edges). For any part \mathcal{R} of S , let $\text{silh}(\mathcal{R})$ be the average length of the part of the silhouette of S that lies in \mathcal{R} .

²⁰Indeed, for any x in P_n , Hypotheses 1 and 3 imply that $h(x) \geq \delta/\sqrt{n}$ for some positive constant δ ; Hypothesis 2 then follows from Hypothesis 4 since $h(x)/\sqrt{n} \geq \delta/n \geq \delta/\gamma \cdot d(x, f_n(x))$.

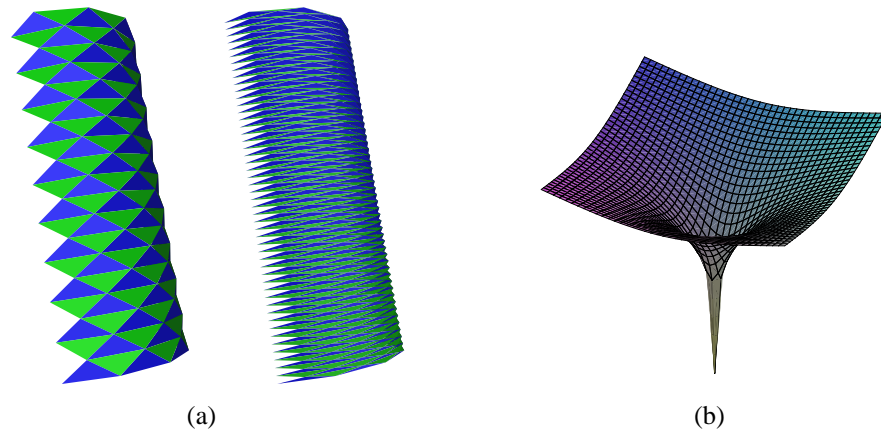


FIG. 12.3 – (a) Two half lanterns of Schwarz (courtesy of Boris Thibert). (b) A surface that cannot be approximated with the right properties.

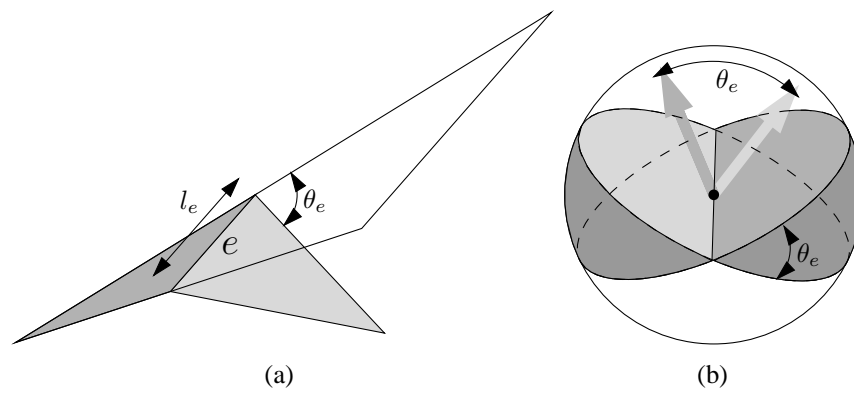


FIG. 12.4 – (a) Length and dihedral angle of an edge ; (b) set of directions for which e is on the silhouette.

We first recall a classical formula on the expected size of silhouettes which can also be found, for instance, in [McG04].

An edge e is on the silhouette if the direction of view is in the dark area of the sphere of directions of Figure 12.4(b). The angular measure of this region is $4\theta_e$, which means that the probability for e to be on the silhouette is θ_e/π . The expected number of edges on the silhouette is thus

$$E = \frac{1}{\pi} \sum_{\text{edge } e} \theta_e.$$

We now state our main lemma. The general idea of this lemma is that under strong hypotheses (S has bounded curvature, the edges have length $\Theta(\frac{1}{\sqrt{n}})$, and Hypotheses 3 and 4 are satisfied), one can prove that $\theta_e \leq \frac{C}{\sqrt{n}}$ for some constant C . In cases where this inequality does not hold, edge e is near some kind of edge of the surface, or at least some feature that will appear quite often on the silhouette and we are going to charge this edge to the silhouette of S .

Lemma 12.3. *For any edge e on P_n ,*

$$\theta_e \leq \frac{C}{\sqrt{n}} + \frac{8\pi}{l_e} \text{silh}(f_n(T_e)) \quad \text{with } C = 31.3\beta.$$

Theorem 12.1 follows from Lemma 12.3. Indeed, since P_n has $\frac{3n}{2}$ edges, each of length at least $\frac{\alpha}{\sqrt{n}}$ (by Hypothesis 1), we get that the expected complexity of the silhouette is

$$E \leq \frac{1}{\pi} \frac{3n}{2} \frac{C}{\sqrt{n}} + 8 \frac{\sqrt{n}}{\alpha} 3 \text{silh}(S),$$

because $\sum_{\text{edge } e} \text{silh}(f_n(T_e)) = 3 \text{silh}(S)$ since the length of the silhouette of S that lies in the image (through f_n) of a triangle is counted three times (once per edge). Hence,

$$E \leq \left(15\beta + \frac{24}{\alpha} \text{silh}(S)\right) \sqrt{n} = O(\sqrt{n}).$$

Proof. [Proof of Lemma 12.3] The idea of the proof is as follows. Consider the set of directions for which e is on the silhouette. We first construct a subset Ω of these directions whose measure is a constant times $\theta_e - \frac{C}{\sqrt{n}}$ (see Figure 12.5). We then prove a lower bound on the length of the silhouette of $f_n(T_e)$ for all these directions, and deduce the result.

Let C be a positive constant, whose value will be defined later (see Equation 12.2). For any edge e on P_n , we can assume that $\theta_e - \frac{C}{\sqrt{n}} > 0$ since, otherwise, $\theta_e \leq \frac{C}{\sqrt{n}}$ and there is nothing else to prove.

The set of directions for which e is on the silhouette is the set of directions between the planes defined by the faces adjacent to e . Rotate each face about e by an angle of $\frac{C}{2\sqrt{n}}$ so that the exterior dihedral angle decreases by $\frac{C}{\sqrt{n}}$ (see Figure 12.5). Ω is defined to be the set of directions between these two new planes that make an angle larger than $\pi/3$ with the line supporting e . The measure of the set of directions between these two planes is $4(\theta_e - \frac{C}{\sqrt{n}})$. Restricting this set of directions to those that make an angle larger than $\pi/3$ with the line supporting e , we get, by integrating on the sphere of directions, that the measure of Ω is $2(\theta_e - \frac{C}{\sqrt{n}})$.

The remaining step uses the property, which we prove in Corollary 12.5, that for all the directions in Ω , the silhouette of $f_n(T_e)$ has length at least $l_e/4$. Assuming this temporarily, we sum this inequality over Ω . The smaller side of the inequality is $2\frac{l_e}{4}(\theta_e - \frac{C}{\sqrt{n}})$. The larger side is the integral of the length of the silhouette of $f_n(T_e)$ over all directions in Ω , which is smaller than this same integral over all directions, that is $4\pi \text{silh}(f_n(T_e))$. Hence $4\pi \text{silh}(f_n(T_e)) \geq \frac{l_e}{2}(\theta_e - \frac{C}{\sqrt{n}})$, which concludes the proof. \square

We now state a lemma and its corollary which we used in the proof of Lemma 12.3 under the hypothesis that $\theta_e - \frac{C}{\sqrt{n}} > 0$. We can thus assume in the sequel that this property holds.

Let e' be the segment obtained by clipping from e all the points at distance less than $\frac{l_e}{4}$ from its extremities. Refer now to Figures 12.6(a)–(b).

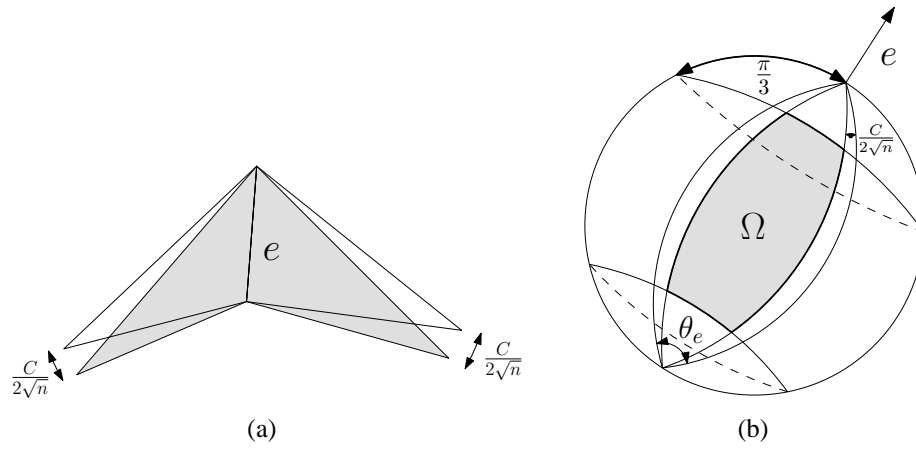
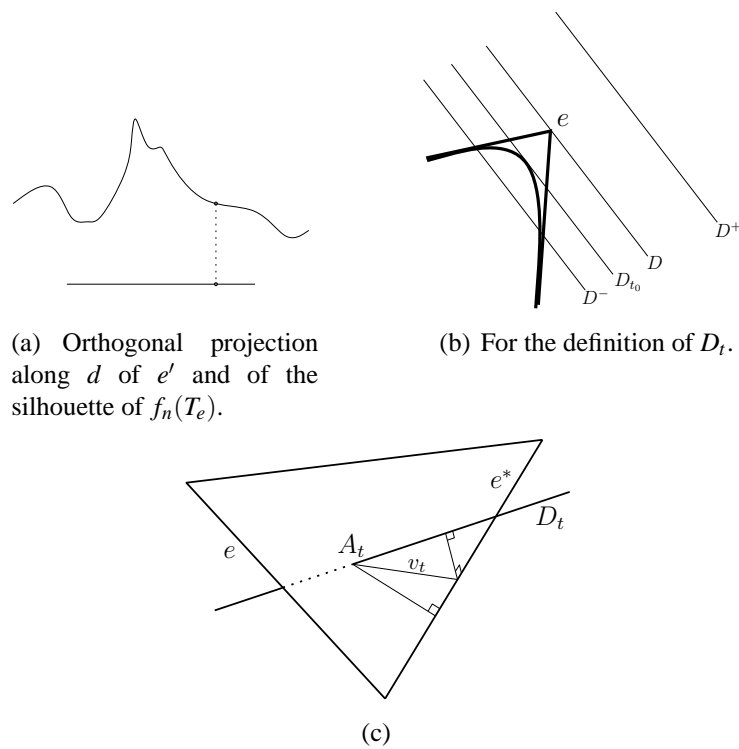


FIG. 12.5 – Construction of Ω .



(a) Orthogonal projection along d of e' and of the silhouette of $f_n(T_e)$.

(b) For the definition of D_t .

(c)

FIG. 12.6 – For the proofs of Lemma 12.4 and Corollary 12.5.

Lemma 12.4. Any line with direction $d \in \Omega$ that intersects e' can be translated in a direction orthogonal to e and d until it becomes tangent to S in $f_n(T_e)$.

Corollary 12.5. For any direction d in Ω , the silhouette of $f_n(T_e)$ has length at least $\frac{l_e}{4}$.

Proof. Consider the projection of e' and of the silhouette of $f_n(T_e)$ onto a plane orthogonal to d (see Figure 12.6(a)). It follows from Lemma 12.4 that, in that plane, each point on the projection of e' maps to a point on the projected silhouette in the direction orthogonal to e' . Hence, the projected silhouette is longer than the projection of e' , which is at least $\frac{\sqrt{3}}{2}$ times the length of e' since d makes an angle of at least $\pi/3$ with e' . Thus the silhouette of $f_n(T_e)$ has length at least $\frac{\sqrt{3}}{2} \frac{l_e}{2} > \frac{l_e}{4}$. \square

Proof. [Proof of Lemma 12.4] Let D denote a line with direction $d \in \Omega$ that intersects e' . Let T_1 and T_2 denote the two triangles adjacent to e and let h_1 and h_2 denote their respective smallest heights. Let $\chi_i = \beta h_i / \sqrt{n}$, $\chi^+ = \max(\chi_1, \chi_2)$, and $\chi^- = \min(\chi_1, \chi_2)$. Refer now to Figure 12.6(b). We call D_t , $t \in [-\chi^-, \chi^+]$, the line obtained by translating D at distance $|t|$ in a direction orthogonal to the plane defined by e and d ; positive values of t correspond to lines in the half-space bounded by the plane defined by e and D , and not containing T_e ; negative values of t correspond to lines in the other half-space. For clarity, we denote $D_{-\chi^-}$ by D^- and D_{χ^+} by D^+ .

By construction, D^+ is at distance χ^+ from T_e . Thus D^+ does not intersect $f_n(T_e)$, by Hypothesis 2. We prove that D^- intersects $f_n(T_e)$ and that no line D_t intersects the boundary of $f_n(T_e)$. This will imply that, sweeping D_t from D^+ to D^- , the first line D_{t_0} that intersects $f_n(T_e)$ is tangent to $f_n(T_e)$ at one of its interior point, which will conclude the proof.

We first prove that no line D_t intersects the boundary of $f_n(T_e)$. In other words, we prove that, for each edge e^* on the boundary of T_e , no line D_t intersects $f_n(e^*)$. Let T_i be the triangle (of T_e) containing e^* . By Hypothesis 2, it is sufficient to prove that the distance between D_t and e^* remains greater than or equal to χ_i for all t .

First notice that it is sufficient to prove that the distance between D_t and e^* remains greater than or equal to χ_i for all $t \in [-\chi^-, 0]$. Indeed, then, the distance between $D_0 = D$ and e^* is at least χ_i , and the distance between D_t and e^* increases for $t \geq 0$ (see Figure 12.6(b)).

Let Γ be the smallest angle d can make with the plane containing T_i and refer to Figure 12.6(c). Let A_t be the point of intersection between D_t and the plane containing T_i and v_t be the distance between A_t and the point on e^* that realizes the distance between D_t and e^* . The distance between D_t and e^* satisfies $d(D_t, e^*) \geq v_t \sin \Gamma \geq d(A_t, e^*) \sin \Gamma$. Hence, for proving that $d(D_t, e^*) \geq \chi_i$ for $t \leq 0$, it is sufficient to prove that $d(A_t, e^*) \geq \frac{\chi_i}{\sin \Gamma}$ for $t \leq 0$. We set $a = \frac{\chi_i}{\sin \Gamma}$ to simplify the notation.

We just proved that $d(A_t, e^*) \geq a$ implies $d(D_t, e^*) \geq \chi_i$ (for all t). Conversely, we have that $d(D_t, e^*) < \chi_i$ implies $d(A_t, e^*) < a$. Similarly, for edge e , we get that $d(D_t, e) < \chi_i$ implies $d(A_t, e) < a$. By definition of D_t , we have that $d(D_t, e) < \chi_i$ for $t \leq 0$, thus $d(A_t, e) < a$ for $t \leq 0$. Furthermore, the angle between e and segment $\{A_t \mid t \in [-\chi^-, \chi^+]\}$ is at least $\pi/3$ because this angle is at least the angle between their orthogonal projection on the plane defined by e and D that is the angle between e and D since all A_t lie in the plane spanned by D_t which projects on D ; the lower bound of $\pi/3$ follows since the angle between e and D is at least $\pi/3$ by definition of Ω . Hence, the locus of points A_t , for $t \leq 0$, lies in a region, denoted Y , shown in dark gray in Figure 12.7(a). For proving that $d(A_t, e^*) \geq a$ for $t \leq 0$, it is thus sufficient to prove that this region does not intersect the set, denoted Y' , of points at distance less than a from e^* (shown in light gray in Figure 12.7(a)).

Referring to Figures 12.7(b)–(c), let p be the endpoint of e' the closest to e^* and s be its projection on the line supporting e^* . If the two regions Y and Y' intersect, there exists a point q in the intersection that is at distance less than or equal to $\frac{2}{\sqrt{3}}a$ from p and at distance less than or equal to a from e^* ; thus $d(p, s) \leq d(p, e^*) \leq d(p, q) + d(q, e^*) \leq (1 + \frac{2}{\sqrt{3}})a$. On the other hand, $d(p, s)$ is one fourth of one of the heights of the triangle T_i and thus is at least $\frac{h_i}{4}$. Hence, if the two regions intersect, then $\frac{h_i}{4} \leq \left(1 + \frac{2}{\sqrt{3}}\right) \frac{\chi_i}{\sin \Gamma}$. We postpone to Lemma 12.7 the proof that, with $C = 31.3\beta$, we have $\frac{h_i}{4} > \left(1 + \frac{2}{\sqrt{3}}\right) \frac{\chi_i}{\sin \Gamma}$, which implies that the two regions Y and Y' are disjoint. This concludes the proof that no line D_t intersects the boundary of $f_n(T_e)$.

We now prove that D^- intersects $f_n(T_e)$. Consider a projection, $p(\cdot)$, along the direction d onto a plane orthogonal to d . We proved that, for any of the two triangles T_i , Y is at distance at least χ_i from each edge $e^* \neq e$ of T_i . It follows that Y lies in triangle T_i and thus that D_t intersects T_i for all $t \leq 0$. Therefore, D^- intersects T_i and is at distance at least χ_i from each edge $e^* \neq e$ of T_i , for $i = 1, 2$. Furthermore, D^- is at distance $\chi^- = \min(\chi_1, \chi_2)$ from

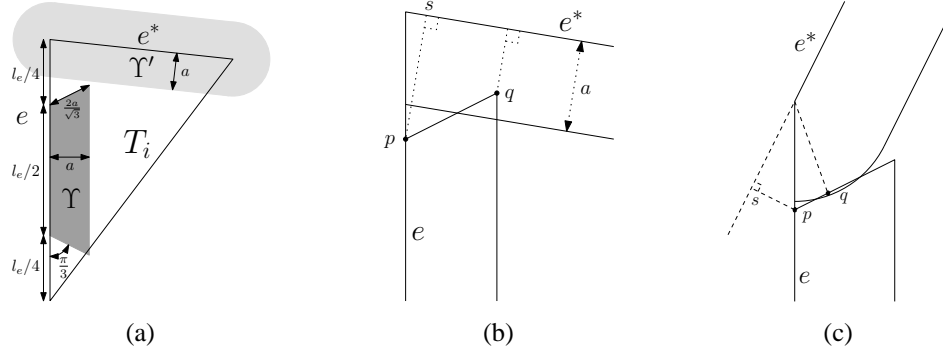


FIG. 12.7 – For the proof of Lemma 12.4.

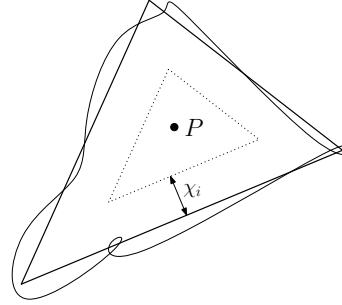


FIG. 12.8 – Projection of ∂T_i , $f_n(\partial T_i)$ and D^- .

e , by definition. We now consider the triangle T_i for which $\chi_i = \chi^-$. It follows that D^- is at distance at least χ_i from all three edges of T_i . Thus D^- projects to a point $P = p(D^-)$ inside triangle $p(T_i)$, at distance at least χ_i from the three edges of $p(T_i)$ (see Figure 12.8).

Roughly speaking, by Hypothesis 2, the curve $f_n(\partial T_i)$ is at distance less than χ_i from ∂T_i (the boundary of T_i) thus its projection $p(f_n(\partial T_i))$ is at distance less than χ_i from the edges of $p(T_i)$. It is thus intuitively clear that $p(D^-)$ intersects $p(f_n(T_i))$, and thus that D^- intersects $f_n(T_i)$ (and thus $f_n(T_e)$).

More formally, consider the application g_n from the triangle $p(T_i)$ to the plane containing it such that, for any point x in T_i , the point $p(x)$ is sent to the point $g_n(p(x)) = p(f_n(x))$. We first prove that the curves $p(\partial T_i)$ and $g_n(p(\partial T_i))$ are homotopic in $\mathbb{R}^2 \setminus P$. Consider the continuous application

$$F : \partial T_i \times [0, 1] \longrightarrow \mathbb{R}^2$$

$$(x, \lambda) \longrightarrow \lambda p(x) + (1 - \lambda) g_n(p(x)) = \lambda p(x) + (1 - \lambda) p(f_n(x)).$$

F is an homotopy between the curves $p(\partial T_i)$ and $g_n(p(\partial T_i))$ in \mathbb{R}^2 . We prove that the image of F does not contain P , which yields the result. The triangle inequality gives

$$d(P, F(x, \lambda)) \geq d(P, p(x)) - d(F(x, \lambda), p(x)).$$

We have already proved that point P is at distance at least χ_i from $p(x)$ for all points x in ∂T_i . On the other hand, the distance between $p(x)$ and $p(f_n(x))$ is larger than or equal to the distance between $p(x)$ and their barycenter $F(x, \lambda)$, for any $\lambda \in [0, 1]$. Hence

$$d(P, F(x, \lambda)) \geq \chi_i - d(p(x), p(f_n(x))).$$

Finally, since $d(p(x), p(f_n(x))) < \chi_i$ for all $x \in T_i$, by Hypothesis 2, we have that $d(P, F(x, \lambda)) > 0$ for all (x, λ) . Hence, the image of F does not contain point P and thus the curves $p(\partial T_i)$ and $g_n(p(\partial T_i))$ are homotopic in $\mathbb{R}^2 \setminus P$.

Now, we can contract $p(\partial T_i)$ to a point while remaining in $p(T_i)$. Composing this with g_n gives a contraction of $g_n(p(\partial T_i))$ in $g_n(p(T_i))$. On the other hand, there is no contraction of $p(\partial T_i)$ in $\mathbb{R}^2 \setminus P$ (since P is in $p(T_i)$), thus

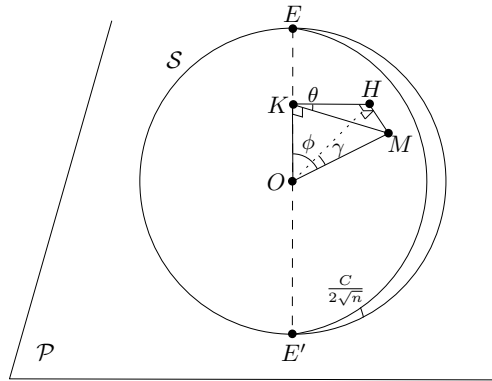


FIG. 12.9 – For the proof of Lemma 12.6.

there is no contraction of its homotopic curve $g_n(p(\partial T_i))$ in $\mathbb{R}^2 \setminus P$. Hence, there exists a curve that is contractible in $g_n(p(T_i))$ but not in $\mathbb{R}^2 \setminus P$. It follows that $g_n(p(T_i))$ is not included in $\mathbb{R}^2 \setminus P$. Hence P is in $g_n(p(T_i)) = p(f_n(T_i))$. Therefore, D^- intersects $f_n(T_i)$ and thus $f_n(T_e)$, which concludes the proof. \square

We finally prove the two following simple technical lemmas which complete the proof of Theorem 12.1. Recall that Γ is the smallest angle a direction $d \in \Omega$ can make with the plane containing T_i .

Lemma 12.6. $\sin \Gamma = \frac{\sqrt{3}}{2} \sin \frac{C}{2\sqrt{n}}$.

Proof. In the following, we identify the sphere of directions with a sphere S embedded in \mathbb{R}^3 ; let O denote its center. We assume that the embedding preserves directions (*i.e.*, for any direction d , the corresponding point M on S is such that d and OM have the same direction).

Refer to Figure 12.9. Let d be a direction in Ω and M be its corresponding point on S . Consider one of the T_i and let \mathcal{P} be the plane containing O and parallel to the plane containing T_i . Let H be the orthogonal projection of M onto plane \mathcal{P} . Let E and E' be the two points on S that correspond to the two (opposite) directions of segment e . Let K be the orthogonal projection of M (and H) onto the line EE' . Finally, let θ be the angle $\angle MKH$, ϕ be the angle $\angle MOK$, and γ be the angle $\angle MOH$.

It follows from these definitions that

$$\sin \gamma = \frac{HM}{OM} = \frac{HM}{KM} \frac{KM}{OM} = \sin \theta \sin \phi.$$

Now, the angle γ is also the angle between direction d and the plane that contains T_i . Thus $\Gamma = \inf_{d \in \Omega} \gamma$, by definition of Γ . The angle θ is the angle between the plane containing T_i and the plane containing e and d . It thus follows from the definition of Ω that $\inf_{d \in \Omega} \theta = \frac{C}{2\sqrt{n}}$ (see Figure 12.5(a)). The angle ϕ is the angle between d and the line containing e . It thus also follows from the definition of Ω that $\inf_{d \in \Omega} \phi = \frac{\pi}{3}$. In addition, since γ , θ and ϕ are in $[0, \frac{\pi}{2}]$, we have

$$\sin \Gamma = \inf_{d \in \Omega} \sin \gamma, \quad \inf_{d \in \Omega} \sin \theta = \sin \frac{C}{2\sqrt{n}} \quad \text{and} \quad \inf_{d \in \Omega} \sin \phi = \sin \frac{\pi}{3}.$$

Furthermore, the constraints on θ and ϕ in the definition of Ω are independent. Thus, the minima of θ and ϕ can be attained for the same direction d in Ω . It follows that

$$\inf_{d \in \Omega} (\sin \theta \sin \phi) = \inf_{d \in \Omega} \sin \theta \cdot \inf_{d \in \Omega} \sin \phi.$$

We can thus conclude that

$$\sin \Gamma = \inf_{d \in \Omega} \sin \gamma = \inf_{d \in \Omega} \sin \theta \sin \phi = \inf_{d \in \Omega} \sin \theta \cdot \inf_{d \in \Omega} \sin \phi = \frac{\sqrt{3}}{2} \sin \frac{C}{2\sqrt{n}}.$$

□

Lemma 12.7. $\frac{h_i}{4} > \left(1 + \frac{2}{\sqrt{3}}\right) \frac{\chi_i}{\sin \Gamma}$ with $C = 31.3\beta$.

Proof. By Lemma 12.6, replacing χ_i and Γ by their values in the inequality $\frac{h_i}{4} > \left(1 + \frac{2}{\sqrt{3}}\right) \frac{\chi_i}{\sin \Gamma}$ gives

$$\frac{h_i}{4} > \left(1 + \frac{2}{\sqrt{3}}\right) \frac{\frac{\beta h_i}{\sqrt{n}}}{\frac{\sqrt{3}}{2} \sin\left(\frac{C}{2\sqrt{n}}\right)}$$

or equivalently

$$4\beta \left(1 + \frac{2}{\sqrt{3}}\right) < \sqrt{n} \frac{\sqrt{3}}{2} \sin\left(\frac{C}{2\sqrt{n}}\right). \quad (12.1)$$

Notice first that for large enough values of n , using the approximation $\sin x \approx x$ in the neighborhood of zero, we derive the sufficient condition

$$C > \frac{16\beta}{\sqrt{3}} \left(1 + \frac{2}{\sqrt{3}}\right) \sim 19.9\beta.$$

Now, since we want our result for all n , the computation is more complicated. Recall first that for any strictly concave function f , such that $f(0) = 0$, $f(x) > \frac{f(x_0)}{x_0} x$ for any $x \in (0, x_0)$. It follows that $\sin x > \frac{2}{\pi} x$ for any $x \in (0, \frac{\pi}{2})$. Since we assumed that $\theta_e - \frac{C}{2\sqrt{n}} > 0$ and thus that $0 < \frac{C}{2\sqrt{n}} < \frac{\theta_e}{2} < \frac{\pi}{2}$, we get

$$\sin\left(\frac{C}{2\sqrt{n}}\right) > \frac{2}{\pi} \frac{C}{2\sqrt{n}}.$$

To guarantee inequality (12.1), it is thus sufficient to have

$$4\beta \left(1 + \frac{2}{\sqrt{3}}\right) \leq \sqrt{n} \frac{\sqrt{3}}{2} \frac{2}{\pi} \frac{C}{2\sqrt{n}}.$$

or equivalently

$$C \geq \frac{8}{3} (2 + \sqrt{3}) \pi \beta \sim 31.27\beta,$$

which concludes the proof. Note that we can set

$$C = 31.3\beta. \quad (12.2)$$

in the definition of Ω (in the proof of Lemma 12.3) since Inequality (12.1) is the only constraint on C . □

12.3.3 Generalizations

We prove here Theorem 12.2. We first show that Theorem 12.1 generalizes to the case where the viewpoint is chosen randomly at finite distance. We then show that considering surfaces with boundary does not change the asymptotic expected complexity of the silhouette.

Point of view at finite distance. We have thus far restricted ourselves to the case where the viewpoint is chosen uniformly at random at infinity. However, our result applies to any distribution of viewpoints such that the probability for an edge e to be on the transparent silhouette is $O(\theta_e)$, where θ_e is the exterior dihedral angle associated to e ; indeed, the expected number of edges on the silhouette is then $\sum_{\text{edge } e} O(\theta_e)$ and we get the result

by applying, as before, Lemma 12.3.²¹ Such a distribution of viewpoints is obtained, in particular, when the point of view is chosen uniformly at random in a ball. This is also the case if S delimits a bounded region \mathcal{O} and the viewpoint is chosen uniformly at random in $B \setminus \mathcal{O}$, for a ball B .

²¹Note that, in Lemma 12.3, $\text{silh}(f_n(T_e))$ always refers to an expected length for a viewpoint chosen randomly at infinity.

Surfaces with boundary. Let S be a 2-manifold with boundary \mathcal{B} . We consider that the boundary is always on the transparent silhouette and so the definition of the transparent silhouette of a 2-manifold S with boundary is exactly that of a 2-manifold without boundary plus the boundary \mathcal{B} .

The surface S is approximated by a triangulated mesh P_n that satisfies Hypotheses 1 and 2, as in the case without boundary, except that now the mesh may not be a polyhedron (some edges may have only one adjacent face rather than two).

To give an upper bound on the number of edges on the silhouette of the mesh, we consider the boundary edges and the other (non-boundary) edges separately. For the non-boundary edges, the same reasoning as before still holds. For the boundary edges, it is easy to see that the length (in 3D) of the boundary of P_n cannot be much larger than the length of \mathcal{B} . Indeed, the two are homeomorphic, and the hypotheses imply that the image of an edge e , of length l_e , is a curve whose extremities lie at distance at least $l_e - 2\beta \cdot \frac{l_e}{\sqrt{n}} = \Omega(\frac{1}{\sqrt{n}})$ apart. This means that the length of \mathcal{B} is at least $\Omega(\frac{1}{\sqrt{n}})$ times the number of boundary edges of P_n . Hence, the number of boundary edges of P_n is at most $O(\sqrt{n})$ times the length of \mathcal{B} . So, if the length of \mathcal{B} is bounded, the expected complexity of the silhouette of P_n is $O(\sqrt{n})$.

12.4 Conclusion

This paper gives an idea of why, and when, the usual claim that the silhouette of a triangulated mesh has size $O(\sqrt{n})$ is valid. In particular, we have given a set of conditions such that any triangulated mesh approximating a surface in a way that satisfies those conditions has a silhouette of expected size $O(\sqrt{n})$. Roughly speaking, the mesh should have no short edges, its faces should be fat, and the distance between it and the surface it approximates should never be too large. The surface itself is not necessarily everywhere differentiable and may have boundaries.

A natural question to ask is whether meshes satisfying those conditions exist. In fact, for smooth surfaces, the meshes produced by Boissonnat and Oudot [BO05] are one such example. The critical property of the meshes they compute is that the ratio between the size of the largest and the smallest triangles remains bounded, although meshes are non-uniform with small triangles in areas of large curvature. However, in order to satisfy our conditions, non-smooth surfaces with curved sharp edges (such as a flying saucer with a sharp equatorial arc) would have to be approximated by small triangles over the whole surface. Such meshes would have silhouettes of expected size $O(\sqrt{n})$ but then n would be much larger than necessary; it would be reasonable to replace the large number of triangles used to mesh large flat portions of the surface with a smaller number of large triangles, which would give a silhouette of size closer to linear. This explains why the observed expected size of silhouettes, as shown in [McG04], is larger than $O(\sqrt{n})$. The fact that non-uniform meshes approximating such surfaces appear, in computer graphics, to have silhouettes of expected size much smaller than n is thus likely due to additional properties of the surfaces or the meshes.

Acknowledgments

The authors wish to thank Gert Vegter, who introduced the problem to us, and gratefully acknowledge fruitful discussions on this topic with Helmut Alt, Olivier Devillers, Hazel Everett, Xavier Goaoc, Bruno Levy and Sylvain Petitjean.

Chapitre 13

Between umbra and penumbra

Une version courte de cet article est parue dans les proceedings de la *23rd European Conference on Computational Geometry* [DDE⁺07].

Abstract

Computing shadow boundaries is a difficult problem in the case of non-point light sources. A point is in the umbra if it does not see any part of any light source ; it is in full light if it sees entirely all the light sources ; otherwise, it is in the penumbra. While the common boundary of the penumbra and the full light is well understood, less is known about the boundary of the umbra. In this paper we prove various bounds on the complexity of the umbra and the penumbra cast by a segment or polygonal light source on a plane in the presence of polygon or polytope obstacles.

In particular, we show that a single segment light source may cast on a plane, in the presence of two triangles, four connected components of umbra and that two fat convex obstacles of total complexity n can engender $\Omega(n)$ connected components of umbra. In a scene consisting of a segment light source and k disjoint polytopes of total complexity n , we prove an $\Omega(nk^2 + k^4)$ lower bound on the maximum number of connected components of the umbra and a $O(nk^3)$ upper bound on its complexity. We also prove that, in the presence of k disjoint polytopes of total complexity n , some of which being light sources, the umbra cast on a plane may have $\Omega(n^2k^3 + nk^5)$ connected components and has complexity $O(n^3k^3)$.

These are the first bounds on the size of the umbra in terms of both k and n . These results prove that the umbra, which is bounded by arcs of conics, is intrinsically much more intricate than the full light/penumbra boundary which is bounded by line segments and whose worst-case complexity is in $\Omega(n\alpha(k) + km + k^2)$ and $O(n\alpha(k) + km\alpha(k) + k^2)$, where m is the complexity of the polygonal light source.

13.1 Introduction

Shadows play a central role in human perception [MKK98, Wan92]. A wide variety of approaches have been considered for simulating and rendering shadows (see, for example, the surveys [Dur00, WPF90]) and many methods make extensive use of graphics hardware (see the survey [HLHS03]). Unfortunately, computing realistic shadows efficiently is a difficult problem, particularly in the case of non-point light sources. A part of this difficulty arises from the complicated internal structure that such shadows may have. In this paper we study this structure.

We say that a point is in the *umbra* if it does not see any part of the light source(s) ; it is in *full light* if it sees entirely all the light source(s) ; otherwise, it is in the *penumbra*. While the boundary between the penumbra and the full light is reasonably well-understood (see Section 13.3), less is known about the boundary of the umbra. Nevertheless, there is an extensive literature concerning the explicit computation of these shadow boundaries ; see, for example, [DF94, DDP97, DDP99, DDP02, Hec92, NN83, SG94, Tel92].

Scene type	Lower bounds	Upper bounds
Segment light source		
2 triangles	4	$O(1)$
2 fat polytopes	$\Omega(n)$	$O(n)$
k polytopes	$\Omega(nk^2 + k^4)$	$O(nk^3)$
n-gon light source		
k polytopes	$\Omega(n^2k^3 + nk^5)$	$O(n^3k^3)$

TAB. 13.1 – Lower bounds on the number of connected components and upper bounds on the complexity of the umbra cast on a plane by k polytopes of total complexity $O(n)$.

In this paper we prove various bounds, summarized in Table 13.1, on the complexity of the umbra cast by a segment or polygonal light source on a plane in the presence of polygon or polytope (*i.e.* convex polyhedral) obstacles. In particular, we show that a single segment light source may cast, in the presence of two triangles, four connected components of umbra. We prove that the umbra defined by one segment light source and two fat convex obstacles of total complexity n can have $\Omega(n)$ connected components. We also prove an $\Omega(nk^2 + k^4)$ lower bound on the maximum number of connected components of the umbra and a $O(nk^3)$ upper bound on its complexity in a scene consisting of a segment light source and k disjoint polytopes of total complexity n . Finally, we prove that the umbra cast on a plane by a polygonal light source and k convex obstacles can have $\Omega(n^2k^3 + nk^5)$ connected components and has complexity $O(n^3k^3)$. These are the first bounds on the size of the umbra in terms of both k and n .

Our results are surprising in the sense that they show that the umbra cast by a single segment light source may have many connected components. The fact that the umbra may have four connected components in the case of two triangle obstacles comes as a total surprise. Our lower bounds of $\Omega(nk^2 + k^4)$ and $\Omega(n^2k^3 + nk^5)$ connected components, for k polytopes of total complexity n , is rather pathological in the sense that most of the obstacles are very long and thin. However, we also present a lower bound example of $\Omega(n)$ connected components in the case of two fat polygons or polytopes of complexity $O(n)$. Concerning our upper bounds of $O(nk^3)$ and $O(n^3k^3)$, even though these bounds are not tight, they substantially improve the only previously known bounds for this problem which were the trivial $O(n^4)$ and $O(n^6)$ upper bounds. Finally, it is interesting to point out that even for the simplest case of non-point light sources, obtaining tight bounds on the complexity of the umbra and understanding its structure is a very challenging problem.

The paper is organized as follows. The next section provides notation and definitions. We give in Section 13.3 almost tight lower and upper bounds on the complexity of the boundary between full light and penumbra cast on a plane by a polygonal light source in the presence of polytope obstacles. We present, in Section 13.4, upper bounds on the complexity of the umbra, in Section 13.5, lower bounds on the maximal number of connected components of umbra and conclude in Section 13.6.

13.2 Preliminaries

Let s be a line segment and p a point. We denote by $\langle s, p \rangle$ the set of line transversals of s through p . Similarly, for any triple of segments s_1, s_2 and s_3 , we denote by $\langle s_1, s_2, s_3 \rangle$ its set of line transversals. It is a well-known fact that $\langle s_1, s_2, s_3 \rangle$ consists of lines belonging to the same regulus of a ruled quadric surface (see *e.g.* [Sal15]). More precisely, the line transversals lie on a hyperboloid of one sheet when the three segments are pairwise skew and not all parallel to the same plane. If the segments are pairwise skew and all parallel to the same plane, then the line transversals lie on a hyperbolic paraboloid. Otherwise, they lie in one or two planes. Hence any set of transversals, whether $\langle s, p \rangle$ or $\langle s_1, s_2, s_3 \rangle$, forms patches of a quadric (possibly degenerating to one or two planes). Moreover, the set of transversals consists of at most three patches, or more formally, at most three connected components in line space [BEL⁺05]. We let $\langle s, p \rangle$ and $\langle s_1, s_2, s_3 \rangle$ denote not just sets of lines but also the surface patches they form.

Let \mathcal{P} be a finite set of disjoint convex polygons or polytopes in \mathbb{R}^3 with $\mathcal{L} \subset \mathcal{P}$ identified as light sources. A surface $\sigma = \langle e, v \rangle$ is called an *EV-surface* if there exist two distinct objects $P, Q \in \mathcal{P}$ so that e is an edge of P , v a vertex of Q and σ intersects a light source. A surface $\sigma = \langle e_1, e_2, e_3 \rangle$ is called an *EEE-surface* if there exist three distinct objects $P, Q, R \in \mathcal{P}$ so that e_1, e_2 and e_3 are respectively edges of P, Q and R and σ intersects a light source.

Any plane Π intersects an EV-surface or an EEE-surface in a set of arcs of a conic (each possibly empty or possibly a line segment). Hence the intersection between Π and *all* the EV and EEE surfaces defines an arrangement of arcs of conics on Π .

Here we are interested in the arcs of conics that correspond to shadow boundaries. In particular, we are interested in arcs resulting from the intersection between Π and maximal free line segments²² that intersect a light source and are supported by a line which is on an EV or EEE surface. The intersection of these free line segments with Π defines an arrangement of arcs of conics on Π which we call the *shadow arrangement* on the *shadow plane* Π .

A point p is in the umbra if for any point q on a light source, the segment pq intersects an object from $\mathcal{P} \setminus \mathcal{L}$. Similarly, p is in full light if for any point q on a light source, the segment pq does not intersect any object from $\mathcal{P} \setminus \mathcal{L}$. Otherwise, p is in the penumbra.

We will make extensive use of the fact that the boundary of the umbra and penumbra consists of arcs of the shadow arrangement (see, for example, [Hec92]). Notice that not all arcs of the shadow arrangement are on the umbra or penumbra boundaries; some arcs correspond to other lighting discontinuities.

Throughout this paper, we consider the regions of umbra and penumbra on a plane cast by a segment light source or polygonal light source(s) in the presence of convex polygons or polytopes.

13.3 The penumbra boundary

We recall here some straightforward and well-known properties of the penumbra and give bounds on the complexity of the common boundary of the penumbra and the full light. In this section we refer to the union of the umbra and penumbra as the *shadow region*.

PROPERTY 1. The shadow region cast by a light source on a plane in the presence of obstacles is the union of all the shadow regions cast by each obstacle.

PROPERTY 2. The shadow region cast on a plane Π by a polygonal light source S in the presence of one polytope P is the intersection of half-planes in Π , each of which is defined as the intersection of Π with a half-space that contains P but not S , is bounded by a plane tangent to both of them, and contains an edge of one of them.

Note that these two properties imply that the boundary of the shadow region is only composed of line segments induced by EV-surfaces.

Theorem 13.1. *The complexity of the shadow region cast on a plane Π by a convex polygonal light source of complexity m in the presence of k convex polyhedra of total complexity n is, in the worst case, in $\Omega(n\alpha(k) + km + k^2)$ and $O(n\alpha(k) + km\alpha(k) + k^2)$, where $\alpha(k)$ denotes the pseudo-inverse of the Ackermann function.*

Proof. By Property 2, the shadow cast on a plane Π by a polygonal light source in the presence of one polytope is a convex polygon. Furthermore, if the light source has m edges and the polyhedron has n_i edges, the shadow region in Π has $O(n_i + m)$ edges. By Property 1, the shadow region in the presence of k polytopes of total complexity n is thus the union of k convex polygons of total complexity $O(n + km)$, which has complexity $O((n + km)\alpha(k) + k^2)$ [AS97].

For the proof of the stated lower bound consider the following collection of examples. In all constructions the shadow plane Π is the plane $z = 0$.

$\Omega(k^2)$ **example.** Refer to Figure 13.1. We consider a point light source at a height z (large enough) and a grid consisting of k thin horizontal and parallel rectangles at height $z = 1$ together with k other thin horizontal and parallel rectangles at height $z = 2$. They form a grid of shadow on plane Π which has size $O(k^2)$.

$\Omega(k\alpha(k))$ **example.** Refer to Figure 13.2. Again, the light source is a point with large positive z -coordinate. We consider a set of k line segments in plane $z = 1$ (with positive y coordinates) having, in that plane, an upper envelope

²²A *maximal free line segment* is a segment that intersects the interior of no object and whose endpoints lie on some object or at infinity.

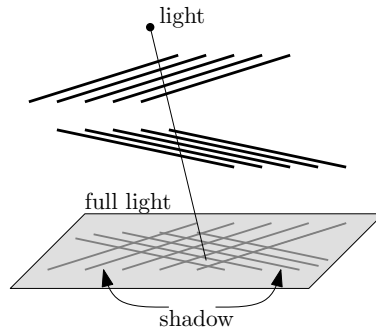


FIG. 13.1 – $\Omega(k^2)$ lower bound.

of size $\Omega(k\alpha(k))$ [WS88]. We transform each line segment into a trapezoid linking it to its projection on the $y = 0$ line (in plane $z = 1$). We get a set of trapezoids whose shadow, in plane $z = 0$, for a point light source at large enough z is basically the upper envelope of the segments. Note that the trapezoids can easily be made disjoint by placing them in different horizontal planes very close to plane $z = 1$.

$\Omega(n\alpha(k))$ **example.** Refer to Figure 13.3. First modify the above $\Omega(k\alpha(k))$ example such that the left “vertical” side of each trapezoid has slope γ and the right “vertical” side has slope $-\gamma$, for some γ large enough. Now, by some suitable scaling, we make all slopes of the vertical walls strictly smaller than $\frac{k\pi}{n}$. Assemble $\frac{n}{k}$ copies of the previous construction into a large regular $\frac{n}{k}$ -gon where each side is, in fact, a $k\alpha(k)$ upper envelope. Finally, this construction can be seen as k convex $3\frac{n}{k}$ -gons by connecting all of the $\frac{n}{k}$ copies of the same trapezoid by extending their walls.

This set of k convex $3\frac{n}{k}$ -gons, embedded in different horizontal planes close to the $z = 1$ plane, engender, in the presence of a point light source at large enough z , shadows of complexity $n\alpha(k)$.

$\Omega(mk)$ **example.** Refer to Figure 13.4. We use a horizontal m -gon as light source and a thin rectangle as obstacle. Then the shadow has $\Omega(m)$ size. Using multiple copies of the obstacle such that the different shadows are disjoint easily gives an $\Omega(mk)$ example. \square

There is still a small gap between the $\Omega(k^2 + mk + n\alpha(k))$ lower bound and the $O(k^2 + mk\alpha(k) + n\alpha(k))$ upper bound. In fact we conjecture that the lower bound is tight; the shadow of the different obstacles have some similarity with homothetic projections of the light, and the union of k convex homothetic m -gons is $\Theta(mk)$ since two convex homothetic polygons intersect in at most two points [KLPS86].

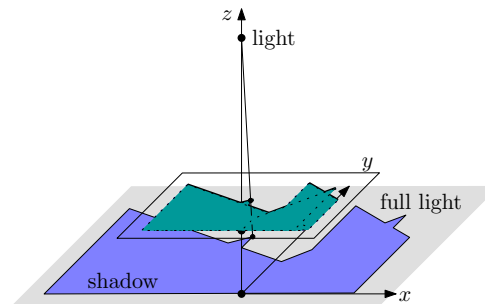


FIG. 13.2 – $\Omega(k\alpha(k))$ lower bound.

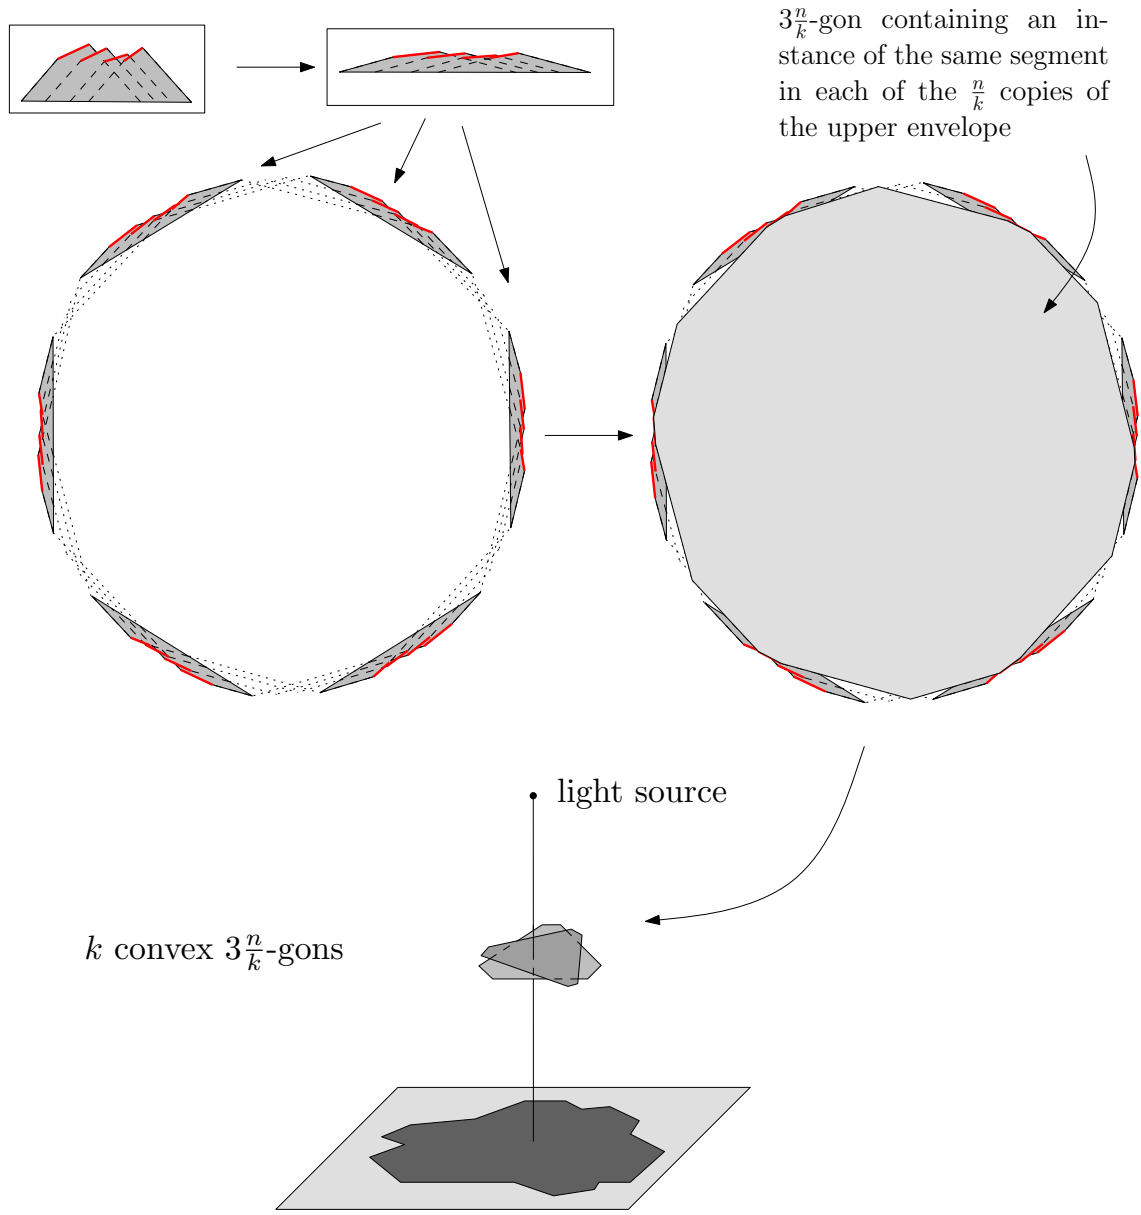


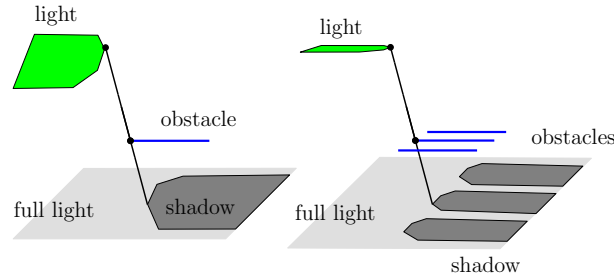
FIG. 13.3 – $\Omega(n\alpha(k))$ lower bound.

13.4 Upper bounds

In this section we prove the following two upper bounds on the complexity of the umbra cast by a segment light source or polygonal light source(s) on a plane.

Theorem 13.2. *The complexity of the umbra cast on a plane by one segment light source in the presence of k disjoint polytopes of total complexity n is $O(nk^3)$.*

Theorem 13.3. *The complexity of the umbra cast on a plane by a set of k disjoint polytopes of total complexity n , some of which are light sources, is $O(n^3k^3)$.*

FIG. 13.4 – $\Omega(mk)$ lower bound.

13.4.1 The umbra cast by a segment light source

We will actually prove an upper bound on the complexity of the shadow arrangement which yields the same bound for the complexity of the umbra. Notice that, in the case of a single segment light source, the EEE -surfaces $\sigma = \langle e_1, e_2, e_3 \rangle$ and EV -surfaces $\sigma = \langle e, v \rangle$ that contribute to the shadow arrangement are such that either one of e_1, e_2 or e_3 is the segment light source and v is one of its endpoints.

We prove Theorem 13.2 by considering a plane rotating about the line supporting the segment light source. First, if the segment light source, s , is not parallel to the shadow plane Π , we apply a projective transformation to the scene, sending to infinity the point of intersection between the line containing s and plane Π ; this does not change the complexity of the shadow arrangement. We thus assume in the rest of this section that the segment light source is parallel to Π . The sweep plane, denoted π , intersects the shadow plane in a line; we will say that, throughout the sweep, this line moves from left to right.

We start with two preliminary lemmas concerning the intersections that one of these sweep planes can have with the shadow arrangement.

Let α be the conic that is the intersection of the shadow plane Π and the EEE -surface $\sigma = \langle s, s_1, s_2 \rangle$ where s is the segment light source, s_1 and s_2 are two other segments, and let π be a plane containing the light source s and intersecting Π .

Lemma 13.4. *If $s, s_1,$ and s_2 are pairwise skew then π intersects α in at most a single point.*

Proof. First, since s is parallel to Π , any transversal to s and to a point in $\alpha \cap \pi$ lies in plane π . If the intersection between π and s_1 or s_2 is empty, there exists no line tangent to the three segments in π and thus π does not cross α . Otherwise, both segments intersect π in two points p_1 and p_2 . The line defined by p_1 and p_2 is the only line that is possibly a transversal to the three segments in π (it may not intersect the segment s). Hence there exists at most one transversal in π which defines exactly one point on α . \square

Notice that the preceding lemma implies that the conic arcs are monotonic in the direction of the sweep (orthogonal to s).

For the rest of the proof we consider an arrangement A of arcs of conics which contains the shadow arrangement. We will establish an upper bound of $O(nk^3)$ on the complexity of A which will yield the bound of Theorem 13.2. The arrangement A consists of the intersection of Π with (i) those lines that are transversal to the light source s , and the edges s_1 and s_2 of two other polytopes and that do not intersect the interior of these polytopes (the connected components of these lines form patches of EEE -surfaces) and (ii) those lines that are transversal to a vertex and an edge of two polytopes, one of which is the segment light source, and that do not intersect the interior of these polytopes (the connected components of these lines form patches of EV -surfaces).

We now count the number of crossings between an instance of the sweep plane π and the arcs in A .

Lemma 13.5. *The plane π properly intersects at most $O(k^2)$ arcs of A .*

Proof. The arcs of A are defined as the intersection with Π of lines ℓ which are (i) transversal to the segment light source s and tangent to two polytopes, (ii) transversal to an endpoint of s and tangent to another polytope, or (iii) transversal to s and to a polytope vertex.

An instance π of the sweep plane never properly intersects an arc of type (iii) (since such an arc is either included in π or does not intersect it). Now, if π intersects an arc of one of the other two types, then π contains the corresponding line ℓ , which is tangent to two polygons of $\mathcal{P} \cap \pi$. Since $\mathcal{P} \cap \pi$ consists of at most k disjoint convex polygons, there are $O(k^2)$ lines in π that are tangent to two polygons of $\mathcal{P} \cap \pi$, hence the result. \square

Proof. [of Theorem 13.2] We consider an orthogonal frame in plane Π whose *vertical* axis is parallel to the segment light source s ; the other axis is called *horizontal*.

We first show that the number of proper intersection points between arcs of A is $O(k^2)$ times the number of arcs. It follows from Lemma 13.4 (along with a simple argument in the case that the three segments are not pairwise skew) that any arc of A is either horizontally monotone or is a vertical line segment. Consider an arc α_0 and its rightmost endpoint p (anyone if α_0 is vertical). We charge to α_0 all points of intersection involving α_0 and all arcs whose rightmost endpoints are strictly to the right of p . Any arc properly intersects α_0 in at most $O(1)$ points so the number of intersection points charged to α_0 is bounded by the number of arcs properly intersected by the sweep plane containing p . By Lemma 13.5, there are at most $O(k^2)$ such arcs. Thus, each arc is charged at most $O(k^2)$ times.

We now bound the number of arcs (and thus the number of arc endpoints) generating A . Each arc corresponds either to a patch of an EV or EEE surface. Consider first the EV-surfaces. Since either the edge or the vertex is on the light source, there are at most $O(n)$ such surfaces.

Now consider the arcs generated by EEE-surfaces. Let n_i be the number of vertices of polytope P_i , $1 \leq i \leq k$. The number of EEE-surfaces involving the light source and edges from polytopes P_i and P_j is $O(n_i + n_j)$ [BDD⁺07, Corollary 2.6] (or [BDD⁺05, Corollary 9]). Then, $\sum_{1 \leq i < j \leq k} O(n_i + n_j) = O(nk)$.

There are at most $O(nk)$ arcs generating A . Since each arc is charged with at most $O(k^2)$ intersection points, there are at most $O(nk^3)$ intersection points. The total complexity of the shadow arrangement, and thus of the umbra, is then $O(nk^3)$. \square

13.4.2 The umbra cast by polygonal light sources

To prove Theorem 13.3 we consider an arrangement B of arcs of conics that, as in the previous section, contains the shadow arrangement. This arrangement B consists of the intersections of Π with (i) the lines that are transversal to a vertex and an edge of two polytopes and that do not intersect the interior of these polytopes (the connected components of these lines form patches of EV-surfaces) (ii) the lines that are transversal to edges of three polytopes and that do not intersect the interior of these polytopes (the connected components of these lines form patches of EEE-surfaces). Notice that B may contain arcs generated by surfaces that do not intersect the light source or possibly by surfaces that intersect the interior of other polytopes in the scene. We will establish a $O(n^3k^3)$ upper bound on the complexity of B which yields the same bound for the complexity of the umbra.

Lemma 13.6. *Any line L in Π properly intersects at most $O(nk^2)$ arcs of B .*

Proof. An intersection point between L and B corresponds to a line transversal which belongs to an EV or EEE surface. Consider first EV-surfaces. The line transversal lies in a plane which contains L and a vertex, say v , of one of the polytopes. There exist $O(n)$ such planes and in each of them there are at most $O(k)$ lines through v that are tangent to a polytope (since we only consider proper intersections between L and the arcs of B). Thus there are at most $O(nk)$ points on $L \cap A$ which correspond to lines in EV-surfaces.

Now we consider EEE-surfaces. Let n_i be the number of vertices of polytope P_i , for $1 \leq i \leq k$. The number of EEE-surfaces generated by three edges of polytopes P_i , P_j and P_l , not intersecting the interior of P_i , P_j and P_l , and that intersect L is $O(n_i + n_j + n_l)$ [BDD⁺05, Main Lemma]. Since $\sum_{1 \leq i < j < l \leq k} O(n_i + n_j + n_l) = O(nk^2)$, there are at most $O(nk + nk^2) = O(nk^2)$ arcs of B which intersect the line L on Π . \square

Proof. [of Theorem 13.3] Here, we introduce an arbitrary coordinate frame Oxy in the plane Π . We call Ox the horizontal axis and Oy the vertical axis.

As in the proof of Theorem 13.2, we first show that the number of intersection points between arcs of B is $O(nk^2)$ times the number of conic arcs. We first break all conic arcs into maximal horizontally monotone pieces. This increases the number of arcs only by a constant factor. Consider a conic arc α_0 and its rightmost endpoint p

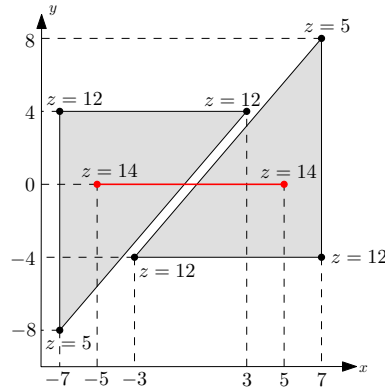


FIG. 13.5 – Two triangles and a segment light source (viewed from above) that cast 4 connected components of umbra on the plane $z = 0$.

along Ox (any endpoint if the arc is vertical). We charge to α_0 all points of intersection between α_0 and all conic arcs whose rightmost endpoints are strictly to the right of p . Any arc properly intersects α_0 in at most $O(1)$ points so the number of intersection points charged to α_0 is bounded by the number of arcs that are properly intersected by a vertical line in Π and containing p . By Lemma 13.6, there are at most $O(nk^2)$ such arcs. Thus, each arc is charged at most $O(nk^2)$ times.

We now bound the number of arcs (and thus the number of arc endpoints) generating B . Let n_i be the number of vertices of polytope P_i , $1 \leq i \leq k$ and e an edge. The number of EEE-surfaces pertinent to B and involving e and edges from polytopes P_i and P_j is $O(n_i + n_j)$ [BDD⁺05, Corollary 9]. Thus, for each edge e , there are, at most, $\sum_{1 \leq i < j \leq k} O(n_i + n_j) = O(nk)$ EEE-surfaces having e as a generating segment. Furthermore, the number of EV-surfaces involving edge e or one of its vertices is $O(n)$. Since there exist n edges, the total number of arcs in B is therefore $O(n^2k)$.

In conclusion, there are at most $O(n^2k)$ arcs generating B , each of them charged with at most $O(nk^2)$ intersection points, hence there are at most $O(n^3k^3)$ intersection points. The total complexity of the B and, thus of the umbra, is $O(n^3k^3)$. \square

13.5 Lower bounds

In this section we present several lower bounds on the complexity of the umbra.

13.5.1 The umbra cast by a segment light source

Here we concentrate on the umbra cast by a segment light source in the presence of various configurations of obstacles.

Theorem 13.7. *A segment light source and two triangles may cast, on a plane, four connected components of umbra.*

Proof. Consider the following scene consisting of a segment light source, s , two triangles, T_1 and T_2 , and a shadow plane, Π , the horizontal plane of equation $z = 0$; see Figure 13.5 and Figure 13.6.

Figure 13.7 shows a superset of the shadow arrangement generated by this configuration (the arrangement A defined in Section 13.4). Although it can be shown that the four shaded regions in the figure are exactly the umbra, we will simply argue here that there are at least four connected components. We do this by illustrating four segments in the umbra and then arguing that they are each in different connected components.

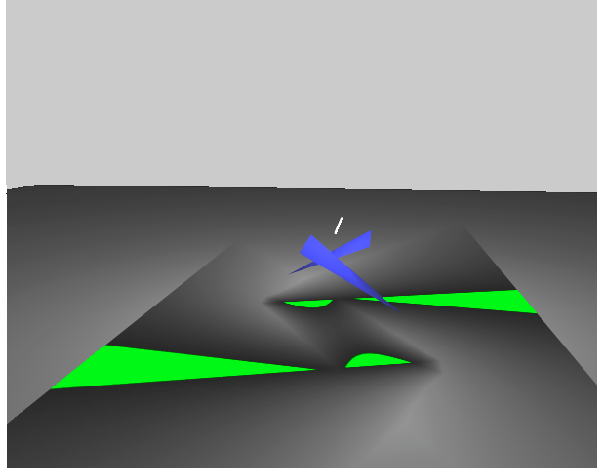


FIG. 13.6 – The scene rendered with the ray tracer OpenRT (the umbra is in light grey); courtesy of Andreas Dietrich.

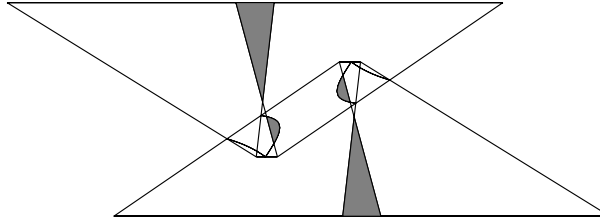


FIG. 13.7 – Superset of the shadow arrangement on plane Π . The four shaded regions are the regions of umbra.

The idea is to consider a series of planes rotating about the segment light source and the intersections of those planes with the two triangles and the shadow plane; Figure 13.8 shows such a sequence. We then examine the umbra in those planes by considering the relevant bitangents.

Let P_+ be one such plane (containing s) and going through the point $(0, 7, 0)$ and L_+ the intersection of P_+ and Π . Figure 13.8(b) shows the segment s , the intersections between P_+ and the two triangles T_1 and T_2 , L_+ and four bitangents that together define the umbra on L_+ . Consider the two segments R_1^+ and R_2^+ as shown in Figure 13.8(b). It is easy to see, by examining the bitangents, that R_1^+ and R_2^+ are in the umbra. Hence there are two segments of umbra on the line L_+ . We obtain two other segments, R_1^- and R_2^- , by taking the symmetric plane P_- with respect to the xz -plane (through point $(0, -7, 0)$ and whose intersections with the scene is shown on Figure 13.8(d)).

Now, we show that the four segments R_1^+ , R_1^- , R_2^+ and R_2^- lie in different connected components of umbra. In order to prove this result, we exhibit two lines on Π which contain no point in the umbra and separate the four segments as shown in Figure 13.9.

First consider the plane $y = 0$ containing the light segment s and orthogonal to the shadow plane Π . This plane intersects Π in a line, δ_1 , as shown in Figure 13.9, and separates R_1^+ and R_2^+ from R_1^- and R_2^- since P_+ and P_- are symmetric with the plane $y = 0$. To show that δ_1 contains no point of the umbra, consider the intersection of the $y = 0$ plane with the segment s and the two triangles T_1 and T_2 ; see Figure 13.8(c). A study of the bitangents reveals that no point of δ_1 lies in the umbra.

Now consider the plane orthogonal to Π , parallel to the two triangle hypotenuses and going through the midpoint of s . Let δ_2 be the intersection of this plane with Π ; see Figure 13.9. Elementary computations show that the line δ_2 separates R_1^+ and R_2^- from R_1^- and R_2^+ . There can be no point of the umbra on δ_2 since the plane intersects the light source but not the triangles (see Figure 13.5). This completes the proof. \square

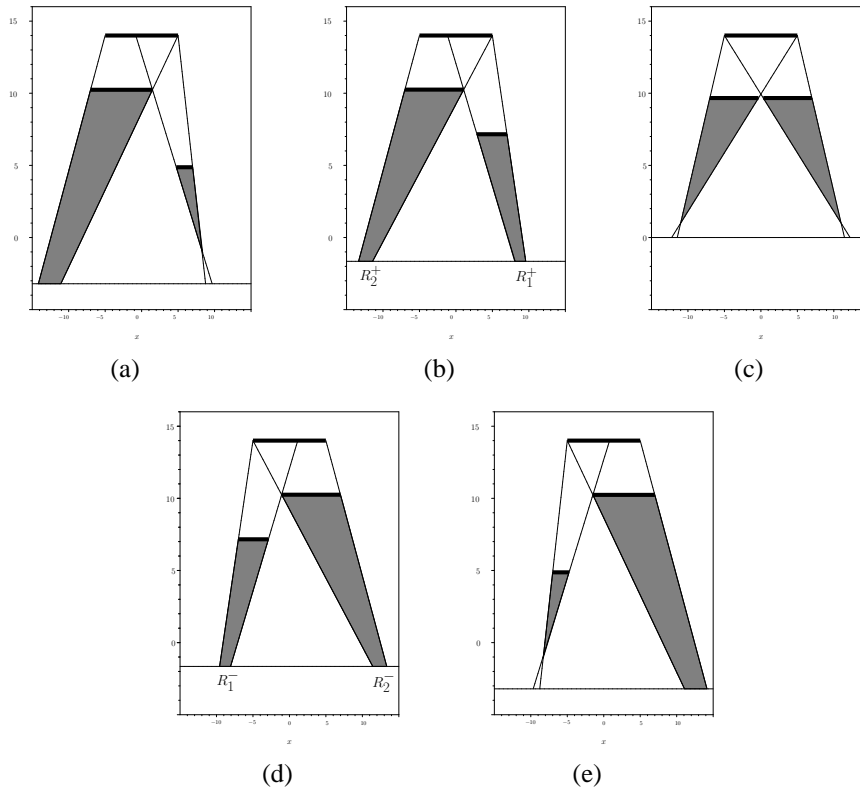


FIG. 13.8 – Views in the sweep plane with bitangents that define the umbra. The number of components of umbra in the intersection of the sweep plane and the plane $z = 0$ is : (a) one (sweep plane through $(0, 10, 0)$), (b) two (sweep plane P_+ through $(0, 7, 0)$), (c) zero (sweep plane $y = 0$), (d) two (sweep plane P_- through $(0, -7, 0)$), (e) one (sweep plane through $(0, -10, 0)$).

Note that the line supporting s and the lines supporting the triangle hypotenuses are pairwise skew and not all parallel to a same plane. Thus the corresponding EEE-surface is a hyperboloid of one sheet which intersects Π in a hyperbola. We determine the equation of this hyperbolic curve to be $41y^2 - 52xy + 928 = 0$. This curve admits two asymptotes which contain no point in the umbra and which separate the connected components of umbra. One of these asymptotes is δ_1 and we could have chosen the other to be δ_2 .

Note also that in our example, the light source is parallel to the shadow plane, and there are also many symmetries. None of this is critical ; the example can be perturbed and the result still holds.

We know prove a lower bound for fat polytopes, polytopes whose aspect ratios are bounded from below by a positive constant when n goes to infinity.

Theorem 13.8. *The umbra cast on a plane by one segment light source in the presence of two fat disjoint polytopes of total complexity n can have $\Omega(n)$ connected components.*

Proof. Our lower-bound example consists of one segment light source s_1 , a polytope Q_2 of size $O(n)$, and another polytope, Q_3 , of constant size. Refer to Figure 13.10.

First we consider three skew lines $l_1 \supset s_1$, l_2 and l_3 and $\sigma = \langle s_1, l_2, l_3 \rangle$ the quadric patch(es) consisting of the lines stabbing s_1 , l_2 and l_3 . In the shadow plane Π , by adding suitable half planes P_2 and P_3 as obstacles limited by the lines l_2 and l_3 , we obtain α , a single conic arc of $\sigma \cap \Pi$, bounding the umbra where the umbra is on the concave side of α (Figure 13.10-left).

We now consider p_1 , one of the endpoints of s_1 , and α_2, α_3 , the intersections of planes P_2, P_3 with the cone of apex p_1 and base α (Figure 13.10-center).

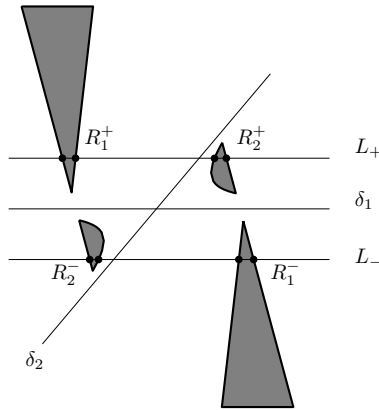


FIG. 13.9 – The four connected components of umbra and the four lines used in the proof of Theorem 13.7.

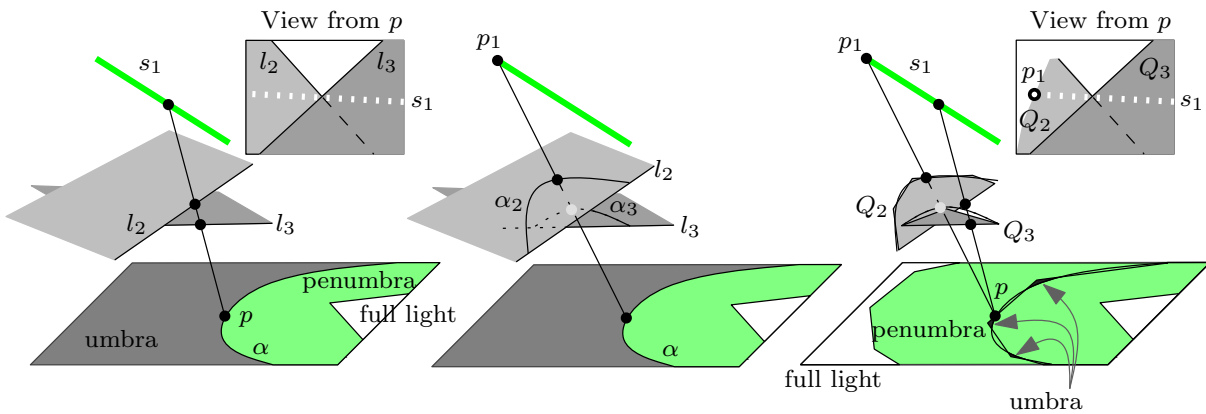


FIG. 13.10 – $\Omega(n)$ lower bound.

Next, we reduce the obstacles P_2 and P_3 to convex polygons Q_2 and Q_3 by limiting them by a polygonal approximation of α_2 and α_3 such that Q_3 remains within (*i.e.*, on the convex side of) α_3 and Q_2 intersects α_2 n times (Figure 13.10-right). The umbra cast by s_1 on Π in the presence of Q_2 and Q_3 then consists of n connected components that are the intersection of the concave region outside α and the convex polygon that is the intersection of the cone of apex p_1 and base Q_2 with the plane Π .

Note that the polygons Q_2 and Q_3 are fat since Q_2 consists of a segment and of an approximation of a conic and Q_3 is of constant size. Finally, polygons Q_2 and Q_3 can be trivially transformed into fat polytopes without changing the umbra. \square

Theorem 13.9. *The umbra cast on a plane by one segment light source in the presence of k disjoint polytopes of total complexity n can have $\Omega(nk^2)$ connected components.*

Proof. Consider three non-parallel segments s_1, l_2 , and l_3 all parallel to the shadow plane Π and planes $P_2 \supset l_2$ and $P_3 \supset l_3$ parallel to Π , refer to Figure 13.11. The surface $\langle s_1, l_2, l_3 \rangle$ intersects Π in a conic arc α .

Now consider the following setup : s_1 is the light source ; P_2 has k narrow rectangular holes (or slits) parallel and arbitrary close to l_2 ; similarly P_3 has k slits parallel and arbitrary close to l_3 . (A plane with k such slits can be modelled by $O(k)$ rectangles.) Each pair of slits, s_2 from P_2 and s_3 from P_3 , together with the light source s_1 induce a piece of penumbra in Π that is essentially a thickened copy of the conic arc α .

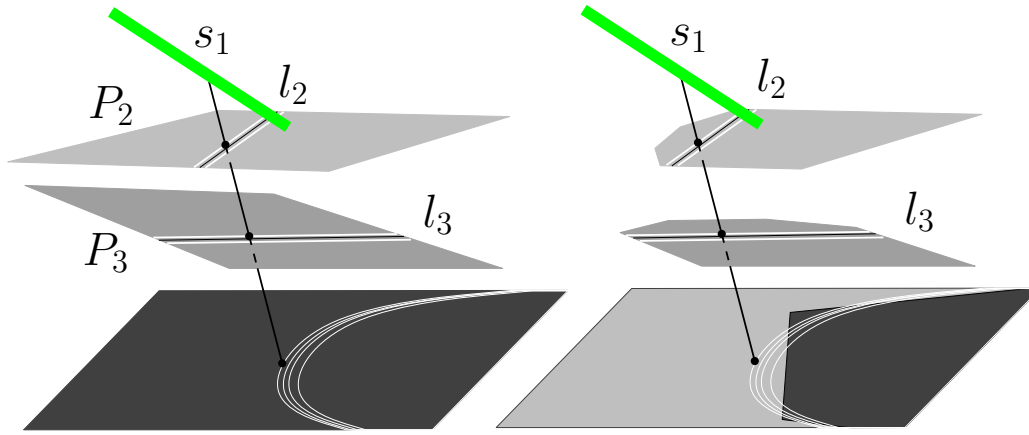


FIG. 13.11 – $\Omega(nk^2)$ lower bound.

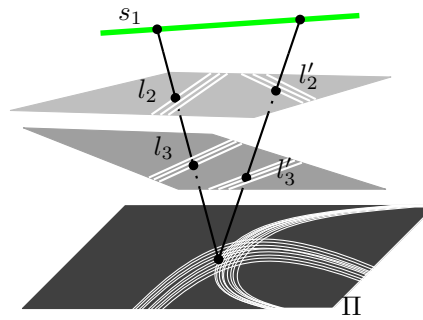


FIG. 13.12 – $\Omega(k^4)$ lower bound.

We thus get that the umbra covers the whole plane Π except for k^2 curves of penumbra that are all close to α (see Figure 13.11-left).

Finally, we trim the two planes P_2 and P_3 , creating an n -sided convex polygon on Π such that the region outside this polygon is in light or penumbra and each edge intersects all the k^2 curves. The umbra then consists of nk^2 regions inside the convex polygon and between the k^2 conics (see Figure 13.11-right). Note that the $O(k)$ convex obstacles can each be transformed into a polytope by the addition of a single vertex without changing the umbra. \square

Theorem 13.10. *The umbra cast on a plane by a segment light source in the presence of k disjoint polytopes can have $\Omega(k^4)$ connected components.*

Proof. Refer to Figure 13.12. As in the previous lower-bound example, we create k^2 curves of penumbra using parallel thin holes. Making a second set of thin holes in each plane, we create a second family of curves of light and penumbra intersecting the first one. The umbra is now the complement of the union of these two sets of curves and it consists of $\Omega(k^4)$ connected components. \square

13.5.2 The umbra cast by a polygonal light source

Note that the lower bound of $\Omega(nk^2 + k^4)$ of Section 13.5.1 for a segment light source can easily be modified into a lower bound of $\Omega(nk^3 + k^6)$ in the case of a polygonal light source (by adding a third plane with $O(k)$ slits

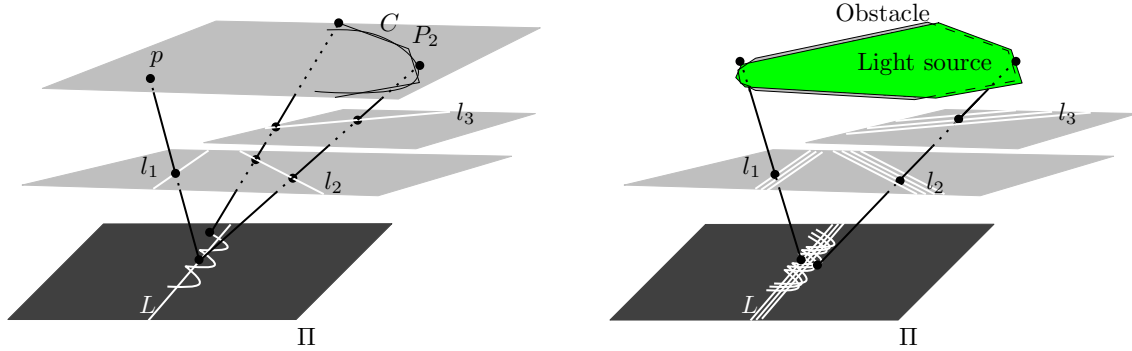


FIG. 13.13 – $\Omega(n^2 k^3)$ lower bound.

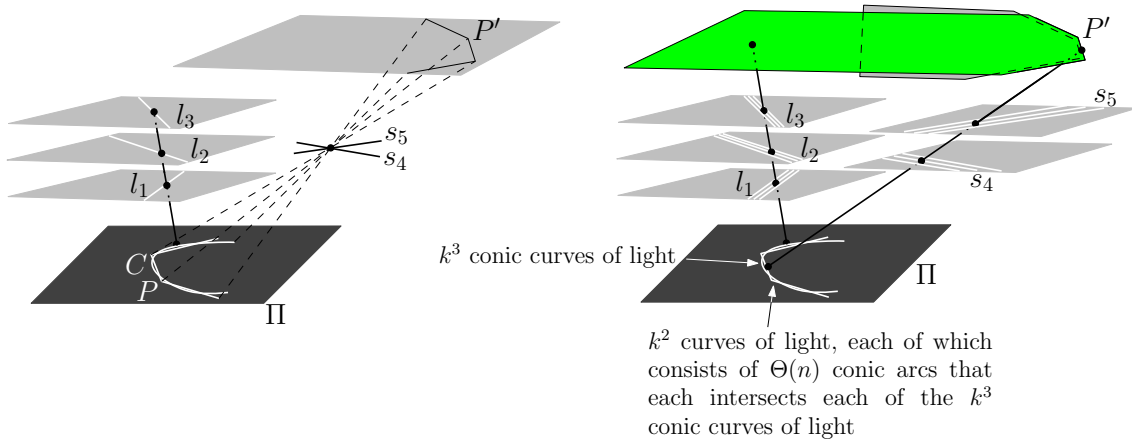


FIG. 13.14 – $\Omega(nk^5)$ lower bound.

and a big polygonal light source). We present here a lower bound of $\Omega(n^2 k^3 + nk^5)$ on the complexity of the umbra cast by a polygonal light source in the presence of k polygonal obstacles.

Theorem 13.11. *The umbra cast on a plane by one polygonal light source in the presence of k disjoint polytopes of total complexity n can have $\Omega(n^2 k^3)$ connected components.*

Proof. Refer to Figure 13.13. Let p be a point and P_1 a small n -gon light source very close to p . Add a n -gon obstacle very close to the light source so that the light source behaves like n point light sources (when viewed from the correct side).

Now consider a plane obstacle with k thin holes parallel to a line l_1 . This creates nk parallel thin lines of light on the shadow plane that can be made arbitrarily close to a line L (by having the k thin holes sufficiently close to each other and the n point light sources sufficiently close to each other). Note that by duplicating this construction (and thus with two polygonal light sources which behave as $2n$ point light sources) we get an arrangement of $2nk$ lines of light with $n^2 k^2$ connected components of umbra.

Now consider two lines l_2 and l_3 that together with L admit a quadric as line transversals. Cut this quadric by a plane and approximate (a piece C of) the resulting conic by a convex n -polyline, P_2 . The set of transversals to the boundary of this polyline with l_2 and l_3 defines a curve on the shadow plane that cuts L order n times. We define a light source as the convex hull of P_2 and put an obstacle very close to it so that the light source behaves as if the polyline P_2 was the light source (when viewed from the right region). Now, replacing l_2 and l_3 by two plane obstacles with order k thin holes close and parallel to l_2 and l_3 , respectively, we get k^2 curves of light, each

of which intersects order n times each of the nk lines of light close to L . This give $\Omega(n^2k^3)$ connected component of the umbra.

Note that the two light sources P_1 and P_2 can be merged into one by considering P_2 in the same plane as P_1 , by noticing that there are enough degrees of freedom on l_1 and l_2 so that the convex hull of P_1 and an arc of the conic contains C on its boundary. \square

Theorem 13.12. *The umbra cast on a plane by one polygonal light source in the presence of k disjoint polytopes of total complexity n can have $\Omega(nk^5)$ connected components.*

Proof. Refer to Figure 13.14. Consider three horizontal pairwise skew lines l_1, l_2, l_3 that lie above a horizontal plane Π and let C be the conic intersection of their common transversal with Π . Replace each of the l_i by a plane obstacle and k thin holes close to l_i and place a large (horizontal) light source S above these planes obstacles.

Consider now a n -gon P that intersects C order n times. Let s_4 and s_5 be two intersecting horizontal segments. Let P' be the symmetric of P with respect of the point of intersection between s_4 and s_5 . We consider P' as a light source and put an obstacle very close to it so that it behaves as a one-dimensional polygonal light source when viewed from C . This induces on the shadow plane a polyline of light that intersects C order n times.

Now perturb segments s_4 and s_5 so that they do not intersect and replace them by (horizontal) plane obstacles with k thin holes close and parallel to s_4 and s_5 , respectively. We hence get k^2 curves of light, each of which consists of order n conic arcs that each intersects C ; hence each of these k^2 curves of light intersects C order n times. By choosing the holes near l_1, l_2 and l_3 sufficiently close to each other, respectively, each of the k^2 curves of light close to P intersects $O(n)$ times each of the k^3 curves of light close to C . We hence get nk^5 connected components of umbra. \square

13.6 Conclusion

The purpose of this paper is to establish the complexity of the boundaries between the umbra, penumbra and fully-lit regions on a plane in a polyhedral scene consisting of k convex objects of total complexity n .

The results presented here constitute a first step toward understanding the intrinsic structure and complexity of the umbra in this setting. We have proved that if the light is reduced to one line segment, then the umbra may have $\Omega(nk^2 + k^4)$ connected components and $O(nk^3)$ complexity. We have also shown that a polygonal light source could generate an umbra with $\Omega(n^2k^3 + nk^5)$ connected components and $O(n^3k^3)$ complexity. In both cases these components of umbra are delimited by arcs of conics. These results prove that the umbra is intrinsically much more intricate than the boundary between full light and penumbra which is bounded by line segments and has complexity $O(n\alpha(k) + km\alpha(k) + k^2)$, where m is the complexity of the light source.

Our upper bounds, in fact, apply to the complexity of the arrangement of the curves where the derivative of the light intensity is discontinuous. These arrangements clearly include the boundary of the umbra, but also a lot of curves inside the penumbra that are not relevant to the umbra. Furthermore, our upper bound on the complexity of these arrangements is tight for a segment light source (see the full paper for details). This perhaps explains why our bounds on the complexity of the umbra are not tight. Notice, however, that we do have tight bounds for small k ($k = O(1)$) and for small n ($n = O(k)$).

Acknowledgments

Many thanks to Jeff Erickson for useful discussions during the early stages of this work.

Troisième partie

Géométrie algorithmique non linéaire sur les quadriques en trois dimensions

Chapitre 14

Near-optimal parameterization of the intersection of quadrics : I. The generic algorithm

Cet article a été accepté dans *Journal of Symbolic Computation* [DLLP07a]. (Voir également le rapport de recherche [DLLP05a].) Une version préliminaire à été publiée dans les proceedings du *19th ACM Annual Symposium on Computational Geometry* [DLLP03] et dans la thèse de L. Dupont [Dup04].

Abstract

We present the first exact and efficient algorithm for computing a parametric representation of the intersection of two quadrics in three-dimensional real space given by implicit equations with rational coefficients. The output functions parameterizing the intersection are rational functions whenever it is possible, which is the case when the intersection is not a smooth quartic (for example, a singular quartic, a cubic and a line, and two conics). Furthermore, the parameterization is near-optimal in the sense that the number of square roots appearing in the coefficients of these functions is minimal except in a small number of cases where there may be an extra square root. In addition, the algorithm is practical : a complete, robust and efficient C++ implementation is described in Lazard et al. [LPP06].

In Part I, we present an algorithm for computing a parameterization of the intersection of two arbitrary quadrics which we prove to be near-optimal in the generic, smooth quartic, case. Parts II and III treat the singular cases. We present in Part II the first classification of pencils of quadrics according to the real type of the intersection and we show how this classification can be used to efficiently determine the type of the real part of the intersection of two arbitrary quadrics. This classification is at the core of the design of our algorithms for computing near-optimal parameterizations of the real part of the intersection in all singular cases. We present these algorithms in Part III and give examples covering all the possible situations in terms of both the real type of intersection and the number and depth of square roots appearing in the coefficients.

14.1 Introduction

The simplest of all the curved surfaces, quadrics (i.e., algebraic surfaces of degree two), are fundamental geometric objects, arising in such diverse contexts as geometric modeling, statistical classification, pattern recognition, and computational geometry. Computing the intersection of two general quadrics is a fundamental problem and an exact parametric representation of the intersection is often desirable. For instance, it is at the basis of such complex geometric operations as computing convex hulls of quadric patches [HI95], arrangements

of sets of quadrics [BHK⁺05, MTT05, SW06, Wol02], and boundary representations of quadric-based solid models [KCF⁺04, Sar83].

Until recently, the only known general method for computing a parametric representation of the intersection between two arbitrary quadrics was due to J. Levin [Lev76, Lev79]. It is based on an analysis of the pencil generated by the two quadrics, i.e. the set of linear combinations of the two quadrics.

Though useful, Levin's method has serious limitations. When the intersection is singular or reducible, a parameterization by rational functions is known to exist, but Levin's pencil method fails to find it and generates a parameterization that involves the square root of some polynomial. In addition, when a floating point representation of numbers is used, Levin's method sometimes outputs results that are topologically wrong and it may even fail to produce any parameterization at all and crash. On the other hand a correct implementation using exact arithmetic is essentially out of reach because Levin's method introduces algebraic numbers of fairly high degree. A good indication of this impracticality is that even for the simple generic example of Section 14.8.2, an exact parametric form output by Levin's algorithm (computed by hand with Maple) fills up over 100 megabytes of space !

Over the years, Levin's seminal work has been extended and refined in several different directions. Wilf and Manor [WM93] use a classification of quadric intersections by the Segre characteristic (see [Bro06]) to drive the parameterization of the intersection by the pencil method. Recently, Wang, Goldman and Tu [WGT03] further improved the method by making it capable of computing structural information on the intersection and its various connected components and able to produce a parameterization by rational functions when it exists. Whether their refined algorithm is numerically robust is open to question.

Another method of algebraic flavor was introduced by Farouki, Neff and O'Connor [FNO89] when the intersection is degenerate. In such cases, using a combination of classical concepts (Segre characteristic) and algebraic tools (factorization of multivariate polynomials), the authors show that explicit information on the morphological type of the intersection curve can be reliably obtained. A notable feature of this method is that it can output an exact parameterization of the intersection in simple cases, when the input quadrics have rational coefficients. No implementation is however reported.

Rather than restricting the type of the intersection, others have sought to restrict the type of the input quadrics, taking advantage of the fact that geometric insights can then help compute the intersection curve [GM91, Mil87, MG95, SJ92, SJ94]. Specialized routines are devised to compute the intersection curve in each particular case. Even though such geometric approaches are numerically more stable than the algebraic ones, they are essentially limited to the class of so-called natural quadrics (i.e., the planes, right cones, circular cylinders and spheres) and planar intersections.

Perhaps the most interesting of the previously known algorithms for computing an explicit representation of the intersection of two arbitrary quadrics is the method of Wang, Joe and Goldman [WJG02]. This algebraic method is based on a birational mapping between the intersection curve and a plane cubic curve. The cubic curve is obtained by projection from a point lying on the intersection. Then the classification and parameterization of the intersection are obtained by invoking classical results on plane cubics. The authors claim that their algorithm is the first to produce a complete topological classification of the intersection (singularities, number and types of connected components, etc.). However, the computation of the center of projection uses (an enhanced version of) Levin's algorithm. Either floating point arithmetic is used and the point will in general not exactly lie on the curve, leading to possibly incorrect classification, or exact arithmetic is used and the parameterizations computed will involve algebraic numbers of very high degree, thereby limiting their practical value.

14.1.1 Contributions

In this series of papers, we present the first exact and efficient algorithm for computing a parametric representation of the intersection of two quadric surfaces in three-dimensional real space given by implicit equations with rational coefficients. As a side product of this algorithm, we also obtain the first classification of pencils of quadrics based on the type of the curve of intersection in real projective space.

Our algorithm (as well as its implementation [LPP06]) has the following main features :

- it computes an exact parameterization of the intersection of two quadrics with rational coefficients of arbitrary size ;
- it places no restriction of any kind on the type of the intersection or the type of the input quadrics ;
- it correctly identifies, separates and parameterizes all the connected components of the intersection and

gives all the information on the incidence between the components, that is where and how (e.g., tangentially or not) two components intersect ;

- the parameterization is rational when one exists ; otherwise the intersection is a smooth quartic and the parameterization involves the square root of a polynomial ;
- the parameterizations are either optimal in the degree of the extension of \mathbb{Q} on which their coefficients are defined or, in a small number of well-identified cases, involve one extra possibly unnecessary square root.

Note that our complete, robust and efficient C++ implementation [LPP06] of this algorithm, which uses arbitrary-precision integer arithmetic, can routinely compute parameterizations of the intersection of quadrics with input integer coefficients having ten digits in less than 40 milliseconds on a mainstream PC.

The above features imply in particular that the output parameterization of the intersection is almost as “simple” as possible, meaning that the parameterization is rational if one exists, and that the coefficients appearing in the parameterization are almost as rational as possible. This “simplicity” is, in itself, a key factor for making the parameterization process both feasible and efficient (by contrast, an implementation of Levin’s method using exact arithmetic is essentially out of reach). It is also crucial for the easy and efficient processing of parameterizations in further applications.

Formally, we prove the following.

Theorem 14.1. *In three-dimensional real space, given two quadrics in implicit form with rational coefficients, our algorithm first computes the type of their intersection in real projective space. If it is a smooth quartic, there does not exist any rational parameterization of the intersection and our algorithm computes a parameterization such that, in projective space, each coordinate belongs to $\mathbb{K}[\xi, \sqrt{\Delta}]$ (the ring of polynomials in ξ and $\sqrt{\Delta}$ with coefficients in \mathbb{K}), where ξ is the (real) parameter, $\Delta \in \mathbb{K}[\xi]$ is a polynomial in ξ , and \mathbb{K} is either the field of the rationals or an extension of \mathbb{Q} by the square root of an integer. If the intersection is not a smooth quartic, our algorithm computes a rational parameterization of each component of the intersection over a field \mathbb{K} of coefficients which is \mathbb{Q} or an extension of \mathbb{Q} of degree 2 or 4 ; this means that each projective coordinate of the component of the intersection is a polynomial in $\mathbb{K}[\xi]$.*

In all cases, either \mathbb{K} is a field of smallest possible degree²³ over which there exists such a parameterization or \mathbb{K} is an extension of such a smallest field by the square root of an integer. In the latter situation, testing if this extra square root is unnecessary and, if so, finding an optimal parameterization are equivalent to finding a rational point on a curve or a surface (which is computationally hard and can even be undecidable).

14.1.2 Overview

Due to the number of contributions and results of this work, this paper has been broken down into three parts. In Part I, we present a first and major improvement to Levin’s pencil method and the accompanying theoretical tools. This simple algorithm, referred to from now on as the “generic algorithm”, outputs a near-optimal parameterization when the intersection is a smooth quartic, i.e. the generic case. However, the generic algorithm ceases to be optimal (both from the point of view of the functions used in the parameterizations and the size of their coefficient field) in several singular situations. Parts II and III refine the generic algorithm by considering in turn all the possible types of intersection. In Part II, we present our classification of pencils of quadrics based on the type of their intersection in real projective space. We also show how to use this classification to compute efficiently the type of the real intersection. In Part III, we present optimal or near-optimal algorithms for each possible type of singular intersection.

Part I is organized as follows. In Section 14.2, we present basic definitions, notation and useful known results. Section 14.3 summarizes the ideas on which the pencil method of Levin for intersecting quadrics is based and discusses its shortcomings. In Section 14.4 we present our generic algorithm. Among the results of independent interest presented in this section are the almost always existence of a ruled quadric with rational coefficients in

²³Recall that, if \mathbb{K} is a field extension of \mathbb{Q} , the *degree* of the extension is defined as the dimension of \mathbb{K} as a vector space over \mathbb{Q} . For instance, if $\mathbb{Q}(\rho)$ is a field extension of \mathbb{Q} (distinct from \mathbb{Q}), then its degree is 2 since there is a one-to-one correspondence between any element $x \in \mathbb{Q}(\rho)$ and $(\alpha_1, \alpha_2) \in \mathbb{Q}^2$ such that $x = \alpha_1 + \alpha_2 \rho$. Similarly, if \mathbb{Q} and two field extensions $\mathbb{Q}(\rho)$ and $\mathbb{Q}(\rho')$ are pairwise distinct, then the degree of $\mathbb{Q}(\rho, \rho')$ is 4 since there is a one-to-one correspondence between any element $x \in \mathbb{Q}(\rho, \rho')$ and $(\alpha_1, \alpha_2, \alpha_3, \alpha_4) \in \mathbb{Q}^4$ such that $x = \alpha_1 + \alpha_2 \rho + \alpha_3 \rho' + \alpha_4 \rho \rho'$.

a pencil (proved in Section 14.5) and new parameterizations of ruled projective quadrics involving an optimal number of radicals in the worst case (a fact proved in Section 14.6). In Section 14.7, we prove the near-optimality of the output parameterization in the generic case, that is when the intersection curve is a smooth quartic, and show that the parameterization is optimal in the worst case, meaning that there are examples in which the possibly extra square root is indeed needed. Then, in Section 14.8, we give several examples and show the result of our implementation on these examples, before concluding.

14.2 Notation and preliminaries

In what follows, all the matrices considered are real square matrices. Given a real symmetric matrix S of size $n + 1$, the upper left submatrix of size n , denoted S_u , is called the *principal submatrix* of S and the determinant of S_u the *principal subdeterminant* of S .

We call a *quadric* associated to S the set

$$Q_S = \{\mathbf{x} \in \mathbb{P}^n \mid \mathbf{x}^T S \mathbf{x} = 0\},$$

where $\mathbb{P}^n = \mathbb{P}(\mathbb{R})^n$ denotes the real projective space of dimension n . (Note that every matrix of the form αS , where $\alpha \in \mathbb{R} \setminus \{0\}$, represents the same quadric Q_S .) When the ambient space is \mathbb{R}^n instead of $\mathbb{P}(\mathbb{R})^n$, the quadric is simply Q_S minus its points at infinity.

In the rest of this paper, geometric objects and parameterizations are assumed to live in projective space. For instance, a point of \mathbb{P}^3 has four coordinates. An object (point, line, plane, cone, quadric, etc.) given by its implicit equation(s) is said to be *rational* over a field \mathbb{K} if the coefficients of its equation(s) live in the field \mathbb{K} . Note that, when talking about parameterizations, some confusion can arise between two different notions : the rationality of the coefficients and the rationality of the defining functions (a quotient of two polynomial functions is often called a rational function). The meaning should be clear depending on the context.

Matrix S being symmetric, all of its eigenvalues are real. Let σ^+ and σ^- be the numbers of positive and negative eigenvalues of S , respectively. The *rank* of S is the sum of σ^+ and σ^- . We define the *inertia* of S and Q_S as the pair

$$(\max(\sigma^+, \sigma^-), \min(\sigma^+, \sigma^-)).$$

(Note that it is more usual to define the inertia as the pair (σ^+, σ^-) , but our definition, in a sense, reflects the fact that Q_S and Q_{-S} are one and the same quadric.) A matrix of inertia $(n, 0)$ is called *definite*. It is *positive definite* if $\sigma^- = 0$, *negative definite* otherwise. Matrix S and quadric Q_S are called *singular* if the determinant of S is zero ; otherwise they are called *nonsingular*.

The inertia of a quadric in \mathbb{P}^3 is a fundamental concept which somehow replaces the usual type of a quadric in \mathbb{R}^3 . For the convenience of the reader we recall in Table 14.1 the correspondence between inertias in \mathbb{P}^3 and types in \mathbb{R}^3 .

In \mathbb{P}^3 , any quadric not of inertia $(3, 1)$ is either a ruled surface or not a surface. Also, the quadrics of inertia $(3, 1)$ are the only ones with a strictly negative determinant. The nonsingular quadrics are those of rank 4, i.e. those of inertia $(4, 0)$, $(3, 1)$ and $(2, 2)$. Quadrics of inertia $(4, 0)$ are however empty of real points. A quadric of rank 3 is called a *cone*. The cone is said to be *real* if its inertia is $(2, 1)$. It is said to be *imaginary* otherwise, in which case its real projective locus is limited to its singular point. A quadric of rank 2 is a *pair of planes*. The pair of planes is real if its inertia is $(1, 1)$. It is called imaginary if its inertia is $(2, 0)$, in which case its real projective locus consists of its singular line, i.e. the line of intersection of the two planes. A quadric of inertia $(1, 0)$ is called a *double plane* and is necessarily real.

Two real symmetric matrices S and S' of size n are said to be *similar* if and only if there exists a nonsingular matrix P such that

$$S' = P^{-1} S P.$$

Note that two similar matrices have the same characteristic polynomial, and thus the same eigenvalues. Two matrices are said to be *congruent* or *projectively equivalent* if and only if there exists a nonsingular matrix P with real coefficients such that

$$S' = P^T S P.$$

Inertia of Q_S	Inertia of S_u	Euclidean canonical equation	Euclidean type of Q_S
(4, 0)	(3, 0)	$x^2 + y^2 + z^2 + 1$	\emptyset (imaginary ellipsoid)
(3, 1)	(3, 0)	$x^2 + y^2 + z^2 - 1$	ellipsoid
	(2, 1)	$x^2 + y^2 - z^2 + 1$	hyperboloid of two sheets
	(2, 0)	$x^2 + y^2 + z$	elliptic paraboloid
(3, 0)	(3, 0)	$x^2 + y^2 + z^2$	point
	(2, 0)	$x^2 + y^2 + 1$	\emptyset (imaginary elliptic cylinder)
(2, 2)	(2, 1)	$x^2 + y^2 - z^2 - 1$	hyperboloid of one sheet
	(1, 1)	$x^2 - y^2 + z$	hyperbolic paraboloid
(2, 1)	(2, 1)	$x^2 + y^2 - z^2$	cone
	(2, 0)	$x^2 + y^2 - 1$	elliptic cylinder
	(1, 1)	$x^2 - y^2 + 1$	hyperbolic cylinder
	(1, 0)	$x^2 + y$	parabolic cylinder
(2, 0)	(2, 0)	$x^2 + y^2$	line
	(1, 0)	$x^2 + 1$	\emptyset (imaginary parallel planes)
(1, 1)	(1, 1)	$x^2 - y^2$	intersecting planes
	(1, 0)	$x^2 - 1$	parallel planes
	(0, 0)	x	simple plane
(1, 0)	(1, 0)	x^2	double plane
	(0, 0)	1	\emptyset (double plane at infinity)

TAB. 14.1 – Correspondence between quadric inertias and Euclidean types.

The transformation sending S to S' is called a *congruence* transformation. Moreover if matrix P has rational coefficients, the congruence is said to be rational. Sylvester's Inertia Law asserts that the inertia is invariant under a congruence transformation [Lam73], i.e. S and S' have the same inertia. Note also that the determinant of S is invariant by a congruence transformation, up to a square factor (the square of the determinant of the transformation matrix).

Let S and T be two real symmetric matrices of the same size and let $R(\lambda, \mu) = \lambda S + \mu T$. The set

$$\{R(\lambda, \mu) \mid (\lambda, \mu) \in \mathbb{P}^1\}$$

is called the *pencil* of matrices generated by S and T . For the sake of simplicity, we sometimes write a member of the pencil $R(\lambda) = \lambda S - T$, $\lambda \in \overline{\mathbb{R}} = \mathbb{R} \cup \{\infty\}$. Associated to it is a pencil of quadrics $\{Q_{R(\lambda, \mu)} \mid (\lambda, \mu) \in \mathbb{P}^1\}$. Recall that the intersection of two distinct quadrics of a pencil is independent of the choice of the two quadrics. We call the binary form

$$\mathcal{D}(\lambda, \mu) = \det R(\lambda, \mu)$$

the *characteristic polynomial* of the pencil.

14.3 Levin's pencil method

Since our solution to quadric surface intersection builds upon the pencil method of [Lev76, Lev79], we start by recalling the main steps of his algorithm for computing a parameterized representation of the intersection of two distinct implicit quadrics Q_S and Q_T of \mathbb{R}^3 . Starting from this short description, we then identify where this algorithm introduces high-degree algebraic numbers and why this is a problem.

The high-level idea behind Levin's algorithm is this : if (say) Q_S is of some "good" type, then Q_S admits a parameterization which is linear in one of its parameters and plugging this parameterization in the implicit equation of Q_T yields a degree 2 equation in one of the parameters (instead of a degree 4 equation) which can be

quadric	canonical equation ($a, b > 0$)	parameterization $\mathbf{X} = [x, y, z], u, v \in \mathbb{R}$
simple plane	$x = 0$	$\mathbf{X}(u, v) = [0, u, v]$
double plane	$x^2 = 0$	$\mathbf{X}(u, v) = [0, u, v]$
parallel planes	$ax^2 = 1$	$\mathbf{X}(u, v) = [\frac{1}{\sqrt{a}}, u, v], \mathbf{X}(u, v) = [-\frac{1}{\sqrt{a}}, u, v]$
intersecting planes	$ax^2 - by^2 = 0$	$\mathbf{X}(u, v) = [\frac{u}{\sqrt{a}}, \frac{u}{\sqrt{b}}, v], \mathbf{X}(u, v) = [\frac{u}{\sqrt{a}}, -\frac{u}{\sqrt{b}}, v]$
hyperbolic paraboloid	$ax^2 - by^2 - z = 0$	$\mathbf{X}(u, v) = [\frac{u+v}{2\sqrt{a}}, \frac{u-v}{2\sqrt{b}}, uv]$
parabolic cylinder	$ax^2 - y = 0$	$\mathbf{X}(u, v) = [u, au^2, v]$
hyperbolic cylinder	$ax^2 - by^2 = 1$	$\mathbf{X}(u, v) = [\frac{1}{2\sqrt{a}}(u + \frac{1}{u}), \frac{1}{2\sqrt{b}}(u + \frac{1}{u}), v]$

TAB. 14.2 – Parameterizations of canonical simple ruled quadrics [Lev76].

easily solved to get a parametric representation of $Q_S \cap Q_T$. When neither Q_S nor Q_T has a “good” type, then one can find a quadric Q_R of “good” type in the pencil generated by Q_S and Q_T , and we are back to the previous case replacing Q_S by Q_R .

The definition of a “good” type is embodied in Levin’s notion of simple ruled quadric²⁴ and the existence of such a quadric Q_R is Levin’s key result :

Theorem 14.2 ([Lev76]). *The pencil generated by any two distinct quadrics contains at least one simple ruled quadric, i.e., one of the quadrics listed in Table 14.2, or the empty set.*

In more details, Levin’s method is as follows.

1. Find a simple ruled quadric in the pencil $\{Q_{R(\lambda)=\lambda S-T} \mid \lambda \in \overline{\mathbb{R}}\}$ generated by Q_S and Q_T , or report an empty intersection. Since simple ruled quadrics have a vanishing principal subdeterminant, this is achieved by searching for a $\lambda_0 \in \overline{\mathbb{R}}$ such that $\det(R_u(\lambda_0)) = 0$ and $Q_R = Q_{R(\lambda_0)}$ is simple ruled ; by Theorem 14.2, such a quadric exists or the pencil contains the empty set. Assume, for the sake of simplicity, that the intersection is not empty and that Q_R and Q_S are distinct. Then $Q_S \cap Q_T = Q_S \cap Q_R$.
2. Determine the orthonormal transformation matrix P_u which sends R_u in diagonal form by computing the eigenvalues and the normalized eigenvectors of R_u . Deduce the transformation matrix P which sends Q_R into canonical form. In the orthonormal frame in which it is canonical, Q_R admits one of the parameterizations \mathbf{X} of Table 14.2.
3. Compute the matrix $S' = P^T S P$ of the quadric Q_S in the canonical frame of Q_R and consider the equation

$$\mathbf{X}^T S' \mathbf{X} = a(u)v^2 + b(u)v + c(u) = 0, \quad (14.1)$$

where \mathbf{X} has been augmented by a fourth coordinate set to 1. (The parameterizations of Table 14.2 are such that $a(u), b(u)$ and $c(u)$ are polynomials of degree at most four in u .)

Solve (14.1) for v in terms of u and determine the corresponding domain of validity of u on which the solutions are defined, i.e., the set of u such that $\Delta(u) = b^2(u) - 4a(u)c(u) \geq 0$. Substituting v by its expression in terms of u in \mathbf{X} , we have a parameterization of $Q_S \cap Q_T = Q_S \cap Q_R$ in the orthonormal coordinate system in which Q_R is canonical.

4. Output $P\mathbf{X}(u)$, the parameterized equation of $Q_S \cap Q_T$ in the global coordinate frame, and the domain of $u \in \mathbb{R}$ on which it is valid.

This method is very nice and powerful since it gives an explicit representation of the intersection of two general quadrics. However, it is far from being ideal from the point of view of precision and robustness since it introduces non-rational numbers at several different places. Thus, if a floating point representation of numbers is used, the result may be wrong (geometrically and topologically) or, worse, the program may crash (especially in Step 1

²⁴In [Lev76, Lev79], Levin refers to these quadrics as to nonelliptic paras.

inertia of S	canonical equation ($a, b, c, d > 0$)	parameterization $\mathbf{X} = [x, y, z, w]$
(4, 0)	$ax^2 + by^2 + cz^2 + dw^2 = 0$	$Q_S = \emptyset$
(3, 0)	$ax^2 + by^2 + cz^2 = 0$	Q_S is point $(0, 0, 0, 1)$
(2, 2)	$ax^2 + by^2 - cz^2 - dw^2 = 0$	$\mathbf{X} = [\frac{ut+avs}{a}, \frac{us-bvt}{b}, \frac{ut-avs}{\sqrt{ac}}, \frac{us+bvt}{\sqrt{bd}}], (u, v), (s, t) \in \mathbb{P}^1$
(2, 1)	$ax^2 + by^2 - cz^2 = 0$	$\mathbf{X} = [uv, \frac{u^2-abv^2}{2b}, \frac{u^2+abv^2}{2\sqrt{bc}}, s], (u, v, s) \in \mathbb{P}^{*2}$
(2, 0)	$ax^2 + by^2 = 0$	$\mathbf{X} = [0, 0, u, v], (u, v) \in \mathbb{P}^1$
(1, 1)	$ax^2 - by^2 = 0$	$\mathbf{X}_1 = [u, \frac{\sqrt{ab}}{b}u, v, s], \mathbf{X}_2 = [u, -\frac{\sqrt{ab}}{b}u, v, s], (u, v, s) \in \mathbb{P}^2$
(1, 0)	$ax^2 = 0$	$\mathbf{X} = [0, u, v, s], (u, v, s) \in \mathbb{P}^2$

TAB. 14.3 – Parameterization of projective quadrics of inertia different from (3,1). In the parameterization of projective cones, \mathbb{P}^{*2} stands for the 2-dimensional real quasi-projective space defined as the quotient of $\mathbb{R}^3 \setminus \{0, 0, 0\}$ by the equivalence relation \sim where $(x, y, z) \sim (y_1, y_2, y_3)$ iff $\exists \lambda \in \mathbb{R} \setminus \{0\}$ such that $(x, y, z) = (\lambda y_1, \lambda y_2, \lambda^2 y_3)$.

when the type of the quadrics $Q_{R(\lambda_0)}$ are incorrectly computed). In theory, exact arithmetic would do, except that it would highly slow down the computations. In practice, however, a correct implementation using exact arithmetic seems out of reach because of the high degree of the algebraic numbers involved.

Let us examine more closely the potential sources of numerical instability in Levin’s algorithm.

- Step 1 : λ_0 is the root of a third degree polynomial with rational coefficients. In the worst case, it is thus expressed with nested radicals of depth two. Since determining if $Q_{R(\lambda_0)}$ is simple ruled involves computing its Euclidean type (not an easy task considering that $Q_{R(\lambda_0-\epsilon)}$ and $Q_{R(\lambda_0+\epsilon)}$ may be and often are of different types), this is probably the biggest source of non-robustness.
- Step 2 : Since Q_R is simple ruled, the characteristic polynomial of R_u is a degree three polynomial having zero as a root and whose coefficients are in the field extension $\mathbb{Q}(\lambda_0)$. Thus, the nonzero eigenvalues of R_u may involve nested radicals of depth three. Since the corresponding eigenvectors have to be normalized, the coefficients of the transformation matrix P are expressed with radicals of nesting depth four in the worst case.

Since the coefficients of the parameterization \mathbf{X} of Q_R are expressed as square roots of the coefficients of the canonical equation Q_{PTRP} (as in Table 14.2), the coefficients of the parameterization of $Q_S \cap Q_T$ can involve *nested radicals of depth five* in the worst case.

- Step 3 : Computing the domain of \mathbf{X} amounts to solving the fourth degree equation $\Delta(u) = 0$ whose coefficients are nested radicals of worst-case depth five in \mathbb{Q} .

Note that this worst-case picture is the generic case. Indeed, given two arbitrary quadrics with rational coefficients, the polynomial $\det(R_u(\lambda))$ will generically have no rational root (a consequence of Hilbert’s Irreducibility Theorem).

14.4 Generic algorithm

We now present a first but major improvement to Levin’s pencil method for computing parametric representations of the intersection of quadrics.

This so-called “generic algorithm” removes most of the sources of radicals in Levin’s algorithm. We prove in Section 14.7 that it is near-optimal in the generic, smooth quartic case. It is however not optimal for all the possible types of intersection and will need later refinements (see the comments in Section 14.9, and Parts II and III). But it is sufficiently simple, robust and efficient to be of interest to many.

We start by introducing the projective framework underlying our approach and stating the main theorem on which the generic approach rests. We then outline our algorithm and detail particular steps in ensuing sections.

From now on, all the input quadrics considered have their coefficients (i.e., the entries of the corresponding matrices) in \mathbb{Q} .

14.4.1 Key ingredients

The first ingredient of our approach is to work not just over \mathbb{R}^3 but over the real projective space \mathbb{P}^3 . Recall that, in projective space, quadrics are entirely characterized by their inertia (i.e., two quadrics with the same inertia are projectively equivalent), while in Euclidean space they are characterized by their inertia and the inertia of their principal submatrix.

In our algorithm, quadrics of inertia different from $(3, 1)$ (i.e., ruled quadrics) play the role of simple ruled quadrics in Levin's method. In Table 14.3, we present a new set of parameterizations of ruled projective quadrics that are both linear in one of their parameters and involve, in the worst case, a minimal number of square roots²⁵, which we prove in Section 14.6. That these parameterizations are faithful parameterizations of the projective quadrics (i.e., there is a one-to-one correspondence between the points of the quadric and the parameters) is proved in the appendix.

Another key ingredient of our approach is encapsulated in the following theorem, which mirrors, in the projective setting, Levin's theorem on the existence of ruled quadrics in a pencil.

Theorem 14.3. *In a pencil generated by any two distinct quadrics, the set \mathcal{S} of quadrics of inertia different from $(3, 1)$ is not empty. Furthermore, if no quadric in \mathcal{S} has rational coefficients, then the intersection of the two initial quadrics is reduced to two distinct points.*

This theorem, which is proved in Section 14.5.2, generalizes Theorem 14.2. Indeed, it ensures that the two quadrics we end up intersecting have rational coefficients, except in one very specific situation. This is how we remove the main source of nested radicals in Levin's algorithm.

The last basic ingredient of our approach is the use of Gauss reduction of quadratic forms for diagonalizing a symmetric matrix and computing the canonical form of the associated projective quadric, instead of the traditional eigenvalues/eigenvectors approach used by Levin. Since Gauss transformation is rational (the elements of the matrix P which sends S into canonical form are rational), this removes some layers of nested radicals from Levin's algorithm. Note, also, that there is no difficulty parameterizing the reduced quadric $S' = P^T S P$ since, by Sylvester's Inertia Law, S and S' have the same inertia.

14.4.2 Algorithm outline

Armed with these ingredients, we are now in a position to outline our generic algorithm.

Let $R(\lambda) = \lambda S - T$ be the pencil generated by the quadrics Q_S and Q_T of \mathbb{P}^3 and $\mathcal{D}(\lambda) = \det(R(\lambda))$ be the characteristic polynomial of the pencil. Recall that, although working in all cases, our generic algorithm is best designed when $\mathcal{D}(\lambda)$ is not identically zero and does not have any multiple root. In the other case, a better algorithm is described in parts II and III. The outline of our intersection algorithm is as follows (details follow in ensuing sections) :

1. Find a quadric Q_R with rational coefficients in the pencil, such that $\det R > 0$ if possible or $\det R = 0$ otherwise. (If no such R exists, the intersection is reduced to two points, which we output.) If the inertia of R is $(4, 0)$, output empty intersection. Otherwise, proceed.

Assume for the sake of simplicity that $Q_S \neq Q_R$, in such a way that $Q_S \cap Q_R = Q_S \cap Q_T$.

2. If the inertia of R is not $(2, 2)$, apply Gauss reduction to R and compute a frame in which $P^T R P$ is diagonal. If the inertia of R is $(2, 2)$, its parameterization contains in general two square roots but one can be eliminated as follows. First, find a rational point close enough to Q_R such that the quadric in the pencil through this point has the same inertia as Q_R . Replace Q_R by this quadric. Then use that rational point to compute a frame in which $P^T R P$ is the diagonal matrix $\text{diag}(1, 1, -1, -\delta)$, with $\delta \in \mathbb{Q}$.

In the local frame, Q_R can be described by one of the parameterizations \mathbf{X} of Table 14.3. Compute the parameterization $P\mathbf{X}$ of Q_R in the global frame.

3. Consider the equation

$$\Omega : (P\mathbf{X})^T S (P\mathbf{X}) = 0. \quad (14.2)$$

²⁵Note that there is necessarily a trade-off between the minimal degree of a parameterization in one of its parameters and the degree of its coefficient field. For instance, Wang, Joe and Goldman [WJG97] give parameterizations of quadrics that have rational coefficients but are quadratic in all of their parameters.

Equation Ω is of degree at most 2 in (at least) one of the parameters. Solve it for this parameter in terms of the other(s) and compute the domain of the solution.

4. Substitute this parameter in PX , giving a parameterization of the intersection of Q_S and Q_T .

14.4.3 Details of Step 1

The detailed description of Step 1 is as follows. Recall that $\mathcal{D}(\lambda) = \det(R(\lambda))$ is the characteristic polynomial of the pencil.

1. a. If $\mathcal{D}(\lambda) \equiv 0$, set $R = S$ and proceed.
- b. Otherwise, compute isolating intervals for the real roots of $\mathcal{D}(\lambda)$ (using for instance a variant of Uspensky's algorithm [RZ04]). Compute a rational number λ_0 in between each of the separating intervals and, for each λ_0 such that $\mathcal{D}(\lambda_0) > 0$, compute the inertia of the corresponding quadrics using Gauss reduction. If one of the inertias is $(4, 0)$, output $Q_S \cap Q_T = \emptyset$. Otherwise, one of these inertias is $(2, 2)$ and we proceed with the corresponding quadric.
- c. Otherwise (i.e. $\mathcal{D}(\lambda) \not\equiv 0$ and $\mathcal{D}(\lambda) \leq 0$ for all λ), compute the greatest common divisor $\gcd(\lambda)$ of $\mathcal{D}(\lambda)$ and its derivative with respect to λ . If $\gcd(\lambda)$ has a rational root λ_0 , proceed with the corresponding quadric $Q_{R(\lambda_0)}$.
- d. Otherwise (i.e. $\mathcal{D}(\lambda)$ has two non-rational double real roots), $Q_S \cap Q_T$ is reduced to two points. The quadric corresponding to one of these two roots is of inertia $(2, 0)$ (an imaginary pair of planes). The singular line of this pair of planes is real and can be parameterized easily, even though it is not rational. Intersecting that line with any of the input quadrics gives the two points.

To assert the correctness of this algorithm, we have several things to prove. First, we make clear why, when looking for a quadric in the pencil (S, T) with inertia different from those of S and T , the right polynomial to consider is $\mathcal{D}(\lambda)$:

Lemma 14.4. *The inertia of $R(\lambda)$ is invariant on any interval of λ not containing a root of $\mathcal{D}(\lambda)$.*

Proof. It suffices to realize that the eigenvalues of $R(\lambda)$ are continuous functions of λ and that the characteristic polynomial of $R(\lambda)$

$$\det(R(\lambda) - lI)$$

is a polynomial in l whose constant coefficient is $\mathcal{D}(\lambda)$, where I is the identity matrix of size 4. Thus the eigenvalues of $R(\lambda)$ may change of sign only at a zero of $\det(R(\lambda))$. \square

Let us now show that Step 1 of our algorithm always outputs empty intersection when $Q_S \cap Q_T = \emptyset$. This, in fact, is a direct consequence of Lemma 14.4 and of the following theorem proved in 1936/1937 by the German mathematician Paul Finsler.

Theorem 14.5 ([Fin37]). *Assume $n \geq 3$ and let S, T be real symmetric matrices of size n . Then $Q_S \cap Q_T = \emptyset$ if and only if the pencil of matrices generated by S and T contains a matrix of inertia $(n, 0)$.*

In Step 1.d, Q_S and Q_T intersect in two points by Theorem 14.3. Furthermore, the quadric corresponding to one of two roots of $\mathcal{D}(\lambda)$ is a real line by the proof of Theorem 14.3.

Finally, note that we can further refine Step 1.b by computing the inertia of the quadrics $Q_{R(\lambda_0)}$ with positive determinant only when the characteristic polynomial has four real roots counted with multiplicities. Indeed, in view of the following proposition, testing for the presence of a definite matrix in the pencil needs to be done only in that case.

Proposition 14.6. *Assume $n \geq 3$ and let S, T be real symmetric matrices of size n . Then $Q_S \cap Q_T = \emptyset$ implies that $\det(\lambda S + \mu T)$ does not identically vanish and that all its roots are real.*

Proof. We use the equivalence provided by Theorem 14.5 of the emptiness of the intersection and the existence of a definite matrix in the pencil. Let U be a definite matrix of the pencil which we choose positive (a similar proof goes for negative definite).

Since U is positive definite, we can apply to it a Cholesky factorization : $U = HH^T$, where H is a lower triangular matrix. Consider the matrix $C = (H^{-1})S(H^{-1})^T$. Since C is real symmetric, it has n pairs of real eigenvalues and eigenvectors (v_i, \mathbf{x}_i) . Let $\mathbf{y}_i = (H^{-1})^T \mathbf{x}_i$. Then we have

$$H(C\mathbf{x}_i) = H(v_i\mathbf{x}_i) \implies S\mathbf{y}_i = v_iU\mathbf{y}_i.$$

Hence all the roots of the characteristic polynomial of $U^{-1}S$ are real, which implies that all the roots of $\det(\lambda S + \mu U) = 0$ are real. It follows that all the roots of $\det(\lambda S + \mu T) = 0$ are also real. \square

14.4.4 Details of Step 2

There are two cases, according to the inertia of R .

14.4.4.1 The inertia of R is not $(2, 2)$

When the inertia of R is different from $(2, 2)$, we use Gauss reduction of quadratic forms and parameterize the resulting quadric, whose associated matrix $P^T R P$ is diagonal. In view of Sylvester's Inertia Law, the reduced quadric $Q_{P^T R P}$ has the same inertia as Q_R . Thus it can be parameterized with at most one square root by one of the parameterizations \mathbf{X} of Table 14.3. Since Gauss reduction is rational (i.e. P is a matrix with rational coefficients), the parameterization $P\mathbf{X}$ of Q_R contains at most one square root.

14.4.4.2 The inertia of R is $(2, 2)$

When the inertia of R is $(2, 2)$, the coefficients of the parameterization of Q_R can live, in the worst case, in an extension $\mathbb{Q}(\sqrt{m}, \sqrt{n})$ of degree 4 of \mathbb{Q} (see Table 14.3). We show here that there exists, in the neighborhood of Q_R , a quadric $Q_{R'}$ with rational coefficients such that

$$Q_S \cap Q_{R'} = Q_S \cap Q_R = Q_S \cap Q_T$$

and the coefficients of the parameterization of $Q_{R'}$ are in $\mathbb{Q}(\sqrt{\det R'})$.

First, apply Gauss reduction to Q_R . If any of \sqrt{ac} or \sqrt{bd} is rational in the parameterization of Q_R (as in Table 14.3), we are done. Otherwise, compute an arbitrary point $\mathbf{p} \in \mathbb{P}^3(\mathbb{R})$ on Q_R by taking any value of the parameters like, say, $(u, v) = (0, 1)$ and $(s, t) = (0, 1)$. Approximate \mathbf{p} by a point $\mathbf{p}' \in \mathbb{P}^3(\mathbb{Q})$ not on $Q_S \cap Q_T$. Then compute $\lambda'_0 \in \mathbb{Q}$ such that \mathbf{p}' belongs to the quadric $Q_{R(\lambda'_0)}$ of the pencil. This is easy to achieve in view of the following lemma.

Lemma 14.7. *In a pencil generated by two quadrics Q_S, Q_T with rational coefficients, there is exactly one quadric going through a given point \mathbf{p}' that is not on $Q_S \cap Q_T$. If \mathbf{p}' is rational, this quadric is rational.*

Proof. In the pencil generated by Q_S and Q_T , a quadric $Q_{R(\lambda, \mu)}$ contains \mathbf{p}' if and only if $\mathbf{p}'^T (\lambda S + \mu T) \mathbf{p}' = 0$, that is if and only if $\lambda(\mathbf{p}'^T S \mathbf{p}') + \mu(\mathbf{p}'^T T \mathbf{p}') = 0$. If \mathbf{p}' is not on $Q_S \cap Q_T$, this equation is linear in $(\lambda, \mu) \in \mathbb{P}^1$ and thus admits a unique solution. Moreover, if \mathbf{p}' is rational, the equation has rational coefficients and thus the quadric of the pencil containing \mathbf{p}' is rational. \square

Note that λ'_0 and the λ_0 such that $R = R(\lambda_0)$ get arbitrarily close to one another as \mathbf{p}' gets close to \mathbf{p} . Thus if \mathbf{p}' is close enough to \mathbf{p} , $R' = R(\lambda'_0)$ has the same inertia $(2, 2)$ as R , by Lemma 14.4. We refine the approximation \mathbf{p}' of \mathbf{p} until R' has inertia $(2, 2)$.

We now have a quadric $Q_{R'}$ of inertia $(2, 2)$ and a rational point on $Q_{R'}$. Consider any rational line through \mathbf{p}' that is not in the plane tangent to $Q_{R'}$ at \mathbf{p}' . This line further intersects $Q_{R'}$ in another point \mathbf{p}'' . Point \mathbf{p}'' is rational because otherwise \mathbf{p}' and \mathbf{p}'' would be conjugate in the field extension of \mathbb{Q} (since $Q_{R'}$ and the line are both rational) and thus \mathbf{p}' would not be rational. Compute the rational transformation P sending $\mathbf{p}', \mathbf{p}''$ onto $(1, \pm 1, 0, 0)$. Apply this transformation to R' and then apply Gauss reduction of quadratic forms. In the local frame, $Q_{R'}$ has equation (up to a constant factor)

$$x^2 - y^2 + \alpha z^2 + \beta w^2 = 0, \tag{14.3}$$

with $\alpha\beta < 0$. Now consider the linear transformation whose matrix is P'

$$P' = \frac{1}{2} \begin{pmatrix} 1 + \alpha & 0 & 1 - \alpha & 0 \\ 1 - \alpha & 0 & 1 + \alpha & 0 \\ 0 & 2 & 0 & 0 \\ 0 & 0 & 0 & 2\alpha \end{pmatrix}.$$

Applying P' to the already reduced quadric of Eq. (14.3) gives the equation

$$x^2 + y^2 - z^2 - \delta w^2 = 0, \quad (14.4)$$

where $\delta = -\alpha\beta > 0$. The quadric of Eq. (14.4) can be parameterized by

$$\mathbf{X}((u, v), (s, t)) = \left(ut + vs, us - vt, ut - vs, \frac{us + vt}{\sqrt{\delta}} \right),$$

with $(u, v), (s, t) \in \mathbb{P}^1$ (see Table 14.3).

The three consecutive transformation matrices have rational coefficients thus $\mathbb{Q}(\sqrt{\delta}) = \mathbb{Q}(\sqrt{\det R'})$ and the product of these transformation matrices with \mathbf{X} is a polynomial parameterization of $Q_{R'}$ with coefficients in $\mathbb{Q}(\sqrt{\delta})$, $\delta \in \mathbb{Q}$.

14.4.5 Details of Step 3

Solving Equation (14.2) can be done as follows. Recall that the content in the variable x of a multivariate polynomial is the gcd of the coefficients of the x^i .

Equation (14.2) may be seen as a quadratic equation in one of the parameters. For instance, if R has inertia $(2, 2)$, Eq. (14.2) is a homogeneous biquadratic equation in the variables $\xi = (u, v)$ and $\tau = (s, t)$. Using only gcd computations, we can factor it in its content in ξ (which is a polynomial in τ or a constant), its content in τ , and a remaining factor. If the content in ξ (or in τ) is not constant, solve it in τ (in ξ); substituting the obtained real values in \mathbf{X} , we have a parameterization of some components of $Q_S \cap Q_T = Q_S \cap Q_R$ in the frame in which Q_R is canonical. If the remaining factor is not constant, solve it in a parameter in which it is linear, if any, or in τ . Substituting the result in \mathbf{X} , we have a parameterization of the last component of the intersection. If the equation which is solved is not linear, the domain of the parameterization is the set of ξ such that the degree 4 polynomial $\Delta(\xi) = b^2(\xi) - 4a(\xi)c(\xi)$ is positive, where $a(\xi), b(\xi)$ and $c(\xi)$ are the coefficients of τ^2, τ and 1 in (14.2), respectively.

14.5 Canonical forms and proof of Theorem 14.3

We now prove Theorem 14.3, the key result stated in the previous section. We start by recalling some preliminary results.

14.5.1 Canonical form for a nonsingular pair of symmetric matrices

We state results, proved by Uhlig [Uhl73, Uhl76], we need for computing the canonical form of a pair of real symmetric matrices. Though only part of this theory is required for the proof of Theorem 14.3 (Section 14.5.2), we will need its full power in Part II of this paper for characterizing real pencils of quadrics.

Let us start by defining the notion of Jordan blocks.

Definition 14.8. Let M be a square matrix of the form

$$(\ell) \quad \text{or} \quad \begin{pmatrix} \ell & e & & 0 \\ & \ddots & \ddots & \\ & & e & \\ 0 & & & \ell \end{pmatrix}.$$

If $l \in \mathbb{R}$ and $e = 1$, M is called a real Jordan block associated with l . If

$$\ell = \begin{pmatrix} a & -b \\ b & a \end{pmatrix}, \quad a, b \in \mathbb{R}, \quad b \neq 0, \quad e = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix},$$

M is called a complex Jordan block associated with $a + ib$.

Now we can state the real Jordan normal form theorem for real square matrices.

Theorem 14.9 (Real Jordan normal form). *Every real square matrix A is similar over the reals to a block diagonal matrix $\text{diag}(A_1, \dots, A_k)$, called real Jordan normal form of A , in which each A_j is a (real or complex) Jordan block associated with an eigenvalue of A .*

The Canonical Pair Form Theorem then goes as follows :

Theorem 14.10 (Canonical Pair Form). *Let S and T be two real symmetric matrices of size n , with S nonsingular. Let $S^{-1}T$ have real Jordan normal form $\text{diag}(J_1, \dots, J_r, J_{r+1}, \dots, J_m)$, where J_1, \dots, J_r are real Jordan blocks corresponding to real eigenvalues of $S^{-1}T$ and J_{r+1}, \dots, J_m are complex Jordan blocks corresponding to pairs of complex conjugate eigenvalues of $S^{-1}T$. Then :*

(a) *The characteristic polynomial of $S^{-1}T$ and $\det(\lambda S - T)$ have the same roots λ_j with the same (algebraic) multiplicities m_j .*

(b) *S and T are simultaneously congruent by a real congruence transformation to*

$$\text{diag}(\varepsilon_1 E_1, \dots, \varepsilon_r E_r, E_{r+1}, \dots, E_m)$$

and

$$\text{diag}(\varepsilon_1 E_1 J_1, \dots, \varepsilon_r E_r J_r, E_{r+1} J_{r+1}, \dots, E_m J_m),$$

respectively, where $\varepsilon_i = \pm 1$ and E_i denotes the square matrix

$$\begin{pmatrix} 0 & & & 1 \\ & \ddots & & \\ & & \ddots & \\ 1 & & & 0 \end{pmatrix}$$

of the same size as J_i for $i = 1, \dots, m$. The signs ε_i are unique (up to permutations) for each set of indices i that are associated with a set of identical real Jordan blocks J_i .

(c) *The sum of the sizes of the blocks corresponding to one of the λ_j is the multiplicity m_j if λ_j is real or twice this multiplicity if λ_j is complex. The number of the corresponding blocks (the geometric multiplicity of λ_j) is $t_j = n - \text{rank}(\lambda_j S - T)$, and $1 \leq t_j \leq m_j$.*

Note that the canonical pair form of Theorem 14.10 can be considered the finest simultaneous block diagonal structure that can be obtained by a real congruence transformation for a given pair of real symmetric matrices, in the sense that it maximizes the number of blocks in the diagonalization of S and T .

14.5.2 Proof of Theorem 14.3

To prove Theorem 14.3, we consider a pencil of real symmetric 4×4 matrices generated by two symmetric matrices S and T of inertia $(3, 1)$. We may suppose that they have the block diagonal form of the above theorem.

If all the blocks had an even size, the determinant of S would be positive, contradicting our hypothesis. Thus, there is a block of odd size in the canonical form of S . It follows that $\det(\lambda S - T)$ has at least one real root and the matrix of the pencil corresponding to this root has an inertia different from $(3, 1)$. This proves the first part.

If $\det(\lambda S - T)$ has a simple real root, there is an interval of values for λ for which $\det(\lambda S - T) > 0$, and we are done with any rational value of λ in this interval. If $\det(\lambda S - T)$ has either a double real root and two complex roots, two rational double real roots or a quadruple real root, the quadrics corresponding to the real root(s) have rational coefficients and have inertia different from $(3, 1)$.

Thus we are left with the case where $\det(\lambda S - T)$ has two non rational double real roots, which are algebraically conjugate. In other words,

$$\det(\lambda S - T) = c(\lambda - \lambda_1)^2(\lambda - \lambda_2)^2,$$

with $\lambda_1, \lambda_2 \in \mathbb{R} \setminus \mathbb{Q}$ and $\lambda_2 = \overline{\lambda_1}$ its (real algebraic) conjugate. Following the notation of Theorem 14.10, we have $m_1 = m_2 = 2$ and $1 \leq t_i \leq 2$, for $i = 1, 2$. In other words, $(t_1, t_2) \in \{(1, 1), (1, 2), (2, 1), (2, 2)\}$.

We can quickly get rid of the case $(t_1, t_2) = (1, 1)$. Indeed, in this case the blocks have an even size and S is not of inertia $(3, 1)$. We can also eliminate the cases $(t_1, t_2) \in \{(1, 2), (2, 1)\}$, because the matrices $\lambda_1 S - T$ and $\lambda_2 S - T$ are algebraically conjugate, and so must have the same rank and the same number of blocks.

We are thus left with the case $(t_1, t_2) = (2, 2)$. In this situation, S and T have four blocks, i.e., they are diagonal :

$$\begin{cases} S = \text{diag}(\varepsilon_1, \varepsilon_2, \varepsilon_3, \varepsilon_4), \\ T = \text{diag}(\varepsilon_1 \lambda_1, \varepsilon_2 \lambda_1, \varepsilon_3 \lambda_2, \varepsilon_4 \lambda_2). \end{cases}$$

The pencil $\lambda S - T$ is generated by the two quadrics of rank 2

$$\begin{cases} S' = \lambda_1 S - T = \text{diag}(0, 0, \varepsilon_3(\lambda_1 - \lambda_2), \varepsilon_4(\lambda_1 - \lambda_2)), \\ T' = \lambda_2 S - T = \text{diag}(\varepsilon_1(\lambda_2 - \lambda_1), \varepsilon_2(\lambda_2 - \lambda_1), 0, 0). \end{cases}$$

We have that

$$\det(S' + T') = \varepsilon_1 \varepsilon_2 \varepsilon_3 \varepsilon_4 (\lambda_1 - \lambda_2)^4$$

is negative since all the quadrics of the pencil have negative determinant except $Q_{S'}$ and $Q_{T'}$. Thus $\varepsilon_1 \varepsilon_2$ and $\varepsilon_3 \varepsilon_4$ have opposite signs. It follows that one of S' and T' has inertia $(2, 0)$ (say S') and the other has inertia $(1, 1)$. Thus $Q_{S'}$ is a straight line, which intersects the real pair of planes $Q_{T'}$. Since $Q_{S'} \cap Q_{T'}$ is contained in all the quadrics of the pencil and since the pencil has quadrics of inertia $(3, 1)$ (which are not ruled), the line $Q_{S'}$ is not included in $Q_{T'}$ and the intersection is reduced to two real points. Since the equations of $Q_{S'}$ and $Q_{T'}$ are $z^2 + w^2 = 0$ and $x^2 - y^2 = 0$ respectively, the two points have coordinates $(1, 1, 0, 0)$ and $(-1, 1, 0, 0)$. They are thus distinct. \square

Remark 14.11. *Pencils generated by two quadrics of inertia $(3, 1)$ and having no quadric with rational coefficients of inertia different from $(3, 1)$ do exist. Consider for instance*

$$\begin{aligned} Q_S &: 2x^2 - 2xz - 2yw + z^2 + w^2 = 0, \\ Q_T &: 4x^2 + 2y^2 - 2yw + z^2 - 6xz + 3w^2 = 0. \end{aligned}$$

Then, $\det(\lambda S - T) = -(\lambda^2 - 5)^2$.

14.6 Optimality of the parameterizations

We now prove that, among the parameterizations of projective quadrics linear in one of the parameters, the ones of Table 14.3 have, in the worst case, an optimal number of radicals. In other words, for each type of projective quadric, there are examples of surfaces for which the number of square roots of the parameterizations of Table 14.3 is required.

More precisely, we prove the following theorem, which will be crucial in asserting the near-optimality of our algorithm for parameterizing quadric intersections.

Theorem 14.12. *In the set of parameterizations linear in one of the parameters, the parameterizations of Table 14.3 are worst-case optimal in the degree of the extension of \mathbb{Q} on which they are defined.*

For a quadric Q of equation $ax^2 + by^2 - cz^2 - dw^2 = 0$ ($a, b, c, d > 0$), the parameterization of Table 14.3 is optimal if Q has no rational point, which is the case for some quadrics. Knowing a rational point on Q (if any), we can compute a rational congruence transformation sending Q into the quadric of equation $x^2 + y^2 - z^2 - abcd w^2 = 0$, for which the parameterization of Table 14.3 is optimal.

For a quadric Q of equation $ax^2 + by^2 - cz^2 = 0$ ($a, b, c > 0$), the parameterization of Table 14.3 is optimal if Q has no rational point other than its singular point $(0, 0, 0, 1)$, which is the case for some quadrics. Knowing such a rational point on Q (if any), we can compute a rational congruence transformation sending Q into the quadric of equation $x^2 + y^2 - z^2 = 0$, for which the parameterization of Table 14.3 is rational (and thus optimal).

For the other types of projective quadrics, the parameterizations of Table 14.3 are optimal in all cases.

We prove this theorem by splitting it into four more detailed propositions : Proposition 14.13 for inertia (1, 1), Proposition 14.14 for inertia (2, 1) and Propositions 14.15 and 14.17 for inertia (2, 2).

Proposition 14.13. *A projective quadric Q of equation $ax^2 - by^2 = 0$ ($a, b > 0$) admits a rational parameterization in \mathbb{Q} if and only if it has a rational point outside the singular line $x = y = 0$, or equivalently iff ab is a square in \mathbb{Q} . If ab is a square in \mathbb{Q} , then the parameterization of Table 14.3 is rational.*

Proof. A point (x, y, z, w) on Q not on its singular line $x = y = 0$ is rational if and only if y/x , z/x , and w/x are rational. Since $(y/x)^2 = \frac{ab}{b^2}$ and z and w are not constrained, there exists such a rational point if and only if ab is a square.

If there exists a parameterization which is rational over \mathbb{Q} , then there exists some rational point outside the line $x = y = 0$, showing *a contrario* that there is no rational parameterization if ab is not a square.

Finally, if ab is the square of a rational number, then the parameterization of Table 14.3 is rational. \square

Proposition 14.14. *A projective quadric Q of equation $ax^2 + by^2 - cz^2 = 0$ ($a, b, c > 0$) admits a rational parameterization in \mathbb{Q} if and only if it contains a rational point other than the singular point $(0, 0, 0, 1)$. Knowing such a rational point, we can compute a rational congruence transformation P sending Q into the quadric of equation $x^2 + y^2 - z^2 = 0$ for which the parameterization of Table 14.3 is rational ; lifting this parameterization to the original space by multiplying by matrix P , we have a rational parameterization of Q .*

On the other hand, there are such quadrics without a rational point and thus without a rational parameterization, for example the quadric of equation $x^2 + y^2 - 3z^2 = 0$.

Proof. If Q has a rational point other than $(x = y = z = 0)$, any rational line passing through this point and not included in Q cuts Q in another rational point. Compute the rational congruence transformation sending these points onto $(\pm 1, 1, 0, 0)$. Applying this transformation to Q gives a quadric of equation $x^2 - y^2 + r$, where r is a polynomial of degree at most one in x and y . Thus Gauss reduction algorithm leads to the form $x^2 - y^2 + dz^2 = (X^2 + Y^2 - Z^2)/d$ where $X = (1+d)x/2 + (1-d)y/2$, $Y = dz$ and $Z = (1-d)x/2 + (1+d)y/2$. The parameterization of Table 14.3 applied to equation $X^2 + Y^2 - Z^2$ is clearly rational. Lifting this parameterization back to the original space, we obtain a rational parameterization of Q .

Reciprocally, if Q has no rational point, then Q does not admit a rational parameterization.

Now, suppose for a contradiction that the quadric with equation $x^2 + y^2 - 3z^2 = 0$ has a rational point (x, y, z, w) different from $(0, 0, 0, 1)$. By multiplying x, y , and z by a common denominator and dividing them by their gcd, we obtain another rational point on the quadric for which x, y and z are integers that are not all even. Note that x^2 is equal, modulo 4, to 0 if x is even and 1 otherwise (indeed, modulo 4, $0^2 = 0$, $1^2 = 1$, $2^2 = 0$ and $3^2 = 1$). Thus, $x^2 + y^2 - 3z^2 \equiv x^2 + y^2 + z^2 \pmod{4}$ is equal to the number of odd numbers in x, y, z , i.e. 1, 2 or 3. Thus $x^2 + y^2 - 3z^2 \neq 0$, contradicting the hypothesis that (x, y, z, w) is a point on the quadric. \square

Proposition 14.15. *Let Q be the quadric of equation $ax^2 + by^2 - cz^2 - dw^2 = 0$ ($a, b, c, d > 0$). Any field \mathbb{K} in which Q admits a rational parameterization, linear in one of its parameters, contains \sqrt{abcd} .*

Proof. Let \mathbb{K} be a field in which Q admits a rational parameterization, linear in the parameter $(u, v) \in \mathbb{P}(\mathbb{K})$. Fixing the value of the other parameter $(s, t) \in \mathbb{P}(\mathbb{K})$ defines a rational line L (in \mathbb{K}) contained in Q . L cuts any plane (in possibly infinitely many points) in projective space. In particular, L cuts the plane of equation $z = 0$. Since $L \subseteq Q$, L cuts the conic of equation $ax^2 + by^2 - dw^2 = z = 0$ in a point $\mathbf{p} = (x_0, y_0, 0, 1)$. Moreover, \mathbf{p} is rational in \mathbb{K} (i.e., $x_0, y_0 \in \mathbb{K}$) because it is the intersection of a rational line and the plane $z = 0$.

The plane tangent to Q at \mathbf{p} has equation $ax_0x + by_0y - dw = 0$. We now compute the intersection of Q with this plane. Since $ax_0^2 + by_0^2 = d$ and $a, b, d > 0$, x_0 or y_0 is nonzero ; assume for instance that $x_0 \neq 0$. Squaring the equation of the tangent plane yields $(ax_0x)^2 = (by_0y - dw)^2$. By eliminating x^2 between this equation and the equation of Q , we get

$$(by_0y - dw)^2 + ax_0^2(by^2 - cz^2 - dw^2) = 0$$

or

$$dw^2(d - ax_0^2) + by^2(ax_0^2 + by_0^2) - 2bdy_0yw - acx_0^2z^2 = 0.$$

It follows from $ax_0^2 + by_0^2 = d$ that $bd(y - y_0w)^2 - acx_0^2z^2 = 0$ or also

$$b^2d^2(y - y_0w)^2 - abcdx_0^2z^2 = 0. \quad (14.5)$$

The intersection of Q and its tangent plane at \mathbf{p} contains the line L which is rational in \mathbb{K} . Thus, Equation (14.5) can be factored over \mathbb{K} into two linear terms. Hence, \sqrt{abcd} belongs to \mathbb{K} . \square

Remark 14.16. $abcd$ is the discriminant of the quadric, i.e., the determinant of the associated matrix, so it is invariant by a change of coordinates (up to a square factor). Thus, if R and R' are two matrices representing the same quadric in different frames, the fields $\mathbb{Q}(\sqrt{\det R})$ and $\mathbb{Q}(\sqrt{\det R'})$ are equal.

Proposition 14.17. A projective quadric Q of equation $ax^2 + by^2 - cz^2 - dw^2 = 0$ ($a, b, c, d > 0$) admits a rational parameterization in $\mathbb{Q}(\sqrt{abcd})$ if and only if it contains a rational point. Knowing such a rational point, we can compute a rational congruence transformation P sending Q into the quadric of equation $x^2 + y^2 - z^2 - abcdw^2 = 0$ for which the parameterization of Table 14.3 is rational over $\mathbb{Q}(\sqrt{abcd})$; lifting this parameterization to the original space by multiplying by matrix P , we have a rational parameterization of Q over $\mathbb{Q}(\sqrt{abcd})$.

On the other hand, there are such quadrics with no rational point and thus without a rational parameterization in $\mathbb{Q}(\sqrt{abcd})$, for example the quadric of equation $x^2 + y^2 - 3z^2 - 11w^2 = 0$.

Proof. If Q admits a rational parameterization in $\mathbb{Q}(\sqrt{abcd})$, then it has infinitely many rational points over this field. If Q has a point (x, y, z, w) that is rational over $\mathbb{Q}(\sqrt{abcd})$, but not rational over \mathbb{Q} , we may suppose without loss of generality that $x = 1$, by permuting the variables in order that $x \neq 0$ and then by dividing all coordinates by x . The conjugate point $(1, y', z', w')$ over $\mathbb{Q}(\sqrt{abcd})$ belongs also to Q . The line passing through these points is rational (over \mathbb{Q}), as is the point $(1, (y+y')/2, (z+z')/2, (w+w')/2)$. Choose a rational frame transformation such that this line becomes the line $z = w = 0$ and this point becomes $(1, 0, 0, 0)$. In this new frame the coordinates of the conjugate points are $(1, \pm e\sqrt{abcd}, 0, 0)$ for some rational number e , and the equation of Q is $abcd e^2 x^2 - y^2 + r = 0$ where r is a polynomial of degree at most 1 in x and y . Gauss reduction thus provides an equation of the form $abcd e^2 x^2 - y^2 + fz^2 - gw^2 = 0$, and the invariance of the determinant (Remark 14.16) shows that fg is the square of a rational number h . Thus $(0, 0, g, h)$ is a rational point of Q over \mathbb{Q} .

Now, if Q has a rational point over \mathbb{Q} , one may get another rational point as the intersection of the quadric and any line passing through the point and not tangent to the quadric. One can compute a rational congruence transformation such that these points become $(1, \pm 1, 0, 0)$. In this new frame the equation of Q has the form $x^2 - y^2 - r$ where r is a polynomial of degree at most 1 in x and y . Gauss reduction provides thus an equation of the form $x^2 - y^2 + ez^2 - fw^2 = (X^2 + Y^2 - Z^2 - efw^2)/e$, with $X = (1+e)x/2 + (1-e)y/2$, $Y = ez$ and $Z = (1-e)x/2 + (1+e)y/2$. By the invariance of the determinant, $ef = g^2abcd$ for some rational number g . Putting $W = gw$, we get the equation $X^2 + Y^2 - Z^2 - abcdW^2 = 0$ for Q , and the parameterization of Table 14.3 is rational over $\mathbb{Q}(\sqrt{abcd})$.

It follows from this proof that, if a quadric of inertia $(2, 2)$ has a rational point, it has a parameterization in $\mathbb{Q}(\sqrt{abcd})$, which is linear in one of the parameters. Conversely, for proving that such a parameterization does not always exist, it suffices to prove that there are quadrics of inertia $(2, 2)$ having no rational point over \mathbb{Q} . Let us consider the quadric of equation $x^2 + y^2 - 3z^2 - 11w^2 = 0$. If it has a rational point (x, y, z, w) , then by multiplying x, y, z and w by some common denominator and by dividing them by their gcd, we may suppose that x, y, z and w are integers which are not all even. As in the proof of Proposition 14.14, $x^2 + y^2 - 3z^2 - 11w^2$ is equal modulo 4 to the number of odd numbers in x, y, z, w . Thus all of them are odd. It is straightforward that the square of an odd number is equal to 1 modulo 8. It follows that $x^2 + y^2 - 3z^2 - 11w^2$ is equal to 4 modulo 8, a contradiction with $x^2 + y^2 - 3z^2 - 11w^2 = 0$. \square

14.7 Near-optimality in the smooth quartic case

In this section, we prove that the algorithm given in Section 14.4 outputs, in the generic (smooth quartic) case, a parameterization of the intersection that is optimal in the number of radicals up to one possibly unnecessary square root. We also show that deciding whether this extra square root can be avoided or not is hard. Moreover, we give examples where the extra square root cannot be eliminated, for the three possible morphologies of a real smooth quartic.

14.7.1 Algebraic preliminaries

First recall that, as is well known from the classification of quadric pencils by invariant factors (see [Bro06] and Part II for more), the intersection of two quadrics is a nonsingular quartic exactly when $\mathcal{D}(\lambda, \mu) = \det R(\lambda, \mu)$ has no multiple root. Otherwise the intersection is singular. Note that the intersection is nonsingular exactly when $\gcd(\frac{\partial \mathcal{D}}{\partial \lambda}, \frac{\partial \mathcal{D}}{\partial \mu}) = 1$.

Moreover, when the intersection is nonsingular, the rank of any quadric in the pencil is at least three ; indeed, all the roots of $\mathcal{D}(\lambda, \mu)$ are simple and thus, in Theorem 14.10(c), $m_j = 1$, thus $t_j = 1$, hence the quadrics associated with the roots of $\mathcal{D}(\lambda, \mu)$ have rank 3.

Whether the intersection of two quadrics admits a parameterization with rational functions directly follows from classical results :

Proposition 14.18. *The intersection of two quadrics admits a parameterization with rational functions if and only if the intersection is singular.*

Proof. First recall that a curve admits a parameterization with rational functions if and only if it has zero genus [Per95].

Assume first that the intersection of the two quadrics is irreducible. In $\mathbb{P}^3(\mathbb{C})$, if two algebraic surfaces of degree d_1 and d_2 intersect in an irreducible curve, its genus is

$$\frac{1}{2}d_1d_2(d_1 + d_2 - 4) + 1 - \sum_{i=1}^k \frac{q_i(q_i - 1)}{2},$$

where k is the number of singular points and $q_{i,i=1,\dots,k}$ their respective multiplicity [Nam84]. The intersection curve has thus genus 1 when it is smooth, 0 otherwise. The result follows.

Assume now that the intersection of the two quadrics is reducible. If the intersection contains only points, lines and conics, which can be parameterized in a classic way by rational functions, we are done. For the remaining case (cubic and line), we use the following result. In $\mathbb{P}^3(\mathbb{C})$, if two algebraic surfaces of degree d_1 and d_2 intersect in two irreducible curves of degree d and d' and of genus g and g' , then [Per95]

$$g' - g = \left(\frac{1}{2}(d_1 + d_2) - 2 \right) (d' - d).$$

For quadrics, $d_1 + d_2 = 4$, so we get $g = g'$. So the genus of the cubic is that of the line, i.e. 0. □

Finally consider the equation $\Omega : \mathbf{X}^T S' \mathbf{X} = 0$, obtained in Step 3 of our algorithm, where \mathbf{X} is the parameterization of Q_R and S' is the matrix of Q_S in the canonical frame of Q_R . Let C_Ω be the curve zero-set of Ω . Depending on the projective type of Q_R , C_Ω is a bidegree (2, 2) curve in $\mathbb{P}^1 \times \mathbb{P}^1$ (inertia (2, 2) or (2, 0)), a quartic curve in \mathbb{P}^{*2} (inertia (2, 1)) or a quartic curve in \mathbb{P}^2 (inertia (1, 1) or (1, 0)). Let C denote the curve of intersection of the two given quadrics Q_S and Q_T . We have the following classical result.

Fact 14.19. *The parameterization of Q_R defines an isomorphism between C and C_Ω . In particular, C and C_Ω have the same genus, irreducibility, and factorization.*

14.7.2 Optimality

Assume the intersection is a real nonsingular quartic. Then $\mathcal{D}(\lambda, \mu)$ has no multiple root, and thus Q_R is necessarily a quadric of inertia (2, 2). After Step 2 of our algorithm, Q_R has a parameterization in $\mathbb{Q}(\sqrt{\delta})$ that is bilinear in $\xi = (u, v)$ and $\tau = (s, t)$. After resolution of Ω and substitution in Q_R , we get a parameterization in $\mathbb{Q}(\sqrt{\delta})[\xi, \sqrt{\Delta}]$ with $\Delta \in \mathbb{Q}(\sqrt{\delta})[\xi]$ of degree 4.

Proposition 14.18 implies that it cannot be parameterized by rational functions, so $\sqrt{\Delta}$ cannot be avoided. The question now is : can $\sqrt{\delta}$ be avoided ? The answer is twofold :

1. deciding whether $\sqrt{\delta}$ can be avoided amounts, in the general case, to finding a rational point on a surface of degree 8,
2. there are cases in which $\sqrt{\delta}$ cannot be avoided.

We prove these results in the following two sections.

14.7.2.1 Optimality test

We first prove two preliminary lemmas.

Lemma 14.20. *If the intersection of two given quadrics has a parameterization involving only one square root (i.e., a parameterization in $\mathbb{Q}(\sqrt{\delta})[\xi]$ or in $\mathbb{Q}[\xi, \sqrt{\Delta}]$ with $\Delta \in \mathbb{Q}[\xi]$), there exists a quadric with rational coefficients in the pencil that contains a rational line.*

Proof. In what follows, call *degree* of a point the degree of the smallest field extension of \mathbb{Q} containing the coordinates of this point.

If the parameterization of the intersection involves only one square root, the intersection contains infinitely many points of degree at most 2, one for any rational value of the parameters. Now we have several cases according to the type of points contained in the intersection.

If the intersection contains a point \mathbf{p} of degree 2, it contains also its algebraic conjugate $\bar{\mathbf{p}}$. The line passing through \mathbf{p} and $\bar{\mathbf{p}}$ is invariant by conjugation, so is rational. Let \mathbf{q} be a rational point on this line. The quadric of the pencil passing through \mathbf{q} is rational (Lemma 14.7). Since it also contains \mathbf{p} and $\bar{\mathbf{p}}$ (the intersection is contained in any quadric of the pencil), this quadric cuts the line in at least 3 points and thus contains it.

If the intersection contains a regular rational point (i.e. a rational point which is not a singular point of the intersection), then the line tangent to the intersection at this point is rational, and is tangent to any quadric of the pencil. The quadric of the pencil passing through a rational point of this tangent line contains the contact point ; thus it contains the tangent line.

If the intersection contains a singular rational point \mathbf{p} , then all the quadrics of the pencil which are not singular at \mathbf{p} have the same tangent plane at \mathbf{p} . Let us consider the quadric of the pencil passing through a rational point \mathbf{q} of this tangent plane (or through any rational point, if none of the quadrics is regular at \mathbf{p}). As above, this quadric contains the rational line \mathbf{pq} . \square

Lemma 14.21. *If a quadric contains a rational line, its discriminant is a square in \mathbb{Q} .*

Proof. If the quadric has rank less than 4, its discriminant is zero. We may thus suppose that the discriminant is not 0 and that the equation of the quadric is $ax^2 + by^2 - cz^2 - dw^2 = 0$. Since this quadric contains a rational line L , and thus a rational point, there is a rational change of frames such that the quadric has equation $x^2 + y^2 - z^2 - abcdw^2 = 0$, by Proposition 14.17. Cut the quadric by the plane $z = 0$. Since the intersection of the plane $z = 0$ and the rational line L is a rational point, the cone $x^2 + y^2 - abcdw^2 = 0$ contains a rational point outside its singular locus. By Proposition 14.14, there is a rational congruence transformation P sending this cone into the cone of equation $x^2 + y^2 - w^2 = 0$. These two cones can be seen as conics in $\mathbb{P}^2(\mathbb{Q})$ and P can be seen as a rational transformation in $\mathbb{P}^2(\mathbb{Q})$. The discriminant $-abcd$ of the conic $x^2 + y^2 - abcdw^2 = 0$ is thus equal to $(\det P)^2$ times -1 , the discriminant of the conic $x^2 + y^2 - w^2 = 0$. Hence $abcd$ is a square in \mathbb{Q} . \square

From these two technical results and the results of Section 14.6, we obtain the following equivalence.

Proposition 14.22. *When the intersection is a nonsingular quartic, it can be parameterized in $\mathbb{Q}[\xi, \sqrt{\Delta}]$ with $\Delta \in \mathbb{Q}[\xi]$ if and only if there exists a quadric of the pencil with rational coefficients having a nonsingular rational point and whose discriminant is a square in \mathbb{Q} .*

Proof. If $\sqrt{\delta}$ can be avoided, there exists, by Lemma 14.20, a quadric of the pencil with rational coefficients containing a rational line. By Lemma 14.21, the discriminant of this quadric is thus a square in \mathbb{Q} . Moreover, since the quadrics of the pencil have rank at least three, the rational line is not the singular line of some quadric (see Table 14.1) and thus contains a nonsingular point.

Conversely, if there exists a quadric of the pencil with rational coefficients having a rational nonsingular point and whose discriminant is a square, then it has a rational parameterization by Theorem 14.12 and thus $\sqrt{\delta}$ can be avoided. \square

Mirroring Proposition 14.22, we can devise a general test for deciding, in the smooth quartic case, whether the square root $\sqrt{\delta}$ can be avoided or not. Consider the equation

$$\sigma^2 = \det((\mathbf{x}^T T \mathbf{x}) S - (\mathbf{x}^T S \mathbf{x}) T), \quad \mathbf{x} = (x, y, z, c)^T, \quad (14.6)$$

where $c \in \mathbb{Q}$ is some constant such that plane $w = c \in \mathbb{Q}$ contains the vertex of no cone (inertia $(2, 1)$) of the pencil. Note that (14.6) has degree 8 in the worst case.

Theorem 14.23. *When the intersection is a nonsingular quartic, it can be parameterized in $\mathbb{Q}[\xi, \sqrt{\Delta}]$ with $\Delta \in \mathbb{Q}[\xi]$ if and only Equation (14.6) has a rational solution.*

Proof. Suppose first that (14.6) has a rational solution $(x_0, y_0, z_0, \sigma_0)$ and let $\mathbf{x}_0 = (x_0, y_0, z_0, c)^T$ and $(\lambda_0, \mu_0) = (\mathbf{x}_0^T T \mathbf{x}_0, -\mathbf{x}_0^T S \mathbf{x}_0)$. The quadric $Q = \lambda_0 Q_S + \mu_0 Q_T$ of the pencil has rational coefficients, contains the rational point $\mathbf{x}_0 = (x_0, y_0, z_0, c)^T$ and its discriminant is a square, equal to σ_0^2 . Moreover, if Q has inertia $(2, 1)$, then \mathbf{x}_0 is not its apex because, by assumption, the plane $w = c$ contains the vertex of no cone of the pencil. It then follows from Theorem 14.12 that our algorithm produces a rational parameterization of Q , and thus a parameterization of the curve of intersection with rational coefficients.

Conversely, if the curve of intersection can be parameterized in $\mathbb{Q}[\xi, \sqrt{\Delta}]$ (with $\Delta \in \mathbb{Q}[\xi]$) there exists a quadric Q of the pencil with rational coefficients containing a rational line and whose discriminant is a square in \mathbb{Q} , by Lemmas 14.20 and 14.21. The quadric Q contains a line and thus intersects any plane. Consider any plane $w = c \in \mathbb{Q}$. Since the intersection of a rational line with a rational plane is (or contains) a rational point, the intersection of Q with plane $w = c$ contains a rational point $\mathbf{x} = (x, y, z, c)^T$. The quadric (Q) of the pencil containing that point has associated matrix $(\mathbf{x}^T T \mathbf{x})S - (\mathbf{x}^T S \mathbf{x})T$ and its determinant is a square. Hence Equation (14.6) admits a rational solution. \square

Unfortunately, the question underlying the above optimality test is not within the range of problems that can currently be answered by algebraic number theory. Indeed, it is not known whether the general problem of determining if an algebraic set contains rational points (known, over \mathbb{Z} , as Hilbert's 10th problem) is decidable [Poo01]. It is known that this problem is decidable for genus zero curves and, under certain conditions, for genus one curves [Poo01], but, for varieties of dimension two or more, very little has been proved on the problem of computing rational points.

The above theorem thus implies that computing parameterizations of the intersections of two arbitrary quadrics that are always optimal in the number of radicals is currently out of reach.

However, in some particular cases, we can use the following corollary to Theorem 14.23 to prove that $\sqrt{\delta}$ cannot be avoided.

Corollary 14.24. *If the intersection C of Q_S and Q_T is a nonsingular quartic and the rational hyperelliptic quartic curve $\sigma^2 = \det(S + \lambda T)$ has no rational point, then the parameterization of C in $\mathbb{Q}(\sqrt{\delta})[\xi, \sqrt{\Delta}]$ with $\Delta \in \mathbb{Q}(\sqrt{\delta})[\xi]$ is optimal in the number of radicals.*

We use this corollary in the next section.

14.7.3 Worst case examples

We prove here that there are pairs of quadrics, intersecting in the different types of real smooth quartic, such that (14.6) has no rational solution.

In [TWW02], Tu, Wang and Wang proved that a real smooth quartic can be of three different morphologies according to the number of real roots of the characteristic polynomial. Recall that a set of points L of \mathbb{P}^3 is called *affinely finite* if there exists a projective plane P such that $P \cap L = \emptyset$; L is called *affinely infinite* otherwise.

Theorem 14.25 ([TWW02]). *Let Q_S and Q_T be two quadrics intersecting in \mathbb{C} in a smooth quartic C . C can be classified as follows :*

- If $\mathcal{D}(\lambda, \mu)$ has four real roots, then C has either two real affinely finite connected components or is empty.
- If $\mathcal{D}(\lambda, \mu)$ has two real roots and two complex roots, then C has one real affinely finite connected component.
- If $\mathcal{D}(\lambda, \mu)$ has four complex roots, then C has two real affinely infinite connected components.

Two real affinely finite components

We first look at the case where the quartic has two real affinely finite components and start with a preliminary lemma.

Lemma 14.26. *The equation*

$$y^2 = ax^4 + bx^2 + c + d(x^3 + x) \quad (14.7)$$

has no rational solution if $a, c \equiv 3 \pmod{8}$, $b \equiv 7 \pmod{8}$ and $d \equiv 4 \pmod{8}$.

Proof. Assume for a contradiction that (x, y) is a rational solution to (14.7). We can write $x = X/Z$ and $y = Y/Z^2$, where X, Y, Z are integers, $Z \neq 0$ and X, Z are mutually prime (so are not both even).

Consider first the reduction of Equation (14.7) modulo 8 :

$$Y^2 \equiv 3X^4 + 7X^2Z^2 + 3Z^4 + 4XZ(X^2 + Z^2) \pmod{8}.$$

If both X and Z are odd, X^2 and Z^2 are equal to 1 (mod 8). Thus $4(X^2 + Z^2) \equiv 0 \pmod{8}$ and $Y^2 \equiv 3 + 7 + 3 \equiv 5 \pmod{8}$, contradicting the fact that $Y^2 \equiv 0, 1$ or $4 \pmod{8}$, for all integers Y .

If X and Z are not both odd, one of X^2 and Z^2 is equal to 0 (mod 4) and the other is equal to 1 (mod 4). The reduction of Equation (14.7) modulo 4 thus gives $Y^2 \equiv 3 \pmod{4}$, contradicting the fact that $Y^2 \equiv 0$ or $1 \pmod{4}$, for all integers Y . \square

Proposition 14.27. *Consider the following pair of quadrics intersecting in a smooth quartic with two real affinely finite components :*

$$\begin{aligned} Q_S : 5y^2 + 6xy + 2z^2 - w^2 + 6zw &= 0, \\ Q_T : 3x^2 + y^2 - z^2 - w^2 &= 0. \end{aligned}$$

Then the square root $\sqrt{\delta}$ is necessary to parameterize the curve of intersection.

Proof. The characteristic polynomial has four simple real roots and we find a quadric of inertia $(2, 2)$ in each of the intervals on which it is positive (in fact Q_S and Q_T are representative quadrics in these intervals). Thus, by Theorem 14.25, the intersection of Q_S and Q_T is a real smooth quartic with two affinely finite components.

We now apply Corollary 14.24 and show that the square root $\sqrt{\delta}$ is necessary to parameterize the curve of intersection. We have :

$$\begin{aligned} \sigma^2 &= \det(S + \lambda T), \\ &= 3\lambda^4 + 12\lambda^3 - 57\lambda^2 - 156\lambda + 99, \\ &\equiv 3\lambda^4 + 7\lambda^2 + 3 + 4(\lambda^3 + \lambda) \pmod{8}, \end{aligned}$$

which has no rational solution by Lemma 14.26, so $\sqrt{\delta}$ cannot be avoided. \square

One real affinely finite component

As above, we prove a preliminary lemma.

Lemma 14.28. *The equation*

$$y^2 = ax^4 + bx^3 + cx^2 + dx + e \quad (14.8)$$

has no rational solution if $a, e \equiv 2 \pmod{4}$, $b, d \equiv 0 \pmod{4}$ and $c \equiv 3 \pmod{4}$.

Proof. As before, we assume for a contradiction that (14.8) has a rational solution (x, y) and write $x = X/Z$ and $y = Y/Z^2$, where X, Y, Z are integers, $Z \neq 0$ and X, Z are mutually prime (so are not both even). We consider the reduction of Equation (14.8) modulo 4 :

$$Y^2 = 2X^4 + 3X^2Z^2 + 2Z^4.$$

If X and Z are not both odd, then $Y^2 \equiv 2 \pmod{4}$. If both X and Z are odd, then $Y^2 \equiv 3 \pmod{4}$. In both cases, we have a contradiction since $Y^2 \equiv 0$ or $1 \pmod{4}$, for all integers Y . \square

We can now prove the following.

Proposition 14.29. *Consider the following pair of quadrics intersecting in a smooth quartic with one real affinely finite component :*

$$\begin{aligned} Q_S : 2x^2 - 2xy + 2xz - 2xw + y^2 + 4yz - 4yw + 2z^2 - 4zw &= 0, \\ Q_T : x^2 - 2xy + 4xz + 4xw - y^2 + 2yz + 4yw + 4zw - 2w^2 &= 0. \end{aligned}$$

Then the square root $\sqrt{\delta}$ is necessary to parameterize the curve of intersection.

Proof. The characteristic polynomial has two simple real roots so it is immediate that the intersection of Q_S and Q_T is a real smooth quartic with one affinely finite component, by Theorem 14.25.

We again apply Corollary 14.24 and show that the square root $\sqrt{\delta}$ is necessary to parameterize the curve of intersection. We have :

$$\begin{aligned} \sigma^2 &= \det(S + \lambda T), \\ &= 22\lambda^4 + 48\lambda^3 - 9\lambda^2 + 60\lambda + 30, \\ &\equiv 2\lambda^4 + 3\lambda^2 + 2 \pmod{4}, \end{aligned}$$

which has no rational solution by Lemma 14.28, so $\sqrt{\delta}$ cannot be avoided. \square

Two real affinely infinite components

We again prove a preliminary result.

Lemma 14.30. *The equation*

$$y^2 = a(x^4 + x + 1) + bx^3 + cx^2 \tag{14.9}$$

has no rational solution if $a \equiv 2 \pmod{4}$, $b \equiv 0 \pmod{4}$ and $c \equiv 1 \pmod{4}$.

Proof. We proceed as in Lemmas 14.26 and 14.28, and consider the reduction of Equation (14.9) modulo 4 :

$$Y^2 = 2X^4 + X^2Z^2 + 2XZ^3 + 2Z^4.$$

If X is even and Z is odd, the equation reduces to $Y^2 = 2XZ + 2 \equiv 2 \pmod{4}$. If X is odd and Z is even, we also have $Y^2 \equiv 2 \pmod{4}$. Finally, if both X and Z are odd, (14.9) reduces to $Y^2 = 1 + 2XZ \equiv 3 \pmod{4}$. In all cases, we have a contradiction since $Y^2 \equiv 0$ or $1 \pmod{4}$, for all integers Y . \square

This is enough to prove the following.

Proposition 14.31. *Consider the following pair of quadrics intersecting in a smooth quartic with two real affinely infinite components :*

$$\begin{aligned} Q_S : x^2 - 2y^2 + 4zw &= 0, \\ Q_T : xy + z^2 + 2zw - w^2 &= 0. \end{aligned}$$

Then the square root $\sqrt{\delta}$ is necessary to parameterize the curve of intersection.

Proof. The characteristic polynomial has four simple complex roots so it is immediate that the intersection of Q_S and Q_T is a real smooth quartic with two affinely infinite components, by Theorem 14.25.

We again apply Corollary 14.24. We have :

$$\begin{aligned} \sigma^2 &= \det(S + \lambda T), \\ &= 2\lambda^4 + 4\lambda^3 + 5\lambda^2 + 2\lambda + 2, \\ &\equiv 2\lambda^4 + \lambda^2 + 2\lambda + 2 \pmod{4}, \end{aligned}$$

which has no rational solution by Lemma 14.30, so $\sqrt{\delta}$ cannot be avoided. \square

Output 1 Execution trace for Example 1.

```

>> quadric 1: 6*x*y + 5*y^2 + 2*z^2 + 6*z*w - w^2
>> quadric 2: 3*x^2 + y^2 - z^2 + 11*w^2

>> launching intersection
>> characteristic polynomial: 33*1^4 - 124*1^3*m + 137*1^2*m^2 - 32*1*m^3 - 11*m^4
>> gcd of derivatives of characteristic polynomial: 1
>> number of real roots: 4
>> intervals: ]-4, 0[, ]0, 1[, ]2/2^1, 3/2^1[, ]3/2^1, 4/2^1[
>> picked test point 1 at [ -4 1 ], sign > 0 -- inertia [ 2 2 ] found
>> picked test point 2 at [ 1 1 ], sign > 0 -- inertia [ 4 0 ] found
>> complex intersection: smooth quartic
>> real intersection: empty
>> end of intersection

>> time spent: 10 ms

```

14.8 Examples

We now give several examples of computing a parameterization of the intersection in case the intersection of two quadrics is a smooth quartic. The examples presented cover the range of morphologies discussed in the previous section and illustrate all aspects of optimality and near-optimality. For more examples, see [LPP06]. All parameterizations have been computed with a C++ implementation of our intersection software (see [LPP06]).

14.8.1 Example 1

Our first example consists of the quadrics given in Output 1. The gcd of the partial derivatives of the characteristic polynomial is 1, so the intersection consists of a (possibly complex) smooth quartic. Since the characteristic polynomial is found to have four real roots, the intersection, over the reals, is either empty or made of two real affinely finite components (Theorem 14.25). We find a sample quadric in each of the intervals on which $\mathcal{D}(\lambda, \mu)$ is positive and compute its inertia. In the first interval, we find a quadric of inertia (2, 2) so we proceed. In the second interval, we find a quadric of inertia (4, 0). By Theorem 14.5, we conclude the intersection is empty of real points.

14.8.2 Example 2

Our second example is as in Output 2. The gcd of the two partial derivatives of the characteristic polynomial is 1, so the intersection (over \mathbb{C}) is a smooth quartic. The fact that the characteristic polynomial has two real roots implies that the smooth quartic is real and that it consists of one affinely finite component (Theorem 14.25). Here, the two input quadrics have inertia (3, 1) and a first quadric Q_R of inertia (2, 2) is found in the pencil between the two roots of \mathcal{D} . A point is taken on Q_R and then approximated by a point with integer coordinates. It turns out that the approximation, i.e. (0, 0, 1, 0), also lies on Q_R . We thus use this quadric to parameterize the intersection. Since the determinant of Q_R is a square, it can be rationally parameterized (Proposition 14.17). The end of the calculation is as in Section 14.4.

14.8.3 Example 3

Our third example is Example 5 from [WJG02]. It is the intersection of a sphere and an ellipsoid that are very close to one another. The output of our implementation on that example is shown in Output 3. Since the characteristic polynomial has four simple real roots, the intersection is either empty or made of two real affinely finite components (Theorem 14.25). Picking a sample quadric in each of the intervals on which $\det R(\lambda, \mu)$ is positive shows that the pencil contains no quadric of inertia (4, 0), so the quartic is real. Here, the determinant of the quadric of inertia (2, 2) used to parameterize the intersection is not a square, so the parameterization of the quartic contains the square root of some integer. It is thus only near-optimal in the sense that this square root can possibly be avoided.

Output 2 Execution trace for Example 2.

```

>> quadric 1: x^2 - x*y - y^2 - y*w + z^2 + w^2
>> quadric 2: 2*x^2 - x*y + y^2 - y*z + y*w + z^2

>> launching intersection
>> characteristic polynomial: - 6*1^4 - 12*1^3*m + 3*1^2*m^2 + 6*1*m^3 - 2*m^4
>> gcd of derivatives of characteristic polynomial: 1
>> complex intersection: smooth quartic
>> real intersection: smooth quartic, one real affinely finite component
>> number of real roots: 2
>> intervals: ]-2, -1[, ]-1, 0[
>> picked test point 1 at [ -1 1 ], sign > 0 -- inertia [ 2 2 ] found
>> quadric (2,2) found: x^2 + 2*y^2 - y*z + 2*y*w - w^2
>> decomposition of its determinant [a,b] (det = a^2*b): [ 2 1 ]
>> a point on the quadric: [ 0 0 1 0 ]
>> param of quadric (2,2): [- s*u + t*v, - 2*s*v, (2*s + 2*t)*u + (- 4*s + 2*t)*v, s*u + t*v]
>> status of smooth quartic param: optimal
>> end of intersection

>> parameterization of smooth quartic, branch 1:
[- 4*u^3 + u^2*v + 6*u*v^2 + 2*v^3 - u*sqrt(Delta), - 6*u^3 - 8*u^2*v - 4*u*v^2, - 4*u^3
+ 2*u^2*v + (2*u + 2*v)*sqrt(Delta), 4*u^3 + 5*u^2*v + 2*u*v^2 + 2*v^3 + u*sqrt(Delta)]
>> parameterization of smooth quartic, branch 2:
[- 4*u^3 + u^2*v + 6*u*v^2 + 2*v^3 + u*sqrt(Delta), - 6*u^3 - 8*u^2*v - 4*u*v^2, - 4*u^3
+ 2*u^2*v + (- 2*u - 2*v)*sqrt(Delta), 4*u^3 + 5*u^2*v + 2*u*v^2 + 2*v^3 - u*sqrt(Delta)]
Delta = - 2*u^4 + 10*u^3*v - 9*u^2*v^2 - 8*u*v^3 - 2*v^4

>> time spent: 10 ms

```

It turns out that in this particular example it can be avoided. Consider the cone Q_R corresponding to the rational root $(\lambda_0, \mu_0) = (-1, 21)$ of the characteristic polynomial :

$$Q_R : -Q_S + 21Q_T = 2x^2 - y^2 - w^2.$$

Q_R contains the obvious rational point $(1, 1, 0, 1)$, which is not its singular point. This implies that it can be rationally parameterized by Proposition 14.14. Plugging this parameterization in the equation of Q_S or Q_T gives a simple parameterization for the smooth quartic :

$$\mathbf{X}(u, v) = \begin{pmatrix} u^2 + 2v^2 \\ 2uv \\ u^2 - 2v^2 \\ 0 \end{pmatrix} \pm \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix} \sqrt{2u^4 + 4u^2v^2 + 8v^4}.$$

14.8.4 Example 4

Our last example is the one of Proposition 14.31. The result is shown in Output 4. Here, again, the gcd of the partial derivatives of the characteristic polynomial is 1, so the intersection curve is, over \mathbb{C} , a smooth quartic. But since $\mathcal{D}(\lambda, \mu)$ has in fact no real root, we know by Theorem 14.25 that the smooth quartic is real and has two affinely infinite components. Here, the intermediate quadric Q_R of inertia $(2, 2)$ found (which is in fact Q_T) is such that its determinant is not a square. So the parameterization of the quartic contains a square root. Our implementation cannot decide whether this square root is needed or not, so outputs that the parameterization is near-optimal. In this particular example, we know in fact that the parameterization is optimal, by Proposition 14.31.

14.9 Conclusion

The generic algorithm introduced in Section 14.4 already represents a substantial improvement over Levin's pencil method and its subsequent refinements. Indeed, we proved that, when the intersection is a smooth quartic

Output 3 Execution trace for Example 3.

```

>> quadric 1: 19*x^2 + 22*y^2 + 21*z^2 - 20*w^2
>> quadric 2: x^2 + y^2 + z^2 - w^2

>> launching intersection
>> characteristic polynomial: - 175560*1^4 - 34358*1^3*m - 2519*1^2*m^2 - 82*1*m^3 - m^4
>> gcd of derivatives of characteristic polynomial: 1
>> number of real roots: 4
>> intervals: ]-14/2^8, -13/2^8[, ]-26/2^9, -25/2^9[, ]-25/2^9, -24/2^9[, ]-3/2^6, -2/2^6[
>> picked test point 1 at [ -13 256 ], sign > 0 -- inertia [ 2 2 ] found
>> picked test point 2 at [ -3 64 ], sign > 0 -- inertia [ 2 2 ] found
>> complex intersection: smooth quartic
>> real intersection: smooth quartic, two real affinely finite components
>> quadric (2,2) found: - 16*x^2 + 5*y^2 - 2*z^2 + 9*w^2
>> decomposition of its determinant [a,b] (det = a^2*b): [ 12 10 ]
>> a point on the quadric: [ 3 0 0 4 ]
>> param of quadric (2,2): [0, - 24*s*u - 24*t*v, 0, 0] + sqrt(10)*[3*t*u + 6*s*v, 0,
    12*s*u - 12*t*v, - 4*t*u + 8*s*v]
>> status of smooth quartic param: near-optimal
>> end of intersection

>> parameterization of smooth quartic, branch 1:
[(72*u^3 + 4*u*v^2)*sqrt(10) + 3*v*sqrt(10)*sqrt(Delta), - 340*u^2*v + 10*v^3
 - 24*u*sqrt(Delta), (- 118*u^2*v + 5*v^3)*sqrt(10) + 12*u*sqrt(10)*sqrt(Delta),
 (96*u^3 - 12*u*v^2)*sqrt(10) - 4*v*sqrt(10)*sqrt(Delta)]
>> parameterization of smooth quartic, branch 2:
[(72*u^3 + 4*u*v^2)*sqrt(10) - 3*v*sqrt(10)*sqrt(Delta), - 340*u^2*v + 10*v^3
 + 24*u*sqrt(Delta), (- 118*u^2*v + 5*v^3)*sqrt(10) - 12*u*sqrt(10)*sqrt(Delta),
 (96*u^3 - 12*u*v^2)*sqrt(10) + 4*v*sqrt(10)*sqrt(Delta)]
Delta = 20*u^4 - 140*u^2*v^2 + 5*v^4

>> time spent: 10 ms

```

(the generic case) our algorithm computes a parameterization which is optimal in the number of radicals involved up to one possibly unnecessary square root. We also showed that deciding (in all cases) whether this extra square root can be avoided is out of reach, and that the parameterization is optimal in some cases. Moreover, for the first time, our algorithms enable to compute in practice an exact form of the parameterization of two arbitrary quadrics with rational coefficients.

Even though this first part of our paper has focused on the generic, smooth quartic case, this algorithm can also be used when the intersection is singular. Assume the intermediate quadric Q_R has inertia $(2, 2)$. When the curve of intersection consists of a cubic and a line, the equation Ω in the parameters has a cubic factor of bidegree $(2, 1)$ and a linear factor of bidegree $(0, 1)$, in view of Fact 14.19. Similarly, when the curve of intersection consists of a conic and two lines, Ω factors in a quadratic factor of bidegree $(1, 1)$ and two linear factors of bidegree $(1, 0)$ and $(0, 1)$. Thus, assuming we know how to factor Ω , we have a way to parameterize each component of the intersection.

Unfortunately, this does not always lead to a parameterization of the intersection that involves only rational functions. When the intersection C is a singular quartic, Ω is irreducible since C itself is, and solving Ω for s in terms of u (or the converse) introduces the square root of a polynomial, while we know that there exists a parameterization of C with rational functions (the genus of the curve is 0).

Always computing parameterizations with rational functions when such parameterizations are known to exist will necessitate rethinking the basic philosophy of our algorithm. Essentially, while the idea of the generic algorithm is to use the rational quadric with *largest* rank as intermediate quadric for parameterizing the intersection, the refined method will instead use the rational quadric with *smallest* rank as intermediate quadric.

Proceeding that way will have the double benefit of always computing the simplest possible parameterizations and much better controlling the size of their coefficients. The price we pay is a multiplicity of cases and the need to write dedicated software for each (real projective) type of intersection. This is the subject of Parts II and III of this paper.

Output 4 Execution trace for Example 4.

```

>> quadric 1: x^2 - 2*y^2 + 4*z*w
>> quadric 2: x*y + z^2 + 2*z*w - w^2

>> launching intersection
>> characteristic polynomial: 2*1^4 + 4*1^3*m + 5*1^2*m^2 + 2*1*m^3 + 2*m^4
>> gcd of derivatives of characteristic polynomial: 1
>> number of real roots: 0
>> complex intersection: smooth quartic
>> real intersection: smooth quartic, two real affinely infinite components
>> quadric (2,2) found: x*y + z^2 + 2*z*w - w^2
>> decomposition of its determinant [a,b] (det = a^2*b): [ 2 2 ]
>> a point on the quadric: [ 1 0 0 0 ]
>> param of quadric (2,2): [4*t*u, - 2*s*v, s*u + t*v, s*u + t*v]
+ sqrt(2)*[0, 0, 0, - s*u + t*v]
>> status of smooth quartic param: near-optimal
>> end of intersection

>> parameterization of smooth quartic, branch 1:
[- 4*u*v^2 + 4*v*sqrt(Delta), - 2*u^3 - 8*u*v^2 + 2*u^3*sqrt(2), 4*v^3 - u^2*v*sqrt(2)
+ u*sqrt(Delta), - 2*u^2*v + 4*v^3 + (u^2*v + 4*v^3)*sqrt(2)
+ (u - u*sqrt(2))*sqrt(Delta)]
>> parameterization of smooth quartic, branch 2:
[- 4*u*v^2 - 4*v*sqrt(Delta), - 2*u^3 - 8*u*v^2 + 2*u^3*sqrt(2), 4*v^3 - u^2*v*sqrt(2)
- u*sqrt(Delta), - 2*u^2*v + 4*v^3 + (u^2*v + 4*v^3)*sqrt(2)
- (u - u*sqrt(2))*sqrt(Delta)]
Delta = 2*u^4 + 10*u^2*v^2 - 4*v^4 + (- 2*u^4 - 4*v^4)*sqrt(2)

>> time spent: 10 ms

```

14.10 Appendix : The parameterizations of Table 14.3 are faithful

We prove in this section that the parameterizations of Table 14.3 are not only faithful parameterizations of the projective quadrics (in the sense that they define one-to-one correspondences between a dense open subset of the space of the parameters and a dense open subset of the quadric) but they are bijections between the space of the parameters and the quadric. The following two lemmas deal with the parameterizations of quadrics of inertia (2, 2) and (2, 1). For other types of quadrics, it is straightforward to show that the parameterizations of Table 14.3 are bijections.

Lemma 14.32. $(u, v), (s, t) \mapsto \left(\frac{ut+a_1vs}{a_1}, \frac{us-a_2vt}{a_2}, \frac{ut-a_1vs}{\sqrt{a_1a_3}}, \frac{us+a_2vt}{\sqrt{a_2a_4}}\right)$ is a bijection from $\mathbb{P}^1 \times \mathbb{P}^1$ onto the surface $\{(x_1, x_2, x_3, x_4) \in \mathbb{P}^3 \mid a_1x_1^2 + a_2x_2^2 - a_3x_3^2 - a_4x_4^2 = 0\}$, where a_1, a_2, a_3, a_4 are positive.

Proof. To prove this lemma, we apply the change of coordinates in \mathbb{P}^3

$$X = \frac{a_1x_1 + \sqrt{a_1a_3}x_3}{2}, Y = \frac{a_1x_1 - \sqrt{a_1a_3}x_3}{2a_1}, Z = \frac{a_2x_2 + \sqrt{a_2a_4}x_4}{2}, W = \frac{-a_2x_2 + \sqrt{a_2a_4}x_4}{2a_2},$$

or equivalently

$$x_1 = \frac{X + a_1Y}{a_1}, x_3 = \frac{X - a_1Y}{\sqrt{a_1a_3}}, x_2 = \frac{Z - a_2W}{a_2}, x_4 = \frac{Z + a_2W}{\sqrt{a_2a_4}}.$$

In the new frame, the equation of the surface is $XY - ZW = 0$ and the map becomes

$$\Phi : (u, v), (s, t) \mapsto (X, Y, Z, W) = (ut, vs, us, vt).$$

The map Φ is clearly a map from $\mathbb{P}^1 \times \mathbb{P}^1$ into \mathbb{P}^3 because $\Phi((\lambda u, \lambda v), (\mu s, \mu t)) = \lambda \mu \Phi((u, v), (s, t))$ and $\Phi((u, v), (s, t)) = (0, 0, 0, 0)$ if and only if $(u, v) = (0, 0)$ or $(s, t) = (0, 0)$. Moreover, the image of Φ is clearly

included in the surface of equation $XY - ZW = 0$. Conversely, if (X, Y, Z, W) is a point of this surface, at least one of its coordinates is non zero (we are in a projective space), and by symmetry we may suppose that $X \neq 0$. Considering $(X, Z, W) = (ut, us, vt)$, we have $ut \neq 0$, $\frac{Z}{X} = \frac{s}{t}$, and $\frac{W}{X} = \frac{v}{u}$. Thus $\frac{Z}{X}$ uniquely defines (s, t) up to a constant factor and similarly for $\frac{W}{X}$ and (u, v) , which shows the injectivity of Φ . Furthermore, $XY - ZW = 0$ implies $Y = \frac{ZW}{X} = \frac{us \cdot vt}{ut} = vs$ which shows that Φ is surjective. \square

Recall that \mathbb{P}^{*2} denotes the quasi-projective space defined as the quotient of $\mathbb{R}^3 \setminus \{0, 0, 0\}$ by the equivalence relation \sim where $(x_1, x_2, x_3) \sim (y_1, y_2, y_3)$ if and only if $\exists \lambda \in \mathbb{R} \setminus \{0\}$ such that $(x_1, x_2, x_3) = (\lambda y_1, \lambda y_2, \lambda^2 y_3)$.

Lemma 14.33. $(u, v, s) \mapsto (uv, \frac{u^2 - a_1 a_2 v^2}{2a_2}, \frac{u^2 + a_1 a_2 v^2}{2\sqrt{a_2 a_3}}, s)$ is a bijection from \mathbb{P}^{*2} onto the surface $\{(x_1, x_2, x_3, x_4) \in \mathbb{P}^3 \mid a_1 x_1^2 + a_2 x_2^2 - a_3 x_3^2 = 0\}$, where a_1, a_2, a_3 are positive.

Proof. For this lemma, we consider the change of coordinates in \mathbb{P}^3

$$X = x_1, Y = \sqrt{a_2 a_3} x_3 + a_2 x_2, Z = \frac{\sqrt{a_2 a_3} x_3 - a_2 x_2}{a_1 a_2}, W = x_4,$$

or equivalently

$$x_1 = X, x_2 = \frac{Y - a_1 a_2 Z}{2a_2}, x_3 = \frac{Y + a_1 a_2 Z}{2\sqrt{a_2 a_3}}, x_4 = W.$$

In the new frame, the equation of the surface is $X^2 - YZ = 0$ and the map becomes

$$\Psi : (u, v, s) \mapsto (X, Y, Z, W) = (uv, u^2, v^2, s).$$

The map Ψ is clearly a map from \mathbb{P}^{*2} into \mathbb{P}^3 because $\Psi(\lambda u, \lambda v, \lambda^2 s) = \lambda^2 \Psi(u, v, s)$ and $\Psi(u, v, s) = (0, 0, 0, 0)$ if and only if $(u, v, s) = (0, 0, 0)$. Moreover, the image of Ψ is clearly included in the surface of equation $X^2 - YZ = 0$. Conversely, if (X, Y, Z, W) is a point of this surface, then we have to prove that its preimage consists in exactly one point of \mathbb{P}^{*2} . If $Y = Z = 0$, we have also $X = 0$ and a point of the preimage should satisfy $u = v = 0$; it is therefore unique (in \mathbb{P}^{*2}) and it exists by $\Psi(0, 0, W) = (0, 0, 0, W)$.

If Y or Z is nonzero, we may suppose by symmetry that $Y \neq 0$. Considering $(X, Y, W) = (uv, u^2, s)$ we have $u \neq 0$, $\frac{X}{Y} = \frac{v}{u}$ and $\frac{W}{Y} = \frac{s}{u^2}$. Thus $\frac{X}{Y}$ and $\frac{W}{Y}$ uniquely define $(u, v, s) \in \mathbb{P}^{*2}$ which implies that Ψ is injective. Furthermore, $YZ = X^2$ implies $Z = \frac{X^2}{Y} = \frac{(uv)^2}{u^2} = v^2$ which shows that Ψ is surjective. \square

Remark 14.34. Although the statements and the proofs of Lemma 14.32 and 14.33 are very similar, there is a big difference between the two bijections : the bijection is an isomorphism and a diffeomorphism in Lemma 14.32 but not in Lemma 14.33 where the space of the parameters is smooth while the surface is singular at $(0, 0, 0, 1)$.

Chapitre 15

Near-optimal parameterization of the intersection of quadrics : II. A classification of pencils

Cet article a été accepté dans *Journal of Symbolic Computation* [DLLP07a]. (Voir également le rapport de recherche [DLLP05b].) Une version préliminaire à été publiée dans les proceedings du *19th ACM Annual Symposium on Computational Geometry* [DLLP03] et dans la thèse de L. Dupont [Dup04].

Abstract

While Part I of this paper was devoted mainly to quadrics intersecting in a smooth quartic, we now focus on singular intersections. To produce optimal or near-optimal parameterizations in all cases, we first determine the the real type of the intersection before computing the actual parameterization. In this second part, we present the first classification of pencils of quadrics based on the type of their intersection in real projective space and we show how this classification can be used to compute efficiently the type of the real intersection. The near-optimal parameterization algorithms in all singular cases will be given in Part III.

15.1 Introduction

At the end of Part I, we saw that the generic algorithm we introduced, while being simple and giving optimal parameterizations in some cases, fails to achieve the stated goal of computing (near-)optimal parameterizations (both in terms of functions and coefficients) of intersections of arbitrary quadrics.

Unfortunately, it turns out that achieving this goal involves more than simple adaptations to the generic algorithm. Reaching optimality implies looking carefully at *each* type of *real* intersection and designing a dedicated algorithm to handle each situation. For this, we need to understand precisely which situations can happen over the reals and thus classify real pencils of quadrics of $\mathbb{P}^3(\mathbb{R})$.

Classifying pencils of quadrics over the complexes was achieved by Segre in the nineteenth century [Seg83]. Its practical value is however limited since its proper interpretation lies in the complex domain (i.e. points on the intersection might be real or complex), whereas our concern is with real parts of the intersection.

Accordingly, we refine the Segre classification of pencils of $\mathbb{P}^3(\mathbb{R})$ by examining the different cases occurring over the reals. This refinement is, in itself, of partial assistance for the parameterization problem : no more than the Segre classification can it be “reverse engineered” to construct explicit representations of the various intersection components. It is however mandatory for the following two reasons : it allows us to obtain structural information on the intersection curve which we use to drive the algorithm for computing a near-optimal parameterization of the intersection curve (Part III) ; it is also a prerequisite for proving the (near-)optimality of our parameterization algorithm.

Contributions. In this second part of our paper, we present a new classification of pencils of quadrics based on the type of their real projective intersection. A summary of this classification is given in Tables 15.4 and 15.5. We then show how to use this classification to compute efficiently the type of the real intersection. In particular we show how computations with non-rational numbers can be avoided for detecting the type of the intersection when the input quadrics have rational coefficients.

It should be stressed that, even though the classification of pencils over the reals is presented here as an intermediate step in a more global process (i.e., parameterization of the intersection), this classification has an interest on its own. It could be used for instance in a collision detection context to predict at which time stamps a collision between two moving objects will happen.

Related work. In the context of the representation of the geometry of Boolean combinations of volumes bounded by quadric surfaces, J. Ocken, J. T. Schwartz, and M. Sharir showed in 1987 showed how two quadrics can be simultaneously diagonalized using a real projective transformation and used this diagonalization to parameterize the intersection of the quadrics. The analysis is however incomplete and some intersection morphologies are overlooked, leading to possible misclassifications. In particular, the cases when the characteristic polynomial of the pencil has two double roots, corresponding to such morphologies as a cubic and a secant line or four lines forming a skew quadrilateral, are missing.

The next result on the classification of pencil of quadrics based on the real type of the intersection was obtained in 2002 by C. Tu, W. Wang, and J. Wang who who classified pencils in the generic case, that is when the intersection is a smooth quartic (in complex space) whose number of connected components (in real space) is two, one, or zero [TWW02]. Note that W. Wang and R. Krasauskas also obtained results on the classification of pencils in the generic case when the pencil is furthermore restricted to be generated by two ellipsoids in affine space [WK04]. Related results have also been obtained by W. Wang, J. Wang, and M. Kim on the separation of two ellipsoids in affine space [WWK01].

In September 2005, Tu, Wang, Mourrain, and Wang published a research report [TWMW05] presenting a classification of pencils very similar to ours. They use the Canonical Pair Form Theorem of F. Uhlig as basic mathematical tool and refine the classification of pencils of quadrics over the complexes in exactly the same way as we do. There are however differences between the two approaches. First, we classify pencils using the inertia of the quadrics at the multiple roots of the characteristic polynomial, except for a small number of cases where simple geometric conditions allow to discriminate. By contrast, Tu et al. classify pencils using the inertia of the quadrics between the roots of the characteristic polynomial (plus the degree of the minimal polynomial of the characteristic polynomial in some cases), and rely on Puiseux expansion to deduce some information at the (multiple) roots. Second, the classification of Tu et al. is limited to non-degenerate pencils (i.e. pencils whose characteristic polynomial does not vanish identically), while ours covers all possible cases. Third, in addition to the enumeration of all real quadric intersection morphologies, we also provide algorithms for exactly and efficiently recovering the real projective type of the intersection of two arbitrary given quadrics.

The rest of this part is organized as follows. Section 15.2 reviews the classical Segre classification of pencils of quadrics over the complexes. We then refine, in Sections 15.3 and 15.4, the Segre classification over the reals with a repeated application of the Canonical Pair Form Theorem for pairs of real symmetric matrices introduced in Part I. In Section 15.3, we consider *regular* pencils, i.e., pencils that contain a non-singular quadric, and, in Section 15.4, *singular* pencils, i.e., pencils that contain only singular quadrics, or, equivalently, pencils with identically vanishing characteristic polynomial. In Section 15.5, we use the results of the classification of pencils over the reals to design an algorithm to quickly and efficiently characterize the complex and real types of the intersection given two input quadrics. Several examples are detailed in Section 15.6, before concluding.

15.2 Classification of pencils of quadrics over the complexes

In this section, we review classical material on the classification of pencils of quadrics. It will serve as the starting point for our classification of pencils over the reals in Sections 15.3 and 15.4.

In the rest of the paper all quadrics are considered in real projective space $\mathbb{P}^3(\mathbb{R})$; their coefficients as well as the coefficients of the characteristic polynomials of pencils are thus real. However, we consider the intersection

of quadrics both in $\mathbb{P}^3(\mathbb{R})$ and in $\mathbb{P}^3(\mathbb{C})$. Accordingly, the classification of pencils is considered both over the complexes and over the reals.

We start in Section 15.2.1 with a proof that the existence of a singularity on the intersection curve is equivalent either to the existence of a multiple root in the characteristic polynomial or to the fact that the characteristic polynomial vanishes identically. Then Section 15.2.2 recalls the basic tenets of the classification of pencils over the complexes. The well-known Segre characteristic is recalled in Section 15.2.2.1 and its relation with the Canonical Pair Form Theorem for pairs of real symmetric matrices (Part I and [Uhl73, Uhl76]) is thoroughly explained in Section 15.2.2.2.

15.2.1 Singular intersections and multiple roots

Recall that, given two 4 by 4 real symmetric matrices S and T , the characteristic polynomial, $\mathcal{D}(\lambda, \mu)$, of the pencil generated by S and T is the determinant of $\lambda S + \mu T$. In the ensuing sections, we use the following equivalence for classifying the singular intersections through the multiplicities of the roots of the characteristic polynomial $\mathcal{D}(\lambda, \mu)$ and the rank of the corresponding quadrics.

Proposition 15.1. *If the intersection of two distinct quadrics Q_S and Q_T has a singular point \mathbf{p} , then*

- either $\mathcal{D} \equiv 0$ and Q_S and Q_T are singular at \mathbf{p} ,
- or $\mathcal{D} \equiv 0$ and there is a unique quadric Q_R of the pencil that is singular at \mathbf{p} ,
- or $\mathcal{D} \not\equiv 0$, there is a unique quadric $Q_R = \lambda_0 Q_S + \mu_0 Q_T$ that is singular at \mathbf{p} and (λ_0, μ_0) is a multiple root of \mathcal{D} .

In the last two cases, all the quadrics of the pencil except Q_R share a common tangent plane at \mathbf{p} .

Proof. First recall that a curve C defined by implicit equations $Q_S = Q_T = 0$ is singular at \mathbf{p} if and only if \mathbf{p} is on C and the rank of the Jacobian matrix J of C is strictly less than 2 when evaluated at \mathbf{p} . J is the matrix of partial derivatives :

$$J = \begin{pmatrix} \frac{\partial Q_S}{\partial x} & \frac{\partial Q_S}{\partial y} & \frac{\partial Q_S}{\partial z} & \frac{\partial Q_S}{\partial w} \\ \frac{\partial Q_T}{\partial x} & \frac{\partial Q_T}{\partial y} & \frac{\partial Q_T}{\partial z} & \frac{\partial Q_T}{\partial w} \end{pmatrix}. \quad (15.1)$$

Let J_S and J_T be the first and second rows of J .

If all the coefficients of J vanish at \mathbf{p} , then \mathbf{p} is a singular point of both Q_S and Q_T and thus of all quadrics of the pencil, implying that $\mathcal{D} \equiv 0$.

Otherwise, J has rank one and there exists a linear relationship between the rows of J evaluated at \mathbf{p} :

$$\lambda_0 J_S|_{\mathbf{p}} + \mu_0 J_T|_{\mathbf{p}} = 0, \quad (\lambda_0, \mu_0) \in \mathbb{P}^1(\mathbb{R}).$$

Also, there is a $\mathbf{v} \in \mathbb{P}^3(\mathbb{R})$ such that $J|_{\mathbf{p}}\mathbf{v}$ is a non-zero multiple of \mathbf{v} . This exactly means that the tangent plane at \mathbf{p} of all the quadrics of the pencil is the plane P of equation $\mathbf{v} \cdot (x \ y \ z \ w)^T = 0$, except for the quadric $\lambda_0 Q_S + \mu_0 Q_T$. For this last quadric, all the partial derivatives at \mathbf{p} vanish, implying that it is singular at \mathbf{p} and has rank at most 3.

Now, we may change the generators of the pencil by taking $\lambda_0 S + \mu_0 T$ as first generator in place of S . This has the effect of translating (λ_0, μ_0) to $(1, 0)$. If we change of frame in order that the coordinates of \mathbf{p} become $(0, 0, 0, 1)$ and that the equation of P becomes $x = 0$, the matrices of the generators of the pencil become

$$S' = \begin{pmatrix} * & * & * & 0 \\ * & \mathbf{A} & 0 & 0 \\ * & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad \text{and} \quad T' = \begin{pmatrix} * & * & * & 1 \\ * & \mathbf{B} & 0 & 0 \\ * & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix},$$

where \mathbf{A} and \mathbf{B} are 2×2 matrices and the stars denote any element. It follows immediately that

$$\det(\lambda S' + \mu T') = -\mu^2 \det(\lambda \mathbf{A} + \mu \mathbf{B}).$$

The case $\det(\lambda \mathbf{A} + \mu \mathbf{B}) \equiv 0$ proves the second assertion. The case $\det(\lambda \mathbf{A} + \mu \mathbf{B}) \not\equiv 0$ proves the last assertion. \square

15.2.2 Classification of pencils by elementary divisors

For the reader's convenience, we review, in this section, the classical classification of pencils of quadrics as originally done by the Italian mathematician Segre [Seg83]. More recent and accessible accounts can be found in [Bro06] or [HP53].

15.2.2.1 Segre characteristic

Assume we are given a pencil $R(\lambda, \mu) = \lambda S + \mu T$ of symmetric matrices of size n such that $\mathcal{D}(\lambda, \mu) = \det R(\lambda, \mu)$ is not identically zero. In general, \mathcal{D} has n complex roots to which correspond n complex projective cones of the pencil. But there can be exceptions to this when a root (λ_0, μ_0) of \mathcal{D} appears with multiplicity larger than 1. It can also happen that (λ_0, μ_0) makes not only the determinant \mathcal{D} vanish but also all its subdeterminants of order $n - t + 1$ say, $t > 0$. This means that the corresponding quadric has as singular set a linear space of dimension $t - 1$.

Let the $(\lambda_i, \mu_i), i = 1, \dots, p$, be the roots of \mathcal{D} and the m_i their respective multiplicities. Indicate by m_i^j the minimum multiplicity with which the root (λ_i, μ_i) appears in the subdeterminants of order $n - j$ of \mathcal{D} . Let $t_i \geq 1$ be the smallest integer such that $m_i^{t_i} = 0$. We see that $m_i^j \geq m_i^{j+1}$ for all j . Define a sequence of indices e_i^j as follows :

$$e_i^j = m_i^{j-1} - m_i^j, \quad j = 1, \dots, t_i,$$

with $m_i^0 = m_i$. The multiplicity m_i of (λ_i, μ_i) is the sum $e_i^1 + \dots + e_i^{t_i}$. We have therefore :

$$\mathcal{D}(\lambda, \mu) = (\lambda\mu_i - \mu\lambda_i)^{m_i} \mathcal{D}^*(\lambda, \mu) = (\lambda\mu_i - \mu\lambda_i)^{e_i^1} \dots (\lambda\mu_i - \mu\lambda_i)^{e_i^{t_i}} \mathcal{D}^*(\lambda, \mu),$$

where $\mathcal{D}^*(\lambda_i, \mu_i) \neq 0$.

The factors $(\lambda\mu_i - \mu\lambda_i)^{e_i^j}$ are called the *elementary divisors* and the exponents e_i^j the *characteristic numbers*, associated with the root (λ_i, μ_i) . Their study goes back to Karl Weierstrass [Wei68]. Segre introduced the following notation to denote the various characteristic numbers associated with the degenerate quadrics that appear in a pencil :

$$\sigma_n = [(e_1^1, \dots, e_1^{t_1}), (e_2^1, \dots, e_2^{t_2}), \dots, (e_p^1, \dots, e_p^{t_p})],$$

with the convention that the parentheses enclosing the characteristic numbers of (λ_i, μ_i) are dropped when $t_i = 1$. This is known as the *Segre characteristic* or *Segre symbol* of the pencil.

The following theorem, essentially due to Weierstrass [Wei68], proves that a pencil of quadrics and the intersection it defines are uniquely and entirely characterized, over the complexes, by its Segre symbol.

Theorem 15.2 (Characterization by Segre symbol). *Consider two pencils of quadrics $R(\lambda_1, \mu_1) = \lambda_1 S_1 + \mu_1 T_1$ and $R(\lambda_2, \mu_2) = \lambda_2 S_2 + \mu_2 T_2$ in $\mathbb{P}^n(\mathbb{R})$. Suppose that $\det R(\lambda_1, \mu_1)$ and $\det R(\lambda_2, \mu_2)$ are not identically zero and let $(\lambda_{1,i}, \mu_{1,i})$ and $(\lambda_{2,i}, \mu_{2,i})$ be their respective roots. Then the two pencils are projectively equivalent if and only if they have the same Segre symbol and there is an automorphism of $\mathbb{P}^1(\mathbb{C})$ taking $(\lambda_{1,i}, \mu_{1,i})$ to $(\lambda_{2,i}, \mu_{2,i})$.*

With the above definition, we see that (λ_i, μ_i) is a root of all subdeterminants of $R(\lambda, \mu)$ of order $n - t_i + k, k > 0$, but not of any subdeterminant of order $n - t_i$. In other words, the rank r_i of $R(\lambda_i, \mu_i)$ is $n - t_i$. In addition, since $m_i = e_i^1 + \dots + e_i^{t_i}$, we have that $n - 1 \geq r_i \geq n - m_i$. Enumerating all possible cases for the e_i^j subject to the constraints induced by its definition gives rise to all possible types of (complex) intersection and accompanying Segre symbols. Tables 15.1, 15.2, and 15.3 list the possible cases for pencils in $\mathbb{P}^3(\mathbb{R})$. Incidentally, we can see that the pair (m_i, r_i) is sufficient to characterize the pencil except in the case $(m_i, r_i) = (4, 2)$.

When the characteristic polynomial $\mathcal{D}(\lambda, \mu)$ vanishes identically, i.e., all the quadrics are singular (see Tables 15.2 and 15.3), the above theory does not apply directly. There are two cases, according to whether the quadrics of the pencil have singular points in common or not :

- When they do not, the pencil can be characterized by a different set of invariants the existence of which was originally proved by Kronecker. We do not detail here how this set is computed (but see [Bro06, p. 55-60]). Suffice it to say that the cases $n = 4$ and $n = 3$ are characterized each by a single set of such invariants, designated by the strings $[1\{3\}]$ and $[\{3\}]$ respectively. In Section 15.4.1, we carry out the analysis of this situation when $n = 4$ without resorting to these special invariants.

Segre characteristic σ_4	roots of $\mathcal{D}(\lambda, \mu)$ in \mathbb{C} and rank of associated quadric	type of intersection in $\mathbb{P}^3(\mathbb{C})$
[1111]	four simple roots	smooth quartic
[112]	one double root, rank 3	nodal quartic
[11(11)]	one double root, rank 2	two secant conics
[13]	triple root, rank 3	cuspidal quartic
[1(21)]	triple root, rank 2	two tangent conics
[1(111)]	triple root, rank 1	double conic
[4]	quadruple root, rank 3	cubic and tangent line
[(31)]	quadruple root, rank 2	conic and two lines crossing on the conic
[(22)]	quadruple root, rank 2	two lines and a double line
[(211)]	quadruple root, rank 1	two double lines
[(1111)]	quadruple root, rank 0	smooth quadric
[22]	two double roots, both rank 3	cubic and secant line
[2(11)]	two double roots, ranks 3 and 2	conic and two lines not crossing on the conic
[(11)(11)]	two double roots, both rank 2	four lines (skew quadrilateral)

TAB. 15.1 – Classification of pencils by Segre symbol in the case where $\mathcal{D}(\lambda, \mu)$ does not identically vanish. When the characteristic polynomial has multiple roots, the additional simple roots are not indicated : they correspond to rank 3 quadrics.

Segre characteristic σ_3	roots of $\mathcal{D}_3(\lambda, \mu)$ in \mathbb{C} and rank of associated conic	type of intersection in $\mathbb{P}^3(\mathbb{C})$
[1{3}]	no common singular point	conic and double line
[111]	three simple roots	four concurrent lines
[12]	double root, rank 2	two lines and a double line
[1(11)]	double root, rank 1	two double lines
[3]	triple root, rank 2	line and triple line
[(21)]	triple root, rank 1	quadruple line
[(111)]	triple root, rank 0	cone
{3}	$\mathcal{D}_3(\lambda, \mu) \equiv 0$	see Table 15.3

TAB. 15.2 – Classification of pencils by Segre symbol in the case where $\mathcal{D}(\lambda, \mu) \equiv 0$. When the quadrics of the pencil have (at least) one singular point \mathbf{p} in common (bottom part), $\mathcal{D}_3(\lambda, \mu)$ is the determinant of the 3×3 upper-left matrix of $R(\lambda, \mu)$ after a congruence transformation sending \mathbf{p} to $(0, 0, 0, 1)$. The conic associated with a root of $\mathcal{D}_3(\lambda, \mu)$ corresponds to the 3×3 upper-left matrix or $R(\lambda, \mu)$.

Segre characteristic σ_2	roots of $\mathcal{D}_2(\lambda, \mu)$ in \mathbb{C} and rank of associated matrix	type of intersection in $\mathbb{P}^3(\mathbb{C})$
{3}	no two common singular points	line and plane
[11]	two simple roots	quadruple line
[2]	double root, rank 1	plane
[(11)]	double root, rank 0	pair of distinct planes
	$\mathcal{D}_2(\lambda, \mu) \equiv 0$	double plane

TAB. 15.3 – Classification of pencils by Segre symbol in the case where $\mathcal{D}(\lambda, \mu) \equiv 0$ and $\mathcal{D}_3(\lambda, \mu) \equiv 0$. When the quadrics of the pencil have (at least) two singular point \mathbf{p} and \mathbf{q} in common (bottom part), $\mathcal{D}_2(\lambda, \mu)$ is the determinant of the 2×2 upper-left matrix of $R(\lambda, \mu)$ after a congruence transformation sending \mathbf{p} and \mathbf{q} to $(0, 0, 0, 1)$ and $(0, 0, 1, 0)$. The matrix associated with a root of $\mathcal{D}_2(\lambda, \mu)$ corresponds to the 2×2 upper-left matrix or $R(\lambda, \mu)$.

Segre string	roots of $\mathcal{D}(\lambda, \mu)$ in \mathbb{C}	rank or inertia of $R(\lambda_1, \mu_1)$	rank or inertia of $R(\lambda_2, \mu_2)$	type of (λ_2, μ_2)	s	type of intersection in $\mathbb{P}^3(\mathbb{R})$
[1111]	4 simple roots					smooth quartic or \emptyset ; see [Fin37] & [TWW02] (or also Th. 14.5 & 14.25)
[112]	1 double root	(3, 0)		real		point
[112]	1 double root	(2, 1)		real	-	nodal quartic; isolated node
[112]	1 double root	(2, 1)		real	+	nodal quartic; convex sing.
[112]	1 double root	rank 3		complex		nodal quartic; concave sing.
[11(11)]	1 double root	(2, 0)		real	+	\emptyset
[11(11)]	1 double root	(2, 0)		real	-	two points
[11(11)]	1 double root	(1, 1)	(2, 1)	real	-	two non-secant conics
[11(11)]	1 double root	(1, 1)	(3, 0)	real	-	\emptyset
[11(11)]	1 double root	(1, 1)		real	+	two secant conics; convex sing.
[11(11)]	1 double root	rank 2		complex	-	conic
[11(11)]	1 double root	rank 2		complex	+	two secant conics; concave sing.
[13]	triple root	rank 3				cuspidal quartic
[1(21)]	triple root	(2, 0)				double point
[1(21)]	triple root	(1, 1)				two tangent conics
[1(111)]	triple root	rank 1	(2, 1)			double conic
[1(111)]	triple root	rank 1	(3, 0)			\emptyset
[4]	quadruple root	rank 3				cubic and tangent line
[(31)]	quadruple root	(1, 1)			-	conic
[(31)]	quadruple root	(1, 1)			+	conic and two lines crossing on the conic
[(22)]	quadruple root	(2, 0)				double line
[(22)]	quadruple root	(1, 1)			+	two single lines & a double line
[(211)]	quadruple root	rank 1			-	point
[(211)]	quadruple root	rank 1			+	two secant double lines
[(1111)]	quadruple root	rank 0				any smooth quadric of the pencil
[22]	2 double roots	rank 3	rank 3	real		cubic and secant line
[22]	2 double roots	rank 3	rank 3	complex		cubic and non-secant line
[2(11)]	2 double roots	(3, 0)	rank 2	real		point
[2(11)]	2 double roots	(2, 1)	rank 2	real	+	conic and two intersecting lines
[2(11)]	2 double roots	(2, 1)	rank 2	real	-	conic and point
[(11)(11)]	2 double roots	(2, 0)	(2, 0)	real		\emptyset
[(11)(11)]	2 double roots	(2, 0)	(1, 1)	real		two points
[(11)(11)]	2 double roots	(1, 1)	(2, 0)	real		two points
[(11)(11)]	2 double roots	(1, 1)	(1, 1)	real		four lines (skew quadrilateral)
[(11)(11)]	2 double roots	rank 2	rank 2	complex		two secant lines

TAB. 15.4 – Classification of pencils in the case where $\mathcal{D}(\lambda, \mu)$ does not identically vanish. (λ_1, μ_1) denotes a multiple root of $\mathcal{D}(\lambda, \mu)$ (if any) and (λ_2, μ_2) another root (not necessarily simple). If (λ_1, μ_1) is a double root then s denotes the sign of $\frac{\det(\lambda S + \mu T)}{(\mu_1 \lambda - \lambda_1 \mu)^2}$ at $(\lambda, \mu) = (\lambda_1, \mu_1)$; if (λ_1, μ_1) is a quadruple root then s denotes the sign of $\det(\lambda S + \mu T)$ for any $(\lambda, \mu) \neq (\lambda_1, \mu_1)$. When the characteristic polynomial has multiple roots, the additional simple roots are not indicated. The Segre characteristic is mentioned for clarity but is not needed for the classification.

Segre string	roots of $\mathcal{D}_3(\lambda, \mu)$ in \mathbb{C}	rank or inertia of $R(\lambda_1, \mu_1)$	inertia of $R(\lambda_2, \mu_2)$	type of (λ_2, μ_2)	type of intersection in $\mathbb{P}^3(\mathbb{R})$
[1{3}]	no common singular point				conic and double line
[111]	3 simple roots	(1, 1)	(1, 1)	real	four concurrent lines meeting at p point p point p two lines intersecting at p
[111]	3 simple roots	(2, 0)	(2, 0)	real	
[111]	3 simple roots			real	
[111]	3 simple roots			complex	
[12]	double root	(1, 1)			2 lines and a double line meeting at p double line
[12]	double root	(2, 0)			
[1(11)]	double root	rank 1	(1, 1)		two double lines meeting at p point p
[1(11)]	double root	rank 1	(2, 0)		
[3]	triple root	rank 2			a line and a triple line meeting at p
[(21)]	triple root	rank 1			a quadruple line
[(111)]	triple root	rank 0			any non-trivial quadric of the pencil
	$\mathcal{D}_3(\lambda, \mu) \equiv 0$				same as in Table 15.3

TAB. 15.5 – Classification of pencils in the case where $\mathcal{D}(\lambda, \mu)$ identically vanishes. In the bottom part, the quadrics of the pencil have a singular point **p** in common. $\mathcal{D}_3(\lambda, \mu)$ is the determinant of the 3×3 upper-left matrix of $R(\lambda, \mu)$ after a congruence transformation sending **p** to $(0, 0, 0, 1)$. The conic associated with a root of $\mathcal{D}_3(\lambda, \mu)$ corresponds to the 3×3 upper-left matrix or $R(\lambda, \mu)$. (λ_1, μ_1) denotes the multiple root of $\mathcal{D}_3(\lambda, \mu)$ (if any) and (λ_2, μ_2) another root. When $\mathcal{D}_3(\lambda, \mu)$ has a multiple root, the additional simple roots are not indicated. The Segre characteristic is mentioned for clarity but is not needed for the classification.

- When the quadrics do have (at least one) singular point in common, say **p**, we may suppose, after a change of frame, that **p** has coordinates $(0, \dots, 0, 1)$. In the new frame, the matrices have their last row and column filled with zeros. To sort out the different types of intersection, we may identify the quadrics with their upper left $(n - 1) \times (n - 1)$ matrices and classify the restricted pencils by looking at the Segre symbol σ_{n-1} of their characteristic polynomial (of degree $n - 1$). This is what we have done in Table 15.2 for the case of quadrics in $\mathbb{P}^3(\mathbb{R})$.

The above process can be repeated by recursing on dimension.

15.2.2.2 From the complexes to the reals

Theorem 15.2 can be used to find a canonical form for a pencil of quadrics when $\mathcal{D}(\lambda, \mu)$ is not identically zero (see [HP53]). Consider the pencil $R(\lambda) = \lambda S - T$ and its characteristic polynomial $\mathcal{D}(\lambda)$, with roots λ_i of multiplicity m_i . Let

$$[(e_1^1, \dots, e_1^{t_1}), (e_2^1, \dots, e_2^{t_2}), \dots, (e_p^1, \dots, e_p^{t_p})]$$

be the Segre symbol of the pencil. Then there exists a change of coordinates in $\mathbb{P}^n(\mathbb{C})$ such that, in the new frame, the pencil writes down as $R'(\lambda) = \lambda S' - T'$, where

$$S' = \text{diag}(E_1^1, \dots, E_1^{t_1}, \dots, E_p^1, \dots, E_p^{t_p}), \quad T' = \text{diag}(E_1^1 J_1^1, \dots, E_1^{t_1} J_1^{t_1}, \dots, E_p^1 J_p^1, \dots, E_p^{t_p} J_p^{t_p})$$

are block diagonal matrices with blocks :

$$E_i^k = \begin{pmatrix} 0 & & 1 \\ & \dots & \\ 1 & & 0 \end{pmatrix} \quad \text{and} \quad J_i^k = \begin{pmatrix} \lambda_i & 1 & 0 \\ & \dots & \\ 0 & & \lambda_i \end{pmatrix}$$

of size e_i^k . The parentheses in the Segre symbol correspond one-to-one to the singular quadrics in the pencil. The root of \mathcal{D} corresponding to a singular quadric of symbol $(e_i^1, \dots, e_i^{t_i})$ has multiplicity $m_i = \sum_{k=1}^{t_i} e_i^k$.

The parallel between the Canonical Pair Form Theorem introduced in Section 14.5.1²⁶ and the decomposition by Segre symbol should now jump to mind : the first is in a sense a real version of the second, i.e. it gives a canonical form that is projectively equivalent by a *real* congruence transformation to the original pencil. In the real form, complex roots of the characteristic polynomial are somehow combined in complex Jordan blocks so that quadric pencils are equivalent by a real projective transformation.

When λ_i is real, the J_i^j are the real Jordan blocks associated with λ_i . The sum of the sizes of the blocks corresponding to λ_i is $\sum_{k=1}^{t_i} e_i^k = m_i$ and the number of those blocks is $t_i = n - \text{rank } R(\lambda_i)$, as in Theorem 14.10.

When λ_i is complex, let λ_j be its conjugate. It is intuitively clear that $t_i = t_j$ in the complex decomposition and that the associated Jordan blocks J_i^k and J_j^k have the same sizes, i.e. $e_i^k = e_j^k$. When the complex roots and their blocks are combined, they give rise to complex Jordan blocks of size $2e_i^k$. In the real canonical form, the number of these blocks is again t_i but the sum of their sizes is $2m_i$.

The Segre symbol can thus serve as a starting point for the study of real pencils using the Canonical Pair Form Theorem. We illustrate this with two examples concerning pencils in $\mathbb{P}^3(\mathbb{R})$. Consider first the Segre symbol $[(211)]$. The associated pencil has a quadruple root, which is necessarily real (otherwise its conjugate would also be a root of the characteristic polynomial of the pencil). In view of the above, the real decomposition of the pencil has three Jordan blocks, one of size 2 and two of size 1. Now consider the Segre symbol $[22]$. The associated pencil has two double roots, which can be either both real or both complex. If they are real, then each of the roots has one Jordan block of size 2. If they are complex, then the two roots appear in the same Jordan block of size 4.

15.3 Classification of regular pencils of $\mathbb{P}^3(\mathbb{R})$ over the reals

We now turn to the classification of pencils of quadrics of $\mathbb{P}^3(\mathbb{R})$ over the reals. A summary of this classification is given in Tables 15.4 and 15.5.

In what follows, we make heavy use of the Canonical Pair Form Theorem for pairs of real symmetric matrices (Part I and [Uhl73, Uhl76]). For each possible Segre characteristic, we examine the different cases according to whether the roots of the characteristic polynomial are real or not and then examine the conditions leading to different types of intersection over the reals.

In each case, we start by computing the canonical form of the pair (S, T) for a given Segre characteristic and type (real or complex) of multiple root(s) of the characteristic polynomial. We then deduce from this canonical form a *normal form* of the pencil over the reals by rescaling and translating the roots to particularly simple values. Recall that the congruence transformation in the Canonical Pair Form Theorem preserves the roots (values and multiplicities) of the characteristic polynomial of the pencil. This normal form is in a sense the “simplest pair” of quadrics having a given real intersection type. The normal pencil is equivalent by a real projective transformation to any pencil of quadrics with the same real and complex intersection type.

A word of caution : the projective transformations involved in the classification of real pencils, if they preserve the real type of the intersection, may well involve irrational numbers. This fact should be kept to mind when interpreting the results.

We treat the first case (nodal quartic) in some detail so that the reader gets accustomed to the techniques we use. For the other cases, we move directly to the normal form without first expliciting the canonical form.

Note that, in the case where the Segre characteristic is $[1111]$, which corresponds to a smooth quartic in $\mathbb{P}^3(\mathbb{C})$, the classification on the type of intersection in $\mathbb{P}^3(\mathbb{R})$ follows from results by Finsler [Fin37] and Tu et al. [TWW02] (see also Theorems 14.5 and 14.25). Also, the case $[(1111)]$ does not necessitate any further treatment : save for the quadric corresponding to the quadruple root (which is $\mathbb{P}^3(\mathbb{R})$), all the quadrics of the pencil are equal and the intersection is thus any of those non-trivial quadrics. Finally, the case of a vanishing characteristic polynomial, $\det R(\lambda, \mu) \equiv 0$, is treated separately in Section 15.4.

Here and in the ensuing sections, a singularity of the intersection will be called *convex* if the branches of the curve are on the same side of the common tangent plane to the branches at the singularity, *concave* otherwise. It should be stressed that there is a close connection between the type of the singularity and the notion of affine finiteness introduced by Tu et al. [TWW02]. Recall that the point set is called *affinely finite* if it does not intersect some

²⁶When reference is made to a section or result in another part of the paper, it is prefixed by the part number.

projective plane and *affinely infinite* otherwise. As we shall see, a convex singularity corresponds to an affinely finite intersection, while a concave one corresponds to an affinely infinite intersection. Furthermore, our classification directly yields the following theorem which provides a global property on the intersection of two quadrics from a property of the pencil; note that this theorem is similar in spirit to the theorem due to Finsler [Fin37] (see Theorem 14.5) which characterizes when two quadrics have an empty intersection.

Theorem 15.3. *If two distinct quadrics have a pencil whose characteristic polynomial does not identically vanish, their intersection is affinely finite if and only if there exists a quadric of inertia (3, 1) in the pencil.*

Proof. Any quadric of inertia (3, 1) is affinely finite, thus if the pencil contains such a quadric, the intersection is affinely finite. Conversely, in the case where the intersection is a smooth quartic (in $\mathbb{P}^3(\mathbb{C})$), the property follows from Tu et al. [TWW02] (see also Theorem 14.25). Otherwise, it follows from our classification that, when there is no quadric of inertia (3, 1) in a pencil generated by two distinct quadrics, the intersection either contains a line (and therefore is affinely infinite) or is a nodal quartic with a concave singularity or two secant conics with a concave singularity. In the last two cases, we show below that the intersection is affinely infinite. \square

An additional benefit of the classification of pencils over the reals is the ability to draw pictures of all possible situations. Such a gallery of intersection cases is given in Figure 15.1. The pictures were made with the *surf* visualization tool [Sur].

15.3.1 Nodal quartic in $\mathbb{P}^3(\mathbb{C})$, $\sigma_4 = [112]$

The characteristic polynomial has a double root λ_1 , which is necessarily real (otherwise its conjugate would also be a double root of $\det R(\lambda)$). Let λ_2 and λ_3 be the other roots. The Segre characteristic implies that the three quadrics $R(\lambda_i)$ have rank 3 (equal to $n - t_i$; see Section 15.2.2.1). The Canonical Pair Form Theorem thus implies that to λ_1 corresponds one real Jordan block of size 2.

There are two cases.

λ_2 and λ_3 are real. $R(\lambda_2)$ and $R(\lambda_3)$ are projective cones. The Canonical Pair Form Theorem gives that S and T are simultaneously congruent to the quadrics of equations

$$\begin{cases} 2\varepsilon_1 xy + \varepsilon_2 z^2 + \varepsilon_3 w^2 = 0, \\ 2\varepsilon_1 \lambda_1 xy + \varepsilon_1 y^2 + \varepsilon_2 \lambda_2 z^2 + \varepsilon_3 \lambda_3 w^2 = 0, \end{cases} \quad \varepsilon_i = \pm 1, i = 1, 2, 3.$$

$\lambda_1 S - T$ and $\lambda_2 S - T$ are thus simultaneously congruent to the quadrics of equations

$$\begin{cases} -\varepsilon_1 y^2 + \varepsilon_2 (\lambda_1 - \lambda_2) z^2 + \varepsilon_3 (\lambda_1 - \lambda_3) w^2 = 0, \\ -\varepsilon_1 y^2 + 2\varepsilon_1 (\lambda_2 - \lambda_1) xy + \varepsilon_3 (\lambda_2 - \lambda_3) w^2 = 0. \end{cases}$$

Let $\varepsilon = \text{sign} \frac{\lambda_2 - \lambda_3}{\lambda_1 - \lambda_3}$ (recall that $\lambda_1 \neq \lambda_3$ and $\lambda_2 \neq \lambda_3$). By multiplying the above two equations by $-\varepsilon_1 \left| \frac{\lambda_2 - \lambda_3}{\lambda_1 - \lambda_3} \right|$ and $-\varepsilon_1$, respectively, we can rewrite them as

$$\begin{cases} \left| \frac{\lambda_2 - \lambda_3}{\lambda_1 - \lambda_3} \right| y^2 - \varepsilon \varepsilon_1 \varepsilon_2 \frac{(\lambda_1 - \lambda_2)(\lambda_2 - \lambda_3)}{\lambda_1 - \lambda_3} z^2 - \varepsilon \varepsilon_1 \varepsilon_3 (\lambda_2 - \lambda_3) w^2 = 0, \\ \sqrt{\left| \frac{\lambda_2 - \lambda_3}{\lambda_1 - \lambda_3} \right|} y \left(\sqrt{\left| \frac{\lambda_1 - \lambda_3}{\lambda_2 - \lambda_3} \right|} y - 2(\lambda_2 - \lambda_1) \sqrt{\left| \frac{\lambda_1 - \lambda_3}{\lambda_2 - \lambda_3} \right|} x \right) - \varepsilon_1 \varepsilon_3 (\lambda_2 - \lambda_3) w^2 = 0. \end{cases}$$

Now, we apply the following projective transformation :

$$\begin{aligned} \sqrt{\left| \frac{\lambda_1 - \lambda_3}{\lambda_2 - \lambda_3} \right|} y - 2(\lambda_2 - \lambda_1) \sqrt{\left| \frac{\lambda_1 - \lambda_3}{\lambda_2 - \lambda_3} \right|} x &\mapsto x, & \sqrt{\left| \frac{\lambda_2 - \lambda_3}{\lambda_1 - \lambda_3} \right|} y &\mapsto y, \\ \sqrt{\left| \frac{(\lambda_1 - \lambda_2)(\lambda_2 - \lambda_3)}{\lambda_1 - \lambda_3} \right|} z &\mapsto z, & \sqrt{|\lambda_2 - \lambda_3|} w &\mapsto w. \end{aligned}$$

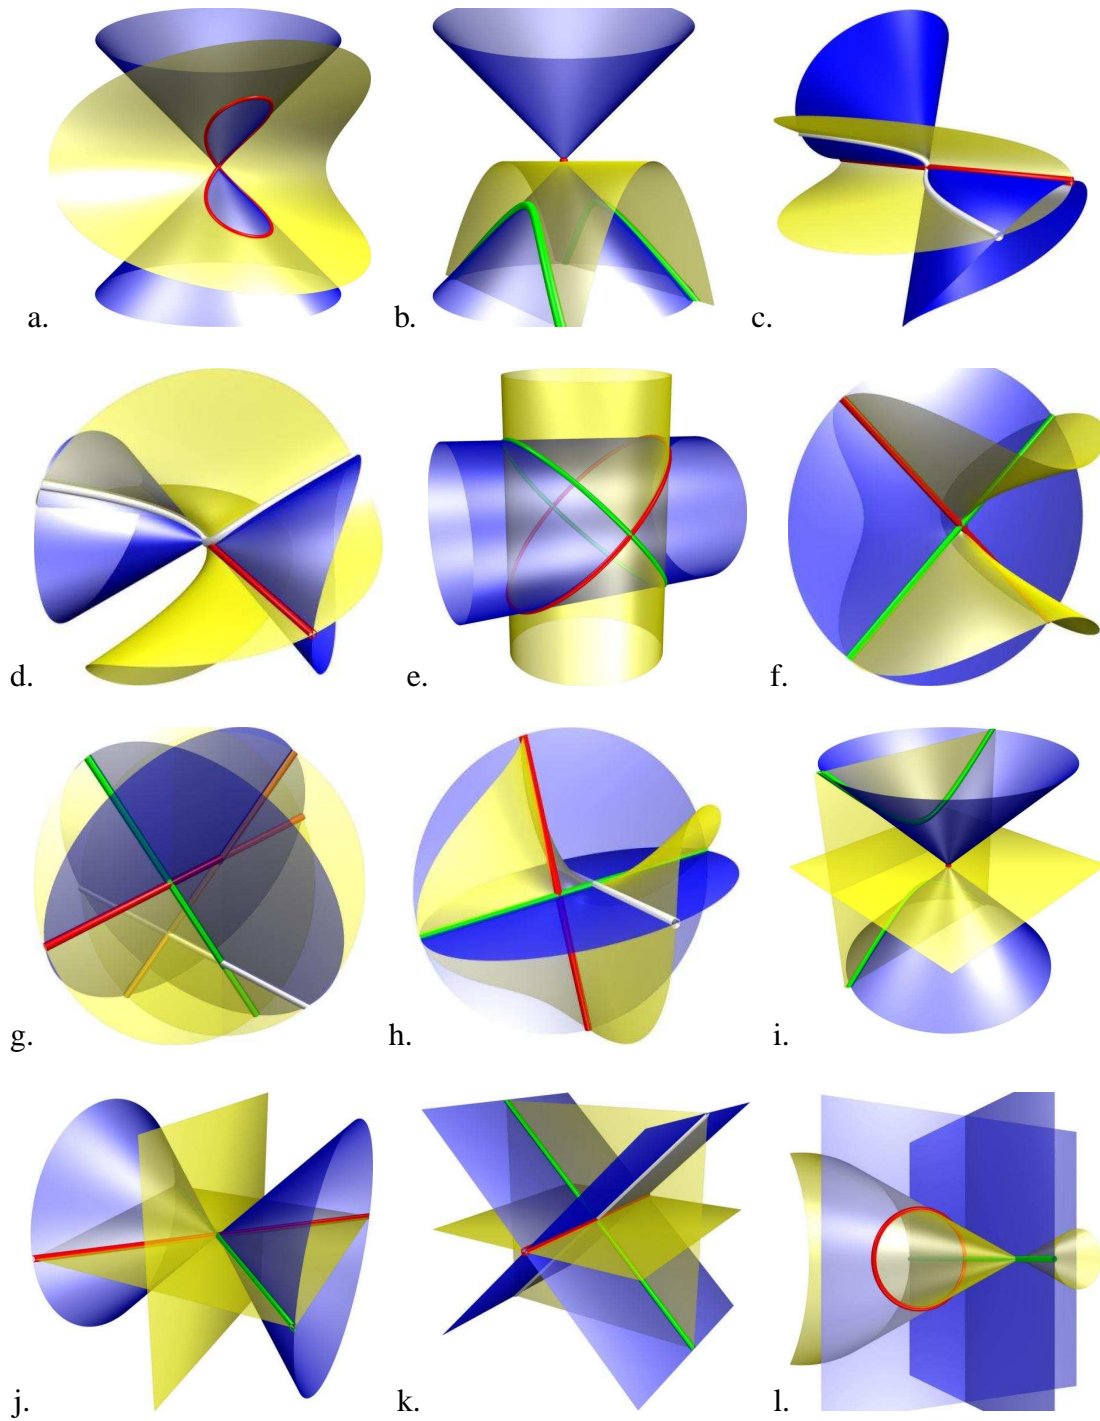


FIG. 15.1 – A gallery of intersections. a. Nodal quartic. b. Nodal quartic with isolated singular point. c. Cubic and secant line. d. Cubic and tangent line. e. Two secant conics. f. Two double lines. g. Four lines forming a skew quadrilateral. h. Two lines and a double line. i. Conic and two lines not crossing on the conic, the two lines being imaginary. j. Four concurrent lines, only two of which are real. k. Two lines and a double line, the three being concurrent. l. Conic and double line.

We obtain that $R(\lambda_1) = \lambda_1 S - T$ and $R(\lambda_2) = \lambda_2 S - T$ are simultaneously congruent, by a real projective transformation P , to the quadrics of equations

$$\begin{cases} P^T R(\lambda_1) P: y^2 + az^2 + bw^2 = 0, \\ P^T R(\lambda_2) P: xy + cw^2 = 0, \end{cases} \quad (15.2)$$

with $a, b, c \in \{-1, 1\}$. One can further assume that $c = 1$ by changing x by $-x$.

From now on we forget about the transformation P and identify $R(\lambda_i)$ with $P^T R(\lambda_i) P$, but it should be kept to mind that things happen in the local frame induced by P .

If a or b is -1 , the cone $R(\lambda_1)$ has inertia $(2, 1)$ and thus is real. Otherwise ($a = b = 1$), the cone $R(\lambda_1)$ is imaginary but for its real apex $\mathbf{p} = (1, 0, 0, 0)$. The other cone $R(\lambda_2)$ is always real and contains the apex \mathbf{p} of $R(\lambda_1)$. We distinguish the three following cases.

- $a = b = 1$. The real part of the nodal quartic is reduced to its node, the apex \mathbf{p} of $R(\lambda_1)$.
- Only one of a and b is 1. Assume for instance that $a = 1, b = -1$ (the other case is obtained by exchanging z and w). By substituting the parameterization of the cone $y^2 + z^2 - w^2 = 0$ (see Table 14.3)

$$\left(s, uv, \frac{u^2 - v^2}{2}, \frac{u^2 + v^2}{2} \right), \quad (u, v, s) \in \mathbb{P}^{*2}(\mathbb{R}),$$

into the other cone $xy + w^2 = 0$, and solving in s , we get the parameterization of the nodal quartic

$$\mathbf{X}(u, v) = ((u^2 + v^2)^2, -4u^2v^2, 2uv(u^2 - v^2), 2uv(u^2 + v^2))^T, \quad (u, v) \in \mathbb{P}^1(\mathbb{R}).$$

The nodal quartic is thus real and its node, corresponding to the parameters $(1, 0)$ and $(0, 1)$, is at \mathbf{p} . The plane tangent to the quadric $Q_{R(\lambda_2)}$ at the quartic's node \mathbf{p} is $y = 0$. In a neighborhood of this node, $x = (u^2 + v^2)^2 > 0$ and $y = -4u^2v^2 \leq 0$ (recall that $\mathbf{X}(u, v)$ is projective, so its coordinates are defined up to a non-zero scalar). We conclude that the two branches lie on the same side of the tangent plane and that the singularity is convex. As can readily be seen, the plane $x = 0$ does not intersect the quartic, so the intersection is affinely finite.

- $a = -1, b = -1$. Parameterizing the nodal quartic as above, we get the parameterization

$$\mathbf{X}(u, v) = (-4u^2v^2, (u^2 + v^2)^2, (u^2 + v^2)(u^2 - v^2), 2uv(u^2 + v^2))^T, \quad (u, v) \in \mathbb{P}^1(\mathbb{R}).$$

It can be checked that the point $\mathbf{p} = (1, 0, 0, 0)$ which is on the intersection is not attained by any real value of the parameter (u, v) (it is only attained with the complex parameters $(1, i)$ and $(i, 1)$). The nodal quartic is thus real with an isolated singular point.

We now argue that we can easily distinguish between these three cases. For this, we first prove the following lemma.

Lemma 15.4. *Given any pencil of quadrics generated by S and T whose characteristic polynomial $\det(\lambda S + \mu T) = 0$ has a double root (λ_1, μ_1) , the sign of $\frac{\det(\lambda S + \mu T)}{(\mu_1 \lambda - \lambda_1 \mu)^2}$ at (λ_1, μ_1) is invariant by a real projective transformation of the pencil and does not depend on the choice of S and T in the pencil.*

Proof. We suppose that $\mathcal{D}(\lambda, \mu) = \det(\lambda S + \mu T)$ has a double root (λ_1, μ_1) . The lemma claims that for any real projective transformation P and any $a_1, \dots, a_4 \in \mathbb{R}$ such that $a_1 a_4 - a_2 a_3 \neq 0$,

$$\mathcal{D}'(\lambda', \mu') = \det(\lambda' P^T (a_1 S + a_2 T) P + \mu' P^T (a_3 S + a_4 T) P)$$

has a double root (λ'_1, μ'_1) such that $\frac{\mathcal{D}(\lambda, \mu)}{(\mu_1 \lambda - \lambda_1 \mu)^2}$ at (λ_1, μ_1) has same sign as $\frac{\mathcal{D}'(\lambda', \mu')}{(\mu'_1 \lambda' - \lambda'_1 \mu')^2}$ at (λ'_1, μ'_1) . We have

$$\mathcal{D}'(\lambda', \mu') = (\det P)^2 \mathcal{D}(a_1 \lambda' + a_3 \mu', a_2 \lambda' + a_4 \mu').$$

Thus $\mathcal{D}'(\lambda', \mu') = 0$ has a double root (λ'_1, μ'_1) defined by

$$\begin{cases} a_1 \lambda'_1 + a_3 \mu'_1 = \lambda_1 \\ a_2 \lambda'_1 + a_4 \mu'_1 = \mu_1 \end{cases} \Leftrightarrow \begin{cases} \lambda'_1 = \frac{a_4 \lambda_1 - a_3 \mu_1}{a_1 a_4 - a_2 a_3} \\ \mu'_1 = \frac{-a_2 \lambda_1 + a_1 \mu_1}{a_1 a_4 - a_2 a_3} \end{cases}.$$

It follows that

$$\frac{\mathcal{D}'(\lambda', \mu')}{(\mu'_1 \lambda' - \lambda'_1 \mu')^2} = (\det P)^2 \frac{\mathcal{D}(a_1 \lambda' + a_3 \mu', a_2 \lambda' + a_4 \mu')}{(\mu_1 (a_1 \lambda' + a_3 \mu') - \lambda_1 (a_2 \lambda' + a_4 \mu'))^2} (a_1 a_4 - a_2 a_3)^2.$$

Hence $\frac{\mathcal{D}'(\lambda', \mu')}{(\mu'_1 \lambda' - \lambda'_1 \mu')^2}$ at (λ'_1, μ'_1) has same sign as $\frac{\mathcal{D}(\lambda, \mu)}{(\mu_1 \lambda - \lambda_1 \mu)^2}$ at (λ_1, μ_1) . \square

Proposition 15.5. *If the characteristic polynomial $\det(\lambda S + \mu T) = 0$ has two simple real roots and one double root (λ_1, μ_1) whose associated matrix $\lambda_1 S + \mu_1 T$ has rank three, then the intersection of S and T in \mathbb{C}^3 is a nodal quartic whose node is the apex of $\lambda_1 S + \mu_1 T$.*

Moreover, if the inertia of $\lambda_1 S + \mu_1 T$ is $(3, 0)$ then the real part of the nodal quartic is reduced to its node. Otherwise the nodal quartic is real; furthermore, if $\frac{\det(\lambda S + \mu T)}{(\mu_1 \lambda - \lambda_1 \mu)^2}$ is negative for $(\lambda, \mu) = (\lambda_1, \mu_1)$, the node is isolated and, otherwise, the singularity is convex.

Proof. The first part of the proposition follows directly from the Segre characteristic (see Section 15.2.2.1 and Table 15.1).

If the inertia of $\lambda_1 S + \mu_1 T$ is $(3, 0)$, then $a = b = 1$ in (15.2) and the result follows as discussed above. Otherwise, considering $S' = P^T(\lambda_1 S - T)P$ and $T' = P^T(\lambda_2 S - T)P$, (15.2) gives that $\det(\lambda S' + \mu T') = -a\lambda(b\lambda + \mu)\mu^2/4$. Evaluating $\frac{\det(\lambda S' + \mu T')}{\mu^2}$ at $(\lambda, \mu) = (1, 0)$, gives by Lemma 15.4 that $-ab$ has same sign as $\frac{\det(\lambda S + \mu T)}{(\mu_1 \lambda - \lambda_1 \mu)^2}$ at (λ_1, μ_1) . The result then follows from the discussion above depending on whether $a = b = -1$ or $ab = -1$. \square

λ_2 and λ_3 are complex conjugate. The reduction to normal pencil form is slightly more involved in this case. Let $\lambda_2 = \alpha + i\beta$, $\lambda_3 = \alpha - i\beta$, $\beta \neq 0$. The Canonical Pair Form Theorem gives that S and T are simultaneously congruent to the quadrics of equations

$$\begin{cases} 2\epsilon xy + 2zw = 0, \\ 2\epsilon \lambda_1 xy + \epsilon y^2 + 2\alpha zw + \beta z^2 - \beta w^2 = 0, \end{cases} \quad \epsilon = \pm 1$$

Through this congruence transformation, $S' = \lambda_1 S - T$ has equation

$$\begin{aligned} 0 &= -\epsilon y^2 + \beta(w^2 - z^2) + 2(\lambda_1 - \alpha)zw, \\ &= -\epsilon y^2 + \beta(w + \xi z) \left(w - \frac{1}{\xi} z \right), \\ &= -\epsilon y^2 + \beta z' w', \end{aligned}$$

where ξ is real and positive. Through the congruence transformation and with the above transformation $(z, w) \mapsto (z', w')$, S has equation

$$\begin{aligned} 0 &= 2\epsilon xy + 2zw \\ &= 2\epsilon xy + \frac{2}{(\xi + \frac{1}{\xi})^2} \left(\frac{1}{\xi} z'^2 - \xi w'^2 + \left(\xi - \frac{1}{\xi} \right) z' w' \right). \end{aligned}$$

Through the above congruence transformations, the quadric of the pencil $T' = \beta S - 2\frac{\xi - \frac{1}{\xi}}{(\xi + \frac{1}{\xi})^2}(\lambda_1 S - T)$ has equation

$$2\epsilon y \left(\beta x + \frac{\xi - \frac{1}{\xi}}{(\xi + \frac{1}{\xi})^2} y \right) + \frac{2\beta}{(\xi + \frac{1}{\xi})^2} \left(\frac{1}{\xi} z'^2 - \xi w'^2 \right) = 0.$$

Finally, by making a shift on x , rescaling on the four axes, and changing the signs of x and z , we get that the two quadrics of the pencil S' and T' are simultaneously congruent to the quadrics of equations

$$\begin{cases} y^2 + zw = 0, \\ xy + z^2 - w^2 = 0. \end{cases} \quad (15.3)$$

As before, we now drop reference to the accumulated congruence transformation and work in the local frame. By substituting the parameterization of the cone $y^2 + zw = 0$ (see Table 14.3)

$$(s, uv, u^2, -v^2), \quad (u, v, s) \in \mathbb{P}^2(\mathbb{R}),$$

into the other quadric $xy + z^2 - w^2 = 0$, and solving in s , we get the parameterization of the nodal quartic

$$\mathbf{X}(u, v) = (v^4 - u^4, u^2v^2, u^3v, -uv^3)^T, \quad (u, v) \in \mathbb{P}^1(\mathbb{R}).$$

The nodal quartic is thus real and its node, corresponding to the parameters $(1, 0)$ and $(0, 1)$, is at $\mathbf{p} = (1, 0, 0, 0)$, the apex of S' . The plane tangent to the quadric $xy + z^2 - w^2 = 0$ at the quartic's node \mathbf{p} is $y = 0$. In a neighborhood of the quartic's node on the branch corresponding to the parameter $(0, 1)$, $x = v^4 - u^4 > 0$ and $y = u^2v^2 \geq 0$. On the other branch corresponding to the parameter $(1, 0)$, $x = v^4 - u^4 < 0$ and $y = u^2v^2 \geq 0$. Hence, the two branches of the quartic are on opposite sides of the tangent plane $y = 0$ in a neighborhood of the node, i.e., the singularity is concave.

Let us briefly show that the intersection is affinely infinite in this case. Consider the plane $\ell_1x + \ell_2y + \ell_3z + \ell_4w = 0$, $(\ell_1, \ell_2, \ell_3, \ell_4) \in \mathbb{P}^3(\mathbb{R})$, which we intersect with the nodal quartic under consideration. This yields a quartic equation E in (u, v) . If $\ell_1 = 0$, E has v in factor, meaning that the point $(1, 0, 0, 0)$ of parameter $(0, 1)$ belongs to the plane. If $\ell_1 \neq 0$, the coefficients of u^4 and v^4 in E have opposite sign, implying by Descartes' Sign Rule that E has at least one real non-trivial solution. The nodal quartic is thus cut by any plane of $\mathbb{P}^3(\mathbb{R})$, implying it is affinely infinite.

To summarize, we have the following result.

Proposition 15.6. *If the characteristic polynomial $\det(\lambda S + \mu T) = 0$ has two simple complex conjugate roots and one double root (λ_1, μ_1) whose associated matrix $\lambda_1 S + \mu_1 T$ has rank three, then the intersection of S and T is a real nodal quartic with a concave singularity at its node, the apex of $\lambda_1 S + \mu_1 T$.*

15.3.2 Two secant conics in $\mathbb{P}^3(\mathbb{C})$, $\sigma_4 = [11(11)]$

The characteristic polynomial has a double root λ_1 and the rank of $R(\lambda_1)$ is 2. λ_1 is necessarily real and there are two Jordan blocks of size 1 associated with it in the canonical form. Let λ_2 and λ_3 be the other (simple) roots, associated with quadrics of rank 3. We have two cases.

λ_2 and λ_3 are real. λ_2 and λ_3 appear in real Jordan blocks of size 1. The normal form of $R(\lambda_1)$ and $R(\lambda_2)$ is :

$$\begin{cases} z^2 + aw^2 = 0, \\ x^2 + by^2 + cw^2 = 0, \end{cases}$$

with $a, b, c \in \{-1, 1\}$.

The two planes of $R(\lambda_1)$ are real if the matrix has inertia $(1, 1)$, i.e. if $a = -1$. The cone $R(\lambda_2)$ is real if its inertia is $(2, 1)$, i.e. if $b = -1$ or $c = -1$. The two conics of the intersection are secant over the reals if the singular line $z = w = 0$ of the pair of planes meets the cone in real points, i.e. if $b = -1$. We have the following cases :

- $a = \pm 1, b = 1, c = 1$: The planes are real or imaginary and the cone is imaginary. The apex of the cone is not on the planes, so intersection is empty.
- $a = 1, b = 1, c = -1$: The planes are imaginary and the cone is real. Their real intersection is the intersection of the singular line $z = w = 0$ of the pair of planes with the cone. The real intersection is thus empty.
- $a = 1, b = -1, c = \pm 1$: The planes are imaginary and the cone is real. The line $z = w = 0$ intersects the cone in two points of coordinates $(1, 1, 0, 0)$ and $(-1, 1, 0, 0)$. The intersection is reduced to these two points.
- $a = -1, b = 1, c = -1$: The planes and the cone are real. The line $z = w = 0$ does not intersect the cone, so intersection consists of two non-secant conics.
- $a = -1, b = -1, c = \pm 1$: The planes and the cone are real. The line $z = w = 0$ intersects the conics. The intersection consists of two conics intersecting in two points \mathbf{p}^\pm of coordinates $(\pm 1, 1, 0, 0)$. All the quadrics of the pencil have the same tangent plane $P^\pm : x \mp y = 0$ at \mathbf{p}^\pm . The two conics of the intersection are on the same side of P^\pm , i.e. the singularity is convex. When $c = 1$, the plane $y = 0$ does not intersect the conics. The same goes for the plane $x = 0$ when $c = -1$. We conclude that the intersection is affinely finite.

Computing the inertia of $R(\lambda_1)$ gives a . Also, in normal form, the characteristic polynomial $\det(\lambda R(\lambda_1) + \mu R(\lambda_2))$ is equal to $b\mu^2\lambda(a\lambda + c\mu)$. Thus, by Lemma 15.4, ab is equal to the sign of $\frac{\det(\lambda S + \mu T)}{(\mu_1\lambda - \lambda_1\mu)^2}$ at (λ_1, μ_1) . Hence we can easily compute a and b . Finally, we need to compute c but only in the case where $a = -1$ and $b = 1$. Then $c = 1$ if the inertia of $R(\lambda_2)$ (or $R(\lambda_3)$) is $(3, 0)$; otherwise $c = -1$ and the inertia of $R(\lambda_2)$ (or $R(\lambda_3)$) is $(2, 1)$.

λ_2 and λ_3 are complex conjugate. There are two complex Jordan blocks of size 2 associated with the two roots. The pencil normal form is obtained as in Section 15.3.1. The end result is :

$$\begin{cases} zw = 0, \\ x^2 + ay^2 + z^2 - w^2 = 0, \end{cases}$$

with $a \in \{-1, 1\}$.

The pair of planes $R(\lambda_1)$ is always real. The intersection consists of the two conics $z = x^2 + ay^2 - w^2 = 0$ and $w = x^2 + ay^2 + z^2 = 0$. We have two cases :

- $a = 1$: One conic is real, the other is imaginary.
- $a = -1$: The two conics are real. They intersect at the points \mathbf{p}^\pm of coordinates $(1, \pm 1, 0, 0)$. All the quadrics of the pencil have the same tangent plane $P^\pm : x \mp y = 0$ at \mathbf{p}^\pm . The two conics of the intersection are on opposite sides of P^\pm , i.e. the singularity is concave.

Let us show that the intersection is affinely infinite. Parameterizing the first conic by $(u^2 + v^2, u^2 - v^2, 0, 2uv)$ and the second by $(u^2 - v^2, u^2 + v^2, 2uv, 0)$, and intersecting with the plane of equation $\ell_1x + \ell_2y + \ell_3z + \ell_4w = 0$, $(\ell_1, \ell_2, \ell_3, \ell_4) \in \mathbb{P}^3(\mathbb{R})$, we obtain two quadratic equations in (u, v) . The product of the coefficients of u^2 and v^2 is $\ell_1^2 - \ell_2^2$ in one case and $\ell_2^2 - \ell_1^2$ in the other case. Therefore, Descartes' Sign Rule implies the existence of a real non-trivial solution to at least one of the quadratic equations if $\ell_1^2 \neq \ell_2^2$. If $\ell_1^2 = \ell_2^2$, then each of the two quadratic equations has either $(u, v) = (0, 1)$ or $(1, 0)$ as real non-trivial solution. We conclude that the intersection is cut by any plane of $\mathbb{P}^3(\mathbb{R})$, i.e. it is affinely infinite.

Note finally that, in normal form, the characteristic polynomial $\det(\lambda R(\lambda_1) + \mu R(\lambda_2))$ is equal to $-\mu^2(\mu^2 + \lambda^2/4)$. Hence a is opposite to the sign of $\frac{\det(\lambda S + \mu T)}{(\mu_1\lambda - \lambda_1\mu)^2}$ at (λ_1, μ_1) (by Lemma 15.4).

15.3.3 Cuspidal quartic in $\mathbb{P}^3(\mathbb{C})$, $\sigma_4 = [13]$

The characteristic polynomial has a triple root λ_1 , which is necessarily real. To it corresponds a real Jordan block of size 3. $R(\lambda_1)$ has rank 3. Let λ_2 be the other root, necessarily real, and $R(\lambda_2)$ the associated cone. The normal form of $R(\lambda_1)$ and $R(\lambda_2)$ is :

$$\begin{cases} w^2 + yz = 0, \\ y^2 + xz = 0. \end{cases}$$

The intersection consists of a cuspidal quartic which can be parameterized (in the local frame of the normal form) by

$$\mathbf{X}(u, v) = (v^4, u^2v^2, -u^4, u^3v), \quad (u, v) \in \mathbb{P}^1(\mathbb{R}).$$

The quartic has a cusp at $\mathbf{p} = (1, 0, 0, 0)$ (the vertex of the first cone), which corresponds to $(u, v) = (0, 1)$. The intersection of $R(\lambda_1)$ with the plane tangent to $R(\lambda_2)$ at \mathbf{p} gives the (double) line tangent to the quartic at \mathbf{p} , i.e. $z = w^2 = 0$.

15.3.4 Two tangent conics in $\mathbb{P}^3(\mathbb{C})$, $\sigma_4 = [1(21)]$

The characteristic polynomial has a triple root λ_1 and the rank of $R(\lambda_1)$ is 2. λ_1 is necessarily real. Attached to λ_1 are two real Jordan blocks, one of size 2, the other of size 1. Let λ_2 be the other simple real root, with $R(\lambda_2)$ of rank 3. The normal forms of $R(\lambda_1)$ and $R(\lambda_2)$ are :

$$\begin{cases} x^2 + aw^2 = 0, \\ xy + z^2 = 0, \end{cases}$$

where $a \in \{-1, 1\}$.

The pair of planes $R(\lambda_1)$ is real when the matrix has inertia $(1, 1)$, i.e. when $a = -1$. The cone $R(\lambda_2)$ is real since its inertia is $(2, 1)$. So we have two cases :

- $a = 1$: The pair of planes is imaginary. Its real part is restricted to the line $x = w = 0$, which intersects the cone in the real double point $(0, 1, 0, 0)$. The intersection is reduced to that point.
- $a = -1$: The planes are real. The intersection consists of two conics intersecting in the double point $(0, 1, 0, 0)$ and sharing a common tangent at that point.

15.3.5 Double conic in $\mathbb{P}^3(\mathbb{C})$, $\sigma_4 = [1(111)]$

The characteristic polynomial has a real triple root λ_1 and the rank of $R(\lambda_1)$ is 1. The Jordan normal form of $S^{-1}T$ contains three blocks of size 1 for λ_1 . Let λ_2 be the other real root, with $R(\lambda_2)$ of rank 3. The normal forms of $R(\lambda_1)$ and $R(\lambda_2)$ are :

$$\begin{cases} w^2 = 0, \\ x^2 + ay^2 + z^2 = 0, \end{cases}$$

where $a \in \{-1, 1\}$.

The cone $R(\lambda_2)$ is real if its inertia is $(2, 1)$, i.e. if $a = -1$. We have two cases :

- $a = -1$: The cone is real. The intersection consists of a double conic lying in the plane $w = 0$.
- $a = 1$: The cone is imaginary. Its real apex does not lie on the plane $w = 0$, so the intersection is empty.

15.3.6 Cubic and tangent line in $\mathbb{P}^3(\mathbb{C})$, $\sigma_4 = [4]$

The characteristic polynomial has a quadruple root λ_1 and the rank of $R(\lambda_1)$ is 3. λ_1 is necessarily real. Associated with it is a unique real Jordan block of size 4. The normal form of $R(\lambda_1)$ and S is :

$$\begin{cases} z^2 + yw = 0, \\ xw + yz = 0. \end{cases}$$

The intersection contains the line $z = w = 0$. The cubic is parameterized by

$$\mathbf{X}(u, v) = (u^3, -u^2v, uv^2, v^3), \quad (u, v) \in \mathbb{P}^1(\mathbb{R}).$$

The cubic intersects the line in the point of coordinate $(1, 0, 0, 0)$, corresponding to the parameter $(1, 0)$. The cubic and the line are tangent at that point.

15.3.7 Conic and two lines crossing on the conic in $\mathbb{P}^3(\mathbb{C})$, $\sigma_4 = [(31)]$

The characteristic polynomial has a quadruple root λ_1 , with $R(\lambda_1)$ of rank 2. λ_1 is necessarily real. To it correspond two real Jordan blocks of size 3 and 1. The normal forms of $R(\lambda_1)$ and S are :

$$\begin{cases} yz = 0, \\ y^2 + xz + aw^2, \end{cases}$$

with $a \in \{-1, 1\}$. $z = 0$ gives two real or imaginary lines. $y = 0$ gives a real conic. The lines cross on the conic at the point $\mathbf{p} = (1, 0, 0, 0)$.

Both the pair of planes and the nonsingular quadric are real. We have two cases :

- $a = 1$: The lines are imaginary. The intersection is reduced to the conic.
- $a = -1$: The lines are real. The intersection consists of a conic and two lines crossing on the conic at \mathbf{p} .

The characteristic polynomial in normal form $\det(\lambda R(\lambda_1) + \mu S) = -a\mu^4/4$ has a quadruple root and thus is always non-negative or non-positive. In this case, it is straightforward to show, similarly as in the proof of Lemma 15.4, that the sign ≥ 0 or ≤ 0 of $\det(\lambda S + \mu T)$ is invariant by real projective transformation and independent of the choice of S and T in the pencil. Hence a is opposite to the sign of $\det(\lambda S + \mu T)$ for any (λ, μ) that is not the quadruple root.

15.3.8 Two lines and a double line in $\mathbb{P}^3(\mathbb{C})$, $\sigma_4 = [(22)]$

The characteristic polynomial has a quadruple root λ_1 , with $R(\lambda_1)$ of rank 2. λ_1 is necessarily real and there are two real Jordan blocks associated with it, both of size 2. The normal forms of $R(\lambda_1)$ and S are :

$$\begin{cases} y^2 + aw^2 = 0, \\ xy + azw = 0, \end{cases}$$

with $a \in \{-1, 1\}$. The intersection consists of the double line $y = w = 0$ and two single lines $y \pm \sqrt{-a}w = x \pm \sqrt{-a}z = 0$ cutting the double line in the points $(\mp\sqrt{-a}, 0, 1, 0)$.

The pair of planes is real if its inertia is $(1, 1)$, i.e. if $a = -1$. We have two cases :

- $a = 1$: The two single lines are imaginary. The intersection is reduced to the double line $y = w = 0$.
- $a = -1$: The intersection consists of the two single lines $y \pm w = x \pm z = 0$ and the double line $y = w = 0$.

Note that the characteristic polynomial $\det(\lambda R(\lambda_1) + \mu S)$ is equal in normal form to $\frac{a^2\mu^4}{16}$. Thus $\mathcal{D}(\lambda, \mu)$ is positive for any (λ, μ) distinct from the quadruple root.

15.3.9 Two double lines in $\mathbb{P}^3(\mathbb{C})$, $\sigma_4 = [(211)]$

The characteristic polynomial has a quadruple root λ_1 , with $R(\lambda_1)$ of rank 1. λ_1 is real and there are three real Jordan blocks associated with it, two having size 1 and the last size 2. The normal forms of $R(\lambda_1)$ and S are :

$$\begin{cases} w^2 = 0, \\ x^2 + ay^2 + zw = 0, \end{cases}$$

with $a \in \{-1, 1\}$. The intersection consists of two double lines $w^2 = x^2 + ay^2 = 0$.

There are two cases :

- $a = 1$: The two double lines are imaginary. The intersection is reduced to their real intersection point, i.e. $(0, 0, 1, 0)$.
- $a = -1$: The two double lines $w^2 = x - y = 0$ and $w^2 = x + y = 0$ are real so the intersection consists of these two lines, meeting at $(0, 0, 1, 0)$.

The characteristic polynomial (in normal form) is equal to $\det(\lambda R(\lambda_1) + \mu S) = -a\lambda^4/4$ thus a is opposite to the sign of $\det(\lambda S + \mu T)$ for any (λ, μ) that is not a root.

15.3.10 Cubic and secant line in $\mathbb{P}^3(\mathbb{C})$, $\sigma_4 = [22]$

The characteristic polynomial has two double roots λ_1 and λ_2 . The associated quadrics both have rank 3. λ_1 and λ_2 are either both real or complex conjugate.

λ_1 and λ_2 are real. There is a real Jordan block of size 2 associated with each root. The normal form of $R(\lambda_1)$ and $R(\lambda_2)$ is :

$$\begin{cases} y^2 + zw = 0, \\ xy + w^2 = 0. \end{cases}$$

The intersection consists of the line $y = w = 0$ and a cubic. The cubic is parameterized by

$$\mathbf{X}(u, v) = (u^3, -uv^2, -v^3, u^2v), \quad (u, v) \in \mathbb{P}^1(\mathbb{R}).$$

The line intersects the cubic in the two points of coordinates $(1, 0, 0, 0)$ and $(0, 0, 1, 0)$, corresponding to the parameters $(1, 0)$ and $(0, 1)$.

λ_1 and λ_2 are complex conjugate. Let $\lambda_1 = \alpha + i\beta, \lambda_2 = \overline{\lambda_1}, \beta \neq 0$. There is complex Jordan block of size 4 associated with the two roots. The normal form of S and $R(\alpha)$ is :

$$\begin{cases} xw + yz = 0, \\ xz - yw + zw = 0. \end{cases}$$

The intersection contains the line $z = w = 0$. The cubic is parameterized by

$$\mathbf{X}(u, v) = (-u^2v, uv^2, u^3 + uv^2, u^2v + v^3), \quad (u, v) \in \mathbb{P}^1(\mathbb{R}).$$

The cubic intersects the line in the points of coordinates $(1, i, 0, 0)$ and $(1, -i, 0, 0)$. Thus, over the reals, the cubic and the line do not intersect.

15.3.11 Conic and two lines not crossing on the conic in $\mathbb{P}^3(\mathbb{C})$, $\sigma_4 = [2(11)]$

The characteristic polynomial has two double roots λ_1 and λ_2 , with associated ranks 3 (a projective cone) and 2 (a pair of planes) respectively. The two roots are necessarily real, otherwise the ranks of the quadrics $R(\lambda_1)$ and $R(\lambda_2)$ would be the same. Associated with λ_1 and λ_2 are respectively a unique real Jordan block of size 2 and two real Jordan blocks of size 1. The pencil normal form is :

$$\begin{cases} y^2 + az^2 + bw^2 = 0, \\ xy = 0, \end{cases}$$

where $a, b \in \{-1, 1\}$. The plane $x = 0$ contains a conic which is real when $a = -1$ or $b = -1$ and imaginary otherwise. The plane $y = 0$ contains two lines which are real if $ab < 0$ and imaginary otherwise. The lines cross at the point $(1, 0, 0, 0)$, the apex of $R(\lambda_1)$, which is not on the conic.

The pair of planes $R(\lambda_2)$ is always real. The cone $R(\lambda_1)$ is real when its inertia is $(2, 1)$, i.e. when $a = -1$ or $b = -1$. We have three cases :

- $a = 1, b = 1$: The lines and the conic are imaginary. The intersection is reduced to the real point of intersection of the two lines, i.e. $(1, 0, 0, 0)$.
- $a = -b$: The lines and the conic are real. The intersection consists of a conic and two intersecting lines, each cutting the conic in a point (at $(0, 0, 1, 1)$ and $(0, 0, -1, 1)$).
- $a = -1, b = -1$: The lines are imaginary, the conic is real. The intersection consists of a conic and the point $(1, 0, 0, 0)$, intersection of the two lines.

To determine in which of the three situations we are, first compute the inertia of $R(\lambda_1)$. If the inertia is $(3, 0)$, this implies that $a = b = 1$. Otherwise, we consider as before the characteristic polynomial in normal form $\det(\lambda R(\lambda_1) + \mu R(\lambda_2)) = -ab\lambda^2\mu^2/4$. By Lemma 15.4, $-ab$ is equal to the sign of $\frac{\det(\lambda S + \mu T)}{(\mu_1\lambda - \lambda_1\mu)^2}$ at (λ_1, μ_1) . If $ab > 0$ then $a = b = -1$, otherwise $a = -b$.

15.3.12 Four lines forming a skew quadrilateral in $\mathbb{P}^3(\mathbb{C})$, $\sigma_4 = [(11)(11)]$

The characteristic polynomial has two double roots λ_1 and λ_2 , with associated quadrics of rank 2. λ_1 and λ_2 can be either both real or both complex conjugate.

λ_1 and λ_2 are real. Each root appears in two real Jordan blocks of size 1. The normal forms of $R(\lambda_1)$ and $R(\lambda_2)$ are :

$$\begin{cases} z^2 + aw^2, \\ x^2 + by^2, \end{cases}$$

where $a, b \in \{-1, 1\}$.

The first pair of planes is imaginary if $a = 1$. The second pair of planes is imaginary if $b = 1$. There are three cases :

- $a = 1, b = 1$: The four lines are imaginary and the intersection is empty.
- $a = -b$: One pair of planes is real, the other is imaginary. If $a = 1$, the intersection consists of the points of intersection of the line $z = w = 0$ with the pair of planes $x^2 - y^2 = 0$, i.e. the points $(1, 1, 0, 0)$ and $(-1, 1, 0, 0)$. Similarly, if $b = 1$ the intersection is reduced to the two points $(0, 0, 1, 1)$ and $(0, 0, -1, 1)$.
- $a = -1, b = -1$: The four lines are real. The intersection consists of four lines forming a skew quadrilateral.

The values of a and b follow from the inertia of $R(\lambda_1)$ and $R(\lambda_2)$. Note also that b directly follows from a because, the characteristic polynomial (in normal form) $\det(\lambda R(\lambda_1) + \mu R(\lambda_2)) = ab\lambda^2\mu^2$ and it is straightforward to show that ab is equal to the sign of $\det(\lambda S + \mu T)$ for any (λ, μ) that is not a root.

λ_1 and λ_2 are complex conjugate. Let $\lambda_1 = a + ib, \lambda_2 = \overline{\lambda_1}, b \neq 0$. The roots appear in two complex Jordan blocks of size 2. The normal forms of S and $aS - T$ are :

$$\begin{cases} xy + zw = 0, \\ x^2 - y^2 + z^2 - w^2 = 0. \end{cases}$$

The intersection consists of two real lines of equations $x \pm w = y \mp z = 0$ and two imaginary lines of equations $x \pm iz = y \mp iw = 0$.

15.4 Classification of singular pencils of $\mathbb{P}^3(\mathbb{R})$ over the reals

We now examine the singular pencils of $\mathbb{P}^3(\mathbb{R})$, i.e. those whose characteristic polynomial vanishes identically.

There are two cases according to whether two arbitrary quadrics of the pencil have a singular point in common or not.

15.4.1 Q_S and Q_T have no singular point in common, $\sigma_4 = [1\{3\}]$

We first prove the following lemma.

Lemma 15.7. *If $\det R(\lambda, \mu) \equiv 0$ and Q_S and Q_T have no singular point in common, then every quadric of the pencil has a singular point such that all the other quadrics of the pencil contains this point and share a common tangent plane at this point.*

Proof. Let Q_R be any quadric of the pencil. First note that R has rank at most 3, otherwise the characteristic polynomial would not identically vanish.

If R has rank 1, it is a double plane in $\mathbb{P}^3(\mathbb{C})$ containing only singular points. Since there is no quadric of inertia $(4, 0)$ in the pencil, the intersection of the double plane with every other quadric of the pencil is not empty in $\mathbb{P}^3(\mathbb{R})$ (by Theorem 14.5). Hence Q_R contains a singular point that belongs to all the quadrics of the pencil.

If R has rank 2, it is a pair of planes in $\mathbb{P}^3(\mathbb{C})$ with a real singular line. By the Segre classification (see Table 15.2) we know that the intersection in $\mathbb{P}^3(\mathbb{C})$ contains a conic and a double line. Furthermore, the line is necessarily real because otherwise its conjugate would also be in the intersection. This line lies in one of the two planes of Q_R and thus cuts any other line in that plane and in particular the singular line of the pair of planes. Hence Q_R contains a singular point that belongs to all the quadrics of the pencil.

If R has rank 3, we apply a congruence transformation so that Q_R has the diagonal form $ax^2 + by^2 + cz^2 = 0$, with $abc \neq 0$. We may also change the generators of the pencil, replacing S by R . After these transformations, the determinant $\mathcal{D}(\lambda, \mu)$ becomes the sum of $\delta abc \lambda^3 \mu$ and of other terms of degree at least 2 in μ , where δ is the coefficient of w^2 in the equation of Q_T . The hypothesis that $\mathcal{D}(\lambda, \mu) \equiv 0$ thus implies that $\delta = 0$. Hence the singular point $(0, 0, 0, 1)$ of Q_R belongs to Q_T and thus to all the quadrics of the pencil.

Thus, in all cases, every quadric of the pencil has a singular point that belongs to all the quadrics of the pencil. Any such point \mathbf{p} lies on the intersection of Q_S and Q_T and is a singular point of the intersection : since \mathbf{p} is a singular point of a quadric of the pencil, the rank of the Jacobian matrix (15.1) is less than two. We conclude on the common tangent plane by applying Proposition 15.1 under the assumption that \mathbf{p} is not a singular point of both Q_S and Q_T . \square

By Lemma 15.7, there exist a singular point \mathbf{s} of Q_S and a singular point \mathbf{t} of Q_T that belong to all the quadrics $Q_{\lambda S + \mu T}$ of the pencil. Quadrics Q_S , Q_T , and Q_{S+T} have rank at most 3 since the characteristic polynomial identically vanishes, and they are not of inertia $(3, 0)$ (see Table 14.1) since they contain \mathbf{s} and \mathbf{t} that are distinct by assumption. Hence Q_S , Q_T , and Q_{S+T} are cones or pairs of (possibly complex) planes. Thus, since \mathbf{s} and \mathbf{t} are singular points of Q_S and Q_T , respectively, the line \mathbf{st} is entirely contained in Q_S and Q_T , and thus is also contained in Q_{S+T} . Moreover, \mathbf{s} and \mathbf{t} are not singular points of Q_{S+T} because otherwise all the quadrics of the pencil would be singular at these points, contradicting the hypothesis. It now follows from the fact that Q_{S+T} is a cone or a pair of planes that its tangent planes at \mathbf{s} and \mathbf{t} coincide. Therefore, by Lemma 15.7, the tangent plane of Q_S at \mathbf{t} is the same as the tangent plane of Q_T at \mathbf{s} .

Now we change of frames in such a way that \mathbf{s} and \mathbf{t} have coordinates $(0,0,0,1)$ and $(0,0,1,0)$ and that the common tangent plane has equation $x = 0$. Then the equations of Q_S and Q_T become $xz + q_1(x,y) = 0$ and $xw + q_2(x,y) = 0$, where q_1 and q_2 are binary quadratic forms. The two equations can thus be expressed in the form $ay^2 + x(by + cx + z) = 0$ and $a'y^2 + x(b'y + c'x + w) = 0$. By a new change of frame, we get equations of the form $ay^2 + xz = 0$ and $a'y^2 + xw = 0$. Replacing the second quadric by a linear combination of the two and applying the change of coordinates $a'z - aw \rightarrow w$ and a scaling on y , gives as normal form for the pencil :

$$\begin{cases} xw = 0, \\ xz + ay^2 = 0, \end{cases}$$

with $a \in \{-1, 1\}$. Furthermore, we can set $a = 1$ by changing z in $-z$.

Therefore, the intersection consists of the double line $x = y^2 = 0$ and the conic $w = xz - y^2 = 0$. The line and the conic meet at $(0, 0, 1, 0)$ in the local frame of the normal form.

15.4.2 Q_S and Q_T have a singular point in common

Let \mathbf{p} be the common singular point. We proceed as already outlined in Section 15.2.2. After a rational change of frame, we may suppose that \mathbf{p} has coordinates $(0,0,0,1)$. In the new frame, the equations of the quadrics are homogeneous polynomials of degree 2 in three variables. To classify the different types of intersection, we may identify the quadrics with their upper left 3×3 matrices and look at the multiple roots of the cubic characteristic polynomial, which we refer to as the *restricted characteristic polynomial*, and the ranks of the associated matrices. We thus apply the Canonical Pair Theorem to pairs of conics.

The case $[(111)]$ is trivial and left aside : in that situation, the cubic characteristic polynomial has a (real) triple root, the associated quadric has rank 0 and all the other quadrics of the pencil are cones. The intersection consists of any cone of the pencil, that is any quadric of the pencil distinct from $\mathbb{P}^3(\mathbb{R})$.

15.4.2.1 Four concurrent lines in $\mathbb{P}^3(\mathbb{C})$, $\sigma_3 = [111]$

The restricted characteristic polynomial has three simple roots. At least one is real : call it λ_1 . Let λ_2 be another root. To these roots correspond quadrics of rank 2.

If λ_2 is real, then the three roots are real. The normal form of $R(\lambda_1)$ and $R(\lambda_2)$ is :

$$\begin{cases} ay^2 + z^2 = 0, \\ bx^2 + z^2 = 0, \end{cases}$$

with $a, b \in \{-1, 1\}$. Note that the equation of the third pair of planes of the pencil is obtained by subtracting the two equations, giving $ay^2 - bx^2 = 0$. We have two cases :

- $a = b = -1$: The intersection consists of four concurrent lines $y - \varepsilon_1 z = x - \varepsilon_2 z = 0$, with $\varepsilon_1, \varepsilon_2 \in \{-1, 1\}$, meeting at \mathbf{p} .
- $a = b = 1$ or $a = -b$: When $a = b = 1$, both $R(\lambda_1)$ and $R(\lambda_2)$ represent imaginary pairs of planes. When $a = -b$, the third pair of planes is imaginary, as well as one of the first two. In both cases, the intersection is reduced to the point \mathbf{p} .

Both a and b are equal to -1 if and only if $R(\lambda_1)$ and $R(\lambda_2)$ have inertia $(1, 1)$.

If $\lambda_2 = \alpha + i\beta$ is complex, $\beta \neq 0$, we obtain the following normal form (proceeding as in Section 15.3.1) :

$$\begin{cases} x^2 + y^2 - z^2 = 0, \\ yz = 0. \end{cases}$$

The intersection consists of two real lines $y = x^2 - z^2 = 0$, intersecting at \mathbf{p} , and two complex lines $z = x^2 + y^2 = 0$.

15.4.2.2 Two lines and a double line in $\mathbb{P}^3(\mathbb{C})$, $\sigma_3 = [12]$

The restricted characteristic polynomial has a double root λ_1 , which is real. The Jordan normal form of $S^{-1}T$ contains one real Jordan block of size 2. Let λ_2 be the other root, also real. The normal forms of $R(\lambda_1)$ and $R(\lambda_2)$ are :

$$\begin{cases} y^2 + az^2 = 0, \\ xy = 0, \end{cases}$$

where $a \in \{-1, 1\}$. There are two cases :

- $a = -1$: The intersection consists of the double line $y = z^2 = 0$ and the two single lines $x = y - z = 0$ and $x = y + z = 0$. The three lines are concurrent at \mathbf{p} .
- $a = 1$: The two single lines are imaginary. Their common point is on the double line, so the intersection consists of this double line $y = z^2 = 0$.

Note that the value of a follows from the inertia of $R(\lambda_1)$.

15.4.2.3 Two double lines in $\mathbb{P}^3(\mathbb{C})$, $\sigma_3 = [1(11)]$

The restricted characteristic polynomial has a double root λ_1 , which is real. The canonical pair form has two real Jordan blocks of size 1 associated with λ_1 . Let λ_2 be the other root, also real. The normal forms of $R(\lambda_1)$ and $R(\lambda_2)$ are :

$$\begin{cases} z^2 = 0, \\ x^2 + ay^2 = 0, \end{cases}$$

where $a \in \{-1, 1\}$. The pair of planes $R(\lambda_2)$ is real when its inertia is $(1, 1)$, i.e. when $a = -1$. We have two cases :

- $a = 1$: The intersection is reduced to the point \mathbf{p} .
- $a = -1$: The intersection consists of the two double lines $x - y = z^2 = 0$ and $x + y = z^2 = 0$, meeting at \mathbf{p} .

Note that the value of a follows from the inertia of $R(\lambda_2)$.

15.4.2.4 Line and triple line in $\mathbb{P}^3(\mathbb{C})$, $\sigma_3 = [3]$

The restricted characteristic polynomial has a triple root λ_1 , which is real. The Jordan normal form of $S^{-1}T$ contains one real Jordan block of size 3. The normal forms of S and $R(\lambda_1)$ are :

$$\begin{cases} xz + y^2 = 0, \\ yz = 0. \end{cases}$$

The intersection consists of the triple line $z = y^3 = 0$ and the simple line $x = y = 0$. The two lines cut at \mathbf{p} , the singular point of all the quadrics of the pencil.

15.4.2.5 Quadruple line in $\mathbb{P}^3(\mathbb{C})$, $\sigma_3 = [(21)]$

The restricted characteristic polynomial has a real triple root λ_1 . The canonical pair form has two real Jordan blocks of size 2 and 1. The normal form of $R(\lambda_1)$ and S is :

$$\begin{cases} y^2 = 0, \\ z^2 + xy = 0. \end{cases}$$

The intersection consists of the quadruple line $y^2 = z^2 = 0$.

15.4.2.6 $\sigma_3 = [\{3\}]$ and remaining cases

In this case, the restricted characteristic polynomial identically vanishes. One can easily prove that if the two conics S and T have no singular point in common, the pencil can be put in the normal form $\lambda xy + \mu xz$. The intersection consists of the plane $x = 0$ and the line $y = z = 0$, which meets the plane at \mathbf{p} .

If the two conics have a singular point in common (call it \mathbf{q}), we can go from 3×3 matrices to 2×2 matrices pretty much as above by sending \mathbf{q} to $(0, 0, 1, 0)$. Consider the new characteristic polynomial, a quadratic polynomial. The cases are :

- Two simple real roots : The pencil can be put in the normal form $\lambda x^2 + \mu y^2$. The intersection consists of the quadruple line $x^2 = y^2 = 0$ which goes through \mathbf{p} and \mathbf{q} .
- Two simple complex roots : A normal form for the pencil is $\lambda xy + \mu(x^2 - y^2)$, giving the quadruple line $x^2 = y^2 = 0$ for the intersection.
- A double real root, with a real Jordan block of size 2 : The normal form is $\lambda xy + \mu y^2$. The intersection consists of the plane $y = 0$.
- Vanishing quadratic equation : The intersection consists of a double plane.

15.5 Classifying degenerate intersections

Our near-optimal algorithm for parameterizing intersections of quadrics works in two stages : first it determines the real type of the intersection and, second, it computes a parameterization of this intersection. The purpose of this section is to detail the first stage, called the *type-detection* phase. The second stage, which consists of case-by-case algorithms for computing (near-)optimal parameterizations of the real part of the intersection, will be presented in Part III.

The splitting in two stages reflects a key design choice of our parameterization algorithm, which sums up as : the sooner you know what is the type of the intersection, the less prone you are of making errors in later stages. Information obtained in the type-detection phase is used to drive the algorithm and make it efficiently compute precisely and uniquely what is needed.

Note however that presenting the type detection distinctly from the parameterization is quite an oversimplification. In the actual implementation, there is no clear cut separation between the two stages, which are largely intertwined. Sometimes detecting the type of the intersection is doing a very small step towards parameterization. And sometimes almost everything takes place in the type-detection phase.

We now turn to a high-level description of the detection phase, which relies heavily on the results of Sections 15.3 and 15.4 on real pencils of quadrics. We start by presenting some global tools, and then outline the type-detection algorithm for each of the possible root patterns (vanishing characteristic polynomial, one double root, one triple root, one quadruple root, two double roots).

15.5.1 Preliminaries

In what follows, we assume that the two input quadrics Q_S and Q_T have rational coefficients. We now briefly describe the basic operations needed for detecting the real type of the intersection $Q_S \cap Q_T$. They essentially fall in two categories : linear algebra routines and elementary algebraic manipulations. Most computations involve rational numbers. We give special attention to the limited number of situations where this is not the case.

Let $R(\lambda, \mu) = \lambda S + \mu T$ be the pencil generated by S and T . Computing the coefficients of the characteristic polynomial $\mathcal{D}(\lambda, \mu) = \det R(\lambda, \mu)$ involves nothing but computing determinants of rational matrices, so there is nothing difficult here. Next, we need to compute the inertia and rank of a matrix $R_0 = R(\lambda_0, \mu_0)$ of the pencil, where (λ_0, μ_0) is a root of \mathcal{D} or a real rational projective number.

Assume first that R_0 has real and rational coefficients. Let $p(\omega) = \det(R_0 - \omega I)$, where I is the identity matrix. Since R_0 is symmetric, all its eigenvalues are real. We can thus compute the number e^+ of positive eigenvalues and the number e^- of negative eigenvalues of R_0 by applying Descartes' Sign Rule to $p(\omega)$ and $p(-\omega)$ respectively. Then the inertia I_0 of R_0 is the pair (e^+, e^-) and its rank r_0 is $e^+ + e^-$.

When the coefficients of R_0 are not rational or real, the worst situation that we have to deal with is when \mathcal{D} has two real or complex conjugate double roots. In both cases, the coefficients of R_0 involve one square root. When the roots are real conjugate, we could use Descartes' Sign Rule again, except that we have to evaluate the signs of coefficients that now involve a square root. In these cases we however propose a more efficient approach based on the rank of R_0 (see Algorithm 6) which only deals with rational numbers. We show below how the rank of R_0 can be computed using only standard linear algebra operations on rational numbers.

Let $(\lambda_0, \mu_0) = (a, b \pm \sqrt{c})$, where $(a, b, c) \in \mathbb{P}^2(\mathbb{Q})$, c is not a square and c is either > 0 (real conjugate roots) or < 0 (complex conjugate roots). Form the 8×8 rational matrix

$$M_0 = \begin{pmatrix} aS + bT & cT \\ T & aS + bT \end{pmatrix}.$$

If the vector $\mathbf{k}_1 + \sqrt{c}\mathbf{k}_2$ is in the kernel of R_0 , then the column vector $(\mathbf{k}_1^T, \mathbf{k}_2^T)^T$ is in the kernel of M_0 . But the vector $(c\mathbf{k}_2^T, \mathbf{k}_1^T)$ also is in the kernel of M_0 . It is not too difficult to realize that the kernel of M_0 has twice the number of elements of the kernel of R_0 , i.e. $\dim \ker R_0 = \frac{1}{2}(\dim \ker M_0 - 1)$. We can thus conclude that

$$r_0 = 3 - \dim \ker R_0 = \frac{1}{2}(7 - \dim \ker M_0).$$

Computing the singular space of a quadric with rational coefficients is another operation we need. This only amounts to computing the kernel of the associated matrix. Also, intersecting the singular spaces of two quadrics is like computing the intersection between two linear spaces : it is a standard linear algebra operation.

Algorithm 1 Main loop for degenerate intersection classification.

Require: a pencil of quadrics $R(\lambda, \mu) = \lambda S + \mu T$

- 1: compute $\mathcal{D}(\lambda, \mu) := \det R(\lambda, \mu)$
- 2: **if** $\mathcal{D} \equiv 0$ **then** // vanishing characteristic polynomial
- 3: execute Algorithm 2
- 4: **else**
- 5: compute $\mathcal{H} := \gcd(\partial \mathcal{D} / \partial \lambda, \partial \mathcal{D} / \partial \mu)$ and let $d := \text{degree}(\mathcal{H}, \{\lambda, \mu\})$
- 6: **if** $d = 0$ **then** // no multiple root
- 7: output: smooth quartic (\mathbb{C}) – see Part I
- 8: **else if** $d = 1$ **then** // double real root
- 9: execute Algorithm 3
- 10: **else if** $d = 2$ **then**
- 11: **if** $\text{discriminant}(\mathcal{H}) = 0$ **then** // triple real root
- 12: execute Algorithm 4
- 13: **else** // two double roots
- 14: execute Algorithm 6
- 15: **end if**
- 16: **else** // $d = 3$: quadruple real root
- 17: execute Algorithm 5
- 18: **end if**
- 19: **end if**

In terms of algebraic computations, we need to be able to compute the gcd of polynomials of degree at most 3 and to isolate the roots of a cubic polynomial (in the four concurrent lines case). This last task can be done using Uspensky's algorithm as in Part I or using special methods for comparing the roots of low-degree polynomials (see [ET04]).

The top level type-detection loop is given in Algorithm 1. It does not necessitate much comment save for the fact that when the gcd \mathcal{H} of the two partial derivatives of \mathcal{D} has degree 2, then either its discriminant vanishes, in which case \mathcal{D} has a triple root, or it does not vanish and \mathcal{D} has two double roots.

15.5.2 \mathcal{D} vanishes identically

The type-detection algorithm when \mathcal{D} is identically zero, outlined in Algorithm 2, is little more than a reprise of the information contained in Section 15.4. The general idea is : determine if two arbitrary quadrics of the pencil have a singular point \mathbf{p} in common. If they do, send \mathbf{p} to infinity and work on the pencil of conics living in the plane $w = 0$. To actually compute the restricted pencil $R_3(\lambda, \mu)$, build the matrix of a real projective transformation P obtained by putting \mathbf{p} as the last column and completing P so that its columns form a basis of $\mathbb{P}^3(\mathbb{R})$. $R_3(\lambda, \mu)$ is then the principal submatrix of the matrix $P^T(\lambda S + \mu T)P$.

Two comments are in order. First, a multiple root of a cubic form – the characteristic polynomial of the pencil of conics – is necessarily real (otherwise its complex conjugate would also be a multiple root) and rational (otherwise its real conjugate would also be a multiple root).

So the only place where we might end up working with non-rational numbers is in the four concurrent lines case. Indeed, in this situation the restricted characteristic polynomial \mathcal{D}_3 is a cubic form with three generically non-rational simple roots. Computing the sign of the discriminant of \mathcal{D}_3 can help us distinguish between the cases when only two lines are real and when the four lines are either all real or all imaginary. But this is not enough to give a complete characterization over the reals. We have thus decided to apply Finsler's Theorem (Theorem 14.5), as in the smooth quartic case, after isolating the roots of the cubic. If the restricted pencil contains a conic of inertia $(3, 0)$, then the intersection of conics is empty in the plane $w = 0$ and the intersection of the two initial quadrics is reduced to \mathbf{p} .

Algorithm 2 Classifying the intersection when the characteristic polynomial vanishes.

Require: $R(\lambda, \mu)$ (from Algorithm 1)

```

1: let  $\gamma := \text{singular}(Q_S) \cap \text{singular}(Q_T)$ ,  $\kappa := \dim \gamma$ 
2: if  $\kappa = -1$  then // conic and double line ( $\mathbb{C}$ )
3:   output: conic and double line
4: else //  $\kappa \geq 0$  : at least one common singular point  $\gamma(1)$ 
5:   send  $\gamma(1)$  to the point  $[0 \ 0 \ 0 \ 1]$ 
6:   compute the restricted pencil  $R_3(\lambda, \mu)$  of upper left  $3 \times 3$  matrices and  $\mathcal{D}_3(\lambda, \mu) := \det R_3(\lambda, \mu)$ 
7:   if  $\mathcal{D}_3 \equiv 0$  then // vanishing restricted characteristic polynomial
8:     either  $\sigma_3 = [\{3\}]$  or repeat restriction
9:   else //  $\mathcal{D}_3 \neq 0$ 
10:    compute  $\mathcal{H}_3 := \gcd(\partial \mathcal{D}_3 / \partial \lambda, \partial \mathcal{D}_3 / \partial \mu)$  and let  $d_3 := \text{degree}(\mathcal{H}_3, \{\lambda, \mu\})$ 
11:    if  $d_3 = 0$  then // no multiple root : four concurrent lines ( $\mathbb{C}$ )
12:      if  $\mathcal{D}_3$  has only one real root then
13:        output: two concurrent lines
14:      else if  $R_3(\lambda, \mu)$  contains a conic of inertia  $(3, 0)$  then
15:        output: point
16:      else
17:        output: four concurrent lines
18:      end if
19:    else // one multiple root
20:      let  $(\lambda_0, \mu_0)$  be the multiple root of  $\mathcal{D}_3$ ,  $I_0$  and  $r_0$  the inertia and rank (resp.) of  $R(\lambda_0, \mu_0)$ 
21:      if  $d_3 = 1$  then //  $\mathcal{D}_3$  has one double root
22:        if  $r_0 = 2$  then // two concurrent lines and a double line ( $\mathbb{C}$ )
23:          if  $I_0 = (1, 1)$  then // pair of planes  $R_3(\lambda_0, \mu_0)$  is real
24:            output: two concurrent lines and a double line
25:          else // pair of planes  $R_3(\lambda_0, \mu_0)$  is imaginary
26:            output: double line
27:          end if
28:        else //  $r_0 = 1$  : two double lines ( $\mathbb{C}$ )
29:          let  $(\lambda_1, \mu_1)$  be the other root of  $\mathcal{D}_3$ ,  $I_1$  the inertia of  $R(\lambda_1, \mu_1)$ 
30:          if  $I_1 = (1, 1)$  then // pair of planes  $R_3(\lambda_1, \mu_1)$  is real
31:            output: two double lines
32:          else // pair of planes  $R_3(\lambda_1, \mu_1)$  is imaginary
33:            output: point
34:          end if
35:        end if
36:      else //  $d_3 = 2$  :  $\mathcal{D}_3$  has one triple root
37:        if  $r_0 = 2$  then // line and triple line ( $\mathbb{C}$ )
38:          output: line and triple line
39:        else if  $r_0 = 1$  then // quadruple line ( $\mathbb{C}$ )
40:          output: quadruple line
41:        else //  $r_0 = 0$  : projective cone ( $\mathbb{C}$ )
42:          if  $S$  or  $T$  (in restricted form) has inertia  $(2, 1)$  then
43:            output: cone
44:          else
45:            output: point
46:          end if
47:        end if
48:      end if
49:    end if
50:  end if
51: end if

```

15.5.3 \mathcal{D} has a single multiple root

The type-detection algorithms when \mathcal{D} has a unique multiple root are given in Algorithms 3 (double root), 4 (triple root) and 5 (quadruple root).

First note that when \mathcal{D} has a single multiple root, it is necessarily real and rational, for the same reasons as above. So the singular quadrics that we deal with all have rational coefficients and their singular set can be parameterized rationally.

In the double real root case, we use the result of Lemma 15.4 and classify the intersections according (among others) the sign

$$s := \text{sign } \mathcal{E}(\lambda_0, \mu_0), \quad \text{with } \mathcal{E}(\lambda, \mu) := \mathcal{D}(\lambda, \mu) / (\mu_0 \lambda - \lambda_0 \mu)^2, (\lambda_0, \mu_0) \text{ double root of } \mathcal{D}.$$

The other slight difficulty occurs in the two secant conics case, when the pair of planes $R(\lambda_0, \mu_0)$ is real and $s = -1$. To separate the two subcases (two non-secant conics or empty set), we can compute the inertia of $R(\lambda_1, \mu_1)$, where (λ_1, μ_1) is a simple root of \mathcal{D} , when this root is rational. But in the general case, we use again Finsler's theorem, looking for a quadric of inertia $(4, 0)$ between and outside the two simple roots of \mathcal{D} . If such a quadric is found, the intersection is empty.

The triple real root case does not necessitate further comment except for noticing that the additional simple root (λ_1, μ_1) of \mathcal{D} is necessarily real and rational, so the associated quadric $R(\lambda_1, \mu_1)$ has rational coefficients.

The type detection in the quadruple real root case is pretty straightforward. The only subtlety is that the case of a quadruple real root (λ_0, μ_0) with associated quadric of rank 2 corresponds to two different Jordan decompositions, with Segre symbols $[(22)]$ and $[(31)]$, as already mentioned in Section 15.2.2. To distinguish between the two, we simply note that in the first situation, the singular line of the pair of planes $R(\lambda_0, \mu_0)$ is entirely contained in all the quadrics of the pencil.

15.5.4 \mathcal{D} has two double roots

The type-detection algorithm when \mathcal{D} has two double roots is given in Algorithm 6. We distinguish between two situations in the algorithm : the roots are either real and rational, or they are not. In the first case, computing the inertia of singular quadrics is easy since computations take place over the rationals. Also, note that if one of the singular quadrics has rank 2 and the other has rank 3, then the associated roots of the characteristic polynomial are necessarily rational.

So assume that the roots are not real or not rational. The ranks of the non-rational singular quadrics are necessarily the same, so we need only compute one of them, in the way indicated above. When this rank is 2 and the roots are real conjugate, we distinguish between the remaining subcases (four real lines forming a skew quadrilateral or the empty set) by testing whether any quadric of the pencil between the two roots has inertia $(4, 0)$ or not.

15.6 Examples

We now give several examples for which the type of the real part of the intersection is determined using the type-detection algorithms of the previous section.

15.6.1 Example 1

Consider the following pair of quadrics :

$$\begin{cases} Q_S : -x^2 - 4xy + 4xz - 6y^2 + 2yz - 4yw + 2zw - 2w^2 = 0, \\ Q_T : -x^2 - 6xy + 4xz - 2xw - 6y^2 - 8yw - 6w^2 = 0. \end{cases}$$

The characteristic polynomial is

$$\mathcal{D}(\lambda, \mu) = \det(\lambda S + \mu T) = -16(2\lambda^4 - 10\lambda^3\mu - 19\lambda^2\mu^2 - 16\lambda\mu^3 - 5\mu^4).$$

Algorithm 3 Classifying the intersection : the double real root case, $d = 1$.

Require: $R(\lambda, \mu), \mathcal{D}$ (from Algorithm 1)

Require: double real root (λ_0, μ_0) , inertia I_0 and rank r_0 of $R(\lambda_0, \mu_0)$

```

1: let  $s := \text{sign } \mathcal{E}(\lambda_0, \mu_0)$  and  $\delta := \text{sign}(\text{discriminant}(\mathcal{E}))$ , where  $\mathcal{E}(\lambda, \mu) := \mathcal{D}(\lambda, \mu) / (\mu_0\lambda - \lambda_0\mu)^2$ 
2: if  $r_0 = 3$  then // nodal quartic ( $\mathbb{C}$ )
3:   if  $\delta = -1$  then // other roots are complex
4:     output: nodal quartic, concave singularity
5:   else // other roots are real
6:     if  $s = +1$  then
7:       output: nodal quartic, convex singularity
8:     else if  $I_0 = (2, 1)$  then // cone  $R(\lambda_0, \mu_0)$  is real
9:       output: nodal quartic with isolated singular point
10:    else // cone  $R(\lambda_0, \mu_0)$  is imaginary
11:      output: point
12:    end if
13:  end if
14: else //  $r_0 = 2$  : two secant conics ( $\mathbb{C}$ )
15:   if  $\delta = -1$  then // other roots are complex
16:     if  $s = +1$  then
17:       output: two secant conics, concave singularities
18:     else
19:       output: one conic
20:     end if
21:   else // other roots are real
22:     if  $I_0 = (1, 1)$  then // pair of planes  $R(\lambda_0, \mu_0)$  is real
23:       if  $s = +1$  then
24:         output: two secant conics, convex singularities
25:       else
26:         if  $R(\lambda, \mu)$  contains a quadric of inertia  $(4, 0)$  then
27:           output:  $\emptyset$ 
28:         else
29:           output: two non-secant conics
30:         end if
31:       end if
32:     else // pair of planes  $R(\lambda_0, \mu_0)$  is imaginary
33:       if  $s = +1$  then
34:         output:  $\emptyset$ 
35:       else
36:         output: two points
37:       end if
38:     end if
39:   end if
40: end if

```

Algorithm 4 Classifying the intersection : the triple real root case, $d = 2$ and discriminant $(\mathcal{H}) = 0$.

Require: $R(\lambda, \mu), \mathcal{D}$ (from Algorithm 1)

Require: triple real root (λ_0, μ_0) , inertia I_0 and rank r_0 of $R(\lambda_0, \mu_0)$

```

1: if  $r_0 = 3$  then // cuspidal quartic ( $\mathbb{C}$ )
2:   output: cuspidal quartic
3: else if  $r_0 = 2$  then // two tangent conics ( $\mathbb{C}$ )
4:   if  $I_0 = (1, 1)$  then // pair of planes  $R(\lambda_0, \mu_0)$  is real
5:     output: two tangent conics
6:   else // pair of planes  $R(\lambda_0, \mu_0)$  is imaginary
7:     output: point
8:   end if
9: else //  $r_0 = 1$  : double conic ( $\mathbb{C}$ )
10:  let  $I_1$  be the inertia of  $R(\lambda_1, \mu_1)$ ,  $(\lambda_1, \mu_1)$  the second root of  $\mathcal{D}$ 
11:  if  $I_1 = (2, 1)$  then // cone  $R(\lambda_1, \mu_1)$  is real
12:    output: double conic
13:  else // cone  $R(\lambda_1, \mu_1)$  is imaginary
14:    output:  $\emptyset$ 
15:  end if
16: end if

```

The gcd of the partial derivatives is equal to $32(\lambda + \mu)$. So, by Algorithm 1, \mathcal{D} has a double real root at $(\lambda_0, \mu_0) = (1, -1)$.

We then follow Algorithm 3. Let $R_0 = \lambda_0 S + \mu_0 T$. We have :

$$\det(R_0 - xI) = x^4 - 4x^3 - 8x^2.$$

Descartes' Sign Rule gives that the inertia of R_0 is $I_0 = (1, 1)$ and the rank is $r_0 = 2$. The intersection thus consists of two secant conics over the complexes. We compute

$$\mathcal{E}(\lambda, \mu) = \frac{\mathcal{D}(\lambda, \mu)}{(\mu_0 \lambda - \lambda_0 \mu)^2} = -16(2\lambda^2 + 6\lambda\mu + 5\mu^2).$$

So $\delta = \text{sign}(\text{discriminant}(\mathcal{E})) = -1$ and $s = \text{sign} \mathcal{E}(\lambda_0, \mu_0) = -1$. We conclude that the intersection consists, over the reals, of a single conic.

This conic can be parameterized by (see Part III)

$$\mathbf{X}(u, v) = (2u^2 - 12uv + 18v^2, -u^2 + 2uv + 3v^2, 8v^2, u^2 - 2uv - 3v^2), \quad (u, v) \in \mathbb{P}^1(\mathbb{R}).$$

15.6.2 Example 2

Consider the following pair of quadrics :

$$\begin{cases} Q_S : -5x^2 - 2xy - 4y^2 - 12yz - 6yw - 8z^2 - 4zw + w^2 = 0, \\ Q_T : -2x^2 - 2xy + 3y^2 + 6yz + 4z^2 + 2zw + w^2 = 0. \end{cases}$$

The characteristic polynomial is

$$\mathcal{D}(\lambda, \mu) = \det(\lambda S + \mu T) = -3(16\lambda^4 - 8\lambda^2\mu^2 + \mu^4).$$

The gcd of the partial derivatives is equal to $\mathcal{H} = 12(\mu^2 - 4\lambda^2)$. Since the discriminant δ of \mathcal{H} is not zero, \mathcal{D} has two double roots, according to Algorithm 1. Further, δ is positive and is a square, so \mathcal{D} has two real rational double roots. These roots are $(\lambda_0, \mu_0) = (-1, -2)$ and $(\lambda_1, \mu_1) = (-1, 2)$.

Algorithm 5 Classifying the intersection : the quadruple real root case, $d = 3$.

Require: $R(\lambda, \mu), \mathcal{D}$ (from Algorithm 1)

Require: quadruple real root (λ_0, μ_0) , inertia I_0 and rank r_0 of $R(\lambda_0, \mu_0)$

```

1: if  $r_0 = 3$  then // cubic and tangent line ( $\mathbb{C}$ )
2:   output: cubic and tangent line
3: else if  $r_0 = 2$  then // conic and two lines crossing or two skew lines and a double line ( $\mathbb{C}$ )
4:   if  $I_0 = (2, 0)$  then // pair of planes  $R(\lambda_0, \mu_0)$  is imaginary
5:     output: double line
6:   else // pair of planes  $R(\lambda_0, \mu_0)$  is real
7:     if  $s = +1$  then
8:       let  $l_0$  be the singular line of  $R(\lambda_0, \mu_0)$ 
9:       if  $l_0$  is contained in  $Q_S$  and  $Q_T$  then
10:        output: two skew lines and a double line
11:       else
12:        output: conic and two lines crossing on conic
13:       end if
14:     else //  $s = -1$ 
15:       output: conic
16:     end if
17:   end if
18: else if  $r_0 = 1$  then // two double lines ( $\mathbb{C}$ )
19:   if  $s = +1$  then
20:     output: two secant double lines
21:   else //  $s = -1$ 
22:     output: point
23:   end if
24: else //  $r_0 = 0$  : smooth quadric ( $\mathbb{C}$ )
25:   if  $S$  or  $T$  has inertia  $(2, 2)$  then
26:     output: smooth quadric
27:   else
28:     output:  $\emptyset$ 
29:   end if
30: end if

```

We now follow Algorithm 6. Let $R_0 = \lambda_0 S + \mu_0 T$ and $R_1 = \lambda_1 S + \mu_1 T$. Applying Descartes' Sign Rule, we find that $r_0 = r_1 = 2$. So the intersection, over the complexes, consists of four lines forming a skew quadrilateral. Since $\mathcal{D}(1, 0) < 0$, the characteristic polynomial is negative outside the roots, and $s = -1$. So the intersection, over the reals, consists of two points.

The two points can be computed with the algorithms of Part III :

$$(-3, -3, 3 + \sqrt{3}, -3 - 4\sqrt{3}) \quad \text{and} \quad (-3, -3, 3 - \sqrt{3}, -3 + 4\sqrt{3}).$$

15.6.3 Example 3

Consider the following pair of quadrics :

$$\begin{cases} Q_S : -2xy + 2xw - y^2 - z^2 + w^2, \\ Q_T : 4xy - 4xw + 2y^2 + z^2 - 2w^2. \end{cases}$$

Algorithm 6 Classifying the intersection : the two double roots case, $d = 2$ and discriminant $(\mathcal{H}) \neq 0$.

Require: $R(\lambda, \mu), \mathcal{D}, \mathcal{H}$ (from Algorithm 1)

Require: double roots (λ_0, μ_0) and (λ_1, μ_1)

```

1: let  $\delta :=$  discriminant  $(\mathcal{H})$  and  $s$  be the sign of  $\mathcal{D}$  outside the roots
2: if  $\delta > 0$  and  $\delta$  is a square then // double roots are real and rational
3:   let  $r_0$  and  $r_1$  ( $r_0 \leq r_1$ ) be the ranks of  $R(\lambda_0, \mu_0)$  and  $R(\lambda_1, \mu_1)$ ,  $I_1$  the inertia of the second
4:   if  $r_0 = 3$  and  $r_1 = 3$  then // cubic and secant line ( $\mathbb{C}$ )
5:     output: cubic and secant line
6:   else //  $r_0 = 2$ 
7:     if  $r_1 = 3$  then // conic and two lines not crossing on conic ( $\mathbb{C}$ )
8:       if  $I_1 = (3, 0)$  then // cone  $R(\lambda_1, \mu_1)$  is imaginary
9:         output: point
10:      else // cone  $R(\lambda_1, \mu_1)$  is real
11:        if  $s = +1$  then
12:          output: conic and two lines
13:        else
14:          output: conic and point
15:        end if
16:      end if
17:    else //  $r_1 = 2$  : four lines forming a skew quadrilateral ( $\mathbb{C}$ )
18:      if  $s = -1$  then
19:        output: two points
20:      else if  $R(\lambda, \mu)$  contains a quadric of inertia  $(4, 0)$  then
21:        output:  $\emptyset$ 
22:      else
23:        output: four lines (skew quadrilateral)
24:      end if
25:    end if
26:  end if
27: else // double roots are complex or real non-rational
28:   let  $r_0 := 3 - \dim(\text{singular}(Q_{R(\lambda_0, \mu_0)}))$ 
29:   if  $r_0 = 2$  then //  $R(\lambda_0, \mu_0)$  and  $R(\lambda_1, \mu_1)$  are pairs of planes
30:     if  $\delta < 0$  then // roots are complex conjugate
31:       output: two skew lines
32:     else //  $\delta > 0$  : roots are real conjugate
33:       if  $s = -1$  then
34:         output: two points
35:       else if  $R(\lambda, \mu)$  contains a quadric of inertia  $(4, 0)$  then
36:         output:  $\emptyset$ 
37:       else
38:         output: four lines (skew quadrilateral)
39:       end if
40:     end if
41:   else //  $r_0 = 3$  :  $R(\lambda_0, \mu_0)$  and  $R(\lambda_1, \mu_1)$  are cones
42:     if  $\delta < 0$  then // roots are complex conjugate
43:       output: cubic and non-secant line
44:     else
45:       output: cubic and secant line
46:     end if
47:   end if
48: end if

```

The characteristic polynomial vanishes identically. We then follow Algorithm 2. Q_S has rank 3, and its singular point \mathbf{p} has coordinates $(-1, 1, 0, 1)$. Q_T has rank 3, and its singular point is again \mathbf{p} . So the dimension κ of the intersection of the singular sets of Q_S and Q_T is 0. Let P be the transformation matrix sending \mathbf{p} to $(0, 0, 0, 1)$, completed as follows :

$$P = \begin{pmatrix} 1 & 0 & 0 & -1 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

Let $S' = P^T S P$ and $T' = P^T T P$, and remove the last line and column of the two matrices. This gives :

$$\begin{cases} Q_{S'} : -2xy - y^2 - z^2, \\ Q_{T'} : 4xy + 2y^2 + z^2. \end{cases}$$

The restricted characteristic polynomial is then

$$\mathcal{D}_3(\lambda, \mu) = \lambda^3 - 5\lambda^2\mu + 8\lambda\mu^2 - 4\mu^3.$$

It has a double root at $(\lambda_0, \mu_0) = (2, 1)$. The associated conic $R'_0 = \lambda_0 S' + \mu_0 T'$ has rank 1. So the intersection consists, over the complexes, of two double lines. \mathcal{D}_3 has a second root at $(\lambda_1, \mu_1) = (1, 1)$. The associated conic has inertia $(1, 1)$, from which we conclude that the intersection consists, over the reals, of two double lines.

The two lines can easily be parameterized as follows :

$$\mathbf{X}_1(u, v) = (v, -u - v, 0, u - v) \quad \text{and} \quad \mathbf{X}_2(u, v) = (u, v, 0, v), \quad (u, v) \in \mathbb{P}^1(\mathbb{R}).$$

The two lines meet at point \mathbf{p} .

15.7 Conclusion

In this second part of our paper, we have shown how the real type of the intersection of two quadrics can be determined by extracting simple information from the pencil of the two quadrics, and in particular its characteristic polynomial. Our type-detection algorithm relies on a classification of real pencils of quadrics of $\mathbb{P}^3(\mathbb{R})$, itself derived from the Canonical Pair Form Theorem for pairs of real symmetric matrices [Uhl73, Uhl76].

In Part III, we will use the structural information gathered in the type-detection phase to drive the parameterization process. In each case, we will show that the parameterization computed is near-optimal.

Chapitre 16

Near-optimal parameterization of the intersection of quadrics : III. Parameterizing singular intersections

Cet article a été accepté dans *Journal of Symbolic Computation* [DLLP07c]. (Voir également le rapport de recherche [DLLP05c].) Une version préliminaire à été publiée dans les proceedings du *19th ACM Annual Symposium on Computational Geometry* [DLLP03] et dans la thèse de L. Dupont [Dup04].

Abstract

In Part II, we have shown, using a classification of pencils of quadrics over the reals, how to determine quickly and efficiently the real type of the intersection of two given quadrics.

For each real type of intersection, we design, in this third part, an algorithm for computing a near-optimal parameterization. We also give here examples covering all the possible situations, in terms of both the real type of intersection and the number and depth of square roots appearing in the coefficients.

16.1 Introduction

Building on the classification of pencils of quadrics of $\mathbb{P}^3(\mathbb{R})$ over the reals achieved in Part II and the type-detection algorithm that we deduced from this classification, we now are ready to present near-optimal parameterization algorithms for all the possible types of real intersection.

Since the smooth quartic case has already been thoroughly studied in Part I, we focus here on the singular cases. For each case, we prove the near-optimality of the parameterization and, when there is possibly an extra square root, we describe the test needed to assert the full optimality which always boils down to finding a rational point on a (possibly non-rational) conic.

In what follows, Q_S, Q_T refer to the initial quadrics and Q_R (assumed to be distinct from Q_S) to the intermediate quadric used to parameterize the intersection C of Q_S and Q_T . As in Section 14.7.1, denote by Ω the equation in the parameters :

$$\Omega : \mathbf{X}^T S \mathbf{X} = 0,$$

where \mathbf{X} is the parameterization of Q_R . Denote also by C_Ω the curve zero-set of Ω . Recall that the parameterization of Q_R defines an isomorphism between C and the plane curve C_Ω . When C is singular, its genus is 0 so it can be parameterized by rational functions (i.e. $\sqrt{\Delta}$ can be avoided).

Our general philosophy is to use for Q_R the rational quadric of the pencil of smallest rank. This will lead us to use repeatedly the results of Section 14.6 on the optimality of parameterizations of projective quadrics and to parameterize cones without a rational point, cones with a rational point, pairs of planes, etc. As will be seen, this

philosophy has the double advantage of (i) avoiding $\sqrt{\Delta}$ in all singular cases, and (ii) minimizing the number of radicals. As an additional benefit, it helps keep the size of the numbers involved in intermediate computations and in the final parameterizations to a minimum (see [LPP06]).

For every type of real intersection, we give a set of worst-case examples where the maximum number of square roots is reached, both in the optimal and near-optimal situations (the best-case examples are those given by the canonical forms of Section 15.3). Examples covering all possible situations are gathered in Appendix C.

A summary of the results of this part is given in Table 16.1.

The rest of this third part is as follows. Section 16.2 gives near-optimal parameterization algorithms for all types of real intersection when the pencil is regular. Section 16.3 does the same for singular pencils (i.e., when the characteristic polynomial vanishes identically). Several examples are detailed in Section 16.4 and it is shown how our implementation fares on these examples. Finally, we conclude in Section 16.5 and give a few perspectives.

16.2 Parameterizing degenerate intersections : regular pencils

In this section, we outline parameterization algorithms for all cases of regular pencils, i.e. when the characteristic polynomial does not identically vanish. Information gathered in the type-detection phase (Part II) is used as input; see in particular Table 15.4 but also the details of the classification of pencils over the reals. In each case, we study optimality issues and give worst-case examples.

In the following, we often need to compute the parameterization of the intermediate quadric Q_R and this is achieved using the normal form of Q_R . Recall that a rational congruence sending a quadric with rational coefficients into normal form can be computed using Gauss reduction of quadratic forms into sums of squares (see Part I).

Recall also that the discriminant of a quadric is the determinant of the associated matrix. In the following, we also call *discriminant of a pair of planes* Q_R the product ab where $ax^2 - by^2 = 0$ is the canonical equation of a pair of planes obtained from Q_R by a real rational congruence transformation; the discriminant is defined up to a rational square factor.

16.2.1 Nodal quartic in $\mathbb{P}^3(\mathbb{C})$, $\sigma_4 = [112]$

If we parameterize C using the generic algorithm (see Part I), we will not be able to avoid the appearance of $\sqrt{\Delta}$ because C_Ω (as C) is irreducible. However, since the intersection curve is singular, we know that $\sqrt{\Delta}$ is avoidable by Proposition 14.18. We thus proceed differently.

16.2.1.1 Algorithms

Let λ_1 be the real and rational double root of the characteristic polynomial. Let Q_R be the rational cone associated with λ_1 . As we have found in Section 15.3, there are essentially two cases depending on the real type of the intersection.

Point. Q_R is an imaginary cone. The intersection is reduced to a point, which is the apex of Q_R . Since λ_1 is rational, this apex is rational (otherwise its algebraic conjugate would also be a singular point of the cone). Thus the intersection in this case is defined in \mathbb{Q} .

Real nodal quartic (with or without isolated singularity). Let P be a real rational congruence transformation sending the apex of Q_R to $(0, 0, 0, 1)$. The parameterization

$$\mathbf{X}(u, v, s) = P(x_1(u, v), x_2(u, v), x_3(u, v), s)^T, (u, v, s) \in \mathbb{P}^{*2}$$

of the cone (see Table 14.3) introduces a square root $\sqrt{\delta}$. Equation Ω in the parameters is

$$as^2 + b(u, v)s + c(u, v) = 0,$$

Segre string	real type of intersection	worst case format of parameterization	worst-case optimality of parameterization
[1111]	nonsingular quartic (see part I)	$\mathbb{Q}(\sqrt{\delta})[\xi, \sqrt{\Delta}]$, $\Delta \in \mathbb{Q}(\sqrt{\delta})[\xi]$	rational point on degree-8 surface
[112]	point	\mathbb{Q}	optimal
	nodal quartic	$\mathbb{Q}(\sqrt{\delta})[\xi]$	rational point on conic
[13]	cuspidal quartic	$\mathbb{Q}[\xi]$	optimal
[22]	cubic and non-tangent line	$\mathbb{Q}[\xi]$	optimal
[4]	cubic and tangent line	$\mathbb{Q}[\xi]$	optimal
[11(11)]	two points	$\mathbb{Q}(\sqrt{\delta})$	optimal
	conic	$\mathbb{Q}(\sqrt{\delta}, \sqrt{\mu})[\xi]$, $\mu \in \mathbb{Q}(\sqrt{\delta})$	optimal if $\sqrt{\delta} \notin \mathbb{Q}$ rational point on conic if $\sqrt{\delta} \in \mathbb{Q}$
	two non-tangent conics	$\mathbb{Q}(\sqrt{\delta}, \sqrt{\delta'})[\xi]$	$\mathbb{Q}(\sqrt{\delta'})$ -rational point on $\mathbb{Q}(\sqrt{\delta'})$ -conic
[1(21)]	point	\mathbb{Q}	optimal
	two tangent conics	$\mathbb{Q}(\sqrt{\delta})[\xi]$	optimal
[1(111)]	double conic	$\mathbb{Q}(\sqrt{\delta})[\xi]$	rational point on conic
[2(11)]	point	\mathbb{Q}	optimal
	conic and point	$\mathbb{Q}(\sqrt{\delta})[\xi]$	rational point on conic
	conic and two lines not crossing on the conic	$\mathbb{Q}(\sqrt{\delta})[\xi]$	rational point on conic
[(31)]	conic	$\mathbb{Q}[\xi]$	optimal
	conic and two lines crossing on the conic	$\mathbb{Q}(\sqrt{\delta})[\xi]$	optimal
[(11)(11)]	two points	$\mathbb{K}[\xi]$, degree(\mathbb{K}) = 4	optimal
	two skew lines	$\mathbb{K}[\xi]$, degree(\mathbb{K}) = 4	optimal
	four lines (skew quadrilateral)	$\mathbb{K}[\xi]$, degree(\mathbb{K}) = 4	optimal
[(22)]	double line	$\mathbb{Q}[\xi]$	optimal
	two simple skew lines cutting a double line	$\mathbb{Q}(\sqrt{\delta})[\xi]$	optimal
[(211)]	point	\mathbb{Q}	optimal
	two double concurrent lines	$\mathbb{Q}(\sqrt{\delta})[\xi]$	optimal
[1{3}]	conic and double line	$\mathbb{Q}[\xi]$	optimal
[111]	point	\mathbb{Q}	optimal
	two concurrent lines	$\mathbb{K}[\xi]$, degree(\mathbb{K}) = 4	optimal
	four concurrent lines	$\mathbb{K}[\xi]$, degree(\mathbb{K}) = 4	optimal
[12]	double line	$\mathbb{Q}[\xi]$	optimal
	two simple and a double concurrent lines	$\mathbb{Q}(\sqrt{\delta})[\xi]$	optimal
[3]	concurrent simple and triple lines	$\mathbb{Q}[\xi]$	optimal
[1(11)]	point	\mathbb{Q}	optimal
	two concurrent double lines	$\mathbb{Q}(\sqrt{\delta})[\xi]$	optimal
[(21)]	quadruple line	$\mathbb{Q}[\xi]$	optimal
[11]	quadruple line	$\mathbb{Q}[\xi]$	optimal

TAB. 16.1 – Ring of definition of the projective coordinates of the parameterization of each component of the intersection and optimality, in all cases where the real part of the intersection is 0- or 1-dimensional. $\delta, \delta' \in \mathbb{Q}$.

with a and the coefficients of b, c defined in $\mathbb{Q}(\sqrt{\delta})$. The nodal quartic passes through the vertex of Q_R and this point corresponds to the value $u = v = 0$ of the parameters. At this point $s \neq 0$ because $(u, v, s) \in \mathbb{P}^{*2}$. Thus $a = 0$, Ω is linear in s and it can be solved rationally for s . This leads to the parameterization of the quartic

$$\mathbf{X}(u, v) = P(b(u, v)x_1(u, v), b(u, v)x_2(u, v), b(u, v)x_3(u, v), -c(u, v))^T.$$

The coefficients of $\mathbf{X}(u, v)$ clearly live in $\mathbb{Q}(\sqrt{\delta})[\xi]$, where $\xi = (u, v)$.

When the node of the quartic is not isolated, the singularity is now reached by two different values of $(u, v) \in \mathbb{P}^1(\mathbb{R})$ which are precisely those values such that $b(u, v) = 0$. When the node of the quartic is isolated, the singularity is not reached by $(u, v) \in \mathbb{P}^1(\mathbb{R})$, i.e. $b(u, v) = 0$ has no real solution. In that situation, the node has to be added to the output. Since this point is the vertex of the cone Q_R , it is rational.

16.2.1.2 Optimality

By Proposition 14.14, if the cone Q_R contains a rational point other than its vertex, it can be parameterized with rational coefficients and thus the parameterization of the nodal quartic is defined over $\mathbb{Q}[\xi]$. Otherwise, if Q_R contains no rational point other than its vertex, then the nodal quartic also contains no rational point other than its singular point. Hence the nodal quartic admits no parameterization over $\mathbb{Q}[\xi]$. Therefore, testing whether $\sqrt{\delta}$ can be avoided in the parameterization of real nodal quartics is akin to deciding whether Q_R has a rational point outside its singular locus; furthermore, finding a parameterization in $\mathbb{Q}[\xi]$ amounts to finding a rational point on Q_R outside its singular locus.

There are cases where $\sqrt{\delta}$ cannot be avoided. Example of these are

$$\begin{cases} x^2 + y^2 - 3z^2 = 0, \\ xw + z^2 = 0 \end{cases}$$

when the singularity is not isolated and

$$\begin{cases} x^2 + y^2 - 3z^2 = 0, \\ zw + x^2 = 0 \end{cases}$$

when the singularity is isolated. In both cases, the projective cone corresponding to the double root of the characteristic polynomial is the first equation. By Proposition 14.14, this cone has no rational point except its singular point and $\sqrt{\delta}$ cannot be avoided in the parameterization of the intersection.

16.2.2 Cuspidal quartic in $\mathbb{P}^3(\mathbb{C})$, $\sigma_4 = [13]$

The intersection in this case is always a real cuspidal quartic. As above, using the generic algorithm is not good idea: it would introduce an unnecessary and unwanted $\sqrt{\Delta}$.

We consider instead the cone Q_R associated with the real and rational triple root of the characteristic polynomial. The singular point of the quartic is the vertex \mathbf{p} of Q_R . The intersection of Q_R with the tangent plane of Q_S at \mathbf{p} consists of the double line tangent to C at the cusp. Since it is double, this line is necessarily rational. So we have a rational cone containing a rational line. By Theorem 14.12, this cone admits a rational parameterization.

So we are left with an equation $\Omega: as^2 + b(u, v)s + c(u, v) = 0$ whose coefficients are defined on \mathbb{Q} . As above, the singularity is reached at $(u, v) = (0, 0)$ and at this point $s \neq 0$, so $a = 0$. Thus Ω can be solved rationally for s and the intersection is in $\mathbb{Q}[\xi]$, $\xi = (u, v)$. This is optimal.

16.2.3 Cubic and secant line in $\mathbb{P}^3(\mathbb{C})$, $\sigma_4 = [22]$

The real intersection consists of a cubic and a line. The cubic and the line are either secant or skew. Note that the line of the intersection is necessarily rational, otherwise its algebraic conjugate would also belong to the intersection.

When the double roots of the characteristic polynomial are real and rational, the pencil contains two rational cones Q_{R_1} and Q_{R_2} . The line of C is the rational line joining the vertices of Q_{R_1} and Q_{R_2} . Also, the vertex of Q_{R_2} is a rational point on Q_{R_1} , and vice versa, so the two cones can be rationally parameterized (see Theorem 14.12). Setting up Ω , we have again that it is linear in s , because the line and the cubic intersect at the vertex of the cone,

corresponding to $(u, v) = (0, 0)$. But here the content of Ω in s is linear in (u, v) and it corresponds to the line of C . The cubic is found after dividing by this content and rationally solving for s . The parameterization of the cubic is defined in $\mathbb{Q}[\xi]$.

When the double roots of the characteristic polynomial are either complex conjugate (the cubic and the line are not secant) or real algebraic conjugate (the cubic and the line are secant), there exists quadrics of inertia $(2, 2)$ in the pencil (by Theorems 14.3 and 14.5). We use the generic algorithm of Part I : first find a quadric Q_R of inertia $(2, 2)$ of the pencil through a rational point. Since C contains a rational line, the discriminant of this quadric is a square by Lemma 14.21 and Q_R can be rationally parameterized by Theorem 14.12. Now compute the bidegree $(2, 2)$ equation Ω . The line of C corresponds to a fixed value of one of the parameters and the contents provide factors of bidegree $(1, 0)$ and $(1, 2)$ (or $(0, 1)$ and $(2, 1)$), which are linear in one of the parameters and thus easy to solve rationally for getting a parameterization of the intersection. The parameterization of C is defined in $\mathbb{Q}[\xi]$.

16.2.4 Cubic and tangent line in $\mathbb{P}^3(\mathbb{C})$, $\sigma_4 = [4]$

The real intersection consists of a cubic and a tangent line. The line is necessarily rational, by the same argument as above. The characteristic polynomial has a real and rational quadruple root. To it corresponds a real rational projective cone. Since this cone contains a rational line, it can be rationally parameterized (by Theorem 14.12). The rest is as in the cubic and secant line case when the two roots are rational. The parameterization of the cubic is defined in $\mathbb{Q}[\xi]$.

16.2.5 Two secant conics in $\mathbb{P}^3(\mathbb{C})$, $\sigma_4 = [11(11)]$

In this case, the characteristic polynomial has a double root corresponding to a rational pair of planes Q_R . There are several cases depending on the real type of the intersection.

Two points. The pair of planes Q_R is imaginary. Its rational singular line intersects any other quadric of the pencil in two points. So parameterize the line and intersect it with any quadric of the pencil having rational coefficients. A square root is needed to parameterize the two points if and only if the equation in the parameters of the line has irrational roots.

This situation can happen as the following example shows :

$$\begin{cases} z^2 + w^2 = 0, \\ x^2 - 2y^2 + w^2 = 0. \end{cases}$$

Clearly, the two points are defined by $z = w = 0$ and $x^2 - 2y^2 = 0$ so they live in $\mathbb{Q}[\sqrt{2}]$.

One conic. In this case, the pair of planes is real, the pencil has no quadric of inertia $(2, 2)$ and only one of the planes of Q_R intersects the other quadrics of the pencil.

The algorithm is as follows. First parameterize the pair of planes and separate the two individual planes. Plugging the parameterization of each plane into the equation of Q_S gives two equations of conics in parameter space, with coefficients in $\mathbb{Q}(\sqrt{\delta})$ where δ is the discriminant of the pair of planes. The conics in parameter space correspond to the components of the intersection, thus one of these conics is real and the other is imaginary. Determine the real conic, that is the one with inertia $(2, 1)$, and parameterize it. Substituting this parameterization into the parameterization of the corresponding plane gives a parameterization of the conic of intersection. The parameterization is in $\mathbb{Q}(\sqrt{\delta}, \sqrt{\mu})$, where δ is the discriminant of the pair of planes Q_R and $\sqrt{\mu}$ is the square root needed to parameterize the conic in parameter space, $\mu \in \mathbb{Q}(\sqrt{\delta})$.

If δ is not a square, the parameterization is optimal. Indeed, if the intersection had a real $\mathbb{Q}(\sqrt{\delta})$ -rational point, the conjugate of that point would be on the conjugate conic which is not real. So such a point does not exist and the parameterization is optimal. If δ is a square, the parameterization is defined in $\mathbb{Q}(\sqrt{\mu})[\xi]$ with $\mu \in \mathbb{Q}$. By Proposition 14.14, the parameterization is optimal if and only if the (rational) conic contains no rational point ; moreover, testing if the parameterization is non-optimal and, if so, finding an optimal parameterization is equivalent to finding a rational point on this rational conic.

The situation where δ is a square but the conic has no rational point (the field of the coefficients is of degree two) can be attained for instance with the following pair of quadrics :

$$\begin{cases} (x-w)(x-3w) = 0, \\ x^2 + y^2 + z^2 - 4w^2 = 0. \end{cases}$$

The two planes of the first quadric are rational. The plane $x-w=0$ cuts the second quadric in the conic $x-w=y^2+z^2-3w^2=0$. By Proposition 14.14, this conic has no rational point, so $\sqrt{\delta}$ cannot be avoided and the parameterization of the conic is in $\mathbb{Q}(\sqrt{3})$.

A field extension of degree 4 is obtained with the following quadrics :

$$\begin{cases} x^2 - 4xw - 3w^2 = 0, \\ x^2 + y^2 + z^2 - w^2 = 0. \end{cases}$$

The pair of planes is defined on $\mathbb{Q}(\sqrt{7})$, so, by the above argument, a field extension of degree 4 is unavoidable.

16.2.5.1 Two (secant or non-secant) conics

By contrast to the one conic case, the pencil now contains quadrics of inertia $(2, 2)$. But going through the generic algorithm and factoring C_Ω directly in two curves of bidegree $(1, 1)$ can induce nested radicals. So we proceed as follows. First, find a rational quadric Q_R of inertia $(2, 2)$ through a rational point. This introduces one square root, say $\sqrt{\delta}$. Independently, factor the pair of planes, which introduces another square root $\sqrt{\delta'}$. Now plug the parameterization of Q_R in each of the planes. This gives linear equations in the parameters of Q_R which can be solved without introducing nested radicals. The two conics have a parameterization defined in $\mathbb{Q}(\sqrt{\delta}, \sqrt{\delta'})$.

Note that when the two simple roots of the characteristic polynomial are rational, an alternate approach is to parameterize one of the two rational cones of the pencil instead of a quadric of inertia $(2, 2)$, and then proceed as above.

In terms of optimality, $\sqrt{\delta'}$ cannot be avoided if the planes are irrational. As for the other square root, it can be avoided if and only if the conics contain a point that is rational in $\mathbb{Q}(\sqrt{\delta'})$ (by Proposition 14.14 in which the field \mathbb{Q} can be replaced by $\mathbb{Q}(\sqrt{\delta'})$); moreover, testing if this square root can be avoided and, if so, finding a parameterization avoiding it is equivalent to finding a $\mathbb{Q}(\sqrt{\delta'})$ -rational point on this conic whose coefficients are in $\mathbb{Q}(\sqrt{\delta'})$.

All cases can happen. We illustrate this in the non-secant case. An extension of \mathbb{Q} of degree 4 is needed to parameterize the intersection of the following pair :

$$\begin{cases} x^2 - 33w^2 = 0, \\ y^2 + z^2 - 3w^2 = 0. \end{cases}$$

Indeed, $\sqrt{\delta'} = \sqrt{33}$ cannot be avoided. In addition, by Proposition 14.17, $y^2 + z^2 - 3w^2 - 11x^2 = 0$ has no rational point on $\mathbb{Q}(\sqrt{33})$, thus its intersection with the plane $x=0$, the conic $y^2 + z^2 - 3w^2 = 0$, also has no rational point on $\mathbb{Q}(\sqrt{33})$; hence the cone $y^2 + z^2 - 3w^2 = 0$ has no rational point on $\mathbb{Q}(\sqrt{33})$ except for its singular locus.

An extension field of degree 2 can be obtained by having conics without rational point, but living in rational planes, as in this example :

$$\begin{cases} x^2 - w^2 = 0, \\ y^2 + z^2 - 3w^2 = 0. \end{cases}$$

It can also be attained by having conics living in non-rational planes but having rational points in the extension of \mathbb{Q} defined by the planes :

$$\begin{cases} x^2 - 3w^2 = 0, \\ y^2 + z^2 - 3w^2 = 0. \end{cases}$$

As can be seen, the points of coordinates $(\sqrt{3}, 0, \pm\sqrt{3}, 1)$ belong to the intersection. So the conic has a parameterization in $\mathbb{Q}(\sqrt{3})[\xi]$.

16.2.6 Two tangent conics in $\mathbb{P}^3(\mathbb{C})$, $\sigma_4 = [1(21)]$

Here, the characteristic polynomial has a real and rational triple root, corresponding to a pair of planes Q_R . The other (real and rational) root corresponds to a real projective cone. There are two types of intersection over the reals.

Point. The pair of planes is imaginary and its rational singular line intersects the cone in a double point, which is the only component of the intersection. This point is necessarily rational, otherwise its conjugate would also be in the intersection. One way to compute it is to parameterize the singular line, plug the parameterization in the rational equation of the cone and solve the resulting equation in the parameters.

Two real tangent conics. The pair of planes is real and each of the planes intersects the cone. The singular line of Q_R is tangent to the cone. As above, the point of tangency of the two conics is rational. So, by Proposition 14.14, the conics have a rational parameterization in the extension of \mathbb{Q} defined by the planes. In other words, the conics have a parameterization defined in $\mathbb{Q}(\sqrt{\delta})[\xi]$, where δ is the discriminant of the pair of planes Q_R , if and only if δ is not a square.

One situation where $\sqrt{\delta}$ cannot be avoided is the following :

$$\begin{cases} x^2 - 2w^2 = 0, \\ xy + z^2 = 0. \end{cases}$$

16.2.7 Double conic in $\mathbb{P}^3(\mathbb{C})$, $\sigma_4 = [1(111)]$

The characteristic polynomial has a real rational triple root, corresponding to a double plane. The other root gives a rational cone. Assume this cone is real (otherwise the intersection is empty).

To obtain the parameterization of the double conic, first parameterize the double plane. Then plug this parameterization in the equation of the cone. This gives the rational equation of the conic (in the parameters of the plane). If the conic has a rational point, then it can be rationally parameterized. Otherwise, one square root is needed.

One worst-case situation where a square root is always needed is the following :

$$\begin{cases} x^2 = 0, \\ y^2 + z^2 - 3w^2 = 0. \end{cases}$$

By Proposition 14.14, the second quadric (a cone) has no rational point outside its vertex. Thus the conic cannot be parameterized rationally.

16.2.8 Conic and two lines not crossing on the conic in $\mathbb{P}^3(\mathbb{C})$, $\sigma_4 = [2(11)]$

The characteristic polynomial has two double roots, corresponding to a cone and a pair of planes which is always real. The two roots are necessarily real and rational, otherwise the quadrics associated with them in the pencil would have the same rank. So both the cone and the pair of planes are rational. Also, the vertex of the cone falls on the pair of planes outside its singular line. Thus, by Proposition 14.13, the discriminant of the pair of planes is a square and each individual plane has a rational parameterization.

Over the reals, there are three cases.

Point. The projective cone is imaginary. The intersection is limited to its real vertex. Since the cone is rational, its vertex is rational.

Point and conic. The cone is now real. One of the planes cuts the cone in a conic living in a rational plane, the other plane cuts the cone in its vertex. The point of the intersection is this vertex and it is rational. To parameterize the conic of the intersection, plug the parameterization of the plane that does not go through the vertex of the cone. This gives a rational conic in the parameters of the plane. One square root is possibly needed to parameterize this conic. It can be avoided if and only if the conic has a rational point.

One example where the square root cannot be avoided is the following :

$$\begin{cases} xw = 0, \\ y^2 + z^2 - 3w^2 = 0. \end{cases}$$

By Proposition 14.14, the projective cone has no rational point other than its vertex $(1, 0, 0, 0)$. So the conic $x = y^2 + z^2 - 3w^2 = 0$ has no rational point.

Two lines and conic. Again, the cone is real and one plane cuts it in a rational nonsingular conic. But now the second plane, going through the vertex of the cone, further cuts the cone in two lines. The parameterization of the conic goes as above. To represent the lines, we plug the second plane in the equation of the cone and parameterize.

Note that if the lines are rational, then the cone contains a rational line and can be rationally parameterized. Since the conic is the intersection of this cone with a rational plane, it has a rational parameterization. So in that case all three components have parameterizations in $\mathbb{Q}[\xi]$. If the lines are irrational, it can still happen that the conic has a rational point and thus a rational parameterization.

We give examples for the three situations we just outlined. First, the pair

$$\begin{cases} xy = 0, \\ y^2 + z^2 - w^2 = 0 \end{cases}$$

gives birth to the rational lines $y = z \pm w = 0$ and the rational conic $x = y^2 + z^2 - w^2 = 0$ which contains the rational point $(0, 0, 1, 1)$ and can be rationally parameterized. Second, the pair of quadrics

$$\begin{cases} xy = 0, \\ 2y^2 + z^2 - 3w^2 = 0 \end{cases}$$

has as intersection the two irrational lines $y = z \pm \sqrt{3}w = 0$ and the conic $x = 2y^2 + z^2 - 3w^2 = 0$ which contains the rational point $(0, 1, 1, 1)$ so can be rationally parameterized. Finally, the lines and the conic making the intersection of the quadrics

$$\begin{cases} xy = 0, \\ y^2 + z^2 - 3w^2 = 0 \end{cases}$$

cannot be rationally parameterized. Indeed, by Proposition 14.14, the cone has no rational point outside the vertex $(1, 0, 0, 0)$, so the conic $x = y^2 + z^2 - 3w^2 = 0$ has no rational point.

16.2.9 Conic and two lines crossing on the conic in $\mathbb{P}^3(\mathbb{C})$, $\sigma_4 = [(31)]$

The characteristic polynomial has a real pair of planes Q_R corresponding to a real and rational quadruple root. The asymmetry in the sizes of the Jordan blocks associated with this root (the two blocks have size 1 and 3) implies that the individual planes of this pair are rational. The conic of the intersection is always real and the two lines (real or imaginary) cross on the conic.

There are two types of intersection over the reals.

Conic. The point at which the two lines cross is the double point that is the intersection of the singular line of Q_R with any other quadric of the pencil. This point is necessarily rational. So the conic can be rationally parameterized by Proposition 14.14.

Conic and two lines. To parameterize the intersection, first compute the parameterization of the two planes of Q_R . Plugging these parameterizations in the equation of any other quadric of the pencil yields a conic on one side and a pair of lines on the other side. As above, the conic can be rationally parameterized. As for the two lines, they have a rational parameterization if and only if the discriminant of the pair of lines is a square.

One situation where this discriminant is not a square is as follows :

$$\begin{cases} yz = 0, \\ y^2 + xz - 2w^2 = 0. \end{cases}$$

The conic is given by $y = xz - 2w^2 = 0$ which contains the rational point $(1, 0, 0, 0)$ and can be rationally parameterized. The lines are defined by $z = y^2 - 2w^2 = 0$. But the pair of planes $y^2 - 2w^2 = 0$ has no rational point outside its singular locus so the lines are defined in $\mathbb{Q}(\sqrt{2})$.

16.2.10 Two skew lines and a double line in $\mathbb{P}^3(\mathbb{C})$, $\sigma_4 = [(22)]$

The characteristic polynomial has a real and rational quadruple root, which corresponds to a pair of planes. The singular line of the pair of planes is contained in all the quadrics of the pencil. There are two cases.

Double line. The pair of planes is imaginary. The intersection is reduced to the rational singular line of the pair of planes. So C is defined in $\mathbb{Q}[\xi]$.

Two simple lines and a double line. The pair of planes is real. We can factor it into simple planes, parameterize these planes and plug them in any other quadric of the pencil. The two resulting equations in the parameters of the planes are pairs of lines, each pair containing the double line of the intersection and one of the simple lines. The simple lines are rational if and only if the discriminant of the pair of planes is a square.

A situation where the two simple lines are irrational is the following :

$$\begin{cases} y^2 - 2w^2 = 0, \\ xy - zw = 0. \end{cases}$$

16.2.11 Two double lines in $\mathbb{P}^3(\mathbb{C})$, $\sigma_4 = [(211)]$

The characteristic polynomial has a real rational quadruple root, which corresponds to a double plane. The double plane cuts any other quadric of the pencil in two double lines in $\mathbb{P}^3(\mathbb{C})$. There are two cases.

Point. Except for the double plane, the pencil consists of quadrics of inertia $(3, 1)$. The two lines are imaginary. The intersection is reduced to their rational intersection point, i.e. the point at which the double plane is tangent to the other quadrics of the pencil.

Two real double lines. Except for the double plane, the pencil consists of quadrics of inertia $(2, 2)$. The two lines are real. To parameterize them, first compute a parameterization of the double plane and then plug it in any quadric of inertia $(2, 2)$ of the pencil. The resulting pair of lines can easily be parameterized. The intersection is thus parameterized with one square root if and only if the lines are irrational.

One case where the square root cannot be avoided is as follows :

$$\begin{cases} w^2 = 0, \\ x^2 - 2y^2 + zw = 0. \end{cases}$$

The lines $w = x^2 - 2y^2 = 0$ have no rational point except for their singular point $(0, 0, 1, 0)$ so their parameterization is in $\mathbb{Q}(\sqrt{2})[\xi]$.

16.2.12 Four lines forming a skew quadrilateral in $\mathbb{P}^3(\mathbb{C})$, $\sigma_4 = [(11)(11)]$

We start by describing the algorithms we use in this case. We then prove the optimality of the parameterizations and conclude the section by giving examples of pairs of rational quadrics for all possible types of real intersections and extension fields.

16.2.12.1 Algorithms

In this case the characteristic polynomial has two double roots that correspond to (possibly imaginary) pairs of planes. It can be written in the form

$$\mathcal{D}(\lambda, \mu) = \gamma(a\lambda^2 + b\lambda\mu + c\mu^2)^2 = 0, \quad (16.1)$$

with γ, a, b , and c in \mathbb{Q} .

In order to minimize the number and depth of square roots in the coefficients of the parameterization of the intersection, we proceed differently depending on the type of the real intersection and the values of γ and $\delta = b^2 - 4ac$.

Note that the roots of the characteristic polynomial are defined in $\mathbb{Q}(\sqrt{\delta})$ and thus the coefficients of the pairs of planes in the pencil also live in $\mathbb{Q}(\sqrt{\delta})$. Let $d^+, d^- \in \mathbb{Q}(\sqrt{\delta})$ be the discriminants of the two pairs of planes, with $d^+ > d^-$. When $d^+ > 0$ (resp. $d^- > 0$), the corresponding pair of planes is real and can be factored into two planes that are defined over $\mathbb{Q}(\sqrt{d^+})$ (resp. $\mathbb{Q}(\sqrt{d^-})$). The algorithms in the different cases are as follows.

Two points. In this case one pair of planes of the pencil is real (the one with discriminant d^+) and the other is imaginary. We factor the two real planes and substitute in each a parameterization of the (real) singular line of the imaginary pair of planes. The singular line is defined in $\mathbb{Q}(\sqrt{\delta})$ and each of the real planes are defined in $\mathbb{Q}(\sqrt{d^+})$. We thus obtain the two points of intersection with coordinates in $\mathbb{Q}(\sqrt{\delta}, \sqrt{d^+})$. The two points are thus defined over $\mathbb{Q}(\sqrt{d^+})$, $d^+ \in \mathbb{Q}(\sqrt{\delta})$, an extension field of degree 4 (in the worst case) with one nested square root.

Two or four lines. Since the intersection is contained in every quadric of the pencil, there are no quadric of inertia $(3, 1)$ in the pencil in this case (such quadrics contain no line) and thus $\gamma > 0$. Furthermore all the non-singular quadrics of the pencil have inertia $(2, 2)$ (by Theorem 14.5) and their discriminant is equal to γ , up to a square factor (by Eq. (16.1)). Hence we can parameterize a quadric Q_R of inertia $(2, 2)$ in the pencil using the parameterization of Table 14.3 with coefficients in $\mathbb{Q}(\sqrt{\gamma})$ (see Part I).

There are three subcases.

$\sqrt{\delta} \in \mathbb{Q}$. The roots of the characteristic polynomial are real (since $\delta > 0$), thus the intersection consists of four real lines and the two pairs of planes of the pencil are real (see Table 15.4). We factor the two pairs of planes into four planes with coefficients in $\mathbb{Q}(\sqrt{d^\pm})$ and intersect them pairwise. We thus obtain a parameterization of the four lines over $\mathbb{Q}(\sqrt{d^+}, \sqrt{d^-})$ with $d^\pm \in \mathbb{Q}$ (since δ is a square), an extension field of degree 4 (in the worst case) with no nested square root.

$\sqrt{\delta} \notin \mathbb{Q}$ and $\sqrt{\gamma\delta} \in \mathbb{Q}$. Here again $\delta > 0$ thus the intersection consists of four real lines and the two pairs of planes of the pencil are real. We factor one of these pairs of planes (say the one with discriminant d^+) in two planes with coefficients in $\mathbb{Q}(\sqrt{d^+})$; if the discriminant of one of the pair of planes is a square, we choose this pair of planes for the factorization. We then substitute the parameterization of the quadric Q_R into each plane. This leads to an equation of bidegree $(1, 1)$ in the parameters with coefficients in $\mathbb{Q}(\sqrt{d^+}, \sqrt{\gamma})$. This field is equal to $\mathbb{Q}(\sqrt{d^+})$ because $d^+ \in \mathbb{Q}(\sqrt{\delta})$ and $\gamma\delta$ is a square. We finally obtain each line by factoring the equation in the parameters into terms of bidegree $(1, 0)$ and $(0, 1)$ and by substituting the solutions of these factors into the parameterization of Q_R . We thus obtain a parameterization of the four lines defined over $\mathbb{Q}(\sqrt{d^+})$, $d^+ \in \mathbb{Q}(\sqrt{\delta})$, an extension field of degree 4 (in the worst case) with one nested square root.

$\sqrt{\delta} \notin \mathbb{Q}$ and $\sqrt{\gamma\delta} \notin \mathbb{Q}$. In this case we apply the generic algorithm of Part I : we substitute the parameterization of Q_R into the equation of another quadric of the pencil (with rational coefficients). The resulting equation in the parameters of bidegree $(2, 2)$ has coefficients in $\mathbb{Q}(\sqrt{\gamma})$. We factor it into two terms of bidegree $(2, 0)$ and $(0, 2)$, whose coefficients also live in $\mathbb{Q}(\sqrt{\gamma})$. We solve each term separately and each real solution leads to a real line. At least one of the two factors has two real solutions, which are defined in an extension field of the form $\mathbb{Q}(\sqrt{\alpha_1 + \alpha_2\sqrt{\gamma}})$, $\alpha_i \in \mathbb{Q}$. If the other factor has real solutions, they are defined in $\mathbb{Q}(\sqrt{\alpha_1 - \alpha_2\sqrt{\gamma}})$. Thus in the case where the intersection consists of two real lines, we obtain parameterization defined over an extension

field $\mathbb{Q}(\sqrt{\alpha_1 + \alpha_2\sqrt{\gamma}})$ of degree 4 (in the worst case), with one nested square root. In the case where the intersection consists of four real lines, the parameterization of the four lines altogether is defined over an extension field of degree 8 (in the worst case) but each of the lines is parameterized over an extension $\mathbb{Q}(\sqrt{\alpha_1 + \alpha_2\sqrt{\gamma}})$ or $\mathbb{Q}(\sqrt{\alpha_1 - \alpha_2\sqrt{\gamma}})$ of degree 4 (in the worst case), with one nested square root.

16.2.12.2 Optimality

We prove that the algorithms described above output parameterizations that are always optimal in the number and depth of square roots appearing in their coefficients. This proof needs some considerations of Galois theory that can be found in Appendix A.

The two input quadrics intersect here in four lines in $\mathbb{P}^3(\mathbb{C})$. The pencil contains two (possibly complex) pair of planes and these lines are the intersections between two planes taken in two different pairs of planes. Let $\mathbf{p}_1, \dots, \mathbf{p}_4$ be their pairwise intersection points of the four lines. These points are the singular points of the intersection. These points are also the intersections of the singular line of a pair of planes with the other pair of planes, and vice versa. Let the points be numbered such that \mathbf{p}_1 and \mathbf{p}_3 are on the singular line of one pair of planes of the pencil ; \mathbf{p}_2 and \mathbf{p}_4 are then on the singular line of the other pair of planes of the pencil. The four lines of intersection are thus $\mathbf{p}_1\mathbf{p}_2$, $\mathbf{p}_2\mathbf{p}_3$, $\mathbf{p}_3\mathbf{p}_4$, and $\mathbf{p}_4\mathbf{p}_1$.

Let \mathbb{K} be the field of smallest degree on which the four points \mathbf{p}_i are rational. The above algorithms show that \mathbb{K} has degree 1, 2, 4 or 8 (since two rational lines in \mathbb{K} intersect in a rational point in \mathbb{K}). Let G be its Galois group, which acts by permutations on the points \mathbf{p}_i . It follows that G is a subgroup of the dihedral group D_4 of order 8 of the symmetries of the square. This group D_4 acts on the four points \mathbf{p}_i and on the lines joining them the way the 8 isometries of a square act on its vertices and edges. We show that the optimal number of square roots needed for parameterizing the four lines and the way this optimal number is reached only depend on G and on its action on the \mathbf{p}_i .

The eight elements of D_4 are the identity, the transpositions τ_{13} and τ_{24} which exchange \mathbf{p}_1 and \mathbf{p}_3 or \mathbf{p}_2 and \mathbf{p}_4 (symmetries with respect to the diagonal), the permutation $\tau_{12,34}$ (resp. $\tau_{14,23}$) of order 2 which exchange \mathbf{p}_1 with \mathbf{p}_2 and \mathbf{p}_3 with \mathbf{p}_4 (resp. \mathbf{p}_1 with \mathbf{p}_4 and \mathbf{p}_2 with \mathbf{p}_3), the circular permutations ρ and ρ^{-1} of order 4, and the permutation $\rho^2 = \tau_{13}\tau_{24} = \tau_{12,34}\tau_{14,23}$ of order 2.

If G is included in the group $G_{\mathcal{L}}$ of order 4 generated by τ_{13} and τ_{24} (symmetries of the lozenge), its action leaves fixed the pairs $\{\mathbf{p}_1, \mathbf{p}_3\}$ and $\{\mathbf{p}_2, \mathbf{p}_4\}$ and thus also the lines $\mathbf{p}_1\mathbf{p}_3$ and $\mathbf{p}_2\mathbf{p}_4$ and the two singular quadrics of the pencil (the two pairs of planes). It follows that the roots of the characteristic polynomial \mathcal{D} are rational. Conversely, if these roots are rational, the singular quadrics and their singular lines are invariant under the action of G , as well as the pairs $\{\mathbf{p}_1, \mathbf{p}_3\}$ and $\{\mathbf{p}_2, \mathbf{p}_4\}$, which implies that G is included in $G_{\mathcal{L}}$. A similar argument shows that G is the identity (resp. is generated by τ_{13} (or τ_{24}), or contains $\tau_{13}\tau_{24}$), if and only if 0 (resp. 1 or 2) of the singular quadrics consist of irrational planes. Moreover, in the case where G contains $\tau_{13}\tau_{24}$, the group is different from $G_{\mathcal{L}}$ if and only if any element which exchanges \mathbf{p}_1 and \mathbf{p}_3 also exchanges \mathbf{p}_2 and \mathbf{p}_4 , i.e. if and only if the conjugations exchanging the planes in each singular quadric is the same (implying that the square roots needed for factoring them are one and the same). As the degree of \mathbb{K} is the order of G , this shows that the number of square roots needed in our algorithm is always optimal if the roots of \mathcal{D} are rational (i.e. δ is a square).

When the roots of \mathcal{D} are not rational, we consider, in the algorithm, a rational quadric Q_R passing through a rational point \mathbf{p} . Let D be the line of Q_R passing through \mathbf{p} and intersecting the lines $\mathbf{p}_1\mathbf{p}_2$ and $\mathbf{p}_3\mathbf{p}_4$ in two points \mathbf{q}_1 and \mathbf{q}_2 . If the discriminant of Q_R is a square (and its parameterization is rational), then D is rational and is fixed by any Galois automorphism. It follows that the lines $\mathbf{p}_1\mathbf{p}_2$ and $\mathbf{p}_3\mathbf{p}_4$ are either fixed or exchanged, which implies that G is included in the group $G_{\mathcal{R}}$ of order 4 generated by $\tau_{12,34}$ and $\tau_{14,23}$ (symmetries of the rectangle). Conversely, if $G \subset G_{\mathcal{R}}$, the lines $\mathbf{p}_1\mathbf{p}_2$ and $\mathbf{p}_3\mathbf{p}_4$ are fixed or exchanged by any Galois automorphism ; the image of D by such an automorphism is D itself or the other line of Q_R passing through \mathbf{p} ; as this image contains the images of \mathbf{q}_1 and \mathbf{q}_2 which are on $\mathbf{p}_1\mathbf{p}_2$ or $\mathbf{p}_3\mathbf{p}_4$, we may conclude that D is fixed by any Galois automorphism, and is rational ; this shows that the discriminant of Q_R is a square by Lemma 14.21. Pushing these arguments a little more, it is easy to see that G is generated by $\tau_{12,34}$ or $\tau_{14,23}$ (or is the identity) if and only if the roots of either or both of the factors of bidegree $(2, 0)$ and $(0, 2)$ of the equation Ω in the parameters are rational.

By similar arguments of invariance, we may also conclude that the group G is generated by the circular permutation ρ if and only if any Galois automorphism which exchanges the lines $\mathbf{p}_1\mathbf{p}_3$ and $\mathbf{p}_2\mathbf{p}_4$ exchanges also the lines of Q_R passing through \mathbf{p} (and if there is such an automorphism). It follows that this case occurs when the square root of the discriminant of Q_R and the roots of \mathcal{D} generate the same field.

Finally, G is of order 8 if none of the preceding cases occur.

Optimality in all cases is proved by checking that, for each possible group, the algorithm involves exactly 0, 1 or 2 square roots for parameterizing the lines if the size of the orbit of a line is 1, 2 or 4 respectively.

16.2.12.3 Examples

We now give examples in all the possible cases outlined in Section 16.2.12.1. These examples are obtained using the following proposition which gives a rational canonical form for pencils having two double roots corresponding to quadrics of rank 2. Its proof is postponed to Appendix B. and needs again elements of Galois theory that can be found in Appendix A.

Proposition 16.1. *Let $R(\lambda, \mu)$ be a rational pencil of quadrics whose characteristic polynomial has two double roots corresponding to two quadrics of rank 2. Then there is a rational change of frame such that the pencil is generated by quadrics (Q_S, Q_T) of equation*

$$\begin{cases} x^2 - \gamma y^2 - 2wz = 0, \\ \alpha x^2 + 2\gamma xy + \alpha \gamma y^2 - z^2 - (\alpha^2 - \gamma)w^2 = 0, \end{cases}$$

where $\alpha, \gamma \in \mathbb{Q}, \delta = \alpha^2 - \gamma \neq 0$. Moreover, a field \mathbb{K} of smallest degree on which the four lines of the intersection are rationally parameterized is generated by the roots of $t^4 - 2\alpha t + \gamma = 0$.

The different cases are as follows :

- $\delta > 0, \gamma > 0$, and $\alpha < 0$: the real intersection is empty.
- $\delta > 0, \gamma > 0$, and $\alpha > 0$: the real intersection consists of four lines.
- $\delta > 0$ and $\gamma < 0$: the real intersection consists of two points.
- $\delta < 0$: the real intersection consists of two lines.

Note that the characteristic polynomial for the reduced pair of quadrics Q_S and Q_T is

$$\mathcal{D}(\lambda, \mu) = \gamma(\lambda^2 - \delta\mu^2)^2 = 0.$$

Its roots are $(\lambda_0, \mu_0) = (\pm\sqrt{\delta}, 1)$ and, when $\delta > 0$, the associated quadrics of the pencil are the pairs of planes of equations

$$\lambda_0 Q_S + \mu_0 Q_T = (\alpha \pm \sqrt{\delta}) \left(x + (\alpha \mp \sqrt{\delta})y \right)^2 - (z \pm \sqrt{\delta}w)^2 = 0$$

and of discriminants $d^\pm = \alpha \pm \sqrt{\delta}$. Note also that, when $\gamma > 0$, the quadric Q_S has inertia $(2, 2)$ and can be parameterized, using the parameterization of Table 14.3, by :

$$\mathbf{X} = \left(ut + vs, \frac{ut - vs}{\sqrt{\gamma}}, vt, 2us \right), \quad (u, v), (s, t) \in \mathbb{P}^1(\mathbb{R}).$$

Plugging \mathbf{X} into the equation of T gives the following biquadratic equation in the parameters :

$$\Omega : (2(\alpha + \sqrt{\gamma})u^2 - v^2)(2(\alpha - \sqrt{\gamma})s^2 - t^2) = 0.$$

We can now give examples in all the possible cases outlined in Section 16.2.12.1. We start with the four real lines case :

- δ is a square :
 - If $(\alpha, \gamma) = (5, 9)$, then $\sqrt{\delta} = 4$, the discriminants of the pairs of planes are $d^\pm = 5 \pm 4$, so $\sqrt{d^\pm} \in \mathbb{Q}$, and the four lines are defined in \mathbb{Q} .
 - If $(\alpha, \gamma) = (3, 5)$, then $\sqrt{\delta} = 2$, the discriminants of the pairs of planes are $d^\pm = 3 \pm 2$, so $\sqrt{d^+} \notin \mathbb{Q}$ and $\sqrt{d^-} \in \mathbb{Q}$, and the four lines are defined in $\mathbb{Q}(\sqrt{5})$.
 - If $(\alpha, \gamma) = (5, 16)$, then $\sqrt{\delta} = 3$, the discriminants of the pair of planes are $d^\pm = 5 \pm 3$, so $\sqrt{d^\pm} \notin \mathbb{Q}$ but $\sqrt{d^-}$ and $\sqrt{d^+}$ generate the same field $\mathbb{Q}(\sqrt{2})$, and the four lines are defined in $\mathbb{Q}(\sqrt{2})$.
 - If $(\alpha, \gamma) = (6, 20)$, then $\sqrt{\delta} = 4$, the discriminants of the pairs of planes are $d^\pm = 6 \pm 4$, so $\sqrt{d^\pm} \notin \mathbb{Q}$, $\sqrt{d^+}$ and $\sqrt{d^-}$ do not generate the same field and the four lines are defined in $\mathbb{Q}(\sqrt{2}, \sqrt{10})$.

- δ is not a square but $\gamma\delta$ is a square : let $(\alpha, \gamma) = (2, 2)$, then $\delta = 2$, $\sqrt{\delta} \notin \mathbb{Q}$ but $\sqrt{\gamma\delta} = 2 \in \mathbb{Q}$. The discriminant of the pairs of planes are $d^\pm = 2 \pm \sqrt{2}$, so the four lines are defined in $\mathbb{Q}(\sqrt{2 + \sqrt{2}})$.
- Neither δ nor $\gamma\delta$ are squares :
 - If $(\alpha, \gamma) = (3, 1)$, then $\delta = \gamma\delta = 8$ is not a square, $\sqrt{\gamma} \in \mathbb{Q}$ so Q_R can be rationally parameterized and the factors of bidegree $(2, 0)$ and $(0, 2)$ of Ω have rational coefficients. Since $2(\alpha - \sqrt{\gamma}) = 4$ is a square, one of those factors splits in two rational linear factors. Thus two lines have a rational parameterization. Since $2(\alpha + \sqrt{\gamma}) = 8$, the other two lines are defined in $\mathbb{Q}(\sqrt{8}) = \mathbb{Q}(\sqrt{2})$.
 - If $(\alpha, \gamma) = (2, 1)$, then $\delta = \gamma\delta = 3$ is not a square, $\sqrt{\gamma} \in \mathbb{Q}$ so Q_R can be rationally parameterized. Since $2(\alpha + \sqrt{\gamma}) = 6$ and $2(\alpha - \sqrt{\gamma}) = 2$, two lines have a rational parameterization in $\mathbb{Q}(\sqrt{2})$ and the other two lines have a rational parameterization in $\mathbb{Q}(\sqrt{6})$.
 - If $(\alpha, \gamma) = (3, 3)$, then $\delta = 6$ and $\gamma\delta = 18$ are not squares, $\sqrt{\gamma} \notin \mathbb{Q}$ so Q_R cannot be rationally parameterized. Two lines are defined in $\mathbb{Q}(\sqrt{6 + 2\sqrt{3}})$ and the other two lines are rational in $\mathbb{Q}(\sqrt{6 - 2\sqrt{3}})$.

Now we give examples for the two real lines case :

- If $(\alpha, \gamma) = (3, 25)$, then $\delta = -16 < 0$, $\sqrt{\gamma} = 5$ and $\sqrt{2(\alpha + \sqrt{\gamma})} = 4$ are rational, so the two lines are rational.
- If $(\alpha, \gamma) = (1, 4)$, then $\delta = -3 < 0$, $\sqrt{\gamma} \in \mathbb{Q}$ and $\sqrt{2(\alpha + \sqrt{\gamma})} = \sqrt{6} \notin \mathbb{Q}$, so the two lines are defined in $\mathbb{Q}(\sqrt{6})$.
- If $(\alpha, \gamma) = (1, 3)$, then $\delta = -2 < 0$, $\sqrt{\gamma}$ and $\sqrt{2(\alpha + \sqrt{\gamma})}$ are not rational, so the two lines are defined in $\mathbb{Q}(\sqrt{\sqrt{3} - 1})$.

Finally, here are examples for the two points case :

- If $(\alpha, \gamma) = (0, -1)$, then $\sqrt{\delta} = 1$ is rational and the discriminant $\alpha + \sqrt{\delta} = 1$ of the real pair of planes is a square, so the two points are in \mathbb{Q} .
- If $(\alpha, \gamma) = (1, -3)$, then $\sqrt{\delta} = 2$ is rational but the discriminant $\alpha + \sqrt{\delta} = 3$ of the real pair is not a square, so the two points are in $\mathbb{Q}(\sqrt{3})$.
- If $(\alpha, \gamma) = (1, -2)$, then $\sqrt{\delta}$ is not rational and the two points are in $\mathbb{Q}(\sqrt{1 + \sqrt{3}})$.

16.3 Parameterizing degenerate intersections : singular pencils

We now turn to singular pencils. Except when the intersection consists of four concurrent lines, the parameterization algorithms are straightforward and therefore only briefly sketched.

16.3.1 Conic and double line in $\mathbb{P}^3(\mathbb{C})$, $\sigma_4 = [1\{3\}]$

As we have seen in Section 15.4.1, the pencil contains in this case one pair of planes. Furthermore each of the planes is rational by Proposition 14.13 because the pair of planes contains a rational point outside its singular locus (by Proposition 15.7). One plane is tangent to all the cones of the pencil, giving a rational double line. The other plane intersects all the cones transversally, giving a conic. The conic contains a rational point (its intersection with the singular line of the planes), so it can be rationally parameterized.

To actually parameterize the line and the conic, we proceed as follows. If Q_S is a pair of planes, replace Q_S by $Q_S + Q_T$. Now, Q_S is a real projective cone whose vertex is on Q_T . Use this rational vertex to obtain a rational parameterization of Q_T . Plug this parameterization into the equation of Q_S . This equation in the parameters factors in a squared linear factor (corresponding to the double line) and a bilinear factor, corresponding to the conic. It can rationally be solved. The parameterization of C is defined in $\mathbb{Q}[\xi]$.

16.3.2 Four concurrent lines in $\mathbb{P}^3(\mathbb{C})$, $\sigma_3 = [111]$

In this case, and in the three following cases, the two quadrics Q_S and Q_T have a singular point in common. So first compute this singular point \mathbf{p} , which is rational, and compute the rational transformation sending this point to $(0, 0, 0, 1)$. In this new frame, Q_S and Q_T are functions of x, y, z only and we can look at the restricted characteristic polynomial of the 3×3 upper left matrices to determine the real type of the intersection.

When the restricted characteristic polynomial has three simple roots (in \mathbb{C}), there are three types of intersection over the reals : a point, two concurrent lines, or four concurrent lines (see Table 15.5). In the first case, the four lines are imaginary and the real part of the intersection consists of their common rational point, i.e. the point \mathbf{p} .

We now look at the two other cases.

16.3.2.1 Algorithm and optimality

The algorithm for computing the lines is as follows. Determine a plane $x = 0, y = 0, z = 0$, or $w = 0$, that does not contain the singular point \mathbf{p} . Substitute the equation of that plane (say $x = 0$) into the equations of $Q_S(x, y, z, w)$ and $Q_T(x, y, z, w)$. This gives a system of two non-homogeneous degree-two equations in three variables having four distinct complex projective solutions \mathbf{q}_i . The real lines of C are then the two or four lines going through \mathbf{p} and one of the \mathbf{q}_i with real coordinates, $i = 1, \dots, 4$.

This algorithm outputs an optimal parameterization of C . Indeed, since the common singular point \mathbf{p} of Q_S and Q_T is rational and the plane ($x = 0$) used to cut Q_S and Q_T is rational, the lines are rational if and only if their intersection with the planes (the points \mathbf{q}_i) are rational.

16.3.2.2 Degree of the extension

The following result shows that the roots of any polynomial of degree 4 without multiple root may be needed to parameterize four real concurrent lines. It uses notions of Galois theory that can be found in Appendix A.

Proposition 16.2. *For any rational univariate polynomial of degree 4 without multiple root, there are rational pencils of quadrics whose intersection is four (real or imaginary) concurrent lines, such that each of them is rational on the field generated by one of the roots of the polynomial and is not rational on any smaller field (for the inclusion and the degree).*

Proof. Let us consider a polynomial of degree 4 with rational coefficients and without multiple factors. Let us consider its four real or complex roots t_1, \dots, t_4 and the four points \mathbf{q}_i of coordinates $(1, t_i, t_i^2, 0)$. Let us consider also two rational points $\mathbf{r}_j = (a_j, b_j, c_j, 0), j = 1, 2$. Exactly one conic exists in the plane $w = 0$, which passes through the four points \mathbf{q}_i and one of the \mathbf{r}_j . Each of these conics has necessarily a rational equation, because, if it were not, the conjugate conics (under the action of the Galois group of the field containing the coefficients) would pass through the same five points. In other words, the equation of the conic is invariant under the Galois group and is thus rational. Now the rational cones containing these conics and having the point $(0, 0, 0, 1)$ as vertex intersect in four (real or imaginary) lines passing through this vertex and the points \mathbf{q}_i .

The equations of these conics are easy to compute explicitly. Consider a conic with generic coefficients. Expressing that it passes through 5 points induces five linear equations in the coefficients of the equation of the conic. Solving this linear system expresses these coefficients as symmetric functions of the t_i , and thus as rational functions of the coefficients of the polynomial of degree 4. \square

16.3.2.3 Examples

The proof of Proposition 16.2 gives a way of constructing examples of pencils of quadrics whose intersection is four concurrent lines for any quartic f without multiple root. Table 16.2 shows an exhaustive list of examples covering the possible degrees of field extension on which the lines of intersection are defined. We here focus on the cases where f has two or four real roots, corresponding, respectively, to the two and four concurrent lines cases.

When f has four real roots, the degree of the extension of \mathbb{Q} needed to parameterize the four lines together is the order of the Galois group of f , in view of Proposition 16.2 and Appendix A. In other words, this degree is either 1, 2, 3, 4, 6, 8, 12 or 24. However, each line is defined individually on an extension of degree at most 4. For instance, when the Galois group is the dihedral group D_4 of order 8, the four lines are collectively defined in an extension of degree 8 but each line is defined in an extension of degree 4.

When f has two real and two complex roots, the degree of the extension on which the four lines are defined is again the order of the Galois group of f , but the degree of the extension on which the two real lines are collectively defined is only half the order of the Galois group. This degree is 1, 2, 3, 4 or 12.

Every extension degree can be attained by picking the right polynomial f . To build examples in all cases, it is sufficient to find the equations of two distinct rational cones containing the four points $\mathbf{q}_i = (1, t_i, t_i^2, 0)$ and having the same vertex. Assume f is given by :

$$f = t^4 + \alpha t^3 + \beta t^2 + \gamma t + \delta.$$

Then the following pair (Q_S, Q_T) satisfies the constraints :

$$\begin{cases} xz - y^2 = 0, \\ \delta x^2 + \gamma xy + \beta y^2 + \alpha yz + z^2 = 0. \end{cases}$$

Any two distinct quadrics of the pencil generated by Q_S and Q_T intersect in four (real or imaginary) concurrent lines defined collectively on an extension of \mathbb{Q} of degree equal to the order of the Galois group of f .

By picking the right polynomial, we can generate pairs of quadrics intersecting in four concurrent lines for all types of Galois groups of a quartic. For instance, taking $f = t^4 + t + 1$, we build the two quadrics

$$\begin{cases} xz - y^2 = 0, \\ x^2 + xy + z^2 = 0. \end{cases}$$

The four real concurrent lines of the intersection are defined on an extension of \mathbb{Q} of degree 24, since the Galois group of f is the group S_4 of permutations of four elements (of order 24). Each line is defined in an extension of degree 4.

16.3.3 Two concurrent lines and a double line in $\mathbb{P}^3(\mathbb{C})$, $\sigma_3 = [12]$

In this case, the restricted pencil has a real and rational double root corresponding to a pair of planes Q_R and a rational simple root corresponding to another pair of planes. The second pair is always real. There are two cases.

Double line. The pair of planes Q_R is imaginary. The intersection is reduced to the singular line of this pair, which is clearly rational.

Double line and two simple lines. The double line is rational (otherwise its conjugate would also be in the intersection). The two simple lines are contained in Q_R and go through the common singular point \mathbf{p} of Q_S and Q_T . They are rational if and only if Q_R has a rational point outside its singular line, i.e. if the discriminant of Q_R is a square.

A simple example where this is not the case is as follows :

$$\begin{cases} xy = 0, \\ y^2 - 2z^2 = 0. \end{cases}$$

16.3.4 Two double lines in $\mathbb{P}^3(\mathbb{C})$, $\sigma_3 = [1(11)]$

The characteristic polynomial has a real and rational double root corresponding to a double plane Q_R . The other root corresponds to a (real or imaginary) pair of planes. There are two cases.

Point. The pair of planes is imaginary. The intersection is reduced to the intersection of its singular line with Q_R , i.e. the rational point \mathbf{p} .

Two real double lines. The two double lines are conjugate. They contain the rational point \mathbf{p} and are rational if they go through another rational point. This happens when the discriminant of the pair of planes is a square.

This situation does not necessarily happen, as the following example shows :

$$\begin{cases} z^2 = 0, \\ x^2 - 2y^2 = 0. \end{cases}$$

16.3.5 Line and triple line in $\mathbb{P}^3(\mathbb{C})$, $\sigma_3 = [3]$

The characteristic polynomial has a real and rational triple root corresponding to a real pair of planes. The intersection consists of the (triple) singular line of this pair of planes, which is clearly rational, and of a single line which is also rational, otherwise its conjugate would also be in the intersection. The simple line is found by intersecting one of the planes with any other cone of the pencil.

16.3.6 Quadruple line in $\mathbb{P}^3(\mathbb{C})$, $\sigma_3 = [(21)]$

Here, the characteristic polynomial has a real and rational triple root, corresponding to a double plane. The intersection consists of the (quadruple) line of tangency of this double plane with a cone of the pencil. This line is clearly rational.

16.3.7 Remaining cases

In the remaining cases, that is when the restricted characteristic polynomial identically vanishes, the description of the possible cases given in Section 15.4.2.6 directly yields algorithms for computing parameterizations of the intersection over $\mathbb{Q}[\xi]$.

16.4 Examples

The algorithm described in this paper for computing a near-optimal parameterization of the intersection of two arbitrary quadrics with integer coefficients has been fully implemented in C++. The implementation details as well as an analysis of the complexity (i.e., the height) of the integer coefficients appearing in the parameterizations can be found in [LPP06].

In this section, we illustrate our algorithm with three examples covering different situations. The output given is the actual output of our implementation, with debug information turned on so as to follow what the algorithm is doing.

16.4.1 Example 1

Our first example is as shown in Output 5. As explained in Section 15.5, we first determine the type of the intersection by looking at the multiple roots of the characteristic polynomial and the ranks of the associated quadrics. Here, we find two real double roots corresponding to quadrics of rank 3 : Algorithm 6 tells us that the real type of the intersection is “cubic and secant line”. We have rational roots, so we can parameterize the intersection using cones. One of the two cones of the pencil is

$$Q_R = -2Q_S + Q_T = -25wx - 30wy - 5wz - 20xy - 30w^2 - 5x^2 - 20y^2.$$

This cone has the point $(-2, 1, 4, 0)$ as vertex and contains the vertex of the second cone, i.e. $(-4, 1, 4, 1)$. The line of the intersection is the line joining these two points. Here, we have applied a simple reparameterization to the line by picking two other representative points with smaller “height” than the original two. On the reparameterized line, a very simple point is $(-2, 0, 0, 1)$ which we use as rational point for parameterizing Q_R . Plugging the parameterization of Q_R in the second cone and leaving aside the linear factor corresponding to the line gives the cubic.

16.4.2 Example 2

Our second example is displayed in Output 6. Here, the characteristic polynomial has a double real root at $(\lambda, \mu) = (0, 1)$, the two other simple roots are real (solution of $\mathcal{E}(\lambda, \mu) = 4\mu^2 - \lambda^2 = 0$), the singular quadric $R = R(0, 1)$ is a real pair of planes and $\mathcal{E}(0, 1) > 0$, so Algorithm 3 tells us that the intersection is two secant conics, the singularities of the intersection being convex. Here, the two planes of R are rational :

$$Q_R = (w - y + 2z)(w - 4x + 3y - 2z).$$

Output 5 Execution trace for Example 1.

```

>> quadric 1:  $3x^2 + 5xy + 2xz + 13xw + 14y^2 - 4yz + 4yw + z^2 + 6zw + 14w^2$ 
>> quadric 2:  $x^2 - 10xy + 4xz + xw + 8y^2 - 8yz - 22yw + 2z^2 + 7zw - 2w^2$ 

>> launching intersection
>> characteristic polynomial:  $4l^4 + 12l^3m + l^2m^2 - 12lm^3 + 4m^4$ 
>> gcd of derivatives of characteristic polynomial:  $2l^2 + 3lm - 2m^2$ 
>> ranks of singular quadrics: 3 and 3
>> two real rational double roots: [ 1 2 ] and [ -2 1 ]
>> complex intersection: cubic and secant line
>> real intersection: cubic and secant line
>> reparameterization of line
>> singular point of cone: [ -2 1 4 0 ]
>> rational point on cone: [ -2 0 0 1 ]
>> parameterization of cone with rational point
>> cubic and line intersect at [ -4 1 4 1 ] and [ -2 1 4 0 ]
>> status of intersection param: optimal
>> end of intersection

>> parameterization of cubic:
[ $4u^3 + u^2v + 2v^3, -u^3 - v^3, -4u^3 - u^2v + uv^2 - 4v^3, -u^3$ ]
>> parameterization of line:
[ $-2u, v, 4v, u - v$ ]

>> time spent: 10 ms

```

We can parameterize each of these planes and plug their parameterization in turn in any other quadric of the pencil. This gives the implicit equations of the two conics which we can parameterize. As explained in Section 16.2.5, we are here in one of the few situations where we cannot guarantee that what we output is optimal : the square root in the parameterization of the conics might well be unnecessary. (It turns out that in this particular example it is necessary.) This explains why the implementation reports that the parameterizations of the two conics are only near-optimal.

16.4.3 Example 3

Our last example is presented in Output 7. Here, the characteristic polynomial vanishes identically and all the quadrics of the pencil share a common singular point, with coordinates $(1, 3, -1, -2)$. We apply to all the quadrics of the pencil a projective transformation sending this point to the point $(0, 0, 0, 1)$. The characteristic polynomial of the pencil restricted to the upper left 3×3 part has a double real root at $(-1, 1)$. Corresponding to this double root is a real pair of planes Q_R and Algorithm 2 tells us that the real type of the intersection is “two concurrent lines and a double line”.

As we have seen in Section 16.3.3, the double line of the intersection is the singular line of Q_R . To parameterize the other two lines, we first parameterize Q_R and plug the result in the equation of the other pair of planes of the pencil, corresponding to the second root $(3, 1)$ of the restricted characteristic polynomial. The result follows.

16.5 Conclusion

We have presented in Parts I, II, and III of this paper a new algorithm for computing an exact parametric representation of the intersection of two quadrics in three-dimensional real space given by implicit equations with rational coefficients. We have shown that our algorithm computes projective parameterizations that are optimal in terms of the functions used in the sense that they are polynomials whenever it is possible and contain the square root of some polynomial otherwise. The parameterizations are also near-optimal in the sense that the number of square roots appearing in the coefficients of these functions is minimal except in a small number of cases (characterized by the real type of the intersection) where there may be an extra square root (see Table 16.1 for a summary). Furthermore, we have shown that in the latter cases, testing whether the extra square root is unnecessary and, if so,

Output 6 Execution trace for Example 2.

```

>> quadric 1: 16*x^2 - 12*x*y + 8*x*z + 4*x*w - y^2 - 20*y*z + 2*y*w - 2*z^2 + 3*w^2
>> quadric 2: 4*x*y - 8*x*z - 4*x*w - 3*y^2 + 8*y*z + 2*y*w - 4*z^2 + w^2

>> launching intersection
>> characteristic polynomial: - 1^4 + 4*1^2*m^2
>> gcd of derivatives of characteristic polynomial: 1
>> double real root: [ 0 1 ]
>> inertia: [ 1 1 ]
>> parameterization of pair of planes
>> complex intersection: two secant conics
>> real intersection: two secant conics, convex singularities
>> parameterization of rational conic
>> status of conic 1 param: near-optimal
>> parameterization of rational conic
>> status of conic 2 param: near-optimal
>> end of intersection

>> parameterization of conic 1:
[- 3*u^2 - 9*v^2 - 14*u*v*sqrt(2), - 12*u^2 - 36*v^2 + (- 7*u^2 - 14*u*v + 21*v^2)*sqrt(2),
- 2*u^2 - 6*v^2, - 8*u^2 - 24*v^2 + (- 7*u^2 - 14*u*v + 21*v^2)*sqrt(2)]
>> parameterization of conic 2:
[(u^2 - 28*u*v - 42*v^2)*sqrt(2), - 2*u^2 - 84*v^2 - 28*u*v*sqrt(2), - 8*u^2 - 336*v^2
+ (- 6*u^2 + 252*v^2)*sqrt(2), - 10*u^2 - 420*v^2 + (- 8*u^2 - 28*u*v + 336*v^2)*sqrt(2)]

>> time spent: 10 ms

```

finding an optimal parameterization are equivalent to finding a rational point on a curve or a surface. Hence, leaving for a moment that well-known problem aside, our algorithm closes the problem of finding parameterizations of the intersection that are optimal in the senses discussed above. It should be emphasized that our algorithm is not only theoretically powerful but is also practical : a complete, robust and efficient C++ implementation is described in [LPP06].

For most applications, the near-optimal parameterizations of intersections of quadrics computed by our algorithm are good enough since they are at most one square root away from being optimal. However, there may be situations where one is interested in fully asserting the optimality of a parameterization and, if a given parameterization is not optimal, in obtaining one. As we have seen, this is akin to deciding whether a given curve or surface has a rational point and to computing such a point. The problem of finding integer (or rational) points on an algebraic variety is known to be hard in general, and many instances of the problem are undecidable [Poo01]. When the intersection is a smooth quartic, deciding whether the extra square root can be avoided amounts to finding a rational point on a surface of degree 8 (see Section 14.7) and very little is known about this problem, to the best of our knowledge. The situation is, however, better for the other near-optimal cases, which boil down to finding a rational point on a (possibly non-rational) conic. Indeed, when the conic is rational, Cremona and Rusin [CR03] recently gave an efficient algorithm for solving this problem, which has been implemented in recent releases of the Magma computational algebra system [Mag]. As an example, this implementation solves the problem for an equation of the form $ax^2 + by^2 = cz^2$, where a, b and c are 200-digit primes, in less than 2 seconds on a mainstream PC. In the future, we plan to use this algorithm in our intersection software.

Finally, it should be stressed that the classification, presented in Part II, of pairs of quadrics depending on the type of their real intersection is of independent interest. For instance, it could be used in a collision detection context to predict when collisions between two moving quadrics will occur.

Acknowledgments

We would like to thank Hazel Everett for her useful comments.

Output 7 Execution trace for Example 3.

```

>> quadric 1: 17*x^2 - 12*x*y + 14*x*z - 8*x*w - 4*y^2 - 18*y*w + 5*z^2 + 2*z*w - 16*w^2
>> quadric 2: - 3*x^2 + 28*x*y + 30*x*z + 24*x*w - 4*y^2 + 8*y*z - 2*y*w + 9*z^2 + 18*z*w

>> launching intersection
>> vanishing 4 x 4 characteristic polynomial
>> dimension of common singular locus: 0
>> common singular point of quadrics: [ 1 3 -1 -2 ]
>> computing matrix sending singular point to [ 0 0 0 1 ]
>> 3 x 3 characteristic polynomial: - 1^3 + 1^2*m + 5*1*m^2 + 3*m^3
>> gcd of derivatives of 3 x 3 characteristic polynomial: 1 + m
>> double real root: [ -1 1 ]
>> second root: [ 3 1 ]
>> complex intersection: two concurrent lines and double line
>> real intersection: two concurrent lines and double line
>> reparameterization of line
>> parameterization of pair of lines
>> reparameterization of lines
>> the three lines meet at [ 1 3 -1 -2 ]
>> status of intersection param: optimal
>> end of intersection

>> parameterization of double line:
[v, 3*v, - 2*u - v, u - 2*v]
>> parameterization of line 1:
[- 3*u + v, 3*v, u - v, - 2*u - 2*v]
>> parameterization of line 2:
[u - v, 3*u + 6*v, - u + 5*v, - 2*u]

>> time spent: 10 ms

```

16.6 Appendix A : A primer on Galois theory

Galois theory was introduced in the 19th century for deciding when a polynomial is solvable by radicals. We give here a brief introduction to this theory, which is especially geared towards geometric objects.

Let K be a finite field extension of the field \mathbb{Q} of the rational numbers. Its dimension as \mathbb{Q} -vector space is called the *degree* of K . In our context, K is usually the smallest field containing the coefficients of the equations or of the parameterization of a geometric object such as a point, a line, a curve or a plane. This field K may always be defined as $\mathbb{Q}(\alpha)$, where α is a root of some polynomial f of degree $n = \text{degree}(K)$.

The *splitting field* K' of K is defined as the smallest field containing all the roots of f . It may be proved that it is independent from the choice of f and α . The *Galois group* G of K and K' is the group of the field automorphisms of K' . It is immediate that the elements of G permute the roots of f , and this allows to identify G to a subgroup of the group S_n of all the permutations of the n roots of f .

The important fact for geometric considerations is that any element g of G acts on any object defined from the elements of K by the four field operations $(+, -, *, /)$ simply by replacing any element of K appearing in the definition of the object by its image by the automorphism g , exactly as the complex conjugation acts on any object defined with complex numbers. The different images of an object under the action of the elements of G are called the *conjugates* of this object.

If H is a subgroup of the Galois group G , one may define K'^H , the field of the elements of K' such that $g(x) = x$ for any $g \in H$. The main result of Galois theory is that the map $H \mapsto K'^H$ is a bijection between the subgroups of G and the subfields of K'^H . Moreover, we have $\text{degree}(K'^H) = \text{order}(G)/\text{order}(H)$. It follows that an element of K which has k conjugates lies in a subfield of K of degree k .

This may be extended to the

Galois principle : *Two conjugate objects are isomorphic and any object which has k conjugates (including itself) may be defined on a field of degree k and may not be defined on a smaller field.*

This principle is behind all our proofs of optimality. We describe briefly some other consequences in the context of intersection of quadrics.

One of its first consequences is that any object which has no conjugate except itself may be rationally defined. For example this is the case for the singular point of a singular quartic appearing as the intersection of two rational quadrics. If the intersection of two quadrics is decomposed in a cubic and a line, both may be rationally defined, because their conjugates have to be components of the intersection and the conjugate of a line is a line. Similarly, if the characteristic polynomial of a pencil of quadrics has two double roots and if the corresponding singular quadrics are a cone and a pair of planes, both are rational because they are not isomorphic.

A more involved application of the Galois principle occurs when the intersection of two quadrics consists in four non concurrent lines. As a conjugate of a line in the intersection may only be another line of the intersection, each line may be defined on a field of degree 4. Moreover, if a point lies on a line, any conjugate of the point lies on the corresponding conjugate of the line. Thus the arrangement of the four lines and their four intersection points is preserved by the action of the Galois group. It follows that the Galois group is included in a group of order 8 which acts on the lines and their intersections as the 8 isometries of a square act on its edges and its vertices. By looking at the subgroups of this group, the Galois principle says that the field K of definition of any of the lines has a degree which divides 4, and that if its degree is 4 it has a subfield of degree 2. Therefore each line may be parameterized with at most 2 square roots. As the splitting field has degree at most 8, it is possibly generated by another square root. This shows, without any explicit computation, that at most 3 square roots are needed to define all the four intersection points and the four lines.

16.7 Appendix B : Rational canonical form for the case of four lines forming a skew quadrilateral

We here give a proof of Proposition 16.1.

Proof. Let \mathbb{K} be a \mathbb{Q} -extension field of smallest degree on which the four lines of the intersection are rationally parameterized, and let \mathbb{L} be the field generated by the roots of the characteristic polynomial. To decompose the singular quadrics of the pencil in two planes, we have to extract the square roots of two elements d_1 and d_2 of \mathbb{L} , which are algebraically conjugate on \mathbb{Q} if \mathbb{L} is different from \mathbb{Q} . Let $t^2 - 2\alpha t + \gamma$ be the (rational) polynomial having d_1 and d_2 as roots. It is easy to verify that \mathbb{K} is generated by the roots of the biquadratic polynomial $t^4 - 2\alpha t^2 + \gamma$.

Let the points \mathbf{p}_i , $i = 1, \dots, 4$, be the singular points of the intersection in $\mathbb{P}^3(\mathbb{C})$ of any two distinct quadrics of the pencil. Now, let us choose four points \mathbf{q}_i on which the Galois group G of this polynomial acts in the same way as on the \mathbf{p}_i . For instance, take for \mathbf{q}_i , $i = 1, \dots, 4$, the points of coordinates $(1, t_i, t_i^2, t_i^3)$, where the t_i are the roots of the biquadratic polynomial, numbered such that $t_1 = -t_3$. It is now easy to compute the equations H_j of the planes containing all the points \mathbf{q}_i except \mathbf{q}_j . The quadrics of equations

$$\begin{cases} H_1 H_3 + H_2 H_4 = 0, \\ \sqrt{\alpha^2 - \gamma}(H_1 H_3 - H_2 H_4) = 0 \end{cases}$$

are rational and their intersection consists in the four (not necessarily real) lines $\mathbf{q}_1 \mathbf{q}_2$, $\mathbf{q}_2 \mathbf{q}_3$, $\mathbf{q}_3 \mathbf{q}_4$, and $\mathbf{q}_4 \mathbf{q}_1$. Now, the change of frame sending the \mathbf{q}_i to the points of coordinates $(t_i, -1/t_i, t_i^2 - \alpha, 1)$ is rational and leads to the equations

$$\begin{cases} x^2 - \gamma y^2 - 2wz = 0, \\ \alpha x^2 + 2\gamma xy + \alpha \gamma y^2 - z^2 - (\alpha^2 - \gamma)w^2 = 0. \end{cases}$$

There is a unique projective transformation sending the points \mathbf{p}_i on the \mathbf{q}_i and leaving fixed some rational point which is not on any of the planes defined by the \mathbf{p}_i or by the \mathbf{q}_i . This transformation is invariant under the action of the Galois group G . Thus it is rational, showing the existence of the rational change of frame which is sought.

We now show how the different types of real intersection follow from the signs of $\delta = \alpha^2 - \gamma$, α , and γ . First note that the characteristic polynomial for the rational canonical pencil above is

$$\mathcal{D}(\lambda, \mu) = \gamma(\lambda^2 - \delta\mu^2)^2.$$

Furthermore, the double roots of \mathcal{D} are $(\lambda_0, \mu_0) = (\pm\sqrt{\delta}, 1)$ and, when $\delta > 0$, the associated pairs of planes have equations

$$\lambda_0 Q_S + \mu_0 Q_T = (\alpha \pm \sqrt{\delta}) \left(x + (\alpha \mp \sqrt{\delta})y \right)^2 - (z \pm \sqrt{\delta}w)^2.$$

The discriminants of the two pairs of planes of the pencil are thus $d^\pm = \alpha \pm \sqrt{\delta}$.

We know from Table 15.4 that, when the real intersection consists of two lines, the double roots are complex, hence $\delta < 0$. In the other cases, the double roots are real and distinct thus $\delta > 0$. Note that d^+d^- is then equal to c . When the real intersection is empty, the inertia of the two pairs of planes is $(2, 0)$, hence $d^\pm < 0$; thus $d^+ + d^- = 2\alpha < 0$ and $d^+d^- = \gamma > 0$. When the real intersection consists of two points, the discriminants of the two pairs of planes have opposite signs, thus $d^+d^- = \gamma < 0$. Finally, when the intersection consists of four lines (forming a skew quadrilateral), the discriminants of the two pairs of planes are both positive, thus $d^+ + d^- = 2\alpha > 0$ and $d^+d^- = \gamma > 0$. This completes the proof since these cases for the signs of δ , α , and γ are disjoint. \square

16.8 Appendix C : Examples in all cases

Table 16.2 gives an exhaustive list of examples covering all possible degrees of extension fields on which the components of the intersection are defined, for all real types of intersection. The next-to-last column gives the optimal ring of definition on which a parameterization of the given example is known to exist (ξ is the parameter of the parameterization). When the parameterization output by our algorithm is optimal, the last column gives the degree of the \mathbb{Q} -extension field on which the coefficients of the parameterization of each real component of the intersection is defined. When our algorithm is only near-optimal, the last column gives both the optimal degree and the near-optimal one.

TAB. 16.2: Exhaustive list of examples when the intersection is 0- or 1-dimensional over \mathbb{C} .

complex type	real type	example	field of definition	degree
smooth quartic, $\sigma_4 = [1111]$	\emptyset	$\begin{cases} 6xy+5y^2+2z^2+6zw-w^2=0 \\ 3x^2+y^2-z^2+11w^2=0 \end{cases}$		
	smooth quartic, two affinely finite components	$\begin{cases} x^2+y^2-z^2-w^2=0 \\ xy-2zw=0 \end{cases}$	quartic in $\mathbb{Q}[\xi, \sqrt{\Delta}]$, $\Delta \in \mathbb{Q}[\xi]$	1 / 2
		$\begin{cases} 6xy+5y^2+2z^2+6zw-w^2=0 \\ 3x^2+y^2-z^2-w^2=0 \end{cases}$	quartic in $\mathbb{Q}(\sqrt{\delta})[\xi, \sqrt{\Delta}]$, $\Delta \in \mathbb{Q}(\sqrt{\delta})[\xi]$	2
	smooth quartic, one affinely finite component	$\begin{cases} x^2+y^2+2zw=0 \\ x^2+z^2+zw-w^2=0 \end{cases}$	quartic in $\mathbb{Q}[\xi, \sqrt{\Delta}]$	1 / 2
		$\begin{cases} 2x^2-2xy+2xz-2xw+y^2 \\ +4yz-4yw+2z^2-4zw=0 \\ x^2-2xy+4xz+4xw-y^2 \\ +2yz+4yw+4zw-2w^2=0 \end{cases}$	quartic in $\mathbb{Q}(\sqrt{\delta})[\xi, \sqrt{\Delta}]$, $\Delta \in \mathbb{Q}(\sqrt{\delta})[\xi]$	2
	smooth quartic, two affinely infinite components	$\begin{cases} xy+zw=0 \\ x^2-y^2+z^2+2zw-w^2=0 \end{cases}$	quartic in $\mathbb{Q}[\xi, \sqrt{\Delta}]$, $\Delta \in \mathbb{Q}[\xi]$	1 / 2
		$\begin{cases} x^2-2y^2+4zw=0 \\ xy+z^2+2zw-w^2=0 \end{cases}$	quartic in $\mathbb{Q}(\sqrt{\delta})[\xi, \sqrt{\Delta}]$, $\Delta \in \mathbb{Q}(\sqrt{\delta})[\xi]$	2
	nodal quartic, $\sigma_4 = [112]$	point	$\begin{cases} y^2+z^2+w^2=0 \\ xy+w^2=0 \end{cases}$	point in \mathbb{Q}
nodal quartic, affinely finite		$\begin{cases} y^2+z^2-w^2=0 \\ xy+w^2=0 \end{cases}$	quartic in $\mathbb{Q}[\xi]$	1 / 2

TAB. 16.2: (continued)

complex type	real type	example	field of definition	degree
		$\begin{cases} y^2+z^2-3w^2=0 \\ xy+w^2=0 \end{cases}$	quartic in $\mathbb{Q}(\sqrt{3})[\xi]$	2
	nodal quartic, affinely infinite	$\begin{cases} y^2+z^2-w^2=0 \\ xy+zw=0 \end{cases}$	quartic in $\mathbb{Q}[\xi]$	1 / 2
		$\begin{cases} y^2+z^2-3w^2=0 \\ xy+zw=0 \end{cases}$	quartic in $\mathbb{Q}(\sqrt{3})[\xi]$	2
	nodal quartic with isolated singularity	$\begin{cases} y^2-z^2-w^2=0 \\ xy+w^2=0 \end{cases}$	quartic in $\mathbb{Q}[\xi]$, point in \mathbb{Q}	1 / 2 1
		$\begin{cases} y^2+z^2-3w^2=0 \\ xw+y^2=0 \end{cases}$	quartic in $\mathbb{Q}(\sqrt{3})[\xi]$, point in \mathbb{Q}	2 1
cuspidal quartic, $\sigma_4 = [13]$	cuspidal quartic	$\begin{cases} yz+w^2=0 \\ xz+y^2=0 \end{cases}$	quartic in $\mathbb{Q}[\xi]$	1
cubic and secant line, $\sigma_4 = [22]$	cubic and secant line	$\begin{cases} y^2+zw=0 \\ xy+w^2=0 \end{cases}$	cubic in $\mathbb{Q}[\xi]$, line in $\mathbb{Q}[\xi]$	1 1
	cubic and non-secant line	$\begin{cases} xw+yz=0 \\ xz-yw+zw=0 \end{cases}$	cubic in $\mathbb{Q}[\xi]$, line in $\mathbb{Q}[\xi]$	1 1
cubic and tangent line, $\sigma_4 = [4]$	cubic and tangent line	$\begin{cases} yw+z^2=0 \\ xw+yz=0 \end{cases}$	cubic in $\mathbb{Q}[\xi]$, line in $\mathbb{Q}[\xi]$	1 1
two secant conics, $\sigma_4 = [11(11)]$	\emptyset	$\begin{cases} z^2-w^2=0 \\ x^2+y^2+w^2=0 \end{cases}$		
	two points	$\begin{cases} z^2+w^2=0 \\ x^2-y^2+w^2=0 \end{cases}$	two points in \mathbb{Q}	1
		$\begin{cases} z^2+w^2=0 \\ x^2-2y^2+w^2=0 \end{cases}$	two points in $\mathbb{Q}(\sqrt{2})$	2
	one conic	$\begin{cases} zw=0 \\ x^2+y^2+z^2-w^2=0 \end{cases}$	one conic in $\mathbb{Q}[\xi]$	1 / 2
		$\begin{cases} x^2-4xw+3w^2=0 \\ x^2+y^2+z^2-4w^2=0 \end{cases}$	one conic in $\mathbb{Q}(\sqrt{3})[\xi]$	2
		$\begin{cases} x^2-4xw-3w^2=0 \\ x^2+y^2+z^2-w^2=0 \end{cases}$	one conic in $\mathbb{Q}(\sqrt{4\sqrt{7}-10})$	4
	two non-secant conics	$\begin{cases} x^2-w^2=0 \\ y^2+z^2-w^2=0 \end{cases}$	two conics in $\mathbb{Q}[\xi]$	1 / 2
		$\begin{cases} x^2-w^2=0 \\ y^2+z^2-3w^2=0 \end{cases}$	two conics in $\mathbb{Q}(\sqrt{3})[\xi]$	2
		$\begin{cases} x^2-3w^2=0 \\ y^2+z^2-3w^2=0 \end{cases}$	two conics in $\mathbb{Q}(\sqrt{3})[\xi]$	2 / 4
		$\begin{cases} x^2-33w^2=0 \\ y^2+z^2-3w^2=0 \end{cases}$	two conics in $\mathbb{Q}(\sqrt{3}, \sqrt{11})[\xi]$	4

TAB. 16.2: (continued)

complex type	real type	example	field of definition	degree
	two secant conics, affinely finite	$\begin{cases} x^2 - y^2 = 0 \\ y^2 + z^2 - w^2 = 0 \end{cases}$	two conics in $\mathbb{Q}[\xi]$	1 / 2
		$\begin{cases} x^2 - y^2 = 0 \\ y^2 + z^2 - 3w^2 = 0 \end{cases}$	two conics in $\mathbb{Q}(\sqrt{3})[\xi]$	2
		$\begin{cases} x^2 - 3y^2 = 0 \\ y^2 + z^2 - 3w^2 = 0 \end{cases}$	two conics in $\mathbb{Q}(\sqrt{3})[\xi]$	2 / 4
		$\begin{cases} x^2 - 33y^2 = 0 \\ y^2 + z^2 - 3w^2 = 0 \end{cases}$	two conics in $\mathbb{Q}(\sqrt{3}, \sqrt{11})[\xi]$	4
	two secant conics, affinely infinite	$\begin{cases} z^2 - w^2 = 0 \\ x^2 - 2y^2 - zw = 0 \end{cases}$	two conics in $\mathbb{Q}[\xi]$	1 / 2
		$\begin{cases} zw = 0 \\ x^2 - 3y^2 + z^2 - 11w^2 = 0 \end{cases}$	two conics in $\mathbb{Q}(\sqrt{3})[\xi]$	2
		$\begin{cases} z^2 - 2w^2 = 0 \\ x^2 - 3y^2 - zw = 0 \end{cases}$	two conics in $\mathbb{Q}(\sqrt{2})[\xi]$	2 / 4
		$\begin{cases} z^2 - 4zw - 3w^2 = 0 \\ x^2 - 3y^2 + z^2 - 11w^2 = 0 \end{cases}$	two conics in $\mathbb{Q}(\sqrt{3}, \sqrt{7})[\xi]$	4
two tangent conics, $\sigma_4 = [1(21)]$	point	$\begin{cases} x^2 + w^2 = 0 \\ xy + z^2 = 0 \end{cases}$	point in \mathbb{Q}	1
	two conics	$\begin{cases} x^2 - w^2 = 0 \\ xy + z^2 = 0 \end{cases}$	two conics in $\mathbb{Q}[\xi]$	1
		$\begin{cases} x^2 - 2w^2 = 0 \\ xy + z^2 = 0 \end{cases}$	two conics in $\mathbb{Q}(\sqrt{2})[\xi]$	2
double conic, $\sigma_4 = [1(111)]$	\emptyset	$\begin{cases} w^2 = 0 \\ x^2 + y^2 + z^2 = 0 \end{cases}$		
	double conic	$\begin{cases} x^2 = 0 \\ y^2 + z^2 - w^2 = 0 \end{cases}$	conic in $\mathbb{Q}[\xi]$	1 / 2
		$\begin{cases} x^2 = 0 \\ y^2 + z^2 - 3w^2 = 0 \end{cases}$	conic in $\mathbb{Q}(\sqrt{3})[\xi]$	2
conic and two lines not crossing, $\sigma_4 = [2(11)]$	point	$\begin{cases} xy = 0 \\ y^2 + z^2 + w^2 = 0 \end{cases}$	point in \mathbb{Q}	1
	conic and point	$\begin{cases} xw = 0 \\ y^2 + z^2 - w^2 = 0 \end{cases}$	point in \mathbb{Q} , conic in $\mathbb{Q}[\xi]$	1 / 2
		$\begin{cases} xw = 0 \\ y^2 + z^2 - 3w^2 = 0 \end{cases}$	point in \mathbb{Q} , conic in $\mathbb{Q}(\sqrt{3})[\xi]$	1 2
	conic and two lines	$\begin{cases} xy = 0 \\ y^2 + z^2 - w^2 = 0 \end{cases}$	two lines in $\mathbb{Q}[\xi]$, conic in $\mathbb{Q}[\xi]$	1 1
		$\begin{cases} xy = 0 \\ 2y^2 + z^2 - 3w^2 = 0 \end{cases}$	two lines in $\mathbb{Q}(\sqrt{3})[\xi]$, conic in $\mathbb{Q}[\xi]$	2 1 / 2
		$\begin{cases} xy = 0 \\ y^2 + z^2 - 3w^2 = 0 \end{cases}$	two lines in $\mathbb{Q}(\sqrt{3})[\xi]$, conic in $\mathbb{Q}(\sqrt{3})[\xi]$	2 2

TAB. 16.2: (continued)

complex type	real type	example	field of definition	degree
conic and two lines crossing, $\sigma_4 = [(31)]$	conic	$\begin{cases} yz=0 \\ xz+y^2+w^2=0 \end{cases}$	conic in $\mathbb{Q}[\xi]$	1
	conic and two lines	$\begin{cases} yz=0 \\ xz+y^2-w^2=0 \end{cases}$	conic in $\mathbb{Q}[\xi]$, two lines in $\mathbb{Q}[\xi]$	1 1
		$\begin{cases} yz=0 \\ xz+y^2-2w^2=0 \end{cases}$	conic in $\mathbb{Q}[\xi]$, two lines in $\mathbb{Q}(\sqrt{2})[\xi]$	1 2
four lines (skew quadrilateral), $\sigma_4 = [(11)(11)]$	\emptyset	$\begin{cases} x^2+y^2=0 \\ z^2+w^2=0 \end{cases}$		
	two points	$\begin{cases} x^2+y^2-2zw=0 \\ 2xy+z^2+w^2=0 \end{cases}$	two points in \mathbb{Q}	1
		$\begin{cases} x^2+3y^2-2zw=0 \\ x^2-6xy-3y^2-z^2-4w^2=0 \end{cases}$	two points in $\mathbb{Q}(\sqrt{3})$	2
		$\begin{cases} x^2+2y^2-2zw=0 \\ x^2-4xy-2y^2-z^2-3w^2=0 \end{cases}$	two points in $\mathbb{Q}(\sqrt{1+\sqrt{3}})$	4
	two skew lines	$\begin{cases} x^2-25y^2-2zw=0 \\ 3x^2+50xy+75y^2-z^2+16w^2=0 \end{cases}$	two lines in $\mathbb{Q}[\xi]$	1
		$\begin{cases} x^2-4y^2-2zw=0 \\ x^2+8xy+4y^2-z^2+3w^2=0 \end{cases}$	two lines in $\mathbb{Q}(\sqrt{6})[\xi]$	2
		$\begin{cases} x^2-3y^2-2zw=0 \\ x^2+6xy+3y^2-z^2+2w^2=0 \end{cases}$	two lines in $\mathbb{Q}(\sqrt{\sqrt{3}-1})[\xi]$	4
	four lines (skew quadrilateral)	$\begin{cases} x^2-9y^2-2zw=0 \\ 5x^2+18xy+45y^2-z^2-16w^2=0 \end{cases}$	four lines in $\mathbb{Q}[\xi]$	1
		$\begin{cases} x^2-y^2-2zw=0 \\ 3x^2+2xy+3y^2-z^2-8w^2=0 \end{cases}$	two lines in $\mathbb{Q}[\xi]$, two lines in $\mathbb{Q}(\sqrt{2})[\xi]$	1 2
		$\begin{cases} x^2-16y^2-2zw=0 \\ 5x^2+32xy+80y^2-z^2-9w^2=0 \end{cases}$	four lines in $\mathbb{Q}(\sqrt{2})[\xi]$	2
		$\begin{cases} x^2-y^2-2zw=0 \\ 2x^2+2xy+2y^2-z^2-3w^2=0 \end{cases}$	two lines in $\mathbb{Q}(\sqrt{2})[\xi]$, two lines in $\mathbb{Q}(\sqrt{6})[\xi]$	2 2
		$\begin{cases} x^2-20y^2-2zw=0 \\ 6x^2+40xy+120y^2-z^2-16w^2=0 \end{cases}$	four lines in $\mathbb{Q}(\sqrt{2}, \sqrt{10})[\xi]$	4
		$\begin{cases} x^2-2y^2-2zw=0 \\ 2x^2+4xy+4y^2-z^2-2w^2=0 \end{cases}$	four lines in $\mathbb{Q}(\sqrt{2+\sqrt{2}})[\xi]$	4
		$\begin{cases} x^2-3y^2-2zw=0 \\ 3x^2+6xy+9y^2-z^2-6w^2=0 \end{cases}$	two lines in $\mathbb{Q}(\sqrt{3-\sqrt{3}})[\xi]$, two lines in $\mathbb{Q}(\sqrt{3+\sqrt{3}})[\xi]$	4 4
two skew lines and a double line, $\sigma_4 = [(22)]$	double line	$\begin{cases} y^2+w^2=0 \\ xy+zw=0 \end{cases}$	double line in $\mathbb{Q}[\xi]$	1
	two skew lines and a double line	$\begin{cases} y^2-w^2=0 \\ xy-zw=0 \end{cases}$	double line in $\mathbb{Q}[\xi]$ two simple lines in $\mathbb{Q}[\xi]$	1 1
		$\begin{cases} y^2-2w^2=0 \\ xy-zw=0 \end{cases}$	double line in $\mathbb{Q}[\xi]$ two simple lines in $\mathbb{Q}(\sqrt{2})[\xi]$	1 2

TAB. 16.2: (continued)

complex type	real type	example	field of definition	degree
two concurrent double lines, $\sigma_4 = [(211)]$	point	$\begin{cases} w^2=0 \\ x^2+y^2+zw=0 \end{cases}$	point in \mathbb{Q}	1
	two double lines	$\begin{cases} w^2=0 \\ x^2-y^2+zw=0 \end{cases}$	two lines in $\mathbb{Q}[\xi]$	1
		$\begin{cases} w^2=0 \\ x^2-2y^2+zw=0 \end{cases}$	two lines in $\mathbb{Q}(\sqrt{2})[\xi]$	2
conic and double line, $\sigma_4 = [1\{3\}]$	conic and double line	$\begin{cases} xw=0 \\ xz+y^2=0 \end{cases}$	conic in $\mathbb{Q}[\xi]$, line in $\mathbb{Q}[\xi]$	1 1
four concurrent lines, $\sigma_3 = [111]$	point	$\begin{cases} x^2+z^2=0 \\ y^2+z^2=0 \end{cases}$	point in \mathbb{Q}	1
	two concurrent lines	$\begin{cases} xz-y^2=0 \\ -x^2+z^2=0 \end{cases}$	two lines in $\mathbb{Q}[\xi]$	1
		$\begin{cases} xz-y^2=0 \\ -2x^2-y^2+z^2=0 \end{cases}$	two lines in $\mathbb{Q}(\sqrt{2})[\xi]$	2
		$\begin{cases} xz-y^2=0 \\ 2xy+z^2=0 \end{cases}$	one line in $\mathbb{Q}[\xi]$, one line in $\mathbb{K}[\xi]$, $\text{degree}(\mathbb{K}) = 3$	1 3
		$\begin{cases} xz-y^2=0 \\ -3x^2+z^2=0 \end{cases}$	two lines in $\mathbb{Q}(\sqrt{3}\sqrt{3})[\xi]$	4
		$\begin{cases} xz-y^2=0 \\ -3x^2-3xy+z^2=0 \end{cases}$	one line in $\mathbb{K}[\xi]$, $\text{degree}(\mathbb{K}) = 4$ one line in $\mathbb{K}'[\xi]$, $\text{degree}(\mathbb{K}') = 4$	4 4
		four concurrent lines	$\begin{cases} xz-y^2=0 \\ 4x^2-5y^2+z^2=0 \end{cases}$	four lines in $\mathbb{Q}[\xi]$
	$\begin{cases} xz-y^2=0 \\ 2x^2-3y^2+z^2=0 \end{cases}$		two lines in $\mathbb{Q}[\xi]$, two lines in $\mathbb{Q}(\sqrt{2})[\xi]$	1 2
	$\begin{cases} xz-y^2=0 \\ -4x^2+8xy-4yz+z^2=0 \end{cases}$		four lines in $\mathbb{Q}(\sqrt{2})[\xi]$	2
	$\begin{cases} xz-y^2=0 \\ xy-3y^2+z^2=0 \end{cases}$		one line in $\mathbb{Q}[\xi]$, three lines l_i in $\mathbb{K}_i[\xi]$, $\text{degree}(\mathbb{K}_i) = 3$	1 3
	$\begin{cases} xz-y^2=0 \\ 2x^2-4y^2+z^2=0 \end{cases}$		four lines in $\mathbb{Q}(\sqrt{2+\sqrt{2}})[\xi]$	4
	$\begin{cases} xz-y^2=0 \\ 4x^2-10y^2+z^2=0 \end{cases}$		four lines in $\mathbb{Q}(\sqrt{6}, \sqrt{14})[\xi]$	4
	$\begin{cases} xz-y^2=0 \\ 2x^2+10xy+15y^2+7yz+z^2=0 \end{cases}$		two lines in $\mathbb{Q}(\sqrt{2})[\xi]$, two lines in $\mathbb{Q}(\sqrt{5})[\xi]$	2 2
	$\begin{cases} xz-y^2=0 \\ xy-4y^2+z^2=0 \end{cases}$		one line in $\mathbb{Q}[\xi]$, three lines l_i in $\mathbb{K}_i[\xi]$, $\text{degree}(\mathbb{K}) = 3$	1 3
	$\begin{cases} xz-y^2=0 \\ 2x^2-5y^2+z^2=0 \end{cases}$		two lines in $\mathbb{Q}(\sqrt{5+\sqrt{17}})[\xi]$, two lines in $\mathbb{Q}(\sqrt{5-\sqrt{17}})[\xi]$	4 4
	$\begin{cases} xz-y^2=0 \\ x^2-3xy-7y^2+z^2=0 \end{cases}$		four lines l_i in $\mathbb{K}_i[\xi]$, $\text{degree}(\mathbb{K}_i) = 4$	4
	$\begin{cases} xz-y^2=0 \\ x^2+xy+z^2=0 \end{cases}$		four lines l_i in $\mathbb{K}_i[\xi]$, $\text{degree}(\mathbb{K}_i) = 4$	4

TAB. 16.2: (continued)

complex type	real type	example	field of definition	degree
two concurrent lines and a double line, $\sigma_3 = [12]$	double line	$\begin{cases} xy=0 \\ y^2+z^2=0 \end{cases}$	double line in $\mathbb{Q}[\xi]$	1
	two concurrent lines and a double line	$\begin{cases} xy=0 \\ y^2-z^2=0 \end{cases}$	double line in $\mathbb{Q}[\xi]$, two simple lines in $\mathbb{Q}[\xi]$	1 1
		$\begin{cases} xy=0 \\ y^2-2z^2=0 \end{cases}$	double line in $\mathbb{Q}[\xi]$, two simple lines in $\mathbb{Q}(\sqrt{2})[\xi]$	1 2
line and triple line, $\sigma_3 = [3]$	line and triple line	$\begin{cases} xz+y^2=0 \\ yz=0 \end{cases}$	simple line in $\mathbb{Q}[\xi]$, triple line in $\mathbb{Q}[\xi]$	1 1
two concurrent double lines, $\sigma_3 = [1(11)]$	point	$\begin{cases} z^2=0 \\ x^2+y^2=0 \end{cases}$	point in \mathbb{Q}	1
	two double lines	$\begin{cases} z^2=0 \\ x^2-y^2=0 \end{cases}$	two lines in $\mathbb{Q}[\xi]$	1
		$\begin{cases} z^2=0 \\ x^2-2y^2=0 \end{cases}$	two lines in $\mathbb{Q}(\sqrt{2})[\xi]$	2
quadruple line, $\sigma_3 = [(21)]$	quadruple line	$\begin{cases} y^2=0 \\ xy+z^2=0 \end{cases}$	line in $\mathbb{Q}[\xi]$	1
quadruple line, $\sigma_2 = [11]$	quadruple line	$\begin{cases} x^2=0 \\ y^2=0 \end{cases}$	line in $\mathbb{Q}[\xi]$	1

Chapitre 17

Near-optimal parameterization of the intersection of quadrics : IV. An efficient and exact implementation

Cet article est paru dans *Computational Geometry : Theory and Applications* [LPP06], une “special issue” du *20th ACM Annual Symposium on Computational Geometry* [LPP04].

Abstract

We present the first complete, exact, and efficient C++ implementation for parameterizing the intersection of two implicit quadrics with integer coefficients of arbitrary size. It is based on the near-optimal algorithm recently introduced by Dupont et al. [DLLP03] and builds upon Levin’s seminal work [Lev76].

Unlike existing implementations, it correctly identifies and parameterizes all the connected components of the intersection in all cases, returning parameterizations with rational functions whenever such parameterizations exist. In addition, the coefficient rings of the parameterizations are either minimal or involve one possibly unneeded square root.

We prove upper bounds on the size of the coefficients of the output parameterizations and compare these bounds to observed values. We give other experimental results and present some examples.

17.1 Introduction

Computing an explicit representation of the intersection of two general quadrics (i.e., quadratic surfaces) is a fundamental problem in areas such as solid modeling, computational geometry, and computer graphics. The range of applications covers well-known problems like computing arrangements [MTT03, SW06], boundary evaluation [Sar83], and convex hull computation [HI95].

Past work. Until recently, the only known general method for computing a parametric representation of the intersection between two arbitrary quadrics was that of J. Levin [Lev76]. This method is based on an analysis of the pencil generated by the two quadrics, i.e., their set of linear combinations.

Though useful for curve tracing, Levin’s method has serious limitations. When the intersection is singular or reducible, a parameterization by rational functions is known to exist, but Levin’s pencil method fails to find it and generates a parameterization that involves the square root of some polynomial. In addition, since it introduces algebraic numbers of very high degree (for instance in the computation of eigenvalues and eigenvectors), a correct implementation using exact arithmetic is essentially out of reach. In addition, when a floating point representation of numbers is used, the method may output results that are wrong (geometrically and topologically) and it may even fail to produce any parameterization at all and crash.

Over the years, Levin's seminal work has been extended and refined in several different directions. Wilf and Manor [WM93] use a classification of quadric intersections by the Segre characteristic (see [Bro06]) to drive the parameterization of the intersection by the pencil method. Recently, Wang, Goldman, and Tu [WGT03] further improved the method making it capable of computing structural information on the intersection and its various connected components and able to produce a parameterization by rational functions when such a parameterization exists. Whether the refined algorithm is numerically robust is open to question.

Another method of algebraic flavor was introduced by Farouki, Neff, and O'Connor [FNO89] for parameterizing the intersection in degenerate situations. In such cases, using a combination of classical concepts (Segre characteristic) and algebraic tools (factorization of multivariate polynomials), the authors show that explicit information on the morphological type of the intersection curve can be reliably obtained. A notable feature of this method is that it can output an exact parameterization of the intersection in simple cases, when the input quadrics have rational coefficients. No implementation is reported however.

Rather than restricting the type of the intersection, others have sought to restrict the type of the input quadrics, taking advantage of the fact that geometric insights can then help compute the intersection curve [MG95, SJ94]. Specialized routines are devised to compute the intersection curve in each particular case. Such geometric approaches are however essentially limited to the class of so-called natural quadrics, i.e., the planes, right cones, circular cylinders, and spheres.

Apart from [DLLP03], perhaps the most interesting of the known algorithms for computing an explicit representation of the intersection of two arbitrary quadrics is the method of Wang, Joe, and Goldman [WJG02]. This algebraic method is based on a birational mapping between the intersection curve and a plane cubic curve. The cubic curve is obtained by projection from a point lying on the intersection. Then the classification and parameterization of the intersection are obtained by invoking classical results on plane cubics. The authors claim that their algorithm is the first to produce a complete topological classification of the intersection (singularities, number, and types of connected components, etc.). Numerical robustness issues have however not been studied and the intersection may not be correctly classified. Also, the center of projection is currently computed using Levin's (enhanced) method : with floating point arithmetic, the center of projection will in general not exactly lie on the curve, which is another source of numerical instability.

Contributions. In this paper, we present the first complete, exact, and efficient implementation of an algorithm for parameterizing the intersection of two arbitrary quadrics, given in implicit form, with integer coefficients. (Note that quadrics with rational or finite floating-point coefficients can be trivially converted to integer form.) This implementation is based on the parameterization method described in [DLLP03], itself built upon Levin's pencil method.

Precisely, our implementation has the following features :

- it computes an exact parameterization of the intersection of two quadrics with integer coefficients of arbitrary size ;
- it places no restriction of any kind on the type of the intersection or the type of the input quadrics ;
- it correctly identifies, separates, and parameterizes all the connected components of the intersection and gives all the relevant topological information ;
- the parameterization is rational when one exists ; otherwise the intersection is a smooth quartic and the parameterization involves the square root of a polynomial ;
- the parameterization is either optimal in the degree of the extension of \mathbb{Z} on which its coefficients are defined or, in a small number of well-identified cases, involves one extra possibly unnecessary square root ;
- the implementation is carefully designed so that the size of the coefficients is kept small ;
- it is fast and efficient, and can routinely compute parameterizations of the intersection of quadrics with input coefficients having ten digits in less than 50 milliseconds on a mainstream PC.

Our code can be downloaded from the LORIA and INRIA web sites²⁷. The C++ implementation can also be queried via a web interface.

The paper is organized as follows. After some preliminaries, we recall in Section 17.3 the main ideas of the

²⁷<http://www.loria.fr>, <http://www.inria.fr>

parameterization algorithm we introduced in [DLLP03]. In Section 17.4, we prove theoretical bounds on the size of the output coefficients when the intersection is generic and compare those bounds to observed values. A similar work is carried out in Section 17.5 for singular intersections and the results are used to validate a key design choice we made in our implementation. After describing our implementation (Section 17.6), we then give experimental results and performance evaluation in Section 17.7, both on random and real data. Finally, we show the output produced by our implementation for some examples in Section 17.8, before concluding.

17.2 Preliminaries

In what follows, all the matrices considered are 4×4 real matrices, unless otherwise specified. We call a *quadric* associated with a symmetric matrix S the set

$$Q_S = \{\mathbf{x} \in \mathbb{P}^3 \mid \mathbf{x}^T S \mathbf{x} = 0\},$$

where $\mathbb{P}^3 = \mathbb{P}^3(\mathbb{R})$ denotes the real projective space of dimension 3 ($\mathbf{x}^T S \mathbf{x}$ is quadratic and homogeneous in the coordinates of \mathbf{x}). In the rest of this paper, points and parameterizations are assumed to live in projective space. Recall that a point of \mathbb{P}^3 has four coordinates.

We define the *inertia* of S and Q_S as the pair

$$\sigma_S = (\max(\sigma^+, \sigma^-), \min(\sigma^+, \sigma^-)),$$

where σ^+ (resp. σ^-) is the number of positive (resp. negative) eigenvalues of S . The *rank* of S is the sum $\sigma^+ + \sigma^-$. Recall that Sylvester's Inertia Law asserts that the inertia of S (and thus the rank) is invariant by a real projective transformation [Lam73].

We call *projective cones* (or simply *cones*) the quadrics of rank 3 and *pairs of planes* the quadrics of rank 2. For the benefit of the reader, we recall that, in affine real space, quadrics of inertia $(4,0)$ are empty, quadrics of inertia $(3,1)$ are ellipsoids, hyperboloids of two sheets, or elliptic paraboloids, and quadrics of inertia $(2,2)$ are hyperboloids of one sheet or hyperbolic paraboloids (see [DLLP03] for a complete characterization of affine quadrics). Also, quadrics of inertia $(2,1)$ are cones or cylinders. All the quadric *surfaces* except those of inertia $(3,1)$ are ruled surfaces, i.e., surfaces that are swept by a one-dimensional family of lines.

Given two matrices S and T , let $R(\lambda, \mu) = \lambda S + \mu T$. The set $\{R(\lambda, \mu) \mid (\lambda, \mu) \in \mathbb{P}^1(\mathbb{R})\}$ is called the *pencil* of matrices generated by S and T . For the sake of simplicity, we sometimes write a member of the pencil $R(\lambda) = \lambda S - T$, $\lambda \in \mathbb{R} \cup \{\infty\}$. Associated to a pencil of matrices is a pencil of quadrics $\{Q_{R(\lambda, \mu)} \mid (\lambda, \mu) \in \mathbb{P}^1\}$. Recall the classical result that the intersection of two distinct quadrics of a pencil is independent of the choice of the two quadrics.

The equation $\det R(\lambda, \mu) = 0$ is called the *determinantal equation* of the pencil. The *singular* quadrics of the pencil are exactly the quadrics $Q_{R(\lambda, \mu)}$ such that $\det R(\lambda, \mu) = 0$. Note that a quadric of the pencil is singular if and only if it has rank less than or equal to 3.

17.3 Algorithm description

In this section, we give a brief presentation of the basic ideas underpinning our near-optimal parameterization method [DLLP03].

From now on, S and T are two symmetric 4×4 matrices with entries in \mathbb{Z} . By abuse of language, we will often talk about (and manipulate) objects with rational coefficients, with the understanding that, in projective space, such coefficients can trivially be converted to integers.

17.3.1 Near-optimal parameterization algorithm

Let $\{Q_{R(\lambda, \mu)} \mid (\lambda, \mu) \in \mathbb{P}^1\}$, with $R(\lambda, \mu) = \lambda S + \mu T$, be a pencil of quadrics. The main idea of existing methods for parameterizing the intersection of two quadrics based on an analysis of their pencil (Levin's and derivatives) is as follows : find a quadric Q_R of some particularly simple form in the pencil generated by Q_S and Q_T (assume

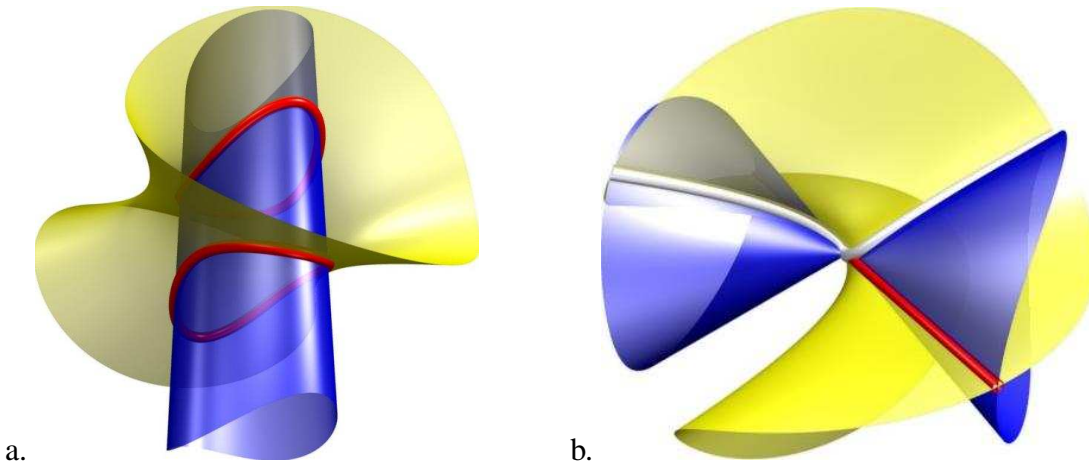


FIG. 17.1 – Examples of intersections (rendered with Surf [Sur]). a. Smooth quartic, with two affinely finite components. b. Cubic and tangent line.

$Q_R \neq Q_S$), parameterize this quadric, plug the parameterization \mathbf{X} of Q_R in the equation of Q_S , solve the resulting equation $\mathbf{X}^T \mathbf{S} \mathbf{X} = 0$, and plug the result in \mathbf{X} , finally giving the parameterization of the intersection.

The key to making this procedure work in practice is to find a quadric Q_R that is ruled and thus admits a parameterization that is linear in one of its parameters so that the equation $\mathbf{X}^T \mathbf{S} \mathbf{X} = 0$ has degree 2. Levin's main result was to prove that a pencil of quadrics always contains at least one "simple" ruled quadric [Lev76]. Furthermore, Levin showed how to compute such a quadric by first finding the zeros of the determinant of the upper left 3×3 submatrix of $R(\lambda, \mu)$, a cubic equation. Since cubic equations have generically no rational root (by Hilbert's Irreducibility Theorem), Levin's algorithm introduces non-rational numbers at an early stage and, in practice, floating-point arithmetic has to be used, resulting in numerical robustness problems.

The principal contribution of [DLLP03] was to show that, by a careful choice of the intermediate quadric Q_R , the appearance of algebraic numbers can be kept to a minimum. One major result is encapsulated in Theorem 3 of [DLLP03] : except when the intersection is reduced to two real points, the pencil contains at least one ruled quadric whose coefficients are rational and such a quadric can be easily computed. In addition, thanks to new worst-case optimal (in the number of square roots) parameterizations of ruled projective quadrics, we can always find such a rational ruled quadric Q_R with a parameterization involving only one square root.

Some of the basic ingredients used in our algorithm to infer information about the intersection are the Segre classification of pencils and its refinement over the reals (the Canonical Form Theorem for pairs of real symmetric matrices – see [Uhl73]), a projective setting, ad hoc projective transformations to compute the canonical form of a projective quadric, and Sylvester's Inertia Law [Lam73].

The basic principles underlying the design of our implementation are as follows :

- compute structural information on the intersection and its various real components as early as possible ;
- use the structural information gathered to drive the parameterization process and make the right choices so that the output is optimal or near-optimal from the point of view of the degree of the extension of \mathbb{Z} on which its coefficients are defined.

In our implementation we were interested not just in optimizing the number of square roots in the output but also in minimizing the size of the output coefficients. For this reason, the basic philosophy is to use as intermediate ruled quadric Q_R a quadric with rational coefficients of the smallest rank that we can easily find, the rationale being, for instance, that the parameterization of a cone involves coefficients of smaller asymptotic size than the coefficients of the parameterization of a quadric of inertia $(2, 2)$. There are essentially two cases : (i) Q_R has rank 4 ; (ii) Q_R has rank 3 or less.

Case (i) : Q_R has rank 4

The main case where Q_R has rank 4 is when the intersection is a smooth quartic (Figure 17.1.a). In this situation, the quartic determinantal equation $\det R(\lambda) = 0$ has no multiple root. It could well be that at least one

of its simple roots is rational and that a Q_R with rank less than 4 could have been used, but checking this via the Rational Root Theorem can be very time consuming²⁸. Since generically a degree-four equation has no rational root, we prefer instead to isolate the real zeros of the determinantal equation using an implementation of Uspensky's algorithm [RZ04]. We then take (at most two) rational test points λ_i outside the isolating intervals in the areas where $\det R(\lambda) > 0$. If one of the quadrics $R(\lambda_i)$ has inertia $(4, 0)$, the intersection is empty (it is a complex smooth quartic), a consequence of Finsler's Theorem (see [DLLP03]). Otherwise, we proceed.

We now have a quadric $R_0 = R(\lambda_0)$ of inertia $(2, 2)$ and a range of values $I = [a, b]$ such that $\lambda_0 \in I$ and $\det R(\lambda) > 0$ for all $\lambda \in I$. In the worst case, the parameterization of Q_R involves two square roots [DLLP03]. We can improve this situation as follows. First, compute a point \mathbf{p}_0 on Q_{R_0} . Approximate this point by a point \mathbf{p} with integer coordinates (recall that \mathbf{p} is a projective point). Find the quadric $Q_R = Q_{R(\lambda_1)}$ through \mathbf{p} . If \mathbf{p} is close enough to \mathbf{p}_0 , then $\lambda_1 \in I$ and $\det R > 0$. We thus have a quadric of inertia $(2, 2)$ containing a point in $\mathbb{P}^3(\mathbb{Z})$: such a quadric can be parameterized with at most one square root [DLLP03].

Plugging the parameterization $\mathbf{X}((u, v), (s, t))$ of Q_R , with $(u, v), (s, t) \in \mathbb{P}^1$, in the equation of any other quadric of the pencil gives a bihomogeneous equation that has degree two in (u, v) and two in (s, t) . Solving this equation for (s, t) in terms of (u, v) and replugging in the parameterization of Q_R gives a parameterization of the smooth quartic:

$$\mathbf{X}(u, v) = \mathbf{X}_1(u, v) \pm \mathbf{X}_2(u, v) \sqrt{\Delta(u, v)},$$

where $\mathbf{X}_1(u, v)$ (resp. $\mathbf{X}_2(u, v)$) is a vector of homogeneous polynomials of degree 3 (resp. 1) and $\Delta(u, v)$ is a homogeneous polynomial of degree 4. Δ and the polynomials of \mathbf{X}_1 and \mathbf{X}_2 have coefficients in $\mathbb{Z}(\sqrt{\det R})$.

If $\det R$ is a square, then all of these polynomials have rational coefficients and the parameterization is *optimal* in terms of the degree of the extension of \mathbb{Z} on which it is defined. If $\det R$ is not a square, then we can only conclude that the parameterization is *near-optimal*: it might well be that there exists another quadric $Q_{R'}$ of inertia $(2, 2)$ in the pencil, containing a rational point, such that $\det R'$ is a square, implying that $\sqrt{\det R}$ could have been avoided in the output (see Section 17.8.2 for an example). Finding such a quadric however implies, in general, finding a rational point on a hyperelliptic curve (see [DLLP03]), a problem known to be very hard.

Case (ii) : Q_R has rank strictly less than 4

Though not generic, the situation where Q_R has rank strictly less than 4 happens quite often in practice since it covers in particular all the types of intersection corresponding, in the Segre characterization, to the determinantal equation having a single multiple root λ_0 . Indeed, in that case, the multiple root is both real (otherwise its complex conjugate would also be a multiple root of $\det R(\lambda) = 0$) and rational (otherwise its algebraic conjugate would also be a multiple root of $\det R(\lambda)$). So the associated quadric $Q_R = Q_{R(\lambda_0)}$ has rational coefficients and has rank 3 or less.

The general philosophy for parameterizing the intersection is to parameterize Q_R , plug the parameterization in any other quadric of the pencil, and solve the resulting equation in the parameters. There are however many situations in which this procedure can be simplified by the fact that we can find a rational point on Q_R outside its singular locus and thus parameterize Q_R rationally, and that we know enough information on the intersection to greatly simplify the solving and factorization of the equation in the parameters.

Let us illustrate this on the example of an intersection consisting of a cubic and a line that are tangent (Figure 17.1.b). The determinantal equation in this case has a quadruple root corresponding to a cone Q_R of inertia $(2, 1)$. By the above argument, Q_R has rational coefficients. So the vertex \mathbf{c} of Q_R has rational coordinates. \mathbf{c} is the point of tangency of the cubic and the line of the intersection. Assume $Q_R \neq Q_S$. The line of the intersection is necessarily rational (otherwise its conjugate would be in the intersection). This line can be found by intersecting the cone Q_R with the plane tangent to Q_S at \mathbf{c} . Picking any point \mathbf{p} with rational coordinates on this line other than \mathbf{c} gives a non-singular rational point on the cone. A projective cone having a rational point \mathbf{p} other than its singular locus can be rationally parameterized. Plugging this parameterization in Q_S gives an equation in the parameters of the cone which factors into two terms of total degree 1 and 3. Each factor can then be solved rationally for one parameter in terms of the other. The linear factor yields the line of the intersection and the cubic factor yields the cubic.

²⁸If however one of the initial quadrics has rank 3, then it should be used to parameterize the intersection. Doing so results in a parameterization having the same algebraic complexity in the worst case, but of smaller coefficient size.

17.4 Height of output coefficients : smooth quartics

In this section and the next, we prove theoretical bounds on the height of the coefficients of the parameterizations computed by our intersection software. We start by defining the notions of height and asymptotic height.

17.4.1 Definition of height

In what follows, we bound the asymptotic height of the coefficients of the parameterization of the intersection of two quadrics S and T with respect to the size of the coefficients of S and T . The height of such a coefficient is roughly its logarithm with base the maximum of the coefficients of S and T (in absolute value); if such a coefficient has a polynomial expression in terms of the coefficients of S and T , its asymptotic height is the (total) degree of this polynomial. However, a precise definition of the height of these coefficients needs care for various reasons. First, we want to compare, and thus define, *observed heights* (the heights computed for specific values of the input) and *theoretical asymptotic heights*.

We face the following problem for computing theoretical asymptotic heights of the coefficients of the parameterizations. Despite being, ultimately, only functions of the input S and T , these coefficients, in the smooth quartic case, are functions of not just S and T but also of an intermediate rational point \mathbf{p} which depends implicitly (and not explicitly) on S, T . Since obtaining a bound on the height of \mathbf{p} is very hard, we chose to express the asymptotic height of the parameterization as a function of the height of \mathbf{p} . As it turns out, the height of \mathbf{p} can, in practice, be neglected, so it is not really a problem (see the discussion at the end of Section 17.4.2).

In what follows, the *size* of an integer e is $\log_{10} |e|$ (assuming $|e| > 1$). The *size* of an algebraic number $e_1 + \sqrt{\delta}e_2$, where e_1, e_2, δ are integers and any two factors of δ are relatively prime, is the maximum of the sizes of e_1, e_2 , and δ . The *size* of a vector or matrix, with at most a constant number of entries, is the maximum size of the entries.

The *height* of an entity E (an integer, a vector, or a matrix) with respect to another entity x (also an integer, a vector, or a matrix) is the size of e over the size of x (assuming that the sizes of e and x are nonzero); note that if E and x are integers, the height is also equal to $\log_{|x|} |E|$. The *asymptotic height* of a function $f(x)$ with respect to an integer x is the limit of the height of $f(x)$ with respect to x when x tends to infinity. If a function f depends on a set X of variables, the *asymptotic height* of $f(X)$ with respect to X is the sum of the asymptotic heights of f with respect to each of the variables of X . For instance, if f is a polynomial in a constant number of variables, the asymptotic height of f with respect to these variables is the (total) degree of f . Finally, if $F(X)$ is matrix of functions depending on a set of variables X , the *asymptotic height* of $F(X)$ with respect to X is the maximum of the asymptotic heights of the entries of the matrix.

We mostly consider in the sequel heights and asymptotic heights with respect to S and T (that is with respect to the set of coefficients of S and T). *Heights* and *asymptotic heights* are thus considered with respect to S and T unless specified otherwise.

17.4.2 Height of the parameterization in the smooth quartic case

We consider now the case of a smooth quartic. This case is important because it is the generic intersection situation (given two random quadrics, a non-empty intersection is a smooth quartic with probability 1) and because it is also the worst case from the point of view of the height of the coefficients involved.

Let Q_R be the quadric of inertia $(2, 2)$ used to parameterize the intersection and \mathbf{p} a point of $\mathbb{P}^3(\mathbb{Z})$ on Q_R , as described in Section 17.3.1.

Proposition 17.1. *The parameterization of a smooth quartic*

$$\mathbf{X}(u, v) = \mathbf{X}_1(u, v) \pm \mathbf{X}_2(u, v) \sqrt{\Delta(u, v)}$$

is such that

- \mathbf{X}_1 has asymptotic height at most $27 + 36h_{\mathbf{p}}$,
- \mathbf{X}_2 has asymptotic height at most $8 + 11h_{\mathbf{p}}$,
- $\Delta(u, v)$ has asymptotic height at most $38 + 50h_{\mathbf{p}}$,

where $h_{\mathbf{p}}$ is the asymptotic height of \mathbf{p} .

Proof. We first show how the parameterization of Q_R is computed and then bound the height of its coefficients.

Let P be a projective transformation sending the point $\mathbf{p}_0 = (1, 0, 0, 0)^T$ to the point \mathbf{p} . Let Y denote the quadric obtained from R through the projective transformation $P : Y = P^T R P$. It follows from Sylvester's Inertia Law [Lam73] that Y has the same inertia as R , i.e. $(2, 2)$. Moreover, the point \mathbf{p}_0 belongs to Q_Y since $P\mathbf{p}_0 = \mathbf{p}$.

Let \mathbf{x} denote the vector $(x_1, x_2, x_3, x_4)^T$. Let L be $1/2$ times the differential of quadric Q_Y at \mathbf{p}_0 (one can trivially show that L is the first row of Y) and let i be such that $Y_{1,i} \neq 0$ (such an i necessarily exists). We compute the polynomial division of $Q_Y = \mathbf{x}^T Y \mathbf{x}$ by $L\mathbf{x}$ with respect to the variable x_i . The result of the division is

$$Y_{1,i}^2 (\mathbf{x}^T Y \mathbf{x}) = (L\mathbf{x})(L'\mathbf{x}) + A, \quad (17.1)$$

where the ξ -th coordinate of L' is equal to $L'_\xi = -Y_{i,i}Y_{1,\xi} + 2Y_{1,i}Y_{i,\xi}$ for $\xi = 1, \dots, 4$ and

$$A = c_j x_j^2 + c_k x_k^2 + 2c_{jk} x_j x_k$$

where j and k are equal to the two values in $\{2, 3, 4\}$ distinct from i , and c_j, c_k , and c_{jk} are coefficients defined as follows :

$$\begin{aligned} c_\xi &= Y_{\xi,\xi} Y_{i,1}^2 + Y_{i,i} Y_{\xi,1}^2 - 2Y_{\xi,1} Y_{i,1} Y_{i,\xi}, \quad \xi \in \{j, k\}, \\ c_{jk} &= Y_{j,k} Y_{i,1}^2 + Y_{j,1} Y_{k,1} Y_{i,i} - (Y_{j,1} Y_{k,i} + Y_{k,1} Y_{j,i}) Y_{i,1}. \end{aligned}$$

We assume in the following that $c_j \neq 0$ (if $c_j = 0$ but $c_k \neq 0$, we exchange the roles of j and k ; otherwise the analysis is different but similar and we omit it here). For clarity we denote in the following

$$c = c_j \quad \text{and} \quad r = Y_{1,i}.$$

We consider the projective transformation M such that, in the new projective frame, the quadric Q_Y has equation (up to a factor)

$$\mathbf{x}'^T M^T Y M \mathbf{x}' = 4x'_1 x'_2 + x'_3{}^2 - c x'_4{}^2.$$

In accordance with Equation (17.1) we choose $x'_1 = L\mathbf{x}$, $x'_2 = L'\mathbf{x}$. We apply Gauss' decomposition of quadratic forms into sum of squares to A and set $x'_3 = c x_j + c_{jk} x_k$ and $x'_4 = x_k$. Precisely, we define M such that its adjoint has its first row equal to L , its second row equal to L' , and the last two rows equal to zero except for the entry $(3, j)$ equal to c , the entry $(3, k)$ equal to c_{jk} , and the entry $(4, k)$ equal to 1.

Straightforward computations show that the four columns of M can be simplified by the factors $rc, r, 2r$, and $2r^2$, respectively. We then get

$$\mathbf{x}^T M^T Y M \mathbf{x} = r^2 c (4x_1 x_2 + x_3^2 - \det(Y) x_4^2). \quad (17.2)$$

If i, j, k are equal to 2, 3, 4 respectively, M is equal to

$$M = \begin{pmatrix} Y_{2,2} & -c & Y_{2,2} Y_{1,3} - r Y_{2,3} & M_{1,4} \\ -2r & 0 & -r Y_{1,3} & M_{2,4} \\ 0 & 0 & r^2 & M_{3,4} \\ 0 & 0 & 0 & rc \end{pmatrix},$$

$$M_{1,4} = r(Y_{1,4}(Y_{2,2} Y_{3,3} - Y_{2,3}^2) + Y_{3,4}(r Y_{2,3} - Y_{2,2} Y_{1,3}) + Y_{2,4}(Y_{1,3} Y_{2,3} - r Y_{3,3})),$$

$$M_{2,4} = r(Y_{1,4}(Y_{1,3} Y_{2,3} - r Y_{3,3}) + Y_{1,3}(r Y_{3,4} - Y_{1,3} Y_{2,4})),$$

$$M_{3,4} = r(-r^2 Y_{3,4} - Y_{2,2} Y_{1,3} Y_{1,4} + r(Y_{1,3} Y_{2,4} + Y_{1,4} Y_{2,3})).$$

We can easily parameterize the quadric of Equation (17.2) and the parameterization of the original Q_R is, with $\delta = \det(Y)$ and (u, v) and (s, t) in $\mathbb{P}^1(\mathbb{R})$,

$$PM \begin{pmatrix} ut\sqrt{\delta} & sv\sqrt{\delta} & (us - tv)\sqrt{\delta} & us + tv \end{pmatrix}^T. \quad (17.3)$$

We now bound the asymptotic height of this parameterization with respect to S, T and \mathbf{p} . For simplicity, asymptotic heights are referred to as *heights* until the end of the proof. First note that the matrix Y is equal to $P^T R P$, where R is the matrix $\lambda_1 S + \mu_1 T$ of the pencil such that $(\lambda_1, \mu_1) \in \mathbb{P}^1$ is solution of

$$\mathbf{p}^T (\lambda_1 S + \mu_1 T) \mathbf{p} = 0. \quad (17.4)$$

So $(\lambda_1, \mu_1) = (-\mathbf{p}^T T \mathbf{p}, \mathbf{p}^T S \mathbf{p})$ has height $1 + 2h_{\mathbf{p}}$ and $R = \lambda_1 S + \mu_1 T$ has height $2 + 2h_{\mathbf{p}}$. Since $P \mathbf{p}_0 = \mathbf{p}$, the first column of P has height $h_{\mathbf{p}}$ and the rest of P has height 0. We can now deduce the heights of the entries of $Y = P^T R P$. Note first that $Y_{1,1}$ is zero because \mathbf{p}_0 belongs to Q_Y . A straightforward computation thus gives that the first line and column of Y have height $2 + 3h_{\mathbf{p}}$ and the other entries have height $2 + 2h_{\mathbf{p}}$. Note that it follows that $\delta = \det Y$ has height $8 + 10h_{\mathbf{p}}$ and that, when δ is a square, $\sqrt{\delta}$ has height $4 + 5h_{\mathbf{p}}$.

It directly follows from the heights of the coefficients of Y and P that the heights of the four columns of PM are, respectively,

$$2 + 3h_{\mathbf{p}}, \quad 6 + 9h_{\mathbf{p}}, \quad 4 + 6h_{\mathbf{p}}, \quad \text{and} \quad 8 + 11h_{\mathbf{p}}.$$

The worst case for the height of the coefficients of the parameterization of Q_R happens when $\sqrt{\delta}$ is a square, because the height of these coefficients is at least the height of PM which is larger than the height of δ . We can thus assume for the rest of the proof that $\sqrt{\delta}$ is a square. It then follows from (17.3) that the coordinates of the parameterization of Q_R are polynomials of the form

$$\rho_1 ut + \rho_2 sv + \rho_3 us + \rho_4 tv. \quad (17.5)$$

The height of ρ_1 is the sum of the heights of the first column of PM and of $\sqrt{\delta}$. Similarly, we get that the heights of ρ_1, \dots, ρ_4 are

$$h_{\rho_1} = 6 + 8h_{\mathbf{p}}, \quad h_{\rho_2} = 10 + 14h_{\mathbf{p}}, \quad \text{and} \quad h_{\rho_3} = h_{\rho_4} = 8 + 11h_{\mathbf{p}}.$$

When substituting the parameterization of Q_R into the equation of one of the initial quadrics (say Q_S), we obtain an equation which can be written as

$$as^2 + bst + ct^2 = 0, \quad (17.6)$$

where a, b , and c depend on (u, v) and whose heights are

$$h_a = 1 + 2 \max(h_{\rho_2}, h_{\rho_3}) = 21 + 28h_{\mathbf{p}},$$

$$h_b = 1 + \max(h_{\rho_2}, h_{\rho_3}) + \max(h_{\rho_1}, h_{\rho_4}) = 19 + 25h_{\mathbf{p}},$$

$$h_c = 1 + 2 \max(h_{\rho_1}, h_{\rho_4}) = 17 + 22h_{\mathbf{p}}.$$

When substituting the solution $(s = 2c, t = -b \pm \sqrt{b^2 - 4ac})$ into each coordinate, of the form (17.5), of the parameterization (17.3) we obtain a parameterization of the smooth quartic in which each coordinate has the form

$$\chi_1(u, v) \pm \chi_2(u, v) \sqrt{\Delta(u, v)}.$$

The height of the coefficients of χ_1, χ_2 , and Δ are

$$h_{\chi_1} = \max(h_{\rho_1} + h_b, h_{\rho_2} + h_c, h_{\rho_3} + h_c, h_{\rho_4} + h_b) = 27 + 36h_{\mathbf{p}},$$

$$h_{\chi_2} = \max(h_{\rho_1}, h_{\rho_4}) = 8 + 11h_{\mathbf{p}},$$

$$\Delta = \max(2h_b, h_a + h_c) = 38 + 50h_{\mathbf{p}}.$$

which concludes the proof. □

Figure 17.2 shows how the observed height of the coefficients of $\Delta(u, v)$ evolves as a function of the input size s for the three variants of our implementation discussed in Section 17.6. For each value of s in a set of samples between 0 and 60, we have generated random quadrics with coefficients in the range $[-10^s, 10^s]$, computed the height of the coefficients of the parameterization of the smooth quartic and averaged the results.

The plots of Figure 17.2 show that the observed height of the coefficients of $\Delta(u, v)$ converges to 38 when no gcd computation is performed for simplifying the output parameterization. Since the asymptotic height of $\Delta(u, v)$ is at most 38 plus 50 times the height of \mathbf{p} , this suggests that the asymptotic height of \mathbf{p} is zero. Indeed, we have observed experimentally that the coordinates of \mathbf{p} are integers between -2 and 2 most of the time. Out of thousands of runs we have encountered no example where the size of \mathbf{p} had a significant impact on the height of the coefficients of the parameterization.

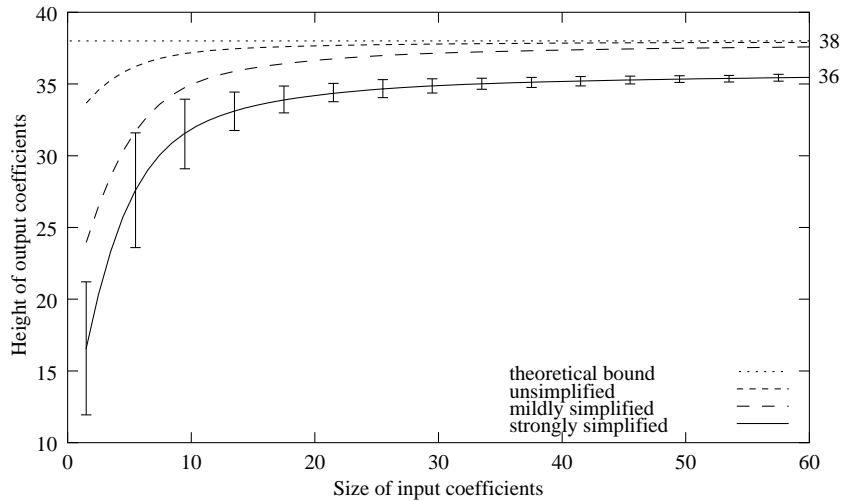


FIG. 17.2 – Evolution of the height of $\Delta(u, v)$ (smooth quartic case) as a function of the size of the input, with the standard deviation displayed on the simplified plot.

Backing this observation by theoretical results is hard, if not out of reach. Let $R = R(\lambda_1, \mu_1)$ be the quadric through \mathbf{p} . By Eq. (17.4), the size of the rational point \mathbf{p} is intimately related to the height of (λ_1, μ_1) . It is intuitively clear that if the size of the interval on which (λ_1, μ_1) is taken is small, then the size of \mathbf{p} will increase. It thus seems natural to look for results on the distance between roots of integer polynomials. Various upper and lower bounds are known as a function of the degree of the polynomial and the height of its coefficients (see, e.g., [BM04]), and pathological examples exhibiting root distances almost matching those bounds can be constructed. However, nothing is known about the average distance between the roots of a polynomial whose coefficients are uniformly distributed between $-h$ and h for some fixed integer h (personal communication with Y. Bugeaud and M. Mignotte).

Figure 17.2 also shows that the observed height of the coefficients of $\Delta(u, v)$ converges to 36 when gcd computations are performed. We ran experiments with inputs of size up to 10,000 and observed the same limit of 36 on the height of the coefficients when gcd computations are performed. We do not have any explanation as to why the bound of 38 is not reached in that case.

17.5 Height of output coefficients : singular intersections

In this section, we analyze two different types of situations to validate a key design choice we made, which is to take the quadric with rational coefficients of lowest possible rank to parameterize the intersection. We first consider the case when the pencil contains a rational cone and then when it contains a rational pair of planes. In both cases, we illustrate the fact that better results are obtained than when using a quadric of inertia $(2, 2)$ as intermediate quadric.

Table 17.1 summarizes the asymptotic heights of the parameterizations in many cases of interest.

17.5.1 Preliminaries

Let Q_R be a singular quadric corresponding to a rational root $(\lambda_0, \mu_0) \in \mathbb{P}^1(\mathbb{Z})$ of multiplicity $d \geq 1$ of the determinantal equation $\det(\lambda S + \mu T) = 0$. Here, we further assume that (λ_0, μ_0) is a representative of the root in \mathbb{Z}^2 such that $\gcd(\lambda_0, \mu_0) = 1$. We also assume that Q_R has rank r (recall that $3 \geq r \geq 4 - d$).

Lemma 17.2. *The asymptotic height of (λ_0, μ_0) is at most $\frac{4}{d}$, and the asymptotic height of $R = \lambda_0 S + \mu_0 T$ is at most $1 + \frac{4}{d}$.*

Proof. We have that

$$\det(\lambda S + \mu T) = C(\mu_0 \lambda - \lambda_0 \mu)^d (\alpha_0 \lambda^{n-d} + \dots + \alpha_{n-d} \mu^{n-d}).$$

real type of intersection	height of parameterization	inertia of Q_R used
smooth quartic	$38 + 50h_{\mathbf{p}}$	(2, 2)
nodal quartic	22	(2, 1) without rational point
cuspidal quartic	38^*	(2, 1) with rational point
cubic and secant line	22 (cubic), 9 (line)	(2, 1) with rational point**
cubic and tangent line	20 (cubic), 11 (line)	(2, 1) with rational point
two tangent conics	$20 + \frac{1}{6}$	(1, 1)
double conic	$13 + \frac{2}{3}$	(1, 0)
conic and two lines crossing	$17 + \frac{1}{2}$ (conic) and 9 (lines)	(1, 1)
two skew lines and a double line	9 (lines) and 4 (double line)	(1, 1)
two double lines	12	(1, 0)

TAB. 17.1 – Asymptotic heights of parameterizations in major cases, when the determinantal equation has a unique multiple root. In the singular cases, these values should be compared to the bound of 27 for each component if a quadric of inertia (2, 2) had been used, keeping in mind that the result could also contain an unnecessary square root. Note : (*) Since 38 is larger than 27, it might seem that using a quadric Q_R of inertia (2, 1) in the cuspidal quartic case is a bad idea and that a quadric of inertia (2, 2) would have given better results. This is in no way the case : since the intersection curve is irreducible, the equation in the parameters using a quadric of inertia (2, 2) would also have been irreducible, therefore producing a parameterization involving the square root of some polynomial. (**) We can easily find a rational point on Q_R here only when the intersection points between the cubic and the line are rational. Otherwise, we need to use a quadric Q_R of inertia (2, 2).

Since the coefficients of $\det(\lambda S + \mu T)$ are integers, we can assume that the α_i are integers and $C \in \mathbb{Q}$. We can also assume that the gcd of all the α_i is one. Recall that an integer polynomial is called *primitive* if the gcd of all its coefficients is one. Since the product of two primitive polynomials is primitive, by Gauss's Lemma (see [DST93, §4.1.2]), C is an integer (equal to the gcd of the coefficients of $\det(\lambda S + \mu T)$). Therefore, since the coefficient $C\mu_0^d\alpha_0 = \det S$ of λ^4 has asymptotic height 4, μ_0 has asymptotic height at most $\frac{4}{d}$, and similarly for λ_0 . It directly follows that $R = \lambda_0 S + \mu_0 T$ has asymptotic height at most $1 + \frac{4}{d}$. \square

Lemma 17.3. *The singular set of Q_R contains a basis of points of asymptotic height at most $r(1 + \frac{4}{d})$.*

Proof. Assume first that R has rank 3, i.e., Q_R has a singular point. Finding this singular point amounts to finding a point $\mathbf{c} \in \mathbb{P}^3(\mathbb{Z})$ in the kernel of R , i.e., such that $R\mathbf{c} = 0$. Since R has rank 3, at least one of its 3×3 minors is non-zero. Assume that the upper left 3×3 minor has this property. We decompose R such that R_u is the upper left 3×3 matrix of R and \mathbf{r}_4 is the first three coordinates of the last column of R , and \mathbf{c} such that \mathbf{c}_u is the first three coordinates and c_4 is the last. Then \mathbf{c} is found by solving

$$R_u \mathbf{c}_u = -c_4 \mathbf{r}_4.$$

A solution is thus $\mathbf{c} = (-R_u^* \mathbf{r}_4, \det R_u)$, where R_u^* is the adjoint of R_u . The asymptotic heights of R_u^* , \mathbf{r}_4 , and $\det R_u$ are the asymptotic height of R times 2, 1, and 3, respectively. The asymptotic height of \mathbf{c} is thus 3 times the asymptotic height of R . Hence, \mathbf{c} has asymptotic height at most $3(1 + \frac{4}{d})$.

The extension to general rank r is similar : Q_R contains in this case a linear space of dimension $3 - r$ of singular points. One can extract a non-singular submatrix of R of size r and points in the kernel of R have asymptotic height r with respect to the coefficients of the matrix. The result follows. \square

17.5.2 When Q_R is a cone

17.5.2.1 Parameterization of a cone

Assume now that Q_R is a real cone with vertex \mathbf{c} containing a rational point $\mathbf{p} \neq \mathbf{c}$. We want to find a rational parameterization of Q_R . First, we apply to R a projective transformation P sending the point $(0, 0, 0, 1)^T$ to \mathbf{c} and

the point $(0,0,1,0)^T$ to \mathbf{p} . We are left with the problem of parameterizing the cone $Q_{P^T R P}$ with apex $(0,0,0,1)^T$ and going through the point $(0,0,1,0)^T$. Such a cone has equation

$$a_1x^2 + a_2xy + a_3y^2 + a_4yz + a_5xz = 0. \quad (17.7)$$

A parameterization of this cone is given by

$$\mathbf{X}'(u, v, s) = \begin{pmatrix} a_5 & 0 & a_4 & 0 \\ 0 & a_4 & a_5 & 0 \\ -a_1 & -a_3 & -a_2 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} u^2 \\ v^2 \\ uv \\ s \end{pmatrix}, \quad (u, v, s) \in \mathbb{P}^{*2}(\mathbb{R}). \quad (17.8)$$

Here, $\mathbb{P}^{*2}(\mathbb{R})$ is the real quasi-projective space defined as the quotient of $\mathbb{R}^3 \setminus \{0,0,0\}$ by the equivalence relation \sim where $(x,y,z) \sim (x',y',z')$ if and only if there exists $\lambda \in \mathbb{R} \setminus \{0\}$ such that $(x,y,z) = (\lambda x', \lambda y', \lambda z')$. Lifting the parameterization to the original space by multiplying by matrix P , we have a parameterization of Q_R .

Let h_R (resp. $h_{\mathbf{p}}, h_{\mathbf{c}}$) denote the asymptotic height of R (resp. of \mathbf{p}, \mathbf{c}). From the above, we can deduce the following.

Lemma 17.4. *The parameterization $\mathbf{X}(u, v, s)$ of Q_R is such that :*

- *the asymptotic height of the coefficients of u^2, v^2, uv is $h_R + h_{\mathbf{p}}$;*
- *the asymptotic height of the coefficients of s is $h_{\mathbf{c}}$.*

Proof. The matrix P has its third column set to \mathbf{p} and its fourth column set to \mathbf{c} . We complete it so that it indeed represents a real projective transformation (i.e., its columns form a basis of \mathbb{P}^3). So the first two columns have height 0 in R, \mathbf{p} , and \mathbf{c} . Computing $P^T R P$, we see that the height of a_1, a_2 , and a_3 is the height of R and the asymptotic height of a_4 and a_5 is the sum of the asymptotic heights of R and \mathbf{p} . From this, we see that the asymptotic height of the coefficients of u^2, v^2, uv in $\mathbf{X}(u, v, s) = P\mathbf{X}'(u, v, s)$ is the sum of the asymptotic heights of R and \mathbf{p} ; also the height of the coefficients of s is the height of \mathbf{c} . \square

17.5.2.2 Cubic and tangent line

We now consider the case of an intersection consisting of a cubic and a tangent line. In this case, we can parameterize the intersection using an intermediate rational quadric Q_R of inertia either $(2,2)$ or $(2,1)$: the pencil contains an instance of both types of quadrics.

We prove the following theoretical bounds on the asymptotic height of the coefficients of the parameterizations of the cubic and the line.

Proposition 17.5. *When a quadric Q_R of inertia $(2,2)$ is used to parameterize the intersection, the parameterizations of the cubic and the line have asymptotic height at most 27 plus 36 times the asymptotic height of the point $\mathbf{p} \in Q_R$ used for parameterizing Q_R .*

Proof. The bounds found in the proof of Proposition 17.1 apply here, and in particular, the bounds $h_{\rho_1}, \dots, h_{\rho_4}, h_a, h_b$, and h_c on the heights of the coefficients of Equations (17.5) and (17.6). Equation (17.6) factors here into two terms, one of degree 0 and the other of degree 2 in, say, (u, v) , and both linear in, say, (s, t) ; Equation (17.6) can thus be written as

$$(\alpha s + \beta t)(\alpha' s + \beta' t) = a s^2 + b s t + c t^2 = 0,$$

where α, β are constants and α', β' are polynomials in (u, v) . Since $\alpha\beta' + \beta\alpha' = b$, α and the coefficients of α' have asymptotic height at most h_b . Similarly, $\beta\beta' = c$ thus β and the coefficients of β' have asymptotic height at most h_c . Substituting the solutions $(s = \beta, t = -\alpha)$ and $(s = \beta', t = -\alpha')$ into the parameterization (17.3), we get parameterizations of the cubic and the line whose coefficients have asymptotic height at most

$$h_c + \max(h_{\rho_2}, h_{\rho_3}) = h_b + \max(h_{\rho_1}, h_{\rho_4}) = 27 + 36h_{\mathbf{p}}$$

where $h_{\mathbf{p}}$ is asymptotic height of \mathbf{p} . \square

Proposition 17.6. *When a quadric Q_R of inertia $(2, 1)$ is used to parameterize the intersection, then asymptotically the parameterization of the line has height at most 11, and the parameterization of the cubic has height at most 20.*

Proof. We follow the algorithm outline given in Section 17.3.1 to determine the asymptotic height of the output.

Here, the determinantal equation has a quadruple root (λ_0, μ_0) corresponding to a quadric Q_R of inertia $(2, 1)$. The asymptotic height h_R of $R = \lambda_0 S + \mu_0 T$ is at most 2, by Lemma 17.2. The asymptotic height h_c of the singular point \mathbf{c} of Q_R is at most 6, by applying Lemma 17.3 with $d = 4$ and $r = 3$.

Since the line of the intersection is the (double) intersection of Q_R and the tangent plane to Q_S at \mathbf{c} , any point \mathbf{p} on this line satisfies

$$R\mathbf{p} = S\mathbf{c}. \quad (17.9)$$

(Observe that if \mathbf{p} is a solution, any $a_1\mathbf{p} + a_2\mathbf{c}$ is also solution.) The right-hand side $S\mathbf{c}$ of (17.9) has asymptotic height at most $6 + 1 = 7$. As in the proof of Lemma 17.3, one can assume that $\det R_u \neq 0$ and there is a unique point \mathbf{p} having zero as last coordinate. Point \mathbf{p} satisfies $\mathbf{p}_u = R_u^*(S\mathbf{c})_u$ and thus, its asymptotic height h_p is at most $4 + 7 = 11$. Overall, the coefficients of the line (\mathbf{c}, \mathbf{p}) have height 11.

We can now compute the asymptotic height of the parameterization $\mathbf{X}(u, v, s)$ of Q_R as defined in Section 17.5.2.1. By Lemma 17.4, the asymptotic height $h_{u,v}$ of coefficients of u^2, v^2, uv in $\mathbf{X}(u, v, s)$ is $h_R + h_p$, and the asymptotic height h_s of the coefficient of s is h_c . Plugging $\mathbf{X}(u, v, s)$ in the equation of any other quadric of the pencil gives an equation in the parameters of the form

$$as^2 + b(u, v)s + c(u, v) = 0, \quad (17.10)$$

where $b(u, v)$ and $c(u, v)$ have asymptotic heights respectively equal to

$$1 + h_{u,v} + h_s = 1 + h_R + h_p + h_c, \quad \text{and} \quad 1 + 2h_{u,v} = 1 + 2(h_R + h_p).$$

Observe that $a = 0$ since the singularity of the cone, which is a point of the intersection, is reached at $(u, v) = (0, 0)$ and at this point $s \neq 0$ necessarily (because $\mathbf{X}(u, v, s)$ is a faithful parameterization of the cone). We also know that (17.10) has a linear factor corresponding to the line of the intersection. By construction (see (17.8)), this line (\mathbf{c}, \mathbf{p}) is represented in parameter space by the line $a_5 u + a_4 v = 0$, where a_4 and a_5 have asymptotic height $h_R + h_p$ (see the proof of Lemma 17.4). So, after factoring out the linear term, (17.10) can be rewritten as

$$b'(u, v)s + c'(u, v) = 0. \quad (17.11)$$

The asymptotic height $h_{b'}$ of $b'(u, v)$ is $1 + h_c$, the difference of the asymptotic heights of $b(u, v)$ and of the linear factor. Similarly, the asymptotic height $h_{c'}$ of $c'(u, v)$ is $1 + h_R + h_p$, the difference of the asymptotic heights of $c(u, v)$ and of the linear factor. We plug the solution of (17.11) in s into the parameterization $\mathbf{X}(u, v, s)$ of Q_R . Multiplying by $b'(u, v)$ to clear the denominators, we get a parameterization of the cubic of asymptotic height

$$\max(h_{u,v} + h_{b'}, h_s + h_{c'}) = 1 + h_R + h_p + h_c \leq 1 + 2 + 11 + 6 = 20.$$

□

The difference in the asymptotic heights of the parameterizations underscored in the above two propositions is vindicated by some experiments we made. Figure 17.3 shows the observed heights of the coefficients of the parameterization of the cubic when a quadric Q_R of inertia $(2, 2)$ or $(2, 1)$ is used. The plots clearly show that the coefficients of the cubic are smaller when a cone is used to parameterize the intersection. The fact that the observed heights are, in the limit, so different from the theoretical bounds (8 instead of 20 when a cone is used) is most likely a consequence of the way the random quadrics are generated : it does not reflect a truly random distribution in the space of quadrics with integer coefficients of given size intersecting in a cubic and a tangent line, as explained in Section 17.6.3.

Figure 17.4 further reinforces our choice of using a cone : the parameterizations have not only smaller coefficients, they are also faster to compute.

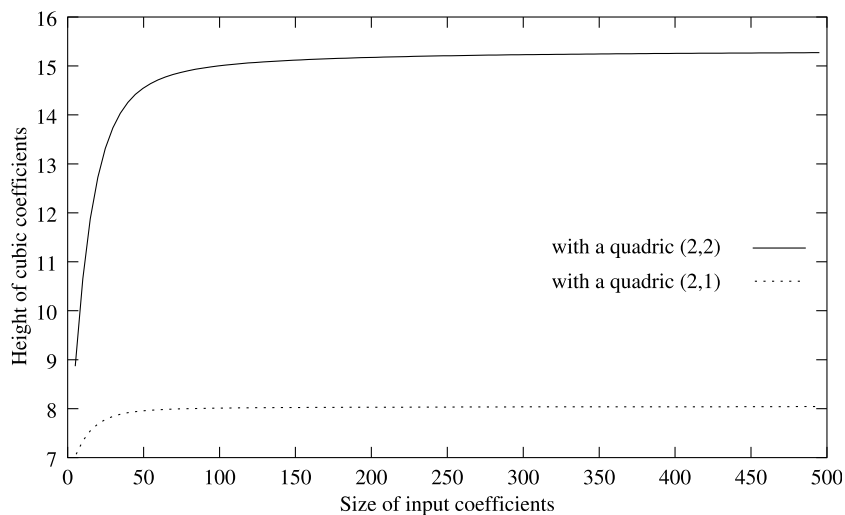


FIG. 17.3 – Observed height of the parameterization of the cubic in the cubic and tangent line case.

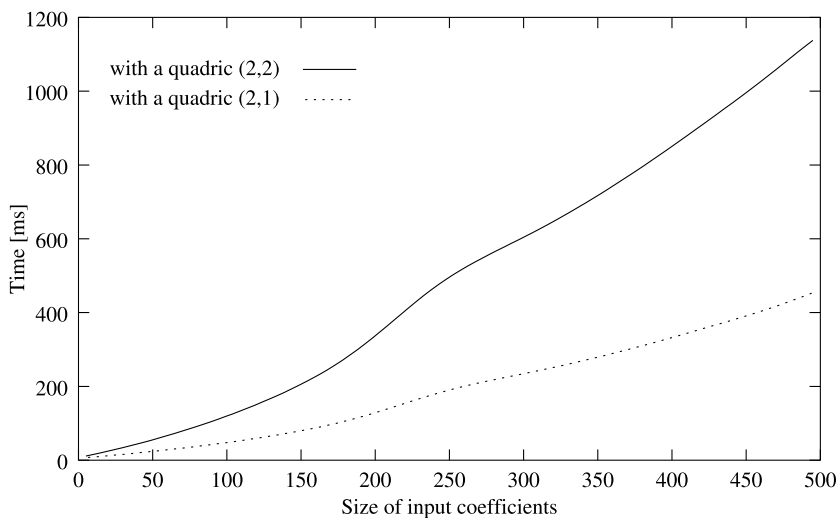


FIG. 17.4 – Computation time for the cubic and tangent line case.

17.5.3 When Q_R is a pair of planes

17.5.3.1 Parameterization of a pair of planes

We now suppose that the singular quadric Q_R corresponding to a root of multiplicity d of the determinantal equation is a pair of planes (i.e., has inertia $(1, 1)$). Let \mathbf{p}_1 and \mathbf{p}_2 two distinct points on the singular line of Q_R . Let P be a projective transformation matrix sending the point $(0, 0, 1, 0)^T$ to \mathbf{p}_1 and the point $(0, 0, 0, 1)^T$ to \mathbf{p}_2 . We are left with the problem of parameterizing the pair of planes $Q_{P^T R P}$ whose singular line contains $(0, 0, 1, 0)^T$ and $(0, 0, 0, 1)^T$. Such a pair of planes has equation

$$a_1x^2 + 2a_2xy + a_3y^2 = 0,$$

and it can be parameterized by $M_{\pm}(u, v, s)^T$ with

$$M_{\pm} = \begin{pmatrix} -a_2 \pm \sqrt{\delta} & 0 & 0 \\ a_1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad \delta = a_2^2 - a_1 a_3, \quad (u, v, s) \in \mathbb{P}^2.$$

Lifting this parameterization to the original space by multiplying by matrix P , we obtain a parameterization of Q_R .

Let h_R (resp. $h_{\mathbf{p}_1}, h_{\mathbf{p}_2}$) denote the asymptotic height of R (resp. of $\mathbf{p}_1, \mathbf{p}_2$). From the above, we deduce the following.

Lemma 17.7. *The asymptotic height of the coefficients a_i in M_{\pm} is h_R . Furthermore, if δ is a square, the parameterization $\mathbf{X}_{\pm}(u, v, s)$ is such that :*

- the asymptotic height of the coefficients of u is h_R ;
- the asymptotic heights of the coefficients of v and s are $h_{\mathbf{p}_1}$ and $h_{\mathbf{p}_2}$, respectively.

Proof. In the parameterization of the pair of planes, the first two columns of P can be completed with 0 and 1 so that it is a non-singular matrix. A straightforward computation then gives that the height of a_1, a_2 , and a_3 is the height of R . Hence, the coefficient of u in $\mathbf{X}_{\pm}(u, v, s)$ has same asymptotic height as R , and the coefficients of v and s have the same heights as \mathbf{p}_1 and \mathbf{p}_2 , respectively. \square

17.5.3.2 Two tangent conics

We now consider the case of two tangent conics. This time, we have three possibilities for Q_R : inertia $(2, 2)$, $(2, 1)$, or $(1, 1)$.

Proposition 17.8. *When the intersection consists of two tangent conics, the parameterization of each of the conics is as follows :*

- when Q_R has inertia $(1, 1)$, the parameterization has asymptotic height at most $20 + \frac{1}{6}$;
- when Q_R has inertia $(2, 1)$, the parameterization has asymptotic height at most $30 + \frac{5}{6}$;
- when Q_R has inertia $(2, 2)$, the parameterization has asymptotic height at most 27 plus 36 times the asymptotic height of the point on Q_R used for parameterizing Q_R ; moreover the coefficients may contain an unnecessary square root.

Proof. The determinantal equation in this case has a real rational triple root corresponding to a pair of planes and a real rational simple root corresponding to a real cone. The pencil also contains quadrics of inertia $(2, 2)$. The rational point of tangency \mathbf{p} of the two conics is the point of intersection of the singular line of the pair of planes with any other quadric of the pencil.

Let us first bound the asymptotic height $h_{\mathbf{p}}$ of point \mathbf{p} . Let $\mathbf{c}_1, \mathbf{c}_2$ be a basis for the singular set of the pair of planes of the pencil. By Lemma 17.3, with $d = 3$ and $r = 2$, \mathbf{c}_1 and \mathbf{c}_2 have asymptotic height $h_{\mathbf{c}_i}$ at most $\frac{14}{3}$. \mathbf{p} is the point of tangency of the line spanned by \mathbf{c}_1 and \mathbf{c}_2 with any quadric of the pencil other than the pair of planes. Let $\mathbf{p} = \alpha_0 \mathbf{c}_1 + \beta_0 \mathbf{c}_2$, where $(\alpha_0, \beta_0) \in \mathbb{P}^1$. Then (α_0, β_0) is the double root of the equation

$$(\alpha_0 \mathbf{c}_1 + \beta_0 \mathbf{c}_2)^T S (\alpha_0 \mathbf{c}_1 + \beta_0 \mathbf{c}_2) = 0.$$

By Lemma 17.2, the asymptotic height of (α_0, β_0) is at most $h_{\mathbf{c}_i} + \frac{1}{2}$. Thus, $h_{\mathbf{p}} \leq 2h_{\mathbf{c}_i} + \frac{1}{2} \leq 2\frac{14}{3} + \frac{1}{2} = \frac{59}{6}$.

Q_R has inertia $(1, 1)$. We consider the case where Q_R is the pair of planes of the pencil. We compute a parameterization $\mathbf{X}_{\pm}(u, v, s) = PM_{\pm}(u, v, s)^T$ of each of the planes of Q_R by sending $(0, 0, 1, 0)^T$ to \mathbf{c}_1 and $(0, 0, 0, 1)^T$ to \mathbf{p} as in Section 17.5.3.1. Plugging each of the $\mathbf{X}_{+}(u, v, s)$ and $\mathbf{X}_{-}(u, v, s)$ in the equation of Q_S gives a degree-two homogeneous equation in u, v , and s (i.e., $\mathbf{X}_{\pm}^T(u, v, s)S\mathbf{X}_{\pm}(u, v, s)$). This projective conic contains the point $(0, 0, 1)^T$ since $PM_{\pm}(0, 0, 1)^T = \mathbf{p}$ by definition of P and M_{\pm} . Such a conic has equation

$$\mathbf{X}_{\pm}^T(u, v, s)S\mathbf{X}_{\pm}(u, v, s) = b_1 u^2 + b_2 uv + b_3 v^2 + b_4 vs + b_5 us = 0 \quad (17.12)$$

which can be parameterized, similarly as for (17.7), by

$$\mathbf{X}'(u', v', s') = \begin{pmatrix} b_5 & 0 & b_4 \\ 0 & b_4 & b_5 \\ -b_1 & -b_3 & -b_2 \end{pmatrix} \begin{pmatrix} u'^2 \\ v'^2 \\ u'v' \end{pmatrix}, \quad (u', v') \in \mathbb{P}^1(\mathbb{R}).$$

Plugging $\mathbf{X}'(u', v', s')$ into the parameterization of Q_R gives $PM_{\pm}\mathbf{X}'(u', v', s')$, the parameterizations of the two conics of intersection.

We now compute the asymptotic height of the parameterizations $PM_{\pm}\mathbf{X}'(u', v', s')$. We assume first that δ in M_{\pm} is a square. Let h_{b_i} , denote the asymptotic height of b_i , and h_a the asymptotic height of $\{a_1, a_2, a_3\}$ in M_{\pm} . The asymptotic height of the three coordinates of $\mathbf{X}'(u', v', s')$ are, respectively,

$$\max(h_{b_4}, h_{b_5}), \quad \max(h_{b_4}, h_{b_5}), \quad \text{and} \quad \max(h_{b_1}, h_{b_2}, h_{b_3}).$$

Thus, the asymptotic height of each of the coordinates of $M_{\pm}\mathbf{X}'(u', v', s')$ are, respectively,

$$h_a + \max(h_{b_4}, h_{b_5}), \quad h_a + \max(h_{b_4}, h_{b_5}), \quad \max(h_{b_4}, h_{b_5}), \quad \text{and} \quad \max(h_{b_1}, h_{b_2}, h_{b_3}).$$

The third and fourth columns of P are \mathbf{c}_1 and \mathbf{p} , and P can be completed with 0 and 1 so that it is a non-singular matrix. Thus, the asymptotic height of $PM_{\pm}\mathbf{X}'(u', v', s')$ is the maximum of

$$h_a + \max(h_{b_4}, h_{b_5}), \quad h_{\mathbf{c}_1} + \max(h_{b_4}, h_{b_5}), \quad \text{and} \quad h_{\mathbf{p}} + \max(h_{b_1}, h_{b_2}, h_{b_3}).$$

Now, the asymptotic height of each b_i is one plus the sum of the asymptotic heights of two of the coefficients of u , v , and s in $\mathbf{X}_{\pm}(u, v, s)$ (by Equation (17.12)). Lemma 17.7 yields

$$h_{b_1} = 1 + 2h_R, \quad h_{b_2} = 1 + h_R + h_{\mathbf{c}_i}, \quad h_{b_3} = 1 + 2h_{\mathbf{c}_i}, \quad h_{b_4} = 1 + h_{\mathbf{c}_i} + h_{\mathbf{p}}, \quad h_{b_5} = 1 + h_R + h_{\mathbf{p}}.$$

Since $h_R \leq 1 + \frac{4}{3} = \frac{7}{3}$ by Lemma 17.2, $h_a \leq \frac{7}{3}$ by Lemma 17.7, $h_{\mathbf{c}_i} \leq \frac{14}{3}$, and $h_{\mathbf{p}} \leq \frac{59}{6}$, we get $h_{b_1} \leq \frac{17}{3}$, $h_{b_2} \leq \frac{24}{3}$, $h_{b_3} \leq \frac{31}{3}$, $h_{b_4} \leq \frac{31}{2}$, and $h_{b_5} \leq \frac{79}{6}$. Hence, if δ is a square, the asymptotic height of the parameterization $PM_{\pm}\mathbf{X}'(u', v', s')$ of the two conics of intersection is at most

$$\max\left(\frac{7}{3} + \frac{31}{2}, \quad \frac{14}{3} + \frac{31}{2}, \quad \frac{59}{6} + \frac{31}{3}\right) = \frac{121}{6} = 20 + \frac{1}{6}.$$

Finally, since this bound is larger than the asymptotic height of δ (which is $2h_a \leq \frac{14}{3}$), the asymptotic height of $PM_{\pm}\mathbf{X}'(u', v', s')$ can only be less than or equal to $20 + \frac{1}{6}$, even if δ is not a square.

Q_R has inertia (2, 1). Let now Q_R be the cone of the pencil with apex \mathbf{c} . By Lemma 17.4, we have a rational parameterization $\mathbf{X}(u, v, s)$ of Q_R whose coefficients in u^2, v^2, uv have asymptotic height $h_R + h_{\mathbf{p}}$ and whose coefficient in s has asymptotic height $h_{\mathbf{c}}$. Plugging this parameterization into the equation of any other quadric of the pencil gives an equation in the parameters of the form

$$as^2 + b(u, v)s + c(u, v) = 0, \tag{17.13}$$

where the asymptotic heights of $a, b(u, v)$, and $c(u, v)$ are, respectively,

$$1 + 2h_{\mathbf{c}}, \quad 1 + h_{\mathbf{c}} + h_R + h_{\mathbf{p}}, \quad \text{and} \quad 1 + 2(h_R + h_{\mathbf{p}}).$$

We know (17.13) factors in two quadratic factors corresponding to the two conics. Also, by construction (see (17.8)), the ruling of Q_R on which \mathbf{p} lies is represented in parameter space by the line $a_5 u + a_4 v = 0$, where a_4, a_5 are as in Section 17.5.2.1. As in the proof of Lemma 17.4, the asymptotic height of a_4 and a_5 is $h_R + h_{\mathbf{p}}$. Point \mathbf{p} must be on each conic on intersection, and \mathbf{p} corresponds in parameter space to (u, v, s) such that $s = a_5 u + a_4 v = 0$. So (17.13) rewrites

$$(\alpha_1 s + (a_5 u + a_4 v)\beta_1(u, v))(\alpha_2 s + (a_5 u + a_4 v)\beta_2(u, v)) = 0,$$

where β_1 and β_2 are linear in u, v (possibly defined over an extension of \mathbb{Z} by the square root of the discriminant of the pair of planes containing the conics). The asymptotic height of $\alpha_1 \beta_2 + \alpha_2 \beta_1$ is $1 + h_{\mathbf{c}}$, the difference of the

asymptotic heights of $b(u, v)$ and of the linear factor. The asymptotic height of $\beta_1\beta_2$ is 1, the difference of the asymptotic height of $c(u, v)$ and of twice the asymptotic height of the linear factor. Hence, the asymptotic height of each β_i is at most 1, and the height of each α_i is at most $1 + h_c$. Solving each factor rationally for s and plugging the solution into the parameterization $\mathbf{X}(u, v, s)$ of Q_R , we get parameterizations of the conics with asymptotic height $1 + h_c + h_R + h_p$. Applying Lemmas 17.2 and 17.3 with $r = 3$ and $d = 1$, and the bound on h_p found above, the asymptotic height of the parameterizations of the conics is at most $1 + 15 + 5 + \frac{59}{6} = 30 + \frac{5}{6}$.

Q_R **has inertia** (2,2). When a quadric Q_R of inertia (2,2) is used, the biquadratic equation (17.6) factors in two factors of bidegree (1, 1) corresponding to the conics. Factoring introduces, as above, the square root of the discriminant of the pair of planes containing the conics. Proceeding as in the proof of Proposition 17.5, we get that the height of each factor is at most 27 plus 36 times the asymptotic height of the point on Q_R used for parameterizing Q_R .

Moreover, we might have an extra square root in the result if the determinant of R is not a square. Consider for instance

$$\begin{cases} Q_S : x^2 - 2w^2 = 0, \\ Q_T : xy + z^2 = 0. \end{cases}$$

Here, the determinantal equation is $2\lambda\mu^3 = 0$. $\sqrt{2}$ (i.e., the discriminant of the pair of planes) cannot be avoided in the result. The point $\mathbf{p} = (-1, 3, 0, 0)$ is contained in the quadric $3Q_S + Q_T$ of inertia (2,2) and determinant 6. If this quadric is used to parameterize the intersection, we have an extra square root, namely $\sqrt{6}$. \square

17.6 Implementation

We now move on to a description of the main design choices we made to implement our near-optimal parameterization algorithm.

17.6.1 Implementation outline

Our implementation builds upon the LiDIA [LiD] and GMP [GMP] C/C++ libraries. LiDIA was originally developed for computational number theory purposes, but includes many types of simple parameterized and template classes that are useful for our application. Apart from simple linear algebra routines and algebraic operations on univariate polynomials, we use LiDIA's number theory package and its ability to manipulate vectors of polynomials, polynomials having other polynomials as coefficients, etc. On top of it, we have added our own data structures. We have compiled LiDIA so that it uses GMP multiprecision integer arithmetic. From now on, we refer to the multiprecision integers as `bigints`, following the terminology of LiDIA.

Our implementation consists of more than 17,000 lines of source code, which is essentially divided into the following chapters :

- *data structures* (1,500 lines) : structures for intersections of quadrics, for components of the intersection, for homogeneous polynomials with `bigint` coefficients (coordinates of components), for homogeneous polynomials with `bigint` polynomials as coefficients, and basic operations on these structures, etc.
- *elementary operations* (2,000 lines) : computing the inertia of a quadric of `bigints`, the coefficients of the determinantal equation, the gcd of the derivatives of the determinantal equation, the adjoint of a matrix, the singular space of a quadric, the intersection between two linear spaces, applying Descartes's Sign Rule, the Gauss decomposition of a quadratic form into a sum of squares, isolating the roots of a univariate polynomial using Uspensky's method, etc.
- *number theory and simplifications* (1,500 lines) : gcd simplifications of the `bigint` coefficients of a polynomial, a vector or a matrix, simplifications of the coefficients of pairs and triples of vectors, reparameterization of lines so that its representative points have small height, ...
- *quadric parameterizations* (2,000 lines) : parameterization of a quadric of inertia (2,2) with `bigint` coefficients going through a rational point, of a cone (resp. conic), of a cone (resp. conic) with a rational point, of a pair of planes, etc.

- *intersection parameterizations* (9,000 lines) : dedicated procedures for parameterizing the components of the intersection in all possible cases, i.e., when the determinantal equation has no multiple root (1,500 lines), one multiple root (3,000 lines), two multiple roots (1,500 lines) or when it vanishes identically (3,000 lines).
- *printing and debugging* (1,000 lines) : turning on debugging information with the `DEBUG` preprocessor directive, checking whether the computed parameterizations are correct, pretty printing the parameterizations, etc.

17.6.2 Implementation variants

Three variants of our implementation are available and using one rather than the other might depend on the context or the application (see Section 17.7). They are :

- *unsimplified* : nothing is done to simplify the coefficients either during the computations or in the parameterizations computed ;
- *mildly simplified* : some gcds are performed at an early stage (optimization of the coefficients and of the roots of the determinantal equation, optimization of the coordinates of singular and rational points, etc.) to avoid hampering later calculations with unnecessarily big numbers ;
- *strongly simplified* : mildly simplified, plus extraction of the square factors of some bigints (like in the smooth quartic case, where $\sqrt{\det R}$ can be replaced by $b\sqrt{a}$ if $\det R = ab^2$) and gcd simplifications of the coefficients of the final parameterizations.

For the extraction of the square factors of an integer n , the strongly simplified variant finds all the prime factors of n up to $\min(\lceil \sqrt[3]{n} \rceil, \text{MAXFACTOR})$, where `MAXFACTOR` is a predefined global variable.

Let us finally mention that we tried a fourth variant of our implementation where the extraction of the square factors is done by fully factoring the numbers (using the Elliptic Curve Method and the Quadratic Sieve implemented in LiDIA [LiD]). But this variant is almost of no interest : for small input coefficients, the strongly simplified variant already finds all the necessary factors, and for medium to large input coefficients, integer factoring becomes extremely time consuming.

17.6.3 Generating random intersections

Our implementation has been tested both on real and random data (see Section 17.7). Generating random intersections of a given type, i.e., random pairs of quadrics intersecting along a curve of prescribed topology, is however difficult. We discuss this issue here.

In the smooth quartic case, random examples can be generated by taking input quadrics with random coefficients. Indeed, given two random quadrics, the intersection is a smooth quartic or the empty set with probability one. (Of course, this does not allow to distinguish between the different morphologies of a real smooth quartic, i.e., one or two, affinely finite or infinite, components.)

When the desired intersection is not a smooth quartic, one way to proceed is to start with a canonical pair of quadrics intersecting in a curve of the prescribed type and apply to this pair a random transformation. More precisely, given a canonical pair S, T , four random coefficients r_1, r_2, r_3, r_4 , with $r_1 r_4 - r_2 r_3 \neq 0$, and a random projective transformation P , we consider the “random” quadrics with matrices S' and T' :

$$S' = P^T (r_1 S + r_2 T) P, \quad T' = P^T (r_3 S + r_4 T) P.$$

If we take the r_i and the coefficients of P randomly in the range $[-\lceil \sqrt[3]{10^s} \rceil, \lceil \sqrt[3]{10^s} \rceil]$, then the quadrics S' and T' have asymptotic expected size s (the size of the canonical pair S, T can be neglected).

The problem here is two-fold. First, since we want the matrices S' and T' to have integer coefficients (because that is what our implementation takes), we have to assume that the r_i and the coefficients of P are integers. But then the above procedure certainly does not reflect a truly random distribution in the space of quadrics with integer coefficients. Indeed, quadrics S' and T' with integer coefficients intersecting in the prescribed curve might exist without P having integer coefficients. Consider for instance the two pairs of quadrics

$$\begin{cases} Q_S : x^2 - w^2 = 0, \\ Q_T : xy + z^2 = 0, \end{cases} \quad \begin{cases} Q_{S'} : x^2 - 2w^2 = 0, \\ Q_{T'} : xy + z^2 = 0. \end{cases}$$

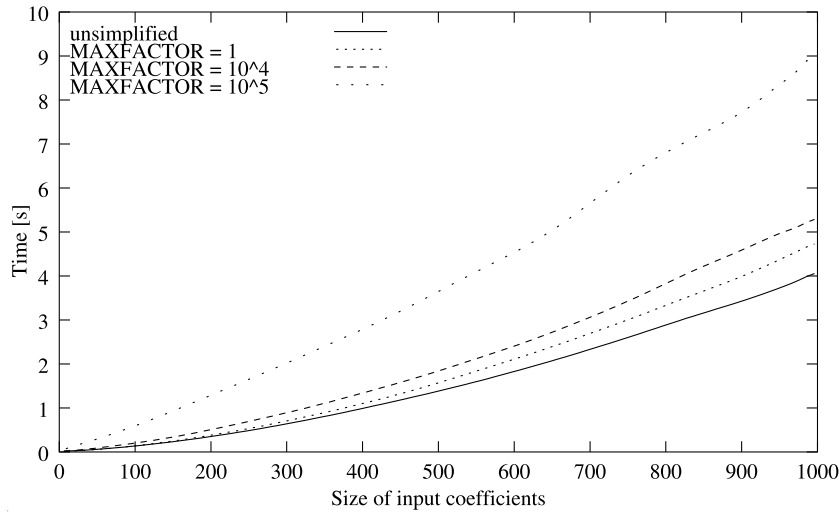


FIG. 17.5 – Evolution of execution time in the smooth quartic case as a function of the size of the input for very large input sizes.

The first pair is a canonical form for the case of an intersection made of two real tangent conics. Both pairs generate an intersection of the same type. But the second form cannot be generated from the first using a transformation matrix P with integer coefficients.

As for the second issue, consider the determinantal equation of the pencil generated by S', T' :

$$\det R'(\lambda, \mu) = \det(\lambda S' + \mu T') = (\det P)^2 \det((\lambda r_1 + \mu r_3)S + (\lambda r_2 + \mu r_4)T).$$

In other words, since P is now assumed to have integer entries, the coefficients of the determinantal equation all have a common integer factor, $(\det P)^2$. So, after simplification by this common factor, the coefficients have asymptotic height $\frac{4}{3}$, instead of 4, with respect to S', T' . This explains why the asymptotic heights are not reached.

Note that the same problems appear when working the reverse way, i.e., start with the canonical parameterization \mathbf{X} of a required type of intersection, apply a random transformation P , recover the pencil of quadrics $R'(\lambda, \mu)$ containing the curve parameterized by $P\mathbf{X}$ and filter them according to the height of their coefficients. Indeed, in that case, $R'(\lambda, \mu) = P^T R(\lambda, \mu) P$, where $R(\lambda, \mu)$ is the pencil of quadrics through the curve parameterized by \mathbf{X} .

Effectively generating random pairs of quadrics with a prescribed intersection type is an open problem.

17.7 Experimental results

We now report on some experimental results and findings from our implementation.

The experiments were made on a Dell Precision 360 with a 2.60 GHz Intel Pentium CPU. LiDIA, GMP and our code were compiled with g++ 3.2.2.

17.7.1 Random data

Let us first discuss the impact of the MAXFACTOR variable (see Section 17.6.2) on the output. Figure 17.5 shows that values of 10^5 and higher have a dramatic impact on computation time while all values less than 10^4 are acceptable. We have determined that the best compromise between efficiency and complexity of the output is obtained by setting MAXFACTOR to 10^3 , which we assume now.

Figure 17.6 shows the evolution of the aggregate computation time in the smooth quartic case, which is the most computationally demanding case, with the three variants outlined above. We infer from these plots that the computation times for the unsimplified and mildly simplified variants are very similar, while we observe (see Figure 17.2) a dramatic improvement in the height of the output coefficients with the mildly simplified variant for

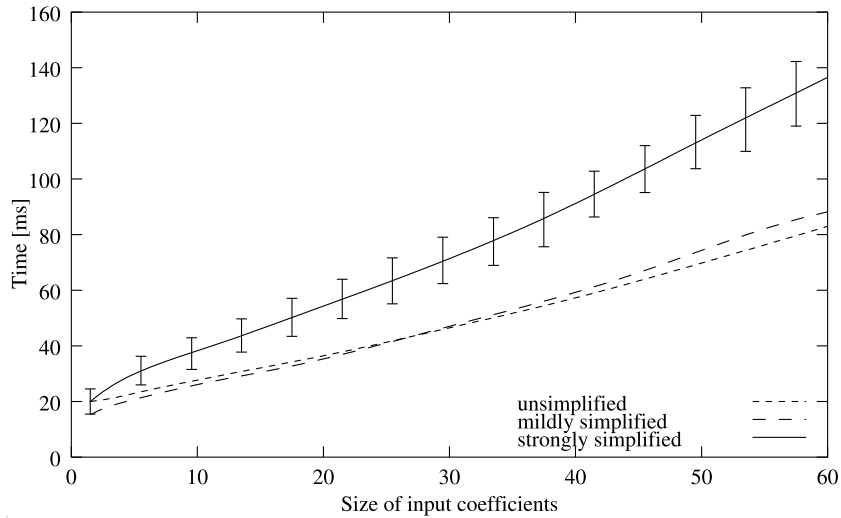


FIG. 17.6 – Evolution of execution time in the smooth quartic case as a function of the input size, with the standard deviation shown on the simplified plot.

reasonably small inputs. This explains our choice of putting the mild simplifications in the form of a preprocessor directive, not a binary argument : they might as well have been called *mandatory simplifications*.

A second lesson to be learned from Figures 17.2 and 17.6 is that for an input with coefficients ranging from roughly 5 to 60 digits, the computation time is roughly 30% larger for the strongly simplified variant than for the mildly simplified. At the same time, the height of the output is on average between 20% (input size of 5) and 5% (input size of 60) smaller. For large values of the input size, the difference in computation time between the mildly simplified and the strongly simplified variants drops to less than 10% (see the two curves in Figure 17.5 with MAXFACTOR equal to 1 and 10^4), but not much is gained in terms of height of the output (see Figure 17.2).

Another interesting piece of information inferred from Figure 17.2 is that the standard deviation of the height of the output coefficients is large for small input size in the strongly simplified variant. This means that in the good cases the height of the output is dramatically smaller than the height in the mildly simplified case, and in the bad cases is similar to it.

Deciding to spend time on simplification essentially depends on the application. For most real-world applications, where the size of the input quadrics is small by construction, we believe simplifying is important : it should be kept in mind that the computed parameterizations are often the input to a later processing step (like in boundary evaluation) and limiting the growth of the coefficients at an early stage makes good sense.

A last comment that can be made looking at Figure 17.5 concerns the efficiency of our implementation. Indeed, those plots show that we can compute the parameterization of the intersection of two quadrics with coefficients having 400 digits in 1 second and 1,000 digits in 5 seconds (on average).

Efficiency can be measured in a different way. In Figure 17.7, we have plotted the total computation time, with the strongly simplified variant, for a file containing 120 pairs of quadrics covering all intersection situations over the reals. The “random” quadrics were generated as in Section 17.5.2.2. For an input size $s = 500$, the total computation time is roughly 72 seconds, on average, for the 120 pairs of quadrics, i.e., 0.6 second per intersection. This should be compared to the 1.7 seconds on average needed to compute the intersection in the smooth quartic case for the same size of input (Figure 17.5). This difference is simply explained by the fact that very degenerate intersections (like when the determinantal equation vanishes identically, which represents 36 of the 120 quadrics in the file) are usually much faster to compute.

Our last word will be on memory consumption. Our implementation consumes very little memory. In the smooth quartic case, the total memory chunks allocated sum up to less than 64 kilobytes for input sizes up to 20. It takes input coefficients of more than 700 digits to get to the 1 MB range of used memory.

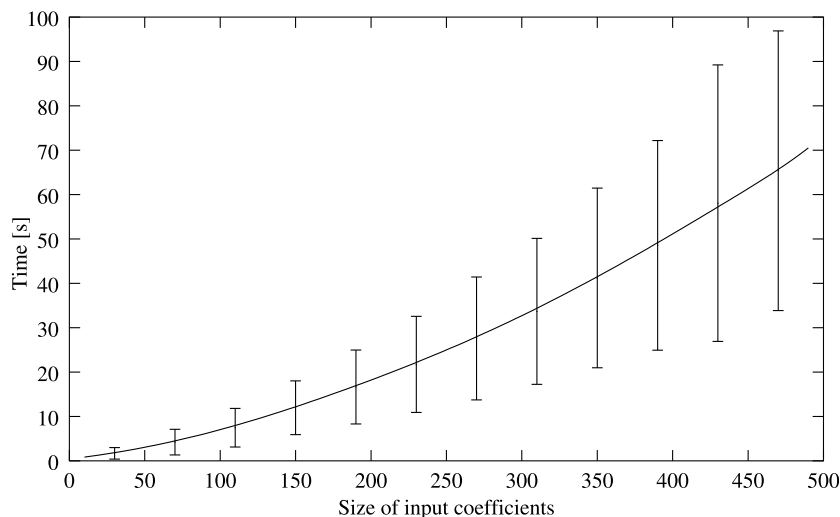


FIG. 17.7 – Computation time for 120 pairs of quadrics covering all intersection cases, with standard deviation.

17.7.2 Real data

Our intersection code has also been tested on real solid modelling data. Our three test scenes are the teapot, the pencil box, and the chess set (Figure 17.8). They were modelled with the SGDL modelling kernel [SGD]. The chess set was rendered with a radiosity algorithm using the virtual mesh paradigm [ACP⁺01]. All computations were made with the strongly simplified variant of our implementation.

The teapot (Figure 17.8.a) is made of 18 distinct quadrics (one hyperboloid of one sheet, one cone, one circular cylinder, two elliptic cylinders, two ellipsoids, four spheres, and seven pairs of planes). The coefficients of each input quadric have between 2 and 5 digits. The 153 intersections (i.e., pairs of quadrics) are computed in 450 milliseconds, or 2.9 ms on average per intersection. They consist in 51 real smooth quartics, 31 nodal quartics, 35 cuspidal quartics, 65 conics, 101 lines, and 9 points. The height of the output never exceeds 6 in terms of the input.

The pencil box (Figure 17.8.b) is made of 61 quadrics, most of which are pairs of planes. The input size for each quadric is between 2 and 5 for most quadrics, with four quadrics having a size of 18. The 1,830 intersections are computed in 6.25 s, or 3.4 ms per intersection on average. They consist in 65 smooth quartics, 356 nodal quartics, 119 cubics, 612 conics, 2,797 lines, and 139 points. The height of the output reaches 11 for some smooth quartics.

In the chess set (Figure 17.8.c), the pawn, the bishop, the knight, the rook, the king, and the queen are respectively made of 12, 14, 20, 18, 19, and 25 quadrics. Most of the quadrics have coefficients with between 2 and 7 digits, except for a small number having 15 digits (the crown of the queen has for instance been generated by rotations of $\pi/10$ applied to a sphere). The intersections were computed for each piece separately. They consist in 86 smooth quartics, 123 nodal quadrics, 360 cuspidal quartics, 284 conics, 484 lines, and 13 points. In total, the 971 intersections were computed in 3.33 s, or 3.4 ms per intersection on average. The height of the output never exceeds 8.

17.8 Examples

We now give four examples of parameterizations computed by our algorithm. Other examples can be tested by querying our parameterization server.

Comparing our results with the parameterizations computed with other methods does not make much sense since our implementation is the first to output exact parameterizations in all cases. However, for the sake of illustration, our first two examples are taken from the paper describing the plane cubic curve method of Wang, Joe, and Goldman [WJG02].

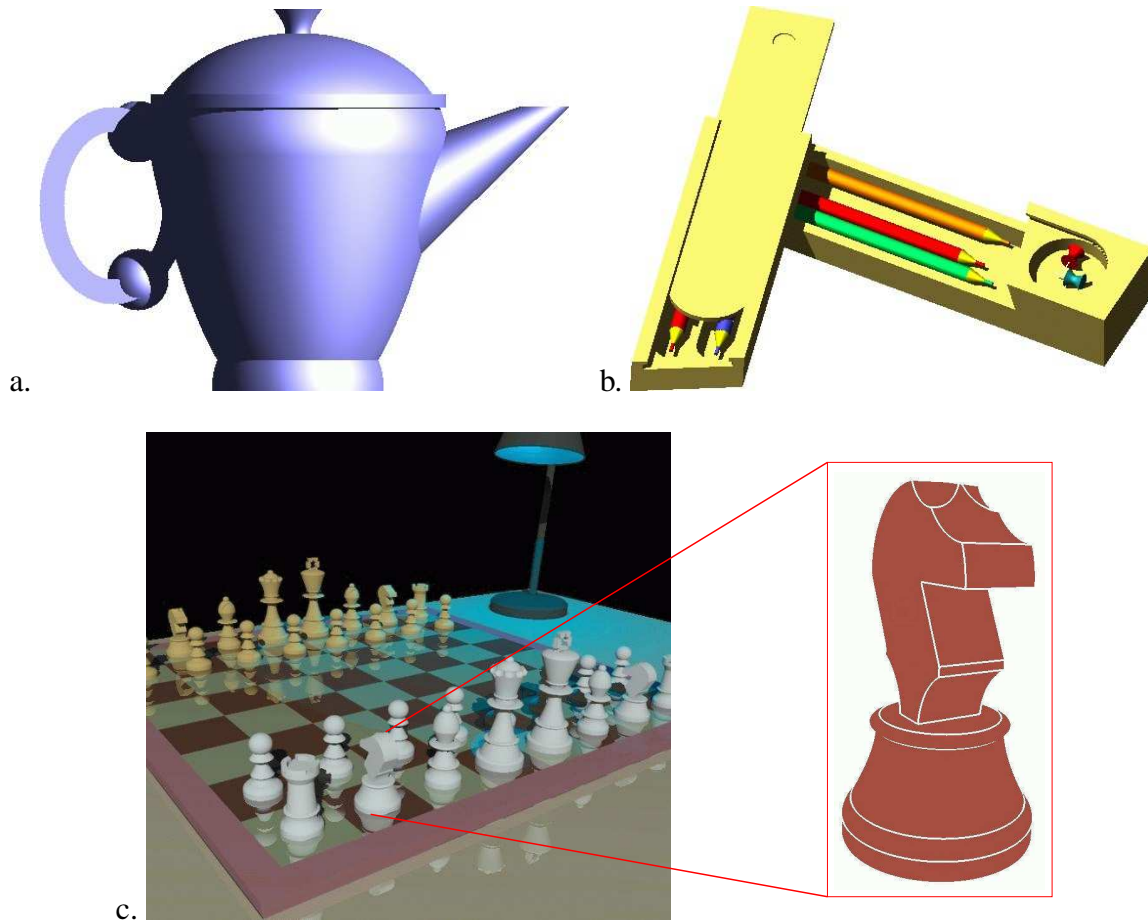


FIG. 17.8 – Three CSG models made entirely of quadrics (models courtesy of SGDL Systems, Inc.). a. A teapot. b. A pencil box. c. A chess set, with a close-up on the knight.

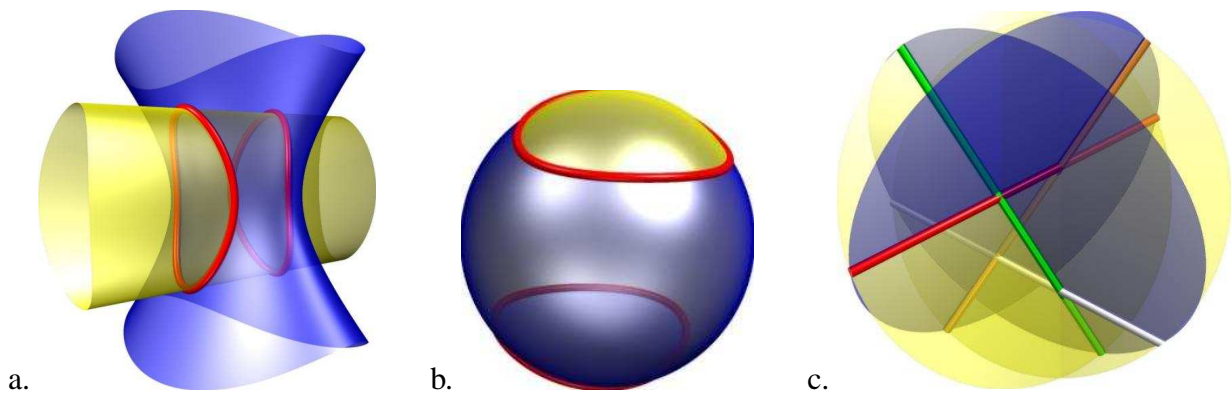


FIG. 17.9 – Further examples of intersection. a. b. Smooth quartics. c. Four skew lines.

17.8.1 Example 1 : smooth quartic

Our first example is Example 4 from [WJG02]. The two quadrics are a quadric of inertia $(2, 1)$ (an elliptic cylinder) and a quadric of inertia $(2, 2)$ (a hyperboloid of one sheet). The curve of intersection C has implicit

equation

$$\begin{cases} 4x^2 + z^2 - w^2 = 0, \\ x^2 + 4y^2 - z^2 - w^2 = 0. \end{cases}$$

A rendering of the intersection is given in Figure 17.9.a.

In [WJG02], the authors find the following parameterization for C :

$$\mathbf{X}(u, v) = \mathbf{X}_1(u, v) \pm \mathbf{X}_2(u, v) \sqrt{\Delta(u, v)}, \quad (u, v) \in \mathbb{P}^1(\mathbb{R}), \quad (17.14)$$

with

$$\mathbf{X}_1(u, v) = \begin{pmatrix} 0.0 \\ 1131.3708u^3 - 5760.0u^2v + 10861.1602uv^2 - 8192.0v^3 \\ -1600.0u^3 + 10861.1602u^2v - 21504.0uv^2 + 11585.2375v^3 \\ 1600.0u^3 + 3620.2867u^2v + 5120.0uv^2 + 11585.2375v^3 \end{pmatrix}, \quad \mathbf{X}_2(u, v) = \begin{pmatrix} -80.0u + 1181.0193v \\ 0.0 \\ 0.0 \\ 0.0 \end{pmatrix},$$

and $\Delta(u, v) = 905.0967u^3v - 3328.0u^2v^2 + 2896.3094uv^3$. The authors report a computation error on this example (measured as the maximum distance from a sequence of sample points on the curve to the input quadrics) of order $O(10^{-7})$.

Our implementation outputs the following exact and simple result in less than 10 ms :

$$\mathbf{X}(u, v) = \begin{pmatrix} 2u^3 - 6uv^2 \\ 7u^2v + 3v^3 \\ 10u^2v - 6v^3 \\ 2u^3 + 18uv^2 \end{pmatrix} \pm \begin{pmatrix} -2v \\ u \\ 2u \\ 2v \end{pmatrix} \sqrt{-3u^4 + 26u^2v^2 - 3v^4}, \quad (u, v) \in \mathbb{P}^1(\mathbb{R}).$$

The polynomials involved in the parameterization are defined in $\mathbb{Z}[u, v]$, which means we are in the lucky case where the intermediate quadric of inertia $(2, 2)$ found to parameterize the intersection has a square as determinant. So the parameterization obtained is optimal (in the extension of \mathbb{Z} on which its coefficients are defined).

17.8.2 Example 2 : smooth quartic

Our second example is Example 5 from [WJG02]. It is the intersection of a sphere and an ellipsoid that are very similar (see Figure 17.9.b) :

$$\begin{cases} 19x^2 + 22y^2 + 21z^2 - 20w^2 = 0, \\ x^2 + y^2 + z^2 - w^2 = 0. \end{cases}$$

In [WJG02], the authors compute the parameterization (17.14) with

$$\mathbf{X}_1(u, v) = \begin{pmatrix} -0.72u^3 - 0.72u^2v + 0.08uv^2 + 0.08v^3 \\ 0.0 \\ 0.72u^3 - 1.2u^2v - 0.72uv^2 - 0.08v^3 \\ 1.0182u^3 + 0.3394u^2v + 0.3394uv^2 + 0.1131v^3 \end{pmatrix}, \quad \mathbf{X}_2(u, v) = \begin{pmatrix} 0.0 \\ 1.697u + 0.5656v \\ 0.0 \\ 0.0 \end{pmatrix},$$

and $\Delta(u, v) = 0.48u^3v - 0.32u^2v^2 - 0.16uv^3$.

Our implementation gives the result displayed in Output 8. Since the polynomials of $\mathbf{X}(u, v)$ involve a square root $\sqrt{10}$, the quadric Q_R of inertia $(2, 2)$ used to parameterize the intersection is such that its determinant is not a square. As explained in Section 17.3.1, the parameterization is thus only near-optimal in the sense that it is possible, though not necessary, that the square root can be avoided in the coefficients.

It turns out that in this particular example it can be avoided. Consider the cone Q_R corresponding to the rational root $(-1, 21)$ of the determinantal equation :

$$Q_R : -Q_S + 21Q_T = 2x^2 - y^2 - w^2.$$

Q_R contains the obvious rational point $(1, 1, 0, 1)$, which is not its singular point. This implies that it can be rationally parameterized. Plugging this parameterization in the equation of Q_S or Q_T gives a simple parameterization

QI output 8 Execution trace for Example 2.

```

>> quadric 1: 19*x^2 + 22*y^2 + 21*z^2 - 20*w^2
>> quadric 2: x^2 + y^2 + z^2 - w^2

>> launching intersection
>> determinantal equation: - 175560*1^4 - 34358*1^3*m - 2519*1^2*m^2 - 82*1*m^3 - m^4
>> gcd of derivatives of determinantal equation: 1
>> number of real roots: 4
>> intervals: ]-14/2^8, -13/2^8[, ]-26/2^9, -25/2^9[, ]-25/2^9, -24/2^9[, ]-3/2^6, -2/2^6[
>> picked test point 1 at [ -13 256 ], sign > 0 -- inertia [ 2 2 ] found
>> picked test point 2 at [ -3 64 ], sign > 0 -- inertia [ 2 2 ] found
>> quadric (2,2) found: - 16*x^2 + 5*y^2 - 2*z^2 + 9*w^2
>> decomposition of its determinant [a,b] (det = a^2*b): [ 12 10 ]
>> a point on the quadric: [ 3 0 0 4 ]
>> param of quadric (2,2): [0, - 24*s*u - 24*t*v, 0, 0] + sqrt(10)*[3*t*u + 6*s*v, 0, 12*s*u
- 12*t*v, - 4*t*u + 8*s*v]
>> status of smooth quartic param: near-optimal
>> end of intersection

>> complex intersection: smooth quartic
>> real intersection: smooth quartic, two real bounded components
>> parameterization of smooth quartic, branch 1:
[(72*u^3 + 4*u*v^2)*sqrt(10) + 3*v*sqrt(10)*sqrt(Delta), - 340*u^2*v + 10*v^3 - 24*u*sqrt(Delta),
(- 118*u^2*v + 5*v^3)*sqrt(10) + 12*u*sqrt(10)*sqrt(Delta), (96*u^3 - 12*u*v^2)*sqrt(10)
- 4*v*sqrt(10)*sqrt(Delta)]
>> parameterization of smooth quartic, branch 2:
[(72*u^3 + 4*u*v^2)*sqrt(10) - 3*v*sqrt(10)*sqrt(Delta), - 340*u^2*v + 10*v^3 + 24*u*sqrt(Delta),
(- 118*u^2*v + 5*v^3)*sqrt(10) - 12*u*sqrt(10)*sqrt(Delta), (96*u^3 - 12*u*v^2)*sqrt(10)
+ 4*v*sqrt(10)*sqrt(Delta)]
>> Delta = 20*u^4 - 140*u^2*v^2 + 5*v^4
>> size of input: 2.3424, height of Delta: 1.3431

>> time spent: < 10 ms

```

of the intersection :

$$\mathbf{X}(u, v) = \begin{pmatrix} u^2 + 2v^2 \\ 2uv \\ u^2 - 2v^2 \\ 0 \end{pmatrix} \pm \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix} \sqrt{2u^4 + 4u^2v^2 + 8v^4}, \quad (u, v) \in \mathbb{P}^1(\mathbb{R}).$$

17.8.3 Example 3 : two tangent conics

Our next two examples illustrate the fact that our implementation is complete in the sense that it computes parameterizations in all possible cases.

Output 9 shows the execution trace for two quadrics intersecting in two conics that are tangent in one point. As can be seen, our implementation gives information about the incidence between the different components of the intersection : for each component, we give the parameter values (“cut parameters”) at which it intersects the other components of the intersection.

17.8.4 Example 4 : four skew lines

Our final example concerns an intersection made of four skew lines, as depicted in Figure 17.9.c. Output 10 shows the execution trace for this example, again illustrating the efficiency and completeness of our implementation.

QI output 9 Execution trace for Example 3.

```

>> quadric 1: - 4*x^2 - 56*x*y - 24*x*z - 79*y^2 - 116*y*z + 70*y*w - 85*z^2 - 20*z*w + 9*w^2
>> quadric 2: 6*x^2 + 84*x*y + 36*x*z + 45*y^2 + 160*y*z - 210*y*w + 131*z^2 + 30*z*w - 45*w^2

>> launching intersection
>> determinantal equation: 8*1^4 - 76*1^3*m + 234*1^2*m^2 - 297*1*m^3 + 135*m^4
>> gcd of derivatives of determinantal equation: 4*1^2 - 12*1*m + 9*m^2
>> triple real root: [ -3 -2 ]
>> inertia: [ 1 1 ]
>> rational point on cone: [ 0 0 0 1 ]
>> parameterization of cone with rational point
>> parameterization of pair of planes
>> the two conics are tangent at [ -39 3 6 -5 ]
>> status of intersection param: optimal
>> end of intersection

>> complex intersection: two tangent conics
>> real intersection: two tangent conics
>> parameterization of conic:
[- 39*u^2 + 443*u*v - 7254*v^2, 3*u^2 - 66*u*v + 1388*v^2, 6*u^2 - 132*u*v + 701*v^2, - 5*u^2
+ 110*u*v - 3005*v^2]
>> cut parameter: (u, v) = [1, 0]
>> size of input: 3.3222, height of output: 1.4631
>> parameterization of conic:
[- 39*u^2 + 443*u*v - 4004*v^2, 3*u^2 - 66*u*v + 1138*v^2, 6*u^2 - 132*u*v + 201*v^2, - 5*u^2
+ 110*u*v - 1205*v^2]
>> cut parameter: (u, v) = [1, 0]
>> size of input: 3.3222, height of output: 1.3854

>> time spent: 10 ms

```

17.9 Conclusion

We have presented a C++ implementation of an algorithm for parameterizing intersections of quadrics. The implementation is exact, efficient and covers all the possible cases of intersection. This implementation is based on the LiDIA library and uses the multiprecision integer arithmetic of GMP.

Future work will be devoted to understanding the gaps between predicted and observed values for the height of the coefficients of the parameterizations, to working out predicates and filters for making the code robust with floating point data (many classes and data structures have already been templated for a future use with floating point coefficients) and to porting our code to the CGAL geometry algorithms library [CGA].

Acknowledgments

The authors wish to acknowledge Laurent Dupont for a preliminary implementation of the parameterization algorithm in MuPAD, Guillaume Hanrot for his C implementation of Uspensky's algorithm, Daniel Lazard for his help in designing the parameterization algorithm, and Etienne Petitjean for his NetTask socket management tool which makes possible the querying of our parameterization software via a web interface.

QI output 10 Execution trace for Example 4.

```

>> quadric 1:  $199x^2 - 4xy + 830xz + 1068xw - 55y^2 - 278yz - 528yw + 587z^2$ 
+  $1146zw + 360w^2$ 
>> quadric 2:  $41x^2 - 64xy + 92xz + 108xw + 23y^2 - 32yz - 24yw + 80z^2$ 
+  $174zw + 72w^2$ 

>> launching intersection
>> determinantal equation:  $49l^4 - 84l^3m + 22l^2m^2 + 12lm^3 + m^4$ 
>> gcd of derivatives of determinantal equation:  $7l^2 - 6lm - m^2$ 
>> ranks of singular quadrics: 2 and 2
>> two real rational double roots: [ -1 -1 ] and [ -1 7 ]
>> status of intersection param: optimal
>> end of intersection

>> complex intersection: four skew lines
>> real intersection: four skew lines
>> parameterization of line:
[- 42*v, 32*u - 78*v, 28*u, - 25*u + v]
>> cut parameter: (u, v) = [- 19, 8]
>> cut parameter: (u, v) = [- 51, - 22]
>> size of input: 4.0592, height of output: 0.71248
>> parameterization of line:
[48*v, 64*u + 176*v, 68*u + 76*v, - 47*u - 69*v]
>> cut parameter: (u, v) = [0, 1]
>> cut parameter: (u, v) = [59, - 25]
>> size of input: 4.0592, height of output: 0.79955
>> parameterization of line:
[6*u, 6*u - 40*v, - 68*v, - 7*u + 111*v]
>> cut parameter: (u, v) = [49, 4]
>> cut parameter: (u, v) = [22, 3]
>> size of input: 4.0592, height of output: 0.75023
>> parameterization of line:
[- 12*v, 4*u, - 52*u - 60*v, 33*u + 41*v]
>> cut parameter: (u, v) = [67, - 49]
>> cut parameter: (u, v) = [39, - 25]
>> size of input: 4.0592, height of output: 0.68441

>> time spent: 10 ms

```

Chapitre 18

The Voronoi diagram of three lines

Une version préliminaire de cet article a été publiée dans les proceedings du *23th ACM Annual Symposium on Computational Geometry* [ELLD07].

Abstract

We give a complete description of the Voronoi diagram, in \mathbb{R}^3 , of three lines in general position, that is, that are pairwise skew and not all parallel to a common plane. In particular, we show that the topology of the Voronoi diagram is invariant for three such lines. The trisector consists of four unbounded branches of either a non-singular quartic or of a cubic and line that do not intersect in real space. Each cell of dimension two consists of two connected components on a hyperbolic paraboloid that are bounded, respectively, by three and one of the branches of the trisector. We introduce a proof technique, which relies heavily upon modern tools of computer algebra, and is of interest in its own right.

This characterization yields some fundamental properties of the Voronoi diagram of three lines. In particular, we present linear semi-algebraic tests for separating the two connected components of each two-dimensional Voronoi cell and for separating the four connected components of the trisector. This enables us to answer queries of the form, given a point, determine in which connected component of which cell it lies. We also show that the arcs of the trisector are monotonic in some direction. These properties imply that points on the trisector of three lines can be sorted along each branch using only linear semi-algebraic tests.

18.1 Introduction

The Voronoi diagram of a set of disjoint objects is a decomposition of the space into cells, one cell per object, such that the cell associated with an object consists of all points that are closer to that object than to any other object. In this paper, we consider the Voronoi diagram of lines in \mathbb{R}^3 under the Euclidean metric.

Voronoi diagrams have been the subject of a tremendous amount of research. For points, these diagrams and their complexities are well understood and optimal algorithms as well as robust and efficient implementations exist for computing them in any dimension (see for instance [Aur91, AK99, BDP⁺02, BDS⁺92, CSY97, CS89a, For97, OBSC00, PT06, Sei81]). Nevertheless, some important problems remain and are addressed in recent papers. The same is true for segments and polygons in two dimensions [Kar04].

For lines, segments, and polyhedra in three dimensions much less is known. In particular, determining the combinatorial complexity of the Voronoi diagram of n lines or line segments in \mathbb{R}^3 is an outstanding open problem. The best known lower bound is $\Omega(n^2)$ and the best upper bound is $O(n^{3+\epsilon})$ [Sha94]. It is conjectured that the complexity of such diagrams is near-quadratic. In the restricted case of a set of n lines with a fixed number, c , of possible orientations, Koltun and Sharir have shown an upper bound of $O(n^{2+\epsilon})$, for any $\epsilon > 0$ [KS03].

There are few algorithms for computing exactly the Voronoi diagram of linear objects. Most of this work has been done in the context of computing the medial axis of a polyhedron, *i.e.*, the Voronoi diagram of the faces

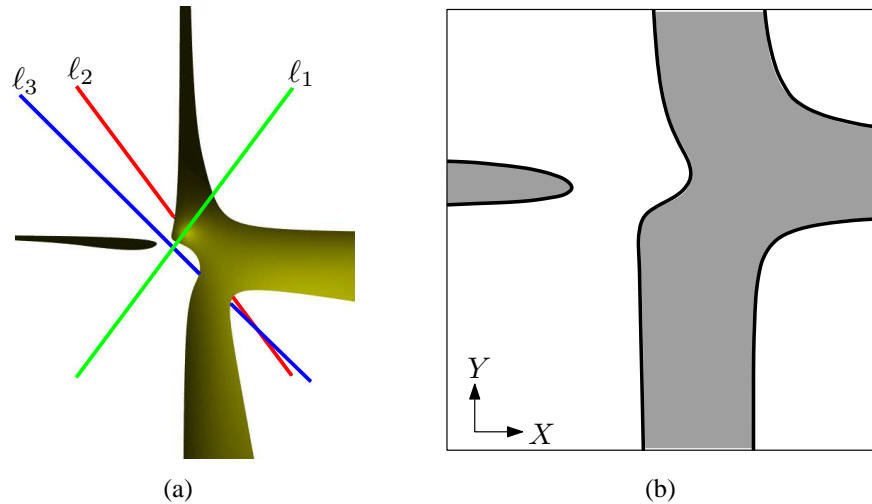


FIG. 18.1 – Voronoi diagram of 3 lines ℓ_1 , ℓ_2 , and ℓ_3 in general position : (a) Voronoi 2D face of ℓ_1 and ℓ_2 , i.e., set of points equidistant to ℓ_1 and ℓ_2 and closer to them than to ℓ_3 . (b) Orthogonal projection of a 2D face on a plane \mathcal{P} with coordinate system (X, Y) ; the plane's normal is parallel to the common perpendicular of ℓ_1 and ℓ_2 and the X and Y -axes are parallel to the two bisector lines (in \mathcal{P}) of the projection of ℓ_1 and ℓ_2 on \mathcal{P} . The 2D face is bounded by four branches of a non-singular quartic.

of the polyhedron [Cul00, Mil93]. Recently, some progress has been made on the related problem of computing arrangements of quadrics (each cell of the Voronoi diagram is a cell of such an arrangement) [BHK⁺05, KKM99, MTT05, SW06, SS97]. Finally, there have been many papers reporting algorithms for computing approximations of the Voronoi diagram (see for instance [DZ02, ER02, HCK⁺99, TT97]).

In this paper, we address the fundamental problem of understanding the structure of the Voronoi diagram of three lines. A robust and effective implementation of Voronoi diagrams of three-dimensional linear objects requires a complete and thorough treatment of the base cases, that is the diagrams of three and four lines, points or planes. We also believe that this is required in order to make progress on complexity issues, and in particular for proving tight worst-case bounds. We provide here a full and complete characterization of the geometry and topology of the elementary though difficult case of the Voronoi diagram of three lines in general position.

Main results. Our main result, which settles a conjecture of Koltun and Sharir [KS03], is the following (see Figure 18.1).

Theorem 18.1. *The topology of the Voronoi diagram of three pairwise skew lines that are not all parallel to a common plane is invariant. The trisector consists of four infinite branches of either a non-singular quartic²⁹ or of a cubic and a line that do not intersect in $\mathbb{P}^3(\mathbb{R})$. Each cell of dimension two consists of two connected components on a hyperbolic paraboloid that are bounded, respectively, by three and one of the branches of the trisector.*

We introduce, for the proof of Theorem 18.1, a new proof technique which relies heavily upon modern tools of computer algebra and which is of interest in its own right. We also provide a geometric characterization of the configurations of three lines in general position whose trisector is not generic, that is, consists of a cubic and a line.

²⁹By non-singular quartic, we mean an irreducible curve of degree four with no singular point in $\mathbb{P}^3(\mathbb{C})$. Recall that a point $p \in \mathbb{P}^3(\mathbb{C})$ of a surface S is said to be singular if its tangent plane is not defined at p , that is, all partial derivatives of the square-free polynomial defining S are zero at p . Similarly, a point $p \in \mathbb{P}^3(\mathbb{C})$ of a curve C defined by the two implicit equations $E_1 = E_2 = 0$ is singular if the rank of the Jacobian matrix of C (the matrix of partial derivatives of E_1 and E_2) is at most 1 when evaluated at p . (Note that the ideal generated by E_1 and E_2 should contain all the polynomials vanishing on C .) A curve is said to be singular if it contains at least a singular point in $\mathbb{P}^3(\mathbb{C})$. A curve is said to be singular in $\mathbb{P}^3(\mathbb{R})$ if it contains at least a singular point in $\mathbb{P}^3(\mathbb{R})$.

Theorem 18.2. *The trisector of three pairwise skew lines that are not all parallel to a common plane consists of a cubic and a line if and only if the hyperboloid of one sheet containing the three skew lines is of revolution.*

This work enables us to prove some fundamental properties of the Voronoi diagram of three lines which are likely to be critical for the analysis of the complexity and the development of efficient algorithms for computing Voronoi diagrams and medial axes of lines or polyhedra. In particular, we obtain the following results.

Monotonicity Property *Given three pairwise skew lines that are not all parallel to a common plane, there is a direction in which all four branches of the trisector are monotonic.*

Theorem 18.3. *Given a point p that lies on a two-dimensional cell of the Voronoi diagram of three pairwise skew lines that are not all parallel to a common plane, deciding on which connected component of the cell point p lies can be done by evaluating the sign of linear forms in the coordinates of p ; similarly, if p lies on the trisector. Furthermore, points on any one branch of the trisector may be ordered by comparing the values of a linear form in the coordinates of the points. Moreover, if the three input lines have rational coefficients, the coefficients of these linear forms may be chosen rational.*

Notice that these tests enable us to answer queries of the form, given a point, determine in which connected component of which cell it lies. Notice also that these tests should be useful for computing the Voronoi diagram of n lines since computing the vertices of such diagrams requires locating the points equidistant to four lines on a Voronoi arc of three of these lines.

The rest of the paper is organized as follows. We first study, in Section 18.2, the trisector of three lines in general position. We then present, in Section 18.3, some fundamental properties of the Voronoi diagram of three lines and prove the Monotonicity Property. We then prove Theorem 18.1 in Section 18.4 and Theorem 18.2 in Section 18.5. Finally, in Section 18.6, we present algorithms for separating the components of each cell of the Voronoi diagram and prove Theorem 18.3.

18.2 Structure of the trisector

We consider three lines in *general position*, that is, pairwise skew and not all parallel to the same plane. The idea of the proof of Theorem 18.1 is to prove that the topology of the trisector is invariant by continuous deformation on the set of all triplets of three lines in general position and that this set is connected. The result will then follow from the analysis of any example.

To prove that the topology of the trisector is invariant by continuous deformation on the set of all triplets of three lines in general position, we first show, in this section, that the trisector of three lines in general position is always homeomorphic to four lines that do not pairwise intersect. To prove this, we show that the trisector is always non-singular in $\mathbb{P}^3(\mathbb{R})$ and has four simple real points at infinity. To show that the trisector is always non-singular in $\mathbb{P}^3(\mathbb{R})$, we study the type of the intersection of two bisectors, which are hyperbolic paraboloids.

We use the classical result that the intersection of two quadrics is a non-singular quartic (in $\mathbb{P}^3(\mathbb{C})$) unless the characteristic equation of their pencil has (at least) a multiple root. In order to determine when this equation has a multiple root, we determine when its discriminant Δ is zero.

This discriminant has several factors, some of which are trivially always positive. We prove that the remaining, so-called “*gros facteur*”, is zero (over the reals) only if a (simple) polynomial F is zero. We provide two proofs of this result. We first give a short direct proof. Although this proof is elegant, it provides no insight into how we discovered the result. We also present a second proof which relies heavily upon sophisticated tools of modern algebra and does not require any detailed understanding of the geometry of the problem. This longer proof is indeed how we originally obtained Theorems 18.1 and 18.2 and only with the geometric insight gained from this process were we able to find the shorter proof. We believe this longer proof to be of interest in its own right because it demonstrates a technique which could be applied to other problems.

This proof goes as follows. We first show that the *gros facteur* is never negative using the RAGLIB Maple package [RAG]. This implies that it is zero only when all its partial derivatives are zero. We thus consider the system that consists of the *gros facteur* and all its partial derivatives, and compute its Gröbner basis. This gives three equations of degree six. We consider separately two components of solutions, one for which the aforementioned polynomial F is zero, the other for which $F \neq 0$. When $F \neq 0$, some manipulations and simplifications, which

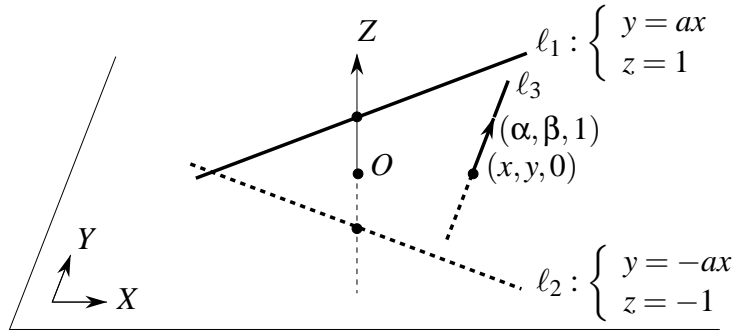


FIG. 18.2 – Three lines in general position.

are interesting in their own right, yield another Gröbner basis, with the same real roots, which consists of three equations of degree four. We show that one of these equations has no real root which implies that the system has no real root and thus that the *gros facteur* is strictly positive on the considered component. We can thus conclude that $\Delta = 0$ only if $F = 0$ and thus that, when $F \neq 0$, the trisector is always a non-singular quartic in $\mathbb{P}^3(\mathbb{R})$.

Then, when the polynomial $F = 0$, we show, by substituting $F = 0$ in Δ and by using the classification of the intersection of quadrics over the reals [DLLP07b], that the trisector is a cubic and a line that do not intersect in $\mathbb{P}^3(\mathbb{R})$.

We can thus conclude that the trisector is always a non-singular quartic or a cubic and a line that do not intersect in real space and thus that the trisector is always non-singular in $\mathbb{P}^3(\mathbb{R})$. We then prove that the trisector always contains four simple real points at infinity and thus that it is always homeomorphic to four lines that do not pairwise intersect.

18.2.1 Preliminaries

Let ℓ_1, ℓ_2 , and ℓ_3 be three lines in general position, *i.e.*, that are pairwise skew and not all parallel to a common plane. Refer to Figure 18.2. Let (X, Y, Z) denote a Cartesian coordinate system. Without loss of generality, we assume that ℓ_1 and ℓ_2 are both parallel to XY -plane, pass through $(0, 0, 1)$ and $(0, 0, -1)$ respectively, and have directions that are symmetric with respect to the XZ -plane. More precisely, we assume that the line ℓ_1 is defined by point $p_1 = (0, 0, 1)$ and vector $v_1 = (1, a, 0)$, and the line ℓ_2 is defined by the point $p_2 = (0, 0, -1)$ and vector $v_2 = (1, -a, 0)$, $a \in \mathbb{R}$. Moreover, since the three lines are not all parallel to a common plane, ℓ_3 is not parallel to the plane $z = 0$, and so we can assume that the line ℓ_3 is defined by point $p_3 = (x, y, 0)$ and vector $v_3 = (\alpha, \beta, 1)$, $x, y, \alpha, \beta \in \mathbb{R}$.

We denote by $\mathcal{H}_{i,j}$ the bisector of lines ℓ_i and ℓ_j and by V_{ij} the Voronoi cell of lines ℓ_i and ℓ_j , *i.e.*, the set of points equidistant to ℓ_i and ℓ_j and closer to them than to ℓ_k , $k \neq i, j$. We recall the following well-known elementary facts. The Voronoi cells are connected and the bisector of two skew lines is a right hyperbolic paraboloid, that is, has equation of the form $Z = \gamma XY$, $\gamma \in \mathbb{R}^*$, in some coordinate system (see for instance [KS03]); for completeness we present a proof of this fact.

Lemma 18.4. *The bisector of two skew lines is a right hyperbolic paraboloid.*

Proof. The bisector of two lines ℓ_i and ℓ_j is the set of points p satisfying the equation

$$\frac{\|(p - p_i) \times v_i\|^2}{\|v_i\|^2} = \frac{\|(p - p_j) \times v_j\|^2}{\|v_j\|^2}. \quad (18.1)$$

It suffices to prove the lemma for the two lines ℓ_1 and ℓ_2 . For these lines, the above equation simplifies into the following equation of a right hyperbolic paraboloid :

$$Z = -\frac{a}{1+a^2}XY. \quad (18.2)$$

□

The trisector of our three lines is the intersection of two right hyperbolic paraboloids, say $\mathcal{H}_{1,2}$ and $\mathcal{H}_{1,3}$. The intersection of two arbitrary hyperbolic paraboloids may be singular; it may be a nodal or cuspidal quartic, two secant conics, a cubic and a line that intersect, a conic and two lines crossing on the conic, etc. We show here that the trisector is always non-singular in $\mathbb{P}^3(\mathbb{R})$ by studying the characteristic polynomial of the pencil of $\mathcal{H}_{1,2}$ and $\mathcal{H}_{1,3}$.

Let $Q_{1,2}$ and $Q_{1,3}$ be matrix representations of $\mathcal{H}_{1,2}$ and $\mathcal{H}_{1,3}$, *i.e.* the Hessian of the quadratic form associated with the surface (see, for instance, [DLLP07a]). The *pencil* of $Q_{1,2}$ and $Q_{1,3}$ is the set of their linear combinations, that is, $P(\lambda) = \{\lambda Q_{1,2} + Q_{1,3}, \forall \lambda \in \mathbb{R} \cup \{\infty\}\}$. The *characteristic polynomial* of the pencil is the determinant, $\mathcal{D}(\lambda) = \det(P(\lambda))$, which is a degree four polynomial in λ . The intersection of any two quadrics is a non-singular quartic, in $\mathbb{P}^3(\mathbb{C})$, if and only if the characteristic equation of the corresponding pencil does not have any multiple roots (in \mathbb{C}) [Seg83] (see also [DLLP07b]). A non-singular quartic of $\mathbb{P}^3(\mathbb{C})$ is, in $\mathbb{P}^3(\mathbb{R})$, either empty or a non-singular quartic. Thus, since the trisector of our three lines cannot be the empty set in \mathbb{R}^3 , the trisector is a smooth quartic in $\mathbb{P}^3(\mathbb{R})$ if and only if the characteristic equation of the pencil does not have any multiple roots (in \mathbb{C}).

The characteristic polynomial of the pencil is fairly complicated (roughly one page in the format of Eq. (18.3)). However, by a change of variable $\lambda \rightarrow 2\lambda(1 + \alpha^2 + \beta^2)$ and by dividing out the positive factor $(1 + \alpha^2 + \beta^2)^3$, the polynomial simplifies, without changing its roots, to the following, which we still denote by $\mathcal{D}(\lambda)$ for simplicity.

$$\begin{aligned} \mathcal{D}(\lambda) = & (\alpha^2 + \beta^2 + 1)a^2\lambda^4 - 2a(2a\beta^2 + ay\beta + a\alpha x - \beta\alpha + 2a + 2a\alpha^2 - \beta\alpha\alpha^2)\lambda^3 \\ & + (\beta^2 + 6a^2\beta^2 - 2\beta xa^3 - 6\beta\alpha a^3 + 6y\beta a^2 - 6a\beta\alpha - 2a\beta x + 6\alpha xa^2 + y^2 a^2 - 2a\alpha y + x^2 a^2 - 2y\alpha a^3 + 6a^2\alpha^2 + a^4\alpha^2 + 4a^2)\lambda^2 \\ & - 2(xa - ya^2 - 2\beta a^2 - \beta + 2a\alpha + \alpha a^3)(xa - y - \beta + a\alpha)\lambda + (1 + a^2)(xa - y - \beta + a\alpha)^2 \end{aligned} \quad (18.3)$$

Let Δ be the discriminant of the characteristic polynomial $\mathcal{D}(\lambda)$ (with respect to λ). Recall that $\Delta = 0$ if and only if $\mathcal{D}(\lambda)$ admits a multiple root, that is, if and only if the trisector is not a smooth quartic. The discriminant Δ , computed with Maple [Map], is equal to

$$16a^4(ax - y - \beta + a\alpha)^2(y + ax - a\alpha - \beta)^2 \quad (18.4)$$

times a factor that we refer to as the *gros facteur* which is a rather large polynomial, of degree 18 in 5 variables with 253 monomials, of which we only show 2 out of 22 lines :

$$\begin{aligned} \text{gros_facteur} = & 8a^8\alpha^4y^2 + 7a^4\beta^2x^4 - 4a\beta^3x + 16a^8\beta^4x^4 + 32a^4\alpha^2y^2 + 2a^6\alpha^2\beta^4x^2 + 38a^8\alpha^2x^2 + 2y^4\beta^2a^4\alpha^2 + 44a^8\alpha^2\beta^2x^2 \\ & \dots + 22a^4y^2\beta^2x^2 + y^6a^6 + \alpha^2y^6a^6 - 2\beta x\alpha y^5a^6 + x^6a^6 + 10\beta x^3a^7\alpha^2 + 2y\alpha^3a^7x^2 - 32a^3\alpha^2y^2\beta x + 28a^3\beta^2x^2\alpha y - 24a^2\beta^3y\alpha x. \end{aligned} \quad (18.5)$$

In the sequel, all polynomials are considered over the reals, that is for $\lambda, a, \alpha, \beta, x, y$ in \mathbb{R} , unless specified otherwise.

18.2.2 The Main Lemma

We find in this section simple algebraic constraints that are satisfied when discriminant Δ is equal to zero. Precisely, we prove the following lemma.

Main Lemma *The discriminant Δ is equal to zero only if $y + a\alpha = 0$ or $ax + \beta = 0$.*

Note that the problem is to prove this lemma but also to obtain these two simple equations which is a difficult problem since Δ is a fairly large polynomial. As discussed in the overview of the proof, we first present a short direct proof of the Main Lemma.

Proof of the Main Lemma. Note first that the discriminant Δ is equal to zero if and only if the *gros facteur* is zero. Indeed, the polynomial (18.4) is not equal to zero under our general position assumption : $a = 0$ is equivalent

to saying that lines ℓ_1 and ℓ_2 are parallel and the two other factors of (18.4) are equal to the square of $\det(p_i - p_3, v_i, v_3)$, for $i = 1, 2$, and thus are equal to zero if and only if ℓ_i and ℓ_3 are coplanar, for $i = 1, 2$.

Now, it can be easily verified (using, for instance, Maple) that the *gros facteur* is, in fact, the discriminant of the characteristic polynomial of the 3×3 top-left submatrix of the matrix representation of the quadric containing ℓ_1, ℓ_2 and ℓ_3 (which is a hyperboloid of one sheet by the general position assumption);³⁰ this 3×3 submatrix corresponds to the quadratic part of the quadric and thus the discriminant is zero if and only if two eigenvalues are equal that is if the hyperboloid is of revolution (since a hyperboloid of one sheet has a canonical equation of the form $\frac{x^2}{\delta_1^2} + \frac{y^2}{\delta_2^2} - \frac{z^2}{\delta_3^2} - 1 = 0$). This directly proves that the *gros facteur* is zero if and only if the hyperboloid containing ℓ_1, ℓ_2 and ℓ_3 is of revolution. Furthermore, this is equivalent to the fact that trisector contains a line; indeed, if the hyperboloid is of revolution then its axis of revolution is on the trisector and, conversely, if the trisector contains a line, the *gros facteur* is zero (since the intersection of the two bisectors is not a non-singular quartic).

We can now prove the Main Lemma. Notice that if the hyperboloid containing ℓ_1, ℓ_2 and ℓ_3 is of revolution then its center of symmetry, O , is equidistant to the three lines. Point O can easily be computed as the intersection of the three planes P_1, P_2 , and P_3 where P_1 is the bisecting plane of ℓ_1 and the line parallel to ℓ_1 and transversal to ℓ_2 and ℓ_3 , and P_2 and P_3 are defined similarly (note that O is the center of the parallelepiped shown in Figure 18.3 and that O can also be easily computed as the point at which the gradient of the equation of the hyperboloid is zero). The constraint that point O is equidistant to lines ℓ_1 and ℓ_2 then reduces into $(y + a\alpha)(ax + \beta) = 0$, which concludes the proof. \square

Note that the above characterization of the *gros facteur* provides a direct proof of Lemma 18.5, which essentially states that the *gros facteur* is non-negative, because it is the discriminant of a polynomial whose roots are all real (since it is the characteristic polynomial of a real symmetric matrix). Alternatively, this also implies that the *gros facteur* is a sum of squares [Lax98] and thus non-negative. Note that we did not succeed to find even an approximation of this sum of square using SOSTOOLS [PS03, SOS].

We now present our original proof of the Main Lemma which relies upon modern tools of computer algebra and does not require any specific insight on the geometric meaning of the *gros facteur* and of the polynomials that appear in the Main Lemma.

Lemma 18.5. *The discriminant Δ is never negative.*

Proof. We prove that the real semi-algebraic set $\mathcal{S} = \{\chi = (a, x, y, \alpha, \beta) \in \mathbb{R}^5 \mid \Delta(\chi) < 0\}$ is empty using the RAGLIB Maple package [RAG] which is based on the algorithm presented in [SED06]. The algorithm computes at least one point per connected component of such a semi-algebraic set³¹ and we observe that, in our case, this set is empty. Before presenting our computation, we first describe the general idea of this algorithm.

Suppose first that $\mathcal{S} \neq \mathbb{R}^6$ and let \mathcal{C} denote any connected component of \mathcal{S} . We consider here Δ as a function of all its variables $\chi = (a, x, y, \alpha, \beta) \in \mathbb{R}^6$. The algorithm first computes the set of generalized critical values³² of Δ (see [SED06] for an algorithm computing them). The image by Δ of \mathcal{C} is an interval whose endpoints³³ are zero and either a negative generalized critical value or $-\infty$. For any v in this interval, there is a point $\chi_0 \in \mathcal{C}$ such that $\Delta(\chi_0) = v$, and the connected component containing χ_0 of the hypersurface $\Delta(\chi) = v$ is included in the connected

³⁰The equation of the hyperboloid containing ℓ_1, ℓ_2 and ℓ_3 can easily be computed by solving a linear system obtained by writing that three points on each of the three lines lie on the quadric.

³¹Note that no certified polynomial-time algorithm (in the number of variables) is known for this problem.

³²Recall that the (real) critical values of Δ are the values of Δ at its critical points χ , *i.e.*, the points χ at which the gradient of Δ is zero. The asymptotic critical values are similarly defined as, roughly speaking, the values taken by Δ at critical points at infinity, that is, the values $c \in \mathbb{R}$ such that the hyperplane $z = c$ is tangent to the surface $z = \Delta(\chi)$ at infinity (this definition however only holds for two variables, *i.e.*, $\chi \in \mathbb{R}^2$). More formally, the asymptotic critical values were introduced by Kurdyka et al. [KOS00] as the limits of $\Delta(\chi_k)$ where $(\chi_k)_{k \in \mathbb{N}}$ is a sequence of points that goes to infinity while $\|\chi_k\| \cdot \|\text{grad}_{\chi_k} \Delta(\chi_k)\|$ tends to zero. The generalized critical values are the critical values and asymptotic critical values. *The set of generalized critical values contains all the extrema of function \mathcal{D} , even those that are reached at infinity.*

³³Since $\mathcal{S} \neq \mathbb{R}^6$, the boundary of \mathcal{C} is not empty and consists of points χ such that $\mathcal{D}(\chi) = 0$. The image of the connected set \mathcal{C} by the continuous function \mathcal{D} is an interval. Hence, zero is an endpoint of the interval $\mathcal{D}(\mathcal{C})$. The other endpoint is either an extremum of \mathcal{D} (and thus a generalized critical value) or $-\infty$.

component \mathcal{C} . Hence, a point in \mathcal{C} can be found by computing a point in each connected component of $\Delta(\chi) = v$. It follows that we can compute at least a point in every connected component of the semi-algebraic set \mathcal{S} defined by $\Delta(\chi) < 0$ by computing at least one point in every connected component of the real hypersurface defined by $\Delta(\chi) = v$ where v is any value smaller than zero and larger than the largest negative generalized critical value, if any. Now, when $\mathcal{S} = \mathbb{R}^6$, that is, $\Delta(p) < 0$ for all p in \mathbb{R}^6 , the above computation returns an empty set of points, so we choose a random point p in \mathbb{R}^6 and return it if $\Delta(p) < 0$.

Notice that computing at least one point in every connected component of a hypersurface defined by $\Delta(\chi) = v$ can be done by computing the critical points of the distance function between the surface and a point, say the origin, that is, by solving the system $\Delta(\chi) = v, \chi \times \text{grad}(\Delta)(\chi) = 0$. This conceptually simple approach, developed in [RRSED00], is, however, not computationally efficient. The efficient algorithm presented in [SED06] computes instead critical points of projections, combining efficiently the strategies given in [SEDS04] and [SEDS03].

The computation of at least one point in every connected component of \mathcal{S} , using the RAGLIB Maple package, gives the empty set, implying that $\Delta(\chi) \geq 0$ for all $\chi \in \mathbb{R}^6$. It should be noted that these computations are time consuming on a polynomial of the size of Δ : they take roughly 10 hours of elapsed time on a standard PC. \square

We now prove that the zeros of Δ are the singular points³⁴ of the *gros facteur*.

Lemma 18.6. *The discriminant Δ is equal to zero if and only if the gros facteur and all its partial derivatives are equal to zero.*

Proof. As we have seen in the direct proof of the Main Lemma, the discriminant Δ is equal to zero if and only if the *gros facteur* is zero. Furthermore, by Lemma 18.5, the *gros facteur* is never negative, thus, if there exists a point where the *gros facteur* vanishes, it is a local minimum of the *gros facteur* and thus all its partial derivatives (with respect to $\{a, x, y, \alpha, \beta\}$) are zero. \square

We now present a simple and direct computational proof of the Main Lemma. As we will see, this proof is, however, based on some polynomials whose origins are discussed in Section 18.2.3.

Computational proof of the Main Lemma. By Lemma 18.6, Δ is zero if and only if the *gros facteur* and all its partial derivatives are zero. We prove below that this implies that $(y + a\alpha)(ax + \beta)(1 + \alpha^2 + \beta^2)\Gamma = 0$, where

$$\Gamma = (2a(y\alpha - \beta x) - a^2 + 1)^2 + 3(ax + \beta)^2 + 3a^2(y + a\alpha)^2 + 3(1 + a^2)^2. \quad (18.6)$$

As the two terms $(1 + \alpha^2 + \beta^2)$ and Γ clearly do not have any real solutions, this proves the lemma. (We discuss later how we found these terms.)

Consider the system in the variables $\{a, x, y, \alpha, \beta, u, v, w, t\}$ that consists of the *gros facteur*, its partial derivatives, and the four equations

$$1 - u(y + a\alpha) = 0, \quad 1 - v(ax + \beta) = 0, \quad 1 - w(1 + \alpha^2 + \beta^2) = 0, \quad 1 - t\Gamma = 0. \quad (18.7)$$

The *gros facteur* and its partial derivatives have a common zero (real or complex) such that $(y + a\alpha)(ax + \beta)(1 + \alpha^2 + \beta^2)\Gamma \neq 0$ if and only if this system has a solution. This follows immediately from the fact that the equations (18.7) are linear in u, v, w, t .

The Gröbner basis of our system is reduced to the polynomial 1 (see Table 18.1) and thus the system has no solution (over the complex numbers). This concludes the proof. \square

The real difficulty in this proof of the Main Lemma is, of course, to find the equations (18.7) that rule out all the components of the set of singular points of the *gros facteur*. Computing these components is the actual key of the computational proof. We believe that the technique we used can be of some interest to the community as it is rather generic and could be applied to other problems. We thus describe in Section 18.2.3 how these components were computed before finishing the study of the algebraic structure of the trisector, in Section 18.2.4.

³⁴Recall that the singular points of a surface are the points where all partial derivatives are zero.

```

> Gamma:=(2*a*(y*alpha-x*beta)-(a^2-1))^2+3*(a*x+beta)^2+3*a^2*(y+a*alpha)^2+3*(a^2+1)^2;
      Γ := (2a(αy - βx) - a2 + 1)2 + 3(xa + β)2 + 3a2(y + aα)2 + 3(1 + a2)2
> [gros_fact, op(convert(grad(gros_fact,[a,x,y,alpha,beta]),list)),
> 1-u*(y+a*alpha), 1-v*(a*x+beta),1-w*(1+alpha^2+beta^2),1-t*Gamma)];
> fgb_gbasis_elim(% ,0,[u,v,w,t],[a,x,y,alpha,beta]);

pack_fgb_call_generic: "FGb: 965.76 sec Maple: 975.98 sec"
[1]

```

TAB. 18.1 – For the proof of the Main Lemma.

18.2.3 About the computational proof of the Main Lemma

We show in this section how we computed, for the proof of the Main Lemma, the equations of (18.7) which correspond to hypersurfaces containing the zeros of the discriminant.

We proceed as follows. We start from the system of equations consisting of the *gros facteur* and all its partial derivatives and use the following techniques to study its set of solutions, or, more precisely, to decompose it into components defined by prime ideals³⁵. This could theoretically be done by a general algorithm computing such a decomposition, however, no currently available software is capable of handling our particular problem and this is, indeed, a significant research challenge in computer algebra.

If the (reduced) Gröbner basis of some system contains a polynomial which has a factor, say F , the solutions of the system splits into two components, one of which such that $F = 0$, the other such that $F \neq 0$. We study separately the two components. One is obtained by adding the equation F to the system and the other is obtained by adding the equation $1 - tF$ and eliminating the variable t ; indeed, there is a one-to-one correspondence between the solutions of the initial system such that $F \neq 0$ and the solutions of the system augmented by $1 - tF$. Sometimes, frequently in our case, the component $F \neq 0$ is empty, which corresponds to the situation where the elimination of t results in the polynomial 1 (inducing the equation $1 = 0$). Note that in some cases the system contains a polynomial which is a square, say F^2 , thus the component such that $F \neq 0$ is obviously empty and we can add F to the system without changing its set of solutions (this however changes the ideal). This operation of adding F to the system frequently adds embedded components to the variety of solutions which explains why, later on in the process, empty components are frequently encountered when splitting into two components.

Our computations, presented in Table 18.2 in the appendix, are performed in Maple [Map] using the Gröbner basis package FGb developed by J.-C. Faugère [FGb]. We use two functions,

$fgb_gbasis(sys,0,vars1,vars2)$ and $fgb_gbasis_elim(sys,0,var1,var2)$ ³⁶,

that compute Gröbner bases of the system sys ; the first uses a degree reverse lexicographic order (DRL) by blocks on the variables of $vars1$ and $vars2$ (where $vars2$ is always the empty set in our computation) and the second one eliminates the variable $vars1$ and uses a reverse lexicographic order on the variables of $vars2$. (The second parameter of the functions refer to the characteristic of the field, here 0.)

We do not show in Table 18.2 the Gröbner bases which are too large to be useful, except in the case where the basis is reduced to 1 (when the system has no solution). We instead only report the first operand of each polynomial of the base; an operand $*$ means that the polynomial is the product of at least two factors; an operand $^$ means that the polynomial is a power of some polynomial; an operand $+$ means that the polynomial is a sum of monomials.

Our computation goes as follows. We first simplify our system by considering $a = 2$ because otherwise the Gröbner basis computations are too slow and use too much memory to be performed successfully. We first see after computing, bs_1 , the Gröbner basis of our system, that $y + 2\alpha$ appears as a factor of one polynomial. This splits the solutions into those such that $y + 2\alpha = 0$ and the others. We will study separately (in Lemma 18.8) the former set of solutions and we only consider here the solutions such that $y + 2\alpha \neq 0$. This is done by adding the polynomial $1 - u(y + 2\alpha)$ to the system, where u is a new variable; indeed there is a one-to-one correspondence between the solutions of the initial system such that $y + 2\alpha \neq 0$ and the solutions of the resulting system.

³⁵An ideal J is prime if $PQ \in J$ implies $P \in J$ or $Q \in J$.

³⁶The function $gbasis(sys,DRL(var1,var2),elim)$ with or without the optional last argument $elim$ can also be used alternatively of these two functions

The term $y + 2\alpha$ corresponds fairly clearly to the polynomial $y + a\alpha$ with $a = 2$, and because of the symmetry of our problem we also study separately the solutions such that $ax + \beta = 0$. Since we assumed $a = 2$, we only consider here the solutions such that $2x + \beta \neq 0$, by adding to the system the polynomial $1 - v(2, x + \beta)$. Finally, we also add $1 - w(1 + \alpha^2 + \beta^2)$ to the system, without changing its set of real roots; we do this because the term $1 + \alpha^2 + \beta^2$ appears in the leading coefficient of $\mathcal{D}(\lambda)$ which suggests that some component of solutions (without any real point) might be included in $1 + \alpha^2 + \beta^2$ (it should be noted that adding this polynomial to the system changes the resulting Gröbner basis, which shows that this indeed removes some imaginary component from the system). We compute the Gröbner basis, bs_2 , of that system, eliminating the variables u, v, w , which gives a system of four polynomials of degree six.

We then compute the Gröbner basis of bs_2 , eliminating the variable x . This gives a basis bs_3 which is reduced to one polynomial of the form P^2 . We thus add P to the system bs_2 (we do not add it to bs_3 since bs_3 does not depend on x). The Gröbner basis, bs_4 , of the new system contains several polynomials that are products of factors. We see that if we add to the system the constraint that the third factor of the first polynomial is not zero, the resulting system has no solution. We thus add this factor to the system and compute its Gröbner basis bs_5 . We operate similarly to get bs_6 . The basis bs_6 contains no product or power and we compute its Gröbner basis, bs_7 , eliminating y (eliminating x gives no interesting basis). The last polynomial of bs_7 is a power and we proceed as before to get bs_8 . We proceed similarly until we get to the basis bs_{12} . (Note that the factor $y + 2\alpha$ reappears in bs_{10} and is removed similarly as in the beginning of the process.)

The basis bs_{12} consists of three polynomials of degree four (which is a simplification over bs_2 which consists of four polynomials of degree six). We observe that the last polynomial of bs_{12} is

$$\Gamma_2 = (4y\alpha - 4\beta x - 3)^2 + 3(2x + \beta)^2 + 12(y + 2\alpha)^2 + 75,$$

which is always positive over the reals.

We have thus proved that all the complex solutions, such that $a = 2$, of the initial system (the *gros facteur* and all its partial derivatives) satisfy $(1 + \alpha^2 + \beta^2)(y + 2\alpha)(2x + \beta)\Gamma_2 = 0$.

Finally, to get the polynomial Γ of Formula (18.6), we performed the same computation with $a = 3$ and $a = 5$ and *guessed* Γ as an interpolation of the polynomials Γ_2, Γ_3 , and Γ_5 .

Note that all the computation for a fixed a takes roughly eight minutes of elapsed time on a regular PC.

Remark 18.7. *All the computations from bs_2 to bs_{12} amounts to finding polynomials that have a power which is a combination of the elements of bs_2 (i.e., which are in the radical of the ideal generated by bs_2 ³⁷). Thus these computations would be advantageously replaced by a program computing the radical of an ideal. Unfortunately, all available such programs fail on the ideal generated by bs_2 either by exhausting the memory or by running unsuccessfully during several days and ending on an error. It is therefore a challenge to improve these programs in order to do this computation automatically.*

We now present another much faster technique to compute bs_{12} , which takes advantage of the structure of bs_2 .

Recall that bs_2 is a Gröbner basis consisting in 4 polynomials of degree 6 (see Table 18.2) and refer to Table 18.3. The Gröbner basis of bs_2 for a block ordering with x in the first block consists of 31 polynomials of degree at least 2 in x , 32 polynomials linear in x and one polynomial, which is independent of x . The latter is a square, P^2 . Let $Q = Rx + S$ be the last linear polynomial of the basis. Clearly, any solution of the system is a common zero of P and Q . Conversely, one may guess that any common zero of P and Q for which $R \neq 0$ is a solution of the system (see [ALMM99]) and we prove this is effectively the case.

We compute the Gröbner basis eliminating t in the system $P, Q, 1 - tR$. This basis consists of 3 polynomials of degree 4, and is equal to bs_{12} . Then we prove that the two systems have the same solutions by showing that the elements of bs_2 are in the ideal generated by bs_{12} and that the square of the elements of bs_{12} are in the ideal generated by bs_2 . This is done by using the function *normalf* which computes the normal form of a polynomial with respect to a Gröbner basis. All these computations need less than eight seconds, instead of eight minutes for the previous method.

Another advantage of this new method is that it shows directly that the ideal generated by bs_{12} is prime³⁵. Indeed, for any polynomial, say F in the ideal, its pseudo-remainder³⁸ by Q (with respect to x) is a multiple of P

³⁷The radical of an ideal \mathcal{J} is the ideal $\{P \mid P^n \in \mathcal{J} \text{ for some } n \in \mathbb{N}\}$.

³⁸Here, the pseudo-remainder of F by Q is the numerator of the expression obtained by substituting x by $-S/R$ in F .

(see, for instance, [ALMM99]). If F is a product, its pseudo-remainder is the product of the pseudo-remainders of the factors. Thus P , which is irreducible, divides one of them, which shows that one of the factors of F is in the ideal, that is that the ideal is prime.

18.2.4 Structure of the trisector : conclusion

We proved in the Main Lemma that the discriminant Δ is equal to zero only if $y + a\alpha = 0$ or $ax + \beta = 0$. We prove in this section that if $\Delta = 0$, the trisector is a cubic and a line that do not intersect. We then show that the trisector always contains four simple real points at infinity and conclude that the trisector is always homeomorphic to four lines that do not pairwise intersect.

Lemma 18.8. *The discriminant Δ is equal to zero if and only if*

$$y = -a\alpha \quad \text{and} \quad x = \frac{\beta(2a^2 + 1) \pm 2\sqrt{a^2(1+a^2)(\alpha^2 + \beta^2 + 1)}}{a}, \quad \text{or} \quad (18.8)$$

$$x = -\frac{\beta}{a} \quad \text{and} \quad y = \frac{\alpha(2+a^2) \pm 2\sqrt{(1+a^2)(\alpha^2 + \beta^2 + 1)}}{a}. \quad (18.9)$$

Proof. We refer to Table 18.4, in the appendix, for the computations. By the Main Lemma, $\Delta = 0$ implies $y + a\alpha = 0$ or $ax + \beta = 0$. Substituting y by $-a\alpha$ in Δ gives an expression of the form $f_0 f_1^2$. Similarly, substituting x by $-\beta/a$ in Δ gives an expression of the form $g_0 g_1^2$ (recall that $a \neq 0$ since the lines are not coplanar, by assumption). It follows that $\Delta = 0$ if and only if $y + a\alpha = f_i = 0$ or $ax + \beta = g_i = 0$, for $i = 0$ or 1 .

The f_i and g_i are polynomials of degree two in x and y , respectively. Solving $f_1 = 0$ in terms of x directly yields that the system

$$y + a\alpha = f_1 = 0 \quad (18.10)$$

is equivalent to (18.8). Similarly, solving $g_1 = 0$ in terms of y yields that the system

$$ax + \beta = g_1 = 0 \quad (18.11)$$

is equivalent to (18.9).

We now show that the solutions of $y + a\alpha = f_0 = 0$ are included in the set of solutions of (18.9). The polynomial f_0 is the sum of two squares. It follows that $y + a\alpha = f_0 = 0$ if and only if

$$y + a\alpha = a^2\alpha^2 - 1 + a\beta x = ax + \beta = 0. \quad (18.12)$$

We show below that the polynomials of (18.11) are included in the ideal generated by the polynomials of (18.12). This implies that (18.11) is, roughly speaking, less constrained than (18.12) and that the set of solutions of (18.11) contains the solutions of (18.12). Hence the solutions of $y + a\alpha = f_0 = 0$ are contained in the set of solutions of (18.11) and thus in the set of solutions of (18.9).

We prove that the polynomials of (18.11) are included in the ideal generated by the polynomials of (18.12) by showing that the normal form of every polynomial of (18.11) with respect to the Gröbner basis of the polynomials of (18.12) is zero. This is done using the function *normalf* (of Maple) which computes the normal form of a polynomial with respect to a Gröbner basis..

We prove similarly that the solutions of $ax + \beta = g_0 = 0$ are included in the set of solutions of (18.10) and thus of (18.8), which concludes the proof. \square

Remark 18.9. *Note that by symmetry with respect to the XY -plane and by changing the sign of a, α , and β , the set of three input lines ℓ_1, ℓ_2, ℓ_3 is invariant, the two components of (18.8) exchange (i.e., the components corresponding to $+2\sqrt{}$ and $-2\sqrt{}$ exchange), and the two components of (18.9) exchange.*

Similarly, by exchanging the X and Y -coordinates, x and y , α and β , and changing a into $1/a$, the set of three input lines is also invariant and each component of (18.8) is changed to a component of (18.9), and conversely.

Lemma 18.10. *If $\Delta = 0$, the trisector of ℓ_1, ℓ_2 , and ℓ_3 consists of a cubic and a line that do not intersect in real space.*

Proof. By Lemma 18.8, $\Delta = 0$ if and only if System (18.8) or (18.9) is satisfied. By symmetry of the problem (see Remark 18.9), we only need to consider one of the components of (18.8) and (18.9). Hence, it is sufficient to show that

$$y = -a\alpha, \quad x = \frac{\beta(2a^2 + 1)}{a} + 2\sqrt{(1 + a^2)(\alpha^2 + \beta^2 + 1)} \quad (18.13)$$

implies that the trisector consists of a cubic and a line that do not intersect. We assume in the following that $\Delta = 0$, that System (18.13) is satisfied. We refer to Table 18.5 for the computations.

We first show that the characteristic polynomial of the pencil generated by the bisectors is always strictly positive. Note first that the characteristic polynomial is not always negative (see [Lev76]). It is thus sufficient to prove that it is never zero, or equivalently, that its product with its algebraic conjugate (obtained by changing the sign of $\sqrt{(1 + a^2)(\alpha^2 + \beta^2 + 1)}$) is never zero. This product is a polynomial T in $a, \alpha, \beta, \lambda$ which can easily be factored in the square of a degree-two polynomial in λ ; furthermore, this degree two polynomial has no real root because its discriminant is the product of a negative term ($-(1 + a^2)^2(1 + \alpha^2 + \beta^2)$) and a term whose sum and product with its algebraic conjugate (obtained, as above, by changing the sign of the square root) is a strictly positive sum of squares. Note that we can also prove that T is always strictly positive by computing, similarly as in the proof of Lemma 18.5, at least one point per connected component of the real semi-algebraic set $\{\chi = (a, \alpha, \beta, \lambda) \in \mathbb{R}^4 \mid T(\chi) - \frac{1}{2} < 0\}$; the resulting set of points is empty, hence $T(\chi)$ is always greater or equal to $1/2$. It thus follows that the characteristic polynomial of the pencil is always strictly positive.

Since the characteristic polynomial $\mathcal{D}(\lambda)$ is always strictly positive and its discriminant Δ is zero, $\mathcal{D}(\lambda)$ admits two (conjugate) double imaginary roots. Let λ_1 and λ_2 denote these two roots. Recall that $\mathcal{D}(\lambda) = \det P(\lambda)$ with $P(\lambda) = \lambda Q_{1,2} + Q_{1,3}$ where $Q_{i,j}$ is the matrix associated with the hyperbolic paraboloid $\mathcal{H}_{i,j}$. It follows from the classification of the intersection of quadrics [DLLP07b, Table 4] that either (i) $P(\lambda_1)$ and $P(\lambda_2)$ are of rank 3 and the trisector $\mathcal{H}_{1,2} \cap \mathcal{H}_{1,3}$ consists of a cubic and a line that do not intersect or (ii) $P(\lambda_1)$ and $P(\lambda_2)$ are of rank 2 and the trisector consists of two secant lines.

We now prove that $P(\lambda_1)$ and $P(\lambda_2)$ are of rank 3. We compute the Gröbner basis of all the 3×3 minors of $P(\lambda)$ and of the polynomial $1 - t\Psi$ with

$$\Psi = (1 + a^2)(1 + \alpha^2 + \beta^2)(ax - y - \beta + a\alpha)(y + ax - a\alpha - \beta).$$

The basis is equal to 1, thus the 3×3 minors of $P(\lambda)$ are not all simultaneously equal to zero when $\Psi \neq 0$. Furthermore, $\Psi \neq 0$ for any x, y, a, α, β in \mathbb{R} such that the lines ℓ_1, ℓ_2 , and ℓ_3 are pairwise skew (see (18.4) and the proof of Lemma 18.6). Thus the rank of $P(\lambda)$ is at least 3. The rank of $P(\lambda_i)$, $i = 1, 2$, is thus equal to 3 since $\det P(\lambda_i) = 0$. We can thus conclude that when $\Delta = 0$ the trisector consists of a cubic and a line that do not intersect in real space. \square

We now state a proposition that shows that the trisector admits four asymptotes that are pairwise skew and gives a geometric characterization of their directions.

Proposition 18.11. *The trisector of ℓ_1, ℓ_2 , and ℓ_3 intersects the plane at infinity in four real simple points. Furthermore, the four corresponding asymptotes are parallel to the four trisector lines of three concurrent lines that are parallel to ℓ_1, ℓ_2 , and ℓ_3 , respectively.*

Proof. The trisector is the intersection of two hyperbolic paraboloids. Any hyperbolic paraboloid contains two lines at infinity. Hence the intersection, at infinity, of any two distinct hyperbolic paraboloids is the intersection of two pairs of lines. The intersection of these two pairs of lines consists of exactly four simple real points unless the point of intersection of the two lines in one pair lies on one line of the other pair. We show that this cannot happen under our assumptions.

The intersection with the plane at infinity of the bisector of lines ℓ_1 and ℓ_2 consists of the lines at infinity in the pair of planes of equation $XY = 0$ (the homogeneous part of highest degree in Eq. (18.2)). This pair of plane is the bisector of the two concurrent lines that are parallel to ℓ_1 and ℓ_2 , respectively. Note that the lines at infinity in this pair of planes are invariant by translation of the planes. We thus get that the lines at infinity of the bisector of any two lines ℓ_i and ℓ_j are the lines at infinity in the pair of planes that is the bisector to any two concurrent lines that are parallel to ℓ_i and ℓ_j , respectively.

It follows that the points at infinity on the trisector of ℓ_1, ℓ_2 , and ℓ_3 are the points at infinity on the trisector lines (the intersection of bisector planes) of three concurrent lines that are parallel to ℓ_1, ℓ_2 , and ℓ_3 , respectively. It remains to show that this trisector consists of four distinct lines.

Let $\ell'_1, \ell'_2,$ and ℓ'_3 be the three concurrent lines through the origin that are parallel to $\ell_1, \ell_2,$ and $\ell_3,$ respectively, and suppose, for a contradiction, that their trisector does not consist of four distinct lines. This implies that the line of intersection of the two bisector planes of two lines, say ℓ'_1 and $\ell'_2,$ is contained in one of the bisector planes of two other lines, say ℓ'_1 and $\ell'_3.$ The intersection of the bisector planes of ℓ'_1 and ℓ'_2 is the Z -axis. It follows that one of the bisector planes of ℓ'_1 and ℓ'_3 is vertical, hence ℓ'_1 and ℓ'_3 are symmetric with respect to a vertical plane and thus ℓ'_3 is contained in the XY -plane. Therefore, $\ell'_1, \ell'_2,$ and ℓ'_3 lie in the XY -plane, contradicting the general position assumption, which concludes the proof. \square

Theorem 18.12. *The trisector of three lines in general position consists of four infinite smooth branches of a non-singular quartic or of a cubic and a line that do not intersect in real space.*

Proof. As mentioned in the beginning of Section 18.2.2, the trisector of three lines consists of a smooth quartic unless the discriminant Δ is zero. Lemma 18.10 and Proposition 18.11 thus yield the result. \square

18.3 Properties of the Voronoi diagram

We present here some fundamental properties of the Voronoi diagram. We will show how the four branches of the trisector of three lines can be labeled and then present two fundamental properties of the trisector.

18.3.1 Preliminaries

We start with the following important proposition.

Proposition 18.13. *The set of triplets of lines in general position is connected.*

Proof. We prove this proposition by proving that there is a one-to-one correspondence between the set of ordered triplets of lines (in general position) and the set of affine frames of positive orientation.

Consider three lines $\ell_1, \ell_2,$ and ℓ_3 in general position and refer to Figure 18.3. For the three choices of pairs of lines $\ell_i, \ell_j,$ consider the plane containing ℓ_i and parallel to $\ell_j,$ the plane containing ℓ_j and parallel to $\ell_i,$ and the region bounded by these two parallel planes. The general position assumption implies that these regions have non-empty interiors and that no three planes are parallel. The intersection of these three regions thus defines a parallelepiped. By construction, each of the lines $\ell_1, \ell_2,$ and ℓ_3 contains an edge of that parallelepiped. These lines are pairwise skew thus exactly two vertices of the parallelepiped are not on the lines. Each of these two points induces an affine frame centered at the point and with basis the three edges of the parallelepiped oriented from the point to the lines $\ell_1, \ell_2,$ and $\ell_3,$ in this order. One of the point (C on the figure) defines a frame of positive orientation, the other defines a frame of negative orientation (C' on the figure). This construction exhibits a one-to-one correspondence between the set of ordered triplets of lines (in general position) and the set of affine frames of positive orientation, which concludes the proof. \square

We consider in the following any three lines $\ell_1, \ell_2,$ and ℓ_3 in general position (pairwise skew and not all parallel to a common plane) and an associated Cartesian coordinate system (X, Y, Z) such that the Z -axis is the common perpendicular of ℓ_1 and $\ell_2,$ the origin is the point on the Z -axis equidistant to ℓ_1 and $\ell_2,$ and such that the X and Y -axes are the two bisector lines, in the plane through the origin and perpendicular to the Z -axis, of the projection of ℓ_1 and ℓ_2 onto this plane.³⁹ Note that the orientations of the axes are not specified (except for the fact that the

³⁹Note that this setting is slightly different than the one described in Section 18.2.1 since, here, any triplet of three lines in general position can be moved continuously into another while the associated frame moves continuously; however, if the initial and final triplets of lines are in the setting of Section 18.2.1, it is not necessarily possible to ensure that, during the motion, all triplets of lines remain in this setting. This is, for instance, the case for the two triplets of lines $(y = x, z = 1; y = -x, z = -1; x = 1, y = 0)$ and $(y = -x, z = 1; y = x, z = -1; x = 0, y = 1)$ for which one triplet can be obtained from the other by a rotation of the frame about the Z -axis (by an angle $\pm\pi/2$) though the triplets of lines cannot be moved continuously from one configuration to the other while remaining (pairwise skew) in the setting of Section 18.2.1.

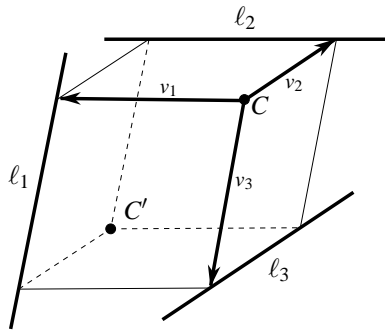


FIG. 18.3 – The parallelepiped formed by ℓ_1 , ℓ_2 , and ℓ_3 and the associated frame (C, w_1, w_2, w_3) of positive orientation.

frame has a positive orientation) and that the X and Y -axes can be exchanged.

18.3.2 Labeling of the four branches of the trisector

We prove here the following proposition which states two properties, one on the asymptotes of the trisector and one on the incidence relations between cells, which, together, yield an unambiguous labeling of the components of the trisector.

Let V_{ij} denote the two-dimensional Voronoi cell of lines ℓ_i and ℓ_j and let U_{ij} and T_{ij} denote the connected components of V_{ij} that are bounded by one and three arcs of the trisector, respectively (see Figure 18.4(a)).

Proposition 18.14. *Exactly one of the four branches of the trisector of three lines in general position admits only one asymptote. Let C_0 denote this branch. Each cell U_{ij} is bounded by a branch distinct from C_0 and every such branch bounds a cell U_{ij} . Let C_k , $k = 1, 2, 3$, denote the branches of the trisector that bound the component U_{ij} , $i, j \neq k$. The labeling of the four branches of the trisector by C_0, \dots, C_4 is unambiguous.*

Note that differentiating between C_1 and C_2 cannot be done, as far as we know, by only looking at the cell V_{12} (see Figure 18.4(a)) but can be done by looking at the other cells V_{13} and V_{23} . More precisely, differentiating between C_1 and C_2 on Figure 18.4(a) can be done by computing (as described in the proof of Lemma 18.16) a vertical ordering of the sheets of the components U_{ij} and T_{ij} ; the branch C_k is then characterized as the branch for which U_{ij} appears only on one of its sides (see Figure 18.4(b)).

We prove two lemmas that, together, prove Proposition 18.14.

Lemma 18.15. *Exactly one of the four branches of the trisector of three lines in general position admits only one asymptote.*

Proof. By Proposition 18.11, the trisector admits four distinct asymptotes, for all triplets of lines in general position. It follows that the property that exactly one of the branches of the trisector has only one asymptote is invariant by continuous deformation on the set of triplets of lines in general position. The result thus follows from Proposition 18.13 and from the observation that the property is verified on one particular example. This property can be observed on Figure 18.4(a) and it can easily be proved as follows. Consider any three lines, in general position, whose trisector consists in a cubic and a line (three such lines exist by Lemmas 18.8 and 18.10). The line is one branch of the trisector that admits only one asymptote. On the other hand, since the cubic consists of only one connected component in projective space and it intersects the plane at infinity in three real simple points (by Proposition 18.11), each of its three branches has two asymptotes. \square

We denote by C_0 the branch of the trisector that admits only one asymptote (see Figure 18.4(a)).

Lemma 18.16. *Each cell U_{ij} is bounded by a branch of the trisector distinct from C_0 and every such branch bounds a cell U_{ij} .*

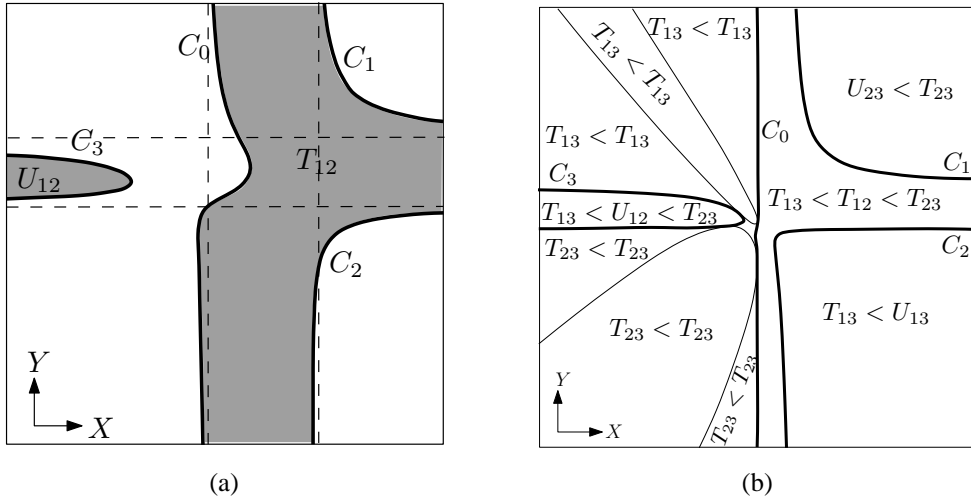


FIG. 18.4 – (a) Projection of the two-dimensional Voronoi cell V_{12} onto the XY -plane. (b) Vertical ordering of the sheets of the connected components of the two-dimensional Voronoi diagram cells above each region induced by the projection of the trisector and the silhouette curves of the bisectors ; the ordering over the small cell in the middle is $T_{13} < T_{13} < T_{23} < T_{23}$ (i.e., a vertical line over that cell intersects twice T_{13} and twice T_{23} in that order).

Proof. This property is invariant by continuous deformation on the set of triplets of lines in general position. It is thus sufficient to prove it for any three given lines in general position, ℓ_1, ℓ_2, ℓ_3 , as defined in Section 18.2.1. We consider in the XY -plane the arrangement of the (orthogonal) projection of the trisector and of the silhouette curves (viewed from infinity in the Z -direction) of the bisectors (see Figure 18.4(b)); these silhouette curves consist of only two parabolas since the bisector of lines ℓ_1 and ℓ_2 has no such silhouette (its equation has the form $Z = \gamma XY$ -see Eq.(18.2)- and thus any vertical line intersects it). By construction, for all vertical lines intersecting one given (open) cell of this arrangement, the number and ordering of the intersection points between the vertical line and all the pieces of the three bisectors that are bounded by the trisector is invariant. For any point of intersection, we can easily determine (by computing distances) whether the point lies on a Voronoi cell V_{ij} . We can further determine whether the point belongs to the component U_{ij} or T_{ij} by using the linear separation test described in Section 18.6. We thus report the ordering of the sheets of the components U_{ij} and T_{ij} above each cell of the arrangement in the XY -plane for a given example ; see Figure 18.4(b).

We can now observe that there is a one-to-one correspondence between the three branches of the trisector distinct from C_0 and the components U_{12}, U_{13} , and U_{23} such that the component appears only on one side of the corresponding branch⁴⁰. It follows that each of the branches distinct from C_0 bounds a cell U_{ij} . \square

Proof of Proposition 18.14. Lemmas 18.15 and 18.16 state the first two properties of Proposition 18.14. Furthermore, since U_{ij} is, by definition, bounded by only one arc of the trisector, Lemmas 18.15 and 18.16 directly yield the labeling of the four branches of the trisector by C_0, \dots, C_4 is unambiguous. \square

18.3.3 Properties of the trisector

We now prove two important properties of trisector of the Voronoi diagram of three lines in general position. In particular, we prove the Monotonicity Property in Proposition 18.18.

⁴⁰Namely, U_{13} (resp. U_{23} and U_{12}) appears on only one side of the lower-right (resp. upper-right and left-most) branch.

Proposition 18.17. *The orthogonal projection of the trisector of ℓ_1 , ℓ_2 , and ℓ_3 onto the XY -plane has two asymptotes parallel to the X -axis and two asymptotes parallel to the Y -axis.*

Proof. By Proposition 18.11, the four asymptotes of the trisector are parallel to the four trisector lines of three concurrent lines parallel to ℓ_1 , ℓ_2 , and ℓ_3 . The bisector to two lines through the origin and parallel to ℓ_1 and ℓ_2 is the pair of planes of equation $XY = 0$. Hence the asymptotes of the trisector are parallel to lines that lie in the pair of planes $XY = 0$. The orthogonal projection of the asymptotes on the XY -plane are thus parallel to the X - or Y -axis. It follows that the number of asymptotes (in projection) that are parallel to the X -axis (resp. Y -axis) is invariant by continuous deformation on any connected set of triplets of lines in general position. The result follows from the fact that, on a particular example (see Figure 18.4(a)), there are two asymptotes parallel to the X -axis and two others parallel to the Y -axis and that the set of triplets of lines in general position is connected (Proposition 18.13). \square

We assume in the following that *the asymptote of C_0 is parallel to the YZ -plane* (as in Figure 18.4(a)) by exchanging, if necessary, the role of X and Y .

Proposition 18.18. *Every branch of the trisector of ℓ_1 , ℓ_2 , and ℓ_3 is monotonic with respect to the Y -direction (or every branch is monotonic with respect to the X -direction).*

Proof. Let \mathcal{P} denote any plane parallel to the XZ -plane. The arc C_0 intersects plane \mathcal{P} an odd number of times (counted with multiplicity) since C_0 has only one asymptote (Proposition 18.14) which is parallel to the YZ -plane. Furthermore, by Proposition 18.17, the trisector has two other asymptotes parallel to the XZ -plane. Hence plane \mathcal{P} intersects the trisector in two points at infinity and C_0 an odd number of times (in affine space). The trisector thus intersects \mathcal{P} in at least three points in real projective space. There are thus four intersection points (in real projective space) since there are four intersection points in complex space (since the trisector is of degree four) and if there was an imaginary point of intersection, its conjugate would also be an intersection point (since the equations of the plane and quadrics have real coefficients) giving five points of intersection.

Therefore the trisector intersects plane \mathcal{P} in two points in \mathbb{R}^3 , one of which lies on C_0 . Since there are an odd number of intersection points on C_0 , plane \mathcal{P} intersects C_0 exactly once and any other branch exactly once. \square

18.4 Topology of the Voronoi diagram

We prove here that the topology of the Voronoi diagram of three lines in general position is invariant. Theorem 18.1 will thus follow from Theorem 18.12 and from the computation of an example of a two-dimensional cell of the Voronoi diagram (for instance the one shown in Figure 18.1).

Theorem 18.19. *The topology of the Voronoi diagram of three lines in general position is constant. More precisely, given two triplets of lines in general position, there is a continuous path between them (in the space of triplets of lines in general position) which induces a continuous deformation of every cell of the Voronoi diagram, preserving the topology of the cells and the incidence relations between them.*

Proof. The general idea of the proof is as follows. Consider three lines in general position and a bisector of two of them. The bisector is a hyperbolic paraboloid which is homeomorphic to a plane. The trisector lies on the bisector and it is homeomorphic to four lines that do not pairwise intersect, by Theorem 18.12. Hence the topology of the regions that lie on the bisector and are bounded by the trisector is invariant by continuous deformation on any connected set of triplets of lines (in general position). The topology of these regions is thus invariant by continuous deformation on the set of all triplets of lines in general position (by Proposition 18.13). It follows that the topology of the two-dimensional cells of the Voronoi diagram is invariant by such a continuous deformation. The Voronoi diagram is defined by the embedding in \mathbb{R}^3 of its two-dimensional cells, hence its topology is also invariant by continuous deformation.

To be more precise, we now show that any continuous path, in the space of triplets of lines in general position, between any two triplets of lines in general position, induces a continuous deformation of every cell of the Voronoi diagram, preserving the topology of the cells and the incidence relations between them.

Consider two triplets of lines in general position (ℓ_1, ℓ_2, ℓ_3) and $(\ell'_1, \ell'_2, \ell'_3)$. Without loss of generality, we may choose for (ℓ_1, ℓ_2, ℓ_3) the triplet of Figure 18.1. As the space of triplets of lines in general position is connected (Proposition 18.13), there is a homotopy between them, *i.e.*, a continuous application $\varphi : t \mapsto \varphi(t) = (\ell_1(t), \ell_2(t), \ell_3(t))$ of the interval $[0, 1]$ into the space of triplets of lines in general position such that $\varphi(0) = (\ell_1, \ell_2, \ell_3)$ and $\varphi(1) = (\ell'_1, \ell'_2, \ell'_3)$.

Consider now an orthonormal frame $\mathcal{F}(t)$ such that the Z -axis is the common perpendicular to $\ell_1(t)$ and $\ell_2(t)$, the origin of the frame is the point of the Z -axis equidistant to $\ell_1(t)$ and $\ell_2(t)$, and the X and Y -axes are the bisectors of the projections of $\ell_1(t)$ and $\ell_2(t)$ onto the plane orthogonal to the Z -axis. Note that this coordinate system is, up to a possible change of orientation of the axes and a possible exchange of the X and Y -axes, the one we considered in Sections 18.2 and 18.3 and which has been used to draw Figure 18.1. When the parameter t of the homotopy varies from 0 to 1, the lines vary continuously, and thus the frame $\mathcal{F}(t)$ can be defined to vary continuously in terms of t .

By Lemma 18.15 and Propositions 18.17 and 18.18, for any t in $[0, 1]$, each of the branches of the trisector is monotonic with respect to either the X or the Y -direction, but not both. Furthermore, the set of t for which each branch is monotonic with respect to the X -direction (resp. the Y -direction) is closed (since the lines and $\mathcal{F}(t)$ vary continuously in terms of t). Hence, each branch of the trisector is monotonic in X for all t or is monotonic in Y for all t . Therefore, by exchanging, if needed, X and Y in all frames $\mathcal{F}(t)$, we may suppose that each of the four branches of the trisector of $\ell_1(t)$, $\ell_2(t)$ and $\ell_3(t)$ is monotonic with respect to the Y -direction.

In the coordinate system $\mathcal{F}(t)$, the bisector of $\ell_1(t)$ and $\ell_2(t)$ has the equation $Z = \alpha(t)XY$ (see the proof of Lemma 18.4). Substituting Z by $\alpha(t)XY$ in the equation of the bisector of $\ell_2(t)$ and $\ell_3(t)$ in the coordinate system $\mathcal{F}(t)$, we get an equation of degree 2 in each of the variables X and Y . Solving it in X , we get a parameterization of the form $X = Y^\pm(Y, t)$ with $Y^\pm(Y, t) = \frac{-P_1(Y, t) \pm \sqrt{P_1(Y, t)^2 - 4P_0(Y, t)P_2(Y, t)}}{2P_2(Y, t)}$, where P_0 , P_1 and P_2 are polynomials of degree 2 in Y , which depend continuously on t (since the frame $\mathcal{F}(t)$ and the equations, in any fixed frame, of the bisectors depend continuously on t).

Notice first that $P_4(Y, t) = P_1(Y, t)^2 - 4P_0(Y, t)P_2(Y, t)$ is always positive. Indeed, it is always non-negative since one of the branches of the trisector of $\ell_1(t)$, $\ell_2(t)$ and $\ell_3(t)$ is defined for all Y in $\mathcal{F}(t)$ (since each branch is monotonic in Y and one of them has only one asymptote, by Lemma 18.15). It thus follows from the fact that the trisector has no real singular point (Theorem 18.12) that $P_4(Y, t)$ is always positive. Notice also that, for any t in $[0, 1]$, $P_2(Y, t)$ has two distinct real roots by Proposition 18.17.

Since $P_4(Y, t)$ is always positive, the branch $C_0(t)$ of $\ell_1(t)$, $\ell_2(t)$ and $\ell_3(t)$ (see Proposition 18.14) is parameterized by $X = Y^-(Y, t)$ or by $X = Y^+(Y, t)$ (but not by a combination of both). Thus, by changing, if needed, the signs of P_0 , P_1 and P_2 , we may suppose that $C_0(0)$ is parameterized by $X = Y^-(Y, 0)$. This implies, by continuity, that the branch $C_0(t)$ is parameterized, in the frame $\mathcal{F}(t)$, by $X = Y^-(Y, t)$, while the other branches are parameterized by $X = Y^+(Y, t)$ and the position of Y with respect to the two roots of $P_2(Y, t)$.

The study of the Voronoi diagram of $\ell_1(0)$, $\ell_2(0)$ and $\ell_3(0)$ (see Figures 18.1 and 18.4(a)) thus implies that the region, denoted $\mathcal{R}_{12}(t)$, of the Voronoi diagram consisting in the points which are at the same distance of the lines $\ell_1(t)$ and $\ell_2(t)$ and closer than to $\ell_3(t)$ consists, when $t = 0$, in two open semi algebraic sets defined in $\mathcal{F}(0)$ by (i) $Z = \alpha(0)XY$, $X < Y^+(Y, 0)$, and Y between the two roots of P_2 and by (ii) $Z = \alpha(0)XY$, $X > Y^-(Y, 0)$ and, when Y is outside the two roots of P_2 , $X < Y^+(Y, 0)$.

Now, since the objects we are considering depend continuously on t , including the distance from a point to one of the lines (note that the distance function is defined independently of $\mathcal{F}(t)$), the Voronoi region $\mathcal{R}_{12}(t)$ is defined, similarly, by the two open semi algebraic sets defined in $\mathcal{F}(t)$ by (i) $Z = \alpha(t)XY$, $X < Y^+(Y, t)$, and Y between the two roots of P_2 and by (ii) $Z = \alpha(t)XY$, $X > Y^-(Y, t)$ and, when Y is outside the two roots of P_2 , $X < Y^+(Y, t)$.

Note that, in the case where the trisector is decomposed, for some value of t , into a cubic and a line, nothing changes in what precedes, the only difference being that the square root is a polynomial and that the parameterization of C_0 simplifies into $X = \text{constant}$.

We thus get that, when t varies, the two-dimensional cells of the Voronoi diagram which are closer to $\ell_1(t)$ and $\ell_2(t)$ than to $\ell_3(t)$ varies continuously, with a constant topology and constant incidence relations with the trisector. As the same study may be done, replacing $\ell_1(t)$ and $\ell_2(t)$ by the other pairs of lines, this proves the theorem for all two-dimensional cells.

Finally, let P be a point of the region of $\ell_1(t)$ (*i.e.*, a point which is closer to $\ell_1(t)$ than the other lines) and Q its orthogonal projection on $\ell_1(t)$. Then, any point of the segment PQ lies also in the region of $\ell_1(t)$. It follows that the region of $\ell_1(t)$ is homeomorphic to a solid cylinder and has thus a constant topology. As this region va-

ries continuously with t , as well as the two-dimensional cells of its border, this finishes the proof of the theorem. \square

18.5 Configurations of three lines whose trisector contains a line

We present here a simple geometric proof of Theorem 18.2 which states that the trisector of three pairwise skew lines that are not all parallel to a common plane consists of a cubic and a line if and only if the hyperboloid of one sheet containing the three skew lines is of revolution. Note that a computational proof is also given by the direct proof of the Main Lemma, in which we proved that the trisector contains a line if and only if the hyperboloid is of revolution, and by Theorem 18.1, which states that the trisector contains a line if and only if it is a cubic and line.

Consider three lines ℓ_1, ℓ_2 and ℓ_3 whose trisector includes a line ℓ . Any point p on ℓ is equidistant to ℓ_1, ℓ_2 and ℓ_3 so p is the center of a sphere that is tangent to all of ℓ_1, ℓ_2 and ℓ_3 . Consider three distinct such points on ℓ and the three corresponding spheres. If these spheres have a common intersection, then this common intersection is a circle (possibly reduced to a point) and all lines tangent to the three spheres lie in the plane of this circle which contradicts the general position assumption. Otherwise, the set of lines tangent to the three spheres are the ruling(s) of a single quadric of revolution with symmetry axis the line through their centers [BGLP06, Lemma 7]. Note that this quadric is a hyperboloid of one sheet since it cannot be a cone or a cylinder by the general position assumption.

Conversely, if three lines lie on a quadric of revolution, any point on the axis of revolution is equidistant to the three lines. Thus the trisector of the three lines contains a line and, by Theorem 18.1, the trisector of three lines in general position is a non-singular quartic or a cubic and a line.

18.6 Algorithms

In this section, we prove Theorem 18.3. We start by presenting an algorithm for determining a plane separating the two components of any two-dimensional Voronoi cell. Refer to Figure 18.5(a). This plane may be non-rational; indeed, as we shall see in Proposition 18.21, it is possible that no rational separating plane exists. We then show how this algorithm can be modified to produce a rational linear test for this problem when the three input lines are rational. As we will see, this algorithm leads directly to another rational linear test for separating the four connected components of the cell of dimension one. Finally, we conclude the proof of Theorem 18.3 by showing how points on a branch of the trisector can be ordered using a linear form with rational coefficients.

Linear test for separating the two connected components of a two-dimensional Voronoi cell.

Input : three lines ℓ_1, ℓ_2 , and ℓ_3 in general position and $i \neq j \in \{1, 2, 3\}$.

Output : a half-space H_{ij} that strictly contains U_{ij} and whose complement strictly contains T_{ij} .

- (i) Determine a Cartesian coordinate system (X, Y, Z) such that the Z -axis is parallel to the common perpendicular of ℓ_i and ℓ_j and such that the X and Y -axes are parallel to the two bisector lines, in a plane perpendicular to the Z -axis, of the projection of ℓ_i and ℓ_j onto that plane.
- (ii) In this frame, compute all the critical values of the trisector with respect to the X -axis. If there is no critical value, exchange the X - and Y -axes (and compute the critical values with respect to the new X -axis).
- (iii) Compute the X -values of the two trisector asymptotes that are parallel to the YZ -plane. If the minimum of these values is smaller than the smaller critical value, then change the orientation of the X -axis. Denote by X_1 the smallest critical value (with respect to the X -axis) of the trisector and by X_2 the smallest of the other critical values and of the two asymptote X -values.
- (iv) Pick a value \tilde{x} in (X_1, X_2) . The half-space, H_{ij} , of equation $X < \tilde{x}$ contains U_{ij} and the half-space $X > \tilde{x}$ contains T_{ij} .

Proof of correctness. Assume without loss of generality that $i = 1$ and $j = 2$. By Proposition 18.18, the trisector has no critical point in the Y -direction after Step (ii).

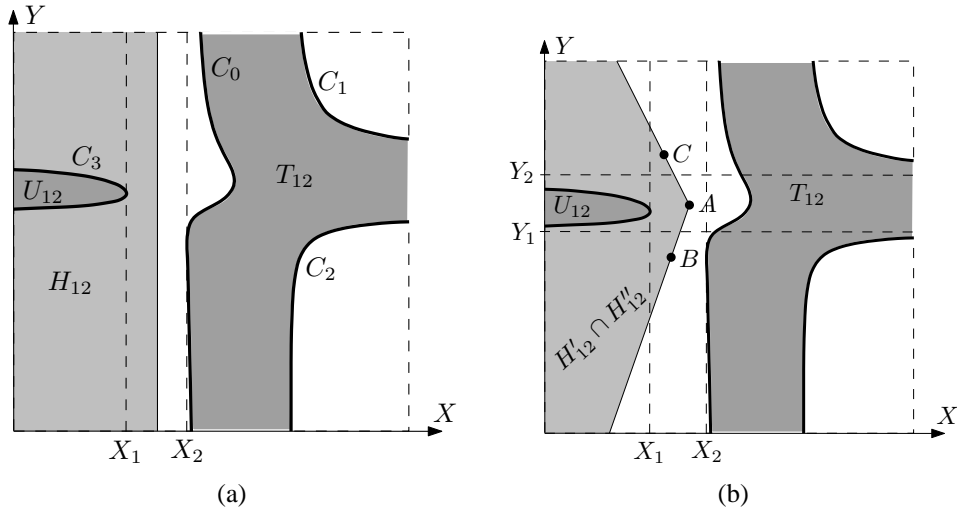


FIG. 18.5 – Separating the two components of a two-dimensional Voronoi cell.

First note that the asymptotes of the trisector are never vertical (*i.e.*, parallel to the Z -axis) because otherwise, by Proposition 18.11 and since ℓ_1 and ℓ_2 are horizontal, the line ℓ_3 would be horizontal (its direction would be the symmetric of the one of ℓ_1 with respect to a vertical plane), contradicting the general position assumption.

It thus follows, since the directions of the asymptotes, projected on the XY -plane, are parallel to the X or Y -axis (by Proposition 18.17) that the oriented directions of the asymptotes of the branches of the projected trisector are invariant (in the direction $\pm X$ or $\pm Y$) by continuous deformation on the set of triplets of lines in general position (which is connected by Proposition 18.13).

Hence, it follows from the analysis of one configuration (see Figure 18.5) that the two projected asymptotes of the branch C_3 have the same oriented direction. Thus C_3 has (at least) a critical point with respect to this direction, which is X since there is no critical point with respect to the Y -axis. We assume, for now, that the oriented asymptotic direction of the two branches of C_3 is the $-X$ direction (as in Figure 18.5), by changing, if necessary, the orientation of the axis. In the sequel of the proof, all the critical points are considered with respect to the X -axis.

Now, the plane, denoted \mathcal{P} , parallel to the YZ -plane through a critical point of the trisector does not intersect the trisector in any other point in \mathbb{R}^3 because the intersection at the critical point has multiplicity two, the plane intersects the trisector in two points at infinity (by Proposition 18.17), and the trisector has degree four (it is the intersection of two quadrics). It thus follows that C_3 has a unique critical point and that this critical point is strictly left (*i.e.*, has smaller X -coordinate) of all the other critical points of the trisector. Furthermore, the plane \mathcal{P} through this leftmost critical point, that is the plane of equation $X = X_1$, separates (strictly, except for the critical point) the branch C_3 from the other branches and leaves C_3 on its left. In other words, the half-space $X < X_1$ contains C_3 except for its critical point and the half-space $X > X_1$ contains the other branches. It then follows from the definition of X_2 that, for any $\tilde{x} \in (X_1, X_2)$, the half-space $X < \tilde{x}$ contains C_3 and the half-space $X > \tilde{x}$ contains the other branches of the trisector. We thus get that the half-space $X < \tilde{x}$ contains U_{12} because U_{12} is bounded by C_3 (by Proposition 18.14) and lies on a hyperbolic paraboloid of equation $Z = \gamma XY$, $\gamma \in \mathbb{R}$ (see Eq. (18.2)). Similarly, the half-space $X > \tilde{x}$ contains T_{12} .

It remains to show that the orientation of the X -axis obtained in Step (iii) of the algorithm is the same as the one we have considered so far. Consider the two X -values of the two trisector asymptotes parallel to the XZ -plane. We prove that the maximum of these values is larger than the largest critical value. This implies the result since, if the orientation of the X -axis was not as assumed above, it would have been changed in Step (iii).

As before, by continuity and by analyzing one particular example, we have that two of the asymptotes of the branches of C_1 and C_2 have direction $+X$ (in projection) and the two others have direction $+Y$ and $-Y$. We consider here the trisector and its asymptotes in projection on the XY -plane and we refer to vertical, right and left in a standard way in the (X, Y) frame. Suppose for a contradiction that there exists a critical point on $C_1 \cup C_2$ that is

right of both their vertical asymptote. Then a vertical line, \mathcal{L} , through this critical point would intersect the trisector at this point, with multiplicity two, and at two other points at infinity (by Proposition 18.17). However, since the critical point is right of the vertical asymptote of C_1 and C_2 , line \mathcal{L} intersects the trisector somewhere else (or with higher multiplicity), which is not possible since the trisector has degree four. \square

The algorithm requires computing the critical values of the trisector with respect to the X and Y -directions. We proved (in Proposition 18.18) that the trisector has no critical values in one of these directions. We show below that the trisector admits at most four critical values with respect to the other direction. We consider below the coordinate system obtained after Step (ii) of the algorithm above.

Lemma 18.20. *The trisector has three or four critical values with respect to the X -direction. Moreover, the trisector has one critical point on C_3 , one on $C_1 \cup C_2$, and either two on C_0 or C_0 is a line perpendicular to the X -axis.*

Proof. We consider here critical points and critical values with respect to the X -direction. First, we proved in the proof of correctness of the algorithm that C_3 has exactly one critical point. A similar study of the directions of the branches of asymptotes of $C_1 \cup C_2$ implies that $C_1 \cup C_2$ has also exactly one critical point. On the other hand, we have that C_0 has two identical asymptotes that are perpendicular to the X -axis (by Propositions 18.14, 18.17 and Step (ii) of the algorithm) and thus C_0 contains at least one critical point.

Consider first the case where C_0 is entirely critical. It then projects on the XY -plane to a line perpendicular to the X -axis. It is planar and thus contained in the intersection of a plane and a quadric (the bisector of any two of the input lines). C_0 is thus a line or an irreducible conic. The trisector never contains an irreducible conic (by Theorem 18.1), thus C_0 is a line that is perpendicular to the X -axis (since its projection on the XY -plane is). This concludes the proof in the case where the trisector contains infinitely many critical points. We assume in the sequel the trisector has finitely many critical points.

Now, the projection (on the XY -plane) of the trisector is a curve of degree four. Furthermore, it has degree two in X and degree two in Y because the curve intersects any line parallel to the X - or Y -axis in at most two points since there are two other points of intersection at infinity (by Proposition 18.17). The projected curve thus has equation $A(X)Y^2 + B(X)Y + C(X) = 0$ where A , B and C are polynomials of degree two in X . The critical points are points on the curve such that the curve's partial derivative with respect to Y is zero. This partial derivative is of degree one in Y and two in X ; it has equation $2A(X)Y + B(X) = 0$. The curve contains no critical point (X_0, Y_0) such that $A(X_0) = 0$ because otherwise $A(X_0) = B(X_0) = C(X_0) = 0$ and thus the line (X_0, Y) is critical, contradicting the above hypothesis. Hence, eliminating Y in the curve's equation gives an equation in X of degree four.

Since $C_1 \cup C_2 \cup C_3$ has exactly two critical points, C_0 has either zero or two critical points, counted with multiplicity. We have shown that C_0 has at least one critical point, thus it has exactly two critical points counted with multiplicity. Finally, C_0 cannot only have one double critical point because its two asymptotes are identical and vertical. Hence, when the trisector has finitely many critical points, exactly two lie on C_0 , one on $C_1 \cup C_2$ and one on C_3 . \square

The following proposition shows that the separating plane computed in the above algorithm may not be rational.

Proposition 18.21. *There exist three rational lines for which the two connected components of any two-dimensional Voronoi cell cannot be separated by a rational plane.*

Proof. Let \mathcal{P} denote any plane separating U_{ij} and T_{ij} . Since \mathcal{P} does not intersect C_0 , it is necessarily parallel to the asymptote of C_0 (see Proposition 18.14).

We now exhibit an example of three rational lines such that there exists no rational plane parallel to an asymptote of their trisector, which will conclude the proof. Consider three lines ℓ_1, ℓ_2 , and ℓ_3 in general position that have direction $(1, 0, 0)$, $(1, 1, 0)$, and $(2, 0, 1)$, respectively. By Proposition 18.11, the four asymptotes of their trisector are parallel to the four trisector lines of three concurrent lines (say, through the origin) with directions those of ℓ_1, ℓ_2 , and ℓ_3 ; let ℓ'_1, ℓ'_2 , and ℓ'_3 denote these lines.

The pair of bisector planes of ℓ'_1 and ℓ'_2 has a square root of 2 in its coefficient: its equation (see Eq. 18.1) factors into $(X - (1 + \sqrt{2})Y)(X - (1 - \sqrt{2})Y)$, which is the equation of a pair of conjugate planes over $\mathbb{Q}(\sqrt{2})$ (the field extension of \mathbb{Q} by $\sqrt{2}$). Similarly, the bisector planes of ℓ'_1 and ℓ'_3 is a pair of conjugate planes over $\mathbb{Q}(\sqrt{5})$

(it has equation $(X - (2 + \sqrt{5})Z)(X - (2 - \sqrt{5})Z)$). It follows that the four lines of intersection of these two pairs of planes are conjugate over $\mathbb{Q}(\sqrt{2}, \sqrt{5})$.

Furthermore, these four lines are not all parallel to a common plane because the intersection of the two planes that are conjugate over $\mathbb{Q}(\sqrt{2})$ is the Z -axis, which properly intersects each of the two other conjugate planes; thus, on each of these latter conjugate planes, the two lines of intersection properly intersect and thus any plane parallel to them is parallel to the plane they define; since the two conjugate planes are not coplanar, no plane is parallel to the four lines of intersection.

Now, any rational plane that is parallel to one of these four lines is also parallel to the three others (since a rational plane is invariant by conjugation over $\mathbb{Q}(\sqrt{2}, \sqrt{5})$). Since this is not possible, there is no rational plane that is parallel to the asymptote of C_0 , which concludes the proof. \square

We now present an algorithm for determining a rational linear test for separating the two components of any two-dimensional Voronoi cell of three rational lines. Refer to Figure 18.5(b).

Rational linear test for separating the two connected components of a two-dimensional Voronoi cell.

Input : three rational lines ℓ_1, ℓ_2 , and ℓ_3 in general position in a coordinate system $(\tilde{X}, \tilde{Y}, \tilde{Z})$ and $i \neq j \in \{1, 2, 3\}$.

Output : two rational half-spaces H'_{ij} and H''_{ij} such that $H'_{ij} \cap H''_{ij}$ strictly contains U_{ij} and its complement strictly contains T_{ij} .

- (i-iii) Idem as in the previous algorithm.
- (iv) Compute the two Y -values of the two trisector asymptotes that are parallel to the XZ -plane. Let $Y_1 < Y_2$ denote these two values.
- (v) Determine a point A with rational coordinates in the original $(\tilde{X}, \tilde{Y}, \tilde{Z})$ -frame such that its X -, Y -, and Z -coordinates in the (X, Y, Z) frame are in (X_1, X_2) , in (Y_1, Y_2) , and equal to 0, respectively; let X_A denote its X -coordinate in the (X, Y, Z) frame.
- (vi) Determine two points B and C with rational coordinates in the original $(\tilde{X}, \tilde{Y}, \tilde{Z})$ -frame such that their X -, Y -, and Z -coordinates in the (X, Y, Z) -frame are, for B , in (X_1, X_A) , in $(-\infty, Y_1)$, and equal to 0, respectively, and for C , in (X_1, X_A) , in $(Y_2, +\infty)$, and equal to 0, respectively.
- (vii) Let P_{ij} (resp. P'_{ij}) be the plane through A and B (resp. C) that is parallel to the Z -axis. Let H'_{ij} (resp. H''_{ij}) be the open half-space bounded plane P_{ij} (resp. P'_{ij}) that contains the point at infinity in the $-X$ -direction.

Remark 18.22. Note that the transformation from the $(\tilde{X}, \tilde{Y}, \tilde{Z})$ -frame to the (X, Y, Z) -frame is not necessarily rational since the X - and Y -axes are not necessarily rational in the $(\tilde{X}, \tilde{Y}, \tilde{Z})$ -frame. Nonetheless, the rational coordinates of the points A, B , and C can easily be computed using interval arithmetic. We however did not study the bit complexity of our algorithm, which requires finding rational values in between roots of constant-degree polynomials whose coefficients are not rational.

Proof of correctness. We assume without loss of generality that i and j are equal to 1 and 2, respectively. We have seen in the proof of correctness of the previous algorithm that the component C_3 has exactly one critical value with respect to the X -axis, no critical value with respect to the Y -axis, and two asymptotes in the $-Y$ -direction. The component C_3 is thus contained in the region defined by $X < X_1$ and $Y_1 < Y < Y_2$. It follows that $H'_{ij} \cap H''_{ij}$ contains U_{ij} .

On the other hand, the complement of $H'_{ij} \cap H''_{ij}$ strictly contains T_{ij} because for any value $\tilde{x} \in (X_A, X_2)$, the half-space $X > \tilde{x}$ contains T_{ij} (as proved above) and this half-space is contained in the complement of $H'_{ij} \cap H''_{ij}$.

Finally, the plane P_{ij} is rational since A and B and are rational as well as the Z -axis (since it is the common perpendicular to ℓ_i and ℓ_j). Similarly, plane P'_{ij} is also rational. \square

Remark 18.23. Note that, if the three input lines are not rational, the above algorithm remains valid except for the fact that the output half-spaces are not necessarily rational anymore (since the common perpendicular to ℓ_i and ℓ_j is not necessarily rational).

Separation of the four connected components of the trisector of three lines.

Consider three lines ℓ_1 , ℓ_2 , and ℓ_3 and the half-space H'_{ij} and H''_{ij} obtained by the above algorithm. Proposition 18.14 (and Remark 18.23) directly yields the following result.

Proposition 18.24. *For any point p on the trisector of ℓ_1 , ℓ_2 , and ℓ_3 , if p belongs to both half-spaces H'_{ij} and H''_{ij} for some $i \neq j \in \{1, 2, 3\}$ then p lies on C_k (with $k \in \{1, 2, 3\}$ distinct from i and j), otherwise p lies on C_0 . Furthermore, if the three input lines are rational, the coefficients of H'_{ij} and H''_{ij} are rational.*

We conclude this section by proving Theorem 18.3.

Proof of Theorem 18.3. First, the algorithms of this section and Proposition 18.24 present some (rational) linear tests for separating the connected components of the Voronoi cells of dimensions one and two. Second, we can compute, as described in Steps (i-ii) of the above algorithms, a direction in which every branch of the trisector is monotonic, which gives a linear test for ordering points on each trisector.

Now, if the three input lines are rational, the Cartesian coordinate system computed in the above algorithms is such that the Z -axis is rational and, if the X -axis is irrational, a slight rotation of the frame around the Z -axis gives a rational frame (*i.e.*, a frame which is defined on the initial frame by a matrix with rational coefficients).

If, as in Figure 4a, there is a critical point on the lower branch C_2 (for the projection on the X -axis) and if the rotation is clockwise, then the projections of C_0 , C_1 and C_2 on the new Y -axis are monotonic. If the critical point is on the upper branch C_1 then a counter-clockwise rotation gives the same result. Thus the points on each of these three branches can be sorted using a linear form with rational coefficients. The same result is obtained for C_3 by doing the same work after a circular permutation of the lines. \square

18.7 Conclusion

We presented a complete description of the Voronoi diagram of three lines that are pairwise skew and not all parallel to a common plane. We also presented some algorithms for determining a rational test for answering queries of the form, given a point, determine in which connected component of which Voronoi cell it lies. We also showed that points on a branch of the trisector of three lines can easily be ordered by comparing their coordinates in a particular direction, which is however not necessarily rational.

Future work includes the characterization of the topology of the Voronoi diagram of three lines that are not in general position. Note that, in this case, the topology of the Voronoi diagram does indeed change; for instance, when three pairwise skew lines are all parallel to a common plane, their bisectors are hyperbolic paraboloids of the form $Z = F_{ij}(X, Y)$ and it follows that their trisector consists of two branches (instead of four) as it is the intersection of one of the bisectors with a hyperbolic cylinder whose axis is parallel to the Z -axis (of equation $F_{12}(X, Y) - F_{13}(X, Y) = 0$). Note also that when two of the lines are coplanar their bisector is one or two planes and the trisector is thus either the intersection of two such bisectors or the intersection of one such bisector with a hyperbolic paraboloid.

A challenging problem is to study Voronoi diagrams of up to six lines; this is of interest for the general case of n lines because the arcs of such diagrams are defined by five lines. Finally, the two major problems remain the determination of the complexity of Voronoi diagrams of n lines and the design of efficient algorithms for computing Voronoi diagrams of lines, segments, triangles, or polyhedra.

Acknowledgments

We would like to thank C. Hillar for discussions about sums of squares.


```

> factor(subs(y=-a*alpha,big_fact));

      (alpha^4 a^4 + 2 beta x alpha^2 a^3 + x^2 a^2 + beta^2 x^2 a^2 - 2 a^2 alpha^2 + 1 + beta^2)
      (beta^2 - 4 a^2 - 4 a^2 alpha^2 - 4 a^4 - 4 a^4 alpha^2 - 2 a beta x - 4 beta x a^3 + x^2 a^2)^2
> f0:=collect(op(1,%),x); f1:=collect(op(1,op(2,%)),x);
      f0 := (a^2 beta^2 + a^2)x^2 + 2 beta x alpha^2 a^3 + alpha^4 a^4 + 1 + beta^2 - 2 a^2 alpha^2
      f1 := x^2 a^2 + (-2 a beta - 4 beta a^3)x + beta^2 - 4 a^2 - 4 a^2 alpha^2 - 4 a^4 - 4 a^4 alpha^2
> factor(subs(x=-beta/a,big_fact));

      (beta^4 - 2 a^2 beta^2 + a^4 + a^4 alpha^2 + 2 beta^2 alpha a y + alpha^2 y^2 a^2 + y^2 a^2)
      (4 + 4 beta^2 + 4 a^2 + 4 a^2 beta^2 - a^4 alpha^2 + 4 a y alpha + 2 y a^3 alpha - y^2 a^2)^2
> g0:=collect(op(1,%),y); g1:=collect(op(1,op(2,%)),y);
      g0 := (a^2 alpha^2 + a^2)y^2 + 2 beta^2 alpha a y + beta^4 - 2 a^2 beta^2 + a^4 + a^4 alpha^2
      g1 := -y^2 a^2 + (4 a alpha + 2 a^3 alpha)y + 4 + 4 beta^2 + 4 a^2 + 4 a^2 beta^2 - a^4 alpha^2

```

Solutions of $f1=0$ in x and of $g1=0$ in y :

```

> map(uu->factor(uu),[solve(f1,x)]);
      [frac(2 a^2 beta + beta + 2 sqrt(a^2 (a^2 + 1) (beta^2 + 1 + alpha^2)), a),
       frac(2 a^2 beta + beta - 2 sqrt(a^2 (a^2 + 1) (beta^2 + 1 + alpha^2)), a)]
> map(uu->factor(uu),[solve(g1,y)]);
      [frac(alpha a^2 + 2 alpha + 2 sqrt((a^2 + 1) (beta^2 + 1 + alpha^2)), a),
       frac(alpha a^2 + 2 alpha - 2 sqrt((a^2 + 1) (beta^2 + 1 + alpha^2)), a)]

```

$f0$ is a sum of square :

```

> (a^2*alpha^2-1+a*beta*x)^2+(a*x+beta)^2;
> simplify(f0-%);

```

$$(a^2 \alpha^2 - 1 + a \beta x)^2 + (x a + \beta)^2$$

0

$a*x+beta$ and $g1$ are in the ideal generated by $y+a*\alpha$, $x*a+beta$, and $a^2*\alpha^2-1+a*beta*x$:

```

> gbasis([y+a*alpha,x*a+beta,a^2*alpha^2-1+a*beta*x],DRL([a,x,y,alpha,beta]));
> normalf(a*x+beta,%), normalf(g1,%);
      0,0

```

$g0$ is a sum of square :

```

> (a*y*alpha+beta^2-a^2)^2+a^2*(y+a*alpha)^2;
> simplify(g0-%);

```

$$(a y \alpha + \beta^2 - a^2)^2 + a^2 (y + a \alpha)^2$$

0

$y+a*\alpha$ and $f1$ are in the ideal generated by $x*a+beta$, $y+a*\alpha$, and $a^2*\alpha^2-1+a*beta*x$:

```

> gbasis([x*a+beta,y+a*alpha,a*y*alpha+beta^2-a^2],DRL([a,x,y,alpha,beta]));
> normalf(y+a*alpha,%), normalf(f1,%);
      0,0

```

TAB. 18.4 – For the proof of Lemma 18.8.

```
> comp1 := [y = -a*alpha, x =
> (2*beta*a^2+beta)/a+2*sqrt((beta^2+1+alpha^2)*(1+a^2))];
```

$$\text{comp1} := [y = -\alpha a, x = \frac{2\beta a^2 + \beta}{a} + 2\sqrt{(1 + \alpha^2 + \beta^2)(1 + a^2)}]$$

We prove that the characteristic equation has no real root on this component.

```
> factor(subs(comp1, Char_eq));
> irrat:=op(2,%):
```

```
a^2(4-4β^2λ^3+8a^2-4λ^3+λ^4-8λ-16α^2λα^2-8β^2λα^2+8α^2+4β^2+12a^2α^2+12a^2β^2+4a^4+8a^4β^2+4a^4α^2-8λα^2-16α^2λ
-8β^2λ+8λ^2+4λ^2a^2+8a^2α^2λ^2+4β^2λ^2a^2+8β^2λ^2-8βαα^3λ-8βαλα+8βαα^3+8aβα+8α√%1-8λα^2α√%1+λ^4β^2
+λ^4α^2+4λ^2β√%1a-8λβ√%1a+12α^2λ^2-4α^2λ^3+8β√%1a+8a^2α√%1+8βα^3√%1-4λ^3α√%1+12λ^2α√%1-16λαα√%1)
%1 := (β^2+1+α^2)(1+a^2)
```

Consider the product of the characteristic polynomial with its algebraic conjugate :

```
> T:=expand(irrat*subs(sqrt((1+a^2)*(alpha^2+beta^2+1))=-sqrt((1+a^2)*(alpha^2
> +beta^2+1)), irrat));
```

The real semi-algebraic set defined by $T-1/2 < 0$ is empty :

```
> sampling_negative(T-1/2, [a, alpha, beta, lambda]);
```

Pre-process.....

Computing critical values of a polynomial mapping from C^4 to C

Computing asymptotic critical values of a polynomial mapping from C^4 to C

*****Enter in internal", [alpha,beta, lambda], [], [], [a]

End of pre-process.....

Computing sampling points in a real hypersurface

Computing Critical Points using FGb (projection on a)

Computing Asymptotic Critical Values of alpha restricted to a hypersurface

Computing Critical Points using FGb (projection on alpha)

Computing Asymptotic Critical Values of alpha restricted to a hypersurface

Computing Asymptotic Critical Values of alpha restricted to a hypersurface

Computing Critical Points using FGb (projection on beta)

Computing Asymptotic Critical Values of beta restricted to a hypersurface

Computing Critical Points using FGb (projection on lambda)

Isolating real solutions of a zero-dimensional system using RS

Isolating real solutions of a zero-dimensional system using RS

Isolating real solutions of a zero-dimensional system using RS

Isolating real solutions of a zero-dimensional system using RS

□

Consider all the 3x3 minors of the matrix $P(\lambda)$ of the pencil :

```
> ldet:=NULL:
> for i to 4 do for j from i to 4 do
> ldet:=ldet,det(minor(P,i,j)):
> od od:
```

The rank of $P(\lambda)$ is always 3 or 4 since there is no common zeros of the minors :

```
> [ldet, 1-t*(1+alpha^2+beta^2)*(1+a^2)*(-beta+y+a*x-a*alpha)*(-beta-y+a*x+a*alpha)]:
> fgb_gbasis_elim(% , 0, [], [t, a, x, y, alpha, beta, lambda]);
```

[1]

TAB. 18.5 – For the proof of Lemma 18.10.

Bibliographie

- [AAKS05] P. K. Agarwal, B. Aronov, V. Koltun, and M. Sharir. On lines avoiding unit balls in three dimensions. *Discrete and Computational Geometry*, 34 :231–250, 2005.
- [AAS99] P. K. Agarwal, B. Aronov, and M. Sharir. Line transversals of balls and smallest enclosing cylinders in three dimensions. *Discrete and Computational Geometry*, 21 :373–388, 1999.
- [AAS04] P. K. Agarwal, B. Aronov, and M. Sharir. On lines avoiding unit balls in three dimensions. In Proceedings of the 20th ACM Annual Symposium on Computational Geometry (SoCG'04), pages 36–45, 2004.
- [ABCC02] B. Aronov, H. Brönnimann, A. Chang, and Y.-J. Chiang. Cost prediction for ray shooting. In Proceedings of the 18th ACM Annual Symposium on Computational Geometry (SoCG'02), pages 293–302, 2002.
- [ACP⁺01] L. Alonso, F. Cuny, S. Petitjean, J.-C. Paul, S. Lazard, and E. Wies. The virtual mesh : A geometric abstraction for efficiently computing radiosity. *ACM Transactions on Graphics*, 20(3) :169–201, 2001.
- [Aga94] P. K. Agarwal. On stabbing lines for convex polyhedra in 3D. *Computational Geometry : Theory and Applications*, 4(4) :177–189, 1994.
- [AGG03] H. Alt, M. Glisse, and X. Goaoc. On the worst-case complexity of the silhouette of a polytope. In Proceedings of the 15th Canadian Conference on Computational Geometry (CCCG'03), pages 51–55, Aug 2003.
- [AK99] F. Aurenhammer and R. Klein. Voronoi diagrams. In J. R. Sack and J. Urrutia, editors, *Handbook of computational geometry*, chapter 5, pages 201–290. Elsevier Publishing House, December 1999.
- [ALMM99] P. Aubry, D. Lazard, and M. Moreno Maza. On the theories of triangular sets. *Journal of Symbolic Computation*, 28(1) :105–124, 1999. Special Issue on Polynomial Elimination.
- [Ang02] P. Angelier. *Algorithmique des graphes de visibilité*. PhD thesis, Université Paris 7, 2002.
- [AP03a] P. Angelier and M. Pocchiola. CGAL-based implementation of visibility complexes. Technical Report ECG-TR-241207-01, Effective Computational Geometry for Curves and Surfaces (ECG), 2003.
- [AP03b] P. Angelier and M. Pocchiola. A sum of squares theorem for visibility complexes and applications. In B. Aronov, S. Basu, J. Pach, and M. Sharir, editors, *Discrete and Computational Geometry*, volume 25 of *Algorithms and Combinatorics*, pages 79–139. Springer-Verlag, 2003.
- [AS96] P. K. Agarwal and M. Sharir. Ray shooting amidst convex polyhedra and polyhedral terrains in three dimensions. *Siam Journal on Computing*, 25 :100–116, 1996.
- [AS97] B. Aronov and M. Sharir. The common exterior of convex polygons in the plane. *Computational Geometry : Theory and Applications*, 8 :139–149, 1997.
- [Aud] Audacity : The Free, Cross-Platform Sound Editor. <http://audacity.sourceforge.net/>.
- [Aur91] F. Aurenhammer. Voronoi diagrams – a survey of a fundamental geometric data structure. *ACM Computing Surveys*, 23(3) :345–405, 1991.
- [BDD⁺02] 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 14th Canadian Conference on Computational Geometry (CCCG'02), pages 113–117, Lethbridge, Canada, 2002.

- [BDD⁺04] H. Brönnimann, O. Devillers, V. Dujmović, H. Everett, M. Glisse, X. Goaoc, S. Lazard, H.-S. Na, and S. Whitesides. The number of lines tangent to arbitrary polytopes in \mathbb{R}^3 . In *Proceedings of the 20th ACM Annual Symposium on Computational Geometry (SoCG'04)*, pages 46–55, Brooklyn, NY, 2004.
- [BDD⁺05] H. Brönnimann, O. Devillers, V. Dujmović, H. Everett, M. Glisse, X. Goaoc, S. Lazard, H.-S. Na, and S. Whitesides. On the number of maximal free line segments tangent to arbitrary three-dimensional convex polyhedra. Research Report n° 5671, INRIA, Sept. 2005. 28 pages.
- [BDD⁺07] H. Bronnimann, O. Devillers, V. Dujmovic, H. Everett, M. Glisse, X. Goaoc, S. Lazard, H.-S. Na, and S. Whitesides. Lines and free line segments tangent to arbitrary three-dimensional convex polyhedra. *SIAM Journal on Computing*, 37(2) :522–551, 2007.
- [BDEG94] M. Bern, D. P. Dobkin, D. Eppstein, and R. Grossman. Visibility with a moving point of view. *Algorithmica*, 11 :360–378, 1994.
- [BDLS04] H. Brönnimann, O. Devillers, S. Lazard, and F. Sottile. Line tangents to four triangles in three-dimensional space. In *Proceedings of the 16th Canadian Conference on Computational Geometry (CCCG'04)*, pages 184–187, 2004.
- [BDLS05] H. Brönnimann, O. Devillers, S. Lazard, and F. Sottile. Line tangents to four triangles in three-dimensional space. Research Report n° 5693, INRIA, 2005. 17 pages.
- [BDLS07] H. Brönnimann, O. Devillers, S. Lazard, and F. Sottile. Lines tangent to four triangles in three-dimensional space. *Discrete and Computational Geometry*, 37(3) :369–380, 2007.
- [BDP⁺02] J.-D. Boissonnat, O. Devillers, S. Pion, M. Teillaud, and M. Yvinec. Triangulations in CGAL. *Computational Geometry : Theory and Applications*, 22 :5–19, 2002.
- [BDS⁺92] J.-D. Boissonnat, O. Devillers, R. Schott, M. Teillaud, and M. Yvinec. Applications of random sampling to on-line algorithms in computational geometry. *Discrete and Computational Geometry*, 8 :51–71, 1992.
- [BEL⁺05] 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) :381–390, 2005.
- [BFM⁺01] C. Burnikel, S. Funke, K. Mehlhorn, S. Schirra, and S. Schmitt. A separation bound for real algebraic expressions. In *Proceedings of the 9th Annual European Symposium on Algorithms*, volume 2161 of *Lecture Notes in Computer Science*, pages 254–265. Springer-Verlag, 2001.
- [BGLP03] C. Borcea, X. Goaoc, S. Lazard, and S. Petitjean. On tangents to quadric surfaces. Preprint, arXiv:math.AG/0402394, 2003.
- [BGLP06] C. Borcea, X. Goaoc, S. Lazard, and S. Petitjean. Common tangents to spheres in \mathbb{R}^3 . *Discrete and Computational Geometry*, 35(2) :287–300, 2006.
- [BHK⁺05] E. Berberich, M. Hemmer, L. Kettner, E. Schömer, and N. Wolpert. An exact, complete and efficient implementation for computing planar maps of quadric intersection curves. In *Proceedings of the 21st ACM Annual Symposium on Computational Geometry (SoCG'05)*, pages 99–115, 2005.
- [BM04] Y. Bugeaud and M. Mignotte. On the distance between roots of integer polynomials. *Proceedings of the Edinburgh Mathematical Society*, 47(3) :553–556, 2004.
- [BO05] J.-D. Boissonnat and S. Oudot. Provably good sampling and meshing of surfaces. *Graphical Models, special issue of Solid Modeling*, 67(5) :405–451, 2005.
- [BP00] J.-D. Boissonnat and F. Preparata. Robust plane sweep for intersecting segments. *Siam Journal on Computing*, 29(5) :1401–1421, 2000.
- [Bro06] T. Bromwich. *Quadratic Forms and Their Classification by Means of Invariant Factors*. Cambridge Tracts in Mathematics and Mathematical Physics, 1906.
- [BY98] J.-D. Boissonnat and M. Yvinec. *Algorithmic Geometry*. Cambridge University Press, 1998.
- [CEK⁺07] O. Cheong, H. Everett, H.-S. Kim, S. Lazard, and R. Schott. Parabola separation queries and their application to stone throwing. *International Journal of Computational Geometry and Applications*, 17(4) :349–360, 2007.

- [CF99] F. Cho and D. Forsyth. Interactive ray tracing with the visibility complex. *Computers and Graphics*, 23(5) :703–717, 1999. Special issue on Visibility - Techniques and Applications.
- [CGA] CGAL : Computational Geometry Algorithms Library. <http://www.cgal.org>.
- [CGN05] O. Cheong, X. Goaoc, and H.-S. Na. Geometric permutations of disjoint unit spheres. *Computational Geometry : Theory and Applications*, 30(3) :253–270, 2005.
- [Cha91] B. Chazelle. Triangulating a simple polygon in linear time. *Discrete and Computational Geometry*, 6 :485–524, 1991.
- [COR] The CORE library. <http://cs.nyu.edu/exact/>.
- [CR03] J. Cremona and D. Rusin. Efficient solution of rational conics. *Mathematics of Computation*, 72(243) :1417–1441, 2003.
- [CS89a] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete and Computational Geometry*, 4 :387–421, 1989.
- [CS89b] R. Cole and M. Sharir. Visibility problems for polyhedral terrains. *Journal of Symbolic Computation*, 7(1) :11–30, 1989.
- [CSY97] T. M. Chan, J. Snoeyink, and C. K. Yap. Primal dividing and dual pruning : Output-sensitive construction of four dimensional polytopes and three-dimensional diagram voronoi diagrams. *Discrete and Computational Geometry*, 18 :433–454, 1997.
- [Cul00] T. Culver. *Computing the Medial Axis of a Polyhedron Reliably and Efficiently*. PhD thesis, University of North Carolina at Chapel Hill, 2000.
- [dBEG98] M. de Berg, H. Everett, and L. Guibas. The union of moving polygonal pseudodiscs – combinatorial bounds and applications. *Computational Geometry : Theory and Applications*, 11 :69–82, 1998.
- [dBHOvK97] M. de Berg, D. Halperin, M. Overmars, and M. van Kreveld. Sparse arrangements and the number of views of polyhedral scenes. *International Journal of Computational Geometry and Applications*, 7(3) :175–195, 1997.
- [DD02] F. Duguet and G. Drettakis. Robust epsilon visibility. In John Hughes, editor, *Proceedings of ACM SIGGRAPH 2002*, pages 567–575. ACM Press / ACM SIGGRAPH, July 2002.
- [DDE⁺03] 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) :1586–1620, 2003.
- [DDE⁺07] J. Demouth, O. Devillers, H. Everett, S. Lazard, and R. Seidel. Between umbra and benumbra. In *Proceedings of the 23rd European Conference on Computational Geometry*, 2007.
- [DDP97] F. Durand, G. Drettakis, and C. Puech. The visibility skeleton : a powerful and efficient multi-purpose global visibility tool. *Computer Graphics Proceedings, Annual Conference Series*, 31 :89–100, 1997. Proceedings of Siggraph’97.
- [DDP99] F. Durand, G. Drettakis, and C. Puech. Fast and accurate hierarchical radiosity using global visibility. *ACM Transactions on Graphics.*, 18(2) :128–170, 1999.
- [DDP02] F. Durand, G. Drettakis, and C. Puech. The 3D visibility complex. *ACM Transactions on Graphics*, 21(2) :176–206, 2002.
- [DF94] George Drettakis and Eugene Fiume. A fast shadow algorithm for area light sources using back-projection. In *SIGGRAPH ’94 : Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 223–230, New York, NY, USA, 1994. ACM Press.
- [DHH00] O. Devillers and O. Hall-Holt. Personal communication, 2000.
- [DK83] D. P. Dobkin and D. G. Kirkpatrick. Fast detection of polyhedral intersection. *Theoretical Computer Science*, 27(3) :241–253, 1983.
- [DLLP03] L. Dupont, D. Lazard, S. Lazard, and S. Petitjean. Near-optimal parameterization of the intersection of quadrics. In *Proceedings of the 19th ACM Annual Symposium on Computational Geometry (SoCG’03)*, pages 246–255, San Diego, 2003.

- [DLLP05a] L. Dupont, D. Lazard, S. Lazard, and S. Petitjean. Near-optimal parameterization of the intersection of quadrics : I. The generic algorithm. Research Report n° 5667, INRIA, Sept. 2005. 36 pages.
- [DLLP05b] L. Dupont, D. Lazard, S. Lazard, and S. Petitjean. Near-optimal parameterization of the intersection of quadrics : II. A classification of pencils. Research Report n° 5668, INRIA, Sept. 2005. 37 pages.
- [DLLP05c] L. Dupont, D. Lazard, S. Lazard, and S. Petitjean. Near-optimal parameterization of the intersection of quadrics : III. Parameterizing singular intersections. Research Report n° 5669, INRIA, Sept. 2005. 34 pages.
- [DLLP07a] L. Dupont, D. Lazard, S. Lazard, and S. Petitjean. Near-optimal parameterization of the intersection of quadrics : I. The generic algorithm, 2007. Accepted in *Journal of Symbolic Computation*. See also INRIA Research Report n° 5667, 36 pages, 2005.
- [DLLP07b] L. Dupont, D. Lazard, S. Lazard, and S. Petitjean. Near-optimal parameterization of the intersection of quadrics : II. A classification of pencils, 2007. Accepted in *Journal of Symbolic Computation*. See also INRIA Research Report n° 5668, 37 pages, 2005.
- [DLLP07c] L. Dupont, D. Lazard, S. Lazard, and S. Petitjean. Near-optimal parameterization of the intersection of quadrics : III. Parameterizing singular intersections, 2007. Accepted in *Journal of Symbolic Computation*. See also INRIA Research Report n° 5669, 34 pages, 2005.
- [DMPT01] O. Devillers, B. Mourrain, F. Preparata, and P. Trebuchet. On circular cylinders by four or five points in space. Technical Report 4195, Inria, 2001.
- [DMPT03] O. Devillers, B. Mourrain, F.P. Preparata, and P. Trebuchet. On circular cylinders by four or five points in space. *Discrete Comput. Geom.*, 29 :83–104, 2003.
- [DR01] O. Devillers and P. Ramos. Personal communication, 2001.
- [DST93] J. H. Davenport, Y. Siret, and E. Tournier. *Computer Algebra. Systems and algorithms for algebraic computation*. Academic Press, 2nd edition, 1993.
- [Dug04] F. Duguet. Shadow computations using robust epsilon visibility. Research Report 5167, INRIA, 2004.
- [Dup04] L. Dupont. *Paramétrage quasi-optimal de l'intersection de deux quadriques : théorie, algorithme et implantation*. Thèse d'université, Université Nancy II, Oct 2004.
- [Dur99] F. Durand. *Visibilité tridimensionnelle : étude analytique et applications*. PhD thesis, Université Joseph Fourier - Grenoble I, 1999.
- [Dur00] F. Durand. A multidisciplinary survey of visibility, 2000. ACM Siggraph course notes, Visibility, Problems, Techniques, and Applications.
- [DZ02] T. K. Dey and W. Zhao. Approximate medial axis as a voronoi subcomplex. In *SMA '02 : Proceedings of the seventh ACM symposium on Solid modeling and applications*, pages 356–366, New York, NY, USA, 2002. ACM Press.
- [Ede87a] H. Edelsbrunner. *Algorithms in Combinatorial Geometry*. Springer-Verlag, Heidelberg, Germany, 1987.
- [Ede87b] H. Edelsbrunner. *Algorithms in Combinatorial Geometry*. Springer-Verlag, Heidelberg, Germany, 1987.
- [EGHHZ00] 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 Symp. on Discrete Algorithms*, pages 910–917, 2000.
- [ELL⁺06] H. Everett, S. Lazard, B. Lenhart, J. Redburn, and L. Zhang. Predicates for line transversals in 3D, 2006. Accepted in *Computational Geometry : Theory and Applications*. (Also in Proceedings of 18th Canadian Conference on Computational Geometry - CCCG'06.).
- [ELLD07] H. Everett, D. Lazard, S. Lazard, and M. Safey El Din. The voronoi diagram of three lines in \mathbb{R}^3 . In Proceedings of the *23rd ACM Annual Symposium on Computational Geometry (SoCG'07)*, S. Korea, 2007.
- [ELPZ07] H. Everett, S. Lazard, S. Petitjean, and L. Zhang. On the expected size of the 2D visibility complex. *International Journal of Computational Geometry and Applications*, 17(4) :361–382, 2007.

- [EMP⁺82] H. Edelsbrunner, H. A. Maurer, F. P. Preparata, A. L. Rosenberg, E. Welzl, and D. Wood. Stabbing line segments. *BIT*, 22 :274–281, 1982. (See also [Ede87a, Chapter 15].).
- [ER02] M. Etzion and A. Rappoport. Computing voronoi skeletons of a 3-d polyhedron by space subdivision. *Computational Geometry : Theory and Applications*, 21(3) :87–120, 2002.
- [ES90] H. Edelsbrunner and M. Sharir. The maximum number of ways to stab n convex non-intersecting sets in the plane is $2n - 2$. *Discrete Comput. Geom.*, 5 :35–32, 1990.
- [ET04] I. Emiris and E. Tsigaridas. Comparing real algebraic numbers of small degree. In *Proc. of European Symposium on Algorithms*, pages 652–663, 2004.
- [FGb] FGb - A software for computing Gröbner bases. J.-C. Faugère. <http://fgbrs.lip6.fr>.
- [Fin37] P. Finsler. über das Vorkommen definitiver und semidefinitiver Formen in Scharen quadratischer Formen. *Comment. Math. Helv.*, 9 :188–192, 1936/1937.
- [FNO89] R. Farouki, C. Neff, and M. O’Connor. Automatic parsing of degenerate quadric-surface intersections. *ACM Transactions on Graphics*, 8(3) :174–203, 1989.
- [For97] S. Fortune. Voronoi diagrams and delaunay triangulations. In *Handbook of discrete and computational geometry*, pages 377–388. CRC Press, Inc., Boca Raton, FL, USA, 1997.
- [Geo] Geomview. <http://www.geomview.org>.
- [GL06] M. Glisse and S. Lazard. An upper bound on the average size of silhouettes, 2006. Submitted to *Discrete and Computational Geometry*.
- [GL07] M. Glisse and S. Lazard. An upper bound on the average size of silhouettes. Research Report n° 6124, INRIA, Sept. 2007. 19 pages.
- [Gla07] L. Graves. An exploration of the 3d visibility complex. Master’s thesis, Polytechnic University, Brooklyn, NY., 2007.
- [Gli06] M. Glisse. An upper bound on the average size of silhouettes. In *Proceedings of the 22nd ACM Annual Symposium on Computational Geometry (SoCG’06)*, pages 105–111, June 2006.
- [Gli07] M. Glisse. Lower bounds on lines and line segments in the presence of balls, 2007. Manuscript.
- [GM91] R. Goldman and J. Miller. Combining algebraic rigor with geometric robustness for the detection and calculation of conic sections in the intersection of two natural quadric surfaces. In *Proc. of ACM Symposium on Solid Modeling Foundations and CAD/CAM Applications*, pages 221–231, 1991.
- [GMP] GMP : The GNU MP Bignum Library. The Free Software Foundation. <http://www.swox.com/gmp>.
- [Goa04] X. Goaoc. *Structures de visibilité globales : tailles, calculs et dégénérescences*. Thèse d’université, Université Nancy 2, May 2004.
- [GPW93a] J. E. Goodman, R. Pollack, and R. Wenger. Geometric transversal theory. In J. Pach, editor, *New Trends in Discrete and Computational Geometry*, pages 163–198. Springer Verlag, Heidelberg, 1993.
- [GPW93b] J. E. Goodman, R. Pollack, and R. Wenger. Geometric transversal theory. In J. Pach, editor, *New Trends in Discrete and Computational Geometry*, pages 163–198. Springer Verlag, Heidelberg, 1993.
- [GS92] G. R. Grimmett and D. R. Stirzaker. *Probability and Random Processes (2nd edition)*. Clarendon Press - Oxford, 1992.
- [HCK⁺99] K. Hoff, T. Culver, J. Keyser, M. Lin, and D. Manocha. Fast computation of generalized Voronoi diagrams using graphics hardware. *Computer Graphics*, 33(Annual Conference Series) :277–286, 1999. Proceedings of ACM SIGGRAPH 1999.
- [HCV52] D. Hilbert and S. Cohn-Vossen. *Geometry and the Imagination*. Chelsea Publishing Company, New York, 1952.
- [Hec92] P. Heckbert. Discontinuity meshing for radiosity. In *Proc. of Eurographics Workshop on Rendering, Bristol*, pages 203–216, 1992.

- [HH02] O. Hall-Holt. *Kinetic Visibility*. PhD thesis, Stanford University, 2002.
- [HI95] C.-K. Hung and D. Ierardi. Constructing convex hulls of quadratic surface patches. In *Proceedings of the 7th Canadian Conference on Computational Geometry (CCCG'95)*, pages 255–260, 1995.
- [HLHS03] J.-M. Hasenfratz, M. Lapierre, N. Holzschuch, and F. Sillion. A survey of real-time soft shadows algorithms. *Computer Graphics Forum*, 22(4) :753–774, 2003. Proceedings of Eurographics'03.
- [Hor06] S. Hornus. *Maintenance de la visibilité depuis un point mobile, et applications*. PhD thesis, Université Grenoble I – Joseph Fourier, 2006.
- [HP53] W. Hodge and D. Pedoe. *Methods of Algebraic Geometry*. Cambridge University Press, 1953. Volumes I and II.
- [HS94] D. Halperin and M. Sharir. New bounds for lower envelopes in three dimensions, with applications to visibility in terrains. *Discrete and Computational Geometry*, 12 :313–326, 1994.
- [HSD94] N. Holzschuch, F. Sillion, and G. Drettakis. An efficient progressive refinement strategy for hierarchical radiosity. In *Proc. of 5th Eurographics Workshop on Rendering*, pages 353–357, 1994. Photorealistic Rendering Techniques, Focus on Computer Graphics Series.
- [HT99] M. Hohmeyer and S. Teller. Determining the lines through four lines. *Journal of Graphics Tools*, 4(3) :11–22, 1999.
- [IFH⁺03] T. Isenberg, B. Freudenberg, N. Halper, S. Schlechtweg, and T. Strothotte. A developer's guide to silhouette algorithms for polygonal models. *IEEE Comput. Graph. Appl.*, 23(4) :28–37, 2003.
- [iMo] iMovie, Apple Computer Inc.
- [Kar04] M. I. Karavelas. A robust and efficient implementation for the segment voronoi diagram. In *International Symposium on Voronoi Diagrams in Science and Engineering*, pages 51–62, 2004.
- [KCF⁺04] J. Keyser, T. Culver, M. Foskey, S. Krishnan, and D. Manocha. ESOLID : A system for exact boundary evaluation. *Computer-Aided Design*, 36(2) :175–193, 2004.
- [KKM99] J. Keyser, S. Krishnan, and D. Manocha. Efficient and accurate B-Rep generation of low degree sculptured solids using exact arithmetic : I – Representations, II – Computation. *Computer Aided Geometric Design*, 16(9) :841–859, 861–882, 1999.
- [KLL92] M. Katchalski, T. Lewis, and A. Liu. The different ways of stabbing disjoint convex sets. *Discrete Comput. Geom.*, 7 :197–206, 1992.
- [KLPS86] K. Kedem, R. Livne, J. Pach, and M. Sharir. On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles. *Discrete and Computational Geometry*, 1 :59–71, 1986.
- [KOS00] K. Kurdyka, P. Orro, and S. Simon. Semialgebraic Sard theorem for generalized critical values. *Journal of Differential Geometry*, 56 :67–92, 2000.
- [KS03] V. Koltun and M. Sharir. Three dimensional euclidean voronoi diagrams of lines with a fixed number of orientations. *SIAM Journal on Computing*, 32(3) :616–642, 2003.
- [KV01] M. Katz and K. Varadarajan. A tight bound on the number of geometric permutations of convex fat objects in \mathbb{R}^d . *Discrete Comput. Geom.*, 26 :546–548, 2001.
- [KW97] L. Kettner and E. Welzl. Contour edge analysis for polyhedron projections. In W. Strasser, R. Klein, and R. Rau, editors, *Geometric Modeling : Theory and Practice*, pages 379–394. Springer, 1997.
- [Lam73] T. Lam. *The Algebraic Theory of Quadratic Forms*. W.A. Benjamin, Reading, MA, 1973.
- [Lau94] A. Laurentini. The visual hull concept for silhouette-based image understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(2) :150–162, February 1994.
- [Lau05] G. Lauvaux. *La réalisation d'œuvres d'art par prototypage rapide avec le procédé de Stratoconception*. Thèse de doctorat, Université de Reims - Champagne Ardennes, Jun 2005.
- [Lax98] P. Lax. On the discriminant of real symmetric matrices. *Communications on pure and applied mathematics*, 51(11-12) :1387–1396, 1998.
- [LED] LEDA : Library of Efficient Data Types and Algorithms. <http://www.algorithmic-solutions.com/>.

- [Lev76] J. Levin. A parametric algorithm for drawing pictures of solid objects composed of quadric surfaces. *Communications of the ACM*, 19(10) :555–563, 1976.
- [Lev79] J. Levin. Mathematical models for determining the intersections of quadric surfaces. *Computer Graphics and Image Processing*, 11(1) :73–87, 1979.
- [LiD] LiDIA : A C++ Library for Computational Number Theory. Darmstadt University of Technology. <http://www.informatik.tu-darmstadt.de/TI/LiDIA>.
- [LLB04] G. Lauvaux, S. Lazard, and C. Barlier. Orientation des pièces artistiques pour le procédé de strato-conception. In *10th European Forum on Rapid Prototyping, Paris, France*, Sept. 2004.
- [LPP04] S. Lazard, L. M. Peñaranda, and S. Petitjean. Intersecting quadrics : An efficient and exact implementation. In *Proceedings of the 20th ACM Annual Symposium on Computational Geometry (SoCG'04)*, pages 419–428, Brooklyn, NY, 2004.
- [LPP06] S. Lazard, L. M. Peñaranda, and S. Petitjean. Intersecting quadrics : An efficient and exact implementation. *Computational Geometry : Theory and Applications*, 35(1-2) :74–99, 2006. (Special issue of the 20th ACM Symposium on Computational Geometry, 2004.).
- [LPY04] C. Li, S. Pion, and C. Yap. Recent progress in exact geometric computation. *J. of Logic and Algebraic Programming*, 64(1) :85–111, 2004. Special issue on “Practical Development of Exact Real Number Computation”.
- [Mag] The Magma Computational Algebra System. School of Mathematics and Statistics of the University of Sydney. <http://magma.maths.usyd.edu.au/magma>.
- [Map] The Maple System. Waterloo Maple Software. <http://www.maplesoft.com>.
- [Mat99] A. M. Mathai. *An Introduction to Geometrical Probability : Distributional Aspects with Applications*. Gordon and Breach Sciences Publishers, 1999.
- [McG04] M. McGuire. Observations on silhouette sizes. *Journal of Graphics Tools*, 9(1) :1–12, 2004.
- [MDC05] K. Mesquita Da Costa. Visibilité 3D : calculs des segments libres tangents à des sphères unitaires. DEA, Université Henri Poincaré Nancy I, 2005.
- [Meg01] G. Megyesi. Lines tangent to four unit spheres with coplanar centers. *Discrete Comput. Geom.*, 26 :493–497, 2001.
- [MG95] J. Miller and R. Goldman. Geometric algorithms for detecting and calculating all conic sections in the intersection of any two natural quadric surfaces. *Graphical Models and Image Processing*, 57(1) :55–66, 1995.
- [Mig82] M. Mignotte. Identification of algebraic numbers. *CJournal of Algorithms*, 3(3) :197–204, 1982.
- [Mil87] J. Miller. Geometric approaches to nonplanar quadric surface intersection curves. *ACM Transactions on Graphics*, 6(4) :274–307, 1987.
- [Mil93] V. J. Milenkovic. Robust construction of the voronoi diagram of a polyhedron. In *Proceedings of the 5th Canadian Conference on Computational Geometry (CCCG'93)*, pages 473–478, 1993.
- [MKK98] P. Mamassian, D. C. Knill, and D. Kersten. The perception of cast shadows. *Trends in Cognitive Sciences*, 2(8) :288–295, 1998.
- [MPT01] I. Macdonald, J. Pach, and T. Theobald. Common tangents to four unit balls in \mathbb{R}^3 . *Discrete and Computational Geometry*, 26(1) :1–17, 2001.
- [MS05] G. Megyesi and F. Sottile. The envelope of lines meeting a fixed line and tangent to two spheres. *Discrete Comput. Geom.*, 33(4) :617–644, 2005.
- [MST03] G. Megyesi, F. Sottile, and T. Theobald. Common transversals and tangents to two lines and two quadrics in \mathbb{P}^3 . *Discrete Comput. Geom.*, 30 :543–571, 2003.
- [MTT03] B. Mourrain, J.-P. Tédécourt, and M. Teillaud. Predicates for the sweeping of an arrangement of quadrics in 3D. Technical Report ECG-TR-242205-01, Effective Computational Geometry for Curves and Surfaces (European project), 2003.
- [MTT05] B. Mourrain, J.-P. Tédécourt, and M. Teillaud. On the computation of an arrangement of quadrics in 3D. *Computational Geometry : Theory and Applications*, 30(2) :145–164, 2005. Special issue, 19th European Workshop on Computational Geometry.

- [Nam84] M. Namba. *Geometry of Projective Algebraic Curves*. Marcel Dekker, Inc., 1984.
- [NFLYCO99] B. Nadler, G. Fibich, S. Lev-Yehudi, and D. Cohen-Or. A qualitative and quantitative visibility analysis in urban scenes. *Computers and Graphics*, 23(5) :655–666, 1999. Special issue on Visibility - Techniques and Applications.
- [NN83] T. Nishita and E. Nakamae. Half-tone representation of 3D objects illuminated by area sources or polyhedron sources. In *IEEE Computer Society's 7th International Computer Software and Applications Conference (COMPSAC '83)*, pages 237–242, 1983.
- [OBSC00] A. Okabe, B. Boots, K. Sugihara, and S. N. Chiu. *Spatial Tessellations - Concepts and Applications of Voronoi Diagrams*. John Wiley, 2nd edition, 2000.
- [O'R98] J. O'Rourke. *Computational Geometry in C*. Cambridge University Press, 2nd edition, 1998.
- [ORDP96] R. Orti, S. Rivière, F. Durand, and C. Puech. Radiosity for dynamic scenes in flatland with the visibility complex. In *Comput. Graph. Forum : Proceedings Eurographics'96*, volume 15, pages 237–248, Poitiers, 1996.
- [OSS87] S. Ocken, J. T. Schwartz, and M. Sharir. Precise implementation of CAD primitives using rational parametrization of standard surfaces. In J. T. Schwartz, M. Sharir, and J. Hopcroft, editors, *Planning, Geometry, and Complexity of Robot Motion*, pages 245–266. Ablex Publishing, 1987.
- [Pap91] A. Papoulis. *Probability, Random Variables, and Stochastic Processes*. McGraw-Hill, 1991. 3rd ed.
- [PD90] H. Plantinga and C. Dyer. Visibility, occlusion, and the aspect graph. *International Journal of Computer Vision*, 5(2) :137–160, 1990.
- [Pel93] M. Pellegrini. Ray shooting on triangles in 3-space. *Algorithmica*, 9 :471–494, 1993.
- [Pel94] M. Pellegrini. On lines missing polyhedral sets in 3-space. *Discrete and Computational Geometry*, 12 :203–221, 1994.
- [Per95] D. Perrin. *Géométrie algébrique : une introduction*. InterEditions, 1995.
- [Pet95] S. Petitjean. The number of views of piecewise-smooth algebraic objects. In E. Mayr and C. Puech, editors, *Proceedings of the Symposium on Theoretical Aspects of Computer Science (STACS'95)*, volume 900 of *Lecture Notes in Computer Science*, pages 571–582. Springer-Verlag, 1995. München, Germany, March 2-4.
- [Poo01] B. Poonen. Computing rational points on curves. In *Number Theory for the Millenium*. A. K. Peters, Boston, 2001. Proc. of Millennial Conference on Number Theory, 2000.
- [PS92] M. Pellegrini and P. Shor. Finding stabbing lines in 3-space. *Discrete Comput. Geom.*, 8 :191–208, 1992.
- [PS03] P. Parillo and B. Sturmfels. Minimizing polynomial functions. In *Algorithmic and quantitative real algebraic geometry*, volume 60 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 83–99. AMS, 2003.
- [PT06] S. Pion and M. Teillaud. 3d triangulations. In CGAL Editorial Board, editor, *CGAL-3.2 User and Reference Manual*. 2006.
- [PV96a] M. Pocchiola and G. Vegter. Topologically sweeping visibility complexes via pseudo-triangulations. *Discrete and Computational Geometry*, 16(4) :419–453, 1996. Proceedings of the 11th ACM Annual Symposium on Computational Geometry (SoCG'95).
- [PV96b] M. Pocchiola and G. Vegter. The visibility complex. *International Journal of Computational Geometry and Applications*, 6(3) :279–308, 1996. Proceedings of the 9th ACM Annual Symposium on Computational Geometry (SoCG'93).
- [PW01] H. Pottmann and J. Wallner. *Computational Line Geometry*. Mathematics and Visualization. Springer-Verlag, Berlin, 2001.
- [RAG] RAG'Lib - A library for real algebraic geometry. M. Safey El Din. <http://www-calfor.lip6.fr/~safey/RAGLib/>.
- [Red03] J. Redburn. Robust computation of the non-obstructed line segments tangent to four amongst n triangles. Master's thesis, B.A. Thesis, Williams College, Massachusetts, May 2003., 2003.

- [Riv97] S. Rivière. Dynamic visibility in polygonal scenes with the visibility complex. In *Proceedings of the 13th ACM Annual Symposium on Computational Geometry (SoCG'97)*, pages 421–423, Nice, June 1997.
- [RRSED00] F. Rouillier, M.-F. Roy, and M. Safey El Din. Finding at least one point in each connected component of a real algebraic set defined by a single equation. *Journal of Complexity*, 16 :716–750, 2000.
- [RZ04] F. Rouillier and P. Zimmermann. Efficient isolation of polynomial's real roots. *Journal of Computational and Applied Mathematics*, 162(1) :33–50, 2004.
- [SA95] M. Sharir and P. K. Agarwal. *Davenport-Schinzel Sequences and their Geometric Applications*. Cambridge University Press, 1995.
- [Sal15] G. Salmon. *A Treatise on the Analytic Geometry of Three Dimensions*, volume II. Dublin, 5th edition, 1915.
- [San76] L. Santaló. *Integral Geometry and Geometric Probability*, volume 1 of *Encyclopedia of Mathematics and its Applications*. Addison-Wesley, 1976.
- [Sar83] R. Sarraga. Algebraic methods for intersections of quadric surfaces in GMSOLID. *Computer Vision, Graphics, and Image Processing*, 22 :222–238, 1983.
- [Sch90] H. A. Schwarz. Sur une définition erronée de l'aire d'une surface courbe. In *Gesammelte Mathematische Abhandlungen*, volume 1, pages 309–311. Springer-Verlag, 1890.
- [Sch00] E. R. Scheinerman. When close enough is close enough. *American Mathematical Monthly*, 107 :489–499, 2000.
- [SED06] M. Safey El Din. Generalized critical values and testing sign conditions on a polynomial. In D. Wang and Z. Zheng, editors, *Proceedings of Mathematical Aspects of Computer and Information Science*, pages 61–84. 2006.
- [SEDS03] M. Safey El Din and E. Schost. Polar varieties and computation of at least one point in each connected component of a smooth real algebraic set. In *Proceedings of the International Symposium on Symbolic and Algebraic Computation (ISSAC'03)*, pages 224–231, Philadelphia PA, 2003. ACM Press.
- [SEDS04] M. Safey El Din and E. Schost. Properness defects of projections and computation of one point in each connected component of a real algebraic set. *Discrete and Computational Geometry*, 32(3) :417–430, 2004.
- [Seg83] C. Segre. Studio sulle quadriche in uno spazio lineare ad un numero qualunque di dimensioni. *Mem. della R. Acc. delle Scienze di Torino*, 36(2) :3–86, 1883.
- [Sei81] R. Seidel. A convex hull algorithm optimal for point sets in even dimensions. M.Sc. thesis, Univ. British Columbia, Vancouver, BC, 1981. Report 81/14.
- [Sei91] R. Seidel. A simple and fast incremental randomized algorithm for computing trapezoidal decompositions and for triangulating polygons. *Computational Geometry : Theory and Applications*, 1(1) :51–64, 1991.
- [SG94] A. J. Stewart and S. Ghali. Fast computation of shadow boundaries using spatial coherence and backprojections. In *Computer Graphics Proceedings, Annual Conference Series, ACM SIGGRAPH*, pages 231–238. ACM Press, 1994.
- [SGD] The SGDL Platform : A Software for Interactive Modeling, Simulation and Visualization of Complex 3D Scenes. SGDL Systems, Inc. <http://www.sgd1-sys.com>.
- [SGG+00] P. V. Sander, Xianfeng Gu, Steven J. Gortler, Hugues Hoppe, and John Snyder. Silhouette clipping. In *SIGGRAPH '00 : Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 327–334, New York, NY, USA, 2000. ACM Press / Addison-Wesley Publishing Co.
- [Sha94] M. Sharir. Almost tight upper bounds for lower envelopes in higher dimensions. *Discrete and Computational Geometry*, 12 :327–345, 1994.
- [Sho98] K. Shoemake. Plücker coordinate tutorial. *Ray Tracing News*, 11(1), 1998.

- [SJ92] C.-K. Shene and J. Johnstone. Computing the intersection of a plane and a natural quadric. *Computer & Graphics*, 12(2) :179–186, 1992.
- [SJ94] C.-K. Shene and J. Johnstone. On the lower degree intersections of two natural quadrics. *ACM Transactions on Graphics*, 13(4) :400–424, 1994.
- [SKHBS02] L. Szirmay-Kalos, V. Havran, B. Balázs, and L. Szécsi. On the efficiency of ray-shooting acceleration schemes. In *Proceedings of the Spring Conference on Computer Graphics*, pages 89–98, Budmerice, Slovakia, 2002.
- [Ski97] S. S. Skiena. *The Algorithm Design Manual*. Mathematics and Visualization. Springer-Verlag, 1997.
- [SKM98] L. Szirmay-Kalos and G. Márton. Worst-case versus average case complexity of ray-shooting. *Computing*, 61(2), 1998.
- [SMS00] S. Smorodinski, J. S. B. Mitchell, and M. Sharir. Sharp bounds on geometric permutations for pairwise disjoint balls in \mathbb{R}^d . *Discrete Comput. Geom.*, 23 :247–259, 2000.
- [SOS] SOSTOOLS - A MATLAB toolbox for sums of squares optimization programs. P. Parillo, A. Papachristodoulou, S. Prajna, and P. Seiler. <http://www.cds.caltech.edu/sostools/>.
- [SS97] O. Schwarzkopf and M. Sharir. Vertical decomposition of a single cell in a three-dimensional arrangement of surfaces and its applications. *Discrete and Computational Geometry*, 18 :269–288, 1997.
- [ST02] F. Sottile and T. Theobald. Lines tangent to $2n - 2$ spheres in \mathbb{R}^n . *Transactions of the AMS*, 354(12) :4815–4829, 2002.
- [Sto91] J. Stolfi. *Oriented Projective Geometry : A Framework for Geometric Computations*. Academic Press, 1991.
- [Sug00] K. Sugihara. How to make geometric algorithms robust. *IEICE Transactions on Information and Systems*, E83-D(3) :447–454, 2000.
- [Sur] Surf : A Tool To Visualize Algebraic Curves and Surfaces. Stephan Endrass, Hans Huelf, Ruediger Oertel, Ralf Schmitt, Kai Schneider and Johannes Beigel. <http://surf.sourceforge.net>.
- [SW06] E. Schömer and N. Wolpert. An exact and efficient approach for computing a cell in an arrangement of quadrics. *Computational Geometry : Theory and Applications*, 33(1-2) :65–97, 2006. Special Issue on Robust Geometric Algorithms and their Implementations.
- [Tel92] S. J. Teller. Computing the antipenumbra of an area light source. In *Computer Graphics Proceedings, Annual Conference Series, Annual Conference Series*, ACM SIGGRAPH, pages 139–148. ACM Press, 1992.
- [The97] T. Theobald. New algebraic methods in computational geometry. Habilitationsschrift, Fakultät für Mathematik, Technische Universität München, 1997.
- [The02] T. Theobald. An enumerative geometry framework for algorithmic line problems in \mathbb{R}^3 . *SIAM Journal on Computing*, 31 :1212–1228, 2002.
- [TT97] M. Teichmann and S. Teller. Polygonal approximation of voronoi diagrams of a set of triangles in three dimensions. Technical Report 766, Laboratory of Computer Science, MIT, 1997.
- [TWMW05] C. Tu, W. Wang, B. Mourrain, and J. Wang. Signature sequence of intersection curve of two quadrics for exact morphological classification. Technical Report Research Report TR-2005-09, The University of Hong-Kong, September 2005.
- [TWW02] C. Tu, W. Wang, and J. Wang. Classifying the nonsingular intersection curve of two quadric surfaces. In *Proc. of GMP'02 (Geometric Modeling and Processing)*, pages 23–32, 2002.
- [Uhl73] F. Uhlig. Simultaneous block diagonalization of two real symmetric matrices. *Linear Algebra and Its Applications*, 7 :281–289, 1973.
- [Uhl76] F. Uhlig. A canonical form for a pair of real symmetric matrices that generate a nonsingular pencil. *Linear Algebra and Its Applications*, 14 :189–209, 1976.
- [Wan92] L. Wanger. The effect of shadow quality on the perception of spatial relationships in computer generated images. *Computer Graphics*, 25(2) :39–42, 1992.

- [Wei68] K. Weierstrass. Zur Theorie der bilinearen und quadratischen Formen. *Monatshefte Akademie der Wissenschaften, Berlin*, pages 310–338, 1868.
- [Wen90] R. Wenger. Upper bounds on geometric permutations for convex sets. *Discrete Comput. Geom.*, 5 :27–33, 1990.
- [Wen98] R. Wenger. Progress in geometric transversal theory. In B. Chazelle, J. E. Goodman, and R. Pollack, editors, *Advances in Discrete and Computational Geometry*, pages 375–393. Amer. Math. Soc., Providence, 1998.
- [WGT03] W. Wang, R. Goldman, and C. Tu. Enhancing Levin’s method for computing quadric-surface intersections. *Computer-Aided Geometric Design*, 20(7) :401–422, 2003.
- [WJG97] W. Wang, B. Joe, and R. Goldman. Rational quadratic parameterizations of quadrics. *International Journal of Computational Geometry and Applications*, 7(6) :599–619, 1997.
- [WJG02] W. Wang, B. Joe, and R. Goldman. Computing quadric surface intersections based on an analysis of plane cubic curves. *Graphical Models*, 64(6) :335–367, 2002.
- [WK04] W. Wang and R. Krasauskas. Interference analysis of conics and quadrics. In R. Goldman and R. Krasauskas, editors, *Topics in Algebraic Geometry and Geometric Modeling : Workshop on Algebraic Geometry and Geometric Modeling, July 29-August 2, 2002, Vilnius University, Lithuania*, volume 334, pages 25–36. AMS Contemporary Mathematics, 2004.
- [WM93] I. Wilf and Y. Manor. Quadric-surface intersection curves : shape and structure. *Computer-Aided Design*, 25(10) :633–643, 1993.
- [Wol02] N. Wolpert. *An Exact and Efficient Algorithm for Computing a Cell in an Arrangement of Quadrics*. PhD thesis, Universität des Saarlandes, Saarbrücken, 2002.
- [WPF90] A. Woo, P. Poulin, and A. Fournier. A survey of shadow algorithms. *IEEE Computer Graphics and Applications*, 10(6) :13–32, 1990.
- [WS88] A. Wiernik and Micha Sharir. Planar realizations of nonlinear Davenport-Schinzel sequences by segments. *Discrete Comput. Geom.*, 3 :15–47, 1988.
- [WWK01] W. Wang, J. Wang, and M.-S. Kim. An algebraic condition for the separation of two ellipsoids. *Computer Aided Geometric Design*, 18(6) :531–539, 2001.
- [Yap97] C. K Yap. Towards exact geometric computation. *Computational Geometry : Theory and Applications*, 7(1-2) :3–23, 1997.
- [ZELW07] L. Zhang, H. Everett, S. Lazard, and S. Whitesides. Towards an implementation of the 3d visibility skeleton. In *Proceedings of the 23rd ACM Annual Symposium on Computational Geometry (SoCG’07)*, S. Korea, 2007. Video, 2 pages.