An Exact and Efficient Approach for Computing a Cell in an Arrangement of Quadrics

Dissertation
zur Erlangung des Grades
Doktor der Ingenieurswissenschaften (Dr.-Ing.)
der Naturwissenschaftlich–Technischen Fakultät I
der Universität des Saarlandes

von

Nicola Wolpert

Saarbrücken
31. Oktober 2002

Datum des Kolloquiums: 31. 10. 2002

Dekan der Naturwissenschaftlich–Technischen Fakultät I:
Professor Dr. Philipp Slusallek

Gutachter:
Professor Dr. Raimund Seidel, Universität des Saarlandes, Saarbrücken
Professor Dr. Kurt Mehlhorn, MPI für Informatik, Saarbrücken

Kurzzusammenfassung

In dieser Doktorarbeit stellen wir eine Ansatz vor, mit dessen Hilfe eine Zelle in einem Arrangement von quadratischen Flächen exakt und effizient berechnet werden kann. Alle Berechnungen basieren auf exakten algebraischen Methoden und führen selbst in degenerierten Fällen zu mathematisch korrekten Ergebnissen. Durch Projektion kann das räumliche Problem darauf reduziert werden, planare Arrangements von Kurven zu bestimmen. Es gelingt uns, alle Ereignispunkte dieser planaren Arrangements einschließlich der tangentialen Schnitte und singulären Punkte zu lokalisieren. Die nicht singulären tangentialen Schnitte bestimmen wir, indem wir eine zusätzliche Kurve, die wir *Jacobi Kurve* nennen, betrachten. Durch eine Verallgemeinerung der Jacobi Kurve sind wir in der Lage, in beliebigen planaren Arrangements nicht-singuläre tangentiale Schnitte zu bestimmen. Wir zeigen, dass die Koordinaten der singulären Punkte in unseren speziellen, durch Projektion entstandenen, planaren Arrangements Nullstellen von quadratischen Polynomen sind. Die Koeffizienten dieser Polynome sind in den meisten Fällen rational und benötigen maximal eine einzelne Quadratwurzel. Eine prototypische Implementierung zeigt, dass unser Ansatz in der Praxis zu einem guten Laufzeitverhalten führt.

Abstract

In this thesis, we present an approach for the exact and efficient computation of a cell in an arrangement of quadric surfaces. All calculations are based on exact rational algebraic methods and provide the correct mathematical results in all, even degenerate, cases. By projection, the spatial problem can be reduced to the one of computing planar arrangements of algebraic curves. We succeed in locating all event points in these arrangements, including tangential intersections and singular points. By introducing an additional curve, which we call the *Jacobi curve*, we are able to find non-singular tangential intersections. By a generalization of the Jacobi curve we are able to determine non-singular tangential intersections in arbitrary planar arrangements. We show that the coordinates of the singular points in our special projected planar arrangements are roots of quadratic polynomials. The coefficients of these polynomials are usually rational and contain at most a single square root. A prototypical implementation indicates that our approach leads to good performance in practice.

Danksagung

An dieser Stelle möchte ich mich bei den Menschen bedanken, die mich in den letzten Jahren unterstützt und zum Gelingen dieser Arbeit beigetragen haben. Zuerst ist dabei natürlich mein Doktorvater Prof. Dr. Raimund Seidel zu nennen. Danke für die gute Betreuung, Dein Engagement und viele hilfreiche Gespräche.

Die Atmosphäre an unserem Lehrstuhl und in der Arbeitsgruppe von Prof. Dr. Kurt Mehlhorn war und ist sehr anregend, kollegial und freundschaftlich. Danke an alle Leute, die diese nette Atmosphäre ausgemacht haben. Insbesondere möchte ich Dr. Elmar Schömer erwähnen. Danke für die gute Zusammenarbeit und die vielen Stunden intensiver Diskussionen.

Abschließend möchte ich noch Markus Denny und Frank Lehmann nennen. Vielen Dank für die schöne Zeit in unserem gemeinsamen Büro, für Eure Unterstützung und für Eure Freundschaft.

This work was partially supported by the IST Programme of the EU as a Shared-cost RTD (FET Open) Project und Contract No IST-2000-26473 (ECG - Effective Computational Geometry for Curves and Surfaces).

Contents

Chapter 1

Introduction

1.1 Problem

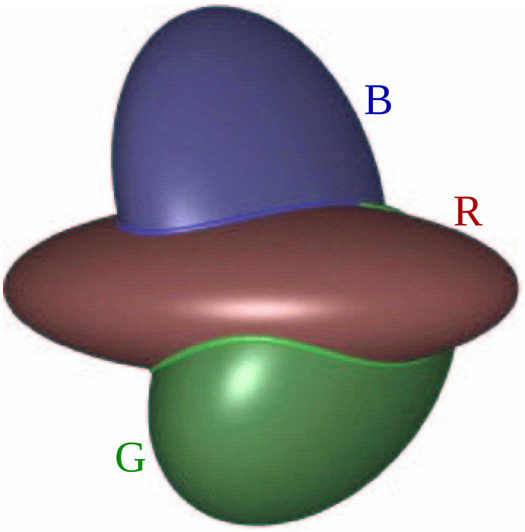

Figure 1.1: Ellipsoids are defined by quadratic polynomials

In this dissertation we consider arrangements induced by quadric surfaces in 3-dimensional space. Quadric surfaces, or quadrics for short, are defined as the set of roots of quadratic polynomials. For example, the red, green, and blue ellipsoids R, G, and B in Figure 1.1 are defined by the following polynomials:

$$
\begin{aligned}
R(x,y,z) &= 27x^2 + 62y^2 + 249z^2 - 10 \\
G(x,y,z) &= 88x^2 + 45y^2 + 67z^2 - 66xy - 25xz \\
&\quad +12yz - 24x + 2y + 29z - 5 \\
B(x,y,z) &= 139x^2 + 141y^2 + 71z^2 - 157xy \\
&\quad +97xz - 111yz - 3x - 6y - 17z - 7.
\end{aligned}
$$

On the surface of a given quadric p, the intersection curves of p with the remaining quadrics build a 2-dimensional subarrangement. In our example, the blue and the green ellipsoid intersect the red ellipsoid. This leads to two intersection curves on the surface of the red ellipsoid, a blue one and a green one (Figure 1.2). Vertices of this subarrangement are common points of two intersection curves, or rather intersection points of three quadrics.

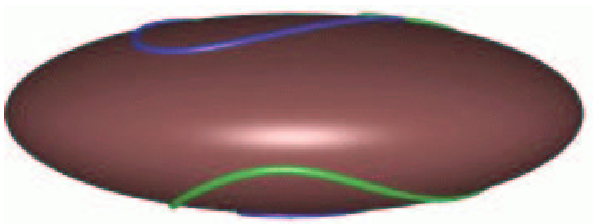

Figure 1.2: The blue and the green ellipsoid intersect the red one in two spatial curves running on the surface of the red ellipsoid

Independ of the special information about the arranged quadrics one may be interested in, for example the topological description of a cell, the basic computation that has to be done in nearly all cases is: For each quadric p, locate and sort all vertices along the intersection curves on the surface of p. This problem is particularly difficult in case one is interested in exact mathematical solution, even for degenerate cases. The coordinates of common points of three quadrics are not expressible as nested square roots of rational numbers. Instead, one has to deal with algebraic numbers. We want to develop an algorithm that is

1. *exact* in the sense that it always computes the mathematical correct result, even for degenerate inputs, and
2. *efficient* in practice concerning its running time.

Moreover, we search for an approach that conceptually is extendible to more complicated surfaces.

In general there is no rational parameterization of the intersection curve of two surfaces in space. Therefore we project for each quadric p all its intersection curves with the other quadrics and additionally its silhouette into the plane. This projection step applied to our example proceeds like shown in Figure 1.3.

We have to compute the planar arrangements resulting from the projection. All curves of the planar arrangements turn out to be defined by polynomials of degree at most 4. For example, the green curve is the set of roots of the polynomial

$$\begin{gathered}408332484x^4 + 51939673y^4 - 664779204x^3y - 24101506y^3x \\ +564185724x^2y^2 - 250019406x^3 + 17767644y^3 \\ +221120964x^2y - 123026916y^2x + 16691919x^2 + 4764152y^2 \\ +14441004xy + 10482900x + 2305740y - 1763465.\end{gathered}$$

In such arrangements of projected curves singular points and tangential intersections appear quite frequently as can be seen in the last picture of Figure 1.3. The main question with respect to exactness and efficiency is how to locate these points. In this thesis we propose a solution to this problem.

1.2 Motivation

Computing arrangements of curves and surfaces is one of the fundamental problems in different areas of computer science like computational geometry, algebraic geometry, and solid modeling. As long as arrangements of lines and planes defined by rational numbers are considered, all computations can be done over the field of rational numbers avoiding numerical errors. In this case algorithms are available that are exact and efficient.

As soon as higher degree algebraic curves and surfaces are considered, instead of linear ones, things becomes more difficult. In general the intersection points of two planar curves or three surfaces in 3-space defined by rational polynomials have irrational coordinates. That means instead of rational numbers one now has to deal with algebraic numbers. One way to overcome this difficulty is to develop algorithms that use floating point arithmetic. These algorithms are quite fast but in degenerate situations they can lead to completely wrong results because of approximation errors, rather than just slightly inaccurate outputs. Assume for example that for two planar curves one is interested in the fact whether they intersect or not. If the curves only have tangential intersection points, numerical inaccuracies can lead to the wrong result "no, they do

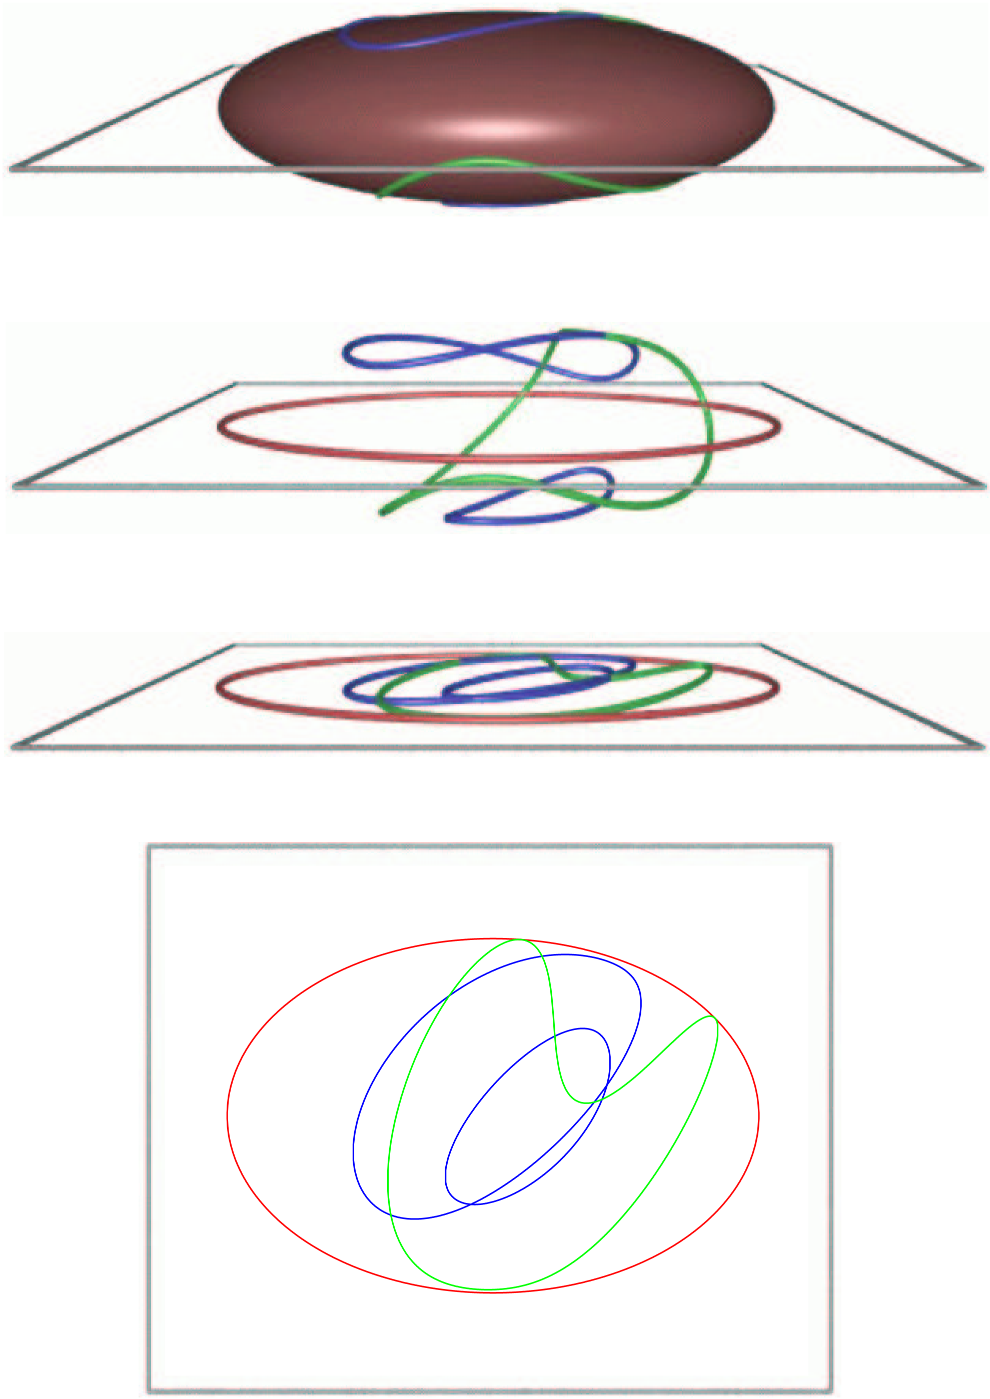

Figure 1.3: Project all intersection curves and the silhouette of p into the plane

not intersect".
A second approach besides using floating point arithmetic is to use exact algebraic computation methods. Then of course the results are correct, but the algorithms in general are very slow. So when dealing with curved surfaces it is difficult to find algorithms that are both: exact and efficient.

We consider 3-dimensional arrangements of surfaces defined by quadratic polynomials. For this special class of curved surfaces we develop an approach for computing the topological description of a cell in the arrangement. Our algorithm is exact and efficient when compared with existing exact algebraic computation methods. As an application we show how to use our results in order to compute the convex hull of ellipsoids.

1.3 Previous work

As mentioned, methods for the calculation of arrangements of algebraic surfaces are an important area of research in different branches of computer science.

1.3.1 Solid modeling

Arrangements of curved surfaces typically arise in solid modeling, see for example [37], when performing boolean operations for quadric surfaces, which play an important role in the design of mechanical parts. Most mechanical parts are build by union and difference of linear and quadratic surfaces. The algorithms in CAD systems have the advantage that they are quite fast. They profit from floating point arithmetic and often use numerical procedures for tracing the intersection curves and then approximate them as spline curves. But just this makes them very sensitive to approximation and rounding errors. Thus they achieve the good running time at the expense of exactness in degenerate situations which are nevertheless frequent in the design of geometric objects. None of these systems are exact. Recently some efforts have been made towards exact and efficient implementations:

MAPC [44] is a library for exact computation and manipulation of algebraic points and curves. It includes an algorithm for computing the arrangement of curves in the plane. Degenerate situations like tangential intersections of two curves or self-intersections of one curve through which another curve cuts through are explicitly not treated. For such cases the use of the gap theorem [15] or multivariate Sturm sequences [53] is proposed. Both methods are not very efficient.

ESOLID [43] performs exact boundary evaluation of low-degree curved solids. It is stated that degenerate cases cannot be handled. For a more detailed description of

MAPC and ESOLID consider the PhD thesis of Keyser [45].

1.3.2 Computational geometry

Also in computational geometry there is a great focus on computing arrangements, but mainly on arrangements of linear objects. For a good and brief overview consider the articles of Halperin [35] and Agarwal and Sharir [2]. The geometric methods have the drawback that nearly all of them are based on an idealized real arithmetic provided by the real RAM model of computation [61]. The assumption is that all, even irrational, numbers are representable and that one can deal with them in constant time. More precisely, it is assumed that the roots of any polynomial of constant degree can be computed exactly in constant time. This postulate is not in accordance with real computers. Mulmuley's algorithm for computing the partition of the plane induced by a set of algebraic segments of bounded degree [55] is based on this model. The same holds for the algorithms of Dobkin and Souvaine [25] and of Snoeyink and Hershberger [70]. The first authors approximate planar curves by what they call splinegons. The second authors provide a sweep line algorithm for arrangements of curves. Further examples for geometric algorithms that are founded on the real RAM model are the convex hull algorithms for planar objects by Bajaj and Kim [8] and by Nielsen and Yvinec [58]. In three dimensions Schwarzkopf and Sharir construct the arrangement of curved surfaces also relying on idealized real arithmetic [67].

Algorithms coping with arrangements of lines can be implemented with exact rational arithmetic and with a good performance, because they only deal with linear algebraic primitives, see for example the implementations in LEDA [48] and CGAL [29]. For general curves the situation is more difficult. Several authors have looked into the question of using restricted predicates to report or compute segment intersections, [12], [11], and [16]. The restriction used in these papers is on the degrees of the predicates used by the algorithms. By restricting to low-degree predicates, one can generally achieve more robust computations. Predicates for arrangements of circular arcs are treated by Devillers et al. in [24]. Recently Wein [76] extended the CGAL implementation of planar maps to conic arcs and Berberich et al. [10] made a similar approach by extending the Bentley-Ottmann sweep-line algorithm [9].

For computing planar arrangements of arbitrary curves very little is known. Milenkovic [49] uses floating point arithmetic in order to compute arrangements of curves. Neagu and Lacolle [57] consider piecewise convex parametric curves and approximate them by polygons. Both approaches cannot handle all degenerate cases.

In higher dimensions Collins [19] introduced the concept of cylindrical algebraic decomposition. Although it is also used and developed further in computational geometry, we will treat it in the next section about computer algebra methods. The special case of quadric surfaces in three dimensions has been studied extensively. The section after the next is separately dedicated to this subject.

1.3.3 Algebraic geometry

Computational real algebraic geometry studies algorithmic questions dealing with real solutions of a system of equalities, inequalities, and inequations of polynomials over real numbers, see for example the overview article of Mishra [54]. Geometrically, a predicate $P < 0$, $P = 0$, or $P > 0$, with $P \in \mathbb{R}[x_1, \ldots, x_d]$ a polynomial, partitions the set of d-dimensional points into three different subsets. For a set of polynomials the real solutions to the Boolean combination of such predicates is called a semi-algebraic set. A cell in an arrangement of surfaces defined by rational polynomials can be interpreted as the connected component of a semi-algebraic set. A cell has the property that the sign of all polynomials is constant for all points of the cell.

In the one-dimensional case, the real roots of univariate polynomials partition the real line in sign-invariant intervals. A generalization of this decomposition to higher dimensions was first provided by Collins [19] as an improvement of the results obtained by Tarski [71] for quantifier elimination. The cylindrical algebraic decomposition (CAD) of Collins partitions the d-dimensional Euclidean space $\mathbb{R}^d$ into connected subsets compatible with the zeros of the polynomials. This is done by projecting all "relevant" points to the lower-dimensional space $\mathbb{R}^{d-1}$. Relevant points are, for example, intersection points of two surfaces defined by the polynomials or the points of one surface that have a tangent plane parallel to the direction of projection. A CAD is computed recursively for the resulting polynomials in $\mathbb{R}^{d-1}$, and this result is extended to a CAD for the polynomials in $\mathbb{R}^d$. Collins introduced the CAD as part of a new quantifier elimination method. Arnon, Collins, and McCallum [4] give an equivalent definition but with more emphasis on the geometric aspect. The CAD construction was extended to reporting pairs of adjacent cells in $\mathbb{R}^2$ [5] and $\mathbb{R}^3$ [6]. The construction of Collins leads to $O(n^{2^d-1})$ cells. This doubly exponential size was improved to singly exponential size by Edelsbrunner et al. [17].

In principle the cylindrical algebraic decomposition can be implemented and our algorithm is based on this method. The problem is that after the projection steps one has to compute the roots of univariate polynomials. This cannot be done explicitly but only by computing isolating intervals, each containing exactly one real root. One has to work with these algebraic numbers during the final extensions to the original space. It is an open problem how to really perform the necessary algebraic primitives in an exact and efficient way. Of course one could use the gap theorem introduced by Canny

[15] or multivariate Sturm sequences discussed by Pedersen [60] and Milne [53] but, as mentioned before, the running time of both methods is quite high.

For bivariate polynomials defining curves in the plane, some specific work has been done. Arnborg and Feng [3] discuss the algebraic decomposition of one regular curve that has no self-intersections, cusps, or isolated points. Arnon and McCallum [7] check whether a curve is non-singular and if so, they compute their topological type. Abhyankar and Bajaj [1] give a polynomial time algorithm that determines the genus of a plane algebraic curve. Their algorithm is based on the real RAM model assuming that common roots of polynomials can be computed explicitly. The only exact approach is the one by Sakkalis [64]. He uses rational arithmetic to compute the topological configuration of a single curve. He determines isolating boxes for the singular points with the help of negative polynomial remainder sequences. This last approach, although it is exact, is not very efficient, at least if singular points occur frequently. Another problem is that it deals only with one single curve. One could interpret intersection points of two curves as singular points of the curve that is the union of both, but this would lead to a loss of important information.

Mathematical algebraic geometry also deals with curves and surfaces. The perspective is more theoretical in the sense that the focus is mainly on classification of points in projective complex space and not on locating them in affine space. Nevertheless some important theorems like the one of Bézout are very useful also from the more practical computational point of view. For a good survey about the geometry of algebraic curves consider the book of Gibson [33].

Of course, in algebraic geometry and computer algebra some effort was made in developing software. For example, LiDIA [59] is a library for computational number theory. APU is [63] a tool for real algebraic numbers. Core [41] and LEDA [48] are libraries that address the issues of robust numerical and geometric computation.

1.3.4 Quadric surface intersection

Quadric surfaces, quadrics for short, are surfaces defined by a quadratic polynomial. They are of great importance because they are the simplest of all curved surfaces and they are widely used in the design of mechanical parts. Levin [46], [47] introduced a pencil method for computing an explicit parametric representation of the intersection between two quadrics. Arguing that Levin's method does not take advantage of the fact that degenerate intersection curves admit a rational parameterization, Farouki, Neff, and O'Connor [28] made a complete study of degenerate cases for arbitrary quadric surfaces.

Miller [51] and Goldman and Miller [50] noticed that Levin's algorithm is numerically sensitive because it is based on solutions of polynomials of degree 4. They developed an approach for natural quadrics, that means spheres, cylinders, cones, and planes, that does not require solutions to polynomials of degree higher than 2. Later they used the pencil method to detect and compute conic intersections between pairs of natural quadrics [52]. The same problem was considered by Shene and Johnstone [69].

For three quadric surfaces, Chionh, Goldman, and Miller [18] used multi-resultants to determine a univariate polynomial of degree at most 8, the real roots of which are the x-coordinates of the common intersection points of the three quadrics.

Wang, Joe, and Goldman [74] investigated the different parameterizations of irreducible quadrics and the relation between them. Later, the same authors computed the intersection curve of only two quadrics [75]. Their algorithm is based on the result that the intersection of two quadrics is birationally related to a plane cubic curve.

Interval arithmetic is used by Geismann, Hemmer, and Schömer [31] to keep track of all occurring rounding and approximation errors that appear in Levin's algorithm while computing a cell in an arrangement of quadrics. If the input does not lie too near to a degenerate configuration, the algorithm will succeed in predicting the correct topological structure of the intersection. Otherwise it can detect the existence of a critical situation.

Recently, Dupont, Lazard, Lazard, and Petitjean [26] improved Levin's method for computing parameterizations for the intersection of two arbitrary implicit quadrics. Their parameterization is nearly as rational as possible, meaning that its coefficients are contained in the smallest possible field extension, up to a unique perhaps unnecessary square root.

1.4 Our contribution

Computing the mathematical correct topology of a cell in an arrangement of curved surfaces efficiently in any case, even a degenerate one, is a challenging task. As far as we know, we are the first who provide such an algorithm for a set of *quadric* surfaces in 3 dimensions [31]. Our algorithm uses exact rational algebraic computation and it can handle each degenerate input. A prototypical implementation shows that the theoretical results promise a good performance in practice.

The 3-dimensional problem of computing arrangements of surfaces defined by quadratic polynomials is more complicated than the corresponding problem in the plane. In the plane transversal intersections of two arbitrary curves can be detected with a known tool called box hit counting. It fails for singular points of one curve and for tangential intersection points of two curves. But a quadratic curve can have at most 1 singular point and two such curves intersect tangentially in at most 2 points. The algorithms in the plane benefit from the fact that the coordinates of these points can be computed as roots of univariate polynomials of degree at most 2. This makes them quite easy to handle. In 3 dimensions the situation is different. A spatial intersection curve can have up to 4 self-intersections and two curves can touch each other in 4 points. Although explicit solutions for univariate polynomials of degree at most 4 exist, the ones for polynomials of degree greater than 2 are based on solving the problem of dividing an angle into three different equal parts (casus irreducibilis). This is known to be not exactly solvable with algebraic methods, and thus the real solutions of polynomials of degree greater than 2 cannot be computed any longer as nested square roots of rational numbers. This makes the problem in 3 dimensions difficult to solve with exact algebraic computation, at least in degenerate cases.

Our approach operates similarly to the cylindrical algebraic decomposition [19]. It reduces the 3-dimensional problem to the one of computing planar arrangements of algebraic curves of degree up to 4. During the calculation we use algebraic techniques like resultants and root separation of univariate polynomials. The reduction is algebraically optimal in the sense that it does not affect the algebraic degree of the problem we consider.

Our contribution, and what is new, is that we succeed in determining *all* event points in the planar arrangement, including tangential intersection points and singular points, while keeping the running time low. The curves in the planar arrangements can have 6 singular points and two of them intersect tangentially in up to 8 points. Locating the event points in the planar arrangements works for two reasons:

1. We show that determining non-singular tangential intersection points can be reduced to the problem of locating transversal intersection points. For the latter we know that simple box hit counting is a suitable tool. This is done by introducing a new curve to the arrangement. To the best of our knowledge we are the first who consider an auxiliary curve in order to solve tangential intersections. We additionally prove that our construction based on auxiliary curves can be extended to general planar arrangements.

2. We succeed in factoring univariate polynomials in a way that the coordinates of singular points are roots of quadratic rational polynomials. Only in the case that a curve consists of four lines, computing the coordinates requires a second square root.

In one sense the work recently done by Dupont, Lazard, Lazard, and Petitjean [26] leads to the same result as ours, namely that computing with quadric surfaces can be done exactly and efficiently in all cases, just working over the rationals with only few additional square roots. But their approach directly works in space and searches for a parameterization of the intersection curves. This way of solving the problem is not extendible to more complicated surfaces. The methods presented in our work can also be applied to arbitrary curved surfaces. Only for computing singular points in the planar arrangement we make use of the fact that we consider quadric input surfaces. In that sense our work is a first and important step towards an efficient and exact algorithm for computing arrangements of arbitrary curved surfaces.

1.5 Outline

The organization of the remaining chapters is as follows:

- In Chapter 2 we provide the notation and mathematical tools we need for our approach. This includes resultants, subresultants, and root separation.
- In Chapter 3 we sketch the overall structure of our algorithm for computing a cell in an arrangement of n quadrics. We show how to reduce the 3-dimensional problem to n planar ones. We introduce simple box hit counting as a tool for determining transversal intersections of two curves.
- Chapter 4 first provides a method for distinguishing transversal intersection points from tangential intersection points and from singular points. Afterwards, we define an auxiliary curve, resulting in a new method for determining non-singular intersections called extended box hit counting. Finally, we prove how to extend this argumentation to non-singular intersections of arbitrary curves.
- Chapter 5 deals with the singular points of the curves we obtain from the reduction. We classify them in two different groups. We prove that one of the groups contains more than 2 singular points only in the case that the spatial intersection curve of two quadrics consists of two lines and another conic curve. The pencil method of Levin is introduced to solve this special case.
- In Chapter 6 we prove our main theorem, namely that every event point in the planar arrangement can be determined.
- Throughout the previous chapters some generality assumptions concerning the planar curves are made. In Chapter 7 we explain how to test and achieve these assumptions for general input.
- In Chapter 8 we show how to apply our results via duality to the convex hull problem of a set of ellipsoids.

- Finally, in Chapter 9 we discuss the results we obtained from a prototype implementation of determining the event points in the plane and give the prospects for further research.

Chapter 2

The mathematical tools

In this chapter we will introduce the algebraic concepts that are behind our computations and that enable us to deal with algebraic surfaces and curves. First, we will clarify the main algebraic terms. Afterwards, we introduce the concept of resultants and subresultants. Finally, we give a short description of the algorithm of Uspensky for root isolation of univariate polynomials. All the ideas presented are fundamental in algebraic geometry and computer algebra, see for example [22], [72], and [77]. We will give some further references to books and articles on this subject in the context. This chapter mainly deals as an introduction of basic algebraic tools, collected from the aspect of computing arrangements of quadric surfaces.

2.1 Notation

2.1.1 Curves and surfaces

The objects we consider and manipulate in our work are algebraic surfaces and curves represented by rational polynomials. More generally, we define an *algebraic hypersurface* in the following way: Let f be a polynomial in $\mathbb{Q}[x_1, \ldots, x_d]$. We set

$$\textsc{zero}(f) := \{(a_1, \ldots, a_d) \in \mathbb{R}^d \mid f(a_1, \ldots, a_d) = 0\}$$

and call $\textsc{zero}(f)$ the *algebraic hypersurface* defined by f. We reserve the terms *algebraic surface* and *algebraic curve* for the special cases $d = 3$ and $d = 2$, respectively.

For example, the red, green, and blue ellipsoids R, G, and B in Figure 2.1 are defined by the following polynomials:

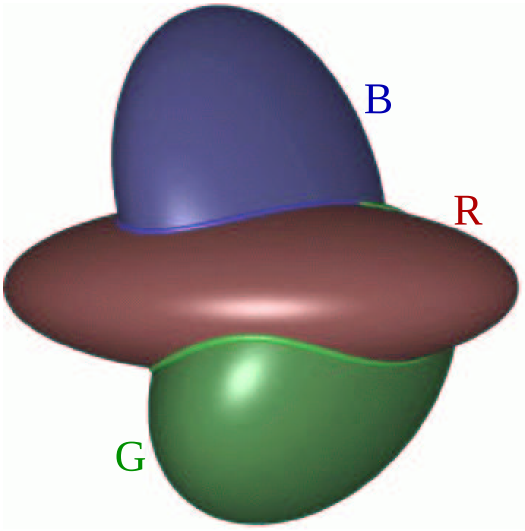

Figure 2.1: Ellipsoids are defined by quadratic polynomials

$$\begin{aligned} R(x,y,z) &= 27x^2 + 62y^2 + 249z^2 - 10 \\ G(x,y,z) &= 88x^2 + 45y^2 + 67z^2 - 66xy - 25xz \\ &\quad +12yz - 24x + 2y + 29z - 5 \\ B(x,y,z) &= 139x^2 + 141y^2 + 71z^2 - 157xy \\ &\quad +97xz - 111yz - 3x - 6y - 17z - 7. \end{aligned}$$

The *total degree* of an algebraic hypersurface is the highest degree of all monomials of its defining polynomial. Thus, ellipsoids are degree 2 algebraic surfaces. We call degree 2 algebraic surfaces *quadric* surfaces, or *quadrics* for short. Our attention is focused on computing with quadric surfaces although some of our ideas can be transfered to algebraic surfaces of arbitrary degree.

If the context is unambiguous, we will often identify the defining polynomial of an algebraic hypersurface with its zero set. For simplification we will speak of the hypersurface f when we mean the hypersurface defined by the polynomial f. We also say that the point $(a_1, \ldots, a_d)$ in d-dimensional real space lies *on* the hypersurface defined by $f \in \mathbb{Q}[x_1, \ldots, x_d]$ if $(a_1, \ldots, a_d)$ is a root of f.

For a hypersurface f the *gradient vector* of f is defined to be

$$\nabla f = \left(\frac{\partial f}{\partial x_1}, \frac{\partial f}{\partial x_2}, \ldots, \frac{\partial f}{\partial x_d} \right) \in (\mathbb{Q}[x_1, \ldots, x_d])^d.$$

With the help of the gradient vector we characterize a point $a = (a_1, \ldots, a_d) \in \mathbb{R}^d$ lying on the hypersurface f. It is named a *singular point* of f if $\nabla f(a) = 0$, otherwise it is *non-singular*. In our work, as we will see later, the main task is computing arrangements of curves in the plane, so consider $d = 2$. Informally speaking, the tangent line at a non-singular point (a, b) of a curve f is perpendicular to $\nabla(f)(a, b)$. The singular points of f are exactly the ones that do not admit a unique tangent line to the curve.

Let (a, b) be a non-singular point of f, i.e. there exists a well defined tangent line in that point. Under certain assumptions we do a further classification of (a, b):

1. We speak of (a, b) having a *vertical tangent* in the case $\nabla f(a, b) = c \cdot (1, 0)$ with c being a non-zero constant.
2. The point (a, b) is called a *turning point* of the curve if the tangent of f at (a, b) crosses f in (a, b). At a turning point the curvature of f changes sign. A necessary condition is that the polynomial

$$f_1(x, y) := (f_{xx} f_y^2 - 2 f_x f_y f_{xy} + f_{yy} f_x^2)(x, y) \in \mathbb{Q}[x, y]$$

   has a root at $(x, y) = (a, b)$: $f_1(a, b) = 0$.
3. If (a, b) is a turning point that additionally has a vertical tangent, then we call it a *vertical turning point*. In our work we consider curves of degree at most 4. A point (a, b) is a vertical turning point of a curve f of degree at most 4 iff

$$f_y(a, b) = 0 \quad \text{and} \quad f_{yy}(a, b) = 0 \quad \text{and} \quad f_{yyy}(a, b) \neq 0 \quad \text{and} \quad f_x(a, b) \neq 0.$$

4. We call (a, b) an *extreme point* if it has a vertical tangent but is not a turning point.

Assume that we want to sweep a planar curve with a sweep-line parallel to the y-axis from left to right. For an introduction to sweep-line algorithms see for example [23]. Then, informally speaking, extreme points are the ones where a new branch of the curve starts or ends. At a vertical turning point a vertical tangent coincides with a change of sign of the curvature. For example consider Figure 2.2. The green curve has three real points with a vertical tangent: two blue extreme points b_1, b_2 and one red vertical turning point r. Also the yellow point y is a turning point, but without a vertical tangent.

Next, we consider common roots of two polynomials $f, g \in \mathbb{Q}[x_1, \ldots, x_d]$. A point $a = (a_1, \ldots, a_d) \in \mathbb{R}^d$ is called an *intersection point* of two hypersurfaces f and g if

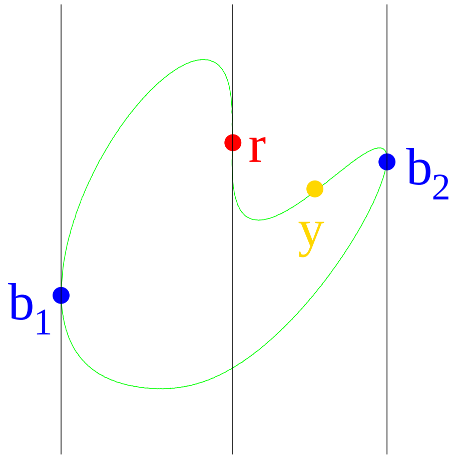

Figure 2.2: The curve has two blue extreme points b_1, b_2, one red vertical turning point r, and a yellow turning point y.

it lies on the hypersurface f as well as on the hypersurface g. It is called a *tangential intersection point* of f and g if additionally the two gradient vectors $\nabla f(a)$ and $\nabla g(a)$ are linearly dependent in a. Otherwise we speak of a *transversal intersection point*. If a is an intersection point of f and g and simultaneously a singular point of f, then of course $\nabla f(a) = 0$ and a is a tangential intersection point of f and g. We call an intersection point (a, b) *non-singular*, if (a, b) is neither a singular point of f nor of g.

The set of all intersection points of two surfaces p and q that do not share a common surface patch in 3-space is named *intersection curve*. A point (a, b, c) on the intersection curve is a tangential intersection point of p and q if and only if the two gradient vectors $(f_x, f_y, f_z)(a, b, c)$ and $(g_x, g_y, g_z)(a, b, c)$ are linearly dependent, which can be expressed algebraically as:

$$\begin{pmatrix} f_y g_z - f_z g_y \\ f_x g_z - f_z g_x \\ f_x g_y - f_y g_x \end{pmatrix} (a, b, c) \;=\; \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}.$$

For two planar curves f and g an intersection point (a, b) is tangential if and only if

$$(f_x g_y - f_y g_x)(a, b) = 0.$$

We are interested in real singular, extreme, and vertical turning points of one curve f and also in real intersection points of two curves f and g. But $\mathbb{R}$ is not algebraically closed and most of the time we have to work over its algebraic closure $\mathbb{C}$. Therefore we transfer all notations and definitions we made for real points also to points in complex d-dimensional space.

2.1.2 Arrangement of quadrics

Consider a set $P = \{p_1, \ldots, p_n\}$ of quadric surfaces. The quadrics form a 3-dimensional arrangement and partition the affine space in a natural way into four different types of maximal connected regions of dimensions 3, 2, 1, and 0, respectively. To state this more formally let $sign : \mathbb{R} \to \{-1, 0, 1\}$ be the function:

$$sign(a) = \begin{cases} -1, & \text{if } a < 0 \\ 0, & \text{if } a = 0 \\ +1, & \text{if } a > 0. \end{cases}$$

Then every point $(a, b, c) \in \mathbb{R}^3$ has a well defined sign sequence:

$$(sign(p_1(a,b,c)), sign(p_2(a,b,c)), \ldots, sign(p_n(a,b,c))) \in \{-1, 0, +1\}^n.$$

The sign sequence defines an equivalence relation on the set of points in space. We say that two points (a_1, b_1, c_1) and (a_2, b_2, c_2) are equivalent, if they have the same sign sequence. Of course the set of points of an equivalence class is not necessarily connected. But every connected subset has a unique sign sequence in $\{-1, 0, +1\}^d$. A maximal connected set of equivalent points is called *cell*, *face*, *edge*, or *vertex*, depending on its dimension:

1. A *cell* is of dimension 3 and on either side of each quadric. It is a connected subset of an equivalence class with no 0-entry in its sign sequence.

2. A *face* is of dimension 2 and has at least one 0-entry in its sign sequence. It lies on the surface of one quadratic.

3. An *edge* is a region of dimension 1 and has a sign sequence containing at least two 0-entries. It is part of the intersection curve of two quadrics.

4. A *vertex* is 0-dimensional and has a sign sequence with at least three 0-entries. It is the intersection point of three or more quadrics or rather the intersection point of at least two intersection curves.

When dealing with arrangements of quadric surfaces there is one basic operation that has to be at our disposal, independent on the kind of information we are interested in: We have to locate intersection points of two or more intersection curves unambiguously along the curves. Algebraically speaking, for two trivariate rational polynomials defining an intersection curve we want to compute common real roots with other polynomials. Our approach for providing this operation is based on a *projection step*, as it also occurs in the cylindrical algebraic decomposition [19]. We do not deal with the intersection curves of two quadrics directly in space. Instead we project them into the plane. The way the projection is computed will be explained in the next section about resultants. For the moment we only want to clarify some notation.

From the point of view of the (x, y)-plane a quadric p consists of three different parts: the *lower part*, the *silhouette*, and the *upper part*.

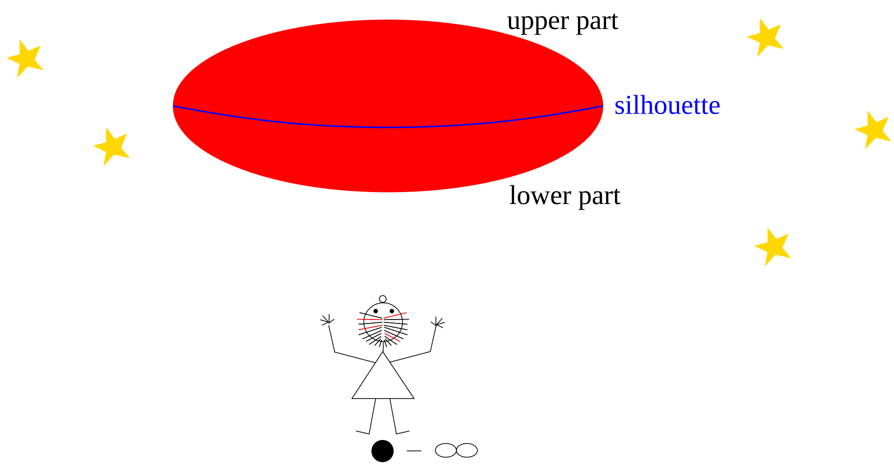

Figure 2.3: A girl observing an ellipsoid in the night

For illustration imagine the following: assume an observer stands at the point $z = -\infty$ and looks upwards to a quadric in the sky directly above him, like the girl in Figure 2.3. Then the *lower part* of the quadric consists of all points that can be seen from the observer, but we exclude the border. The border is called the *silhouette*. The set of points that cannot be seen from the observer is called the *upper part*. More mathematically, we define the three parts of a quadric p the following way:

1. The *lower (upper) part* of the quadric p consists of all points $(a, b, c) \in \mathbb{R}^3$ such that
   (a) $p(a, b, z) \in \mathbb{R}[z]$ has two different real roots and
   (b) c is the smaller (bigger) root.
2. The *silhouette* of the quadric p consists of all points $(a, b, c) \in \mathbb{R}^3$ such that

(a) $p(a, b, z) \in \mathbb{R}[z]$ has one root of multiplicity 2 and

(b) c is this root.

Now, after we have introduced the most important terms, we will answer the question how to project the intersection curve of two quadrics into the plane. Besides the intersection curves we will also project the silhouette of each quadric. As we will see, both projections can be algebraically realized by the same tool: resultants.

2.2 Resultants

The intersection curve of two quadrics is given by the set of common roots of their defining polynomials. So we are interested in common roots of two polynomials.

2.2.1 Resultants of univariate polynomials

Let us first consider univariate polynomials. Let k denote an arbitrary field. Most of the time we work with polynomials over the rational numbers or their algebraic closure. So imagine $k = \mathbb{Q}$ or $k = \mathbb{C}$ in the following. We will denote the *degree* of a polynomial $f \in k[x]$ by $\deg(f)$ and its *leading coefficient* by $\text{ldcf}(f)$. If f is the zero polynomial, we write $f \equiv 0$. The degree of the zero polynomial is equal to $-\infty$: $\deg(0) = -\infty$. We call a polynomial $h \in k[x]$ a *factor* of f if there exists another polynomial $f_1 \in k[x]$ with $f = h \cdot f_1$. In particular, if f is the zero polynomial, then every $h \in k[x]$ is a factor of f.

For every two univariate polynomials f and g each constant number $c \in k$ is a common factor of f and g. But when do two univariate polynomials f and g of positive degree have a non-constant common factor?

Theorem 2.1: Let $f, g \in k[x]$ be two polynomials of degrees $\deg(f) = n > 0$ and $\deg(g) = m > 0$. Then f and g have a non-constant common factor if and only if there are polynomials $A \in k[x]$ and $B \in k[x]$ with $\deg(A) < m$ and $\deg(B) < n$ which are not both zero such that $Af + Bg = 0$.

Proof. First assume that f and g have a non-constant common factor h. Then we can write $f = hf_1$ and $g = hg_1$, where $f_1, g_1 \in k[x]$. It is easy to check that $A = g_1$ and $B = -f_1$ fulfill the required properties. Conversely assume that there are polynomials A and B of degree at most $n-1$ and $m-1$ which are not both zero such that $Af = -Bg$. Then all irreducible factors of $A(x)f(x)$ divide $B(x)g(x)$. But B has degree at most $n - 1$ and therefore cannot contain all factors of f with their multiplicities. $\square$

Now in order to determine the existence of a common factor of

$$\begin{aligned} f(x) &= f_n x^n + f_{n-1}x^{n-1} + \cdots + f_0, \quad f_n \neq 0,\ n > 0 \quad \text{and} \\ g(x) &= g_m x^m + g_{m-1}x^{m-1} + \cdots + g_0, \quad g_m \neq 0,\ m > 0 \end{aligned}$$

we have to decide whether two polynomials A and B with the required properties can be found. This questions can be answered with the help of linear algebra: A and B are polynomials of degree at most $m-1$ and $n-1$, and therefore there are all in all $m+n$ unknown coefficients $a_{m-1}, \ldots, a_0, b_{n-1}, \ldots, b_0$ of A and B:

$$\begin{aligned} A(x) &= a_{m-1}x^{m-1} + \cdots + a_1 x + a_0 \\ B(x) &= b_{n-1}x^{n-1} + \cdots + b_1 x + b_0. \end{aligned}$$

The polynomial $A(x)f(x) + B(x)g(x)$ has degree at most $n+m-1$ in x. Each of its coefficients has to be zero in order to achieve $Af + Bg = 0$:

$$\begin{array}{ccccl} f_n a_{m-1} & + & g_m b_{n-1} & = 0 & \text{(coefficient of } x^{m+n-1}) \\ f_{n-1}a_{m-1} + f_n a_{m-2} & + & g_{m-1}b_{n-1} + g_m b_{n-2} & = 0 & \text{(coefficient of } x^{m+n-2}) \\ \ddots & & \ddots & \vdots & \\ f_0 a_0 & + & g_0 b_0 & = 0 & \text{(coefficient of } x^0). \end{array}$$

We get $n+m$ linear equations in the unknowns a_i, b_i with coefficients in f_i, g_i. Written in matrix style this system of linear equations has the form:

$$(a_{m-1}, \ldots, a_0, b_{n-1}, \ldots, b_0) \cdot \begin{pmatrix} f_n & f_{n-1} & \cdots & f_0 & & & \\ & f_n & f_{n-1} & \cdots & f_0 & & \\ & & \ddots & \ddots & & \ddots & \\ & & & f_n & f_{n-1} & \cdots & f_0 \\ g_m & g_{m-1} & \cdots & g_0 & & & \\ & g_m & g_{m-1} & \cdots & g_0 & & \\ & & \ddots & \ddots & & \ddots & \\ & & & g_m & g_{m-1} & \cdots & g_0 \end{pmatrix} = (0, \ldots, 0)$$

where the empty places are filled with zeros. We know from linear algebra that this system of linear equations has a non-zero solution if and only if the determinant of the coefficient matrix is equal to zero. This leads to the following definition:

Definition 2.2: The $(m+n) \times (m+n)$ coefficient matrix with m rows of f-entries and n rows of g-entries

$$\mathrm{Syl}(f,g) := \begin{pmatrix} f_n & f_{n-1} & \cdots & f_0 & & & \\ & f_n & f_{n-1} & \cdots & f_0 & & \\ & & \ddots & \ddots & & \ddots & \\ & & & f_n & f_{n-1} & \cdots & f_0 \\ g_m & g_{m-1} & \cdots & g_0 & & & \\ & g_m & g_{m-1} & \cdots & g_0 & & \\ & & \ddots & \ddots & & \ddots & \\ & & & g_m & g_{m-1} & \cdots & g_0 \end{pmatrix}$$

where the empty spaces are filled by zeros is called the *Sylvester matrix* of f and g. The determinant of the matrix is called the *resultant* of f and g: $\mathrm{res}(f,g) := \det(\mathrm{Syl}(f,g))$.

From the above observations we immediately obtain a criterion for testing whether two polynomials f and g have a non-constant common factor.

Proposition 2.3: Given $f, g \in k[x]$ of positive degree, the resultant $\mathrm{res}(f,g) \in k$ is equal to zero if and only if f and g have a non-constant common factor. For $k = \mathbb{C}$ the equality $\mathrm{res}(f,g) = 0$ holds if and only if f and g have a common complex root.

So far we have only considered polynomials f and g of positive degree. What about exactly one of them having degree 0 or $-\infty$? In this case we can analogously derive the resultant of f and g. If f has degree $n > 0$ and $g = g_0$ is constant (possibly 0), then the Sylvester matrix of f and g is the $n \times n$ matrix with g_0 on the main diagonal and 0's elsewhere. From that we conclude $\mathrm{res}(f, g_0, x) = g_0^n$. If $f = f_0$ is constant and g of positive degree m then $\mathrm{res}(f_0, g, x) = f_0^m$. The case that both polynomials are constant is of no interest for the following considerations.

2.2.2 Properties of resultants

There are some useful properties of resultants. For an arbitrary field k let $\bar{k}$ denote its algebraic closure.

Lemma 2.4: Let $f, g \in k[x]$ and $\alpha \in \bar{k}$.

1. For $\deg(f) > 0$ and $\deg(g) = m > 0$ we have $\mathrm{res}(\alpha \cdot f, g) = \alpha^m \cdot \mathrm{res}(f,g)$.

2. If $\deg(g) > 0$, then $\mathrm{res}((x-\alpha) \cdot f, g) \;=\; g(\alpha) \cdot \mathrm{res}(f,g)$.

Proof. 1. With $\deg(f) = n$ and $\deg(g) = m$ we have

$$\alpha \cdot f(x) = \alpha f_n x^n + \alpha f_{n-1} x^{n-1} + \cdots + \alpha f_0.$$

It is clear from the definition of the resultant of f and g that one can extract from each of the first m rows of the Sylvester matrix a factor α, leading to the stated result.

2. Again let $f(x) = \sum_{i=0}^{n} f_i x^i$ and $g(x) = \sum_{i=0}^{m} g_i x^i$. Then the coefficients of $(x - \alpha)f$ in order of decreasing powers of x are

$$(f_n\ ,\ f_{n-1} - \alpha f_n\ ,\ f_{n-2} - \alpha f_{n-1}\ ,\ \ldots\ ,\ f_0 - \alpha f_1\ ,\ -\alpha f_0).$$

We do some transformations with the Sylvester matrix of $(x - \alpha)f$ and g. First we add α times column i to column $(i+1)$ for all $i = 1, \ldots, m+n$. This puts the rows with the f-entries into the right form:

$$\begin{matrix} f_n & f_{n-1} & \cdots & f_0 & 0 & & & & \\ & f_n & f_{n-1} & \cdots & f_0 & 0 & & & \\ & & \ddots & \ddots & & & \ddots & & \\ & & & f_n & f_{n-1} & \cdots & f_0 & 0 & \end{matrix}.$$

Each non-zero entry of the last $(n+1)$ rows is now equal to $g(\alpha)/\alpha^i$ for a suitable i. By $g(\alpha)/\alpha^i$ we mean the substitution of α for x into the integral part of $g(x)$ divided by x^i. For example, the former entry b_i is replaced by $g(\alpha)/\alpha^i = g_m \alpha^{m-i} + \cdots + g_{i+1}\alpha + g_i$, for all $1 \le i \le m$. To simplify the matrix we subtract α times row j from row $j-1$ for $j = m+2, \ldots, n+m+1$. This leads to the following matrix, which has the same determinant as the Sylvester matrix of $(x - \alpha)f$ and g:

$$\begin{pmatrix} f_n & f_{n-1} & \cdots & f_0 & 0 & & & & \\ & f_n & f_{n-1} & \cdots & f_0 & 0 & & & \\ & & \ddots & \ddots & & \ddots & & & \\ & & & f_n & f_{n-1} & \cdots & f_0 & 0 \\ g_m & g_{m-1} & \cdots & g_0 & & & & \\ & g_m & g_{m-1} & \cdots & g_0 & & & \\ & & \ddots & \ddots & & \ddots & & \\ & & & g_m & g_{m-1} & \cdots & g_0 & 0 \\ & & & & g(\alpha)/\alpha^m & \cdots & g(\alpha)/\alpha^1 & g(\alpha) \end{pmatrix}.$$

The last column has only one non-zero entry $g(\alpha)$ and the cofactor of this entry is equal to the resultant of f and g. This proves our claim.

□

The Lemma leads to the following important characterization of resultants:

Theorem 2.5: Let $f, g \in k[x]$, $f_n = \mathrm{ldcf}(f)$, $g_m = \mathrm{ldcf}(g)$, $\deg(f) = n > 0$, $\deg(g) = m > 0$, with roots

$$\alpha_1, \ldots, \alpha_n, \beta_1 \ldots, \beta_m \in \bar{k}.$$

For the resultant of f and g the following holds:

$$\mathrm{res}(f,g) \;=\; f_n^m g_m^n \prod_{i=1}^{n} \prod_{j=1}^{m} (\alpha_i - \beta_j).$$

Proof. Writing $f(x) = f_n \cdot \prod_{i=0}^{n}(x - \alpha_i)$ and $g(x) = g_m \cdot \prod_{j=0}^{m}(x - \beta_j)$ we get from Lemma 2.4

$$\begin{aligned}
\mathrm{res}(f,g) &= f_n^m \mathrm{res}\left(\prod_{i=1}^{n}(x-\alpha_i), g\right) \\
&= f_n^m g(\alpha_1) \mathrm{res}\left(\prod_{i=2}^{n}(x-\alpha_i), g\right) \\
&= \ldots \\
&= f_n^m g(\alpha_1) \ldots g(\alpha_n) \\
&= f_n^m g_m^n \prod_{i=1}^{n} \prod_{j=1}^{m} (\alpha_i - \beta_j).
\end{aligned}$$

□

A direct conclusion is that the resultant is multiplicative:

Corollary 2.6: Let $f, g, h \in k[x]$ be polynomials with positive degree. Then the equalities $\mathrm{res}(f \cdot h, g) = \mathrm{res}(f,g) \cdot \mathrm{res}(h,g)$ and $\mathrm{res}(f, g \cdot h) = \mathrm{res}(f,g) \cdot \mathrm{res}(f,h)$ hold.

Another immediate surprising consequence is that the resultant of two polynomials is invariant under translation:

Corollary 2.7: Let $f, g \in k[x]$ with $\deg(f) > 0$ and $\deg(g) > 0$. For a number $c \in \bar{k}$ we define $\tilde{f}(x) := f(x+c)$ and $\tilde{g}(x) := g(x+c)$. Then $\mathrm{res}(f,g) = \mathrm{res}(\tilde{f}, \tilde{g}) \in k$.

2.2.3 Computation of resultants

The computation of resultants can be performed more efficiently than using the obvious determinant computation. This whole section is taken from [77]. Let $f(x) = q(x)g(x) + r(x)$ where $\deg(f) = n$, $\deg(g) = m$, $\deg(r) = l$, and $n \geq m > l$. The basic idea is the following: One can show that

$$
\begin{aligned}
(\star) \qquad \operatorname{res}(f,g) &= (-1)^{m(n-l)} g_m^{n-l} \operatorname{res}(r,g) \qquad (g_m = \operatorname{ldcf}(g)) \\
&= (-1)^{nm} g_m^{n-l} \operatorname{res}(g,r).
\end{aligned}
$$

Thus, the resultant of f and g can be expressed in the terms of the resultant of g and r. Since r is the remainder of f divided by g, we can apply a Euclidean–like algorithm: Given f and g, we construct the Euclidean remainder sequence

$$s_0 = f \ , \ s_1 = g \ , \ \ldots \ , \ s_h$$

where $s_{i+1} = (s_{i-1} \mod s_i)$ and $s_{h+1} = 0$. If $\deg(s_h)$ is non-constant, then $\operatorname{res}(f,g) = 0$. Otherwise, we can repeatedly apply the formula $(\star)$ above until the basic case given by $\operatorname{res}(s_{h-1}, s_h) = s_h^{\deg(s_{h-1})}$. This computation can be sped up further as shown in [66].

So far we have seen that resultants are related to the question whether two polynomials f and g share a common factor. Sometimes this information is not enough and we are also interested in the degree of their greatest common divisor. For this we take a glance at subresultant theory, which is a generalization of the theory of resultants.

2.2.4 Subresultants of univariate polynomials

The underlying Theorem 2.1 of the previous section can be generalized to

Theorem 2.8: Let $f, g \in k[x]$ be two polynomials of degrees $\deg(f) = n > 0$ and $\deg(g) = m > 0$. Then f and g have a common factor of degree greater than $l \geq 0$ if and only if there are rational polynomials A and B with $\deg(A) < m - l$ and $\deg(B) < n - l$ which are not both zero such that $Af + Bg = 0$.

Proof. Theorem 2.1 treats the case $l = 0$. For every other $l > 0$ the proof works the same way. □

As an immediate consequence we obtain a statement about the degree of the gcd of f and g:

Corollary 2.9: The degree of the gcd of two polynomials $f, g \in k[x]$ is equal to the smallest index h such that for all rational polynomials A and B with $\deg(A) < m - h$ and $\deg(B) < n - h$: $Af + Bg \neq 0$.

This Corollary can be reformulated the following way:

Corollary 2.10: The degree of the gcd of two polynomials $f, g \in k[x]$ is equal to the smallest index h such that for all rational polynomials A and B with $\deg(A) < m - h$ and $\deg(B) < n - h$: $\deg(Af + Bg) \geq h$.

Proof. Corollary 2.10 follows from Corollary 2.9, because $\gcd(f, g)$ is a generator of the ideal $\langle f, g\rangle$ defined by f and g. For every two rational polynomials A and B the polynomial $Af + Bg$ is a member of $\langle f, g\rangle$. That means either $Af + Bg = 0$ or $\deg(Af + Bg) \geq \deg(\gcd(f, g)) = h$. □

We are interested in determining the degree of the greatest common divisor of two polynomials f and g. According to Corollary 2.10 we have to test in succession whether for $l = 1, 2, 3, \ldots$ there exist polynomials A and B, with the claimed restriction of the degrees such that the degree of $Af + Bg$ is strictly smaller than l. The first index h, for which this test gives a negative answer, is equal to the degree of the gcd. How can we perform such a test? We have seen in the previous section that the test for $l = 0$ can be made by testing whether the resultant of f and g is equal to zero. For $l = 1, 2, 3, \ldots$ we proceed in a similar way. Let l be a fixed index and let

$$\begin{aligned} f(x) &= f_n x^n + f_{n-1}x^{n-1} + \cdots + f_0, \quad f_n \neq 0, \quad \text{and} \\ g(x) &= g_m x^m + g_{m-1}x^{m-1} + \cdots + g_0, \quad g_m \neq 0. \end{aligned}$$

We are looking for two polynomials

$$\begin{aligned} A(x) &= a_{m-l-1}x^{m-l-1} + \cdots + a_1 x + a_0 \\ B(x) &= b_{n-l-1}x^{n-l-1} + \cdots + b_1 x + b_0, \end{aligned}$$

such that $\deg(Af + Bg) < l$. There are $m + n - 2l$ unknown coefficients $a_{m-l-1}, \ldots, a_0, b_{n-l-1}, \ldots, b_0$. The polynomial $A(x)f(x) + B(x)g(x)$ has degree at most $n + m - l - 1$. The $m + n - 2l$ coefficients of $x^l, x^{l+1}, \ldots, x^{m+n-l-1}$ have to be zero in order to achieve $\deg(Af + Bg) < l$. This leads to a linear system

$$(a_{m-l-1}, \ldots, a_0, b_{n-l-1}, \ldots, b_0) \cdot S_l = (0, \ldots, 0)$$

where S_l is the submatrix of the Sylvester matrix of f and g obtained by deleting the last $2l$ columns, the last l rows of f-entries, and the last l rows of g-entries. We call $\mathrm{sres}_l(f,g) = \det S_l$ the lth *subresultant* of f and g. For $l = 0$ the equality $\mathrm{res}(f,g) = \mathrm{sres}_0(f,g)$ holds. In fact, S_l is a submatrix of S_i for $l > i \geq 0$. So here is our main Proposition:

Proposition 2.11: Two polynomials f and g of positive degree have a gcd of degree h if and only if h is the least index l for which $\mathrm{sres}_l(f,g) \neq 0$. If we work over the complex numbers, then f and g have exactly h common roots if and only if h is the least index l for which $\mathrm{sres}_l(f,g) \neq 0$.

Also for subresultants it is true that they are invariant under translation. A proof can be found in [38]. There the more general case of arbitrary composition is discussed.

Proposition 2.12: Let $f, g \in k[x]$ with $\deg(f) > 0$ and $\deg(g) > 0$. For a number $c \in \bar{k}$ in the algebraic closure of k we define $\tilde{f}(x) := f(x+c)$ and $\tilde{g}(x) := g(x+c)$. For $0 \leq l \leq \deg(f,g)$ we have

$$\mathrm{sres}_l(f,g) = \mathrm{sres}_l(\tilde{f},\tilde{g}).$$

It is an open question whether a direct proof of this Proposition exists, reasoning with roots as in Theorem 2.5 and Corollary 2.7 for resultants. Unlike the resultant case, no nice expression of subresultants in terms of roots is known.

Our next task is to adapt the theory of resultants and subresultants to the case of multivariate polynomials $f, g \in k[x_1, x_2, \ldots, x_d]$ of positive degree in x_d. For our computations we are especially interested in the cases $d = 3$ and $d = 2$.

2.2.5 Resultants of multivariate polynomials

Suppose we are given $f, g \in \mathbb{Q}[x_1, \ldots, x_d]$ with positive degree in x_d. We write

$$\begin{aligned} f &= f_n x_d^n + \cdots + f_0 x_d^0, \quad f_n \not\equiv 0 \\ g &= g_m x_d^m + \cdots + g_0 x_d^0, \quad g_n \not\equiv 0. \end{aligned}$$

Thus we regard f and g as polynomials in x_d with coefficients f_i and g_j that are polynomials in $k[x_1, \ldots, x_{d-1}]$. We define the *resultant* of f and g with respect to x_d similarly to the one in Definition 2.2:

$$\operatorname{res}(f,g,x_d) := \det\begin{pmatrix} f_n & f_{n-1} & \cdots & f_0 & & & \\ & f_n & f_{n-1} & \cdots & f_0 & & \\ & & \ddots & \ddots & & \ddots & \\ & & & f_n & f_{n-1} & \cdots & f_0 \\ g_m & g_{m-1} & \cdots & g_0 & & & \\ & g_m & g_{m-1} & \cdots & g_0 & & \\ & & \ddots & \ddots & & \ddots & \\ & & & g_m & g_{m-1} & \cdots & g_0 \end{pmatrix} \in k[x_1,\ldots,x_{d-1}].$$

For the resultant of two complex univariate polynomials f and g we have seen that its vanishing is a necessary and sufficient condition for the two polynomials to have a common root. What can we say about the resultant of multivariate polynomials?

Proposition 2.13: Given $f, g \in \mathbb{C}[x_1,\ldots,x_d]$ regarded as polynomials in x_d with coefficients in $\mathbb{C}[x_1,\ldots,x_{d-1}]$ and leading coefficients f_n and g_m not equal to the zero polynomial. If $(c_1,\ldots,c_d) \in \mathbb{C}^d$ is a common root of the multivariate polynomials f and g, then the equality $\operatorname{res}(f,g,x_d)(c_1,\ldots,c_{d-1}) = 0$ holds.

Proof. Let $c := (c_1,\ldots,c_{d-1})$ and let us denote $f(c_1,\ldots,c_{d-1},x_d) =: f(c,x_d)$. In the case that neither f_n nor g_m vanishes at c, the determinant

$$\operatorname{res}(f,g,x_d)(c) = \det\begin{pmatrix} f_n(c) & \cdots & f_0(c) & & \\ & \ddots & \ddots & \ddots & \\ & & f_n(c) & \cdots & f_0(c) \\ g_m(c) & \cdots & g_0(c) & & \\ & \ddots & \ddots & \ddots & \\ & & g_m(c) & \cdots & g_0(c) \end{pmatrix}$$

is equal to the resultant of the univariate polynomials $f(c,x_d)$ and $g(c,x_d)$. By assumption c_d is a common root of $f(c,x_d)$ and $g(c,x_d)$ and therefore by Proposition 2.3 we have $\operatorname{res}(f,g,x_d)(c) = 0$. In the case $f_n(c) = g_m(c) = 0$ it is easy to verify that $\operatorname{res}(f,g,x_d)(c) = 0$. The last case we have to consider is that either f_n or g_m vanishes at c, without loss of generality g_m. Then, up to a multiple of the factor $f_n(c)$, $\operatorname{res}(f,g,x_d)(c)$ equals $\operatorname{res}(f(c,x_d), g(c,x_d), x_d)$ and by assumption the latter is equal to zero. □

So the vanishing of the resultant of f and g with respect to x_d at a point $(c_1,\ldots,c_{d-1}) \in \mathbb{C}^{d-1}$ is a necessary condition for the extendibility to a common solution $(c_1,\ldots,c_{d-1},c_d) \in \mathbb{C}^d$ of f and g. Does this criterion is also sufficient, like in the case for univariate polynomials? Does each root $(c_1, \ldots, c_{d-1}) \in \mathbb{C}^{d-1}$ of $\mathrm{res}(f, g, x_d)$ extend to a common root $(c_1, \ldots, c_{d-1}, c_d) \in \mathbb{C}^d$ of f and g? Not in all cases:

Proposition 2.14: Given $f, g \in \mathbb{C}[x_1, \ldots, x_d]$ regarded as polynomials in x_d with coefficients in $\mathbb{C}[x_1, \ldots, x_{d-1}]$ and leading coefficients f_n and g_m not equal to the zero polynomial. If $\mathrm{res}(f, g, x_d)$ vanishes at $(c_1, \ldots, c_{d-1}) \in \mathbb{C}^{d-1}$, then either

1. f_n or g_m vanishes at $(c_1, \ldots, c_{d-1})$, or
2. there is a number $c_d \in \mathbb{C}$ such that f and g vanish at $(c_1, \ldots, c_d) \in \mathbb{C}^d$.

Proof. Let $c := (c_1, \ldots, c_{d-1})$ and let us denote $f(c_1, \ldots, c_{d-1}, x_d) =: f(c, x_d)$. It suffices to show that $f(c, x_d)$ and $g(c, x_d)$ have a common root when $f_n(c)$ and $g_m(c)$ are both non-zero. To prove this look at

$$\begin{aligned} f(c, x_d) &= f_n(c)x_d^n + \cdots + f_1(c)x_d + f_0(c) \\ g(c, x_d) &= g_m(c)x_d^m + \cdots + g_1(c)x_d + g_0(c). \end{aligned}$$

By hypothesis we know that the resultant $\mathrm{res}(f, g, x_d)$ of the multivariate polynomials f and g vanishes at c:

$$0 = \mathrm{res}(f, g, x_d)(c) = \det \begin{pmatrix} f_n(c) & \ldots & f_0(c) & & \\ & \ddots & \ddots & \ddots & \\ & & f_n(c) & \ldots & f_0(c) \\ g_m(c) & \ldots & g_0(c) & & \\ & \ddots & \ddots & \ddots & \\ & & g_m(c) & \ldots & g_0(c) \end{pmatrix}.$$

By assumption we have $f_n(c) \neq 0 \neq g_m(c)$ and because of that the resultant of the univariate polynomials $f(c, x_d)$ and $g(c, x_d)$ is exactly the same determinant. It follows $\mathrm{res}(f(c, x_d), g(c, x_d)) = 0$ and Proposition 2.3 implies that f and g have a common root $(c_1, \ldots, c_{d-1}, c_d)$ for some $c_d \in \mathbb{C}$. □

We easily derive the following statement:

Corollary 2.15: Let $f, g \in \mathbb{C}[x_1, \ldots, x_d]$ be two polynomials with constant leading coefficients when regarded as polynomials in x_d. Then a point $(c_1, \ldots, c_{d-1}) \in \mathbb{C}^{d-1}$ is extendible to a common solution $(c_1, \ldots, c_{d-1}, c_d) \in \mathbb{C}^d$ of f and g if and only if $(c_1, \ldots, c_{d-1})$ is a root of the resultant $\mathrm{res}(f, g, x_d)$.

In our application we are interested in the intersection curve of two quadrics. We just derived the geometric result that the resultant of two multivariate polynomials f and

g can be interpreted as the projection of their intersection curve. The set of roots of $\text{res}(f, g, x_d)$ corresponds, with some exceptions, to the set of $(x_1, \ldots, x_{d-1})$-coordinates of common complex roots of f and g. This is a very useful result but the drawback is that it is a statement over the *complex* numbers and we are only interested in common *real* roots: The intersection curve of two quadrics p and q is defined as the set of common real roots. It can happen that a real root of $\text{res}(p, q, z)$ derives from a complex common root of p and q. For example let

$$\begin{aligned} p(x, y, z) &= (x - 1/4)^2 + (y - 1/4)^2 + z^2 - 1 \\ q(x, y, z) &= (x + 1/4)^2 + (y + 1/4)^2 + z^2 - 1. \end{aligned}$$

Then $\text{res}(p, q, z) = (x+y)^2$ and the point $(1, -1)$ in the real plane is a root of the resultant. But p and q have the common complex roots $(1, -1, i\sqrt{9/8})$ and $(1, -1, i\sqrt{9/8})$.

The polynomial $\text{res}(f, g, x_d)$ is a polynomial in $k[x_1, \ldots, x_{d-1}]$ and a natural question is to bound its degree. Let $\deg(f)$ denote the *total degree* of a multivariate polynomial f. The total degree of $\text{res}(f, g, x_d)$ can be bounded from above in terms of the total degrees of f and g.

Proposition 2.16: Assume that f and g are polynomials with $\deg(f) = n$ and $\deg(g) = m$. Then $\text{res}(f, g, x_d)$ is a polynomial of total degree at most $n \cdot m$.

Proof. We first homogenize the polynomials f and g by introducing a new variable z. Then each monomial of the resulting homogeneous polynomials F and G has degree n or m, respectively:

$$\begin{aligned} F &= F_n x_d^n + F_{n-1} x_d^{n-1} + \cdots + F_1 x_d + F_0 \\ G &= G_m x_d^m + G_{m-1} x_d^{m-1} + \cdots + G_1 x_d + G_0. \end{aligned}$$

Each F_i is a homogeneous polynomial in $k[x_1, \ldots, x_{d-1}, z]$ of degree $n - i$ and each G_j is homogeneous of degree $m - j$.

We will show that $\text{res}(F, G, x_d) \in k[x_1, \ldots, x_{d-1}, z]$ either is the zero polynomial or it is homogeneous of degree mn, see also [73]. Then we can conclude, by substituting $z = 1$, that $\text{res}(f, g, x_d) \not\equiv 0$ is a polynomial of total degree at most $n \cdot m$.

It is easy to see that a non-zero polynomial $H(x_1, \ldots, x_d)$ is homogeneous of degree n if and only if $H(tx_1, \ldots, tx_d) = t^n H(x_1, \ldots, x_d)$ for every $t \in k$

Now look at the resultant of the two homogeneous polynomials $F, G \in k[x_1, \ldots, x_d, z]$:

$$
\begin{aligned}
&\operatorname{res}(F,G,x_d)(tx_1,\dots,tx_{d-1},tz) \\
&= \det\begin{pmatrix}
F_n & tF_{n-1} & \dots & t^nF_0 & & & \\
 & \ddots & \ddots & & \ddots & & \\
 & & F_n & tF_{n-1} & \dots & t^nF_0 \\
G_m & tG_{m-1} & \dots & t^mG_0 & & \\
 & \ddots & \ddots & & \ddots & \\
 & & G_m & tG_{m-1} & \dots & t^mG_0
\end{pmatrix}.
\end{aligned}
$$

Multiply the ith row of F's by t^i and the jth row of G's by t^j. We obtain

$$
\begin{aligned}
&t^p\cdot\operatorname{res}(F,G,x_d)(tx_1,\dots,tx_{d-1},tz) \\
&= \det\begin{pmatrix}
tF_n & t^2F_{n-1} & \dots & t^{n+1}F_0 & & & \\
 & t^2F_n & t^3F_{n-1} & \dots & t^{n+2}F_0 & & \\
 & & \ddots & \ddots & & \ddots & \\
 & & & t^mF_n & t^{m+1}F_{n-1} & \dots & t^{n+m}F_0 \\
tG_m & t^2G_{m-1} & \dots & t^{m+1}G_0 & & & \\
 & t^2G_m & t^3G_{m-1} & \dots & t^{m+2}G_0 & & \\
 & & \ddots & \ddots & & \ddots & \\
 & & & t^nG_m & t^{n+1}G_{m-1} & \dots & t^{n+m}G_0
\end{pmatrix} \\
&= t^q\cdot\operatorname{res}(F,G,x_d)(x_1,\dots,x_{d-1},z)
\end{aligned}
$$

where $p = m(m+1)/2 + n(n+1)/2$, and $q = (m+n)(m+n+1)/2$. Hence

$$\operatorname{res}(F,G,x_d)(tx_1,\dots,tx_{d-1},tz) \;=\; t^{nm}\operatorname{res}(F,G,x_d)(x_1,\dots,x_{d-1},z).$$

□

The last proposition bounds the degree of the resultant of two polynomials from above. Under what circumstances does the resultant of f and g become the zero polynomial? What does the vanishing of the resultant at any point $(c_1,\dots,c_{d-1}) \in k^{d-1}$ mean geometrically?

Proposition 2.17: Let $f, g \in k[x_1,\dots,x_d]$ be two polynomials with positive degrees in x_d. The resultant $\operatorname{res}(f,g,x_d)$ is equal to the zero polynomial if and only if f and g have a common factor in $k[x_1,\dots,x_d]$ that has positive degree in x_d, i.e. if and only if there exist three polynomials $h, f_1, g_1 \in k[x_1,\dots,x_d]$, $\deg(h) > 1$, with

$$f = h\cdot f_1 \quad\text{and}\quad g = h\cdot g_1.$$

Proof. In the last sections about resultants of univariate polynomials we worked with polynomials in one variable with coefficients from an arbitrary field. Since f and g are polynomials in x_d with coefficients in $k[x_1, \ldots, x_{d-1}]$, the field these coefficients lie in is the field extension $k(x_1, \ldots, x_{d-1})$ of k. Then Proposition 2.3, applied to $f, g \in k(x_1, \ldots, x_{d-1})[x_d]$, tells us that $\operatorname{res}(f, g, x_d) = 0$ if and only if f and g have a common factor in $k(x_1, \ldots, x_{d-1})[x_d]$ with positive degree in x_d. What remains to prove is that this is equivalent to having a common factor in $k[x_1, \ldots, x_d]$ of positive degree in x_d. This will be shown in the next proposition. □

Proposition 2.18: Suppose that $f, g \in k[x_1, \ldots, x_d]$ have positive degree in x_d. Then f and g have a common factor in $k[x_1, \ldots, x_d]$ of positive degree in x_d if and only if they have a common factor in $k(x_1, \ldots, x_{d-1})[x_d]$ of positive degree in x_d.

Proof. One direction is easy: If f and g have a common factor in $k[x_1, \ldots, x_d]$ of positive degree in x_d, then of course they have a common factor in the larger ring $k(x_1, \ldots, x_{d-1})[x_d]$. For the other direction suppose that f and g have a common factor $h \in k(x_1, \ldots, x_{d-1})[x_d]$. Then there exist $f_1, g_1 \in k(x_1, \ldots, x_{d-1})[x_d]$ with

$$f = h \cdot f_1 \quad \text{and} \quad g = h \cdot g_1.$$

But the h, f_1, and g_1 may have denominators in the field $k(x_1, \ldots, x_{d-1})$. Let $d \in k[x_1, \ldots, x_{d-1}]$ be a common denominator of them and look at the polynomials $\tilde{h} = dh$, $\tilde{f}_1 = df_1$, and $\tilde{g}_1 = dg_1$ in $k[x_1, \ldots, x_d]$. If we multiply the above equations by d^2 we obtain

$$d^2 f = \tilde{h} \cdot \tilde{f}_1 \in k[x_1, \ldots, x_d] \quad \text{and} \quad d^2 g = \tilde{h} \cdot \tilde{g}_1 \in k[x_1, \ldots, x_d].$$

Now let $\tilde{h}_1$ be an irreducible factor of $\tilde{h}$ of positive degree in x_d. Irreducible means that $\tilde{h}_1$ is not the product of two non-constant polynomials in $k[x_1, \ldots, x_d]$. Since $h = \tilde{h}/d$ has positive degree in x_d, such an $\tilde{h}_1$ must exist. Then $\tilde{h}_1$ divides $d^2 f$. Since $\tilde{h}_1$ is irreducible, one can proof that it either divides d^2 or f. The former is impossible because $d^2 \in k[x_1, \ldots, x_{d-1}]$ and that is why $\tilde{h}_1$ must divide f in $k[x_1, \ldots, x_d]$. A similar argument shows that $\tilde{h}_1$ divides g and thus $\tilde{h}_1$ is the required common factor. □

We have shown in Corollary 2.6 that the resultant of two univariate polynomials is multiplicative. This remains true for multivariate polynomials:

Corollary 2.19: Let $f, g, h \in k[x_1, \ldots, x_d]$ be polynomials with positive degree when regarded as polynomials in x_d. Then the equalities $\operatorname{res}(f \cdot h, g, x_d) = \operatorname{res}(f, g, x_d) \cdot \operatorname{res}(h, g, x_d)$ and $\operatorname{res}(f, g \cdot h, x_d) = \operatorname{res}(f, g, x_d) \cdot \operatorname{res}(f, h, x_d)$ hold.

Proof. All polynomials $f, g, h \in k[x_1, \ldots, x_d]$ can be interpreted as univariate polynomials in the larger ring $k(x_1, \ldots, x_{d-1})[x_d]$. From Corollary 2.6 we know that in this larger ring the equalities $\operatorname{res}(f \cdot h, g, x_d) = \operatorname{res}(f, g, x_d) \cdot \operatorname{res}(h, g, x_d)$ and $\operatorname{res}(f, g \cdot h, x_d) = \operatorname{res}(f, g, x_d) \cdot \operatorname{res}(f, h, x_d)$ hold. But in both equalities the factors on the right side are elements in $k[x_1, \ldots, x_d]$ and therefore the equalities also hold in $k[x_1, \ldots, x_d]$. □

For resultants and subresultants of univariate polynomials we have proven that they are invariant under translation, remember Corollary 2.7 and Proposition 2.12. This nice property can be extended to resultants of multivariate polynomials:

Corollary 2.20: Given $f, g \in k[x_1, \ldots, x_d]$ regarded as polynomials in x_d with coefficients in $k[x_1, \ldots, x_{d-1}]$ and leading coefficients not equal to the zero polynomial. For a number $c = (c_1, \ldots, c_d) \in \bar{k}^d$ in the algebraic closure of k we define $\tilde{f}(x) := f(x_1 + c_1, \ldots, x_d + c_d)$ and $\tilde{g}(x) := g(x_1 + c_1, \ldots, x_d + c_d)$. Then the following equality holds:

$$\operatorname{res}(\tilde{f}, \tilde{g}, x_d) \;=\; \operatorname{res}(f, g, x_d)(x_1 + c_1, \ldots, x_{d-1} + c_{d-1}).$$

Proof. A translation along the vector c can be split into d translations $x_i \longrightarrow x_i + c_i$ along the coordinate axes $x_1, \ldots, x_d$. The resultant is invariant under translation along the x_d-axis, because we can apply Proposition 2.12 for univariate polynomials. A translation of $\operatorname{res}(f, g, x_d)$ along one of the other axes x_i, $i \in \{1, \ldots, d-1\}$, by a factor c_i causes a translation of each coefficient entry in the Sylvester matrix of f and g. Since the translation does not decrease the total degrees of the coefficients of f and g, the determinant of this matrix is equal to the resultant of $\tilde{f}$ and $\tilde{g}$ with respect to x_d. □

2.2.6 Subresultants of multivariate polynomials

The same way as for resultants we also transfer the notion of the lth *subresultant* to multivariate polynomials. We consider two polynomials $f, g \in k[x_1, \ldots, x_d]$ as polynomials in one variable, for example in x_d. Their coefficients are polynomials in $k[x_1, \ldots, x_{d-1}]$. We delete the last $2l$ columns of the Sylvester matrix of f and g and the last l rows of f entries and the last l rows of g entries. The determinant of this matrix is again a polynomial in $k[x_1, \ldots, x_{d-1}]$ and we call it the lth *subresultant* of f and g with respect to x_d: $\operatorname{sres}_l(f, g, x_d) \in k[x_1, \ldots, x_{d-1}]$. For univariate polynomials we have shown that the subresultants can help us to determine the degree of the greatest common divisor of two polynomials. In Corollary 2.15 we have derived an important fact for two multivariate polynomials f and g that have a constant leading coefficient when regarded as polynomials in x_d: The vanishing of the resultant of these two polynomials over $\mathbb{C}$ is a necessary and sufficient condition for extending a $(d-1)$-dimensional point $(c_1, \ldots, c_{d-1})$ to a common solution $(c_1, \ldots, c_d)$ of f and g. Analogously the vanishing of the first h subresultants for $(c_1, \ldots, c_{d-1}) \in \mathbb{C}^{d-1}$ is a necessary and sufficient condition for the existence of h complex numbers $c_{d_1}, \ldots, c_{d_h}$, not necessarily different, with $f(c_1, \ldots, c_{d-1}, c_{d_i}) = 0 = g(c_1, \ldots, c_{d-1}, c_{d_i})$:

Lemma 2.21: Given $f, g \in \mathbb{C}[x_1, \ldots, x_d]$ regarded as polynomials in x_d with coefficients in $\mathbb{C}[x_1, \ldots, x_{d-1}]$ and constant leading coefficients f_n and g_m. The polynomials $f(c_1, \ldots, c_{d-1}, x_d) \in \mathbb{C}[x_d]$ and $g(c_1, \ldots, c_{d-1}, x_d) \in \mathbb{C}[x_d]$ have a greatest common divisor of degree h for a point $(c_1, \ldots, c_{d-1}) \in \mathbb{C}^{d-1}$ if and only if h is the least index l for which $\mathrm{sres}_l(f, g, x_d)$ does not vanish at $(c_1, \ldots, c_{d-1})$.

Proof. The proof works in the same way as for the resultant of two multivariate polynomials. □

The subresultant of two multivariate polynomials f and g with respect to x_d is invariant under translation along the x_d-axis. Or, more precisely:

Corollary 2.22: Let $f, g \in k[x_1, \ldots, x_d]$ be two polynomials with leading coefficients not equal to the zero polynomial when regarded as polynomials in x_d with coefficients in $k[x_1, \ldots, x_{d-1}]$. For a number $c = (c_1, \ldots, c_d) \in \bar{k}^d$ in the algebraic closure of k we define $\tilde{f}(x) := f(x_1 + c_1, \ldots, x_d + c_d)$ and $\tilde{g}(x) := g(x_1 + c_1, \ldots, x_d + c_d)$. For $0 \leq l \leq \deg(f, g)$ we have

$$\mathrm{sres}_l(\tilde{f}, \tilde{g}) \ = \ \mathrm{sres}_l(f, g)(x_1 + c_1, \ldots, x_{d-1} + c_{d-1}).$$

Proof. The proof works in the same way as for the resultant of two multivariate polynomials in Corollary 2.20. □

Another important statement is as follows: If for an index $h \in \mathbb{N}$ all indices $1 \leq l \leq h$ lead to $\mathrm{sres}_l(f, g, x_d) \equiv 0$, then f and g have a common factor of degree at least h.

Corollary 2.23: Let $f, g \in k[x_1, \ldots, x_d]$ be two polynomials with positive degree in x_d. The first h subresultants $\mathrm{res}_l(f, g, x_d)$, $0 \leq l < h$, are equal to the zero polynomial if and only if f and g have a common factor in $k[x_1, \ldots, x_d]$ which has degree at least h in x_d.

Proof. Also this proof works analogously to the one for resultants of multivariate polynomials. Since f and g are polynomials in x_d with coefficients in $k[x_1, \ldots, x_{d-1}]$, the field the coefficients lie in is $k(x_1, \ldots, x_{d-1})$. We apply Proposition 2.11 to $f, g \in k(x_1, \ldots, x_{d-1})[x_d]$: $\mathrm{sres}_l(f, g, x_d) \equiv 0$ for $0 \leq l < h$ if and only if f and g have a common factor in the ring $k(x_1, \ldots, x_{d-1})[x_d]$ over the field extension $k(x_1, \ldots, x_{d-1})$ with degree at least h in x_d. What again remains to prove is that this is equivalent to having a common factor in $k[x_1, \ldots, x_d]$ of degree at least h in x_d. This immediately follows from Proposition 2.18 by induction on the degree of the common factor of f and g and by successively factoring the two polynomials. □

2.3 Root isolation

In the previous section we have developed a method to project the intersection curve of two quadrics into the plane. Let p, q, and r be quadrics in space. Projecting the intersection curves of several pairs of quadrics leads to a planar arrangement of algebraic curves. For example we may obtain the two curves $f = \mathrm{res}(p, q, z)$ and $g = \mathrm{res}(p, r, z)$. Again with the help of resultants we can project the intersection points of f and g onto the x- and y-axis. This yields univariate polynomials $X = \mathrm{res}(f, g, y)$ and $Y = \mathrm{res}(f, g, x)$ in x and y, respectively. We are interested in the real roots of X and Y because they are the candidates for x- and y- coordinates of intersection points of f and g. The two curves f and g are projected intersection curves of p with q and with r. Therefore the roots of X and Y are also the candidates for x- and y-coordinates of common intersection points of p, q, and r.

2.3.1 Algebraic numbers

Consider a rational polynomial $u(x) = \sum_{i=0}^{n} a_i x^i$. We are interested in the real roots of u. That means we want to solve the equation $u(x) = 0$ and find real numbers α with $u(\alpha) = 0$. By definition, an *algebraic number* is a root of some polynomial $u \in \mathbb{Q}[x]$. The Fundamental Theorem of Algebra states that every non-constant polynomial $u(x) \in \mathbb{C}[x]$ has a root $\alpha \in \mathbb{C}$. The question is to find these roots. For $\deg(u) > 2$ there is no general way via radicals to explicitly compute the roots in every case. But we are only interested in the real ones and, as we will see, it is sufficient to know an *isolating interval* for each real root α of u. That means we compute two rational numbers a and b such that α is the one and only real root of u in $[a, b]$. There are various methods of determining these isolating intervals like the algorithm of Uspensky, Sturm sequences, Kronecker's Algorithm, or isolation by differentiation. For a good survey see [21]. We will just quote Uspensky's algorithm because it is quite fundamental and often used in praxis.

2.3.2 Uspensky algorithm

First let us state a basic tool for real root counting, namely Descartes' rule of sign. Let $(u_1, \ldots, u_s)$ be a sequence of real numbers and $(u'_1, \ldots, u'_t)$ the subsequence of all non-zero real numbers. Then $\mathrm{var}(u_1, \ldots, u_s)$, the number of *sign variations*, is the number of indices i, $1 \leq i < t$, such that $u'_i u'_{i+1} < 0$. With this notation Descartes' rule of signs is as follows:

Theorem 2.24: The sign variation $\mathrm{var}(u)$ in the sequence $(u_n, u_{n-1}, \ldots, u_1, u_0)$ of coefficients of the polynomial $u(x) = \sum_{i=0}^{n} u_i x^i$ exceeds the number of positive real roots of $u(x)$ by some non-negative even number.

Let $u(x)$ be a squarefree polynomial. This can be achieved by factoring out the gcd of u and its derivative u'. Since the negative real zeros of $u(x)$ are the positive real zeros of $u(-x)$, and since $u(0) = 0$ iff $u_0 = 0$, we can restrict ourselves to isolate the positive zeros of u. From Descarte's rule of signs we can conclude that if $\mathrm{var}(u) = 0$, u has no positive zeros, and if $\mathrm{var}(u) = 1$, u has exactly one positive zero. Uspensky shows that after a finite number of transformations $v(x) = u(x+1)$ and $w(x) = (x+1)^n u(1/(x+1))$ one arrives at polynomials having sign variation 1 or 0. This is the basic idea of his algorithm. His method, however, suffers from an exponential computing time since the bisection into subintervals is not balanced. This drawback was eliminated in the modified Uspensky algorithm by Collins and Akritas [20] using bisection by midpoints.

2.3.3 Decreasing the width of an interval

Once we have an isolating interval $[a, b]$ of a root of a polynomial $u \in \mathbb{Q}[x]$, we can easily decrease its width using bisection by midpoints. Let us assume without loss of generality that u is squarefree. We know that u has exactly one root in $[a, b]$, so $u(a) \cdot u(b) < 0$. For the rational midpoint $c := \frac{a+b}{2}$ of the interval there are two possibilities:

1. either $u(c) = 0$, in which case we are lucky because we have found a rational root,

2. or exactly one of the intervals $[a, c]$ or $[c, b]$ contains a root of u and the other one is empty. The one containing the root can simply be determined by testing $u(a) \cdot u(c) < 0$. If the inequality is true, the root is contained in $[a, c]$, otherwise it is in $[c, b]$.

2.3.4 Pairwise root separation

Sometimes it is necessary to pairwise separate the root isolating intervals of two polynomials $u, v \in \mathbb{Q}[x]$. We want to achieve that a root isolating interval $[a, b]$ of a root α of u and a root isolating interval $[c, d]$ of a root β of v either are identical, in the case $\alpha = \beta$, or disjoint otherwise. Again assume without loss of generality that u and v are squarefree. Of course, we can compute such isolating intervals by applying Uspensky's algorithm to the polynomial $u \cdot v$. This has the disadvantage that we deal with a polynomial of higher degree than u and v. After computing the isolating intervals for $u \cdot v$, we have to assign each interval to u or v or to both. Another problem is that in the case we already know isolating intervals for u and v, we do not use this information. We can avoid this by computing $g := \gcd(u, v)$, $\tilde{u} := u/g$, and $\tilde{v} := v/g$. We first determine the intervals of u and v that contain a root of g. An interval $[a, b]$ can again be tested by computing $g(a) \cdot g(b) < 0$. The intervals $[a, b]$ of u and $[c, d]$ of v containing the same root of g are replaced by $[a, b] \cap [c, d]$. The remaining intervals of u and v have to be shrunk until they are pairwise disjoint using bisection by midpoints as described before.

Chapter 3

The basic algorithmic ideas

We consider the arrangement $\mathcal{A}(P)$ of a set of quadric input surfaces $P = \{p_1, \ldots, p_l\}$. For example $\mathcal{A}(P)$ for the set $P = \{R, G, B\}$ of ellipsoids in Figure 3.1.

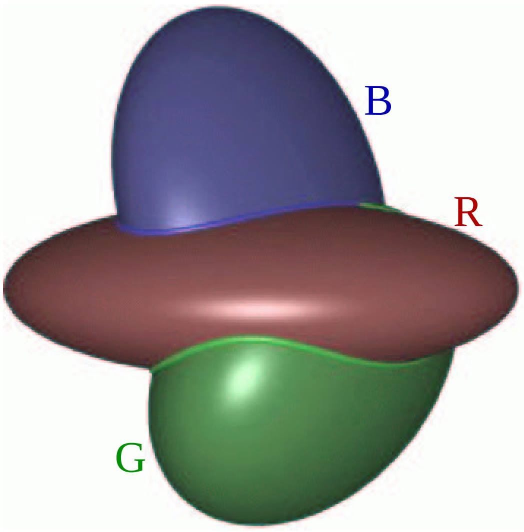

Figure 3.1: Arrangement of three red, blue, and green ellipsoids R, B, and G, respectively

Our initial question is how to compute the topological description of a cell in the arrangement. But independent of the special information and application one may be interested in, the basic computation that has to be done in nearly all cases is: Given a quadric p, locate and sort all vertices along the intersection curves running on the surface of p. Each such curve is an intersection curve of p with another quadric q. A vertex is an intersection point of p and two other quadrics q and r. Equivalently we can say that it is an intersection point of two spatial intersection curves, namely the intersection curve of p and q and the one of p and r. If, for each quadric p, we can compute the 2-dimensional subarrangement on its surface, combining these results to the desired description of the cell is a problem of discrete combinatorics and data structures. We are mainly interested in the aspect of exact algebraic computation and therefore do not treat this last step.

For each quadric p we want to compute the 2-dimensional sub-arrangement that is induced by $\mathcal{A}(P)$ on the surface of p. For example, consider the red ellipsoid R and the blue ellipsoid G. Their intersection curve is the blue spatial curve lying on the red ellipsoid and consisting of two connected components, see Figure 3.2. Its intersection points with the remaining quadrics, in our example simply the ones with the green ellipsoid G, divide the blue curve into edges of the arrangement. The green curve is the intersection curve of the red and the green ellipsoid. We want to locate and sort the intersection points of the blue curve with the green ellipsoid along the different branches of the blue curve, or equivalently the intersection points of the blue and the green spatial intersection curve. We have marked their intersection points by small arrows.

3.1 Reduction to planar arrangements

3.1.1 The projection phase

The question we have to answer is how to locate and topologically sort, for a fixed quadric p, the intersection points of curves running on its surface. The bad news is that in general there is no rational parameterization of the intersection curve of two quadrics, see [28]. Moreover, our aim is to develop an approach that might be extendible to more general surfaces. Therefore we cannot work directly in space. But we can reduce the problem to one of computing intersection points of planar algebraic curves of degree at most 4 by projecting the intersection curves into the plane. As we have seen in the previous chapter, under some assumptions the projection can be algebraically realized via resultants. This reduction also works for arbitrary surfaces. Only the algebraic degree of the resulting curves differs. All ideas developed in this chapter are applicable to arbitrary surfaces.

In our example in Figure 3.2, the polynomial

$$\begin{gathered}408332484x^4 + 51939673y^4 - 664779204x^3y - 24101506y^3x + 564185724x^2y^2 \\ -250019406x^3 + 17767644y^3 + 221120964x^2y - 123026916y^2x \\ +16691919x^2 + 4764152y^2 + 14441004xy + 10482900x + 2305740y - 1763465\end{gathered}$$

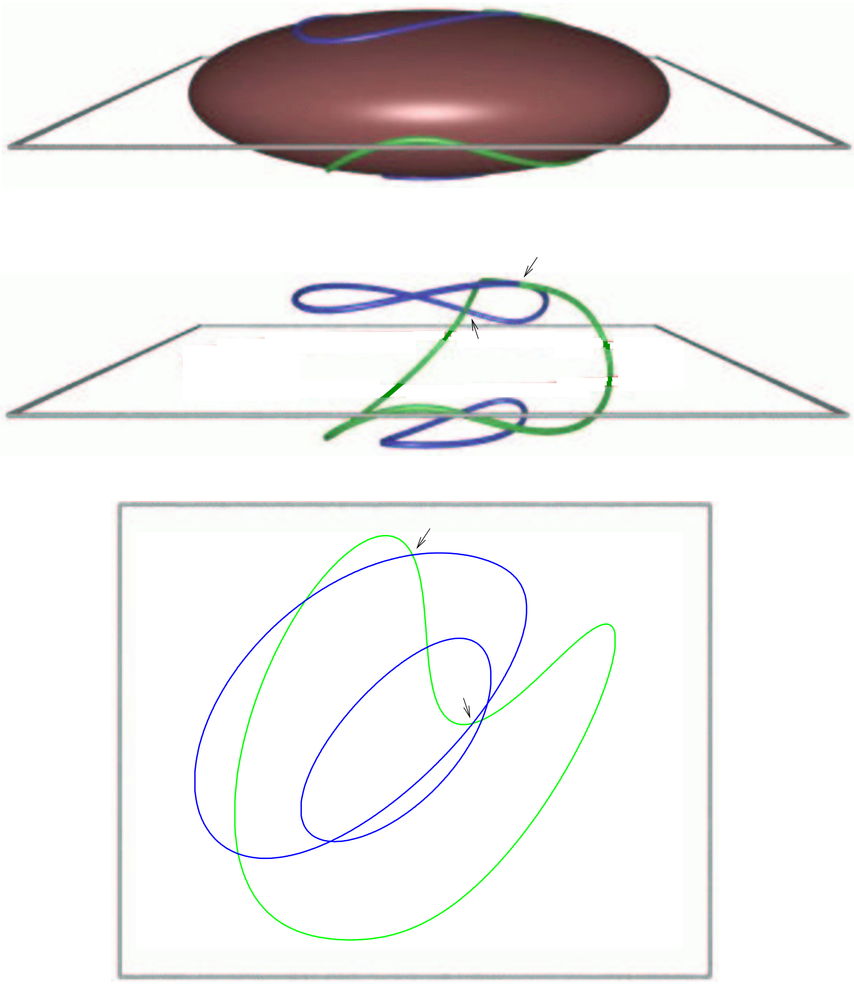

Figure 3.2: Project the intersection curves into the plane

defines the projected green curve. We project the blue and the green intersection curve into the plane.

In order to correctly interpret the resultant as the projected intersection curve, we assume throughout this and the next chapters that the quadratic input polynomials are *squarefree* and *generally aligned*, and that each two of them have a *disjoint factorization*. These terms are defined as follows:

Definition 3.1: 1. A polynomial $f \in \mathbb{Q}[x_1, \ldots, x_d]$ is *squarefree* if the resultant $\text{res}(f, f_{x_d}, x_d)$ is not equal to the zero polynomial. That means f and f_{x_d} do not share a common factor, or, in other words, the factorization of f is squarefree.

2. A polynomial $f \in \mathbb{Q}[x_1, \ldots, x_d]$ is *generally aligned* if it has a constant leading

coefficient when regarded as a polynomial in one variable x_d. This removes the negative effect that the resultant of two polynomials f and g with respect to x_d may vanish for a point $a = (a_1, \ldots, a_{d-1})$, although a is not extendible to a common solution of f and g.

3. Two polynomials $f, g \in \mathbb{Q}[x_1, \ldots, x_d]$ have a *disjoint factorization* if $\text{res}(f, g, x_d) \not\equiv 0$, i.e. if f and g do not share a common factor.

These assumptions constitute no restriction on the input quadrics. In Chapter 7 we will show how to realize them for each kind of input.

During the projection we loose the spatial information. Points of intersection curves on the upper part of p and on the lower part of p are projected on top of each other. This can cause self-intersections. The two branches of the blue curve, one running on the upper and one on the lower part of the red ellipsoid, are projected on top of each other generating two self-intersections, see Figure 3.2. Moreover, in space the green and the blue curve had 2 intersection points. The projected curves in comparison have 6 intersection points, 4 of them resulting from the loss of spatial information.

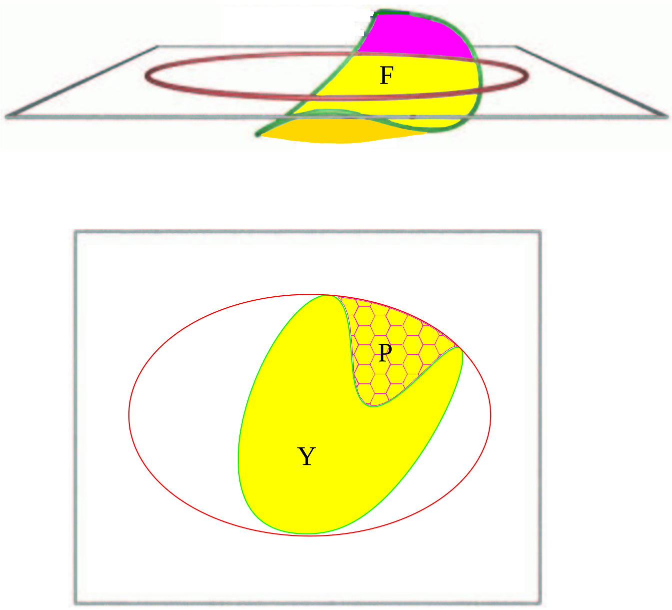

Figure 3.3: Only projecting the intersection curves is not sufficient

Another negative side effect of the projection is the following: Look at the green and red spatial curves in Figure 3.3. The red curve is the silhouette of the red ellipsoid in the example above. The green curve is again the intersection curve of the red and the green ellipsoid. The latter consists of one connected closed branch, bounding a face

F colored pink and yellow. The pink portion of F lies on the upper part of the red ellipsoid, the yellow portion on the lower part. The projected green curve also consists of one connected closed branch bounding the yellow region Y. But not all points of F are projected onto points of Y. The points on the pink portion of F and the ones lying vertically underneath on the yellow portion are projected onto points of the pink–yellow region P. P is bounded by the green and red curve. The example makes clear that only projecting the intersection curve into the plane is not sufficient. We have to identify the region P and for that we additionally need the projection of the silhouette of p.

The projection of the silhouette can be computed with the help of the partial derivative $p_z := \partial p / \partial z$ of p with respect to z. The interpretation of the resultant $\mathrm{res}(p, p_z, z)$ as the projection of the silhouette is permissible because p is generally aligned:

Proposition 3.2: Let $p \in \mathbb{C}[x, y, z]$ be a polynomial with leading coefficient a not equal to the zero polynomial when regarded as polynomial in z with coefficients in $\mathbb{C}[x, y]$. If $\mathrm{res}(p, p_z, z)$ vanishes at $(c_x, c_y) \in \mathbb{C}^2$, then either

1. a vanishes at (c_x, c_y), or
2. there is $c_z \in \mathbb{C}$ such that c_z is a root of multiplicity 2 of $p(c_x, c_y, z)$.

Proof. We have shown in Proposition 2.14 that $p(c_x, c_y, z)$ and $p_z(c_x, c_y, z)$ have a common root when $a(c_x, c_y)$ is non-zero. So $p(c_x, c_y, z)$ has a root of multiplicity 2 when $a(c_x, c_y)$ is non-zero. □

A third problem, which we already mentioned in the last chapter, is the following: It can happen that a point of a curve $f = \mathrm{res}(p, q, z)$ in the real plane originates from a complex common point of p and q. While examining the planar curves we will not take care of this event. Such projected non-real intersection points just refine the planar arrangements of all projected real spatial intersection curves we are interested in.

The algorithmic part of projecting intersection curves of quadrics and the silhouette into the plane we call *projection step* or *projection phase*. It is easy to see that the projection step makes no sense for linear input polynomials. For a linear one we just substitute its coefficients into all other input quadrics and directly compute the arrangement of the rational quadratic intersection curves lying on this plane. Because each resulting bivariate polynomial has total degree at most 2, computing their planar arrangement is much easier than the problem we will discuss in the next two chapters. So we will ignore it in the following. Of course, our method can also be applied to this kind of arrangement.

The whole projection step for our example proceeds like shown in Figure 3.4.

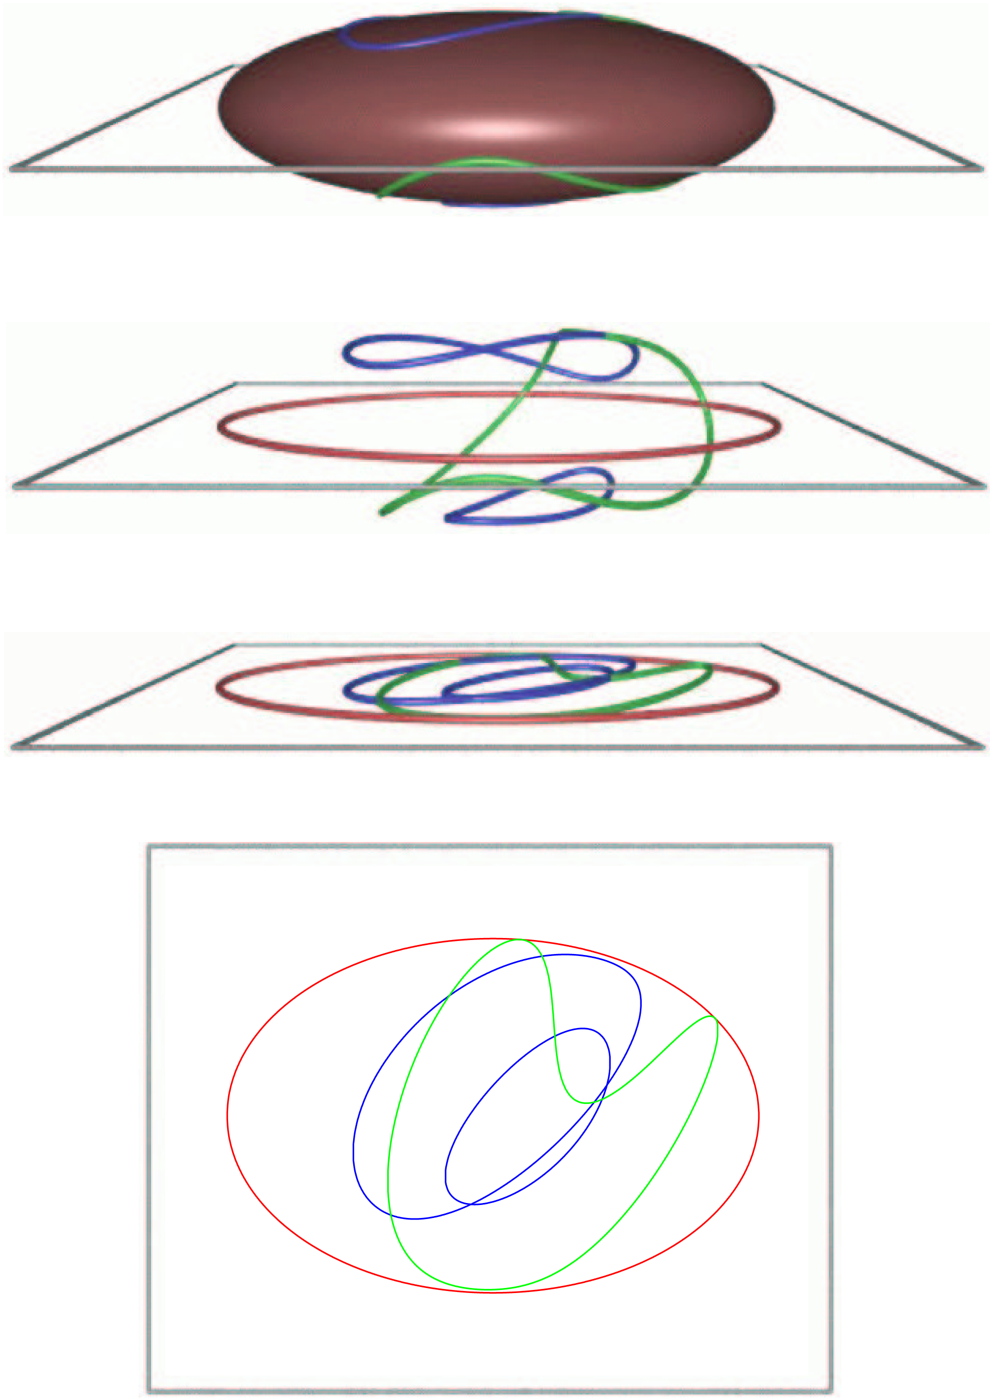

Figure 3.4: Project all intersection curves and the silhouette of p into the plane

3.1.2 The planar arrangement

In the planar arrangement, we obtain from the projection phase, there are two different types of planar curves and exactly one curve is of the first type:

silhouettecurve: The projection of the silhouette of p. The planar curve is the set of roots of $\text{res}(p, p_z, z)$ and its algebraic degree is bounded from above by $\deg(p) \cdot \deg(p_z) = 2$. This is the red curve in our example.

cutcurve: The projection of the spatial intersection curve of p with another quadric q. The planar curve is the set of roots of $\text{res}(p, q, z)$ and its algebraic degree is at most 4.

3.2 Computing planar arrangements

We have to compute the planar arrangements we obtain from the projection phase. As we have seen, each arrangement consists of one silhouettecurve and a set of cutcurves.

3.2.1 Event points

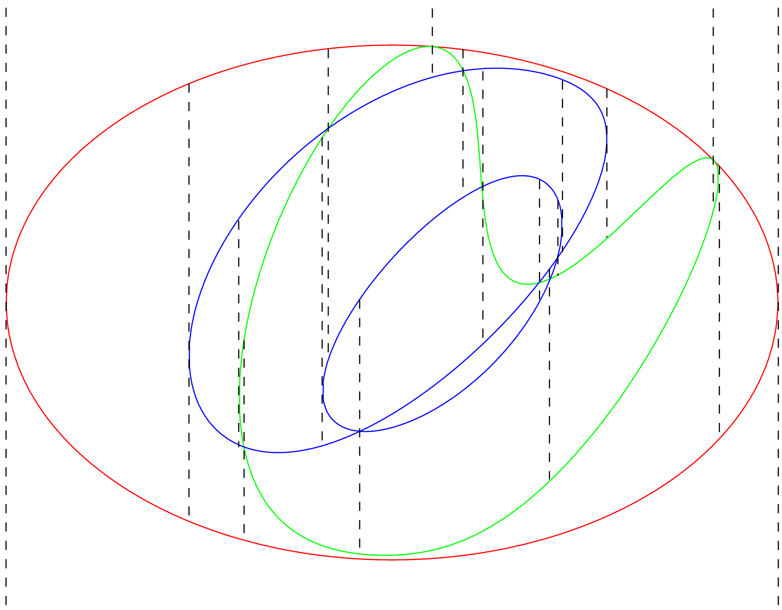

Figure 3.5: A trapezoidal decomposition of the planar arrangement

One way of representing the arrangement would be to store its trapezoidal decomposition [56]. The motivation and mathematical definition for such a decomposition is as follows: The curves decompose the plane into several regions. Unfortunately, the regions can have complicated shapes. It is convenient to refine this partition further, as in Figure 3.5, by passing a vertical attachment through all

1. intersection points of two curves,
2. singular points of one curve, for example self-intersection points,
3. and extreme points.

Each vertical attachment extends upwards and downwards until it hits another curve, and if no such curve exists, then it extends to infinity. Of course we cannot explicitly compute the vertical attachments because in general they have irrational endpoints. But we are interested in the topology of the trapezoidal decomposition and this information does not depend on the exact computation of the attachments. The important step is to determine and locate all intersection points, extreme and singular points, consider Figure 3.6.

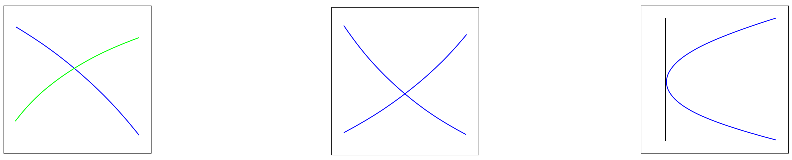

Figure 3.6: Intersection, singular, and extreme points

In the mathematical preliminaries we also defined vertical turning points. At such a point only the sign of the curvature changes, but it does not cause a vertical line of the decomposition. Because of that, there is no need to consider them. Why did we define them nevertheless? Because we would like to interpret also singular and extreme points of a curve f as intersection points of two curves, namely of f and f_y. The two curves f and f_y intersect in singular, extreme, and unfortunately also in vertical turning points. The good news is that these three kinds of points remain the only intersection points:

Theorem 3.3: The singular points and the points with a vertical tangent are exactly the intersection points of f and f_y.

Proof. By definition a point (a, b) is a point with a vertical tangent or a singular point of f if and only if $f(a, b) = 0$ and $f_y(a, b) = 0$, i.e if (a, b) is an intersection point of f and f_y. □

We would like to reduce the question of determining points causing a vertical line of the trapezoidal decomposition to the question of locating intersection points of two curves. This algorithmic aim has the consequence that we include vertical turning points to the set of event points, although there is no geometric need to do that.

Definition 3.4: The *event points* of a planar arrangement induced by a set F of planar curves are defined as the intersection points of each two curves $f, g \in F$ and the intersection points of f and f_y for all $f \in F$.

For illustration of the definition above consider Figure 3.7. The singular and extreme points of f are marked by the small boxes. In our special example the singular points are all self-intersections and there is no vertical turning point. In the second picture of Figure 3.7 one can see that the magenta curve $g = f_y$ exactly cuts through the singular points and the points with a vertical tangent of f.

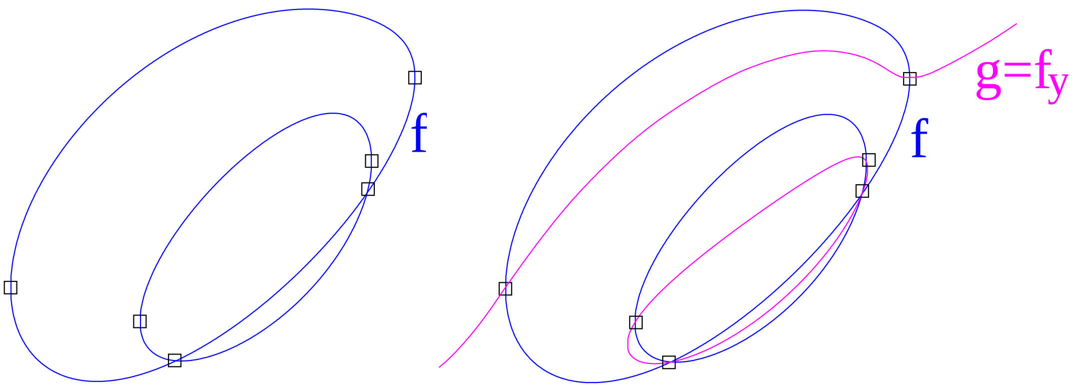

Figure 3.7: Singular and extreme points of f are common roots of f and $g = f_y$

3.2.2 Computing intersection points of two curves

With our last observations we have reduced the problem of computing the event points in the planar arrangement to the question of determining intersection points of two curves f and g. For F being the set of all algebraic curves of the planar arrangement, we have to locate the intersection points of all pairs of algebraic curves f and g, whereby either $f, g \in F$ or $f \in F$ and $g = f_y$. Remember that our planar arrangements result from the projection phase. Therefore the set F consists of exactly one silhouettecurve and a set of cutcurves. Keeping this in mind, we can distinguish four different types of pairs of curves, the intersection points of which we want to locate:

1. $f \in F$ and $g = f_y$, whereby f is the silhouettecurve.

2. $f \in F$ and $g = f_y$, whereby f is a cutcurve.

3. $f, g \in F$ and one of the two curves is the silhouettecurve and the other one is a cutcurve.

4. $f, g \in F$ and both curves are cutcurves.

For illustration look at our small example in Figure 3.8.

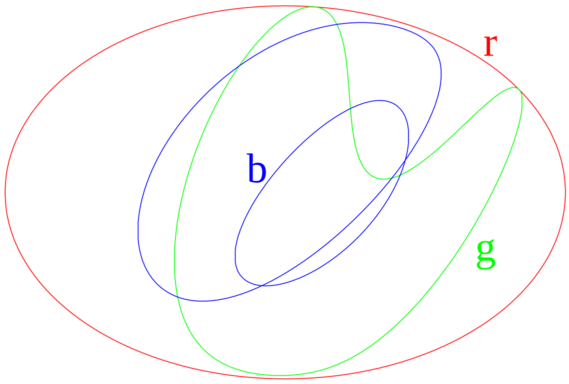

Figure 3.8: Compute all intersection points between (r and r_y), (b and b_y), (g and g_y), (r and b), (r and g), and (b and g)

We want to compute the intersection points of two curves f and g. In general common points of f and g will have irrational, or even complex, coordinates. We are only interested in the real common roots, so we face two problems:

1. Although all mathematical tools work over $\mathbb{C}$, we have to distinguish real roots from complex roots.

2. Although the intersection points will have irrational coordinates, we have to locate and characterize them exactly and unambiguously.

Our solution again works in the spirit of cylindrical algebraic decomposition. With the help of resultants we compute the two univariate polynomials $X = \text{res}(f, g, y) \in \mathbb{Q}[x]$ and $Y = \text{res}(f, g, x) \in \mathbb{Q}[y]$. Their real roots contain the x- and y-coordinates of all intersection points, respectively. Like in the projection phase from space to the plane, we assume in the following that all polynomials $f \in F$ are generally aligned and squarefree and that each two polynomials $f, g \in F$ have a disjoint factorization. That means each polynomial $f \in F$ has a constant leading coefficient regarded as a polynomial in y as well as regarded as a polynomial in x. Furthermore, for two polynomials f and g, whereby either $f \in F$ and $g = f_y$ or $f \in F$ and $g \in F$, their resultants with respect to y and x are not equal to the zero polynomial. In Chapter 7 we describe how to establish these conditions.

Let $\mathcal{R}(X)$ be the set of real roots of X and $\mathcal{R}(Y)$ be the ones of Y. Each real intersection point of f and g is a member of the grid

$$\textsc{Grid}(X,Y) \;:=\; \mathcal{R}(X) \times \mathcal{R}(Y) \;=\; \{(r_x, r_y) \mid r_x \in \mathcal{R}(X)\ ,\ r_y \in \mathcal{R}(Y)\}.$$

We conclude that the search for real intersection points is limited to examining only the points on the grid $\textsc{Grid}(X,Y)$. But the coordinates of the points on the grid are only given as roots of the polynomials X and Y. The degrees of each of these polynomials can be 16. Therefore we cannot compute the points of $\textsc{Grid}(X,Y)$ explicitly. Instead, with the help of a root isolation algorithm, we determine rational interval representations for the real algebraic numbers of $\mathcal{R}(X)$ and $\mathcal{R}(Y)$. For the x-axis this gives us rational intervals, each containing one real root of X. Every interval $[a,b]$, $a,b \in \mathbb{Q}$, can be vertically extended to a stripe in the plane consisting of all points (x,y) with $a \le x \le b$ and $y \in \mathbb{R}$. The same way each interval on the y-axis can be extended to a horizontal stripe. So the intervals on the x- and on the y-axis define narrow rational stripes parallel to the y- and parallel to the x-axis, respectively. The intersection of the stripes yields disjoint boxes with rational corners containing the points of $\textsc{Grid}(X,Y)$. Each box contains at most one real intersection point of f and g.

We want to avoid that a box entirely contains a loop of one of the curves like in Figure 3.9. Therefore we additionally choose all boxes small enough to guarantee that at most one intersection point between f, f_y and g, g_y takes place. This can easily be done by pairwise separating the real roots of X, $\mathrm{res}(f, f_y, y)$, and $\mathrm{res}(g, g_y, y)$, and of Y, $\mathrm{res}(f, f_y, x)$, and $\mathrm{res}(g, g_y, x)$. Now the isolated roots of X and Y define boxes that contain at most one intersection point of f and g and at most one singular or extreme point of f and of g. In the case that two of these event points take place in one box, they take place in the same point.

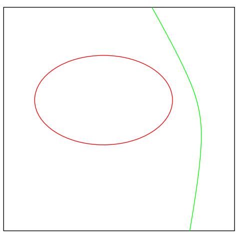

Figure 3.9: We want to avoid the case that the box entirely contains the red loop

It remains to test each box for a real intersection point. Unfortunately, the number of boxes is nearly quadratic in the number of intersection points. In the example in Figure 3.10 we have to distinguish the empty yellow boxes from the red ones that contain an intersection point.

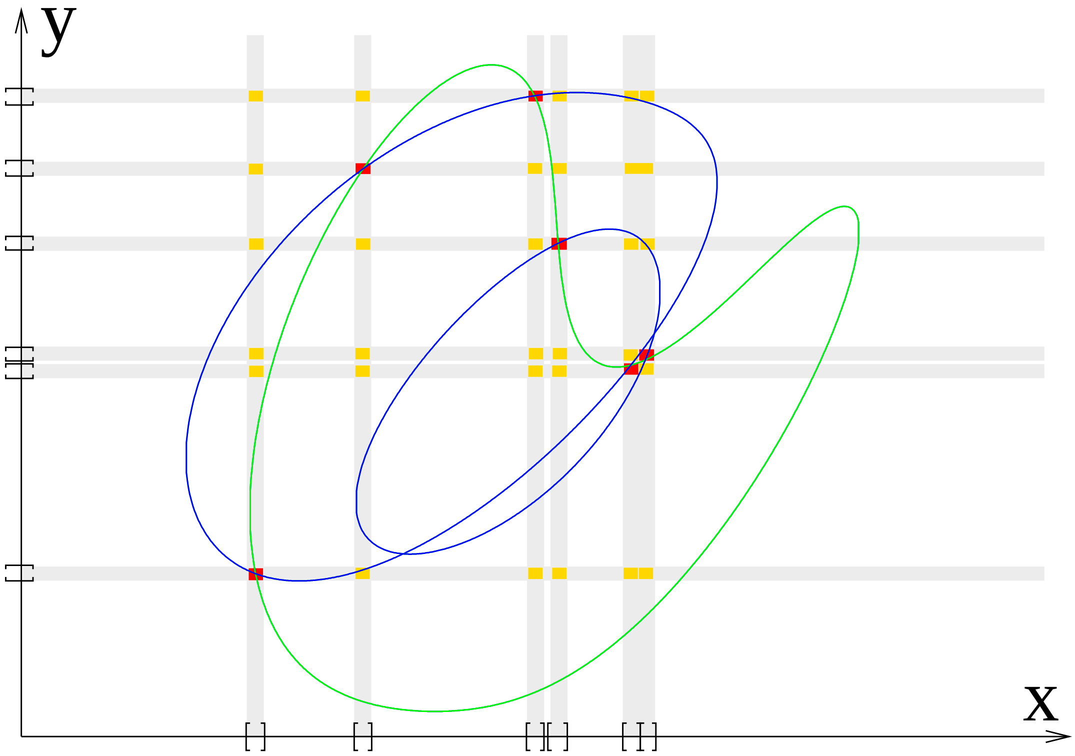

Figure 3.10: Distinguish the empty yellow boxes from the red ones

3.2.3 Testing a box for an intersection point

We have to answer the question whether a box with rational corners contains an intersection point of f and g or not. This reduction from a problem in the complex plane to one that is locally limited to a small box in the real plane answers our first question how to distinguish real roots from complex ones. The root isolation algorithm applied to $\mathrm{res}(f, g, y)$ and to $\mathrm{res}(f, g, x)$ separates their *real* roots and therefore is responsible for the distinction. We only consider boxes in the real plane and ask for intersection points inside them. The second problem we formulated is that intersection points can have irrational coordinates and this of course is still true. We have no information about what is happening inside the box. The only thing we can obtain is some information about the boundary of the box. We compute the sequence of hits of the curves with the boundary: Let $[a, b]$ and $[c, d]$ be the defining intervals of the box on the x- and y-axis, see Figure 3.11. Then the box has the rational corners (a, c), (a, d), (b, d), and (b, c).

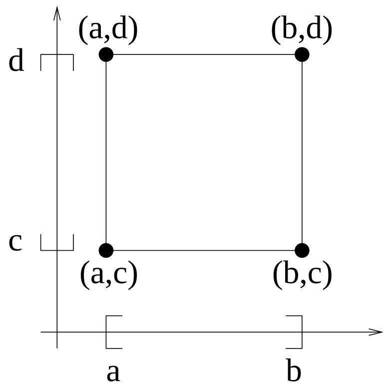

Figure 3.11: A box with its rational corners

We substitute the rational number a for x in f and g, yielding two polynomials $f(a, y), g(a, y) \in \mathbb{Q}[y]$. For these polynomials we pairwise isolate the real roots in the interval $[c, d]$ and determine the multiplicity of each root. This leads to the sequence of hits of the curves with the left edge of the box, counted with multiplicities. We analogously do the same for the upper, right, and lower edge of the box.

Can the sequence of hits of the curves with the boundary of the box help us to determine the behavior of the curves inside the box? If there are exactly two hits with each curve, counted with multiplicities, and the hits alternate, then we can be sure that there is an intersection point inside the box at which the two curves cross each other, see Figure 3.12. This method of locating for example transversal intersection points we call *simple box hit counting*. It is also discussed in [44].

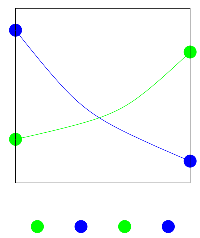

Figure 3.12: Transversal intersections can be solved with simple box hit counting

Simple box hit counting:

```
begin
  determine sequence of hits of f and g with the box
  while (#hits(f) >= 2) or (#hits(g) >= 2)
     shrink the box
     determine sequence of hits of f and g with the box
  if (#hits(f) < 2) or (#hits(g) < 2)
     output: 0   // empty box
  else
     if (hits alternate)
        output: 1 // intersection point
     else
        output: ???
end
```

It is easy to see that simple box hit counting has the output `1` if and only if the box contains an intersection point at which f and g cross each other:

Proposition 3.5: Let f and g be two curves in the plane. Let furthermore B be a box with rational corners such that there is at most one point inside the box that is an intersection point of f and g, or an intersection point of f and f_y, or an intersection point of g and g_y. The point can be an intersection point of several of these curves. Moreover assume that f and g have exactly two hits with the boundary of B, counted with multiplicities.
The hits of f and g on the boundary of B alternate if and only if B contains an intersection point at which f and g cross each other.

Proof. The Fundamental Theorem of Algebra implies that for a fixed real number x_0 the roots of $f(x_0, y) \in \mathbb{R}[y]$ are either real or occur in complex conjugate pairs. We can count the real roots along this line $x = x_0$. By continuously moving x_0 along the x-axis also the roots of $f(x_0, y)$ change continuously. We call the union of all i-th real roots the i-th *level* of f. For example consider the curves in Figure 3.13. In both cases the red points form the first level of f.

As can be seen in the examples, the levels of f may be non-differentiable at some points, they may have discontinuous jumps or even may disappear completely. There is a non-differentiable point if and only if two real roots of f coincide. A level makes a jump or disappears if and only if two real roots change to two complex ones. Between such points every level of the curve is the graph of a differentiable real function.

Due to our assumptions, a box contains at most one intersection point of f and f_y, so there is at most one non-differentiable point of the levels inside the box. If the curve f

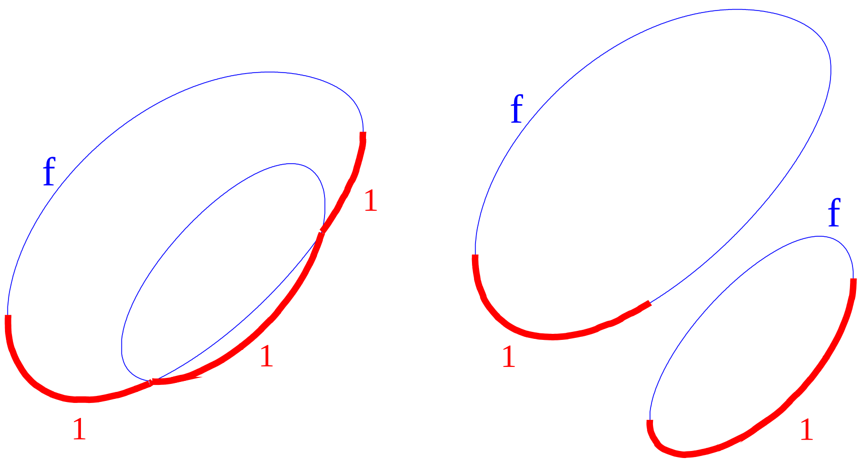

Figure 3.13: The red points form the first level of f

hits the boundary of the box exactly twice, then inside the box these points are either continuously connected by the same level or they are part of two distinct levels that obtain a common point and make discontinuous jumps to points outside the box or disappear completely. For illustration consider Figure 3.14.

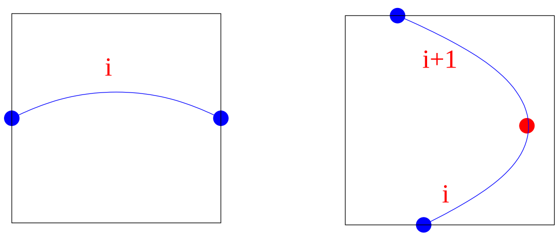

Figure 3.14: Either the points on the boundary are connected by the same level or by two distinct levels that have a common point inside the box.

In both cases the two points on the boundary are connected by a single real branch of the curve f. This branch divides the box into two regions. Assume that the hits of f with the boundary of the box alternate with the ones of a second curve g. Also the two points of g are connected by a branch and the two points lie on different sides of the branch of f. So f and g have to cross inside the box.

The other direction is even more easy. If f and g cross inside the box, then locally around this intersection point the curves appear in alternating order. Due to our assumptions, except in the intersection point of f and g all levels are graphs of differentiable functions and there is no further intersection point of f and g inside the box. So the order the branches does not change any more and we obtain alternating order of the hits of f and g at the boundary of the box. □

The problem of simple box hit counting is that it cannot detect non-singular tangential intersection points of f and g, see Figure 3.15. In the first box f and g have a non-singular tangential intersection, in the second box they have not. But the sequence of hits is identical in both cases. For non-singular tangential intersections our simple box hit counting algorithm ends up with the output ???. Both scenarios shown in Figure 3.15 can actually arise. Look at the two curves in Figure 3.16. The isolating intervals on the two coordinate axes lead to 4 boxes. In the upper left box there is a tangential intersection point of the two curves, whereas in the lower left box there is none. At the boundary of the boxes both situations lead to the same sequence of intersections.

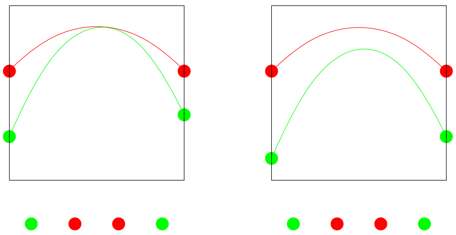

Figure 3.15: Tangential intersections cannot be solved with simple box hit counting

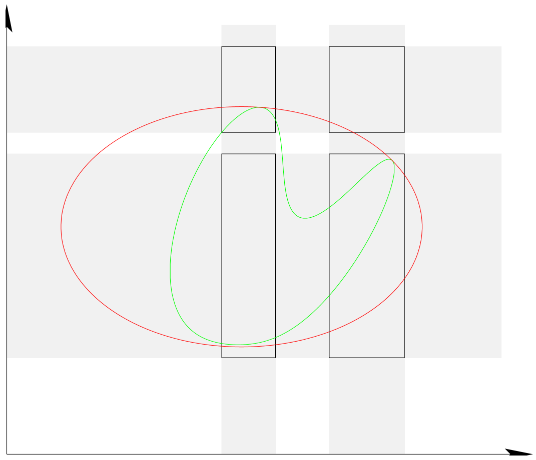

Figure 3.16: In the upper left box there is a tangential intersection whereas in the lower left box there is none

Other problematic events are self-intersections or isolated points of f, i.e. intersection points of f and $g = f_y$ that are singular points of f. For example consider Figure 3.17. In the first box the blue curve has a self-intersection, whereas in the second box it does not. The sequence of hits of f and f_y with the boundary of the box is the same in both cases. In the third box the two curves f and f_y intersect in an isolated blue point, but there are no hits of f with the boundary of the box. For self-intersections the first while-loop of the simple box hit counting algorithm runs forever. For the isolated point it gives the wrong answer `empty box`.

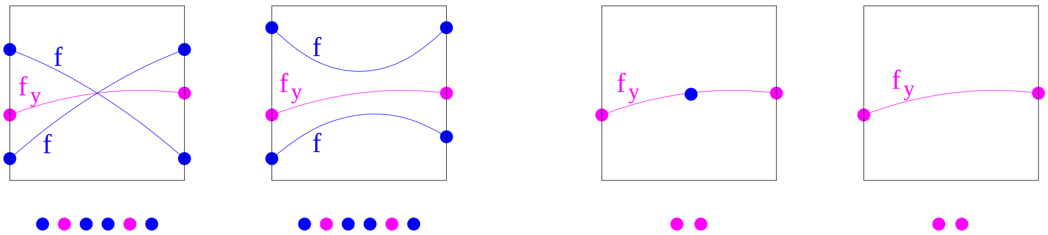

Figure 3.17: Some singular points cannot be solved with simple box hit counting

That means concerning the examination of the boxes we have to solve two problems in the following:

1. Find a method to avoid applying simple box hit counting to boxes that contain a tangential intersection in order to avoid infinite loops and wrong results.
2. Find methods to solve
   (a) non-singular tangential intersections and
   (b) singular points.

The answer to these two questions is crucial, because, as one can see in Figure 3.18, non-singular tangential intersections and singular points appear quite often in our arrangements. This is the reason why classical methods like the gap theorem [15] or multivariate Sturm calculation [53] are too expensive. In the next two chapters we will develop a new method that treats these cases in a fast and robust way.

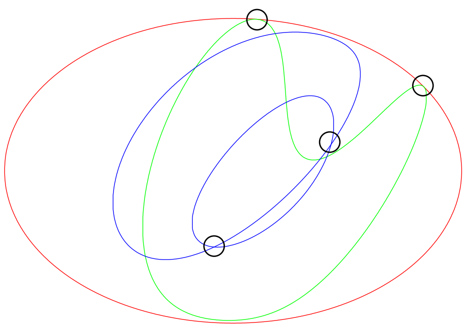

Figure 3.18: Tangential intersections and self-intersections appear quite often

Chapter 4

Intersection points of general curves

In this chapter we will give a first answer to the question, how to determine whether two curves f and g intersect inside a given box under the assumption that there is no singular point of one of the curves inside the box. Remember, that we obtain the boxes from computing the resultants of f and g on both coordinate axes and then separating their real roots. We will show how to predict the case that a box cannot contain a tangential intersection. With this pre-information, simple box hit counting is a suitable tool to obtain the correct result. This will answer our first question we posed at the end of the last chapter.

Next, we give a solution for determining *non-singular* tangential intersections inside a box. Thereby we assume that we know in advance that there is no singular point inside the box. This is a necessary and of course big assumption. The method we will develop for non-singular tangential intersections uses modified box hit counting.

At last, we describe a method how to factor the resultants on the coordinate axes. If any kind of factorization leads to polynomials on the x- and y-axis of degree at most 2, we can explicitly compute whether there are intersection points on their corresponding grid points or not.

All results in this chapter do not only apply to our special arrangements we obtain from projecting quadric intersection curves into the plane, but to arbitrary arrangements of planar curves. We can locate each intersection point between two arbitrary planar curves inside a box, but only if we know in advance that it is not a singular point of one of them.

4.1 Transversal intersections

In this section we will first provide some mathematical foundations and then answer the question of how to avoid applying simple box hit counting to boxes that contain a tangential intersection. The boxes are defined by the roots of the resultants on the x- and on the y-axis. The roots can have different multiplicities. Do these multiplicities have a geometric interpretation? The answer is positive. There is a strong connection between the kind of intersection two planar curves f and g have at a common point (a, b) and the multiplicity of the root a of the resultant $X = \text{res}(f, g, y)$. Of course, the considerations symmetrically hold for a root b of the resultant $Y = \text{res}(f, g, x)$.

4.1.1 Multiple roots of the resultant

For two curves f and g we want to investigate the roots of $X = \text{res}(f, g, y) \in \mathbb{Q}[x]$ and especially their multiplicities. Remember our overall assumptions that the polynomials f and g are generally aligned and squarefree and that they have a disjoint factorization. We assume that this is valid in the whole chapter. Therefore we know that $a \in \mathbb{C}$ is a root of X if and only if there exists a complex number b with $f(a, b) = g(a, b) = 0$. Under what circumstance does the factor $(x - a)$ of X has a multiplicity greater than 1?

Without loss of generality let us in the following assume $(a, b) = (0, 0)$. According to Corollary 2.20 this is not a restriction, because the multiplicities of the roots of X are invariant under translation of f and g: a translation of the two curves in x-direction only causes the same translation of the roots of the resultant. A translation in y-direction keeps the resultant unchanged.

Let f and g be of total degree n and m, respectively:

$$\begin{aligned} f &= f_n y^n + f_{n-1} y^{n-1} + \cdots + f_1 y + f_0, \quad f_n \not\equiv 0 \\ g &= g_m y^m + g_{m-1} y^{m-1} + \cdots + g_1 y + g_0, \quad g_m \not\equiv 0. \end{aligned}$$

Each f_i and g_j is a polynomial in x of degree at most $n - i$ and $m - j$. By successively taking the partial derivative of f with respect to y and substituting $y = 0$, we get the equalities

$$f_i = \frac{1}{i!} f_{y^i}(x, 0)$$

for all $0 \leq i \leq n$. By $f_{y^i} \in \mathbb{Q}[x, y]$ we denote the polynomial we obtain from f by computing i-times the partial derivative with respect to y. Substituting $y = 0$ leads to the polynomial $f_{y^i}(x,0) \in \mathbb{Q}[x]$. We know that each f_i is a polynomial in x of degree at most $n-i$ and therefore

$$f_i(x) = \sum_{j=0}^{n-i} \frac{1}{i!j!} f_{x^j y^i}(0,0) \cdot x^j.$$

The same considerations hold for the polynomial g with $\deg(g) = m$. For the sake of simplicity let in the following $f_{x^i y^j}$ denote $f_{x^i y^j}(0,0)$ and $g_{x^i y^j}$ denote $g_{x^i y^j}(0,0)$. Then especially f_0, f_1, g_0, and g_1 are of the form

$$\begin{aligned}
f_0(x) &= x^2 \cdot \underbrace{(\text{polynomial of degree} \leq n-0-2)}_{=:A_0(x)} + f_{x^1y^0}x + f_{x^0y^0} \\
&= x^2 A_0(x) + f_x x + f \\
f_1(x) &= x \cdot \underbrace{(\text{polynomial of degree} \leq n-1-1)}_{=:A_1(x)} + f_{x^0y^1} \\
&= x A_1(x) + f_y \\
g_0(x) &= x^2 B_0(x) + g_x x + g \\
g_1(x) &= x B_1(x) + g_y.
\end{aligned}$$

The resultant of f and g with respect to y is the determinant of the following $(m+n) \times (m+n)$ matrix:

$$X = \operatorname{res}(f,g,y) = \det \begin{pmatrix}
* & \dots & * & 0 & 0 \\
\vdots & & \vdots & \vdots & \vdots \\
* & \dots & * & 0 & 0 \\
* & \dots & * & x^2A_0(x) + f_x x + f & 0 \\
* & \dots & * & xA_1(x) + f_y & x^2A_0(x) + f_x x + f \\
* & \dots & * & 0 & 0 \\
\vdots & & \vdots & \vdots & \vdots \\
* & \dots & * & 0 & 0 \\
* & \dots & * & x^2B_0(x) + g_x x + g & 0 \\
* & \dots & * & xB_1(x) + g_y & x^2B_0(x) + g_x x + g
\end{pmatrix}.$$

We are interested in the multiplicity of x as a factor of the resultant. Because f and g vanish at $(0,0)$, in the last column of the resultant $f = 0 = g$ holds and we can factor out x. That means x is a factor of X. Equivalently we can say that 0 is a root of the resultant. This of course is what we expect. Under which conditions does this root has a multiplicity greater than 1?

We know that X is a polynomial of degree at most mn and x is a factor of X. That means X has the following form:

$$X = \operatorname{res}(f, g, y) = x(\alpha_1 + \alpha_2 x + \cdots + \alpha_{mn} x^{mn-1})$$

for some rational α_i, $1 \leq i \leq mn$. The resultant has a root of multiplicity greater than 1 in $x = 0$ if and only if the coefficient α_1 is equal to zero. When does this happen? Let us look what the coefficient α_1 looks like.

Lemma 4.1:

$$\alpha_1 = (-1)^n \cdot (f_x g_y - f_y g_x) \cdot \operatorname{sres}_1(f, g, y)(0)$$

Proof. We have

$$X = \operatorname{res}(f, g, y) = x \underbrace{(\alpha_1 + \alpha_2 x + \cdots + \alpha_{mn} x^{mn-1})}_{=:\, v(x)}$$

and $\alpha_1 = v(0)$. Moreover, because $f(0,0) = g(0,0) = 0$, we have

$$X = x \cdot \det \begin{pmatrix} * & \dots & * & 0 & 0 \\ \vdots & & \vdots & \vdots & \vdots \\ * & \dots & * & 0 & 0 \\ * & \dots & * & x^2 A_0(x) + f_x x & 0 \\ * & \dots & * & x A_1(x) + f_y & x A_0(x) + f_x \\ * & \dots & * & 0 & 0 \\ \vdots & & \vdots & \vdots & \vdots \\ * & \dots & * & 0 & 0 \\ * & \dots & * & x^2 B_0(x) + g_x x & 0 \\ * & \dots & * & x B_1(x) + g_y & x B_0(x) + g_x \end{pmatrix} =: x \cdot \det V.$$

We conclude $v(x) = \det V$ and therefore $\alpha_1 = \det V(0)$. By substituting $x = 0$ into V and applying the definition of the first subresultant, we obtain the statement we want to prove:

$$\begin{aligned}
\alpha_1 &= \det V(0) = \begin{pmatrix} & & & 0 & 0 \\ & A & & \vdots & \vdots \\ & & & 0 & 0 \\ & & & 0 & 0 \\ * & \dots & * & f_y & f_x \\ & & & 0 & 0 \\ & B & & \vdots & \vdots \\ & & & 0 & 0 \\ & & & 0 & 0 \\ * & \dots & * & g_y & g_x \end{pmatrix} \\
&= (-1)^n \cdot (f_x g_y - f_y g_x) \cdot \det \begin{pmatrix} A \\ B \end{pmatrix} \\
&= (-1)^n \cdot (f_x g_y - f_y g_x) \cdot \mathrm{sres}_1(f, g, y)(0).
\end{aligned}$$

□

The Lemma immediately leads to the following statement:

Corollary 4.2: Let $f, g \in \mathbb{Q}[x, y]$ be two generally aligned and squarefree polynomials with disjoint factorizations. Let furthermore $a \in \mathbb{C}$ be such that a is no common root of $\mathrm{res}(f, g, y)$ and $\mathrm{sres}_1(f, g, y)$. Then $\mathrm{res}(f, g, y)$ has a root of multiplicity ≥ 2 at $x = a$ if and only if there exists a point $b \in \mathbb{C}$ such that (a, b) is a tangential intersection point of f and g.

Under the assumption $\gcd(\mathrm{sres}_1(f, g, y), \mathrm{res}(f, g, y))(a) \neq 0$ we can conclude that the multiple root $x = a$ of the resultant is caused by a tangential intersection of f and g. The tangential intersection of course may be complex. What can we say in the other case that $a \in \mathbb{C}$ is a common root of the resultant as well as the first subresultant? Then we know that the two polynomials $f(a, y) \in \mathbb{C}[y]$ and $g(a, y) \in \mathbb{C}[y]$ have at least two common roots:

(a) Either there are at least two different common roots of $f(a, y)$ and $g(a, y)$

(b) or all common roots of $f(a, y)$ and $g(a, y)$ coincide.

In the second case (b) there exists one complex number b with $\gcd(f(a, y), g(a, y)) = (y - b)^i$ for some $i \geq 2$. Then of course $(y - b)$ is a factor of $\gcd(f_y(a, y), g_y(a, y))$. We conclude $f_y(a, b) = 0$ and $g_y(a, b) = 0$ and because of that $f(a, b) = g(a, b) =$

$(f_x g_y - f_y g_x)(a,b) = 0$. That means also in this case (a,b) is a tangential intersection point of f and g.

What we would like to do in advance is excluding case (a) that two distinct intersection points of f and g share the same x-coordinate.

Definition 4.3: We say that two polynomials f and g are in *general relation with respect to* x, if they have no two common roots with the same x-value. If two polynomials f and g are in general relation with respect to x as well as with respect to y, they are in *general relation*.
We call a pair of polynomials f and g *well-behaved*, if they are generally aligned and squarefree and if they have disjoint factorizations and are in general relation.

With this notation we obtain our desired characterization of multiple roots of the resultant:

Theorem 4.4: Let f and g be two well-behaved polynomials. Then *every* multiple root of $X = \text{res}(f, g, y)$ is in 1-1 correspondence to one tangential intersection point of the curves defined by f and g.

Like general alignment, squarefreeness, and disjoint factorization, also general relation can always be realized and we describe the corresponding algorithm in Chapter 7. In the following we assume that each two polynomials f and g we consider during our computation are well-behaved.

4.1.2 Simple box hit counting

Now let us have a look at how we can use this result in order to determine transversal intersections. From Theorem 4.4 it easily follows:

Corollary 4.5: Let f and g be two well-behaved polynomials and (a,b) an intersection point of the curves defined by f and g. This point is a transversal intersection point of f and g if and only if $(x-a)$ divides $\text{res}(f, g, y)$ and $(x-a)^2$ does not.

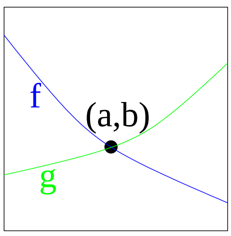

Figure 4.1: A transversal intersection

Transversal intersections are exactly the ones that cause simple roots of the resultants. We have shown in the Chapter 3 that transversal intersections can be solved with the help of simple box hit counting. So here is our criterion for how to distinguish boxes that can be solved by simple box hit counting from the ones for which this tool is not suitable. The only thing we have to do is factoring the resultant $X = \mathrm{res}(f, g, y)$ of two curves f and g into one polynomial u_1 containing all simple roots and one polynomial u_2 containing all multiple roots: $X = u_1 \cdot u_2$. The same has to be done for $Y = \mathrm{res}(f, g, y)$: $Y = v_1 \cdot v_2$. We will explain in Section 4.3 how to achieve this. The boxes defined by the real roots of u_1 and v_1 can be solved by simple box hit counting. This answers our first question.

4.2 Non-singular tangential intersections

In the last section we saw that the resultant of two bivariate well-behaved polynomials f and g has a root of multiplicity ≥ 2 at $x = a$ if and only if there exists a point $b \in \mathbb{C}$ such that f and g have a tangential intersection in (a, b). We concluded that a simple root of the resultant originates from a transversal intersection point of f and g. Boxes defined by simple roots of the resultants on the x- and y-axis can be solved with simple box hit counting.

Next we have to consider boxes that are defined by the intervals of two multiple roots. We know that both intervals are caused by tangential intersection points of f and g. What can we say about tangential intersection points? First of all, there are two situations that cause (a, b) to be a tangential intersection point:

1. either both curves have the same uniquely defined tangential line in (a, b)
2. or the intersection point (a, b) is a singular point of f or g.

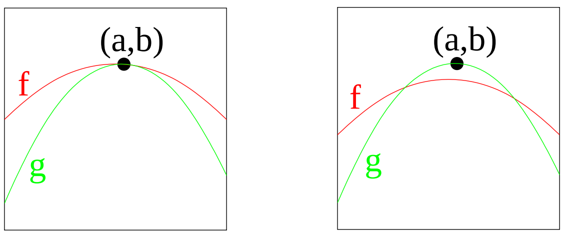

Figure 4.2: Non-singular tangential intersection and a slight move of the red curve f

For a real non-singular tangential intersection point it is easy to visualize informally that it causes a root of multiplicity ≥ 2 of the resultant. If f and g touch with the same tangent in the real point (a, b), as in the left box of Figure 4.2, then moving the red curve f in negative y-direction causes two intersection points of f and g, see the right box of Figure 4.2. The two intersection points are responsible for two roots of the resultant. Now by continuously moving the curve back to its starting position, the two roots of the resultant on the x-axis become closer and closer. The resultant is, as a determinant of polynomial coefficients, continuous and because of that we end up with a double root of the resultant.

4.2.1 The Jacobi curve

We are interested in tangential intersection points of two curves f and g. By definition a point (a, b) is a tangential intersection of f and g if and only if (a, b) is a root of the polynomial $f_x g_y - f_y g_x$. This polynomial and the curve it defines will play an important role in our future investigations. Therefore we will give it a name:

Definition 4.6: Let $f \in \mathbb{Q}[x, y]$ and $g \in \mathbb{Q}[x, y]$ be two bivariate polynomials. We define a third polynomial $h \in \mathbb{Q}[x, y]$ by

$$h := f_x g_y - f_y g_x.$$

The set of real roots of this polynomial h we call *Jacobi curve* of f and g.

We remark that the algebraic degree of h is bounded from above by $\deg(f) + \deg(g) - 2$. With the help of the Jacobi curve we reformulate:

Corollary 4.7: Let $f, g \in \mathbb{Q}[x, y]$ be well-behaved polynomials. The point $a \in \mathbb{C}$ is a root of multiplicity ≥ 2 of the resultant $\mathrm{res}(f, g, y)$ if and only if there exists a number $b \in \mathbb{C}$ such that (a, b) is a common root of f, g, and $h = f_x g_y - f_y g_x$.

The Jacobi curve cuts exactly through all tangential intersection points of f and g. We will see that this fact leads to a new test for *non-singular* tangential intersections inside a given box. The drawback of this test will be that it works only if we know in advance that the tested box contains no singular point of f or g. So in the following we assume that there exists an oracle that works in the following way:

input: Two curves f and g, a coordinate axis x or y, and an isolating interval $[a, b]$ of a multiple root of the resultant on this axis,

output: Answer to the question, whether the root inside $[a, b]$ originates from a *non-singular* tangential intersection of f and g.

Next we will show how to test a box that is defined by two intervals for which the oracle gives a positive answer. These considerations are valid for general curves. Singular points and the way the oracle works for curves that originate from quadric intersection are treated in detail in the next chapters.

We promised that the Jacobi curve h will help us to detect non-singular tangential intersections. What can we say about the behavior of h in the non-singular tangential intersection point (a, b) of f and g? There are two possibilities:

1. Either h cuts transversally through f
2. or h and f have a tangential intersection.

This characterization of h is made with respect to f. It is easy to see that we can exchange f by g:

Corollary 4.8: Let (a, b) be a non-singular tangential intersection point of f and g. The Jacobi curve $h = f_x g_y - f_y g_x$ cuts f transversally if and only if it cuts g transversally.

Proof. By assumption, there exists a number $0 \neq c \in \mathbb{Q}$ with $\nabla f(a, b) = c \cdot \nabla g(a, b)$. The Jacobi curve h cuts f transversally if and only if $\nabla f(a, b)$ and $\nabla h(a, b)$ are not linearly dependent or, equivalently, if $c \cdot \nabla f(a, b) = \nabla g(a, b)$ and $\nabla h(a, b)$ are not linearly dependent. □

Do tangential intersections of f, g, and h (as shown in the first picture of Figure 4.3) differ form transversal intersection of f and h and g and h (shown in the second picture of Figure 4.3) in the way that they cause different multiplicities in $\text{res}(f, g, y)$? The following Theorem gives the answer:

Figure 4.3: The Jacobi curve h either cuts tangentially or transversally through a non-singular tangential intersection point of f and g.

Theorem 4.9: Let $f, g \in \mathbb{Q}[x, y]$ be two well-behaved bivariate polynomials, the planar curves of which have a non-singular tangential intersection in the point (a, b). Then the Jacobi curve h intersects f as well as g transversally in (a, b), or a is a root of multiplicity ≥ 3 of $\text{res}(f, g, y)$.

The Theorem immediately leads to the statement that will serve as a basis for testing boxes that are defined by roots of multiplicities exactly 2 and candidates for non-singular tangential intersections:

Corollary 4.10: Let $f, g \in \mathbb{Q}[x, y]$ be two well-behaved bivariate polynomials the planar curves of which have a non-singular tangential intersection in the point (a, b). If $\mathrm{res}(f, g, y)$ has a root of multiplicity exactly 2 in $x = a$, then the Jacobi curve cuts transversally through f and g in (a, b).

Proof. (of Theorem 4.9) We again assume without loss of generality that $(a, b) = (0, 0)$. Further let f be a polynomial of total degree n and g be a polynomial of total degree m. Remember our remarks on resultants in the last section:

$$\begin{aligned} f(x,y) &= f_n(x)y^n + f_{n-1}(x)y^{n-1} + \cdots + f_1(x)y + f_0(x) \\ g(x,y) &= g_m(x)y^n + g_{m-1}(x)y^{m-1} + \cdots + g_1(x)y + g_0(x) \end{aligned}$$

with

$$\begin{aligned} f_i(x) &= \sum_{j=0}^{n-i} \frac{1}{i!j!} f_{x^j y^i} x^j, \qquad 0 \le i \le n \\ g_i(x) &= \sum_{j=0}^{m-i} \frac{1}{i!j!} g_{x^j y^i} x^j, \qquad 0 \le i \le m \end{aligned}$$

where $f_{x^i y^j}$ denotes the rational number $f_{x^i y^j}(0, 0)$ for all indices i and j and analogously $g_{x^i y^j} = g_{x^i y^j}(0, 0)$. Let us look at the resultant $\mathrm{res}(f, g, y)$ of f and g. We only write down the important parts of the $(m + n) \times (m + n)$ matrix we need for our proof:

$$\mathrm{res}(f,g,y) = \det \begin{pmatrix} & \cdots\cdots\cdots & & \\ \ldots & f_0(x) & 0 & 0 \\ \ldots & f_1(x) & f_0(x) & 0 \\ \ldots & f_2(x) & f_1(x) & f_0(x) \\ & \cdots\cdots\cdots & & \\ \ldots & g_0(x) & 0 & 0 \\ \ldots & g_1(x) & g_0(x) & 0 \\ \ldots & g_2(x) & g_1(x) & g_0(x) \end{pmatrix}$$

$$
\begin{aligned}
&= \quad x \cdot \det \begin{pmatrix}
\multicolumn{4}{c}{\cdots\cdots\cdots\cdots\cdots\cdots\cdots\cdots\cdots\cdots\cdots\cdots\cdots\cdots} \\
\ldots & x(*) & 0 & 0 \\
\ldots & x(*) + f_y & x^2(*) + f_x x & 0 \\
\ldots & x(*) + \frac{1}{2!} f_{yy} & x^2(*) + f_{xy} x + f_y & x^2(*) + \frac{1}{2!} f_{xx} x + f_x \\
\multicolumn{4}{c}{\cdots\cdots\cdots\cdots\cdots\cdots\cdots\cdots\cdots\cdots\cdots\cdots\cdots\cdots} \\
\ldots & x(*) & 0 & 0 \\
\ldots & x(*) + g_y & x^2(*) + g_x x & 0 \\
\ldots & x(*) + \frac{1}{2!} g_{yy} & x^2(*) + g_{xy} x + g_y & x^2(*) + \frac{1}{2!} g_{xx} x + g_x
\end{pmatrix} \\
&=: \quad x \cdot \det V \\
&= \quad x \cdot (\alpha_1 + \alpha_2 x + \cdots + \alpha_{mn} x^{mn-1}).
\end{aligned}
$$

Note that all other entries in the last three columns of the determinant are zero. We know from the previous chapter that $\alpha_1 = 0$, because f and g intersect tangentially in $(0,0)$. It remains to show that $\alpha_2 = 0$ if h does not intersect f and g transversally. If h does not intersect f and g transversally, then all three gradient vectors

$$
\begin{pmatrix} f_x(0,0) \\ f_y(0,0) \end{pmatrix} = \begin{pmatrix} f_x \\ f_y \end{pmatrix}, \quad \begin{pmatrix} g_x(0,0) \\ g_y(0,0) \end{pmatrix} = \begin{pmatrix} g_x \\ g_y \end{pmatrix}, \quad \begin{pmatrix} h_x(0,0) \\ h_y(0,0) \end{pmatrix} := \begin{pmatrix} h_x \\ h_y \end{pmatrix}
$$

are linearly dependent. With

$$
\begin{aligned}
h_x &= (f_x g_y - f_y g_x)_x = f_{xx} g_y + f_x g_{xy} - f_{xy} g_x - f_y g_{xx} \\
h_y &= (f_x g_y - f_y g_x)_y = f_{xy} g_y + f_x g_{yy} - f_{yy} g_x - f_y g_{xy}
\end{aligned}
$$

we obtain the three properties

$$
\begin{aligned}
1) \quad 0 &= \det \begin{pmatrix} f_x & g_x \\ f_y & g_y \end{pmatrix} = f_x g_y - f_y g_x \\
2) \quad 0 &= \det \begin{pmatrix} h_x & f_x \\ h_y & f_y \end{pmatrix} = h_x f_y - h_y f_x \\
&= f_y(f_{xx} g_y - f_y g_{xx}) + 2 f_y (f_x g_{xy} - f_{xy} g_x) - f_x (f_x g_{yy} - f_{yy} g_x) \\
3) \quad 0 &= \det \begin{pmatrix} h_x & g_x \\ h_y & g_y \end{pmatrix} = h_x g_y - h_y g_x \\
&= g_y(f_{xx} g_y - f_y g_{xx}) + 2 g_y (f_x g_{xy} - f_{xy} g_x) - g_x (f_x g_{yy} - f_{yy} g_x).
\end{aligned}
$$

We will show that under these three conditions $\alpha_2 = 0$ holds. We have $\det V \in \mathbb{Q}[x]$ and we know that $\alpha_2 = (\det V)'(0)$. Let V_i be the matrix we obtain by taking the derivative of each polynomial entry in the i-th column of V. It is easy to see that

$$(\det V)' = \sum_{i=1}^{n+m} \det V_i$$

and we conclude

$$\alpha_2 = \sum_{i=1}^{n+m} \det V_i(0).$$

Let us first have a look on $\det V_i(0)$ for $i = 1, \ldots, n+m-2$. For such a V_i we take the derivative of a column of V that is not one of the last two. For all these determinants there exist submatrices A_i and B_i with

$$\begin{aligned}(\det V_i)(0) &= \det \begin{pmatrix} & & & 0 & 0 \\ & A_i & & \vdots & \vdots \\ & & & 0 & 0 \\ * & \cdots & * & f_y & f_x \\ & & & 0 & 0 \\ & B_i & & \vdots & \vdots \\ & & & 0 & 0 \\ * & \cdots & * & g_y & g_x \end{pmatrix} \\ &= (-1)^n \cdot (f_x g_y - f_y g_x) \cdot \det \begin{pmatrix} A_i \\ B_i \end{pmatrix} \\ &= 0.\end{aligned}$$

by our assumption 1). It follows

$$\begin{aligned}\alpha_2 &= \det V_{m+n-1}(0) + \det V_{m+n}(0) \\ &= \det \begin{pmatrix} \multicolumn{4}{c}{\cdots\cdots\cdots\cdots} \\ \cdots & f_y & f_x & 0 \\ \cdots & \frac{1}{2!} f_{yy} & f_{xy} & f_x \\ \multicolumn{4}{c}{\cdots\cdots\cdots\cdots} \\ \cdots & g_y & g_x & 0 \\ \cdots & \frac{1}{2!} g_{yy} & g_{xy} & g_x \end{pmatrix} + \det \begin{pmatrix} \multicolumn{4}{c}{\cdots\cdots\cdots\cdots} \\ \cdots & f_y & 0 & 0 \\ \cdots & \frac{1}{2!} f_{yy} & f_y & \frac{1}{2!} f_{xx} \\ \multicolumn{4}{c}{\cdots\cdots\cdots\cdots} \\ \cdots & g_y & 0 & 0 \\ \cdots & \frac{1}{2!} g_{yy} & g_y & \frac{1}{2!} g_{xx} \end{pmatrix}\end{aligned}$$

with all other entries in the last 3 columns being 0. Let us have a look at these last 3 columns of V_{m+n-1} and V_{m+n}. For an $n \times n$ matrix M let $M(i,j,k)$ denote the 3×3 submatrix the entries of which are taken from the last 3 columns and the rows i, j, and k of M. If the determinant of each of these submatrices is equal to zero, that means if for each triple (i,j,k) with $1 \leq i < j < k \leq n$ we have $\det M(i,j,k) = 0$, then we can easily conclude $\det M = 0$. The two matrices V_{m+n-1} and V_{m+n} are of the same size and all columns are identical, except the last 2. So a similar argumentation about developing the two determinants with respect to the last 3 columns leads to the following: If $\det V_{m+n-1}(i,j,k) + \det V_{m+n}(i,j,k) = 0$ for all $1 \leq i < j < k \leq n+m$, then $\alpha_2 = \det V_{m+n-1} + \det V_{m+n} = 0$. In order to finish the proof it remains to show

$$\det V_{m+n-1}(i,j,k) + \det V_{m+n}(i,j,k) = 0 \quad \text{for all } 1 \leq i < j < k \leq n+m.$$

Remember our properties 1), 2), and 3) we made before.

1. $\{i,j,k\} \not\subset \{m-1, m, n+m-1, n+m\}$
   In this case we know that both matrices V_{m+n-1} and V_{m+n} have one row with only zero entries.

2. $(i,j,k) = (m-1, m, n+m-1)$

$$\begin{aligned}
&\det \begin{pmatrix} f_y & f_x & 0 \\ \frac{1}{2}f_{yy} & f_{xy} & f_x \\ g_y & g_x & 0 \end{pmatrix} + \det \begin{pmatrix} f_y & 0 & 0 \\ \frac{1}{2}f_{yy} & f_y & \frac{1}{2}f_{xx} \\ g_y & 0 & 0 \end{pmatrix} \\
&= -f_x(f_y g_x - f_x g_y) - \frac{1}{2} f_{xx}(f_y \cdot 0 - 0 \cdot g_y) \\
&\overset{1)}{=} 0
\end{aligned}$$

3. $(i,j,k) = (m-1, m, n+m)$

$$\begin{aligned}
&\det \begin{pmatrix} f_y & f_x & 0 \\ \frac{1}{2}f_{yy} & f_{xy} & f_x \\ \frac{1}{2}g_{yy} & g_{xy} & g_x \end{pmatrix} + \det \begin{pmatrix} f_y & 0 & 0 \\ \frac{1}{2}f_{yy} & f_y & \frac{1}{2}f_{xx} \\ \frac{1}{2}g_{yy} & g_y & \frac{1}{2}g_{xx} \end{pmatrix} \\
&= f_y((f_{xy} g_x - f_x g_{xy}) - \frac{1}{2} f_x(f_{yy} g_x - f_x g_{yy}) + \frac{1}{2} f_y(f_y g_{xx} - f_{xx} g_y) \\
&= -\frac{1}{2}\left(f_y(f_{xx} g_y - f_y g_{xx}) + 2 f_y(f_x g_{xy} - f_{xy} g_x) - f_x(f_x g_{yy} - f_{yy} g_x)\right) \\
&\overset{2)}{=} 0
\end{aligned}$$

4. $(i,j,k) = (m-1, n+m-1, n+m)$

$$\det\begin{pmatrix} f_y & f_x & 0 \\ g_y & g_x & 0 \\ \frac{1}{2}g_{yy} & g_{xy} & g_x \end{pmatrix} + \det\begin{pmatrix} f_y & 0 & 0 \\ g_y & 0 & 0 \\ \frac{1}{2}g_{yy} & g_y & \frac{1}{2}g_{xx} \end{pmatrix}$$

$$\begin{aligned} &= g_x(f_y g_x - f_x g_y) + 0 \\ &\overset{1)}{=} 0 \end{aligned}$$

5. $(i,j,k) = (m, n+m-1, n+m)$

$$\det\begin{pmatrix} \frac{1}{2}f_{yy} & f_{xy} & f_x \\ g_y & g_x & 0 \\ \frac{1}{2}g_{yy} & g_{xy} & g_x \end{pmatrix} + \det\begin{pmatrix} \frac{1}{2}f_{yy} & f_y & \frac{1}{2}f_{xx} \\ g_y & 0 & 0 \\ \frac{1}{2}g_{yy} & g_y & \frac{1}{2}g_{xx} \end{pmatrix}$$

$$\begin{aligned} &= -g_y(f_{xy}g_x - f_x g_{xy}) + \frac{1}{2}g_x(f_{yy}g_x - f_x g_{yy}) - \frac{1}{2}g_y(f_y g_{xx} - f_{xx}g_y) \\ &= \frac{1}{2}\left(g_y(f_{xx}g_y - f_y g_{xx}) + 2g_y(f_x g_{xy} - f_{xy}g_x) - g_x(f_x g_{yy} - f_{yy}g_x)\right) \\ &\overset{3)}{=} 0 \end{aligned}$$

□

For illustration have a look at the red silhouettecurve f and the green cutcurve g in Figure 6.4. The resultant with respect to y leads to the polynomial

$$\mathrm{res}(f,g,y) = (25278025x^4 - 13063152x^3 + 707600x^2 - 252960x + 17120)^2.$$

The polynomial X has two real roots, each of multiplicity 2. The multiplicities arise from the two non-singular tangential intersections at the points marked with circles. The last theorem implies that the Jacobi curve cuts the red and the green curve transversally in both points. So we additionally consider the Jacobi curve h. This is the blue one consisting of four connected components. And indeed, the Jacobi curve has transversal intersections with f as well as with g in both marked points.

Based on this result, the Jacobi curve h leads to a new test for tangential intersection points.

4.2.2 Extended box hit counting

The idea is the following: In order to determine the intersection points of two curves f and g, we partially factor their resultant $X = \mathrm{res}(f,g,y)$ into one polynomial u_1

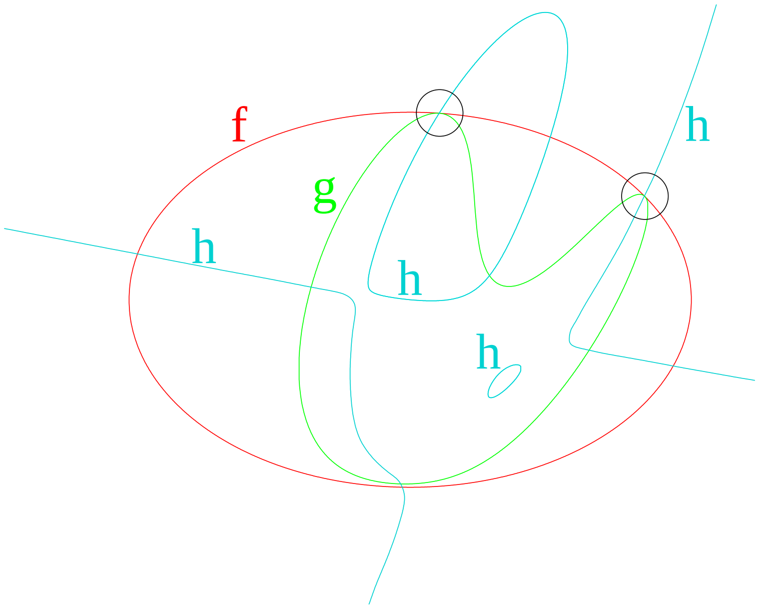

Figure 4.4: Introduce the Jacobi curve in order to solve simple tangential intersections

containing all simple roots, one polynomial u_2 containing all double roots, and one polynomial u_3 containing the rest: $X = u_1 \cdot u_2^2 \cdot u_3$. The same has to be done for $Y = \mathrm{res}(f, g, y)$: $Y = v_1 \cdot v_2^2 \cdot v_3$. For the factorization we again refer to Section 4.3. Then the boxes defined by the real roots of u_1 and v_1 can be solved by simple box hit counting. For the boxes defined by u_2 and v_2 we know that, if they contain an intersection point, then it is a tangential one. If moreover the oracle says that there is no singular point inside the boxes, we can test with simple box hit counting whether f and h and whether g and h intersect transversally in the same box. If both answers are positive, we would like to conclude that there is a non-singular tangential intersection inside the box.

In order to get a correct result, we first have to make the box small enough to guarantee that there is exactly one intersection point between f, g, and h inside the box. See for example the left upper box in Figure 4.5. In this box the red curve f and the green curve g have no tangential intersection point. If we simply test f and the blue Jacobi curve h and the curves g and h for transversal intersections, then both answers are positive. The Jacobi curve intersects transversally through f as well as through g inside the box. So applying simple box hit counting to f and h and to g and h and combining the two results would lead to a wrong positive answer. In the right upper box f and g have a tangential intersection through which h transversally cuts through. But the points of f and h on the boundary of the box do not alternate. The same holds for the points of g and h. So simple box hit counting does not detect the transversal intersection points.

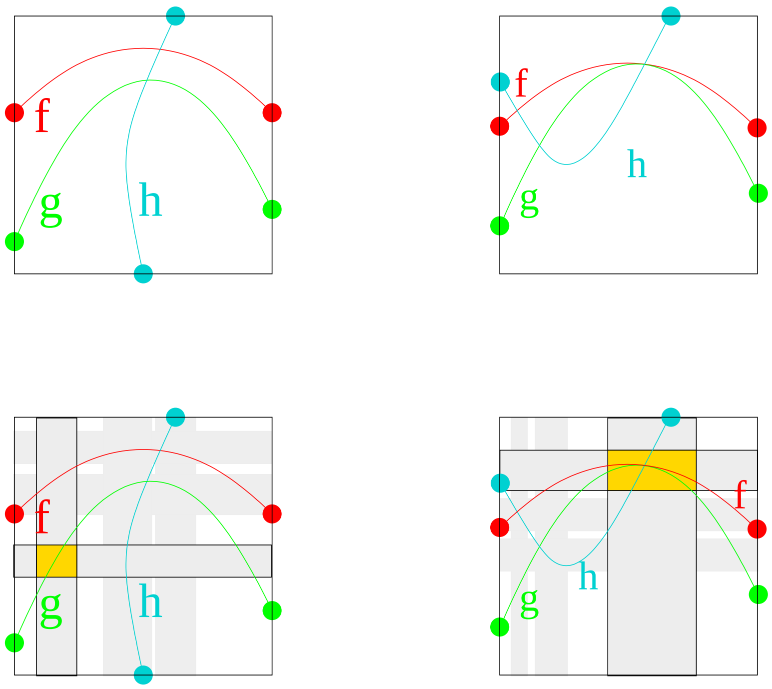

Figure 4.5: Eventually the box has to be shrunken

The problem in both cases is that there is more than 1 intersection point between f, g, and h inside the box. The box is defined by an isolating interval of a root $\alpha \in \mathbb{R}$ of $X = \text{res}(f, g, y)$ and a root $\beta \in \mathbb{R}$ of $Y = \text{res}(f, g, x)$. To avoid multiple intersection points between f, g, and h inside the box, we additionally compute $X_1 := \text{res}(f, h, y)$ and $X_2 := \text{res}(g, h, y)$. Then we separate pairwise the roots of X, X_1, and X_2 as described in Chapter 2. The root α now potentially has a smaller isolating interval than before. We do the same with the resultants Y, $Y_1 := \text{res}(f, h, x)$, and $Y_2 := \text{res}(g, h, x)$ on the y-axis and determine the isolating interval of β. The box that is defined by these two new subintervals contains at most one intersection point between f, g, and h.

In our example in Figure 4.3 the pairwise root separation leads to the grey horizontal and vertical stripes shown in the lower two boxes. The ones containing the root of $X = \text{res}(f, g, y)$ and $Y = \text{res}(f, g, x)$, respectively, are bordered in black leading to the new yellow box with which we can proceed. Notice that the box we consider is only a small fraction of the arrangement of the two curves. So possibly the roots α and β arise from two different tangential intersection points of f and g leading to a yellow box that contains no intersection point of f and g.

So here is our test:

Extended box hit counting algorithm:

```
begin
  if (# different intersection points f,g,h >1)
    make box small enough
  if (  simple box hit counting (f,h) = 1
      & simple box hit counting (g,h) = 1 )
    output: 1   // tangential intersection
  else
    output: 0
end
```

We have shown that for a non-singular tangential intersection of f and g the Jacobi curve cuts transversally through f if and only if it cuts transversally through g. In the extended box hit counting algorithm we nevertheless have to test both intersections, because we do not know whether f and g intersect at all.

4.2.3 Generalization of the Jacobi curve

Until now, we have two results on how the roots of the resultant $\operatorname{res}(f, g, y)$ indicate the kind of intersections f and g have in the plane. If there is a root of multiplicity 1, then we know that f and g have a transversal intersection. If the multiplicity is at least 2, we know that both curves have a tangential intersection. In case of non-singular points, multiplicity exactly 2 tells us that the Jacobi curve h cuts transversally through f and g. Can we generalize this concept of the Jacobi curve in the way that for each kind of non-singular intersection point of two curves f and g we can find a third curve that cuts f as well as g transversally?

For real intersection points, and this are in particular the ones we are interested in, we will give a positive answer to this question with the help of real Analysis, the Theorem of Implicit Functions. Of course, such considerations are only necessary if f and g intersect the boundary of the box not in alternating order. Otherwise simple box hit counting is a suitable tool.

Let $(a, b) \in \mathbb{R}^2$ be a real intersection point of $f, g \in \mathbb{Q}[x, y]$ that neither is a singular point of f nor a singular point of g. We will iteratively define a sequence of polynomials $\tilde{h}_1, \tilde{h}_2, \tilde{h}_3, \ldots$ such that $\tilde{h}_k$ cuts transversally through f in (a, b) for some index k. If f and g are well-behaved, the index k is equal to the degree of a as a root of $\operatorname{res}(f, g, y)$. This last observation will easily lead to a generalization of extended box hit counting to arbitrary non-singular tangential intersection points.

Definition 4.11: Let f and g be two planar curves. We define *generalized Jacobi curves* the following way:

$$\begin{aligned}\tilde{h}_1 &:= g \\ \tilde{h}_{i+1} &:= (\tilde{h}_i)_x f_y - (\tilde{h}_i)_y f_x.\end{aligned}$$

Here is our main Theorem that will provide all necessary mathematical tools:

Theorem 4.12: Let (a,b) be an intersection point of two different curves f and g, that means $f(a,b) = 0 = g(a,b)$. Furthermore let $(f_x, f_y)(a,b) \neq (0,0)$ and $(g_x, g_y)(a,b) \neq (0,0)$. There exists an index $k \geq 1$ such that $\tilde{h}_k$ cuts transversally through f in (a,b).

Proof. In the case g cuts through f in the point (a,b), especially if (a,b) is a transversal intersection point of f and g, this is of course true for $h_1 = g$. So assume in the following that $(g_x f_y - g_y f_x)(a,b) = \tilde{h}_2(a,b) = 0$. From now on we will only consider the polynomials $\tilde{h}_i$ with $i \geq 2$.

Our assumption is that (a,b) is a non-singular point of f: $(f_x, f_y)(a,b) \neq 0$. We only consider the case $f_y(a,b) \neq 0$. In the case $f_x(a,b) \neq 0$ and $f_y(a,b) = 0$ we would proceed the same way as described in the following by just exchanging the two variables x and y. The property $f_y(a,b) \neq 0$ leads to $(\frac{f_x}{f_y} g_y)(a,b) = g_x(a,b)$ and because $(g_x, g_y)(a,b) \neq (0,0)$ we conclude $g_y(a,b) \neq 0$. From the Theorem of Implicit Functions we derive that there are real open intervals $I_x, I_y \subset \mathbb{R}$ with $(a,b) \in I_x \times I_y$ such that

1. $f_y(x_0, y_0) \neq 0$ and $g_y(x_0, y_0) \neq 0$ for all $(x_0, y_0) \in I_x \times I_y$,
2. there exists a continuous function $F : I_x \to I_y$ with the two properties
   (a) $f(x, F(x)) = 0$ for all $x \in I_x$
   (b) $(x,y) \in I_x \times I_y$ with $f(x,y) = 0$ leads to $y = F(x)$,
3. and there exists a continuous function $G : I_x \to I_y$ with
   (a) $g(x, G(x)) = 0$ for all $x \in I_x$
   (b) $(x,y) \in I_x \times I_y$ with $g(x,y) = 0$ leads to $y = G(x)$.

Locally around the point (a,b) the curve defined by the polynomial f is equal to the graph of the function F. The same holds for g and G. Especially we have $b = F(a) = G(a)$. Moreover, the Theorem of Implicit Holomorphic Functions implies that F as well as G are holomorphic and thus developable in a Taylor series around the point (a,b) [34].

In the following we will sometimes consider the functions $h_i : I_x \times I_y \to \mathbb{R}$, $i \geq 2$, with

$$h_2 := \frac{g_x}{g_y} - \frac{f_x}{f_y} = \frac{\tilde{h}_2}{g_y f_y}$$
$$h_{i+1} := (h_i)_x - (h_i)_y \cdot \frac{f_x}{f_y}$$

instead of the polynomials $\tilde{h}_i$. Each h_i is well defined and for $(x, y) \in I_x \times I_y$ we have

$$\tilde{h}_i(x, y) = 0 \quad \Leftrightarrow \quad h_i(x, y) = 0.$$

The kernel of h_i is equal to the set of roots of $\tilde{h}_i$ inside the interval $I_x \times I_y$. So locally around the point (a, b) the polynomial $\tilde{h}_i$ and the function h_i define the same curve.

Let us assume we know the following proposition:

Let $k \geq 1$. If $F^{(i)}(a) = G^{(i)}(a)$ for all $0 \leq i \leq k-1$, then $h_{k+1}(a, b) = F^{(k)}(a) - G^{(k)}(a)$.

We know that the two polynomials f and g are different. That means the Taylor series of F and G differ in some term. There is an index k such that $F^{(i)}(a) = G^{(i)}(a)$ for all $0 \leq i \leq k-1$ and $F^{(k)}(a) \neq G^{(k)}(a)$. According to the statement we have $h_k(a, b) = F^{(k-1)}(a) - G^{(k-1)}(a) = 0$, that means $\tilde{h}_k$ intersects f in (a, b). This intersection is transversal if and only if

$$((\tilde{h}_k)_x f_y - (\tilde{h}_k)_y f_x)(a, b) = \tilde{h}_{k+1}(a, b) \neq 0.$$

This follows easily from $h_{k+1}(a, b) = F^{(k)}(a) - G^{(k)}(a) \neq 0$. □

It remains to state and prove the proposition:

Proposition 4.13: Let $k \geq 1$. If $F^{(i)}(a) = G^{(i)}(a)$ for all $0 \leq i \leq k-1$, then $h_{k+1}(a, b) = F^{(k)}(a) - G^{(k)}(a)$.

Proof. For each $i \geq 0$ we define a function $H_i : I_x \to \mathbb{R}$ by

$$H_i(x) := h_i(x, F(x)).$$

For $x = a$ we derive $H_i(a) = h_i(a, b)$. So in terms of our new function we want to prove that $H_{k+1}(a) = F^{(k)}(a) - G^{(k)}(a)$ holds if $F^{(i)}(a) = G^{(i)}(a)$ for all $0 \leq i \leq k-1$.

By definition $f(x, F(x)) : I_x \to \mathbb{R}$ and $f(x, F(x)) = 0$ for all $x \in I_x$. That means $f(x, F(x))$ is constant and therefore its derivative is equal to zero:

$$\frac{d}{dx}f(x,F(x)) = f_x(x,F(x)) + F'(x)f_y(x,F(x)) = 0.$$

We conclude

$$F'(x) = -\frac{f_x(x,F(x))}{f_y(x,F(x))}.$$

For the functions H_i the equality $H_i'(x) = H_{i+1}(x)$ holds, because

$$\begin{aligned} H_i'(x) &= (h_i)_x(x,F(x)) + F'(x)\cdot(h_i)_y(x,F(x)) \\ &= (h_i)_x(x,F(x)) - \frac{f_x(x,F(x))}{f_y(x,F(x))}\cdot(h_i)_y(x,F(x)) \\ &= h_{i+1}(x,F(x)) \\ &= H_{i+1}(x). \end{aligned}$$

Inductively we obtain $H_k(x) = H_0^{(k)}(x)$ for all $k \geq 0$ and $H_k(x) = H_2^{(k-2)}(x)$ for all $k \geq 2$. In order to prove the proposition it is sufficient to show the following: Let $k \geq 1$. If for all $0 \leq i \leq k-1$ we have $F^{(i)}(a) = G^{(i)}(a)$, then $H_2^{(k-1)}(a) = (F'-G')^{(k-1)}(a)$. We will derive this result by induction on k.

1. Let $k = 1$. Our assumption is $F(a) = G(a)$ and we have to show $H_2(a) = (F'-G')(a)$. We have

$$\begin{aligned} H_2(x) = h_2(x,F(x)) &= \frac{g_x(x,F(x))}{g_y(x,F(x))} - \frac{f_x(x,F(x))}{f_y(x,F(x))} \\ (G'-F')(x) &= \frac{g_x(x,G(x))}{g_y(x,G(x))} - \frac{f_x(x,F(x))}{f_y(x,F(x))} \end{aligned}$$

and both functions just differ in the functions that are substituted for y in $\frac{g_x(x,y)}{g_y(x,y)}$. In the equality of $H_2(x)$ we substitute $F(x)$, whereas in the one of $(G'-F')$ we substitute $G(x)$. But of course $F(a) = G(a)$ leads to

$$H_2(a) = \frac{g_x(a,F(a))}{g_y(a,F(a))} - \frac{f_x(a,F(a))}{f_y(a,F(a))} = \frac{g_x(a,G(a))}{g_y(a,G(a))} - \frac{f_x(a,F(a))}{f_y(a,F(a))} = (G'-F')(a).$$

2. Let $k > 1$. We know that $F^{(i)}(a) = G^{(i)}(a)$ for all $0 \le i \le k-1$. We again use the equations

$$\begin{aligned} H_2(x) &= \frac{g_x(x,F(x))}{g_y(x,F(x))} - \frac{f_x(x,F(x))}{f_y(x,F(x))} \\ (G'-F')(x) &= \frac{g_x(x,G(x))}{g_y(x,G(x))} - \frac{f_x(x,F(x))}{f_y(x,F(x))} \end{aligned}$$

and the fact that $H_2(x)$ and $(G'-F')$ only differ in the functions that are substituted for y in $\frac{g_x(x,y)}{g_y(x,y)}$.

By taking $(k-1)$ times the derivative of $H_2(x)$ and $(G'-F')$, we structurally obtain the same result for both functions. The only difference is that some of the terms $F^{(i)}(x)$, $0 \le i \le k-1$, in H_2 are exchanged by $G^{(i)}(x)$ in $(G'-F')$. But due to our assumption we have $F^{(i)}(a) = G^{(i)}(a)$ for all $0 \le i \le k-1$ and we obtain

$$H_2^{(k-1)}(a) = (G'-F')^{(k-1)}(a).$$

□

We have proven that for every non-singular tangential intersection of f and g, not only the ones causing a double root of the resultant like in the previous section, there exists a curve $\tilde{h}_k$ that cuts both curves transversally in this point. The index k depends on the degree of similarity of the functions that describe both polynomials in a small area around the given point. The degree of similarity is measured by the number of successive matching derivatives in this point. A useful result would be the following: If k is the multiplicity of a in the resultant $\text{res}(f,g,y)$, then $\tilde{h}_k$ cuts transversally through f in the corresponding point (a,b). We have shown this for $k=1$ and $k=2$. For arbitrary k it follows as a consequence of the previous theorem:

Corollary 4.14: Let $f, g \in \mathbb{Q}[x,y]$ be two well-behaved polynomials and let (a,b) be a non-singular intersection point of the curves defined by f and g. If k is the degree of a as a root of the resultant $\text{res}(f,g,y)$, then $\tilde{h}_k$ cuts transversally through f and g.

Proof. In Chapter 2 we have seen that the resultant of two univariate polynomials is equal to the product of the differences of their roots. So if we compute the resultant $X = \text{res}(f,g,y)$ for two bivariate polynomials f and g, the value of X for each fixed x_0 equals the product of the differences of the roots of $f(x_0,y)$ and $g(x_0,y)$.

In the previous theorem we have proven that, in a region $I_x \times I_y$ locally around the point (a,b), we can develop f and g in Taylor series $y = F(x) = \sum_{i=0}^{\infty} \frac{F^{(i)}(a)}{i!}(x-a)^i$

and $y = G(x) = \sum_{i=0}^{\infty} \frac{G^{(i)}(a)}{i!}(x-a)^i$, respectively. The index i for which $\tilde{h}_i$ cuts transversally through f and g equals the index for which $F^{(i)}(a) \neq G^{(i)}(a)$ holds the first time.

For a point $x_0 \in I_x$ we know the root $F(x_0)$ of the univariate polynomials $f(x_0, y)$ and the root $G(x_0)$ of $g(x_0, y)$. So for all $x_0 \in I_x$ the term $F(x_0) - G(x_0)$ is a factor of $X(x_0)$. We conclude that all roots of $F(x) - G(x)$ in I_x are roots of X, together with their multiplicities. By assumption $(a, F(a))$ is the only intersection point of f and g at $x = a$. That means the degree k of $(x-a)$ in X equals the degree of $(x-a)$ in $F(x) - G(x)$. We obtain our desired result that $\tilde{h}_k$ cuts transversally through f and g. □

Finally, we give a generalized extended box hit counting algorithm for determining all non-singular tangential intersections of f and g. We first apply simple box hit counting to f and g in order to test whether g cuts through f. If this test fails, we compute extended box hit counting of f, g, and h_k. Thereby k is the degree of the root that defines the box on the coordinate axes.

Generalized extended box hit counting:

```
begin
  if (simple box hit counting (f,g) = 1) return 1;
  k := deg(root);
  extended box hit counting (f,g,hk)
end
```

4.3 Partial factorization

In the last sections we used partial factorization of univariate polynomials with respect to the multiplicities of their roots. We will now describe how to realize this factorization.

4.3.1 Multiplicity factorization

When we deal with a univariate polynomial $u \in \mathbb{Q}[x]$, we can partially factor it according to the multiplicities of its roots by using multiple differentiation and division of polynomials. Without loss of generality we assume $\mathrm{ldcf}(u) = 1$. For example the polynomial

$$u(x) = x^9 - x^7 - x^5 + x^3$$

can be partially factored into three polynomials

$$\begin{aligned} u_1(x) &= x^2+1 \\ u_2(x) &= x^2-1 \\ u_3(x) &= x \end{aligned}$$

such that

$$u(x) = u_1(x)^1 \cdot u_2(x)^2 \cdot u_3(x)^3$$

and each u_i only has simple roots. This can be done by iteratively computing

$$\begin{aligned} w_1 &:= u \\ w_{i+1} &:= \gcd(w_i, w_i') \end{aligned}$$

until $\gcd(w_i, w_i')$ is a constant. Let this happen for $i = d$. In each step the multiplicity of the roots of u is decreased by 1. Therefore each polynomial w_i contains exactly all roots of u of multiplicity at least i. A root of multiplicity $j \geq i$ of u is a root of multiplicity $j - i + 1$ of w_i and vice versa. We compute the polynomials u_i that exactly contain all roots of u of multiplicity i as simple roots in order $i = d \ldots 1$:

$$u_i(x) = \frac{w_i(x)}{\prod_{j=i+1}^{d} u_j^{j-i+1}(x)}.$$

Of course, we can apply multiplicity factorization to the two polynomials $X = \mathrm{res}(f, g, y)$ and $Y = \mathrm{res}(f, g, x)$ on the x- and the y-axis. One advantage of this factorization is that afterwards we deal with polynomials of smaller degree than before. Another advantage is the one we already mentioned before, namely we know that boxes defined by the simple roots of X and Y can be solved with simple box hit counting and double roots with extended box hit counting.

4.3.2 Bi–factorization

We want to compute whether a box contains an intersection point or not. The boxes are produced by the rational interval representations of the real roots of the resultants on the x- and on the y-axis. For a clarification of our terms, we repeat and complete a notation we introduced in Chapter 3:

Definition 4.15:

1. Let $u \in \mathbb{Q}[x]$ and $v \in \mathbb{Q}[y]$. By $\mathcal{R}(u)$ we denote the set of real roots of u. By $\textsc{Grid}(u, v)$ we mean the grid $\mathcal{R}(u) \times \mathcal{R}(v)$.

2. Let $f, g \in \mathbb{Q}[x, y]$, $X = \mathrm{res}(f, g, y)$, and $Y = \mathrm{res}(f, g, x)$. We call the pair (X, Y) the *bi-resultant* of f and g.

For both resultants we can make a multiplicity factorization as described in the section before. For example look at figure 4.6. There are transversal intersections of f and g and self-intersections of f through which also g cuts through, i.e tangential intersections of f and g.

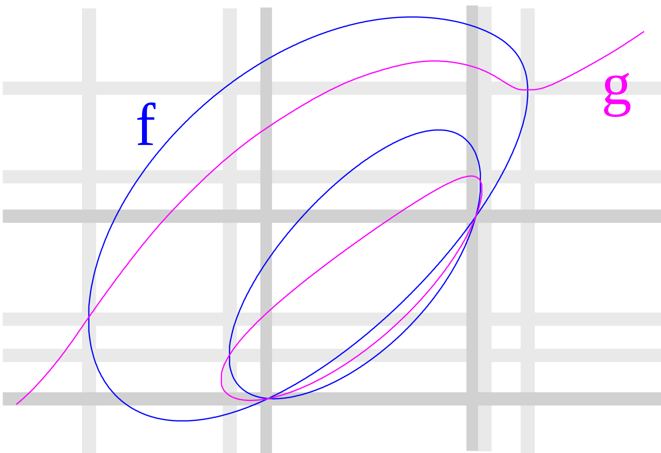

Figure 4.6: A multiplicity factorization

We can factor both resultants according to multiplicities

$$\begin{aligned} \mathrm{res}(f, g, y) &= X = u_1 u_2 \\ \mathrm{res}(f, g, x) &= Y = v_1 v_2 \end{aligned}$$

such that u_1 and v_1 contain exactly all simple roots of X and Y and u_2 and v_2 contain all roots of multiplicity at least 2. The simple roots cause the light grey stripes and the multiple roots cause the dark grey stripes. We have seen that for two well-behaved curves, transversal intersections exactly give rise to simple roots whereas tangential intersections cause roots of multiplicity at least 2. That means all transversal intersections are contained in the light grey stripes and tangential intersections are contained in the dark grey stripes. We have to test all boxes that are intersections of two light grey stripes for transversal intersections and all boxes that are intersections of two dark grey stripes for tangential intersections. We do not need to consider boxes that are intersections of light grey with dark grey stripes. For illustration see also Figure 4.7.

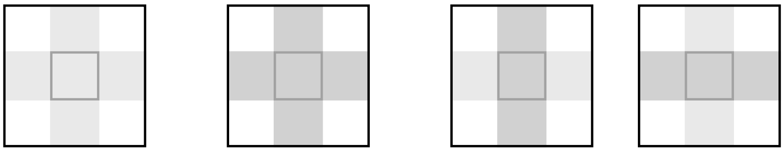

Figure 4.7: The first box is a candidate for a transversal intersection point, the second one for a tangential intersection point, and the last two ones are no candidates for event points

Of course the idea of distinguishing boxes by factoring the resultants can also be applied to other criteria than the multiplicity of the roots of the resultants. Therefore our notation is in a more general way:

Definition 4.16: Let $u_1, u_2 \in \mathbb{Q}[x]$ and $v_1, v_2 \in \mathbb{Q}[y]$. The expression $(u_1, v_1) \cdot (u_2, v_2)$ is called *bi-factorization* of the bi-resultant $(X, Y) = (\text{res}(f, g, y), \text{res}(f, g, x))$ iff

1. $X = u_1 \cdot u_2$, $Y = v_1 \cdot v_2$,
2. and all intersection points of f and g lie on $\textsc{Grid}(u_1, v_1) \cup \textsc{Grid}(u_2, v_2)$.

The pairs (u_1, v_1) and (u_2, v_2) are called *bi-factors*.

4.4 Explicit solutions

In this last section we give a criterion how to determine general intersections, including singular points, of two curves under special circumstances. Let again f and g be two bivariate polynomials. Let additionally (u, v) be a bi-factor of $(\text{res}(f, g, y), \text{res}(f, g, x))$, u and v being polynomials of degree at most 2:

$$\begin{aligned} u(x) &= a_u x^2 + b_u x + c_u \\ v(x) &= a_v y^2 + b_v y + c_v. \end{aligned}$$

We assume that some kind of bi-factorization gave us these two rational polynomials u and v. The roots $x_{1,2}$ of u and $y_{1,2}$ of v define four points (x_i, y_j), $i, j = 1, 2$, in the complex plane. We can give explicit terms for them involving only one square root per coordinate. We are only interested in real intersection points. That is why we only have to consider solutions with a non-negative term under the square root.

$$\begin{aligned}
x_{1,2} &= -\frac{1}{2a_u} \cdot (b_u \pm \sqrt{b_u^2 - 4a_u c_u})) \\
&=: -\frac{1}{2a_u} \cdot (b_u \pm \sqrt{a}) \\
y_{1,2} &= -\frac{1}{2a_v} \cdot (b_v \pm \sqrt{b_v^2 - 4a_v c_v})) \\
&=: -\frac{1}{2a_v} \cdot (b_v \pm \sqrt{b}).
\end{aligned}$$

Assume we know a bi-factor (u, v) of (X, Y) describing the two self-intersection points in Figure 4.8. If both polynomials u and v have degree 2, then we can explicitly compute their roots. The self-intersections lie on the grid points $\textsc{Grid}(u, v)$.

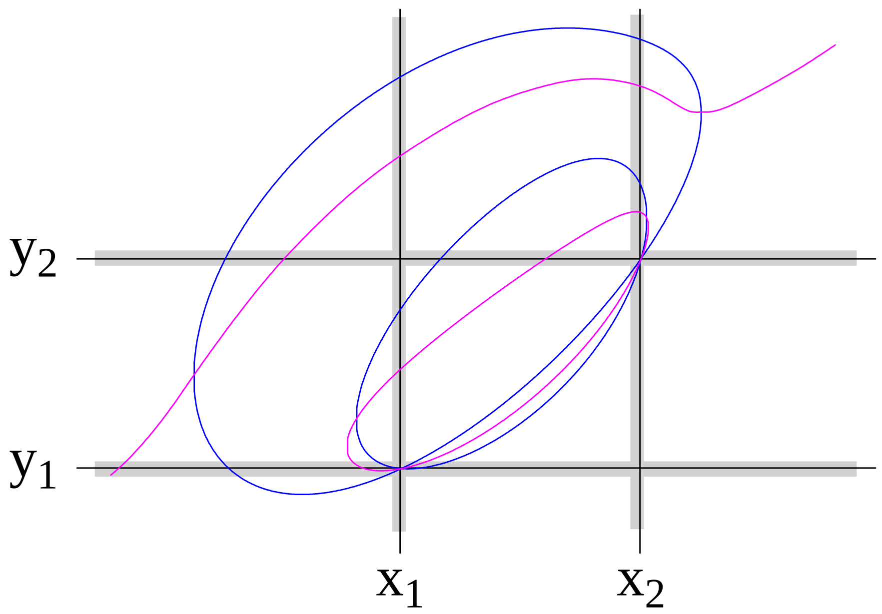

Figure 4.8: We can explicitly compute the grid points of $\textsc{Grid}(u, v)$

Substituting the terms x_i and y_j into f and g reduces the question whether f and g intersect at one of the real points to the question whether two terms involving only simple square root expressions both evaluate to zero:

$$f(x_i, y_j) \stackrel{?}{=} 0 \stackrel{?}{=} g(x_i, y_j).$$

Both tests can be made by using root separation bounds, for example realized in the LEDA real class [48].

Another way is the following: The expressions $f(x_i, y_j) \in \mathbb{R}$ and $g(x_j, y_j) \in \mathbb{R}$ we have to test for zero are of the form

$$\alpha\sqrt{a} + \beta\sqrt{b} + \gamma\sqrt{a \cdot b} + \delta.$$

with $\alpha, \beta, \gamma, \delta, a, b \in \mathbb{Q}$. Let sign : $\mathbb{R} \to \{-1, 0, 1\}$ denote the function with

$$\operatorname{sign}(a) = \begin{cases} -1 & \text{, if } a < 0 \\ 0 & \text{, if } a = 0 \\ +1 & \text{, if } a > 0. \end{cases}$$

Now the following equivalence relations hold:

$$\begin{aligned}
& \alpha\sqrt{a} + \beta\sqrt{b} + \gamma\sqrt{a \cdot b} + \delta &=& \ 0 \\
\Longleftrightarrow\ & \alpha\sqrt{a} + \beta\sqrt{b} &=& \ -\gamma\sqrt{a \cdot b} - \delta \\
\Longleftrightarrow\ & \operatorname{sign}(\alpha\sqrt{a} + \beta\sqrt{b}) &=& \ \operatorname{sign}(-\gamma\sqrt{a \cdot b} - \delta) \quad \text{and} \\
& \alpha^2 a + 2\alpha\beta\sqrt{ab} + \beta^2 b &=& \ \gamma^2 ab + 2\gamma\delta\sqrt{ab} + \delta^2 \\
\Longleftrightarrow\ & \operatorname{sign}(\alpha\sqrt{a} + \beta\sqrt{b}) &=& \ \operatorname{sign}(-\gamma\sqrt{a \cdot b} - \delta) \quad \text{and} \\
& \alpha^2 a + \beta^2 b - \gamma^2 ab - \delta^2 &=& \ (-2\alpha\beta + 2\gamma\delta)\sqrt{ab} \\
\Longleftrightarrow\ & \operatorname{sign}(\alpha\sqrt{a} + \beta\sqrt{b}) &=& \ \operatorname{sign}(-\gamma\sqrt{a \cdot b} - \delta) \quad \text{and} \\
& \operatorname{sign}(\alpha^2 a + \beta^2 b - \gamma^2 ab - \delta^2) &=& \ \operatorname{sign}(-2\alpha\beta + 2\gamma\delta) \quad \text{and} \\
& (\alpha^2 a + \beta^2 b - \gamma^2 ab - \delta^2)^2 &=& \ ab(-2\alpha\beta + 2\gamma\delta)^2
\end{aligned}$$

The last two equations only involve rational numbers and therefore are easy to test. In principle, the first one can be solved by computing the sign of a real number $\sqrt{c} + d$ for rational numbers c and d. We have

$$\operatorname{sign}(\sqrt{c} + d) = \begin{cases} +1 & \text{, if } \operatorname{sign}(d) = +1 \\ +1 & \text{, if } \operatorname{sign}(d) = -1 \text{ and } \operatorname{sign}(c - d^2) = +1 \\ 0 & \text{, if } \operatorname{sign}(d) = -1 \text{ and } \operatorname{sign}(c - d^2) = 0 \\ -1 & \text{, if } \operatorname{sign}(d) = -1 \text{ and } \operatorname{sign}(c - d^2) = -1 \end{cases}$$

and all computations can be done over the rational numbers.

Next, consider the case that the two univariate polynomials u and v that define the grid points are not polynomials over $\mathbb{Q}$, but over a field extension $\mathbb{Q}(\sqrt{\rho})$ for $\rho \in \mathbb{Q}$:

$u, v \in \mathbb{Q}(\sqrt{\rho})[x]$. Then all considerations so far also hold for these polynomials with $\alpha, \beta, \gamma, \delta, a, b, c, d \in \mathbb{Q}(\sqrt{\rho})$. The remaining tests are either comparing two numbers in $\mathbb{Q}(\sqrt{\rho})$ or comparing their signs. But of course, sorting and squaring the involved terms again reduces also this problem over $\mathbb{Q}(\sqrt{\rho})$ into a rational one.

All in all, if there are quadratic polynomials $u, v \in \mathbb{Q}$ or $u, v \in \mathbb{Q}(\sqrt{\rho})$ the roots of which define boxes in the plane, then by only using rational computations we can explicitly test each box for an intersection point. We call this method *explicit solutions*. We will see in the next two chapters that computing explicit solutions is the last missing tool for computing singular points in our arrangements of silhouettecurves and cutcurves.

Chapter 5

Singular points of cutcurves

In the last chapter we have seen how to locate non-singular intersection points of general curves in planar arrangements with box hit counting methods. For non-singular tangential intersections we need in advance the information that the box contains no singular point of one of the curves. What remains to do is

1. distinguishing in advance the boxes that potentially contain singular points from the ones that cannot contain such a point and

2. determining singular points inside the boxes.

We will attack the two problems only for the curves in our arrangements we obtain from projecting quadric intersection curves into the plane. The focus is on singular points of cutcurves, because, as we will see in the next chapter, singular points of silhouettecurves pose no problem.

5.1 The origin of singular points

We want to investigate the singular points of cutcurves. As in the previous chapters we will assume that all pairs of quadrics we consider are generally aligned and squarefree and have a disjoint factorization. Moreover we assume that all pairs of curves are well-behaved. Consider Chapter 7 for the exact definitions. Let us first have a look at the origin of cutcurves. They are a result of the projection phase, namely they are given by the resultant of two input quadrics as described in Chapter 3. We only consider input polynomials of total degree 2, because intersecting a plane with a quadric leads to a planar quadratic curve that can be handled the same way as a silhouettecurve.

So let $p \in \mathbb{Q}[x, y, z]$ and $q \in \mathbb{Q}[x, y, z]$ be two trivariate quadratic polynomials. The polynomials are of the form

$$\begin{aligned} p &= z^2 + p_1 z + p_0 \\ q &= z^2 + q_1 z + q_0 \end{aligned}$$

with $p_1, p_0, q_1, q_0 \in \mathbb{Q}[x, y]$. The polynomials are generally aligned and therefore we know that p and q have a constant non-zero coefficient of z^2. Without changing the set of roots of the polynomials we can assume that this leading coefficient is equal to 1. The resultant of p and q with respect to z is of the form

$$\begin{aligned} f = \operatorname{res}(p, q, z) &= \begin{pmatrix} 1 & p_1 & p_0 & 0 \\ 0 & 1 & p_1 & p_0 \\ 1 & q_1 & q_0 & 0 \\ 0 & 1 & q_1 & q_0 \end{pmatrix} \\ &= (p_0 q_1 - p_1 q_0) \cdot (q_1 - p_1) + (p_0 - q_0)^2 \\ &= (p_0 q_1 - p_1 q_0) \cdot \operatorname{sres}_1(p, q, z) + (p_0 - q_0)^2. \end{aligned}$$

That means each cutcurve in our planar arrangements is defined by a polynomial of this special form. We want to avoid applying box hit counting to boxes that contain a singular point. Which situations in space cause a point (a, b) to be a singular point of f?

If $(a, b) \in \mathbb{C}^2$ is a point on the cutcurve $f := \operatorname{res}(p, q, z)$, then general alignment of p and q guarantees the existence of a number $c \in \mathbb{C}$ with $p(a, b, c) = q(a, b, c) = 0$, remember Corollary 2.15. We will show next that $(0, 0)$ is a singular point of f only

- if p and q share another common root $(a, b, c') \neq (a, b, c)$
- or if p and q intersect tangentially in (a, b, c).

For illustration first have a look at Figure 5.2. The blue spatial curve running on the red ellipsoid is the intersection curve of the red and the blue quadric in our overall example. The spatial curve has well defined tangents in every point. It consists of two branches, one on the upper and one on the lower part of the red ellipsoid. The two branches are projected on top of each other causing two singular, namely self-intersection, points.

In Figure 5.2 the red and the green ellipsoid have a tangential intersection in space, that means the spatial intersection curve already has a singular point. This singular point will be projected into the plane.

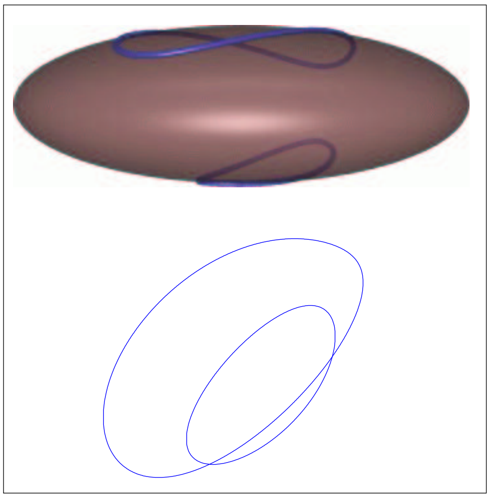

Figure 5.1: Top-bottom points result from the projection

We will show that these two events are the only ones that cause singular points. Due to this distinction we classify singular points in the plane:

Definition 5.1: Let $p, q \in \mathbb{Q}[x, y, z]$ be quadratic polynomials and $(a, b) \in \mathbb{C}^2$ be a point on the cutcurve defined by $f = \text{res}(p, q, z)$. If the two complex numbers c and c' are the roots of $p(a, b, z) \in \mathbb{C}[z]$ and we have $p(a, b, c) = q(a, b, c) = 0$ and $p(a, b, c') = q(a, b, c') = 0$, then we call (a, b) a *top-bottom point*. If (a, b) is the projection of a tangential intersection point of p and q, we call it *genuine*.

Top-bottom points are common roots of $f = \text{res}(p, q, z)$ and $\text{sres}_1(p, q, z)$ as shown in Lemma 2.21. Here is our theorem:

Theorem 5.2: Every singular point of a cutcurve which originates from two generally aligned quadrics is top-bottom or genuine.

Proof. Let (a, b) be a singular point of a cutcurve $f = \text{res}(p, q, z)$ and let (a, b, c) be the corresponding common root of $p = z^2 + p_1 z + p_0$ and $q = z^2 + q_1 z + q_0$ in space. We assume without loss of generality that $(a, b, c) = (0, 0, 0)$. As shown in Corollary 2.20, the resultant computation is invariant under translation along the z-axis. A translation of p and q along the x- or y-axis just causes the same translation of the resultant.

It suffices to show the following: If $(0, 0)$ is a singular point of f and $\text{sres}_1(p, q, z)(0, 0) \neq 0$, then p and q intersect tangentially in $(0, 0, 0)$. So we want to show that $f_x(0, 0) = f_y(0, 0) = 0$ and $\text{sres}_1(p, q, z)(0, 0) \neq 0$ imply

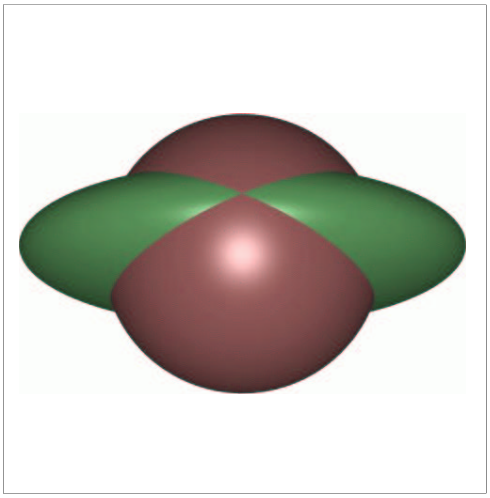

Figure 5.2: Genuine points are caused by tangential intersections in space

$$0 = (p_y q_z - p_z q_y)(0,0,0) = (p_x q_z - p_z q_x)(0,0,0) = (p_x q_y - p_y q_x)(0,0,0).$$

For the coefficients of p and q we know the equalities

$$\begin{array}{lllllll} p_0 & = & p(x,y,0) & =: & p|_{z=0} & \quad q_0 & = & q(x,y,0) & =: & q|_{z=0} \\ p_1 & = & p_z(x,y,0) & =: & p_z|_{z=0} & \quad q_1 & = & q_z(x,y,0) & =: & q_z|_{z=0}. \end{array}$$

We just computed the resultant of p and q in terms of their coefficients and conclude

$$f = \operatorname{res}(p,q,z) = (pq_z - p_z q)|_{z=0} \cdot \operatorname{sres}_1(p,q,z) + ((p-q)|_{z=0})^2.$$

For a polynomial $p \in \mathbb{C}[x,y,z]$ it is obvious that taking the partial derivative with respect to a variable $x \neq z$ and then substituting $z = 0$ is the same as first substituting $z = 0$ and then taking the partial derivative:

$$(p_x)|_{z=0} = (p|_{z=0})_x.$$

This leads to

$$\begin{aligned} f_x &= ((pq_z - p_z q)|_{z=0} \cdot \mathrm{sres}_1(p,q,z) + ((p-q)|_{z=0})^2)_x \\ &= (p_x q_z - p_z q_x)|_{z=0} \cdot \mathrm{sres}_1(p,q,z) + \\ &\quad (pq_{xz} - p_{xz} q)|_{z=0} \cdot \mathrm{sres}_1(p,q,z) + \\ &\quad (pq_z - p_z q)|_{z=0} \cdot (\mathrm{sres}_1(p,q,z))_x + \\ &\quad 2(p-q)|_{z=0} \cdot (p_x - q_x)|_{z=0}. \end{aligned}$$

Due to our assumptions we have $f_x(0,0) = 0$ and $p(0,0,0) = q(0,0,0) = 0$. This leads to

$$0 \;=\; f_x(0,0) \;=\; ((p_x q_z - p_z q_x)|_{z=0} \cdot \mathrm{sres}_1(p,q,z))(0,0).$$

Analogously, because of $0 = f_y(0,0)$, we get

$$0 \;=\; f_y(0,0) \;=\; ((p_y q_z - p_z q_y)|_{z=0} \cdot \mathrm{sres}_1(p,q,z))(0,0).$$

We assumed $\mathrm{sres}_1(p,q,z)(0,0) = (q_z - p_z)|_{z=0}(0,0) \neq 0$ and conclude

$$0 \;=\; (p_x q_z - p_z q_x)(0,0,0) \;=\; (p_y q_z - p_z q_y)(0,0,0).$$

It remains to show that also $0 = (p_x q_y - p_y q_x)(0,0,0)$ holds. In the case that at least one of q_z or p_z does not vanish at $(0,0,0)$, without loss of generality p_z, this is easy to see: From $q_z(0,0,0) \neq 0$, $0 = (p_x q_z - p_z q_x)(0,0,0)$, and $0 = (p_y q_z - p_z q_y)(0,0,0)$ it follows

$$\begin{aligned} p_x(0,0,0) &= \left(\frac{p_z}{q_z} q_x\right)(0,0,0) \\ p_y(0,0,0) &= \left(\frac{p_z}{q_z} q_y\right)(0,0,0). \end{aligned}$$

This immediately leads to

$$(p_x q_y - p_y q_x)(0,0,0) \;=\; \left(\frac{p_z}{q_z} q_x q_y - \frac{p_z}{q_z} q_y q_x\right)(0,0,0) = 0.$$

If the case $p_z(0,0,0) = 0 = q_z(0,0,0)$ occurs, then the first subresultant of p and q vanishes at $(0,0)$, but we excluded this case in the assumptions. This proves our claim. □

We have seen that every singular point of a cutcurve is top-bottom or genuine. But also the other inclusion holds:

Corollary 5.3: Every top-bottom and every genuine point is a singular point.

Proof. Without loss of generality let $(0,0)$ be a top-bottom point of $f = \text{res}(p,q,z)$. That means there exist numbers $c, c' \in \mathbb{C}$, again without loss of generality $c = 0$, with $p(0,0,0) = q(0,0,0)$ and $p(0,0,c') = q(0,0,c') = 0$. So $(0,0)$ is a root of the first subresultant $\text{sres}_1(p,q,z)$. We have to show $f_x(0,0) = f_y(0,0) = 0$. From the previous proof we know

$$\begin{aligned}(*)\quad f_x(0,0) &= (((pq_z - p_zq)|_{z=0} \cdot \text{sres}_1(p,q,z) + ((p-q)|_{z=0})^2)_x)(0,0)\\ &= ((p_xq_z - p_zq_x)|_{z=0} \cdot \text{sres}_1(p,q,z))(0,0) + \\ &\quad ((pq_{xz} - p_{xz}q)|_{z=0} \cdot \text{sres}_1(p,q,z))(0,0) + \\ &\quad ((pq_z - p_zq)|_{z=0} \cdot (\text{sres}_1(p,q,z))_x)(0,0) + \\ &\quad (2(p-q)|_{z=0} \cdot (p_x - q_x))(0,0).\end{aligned}$$

With our assumptions $\text{sres}_1(p,q,z)(0,0) = 0$ and $p(0,0,0) = 0 = q(0,0,0)$, we immediately derive $f_x(0,0) = 0$. The proof for $f_y(0,0)$ works analogously.

If $(0,0)$ is a genuine point, then by definition

$$\begin{aligned}(p_xq_z - p_zq_x)(0,0,0) &= 0 = (p_yq_z - p_zq_y)(0,0,0)\\ p(0,0,0) &= 0 = q(0,0,0)\end{aligned}$$

and again from equation $(*)$ it easily follows $f_x(0,0) = f_y(0,0) = 0$. □

5.2 Classification of singular points

Projecting two intersection points in space on top of each other or projecting a tangential intersection point in space are exactly the events that cause singular points of a cutcurve. Of course, a singular point can be top-bottom as well as genuine. We will see next that under our assumptions of general alignment and squarefreeness a cutcurve can have at most 2 top-bottom points and at most 4 genuine points.

5.2.1 Top-bottom points

We will first prove that under our conditions of general alignment and squarefreeness a cutcurve f can have at most 2 top-bottom points. These 2 points can be determined using explicit computation as described in the previous chapter.

Theorem 5.4: Let $f \in \mathbb{Q}[x,y]$ be a generally aligned and squarefree polynomial. The cutcurve defined by f can have at most 2 top-bottom points.
Moreover, one can compute two at most quadratic polynomials $u_{\text{tb}} \in \mathbb{Q}[x]$ and $v_{\text{tb}} \in \mathbb{Q}[y]$ such that the top-bottom points lie on $\textsc{Grid}(u_{\text{tb}}, v_{\text{tb}})$.

Proof. Let $p \in \mathbb{Q}[x,y,z]$ and $q \in \mathbb{Q}[x,y,z]$ be two quadratic polynomials and $f = \text{res}(p,q,z)$. Without loss of generality we denote

$$\begin{aligned} p &= z^2 + p_1 z + p_0 \\ q &= z^2 + q_1 z + q_0 \end{aligned}$$

where $p_i \in \mathbb{Q}[x,y]$ and $q_j \in \mathbb{Q}[x,y]$ are polynomials of degree at most $2-i$. The first subresultant of p and q with respect to z is of the form:

$$\text{sres}_1(p,q,z) = q_1 - p_1.$$

That means $\text{sres}_1(p,q,z)$ is a polynomial of degree at most 1. The case $\text{sres}_1(p,q,z) \equiv 0$ is impossible, because then every point of the cutcurve f would be top-bottom and therefore singular. The curves f and f_y would have infinitely many intersections, contradicting our assumption of f being squarefree. So the first subresultant is not equal to the zero polynomial. That means

1. either it is constant and non-zero, in which case there is no top-bottom point,
2. or $l := \text{sres}_1(p,q,z)$ is a polynomial of total degree 1 defining a line.

In the first case nothing has to be done. In the second case all common roots of f and l are top-bottom and therefore singular points. We assume without loss of generality that l is generally aligned. The resultant $\text{res}(f,l,y) \in \mathbb{Q}[x]$ has degree at most 4. It cannot be the zero polynomial, because otherwise l would be a factor of f and all points on l would be singular points of f. Therefore l would also divide f_y, contradicting f being squarefree.

Each root of the resultant $\mathrm{res}(f, l, y)$ has multiplicity ≥ 2, because it results from a singular point of f through which also l cuts, see Theorem 4.4. We conclude that the resultant has at most 2 different roots. Therefore f can have at most two top-bottom points. The polynomial

$$u_{\mathrm{tb}} = \gcd(\mathrm{res}(f, l, y), \frac{d}{dx}\mathrm{res}(f, l, y))$$

has degree at most 2 and contains the x-coordinates of the top-bottom points. Analogously one can compute the polynomial v_{tb} using $\mathrm{res}(f, l, x)$. □

We have classified the singular points of a cutcurve f into top-bottom and genuine ones. Each singular point gives rise to a root of multiplicity at least 2 of the resultant $X = \mathrm{res}(f, f_y, y)$. The degree of X is at most 12, so a cutcurve can have at most 6 singular points. We have just shown, that at most 2 of them are top-bottom. Moreover we are able to determine a quadratic bi-factor (u_{tb}, v_{tb}) of the bi-resultant of f and f_y. With the help of this bi-factor explicit solutions for top-bottom points can be computed. What about the genuine points? We will prove next that their number is bounded by 4 and give the algorithmic ideas how to determine them.

5.2.2 Genuine points

The genuine points of a cutcurve $f = \mathrm{res}(p, q, z)$ arise from tangential intersections of p and q. Have again a look at Figure 5.2. The two quadrics intersect with the same tangential plane to both quadrics, causing a self-intersection of the spatial intersection curve. Generally, the tangent vector to a point on the spatial intersection curve is embedded in the tangential planes to both quadrics. The vector is well defined if the tangential planes are different. The tangential intersection points of p and q are exactly the singular points of the spatial intersection curve. For a classification of spatial intersection curves of two quadrics in projective space have a look at the one given in [28]:

1)	a non-singular space quartic
2)	a nodal space quartic
3)	a cuspidal space quartic
4)	a space cubic and a line, intersecting in 2 distinct points
5)	a space cubic and a line, touching at 1 point
6)	two conics intersecting in 2 distinct points
7)	two conics, touching at 1 point
8)	a conic and two lines, all intersecting in 1 point
9)	a single conic counted twice
10)	two lines, each counted twice
11)	a conic and a line counted twice
12)	three lines, one counted twice
13)	a conic and two lines, intersecting pairwise in 3 points
14)	four lines intersecting pairwise in 4 non-coplanar points

Each of the 14 cases is illustrated in the Figures 5.3 up to 5.9. The pictures are taken from [36].

Figure 5.3: A non-singular space quartic 1) and a nodal space quartic 2).

We are interested in the maximal number of genuine points and in a method to locate them. In the following theorem we will provide the upper bound. The proof will explicitly lead to the result that only in the cases 13) and 14) of the classification the cutcurve has more than 2 genuine points. In the situations

1) - 8) the cutcurve has at most 2 genuine points. This will lead to a quadratic bi-

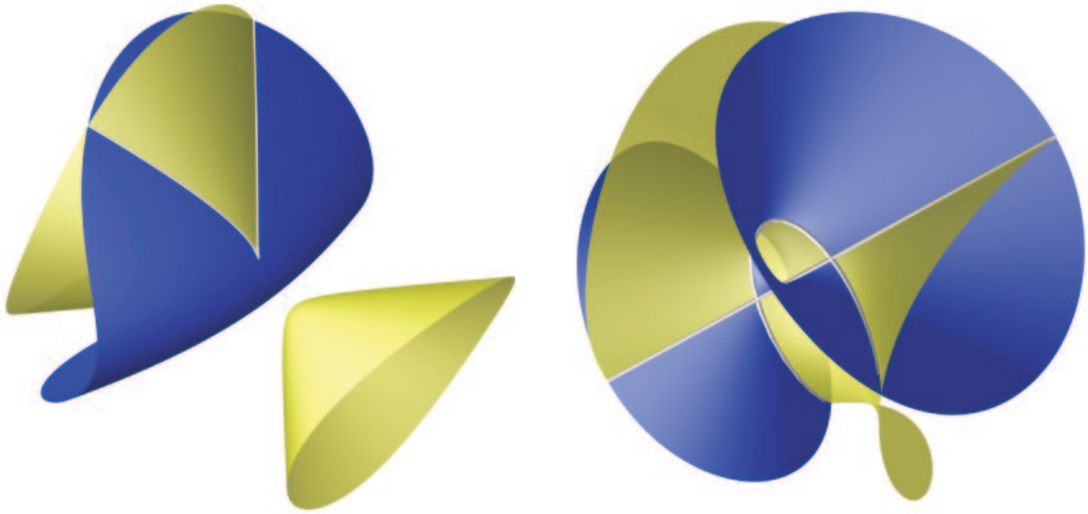

Figure 5.4: A cuspidal space quartic 3) and a space cubic and a line, intersecting in 2 distinct points 4).

factor solvable with explicit solutions.

9) - 12) the cutcurve is not squarefree, against our assumptions. It can be factored into two bivariate rational polynomials of total degree ≤ 2. Singular points of quadratic curves are easy to handle using explicit solutions.

13), 14) there are 3 or 4 discrete genuine points. These cases are the most complicated ones, because explicit solutions cannot be applied. The next section is dedicated to the subject of separating the genuine points in two groups. This will again lead to the problem of solving quadratic polynomials only.

But first we prove the upper bound on the number of genuine points:

Theorem 5.5: Let $p \in \mathbb{Q}[x, y, z]$ and $q \in \mathbb{Q}[x, y, z]$ be two generally aligned and squarefree quadratic polynomials with disjoint factorization such that $f := \text{res}(p, q, z)$ is squarefree and generally aligned. The cutcurve defined by f can have at most 4 genuine points. If it has more than 2 genuine points, f consists of two distinct lines and another quadratic curve, all of them not necessarily rational.

Proof. Let $A = (\alpha_x, \alpha_y, \alpha_z)$, $B = (\beta_x, \beta_y, \beta_z)$, and $C = (\gamma_x, \gamma_y, \gamma_z)$ be three distinct tangential intersection points of p and q. We will show that $f = \text{res}(p, q, z)$ consists of two lines and another quadratic curve and has at most 4 genuine points.

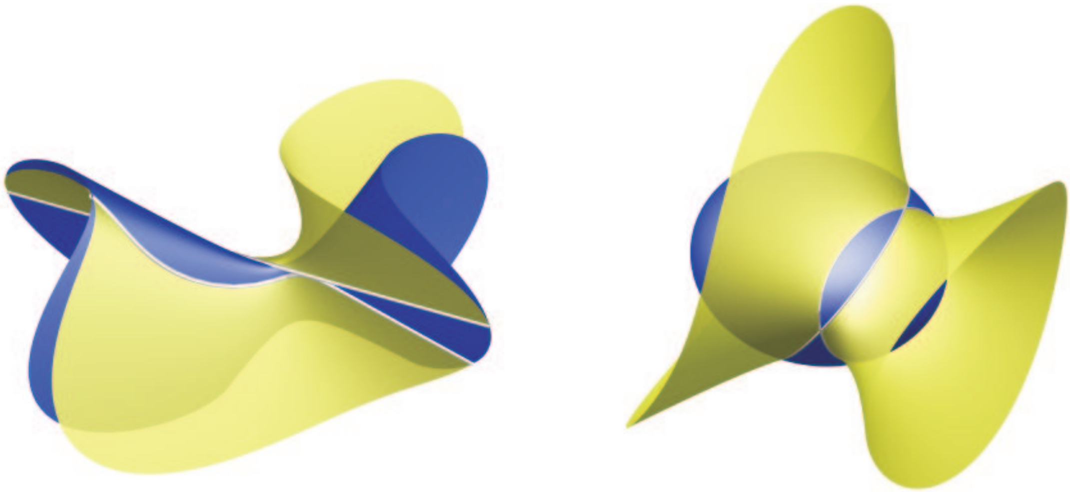

Figure 5.5: A space cubic and a line, touching at 1 point 5) and two conics intersecting in 2 distinct points 6).

There exists a plane defined by a polynomial $h \in \mathbb{C}[x,y,z]$ through these three points A, B, and C. If they do not lie on a single line l, plane h is uniquely defined.

A, B, C collinear

Let us first consider the case that the three points A, B, and C are collinear. Let l be the line passing through them. We will prove that in this case each point on l is a tangential intersection point of p and q. That means every point on the projection of l is genuine and therefore a singular point of f. We conclude that f and f_y have infinitely many intersections. This can only happen if f and f_y share a common factor, but this is against our assumption of f being squarefree. So due to our assumptions it is impossible that A, B, and C lie on a common line l.

We have to show that each point on l is a tangential intersection point of p and q. For this investigation, the location of l in space is not important. So we assume without loss of generality that the line l passing through the three points is the x-axis: $A = (\alpha_x, 0, 0)$, $B = (\beta_x, 0, 0)$, and $C = (\gamma_x, 0, 0)$. We are interested in the points $D = (x, 0, 0)$ on l and want to prove

$$0 = (p_x q_y - p_y q_x)(D) = (p_x q_z - p_z q_x)(D) = (p_y q_z - p_z q_y)(D).$$

We already know three different roots of the quadratic univariate polynomials $p(x,0,0)$ and $q(x,0,0)$, namely α_x, β_x, and γ_x. We conclude $p(x,0,0) \equiv 0 \equiv q(x,0,0)$. It easily follows $p_x(x,0,0) \equiv 0 \equiv q_x(x,0,0)$. That means every $D = (x,0,0)$ is a root of p, q,

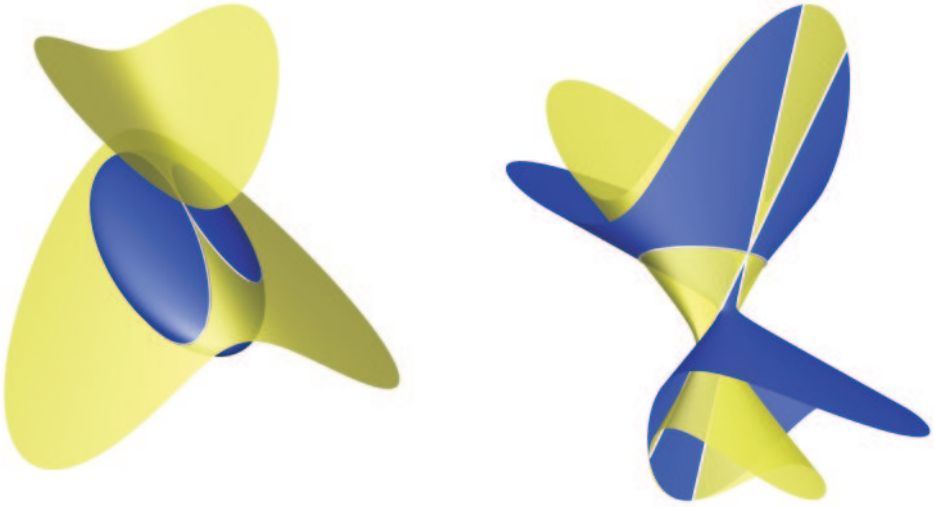

Figure 5.6: Two conics, touching at 1 point 7) and a conic and two lines, all intersecting in 1 point 8).

p_x, and q_x leading to

$$0 \; = \; (p_x q_y - p_y q_x)(D) \; = \; (p_x q_z - p_z q_x)(D).$$

The univariate polynomial $(p_y q_z - p_z q_y)(x, 0, 0) \in \mathbb{Q}[x]$ has degree at most 2. We again know three roots: α_x, β_x, and γ_x. So also $(p_y q_z - p_z q_y)(x, 0, 0)$ is the zero polynomial and of course

$$(p_y q_z - p_z q_y)(D) \; = \; 0.$$

Every D is a tangential intersection point of p and q.

A, B, C not collinear and h is a factor of p or q

Next consider the case of A, B, and C not being collinear. What happens if h is a factor of p or q? It cannot be a factor of both, because by assumption p and q have disjoint factorizations. Assume without loss of generality $p = h \cdot \tilde{h}$. The spatial intersection curve consists of the two quadratic curves embedded in h and $\tilde{h}$, respectively. Corollary 2.19 leads to

$$f \; = \; \operatorname{res}(h, q, z) \; \cdot \; \operatorname{res}(\tilde{h}, q, z) \; := \; f_1 \cdot f_2.$$

By assumption f is squarefree. Therefore f_1 and f_2 are squarefree and have a disjoint factorization. The singular points of f are either intersection points of f_1 and f_2 or singular points of f_1 or singular points of f_2.

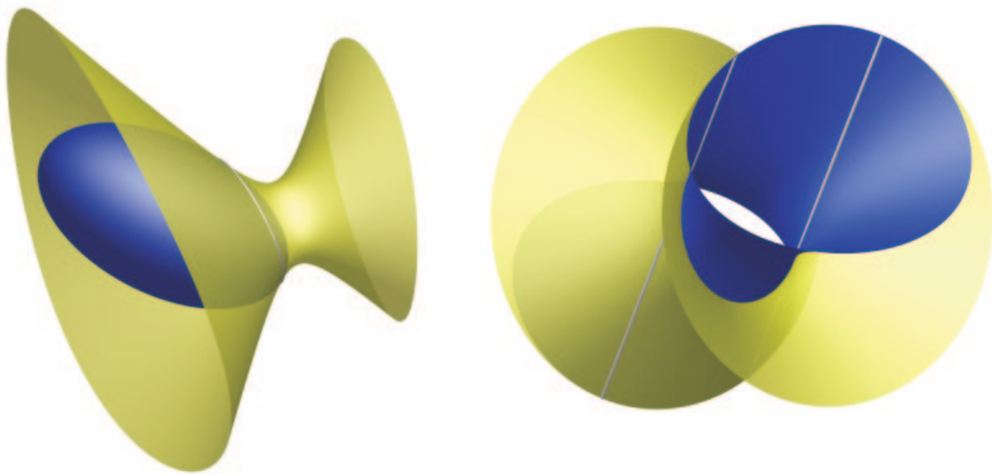

Figure 5.7: A single conic counted twice 9) and two lines, each counted twice 10).

At most 2 intersection points of f_1 and f_2 are genuine points of f. This can be easily seen, because an intersection point (a, b) of f_1 and f_2 is genuine if there exists a point $(a, b, c) \in \mathbb{C}^3$ such that $h(a, b, c) = \tilde{h}(a, b, c) = q(a, b, c) = 0$. The set of points $\{(x, y, z) \in \mathbb{C}^3 \mid h(x, y, z) = \tilde{h}(x, y, z) = 0\}$ defines a line in space. A line intersects a quadric in at most 2 points.

The curve f has 3 genuine points, so we conclude that at least one of f_1 or f_2 must have a singular point. We will show that a planar quadratic curve has a singular point if and only if it consists of two intersecting lines. We conclude that

1. f can have at most 4 genuine points, because each planar curves f_1 and f_2 has at most 1 singular point and at most 2 intersection points of f_1 and f_2 are genuine points of f.

2. f consists of two intersecting lines and another quadratic curve.

We still have to show that a planar quadratic curve has a singular point if and only if it consists of two intersecting lines. One direction is easy: If a curve consists of two intersecting lines, then of course it has a singular point. For the other direction assume that there exists a quadratic planar curve g not consisting of two lines that has a singular point $a = (a_x, a_y)$. We can find two points $b = (b_x, b_y)$ and $c = (c_x, c_y)$ lying on g such that a, b, and c are not collinear. We choose a line l_1 cutting through a and b and a line l_2 cutting through a and c. By assumption $l_1 \cdot l_2$ and g have a disjoint factorization and $\text{res}(l_1 l_2, g, y) \not\equiv 0$ has degree at most 4. We know the roots

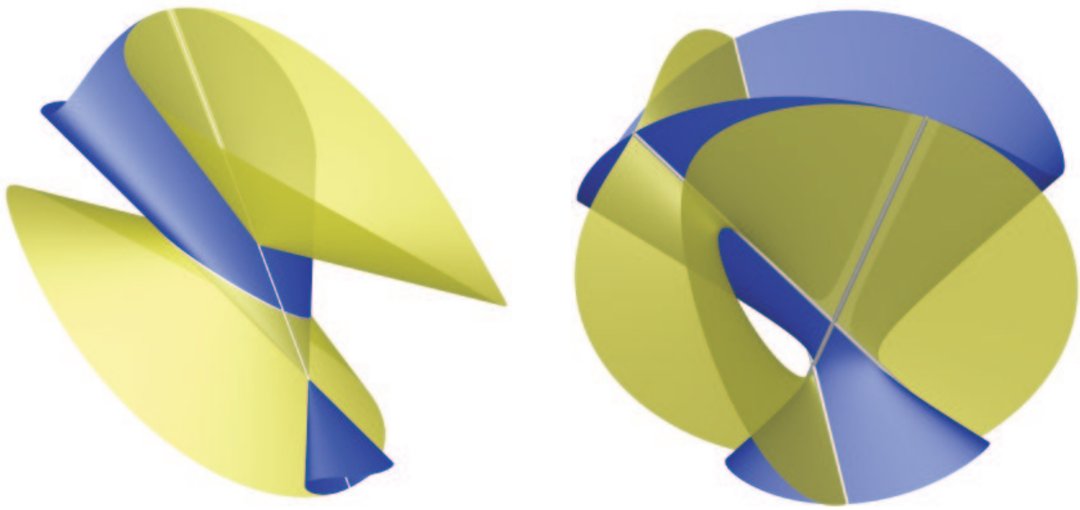

Figure 5.8: A conic and a line counted twice 11) and three lines, one counted twice 12).

of the resultant, namely a_x, b_x, and c_x. But (a_x, a_y) is a singular point of g as well as of $l_1 l_2$. Therefore a_x is a root of multiplicity at least 4 of the resultant, leading to a contradiction.

A, B, C not collinear and h is neither a factor of p nor of q

The case we have not discussed so far is that the tangential intersection points $A = (\alpha_x, \alpha_y, \alpha_z)$, $B = (\beta_x, \beta_y, \beta_z)$, and $C = (\gamma_x, \gamma_y, \gamma_z)$ of p and q are not collinear and h is neither a factor of p nor of q. Then the spatial intersection curves of h and p and of h and q are quadratic. So on h there are two quadratic planar curves that have 3 tangential intersection points. This can only happen if both curves are identical:

$$\{(x, y, z) \mid p(x, y, z) = 0 = h(x, y, z)\} \;=\; \{(x, y, z) \mid q(x, y, z) = 0 = h(x, y, z)\} \;:=\; S.$$

We easily conclude that p and q intersect in quadratic spatial curves, the number of which is at most two.

Under our assumption of $f = \operatorname{res}(p, q, z)$ being squarefree, at least one of the two quadratic spatial curves consists of two lines. We will prove this via contradiction by showing that otherwise p and q would have the same tangential plane in every point of S. Then f would have infinitely many singular points, contradicting the squarefreeness of f. That means p and q intersect in two quadratic curves, at least one consisting of two lines. So there are at most 4 genuine points of f. This leads to our desired result.

It remains to prove the following: If the spatial intersection curve of p and h (and this is equal to the one of q and h) does not consist of two lines, then p and q have the same tangential plane in every point of S. For the proof, the location of h in space is

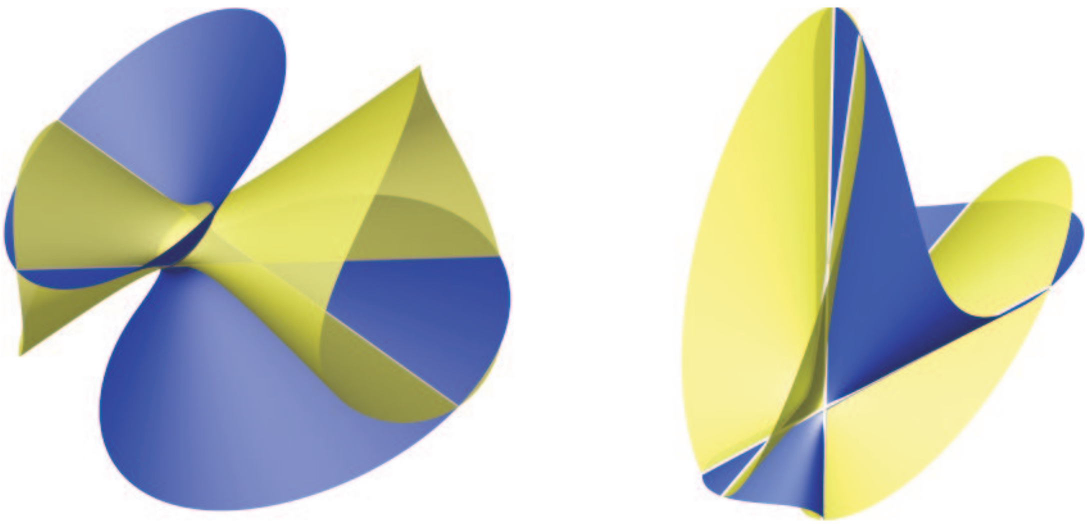

Figure 5.9: A conic and two lines, intersecting pairwise in 3 points 13) and four lines intersecting pairwise in 4 non-coplanar points 14).

not important. So we assume without loss of generality that h is the (x, y)-plane, i.e. $\alpha_z = \beta_z = \gamma_z = 0$.

Let $\tilde{p}(x, y) := p(x, y, 0)$, $\tilde{q}(x, y) := p(x, y, 0)$, and $D = (x, y, 0)$ with $p(D) = 0$ an arbitrary point on the spatial intersection curve of p and h. We already know that $\tilde{p}$ and $\tilde{q}$ define the same quadratic planar curve. It easily follows $(\tilde{p}_x\tilde{q}_y - \tilde{p}_y\tilde{q}_x) \equiv 0$ and therefore

$$(p_xq_y - p_yq_x)(D) \;=\; 0.$$

What can we say about $(p_zq_y - p_yq_z)(D)$ and $(p_zq_x - p_xq_z)(D)$? Remember that $\tilde{p}$ and $\tilde{q}$ define the same curve, so there exists a constant $c \in \mathbb{R}$ such that $c\tilde{p}(x, y) = cp(x, y, 0) = \tilde{q}(x, y) = q(x, y, 0)$:

$$\begin{aligned} h_1(x, y) &:= (p_zq_y - p_yq_z)(x, y, 0) \\ &= p_y(x, y, 0) \cdot (cp_z - q_z)(x, y, 0) \\ h_2(x, y) &:= (p_zq_x - p_xq_z)(x, y, 0) \\ &= p_x(x, y, 0) \cdot (cp_z - q_z)(x, y, 0). \end{aligned}$$

We know that (α_x, α_y), (β_x, β_y), and (γ_x, γ_y) are roots of h_1 and h_2. The polynomial $\tilde{p}$ does not define two lines, so no point among them is a singular point of $\tilde{p}$. We conclude that we have three non-collinear points on the line defined by the polynomial $(cp_z - q_z)(x, y, 0)$. That means h_1 and h_2 both are equal to the zero polynomial and all $D = (x, y, 0)$ with $p(D) = 0$ are tangential intersection points of p and q. □

The statement of the theorem is that in most cases a squarefree cutcurve has at most 2 genuine points. Bad things only happen in the situations 13) and 14) resulting in one of the two cutcurves shown in Figure 5.10. There are at least 3 genuine points, marked by the red circles. The remaining singular points are top-bottom points and marked by the blue squares. What remains to do is further classifying the genuine points in a way that computing explicit solutions becomes a suitable tool also in these cases.

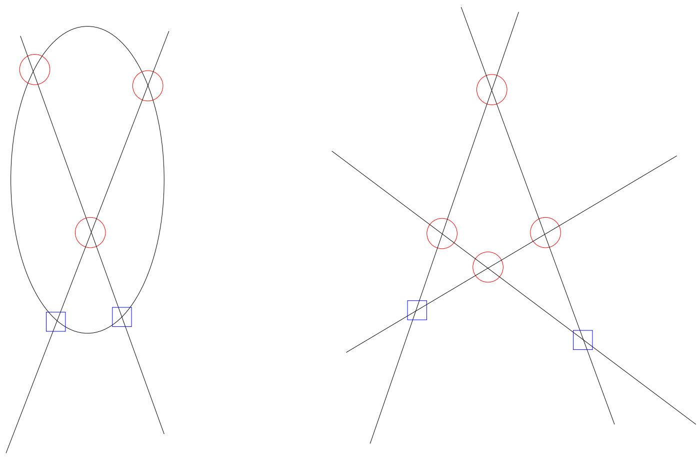

13) 14)

Figure 5.10: The cutcurve consists of two intersecting lines and another quadric

A close look at the proof of the theorem leads to the following statement: A cutcurve f consists of two lines and another quadratic curve if and only if the corresponding quadrics p and q in space intersect in two lines and another quadratic spatial curve, each embedded in a plane. Let h and $\tilde{h}$ be these two planes and $r := h \cdot \tilde{h}$. For h and $\tilde{h}$ we have that

a) at most 2 tangential intersection points lie on the intersection line l of h and $\tilde{h}$,

b) at most 1 lies on h and not on l and

c) at most 1 lies on $\tilde{h}$ and not on l.

It would be quite useful to know r, because the line defined by the polynomial $\text{res}(r, r_z, z)$ is the projection of l. It cuts through the genuine points caused by the tangential points in a). With the help of this line the genuine points could be classified: the ones that lie on the line and the ones that do not. Each group has at most two members. As we did for top-bottom points we could factor the resultants on the x- and y-axis according to this distinction leading to quadratic polynomials as desired.

In the rest of this chapter we show that we really can compute r. The polynomial r either is rational or its coefficients are in $\mathbb{Q}(\sqrt{\rho})$ for some $\rho \in \mathbb{Q}$. The way the different factorization criteria just developed fit together and lead to an algorithm for determining singular points will be described in details in the next chapter.

5.2.3 Quadric pencils

We want to prove that, in the case the cutcurve $f = \text{res}(p, q, z)$ consists of two lines and another quadratic curve, we can compute a quadric r consisting of two planes such that $f = \text{res}(p, r, z)$. This result is not new. It is part of the theory of quadric pencils and can be found for example in [28]. The following survey about quadric pencils leading to our desired result is taken from this article. Here is the statement we want to prove:

Theorem 5.6: Let p and q be two quadrics in space.

1. If $f = \text{res}(p, q, z)$ consists of two lines and another quadratic curve not equal to two lines, then there exists a polynomial $r \in \mathbb{Q}[x, y, z]$ defining two planes such that $f = \text{res}(p, r, z)$.

2. If $f = \text{res}(p, q, z)$ consists of four lines, then there exists a polynomial $r \in \mathbb{Q}(\sqrt{\rho})[x, y, z]$, for some $\rho \in \mathbb{Q}$, defining two planes such that $f = \text{res}(p, r, z)$.

In both cases we can compute r.

For the proof we have to clarify some notation. Up to now we used the implicit description of a quadric: we defined a quadric to be the set of roots of a rational polynomial

$$p(x, y, z) \;=\; ax^2 + by^2 + cz^2 + 2dxy + 2exz + 2fyz + 2gx + 2hy + 2kz + l.$$

Another way of representing p is just to give the symmetric matrix

$$P \;=\; \begin{pmatrix} a & d & e & g \\ d & b & f & h \\ e & f & c & k \\ g & h & k & l \end{pmatrix}.$$

Of course both notations are equivalent:

$$p = (x, y, z, 1) \cdot P \cdot (x, y, z, 1)^T.$$

For a quadric p let the capital P denote its matrix representation. The determinant $\Delta = \det P$ is called the *discriminant* of the quadric surface. A quadric surface with $\Delta = 0$ is classified *singular* and called a *cone*. If moreover $\Delta = 0$ and P only has rank 2, the quadric cone degenerates into two distinct real or complex conjugate planes. If the rank of P is 1, the two planes are rational and identical. That means the quadric is not squarefree.

Two distinct squarefree quadric surfaces P_0 and P_1 define an infinite linear family of quadrics, or *quadric pencil*, given by

$$P_\lambda \;=\; (1-\lambda)P_0 + \lambda P_1$$

for $-\infty < \lambda < \infty$. In particular, if C denotes the intersection curve of P_0 and P_1, then any two distinct members P_r and P_s $(r \neq s)$ of the pencil also have C for their intersection. Like in traditional approaches to computing quadric intersections ([46] and [47]), we will search the pencil for a simple member which facilitates our calculation. We are interested in a λ such that P_λ defines a surface consisting of two planes.

Consider the 4×4 determinant of a generic member of the pencil:

$$\begin{aligned} \Delta(\lambda) &= \det((1-\lambda)P_0 + \lambda P_1) \\ &= \Delta_4\lambda^4 + \Delta_3\lambda^3 + \Delta_2\lambda_2 + \Delta_1\lambda + \Delta_0. \end{aligned}$$

This rational polynomial in $\mathbb{Q}[\lambda]$ is called the *discriminant* of the pencil. Its distinct roots identify the cones of the pencil. Let $\lambda_1, \ldots, \lambda_n$, $n \leq 4$, denote the different roots of $\Delta(\lambda)$ and let $m_1, \ldots, m_n$ be their multiplicities, so that

$$\Delta(\lambda) \;=\; c\prod_{i=1}^{n}(\lambda-\lambda_i)^{m_i}, \qquad \text{where } \sum_{i=1}^{n} m_i \;\leq\; 4, \quad c \neq 0.$$

A quadric pencil may be uniquely characterized by at most 4 sequences of positive integer numbers, one sequence for each root λ_i of $\Delta(\lambda)$. This classification is called *Segre characteristic* [14]. If λ_i is a simple root, the sequence contains exactly one element: (1). For a multiple root, however, there may be more elements in the sequence. From the 4×4 determinant P_λ we may form sixteen 3×3 first minors by deleting single rows and columns, thirty-six 2×2 second minors by deleting pairs of rows and columns, and finally sixteen 1×1 third minors by deleting triples of rows and columns. These first, second, and third minors are generally cubic, quadratic, and linear in λ. If λ_i is a root of multiplicity $m_i \geq 2$ we ascertain its multiplicities m_i', m_i'', and m_i''' as a root of every first, second, and third minor, respectively, and compute the differences

$$d_i^1 = m_i - m_i', \qquad d_i^2 = m_i' - m_i'', \qquad d_i^3 = m_i'' - m_i'''.$$

The sequence associated with the multiple root λ_i is given by the sequence of differences d_i^j, terminating when a zero is encountered. The table of classification we have just seen is according to this Segre characteristic.

In the cases we are interested in, namely the spatial intersection curve consists of two lines and a cone intersecting in projective space in 3 or 4 points, the Segre characteristic is as follows:

13)	a conic and two lines, intersecting pairwise in 3 points	$[(2),(1,1)]$
14)	four lines intersecting pairwise in 4 non-coplanar points	$[(1,1),(1,1)]$

With the help of this characterization we can proof our theorem:

Proof. (of Theorem 5.6)

1. Let $f = \text{res}(p,q,z)$ define two lines and another quadratic curve not equal to two lines. That means the two quadrics P and Q have a spatial intersection curve as described in 13). The first information we obtain from the Segre characteristic is that the discriminant of the pencil has 2 roots: λ_1 and λ_2.

   The sequence for λ_1 is (2). That means $d_1^1 = m_1 - m_1' = 2$ and $m_1' = 0$. We conclude $m_1 = 2$ and λ_1 is a double root of the discriminant of the pencil, but not a common root of all first minors of $\Delta(\lambda)$. The quadric $P_{\lambda_1} = (1-\lambda_1)P + \lambda_1 Q$ is a non-degenerate cone because it has rank 3.

   The second root λ_2 has the sequence $(1,1)$. We have $m_2 = 2$ and $m_2' = 1$ and therefore $P_{\lambda_2} = (1-\lambda_2)P + \lambda_2 Q$ has rank 2.

   The two roots λ_1 and λ_2 of $\Delta(\lambda)$ both have multiplicity 2 and are the only roots of the degree 4 polynomial $\Delta(\lambda)$. Moreover λ_2 is a root of each 3×3 minor of P_λ and λ_1 is not. So by computing gcd's of the discriminant and the first minors we obtain the rational value of λ_2. We conclude that there exists a rational quadric $R = P_{\lambda_2}$ that consists of two planes such that P and Q have the same spatial intersection curve as P and R.

2. Let $f = \operatorname{res}(p, q, z)$ define four lines. The two quadrics P and Q have a spatial intersection curve as described in 14). The discriminant of the pencil again has two roots λ_1 and λ_2 of multiplicity 2. In this case we cannot distinguish them, because they both are roots of every first minor of P_λ. So we only know that they are roots of the degree 2 polynomial $\gcd(\Delta(\lambda), \frac{\partial}{\partial\lambda}\Delta(\lambda)) \in \mathbb{Q}[\lambda]$. But we can compute λ_1 and λ_2 as elements of $\mathbb{Q}(\sqrt{\rho})$ for some $\rho \in \mathbb{Q}$. We determine a quadric $R = P_{\lambda_1}$ with implicit representation $r \in \mathbb{Q}(\sqrt{\rho})[x, y, z]$ such that R consists of two distinct planes and the intersection curves of P and Q and P and R are identical.

□

Chapter 6

Computing the planar arrangements

In this chapter we consider arrangements that arise from projecting the silhouette of a quadric as well as all its intersection curves with other quadrics into the plane. We will prove our main theorem, namely that we can determine all extreme and singular points of the resulting planar curves. Additionally, all intersection points of every pair of curves can be located. As we have already seen, we can compute boxes with rational corners such that each real event point is contained in such a box. The number of boxes exceeds the number of real event points. We show how to distinguish the boxes that contain a real event point from the empty ones.

Each planar arrangement we obtain consists of exactly one silhouettecurve and of many cutcurves. The restriction to one silhouettecurve is only a result of our algorithmic approach in three dimensions. Our algorithm can be easily extended to arrangements with additional silhouettecurves.

Let $\mathcal{F}$ be the set of curves in a planar arrangement. As we have already seen, in order to compute the planar arrangement we have to locate for all $f \in \mathcal{F}$ the intersection points of f and $g = f_y$. Furthermore we have to compute all intersection points between two curve $f \in \mathcal{F}$ and $g \in \mathcal{F}$. According to the distinction of silhouettecurve and cutcurve we obtain four different kinds of pairs of curves f and g:

1. f is the silhouettecurve and $g = f_y$,
2. f is a cutcurve and $g = f_y$,
3. f is the silhouettecurve and g is a cutcurve, and
4. f and g are both cutcurves.

We have to show that, in each case, we can locate all intersection points. This is the main topic of this chapter and we will use the mathematical tools developed earlier. But first we state our main theorem:

Theorem 6.1: Let $\mathcal{P} = \{p_1, \ldots, p_n\}$ be a set of trivariate quadratic polynomials. For an arbitrary $1 \leq i \leq n$ let furthermore

$$\mathcal{F}_i \ := \ \{\mathrm{res}(p_i, (p_i)_z, z)\} \ \cup \ \bigcup_{i \neq j} \{\mathrm{res}(p_i, p_j, z)\}.$$

For each pair of polynomials f and g, either $f, g \in \mathcal{F}_i$ or $f \in \mathcal{F}_i$ and $g = f_y$, we can compute a set of k 4-tuples of rational numbers (a_j, b_j, c_j, d_j), $1 \leq j \leq k$, that is in 1-1 correspondence to the set of real intersection points of the curves defined by f and g. The jth real intersection point (α_j, β_j) is the only one for which $a_j \leq \alpha_j \leq b_j$ and $c_j \leq \beta_j \leq d_j$ holds. Moreover we can determine whether the intersection point represented by (a_j, b_j, c_j, d_j) is transversal or tangential and whether it is a singular point of one of the curves.

The theorem states that all boxes containing real intersection points between two curves $f \in \mathcal{F}$ and $g = f_y$ and between $f \in \mathcal{F}$ and $g \in \mathcal{F}$ can be determined. In the next four sections we will give the proof. We show that in all four cases listed before either applying box hit counting arguments or partially factoring the bi-resultant of f and g and computing explicit solutions leads to the desired result. We are able to distinguish empty boxes from the ones that contain a real intersection point.

We assume without loss of generality that all pairs of curves are well-behaved. That means all curves are generally aligned and squarefree and every pair has a disjoint factorization and is in general relation. The way these conditions are tested and realized is described in the next chapter.

Let in the following $X := \mathrm{res}(f, g, y)$ and $Y := \mathrm{res}(f, g, x)$.

6.1 f is the silhouettecurve and $g = f_y$

If f is the silhouettecurve and $g = f_y$, then we know that both resultants X and Y have degree at most 2:

$$\begin{aligned} \deg(X) \ &\leq \ \deg(f) \cdot \deg(g) \ \leq \ 2 \cdot 1. \\ \deg(Y) \ &\leq \ \deg(f) \cdot \deg(g) \ \leq \ 2 \cdot 1. \end{aligned}$$

That enables us to compute explicit solutions. Consider also Figure 6.1.

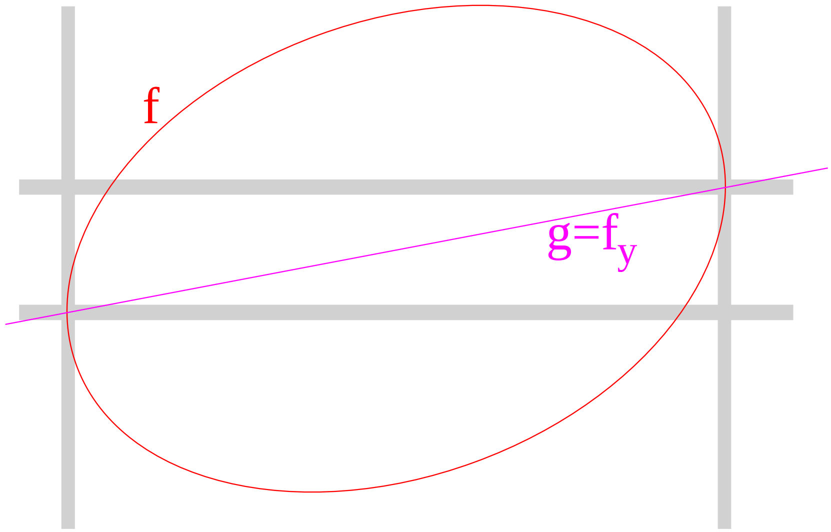

Figure 6.1: We can compute explicit solutions

6.2 f is a cutcurve and $g = f_y$

In the case f is a cutcurve and $g = f_y$ is its partial derivative, the resultants X and Y have degree at most 12 because $\deg(f) \leq 4$ and $\deg(g) \leq 3$.

We compute a multiplicity bi-factorization $(u_1, v_1) \cdot (u_2, v_2)$ of (X, Y) such that all intersection points with multiplicity 1, i.e. all transversal intersection points, lie on $\textsc{Grid}(u_1, v_1)$. All intersection points with multiplicity ≥ 2, i.e. all tangential intersection points, lie on $\textsc{Grid}(u_2, v_2)$. For illustration have a look at Figure 6.2.

The light grey boxes around $\textsc{Grid}(u_1, v_1)$ can be handled with simple box hit counting as shown in Chapter 4.

The dark grey boxes around $\textsc{Grid}(u_2, v_2)$ are the candidate boxes for tangential intersections. Each root of u_2 and v_2 occurs with multiplicity at least 2. Using multiplicity factorization, we compute rational polynomials u_2' and v_2' such that u_2' and v_2' only have simple roots and $\mathcal{R}(u_2) = \mathcal{R}(u_2')$ as well as $\mathcal{R}(v_2) = \mathcal{R}(v_2')$. Then $\textsc{Grid}(u_2, v_2) = \textsc{Grid}(u_2', v_2')$ and we have to test the boxes defined by u_2' and v_2' whether they contain a tangential intersection or not.

What does a tangential intersection of f and $g = f_y$ geometrically mean? By the mathematical definition, a point $(a, b) \in \mathbb{C}^2$ is a tangential intersection of f and g if and only if $f(a, b) = 0$, $g(a, b) = f_y(a, b) = 0$, and

$$(f_x g_y - f_y g_x)(a, b) = (f_x f_{yy} - f_y f_{xy})(a, b) = (f_x f_{yy})(a, b) = 0.$$

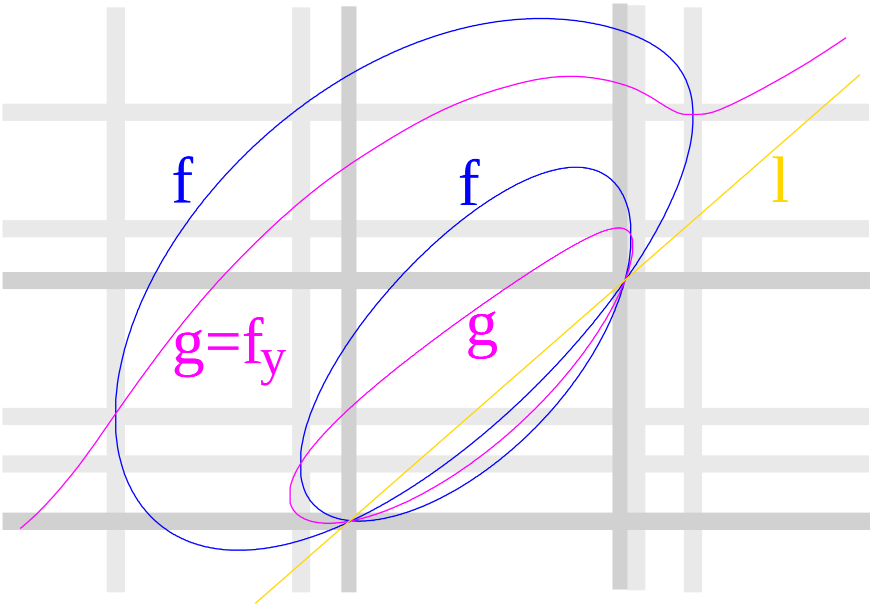

Figure 6.2: Transversal intersections lie inside light grey boxes, tangential intersection points lie inside dark grey boxes. The yellow line cuts through the top-bottom points.

The last equation is equal to zero if and only if $f_x(a,b) = 0$ or $f_{yy}(a,b) = 0$. We conclude that there are two kinds of tangential intersections:

1. Singular points of f, that means $f_x(a,b) = 0$, and
2. non-singular points of f with $f_{yy}(a,b) = 0$. We call these points *vertical flat points*. A vertical flat point is either a vertical turning point of f ($f_{yyy}(a,b) \neq 0$) or an extreme point of f ($f_{yyy}(a,b) = 0$).

According to Theorem 5.4 we compute a bi-factor (u_{tb}, v_{tb}) of (u_2', v_2') splitting off the top-bottom points of f: $(u_2', v_2') = (u_{tb}, v_{tb}) \cdot (u_g, v_g)$. Both polynomials in (u_{tb}, v_{tb}) have degree at most 2 and therefore it is possible to compute explicit solutions. In our example all tangential intersection points are top-bottom. The yellow line l, which is the first subresultant of the involved spatial quadrics, cuts through them, see Figure 6.2.

In the case that u_g as well as v_g are at most quadratic polynomials, as in our example, everything is fine and we compute explicit solutions also for (u_g, v_g). Of course it can happen that a tangential intersection point (a,b) we computed explicitly this way is a vertical flat point instead of a singular point. In order to recognize this, we substitute (x,y) by (a,b) in f_x and explicitly test $f_x(a,b)$ for zero. If $f_x(a,b) = 0$, (a,b) is a singular point of f. Otherwise it is a vertical flat point and we explicitly test $f_{yyy}(a,b)$ for zero in order to distinguish vertical turning points of f from extreme points of f.

Now consider the case that u_g or v_g or both have degree > 2. We assumed that all curves we consider are squarefree. We conclude according to Theorem 5.5 that

1. either the cutcurve consists of two intersecting lines and another quadratic curve
2. or the cutcurve has at least one vertical flat point.

We first test whether the cutcurve consists of two intersecting lines and another quadratic curve. Let p_0 and p_1 be the quadrics with $f = \text{res}(p_0, p_1, z)$ and let P_0 and P_1 be their matrix representations introduced in the previous chapter, respectively. We compute the discriminant

$$\Delta(\lambda) = \det((1-\lambda)P_0 + \lambda P_1) := \det(P_\lambda).$$

and determine whether the discriminant has two roots, each of multiplicity 2. We compute these roots λ_1 and λ_2 explicitly. If for at least one root λ_i the matrix P_{λ_i} has rank less than 3, the cutcurve consists of two intersecting lines and another quadratic curve. The statement of Theorem 5.6 is as follows: If exactly one root fulfills this property, without loss of generality λ_1, then λ_1 is rational. Moreover $Q := P_{\lambda_1}$ is the matrix representation of a rational polynomial $q \in \mathbb{Q}[x,y,z]$ such that q defines two, not necessarily rational, planes q_1 and q_2, $q = q_1 \cdot q_2$, and $\text{res}(p_0, p_1, z) = \text{res}(p_0, q, z)$. If both roots fulfill the property, we cannot distinguish them and we have $q \in \mathbb{Q}(\sqrt{\rho})[x,y,z]$ with $\rho \in \mathbb{Q}$.

Let l denote the projected intersection line of q_1 and q_2. In the discussion after the proof of Theorem 5.5 we noticed that at most 2 genuine points lie on l and at most 2 genuine points do not lie on l. Although we do not know q_1 and q_2 explicitly, we can compute l:

$$l^2 \;:=\; \text{res}(q, q_z, z).$$

We know that l is a bivariate polynomial in $\mathbb{Q}[x,y]$ or in $\mathbb{Q}(\sqrt{\rho})[x,y]$ and it cuts through at most two genuine points. As in the case of top-bottom points, we use l to factor off these points: $(u_g, v_g) = (u_{g1}, v_{g1}) \cdot (u_{g2}, v_{g2})$ with

$$\begin{aligned} u_{g1}^2 &= \text{res}(l, f, y) \qquad & v_{g1}^2 &= \text{res}(l, f, x) \\ u_{g2} &= u_g / u_{g1} & v_{g2} &= v_g / v_{g1}. \end{aligned}$$

The polynomials u_{g1} and v_{g1} are at most quadratic and again we compute explicit solutions. We additionally test f_x for zero in order to distinguish singular points from the ones with a vertical tangent. What remains is the bi-factor (u_{g2}, v_{g2}). If both polynomials are at most quadratic we are fine: compute explicit solutions and test f_x for zero in order to distinguish singular points from the ones with a vertical tangent. In the later case we additionally test f_{yyy} for zero. If at least one of u_{g2} or v_{g2} has degree > 2, we know that this is caused by vertical flat points.

At this time of our algorithm, we have partially factored the bi-polynomial (u_2', v_2') in such a way that the grid of each bi-factor (u_{bi}, v_{bi}) covers at most two singular points. If there is still a bi-factor (u_{bi}, v_{bi}) with one of the polynomials u_{bi} or v_{bi} having degree greater than 2, we know that this is caused by vertical flat points. If we detect this, we apply a random shear to the curves f and g in order to get rid of this situation and restart, see Chapter 7 for further information on this subject. What we obtain at last is a bi-factorization of (u_2', v_2') into ≤ 3 at most quadratic bi-factors for which we compute explicit solutions.

For illustration let us again have a look at the cutcurves consisting of a pair of lines and another quadratic curve in Figure 6.3.

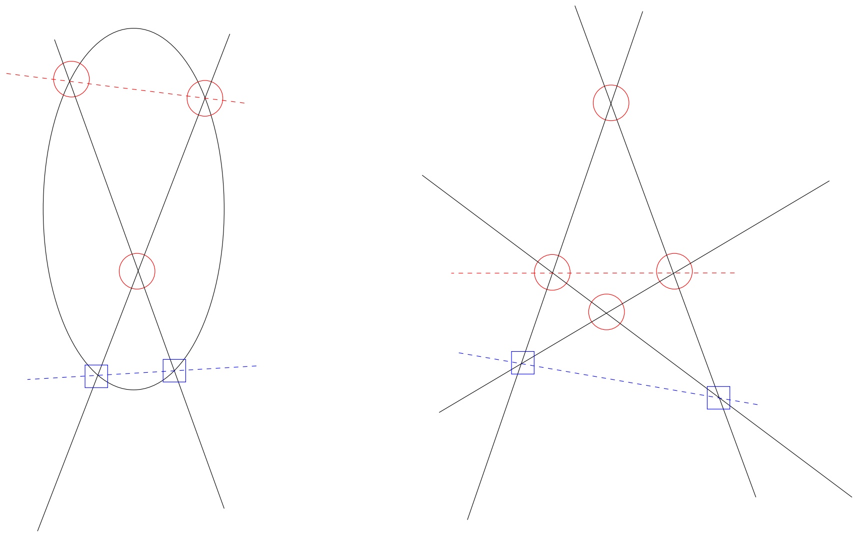

Figure 6.3: The cutcurve consists of two intersecting lines and another quadric

In the first picture the cutcurve, consisting of a pair of lines and a quadratic curve, has 5 singular points. The x- and y-coordinates are simple roots of a polynomial of degree 5, respectively. Two of the singular points are top-bottom, resulting from the projection phase. They are marked by the blue boxes and we know the line that passes through them: the first subresultant of the corresponding spatial quadrics. The remaining 3 singular points marked by the red circles are genuine. We compute a third rational quadric for which we know that it consists of two planes. We project the common line of these two planes and obtain a rational line running through 2 of the genuine points. With the help of the two lines we factor the polynomial of degree 5 in three polynomials: two of degree 2 and one of degree 1.

In the second picture the cutcurve consists of four lines intersecting pairwise in 6 singular points. Now the x- and y-coordinates are simple roots of a polynomial of degree 6, respectively. Again two points are top-bottom and we know a rational line through them. For the 4 genuine points we compute a third quadric in space that consists of two planes. Projecting their common line results in a linear polynomial with coefficients in $\mathbb{Q}(\sqrt{\rho})$, $\rho \in \mathbb{Q}$, for which 2 of the genuine points are roots. Again with the help of the two lines we partially factor the polynomial of degree 6 in three polynomials, each of degree 2.

6.3 f is the silhouettecurve and g is a cutcurve

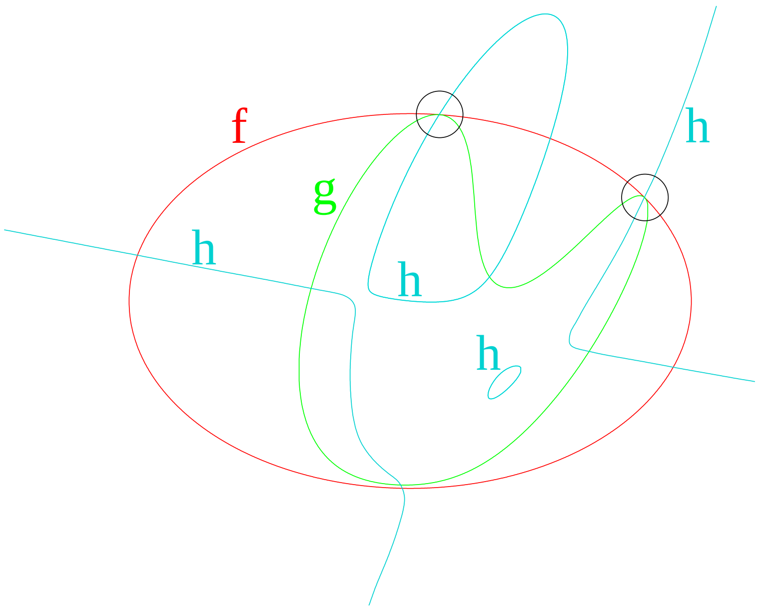

Figure 6.4: Consider the Jacobi curve in order to solve simple tangential intersections

Let f be the silhouettecurve and g be a cutcurve. Then the algebraic degree of f is at most 2. The degree of g is at most 4. Therefore the polynomials X and Y have degree at most 8. This implies that there are at most two roots of multiplicity ≥ 3. We compute a bi-factorization $(u_1, v_1) \cdot (u_2, v_2)^2 \cdot (u_3, v_3)$ of (X, Y) such that u_1, v_1 contain all simple roots, u_2, v_2 all roots of multiplicity 2, and u_3, v_3 all roots of multiplicity ≥ 3.

All transversal intersections points lie on GRID(u_1, v_1) and can be solved with simple box hit counting.

The ones lying on GRID(u_2, v_2)

1. either are singular points of f or g
2. or transversal intersections of the Jacobi curve $h = f_x g_y - f_y g_x$ and f and of h and g, according to Corollary 4.9.

We would like to apply extended box hit counting to these boxes, but first we have to be sure that there is no singular point inside the tested box. In the last section we have shown how to compute quadratic bi-factors (u_{tb}, v_{tb}), (u_{g1}, v_{g1}), and (u_{g2}, v_{g2}) for all singular points and compute explicit solutions. If any of these bi-factors has a common bi-factor with (u_2, v_2), we split off this common bi-factor. What remains is a bi-factor (u_2', v_2') with only non-singular tangential intersections of f and g on its grid. We apply extended box hit counting to the boxes defined by u_2' and v_2'.

Because of the degree of X and Y the bi-polynomial (u_3, v_3) has at most two different roots. With the help of gcd-computation we compute two at most quadratic polynomial u_3' and v_3' containing them and apply explicit solutions.

6.4 f and g both are cutcurves

Let f and g both be cutcurves. They are the result of intersecting a quadric p with other quadrics q and r, respectively. Each cutcurve has algebraic degree ≤ 4 and therefore the polynomials X and Y have degree at most 16. We would like to compute a bi-factorization $(X, Y) = (u_s, v_s) \cdot (u_a, v_a)$ such that all polynomials u_s, v_s, u_a, v_a have degree at most 8 and the polynomials u_s and u_a and the polynomials v_s and v_a share no common factor. For example consider Figure 6.5. There are two at most 8×8 grids, one covered by light grey and the other covered by dark grey stripes.

Let us assume we have such a bi-factorization. Then for (u_s, v_s) and for (u_a, v_a) we can proceed exactly like in the case of a silhouettecurve and a cutcurve described in the previous section. We perform a bi-factorization according to the multiplicities 1, 2, and ≥ 3. Again a polynomial of degree 8 can have at most two roots of multiplicity ≥ 3. According to the assumption that u_s, u_a and v_s, v_a have no common roots, the boxes belonging to multiplicity 1 can be handled with simple box hit counting. The boxes defined by roots of multiplicity 2 are solved with extended box hit counting, after we

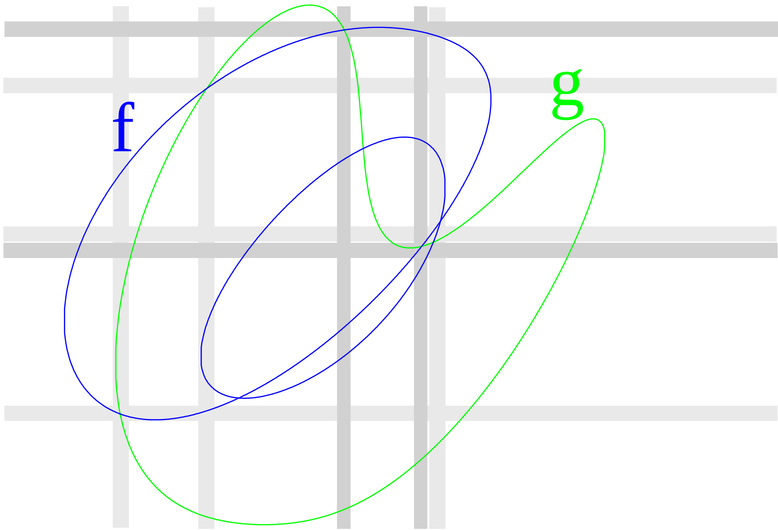

Figure 6.5: A bi-factorization of (X, Y) such that each bi-factor has degree ≤ 8

split off the singular points of f or g. The remaining bi-factor belonging to roots of multiplicity ≥ 3 defines at most 4 grid points that are solvable with explicit solutions.

What remains to do is to establish the bi-factorization $(X, Y) = (u_s, v_s) \cdot (u_a, v_a)$ such that each involved polynomial has degree ≤ 8. As for singular points, we can distinguish two different types of intersection points f and g have: *spatial* and *artificial*.

Definition 6.2: Let $(a, b) \in \mathbb{C}^2$ be an intersection point of two curves $f = \text{res}(p, q, z)$ and $g = \text{res}(p, r, z)$ with p, q, r being quadratic trivariate polynomials. We call (a, b) *spatial*, if for a root c of $p(a, b, z) \in \mathbb{C}[z]$ we have $p(a, b, c) = q(a, b, c) = r(a, b, c) = 0$. If $p(a, b, z) \in \mathbb{C}[z]$ has the two roots c and c' and it holds $p(a, b, c) = q(a, b, c) = 0$ and $p(a, b, c') = r(a, b, c') = 0$, then we call (a, b) *artificial*.

Spatial points are projected common intersection points of p, q, and r. Equivalently we can say that they are projections of common intersection points of the spatial intersection curves of p and q and of p and r. Of course also the intersection curve of q and r cuts through this common point. Artificial points are a result of the projection phase. One spatial intersection curve, for example the one of p and q, runs on the upper part of p, whereas the other one of p and r on its lower part. Both space curves are projected on top of each other causing an intersection point. For illustration have a look at Figure 6.6. The green and blue spatial intersection curves are as always the intersection curves of the red and the green and the red and the blue ellipsoids in our permanent example, respectively. They have two common points, marked by the arrows. Projecting both curves into the plane results in a green and a blue cutcurve

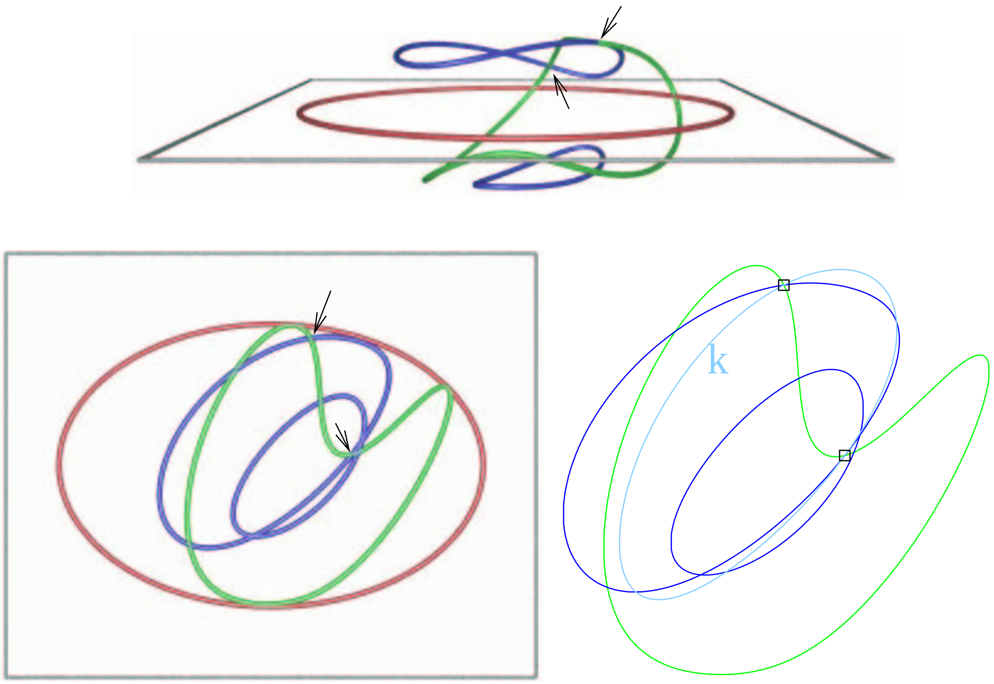

Figure 6.6: Projections of common intersection points of the red, green, and blue curve are called spatial.

that have 6 real intersection points: 2 spatial and 4 artificial. The light blue curve k is the projection of the intersection curve of the blue and the green ellipsoid. We know that it cuts through the spatial points. Of course it can happen that points are both spatial and artificial. For example if the spatial intersection curves of p and q and of p and r intersect transversally in a point (a, b, c) on the silhouette of p. Then the projected point (a, b) is a tangential intersection point of the projected curves. The multiplicity 2 is caused by the meeting of a spatial and an artificial point.

The question is how many spatial and artificial intersection points two cutcurves can have, counted with multiplicities. Both numbers have to be bounded by 8 in order to make this distinction suitable for our bi-factorization into polynomials of degree at most 8. In the next theorem we will show that this is really true:

Theorem 6.3: Two cutcurves have at most 8 spatial and at most 8 artificial intersections, counted with multiplicities.

Proof. Let as before $f = \text{res}(p, q, z)$ and $g = \text{res}(p, r, z)$ be the cutcurves. The bound for their spatial intersections immediately follows by the theorem of Bézout. Three quadrics in space can have at most 8 discrete common intersection points, counted with multiplicities.

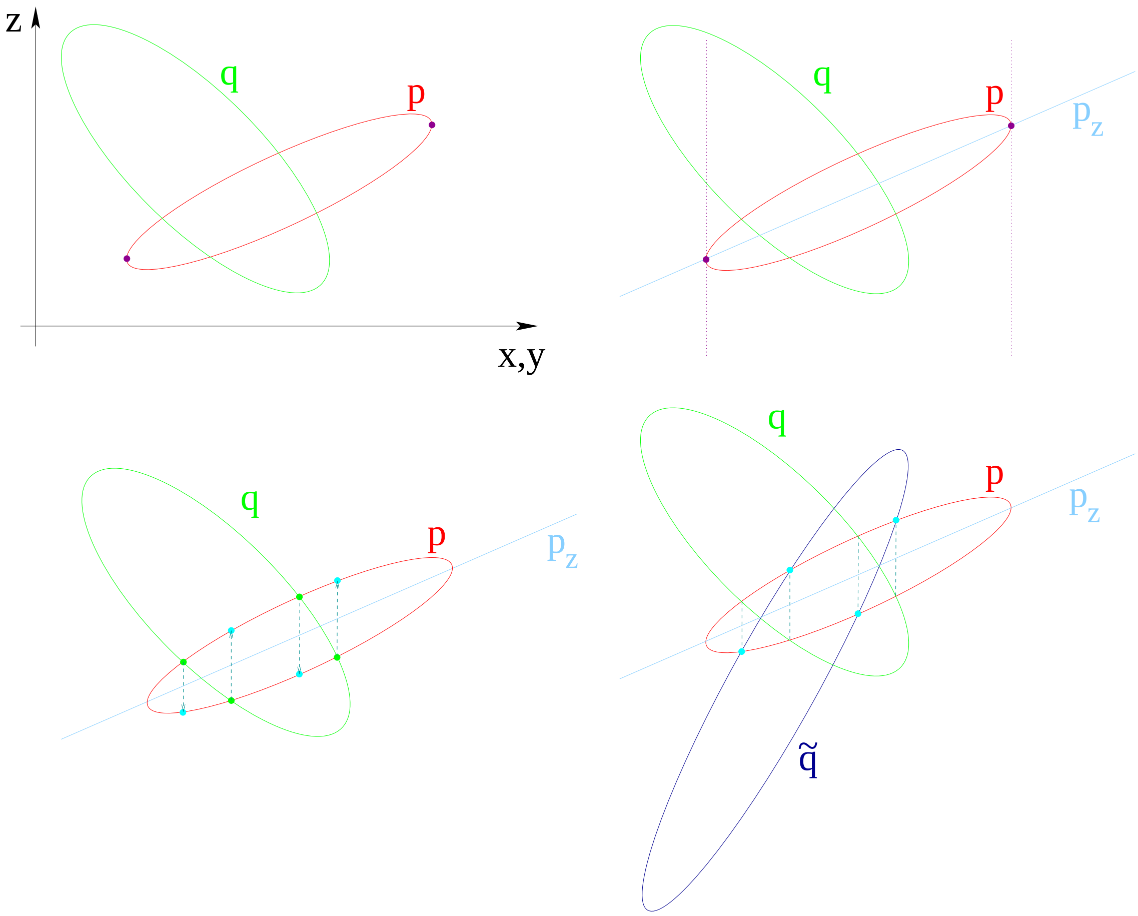

Figure 6.7: Projections of common intersection points of the red, green, and blue curve are called spatial.

The idea of proving the second bound is the following: We mirror q parallel to the z-axis at the plane $p_z = 0$. For illustration have a look at the four pictures in Figure 6.7. In the first picture, there are the two ellipsoids p and q and the coordinate axes. Assume that the z-axis is the vertical one. In the next picture, additionally p_z is shown, which exactly cuts through the points of vertical tangency of p and separates the lower part of p from its upper part. Now we take every point on the green ellipsoid q and move it vertically to the other side of p_z, keeping the vertical distance to p_z the same. Let $\tilde{q}$ be the resulting surface. We will show that $\tilde{q}$ again is a rational quadric, in our example again an ellipsoid. As one can see in the last picture, this transformation has the effect that all intersection points of p and $\tilde{q}$ have the same (x, y)-coordinates as the intersection points of p and q. That means $\text{res}(p, \tilde{q}, z) = \text{res}(p, q, z)$ holds. But the intersection points of p and q that lie on the top of p now lie on its bottom and vice versa. So the spatial and artificial intersections have changed place and we can again apply the theorem of Bézout to p, r and $\tilde{q}$.

Now here is the technical part of the proof. Without loss of generality let the polynomial p be of the form

$$p(x,y,z) \;=\; z^2 + p_1(x,y)\cdot z + p_0(x,y).$$

All quadrics we consider are generally aligned and therefore have a constant leading coefficient. We are looking for a function $f : \mathbb{C}^3 \to \mathbb{C}^3$ that mirrors points (a,b,c) vertically at the plane $p_z = 2z + p_1(x,y)$. We want a vertical translation of (a,b,c), so we move the point along a line l parallel to the z-axis through the point (a,b) on the (x,y)-plane. The intersection point α of l and p_z is the root of $p_z(a,b,z) = 2z + p_1(a,b) \in \mathbb{Q}[z]$: $\alpha = (a,b,-p_1(a,b)/2)$. The point (a,b,c) has to be moved along l to the other side of α, keeping the distance $|c + p_1(a,b)/2|$ the same. All in all we obtain

$$f(a,b,c) \;=\; (a,b,-p_1(a,b) - c).$$

If we apply this function to every point of a quadric

$$q(x,y,z) \;=\; z^2 + q_1(x,y)z + q_0(x,y),$$

this leads to

$$\begin{aligned}\tilde{q}(x,y,z) &= (-p_1(x,y) - z)^2 + q_1(x,y)\cdot(-p_1(x,y) - z) + q_0(x,y)\\ &= z^2 + (2p_1(x,y) - q_1(x,y))z + (p_1^2 - p_1q_1 + q_0)(x,y).\end{aligned}$$

This is again a quadric. For $q = p$ we obtain $\tilde{q} = p$. That means mirroring p vertically on p_z leads to p again. It remains to show for $q \neq p$ that

$$\mathrm{res}(p,q,z) \;=\; \mathrm{res}(p,\tilde{q},z).$$

We know that

$$\mathrm{res}(p,q,z) \;=\; (p_0q_1 - p_1q_0)(q_1 - p_1) + (p_0 - q_0)^2.$$

On the other hand we have

$$\begin{aligned}
\operatorname{res}(p,\tilde{q},z) &= (p_0\tilde{q}_1 - p_1\tilde{q}_0)(\tilde{q}_1 - p_1) + (p_0 - \tilde{q}_0)^2 \\
&= (p_1q_0 + p_1^2(p_1 - q_1) + p_0q_1 - 2p_1p_0)(p_1 + q_1 - 2p_1) \\
&\quad +(p_0 - q_0 - p_1(p_1 - q_1))^2 \\
&= ((p_0q_1 - p_1q_0) + (2p_1q_0 + p_1^2(p_1 - q_1) - 2p_1p_0))(q_1 - p_1) \\
&\quad +((p_0 - q_0) - p_1(p_1 - q_1))^2 \\
&= (p_0q_1 - p_1q_0)(q_1 - p_1) + (p_0 - q_0)^2 \\
&\quad +(-2q_0 - p_1(p_1 - q_1) + 2p_0)p_1(p_1 - q_1) \\
&\quad +2(q_0 - p_0)p_1(p_1 - q_1) + p_1^2(p_1 - q_1)^2 \\
&= \operatorname{res}(p,q,z).
\end{aligned}$$

This ends our proof. □

Now we know that there at most 8 spatial and artificial intersection points, we would like to compute a bi-factorization $(u_s, v_s)\cdot(u_a, v_a)$ of (X, Y) according to this distinction. We want the roots of u_s to be the x-coordinates of common intersection points of p, q, and r. A first idea is to additionally compute the resultant $k = \operatorname{res}(q, r, z)$ and to perform a greatest common divisor computation between X, $\operatorname{res}(f, k, y)$, and $\operatorname{res}(g, k, y)$. But caused by the projection from space to the plane it can happen that k cuts through an artificial intersection point of f and g. Then the x-coordinate of this artificial point would be a root of u_s, contradicting our goal. This would not disturb our following algorithm as long as the degree of u_s would still be at most 8. Otherwise, similar to the methods described in the next chapter, we could shear the spatial arrangement in order to remove this effect. But an alternative, correct, and even more efficient way to compute u_s is to use the results of [18]. There a method for computing u_s directly from the spatial quadrics p, q, and r with the help of multivariate resultants is provided. Of course all the considerations symmetrically hold for v_s. As a result we obtain a bi-factorization $(X, Y) = (u_s, v_s)(u_a, v_a)$ and the degree of each involved polynomial is bounded by 8. In the case the bi-factors u_s, u_a and v_s, v_a have no common root we already saw how to locate the intersection points of f and g.

The remaining problem is that the polynomials $u_c = \gcd(u_s, u_a)$ and $v_c = \gcd(v_s, v_a)$ might have a positive degree. In this case we handle the boxes defined by u_c and v_c separately. The roots of these polynomials are x- and y-coordinates of points that are spatial as well as artificial, respectively. By definition a point (a, b) is both spatial and artificial if for the roots c and c' of $p(a, b, z) \in \mathbb{C}[z]$ we have $p(a, b, c) = q(a, b, c) = r(a, b, c) = 0$ and $p(a, b, c') = q(a, b, c') = 0$ or $p(a, b, c') = r(a, b, c') = 0$. There are two situations that cause a point to be spatial and artificial. If we have $c = c'$, the common intersection point takes place on the silhouette of p. If we have $c \neq c'$, at least one of the planar curves has a top-bottom point in (a, b). We know how to compute explicit solutions for top-bottom points of a curve. So what remains to do is determining projections of common intersection points of p, q, and r on the silhouette of p. These points are projected common intersection points of p, q, r, and p_z or equivalently common intersection points of the quadratic curves $\text{res}(p, p_z, z)$, $\text{res}(q, p_z, z)$, and $\text{res}(r, p_z, z)$. The resultant of two quadratic curves is a polynomial of degree at most 4 and therefore common intersection points can be easily computed by simple box hit counting or by computing explicit solutions.

This finishes the proof of our main theorem. We considered planar arrangements that are the result of projecting intersection curves and silhouettes of spatial quadrics into the plane. We are able to determine every event point of such an arrangement exactly and unambiguously by only using rational arithmetic.

Chapter 7

Establishing the generality assumptions

In the previous chapters we made some assumptions on the location of the quadrics in space and the curves in the plane in order to simplify the argumentation. The assumptions were general alignment, squarefreeness, disjoint factorization, and general relation. A pair of bivariate polynomials that fulfills all these criteria we have called well-behaved. For each criterion we will present a method for detecting whether it is violated or not. We will give randomized shear algorithms that accomplish all assumptions with probability at least 1/2.

A troublesome side-effect of the resultant computation of two rational polynomials f_1 and f_2 with respect to a variable x_i is that there can be roots of the resultant which only appear because the leading coefficient of f_1 or f_2 vanishes. This problem does not occur if both polynomials are generally aligned, because then they have a constant leading coefficient. We will generate this condition for each trivariate and bivariate polynomial we consider.

A second problem that we encounter is that two polynomials f_1 and f_2 might share a common factor. We have seen in Proposition 2.17 that this is the case if and only if $\mathrm{res}(f_1, f_2, x_d) \equiv 0$. For example consider the two quadratic bivariate polynomials $f = h \cdot f_1$ and $g = h \cdot g_1$ where $h(x,y) = (x - y)$, $f_1(x,y) = (\frac{1}{2}x - y + 5)$, and $g_1(x,y) = (2x - y - 12)$. They define the red and the blue curve consisting of two lines each in Figure 7.1. The resultant $\mathrm{res}(f, g, y)$ is equal to the zero polynomial. It gives no information about event points in the planar arrangement, although there are interesting intersection points, namely the intersection points of h and f_1, h and g_1, and finally the ones of f_1 and g_1. We will show how to avoid this problem and achieve squarefreeness for each polynomial f and disjoint factorization for each two polynomials

f_1 and f_2 we consider.

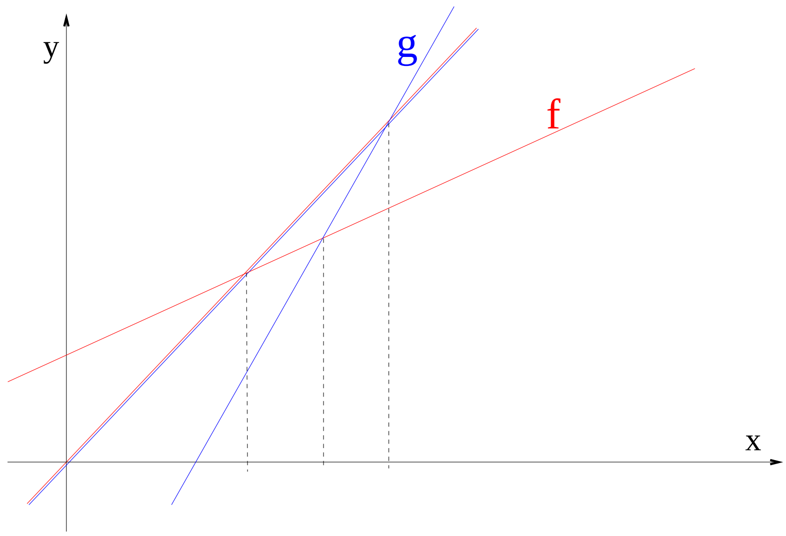

Figure 7.1: The resultant of f and g is equal to the zero polynomial

For a pair of planar curves f and g we want to draw conclusions from the multiplicities of the roots of the resultant $\mathrm{res}(f, g, y) \in \mathbb{Q}[x]$ to the kind of intersection point f and g have. Especially we want to state that a root of multiplicity at least 2 derives from a tangential intersection point of f and g. But if there are two intersection points with the same x-coordinate, then this conclusion is not valid. We will show how to overcome the problem of common x- or y-coordinates. Afterwards each pair of bivariate polynomials is in general relation.

At last a curve f may have vertical flat points that prevent characterizing its singular points by at most quadratic polynomials. We will see how to avoid such a situation.

Let P be the set of input quadrics and F the set of planar curves we obtain from the projection phase. The questions we have to answer are:

1. How can we achieve general alignment for each $p \in P$ when regarded as a polynomial in z?

2. Analogously we would like each $f \in F$ to have constant leading coefficients in x as well as in y.

3. What shall we do if the silhouettecurve $\mathrm{res}(p, p_z, z)$ for $p \in P$ or the resultant $\mathrm{res}(f, f_y, y)$ for a curve $f \in F$ is the zero polynomial? How can we achieve that each $p \in P$ and each $f \in F$ is squarefree?

4. How can we deal with $\operatorname{res}(p, q, z) \equiv 0$ or $\operatorname{res}(f, g, y) \equiv 0$ for $p, q \in P$ and $f, g \in F$? That means what can we do when two polynomials do not have a disjoint factorization?

5. When determining the intersection points of two planar curves, how can we manage to do the computation with both defining polynomials being in general relation?

6. If we detect a vertical flat point, what can we do to remove it?

7.1 General alignment for trivariate polynomials

Let $P \subset \mathbb{Q}[x, y, z]$ be a set of n quadrics. We want each trivariate polynomial to have a constant non-zero coefficient of z^2. This of course is easy to test by just examining the coefficient of z^2 for each $p \in P$. If the coefficient is non-constant for at least one p, we will shear all quadrics. The main idea of a shear method is described in [62].

Let us assume that we can find a vector $u = (u_1, u_2) \in \mathbb{Q}^2$ which fulfills the following property: The total degree of each input polynomial $p \in P$ equals the degree of the polynomial $p(u_1 \cdot z, u_2 \cdot z, z) \in \mathbb{Q}[z]$.

For such a vector u, we consider the shear function

$$\phi(x, y, z) \;=\; (x + u_1 \cdot z, y + u_2 \cdot z, z)$$

and compute $p \circ \phi$ for each $p \in P$. Substituting $x = x + u_1 z$ and $y = y + u_2 z$ in p cannot increase the total degree of p. We conclude

$$\deg(p) \geq \deg(p \circ \phi).$$

Analogously, substituting $x = 0$ and $y = 0$ in $p \circ \phi(x, y, z) = p(x + u_1 \cdot z, y + u_2 \cdot z, z)$ cannot increase the total degree of $p \circ \phi$:

$$\begin{aligned} \deg((p \circ \phi)(x, y, z)) &= \deg(p\ (x + u_1 \cdot z, y + u_2 \cdot z, z)) \\ &\geq \deg(p\ (u_1 \cdot z, u_2 \cdot z, z)) \\ &= \deg((p)(x, y, z)). \end{aligned}$$

We conclude that the affine transformation does not change the total degree of the polynomial at all:

$$\deg(p \circ \phi) \;=\; \deg(p).$$

By assumption we have $\deg(p) = \deg(p(u_1 z, u_2 z, z))$. So each sheared polynomial $p \circ \phi$ has a constant leading coefficient when regarded as a polynomial in z and that was our aim.

The remaining question is how to find the vector u. Let us look at the leading coefficients of the n input polynomials $p \in P$ when substituting (x, y, z) by $(u_1 z, u_2 z, z)$. The leading coefficients are at most quadratic polynomials in $\mathbb{Q}[u_1, u_2]$, let us call them $l_1, \ldots, l_n \in \mathbb{Q}[u_1, u_2]$. The polynomials l_i define at most quadratic planar curves and the point u we choose shall not lie on any of these curves. If we substitute $u_1 = a \in \mathbb{Q}$ in the polynomials l_i, $1 \le i \le n$, then it might happen that a leading coefficient l_j becomes the zero polynomial. This happens if $(u_1 - a)$ is a factor of l_j. But each l_i has total degree at most 2 and because of that this unwelcome effect can occur for at most $2n$ values of u_1. If we choose by random a value from the set $\{-2n, \ldots, 2n\}$, then with probability at least $\frac{1}{2}$ we have a good choice for u_1. Now we have found a value $u_1 = a$, we substitute it into the l_i and obtain n univariate polynomials $l_1(a, u_2), \ldots, l_n(a, u_2) \in \mathbb{Q}[u_2]$. Each $l_i(a, u_2)$ has at most 2 roots, so there are all in all $\le 2n$ values we are not allowed to choose for u_2. Again a random choice from the set $\{-2n, \ldots, 2n\}$ gives us a good value b for u_2 with probability $\ge \frac{1}{2}$. Of course we can increase the probability by choosing a larger range.

Now each input polynomial has a constant non-zero coefficient of z^2. The resultant of two polynomials p and q with respect to z defines a planar curve such that each point (a, b) on this curve gives rise to a, maybe complex, intersection point of p and q. The drawback is that we have to buy for this operation with a slightly larger coefficient size. This can have an impact on all following resultant and root isolation algorithms. The geometry of the spatial arrangement is not effected by the shear. Intersection points of quadrics remain intersection points. They only change their x- and y-coordinate.

7.2 General alignment for bivariate polynomials

We also want to have the property that all bivariate polynomials we consider have a constant non-zero leading coefficient with respect to each variable. Let F be the set of bivariate polynomials defining the curves in a planar arrangement. For all $f \in F$ the property can again be easily tested by examining the coefficients. If it is violated for at least one f, we will apply a shear to all curves in F. Of course another approach would be to go back to space, shear the quadrics, and start all the computation leading to the bivariate polynomials from the beginning.

Let us look at the polynomials in F as polynomials in the variable y. Like in the case for the trivariate polynomials, we assume that we can find a rational number $v \in \mathbb{Q}$ such that $\deg(f(x,y)) = \deg(f(v \cdot y, y))$ for all $f \in F$. Applying the affine transformation

$$\psi(x,y) \;=\; (x + v \cdot y, y)$$

to each polynomial $f \in F$ will result in a set of polynomials that have a constant leading coefficient of y.

For each $f \in F$ the leading coefficient of the polynomial $f(v \cdot y, y)$ is a polynomial in $\mathbb{Q}[v]$. It should not vanish for our choice of v. Let $l_1, \ldots, l_n \in \mathbb{Q}[v]$ be the leading coefficients. In our application, all bivariate polynomials $f \in F$ have total degree at most 4 and we conclude $\deg(l_i) \leq 4$ for $1 \leq i \leq n$. So all in all there are at most $4n$ different roots of the l_i which we shall not use for our choice of v. Again a random choice from the set $\{-4n, \ldots, 4n\}$ gives us a good value for v with probability $\geq \frac{1}{2}$.

We also want to have a constant leading coefficient of x. So in the same way, we randomly choose a rational number w and apply the shear

$$\tilde{\psi}(x,y) \;=\; (x, y + w \cdot x)$$

to each polynomial. It is easy to see that the constant leading coefficient of y is not effected by this shear and so the second shear does not destroy the effect of the first shear.

7.3 Squarefreeness

We want all polynomials $p \in P$ and $f \in F$ to be squarefree. Let us first consider the trivariate polynomials $p \in P$. If the resultant $\mathrm{res}(p, p_z, z)$ of an input quadric p becomes the zero polynomial, we know that $p = (p_z)^2$. The surface defined by p consist of two equal planes. We want to compute the topology of the arrangement. The result does not change if, instead of two identical planes, we consider only one of them. So we delete p from the set P of input quadrics and insert p_z.

We also want to establish squarefreeness for all polynomials $f \in F$. For a polynomial $f \in \mathbb{Q}[x,y]$ we recognize the lack of squarefreeness if the resultant of f and f_y with respect to y equals the zero polynomial. In this case we even know that f and f_y share a common *rational* factor h as shown in Proposition 2.17. In order to compute h, we apply the Euclidean gcd-algorithm using pseudo division to f and f_y as described in [42]. The polynomial $\tilde{f} = f/h$ is squarefree and has the same set of roots as f. We delete f from F and insert $\tilde{f}$. The new polynomial $\tilde{f}$ needs the attribute silhouettecurve or cutcurve. It inherits this attribute from f.

Of course if $\tilde{f}$ and h have a common factor, one could further factor $\tilde{f}$ with the help of h. This would lead to polynomials of smaller degrees and therefore faster computations. We omit the details.

7.4 Disjoint factorization

We have shown how to achieve squarefreeness for each polynomial $p \in P$ and $f \in F$. What about disjoint factorization for each pair of squarefree polynomials $p, q \in P$ and $f, g \in F$? We proceed the same way as shown in the previous section. If for two polynomials $f, g \in \mathbb{Q}[x_1, \ldots, x_d]$ their resultant with respect to any variable vanishes, we compute $h = \gcd(f, g)$. We obtain $f = h \cdot f_1$ and $g = h \cdot g_1$ where $f_1 = f/h$ and $g_1 = g/h$. By assumption f and g are squarefree and therefore also h, f_1, and g_1 are squarefree.

Now we remove f and g from P respectively F and insert the polynomials h, f_1, and g_1 in the case they are non-constant. We keep in mind that the polynomial f_1 derives from f, g_1 from g, and h from both f and g.

For $d = 3$ the defining polynomials of two input quadrics share a common non-constant rational factor if they are equal up to a non-zero constant factor or if they have a common factor of degree 1. Geometrically the polynomials define the same surface or both surfaces share a common plane. In the first case we delete one of the quadrics from the set P. In the second case we delete both quadrics and instead insert the three rational planes. In Chapter 3 we already mentioned that dealing with linear input polynomials is much easier than dealing with quadratic input polynomials.

In the plane the polynomials that will be inserted into F need the attribute silhouettecurve or cutcurve. If both polynomials f and g are cutcurves, then of course all factors inherit this attribute. If one curve f is the silhouettecurve, then f_1 and h get the attribute silhouettecurve and g_1 will be marked as a cutcurve.

7.5 General relation of two planar curves

In our algorithm we want to compute planar arrangements of curves we obtain from the projection phase. Let F be the set of curves of such a planar arrangement. In order to compute the arrangement we have to determine all intersection points between pairs of curves f and g with either $f, g \in F$ or $f \in F$ and $g = f_y$. So the main computation is the following: given two curves f and g, determine all their intersection points. The polynomials f and g are defined to be in general relation if there are no two common roots with the same x- or y-value. In this case we can draw conclusions from the multiplicity of a root of the resultant $\mathrm{res}(f, g, y)$ to the kind of intersection of f and g. How can we guarantee general relation of the polynomials?

The first observation is that common coordinates of intersection points are not intrinsic to the two curves. We can avoid it by choosing a different direction of projection or by equivalently shearing the curves. For example look at the red and the blue curve in the left picture of Figure 7.2. There a two pairs of points having the same x-coordinate and two pairs sharing the same y-coordinate. But if f and g are slightly sheared along the x-axis, like in the right picture, then all intersection points have different x-coordinates.

So if there are common x-coordinates, we will choose a shear in x-direction. If there are common y-coordinates, we do a shear along the y-axis. We have to choose the shears such that afterwards the curves are in general relation, but without destroying general alignment. Then we compute the intersection points of the sheared polynomials as described in the previous chapters. As a result we obtain rational boxes around each real intersection point. By applying the inverse shears we transform the curves together with the computed boxes back to the starting position. The sheared boxes are no longer parallel to the coordinate axes. So for each sheared box we have to compute another one parallel to the coordinate axes such that both boxes contain exactly the same intersection point. Proceeding this way we only have to shear a pair of curves that is not in general relation. We avoid shearing the whole planar arrangement.

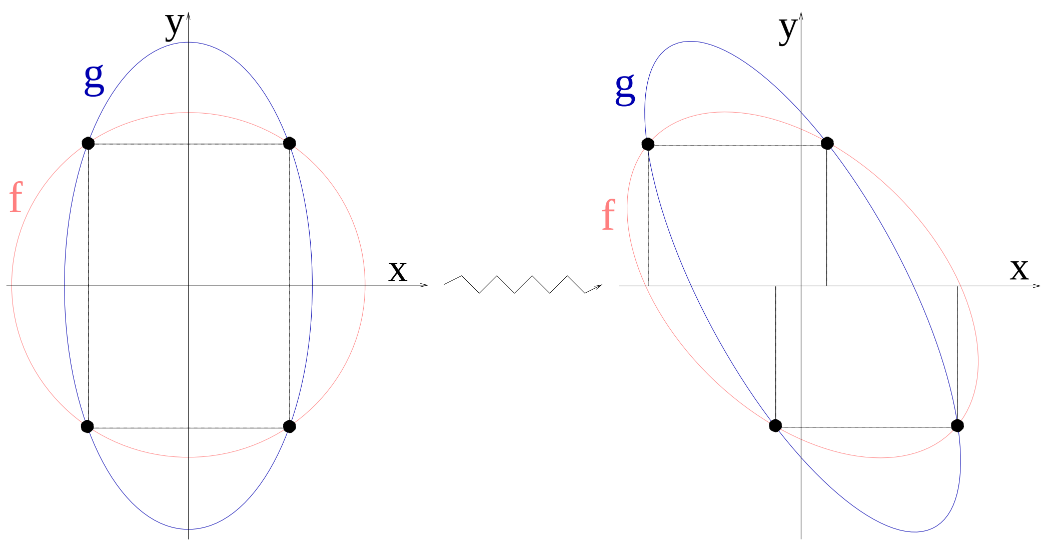

Figure 7.2: A shear along the x-axis

Before going into details on how to test for general relation and how to choose the shears, we will first do some helpful remarks on the way the multiplicities of the roots of $\mathrm{res}(f, g, y)$ change when we shear the two curves along the x-axis. A similar algorithm to the one we we will explain is described in [65].

7.5.1 Multiplicity of intersection points

We will prove that every intersection point (a, b) is responsible for a fixed multiplicity i of the root $x = a$ of the resultant. This number i is not effected by a shear. If there are exactly two intersection points (a, b) and (a, b') at $x = a$ with associated multiplicities i and j, then the multiplicity of a as a root of the resultant is equal to $i + j$. It makes sense to definite the following (remember also Corollary 4.14):

Definition 7.1: Assume that for two curves f and g no two intersection points share the same x-coordinate. Furthermore let f and g have constant leading coefficients when regarded as polynomials in y. Then we define the *multiplicity* of a common root $(a, b) \in \mathbb{C}^2$ of f and g to be the multiplicity of the root a of $\mathrm{res}(f, g, y)$.

We will show next that the multiplicity of an intersection point is invariant under a shear along the x-axis

$$\psi(x, y) \;=\; (x + vy, y).$$

Proposition 7.2: Assume that for two curves f and g no two intersection points share the same x-coordinate. Furthermore let f and g have constant leading coefficients when regarded as polynomials in y. Let $v \in \mathbb{R}$ such that both assumptions also hold for

$$\begin{aligned} f^v(x, y) &= f(x + vy, y) \\ g^v(x, y) &= g(x + vy, y). \end{aligned}$$

Then the multiplicity of a common root $(a, b) \in \mathbb{C}^2$ of f and g is equal to the multiplicity of $(a - vb, b)$ of f^v and g^v.

Proof. Let $(a_1, b_1), \ldots, (a_l, b_l)$ be all complex intersections points of f and g and let $i_1, \ldots, i_l$ be their corresponding multiplicities. Then by definition we have

$$X \;=\; \mathrm{res}(f, g, y) \;=\; (x - a_1)^{i_1} \cdot \ldots \cdot (x - a_l)^{i_l}.$$

We regard f^v and g^v as polynomials in the three variables x, y, and v. By assumption the polynomials $f = f^0$ and $g = g^0$ have constant non-zero leading coefficients with respect to the variable y. Then the leading coefficients of f^v and g^v are polynomials in

v with only finitely many roots. Let R be the set of these roots. By definition of the resultant of two multivariate polynomials we obtain for all $v_0 \notin R$

$$X^{v_0} \;=\; \mathrm{res}(f^{v_0}, g^{v_0}, y) \;=\; \mathrm{res}(f^v, g^v, y)|_{v=v_0} \;=\; (x - a_1 + v_0 b_1)^{j_1(v_0)} \cdot \ldots \cdot (x - a_l + v_0 b_l)^{j_l(v_0)}.$$

It remains to prove $i_1 = j_1(v_0), \ldots, i_l = j_l(v_0)$ for all $v_0 \notin R$.

Of course for some v_0 it can happen that $a_i - v_0 b_i = a_j - v_0 b_j$ for some $1 \leq i, j \leq l$. But this set of points is obviously finite and we denote it by S. For all $v_0 \notin S$ the two curves defined by the polynomials f^{v_0} and g^{v_0} have no two intersection points with the same x-coordinate or equivalently the roots $a_i - v_0 b_i$ of the resultant are pairwise different.

The sets R and S are finite and we conclude that $R \cup S$ is finite. Due to our assumptions that $f = f^0$ and $g = g^0$ have constant leading coefficient and no two intersection points share the same x-coordinate we have $0 \notin R \cup S$. Then there exists a $\delta > 0$ such that for all $|v_0| < \delta$ we have $v_0 \notin R \cup S$. That means for all small v_0 we have $v_0 \notin R$ and therefore

$$X^{v_0} \;=\; (x - a_1 + v_0 b_1)^{j_1(v_0)} \cdot \ldots \cdot (x - a_l + v_0 b_l)^{j_l(v_0)}.$$

Moreover for every $|v_0| < \delta$ we have $v_0 \notin S$ and all roots $a_i - v_0 b_i$, $1 \leq i \leq l$, are pairwise different. For $v_0 = 0$ we know $j_1(0) = i_1, \ldots, j_l(0) = i_l$. The resultant X^v is as a polynomial in x and v continuous in both variables. We conclude $i_1 = j_1(v_0), \ldots, i_l = j_l(v_0)$ for all $v_0 < \delta$. But then of course it must hold for every $v_0 \notin R$. $\square$

For illustration have a look at Figure 7.3.

In the first picture there are two plotted curves. Namely the ones defined by the polynomials

$$\begin{aligned} f &= y^2 + x^2 - 2xy + 2x - 7 \\ g &= y^2 + 7x^2 - 8xy + x - y + 1/8. \end{aligned}$$

On the x-axis the roots of the resultant $\mathrm{res}(f, g, y) = 72x^3 - 224x^2 - 221/2x + 2801/64$ are marked by the filled dots. Now let us look how these dots move when we shear the curves along the x-axis. The sheared polynomials are of the form

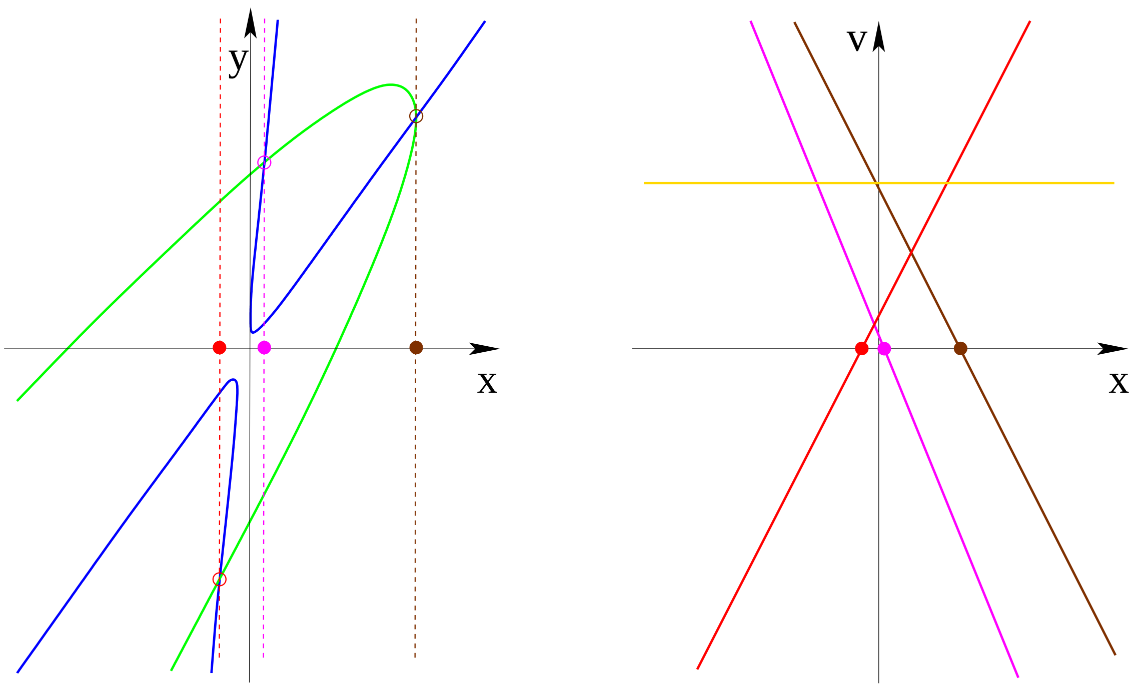

Figure 7.3: The x-coordinates of intersection points of two curves change linearly during the shear.

$$\begin{aligned} f^v &= y^2 + (x+vy)^2 - 2(x+vy)y + 2(x+vy) - 7 \\ &= (v^2 - 2v + 1)y^2 + (2vx - 2x + 2v)y + (x^2 + 2x - 7) \\ g^v &= y^2 + 7(x+vy)^2 - 8(x+vy)y + (x+vy) - y + 1/8 \\ &= (1 + 7v^2 - 8v)y^2 + (14vx + v - 1)y + (7x^2 - 7x + 1/8). \end{aligned}$$

In the second picture the bivariate polynomial $\mathrm{res}(f^v, g^v, y) \in \mathbb{Q}[x, v]$ is plotted. At each horizontal line one can see the x-coordinates of the intersection points at one time during the shear. We have chosen a linear shear and so the roots of the resultants move along lines. Of course, if an intersection point has multiplicity i, then the complex polynomial defining its moving-line is a factor of multiplicity i of $\mathrm{res}(f^v, g^v, y)$. There are intersection points of these lines when the x-coordinates of intersection points coincide. In our example there is a horizontal line for $v_0 = 1$. This line occurs because the leading coefficients of f^v and g^v when regarded as polynomials in y vanish for $v = 1$.

We will see in the next section that these considerations enable us to test whether there are two intersection points with the same x-coordinate.

7.5.2 Testing for general relation

We want to develop a test that answers the question whether two curves f and g have two distinct intersections with the same x-coordinate. Of course a first and easy test is

$$\gcd(\mathrm{res}(f,g,y), \mathrm{sres}_1(f,g,y)) \;=\; \text{constant ?}$$

If the answer is "yes", then there exists no $a \in \mathbb{C}$ with $f(a,y)$ and $g(a,y)$ having two common roots. That means the two curves are in general relation and nothing has to be done. But if the answer is negative we unfortunately cannot conclude that there are two intersection points with the same x-coordinate. Look at the two examples in Figure 7.4. If the resultant and the first subresultant have a common root a, then we know that $f(a,y)$ and $g(a,y)$ have a common factor $c(y)$ of degree at least 2. In the first picture we have $c(y) = (y-b_1)(y-b_2)$ and the intersections points (a,b_1) and (a,b_2) have the same x-coordinate. In the second picture we have $c(y) = (y-b)^2$ and the curves are in general relation.

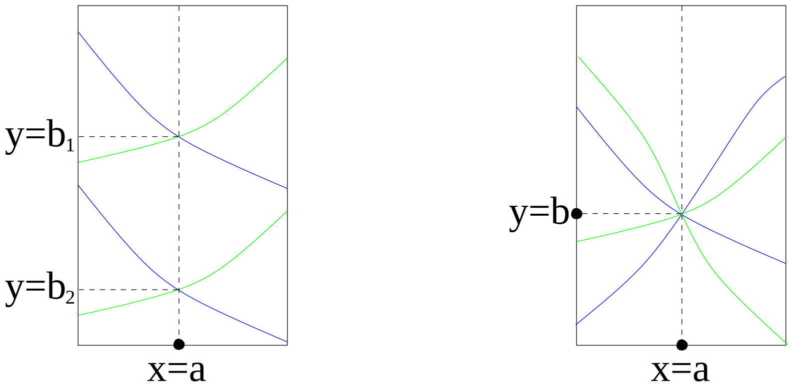

Figure 7.4: A degenerate case and a non-degenerate one, both effecting the vanishing of the first subresultant

So a non-constant gcd of the resultant and the first subresultant does not permit decision. But we saw in the last section that each intersection point has a fixed multiplicity that does not change when we shear the arrangement. For the two curves f and g we again define the polynomials

$$\begin{aligned} f^v(x,y) &:= f(x+vy,y) \\ g^v(x,y) &:= g(x+vy,y) \end{aligned}$$

and regard them as polynomials in $\mathbb{Q}[x,y,v]$. Of course $f^0 = f$ and $g^0 = g$. The resultant $\text{res}(f^v, g^v, y) \in \mathbb{Q}[x, v]$ defines an arrangement of lines the x-coordinates of the intersection points of f and g walk along during the shear. That means the resultant factors over $\mathbb{C}$ in linear parts $l_1, \ldots, l_k \in \mathbb{C}[x, v]$ with k being an upper bound on the number of intersection points of f and g:

$$\text{res}(f^v, g^v, y) \;=\; l_1^{i_i} \cdot \cdots \cdot l_k^{i_k}.$$

For each intersection point $(a, b) \in \mathbb{C}^2$ of f and g there is an index j such that during the shear a moves along the line l_j. The multiplicity of l_j in the factorization of the resultant is exactly the multiplicity of (a, b) as we have shown in the proof of Proposition 7.2. The two curves $f = f^0$ and $g = g^0$ have no two intersection points with common x-coordinates if and only if there is no intersection point of two of the lines on the x-axis. This is equivalent to the statement that there are no factors l_a and l_b in the complex factorization of $\text{res}(f^v, g^v, y)$, $a \neq b$, with $l_a(x, 0) = l_b(x, 0)$. The polynomials f and g are both generally aligned and because of that

$$(\text{res}(f^v, g^v, y))|_{v=0} \;=\; \text{res}(f, g, y).$$

There is an intersection point on the x-axis if and only if substituting $v = 0$ in the factorization of the resultant $(l_1^{i_i} \cdot \cdots \cdot l_k^{i_k})|_{v=0}$ differs from the factorization over $\mathbb{C}$ of $\text{res}(f, g, y)$.

In practice we cannot perform a complex factorization of $\text{res}(f^v, g^v, y)$. But we can compute its multiplicity factorization over $\mathbb{Q}$ analogously to the one described in Chapter 4 for univariate polynomials. Only gcd-computations are performed and they can also be realized for bivariate polynomials using pseudo-division as mentioned before. And of course we can do a multiplicity factorization of the univariate polynomial $\text{res}(f, g, y)$. With our previous remarks it is easy to see that f and g are in general position if and only if both rational factorizations are equal.

Now we have a test that answers the question of disjoint x-coordinates of all intersection points. What can we do if our test gives the answer "no, f and g are not in general relation with respect to the x-coordinates"?

7.5.3 Shearing the pair of planar curves

The remaining question is how to choose a shear in order to establish general relation with respect to the x-coordinates and at the same time not to destroy general alignment. For our two curves f and g we are looking for a v_0 such that

1. all leading coefficients of f^v and g^v regarded as polynomials in y do not vanish for v_0 and
2. no two intersections points of f^v and g^v have the same x-coordinate.

In our application the polynomials f^v and g^v have total degree at most 4. The leading coefficients of them are polynomials in v of degree at most 4. The curves f and g can have at most 16 intersection points. So there are at most $\binom{16}{2} = 120$ different shears that make two of them get the same x-coordinate. All in all there are at most $4 + 120 = 124$ values which should not be chosen for v_0. By randomly taking a value from the set $\{-124, \ldots, 124\}$ we have a good one with probability at least 1/2.

We also want to establish different y-coordinates for all intersection points and analogously a random shear $\tilde{\psi}$ from the set $\{-124, \ldots, 124\}$ realizes it with probability 1/2. By definition it is clear that the shear ψ prevents the $<$ relation of the y-coordinates of all points. Analogously $\tilde{\psi}$ prevents the $<$ relation of the x-coordinates and therefore the second shear does not destroy the effect of the first shear. That is why applying both affine transformations results in polynomials that are in general relation.

7.6 Avoiding vertical flat points

At last we have to consider the situation that for a curve f the bi-resultant of f and f_y cannot be factored into quadratic bi-factors because of the existence of vertical flat points. Every vertical flat point of f is an intersection point of f and $f_1 = f_{xx}f_y^2 - 2f_xf_yf_{xy} + f_{yy}f_x^2$. Because $\deg(f) \leq 4$ and $\deg(f_1) \leq 8$ we conclude that f and f_1 can have at most 32 intersection points. We use the notation introduced in the last section. We want to determine a shear factor v_0 such that

1. the leading coefficient of f^v regarded as polynomials in y does not vanish for v_0,
2. no two intersections points of f^{v_0} and $f_y^{v_0}$ have the same x-coordinate, and
3. f^{v_0} has no vertical flat point.

There are at most $4 + \binom{12}{2} + 32 = 102$ shears that violate one of the conditions. By randomly taking a value for v_0 from the set $\{-102, \ldots, 102\}$, we have a good one with probability at least 1/2. The shear prevents the $\leq$ relation of the y-coordinates and therefore has no effect on the general relation concerning the y-coordinates.

Chapter 8

The convex hull of ellipsoids

Computing the convex hull of a set of points in any dimension d is a well known area of research, see for example the overview article of Seidel [68]. A few results are about the convex hull of algebraic curves in $\mathbb{R}^2$, see for example [8] and [25]. But very little is known about the convex hull of algebraic surfaces and surface patches in space. In 1993 Boissonat et al. [13] discovered a method for computing the convex hull of a set of n spheres even in d-dimensional space. They showed that the same method applies to a set of n homothetic convex objects in $\mathbb{R}^d$ of constant complexity. In the previous chapters we have presented a method of computing the topology of a cell in an arrangement of quadric surfaces. We will show how to use this result via duality for computing the topology of the convex hull of n ellipsoids in space. The dual correspondence between the convex hull of quadric surface patches and computing the arrangement of quadrics has first been considered by Hung and Ierardi [39]. For an animated illustration of our algorithm have a look at the video [32].

An ellipsoid is defined by a quadratic polynomial in $\mathbb{Q}[x_1, \ldots, x_d]$ with special properties on the coefficients:

Definition 8.1: In d-dimensional space an *ellipsoid* e is the set of roots of the polynomial

$$e(x) \;=\; (x-c)^T M (x-c) - 1 \;\in\; \mathbb{R}[x_1, \ldots, x_d].$$

Thereby M is a positive definite and symmetric real $d \times d$ matrix and $c \in \mathbb{R}^d$ is a vector of translation.

We are interested in the convex hull of a set of ellipsoids defined by rational polynomials. So we only consider ellipsoids for which all entries of M and c are rational numbers. Generally, the convex hull of a point set is defined in the following way:

Definition 8.2: A set of points $X \subset \mathbb{R}^d$ is said to be *convex*, if for every two points

$x_1, x_2 \in X$ all the points

$$\lambda_1 x_1 + \lambda_2 x_2$$

with $\lambda_1 \geq 0$, $\lambda_2 \geq 0$, $\lambda_1 + \lambda_2 = 1$ also belong to X.
If X is any set of points in $\mathbb{R}^d$, then by the *convex hull* of X, denoted $CH(X)$, we mean the set of points in the intersection of all the convex sets that contain X.

By induction we immediately derive the following well known characterization for the convex hull of a set of points:

Proposition 8.3: Let $X \subset \mathbb{R}^d$. For any finite set $\{x_1, \ldots, x_k\} \subset X$ and parameters $\lambda_1, \ldots, \lambda_k \geq 0$ with $\lambda_1 + \cdots + \lambda_k = 1$, the convex hull of X must contain the point $\lambda_1 x_1 + \cdots + \lambda_k x_k$:

$$CH(X) = \{\lambda_1 x_1 + \cdots + \lambda_k x_k \mid \{x_1, \ldots, x_k\} \subset X, \lambda_i \geq 0, \sum_{i=1}^{k} \lambda_i = 1\}.$$

The convex hull of one ellipsoid e consists of all points $x \in \mathbb{R}^d$ for which $e(x) \leq 0$. So computing the convex hull of a set of ellipsoids $e_1, \ldots, e_n$ is equivalent to computing the convex hull of the sets $X_i = \{x \mid e_i(x) \leq 0\}$ for $1 \leq i \leq n$. Each X_i is closed and bounded.

Proposition 8.4: The convex hull of a set $X = \{X_1, \ldots, X_n\}$ of closed bounded sets $X_i \subset \mathbb{R}^d$ consists of all points that are in the intersection of all closed halfspaces containing X.

A halfspace is bounded by a plane. Each such plane h is uniquely determined by three real numbers $x_h = (a, b, c) \in \mathbb{R}^3$, namely the ones for which $h = \{(x, y, z) \mid z = ax + by + c\}$. For each plane h tangential to the convex hull, we define a vector $v^h = (v_1^h, \ldots, v_n^h) \in \{0, 1\}^n$ with

$$v_i^h = \begin{cases} 1 & \text{, if } h \text{ has a tangential intersection point with ellipsoid } e_i \\ 0 & \text{, otherwise.} \end{cases}$$

We say that two tangential planes h' and h'' are *equivalent* if $v^{h'} = v^{h''}$. Moreover we define them to be *strongly equivalent*, if additionally there exists a continuous function

$$\phi : [0, 1] \to \mathbb{R}^3 \quad , \quad \phi(0) = x_{h'} \quad , \quad \phi(1) = x_{h''}$$

such that h' is equivalent to all planes represented by $\phi(i)$, $i \in [0, 1]$.

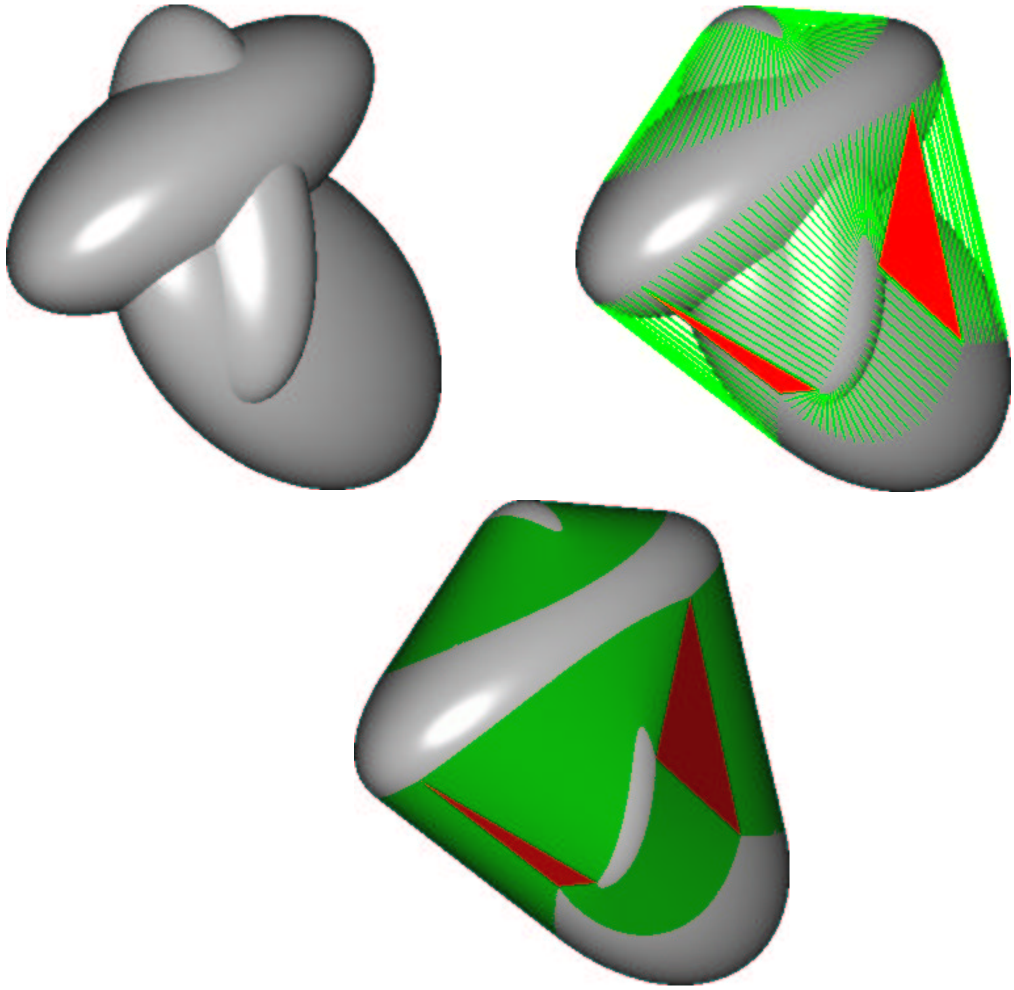

Figure 8.1: The convex hull of the three ellipsoids contains red triangular planar facets, green tunnel facets, and grey elliptical facets.

The boundary of the convex hull of a set of n ellipsoids $\{e_1, \ldots, e_n\}$ in 3-dimensional space consists of 2-dimensional facets, 1-dimensional edges, and 0-dimensional vertices. With the help of the equivalence relation we can distinguish three different kinds of facets:

1. A *planar facet* is embedded on one plane h to which no other plane is strongly equivalent. In this one-elementary case, h has at least 3 tangential intersection points with different ellipsoids. The planar facet consists of all points in the relatively open convex hull of these intersection points.

2. A *tunnel facet* is build by a set of strongly equivalent planes, each contributing a line segment to the convex hull. The endpoints of each line segment are tangential

intersection points of the plane and the ellipsoids. The relatively open union of all these line segments forms the tunnel facet.

3. An *elliptical facet* is defined by a set of strongly equivalent planes, each plane intersecting exactly one ellipsoid. The relatively open union of the intersection points lies on one ellipsoid building an elliptical facet.

For illustration let us have a look at Figure 8.1. The convex hull of the three ellipsoids consists of red planar facets, green tunnel facets, and grey elliptical facets.

We have defined facets to be relatively open sets of points. An edge is a maximal open connected subset of points in the intersection of the closure of two facets. Analogously a vertex is a common point in the closure of two edges.

8.1 A short remark on the lower bound of the complexity

Before applying our previous results to compute the convex hull of a set of ellipsoids, we will first make a short remark on the topological complexity of the convex hull of ellipsoids in space. We will show that $3n$ disjoint ellipsoids can have a convex hull of complexity $\Omega(n^2)$, measured without loss of generality in the number of facets.

Theorem 8.5: The convex hull of $3n$ disjoint ellipsoids can have $\Omega(n^2)$ facets.

The construction we work out for the proof seems to be extendible to arbitrary dimension d. This would lead to the statement that the convex hull of $n(2^{d-1}-1)$ disjoint ellipsoids can have $\Omega(n^{d-1})$ facets.

In the following we will interpret points as ellipsoids: degenerate ellipsoids with infinitesimally small extension. The same way we say that a $(d-1)$-dimensional ellipsoid is also a d-dimensional ellipsoid with infinitesimally small extension in x_d-direction. This does not fit algebraically in our definition of ellipsoids, but topologically this is no loss of generality.

8.1.1 A lower bound in 3-space

Our constructive proof works in the spirit of the one for a set of spheres described in [13]. We choose n points on a circle c, see the dashed circle in Figure 8.2. The convex hull of these n disjoint ellipsoids is a polygon with n sides. We move one copy of this polygon along the z-axis in positive direction and one copy in negative direction. This gives us $2n$ points in $\mathbb{R}^3$, the convex hull of which is a cylinder with n lateral facets, one top facet, and one bottom facet.

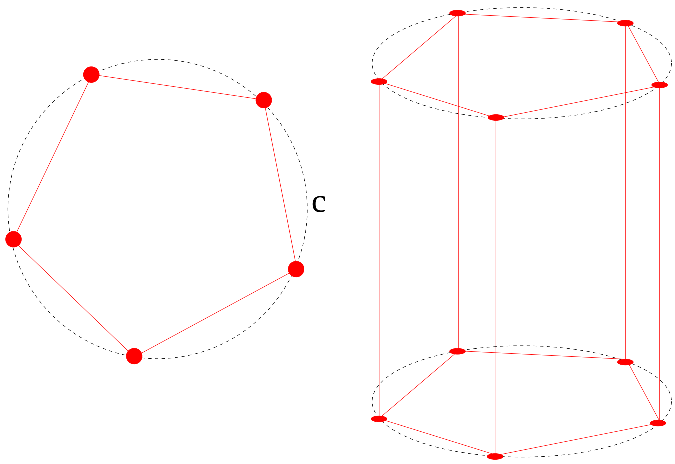

Figure 8.2: Construction of a cylinder with $n + 2$ facets

We perform an auxiliary construction. Imagine a 3-dimensional ellipsoid e with center at the origin and orthogonal axes such that e intersects the (x, y)-plane in the circle c. While building the convex hull, we will not take account of e. Of course, e breaks through every facet of the cylinder near the (x, y)-plane and then, in both directions along the z-axis, continuously dives into the cylinder, see Figure 8.3.

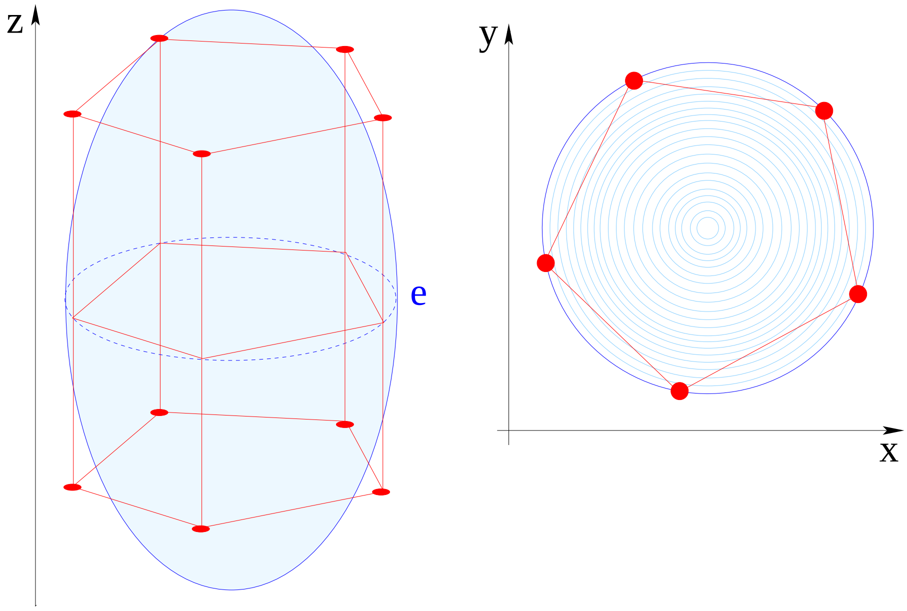

Figure 8.3: The auxiliary ellipsoid e and the cylinder: side view and top view

Inscribe n tangential ellipsoids to e with center on the strictly positive z-axis such that

1. the intersection of each ellipsoid with a plane parallel to the (x, y)-plane is a circle

and

2. each ellipsoid breaks through every facet of the cylinder. That means each intersection of an ellipsoid with a facet is an ellipse.

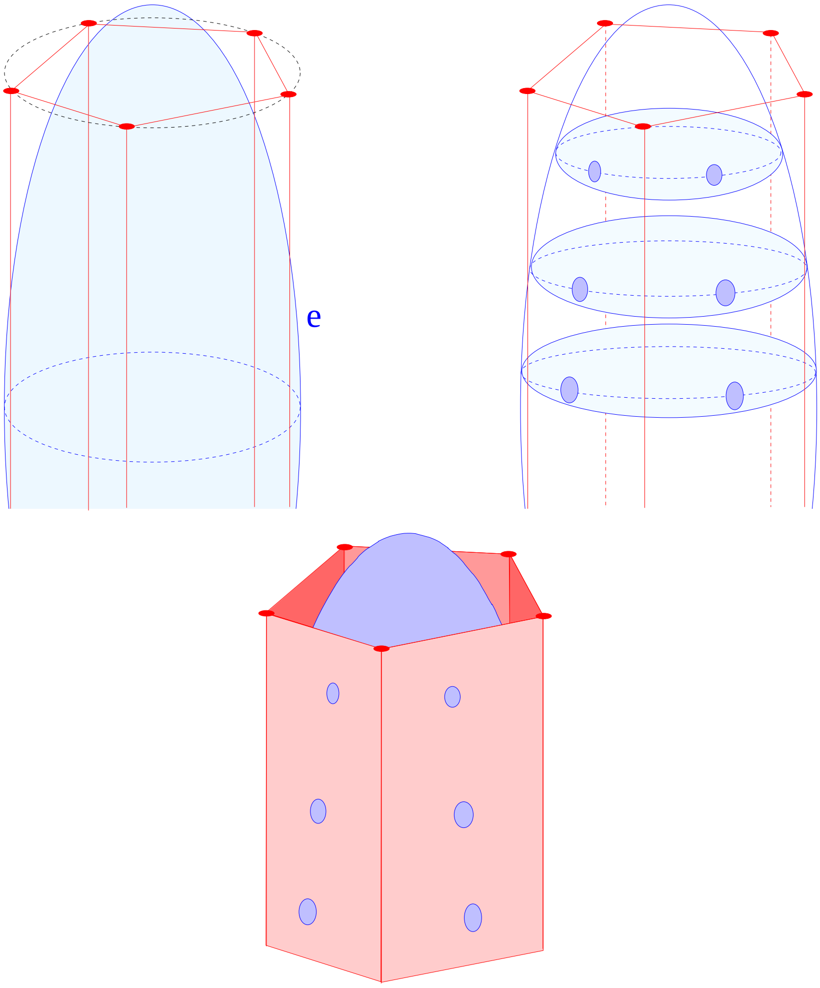

Figure 8.4: Each inscribed ellipsoid cuts through every face of the cylinder

For illustration have a look at Figure 8.4.

Cutting the convex hull of the $3n$ ellipsoids parallel to the (x, y)-plane through the center of one of the inscribed ellipsoids yields a situation as in Figure 8.5. The convex hull consists of n vertices belonging to the edges of the cylinder, n arcs that are parts of the ellipsoid, and $2n$ connecting segments. Every arc contributes to the 3-dimensional convex hull as a part of an elliptical facet. This is guaranteed by tangentially inscribing the ellipsoids in the convex ellipsoid e. So every inscribed ellipsoid contributes

$\Omega(n)$ facets to the convex hull. Notice that the elliptical facets and the tunnel facets connecting two inscribed ellipsoids are entirely contained in the auxiliary ellipsoid e.

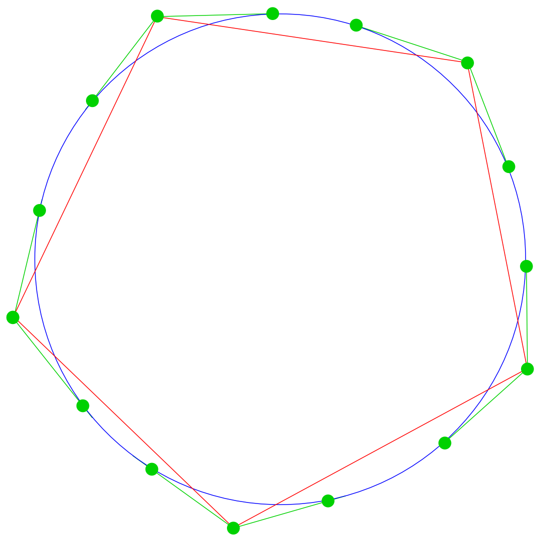

Figure 8.5: The convex hull cut by a plane parallel to the (x, y)-plane

All in all we have arranged $3n$ disjoint ellipsoids. Each of the n inscribed ellipsoids contributes $\Omega(n)$ facets to the convex hull. So the convex hull of this set has the complexity $\Omega(n^2)$.

8.2 Reduction by means of Duality

We want to construct the incidence graph of the convex hull of a set of n ellipsoids in 3-dimensional space. We solve this problem by means of duality. We reduce it to the problem of constructing the intersection cell containing the origin in an arrangement of n ellipsoids, paraboloids, and two sheet hyperboloids. In the previous chapters we have shown how to perform the basic algebraic primitives in order to compute a cell in an arrangement of quadrics.

There are some terms that need to be explained. First we define what we mean by the dual of a set of points, see also [27]. Next we show that the convex hull of a point set and the intersection cell of the dual set are topologically equivalent. At least we show how to dualize a set of ellipsoids.

8.2.1 The dual of a point set

We want to dualize sets of points:

Definition 8.6: Let $X \subset \mathbb{R}^d$ be a set of points. We define the dual of X to be the set

$$X^* = \{y \in \mathbb{R}^d \mid x^T y \leq 1 \ \forall x \in X\}.$$

If X consists of a single point not equal to zero, then X^* is a closed halfspace. The dual of $X = \{0\}$ is the whole space $X^* = \mathbb{R}^d$. Since the dual of $\cup_{i \in I} X_i$ is the intersection of the duals of the sets X_i, it follows that for any X, X^* is a convex set. It is clear from definition that for every set X the dual contains the origin and that if $y \in X^*$ so does λy, $0 \leq \lambda \leq 1$.

By proceeding in a similar way, we can define the dual of X^* as a subset $X^{**} \in \mathbb{R}^d$. We are interested in the relationship between X and X^{**} on the one hand and X^* and X^{***} on the other hand:

Theorem 8.7: For a subset $X \subset \mathbb{R}^d$ let $\overline{X}$ denote its closure. The following equalities hold for X:

$$X^{**} = CH(\overline{X} \cup \{0\}) \quad \text{and} \quad X^{***} = X^*.$$

Proof. The proof proceeds in two steps. In a first one we will show that the dual sets of X and of $CH(\overline{X} \cup \{0\})$ are equal. In a second one we proof the equality $X^{**} = X$ for every closed convex set X that contains 0. The theorem follows easily from these two statements.

1. $X^* = CH(\overline{X} \cup \{0\})^*$
   The inclusion "$\supset$" is clear. For "$\subset$" let y be a point of X^*, that means $x^T y \leq 1$ for all $x \in X$. We have to prove that $x^T y \leq 1$ for all $x \in CH(\overline{X} \cup 0)$ or equivalently

   $$\lambda_1 x_1^T y + \cdots + \lambda_k x_k^T y \leq 1$$

   for all $\{x_1, \ldots, x_k\} \subset \overline{X} \cup \{0\}$, $\lambda_i \geq 0$, $\sum_{i=1}^k \lambda_i = 1$, according to Proposition 8.3. The inequality is true for $\{x_1, \ldots, x_k\} \subset X$ by assumption. If some of the $x_i \in X$ are replaced by 0, then the inequality certainly remains true. How about the limit points of X? Let p be a limit point of X: $p \in \overline{X} \backslash X$. Since there is no $z \in \mathbb{R}^d$ for which $x^T z \leq 1$ for all $x \in X$ and $p^T z > 1$, this is especially valid for $z = y$ for our special point y. We conclude $p^T y \leq 1$ and this proves the first assumption.

2. $X^{**} = X$ for a closed convex set X that contains 0
   If $x \in X$, then $x^T y \leq 1$ for all $y \in X^*$ and so $x \in X^{**}$ and this proves $X^{**} \supset X$. If on the other hand $a \in \mathbb{R}^d$ does not belong to X, there exists a hyperplane that strictly separates a from X. This hyperplane does not contain the origin since the origin lies in X. Then we can find $y \in \mathbb{R}^d$ such that

   $$x^T y < 1 \ \ \forall x \in X \quad \text{and} \quad a^T y > 1.$$

We conclude $y \in X^*$ and thus $a \notin X^{**}$. It follows $X^{**} \subset X$ and we have $X^{**} = X$.

□

It is sometimes convenient to regard duality as one between points and hyperplanes. If $x \neq 0$ is a point of $\mathbb{R}^d$, then the dual set x^* is a halfspace bounded by a hyperplane which may be denoted by h_x: $h_x(x, y) = x^T y - 1$. We can say that this hyperplane is *dual* to x. Similarly, if h is a hyperplane in $\mathbb{R}^d$, its dual is a segment joining 0 to a point x of $\mathbb{R}^d$, unless h passes through the origin. In the latter case its dual is an infinite line. If we exclude this last case, then we can say that x is the dual of h and write it as x_h.

If X is a closed bounded and convex set with the origin as an interior point, then the points on its boundary are dual to tangential hyperplanes to X^* and vice versa. Similarly, the points exterior to X are dual to hyperplanes that intersect X^*. The ones interior to X are dual to hyperplanes that have no common point with X^*.

8.2.2 Convex hull versus intersection cell

Let $X_1, \ldots, X_n \subset \mathbb{R}^d$ be closed bounded and convex sets and let the origin be inside an X_i. For $X = \cup_{i=1}^n X_i$ we know by Theorem 8.7 that

$$X^{**} \;=\; (X^*)^* \;=\; (\bigcap_{i=1}^{n} X_i^*)^* \;=\; CH(X).$$

We are interested in the convex hull $CH(X) = (X^*)^*$. The set X^* is closed bounded and convex and it contains the origin. Exactly the points on the boundary of X^* are dual to tangential hyperplanes to $CH(X)$. A boundary point of X^* that is a boundary point of X_i^* and X_j^* for $1 \leq i, j \leq n$ dualizes to a plane that supports the convex hull and is tangential to X_i and X_j. So if we know the topology of X^*, then we know the one of $CH(X)$. What we have to do is computing the topology of the boundary of X^*.

We want to compute the boundary of the intersection cell build by $X_1^*, \ldots, X_n^*$. Unfortunately, the origin is not necessarily contained in each X_i. Not every dual point of a tangential hyperplane to X_i is a member of X_i^*. A hyperplane not containing the origin that supports a point of X_i in the interior of $(CH(X_i \cup \{0\})$ dualizes to a point in the exterior of X_i^*. But we know that X_i^* contains the origin and is closed bounded and convex. So even if we dualize for each set X_i each supporting hyperplane not containing the origin and call these set of points X_i', then the topology of the cell in the arrangement of $X_1', \ldots, X_n'$ containing the origin is equal to the one of $X_1^*, \ldots, X_n^*$.

So here is our algorithm for computing the topology of the convex hull of a set of ellipsoids $e_1, \ldots, e_n$: First make sure that at least one ellipsoid contains the origin. This is easily done by a rational translation. Next for each ellipsoid e_i take each tangential plane not containing the origin and dualize it. We will show that this set is again a quadric. We compute its defining polynomial. Depending on the location of the origin with respect to e_i, the resulting quadric is an ellipsoid, an elliptic paraboloid, or a two sheet hyperboloid. Then with the help of our previous results we compute the topology of the cell containing the origin. This directly gives us our desired result: the topology of the convex hull of $e_1, \ldots, e_n$.

8.2.3 The dual of an ellipsoid

Let a be a point on the ellipsoid e. We want to dualize each tangential hyperplane to e not containing the origin. So first we have to think about a formula for the tangential plane at e in the point a, denoted by $T_a(e)$. All the following considerations about tangential planes and duality generally hold for ellipsoids in d-dimensional space. So we will state all results for the general case but keep in mind that we are interested in $d = 3$.

Theorem 8.8: Let $e = (x - c)^T M(a - c) - 1$ be a d-dimensional ellipsoid and $a = (a_1, \ldots, a_d) \in \mathbb{R}^d$ a point with $e(a) = 0$. Then we have

$$T_a(e) \; = \; \{x \in \mathbb{R}^d \mid (x - c)^T M(a - c) - 1 = 0\}.$$

Proof. The tangent space $T_a(e)$ is given by the variety

$$T_a(e) = \{x \in \mathbb{R}^d \mid \frac{\partial e}{\partial x_1}(a)(x_1 - a_1) + \cdots + \frac{\partial e}{\partial x_d}(a)(x_d - a_d) = 0\},$$

see for example [22]. In order to simplify the formula we first look at the derivative $\frac{\partial e}{\partial x_k}(a)$ for $k = 1, \ldots, d$. A short calculation leads to

$$\frac{\partial e}{\partial x_k}(a) \; = \; 2 \sum_{j=1}^{d} m_{kj}(a_j - c_j).$$

with m_{ij} and c_i being the entries of M and c, respectively. This yields the result

$$\begin{aligned} T_a(e) &= \{x \in \mathbb{R}^d \mid \sum_{k=1}^{d} \sum_{j=1}^{d} m_{kj}(a_j - c_j)(x_k - a_k) = 0\} \\ &= \{x \in \mathbb{R}^d \mid (x - a)^T M(a - c) = 0\} \\ &= \{x \in \mathbb{R}^d \mid (x - c)^T M(a - c) - 1 = 0\}. \end{aligned}$$

We conclude the last equality from

$$0 = (a-c)^T M(a-c) - 1 = a^T M(a-c) - c^T M(a-c) - 1.$$

□

Due to our observations in the last section we want to dualize for each ellipsoid $e = (x-c)^T M(x-c) - 1$ every tangential plane not containing the origin. For simplicity we call this set of points the *dual* of e although it does not fits precisely the definition we gave before.

Theorem 8.9: The dual of an ellipsoid $e = (x-c)^T M(x-c) - 1$ is the quadric surface defined by the polynomial

$$q \;=\; y^T(M^{-1} - cc^T)y + 2c^T y - 1 \;=\; y^T M^{-1} y - (1 - c^T y)^2.$$

Proof. Let $a \in \mathbb{R}^d$ be a point on e such that the tangential plane at a does not contain the origin: $0 \neq (0-c)^T M(a-c) - 1 - c^T M(a-c) - 1$. The tangential hyperplane in a is of the form:

$$(x-c)^T M(a-c) = 1 \qquad \Leftrightarrow \qquad x^T \frac{M(a-c)}{1 + c^T M(a-c)} = 1.$$

By the choice of a, the denominator $1 + c^T M(a-c)$ is not equal to zero. Thus the dual set of all tangential planes not containing the origin is equal to

$$\begin{aligned} & \{\frac{M(a-c)}{1 + c^T M(a-c)} \mid (a-c)^T M(a-c) = 1, \;\; -c^T M(a-c) \neq 1\} \\ = \;\; & \{y \mid y^T M^{-1} y - (1 - c^T y)^2 = 0\}. \end{aligned}$$

It remains to show that the last equality holds:

"$\subset$" Let a be such that $(a-c)^T M(a-c) = 1$ and $-c^T M(a-c) \neq 1$. Then we have

$$\begin{aligned} & (\frac{M(a-c)}{1 + c^T M(a-c)})^T M^{-1} \frac{M(a-c)}{1 + c^T M(a-c)} - (1 - c^T \frac{M(a-c)}{1 + c^T M(a-c)})^2 \\ = \;\; & \frac{(a-c)^T M(a-c) - 1}{(1 + c^T M(a-c))^2} \\ = \;\; & 0. \end{aligned}$$

"$\supset$" Let y be such that $y^T M^{-1} y - (1 - c^T y)^2 = 0$. It cannot happen that additionally $c^T y = 1$, because then $y^T M^{-1} y = 0$ would imply $y = 0$ because of M and M^{-1} being positive definite. This is a contradiction to $c^T y = 1$. So we are allowed to define

$$a := \frac{M^{-1}y}{1 - c^T y} + c.$$

The point a is chosen in a way that $(M(a-c))(1 + c^T M(a-c)) = y$. It is easy to prove that a fulfills both criteria $(a-c)^T M(a-c) = 1$ and $-c^T M(a-c) \neq 1$.

□

Now we know that all dual sets are again quadrics, we are done and can use our previous algorithm in order to compute the intersection cell in 3-space.

8.2.4 Classification of the dual of ellipsoids

At last we make a short remark on the kind of quadrics that arise as the dual sets of ellipsoid. For an ellipsoid $e = (x-c)^T M(x-c) - 1$ the dual quadric is defined by a polynomial of the form

$$q = y^T(M^{-1} - cc^T)y + 2c^T y - 1.$$

Due to the classification of quadrics in d-dimensional space given in [40], the kind of q is determined by the sign of the eigenvalues of $(M^{-1} - cc^T)$ and the relation between the rank of the following two matrices:

$$\operatorname{rank}\begin{pmatrix} M^{-1} - cc^T & c \\ c^T & -1 \end{pmatrix} \quad \text{and} \quad \operatorname{rank}(M^{-1} - cc^T).$$

It is easy to see that the first matrix always has rang $d+1$. That means the kind of the dual set is only determined by the eigenvalues of $(M^{-1} - cc^T)$. The eigenvalues depend on the vector of translation c. If we start with $c = 0$ and then continuously move e away from the origin, the dual set behaves in the following way:

1. As long as the origin is situated in the interior of e, $e(0) < 0 \Leftrightarrow c^T Mc < 1$, the dual is again an *ellipsoid*, because all eigenvalues of $(M^{-1} - cc^T)$ are positive.

2. The dual is a *paraboloid* if the origin is on e, $e(0) = 0 \Leftrightarrow c^T Mc = 1$, because then one of the eigenvalues becomes zero, the others are positive.

3. In the case the origin is outside e, $e(0) > 0 \Leftrightarrow c^T Mc > 1$, the dual is a *hyperboloid*, because one eigenvalue is negative, the others all positive.

If the origin is not on e, that means $c^T Mc \neq 1$, the matrix $(M^{-1} - ccT)$ is invertible:

$$(M^{-1} - cc^T)^{-1} = \frac{M + Mcc^T M - M(c^T Mc)}{1 - c^T Mc}.$$

All eigenvalues are non-zero. They are strictly positive if and only if $(M^{-1} - cc^T)^{-1}$ is positive definite, that means if for all $y \neq 0$ the inequality $y^T (M^{-1} - cc^T)^{-1} y > 0$ holds. This is true for $c^T Mc < 1$:

$$y^T (M^{-1} - cc^T)^{-1} y \;=\; \underbrace{\frac{1}{1 - c^T Mc}}_{>0} (y^T M \underbrace{(1 - c^T Mc)}_{>0} y + \underbrace{y^T Mcc^T My}_{(c^T My)^2 \geq 0}) \;>\; 0.$$

In the case $c^T Mc > 1$ at least one eigenvalue is negative: by assumption the matrix M only has positive eigenvalues and thus the signs of the eigenvalues of $(M^{-1} - cc^T)$ are the same as for $M(M^{-1} - cc^T) = (E - Mcc^T)$. The vector c is an eigenvector of $(E - cc^T M)$ with eigenvalue $1 - c^T Mc < 0$: $(E - cc^T M)c = c - cc^T Mc = c(1 - c^T MC)$.

In the case $c^T Mc = 1$ we know that Mc is an eigenvector of $(M^{-1} - cc^T)$ with corresponding eigenvalue 0: $(M^{-1} - cc^T)Mc = c - cc^T Mc = c - c = 0$. That means the matrix has rank $< d$. If we can show that in this case the rank is equal to $d - 1$, we are done. By a continuity argument all other eigenvalues, as for $c^T Mc < 1$, have to be positive and so do the remaining ones in the case of $c^T Mc > 1$.

What remains to do is proving that for $c^T Mc = 1$ the matrix $(M^{-1} - cc^T)$, or equivalently $(E - cc^T M)$, has rank at least $d - 1$:

$$(E - cc^T M) = \begin{pmatrix} 1 - c_1 \sum_{k=1}^d c_k m_{k1} & -c_1 \sum_{k=1}^d c_k m_{k2} & \dots & -c_1 \sum_{k=1}^d c_k m_{kd} \\ -c_2 \sum_{k=1}^d c_k m_{k1} & 1 - c_2 \sum_{k=1}^d c_k m_{k2} & \dots & -c_2 \sum_{k=1}^d c_k m_{kd} \\ \vdots & \vdots & \ddots & \vdots \\ -c_d \sum_{k=1}^d c_k m_{k1} & -c_d \sum_{k=1}^d c_k m_{k2} & \dots & 1 - c_d \sum_{k=1}^d c_k m_{kd} \end{pmatrix}.$$

By assumption the vector $c = (c_1, \dots, c_d)$ is not equal to the zero-vector, so assume without loss of generality $c_d \neq 0$. For all $1 \leq i \leq d - 1$ multiply the last row by c_1/c_d

and subtract it from the i-th row. The rank of the matrix does not change and we obtain a matrix of the form

$$\begin{pmatrix} & & * \\ & E & \vdots \\ * & \dots & * \end{pmatrix}.$$

The upper left $(d-1) \times (d-1)$ matrix is the unit matrix and the rank of the whole matrix therefore at least $d-1$.

Chapter 9

Summary and outlook

In this thesis we have developed a method for computing a cell in an arrangement of quadric surfaces. It uses exact algebraic computation and provides the correct mathematical result in every case, even a degenerate one.

9.1 Theoretical results

The most important problem we faced was founded in the algebraic degree of the intersection curve of two quadrics. In general, the intersection points of such a curve with a second one cannot be computed as nested square roots of rational numbers. So we cannot access the coordinates of common intersection points of three quadrics explicitly. Instead, we have to deal with algebraic numbers implicitly.

The analogous quadric intersection-problem in 2-dimensional space does not bring up this problem. As implied in our investigations, the problematic event points in a planar arrangement of curves are singular points of a curve and tangential intersection points of two curves. For quadratic curves the coordinates of both kinds of event points are computable as simple square root expressions.

In order to solve the 3-dimensional problem we have chosen an approach that works by reduction to planar arrangements. Our aim was to find a method that can also be applied to more general surfaces. Usually no rational parameterization of the intersection curve of two spatial surfaces exists. So there is no hope for general surfaces to do the computation directly in space.

With the help of resultants we have projected the silhouette of each quadric and all intersection curves between two quadrics into the plane. This reduction is degree optimal in the sense that the projection does not change the algebraic degree of the intersection curves. The prize we have to pay for the projection is an increase of the number of singular points and of tangential intersections.

Using partial derivatives, resultant computation, and root isolation, the problem of computing the planar arrangements has been reduced to the problem of answering questions of the following form: Given two curves f and g and a box with rational corners, determine whether f and g have an intersection point inside this box. It turned out that transversal intersection points are easy to compute by determining the sequence of hits of the curves with the boundary of the box. This previously known method we have called simple box hit counting. The more difficult points are tangential intersections and singular points. For these points the sequence of hits provides no information about the behavior of the curves inside the box.

We have presented a new method to test boxes for non-singular tangential intersection points by introducing an additional curve, which we have called the *Jacobi curve*, to the arrangement. The Jacobi curve reduces the problem of determining tangential intersections to the one of testing transversal intersections inside the box.
We have generalized our method of introducing a new curve to the arrangement. This has led to an algorithm for arbitrary planar arrangements that determines all boxes containing tangential intersections. The drawback of the whole approach is that it works only in the case where we know in advance that a box contains no singular point of one of the curves.

For computing the singular points we have taken advantage of the fact that the curves in the planar arrangements are projected intersection curves of two quadrics. We have proven that at most two singular points can result from the projection in the sense that two non-intersecting branches of the spatial intersection curve are projected on top of each other. We have shown how to compute the coordinates of these singular points as roots of quadratic rational polynomials. In most cases we have succeeded in expressing the coordinates of the remaining singular points as roots of quadratic rational polynomials. Only in the case that the spatial intersection curve consists of four lines, computing the coordinates requires a second square root.

9.2 Experimental results

We claimed that our theoretical results for computing arrangements of quadric surfaces promise a good performance in practice. In order to justify this statement, we made some experiments in implementing and testing our ideas. Our prototypical implementation determines event points in the planar arrangements induced by three quadrics. It uses the basic data types of LEDA [48] and the rational polynomial class as well as the resultant and Sturm sequence computation of MAPC [44].

Consider the screen shots in Figures 9.1 – 9.6 we made from the output for computing the event points for the three input quadrics

$$\begin{aligned} p(x,y,z) &= 7216x^2 - 11022xy - 12220xz + 15624y^2 + 15168yz \\ &\quad +11186z^2 - 1000 \\ q(x,y,z) &= 4854x^2 - 3560xy + 4468xz + 658x + 5040y^2 + 32yz + 1914y \\ &\quad +10244z^2 + 3242z - 536 \\ r(x,y,z) &= 8877x^2 - 10488xy + 9754xz + 1280x + 16219y^2 - 16282yz \\ &\quad -808y + 10152z^2 - 1118z - 796. \end{aligned}$$

The running time of our implementation for this special example on an Intel Pentium 700 is about 18 seconds. Of course the running time mainly depends on the number of decimal digits of the three input quadrics as can be seen in the following table:

number of digits	5	10	15	20	25	30
running time in seconds	18	33	56	92	126	186

The only mathematical tools that are used during the calculation are resultants and subresultants, root separation, gcd of univariate and bivariate polynomials, and solving quadratic univariate polynomials. The size of the coefficients of the polynomials has a great impact on the behavior of all these computations. In our example about half of the running time is spent on computing all necessary resultants. Isolating the real roots on the coordinate axes with Uspensky's algorithm is quite fast. The rest of the time is needed to test the more than 100 boxes for intersection points.

In order to judge the running time of our algorithm, we have run our example in Maple [30] using the plot_real_curves function in the algcurves package. As can be seen in the screen shots in Figures 9.1-9.6, our program outputs a plotting of the real branches of the curves and marks their extreme and singular points and all intersection points between two curves by small boxes. The plot_real_curves function applied to a bivariate polynomial nearly does the same. It also plots the defined curve and inter alia marks all extreme and singular points. Computing and plotting the bivariate resultants defining the silhouettecurves and cutcurves in our example with Maple took 60 seconds on an Intel Pentium III with 850 MHz. Thereby Maple missed the isolated point of $\mathrm{res}(p,q,z)$. Computing the intersection points of two curves was even slower. We have tested the plot_real_curves function with the input $\mathrm{res}(p,q,z) \cdot \mathrm{res}(p,r,z)$ and it ran for 70 seconds. Of course the comparison is not really fair, because Maple makes no use of the fact that the algebraic curves, as projected intersection curves of quadrics, have a special form. But it nevertheless shows that locating singular points and intersection points of two curves is difficult and time-consuming and that our approach leads to promising running times.

So far our implementation cannot handle all degenerate cases, especially it cannot detect genuine points. But, as shown before, all singular points, either top-bottom or genuine, conceptually are treated the same way. A line cutting through two of them is computed that helps to factor the univariate resultants. The remaining polynomials are at most quadratic and explicit solutions can be calculated. So we have proven in the previous chapters that degenerate cases definitely do not increase the running time of our algorithm.

9.3 Further research

The prototypical implementation shows that our algorithm is a first and important step towards exact and efficient computation of arrangements of curves and surfaces. Until now we have made no special efforts to optimize the running time of our implementation. Choosing the algorithms for the mathematical computations more carefully and making use of filtering techniques will surely lead to a better performance. Above that, of course, there is still some work in practical as well as in theoretical sense. So far our implementation only determines event points in the plane. Part of our future work will be to implement a sweep line algorithm that computes the trapezoidal decomposition of each planar arrangement and to combine the results to an overall description of the arrangement in space.

Our approach provides an efficient and exact algorithm for computing a cell in an arrangement of quadric surfaces, even in degenerate cases. It is general in the sense that it could be applied to every kind of spatial surfaces defined by rational polynomials. Only at one point we made use of the fact that the input surfaces are restricted to ones defined by at most quadratic polynomials. Computing the singular points of the planar arrangements was restricted to this assumption. So another topic of our future research will be locating singular points of arbitrary planar curves in order to enlarge our approach to more general algebraic surfaces.

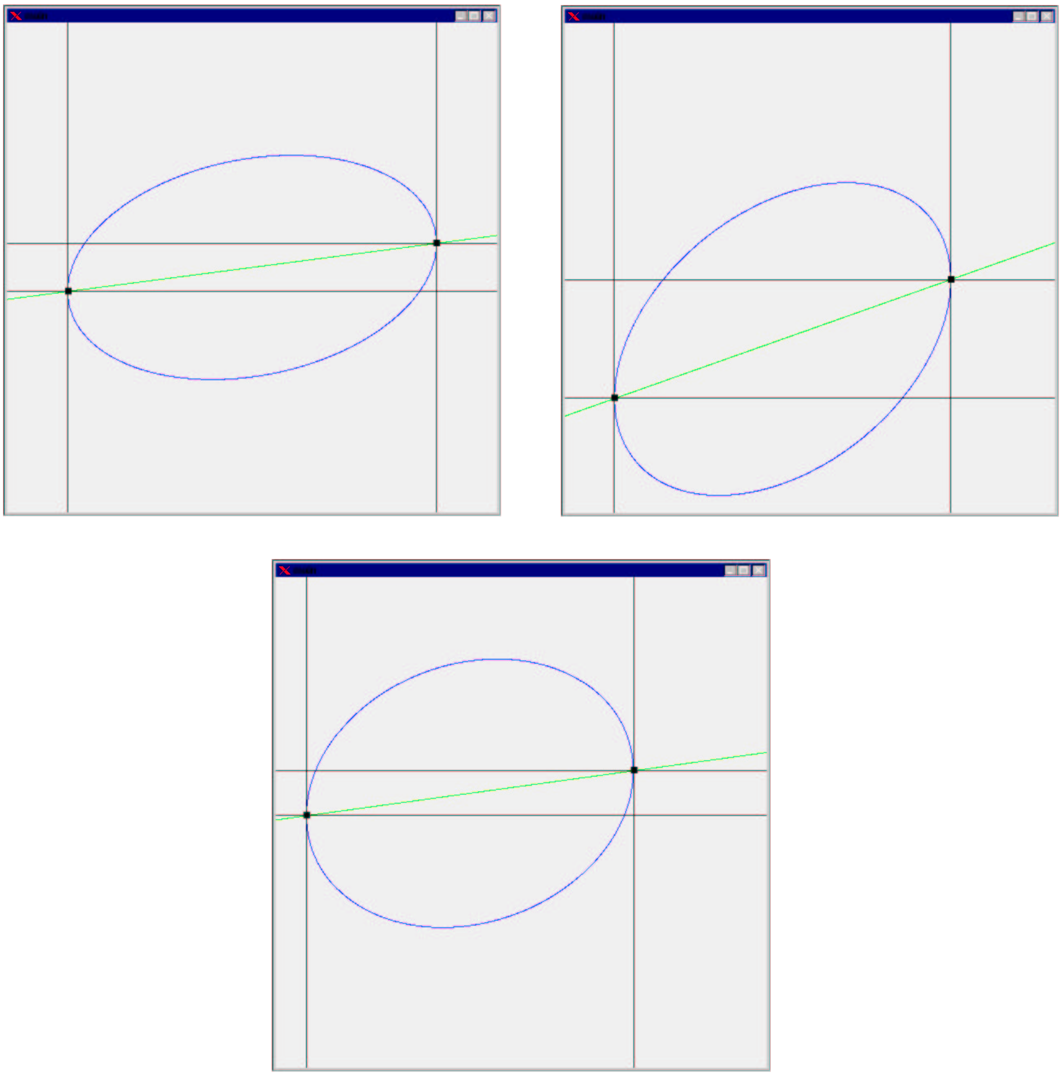

Figure 9.1: The blue ellipses are the projected silhouettes of the input ellipsoids p, q, and r. We denote them by $sil_p = \text{res}(p, p_z, z)$, $sil_q = \text{res}(q, q_z, z)$, and $sil_r = \text{res}(r, r_z, z)$, respectively. An extreme point of a blue silhouettecurve is an intersection point with the green line. The green line is defined by the partial derivative with respect to y. All extreme points are determined correctly and marked by small boxes.

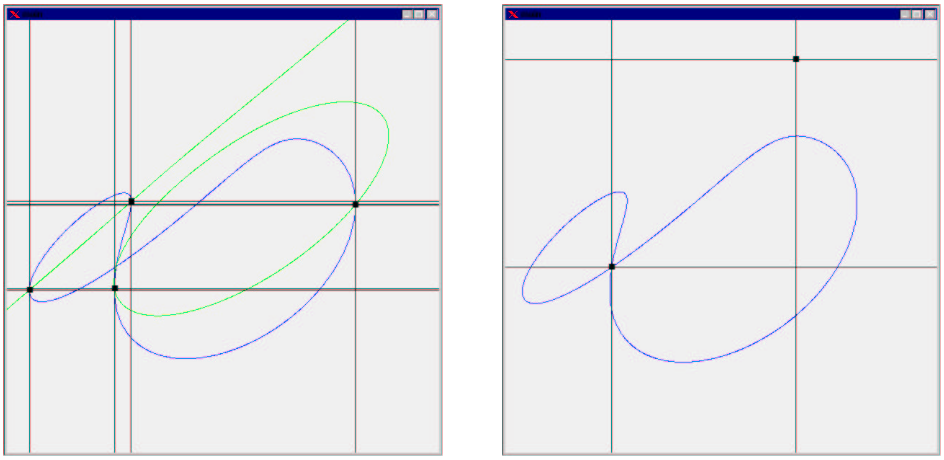

Figure 9.2: The cutcurve $cut_{p,q} = \mathrm{res}(p, q, z)$, besides extreme points, has two top-bottom singular points. In the left picture the extreme points of $cut_{p,q}$ are determined by simple box hit counting. In the right picture one can see that our program locates two singular points. This is done by computing explicit solutions. One singular point is a self-intersection point of $cut_{p,q}$. The second one is an isolated point that leads to two complex common points of p and q.

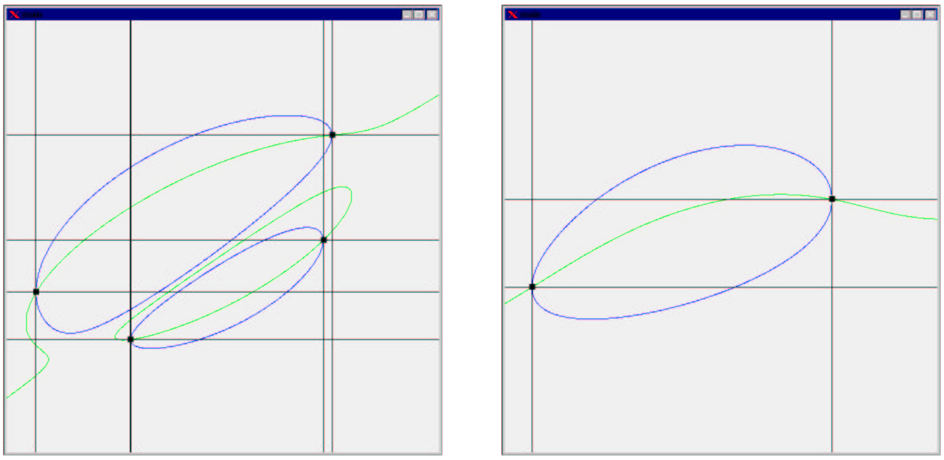

Figure 9.3: The cutcurves $cut_{p,r} = \mathrm{res}(p, r, z)$ and $cut_{q,r} = \mathrm{res}(q, r, z)$ have no singular points. Their extreme points are determined correctly.

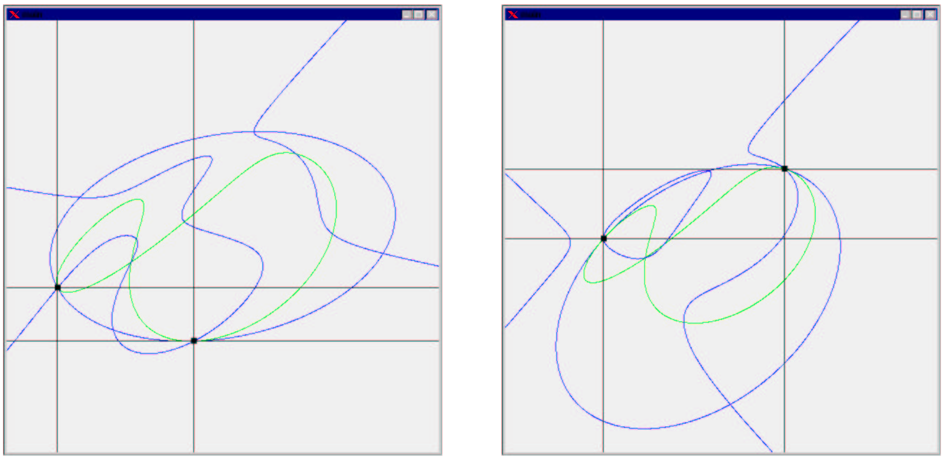

Figure 9.4: For each projected intersection curve of two quadrics p_1 and p_2 our program has to locate the tangential intersection points with the projected silhouettes of p_1 and of p_2. The green cutcurve $cut_{p,q}$ has tangential intersection points with the blue silhouettecurve sil_p in the left picture and with the blue silhouettecurve sil_q in the right picture. The additional blue curve in each picture is the Jacobi curve we need for applying extended box hit counting.

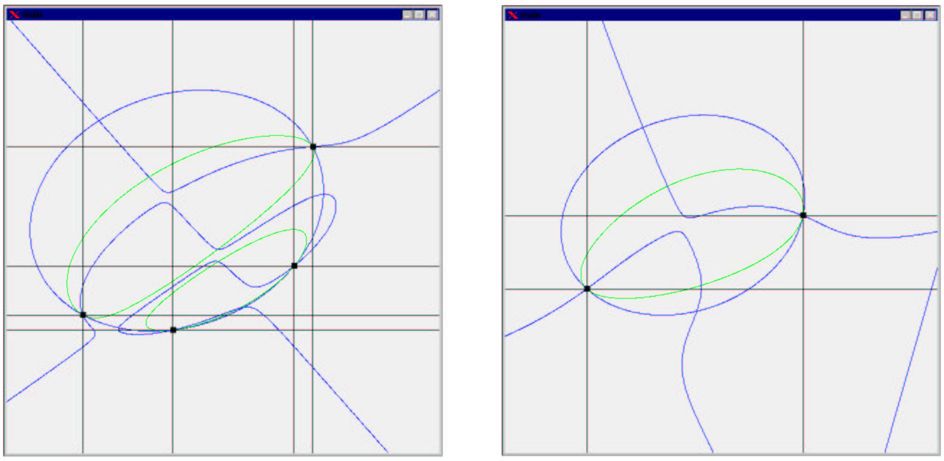

Figure 9.5: The cutcurve $cut_{p,r}$ has intersection points with the silhouettecurve sil_r, consider the left picture. There are no intersection points with sil_p. Also the cutcurve $cut_{q,r}$ has tangential intersections with sil_r, see the right picture. There are no intersection points of $cut_{q,r}$ and sil_q.

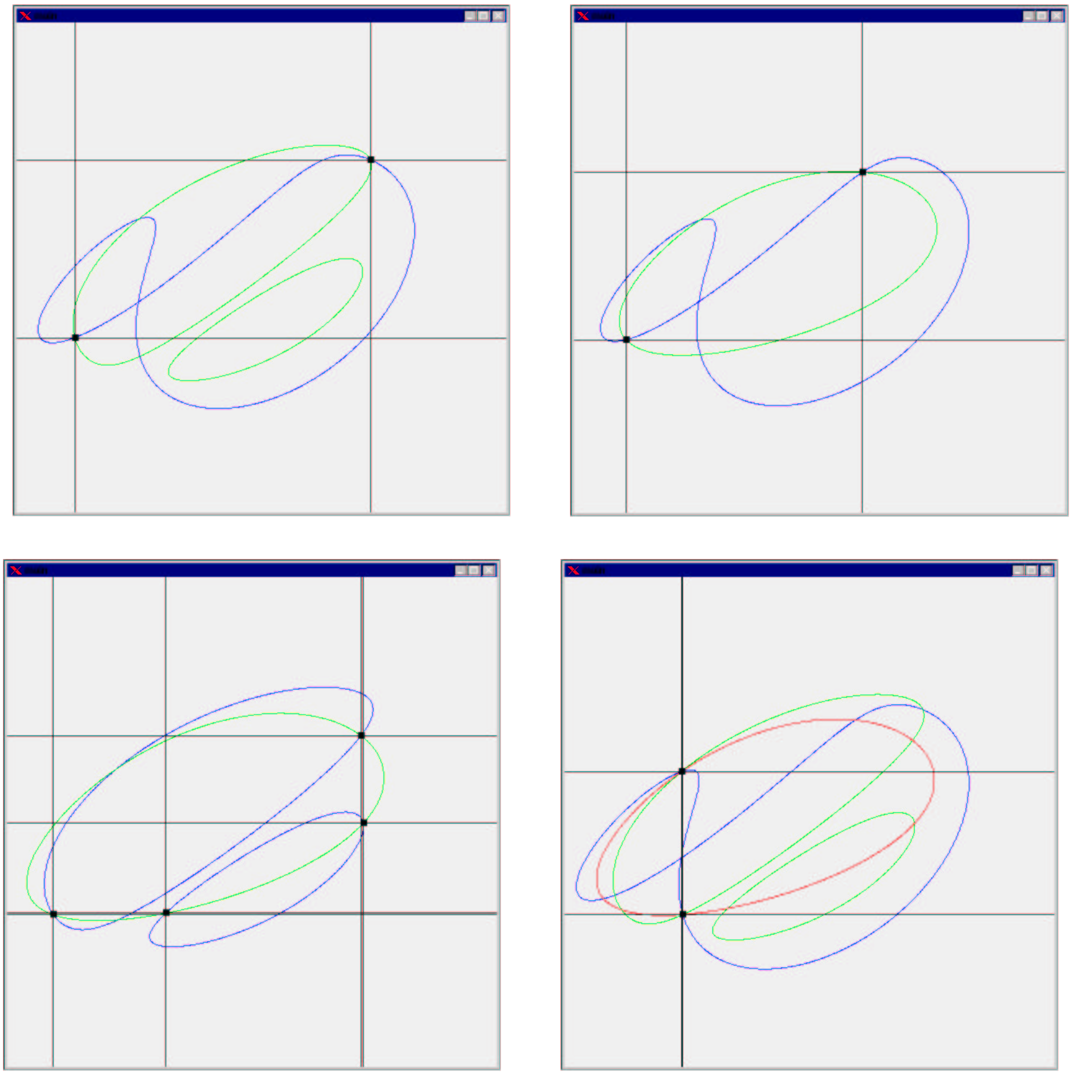

Figure 9.6: At last, artificial and spatial intersection points between two cutcurves have to be computed. In the first 3 pictures one can see the artificial intersection points of $cut_{p,q}$ and $cut_{p,r}$, of $cut_{p,q}$ and $cut_{q,r}$, and of $cut_{p,r}$ and $cut_{q,r}$, respectively. Spatial intersection points are common intersection points of $cut_{p,q}$, $cut_{p,r}$, and $cut_{q,r}$. They are shown in the lower right picture.

Bibliography

[1] S. Abhyankar and C. Bajaj. Computations with algebraic curves. In *Proc. Internat. Sympos. on Symbolic and Algebraic Computation*, volume 358 of *Lecture Notes Comput. Sci.*, pages 279–284. Springer-Verlag, 1989.

[2] P. K. Agarwal and M. Sharir. Arrangements and their applications. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 49–119. Elsevier Science Publishers B.V. North-Holland, Amsterdam, 2000.

[3] S. Arnborg and H. Feng. Algebraic decomposition of regular curves. *J. Symbolic Comput.*, 15(1):131–140, 1988.

[4] D. S. Arnon, G. E. Collins, and S. McCallum. Cylindrical algebraic decomposition I: The basic algorithm. *SIAM J. Comput.*, 13(4):865–877, 1984.

[5] D. S. Arnon, G. E. Collins, and S. McCallum. Cylindrical algebraic decomposition II: The adjacency algorithm for the plane. *SIAM J. Comput.*, 13(4):878–889, 1984.

[6] D. S. Arnon, G. E. Collins, and S. McCallum. An adjacency algorithm for cylindrical algebraic decomposition in three-dimensional space. *J. Symbolic Comput.*, 5(1–2):163–187, 1988.

[7] D. S. Arnon and S. McCallum. A polynomial time algorithm for the topological type of a real algebraic curve. *J. Symbolic Comput.*, 5:213–236, 1988.

[8] C. Bajaj and M. S. Kim. Convex hull of objects bounded by algebraic curves. *Algorithmica*, 6:533–553, 1991.

[9] J. L. Bentley and T. Ottmann. Algorithms for reporting and counting geometric intersections. *IEEE Trans. Comput.*, C-28:643–647, 1979.

[10] E. Berberich, A. Eigenwillig, M. Hemmer, S. Hert, K. Mehlhorn, and E. Schömer. A computational basis for conic arcs and boolean operations on conic polygons. submitted to ESA, 2002.

[11] J.-D. Boissonat and J. Snoeyink. Efficient algorithms for line and curve segment intersection using restricted predicates. In *Proc. 15th Annu. ACM Sympos. Comput. Geom.*, pages 370–379, 1999.

[12] J. Boissonnat and F. P. Preparata. Robust plane sweep for intersecting segments. *SIAM Journal on Computing*, 23:1401–1421, 2000.

[13] J.-D. Boissonnat, A. Cérézo, O. Devillers, J. Duquesne, and M. Yvinec. An algorithm for constructing the convex hull of a set of spheres in dimension d. *Comput. Geom. Theory Appl.*, 6:123–130, 1996.

[14] T. J. Bromwich. *Quadratic forms and their classification by means of invariant factors.* Cambridge Tracts in Mathematics and Mathematical Physics, No. 3, 1906. Reprint. Hafner, New York.

[15] J. Canny. *The Complexity of Robot Motion Planning.* MIT Press, Cambridge, MA, 1987.

[16] T. Chan. Reporting curve segment intersection using restricted predicates. *Computational Geometry*, 16:245–256, 2000.

[17] B. Chazelle, H. Edelsbrunner, L. Guibas, and M. Sharir. A singly exponential stratification scheme for real semi-algebraic varieties and its applications. *Theoretical Computer Science*, 84:77–105, 1991.

[18] E. Chionh, R. Goldman, and J. Miller. Using multivariate resultants to find the intersection of three quadric surfaces. *Transactions on Graphics*, 10:378–400, 1991.

[19] G. E. Collins. Quantifier elimination for real closed fields by cylindrical algebraic decomposition. In *Proc. 2nd GI Conf. on Automata Theory and Formal Languages*, volume 6, pages 134–183. Lecture Notes in Computer Science, Springer, Berlin, 1975.

[20] G. E. Collins and A. G. Akritas. Polynomial real root isolation using descartes' rule of signs. In *SYMSAC*, pages 272–275, 1976.

[21] G. E. Collins and R. Loos. Real zeros of polynomials. In B. Buchberger, G. E. Collins, and R. Loos, editors, *Computer Algebra: Symbolic and Algebraic Computation*, pages 83–94. Springer-Verlag, New York, NY, 1982.

[22] D. Cox, J. Little, and D. O'Shea. *Ideals, Varieties, and Algorithms.* Springer, New York, 1997.

[23] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf. *Computational Geometry: Algorithms and Applications.* Springer-Verlag, 2000, 2nd rev. edition.

[24] O. Devillers, A. Fronville, B. Mourrain, and M. Teillaud. Exact predicates for circle arcs arrangements. In *Proc. 16th Annu. ACM Symp. Comput. Geom.*, 2000.

[25] D. P. Dobkin and D. L. Souvaine. Computational geometry in a curved world. *Algorithmica*, 5:421–457, 1990.

[26] L. Dupont, D. Lazard, S. Lazard, and S. Petitjean. A new algorithm for the robust intersection of two general quadrics. submitted to Solid Modelling, 2002.

[27] H. G. Eggleston. *Convexity.* Cambridge Tracts in Mathematics and Mathematical Physics, Cambridge University Press, 1958.

[28] R. T. Farouki, C. A. Neff, and M. A. O'Connor. Automatic parsing of degenerate quadric-surface intersections. *ACM Trans. Graph.*, 8:174–203, 1989.

[29] E. Flato, D. Halperin, I. Hanniel, and O. Nechushtan. The design and implementation of planar maps in cgal. In *Proceedings of the 3rd Workshop on Algorithm Engineering*, Lecture Notes Comput. Sci., pages 154–168, 1999.

[30] F. Garvan. *The MAPLE book.* CRC Press, 2001.

[31] N. Geismann, M. Hemmer, and E. Schömer. Computing a 3-dimensional cell in an arrangement of quadrics: Exactly and actually! In *SOCG*, 2001.

[32] N. Geismann, M. Hemmer, and E. Schömer. The convex hull of ellipsoids. In *SOCG video track*, 2001.

[33] C. G. Gibson. *Elementary Geometry of Algebraic Curves.* Cambridge University Press, 1998.

[34] R. Gunning and H. Rossi. *Analytic functions of several complex variables.* Prentice-Hall, Inc., Englewood Cliffs, N.J., 1965.

[35] D. Halperin. Arrangements. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 21, pages 389–412. CRC Press LLC, Boca Raton, FL, 1997.

[36] M. Hemmer. *Realiable computation of planar and spatial quadric arrangements.* Universität des Saarlandes, Saarbrücken, 2002. Master Thesis.

[37] C. Hoffmann. *Geometric and Solid Modeling.* Morgan-Kaufmann, San Mateo, CA, 1989.

[38] H. Hong. Subresultant under composition. *Journal of Symbolic Computation*, 23(4):355–365, 1997.

[39] C.-K. Hung and D. Ierardi. Constructing convex hulls of quadratic surface patches. In *Proc. 7th Canad. Conf. Comput. Geom.*, pages 255–260, 1995.

[40] K. Itô. *Encyclopedic Dictionary of Mathematics.* MIT Press, Cambridge, Massachusetts and London, England, 1996. Second Edition.

[41] V. Karamcheti, C. Li, I. Pechtchanski, and C. Yap. *The CORE Library Project*, 1.2 edition, 1999. http://www.cs.nyu.edu/exact/core/.

[42] G. L. Keith O. Geddes, Stephen R. Czapor. *Algorithms for Computer Algebra.* Kluwer Academic Publishers, 1992.

[43] J. Keyser, T. Culver, M. Foskey, D. Manocha, and S. Krishnan. Esolid - a system for exact boundary evaluation. submitted to Solid Modeling, 2002.

[44] J. Keyser, T. Culver, D. Manocha, and S. Krishnan. MAPC: A library for efficient and exact manipulation of algebraic points and curves. In *Proc. 15th Annu. ACM Sympos. Comput. Geom.*, pages 360–369, 1999.

[45] J. C. Keyser. *Exact boundary evaluation for curved solids.* Univ. of North Carolina at Chapel Hill, Chapel Hill, 2000. Ph.D. dissertation.

[46] J. Levin. A parametric algorithm for drawing pictures of solid objects composed of quadric surfaces. *Commun. ACM*, 19(10):555–563, Oct. 1976.

[47] J. Levin. Mathematical models for determining the intersections of quadric surfaces. *Comput. Graph. Image Process.*, 11:73–87, 1979.

[48] K. Mehlhorn and S. Näher. *LEDA – A Platform for Combinatorial and Geometric Computing.* Cambridge University Press, 1999.

[49] V. Milenkovic. Calculating approximate curve arrangements using rounded arithmetic. In *Proc. 5th Annu. ACM Sympos. Comput. Geom.*, pages 197–207, 1989.

[50] J. Miller and R. Goldman. Combining algebraic rigor with geometric robustness for the detection and calculation of conic sections in the intersection of two quadric surfaces. In *Proceedings of the Symposium on Solid Modeling and Applications*, pages 221–231, 1991.

[51] J. R. Miller. Geometric approaches to nonplanar quadric surface intersection curves. *ACM Trans. Graph.*, 6:274–307, 1987.

[52] J. R. 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*, 57:55–66, 1995.

[53] P. S. Milne. On the solutions of a set of polynomial equations. In *Symbolic and Numerical Computation for Artificial Intelligence*, pages 89–102. 1992.

[54] B. Mishra. Computational real algebraic geometry. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 29, pages 537–558. CRC Press LLC, Boca Raton, FL, 1997.

[55] K. Mulmuley. A fast planar partition algorithm, II. *J. ACM*, 38:74–103, 1991.

[56] K. Mulmuley. *Computational Geometrie.* Prentice Hall, Englewood Cliffs, NJ, 1994.

[57] M. Neagu and B. Lacolle. Computing the combinatorial structure of arrangements of curves using polygonal approximations. In *Proceedings of the European Workshop on Computational Geometry*, 1998.

[58] F. Nielsen and M. Yvinec. An output-sensitive convex hull algorithm for planar objects. Technical Report 2575, Institut nationale de recherche en informatique at en automatique, INRIA Sophia-Antipolis, 1995.

[59] T. Papanikolaou. *LiDIA Manual - A Library for Computational Number Theory.* Universität des Saarlandes, Saarbrücken, 1995.

[60] P. Pedersen. Multivariate sturm theory. In *Proceedings of Applied Algebra, Algebraic Algorithms and Error-Correcting Codes*, pages 318–323, 1991.

[61] F. P. Preparata and M. I. Shamos. *Computational geometry and introduction.* Springer-Verlag, New York, 1985.

[62] D. Prill. On approximations and incidence in cylindrical algebraic decomposition. *Siam J. Comput.*, 15(4):972–993, 1986.

[63] A. Rege. *A Toolkit for Algebra and Geometry.* Univ. of California at Berkely, Berkely, California, 1996. Ph.D. dissertation.

[64] T. Sakkalis. The topological configuration of a real algebraic curve. *Bulletin of the Australian Mathematical Society*, 43:37–50, 1991.

[65] T. Sakkalis and R. T. Farouki. Singular points of algebraic curves. *Journal of Symbolic Computation*, 9:405–421, 1990.

[66] J. T. Schwartz. Fast probabalistic algorithms for verification of polynomial identities. *J. of the ACM*, 27:701–717, 1980.

[67] O. Schwarzkopf and M. Sharir. Vertical decomposition of a single cell in a three-dimensional arrangement of surfaces and its applications. *Discrete Comput. Geom.*, 18:269–288, 1997.

[68] R. Seidel. Convex hull computations. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 19, pages 361–375. CRC Press LLC, Boca Raton, FL, 1997.

[69] C.-K. Shene and J. K. Johnstone. On the planar intersection of natural quadrics. In *Proc. ACM Sympos. Solid Modeling Found. CAD/CAM Appl.*, pages 233–242. Springer-Verlag, 1991.

[70] J. Snoeyink and J. Hershberger. Sweeping arrangements of curves. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 6:309–349, 1991.

[71] A. Tarski. *A Decision Method for Elementary Algebra and Geometry.* Univ. of California Press, Berkely, 1951. second ed., rev.

[72] J. von zur Gathen and J. Gerhard. *Modern Computer Algebra.* Cambridge University Press, 1999.

[73] R. J. Walker. *Algebraic curves.* Springer–Verlag, New York, Heidelberg, Berlin, 1978.

[74] W. Wang, B. Joe, and R. Goldman. Rational quadratic parameterizations of quadrics. *Internat. J. Comput. Geom. Appl.*, 7:599–619, 1997.

[75] W. Wang, B. Joe, and R. Goldman. Computing quadric surface intersections based on an analysis of plane cubic curves. *submitted to Trans. on Graphics*, 2000.

[76] R. Wein. High-level filtering for arrangements of conic arcs. submitted to ESA, 2002.

[77] C. K. Yap. *Fundamental Problems of Algorithmic Algebra.* Oxford University Press, New York, Oxford, 2000.

Summary

Computing arrangements of curves and surfaces is one of the fundamental problems in different areas of computer science like computational geometry, algebraic geometry, and solid modeling. As long as arrangements of lines and planes defined by rational numbers are considered, all computations can be done over the field of rational numbers avoiding numerical errors. In this case, algorithms are available that are *efficient* in practice concerning their running time, and *exact* in the sense that they always compute the mathematical correct result, even for degenerate inputs.

When higher-degree algebraic curves and surfaces are considered instead of linear ones, things are more difficult. In general, the intersection points of two planar curves or three surfaces in 3-space defined by rational polynomials have irrational coordinates. That means, instead of rational numbers, one has to deal with algebraic numbers. One way to overcome this difficulty is to develop algorithms that use floating point arithmetic. These algorithms are quite fast but can, because of approximation errors introduced in the floating point computations, produce not just slightly inaccurate output but completely wrong results. A second approach is to use exact algebraic computation methods. Then of course the results are correct but the algorithms are in general very slow.

Computing the mathematically correct topology of a cell in an arrangement of curved surfaces efficiently is a challenging task. As far as we know, we are the first who provide such an algorithm for a set of *quadric* surfaces in 3 dimensions [31]. Our algorithm uses exact rational algebraic computation and it can handle every degenerate input. A prototypical implementation shows that the theoretical results promise good performance in practice.

Our approach operates similar to the cylindrical algebraic decomposition [19]. By projection, it reduces the 3-dimensional problem to the one of computing planar arrangements of algebraic curves of degree up to 4. The reduction is algebraically optimal in the sense that it does not affect the algebraic degree of the problem we consider. The curves in the planar arrangements can have 6 singular points and two curves can intersect in up to 16 points. The coordinates of these event points are given as algebraic numbers. The most important problem we face is founded in the high degree of these polynomials.

Using partial derivatives, resultant computation, and root isolation, the problem of computing the event points of the planar arrangements is reduced to the problem of answering questions of the following form: Given two curves f and g of degree at most

4 and a box with rational corners containing at most one intersection point of f and g, determine whether f and g have an intersection point inside the box. Transversal intersection points are easy to compute by determining the sequence of hits of the curves with the boundary of the box. The more difficult points are tangential intersections and singular points. For these points the sequence of hits provides no information about the behavior of the curves inside the box.

We present a new method to test boxes for non-singular tangential intersections. This is done by introducing an additional curve, which we call the *Jacobi curve*, in the arrangement. To the best of our knowledge, we are the first who consider an auxiliary curve in order to solve tangential intersections. We prove that at an intersection point of multiplicity 2, the Jacobi curve cuts transversally through both involved curves. This fact enables us to reduce the problem of detecting tangential intersections of multiplicity 2 to the one of locating transversal intersections.
We generalize our method of introducing a new curve to the arrangement. This leads to an algorithm for arbitrary planar arrangements that determines all boxes containing tangential intersections. The drawback is that it works only if we know in advance that a box contains no singular point of one of the curves.

For locating the singular points of the curves we make use of the fact that we consider projected intersection curves of quadrics. We prove that at most two singular points result from the projection, in the sense that two non-intersecting branches of the spatial intersection curve are projected on top of each other. We show how to compute the coordinates of these singular points as roots of quadratic rational polynomials. In most cases we also succeed in expressing the coordinates of the remaining singular points as roots of quadratic rational polynomials. Only in the case that the spatial intersection curve consists of four lines does computing the coordinates require a second square root.

In one sense the work recently done by Dupont, Lazard, Lazard, and Petitjean [26] leads to the same result as ours, namely that computing with quadric surfaces can be done exactly and efficiently in all cases, by working over the rationales with only few additional square roots. But their approach directly works in space and searches for a parameterization of the intersection curves. This way of solving the problem is not extendible to more complicated surfaces. The methods presented in our work can also be applied to arbitrary algebraic curved surfaces. Only for computing singular points in the planar arrangement do we make use of the fact that we consider quadric input surfaces. In that sense, our work is a first and important step towards an efficient and exact algorithm for computing arrangements of arbitrary algebraic surfaces.

Zusammenfassung

In verschiedenen Bereichen der Informatik, wie zum Beispiel der Algorithmische Geometrie, der Computer Algebra und des Solid Modelings, ist das Berechnen von Arrangements, die von Kurven und Flächen erzeugt werden, ein grundlegendes Problem. Solange Arrangements von Linien und Ebenen, die durch rationale Zahlen definiert sind, betrachtet werden, können alle Berechnungen über dem Körper der rationalen Zahlen durchgeführt werden. Das verhindert das Auftreten numerischer Fehler. In diesem Fall sind Algorithmen bekannt, die sowohl *effizient* bezüglich ihrer realen Laufzeit sind, als auch *exakt* das jeweils richtige mathematische Ergebnis berechnen, selbst für degenerierte Eingaben.

Das Problem wird schwieriger, falls man statt linearer Kurven und Flächen solche mit höherem algebraischen Grad betrachtet. Im Allgemeinen haben Schnittpunkte von zwei Kurven oder drei Flächen, selbst wenn sie durch rationale Polynome definiert sind, irrationale Koordinaten. Anstelle rationaler Zahlen muss man mit algebraischen Zahlen rechnen. Eine Möglichkeit, das Problem zu lösen, stellt die Verwendung von Fließkomma-Arithmetik dar. Algorithmen, die auf dieser Grundlage arbeiten, sind sehr schnell. Aber in degenerierten Situationen kann es passieren, dass sie statt einer leicht ungenauen Ausgabe sogar ein falsches Ergebnis liefern. Ein zweiter Ansatz basiert auf exakten algebraischen Berechnungen. Algorithmen, die darauf basieren, berechnen alle Ausgaben korrekt, sind aber im allgemeinen sehr langsam.

Die mathematisch korrekte Topologie einer Zelle in einem Arrangement von gekrümmten Flächen effizient zu berechnen, stellt ein schwieriges Problem dar. Soweit uns bekannt ist, sind wir die ersten, die dieses Problem für eine Menge von quadratischen Flächen im Raum lösen [31]. Unser Algorithmus basiert auf exakten algebraischen Berechnungen und kann jede degenerierte Eingabe behandeln. Eine prototypische Implementierung zeigt, dass unsere theoretischen Ergebnisse eine gute Laufzeit in der Praxis versprechen.

Unser Ansatz arbeitet ähnlich wie die zylindrische algebraische Zerlegung [19]. Durch Projektion wird das 3-dimensionale Problem auf mehrere 2-dimensionale reduziert, jedes bestehend aus der Berechnung eines planaren Arrangements algebraischer Kurven vom Maximalgrad 4. Die Reduktion ist also algebraisch optimal in dem Sinne, dass der algebraische Grad des Problems nicht verändert wird. Die Kurven in einem planaren Arrangement können bis zu 6 singuläre Punkte aufweisen und zwei Kurven können sich in bis zu 16 Punkten schneiden. Wir kennen die Koordinaten dieser Punkte nur als algebraische Zahlen. Der hohe Grad der auftretenden Polynome stellt dabei das größte

Problem dar.

Durch das Berechnen von partiellen Ableitungen, Resultanten und isolierenden Intervallen von Nullstellen kann die Berechnung der Ereignispunkte eines planaren Arrangements reduziert werden auf die Beantwortung von Fragen folgenden Typs: Gegeben sind zwei Kurven f und g vom Maximalgrad 4 und ein Rechteck in der Ebene mit rationalen Eckpunkten, das höchstens einen Schnittpunkt von f und g enthält. Bestimme, ob f und g einen Schnittpunkt in dem Rechteck haben. Transversale Schnittpunkte können leicht erkannt werden, indem man die Reihenfolge bestimmt, in der die Kurven den Rand des Rechtecks schneiden. Tangentiale Schnitte und singuläre Punkte sind weit schwieriger zu bestimmen. Für diese Punkte sagt die Reihenfolge der Kurven auf dem Rand des Rechtecks nichts über ihr Verhalten im Inneren aus.

Wir präsentieren eine neue Methode, mit deren Hilfe man Rechtecke auf nicht singuläre tangentiale Schnitte überprüfen kann. Dieses wird durch das Hinzufügen einer neuen Kurve, die wir Jacobi Kurve nennen, zu dem Arrangement realisiert. Soweit wir wissen, sind wir die ersten, die eine zusätzliche Kurve betrachten, um tangentiale Schnitte zu lösen. Wir zeigen, dass in einem Schnittpunkt der Multiplizität 2 die Jacobi Kurve beide an dem Schnitt beteiligten Kurven transversal schneidet. Dieses Ergebnis erlaubt es, die Bestimmung tangentialer Schnitte der Multiplizität 2 auf die transversaler Schnitte zurückzuführen.
Wir verallgemeinern das Prinzip, tangentiale Schnitte anhand einer zusätzlichen Kurve zu bestimmen. Dieses führt zu einem Verfahren, mit dessen Hilfe man in allgemeinen planaren Arrangements algebraischer Kurven alle Rechtecke mit tangentialen Schnitten bestimmen kann. Unser gesamter Ansatz hat jedoch einen Nachteil. Er ist nur dann anwendbar, wenn man ausschließen kann, dass das zu untersuchende Rechteck einen singulären Punkt einer der Kurven enthält.

Um die singulären Punkte der Kurven zu bestimmen, benutzen wir die Tatsache, dass es sich um projizierte Schnittkurven von Quadriken handelt. Wir zeigen, dass höchstens zwei singuläre Punkte von der Projektion herrühren, in dem Sinne, dass zwei sich nicht schneidende Äste einer räumlichen Schnittkurve übereinander projiziert werden. Wir zeigen, wie man die Koordinaten dieser singulären Punkte als Nullstellen von quadratischen rationalen Polynomen berechnet. In fast allen Fällen können wir auch die Koordinaten der übrigen singulären Punkte als Nullstellen von quadratischen rationalen Polynomen angeben. Nur für den Fall, dass die räumliche Schnittkurve aus vier Geraden besteht, benötigen wir zum Berechnen der Koordinaten eine zusätzliche Quadratwurzel.

In gewisser Hinsicht führt das kürzlich von Dupont, Lazard, Lazard und Petitjean [26] vorgestellte Resultat zu dem gleichen Ergebnis wie unsere Arbeit. Man kann Berechnungen mit quadratischen Flächen exakt und effizient in allen Fällen nur unter Verwendung weniger Quadratwurzeln durchführen. Der Ansatz der Franzosen arbeitet im Raum und bestimmt die Parametrisierungen der Schnittkurven. Diese Vorgehensweise ist nicht auf kompliziertere Flächen verallgemeinerbar. Die von uns vorgestellten Methoden können auf beliebige gekrümmte algebraische Fächen angewandt werden. Lediglich für das Berechnen der singulären Punkte haben wir Gebrauch von der Tatsache gemacht, dass es sich bei der Eingabe um Quadriken handelt. In diesem Sinne ist unsere Arbeit ein erster und wichtiger Schritt in Richtung eines effizienten und exakten Verfahrens zur Berechnung von Arrangements beliebiger algebraischer Flächen.