Isotopic triangulation of a real algebraic surface

Lionel Alberti, Bernard Mourrain, Jean-Pierre Técourt

Isotopic triangulation of a real algebraic surface

Lionel Alberti, Bernard Mourrain, Jean-Pierre Técourt

GALAAD
INRIA Méditerranée
2004 route des Lucioles, B.P. 93
06902 Sophia-Antipolis, Cedex FRANCE

Abstract

We present a new algorithm for computing the topology of a real algebraic surface S in a ball B, even in singular cases. We use algorithms for 2D and 3D algebraic curves and show how one can compute a topological complex equivalent to S, and even a simplicial complex isotopic to S by exploiting properties of the contour curve of S. The correctness proof of the algorithm is based on results from stratification theory. We construct an explicit Whitney stratification of S, by resultant computation. Using Thom's isotopy lemma, we show how to deduce the topology of S from a finite number of characteristic points on the surface. An analysis of the complexity of the algorithm and effectiveness issues conclude the paper.

Key words: meshing, implicit algebraic surfaces, isotopy, Thom's lemma, Whitney stratification, singularity

1 Introduction

The study of algebraic surfaces is a fascinating area where important developments of Mathematics such as singularity theory interact with visualization problems and the rendering of mathematical objects. The classification of singularities [5] provides simple algebraic formulas for complicated shapes, which may be difficult to handle geometrically. Such models can be visualized through techniques such as ray-tracing [1] in order to produce beautiful pictures of these singularities. Unfortunately, this approach does not allow to exploit the singularity models in applications other than static visualization.

[1] see e.g. `http://www.algebraicsurface.net/`

The aim of this paper is to describe an algorithm which produces a mesh of an algebraic surface S, with the guarantee that the topology of the surface is caught correctly. Such a piecewise linear model of a singular surface can be used in geometric modeling, coupled with refinement methods, for approximation or simulation purposes. As a by-product, it yields important topological information such as the number of connected components of the surface in a ball or a box, the Euler characteristic, ...

The related problem of determining the connected components of a semi-algebraic set and a path between two points of the same connected component has been investigated for instance in [11], [33]. Properties of the polar variety or silhouette curve of the semi-algebraic set and non-explicit Whitney stratifications were used to define these so-called roadmaps which provide a path between two given points.

The problem of triangulating a (semi)-algebraic set has been studied in the literature [23], [24], mainly from a theoretical point of view. See also [8], [13], [6] for a more introductory and computational point of view.

The special case of surfaces in $\mathbb{R}^3$ already received a lot of attention: we refer in particular to [18], [17], [9], [1], but these works deal only with smooth surfaces.

Another trend for tackling this triangulation problem is via Cylindrical Algebraic Decomposition [12]. It consists in decomposing a semi-algebraic set S into cells, defined by sign conditions on polynomial sequences. Such polynomial sequences are obtained by (sub)-resultant computations, corresponding to successive projections from $\mathbb{R}^{k+1}$ to $\mathbb{R}^k$. The degree of the polynomials in these sequences is bounded by $\mathcal{O}(d^{2^{n-1}})$ and their number by $\mathcal{O}((md)^{3^{n-1}})$, where m is the number of polynomials defining the semi-algebraic set S, d is a bound on the degree of these polynomials and n the number of different variables appearing in these polynomials [6]. For the case of implicit surfaces in $\mathbb{R}^3$ $(m = 1, n = 3)$, this yields a bound of $\mathcal{O}(d^4 \times d^9) = \mathcal{O}(d^{13})$ points to compute in order to get the topology of the surface. This Cylindrical Algebraic Decomposition does not directly yield a triangulation, nor any global topological information on the set S because the representation lacks information about the adjacency of the cells. Additional work is required to obtain a triangulation of S (see [13], [6], [29]). See also [7] for a new variant of the adjacency algorithm. Recently, this Cylindrical Algebraic Decomposition approach has been further investigated in [3]. It is shown how to analyze the topology of critical sections of an implicit surface, by exploiting the properties of delineability.

Our aim here is to describe an effective (and efficient) method for the triangulation of the part of a real algebraic surface S of $\mathbb{R}^3$ that lies inside a sphere. It can be generalized to other bounding shapes than spheres such as boxes, but for the sake of clarity we will stick to a spherical bounding shape throughout the rest of the article. The method is based on the computation of characteristic points on this surface. As we will see, it requires the computation of $\mathcal{O}(d^7)$ points. We follow a sweeping plane approach and exploit the following idea: after choosing a generic sweeping plane direction, the topology of the sections of the surface with this plane only changes for a discrete set of positions C. Computing this set of critical values (or more precisely a sup-set $C' \supset C$) and the topology of the sections at these critical values, will allow us to recover the topology of the surface. For this purpose, we will use the contour curve of S, which is a 3D curve on S. Our approach exploits results from stratified Morse theory. We give an explicit Whitney stratification of S, involving resultant computation and prove its correctness using equi-singularity arguments. This ensures the cylindrical structure of the surface between the critical sections that we have computed and yields a way to connect them, by "following" the contour curve.

The paper is organized as follows. In the next section, we recall basic definitions and describe the set of interesting points on the surface that we use to deduce its topology. In section 3, we describe how we treat the critical section of surface. In section 4, we describe how we compute the topology of the polar curve. In part 5, we describe the algorithm for surfaces and in particular how to connect two consecutive sections and obtain a triangulation of this connection while keeping safe the topology. We will see, in particular, how a discrete description of the polar variety allows us to recover the two dimensional faces of a triangulation of the surface. In section 6, we prove the correctness of the algorithm, showing as a new result, how a resultant computation yields a Whitney stratification of the surface. The proof of correctness of the connection algorithm is given and the isotopy between the surface S and its triangulation is made explicit. An example is given in section 7. Finally, we detail effectiveness and complexity issues in section 8.

2 Notations

We consider an algebraic surface S defined by the equation $f(x, y, z) = 0$ (with $f \in \mathbb{R}[x, y, z]$) in a given ball B for the Euclidian distance (instead of a ball B, we could also consider a box, but the description of the method is less simple). Hereafter, to simplify the presentation, we will assume that the boundary of B is not included in S. We denote by $S_B = S \cap B$ the intersection of S with the closed volume defined by the ball B. Our objective is to compute a simplicial complex, isotopic to the surface S_B.

We denote by π_y (resp. π_z, $\pi_{y,z}$), the projection of $\mathbb{R}^3$ along the y direction (resp. the z direction, the (y, z) plane) on the (x, z) plane (resp. the (x, y)

plane, the x axis). A (x, y) plane section of a variety V of $\mathbb{R}^3$ will be the intersection of V, with a plane parallel to the (x, y) plane (and similarly for the other variables).

A point $p \in V$ for which the matrix $[\nabla(f_1)(p), \ldots, \nabla(f_s)(p)]$ is not of maximal rank is called a *singular* point of V, where $\nabla(f) = [\partial_{x_1}(f), \ldots, \partial_{x_n}(f)]$.

The notion of critical point for a projection is a key notion of our approach, this is how we define it.

Definition 2.1 (Critical point) *For any algebraic variety V in $\mathbb{R}^n$ defined by equations $f_1 = 0, \ldots, f_s = 0$ and any linear map $\pi : \mathbb{R}^n \to \mathbb{R}^m$, a point p of V is said to be critical for the map π if the matrix $[\nabla(f_1)(p), \ldots, \nabla(f_s)(p), \pi]$ is not of maximal rank.*

A point $p \in \mathbb{R}^3$ of an algebraic variety $V \subset \mathbb{R}^3$ is x-critical (resp. (x, y)-critical) if it is critical for the projection $\pi_{y,z}$ (resp. π_z) on the x axis (resp. (x, y) plane). If $V \subset \mathbb{R}^3$ is defined by the polynomial equations $f_1 = 0, \ldots, f_s = 0$, a x-critical point of V is either a singular point or a point where the tangent space of V at this point is in a plane parallel to the (y, z) plane i.e the multiplicity of the intersection of the plane with the ideal $(f_1, \ldots, f_s)$ at p is greater or equal to 2. The corresponding x-coordinate of p is called a *x-critical value*. If a value is not x-critical, it is called x-regular. We use similar notations for the other variables.

2.1 The contour curve

Hereafter, we will use the properties of the contour curve of $S_B = S \cap B$. The contour curve is in fact a polar curve of S augmented with information to take into account the interference of S with the ball B.

Definition 2.2 *The polar curve of S for the projection π_z in the z-direction is the locus of the critical points of S for the projection along the direction z.*

If S is defined by $f(x, y, z) = 0$, this polar curve is defined by the equations $f(x, y, z) = \partial_z f(x, y, z) = 0$.

In order to take into account the restriction of S to B, we use the following definition for the contour curve:

Definition 2.3 *We denote by $\mathcal{C} := \mathfrak{C}_z(S_B)$ the union of*

- *the set of points $p \in B$ on the polar curve of S in the z-direction,*
- *the intersection of S with the boundary ∂B of the ball B.*

In other words $\mathfrak{C}_z(S_B) = (V(f, \partial_z f) \cap B) \cup (V(f) \cap \partial B)$. *We will call it the contour curve of* S.

The equations of the intersection of S with the boundary of B are $f(x, y, z) = 0$ and $q(x, y, z) = 0$, where q is the quadratic polynomial of the sphere associated to B. How to compute the topology of the contour curve is described in section 4. If we had used a box instead of a ball for the domain B, it would have been necessary to use the restrictions of $f(x, y, z)$ to the faces of the box B and the 2D algorithm (see section 3) to compute the topology of the corresponding planar curves.

2.2 Characteristic points on the surface

Let $f \in \mathbb{R}[x, y, z]$ be a square-free polynomial and let $S = V(f)$ be the surface it defines. Let $q(x, y, z)$ be the quadratic polynomial associated to the ball B. We denote by $R(x, y) = \mathrm{Res}_z(f(x, y, z), q(x, y, z)\, \partial_z f(x, y, z))$ and by $\Delta(x, y)$ its square-free part. Let $\mathcal{C}_{x,y} \subset \mathbb{R}^2$ be the planar curve defined by $\Delta(x, y) = 0$. For any $x_0 \in \mathbb{R}$, the points of $\mathcal{C}_{x,y} \cap V(x - x_0)$ are the projections on the (x, y) plane of the points of $S \cap V(x - x_0)$ that are either singular or smooth with a vertical tangent or in ∂B.

The algorithm for computing the topology of S will isolate the real solutions of the following system:

$$\begin{cases} \Delta(x, y') = 0 \\ \partial_y \Delta(x, y') = 0 \\ f(x, y, z') = 0 \\ q(x, y, z')\, \partial_z f(x, y, z') = 0 \\ f(x, y, z) = 0 \end{cases} \tag{1}$$

We denote by $\Xi(f)$ the set of real solutions of this system. This system can be assumed to be 0-dimensional over the complex field (since $V(q)$ is not in S) as we can perform a change of coordinates to put it in general position. An alternative way to say this, is that one can use a coordinate system different from x, y and z. The invariance of the sphere under rotations makes this step easy, but there is no substantial obstruction to developing the same algorithm with another bounding shape provided one is able to take into account the effect of the coordinate change on the bounding shape.

Notice that if $(\alpha, \beta, \gamma, \beta', \gamma') \in \Xi(f)$, then

- (α, β, γ) is a point on S,

- (α, β, γ') is a point on the contour curve of S,
- (α, β') is a critical point of $\mathcal{C}_{x,y} \subset \mathbb{R}^2$ for the projection to the x axis.

We associate to a solution $(\alpha, \beta, \gamma, \alpha', \beta')$ of the system (1) the following index:

- x if $\gamma = \gamma'$ and $\beta = \beta'$,
- c if $\gamma = \gamma'$ and $\beta \neq \beta'$,
- r otherwise.

A point $(\alpha, \beta, \gamma, \alpha', \beta')$ has index x if and only if (α, β, γ) is a point of the contour curve of S, which projects onto a x-critical point of $\mathcal{C}_{x,y}$.

A point $(\alpha, \beta, \gamma, \alpha', \beta')$ has index c if and only if (α, β, γ) is a point of the contour curve $\mathcal{C}$ of S, which projects onto a regular point of $\mathcal{C}_{x,y}$. Thus it is also smooth on $\mathcal{C}$.

For a point $(\alpha, \beta, \gamma, \alpha', \beta')$ with index r, (α, β, γ) is a smooth point of S, on the same vertical line as a point of $\mathcal{C}_{x,y}$ but not on the contour curve.

The intersection of the surface S with a plane $x = \alpha$ where α is the first coordinate of a solution of the system (1), will be called a x-critical section of S (at $x = \alpha$) and denoted by S_α.

3 Topology of the x-critical sections

In this section, we describe how we compute the topology of the x-critical section $S_\alpha = S \cap V(x - \alpha)$ at a x-critical value α. We use the subdivision approach presented in [2] to determine the topology of S_α. In the following we outline shortly the strong points of the method and how it proceeds.

The algorithm works on a square-free polynomial and $f(\alpha, y, z)$ will always be square-free in our algorithm. Otherwise the contour curve $\mathcal{C}$ would contain the x-critical section S_α, but this can't happen in the generic positions we allow (see definition 5.1 later to see how we enforce that condition). To make explanations clearer in the rest of this section we drop the first component of f, and consider it as function in the y and z variables (i.e. "$f(y, z) = f(\alpha, y, z)$"). The algorithm works by covering the disk $(V(x - \alpha) \cap B)$ where we want the topology, with rectangular boxes in which we know how to compute the topology. An important feature of this subdivision approach is that, unlike sweeping methods, it does not require any genericity assumptions, and will work in the given coordinate system, taking advantage of the potential sparsity of its input. This feature is important for our usage because when cutting over a singular point of $\Delta(x, y)$ there is no need for an additional change of variables (over such points there are often two singular points with same y coordinate).

The two categories of boxes for which we can determine the topology are the following:

- Regular boxes where either $\partial_y f(\alpha, y, z)$ or $\partial_z f(\alpha, y, z)$ does not vanish.
- Star-singular boxes where the topology in the box is star like. This means that to triangulate the portion of curve inside the box it suffices to pick any point in the interior of the box and to connect by a straight line to all the points of the curve that lie on the boundary of the box.

For the algorithm to be complete, we need to explain how we can effectively cover the disk with such boxes and explain how we manage to recover the topology for a box when it falls into one of the two above categories.

In the first step, we consider the points of $\Xi_\alpha(f)$ of index c, x and refine their isolating boxes until all the extremal points of f which are not on S lie outside the box. To determine that a box is star-singular we use the following criterion:

Lemma 3.1 (Star-singular box) *Let* $\deg(F_1, F_2, D)$ *denote the topological degree of a continuous map* $F : \mathbb{R}^2 \to \mathbb{R}^2$ *in a connected domain* $D \subset \mathbb{R}^2$ *[27]. If* D *is a box containing a singular point* p *such that* p *is the only extremal point of* f *in the box (i.e.* $\forall u \in D,\ \partial_y f(u) = \partial_z f(u) = 0 \Rightarrow p = u$ *and* $f(p) = \partial_y f(p) = \partial_z f(p) = 0$*) and if in addition the number of zeros of* f *on the boundary of* D *is* $2(1 - \deg(\partial_y f, \partial_y f, D))$*, then the topology in* D *is star like.*

This stems from the fact that $2(1 - \deg(\partial_y f, \partial_y f, D))$ is the number of half-branches at p [26]. Computing the topological degree is made possible by a formula that expresses it as a function of the values of f on the boundary of D, therefore it is possible to compute it using univariate solving on the segments of the boundary of D.
The second step is then to refine the isolation of the singular points until the topological degree in the box matches the number of zeros of f on the boundary.
The third step is quite straightforward, we refine the star-singular boxes so that they do not intersect the boundary of the disk where we want the topology and then we create a subdivision that contains all the isolation boxes we have computed so far. Therefore we end up with a set of regular boxes and star-singular boxes. If a singular point unluckily happens to be on the boundary of the disk, it is not a problem as it is possible to handle this case by counting the number of half-branches that lie inside the disk.
In the final step, we compute the topology in star-singular boxes by connecting a point inside the box to the point of the curve on the boundary. For regular boxes, we explain how the connection algorithm works if $\partial_z f$ does not vanish, the treatment of the case where $\partial_y f$ does not vanish is symmetrical. If $\partial_z f$ does not vanish, then there is no vertical tangent, therefore we can orient the

curve segments that lie in the box from left to right (i.e. according to their y component). This gives a formal meaning to "entering" the box (leftmost endpoint) and "exiting" the box (rightmost endpoint). Finally, because the curve segments cannot intersect each other, if we take the leftmost exit point, the corresponding entry point has to be the first point encountered to its left on the boundary of the box (because there are only entry points to its left). So we just connect these two points together, remove them from the list of points to be connected, and repeating this process recursively eventually gives the topology of the curve in the box. For more details, see [2].

4 Topology of the contour curve

The algorithm to compute the topology of the contour curve, exploits the 2D algorithm [2] described in the previous section, combined with the algorithm in [19]. We use two projections of the 3D curve to recover the connection of the points above these projected planar curves and the points in $\Sigma(f)$ to analyze the critical points of the projected curves. The restriction in [19] that $(f, \partial_z f)$ has to be a radical ideal can be removed, since we deduce the critical points of the 3D curve from the points of $\Sigma(f)$.

Other approaches can be used here to compute the topology of the 3D contour curve. One can use for instance the algorithm in [4] (if $(f, \partial_z f)$ is radical), the main difference being the genericity conditions which are required and the technique used to lift points from the (x, y) or (x, z) plane to 3 dimensional space. In [19], the genericity conditions are related to the projection of the curves on the x axis, whereas in [4] they are related to the projection on the (x, y) plane and to the projection of this projection on the x axis (which is more restrictive). The effective techniques described in [4] to check this genericity condition involve delicate computation such as approximate gcd or absolute factorization, in particular in the presence of singularities. In another recent approach [16], the 3D curve is described by its projection on the plane (x, y) and by a reduced "monoid" equation $a(x)\,z - b(x, y) = 0$. This allows to lift the planar curve and to deduce the connection above the critical points, under some genericity conditions. The polynomial $a(x)$ is obtained from an iterated resultant and may be huge. In another recent work [30], non-reduced curves are treated using rational lifting maps deduced directly from the decomposition of subresultants with respect to the variable z.

As opposed to the aforementioned methods, the approach that we are going to describe here does not require genericity conditions on the projected curves but only pseudo-generic conditions (two branches of the contour curve do not project on the same branch in the (x, y) or (x, z) plane).

The general idea is to project the contour curve onto the (x, y) plane and (x, z) plane, and to compute the topology of the projected curves in order to recover the topology of the 3D contour curve.

We will use the subset $\Xi_{\mathcal{C}}(f)$ of points of $\Xi(f)$ that have index c or x. Points of $\Xi(f) \in \mathbb{R}^5$ naturally project to points of S by taking their first three components. The role of the fourth and fifth components is to allow us to label them as x, c, and r points. Once we have this information we can discard the two last components, and to simplify the discussion we will in the following, consider the points in $\Xi_{\mathcal{C}}(f)$ as the points on $S \subset \mathbb{R}^3$ to which they project. In this way, points with index c are smooth points on $\mathcal{C}$ (since their projection to $\mathcal{C}_{x,y}$ is smooth). We will also use the points of $\mathcal{C}$ at intermediate sections $x = \mu$, chosen adequately as we describe now.

Let $\Delta(x, y)$ be the square-free part of $\mathrm{Res}_z(f, q\,\partial_z f)$ and $\mathcal{C}_{x,y}$ the curve it defines in the plane (x, y). We also denote by $\Psi(x, z)$ be the square-free part of $\mathrm{Res}_y(f, q\,\partial_z f)$ and by $\mathcal{C}_{x,z}$ the curve it defines in the plane (x, z).

In a first step, we compute the topology of the curve $\mathcal{C}_{x,y}$ (see section 3) in the projection of the bounding ball B where we want to determine the topology.

Let Σ be the x-critical values of $\mathcal{C}_{x,y}$ and Σ' the x-critical values of $\mathcal{C}_{x,z}$: $\Sigma = \{\sigma_1, \ldots, \sigma_s\}$ with $\sigma_1 < \cdots < \sigma_s$. For each $\sigma_i \in \Sigma$, we compute two (rational) values μ_i, μ_i' such that $\sigma_{i-1} < \mu_i < \sigma_i < \mu_i' < \sigma_{i+1}$ and $\Sigma' \cap [\mu_i, \sigma_i[= \Sigma' \cap]\sigma_i, \mu_i'] = \emptyset$. Note that Σ and Σ' can have points in common, that's why the intervals are open in σ_i (in fact if there is a y, z-critical point they will have a σ_i in common).

In the following, we denote by $\mathcal{C}_{\mu_i}$ the section $\mathcal{C} \cap V(x - \mu_i)$. By construction, above the interval $[\mu_i, \sigma_i[$ the curves $\mathcal{C}_{x,y}$ and $\mathcal{C}_{x,z}$ have no x-critical points. If two points of $\mathcal{C}_{\mu_i}$ have the same y-coordinate, and if the projection $\mathcal{C}_{x,y}$ has no critical point at $x = \mu_i$, then two branches of $\mathcal{C}$ project onto the same branch of $\mathcal{C}_{x,y}$. By a generic change of coordinates, we can avoid this situation. We proceed similarly, if two points of $\mathcal{C}_{\mu_i}$ have the same z-coordinates. Thus we can assume that $\mathcal{C}_{\mu_i}$ projects injectively on the (x, y) and (x, z) planes.

In order to connect the points of $\mathcal{C}_{\mu_i}$ to those of $\mathcal{C}_{\sigma_i}$, we also compute the topology of $\mathcal{C}_{x,z}$ above the interval $[\mu_i, \mu_i']$ using $\Sigma(f)$ (see section 3). Notice that by construction of μ_i, μ_i', the projection of $\Sigma(f)$ on the (x, z) plane contains the z-critical of the projected curve, above the interval $[\mu_i, \mu_i']$. Using the computed topological graph of $\mathcal{C}_{x,y}$ and $\mathcal{C}_{x,z}$, we proceed as follows.
Given a point $p = (\mu_i, v, w) \in \mathcal{C}_{\mu_i}$, its projection (μ_i, v) is connected to a point (σ_i, β) by the topological graph that we have computed for $\mathcal{C}_{x,y}$. Its projection (μ_i, w) is connected to a point (σ_i, γ) by the topological graph of $\mathcal{C}_{x,z}$. As the projections of $\mathcal{C}_{\mu_i}$ onto the planes (x, y) and (x, z) are injective, there is a (unique) branch of $\mathcal{C}$, which connects p to the point $(\sigma_i, \beta, \gamma) \in \mathcal{C}_{\sigma_i}$.

The connection above the interval $[\mu'_{i-1}, \mu_i]$ proceeds similarly by using only the topological graph of the curve $\mathcal{C}_{x,y}$, which has no x-critical values in $[\mu'_{i-1}, \mu_i]$, since $\mathcal{C}_{\mu'_{i-1}}$ and $\mathcal{C}_{\mu_i}$ project injectively to the (x, y) plane.

Let us summarize the main steps of this algorithm:

Algorithm 4.1 (Topology of $\mathcal{C}$ defined by $f_1(x, y, z) = 0, f_2(x, y, z) = 0$)

INPUT: *Polynomials $f_1, f_2 \in \mathbb{Q}[x, y, z]$ and a box $D_0 \subset \mathbb{R}^3$*

- *Compute the square-free part $\Delta(x, y)$ of $\mathrm{Res}_z(f_1, f_2)$, defining the projected curve $\mathcal{C}_{x,y} \subset \mathbb{R}^2$.*
- *Compute the square-free part $\Psi(x, z)$ of $\mathrm{Res}_y(f_1, f_2)$, defining the projected curve $\mathcal{C}_{x,z} \subset \mathbb{R}^2$.*
- *Compute the topology of $\mathcal{C}_{x,y}$ in the projection of D_0 on the plane (x, y).*
- *Compute the x-critical values $\Sigma := \{\sigma_1, \dots, \sigma_k\}$ with $\sigma_1 < \cdots < \sigma_k$. of $\mathcal{C}_{x,y}$ and the critical values Σ' of $\mathcal{C}_{x,z}$.*
- *Choose a (rational) $\mu_i \in]\sigma_{i-1}, \sigma_i[$ (resp. $\mu'_i \in]\sigma_i, \sigma_{i+1}[$) such that $[\mu_i, \sigma_i[\cap \Sigma' = \emptyset$ (resp. $]\sigma_i, \mu'_i] \cap \Sigma' = \emptyset$).*
- *Compute the topology of $\mathcal{C}_{x,z}$ above $[\mu_i, \mu_i]$ in the projection of D_0 on the plane (x, z).*
- *Compute the set $\mathcal{C}_{\mu_i}$ of real points of $\mathcal{C}$ at $x = \mu_i$ and check that it is finite and that two points do not have the same y-coordinates (resp. z-coordinates). If this is the case, raise the exception "non-generic position".*
- *Use the topology of $\mathcal{C}_{x,y}$ and $\mathcal{C}_{x,z}$ above $[\mu_i, \sigma_i]$ (resp. $[\sigma_i, \mu'_i]$ to connect the points of $\mathcal{C}_{\sigma_i}$ to $\mathcal{C}_{\mu_i}$ (resp. $\mathcal{C}_{\mu'_i}$).*
- *Use the topology of $\mathcal{C}_{x,y}$ above $[\mu'_{i-1}, \mu_i]$ to connect the points $\mathcal{C}_{\mu'_{i-1}}$ to $\mathcal{C}_{\mu_i}$.*

OUTPUT: *The graph of 3D points of the curve connected by segments, isotopic to the curve $\mathcal{C}$ or the exception "non-generic position".*

This algorithm is applied for $f_1 = f$, $f_2 = q\, \partial_z f$ where $q(x, y, z) = 0$ is the equation of the boundary of B, to get the topology $\mathcal{C} = \mathfrak{C}_z(S_B)$. We need the topology of the contour curve because there are topology changes that come from the interference of the bounding sphere B with S, and this is taken into account by adding $S \cap B$ into the contour curve. The way this comes into play is explained in the next section.

Since we are interested in the topology of $S \cap B$, we only need to compute the topology of the curves $\mathcal{C}_{x,y}$ or $\mathcal{C}_{x,z}$ in boxes of $\mathbb{R}^2$ which are the projection of a box in $\mathbb{R}^3$ containing B.
In order to compute the x-critical values of $\mathcal{C}_{x,y}$ or $\mathcal{C}_{x,z}$, we apply iterated resultant computations. Though the degree of these resultant polynomials can grow quickly, they can be decomposed into explicit factors in order to simplify the computation (see [10]).

5 The algorithm for singular algebraic surfaces

5.1 The algorithm

We will assume hereafter that the surface is in a generic position:

Definition 5.1 *We say that the surface is in generic position if*

- *the system* (1) *has a finite number of (complex) solutions.*
- *two distinct arcs of the contour curve do not have the same projection in the* (x, y) *(resp.* (x, z)*) plane.*

The first point is checked while solving system (1). If it is a zero dimensional system, we assume that the polynomial solver over the complex field yields isolating boxes containing one and only one real root of $\Xi(f)$. The second point is checked while applying the algorithm 4.1 to $f(x, y, z)$ and $q(x, y, z)\,\partial_z f(x, y, z)$. If these conditions are not fulfilled, we perform a random change of coordinates and restart the algorithm. There is a high probability to be in generic position after a change of coordinate, and therefore this process eventually stops and yields a surface in generic position.

Let us first outline briefly the algorithm, before going into the details.

The first step consists in computing the contour curve for the projection in the z-direction. We apply algorithm 4.1 for 3D-curves with $f_1 = f$, $f_2 = q\,\partial_z f$, which computes a polygonal approximation of the contour curve which is isotopic to it. Doing this, the algorithm computes x-critical values corresponding to x-critical points of the 3D curve and singular points of the projection of the contour curve on the (x, y) plane and the (x, z) plane. Let us call Σ this set of x-critical values.

For each σ of Σ, we compute the topology of the corresponding sections of the surface, by applying algorithm for the topology of 2D curves (see section 3).

Next, we compute regular values between two critical values and the topology of the corresponding sections. Here again, we use the 2D algorithm for implicit curves (see section 3).

The following step consists in connecting two consecutive sections, using the topology of the contour curve (see details in section 5.2).

Finally, we mesh the resulting patches of the surface, by computing a set of points, open segments and open triangles, which are not self-intersecting, and which defines a simplicial complex isotopic to the surface (see details in section

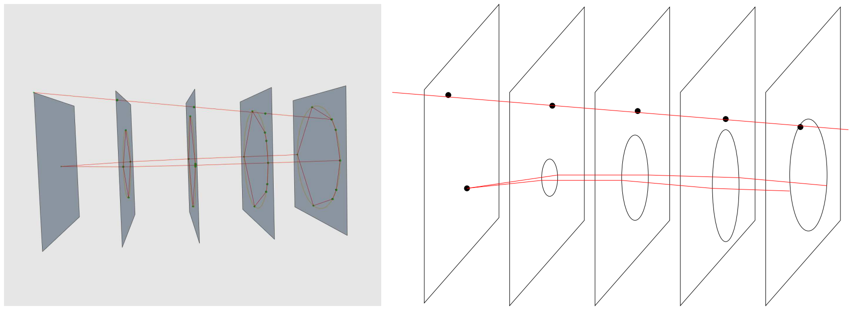

Fig. 1. Polar variety and first connections for the union of a sphere and a line defined by one equation.

5.3).

We summarize the algorithm as follows:

Algorithm 5.2 (Topology of an algebraic surface S in a ball B)
INPUT: *A polynomial $f(x, y, z)$ defining S and a ball B.*

- *Compute the topology of the contour curve for the projection in the z-direction, using algorithm 4.1.*
- *Compute the topology of the sections, using algorithm 3.*
- *Connect two consecutive sections, by exploiting the topology of the contour curve.*
- *Triangulate the resulting surface patches, avoiding self-intersection of segments and triangles.*

OUTPUT: *A simplicial complex isotopic to S_B.*

Let us now detail the two last steps of this algorithm.

5.2 Connection algorithm

We denote by $\mathcal{V}$ the topological description of $\mathcal{C} = \mathfrak{C}_z(S_B)$ and $\mathcal{K} := \mathcal{V}$ the initial value of the topological complex describing S. The initial value for $\mathcal{K}$ is the result of the curve topology computation done for $\mathfrak{C}_z(S_B)$ by algorithm 4.1. We are going to update this complex by explaining how we define the connections between the points of two successive sections of S, a regular one which is regular S_r and a critical one S_c which contains a x-critical point of $\mathfrak{C}_z(S_B)$. By section of S we mean a set $S \cap V(x - \alpha)$ where α is such that $V(x - \alpha)$ contains no x-critical point of $\mathfrak{C}_z(S_B)$ for regular sections S_r and does contain such a point for critical sections S_c.

Let us denote by $p_1, \ldots, p_r$ (resp. $q_1, \ldots, q_s$) the points of $\pi_z(\mathcal{V} \cap S_r)$ (resp. $\pi_z(\mathcal{V} \cap S_c)$) ordered by increasing y-coordinates, which are also on the projec-

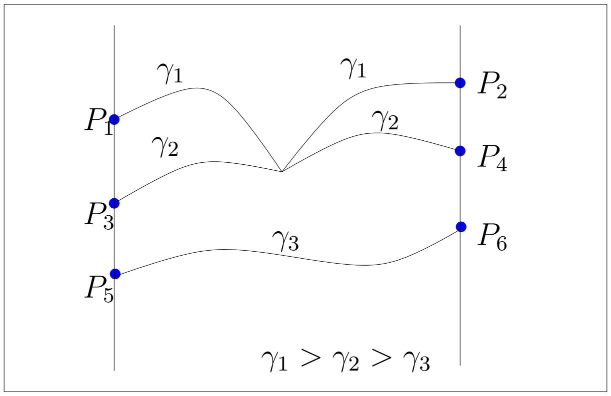

Fig. 2. Order on the arcs.

tion of an arc of $\mathcal{V}$ connecting S_r and S_c.

Hereafter, we use the convention that $p_0, p_{l+1}, q_0, q_{m+1}$ are points on the boundary of the ball B. We denote by $\mathcal{A}_i$ $(i = 0, \ldots, l)$ the set of arcs of S_r which projects onto $[p_i, p_{i+1}]$. We denote by $\mathcal{B}_j$ $(j = 0, \ldots, m)$ the set of arcs of S_c, which connect a point projecting at q_j to a point projecting at q_{j+1}. If, moreover there is a critical point U in between, we require that if this arc is to the k^{th} branch arriving at U on the left, then it is also the k^{th} branch starting from U on the right, if this branch exists.

The arcs in $\mathcal{A}_i$ (resp. $\mathcal{B}_j$) are naturally ordered according to their z-position: an arc is bigger than another if it is above the other (see Figure 2). We treat incrementally the points p_i, starting from p_0. Let us denote by $q_{\nu(i)}$ the point connected to p_i by an arc of $\pi_z(\mathcal{K}) \supset \pi_z(\mathcal{V})$.

- If p_{i+1} is connected to $q_{\nu(i)}$ by an arc of $\pi_z(\mathcal{K})$, for any arc $\gamma = (P, P')$ of $\mathcal{A}_i$, such that P is connected by $\mathcal{K}$ to Q, we add the arc (P', Q) and the face (P, P', Q) to $\mathcal{K}$.
- If p_{i+1} is not connected to $q_{\nu(i)}$, it is connected to $q_{\nu(i)+1}$. We consider the smallest arc $\gamma = (P, P')$ of $\mathcal{A}_i$, and the smallest arc $\eta = (Q, Q')$ of $\mathcal{B}_{\nu(i)}$. The arc (P, Q) is in $\mathcal{K}$. We add the arc (P', Q') and the face (P, P', Q', Q) to $\mathcal{K}$. Then we remove these smallest arcs γ and δ, respectively from $\mathcal{A}_i$ and $\mathcal{B}_{\nu(i)}$ and go on until $\mathcal{A}_i$ is empty.

This procedure is applied iteratively, until we reach the point p_r, so that we move to the next section S'_r, S'_c.

5.3 Triangulation algorithm

The final step is the triangulation of the different faces computed previously.

Assume that in the connection algorithm (section 5.2), we have connected an arc $\gamma = (P_1, P_2)$ of S_r to an arc (or point) $\eta = (Q_1, Q_2)$ in S_c, by a face of $\mathcal{K}$.

The triangulation algorithm works as follows (see Figure 3):

If $Q_1 = Q_2$ then we connect successively all the points between P_1 and P_2 to Q_1, creating the triangles of our triangulation.

If $Q_1 \neq Q_2$, there can exist at most one critical point U on η. The situation to avoid is described in Figure 4. If we do not pay attention to the connections that are created during the triangulation, we can create intersection curves between two faces that do not exist. We will quickly explain on an example what we have to do to avoid that before going back to the general algorithm. We see on Figure 4 a situation where the arc γ_1 (resp. γ_2) is connected to an arc η_1 (resp. η_2). The two patches defined respectively by γ_1 and η_1 and by γ_2 and η_2 connects to the arcs P_1Q_1 and P_2Q_2 but do not intersect. To create an intersection, we would need to connect a point shared by γ_1 and γ_2 to a point shared by δ_1 and δ_2. This case corresponds to the drawing in Figure 4. So what has to be done is simply to connect the point U (the only point different from Q_1 and Q_2 belonging to the two arcs η_1 and η_2) to a point different from P_1 and P_2. This can always be done because there exist intermediate points between P_1 and P_2.
If there exists a y-critical point U on the arc $\eta = (Q_1, Q_2)$ then, we connect the point U to an intermediate point T_1 of γ between P_1 and P_2 (see Figure 4). Let us now consider the two sub-arcs (P_1, T_1) and (Q_1, U). We start simultaneously from P_1 and Q_1. The two points are connected by an arc of $\mathcal{K}$. We consider the next point on the arc (P_1, U) and the next point on the arc (Q_1, U). We connect them. This process goes on until there are no more points on one of the two arcs. If there is less points on an arc than on the other, we connect the remaining points on one arc by adjacent triangles sharing the same vertex (see Figure 5). After this step, we obtain triangles or quadrangles, which we subdivide in order to obtain the final triangulation.

6 Why we get the topology

As mentioned previously, the general idea of this sweeping algorithm is to detect where *some topological changes* in the intersection of S with the sweeping plane happen so that in-between the topology is fixed. We are going to prove that in-between the events that we have computed in the previous section, the topology is locally trivial and we use this result to describe explicitly the isotopy between the mesh and the surface.

To prove the correctness of the algorithm we will use results from stratified

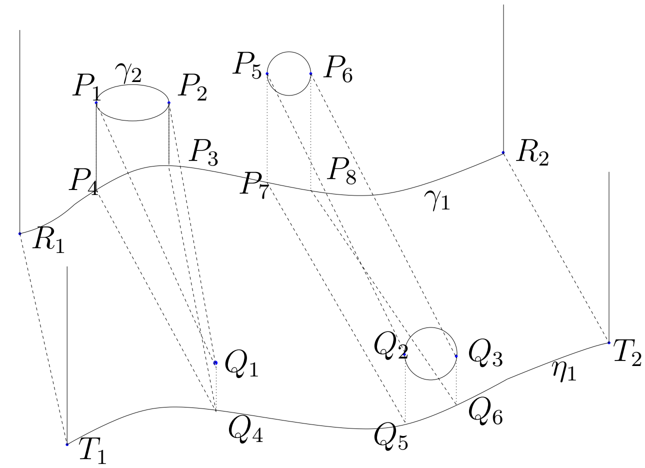

Fig. 3. Division of the space with vertical walls.

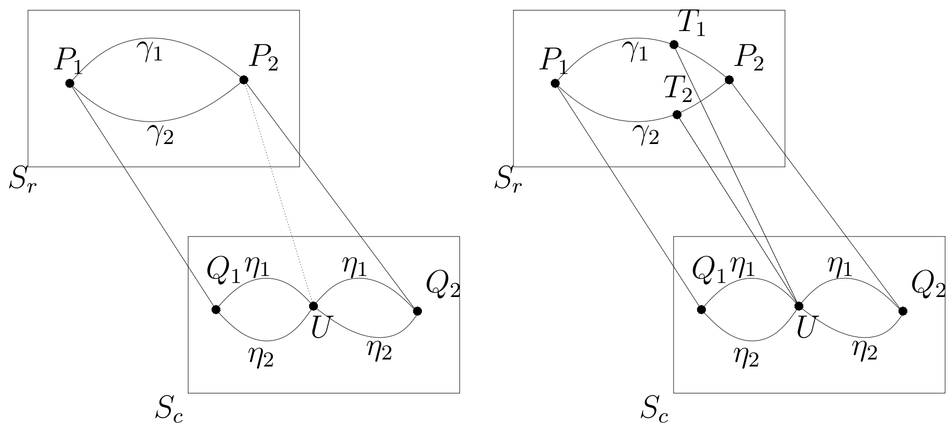

Fig. 4. Incorrect connection and its correction

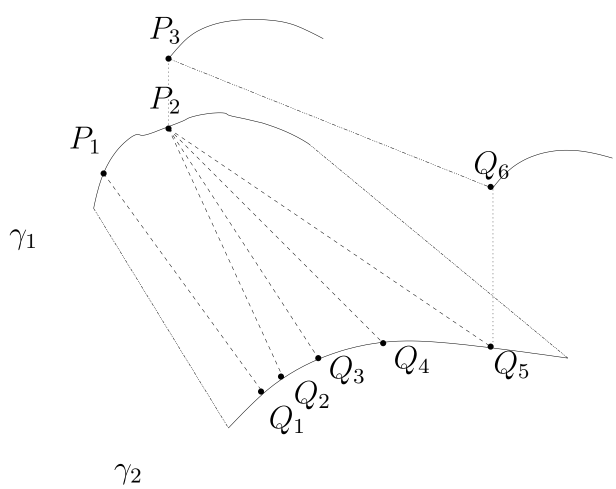

Fig. 5. Meshing.

Morse theory. We refer to [22], [14] for more details.

The fundamental notion is of Whitney stratification. It is a decomposition of the variety into smooth parts that fit together "regularly". Here are some definitions:

Definition 6.1 *[14, Def. 1.5, p.3] A stratification of a (semi-algebraic) variety $A \subset \mathbb{R}^n$ is a locally finite partition of A into smooth submanifolds.*

Definition 6.2 *[14, Def. 1.6, p.3] Let (X, Y) be two strata and $p \in \overline{X} \cap Y \subset \mathbb{R}^n$. X is Whitney-regular at p along Y if for any sequences $x_n \in X$, $y_n \in Y$ converging to p, $l = \lim_{n\to+\infty} \overline{x_n y_n} \subset \mathcal{T} = \lim_{n\to+\infty} T_{x_n}X$, where T_xX is the tangent space of X at the point x.*

A Whitney stratification of a variety S is a stratification of S so that all pairs of strata are Whitney-regular.

Remark 6.3 *The Whitney condition implies what is usually referred to as the boundary condition. That is, if the closure of two strata intersect, then one is included in the boundary of the other (see [28]).*

Theorem 6.4 *[14, Cor. 1.12, p.6] All semi-algebraic varieties $A \subset \mathbb{R}^n$ admit a Whitney stratification.*

Proposition 6.5 *Any semi-algebraic stratum S is Whitney regular along a zero-dimensional stratum.*

See [14, Lemma 1.10, p.5] for a proof, using the Curve Selection Lemma.

Definition 6.6 *For Z and W two stratified sets, a differential map $f : Z \to W$ is a stratified submersion at a point p of Z if the differential map at p of f, $Df : T_p(Z_\sigma) \to T_{f(p)}(W_\tau)$ is surjective. Where Z_σ and W_τ are the strata of Z and W containing p and $f(p)$.*

Definition 6.7 *If Z and W are two stratified sets, a continuous map $f : Z \to W$ is proper if the inverse image of any compact set of W is a compact of Z.*

The main theorem that we will use is Thom's lemma [22, Sec. 1.5, p.41].

Theorem 6.8 (Thom's first isotopy lemma) *Let Z be a Whitney stratified subset of $\mathbb{R}^m$ and $\pi : Z \to \mathbb{R}^n$ be a proper stratified submersion. Then there is a stratum preserving homeomorphism*

$$h : Z \to (\pi^{-1}(0) \cap Z) \times \mathbb{R}^n$$

which is smooth on each stratum and such that π factorizes via the projection *to the second component* $\mathbb{R}^n$.

$$\begin{tikzcd} Z \arrow[r, "\pi"] \arrow[d, "h"'] & \mathbb{R}^n \\ (\pi^{-1}(0) \cap Z) \times \mathbb{R}^n \arrow[ur, "\nu"] & \end{tikzcd}$$

This means that Z is homeomorphic to the cylinder with base $\pi^{-1}(0) \cap Z$.

In our case, we will apply the theorem with $Z = S_B$, $m = 3$, $n = 1$ and π the projection on the x axis which is automatically proper as we work in a ball B which is compact.

6.2 Computation of a Whitney stratification

For a projection π_z in the direction z on the (x, y) plane, we recall that $\Delta(x, y)$ is the square-free part of the resultant $\mathrm{Res}_z(f, q\,\partial_z f)$ and that its associated zero set $V(\Delta(x, y))$ is $\mathcal{C}_{x,y}$.

Theorem 6.9 *For a generic projection π_z, let*

- *S^0 be the set of points of $\mathfrak{C}_z(S_B)$ that projects by π_z onto singular points of $\mathcal{C}_{x,y}$, each point is considered as a stratum,*
- *S^1 the set of the connected components of $\mathfrak{C}_z(S_B) - S^0$, (each connected component is a stratum),*
- *S^2 the set of the connected components of $S - \mathfrak{C}_z(S_B)$ (each connected component is a stratum).*
- *S^3 the set of connected components of $\mathbb{R}^3 - S$ (each connected component is a stratum).*

Then (S^0, S^1, S^2, S^3) is a Whitney stratification of $\mathbb{R}^3$ compatible with S.

From Proposition 6.5 and as the Whitney regularity of any stratum of S^1 or S^2 along a stratum of S^3 is always fulfilled, we deduce that showing that (S^0, S^1, S^2, S^3) is a Whitney stratification of $\mathbb{R}^3$ compatible with S boils down to showing that (S^1, S^2) is Whitney-regular.

Depending on whether we consider the polynomial f defining S over $\mathbb{R}$ or $\mathbb{C}$, we obtain a real variety $S = S_{\mathbb{R}}$ or a complex variety $S_{\mathbb{C}}$, as the set of zeros of f. We will use the results of equi-singularity along $\mathbb{C}$ and the notion of "permissible" projection to prove the proposition. Speder gave in [32] a definition of permissible projection, stronger than the original one of Zariski [35]. We will consider only the case of codimension 1, for which both definitions coincide, so hereafter we will just consider the definition of permissible projection of Zariski:

Definition 6.10 *A permissible direction of projection for the pair (X, Y) with $Y \subset \overline{X}$ at $Q \in Y$ is an element of $\mathbb{PC}^3$ so that the line passing through Q defined by this direction is neither included in a neighborhood of Q nor in the tangent space to Q at Y.*

Proposition 6.11 *For a given algebraic surface S, a generic direction of projection is permissible for (S^1, S^2) at every point of S^1.*

Proof. For an algebraic variety, the local inclusion of a line into the surface is equivalent to a global inclusion. We deduce that the directions of projection to avoid are included into the union of:

- directions of lines included into the surface
- directions of the tangents to the smooth part of the singular locus of the variety.

We consider the first set of directions of lines included in the surface S, defined by one equation $f(x, y, z) = 0$. We consider the surface embedded in projective space. The directions of lines included in S considered as points of projective space are included in the intersection of S with the hyperplane at infinity which is a projective curve. Thus the directions corresponding to the first set are included in a set of dimension 1 and are generically avoided.

Now let us consider the second set. We consider an arc of the smooth part of the contour curve (there exists a finite number of such arcs for an algebraic surface). We consider a semi-algebraic parameterization of this arc $(x(s), y(s), z(s))$. Thus we obtain a semi-algebraic parameterization $(x'(s), y'(s), z'(s))$ of a set of unit vectors corresponding to the directions of the tangents to the curve. We deduce that the set to avoid (for the tangency condition) corresponds to a semi-algebraic curve on the unit sphere of $\mathbb{R}^3$ and is generically avoided. □

Proposition 6.12 *If the surface S is in generic position (see definition 5.1) then the projection π_z along the z axis is a permissible projection.*

Proof. First, there is no line parallel to the z axis in S because if this were the case, all the vertical line would be included in the contour curve and we would not be in generic position. The second point to check is that the z-direction is not a direction of a tangent of S^1. This is the case as by construction the points of the contour curve with vertical tangents project onto singular points of $\mathcal{C}_{x,y}$ and are thus in S^0. □

We also recall the notion of equi-singularity (which is defined inductively):

Definition 6.13 *[32, p. 577], [35, Def. 4.1 p.981]. Let $X \subset \mathbb{C}^n$ a hypersurface, Y a smooth submanifold of X of codimension c, P be a point of Y. We say that X is equi-singular at P along Y if either $c = 0$ and X is smooth at P or $c > 0$, $Y \subset X_{\mathrm{sing}}$ and there exists a so-called permissible projection π such that the critical locus of $\pi|_X$ is equi-singular at $\pi_z(P)$ along $\pi_z(Y)$.*

The main result that we use is the following:

Proposition 6.14 *[32, Def. 4.1 p.981] If the hypersurface $X \subset \mathbb{C}^n$ is equi-singular along Y and if the codimension of Y in X is 1, then the pair (X_{smooth}, Y) fulfills the Whitney conditions along Y.*

This allows us to check the Whitney condition over the complex field. We need to check it on $\mathbb{R}$:

Proposition 6.15 *If X and Y are two strata of a Whitney stratification of $S_{\mathbb{C}}$ with $\dim X = 2$ and $\dim Y = 1$, then $(X_{\mathbb{R}}, Y_{\mathbb{R}})$ is Whitney regular, where $X_{\mathbb{R}} = X \cap \mathbb{R}^3$ and $Y_{\mathbb{R}} = Y \cap \mathbb{R}^3$.*

Proof. Let P be a point in $Y_{\mathbb{R}} \cap \overline{X_{\mathbb{R}}}$. Consider a sequence x_n of points of $X_{\mathbb{R}}$ and y_n of points of $Y_{\mathbb{R}}$, both sequences converging to P. Note these sequences exist because P is in $Y_{\mathbb{R}} \cap \overline{X_{\mathbb{R}}}$ which means there are points of $X_{\mathbb{R}}$ in a neighborhood of P (and $P \in Y_{\mathbb{R}}$). Of course $Y_{\mathbb{R}}$ and $X_{\mathbb{R}}$ are disjoint sets because Y and X were already disjoint. We assume that the sequence of secants $\overline{x_n y_n}$ converges to a limit $l \in \mathbb{R}^3$ and the sequence of tangent planes $T_{x_n} X_{\mathbb{R}}$ converges to a limit T. If we consider x_n and y_n in $\mathbb{C}^3$, the sequence of secants converges also to a complex line $l_{\mathbb{C}}$ because $x_n \wedge y_n$ converges to a limit L in $\mathbb{P}(\Lambda^2 \mathbb{R}^4)$ which is embedded in $\mathbb{P}(\Lambda^2 \mathbb{C}^4)$. The convergence of the sequence $T_{x_n} X_{\mathbb{R}}$ is equivalent to the convergence of $T_{x_n} X$: the sequence of normals defined by the orthogonal vectors ∇f converges equivalently in $\mathbb{R}$ or $\mathbb{C}$. Thus $\lim_{n\to\infty} \overline{x_n y_n}_{\mathbb{R}} \subset \lim_{n\to\infty} \overline{x_n y_n}_{\mathbb{C}} \subset \lim_{n\to\infty} T_{x_n} X$ (since (X, Y) is Whitney regular). We deduce that $\lim_{n\to\infty} \overline{x_n y_n}_{\mathbb{R}} \subset \lim_{n\to\infty} T_{x_n} X \cap \mathbb{R}^3$. We know that $\lim_{n\to\infty} T_{x_n} X_{\mathbb{R}} \subset \lim_{n\to\infty} T_{x_n} X \cap \mathbb{R}^3$. As x_n is a sequence of real points, $\lim_{n\to\infty} T_{x_n} X$ is defined as the orthogonal in $\mathbb{C}$ of a real vector. We deduce that $\lim_{n\to\infty} T_{x_n} X \cap \mathbb{R}^3$ is a real space of dimension less or equal to 2 containing $\lim_{n\to\infty} T_{x_n} X_{\mathbb{R}}$ which is of dimension 2, thus the two linear spaces are equal. So we deduce that $\lim_{n\to\infty} \overline{x_n y_n}_{\mathbb{R}} \subset \lim_{n\to\infty} T_{x_n} X_{\mathbb{R}}$ and that $(X_{\mathbb{R}}, Y_{\mathbb{R}})$ is Whitney regular. □

Proof of Theorem 6.9. The stratification defined in (6.9) over the complex field, yields a stratification of $S_{\mathbb{C}}$. We consider its restriction to $\mathbb{R}^3$. By Proposition 6.5, we only need to check the Whitney condition for the 1-dimensional strata $S^1_{\mathbb{R}}$ and the 2-dimensional strata $S^2_{\mathbb{R}}$ of $S_{\mathbb{R}}$. Let $p \in S^1_{\mathbb{R}} \cap \overline{S^2_{\mathbb{R}}}$. If p is a smooth point of S, the Whitney condition is trivially satisfied. If p is singular, by Proposition 6.14, we have the Whitney condition for $(S^2_{\mathbb{C}}, S^1_{\mathbb{C}})$ at p. And by Proposition 6.15, we deduce the Whitney condition for $(S^2_{\mathbb{R}}, S^1_{\mathbb{R}})$ at p. This proves that $(S^0_{\mathbb{R}}, S^1_{\mathbb{R}}, S^2_{\mathbb{R}}, S^3_{\mathbb{R}})$ is a Whitney stratification of $\mathbb{R}^3$ compatible with

$S_{\mathbb{R}}$. □

6.3 Connection of the sections

We have described in section 5 an algorithm to connect two successive sections. Now we are going to justify what this algorithm does.

By Theorem 6.9 and using Thom's lemma (Theorem 6.8), we deduce that in between two consecutive critical sections, the topology of the sections is constant. We have computed the topology of regular sections, in between two successive critical ones. So now, in order to prove the isotopy of the surface and the mesh, we have three things to verify:

a) From a topological point of view, we define the good connections between the sections.
b) The triangulation that we construct is valid (i.e. the simplices of the complex do not intersect). Or in other words, the embedding in $\mathbb{R}^3$ of the simplicial complex we have constructed is injective.
c) The mesh is isotopic to the surface.

The point c) will be made explicit in subsection 6.4. We now prove the first two points:

a) We are going to justify the connection algorithm described in section 5.2. Let us recall the notations of section 5.2.

We denote by $p_1, \ldots, p_l$ (resp. $q_1, \ldots, q_m$) the points of $\pi_z(\mathcal{V} \cap S_r)$ (resp. $\pi_z(\mathcal{V} \cap S_c)$) ordered by increasing y-coordinates, which are on the projection of an arc of $\mathcal{V}$ connecting S_r and S_c. Notice that we have $s \leq r$.

We denote by $\mathcal{A}_i$ $(i = 0, \ldots, l)$ the set of arcs of S_r which project onto $[p_i, p_{i+1}]$ and by $\mathcal{B}_j$ $(j = 0, \ldots, m)$ the set of arcs of S_c which project onto $[q_j, q_{j+1}]$, with the convention that $p_0, p_{l+1}, q_0, q_{m+1}$ are on the boundary of the ball B.

The point p_i is connected to $q_{\nu(i)}$ by an arc δ_i of the projection of $\mathcal{K}$. We note Θ_i the open planar domain between δ_i and δ_{i+1} (dashed part in Figure 6).

Proposition 6.16 *If the topology of $\pi_z^{-1}(\delta_i) \cap S$, S_r, S_c is determined, then algorithm 5.2 computes the topology of $\pi_z^{-1}(\delta_{i+1}) \cap S$ and of $\pi_z^{-1}(\Theta_i) \cap S$.*

Proof. Let us consider an arc γ in $\mathcal{A}_i$ connecting a point P to a point P'. If we apply Thom's lemma to $S \cap \pi_z^{-1}(\Theta)$, we deduce that $S \cap \pi_z^{-1}(\Theta)$ is topologically trivial (i.e. made of a family of patches lying one above the other) and that the boundary of each patch contains an arc θ_i in $\pi_z^{-1}(\delta_i)$ and an arc θ_{i+1} in

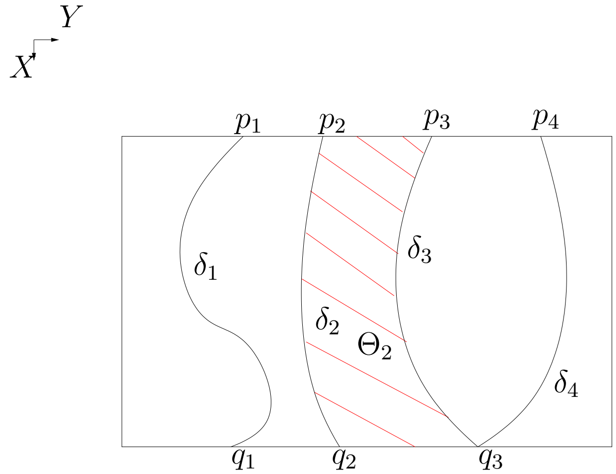

Fig. 6. Projection of the contour curve

$\pi_z^{-1}(\delta_{i+1})$. We denote hereafter by F the patch associated to γ.

There are two cases to consider:

(1) δ_i and δ_{i+1} intersect in $q_{\nu(i)}$.
(2) δ_i and δ_{i+1} do not intersect.

In the first case, we denote by $Q = \theta_i \cap \theta_{i+1}$ the point of $\overline{F}$ which projects onto $q_{\nu(i)}$. By induction hypothesis, as the topology of $\overline{F_i} \cap \pi_z^{-1}(\delta_i)$ is determined by algorithm 5.2, the arc θ_i is represented in $\mathcal{K}$ as the connection of P to Q. The arc θ_{i+1} corresponds to the connection (P', Q), produced by the algorithm, as well as the face (P, P', Q) corresponding to F.

We have

$$\pi_z^{-1}(\delta_{i+1}) \cap S = \left(\mathfrak{C}_z(S_B) \cap \pi_z^{-1}(\delta_{i+1})\right) \cup \left(\overline{\pi_z^{-1}(\Theta_i) \cap S} \cap \pi_z^{-1}(\delta_{i+1})\right)$$

According to the previous paragraph, the arcs of $\overline{\pi_z^{-1}(\Theta_i) \cap S} \cap \pi_z^{-1}(\delta_{i+1})$ are thus obtained by algorithm 5.2. The arcs of $\mathfrak{C}_z(S_B) \cap \pi_z^{-1}(\delta_{i+1})$ are obtained by algorithm 4.1. Thus the algorithm 5.2 compute the topology of $\pi_z^{-1}(\delta_i) \cap S$.

In the second case, we denote again by $Q = \theta_i \cap S_c$ the point of $\overline{F}$ which projects onto $q_{\nu(i)}$ and by $Q' = \theta_{i+1} \cap S_c$ the point of $\overline{F}$ which projects onto $q_{\nu(i)+1}$. The intersection $\overline{F} \cap S_c$ is an arc connecting Q to Q', which exists in $\mathcal{K}$, by hypothesis.

Conversely, as the surface is in generic position (see definition 5.1), an arc of $S_c \cap \pi_z^{-1}([q_{\nu(i)}, q_{\nu(i)+1}]) = \mathcal{B}_{\nu(i)}$ is in the closure of only one patch defined by an arc in $S_r \cap \pi_z^{-1}([p_i, p_{i+1}]) = \mathcal{A}_i$. So there is a one to one correspondence between the arcs in $\mathcal{A}_i$ and the arcs in $\mathcal{B}_{\nu(i)}$. Moreover, this correspondence respects the z-order on the arcs, since there is no point of $\mathfrak{C}_z(S_B)$ above Θ_i.

In particular, the smallest arc $\gamma = (P, P')$ in $\mathcal{A}_i$ is connected to the smallest arc $\eta = (Q, Q')$ in $\mathcal{B}_{\nu(i)}$ by a face (P, P', Q', Q) corresponding to F, as computed by algorithm 5.2.

The arc θ_{i+1} connects the point P' to Q', as computed by the algorithm 5.2, so that the topology of $\pi_z^{-1}(\delta_{i+1}) \cap S$ is determined by the algorithm.

This proves that if the topology of $\pi_z^{-1}(\delta_i) \cap S$, S_r, S_c are determined, then algorithm 5.2 compute the topology of S above $\overline{\Theta_i}$. $\square$

b) We have to ensure that our triangulation is valid. It is clear that the triangulation we compute does not create holes, because the triangulation refines the topological complex $\mathcal{K}$. Let us check now that we do not create intersection of the open segments and open triangles.

As the algorithm proceeds iteratively on the cylinders $\overline{\pi_z^{-1}(\Theta_i)}$, we have only to check this property above $\overline{\Theta_i}$. By construction, the projection by π_z of open segments and open triangles are either disjoint or included one in the other.

If these projections are disjoint, they cannot self-intersect.

Otherwise, since these are linear objects, their intersection would imply an inversion of the z-position of the corresponding arcs (resp. points) in the section S_r and S_c, which is not possible by Thom's isotopy lemma.

This shows that the triangulation of S is valid.

6.4 The isotopy

We are going to detail an explicit isotopy between the original surface and its polygonal approximation.

Definition 6.17 *We say that two surfaces S and S' are isotopic if there exist an application $F : \mathbb{R}^3 \times [0, 1] \longrightarrow \mathbb{R}^3$ such that:*

(1) F is continuous
(2) $F(.,0)$ = Identity
(3) $F(S,1) = S'$
(4) $\forall t \in [0, 1]$, $F(.,t)$ is a homeomorphism onto its image.

We have seen that the projection of the contour curve on the (x, y) plane (parallel to the z-direction), partitions the part of the (x, y) plane between S_r and S_c (see Figure 6).

The region $\overline{\Theta_i}$ is defined by the projection of two arcs of the contour curve. We will call δ_{i-1} and δ_i the two projected arc. They correspond to graphs of semi-algebraic functions of x on $[a,b]$, as the restriction of $\pi_{y,z}$ to $\mathfrak{C}_z(S_B)$ is submersive over $[a,b[$ and the arcs are of dimension 1.

The vertical cylinder with base $\overline{\Theta_i}$, cuts the variety along a family of patches and possibly curves which are not included in the closure of a dimension 2 patch. By construction, to each patch corresponds a sub-part of the triangulation with the particular property that the set of all those parts of the triangulation has also a cylindrical structure. More exactly, the patches of the original surface are projected onto $\overline{\Theta_i}$ and the corresponding triangulations project on the same quadrangle or triangle that will be denoted by Δ_i. This is a consequence of the division of the space with vertical walls that we have made (see Figure 3).

Let us now consider two families $\phi_k, k=1,\dots,n$ and $\psi_k, k=1,\dots,n$ of graphs of continuous functions defined on the interval $[0,1]$. Those graphs verify :

(1) $\forall x \in]0,1[\ \phi_1(x) < \cdots < \phi_n(x)\ ,\ \psi_1 < \cdots < \psi_n(x)$.
(2) For $x \in \{0,1\}$, if $\phi_k(x) = \phi_{k+1}(x)$ then $\psi_k(x) = \psi_{k+1}(x)$.

Then there exists an isotopy from $[0,1]\times\mathbb{R}$ that send each $\phi_k, k=1,\dots,n$ on ψ_k. One can easily verify that the following application is suitable: $(x,y,t) \mapsto (x, g(x,y,t))$ with:

$$
\begin{aligned}
&g(x,y,t) = \mathbf{1}_{]-\infty,\phi_1(x)]}(y + t(\psi_1(x) - \phi_1(x)))\\
&+\mathbf{1}_{]-\Phi_1(x),\phi_2(x)]}((1-t)y + t(\tfrac{y-\phi_1(x)}{\phi_2(x)-\phi_1(x)}\psi_2(x) + \tfrac{\phi_2(x)-y}{\phi_2(x)-\phi_1(x)}\psi_1(x))) + \cdots\\
&+\mathbf{1}_{]-\Phi_{n-1}(x),\phi_n(x)]}((1-t)y + t(\tfrac{y-\phi_{n-1}(x)}{\phi_n(x)-\phi_{n-1}(x)}\psi_n(x) + \tfrac{\phi_n(x)-y}{\phi_n(x)-\phi_{n-1}(x)}\psi_{n-1}(x))))\\
&+\mathbf{1}_{]\phi_n(x),+\infty[}(y + t(\psi_n(x) - \phi_n(x))).
\end{aligned}
$$

For a fixed x, the application sends intervals on intervals.

Let us consider for ϕ_k, the family of arcs defining the Θ_i and for ψ_k, those defining the Δ_i. We deduce from the previous result that applying an isotopy of the form : $(x,y,z,t) \overset{H_1}{\mapsto} (x, g(x,y,t), z)$, we make the projections of the patches and their triangulations on the plane (x,y) coincide. As illustrated in Figure 8, we have transformed the Θ_i into the Δ_i. Moreover, applying this result on each interval between a regular and a critical section, the isotopies glue together into a global one.

More precisely :

(1) In the first step, we have considered a transformation of the form $(x,y,z,t) \overset{H_1}{\mapsto} (x, g(x,y,t), z)$ which does not modify the coordinates x and z. This transformation makes the projections on the plane (x,y) of ϕ_k (patches

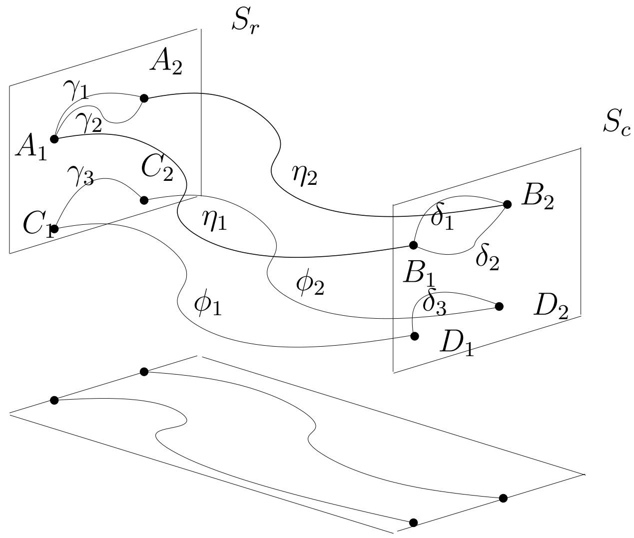

Fig. 7. Family of patches

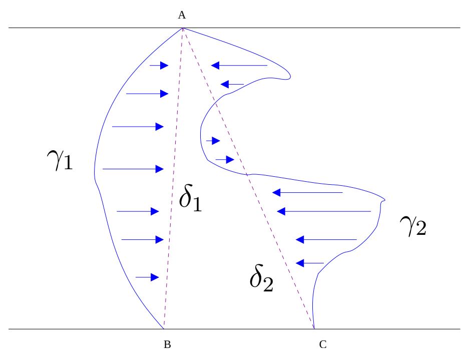

Fig. 8. First step of the isotopy

of surfaces) and of ψ_k (patches of triangulation) coincide. This transformation does not modify the relative order of the graphs.

(2) The second step of the isotopy is a vertical transformation of the form $(x, y, z, t) \overset{H_2}{\mapsto} (x, y, h(x, y, z, t))$ which sends $H_1(\phi_k)$ on ψ_k. It is similar to the previous one, we will not go into any further details here. Above $\overline{\Theta_i}$, there are patches and possibly isolated arcs of the contour curve. If such an isolated arc γ is in between two patches $H_1(\phi_k)$ and $H_1(\phi_{k+1})$, we add a term in the isotopy transformation, corresponding to a virtual patch F with γ on its boundary and which lies between the two patches $H_1(\phi_k)$ and $H_1(\phi_{k+1})$.

7 Example

In this section, we illustrate what the algorithm does on two examples. The first example we chose is known as Whitney's umbrella and the classical normal form for it is $zx^2 - y^2$. We ask the algorithm to compute the topology of this surface in the unit ball.

In the first step the algorithm determines that $z\,x^2 - y^2$ is not in generic position because the line $x = 0$, $y = 0$ is included in the surface. It performs a random change of variable and the surface is now in generic position. The algorithm then computes the projection of the intersection of the unit ball and the umbrella and of the polar variety in the new coordinate system and identifies the x-critical points. Then it performs the connection between the points on the surface. These points have been plotted in green on the pictures below, the blue lines show the polar variety (the vertical z axis is in it, but it is hidden by the red line), and the red lines that connect the green points illustrate the arcs that connect the green points in the output of the algorithm. In other words, two critical points are connected by a red line if there is a direct path that connect them in the output of the algorithm. We have only represented the critical points in the picture to keep it clear. As a matter of fact on this example the algorithm has to compute more points but the underlying connection structure is the one represented here. The structure has two symmetrical "chip"-like parts, and the stick of the umbrella separates them. Notice that the isolated part of the stick is correctly handled by the algorithm, but the lower "chip"-like part partially hides its endpoint, so one has to look carefully to see the whole stick.

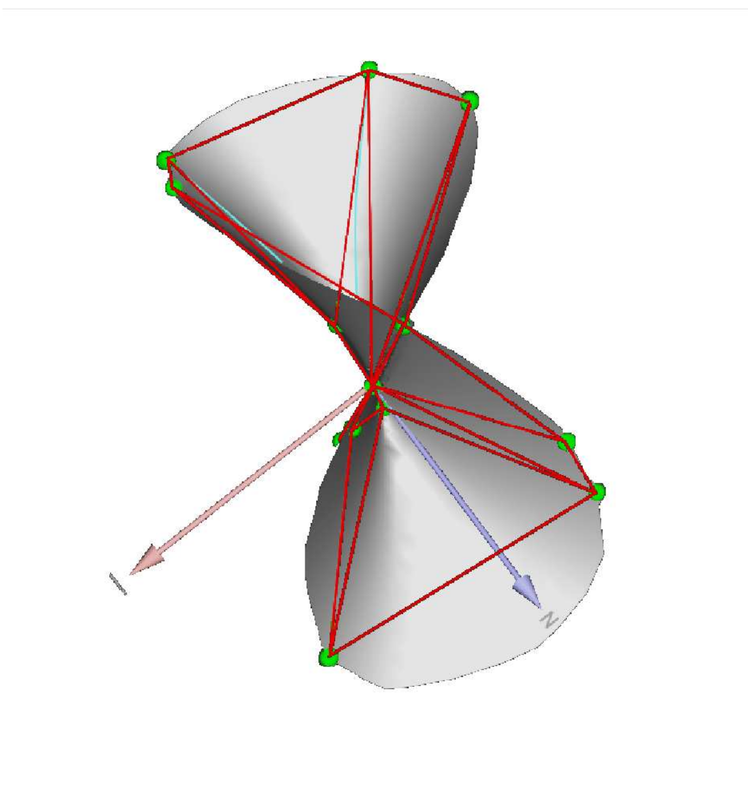

The second example comes from the following equation:

$$4\,z^4x^4 + 8\,z^4x^2y^2 + 4\,z^4y^4 + 8\,z^3x^4 + 16\,z^3x^2y^2 + 8\,z^3y^4 + 19\,z^4x^2 + 19\,z^4y^2 \\ +8\,z^2x^4 + 8\,z^2y^4 + 16\,y^2\,z^2x^2 - 72\,z^3x^2 - 72\,z^3y^2 + 4\,zx^4 + 8\,zx^2y^2 + 4\,zy^4 \\ +121\,z^4 - 28\,z^2x^2 - 28\,z^2y^2 + x^4 + 2\,x^2y^2 + y^4 - 308\,z^3 + 20\,zx^2 + 20\,zy^2 \\ +262\,z^2 - 3\,x^2 - 3\,y^2 - 84\,z + 9 = 0.$$

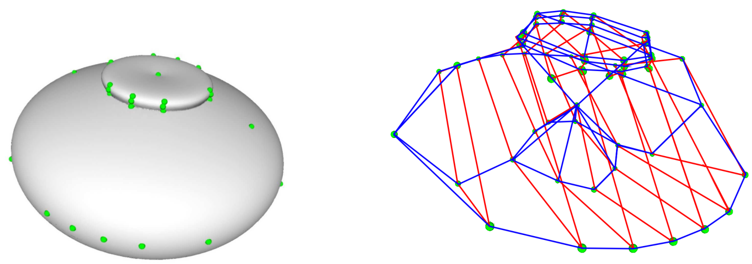

The first picture shows what the surface looks like. The second picture is an illustration of the connections. The green points are the same as in the previous picture: they are the characteristic points that the algorithm uses to recover the topology. They have been computed by the algorithm then displayed. The red lines were added to show the topology of the slices. The blue line represent *some* of the connections between the slices, drawing them all would have made the picture too messy.
From the output of the algorithm one sees that the surface is self-intersecting with a cone, which was not obvious on the first picture.

8 Complexity and effectiveness

The algorithm of Cylindrical Algebraic Decomposition (CAD) computes in the case of a polynomial $f(x, y, z) = 0$, at most $\mathcal{O}(d^{3^2})$ polynomials of degree at most $\mathcal{O}(d^{2^2})$ [6, Chap. 11], which yields at most $\mathcal{O}(d^{13})$ points to compute.

With our algorithm, we have the following result.

Proposition 8.1 *At most $\mathcal{O}(d^7)$ points on an algebraic surface S of degree d are enough to determine a simplicial complex isotopic to it.*

Proof. As described in the previous sections, we are able to deduce the topology of the surface from the solution of system (1) and from the intersection points of the polar curve with planes $V(x - \alpha)$ where the α's lie in between the x-critical values of the planar curves defined by the polynomials $\Delta(x, y) = \mathrm{Res}_z(f(x, y, z), q(x, y, z)\,\partial_z f(x, y, z))$, $\Psi(x, z) = \mathrm{Res}_y(f(x, y, z), q(x, y, z)\,\partial_z f(x, y, z))$.

As $\deg(f) = d$ and $\deg(q\,\partial_z f) = d+1$, the degree of $\Delta(x,y)$ is bounded by $d\,(d+1)$. By Bezout theorem, the number of (real) solutions of the system 1 is bounded by

$$d\,(d+1)\,(d\,(d+1)-1)\,d\,(d+1)\,d = \mathcal{O}(d^7).$$

As there are at most $\mathcal{O}(d^4)$ critical values for $\Delta(x,y)$ and $\Psi(x,z)$ (which are of degree d^2), and as the polar curve is of degree $d\,(d-1)$, there are at most $\mathcal{O}(d^6)$ additional points to insert to get the topology of the polar curve and to deduce an isotopic triangulation of the surface. □

Notice that this bound is bigger than the size of a minimal cell decomposition, since several non-isotopic curves or surfaces yield the same size for the minimal decomposition (eg. just take distinct configurations of ovals in the plane) and does not compare with the bounds on connected components (see eg. [8]) or the complexity of the semi-algebraic set [34] or with output size bounds in [7].

From an effectiveness point of view, we have to compute an approximate or exact representation of the real roots of system (1) and then to compare their coordinates in order to deduce the connections. This can be performed effectively by using a rational univariate representation of the roots and Sturm (Habicht) sequences [21], [6], [15].

In [31], [25] an analysis of the number of isotopy types of a smooth plane algebraic curve of degree d is given. It is shown that this number is exponentially weakly equivalent [2] to e^{d^2} when $d \to \infty$.

Using the sweeping algorithm in 2D [20], we can prove that the number of isotopy classes for general planar curves of degree d is exponentially weakly bounded by e^{d^3}. The proof is similar to the one that we detail now for surfaces:

Proposition 8.2 *The number of isotopy types of an algebraic surface of degree d is exponentially weakly bounded by e^{d^7}.*

Proof. Assume the surface is in generic position (see definition 5.1) and that moreover the projected curve $\mathcal{C}_{x,y}$ has at most one x-critical point for each x and that the number of points on $\mathcal{C}$ above a x-critical point of $\mathcal{C}_{x,y}$ is ≤ 2. These conditions can be satisfied by a generic change of coordinates.

As the degree of $\mathcal{C}_{x,y}$ is $\leq d^2$, it as at most d^4 x-critical points. We consider d^4 x-critical sections which intersect $\mathcal{C}_{x,y}$ in at most d^2 points, above which we have at most d points on S. This yields a total of d^7 points. To each of these points, we associate the value

[2] A function f is said to be exponentially bounded by (resp. weakly equivalent to) g if $\log(f) = \mathcal{O}(\log(g))$ (resp. $\log(f) = \mathcal{O}(\log(g))$ and $\log(f)^{-1} = \mathcal{O}(\log(g)^{-1})$).

- 0 if it is not in the section of S,
- r if it is a regular point of the section of S,
- c if it is on the contour curve $\mathcal{C}$ and projects onto a regular point of $\mathcal{C}_{x,y}$,
- x if it is on the contour curve and projects onto a x-critical point of $\mathcal{C}_{x,y}$.

By the genericity assumption, there is at most two points with index x on a x-critical section. Similarly, we insert regular sections in-between these x-critical sections and regular vertical lines in between the points of $\mathcal{C}_{x,y}$ at x-critical section. This gives $\mathcal{O}(d^7)$ additional points to which we associate the index 0 if the point is not on S and r otherwise.

From this information, the algorithm determines in a unique way the connections between the points of x-critical section and a consecutive regular section, if there is only one point with index x in the x-critical section. Otherwise, there are $\mathcal{O}(d)$ choices to connect the two points of index x with the other in the x-critical section and $\mathcal{O}(d^2)$ choices to connect them in the next regular sections of $\mathcal{C}$. Once these connections are chosen, they determine a unique topological complex equivalent to the surface. This shows that the number of isotopy classes of algebraic surfaces of degree d is bounded by $d^3\, 4^{\mathcal{O}(d^7)}$, which proves the proposition. □

Acknowledgments: We would like to thank Georges Comte and Michel Merle for very helpful discussions on stratification theory. We thank V. Kharlamov for suggesting this problem on the asymptotic of the number of isotopy classes of curves and surfaces. We also thank the reviewers for their precise and helpful comments.

References

[1] L. Alberti, G. Comte, and B. Mourrain. Meshing implicit algebraic surfaces: the smooth case. In L.L. Schumaker M. Maehlen, K. Morken, editor, *Mathematical Methods for Curves and Surfaces: Tromso'04*, pages 11–26. Nashboro, 2005.

[2] L. Alberti and B. Mourrain. Visualisation of algebraic curves. In *The 15th Pacific Graphics*, pages 303–312. IEEE Computer Society, 2007.

[3] J. G. Alcázar, J. R. Sendra and J. Schicho. A delineability-based method for computing critical sets of algebraic surfaces. *J.Symb.Comp.*, 42(6):678–691, 2007.

[4] J. G. Alcázar and J. R. Sendra. Computing the topology of real algebraic space curves. *J. Symbolic Comput.*, 39:719–744, 2005.

[5] V. I. Arnol'd, S. M. Guseĭn-Zade, and A. N. Varchenko. *Singularities of differentiable maps. Vol. I*, volume 82 of *Monographs in Mathematics*. Birkhäuser Boston Inc., Boston, MA, 1985. The classification of critical points,

caustics and wave fronts, Translated from the Russian by Ian Porteous and Mark Reynolds.

[6] S. Basu, R. Pollack, and M.-F. Roy. *Algorithms in Real Algebraic Geometry*. Springer-Verlag, Berlin, 2003. ISBN 3-540-00973-6.

[7] E. Berberich, M. Kerber, and M. Sagraloff. Exact geometric-topological analysis of algebraic surfaces. In *SoCG*. ACM Press, 2008. to appear.

[8] J. Bochnak, M. Coste, and M.-F. Roy. *Géométrie Algébrique Réelle*. Springer-Verlag, Heidelberg, 1987.

[9] J.D. Boissonnat, D. Cohen-Steiner, and G. Vegter. Isotopic implicit surface meshing. In *Proceedings of STOC 2004*, pages 301–309, 2004.

[10] L. Busé and B. Mourrain. Explicit factors of some iterated resultants and discriminants. 2008.

[11] J. Canny. *The Complexity of Robot Motion Planning*. M.I.T. Press, Cambridge, Mass., 1988.

[12] G. E. Collins. Quantifier elimination for real closed fields by cylindrical algebraic decomposition. In *Proc. 2nd GI Conference on Automata Theory and Formal Languages*, volume 33 of *Lecture Notes Comput. Sci.*, pages 134–183. Springer-Verlag, 1975.

[13] M. Coste. An introduction to semi-algebraic geometry. RAAG network school, 2002.

[14] A. Dimca. *Singularities and topology of hypersurfaces*. Universitext. Springer-Verlag, New York, 1992.

[15] M. Elkadi and B. Mourrain. *Introduction à la résolution des systèmes d'équations algébriques*, volume 59 of *Mathématiques et Applications*. Springer-Verlag, 2007.

[16] M. El Kahoui. Topology of a real algebraic space curve. *J.Symb.Comp.*, 43:235–258, 2008.

[17] E. Fortuna, P. M. Gianni, and D. Luminati. Algorithmical determination of the topology of a real algebraic surface. *J. Symb. Comput.*, 38(6):1551–1567, 2004.

[18] E. Fortuna, P. M. Gianni, P. Parenti, and C. Traverso. Algorithms to compute the topology of orientable real algebraic surfaces. *J. Symb. Comput.*, 36(3-4):343–364, 2003.

[19] G. Gatellier, A. Labrouzy, B. Mourrain, and J.-P. Técourt. Computing the topology of 3-dimensional algebraic curves. In *Computational Methods for Algebraic Spline Surfaces*, pages 27–44. Springer-Verlag, 2004.

[20] L. González-Vega and I. Necula. Efficient topology determination of implicitly defined algebraic plane curves. *Comput. Aided Geom. Design*, 19(9):719–743, 2002.

[21] L. Gonzalez-Vega, F. Rouillier, and M.F. Roy. *Symbolic Recipes for Polynomial System Solving.* Some Tapas of Computer Algebra. Springer, 1997.

[22] M. Goresky and R. MacPherson. *Stratified Morse Theory.* Springer-Verlag, 1988.

[23] Robert M. Hardt. Triangulation of subanalytic sets and proper light subanalytic maps. *Invent. Math.*, 38(3):207–217, 1976/77.

[24] H. Hironaka. Triangulations of algebraic sets. In *Algebraic geometry (Proc. Sympos. Pure Math., Vol. 29, Humboldt State Univ., Arcata, Calif., 1974)*, pages 165–185. Amer. Math. Soc., Providence, R.I., 1975.

[25] V. M. Kharlamov and S. Yu. Orevkov. The number of trees half of whose vertices are leaves and asymptotic enumeration of plane real algebraic curves. *J. Combin. Theory Ser. A*, 105(1):127–142, 2004.

[26] G. N. Khimšiašvili. The local degree of a smooth mapping. *Sakharth. SSR Mecn. Akad. Moambe*, 85(2):309–312, 1977.

[27] N. G. Lloyd. *Degree Theory.* Cambridge University Press, Cambridge, 1978.

[28] J. Mather. Lecture notes on topological stability. Havard booklet, 1971.

[29] S. McCallum and G.E. Collins. Local box adjacency algorithms for cylindrical algebraic decompositions. *J. Symbolic Comput.*, 33(3):321–342, 2002.

[30] D. N'Diatta, B. Mourrain, and O. Ruatta. On the computation of the topology of a non-reduced implicit space curve. In D. Jeffrey, editor, *ISSAC*. ACM Press, 2008. to appear.

[31] S. Yu. Orevkov and V. M. Kharlamov. Growth order of the number of classes of real plane algebraic curves as the degree grows. *Zap. Nauchn. Sem. S.-Peterburg. Otdel. Mat. Inst. Steklov. (POMI)*, 266(Teor. Predst. Din. Sist. Komb. i Algoritm. Metody. 5):218–233, 339, 2000.

[32] J. P. Speder. Équisingularité et conditions de Whitney. *Amer. J. Math.*, 97(3):571–588, 1975.

[33] D. Trotman. On Canny's roadmap algorithm: orienteering on semialgebraic sets (an application of singularity theory to theoretical robotics). In *Proceedings of the 1989 Warwick Singularity Theory Symposium*, Lecture Notes in Math., pages 320–339. Springer, 1991.

[34] Y. Yomdin and G. Comte. *Tame geometry with applications in smooth analysis.* LNM 1834. Springer-Verlag, 2004.

[35] O. Zariski. Studies in equisingularity. II. Equisingularity in codimension 1 (and characteristic zero). *Amer. J. Math.*, 87:972–1006, 1965.