On the topology of real algebraic plane curves

Jinsan Cheng* Sylvain Lazard* Luis Peñaranda*
Marc Pouget* Fabrice Rouillier† Elias Tsigaridas‡

April 5, 2010

Contents

1 Introduction **2**

2 Preliminaries **7**

3 Algorithm **10**
3.1 Curve without vertical lines 10
3.2 Curve with vertical lines 13

4 Complexity Analysis **13**
4.1 Preliminaries 14
4.2 Computation of the isolating boxes 15
4.3 Refinement with respect to the isolating curve 16
4.4 Refinement of the singular points to avoid top/bottom crossings 17
4.5 Overall complexity 18
4.6 Technical details 18
4.6.1 Interval arithmetic 18
4.6.2 Separation bounds 19

5 Experiments **20**

6 Conclusion **21**

Abstract

We revisit the problem of computing the topology and geometry of a real algebraic plane curve. The topology is of prime interest but geometric information, such as the position of singular and critical points, is also relevant. A challenge is to compute efficiently this information for the given coordinate system even if the curve is not in generic position.

*LORIA (INRIA, CNRS, Nancy Université) and INRIA Nancy - Grand Est, Nancy, France. Firstname.Name@loria.fr

†LIP6 (Université Paris 6, CNRS) and INRIA Paris-Rocquencourt, Paris, France. Firstname.Name@lip6.fr

‡INRIA Sophia-Antipolis Méditerranée. Firstname.Name@inria.fr

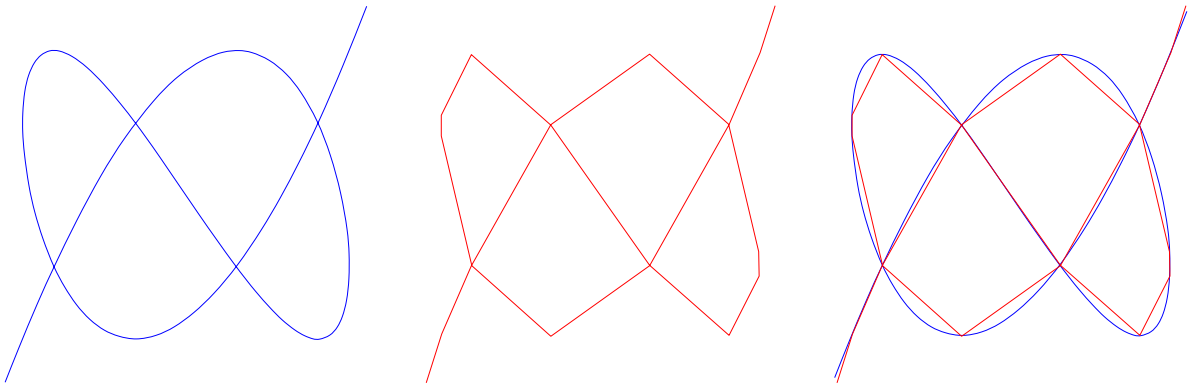

(a) (b) (c)

Figure 1: (a) Curve of equation $16\,x^5 - 20\,x^3 + 5\,x - 4\,y^3 + 3\,y = 0$ plotted in MAPLE and (b) its isotopic graph computed by ISOTOP.

Previous methods based on the cylindrical algebraic decomposition use sub-resultant sequences and computations with polynomials with algebraic coefficients. A novelty of our approach is to replace these tools by Gröbner basis computations and isolation with rational univariate representations. This has the advantage of avoiding computations with polynomials with algebraic coefficients, even in non-generic positions. Our algorithm isolates critical points in boxes and computes a decomposition of the plane by rectangular boxes. This decomposition also induces a new approach for computing an arrangement of polylines isotopic to the input curve. We also present an analysis of the complexity of our algorithm. An implementation of our algorithm demonstrates its efficiency, in particular on high-degree non-generic curves.

1 Introduction

⟶ MARC SAYS: Let $\mathcal{C}$ be a real algebraic plane curve defined in a Cartesian coordinate system by a bivariate polynomial f with rational coefficients, *i.e.*, $\mathcal{C} = \{(x, y) \in \mathbb{R}^2 \mid f(x, y) = 0\}$ with $f \in \mathbb{Q}[x, y]$. We consider on $\mathcal{C}$ the topology induced by the euclidean plane. The article addresses the encoding of the topology of $\mathcal{C}$ with additional geometric information associated to the given coordinate system. We first give the formal definition of the topological encoding we chose before discussing the relevance of the geometric issues.

As $\mathcal{C}$ is embedded in the plane, the more precise way to encode the topology of $\mathcal{C}$ is not only by describing a homeomorphism with a simpler object, but also to require that this homeomorphism extends to the whole plane. Clearly, it seems not natural to consider two plane curves topologically equivalent if one is two nested circles and the other is also two circles but not nested (note that these two curves however are homeomorphic). Hence, we aim at constructing an "extended" plane straight line graph $\mathcal{G}$ that encodes the embedded topology of $\mathcal{C}$. More precisely, we construct $\mathcal{G}$ (its topology being induced by the euclidean plane) such that it is ambient homeomorphic to $\mathcal{C}$, in other words, there exists an homeomorphism of the plane that maps $\mathcal{C}$ onto $\mathcal{G}$. We need an "extended" notion of plane graph because the usual definition of a plane graph implies compactness [**?**] whereas an algebraic curve may not be compact. We solve this technical detail by allowing an "extended" plane graph to have (a) vertices with one of its coordinates equal to $+\infty$ or $-\infty$.

The graph $\mathcal{G}$, without these infinite vertices, is a subset of the euclidean plane with a well defined induced topology, and this is formally this object that is considered for topological consideration[1].

We now discuss the relevance of adding geometric (or algebro-geometric) information to the graph $\mathcal{G}$. From the topological point of view, the graph $\mathcal{G}$ must contain vertices that correspond to self-intersections and isolated points of the curve. However, in order to avoid separating such relevant points from other singularities (*e.g.*, cusps), all *singular points* of $\mathcal{C}$, that is, points at which the tangent is not well defined, are chosen to be vertices of the graph. ⟵

⟶ MARC SAYS: NOTE: technical
Again another approach would have been to define a bounding box outside of which the curve is homeomorphic to a set of non crossing rays, exactly one for each point of $\mathcal{C}$ of the boundary of this box and then just consider topological issues inside the box. In fact this is exactly what we do in practice, but it doesnt seem easier nor more elegant to formulate it that way.... any comments?

ANOTHER MORE TECHNICAL AND THEORETICAL NOTE:
The definition I gave above of ambient homeomorphism (assuming it is orientation preserving) implies that we even have that $\mathcal{C}$ and $\mathcal{G}$ are ambient isotopic, i.e. there exists a continuous map $F : \mathbb{R}^2 \times [0,1] \longrightarrow \mathbb{R}^2$, such that F_t is an homeomorphism for any $t \in [0,1]$, F_0 is the identity of $\mathbb{R}^2$ and $F_1(\mathcal{C}) = \mathcal{G}$. There are two ways to see this: (a) actually construct F which is easy (i will sketch it), or (b) just use the Alexandrov theorem that asserts the any orientation preserving homeo of the plane is isotopic to the identity. Anyway i think these details are not relevant here, the important point is that the reader understands the relevance of the ambient nature of the topological issue.

NOTE ON TOPOLOGY IN THE PROJECTIVE PLANE
I checked the details, and it is easy to also compute the embedded topology in $\mathbb{P}^2$, we may mention this in conclusion and it may be interesting for some applications. For instance, Laurent Dupont asked us the number of projective connected components. ⟵

While singular points are needed for computing the topology of a curve, the extreme points of a curve are also very important for representing its geometry. Precisely, the *extreme points* of $\mathcal{C}$ for a particular direction, say the direction of the x-axis, are the non-singular points of $\mathcal{C}$ at which the tangent line is vertical (*i.e.*, parallel to the y-axis); the extreme points in the direction of the x-axis are called *x-extreme*. These extreme points are crucial for various applications and, in particular, for computing arrangements of curves by a standard sweep-line approach [EK08]. Of course, one can theoretically compute an arrangement of algebraic curves by computing the topology of their product. However, this approach is obviously highly inefficient compared to computing the topology of each input curve and, only then, computing the arrangement of all the curves with a sweep-line algorithm. Note that the x-extreme and singular points of $\mathcal{C}$ form together the *x-critical points* of the curve (the x-coordinates of these points are exactly the positions of a vertical sweep line at which there may be a change in the number of intersection points with $\mathcal{C}$).

It is thus useful to require that all the x-critical points of $\mathcal{C}$ are vertices of the graph we want to compute. To our knowledge, almost all methods for computing the topology of a curve compute the *critical points* of the curve and associate corresponding vertices in the graph. (Refer to [AMW08, BSGY08] for recent subdivision methods that avoid the computation of non-singular critical points.) However, it should be stressed that almost all methods do not necessarily compute the critical points for the specified x-direction. Indeed, when the curve is not in *generic position*, that is, if two x-critical points have the same x-coordinate or if the curve admits a vertical asymptote, most

[1]Note that an alternative would have been to compute the topology in the real projective plane. On one hand this solves the compactness issue but on the other hand this changes the topology and hence the problem itself.

algorithms shear the curve so that the resulting curve is in generic position. This is, however, an issue for several reasons. First, determining whether a curve is in generic position is not a trivial task and it is time consuming [GVEK96, SF90]. Second, if one wants to compute arrangements of algebraic curves with a sweep-line approach, the extreme points of all the curves have to be computed for the same direction. Finally, if the coordinate system is sheared, the polynomial of the initial curve is transformed into a dense polynomial, which slows down, in practice, the computation of the critical points.

In this paper, given a curve $\mathcal{C}$ which is not necessarily in generic position, we aim at computing efficiently its topology, including all the critical points for the specified x-direction. In terms of efficiency, our primary goal is the practical efficiency (rather than worst-case complexity). In particular, we want to avoid computations with non-rational algebraic numbers or, equivalently in this context, algebraic computations such as Sturm sequences, Sturm-Habicht sequences (which are a generalization of Sturm sequences, with better specialization properties [GVLRR89]), and sub-resultant sequences.

After an brief overview of our algorithm, we review below previous work on the problem and then present our contributions.

Overview of the algorithm. $\longrightarrow$ MARC SAYS: To get a clearer idea of what is actually computed, we sketch the main steps of the algorithm. For simplicity, we restrict here to a compact curve so that $\mathcal{G}$ does not need to be "extended", the complete algorithm is detailed in Section 3. The general idea is to decompose the plane by rectangles such that in each rectangle the topology of $\mathcal{C}$ can be easily recovered by the knowledge of its intersections with the boundary of the rectangle.

- The first step isolates critical points in two dimensional rectangles, called critical boxes in the following. Each critical box is such that the restriction of $\mathcal{C}$ to the box is a set of x-monotone (non-crossing) arcs connecting the critical point to a point on the boundary of the box.

- In the next step, a rectangular decomposition is constructed by extending the vertical sides of the critical boxes either to infinity or to the first encountered critical box (see Figure 2(a) or give a simplified picture with a compact curve?). Let $\mathcal{R}$ be the edges and vertices of this decomposition. In each connected component of $\mathbb{R}^2 \setminus \mathcal{R}$, the topology of $\mathcal{C}$ is simple: this domain is either a critical box or a rectangle with x-monotone (non-crossing) arcs of $\mathcal{C}$ connecting two points on its boundary.

- Finally, it only remains to define the graph $\mathcal{G}$ from the rectangular decomposition. The vertices of $\mathcal{G}$ are critical points and points of $\mathcal{R} \cap \mathcal{C}$, and edges are straight line segments joining two vertices of the same critical box or rectangle if and only if they are connected by an arc of $\mathcal{C}$. We will show that $\mathcal{G}$ is ambient homeomorphic to $\mathcal{C}$ by explicitly constructing such a homeomorphism.

- In addition, geometrical information of $\mathcal{C}$ relative to the coordinate system is encoded as attributes to vertices. For instance a vertex have as attribute a refinable representation of its coordinates, a type such as singular or extreme, or the multiplicity of the point on the curve.

It is worth mentioning that for visualization purposes, a more geometrically accurate picture of $\mathcal{C}$ can be easily obtained by refining the graph $\mathcal{G}$: it reduces to subdivide vertically each rectangle, compute intersections with $\mathcal{C}$ and add corresponding vertices and edges in $\mathcal{G}$. $\longleftarrow$

Previous Work. There have been many papers addressing the problem of computing the topology of algebraic plane curves (or closely related problems) defined by a bivariate polynomial with rational coefficients [AMW08, AM88, CR88, BPR06, EKW07, Fen92, GVEK96, GVN02, Hon96, MPS$^+$06, Sak91, Bro01, ACM84, MC02, Str06, DET07]. Most of the algorithms assume generic position for the input curve. As mentioned above, this is without loss of generality since we can always shear a curve into generic position [SF90, GVEK96] but this has a substantial negative impact on the time computation. All these algorithms perform the following phases. (1) Project the x-critical points of the curve on the x-axis, using sub-resultants sequences and isolate the real roots of the resulting univariate polynomial in x. This gives the x-coordinates of all the x-critical points. (2) For each such value x_i, compute the intersection points between the curve $\mathcal{C}$ and the vertical line $x = x_i$. (3) Through each of these points, determine the number of branches of $\mathcal{C}$ coming from the left and going to the right. (4) Connect all these points appropriately.

The main difficulty in all these algorithms is to compute efficiently all the critical points in Phase 2 because the x-critical values in Phase 1 are, a priori, non-rational thus computing the corresponding y-coordinates in Phase 2 amounts, in general, to solving a univariate polynomial with non-rational coefficients and at least a multiple root (corresponding to the critical point). To this end, most algorithms [AM88, CR88, BPR06, Fen92, GVEK96, GVN02, Hon96, MPS$^+$06, Sak91, ACM84, Str06] rely heavily on computations with real algebraic numbers, Sturm sequences or sub-resultant sequences.

An approach using a variant of sub-resultant sequences for computing the critical points in Phase 2 was proposed by Hong [Hon96]. He computed this way (xy-parallel) boxes with rational endpoints and separating the critical points. Counting the branches in Phase 3 can then be done by intersecting the boundary of the boxes with the curve which only involves univariate polynomials with rational coefficients. This approach, see also [ACM84, MC02, Bro01] and the software package Cad2D[2], does not assume that the curve is in generic position.

In a more recent paper, González-Vega and Necula [GVN02] use Sturm-Habicht sequences which allow them to express the y-coordinate of the each critical point as a rational function of its x-coordinate. They implemented their algorithm in MAPLE with symbolic methods modulo the fact that the algebraic coefficients of the polynomials in Phase 2 are approximated in fixed-precision arithmetic. The algorithm takes as a parameter the initial precision for the floating-point arithmetic and recursively increases the precision. This approach is however not certified in the case where the curve is not in generic position because the algorithm checks for the equality of pairs of polynomials whose coefficients are evaluated (incorrect results have been reported in [SW05]). Note that there exists one variant of González-Vega and Necula algorithm that handles, without shearing, curves that are not in generic position [MPS$^+$06]. This approach, however, requires substantial additional time-consuming symbolic computations such as computing Sturm sequences.

Recently, Seidel and Wolpert [SW05] presented an alternate approach for computing the critical points avoiding most costly algebraic computations but to the expense of computing several projections of the critical points. They project, in Phase 1, the x-critical points on both x and y-axes and also on a third random axis. After isolating the roots on each axis, they can recover (xy-parallel) boxes with rational endpoints that contain each exactly one critical point. From there, all computations only involve rational numbers but they, however, still need to compute Sturm-Habicht sequences for refining the boxes containing the singular points until each box interests only the branches of the curve incident to the singular point. Their approach assumes that the curve

[2] http://www.cs.usna.edu/~qepcad/B/user/cad2d.html

is in generic position by a pre-processing phase in which the curve is sheared if needed. They also present a MAPLE implementation, INSULATE, which is an implementation of a certified algorithm for curves in arbitrary positions. Note that their implementation does not report x-extreme points in the original system when the curved is sheared.

Even more recently, Eigenwillig et al. [EKW07] (see also [Ker06]) presented a variant of González-Vega and Necula approach, in which the roots of the polynomials with non-rational coefficients are efficiently isolated using an implementation of a variant of an interval-based Descartes algorithm [EKK$^+$05]. This variant, as [GVN02], does not assume that the curve is in generic position but detects such configurations online. More precisely, if the bit-stream Descartes algorithm is in, a sense, unlucky then, rather than refining down to a separation bound (*e.g.* [BFM$^+$01, LPY04]), the algorithm shears the input curve and starts again. Note that this approach still computes Sturm-Habicht sequences for determining the polynomials appearing in Phase 2 and the multiplicity of its multiple root. Also, if the curve is sheared to a (x', y') coordinate system, they compute extreme points both for the x'-direction and the direction corresponding to the x-axis. This approach has been implemented in C++ and is an implementation of a certified algorithm that handles curves that are not necessarily in generic positions and that reports x-extreme points for the original coordinate system.

Note finally that another approach that avoids expensive algebraic computations is to compute the critical or singular points using subdivision methods [AMW08, BSGY08]. The major drawback of these methods is that, in order to certify the results, the subdivision has, in general, to reach a separation bound (certification can also be achieved by solving algebraic systems by other means). It follows that, if no certification is required, these methods are very fast in practice, however, they can become very slow on difficult instances, if certification is required. To our knowledge, no implementation of such certified algorithm is available.

As for the complexity of the problem, the best known bound so far is $\widetilde{\mathcal{O}}(N^{12})$ [DET07], where N is essentially the maximum of the degree of f and of the maximum number of bits needed for representing the input coefficients and the notation $\widetilde{\mathcal{O}}$ denotes that the poly-logarithmic factors are omitted. This assumes that the real algebraic numbers are represented by isolating intervals. This complexity is based on theoretically fast computations with Sturm-Habicht sequences and algorithms for univariate real root isolation.

Our contributions.

We present an algorithm for computing the topology of an algebraic plane curve which is possibly in non-generic position. The algorithm handles curves in non-generic positions in the Cartesian coordinate system in which they are defined. In particular, the algorithm never shears the coordinate system, which avoids the associated extra costs discussed above.

The other originality of our approach is that we succeed to avoid, in all cases, the computation of sub-resultant sequences and computations with algebraic numbers; we use instead Gröbner bases and the rational univariate representation. One can argue that these are just two different tool sets, this is true from a theoretical point of view but we show in the experiments the benefit of our choice when computing with non-generic curves. Furthermore, the philosophy of our approach is to avoid, as much as possible, computations that are time consuming in practice. This leads to various algorithmic choices such as avoiding the computation of y-critical points and allowing the curve to intersect the top and bottom sides of boxes isolating x-extreme points (which avoids substantial subdivisions since the tangent at an extreme point is vertical).

The novelty of our algorithm mainly relies upon the use of three new ingredients for this problem.

First, we use some of the state-of-the-art techniques to isolate the roots of bivariate systems, *i.e.,* we use (i) Gröbner basis [Fau02], (ii) Rational Univariate Representations (RUR) [Rou99] which represent the roots of the system as rational functions of the roots of a univariate polynomial, and (iii) a subdivision technique based on Descartes' rule of signs (and filtered with interval arithmetic for efficiency) for isolating the roots of the univariate polynomials [RZ03, ESY06, EMT08]. Even though this approach is well known for system solving, it was not used before for computing the topology of algebraic curves.

Second, we compute and use the multiplicities in fibers (see Definition 1) of critical points to compute the topology at singular points and to do the connection at extreme points. For extreme points, we get these multiplicities by the RUR and a special case of a formula of Teissier [Tei73]. For singular points, we solve the system of singular points with additional constraints. Note that the overall method to compute the topology at critical points is not new, it is described in full details in [SW05], see also for example [ACM84, Hon96, Bro01, MC02, MPS$^+$06, AMW08] for closely related approaches for curves in non-generic position. The novelty appears in the way we compute multiplicities in this context; once again we avoid computing sub-resultant sequences.

Third, we present a variant of the standard combinatorial part of the algorithm for computing the topology. We compute a decomposition of the plane (which is not a cylindrical algebraic decomposition) by rectangular boxes containing at most one critical point. Since we allow crossings on the top and bottom of extreme point boxes to avoid costly refinement of boxes, the connection is not always straightforward. To achieve the connection near such points, we use the multiplicities in fibers of extreme points. Note that one advantage of this variant is that, even when the curve is in generic position, the algorithm does not require to refine these boxes until they do not overlap in x.

We also analyze the worst-case bit complexity of our algorithm. To the best of our knowledge, this is, for the problem considered here, the first time that the complexity of a (certified) algorithm based on refinements and approximations is analyzed. Our technique, even though not novel, could presumably also be used to analyze the complexity of the algorithms in [GVN02, Hon96, SW05, EKW07].

Finally, we ran large scale benchmarks on over a thousand of curves during several weeks for comparing the different state-of-the-art available implementations.

The rest of the paper is organized as follows. First, we recall in Section 2 some basic material, in particular, on the multiplicity of the intersection of two curves and on Rational Univariate Representations of the roots of system of (algebraic) equations. We present our algorithm in Section 3 and its complexity in Section 4. Finally, we present in Section 5 extensive experiments comparing various implementations.

2 Preliminaries

Let $\mathcal{C}$, also denoted $\mathcal{C}_f$, be a real algebraic plane curve defined by a bivariate polynomial f in $\mathbb{Q}[x, y]$. Since the geometry of the curve is not modified by taking the square-free part of f, we can assume without loss of generality that f is square-free. Note that $\mathcal{C}$ may consist of several algebraic components, that is, f is not necessarily irreducible in $\mathbb{R}[x, y]$. The algebraic components of the curve that are vertical lines (*i.e.*, lines parallel to the y-axis) can be easily computed since their abscissa correspond to the real roots of the polynomial in x obtained as the gcd of the coefficients of f seen as a polynomial in y.

Partial derivatives are denoted with subscripts: for instance, f_x denotes the derivative of f with respect to x and f_{y^k} (sometimes also f_k) denotes the k^{th} derivative with respect to y. A point $\mathbf{p} = (\alpha, \beta) \in \mathbb{C}^2$ is called x-critical if $f(\mathbf{p}) = f_y(\mathbf{p}) = 0$, singular if $f(\mathbf{p}) = f_x(\mathbf{p}) = f_y(\mathbf{p}) = 0$, and x-extreme if $f(\mathbf{p}) = f_y(\mathbf{p}) = 0$ and $f_x(\mathbf{p}) \neq 0$ (*i.e.*, it is x-critical and non-singular). Similarly are defined y-critical and y-extreme points. As x-critical and x-extreme points are more often used in the following, we often simply refer to them as critical and extreme points.

The ideal generated by polynomials $P_1, \ldots, P_i$ is denoted $\mathbb{I}(P_1, \ldots P_i)$. In the following, we often identify the ideal and the system of equations $\{P_1 = 0, \ldots, P_i = 0\}$ (or any equivalent system induced by a set of generators of the ideal). We consider, in particular, the ideals $I_c = \mathbb{I}(f, f_y)$ and $I_s = \mathbb{I}(f, f_x, f_y)$; their roots are, respectively, the x-critical and singular points of $\mathcal{C}$.

Multiplicities. We now recall the notion of multiplicity of the roots of an ideal, then we state two lemmata using this notion for studying the local topology at critical points. Geometrically, the notion of multiplicity of intersection of two regular curves is intuitive. If the intersection is transverse, the multiplicity is one; otherwise, it is greater than one and it measures the level of degeneracy of the tangential contact between the curves. Defining the multiplicity of the intersection of two curves at a point that is singular for one of them (or possibly both) is more involved and an abstract and general concept of multiplicity in an ideal is needed.

Definition 1 ([CLO05, §4.2])**.** *Let I be an ideal of $\mathbb{Q}[x, y]$ and denote $\overline{\mathbb{Q}}$ the algebraic closure of $\mathbb{Q}$. To each zero (α, β) of I corresponds a local ring $(\overline{\mathbb{Q}}[x, y]/I)_{(\alpha,\beta)}$ obtained by localizing the ring $\overline{\mathbb{Q}}[x, y]/I$ at the maximal ideal $\mathbb{I}(x-\alpha, y-\beta)$. When this local ring is finite dimensional as $\overline{\mathbb{Q}}$-vector space, we say that (α, β) is an isolated zero of I and this dimension is called the* **multiplicity of** (α, β) **as a zero of I**.

Let $f, g \in \mathbb{Q}[x, y]$ be such that the intersection of $\mathcal{C}_f$ and $\mathcal{C}_g$ in $\mathbb{C}^2$ contains a zero-dimensional component equal to point $\mathbf{p} = (\alpha, \beta)$. Then (α, β) is an isolated zero of $\mathbb{I}(f, g)$ and its multiplicity, denoted by $Int(f, g, \mathbf{p})$, is called the **intersection multiplicity of the two curves at this point**.

We call a **fiber** *a vertical line of equation $x = \alpha$. For a point $\mathbf{p} = (\alpha, \beta)$ on the curve $\mathcal{C}_f$, we call the multiplicity of β in the univariate polynomial $f(\alpha, y)$ the* **multiplicity of p in its fiber** *and denote it as $mult(f(\alpha, y), \beta)$.*

The next lemma, due to Teissier [Tei73], relates the multiplicity of a point in a fiber with the multiplicity in the critical ideal. We will use it to deduce the multiplicity in the fiber knowing multiplicity in the ideal. More precisely, we will use the multiplicity in fibers of extreme points during the connection step of our algorithm.

Lemma 2 ([Tei73][BR90, Lemma D.3.4 p.314])**.** *For an x-extreme point $\mathbf{p} = (\alpha, \beta)$ of f one has*

$$mult(f(\alpha, y), \beta) = Int(f, f_y, \mathbf{p}) + 1. \tag{1}$$

To compute the local topology of the curve at a singular point, we aim at isolating the singular point in a box so that the intersection of its border and the curve determines the topology. Indeed for a small enough box, the topology is given by the connection of the singular point with all the intersections on the border. So the box shall avoid parts of the curve not connected to the singular point. Knowing the multiplicity of the singular point in the fiber enables to isolate the singular point from other crossings of the curve in this fiber. Requiring in addition that intersections with the curve only occur on the left or right sides of the box leads to the following.[3]

[3]The proof is based on a recursive application of the mean value theorem stating that the roots of the derivative of a polynomial P lie between those of P.

Lemma 3 ([SW05]). *Let $\mathbf{p} = (\alpha, \beta)$ be a real singular point of the curve $\mathcal{C}_f$ of multiplicity k in its fiber. Let B be a box satisfying (i) B contains $\mathbf{p}$ and no other x-critical point, (ii) the function f_{y^k} does not vanish on B, and (iii) the curve $\mathcal{C}_f$ crosses the border of B only on the left or the right sides. Then the topology of the curve in B is given by connecting the singular point with all the intersections on the border.*

We finally describe the algebraic tools we use for isolating the roots of univariate and bivariate ideals.

Univariate Root Isolation We need to count and isolate the real roots of univariate polynomials, possibly in a given interval. This is, in particular, needed for computing the intersections between $\mathcal{C}$ and the sides of the boxes isolating the critical points. Only polynomials with rational coefficients will be considered. The square-free part of the considered polynomials is first computed. The real roots are then isolated using recursive subdivision and the Descartes' rule of signs (see [ESY06, EMT08, RZ03] for details and [RZ03] for the way interval arithmetic can be used to speed up computations). In our implementation, we use the RS software [RS].

Rational Univariate Representation [Rou99]. In our algorithm, we need to represent solutions of zero-dimensional ideals depending on two variables by boxes containing them. We use the so-called Rational Univariate Representation (RUR) of the roots, which can be viewed as a univariate equivalent to the studied ideal. The key feature of this RUR is the ability to isolate solutions in easily refinable boxes and to compute multiplicities.

Given a zero-dimensional ideal $I = \mathbb{I}(P_1, \ldots, P_s)$ where the $P_i \in \mathbb{Q}[x_1, \ldots, x_n]$, a Rational Univariate Representation of the solutions $\mathrm{V}(I)$ is given by $F(t) = 0, x_1 = \frac{G_1(t)}{G_0(t)}, \ldots, x_n = \frac{G_n(t)}{G_0(t)}$, where $F, G_0, \ldots, G_n$ are univariate polynomials in $\mathbb{Q}[t]$ (where t is a new variable). All these univariate polynomials, and thus the RUR, are uniquely defined with respect to a given polynomial $\gamma \in \mathbb{Q}[x_1, \ldots, x_n]$ which is injective on $V(I)$; γ is called the *separating polynomial* of the RUR. [4] Note that a random degree-one polynomial in $x_1, \ldots, x_n$ is a separating polynomial with probability one. The RUR defines a bijection between the (complex and real) roots of the ideal I and those of F. Furthermore, this bijection preserves the multiplicities and the real roots. Computing a box for a solution of the system is done by isolating the corresponding root of the univariate polynomial F and evaluating the coordinate functions with interval arithmetic (see Section 4 for details on interval analysis). To refine a box, one just needs to refine the corresponding root of F and evaluate the coordinates again.

There exists several ways for computing a RUR. One can use the strategy from [Rou99] which consists of computing a Gröbner basis of I and then to perform linear algebra operations to compute a separating element as well as the full expression of the RUR. The Gröbner basis computation can also be replaced by the generalized normal form from [MT05]. There exists more or less certified alternatives such as the Geometrical resolution [GLS01] (it is probabilistic since the separating element is randomly chosen and its validity is not checked, one also loses the multiplicities of the roots) or resultant based strategies such as [KOR05]. In our implementation, we use the strategy from [Rou99] and compute Gröbner bases using the algorithm F5 [Fau02].

[4] The polynomial F is the characteristic polynomial of m_γ, the multiplication operator by the polynomial γ, in $\mathbb{Q}[x_1, \ldots, x_n]/I$.

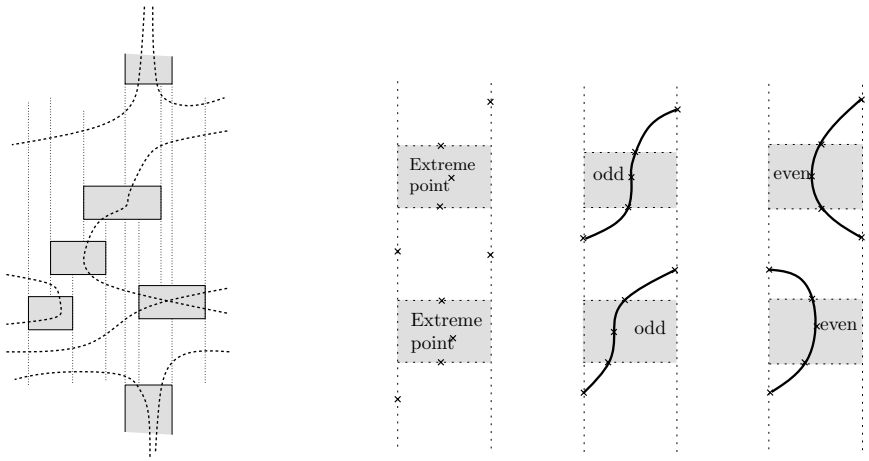

(a) (b)

Figure 2: (a) Example of rectangle decomposition of the plane induced by the isolating boxes (critical and asymptotes). There are boxes for the asymptote, the singular point, and the three extreme points, two of them with even multiplicities and one with odd multiplicity. (b) Possible connections involving extreme points depending on their multiplicities.

3 Algorithm

Our algorithm neither our implementation require assumptions about the existence of vertical lines, but the processing of vertical lines is rather technical and we choose to show first the algorithm for curves without vertical lines and postpone the treatment of vertical lines to section 3.2).

3.1 Curve without vertical lines

As discussed in Section 2, we consider, as input, a curve $\mathcal{C}$ without vertical lines and defined by a square-free polynomial $f \in \mathbb{Q}[x, y]$ In a few words, the algorithm first focuses on critical points, their rational univariate representations enable to compute multiplicities and boxes isolating each point with known topology inside the box. Then a sweep method computes a rectangular decomposition of the plane induced by the boxes of critical points. Eventually the connection is processed in all rectangles with a greedy method using multiplicities in fibers for extreme points. We describe more precisely our algorithm in six steps.

Step 1. Isolating boxes of the singular points and of the x-extreme points. As a general practical rule, the smaller the number of solutions of a system, the easier it is to work with. Hence we split the system of critical points into the system of singular points and the system of extreme points. The system of singular points is the one of critical with in addition the equation $f_x = 0$. The system of extreme points, denoted I_e, is computed by saturation. Indeed, the extreme points are critical for which $f_x \neq 0$, thus we add to the critical system the equation $1 - uf_x = 0$ with a new variable u that we eliminate afterwards. We then compute the RURs of these systems I_s and I_e and isolating boxes for the solutions (see Section 4 for details on interval analysis). We may need to refine the boxes of extreme and singular points to avoid overlaps.

Step 2. Compute multiplicities of critical points in fibers. For extreme points, we use the

Teissier formula: the multiplicity of an extreme point in I_c is the same than in I_e because precisely f_x does not vanish at these points. The multiplicity in I_e is given by the RUR, and hence the multiplicity of an extreme point in its fiber is this number plus one according to the equation of Lemma 2.

For singular points, we use the definition of univariate multiplicity, namely the smallest integer k such that the k^{th} derivative no longer vanishes. Let $I_{s,k}$ be the system of singular points with in addition the equations $f_{y^i} = 0$ for i from 2 to k. Hence we solve, for k increasing from 2, the systems $I_{s,k}$ until it has no solutions. At each step, a singular point which was a solution of $I_{s,k-1}$ but is no longer solution of $I_{s,k}$ has its multiplicity in fiber equal to k.[5]

Step 3. Refinement of the isolating boxes of the x-extreme points. Consider each such box, B, in turn. For each vertical or horizontal side of B, isolate its intersections with $\mathcal{C}$ and refine the box until there are two intersection points. We further refine until there is at most one crossing on the top (resp. bottom) side of B. Note that, unlike comparable algorithms, we do not require that $\mathcal{C}$ intersects the boundary of B on its vertical sides. This is important in practice because, since the curve has a vertical tangent at an x-extreme point, refining until the curve intersects the vertical side is time consuming.

Step 4. Refinement of the isolating boxes of the singular points. We refine these boxes exactly as in [SW05] (see Lemma 3) except for the way the multiplicity k of each singular point in its fiber is computed. In [SW05], k is computed using Sturm-Habicht sequences under the assumption of generic position while we deduce k as explained in Step 2. Then, as in [SW05], every box is refined until the evaluation of f_{y^k} with interval arithmetic does not contain 0 (see Section 4 for details on interval analysis). Further refine the x-coordinates of the box until $\mathcal{C}$ only intersects the vertical boundary of the box.

Step 5. Vertical asymptotes. The x-coordinates of vertical asymptotes are the roots of the leading coefficient $V_a(x)$ of the polynomial $f(x, y)$ considered as a polynomial in y. To deal with an asymptote $x = \alpha$, the idea is, informally, to isolate the point (α, ∞) in a box $[a, b] \times [M, ...\infty..., -M]$ whose vertical sides do not intersect the curve $\mathcal{C}$. The number of crossings with the horizontal sides then gives the number of branches going to $\pm\infty$ with this asymptote (see Figure 2(a)). First, compute an upper bound M_y on the absolute value of the y-coordinates of the y-critical points (this is of course done without computing these critical points, but only the discriminant with respect to y and an upper bound of the roots of this univariate polynomial). Compute also a bound M_x on the absolute value of the y-coordinates of the x-critical points (for which we have already computed boxes). Isolate the roots of the polynomial V_a, hence each root α has an isolating interval $[a, b]$. Substitute $x = a$ (resp. $x = b$) in f and deduce an upper bound, M_1, on the absolute value of the y-coordinates of the intersection of $\mathcal{C}$ and $x = a$ (resp. $x = b$). Set $M = \max(M_1, M_x, M_y)$. Then, a branch crossing the segment $]a, b[\times M$ (resp. $]a, b[\times -M$) goes to $+\infty$ (resp. $-\infty$) with asymptote $x = \alpha$. Finally, we determine whether a given branch is to the left or to the right of the asymptote by comparing the x-coordinates of the asymptote and the crossing point. ⟶ MARC SAYS: This last point is not theoretically needed for the topology but it is needed for the next connection step. It also adds some geometric information. ⟵

Step 6. Connections. For simplicity, all the boxes computed above are called critical boxes and the points at infinity on vertical asymptotes are also called critical. First compute, with a

[5] We also refer the interested reader to a more elegant way to compute the multiplicity in fibers with the Teissier formula [CLP+08]. Experimentally, it appears that this alternative was less efficient because even if it usually needs to work with less systems, these systems are larger (i.e. with more solutions).

sweep-line algorithm, the vertical rectangular decomposition obtained by extending the vertical sides of the critical boxes either to infinity or to the first encountered critical box (see Figure 2(a)). On each of the edges of the decomposition, isolate the intersections with $\mathcal{C}$.[6] Create vertices in the graph corresponding to these intersection points and to the critical points. For describing the arcs connecting these vertices in the graph, we assimilate, for simplicity, the points and the graph vertices. In each critical box the topology is simple: the critical point is connected to each of the intersection points on the boundary of the box.

There are several approaches to do the connections in the other rectangles of the decomposition. The usual and conceptually simplest is to refine boxes of extreme points so as to avoid top and bottom crossings; then, the number of left and right crossings in rectangles always match and the connection is one-to-one. Since we allow top/bottom crossings for efficiency (see above), this straightforward method does not apply. Another approach (see [AMW08, Sak91]) is to compute the sign of the slope of the tangent to the curve at the top/bottom crossings (this yields whether the top/bottom crossing should be connected to vertex to the right or to the left of its rectangle). We however want to avoid such additional computations.

For computing the connections in the non-critical rectangles of the decomposition, we use the multiplicities in fibers of the extreme points and a greedy algorithm. The geometric meaning of the parity of this multiplicity is the following: if it is even, the curve makes a U-turn at the extreme point, else it is odd and the curve is x-monotone in the neighborhood of the extreme point. Still, there are some difficulties for connecting the vertices, as illustrated on Figure 2(b): on the left is the information we may have on the crossings for two extreme points with x-overlapping boxes; the second and third drawings are two possible connections *in the middle rectangle* for different parities of the multiplicities. To distinguish between these two situations we compute the connections in rectangles starting from the top such that the connections in a rectangle below a critical box are computed once the connections in all the rectangles above the box are done.

The connections can easily be computed as follows. We assume here for simplicity that the curve has no vertical asymptote (similar arguments can be used otherwise). Then, due to requirements on extreme point boxes, there is at most one vertex on the top or bottom of any rectangle. The problem is to determine if such a vertex should be connected to another on the right or left side in the rectangle. First, the connections in the unbounded rectangles above critical boxes are straightforward: the connections between the vertices on the two vertical sides are one-to-one starting from infinity and if a vertex remains on a vertical side, there is a vertex on the horizontal side which it has to be connected with. Now, once all the connections have been computed in the rectangle(s) above the box of an extreme point, these connections and the multiplicity of the extreme point yield the connections in the rectangle(s) below, see Figure 2(b). Indeed, if there is a vertex on the bottom side of the critical box, it lies on the top side of a rectangle and, inside this rectangle, the vertex is connected to the topmost vertex on the left or right side, depending on the multiplicity of the extreme point and on the side of the connection of the branch above the extreme point; the other connections in this rectangle and in the possible other rectangles below the critical box are performed similarly as for unbounded rectangles. Note that the two unbounded rectangles (the leftmost and the rightmost) that are vertical half-planes are treated separately in a straightforward manner: for each vertex on the vertical side is associated an arc going to infinity.

[6] For simplicity, we ensure, thanks to refinements, that the curve never intersects an endpoint of an edge, that is, a corner of a rectangle. Note also that the intersections are already isolated on the sides of the critical boxes; an isolating interval may, however, need to be refined if it contains a vertex of the rectangle decomposition.

Output. A graph isotopic to the curve is output. In addition, x-extreme points, singular points and vertical asymptotes are identified and their position is approximated by boxes.

3.2 Curve with vertical lines

We assumed in the above that the curve $\mathcal{C}$ has no vertical line. In order to generalize the algorithm, we explain how to calculate the topology of a curve $\mathcal{C}_F$ defined by a rational bivariate polynomial $F(x, y) = 0$, which has vertical lines. The idea is first to process the curve without its vertical lines, then study the intersections of this curve with the vertical lines. Technically, these two processings are intertwined:

- The vertical lines are points whose x-coordinates are the roots of V_l, the gcd of the coefficients of F seen as a polynomial in y. V_l is an univariate polynomial in x, its roots are isolated (vertical lines can be thought of as vertical *stripes*).
- The curve $\mathcal{C}_f$ with $f = \frac{F}{V_l}$ which has no vertical line, is processed as explained before until Step 4 included.
- We identify the critical points of $\mathcal{C}_f$ that are on vertical lines (this implies solving the singular and extreme ideals with the additional polynomial V_l).
- The x-intervals of critical points of $\mathcal{C}_f$ and vertical lines stripe are refined until a vertical line crosses a critical box if and only if this critical point is on the line.
- We create new boxes, called *vertical*, that contain every point that is an intersection between a vertical line and $\mathcal{C}_f$ and that is not a critical point of $\mathcal{C}_f$.
- We refine the extreme point boxes of $\mathcal{C}_f$ and vertical boxes until there is at most one crossing with $\mathcal{C}_F$ on top (and bottom).
- Then Step 5 of the previous algorithm is performed with the following modifications:
  - For the computation of the bound M_x, we consider all vertical boxes in addition to critical boxes.
  - We identify which vertical lines are also asymptotes by computing $gcd(V_l, V_a)$. We refine vertical lines stripes and asymptote boxes so that a vertical line stripe, whose line is not an asymptote, does not overlap any asymptote box.
  - We add crossing points on asymptote boxes whose asymptote is also a vertical line.
- Finally, Step 6 of the previous algorithm computes a graph isotopic to the curve $\mathcal{C}_F$.

4 Complexity Analysis

This section details the proof of the following theorem stating the bit complexity of our algorithm for the computation of the topology of a curve. We consider the Turing machine model of computation and $\widetilde{\mathcal{O}}_B$ denotes the bit complexity where polylogarithmic factors are omitted. Let an algebraic curve $\mathcal{C}$ be given by a square-free polynomial $f \in \mathbb{Z}[x, y]$ of total degree bounded by d and coefficients of bitsize bounded by τ. Let R be the number of critical points of the curve.

Theorem 4. *The bit complexity of our algorithm for the computation of the topology of the curve $\mathcal{C}$ is $\widetilde{\mathcal{O}}_B(R\,d^{22}\,\tau^2)$, which is $\widetilde{\mathcal{O}}_B(N^{26})$, where $N = \max\{d, \tau\}$.*

Definitions and notation. The bitsize of a rational number is defined as the maximum bitsize of its numerator and denominator. The bitsize of a polynomial is the maximum bitsize of its coefficients. In order to simplify notation, we assume in the sequel that $d = \mathcal{O}(\tau)$. However, we still express complexities in terms of d and τ when it is simple enough since in practice d may be much smaller than τ. We may also assume that the univariate polynomials that we compute with, are square free. This assumption does not change the complexity since the computation of their square-free part and computations with their square-free part, is of no extra cost. Indeed, for a polynomial of degree d and bitsize τ, its square-free part has degree $\mathcal{O}(d)$ and bitsize $\mathcal{O}(d+\tau) = \mathcal{O}(\tau)$ and it can be computed in $\widetilde{\mathcal{O}}_B(d^2\tau)$ [LR01]. We use the notion of separation bound of a polynomial (or of a zero-dimensional system of polynomial equations) which is the minimum distance between its (complex) roots. We call the bitsize of a separation bound s the minimum integer $\sigma \geqslant 0$ such that $s > 2^{-\sigma}$. For simplicity, we do not consider poly-logarithmic factors in the complexity bounds that we denote $\widetilde{\mathcal{O}}$.

Isolations of roots of systems via a RUR require some machinery from interval analysis, we briefly recall the basics and refer to [AH83] for additional details. For an interval $A = [a_1, a_2]$, let its width be $w(A) = a_2 - a_1$ and its absolute value be $|A| = \max(|a_1|, |a_2|)$. For a two-dimensional box $A \times B$, let $w(A \times B) = \max(w(A), w(B))$ and $|A \times B| = \max(|A|, |B|)$. We denote by $I(\mathbb{R}^n)$ the set of products of n real intervals. A polynomial has a natural extension to a function over $I(\mathbb{R}^n)$ by a process we call evaluation with interval arithmetic. More precisely, for a polynomial f, we denote by the bold letter $\mathbf{f}$ the corresponding interval function defined over $I(\mathbb{R}^n)$ by replacing the usual operations $+, -, \times$ by interval operations. In the sequel, we bound the number of times we refine the isolating intervals of the roots of the univariate polynomial of the RUR. We assume that every refinement divides by at least two the interval width.

Overview of the section. In order to derive the bit complexity of the algorithm presented in Section 3 we analyze the complexity of each step. The first two steps consists of computing the isolating boxes of the critical points and multiplicities. The analysis of these steps is presented in Section 4.2. In the third step, we refine the isolating boxes of the x-extreme points. The complexity of this step is bounded by the complexity of the next one, so we do not consider it explicitly. During the first part of the fourth step we refine the isolating boxes of the singular points with respect to their isolating curve. The analysis of this step is done in Section 4.3. During the second part, we refine the isolating boxes of the singular points, until the curve does not intersect them on the top and on the bottom. The analysis of this operation is done in Section 4.4. We conclude the proof of Theorem 4 in Section 4.5. Some technical details are postponed to Section 4.6.

4.1 Preliminaries

Before analyzing each step of the algorithm, we state the bitsize complexity of some recurrent basic computations. Let f be a univariate integer polynomial of degree d and bitsize τ, and x be a rational of bitsize σ.

Lemma 5 ([Yap00])**.** *The bitsize of the separation bound of f is in $\mathcal{O}(d\tau)$. Similarly, the bitsize of the endpoints of isolating intervals of the roots of f is in $\mathcal{O}(d\tau)$. Moreover, the absolute value of the roots of f is in $\mathcal{O}(2^\tau)$.*

We omit the proof of the following lemma which is straightforward.

Lemma 6. *The evaluation of f over x has complexity $\widetilde{\mathcal{O}}_B(d(\tau + d\sigma))$, while the number $f(x)$ has bitsize $\mathcal{O}(\tau + d\sigma)$.*

Since the interval evaluation of the operations $+, -, \times$ are a constant time more expansive that their usual counterparts we obtain the corollary:

Corollary 7. *The evaluation of f using interval arithmetic over an interval I with end points of the same bitsize as x, i.e. the computation of $\mathbf{f}(I)$, has complexity $\widetilde{\mathcal{O}}_B(d(\tau + d\sigma))$, while the interval $\mathbf{f}(I)$ has end points with bitsize $\mathcal{O}(\tau + d\sigma)$.*

We postpone to Section 4.6 the proof of the following lemma bounding the increase of the width of an interval by evaluation by interval arithmetic of a polynomial.

Lemma 8. *Let P be a univariate rational polynomial of degree d and bitsize τ, and A be an interval such that $|A| \leqslant 2^\sigma$ with $\sigma \geqslant 0$; then*

$$w(\mathbf{P}(A)) \leqslant 2^{\tau + d\sigma} d^2 w(A).$$

Let Q be a bivariate rational polynomial of total degree d and bitsize τ, and B be an interval such that $|B| \leqslant 2^\sigma$ with $\sigma \geqslant 0$; then

$$w(\mathbf{Q}(A, B)) \leqslant 2^{\tau + d\sigma + 1} d^3 w(A \times B).$$

4.2 Computation of the isolating boxes

The first step of the algorithm is the computation of the isolating boxes of the singular and the extreme points. For the complexity of this step it suffices to compute the complexity of solving the system of critical points $I_c = \mathbb{I}(f, f_y)$.

The steps for solving the system are the followings. We compute the Gröbner basis of I_c and the RUR of the system. We solve the univariate polynomial of the RUR and, using the isolating intervals of its real roots, we compute boxes that contain the real solutions of the systems. Finally, we refine the boxes until they are all disjoint.

Gröbner basis and RUR. We compute the Gröbner basis of I_c with the degree reverse lexicographic order in $\widetilde{\mathcal{O}}_B(d^8\tau)$ [Laz83]. Next we compute the RUR representation of the solution of the system and multiplicities in $\widetilde{\mathcal{O}}_B(d^{13}\tau)$, the reader may refer to [Rou99] for more details.

The RUR representation has the following form: $h(T) = 0, x = \frac{g_x(T)}{g_0(T)}, y = \frac{g_y(T)}{g_0(T)}$, where $h, g_0, g_x, g_y \in \mathbb{Q}[T]$. The polynomial $h(T)$ is actually the so called *u-resultant* [Can88]. It has degree $\mathcal{O}(d^2)$ and bitsize $\mathcal{O}(d^2 + d\tau) = \mathcal{O}(d\tau)$. One way to see this is to consider a curve (and thus the system I_c) in generic position, this is without loss of generality since shearing the curve in generic position does not increase the bitsizes more than by a factor in $\mathcal{O}(log(d))$. In this case we can consider as the polynomial $h(T)$ the projection of the system on the x-axis or, in other words, the resultant of the system with respect to y. Under this notion, g_0 and g_x, respectively g_0 and g_y could be seen as the coefficients of the first sub-resultant of f and f_y, with respect to y, respectively with respect to x. Thus, the degree of these polynomials is $O(d^2)$ and their bitsize is $\mathcal{O}(d\tau)$.

We can simplify the expressions of x and y, if we take into account that $h(T)$ and $g_0(T)$ are relative prime. Indeed, we can compute a polynomial $g(T)$ such that $g \cdot g_0 = 1 \mod h$ and thus

express the coordinates of the solutions as $x = h_x(T) = g_x(T)\, g(T)$ and $y = h_y(T) = g_y(T)\, g(T)$. The degree of g is $\mathcal{O}(d^2)$ and its bitsize is $\mathcal{O}(d^3\tau)$, as g is a Bézout coefficient of the extended Euclidean division of g_0 and h [vzGG03]. The same bounds hold for $h_x(T)$ and $h_y(T)$.

Computation of the boxes. In order to solve the system, it suffices to compute the isolating intervals of the real roots of h, and substitute them in h_x and h_y using interval arithmetic. This gives isolating boxes that contain the real roots of the system. Finally, we refine these boxes until they become disjoint.

Lemma 9. *To ensure that the boxes of the critical points are disjoint, it is sufficient to refine $\mathcal{O}(d^3\tau)$ times each isolating interval of the corresponding root of the univariate polynomial of the RUR of the critical points.*

Proof. One needs to refine the isolating intervals of the real roots of h until the corresponding boxes of the system, computed by interval evaluation with h_x and h_y, become disjoint. In other words, the isolating boxes should have width smaller than 2^{-s_c}, where s_c is the separation bound of the system of critical points. If $2^{-\mu}$ is a lower bound on the width of the isolating intervals of h, Lemma 8 yields a value μ, so that the evaluations by h_x and h_y give intervals of width at most 2^{-s_c}. We consider only the polynomial h_x, since the computation is similar for h_y. Lemma 8 applied with h_x of bitsize $\mathcal{O}(d^3\tau)$ and degree $\mathcal{O}(d^2)$ evaluated at the roots of h of absolute value $\mathcal{O}(2^{d\tau})$ (by Lemma 5) yields:

$$2^{\mathcal{O}(d^3\tau)+\mathcal{O}(d^2)\mathcal{O}(d\tau)} d^2 2^{-\mu} \leqslant 2^{-s_c}.$$

Thus, it suffices to consider μ in $\mathcal{O}(s_c + d^3\tau)$.

On the other hand, the separation bound of the critical points is larger than the separation bounds of the x (or y) coordinates of these points. The coordinates are roots of the resultant of f and f_y with respect to y (or x), which is a polynomial of degree $\mathcal{O}(d^2)$ and bitsize $\mathcal{O}(d\tau)$. Hence the separation bound of the coordinates is of bitsize $\mathcal{O}(d^3\tau)$ (by Lemma 5). This is also a bound for the separation bound of the critical points thus s_c is in $\mathcal{O}(d^3\tau)$ and $\mathcal{O}(s_c + d^3\tau)$ is also in $\mathcal{O}(d^3\tau)$.

By Lemma 5, the isolating intervals of h have initial width $\mathcal{O}(2^{d\tau})$. Thus, to reduce them to the width $2^{-\mu}$, it suffices to perform $\mathcal{O}(d\tau + d^3\tau) = \mathcal{O}(d^3\tau)$ refinements, provided that the refinements divide the interval widths by at least two. □

4.3 Refinement with respect to the isolating curve

We require that the isolating boxes of singular points avoid their associated curve $f_k = \partial^k f/\partial y^k$. This is ensured by refining the isolating boxes of the singular points so that the evaluation of f_k, using interval arithmetic, results an interval that does not contain zero.

Lemma 10. *To ensure that the boxes of singular points do not overlap their associated curve f_k, it is sufficient to refine $\widetilde{\mathcal{O}}(d^4\tau)$ times the corresponding roots of the univariate polynomial of the RUR.*

Proof. Consider a box $B = I \times J$ that isolates a singular point p. Assume, without loss of generality, that the width of B is $2^{-\mu}$. Let k be such that f_k is the isolating curve for p. We need to ensure that the evaluation of f_k over B does not contain 0. A sufficient condition is that $w(\mathbf{f_k}(B)) < \delta$ with $\delta < |f_k(p)|$. Defining s_v as the maximum bitsize of the values of f_k at the singular points of f

(more precisely $s_v = \max_{s\in\mathbb{N}^*}\{|f_k(p)| > 2^{-s}, p \text{ singular point of } f \text{ and } f_k(p) \neq 0\}$), we can choose $\delta = 2^{-s_v}$.

f_k is of degree $\mathcal{O}(d)$ and bitsize $\widetilde{\mathcal{O}}(d\tau)$ (by Stirling formula). The absolute value of a box of a singular point is $\mathcal{O}(2^{d\tau})$ since the x or y coordinates of a box are roots of the resultant of f and f_y with respect to y or x, and such a resultant has bitsize $\mathcal{O}(d\tau)$ (Lemma 5). Lemma 8 applied for the evaluation of f_k over a singular point box yields:

$$2^{\widetilde{\mathcal{O}}(d\tau)+\mathcal{O}(d)\mathcal{O}(d\tau)+1} d^3 2^{-\mu} \leqslant 2^{-s_v}.$$

Thus, it suffices to consider μ in $\mathcal{O}(s_v + d^2\tau)$.

On the other hand, Lemma 13 gives that s_v is in $\widetilde{\mathcal{O}}(d^4\tau)$, hence μ is in $\widetilde{\mathcal{O}}(d^4\tau)$.

Let $2^{-\nu}$ be the width of the isolating interval of h that corresponds to the singular point via the RUR. Then ν should be such that the box computed via the RUR has width $2^{-\mu}$. Applying Lemma 8 again, for the evaluation of h_x of degree $\mathcal{O}(d^2)$ and bitsize $\mathcal{O}(d^3\tau)$ over this isolating interval of absolute value $\mathcal{O}(2^{d\tau})$, ν should satisfy:

$$2^{\mathcal{O}(d^3\tau)+\mathcal{O}(d^2)\mathcal{O}(d\tau)} d^2 2^{-\nu} \leqslant 2^{-\mu}.$$

Thus, it suffices to consider ν in $\widetilde{\mathcal{O}}(d^4\tau)$.

We conclude, as in the proof of Lemma 9, that it suffices to perform $\widetilde{\mathcal{O}}(d^4\tau)$ refinements. □

4.4 Refinement of the singular points to avoid top/bottom crossings

The last step that we need to consider is the analysis of the refinement of the isolating boxes of the singular points, until there is no intersection between the curve and their top and bottom sides.

Lemma 11. *To avoid top/bottom intersection between the curve and the boxes of singular points, it is sufficient to refine $\widetilde{\mathcal{O}}(d^9\tau)$ times the corresponding roots of the univariate polynomial of the RUR.*

Proof. Consider a singular point and its isolating box refined according to Lemmata 9 and 10. We further refine the x-coordinate of the box until the top (or equivalently the bottom) side does not intersect the curve. The line supporting the top side is of the form $y = c$, for some constant c. Hence, the x-coordinates of the intersections of the curve with the top side are among the roots of $f(x, c)$. Consider the polynomial P whose roots are the roots of $f(x, c)$ and the x-coordinates of the critical points. A sufficient condition to avoid intersection on the top side is to ensure that the x-width of the box is smaller than the separating bound of P.

The bitsize of c is the same as that of the evaluation of h_x over an end-point a of an isolating interval of a root of h. From Lemmata 9 and 10, the bitsize of a is in $\widetilde{\mathcal{O}}(d^4\tau)$. Since h_x is of degree $\mathcal{O}(d^2)$ and bitsize $\mathcal{O}(d^3\tau)$, c has bitsize $\widetilde{\mathcal{O}}(d^6\tau)$ (Lemma 6). Hence $f(x, c)$ is a polynomial of degree $\mathcal{O}(d)$ and bitsize $\widetilde{\mathcal{O}}(d^7\tau)$. The polynomial P is the product of $f(x, c)$ and the resultant with respect to y of f and f_y, thus its degree is in $\mathcal{O}(d^2)$ and its bitsize is in $\widetilde{\mathcal{O}}(d^7\tau)$. The bitsize of the separation bound of P is thus δ in $\widetilde{\mathcal{O}}(d^9\tau)$

Let $2^{-\mu}$ be the width of the isolating intervals of h corresponding to the singular point by the RUR. Lemma 8 applied with h_x of bitsize $\mathcal{O}(d^3\tau)$ and degree $\mathcal{O}(d^2)$ yields:

$$2^{\mathcal{O}(d^3\tau)+\mathcal{O}(d^2)\mathcal{O}(d\tau)} d^2 2^{-\mu} \leqslant 2^{-\delta}$$

Thus, it suffices to consider μ in $\widetilde{\mathcal{O}}(d^9\tau)$.

We conclude, as in the proof of Lemma 9, that it suffices to perform $\widetilde{\mathcal{O}}(d^9\tau)$ refinements.

□

4.5 Overall complexity

Combining the results from Section 4.2, 4.3 and 4.4 we can now prove Theorem 4.

Proof of Theorem 4. Combining the results of Lemmata 9, 10, and 11, we prove that the algorithm performs $\widetilde{\mathcal{O}}(d^9\tau)$ refinements. Each refinement consists of an evaluation of h, h_x and h_y over a rational number of bitsize $\widetilde{\mathcal{O}}(d^9\tau)$. Using Horner's rule, each evaluation of these polynomials of degree $\mathcal{O}(d^2)$ and bitsize $\mathcal{O}(d^3\tau)$ has complexity $\widetilde{\mathcal{O}}_B(d^2(d^3\tau + d^2d^9\tau)) = \widetilde{\mathcal{O}}_B(d^{13}\tau)$ (Lemma 6). The complexity of the $\widetilde{\mathcal{O}}(d^9\tau)$ refinements is thus in $\widetilde{\mathcal{O}}_B(d^{13}\tau d^9\tau) = \widetilde{\mathcal{O}}_B(d^{22}\tau^2)$. If there are R singular points the total cost is thus $\widetilde{\mathcal{O}}_B(Rd^{22}\tau^2)$.

Note that the costs of Gröbner and RUR computations are dominated. Finally, the complexity of dealing with vertical asymptotes (Step 5), vertical lines and the connection part of the algorithm (Step 6) is dominated by the complexity of the other steps.

Finally, if $N = \max\{d, \tau\}$, note that R is in $\mathcal{O}(d^2) = \mathcal{O}(N^2)$, and so the total complexity of the algorithm is $\widetilde{\mathcal{O}}_B(N^{26})$. □

4.6 Technical details

4.6.1 Interval arithmetic

In our algorithm, we need, several times, to evaluate univariate and bivariate polynomials over intervals. This is done using classical interval arithmetic operations. Theoretically, we need to control how large an interval becomes when a polynomial operation is performed. We prove here Lemma 8, which we recall for convenience.

> *Let P be a univariate rational polynomial of degree d and bitsize τ, and A be an interval such that $|A| \leqslant 2^\sigma$ with $\sigma \geqslant 0$; then $w(\mathbf{P}(A)) \leqslant 2^{\tau+d\sigma}d^2w(A)$.*
>
> *Let Q be a bivariate rational polynomial of total degree d and bitsize τ, and B be an interval such that $|B| \leqslant 2^\sigma$ with $\sigma \geqslant 0$; then $w(\mathbf{Q}(A,B)) \leqslant 2^{\tau+d\sigma+1}d^3\max(w(A), w(B))$.*

Proof of Lemma 8. We apply the basic formulas for the sum and the product of intervals [AH83, Theorem 9, p.15], which are for any real number a and integer $n \geqslant 1$:

$$\begin{aligned} &w(A \pm B) = w(A) + w(B), && w(aA) = |a|w(A), \\ &w(AB) \leqslant w(A)|B| + |A|w(B), && w(A^n) \leqslant n|A|^{n-1}w(A). \end{aligned}$$

Let $P(x) = \sum_{i=0}^{d} c_i x^i$ with $|c_i| \leqslant 2^\tau$ and $Q(x,y) = \sum_{i,j\geqslant 0}^{i+j\leqslant d} c_{ij} x^i y^j$ with $|c_{ij}| \leqslant 2^\tau$. We have:

$$\begin{aligned}
w(\mathbf{P}(A)) &= \sum_{i=0}^{d} |c_i| w(A^i) \quad \leqslant \quad 2^\tau \sum_{i=0}^{d} i|A|^{i-1} w(A) \quad \leqslant \quad 2^\tau w(A) d \sum_{i=1}^{d} |A|^{i-1} \\
&\leqslant 2^\tau w(A) d^2 \max(1, |A|^{d-1}) \quad \leqslant \quad 2^\tau w(A) d^2 2^{d\sigma} \quad \leqslant \quad 2^{\tau + d\sigma} d^2 w(A). \\
w(\mathbf{Q}(A,B)) &= \sum |c_{ij}| w(A^i B^j) \quad \leqslant \quad 2^\tau \sum w(A^i)|B|^j + w(B^j)|A|^i \\
&\leqslant 2^\tau \sum i w(A)|A|^{i-1}|B|^j + j w(B)|A|^i |B|^{j-1} \\
&\leqslant 2^\tau d^3 2 w(A \times B) \max(1, 2^{\sigma(d-1)}) \quad \leqslant \quad 2^{\tau + d\sigma + 1} d^3 w(A \times B).
\end{aligned}$$

□

4.6.2 Separation bounds

In this section, we compute the separation bound needed for the proof of Lemma 10. We first need a refinement, due to Yap [Yap00], of Gap theorem [Can88].

Theorem 12 ([Yap00] Gap theorem 11.45)**.** *Let $\Sigma = \{A_1, \ldots, A_n\} \subseteq \mathbb{Z}[x_1, \ldots, x_n]$ be a system of n polynomials, not necessarily homogeneous. Suppose that Σ has finitely many complex zeros and $(\xi_1, \ldots, \xi_n)$ is one of these zeros. Assume $d_i = \deg(A_i)$ and*

$$K := \max\{\sqrt{n+1}, \max\{\|A_i\|_2 \,|\, 1 \leqslant i \leqslant n\}\},$$

where $\|A_i\|_2$ is the usual Euclidean norm of the vector of coefficients of the polynomial A_i. If $|\xi_i| \neq 0$, $i = 1, \ldots, n$, then

$$|\xi_i| > (2^{2/3} N K)^{-D}\, 2^{-(n+1) d_1 \cdots d_n},$$

where

$$N := \binom{1 + \sum_{i=1}^{n} d_i}{n}, \quad D := \left(1 + \sum_{i=1}^{n} \frac{1}{d_i}\right) \prod_{i=1}^{n} d_i.$$

We now prove, following closely [Yap06], the lemma we use in the proof of Lemma 10.

Lemma 13. *Let $\mathbf{p}$ be a critical point of a curve $\mathcal{C}_f$, without vertical lines, defined by a square-free polynomial f of degree d and bitsize τ. Let $f_k = \partial^k f / \partial y^k$, where $2 \leqslant k \leqslant d$. If $\mathbf{p}$ is not a point on f_k, then $|f_k(\mathbf{p})| > 2^{-s_v}$ with s_v in $\widetilde{\mathcal{O}}(d^4\tau)$.*

Proof. We consider the following system for $k \in \{2, \ldots, d\}$

$$\begin{aligned}
A_1 &: f(\mathbf{p}) = 0, \\
A_2 &: f_y(\mathbf{p}) = 0, \\
A_3 &: h - f_k(\mathbf{p}) = 0.
\end{aligned}$$

The system is zero-dimensional because the number of critical point is finite, and we can apply Theorem 12, where

$$d_1 = d, \quad d_2 = d - 1, \quad d_3 \leqslant 2(d-k),$$

$$N \leqslant \binom{4d - 2k}{3} \leqslant 16 d^3, \quad D \leqslant (d+2)(d-1)(2d-1) \leqslant 4d^3.$$

The bitsize of the norm of a polynomial is the bitsize if the polynomial itself. Since the bitsize of f and f_y is in $\mathcal{O}(\tau)$ and that of f_k is in $\widetilde{\mathcal{O}}(d\tau)$, we conclude that the bitsize of K is in $\widetilde{\mathcal{O}}(d\tau)$. The factor that gives the bitsize in the lower bound of Theorem 12 is thus K^{-D} and h is bounded by a value of bitsize in $\widetilde{\mathcal{O}}(d^4\tau)$. □

5 Experiments

We implemented our algorithm, ISOTOP, in MAPLE using the Gb/RS MAPLE package [FGb, RS], implemented in C, for computing Gröbner bases, RURs and isolating roots. We believe that comparing MAPLE and C/C++ implementations is fair for our problem when the running time is not too small because then, most of the time is usually spent on algebraic computations which are coded in C/C++ (possibly in the kernel of MAPLE). When the running time is too small, the MAPLE part of the code is not negligible and comparing MAPLE and C/C++ implementations becomes meaningless. This is why we focused our tests on examples for which the running time exceeds 1 second. We measure the running time for computing the isotopic graph, but not the drawing. All the experiments were performed using 2.6 GHz single-core Pentium 4 with 1.5Gb of RAM and 512kb of cache, running 32-bit linux.

We compared our code, ISOTOP, with two C++ implementations Alcix [EKW07] and Cad2d [Bro02] and two MAPLE implementations, Top [GVN02] and Insulate [SW05]. Another promising software is Axel, [AMW08] but no implementation of the certified subdivision algorithm is currentlty available.

Alcix is a C++ code, part of the CGAL library [CGA].[7] [8] Cad2d is a standalone C++ code which can also be compiled in combination with the Singular library [GPS01] (used for polynomial factorization). In our tests, Cad2d appears to be much more efficient when ran with Singular (and we report these tests). Finally, recall that Top requires an initial precision, which we set to 50.

As discussed in Section 1, the various implementations do not compute exactly the same thing and comparisons should thus be taken with care. Recall that when the curve is not in generic position, Top and Insulate do not compute the critical points (and, in particular, the x-extreme points) in the original coordinate system. ISOTOP, Alcix and Cad2d always output the critical points in the original coordinate system.

We ran large scale benchmarks on over a thousand of curves during several weeks. In particular, we considered curves suggested in [Lab08, Bro02, GVN02] and several classes of non-generic curves.[9] We considered about 1300 curves from [Lab08], which are classified in 18 challenges covering a large variety of interesting cases such as isolated points, high multiplicity of tangency at singularities, large number of branches at singularities or many singularities. This set contains curves of degree up to 90 that are both in generic and non-generic position. As suggested in [Bro02], typical curves in generic position can be generated (i) as a random bivariate polynomial (which usually do not have singular points) or (ii) as resultants of two random trivariate polynomials (which usually have singular points, including isolated points). In both cases, we considered random polynomials with 50% non-zero coefficients of bitsize 32 in Case (i) and initial bitsize 8 in Case (ii). We generated such curves with degrees up to 25. We also generated classes of curves in non-generic position

[7] Following the recommendation of M. Kerber, we ran two versions of the code with the flag CGAL_ACK_RESULTANT_FIRST_ STRATEGY set to 1 and 0. One being optimized for generic cases, while the other is optimized for singular curves. We always compare to the better running time.

[8] Note to the reviewer: Alcix is not yet in the official distribution of CGAL but should be soon.

[9] The logs are available at `http://www.loria.fr/equipes/vegas/isotop/benchmarks`.

in two different ways. First, we considered products of a curve with one or several of its vertical translates. Second, we considered curves of the type $g = f^2(x, y) + f^2(x, -y)$; such curves are usually irreducible and consist of isolated points which are the intersections of the curve C_f with its symmetric with respect to the x-axis. We generated such curves with degrees up to 24.

We set in our experiments a limit of 30 minutes for the computation of the topology of one curve. We report as *time out* instances that exceed this running time. Also, Cad2d which uses Singular for modular arithmetic, often reports on difficult instances that the table of primes has been exhausted, which results in an interruption of computation; this is reported in the tables as *aborted*.

In summary, we ran our benchmarks on a total of 1500 curves. As mentioned above, it is not significant to compare C++ and Maple implementations when the running time is too small. We thus only report experiments on 650 curves whose running times exceeded 1 second for ISOTOP. The distribution of degrees and number of critical points of these 650 curves is shown in Figure 3.

Figure 4 shows the ratio of running times between each of the competing implementations and ISOTOP over our set of 650 curves. It appears difficult to analyze the benchmarks globally because there are always particular examples that are processed faster by a given implementation. We note, however, that Insulate is almost always slower than ISOTOP, except for random curves with no singular points. In addition, Insulate and Top reached the time limit on more than half of the examples and, in particular, on difficult examples. We can, nevertheless, comment on the general behavior of the different approaches depending on the classes of examples.

To illustrate the behavior on curves in generic positions, we report the running times for random curves in Table 1 and for resultants of surfaces in Table 2. Random curves have no singular points and few extreme points. In this case, we observe that ISOTOP is the least efficient implementation. This can be explained by the fact that ISOTOP computes the Gröbner basis of a large system without multiplicities, which is the worst case in practice. On the other hand, the other implementations benefit from interval arithmetic filters in the lifting phase, which speed up computations by avoiding expensive symbolic computations, see for example [Bro02]. Generic curves generated as resultants have many singularities and extreme points. ISOTOP benefits from splitting the critical system in two smaller (singular and extreme) systems and hence it performs relatively better than in the completely random case. We observe that ISOTOP is typically a bit slower than Alcix but faster than Top, and that Cad2d aborts.

To illustrate the behavior on curves in non-generic position, we consider different classes of curves. The first class of non-generic curves are constructed with one curve multiplied by one or several of its vertical translates. The initial curve is taken either randomly, in Table 3, or it is a resultant of two surfaces, in Table 4. Table 5 reports results on the second class of non-generic curves of the type $f^2(x, y) + f^2(x, -y)$ for random polynomials f. For these non-generic curves, ISOTOP is typically faster than other implementations.

As a general rule, we observe that, except for random curves, that is, curves in generic position and without singular point, the ratio of the running times between other implementations and ISOTOP is increasing with the degree of the curve. In other words, except for random curves, ISOTOP tends to perform better, compared to others, when the degree increases.

6 Conclusion

We presented a new algorithm and implementation for computing the topology of plane algebraic curves. Instead of a CAD (cylindrical algebraic decomposition) based method, our algorithm relies

upon Gröbner bases, Rational Univariate Representations and hence avoids computations with algebraic numbers even in non-generic cases. A strength of our approach is to be insensitive to the non-genericity of the curve. As demonstrated by the experiments, our implementation is competitive with state-of-the-art C/C++ implementations in the case of generic curves and faster for high-degree non-generic ones. Future work includes taking advantage of the possible decomposition of the curve into factors. We already decompose the system of critical points into systems of extreme and singular points. One natural step further would be to consider the primary decompositions of the ideals and thus work with systems of lower complexity.

Acknowledgments

We would like to thank the authors of Alcix, Cad2d, Insulate, and Top for supplying their code and O. Labs for his list of curve challenges.

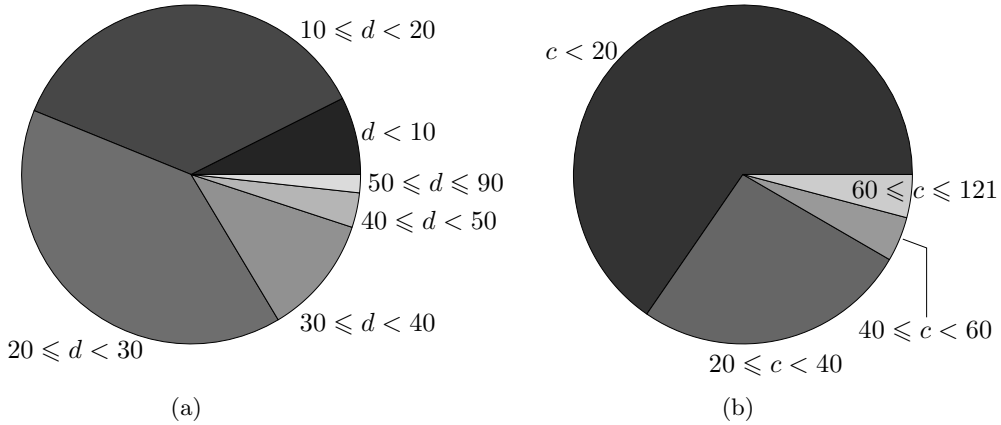

Figure 3: Distributions of (a) degrees and (b) number of critical points in our 650 examples.

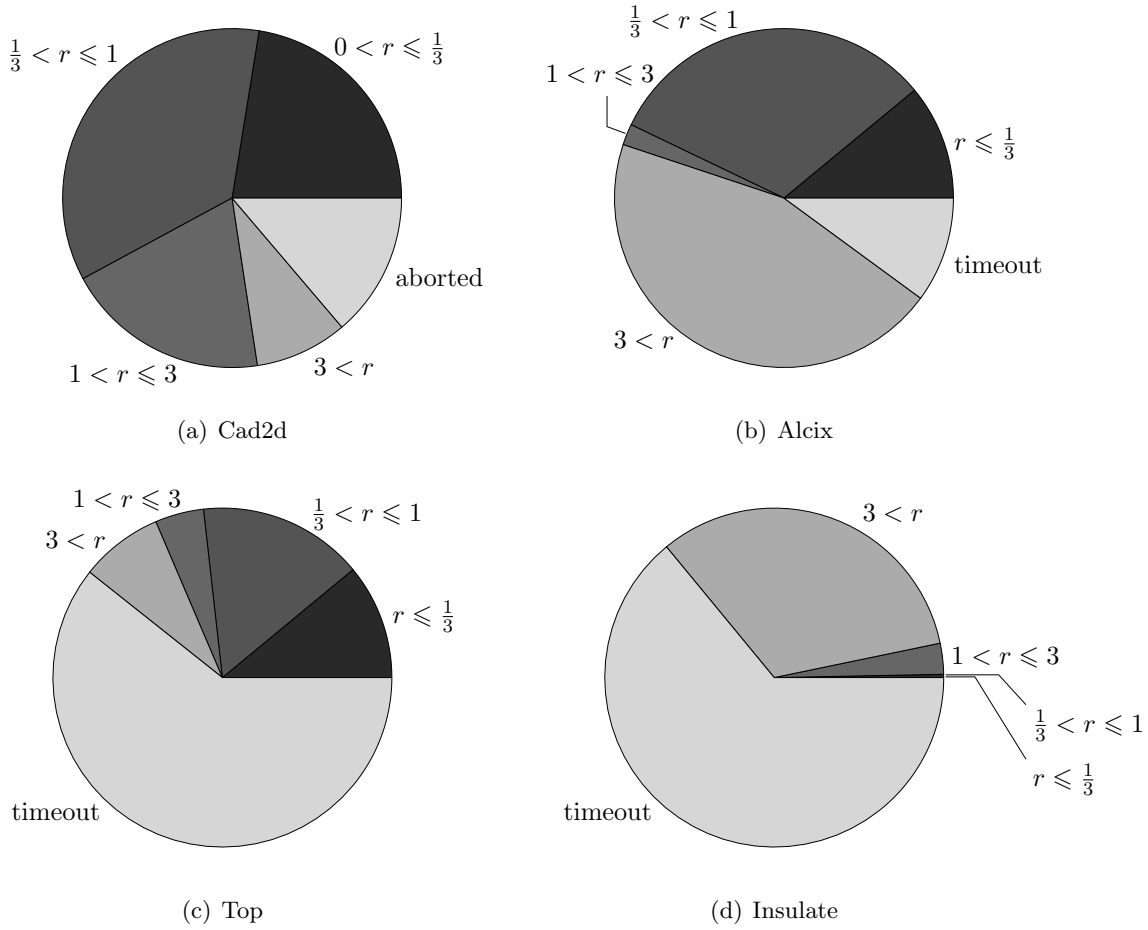

Figure 4: Distributions of running time ratios for (a) Cad2d, (b) Alcix, (c) Top, and (d) Insulate over ISOTOP. Timeout means that the limit of 30 minutes was reached.

degree	bitsize	ISOTOP	Alcix	Cad2d	Top	Insulate	$\frac{\text{Alcix}}{\text{ISOTOP}}$	$\frac{\text{Cad2d}}{\text{ISOTOP}}$	$\frac{\text{Top}}{\text{ISOTOP}}$	$\frac{\text{Insulate}}{\text{ISOTOP}}$
10	32	3.8	0.42	1.6	1.4	4.3	0.1	0.4	0.4	1.1
12	32	9.7	2.2	2.3	4.8	8.1	0.2	0.2	0.5	0.8
14	32	22	3.7	4.3	15	26	0.2	0.2	0.7	1.2
16	32	72	4.5	6.1	38	39	0.1	0.1	0.5	0.5
18	32	160	16	12	94	98	0.1	0.1	0.6	0.6
20	32	320	31	25	210	120	0.1	0.1	0.7	0.4

Table 1: Running times in seconds (averaged over 5 runs) for random bivariate polynomials with 50% non-zero coefficients of bitsize 32.

degree	bitsize	ISOTOP	Alcix	Cad2d	Top	Insulate	$\frac{\text{Alcix}}{\text{ISOTOP}}$	$\frac{\text{Cad2d}}{\text{ISOTOP}}$	$\frac{\text{Top}}{\text{ISOTOP}}$	$\frac{\text{Insulate}}{\text{ISOTOP}}$
16	64	29	14	aborted	52	>600	0.5	-	1.8	>20
25	80	590	410	aborted	>1800	>1800	0.7	-	>4.3	>3

Table 2: Running times (averaged over five runs) in seconds for resultants of two random trivariate polynomials, both of total degree 4 or 5, with 50% non-zero coefficients of bitsize 8.

degree	τ	ISOT.	Alcix	Cad2d	Top	Insulate	$\frac{\text{Alcix}}{\text{ISOTOP}}$	$\frac{\text{Cad2d}}{\text{ISOTOP}}$	$\frac{\text{Top}}{\text{ISOTOP}}$	$\frac{\text{Insulate}}{\text{ISOTOP}}$
12	96	4.1	6.5	1.6	9 or >600	>600	1.6	0.4	2.2 or >92	>140
15	96	15	38	10	40 or >600	>600	2.5	0.7	2.7 or >16	>40
18	96	49	120	66	>600	>600	2.4	1.3	> 12	>12
21	96	140	510	510	>1200	>1200	3.5	3.6	> 8.6	>9

Table 3: Running times (averaged over five runs) in seconds for non-generic curves generated by the product of a curve f with two of its translates $f(x, y+1)$ and $f(x, y+2)$. The curve f is chosen randomly with degree between 4 and 7 and 50% non-zero coefficients of bitsize 32, resulting in curves of bitsize τ. The running-time discrepancy of Top is large and is not averaged.

d_1	d_2	degree	τ	ISOT.	Alcix	Cad2d	Top	Insulate	$\frac{\text{Alcix}}{\text{ISOTOP}}$	$\frac{\text{Cad2d}}{\text{ISOTOP}}$	$\frac{\text{Top}}{\text{ISOTOP}}$	$\frac{\text{Insulate}}{\text{ISOTOP}}$
3	3	18	96	26	93	160	>600	>600	3.6	6.3	>20	>20
3	4	24	112	250	510 or >1800	abort.	>1800	>1800	2 or >7	-	>20	>20

Table 4: Running times (averaged over five runs) in seconds for non-generic curves generated by the product of a curve f with its translate $f(x, y+1)$. The curve f is the resultant of two random trivariate polynomials of total degree d_1 and d_2 and 50% non-zero coefficients of bitsize 8. The running-time discrepancy of Alcix is large for degree 24 and is not averaged.

degree	bitsize	ISOTOP	Alcix	Cad2d	Top	Insulate	$\frac{\text{Alcix}}{\text{ISOTOP}}$	$\frac{\text{Cad2d}}{\text{ISOTOP}}$	$\frac{\text{Top}}{\text{ISOTOP}}$	$\frac{\text{Insulate}}{\text{ISOTOP}}$
10	64	39	180	25	>600	>600	4.5	0.6	>15	>15
12	64	240	300	aborted	>1800	>1800	1.2	-	>7	>7
14	64	350	>1800	aborted	>1800	>1800	> 5	-	>5	>5

Table 5: Running times (averaged over five runs) in seconds for polynomials of the form $g = f^2(x, y) + f^2(x, -y)$. The random polynomials f have bitsize 32, term density 50% and degrees varying between 5 and 7.

References

[ACM84] D. S. Arnon, G. E. Collins, and S. McCallum. Cylindrical algebraic decomposition ii: An adjacency algorithm for the plane. *SIAM J. Comput.*, 13(4):878–889, 1984.

[AH83] G. Alefeld and J. Herzberger. *Introduction to Interval Computations*. Academic Press, New York, 1983.

[AM88] D. Arnon and S. McCallum. A polynomial time algorithm for the topological type of a real algebraic curve. *J. Symbolic Computation*, 5:213–236, 1988.

[AMW08] L. Alberti, B. Mourrain, and J. Wintz. Topology and arrangement computation of semi-algebraic planar curves. *Comput. Aided Geom. Des.*, 25(8):631–651, 2008.

[BFM+01] C. Burnikel, S. Funke, K. Mehlhorn, S. Schirra, and S. Schmitt. A separation bound for real algebraic expressions. In *Proc. 9th Annual European Symposium on Algorithms*, volume 2161 of *LNCS*, pages 254–265. Springer-Verlag, 2001.

[BPR06] S. Basu, R. Pollack, and M.-R. Roy. *Algorithms in Real Algebraic Geometry*, volume 10 of *Algorithms and Computation in Mathematics*. Springer-Verlag, 2nd edition, 2006.

[BR90] R. Benedetti and J.J. Risler. *Real Algebraic and Semi-algebraic Sets, Actualites Mathematiques*. Hermann, 1990.

[Bro01] C. W. Brown. Improved projection for cylindrical algebraic decomposition. *J. Symb. Comput.*, 32(5):447–465, 2001.

[Bro02] C. W. Brown. Contructing cylindrical algebraic decomposition of the plane quickly, 2002. Manuscript, `http://www.cs.usna.edu/~wcbrown/`.

[BSGY08] M. Burr, S.W.Choi, B. Galehouse, and C. Yap. Complete subdivision algorithms, ii: Isotopic meshing of singular algebraic curves. In *Proc. Intl. Symp. on Symbolic & Algebraic Computation (ISSAC 2008)*, 2008.

[Can88] J. Canny. *The Complexity of Robot Motion Planning*. MIT, 1988.

[CGA] CGAL: Computational Geometry Algorithms Library. `http://www.cgal.org`.

[CLO05] D. Cox, J. Little, and D. O'Shea. *Using Algebraic Geometry*. Number 185 in Graduate Texts in Mathematics. Springer, New York, 2nd edition, 2005.

[CLP+08] J. Cheng, S. Lazard, L. Peñaranda, M. Pouget, F. Rouillier, and E. Tsigaridas. On the topology of planar algebraic curves. In *Proc. of the 24th European Workshop on Computational Geometry*, March 2008.

[CR88] M. Coste and M. F. Roy. Thom's lemma, the coding of real algebraic numbers and the computation of the topology of semi-algebraic sets. *J. Symb. Comput.*, 5(1/2):121–129, 1988.

[DET07] D. I. Diochnos, I. Z. Emiris, and E. P. Tsigaridas. On the complexity of real solving bivariate systems. In C. W. Brown, editor, *Proc. Int. Symp. Symbolic and Algebraic Computation*, pages 127–134, Waterloo, Canada, 2007.

[EK08] A. Eigenwilling and M. Kerber. Exact and efficient 2d-arrangements of arbitrary algebraic curves. In *Proc. 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA08)*, pages 122–131, San Francisco, USA, January 2008. ACM-SIAM, ACM/SIAM.

[EKK+05] A. Eigenwillig, L. Kettner, W. Krandick, K. Mehlhorn, S. Schmitt, and N. Wolpert. A Descartes Algorithm for Polynomials with Bit-Stream Coefficients. In V. Ganzha, E. Mayr, and E. Vorozhtsov, editors, *CASC*, volume 3718 of *LNCS*, pages 138–149. Springer, 2005.

[EKW07] A. Eigenwillig, M. Kerber, and N. Wolpert. Fast and exact geometric analysis of real algebraic plane curves. In C. W. Brown, editor, *Proc. Int. Symp. Symbolic and Algebraic Computation*, pages 151–158, Waterloo, Canada, 2007. ACM.

[EMT08] I. Z. Emiris, B. Mourrain, and E. P. Tsigaridas. Real Algebraic Numbers: Complexity Analysis and Experimentation. In P. Hertling, C. Hoffmann, W. Luther, and N. Revol, editors, *Reliable Implementations of Real Number Algorithms: Theory and Practice*, volume 5045 of *LNCS*, pages 57–82. Springer Verlag, 2008. also available in www.inria.fr/rrrt/rr-5897.html.

[ESY06] A. Eigenwillig, V. Sharma, and C. K. Yap. Almost tight recursion tree bounds for the Descartes method. In *Proc. Int. Symp. on Symbolic and Algebraic Computation*, pages 71–78, New York, NY, USA, 2006. ACM Press.

[Fau02] J.-C. Faugère. A new efficient algorithm for computing gröbner bases without reduction to zero f_5. In *International Symposium on Symbolic and Algebraic Computation Symposium - ISSAC 2002, Villeneuve d'Ascq, France*, Jul 2002.

[Fen92] H. Feng. *Decomposition and Computation of the Topology of Plane Real Algebraic Curves*. Ph.d. thesis, The Royal Institute of Technology, Stockholm, 1992.

[FGb] FGb - A software for computing Gröbner bases. J.-C. Faugère. `http://fgbrs.lip6.fr`.

[GLS01] M. Giusti, G. Lecerf, and B. Salvy. A Gröbner free alternative for solving polynomial systems. *J. of Complexity*, 17(1):154–211, 2001.

[GPS01] G.-M. Greuel, G. Pfister, and H. Schönemann. SINGULAR 3.0 — a computer algebra system for polynomial computations. In M. Kerber and M. Kohlhase, editors, *Symbolic computation and automated reasoning, The Calculemus-2000 Symposium*, pages 227–233. A. K. Peters, Ltd., Natick, MA, USA, 2001.

[GVEK96] L. González-Vega and M. El Kahoui. An improved upper complexity bound for the topology computation of a real algebraic plane curve. *J. Complexity*, 12(4):527–544, 1996.

[GVLRR89] L. González-Vega, H. Lombardi, T. Recio, and M-F. Roy. Sturm-Habicht Sequence. In *Proc. Int. Symp. on Symbolic and Algebraic Computation*, pages 136–146, 1989.

[GVN02] L. González-Vega and I. Necula. Efficient topology determination of implicitly defined algebraic plane curves. *Computer Aided Geometric Design*, 19(9), 2002.

[Hon96] H. Hong. An efficient method for analyzing the topology of plane real algebraic curves. *Mathematics and Computers in Simulation*, 42(4–6):571–582, 1996.

[Ker06] M. Kerber. Analysis of real algebraic plane curves. Master's thesis, MPII, 2006.

[KOR05] J. Keyser, K Ouchi, and M. Rojas. The exact rational univariate representation for detecting degeneracies. In *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*. AMS Press, 2005.

[Lab08] O. Labs. A list of challenges for real algebraic plane curve visualization software. Manuscript, 2008.

[Laz83] D. Lazard. Gröbner bases, Gaussian elimination, and resolution of systems of algebraic equations. In *EUROCAL' 83 European Computer Algebra Conference*, volume 162 of *LNCS*, pages 146–156. Springer, 1983.

[LPY04] C. Li, S. Pion, and C. Yap. Recent progress in exact geometric computation. *J. of Logic and Algebraic Programming*, 64(1):85–111, 2004. Special issue on"Practical Development of Exact Real Number Computation".

[LR01] T. Lickteig and M-F. Roy. Sylvester-Habicht Sequences and Fast Cauchy Index Computation. *J. Symb. Comput.*, 31(3):315–341, 2001.

[MC02] S. McCallum and G. E. Collins. Local box adjacency algorithms for cylindrical algebraic decompositions. *J. Symb. Comput.*, 33(3):321–342, 2002.

[MPS+06] B. Mourrain, S. Pion, S. Schmitt, J.-P. Técourt, E. P. Tsigaridas, and N. Wolpert. Algebraic issues in Computational Geometry. In J.-D. Boissonnat and M. Teillaud, editors, *Effective Computational Geometry for Curves and Surfaces*, Mathematics and Visualization, chapter 3. Springer, 2006.

[MT05] B. Mourrain and P. Trébuchet. Generalized normal forms and polynomial system solving. In *Proc. Int. Symp. Symbolic and Algebraic Computation*, pages 253–260, 2005.

[Rou99] F. Rouillier. Solving zero-dimensional systems through the rational univariate representation. *J. of Applicable Algebra in Engineering, Communication and Computing*, 9(5):433–461, 1999.

[RS] RS - A software for real solving of algebraic systems. F. Rouillier. `http://fgbrs.lip6.fr`.

[RZ03] F. Rouillier and P. Zimmermann. Efficient isolation of polynomial real roots. *J. of Computational and Applied Mathematics*, 162(1):33–50, 2003.

[Sak91] T. Sakkalis. The topological configuration of a real algebraic curve. *Bull. Austrl. Math. Soc.*, 43:37–50, 1991.

[SF90] T. Sakkalis and R. Farouki. Singular points of algebraic curves. *J. Symb. Comput.*, 9(4):405–421, 1990.

[Str06] A. Strzebonski. Cylindrical algebraic decomposition using validated numerics. *J. Symb. Comput.*, 41(9):1021–1038, 2006.

[SW05] R. Seidel and N. Wolpert. On the exact computation of the topology of real algebraic curves. In *Proc 21st ACM Symposium on Computational Geometry*, pages 107–115, 2005.

[Tei73] B. Teissier. Cycles évanescents, sections planes et conditions de Whitney. (french). In *Singularités à Cargèse (Rencontre Singularités Géom. Anal., Inst. Études Sci., Cargèse, 1972)*, number 7–8 in Asterisque, pages 285–362. Soc. Math. France, Paris, 1973.

[vzGG03] J. von zur Gathen and J. Gerhard. *Modern Computer Algebra*. Cambridge Univ. Press, Cambridge, U.K., 2nd edition, 2003.

[Yap00] C.K. Yap. *Fundamental Problems of Algorithmic Algebra*. Oxford University Press, Oxford-New York, 2000.

[Yap06] C.K. Yap. Complete subdivision algorithms, I: intersection of Bézier curves. In N. Amenta and O. Cheong, editors, *Proceedings of the 22nd ACM Symposium on Computational Geometry, Sedona, Arizona, USA, June 5-7, 2006*, pages 217–226. ACM, 2006.